

Modeling frameworks for resilience in socio-technical systems

by

Adaeze Okeukwu-Ogbonnaya

B.Tech., Federal University of Technology, Minna, 2012

AN ABSTRACT OF A DISSERTATION

submitted in partial fulfillment of the
requirements for the degree

DOCTOR OF PHILOSOPHY

Department of Computer Science
Carl R. Ice College of Engineering

KANSAS STATE UNIVERSITY
Manhattan, Kansas

2025

Abstract

In the face of growing disaster frequency and severity, we aim to strengthen resilience across socio-technical systems. We develop several modeling frameworks to understand how communities prepare for and recover from disruptions that span interdependent infrastructures and communication channels. Although communities differ in structure and vulnerability, disasters uniformly strain both the physical networks and the social systems that support them.

This dissertation examines resilience across two fronts: the operational resilience of communities within interdependent infrastructure systems and the communicative resilience of information flow within and between them.

In the first part, we use a Markov Decision Process (MDP) framework to model failure-and-repair dynamics in binary hetero-functional graphs (HFGs) that couple power, water, and community infrastructures. We show that communities experiencing power outages respond more sensitively to prioritization strategies based on vulnerability or criticality, whereas those affected by water shortages are primarily influenced by service demand. To overcome the limits of binary HFGs, we introduce the Hetero-functional Agent-Based Infrastructure Toolkit (HABiT). HABiT enables fine-grained simulation of heterogeneous infrastructures and their interdependencies. We apply HABiT to a synthetic model of three communities with varying infrastructure access and social vulnerability. The simulation reproduces normal operations, disruptions, and recovery influenced by scarce resources, vehicle routing, and mobile repair crews. By incorporating stochasticity, we uncover variations in cascading failures and recovery patterns that deterministic models can miss. HABiT thus enables rapid evaluation of disruption scenarios and guides resource allocation and recovery planning under uncertainty.

Recognizing that technical criteria alone may diverge from community priorities, we integrate community preferences into decision-making. We employ Large Language Models (LLMs) as proxy survey tools, generating simulated personas with diverse disaster experiences to obtain infrastructure-repair preferences. We aggregate these responses through a learning-to-rank algorithm, producing a total repair order that balances technical feasibility with social priorities.

In the second part, we shift from infrastructure to information, examining how communication systems support or hinder resilience. Using survey data from three Midwest counties, we develop stochastic diffusion models that trace how information spreads through communities during both normal and disaster conditions. We integrate neighbors, online social networks, local news, cable news, and local government, mapping trust-based and interaction-based ties in each community. We find that trust in local government, frequent interaction with cable news, and social media strongly shape diffusion, with diffusion speed being independent of community size. We then extend this analysis to examine manipulation of information in online social networks. Using Twitter data from the 2016 United States elections, we analyze interactions between Russian troll tweets and public replies. We cluster historical user-interaction sequences and apply statistical tests to reveal diverse, unpredictable engagement behaviors. We find that trolls lacked a consistent strategy for provoking responses, although content with political figures generally attracted more attention.

Modeling frameworks for resilience in socio-technical systems

by

Adaeze Okeukwu-Ogbonnaya

B.Tech., Federal University of Technology, Minna, 2012

A DISSERTATION

submitted in partial fulfillment of the
requirements for the degree

DOCTOR OF PHILOSOPHY

Department of Computer Science
Carl R. Ice College of Engineering

KANSAS STATE UNIVERSITY
Manhattan, Kansas

2025

Approved by:

Major Professor
George Amariuca

Copyright

© Adaeze Okeukwu-Ogbonnaya 2025.

Abstract

In the face of growing disaster frequency and severity, we aim to strengthen resilience across socio-technical systems. We develop several modeling frameworks to understand how communities prepare for and recover from disruptions that span interdependent infrastructures and communication channels. Although communities differ in structure and vulnerability, disasters uniformly strain both the physical networks and the social systems that support them.

This dissertation examines resilience across two fronts: the operational resilience of communities within interdependent infrastructure systems and the communicative resilience of information flow within and between them.

In the first part, we use a Markov Decision Process (MDP) framework to model failure-and-repair dynamics in binary hetero-functional graphs (HFGs) that couple power, water, and community infrastructures. We show that communities experiencing power outages respond more sensitively to prioritization strategies based on vulnerability or criticality, whereas those affected by water shortages are primarily influenced by service demand. To overcome the limits of binary HFGs, we introduce the Hetero-functional Agent-Based Infrastructure Toolkit (HABiT). HABiT enables fine-grained simulation of heterogeneous infrastructures and their interdependencies. We apply HABiT to a synthetic model of three communities with varying infrastructure access and social vulnerability. The simulation reproduces normal operations, disruptions, and recovery influenced by scarce resources, vehicle routing, and mobile repair crews. By incorporating stochasticity, we uncover variations in cascading failures and recovery patterns that deterministic models can miss. HABiT thus enables rapid evaluation of disruption scenarios and guides resource allocation and recovery planning under uncertainty.

Recognizing that technical criteria alone may diverge from community priorities, we integrate community preferences into decision-making. We employ Large Language Models (LLMs) as proxy survey tools, generating simulated personas with diverse disaster experiences to obtain infrastructure-repair preferences. We aggregate these responses through a learning-to-rank algorithm, producing a total repair order that balances technical feasibility with social priorities.

In the second part, we shift from infrastructure to information, examining how communication systems support or hinder resilience. Using survey data from three Midwest counties, we develop stochastic diffusion models that trace how information spreads through communities during both normal and disaster conditions. We integrate neighbors, online social networks, local news, cable news, and local government, mapping trust-based and interaction-based ties in each community. We find that trust in local government, frequent interaction with cable news, and social media strongly shape diffusion, with diffusion speed being independent of community size. We then extend this analysis to examine manipulation of information in online social networks. Using Twitter data from the 2016 United States elections, we analyze interactions between Russian troll tweets and public replies. We cluster historical user-interaction sequences and apply statistical tests to reveal diverse, unpredictable engagement behaviors. We find that trolls lacked a consistent strategy for provoking responses, although content with political figures generally attracted more attention.

Table of Contents

List of Figures	xiii
List of Tables	xvii
Acknowledgements	xix
Dedication	xxi
1 Introduction	1
1.1 Overview of Subsequent Chapters	3
2 Dynamic Hetero-Functional Graph Modeling for Enabling Community-Value-Informed Operational Decisions in Utility Infrastructures	6
2.1 Introduction	6
2.1.1 Acknowledgments	9
2.2 Literature Review	9
2.3 Data	11
2.3.1 Integrated Infrastructure and Community Network	12
2.3.2 Metrics	15
2.3.3 Community Features	16
2.4 Methodology	17
2.4.1 Binary Hetero-functional Graph	17
2.4.2 Modeling Failure and Repair with a Stochastic Binary Hetero-Functional Graph Model	18
2.4.3 Repair Orderings	20

2.4.4	Simulation	21
2.5	Experiments	23
2.6	Results	24
2.6.1	Dependencies between Power and Community Functionalities	24
2.6.2	Dependencies Water and Community Functionalities	28
2.6.3	Correlation of Repair Ordering Rankings	31
2.7	Conclusion	33
3	HABiT: A Hetero-functional Agent-Based Infrastructure Toolkit for Modeling Com-	
	munity Resilience	35
3.1	Introduction	35
3.1.1	Acknowledgments	37
3.2	Literature Review	38
3.3	Methodology	41
3.3.1	Modeling Approach	41
3.3.2	Model Assumptions	42
3.3.3	Hetero-Functional Graph Modeling	43
3.3.4	Architecture	44
3.4	Implementation	44
3.4.1	Graph-Driven Scheduling	46
3.4.2	Power Classes	48
3.4.3	Water Classes	50
3.4.4	Transportation Classes	52
3.4.5	Community-Related Classes	53
3.4.6	Repair and Event Classes	54
3.5	Case Study: A Synthetic Multi-Infrastructure System of Three Communities	55
3.5.1	Overview of Components	56
3.5.2	Graph Layout and Composition	57

3.6	Simulation Results	62
3.6.1	Scenario A	62
3.6.2	Scenario B	63
3.6.3	Scenario C	66
3.7	Discussion	69
3.8	Conclusion	71
4	LLM-Based Community Surveys for Operational Decision Making in Intercon- nected Utility Infrastructures	73
4.1	Introduction	73
4.1.1	Acknowledgments	75
4.2	Literature Review	75
4.3	Methodology	79
4.3.1	Hetero-Functional Graph (HFG) Model	79
4.3.2	LLM-based Simulated Community Preferences	81
4.3.3	Feature Extraction & Data Processing	84
4.3.4	Pairwise Comparator and Chainization Model	85
4.3.5	Random Subset Sampling and Fine-Tuning with Pseudo-Pairs	86
4.3.6	Validation and Sensitivity Analysis	87
4.4	Results	89
4.4.1	Impact of Community Awareness in Simulated Preferences	89
4.4.2	Accuracy with Limited Pairwise Comparisons	90
4.4.3	Sensitivity to Prompt Variations	92
4.5	Discussion	93
4.6	Conclusion	95
5	Towards Quantifying the Communication Aspect of Resilience in Disaster-Prone Communities	98

5.1	Introduction	98
5.1.1	Acknowledgments	102
5.2	Literature Review	102
5.3	Methodology	105
5.3.1	Data	105
5.3.2	Graph Generation	107
5.3.3	Stochastic Model	109
5.3.4	Network Metrics	111
5.3.5	Diffusion Time Prediction with Gaussian Process Regression	112
5.4	Results	113
5.4.1	Network Properties of Information Networks	113
5.4.2	Information Diffusion in Normal and Disaster Times	115
5.4.3	Resilience to Information Hub Failure	118
5.4.4	Impact of Population Size on Time of Diffusion	119
5.4.5	Gaussian Process Regression Model Results	121
5.5	Discussion	123
5.6	Conclusion	125
6	Predicting Topic Engagement during Russian Troll Activity around the 2016 US Political Elections	127
6.1	Introduction	127
6.2	Literature Review	128
6.2.1	Theoretical Models of Manipulation	129
6.2.2	Empirical Models of Manipulation	129
6.2.3	Political Discourse Manipulations on Twitter	129
6.3	Methodology	130
6.3.1	Data Collection	131
6.3.2	Data Preprocessing	131

6.3.3	Feature Extraction	132
6.3.4	Clustering	133
6.3.5	Hypothesis Testing	135
6.4	Results	136
6.4.1	Analysis of Clustering Results	136
6.4.2	Median Response Time of Tweet Sequence	137
6.4.3	Analyzing Response Time Trends	139
6.4.4	Content Analysis of Troll Tweets	141
6.5	Discussion	143
6.6	Conclusion	144
7	Conclusion	145
7.1	Contributions	145
7.2	Future Work	148
	Bibliography	150
A	Chapter 4 Appendix	171
A.0.1	Prompt Templates	171
A.0.2	Sample Persona and Response	176
B	Chapter 5 Appendix	179
B.0.1	Stochastic Model	179
B.0.2	Generating Trust and Interaction Values	183
B.0.3	Gradient Results for the Budget Allocation Method	185

List of Figures

2.1	The Integrated Infrastructure and Community Network Model	13
2.2	The Dependency Graph of Interconnected Infrastructure. Directed edges (arrows) indicate functional dependencies from one node to another.	14
2.3	Functionality Loss (Average Number of Functionalities Affected) Over Simulation Timestep (Total Repair Order by Criticality)	25
2.4	Functionality Loss (Average Number of Functionalities Affected) Over Simulation Timestep (Total Repair Order by SVS)	26
2.5	Functionality Loss (Average Number of Functionalities Affected) Over Simulation Timestep (Total Repair Order by SVS and Criticality)	26
2.6	Average Time of Outages by Location Before Repair (Ordering by Criticality)	27
2.7	Average Time of Outages by Location Before Repair (Ordering by SVS) . . .	28
2.8	Average Time of Outages by Location Before Repair (Ordering by SVS and Criticality)	29
2.9	Proportion of DNS over Total Load (Repair Order by Criticality)	29
2.10	Proportion of DNS over Total Load (Repair Order by SVS)	30
2.11	Proportion of DNS over Total Load (Repair Order by Weighted Criticality and SVS)	30
2.12	Average Time to Repair Each Node in 1000 simulations	31
2.13	Average Water Storage Levels Over Time	32
2.14	Ratio of Total water over Total Demand for Each Community	32

3.1	The HABiT Architecture. This architecture shows the interdependencies across transport, water, and power infrastructures. Arrows indicate resource or control flow between agent types. Repair crews are dispatched via the transportation layer to failed components marked by RepairTasks. Resources are generated, transported, and consumed. EnableService evaluates the quality of service delivered to community-facing infrastructure such as hospitals, schools, and residential buildings.	45
3.2	A Synthetic Network of 3-Community Interdependent Infrastructure Network. The diagram shows roads, power lines, and water pipelines connecting residential, school, and hospital buildings across three communities.	58
3.3	Community-level DNS under three load-shedding policies (SVS, Criticality, Weighted) when a generator's capacity generation is reduced and less than demand; demands are stochastic.	64
3.4	Building-level breakdown of DNS for each community (SVS = 0.2, 0.3, 0.7). Each bar shows the fraction of unmet demand under SVS, Criticality, and Weighted policies. The scenario is under stochastic loads.	65
3.5	Community-level DNS under three load-shedding policies (SVS, Criticality, Weighted) when a generator's capacity generation is reduced and less than demand; demands are deterministic.	66
3.6	Building-level breakdown of DNS for each community (SVS = 0.2, 0.3, 0.7). Each bar shows the fraction of unmet demand under SVS, Criticality, and Weighted policies. The scenario is under deterministic loads.	67
3.7	Three Time series of Average Service Score showing System-Wide Service recovery curve under three pump failure and repair scenarios, aggregated over 100 simulation runs. Each plot shows the results under a different repair prioritization policy.	68

3.8	Three Time series of Average Service Score comparing service recovery for 3 communities under three pump failure and repair scenarios, aggregated over 100 simulation runs. Each plot shows the results under a different repair prioritization policy. C1, C2, C3 represents Community1, Community2 and Community3 respectively.	69
3.9	A Bar Chart of Water Demand Not Served per community under three pump failure and repair scenarios, aggregated over 100 simulation runs. Each colored bar represents results under a different repair prioritization policy.	70
4.1	The Integrated Infrastructure and Community Network Model	80
4.2	A subset of the hetero-functional graph (HFG) showing dependencies across power and water systems. The final layer consists of consumption functionalities that can be repaired in any order. Arrows represent functional dependencies. The highlighted functionalities are damaged while the links marked $\times$ are paths requiring repair before the associated community-level services (bottom layer) can be restored.	81
4.3	Six-fold cross-validation curves for the pairwise comparator at $LR = 1 \times 10^{-3}$. Solid lines show training loss; dashed lines show validation loss for each fold.	86
4.4	Community Preferences by Community-Aware Personas	91
4.5	Community Preferences by Community-Unaware Personas	92
4.6	Kendall's τ correlation versus percentage of comparisons used.	93
4.7	Top-K overlap on Rankings obtained via Random SubSampling and Pseudo-Pairs	94
4.8	Prompts Variations Overall Agreement with Original using Kendall's τ and Top-K overlaps.	95
5.1	Classification of Areas in County. These maps were created using Adobe Photoshop (version 22.1, Adobe Systems, https://adobe.com/products/photoshop.html)	105

5.2	The Features of a Community	108
5.3	Diffusion Times reaching to 90% of Population	114
5.4	Sensitivity Analysis of the Average Diffusion Time during Disaster (1a, 1b, 1c) and Normal (2a, 2b, 2c) Times	117
5.5	Average diffusion times comparing scenarios with one hub failure against all hubs working, under disaster and normal conditions.	120
6.1	Frequency of Tweets in Each Cluster	137
6.2	Statistical Median and Sample Median Response Time for Each Cluster . . .	137
6.3	Sequence of Median Response Time of Tweet Sequences 397 and 3084	139
6.4	Sequence of Median Response Time of Tweet Sequences 3424 and 3479 . . .	139
6.5	Sequence of Mean Response Time of Tweet Sequences 3489 and 3491	139
6.6	Response Time of First Two Responses to Troll's Tweets Over Sequence (De- creasing Pattern)	140
6.7	Response Time of First Two Responses to Troll's Tweets Over Sequence (In- creasing Pattern)	141
6.8	Response Time of First Two Responses to Troll's Tweets Over Sequence (Sta- ble Pattern)	141
A.1	Original Prompt	172
A.2	Changing Instruction Format	173
A.3	Shifting Tone Prompt	174
A.4	Rewording Question	175
A.5	Rewording Reasoning	176
A.6	Full specification of one synthetic persona used in the simulation	177
A.7	Sample Response	178

List of Tables

2.1	Simple Hetero-functional Matrix	18
2.2	Correlation of the Repair Order Rankings	33
3.1	Node Attributes in the GraphML Model	59
3.2	Average Repair Time (in Steps) for Each Pump by Scenario	69
4.1	Attributes Used to Generate Synthetic Personas	83
4.2	Infrastructure Repair Option Features	85
4.3	Final Global Repair Prioritization (Same for Community-Aware and Community-Unaware Datasets)	90
4.4	Global Repair Prioritization Across Prompt Variations	96
5.1	Distribution of Responses of Communities. Note: * removed due to insufficient data	106
5.2	Network Properties of Information Network	113
5.3	Reliance of Urban and Rural Communities on Online Social Networks	120
5.4	Comparison Using Different Number of Nodes in Network	121
5.5	Gaussian Process Model Predictions for the 10 community groups.	122
6.1	Cummulative Features	133
6.2	Cluster Distribution Test-Pareto	136
6.3	Cluster Distribution Test-Exponential	136
6.4	Analysis of Median Response Time Trends	140
6.5	Identified Top Words with One Response	142
6.6	Identified Topics for Troll Tweets with Many Responses	142

B.1 Table of Gradient Results 187

Acknowledgments

I am sincerely grateful to my advisor, George Amariuca, for his patient guidance and unwavering support throughout this Ph.D. journey. His encouragement and mentorship were pivotal to the success of my journey, especially during challenging times when I doubted myself. I feel fortunate to have been his student.

I also thank my Ph.D. committee members, Eugene Vasserman, Doina Caragea, and Bala Natarajan, for their valuable feedback and support, which helped shape my research and personal growth. I am grateful for the opportunity to have worked with each of you.

To all my collaborators (colleagues in PITS Lab and ARISE project), thank you for enriching my academic journey with your insights and contributions. I appreciate the faculty and staff of the Computer Science Department at Kansas State University for creating such a supportive environment.

I gratefully acknowledge the support of the American Association of University Women (AAUW) fellowship, for awarding me a one-year fellowship that provided living support during my studies.

Words cannot fully express my appreciation to my husband, Onyekachi Ogbonnaya. His unwavering support, patience, and understanding carried me through the most challenging periods of this journey. Thank you for being the best husband and a wonderful father to our child, Chinalurum. I could not have come this far without you. I am also thankful to my mother, Ijeoma Anaegbu-Okeukwu, for laying the foundation that made this possible. She is the strongest woman I know, and I am forever grateful for her prayers, love, and constant support. My heartfelt thanks go to my siblings, Obinna and Onyinyechi Okeukwu, for always being there for me. I also thank my in-laws for their encouragement, especially my parents-in-law, Iroakazi and Comfort Ogbonnaya. Special thanks to Hon. Vernatius Okeukwu and Andrew Okeukwu for stepping into a fatherly role in my life. I am also grateful to my aunt and uncle, Godson and Adachukwu Ag Umeh, for their prayers over the years.

Thank you, Folashade Adebayo, for your steadfast support. I remember the late Lilian Atobatele with deep gratitude. May your soul rest in peace. I am grateful to my Westview Community Church family, especially Brian Smith, Dick and Joyce Nelson, and John and Jo Wolters, and to the African community at Kansas State University. Thank you, Gabriel Inalegwu, Aunty Hope Ezuruonye, and Chinemerem Onah, for all your support.

Above all, I thank God for His guidance and blessings throughout this journey.

Dedication

To my mom, for your prayers and strength.

To my husband, for your love and sacrifice.

To my son, for your laughter and babbles.

Chapter 1

Introduction

In recent years, there has been a growing focus on strengthening infrastructure resilience in light of global challenges like population growth and climate change. This increased attention has led to studies aimed at improving the reliability and robustness of infrastructure. One significant discovery is the interconnectedness of infrastructures and their relationship with surrounding communities^{1;2}. Therefore, the design and development of resilient infrastructure becomes complicated by the need to consider its heterogeneous and dynamic states, as well as uncertainties stemming from human behavior and environmental conditions. This complexity makes it relevant to understand how to model resilient infrastructure and community systems, especially when decisions surrounding infrastructure can affect the resilience of their connected communities.

Not only are critical infrastructure systems such as power, water, and transportation deeply interdependent, but they can also operate under uncertain conditions that can change rapidly during disruptions¹. And when such failures occur, the consequences often cascade across systems, and affect not just physical assets but also the communities that rely on them^{3;4}. For this reason, modeling these interdependencies is essential to understand how disruptions propagate and to design effective strategies to restore services⁵⁻⁷ back to communities.

But what services do we prioritize in communities when natural disasters damage multi-

ple critical infrastructures across communities? This situation creates complex prioritization challenges for response efforts. Traditional approaches for prioritizing infrastructure repairs rely heavily on expert opinions, which potentially overlook diverse community perspectives and public demand urgency. A recent study by Doorn et al⁸ shows that restoring damaged infrastructure during disasters following the technical approach is likely ineffective in addressing the immediate needs of affected communities. Also, studies have shown that vulnerable communities have higher immediate needs and are most likely to be affected if such needs are not met^{9;10}. With this understanding, it becomes more important to integrate community preferences into infrastructure response and post-disaster efforts.

Infrastructure resilience goes beyond the physical systems themselves and is tied to communities that need to communicate with each other on different levels to support the resilience effort. This understanding that communities rely on communication leads us to focus on the communication resilience of information flow within and outside a community. Effective communication of life-saving information during a disaster is critical. However, the unique characteristics of communities, whether rural, suburban, or urban, can affect the speed at which such information spreads, potentially impacting a community's resilience. Therefore, we need to understand how the process of information diffusion, which is indicative of a community's communication resilience, depends on the community's characteristics.

Digging deeper, how can we ensure that the information being diffused is authentic and free from manipulation? This question leads us to investigate manipulation strategies. Given that online social networks, such as Facebook, X (formerly Twitter), and Instagram, have become an essential part of people's daily lives for getting real-life information, it is important that such information is organic and free from manipulation. However, the ease of access and anonymity offered by these platforms have also led to the emergence of online trolls, who seek to manipulate, harass, and intimidate others for their own personal gain.

Manipulation by trolls in online social networks has become an increasingly concerning issue, particularly due to its potential to cause harm to individuals, damage reputations, and create a toxic online environment. Various manipulation techniques exist, including online astroturfing, misinformation, disinformation, clickbait, deepfake, gaming, trending

and ranking algorithms, and coordinating attacks^{11–17}. Given these challenges, identifying and understanding the strategies used to manipulate legitimate users on social media has become imperative.

Given these considerations, this dissertation has two primary goals: first, to operationalize community resilience within interdependent infrastructure systems, and second, to understand the communicative resilience of information flow within and between these systems. We address these goals from five distinct perspectives, detailed in Section 1.1.

1.1 Overview of Subsequent Chapters

In Chapter 2, we propose a method to enhance the resilience of interconnected infrastructure by utilizing a stochastic binary hetero-functional graph to model the dependencies among elements of an infrastructure network. Using a Markov Decision Process (MDP), we simulate failures and repair strategies and identify optimal repair strategies tailored to interdependent systems. Our developed algorithms incorporate equity-based metrics to ensure that we consider equity in the repair processes within the infrastructure network containing power, water, transportation, and community services. We aim to ensure a fair recovery process during disruptions by optimizing repair sequences based on interconnectedness and functionalities.

In Chapter 3, to overcome the limits of binary HFGs, we introduce the Hetero-functional Agent-Based Infrastructure Toolkit (HABiT). HABiT models infrastructure as a network of functionalities, such as generating power, transporting water, or consuming vehicles. Each functionality is represented by an autonomous agent with defined behavior, location, and evolving states. The simulation operates over discrete time steps and includes stochastic elements such as random failures, uncertain repair durations, and probabilistic transport availability. We apply HABiT to a simplified model of three communities with different levels of infrastructure access and social vulnerability. The simulation captures both normal operations and disruptions, including recovery influenced by constrained resources, vehicle routing, and mobile repair agents. HABiT uses a hetero-functional graph to represent in-

terdependencies across functionalities. This structure allows components to be added or removed without altering the system configuration, supporting flexible scenario design.

In Chapter 4, we represent interdependent infrastructure systems and communities alike with a hetero-functional graph (HFG) that encodes the dependencies between functionalities. This graph naturally imposes a partial order of functionalities that can inform the sequence of repair decisions to be made during a disaster across affected communities. However, using such technical criteria alone provides limited guidance at the point where the functionalities directly impact the communities, since these can be repaired in any order without violating the system constraints. To address this gap and improve resilience, we integrate community preferences to refine this partial order from the HFG into a total order. Our strategy involves getting the communities’ opinions on their preferred sequence for repair crews to address infrastructure issues, considering potential constraints on resources. Due to the delay and cost associated with real-world survey data, we utilize a Large Language Model (LLM) as a proxy survey tool. We use the LLM to craft distinct personas representing individuals, each with varied disaster experiences. We construct diverse disaster scenarios, and each simulated persona provides input on prioritizing infrastructure repair needs across various communities. Finally, we apply learning algorithms to generate a global order based on the aggregated responses from these LLM-generated personas.

A version of Chapter 4¹⁸ was accepted for publication [Adaeze Okeukwu-Ogbonnaya, Rahul Amatapu, Jason Bergtold, and George Amariuca. “LLM-Based Community Surveys for Operational Decision Making in Interconnected Utility Infrastructures.” Proceedings of the 2025 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM), 2025]. I served as the primary investigator and led the development of the research concept, experimental design, data analysis, and manuscript preparation. Rahul Amatapu contributed to the modeling and data preprocessing efforts. Jason Bergtold provided guidance on socioeconomic framing and interpretation of results. George Amariuca supervised the research and contributed to the conceptual development and manuscript revisions.

In Chapter 5, we investigate the communication networks of urban, suburban, and rural

communities from three US Midwest counties through a stochastic model that simulates the diffusion of information over time in disaster and in normal situations. To understand information diffusion in communities, we investigate the interplay of information that individuals get from online social networks, local news, government sources, mainstream media, and print media. We utilize survey data collected from target communities and create graphs of each community to quantify node-to-node and source-to-node interactions, as well as trust patterns. We run Monte Carlo simulation to obtain the average time it takes for information to propagate to 90% of the population for each community.

A version of Chapter 5¹⁹ of this dissertation has been published in Scientific Reports [Okeukwu-Ogbonnaya, A., Amariuca, G., Natarajan, B., & Kim, H. J. (2024). Towards quantifying the communication aspect of resilience in disaster-prone communities. Scientific Reports, 14, 8837]. I was the lead investigator, contributed to the development of the concept, participated in data collection and analysis, conducted the experiments, and performed the result analyses. Balasubramaniam Natarajan and Hyung Jin Kim contributed to the conception of ideas, acquisition, analysis, and interpretation of data. George Amariuca was the supervisory author on this project and was involved throughout the project from concept formation, acquisition, analysis, and interpretation of data to manuscript composition and edits. All authors participated in reviewing the manuscript.

In Chapter 6, we examine manipulation tactics used by Russian trolls during the 2016 U.S. General Election by analyzing Twitter data from February 8, 2012, to May 30, 2018. To understand the interaction dynamics between trolls and legitimate users, we link troll tweets with user replies and extract cumulative features that may influence user response behavior. Assuming that users respond to troll tweets based on tweet characteristics and their own interaction histories, we cluster user response histories to identify recurring behavioral patterns. We then perform statistical analysis of response time distributions within and across clusters to assess whether trolls employed coordinated strategies to prompt timely user engagement.

In Chapter 7, we focus on contributions, limitations, and future works of the dissertation.

Chapter 2

Dynamic Hetero-Functional Graph Modeling for Enabling Community-Value-Informed Operational Decisions in Utility Infrastructures

2.1 Introduction

In recent years, there has been a growing focus on bolstering infrastructure resilience in light of global challenges like population growth and climate change. This increased attention has led to studies aimed at improving the reliability and robustness of infrastructure. One significant discovery is the interconnectedness of infrastructures and their relationship with surrounding communities. Developing infrastructure poses a challenge due to these connections requiring consideration of the heterogeneous and dynamic infrastructure systems, along with uncertainties related to human actions and environmental factors.

Recently, to address the issues of the heterogeneous nature of infrastructure, a hetero-

functional graph (HFG) model has been proposed²⁰ to model unique functionalities in the infrastructure. This use of hetero-functional graphs eliminates one layer of complexities, hence we can view infrastructures as a network of functionalities that connect to each other and to communities.

A hetero-functional graph model representation of interdependent infrastructure maps functionalities of infrastructure resources to one another as a way to show their logical relationships. We define functionalities as the representations of processes that capture the capabilities of the resources. Some examples of processes are: Consume cold water, Enter electric vehicle, Exit electric vehicle, Transport Water from Point A to B ²⁰. Resources are objects that have functions such as transforming, transporting, or storing things or operands. Some examples of resources are houses, parking lots, water pipelines, roads, and power lines. Operands are inputs or outputs, that is, things that flow through the infrastructure and community. Some examples of operands are water, electric power, vehicles, and people.

In a hetero-functional graph, each vertex represents the functionality of the resources, while the edges represent the dependencies between functionalities. An edge AB between two nodes signifies a dependency between these nodes, meaning that functionality A must be operational for functionality B to work properly.

In our research, we aim to develop a framework to analyze how interdependent infrastructure changes over time. We use binary stochastic hetero-functional graphs for this purpose. These graphs enable us to illustrate the relationships among components within a system, such as the links between critical functionalities. In this model, we use this representation to show interdependencies; a “0” signifies no dependence, while a “1” shows there is a dependence.

To improve our insights regarding these dynamic systems, we integrate Markov Decision Processes (MDPs). Using MDPs to model interdependent infrastructures allows us to take into account the stochastic nature of the system, where outcomes are random and influenced by previous states. By utilizing MDPs, we can depict the dynamics of interconnected systems, make informed decisions based on probabilities, and devise strategies for improving infrastructures while taking various scenarios into account. Moreover, MDPs provide a

method for making decisions when dealing with uncertainties, paving the way for flexible approaches to infrastructure management.

When we develop resilience models, we consider three decision planes: the repair decision plane, the distribution of limited resources decision plane, and the investments decision plane. The repair decision plane focuses on making decisions about repair and restoration activities after events. It involves identifying which components or systems need attention, how resources should be assigned for repairs, and planning restoration efforts. A stochastic resilience model would take into account uncertainties related to repair times, resource availability, and the dynamic relationships between infrastructure systems. The distribution of limited resources decision plane deals with allocating and distributing resources like labor supply, materials, and financial resources to address resilience needs across infrastructure systems. The stochastic model would consider fluctuations in resource availability over time, competing resource demands from other infrastructures and sectors, and how resource allocation decisions affect system resilience. The investment decision plan includes decisions regarding long-term investments in infrastructure resilience, such as upgrades, maintenance, and implementing technologies or strategies to improve system robustness. A stochastic model would assist in evaluating the return on investment for measures that enhance resilience while taking into account uncertainties in future hazards, technological advancements, and socio-economic factors.

In summary, our research models interdependent infrastructures as binary stochastic hetero-functional graphs and uses MDP to comprehensively model and manage the stochastic nature of interdependent infrastructures. This integrated approach allows for a deeper understanding of dynamic system behavior and empowers us to make decisions that account for the inherent unpredictability and variability associated with complex, interdependent systems. Since our interest is in the repair decision plane, the objective of this chapter is to understand how different decision-making impacts communities. The contributions of this chapter are as follows:

1. We develop a Markov Decision framework to model the dynamics of failures and repairs

in a hetero-functional graph of multiple infrastructures.

2. We examine the implications of the various repair decisions ordered by community metrics such as vulnerability and criticality.
3. Our observations reveal distinct impacts on communities affected by power systems versus those affected by water failures. Decision-making based on vulnerability or criticality influences communities affected by power failure differently. In communities affected by water failure, impacts are mostly based on the water demand of each community.

2.1.1 Acknowledgments

This research was supported by funding from the USA National Science Foundation (NSF) under Grant Number 2148878.

2.2 Literature Review

The state of research in infrastructure resilience has evolved over time, moving from component-level optimization toward frameworks that account for uncertainty, interdependence, and social factors.

Some early works focused on minimizing downtime in individual infrastructure networks through optimization models, such as those in Zhang and Cheng²¹, who used an MDP to schedule post-disaster maintenance and showed that breaking large networks into smaller sub-networks can reduce repair delays. Similarly, Fang and Sansavini²² developed a two-stage stochastic programming approach to improve network recovery under uncertain repair durations and resource constraints. Alkhaleel et al.²³ extended this paradigm with scenario-based optimization and conditional value-at-risk metrics to account for uncertainty in repair and travel times, and Wu and Wang²⁴ introduced a performance-driven framework for interconnected networks using a combination of mathematical programming and heuristic agent coordination.

Next, we observe the shift toward modeling interdependent and multi-layered infrastructure systems. Ghorbani-Renani et al.²⁵ proposed a tri-level protection-interdiction-restoration model that simultaneously addresses vulnerability and recoverability in interconnected utilities. Almoghatawi et al.²⁶ and Liu et al.²⁷ presented restoration models for interdependent and transportation networks, respectively, leveraging mixed-integer optimization to provide prioritized repair schedules under limited crew availability. Vugrin et al.²⁸ formulated a multi-mode project scheduling framework for post-disaster recovery sequencing, emphasizing efficiency and service restoration. These studies illustrate the trend from single-system optimization to integrated approaches that account for cascading failures and cross-network effects.

Some research efforts have explored how functional importance and network topology inform resilience. Zhang et al.²⁹ introduced a method for balancing cost and resilience via stochastic ranking, while Fang et al.³⁰ proposed robust pre- and post-disruption planning to enhance network performance. Approaches such as genetic algorithms for bridge restoration³¹, resilience-aware transportation planning³², and network science heuristics like Cent Restore³³ demonstrate the use of functional criticality and topological measures in guiding recovery. Li et al.³⁴ and Ye and Ukkusuri³⁵ showed how multi-stage design and scheduling strategies improve post-disaster recovery in transportation systems.

In recent years, researchers started working on incorporating vulnerability and equity considerations into resilience planning. Giovannettone et al.³⁶ identified a persistent gap in embedding equity in infrastructure decision-making, warning of potential innovation traps that overlook historical injustices. Meerow et al.³⁷ characterized urban resilience planning across distributional, recognitional, and procedural equity, while Karakoc et al.³⁸ combined vulnerability indices with component importance measures in a multi-criteria restoration framework. Gunda et al.³⁹ emphasized the role of stakeholder-driven logic models for integrating equity into infrastructure planning, and Wang and Monache⁴⁰ demonstrated equity-focused recovery using reinforcement learning for urban traffic networks. Yap et al.⁴¹ and Baghersad et al.⁴² highlighted the trade-offs between efficiency and equity in infrastructure prioritization.

In addition to these planning approaches, functional and system-of-systems models have emerged to capture dynamic interactions and cascading failures. Cavallini et al.⁴³ developed a system dynamics framework to forecast the evolution of disruptions, while Filippini and Silva⁴⁴ analyzed functional dependencies to inform resilience assessments. Agent-based and hybrid simulation approaches, such as those by Münzberg et al.⁴⁵ and Rasoulkhani and Mostafavi⁴⁶, explicitly model human–infrastructure dynamics to capture emergent recovery behaviors. Guidotti et al.⁴⁷ provided a probabilistic framework for modeling interdependent networks, and Galbusera et al.⁴⁸ introduced Boolean network models to simulate failure and recovery phases under resource constraints. Turnquist and Vugrin⁴⁹ advanced stochastic optimization for distribution systems by integrating pre- and post-disruption strategies, and Argyroudis et al.⁵⁰ demonstrated quantitative resilience assessment for multi-hazard transport infrastructure.

