Master's Thesis (Masters in Computer Science)

Bridging Analytical Gaps in Human Trafficking Detection: An Integrated Pipeline Utilizing NLP, GNNs, and GIS based visualizations

Mehul Bhojraj Upase (s4017633)
m.b.upase@umail.leidenuniv.nl

August 2025

1st supervisor: Prof. Dr. Mirjiam Van Reisen
2nd supervisor: Prof. Dr. Lu Cao

Leiden Institute of Advanced Computer Science (LIACS)
Faculty of Science

liacs.leidenuniv.nl

Contents

Abstract

Human trafficking analysis needs to work with the data we have. This data consists of sensitive, qualitative narratives that should not be exposed. Existing studies provide useful elements like text analytics, network views, and operational research models. However, there is little information on a governance-first, end-to-end pipeline that converts raw uploads into standard, reproducible artifacts that others can use and regenerate without access to the raw content. This thesis aims to fill that gap with a privacy-preserving, artifact-centric system that (a) projects uploads into a single, tidy processed schema with only essential, non-sensitive fields like serialized IDs, normalized locations, token lists, and optional durations, and (b) creates three types of interoperable exports: typed and weighted social networks (PNG/HTML), semantic representations (RDF/TTL), and geospatial layers (GeoJSON). Each export is marked with parameters and linked to a minimal registry entry that includes the processed dataset ID, schema version, tool versions, and thresholds. The outcome is a workflow where figures and maps become executable artifacts. They are auditable, portable, and reproducible from the sanitized projection alone. Meanwhile, the raw narratives are never stored and only visible to their uploader.

Alignment to Research Questions (RQs) and Objectives, The pipeline is designed and evaluated specifically against the thesis RQs and Objectives:

1. **RQ1 — Interoperability (Objective O1/O3)** The processed schema serves as a single contract for all modules. Exports use open standards—RDF/TTL for semantic reuse and GeoJSON for cartographic reuse—so external tools can load them without needing adapters. Minimal, machine-actionable provenance (dataset ID, schema version, parameters) accompanies every export, allowing for traceable reuse and exact or near-exact regeneration.

2. **RQ2 — Governance: privacy, residency, access, accountability (Objective O2)** The system ensures processed-only persistence (no raw-file storage), owner-only raw visibility through role-based access controls, and non-egress operation (maps render from local GeoJSON without third-party geocoding). Provenance records support accountability without revealing raw content and meet data minimization and purpose-limitation principles.

3. **RQ3 — Pipeline Utility for practitioner tasks (Objective O4)** From the sanitized table alone, the application supports route reconstruction (ordered visits and optional time playback), hub identification (distinct-victim loads with seeded, reproducible layouts), and cross-case consolidation (shared places and standardized actor tokens). Outputs are decision-ready artifacts that can be inspected, shared, and regenerated.

4. **Portability and transparent limits (Objectives O5/O6)** Deterministic preprocessing (normalizing headers, type coercions, token rules) and careful toponym resolution (exact lookup to high-threshold fuzzy with unresolved flags) apply across similar datasets while avoiding false precision. Remaining constraints—such as curation of header maps, unresolved locations, and absence of bundled SHACL shapes—are documented with safe upgrade paths.

5. **Evaluation mapping to RQs (O1–O4)** Instead of focusing on model accuracy, the thesis provides artifact-centered evidence: (i) literature concordance, showing that behaviors align with established guidelines (interoperability, provenance, reproducibility, governance); (ii) a gap-closure analysis connecting implemented methods to identified research gaps; and (iii) self-audits that demonstrate storage and residency, access limitations, completeness of provenance, and deterministic regeneration. Each figure and table references the processed dataset ID, schema version, and parameters so readers can verify results without needing raw data.

0.1 Findings and contribution

The system answers the RQs positively within its scope. It produces interoperable (RDF/GeoJSON), privacy-preserving (processed-only, owner-scoped, non-egress), and reproducible (parameter-stamped, regenerable) artifacts that support the three practitioner tasks. Methodologically, the contribution is a compact pattern—deterministic transforms over a single processed schema with standards-based exports and provenance—that provides auditable analytics without raw-data storage. Practically, it offers a functional end-to-end workflow and simple design guidelines that teams can follow under strict governance constraints. Predictive prototypes (LSTM/GNN) were explored experimentally but excluded from the validated contribution due to unstable accuracy and ethical risks. They are documented transparently elsewhere in the thesis, but not used as evidence.

1 Introduction

1.1 Background

Human trafficking is one of the most widespread transnational crimes 1. It affects millions of people and generates around $150 billion in annual profit. Human trafficking includes many kinds of crimes—forced labor, sexual exploitation, and servitude, as well as trafficking in children, for instance—that span a very large and complex problem. Law enforcement across the world, from every kind of government, is struggling to understand not just the networks of modern trafficking but also the many operations that make up this very large problem. To do this effectively, lasting interventions, they need new, scalable data-driven analytical solutions.

Figure 1 shows how large & widespread human trafficking networks can get, and how some locations are more active than others, and why they become nodes or hotspots of human trafficking. Also it visualizes more common routes and active areas

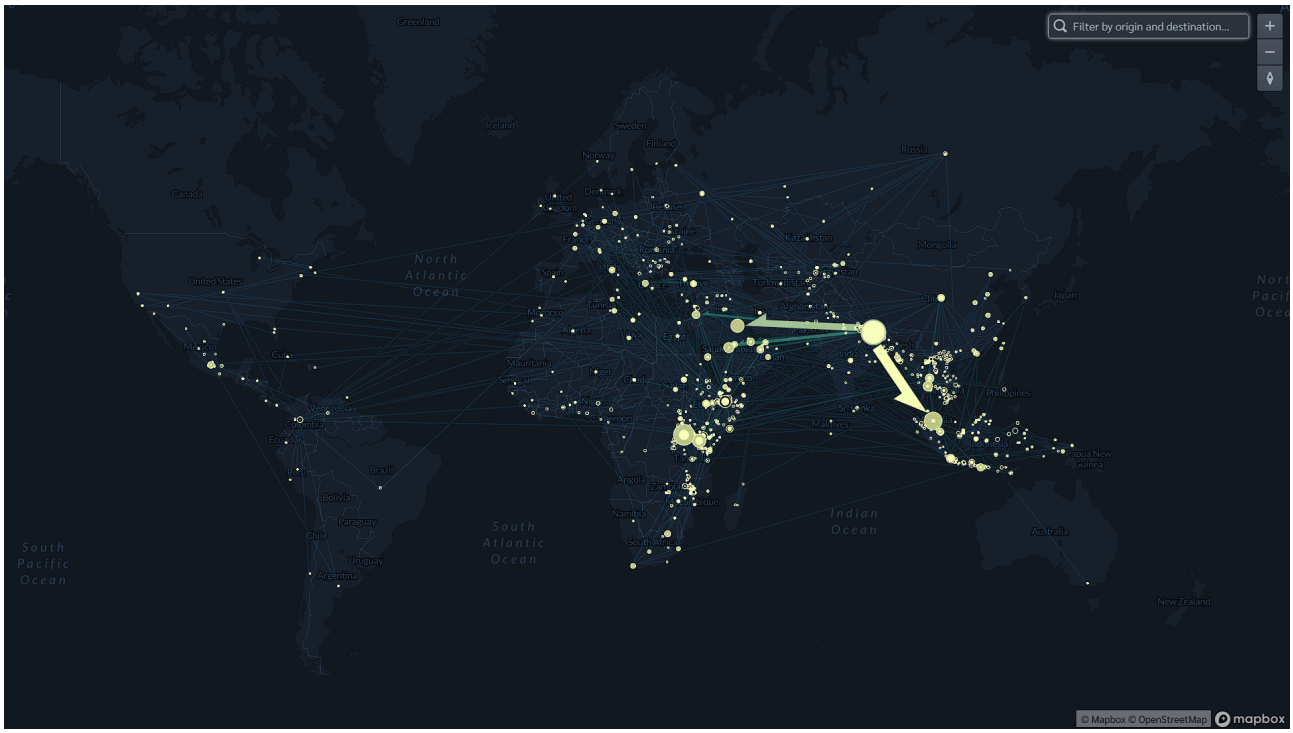

Figure 1. *Human Trafficking Network on a Large Public Dataset as visualized by https://www.freedomcollaborative.org/data-dashboard-vjm*

Artificial Intelligence (AI), Machine Learning (ML), and network analytics have transformed the detection of human trafficking into one that can find trafficked individuals hidden among the vast amounts of structured and unstructured data we now have. In recent studies, Dimas et al. (2021) and Konrad et al. (2022) effectively used Operations Research and Analytics to solve some problems related to human trafficking. They, along with other scholars, have shown the value of statistical modeling, network science, and decision-support systems in the fight against human trafficking. Similarly, Poelmans et al. (2012) have used semi-automated

knowledge discovery methods to unravel hidden connections between traffickers and victims. Szakonyi et al. (2021) on the other hand, have focused on technological interventions to detect such patterns. However, significant limitations in these studies necessitate a shift towards more integrated, real-time and scalable solutions.

1.2 Analytical Challenges in Current Human Trafficking Research

The analytical problems linked to human trafficking mean that effective interventions and responses are currently hampered. This section pulls together the most important and relevant discussions to provide an overview of the situation in current international and national trafficking research:

- **Fragmentation and inconsistency of available data sources:**

  Data associated with trafficking comes from numerous, disparate places, including law enforcement reports, NGO databases, online platforms, and victim testimonies. These sources significantly differ from each other, creating major inconsistencies and posing challenges to integration Konrad et al. (2022); Dimas et al. (2021). Incomplete, outdated, or biased data adds another layer of complexity to the problem. A lack of accurate information by trafficking activity and the under-reporting of trafficking by its victims leads to an apparent inconsistency in the data of trafficking United Nations Office on Drugs and Crime (2024); International Organization for Migration (2018); Ediae et al. (2024).

- **Under-utilization of qualitative data such as victim interviews and narratives:**

  Qualitative data, especially survivor testimonies, offer valuable insights into trafficker behaviors, victim dynamics, and recruitment strategies. However, these narratives often remain under-utilized due to analytical complexities and a lack of standardized methodologies to systematically structure qualitative insights Poelmans et al. (2012); Szakonyi et al. (2021). Despite their potential to enrich quantitative analyses, these datasets frequently remain isolated Latonero et al. (2012); Ibanez and Suthers (2014b); Konrad et al. (2022). Survivor testimonies that offer qualitative data provide most valuable insights into trafficker behaviors, victim dynamics, and recruitment strategies, because these are not covered in other forms of intelligence. Yet despite these insights, the narratives often remain under-utilized because of analytical complexities and a lack of standardized methods for systematizing the structuring of such qualitative insights Poelmans et al. (2012); Szakonyi et al. (2021). Hence, we have the qualitatively rich datasets that can enrich quantitative analyses, but still remain isolated Latonero et al. (2012); Ibanez and Suthers (2014b); Konrad et al. (2022).

- **Lack of fully integrated end-to-end analytical pipelines:**

  Current research typically employs isolated methodologies such as machine learning classification Ibanez and Suthers (2014b), predictive analytics **?**, or social network analysis Konrad et al. (2022); Dimas et al. (2021). Few studies integrate these methodologies into comprehensive end-to-end workflows. This gap results in manual data transfers between

disparate tools, reduces analytical responsiveness, and limits accessibility to non-specialist stakeholders Mansour et al. (2016); Wilkinson et al. (2016b).

- **Limited interoperability and semantic standardization:**

  Meaningful interoperability between organizations and across jurisdictions can be significantly restricted by human trafficking dataset semantic heterogeneity. Datasets collected by different agencies or in different geographic contexts often use incompatible terminologies and standards, thus hindering integration and collaborative analysis Wilkinson et al. (2016b); Mansour et al. (2016). In the absence of standardized semantic frameworks, efforts towards collaborative international action remain limited Konrad et al. (2022); Binu et al. (2019).

- **Ethical and privacy concerns restricting effective data sharing:**

  Confidentiality, informed consent, and victim protection are ethical matters that significantly limit data sharing and collaborative analytics. In the case of human trafficking, the data involved is to be highly sensitive. Therefore, strict privacy protocols are in place to prevent not only the further victimization of individuals but also the direct and indirect harm that may come from the misuse of this information United Nations Office on Drugs and Crime (2024); International Organization for Migration (2018); Mansour et al. (2016). Current practices, however, do not yet seem to adequately ensure all of these ethical challenges are properly addressed, limiting stakeholder collaboration and analytical effectiveness Binu et al. (2019); Wilkinson et al. (2016b).

- **Importance of provenance, traceability, and trustworthiness of the data:**

  In global criminal justice, data provenance, i.e., determining the how, what, when, where, and why of data, ensures the integrity and authenticity of the data required for legal and analytical purposes. For data to be useful in legal processes or credible for analytical functions, it must be trustworthy. Provenance mechanisms help ensure this by making compliance with international standards more likely, increasing trust among users, and making it easier to establish accountability among global anti-trafficking stakeholders Hasan and Myagmar (2022).

- **Significance of data sensitivity and high security risks:**

  Sensitive information is contained within human trafficking datasets, the compromise of which could also endanger victims, witnesses, or investigators and jeopardize ongoing interventions. Such datasets are particularly important to organized crime entities, making robust cybersecurity measures to safeguard their information against unauthorized access and misuse necessary The White House (2011); International Federation of Accountants (2021).

- **Personal data protection to safeguard involved individuals:**

  Critical ethical and legal issues are involved in the way personal data are handled in human trafficking research. Breaches could result in severe repercussions, including victim retaliation and social stigma. Compliance with a data protection regime that is rigorous, for example anonymization practices and secure storage solutions, is necessary to uphold

the rights and safe conditions of all individuals involved at each stage of the research Regent University (2022); World Bank (2021).

1.3 Literature Review

Introduction to Human Trafficking Research

The human trafficking research field has progressed rapidly in recent years, especially in the past two decades. Data science, artificial intelligence, and advances in network analysis have all been enlisted in the fight against trafficking, helping to modernize both detection and intervention efforts, as well as policy formation. But despite all this progress, literature on human trafficking remains fractured. Research often consists of solitary studies that focus on a particular set of analytical methods, or what we might call "work in the dark," where studies either don't employ any structured data or rely on datasets with such limited contextual depth that any findings are essentially meaningless.

Data-Driven Approaches in Trafficking Research

Technology's Role in Human Trafficking Detection

The spread of digital platforms and mobile technology has completely transformed human trafficking, enabling speedier recruitment, anonymous communication, and extensive victim exploitation. Latonero et al. (2012) provided one of the first examinations of this change, analyzing how mobile apps, gaming platforms, and social media are used to facilitate the crime. They found, not surprisingly, that traffickers use prepaid phones, but also found that the organized networks of traffickers communicate and avoid detection via social media and online classified ads. Despite these and other keen insights, the authors did not have the benefit of real-time analytical tools to track, with any precision, the God's-eye movement of victims forced into the trade across digital and physical domains.

