The author(s) shown below used Federal funding provided by the U.S. Department of Justice to prepare the following resource:

Document Title:	**Human-Centered Approach To Technology To Combat Human Trafficking**
Author(s):	**Julia Deeb-Swihart**
Document Number:	**309486**
Date Received:	**September 2024**
Award Number:	**2018-75-CX-0031**

HUMAN-CENTERED APPROACH TO TECHNOLOGY TO COMBAT HUMAN TRAFFICKING

A Dissertation
Presented to
The Academic Faculty

By

Julia Deeb-Swihart

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy in the
School of Interactive Computing
College of Computing

Georgia Institute of Technology

Dec 2022

HUMAN-CENTERED APPROACH TO TECHNOLOGY TO COMBAT HUMAN TRAFFICKING

Thesis committee:

Dr. Amy Bruckman
School of Interactive Computing
Georgia Institute of Technology

Dr. Alex Endert
School of Interactive Computing
Georgia Institute of Technology

Dr. Keith Edwards
School of Interactive Computing
Georgia Institute of Technology

Dr. Polo Chau
School of Computational Science & Engineering
Georgia Institute of Technology

Dr. Ross Maciejewski
School of Computing, Informatics & Decision Systems Engineering
Arizona State University

Date approved: May 26, 2022

”Enslave the liberty of but one human being and the liberties of the world are put in peril”

William Lloyd Garrison, Abolitionist

ACKNOWLEDGMENTS

First and foremost I would like to thank my advisers, Amy Bruckman and Alex Endert. Without their patience and constant support, I would have been doomed to never complete my thesis. Words cannot adequately express my gratitude – I am forever thankful for everything they have done. Both Amy and Alex showed me how to be a the best mentor and teacher, and I continue to use their example in my own teaching.

I would also like to thank my committee members – Keith Edwards, Polo Chau, and Ross Maciejewski – all of whom provided invaluable advice and feedback.

Special thanks are due to the friends and colleagues who made this work possible. Eric Gilbert, who played a huge role in formation of this research. My collaborators at Marinus Analytics, Thorn, the Atlanta Police Department, and the Dekalb County Police Department all of whom provided insights and expertise for my research. My extraordinary lab mates – Chaya Hiruncharoenvate, Tanu Mitra, Michaelanne Dye, Eshwar Chandrasekharan, Sucheta Ghosal, and Shagun Jhaver – whose friendship made the PhD process a joy. My colleagues at the University of Washington who provided advice and support for finishing a thesis while starting a new faculty position. I am also thankful for the incredible mentors and friends I have worked with at Georgia Tech – including Munmun de Choudhury, Tanu Mitra, Irfan Essa, and Stevie Chancellor.

Finally, I would like to thank my husband and partner Tim. I am incredibly grateful for his emotional and technical support. Without his efforts, this research would not be possible.

The author also gratefully acknowledges the support for this work offered by National Institute of Justice under grant award number 2018-75-CX-0031 and by their Graduate Research Fellowship in Science, Technology, Engineering, and Mathematics. Any views and conclusions contained herein are those of the author, and do not necessarily represent the official positions, express or implied, of the funders.

TABLE OF CONTENTS

Acknowledgments iv

List of Tables xi

List of Figures xii

List of Acronyms xiv

Summary xv

Chapter 1: Introduction 1

1.1 Thesis Overview 2

1.2 Research Significance and Contributions 3

Chapter 2: Background 7

2.1 Human Trafficking Overview 7

2.1.1 Human Trafficking Defined 7

2.1.2 Sex Trafficking vs. Consensual Sex Work 8

2.1.3 Extent of Human Trafficking 10

2.1.4 Human Trafficking Ecosystem 10

2.1.5 Human Trafficking Circuits 12

2.1.6 Sex Trafficking and Technology 13

2.2 Industrial Context 14
2.2.1 American Law Enforcement 14
2.2.2 Role of Law Enforcement 15
2.2.3 Linkage Blindness 16
Chapter 3: Related Work 18
3.1 Detecting Collective Movement 18
3.1.1 Inferring Social Behavior 19
3.1.2 Working with Social Media & Online Data 20
3.2 Visualizing Collective Movement 21
3.2.1 Visualizing Flow 21
3.2.2 Visualizing Group Behavior 22
Chapter 4: Understanding Law Enforcement Strategies and Needs for Combating Human Trafficking 24
4.1 Introduction 24
4.2 Related Works 25
4.2.1 ICT and Law Enforcement 25
4.2.2 Big Data and Law Enforcement 26
4.3 Methods 27
4.3.1 Participant Backgrounds 29
4.3.2 Limitations 30
4.4 Findings 30
4.4.1 Overview of Investigation Process 30

4.4.2 Tools Used During an Investigation 38

4.4.3 Collaboration During an Investigation 40

4.4.4 Sociotechnological Needs 42

4.5 Discussion 47

4.5.1 Design Challenges for Building Solutions for Law Enforcement . . 47

4.5.2 Information Visualization Implications 47

4.5.3 Policy Implications 48

4.5.4 Privacy Implications 49

4.5.5 Design Limitation Implications 50

4.6 Conclusion 51

Chapter 5: Ethical Tensions in Computing Applications for Addressing Human Trafficking – A Human Rights Perspective 52

5.1 Introduction 52

5.2 Background and Definitions 54

5.2.1 Defining AI 54

5.2.2 Overview of Human Rights-Based Approaches (HRAs) 55

5.3 Related Work 56

5.4 Methods 58

5.4.1 Search and Screening Process 59

5.4.2 Data Extraction and Synthesis 60

5.4.3 Limitations 62

5.5 Findings 62

5.5.1 Problem Areas & Use Cases 63

5.5.2 Intended Users 70

5.5.3 Discussion of Ethics 70

5.6 Application of the Human Rights Framework 71

5.6.1 Privacy, Safety, & Security 71

5.6.2 Accountability 75

5.6.3 Transparency and Explainability 76

5.6.4 Fairness and Non-Discrimination 78

5.6.5 Human Control of Technology 81

5.6.6 Professional Responsibility 82

5.6.7 Promotion of Human Values 83

5.7 Discussion 84

5.7.1 Broader Use of Participatory Design 84

5.7.2 Broadening Research to Address Gaps and Engage with Other Forms of Trafficking 86

5.7.3 Development of Best Practices for Harm Prevention 87

5.7.4 Broader Inclusion of Ethics Disclosures in Research and Transparent Discussions of Limitations 88

5.7.5 Limitations of AI for Human Trafficking 89

5.8 Conclusion 90

Chapter 6: Visual Analytics Prototype for Uncovering Human Trafficking Operations using Spatiotemporal Matching 91

6.1 Introduction 91

6.2 Related Work 93

6.2.1 Trajectory Mining 93

6.2.2 Trajectory Analysis in Criminal Justice 93

6.2.3 Visual Analytics in the Criminal Justice Field 94

6.3 Design Process, Analytical Tasks, and Requirements 95

6.3.1 User Tasks and Analytics Questions 96

6.3.2 Design Goals 97

6.4 Datasets Overview 100

6.4.1 CityXGuide Dataset 100

6.4.2 Mock Dataset 101

6.5 User Interface Design and Implementation 103

6.5.1 Database and Processing 103

6.5.2 Trajectory Querying and Matching 104

6.5.3 Inspect View 105

6.5.4 Search View 107

6.5.5 Analyze View 110

6.6 Usage Scenarios 112

6.6.1 Usage Scenario 1: Finding an Interstate Commercial Sex Trafficking Ring 113

6.6.2 Usage Scenario 2: Finding an Interstate Illicit Massage Parlor Ring 117

6.7 Case Studies 123

6.7.1 Case Study 1: Investigating Frequent Movement 124

6.7.2 Case Study 2: Finding matches with Limited to No Movement . . . 128

6.8 Discussion and Limitations 131

6.9 Ethics Analysis 134

6.10 Conclusion 136

Chapter 7: Reflections and Conclusions 139

7.1 Open Challenges with Combating Human Trafficking 140

7.2 Open Questions & Future Research 142

7.2.1 Addressing Data Problems and Access 142

7.2.2 Addressing Resource Constraints and Allocation 142

7.2.3 Addressing Collaboration and Partnership 143

7.2.4 Addressing Prevention 143

7.2.5 Addressing Protection and Survivor Support 144

7.3 Takeaways 145

Appendices 146

References 147

Vita 171

LIST OF TABLES

1.1 Overview of dissertation projects 4

4.1 Overview of interview study participants 28

5.1 Table showing the breakdown of papers by intended use case and targeted area of intervention 65

LIST OF FIGURES

2.1 Spectrum of Sexual Exploitation 10

4.1 Overview of general investigation process 31

4.2 Example of link chart in IBM's i2 Analyst Notebook 40

4.3 Investigation collaborators and their roles 41

5.1 Histogram of publications per year 63

5.2 Distribution of papers in corpus across 5 problem areas 64

6.1 Examples of converging and diverging trajectories 102

6.2 Inspect View 106

6.3 Comparison for likely spam vs unlikely spam 107

6.4 Example of posting to multiple cities and traveling via airfare 108

6.5 Search View 109

6.6 Search options control panel 109

6.7 Analyze View 110

6.8 Detailed Comparison View 111

6.9 Trajectory for Phone A vs Interstate I85 114

6.10 Comparison of user specified search radii for case study 115

6.11 Case Study's Analyze View Results 116

6.12 Pairwise Comparison View 117

6.13 User Scenario 2 Inspect View Results 118

6.14 User Scenario 2 Search View Results 120

6.15 Case Study 2 Analyze View Results 121

6.16 Case Study 2 Detailed Comparison in Analyze View 122

6.17 Trajectory Details for Case Study 125

6.18 Overview of Matched Details 127

6.19 Hyper-local Trajectory Details 129

6.20 Overview of Matched Details 131

6.21 Example trajectories in reverse order from each other 132

6.22 Comparison of user specified order search options 132

LIST OF ACRONYMS

LCS Longest Common Subsequence

TVPA Trafficking Victim Protection Act of 2000

SUMMARY

Human trafficking is a serious crime that continues to plague the United States. With the rise of computing technologies, the internet has become one of the main mediums through which this crime is facilitated. Fortunately, these online activities leave traces which are invaluable to law enforcement agencies trying to stop human trafficking. However, identifying and intervening with these cases is still a challenging task. The sheer volume of online activity makes it difficult for law enforcement to efficiently identify any potential leads. To compound this issue, traffickers are constantly changing their techniques online to evade detection. Thus, there is a need for tools to efficiently sift through all this online data and narrow down the number of potential leads that a law enforcement agency can deal with. While some tools and prior research do exist for this purpose, none of these tools adequately address law enforcement user needs for information visualizations and spatiotemporal analysis. Thus to address these gaps, this thesis contributes an empirical study of technology and human trafficking. Through in-depth qualitative interviews, systemic literature analysis, and a user-centered design study, this research outlines the challenges and design considerations for developing sociotechnical tools for anti-trafficking efforts. This work further contributes to the greater understanding of the prosecution efforts within the anti-trafficking domain and concludes with the development of a visual analytics prototype that incorporates these design considerations.

CHAPTER 1

INTRODUCTION

Human trafficking is a complex crime where victims are forced to engage in commercial sex, labor, or other services for the financial gain of others [1]. Throughout the world, including within the United States, millions of people fall victim to this type of crime [2]. Despite efforts to combat this crime, human trafficking remains a significant global humanitarian crisis and the number of reported cases continues to rapidly increase every year [2, 3].

Law enforcement plays a significant and critical role in anti-trafficking efforts. They are usually the first to identify and respond to human trafficking cases and are often in the best position to connect victims to appropriate social service providers[4]. Many survivors mentioned that they had some contact with law enforcement while they were being trafficked and further noted that their interactions with law enforcement helped them leave their traffickers [5]. As a result, designing new tools to support law enforcement in this effort has a major impact on combating human trafficking.

However, the majority of human trafficking survivors are never identified. The US State Department estimates that less than 1% of victims have ever been identified [6]. Law enforcement continues to struggle with identifying victims of sex trafficking - with many describing feeling like they are only just scratching the surface with current strategies and tools [7].

Thus, social computing researchers have the unique opportunity to assist with combating this crime. Increasingly, traffickers have adopted internet technologies to facilitate their criminal activities [8]. Fortunately, these online activities leave behind traces which can be used to uncover their operations and identify their victims [8, 9, 10]. However, despite increasing research into leveraging these online traces to combat human trafficking, identifying victims remains a deeply challenging task. The sheer volume of online activity makes it difficult for law enforcement to efficiently identify any potential leads using traditional methods; depending on the site, there are as many as thousands of new ads posted per day [11]. To compound this issue, traffickers are constantly changing their techniques online to evade detection meaning that law enforcement have to constantly look for new indicators of trafficking. Thus, there is an urgent need for tools to efficiently pair down the number of posts for law enforcement to begin investigations from. Additionally, such analytics tools also need to support law enforcement users' understanding of the underlying models and support the ability for users to directly interact with those models. As informed by prior work on law enforcement needs [12, 13], it is essential for these analytic tools to include information visualizations so that law enforcement users can understand the data patterns and the models that underlie the tools.

developing visual analytics prototypes that use computational techniques to assist with identifying potential victims of trafficking. The interactive visualizations will enable users to understand and interact with those complex models to understand their data, gain insight, and generate leads from which to begin investigations.

1.1 Thesis Overview

To address these gaps, my work focuses on three areas to study. First, the understanding of the unique computational needs of law enforcement working human trafficking cases. Second, the understanding of the political and ethical tensions present in design and development of future computational tools. And third, how such future computational work can both be designed to address the unique needs for law enforcement while also incorporating ethics into the design process. To this end, ,my dissertation consists of three interrelated projects. First, an interview study with law enforcement for to purposes of understanding their specific sociotechnical needs (chapter 4). Second, a systematic literature review to best understand how to ethically develop tools within this field (chapter 5). And finally, the development of a visual analytics prototype designed to address the needs we identified in both prior chapters (chapter 6). Ultimately, the overarching goal of this work is to uncover the complexities and nuance present when developing computational tools for such challenging subject area. Further, this work also demonstrates the potential utility in incorporating ethical analysis within the development process through developing a preliminary prototype for investigating human trafficking cases. An overview of each of these projects can be seen in Table 1.1.

At a high-level, these projects each address the following research questions:

- **RQ1**: What are the sociotechnical needs of law enforcement working human trafficking cases?
- **RQ2**: What are the ethical tensions present in current efforts to apply computing methods within the anti-trafficking domain? How can future research within this context be conducted thoughtfully and ethically?
- **RQ3**: How can we build models to find patterns in data that law enforcement currently uses? How can we iteratively improve a tool with these models to ensure that it supports law enforcement needs in responding to and investigating human trafficking? How can we use existing visualization approaches to help users understand the model and incorporate their domain expertise?

1.2 Research Significance and Contributions

My work contributes to the existing efforts to combat trafficking while incorporating the unique needs of criminal justice practitioners. As noted earlier, the overarching goal of this thesis is the development of a prototype for a tool to assist law enforcement with identifying connected cases of human trafficking. By enhancing law enforcement's ability to identify potential victims through the use of technology, this research helps meet the goals outlined

Table 1.1: An overview of the three projects that make up this dissertation.

Research Project	Methods	Research Goals	Research Questions
P1: Interview study with law enforcement	Semi-structured interviews Thematic analysis	Gather law enforcement socio-technical needs Understand the investigation process and the role of technology	What socio-technical solutions would be most useful in an investigation?
P2: Literature Review and Analysis	Systemic Literature Review Data extraction and synthesis through qualitative analysis Human Rights-Based Analysis	Examine the current state of the art Uncover the ethical tensions with using AI in anti-trafficking research	What are the ethical tensions when applying AI in the anti-trafficking domain? How can future AI research be conducted ethically?
P3: Building Visual Analytics Prototypes	Spatiotemporal trajectory matching and querying User-Centered Design	Develop models to measure the similarity of movement trajectories Develop a functioning prototype	How can we automatically match trajectories with imprecise data points? How can we design a visualization that helps users understand the model and incorporates feedback in real-time?

in the Federal Strategic Action Plan on Services for Victims of Human Trafficking [14] and the goals for law enforcement investigations outlined in the U.S Department of Justice's National Strategy for Child Exploitation Prevention and Interdiction [15]. Additionally, by examining and discovering patterns in real-word data this work can improve our understanding of the online human trafficking ecosystem which will enable policy makers and law enforcement to use empirical models to inform decision making.

This work also contributes to the field of spatiotemporal analysis and visualization through the design and development of a tool that incorporates movement trajectory querying processes that overcomes the challenges of working with messy, incomplete online data. This has important applications not only for the domain of human trafficking but also for other domains such as city planning, tourism, and crisis management. For example, recent efforts to model real-world movement patterns of tourists from online behavior has resulted in the development of tools for tour guides and city planners. Finally, our work contributes to the field of visualization by demonstrating a novel application of visual analytics for interactively explore geospatial and temporal sequences.

In summary, this thesis offers the following research contributions:

- Through an interview study with 16 law enforcement personnel, we detail the unique technological and design needs of law enforcement investigating human trafficking cases. This work further outlines their investigation process, the complex political and technological challenges associated with designing for law enforcement, and uncovers areas where future HCI research can contribute. Finally, our last contribution is a detailed account of the path forward for the development of new tools for anti-trafficking efforts.

- Through conducting a systematic literature review of computing research and applying a framework based in international human rights law, our work contributes the categorization and understanding of the ethical risks associated with developing AI tools for anti-trafficking efforts. This work uncovers a number of ethical tensions including bias endemic in datasets, privacy risks stemming from data collection and reporting, and issues concerning potential misuse. This work further contributes to this area by highlighting four suggestions for future research: broader use of participatory design; engaging with other forms of trafficking; developing best practices for harm prevention; and including transparent ethics disclosures in research.

- Through the development of a visual analytics prototype, this work contributes a series of unique design goals and considerations for developing computational techniques for law enforcement investigations into human trafficking. This work also contributes case studies, usage scenarios, and ethical analysis of the prototype to demonstrate how such techniques could be used to identify connected cases of human trafficking while incorporating the user's expert domain knowledge. Finally, the design contributes a novel approach that combines existing techniques for similarity matching and visualization in order to allow expert users to flexibly define group behavior according to their varied case needs.

CHAPTER 2
BACKGROUND

The first half of this chapter will give an overview of sex trafficking while highlighting the role technology plays in the ecosystem. In this section I will also discuss what current technological interventions have been proposed and argue for further data-driven approaches to combat this complex problem. Then, I give an overview of American Law Enforcement and their response to sex trafficking to give context for the user group.

2.1 Human Trafficking Overview

2.1.1 Human Trafficking Defined

Human Trafficking is a complex crime that takes a variety of forms. The circumstances and situations each trafficking victim faces varies drastically from person to person. Thus, historically human trafficking has been difficult even for experts to define. For example, child marriages are defined as a form of human trafficking by the United Nations but is not defined as trafficking per se under US federal laws [16].

Because our work is situated within a US context, we use the US federal definition established by the Trafficking Victim Protection Act of 2000 (TVPA). This definition categorizes human trafficking into two types: labor trafficking and sex trafficking. Under the TVPA, sex trafficking refers to the crime of "recruiting, harboring, transporting or obtaining a person for the purposes of a commercial sex act [1] induced by force, fraud, or coercion, or in which the person induced to perform such an act has not attained 18 years of age" [17]. This definition of sex trafficking makes a makes a clear distinction between adult and child victims. Because children cannot consent, no force, fraud, or coercion needs to be proved

[1] Under this definition, commercial sex acts include trafficking for sexual purposes, sex tourism, prostitution, and pornography.

for a child to be considered a trafficking victim - though, these elements are often present. Note that sex trafficking is distinct from cases in which adults willingly engage in sex work without coercion - victims of human trafficking do not consent to their exploitation.

Labor trafficking on the other hand refers cases where an person is forced against their will to work in typically an otherwise lawful industries such as agriculture or domestic work [18]. The TVPA defines labor trafficking as the crime of ”recruiting, harboring, transporting or obtaining a person for labor or services, through the use of force, fraud, or coercion for the purposes of subjection to involuntary servitude, peonage, debt bondage, or slavery” [17]. Labor traffickers often lure victims through false promises of gainful employment and then force them to work long hours for little to no pay in unsafe working environments.

Regardless of the type of trafficking, no movement needs to occur for a victim to be considered a trafficking victim. However for law enforcement and experts, frequent movement between locations is a strong indicator of trafficking [19, 20].

2.1.2 Sex Trafficking vs. Consensual Sex Work

The TVPA's definition also separates consensual sex work where an adult engages in sex work voluntarily from sex trafficking. However, drawing the line between what constitutes consensual versus exploitation is a complicated task even within these definitions. Many cases, especially those concerning survival sex (where a person engages in sex work due to extreme need like homelessness), end up falling in a grey area between consensual sex work and trafficking. For example, consider a case where a person engages in sex work due to lack of other viable jobs and uses a "pimp" for protection who then takes part of the wages. There is clearly some form of exploitation happening both socioeconomically and inter-personally with the "pimp", but it complicated to say whether that case would necessarily be considered trafficking and requires investigating a lot of nuance. This is one of the reasons identifying trafficking victims is a challenging and labor-intensive task for law enforcement.

Additionally, distinguishing between trafficking and sex work is complicated by the fact that in the US voluntary sex work is a crime[2]. However, many criminal justice practitioners including law enforcement will decline to arrest and prosecute adults who engage in sex work. Our prior work supports this as all of the law enforcement officers we spoke to mentioned that they are not interested in arresting adult sex workers and believe that arresting these adults is a harmful practice [13].

Instead of thinking of trafficking and consensual sex work as discreet and separate categories, many researchers and activists think of sexual exploitation as a spectrum with adults who willingly and freely consent in sex work on one end of the spectrum and those who are forced and therefore trafficked on the other end. Figure 2.1 highlights the exploitation spectrum with example cases. Note however, that the case of a minor engaging in survival sex is considered a trafficking victim under the TVPA regardless of whether or not that minor "willingly" engages in sex work. In fact, many juvenile trafficking cases fit into this category where the child is homeless because they ran away from home or were kicked out and subsequently turn to sex work to survive. However as that child is underage, they cannot consent to engage in sex work and thus their customers are one of the exploiters in this situation.

Our work uses the strict definition of trafficking as defined by the TVPA. Our work also doesn't identify cases of trafficking for law enforcement, but instead the role of my thesis work is to help law enforcement generate leads for trafficking cases. Without a full investigation by law enforcement, there is no certain way to determine if a case is trafficking or not.

[2]Except in two counties in Nevada, neither of which contain Las Vegas

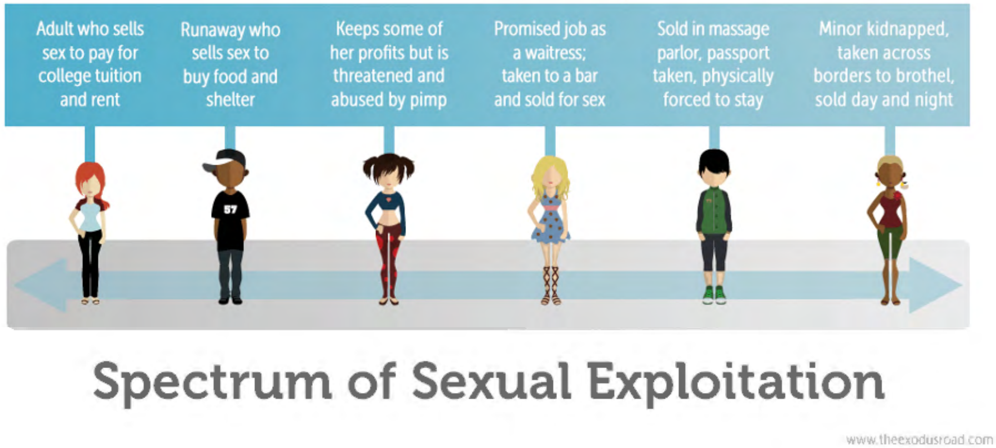

Figure 2.1: Examples of cases on the sexual exploitation spectrum. Image taken from materials from The Exodus Road, an NGO dedicated to combating human trafficking

2.1.3 Extent of Human Trafficking

Because of the underground and hidden nature of sex trafficking, it is difficult to estimate the number of victims or the scope of the problem [21]. As such there currently is no reliable estimate for the number of trafficking victims at the federal level [5, 7]. However from reports from the National Human Trafficking Hotline, we know that sex trafficking has occurred in all 50 states and affects both citizens and non-citizens alike [5]. Further since 2007, the hotline has identified 63,380 situations of human trafficking, of which 22,326 victims were identified in the US in 2019 alone [22].

2.1.4 Human Trafficking Ecosystem

Much of the research to date on sex trafficking has focused on understanding the human trafficking ecosystem and the role technology plays in facilitating these crimes. When discussing the roles individuals play in the crime of sex trafficking, prior literature often separates the individuals involved into three categories: traffickers – sometimes referred colloquially as “pimps” or “madams”, customers – sometimes referred colloquially as “johns”, and victims. We continue to separate these roles is our work; however, it is important to note that these categories are not necessarily distinct, which can make the tasks of teasing apart a person's involvement in a trafficking operation more difficult. Some victims of human trafficking are, for example, forced by their trafficker to recruit new victims making them both a victim and a trafficker [7].

Traffickers employ a variety of techniques to maintain control over their victims. Some use overt violence towards the victims or the victim's families; others use more subtle techniques to exert control including promises of love and marriage, false promises of employment, or ever-increasing debt [18]. Many traffickers also use illicit drugs to control and coerce their victims by using drugs to make the victim more compliant and dependent on the trafficker and then using the threat of withdrawal to force the victim to continue to work [23]. Even the trafficker's operations vary drastically; some exploiters operate as a sophisticated agency with dedicated staff for recruitment and security, while others operate alone or part-time [24, 18, 7].

Human trafficking also intersects with other areas of crime because it is an especially profitable crime. In a report for the National Institute of Justice, The Urban Institute estimated that this crime generated anywhere from $39.9 million to $290 million in 2014 alone depending on the city [7]. Children are especially vulnerable to sex trafficking because they are perceived by traffickers as more profitable and easier to control [15]. Thus, people engaged in other areas of crime have shifted towards trafficking instead or have included sex trafficking as part of their criminal organization [18]. In addition, many traffickers perceive human trafficking as a lower risk crime especially compared to drug trafficking [7]. Thus, researchers and criminal justice practitioners are finding more and more that sex trafficking exists within other areas of crime including human smuggling, drug trafficking, organized crime, and gang activity. The Urban Institute that found that gang involvement with sex trafficking is increasing [7]. Other reports such as Polaris' recent one on cantina cases have also pointed towards the involvement of organized crime in sex trafficking with human smuggling as a component [25]. When it comes to child victims, the FBI noted in a report for the Department of Justice that most traffickers who victimize children have prior criminal records particularly associated with violent crimes and that there has been an increase in the number of traffickers associated with gangs [18]. This is important to note because it means any interventions designed to combat human trafficking can also help combat other crimes.

2.1.5 Human Trafficking Circuits

Traffickers may move victims from city to city within the US - traveling using the interstate highway system or via local airports [26]. This strategy of frequent movement is employed by traffickers to evade detection, increase access to larger customer-bases, and prevent victims from escaping [21].

Some traffickers follow predictable paths, but most are highly unpredictable. What paths a trafficker takes largely depends on the scale of their operation, access to financial and logistical resources, and what clientele are available [26, 21]. For example, there have been several documented cases [25, 27, 21, 28] where a trafficking operation owns several locations that the traffickers ”rotate” victims between. Since 2007 over 1,300 victims have been trafficked in the US to provide commercial sex for these kind of organizations [25]. These operations tend to be larger-scale, international organizations with multi-state criminal networks that operate out of bars, nightclubs, illicit massage parlors, and/or strip clubs that are managed by the criminal organization. Because the traffickers own the locations, the victims move between known locations and follow predictable stops, but the ordering and intervals between stops are irregular.

With other human trafficking operations, sometimes only the origin and destinations cities are planned in advance by traffickers. Movement between these locations is highly irregular but tends to follow stops along a predictable path such as traveling back-and-forth along a highway system and stopping at major cities [26]. In the literature, some researchers and experts refer to these kinds of interstate paths as ”circuits” or ”pipelines”. Famous examples of these circuits include the Midwest or Minnesota pipeline (where traffickers move victims across the US between Washington State and New York State [7]), the West Coast Circuit (which travels up and down the US West Coast from San Diego, CA to Portland, OR [7, 20]), and the Eastern Circuit (which travels up and down the east coast from Miami, FL to New York City [20, 21]). One example of this type of operation that followed the ”east-coast circuit” involved two traffickers who moved victims between Georgia and Virginia using greyhound buses and opportunistically stopped at major cities to both recruit more victims and exploit their existing victims [29].

2.1.6 Sex Trafficking and Technology

Sex traffickers in the US often use internet technologies to facilitate human trafficking. With the widespread availability of internet technologies, many traffickers have moved away from traditional avenues of advertising and selling victims and instead advertise online [8]. Previous work has found that the majority of sex trafficking victims in the US are advertised online [30]. These sites provide a faster and easier way for customers to find victims and pay for services than traditional means. Other websites provide an online portal for these customers to discuss their activities and to review sex workers. These technologies provide easier avenues for traffickers to connect with customers at little to no cost to themselves.

While these online sites facilitate sex trafficking, these sites also provide vital information to assist law enforcement and NGOs in identifying and rescuing victims. Online advertisements and post leave behind digital traces that allow law enforcement to track these activities. For example In 2011, the FBI and the Jacksonville Sheriffs’ Office investigated and successfully took down Tyrone Townsend’s operations. Townsend had been forcing several women to perform commercial sex acts up and down the east coast – beating them if they didn’t comply or get enough customers. The 28 advertisements for the victims as well as the GPS information from a personal navigation device were all crucial to getting Townsend convicted [29].

While these ads are often invaluable, the sheer volume of ads posted on these sites makes it difficult for these organizations to identify cases of trafficking. Depending on the site, there are as many as thousands of new ads posted per day [11]. Much of the prior work seeks to address this issue by utilizing artificial intelligence to detect instances of trafficking on individual advertisements [8, 31, 32]. One of the earliest approaches for identifying potential victims was developed by Mark Latonero [8] which provided a preliminary software tool that filtered online advertisements based on a list of key words generated by law-enforcement. This work focused primarily on Backpage and aimed to help law enforcement narrow their returned searches on this site for potential trafficking victims. This work eventually went towards a system called TrafficBot which featured a searchable database for law enforcement [32]. More recent work [33] and [19] build on this idea by using more sophisticated machine learning and natural language processing techniques to filter escort advertisements that are likely are for human trafficking victims. Some of the work mentioned above has also been integrated into commercially available tools: notably Thorn's tool Spotlight and Marinus Analytics's tool TrafficJam

2.2 Industrial Context

Our intended user group is American law enforcement working on sex trafficking cases.

2.2.1 American Law Enforcement

In the US, there are over 18,000 independent police departments – each of which operate at different yet overlapping jurisdictional levels (i.e. federal, county, municipal, tribal, and specialist police departments) and have their own independent powers, responsibility, and funding. For human trafficking investigations, this diversity in police departments results in a lack of universal practices for handling investigations because there is no uniformity in training, education, and access to resources.

Human trafficking cases are investigated at all jurisdictional levels, but are most commonly prosecuted at the federal level. As a result, most investigations involve some level of collaboration between police departments. Though the degree to which each department collaborates often depends on department funding, access to shared training and resources, as well as a prior history of collaboration (such as through mutual membership in a task force).

2.2.2 Role of Law Enforcement

Survivors of sex trafficking have mentioned that they routinely were in contact with people outside of their trafficking situation who had the ability to provide some form of assistance [5]. Thus, these individuals who had contact with victims had the unique potential to provide assistance for them to leave their trafficking situation; however, their help relies on their ability to identify cases of human trafficking [5].

In particular, law enforcement officials are often in this position. Many survivors of sex trafficking have mentioned that they had some contact with the criminal justice system during their trafficking situation [5, 22]. A study for the US Department of Justice noted that the majority of victims who were referred to service providers were initially identified by law enforcement [4]. Further, the National Human Trafficking Hotline found that "many of these [survivors] cited their interaction with law enforcement as the reason they were able to leave their controller" [5] and that law enforcement was the 2nd most common point of access victims reported for getting help [22].

While law enforcement is uniquely poised to identify and intervene with potential victims, law enforcement officials often do not recognize if a person in their custody is a victim of trafficking. Despite training, identifying cases of human trafficking is very difficult. Victims of human trafficking do not often self-identify as such and are usually reluctant to speak with any outsiders [24, 18, 21]. Some traffickers will also coerce their victims to not cooperate with service providers or law enforcement officials [24]. When law enforcement misses the signs that a person is a trafficking victim, that victim is often arrested, convicted, and in some cases deported without getting any help. Long-term, these convictions make it increasingly difficult for victims to find opportunities necessary (such as housing and employment) to leave a trafficking situation [24]. Some victims may be arrested multiple times before someone realizes they are a victim of human trafficking. Thus, there is a real need to develop tools that assist law enforcement with recognizing when they are dealing with a trafficking situation.

2.2.3 Linkage Blindness

One common investigation challenge when investigating serial offenses[3] like human trafficking, is overcoming linkage blindness. Linkage blindness refers to situations where an investigator fails to connect related crimes committed by the same perpetrator [34, 35]. Causes of linkage blindness include a lack of effective resources and techniques for linking cases across jurisdictional boundaries; a lack of information sharing technology especially between police departments; and a lack of established avenues for multi-sector collaboration [34]. Failure to recognize linked cases can lead to lower rates of successful prosecution, lower quality and quantity of evidence, and delays in access to necessary resources and collaborators [34]. With human trafficking investigations, linkage blindness can make it more difficult to identify and recover all victims; to uncover the identities of the traffickers; to recognize connections to related other criminal sectors; and to recognize larger criminal organization involvement. Linkage blindness can also lead to investigators duplicating work across police departments resulting in delays in recovering victims and prosecuting the traffickers.

A number of technical solutions have been proposed to overcome linkage blindness and include augmenting investigation capabilities through behavior pattern matching. Approaches to link crimes use techniques from data science and generally involve querying via similarity of criminal behavior patterns [36, 37]. In particular, matching geospatial and

[3] where one person or group commits two or more connected crimes

temporal behavior patterns has been shown to be one the most effective techniques for uncovering connected cases [38]. In practice, case linkage through geographic and temporal patterns have been used to match serial arsonists, serial rapists, and other violent serial crime [38, 39].

Our work build on this theory by using geospatial and temporal matching to connect potential cases of human trafficking. Further our work incorporates visualizations to help make the underlying models more accessible and understandable to law enforcement.

CHAPTER 3

RELATED WORK

Within the broader fields of spatio-temporal visual analytics, my work draws from scholarship relating to detecting group movement, visualizing flow, and visualizing collective movement.

3.1 Detecting Collective Movement

Collective movement[1] refers to a set of moving objects or people whose movement is inherently coordinated. Examples of collective movement include flocks of geese that migrate together, series of cars that follow the same route, and solders moving across a battlefield in a particular formation [40]. As seen those examples, collective movement can exhibit a number of behavior patterns. While much of the prior work assumes collective movement follows a moving cluster type pattern (where each member of the group moves along the same path give or take some distance from the center), more recent research has shown that collective movement can involve more complex patterns. Examples include behavior where some members deviate from the cluster (ex: a goose during migration misses the flock and travels alone) or where coordination is for a limited time frame (ex: individuals dropping cars off at the airport)[40]. Because of this complexity, techniques to detect collective movement have to account for a wide range of possible coordination patterns. However, techniques designed to detect the full range of possible patterns remains an open research problem [41].

Inferring collective behaviors from trajectory datasets generally involves looking for evidence of coordinated behavior such as finding overlapping trajectory traces, matching

[1]The term collective movement has gone by a wide range of different names in prior work including but not limited to ”coordinated movement”, ”moving cluster”, ”group movement”, and ”flocking behavior”.

trajectories to known coordinated movement patterns, or finding clusters based on similar movement behavior. A critical component of this process is establishing a similarity metric which measures how similar two trajectories are. There are a number of existing metrics to measure similarity, though most typically define similarity as some kind of ratio of the amount of overlap between any two trajectories. The most common metrics involve euclidean distance/P-Norm [42, 43], dynamic time warping (DTW) [44, 45, 46], Longest Common Subsequence (LCS) [47, 48], and hausdorff distance [43, 49]. However, all of these approaches do not allow the user to query based on multiple scales of similarity or allow for matching based on both the similarity of the trajectory points as well as the trajectory descriptors (for example, matching based on both time-location similarity as well as sampling cadence) [50].

3.1.1 Inferring Social Behavior

Detecting collective movement is commonly used to infer social relationships. Because collective movement requires some form of coordination, these patterns are often indicative of some kind of social organization such as herd membership. For example, [51] proposed a coordinated small-multiples system to examine ”co-occurrence” patterns (where two people visit the same location in the same time frame) in mobile phone data through the use of bi-clustering. Similarly, a recent paper by [52] used the frequency of meeting patterns to model social relationships and detect social groups using community detection algorithms. Other prior work by [53, 50, 54] focused instead on detecting ”trend-setter” or ”leader-follower” patterns where a single group member leads the rest. More commonly, the prior work has focuses on identifying moving clusters[2] [55, 56, 57, 58].