Our work builds on these developments by integrating dynamic repair modeling, interdependence, and equity into a unified framework. Specifically, we model interdependent infrastructures as binary stochastic hetero-functional graphs and simulate repair decisions using MDP. Unlike prior studies that evaluate equity post hoc or treat it as a secondary objective, we include the community’s social vulnerability scores⁵¹ directly into the repair prioritization process alongside infrastructure criticality. By comparing repair strategies based on individual and combined metrics, our approach provides an operationally actionable and socially responsive framework for post-disaster infrastructure recovery.

2.3 Data

In this section, we present the notional infrastructure model utilized in our study and the metrics employed in our analysis.

2.3.1 Integrated Infrastructure and Community Network

In the proposed infrastructure model, we outline a comprehensive utility network encompassing electricity generation, distribution, water supply, transportation, and community network systems. This model shows three distinct communities, each featuring dedicated subsystems for electrical power distribution and water supply, as illustrated in Figure 4.1. In the diagram, undirected edges represent bidirectional connections, such as two-way roads, while directed edges represent one-way dependencies like the flow of power or water.

In our work, we focus on the subset of the infrastructure model, which is power, water, and community networks. The model centers around a main power generation unit, Generator 1 (Gen1), which acts as the core of the electricity distribution network. From Generator 1 (Gen1), power flows to a primary substation, SubStation 1 (SS1), then to four nodes T1-T4, which we refer to as transformers in this work. In practice, these represent the distribution transformers within the substations that deliver power to the three communities. Water Treatment Plant 1 (WTP1) gets electricity from T4 through Powerline 16 (P16). See the dependency graph in Figure 2.2. We discuss the network of three communities below:

- Community 1: Community 1 is served by T1, receiving electrical power from the central generator, Generator 1 (Gen1) facility. Power flows through Powerline 1 (P1) to SubStation 1 (SS1) and then via Powerline 2 (P2) to Transformer 1 (T1). Powerline 3 (P3), Powerline 4 (P4), Powerline 5 (P5), and Powerline 6 (P6) give power to the Elementary School 1 (ES1), Residential 1 (RES1), Commercial Area 1 (COM1), and Hospital 1 (H1) respectively. Simultaneously, potable water processed at Water Treatment Plant 1 (WTP1) is stored at the water storage tank Water Storage Tank 1 (WS1) and distributed to corresponding end-users through water pipelines Water Pipelines 2 (WP2), Water Pipelines 3 (WP3), Water Pipelines 4 (WP4), and Water Pipelines 5 (WP5). The transportation network within Community A comprises roads (R1, R23, R22, R28, R24, R29, R26, and R27), connecting residential, commercial, and household sectors with intersections (I4, I5, I6, I3, I1).
- Community 2: Transformer T2 serves Community 2's electrical needs, with power dis-




Figure 2.1: The Integrated Infrastructure and Community Network Model

tributed to the Elementary School 2 (ES2), Residential 2 (RES2), and Commercial Area 2 (COM2) through powerlines, Powerline 8 (P8), Powerline 10 (P10), and Powerline 11 (P11), respectively. Potable water flows from Water Treatment Plant 1 (WTP1) to the water storage tank, Water Storage Tank 2 (WS2), and is distributed through




Figure 2.2: The Dependency Graph of Interconnected Infrastructure. Directed edges (arrows) indicate functional dependencies from one node to another.

pipelines Water Pipelines 6 (WP6), Water Pipelines 7 (WP7), and Water Pipelines 8 (WP8) to meet the water consumption demands of the community nodes. The transportation infrastructure features includes roads (R5, R33, R32, R38, R35, R34, R39,

R36, and R37) which connects intersections (I11, I12, I16, I15, I14, I13).

- Community 3: Transformer T3 distributes electricity through powerlines P13, P14, and P15 to the Elementary School 3 (ES3), Residential 3 (RES3), and Hospital 2 (H2). Water distribution from Water Storage Tank 3 (WS3) is carried out through Water Pipelines 9 (WP9), Water Pipelines 11 (WP11), and Water Pipelines 10 (WP10) to these end-users. The transportation infrastructure comprises of roads (R8, R43, R42, R48, R45, R44, R49, R46, and R47), and intersections (I26, I21, I22, I25, I24, I23) enabling vehicle movement within communities.

In addition, road R100 connects Community 1 to Community 2, and roads R200 and R300 connect Community 2 and Community 3.

2.3.2 Metrics

We define four metrics that we use in decision-making and evaluating the results of those decisions.

- Social Vulnerability Score (SVS): The SVS is a scalable metric created to gauge the varying levels of vulnerability across different segments of a community during disasters or crises⁵¹. This index is tailored to capture the diverse starting conditions and experiences of these groups by utilizing demographic information from U.S. Census data. The SVS is derived from these demographic variables in a similar manner to the CDC/ATSDR Social Vulnerability Index (SVI)⁵², but is scaled for use in simulation. The SVS uses Zone 1 to Zone 5 to indicate the degree of vulnerability, with 5 showing the highest vulnerability. In our work, we normalize to the 0-1 range, with 1 showing the highest vulnerability.
- Criticality: Criticality measures the significance of the effects that a compromised asset can have on the wider infrastructure network. This measure is commonly quantified through structural metrics that determine a node’s connectivity and significance within the network or through reductions in functionality, impacting the delivery of essential

goods, services, or utilities through the node^{53;54}. To represent the level of criticality of functionalities, we assign values ranging from 0 to 1, where 1 indicates the highest level of criticality.

- **Weighted Social Vulnerability Score and Criticality Score:** The Weighted Social Vulnerability Score and Criticality Score merge social vulnerability and criticality metrics into a unified weighted measure through a linear combination. In this combination, each weight is a value between 0 and 1, where the total of all weights equals 1. We apply this method to provide a more comprehensive assessment by considering both the vulnerability of the community and the critical importance of the resource or node. We do this because our repair algorithm requires a single scalar value to produce a total order of repair priorities, we define:

$$\text{Priority}_i = w_{\text{SVS}} \text{SVS}_i + w_{\text{Crit}} \text{Crit}_i, \quad (2.1)$$

where w_{SVS} and w_{Crit} are the respective weights. The combined score is a ranking tool, not a physical quantity, and the weights can be tuned to reflect stakeholder preferences.

- **Demand Not Served (DNS):** DNS is a metric used to assess the consequences of a cascading event in power systems. It quantifies the amount of electrical load that cannot be supplied due to disruptions within the power network. This measure assumes that the customer’s load demand remains constant throughout the event⁵⁵. We compute DNS and Total Load as cumulative sums over the entire simulation period (i.e., total across all discrete time steps). This allows us to compare the impact of repair strategies on overall service delivery.

2.3.3 Community Features

We have three different communities, each with distinct characteristics. In our model, Community 1 has a social vulnerability score of 0.2, Community 2 scores 0.3, and Community 3 scores 0.7, indicating that Community 3 is the most vulnerable among them. Following this,

we assign criticality values by first looking at the dependency graph. The nodes that support more downstream functionalities get higher scores, which we determine using a depth-first traversal of the hetero-functional graph.

Gen1 is assigned the maximum criticality score (1.0). Nodes on the same level in a depth-first search tree share the same score, such as transformers at 0.87 and water storage tanks at 0.71. We assign hospitals a criticality score of 0.9 to reflect their essential role during emergencies, while residential, commercial, and school nodes are assigned lower scores (0.25–0.45) according to their position in the network. Additionally, we categorize each node within our communities as consuming nodes and assign them specific total demands for water and power. The data, code, and complete set of assigned scores for all nodes are available in the chapter’s GitHub repository⁵⁶.

2.4 Methodology

We present the binary hetero-functional graph and its application in modeling failure and repair dynamics via MDP. Next, we explore the process of prioritizing repairs using different metrics.

2.4.1 Binary Hetero-functional Graph

Binary Hetero-functional graphs are useful for modeling systems where dependencies are either present or absent without considering degrees or levels of dependency. In a binary hetero-functional graph, the edges between nodes represent dependencies between functionalities. Each node in a binary hetero-functional graph represents the functionality of the system. An edge AB between two nodes signifies a dependency between these nodes so that functionality A must be operational for functionality B to function effectively.

We apply hetero-functional graph theory to construct hetero-functional graphs of inter-dependent infrastructure from a multilayer network. Typically, in such a multilayer inter-dependent infrastructure network, represented as $G = (V, E)$, there are n vertices and m

edges. Here, the n nodes correspond to components within power, water, and transportation infrastructures, as well as community structures, while the edges represent the connecting pipelines, power lines, or road networks. However, in the hetero-functional graph, the vertices and edges in a multi-layer graph are nodes. The edges show the dependency relationship between the nodes. We show a simple hetero-functional graph in Table 2.1. For the hetero-functional graph of our model, we have 141 nodes and 208 edges.

Table 2.1: Simple Hetero-functional Matrix

	Treat Water at Water Treatment Plant	Store Water at Water Storage Tank
Consume Power at Water Treatment Plant	1	0
Transport Water from Water Treatment Tank to Water Storage Tank	0	1

From the graph, we observe that functionalities adhere to a 'Process + Resource' format. The functionalities represented in the columns depend on those listed in the rows. For example, the functionality 'Consume Power at Water Treatment Plant' precedes Treat Water at Water Treatment Plant. Typically, these functionalities represent the capabilities of the infrastructure, with nodes being defined according to the resource functions across various infrastructures.

2.4.2 Modeling Failure and Repair with a Stochastic Binary Hetero-Functional Graph Model

In this study, we explore the dynamics of component failures within interdependent infrastructure systems, particularly those damages that arise from disasters or disruptions. We use small case letters to denote scalars, small case bold letters for vectors, capital letters for matrices, and capital script letters for sets, which establishes the notational conventions that will be employed throughout our analysis.

Repairs are frequently required due to the interconnected nature of these systems. Our study concentrates on repairs linked to disasters or disruptions and explores the dynamics of component failures within interdependent infrastructures. Given this interdependence, a failure in one component can lead to cascading failures across the system. Hence, we

aim to model the decision-making process on which repair process to undertake, given the interdependent nature of our system. To accomplish this, we present a binary stochastic hetero-functional graph as a model of MDP. For an infrastructure system depicted by a hetero-functional graph, we define the MDP framework as follows:

1. $\mathcal{S}$: The state space of the system, denoted by $\mathcal{S}_i$ where each $\mathcal{S}_i$ lies within the range $(0, 1)$
2. $\mathcal{A}$: The action space where the action is stochastic, $a(k) \in \mathcal{A}$ is an n -tuple binary sequence of n actions at timestep k
3. $\mathcal{D}$: The set of damages, $d(k) \in \mathcal{D}$ is an n -tuple binary sequence of n damages at timestep k
4. G : The hetero-functional graph
5. $\mathbf{s}(k)$: The state of functionalities, $\mathbf{s}(k)$ is an n -tuple binary sequence of n functionalities in the system at timestep k
6. P : The transition probability matrix
7. r : The number of repair crews

We use this framework to model decision-making in infrastructure systems facing disruptions. The state space $\mathcal{S}$ and an action space $\mathcal{A}$ are represented by n -tuple binary sequences that enable the capture of system functionalities and stochastic actions. The damages, $\mathcal{D}$, enable the assessment of repair strategies under various scenarios. By integrating the hetero-functional graph, G , we ensure adaptive responses by directing repair actions based on the system's state and dependence. Utilizing a transition probability matrix P aids in modeling the stochastic nature of state transitions, facilitating long-term strategy assessment. Including the number of repair crews, r introduces an additional constraint, ensuring that at each time step, the number of system functionalities to be fixed aligns with the available repair crews. This introduces a queue-like mechanism into the framework. With this framework,

we can track the evolution of system states over time and make informed decisions regarding repair strategies.

Given that damage $\mathcal{D}$ and action $\mathcal{A}$ are stochastic, to find a new state, we solve the following:

$$d(k+1) \sim P(d(k+1)|d(k), a(k)) \quad (2.2)$$

$$s(k+1) = f(s(k), d(k+1), G), \quad (2.3)$$

where $d(k)$ is the current damage on the system and $d(k+1)$ is the next state of damage when actions, $a(k)$ are taken on the system. The function f is the mapping from the previous state, $s(k)$, to the new state, $s(k+1)$, given the new damage, $d(k+1)$ on the system, and the graph G which determines dependencies between functionalities.

2.4.3 Repair Orderings

In the previous subsection, we established a natural order of repairs based on the dependencies among nodes in the hetero-functional model and the extent of damages incurred. This ordering, referred to as a *partial order* on actions, is determined using the depth-first search algorithm, prioritizing the sequence of functionalities for repair. However, while this partial ordering ensures an organized sequence of repairs, it does not prioritize individual repair actions within functionalities on the same level.

To refine the repair ordering further, our aim is to establish a *total order* of actions. In addition to considering dependent relationships and damages, this approach incorporates various metrics to enhance the ordering process. Specifically, as our infrastructure comprises interconnected communities, we introduce additional metrics such as social vulnerability score, criticality, and weighted social vulnerability score and criticality. These metrics serve to prioritize repairs based on the broader impact within the community and the criticality of the infrastructure component.

We note that multi-objective and multi-criteria optimization methods are commonly used

for post-disaster repair planning. In this study, we evaluate each decision metric separately to preserve interpretability. SVS and criticality represent distinct prioritization philosophies, community-focused versus infrastructure-focused, that we aim to compare directly. Collapsing these into a single multi-objective framework could obscure the trade-offs that our analysis is designed to highlight.

We use a breadth-first search (BFS)⁵⁷ to traverse the dependency graph, starting from a given node and moving through neighboring nodes level by level. Our only change is that within each BFS layer, we sort the nodes based on a user-defined priority metric, such as criticality or social vulnerability score, so higher-priority nodes are handled first. This helps ensure that the traversal respects the dependency structure, avoids getting stuck in cycles or deep branches, and keeps the repair sequence aligned with the chosen priorities.

By integrating these metrics into our repair ordering process, our aim is to develop a framework that optimizes repair sequences and ensures efficient allocation of resources in response to infrastructure disruptions.

2.4.4 Simulation

Our study implements a detailed simulation model to track the evolution of damage within the infrastructure systems, focusing on the interactions between damage, repair processes, and infrastructure dependencies.

During each cycle of the simulation, the model assesses and updates the damage status of the infrastructure components based on a set of probabilistic rules. These rules simulate the stochastic nature of damage incidents and repair outcomes, influencing how each component’s state might change in response to operational decisions and external factors. To quantify the dynamics of repair and failure probabilities over time, the model incorporates the transition probability matrix that provides the probabilities of transitioning from one state to another within the system. The probability that a component remains undamaged is set to 1, implying certainty in the absence of external damage influences. The probability that a component remains in a slightly damaged but stable state is maintained at nearly

1 (0.999999), indicating that the component is almost certain to remain in this state without further intervention. For components that are actively being repaired, the transition probability matrix holds probabilities calculated for each time step based on an exponential distribution. These probabilities are crucial as they reflect the likelihood of a component being repaired over time, adjusting for the increasing complexity of repairs as time without intervention extends. The probabilities of repair are of the form $1 - e^{-\lambda t}$, where t is the time since the last transition and λ is a parameter that adjusts the rate at which the exponential decays, which emphasizes how repairs become more likely to be completed as time passes. In our simulation, time moves in discrete steps k , so we treat t as proportional to k .

Within the simulation, repair decisions are enacted according to the *total order* of actions, which directs the deployment of repair crews to specific components based on their current states and priority rules. Each decision impacts the system’s state, and the outcomes are determined during each simulation cycle, where a random number is generated for each functionality and compared against its corresponding transition probability from the transition matrix. If the functionality is under repair and the random number exceeds this probability for the given time step, indicating a successful repair according to our stochastic model, the component’s damage state is updated to repaired (0). If not, it remains damaged (1). This method enables the simulation to realistically replicate the variability and uncertainty typical of real-world repair outcomes, reflecting the dynamic nature of infrastructure maintenance.

In the water infrastructure, we model the failure and repair process of a pump in a pipeline connected to water storage as follows. When the pump is operational, and damage has been repaired (indicated by the status 'fixed'), the system calculates the total water usage by aggregating the demand across all consumption functionalities downstream. If the available water storage is sufficient to meet this demand, the water amount is appropriately reduced and then increased by a predetermined refill quantity. This addition helps maintain optimal water levels, ensuring continuous supply without exceeding the storage’s initial capacity, which is safeguarded by adjusting the storage to the lesser of the current or initial levels. In instances where the water storage does not meet the total demand, the

system prioritizes distribution based on a predefined criticality of nodes. Water is allocated to these critical nodes first, ensuring that the most important or vulnerable nodes receive water before others. If the storage can satisfy a node’s full demand, the amount is deducted and recorded. However, if the storage is insufficient, it allocates the remaining water to the node, exhausts the storage, and halts further distribution. After addressing these critical needs, the storage is refilled by the set refill quantity.

However, if the pump has not been fixed, the approach to managing water storage changes slightly. The system still sums up the total demand from consumption functionalities downstream but directly deducts this from the water storage. However, there is no refilling of the water storage since the pipes that supply water are not functional due to pump failure.

2.5 Experiments

In this section, we present the simulation setup for the experiment. We define two different scenarios considering failures in the power and water infrastructure, as follows.

1. We simulate the failure of three powerlines, P2, P7, and P12, which connect one substation to three transformers, which connect to the three different communities. Our simulation parameters include the number of repair crews given as 2 with the rate parameter of the exponential distribution (λ) as 0.3. These settings are implemented across four different repair orderings. Each experiment is executed through a Monte Carlo simulation, repeated 1000 times. We use different repair orderings based on criticality and social vulnerability scores. When we use the weighted scores, we use weights on (SVS: Criticality) metrics as (0.4:0.6). The power demand not served across the consumption functionalities, consume power at ES1, RES1, COM1, H1, RES2, ES2, COM2, RES3, ES3, and H2 are given as 100MW, 200MW, 500MW, 1000MW, 200MW, 100MW, 500MW, 200MW, 100MW, and 1000MW respectively.
2. We simulate the failure of two pumps, pumps 20 and pumps 30, on two water pipelines, WP12 and WP13. These water pipelines connect two water storage tanks, WS2 and

WS3, which connect to other pipelines and finally to Community 2 and Community 3 with SVS of 0.3 and 0.7, respectively. The water demands from Community 2 and Community 3 are 3300 liters and 2100 liters, respectively. Our simulation parameters include the number of repair crews given as 1 with the rate parameter of the exponential distribution (λ) as 0.3. Each experiment is executed through a Monte Carlo simulation, repeated 1000 times. We use different repair orderings based on criticality and social vulnerability scores. When we use the weighted scores, we use weights on (SVS: Criticality) metrics as (0.4:0.6).

2.6 Results

We compare the outcomes achieved by employing different repair orderings and also evaluate the effects of integrating equity metrics into the repair prioritization process.

2.6.1 Dependencies between Power and Community Functionalities

Our study explores how different metrics, Criticality, SVS, and a combination of both, impact the order of repair and thereby affect the functionality of essential infrastructure over time.

In this study, we use discrete simulation timesteps k rather than real-world time units. Each timestep represents a simulation step in which the system transitions to the next steady state after repairs or failures are processed. Our purpose for the simulation is to analyze recovery sequences and steady-state impacts, not to quantify transient dynamics. With this abstraction, each infrastructure can evolve on its own effective timescale. For example, power system processes stabilize faster than water system processes.

Figures 2.3, 2.4, and 2.5 illustrate the distinct dynamics of functionality loss and recovery across three power lines influenced by the metric used to prioritize repairs. The curves do not start at zero because the three powerlines, P2, P12, and P7, are already damaged at the start of the simulation. Functionality loss increases gradually because failures move through

the dependency graph one step at a time, and downstream nodes only become non-functional after their upstream parent has failed. Since we run 1000 Monte Carlo simulations, some failures happen in different steps across runs, and averaging over these runs produces the smooth rise seen in these figures.

In Figures 2.3–2.5, we quantify recovery by calculating the average rate at which functionality loss decreases between timesteps 3 and 6 where a steeper downward slope over this interval corresponds to a faster restoration of service. In Figure 2.3, where repairs are prioritized based on Criticality, functionality loss declines most rapidly for P2, with a slope of -0.57 functionalities per timestep, followed by P7 (-0.37) and P12 (-0.33). This outcome aligns with the presence of a critical facility (a hospital) in Community 1, which is served by P2. In Figure 2.4, which uses SVS-based prioritization, P12 shows the steepest decline in loss (-0.47), while P2 and P7 decline more gradually (-0.37 each). This result is consistent with P12 having the highest SVS, hence the fastest recovery. Under the combined SVS and Criticality strategy in Figure 2.5, P12 again experiences the most rapid decrease in loss (-0.50), followed by P7 (-0.47) and P2 (-0.23). Our results show that the prioritization strategy directly affects which functionality is restored.




Figure 2.3: Functionality Loss (Average Number of Functionalities Affected) Over Simulation Timestep (Total Repair Order by Criticality)

Furthermore, we investigate the average outage time across specific facilities within three communities, exploring the effects of different prioritization metrics on outage durations.




Figure 2.4: Functionality Loss (Average Number of Functionalities Affected) Over Simulation Timestep (Total Repair Order by SVS)




Figure 2.5: Functionality Loss (Average Number of Functionalities Affected) Over Simulation Timestep (Total Repair Order by SVS and Criticality)

Figures 2.6, 2.7, and 2.8 illustrate these variations in outage times influenced by the metrics of Criticality, SVS, and a combined weighting of SVS and Criticality. Figure 2.6 reveals a focused prioritization of hospitals in all communities, where the average outage time is consistently lower than for other facilities such as elementary schools, residential buildings, and commercial buildings. This demonstrates the high priority given to healthcare facilities due to their critical nature, ensuring that they experience the shortest repair delays to minimize impact on essential health services. Figure 2.7 shows a distinct prioritization

pattern where communities with higher SVS, presumably Community 3 and Community 2, experience shorter outage times across all facilities, not just hospitals. The SVS-driven prioritization emphasizes reducing vulnerabilities in socially sensitive areas, which leads to quicker recovery times in communities facing higher social risks. Combining SVS and Criticality leads to a nuanced prioritization strategy, as shown in Figure 2.8. Hospitals continue to receive prompt attention, indicative of their criticality, while the SVS component ensures that other essential services in high-vulnerability areas also see reduced outage durations. This dual-factor prioritization effectively balances the urgency of infrastructure repairs with the social vulnerability of the community, leading to a more equitable distribution of repair resources.




Figure 2.6: Average Time of Outages by Location Before Repair (Ordering by Criticality)

Also, we use physical usage to analyze the impacts of these repair orderings, as shown in Figures 2.9, 2.10, 2.11. We use the ratio of DNS over Total Load to see the effect of the different prioritization techniques. In Figure 2.9, repairs are prioritized based on the criticality of the communities. Community 1 and Community 3 are deemed the highest priority because they have hospitals, and Community 1 has the lowest electricity demand not served. Community 2 has the highest electricity demand not served, suggesting lower criticality levels than Community 1 and Community 3. In Figure 2.10, the order of DNS




Figure 2.7: Average Time of Outages by Location Before Repair (Ordering by SVS)

differs from the criticality-based plot, suggesting that the vulnerabilities do not exactly correlate with the critical infrastructure needs. In Figure 2.11, this plot combines both the criticality and SVS to determine the repair order. It integrates both the infrastructural importance and the community’s social vulnerability to produce a more holistic prioritization metric. This plot is similar to the plot where repairs are prioritized by criticality which shows that criticality has a stronger influence than SVS on the combined metric since the weight of criticality (0.6) is greater than SVS in the metric (0.4).

2.6.2 Dependencies Water and Community Functionalities

Figure 2.12 displays the average time to repair two critical pipelines, WP12 and WP13 when their failed pumps are fixed based on data from 1000 simulations. The plot distinctly shows that the repair time for WP13 is consistently shorter than for WP12. This observation can be linked to several operational factors influencing repair prioritization. WP13 serves Community 3, which includes critical infrastructure such as Hospital 2. Despite similar criticality scores from the dependency graph, WP13 receives prioritized repairs due to the presence of such an essential service to ensure minimal disruption to crucial healthcare services. Using SVS, WP13 receives prioritized repairs due to the higher SVS associated




Figure 2.8: Average Time of Outages by Location Before Repair (Ordering by SVS and Criticality)




Figure 2.9: Proportion of DNS over Total Load (Repair Order by Criticality)

with Community 3. This higher vulnerability score compels more immediate action to restore services in Community 3. Consequently, despite utilizing various parameters such as SVS, criticality, and a weighted combination of both in determining repair priorities, the repair time for WP13 remains notably shorter.

Figure 2.13 illustrates the average water storage levels over time for two water storage tanks designated for Community 2 and Community 3. When the pumps are fixed, we assume that the refill levels from both pipelines are 4500 liters of water at each time step.




Figure 2.10: Proportion of DNS over Total Load (Repair Order by SVS)




Figure 2.11: Proportion of DNS over Total Load (Repair Order by Weighted Criticality and SVS)

As the plot shows, WS3 consistently maintains higher water levels than WS2 throughout the observed period. This persistent disparity can be attributed primarily to the differential water demands from the communities they serve, where Community 3 requires 2100 liters, significantly less than the 3300 liters required by Community 2. During the period of pump failure and subsequent repair, the storage levels in both tanks demonstrate distinct trajectories. WS2 experiences a dramatic decline in water levels around the tenth time step, followed by a recovery. In contrast, WS3 shows a gradual decline and a steadier recovery.

The minimal variation in water storage levels, irrespective of the sequence of repairs de-

pending on whether criticality, SVS, or weighted SVS and criticality was used, underscores system robustness where consistent community demands and stable water inflow from unaffected pipelines mitigate potential fluctuations. This stability indicates that variations in repair sequences have a limited impact on overall water availability when demand and supply factors remain constant.

The bar chart in Figure 2.14 illustrates the ratio of total water available to total demand for two communities, evaluated across three distinct metrics: Criticality, SVS, and a weighted combination of both SVS and Criticality. Both communities exhibited similar patterns across the three metrics, suggesting a demand patterns play a more significant role than repair sequences in this scenario. Criticality and the combined SVS and Criticality metrics showed nearly identical values, aligning with the higher weight of criticality in the weighted metrics.




Figure 2.12: Average Time to Repair Each Node in 1000 simulations

2.6.3 Correlation of Repair Ordering Rankings

We employ Kendall's Tau to assess the correlation between rankings of repair orderings based on different metrics. Table 2.2 reports Kendall's τ correlation among the three total repair orders (SVS, Criticality, and weighted SVS and Criticality). As shown, the SVS ordering has a higher correlation with weighted SVS and Criticality ($\tau = 0.69$) than with Criticality alone ($\tau = 0.32$). At first glance, this might suggest that the weighted approach (SVS and Criticality) more closely resembles the SVS ranking at the individual node level. However, in




Figure 2.13: Average Water Storage Levels Over Time




Figure 2.14: Ratio of Total water over Total Demand for Each Community

our broader impact analysis (e.g., measuring demand not served or total time before repairs), the system-wide performance of weighted SVS and Criticality tends to track more closely with the Criticality approach.

The apparent contradiction stems from the distinction between ranking correlations across all nodes and system-level impact dominated by a handful of highly critical nodes. In calculating Kendall's τ , the Weighted ($0.4 \times \text{SVS} + 0.6 \times \text{Criticality}$) and SVS repair orders correlate strongly because many mid-priority nodes share relatively high SVS and moderate Criticality, resulting in similar within-layer sorting. However, at the impact stage, measured by average outage times and proportion of demand not served, both Weighted and

Table 2.2: Correlation of the Repair Order Rankings

	Total Order by SVS	Total Order by Criticality	Total Order by SVS and Criticality
Total Order by SVS	1	0.32	0.69
Total Order by Criticality	0.32	1	0.33
Total Order by SVS and Criticality	0.69	0.33	1

pure Criticality strategies focus heavily on the same very high criticality nodes (e.g., hospitals), leading to nearly identical macroscopic outcomes. Hence, Weighted’s partial emphasis on SVS raises its overall ordinal correlation with SVS, but its emphasis on high-criticality infrastructure drives performance metrics into alignment with pure Criticality.

Consequently, no error is indicated; rather, the data reflect two distinct phenomena: first, an overall similarity in how a large subset of nodes are ranked (hence the higher τ between SVS and Weighted), and second, a strong emphasis on high-criticality nodes that drive final system-wide performance into alignment with the repair by Criticality-only strategy.

2.7 Conclusion

In this chapter, we show the effect of repair prioritization on three communities affected by disaster. We show the effects of using criticality, social vulnerability score, and a combination of both in ordering repairs and its effects on community outcomes. We show that prioritization of repairs in the power infrastructure is important, and a combination of metrics may give a balanced approach to ordering repairs during outages. For the water infrastructure, we conclude that though ordering of repairs plays a part in equitable restoration, in the long run, the demand for water from communities affected surpasses any effects of this prioritization.

This work has several limitations. First, criticality values are assigned using a simplified hybrid approach that balances graph structure and domain knowledge. Future works need to

apply rigorous graph-based or impact-of-removal measures. Second, our model uses abstract timesteps rather than real-world time, so it does not capture detailed transient dynamics. Third, the case study focuses on a notional network, while our framework is conceptually scalable to larger interdependent networks. For us to apply it to real-world infrastructure, it will require further validation and computational optimization, which we leave for future work. Also, we do not model the geographic distribution of infrastructure or repair crew travel times. Future work will incorporate spatial layouts and dispatch delays to improve the practical realism of the recovery plan.

Chapter 3

HABiT: A Hetero-functional Agent-Based Infrastructure Toolkit for Modeling Community Resilience

3.1 Introduction

Critical infrastructure systems, such as power, water, and transportation, are deeply interdependent and operate under uncertain conditions that can change rapidly during disruptions¹. When failures occur, the consequences often cascade across systems, and affect not just physical assets but also the communities that rely on them^{3;4}. As a result, modeling these interdependencies is essential to understand how disruptions propagate and to design effective strategies to restore service⁵⁻⁷.

While existing agent-based infrastructure models attempt to capture adaptive behaviors, stochastic disruptions, and community-level responses, several limitations remain. Most ABM studies either represent only one or two infrastructure networks or integrate additional systems in a simplified, static manner⁵⁸⁻⁶⁰. This limited cross-system representation makes it difficult to explore dynamic scenarios involving multiple simultaneous disruptions and shifting interdependencies. Additionally, traditional ABMs typically adopt asset-centric

representations, which focus primarily on physical components rather than directly modeling the functionalities that communities depend upon^{60;61}. Consequently, these approaches struggle to represent scenarios that involve service substitutions, alternative delivery mechanisms, or dynamic reconfiguration during failures and recoveries. Furthermore, although recent works show a gradual introduction of some sort of social metrics, few provide a generalized framework for embedding diverse social and community-centric attributes directly within infrastructure models^{62;63}.

To address these gaps, our work introduces HABiT, a flexible agent-based simulation toolkit built around a hetero-functional graph (HFG). HABiT models interdependent functionalities rather than physical components alone, supports flexible addition or removal of system layers during runtime, and integrates social metrics such as Social Vulnerability Scores (SVS) and Criticality metrics directly into the infrastructure resilience assessment process.

This is a proof-of-concept work that demonstrates the feasibility of modeling interdependent infrastructure with an agent-based, stochastic hetero-functional graph. Our current prototype established the core mechanisms such as resource flows, cascading failures, priority-driven repairs, and stochastic response. Other key capabilities will be integrated in future works. Our work can be interpreted as evidence of validity, and ongoing extensions will refine the model and broaden its quantitative scope.

We apply HABiT to a synthetic network involving three communities with different infrastructure needs and social vulnerability scores. Each community interacts with shared infrastructure systems but experiences disruptions differently based on repair timing and access to critical services. The model supports both normal operations and recovery phases, which allows us to explore how different repair policies affect community-level outcomes and overall system resilience.

This study contributes to the broader effort of modeling infrastructure interdependencies in a way that balances realism with computational tractability. While the model is not calibrated to real data, it demonstrates key mechanisms that are relevant for resilience planning. Future work will extend this prototype to more realistic datasets, integrate policy constraints, and incorporate stakeholder-defined preferences in recovery planning. Our main

contributions are:

1. We develop HABiT, an agent-based simulation framework for modeling interdependent infrastructure using a hetero-functional graph structure. Each node in the graph represents an autonomous functionality with defined behavior, location, and evolving state.
2. We demonstrate HABiT’s capabilities through a synthetic three-community scenario, which analyzes how failures of key functionalities affect infrastructure performance and recovery. We also evaluate multiple repair prioritization strategies and load-shedding decisions under failure.
3. Our results reveal that although overall system recovery trajectories may appear similar under various prioritization policies, the choice of repair ordering significantly influences how unmet demands distribute across communities, which highlights important equity and fairness implications. This result points to the value of incorporating distributional objectives into infrastructure resilience planning and policy design.

The rest of this chapter is organized as follows: Section 2 reviews related work on agent-based simulations, interdependent infrastructure modeling, and hetero-functional graphs. Section 3 describes the methodology for this research. Section 4 presents the implementation details of our work. We discuss the synthetic network in Section 5. Section 6 discusses the simulation results. We discuss the implications of the result in Section 7, and in Section 8, we conclude with limitations and suggest future work.

3.1.1 Acknowledgments

This research was supported by funding from the USA National Science Foundation (NSF) under Grant Number 2148878.

3.2 Literature Review

We present our literature review focusing on four interconnected themes: agent-based modeling for infrastructure systems, looking into stochastic-based models in ABMs, and the use of social metrics in ABM decision making, and finally, on current research on hetero-functional graph theory.

Agent-based modeling (ABM) has been widely applied to infrastructure management due to its ability to capture localized decision-making and emergent behaviors. Macal and North^{64;65} have written multiple papers on agent-based modeling and simulation (ABMS) presented at the Winter Simulation Conference (WSC) from 2005 to 2018. These papers offer guidelines on how to set up ABMS, with details on how to choose the right concepts, design agents and their environment, and select implementation tools. They also provide guidance on scheduling model events, calibrating and validating models, and documenting the modeling process.

Recent studies have extended ABM frameworks to include both physical and social infrastructure layers. Esmalian et al.⁶¹ developed a multi-agent simulation framework linking hazard exposure, power infrastructure disruptions and restorations, and household behaviors under hurricane-induced outages. Their approach integrates empirical household survey data to capture community hardship and incorporates equity considerations and scenario analyses, such as effects from early warnings and social networks.

Jamal et al.⁶⁰ developed an ABM specifically linking electric power and transportation networks to study physical infrastructure dependencies. Their model simulates how disruptions in one network influence restoration efforts in another, evaluating various strategies using real hurricane data from Miami-Dade County, Florida. Sun et al.^{58;59;66} developed agent-based models to simulate post-earthquake recovery for interconnected community, power supply, and transportation systems. Their framework models behaviors of system operators and administrators, quantifying resilience through supply and demand imbalances and cascading recovery processes.

From a community perspective, Aghababaei and Koliou⁶³ proposed an ABM framework

focused on tornado disruption scenarios, modeling interactions among businesses, healthcare facilities, educational institutions, utility services, and households. Their work evaluates the impacts of different mitigation strategies on community resilience.

ABM has expanded into multilayer frameworks for systemic risk assessment. Dubanowski et al.^{67–70} developed agent-based models to assess systemic risks in complex sociotechnical networks, integrating infrastructure, business, and household agents. These studies introduce new metrics to measure resilience. Likewise, Xue et al.⁶² proposed PostDisasterSim, an ABM framework that integrates human behaviors with social and physical infrastructure layers following Hurricane Harvey. Their study used empirical data to reveal distinct recovery patterns by agent type, geographic location, and socioeconomic status.

Further integrating multiple system layers, Valinejad and Mili^{71;72} developed a cyber-physical-social ABM that integrates power systems, emergency services, social interactions, and cyberspace. Their model quantifies community social metrics such as anxiety, risk perception, and cooperation, validated using real hurricane data.

Several papers include stochastic elements in modeling infrastructure disruptions and recovery^{58;60;61;72}. Esmalian et al.⁶¹ and Jamal et al.⁶⁰ incorporate stochastic hazard exposures and probabilistic restoration processes. Valinejad et al.⁷² assess uncertainty in restoration outcomes considering community cooperation and wellbeing under stochastic influences.