In a like manner, Ibanez and Suthers (2014b) targeted automated web scraping techniques to extract trafficking-related indicators from escort advertisements and classified postings. The bulk of their work consisted of extracting these kinds of indicators from these postings using automated web scraping techniques and then performing some sorts of analysis on the indicators that they managed to scrape together. Phone number tracking methodologies were introduced to map the kinds of trends that are being missed by appearances in the periodic census of street prostitutes and those contained in the escort service advertising business.

Predictive Analytics and Machine Learning Models

Current research utilizes AI-driven machine learning models for more efficient analyses of human trafficking patterns. In their 2019 study, Sabo (2019) created an AI framework using the SAS Viya system to detect trafficking discourse in social media posts and classified advertisements. They used Natural Language Processing (NLP) techniques as part of their work and demonstrated the power of AI to sort through and classify vast quantities of text with a high degree of accuracy. They taught the AI to recognize the subtle signals that, in conjunction with other signals, could be used to indicate that human trafficking was actually happening. This is the solid foundation upon which we can build next-generation trafficking analyses.

In a separate piece of research, Ediae et al. (2024) looked into predictive analytics and how it could be used in the victim reintegration and trafficking prevention space. They took risk factors associated with trafficking victim post-rescue reintegration and applied them to Random Forest models and Decision Trees to assess predictive outcomes. The study had useful implications; however, like many others discussed in this chapter, it fell short on some level because it did not consider the way many law enforcement agencies fail to consider how traffickers might adjust their strategies in response to the kinds of interventions studied.

Network Analysis and Trafficking Ecosystem Modeling

Methodologies based on graphs have come to be very important in unveiling the structural makeup of trafficking networks. Konrad et al. (2022) explored the application of operations research techniques to the procurement of trafficking supply chains, using network interdiction models to optimize resource allocation for investigations into trafficking. However, these same authors did not venture into the use of network structures that change over time, which would be a way to effectively address trafficking networks that change in order to avoid being caught.

Dimas et al. (2021) also provided a comprehensive analysis of trafficking networks, but they used for this analysis social network analysis (SNA) techniques. They found that traffickers act as close to central nodes in the graph structure that represent trafficking rings. These findings do illustrate the power of graph analytics when used to identify where and who close to the top of the organizational structure.

Conversely, Binu et al. (2019) examined the application of fuzzy graph models to study the movement patterns of trafficking routes, integrating indices assessing connectivity to pinpoint trafficking hotspots. While offering an undoubtedly unique mathematical approach, their study lacks the empirical validation of real-world trafficking datasets, which would make their findings more relatable and operational.

Integration of AI & Deep Learning in Human Trafficking Detection

The recent development in artificial intelligence (AI) and deep learning has transformed human trafficking research into automated detection techniques that can process huge datasets. Sabo (2019) demonstrated the effectiveness of machine learning classifiers in detecting trafficking patterns across multiple online platforms and used Natural Language Processing (NLP) algorithms to identify coded trafficking language in online ads. Despite its successes, the study was limited to text-based sources and lacked integration with network-based relationship modeling; as a consequence, it could not recover the dynamics between traffickers and victims.

The same way, Ediae et al. (2024) used Random Forest, Decision Trees, and Neural Networks to analyze the results of rehabilitation for trafficking victims, pulling out some of the key predictors that make for successful reintegration. They worked on a victim-centric basis, and so didn't model the trafficking supply chain itself or pursue the movement of perpetrators across geography. Still, this kind of work very much establishes the realization of AI-driven analytics, and how they could serve a useful purpose if harnessed in the right sort of integrated

framework.

Pattern recognition and anomaly detection have been powerful use cases for deep learning. One specific application of these uses is for identifying online advertisements that depict human trafficking. Szakonyi et al. (2021) utilized computer vision techniques to detect delinking imagery in online ads. Their work, however, lacks an analytical model that can associate the visual patterns they've detected with a larger trafficking network, meaning this work is only applicable at the level of individual ad detection. Moving forward, research should harness the true power of multimodal analytics, integrating text, image, and network analytics to produce a more comprehensive detection capability.

Spatio-Temporal Modeling & Geographic Intelligence

It is crucial for law enforcement and policy interventions to understand how trafficking networks develop through time and across space.Dimas (2021) Dimas et al. (2021) applied social network analysis (SNA) to map trafficking rings but did not incorporate spatial intelligence—a key limitation that restricts the model's applicability to the real world. When combined with AI-based risk prediction models, geospatial analytics can significantly boost the detection of emerging trafficking routes.

Binu et al. (2019) Binu et al. (2019) spatially and temporally modeled something in an attempt to understand trafficking networks better. The authors understand that, even with the models they created, the trafficking networks could still be figured out by bad actors, and, as a result, those networks could still work and be detrimental to people. Binu et al. (2019) Binu et al. (2019) work is smart, and the work they did, even if not fully validated, contributes to understanding trafficking on a global scale better. Another significant research effort, Konrad et al. (2022), focused on deriving an approach from Operations Research (OR) to optimize counter-trafficking interventions. Their work centered on the challenge that law enforcement agencies have in allocating resources when the agencies want to make counter-trafficking work more effective. And what they came up with did a good work of identifying places that are likely to be high-risk trafficking zones and even forecasted somewhat reasonably where trafficking could go next. We think, however, that future research could capitalize more on this work and make it even more useful.

Ethical Considerations & Data Privacy in AI-Based Trafficking Detection

Ethical issues related to privacy, the potential for surveillance, and the biases that AI-driven data analytics might have when it comes to detection are becoming even more critical to address Bor and Koech (2023). We are now relying on these systems more and more, and their use is spreading rapidly—especially at the border. AI systems can provide useful information, especially when used in conjunction with other types of data, and particularly when it comes to prediction and human trafficking. But there's a dangerous notion of AI as the "new border." Because if we can't see through the algorithms, how can we guide such fundamentally transformative systems in secure and reliable ways? Another problem associated with this is managing victim data, especially sensitive survivor interviews used in machine learning pipelines. Poelmans

et al. (2012) Poelmans et al. (2012) stressed the need for tightly controlled data-sharing protocols to avoid re-traumatizing and securing the identities of those who survive such events. They also propose some solutions—federated learning and blockchain-based encryption—that ensure both privacy and data integrity.

Additionally, Sabo et al. (2019) Sabo (2019) examined automated AI systems that misidentify trafficking cases and concentrate on understanding social media activities. Their research revealed that word embeddings within natural language processing models might misclassify legal escort services as trafficking. Indeed, this work indicates to us that AI systems need greater refinements in both their linguistic models and, just as importantly, in the bias-mitigation techniques that are employed. For future endeavors, Sabo and her team suggested using explainable AI frameworks to keep AI-driven trafficking detection systems accountable.

Cross-Sector Collaboration and Multi-Agency Data Integration

One of the most significant challenges in combatting human trafficking is the lack of standardized data-sharing mechanisms across law enforcement, NGOs, financial institutions, and government agencies. Studies such as Konrad et al. Konrad et al. (2022) (2022) have underscored the importance of Operations Research (OR) models in optimizing resource allocation for counter-trafficking interventions, yet they do not address how real-time intelligence sharing between multiple entities can improve enforcement efforts. Dimas et al. (2021) Dimas et al. (2021) highlight the potential of federated learning models, which allow agencies to train AI models on decentralized datasets without compromising data privacy, though their implementation remains in early stages.

Multiple studies stress the very real need for cross-border intelligence sharing, especially when it comes to trafficking networks that operate across several jurisdictions. Notably, this includes the work of Bor and Koech Bor and Koech (2023). They, along with others, argue that, for the problem of trafficking to be effectively addressed, intelligence sharing between agencies will have to improve. And to do so, they recommend that future research focus on developing new ways to share information securely and across institutional boundaries, a problem they see as solvable with blockchain technology.

Financial Crime Analysis and Trafficking Networks

Human trafficking frequently intersects with money laundering and illicit financial transactions. The role of banking data, cryptocurrency transactions, and cross-border financial flows in trafficking networks remains underexplored in empirical research, despite the availability of helpful financial forensic techniques. Szakonyi et al. (2021) Szakonyi et al. (2021)suggested that financial institutions should be more active in detecting unusual transaction patterns indicative of trafficking operations. Such patterns might include the frequent making of small deposits, followed by bulk withdrawals in known trafficking hubs. However, the authors of this study did not propose a machine learning-based risk scoring system that can systematically flag these suspicious transactions.

Ediae et al. Ediae et al. (2024) scrutinized using AI-powered anomaly detection models to trace illicit financial transactions. They uncovered that traffickers conceal their financial footprints, often employing cryptocurrency tumblers and shell companies. Binu et al. Binu et al. (2019) likewise explored using fuzzy logic models to flag unusual transaction behaviors linked to trafficking operations. Although these studies add to the nexus of financial intelligence and trafficking analytics, they do not provide a path for integrating law enforcement forensic analysis with banking transaction monitoring.

The Role of Dark Web Intelligence in Trafficking Prevention

Although most existing research concentrates on surface web data, very few studies have investigated trafficking operations on the dark web. One of the few to even attempt such a study, Sabo (2019) **?**, introduced early-stage machine-learning models for this purpose. However, while discussing their use and the authors' results, it became evident that the models struggle with operationalization for several key reasons. the lack of real-time monitoring and inadequate classification techniques made it difficult to operationalize their findings.

A study by Poelmans et al. Poelmans et al. (2012) leveraged Formal Concept Analysis (FCA) to structure unclassified dark web discussions into meaningful categories, revealing how traffickers use coded language and hidden communication channels. Despite this, their study lacked real-time analytical tools to monitor evolving dark web transactions. Future research must focus on integrating AI-driven Natural Language Processing (NLP) models with forensic web crawling techniques to detect emerging trafficking hotspots in darknet forums.

Addressing Survivors' Reintegration & Post-Trafficking Support

An area that has been mostly overlooked in trafficking research is the long-term effect on survivors and their integration back into society. While studies such as Ediae et al. (2024) Ediae et al. (2024) focus on predicting victim outcomes during rehabilitation, there is scant research on how trafficked persons manage the unsafe challenges they face after rescue, including court cases, slippery social reputations, and money problems.

A resource optimization model for allocating funds to survivor support programs was introduced by Konrad et al. Konrad et al. (2022) in 2022, but it did not account for psychological recovery pathways. Dimas et al. (2021) Dimas et al. (2021) highlighted the importance of vocational training and education for survivors, but their research lacked longitudinal tracking mechanisms to assess the success of reintegration programs over time.

1.4 Research Gap

For a long time, an overwhelming amount of research on human trafficking relied on very structured datasets like the Counter-Trafficking Data Collaborative (CTDC) or the reports of the UNHCR. Most of this work paid little attention to the qualitative data found in victim interviews. But this is a crucial piece of the puzzle if we want to understand human trafficking. What we too often miss is that the dynamics at play between the trafficker and the victim are tough to grasp without this kind of close, narrative-based understanding. Migrant testimonies give us clues about trafficker relationships, recruitment strategies, and movement patterns. But our research is fundamentally limited because we are not using these narratives for Network Discovery. Poelmans et al. (2012); Sabo (2019).

Also, an end-to-end analytical pipeline is not yet in place that connects the processes of data preprocessing, social network based trafficking network modeling, and dynamic spatio-temporal visualization, query & insights and ontology creation, all this while strictly adhering to data security and residency requirements, since we are handling very sensitive data and data should reside with the data owner, this approach uses the data to extract information or keywords using NLP and then store the processed 'essential' data that doesn't have any sensitive information and the original dataset is instantly flushed from system. Current efforts apply what can best be described as isolated techniques: social network analysis (SNA), machine learning classification, and predictive analytics. They are sacrificing comprehensiveness and unity for, quite ironically, some dubiously attained virtues of real-time intelligence processing. Furthermore, the graph network models employed in most studies prior to ours are static and fail to adapt to evolving networks. For example, in modeling trafficker-victim interactions, Dimas et al. (2021) Dimas et al. (2021) and Konrad et al. (2017) Konrad et al. (2017) constructed frameworks lacking both temporal adaptability and predictive capabilities. Yet, traffickers continuously alter their routes, have recruits on the move, and shift strategies Binu et al. (2019);Konrad et al. (2017).

The trafficking research field lacks methods and tools of spanning real-time, cross-border intelligence—leading to significant gaps in knowledge of the spatial and temporal dimensions of the trafficking phenomenon.Though some studies (e.g., Binu et al. (2019)) have attempted to use GIS or other mapping techniques to visualize the trafficking phenomenon and identify emerging hotspots, trafficking research has not yet harnessed the full power of these tools for intelligence purposes.Lastly, predictive analytics in trafficking detection remains limited. While Ediae et al. Ediae et al. (2022) and Szakonyi et al. Szakonyi et al. (2021) applied AI-driven classification models, deep learning-based forecasting techniques such as LSTMs and Transformers have yet to be fully utilized to anticipate emerging trafficking routes Latonero (2012); Ibanez and Suthers (2014a).

Law enforcement and policymakers clearly have a need for such tools, especially dashboards that can indicate movements across borders as well as predict where the trafficking problem will emerge next. If trafficking research was able to utilize these tools to span the projected trafficking problem, it could be anticipating the traffickers' next moves and serving the law enforcement community and trafficking victim NGOs in a proactive manner.

These problems are made worse by the large-scale ethical and privacy issues associated

with AI-driven detection systems. Poorly constructed algorithms can lead to biased results; for example, if a system determines that the overall 'normal' pattern of behavior for a particular state is that 22% of its citizens are armed and dangerous, then any profiling the system does will carry the bias of that false positive Sales et al. (2023). This gap highlights the necessity of developing transparent, accountable, and ethically sound AI models to ensure compliance with rigorous privacy standards, such as GDPR, thus safeguarding sensitive victim data Sales et al. (2023);Olisah et al. (2024).

What is an ontology? An ontology is a formal representation of a set of concepts within a domain and the relationships between those concepts. An ontology provides a commonly understood vocabulary that can be used to underpin a knowledge representation in that domain. In computer science and information science, an ontology can be used for semantic reasoning. Sales et al. (2023)

In addition, centralized data storage is now raising very serious trust and privacy issues, which is hobbling all kinds of collaboration across different sectors. This collaboration is essential if we're going to address the many intertwined challenges—and the potential solutions to them—that involve and affect very different parts of society. These are just some of the reasons why decentralized storage systems like SOLID Pods offer a promising alternative. Solid Project (2024).

In conclusion, the identified gaps collectively indicate the critical need for an integrated, end-to-end analytical framework that systematically addresses data fragmentation, incorporates dynamic modeling and real-time predictive analytics, resolves semantic heterogeneity through ontology-driven data structuring, and employs secure decentralized storage solutions. Addressing these interconnected gaps is essential to developing actionable, scalable, and ethically responsible solutions for global anti-trafficking efforts.