One main challenge with these existing approaches for inferring social relationships is that they all rely on the underlying assumption that only a high degree of overlap between two trajectories is indicative of some kind of relationship. Most of these approaches cannot,

[2] also described in the prior work as flocking, swarming, or convoy patterns

for example, detect clusters with inconsistent temporal lag or with group members whose movement deviates from the rest of the group [41, 58]. In addition, other behavior patterns such as collision avoidance, velocity matching, flock centering, attraction, repulsion, and alignment have yet to be detected from natural data [59]. Finally, many of these approaches are also designed to find specific group patterns and cannot find multiple kinds of patterns simultaneously [41].

3.1.2 Working with Social Media & Online Data

With the widespread adoption of social media, extracting movement patterns from online sources has emerged as a new area of study. Social media and other online sites provide a rich source of data on human behavior patterns and the widespread availability of location and time data within these online sites means that we can more broadly understand movement patterns than ever before. Inferring location from online posts has been used in applications such as improving awareness for disaster response [60, 61], modeling disease control [62], and understanding travel patterns within city areas [63, 64]

However, this kind of data presents a number of unique challenges for fundamental research [63]. Unlike with traditional movement datasets which consist of data from tracking devices, trajectories have to be extracted or inferred from the social media data. Typically this is done by defining posts as events and then extracting spatial, temporal, and semantic information from the text and metadata within each post. As a result, the spatial and temporal information extracted from social media posts tends to have a high degree of uncertainty [63, 51]. Further while some social media sites incorporated geo-tags, others require the researcher to either geocode locations based on user specified locations [65] or from inferring the location using other information such as the objects in an image [66]. This process can often introduce even more uncertainty into the extracted trajectories. As a result models that incorporate social media data have to not only be robust to noise, but also account for the uncertainty associated with the data. While some prior work exists that tackles this issue [63, 58], many of these approaches lack the ability to account for long-term temporal trends as these approaches are typically only applied on a snapshot of data. Thus even within my own work, we have to be cognizant of how the combination models to detect group behavior behave when the data is from noisy, incomplete, and uncertain online data.

3.2 Visualizing Collective Movement

Prior work on visualising collective movement usually takes the form of either visualizing flow or on visualizing group behavior more broadly. The emphasis with flow data is on visualizing the transitions between states which can be used to understand collective trends between events or locations. Whereas visualizing group behavior tend to focus less on transitions and instead on movement patterns more broadly.

3.2.1 Visualizing Flow

For flow data, there are two general approaches for visualizing trajectories: non-cartographic flow maps and cartographic flow maps. Non-cartographic flow maps show movement between states in a non-map based representation and can include approaches like origin-destination matrices[67]. For example, FlowStrates uses a matrix heatmap to encode the temporal trends of movement between origin-destination paired data [68]. Another approach by [69] uses parallel sets where the time-ordered vertical axes represent the locations. To show longer trajectories, graphical/node-link approaches have been proposed. TrajGraph uses a node-link view to show movement between nodes that represent abstracted locations [70]. Nodes in this view are automatically generated using clustering algorithms based on user input. For example, users can specify to the system to cluster based on time periods, trajectory speed, etc. Similarly, [71] proposes an approach that uses a state-transition graph where states in the graph represent semantic categories such as “home” or “work”.

Cartographic-based representations typically use either animation, multiple map views, or the space-time cube [72] to encode the temporal patterns [73]. Another possible approach is to use color or other graphical indicators to indicate temporal changes; however, this approach performs poorly when the node-link graph is non-planar as is the case with most movement datasets [74]. For trajectories that span a long period of time and contain a high number of unique locations, animation tends to be an inappropriate choice as the user will be unable to perceive much of the long-term trends across multiple maps [75, 76, 77]. Animation is instead a better choice when users want to compare changes very short time frames or across a small number specific locations [78]. The more popular approach is to use variations of small-multiples. For example, [79] proposed using multiple maps arranged for pairwise comparisons between chosen time frames. [63] allowed for more complex analysis with a multi-view system containing maps, node-link views, and space-time matrix views. However only a small number of maps or charts can be shown simultaneously before the overall visualization is cluttered and ineffective [77]. Thus for longer trajectories or large datasets, some spatial or temporal simplification has to be done [79, 80].

Overall, flow maps visualized cartographically can have problems with occlusion, visual clutter, and lots of intersections that make interpreting the results difficult. On the other hand, non-cartographic visualizations lose the geographic context which may be important for certain user-tasks. To solve the clutter and occlusion problem, some of the prior work has proposed techniques to simplify the spatial and/or temporal data. These approaches can involve interactive filtering [81, 82], edge bundling [83, 84, 85], and data aggregation/clustering [77, 79, 86].

3.2.2 Visualizing Group Behavior

As [87, 59, 41] note, there are very few papers addressing the problem of visualizing group behavior. There is a gap on visualization research designed to answer questions such as "what is the course of the whole group?", "how do the movements of both the group and its members change over time?", "how coherent are the members movements within a group?", and "which members have have particular roles say leader or follower?" [59].

Work by [88] developed a visual analytics system to understand group movement. Their approach used metrics to characterize the positions of each group member with respect to the rest of the group and then these metrics were visualized on a set of connected visualizations including a time graphs and a map view. To compare across groups, the visualization uses color to indicate group membership. Work by [59] uses a similar visualization approach, but incorporates a unique computational method to uncover new attributes within the data to detect group organization structure. Their technique applies a space transformation to convert the trajectories from a cartographic space to an abstract "group space". More recent work by [87], created a visualization called MotionRugs to analyze fish swarming behavior over time. Color in this visualization is used to encode trajectory descriptors such as speed. Unlike the previous two approaches, [87]'s approach loses the geographic context and this approach can not be reliably used to compare groups (an important analytics task for my work). Further, all three of these approaches assume that group membership is known ahead of time and does not incorporate the process of detecting group membership into their systems. This gap represents an opportunity for future research.

CHAPTER 4
UNDERSTANDING LAW ENFORCEMENT STRATEGIES AND NEEDS FOR COMBATING HUMAN TRAFFICKING

In this chapter, I performed a semi-structured interview study with law enforcement who work on human trafficking cases to understand their socio-technical needs. The insights from this study form the backbone of the design goals described in chapter 6.

4.1 Introduction

Law enforcement plays a significant and critical role in anti-trafficking efforts. They are often the first to identify and respond to human trafficking cases and they play a critical role in connecting survivors to necessary support services [4]. Many survivors mentioned that they had some contact with law enforcement while they were being trafficked and further noted that their interactions with law enforcement helped them leave their traffickers [5]. As a result, designing new tools to support law enforcement investigations will have a major impact on combating human trafficking.

However, little is known about what the current needs of law enforcement personnel are - particularly with respect to their technological needs. Furthermore, there is a gap in research concerning the role technology plays in an investigation (despite current research efforts into the development of new tools) and a gap in understanding how collaboration occurs and the role technology plays in facilitating these interactions.

Thus to address these gaps, I conducted an interview study with sixteen law enforcement personnel working on human trafficking cases. Our analysis focused on uncovering their unique technological and design needs. Through analyzing their investigation process and understanding their social technical needs, we highlight areas where future research in HCI can assist and discuss the complex challenges associated with designing for law enforcement. Our contributions outline a path forward for the development of new tools for anti-trafficking efforts.

4.2 Related Works

There is limited research on the role of law enforcement in anti-trafficking efforts; much of the prior work has focused on law enforcement attitudes towards human trafficking [89, 90], the role officers play in victim identification [4, 91], and understanding law enforcement training needs [91]. To the best of our knowledge, there are no studies concerning the role of technology in a law enforcement investigation of human trafficking. Instead, much of the prior work has focused on the role technology plays in facilitating trafficking [9, 92] and the potential methods law enforcement can use to exploit such technology [19, 10]. Additionally, prior work has lead to the development of analytics tools to support law enforcement investigations for sex trafficking [93, 32]. Note that there are no complementary tools designed to combat labor trafficking [94]. As such, we draw on the body of work on information and communication technologies (ICTs) and big data with respect to law enforcement to situate our work.

4.2.1 ICT and Law Enforcement

With efforts to reform policing strategies in the US, law enforcement departments over the past several decades have increasingly adopted ICTs [95, 96]. Adoption of ICTs has the potential to increase the effectiveness and efficiency of a police force while lowering costs [95]. However, recent research has indicated that adoption of ICTs alone does not improve productivity of a particular police department unless adoption is paired with support from management and strategic deployment of such tools [97]. In particular, community policing and data-driven policing are strategies that most often correlate with increases in successful technology adoption [97, 96]. Adoption of ICTs is also associated with certain organizational changes within a police department; departments who adopted ICTs are associated with increases in hiring specialized staff (particularly staff with IT backgrounds), increases in the number of specialized units, and increases in the educational and training requirements for staff across the board [97].

Despite these advantages, law enforcement adoption of ICTs is not universal. Police departments in the US face a serious lack of funding and training for new technology, which is a barrier for effective investigations. This is especially true for human trafficking investigations which (as discussed in our findings section) rely heavily on technology for data mining and analytics. As noted in a study by the non-profit research institute RTI International, many police departments lack technology capable of those tasks. In 2017 only 14% of police departments in the US had technology to share and search for data across silos [96]. Additionally, only 10% had tools for data-mining and only 5% had tools to uncover connections between data points [96].

Furthermore, police departments continue to struggle with collaboration and information sharing despite the increased efforts to use ICTs for this purpose [98, 99, 100, 101]. While departments have increasingly adopted policies to support information sharing, the lack of software standards across multiple platforms makes sharing almost impossible [98]. As discussed in our findings and discussion section, this remains a serious barrier for human trafficking investigations.

4.2.2 Big Data and Law Enforcement

Prior work has also looked at the role of big data in policing contexts. Technology to support data collection and analysis for police departments has been discussed by researchers as a solution towards police accountability and as a solution to help mitigate policing biases like racial profiling [102, 103]. However, recent work by Verma and Dombrowski has found that the technology to support big data in its current form is not enough to support police action and judgment [104]. Furthermore, data collected and analyzed has the potential to either confirm or counter an officer's bias depending entirely on how the data is collected and displayed [104].

In practice, adopting big-data oriented practices in law enforcement faces a number of technical issues - from dealing with siloed databases to issues with mistyped data entries [105] - and faces a number of important social criticisms. Police departments are increasingly working with new data sources including social media, which leads to new questions concerning the potential for mass surveillance [106, 103]. Additionally as the collected data is inherently not objective in nature, researchers have noted that data-driven policing has the potential to perpetuate biases rather than mitigate them [104]. With human trafficking cases, there is the concern that data collection at a national scale will reinforce stereotypes (such as victims of sex trafficking are more likely to be foreign nationals despite domestic victims being more common), and lead to continued problems with victim identification [107]. Additionally, researchers warn that unchecked data collection could infringe on the rights of individuals [8, 101]. With the development of modern tools, careful attention must be paid to ensure that such tools do not cause harm. Notably, data about voluntary sex workers is intermixed with data about trafficking victims. The rights of voluntary sex workers remains a controversial issue; however, as we will see, our law enforcement informants have little to no interest in arresting them (and often actively help them, connecting them with social services). As human trafficking investigations often rely on big data practices, our work examines the role of technology noting issues with current implementations. Additionally, we discuss potential strategies to mitigate these issues drawing from lessons from prior work that notes that efforts require a mix of social and technical solutions [104, 107].

4.3 Methods

For this study, we performed semi-structured interviews with 16 law enforcement personnel working on human trafficking cases. Our participants worked in a variety of roles within law enforcement agencies including analyst, detective, and senior personnel (see Table **??**).

Table 4.1: Overview of study participants. Note that we interviewed one Canadian participant. Police departments in Canada and the US are largely comparable and the investigation process described by our Canadian participant matched the process described by our US participants.

	Sex	Job Description	Region	Department Type	Recruitment
P1	M	Detective in a VICE unit	Southeast	Mid-size county police department	Personal Contact
P2	M	Detective in a Tech unit	Southeast	Mid-size county police department	Referal
P3	M	Detective in an SVU unit	Southeast	Large city police department	Personal Contact
P4	M	Analyst	Midwest	Smaller city police department	Mailing List
P5	F	Senior Analyst	West	Large city police department	Mailing List
P6	F	Analyst	Southeast	Large city police department	Mailing List
P7	M	Retired lieutenant, specializes in Human Trafficking investigations	West	large city police department	Referral
P8	M	Retired US Marshall	Southeast	Federal police department	Personal Contact
P9	F	Detective in a Human Trafficking Unit	Southwest	Large city police department	NGO contact list
P10	M	Human Trafficking Investigator	Canada	Mid-sized municipal department	NGO contact list
P11	M	Analyst	Southwest	State Police organization	NGO contact list
P12	M	Retired Detective, current director of a related non-profit	Southwest	Large city police department	Referal
P13	F	Investigator specializing in Human Trafficking	West	Large county police department	NGO contact list
P14	M	Analyst specializing in Human Trafficking	Southwest	State police organization	Referal
P15	M	Retired Major, currently Coordinator for Human Trafficking	Southeast	State police department	NGO contact list
P16	M	Human Trafficking Investigator	Southwest	State police department	NGO contact list

Drawing methods from prior work [108], we used convenience sampling [109] to recruit initial participants and used snowball sampling to gather additional participants. We recruited participants for the initial 3 interviews through personal and professional contacts. Additionally, 3 were recruited through a relevant mailing list, and the remaining were recruited through a random sample of a contact list provided by an NGO partner (see Table 4.1 for complete breakdown). We continued to recruit participants for this study until we felt that the data was saturated [110].

The interviews were conducted over the phone or in person, lasting between 30-90 minutes. During the interviews, the participants were asked questions about their background, investigation process, technology use, and collaboration. We also asked participants to explain their technological needs and how researchers could design better tools for them. One researcher qualitatively coded the interviews using inductive thematic analysis [111], and throughout the process, emerging themes were discussed by all the authors.

4.3.1 Participant Backgrounds

Our inclusion criteria was that our participants had to have worked at least one human trafficking case prior to the interview. However, all of our participants had worked on human trafficking cases for at least a year before the interview. When we reached out to police departments, we were often directed to speak to the more senior person on a unit because that person had worked more cases and was more familiar with what the overall needs were for their unit. In particular, P7 and P15 had both worked over 10 years on human trafficking and are experts on the topic.

In terms of familiarity with Human Trafficking, all our participants had received some training about trafficking and were largely familiar with recent policy and research on trafficking. Participants who specialized in trafficking received more training than their other counterparts and our participants who worked as analysts typically had the least training. Training for human trafficking tends to happen after a person joins or is assigned to a human trafficking case. Officers tend to move between units frequently mostly as a mechanism to gain a promotion (i.e., a person might join the VICE unit because there was an opening for a detective). However, our participants noted that units that deal with human trafficking cases tended to self-select more (i.e., officers choose to join those units specifically). As such, we find that our participants are self-motivated to investigate these cases.

Expertise and comfort with technology varied between participants, with all being at least familiar with the majority of the tools discussed in the findings section and highly proficient with using the case management tools and police databases. More experienced participants (P2, P9, P15) used the more complicated visual analysis tools. How experienced a person was with technology depended on how much training they had access to. Our participants described getting training for the various tools through a mix of formal trainings provided by the tool's manufacturer and informal trainings where a colleague trained them. Access to sufficient training was mentioned by our participants as a serious barrier to technology adoption.

4.3.2 Limitations

Our participants are not representative of law enforcement in general. We limited our study to include only those who have worked on human trafficking cases. Officers with experience with human trafficking cases tend to work in larger police departments and tend to be more specialized than might be otherwise be typical. Additionally, our participants have access to more technology than is typical for police departments. As mentioned in the related works section, only 10% of police departments have access to data mining tools; whereas the majority of our participants describe having access to some tools in this space. Like other qualitative studies on law enforcement working in human trafficking [108], our participants were mostly from the southeast and southwest regions. Many states in these areas have policies that list human trafficking as a high priority and have a long history of participating in human trafficking task forces. As such, we did not interview participants in departments who have only recently begun to investigate human trafficking cases.

4.4 Findings

In this section, we first describe the overall investigation process our participants described. It is important to understand how an investigation progresses in order to contextualize the role technology and collaboration play in this process. We then discuss participants' technology use and collaboration process. Finally, we present our main findings, the sociotechnical needs of our participants.

4.4.1 Overview of Investigation Process

At a high level, most investigations as described by our participants go through the same general steps (see Figure Figure 4.1). Note that these stages don't always go in this order

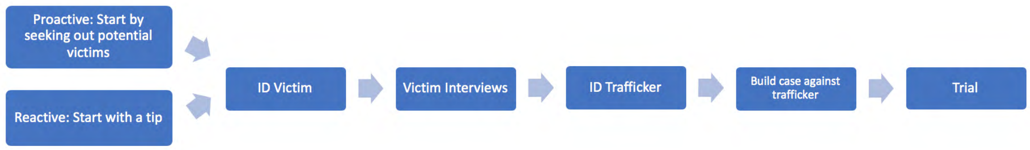

Figure 4.1: Overview of general investigation process

and some investigations perform these stages in parallel. Also investigators often work multiple cases at once. In terms of timeline for this process, our participants noted that typical investigations take a long time - ranging from several months to several years. Human trafficking investigations are often more complex and take more time than other criminal investigations [91].

Starting a New Investigation

There are two main processes through which an investigator starts a case: proactive and reactive [21]. The proactive process starts with the investigator actively seeking out new cases. The reactive process starts with a case that has already been identified by someone else. Seven of our participants noted doing a mixture of both processes (with P1 and P15 working proactive cases almost exclusively) and nine mentioned working reactively only.

In a typical proactive process, an investigator starts a case by examining online advertisements for sex work to look for indicators of trafficking. Investigators use tools such as Marinus Analytics' Traffic Jam[1] and Thorn's Spotlight[2] to assist with this process. These tools provide users with a searchable database of these advertisements and an interface to explore connections across selected advertisements. P16 describes this process noting that he uses his experience to identify trafficking indicators: *"I can literally go get on Traffic Jam right now, find a girl[3] or find a posting that looks like she looks kinda young or she looks like she might not be doing this on her own or something like that."* With the

[1] http://www.marinusanalytics.com/
[2] https://www.wearethorn.org/spotlight/
[3] Our participant is referring to a child victim. However, law enforcement do sometimes refer to adult female sex workers as "girls."

rise in adoption of internet sources to facilitate trafficking, combating human trafficking is increasingly a big data problem [8] and investigators rely on technology to support this proactive process. Much of the software designed specifically for human trafficking cases targets this investigation stage.

In contrast, reactive cases rely far less on technology and instead begin with collaboration. In a reactive process, an investigation starts with some form of a tip from established collaborators, civilians, or other police departments. As P9 describes cases can begin from many different sources: *"It's all different 'cause we get tips from so many areas, we might have a report from patrol, we might have a civilian tip, we might have an anonymous tip."*

All of our participants noted working with a wide range of individuals as collaborators. For example, several participants mentioned working with human trafficking hotlines to get tips for cases. Others mentioned working with government agencies like Child Protective Services or non-profits like the National Center for Missing and Exploited Children. Additionally, a few participants mentioned working with the medical community.

Tips may also come from civilians, sometimes anonymously through programs like Crime Stoppers that allow members of a community to pass along information about a crime through a neutral party. Some departments use civilian informants such as cab drivers to get tips on recent trafficking-related crimes. For example, P1 described starting many of his cases based on tips generated by his local community:

"There are some cases that will get assigned to us. There are certain reports that they'll say "Hey, we want you guys to take a look at this." Or they'll assign us what would be referred to as a suspicious activity report. Then go out there, a citizen saw this type of activity. I'll go out there and take a look at it and see if there's anything to it. We'll conduct our own independent investigation and we'll go from there."

Social media also plays a key role in gathering civilians tips as noted by P5 : *"Social media's driving a lot of things, anywhere from litter, almost like a crime stoppers real time sort of. So you see a lot more crime tips and stuff come in off Twitter and Instagram and*

Facebook."

More often, however, our participants described scenarios where a case was identified by a different police department investigating other crimes. For example, P3 described an investigation that began when a patrol officer responded to a noise violation and accidentally uncovered a trafficking ring. Other participants such as P13 and P8 work in request-based agencies - meaning they take on or assist with cases when requested to by other departments. Other departments may request help if they either do not have the experience or do not have the resources to investigate human trafficking cases. P13 describes this process below, noting that other officers will call her when they need help investigating a suspected trafficking case:

"They'll call us, hey I came out to a hotel. Initially came out as a domestic violence. It kind of looks maybe more like human trafficking. I need some help. I need some direction. We'll go out and help them ... Or someone just on scene right then and there and they call us, like hey I need help. I need some direction. I'm not really sure what I have. We'll go out there and assess."

These cases highlight an important issue with reactive cases: victim identification is largely dependent on collaborators and other police officers recognizing the signs of trafficking. These cases highlight the need for more pervasive training for law enforcement beyond those who investigate human trafficking cases and the need for more widespread public awareness of the realities of trafficking. Additionally, as supported by the findings of prior work, these findings point towards the importance for police communication with the public [101, 100].

Identifying a Potential Victim

At this stage, investigators begin a process of evidence gathering with the goal of getting enough information to make contact with the victim and to access the situation the person is in. Investigators are looking for information concerning the victim's real identity, location, and/or contact information. Additionally, investigators look for information that may confirm the victim is in a human trafficking situation. P9 describes this process noting that she uses several databases to find this information: *"We try to figure out as much as we can from our databases before we go out in the field to do anything so we're trying to use Traffic Jam to find ads, we are on our portals ... and so on to try to research addresses and names to just try to build up the best idea of what's going on that we can."*

This process involves using a large number of tools at once to look for this information across disparate data sources. A typical process for this is described in detail by two participants:

"We will look on our database to find reports for our victim. There may be a report out there that they don't know about, or they forgot about. We'd look on the computer, on police computer databases for those reports ... Then another option would be through Traffic Jam. We would put in whatever information we had about the person. For example, their phone number, punch that in and then get the list of all the locations that they had been working in. Usually you'll find that they're transient, and moving about. Who knows where they were, but then you could follow up on that based on the information from Traffic Jam. Maybe then trying to get hotel records, or video surveillance." - P10

"I'll take that phone number, I'll search it in Facebook, I'll search it in Google, see if it comes back to any place or any thing. I'll do the reverse image search of her image on Traffic Jam to see if she posted it on other websites or what other phone numbers she's posted under. I'll then check it against Facebook or Instagram or anything like that ... Some people put their phone number on their Facebook page. I've had a girl to use the same number and post on Backpage, so that was a super easy find. I just plugged it in there and we found her. Then we actually found out who she was and her actual name, and then we can use our ... law enforcement databases to search their name, date of birth. We can find addresses and things like that." - P16

Note the sheer number of tools both participants described using. This process requires the investigators to keep track of a number of new leads and past search histories at once. Our participants all emphasized the importance of keeping case information organized, and some mentioned using tools like note-taking systems or Microsoft Excel to keep track of case information. Others mentioned relying on memory to keep track of a case's progress.

Victim Interviews

Once a victim is identified, investigators will try to make contact with the person and get an interview. Victim interviews are critical for an investigation because often the information in the interviews is the basis for the entire case at a trial. P2, a former SVU detective that worked on cases relating to sexual assault including trafficking, noted the importance of interviews: *"Special victims, particularly dealing with children, a lot of the crimes you won't find out about until afterwards, there is no physical evidence, so everything relies on our interviews. The majority of our cases, it comes down to talking to the victims and talking to the suspect and it's purely interviews."*

Additionally, our participants noted that they needed to perform multiple interviews with the victim before they can fully understand the situation. This requires investigators to locate a victim multiple times before they can help the person. P9 describes the difficulty with this below:

"They lie a lot and there's reasons for that, there's a common bond, and psychological issues and so on but it can be real difficult to get to the truth of what they're saying especially often it will take three or four interviews before you kind of get the whole story of what went on and in between that time, they'll often have run away or disappeared so you never quite know what part of what they told you know was the total truth and what wasn't."

Many participants mentioned that the first goal of the interviewing process is to get the person the help they need. They work with victim services or victim advocates as part of this process to ensure the needs of the victims are met. P13 describes below that before a formal interview, she first makes sure the person's basic needs are met:

"First and foremost, we don't care about the situation in a sense. We don't want to know the story. We're not going to grill you. Do you want to eat? When was the last time you ate? Are you hungry? Are you cold? We have backpacks that we carry in our cars and it comes with clothes, toiletries and socks of different sizes. We want to meet their needs number one. Medical care. Just show care and concern." - P13

Building rapport with the victim is absolutely crucial - not just for getting the necessary information for a case but to also ensure that the person gets the help they need when they are ready for it. For example, P13 notes that being non-judgmental across multiple interviews is critical: *"But just being there for them really consistently and always and not judging them and listening to them and you know taking them back again, again no matter how many times they've run away and really showing that you care for them is a big part."* P10 additionally notes that building rapport is important for getting the person the help they need when they are ready: *"Often times the person doesn't want to come forward at the time. What you do is develop some rapport, and sometimes they'll come back to you a month later and say, "I met you a couple months ago. You probably don't remember me." Most of the time you do remember them, but say, "Yeah. At the time ... Yeah, I want to tell you about a situation where things aren't okay." Then we would investigate it."*

Identifying the Trafficker

From the information gathered in the victim interviews, the next stage requires investigators to identify the trafficker. Our participants often relied on subpoenas for the victim's phone and social media to do this. As P11 described, investigators look at information in conversations or look at who owns the phone to identify the trafficker: *"We can tell they're in control because of conversions they're having where, a lot of times the trafficker's directing what the traffic person should do."*

Using a similar evidence-gathering process as described in the earlier stages, investigators search to uncover the trafficker's name and address. Some participants noted that they often had to be resourceful to get this information. P15, for example, noted that he sometimes uses his connection with a local restaurant to get the address of a trafficker: *"If I'm looking for somebody and I think they're local, I'll call a friend of mine that works at a pizza place and see if they get food delivered to them, because generally they give you the right address"*. This process requires investigators to keep track of the case's progress and organize the different aspects of the case. This is even more challenging at this stage because investigators are working with more information from multiple sources. As some of theses sources do not come from technology (like interview data), investigators often rely on physical systems like sticky notes and folders to keep track of a case's progress.

Building a Case Against the Trafficker

Once the trafficker is identified, then the investigator must build a case against that trafficker and issue arrest warrants. This process requires corroborating information from the victim interviews and evidence gathered in earlier stages. In contrast to the previous evidence-gathering stages, this process focuses on proving the links between the evidence so as to build a strong case that convinces a jury. For example, investigators like P3 might gather evidence like hotel receipts or cash payments records to prove the trafficker is behind the operation:

"The goal is to try and corroborate some of the things that we are hearing ... let's say we've got someone who is running an organization dealing with human trafficking, well they're keeping tabs on their clients and the money and all of that. We would like to take a look at that to prove that this is their organization, that they're running it. This is who we're seeing. Telephone numbers that match some of the victims that we've recovered. That we're seeing cash payments that match some of the payments that received on this particular day ... Everything has to correlate."- P3

Additionally, prosecutors in collaboration with investigators decide what charges can be brought against the trafficker. The charges will depend on what can be proven and how likely the victim is to go to trial. The process of a trial can be traumatizing for the victims; for this reason, some of our participants have recently pursued other charges besides human trafficking to avoid having to rely on victim testimony. P12 noted this shift in his own investigations where he now typically pursues a financial charge rather than a traditional human trafficking charge:

"We're going after [a financial case]. And then at that point we don't need the testimony as much, which they tend to back out on, for a lot of different reasons. So we're focused on following the money, and see if we can find who's in control of that, and go after money laundering charges, have good strong case on that." - P12

4.4.2 Tools Used During an Investigation

As discussed above, our participants use a wide range of tools during an investigation. Broadly speaking, the tools described by our participants fall into the following categories: police databases, human-trafficking-specific tools, visual-analysis tools, case-organization tools, and general websites.

Police Databases give investigators access to a searchable database for government records, criminal records, or prior case records. Investigators have access to multiple databases at once. Most of our participants named three or more databases when describing their investigation process. Investigators use these databases to find information such as name, address, and criminal records on a victim or trafficker. Other tools mentioned by our participants in this area include LexisNexis's Accurint and Kaseware that allow for more advanced searches on their databases and some basic visualizations of connections between data points returned by a search. Most participants who used either of these tools noted that as part of the searching process, they had to verify the results using other data sources, noting that these databases can often be inaccurate.

Human-Trafficking-Specific Tools mentioned by our participants include Thorn's Spotlight and Marinus Analyitics' Traffic Jam. Both tools provide a web-based platform for investigators to use to search through records of sex-work advertisements. Additionally, these tools allow investigators to run facial recognition search on images in the advertisements. In an investigation, these tools are predominantly used either to proactively begin a new case or to help investigators gather evidence on a victim's background.

Visual Analysis Tools most often mentioned by our participants were IBM's i2 Analyst Notebook, ArcGIS, and Microsoft's PowerPoint. Investigators use i2 Analyst Notebook to visualize social networks for a case. For report writing and quick visualizations for court hearings, many participants mentioned using PowerPoint. Only two of our participants mentioned using ArcGIS, a geographic information system for creating maps, which they mostly use to visualize heatmaps of general crime trends for reports or their department's web-page. Investigators used visualizations to explain their insights into a case to co-workers and prosecutors working with them. Others mentioned using visualizations to explain complicated aspects of a case in court like link charts to explain the connections between different call records (an example of such a visualization can be seen in Figure Figure 4.2).

Case Organization Tools are used by investigators to organize their case notes, keep track of connections, and track the overall progress of their investigation. These tools include a mix of physical systems (like note taking systems, sticky notes, and printouts in binders) and software tools. Software tools again include software used to create link charts. Some also used Microsoft Excel to create checklists to keep track of their progress. As investigators work multiple cases at once, our participants noted that it was absolutely important that they were organized about each case so that leads are not forgotten. Many of our participants described using mostly physical systems to organize their case notes.

General Websites that investigators use include social media sites and search engines. Investigators use these sites to uncover a victim's identity and examine their social network to identify their trafficker. The majority of our participants noted that they searched across

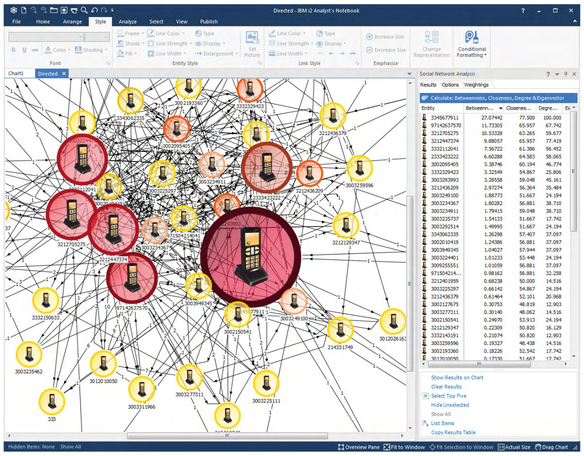

Figure 4.2: Example screen shot of a link chart made with IBM's i2 Analyst Notebook. Image courtesy of IBM

social media frequently to uncover new leads.

4.4.3 Collaboration During an Investigation

Because victims of human trafficking travel across multiple jurisdictions, collaboration between police departments is essential[7]. Furthermore, multi-sector collaboration between law enforcement, service providers, and other disciplines is critical for the identification and support for victims [21]. In regular reports, the US government emphasizes the importance of collaboration for combating human trafficking [18, 14, 112].

All our participants echoed these sentiments in our interviews. As mentioned above, our participants all described working with individuals and groups across multiple disciplines. In this section, we describe the roles of each of these collaborators with respect to an investigation.

In the start of a case, collaborators play the critical role of identifying the potential victims of trafficking. These collaborators include health professionals, social services, governmental agencies, other local police departments, civilians, and even hotel employees. P7 noted the importance of building relationships with collaborators as the primary mechanism for identifying new cases: *"you have to have local partnerships. You all have*

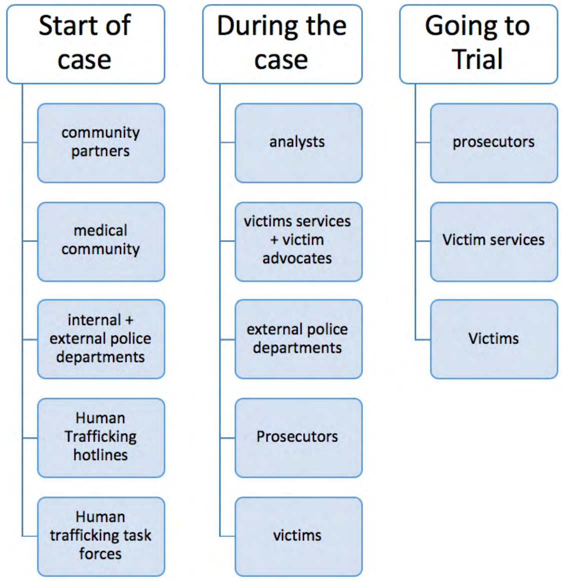

Figure 4.3: Examples of collaborators and their roles in the investigation process

to know how to work together and be working together and have protocols in place to do this work. Because if you have that, then the cases will present themselves".

Additionally, collaborators assist with other stages of the investigation. Analysts, for example, assist with the evidence-gathering and case-building process acting as support for the investigator. Analysts can perform some of the data analysis and visualizations tasks and assist investigators by providing distilled insights from their analysis in reports. Victim advocates assist with the interviewing process - ensuring that the needs of the victim are met. Law enforcement also regularly collaborate with various victim service providers which they rely on to help meet the victim's needs as P10 notes:

"With Victim Services, they provide more resources for the victim than we would. They're a partner ... let's say a victim comes forward and they don't have a phone for example. They had it taken away by the trafficker, or they don't have enough money because they lost the money through the crime, then Victim Services can arrange to get them a phone. There's a whole lot of other resources that they provide, but that'd be an example of something that they could do. They have some funding for tattoo removal as well ... branding is sometimes used by traffickers to indicate that somebody's working for them ...

Getting that branding or tattoo removed or covered, can sometimes be of huge value for the victim."

Throughout the entire investigation, several participants noted that they had to work with other police departments because victims move frequently across different jurisdictions. Other police departments can share information relating to the investigator's case or even agree to work the case together.

4.4.4 Sociotechnological Needs

Tools to Support Collaboration

We see two main problems with collaboration that a socio-technological solution could improve: the problem of a lack of shared data and the problem in identifying key partnerships especially across state lines.

Many participants noted that a lack of shared data between law enforcement departments is a major issue in solving human trafficking cases. Because trafficking victims often move between jurisdictions, multiple departments will accidentally work the same case, thus duplicating efforts. Compounding this issue is the fact that different departments have access to different information and tools. Consequently, some cases can only be solved when those departments collaborate as P9 describes when discussing the main barrier in her cases: *"It's incredible the information is out there that we have no access to just 'cause there's no automated sharing so it's kind of ridiculous ... I feel that could probably solve most of the crime in the world if we had access to all this information that everyone has. Probably would solve a lot of these a lot faster"*.

Technology can assist with this problem in a number of different ways. One solution could be to build a central repository that consolidates case information for human trafficking cases as described by P15: *"The type of technology that would really be needed is the type to have a central repository that is searchable, where everybody could put information in and then there would be deeper access levels for people."* Another solution is to add features to existing tools that let users know who else is looking at the same information. For example, when an investigator searches in a database for a particular record, the system could display the name and contact information for the other officers who searched for that record.

Related to this issue is the problem in identifying key partnerships. In particular, our participants noted having trouble figuring out who would be best to collaborate with across state lines. Currently, most partnerships are established through word-of-mouth and interpersonal relationships. However, victims may move to areas where the investigator has no contacts, leading to issues with continuing the investigation.

"If you have a personal contact, that's always going to be better of course because it's just the personal relationship there that makes it more likely for you to get the action done. If not, then you have to refer to different directories or go to an agency's website and find the phone number of a particular division within the department and hopefully get the right person on the phone to pitch the information to." - P14

Technology can help mitigate this problem in a number of ways. Developers can create platforms such as apps or online forums that support communications between law enforcement personnel working trafficking cases. Recently, some solutions have come out such as the app BLUE[4] for communicating within a law enforcement department. However, none of these solutions are universally adopted yet. Additionally, researchers can develop methods to find common pathways victims travel between areas to discover connected regions. Uncovering these regions will help law enforcement and policy makers form inter-state task forces to tackle those cases. Similar methods have been applied in the past to tackle drug trafficking in the US through the High Intensity Drug Trafficking Areas (HIDTA) program.

It is important to note that while technology has the potential to assist with maintaining and forming these connections, for any collaboration to be effective, departments and policy makers should continue to prioritize collaboration and information sharing.

[4]http://bluetheapp.com/

Tools to Support Case Building and Organization

Many of our participants noted using physical systems to keep track of an investigation such as sticky note systems to organize key information or note-taking strategies to track what has been discovered. With investigators working multiple complex cases at the same time, it can be hard to keep track of all the relevant information for a particular case with these existing systems. Additionally, some investigators rely on memory to to recall connections to previous cases which means that there is the potential for investigators to forget key links. P16 describes using physical systems like notepads to keep track of multiple ongoing investigations:

"Sometimes I grab a notepad and write it down as I go so I don't forget, but I think it becomes the toughest when let's say you needed all that information for court. Like you needed a printout of their Facebook, you needed the law enforcement database search for that person, you needed the Backpage ad. Well that's when it becomes kind of confusing. You just want to make sure you have all of the right information, and also that you're not printing out excessive information that's not relevant."