Only a small fraction of studies incorporates equity or social vulnerability into ABM-based infrastructure resilience analyses. Esmalian et al.⁶¹ use household survey data to gauge hardship levels and social-network influences. Xue et al.⁶² explore recovery dynamics across household income brackets and locations, while Aghababaei and Koliou⁶³ model community, healthcare, and educational agents to capture key social dimensions. Yet most of the literature still centers on physical recovery metrics, service gaps, node criticality, and restoration time, rather than social impacts, and rarely embeds equity criteria in planning. To bridge this shortcoming, our work on HABiT adds vulnerability modeling by simulating communities with varying vulnerability, which brings social factors into resilience assessment and clarifies distributional impacts across groups.

Most agent-based work models one or two networks at a time, or treats extra layers only

in a simplified way. Sun et al.^{58;59;66} link community demand with electric power supply yet leave water and transport outside the model scope. Jamal et al.⁶⁰ show how power outages disturb road traffic, but still focus on a two-layer setting. This leaves no common platform to add or remove whole subsystems during analysis. HABiT loads a heterofunctional graph that already contains power, water, and transportation functions. Each function becomes an autonomous agent with its own behaviour, location, and state. Because the graph is read at start up, agents can be taken out of service or brought online when needed via dynamic links during run time, which provides practical adaptability without deleting objects from the graph. Also, users can revise the file to add or delete components and then rerun it, which makes it a flexible multilayer test bed.

Most resilience studies still build asset-centred graphs. Nodes are pipes, roads, transformers, and edges show physical links^{58–60;66}. Failures and repairs are counted at the component level, so results show how fast hardware is restored rather than how long power, water, or mobility stay available. When analysts want to test a workaround, such as moving water by truck after a pipeline break, they must redraw the network or write sector-specific code. This asset focus limits insight into cross-sector substitution and makes it hard to compare different layouts. HABiT turns the view around. Every node is a service functionality, and physical equipment is just an attribute. The same transport water functionality can be delivered by a pipeline, a tanker convoy, or any new mechanism added later, with no change to the code. One agent base class gives uniform behaviour across power, water, and transport, so cascades flow naturally from service to service. Metrics track service continuity and unmet demand, which lets analysts compare very different configurations on the same functional scale. By making services, not hardware, the primary objects, HABiT closes the gap between fixing components and delivering real service.

Previous literature on hetero-functional graphs, such as Schoonenberg et al.²⁰ and Munikoti et al.⁷³, has typically employed binary graph representations. These binary graphs primarily denote only the presence or absence of functional interdependencies between infrastructure components, without any system characteristics. However, our research extends the foundational binary HFG framework by introducing detailed node and edge attributes. We

introduce attributes that represent system characteristics such as capacity, resource availability, and dynamic operational states. Also, our HFG graph includes community-centric metrics of fairness and equity such as Social Vulnerability Scores(SVS)⁵¹ and criticality, with more opportunities for expansion as needed. The goal of this attribute-based hetero-functional graph is to provide a more nuanced tool for infrastructure resilience assessment.

3.3 Methodology

We present the Hetero-functional Agent-Based Infrastructure Toolkit (HABiT), a simulation platform for modeling interdependent infrastructure systems. HABiT combines agent-based modeling (ABM) with a hetero-functional graph (HFG) representation to simulate dynamic behavior and cascading dependencies across power, water, and transportation networks, as well as their interaction with community-facing functionalities. This section describes the conceptual modeling framework, underlying assumptions, and how the HFG and ABM layers interact to enable flexible, modular infrastructure simulation.

3.3.1 Modeling Approach

HABiT integrates agent-based modeling (ABM) with a hetero-functional graph (HFG) structure to represent both localized decisions and system-wide dependencies in interdependent infrastructure networks. Each node in the HFG represents a functionality, such as generating power, transporting water, or directing traffic, rather than a physical asset. Directed edges denote functional precedence; for instance, power must be generated before it can be distributed. This abstraction supports multi-layered interdependencies while keeping the model extensible and modular.

Within the ABM framework, each functionality is modeled as an autonomous agent that tracks its operational status, dependencies, and dynamic behavior. Agents activate based on upstream availability of required resources (e.g., power, water, vehicles) and respond to failures or recovery events. The agent-based layer allows for detailed simulation of local

rules, repair logistics, and emergent network behavior over time.

3.3.2 Model Assumptions

We design HABiT to simulate any interdependent infrastructure system represented as an attribute-based hetero-functional graph. The model reads a pre-processed input graph where nodes define functional roles and associated metadata, and edges define resource or service dependencies. Once loaded, each node is instantiated as an agent. If the functionality is known and implemented in code, no additional configuration is needed. Otherwise, the user must implement a new agent class with appropriate logic before simulation. With this flexibility, HABiT can be applied to systems with varying complexity. We made the following assumptions during the design of HABiT:

- Infrastructure behavior is governed by localized rules and the directed dependencies encoded in the graph.
- Each functionality is tied to a physical location and can fail or recover based on its own internal logic.
- The physical network does not rewire during the simulation.
- Demand for power, water, and vehicles is deterministic unless stochastic demand is enabled, in which case it varies at every time step.
- Each transport edge forwards at most its rated capacity per timestep.
- Failures can be random or scheduled based on disaster inputs.
- Repair times follow an exponential distribution with rate $\lambda = f(d)$ where d is path length. Demand noise is Gaussian with standard deviation $\sigma = \eta \bar{D}$.
- Repairs are carried out by mobile repair crews whose arrival time is determined by road network distance (travel time) and by the number of available crews.

- Simulations proceed in discrete time steps, and all agents update simultaneously at each timestep.
- The simulation stops after T_{max} time steps or when all tasks are repaired.

3.3.3 Hetero-Functional Graph Modeling

We adopt a hetero-functional graph (HFG) approach to model the interactions between infrastructure functionalities. Unlike traditional infrastructure modeling that focuses on physical components and topology, HFG captures system behavior in terms of interactions among *functional nodes*. Each node in our model represents an autonomous agent responsible for a specific functionality, such as power generation, water treatment, vehicle transport, or storage management.

We define the infrastructure system as a directed hetero-functional graph $G = (V, E)$, where each node $v \in V$ represents a functionality (e.g., generate power using X, consume vehicle at Y) rather than a static asset. Each edge $e = (v_i, v_j) \in E$ encodes a dependency relationship between functionalities, indicating that functionality v_j requires v_i to be operational. Each node $v \in V$ is assigned attributes such as:

- a resource label, e.g., "Transformer4" or "Hospital1",
- a spatial location $\ell(v) \in \mathbb{R}^2$,
- an infrastructure type $d(v) \in \{\text{power, water, transport}\}$,
- and optionally, community affiliation, criticality, and social vulnerability score (SVS).

The use of HFG provides several modeling advantages, such as:

1. Clear representation of functional dependencies: The dependencies between production, transport, and consumption are made clear via directed edges.
2. Modularity and flexibility: We can add new functionalities (e.g., 'enable service', 'repair resource') without altering the underlying graph structure.

3. Multi-infrastructure integration: Different infrastructure systems, such as power, water, and transportation, can be modeled in a unified framework in a single-layer graph.

However, HFG also introduces certain tradeoffs. The graph abstracts away physical layouts, which may hinder interpretability. Also, the approach remains relatively new and lacks widespread use.

3.3.4 Architecture

We build the Hetero-Functional Agent-Based Infrastructure Toolkit (HABiT) using the MESA agent-based modeling library in Python. The architecture defines each infrastructure function as an agent-based functionality class that inherits from a base `FunctionalityAgent` superclass. This base class includes attributes such as physical status, operational status, and outage duration. We load additional properties from the graph, which we discuss later. Figure 4.1 shows the HABiT architecture.

Functionality agents are initialized directly from the hetero-functional graph structure, where each node represents a functionality and each edge denotes a directional dependency. The simulation interprets these relationships to resolve activation, propagate flow, and trigger cascading effects. Each agent evaluates its state and dependencies during discrete time steps. Agents act on local information, which enables fine-grained control and behavior. This locality supports autonomous, decentralized responses to failures, congestion, or resource scarcity, allowing complex system-wide behavior to emerge from simple rules.

3.4 Implementation

Each infrastructure function in HABiT is implemented as a class that inherits from the base `FunctionalityAgent`. These subclasses override the update method to encode custom logic, such as capacity checks, flow distribution, or repair status. The model defines a distinct set of agents for each infrastructure domain.

We use a custom graph-driven scheduler that iterates strongly-connected components(SCCs)




Figure 3.1: The HABiT Architecture. This architecture shows the interdependencies across transport, water, and power infrastructures. Arrows indicate resource or control flow between agent types. Repair crews are dispatched via the transportation layer to failed components marked by RepairTasks. Resources are generated, transported, and consumed. EnableService evaluates the quality of service delivered to community-facing infrastructure such as hospitals, schools, and residential buildings.

in topological order. Each SCC is run in three passes where pass one flushes edge buffers from the previous time step into their downstream nodes, pass 2 lets every node compute its new state from those inputs and pass 3 writes the resulting outflows back into the outgoing edge buffers so they are available to other components and to the next simulation timestep. We chose to implement a custom scheduler because the standard Mesa schedulers step agents in an arbitrary or fixed sequence, which led to race conditions and deadlocks when applied to such tightly coupled SCCs. Because the scheduler decomposes the graph into SCCs and processes them sequentially, the scheduler enforces a causally correct order while still allowing parallelism across independent components. The simulation operates on a discrete timestep basis. This ensures that all agents evaluate their state, process updates, and interact with their neighbors synchronously at each step.

HABiT applies stochastic initialization by sampling vehicle availability from a multinomial distribution at the start of each simulation. In addition, repair durations for failed agents are drawn from an exponential distribution that reflects spatial distance to nearby residential nodes. All codes are available at [74](#).

3.4.1 Graph-Driven Scheduling

To coordinate interdependent infrastructure agents, we use a custom scheduler that executes the model’s directed dependency graph in strongly connected component (SCC) order. As shown in Algorithm 1, the scheduler first condenses the full agent graph G into a directed acyclic graph (DAG) of SCCs. It then performs a topological sort over the SCCs to determine a valid execution order. This ensures that, even if cycles exist within a subset of agents (e.g., water pipelines with feedback), execution still respects causality between higher-level components.

Each SCC is processed in three phases. In the first phase, edge agents move any residual flow from the previous time step. Next, node agents transform, generate, or consume resources, and finally, edges push newly generated flow to downstream nodes. After sweeping all utility-dependency SCCs, the scheduler runs any transient agents (such as completed

RepairTask).

Algorithm 1 GraphDrivenScheduler one-time step routine

```

1:  $C, \text{cmap} \leftarrow \text{CONDENSEToSCCDAG}(G)$ 
2: for all component  $S$  in  $\text{TOPOLOGICALSORT}(C)$  do
3:    $E \leftarrow$  edge agents in  $S$ ;  $N \leftarrow S \setminus E$ 
4:   for all  $e \in E$  do ▷ pass 1: move old flow
5:      $e.\text{STEP}()$ 
6:   end for
7:   for all  $n \in N$  do ▷ pass 2: transform or generate
8:      $n.\text{STEP}()$ 
9:   end for
10:  for all  $e \in E$  do ▷ pass 3: push new flow
11:     $e.\text{STEP}()$ 
12:  end for
13: end for
14: for all  $a \in \text{TRANSIENTPOOL}$  do
15:    $a.\text{STEP}()$  ▷ clean-up stray agents
16: end for

```

This design ensures that feedback cycles are resolved locally within each SCC while preserving the larger system’s acyclic execution semantics. It is particularly well-suited to interdependent infrastructures, where cycles may arise from, for example, water feedback loops or bidirectional transportation links, but higher-level system flows remain mostly feed-forward.

Efficiency Considerations Although the SCC condensation introduces an $\mathcal{O}(|V| + |E|)$ overhead each time step, the resulting execution plan ensures logical consistency and modular execution. In our model, which includes 157 SCCs, most components are small and modular. This fragmentation is not a limitation but reflects the decoupled nature of subsystems (e.g., edge-only vehicle routes or localized loops). Each SCC contains a small group of tightly coupled agents, which allows for fast intra-SCC scheduling and clean separation of concerns between functional regions.

Additionally, the three-phase pass inside each SCC helps to separate pre-existing flow from new decisions, making time-dependent propagation (such as vehicle and water transport) more accurate. Transient agents—those not wired into the static graph, such as com-

pleted RepairTasks, are updated afterward in a dedicated clean-up step, preserving correctness without polluting the main dependency flow.

Limitations This approach does not yet exploit fine-grained parallelism; all agents within an SCC are still stepped sequentially. In large-scale models, repeated SCC decomposition may become a bottleneck. Finally, the scheduler relies on NetworkX’s in-memory graph representation, which may limit scalability on extremely large graphs or when streaming data is required. Future work could explore incremental SCC updates or integration with out-of-core graph engines.

3.4.2 Power Classes

The power layer models the delivery of electricity from generation, through substation and transformers, to low-voltage distribution downstream. Four core agent classes capture these stages: GeneratePower, TransportPower, DistributePower, and PassPower. Together, they represent sources, capacity-constrained links, load-shedding transformers, and switches, respectively. We discuss their state variables and interaction logic in more detail.

- **GeneratePower:** This class represents a generator’s generation functionality. It produces a fixed amount of electrical power when operational. Its output is passed to downstream transport agents. The agent resets to zero output if it fails. GeneratePower agents optionally apply stochastic generation logic at each timestep. Power output is sampled from a distribution centered around the nominal capacity, normal to simulate operational variability such as fuel availability or intermittent generation. Its state variables are `power_capacity` and `power_output`.
- **DistributePower:** This class models the distribution capabilities of transformers and substations. Aggregates upstream power and applies load-shedding if demand exceeds capacity. It uses three kinds of metrics (SVS, Criticality, and Weighted Sum of SVS and Criticality) to rank downstream needs and distribute available power. It maintains an internal switch state per downstream line. The class also keeps a Boolean switch

per outgoing line that opens whenever allocated flow drops to zero. The Social Vulnerability Score (SVS) describes the vulnerability of the node. Usually, infrastructures do not have SVS, but we assigned SVS of communities via propagation to infrastructures tied to those communities. Criticality is a technical metric that defines how important a node is in a network. Weighted Sum of SVS and Criticality is user-defined and calculated as $\alpha \times SVS + \beta \times Criticality$ where α and $\beta = 1$ and are user-assigned. This kind of balances social versus technical priorities. Its state variables are power_capacity, aggregated power_demand from children nodes, and switches.

- **TransportPower:** This class represents the transport functionalities of powerlines. It receives power from upstream functionalities, enforces its capacity constraint(power_capacity), and distributes the resulting flow to downstream nodes. It updates its power output each step and can enter a failed state independently, cutting off the flow. It also has a switch that, when it is on, turns off the power. The agent monitors its physical integrity and may fail independently. When this occurs, it opens its internal switch, immediately cutting off power transfer. If downstream demand exceeds its capacity, it raises an overload flag, enabling upstream transformers to shed load accordingly. Its state variables are power_capacity, power_outflow, and switch.
- **PassPower:** The agents of this class act like an unlimited switch, such that whatever power comes in goes straight out, with no capacity limits or priority rules. Its only state is an on/off switch that upstream protection or repair tasks can open; otherwise, it just passes power through with no losses. We acknowledge that this is a design simplification and not a realistic mode of operation. But we did this to keep the electrical graph co-located with the road graph. We drop a "ValveIntersection" resource implemented with PassPower at every road intersection where overhead lines must follow street geometry. The extra node maps the line along a road edge without throttling flow or altering load-shedding decisions.

Algorithm 2 Priority-based Load Shedding in DistributePower

Require: demand map D_i , priority scores s_i , total capacity P_{cap} , load set L

Ensure: allocation map A_i

```
1:  $P_{\text{left}} \leftarrow P_{\text{cap}}$ 
2:  $A_i \leftarrow 0 \quad \forall i \in L$  ▷ initialise all allocations
3: for each  $i$  in  $L$  sorted by  $s_i$  (desc.) do
4:    $g \leftarrow \min(D_i, P_{\text{left}})$ 
5:    $A_i \leftarrow g$ 
6:    $P_{\text{left}} \leftarrow P_{\text{left}} - g$ 
7:   if  $P_{\text{left}} = 0$  then
8:     break ▷ no capacity left
9:   end if
10: end for
11: if  $P_{\text{left}} > 0$  and  $\exists i \in L : A_i < D_i$  then
12:    $U \leftarrow \{i \in L \mid A_i < D_i\}$  ▷ unsatisfied loads
13:    $s \leftarrow P_{\text{left}}/|U|$  ▷ equal share
14:   for each  $i \in U$  do
15:      $A_i \leftarrow A_i + s$ 
16:   end for
17: end if
```

3.4.3 Water Classes

The water system starts by treating raw water, moves it through pipelines with the help of pumps, stores it in tanks, and finally supplies customers. Because treatment plants and pumps need electricity, the water layer relies on the power layer, creating strong interdependence between the two networks. During simulation, when demand is higher than production, the model can split the limited flow using proportional and priority rules, just like in the

power layer. There are five classes covering the functionalities of the water layer. They are TreatWater, PumpWater, TransportWater, StoreWater, and PassWater.

- TreatWater: This class simulates the treatment of water functionality of a water treatment plant. Its agent starts each run with a partially filled storage tank. Each time step, if power is available, it first distributes treated water to operational downstream pipelines, splitting the volume in proportion to their demand, then produces fresh water up to production capacity to refill storage capped by storage capacity.
- PumpWater: This class simulates a simple pump functionality. The agent of the class buffers the water received from upstream. When both inflow and power are present, it pushes the buffered volume to downstream pipes, allocating flow according to each pipe's design capacity share. If either resource is missing, it flushes the buffer, sends zero, and tallies an outage time step.
- TransportWater: This class represents the transport functionality of water pipelines. Its agent is a pipe segment with fixed capacity and a valve. After buffering inflow, it runs by forwarding water to children using the global allocation rule (proportional or priority) while respecting its capacity. Closed valves or failures make the pipe report zero demand upstream.
- StoreWater: This class represents the behavior of storage tanks. It operates as a reservoir that receives inflow until the water level reaches capacity. Each timestep, it reports demand, so upstream pipes know how much more water they can accept. When empty or failed, it sends nothing.
- PassWater: This represents an unlimited capacity valve node used to anchor pipes to road intersections. It forwards all incoming water unchanged. Its only state is a valve flag that can be closed during repairs.

Water Sharing

When multiple pipes draw water from the same TransportWater segment, the model decides how to share the limited flow using one of three rules saved in the model.

- Proportional (default): Each downstream pipe gets the same fraction of its demand, scaled so that the total flow does not exceed the pipe’s capacity. For example, if the pipe can only carry half of what is requested in total, each consumer gets half of what they asked for.
- Greedy: Consumers are sorted by their demand, and the biggest ones are served first. Then water is given out starting from the largest demand down to the smallest, until there is no water left. Smaller consumers may get nothing.
- Priority: Each pipe is scored based on a weighted SVS and Criticality, where $\text{score} = \alpha \times SVS + \beta \times Criticality$. The pipes with higher scores are served first, just like in the greedy rule, but here the order reflects social value rather than just demand.

We use these to explore how different allocation strategies affect who gets water and how much, across the whole system.

3.4.4 Transportation Classes

This layer governs the flow of vehicles, which act as shared resources required by both repair crews and service agents. The total vehicle count is fixed at initialization and conserved throughout the simulation. Vehicles traverse one-way road segments and are dynamically reallocated across the network.

- GenerateVehicle: These agents act as depots that inject vehicles into the system. Each depot draws from the global pool and releases vehicles onto connected road segments, subject to available power, and optionally, water. Dispatch is constrained by a maximum transport capacity.

- **TransportVehicle:** These agents represent directed road segments. They accumulate incoming vehicles and forward them downstream within their capacity limits. Excess vehicles are retained in the queue, contributing to local congestion.
- **DirectTraffic:** This agent simulates an intersection controller. It gathers vehicles from upstream roads and redistributes them across outgoing links, up to a specified routing capacity. When multiple paths are available, it favors links with lower congestion, ensuring efficient load balancing without altering the global vehicle count.

Overall, this layer manages routing, queuing, and congestion dynamics, maintaining a constant vehicle population and enabling mobility across the infrastructure network.

3.4.5 Community-Related Classes

This group includes demand-side agents that consume electricity, water, and vehicles, as well as service-level aggregators that evaluate community well-being based on resource delivery.

- **ConsumePower:** This class represents a power-consuming facility, such as a household or service node. Each timestep, it checks whether the local vehicle buffer is sufficient and whether power is available from upstream sources. If both conditions are met, it computes a power efficiency ratio ($\text{received} \div \text{demanded}$) and increments an outage counter whenever power is unavailable.
- **ConsumeWater:** This class models nodes that require water, including residential or commercial consumers. It activates only if at least one vehicle is present and water inflow is available. The class incorporates stochastic demand variation around a nominal mean. It records water received, calculates efficiency, and tracks outage durations.
- **ConsumeVehicle:** This class represents a consumer-side vehicle buffer, such as a parking bay or delivery dock. It accumulates incoming vehicles and provides them to nearby power and water consumers. When empty, it induces vehicle-related outages. Vehicles are returned to the global pool after use.

- **EnableService:** This class represents a community-level service that depends on both power and water availability. At each timestep, it receives efficiency scores from its linked `ConsumePower` and `ConsumeWater` agents. Based on the lower of the two scores, it assigns a quality level (none, low, medium, or high) and updates a service score used in system-wide quality-of-life metrics.

3.4.6 Repair and Event Classes

This part of the model introduces failures, dispatches mobile crews, and restores functionality. It links the failure process to recovery resources. With this, we can analyze how the different repair policies, crew mobility, and road congestion shape restoration timelines.

- **EventManager:** We use the `EventManager` as the scenario controller responsible for introducing disruptions. It reads a JSON-based schedule of predefined failure events and, at each specified timestep, triggers those failures in the targeted agent. This includes disabling the agent, opening any associated switches, and generating a corresponding `RepairTask` tied to the affected functionality. It can also inject background faults at a user-defined rate to simulate random failures.
- **RepairManager:** The `RepairManager` coordinates the deployment of repair resources. At each timestep, it identifies all failed agents, ranks them using the current prioritization policy, such as SVS, Criticality, or a weighted combination. If multiple agents share the same value for the primary score, the model applies tie-breaking rules. For SVS-first policies, ties are broken using Criticality; for Criticality-first, ties fall back to SVS. After ranking, the top-priority task is assigned to the nearest idle crew based on shortest path routing. We compute and log both travel and repair durations for later analysis. Also, it tracks crew assignments, monitors ongoing repairs, and marks tasks as complete once the repairs are finished.
- **RepairCrew:** Each `RepairCrew` agent represents a mobile unit. Once dispatched, it follows the shortest path along the transportation network to reach the assigned Re-

pairTask. Travel time is determined by the number of road segments divided by a fixed traversal rate. Upon arrival, the crew performs repairs for a duration drawn from an exponential distribution scaled by path length, after which the target agent is restored and the crew becomes available for reassignment.

- **RepairTask:** This class represents a temporary point created whenever an agent fails. It shares the failed functionality location and is connected to the road network via a `TransportVehicle` agent. While open, it acts as a `ConsumeVehicle` sink that receives repair crew vehicles, and once the crew finishes their task, it flags itself as complete and is removed from the graph.
- **NetworkModel:** The `NetworkModel` is the top-level controller for the simulation. It loads the infrastructure network from a `GraphML` file, initializes all agents, and sets global parameters such as noise levels, prioritization rules, and the total vehicle pool. During each simulation time step, it advances the scenario by activating the event schedule, processing repair operations, updating all agents based on network dependencies, recording relevant metrics, and evaluating termination conditions.

3.5 Case Study: A Synthetic Multi-Infrastructure System of Three Communities

We demonstrate the capabilities of the HABiT framework using a synthetic infrastructure model comprising three interdependent communities. Each community is composed of critical infrastructure functionalities across power, water, and transportation domains, along with essential community resources such as residential buildings, schools, hospitals, and commercial areas.

The model is encoded as a hetero-functional graph, represented in a `GraphML` file where each node defines a specific functionality (e.g., power generation, water transport, traffic control), and each edge encodes a directional dependency between functionalities. This

design supports modular extension and allows new components or dependencies to be added without altering simulation logic.

3.5.1 Overview of Components

The model distinguishes between physical infrastructure resources and the functionalities they support. Physical resources are entities like substations, water treatment plants, roads, and buildings. Functionalities represent the operations performed by these resources and are implemented as simulation agents.

Resource Resources transform, transport, or store things or operands.

- Power System: Generator, Substation, Transformer, Powerline
- Water System: Water Treatment Plant, Pump, Storage Tank, Pipeline
- Transport System: Road, Intersection
- Community Resources: Residential, Commercial, Elementary School, Hospital

Functionalities Functions are the representations of processes that capture the capabilities of the resources. Each functionality is represented as a node in the graph, labeled with the format "action at/using resource". We extend the typical hetero-functional graph to become attribute-rich. Nodes include metadata such as spatial location, infrastructure type, social vulnerability score (SVS), and domain-specific capacities (e.g., water demand or power load). This setup allows for runtime activation logic, repair propagation, and interdependence resolution during simulation. Sample functionalities in the graph are:

- Generate power using Generator1
- Treat water at WaterTreatmentPlant1
- Distribute power using Transformer12
- Transport water from WaterStorageTank3 to Hospital1 using WaterPipeline7

- Consume power at Hospital1
- Consume water at ElementarySchool2
- Enable household services at ResidentialArea3
- Direct traffic at Intersection12

Operands Operands are inputs or outputs, that is, things that flow through the infrastructure and community. In our synthetic model, we have four operands which are power, water, vehicle and repair crew.

3.5.2 Graph Layout and Composition

Figure 3.2 illustrates the graphical layout of the three-community synthetic model, highlighting functional clusters, transport routes, and spatial organization of the resource nodes. the graph contains 406 nodes and 610 nodes (functionalities), 113 unique resources, 14 unique functionalities. The functionality nodes in the graph include a rich set of attributes that support simulation are in Table 3.1.

The synthetic model partitions the infrastructure system into three spatially and functionally distinct communities. These communities reflect diverse asset configurations and are encoded using binary flags in the graph metadata.

The network connects a single generator and a single water treatment plant to three communities via a mix of pumps, storage tanks, substations, transformers, switches, and road and pipeline networks. Each community includes a residential building, elementary school, commercial building, and hospital. These consumer nodes are connected to distinct sets of intermediate transport functionalities such as intersections, valves, and switching stations. The breakdown of the resources in each community are described below.

Community 1

- ResidentialBuilding1, ElementarySchool1, Hospital1, CommercialBuilding1




Figure 3.2: A Synthetic Network of 3-Community Interdependent Infrastructure Network. The diagram shows roads, power lines, and water pipelines connecting residential, school, and hospital buildings across three communities.

- Transformers (e.g., Transformer1, Transformer12), Pump1, WaterStorageTank3
- Roads and intersections: Road21-29, Intersection3, Intersection4, Intersection6

Community 2

- ResidentialBuilding2, CommercialBuilding2, ElementarySchool2

- Infrastructure components: Pump2, WaterStorageTank2, PowerLine17-25
- Roads and intersections: Road30-39, Intersection11-16

Community 3

- ResidentialBuilding3, Hospital3, ElementarySchool3
- Pump3, WaterStorageTank1, Transformer3
- Roads and intersections: Road41-49, Intersection21-26

Table 3.1: Node Attributes in the GraphML Model

Attribute	Description
resource	Name of the physical infrastructure entity hosting the functionality (e.g., Transformer1, Hospital3).
Functionality	Type of action the node performs, such as GeneratePower or ConsumeWater.
type	Functional role of the node (e.g., Generator, Consumer, Transport).
infrastructure	Domain classification: Power, Water, or Transport.
location	String of spatial coordinates in the format x,y.
x, y	Parsed numerical coordinates used for layout or distance calculations.
index	Internal indexing used for node ordering.
SVS	Social Vulnerability Score, used to prioritize repairs or load shedding.
Criticality	A normalized value representing the node's systemic importance.
community	Boolean flag indicating whether the node belongs to a community layer.
primary_community	Integer label for the primary community assignment (e.g., 1, 2, or 3).
Community_1, Community_2, Community_3	Binary indicators for whether a node belongs to each community. Supports overlapping memberships.
Power_Capacity	Capacity of power generation, distribution, or transport elements.
Power_Total_Load	Total power demand for consuming nodes.
Pipeline_Capacity	Flow limit for water pipelines.
Production_Capacity	Maximum water output at treatment plants per timestep.
Storage_Capacity	Buffer capacity for water storage tanks.
Water_Total_Load	Water demand for consumers such as buildings or hospitals.
Transport_Capacity	Capacity for transporting vehicles along a road segment.
vehicle_demand	Maximum number of vehicles needed per timestep by a consumer.

Criticality and SVS Propagation To assign SVS and Criticality scores across the network, we propagate values from the consumer functionalities they support downstream. The method is the same for both metrics and applies at the level of functionalities, not physical nodes. We begin at the power generation functionality of Generator1 and perform a depth-first search through the hetero-functional graph, moving downstream but stopping the walk when we reach any consumer functionality, such as consuming water, power, or vehicles at hospitals, schools, or residences. This produces a trimmed tree that includes all directed paths from the generator to consumer targets, without traversing further once a terminal consumer functionality is reached.

We then work bottom-up through this trimmed functionality tree. Each intermediate functionality copies the highest SVS or Criticality score among its immediate children, provided at least one of those children is a consumer. This propagation stops at any functionality that represents power generation, power distribution at substations, or water treatment, since these are considered supply-side sources and are excluded from inheriting downstream scores. These functionalities retain their original Criticality scores and are excluded from SVS propagation. This enables upstream prioritization not just by structural position but by the societal vulnerability or infrastructure importance embedded in the consumer functionalities they connect to.

Service Metrics We track a set of metrics that quantify both the availability of services and the proportion of demand that is met across the network. We use this to evaluate system performance and equity of recovery. These metrics are computed at each timestep and aggregated for analysis over time.

- **Service Score:** This is a score that represents the operational quality at the node and is used for system-wide and community-level aggregation. Each EnableService agent represents a community-facing facility such as a hospital, school, or commercial building. These agents assess upstream delivery from ConsumePower and ConsumeWater nodes. For each resource type, efficiency is defined as the ratio of received to demanded

quantity:

$$\text{efficiency} = \frac{\text{resource received}}{\text{resource demanded}}.$$

Based on these efficiencies, the EnableService agent assigns itself a service score according to the following rule:

- If both power and water efficiency ≥ 0.9 , then score = 1.0 (high quality).
 - If both ≥ 0.6 , then score = 0.66 (medium quality).
 - If both ≥ 0.3 , then score = 0.33 (low quality).
 - Otherwise, score = 0.0 (no service).
- Average Service Score: At each timestep, we compute the system-wide average of service scores across all EnableService agents:

$$\text{AvgServiceScore} = \frac{1}{N} \sum_{i=1}^N \text{service_score}_i,$$

where N is the number of service agents. This metric summarizes global service performance at each point in time.

- Average Service Score per Community: Each EnableService agent is tagged with a corresponding community label. Community-level scores are computed by restricting the average to agents within each community.
- Power Demand Not Served (%): This metric tracks the shortfall in power delivery across the network. It is calculated as:

$$\text{UnservedPower}\% = \left(1 - \frac{\text{TotalPowerServed}}{\text{TotalPowerDemand}}\right) \times 100.$$

Demand and served values are aggregated from all ConsumePower agents.

- Water Demand Not Served (%): Similarly, we compute unmet water demand as:

$$\text{UnservedWater}\% = \left(1 - \frac{\text{TotalWaterServed}}{\text{TotalWaterDemand}}\right) \times 100.$$

This value captures how disruptions in the water layer affect consumers.

3.6 Simulation Results

3.6.1 Scenario A

In this scenario, we reduce the available generator capacity and observe how the system prioritizes power distribution under three different priority rules: SVS, Criticality, and Weighted. Generator1 is the only supply source to Substation1 and all downstream loads. We reduced its capacity from 600kW to 250kW below the total system demand to simulate a load-shedding situation. The demand is from consumers is stochastic across 100 runs. Our goal is to evaluate how different prioritization rules affect power distribution and the resulting demand-not-served (DNS) across communities and consumer types.

For each run, we use the same network structure, stochastic demand levels, and the same failure conditions. We only change the prioritization policy. SVS should prioritize based on the consumers' social vulnerability score, criticality prioritizes based on the infrastructure's operational importance, and the Weighted rule combines both in different ratios, where α is 0.6 and β is 0.4.

Results

Figure 3.3 compares the three load-shedding policies (SVS, Criticality, and Weighted, where α is 0.6 and β is 0.4) at the community level when a generator supply of power is lower than demand from loads under stochastic demand. Under the SVS rule, Community 3 and Community 2 exhibit the lowest fraction of unserved load. This outcome matches expectations, since Community 3 (with SVS = 0.7) is the most vulnerable and thus receives

top priority, and Community 2 ($SVS = 0.3$) still benefits relative to Community 1 ($SVS = 0.2$). By contrast, when using the purely Criticality-based policy, Community 1 and Community 3 both show much lower unserved demand than Community 2; critical assets in Communities 1 and 3 (specifically Hospital 1 and Hospital 3) compel the algorithm to allocate available capacity to those communities, effectively shedding more of Community 2's load. The Weighted policy mirrors this same behavior, because when combining SVS and Criticality in our weighting, the critical hospitals in Communities 1 and 3 dominate the ranking, such that Community 2 again suffers the worst service.

Figure 3.4 provides a breakdown that explains these community-level trends. In every community, the Criticality and Weighted rules prioritize the hospital above all other nodes, forcing Elementary and Residential buildings into deeper cuts. This is particularly visible in Community 2, where the elementary school and residential block are nearly completely unserved. In contrast, the SVS policy shifts priority toward Community 3's buildings first (since $SVS = 0.7$ places those facilities at the top), then to Community 1 buildings, and lastly to Community 2. As a result, Community 3's load remains largely served at the building level under SVS, driving its low community-level unserved percentage. In short, SVS concentrates restoration on the most vulnerable community (Community 3) and then on Community 2, whereas Criticality (and consequently Weighted) funnel capacity almost exclusively to hospitals in Communities 1 and 3, which explains why Community 2 fares worst under those two policies.

3.6.2 Scenario B

In this scenario, we repeat the same scenario as Scenario A. The only difference is that we run this with deterministic demand; here, all consumers draw fixed, deterministic power loads across runs. The weights still remain α is 0.6 and β is 0.4.




Figure 3.3: Community-level DNS under three load-shedding policies (SVS, Criticality, Weighted) when a generator’s capacity generation is reduced and less than demand; demands are stochastic.

Results

Figure 3.5 and Figure 3.6 compare the three load-shedding policies (SVS, Criticality, and Weighted) at the community and building level, respectively, when a generator supply of power is lower than demand from loads under deterministic demand. We start by viewing the results from the building level. Under deterministic demand, each building’s fixed load means that SVS and Criticality alone determine allocation. In Community 1 ($SVS = 0.2$), SVS-based shedding cuts Hospital 1 entirely, whereas both Criticality retains some, and the hybrid Weighted approach meets 100% of its demands. In Community 2 ($SVS = 0.3$), SVS prioritizes Commercial Building 2 and Residential Building 2 over the Elementary School 2, while pure Criticality protects sheds most of the consumers, forcing Commercial Building 2, Elementary School 2, and Residential Building 2 into near-total blackout. The Weighted policy behaves almost similarly to Criticality. In Community 3 ($SVS = 0.7$), weighted is the best, meeting 100% demands of all nodes, SVS comes in second, while Criticality is the worst overall. However, criticality comes second in meeting the demands of Hospital3, which




Figure 3.4: Building-level breakdown of DNS for each community (SVS = 0.2, 0.3, 0.7). Each bar shows the fraction of unmet demand under SVS, Criticality, and Weighted policies. The scenario is under stochastic loads.

makes sense because of its high criticality value.