1.5 Synthesis of Research Gaps

The scholarly and practitioner literature agrees that progress in anti-trafficking analytics has been hindered more by systemic and architectural gaps than by isolated technical limitations. Studies across different fields often focus on optimizing single components of the pipeline, such as text classification, static social network analysis, or map-based hotspot detection. However, they do not ensure interoperability, temporal accuracy, or governance-aware deployment. The following subsubsections outline the main gaps that inspire this thesis and shape the research questions and methods that follow.

1.5.1 Fragmented analytical approaches and the absence of end-to-end integration

Current work often views ingestion, normalization, semantic modeling, inference, and visualization as separate processes. This split limits cumulative science and reproducibility. Without a clear path from unstructured narratives to semantically governed knowledge representations and then to decision-making analytics, results are hard to validate and even harder to implement across institutions. A review of past systems shows that end-to-end integration is more the exception than the norm. There is limited focus on artifact consistency, configuration management, and re-runnable systems needed for independent verification United Nations Office on Drugs and Crime (2023a).

1.5.2 Under-utilization of unstructured, narrative sources

Qualitative, text-rich sources, such as survivor testimonies, practitioner reports, and investigative briefs, contain important context, timing, and uncertainty needed to understand methods of operation and movement patterns. However, modern transformer-based NLP and domain ontologies are not consistently used to extract, standardize, and align entities, relations, and time markers at scale. This creates a significant gap in evidence for both research and practice and weakens connections with structured registries Latonero et al. (2012); Poelmans et al. (2012).

1.5.3 Semantic interoperability and ontology-driven integration

Trafficking analytics involve various custodians and jurisdictions with different schemas and vocabularies. Without a semantic layer based on ontology, connections become fragile, variable meanings are lost, and cross-border comparisons are unreliable. FAIR-aligned semantic modeling and tracking of data provenance are only occasionally applied instead of being fundamental, which limits reuse, comparability, and machine-actionable transparency Wilkinson et al. (2016b). A well-structured, domain-specific ontology is necessary for stable integration of narratives, administrative records, OSINT, and geospatial information.

1.5.4 Data sharing, governance, and provenance as first-class design requirements

Collaboration suffers from inconsistent formats, unclear legal bases for processing, and weak provenance trails. Prototypes often push governance aside, leading to limited auditability and trust issues within institutions. Methodological advancements must be paired with governance-by-design. This includes role-based access, minimization, consent tracking, encryption during

transmission and storage, and complete provenance that makes every change and decision traceable Wilkinson et al. (2016b). Decentralized or institutionally managed data stewardship can improve control and accountability when backed by strict access policies and verifiable logs.

1.5.5 From static graphs to time-aware network modeling

Most network analyses are static snapshots that overlook temporal changes in roles, routes, and coordination strategies. Without timestamped connections, episode modeling, and trajectory-aware inference, systems cannot identify new configurations or predict near-term shifts. Time-aware knowledge graphs, enhanced by temporal markers and event sequencing, are essential for moving from description to explanation and early warning **?**.

1.5.6 Insufficient geospatial integration and uncertainty-aware cartography

Even when locations are noted, solid resolution of place names, quantification of uncertainty, and layered, interactive mapping are uncommon. Maps without provenance for geocoding accuracy or time-based animation obscure movement patterns and reduce operational usefulness. Connecting narrative extraction to GIS requires thorough geocoding, disambiguation signals, density and clustering layers to address overlap, and time-aware paths to show dynamics important for screening, interdiction, and assistance planning United Nations Office on Drugs and Crime (2023a).

1.5.7 Real-time adaptability and streaming readiness

Stakeholders need timely updates on routes, hubs, and tactics. Batch-centric systems trained on static data fail to incorporate and analyze new evidence quickly enough for operations. Frameworks that support streaming data ingestion, incremental inference, and controlled cache invalidation while maintaining provenance and reproducibility are still underdeveloped in trafficking analytics. These are vital for generating actionable intelligence.

1.5.8 Evaluation design and the evidentiary standard for usefulness

Metrics that focus solely on models (like F1 for NER or AUROC for classification) do not demonstrate practical usefulness. Evaluation must be modular, assessing extraction accuracy, graph structure, and reasonable spatiotemporal claims while measuring uncertainty and identifying failure modes. Most importantly, assessments must connect to user tasks, decision benchmarks, and potential risks to prevent misleading performance claims and ensure procurement-grade comparability United Nations Office on Drugs and Crime (2023a).

1.5.9 Cross-regional interoperability and transfer

Since trafficking is a transnational issue, tools must work across different regions and legal contexts. However, schemas, ontologies, and models are seldom designed for portability between jurisdictions or tested for cross-region effectiveness. Without focus on clear, mappable vocabularies and unified metadata, comparative analysis and multinational collaboration remain challenging.

1.5.10 Ethics, privacy, and risk mitigation as testable properties

Sensitive humanitarian data requires clear, testable protections. Many artifacts lack documented threat models, de-identification trade-offs, or verifiable access controls. Responsible use calls for privacy by design, minimization, differential risk assessments, and transparent model governance, including checks for bias and explainability compliant with regulations like GDPR. Ethics cannot be an afterthought; it should be integrated into the system and demonstrated through evaluation Wilkinson et al. (2016b).

1.5.11 Alignment with practitioner requirements and traceability to features

There is a continuous gap between academic problem-solving and the real needs of investigators, service providers, and policy staff. Requirements should come from the literature and domain expertise, then link back to implemented features and evaluation elements that show practical utility. Without this connection, integrated systems may seem like patchwork solutions instead of a unified answer to well-defined problems United Nations Office on Drugs and Crime (2023a).

1.5.12 Synthesis.

These gaps are interconnected. The under-use of narratives affects ontology-driven integration, the lack of temporal modeling limits geospatial insights, weak governance undermines trust and reproducibility, and insufficient evaluation does not prove usefulness. To tackle these issues, this thesis proposes an ontology-anchored, end-to-end pipeline that converts unstructured texts into a categorized, time-aware knowledge graph and produces spatiotemporal intelligence through GIS within a governance-aware, reproducible framework. This synthesis directly informs the research questions and methods to follow.

1.6 Research Questions & Sub-Questions

This section outlines the inquiry that guides the thesis. The research questions come from a combination of the anti-trafficking literature and data governance practices. They relate to established requirements engineering to ensure each question ties back to stakeholder needs and tangible artifacts produced by the system. The main goal is to connect unstructured narratives and various administrative sources with a governance-aware, ontology-based analytics pipeline. This will produce decision-ready, spatiotemporal intelligence that can withstand both scientific and operational scrutiny.

1.6.1 Primary Research Question (PRQ)

How can an ontology-based, governance-aware, end-to-end pipeline transform unstructured narratives and various administrative records into a structured, time-aware knowledge graph and geospatial intelligence that effectively supports practitioner tasks in anti-trafficking contexts?

The PRQ connects three main aspects: (i) semantic interoperability based on a domain ontology, (ii) data management that considers privacy, residency, and provenance while following widely accepted frameworks, and (iii) integrated NLP, graph, and GIS analytics that produce results which are reproducible, auditable, and understandable by investigators and humanitarian workers.

1.6.2 RQ1 — Semantic interoperability for heterogeneous sources.

What domain ontology and semantic validation procedures are needed to achieve interoperability across narrative testimonies, administrative registries, and open-source intelligence, while preserving meaning for downstream graph and GIS analytics?

The thesis argues that an ontology-driven layer is essential for integrating diverse sources and making fair comparisons across different regions and institutions. Specifically, the work looks into how controlled vocabularies and constraint languages can stabilize representations of recruitment, transit, exploitation, actor roles, time anchors, and geo-references, so the resulting knowledge graph can support both structural and spatiotemporal queries.

- *SQ1.1. What core classes, relations, and axioms must the ontology include to accurately represent events and actors in trafficking narratives while remaining applicable across jurisdictions?*
- *SQ1.2. How should extracted entities and events be standardized and validated to ensure reliable, auditable graph construction (for example, shape constraints and integrity rules for time, role, and location fields)?*
- *SQ1.3. To what extent does the semantic layer allow transfer across regions (for example, East to North Africa routes), and what mappings are necessary to accommodate local taxonomies while keeping them comparable?*

1.6.3 RQ2 — Data security, residency, privacy, FAIR governance, and federated analytics

Which data governance structure allows for secure, legal, and practically useful collaboration under privacy and residency rules while maintaining FAIR (Findable, Accessible, Interoperable, Reusable) properties and supporting federated analytics when data cannot be centralized?

The thesis views governance as a key engineering problem. It addresses access control, data minimization, encryption at rest and in transit, key management, audit trails, provenance capture, and documentation (like model and data cards). It evaluates whether federated or distributed analytics can limit cross-border movement of sensitive data while still producing comparable, scientifically valid results.

- *SQ2.1. What mix of policy controls (role-based access, minimization, consent tracking) and technical controls (encryption, key management, tamper-evident logging) meets GDPR requirements for narrative and case data while keeping it analytically useful?*
- *SQ2.2. How can data residency rules and institutional boundaries be respected through federated analytics or learning, while maintaining thorough provenance, comparable results, and FAIR-compliant metadata?*
- *SQ2.3. What provenance standards and documentation artifacts (for instance, PROV-compatible records, model cards, data-use agreements) are necessary to make the pipeline auditable and trustworthy for external reviewers?*
- *SQ2.4. What trade-offs arise between privacy safeguards (like aggregation thresholds and noise mechanisms) and the detail required for practitioner tasks, and how should these be openly addressed in reports and interfaces?*

1.6.4 RQ3 —Integrated NLP-Graph-GIS pipeline with governance-by-design

How can an end-to-end pipeline that combines modern NLP, knowledge graph construction, and GIS visualization be designed to be modular, reproducible, and governed by design, and to what extent does it offer utility validated by practitioners on relevant tasks?

This question combines pipeline integration with ethical and provenance-aware operation. It treats front-end orchestration and artifact generation (such as interactive graph exports, layered maps, and animated trajectories) as research objects, requiring versioned configurations, repeatable runs, and traceability from inputs to outputs.

- *SQ3.1. What is the minimal yet sufficient NLP stack (including tokenization, NER, canonicalization, relation/event extraction) that produces graph-ready, time-stamped records that align with the domain ontology?*
- *SQ3.2. How should the knowledge graph be structured (node and edge types, time qualifiers, provenance fields) to support structural queries, time-aware reasoning, and easy integration into geospatial views?*

- *SQ3.3. What geocoding, toponym resolution, and uncertainty-handling methods are necessary for creating density maps, clusters, and animated trajectories that accurately represent narrative-derived locations over time, without exaggerating precision?*
- *SQ3.4. How can the pipeline's orchestration and user interface be designed so that artifacts are repeatable, versioned, exportable, and understandable to non-technical users, allowing for independent verification and incorporation into reports?*
- *SQ3.5. In evaluations focused on practitioners, to what extent do the pipeline's outputs enhance task performance on relevant questions (such as route triage, hub identification, cross-case consolidation), and what limitations do uncertainty and failure modes impose on this utility?*

1.6.5 Traceability and evaluation logic.

Each research question corresponds to specific gaps and evidence the system must produce. RQ1 addresses semantic fragmentation by committing to an ontology and clear validation rules. RQ2 deals with lawful and ethical collaboration under privacy and residency constraints by defining governance controls, provenance, FAIR-aligned metadata, and federation when centralization isn't possible. RQ3 focuses on creating an integrated analytics pipeline and connects its outputs to user tasks. Evaluation follows a modular framework: (i) extraction accuracy and standardization, (ii) structural and temporal validity of the graph under constraint checks, (iii) geospatial realism and uncertainty reporting, (iv) governance compliance through complete provenance and access control checks, and (v) user-centered evaluations of task usefulness with noted limitations. The research questions, sub-questions, and evaluation artifacts are collaboratively designed to ensure that the thesis is both scientifically sound and practically useful.

1.7 Research Objectives & Sub-Objectives

This section presents the research objectives that address the earlier questions. Each objective is framed as a required outcome for the thesis. Each one includes sub-objectives that make the goal specific, testable, and auditable. This approach ensures the work clearly moves from a defined problem to measurable results. It helps connect academic theory with practical use. The objectives relate directly to the integrated pipeline and the governance structure developed in this thesis. This relationship guarantees that the contributions are methodologically sound, ethically justifiable, and practically useful.

1.7.1 Objective 1: Establish a domain ontology and semantic validation regime that achieves cross-source, FAIR-aligned interoperability

The first objective is to create and implement a semantic layer that allows various sources, such as narratives, administrative records, and open-source intelligence, to merge without losing meaning while ensuring machine-actionable transparency. The ontology must define key actor classes, roles, events, time markers, and geographic references in a way that is not linked to any jurisdiction but detailed enough to address the investigative questions of stakeholders. Sub-objectives include specifying classes and relationships with formal constraints, stabilizing identifiers through standardization, and defining integrity rules to make graph construction reliable and auditable. The ontology will work with a shapes-based validation method to ensure that records and constructed subgraphs meet the integrity rules before analysis, thus improving reproducibility, comparability, and error detection later (Wilkinson et al., 2016b; Knublauch and Kontokostas, 2017).

A first sub-objective is to define the vocabulary and framework needed to represent recruitment, transit, and exploitation events; actor roles and relationships; time qualifiers; and location references at an abstraction level that enables cross-regional application. The deliverable is an ontology file along with human-readable documentation explaining the modeling choices, including mappings to commonly used trafficking and migration taxonomies where applicable. The second sub-objective is to develop a shapes-based validation method so that extracted records and assembled subgraphs can be checked against the ontology's integrity rules before ingestion or analysis. The deliverable includes a set of shapes covering time, role, location, and source fields, plus a validation report that measures conformance and highlights consistent violations. The third sub-objective is to implement reliable standardization of entities and events, leading to stable identifiers that ensure reproducibility across different environments and runs. Together, these sub-objectives confirm that interoperability is shown, in line with FAIR principles for data, tools, and workflows (Wilkinson et al., 2016b; **?**).

1.7.2 Objective 2: Engineer a governance architecture that satisfies privacy, residency, and accountability while preserving FAIR properties and enabling federated analytics

The second objective regards data governance as a key engineering issue. It involves creating a system where sensitive records can be processed lawfully and ethically, with access

controls that respect institutional limits and residency rules while allowing for meaningful collaborative analysis. Sub-objectives include policy controls (such as purpose limitation, minimization, role-based access, and consent tracking), technical controls (including encryption during transmission and storage, key management, and tamper-evident audit trails), and capturing provenance to make every transformation and analytic decision traceable. The goal is to establish a defensible baseline that complies with regulations while maintaining FAIR-compliant metadata for discovery and reuse (EU-, 2016; Wilkinson et al., 2016b; Lebo et al., 2013).