Investigators need tools to help conceptualize their case-building process including the process of organizing the results of their investigations and keeping track of the case's progress. Some software exists for this purpose. Notably, participants mentioned creating link charts in IBM's i2 Analyst Notebook or in Kaseware to organize case information. However, participants found the process of creating link charts was too cumbersome for routine investigations and instead only used the software for unusually large cases. Further, prior work has noted that existing software still struggles with displaying uncertain, incomplete, or dynamic data [113], which is problematic as human trafficking data tends to have all three of those attributes.

Borrowing from other techniques developed by the information visualization community, researchers can develop new systems to support investigators. As mentioned above, there are some notable challenges that make this space an active area for future research.

Tools to Support Pattern Identification and Forecasting in Geospatial-Temporal Data

As we have seen, human trafficking victims often move across multiple locations. As part of the evidence gathering process, investigators need to map out where a person has been at what times for a number of reasons including proving travel in court. P16 comments:

"With that map, that is super useful for us to be able to go to court and say, "Hey, this is where they were posting at. This is the time, the day, and then if you look on the timestamps, two days later they were posting in a different city." So, for us that kind of proves the travel. We can get a better federal charge if we can show that they were traveling traffickers." - P16

Many participants mentioned working with geospatial and temporal data as part of their investigation but lacked tools to visualize and identify patterns in their data. While some had access to tools like ArcGIS for this purpose, the data the investigators have access to is not in a format that can easily be ported into ArcGIS. Most of their data comes from online advertisements, so the locations are often vague (i.e. an ad might list a region or state) and the times uncertain (i.e. ads can be posted in advance of traveling to a location).

Additionally, participants wanted tools to help them understand the victim's movement patterns. Two main reasons were given for this. First, investigators need to to identify where the victim is headed so that they can alert other law enforcement officers who can intervene as P11 describes below: *"The issue that we have is figuring out where these girls are going to be, what times they're gonna be at. How long they're there. Figuring out patterns of the activity, see if there's any discernible patterns so we can try and get a hold of law enforcement in some other state to help us interdict the person, so that we can then follow back."* - P11

Second, investigators mentioned needing such a tool to discover the identity of victims based on known travel information. P14 describes this need when discussing what kinds of cases he struggles to investigate:

"So, it's just about a layered attack, basically, it's trying to find the ads. I don't have a

phone number, I don't know her alias, so I can't queue by phone. I can't queue by alias. I might have one city that I know she was at but there's just too many ads in that city so it's just impossible. Facial recognition doesn't work because she's using fake images. So what else can we do? I think that's the next step in terms of trying to find ads for someone who you know is probably out there and using the geography and the time frames to do that." - P14

Tools that Unify Existing Software and Methods

Our participants all noted using a large number of unconnected tools at the same time as part of their evidence gathering processes. Because all these tools are disconnected, investigators end up spending a lot of time keeping track of the connections across each of the tools.

"It can be time consuming. It can definitely become kind of a hassle ... I'm not exaggerating about this, but normally when I'm searching for ads I have Thorn's Spotlight pulled up, I have Traffic Jam pulled up, I have my undercover Facebook page pulled up. Sometimes I'll have the AdultFriendFinder because somehow everyone ends up there in some kind of way. And then sometimes I'll have the actual website like SkipTheGames, CityVibe, you know whatever they're using ... that's like five different programs to find out who one person is. It can kind of overlap and get confusing and time consuming." - P16

Investigators have a real need for their tools to integrate with one another. Many of our participants note that they would like a platform that unifies these tools or at least connects the results from one tool into another.

"For me, I forget a lot. Like I'll look at a phone number or name and then I'll go to the next tab over, a few tabs over, and I'll have no clue what I just looked at ... if there was a way we could like kind of combine a couple programs like SocialNet[5] along with Traffic Jam to where some of those searches could be done automatically that would be beneficial.

[5]SocialNet is a tool that helps investigators visualize a person's social network based on social media relationships

It would save a lot of time too as far as I'm concerned." - P16

4.5 Discussion

4.5.1 Design Challenges for Building Solutions for Law Enforcement

When designing new tools for law enforcement there are a few notable challenges to consider. Police officers may need to explain the results of an algorithm or how a tool works in court. When designing tools for law enforcement, it is important to choose algorithms that are human interpretable and design visualizations that help officers get an intuition for how the process works. Recent work in [114] and [115] demonstrate some techniques for making complex machine learning models more interpretable by users. Additionally, training, funding, and experience with advanced technology is unequally distributed across police departments. Tools to support law enforcement have to support both tech-savy users and users with limited computer skills. To the greatest extent possible, we should design tools that don't need a lot of training to use. Finally as part of the case building process, investigators have to verify the results returned by a tool. Therefore, we need to build interactions in the tool that allow the investigator to verify and manipulate the results.

4.5.2 Information Visualization Implications

Designing information visualization systems for law enforcement presents a number of hard challenges including visualizing uncertain data and building systems that law enforcement can trust. This application area pushes the current boundaries of information visualization research making this a new area for current practitioners to explore.

Challenges with regards to uncertainty can arise through the data itself, or from analytic models computing potential outcomes. For instance, the data collected range from semi-structured databases to coded transcriptions of interviews. Merging these multiple data types is challenging, in part due to the variability in data quality given the source. Further, cases often leverage a significant amount of latent domain expertise of investigators. This poses challenges to the visualization community about how to integrate this knowledge into more traditional datasets being visualized.

Many visualization tools rely on computational or analytic models to compute more complex relationships among data. However, many such models have an inherent amount of uncertainty associated with the results generated. Often, this uncertainty increases when the quality (or perceived accuracy) of the data decreases. Specific to law enforcement, this may lead to significantly high uncertainty. For visualization, this raises challenges not only in how to visualize the potential uncertainty, but also how to create interactive interfaces that allow investigators to make well-informed decisions.

4.5.3 Policy Implications

Many of the problems that our participants shared with us are not hard to solve from a research perspective. The officers we spoke with used a patchwork of miscellaneous tools mostly not designed for their needs, and not designed to inter-operate. When one of our participants needed a professional contact in the Las Vegas area, he stood up in the squad room and yelled, ”Anyone know someone in Vegas?” Law enforcement could be more effective if their computing and communications infrastructure was brought up to the standard for corporations for ten or even twenty years ago. Consequently, many of the barriers between law enforcement and more effective attempts to rescue human trafficking victims are a policy issue–the lack of resources for development of quality information technology tools. As discussed in our related works section, few police departments in the US have access for data mining and information sharing. Without these tools, many stages of the investigation process described in our findings would be extremely difficult if not impossible to complete. In light of these results, we urge policy makers and practitioners at all levels of government to consider these issues and to continue to prioritize collaboration and software development.

4.5.4 Privacy Implications

Our research subjects would like better, unified data solutions–national or better yet international databases to help them rescue victims of human trafficking. While this would indeed help them in their critical work, it has troubling privacy implications. Such a comprehensive database can invade the privacy of law-abiding citizens. In the case of sex trafficking, we also have the delicate issue of the rights of voluntary sex workers, whose data will inevitably be represented in those databases.

The rights of voluntary sex workers is a complex issue. Prostitution is indeed a crime in the US[6]. However, many non-governmental organizations, including Amnesty International, advocate for the right of voluntary sex workers to pursue their chosen profession. Some efforts to stop human trafficking have made voluntary sex workers less safe [**burns˙2018**]. Reasonable people can disagree on how to approach the issue of the rights of voluntary sex workers. However, our law enforcement participants are unanimous in their lack of interest in prosecuting voluntary sex workers, and their desire to help those workers who need assistance.

The current byzantine state of information databases to help combat human trafficking, paradoxically, protects privacy. If it is so much harder to do anything, then it is harder to violate people's rights. This then raises a key research issue in the growing field of usable privacy and security: *how can we design databases that help law enforcement be more effective, while preserving citizens' privacy rights?* Incoherent software systems do indeed provide some privacy protection by accident; however, as the research discipline of usable security and privacy evolves, we should be able to provide better privacy protection by design.

Prior abuses of databases such as the National Crime Information Center (NCIC) raise the concerns for creating new unified databases or augmenting existing ones. While the NCIC has been instrumental in solving crimes such as the assassination of Martin Luther

[6] With the exception of regulated brothels in certain Nevada counties

King, the NCIC has also been used inappropriately by police officers to stalk former romantic partners, get phone numbers of romantic interests, and even look up information on reporters who wrote unflattering articles about their department [116]. While many of these officers were disciplined and in some cases prosecuted, abuses still persist and there are no clear methods to track how often such abuses happen [116]. It is clear from these cases that simply keeping records of a person's searches is not enough to safeguard against abuse. With the collection of sensitive and personally identifiable information, external oversight should be required to minimize cases of abuse and investigate suspicious search histories. Additionally, we urge policy makers to identify procedures and regulations to ensure that appropriate safeguards are put in place to protect against abuse.

Additionally, large-scale data collection concerning human trafficking investigations needs to address serious concerns about data protection. Researchers have continued to highlight the need for better data on human trafficking, especially as current efforts and policies to combat human trafficking often have to rely on anecdotal evidence [3, 107, 92]. However, human trafficking victims and survivors are an especially vulnerable population. Data breaches concerning human trafficking data can lead to further victimization and carries the risk of endangering the lives of victims and survivors [107]. Careful attention must be paid to ensure that data-collection efforts do not cause more harm than good. Best practices should be taken to ensure any solutions are privacy preserving and that the rights of the individuals are not infringed upon. Work by Mark Latonero [8] and by Felicity Gerry, Julia Muraszkiewicz, and Niovi Vavoula [107] provides guidelines. More research on best practices is needed.

4.5.5 Design Limitation Implications

Finally, while our research subjects had many real computational needs for their work, we must also consider what the limitations should be considered for any future development. In particular, while some of our participants expressed an interest in using facial recognition, we argue that given the potentials for harm facial recognition poses that such tools should not be developed until more research has been done to understand and circumvent the harms. We discuss this point in more detail in the next chapter.

As discussed in the section above, while many officers want larger and more unified databases, the process of both collecting and storing such data is rife with concerning security and privacy implications. Future work must contend with the often competing needs by law enforcement for more information with the rights of individuals for privacy and their need for less surveillance. Thus we argue that until more research is done to both understand and provide viable solutions that are rights preserving, such unified databases should not be implemented in practice. Further if and when such databases are put into practice, we argue for the necessary inclusion of proactive auditing systems. Without proactive audits, it would be unlikely to know if such a system is being misused until it is too late.

4.6 Conclusion

We began this research with a key question: how can HCI researchers help combat the problem of human trafficking? In this initial qualitative investigation, we have uncovered a host of questions that the research community can address. Some challenges are policy oriented: why don't professionals charged with a task as important as stopping modern-day slavery have better tools? If you imagine giving law enforcement the kind of IT support that corporations use for everyday business tasks, tremendous progress could be made. Other challenges push the limits of the state of the art of usable security and privacy, and information visualization. These findings are a road map for our research group's efforts to build tools to assist law enforcement in rescuing victims of human trafficking.

CHAPTER 5

ETHICAL TENSIONS IN COMPUTING APPLICATIONS FOR ADDRESSING HUMAN TRAFFICKING – A HUMAN RIGHTS PERSPECTIVE

5.1 Introduction

Over the last two decades, interest in applying computational methods including artificial intelligence (AI) to combat human trafficking has increased – driven in part by increased awareness of the many ways in which human trafficking operations use online technology. This work incorporates methods like machine learning, computer vision, and social network analysis in a variety of anti-trafficking contexts including automatic victim identification [117, 118, 119, 120, 121, 122, 123, 124, 125, 126, 32], evaluating policy outcomes [127, 128], identifying organization networks [129, 130], assisting with forensic investigations [131, 132, 133, 134, 135], and automatically detecting grooming behavior and child exploitation [136, 137, 138, 139, 140].

These approaches are increasingly being used in real investigations and are embedded in commercially available tools. For example, Amazon Rekognition, one of the most popular computer vision services, has been used by law enforcement since 2017 for cases relating to human trafficking [141]. This software was even credited as specifically instrumental in the 2019 arrest and prosecution of a trafficker [142]. While this work has the potential to save lives, prevent exploitation, and ultimately help end slavery worldwide, this research also faces a number of criticisms concerning potential dangers. A number of civil rights groups have noted major concerns that this work could further cause harm to marginalized communities, increase unwarranted governmental surveillance capacities, and violate privacy rights of citizens.

As these computing systems become more widely adopted and are therefore fundamentally embedded into our social systems, we must examine the implications of their development particularly with respect to their potential to cause harm. Additionally, more of the computing research community is starting to get involved in these efforts. In March of 2020, the Computing Community Consortium (CCC) hosted a workshop to create a research road map for efforts to apply AI to combating Human Trafficking. It is critical that we examine best practices as well as safeguards and strategies for harm prevention. By examining these efforts through the lens of human rights we can begin to uncover the ethical tensions associated with this research and highlight new avenues for future research centered on ethics. The human rights perspective uses the principles enshrined in international law by governing agencies like the UN, and this perspective affords a view that is centered on the individuals affected over technological advancements. Further, this framework uses a shared universal language which allows for cross-disciplinary understanding of our analysis.

The goal of this work is to examine the current state of the art and understand the ethical tensions surrounding applications of computational methods and AI in anti-trafficking research through a Systematic Literature Review. We use principles from human rights law as a lens to examine the prior work and guide our recommendations for future work. Additionally, this work represents a unique case study for examining the ethics of applied computational methods in domains that intersect the criminal justice systems and vulnerable populations.

This work presents an overview of the ethical issues present in applied computational methods research for human trafficking. We discuss these issues and gaps in work across eight human rights principles – Privacy, Accountability, Safety & Security, Transparency & Explainability, Fairness & Non-Discrimination, Human Control of Technology, Professional Responsibility, and Promotion of Human Values – and contextualize our analysis by drawing from prior work on ethics and AI. We conclude by presenting a series of ”calls to action” for future research to address the ethical issues we found in our analysis. We call for future work to include a broader use of participatory design; engage with other forms of trafficking; develop best practices for harm prevention; and include transparent ethics disclosures in their research. Our calls to action further offer insight as a case study on how to ethically apply computational methods to sensitive domains.

5.2 Background and Definitions

To better contextualize our work, we first define the term "Artifical Intelligence" or AI so as to explicitly lay out the scope of this work. Then we give an overview of Human Rights-Based Approaches (sometimes referred to as HRAs) and situate the specific human rights framework we use for our analysis within the broader field of HRAs.

5.2.1 Defining AI

Defining Artificial Intelligence (AI) has long been a contentious issue among researchers and practitioners [143, 144]. The field of AI is evolving at an unprecedented rate and prior definitions have not kept up-to-date with modern conceptualizations of AI. To this day there remains no consensus on a universal definition of AI [143].

To accommodate the many possible conceptualizations of AI, we take a broad and descriptive definition of AI that encompasses a wide range of sub-fields including machine learning, visual analytics, and data science. Papers included in our analysis are not limited to those that are explicitly labeled AI, but instead involved either A) some sort of machine-assisted processing of data and/or B) the creation of a computing system that assists with decision making. Our goal with taking such a broad view of AI was to ensure that our work included a wide range of methods and approaches. Our definition of AI is heavily inspired by definitions listed in AI governance policy documents from AccessNow [145] and the European Commission's High Level AI Expert Group [146] and only excludes methods that are exclusively qualitative or exclusively hardware oriented.

5.2.2 Overview of Human Rights-Based Approaches (HRAs)

Human-Rights based approaches use the values and principles that underlie international law to guide and inform development; evaluate current research efforts; and guide public policy discussions [147]. These principles are derived from internationally recognized legal frameworks, such as the United Nation's Declaration of Human Rights (UNDHR) [148], and include fundamental principles like dignity, fairness, liberty, and equality. As a framework, human-rights approaches take a people-centered approach to analysis by drawing attention to the rights of those affected by development and the responsibility for researchers to examine practices for preserving human rights.

In this work, we contextualize our analysis using the 8 human rights principles – *Privacy, Accountability, Safety & Security, Transparency & Explainability, Fairness & Non-Discrimination, Human Control of Technology, Professional Responsibility, and Promotion of Human Values* – outlined by legal scholars in [144] and described in the list below. This work derived its principles by examining shared themes found in over 20 documents for AI governance put forth by civil society, governments, inter-governmental organizations, and the private sector. As this framework included a broad geographic representation not centered specifically on American or European laws, these 8 principles represent an international perspective to rights-based approaches for AI development.

- *Privacy* - AI systems should respect a person's right to privacy both in the data collection process and in the design and deployment of an AI system. People have the right to control the use of data about them and AI systems should respect their agency in how their data is used.

- *Accountability* - This principle is concerned with mechanisms for determining responsibility for the AI systems and ensuring said responsibility is appropriately distributed. Researchers must consider the impacts of an AI system and provide adequate remedies for those impacts.

- *Safety & Security* - AI systems need to reliably perform as expected without causing harm and not be vulnerable to attacks.

- *Transparency & Explainability* - Systems need to be designed for oversight through accessible and transparent design and through human understandable outputs.

- *Fairness & Non-discrimination* - Systems must be designed to prevent algorithmic bias and discrimination such as through developing representative datasets and through designing systems to be inclusive.

- *Human Control of Technology* - AI systems in critical decision making contexts must be subject to human review.

- *Professional Responsibility* - Researchers must ethically design AI systems and consider the long-term consequences of their designs. This principle emphasizes the values of accuracy, responsible design, multi-stakeholder collaboration, and scientific integrity.

- *Promotion of Human Values* - Systems should reflect core values and promote human well-being. AI should used to benefit society and not to promote harmful practices, discrimination, and unequal conditions.

5.3 Related Work

There has been a long history examining the ethics of computing research across multiple domains, with more recent efforts focusing on applications of AI [149]. Prior work has, for example, examined the ethical tensions that arise from using large-scale social media datasets to train computational models [150, 151, 152, 153], examined how AI systems embed existing discriminatory values and biases [154, 155], and analyzed the human impact of large scale data collection such as issues of data privacy and algorithmic surveillance [156, 153, 157]. Others, such as the work in [158, 159, 158], further examine the role of

AI within criminal justice settings and note that AI systems are fundamentally changing decision making processes.

Alongside these efforts, scholars have also looked to develop principles and frameworks to use as a critical lens to tease apart the ethical tensions present in AI research. In particular, there has been an increased attention on using the language of human rights as a framework to guide AI development. Thus far much of this work has focused at the macro-level – examining governing and policy practice for AI development [144] – but recently authors have called for using this framework to guide academic work as well [147, 160, 161, 162].

Several authors have noted that using human rights has several advantages as a framework over traditional ethics-based frameworks [147, 160, 163]. The human rights framework is based on internationally recognized legal standards and thus affords a shared, universal understanding. The principles that underlie this framework tend to be more concrete and center the analysis on the rights of the individuals affected by technology [161, 160]. This in turn forces the analysis to focus on people first over technological considerations – much in the same way as Value Sensitive Design and Participatory Design [147, 162]. Additionally, there is a natural overlap between human rights and other ethics-based frameworks [147]. The principles [144] we use in our analysis includes the 3 represented by the FAT model (Fairness, Accountability, Transparency). In practice, rights-based approaches have been used by scholars and policy experts to recommend best practices for preserving data privacy during pandemic response [164], highlight ethical considerations when using big data in healthcare settings [165], and to access harm cause by algorithms used in key decision-making contexts [163].

However, human-rights based approaches are not without their limitations. Using human rights as a lens tends to focus more on a macro-level and thus can miss nuance found at smaller scales [166]. The language of human rights is taken from legal documents and public policy briefs and thus presents barriers to participation and understanding especially outside the legal field [167]. Human-rights-based policies often feature unclear enforcement mechanisms for systems that span multiple states and involve non-state participants – a scenario found more often than not in AI development [167]. Human-rights-based policies further tend to feature a heavy emphasis on legal mechanisms to combat inequality and discrimination – which in some circumstances has a positive effect, but in others can secure political power for certain groups at the expense of others [168]. Finally, many of the values – including those used in this paper – heavily feature industry perspectives including from organizations with histories of problematic AI research. Many researchers and policy makers have expressed skepticism on principles derived and implemented by industry participants who may be motivated to incorporate human rights and other ethics-based principles for the purposes of "ethics washing" their brand [169, 170]. Thus, prior works has suggested that applications of human rights be careful not to de-emphasize the principles of accountability, transparency, and participation [169].

Finally, it should be noted that human rights and ethics are complementary frameworks. There is a real benefit to synthesizing these perspectives. To this end, while we explicitly highlight a human-rights lens in our analysis, our discussion and interpretation of these principles is informed by our experiences with ethics frameworks. We particularly tried to emphasize values of equal participation and just design.

5.4 Methods

Our work performs a Systemic Literature Review (SLR) using the methods outlined in [171] and [172]. At a high level, this process involved developing a set of keywords to search for papers, gathering the manuscripts captured by those terms from computing related publication databases, systematically pruning the resulting corpus to remove irrelevant material, and then analyzing the final corpus using a systematic data extraction and synthesis process.

5.4.1 Search and Screening Process

To construct our corpus, we conducted a literature search using a set of human trafficking related keywords to identify potentially relevant work. Because our the primary focus of our search was to identify work specifically at the intersection of AI and Human Trafficking, our search process targeted 3 main computer science publication databases (IEEExplore[1], The ACM Digital Library[2], Springer- Link[3]), and we used the following keywords and their derivatives to perform the queries: ”human trafficking”, ”sex trafficking”, ”labor trafficking”, ”modern slavery”, ”sexual exploitation”, ”labor exploitation”, ”forced prostitution”, ”forced labor”, ”forced marriage”, ”trafficking survivors”. Finally, we limited our search to only English peer-reviewed results that were published after 2000. We impose this time restriction to ensure a consistent definition of human trafficking; note that 2000 is the year when the US passed the TVPA and United Nations passed the UN Trafficking in Persons Protocol (Palermo Protocol). In total, 616 unique results were collected - 213 from the ACM Digital Library, 316 from IEEE Xplore, and 87 from Springer-Link.

We then constructed a practical screen using the criteria outlined in [171] to ensure that our final corpus only contained relevant material. Our screen took the form of a check-list and one of the authors was tasked with systematically classifying the articles as relevant or irrelevant to the survey. Articles are excluded if the answer to any of the questions below was no, otherwise the article was included in the final corpus:

1. Does the study have a clearly stated or heavily implied application for human trafficking?

2. Does the study have methods which involved either fully automated or machine-assisted processing of data? OR Does the study involve the creation or design of some computing tool?

[1] https://ieeexplore.ieee.org/Xplore/home.jsp
[2] https://dl.acm.org/
[3] https://link.springer.com/

Our inclusion and exclusion criteria had the effect of excluding papers whose methods are exclusively qualitative (such as ethnographic studies) and papers exclusively concerning hardware design. Because our emphasis was on examining AI practices, the papers who used these methods fell outside of the scope of our study. We did however include studies that used mixed-methods approaches as long as those studies involved the design or implementation of an AI system. Examples of papers that were excluded because they did not fit our definition of AI included interview studies to understand the socio-technical needs of anti-trafficking non-profits [173, 174], sex worker charities [175], law enforcement working human trafficking cases [13], and Nepalese survivors of sex trafficking [176]. However while these papers were excluded, the analysis and lessons learned from this work helped inform and shape our analysis. In particular, we found that many of these studies had valuable insights for how future research could be developed in a way that centers the needs of marginalized groups and we include these insights within our discussion. We also excluded papers which only incidentally matched our keywords which included papers discussing the "master-slave" computer architecture and papers focusing on automotive accidents. This results in a final corpus of 69 papers.

As part of the screening process, the author also noted if any of the papers were developed by the same team of authors concerning the same AI system. Our goal was to group papers written about the same system and treat those papers as one unit for analysis. However, we found that while two papers in our corpus were written about the same system, these papers described discrete and separate components of a larger system. After careful consideration, we decided to not group these two papers because each of the papers were substantially different enough to consider each of these studies their own AI systems.

5.4.2 Data Extraction and Synthesis

Our process involved two stages of analysis: 1) data extraction and 2) qualitative analysis. We use the results of the data extraction process to summarize the current state of the art and uncover gaps in the existing literature. We use the results of the qualitative analysis to form the taxonomy of ethical tensions described in the discussion section.

For data extraction, we use methods described in [177] and [171] to create a data extraction form to analyze the texts. This form contained qualitative criteria including: 1) Who is the intended user? 2)What is the intended use-case? 3) What type of human trafficking does this paper address? 4) What methods did the paper use? 5) What data did the paper use and where did the data come from? The form also tracked meta-information about the publication including the year of publication, the publication venue, the location of the study, and the institutions of the authors. Once the form was created, one of the authors performed the data extraction and used the ”test-retest” process [177] to ensure consistency. The ”test-retest” process involves the researcher re-extracting data from a random sample of the texts in the corpus and checking for consistency in answers to the form.

Finally, we performed a meta-analysis of the corpus using framework synthesis [172]. Framework synthesis takes a structured approach to qualitative data analysis, and uses an iterative, deductive process where codes are generated against an chosen framework [172]. For our analysis, we use the 8 human rights principles – *Privacy, Accountability, Safety & Security, Transparency & Explainability, Fairness & Non-Discrimination, Human Control of Technology, Professional Responsibility, and Promotion of Human Values* – described in [144] as our framework. The authors iteratively generated codes from the text based on relevance and/or absence of discussion relating to 8 human rights principles. As mentioned earlier, the goal of this analysis was to uncover themes around the ethical tensions noted in the papers as well as tensions not present in the papers. During this synthesis process, the authors also took detailed notes about the any ethical concerns or areas where future research is warranted. Results from the framework synthesis process were then used to narrative describe the existing research space and the structure the discussion of the ethical tensions.

5.4.3 Limitations

We note the following limitations in our methods choices. First due to practical considerations, our approach only uses one reviewer to judge the inclusion/exclusion of papers in the final corpus. While other papers utilize the same approach [178], using only 1 reviewer does increase the risk of bias and human error affecting the results our analysis.

The range of search criteria (both keywords and time frame) was limited to maintain a manageable set of papers to consider for the final corpus. Thus, some relevant papers may have been excluded that did not meet this criteria. In addition, because we limited our keywords to only those directly referencing human trafficking, we may have missed papers that might be relevant to human trafficking but do not explicitly mention human trafficking within the text of the publication. Finally, because we are only considering academic work, we must consider that some applications of AI for human trafficking exist outside of academic work (for example, in commercially available tools) and are therefore excluded from our analysis.

Finally our analysis uses human-rights as framework to uncover ethical tensions. Thus, our analysis might miss ethical tensions that don't align well with any of the principles. Alternative frameworks such as transformative justice [179] could provide different views on this issue.

5.5 Findings

We note two major trends in publication timing: 1) there was a surge in publications post 2016 and 2) that the popularity of this topic is continuing. At the same time, this remains an understudied issue. Our corpus found only 69 papers published in the last two decades. While human trafficking has gained more awareness over the years, this field remains relatively niche.

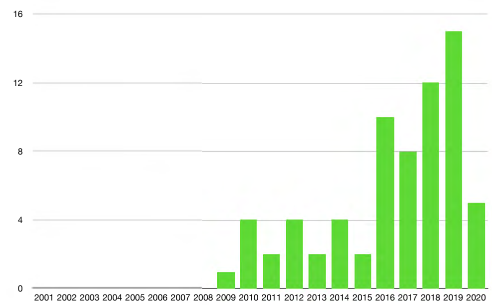

Figure 5.1: Histogram of publications per year, 2020 includes partial data

5.5.1 Problem Areas & Use Cases

We identified five sub-domains of human trafficking that the papers address:

1. *Labor Trafficking* - covers papers exclusively concerning forced labor, labor exploitation, and/or labor trafficking.

2. *Sex Trafficking* - covers papers exclusively concerning forced prostitution, sex trafficking, and/or sexual exploitation but was not specific to only child cases.

3. *Commercial Sexual Exploitation of Children (CSEC), Child Trafficking, and/or Child Pornography* - This category is shortened in Table Table 5.1 as "CSEC" and covers papers targeting victims of sex trafficking that are minors. Note that in our corpus, there were no papers targeting specifically child labor trafficking.

4. *Broadly Human Trafficking* - covers papers whose applications either specifically targeted cases of both labor and sex trafficking or whose methods are intended to apply in circumstances where both labels would apply (such as Cantina workers who are both forced to provide labor for the restaurant and provide sexual services for the customers).

5. *Human Trafficking as one of many possible use-cases* - covers papers where human trafficking was listed as one of many possible use cases. Many of these papers were designed to target other crimes that occur alongside human trafficking such as gang

violence, organized crime, and drug trafficking. This category is shortened in Table Table 5.1 as ”One of Many.”

Results for the distribution of papers across these sub-domains can be seen in Figure Figure 5.2. Combined the sex trafficking and CSEC categories represent roughly 78% of papers in the corpus - meaning that overwhelming majority of papers targeted sex trafficking more broadly. Only 3 papers listed labor trafficking as the primary focus.

Figure 5.2: Bar chart showing the distribution of papers in the corpus across the 5 problem areas. 3 Papers specifically targeted Labor, 28 Sex trafficking, and 26 CSEC. 5 target human trafficking more broadly, and 7 target human trafficking as one of many application areas. Combined, sex trafficking of adults and minors represents the most common focus of papers in the corpus

We also analyzed what the intended use case was for the results and tools developed in the paper. At a high level, these use-cases fit into the following categories (Note that we also included the category of ”unclear” to capture papers who had unclear use-cases):

1. *Aiding Criminal Investigations*
2. *Informing Policy Decisions*
3. *Preventing Human Trafficking and Educating the Public*
4. *Supporting Survivors of Human Trafficking*
5. *Unclear*

The categories we describe above broadly line up with the ”4P Paradigm” - a framework established in both the TVPA and Palermo Protocol that categorizes anti-trafficking efforts [180]. This paradigm classifies anti-trafficking activities into 4 categories: prevention, protection, prosecution, and partnership [181, 180]. Prevention refers to activities centered on educating the general public, identifying vulnerable populations, and providing services

Table 5.1: Table showing the breakdown of papers by intended use case and targeted area of intervention. Categories in this table are not exclusive - meaning papers can be assigned to multiple categories (ex: papers that address both labor and sex trafficking) and can be assigned to multiple use cases (ex: papers that target both aiding criminal investigations and inform policy). The number in parenthesis across the top represents the total number of papers that were labor, sex, ect.

	Labor (3)	Sex (28)	CSEC (26)	Broadly HT (5)	One of Many (7)	Total
Aid Criminal Investigation	1	25	19	1	10	56
Inform Policy	2	2		2	2	8
Prevention		1	8	1		10
Support Survivors				1		1
Unclear					1	1

and vocational alternatives for vulnerable groups. Protection refers to actives centered on rescuing, rehabilitating, and reintegrating survivors of human trafficking. Prosecution refers to activities centered around creating and enforcing anti-trafficking laws. This category also includes efforts to support existing criminal justice practices. Finally, partnership refers to activities centered on identifying and fostering collaboration and information sharing between anti-trafficking groups.

The 4p paradigm is used to evaluate anti-trafficking efforts and highlight areas where more work is needed [3]. We can use this model to see high-level patterns in what researchers have focused on and where future research is needed. Thus, in the sections below we note the overlaps between our findings and the categories in the paradigm.

Aiding Criminal Investigations

This category represents by far the largest use case with roughly 75% of papers (56 total out of 75) explicitly indicating this use case. Papers in this group predominantly dealt with applications of machine learning for automatic victim identification [117, 118, 119, 120, 121, 122, 123, 124, 125, 126, 32] including specifically identifying child pornography [139], characterizing human trafficking activity [118, 182, 183], and detecting illicit behavior [138, 184, 185] or organizations [129, 130]. Using the 4P paradigm, the papers in this category broadly address ”prosecution” as these papers were intended to support criminal justice efforts.

The most common task within this category was victim identification either through developing computational models to detect trafficking online or through developing public reporting systems [186, 187]. Research focusing on computational approaches for automatic victim identification use online sites like social media, dark web forums, and sex work ad sites to train models to detect cases of human trafficking. The goal of this work is to narrow down possible leads for law enforcement by differentiating posts about human trafficking from posts about other topics. There are a variety of approaches used for this task, though most borrow methods from Natural Language Processing (NLP) and machine learning. Many of these approaches rely on expert generated keywords to train models [120, 139]. For example, [119, 121] used supervised learning models trained on keywords and [117, 188, 124] used a mix of keywords and expert labels to both train and evaluate semi-supervised learning models. Beyond keywords and text-based approaches for victim identification, other work in this corpus used metadata present in the posts, images present alongside posts, and even payment data [130] to detect instances of human trafficking. For example, [121, 122, 123, 125] used the location metadata present in Backpage[4] posts alongside keywords to detect instances of sex trafficking. Both [184, 185] use a mixture of NLP and computer vision to detect cases of child sex trafficking and use Convolutional Neural Networks (CNNs)[5] to estimate age and gender in images accompanying suspicious social media posts. Similarly, [189, 190, 32, 191] all use a mix of textual based approaches (like keywords), computer vision approaches, and metadata analysis to detect sex trafficking in online sites. [192, 193] use computer vision to identify distinguishing visual features present in the background of images posted alongside advertisements to geolocate the victims by matching those features to photos of hotel rooms.

[4]Backpage was an online classified site similar to Craigslist that included sections where people could advertise adult services and solicit sex work. However, traffickers have also used this site to advertise their victims [30, 8]. As a result in 2018, the website was seized by the US Department of Justice and was shut down.

[5]CNNs is a type of deep neural network used for image classification. This method falls broadly under the area of "Deep Learning"

Within the broader category of "Aiding Criminal Investigations" there were also papers that focused on the task of detecting trafficking organizations/networks. Many of these papers used similar data sources and methods as the research focused on the task of victim identification. For example, [118] used unsupervised template matching to uncover human trafficking organizations and connections between advertisements posted to Backpage. Further, [130] used patterns in bitcoin payment information to uncover human trafficking organizations. Many of the approaches used various social network analysis (SNAs), Natural Language Processing (NLP) and community detection methods to uncover the groups [129, 183, 126, 194, 195, 196, 197] . The goal of this work is to support law enforcement investigations by allow investigators to uncover connected profiles based on similarities, interactions, and relationships with identified instances of trafficking.

Other papers dealt specifically with identifying grooming behaviors and child exploitation online. These papers used NLP and machine learning methods such as formal concept analysis[136], and various supervised learning approaches [137, 138, 139, 140] to categorize grooming behaviors in chat logs. The goal of this work is to assist law enforcement with assessing threat levels related to sexual abuse of minors.

Other papers in this category related to applications of computer vision models for forensics - including matching identifiable features in sexual abuse imagery [131, 132, 133, 134, 135] and tracking unique characteristics present in other human trafficking related images [192, 193]. Papers [131, 133, 198] all tackled the problem of vein pattern visualization where computer vision techniques are used to locate vein patterns seen on body parts to use for identification purposes. This application is particularly important for identifying the perpetrators in child pornography as faces tend to be obscured [131]. Additionally, to aid in the identification of child pornography and sexual abuse imagery, papers used various deep-learning models to estimate ages [132, 135, 199] and assist with generating age progression photographs [134].

Finally, this category also included papers that designed tools to support law enforcement investigations through building tools to organize leads [200, 201, 202], visualizing case data [203, 190], and improve search capabilities[120, 204, 205, 206, 207, 191, 208].

Informing Policy Decisions

Papers in this category used computational approaches to evaluate current policy and programs. Many papers borrowed methods from operations research and used simulations to evaluate the efficacy of current policies. For example, [127] used system dynamics simulation models[6] to evaluate outcomes based on different labor policy implementations and [128] similarly used system dynamics models but focused on policies targeting sex trafficking. Much of this work relies on data from national reports. For example, [210] constructed transnational flow models using both quantitative and qualitative data within the US State department's annual Trafficking in Persons (TIP) report.

Others like [211] focused on improving risk assessments that identify vulnerable populations. Their goal was to better inform policy and governmental practices for preventing instances of trafficking. Additionally, [182] sought to evaluate the impact of environmental disasters (like hurricanes) on human trafficking patterns and whether these event increase the risk of victimization within already vulnerable populations. The goal is to both better understand the human trafficking ecosystem and to link the potential relationship between climate change and the increased risk for victimization.

Finally, papers in this category included [212, 213] which sought to use machine learning to improve prevalence estimation. For example, [212] specifically targets labor trafficking and uses time series analysis to estimate the scale of human trafficking in particular regions in India. Prevalence estimation is important for policy makers as these estimates provide context for decision-making and resource allocation.

Papers within this use case could align across all the categories in the 4P paradigm depending on the use case. Much of the work for this use case emphasizes evaluating

[6] System dynamics simulations use models to represent cause-and-effect relationships and equations to simulate system behavior. Example models include causal-loop diagrams, and flow diagrams [209]

current policy practices or estimating current population to better inform policy decisions, while placing less of an emphasis on policing or establishing laws. Thus, while some papers mention how their work could improve criminal justice policy and prevent trafficking, these papers perhaps better fit somewhere between prosecution and prevention.