When we viewed the results at the community level, Weighted-based shedding tends to minimize overall unmet demand compared to Criticality or SVS. In Community 1, SVS leaves more total demand unserved than the other policies because deprioritizing the hospital drives excess blackout, whereas Criticality holds the community shortfall to about one-third, and Weighted meets all demands. In both Community 2, SVS generates the lowest community-wide shortfall (around one-quarter to one-third of total load), while Criticality (and, to a lesser degree, Weighted) concentrates capacity on single critical nodes and leaves most other consumers dark. Because loads never fluctuate, Weighted behaves like Criticality when hospital importance dominates and like SVS when vulnerability is strongest. In Community 3, weighted meets 100% of demands while Criticality leaves most unserved, followed by SVS.

Because loads never fluctuate, the weighted rule behaves like pure Criticality whenever hospital presence dominates and like pure SVS when vulnerability dominates. The result is extreme: Weighted achieves perfect service in the two hospital communities yet almost completely sacrifices the community without a hospital. SVS, in contrast, never reaches 100




Figure 3.5: Community-level DNS under three load-shedding policies (SVS, Criticality, Weighted) when a generator’s capacity generation is reduced and less than demand; demands are deterministic.

% service anywhere but avoids a total blackout in every community, producing the most balanced outcome across the three deterministic runs.

3.6.3 Scenario C

In this scenario, we simulate the failure of the three pumps’ functionalities, connecting the TreatWater functionality to the 3 StoreWater Functionalities. We simulate these three pumps failing simultaneously at time step 3. We run 100 Monte Carlo replications. Two repair crews are available. Each crew’s travel and on-site repair durations are drawn from $\text{Exp}(\lambda = 0.10)$. Our objective is to see the impact of the four different repair policies, which are SVS (with Criticality tie break) and Criticality (with SVS tie break), and Weighted with $\alpha = 0.6$ and $\beta = 0.4$.




Figure 3.6: Building-level breakdown of DNS for each community (SVS = 0.2, 0.3, 0.7). Each bar shows the fraction of unmet demand under SVS, Criticality, and Weighted policies. The scenario is under deterministic loads.

Results

Figure 3.7 presents the average service score for the entire system following the simultaneous failure of three pumps, with repairs prioritized using three different policies: Criticality, SVS, and Weighted. This figure represents overall system recovery, computed by averaging the service score across all service functionalities. The three recovery curves are identical, with full restoration occurring at roughly the same time across all policies. This similarity is largely driven by the fact that all pumps are repaired within a narrow time window, so no policy outperforms the others.

We then examine recovery at the community level. Figure 3.8 presents the average service score for each community. We observe that while the repair order shifts depending on the policy, all three curves show a similar pattern. Their failure causes a sharp collapse in service, followed by a staggered recovery that ultimately brings all communities back to full service. The difference is which community recovers first. Under SVS, Community 1 (C1) recovers last, while under Criticality and Weighted, it is Community 2 (C2). Despite these differences in ordering, all communities eventually return to full service.

The impact of this ordering is clearer in Figure 3.9, which shows the percentage of total

water demand not served in each community. Whichever pump is fixed last leads to the highest loss for that community. For example, C1 suffers the most under SVS, while C2 sees the highest losses under Criticality and Weighted. The Weighted policy slightly balances the trade-off, but the differences remain small. Table 3.2 explains this result where the pump repaired last in each policy is directly linked to the community that experiences the highest unmet demand. It shows Pump 1 under SVS, Pump 2 under Criticality and Weighted.

This result highlights that while the repair order shapes short-term recovery and determines which community experiences the longest outage, it does not substantially alter long-term system performance. When failures occur simultaneously and repair resources are sufficient, prioritization affects who recovers first, but not the overall speed of system recovery.

These results underline a key policy tension. A decision-maker who looks only at a system-level resilience curve might conclude that all three strategies are equivalent, yet the choice of priority metric sharply alters which community is disadvantaged. If the objective is to minimise the maximum community shortfall to protect critical lifelines such as the hospital in C3 then the repair rule must be selected with that distributional goal in mind, not merely on the basis of aggregate recovery speed.




Figure 3.7: Three Time series of Average Service Score showing System-Wide Service recovery curve under three pump failure and repair scenarios, aggregated over 100 simulation runs. Each plot shows the results under a different repair prioritization policy.




Figure 3.8: Three Time series of Average Service Score comparing service recovery for 3 communities under three pump failure and repair scenarios, aggregated over 100 simulation runs. Each plot shows the results under a different repair prioritization policy. C1, C2, C3 represents Community1, Community2 and Community3 respectively.

Table 3.2: Average Repair Time (in Steps) for Each Pump by Scenario

Scenario	Pump1 (N=100)	Pump2 (N=100)	Pump3 (N=100)
SVS	36.50	21.90	21.75
Criticality	24.09	37.77	21.05
Weighted	23.53	36.62	21.83

3.7 Discussion

Across the three scenarios, a consistent picture emerges. The network’s aggregate resilience curve often masks important distributional effects that unfold underneath.

When the generation capacity of a generator is less than demand, load shedding occurs. The load-shedding results show that while all three prioritization rules keep the system-wide unmet demand within a narrow range, they do so by distributing that shortfall differently across communities. The SVS rule protects the most vulnerable community first, often at the expense of others. Criticality shifts supply toward infrastructure like hospitals, no matter where they are, leading to a different pattern of who gets served and who doesn’t. The weighted rule moves between the two, it acts like SVS when vulnerability dominates and like Criticality when infrastructure ranks higher. This trade-off matters. A decision-maker looking only at system-wide averages may believe all three rules perform equally well, but on the ground, they produce very different experiences. One community may be fully served while another is left with little or no power or water. These differences in distribution are




Figure 3.9: A Bar Chart of Water Demand Not Served per community under three pump failure and repair scenarios, aggregated over 100 simulation runs. Each colored bar represents results under a different repair prioritization policy.

what make the choice of priority rule consequential.

In the pump failure scenario, we see that system-level recovery curves remain nearly identical regardless of the repair rule. This is expected since all three pumps are fixed within a tight window, and there are enough repair crews to work efficiently. However, the recovery paths for each community tell a different story. The order in which pumps are repaired directly determines which community recovers first and which one suffers the longest outage. The link is clear: the community served by the pump that was last repaired consistently sees the highest share of unmet demand. The impact is temporary. Eventually, all communities return to full service, but the period of disadvantage matters, especially when critical facilities like hospitals are affected. This highlights a tension in resilience planning: fast system-level recovery does not mean fair recovery. The choice of repair order creates a short-term burden that falls hardest on whichever community is served last.

A final observation concerns how the three local ranking rules shape community-wide outcomes. With deterministic loads, an SVS-first rule spreads the reduction in proportion

to social vulnerability and therefore produces the smallest total demand not served across the three communities. A pure Criticality rule achieves its intended goal, full protection of the hospital loads, but does so at the cost of almost complete blackout in the non-hospital community, amplifying inequities among other consumers. The weighted formulation offers a tunable compromise: by choosing larger or smaller Criticality weight (β) the operator can shift the solution continuously from “protect every hospital at any cost” toward “spread cuts by vulnerability.”

3.8 Conclusion

This study shows that while all three prioritization rules, SVS, Criticality, and Weighted, can restore full service to the network in roughly the same amount of time, they differ substantially in how recovery is distributed across communities. When multiple failures occur and resources are limited, the rule used to allocate repairs or shed load determines which groups bear the brunt of the outage. SVS protects the most socially vulnerable. Criticality secures infrastructure. The Weighted rule offers a tunable middle ground but still depends on how the two scores are combined.

These results suggest that resilience assessments must go beyond system-wide averages. A fast recovery curve does not answer the deeper question of who recovers first and who waits longest. Decision-makers need to clarify whether their priority is minimizing total unmet demand, protecting critical services, or distributing recovery more equitably. Each rule reflects a different objective and leads to different outcomes. A network can be fully restored and still fail its most vulnerable population if the wrong repair policy is used.

There are several directions to expand this work. First, the current model assumes fixed travel and repair times drawn from an exponential distribution. Introducing more realistic routing delays, congestion effects, or blocked access could shift how the repair order plays out. Second, this model already includes interdependence, for example, pumps and water treatment plants require power to operate. Future work could explore more complex interactions, such as power systems depending on water, or cascading failures that move

between layers. Third, this study used synthetic values for SVS and Criticality. Future work could integrate real community data or co-design these metrics with local stakeholders to reflect realistic priorities. Fourth, it would be valuable to test these rules under extended disruptions, partial observability, or feedback-enabled demand responses to understand how they perform in more dynamic and uncertain environments. There is also a need for more stochastic representation beyond random demand, including variability in failure timing, repair success, and network conditions. These additions would better reflect real-world uncertainty and system complexity. Finally, there is a need for external validation using a cosimulation framework. Coupling this agent-based model with established infrastructure simulators would allow for more rigorous testing of assumptions, calibration against real system dynamics, and assessment of whether the observed patterns hold under more physically grounded representations.

Chapter 4

LLM-Based Community Surveys for Operational Decision Making in Interconnected Utility Infrastructures

4.1 Introduction

Natural disasters often damage multiple critical infrastructures across communities, which creates complex prioritization challenges for response efforts. Traditional approaches for prioritizing infrastructure repairs rely heavily on expert opinions, which potentially overlook diverse community perspectives and public demand urgency. A recent study by Doorn et al⁸ shows that restoring damaged infrastructure during disasters following the technical approach is likely ineffective in addressing the immediate needs of affected communities. Also, studies have shown that vulnerable communities have higher immediate needs and are most likely to be affected if such needs are not met^{9;10}.

Due to this, the current research trend is focused on integrating the social needs of communities into the technical aspects of infrastructure restoration to achieve a fair allocation of limited resources during disasters^{75–78}. Building upon these efforts, this chapter proposes additional prioritization schemes that incorporate community perspectives, thereby com-

plementing existing approaches and enhancing the overall fairness and effectiveness of the infrastructure recovery process. We leverage the use of Large Language Models (LLM) to simulate diverse community opinions and generate synthetic data on the community’s preferred orders of repair. Afterward, we use a ranking algorithm to develop a total order for prioritizing infrastructure repairs in communities.

While fairness may often be treated as a technical goal, for instance, like ensuring access to all or prioritizing the most vulnerable, it is also a socially constructed and context-dependent concept. What a community may perceive as fair may differ based on their values, lived experiences, or crisis-specific needs. In using LLM, we heavily rely on our prompt design to encode assumptions about fairness, urgency, or moral trade-offs. This might reflect in our results, such that prompts that emphasize vulnerability may lead to prioritizing high-risk areas, and prompts that focus on service restoration might focus on high-density or populated areas. This means that the fairness in our results is not an objective truth rather a reflection of the values embedded in our prompts. Hence, we acknowledge that prompt design may influence what the model learns to treat as fair.

This study serves as proof of concept in incorporating community preferences and priorities into recovery and restoration after a natural disaster or hazard, demonstrating how LLM-generated preferences can complement infrastructure prioritization models. In future work, we plan to validate this approach using real community surveys or explore using prior survey data to generate synthetic preferences for new contexts, particularly where survey response rates are low. Our contributions include the following:

1. We introduce a hetero-functional graph of a toy infrastructure model comprising power, water, and community infrastructures. We use this graph to establish a partial order showing the sequence of repair decisions to be made during a disaster or post-disaster across affected communities.
2. We use a Large Language Model (LLM) as a survey tool and obtain community preferences as ordered pairs.
3. We use a neural network-based ranking algorithm to obtain a total order of repair

prioritization, which is a consensus of the community preferences.

The rest of this chapter is organized as follows: Section 2 reviews related work on the use of hetero-functional graphs, infrastructure repair prioritization, the use of LLM in surveys, and order learning. Section 3 describes the proposed methodology, including the LLM-based simulated community preferences, the comparator and chainization model, random subset sampling, and the sensitivity analysis. Section 4 describes the data collected during the experimentation. Section 5 presents experimental results and key findings. Finally, in Sections 6 and 7, we discuss the findings in our research and conclude the chapter, outlining directions for future research.

4.1.1 Acknowledgments

This research was supported by funding from the USA National Science Foundation (NSF) under Grant Number 2148878. This chapter is based on the paper¹⁸, which has been accepted for publication.

4.2 Literature Review

Hetero-functional graphs (HFGs) have been employed to model complex interdependent infrastructure systems, particularly in the context of smart cities and utility networks^{20;73;79}. For instance, Schoonenberg et al.²⁰ developed a comprehensive HFG model to represent interdependent smart city infrastructures, demonstrating its applicability in capturing the multifaceted interactions among urban systems. Similarly, Munikoti et al.⁷³ utilized HFGs to assess the robustness of interdependent urban utility networks, providing insights into system vulnerabilities under various failure scenarios. Despite these advancements, the application of HFGs in disaster recovery, specifically for prioritizing infrastructure repair, remains unexplored. Our work extends HFG modeling into this domain, applying it as a tool for extracting partial orders of functionality during disaster events. This allows us to formalize system-level dependencies and provide structure to otherwise ambiguous recovery decisions.

Recent research has explored how Large Language Models can be used in persona generation and generating opinions in surveys or stand-ins for community feedback. This approach becomes more important, especially when real-time or large-scale surveys are difficult to execute. These studies lay the groundwork for synthetic survey and preference generation using LLMs. Dolant and Kumar⁸⁰ proposed a multi-agent LLM in which simulated personas assume distinct stakeholder roles to support decision discourse under disaster conditions. They explored tradeoffs in decision-making and in strategy generation by using a flood scenario in a MidWestern township where personas exhibit evolving priorities. Li et al.⁸¹ conducted large-scale empirical evaluations of LLM-generated personas for simulating public feedback, including in use cases such as U.S. election forecasting. Their study critically identified systematic biases and deviations from real-world data, which they attribute to ad hoc generation techniques. Shi et al.⁸² used LLM-generated personas in structured multi-perspective debates to simulate community perspectives on complex topics. Though not explicitly survey-based, the study showed how synthesized viewpoints can approximate population-level opinion diversity and influence the ranking of decisions by exposing users to multiple attitudes and concerns. These papers demonstrate the potential of LLMs to produce community inputs through simulations that can substitute for or augment surveys.

Xie et al.⁸³ presented WildfireGPT, a multi-agent LLM system that tailors decision support to different stakeholders during wildfire scenarios. In this study, community-based feedback is used to inform risk reports via retrieval-augmented LLMs. Alqitham⁸⁴ introduced a system that uses real-time public sentiment from social media along with reinforcement learning to improve equity in disaster resource distribution. In this work, community feedback is ranked and weighted using optimization to guide the allocation of utility resources such as water, electricity, and medical aid. Chen et al.⁸⁵ designed a decision-support system that combines LLMs with multi-objective optimization algorithms and reinforcement learning for dispatch and resource allocation in urban emergencies. In this work, user feedback is collected iteratively and used for personalized allocation in urban emergency response. Chen et al.⁸⁶ built a structured emergency system using knowledge graphs to support LLM reasoning during crises. Though these works are not focused on graphs of interdependent in-

frastructures, they applied LLM-based systems using community feedback to guide real-time or resource allocation decisions in disaster management.

As LLMs are used more often to simulate people, concerns about reliability and fairness have increased. Li et al.⁸¹ pointed out that ad hoc persona generation can introduce systemic biases. They argue that a lack of control procedures can cause downstream distortions, especially in models meant for decision support. Shi et al.⁸² tested whether synthetic views influence human judgment by tracking eye movements and engagement metrics. The study raises critical concerns about the power of synthetic viewpoints to shape judgment.

These studies emphasize the need to evaluate synthetic feedback carefully. Our approach addresses this by performing sensitivity analyses by varying the prompts to the LLM. This allows us to detect and explain potential preference bias in the LLM-generated rankings.

Recent works show that infrastructure repair and investment prioritization rely on both technical and social criteria. Technical factors are usually based on asset condition, efficiency, cost-benefit, and performance objectives, while social factors include equity, social vulnerability, and stakeholder input^{87–95}. These solutions often use tools like multi-criteria decision analysis (MCDA) and analytic hierarchy process (AHP) to combine different types of criteria into a final score or ranking.

Even though social criteria are included, technical factors tend to carry the most weight and usually drive the final decision. Pramesti et al.⁹⁶ show that while economic and social factors are considered, they are limited in scope and mainly reflect the views of stakeholders and regulators. Karasneh and Moqbel⁹⁴ find that technical factors dominate, while the influence of socioeconomic and environmental factors is minor. Das and Nakano⁸⁹ operationalize social impact metrics alongside technical ones, but actual community participation is not included. Bauer et al.⁹⁷ show that although customer satisfaction and external benefits are considered, community preferences are not directly built into the decision-making process for the annual rehabilitation programs in drinking water networks. The focus remains on technical performance and service delivery.

Some works show that including social aspects can shift the rankings when done well. The study by Mohamadiazar et al.⁸⁷ shows that incorporating social and hazard vulnerability

changes which bridges are seen as most critical. In Dhansinghani et al.⁹³, both stakeholder input and technical criteria are used. One result of stakeholder involvement was that separating cost from the priority score gave more weight to social goals like walkability and school access. Dell’Anna et al.⁹¹ also show that adding vulnerability-related criteria can re-rank urban regeneration projects compared to a technical-only approach. Although they do not collect community surveys, they still reflect equity concerns through weighted criteria.

Although many prioritization methods include social criteria, most do not gather or use community preferences to guide final decisions. We observe that technical scores still take the lead in most cases. In some works where social criteria are weighted carefully or when stakeholder input is added, it affects the priorities being studied. In our work, we focus on cases where there is no technical total order, which is when functionalities can be repaired in any order. We begin with a technical prioritization, which is the partial order provided by the hetero-functional graph, and resolve the remaining ambiguity by including community preferences from LLM simulations.

Recent studies in order learning focus on how to create a full ranking from a limited number of pairwise comparisons. Some of these approaches treat the problem like matrix completion, where missing preferences are filled in based on patterns in the data. For example, Park et al.⁹⁸ model preference data as a low-rank score matrix and show that a small number of comparisons can still give good results. Gunasekar et al.⁹⁹ go further by recovering the full ranking from a partial order (represented as a directed graph) without needing exact scores, using a method called nuclear-norm regularization.

Other researchers use semi-supervised learning, which combines labeled and unlabeled data. Szummer and Yilmaz¹⁰⁰ add a graph-based regularizer that ensures similar items get similar ranks. More recently, neural network models have been applied to this problem. He et al.¹⁰¹ propose GNNRank, a graph neural network that learns from comparison graphs and penalizes incorrect rankings, achieving strong results even with limited data. Lee and Kim¹⁰² introduced the chainization method, which builds a full ranking by starting with a partial order and adding new “pseudo” comparisons to improve training.

We build on these ideas by adapting chainization¹⁰² to infrastructure planning. We start

with pairwise preferences generated by LLMs, train a model to compare options, and then use pseudo-pairs to improve the ranking. This helps us build a complete repair priority list even when we cannot collect all possible comparisons from the community.

4.3 Methodology

In this section, we discuss the approach of modeling infrastructure dependencies and generating a partial order using a hetero-functional graph representation. Next, we discuss the use of Large Language Models (LLMs) to generate community preferences. We describe the aggregation of these preferences and, finally, the machine learning-based chainization approach for ordering functionalities to produce a total repair order.

4.3.1 Hetero-Functional Graph (HFG) Model

We use a hetero-functional graph to represent the infrastructure system and its interdependencies. An HFG allows us to model different types of functionalities, such as generation, transportation, storage, and consumption of operands like power, water, and vehicles. Each functionality is represented as a node in the graph, and directed edges capture the dependency relationships between them.

The HFG includes infrastructure functionalities across multiple communities. These communities are linked to the infrastructure nodes through consumption and service-related functionalities. We use this structure to capture both physical infrastructure flow and its connections to community-level services. Figure 4.1 shows three distinct communities we use in this study, each featuring dedicated subsystems for electrical power distribution and water supply. The hetero-functional graph used in this study is publicly available at [103](#).

From the HFG, we extract a partial order of functionalities. This partial order tells us which functionalities must be repaired before others based on system dependencies. A total order, on the other hand, is a complete ranking of all functionalities using some metric or preference criteria. In the final layer of the graph, where consumption nodes like “Consume




Figure 4.1: The Integrated Infrastructure and Community Network Model

Water in Residential Area 1” or “Consume Power in Residential Area 2” are located, we notice that many of these do not depend directly on each other. That means they can technically be repaired in any order as long as their upstream dependencies are already working. Figure 4.2 shows the bottom part of the HFG used in our model. The final layer contains functionalities directly tied to community-level services, where technical prioritization is not sufficient.

In our study, we focus on the fact that consumption functionalities may stop working

at the time of disaster because their associated upstream functionality, such as transport power to Residential Area 2 using PowerlineX or transport water to Commercial Area 1 using WaterPipelineX, may be damaged. Therefore, the consumption functionalities depend on repairs being made to their upstream dependencies before they come on. Once those dependencies are fixed, the consumption nodes become available. Since consumption functionalities do not depend on one another, the model gives no guidance on which upstream dependency should be repaired first.




Figure 4.2: A subset of the hetero-functional graph (HFG) showing dependencies across power and water systems. The final layer consists of consumption functionalities that can be repaired in any order. Arrows represent functional dependencies. The highlighted functionalities are damaged while the links marked $\times$ are paths requiring repair before the associated community-level services (bottom layer) can be restored.

This creates ambiguity when resources like repair crews are limited. The technical model tells us what is possible and necessary, but it does not tell us what should take priority when multiple actions are equally feasible. To solve this, we turn to community preferences. In the next stage of our method, we bring in input from simulated personas to help resolve this ambiguity.

4.3.2 LLM-based Simulated Community Preferences

We simulate community input and generate responses from synthetic personas using large language models (LLMs). The personas we create represent a diverse set of individuals from different communities with varied backgrounds, social contexts, and disaster experience. Our goal is to approximate how people in communities might prioritize infrastructure repairs during a disaster when multiple repair actions are possible.

Persona Generation

We generate 200 personas by applying prompting using few-shot learning in LangChain. We define each persona using 44 attributes from the American Census Survey (ACS), such as age, occupation, household composition, and several others. We add some additional attributes, such as disaster experience and geographical class. The full list of attributes is in Table 4.1.

We use the GPT-3.5-turbo model with a temperature of 0.9. All personas were generated on April 2nd, 2025, between 2:00 PM and 4:00 PM Central Time using the OpenAI API.

We use two persona datasets: one where each persona is assigned to a community (*community-aware*) and another where community identity is removed (*community-unaware*). Both share the same attributes. We generate the community-aware set first and then create the community-unaware version by removing community labels. This allows us to test whether community awareness affects persona preferences. See generated persona data [here](#)¹⁰³.

Simulation Setup

We assign each persona a disaster scenario with infrastructure disruptions across three communities. For each scenario, we specify which infrastructure functionalities are working and which are damaged. The damaged functionalities are selected to create a mix of affected services and sectors, requiring personas to choose between, for example, water access in residential areas and power in commercial areas.

The information provided about each community includes the list of working and damaged functionalities, its geographical classification (urban, suburban, or rural), and its social vulnerability score. For example, Community 1 is classified as rural and has a social vulnerability score of 2.0 out of 10. In this community, water in the school is damaged, and both water and power in the residential area are down, while the hospital and commercial area remain functional. Community 2 is suburban with a vulnerability score of 3.0. Here, the residential area and water access in the commercial area are damaged, while the school and other infrastructures are operational. Community 3 is urban with a high vulnerability score

Table 4.1: Attributes Used to Generate Synthetic Personas

Attribute	Description
Persona ID	Unique identifier for each persona
Age	Age in years
Gender	Gender identity
Race/Ethnicity	Self-identified racial or ethnic group
Marital Status	Marital or relationship status
Citizenship Status	Citizenship classification
Birth Place	Country or region of birth
Year of Immigration	Year of immigration (if applicable)
Primary Language	Main language spoken
English Proficiency	Level of fluency in English
Education Level	Highest level of education attained
School Enrolment	Current enrolment status
Employment Status	Employment situation
Occupation	Job or role in the workforce
Industry of Employment	Sector or field of employment
Class of Worker	Employment classification
Work Hours Per Week	Average hours worked weekly
Income Bracket	Annual income range
Poverty Status	Position relative to poverty line
Government Assistance	Participation in support programs
Household Relationship	Role in household (e.g., head, dependent)
Household Size	Number of people in the household
Number of Children	Number of children in the household
Responsible for Grandchildren	Caretaking responsibility
Home Ownership	Own or rent housing
Housing Type	Type of dwelling
Utilities Access	Availability of basic utilities
Utilities Housing Costs	Cost of utilities and housing
Transportation Mode	Main mode of transport
Commute Duration	Travel time to work
Work From Home	Whether the persona works remotely
Health Insurance	Type of health coverage
Disability Status	Whether the persona has a disability
Disability Type	Type of disability (if applicable)
Military Service	Military background
Military Service Period	Time served in the military
Disability Rating	Government-assigned disability rating
Internet Access	Type of internet connection at home
Device Access	Devices available to the persona
Recent Birth	Whether persona recently had a child
Moved in Last Year	Whether persona changed residence
Previous Residence	Where persona lived before moving
Disaster Experience	Experience with past disasters
Community Description	Infrastructure status across 3 communities
Geographical Class	Urban, suburban, or rural location
Community*	Community assignment

of 7.0. In this case, water access is disrupted in school and residential areas, and power is also down in the residential area, while the rest of the infrastructure is functional.

We include this information in the prompt so that the persona can consider both system-wide infrastructure conditions and the community context when making decisions. This setup allows us to simulate trade-offs in prioritization.

Question and Response Generation

For questions, we create all possible pairwise comparison questions by taking all unique pairs of damaged functionalities across the three communities. To collect repair preferences, each persona is presented with a disaster scenario where multiple infrastructure functionalities are damaged. Then, the persona is asked to choose between two feasible repair options and to provide a justification for their choice. Each prompt is designed to capture both the decision and the reasoning behind it. In every prompt, we make sure the persona has access to information about the status of all three communities. We repeat this process across many personas and question combinations to cover all pairwise preferences we created. We use the GPT-3.5-turbo model with a temperature of 0.7. We did these generations on April 3rd, 2025, between 2:00 PM and 4:00 PM Central Time using the OpenAI API. Previously, we mentioned that we used a temperature of 0.9 in the persona generation. This was to promote diversity in characteristics such as age, occupation, and disaster experience. Now, for the response generation, we set the temperature to 0.7 to encourage more stable, realistic, and context-grounded reasoning. With this configuration, we try to balance creative variability during persona creation with more deterministic and coherent outputs during response collection. The specific prompt templates used to generate the personas and collect their responses are provided in the Appendix [A](#).

4.3.3 Feature Extraction & Data Processing

We group the data by question and calculate the percentage of personas who chose each option. These percentages serve as soft labels, a two-dimensional probability vector $[p_1, p_2]$ where p_1 is the number of personas who chose option 1 and p_2 is $1-p_1$. For each pairwise question, the dataset is a pair of feature vectors and the corresponding soft label.

Next, we extract features for each repair option, which represent damaged functionalities in specific communities. Each feature vector includes information about the infrastructure type and its community. Table [4.2](#) lists the features we use for each option. The SVS⁵¹ is a metric we use to identify the different vulnerability levels across a community during

disasters, originally defined on a 1–5 zone scale. In our work, we normalize this to a 0–1 scale, with 1 representing the highest vulnerability. In total, we aggregate 36 distinct pairwise questions drawn from 7,200 persona responses.

Table 4.2: Infrastructure Repair Option Features

Feature	Description
Infrastructure Type	One-hot encoded: Water, Power
Facility Type	One-hot encoded: Hospital, Residential, Commercial, School
Community ID	One-hot encoded: Community 1, 2, or 3
Geographical Class	One-hot Encoded: Urban, Suburban, or Rural
Social Vulnerability Score (SVS)	Continuous value (0 to 10)

4.3.4 Pairwise Comparator and Chainization Model

We train a neural network model to predict preferences between pairs of infrastructure repair options. The architecture consists of a feature extractor that processes each option independently, followed by a comparison layer that combines the extracted features. The output is a softmax probability distribution over the two choices.

We train the model with a hybrid loss function, where cross-entropy loss is for when the majority of personas strongly prefer one option ($P_1 \neq P_2$), while we use KL-divergence loss when both options are equally preferred ($P_1 = P_2$).

For optimisation, we use Adam and a batch size of 8. We select the learning rate by using a six-fold validation of the full set of 36 pairwise questions. For each learning rate tested (1×10^{-3} , 5×10^{-4} , 1×10^{-4}), we obtained the best validation loss reached for every fold before early stopping. Then we averaged these six values. The mean best-fold losses were 0.4267, 0.4842, and 0.6584, respectively. We adopt the learning rate of 1×10^{-3} for all subsequent experiments because that gave us the lowest average. See the training and validation curves for this setting in Figure 4.3.

Once pairwise comparisons are learned, we construct a global ranking of infrastructure repairs using the Chainization Algorithm. To construct the pairwise preference graph, we use a directed graph, $G = (V, E)$ where the nodes (V) represent infrastructure repair options and edges (E) show preference such that when option 1 is preferred over option 2 based on




Figure 4.3: Six-fold cross-validation curves for the pairwise comparator at $LR = 1 \times 10^{-3}$. Solid lines show training loss; dashed lines show validation loss for each fold.

model prediction, an edge (*option2* $\rightarrow$ *option1*) is added.

For the ranking, we apply PageRank instead of Kahn’s topological sorting to order infrastructure repairs. In all PageRank computations, we assign edge weights based on the model’s predicted preference confidence. We use PageRank because it handles cycles cleanly by assigning global importance scores. Then we rank infrastructure options by score to produce the final prioritization.

4.3.5 Random Subset Sampling and Fine-Tuning with Pseudo-Pairs

We adopt a training and refinement approach to investigate the feasibility of reducing survey burden while maintaining prioritization accuracy. First, we randomly sample a fraction (20%, 30%, 40%, 50%, 60%, 75%) of the available pairwise comparisons from the full dataset. This simulates realistic scenarios where it is impractical to collect all possible preference

judgments.

Then we train a pairwise comparator model using only the sampled subset. We still use the same hybrid loss formulation, that is, cross-entropy loss for strong preferences and KL-divergence loss for equal cases.

After this initial training, the comparator model predicts pairwise preferences for all infrastructure options. We use these predictions to build a directed preference graph and apply chainization to infer a preliminary global ranking.

Based on the recovered initial global rank, we generate synthetic pseudo-pairs. We do this by assuming that for any two facilities that have a clear ranking relationship (e.g., $A > B$), a pseudo-comparison is added with a hard label indicating definite preference, i.e., $[1.0, 0.0]$. We use this to enrich the available training data. Then, we fine-tune the comparator model on a combined dataset that comprises both the original sampled comparisons and the generated pseudo-pairs. Following fine-tuning, final pairwise predictions are used to reconstruct the preference graph and produce a global repair prioritization.

We present the results of Kendall’s τ correlation between the full ranking and the sampled rankings, as well as the Top-K-Lists overlap (Top-3 and Top-5)^{104;105} in Section 5.4.

4.3.6 Validation and Sensitivity Analysis

To assess the robustness of our prioritization approach, we conduct two complementary evaluations. First, we examine the influence of community identity on repair prioritization by comparing the global rankings derived from two distinct datasets: one generated using community-aware personas and the other using community-unaware personas. Second, we conduct a sensitivity analysis to test the stability of our results with respect to variations in prompt formulation. Prior studies have shown that large language models (LLMs) can be highly sensitive to prompt rewording, tone, and instruction style^{106–109}. To evaluate this, we create six variations of the original prompt:

1. Reworded Question: We changed the phrasing of the question from "Which should be repaired first: 'option1' or 'option2'?" to "Between the two options, 'option1' and

'option2', which do you deem most critical for immediate repair?"

2. Reworded Reasoning: We changed the reasoning from "Your one-sentence explanation" to "Explain your selection with a concise rationale, highlighting relevant factors influencing your decision."
3. Tone Shift: We introduced empathetic framing like "Keep in mind the emotional and practical consequences of the outage for residents living in these communities as you evaluate what should be repaired first."
4. Reformatted Instruction Style: We change the structural layout of the prompt by adding steps for each information in the prompt.
5. Choice Before Reasoning: We restructured the prompt so that the LLM first states its repair choice and only then provides its justification. All other variants (including the original) require the model to explain its reasoning before revealing its selection.
6. SVS Removed: We dropped the explanation of SVS to test its influence. In exploratory runs, we observed that when SVS was not defined in the prompt, the model sometimes misread the community with a low vulnerability score of 0.2 as having high vulnerability, so removing that detail helped to understand its impact.

We use each variant to regenerate the persona responses for the same infrastructure repair comparisons. These new responses are then processed through the full modeling pipeline. To quantify robustness, we compute the Kendall τ correlation between each prompt-variant ranking and the baseline ranking from the original prompt. This allows us to determine whether minor variations in prompt design affect prioritization outcomes. Also, we look at the Top-K-Lists overlap (Top-3 and Top-5). All code and generated data used in this chapter are publicly available at the GitHub repository

4.4 Results

We focus our results on three different aspects: the impact of community awareness in persona generation, the effect of limiting the number of preference queries using chainization and pseudo-pairs, and the sensitivity of prioritization outcomes to variations in prompt design.

4.4.1 Impact of Community Awareness in Simulated Preferences

To investigate whether awareness of community assignment affects infrastructure repair prioritization, we compare the results from community-unaware personas and community-aware personas. In the community-aware setting, personas were told which community they belonged to, which allowed for potential bias toward their own community. In contrast, in the community-unaware setting, personas made decisions without this information.

Table 4.3 shows the final repair ranking, which is identical for both the community-aware and community-unaware personas. We initially expected that making personas aware of their own community would introduce some level of in-group bias and shift repair priorities towards their own group. However, the consistency in rankings suggests that personas weighed other factors more heavily in their decision-making. In this scenario, community identity did not significantly influence the overall prioritization.

Also, the final ranking shows that repairs related to schools and residential areas appear most frequently near the top. Water in a school in Community 3 is ranked first, followed by Water School in Community 1. School-related repairs appear in two of the top three positions.

Residential infrastructure fills most of the middle rankings, with power generally ranked above water within the same community. We notice that Water in Residential in Community 3 ranks above Power Residential in Community 1. This might be linked to the high vulnerability of Community 3. The only commercial repair, Water Commercial in Community 2, is ranked last.

Figure 4.4 and Figure 4.5 show how personas from each community made choices across all

Table 4.3: Final Global Repair Prioritization (Same for Community-Aware and Community-Unaware Datasets)

Rank	Repair Option
1	Repair Water School in Community 3
2	Repair Water School in Community 1
3	Repair Power Residential in Community 2
4	Repair Power Residential in Community 3
5	Repair Water Residential in Community 3
6	Repair Power Residential in Community 1
7	Repair Water Residential in Community 2
8	Repair Water Residential in Community 1
9	Repair Water Commercial in Community 2

three communities, in both the community-aware and community-unaware persona datasets. While we expected stronger in-group preference when community identity was made explicit, the results show only a minor difference. For example, Community 3 personas chose their own community 41.6% of the time when aware, and 42.8% when unaware. Similarly, Communities 1 and 2 showed nearly identical in-group selection across both datasets.

We observe that Community 3 is consistently preferred even by out-group personas. This may suggest Community 3 is perceived as most in need, possibly due to its higher SVS of 0.7. These results suggest that community awareness had minimal impact on in-group preference, and that other factors may have influenced decision-making more strongly. Given the high SVS of Community 3, we suspect that personas may have considered broader vulnerability rather than personal affiliation when making prioritization decisions.

4.4.2 Accuracy with Limited Pairwise Comparisons

We evaluated the possibility of achieving a good global ranking by using the subset of possible pairwise comparisons. We used random sampling to select some samples of the pairwise data, applied chainization and fine-tuning on pseudo pairs to obtain a global rank. Then, we applied Kendall’s τ correlation between the final inferred ranking and the ranking obtained from the full dataset.

Figure 4.6 shows the trend of mean Kendall’s τ scores across different percentages of comparisons used. Results are averaged over 100 random trials for each sampling fraction. We observe that as the percentage of available comparisons increases, the agreement between




Figure 4.4: Community Preferences by Community-Aware Personas

the predicted and ground-truth rankings improves. We notice that using only 20–30% of the comparisons achieves a low Kendall’s τ correlation (<0.6), but once the sampled comparisons reach 50%, the ranking performance improves substantially, reaching a Kendall’s τ above 0.7. These results show that we can recover good global rankings from limited preference data, which offers a practical way to reduce the burden on survey participants while maintaining decision quality.