The first sub-objective is to specify and implement a model for access and purpose control that details permissible uses of narrative and case data, with clear role definitions and audit mechanisms. The second sub-objective is to enable complete provenance tracking so that the history of every artifact—from raw input to normalization, extraction, graph creation, and visualization—can be reconstructed, queried, and exported in a standards-aligned format. The deliverable will contain provenance records that are comprehensive enough to support external review and independent re-analysis (Lebo et al., 2013). The third sub-objective is to assess the feasibility of federated or distributed analytics where centralization isn't possible, allowing data to be kept in place while only exchanging models or aggregate data. The deliverable will demonstrate that these methods honor residency requirements without compromising validity or comparability of results, along with an analysis of communication, privacy, and governance trade-offs (McMahan et al., 2017). A fourth sub-objective is to create documentation artifacts—model cards and data documentation—that clarify intended use, limitations, evaluation conditions, and known risks, thus improving accountability and interpretability for non-technical reviewers (Mitchell et al., 2019).

1.7.3 Objective 3: Build an integrated, governance-by-design NLP–Graph–GIS pipeline whose artifacts are deterministic, versioned, and practitioner-comprehensible

The third objective combines the implementation goals discussed in previous sections. It aims to deliver a flexible system that transforms unstructured narratives and structured records into a typed, time-aware knowledge graph and spatiotemporal intelligence that investigators can understand and act upon. Sub-objectives start with defining a minimal yet sufficient natural language processing stack—tokenization, normalization, named entity recognition, standardization, and relation/event extraction—explicitly linked to the ontology to ensure outputs are ready for graphs and time-stamped. This continues with a graph construction process that establishes node and edge types, time qualifiers, and provenance fields in a way that accommodates structured and time-sensitive queries, culminating in a geospatial module that translates place names into coordinates while accounting for uncertainty, aggregates spatial information, and produces layered, interactive maps, including density, clustering, and animated trajectories.

The first sub-objective is to ensure that each module has a stable interface and that the overall orchestration remains consistent under fixed settings. The deliverable consists of a set of versioned configurations and exported artifacts—interactive graph visuals and map outputs—that can be regenerated on demand to match the same checksums, plus automated reports summarizing parameter settings for inclusion in the manuscript. The second sub-objective is to design the user interface so decision-makers can export figures and tables directly from the

application with relevant context, reducing the risk of misinterpretation and facilitating independent verification. The third sub-objective is to integrate governance directly into the pipeline: access checks at ingestion, validation gates before graph creation, uncertainty annotations for geocodes and trajectories, and provenance links on every output. The measurable outcome is a system where functionality aligns smoothly with ethical and accountability standards, meeting contemporary expectations for sensitive data analytics (United Nations Office on Drugs and Crime, 2023a; Counter-Trafficking Data Collaborative, 2025).

1.7.4 Objective 4: Demonstrate scientific validity and practitioner usefulness through a modular evaluation design

The fourth objective is to establish a standard for both scientific validity and practical use. Sub-objectives specify, before any experiments, the tests and user assessments at the module level that will determine success. On the scientific side, evaluation includes checking extraction fidelity for NLP components, structural and temporal validity for the knowledge graph under the ontology and shapes constraints, and demonstrating spatiotemporal plausibility for the geospatial products with clear reporting of uncertainty to avoid misleading impressions of precision. On the operational side, evaluation connects outputs to typical practitioner tasks—route triage, hub identification, and cross-case consolidation—with metrics focused on decision support rather than model-only measures. The deliverable consists of an evaluation protocol and supporting materials that allow external reviewers to understand what was measured, its relevance, and how limitations were addressed or disclosed (Mitchell et al., 2019; United Nations Office on Drugs and Crime, 2023a).

A first sub-objective is to formalize acceptance criteria that fit the domain and are realistic for the available data, such as addressing ontology constraints in validation reports, reproducibility of exported materials across different runs and environments, and completeness of provenance chains from raw input to published figures. A second sub-objective is to outline uncertainty thresholds and disclosure guidelines for geospatial rendering and temporal inference, ensuring stakeholders are not misled by apparent precision. A third sub-objective is to document negative outcomes and failure modes, including scenarios where federation, privacy constraints, or data limitations restrict analytic depth, thus aligning evaluation practice with responsible research norms.

1.7.5 Synthesis and traceability.

Together, these objectives and their sub-objectives turn the research questions into a practical work program with outputs that are both technically precise and aware of governance concerns. Objective 1 addresses semantic fragmentation by providing an ontology and validation method that make interoperability measurable and clear. Objective 2 tackles privacy, residency, and accountability by delivering a governance framework that enforces lawful processing, provenance, and, when suitable, federation. Objective 3 brings the integrated NLP–Graph–GIS pipeline to fruition and incorporates governance into its functioning and user experience. Objective 4 creates an evidence-based framework that confirms scientific validity and practical

usefulness. The overall result is a thesis where the system's architecture, ethical commitments, and evaluative assertions support each other rather than being incidental.

1.8 Theoretical Concepts, Background, Framework, & Assumptions

This research builds on an interdisciplinary foundation that www.findability.org integrates network science, natural language processing (NLP), and geospatial intelligence to construct an AI-driven framework for analyzing and predicting human trafficking patterns. The work is guided by the FAIR data principles—Findability, Accessibility, Interoperability, and Reusability—which ensure that both the datasets and analytical outcomes remain transparent, reproducible, and ethically governed Wilkinson et al. (2016b); Vogt et al. (2024). By aligning the pipeline with FAIR, the study addresses one of the recurring challenges in human trafficking research: the fragmentation of heterogeneous data across jurisdictions, institutions, and modalities Prince Sales et al. (2023).

The framework fundamentally uses graph theory and network science to characterize human trafficking as a complex, dynamic system of relationships between victims, traffickers, locations, and transit hubs. Graph Neural Networks (GNNs), such as Graph Convolutional Networks Kipf and Welling (2016) and Graph Attention Networks Veličković and Cucurull (2018), are implemented using PyTorch Geometric Fey and Lenssen (2019) to capture higher-order dependencies within these networks. This enables the framework to pinpoint central actors, uncover hidden communities, and perform link prediction tasks that forecast potential future trafficking routes. Compared to static social network analysis approaches, GNNs are much more capable of real-time adaptation as new data is ingested. This is especially beneficial for human trafficking, which is an evolving problem that requires dynamic solutions Dimas et al. (2021); Konrad et al. (2022).

Concurrently, the pipeline takes advantage of Natural Language Processing (NLP) to convert unstructured narratives—such as victim statements, police reports, and open-source intelligence—into a structured form. We use state-of-the-art NLP methods for this, relying on transformer-based models like BERT Devlin et al. (2019) and its derivatives. These are used for tasks such as Named Entity Recognition (NER), which figures out what types of things (considered entities) are mentioned in the text (e.g., individuals, locations, events). NER is not perfect, but current models have reached a point where they can achieve a significant degree of accuracy—not just with the types of entities that are commonly recognized (this is the near-synonym mapping that NER does) Ibanez and Suthers (2014b); Poelmans et al. (2012), but also with many more obscure or private entities, including with entities in collections (i.e., groups of things that 'look' the same without being a direct copy of another entity). These mapped entities are fed back into the model architecture.

Another vital facet of the framework is the spatio-temporal analysis enabled by Geographic Information Systems (GIS). Building on well-documented trafficking hotspots in East Africa and North Africa National Crime Research Centre (2018); International Organization for Migration (2018); United Nations Office on Drugs and Crime (2024), GIS is used to map trafficking flows over time, yielding not just a recruitment map but also a transit map—the two together giving a sense of how a trafficking operation unfolds in real time. These spatial layers provide a means to detect route convergence, cross-border corridors, and urban exploitation centers. Integrating GIS with GNNs enables the system to contextualize relational patterns in geographic space, producing actionable insights for stakeholders who require not just network analysis but also

territorial intelligence Yagci Sokat et al. (2024).

These two maps in tandem provide the basis for figuring out how a trafficking operation is vectoring across space. Together with GNNs, GIS becomes an intelligence platform for the trafficking analyst.

The technical pillars are undergirded by a set of guiding assumptions.

- The system assumes that both structured and unstructured data sources contain latent but complementary insights; when integrated, these provide a more holistic view of trafficking dynamics.
- The framework assumes that advanced deep learning methods can generalize from partial, noisy, and incomplete datasets—a common reality in trafficking research.Little and Rubin (2019)
- The study operates under the ethical assumption that all data, particularly victim narratives, must be processed with strict attention to privacy, provenance, and regulatory compliance Hasan and Myagmar (2022); Regent University (2022).
- Finally, the research assumes that interoperability between datasets from North Africa, and global repositories like CTDC is achievable through semantic harmonization and ontology-driven integration Kejriwal and Szekely (2022).

These foundational elements—graph theory, natural language processing, geographic information systems, and FAIR-governed data—provide a coherent framework that can cut across the siloed methods of human trafficking analytics. By integrating relational, textual, and spatial forms of analytical human intelligence, the system offers an alternative to the traditional, fragmented models of human trafficking and advances toward an intervention model that is scalable, adaptive, and ethically robust.

1.9 Ethical Considerations and Data Management

Ethical considerations form a central pillar of this research, as human trafficking data is among the most sensitive forms of information, often involving survivors whose identities and experiences must be safeguarded at all times. The handling of such data requires adherence to strict ethical frameworks that prioritize privacy, informed consent, and compliance with international data protection laws. This study follows the FAIR principles Wilkinson et al. (2016b), ensuring that datasets are findable, accessible, interoperable, and reusable, while simultaneously embedding rigorous privacy and security safeguards. The extension of FAIR to semantic interoperability through FAIR 2.0 further strengthens cross-jurisdictional data integration and ethical governance Vogt et al. (2024).

Data provenance is a matter of tracking data from its origins, through its transformations, and then to its applications across an analytical pipeline. It is all about ensuring accountability, transparency, and reproducibility—and also about enabling the very diverse range of stakeholders with whom we interact to assess whether we are handling data in a way that aligns with relevant regulatory and ethical standards Hasan and Myagmar (2022).Provenance, in short, ensures that we are doing science not only with rigor but also with trust—and a palpable sense of ethics. This is something we take seriously, particularly when working across the very diverse regional contexts of East and North Africa.

The study also incorporates robust data protection measures shaped by worldwide privacy structures such as the European Union's General Data Protection Regulation and other international frameworks that guide personal data protection in humanitarian contexts **?**. Sensitive personal information—like victim testimonies—is anonymized and pseudonymized before any analysis is done. Furthermore, access to the raw datasets is restricted to just a few people, and all data sharing is conducted under agreements that make the data confidential and follow ethical standards that prioritize the welfare of the survivors Regent University (2022).

This project emphasizes secure, institutionally managed storage solutions where strict access controls and encryption protocols keep sensitive information safe from unauthorized access. Compared with some past approaches, this project's methods make it more likely that sensitive information will be kept secure while also allowing for robust interoperability and secure collaboration between various research teams Lundberg and Lee (2017). Additionally, the project evaluates the AI models developed within it for fairness and potential bias.

In short, this study combines FAIR-compliant governance, good data management, and robust privacy protections into an ethical framework. With this setup, researchers can analyze huge data sets and find subtle, previously unnoticed signals that might indicate human trafficking. What makes these "signals" comprehensible and trustworthy is that the study also provides big-picture context for what's going on with human trafficking globally. And then there's something even more important: This study doesn't compromise the dignity, privacy, or safety of the people it studies.

1.10 Research Design

This research uses a mixed-method design where qualitative and quantitative approaches are integrated into a computational pipeline. The computational pipeline is necessary because human trafficking is a complex, interdisciplinary problem that requires not only the integration of kinds of knowledge but also the kinds of narratives that tell us how human trafficking occurs in specific places at specific times. Advanced computational modeling permits the kinds of integration required. And because we are talking about a social problem with serious consequences for individuals, the work must also be scientifically and socially responsible.

The overarching design is guided by the FAIR principles Wilkinson et al. (2016b), ensuring that datasets and analytical outputs are findable, accessible, interoperable, and reusable. These principles address the frequent fragmentation of trafficking data across institutions and regions, a challenge that hampers collaboration and comparative research Prince Sales et al. (2023). The design further extends to FAIR 2.0, which incorporates semantic interoperability to ensure that knowledge extracted from diverse data sources can be consistently aligned and reused across different jurisdictions Vogt et al. (2024).

At the technical level, the framework is arranged into sequential yet interconnected stages. The first stage involves the development of an ontology and a semantic model, where trafficking-related entities—like victims, perpetrators, locations, and routes—are formally defined. Ontology-driven methods tackle the problem of semantic heterogeneity and furnish a consistent foundation for data integration. They allow for the combination of data from sources with different structures—like the CTDC, which contains mostly structured data, versus which mostly unstructured data comes from interviews and testimonies Kejriwal and Szekely (2022).

Phase two targets the qualitative narrative's natural language processing. For this, we utilize models that are transformer-based, like BERT Devlin et al. (2019) and its variants. Using these, we complete tasks such as Named Entity Recognition (NER), which help us in extracting key entities and relationships from our unstructured text. This unstructured text is whirlpooling with the contextual richness of survivor testimonies. Through this unstructured text, we are transforming the aforementioned statements into structured, machine-readable context—without ruining their essence. This is as per Jurafsky and Martin (2020); Ibanez and Suthers (2014b). The structured outputs are aligned with our ontology from phase one, ensuring that our output is semantically consistent.

The operationalization of structured data is achieved through the construction of graphs and spatiotemporal modeling. Graph Neural Networks (GNNs)—for example, Graph Convolutional Networks (GCNs), Kipf and Welling (2016), and Graph Attention Networks (GATs), Veličković and Cucurull (2018) are implemented using PyTorch Geometric Fey and Lenssen (2019). These models are utilized to perform three tasks: identify central actors; detect hidden communities; and predict potential links in the trafficking networks, specifically trafficker-victim relationships Dimas et al. (2021); Konrad et al. (2022). Simultaneously, Relational Geographic Information Systems (GIS) plot these relational dynamics in space and time, allowing a visualization of the networks as they operate across place and duration Yagci Sokat et al. (2024); United Nations Office on Drugs and Crime (2024).

The fourth phase centers on predictive modeling. Recurrent neural networks like Long Short-Term Memory (LSTM) models citehochreiter1997lstm and transformer-based time-series models Vaswani et al. (2017) are used to forecast new trafficking patterns that are likely to emerge, such as future hotspots or new routes. This ensures that the framework does not merely characterize existing dynamics but also anticipates future risks in a way that is quite similar to how most forecasting works in law enforcement.

All stages of the research are protected by ethical safeguards that are integrated into its design Hasan and Myagmar (2022). The accountability and transparency of the data handling are ensured with provenance techniques. Sensitive data is managed under strict privacy regulations that comply with international humanitarian data guidelines and the GDPR **?**Regent University (2022). The research outputs, in the form of complex models, are ensured to be interpretable and valid by domain experts through the use of model explainability methods Lundberg and Lee (2017).

This research takes on the ambitious task of addressing the substantial gaps identified in the human trafficking literature. It leverages ontology-driven semantic modeling and narrative structure from natural language processing to effectively analyze the kinds of large, unstructured data that are typical in trafficking cases. It then uses predictive modeling to identify likely trafficking scenarios. That might be enough for a usual human trafficking research project. But what is unusual about this project is what comes next.