Preventing Human Trafficking and Educating the Public

As the name of this category implies, papers in this category focused on preventing trafficking from occurring and explicitly fell under ”Prevention” within the 4P paradigm. Papers that focused on preventing cases of trafficking tended to focus on either developing computational models to predict high risk individuals [214, 215] or on developing models to detect grooming behavior[7] so that interventions can occur sooner [136, 216, 217, 218]. This is distinct from the risk assessments described in the section above because the papers in this category are specifically intended to improve online moderation rather than inform governmental policy decision makers. Further, these predictions are intended to only be used in online settings and cannot be used to predict risk of victimization outside of an online context.

Additionally, some of the papers in this category further focused on developing systems to educate the general public about internet safety and human trafficking indicators. For example, [219] built a video game to teach children about harms online.

With the exception of [214], all of the papers in this category exclusively focused on preventing the trafficking of children and none focused on preventing labor trafficking.

Supporting Survivors of Human Trafficking

Only one paper fell within this category – [220] built a web-application to support Nepalese survivors needs for peer-bonding and social support. This work corresponds with ”Protection” within the 4P paradigm.

[7] grooming is a term used to describe the manipulative tactics adult sex offenders use to form connections with children for the purposes of future victimization

Unclear

We labeled one paper, [221], as unclear. The authors of [221] built the tool, "Cyber Trafficking Surveillance System (CyTrass)", to analyze discussions of cyber-trafficking related discussion on social media. However, it was unclear how the authors intended this system and its insights to be used in practice.

5.5.2 Intended Users

The overwhelming majority of papers (73%) explicitly list law enforcement as an intended user of the system. Interestingly, very few of these papers mentioned engaging with law enforcement as part of the design and development process despite the focus on law enforcement applications.

Other users mentioned included governments, NGOs & non-profit organizations, policy makers, military, the general public and parents of young children.

5.5.3 Discussion of Ethics

In our analysis, we also analyzed the discussions of ethics and limitations within the paper's text (if present). We took a broad definition of ethics for this analysis which included any discussion of impacts and potential harms and any discussion that considered ethics principles in the research process.

Only 11 papers had explicit mentions of ethics either as a distinct section or embedded within another section. Within these papers, most focused primarily on issues of fairness and bias stemming from datasets. For example, papers [199, 132, 134] assessed the accuracy and bias present in facial recognition software used by law enforcement. These papers all note that many of the datasets that underlie commercially available computer vision software are unbalanced with respect to race, gender, and age; The result is that many of these tools have lower accuracy on faces that are young, feminine, and/or have darker skin [134, 131].

Additionally, only 2 papers explicitly mentioned using an ethics framework to guide their analysis. [222] examined the privacy implications of image classification used to detect child pornography and used the ”Personal Information Protection and Electronic Documents Act” or PIPEDA principles[8] to guide their analysis and research process. Using principles from prior work on bias and AI systems, [126] proposed a bias mitigation plan as an inherent part of their research approach and tool design. Their plan included steps to both diagnose potential sources of bias and steps to mitigate those biases.

To some extent, the overall lack of ethical discussions can be explained by paper length limitations and the differing conventions for paper formats at different computing conferences. While limitation sections are common, distinct ethics sections are not.

5.6 Application of the Human Rights Framework

In this section, we use the human rights principles established in [144] to analyze our data and provide a series of recommendations for future work. This paper seeks to answer what it means to address these 8 principles – Privacy, Accountability, Safety & Security, Transparency & Explainability, Fairness & Non-Discrimination, Human Control of Technology, Professional Responsibility, and Promotion of Human Values – with respect to the development of AI for anti-trafficking. We use this lens because it provides a universal vocabulary, well-developed standards and principles, and a concrete framework for solutions. Note that due to the high degree of overlap in our analysis, we condensed the principles ”Privacy” and ”Safety & Security” into one section.

5.6.1 Privacy, Safety, & Security

Privacy, Safety and Security have long been established in human rights law, and privacy is considered a fundamental human right under the UN’s Declaration of Human Rights (UNDHR) [148]. In the context of AI research, these principles are concerned with the

[8]PIPEDA is federal data privacy legislation in Canada

idea that AI systems should respect an individual's right to privacy – particularity with respect to data privacy, and that AI systems and their underlying datasets should be secure and resistant as much as possible to compromise.

The privacy and security concerns we found in our analysis using these principles stemmed from the datasets that underlie the models – raising issues around the storage, collection, and access to data as well as concerns about the scope of data collection and its impact on marginalized groups. In this section, we will first focus on issues of data collection scope and impact, and then focus on issues relating to long-term data stewardship and security practices. We noted three themes in the gaps in how papers addressed privacy. First through their dataset collection process and tool designs, we noted that several papers advocated either directly or indirectly for mass-scale surveillance and are thus potentially violating peoples' right to privacy. Second, we noted instances of researchers potentially exposing identifiable data either within the paper itself or through the release of non-anonymized public datasets. Finally, we find that this research raises questions about data access and protection.

We found a number of papers whose work – either directly or indirectly – advocated for mass-scale surveillance and used crime prevention as justification. This is a concerning theme to find because some cases of human trafficking have been linked to corruption and exploitation by user groups identified in the corpus [187, 223]. Additionally, surveillance for the purposes of crime prevention historically has also been used for human rights abuses and remains a contentious civil rights debate [224, 225, 226, 227]. While this work represents avenues for social good, we have to examine what limitations should be considered for collecting data and what impact this work will have on different groups. These privacy concerns were especially prominent in models designed to detect instances of human trafficking in online activity. For example, [217] developed a tool for parents that monitors a child's online activity and uses AI to detect instances of grooming behavior. Similarly, [215] developed a system of connected IoT tools that tracks children's movement patterns via GPS to prevent cases of child abduction for the purposes of sexual abuse. While human rights law is murky concerning children's rights to privacy, many of these applications – through the collection and storage of data concerning these children's activities – create potential avenues for harm. This work often ignores the potential for adults to use these systems for abuse and control and also ignores the very real cases where human trafficking is facilitated by a parent or close relative. Based on reports to service providers, between 10 and 30% of human trafficking cases are facilitated by family member [223]. Further from a security perspective, the collection of data on children's activities if compromised might further put vulnerable children at risk for exploitation. Recent cases like IoT baby monitors being hacked by pedophiles have driven security researchers to caution using these tools to track children [228].

We found concerning privacy implications even in papers that didn't target children specifically. This raises concerns over how this kind of data collection could disproportionately impact other vulnerable populations. For example, papers [117, 118, 119, 120, 121, 122, 123, 124, 125, 126, 32] collected data from sex work sites and provided this data to law enforcement and government entities. These datasets contain information on both human trafficking victims and voluntary sex workers; as a result, this data raises a number of privacy concerns beginning with the issue of misidentification. False positives generated by AI systems trained on these datasets could result in increased police and governmental intervention on groups historically affected by police and state violence. Further, these datasets tend to disproportionately contain data on marginalized identities – including BIPOC (Black, Indigenous People of Color) and LGBTQ+ people – who are already over-surveilled and subject to discriminatory practices resulting from increased surveillance [229, 230]. The use of satellite imagery as seen in papers [212] and [213] further sharpens our concerns. While tracking behaviors using satellite imagery have been used in a number of positive circumstances such as holding governments accountable and detecting cases of human rights abuses [231, 232], there are concerns that these systems will be used by nation-states to invasively monitor citizen's behaviors and could be used to systematically target marginalized groups [233].

In addition to data privacy, we as researchers must examine the ways publishing materials can potentially violate the right to privacy. In our corpus, most papers used mockups or synthetic data in examples to protect the privacy of victims. However, we found a few notable examples where researchers used real data within the text of the paper and/or used a public dataset with Personally identifiable information (PII) present. In all these cases, the PII concerned the perpetrators of sexual abuse or human trafficking. We caution against this practice and urge researchers to protect the privacy of all subjects in a dataset including perpetrators, even though the laws surrounding the privacy rights of those accused of crimes is in many cases unclear. This is especially important because even within the corpus how a person is being labeled a trafficker can be unclear (eg is a person labeled a trafficker upon conviction? After being accused? etc.). We further argue against releasing datasets without anonymizing this information and caution against using public datasets who engage in this practice as this has the potential to legitimize bad privacy practices (and in the case of one paper, potentially draw more attention to the social media accounts for suspected terrorists).

Finally, the issue of data access and sharing raises a number of questions concerning privacy. In our corpus, the datasets tended to fit into three categories:

1. Public facing datasets that the authors did not create. For example, many papers use the IMDB-WIKI dataset [234] for training computer vision models for age estimation.

2. self-collected datasets that an author created by scraping public webpages

3. datasets provided to the authors by a third-party.

For public datasets, most papers used existing public-facing datasets and tended to have no data on human trafficking. This was particularly present in computer vision papers concerning age estimation who use human trafficking as a motivating reason to pursue this line of work but didn't exclusively examine human trafficking imagery. Researchers releasing public datasets about trafficking brings up the question of how their anonymization process guarantees privacy.

In cases where researchers collected their own datasets or were granted access to an existing private database, many of these datasets nominally collected no ”identifiable data”. Thus, it is unclear if laws such as GDPR[9] – which underlie most current human right approaches for protecting privacy rights – apply to many of the datasets described in this corpus. This brings up a number of questions concerning long-term data stewardship practices. Who is allowed to view the dataset including after the project is complete? What is the process for guaranteeing that only those individuals will have access to that data? Are there existing methods to proactively check for data-breaches? Long-term who retains control of these databases?

While answering these questions are typically is beyond the scope of most research papers, addressing these questions within publications represents an avenue for researchers to lead by example by demonstrating their principles for privacy protections and data stewardship. Additionally, by not addressing privacy concerns, it becomes difficult for reviewers and readers to evaluate the tool with respect to ethics.

5.6.2 Accountability

Accountability concerns the mechanisms that exist to monitor the impacts of AI systems over time as well as the processes to determine who is responsible when things go wrong. In practice, this principle has been applied to AI research through recommending new policy and regulations, requiring routine impact assessments, and establishing auditing requirements thorough an AI system's life-cycle.

However, conceptualizing accountability with respect to academic work can be chal-

[9] GDPR stands for the ”General Data Protection Regulation” and this the European Union's data privacy and protection law

lenging. In AI governance, Accountability is often discussed with respect to legal risk and responsibility. But this is not often applicable in academic settings because not all of the papers involve a fully deployed AI system. Thus, we instead focus our analysis on what mechanisms exist to address accountability aimed at the pre-deployment stage of AI development. This is most commonly through impact assessments which emphasize authors identifying and documenting potential harms and risks.

As all the papers in our corpus were peer-reviewed, we can assume that some level of this analysis was performed – ensuring a certain level of verifiability and replicability of the work. However, not all peer-review systems explicitly require reviewers to consider potential impacts and risks. Though this is becoming more common with more academic venues requiring authors to include analysis of harms and benefits in their work and to require reviewers to reflect on those points.

We found that documentation on the harms and benefits directly within the text of the papers was also inconsistent across the corpus. As mentioned in our findings section, very few papers explicitly mentioned issues of potential harm or identified risks. Additionally, few authors mention using impact assessment directly within the publications and even fewer authors address issues of long-term evaluations or mechanism for redress. There were some notable exceptions in our corpus however. [187] considered the specific needs of survivors and included specific design considerations to prevent data misuse, while [126] included an impact statement within the paper and emphasized design considerations for bias mitigation and redress.

5.6.3 Transparency and Explainability

Transparency and Explainability are concerned with the ability to understand and evaluate AI systems. Together, the relate to the idea that AI systems need to be designed so that oversight is possible and that technical concepts about an AI system should be translated into human-intelligible formats. Historically, this has been thought of as a binary (i.e.

a system is either fully transparent or fully opaque), but more recent attention has been called to the idea that transparency – throughout the development life-cycle – exists on a continuum. Explainability is concerned with how technical concepts about an AI system can translated into human-intelligible formats.

Academic work can be transparent through open data and code, accessible documentation of the design process and results, and mechanisms to evaluate how the work will impact the general public. As discussed earlier in the privacy section, open data and code is not always a feasible option for research in this field. The need for transparency has to be balanced with the right to privacy and the need to protect investigative procedures. Though, some of the papers in the corpus have suggested dataset alternatives to ensure privacy while also promoting transparency. For example, some authors test their models using datasets that are related to the task but don't contain actual entries of human trafficking. For examples, papers [131] and [133] trained computer vision algorithms to match vein patterns in thighs and arms for forensic investigations of child pornography. Both papers use a dataset that collected images of these body parts from willing participants and not from sexual abuse imagery. Similarly, papers [134, 132, 135, 222] (which sought to improve age estimation of images for detecting cases of sex trafficking) online used public computer vision datasets of faces unrelated to trafficking to train their models. Similarly, work from related fields has suggested the use of "semi-synthetic" datasets as a workaround for this issue [235].

With respect to transparent documentation, all of the papers in the corpus clearly lay out their data collection, design process, and methods within the text. However, evaluating the transparency of this reporting is difficult when there are differing standards across publication venues for the reporting of results, limitations, etc. For example, only 20 papers (roughly one third) included a specific limitations section or the explicit discussion of limitations with another section such as the discussion section. Those that did discuss limitations tended to somewhat narrowly focus on dataset bias. Further, as mentioned in the accountability section, rarely did authors discuss the impact on stakeholders. The degree to which authors address potential harms and impact is somewhat determined by the norms of specific academic communities and page lengths.

Another issue affecting the transparency of this work is that academic papers are not the most accessible to all audiences. Almost all of the papers are hidden behind a paywall which limits who can easily access the information in the paper and evaluate the methods and limitations. Furthermore only 11 of the papers (roughly 15%) included an explicit statement about the availability of the data and/or code relevant to the study (including those that justified not making the data or code available) with only 9 of those papers including links to either the code or dataset. While some of the papers outside of these 11 used publicly available datasets, the authors did not mention this within the text on the paper.

5.6.4 Fairness and Non-Discrimination

Perhaps the most discussed topic within the corpus was the issue of bias. Most papers acknowledged the bias present in datasets, but none discussed bias in the model itself or in the focus of the applications. Discussions of bias are critical in this context as issues of discrimination in Human Trafficking applications can have profound life-or-death implications.

There are a number of areas where bias is potentially introduced into these systems, including bias present in the datasets, the labeling schemas, and the models trained as a result. There may also be bias present in what areas that computing researchers tended to focus on.

Within our corpus, we identified several incidents of bias with respect to gender, age, disability, and socioeconomic status that were present in the datasets that underlie these AI systems. In many cases these biases stems from a lack of sufficient training data. As papers from the corpus note [134, 199], many of the systems – especially those that employ facial recognition or age estimation techniques – lack sufficient training data on darker skin tones and thus are less accurate at identifying faces of BIPOC individuals. This issue of racial bias in computer vision applications is not unique to human trafficking efforts, but in this context these biases result in the development of real-world systems that are less accurate at identifying BIPOC victims and reinforces existing disparities in the criminal justice system. [199] further notes that many of the commercially available facial recognition algorithms used in identifying cases of child sex trafficking are less accurate on younger and female faces. This indicates that these systems – which are currently used today to specifically to identify cases of child sex trafficking – might not work as intended in that exact context. In practice, these biases translate to systems limiting which cases are investigated, which victims are identified, and which cases go to trial all while reinforcing existing racial and gender disparities in the criminal justice system.

It is also unclear how effective many of these computer vision models are on identifying faces with unique facial features such as those with Down's Syndrome, facial scarring, or prominent birthmarks. Given that none of the training datasets explicitly mentioned collecting images of those with disabilities or facial differences, it is likely that the accuracy of these models are limited in this area. Individuals with disabilities are at increased risk of exploitation [223]. Again these systems have the effect of steering investigations towards only certain groups.

Related to this issue of dataset bias, is the issue of gathering valid and representative "ground truth labels" as to a person's status as a victim of human trafficking or relating to a person's likelihood of being trafficked in the future. In this space, the most common application of AI is for automatic identification of human trafficking victims and/or the prediction of the likelihood of a person being trafficked in the future. Ground truth data labels are directly fed into most systems and used to measure the accuracy of models developed. Thus, establishing the validity of these labels is a critical process to research in this space and directly impacts the quality and effectiveness of the resulting systems. We identified a number of challenges in gathering ground truth labels revealed by our corpus. First, there is a general lack of consensus on a clear definition of human trafficking. As a result, some authors conflate instances of voluntary work with human trafficking or conflate certain behaviors as being high risk for trafficking. For example, one paper collected data from an online fetish website and conflated interest in consensual taboo sexual preferences for involuntary human trafficking. Compounding this is the general ”fuzziness” with distinguishing between exploitation and trafficking. As many researchers note, human trafficking exists on a spectrum and drawing a firm line is a complicated task [8]. Even law enforcement are reluctant to make firm judgements on what is or isn't human trafficking - saying that it takes multiple in-person interviews to determine that [13]. Thus, this process of labeling ground truth injects bias into the dataset as these labels will reflect the cultural values of the researchers and experts assisting with labeling. Additionally, the models themselves may learn racist human trafficking indicators was the case with [117] which found that their machine learning model learned that the phrase "Asian" indicates sex trafficking.

There is also a general lack of established ground truth data sets - especially with respect to distinguishing between consensual sex work and human trafficking. There are no standardized techniques to label activity as human trafficking or as high risk of trafficking. Thus, researchers rely on proxy indicators of trafficking such as suspicious key words that are suggested to be indicative of trafficking. Most commonly, researchers have used law enforcement or NGO-generated key-words to label advertisements as likely involving human trafficking as seen in [117, 118, 121, 122, 123, 125, 130, 126, 188]. However, experts often disagree as to how accurate these key-words are at predicting human trafficking cases [8, 236, 237, 33]. As one researcher put it "no researcher or investigator can ascertain with 100% confidence that a particular online advertisement is a positive case of sex trafficking, just as one cannot be completely certain whether an advertisement is a negative case" [8]. To date no research has empirically evaluated the reliability of these keywords.

Finally, as noted earlier in our results section, the overwhelming majority of papers focused on sex trafficking - with a nearly even split of adult and child cases – and focused heavily on building systems for investigative purposes. There is a serious gap in work addressing labor trafficking and work that addresses the needs of survivors and service providers. The principles of fairness and non-discrimination extend beyond algorithmic bias and also include the principles of inclusiveness in design and in impact. Only two papers took into account the needs of survivors [187, 220] and none explicitly included survivors as authors. Only a few papers mentioned working directly with their intended users as part of the design process.

5.6.5 Human Control of Technology

The principle of Human Control of Technology deals with the ethical tensions that arise as a result of shifts in control away from people to AI systems. Both the individuals who use an AI system and those who are impacted by said system should be able to review decisions and remedy any objectionable results. In practice, human control of technology takes on a variety of forms – such as designs that integrate ex post expert reviews, models that incorporate direct human input, and even designs that include models that predict when humans should intervene.

In our corpus, we found a spectrum of how authors included human input/reviews – from fully autonomous designs where expert input was limited to labeling and evaluation tasks, to designs that fully integrated direct human-input. Where a paper fell on this spectrum tended to depend heavily on what computing sub-discipline the authors drew their methods from. Papers using methods from data science and machine learning tended only include human input in dataset labeling and model evaluation tasks, whereas papers who used methods from HCI tended to include more involved human-in-the-loop elements in their designs. However, the overwhelming majority of papers in our corpus limited human input to only dataset labeling and rarely included expert evaluation of the results after the fact. Papers [138, 212, 131, 117, 118, 122, 123, 193, 238] all used experts to label their training data and used these labels to evaluate the accuracy of their model's decisions; but none of these papers included expert evaluation of these decisions afterwards. This gap highlights opportunities for future research to include more avenues for human interaction.

This also brings up the concern that incorporating AI systems is inherently changing human behavior. Rather than human users acting as safeguards, these systems may influence their users to instead change behaviors [158]. Further research is needed to understand how in-situ uses of AI are affecting excising decision making process, especially in the context of criminal justice systems and policy oversight.

5.6.6 Professional Responsibility

This principle emphasizes the personal responsibility researchers have when designing any AI system and argues that researchers should ensure that appropriate stakeholders are included and that long-term impacts are accounted for. In our analysis, we examined this principle with respect to themes mentioned in the framework: 1) responsible design, 2) consideration for long-term effects, and 3) multi-sector collaboration.

Responsible design calls for researchers to directly engage with how AI impacts society and encode values that align with social norms. This highlights the need for researchers to center their design around considerations for potential harms and benefits of their work. This ties into the next principle of "considerations for long-term effects" as researcher should consider the benefits and harms beyond the immediate. Long-term effects of research are difficult to evaluate as academic papers represent snap-shots. However, a somewhat troubling trend is how few papers directly mention long-term plans or the evaluation of future harms. Related to the discussion of long-term effects, is the also question of long-term data stewardship. What happens to these datasets after a research project is complete? What should the long-term data plan be for projects using human trafficking data?

Finally, multi-sector collaboration is a necessity with this research. However beyond labeling and data access, we need to ask what it means to partner with external groups on research projects. These partnerships can be a mechanism to include voices typically excluded in academic settings. But note that these groups are not necessarily bound by the same standards as academics; there are little regulations with how private organization operate, collect data, or use results from research. Additionally, this research could also be unintentionally lending credibility to harmful industries and practices. For example, [205] partnered with an organization comprised of "psychic detectives", a well-known predatory industry.

5.6.7 Promotion of Human Values

This principle is concerned with the idea that AI should be developed and used to promote "human flourishing" and leveraged to benefit society. All of the papers in our corpus intended to do this. The goal of these papers is to use computing to combat human trafficking and support existing anti-trafficking efforts. They also represent avenues to inform the broader public of the issue of human trafficking.

However, given recent controversies surrounding law enforcement use of AI and its ability to cause harm, it is worth examining how these tools might be used in ways that could be detrimental to society. Facial recognition is an interesting case study to examine in this context - especially as many of the papers are either developing systems for facial recognition or use facial recognition as a component within a tool used by law enforcement. We recognize that facial recognition is an important tool for law enforcement; As noted by papers in our corpus, facial recognition has, for example, been used to search for missing persons suspected to be trafficked [134, 135], identify perpetrators in child abuse imagery [131, 133, 132], and forensically link evidence over time [239, 135, 240]. There have been several notable cases where the use of commercially available facial recognition APIs have lead to the arrest of a trafficker [142, 241].

However, these same APIs have also been used by law enforcement and governments in other contexts with little oversight. For example, during the US Black Lives Matter Protests in 2015 and 2020, facial recognition was used by law enforcement to identify and arrest protesters [242]. In China, facial recognition has long been a component of the systematic processes used to surveil and oppress ethnic minorities [224]. Additionally, because of bias present in the models and data, facial recognition has resulted in a number of wrongful arrests caused solely by algorithmic miss-identification [225]. Together these concerns have lead many civil rights organizations to call for the ban of facial recognition used for law enforcement purposes until better safeguards can be put in place [242, 243].

This brings up the concern that the development of AI for human trafficking purposes can cause unintended harm especially towards marginalized communities. More research is needed to understand the impact of AI in criminal justice and governmental settings. This research should further include mechanisms for community input and control with particular attention given to including marginalized communities.

5.7 Discussion

As research in this area continues to grow, we as a research community should consider how this work will affect the world around us. We believe this paper marks a first-step towards examining the ethical tensions present in this line of work and further highlights areas for future research. In this section, we briefly summarize some of the challenges we identified and propose the following calls to action and future research directions.

5.7.1 Broader Use of Participatory Design

Responding to the problem of human trafficking necessarily requires interdisciplinary and multi-sector collaboration. The research community has already responded to this need through partnerships with NGOs, law enforcement, the civil sector as well as through cross-discipline academic research. However as discussed in our analysis, there is a need for researchers to include a broader range of stakeholders in the research process – with special care taken to include those most impacted by the research. AI research on human trafficking often impacts marginalized groups including sex workers, migrant workers, queer identities, and people of color. In addition, data collection has the potential to further harm and stigmatize survivors of human trafficking.

Despite how this work may impact these groups, we found that few papers directly included survivors, sex workers, or migrant workers in the research process. There is a strong need for future work to directly incorporate a broader range of stakeholders within the research process. As the products of AI research can further marginalize already vulnerable groups, future AI development for human trafficking should ensure that survivors are involved in the development and research process. Participation from these groups can take many forms: survivors could be involved to vet ground-truth data used to train these systems and evaluate research outcomes, through processes to evaluate and inform research questions and applications, and through direct inclusion as researchers themselves. Future work could also benefit from action research performed in partnership with survivor groups. Prior work (found both in our corpus [187, 220] and elsewhere[166]) has shown that partnerships with survivor groups helps center survivor voices in the research. In both [187] and [220] the use of focus group discussions and frequent feedback sessions with survivors was found to be helpful for uncovering key values and critical needs that underlie their design choices.

However, how participation is included in practice will have profound effects on future AI development. Participatory research is not without its flaws [244, 245] and we encourage researchers to examine the power structures embedded in participation [166]. Particularly when including survivors in the research process, we must be cognizant of the fact that participation must provide benefits to all participants. Some survivors who have shared their experiences have felt exploited by the researcher and anti-trafficking community [246, 166]. Researchers need to ensure that survivors are not further exploited by the research process and that survivors are recognized and compensated for their expertise and knowledge. The goal with engaging with practices like action research and participatory research should be emancipation and democratization [247]. To this end, we argue that survivors and other marginalized groups should be directly and equally included in the research process – not as subjects or users, but instead as project leaders, researchers, and PhD students. Towards this goal of direct inclusion, future work should also examine what existing barriers are in place that have historically prevented direct and equal participation from these groups. This represents an opportunity for researchers to examine the ways in which technology can further participation and cross-disciplinary research with under-served populations.

5.7.2 Broadening Research to Address Gaps and Engage with Other Forms of Trafficking

As discussed earlier in the "Fairness and Non-Discrimination" section, there are a number of areas that have seen little attention from the research community. Thus far, the research community has disproportionately focused on sex trafficking and in particular child sex trafficking. As a result, there is a distinctive lack of research addressing labor trafficking – despite estimates pointing towards labor trafficking being the most common form of trafficking worldwide [248]. Further within research specifically focusing on sex trafficking, there is a lack of research addressing sex trafficking experiences beyond forced prostitution and pornography. None of the papers we have reviewed have addressed the issue of forced marriage and only one of the papers [221] addressed cybersex trafficking. Additionally, the existing work on sex trafficking tends to focus heavily on women, girls, and trans-feminine victims. Notably, many of the data sources used by researchers exclude men and trans-masculine individuals; thus to date, little research has focused specifically on addressing sex trafficking of men and non-binary individuals. Finally, the existing work has more broadly focused on assisting "prosecution"-aligned efforts and in particular on victim identification rather than addressing the other 3Ps: Protection, Prevention, Partnership.

Together, these gaps highlight the need for future work to examine the larger human trafficking ecosystem. Prior work has noted that labor trafficking is increasingly intersecting with technology which represents avenues for researchers to understand the labor trafficking ecosystem using similar methods as the work done to understand the sex trafficking ecosystem. These gaps also highlight the need for further research on the ways in which technology is changing sex trafficking and mechanisms to prevent internet-facilitated exploitation. We urge researchers to use their platform and audience to draw attention to other forms of trafficking.

In addition to broadening the focus of AI research, we also advocate for further research focusing on the impacts of deploying this research. Particularly as much of the work intersects with the criminal justice system, we urge researchers to examine how these tools are impacting decision-making processes. We call for more research focused on examining the use of AI in-situ perhaps using similar approaches as those used in understanding AI impacts in healthcare, etc [149].

5.7.3 Development of Best Practices for Harm Prevention

Research directed towards human trafficking applications present unique avenues for potential harm. There is the potential to cause harm through the data collection process, through the design choices made by researchers, through misuse of existing tools, and even through harm to the researchers caused from interacting with disturbing materials. To this end, we urge the research community to form a series of best practices targeting harm prevention with a particular focus on data privacy and security practices, long-term data stewardship practices, and transparent design documentation. Careful attention must be paid towards ensuring that research efforts do not cause more harm than good and that the work does not infringe upon the rights of individuals impacted by the research. When developing these best practices, the research community should consult with those most impacted by this research – including survivors, migrant workers, and sex workers – to ensure that these best practices reflect their needs.

In addition, these best practices can draw from guidelines developed for computing research that handles similarly sensitive topics. Borrowing from guidelines developed for computing research that handles sensitive data, future work should include mechanisms in their designs for proactive monitoring, use data sharing agreements that limit access to sensitive data, and include policies to delete datasets after a project is complete [150, 249]. With respect to data privacy, future work should consider using mock datasets alongside other datasets to improve the transparency and accountability of the work while also preserving the privacy of survivors. Finally throughout the design process and within the text of the publication, researchers should document potential harms, benefits, and most importantly strategies for mitigating harm.

5.7.4 Broader Inclusion of Ethics Disclosures in Research and Transparent Discussions of Limitations

Our analysis highlighted the need for a standardized practice of including ethics discussions in academic work and for researchers to include transparent and accessible documentation of terms of use and limitations of their work. The practice of including ethics discussion – particularly an ethics discussion focusing on the potential harms and benefits – would help ensure that researchers consider the impacts of their work and would further help the broader community asses the impact of these AI systems. However, there is a real asymmetry in power and knowledge between those who create these AI systems and those who use and are impacted by said systems. Researchers must be cognizant of this asymmetry in designing and writing their limitations and careful attention must be paid to ensuring that the broader public fully understands the impact of their work.

Researchers should also include similar discussion of limitations and biases for their datasets. Especially for datasets that are shared beyond those who collected it, datasets need to include documentation detailing the data's provenance, limitations for usage, and any known biases. We encourage researchers to consider using tools such as the one provided by the Data Nutrition Project[10] to detail these data limitations in an accessible format [250].

Finally, researchers should consider ways to include the discussion of limitations within the design of the tool itself. Tool interfaces could, for example, include disclaimers written in clear language detailing the limitations of the tool. Additionally drawing from fields like information visualization, future research could look at how can we can design interfaces so that users can clearly understand the limitations of an AI system and the biases of the datasets used [251, 252].

5.7.5 Limitations of AI for Human Trafficking

Human trafficking is inherently linked to inequality and marginalization [253, 223]. Human trafficking is a deeply complex problem that is difficult to research and requires navigating intersecting complex social, economic, cultural, and political structures that facilitate exploitation. At the root of it all, those most vulnerable to exploitation are marginalized groups who have unequal economic opportunities. Thus, dismantling systems of oppression and implementing necessary social programs (like increasing the number of long-term shelters and support for equal education and employment opportunities) will have profound effects towards ending human trafficking. To this end, we must be cognizant of where and when AI should and should not be used. Technology and AI systems should not replace or be used instead of much needed social programs, training, and education. We must be careful as computing researchers not to divert resources from social programs towards technology development.

Further, AI systems should not be marketed as "bias-free" or "objective" tools that overcome the limitations of human decision making. Especially in the context of criminal justice applications, AI systems have been shown time and time again to both reproduce and magnify existing inequalities [254]. Researchers must pay attention to the language they use when describing their tool's intended use cases to watch for descriptions that imply

[10]https://datanutrition.org/

this. As discussed earlier, the limitations of any system must be understood by its users.

5.8 Conclusion

With the availability of new datasets and the increased awareness, applying AI to combat human trafficking is an emerging area for new research. In this paper, we analyzed a corpus on this topic using human rights as a lens to highlight ethical tensions that arise from this work. We further propose five calls to action that highlight avenues for ethically-driven future work. We hope this work provides inspiration for thoughtful future research aimed at tackling human trafficking.

CHAPTER 6

VISUAL ANALYTICS PROTOTYPE FOR UNCOVERING HUMAN TRAFFICKING OPERATIONS USING SPATIOTEMPORAL MATCHING

6.1 Introduction

In this chapter, we present a visual analytics prototype designed to help law enforcement overcome linkage blindness within human trafficking investigations. As described earlier, linkage blindness is a common investigation challenge where officers either fail to recognize that cases are connected or are unable to establish connections between ongoing cases. In our prior work [13], we found that one of the investigation challenges officers frequently encounter is connecting human trafficking cases using existing tools. Officers consistently described struggling to generate leads when the only available information on a case was location based. This investigation roadblock frequently occurs when the primary data source for an investigation lead is from social media meta-data, hotel and taxi receipts, and/or phone records. In these circumstances, officers may not have confirmed phone numbers or full legal names for a victims and thus cannot use existing software for connecting cases because these tools require the officer to know either a name or phone number. As a result, our participants specifically described a need for tools to support the ability to connect cases based on geospatial-temporal behavioral similarity.

In addition, human trafficking cases (especially those that involve criminal networks) often rotate or frequently move victims between locations as a strategy to avoid detection [90, 19]. In these cases with frequent movement, traffickers will often re-visit locations or follow a set pathway to transport their victims – which leaves behind a potentially traceable movement pattern. These may be circuits, which are common routes used by multiple organizations/traffickers, or they may follow organization/trafficker specific routes between owned locations [20]. If we can isolate that pattern of movement, we can identify groups of victims that otherwise are hard to detect using existing strategies. Knowing the information can assist law enforcement in disrupting these networks and empower them to form effective collaborations between agencies.

Thus, given these officers' needs and the established practice within criminal justice of using movement behavior to link cases, we designed a tool that allows officers to search based on similar geospatial movement behavior to provide officers with the much-needed insight to connect cases. Our tool was designed to work within existing investigation practices and potentially help replace slower existing methods that rely on manual processes (such as analyzing movement behavior on a hand-generated map).

In designing our tool, we took a user-centered approach. We gathered design goals and user tasks through interviewing police officers and other anti-trafficking practitioners, performing a human-rights centered systematic literature review, and attending anti-trafficking workshops and meetings. The insights from these activities comprise our prior work described in chapter 4 and chapter 5. To present our prototype, we provide examples using a mix of synthetic data and real data based on a dataset used currently by law enforcement. The synthetic data was generated based on police reports and press releases on two famous human trafficking cases. These two cases were chosen to show the diverse circumstances human trafficking can take and highlight the flexibility of our tool in working with a range of human trafficking investigations.

The major contributions of this work include: 1) the design goals and challenges for supporting human trafficking investigations, 2) a prototype for a visual analytics tool designed to help law enforcement explore movement behavior and query for potential leads using behavior patterns, and 3) example usage scenarios illustrating how the tool can be used to support current human trafficking investigations. Further this work represents a preliminary attempt to put the ethical principles discussed in the previous chapter in practice while also addressing law enforcement needs. Through our example usage scenarios, we find that our concept shows promising utility for augmenting existing investigation practices and overcoming linkage blindness problems.

6.2 Related Work

6.2.1 Trajectory Mining

As noted in the Related Works (chapter 3), there is a wealth of existing literature focused on developing techniques for trajectory mining for a variety of analytics tasks. Our work specifically draws inspiration from movement and group behavioral analysis [255, 57] including approaches for efficient k-nearest neighbor queries [256].

Within this body of work, there are a number of open research challenges that intersect with our work. Notably, existing similarity metrics do not work across multiple scales; thus existing approaches cannot trajectories across different scales of movement, in terms of both geographic and temporal scale [255]. Further, existing techniques for detecting collective movement can only find specific patterns of movement and requires that these patterns are known in advance [41]. Thus in circumstances where the data contains multiple unknown movement patterns of interest (ex: both diverging and converging behavior), existing approaches may not be able to simultaneously detect both patterns. Finally, there is a gap in approaches that inherently incorporate on-going expert input within the detection process. There is a need for approaches designed to work in cases where the knowledge about group membership is an ongoing and interactive process, as is frequently seen in law enforcement contexts [36].

6.2.2 Trajectory Analysis in Criminal Justice

With the rise of data-oriented policing policies, there has been a significant interest for extending spatial-temporal analysis to crime-related data. Much of this work has focused on identifying hot-spots given spatial crime datasets [257, 258]. In addition, there is some promising work using trajectory mining for inferring social networks (in the form of gang membership) from spatial patterns [259].

However, there are ongoing and unique challenges with applying spatial analysis in criminal justice contexts. There is a strong need for these methods to be more flexible in order to support the wide range of investigation tasks [260]. Further because of the ever-changing nature of crime and investigation strategies, these metrics need to be adjustable over time to work with new insights about criminal behavior and criminology theory [260].

Finally, these metrics have to adapt to newer data sources. Existing metrics rely on GPS-tagged trajectories, but increasingly law enforcement is relying on alternative data sources for trajectory data. Roughly 70% of officers reported regularly using social media data in their investigations [261]. In the context of human trafficking investigations, many officers use online data sources as a significant part of the investigation process [13]. However, social media and online data presents some unique challenges for trajectory extraction and analysis. These newer data sources tend to have a wider range of location specificity and uncertainty. Social media data (especially in the context of human trafficking) is rarely geo-coded and the locations provided may be non-specific (ex: location is listed as 'near my house').

6.2.3 Visual Analytics in the Criminal Justice Field

We draw inspiration from prior work that designed systems within this industry. Much of this prior work was mostly designed to analyze data associated with specific crime types (such as traffic violations and serial offender cases [262]) or for supporting administrative tasks like resource allocation and report generation [258]. None, however, were designed specifically for human trafficking investigations; which, as prior work notes [13], has unique data science challenges associated with it. In addition, there exists a number of commercially available tools used by law enforcement for these purposes. Though as established in our prior work [13], many of these tools do not adequately meet their needs for investigating more complex crimes like human trafficking.

6.3 Design Process, Analytical Tasks, and Requirements

We adopted a user-centered approach in our design process, and through a series of design activities learned about the needs of our users, the ethical considerations for tool development, the user tasks and requirements, and a set of design goals that form the basis of our final designs. The design activities included attending human trafficking related workshops and policy meetings[1], informal design sessions with experts[2], and iterative prototype development. This process further helped us gain a more complete understanding of the existing technical ecosystem and policy landscape.