We also present the results of the Top-3 and Top-5 overlaps in Figure 4.7. The overlap measures how consistently the highest-ranked repair options match the ground truth rankings as we increase the percentage of pairwise comparisons used for training. We observe that the overlap steadily improves as the sample percentage increases. In the range of 40% to 60%, the overlaps for both Top-3 and Top-5 remain relatively stable, showing incremental alignment with the ground truth. In fact, the agreement is as high as 80%. This suggests that a relatively small fraction of the total data is sufficient to capture the most critical repair priorities, which is valuable when collecting survey data is costly or time-constrained.




Figure 4.5: Community Preferences by Community-Unaware Personas

4.4.3 Sensitivity to Prompt Variations

We evaluated six prompt variations against our reference ranking of nine repair options. See Table 4.4 for each item’s numeric position. Figure 4.8 summarizes the overall agreement via Kendall’s Tau and Top- k overlaps. We observe that changing instruction format gives $\tau = 0.50$ ($p=0.075$) and rewording the question in the prompts yields $\tau = 0.44$ ($p=0.12$), rewording the reasoning gives $\tau = 0.39$ ($p=0.18$), shifting the tone of prompt yields $\tau = 0.56$ ($p=0.045$), choice first gives $\tau = 0.44$ ($p=0.12$) and removing SVS explanation gives $\tau = 0.72$ ($p=0.0059$). We observe that only the variant where we remove SVS shows high agreement with the original. Using the Top- K -Lists overlap, we observe that SVS-removed gives 100% Top-5 overlap, rewording the question yields 100% Top-3 overlap, and other variants show a high level of Top-5 overlap. Our results show that the lower-ranked repairs are more sensitive to prompt changes, and the very top priorities remain stable.

More so, our results indicate that SVS strongly influences persona prioritization and outweighs community affiliation in repair choices. This raises questions about whether the LLM separates vulnerability from affiliation when reasoning, or treats them as interchangeable signals of need. Additionally, our sensitivity analysis shows that the repair rankings




Figure 4.6: Kendall’s τ correlation versus percentage of comparisons used.

remain highly dependent on the prompts provided, with SVS as a dominant factor.

4.5 Discussion

This study presents an approach to prioritize infrastructure repairs by combining hetero-functional graph modeling with community input generated using large language models (LLMs). Our findings show that combining synthetic persona preferences with repair dependencies can help resolve ambiguity in infrastructure recovery in a scalable and people-centered way.

One of the most important findings in this study is that we can recover accurate global repair rankings using only part of the data. By selecting just 60 percent of the pairwise comparisons and applying chainization with added pseudo-pairs, we achieve high agreement with the full-data rankings. The Kendall Tau correlation goes above 0.7. This is a useful result for disaster planning, where one may lack time or resources to gather full data. Our method reduces the burden on people while still producing reliable prioritization.

We also tested how sensitive our method is to changes in the way we ask questions.




Figure 4.7: Top-K overlap on Rankings obtained via Random SubSampling and Pseudo-Pairs

When we reworded the prompt, changed its tone, or adjusted the instruction format, the final rankings changed slightly. However, we saw more consistency among higher priority items, while lower ranked repairs were more sensitive to prompt changes. This demonstrates a strong consensus among personas on identifying the most critical tasks.

While our results are promising, there are some limitations. First, the preferences generated by the model may not fully match how real people behave, especially in stressful or emotional situations. Second, the infrastructure model we use is simplified and includes a fixed set of repair tasks. Future research could apply this method to more complex systems with more diverse failure scenarios.

Third, in our approach, we treat all functionalities, regardless of infrastructure type, as part of a unified repair priority list. We acknowledge that this might not necessarily reflect how infrastructure is usually restored in the field. With the exception of certain small communities, water, power, and transportation are usually handled by specialized repair crews and agencies. Nevertheless, our method can be easily extended to deal with these cases. Moreover, our method can also support broader planning questions, such as where to




Figure 4.8: Prompts Variations Overall Agreement with Original using Kendall’s τ and Top-K overlaps.

invest in resilience or how to include community input in upgrade decisions. In such high-level decisions, cross-sector rankings can be helpful, especially when guided by community preferences.

Also, we do not consider social cognitive factors such as measurements of community connectedness and personality traits.

Lastly, our experiments with sampling and prompt changes are only a first step. Future work can explore more adaptive strategies that learn which comparisons are most useful as the model is trained.

4.6 Conclusion

Our work presents a new approach to infrastructure repair prioritization by combining hetero-functional graphs, community preferences generated using large language models (LLMs), and a machine learning model that learns from pairwise comparisons. Our method focuses

Table 4.4: Global Repair Prioritization Across Prompt Variations

Repair Option	Instruction Format	Reworded Question	Reworded Reasoning	Shifting Tone	Choice First	Without SVS
Repair Water School in Community 3	1	1	1	1	3	2
Repair Water School in Community 1	4	2	6	2	5	5
Repair Power Residential in Community 2	6	3	5	6	4	1
Repair Power Residential in Community 3	3	8	2	5	1	3
Repair Water Residential in Community 3	2	7	3	3	2	4
Repair Power Residential in Community 1	8	6	9	8	8	6
Repair Water Residential in Community 1	7	5	7	7	7	8
Repair Water Residential in Community 2	5	4	4	4	6	7
Repair Water Commercial in Community 2	9	9	8	9	9	9

on a common challenge in disaster recovery, which is what to repair first when technical dependencies allow for multiple valid options. By simulating community input through synthetic personas and applying a comparator model with chainization, we are able to produce repair priorities that reflect both system constraints and community preferences.

Our results show that full-data rankings can be approximated using a reduced set of comparisons, lowering the burden of data collection. And we expect that with actual survey data, the selection will likely be more accurate. We also tested how prompt changes affect the outcomes and found that while lower-ranked items shift, top priorities remain stable across variations.

This work opens new directions for integrating synthetic decision-making tools into critical infrastructure resilience frameworks. Based on this, we have identified several future directions for this research.

First, future research should validate these methods with real-world stakeholder input and extend them to more complex systems. We also plan to explore better ways to select informative comparisons and apply the method to a broader range of infrastructure types. Additionally, separating infrastructure types and applying domain-specific repair constraints, as well as exploring more complex ranking situations, will be important extensions. However, our current unified ranking approach remains valuable for investment planning contexts,

where decision-makers must choose across multiple systems.

The next step is to collect survey data and formally test our proof of concept in a real-world setting using actual community data. We also plan to use prior survey data or data from another community and leverage the LLM to generate data for a different community, which may be advantageous given the limited survey response rates in recent years.

An important direction is to address potential biases introduced by the generation of synthetic public viewpoints and responses. We plan to extend this work by validating our approach with real-world survey data that incorporates community demographics, personality characteristics, social cognitive factors, and perception data to enhance the realism and representativeness of the simulated preferences.

We also plan to explore infrastructure consumption nodes as interdependent entities based on resident or household utility demand, allowing for a more integrated modeling of how service disruptions impact communities. This will be a focus in the next phase of our survey-based studies.

In addition, we aim to incorporate community attachment explicitly, as it may significantly influence persona preferences. The community-aware dataset will be expanded and validated using additional metrics of community connectedness in subsequent survey efforts.

Finally, another area for future research is to investigate how repair prioritization rankings might be established pre-disaster, recognizing that real-time or post-disaster decision-making may be challenging. This would allow our framework to support proactive planning scenarios in addition to recovery contexts.

Chapter 5

Towards Quantifying the Communication Aspect of Resilience in Disaster-Prone Communities

5.1 Introduction

Disasters can occur in any community, whether rural, suburban, or urban. However, disaster preparedness, response, and recovery differ across communities. Rural communities typically have low population density and access to essential critical services may be challenging. On the other hand, urban areas may have better access to and more robust critical infrastructures. However, since urban communities are much more heavily populated, any disaster-related disruption could have a significant impact. This contrast between urban and rural communities presents challenges in determining the optimal allocation of limited resources during disasters. Each community has to deal with its own local factors that can hamper its ability to respond to disaster. Amongst these factors, effective communication of life-saving information during a disaster is critical, and the unique characteristics of communities can affect the speed at which such information spreads, potentially impacting a community's resilience. Our research investigates how the process of information diffusion,

which is indicative of a community’s communication resilience, depends on the community’s characteristics.

In our research, ”disaster” conditions refer to significant disruptions or emergencies resulting from natural disasters such as floods, hurricanes, or severe storms. These conditions can lead to widespread damage to infrastructure, displacement of populations, and interruption of essential services. However, in this study, we focus on communities that are prone to flooding and how the diffusion time during normal and disaster times differs among these communities.

In the context of flood disasters, the provided disaster times signify the duration it takes for critical information to reach various segments of the population. During these times, people are likely to receive information through multiple channels, including news broadcasts, social media updates, messages from friends and family, and official government alerts. However, the timing of when individuals receive these messages may vary depending on factors such as the availability of communication channels, individuals’ media preferences, and the effectiveness of government alert systems. Understanding the implications of disaster times on communication channels and information dissemination is indeed crucial, as it sheds light on how communities cope with and respond to emergencies. For example, during a flood, individuals might rely more on immediate sources like social media and friends for real-time updates on the situation, especially if traditional communication infrastructure like television and radio are disrupted. In contrast, on a normal day, people might primarily receive news through traditional media channels. Additionally, the types of alerts issued during a disaster, such as evacuation orders or safety advisories, would also influence people’s behaviors and the urgency with which they seek information.

Traditionally, we may define a *community* as a group of people living in close geographical proximity. Multiple studies highlight the importance of a community-based approach in enhancing disaster resilience^{110–114}. However, with the recent increase in the use of digital devices and social media for communication, it is necessary to redefine the notion of *community* as people build interpersonal relationships beyond their physical neighborhoods using digital devices. People use social network platforms such as Facebook and Twitter to share

information concerning disasters^{115–117}. Hence, it makes sense to define the community-based social network approach to disaster resilience by augmenting face-to-face networks with other forms of social networking. Our work expands on the definition of community by including face-to-face and online social networks.

The usage of different communication mechanisms differs across communities – for instance, in communities in the United States, studies show that adults in urban and suburban communities are more likely to have access to broadband internet than their rural counterparts^{118;119}. Individuals in these communities are exposed to information from diverse sources and rely on different communication methods in different ways, as various factors influence their preference for communication techniques. Our approach considers that community structure varies in different contexts, including the relationship and dependence between people and specific infrastructures, which affects the time of information diffusion. Network structure, trust, and interaction rates between people and communication infrastructures play a crucial role.

The study of resilience in complex networks is crucial for assessing the impact of failures and implementing mitigation techniques¹²⁰. Understanding resilience is especially important for communities during natural disasters, as communication infrastructure failures can cause information disconnection. Different communities face different challenges, such as face-to-face communication difficulty when roads or bridges are inaccessible or TV and digital device failure during power outages. Over-reliance on a few methods of information diffusion creates problems when infrastructures fail. Our study looks at the impact of communication infrastructure failure on each community and the ways to improve a community’s resilience to such disruptions.

Previous studies have explored the evolution of disaster information on online social networks like Twitter and Facebook^{121;122}, but little work exists for face-to-face networks or other offline communication methods. Our approach investigates the complex spread of information in disaster-affected communities by analyzing the interplay between five information sources: online social networks, local news, government sources, mainstream media, and print media. In this study, we have specific objectives as follows: First, we investi-

gate whether there is significant variability in the information flow graph between rural, suburban, and urban communities as well as between disaster and non-disaster scenarios. Second, we investigate whether there is more information diffusion in online social networks between different community settings where spatial connectivity and infrastructure provision are different. Finally, we explore whether there is significant variability in network properties within the disaster and regular (non-disaster) information networks from rural to urban communities.

Our study finds that urban communities rely more on online social networks than rural communities for information dissemination. Furthermore, the information network properties and flow graphs of urban, suburban, and rural communities differed in normal and disaster situations, with variations in network centrality but no significant variability in the average shortest path lengths. We developed a stochastic model using survey responses to simulate information diffusion in a multilayered network of three Midwest US counties. We inferred interaction and trust patterns among people in rural, suburban, and urban communities and analyzed the effects of different sources of information on information spread. We found that community structure, the degree of sparsity, and connectedness influence the average time of diffusion of information. Moreover, the failure of one communication infrastructure affects each community differently. To overcome the need for accurate network information, we learned the distributions of the node-to-node relationships from survey responses and created synthetic graphs using these distributions.

Our contributions include the development of a stochastic model that predicts the average time of information diffusion in urban, suburban, and rural areas and districts in normal and disaster times using multiple social networks. This model integrates various factors such as community degree distribution, trust relationships, and interaction rates between individuals and communication infrastructure. Additionally, we demonstrate the independence of diffusion time from community size, even with a community graph ranging from 1000 to 25000 nodes, assuming the same type of information is being propagated in each community in normal times or disaster times. Furthermore, we analyze information flow graphs to understand how different network structures impact diffusion time and we evaluate com-

munity resilience in scenarios of communication infrastructure failure. Finally, we develop a framework for extracting policy recommendations, enabling communities to make informed, budget-conscious investments to enhance information diffusion during disasters.

5.1.1 Acknowledgments

This research was supported by funding from the USA National Science Foundation (NSF) under grant numbers 1952206 and 2148878. A portion of this work is published in [19](#).

5.2 Literature Review

The independent cascade model [123;124](#) and linear threshold model [124](#) are the two main models used in information diffusion theory to describe the diffusion of information in a network. In the linear threshold model, every node i has a threshold τ_i . For every edge (i, j) between i and its neighbor j there exists an influence (of i from j) weight of $w_{ij} \geq 0$ such that $\sum_j w_{ij} \leq 1$. A node i becomes active when the sum of the influence weights from its active neighbors surpasses τ_i . By contrast, in the independent cascade model, each edge (i, j) has a certain probability in the directed graph, such that an active node j can influence node i with probability P_{ij} . Through this activation process, influence spreads in the network. In our work, we use a variant of these models since the propagation of information depends on two factors, trust and the rate of interactions between nodes. Hence, our work incorporates both neighbors' influences and a trust threshold value to propagate information. Our network has rates of interaction and trust values between all the participants and their neighbors. We assume that the neighbors of an uninformed node will share information with this node during a time step. With a certain probability, informed by the nodes' rate of interaction, an uninformed node meets an informed node during a time equal to one simulation step. During the meeting, the informed node communicates the relevant information to the uninformed node, who takes it into consideration through a mechanism that depends on the trust relationship between the two nodes, and previous exposures to the same information.

When the uninformed node's trust in the new information exceeds a certain threshold, the uninformed node becomes informed.

Past research has extensively investigated the evolution of disaster information on online social networks like Twitter and Facebook. Yoo et al.¹²¹ characterized diffusion rates during a disaster by examining key elements of information propagation rates on Twitter data. Fan et al.¹²⁵ studied the influence of different users on the diffusion of disaster information on Twitter, finding that early interventions from hubs increased the speed of information propagation. Kim et al.¹²⁶ analyzed interactions between news agencies, weather agencies, and the public on Twitter during the 2017 Storm Cindy, identifying news and weather agencies as dominant sources of information. Kim and Hastak¹¹⁶ studied the role of social media data on Facebook in information propagation during disasters, identifying specific individuals as hubs. Yang et al.¹²⁷ investigated influential Twitter users during Hurricane Harvey, highlighting the role of objective information sharing in disaster-related follower growth. Fan et al.¹²² studied dynamic changes in network structure during information diffusion of disaster information on social networks. Dong et al.¹²⁸ created an information flow model from Weibo data to characterize information diffusion patterns during disasters. Sharma et al.¹²⁹ investigated challenges of integrating social network data and emergency management to maximize information diffusion of disaster information. Zhang et al.¹³⁰ explored the potential of social media in enhancing public information and warning systems for disaster management, emphasizing its role in acquiring situational awareness and enabling communication with disaster management agencies. Zhai et al.¹³¹ characterized disaster tweets to understand different perspectives of shared data during information diffusion. Xu and Qiang¹³² used retweets to model information diffusion and analyze the geographic distribution of topics during information flow. Nagar et al.¹³³ characterized social media response by studying how disaster news spread on Twitter and the countries where tweets emerged. Zhu et al.¹³⁴ built a probabilistic model for retweeting during disasters, identifying factors affecting retweeting behaviors. Altay and Pal¹³⁵ used an agent-based model to demonstrate the importance of information hubs in the speed of information during disaster response. Liu et al.¹³⁶ developed a theoretical model for modeling information diffusion in online social

networks during emergencies.

From the review of past studies, we deduce that most work on understanding information diffusion in social networks is focused on online social networks with little work on face-to-face networks or other offline communication methods. This research aims to address this gap by exploring how humans communicate using different communication mechanisms (local government, online social networks, local news, cable news, and print media) across different development contexts (urban, suburban, and rural) and identifying opportunities for enhancing communication strategies to improve disaster response and resilience. Our model incorporates a person’s relationship in their social network by including trust in neighbors and different communication mechanisms. We take this approach since research has shown that trust is vital in information diffusion and people’s decisions during disasters. Adali et al.¹³⁷ emphasized the significance of trust in interpersonal relationships, particularly in the context of social networks, and introduced algorithmic methods to measure and quantify trust based on communication behavior, highlighting its relevance for understanding information flow dynamics in networks like Twitter. Wu et al.¹³⁸ underscored the importance of trust in information diffusion processes within network theory, highlighting the influence of rational decisions based on trust levels in acquaintances, and explored how trust dynamics impact information propagation in a two-layer multiplex network, revealing that memory span and trustable acquaintances significantly affect information spreading dynamics. Moreover, Fridman et al.¹³⁹ found that trust in government sources leads to accurate knowledge and application of preventive measures during COVID-19. In contrast, a negative relationship exists between trust in cable news and social networks. Additionally, an individual’s social network can affect their decisions during a disaster^{140;141}. Specifically, Widener et al.¹⁴⁰ used an agent-based model to show that social networks can influence evacuation decisions, while Haer et al.¹⁴¹ found that communication strategies and social networks affect whether or not people take protective actions during disasters. Our stochastic model incorporates people’s trust relationships within the community and with different communication hubs.

5.3 Methodology

Here we discuss our survey methodology, feature extraction, graph generation, stochastic model, and model assumptions.

5.3.1 Data

A total of 2,736 households were selected using a stratified cluster sampling method from urban, suburban, and rural residents within Riley County, Kansas, Buchanan County, Missouri, and Platte County, Nebraska which were recently exposed to flooding. These areas are in the high-risk Midwest watershed areas.

In our selection methodology, we considered the extent and direction of the urban-to-rural gradient within each setting, utilizing zip codes as the unit of analysis. We specifically incorporated demographic, economic, and physical characteristics to define the study areas, including population density, household income, street networks, urbanization patterns, and flood impact areas along this gradient. The resulting classification of each county into urban, suburban, and rural zones is depicted in Figure 5.1. This approach allowed for a comprehensive examination of the diverse urban and rural landscapes present in our study areas.




Figure 5.1: Classification of Areas in County. These maps were created using Adobe Photoshop (version 22.1, Adobe Systems, <https://adobe.com/products/photoshop.html>)

The survey was conducted from March to April 2021 through the USPS every door direct mail (EDDM) service and 155 complete responses were obtained with a response rate of 5.6%.

Before the survey was conducted, this study was approved by the Kansas State University’s Institutional Review Board. Table 5.1 shows the distribution of responses.

Table 5.1: Distribution of Responses of Communities. Note: * removed due to insufficient data

Area Type	Response	Districts	Response
Rural	49	Buchanan Rural	21
		Platte Rural*	5*
		Riley Rural	23
Suburban	62	Buchanan Suburban	16
		Platte Suburban	13
		Riley Suburban	33
Urban	44	Buchanan Urban*	4*
		Platte Urban	21
		Riley Urban	19

We focus on survey responses relevant to our study of developing a network of nodes (people) with edges that describe interaction patterns and trust relationships. Although our data is limited because we do not have complete network information, we assume that survey results from random samples of nodes from different communities provide a close estimation of the interactions between community members. Since we are interested in people’s interaction and behavior in normal daily life and disaster times, our survey questions cover both situations. These responses allow us to build different information networks for disaster and normal daily life. The data includes: (i) the number of people with whom respondents interact regularly within their neighborhood, which represents the degree distribution of communities, ranging from 0 to 160; (ii) the five people respondents trust the most outside their households, from whom they prefer to obtain information during normal daily life and disaster situations - ‘the frequency of interaction with them’ (a five-point Likert scale: Never =1 to Always =5)’ as well as ‘the level of trust in them’ (0-10 scale); and, (iii) respondent’s preferred information source during normal and disaster times among the local government, online social networks, local news, cable news, and print media - ‘the frequency of weekly access to these sources (a five-point Likert scale: Never =1 to Always =5)’ and ‘the level of trust in these sources’ (0-10 scale).

After the survey data cleaning, we have ten communities to explore, as shown in Table 5.1. We dropped two districts (Buchanan Urban and Platte Rural) due to insufficient responses

but included their data when considering total urban and rural areas. To create features for each community, we convert survey responses into two sets of thirteen features - one for normal times and one for disaster times. The features are in Table 5.2.

We obtained each community’s degree sequence by using the node’s degree, which is the number of connections a node has with others¹⁴². We cleaned up the data by dropping one missing data point and outliers, defined as data points above 100. For the frequency of interaction with each information source and interactions with neighbors, we dropped the missing data points and encoded responses with a scale of 1-5 (Never: 1, Rarely: 2, Sometimes: 3, Frequently: 4, Very Frequently: 5). Trust values for neighbors and information sources range from 0 to 10 and we dropped missing data points for these as well.

5.3.2 Graph Generation

To build a stochastic model of information diffusion in each community, we need a graph representation of the community. We use the thirteen features obtained in the previous section to build graphs of the ten communities. Our assumptions to generate these graphs and assign attributes to the communities are the obtained degree distribution from the degree sequence represents the degree distribution in each community, the frequency of interaction and trust relationships indicate the strength of ties between nodes. and each person in a community can obtain information directly from five sources: cable news, online social networks, local government agencies, local news, and print media.

In reality, the populations of the communities of interest are much larger than our available data. To overcome the small sample size, we assume that the number of nodes in our graph, denoted as n , represents the size of each community. Our graph $G = (V, E)$ has n vertices and m edges. We generate graphs that only contain community members (without hubs) by using the kernel density estimation (KDE)¹⁴³ to learn the degree distribution of each community. We sample and generate a degree sequence of size n from this distribution. We use the Markov chain Monte Carlo (MCMC) algorithm^{144;145} to sample 100 graphs from the configuration model¹⁴⁶ for each community during the simulation process. We add five

information source nodes (hubs) to the generated graphs – to represent the local government, online social networks, local news, cable news, and print news – increasing the number of nodes to $n + 5$. Each node can be in one of two states: with information (state 1) or without information (state 0). Before the simulation, we assume all information sources are in state 1, and all other nodes are in state 0. We build the graph using the NetworkX¹⁴⁷ Python library.

Figure 5.2: The Features of a Community

Community Features
1. Features relating to node degree:
a. The number of people each person interacts with daily
2. Features relating to degree of interaction
b. The frequency of interaction with 5 most trusted neighbors
c. The frequency of interaction with local government
d. The frequency of interaction with online social networks
e. The frequency of interaction with cable news
f. The frequency of interaction with local news
g. The frequency of interaction with print news
3. Features relating to degree of trust
h. Trust value for neighbors for 5 most trusted neighbors
i. Trust value for local government
j. Trust value for online social networks
k. Trust value for cable news
l. Trust value for local news
m. Trust value for print news

Since information propagation in a network depends on tie strength¹⁴⁸, we focused on understanding node interaction rates and trusts to achieve our goal of propagating information through the network. We consider two types of edge relationship attributes: interactions and trust. For interactions, we have directed edges from the sources of information to all other nodes and undirected edges between nodes. The frequency of interaction between two nodes i and j is denoted by λ_{ij} . For trust relationships, we have directed edges from nodes to sources of information and undirected edges between nodes. The degree of trust between two nodes i and j is denoted by θ_{ij} .

We use kernel density estimation to learn the trust distribution from each community's trust sequence to each source of information. Then we randomly sample n values from this distribution to assign trust values to edges connecting the nodes to the hubs. This results in five different trust values for each node corresponding to their ties with the five hubs. Next,

we use kernel density estimation to learn the distribution of trust levels between people in a community and assign trust values to five randomly selected edges of the graph. We assign these sampled trust values to the edges representing the ties between a node and its five most trusted neighbors, and we set a constant value of 5 for the remaining edges between a node and its other neighbors.

As mentioned above, we converted the weekly interaction rate to numbers ranging from 1 to 5. We use kernel density estimation to learn the distribution for the frequency of interactions with sources of information and people in each community and sample from these distributions. For the sources of information, we get n samples from the learned distribution. We get $n \times 5$ random samples from the learned distribution for interactions with people. We convert these frequencies to a weekly interaction rate by assuming a typical range of days for talking with friends between $[0,12]$ days. We assign a rate to each edge by uniformly selecting a value from a corresponding range: 0 days for frequency 0, 1-2 days for frequency 2, 3-4 days for frequency 3, 5-7 days for frequency 4, and 8-12 days for frequency 5. We convert these rates to hourly rates by dividing the number of days by 24×7 . We assign the sampled rate values to only five random edges that depict an individual's five most trusted friends and a constant rate, $\frac{5}{24 \times 7}$, to the remaining edges between node i and its neighbors.

5.3.3 Stochastic Model

We are interested in estimating the average diffusion time, defined as the time that elapses between the start of an information event, and the time when 90% of individuals in the community have accepted the information (or have become *informed*). We use a stochastic model to simulate information passing within the built graphs. We start with an initial graph with $n + 5$ nodes. Initially, we assume there is a piece of information to be shared, and only the five hubs have this information. Let X be the number of interactions a node has with a specific neighbor – another person or a source of information – in a week, where $X \sim \text{Poisson}(\lambda)$, and λ is an edge value between the two nodes. Using this rate, we calculate

the probability of 0 meetings between a node and its neighbors. Zero meetings between node i and a neighbor j mean that node i and j have not met yet, so node i does not learn any information from node j . We begin by iterating through all the nodes in the graph. For each time step, we check the state of the node. If a node has the information, we do not take any action. However, if a node does not have the information, we check its informed neighbors. We iterate through the informed neighbors of each node i . To determine the probability that node i meets node j , P_0 , let the rate of interactions between i and j be λ_{ij} . We calculate P_0 as 1 minus the probability of zero meetings between node i and node j , that is, $P_0 = 1 - \text{Poisson}(X = 0, \lambda_{ij}) = 1 - e^{-\lambda_{ij}}$.

We draw once from a Bernoulli distribution with parameter P_0 . If the outcome is 1, we update the uninformed node's total trust in the new information. To update this total trust, we apply a discount factor to the trust value between node i and j . The discount factor accounts for the number of times node j passed the same information to node i before. Previous interactions with node j , as well as with other nodes, are discounted by a forgetting factor that accounts for the length of time that passed since those interactions. We calculate the trust as follows:

$$T_c = \sum_{j=1}^N \sum_{k=0}^{M_{ij}} \theta_{ij} d^k f^{t_c - t_k} \quad (5.1)$$

T_c is the total trust in the new information, N is the number of neighbors of i with a state of 1, M_{ij} is the number of meetings between i and j , θ_{ij} is the trust between node i and j , d is the discount factor and f is the forgetting factor, t_c is current time step and t_k is the time of meeting number k between i and j . Node i becomes *informed* when T_c exceeds a threshold Θ . In this experiment, we assign a discount factor $d = 0.5$, a forgetting factor $f = 1.3$, and a trust threshold value $\Theta = 30$. Next, we conduct a sensitivity analysis for these parameters, and the results are in the [Results](#) section. The simulation ends when 90% of the population receives the information. We repeat the simulation 1000 times for each of the 100 graphs and take an average of the diffusion time. See Algorithm 1 in the [Appendix B](#) for the pseudocode.

5.3.4 Network Metrics

We chose several key network metrics for our analysis, each providing insights into the structure and dynamics of community networks and their significance in understanding information propagation within communities.. We computed these network metrics using algorithms implemented in NetworkX¹⁴⁷ and Gephi¹⁴⁹. We discuss the metrics in the following paragraphs.

The number of edges in a network reflects the connections between nodes, indicating overall connectivity and the density of relationships within the community¹⁵⁰. More edges suggest a denser network with increased communication channels, fostering efficient information exchange and collaboration among members. Similarly, the number of triangles signifies closed loops formed by interconnected nodes, revealing the level of clustering within the network^{142;151}. A high level of clustering can facilitate local information sharing and collaboration among closely-knit groups.

The average clustering coefficient quantifies how closely nodes cluster in the network, indicating the prevalence of tightly interconnected groups¹⁴². A higher coefficient suggests a greater tendency for clustering, facilitating efficient information diffusion and collaboration. Meanwhile, the average path length measures the average shortest distance between all node pairs¹⁵²,signifying information spread efficiency. A shorter path length enables faster communication and coordination, crucial for community disaster resilience as critical information spreads quickly among individuals, enhancing overall resilience.

The average degree signifies the average connections per node¹⁴⁹, reflecting network connectivity and influence. A higher average degree suggests stronger communication networks, aiding information dissemination during emergencies. On the other hand,the average weighted degree offers a nuanced measure, considering both connection number and strength¹⁴⁹. It highlights nodes with stronger relationships, facilitating effective information exchange and collaboration within the community network.

5.3.5 Diffusion Time Prediction with Gaussian Process Regression

Communities need to identify target areas for the improvement of communication resilience during disasters. We can approach this problem as a regression task where community features help to predict diffusion time. However, with only ten communities, we need more data points. We generate synthetic features using the original ten communities to increase the number of data points, focusing on features during disaster events. We generate 1000 data points using the algorithms explained in the following subsection. After generating the synthetic graphs and features, we simulate them using Monte Carlo simulations and obtain the mean and standard deviation of diffusion time from 1000 simulation runs. Afterward, we use a Gaussian Process Regression Model to predict the diffusion time for the ten original communities using features obtained through synthetic data generation. We split the dataset into train and test by a ratio of 80:20. We use the sum of DotProduct and WhiteKernel as the kernel of the regression model. We discuss the results in the Results section.

We utilize two algorithms for generating synthetic graphs that closely resemble the original data. Both employ kernel density estimation to learn the degree distribution from the initial survey data. Algorithm 2 applies slight perturbations subsequently to the mean and standard deviation of the degree sequence to introduce variability. Algorithm 3 adjusts histogram bin heights. Sampling from these distributions generates a degree sequence of size $d = 1000$. Further details and implementation steps can be found in the Appendix B. Furthermore, we generate synthetic trust and interaction data to label graph edges, aiming to mimic the original community features closely. Initially, we adjust the relative frequencies of ordinal variables by ± 0.1 and renormalize them. Trust and interaction values are then generated based on these adjustments. For detailed implementation, see Algorithm 4 in the Appendix B. All codes are available on this GitHub repository¹⁵³.

5.4 Results

We discuss the results and provide answers to the questions in the [Introduction](#) using the data and experimental results.

5.4.1 Network Properties of Information Networks

The structure of community networks serves as a crucial determinant in understanding the dynamics of different communities, particularly during normal and disaster situations. Throughout our analysis of the network structures, we observe variations in network properties across different community types and between normal and disaster situations. These variations highlight the nuanced dynamics of community networks and their implications for information diffusion and community resilience.

Table 5.2: Network Properties of Information Network

Community (5005 nodes)	Normal Times						Disaster Times					
	Number of Edges	Number of Triangles	Average Clustering Coefficient	Average Path Lengths	Average Degree	Average Weighted Degree	Number of Edges	Number of Triangles	Average Clustering Coefficient	Average Path Lengths	Average Degree	Average Weighted Degree
Buchanan Rural	45751	87723	0.277	1.997	18.282	0.46	46548	95353	0.294	1.996	18.601	0.468
Buchanan Suburban	25322	19388	0.262	2.033	10.119	0.218	25941	20789	0.282	2.011	10.366	0.252
Platte Suburban	39272	62845	0.345	2.002	15.693	0.431	38402	62158	0.372	2	15.345	0.421
Platte Urban	38185	50716	0.239	2.027	15.259	0.373	38587	57622	0.263	2.006	15.345	0.372
Riley Rural	46361	83500	0.215	2.006	18.526	0.485	48307	94272	0.231	1.999	15.419	0.511
Riley Suburban	35351	47340	0.325	2.005	14.126	0.341	36044	53147	0.355	1.997	19.303	0.372
Riley Urban	30538	32017	0.279	2.031	12.203	0.293	30664	34468	0.312	2.002	14.403	0.314
Rural	45736	83108	0.254	2.002	18.276	0.469	47336	95551	0.269	1.998	12.253	0.493
Suburban	33577	42120	0.312	2.01	13.417	0.324	33680	44299	0.340	1.999	18.915	0.347
Urban	34511	42839	0.268	2.023	13.791	0.335	34542	46396	0.296	2.003	13.803	0.348

Our analysis encompasses several key network properties, as outlined in [Table 5.2](#). Across the various communities, we discern the following trends:

In Buchanan Rural, during disaster times, there is a slight increase in both the number of edges and triangles compared to normal times. This suggests that there might be more interactions or connections among nodes in the network during disasters, leading to higher

Figure 5.3: Diffusion Times reaching to 90% of Population

Community	Normal Time		Disaster Time	
	mean (hours)	std	mean (hours)	std
Buchanan Rural	22.366	0.12	20.447	0.111
Buchanan Suburban	38.341	0.279	32.016	0.208
Platte Suburban	22.091	0.059	22.069	0.057
Platte Urban	27.091	0.183	26.767	0.163
Riley Rural	27.784	0.21	24.360	0.158
Riley Suburban	25.784	0.118	22.977	0.069
Riley Urban	30.038	0.192	25.843	0.155
Rural	24.395	0.152	21.881	0.091
Suburban	27.469	0.145	24.600	0.115
Urban	27.966	0.169	25.091	0.135

clustering coefficient. Similarly, Buchanan Suburban experiences an increase in the number of edges and triangles during disasters, alongside a more noticeable rise in the average clustering coefficient. This indicates a denser network structure and potentially stronger local connections among nodes.

Platte Suburban shows a relatively stable number of edges between normal and disaster times, with a slight decrease in the number of triangles during disasters. Nonetheless, the average clustering coefficient increases, implying a more clustered network structure during disasters. Platte Urban witnesses a slight increase in both the number of edges and triangles during disaster times. Notably, the average clustering coefficient sees a significant increase, suggesting a more tightly knit network with stronger local connections among nodes.

Riley Rural experiences an increase in both the number of edges and triangles during disasters, indicating a denser network structure. However, the average clustering coefficient remains relatively stable, suggesting that the increase in connections may not necessarily result in stronger local clustering. Similarly, Riley Suburban sees an increase in both the number of edges and triangles during disasters. However, the average clustering coefficient sees a significant increase, indicating a more clustered network structure with stronger local connections. In Riley Urban, there is a slight increase in both the number of edges and

triangles during disaster times. Notably, the average clustering coefficient witnesses a notable increase, suggesting a more tightly clustered network with stronger local connections among nodes.

In rural areas, during disaster times, there is an increase in both the number of edges and triangles, indicating a denser network structure. However, the increase in the average clustering coefficient is relatively small, suggesting a moderate increase in local clustering. Similarly, in suburban areas, there is an increase in both the number of edges and triangles during disaster times. However, the average clustering coefficient sees a more significant increase, indicating a denser and more clustered network structure during disasters. Finally, in urban areas, there is a slight increase in both the number of edges and triangles during disaster times. The increase in the average clustering coefficient is notable, suggesting a more tightly knit network with stronger local connections among nodes during disasters.

In general, average path lengths remain relatively stable across communities during both normal and disaster times, indicating consistent information transmission efficiency regardless of conditions. However, differences in community structures are evident across area types in our dataset. Variations in average path lengths suggest differing efficiency levels in information transmission within each area type. The number of edges varies, indicating denser networks in specific regions, while significant differences in the number of triangles suggest unique clustering patterns in each area type. Additionally, variations in average clustering coefficient and average degree reflect disparities in network cohesion and node connectivity specific to different areas. Similarly, variations in average weighted degree highlight diverse connection strengths inherent to specific regions.