1.11 Relevance

One of the most pressing humanitarian and security challenges worldwide is human trafficking. It affects an estimated 49.6 million individuals and is a nasty business that generates illicit profits exceeding $150 billion annually OUR Rescue (2024); International Labour Organization (2014). When it comes to human trafficking, one must think not only of those ensnared in forced labor but also of the many who are or have been sold for sexual exploitation. Ample government and NGO resources are directed toward combating the problem, and traditional detection methods often rely on structured administrative or law enforcement records. But many of us need to have a much sharper and clearer picture in our minds of just what trafficking looks like—not only at home but also in the many foreign countries to which we direct our business Allais (2005); United Nations Office on Drugs and Crime (2024).

A significant shortcoming in present methods is the scant use of unstructured, qualitative data like survivor testimonies, police narratives, and open-source intelligence. These materials often hold nuanced, crucial details about recruitment tactics, victim–perpetrator relationships, and local context that we miss when we confine our analysis to structured datasets like the Counter Trafficking Data Collaborative (CTDC). Studies have highlighted the essentialness of qualitative sources for a comprehensive understanding of human trafficking analytics Poelmans et al. (2012); Szakonyi et al. (2021), yet almost no frameworks have ventured to operationalize these materials with cutting-edge computational techniques. My research directly addresses that dearth by using some of the most advanced NLP models—featuring the human-like comprehension of the BERT model Devlin et al. (2019) to extract, structure, and integrate, at scale, the crucial insights from those unstructured narratives.

Likewise, the methodologies that currently exist are still too fragmented, with network analysis, machine learning classifiers, or geospatial visualization applied in isolation without much integration between them and rarely achieving anything resembling the formal structure of a dynamic, scalable pipeline Dimas et al. (2021); Konrad et al. (2022). Graph Neural Networks (GNNs) appear to be a promising way of achieving the necessary integration. They can identify the central traffickers in relational networks, the hidden communities, and potential future links that form within these networks Kipf and Welling (2016); Veličković and Cucurull (2018). Moreover, GNNs can do this in a way that is both "relational" and "geographic" Yagci Sokat et al. (2024); National Crime Research Centre (2018); International Organization for Migration (2018). This is precisely what we need in order to achieve multilayered intelligence.

Predictive modeling is another relevant area. While the existing literature tends to focus on retrospective analysis, trafficking networks are dynamic and continuously adapting to stay one step ahead of law enforcement. This research tries to do something much harder—predicting attempts to evade detection that haven't been made yeta although the motive in this research is purely experimental. To do that, the researchers employ a range of temporal deep learning models, including long short-term memory networks Hochreiter and Schmidhuber (1997b) and transformer-based architectures Vaswani et al. (2017). This research models trafficking networks as a type of Evolving Graph Structure that can offer predictive insights. These insights can hopefully be used to allocate scarce resources to the emerging hotspots or just-in-time shifts in trafficking routes that these networks make.

In the end, this study gives priority to the ethical and legal aspects. It guarantees the responsible handling of sensitive data Hasan and Myagmar (2022). Within the study, survivor narratives have been rendered anonymous, a move that not only protects individual identities but also upholds the principle of data provenance—ensuring that we can track where our data comes from and, as importantly, where it goes. With that foundation, let us explore the guarantees of transparency and responsible data governance that this framework provides **?**Regent University (2022).

This research is highly relevant to global anti-trafficking work and holds the potential to significantly inform and influence the field's trajectory for the better. Its relevance stems from four main data gaps that the research work attempts to address. These gaps center around (1) data usage, (2) integrations of different methodologies, (3) predictive adaptability, and (4) ethical compliance. The research attempts to address these gaps and thus has the potential to generate huge payoffs for the anti-trafficking field.

1.12 Location of the Study

The primary geographical focus of this research is North Africa, which plays multiple roles in the global human trafficking issue as a source, transit, and destination hub National Crime Research Centre (2018). North Africa links to key trafficking routes that lead to the Middle East and Europe. The coastal regions of various countries, in particular, have been identified as areas where trafficking is common. Here, victims are recruited or moved before crossing borders International Organization for Migration (2018). These regions show the mix of economic hardship, open borders, and weak institutional oversight, making them crucial sites for intervention and study.

Given the international nature of human trafficking, a secondary but equally important focus of this study is the compatibility of analytical models with North African contexts. Trafficking routes from East Africa often come together in North African transit hubs before moving into European and Middle Eastern markets United Nations Office on Drugs and Crime (2024). The cross-border relationships between East and North Africa highlight the importance of creating models that can combine different datasets from various regions. By developing a framework that is clear in meaning and follows FAIR principles, this research facilitates the comparative analysis of trafficking patterns beyond national borders, supporting more coordinated regional efforts.

The research also considers the practical challenges of conducting direct fieldwork in North Africa. As an international master's student at Leiden University, the author depends on academic partnerships rather than gathering primary field data. Datasets are accessible through established collaborations under the guidance of Prof. Mirjam van Reisen at the Leiden Institute of Advanced Computer Science (LIACS). Doctoral researchers in this network, such as Kai Smiths, have carried out significant fieldwork in North Africa and nearby areas, producing qualitative victim accounts and organized case datasets that form the empirical basis for this study.

By focusing on North Africa as the main case study and expanding interoperability to other African datasets, the research finds a balance between being grounded in data and having regional applicability. This dual focus ensures that the framework reflects both the specific local dynamics of East African trafficking and the wider international patterns that connect to North Africa. Thus, the choice of location not only boosts the study's methodological strength but also enhances its practical relevance for those working to combat trafficking across related regions.

1.13 Timeline of the Study

Table 1. *Timeline of the Study (February 2025 – August 2025)*

Month	Phase	Activities and Deliverables
Feb 2025	Problem Analysis and Literature Review (Phase I)	• Conducted extensive literature review on human trafficking, semantic technologies, graph networks, and NLP. • Finalized research gaps and formulated refined research questions & objectives with supervisor feedback. • Reviewed datasets (CTDC, UNODC, anonymized interview data) and assessed their suitability. • Began drafting Chapter 1 (Introduction) and Chapter 2 (Theoretical Framework).
Mar 2025	Methodological Design and System Architecture (Phase II)	• Designed the modular analytical pipeline integrating NLP, ontology, GNNs, GIS, and predictive analytics. • Implemented initial data preprocessing and NLP modules (NER, coreference resolution). • Finalized methodology chapter structure; began drafting initial methodology sections. • Participated in mid-term review with supervisor; feedback integrated.

Month	Phase	Activities and Deliverables
Apr 2025	Ontology Development and Symposium Presentation (Phase III)	• Developed OWL-based human trafficking ontology using Protégé. • Validated ontology using SHACL constraints and integrated with NLP outputs. • Presented research at the *International SOLID Symposium 2025* (April 24–25, LIACS). • Revised Chapters 1 and 2 based on symposium feedback; enhanced clarity and rigor.
May 2025	Graph Construction and GNN Integration (Phase IV)	• Constructed multi-relational knowledge graphs from preprocessed and ontology-tagged data. • Implemented Graph Neural Networks using PyTorch Geometric. • Designed experiments for role prediction and route inference; finalized evaluation strategy. • Continued development of thesis draft (Chapter 3: Methodology).
Jun 2025	GIS and Predictive Modules (Phase V)	• Built GIS visualizations using GeoPandas, Leaflet, and Folium. • Integrated NLP+Graph outputs into dynamic spatio-temporal maps. • Implemented LSTM/Transformer-based predictive models for route forecasting. • Conducted internal testing and debugging of pipeline modules.

Month	Phase	Activities and Deliverables
Jul 2025	Evaluation and Final Results (Phase VI)	• Executed full experimental pipeline; evaluated model performance (F1, AUC, Precision@k, RMSE). • Interpreted results and generated visualizations and comparative tables. • Drafted Chapter 4 (Experiments & Results) and Chapter 5 (Discussion). • Conducted feedback loop with supervisor; incorporated changes iteratively.
Aug 2025	Thesis Finalization and Defense Preparation (Phase VII)	• Polished all chapters and finalized full thesis document. • Wrote Chapter 6 (Conclusion) and supplementary materials (Abstract, Acknowledgments, Appendices). • Submitted thesis for examination; prepared presentation slides and defense script. • Final oral defense conducted before committee.

2 Theoretical Framework

This research explores the intersection of semantic interoperability, responsible artificial intelligence, data provenance, and data ethics focused on human rights. The primary goal is to design and build a complete analytical pipeline for detecting human trafficking. This pipeline aims to be technically robust, ethically responsible, and practically scalable. This chapter discusses the theoretical principles, normative assumptions, and computational models that guide the development of the proposed system.

2.1 Foundations in Semantic Interoperability and Ontologies

A key concept behind this research is semantic interoperability. Human trafficking data is often scattered due to different sources like NGOs, law enforcement agencies, and international organizations using various formats and terms. Semantic interoperability unites these diverse data sources by mapping them to a common framework called an ontology. Ontologies offer machine-readable and logically structured vocabularies that define important entities, relationships, and constraints within a specific area Wilkinson et al. (2016b).

In this project, we develop a domain-specific ontology for human trafficking that follows the FAIR principles Wilkinson et al. (2016b); Vogt et al. (2024). The ontology allows integration of structured datasets, like administrative records, and unstructured sources, like survivor testimonies or police reports, into a unified semantic layer. This foundational layer enables meaningful data linking, reasoning, and downstream analytics using graph-based models. Additionally, the ontology follows the FAIR 2.0 recommendations Vogt et al. (2024), ensuring interoperability and semantic richness that supports international cooperation and reuse across various jurisdictions.

2.2 Data Ownership, Sovereignty, and Control

An important assumption of this research is that data custodians, especially survivors and front-line organizations, should have control over their data. Traditional anti-trafficking systems often rely on centralized repositories where sensitive information is pooled. This practice can risk misuse or loss of agency for those providing the data. This study moves away from such extraction methods by emphasizing institutionally managed but ethically governed data systems.

Instead of decentralization through personal data storage, the framework uses secure, institutionally managed storage solutions backed by strict access controls, encryption, and anonymization protocols. Data custodians and contributing organizations maintain control over how their data is used, with access granted only through agreements that respect survivor dignity and comply with international privacy laws **??**. By including sovereignty in governance practices, this pipeline ensures responsible analytics while protecting survivors and communities from potential risks.

2.3 Trust as a Design Principle

Trust is crucial for adopting any analytical framework in sensitive situations like anti-trafficking work. In this project, trust is established through three connected methods. First, transparency and traceability are achieved using data provenance methods that track the source, changes, and uses of data throughout the pipeline. Provenance confirms that all analytical outputs can be verified for accountability Hasan and Myagmar (2022).

Second, we integrate explainable AI techniques across the pipeline. Outputs from natural language processing models and graph neural networks go through interpretability frameworks like SHAP Lundberg and Lee (2017), allowing stakeholders to evaluate the basis for any prediction or classification. Third, we embed ethical safeguards from the outset. Survivor data is anonymized, sensitive attributes are minimized, and informed consent supports all partnerships, ensuring users can trust both the process and the outcomes of the system. Hence, trust in this framework is built into the technical and ethical design. By creating systems that are clear, predictable, and focused on survivor needs, the pipeline encourages sustainable collaboration among multiple stakeholders.

2.4 Responsible AI and Distributed Processing Paradigms

This research also draws on distributed processing methods such as federated learning and edge analytics. While it does not involve entirely decentralized storage, the pipeline is designed to be modular and federated. Analytical processes are distributed across interconnected components, allowing different organizations to join without needing a single area of aggregation. This approach has several advantages. It reduces the need for sensitive cross-border data transfers, ensures compliance with local laws, and allows for responsiveness by enabling local processing near the data source.

Moreover, the system's design allows for the modular integration of advanced AI models, enabling future extensions where organizations can implement lightweight versions of entity recognition or predictive models in secure local settings. This distributed yet federated approach ensures scalability, legal compliance, and flexibility across various contexts in North Africa, where trafficking patterns change and data governance varies significantly.

2.5 Ethical Framework and Human Rights–Based Design

The final component of the theoretical framework emphasizes a strong ethical and human rights perspective. This research is grounded in the belief that data collection, processing, and analysis must prioritize the dignity, autonomy, and safety of survivors and vulnerable communities. Mishandling trafficking data can lead to retraumatization, stigmatization, or even legal or physical danger. To reduce these risks, the pipeline uses privacy-preserving techniques, adheres to data minimization principles, and processes only what is essential for research goals.

All datasets are collected through partnerships that follow rigorous ethical review processes

and informed consent protocols. Sensitive identifiers are anonymized or pseudonymized, and only ethically approved data is included in the pipeline. We ensure transparency by informing stakeholders about how data is collected, processed, and analyzed. Additionally, security is a key part of the system's design, with strong encryption, controlled access, and secure processing environments.

By incorporating these rights-based safeguards, the framework not only complies with international data protection laws like the GDPR **?** but also gains social legitimacy in humanitarian settings. Ethics is integral to the technical framework, ensuring the system is reliable and trustworthy for long-term use in North Africa and beyond.

3 Methodology

3.1 System Overview and Scope

This Thesis describes a pipeline from start to finish that is aware of governance issues and that converts heterogeneous, qualitative trafficking research data from its original state into an interoperable, time-aware representation of knowledge and geospatial intelligence while exposing as little data as possible in order to keep it secure. This pipeline was designed with interoperability and accountability as the top priorities, including FAIR-aligned stewardship and machine-aligned actionable provenance (Wilkinson et al., 2016a; Lebo et al., 2013), and adheres to regulatory expectations for privacy, minimization, and accountability under the GDPR (EU-, 2016).