We first conducted an interview study, the results of which are described in chapter 4, to understand the investigation process, understand existing technological gaps, and to gather our initial design requirements. From this interview study, we narrowed down the scope of our project to address the one of the most challenging aspects of human trafficking investigations: solving the problem linkage blindness (especially cross-jurisdictional) when officers primarily have only have geospatial and temporal information as a case lead. We then conducted a systematic literature review centered around ethical analysis to further gather design requirements.

Then we used various design exercises to explore potential visualizations and interaction techniques. We sketched these designs on paper and then used a wireframe tool called Excalidraw. Some of these designs were shared with experts in informal design sessions to gather feedback, narrow down user scenarios, and further situate the design to match the users' investigation processes. These sessions guided the development of our final set of user scenarios, user tasks, and design goals. Finally we constructed the final prototype of our tool.

---

[1] these included attending the Police Executive Research Forum's (PERF) session on Human Trafficking, the Computing Community Consortium (CCC)'s workshop on "Applying AI in the Fight Against Modern Slavery", and the Code 8.7 symposium hosted at the U.N. on "Using Computational Research and Artificial Intelligence to end Modern Slavery".

[2] this included showing early prototypes to local law enforcement and relevant experts working at NGOs for feedback. In addition, we presented our analysis and early stage prototypes at the American Society of Criminology Meeting in 2019 and 2020 to get feedback and advice from other criminal justice researchers.

To demonstrate its utility in practice, we also include detailed usage scenarios based on real human trafficking cases that follows the investigation process described by law enforcement. These usage scenarios provide detailed descriptions of how our design could be used in practice and are intended to communicate the utility and promise of our design. Note that while these scenarios draw on real investigations and use data inspired by real cases, they are not descriptions of an in-situ uses of the tool. No user evaluations were done for this tool. Our prototype represents an early phase concept and is not a complete artifact that has been deployed.

6.3.1 User Tasks and Analytics Questions

Through the design process described above, we find that our tool should support 3 main user tasks:

Task 1: Forming Case Profiles – Generating "profiles" of case data is crucial to the crime investigation process [39]. A profile in this context refers to how an investigator identifies key behavior characteristics, crime signatures, and geographic movement patterns about a particular lead. In the specific context our our work, this tasks has users characterize patterns (or lack of patterns) present in the behavior seen in case data or groups of case data. Leads in this context begins with a phone number but a lead could also include other case data, a priori knowledge about a case, or even the trajectory of an identified individual. Given a case lead, the profile generated might include identified patterns/signatures in posting behavior, geographic profile, and unique characteristics present in temporal profile. This task asks the analytics questions: "given a phone number, what does a person's travel path look like?", "are there patterns within that person's movement and posting behavior?", and "are those patterns stable, or are they changing/diverging over time".

Task 2: Identify Potentially Connected Cases – The process of identifying connected cases requires investigators to examine movement behavior and profiles to use to search for connections between cases. Investigators use internally generated heuristics to match cases such as looking for patterns indicating potential coordination between individuals or other indicators of group dynamics [39]. This task asks the analytics questions: “who else shares a similar posting signature and/or geographic profile?”, “who else moved similarly to that person?, and “who else is connected to that person?”.

Task 3: Iteratively Define Heuristics for Matching Cases – As part of the process for connecting cases, investigators have to define what constitutes similar cases, behaviors, and profiles for their specific investigation. Further investigators have to define how these similarities indicate group and coordinated behavior. This process requires navigating challenges such as identifying connections even as crime behavior and profile of a particular individual changes over time. In addition because investigations evolve and develop over time, the process of defining a heuristic is necessarily iterative. This task asks the analytics questions: “what do I mean by similar/connected?”, “what metrics should I use to find signs of group behavior?”, and “how can I use my a priori knowledge and expert intuition about the case to narrow down potential connections?”.

6.3.2 Design Goals

Based on the 3 identified users tasks, we developed the following design goals:

DG1: Discover Patterns in Posting Behavior – The tool should allow users to form a profile for an individual through uncovering patterns/trends/signatures in that individual’s movement and post history. For example, users should be able to use this tool to 1) identify movement patterns such as cyclical movement between locations and frequent travel to unusual locations, 2) to identify posting behaviors such as those indicative of a trafficker posting on behalf of the victim and those indicative of a person using a bot to post ads, and 3) to identify typical behavior as well as any sudden changes or deviations from typical behavior. The tool should also support fast judgements to determine if a phone number is legitimate or spam.

Additionally, the tool should allow users to identify trends in the individual's movement with the goal of identifying other police departments who may also be investigating the same individual. If an investigator sees that a individual frequently visits particular locations, then it is likely that the individual may have already been investigated by a different police department. Thus, visualizing locations may help reduce overlap in work and help with "linkage blindness" caused by inter-jurisdictional data-loss.

Further as noted in prior literature[39], the process of generating profiles for individuals helps investigators connect information present outside of the tool and dataset. These profiles often help jog an investigators memory about prior cases and alternative leads. They also help investigators identify best policing strategies for targeting a particular case and help the investigator suggest alternative approaches and information sources for their investigation [39].

DG2: Discover Coordinated/Group Behavior – The tool should allow users to uncover indicators of coordinated or group behavior by allowing users to search for similar profiles. For example, the tool should be able to find individuals whose movement and posting profiles frequently overlap. Frequent posting to the same locations at the same time may indicate that those individuals are in contact with each other and/or traveling together. This tool should also be able to find coordinated behavior that occurred over time such as cases where two individuals profiles only overlap for a short period of time.

DG3: Support Flexible Definitions of Coordination –

The tool should allow for flexible definitions of similarity. Because human trafficking cases vary and are investigated by different units, the tool needs to be flexible enough to account for these differences. For example, in one feedback session, one officer informed us that in child cases, he looks for people who run away together, essentially looking for a "converging pattern" where posts are spaced out similarly. Whereas a different officer, who investigates cases tied to organized crime, was more concerned with overlapping locations rather than order of posts. Criminal organizations might own several locations that they move victims between, but not always together and not always in the same order; so this officer looks for patterns with a high degree of overlap between locations but does not care about the order of locations.

DG4: Support Different Starting Points to an Investigation –

From our discussion with law enforcement, we found that there were two common starting points with cases our tool was designed to support. Either the investigator had a phone number or they had a list of locations and times from evidence (such as hotel receipts). Our tool should support both of these starting points.

DG5: Support Fast and Flexible Searching –

The search process needs to be fast enough that the officers can find possible leads within minutes to use in an interview; thus, the tool should support fast judgement of group dynamics and fast judgement of phone numbers that only incidentally match. The searching process should also allow users to fluidly and iteratively switch between views of individuals and views of groups to support hypothesis formation and interactive refinement.

DG6: Intuitive and Simple Design –

The tool must be easy-to-use and support high level understanding of the models by users with little background in data science. Further because of the diverse range of backgrounds and familiarity with technology found in our user group, our tool should support a diverse range of possible users.

6.4 Datasets Overview

Our tool uses two datasets. The first is a real-world dataset scraped from a website with known human trafficking activity. The second consists of mock-data constructed from human trafficking case studies and is designed to highlight the full range of our tool's capabilities.

6.4.1 CityXGuide Dataset

Our primary dataset was constructed by scraping the website CityXGuide. Similar to Backpage and Craigslist, this website allows users to post classified style advertisements for specifically sex work and has also been used by traffickers to advertise their victims [263].

Ads on this website consist of three main components: Title, free-text body, and structured information about the sex worker being advertised. Structured information always includes a contact phone number and the date and time when the ad was posted to the website. Users post these ads to city specific boards, but within the structured data can provide more specific location information. Board names are formatted on this site as city name, state-name, country to avoid confusion between cities with the identical name such as Greenville, South Carolina and Greenville, Texas. In addition, we limit our data collection to only ads posted to US specific cites so as to match our user's needs.

We use the location information provided by the board name to geocode the locations using the MapQuest API[3]. We found that using the board name over the free-text location information was more accurate when geocoded. The free-text locations were prone to misspellings, slang, or underspecified city information (i.e. providing a common city name with no state information or labeling the location as "near the highway" or "my place") and thus tended to be geocoded incorrectly by the API. However, this means our GPS locations are accurate only up to a city location and should not be used to match within more specific regions like neighborhood or city blocks.

[3] https://developer.mapquest.com/documentation/geocoding-api/

The scraper extracted ads posted between Jan 2018 and Jun 2020. Our data collection process began at the beginning of the pandemic, thus was only minimally affected by lockdown. The final dataset consists of approximately 200,000 ads for 63,903 unique phone numbers.

6.4.2 Mock Dataset

We use case studies developed from insights from our experts and real human trafficking cases to illustrate the effectiveness of our tool. These case studies were constructed using mock data to address the ethical considerations raised by our systematic literature review (see chapter 5). Mock data allows for our tool to be accessible to the general public (addressing the need for accountability and transparency of the design and algorithms) while also protecting the privacy, safety and security of survivors and sex workers.

This dataset consists of trajectories for two case studies chosen to represent common human trafficking operations that involve movement. The first is based on a 2013 human trafficking case [29] and the second is based on Polaris' reports on Cantina cases [25]. In the following subsections I will describe these case studies in more detail and describe the process for generating the mock data for each of these cases.

Interstate Mock Dataset

Data for our first case study was modeled from case data on an interstate human trafficking enterprise that operated along the east coast in the US [29, 21]. This enterprise involved two primary traffickers who operated out of cities within 7 states (Florida, Georgia, Tennessee, South Carolina, North Carolina, Virginia, and Maryland) and used greyhound buses to transport victims between locations. The traffickers frequently relocated their operation along their route to maintain a steady access to potential customers and to avoid detection from law enforcement. In each of the locations, the traffickers would use Backpage.com to advertise their victims and to recruit more victims for their operation. In total, these two

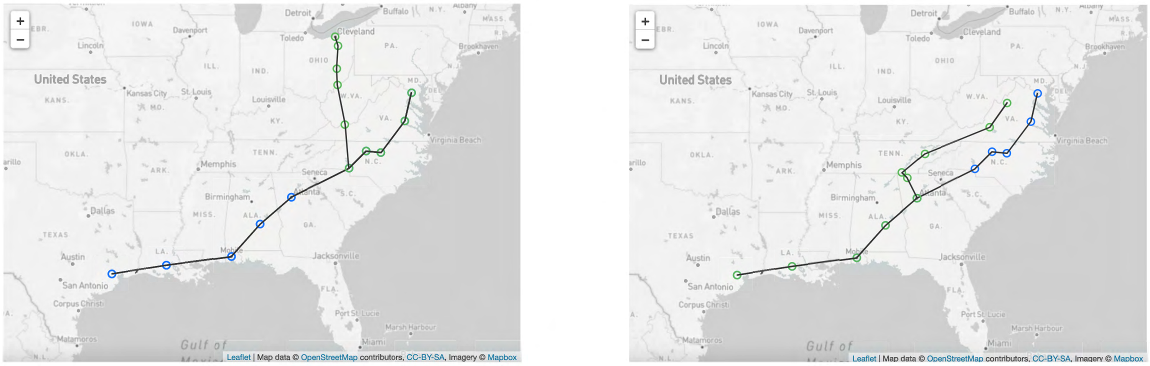

(a) Two trajectories modeling converging behavior

(b) Two trajectories modeling diverging behavior

Figure 6.1: Example of converging and diverging movement behavior patterns modeled in the case study.

traffickers victimized at least 7 minors and 23 adults.

Based on this case data, we construct a mock dataset that shows similar ad posting behavior as was seen in this case. This dataset consists of 12 trajectories encompassing a wide range of group behavior dynamics seen both in the case data and in the literature. These included diverging behavior (where two trajectories diverge overtime, see Figure 6.1b), converging behavior (where two trajectories merge overtime, see Figure 6.1a), fixed and variable lag coincidence (where two trajectories visit similar locations but at different times), reverse ordering, and swapping between phone numbers (where two trajectories represent one person who rotates between phone numbers).

Networked Mock Dataset

Our second case study is modeled on data from a 2016 case involving a network of illicit massage parlors run by an organized crime operation [25]. These illicit massage parlors operated under the guise of legitimacy by masquerading as a massage therapy business. However unlike with legitimate massage business, the workers are forced through economic and social pressure to provide sexual services to customers. Owners of illicit massage parlors will typically use a mix of debts and fraud to lure victims into working at these parlors and then use increasing debts to coerce them into remaining in the industry [264].

This particular organized-crime operation lured victims from Mexico to work in their massage parlors primarily within the state of California. Victims were rotated between businesses per customer demand. Sometimes victims were moved between locations as a group, other times they were moved alone. As a result, this case study represents a more complicated dataset to find similar behavior patterns within. Trajectories in this dataset were longer overall and included different start and end dates. These operations also tend to operate in major cities because of of customer availability. As a result, the return results tend to include more spurious connections. Further because victims aren't always moved as a group, percent overlap as a similarity score might not be as meaningful. Finally, the order of locations visited might not matter because victims are not moved along a set path but instead are rotated between a series of known locations.

6.5 User Interface Design and Implementation

This section describes the design and functionality of the tool. The final design has three primary views – 1) Inspect View, 2) Search View, and the 3) Analyze View – which correspond to the three user tasks described earlier. The views in the design are separated using a tab layout allowing users to fluidly move between stages (**DG5**).

Our tool is implemented as a web-application using D3.js[4] and Leaflet[5] on the front-end. The data including the mockdata is stored in an instance of PostgreSQL [6].

6.5.1 Database and Processing

We use a relational database model that links individual ads together based on shared phone numbers. This approach allows users to search for all ads advertised using a particular phone number which matches the functionality of existing tools on the market. This design choice would further allow our tool to be integrated into existing infrastructure in the fu-

[4]https://d3js.org/
[5]https://leafletjs.com/
[6]https://www.postgresql.org/

ture. While this is not a specific design goal, we do note that based on feedback with law enforcement that it helpful to consider designs that allow for the ability to unify existing tools and practices.

We define a trajectory in this database as an ordered list of GPS coordinates and timestamps representing the route a person took.

One limitation with this approach is that it assumes that phone numbers correlate one-to-one to a unique identity. However, other work has noted that while phone numbers are a reasonable approach, this assumption is complicated by the fact that ads may be for more than one person at a time and may be written by a different person that the person being advertised [93]. Thus, phone numbers are often still the best links between advertisements [19, 20].

6.5.2 Trajectory Querying and Matching

The primary goal of our tool is to allow users to find individuals who travel similarly. Thus, the key component of our back-end processing is the trajectory matching and querying approaches. To allow for flexibility in defining similarity, our design uses three main trajectory matching algorithms: Longest common subsequence and two variations of set intersection.

The LCS metric returns the length of the maximum subsequence shared between two trajectories. For example, say we have two people who traveled along the following paths: New York –> Chicago –> Atlanta and New York –> Atlanta –> San Francisco. The LCS between these two paths would be New York, Atlanta and so this metric would return 2. Essentially, this metric measures how much overlap there exists between two sequences. Two tracks that have significant overlap are likely similar long-term and may imply that the two individuals are connected in some way [58]. LCS is used when relative ordering matters but not exact temporal matching. LCS also does not require matched sequences to be consecutive, thus this metric also allows for sequence to match with gaps.

The first set intersection metric is used when order does not matter. This metric returns the length of the set of unique locations shared between two trajectories. Using the same example trajectories as above, one person has been to New York, Chicago, and Atlanta and the other has been to New York, Atlanta, and San Francisco. The set intersection based on location is New York and Atlanta, and thus the metric would return 2.

For both this set intersection and LCS, users specify a spatial threshold, ϵ which affects how point wise matches are made. For example, let $L_i = \{a_1, a_2, \cdots a_N\}$ and $L_j\{b_1, b_2, \cdots b_M\}$ be two trajectories. A point, a_i, in L_i is said to be equivalent to a point, b_j in L_j if the distance (as measured using Haversine or great-circle distance) between a_i and b_j is less than or equal to ϵ.

The final set intersection metric is used when exact time matching is required. This metric returns the length of the ordered set of points shared between two trajectories where points match both in time and space. For example, let $L_i = \{a_1, a_2, \cdots a_N\}$ and $L_j\{b_1, b_2, \cdots b_M\}$ be two trajectories and let ϵ be the user specified spacial threshold and λ be the user specified temporal threshold. A point, a_i, in L_i is said to be equivalent to a point, b_j in L_j if and only if the distance between a_i and b_j is less than or equal to ϵ and the time between a_i and b_j is less than or equal to $\pm\lambda$.

Once the user specifies which of the three similarity metrics, the spacial and temporal thresholds, and the total number of matches they want to consider, our tool queries for the top k matches using the k-nearest neighbor algorithm. Nearest Neighbor queries efficiently retrieve trajectories that are closest to either a specified trajectory or a set of specified locations and uses the specified similarity metric to determine closeness.

6.5.3 Inspect View

The inspect view, shown in Figure 6.2, is the landing page and shows the user the trajectory details for a specified phone number. Users specify the phone number using the search box on the right hand side. This view consists of two side-by-side panels labeled "Trajectory" and “Details”. The trajectory panel consists of two viz components, a map and timeline scatterplot, and shows the overview of the trajectory information. The map shows the location information and general ordering for the trajectory and the scatterplot view below gives the user the timeline posting behavior. The details panel is populated when users click or brush over locations on the map or points in the scatter plot. This panel shows the details from the ad text (note that in Figure 6.2 these details are shown with placeholder text to preserve the privacy of survivors).

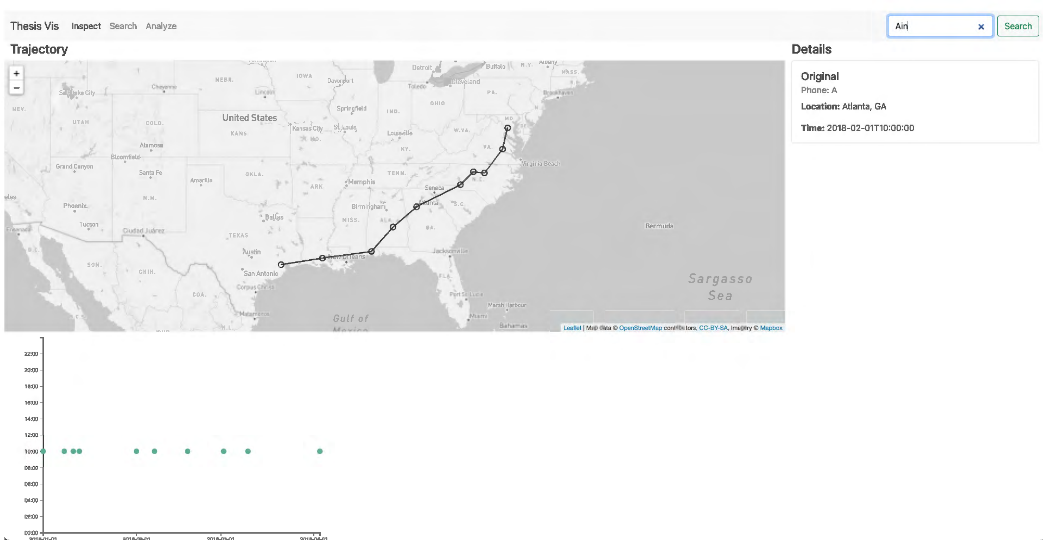

Figure 6.2: the **Inspect View** shows the user the trajectory details for a specified phone number (shown here as ”A”).

The inspect view was designed to help users perform general exploration of an individual’s posting behaviors and movement (**DG1**). Users can use this view to uncover behavior like cyclical movement patterns, gaps in posting history, or even if the ads are for more than one victim. For example, users can use this view to determine if the ads are spam or posted by a bot based on the patterns seen in the frequency of posts and movement. If a phone number was used to post ads to an impossible number of locations at the same time, the user might conclude that a bot must have posted the ads (See Figure 6.3a).

Other examples of movement behaviors users can investigate through this view include determining likely mode of travel. Using the ordering of specific locations and proximity to

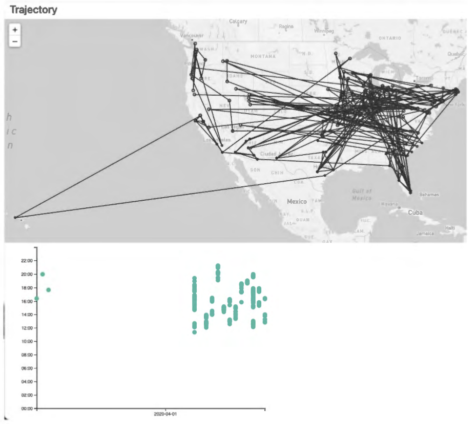

(a) Example of a phone number likely corresponding to spam or bot-assisted posting. Ads were posted to over 150 different unique locations all within a one week period.

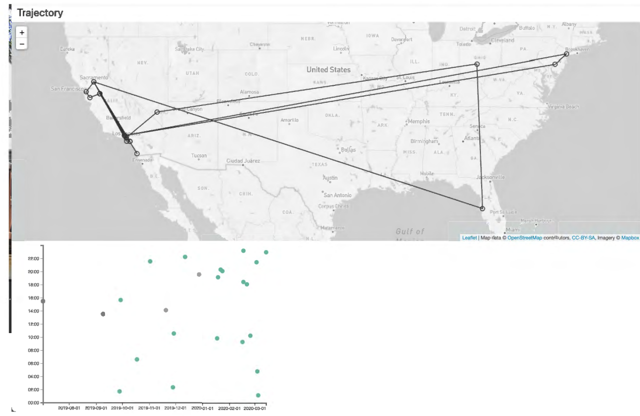

(b) Example of a phone number unlikely to be spam. Ads were posted in the right time-zones for each location and the locations were spaced out reasonably. In total, ads were posted to 14 unique locations over an 8 month period

Figure 6.3: Comparison showing trajectories likely to be spam (left) versus unlikely to be spam (right).

transport routes, users can determine if a person appears to be traveling along an interstate or if a person is likely flying between locations. Similarly as seen in Figure 6.4, users can also determine if this is a case where a trafficker posts ads to multiple locations at once and based on demand decides to move to particular locations.

6.5.4 Search View

The search view, shown in Figure 6.5, allows users to search for matching trajectories within the database. This view is initially populated using the phone number specified in the individual view. However, users can opt for a different phone number using the same search box in the upper right hand corner.

The search view is comprised of three panels: search sidebar, trajectory panel, and details panel. The search sidebar provides users with search parameter options in order to allow users to both iteratively define similarity (**DG3**) and refine their search (**DG5**). As users specify parameters for their search, the two plots within the center trajectory panel update which allows users to rapidly see how their parameter choices affect the returned

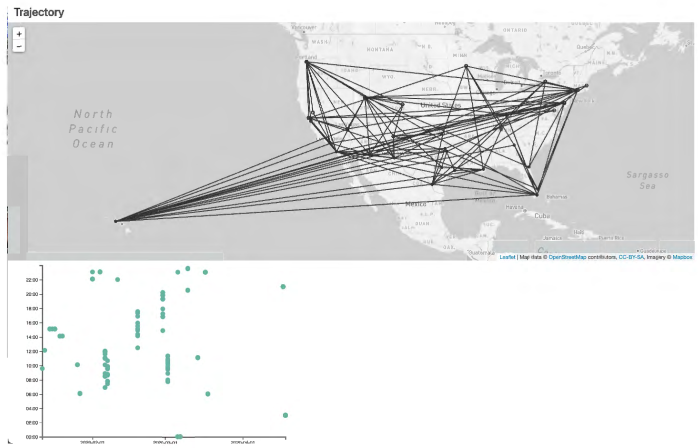

Figure 6.4: Example where a person posted ads to multiple cities simultaneously likely to decide where to move next. In addition this example also likely involved flying because all of the locations where near airports and the person began traveling from Hawaii

results. The details panel has the same functionality and content as the Details panel in the previous inspect view.

The "Search" panel has the following options for users to specify: spatial threshold (denoted earlier as ϵ), temporal threshold (denoted earlier as λ), similarity metric options, and number of total search results returned (denoted earlier as k). For the spatial threshold, users are provided an option to enter a value in kilometers via a numeric input form with the prompt "mark distances as similar when within". As users adjust the value (either by typing in a new number or by using the stepper arrows), the map view in the center "Trajectory" panel updates to show the change in search radius for each location (see Figure 6.10 for an example of this interaction). For spacial similarity, users can select between the three options using a drop down menu as shown in Figure 6.6. These options are phrased for users in terms of preferences for order matching to make the options understandable for users with no prior knowledge of spatiotemporal algorithms (**DG5**). The three options – Relative, Exact, and Ignore – correspond respectively to the three similarity metrics defined earlier: LCS, ordered set intersection with temporal and spatial point-wise matching, and set intersection with only spatial matching. For total search results, users can enter a value using the numeric input form with the prompt "Number of similar people to find". As

Figure 6.5: The **search view** allows users to search for matching trajectories within the database.

users specify parameters for their search, the Matches Preview plot within the "Trajectory" center panel updates giving users a preview of the returned search results.

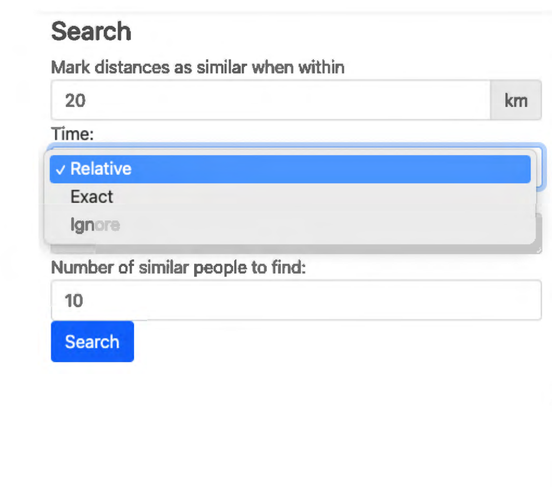

Figure 6.6: The search option control panel allows users to flexibly define similarity.

The center "Trajectory" panel consists of two main plots. The first is a map that mirrors the map shown in the previous "Inspect View". The second is a plot labeled "Match Preview" that shows a preview of the main view shown in the Analyze tab. Repeating these important plots and panels between tabs allows the users to flexibly switch between views without loosing context (**DG5**).

6.5.5 Analyze View

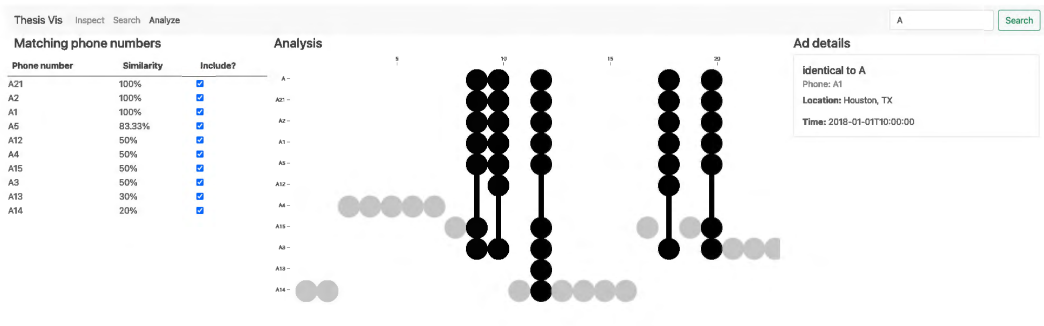

Figure 6.7: the **Analyze View** allows users to iteratively refine their search results and compare trajectories.

The analysis view, shown in Figure 6.7 provides the major analysis capabilities and allows the users to refine search results, compare trajectories, and narrow down potential leads (**DG2**). This center panel contains the same visualization as the "Match Preview" from the search tab which functions as an overview of all the matches. This plot is inspired by "UpSet" plots [265] for showing set intersections. Rows in this plot correspond to each phone number returned as a match and ordered in increasing order from top to bottom based on similarity score. Columns represent locations visited with the trajectories and are ordered left-to-right by post date from oldest to newest. Each circle within the plot corresponds to a unique ad and is colored based on whether that point matches other trajectories. Black circles indicate that the points (each ad) along the trajectory matched other points on other trajectories. Grey circles indicate that this point does not match any of the points in any of the other trajectories. When users hover over the dots in this view, the "ad details" panel on the right side is populated with the respective ad text, location, and time for the ad that the circle represents.

The "matching phone numbers" panel on the left lists all the matches returned by search query in decreasing order according to the similarity score. Next to each phone number in this panel is a checkbox allowing users to select/deselect numbers to include as matches. These check boxes allow the users to keep track of which phone numbers they have eliminated for consideration. When a phone number is unchecked, the phone number is removed from the "Matches Preview" panel; however, the user can re-add that number by checking the box once again.

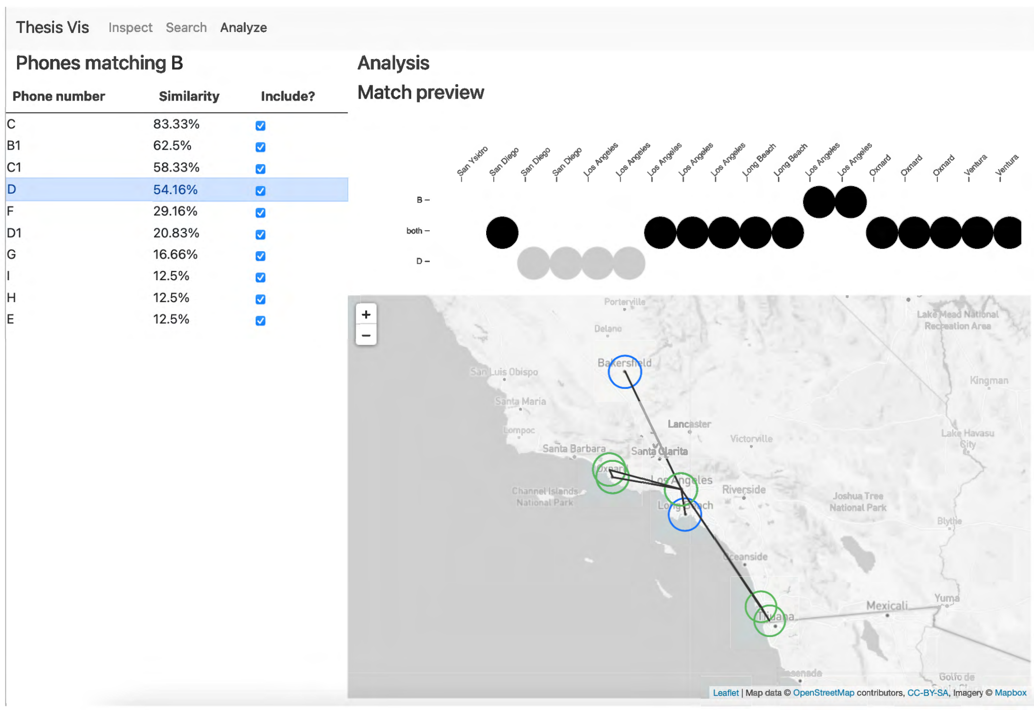

Figure 6.8: The detailed **comparison view** allows users to compare two trajectories. This view appears after a user clicks a phone number (user selection is shown using a blue highlight) in the matching phone number tables within the Analyze view tab

Users also have to option to click the phone number itself. When users click on the phone numbers in this panel, a new center panel appears allowing for pairwise comparisons (shown in Figure 6.8). The pairwise comparison center panel shows detailed information for the initially matched phone number and whichever phone number was selected from the table. As users click through the numbers in the left panel, their selection is shown with a blue highlight and the comparison view changes to reflect the new selection. When the user clicks away from the table, the center panel is restored to show the original analyze view (i.e. Figure 6.7).

This comparison view has two components: a smaller "Match Preview" view and below that a map. The smaller match preview shows the degree of specific overlap between these phone numbers in a timeline-like view. Each circle in this view corresponds to a unique ad and are ordered (left to right) by postdate from oldest to newest. The rows within this view correspond (in top-to-down order) to the initial phone number used in the search, points both phone numbers share in common, and the user-specified phone number for comparison. Points that both phone numbers share in common is determined by the user specified search parameters. Circles shown in grey (as opposed to black) in this view also indicate that that point on the user specified trajectory does not match any of the points on the original trajectory.

The map in the comparison view shows both trajectories, but colors each circle on the map in order to label which phone number each point belongs to. Circles shown in green correspond to the initial phone number (phone number “B” in Figure 6.8) and circles in blue correspond to the user selected phone number (phone number “D” in Figure 6.8).

6.6 Usage Scenarios

In this section, I will illustrate how our design can help experts investigate human trafficking cases by detailing two usage scenarios. As mentioned earlier, these scenarios are examples of how our prototype – if deployed – could be used in a real investigation. Note that these are not example narratives and do not represent a real-world deployment of the tool. However, these scenarios do draw on information from real investigations and follow the general investigation process described by law enforcement.Both scenarios begin with an investigator attempting to find an unknown number of connected victims with the goal of eventually recovering those victims. In both scenarios, the investigator begins some prior knowledge about the case – including that the case does in fact concern human trafficking, and that there are multiple victims being trafficked by the same person or group.

As described earlier, these two usage scenarios use anonymized, synthetic data in order to protect the privacy of the victims. Thus the example screenshots shown in the descriptions below do not contain any information about a real person. The first scenario uses the Interstate Mock Dataset and the second uses the Network Mock Dataset. The goal of in choosing these specific scenarios was to illustrate how our design can support different investigation types and find common human trafficking behavior patterns while also highlighting any limitations or flaws within our design.

6.6.1 Usage Scenario 1: Finding an Interstate Commercial Sex Trafficking Ring

A Virginia state police officer is investigating a suspected human trafficking case that involves a criminal organization. The officer began this investigation after a former driver for this criminal organization was arrested for an unrelated offense. During the post-arrest interviews, the driver admitted to previously being involved with a trafficking organization and gave the officer some information about the organization. The driver told the officer that the organization recruits victims from Houston, Texas and Washington D.C. and then uses a mix of greyhound busses and a network of drivers to move victims along interstates between those locations based on customer demand solicited through online ads. The driver is unsure of the total number of victims exploited by trafficking ring, but knows the phone number and alias for one of the victims he drove around in the past. In this case study this known phone number is denoted as ”A”. Based on this information, the officer’s goal is to find all the victims associated with this trafficking ring and get information on the organization’s current operation.

This case study shows how the tool can be used to find “flow” movement where traffickers move between two established locations and opportunistically stop along the way. This pattern has commonly been seen in multiple interstate human trafficking investigations. However, the data shown for this cast study is specifically based on a 2013 human trafficking case [29].

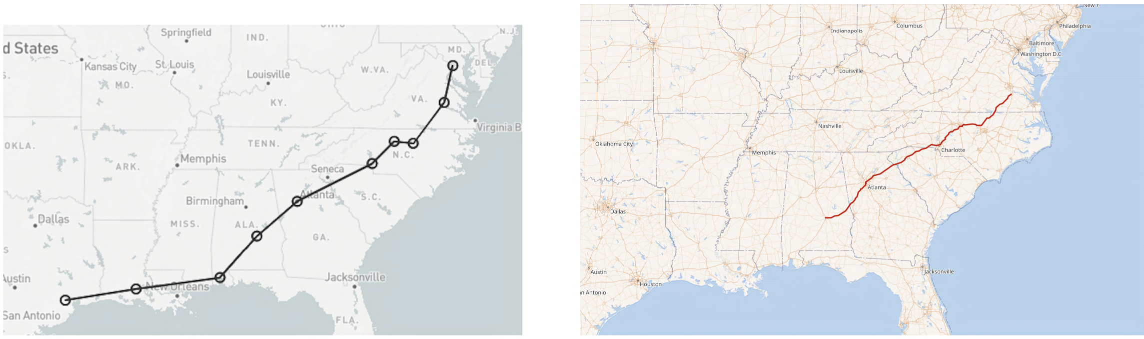

Figure 6.9: Comparison showing the trajectory path followed by phone A (left) versus Interstate I85 (right) shown in red. Interstate 85 map courtesy of OpenStreetMap and Wikimedia Commons

Analyze Initial Phone Number

The officer begins by examining phone number "A" in the inspect view, which results in the view shown in Figure 6.2. Based on the trajectory shown in the map view, the officer notes that the person appears to be traveling starting in Texas and ending in the Washington D.C. area. The route seems to follow a common interstate route between these locations and stops at major cities along the route. The officer deduces this information based on the ordering of locations and which cities the person posted from (the cities all appear in order along the same highway system as seen in Figure 6.9). This all matches the known case information provided by the driver. Using this information, the officer concludes that this phone number appears to be relevant for the case and so the officer proceeds to the next step.

Specify Initial Search Parameters and Similarity

Next, the officer navigates to the search view which is already populated with the information for phone number A. Within the search view, the officer's first task is to specify the initial search parameters. The search view asks the officer to specify search radius, search options, and the number of search results. The officer specifies each of these using context from the case.