5.4.2 Information Diffusion in Normal and Disaster Times

We measure a community’s communication resilience to disasters by the time it takes for information to diffuse to 90% of the population. Our study aims to answer whether there is a significant difference in the information flow graph between rural, suburban, and urban communities and between disaster and non-disaster scenarios. During the simulation, we ob-

tain the diffusion times and their standard deviations for all communities under study. Table 5.3 displays the results for different area types and districts, which show that information diffuses at varying speeds across all communities and diffuses faster in disaster situations than in normal daily life. Interestingly, rural communities have a faster diffusion rate than urban communities, even in disaster scenarios. These findings suggest significant variability in the information flow graph between communities.




Figure 5.4: Sensitivity Analysis of the Average Diffusion Time during Disaster (1a, 1b, 1c) and Normal (2a, 2b, 2c) Times

We conduct a sensitivity analysis to examine how sensitive our results are to changes in the discount, forgetting, and trust threshold values. The analysis results are presented in Fig. 5.4. We observe that the trends of the diffusion time plots are consistent for the discount, forgetting, and trust threshold values in both disaster and normal scenarios. Rural

communities have a significantly different diffusion time than urban and suburban communities. Interestingly, for certain forgetting factor values in normal times, we observed similar diffusion times in urban and suburban communities. Similarly, for the plot of the diffusion times using different trust thresholds in normal scenarios, urban and suburban communities are close.

5.4.3 Resilience to Information Hub Failure

We recognize that infrastructures are vulnerable, especially during disasters. Therefore, it is crucial to incorporate redundancies in infrastructure networks. For communication networks such as the one we are studying, it is essential to have multiple communication paths when some of the infrastructure fails. To investigate the impact of infrastructure failure on communities, we adjust our model by dropping one of the five hubs and observing the resulting diffusion time. We examine rural, suburban, and urban communities' response to the loss of a particular infrastructure since communities rely on specific communication infrastructure.

We compare these simulation results to our initial results when all hubs participate in the diffusion process. The results are in Fig. 5.5, which indicates that the diffusion time increases when a source of information fails, both in normal and disaster times. In normal times, the failure of local news seems to have the most significant impact on the time of information diffusion across all area types, with the most pronounced effect in urban areas, suggesting heavy reliance on social media in these areas. This is followed by the Government hub failure again, with urban areas showing the highest increase in diffusion time, indicative of their reliance on government communication channels. It is interesting to see that Online social networks and print news failures show a relatively less impact on the diffusion time, with the least effect seen in rural areas, potentially due to alternative information sources being utilized.

During disaster times, the failure of each communication hub generally results in a decrease in diffusion time compared to that of normal times, when a hub fails. We have

already shown that the diffusion speed decreases during disasters as compared to normal times. Hence, this situation follows the same pattern. The failure of Local News still has the most substantial impact as it shows the greatest increase, particularly in suburban areas. This is followed by the Government hub failure.

From a resilience standpoint, the findings imply that safeguarding Local News infrastructure and providing robust fail-over mechanisms for Government messaging will yield the greatest reductions in worst-case diffusion delays. Social media and print outlets should be treated as amplifiers within a layered communication strategy rather than as primary conduits. Prioritizing investments accordingly will both speed routine information flows and mitigate cascading delays during emergencies.

The greater spatial disconnection in urban areas compared to rural communities implies that urban dwellers may rely more on online social networks (OSN) than rural residents. To investigate this, we compare the diffusion time when the OSN fails while other hubs are active to the time when all hubs are functioning. Table 5.3 shows that the mean difference in diffusion time between all hubs working and when the online social network hub fails in urban communities during a disaster is higher than in rural or suburban communities. For disaster scenarios, absolute time in hours matters the most, more than the percentage change in time. Hence, we interpret this result to mean that a high difference indicates a high reliance on OSN. With this perspective we observe that, in a disaster scenario, urban communities rely on OSN more than other communities. However, in normal times, suburban communities are most affected when OSN fails. These findings suggest that urban communities tend to rely more on OSN compared to rural and suburban areas, with suburban communities being most affected during normal times when the OSN hub fails.

5.4.4 Impact of Population Size on Time of Diffusion

We tested the impact of the number of nodes on information diffusion times to ensure that our simulation results were not biased. We varied the number of nodes to include 1000, 5000, 10000, 15000, 20000, and 25000 nodes and kept the graph features constant for each




Figure 5.5: Average diffusion times comparing scenarios with one hub failure against all hubs working, under disaster and normal conditions.

Table 5.3: Reliance of Urban and Rural Communities on Online Social Networks

Community	Normal (All hubs)		Normal (OSN Fails)		Mean Difference	Disaster (All hubs)		Disaster (OSN Fails)		Mean Difference
	mean (hours)	std	mean (hours)	std		mean (hours)	std	mean (hours)	std	
Rural	24.395	0.152	26.392	0.161	1.997	21.881	0.091	23.698	0.122	1.817
Suburban	27.469	0.145	30.231	0.189	2.762	24.600	0.115	26.472	0.127	1.871
Urban	27.966	0.169	29.809	0.202	1.843	25.091	0.135	27.143	0.149	2.052

community. Our findings showed no significant differences (p-value=0.1) in the diffusion times across different population sizes of similar communities, as presented in Table 5.4.

Table 5.4: Comparison Using Different Number of Nodes in Network

Size	Platte Suburban (Disaster and No Government Hub)		Rural Disaster		Riley Urban Normal	
	mean (hours)	std	mean (hours)	std	mean (hours)	std
1000	25.409	0.293	22.038	0.222	29.771	0.407
5000	24.860	0.242	21.881	0.091	30.038	0.192
10000	24.901	0.059	22.015	0.033	29.723	0.13
15000	24.913	0.107	21.963	0.027	29.835	0.101
25000	25.000	0.003	21.986	0.01	29.635	0.135

5.4.5 Gaussian Process Regression Model Results

The Gaussian Process Regression model provides us with a reliable estimate of the time it takes for information to spread, along with the uncertainty surrounding this estimate. We obtain the average diffusion time for test data with corresponding standard deviations. The mean absolute percentage error (MAPE) for the regression is 4.574%, which suggests a reasonable forecasting performance for the regression model. Using our model, we make predictions for the ten communities as shown in Table 5.5. We observe that our predictions closely align with actual diffusion times.

For good response planning, we need to understand the speed at which information disseminates during a disaster. Accurate predictions enable community leaders and emergency responders to make informed decisions, prioritizing areas that require immediate attention. The use of original community data ensures the applicability of our predictions to real-world scenarios. This direct connection to community data enhances the value of our forecasts for decision-makers, providing insights grounded in actual community dynamics. Additionally, our predictions can be utilized to simulate various disaster scenarios, offering communities the opportunity to assess and refine their response strategies. By comprehensively understanding how different scenarios influence communication and response times, communities

Table 5.5: Gaussian Process Model Predictions for the 10 community groups.

Community	Average Time of Diffusion (Hours)	Standard Deviation	Predicted Time of Diffusion (Hours)	Standard Deviation
Buchanan Rural	20.660	1.393	20.447	0.111
Buchanan Suburban	33.007	1.402	32.016	0.208
Platte Suburban	22.995	1.395	22.069	0.057
Platte Urban	27.334	1.394	26.767	0.163
Riley Rural	25.222	1.393	24.360	0.158
Riley Suburban	24.866	1.398	22.977	0.069
Riley Urban	26.625	1.395	25.843	0.155
Rural	21.623	1.394	21.881	0.091
Suburban	25.444	1.405	24.600	0.115
Urban	25.744	1.396	25.091	0.135

can bolster their preparedness efforts.

Budget Allocation Methodology using Gradients

A community’s unique characteristics may be crucial to improve its disaster resiliency strategies. Using the resilience metrics of diffusion time, we provide a budget allocation method for a community interested in improving the time of diffusion during a disaster. The average diffusion time T is a multi-variable function $\bar{f}(p_1, p_2, \dots, p_n)$, learned through the Gaussian Regression model, with parameter vector $p = [p_1, p_2, \dots, p_n]^T$, which is the feature vector of each community.

We use a gradient approach to develop the budget allocation method. We calculate the partial derivative, $\nabla \bar{f}(p)$, as below:

$$\nabla \bar{f}(p) = \left[\frac{\partial \bar{f}(p)}{\partial p_1} \quad \frac{\partial \bar{f}(p)}{\partial p_2} \quad \dots \quad \frac{\partial \bar{f}(p)}{\partial p_n} \right]^T = \left[\frac{\bar{f}(p_1 + \delta p_1, p_2, \dots, p_n) - \bar{f}(p_1, p_2, \dots, p_n)}{\delta p_1} \quad \dots \quad \frac{\bar{f}(p_1, p_2, \dots, p_n + \delta p_n) - \bar{f}(p_1, p_2, \dots, p_n)}{\delta p_n} \right]^T \quad (5.2)$$

We used $\delta p_i = 1$ for each feature p_i in the vector space, except in one case, where the feature’s magnitude was high and we used a value of 10. We calculate the gradient at the values of the feature vectors specific to each community. To evaluate $\bar{f}$ in the modified values of the feature vector (i.e., $\bar{f}(p_1 + \delta p_1, p_2, \dots, p_n)$, $\bar{f}(p_1, p_2 + \delta p_2, \dots, p_n)$, $\dots$, $\bar{f}(p_1, p_2, \dots, p_n +$

δp_n)) we use the developed Gaussian Regression model.

See Appendix B, Table 1 for the gradient results. From our results, we find that the five most critical features, in order of importance, are the Median Frequency of Interaction with Cable News, the Median Trust for Online Social Networks, the Median Frequency of Interaction with the Local Government, the Mean of the Degree distribution of the Community Network, and the Median Trust for People. These listed features correspond to the highest (in absolute value) gradient components for each community and could be where the communities might want to focus.

To guide community investments based on the calculated gradient, we propose the following simple strategy. Let the vector $\Delta = [\Delta_1, \Delta_2, \dots, \Delta_n]^T$ denote the changes in the community feature vector achievable by the investment of one monetary unit. – e.g., Δ_i can be the change in p_i achieved by investing one dollar in changing p_i (increasing it or decreasing it, depending on whether the gradient component $\frac{\partial f(p)}{\partial p_i}$ is negative or positive, respectively). Then the gradient with respect to investments, and which, after normalization (so that the sum of the components becomes 1) should inform how one dollar is spent, is given by $\nabla \bar{f}(p) * \Delta$, where $*$ denotes the Hadamard product. One single dollar should then be split amongst the community feature components as indicated by the vector $\frac{1}{|\nabla \bar{f}(p) * \Delta|_1} \nabla \bar{f}(p) * \Delta$, where $|\cdot|_1$ denotes the L_1 norm.

5.5 Discussion

The diffusion times during normal and disaster scenarios, as presented in Table 5.3, highlight variations among different community types. In regular circumstances, diffusion times range from approximately 22.09 to 38.34 hours, indicating gradual information spread. During disasters, diffusion times decrease across all community types, suggesting faster information dissemination, albeit with some variability based on community context. These findings emphasize the intricate relationship between community attributes and information diffusion dynamics, underscoring the need for tailored communication strategies.

The Gaussian Process Regression (GPR) model results provide reliable estimates of dif-

fusion time, with a low mean absolute percentage error (MAPE) of 4.574%, indicating reasonable forecasting performance. By predicting diffusion times for different communities, valuable insights for response planning are provided, enabling leaders and responders to prioritize areas needing immediate attention. Moreover, the model's accuracy is demonstrated through close alignment with observed diffusion times.

The gradient analysis identifies critical factors influencing diffusion time across communities. Key influencers include interaction frequency with cable news and online social networks, trust in the local government and people, and network properties. These findings highlight the importance of social dynamics and communication channels in shaping community resilience during disasters. For instance, communities with higher interactions with local government may exhibit faster information diffusion, as residents are more likely to heed official warnings and instructions. Similarly, frequent interactions with cable news and online social networks can facilitate the rapid dissemination of critical information, enabling residents to stay informed and take appropriate actions during emergencies. Communities can utilize this information to prioritize investments or interventions targeting specific features to enhance communication resilience. By allocating resources based on gradient analysis, communities can optimize preparedness efforts, strengthening overall resilience.

Analyzing the relationship between network properties and diffusion time reveals valuable insights for disaster resilience strategies. Rural communities generally exhibit smaller diffusion times compared to suburban and urban counterparts, indicating potentially more efficient information dissemination processes during disasters. Certain network features, such as degree distribution mean, interaction frequency with cable news, and trust in online social networks, play crucial roles in shaping diffusion dynamics across different community types. Rural areas often display negative gradient values for these features, suggesting that increasing interaction frequency or trust levels in specific communication channels could expedite information spread. Conversely, suburban and urban communities show mixed gradient values, implying a more nuanced relationship between network properties and diffusion time. These findings stress the importance of tailoring communication strategies to community-specific characteristics. By leveraging insights from diffusion time analysis and gradient

results, stakeholders can prioritize investments in communication infrastructure and information exchange channels to mitigate the impact of disasters on vulnerable communities. Accurate modeling and analysis are crucial for informing effective disaster response and resilience strategies, empowering communities to enhance their resilience and minimize the impact of disasters on their populations.

5.6 Conclusion

This study highlights the need for tailored communication strategies because of these differences in communication patterns among different communities. We further emphasize that certain features are more important for each community, and investments in social programs and infrastructure should focus on improving interactions with cable news and local government and trust for online social networks. Also, we find that urban communities rely more on online social networks than rural communities for information dissemination. Moreover, the failure of one communication infrastructure affects each community differently, while the time of diffusion is independent of the size of the community.

Our study provides valuable insights for community partners and policymakers to develop effective disaster communication strategies and improve community resilience. However, it has some limitations. Firstly, the study only focuses on one type of disaster event within three US Midwestern regions, which may limit the generalizability of the findings. Future studies could explore the influence of community features in different disaster contexts and geographical locations to enhance the generalizability of the results. Secondly, this study relies on self-reported data, which may be subject to bias and inaccuracies. Future studies could use more objective measures to collect data, such as social media activity or mobile phone usage patterns. Thirdly, there was a low response rate for the survey during the COVID-19 pandemic, which might violate statistical validity based on the lack of sample size. Additionally, the absence of data on shared contacts between individuals limits the depth of our analysis. Future research could aim to address this limitation by exploring novel methodologies for collecting such data. Also, we note that another limitation is that

the single-step gradient approach, while useful for developing a budget allocation method for policy recommendations, may oversimplify the complexity of policy implementation, potentially neglecting essential factors critical for effective execution. Finally, future studies could delve deeper into the mechanisms to better understand how specific community features influence information diffusion during a disaster. Despite these limitations, this study's findings offer valuable insights into how to improve communication strategies for disaster resilience in communities.

Chapter 6

Predicting Topic Engagement during Russian Troll Activity around the 2016 US Political Elections

6.1 Introduction

Online social networks have become an essential part of people's daily lives. These platforms, such as Facebook, Twitter, and Instagram, have provided an avenue for individuals to express themselves, connect with others, and share information. However, the ease of access and anonymity offered by these platforms has also led to the emergence of online trolls, who seek to manipulate, harass, and intimidate others for their own personal gain.

Trolls are individuals who engage in online behavior that is intended to provoke, upset, or offend others. They may use various tactics, such as spreading false information, posting offensive comments or images, and engaging in cyberbullying. Trolls can operate under different personas and may even create fake accounts or use multiple accounts to hide their true identity.

Manipulation by trolls in online social networks has become an increasingly concerning issue, particularly due to its potential to cause harm to individuals, damage reputations, and

create a toxic online environment. Manipulation occurs when someone is taking deliberate actions that subvert the individual’s ability to make a well-informed decision¹⁵⁴ or when someone’s prior opinions and cognitive states are altered severely by another person¹⁵⁵. In this work, we are interested in online manipulation that happens on social networks. Susser et al.¹⁵⁶ define online manipulation as using information technology to influence people’s behavior by targeting their weaknesses. Various manipulation techniques exist, including online astroturfing, misinformation, disinformation, clickbait, deepfake, gaming trending and ranking algorithms, and coordinating attacks^{11–17}. Online astroturfing involves distributing coordinated information to create false impressions of support, while misinformation and disinformation entail spreading incorrect or misleading information intentionally. Clickbait employs sensational headlines to attract clicks, and deepfake alters original content using software and AI tools. Gaming trending and ranking algorithms manipulate algorithms for visibility, and coordinating attacks involves synchronized efforts to achieve malicious goals.

Our study contributes to understanding the dynamics of user engagement with troll tweets on social media platforms, utilizing cluster analysis and median response time calculations to uncover interaction patterns. We find a diverse landscape of tweet distribution across clusters, indicating varying user engagement patterns. We find clusters characterized by low median response times were not consistently prevalent across sequences, suggesting no consistent strategy by trolls to elicit quick responses. Additionally, mentions of political figures by trolls tend to garner more responses.

6.2 Literature Review

Manipulation of people’s opinions is a timely topic in online social networks, especially in political circles. Hence, several studies have investigated manipulation techniques on different social media platforms. In this section, we look at some of the related works on manipulation

6.2.1 Theoretical Models of Manipulation

There are several theoretical models of manipulation on social networks. Förester et al.¹⁵⁷ provided a theoretical model using DeGroot’s opinion dynamics model, which shows how agents manipulate social trust to increase their influence on long-run beliefs, thus gaining opinion leadership. Etesami et al.¹⁵⁸ investigated manipulative agents’ impact on public opinion using a game-theoretic approach to model and simulate the change in public opinion when individuals display conformist or manipulative actions. The downside of these theoretical models is a lack of validation with real data, as manipulation in online social networks can be too complex for simple theoretical models.

6.2.2 Empirical Models of Manipulation

We explore additional manipulation models in online social networks. Nimmo¹⁵⁹ introduced the coefficient of traffic manipulation (CTM) to assess manipulation within Twitter traffic, finding that manipulated traffic has a higher CTM compared to normal patterns.

Benigni et al.¹⁶⁰ analyzed social influence bots’ impact on Twitter communities, noting their engagement in repetitive mentions, leading to data manipulation. Paquet-Clouston et al.¹⁶¹ investigated Social Media Fraud (SMF) on online social networks, involving the creation of fake accounts or using compromised real accounts to falsify social media data.

Hussain et al.¹⁶² utilized social network analysis to study commenter behavior on YouTube channels promoting conspiracy theories, identifying patterns of manipulated data and bot-like behaviors in communities. Assenmacher et al.¹⁶³ proposed a two-phase framework for detecting and analyzing manipulation campaigns and trending topics, emphasizing human-in-the-loop decision-making for transparent analysis and countering manipulation.

6.2.3 Political Discourse Manipulations on Twitter

During the 2016 elections on Twitter, identifying common features among these accounts and their content. They found that topics of tweets were effective in identifying a significant portion of these accounts, which were responsible for the majority of English-language content.

Kim et al.¹⁶⁴ introduced the time-sensitive semantic edit distance (t-SED) method, that emphasized temporal context in the classification of Russian trolls in US political discourse on Twitter. Monakhov¹⁶⁵ proposed a language-centric measure for swift and accurate troll identification across languages. Chun et al.¹⁶⁶ utilized various machine learning models to detect Russian trolls’ impact on the 2016 US Presidential elections and their model achieved high accuracy in distinguishing troll from non-troll content. Linvill et al.¹⁶⁷ analyzed Russian trolls’ communication strategies prior to the 2016 US Presidential elections, identifying seven distinct behaviors. Badawy et al.¹⁶⁸ examined the influence of Russian trolls on the 2016 US Presidential elections, finding varying retweet frequencies among political affiliations. Ferrara et al.¹⁶⁹ characterized bot influence on the 2020 US Presidential elections, revealing organized efforts to manipulate political opinions on Twitter. Deb et al.¹⁷⁰ analyzed Twitter conversations during the 2018 US Midterm elections to detect interference in political discourse. They discovered statistical irregularities between election day tweets and actual voting results, cautioning against drawing conclusions solely from social media data due to biases like population bias and bot activity. Cardoso et al.¹⁷¹ investigated bot behavior leading up to the 2018 US Midterm Election, finding that bots strategically mimicked human temporal patterns and engaged with both human and bot populations. They also observed a significant proportion of human retweets originated from bot-generated content.

6.3 Methodology

This study employs a longitudinal study and clustering and statistical analysis methods to analyze interactions between Russian trolls and legitimate Twitter users during the 2016 US presidential elections. We detail our methodology, including data collection, cleaning, and preprocessing.

6.3.1 Data Collection

Our primary data source is a public dataset from Clemson University researchers, made available by FiveThirtyEight on GitHub¹⁷². This dataset contains 2,973,371 tweets from 2,848 accounts linked to the Internet Research Agency, identified as Russian trolls. We utilized twelve files, focusing on tweets that provided comprehensive metadata.

6.3.2 Data Preprocessing

From the dataset, we extracted relevant information from six specific columns: external author ID, the author’s username, tweet ID, tweet content, language, and publication date and time. We filtered the dataset to include only English-language tweets and extracted hashtags from the content for further analysis. This process identified 1,173 unique troll accounts.

Due to the removal of these accounts and their tweets from Twitter, direct data on responses and network interactions was inaccessible. To overcome this limitation, we used the Twint Python library to search for mentions of these trolls on Twitter by legitimate users, based on names derived from the troll dataset. This search yielded 39,149 tweets from what we assume to be legitimate users, spanning from February 8, 2012, to May 30, 2018. We stored this supplementary data in MongoDB.

To align tweets from legitimate users with those from trolls, we filtered tweets to include only those with references in the "referenced_tweets_id" and "in_reply_to_user_id" columns. We further narrowed this dataset by ensuring that the user IDs in "in_reply_to_user_id" matched a troll ID from the initial dataset. This resulted in a subset of 16,666 tweets.

We then applied a temporal filter to only retain tweets posted after the corresponding troll tweets, focusing on 3,499 tweets that potentially responded to initial troll activity.

Following this, we computed the popularity of hashtags across our entire dataset, encompassing both trolls and legitimate users, as this would serve as a key feature in our model. This involved collecting hashtags from both user groups and their respective tweet dates. We calculated the hashtag popularity by computing the number of times these hashtags

occurred in tweets during the seven days preceding the user’s post on the same topic.

For text preprocessing, we removed URLs, user mentions, and punctuation and converted all text to lowercase. We also used the emoji library in Python to replace emojis with corresponding text descriptions, which assists in more accurate sentiment and emotion analysis of the tweets.

6.3.3 Feature Extraction

This section details the features extracted from the dataset to analyze the dynamics of interactions between legitimate users and trolls. We hypothesize that responses from legitimate users to trolls are influenced by a combination of message content and the cumulative features that accumulate to a threshold prompting a user response. Each troll tweet is considered an originating message in this context.

The following features were extracted to model the potential of a user’s response to a troll’s tweet:

1. **Sentiment Analysis:** Using the Natural Language Processing Toolkit (NLTK), we assessed the sentiment of each tweet. This analysis produced three sentiment scores—negative, neutral, and positive—for each tweet, ensuring that the sum of these scores equals 1.
2. **Emotion Detection:** We employed an emotion prediction model referenced in [173](#) to classify tweets into one of six emotions: anger, disgust, fear, joy, sadness, or surprise.
3. **Tweet Length:** We measured the length of each tweet by counting the number of words it contained, which can indicate the intensity or complexity of the communication.
4. **Popularity of Hashtags:** To understand the influence of hashtags, we tracked the frequency of hashtags used by trolls and legitimate users within the seven days prior to each tweet. This metric reflects the reach and relevance of the hashtags.
5. **Time to Reply:** This metric quantifies the delay in response by measuring the time elapsed between a troll’s tweet and the subsequent response from a legitimate user. This feature helps identify the immediacy of the interactions.

Additionally, we considered the user’s tweet history as an accumulation of exposure to various tweets, including past interactions with trolls and responses from other legitimate users. This historical context forms a base for each legitimate user’s potential to respond. From these considerations, we derived a total of 12 new features, which are comprehensively listed and described in Table 6.1.

Table 6.1: Cumulative Features

Feature	Description
History of Negative Sentiments	The sum of the negative sentiments from tweets of the original troll to legitimate users that replied to that troll before the current user.
History of Neutral Sentiments	The sum of the neutral sentiment from tweets of the original troll to legitimate users that replied to that troll before the current user.
History of Positive Sentiments	The sum of the positive sentiment from tweets of the original troll to legitimate users that replied to that troll before the current user.
History of Length of Tweets	Uses the sum of the length of tweets from the original troll to legitimate users that replied to that troll before the current user.
History of Popularity of Hashtags	Uses the popularity of hashtags contained in tweets from original troll to legitimate users that replied to that troll before the current user.
History of Anger Emotion	The Sum of the anger emotion from tweets of the original troll to legitimate users that replied to that troll before the current user.
History of Fear Emotion	Sum of fear emotion from tweets of the original troll to legitimate users that replied to that troll before the current user.
History of Disgust Emotion	Sum of disgust emotion from tweets of original troll to legitimate users that replied to that troll before current user.
History of Joy Emotion	Sum of joy emotion from tweets of original troll to legitimate users that replied to that troll before current user.
History of Sadness Emotion	Sum of sadness emotion from tweets of original troll to legitimate users that replied to that troll before current user.
History of Surprise Emotion	Sum of surprise emotion from tweets of original troll to legitimate users that replied to that troll before current user.
Time to Reply	The difference between the time the current user and the last user (maybe troll or legitimate) tweeted—time in hours.

6.3.4 Clustering

In our analysis, each legitimate user has history-based features that we treat independently. This approach allows us to apply clustering to identify groups of data points of tweet features with similar histories. Our aim is to uncover underlying patterns and trends within the data

that relate to how users respond to interactions with trolls on Twitter. We hypothesize that legitimate users with similar historical exposure to trolls might exhibit common behaviors and responses. Therefore, by grouping data points of tweet features based on their historical patterns, we can draw clearer and more interpretable conclusions about user reactions to troll interactions.

For our clustering, we make use of 11 of the features, excluding the "Time to Reply" feature. During our initial data exploration, we identified outliers in the dataset using the Interquartile Range (IQR) method. We applied the k-means clustering and considered variations that included or excluded the first Principal Component Analysis (PCA) component and outliers. To ensure that clustering was not solely influenced by the length of histories but instead focused on their attributes, we opted to decorrelate the PCA by removing its first component. We then performed dimensionality reduction using PCA, reducing the dataset to six principal components, which simplified the data and provided a clearer foundation for subsequent clustering.

We chose to test clustering with 40, 45, and 50 clusters to capture nuanced insights into the patterns associated with similar historical data. The decision to use these specific numbers of clusters was guided by preliminary tests, which suggested that these sizes were optimal for revealing distinct patterns without overfitting or excessive fragmentation of the data. Smaller clusters help identify more detailed patterns that might be obscured in larger clusters.

To further refine our clustering strategy, we implemented four different methods, resulting in a total of twelve combinations to analyze. The four varieties are:

1. Method 1: cluster with the first principal component with outliers
2. Method 2: cluster with the first principal component without outliers
3. Method 3: cluster without the first principal component with outliers
4. Method 4: cluster without the first principal component without outliers

Each method was selected to explore how the presence or absence of principal components

and outliers affects the clustering outcome, thus allowing us to determine the most effective approach for analyzing our specific dataset.

6.3.5 Hypothesis Testing

In our study, we utilized a simulation-based Kolmogorov-Smirnov goodness-of-fit test to identify the most suitable distribution for modeling the time to reply in our experiment. We considered different distributions: exponential and Pareto. For this purpose, we generated random numbers matching the size m of each cluster from both the Exponential and Pareto distributions, using parameters that align with their hypothesized characteristics. For the Pareto distribution, these parameters were the shape and minimum value, while for the exponential distribution, it was the mean. We repeated this simulation process 10,000 times to ensure accuracy.

For each iteration, we calculated the Kolmogorov-Smirnov statistic by comparing the empirical cumulative distribution function (CDF) of the simulated data set to the theoretical CDF. Following this, we computed the KS statistic for our actual data set and compared it with the KS statistics from the simulations. We derived the p-value by identifying the proportion of simulated KS statistics that exceeded the KS statistic of the actual data. A p-value threshold of 0.05 was used to determine if the fit was statistically significant.

We use the results from these distribution tests as a metric in selecting the most effective clustering method for our analysis. Our findings show that the Pareto Distribution was the most successful suggesting it was a better fit than the exponential distribution. In particular, Clustering Method 3 with Cluster 40 achieved the highest pass rate at 95%.

The detailed results for both the Pareto and exponential distribution tests across various clustering methods are presented in Tables 6.2 and 6.3. The percentages indicate the proportion of clusters, out of groups containing 40, 45, and 50 clusters, that passed each respective test.

Some clusters did not pass the initial distribution tests for the Pareto and exponential distributions. However, the Pareto distribution was the best fit for our data. After compre-

Table 6.2: Cluster Distribution Test-Pareto

Clustering Method	Clusters Passed		
	40	45	50
Method 1 Clustering	87.5%	93.3%	90%
Method 2 Clustering	87.5%	88.9%	92%
Method 3 Clustering	95%	91.1%	88%
Method 4 Clustering	87.5%	80%	86%

Table 6.3: Cluster Distribution Test-Exponential

Clustering Method	Clusters Passed		
	40	45	50
Method 1 Clustering	27.5%	26.7%	32%
Method 2 Clustering	32.5%	31.1%	40%
Method 3 Clustering	22.5%	26.7%	32%
Method 4 Clustering	22.5%	22.2%	34%

hensive testing, we selected Clustering Method 3, specifically Cluster 40 using the Pareto distribution, as our primary metric due to its better performance. Two clusters that did not fit the Pareto. These clusters were clusters 39 and 2, which had 426 and 94 data points, respectively. Cluster 39 represents a significant portion of our data and mostly contains tweets with only one response. Since our objective is to understand the broader patterns of user behavior, we chose to keep these clusters. Following this, we calculated the median time to reply for each cluster. We use the median rather than the mean because the median is less sensitive to extreme values. We also applied additional tests using log-normal and beta distributions, but Pareto seems to be the best for our data. The data and code are available in the project’s GitHub repository¹⁷⁴.

6.4 Results

6.4.1 Analysis of Clustering Results

Figure 6.1 shows the distribution of tweets across all clusters. We observe that cluster 39 has the highest frequency, followed by cluster 21. We observe that cluster 39 comprises mostly tweets where only one user responded to a troll’s tweet. From our data, we see that many tweets fall into this category.

Next, we observe the sample median and statistical median (from the fitted distribution)




Figure 6.1: Frequency of Tweets in Each Cluster

response times for each cluster in Figure 6.2. The blue bars, representing the statistical median, generally appear slightly higher than the orange bars (sample median). There is considerable variation in response times across clusters. Some clusters exhibit very rapid responses, with median times under an hour. Others show much slower response times, exceeding 5 hours. This suggests different user behaviors and engagement patterns across the clusters. Also, across most clusters, the statistical median derived from the fitted Pareto distribution tends to be slightly higher than the sample median calculated directly from the observed data. This consistent difference hints at the influence of the Pareto distribution’s characteristics on the statistical median calculation.




Figure 6.2: Statistical Median and Sample Median Response Time for Each Cluster

6.4.2 Median Response Time of Tweet Sequence

To begin, our objective is to generate a super-sequence comprising all tweets from legitimate users. A super-sequence refers to a sequence that encompasses or contains multiple other

sequences within it. Here, a super-sequence is a sequence of tweets that encapsulates or combines the sequences of clusters observed within tweet sequences.

For every tweet initiated by a troll, we trace the subsequent responses from legitimate users. These responses are then followed by additional replies from other users, forming a chronological sequence. This process captures the progression of interactions from trolls to legitimate users and subsequently includes contributions from other participants over time. Afterward, we construct another sequence indicating the succession of clusters to which each tweet belongs. This is feasible as we have assigned every legitimate user's tweet to a cluster. This results in a sequence of clusters. By utilizing these cluster sequences, we generate super-sequences that encompass the observed tweet cluster sequences. As a result, we obtained 63 super-sequences of tweets, providing a comprehensive view of the clustering patterns observed within the tweet sequences. Figure 6.3, 6.4 and 6.5 depict 6 such super-sequences where the x-axis represents the tweets identifier over time, and the y-axis indicates the sample and statistical median response time, where the statistical median response time is based on the distribution of the clusters.

Across all sequences, we observe that response times (measured as both statistical and sample medians) exhibit similar trends. While some sequences display predictable and consistent patterns, others are highly erratic, with sudden spikes or drops. Response times range from near-zero values to peaks exceeding 4 hours, reflecting a diverse engagement dynamic. Also, we observe that the response time varies quite significantly across different tweets. This is evident from the large fluctuations in the y-axis values across all graphs. This suggests a wide disparity in how quickly users react to different tweets.

Some graphs show a more regular pattern (e.g., periodic spikes), while others are more erratic. The regular patterns might indicate a consistent behavior in response times for the associated tweets. Some graphs exhibit high isolated peaks, suggesting a sudden increase in response time that quickly returns to a lower level. This could correspond to tweets that people did not engage with initially but later received attention. If we compare the tweets, we can deduce that some tweets have a more significant impact on response times than others. For example, sequence 3424 follows a sequence of low response time after an initial

high peak.

This analysis reveals a lack of consistent behavioral patterns, indicating the lack of a clear strategy by trolls to prompt quick responses from users. Some clusters exhibit rapid responses regardless of interaction frequency, while others demonstrate longer response times despite encountering trolls less frequently. This variance suggests that user responses to trolls may be influenced by factors beyond just the volume of troll interactions.





Figure 6.3: Sequence of Median Response Time of Tweet Sequences 397 and 3084





Figure 6.4: Sequence of Median Response Time of Tweet Sequences 3424 and 3479





Figure 6.5: Sequence of Mean Response Time of Tweet Sequences 3489 and 3491

6.4.3 Analyzing Response Time Trends

This subsection explores whether trolls implemented strategies to enhance legitimate user's responsiveness to tweets over time. By comparing statistical median and average response times, we seek to uncover patterns indicative of deliberate engagement tactics. To achieve

this, we identify troll sequences—series of tweets posted by a troll account over a specific timeframe—along with the responses to those tweets.

Within each troll sequence, we analyze the average time it takes for the first two responses to a troll’s tweet for every tweet they post. Additionally, we examine the average statistical median response time of the first two responses to a troll’s tweets throughout the sequence. We anticipate that over time, the troll’s proficiency in crafting tweets that elicit higher engagement or quicker responses will improve. Each point within this sequence represents the average response time of the first two responses to a troll’s tweet or the median response time to a troll’s tweet. Our analysis shows more sporadic movements using the average response time with no clear trends; hence, we focus on the median response time. Figures 6.6, 6.7, 6.8 illustrate three troll sequences displaying decreasing, increasing, or no change patterns over time using the median response time. In Table 6.4, we observe that only one sequence exhibited a decreasing trend, while four showed an increasing trend, with the majority remaining stable or showing no apparent changes. Additionally, we identified 90 troll sequences with only a single tweet following the troll’s initial tweet.




Figure 6.6: Response Time of First Two Responses to Troll’s Tweets Over Sequence (Decreasing Pattern)

Table 6.4: Analysis of Median Response Time Trends

Trend	Number of Sequences
Decreasing	1
Increasing	4
Stable/No Change	58




Figure 6.7: Response Time of First Two Responses to Troll’s Tweets Over Sequence (Increasing Pattern)




Figure 6.8: Response Time of First Two Responses to Troll’s Tweets Over Sequence (Stable Pattern)

While our analysis of median and average response times and examination of trends within troll sequences revealed varying patterns, it is challenging to conclusively determine whether trolls implemented deliberate engagement tactics. While some troll sequences exhibited decreasing, increasing, or stable response time trends, the presence of troll sequences with only a single subsequent tweet suggests potential limitations in the effectiveness of such strategies.

6.4.4 Content Analysis of Troll Tweets

In this subsection, we undertake a content analysis of troll tweets, specifically the text, to discern the prevalent topics utilized by trolls. We employed Latent Dirichlet Allocation (LDA) to uncover distinct topics within troll tweets, distinguishing between those with numerous

responses and those with just one response. The results of this analysis, including the top words associated with each identified topic, are presented in Table 6.5 and Table 6.6.