Figure 2 Shows the Dashboard and the menu on the side

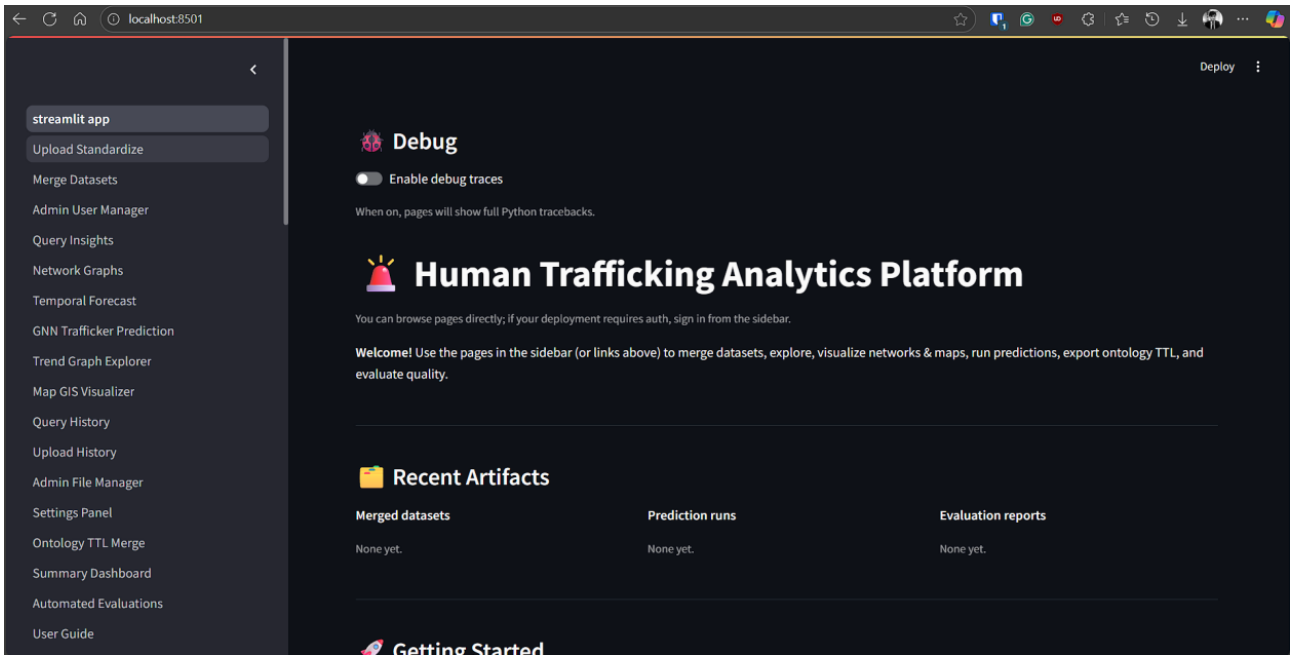

Figure 2. *The Dashboard / Home Page of the Application*

At runtime, a data owner uploads a dataset containing qualitative records. The application loads the file into memory, performs deterministic preprocessing and normalization, and creates a *processed long-format* table. This table retains only the columns essential for analysis, such as serialized identifiers, normalized locations with route order, and coarse durations when present. Also included are cleaned, non-sensitive attributes that are needed for aggregation. The upload is not persisted. It is visible only to the uploader during the active session. After that, it is discarded. All the downstream modules operate exclusively on the processed table. They construct a typed graph in memory, which then is exported as a single static image. Also, they make an interactive HTML visualization and a geospatial module that resolves toponyms and synthesizes map layers for exploratory analysis. Since every artifact can be regenerated

deterministically from the processed table, results remain reproducible without re-exposing to the uploader the raw content (Lebo et al., 2013; Wilkinson et al., 2016a).

Figure 3 Shows the Upload page where user can upload the dataset and once the user uploads the dataset the NLP pipeline is triggered

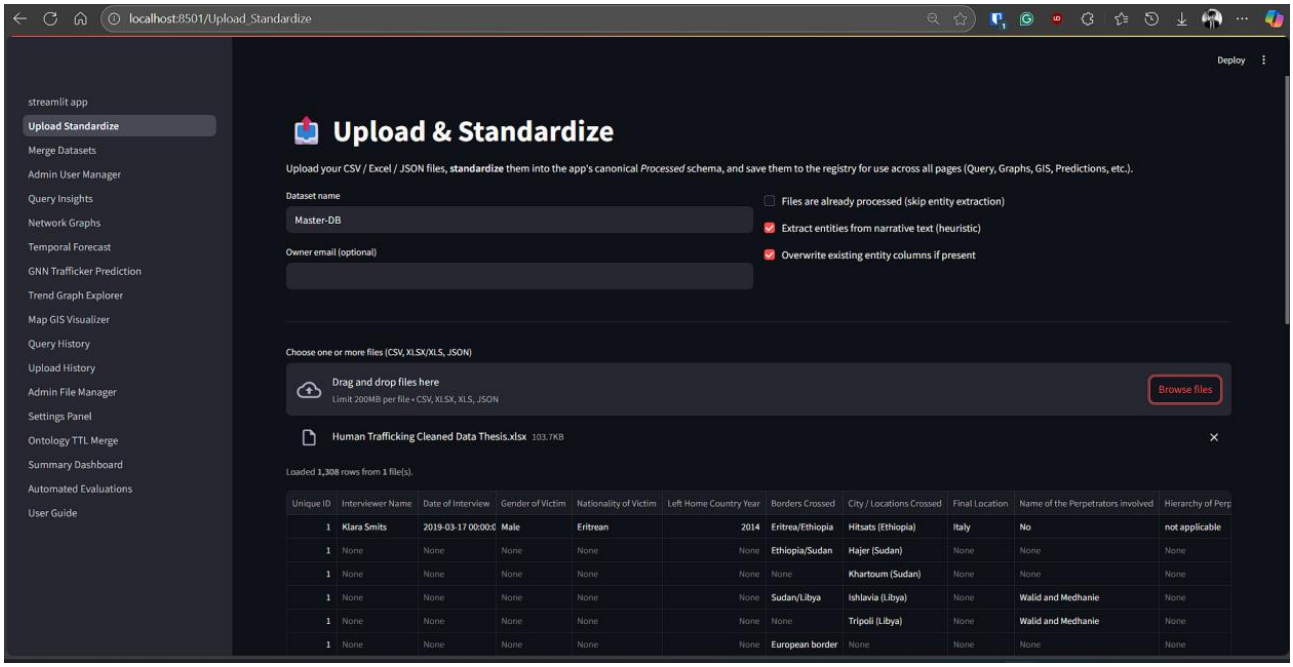

Figure 3. *The upload page of the Application*

The internal representation supports clear structures that are easy to check. Graph construction uses NetworkX, featuring typed nodes such as victim, location, perpetrator, and chief, with labeled edges derived directly from the processed table. It also includes standard descriptive measures that help with understanding (Hagberg et al., 2008). To enable interoperability, an RDF export created with RDFLib sets up classes and properties that match this schema. It also attaches temporal order and WGS84 coordinates to location nodes when available. This setup allows for reuse by semantic tools and meets FAIR objectives (Lebo et al., 2013; Wilkinson et al., 2016a).

The geospatial subsystem turns location strings into coordinates and collects them into layers that reveal spatial concentration and movement. Resolution depends on looking up known place names, with conservative fuzzy matching based on classical edit-distance methods to handle minor spelling variations (Levenshtein, 1966). Mapping is done using Folium, which uses Leaflet to display a base map, a density heatmap, marker clustering to reduce overlap, and animated paths when route order is available. This approach provides a time-aware view of flows without overstating precision (Leaflet, 2025; Folium Project, 2025). Locations that cannot be resolved with confidence are either excluded or flagged to prevent misleading details.

Data handling strictly separates roles and minimizes data use. Raw uploads stay only in

the uploader's session and are not stored. The only saved items are the processed projection and derived exports, which are owned by the uploader and remain private unless shared. Administrative tasks are limited to configuration and health checks, without access to raw data. Storage is kept within the controlled environment; the system does not automatically transfer data across borders. These practices ensure data minimization and purpose limitation, while also reducing potential vulnerabilities by design (EU-, 2016). Provenance links accompany each derived artifact, allowing for audit of analytical lineage without exposing sensitive inputs (Lebo et al., 2013).

The main goal is to change messy, qualitative evidence into a semantically stable and spatiotemporally clear representation. This supports investigative and humanitarian tasks like reconstructing routes, identifying activity hubs, and linking related cases through shared places and actors. This focus aligns with sector guidance that emphasizes integrated, multi-source analysis and clear methods, while also considering the sensitivity of humanitarian data (United Nations Office on Drugs and Crime, 2023b; Counter-Trafficking Data Collaborative, 2025).

The project scope is intentionally limited to avoid overstating claims. Predictive analytics, such as sequence forecasting with recurrent networks, graph neural prediction, or time-of-arrival estimation, are not included. The contribution is descriptive and integrative: an interoperable, reproducible, and privacy-conscious pipeline, with artifacts that can be reproduced and audited independently (Wilkinson et al., 2016a; EU-, 2016).

The following sections explain, in order, the dataset intake and compatibility layer, the governance-aware data handling that puts processes into practice, the preprocessing and canonicalization processes that create the long-format table, graph construction and RDF export, geospatial synthesis, and the orchestration that ensures consistent artifact results. Each module connects to the research questions and goals, and each produces verifiable outputs that support the evaluation plan.

3.2 User Requirements and Task Grounding

The system's features and limits come from a combination of existing literature and accepted practices in requirements engineering. The goal is to provide decision-ready analytics without risking the confidentiality of sensitive trafficking data. The requirements highlight the need for semantic interoperability, privacy, data minimization, strict control over access to raw content, and outputs that can be audited and reproduced to support specific tasks for practitioners. The analysis follows a goals- and artifacts-focused approach, ensuring that every function connects to a stakeholder need and a verifiable system artifact (ISO/IEC/IEEE, 2018; van Lamsweerde, 2001). Governance and stewardship considerations are included from the beginning using principles that promote reuse and transparency, along with machine-actionable provenance and privacy obligations set out in the GDPR (Wilkinson et al., 2016a; Lebo et al., 2013; EU-, 2016).

3.2.1 Primary user tasks

The system is meant to help with three common tasks identified in literature and practice. These are: reconstructing movement across locations based on qualitative records, identifying active spatial hubs, and consolidating potentially related cases through shared places and people. These tasks are emphasized in international monitoring and research synthesis, which highlight the need for integrated, multi-source situational awareness while recognizing the sensitivity of victim-centered data (United Nations Office on Drugs and Crime, 2023b; Counter-Trafficking Data Collaborative, 2025). The implemented process—processed table to typed graph to layered maps—facilitates these tasks without revealing raw content to unintended parties.

3.2.2 Requirement: exclusive control over raw uploads

Data owners need raw datasets to remain inaccessible to other users, including administrators and viewers. They also require that raw content not be stored beyond the uploader's active session. This mirrors GDPR principles of data minimization and purpose limitation, along with the security principle of least privilege Mechanism (EU-, 2016; Saltzer and Schroeder, 1975). The application loads raw files into memory for transformation, only creating a sanitized processed version, while discarding the raw upload afterwards. Role checks in the interface and backend ensure that only the uploader can view the raw content during the active session. No storage or background export of raw fields occurs.

3.2.3 Requirement: processed-only persistence and data minimization

Stakeholders require the system to keep only a reduced, meaningful version of the data for analysis, excluding sensitive fields. This supports FAIR's focus on machine-actionable metadata and aligns with GDPR's data minimization rules Mechanism (Wilkinson et al., 2016a; EU-, 2016). The preprocessing layer generates a long-format table with serialized identifiers, normalized locations in order, approximate durations when possible, and cleaned, non-sensitive attributes needed for aggregation. All derived analytics—graphs, maps, and RDF exports—come strictly from this processed version.

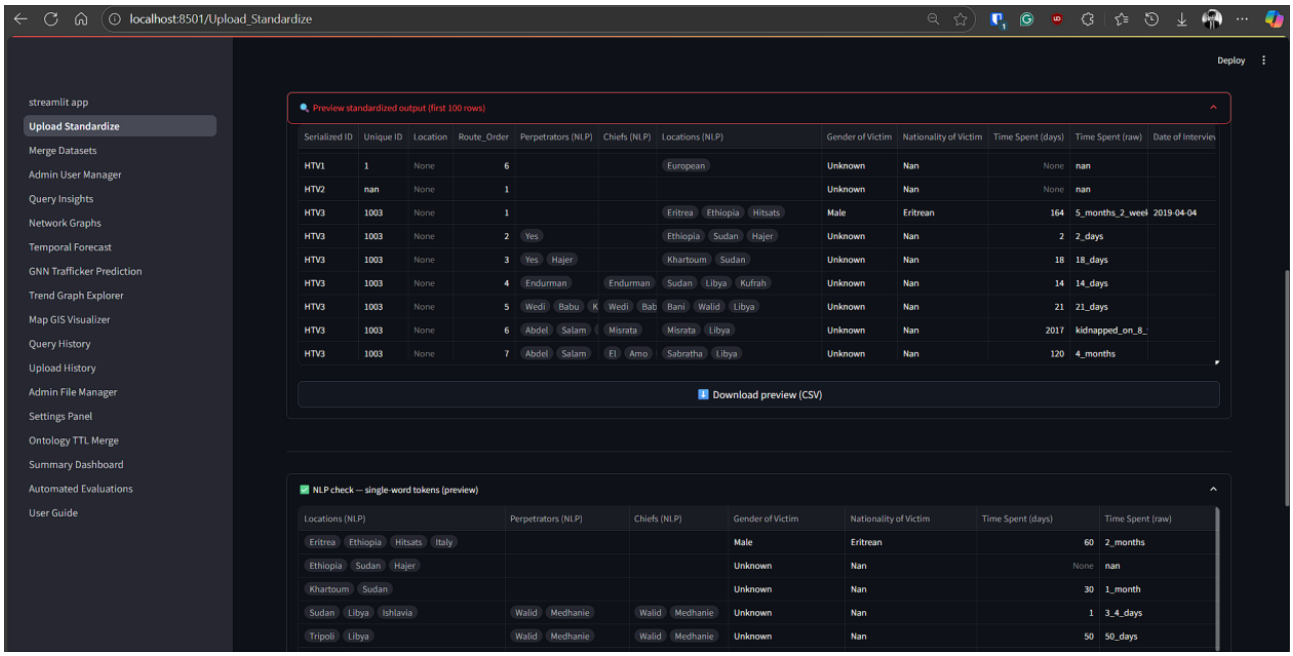

Figure 4. *NLP Extracted infromation from RAW dataset is stored and the RAW dataset is instantly flushed from the system*

3.2.4 Requirement: strict role separation and least privilege

Administrative tasks should not allow access to user data. Viewers and non-owners must not access another user's raw or unshared processed artifacts. Maintaining least privilege and compartmentalization is important for reducing risk and preventing accidental disclosure Mechanism (Saltzer and Schroeder, 1975) . Roles are confined to their designated areas: administrators manage configuration and system health without read access to raw uploads; processed artifacts belong to the uploader and stay private until explicitly shared; non-owners cannot view or access another user's artifacts.