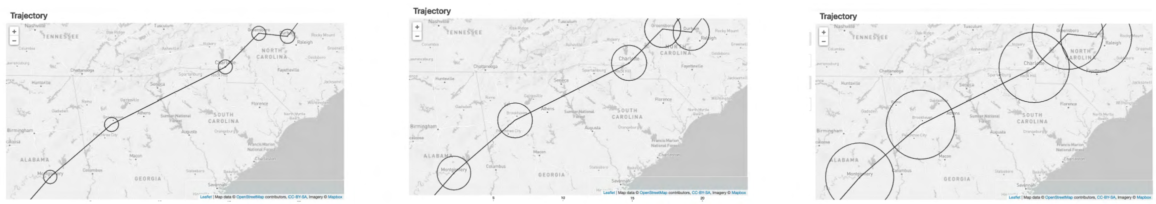

(a) Radius set to 20km (b) Radius set to 50km (c) Radius set to 100km

Figure 6.10: Comparison of different user specified search radii (20km, 50km, 100km) as seen in the Search View tab. Trajectories shown here are zoomed in portions for phone number A.

The officer knows from context that this organization moves between two states using drivers, so the officer is looking for flow like movement that follow interstates. However, The officer isn't sure of stops between then start and end points. Based on posting behavior, the officer suspects that stops are chosen opportunistically rather than planned in advance. The officer also notes (using his personal knowledge of US travel paths and routes suggested by Google Maps) that there are several possible highway systems the organization might use.

Together, this information suggests to the officer that he may want to cast a wide geographic search radius so he sets the search radius wide enough to include surrounding cities in the search. The officer uses the map view to iteratively find an appropriate radius. He starts with the initial radius of 20km, then tries 50km and 100km. As he types different values into the control panel in the search view, the map view immediately updates to show the effect these changes (see Figure 6.10 for the results of his search). He sees that 20km matches at the metro area level, 50km matches all nearby metro areas, and 100km tends to include regions. He concludes that he should start with the search radius set to 50km but notes that he might want to increase this later if he needs to cast an even wider net.

This background information also suggests that for this case, relative ordering of the locations matters. The officer knows that the victims likely all followed similar paths but not all the victims were moved together at the exact same time. As a result, he sets the search option to ”Relative”.

Finally, the officer needs to specify the total number of results the tool should return. Because the officer does not know the total number of victims and he is looking to cast a wide net in his initial search, the officer sets the value to 50. As he changes each of the parameters, the Matches Preview view updates and provides the officer with immediate feedback on how his parameter choices affect the searches.

Iteratively Refine Search to Generate Initial Hypothesis

The officer navigates to the Analyze tab where he is shown a more in-depth view of his search results (shown in Figure 6.11). His goal is to find an initial set of phone numbers to begin investigating from.

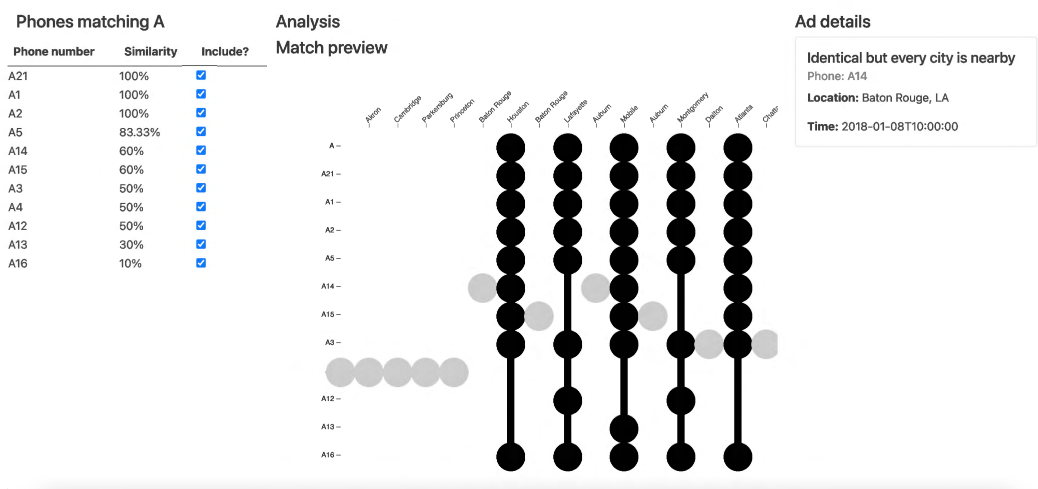

Figure 6.11: Analyze View tab showing the results of the search results described in the first user scenario

By selecting phone numbers on the left hand side, he can view pairwise comparisons between phone numbers. For example, when he selects phone number A5 a more detailed view (shown in Figure 6.12) appears. Using these comparisons, the officer concludes that phone numbers with less than 50% match score seem to be incidental matches. He unchecks them from the matches, leaving him with 6 phone numbers he feels confident beginning an investigation from.

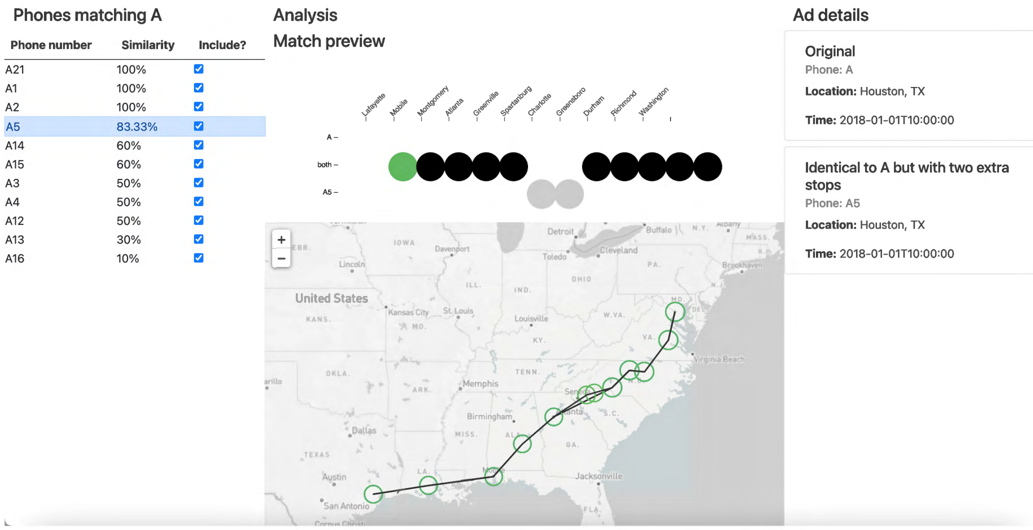

Figure 6.12: Analyze View tab showing the results of the search results described in User Scenario 1

6.6.2 Usage Scenario 2: Finding an Interstate Illicit Massage Parlor Ring

An officer in the Human Trafficking Unit in the Los Angeles Police Department is investigating reports from a hotline concerning a suspicious local massage parlor business. As part of her initial investigation, she found that the massage parlor advertised online and that the customers reviewed workers in an online forum. Using this information, she uncovered a phone number B used in a number of advertisements. Given this phone number, she wants to find all advertisements that have used this phone number as well as any other phone numbers that might be associated with the business. She knows that the more evidence she collects, the better a chance someone within the organization made a mistake and she can find out who is behind all this.

This case study shows how the tool can be used to find a ”networked” movement pattern where traffickers move victims between known established locations based on customer demand. Unlike in the interstate example above, movement between locations in this case study does not necessarily follow a particular order. The particular example shown here is based on reports from a 2016 case where an international human trafficking ring operating out Mexico and the west coast region of the US forced victims to work in massage parlors and provide sexual services to customers [25].

Figure 6.13: the inspect View shows the user the trajectory details for the specified phone number “B” from User Scenario 2

Analyze Initial Phone Number

The officer first begins by examining the phone number “B” that she uncovered through her initial investigation. She searches for “B” in the inspect view which returns the view shown in Figure 6.13. When examining the trajectory of “B” in the map view, the officer notes that the phone number was used to advertise in locations within a relatively small geographic area. All posts were to cities within the coastal region of southern California and in particular to regions within a short distance of Los Angeles, California. She knows from the hot-line tip that this operation involves a massage parlor, but she was unaware that this operation might involve multiple locations. Based on her experience investigating human trafficking cases, the trajectory taken by “B” tells her a few things about this operation. First, based on the locations and frequency of movement between locations, the organization likely owns or partners with businesses in the greater LA area as well as in Bakersfield, CA and in the southern San Diego area. Second because this trajectory has so many locations, she is likely investigating a larger scale organized crime operation. Finally, because the first advertisement appeared near the border of the US and Mexico, she hypothesizes that this case may involve border crossings and that some of the victims may not be US citizens.

Using all this information, she decides her next step is to reach out to her contacts to see if anyone else has investigated this case. She screenshots the map view to send along in her email knowing that her colleagues will find the map useful. While she waits to hear back, she decides to continue investigating this phone number and see if she can uncover any more information about the organization.

Initial Search Process

She navigates to the search view hoping to find out if other phone numbers were used by the organization. She knows from experience that traffickers tend to use multiple phone numbers at different times to advertise their victims. If the organization did use multiple phone numbers, then they likely would have posted ads to similar locations as “B” at some point. Given this context, her goal is to use geospatial similarity search to see if there are any phone numbers posted to similar locations. She hypothesizes that any phone numbers that have a high degree of similarity as phone number “B” are more likely to be relevant to her case and thus be good candidates for further investigation.

Using her background knowledge and context from the case, her first task using the search view is to specify the initial search parameters. As described in the previous user scenario, the search view asks users to specify search radius, temporal matching options, and the number of search results.

For the search radius, the officer decides to use the default option of 20km. She sees on the map that a radius of 20km results in a circle that is large enough to encompass each respective metro area while not overlapping to include other nearby cities. She experiments with increasing the radius to 40km, but finds that when she does, the resulting circles begin to have a significant overlap with each other. As a result, she decides against using the larger

40km radius because she wants her matches to contain the specific regional movements.

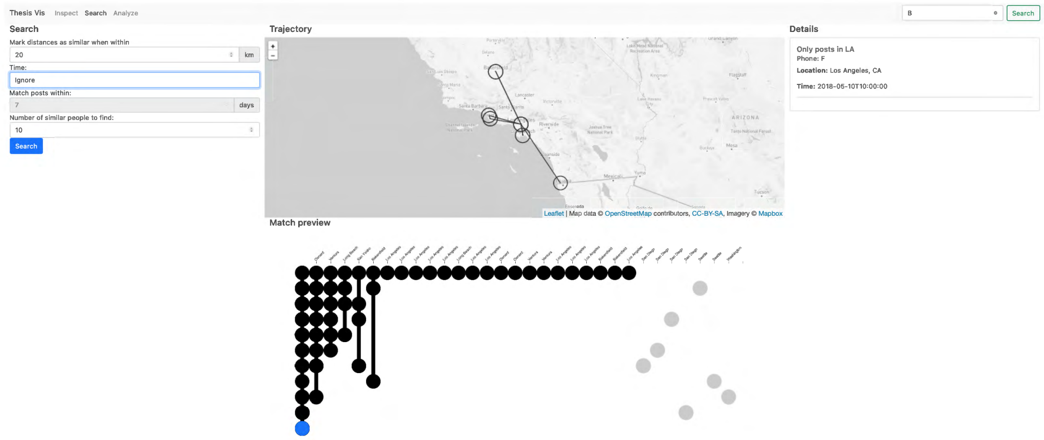

Figure 6.14: the Search View shows the user the trajectory details for the specified phone number "B" from User Scenario 2. Note that in this image, the user is hovering over one of the points in the match preview (shown as a blue highlight)

The next search parameter she needs to specify is the temporal search option. Because her case involves a network of physical locations owned by an organization, she suspects that the traffickers move victims between locations based on availability and other logistical factors. Thus, the victims are likely moving between a set of known locations but in no particular order. This also means that she isn't looking for flow-like movement (as was seen in User Scenario 1) where all the victims follow a somewhat predetermined order of locations. Given this information, she decides that her search should only take into account location similarity rather than using both location and time similarity for matching. Thus, she opts for the "ignore" order option in the search preferences (as opposed to the selecting either "exact" or "relative").

As she adjusts these settings, she notices in the match preview (shown in Figure 6.14) that this phone number is matching to a number of other phone numbers most of which tended to overlap in the Los Angeles area. This tells her two things: 1) Los Angeles may be the primary location that this organization operates out of (meaning many of the victims will have traveled to the Los Angeles locations at some point) and/or 2) because Los Angeles is in general a very popular travel destination, her search may include a number of spurious connections. Given this information, she decides her next step is to examine these matches in detail to eliminate any spurious connections.

Examining Matches and Filtering out Spurious Connections

The officer navigates to the Analyze tab to examine her search results in more detail. Her goal is to narrow down the possible matches to a reasonable set of leads she can run against a database. Not only does she need to filter out any spurious connections, but she also needs to find a set of numbers she is confident are likely associated with the criminal organization.

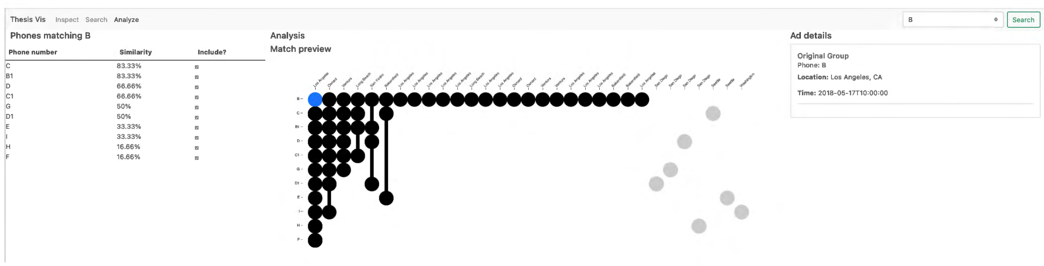

Figure 6.15: Analyze View tab showing the results of the search results described in User Scenario 2. Note that in this image, the user is hovering over one of the points in the match preview (shown as a blue highlight)

When she navigates to the Analyze Tab (shown in Figure 6.15), she is shown an overview of the matches that her search results returned. While she specified 20 matches, she notes that she is only shown 11 matches. This indicates that the tool could not find any more possible matches and she can be confident that her initial search parameters captured all of the data available.

The officer then begins her analysis by hovering over each circle (representing each unique ad) in the Match Preview plot. As she does this, she notices that some of the phone numbers in her matches look interesting as potential matches. The first four results all have very high scores for similarity and share more than half of the same location posting behavior. In addition, she notes that three other phone numbers (two of which are top matches) all posted to the same location San Ysidro, CA in their first posted advertisement. Based on this information, she forms an initial candidate set of phone numbers (C, B1, D,

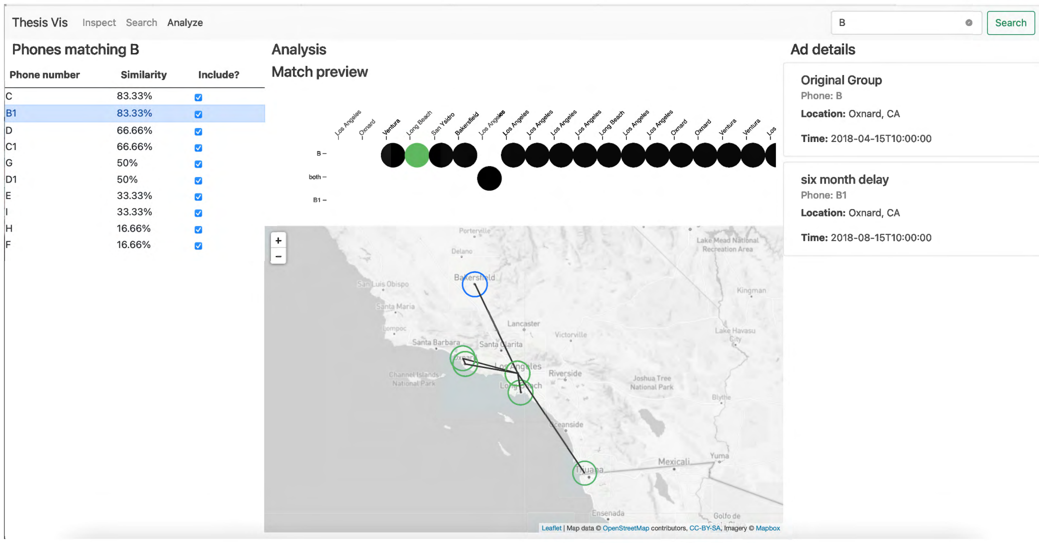

Figure 6.16: Example of the detailed comparison views between phone number "B" and phone number "B1" from User Scenario 2. User selection from the table is indicated by the blue highlight.

and C1) to consider as potential matches. She further examines each of these numbers to confirm the high overlap by clicking through each phone number in the left hand table to compare each match to the original phone number. An example of this comparison can be seen in Figure 6.16 showing the comparison between phone number "B" and "B1".

After forming her candidate set, she then examines the phone numbers that have low scores to decide whether or not to eliminate them as potential leads. She sees that both phone number H and F have very low matching scores and tended to only match posts within major cities. In addition, neither H nor F posted to San Ysidro, California. Thus, she suspects that these two phone numbers are likely spurious connections and so removes them from consideration.

Restarting Search

After removing H and F, she decides to restart the search process using each of top four candidate phone numbers (C, B1, D, and C1) to see what each of them match. The idea is that if each of the numbers is a good candidate then they likely match to similar sets of candidates and may match to other phone numbers that should be considered. She decides that she should add any phone number to her candidate set if multiple phone numbers also returned that number as a result. This process lets her refine her candidate list and double check if H and F are in fact spurious connections.

She navigates back to the Search tab, and restarts the search process using each of the four top candidate phone numbers and identical search parameters. As she does this, she jots down the top matches shown in the preview tab for each of the phone numbers. Looking at the overlaps between the top matches between each of the returned results, she notes that phone numbers D, B1, and C1 all match both the initial 4 candidate set as well as to phone G and D1 with a high similarity score. This tells the officer that she should expand her potential candidate set to include phone numbers D1 and G. As a result of this process, she feels confidant that all the phone numbers in her candidate set (C, B1, D, C1, D, and G) are very likely connected to her initial phone number B and are her top priorities for investigating further. She also notes that phone numbers E and I sometimes appeared as potential matches, but more often with a lower similarity score than those in the candidate set. She makes a note that theses numbers are possibly connected to her case, but that they possibly a lower priority for investigating compared to her candidate list. Finally, she finds that no phone numbers in her list matched H and F with high similarity scores and so she is confident that these two numbers are in fact not connected and can be eliminated for consideration.

6.7 Case Studies

In addition to providing detailed usage scenarios, we also performed two informal test cases using the design in order to understand how the the tool might work in practice. In these test cases, we asked another researcher to try to use the tool to answer a series of questions. The goal was to get an idea of whether our approach would even be capable of finding matches using more complex data than was shown in the usage scenarios. In this section, we include a description of these two test cases (which we refer to as case studies) to highlight the potential promise the design holds. However, these case studies are not user evaluations but are instead illustrations of the tool's potential functionality with respect to finding potential matches and creating "profiles".

In each of these test cases, our volunteer "user" was asked to act like a potential investigator working a human trafficking case and was then given a phone number to investigate alongside a series of constrained tasks to perform as part of a simulated investigation. Phone numbers for each case study were chosen at random from subsets of the overall dataset that met certain criteria. These criteria were deliberately chosen to highlight the range of possible scenarios a real investigation might encounter while also highlighting both the strengths and limitations of our design. To guide their investigation process, the user was given a set of open-ended questions to answer about the phone number and were tasked with identifying if there exists any connections to other phone numbers in the database. The user for the case studies was not aware ahead of time if there were any connections, nor did they have any prior knowledge about the particular phone numbers.

In this section we will show screenshots of the tool using real data, but the screenshots are cropped and censored to remove any possible identifiable information. We try to limit the number of screenshots needed to explain the analysis and chose only to show specific examples that do not have immediately identifiable behavior patterns. Further to protect the privacy and safety of those represented within our dataset, we will not reveal the real phone numbers used in case studies.

6.7.1 Case Study 1: Investigating Frequent Movement

The initial phone number for this case study was chosen at random from a list of phone numbers found in the dataset that advertised in 3 or more unique locations. The goal with this case study is to showcase how the tool can be used for more exploratory analysis on the dataset. Because the phone number was chosen at random, we had no prior knowledge about the phone number and relied exclusively on the information in this tool and our own knowledge of trafficking for the analysis.

Initial Investigation

The goal in our initial investigation was to build a profile of the phone number. For this process, we loaded the phone number into the Inspect View in the tool and took note of any initial observations we had. In addition to noting our general observations, we also specifically sought to answer five questions about the phone number inspired by those used by law enforcement in an investigation. These questions were: 1) ”are there any patterns in movement or posting behavior”, 2) “Does the phone number look legitimate versus spam”, 3)“are there any human trafficking indicators”, “can you tell anything about 4) mode of transport and 5) trafficking operation type?”. I give a summary of our observations below.

Figure 6.17: The trajectory details as seen in the Inspect View for the specified phone number from the database.

The phone number used for this case study had a California area-code, but was used in multiple cities across the US. The trajectory as shown in Figure 6.17 seemed somewhat unique - the locations advertised at were mostly in smaller suburban cities rather than major cities within each state. The trajectory path started in Utah, traveled to southern California, then to Florida, then Texas, and ended in the Connecticut and New Jersey area. The time between cities was on average two days apart and across a larger geographic distance, which indicates that the person could have possibly used airfare for transportation. However as most of the cities in the trajectory were only near regional airports, it seems less likely that the person flew between locations. Instead its seems more likely that the person used busses to move between locations because most of the cities in the trajectory were near major Greyhound bus stations.

We also noted that this trajectory had an incorrect GPS coordinate assigned to one of the locations. MapQuest's geocoding process assigned the location "Central, Jersey" to Jersey, Arkansas. Based on the text within the ad, it seems more likely that the intended location was within the state of New Jersey.

The posting behavior pointed towards the phone number being legitimate. The locations were reasonably spaced out, posted at variable times and at a relatively low frequency. Beyond the frequent movement, there were no other clear indicators of trafficking.

Finding Possible Matches

Because the trajectory seemed relatively unique, we opted to first search using relative ordering and return the top 50 matches. We were hoping by using relative ordering instead of exact matches, we would have a higher chance of finding a match (if there was one) despite the location error within the trajectory. Because this phone number was chosen at random, we had no prior knowledge about potential matches for this trajectory.

The overall match scores shown to us on the Analyze view were relatively low - ranging from 23.52% to 14.28%. Based on these score, we initially concluded that this phone number was probably not similar to any other number in the dataset. To confirm our hypothesis, we examined each of the numbers in the Analyze tab and looked for indications that the numbers were likely connected (such as using the same text in ads across phone numbers or referring to multiple phone numbers in the ad text) or likely not connected (such as ads referring to very different people or that the trajectories only had incidental overlaps). One phone number in the returned results however had identical ad text to the ads posted with the original phone number. This indicates that these ads were written by the same person and therefore these phone numbers are connected.

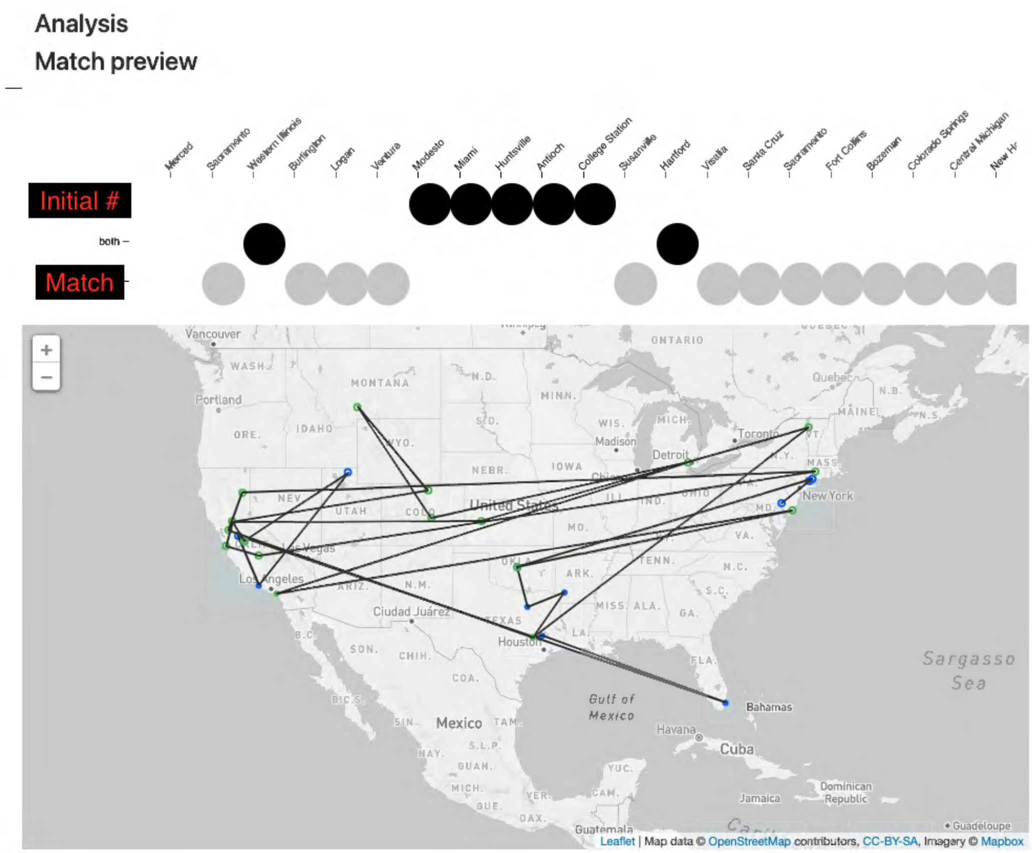

Figure 6.18: Overview of the match details between the initial trajectory and the matched trajectory.

The overall similarity score was 15.87% which seemed fairly low. We did note that this is likely because the matched phone number posted to more cities across the US so the initial phone number only has a partial overlap with the matched trajectory (see Figure 6.18). Given this information, we concluded that our next step should be to examine the matched phone number in more detail.

Examining Matched Phone Number

We examined this phone number in the inspect view, and observed that this phone number had quite a different posting behavior from the initial phone number. This phone number was used to post to a number of different cities at the same time, then a gap in posting, and then a number of posts to multiple cities at the same time. We compared the two timelines by switching back to the Analyze tab (see top half of Figure 6.18), and noted that the phone numbers posted on different days with a leader-follower style delay. The matched number would post to multiple cities, then the initial phone number, and then back to the matched number.

Based on this information, we came up with some possible explanations to explain this behavior. The first was the this matched number might be spam that is harvesting and re-posting other peoples ads. The ads all posted by the matched number were exact copies of only the ads posted by the initial phone number. Another possible explanation is that both phone numbers were used by the same person, but they used the matched number to scope out locations before traveling. That is similar to behavior noted by law enforcement, this person might be using the matched number to gauge interest before buying tickets and making arrangements to travel to locations.

6.7.2 Case Study 2: Finding matches with Limited to No Movement

For this case study, we randomly chose a phone number that visited 2 or fewer unique locations. This case study showcases how the tool can still be used to find matches even in cases with little to no movement. Because this phone number was chosen at random, we had no prior knowledge about the phone number or if the phone number should have matches within the dataset.

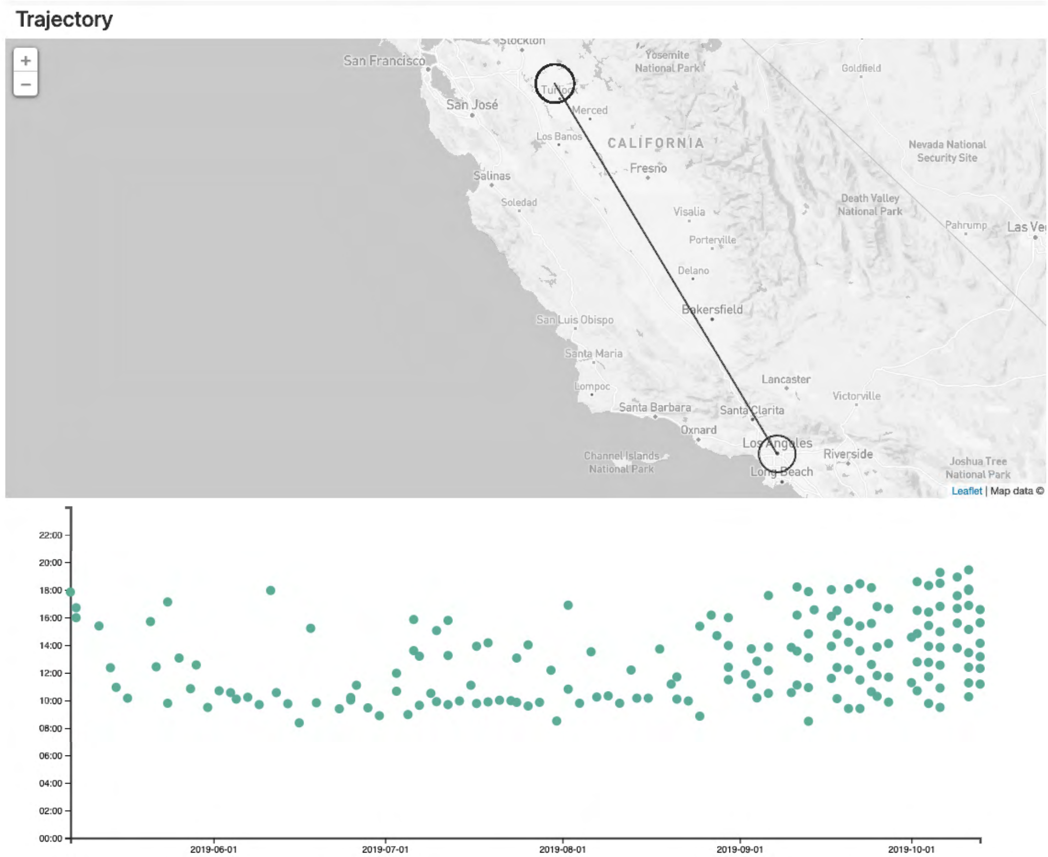

Figure 6.19: Overview of the details for the trajectory used in Case Study 4 showing exclusively hyper-local movement.

Initial Investigation

For this case study we used the same questions and processes to guide our observations as was used in the previous case study. As seen in Figure 6.19, the phone number was only used in advertisements in the Los Angles area and briefly elsewhere in California. The person posting seems very active and posts ads multiple times per day during typical work hours (though tends to post most of the ads between 8AM and 10AM).

The ad text is suggestive of an operation involving a set physical location. The ads feature phrases like ”come visit us” which points to the person using a set-up where the customers visit an established address rather than the person driving to visit the customers. In addition, the ads mention the name of the organization/business running the ads and indicates that these ads are for multiple women working within that organization. The use of a shared management coupled with phrases like “fresh girls everyday” are indicative of human trafficking, but aren’t necessarily specific to just human trafficking operations. While shared management is often associated with human trafficking, some voluntary sex workers will employ a manager to help with scheduling and logistics.

Finding Possible Matches

Because this trajectory involved very little movement and posted mostly to a very popular major US city, we expected that finding matches (if there are any) would be more difficult. As a result, we decided that we should try to cast a very wide net. We initially set the search option to ”ignore” and the number of matches to 50. This meant the similarity metric would give high scores to anyone who posted very frequently to the LA area. This would also mean however, that the matched results would also likely include a high number coincidental overlap.

We later set the search option to “Exact” to see how those options affected the results. The “Exact” option would give high similarity scores to phone numbers that posted to LA during the same time period as our initial phone number so would have the effect of filtering out any older phone numbers. However, we found that there weren’t many differences between these two search options for this particular phone number. This is likely because this particular phone number has been active for a number of months so it would match to roughly the same set of phone numbers.

By clicking through each phone number in the analyze tab, we were able in under 10 minutes to narrow down the list of potential matches from 50 to just 4 based on posting behavior similarity – one of which is highly likely to be a match as the ad text. That matched number names the same organization as the original and posted some identical ads as the original. The similarity score between the matched number and original phone number was 13.06% which indicates that as expected, geographic similarity when there is limited movement may not be as meaningful. However, we were still able to find match in a short period of time with relative ease even-though the person didn't move. This demonstrates that there still is some value to examining posting and movement behavior even in cases with the absence of movement.

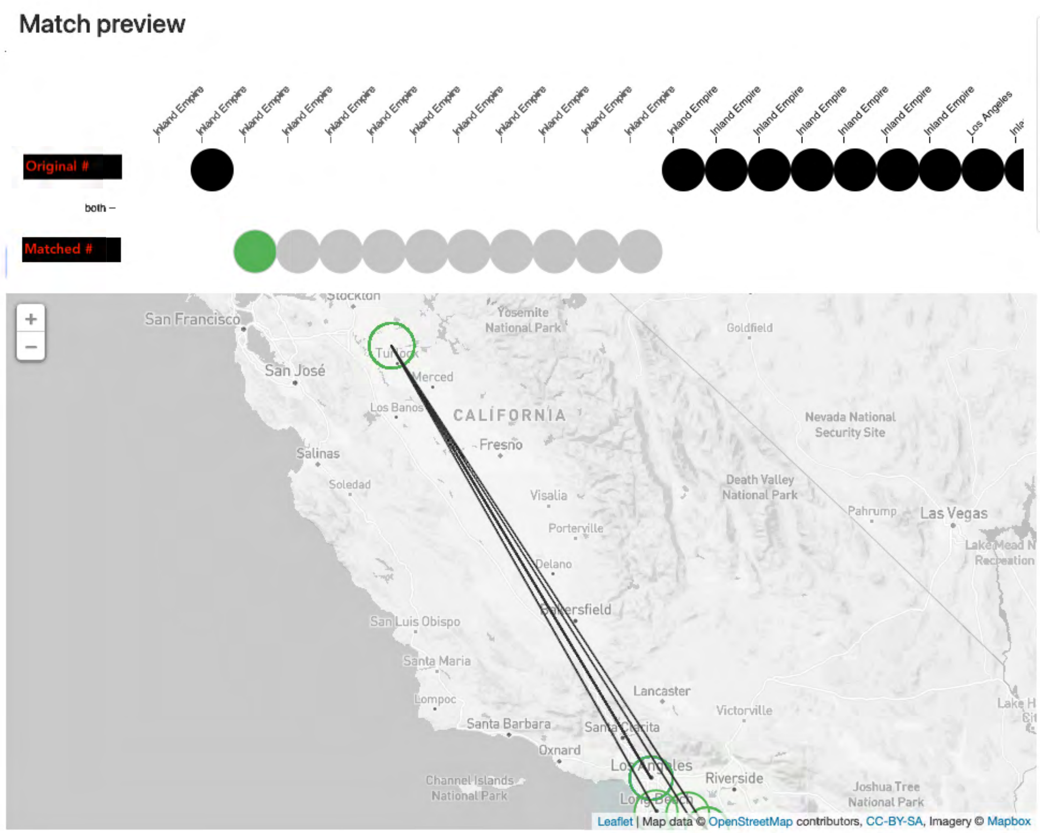

Figure 6.20: Overview of the match details between the initial trajectory and the matched trajectory for the fourth case study.

When examining the matched phone number in more detail (as seen in Figure 6.20) we found that the the matched number tended to post to specific areas in the greater Los Angeles area rather than just to the main city. This means that our choice of radius ended up having a significant impact on whether or not we would find a match.

6.8 Discussion and Limitations

In this section, we examine the limitations of our design and examine areas for future work. As both user scenarios used data containing a controlled set of matches, we can use

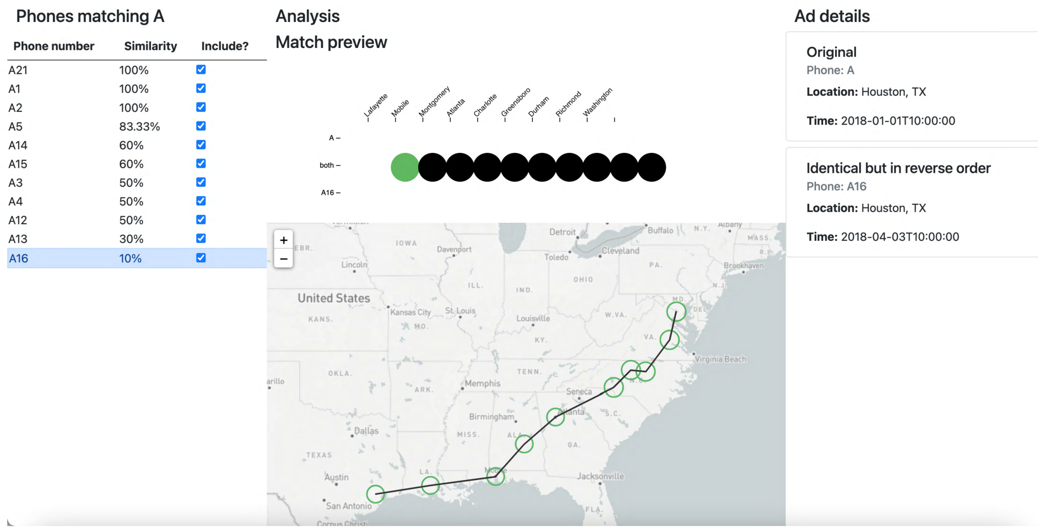

Figure 6.21: Example of the detailed comparison views between phone number "A" and phone number "A16" from User Scenario 1. Trajectories for each are identical but in reverse order form each other

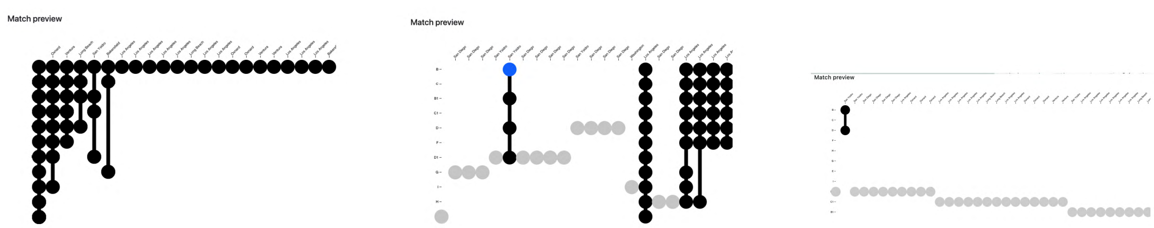

(a) ignore ordering (b) relative ordering (c) exact ordering

Figure 6.22: Comparison showing the effect of different user-specified temporal ordering options (ignore ordering, relative ordering, exact ordering) as seen in the Search View tab using phone number B from User Scenario 2.

the results to analyze what patterns the tool was capable of finding and where the tool is less capable. In both user scenarios the tool was able to find all possible matches, but some movement patterns proved more challenging to find. In particular, trajectories that are in reverse order from each other could only be found in limited scenarios. In User Scenario 1, the trajectories for phone A and A16 were identical but in reverse order from each other (see Figure Figure 6.21). Phone number A16 only would be returned as a high match when the user selected the "ignore" time search option. In all other combinations of search options, reverse order trajectories would not be found as a match.