From the two tables, we observe that the topics suggest the trolls are engaging in divisive, sensitive, and highly debated social and political issues, potentially amplifying existing tensions. We notice there is an overlap in themes related to "blacklivesmatter" across the tables, indicating its significant presence in the discourse within troll tweets. However, we notice Table 6.5 contains more specific issue-oriented terms showing intent to tap into existing social issues, while Table 6.6 has a broader political focus with repeated mentions of political figures' names showing an attempt to engage users in political discussions or influence their perceptions of political figures and the media.

We conclude that the analysis of troll tweets reveals an interesting trend: the mention of political figures appears to elicit more responses from users compared to tweets focused on specific social issues.

Table 6.5: Identified Top Words with One Response

Topic	Top Words
Topic 1	islamkills, stopislam, isis, thingspeopleontwitterlike, brussels fergusonremembers, black
Topic 2	blacklivesmatter, mikebrown, gopdebate, fergusonremembers,tcot, police
Topic 3	blacklivesmatter,stopthegop, gopdebate,ccot

Table 6.6: Identified Topics for Troll Tweets with Many Responses

Topic	Top Words
Topic 1	trump, hillary, blacklivesmatter, obama, right, black, muslim
Topic 2	trump, Donald, women, police, Muslim, hillary america, refugees, democrats
Topic 3	trump, clinton, obama, hillary, cnn

6.5 Discussion

The study explores the dynamics of user engagement and response behaviors within the realm of troll tweets on social media platforms. Employing cluster analysis and median response time calculations, the investigation aimed to uncover patterns and trends in user interactions with troll content. The analysis of clusters revealed a diverse landscape of tweet distribution across clusters, each representing distinct user engagement patterns. Cluster 39 emerged as the most prevalent cluster, with tweets with minimal user engagement. Conversely, other clusters exhibited varying degrees of activity, with some clusters garnering higher frequencies of interactions.

We observe the temporal behavior of the median response times to give us insights into the temporal dynamics of user responses to troll tweets. Across different clusters and tweet sequences, significant variations in median response times were observed. While some clusters displayed prompt responses, others exhibited longer response times, indicating varying levels of user engagement and interaction with troll content.

The findings highlight the complexity of user behaviors in response to troll tweets. Despite the prevalence of certain clusters, they did not result in a reduction in the median time of reply. Interestingly, clusters characterized by a low median time of reply were not consistently prevalent across different sequences.

From our analysis, we do not see a pattern of behavior, and it seems that there is no consistent strategy by trolls to influence users to respond quickly to trolls' tweets. Some clusters show quick responses regardless of the frequency of interactions, while others display longer response times despite encountering trolls less frequently. This variance suggests that user responses to trolls might be influenced by factors other than the sheer volume of troll interactions. However, our analysis reveals that mentions of political figures tend to produce more responses from users compared to tweets addressing specific social issues.

6.6 Conclusion

While the study provides valuable insights into user behaviors in response to troll tweets, several limitations should be acknowledged. First, the dataset may suffer from sampling bias and missing context (e.g., deleted accounts, private interactions). Second, historical tweet metadata can contain inconsistencies that affect feature accuracy. Third, because our analysis is confined to one platform and time frame, results may not generalize to other social networks or election cycles.

Future work could embed these anomalies in the underlying follower-reply network to trace their diffusion paths, second, extend the analysis longitudinally to see whether anomalous patterns persist or evolve, and third, pair quantitative signals with qualitative content coding to better interpret intent.

Despite these constraints, our model offers a method to deepen our understanding of how manipulation attempts play out in real-world social media interactions.

In conclusion, the discussion highlights the nuanced nature of user engagement and response behaviors in the context of troll tweets. By leveraging cluster analysis and mean response time calculations, the study offers valuable insights into the dynamics of online interactions.

Chapter 7

Conclusion

This dissertation explores the resilience of interdependent infrastructure systems and the communities they serve, with a particular focus on decision-making under uncertainty, equitable recovery, and information integrity. Each chapter addresses a distinct aspect of this broader goal using a variety of modeling approaches, from stochastic models to agent-based simulations and social data analysis. We list our contributions below.

7.1 Contributions

In Chapter 2, we proposed a method to enhance the resilience of interconnected infrastructure by utilizing a stochastic binary hetero-functional graph to model the dependencies among elements of an infrastructure network. Using a Markov Decision Process (MDP), we simulated failures and repair strategies and identify optimal repair strategies tailored to interdependent systems. Our developed algorithms incorporate equity-based metrics to ensure that we consider equity in the repair processes within the infrastructure network containing power, water, transportation, and community services. Our aim was to ensure a fair recovery process during disruptions by optimizing repair sequences based on interconnectedness and functionalities. Our simulations indicated that our algorithms show the impacts of repair prioritization and help to enhance resilience in interdependent systems. Our contributions

are developing a Markov Decision framework to model the dynamics of failures and repairs in a heterofunctional graph of multiple infrastructures and examine the implications of the various repair decisions ordered by community metrics such as social vulnerability scores and criticality. Our results showed that prioritization of repairs in the power infrastructure is important, and a combination of metrics may give a balanced approach to ordering repairs during outages. For the water infrastructure, we conclude that though ordering of repairs plays a part in equitable restoration, in the long run, the demand for water from communities affected surpasses any effects of this prioritization.

In Chapter 3, we presented HABiT, the Hetero-functional Agent-Based Infrastructure Toolkit, for simulating interdependent power, water, and transportation systems. HABiT models infrastructure as a network of functionalities, such as generating power, transporting water, or consuming vehicles. Each functionality is represented by an autonomous agent with defined behavior, location, and evolving states. The simulation operates over discrete time steps and includes stochastic elements such as random failures, uncertain repair durations, and probabilistic transport availability. We apply HABiT to a simplified model of three communities with different levels of infrastructure access and social vulnerability. The simulation captures both normal operations and disruptions, including recovery influenced by constrained resources, vehicle routing, and mobile repair agents. HABiT uses a heterofunctional graph to represent interdependencies across functionalities. This structure allows components to be added or removed without altering the system configuration, supporting flexible scenario design. Results show that incorporating stochasticity reveals variations in failure cascades and recovery patterns that deterministic models may miss. HABiT enables rapid evaluation of infrastructure disruptions and supports decisions on resource allocation, recovery planning, and resilience analysis under uncertainty. Our results suggest that resilience assessments must go beyond system-wide averages. A fast recovery curve does not answer the deeper question of who recovers first and who waits longest. Decision makers need to clarify whether their priority is minimizing total unmet demand, protecting critical services, or distributing recovery more equitably. Each rule reflects a different objective and leads to different outcomes. A network can be fully restored and still fail its most vulnerable

population if the wrong repair policy is used.

In Chapter 4, we represented interdependent infrastructure systems and communities alike with a hetero-functional graph (HFG) that encodes the dependencies between functionalities. This graph naturally imposes a partial order of functionalities that can inform the sequence of repair decisions to be made during a disaster across affected communities. However, using such technical criteria alone provides limited guidance at the point where the functionalities directly impact the communities, since these can be repaired in any order without violating the system constraints. To address this gap and improve resilience, we integrated community preferences to refine this partial order from the HFG into a total order. Our strategy involves getting the communities' opinions on their preferred sequence for repair crews to address infrastructure issues, considering potential constraints on resources. Due to the delay and cost associated with real-world survey data, we utilize a Large Language Model (LLM) as a proxy survey tool. We used an LLM to craft distinct personas representing individuals, each with varied disaster experiences. We constructed diverse disaster scenarios, and each simulated persona provides input on prioritizing infrastructure repair needs across various communities. Finally, we applied learning algorithms to generate a global order based on the aggregated responses from these LLM-generated personas. Our results show that this method is promising in the attempt to augment a real-life survey with synthetic data.

In Chapter 5, we investigate the communication networks of urban, suburban, and rural communities from three US Midwest counties through a stochastic model that simulates the diffusion of information over time in disaster and in normal situations. To understand information diffusion in communities, we investigate the interplay of information that individuals get from online social networks, local news, government sources, mainstream media, and print media. We utilize survey data collected from target communities and create graphs of each community to quantify node-to-node and source-to-node interactions, as well as trust patterns. Monte Carlo simulation results show the average time it takes for information to propagate to 90% of the population for each community. We conclude that rural, suburban, and urban communities have different inherent properties, promoting the varied flow of information. Also, information sources affect information spread differently, causing

degradation of information speed if any source becomes unavailable. Finally, we provide insights on the optimal investments to improve disaster communication based on community features and contexts. Our results show that there is need for tailored communication strategies because of these differences in communication patterns among different communities. We further emphasize that certain features are more important for each community, and investments in social programs and infrastructure should focus on improving interactions with cable news and local government and trust for online social networks. Also, we find that urban communities rely more on online social networks than rural communities for information dissemination. Moreover, the failure of one communication infrastructure affects each community differently, while the time of diffusion is independent of the size of the community.

In Chapter 6, we examine manipulation tactics used by Russian trolls during the 2016 U.S. General Election by analyzing Twitter data from February 8, 2012, to May 30, 2018. To understand the interaction dynamics between trolls and legitimate users, we link troll tweets with user replies and extract cumulative features that may influence user response behavior. Assuming that users respond to troll tweets based on tweet characteristics and their own interaction histories, we cluster user response histories to identify recurring behavioral patterns. We then perform statistical analysis of response time distributions within and across clusters to assess whether trolls employed coordinated strategies to prompt timely user engagement. Our results indicate a lack of systematic behavioral patterns among trolls; the observed response sequences comprise both fast and slow reaction clusters, suggesting that manipulation tactics were opportunistic rather than strategically timed. Our results show that there is no clear or consistent strategy by trolls to influence users to respond quickly to trolls' tweets using our model.

7.2 Future Work

Despite progress made in modeling and analyzing, several challenges remain unsolved. For Chapter 2, future work could focus on extending the model to larger-scale infrastructure net-

works using real-world data to validate the performance and robustness of your equity-based metrics and repair prioritization. Also, future work can integrate multi-hazard analysis. For Chapter 3, future works can introduce more realistic routing delays, congestion effects, or blocked access could shift how the repair order plays out. Also, more work could explore more complex interactions, such as power systems depending on water, or cascading failures that move between layers and integrate real community data or co-design these metrics with local stakeholders to reflect realistic priorities. Finally, there is a need for external validation using a cosimulation framework. Coupling this agent-based model with established infrastructure simulators would allow for more rigorous testing of assumptions, calibration against real system dynamics, and assessment of whether the observed patterns hold under more physically grounded representations.

For Chapter 4, future research should validate these methods with real-world stakeholder input and extend them to more complex systems. Also, more work can explore better ways to select informative comparisons and apply the method to a broader range of infrastructure types. Additionally, separating infrastructure types and applying domain-specific repair constraints, as well as exploring more complex ranking situations, will be important extensions. However, our current unified ranking approach remains valuable for investment planning contexts, where decision-makers must choose across multiple systems. Another step is to collect survey data and formally test the proof of concept in a real-world setting using actual community data. For Chapter 5, future works can construct optimization models that recommend tailored communication infrastructure investments based on unique community characteristics, enhancing targeted disaster communication effectiveness, and integrate some uncertainty quantification studies. And finally, for Chapter 6, future works can research more on novel approaches to detect trolling in online social networks and how to prevent users from being victims.

Bibliography

- [1] Steven M Rinaldi, James P Peerenboom, and Terrence K Kelly. Identifying, understanding, and analyzing critical infrastructure interdependencies. *IEEE control systems magazine*, 21(6):11–25, 2001.
- [2] Kathleen J Tierney. From the margins to the mainstream? disaster research at the crossroads. *Annu. Rev. Sociol.*, 33(1):503–525, 2007.
- [3] Sergey V Buldyrev, Roni Parshani, Gerald Paul, H Eugene Stanley, and Shlomo Havlin. Catastrophic cascade of failures in interdependent networks. *Nature*, 464(7291):1025–1028, 2010.
- [4] Min Ouyang. Review on modeling and simulation of interdependent critical infrastructure systems. *Reliability engineering & System safety*, 121:43–60, 2014.
- [5] Alessandro Vespignani. The fragility of interdependency. *Nature*, 464(7291):984–985, Apr 2010. ISSN 1476-4687. doi: 10.1038/464984a. URL <https://doi.org/10.1038/464984a>.
- [6] Samuel A Markolf, Mikhail V Chester, Daniel A Eisenberg, David M Iwaniec, Cliff I Davidson, Rae Zimmerman, Thaddeus R Miller, Benjamin L Ruddell, and Heejun Chang. Interdependent infrastructure as linked social, ecological, and technological systems (setss) to address lock-in and enhance resilience. *Earth’s Future*, 6(12):1638–1659, 2018.
- [7] Mert Korkali, Jason G Veneman, Brian F Tivnan, James P Bagrow, and Paul DH Hines. Reducing cascading failure risk by increasing infrastructure network interdependence. *Scientific reports*, 7(1):44499, 2017.

- [8] Neelke Doorn, Paolo Gardoni, and Colleen Murphy. A multidisciplinary definition and evaluation of resilience: The role of social justice in defining resilience. *Sustainable and Resilient Infrastructure*, 4(3):112–123, 2019.
- [9] David P Eisenman, Kristina M Cordasco, Steve Asch, Joya F Golden, and Deborah Glik. Disaster planning and risk communication with vulnerable communities: lessons from hurricane katrina. *American journal of public health*, 97(Supplement_1):S109–S115, 2007.
- [10] Douglas Paton and David Johnston. Disasters and communities: vulnerability, resilience and preparedness. *Disaster Prevention and Management: An International Journal*, 10(4):270–277, 2001.
- [11] Jovy Chan. Online astroturfing: A problem beyond disinformation. *Philosophy & Social Criticism*, 50(3):507–528, 2024.
- [12] James H Fetzer. Information: Does it have to be true? *Minds and Machines*, 14: 223–229, 2004.
- [13] Yimin Chen, Niall J Conroy, and Victoria L Rubin. Misleading online content: recognizing clickbait as” false news”. In *Proceedings of the 2015 ACM on workshop on multimodal deception detection*, pages 15–19, 2015.
- [14] Brian Dolhansky, Russ Howes, Ben Pflaum, Nicole Baram, and Cristian Canton Ferrer. The deepfake detection challenge (dfdc) preview dataset. *arXiv preprint arXiv:1910.08854*, 2019.
- [15] Samantha Bradshaw. Disinformation optimised: Gaming search engine algorithms to amplify junk news. *Internet policy review*, 8(4):1–24, 2019.
- [16] Savvas Zannettou, Tristan Caulfield, Emiliano De Cristofaro, Michael Sirivianos, Gianluca Stringhini, and Jeremy Blackburn. Disinformation warfare: Understanding state-sponsored trolls on twitter and their influence on the web. In *Companion proceedings of the 2019 world wide web conference*, pages 218–226, 2019.

- [17] Samantha Bradshaw and Philip Howard. Troops, trolls and troublemakers: A global inventory of organized social media manipulation. *Computational Propaganda Research Project*, 2017.
- [18] Adaeze Okeukwu-Ogbonnaya, Rahul Amatapu, Jason Bergtold, and George Amariuca. Llm-based community surveys for operational decision making in interconnected utility infrastructures. *arXiv preprint arXiv:2507.13577*, 2025. URL <https://arxiv.org/abs/2507.13577>.
- [19] Adaeze Okeukwu-Ogbonnaya, George Amariuca, Balasubramaniam Natarajan, and Hyung Jin Kim. Towards quantifying the communication aspect of resilience in disaster-prone communities. *Scientific Reports*, 14(1):8837, 2024.
- [20] Wester CH Schoonenberg, Inas S Khayal, and Amro M Farid. *A Hetero-functional graph theory for modeling interdependent smart City infrastructure*. Springer, 2019.
- [21] Zilin Zhang and Yanan Cheng. Critical infrastructure network maintenance scheduling optimization based on random resource requirement. In *2022 4th International Conference on System Reliability and Safety Engineering (SRSE)*, pages 104–108, 2022. doi: 10.1109/SRSE56746.2022.10067841.
- [22] Yi-Ping Fang and Giovanni Sansavini. Optimum post-disruption restoration under uncertainty for enhancing critical infrastructure resilience. *Reliability Engineering & System Safety*, 185:1–11, 2019.
- [23] Basem A Alkhaleel, Haitao Liao, and Kelly M Sullivan. Risk and resilience-based optimal post-disruption restoration for critical infrastructures under uncertainty. *European Journal of Operational Research*, 296(1):174–202, 2022.
- [24] Jiaxin Wu and Pingfeng Wang. Post-disruption performance recovery to enhance resilience of interconnected network systems. *Sustainable and Resilient Infrastructure*, 6(1-2):107–123, 2021.

- [25] Nafiseh Ghorbani-Renani, Andrés D González, Kash Barker, and Nazanin Morshedlou. Protection-interdiction-restoration: Tri-level optimization for enhancing interdependent network resilience. *Reliability Engineering & System Safety*, 199:106907, 2020.
- [26] Yasser Almoghathawi, Kash Barker, and Laura A Albert. Resilience-driven restoration model for interdependent infrastructure networks. *Reliability Engineering & System Safety*, 185:12–23, 2019.
- [27] Kezhi Liu, Changhai Zhai, and You Dong. Optimal restoration schedules of transportation network considering resilience. *Structure and Infrastructure Engineering*, 17(8):1141–1154, 2021.
- [28] Eric D Vugrin, Mark A Turnquist, and Nathanael JK Brown. Optimal recovery sequencing for enhanced resilience and service restoration in transportation networks. *International Journal of Critical Infrastructures*, 10(3-4):218–246, 2014.
- [29] Xiaoge Zhang, Sankaran Mahadevan, Shankar Sankararaman, and Kai Goebel. Resilience-based network design under uncertainty. *Reliability Engineering & System Safety*, 169:364–379, 2018.
- [30] Yi-Ping Fang, Chao Fang, Enrico Zio, and Min Xie. Resilient critical infrastructure planning under disruptions considering recovery scheduling. *IEEE Transactions on Engineering Management*, 68(2):452–466, 2019.
- [31] Aman Karamlou and Paolo Bocchini. Optimal bridge restoration sequence for resilient transportation networks. In *Structures Congress 2014*, pages 1437–1447, 2014.
- [32] Tsai-Yun Liao, Ta-Yin Hu, and Yi-No Ko. A resilience optimization model for transportation networks under disasters. *Natural hazards*, 93:469–489, 2018.
- [33] Aybike Ulsan and Ozlem Ergun. Restoration of services in disrupted infrastructure systems: A network science approach. *PloS one*, 13(2):e0192272, 2018.

- [34] Zhaolong Li, Chun Jin, Pan Hu, and Cong Wang. Resilience-based transportation network recovery strategy during emergency recovery phase under uncertainty. *Reliability Engineering & System Safety*, 188:503–514, 2019.
- [35] Qing Ye and Satish V Ukkusuri. Resilience as an objective in the optimal reconstruction sequence for transportation networks. *Journal of Transportation Safety & Security*, 7(1):91–105, 2015.
- [36] Jason P Giovannettone, Gregg P Macey, Amir AghaKouchak, Michele Barbato, William J Capehart, Auroop R Ganguly, Mital Hall, Jennifer F Helgeson, Si Han Li, Teng Wu, et al. Equitable infrastructure: Achieving resilient systems and restorative justice through policy and research innovation. *PNAS nexus*, 3(5):pgae157, 2024.
- [37] Sara Meerow, Pani Pajouhesh, and Thaddeus R Miller. Social equity in urban resilience planning. *Local Environment*, 24(9):793–808, 2019.
- [38] Deniz Berfin Karakoc, Kash Barker, Christopher W Zobel, and Yasser Almoghathawi. Social vulnerability and equity perspectives on interdependent infrastructure network component importance. *Sustainable Cities and Society*, 57:102072, 2020.
- [39] Thushara Gunda, Amanda Wachtel, Shruti Khadka Mishra, and Emily Moog. Quantitative approaches for including equity in risk and resilience infrastructure planning analyses. *Risk Analysis*, 2023.
- [40] Han Wang and Maria Laura Delle Monache. Urban network resilience analysis and equity emphasized recovery based on reinforcement learning. In *2022 European Control Conference (ECC)*, pages 01–06. IEEE, 2022.
- [41] Jin Rui Yap, Bramka Arga Jafino, and Trivik Verma. Equity principles highlight variations in road network criticality. *Findings*, 2021.
- [42] Milad Baghersad, Christopher W Zobel, Mehdi H Farahani, and Ravi S Behara. Trade-offs between equity and efficiency in prioritizing critical infrastructure investments: A

- case of stormwater management systems. *Production and Operations Management*, page 10591478241243387, 2024.
- [43] Simona Cavallini, Cristina d’Alessandro, Margherita Volpe, Stefano Armenia, Camillo Carlini, Elisabeth Brein, and Pierluigi Assogna. A system dynamics framework for modeling critical infrastructure resilience. In *Critical Infrastructure Protection VIII: 8th IFIP WG 11.10 International Conference, ICCIP 2014, Arlington, VA, USA, March 17-19, 2014, Revised Selected Papers 8*, pages 141–154. Springer, 2014.
- [44] Roberto Filippini and Andrés Silva. A modeling framework for the resilience analysis of networked systems-of-systems based on functional dependencies. *Reliability Engineering & System Safety*, 125:82–91, 2014.
- [45] Thomas Münzberg, Tim Müller, and Wolfgang Raskob. A future-oriented agent-based simulation to improve urban critical infrastructure resilience. *Urban Disaster Resilience and Security: Addressing Risks in Societies*, pages 261–284, 2018.
- [46] Kambiz Rasoulkhani and Ali Mostafavi. Resilience as an emergent property of human-infrastructure dynamics: A multi-agent simulation model for characterizing regime shifts and tipping point behaviors in infrastructure systems. *PloS one*, 13(11):e0207674, 2018.
- [47] Roberto Guidotti, Hana Chmielewski, Vipin Unnikrishnan, Paolo Gardoni, Therese McAllister, and John van de Lindt. Modeling the resilience of critical infrastructure: The role of network dependencies. *Sustainable and resilient infrastructure*, 1(3-4): 153–168, 2016.
- [48] Luca Galbusera, Georgios Giannopoulos, Sotiris Argyroudis, and Kalliopi Kakderi. A boolean networks approach to modeling and resilience analysis of interdependent critical infrastructures. *Computer-Aided Civil and Infrastructure Engineering*, 33(12): 1041–1055, 2018.

- [49] Mark Turnquist and Eric Vugrin. Design for resilience in infrastructure distribution networks. *Environment Systems & Decisions*, 33:104–120, 2013.
- [50] Sotirios A Argyroudis, Stergios A Mitoulis, Lorenzo Hofer, Mariano Angelo Zanini, Enrico Tubaldi, and Dan M Frangopol. Resilience assessment framework for critical infrastructure in a multi-hazard environment: Case study on transport assets. *Science of the Total Environment*, 714:136854, 2020.
- [51] S Amin Enderami and Elaina Sutley. Social vulnerability score: a scalable index for representing social vulnerability in virtual community resilience testbeds. *Natural Hazards*, 120(7):6457–6480, 2024.
- [52] CDC. Social Vulnerability Index, October 2024. URL <https://www.atsdr.cdc.gov/place-health/php/svi/index.html>.
- [53] Alexander Fekete. Common criteria for the assessment of critical infrastructures. *International Journal of Disaster Risk Science*, 2:15–24, 2011.
- [54] Xian Bin Wee, Manuel Herrera, Georgios M Hadjidemetriou, and Ajith Kumar Parlikad. Simulation and criticality assessment of urban rail and interdependent infrastructure networks. *Transportation Research Record*, 2677(1):1181–1196, 2023.
- [55] Bing Li, Kash Barker, and Giovanni Sansavini. Measuring community and multi-industry impacts of cascading failures in power systems. *IEEE Systems Journal*, 12(4):3585–3596, 2017.
- [56] Adaeze Okeukwu-Ogbonnaya. Dynamichfg-modeling. <https://github.com/adaezy/DynamicHFG-Modeling>.git, 2024.
- [57] Thomas H Cormen, Charles E Leiserson, Ronald L Rivest, and Clifford Stein. *Introduction to algorithms*. MIT press, 2022.
- [58] Li Sun, Bozidar Stojadinovic, and Giovanni Sansavini. Agent-based recovery model for

- seismic resilience evaluation of electrified communities. *Risk Analysis*, 39(7):1597–1614, 2019.
- [59] Li Sun, Bozidar Stojadinovic, and Giovanni Sansavini. Resilience evaluation framework for integrated civil infrastructure–community systems under seismic hazard. *Journal of Infrastructure Systems*, 25(2):04019016, 2019.
- [60] Tasnuba Binte Jamal, Samiul Hasan, Omar I Abdul-Aziz, Pallab Mozumder, and Rounak Meyur. Strengthening infrastructure resilience to hurricanes by modeling transportation and electric power network interdependencies. *Natural Hazards Review*, 26(3):04025024, 2025.
- [61] Amir Esmalian, Wanqiu Wang, and Ali Mostafavi. Multi-agent modeling of hazard–household–infrastructure nexus for equitable resilience assessment. *Computer-Aided Civil and Infrastructure Engineering*, 37(12):1491–1520, 2022.
- [62] Jiawei Xue, Sangung Park, Washim Uddin Mondal, Sandro Martinelli Reia, Tong Yao, and Satish V Ukkusuri. An agent-based model of post-disaster recovery in multilayer socio-physical networks. *Sustainable Cities and Society*, 115:105863, 2024.
- [63] Mohammad Aghababaei and Maria Koliou. Community resilience assessment via agent-based modeling approach. *Computer-Aided Civil and Infrastructure Engineering*, 38(7):920–939, 2023.
- [64] Charles M Macal and Michael J North. Tutorial on agent-based modeling and simulation. In *Proceedings of the Winter Simulation Conference, 2005.*, pages 14–pp. IEEE, 2005.
- [65] Charles M Macal. Tutorial on agent-based modeling and simulation: Abm design for the zombie apocalypse. In *2018 Winter Simulation Conference (WSC)*, pages 207–221. IEEE, 2018.
- [66] Li Sun, Bozidar Stojadinovic, and Giovanni Sansavini. Seismic resilience of integrated critical infrastructure network system. In *Proceedings of the 16th World Conference*

- on *Earthquake Engineering (WCEE 2017)*, page 1783. National Information Centre of Earthquake Engineering, Indian Institute of . . . , 2017.
- [67] Mateusz Iwo Dubaniowski and Božidar Stojadinović. Agent-based framework for assessing systemic risk of interdependent sociotechnical and infrastructure systems. In *2022 6th International Conference on System Reliability and Safety (ICSRS)*, pages 196–201. IEEE, 2022.
 - [68] Mateusz Iwo Dubaniowski and Hans Rudolf Heinimann. Framework for modeling interdependencies between households, businesses, and infrastructure system, and their response to disruptions—application. *Reliability Engineering & System Safety*, 212:107590, 2021.
 - [69] Mateusz Iwo Dubaniowski and Hans Rudolf Heinimann. A framework for modeling interdependencies among households, businesses, and infrastructure systems; and their response to disruptions. *Reliability Engineering & System Safety*, 203:107063, 2020.
 - [70] Mateusz Iwo Dubaniowski. *Modeling the behavior of interdependent infrastructure, business unit and household systems under multiple disruptions*. PhD thesis, ETH Zurich, 2019.
 - [71] Jaber Valinejad and Lamine Mili. Cyber–physical–social model of community resilience by considering critical infrastructure interdependencies. *IEEE Internet of Things Journal*, 10(19):17530–17543, 2023.
 - [72] Jaber Valinejad, Lamine Mili, and C Natalie Van Der Wal. Multi-agent based stochastic dynamical model to measure community resilience. *Journal of Social Computing*, 3(3):262–286, 2022.
 - [73] Sai Munikoti, Kexing Lai, and Balasubramaniam Natarajan. Robustness assessment of hetero-functional graph theory based model of interdependent urban utility networks. *Reliability Engineering & System Safety*, 212:107627, 2021.

- [74] Adaeze Okeukwu-Ogbonnaya. Habit: Hetero-functional agent-based infrastructure toolkit. <https://github.com/adaezy/HABiT>, 2025.
- [75] Wei Zhai, Zhong-Ren Peng, and Faxi Yuan. Examine the effects of neighborhood equity on disaster situational awareness: Harness machine learning and geotagged twitter data. *International Journal of Disaster Risk Reduction*, 48:101611, 2020.
- [76] Oronde Drakes, Eric Tate, Jayton Rainey, and Samuel Brody. Social vulnerability and short-term disaster assistance in the united states. *International Journal of Disaster Risk Reduction*, 53:102010, 2021.
- [77] Simone J Domingue and Christopher T Emrich. Social vulnerability and procedural equity: Exploring the distribution of disaster aid across counties in the united states. *The American Review of Public Administration*, 49(8):897–913, 2019.
- [78] Yudi Chen and Wenying Ji. Public demand urgency for equitable infrastructure restoration planning. *International Journal of Disaster Risk Reduction*, 64:102510, 2021.
- [79] Dakota J Thompson and Amro M Farid. A hetero-functional graph structural analysis of the american multi-modal energy system. *Sustainable Energy, Grids and Networks*, 38:101254, 2024.
- [80] Antoine Dolant and Praveen Kumar. Agentic llm framework for adaptive decision discourse. *arXiv preprint arXiv:2502.10978*, 2025.
- [81] Ang Li, Haozhe Chen, Hongseok Namkoong, and Tianyi Peng. Llm generated persona is a promise with a catch. *arXiv preprint arXiv:2503.16527*, 2025.
- [82] Li Shi, Houjiang Liu, Yian Wong, Utkarsh Mujumdar, Dan Zhang, Jacek Gwizdka, and Matthew Lease. Argumentative experience: Reducing confirmation bias on controversial issues through llm-generated multi-persona debates. *arXiv preprint arXiv:2412.04629*, 2024.

- [83] Yangxinyu Xie, Bowen Jiang, Tanwi Mallick, Joshua David Bergerson, John K. Hutchison, Duane R. Verner, Jordan Branham, M. Ross Alexander, Robert B. Ross, Yan Feng, Leslie-Anne Levy, Weijie Su, and Camillo J. Taylor. A rag-based multi-agent llm system for natural hazard resilience and adaptation, 2025. URL <https://arxiv.org/abs/2402.07877>.
- [84] Saad Alqithami. Integrating sentiment analysis and reinforcement learning for equitable disaster response: A novel approach. *Sustainability (2071-1050)*, 17(3), 2025.
- [85] Siyao Chen, Guofeng Luo, and Longbiao Chen. Intelligent urban emergency response: Integrating large language models, multi-objective optimization, and reinforcement learning. In *2024 IEEE Smart World Congress (SWC)*, pages 1469–1476. IEEE, 2024.
- [86] Minze Chen, Zhenxiang Tao, Weitong Tang, Tingxin Qin, Rui Yang, and Chunli Zhu. Enhancing emergency decision-making with knowledge graphs and large language models. *International Journal of Disaster Risk Reduction*, 113:104804, 2024.
- [87] Nasim Mohamadiazar, Ali Ebrahimian, and Atorod Azizinamini. Equitable prioritization of bridge rehabilitation projects using a spatial multi-criteria decision support framework. *Transportation Research Record*, page 03611981241254110, 2024.
- [88] Makarand Hastak, Qingbin Cui, Bhavin Safi, and Sanjiv Gokhale. A decision support system for infrastructure rehabilitation planning. In *Computing in Civil Engineering (2005)*, pages 1–10. 2005.
- [89] Rubel Das and Masaaki Nakano. A multi-criteria decision-making model using socio-technical attributes for transportation bridge maintenance prioritisation. *International Journal of Construction Management*, 23(4):579–585, 2023.
- [90] Marianna Garfi and Laia Ferrer-Martí. Decision-making criteria and indicators for water and sanitation projects in developing countries. *Water Science and Technology*, 64(1):83–101, 2011.

- [91] Federico Dell’Anna, Mauro Berta, Marta Bottero, Gloria Mallia, and Vito Morgese. Multicriteria-decision support for master plan scheduling: urban regeneration of an industrial area in northern italy. *Construction Management and Economics*, 42(5): 476–501, 2024.
- [92] Darwin Marcelo, Cledan Mandri-Perrott, Schuyler House, and Jordan Schwartz. Prioritizing infrastructure investment: a framework for government decision making. *World Bank Policy Research Working Paper*, (7674), 2016.
- [93] Rahul Dhansinghani, Ayman Ibrahim, Aditya Kannothe, Claire Miller, Lena Nguyen, Steven Pham, and Reid Bailey. Design of a prioritization methodology for equitable infrastructure planning. In *2022 Systems and Information Engineering Design Symposium (SIEDS)*, pages 315–320. IEEE, 2022.
- [94] Sura Karasneh and Shadi Moqbel. Priority-based decision model for rehabilitation of water networks using fahp. *Water Conservation & Management*, 8(1), 2024.
- [95] Hana Sherin Galal. *Integrating sustainability in municipal wastewater infrastructure decision-analysis using the analytic hierarchy process*. PhD thesis, University of British Columbia, 2013.
- [96] Nectaria Putri Pramesti, Henricus Priyosulistyo, Akhmad Aminullah, and Ambrosius Koesmargono. The bridge maintenance factors model: a pls-sem approach. *Age*, 22: 18, 2021.
- [97] Rolf Baur, Pascal Le Gauffre, and Sveinung Sægrov. Multi-criteria decision support for annual rehabilitation programmes in drinking water networks. *Water Science and Technology: Water Supply*, 3(1-2):43–50, 2003.
- [98] Dohyung Park, Joe Neeman, Jin Zhang, Sujay Sanghavi, and Inderjit Dhillon. Preference completion: Large-scale collaborative ranking from pairwise comparisons. In *International Conference on Machine Learning*, pages 1907–1916. PMLR, 2015.

- [99] Suriya Gunasekar, Oluwasanmi O Koyejo, and Joydeep Ghosh. Preference completion from partial rankings. *Advances in Neural Information Processing Systems*, 29, 2016.
- [100] Martin Szummer and Emine Yilmaz. Semi-supervised learning to rank with preference regularization. In *Proceedings of the 20th ACM international conference on Information and knowledge management*, pages 269–278, 2011.
- [101] Yixuan He, Quan Gan, David Wipf, Gesine D Reinert, Junchi Yan, and Mihai Cucuringu. Gnnrank: Learning global rankings from pairwise comparisons via directed graph neural networks. In *international conference on machine learning*, pages 8581–8612. PMLR, 2022.
- [102] Seon-Ho Lee and Chang-Su Kim. Order learning using partially ordered data via chainization. In *European Conference on Computer Vision*, pages 196–211. Springer, 2022.
- [103] Adaeze Okeukwu-Ogbonnaya. CommRank-HFG. <https://github.com/adaezy/CommRank-HFG>, 2025.
- [104] Or Zuk, Liat Ein-Dor, and Eytan Domany. Ranking under uncertainty. *arXiv preprint arXiv:1206.5280*, 2012.
- [105] Pranay Mundra, Jianhao Zhang, Fatemeh Nargesian, and Nikolaus Augsten. Koios: Top-k semantic overlap set search. In *2023 IEEE 39th International Conference on Data Engineering (ICDE)*, pages 1531–1543. IEEE, 2023.
- [106] Moran Mizrahi, Guy Kaplan, Dan Malkin, Rotem Dror, Dafna Shahaf, and Gabriel Stanovsky. State of what art? a call for multi-prompt llm evaluation. *Transactions of the Association for Computational Linguistics*, 12:933–949, 2024.
- [107] Anwoy Chatterjee, HSVNS Kowndinya Renduchintala, Sumit Bhatia, and Tanmoy Chakraborty. Posix: A prompt sensitivity index for large language models. *arXiv preprint arXiv:2410.02185*, 2024.