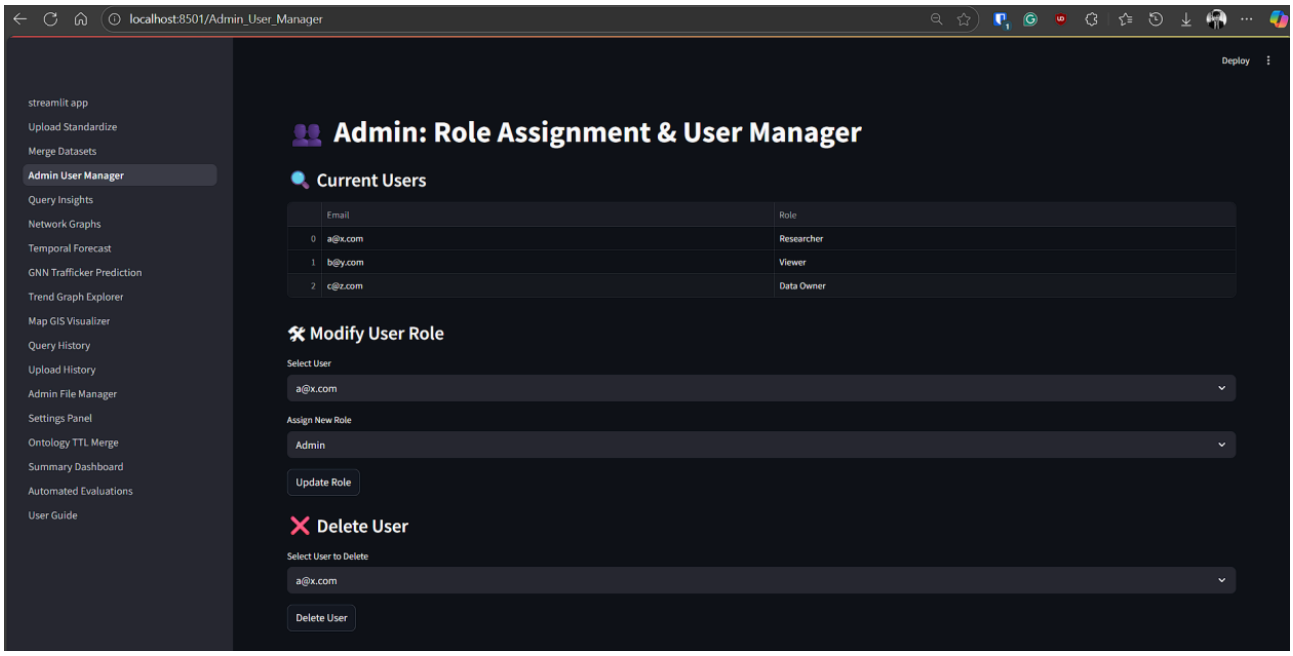

Figure 5. *Admin dashbaord with Role assignment and to further enforce Role Based Access Control (RBACS) and also application ensures that the RAW data is never stored and even the processed version has heavy access control*

3.2.5 Requirement: clear data residency and non-egress by default)

Deployments should not send data outside the hosting environment without the owner's explicit action to prevent cross-border issues and maintain institutional control Mechanism. All persistence targets stay local to the controlled environment. There is no automatic egress, synchronization, or third-party upload. Exports are created as local files for the owner's distribution.

3.2.6 Requirement: semantic coherence and interoperability

Analysts need a stable, auditable representation that retains meaning across records and supports both structural (graph) and spatiotemporal (GIS) queries. FAIR promotes machine-actionable integration, while provenance standards allow for analysis inspection Mechanism (Wilkinson et al., 2016a; Lebo et al., 2013). The transformation step yields a normalized long-format schema; the graph layer creates typed nodes and labeled edges directly from this schema; the RDF export provides a standards-compliant serialization that reflects the same structure for use with semantic tools.

3.2.7 Requirement: reproducibility and auditability without re-exposing raw content

Reviewers should be able to recreate figures, tables, and interactive artifacts from the processed version and examine analytical lineage without accessing raw inputs. The orchestration follows a fixed sequence. Parameters and identifiers accompany every derived artifact, ensuring they can be recreated to achieve the same results given identical inputs. Minimal run metadata

shows lineage while preventing leakage of raw fields (Lebo et al., 2013).

3.2.8 Requirement: uncertainty-aware geospatial rendering

Maps should indicate the uncertainty present in place names from qualitative narratives and avoid giving false precision. International reports highlight the risks of misinterpreting sparse or noisy spatial signals (United Nations Office on Drugs and Crime, 2023b). Locations are determined through explicit lookups, using cautious fuzzy matching only for minor variations. Unresolved or low-confidence entries are either excluded or flagged. Density, clustering, and trajectory layers are provided to support exploratory analysis without implying unwarranted accuracy.

3.2.9 Requirement: simplicity of operation and export

Practitioners need artifacts that can be exported and embedded into reports without requiring specialized tools. The user interface displays an interactive graph for inspection and offers exports as static images and HTML. Geospatial views are created with web-native layers for easy sharing. The processed table and RDF serialization are available as files that the owner can archive or transfer.

3.2.10 Traceability to research questions and evaluation

Each requirement connects to one or more research questions and to corresponding evaluation evidence. Exclusive control over raw uploads, processed-only persistence, role separation, and data residency aligns with governance-focused questions, with evidence from configuration, storage layout, and access tests. Semantic coherence, reproducibility, and exportability align with interoperability and pipeline questions, supported by consistent checks in artifact regeneration and schema/graph/RDF. Uncertainty-aware mapping aligns with the geospatial analysis question, evidenced by resolution diagnostics and exclusion policies. Therefore, this chapter moves from these requirements to the specific design and implementation choices that make them possible in software and to the evaluation plan that ensures they hold in practice.

3.3 Data Sources, Intake, and Compatibility

The pipeline ingests researcher-supplied datasets containing qualitative records and associated tabular fields. Uploads are accepted in common spreadsheet formats (CSV/Excel) and are read into memory for immediate transformation using column-normalization and schema checks. To ensure that downstream analytics operate on a coherent representation, the intake step validates a small set of mandatory fields and normalizes common header variants to canonical names. At minimum, the schema requires a stable victim identifier and a route-bearing location field; in practice, it also expects demographic fields that are non-sensitive in aggregate form and that support descriptive analyses. Concretely, the intake layer normalizes headers and verifies the presence of key columns, including *Unique ID*, *City / Locations Crossed*, and, when available, *Time Spent in Location / Cities / Places*, along with basic descriptors such as *Gender of Victim* and *Nationality of Victim*. Header normalization collapses whitespace and maps frequent aliases (for example, *City/Locations Crossed* $\rightarrow$ *City / Locations Crossed*) so that

semantically equivalent inputs are treated uniformly. Datasets that fail validation are rejected with a clear message indicating missing fields, which prevents silent misalignment at later stages.

Intake is designed to preserve confidentiality while maximizing reproducibility. Raw files are parsed into an in-memory `DataFrame` for transformation and are not persisted. The transformation yields a *processed long-format* table in which each row represents a single step of a victim's route through a location, with an explicit order index and a serialized identifier propagated to all steps for the same case. This design follows well-established principles of "tidy" data organization—for which each variable has its own column and each observational unit its own row—thereby simplifying downstream joins, aggregation, and validation (Wickham, 2014). The use of mature tabular tooling (`pandas`) ensures stable, explicit type coercions and deterministic operations necessary for auditability and reruns (McKinney, 2010).

Because additional datasets may need to be integrated over time, the intake module provides a compatibility layer that maps project-specific or partner-supplied headers into the canonical schema and coerces list-like cells into normalized Python lists. For example, narrative-derived lists such as perpetrators or chiefs are split using conservative tokenization rules that tolerate commas, semicolons, and bracketed forms, followed by trimming and de-duplication under case-insensitive comparison with optional near-duplicate collapse when a fast similarity function is available. Locations are cleaned as strings at intake but not resolved to coordinates until the geospatial module executes; this separation keeps the compatibility shim focused on structural harmonization rather than inference.

When multiple sources are loaded for comparative analysis, the intake layer assigns each dataset a source identifier and records a user-defined precedence order for conflict resolution. Downstream utilities can then merge sources deterministically using simple, documented strategies, such as keeping values from higher-priority sources, uniting list-valued fields where appropriate, or retaining the last declaration by precedence. Prior to merge, a grouping key built from victim identifier, route order, and location allows the system to detect potential conflicts across sources for manual review. This approach acknowledges that anti-trafficking datasets are often assembled under heterogeneous practices and that transparent, reproducible merge rules are preferable to opaque heuristics.

The result of intake and compatibility is a single processed table that is semantically coherent, explicit in its observational unit, and stripped of fields that are not required for analysis. Only this processed projection is eligible for persistence; the original upload remains confined to memory during the session and is discarded thereafter. By coupling tidy-data organization with strict minimization and by avoiding raw-file persistence, the intake design supports both FAIR-aligned reuse of derived artifacts and privacy-by-design handling of sensitive content (Wilkinson et al., 2016a). In addition, by enforcing explicit, documented conversion steps at the boundary, the system reduces the likelihood of silent drift between partners' data definitions and the analysis-ready schema used throughout the thesis, reflecting best practice for transparent, auditable data preparation in computational social science.

3.4 Governance-by-Design Data Handling

Data handling in the pipeline revolves around three core principles: data minimization, strict separation of roles, and accountable derivation. The system processes raw uploads in memory to create a reduced, analysis-ready version and only saves this *processed* table along with its downstream artifacts. This design meets GDPR requirements for limiting purpose and minimizing data, while also avoiding unnecessary retention of sensitive information (EU-, 2016). To maintain tracking of analysis without exposing raw inputs, every derived artifact includes machine-actionable provenance links to the processed version instead of the original upload. This aligns with FAIR and W3C PROV guidance (Wilkinson et al., 2016a; Lebo et al., 2013).

3.4.1 Processed-only persistence and minimization.

When a dataset is uploaded, it is parsed into memory, validated, and transformed into a long-format table. This table contains only the fields necessary for analysis: serialized identifiers, normalized locations in route order, approximate durations when available, and a small set of non-sensitive descriptors. Sensitive columns from the original file are excluded by design. The raw file is not saved to disk and becomes inaccessible after the transformation finishes. Saved artifacts—processed table, graph exports, RDF/TTL serializations, and GeoJSON—are projections that do not contain any raw text or direct identifiers. This enforces minimization at the storage level and ensures downstream operations work with only the necessary data (EU-, 2016).

3.4.2 Role separation and least privilege.

Access to the data and artifacts follows the principle of least privilege (Saltzer and Schroeder, 1975). The uploader, or data owner, is the only person who can see the raw upload during their active session; no other users, including administrators, can access the raw information. Administrative functions are limited to configuration and health monitoring and do not include access to user data. Processed artifacts belong to the uploader and cannot be discovered or read by anyone else unless sharing is explicitly allowed. The application performs checks when reading or exporting data to ensure that even metadata listings do not reveal raw field values or file paths. These decisions reduce potential vulnerabilities and match operational actions with privacy-by-design goals (EU-, 2016; Cavoukian, 2010).

3.4.3 Residency and non-egress by default.

All saved artifacts are kept within the controlled deployment environment; there is no automatic transfer to external services, and no background synchronization occurs. Exports, like the interactive graph HTML or map data, are created as local files linked to the owner's account. This approach allows the data owner to make decisions about releasing their data and avoids cross-border transfers by default. When collaboration is necessary, owners share *processed* artifacts, not raw inputs, which keeps the same level of risk outside the application.

3.4.4 Provenance without raw re-exposure.

To support verification and repeated analyses, the system logs machine-actionable provenance related to the transformations applied to the processed table and the parameters that generated each artifact. Provenance records include stable identifiers for the processed dataset, transformation steps, and configuration hashes while intentionally excluding raw-field echoes or free text. This allows for the independent reproduction of figures and tables from the processed version without including sensitive inputs in logs or audit trails (Lebo et al., 2013; Wilkinson et al., 2016a).

3.4.5 Logging and operational metadata.

Operational logs capture minimal, non-content metadata such as timestamps, status codes, and anonymized dataset IDs. Error messages are crafted to provide helpful information for debugging but do not include raw values; for example, they may truncate or redact user-provided strings. This meets needs for accountability and forensic review while preserving data minimization (EU-, 2016).

3.4.6 Threat model boundaries and non-goals.

The governance approach aims to address practical risks in a research setting, including unintentional disclosures through storage or logs, excessive administrative read access, and uncontrolled sharing of detailed raw text. The system does not claim to provide specialized defenses against nation-state threats or to implement advanced cryptographic techniques; these features would require different constraints and are outside its scope. Instead, the focus is on strong defaults, such as processed-only persistence, owner-only visibility of raw data, and non-egress storage. It also emphasizes clear accountability through provenance on every artifact and simplicity in operations to ensure privacy safeguards are effective in everyday use (Cavoukian, 2010; EU-, 2016).

3.5 Dataset Registry and Provenance Artifacts

The pipeline keeps a simple registry that only records analysis-ready artifacts and the minimal metadata needed to reproduce them. Its purpose is twofold. First, it makes exports easy to find and rerun without keeping raw inputs. Second, it attaches machine-readable lineage to every figure or table that appears later in the thesis. The design follows principles of good data stewardship, such as persistent identifiers, clear labeling of artifacts, and enough metadata for reuse. It represents derivations in a way that is compatible with W3C PROV. This makes analytical steps clear and allows for audits (Wilkinson et al., 2016a; Lebo et al., 2013). In line with good practices for reproducible computational research, the registry values simplicity and predictability over complicated systems. It captures just enough structure to support independent verification (Sandve et al., 2013; Wilson et al., 2017).

3.5.1 Registry scope and entities.

Entries are created only for the *processed* dataset and its derivatives. For each upload session, the system creates a processed long-format table that excludes sensitive fields. This

table is registered as a dataset entity with a stable identifier and a brief description of its schema version. From this single source, the application produces a small group of derived artifacts: a static graph image, an interactive graph HTML view, an RDF serialization (TTL) that mirrors the graph structure, and a GeoJSON stream for map rendering. Each derivative has its own registry entry that points back to the original processed dataset. No raw uploads are registered or saved, and no registry field stores raw text or direct identifiers.

3.5.2 Identifiers, locations, and minimal metadata.

Every registry entry holds three types of information: identity, storage, and lineage. Identity includes a local unique identifier, a human-readable label, the artifact type (processed table, graph PNG/HTML, RDF/TTL, GeoJSON), and timestamps for creation and last access. Storage notes the file path within the controlled environment and the owning account, allowing the system to manage access without revealing content. Lineage connects each derivative to its original processed dataset and captures a compact parameter snapshot needed to recreate the artifact. For instance, it records the chosen layout seed for the graph view or the map layer switches that were enabled. The parameter snapshot is intentionally narrow and avoids repeating any raw fields; its only purpose is to ensure that "same inputs, same outputs" is a documented feature of the workflow (Sandve et al., 2013; Wilson et al., 2017).

3.5.3 Provenance model and audit trail.

In concept, the registry's lineage fields match the PROV-O pattern of *Entity*, *Activity*, *Entity*. A processed dataset (*Entity*) is used by a rendering or export step (*Activity*) to create a graph image, HTML view, RDF file, or GeoJSON stream (*Entity*). The registry captures this with explicit back-links from each derivative to its source dataset and a brief activity description (tool, version, and key parameters). Since all derivations start from the same processed table, auditors can check that no export relies on unregistered inputs. This level of provenance is enough to rerun the pipeline and to explain in the manuscript how each figure or table was made (Lebo et al., 2013).