This problem with reverse ordering points to a larger pattern we noticed – which is that search options very much matter in determining which phone numbers are found. As seen in figure Figure 6.22, what option a user selected between the different user-specified temporal ordering had a profound effect on which phone numbers are returned by the tool. These difference mean that depending on which option was selected, a user will have very different conclusions about which numbers are matches and even which are possibly good matches. Further what phone number is selected as an initial starting point can have an effect on which numbers are returned. Doing re-starts as was shown in User Scenario 2 is one possible strategy for overcoming some of those limitations and we find the flexibility of the tabs support this process. However, future designs might need to take into account how to support other strategies for examining overlaps between multiple searches.

Finally, our hypothesis was that the system would perform well when the initial trajectory was long and had lots of movement. We also suspected that it would perform poorly at best (if not at all) when the trajectory was very short or had little movement. To our surprise, the tool was still useful at uncovering matches even when when movement was not particularly present. In both User Scenario 2 and Case Study 2, our tool was still able to find potential matches fairly quickly even when the initial trajectory had little overall movement. However, we still found that the tool was not as good at finding cases where both trajectories had limited overlap. For example, in User Scenario 2 phone number E was was identified as a potential match but was given a much lower priority for consideration due to the low degree of overlap between it and the other phone numbers. A fundamental assumption our tool makes is that overlap is a strong indicator of connection. This assumption may or may not hold in all human trafficking cases. As a result, our tool would not be able to find cases where a person posted to different cities each time with a different phone number or in cases with diverging/converging behavior where the overlap was only for a very short period of time. Whereas methods like text or image similarity might potentially be able to make those connections because they do not rely on trajectory similarity for determining matches. Future work should look at examining how combining these approaches for determining similarity (i.e. using a mix of text, image, and movement behavior similarity) can complement each other and improve search outcomes. Many of the tools that are established in practice include mechanisms for matching phone numbers using a mix of text and image similarity, but none use geospatial or temporal similarity as a part of this process. Given the apparent utility that our approach demonstrates, we imagine that if combined this could be a really powerful way of finding connected cases.

6.9 Ethics Analysis

In line with the suggestions we noted in chapter 5, our research should incorporate ethics-driven considerations and document the harms and benefits of our work. To this end, we use our prior analysis as well as the ethics principles established in other prior work [8, 107] to understand the limitations and ethical tensions our design could potentially face. We structure our analysis by first documenting the benefits of this research and our design, and then we conclude with a discussion of the potential harms and any strategies to address those harms.

One of primary benefits of this work is that the methods and design could potentially be used to combat human trafficking and increase the efficiency and efficacy of existing investigative strategies. There is a strong need for tools to assist with computational tasks like linking cases based on movement behavior similarity [13]. As seen in both the user scenarios and in the case studies, our design can speed up this linking process and as a result enhance existing practices for identifying victims and key partnerships relevant to an investigation. The development of technology for these purposes were key strategic goals outlined in both the Federal Strategic Action Plan on Services for Victims of Human Trafficking [14] and the U.S Department of Justice's National Strategy for Child Exploitation Prevention and Interdiction [15].

In addition, our design is intended to addressing the investigation needs for more challenging cases of human trafficking. As discussed in both user scenarios, our tool is particularly beneficial for finding cases involving organized crime, gang activity, or other forms of organized human trafficking networks because our tool is specifically used for identifying for coordinated behavior. This means that our work can potentially have an impact on disrupting these organized crime networks which currently remain under-investigated despite committing more egregious offences [13, 90].

Finally, another benefit of our tool is that it can serve as an example of cross-discipline collaboration in the design of technology. As noted by the ethics principles for human trafficking research established in [8], technological research should involve collaboration across across government, non-profit and non-governmental, academic, and private sectors. Throughout our design process, our work required collaboration with law enforcement, non-profit, and policy organizations; the lessons and design implications we learned through working with our collaborators can serve to inform future technological designs in this space.

In terms of potential harms, we note that our dataset very likely includes data on voluntary sex work. Both voluntary sex-workers and trafficking victims are advertised and discussed on the same sites and there is no way to cleanly separate out advertisements for one group from the other [8]. As a result, there is the concern that collecting and using such data could lead to increased surveillance on vulnerable populations [229, 230]. Thus we must take every precaution to both protect the right to privacy of individuals within our dataset and discuss safe-guards against misuse of our design.

In order to protect privacy rights, the dataset used to design and evaluate the tool is not publicly available and will be deleted after the project has concluded, In addition, we tried to be as careful as possible when selecting case studies and screenshots to include in order to not inadvertently reveal identities. Instead, we use simulated data to show the full capabilities of the tool. This simulated dataset can be available and would allow for demonstrations and public input into the design.

Finally, we note that our tool and future versions of our tool if deployed would come with limitations of use. Our design is only intended to be used for cases directly related to human trafficking and by law enforcement that has received some human trafficking training. Based on both our prior work [13] and discussions with our collaborators, we limit the use to that specific population because law enforcement with human trafficking training typically use trauma-informed policing practices and demonstrate a nuanced understanding of exploitation and sex-worker rights. Our design is not intended to be used to investigate other crimes nor is it intended to be used widely by law enforcement in general. Our design is necessarily limited in scope as we have only studied that specific population (law enforcement investigators with human trafficking training), and so we only have insight into their practices and potentials for harm.

6.10 Conclusion

Human trafficking investigations require complex investigation strategies and increasingly involve more data science tasks. One challenge investigators currently face is recognizing connected cases. As uncovered in our prior work, there is a strong need for computing tools to support methods for searching for connected cases.

In this chapter, we introduced a visual analytics tool that tackles one of these challenges identified in our prior work. Officers were often in need of tools to support geographic and temporal search especially for data where GPS coordinates are not immediately available and involve some level of uncertainty. While the focus of our work was specifically on human trafficking cases, we envision that this work would be effective for other domains. As noted in the prior criminal justice literature, geographic temporal similarity search is an effective strategy for uncovering connected cases concerning other crime areas. More generally, our design can be applied to any task that involves searching for similar trajectory pattern behavior. We envision that our tool could also provide benefit to investigations using crime data such as those on drug and arms trafficking (which also often include social media data).

We envision 4 main contributions with this work. Our first contribution is the design goals and challenges we identified to support human trafficking investigations. These design challenges provide a unique insight into the design considerations when working with law enforcement.

Our second contribution lies in the development of case studies and usage scenarios. We demonstrate the efficacy of this tool through two user scenarios and two cases studies with detailed descriptions of how this tool can be used to identify connected cases of human trafficking. In addition we tested our tool on real data as demonstrated in the case studies. Further these case studies demonstrate how the tool can be used to get unique insights into the data itself and discover important patterns like likely spam, likely human trafficking, likely modes of transportation used, and even what routes or highway systems the trafficker tends to frequent. This demonstrates the utility of examining movement behavior patterns for specifically human trafficking cases. These scenarios provide evidence that our tool would be an effective addition to real investigations. Together, these scenarios illustrate the effectiveness of our design in real-world applications.

Our third contribution is an ethical analysis of the design an usage of the tool itself. In the prior chapter, we noted that AI development in this area would benefit from including ethical analysis within the design process. To this end, we included a explicit ethics and limitations section to discuss these issues.

Finally, the main contribution of this work is the design and development of a visual analytics tool specifically designed to tackle the needs of law enforcement. Our design takes a novel approach that combines existing techniques for similarity matching and visualization that allows users direct and flexible input in defining what constitutes group behavior. Further our design supports multiple kinds of investigations and behavioral analysis. As seen in the case studies and user scenarios, our tool can be used to quickly narrow down potential leads given some prior knowledge about a case. Further, our tool still works even in scenarios where movement is absent.

Future work should look to involve heterogeneous data sources especially those with even more geographic uncertainty. Future work should also look to integrate this design into existing tools. And finally, future work should look to perform an in-situ evaluation of the tool. Unfortunately due to COVID restrictions, we were not able to perform a live evaluation.

CHAPTER 7
REFLECTIONS AND CONCLUSIONS

In my thesis, I have examined the ways in which computing researchers can assist anti-trafficking efforts. First through an interview study with law enforcement, we gained a better understanding of the criminal justice response and the technological needs of this group. This work sought to augment our understanding of the prosecution aspect of anti-trafficking efforts. Then we examined the prior work through the lens of human rights to get a better understanding of the nuance and ethical practices of developing AI in this space. Finally using the insights gained from both projects, we constructed a prototype of a tool designed to fill some of the needs brought forth by our interview participants.

Ultimately my thesis work's contributions to this space are as follows:

- A detailed account of the unique technological and design needs of law enforcement investigating human trafficking cases. Our work notes the complex social, political and technological challenges associated with designing for law enforcement and highlights avenues for future computing research.

- An outline of the law enforcement investigation process and the role technology plays within an investigation. Our work further describes the current investigation roadbloacks and any opportunities for technology to assist with overcoming said roadblocks.

- The categorization and understanding of the ethical risks associated with developing AI tools for anti-trafficking efforts that uses the language of international human rights law. Our work notes five suggestions for future research: broader use of participatory design; engaging with other forms of trafficking; developing best practices for harm prevention; including transparent ethics disclosures in research; and a discussion on the what limitations should exist for of AI in anti-trafficking.

- A series of design goals, case studies, usage scenarios, and ethical analysis of a visual analytics prototype that demonstrate how computing methods can be used to identify connected cases of human trafficking and overcome known investigation challenges.

- Finally, the development and implementation of a visual analytics prototype that presents a novel approach that combines existing techniques for similarity matching and group behavior visualization in a uniquely challenging domain.

As I reflect on my work in this chapter, I want to draw attention to future research and open questions brought forth by this research, and then I will conclude with lessons I hope researcher take away from my work.

7.1 Open Challenges with Combating Human Trafficking

There are a number of complex challenges that anti-trafficking efforts still face. Through my research, I have encountered four main challenges: 1) lack of representative data, 2) lack of resources for interventions, 3) lack of collaboration across jurisdictions, and 4) lack of a shared understanding of human trafficking.

1) There is a lack of existing data on human trafficking and what little data does exist tends to be incomplete or missing in some cases. Researchers and anti-trafficking groups have fragmented access to human trafficking data - in part because there is no centralized, shared location for data on human trafficking. Data on human trafficking tends to be siloed in a combination of law enforcement, NGO, and governmental databases each of which have (understandably) different requirements and protections for access. Further, these databases are fragmented geographically. Thus given that victims often move between jurisdictions, researchers often have to get access to datasets both across different sectors and across different geographic regions to get more complete picture. As a result of these data

access challenges, it is difficult for researchers to get a complete, holistic picture of human trafficking and get datasets of sufficient size to use for certain computational methods.

Finally, there is also a near complete lack of ground truth data on human trafficking which makes evaluating models more difficult. For example, much of the prior (including our own) use online sources of data that likely contain information on trafficking victims. However as sex trafficking inherently exists as a component within the consensual sex-work industry, distinguishing between consenting sex workers and trafficked victims is non-trivial task and there are no standardized techniques to separate activity for human trafficking versus consensual sex-work. Some of the prior work including in [8], [93], and [20] rely on law enforcement or NGO-generated key-words to label advertisements as likely involving human trafficking. However, there are often disagreements among experts as to how accurate these key-words are at predicting human trafficking cases. To date no research has empirically evaluated the reliability of these keywords.

2) Anti-trafficking efforts are also incredibly resource constrained. There is a distinct lack of sufficient funding for anti-trafficking initiatives both globally and within the United States. Additionally, law enforcement and victim advocates are time-constrained and can only assist with a select number of cases. As a result, many officers describe feeling like they are only only scratching the surface of human trafficking cases because they are so resource constrained.

3) Multilateral collaboration remains an open problem for anti-trafficking initiatives especially at a global scale.

4) Finally, there is a lack of shared understanding of human trafficking. There are a number of pervasive myths and even more recently conspiracy theories that prevent effective interventions. Misinformation remains an ongoing problem especially with the spread of QAnon beliefs.

In addition to these challenges, there also remains a major gap in work addressing non-prosecution efforts within anti-trafficking efforts. There is a need to address human trafficking from a more holistic perspective that addresses efforts associated with the other 3Ps: Protection, Prevention, Partnership.

7.2 Open Questions & Future Research

From these existing challenges and gaps, I see a number of future research opportunities.

7.2.1 Addressing Data Problems and Access

Addressing the problem of incomplete datasets represents a unique opportunity for computational researchers because there already exists a wealth of established techniques designed to navigate this concern. For example, researchers can use techniques from operations research such as matrix completion [18] and network completion [19] to infer missing information. From data science, researchers can use techniques like entity unification to associate data points across disparate data sources.

However, applying these techniques to human trafficking datasets brings up a number of open research questions. For example, how can we unify databases in a privacy preserving manner? With inference models, how can we ensure accuracy of the inferred data points? Especially in criminal justice settings where data provenance is critical, how can we convey to users that the datasets contain a mix of sources and/or inferred data points?

7.2.2 Addressing Resource Constraints and Allocation

Because anti-trafficking efforts are so resource constrained, interventions have to be designed to most efficiently use the time of the individuals involved. This problem highlights an avenue for future research using computational techniques to support efficient resource allocation for anti-trafficking efforts. For example building off of the geospatial-temporal analysis shown in chapter 6, future research can look to identify spatial and temporal regions where trafficking increases so that anti-trafficking efforts can more efficiently respond. Future work could also look to identify individuals who are most at risk of being trafficked so that early stage interventions can be targeted for those individuals and be most effective at preventing trafficking.

7.2.3 Addressing Collaboration and Partnership

Collaboration across all levels remains an open problem and there are unique ways computational research can help. This area represents an opportunity to revisit classic CSCW problems but instead with a focus on the needs of government sectors. For example, there is a need for tools to support information sharing both within organizations and between different sectors. Both law enforcement and victims need tools to help vet and connect them to appropriate service providers. Research in this area can also look at identifying and evaluating existing barriers to collaboration and information sharing.

Finally, computation models can also be used in future research to identify potential collaboration opportunities within anti-trafficking networks. For example, by identifying regions traffickers often travel to and from computational models can identify necessary collaborators across geographic regions by ranking each area's connectedness.

7.2.4 Addressing Prevention

As noted in chapter 5, there is a distinctive gap in work designed to prevent human trafficking. There are opportunities for researchers to assist with existing prevention strategies by designing better targeted interventions. Current prevention strategies like education and awareness campaigns tend be limited in scope and tend to miss critical vulnerable groups [12, 31]. Notably despite being disproportionately affected by trafficking, black and Hispanic communities in the US are not addressed in current campaign strategies [12, 31]. There is a need for future research to design and evaluate what strategies would work to target outreach to these marginalized populations that take into account the needs of these communities. Additionally, certain geographic regions are not addressed by current strategies. Notably, regions close to extractive industries see increases in both sex and labor trafficking [31]. Both the rural nature of these regions and the lack of existing infrastructure make these regions particularly vulnerable to trafficking and difficult to reach with broad campaign strategies [31]. Borrowing from related work on internet campaigns in low-internet regions, Future research could look into how to design campaigns that effectively reach rural regions. In addition, researchers could evaluate the effectiveness of existing prevention strategies and model which combinations of these strategies would be most effective in different geographic regions and across different sub-populations.

From a social computing perspective, future research could also look to design prevention strategies aimed at targeting labor exploitation facilitated by online job boards. In particular there is a particular need for interventions designed to help vulnerable populations like former convicts and migrant workers who are both vulnerable to exploitation and tend to be less computer literate.

7.2.5 Addressing Protection and Survivor Support

More research is also needed to understand the needs of survivors particularly with respect to how technology can be used for support. There is a need for tools to assist with accessing resources such as finding employment and housing; with accessing education like job training; and with finding resources for healthcare needs.

In addition, we also need to understand how human trafficking interventions impact survivors with special attention given to understand and prevent further victimization. This brings up a number of open research questions including: "how can we design for harm prevention"; "how can we design tools that are robust against missuse and corruption?"; "how can we design tools to proactively (rather than reactively) search for misuses?". Future research should examine the role of technology for both harm prevention and for perpetuating violence.

7.3 Takeaways

Finally, I want to highlight some takeaways I hope to leave readers with. First and foremost, a better understanding of human trafficking that dispels common myths. Perhaps the most important lesson I learned through this work was the number of socially pervasive myths about human trafficking. Cultural narratives surrounding human trafficking tend to focus on a narrow subset of human trafficking - namely sex trafficking cases where the victim is young, white, and female and where her trafficker is an adult male stranger. The reality is that human trafficking takes many forms and both victims and traffickers come from incredibly diverse backgrounds.

Further I hope to leave readers with an understanding that human trafficking is inherently linked to inequality and marginalization. Human trafficking is a deeply complex problem that is difficult to research and requires navigating intersecting complex political and social issues. Our understanding of human trafficking requires examining the complex social, economic, cultural, and political factors that each play a role in facilitating this crime. But at the root of it, those most vulnerable to exploitation are marginalized groups who have unequal economic opportunities. Thus, dismantling systems of oppression and implementing necessary social programs (like increasing the number of long-term shelters and support for equal education opportunities) will impact human trafficking. To this end, we also must be cognizant of where and when technology should and should not be used. Technology should not replace or be used instead of much needed social programs, training, and education.

Finally, I hope my research inspires more computational researchers to center ethics within their own work. Ethics provides an invaluable perspective to computational research and ensures that we examine how our work impacts the communities around us.

Appendices

REFERENCES

[1] K. Kangaspunta, F. Sarrica, R. Johansen, J. Samson, A. Rybarska, and K. Whelan, "United nations office on drugs and crime: Global report on trafficking in persons 2018," 2018.

[2] International Labour Organization, *Global estimates of modern slavery: Forced labour and forced marriage*, 2017.

[3] R. A. Konrad, A. C. Trapp, T. M. Palmbach, and J. S. Blom, "Overcoming human trafficking via operations research and analytics: Opportunities for methods, models, and applications," *European Journal of Operational Research*, vol. 259, no. 2, pp. 733–745, 2017.

[4] Caliber, " Evaluation of Comprehensive Services for Victims of Human Trafficking: Key Findings and Lessons Learned," US Department of Justice, Tech. Rep., Apr. 2007.

[5] Polaris, "Sex Trafficking in the U.S.: A Closer Look at U.S. Citizen Victims," Polaris, Tech. Rep., 2015.

[6] US Dept of State, "Trafficking in persons report 2013," 2013.

[7] M. L. Dank *et al.*, "Estimating the size and structure of the underground commercial sex economy in eight major us cities," 2014.

[8] M. Latonero, "Human trafficking online: The role of social networking sites and online classifieds," *Available at SSRN 2045851*, 2011.

[9] M. Latonero *et al.*, *The rise of mobile and the diffusion of technology-facilitated trafficking*. University of Southern California, Center on Communication Leadership . . ., 2012.

[10] A. Dubrawski, K. Miller, M. Barnes, B. Boecking, and E. Kennedy, "Leveraging publicly available data to discern patterns of human-trafficking activity," *Journal of Human Trafficking*, vol. 1, no. 1, pp. 65–85, 2015.

[11] D. Roe-Sepowitz, J. Gallagher, L. Martin, G. Snyder, K. Hickle, and J. Smith, "One-day sex trafficking snapshot of an internet service provider: Research brief," *Arizona State University*, 2012.

[12] R. S. Silberglitt, B. G. Chow, J. S. Hollywood, D. Woods, M. Zaydman, and B. A. Jackson, *Visions of law enforcement technology in the period 2024-2034: report of the Law Enforcement Futuring Workshop*. 2015.

[13] J. Deeb-Swihart, A. Endert, and A. Bruckman, "Understanding law enforcement strategies and needs for combating human trafficking," in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 2019, pp. 1–14.

[14] Office for Victims of Crime (OVC) and US Dept of Justice and Office of Justice Programs, "Coordination, collaboration, capacity: Federal strategic action plan on services for victims of human trafficking in the united states," 2016.

[15] U.S Department of Justice, "The national strategy for child exploitation prevention and interdiction," 2016.

[16] "Protocol to prevent, suppress and punish trafficking in persons, especially women and children, supplementing the united nations convention against transnational organized crime," in *Vienna: United Nations Office on Drugs and Crime, United Nations Convention against Transnational Organized Crime and the Protocols Thereto*, 2000.

[17] *United States of America: Victims of Trafficking and Violence Protection Act of 2000*, Public Law 106-386 [H.R. 3244], 2000.

[18] ——, "National Strategy to Combat Human Trafficking," US Department of Justice, Tech. Rep., 2017.

[19] R. S. Portnoff, D. Y. Huang, P. Doerfler, S. Afroz, and D. McCoy, "Backpage and bitcoin: Uncovering human traffickers," in *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ACM, 2017, pp. 1595–1604.

[20] M. Ibanez and D. D. Suthers, "Detection of domestic human trafficking indicators and movement trends using content available on open internet sources," in *System Sciences (HICSS), 2014 47th Hawaii International Conference on*, IEEE, 2014, pp. 1556–1565.

[21] A. J. Nichols, *Sex trafficking in the United States: Theory, Research, Policy, and Practice*. Columbia University Press, 2016.

[22] "2019 Data Report, The U.S. National Human Trafficking Hotline: More victims and survivors speaking for themselves," Polaris, Tech. Rep., Sep. 2019.

[23] L. N. Roberson, "She leads a lonely life: When sex trafficking and drug addiction collide," *Wake Forest L. Rev.*, vol. 52, p. 359, 2017.

[24] S. Phillips, C. Coates, C. Ortiz, L. Rast, J. Sheltry, and K. Thomas, "Clearing the Slate- Seeking Effective Remedies for Criminalized Trafficking Victims," Interna-

tional Women's Human Rights Clinic (IWHRC) at the City University of New York School of Law, Tech. Rep., 2014.

[25] T. Couture, "More than Drinks for Sale: Exposing Sex Trafficking in Cantinas & Bars in the U.S.," Polaris, Tech. Rep., Sep. 2016.

[26] D. Mletzko, L. Summers, and A. N. Arnio, "Spatial patterns of urban sex trafficking," *Journal of criminal justice*, vol. 58, pp. 87–96, 2018.

[27] U.S. Attorney's Office, FBI Houston Field Office, *Federal jury convicts sex trafficking ring leader*, Apr. 2015.

[28] K. Bales, S. Lize, and C. I. for International Studies, *Trafficking in persons in the United States, final report*. BiblioGov, 2012.

[29] U.S. Attorney's Office, FBI Washington Field Office, *Atlanta man admits his role in operating an interstate juvenille sex trafficking enterprise*, URL: https://archives.fbi.gov/archives/washingtondc/press-releases/2013/atlanta-man-admits-his-role-in-operating-an-interstate-juvenille-sex-trafficking-enterprise, Mar. 2013.

[30] V. Bouché *et al.*, "A report on the use of technology to recruit, groom and sell domestic minor sex trafficking victims," 2015.

[31] H. Wang, C. Cai, A. Philpot, M. Latonero, E. H. Hovy, and D. Metzler, "Data integration from open internet sources to combat sex trafficking of minors," in *Proceedings of the 13th Annual International Conference on Digital Government Research*, 2012, pp. 246–252.

[32] D. R. Silva, A. Philpot, A. Sundararajan, N. M. Bryan, and E. Hovy, "Data integration from open internet sources and network detection to combat underage sex trafficking," in *Proceedings of the 15th Annual International Conference on Digital Government Research*, 2014, pp. 86–90.

[33] E. Tong, A. Zadeh, C. Jones, and L.-P. Morency, "Combating human trafficking with multimodal deep models," in *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2017, pp. 1547–1556.

[34] H. B. Grant and K. J. Terry, *Law enforcement in the 21st century*. Pearson, 2012.

[35] J. Woodhams, C. R. Hollin, and R. Bull, "The psychology of linking crimes: A review of the evidence," *Legal and Criminological Psychology*, vol. 12, no. 2, pp. 233–249, 2007.

[36] N. Qazi and B. W. Wong, "An interactive human centered data science approach towards crime pattern analysis," *Information Processing & Management*, vol. 56, no. 6, p. 102 066, 2019.

[37] W. Jentner, D. Sacha, F. Stoffel, G. Ellis, L. Zhang, and D. A. Keim, "Making machine intelligence less scary for criminal analysts: Reflections on designing a visual comparative case analysis tool," *The Visual Computer*, vol. 34, no. 9, pp. 1225–1241, 2018.

[38] M. Tonkin, J. Woodhams, R. Bull, J. W. Bond, and E. J. Palmer, "Linking different types of crime using geographical and temporal proximity," *Criminal Justice and Behavior*, vol. 38, no. 11, pp. 1069–1088, 2011.

[39] O. Ribaux and P. Margot, "Case based reasoning in criminal intelligence using forensic case data," *Science & Justice*, vol. 43, no. 3, pp. 135–143, 2003.

[40] S. Dodge, R. Weibel, and A.-K. Lautenschütz, "Towards a taxonomy of movement patterns," *Information visualization*, vol. 7, no. 3-4, pp. 240–252, 2008.

[41] J. Buchmüller, E. Cakmak, N. Andrienko, G. Andrienko, J. W. Jolles, and D. A. Keim, "Moving together: Towards a formalization of collective movement," in *EuroVis Workshop on Visual Analytics (EuroVA)*, 2019.

[42] C. Piciarelli, C. Micheloni, and G. L. Foresti, "Trajectory-based anomalous event detection," *IEEE Transactions on Circuits and Systems for video Technology*, vol. 18, no. 11, pp. 1544–1554, 2008.

[43] Z. Sun and F. Wang, "A comparative study of features and distance metrics for trajectory clustering in open video domains," in *New Research in Multimedia and Internet Systems*, Springer, 2015, pp. 47–56.

[44] E. J. Keogh and M. J. Pazzani, "Scaling up dynamic time warping for datamining applications," in *Proceedings of the sixth ACM SIGKDD international conference on Knowledge discovery and data mining*, 2000, pp. 285–289.

[45] L. Chen, M. T. Özsu, and V. Oria, "Robust and fast similarity search for moving object trajectories," in *Proceedings of the 2005 ACM SIGMOD international conference on Management of data*, 2005, pp. 491–502.

[46] S. Salvador and P. Chan, "Toward accurate dynamic time warping in linear time and space," *Intelligent Data Analysis*, vol. 11, no. 5, pp. 561–580, 2007.

[47] M. Vlachos, G. Kollios, and D. Gunopulos, "Discovering similar multidimensional trajectories," in *Proceedings 18th international conference on data engineering*, IEEE, 2002, pp. 673–684.

[48] G. Zhu *et al.*, "Trajectory based event tactics analysis in broadcast sports video," in *Proceedings of the 15th ACM international conference on Multimedia*, 2007, pp. 58–67.

[49] Z. Y. Zheng Kai, N. YUAN, *et al.*, "On discovery of gathering patterns from trajectories," *ICDE, Tianjin, China*, 2013.

[50] P. Laube, S. Imfeld, and R. Weibel, "Discovering relative motion patterns in groups of moving point objects," *International Journal of Geographical Information Science*, vol. 19, no. 6, pp. 639–668, 2005.

[51] W. Wu *et al.*, "Telcovis: Visual exploration of co-occurrence in urban human mobility based on telco data," *IEEE transactions on visualization and computer graphics*, vol. 22, no. 1, pp. 935–944, 2015.

[52] L. Zhao, F. Ni, X. Jin, and L. Qin, "Social relationship modeling and community detection based on trajectory data," in *2019 IEEE 21st International Conference on High Performance Computing and Communications; IEEE 17th International Conference on Smart City; IEEE 5th International Conference on Data Science and Systems (HPCC/SmartCity/DSS)*, IEEE, 2019, pp. 2681–2687.

[53] P. Laube, M. van Kreveld, and S. Imfeld, "Finding remo—detecting relative motion patterns in geospatial lifelines," in *Developments in spatial data handling*, Springer, 2005, pp. 201–215.

[54] M. Andersson, J. Gudmundsson, P. Laube, and T. Wolle, "Reporting leadership patterns among trajectories," in *Proceedings of the 2007 ACM symposium on Applied computing*, 2007, pp. 3–7.

[55] P. Kalnis, N. Mamoulis, and S. Bakiras, "On discovering moving clusters in spatio-temporal data," in *International symposium on spatial and temporal databases*, Springer, 2005, pp. 364–381.

[56] Y. Huang, C. Chen, and P. Dong, "Modeling herds and their evolvements from trajectory data," in *International Conference on Geographic Information Science*, Springer, 2008, pp. 90–105.

[57] M. Steptoe, R. Krüger, R. Garcia, X. Liang, and R. Maciejewski, "A visual analytics framework for exploring theme park dynamics," *ACM Transactions on Interactive Intelligent Systems (TiiS)*, vol. 8, no. 1, pp. 1–27, 2018.

[58] J. Kim and H. S. Mahmassani, "Spatial and temporal characterization of travel patterns in a traffic network using vehicle trajectories," *Transportation Research Procedia*, vol. 9, pp. 164–184, 2015.

[59] N. Andrienko, G. Andrienko, L. Barrett, M. Dostie, and P. Henzi, "Space transformation for understanding group movement," *IEEE transactions on visualization and computer graphics*, vol. 19, no. 12, pp. 2169–2178, 2013.

[60] J. Yin, S. Karimi, A. Lampert, M. Cameron, B. Robinson, and R. Power, "Using social media to enhance emergency situation awareness," in *Twenty-fourth international joint conference on artificial intelligence*, 2015.

[61] S. Vieweg, A. L. Hughes, K. Starbird, and L. Palen, "Microblogging during two natural hazards events: What twitter may contribute to situational awareness," in *Proceedings of the SIGCHI conference on human factors in computing systems*, 2010, pp. 1079–1088.

[62] M.-H. Hwang, S. Wang, G. Cao, A. Padmanabhan, and Z. Zhang, "Spatiotemporal transformation of social media geostreams: A case study of twitter for flu risk analysis," in *Proceedings of the 4th ACM SIGSPATIAL International Workshop on GeoStreaming*, 2013, pp. 12–21.

[63] S. Chen *et al.*, "Interactive visual discovering of movement patterns from sparsely sampled geo-tagged social media data," *IEEE transactions on visualization and computer graphics*, vol. 22, no. 1, pp. 270–279, 2015.

[64] Y.-T. Zheng, Z.-J. Zha, and T.-S. Chua, "Mining travel patterns from geotagged photos," *ACM Transactions on Intelligent Systems and Technology (TIST)*, vol. 3, no. 3, pp. 1–18, 2012.

[65] B. Hecht, L. Hong, B. Suh, and E. H. Chi, "Tweets from justin bieber's heart: The dynamics of the location field in user profiles," in *Proceedings of the SIGCHI conference on human factors in computing systems*, 2011, pp. 237–246.

[66] J. Hays and A. A. Efros, "Im2gps: Estimating geographic information from a single image," in *2008 ieee conference on computer vision and pattern recognition*, IEEE, 2008, pp. 1–8.

[67] J. Wood, J. Dykes, and A. Slingsby, "Visualisation of origins, destinations and flows with od maps," *The Cartographic Journal*, vol. 47, no. 2, pp. 117–129, 2010.

[68] I. Boyandin, E. Bertini, P. Bak, and D. Lalanne, "Flowstrates: An approach for visual exploration of temporal origin-destination data," in *Computer Graphics Forum*, Wiley Online Library, vol. 30, 2011, pp. 971–980.

[69] S. Bremm, G. Andrienko, N. Andrienko, T. Schreck, and T. v. Landesberger, "Interactive analysis of object group changes over time," in *EuroVA 2011*, 2011.

[70] X. Huang, Y. Zhao, C. Ma, J. Yang, X. Ye, and C. Zhang, "Trajgraph: A graph-based visual analytics approach to studying urban network centralities using taxi trajectory data," *IEEE transactions on visualization and computer graphics*, vol. 22, no. 1, pp. 160–169, 2015.

[71] N. Andrienko and G. Andrienko, "State transition graphs for semantic analysis of movement behaviours," *Information Visualization*, vol. 17, no. 1, pp. 41–65, 2018.

[72] B. Bach, P. Dragicevic, D. Archambault, C. Hurter, and S. Carpendale, "A review of temporal data visualizations based on space-time cube operations," 2014.

[73] G. Andrienko, N. Andrienko, P. Bak, D. Keim, S. Kisilevich, and S. Wrobel, "A conceptual framework and taxonomy of techniques for analyzing movement," *Journal of Visual Languages & Computing*, vol. 22, no. 3, pp. 213–232, 2011.

[74] F. Reitz, "A framework for an ego-centered and time-aware visualization of relations in arbitrary data repositories," *arXiv preprint arXiv:1009.5183*, 2010.

[75] D. Archambault, H. Purchase, and B. Pinaud, "Animation, small multiples, and the effect of mental map preservation in dynamic graphs," *IEEE transactions on visualization and computer graphics*, vol. 17, no. 4, pp. 539–552, 2010.

[76] B. Tversky, J. B. Morrison, and M. Betrancourt, "Animation: Can it facilitate?" *International journal of human-computer studies*, vol. 57, no. 4, pp. 247–262, 2002.

[77] G. Andrienko, N. Andrienko, P. Bak, D. Keim, and S. Wrobel, *Visual analytics of movement*. Springer Science & Business Media, 2013.

[78] I. Boyandin, E. Bertini, and D. Lalanne, "A qualitative study on the exploration of temporal changes in flow maps with animation and small-multiples," in *Computer Graphics Forum*, Wiley Online Library, vol. 31, 2012, pp. 1005–1014.

[79] T. Von Landesberger, F. Brodkorb, P. Roskosch, N. Andrienko, G. Andrienko, and A. Kerren, "Mobilitygraphs: Visual analysis of mass mobility dynamics via spatio-temporal graphs and clustering," *IEEE transactions on visualization and computer graphics*, vol. 22, no. 1, pp. 11–20, 2015.

[80] S. van den Elzen, D. Holten, J. Blaas, and J. J. van Wijk, "Reducing snapshots to points: A visual analytics approach to dynamic network exploration," *IEEE transactions on visualization and computer graphics*, vol. 22, no. 1, pp. 1–10, 2015.

[81] Y. Hu, H. Mao, and G. McKenzie, "A natural language processing and geospatial clustering framework for harvesting local place names from geotagged housing advertisements," *International Journal of Geographical Information Science*, vol. 33, no. 4, pp. 714–738, 2019.

[82] R. Beecham and J. Wood, "Characterising group-cycling journeys using interactive graphics," *Transportation Research Part C: Emerging Technologies*, vol. 47, pp. 194–206, 2014.

[83] K. Buchin, B. Speckmann, and K. Verbeek, "Flow map layout via spiral trees," *IEEE transactions on visualization and computer graphics*, vol. 17, no. 12, pp. 2536–2544, 2011.

[84] O. Ersoy, C. Hurter, F. Paulovich, G. Cantareiro, and A. Telea, "Skeleton-based edge bundling for graph visualization," *IEEE transactions on visualization and computer graphics*, vol. 17, no. 12, pp. 2364–2373, 2011.

[85] D. Holten and J. J. Van Wijk, "Force-directed edge bundling for graph visualization," in *Computer graphics forum*, Wiley Online Library, vol. 28, 2009, pp. 983–990.

[86] S. Kisilevich, F. Mansmann, M. Nanni, and S. Rinzivillo, "Spatio-temporal clustering," in *Data mining and knowledge discovery handbook*, Springer, 2009, pp. 855–874.

[87] J. Buchmüller, D. Jäckle, E. Cakmak, U. Brandes, and D. A. Keim, "Motionrugs: Visualizing collective trends in space and time," *IEEE transactions on visualization and computer graphics*, vol. 25, no. 1, pp. 76–86, 2018.

[88] T. Von Landesberger, S. Bremm, T. Schreck, and D. W. Fellner, "Feature-based automatic identification of interesting data segments in group movement data," *Information Visualization*, vol. 13, no. 3, pp. 190–212, 2014.

[89] D. G. Wilson, W. F. Walsh, and S. Kleuber, "Trafficking in human beings: Training and services among us law enforcement agencies," *Police Practice and Research*, vol. 7, no. 02, pp. 149–160, 2006.