- [108] Vishal Gandhi and Sagar Gandhi. Prompt sentiment: The catalyst for llm change. *arXiv preprint arXiv:2503.13510*, 2025.
- [109] Jia He, Mukund Rungta, David Koleczek, Arshdeep Sekhon, Franklin X Wang, and Sadid Hasan. Does prompt formatting have any impact on llm performance? *arXiv preprint arXiv:2411.10541*, 2024.
- [110] Rajib Shaw. *Community based disaster risk reduction*. Emerald Group Publishing, 2012.
- [111] A. Fema. *Whole community approach to emergency management: Principles, themes, and pathways for action*. US Department of Homeland Security, Washington DC, 2011.
- [112] A. Plough, J. E. Fielding, A. Chandra, M. Williams, D. Eisenman, K. B. Wells, G. Y. Law, S. Fogleman, and A. Magaña. Building community disaster resilience: perspectives from a large urban county department of public health. *American journal of public health*, 103(7):1190–1197, 2013.
- [113] S. Romac. The importance of community resilience: Developing the american red cross international services department in the new hampshire region. 2014.
- [114] A. Chandra, J. Acosta, S. Howard, L. Uscher-Pines, M. Williams, D. Yeung, J. Garnett, and L. S. Meredith. Building community resilience to disasters: A way forward to enhance national health security. *Rand health quarterly*, 1(1), 2011.
- [115] P. K. Jayasekara. Role of facebook as a disaster communication media. *International Journal of Emergency Services*, 2019.
- [116] J. Kim and M. Hastak. Social network analysis: Characteristics of online social networks after a disaster. *International journal of information management*, 38(1):86–96, 2018.
- [117] M. Martínez-Rojas, M. del Carmen Pardo-Ferreira, and J. C. Rubio-Romero. Twitter

- as a tool for the management and analysis of emergency situations: A systematic literature review. *International Journal of Information Management*, 43:196–208, 2018.
- [118] E. A. Vogels. Some digital divides persist between rural, urban and suburban america. 2021.
- [119] E. A. Dobis, T. P. Krumel, J. Cromartie, K. L. Conley, A. Sanders, and R. Ortiz. *Rural america at a glance: 2021 edition*. Rep, Tech, 2021.
- [120] W. Li, Y. Li, Y. Tan, Y. Cao, C. Chen, Y. Cai, K. Y. Lee, and M. Pecht. Maximizing network resilience against malicious attacks. *Scientific reports*, 9(1):2261, 2019.
- [121] E. Yoo, W. Rand, M. Eftekhari, and E. Rabinovich. Evaluating information diffusion speed and its determinants in social media networks during humanitarian crises. *Journal of Operations Management*, 45:123–133, 2016.
- [122] C. Fan, J. Shen, A. Mostafavi, and X. Hu. Characterizing reticulation in online social networks during disasters. *Applied Network Science*, 5:1–20, 2020.
- [123] J. Goldenberg, B. Libai, and E. Muller. Talk of the network: A complex systems look at the underlying process of word-of-mouth. *Marketing letters*, 12(3):211–223, 2001.
- [124] D. Kempe, J. Kleinberg, and É. Tardos. Maximizing the spread of influence through a social network. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 137–146, 2003.
- [125] C. Fan, Y. Jiang, Y. Yang, C. Zhang, and A. Mostafavi. Crowd or hubs: information diffusion patterns in online social networks in disasters. *International journal of disaster risk reduction*, 46:101498, 2020.
- [126] J. Kim, J. Bae, and M. Hastak. Emergency information diffusion on online social media during storm cindy in us. *International Journal of Information Management*, 40:153–165, 2018.

- [127] Y. Yang, C. Zhang, C. Fan, W. Yao, R. Huang, and A. Mostafavi. Exploring the emergence of influential users on social media during natural disasters. *International Journal of Disaster Risk Reduction*, 38:101204, 2019.
- [128] R. Dong, L. Li, Q. Zhang, and G. Cai. Information diffusion on social media during natural disasters. *IEEE transactions on computational social systems*, 5(1):265–276, 2018.
- [129] D. Sharma, R. Kumar, et al. Emergency management using social networks. In *TENCON 2019-2019 IEEE Region 10 Conference (TENCON)*, pages 721–726. IEEE, 2019.
- [130] C. Zhang, C. Fan, W. Yao, X. Hu, and A. Mostafavi. Social media for intelligent public information and warning in disasters: An interdisciplinary review. *International Journal of Information Management*, 49:190–207, 2019.
- [131] W. Zhai. A multi-level analytic framework for disaster situational awareness using twitter data. *Computational Urban Science*, 2(1):23, 2022.
- [132] J. Xu and Y. Qiang. Analysing information diffusion in natural hazards using retweets-a case study of 2018 winter storm diego. *Annals of GIS*, 28(2):213–227, 2022.
- [133] S. Nagar, A. Seth, and A. Joshi. Characterization of social media response to natural disasters. In *Proceedings of the 21st international conference on world wide web*, pages 671–674, 2012.
- [134] J. Zhu, F. Xiong, D. Piao, Y. Liu, and Y. Zhang. Statistically modeling the effectiveness of disaster information in social media. In *2011 IEEE Global Humanitarian Technology Conference*, pages 431–436. IEEE, 2011.
- [135] N. Altay and R. Pal. Information diffusion among agents: Implications for humanitarian operations. *Production and Operations Management*, 23(6):1015–1027, 2014.

- [136] X. Liu, D. He, and C. Liu. Information diffusion nonlinear dynamics modeling and evolution analysis in online social network based on emergency events. *IEEE Transactions on Computational Social Systems*, 6(1):8–19, 2019.
- [137] S. Adali, R. Escriva, M. K. Goldberg, M. Hayvanovych, M. Magdon-Ismael, B. K. Szymanski, W. A. Wallace, and G. Williams. Measuring behavioral trust in social networks. In *2010 IEEE international conference on intelligence and security informatics*, pages 150–152. IEEE, 2010.
- [138] H. Wu, A. Arenas, and S. Gómez. Influence of trust in the spreading of information. *Physical Review E*, 95(1):012301, 2017.
- [139] I. Fridman, N. Lucas, D. Henke, C. K. Zigler, et al. Association between public knowledge about covid-19, trust in information sources, and adherence to social distancing: cross-sectional survey. *JMIR public health and surveillance*, 6(3):e22060, 2020.
- [140] M. J. Widener, M. W. Horner, and S. S. Metcalf. Simulating the effects of social networks on a population’s hurricane evacuation participation. *Journal of geographical systems*, 15:193–209, 2013.
- [141] T. Haer, W. W. Botzen, and J. C. Aerts. The effectiveness of flood risk communication strategies and the influence of social networks—insights from an agent-based model. *Environmental Science & Policy*, 60:44–52, 2016.
- [142] M. Posfai and A.-L. Barabasi. Network science. , 1em plus 0.5em minus 0.4emCiteseer, 2016.
- [143] J. Moss and M. Tveten. kdensity: An r package for kernel density estimation with parametric starts and asymmetric kernels. *Journal of Open Source Software*, 4(42):1566, 2019.
- [144] B. K. Fosdick, D. B. Larremore, J. Nishimura, and J. Ugander. Configuring random graph models with fixed degree sequences. *Siam Review*, 60(2):315–355, 2018.

- [145] U. Dutta, B. K. Fosdick, and A. Clauset. Sampling random graphs with specified degree sequences. *arXiv e-prints*, pages Xiv–2105, 2021. pp. ar.
- [146] M. E. Newman. The structure and function of complex networks. *SIAM review*, 45(2):167–256, 2003.
- [147] A. Hagberg, P. Swart, and D. S. Chult. *Exploring network structure, dynamics, and function using networkx*. Los Alamos National Lab.(LANL, Los Alamos, NM (United States)), Tech. Rep, 2008.
- [148] V. Arnaboldi, M. Conti, A. Passarella, and R. I. Dunbar. Online social networks and information diffusion: The role of ego networks. *Online Social Networks and Media*, 1: 44–55, 2017.
- [149] M. Grandjean. Gephi: Introduction to network analysis and visualisation. 2015.
- [150] Stanley Wasserman. Social network analysis: Methods and applications. *The Press Syndicate of the University of Cambridge*, 1994.
- [151] M. Al Hasan and V. Dave. Triangle counting in large networks: a review. *Wiley Interdisciplinary Reviews: Data Mining And Knowledge Discovery*, 8:e1226, 2018.
- [152] M. Newman. *Networks: An Introduction*. Oxford University Press, Oxford, 2010.
- [153] Adaeze Okeukwu. Smart and Connected Communities. https://github.com/adaezy/Smart_and_Connected_Communities, 2023. GitHub repository.
- [154] Tal Z Zarsky. Privacy and manipulation in the digital age. *Theoretical Inquiries in Law*, 20(1):157–188, 2019.
- [155] Robert Noggle. The ethics of manipulation. 2018.
- [156] Daniel Susser, Beate Roessler, and Helen Nissenbaum. Technology, autonomy, and manipulation. *Internet Policy Review*, 8(2), 2019.

- [157] Manuel Förster, Ana Mauleon, and Vincent J Vannetelbosch. Trust and manipulation in social networks. *Network Science*, 4(2):216–243, 2016.
- [158] S Rasoul Etesami, Sadegh Bolouki, Angelia Nedić, Tamer Başar, and H Vincent Poor. Influence of conformist and manipulative behaviors on public opinion. *IEEE Transactions on Control of Network Systems*, 6(1):202–214, 2018.
- [159] Ben Nimmo. Measuring traffic manipulation on twitter. *University of Oxford Internet Institute*, 2019.
- [160] Matthew C Benigni, Kenneth Joseph, and Kathleen M Carley. Bot-ivism: assessing information manipulation in social media using network analytics. *Emerging research challenges and opportunities in computational social network analysis and mining*, pages 19–42, 2019.
- [161] Masarah Paquet-Clouston, Olivier Bilodeau, and David Décary-Héту. Can we trust social media data? social network manipulation by an iot botnet. In *Proceedings of the 8th international conference on social media & society*, pages 1–9, 2017.
- [162] Muhammad Nihal Hussain, Serpil Tokdemir, Nitin Agarwal, and Samer Al-Khateeb. Analyzing disinformation and crowd manipulation tactics on youtube. In *2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, pages 1092–1095. IEEE, 2018.
- [163] Dennis Assenmacher, Lena Clever, Janina Susanne Pohl, Heike Trautmann, and Christian Grimme. A two-phase framework for detecting manipulation campaigns in social media. In *Social Computing and Social Media. Design, Ethics, User Behavior, and Social Network Analysis: 12th International Conference, SCSM 2020, Held as Part of the 22nd HCI International Conference, HCII 2020, Copenhagen, Denmark, July 19–24, 2020, Proceedings, Part I 22*, pages 201–214. Springer, 2020.
- [164] Dongwoo Kim, Timothy Graham, Zimin Wan, and Marian-Andrei Rizoіu. Analysing

- user identity via time-sensitive semantic edit distance (t-sed): a case study of russian trolls on twitter. *Journal of Computational Social Science*, 2:331–351, 2019.
- [165] Sergei Monakhov. Early detection of internet trolls: Introducing an algorithm based on word pairs/single words multiple repetition ratio. *PloS one*, 15(8):e0236832, 2020.
 - [166] Soon Ae Chun, Richard D Holowczak, Kannan Dharan, Ruoyu Wang, Soumaydeep Basu, and James Geller. Detecting political bias trolls in twitter data. In *WEBIST*, pages 334–342, 2019.
 - [167] Darren L Linvill, Brandon C Boatwright, Will J Grant, and Patrick L Warren. “the russians are hacking my brain!” investigating russia’s internet research agency twitter tactics during the 2016 united states presidential campaign. *Computers in Human Behavior*, 99:292–300, 2019.
 - [168] Adam Badawy, Emilio Ferrara, and Kristina Lerman. Analyzing the digital traces of political manipulation: The 2016 russian interference twitter campaign. In *2018 IEEE/ACM international conference on advances in social networks analysis and mining (ASONAM)*, pages 258–265. IEEE, 2018.
 - [169] Emilio Ferrara, Herbert Chang, Emily Chen, Goran Muric, and Jaimin Patel. Characterizing social media manipulation in the 2020 us presidential election. *First Monday*, 2020.
 - [170] Ashok Deb, Luca Luceri, Adam Badaway, and Emilio Ferrara. Perils and challenges of social media and election manipulation analysis: The 2018 us midterms. In *Companion proceedings of the 2019 world wide web conference*, pages 237–247, 2019.
 - [171] Felipe Cardoso, Luca Luceri, and Silvia Giordano. Digital weapons in social media manipulation campaigns. In *Workshop Proceedings of the 14th International AAAI Conference on Web and Social Media*, volume 3, 2020.
 - [172] FiveThirtyEight. 3 million russian troll tweets. <https://github.com/fivethirtyeight/russian-troll-tweets>, 2018.

- [173] Niko Colnerič and Janez Demšar. Emotion recognition on twitter: Comparative study and training a unison model. *IEEE Transactions on Affective Computing*, 11(3):433–446, 2020. doi: 10.1109/TAFFC.2018.2807817.
- [174] Adaeze Okeukwu-Ogbonnaya. Manipulation-with-Russian-Trolls: Predicting Topic Engagement during Russian Troll Activity around the 2016 US Political Elections. https://github.com/adaezy/manipulation_russian_trolls.

Appendix A

Chapter 4 Appendix

Here, we provide the original prompts and the different prompt variants. Also, we provide a sample persona and response to the original prompt. Full data, prompt templates, and code are available at [103](#).

A.0.1 Prompt Templates

We show the prompt templates we use to get the repair preferences from LLM-generated personas in our study. The variables in curly brackets({}) are replaced dynamically at runtime with persona attributes and scenario-specific information.

Figure [A.1](#) shows the original prompt we use. This prompt format differs from the “Choice Before Reasoning” prompt, where the persona makes a choice before reasoning, as seen in this example response [A.7](#).

Figure A.1: Original Prompt

You are a {full.persona}, and you are aware that a tornado has impacted three communities (1, 2, and 3), causing damage to critical infrastructure.

You are familiar with the conditions in all communities, described as follows: {all.community.descriptions}. Social Vulnerability Scores (SVS) range from 0 (least vulnerable) to 1 (most vulnerable), with higher values indicating greater vulnerability.

Consider the situation, taking into account any details you find relevant.

Question: Which should be repaired first: '{option1}' or '{option2}'?

Before choosing between '{option1}' and '{option2}', first explain your reasoning.

Provide your answer in the following format:

Reasoning: [Your one-sentence explanation]

Priority Choice: [Your chosen option]

Figure [A.2](#) shows the prompt for changing the instruction format.

Figure A.2: Changing Instruction Format

Step 1: You are a {full_persona}, and you are aware that a tornado has impacted three communities (1, 2, and 3), causing damage to critical infrastructure.

Step 2: You are familiar with the conditions in all communities, described as follows: {all_community_descriptions}. Social Vulnerability Scores (SVS) range from 0 (least vulnerable) to 1 (most vulnerable), with higher values indicating greater vulnerability.

Step 3: Consider the situation, taking into account any details you find relevant.

Step 4: Question: Which should be repaired first: '{option1}' or '{option2}'?

Step 5: Before choosing between '{option1}' and '{option2}', first explain your reasoning.

Step 6: Provide your answer in the following format:

Reasoning: [Your one-sentence explanation]

Priority Choice: [Your chosen option]

Figure A.3 shows the prompt for shifting tone.

Figure A.3: Shifting Tone Prompt

You are a {full.persona}, and you are aware that a tornado has impacted three communities (1, 2, and 3), causing damage to critical infrastructure.

You are familiar with the conditions in all communities, described as follows: {all.community.descriptions}. Social Vulnerability Scores (SVS) range from 0 (least vulnerable) to 1 (most vulnerable), with higher values indicating greater vulnerability.

Keep in mind the emotional and practical consequences of the outage for residents living in these communities as you evaluate what should be repaired first.

Consider the situation, taking into account any details you find relevant.

Question: Which should be repaired first: '{option1}' or '{option2}'?

Before choosing between '{option1}' and '{option2}', first explain your reasoning.

Provide your answer in the following format:

Reasoning: [Your one-sentence explanation]

Priority Choice: [Your chosen option]

Figure A.4 shows the prompt for rewording the question.

Figure A.4: Rewording Question

You are a {full.persona}, and you are aware that a tornado has impacted three communities (1, 2, and 3), causing damage to critical infrastructure.

You are familiar with the conditions in all communities, described as follows: {all.community.descriptions}. Social Vulnerability Scores (SVS) range from 0 (least vulnerable) to 1 (most vulnerable), with higher values indicating greater vulnerability.

Consider the situation, taking into account any details you find relevant.

Question: Between the two options, '{option1}' and '{option2}', which do you deem most critical for immediate repair?

Before choosing between '{option1}' and '{option2}', first explain your reasoning.

Provide your answer in the following format:

Reasoning: [Your one-sentence explanation]

Priority Choice: [Your chosen option]

Figure [A.5](#) shows the prompt for rewording the reasoning.

Figure A.5: Rewording Reasoning

```
You are a {full.persona}, and you are aware that a tornado has impacted three communities (1, 2, and 3),
causing damage to critical infrastructure.

You are familiar with the conditions in all communities, described as follows: {all.community.descriptions}.
Social Vulnerability Scores (SVS) range from 0 (least vulnerable) to 1 (most vulnerable), with higher
values indicating greater vulnerability.

Consider the situation, taking into account any details you find relevant.

Question: Which should be repaired first: '{option1}' or '{option2}'?

Before choosing between '{option1}' and '{option2}', first explain your reasoning.

Provide your answer in the following format:
Reasoning: [Explain your selection with a concise rationale, highlighting relevant factors influencing
your decision.]
Priority Choice: [Your chosen option]
```

A.0.2 Sample Persona and Response

Figure A.6 is an example persona created by the LLM and used in the simulation. Figure A.7 is an example response to the pairwise prioritization question between choosing Power in Residential in Community 1 and Power in Residential in Community 3.

```
"persona_id": "P002",
"age": 30,
"gender": "Male",
"race_ethnicity": "Hispanic",
"marital_status": "Married",
"citizenship_status": "Naturalized Immigrant",
"birth_place": "Mexico",
"year_of_immigration": null,
"primary_language": "Spanish",
"english_proficiency": "Fluent",
"education_level": "Master's Degree",
"school_enrolment": "No",
```

```

"employment_status": "Employed",
"occupation": "Software Engineer",
"industry_of_employment": "Technology",
"class_of_worker": "Private Sector",
"work_hours_per_week": 40,
"income_bracket": "$50K-$75K",
"poverty_status": "Above Poverty Line",
"government_assistance": "None",
"household_relationship": "Married",
"household_size": 2,
"number_of_children": 2,
"responsible_for_grandchildren": false,
"home_ownership": "Own",
"housing_type": "Single-Family Home",
"utilities_access": "Full access to utilities",
"utilities_housing_costs": "$1500\_per\_month",
"transportation_mode": "Drives",
"commute_duration": 30,
"work_from_home": false,
"health_insurance": "Employer-provided",
"disability_status": "No",
"disability_type": null,
"military_service": "None",
"military_service_period": null,
"disability_rating": null,
"internet_access": "Broadband",
"device_access": "Laptop",
"recent_birth": true,
"moved_in_last_year": false,
"previous_residence": "Same city",
"disaster_experience": "Experienced an earthquake in 2018",
"community_description": "In Community 1: ... (text truncated for brevity)",
"geographical_class": "Suburban"

```

Figure A.6: Full specification of one synthetic persona used in the simulation

Figure A.7: Sample Response

Question:

Which should be repaired first: Power Residential in Community 1
or Power Residential in Community 3?

Response:

Priority Choice: Repair Power Residential in Community 3

Reasoning: Community 3 has a higher social vulnerability score
of 7.0/10 compared to Community 1's 2.0/10, indicating a greater
need for immediate assistance to ensure safety in the more
vulnerable urban area.

Appendix B

Chapter 5 Appendix

B.0.1 Stochastic Model

In this section, Algorithm 3 presents the pseudocode for the stochastic model. We are interested in estimating the average diffusion time, defined as the time that elapses between the start of an information event, and the time when 90% of individuals in the community have accepted the information (or have become *informed*). We use a stochastic model to simulate information passing within the built graphs. We start with an initial graph with $n + 5$ nodes. Initially, we assume there is a piece of information to be shared, and only the five hubs have this information. Let X be the number of interactions a node has with a specific neighbor – another person or a source of information – in a week, where $X \sim \text{Poisson}(\lambda)$, and λ is an edge value between the two nodes. Using this rate, we calculate the probability of 0 meetings between a node and its neighbors. Zero meetings between node i and a neighbor j mean that node i and j have not met yet, so node i does not learn any information from node j . We begin by iterating through all the nodes in the graph. For each time step, we check the state of the node. If a node has the information, we do not take any action. However, if a node does not have the information, we check its informed neighbors. We iterate through the informed neighbors of each node i . To determine the probability that node i meets node j , P_0 , let the rate of interactions between i and j be λ_{ij} . We calculate P_0

as 1 minus the probability of zero meetings between node i and node j , that is,

$$P_0 = 1 - \text{Poisson}(X = 0, \lambda_{ij}) = 1 - e^{-\lambda_{ij}}. \quad (\text{B.1})$$

We draw once from a Bernoulli distribution with parameter P_0 . If the outcome is 1, we update the uninformed node's total trust in the new information. To update this total trust, we apply a discount factor to the trust value between node i and j . The discount factor accounts for the number of times node j passed the same information to node i before. Previous interactions with node j , as well as with other nodes, are discounted by a forgetting factor that accounts for the length of time that passed since those interactions. We calculate the trust as follows:

$$T_c = \sum_{j=1}^N \sum_{k=0}^{M_{ij}} \theta_{ij} d^k f^{t_c - t_k} \quad (\text{B.2})$$

T_c is the total trust in the new information, N is the number of neighbors of i with a state of 1, M_{ij} is the number of meetings between i and j , θ_{ij} is the trust between node i and j , d is the discount factor and f is the forgetting factor, t_c is current time step and t_k is the time of meeting number k between i and j . Node i becomes *informed* when T_c exceeds a threshold Θ .

Graph Generation Algorithms

For graph generation, we use two methods to generate synthetic graphs that closely resemble the original. In the first graph generation method, Algorithm 4, we first use kernel density estimation to learn the degree distribution of the degree sequence of initial survey data and sample from it to generate a degree sequence of size $d = 1000$. We shift the mean and standard deviation of the degree sequence of the initial survey data by small amounts, randomly sampled from intervals centered around zero as shown in Algorithm 4. In the second method, as shown in Algorithm 5, we change the heights of the bins of the histograms of the degree distribution. We again use kernel density estimation to learn the degree distribution of the

degree sequence of the initial survey data. By sampling from this distribution, we generate a degree sequence of size $d = 1000$. We then create a histogram of the distribution and add or subtract 0.1 from the weighted frequency of each bin to get a new frequency for each bin. We normalize the new weighted frequencies based on the sum of the original frequency of each bin in the histogram and generate points from a uniform distribution between the original bin widths. The number of points generated from each bin equaled the calculated normalized frequencies. Finally, we perform some cleanup to ensure the length of the generated data is equal to the size needed and to round up to zero if any data point is less than zero.

Algorithm 3 The Monte Carlo Simulation Algorithm

Require: Graph G , exposedNodes, time

```

1: initialize populationPercent = 90% of nodes in  $G$ 
2: initialize time = 0
3: while len(exposedNodes) > 0 do
4:    $d \leftarrow 0.5$  ▷ discount value
5:    $f \leftarrow 1.3$  ▷ forgetting value
6:    $\Theta \leftarrow 30$  ▷ trust threshold
7:   for  $i = 1$  to  $N$  nodes in  $G$  do
8:     if state of  $i = 0$  then
9:       Get all the neighbors of node  $i$ , call them neighborNodes
10:      newTrust = 0
11:      for  $j$  in neighborNodes do
12:        if state of  $j = 1$  then
13:          Get zeroProbability
14:          probabilityInfection =  $1 - \text{zeroProbability}$ 
15:          if probabilityInfection > Uniform[0,1] then
16:            if key  $j$  exists in the meetings-dictionary between  $i$  and  $j$  then
17:              Update  $M \leftarrow M + 1$  ▷  $M$  = number of meetings of  $i$  and  $j$ 
18:              Update array timeStep with current time
19:              Get the  $\theta$  between node  $j$  and  $i$  ▷ (trust-value)
20:              trustInStep = 0
21:              for  $m = 0$  to  $M$  do
22:                Get  $t$ 
23:                trustInStep +=  $d^m \times f^{(\text{time}-t)}$ 
24:              end for
25:              newTrust += (trustInStep  $\times \theta$ )
26:            else
27:               $M, \text{time}, \text{trustInStep} = 0, 0, 0$ 
28:              Get the  $\theta$  between node  $j$  and  $i$ 
29:              Update  $M \leftarrow M + 1$ 
30:              Get  $M$  ▷ (last-meeting index of  $i$  and  $j$ )
31:               $t = 0$ 
32:              for  $m = 0$  in  $M$  do
33:                trustInStep += ( $d^m \times f^{(\text{time}-t)}$ )
34:              end for
35:              newTrust += (trustInStep  $\times \theta$ )
36:            end if
37:          end if
38:        end if
39:      end for
40:    end if
41:    if newTrust >  $\Theta$  then
42:      Update State of  $i = 0$ 
43:    else
44:      State of  $i = 0$ 
45:    end if
46:    if State of  $i == 1$  then
47:      Add node  $i$  to exposedNodes
48:    end if
49:  end for
50:  time += 1
51: end while
52: timeDiffusion.append(time) return timeDiffusion

```

Algorithm 4 Graph Generation 1

Require: Original data d , Length L

```
1: newData  $\leftarrow$  KDE( $L, d$ ) ▷ Kernel Density Estimation (KDE) requires the graph length to generate
2: oldMean  $\leftarrow$  Mean of  $d$ 
3: oldStd  $\leftarrow$  Standard deviation of  $d$ 
4: meanChange  $\leftarrow$  oldMean  $\times$  0.1
5: stdChange  $\leftarrow$  oldStd  $\times$  0.05
6: meanChange  $\leftarrow$  Randomly select one of  $\{-\text{meanChange}, \text{meanChange}\}$ 
7: stdChange  $\leftarrow$  Randomly select one of  $\{\text{stdChange}, (\text{stdChange})^{-1}\}$ 
8: for all  $i$  in newData do
9:    $i \leftarrow (i + \text{meanChange}) \times \text{stdChange}$ 
10:   Convert  $i$  to integer; if the result is less than 0, then set  $i \leftarrow 0$ 
11:   newData.append( $i$ )
12: end for
13: return newData
```

B.0.2 Generating Trust and Interaction Values

We discuss the algorithm for generating trust and interaction values as shown in pseudocode Algorithm 6. Since, we need to get synthetic trust and frequency of interaction data to label the edges of the synthetic graphs. We want the artificial trust and interaction data to be close to the original features of the ten communities. We first get the frequencies of each of the ordinal variables. Next, we change the sizes of the relative frequencies by adding or subtracting 0.1 from each of them. Then we renormalize the new relative frequencies and generate the new trust and interaction values according to these values. Then, the remaining code is for cleanup to ensure the length of data generated is equal to the expected length.

Algorithm 5 Graph Generation 2

Require: Original data d , Length L

```
1: newData  $\leftarrow$  KDE( $L, d$ ) ▷ Kernel Density Estimation (KDE) requires the graph length to generate
2: ( $n, \text{bins}$ )  $\leftarrow$  Histogram(newData) ▷  $n$  is the count of samples in each bin; bins are the bin edges
3: binWidth  $\leftarrow$  bins[2] - bins[1]
4: Get array weightfreq = weighted frequency of each count in  $n$ 
5: oldSum  $\leftarrow \sum n$ 
6: change  $\leftarrow$  0.1
7: Initialize empty array newFreqList
8: for all  $i, j$  in bins and weightfreq do
9:   if  $j \neq 0$  then
10:    newFreqBin  $\leftarrow j + \text{UniformRV}[-\text{change} \times i, \text{change} \times i]$ 
11:   else
12:    newFreqBin  $\leftarrow j + \text{UniformRV}[-\text{change}, \text{change}]$ 
13:   end if
14:   newFreqList.append(newFreqBin)
15: end for
16: Initialize empty array newData
17: for all  $a, b$  in newFreqList and bins do
18:   numVal  $\leftarrow (a / \sum(\text{newFreqList})) \times 100$ 
19:   newFreq  $\leftarrow (\text{numVal} / 100) \times \text{oldSum}$ 
20:   gen  $\leftarrow \text{UniformRV}[b, b + \text{binWidth}]$  of size newFreq
21:   newData.append(gen)
22: end for
23: if  $|\text{newData}| < L$  then ▷ Cleanup code to ensure  $|\text{newData}| = L$ 
24:   deficit  $\leftarrow L - |\text{newData}|$ 
25:   minVal, maxVal  $\leftarrow \min(\text{newData}), \max(\text{newData})$ 
26:   r  $\leftarrow \text{DiscreteUniformRV}[\text{minVal}, \text{maxVal}]$  of length deficit
27:   newData.append(r)
28: else
29:   k  $\leftarrow ||\text{newData}| - L|$ 
30:   for  $\_ = 1$  to k do
31:     newData.pop()
32:   end for
33: end if
34: Set any element in newData to 0 if it is less than 0
35: return newData
```

Algorithm 6 Generation of Trust and Interaction Values

Require: Original data d

```
1: Let array  $\mathbf{n}$  contain the frequency of all data points in  $d$ 
2: Let array  $\mathbf{keys}$  contain the set of unique points in  $d$ 
3: Get array  $\mathbf{relativeFreq}$  = relative frequency of each count in  $\mathbf{n}$ 
4:  $\mathbf{oldSum} \leftarrow \sum(\mathbf{n})$ 
5:  $\mathbf{change} \leftarrow 0.1$ 
6: Initialize empty array  $\mathbf{newFreqList}$ 
7: for all  $i$  in  $\mathbf{relativeFreq}$  do
8:   if  $i \neq 0$  then
9:      $\mathbf{newFreq} \leftarrow i + \text{UniformRV}[-\mathbf{change} \times i, \mathbf{change} \times i]$ 
10:   else
11:      $\mathbf{newFreq} \leftarrow i + \text{UniformRV}[-\mathbf{change}, \mathbf{change}]$ 
12:   end if
13:  $\mathbf{newFreqList.append}(\mathbf{newFreq})$ 
14:   for all  $k, f$  in  $\mathbf{keys}$  and  $\mathbf{newFreqList}$  do
15:      $\mathbf{numVal} \leftarrow (f / \sum(\mathbf{newFreqList})) \times 100$ 
16:      $\mathbf{newFreq} \leftarrow (\mathbf{numVal} / 100) \times \mathbf{oldSum}$ 
17:      $\mathbf{gen} \leftarrow$  array of value  $k$  repeated  $\lfloor \mathbf{newFreq} \rfloor$  times
18:      $\mathbf{newData.append}(\mathbf{gen})$ 
19:   end for
20: end for

21: if  $|\mathbf{newData}| < L$  then
22:    $\mathbf{deficit} \leftarrow L - |\mathbf{newData}|$ 
23:    $\mathbf{minVal}, \mathbf{maxVal} \leftarrow \min(\mathbf{newData}), \max(\mathbf{newData})$ 
24:    $\mathbf{r} \leftarrow \text{DiscreteUniformRV}[\mathbf{minVal}, \mathbf{maxVal}]$  of length  $\mathbf{deficit}$ 
25:    $\mathbf{newData.append}(\mathbf{r})$ 
26: else
27:    $\mathbf{k} \leftarrow \lfloor L - |\mathbf{newData}| \rfloor$ 
28:   for  $=1$  to  $\mathbf{k}$  do
29:      $\mathbf{newData.pop}()$ 
30:   end for
31: end if
32: Set any element in  $\mathbf{newData}$  to 0 if it is less than 0
33: return  $\mathbf{newData}$ 
```

▷ Cleanup code to ensure $|\mathbf{newData}| = L$

B.0.3 Gradient Results for the Budget Allocation Method

We outline a budget allocation method for communities aiming to enhance information diffusion time during disasters.

The average diffusion time (T) is determined by various factors, represented by a feature vector ($\mathbf{p}$), learned through a Gaussian Regression model. To develop the budget allocation method, we utilize a gradient-based approach. The gradient, which indicates the rate of change of diffusion time with respect to each feature in the vector, is calculated by comparing the average diffusion time at slightly modified feature values. Specifically, we increment each

feature by a small amount (δ) and observe the resulting change in diffusion time. For most features, we set δ to 1, but for features with high magnitudes, we use a larger value, typically 10. This allows us to capture the impact of each feature on diffusion time accurately.

Using the developed Gaussian Regression model, we evaluate the average diffusion time at these modified feature values to compute the gradient. This approach guides the allocation of resources by highlighting which features have the most significant influence on improving information diffusion time during disasters.

We provide details on the Gradient Results here in Table [B.1](#).

Table B.1: Table of Gradient Results

Features	Buchanan Rural	Buchanan Suburban	Platte Suburban	Platte Urban	Riley Rural	Riley Suburban	Riley Urban	Rural	Suburban	Urban
Mean of Degree Distribution	-0.986	-0.021	-0.273	-0.633	-0.337	0.689	-0.417	-1.458	-0.355	-0.546
Variance of Degree Distribution	0.436	0.121	1.149	0.789	1.085	2.111	1.005	-0.036	1.067	0.876
Skewness of Degree Distribution	3.712	0.449	4.425	4.065	4.361	5.387	4.281	3.241	4.343	4.152
Kurtosis of Degree Distribution	-0.256	0.052	0.457	0.097	0.393	1.420	0.313	-0.727	0.375	0.184
Mode of Interaction with People	0.130	0.091	0.843	0.484	0.779	1.806	0.699	-0.341	0.761	0.570
Median Interaction with People	-3.526	-0.275	-2.813	-3.173	-2.877	-1.851	-2.957	-3.998	-2.895	-3.086
Mode of Trust with People	0.982	0.176	1.695	1.335	1.631	2.658	1.551	0.511	1.613	1.422
Median Trust with People	-0.662	0.012	0.051	-0.308	-0.013	1.014	-0.093	-1.133	-0.031	-0.222
Mode of Interaction with Local Government	0.624	0.140	1.337	0.977	1.273	2.299	1.193	0.153	1.255	1.064
Median Interaction with Local Government	-2.045	-0.127	-1.331	-1.691	-1.396	-0.369	-1.475	-2.516	-1.413	-1.605
Mode of Interaction with Cable News	0.138	0.091	0.851	0.491	0.787	1.813	0.707	-0.334	0.769	0.577
Median Interaction with Cable News	-2.615	-0.184	-1.902	-2.262	-1.966	-0.939	-2.046	-3.086	-1.984	-2.175
Mode of Interaction with Online Social Network	-0.480	0.030	0.233	-0.127	0.169	1.195	0.089	-0.952	0.151	-0.040
Median Interaction with Online Social Network	0.508	0.129	1.221	0.861	1.157	2.184	1.077	0.037	1.139	0.948
Mode of Interaction with Local News	-0.517	0.026	0.196	-0.164	0.132	1.159	0.052	-0.988	0.114	-0.077
Median Interaction with Local News	1.911	0.269	2.624	2.264	2.560	3.587	2.480	1.440	2.542	2.351
Mode of Interaction with Print News	-0.263	0.051	0.450	0.090	0.386	1.412	0.306	-0.735	0.368	0.177
Median Interaction with Print News	0.596	0.137	1.309	0.949	1.245	2.272	1.165	0.125	1.227	1.036
Mode of Trust with Local Government	1.546	0.232	2.259	1.899	2.195	3.221	2.115	1.075	2.177	1.986
Median Trust with Local Government	0.925	0.170	1.638	1.278	1.574	2.601	1.494	0.454	1.556	1.365
Mode of Trust with Cable News	0.381	0.116	1.094	0.735	1.030	2.057	0.950	-0.090	1.012	0.821
Median Trust with Cable News	-2.087	-0.131	-1.374	-1.733	-1.438	-0.411	-1.518	-2.558	-1.456	-1.647
Mode of Trust with Online Social Network	0.268	0.105	0.981	0.621	0.917	1.943	0.837	-0.204	0.899	0.707
Median Trust with Online Social Network	-0.828	-0.005	-0.115	-0.474	-0.179	0.848	-0.259	-1.299	-0.197	-0.388
Mode of Trust with Local News	0.535	0.131	1.248	0.889	1.184	2.211	1.104	0.064	1.166	0.975
Median Trust with Local News	-1.707	-0.093	-0.994	-1.354	-1.058	-0.032	-1.138	-2.179	-1.076	-1.268
Mode of Trust with Print News	0.062	0.084	0.775	0.415	0.711	1.737	0.631	-0.410	0.693	0.502
Median Trust with Print News	2.638	0.341	3.351	2.991	3.286	4.313	3.207	2.166	3.269	3.077