[90] A. J. Nichols and E. C. Heil, "Challenges to identifying and prosecuting sex trafficking cases in the midwest united states," *Feminist criminology*, vol. 10, no. 1, pp. 7–35, 2015.

[91] A. Farrell, J. McDevitt, and S. Fahy, "Understanding and improving law enforcement responses to human trafficking," 2008.

[92] K. J. Mitchell and D. Boyd, "Understanding the role of technology in the commercial sexual exploitation of children: The perspective of law enforcement.," 2014.

[93] E. Kennedy, "Predictive patterns of sex trafficking online," M.S. thesis, Carengie Mellon University, Apr. 2012.

[94] M. Latonero, B. Wex, and M. Dank, "Technology and labor trafficking in a network society: General overview, emerging innovations, and philippines case study," 2015.

[95] B. A. Jackson, V. Greenfield, A. R. Morral, and J. S. Hollywood, "Police department investments in information technology systems," 2014.

[96] K. Strom, *Research on the impact of technology on policing strategy in the 21st century, final report*, 2017.

[97] L. Garicano and P. Heaton, "Information technology, organization, and productivity in the public sector: Evidence from police departments," *Journal of Labor Economics*, vol. 28, no. 1, pp. 167–201, 2010.

[98] J. S. Hollywood and Z. Winkelman, *Improving Information-Sharing Across Law Enforcement: Why Can't We Know?* JSTOR, 2015.

[99] J. S. Hollywood, J. E. Boon Jr, R. Silberglitt, B. G. Chow, and B. A. Jackson, *High-priority information technology needs for law enforcement*. Rand Corporation, 2015.

[100] J. S. Hollywood *et al.*, *Fostering Innovation in US Law Enforcement*. Rand Corporation, 2017.

[101] J. S. Hollywood, D. Woods, A. Lauland, B. A. Jackson, and R. Silberglitt, *Addressing Emerging Trends to Support the Future of Criminal Justice*. RAND Corporation, 2018.

[102] J. B. Chan, "The technological game: How information technology is transforming police practice," *Criminal Justice*, vol. 1, no. 2, pp. 139–159, 2001.

[103] E. E. Joh, "Policing by numbers: Big data and the fourth amendment," *Wash. L. Rev.*, vol. 89, p. 35, 2014.

[104] N. Verma and L. Dombrowski, "Confronting social criticisms: Challenges when adopting data-driven policing strategies," in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, ACM, 2018, p. 469.

[105] S. Nunn, "Police information technology: Assessing the effects of computerization on urban police functions," *Public Administration Review*, vol. 61, no. 2, pp. 221–234, 2001.

[106] S. Brayne, "Big data surveillance: The case of policing," *American Sociological Review*, vol. 82, no. 5, pp. 977–1008, 2017.

[107] F. G. QC, J. Muraszkiewicz, and N. Vavoula, “The role of technology in the fight against human trafficking: Reflections on privacy and data protection concerns,” *Computer Law & Security Review*, vol. 32, no. 2, pp. 205–217, 2016.

[108] H. Clawson, N. Dutch, and M. Cummings, “Law enforcement response to human trafficking and the implications for victims: Current practices and lessons learned,” *Washington, DC: US Department of Justice, National Institute of Justice*, 2006.

[109] M. N. Marshall, “Sampling for qualitative research,” *Family practice*, vol. 13, no. 6, pp. 522–526, 1996.

[110] K. Herr and G. L. Anderson, *The action research dissertation: A guide for students and faculty*. Sage publications, 2014.

[111] R. E. Boyatzis, *Transforming qualitative information: Thematic analysis and code development*. sage, 1998.

[112] US Department of State, “Trafficking in persons report,” Tech. Rep., 2017.

[113] J. Xu and H. Chen, “Criminal network analysis and visualization,” *Communications of the ACM*, vol. 48, no. 6, pp. 100–107, 2005.

[114] M. T. Ribeiro, S. Singh, and C. Guestrin, “Why should i trust you?: Explaining the predictions of any classifier,” in *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, ACM, 2016, pp. 1135–1144.

[115] Z. Lipton, “The mythos of model interpretability,” in *ICML Workshop on Human Interpretability in Machine Learning (WHI 2016)*, New York, NY, USA, 2016.

[116] S. Gurman, “Across us, police officers abuse confidential databases,” *Associated Press*, Sep. 2016.

[117] H. Alvari, P. Shakarian, and J. K. Snyder, “A non-parametric learning approach to identify online human trafficking,” in *2016 IEEE Conference on Intelligence and Security Informatics (ISI)*, IEEE, 2016, pp. 133–138.

[118] L. Li, O. Simek, A. Lai, M. Daggett, C. K. Dagli, and C. Jones, “Detection and characterization of human trafficking networks using unsupervised scalable text template matching,” in *2018 IEEE International Conference on Big Data (Big Data)*, IEEE, 2018, pp. 3111–3120.

[119] J. Zhu, L. Li, and C. Jones, “Identification and detection of human trafficking using language models,” in *2019 European Intelligence and Security Informatics Conference (EISIC)*, IEEE, 2019, pp. 24–31.

[120] M. Hultgren, M. E. Jennex, J. Persano, and C. Ornatowski, “Using knowledge management to assist in identifying human sex trafficking,” in *2016 49th Hawaii International Conference on System Sciences (HICSS)*, IEEE, 2016, pp. 4344–4353.

[121] M. Ibanez and D. D. Suthers, “Detecting covert sex trafficking networks in virtual markets,” in *2016 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, 2016, pp. 876–879.

[122] M. Ibanez and R. Gazan, “Detecting sex trafficking circuits in the u.s. through analysis of online escort advertisements,” in *2016 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, 2016, pp. 892–895.

[123] M. Ibanez and D. D. Suthers, “Detection of domestic human trafficking indicators and movement trends using content available on open internet sources,” in *2014 47th Hawaii International Conference on System Sciences*, 2014, pp. 1556–1565.

[124] M. Hernández-Álvarez, “Detection of possible human trafficking in twitter,” in *2019 International Conference on Information Systems and Software Technologies (ICI2ST)*, 2019, pp. 187–191.

[125] M. Ibanez and R. Gazan, “Virtual indicators of sex trafficking to identify potential victims in online advertisements,” in *2016 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, 2016, pp. 818–824.

[126] K. Hundman, T. Gowda, M. Kejriwal, and B. Boecking, “Always lurking: Understanding and mitigating bias in online human trafficking detection,” in *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, ser. AIES ’18, New Orleans, LA, USA: Association for Computing Machinery, 2018, pp. 137–143, ISBN: 9781450360128.

[127] J. Brelsford and S. Parakh, “A systems modeling approach to analyzing human trafficking,” in *2018 Winter Simulation Conference (WSC)*, IEEE, 2018, pp. 12–21.

[128] E. Senft, B. Weeks, J. Palmer, B. Neely, B. Turner, and J. Caddell, “A systems dynamics approach to human trafficking in maharashtra, india,” in *2019 IEEE International Systems Conference (SysCon)*, IEEE, 2019, pp. 1–7.

[129] A. Tundis, A. Jain, G. Bhatia, and M. Muhlhauser, “Similarity analysis of criminals on social networks: An example on twitter,” in *2019 28th International Conference on Computer Communication and Networks (ICCCN)*, IEEE, 2019, pp. 1–9.

[130] R. S. Portnoff, D. Y. Huang, P. Doerfler, S. Afroz, and D. McCoy, “Backpage and bitcoin: Uncovering human traffickers,” in *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ser. KDD

'17, Halifax, NS, Canada: Association for Computing Machinery, 2017, pp. 1595–1604, ISBN: 9781450348874.

[131] H. R. Sharifzadeh, H. Zhang, and A. W.-K. Kong, "Vein pattern visualization through multiple mapping models and local parameter estimation for forensic investigation," in *2014 22nd International Conference on Pattern Recognition*, IEEE, 2014, pp. 160–165.

[132] F. Anda, D. Lillis, N.-A. Le-Khac, and M. Scanlon, "Evaluating automated facial age estimation techniques for digital forensics," in *2018 IEEE Security and Privacy Workshops (SPW)*, IEEE, 2018, pp. 129–139.

[133] S. Varastehpour, H. Sharifzadeh, I. Ardekani, and X. Francis, "Vein pattern visualisation and feature extraction using sparse auto-encoder for forensic purposes," in *2019 16th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS)*, IEEE, 2019, pp. 1–8.

[134] P. K. Chandaliya and N. Nain, "Conditional perceptual adversarial variational autoencoder for age progression and regression on child face," in *2019 International Conference on Biometrics (ICB)*, IEEE, 2019, pp. 1–8.

[135] F. Anda *et al.*, "Improving borderline adulthood facial age estimation through ensemble learning," in *Proceedings of the 14th International Conference on Availability, Reliability and Security*, ser. ARES '19, Canterbury, CA, United Kingdom: Association for Computing Machinery, 2019, ISBN: 9781450371643.

[136] P. Elzinga, K. E. Wolff, and J. Poelmans, "Analyzing chat conversations of pedophiles with temporal relational semantic systems," in *2012 European Intelligence and Security Informatics Conference*, IEEE, 2012, pp. 242–249.

[137] T. R. Ringenberg, K. Misra, and J. T. Rayz, "Not so cute but fuzzy: Estimating risk of sexual predation in online conversations," in *2019 IEEE International Conference on Systems, Man and Cybernetics (SMC)*, IEEE, 2019, pp. 2946–2951.

[138] M. A. Fauzi and P. Bours, "Ensemble method for sexual predators identification in online chats," in *2020 8th International Workshop on Biometrics and Forensics (IWBF)*, IEEE, 2020, pp. 1–6.

[139] B. Westlake, M. Bouchard, and R. Frank, "Comparing methods for detecting child exploitation content online," in *2012 European Intelligence and Security Informatics Conference*, IEEE, 2012, pp. 156–163.

[140] A. Vartapetiance and L. Gillam, "" our little secret": Pinpointing potential predators," *Security Informatics*, vol. 3, no. 1, pp. 1–19, 2014.

[141] A. Staff, *How amazon rekognition helps in the fight against some of the worst types of crime*, Nov. 2019.

[142] Marinus Analytics, *Traffic jam for facial recognition: Sex trafficking victim found online from 2 year old photo*, 2019.

[143] P. Krafft, M. Young, M. Katell, K. Huang, and G. Bugingo, "Defining ai in policy versus practice," in *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 2020, pp. 72–78.

[144] J. Fjeld, N. Achten, H. Hilligoss, A. Nagy, and M. Srikumar, "Principled artificial intelligence: Mapping consensus in ethical and rights-based approaches to principles for ai," *Berkman Klein Center Research Publication*, no. 2020-1, 2020.

[145] L. Andersen *et al.*, "Human rights in the age of artificial intelligence," *Access Now*, p. 29, 2018.

[146] A. HLEG, "High-level expert group on artificial intelligence: Ethics guidelines for trustworthy ai," *European Commission, 09.04*, 2019.

[147] R. Kirkham, "Using european human rights jurisprudence for incorporating values into design," in *Proceedings of the 2020 ACM Designing Interactive Systems Conference*, 2020, pp. 115–128.

[148] U. G. Assembly *et al.*, "Universal declaration of human rights," *UN General Assembly*, vol. 302, no. 2, pp. 14–25, 1948.

[149] T. Gillespie and N. Seaver, "Critical algorithm studies: A reading list," *Social Media Collective*, 2016.

[150] S. Chancellor, M. L. Birnbaum, E. D. Caine, V. M. Silenzio, and M. De Choudhury, "A taxonomy of ethical tensions in inferring mental health states from social media," in *Proceedings of the conference on fairness, accountability, and transparency*, 2019, pp. 79–88.

[151] M. Zimmer, ""but the data is already public": On the ethics of research in facebook," *Ethics and information technology*, vol. 12, no. 4, pp. 313–325, 2010.

[152] A. Bruns, K. Weller, M. Zimmer, and N. J. Proferes, "A topology of twitter research: Disciplines, methods, and ethics," *Aslib Journal of Information Management*, 2014.

[153] D. Boyd and K. Crawford, "Critical questions for big data: Provocations for a cultural, technological, and scholarly phenomenon," *Information, communication & society*, vol. 15, no. 5, pp. 662–679, 2012.

[154] P. Baker and A. Potts, "'why do white people have thin lips?' google and the perpetuation of stereotypes via auto-complete search forms," *Critical discourse studies*, vol. 10, no. 2, pp. 187–204, 2013.

[155] H. Nissenbaum, "How computer systems embody values," *Computer*, vol. 34, no. 3, pp. 120–119, 2001.

[156] J. Metcalf and K. Crawford, "Where are human subjects in big data research? the emerging ethics divide," *Big Data & Society*, vol. 3, no. 1, p. 2 053 951 716 650 211, 2016.

[157] L. Introna and D. Wood, "Picturing algorithmic surveillance: The politics of facial recognition systems," *Surveillance & Society*, vol. 2, no. 2/3, pp. 177–198, 2004.

[158] A. Završnik, "Algorithmic justice: Algorithms and big data in criminal justice settings," *European Journal of Criminology*, p. 1 477 370 819 876 762, 2019.

[159] J. Kleinberg, H. Lakkaraju, J. Leskovec, J. Ludwig, and S. Mullainathan, "Human decisions and machine predictions," *The quarterly journal of economics*, vol. 133, no. 1, pp. 237–293, 2018.

[160] M. Latonero, "Governing artificial intelligence: Upholding human rights & dignity," *Data & Society*, pp. 1–37, 2018.

[161] N. A. Smuha, "Beyond a human rights-based approach to ai governance: Promise, pitfalls, plea," *Philosophy & Technology*, pp. 1–14, 2020.

[162] E. Aizenberg and J. van den Hoven, "Designing for human rights in ai," *Big Data & Society*, vol. 7, no. 2, p. 2 053 951 720 949 566, 2020.

[163] L. McGregor, D. Murray, and V. Ng, "International human rights law as a framework for algorithmic accountability," *International & Comparative Law Quarterly*, vol. 68, no. 2, pp. 309–343, 2019.

[164] E. Massé, "Recommendations on privacy and data protection in the fight against covid-19," Access Now, Tech. Rep., 2020.

[165] J. Salerno, B. M. Knoppers, L. M. Lee, W. M. Hlaing, and K. W. Goodman, "Ethics, big data and computing in epidemiology and public health," *Annals of Epidemiology*, vol. 27, no. 5, pp. 297–301, 2017.

[166] M. K. Twis and K. Preble, "Intersectional standpoint methodology: Toward theory-driven participatory research on human trafficking," *Violence and victims*, vol. 35, no. 3, pp. 418–439, 2020.

[167] J. Pielemeier, “Ai & global governance: The advantages of applying the international human rights framework to artificial intelligence,” *Centre for Policy Research at United Nations University. Retrieved*, vol. 30, 2019.

[168] M. Broberg and H.-O. Sano, “Strengths and weaknesses in a human rights-based approach to international development–an analysis of a rights-based approach to development assistance based on practical experiences,” *The International Journal of Human Rights*, vol. 22, no. 5, pp. 664–680, 2018.

[169] S. Fukuda-Parr and E. Gibbons, “Emerging consensus on ‘ethical ai’: Human rights critique of stakeholder guidelines,” *Global Policy*, vol. 12, pp. 32–44, 2021.

[170] F. Wettstein, “Normativity, ethics, and the un guiding principles on business and human rights: A critical assessment,” *Journal of Human Rights*, vol. 14, no. 2, pp. 162–182, 2015.

[171] C. Okoli, “A guide to conducting a standalone systematic literature review,” *Communications of the Association for Information Systems*, vol. 37, no. 1, p. 43, 2015.

[172] A. Booth, A. Sutton, and D. Papaioannou, “Systematic approaches to a successful literature review,” 2016.

[173] J. Stoll, W. K. Edwards, and E. D. Mynatt, “Interorganizational coordination and awareness in a nonprofit ecosystem,” in *Proceedings of the 2010 ACM conference on Computer supported cooperative work*, 2010, pp. 51–60.

[174] ——, “Informal interactions in nonprofit networks,” in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2010, pp. 533–536.

[175] A. Strohmayer, M. Laing, and R. Comber, “Technologies and social justice outcomes in sex work charities: Fighting stigma, saving lives,” in *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, 2017, pp. 3352–3364.

[176] A. Gautam, C. Shrestha, A. Kulak, S. Harrison, and D. Tatar, “Participatory tensions in working with a vulnerable population,” in *Proceedings of the 15th Participatory Design Conference: Short Papers, Situated Actions, Workshops and Tutorial-Volume 2*, 2018, pp. 1–5.

[177] B. Kitchenham and S. Charters, “Guidelines for performing systematic literature reviews in software engineering,” 2007.

[178] M. Edwards, A. Rashid, and P. Rayson, “A systematic survey of online data mining technology intended for law enforcement,” *ACM Computing Surveys (CSUR)*, vol. 48, no. 1, pp. 1–54, 2015.

This resource was prepared by the author(s) using Federal funds provided by the U.S. Department of Justice. Opinions or points of view expressed are those of the author(s) and do not necessarily reflect the official position or policies of the U.S. Department of Justice.

[179] M. E. Kim, "Transformative justice and restorative justice: Gender-based violence and alternative visions of justice in the united states," *International review of victimology*, vol. 27, no. 2, pp. 162–172, 2021.

[180] U. D. of State, *Four "ps": Prevention, protection, prosecution, partnerships*, 2010.

[181] C. M. Ribando, "Trafficking in persons: Us policy and issues for congress," LIBRARY OF CONGRESS WASHINGTON DC CONGRESSIONAL RESEARCH SERVICE, 2007.

[182] S. Tomkins, G. Farnadi, B. Amanatullah, L. Getoor, and S. Minton, "The impact of environmental stressors on human trafficking," in *2018 IEEE International Conference on Data Mining (ICDM)*, IEEE, 2018, pp. 507–516.

[183] N. A. Giacobe *et al.*, "Characterizing sex trafficking in pennsylvania for law enforcement," in *2016 IEEE Symposium on Technologies for Homeland Security (HST)*, 2016, pp. 1–5.

[184] S. L. Granizo, Á. L. V. Caraguay, L. I. B. López, and M. Hernández-Álvarez, "Detection of possible illicit messages using natural language processing and computer vision on twitter and linked websites," *IEEE Access*, vol. 8, pp. 44 534–44 546, 2020.

[185] E. Torres, S. L. Granizo, and M. Hernandez-Alvarez, "Gender and age classification based on human features to detect illicit activity in suspicious sites," in *2019 International Conference on Computational Science and Computational Intelligence (CSCI)*, IEEE, 2019, pp. 416–419.

[186] S. Roshan, S. V. Kumar, and M. Kumar, "Project spear: Reporting human trafficking using crowdsourcing," in *2017 4th IEEE Uttar Pradesh Section International Conference on Electrical, Computer and Electronics (UPCON)*, 2017, pp. 295–299.

[187] H. Thinyane and K. S. Bhat, "Apprise: Supporting the critical-agency of victims of human trafficking in thailand," in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, ser. CHI '19, Glasgow, Scotland Uk: Association for Computing Machinery, 2019, pp. 1–14, ISBN: 9781450359702.

[188] H. Alvari, P. Shakarian, and J. K. Snyder, "Semi-supervised learning for detecting human trafficking," *Security Informatics*, vol. 6, no. 1, p. 1, 2017.

[189] M. Kejriwal and P. Szekely, "Information extraction in illicit web domains," in *Proceedings of the 26th International Conference on World Wide Web*, ser. WWW '17, Perth, Australia: International World Wide Web Conferences Steering Committee, 2017, pp. 997–1006, ISBN: 9781450349130.

[190] ——, “Technology-assisted investigative search: A case study from an illicit domain,” in *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems*, ser. CHI EA ’18, Montreal QC, Canada: Association for Computing Machinery, 2018, pp. 1–9, ISBN: 9781450356213.

[191] H. Wang, C. Cai, A. Philpot, M. Latonero, E. H. Hovy, and D. Metzler, “Data integration from open internet sources to combat sex trafficking of minors,” in *Proceedings of the 13th Annual International Conference on Digital Government Research*, ser. dg.o ’12, College Park, Maryland, USA: Association for Computing Machinery, 2012, pp. 246–252, ISBN: 9781450314039.

[192] A. Stylianou, A. Norling-Ruggles, R. Souvenir, and R. Pless, “Indexing open imagery to create tools to fight sex trafficking,” in *2015 IEEE Applied Imagery Pattern Recognition Workshop (AIPR)*, IEEE, 2015, pp. 1–6.

[193] A. Stylianou, J. Schreier, R. Souvenir, and R. Pless, “Traffickcam: Crowdsourced and computer vision based approaches to fighting sex trafficking,” in *2017 IEEE Applied Imagery Pattern Recognition Workshop (AIPR)*, 2017, pp. 1–8.

[194] D. Burbano and M. Hernández-Alvarez, “Illicit, hidden advertisements on twitter,” in *2018 International Conference on eDemocracy & eGovernment (ICEDEG)*, IEEE, 2018, pp. 317–321.

[195] S. Andrews, B. Brewster, and T. Day, “Organised crime and social media: A system for detecting, corroborating and visualising weak signals of organised crime online,” *Security Informatics*, vol. 7, no. 1, pp. 1–21, 2018.

[196] M. Kejriwal and R. Kapoor, “Network-theoretic information extraction quality assessment in the human trafficking domain,” *Applied Network Science*, vol. 4, no. 1, pp. 1–26, 2019.

[197] M. Kejriwal and Y. Gu, “Network-theoretic modeling of complex activity using uk online sex advertisements,” *Applied Network Science*, vol. 5, no. 1, pp. 1–23, 2020.

[198] G. J. Bala and S. L. Fernandes, “Developing novel skin detection on odroid xu4 heterogeneous multi-processing device,” in *Proceedings of the Second International Conference on Information and Communication Technology for Competitive Strategies*, ser. ICTCS ’16, Udaipur, India: Association for Computing Machinery, 2016, ISBN: 9781450339629.

[199] F. Anda, B. A. Becker, D. Lillis, N.-A. Le-Khac, and M. Scanlon, “Assessing the influencing factors on the accuracy of underage facial age estimation,” in *2020 International Conference on Cyber Security and Protection of Digital Services (Cyber Security)*, IEEE, 2020, pp. 1–8.

[200] J. Poelmans, P. Elzinga, S. Viaene, and G. Dedene, "Concept discovery innovations in law enforcement: A perspective," in *2010 International Conference on Intelligent Networking and Collaborative Systems*, IEEE, 2010, pp. 473–478.

[201] M. Liberatore, B. N. Levine, and C. Shields, "Strengthening forensic investigations of child pornography on p2p networks," in *Proceedings of the 6th International COnference*, ser. Co-NEXT '10, Philadelphia, Pennsylvania: Association for Computing Machinery, 2010, ISBN: 9781450304481.

[202] J. Shin, C. Ré, and M. Cafarella, "Mindtagger: A demonstration of data labeling in knowledge base construction," *Proc. VLDB Endow.*, vol. 8, no. 12, pp. 1920–1923, Aug. 2015.

[203] H. Z. Alshammari and K. S. Alghathbar, "Clogvis: Crime data analysis and visualization tool," in *Proceedings of the Second International Conference on Advanced Wireless Information, Data, and Communication Technologies*, ser. AWICT 2017, Paris, France: Association for Computing Machinery, 2017, ISBN: 9781450353106.

[204] C. Zhang, J. Shin, C. Ré, M. Cafarella, and F. Niu, "Extracting databases from dark data with deepdive," in *Proceedings of the 2016 International Conference on Management of Data*, ser. SIGMOD '16, San Francisco, California, USA: Association for Computing Machinery, 2016, pp. 847–859, ISBN: 9781450335317.

[205] E. Shaabani, H. Alvari, P. Shakarian, and J. K. Snyder, "Mist: Missing person intelligence synthesis toolkit," in *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management*, ser. CIKM '16, Indianapolis, Indiana, USA: Association for Computing Machinery, 2016, pp. 1843–1867, ISBN: 9781450340731.

[206] K. Pham, A. Santos, and J. Freire, "Learning to discover domain-specific web content," in *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*, ser. WSDM '18, Marina Del Rey, CA, USA: Association for Computing Machinery, 2018, pp. 432–440, ISBN: 9781450355810.

[207] R. Rabbany, D. Bayani, and A. Dubrawski, "Active search of connections for case building and combating human trafficking," in *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, ser. KDD '18, London, United Kingdom: Association for Computing Machinery, 2018, pp. 2120–2129, ISBN: 9781450355520.

[208] R. Kapoor, M. Kejriwal, and P. Szekely, "Using contexts and constraints for improved geotagging of human trafficking webpages," in *Proceedings of the Fourth International ACM Workshop on Managing and Mining Enriched Geo-Spatial Data*, ser. GeoRich '17, Chicago, Illinois: Association for Computing Machinery, 2017, ISBN: 9781450350471.

[209] M. Sastry and J. D. Sterman, “Desert island dynamics: An annotated survey of the essential system dynamics literature,” in *Proceedings of the 1993 Internatioanl System Dynamics Conference, Cancun, Mexico*, 1992, pp. 466–475.

[210] M. Goist, T. H. Y. Chen, and C. Boylan, “Reconstructing and analyzing the transnational human trafficking network,” in *2019 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, IEEE, 2019, pp. 493–500.

[211] M. da Silva Santos, M. Ladeira, G. C. G. Van Erven, and G. Luiz da Silva, “Machine learning models to identify the risk of modern slavery in brazilian cities,” in *2019 18th IEEE International Conference On Machine Learning And Applications (ICMLA)*, 2019, pp. 740–746.

[212] X. Li, G. M. Foody, D. S. Boyd, and F. Ling, “Aging brick kilns in the asian brick belt using a long time series of landsat sensor data to inform the study of modern day slavery,” in *IGARSS 2019-2019 IEEE International Geoscience and Remote Sensing Symposium*, IEEE, 2019, pp. 130–133.

[213] J. Shao, B. Du, C. Wu, and L. Zhang, “Tracking objects from satellite videos: A velocity feature based correlation filter,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 57, no. 10, pp. 7860–7871, 2019.

[214] P. Kostakos, L. Špráchalová, A. Pandya, M. Aboeleinen, and M. Oussalah, “Covert online ethnography and machine learning for detecting individuals at risk of being drawn into online sex work,” in *2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, IEEE, 2018, pp. 1096–1099.

[215] V. Rengaraj and K. Bijlani, “A study and implementation of smart id card with m-learning and child security,” in *2016 2nd International Conference on Applied and Theoretical Computing and Communication Technology (iCATccT)*, IEEE, 2016, pp. 305–311.

[216] D. Michalopoulos and I. Mavridis, “Towards risk based prevention of grooming attacks,” in *2010 International Conference on Security and Cryptography (SECRYPT)*, IEEE, 2010, pp. 1–4.

[217] ——, “Utilizing document classification for grooming attack recognition,” in *2011 IEEE Symposium on Computers and Communications (ISCC)*, IEEE, 2011, pp. 864–869.

[218] D. Michalopoulos, E. Papadopoulos, and I. Mavridis, “Artemis: Protection from sexual exploitation attacks via sms,” in *2012 16th Panhellenic Conference on Informatics*, IEEE, 2012, pp. 19–24.

[219] J. Mikka-Muntuumo, A. Peters, and H. Jazri, "Cyberbullet - share your story: An interactive game for stimulating awareness on the harm and negative effects of the internet," in *Proceedings of the Second African Conference for Human Computer Interaction: Thriving Communities*, ser. AfriCHI '18, Windhoek, Namibia: Association for Computing Machinery, 2018, ISBN: 9781450365581.

[220] A. Gautam, D. Tatar, and S. Harrison, "Crafting, communality, and computing: Building on existing strengths to support a vulnerable population," in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, ser. CHI '20, Honolulu, HI, USA: Association for Computing Machinery, 2020, pp. 1–14, ISBN: 9781450367080.

[221] W. Chung, E. Mustaine, and D. Zeng, "Criminal intelligence surveillance and monitoring on social media: Cases of cyber-trafficking," in *2017 IEEE International Conference on Intelligence and Security Informatics (ISI)*, IEEE, 2017, pp. 191–193.

[222] A. Ibrahim and M. V. Martin, "Addressing privacy constraints for efficient monitoring of network traffic for illicit images," in *2009 IEEE Toronto International Conference Science and Technology for Humanity (TIC-STH)*, IEEE, 2009, pp. 302–308.

[223] A. J. Nichols and A. Nichols, "Sex trafficking in the united states," in *Sex Trafficking in the United States*, Columbia University Press, 2017.

[224] N. Chokshi, "Facial recognition's many controversies, from stadium surveillance to racist software'," *New York Times*, vol. 15, 2019.

[225] Amnesty International, *Ban dangerous facial recognition technology that amplifies racist policing*.

[226] United Nations General Assembly, "Resolution adopted by the general assembly on 19 december 2016: The right to privacy in the digital age," *United Nations*, 2017.

[227] R. F. Jørgensen, *Human rights in the age of platforms*. The MIT Press, 2019.

[228] O. Shwartz, Y. Mathov, M. Bohadana, Y. Elovici, and Y. Oren, "Opening pandora's box: Effective techniques for reverse engineering iot devices," in *International Conference on Smart Card Research and Advanced Applications*, Springer, 2017, pp. 1–21.

[229] A. Toor, "Our identity is often what's triggering surveillance: How government surveillance of# blacklivesmatter violates the first amendments freedom of association," *Rutgers Computer & Tech. LJ*, vol. 44, p. 286, 2018.

[230] D. Lyon, *Surveillance society: Monitoring everyday life*. McGraw-Hill Education (UK), 2001.

[231] C. Beyrer and A. Kamarulzaman, “Ethnic cleansing in myanmar: The rohingya crisis and human rights,” *The Lancet*, vol. 390, no. 10102, pp. 1570–1573, 2017.

[232] S. Wolfinbarger, J. Drake, and E. Ashcroft, “Geospatial technologies and human rights project: Satellite imagery assessment of forced relocation near luiswishi mine,” 2014.

[233] A. Lichtman and M. Nair, “Humanitarian uses of drones and satellite imagery analysis: The promises and perils,” *AMA journal of ethics*, vol. 17, no. 10, pp. 931–937, 2015.

[234] R. Rothe, R. Timofte, and L. V. Gool, “Deep expectation of real and apparent age from a single image without facial landmarks,” *International Journal of Computer Vision*, vol. 126, no. 2-4, pp. 144–157, 2018.

[235] M. Young *et al.*, “Beyond open vs. closed: Balancing individual privacy and public accountability in data sharing,” in *Proceedings of the Conference on Fairness, Accountability, and Transparency*, 2019, pp. 191–200.

[236] J. D. Moorman and K. Harrison, “Gender, race, and risk: Intersectional risk management in the sale of sex online,” *The Journal of Sex Research*, vol. 53, no. 7, pp. 816–824, 2016.

[237] M. Kejriwal, J. Ding, R. Shao, A. Kumar, and P. Szekely, “Flagit: A system for minimally supervised human trafficking indicator mining,” *arXiv preprint arXiv:1712.03086*, 2017.

[238] R. Hurley *et al.*, “Measurement and analysis of child pornography trafficking on p2p networks,” in *Proceedings of the 22nd International Conference on World Wide Web*, ser. WWW ’13, Rio de Janeiro, Brazil: Association for Computing Machinery, 2013, pp. 631–642, ISBN: 9781450320351.

[239] T. Ali, R. Veldhuis, and L. Spreeuwers, “Forensic face recognition: A survey,” *Centre for Telematics and Information Technology, University of Twente, Tech. Rep. TR-CTIT-10-40*, vol. 1, 2010.

[240] D. Chaves, E. Fidalgo, E. Alegre, F. Jánez-Martino, and J. Velasco-Mata, “Cpu vs gpu performance of deep learning based face detectors using resized images in forensic applications,” in *9th International Conference on Imaging for Crime Detection and Prevention (ICDP-2019)*, 2019, pp. 93–98.

[241] U.S. Congress, *The Role of Technology in Countering Trafficking in Persons*, House Hearing before the Committee on Science, Space, and Technology in 116 Congress Second Session, 2020.

[242] K. Ringrose, "Law enforcement's pairing of facial recognition technology with body-worn cameras escalates privacy concerns," *Va. L. Rev. Online*, vol. 105, p. 57, 2019.

[243] J. Purshouse and L. Campbell, "Privacy, crime control and police use of automated facial recognition technology," *Criminal Law Review*, vol. 2019, no. 3, pp. 188–204, 2019.

[244] M. Minkler, "Ethical challenges for the "outside" researcher in community-based participatory research," *Health Education & Behavior*, vol. 31, no. 6, pp. 684–697, 2004.

[245] F. Kensing and J. Blomberg, "Participatory design: Issues and concerns," *Computer supported cooperative work (CSCW)*, vol. 7, no. 3, pp. 167–185, 1998.

[246] K. Countryman-Roswurm and B. Patton Brackin, "Awareness without re-exploitation: Empowering approaches to sharing the message about human trafficking," *Journal of Human Trafficking*, vol. 3, no. 4, pp. 327–334, 2017.

[247] G. R. Hayes, "Inclusive and engaged hci," *Interactions*, vol. 27, no. 2, pp. 26–31, 2020.

[248] International Labour Office, *Global Estimates of Modern Slavery: Forced Labour and Forced Marriage*. 2017.

[249] A. Benton, G. Coppersmith, and M. Dredze, "Ethical research protocols for social media health research," in *Proceedings of the First ACL Workshop on Ethics in Natural Language Processing*, 2017, pp. 94–102.

[250] K. S. Chmielinski *et al.*, "The dataset nutrition label (2nd gen): Leveraging context to mitigate harms in artificial intelligence," in *Proceedings of the NeurIPS 2020 Workshop on Dataset Curation and Security, Online*, vol. 11, 2020.

[251] E. Wall, L. M. Blaha, L. Franklin, and A. Endert, "Warning, bias may occur: A proposed approach to detecting cognitive bias in interactive visual analytics," in *2017 IEEE Conference on Visual Analytics Science and Technology (VAST)*, IEEE, 2017, pp. 104–115.

[252] E. Wall, J. Stasko, and A. Endert, "Toward a design space for mitigating cognitive bias in vis," in *2019 IEEE Visualization Conference (VIS)*, IEEE, 2019, pp. 111–115.

[253] L. Vollinger, "Concretizing intersectional research methods: Incorporating social justice and action into united states sex trafficking research," *Journal of Human Behavior in the Social Environment*, vol. 31, no. 5, pp. 599–625, 2021.

[254] S. Brayne, *Predict and surveil: Data, discretion, and the future of policing*. Oxford University Press, USA, 2020.

[255] M. Birkin, S. Dodge, B. T. Fasy, and R. P. Mann, "From observations to prediction of movement (dagstuhl seminar 18282)," in *Dagstuhl Reports*, Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, vol. 7, 2018.

[256] C. Ma, H. Lu, L. Shou, and G. Chen, "Ksq: Top-k similarity query on uncertain trajectories," *IEEE Transactions on Knowledge and Data Engineering*, vol. 25, no. 9, pp. 2049–2062, 2012.

[257] R. Maciejewski *et al.*, "A visual analytics approach to understanding spatiotemporal hotspots," *IEEE Transactions on Visualization and Computer Graphics*, vol. 16, no. 2, pp. 205–220, 2009.

[258] A. Malik, "Assisted decision making using multivariate spatiotemporal data through the application of visual analytics," Ph.D. dissertation, Purdue University, 2014.

[259] S. M. Radil, C. Flint, and G. E. Tita, "Spatializing social networks: Using social network analysis to investigate geographies of gang rivalry, territoriality, and violence in los angeles," *Annals of the Association of American Geographers*, vol. 100, no. 2, pp. 307–326, 2010.

[260] L. Anselin, S. J. Rey, and J. Kochinsky, "Flexible geospatial visual analytics and simulation technologies to enhance criminal justice decision support systems," *Crime Mapp*, vol. 3, no. 1, 2010.

[261] K. Kim, A. O.-N. E. Mohr, and A. Oglesby, "2016 law enforcement use of social media survey," *Report, International Association of Chiefs of Police and the Urban Institute*, 2017.

[262] N. Levine, "Crime mapping and the crimestat program," *Geographical analysis*, vol. 38, no. 1, pp. 41–56, 2006.

[263] U.S. Attorney's Office, Northern District of Texas, *U.S. Attorney's Office Shuts Down Website Promoting Prostitution and Sex Trafficking, Indicts Owner*, URL: https://www.justice.gov/usao-ndtx/pr/us-attorney-s-office-shuts-down-website-promoting-prostitution-and-sex-trafficking, Jun. 2020.

[264] Polaris, "Hidden in Plain Sight: How Corporate Secrecy Facilitates Human Trafficking in Illicit Massage Parlors," Polaris, Tech. Rep., Apr. 2018.

[265] A. Lex, N. Gehlenborg, H. Strobelt, R. Vuillemot, and H. Pfister, “Upset: Visualization of intersecting sets,” *IEEE transactions on visualization and computer graphics*, vol. 20, no. 12, pp. 1983–1992, 2014.

VITA

Julia Deeb-Swihart was born on July 25 in Newark, Delaware. She currently resides in Seattle and works at the University of Washington as an Assistant Teaching Professor in the School of Information. There she teaches data science, information visualization, and machine learning with a strong emphasis on ethics. She attended Georgia Tech and received a Bachelors of Science in Computer Science in 2012.