3.5.4 Registration workflow and retrieval.

Registration happens when a user creates an artifact through the user interface. When a user exports the processed table, generates a graph view, saves the RDF/TTL, or writes GeoJSON for the map, the application stores the file in a set location under the owner's workspace and adds a registry entry with the identity, storage, and lineage information. Later components find artifacts by querying the registry using the processed dataset identifier, ensuring that only compatible exports appear together. For example, the graph and map generated from the same processed table will surface next to each other. Because the registry only knows about processed and derived artifacts, a search can never return raw content.

3.5.5 Privacy-preserving metadata and retention.

To limit risk, the registry keeps only operational metadata. Labels are short and neutral, such as a user-supplied dataset nickname, and paths are restricted to the owner's area. There is no inclusion of free text or column values. Retention aligns with the usefulness of artifacts:

processed tables and their exports remain until the owner deletes them, allowing for the reproduction of manuscript figures. Meanwhile, temporary UI previews that are not exported leave no trace. This method meets data minimization and purpose limitation requirements while keeping enough information for recomputation and review (Wilkinson et al., 2016a; Wilson et al., 2017).

3.5.6 Consistency and failure handling.

To prevent partial writes, such as those from an interrupted export, the application registers an artifact only after the file operation succeeds. Each entry includes a simple integrity check, like file size and modification time, which later components check before loading. This reduces the chance of using outdated or incomplete exports in analysis. If a referenced file is missing or fails the integrity check, the registry marks the entry as invalid and provides a clear path for repair: regenerate the artifact from the same processed dataset using the recorded parameters. This process keeps the registry consistent without adding complicated transactional systems, aligning with the project's focus on transparent, auditable workflows (Sandve et al., 2013).

3.5.7 Implications for the thesis.

Since every graph image, interactive view, RDF file, and GeoJSON stream in later chapters is represented in the registry with clear lineage to a single processed dataset, the thesis can cite stable identifiers and parameter snapshots alongside figures and tables. Readers and reviewers gain a direct way to recompute results: load the processed dataset and rerun the recorded export without needing access to raw inputs. In this way, the registry supports both result reproducibility and source confidentiality, reflecting the methodological approach taken throughout the work (Wilkinson et al., 2016a; Lebo et al., 2013).

3.6 NLP Preprocessing, Normalization, and Canonicalization

This subsection describes the clear text and table transformations that turn an uploaded spreadsheet into the *processed long-format* table used by all following modules. The procedures are intentionally simple, including rule-based normalization, list tokenization, canonicalization, and route expansion. This design allows for auditing and reproducing the process without showing raw free text. Implementation uses established tabular tools for clear typing and stable sorting (McKinney, 2010) and organizes the outcome in a "tidy" observational structure that is suitable for joins, validation, and export (Wickham, 2014).

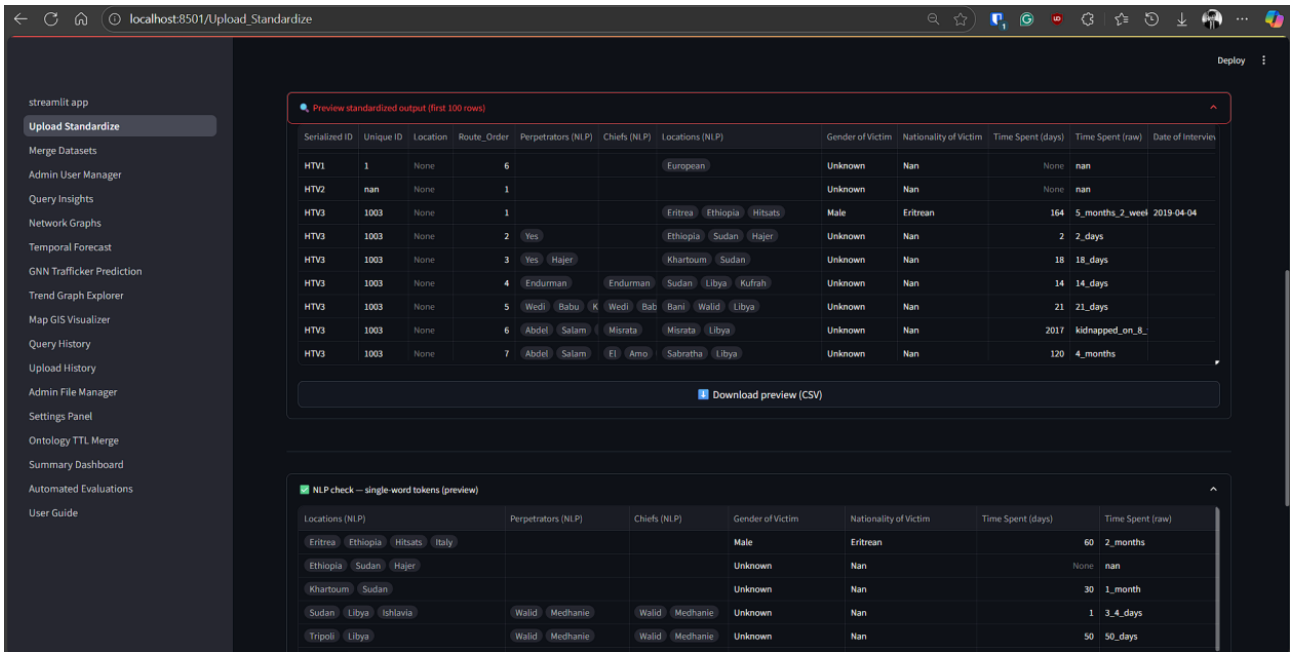

Figure 6. *The upload page of the application post NLP stage where it extracts all vital information and drops all the RAW Columns, it never stores the original daatset whatsoever, instead a very porcessed version that doesn't have any sensitive data is stored*

3.6.1 Input notation and goal.

Let the raw upload be a table $D = \{r_i\}_{i=1}^{n}$ with headers H and rows r_i; each row contains at least a stable case identifier and a route-bearing location field. The goal is to produce a *processed* table

$$\widehat{D} = \{\,(\texttt{SerializedID}, \texttt{UniqueID}, \texttt{Location}, \texttt{Route\_Order},\,\}$$

$$\widehat{D} = \{\,(\texttt{Time\_Spent}, \texttt{Perpetrators}, \texttt{Chiefs}\,\}$$

where each row represents one step in a route, $\texttt{Route\_Order} \in \mathbb{N}$ gives the within-case sequence, and list-valued fields are normalized to de-duplicated lists of canonical strings. Sensitive columns from D are *not* propagated.

3.6.2 Header normalization and schema validation.

A normalization map $\phi : H \to \widetilde{H}$ coerces common header variants to canonical names (e.g., *City/Locations Crossed* $\mapsto$ `City / Locations Crossed`). Minimal schema validation requires the presence of `Unique ID` and `City / Locations Crossed`, and optionally `Time Spent in Location / Cities / Places` plus safe demographic descriptors. Rows failing mandatory-field checks are rejected with explicit diagnostics to prevent silent drift.

3.6.3 Identifier serialization.

A stable serializer $s : \Sigma^{\to\Sigma}$ produces `SerializedID` by (i) Unicode NFC normalization, (ii) trimming, (iii) internal whitespace collapse, and (iv) case folding:

$$s(x) = \mathrm{lower}\big(\mathrm{collapse\_ws}(\mathrm{trim}(\mathrm{nfc}(x)))\big).$$

Serialized IDs propagate to all exploded route rows, providing a deterministic key for joins and exports.

3.6.4 List tokenization and canonicalization.

For list-valued cells (e.g., perpetrators or chiefs), tokenization uses a conservative regular expression that tolerates common delimiters while avoiding token splits inside bracketed substrings:

$$\mathcal{T} = \texttt{/[,;|]/} \quad \text{with guarded splits on } \texttt{[()]}.$$

Each token t is canonicalized by

$$c(t) = \text{lower}\big(\text{strip\_punct}(\text{collapse\_ws}(\text{trim}(t)))\big),$$

and the cell value becomes an ordered, de-duplicated list $\text{uniq}([\, c(t)\,]_{t\in\mathcal{T}(x)})$. To mitigate near-duplicates differing by minor edits, an optional bounded edit-distance merge replaces tokens u, v with a representative when

$$\text{lev}(u, v) \le \tau,$$

where lev is the Levenshtein distance (Levenshtein, 1966; Wagner and Fischer, 1974) and $\tau \in \{1, 2\}$ for conservative collapse. This step is deterministic given a total order on candidates (e.g., keep the lexicographically smaller form).

3.6.5 Location string normalization and route expansion.

Location cells are cleaned by the same $c(\cdot)$ and then exploded into a sequence $L_i = [\ell_{i,1}, \ldots, \ell_{i,k_i}]$ using $\mathcal{T}$. The *route expansion* is the cartesian composition of case identifiers with ordered locations:

$$\text{expand}(r_i) \mapsto \big\{(\texttt{SerializedID}_i, \texttt{UniqueID}_i, \ell_{i,j},\ \texttt{Route\_Order} = j)\big\}_{j=1}^{k_i}.$$

If the source provides a final terminal place separate from the crossing list, the expansion appends it as ℓ_{i,k_i+1} with $\texttt{Route\_Order} = k_i + 1$. This representation directly supports graph construction and geospatial synthesis while keeping the original free-text field out of persistence.

3.6.6 Duration parsing (descriptive, not predictive).

When present, narrative durations are parsed by a simple finite set of unit patterns and coerced to days:

$$\text{days}(q, u) = q \cdot \gamma(u), \quad \text{where } \gamma(\text{day}) = 1,\ \gamma(\text{week}) = 7,\ \gamma(\text{month}) \approx 30,\ \gamma(\text{year}) \approx 365.$$

Extraction uses regexes such as `/(\d+)*(day|week|month|year)s?/i`. The result is stored as a coarse integer `Time_Spent` for descriptive summaries only; no ETA or forecasting is performed.

3.6.7 Type coercions and invariants.

Columns in $\widehat{D}$ have specific types: identifiers are strings, `Route_Order` is an integer, list-valued fields are JSON-serializable arrays, and durations are integers representing days. Invariants checked after transforming include: (i) the uniqueness of (`SerializedID`, `Route_Order`), (ii) non-empty `Location`, and (iii) list elements that are already canonicalized, which means no leading or trailing spaces and all in lowercase. Any violations trigger actionable diagnostics.

3.6.8 Deterministic caching and run fingerprints.

To allow recomputation without saving raw inputs, the system keeps a run fingerprint $h(\widehat{D})$. This fingerprint is a cryptographic hash of a canonical serialization of the processed table, which includes the schema version, sorted rows, and stable encodings. With the same inputs and parameters, the pipeline is idempotent. It creates the same $\widehat{D}$ and identical downstream outputs, such as the graph, RDF, and GeoJSON. This ensures that "same inputs, same outputs" reproducibility is maintained (Wilson et al., 2017; Sandve et al., 2013).

3.6.9 Algorithmic specification.

Algorithm 1 summarizes the transformation. Its worst-case time is $O(NL + M\,\tau)$ where N is the number of rows, L the average number of location tokens per row, M the total list tokens across perpetrator/chief fields, and τ the small constant threshold used in bounded edit-distance checks.

Algorithm 1 Deterministic preprocessing, normalization, and canonicalization

Require: Raw table D, header map ϕ, delimiter set $\mathcal{T}$, edit-distance threshold $\tau \in \{0, 1, 2\}$
Ensure: Processed long-format table $\widehat{D}$

$D \leftarrow$ normalize_headers(D, ϕ); **assert** required columns exist
for all row r in D sorted by `Unique ID` **do**
  sid $\leftarrow s(r[$`Unique ID`$])$
  $L \leftarrow$ explode_locations$(r[$`City / Locations Crossed`$], \mathcal{T})$
  $L \leftarrow [\, c(\ell) \mid \ell \in L \,]$ {canonicalize locations}
  Perp $\leftarrow$ uniq$([\, c(t) \,])$ from tokenizing $r[$`Perpetrators`$]$
  Chiefs $\leftarrow$ uniq$([\, c(t) \,])$ from $r[$`Chiefs`$]$
  if $\tau > 0$ **then**
    Perp $\leftarrow$ merge_near_dups(Perp, lev, τ)
    Chiefs $\leftarrow$ merge_near_dups(Chiefs, lev, τ)
  end if
  $T \leftarrow$ parse_duration_days$(r[$`Time Spent in Location / Cities / Places`$])$ {optional}
  for $j \leftarrow 1$ **to** $|L|$ **do**
    append to $\widehat{D}$: (sid, $r[$`Unique ID`$], L[j], j, T$, Perp, Chiefs, ...)
  end for
end for
enforce types and invariants on $\widehat{D}$; compute $h(\widehat{D})$; **return** $\widehat{D}$ =0

3.6.10 Quality checks (post-transform).

Simple, fast diagnostics run on $\widehat{D}$ to catch inconsistencies before the graph or GIS stages. These include counts of unresolved or empty locations, checks for SID-UID pairs, duplicate route orders for each SID, and summaries of coverage for list fields. The user interface reports these checks for corrective action without showing any raw free text.

3.6.11 Discussion.

The selected operations promote transparency. Regex-based parsing, conservative edit-distance collapse, and explicit typing are easy to review and repeat. This aligns with best practices in reproducible computational research (Sandve et al., 2013; Wilson et al., 2017). Each transformation is expressed as a total function on well-defined inputs—headers, tokens, route lists. The pipeline ensures that downstream artifacts (graph exports, RDF/TTL, GeoJSON) depend only on the processed projection, which helps maintain privacy while improving audit processes.

3.7 Processed Schema and Sanitized Projection

This section describes the processed table that all downstream modules use. The schema is designed to be (i) minimal, containing only essential, non-sensitive fields; (ii) tidy, with each variable in its own column and each observation in its own row; and (iii) reproducible, featuring stable types, explicit constraints, and provenance links to derived artifacts (Wickham, 2014; Lebo et al., 2013; Wilkinson et al., 2016a). To follow data-minimization and purpose-limitation principles, raw qualitative fields and direct identifiers are not kept. Analysis works only with this cleaned projection (EU-, 2016).

3.7.1 Authoritative column dictionary.

Column	Type	Semantics (notes)
`SerializedID`	string	Deterministic, hashed/normalized case identifier propagated to all exploded rows for the same case; safe for joins and exports.
`UniqueID`	string	Original case key (cleaned). Retained as a stable reference; never used to reconstruct raw text.
`Location`	string	Canonicalized toponym (lowercased, trimmed, punctuation-normalized). Represents one step in the route.
`Route_Order`	int	1-based within-case sequence index for the route; strictly increasing per `SerializedID`.
`Time_Spent` (opt.)	int (days)	Coarse duration parsed from narrative units when present; used for descriptive summaries only, not prediction.
`Perpetrators`	list[string]	Canonicalized tokens (keyword-like) for alleged perpetrators; de-duplicated, optionally near-duplicate merged.
`Chiefs`	list[string]	Canonicalized tokens for chiefs/handlers; same normalization as `Perpetrators`.
`Gender_Victim` (opt.)	string	Coarse label after normalization (e.g., `female`/`male`/`unspecified`).
`Nationality_Victim` (opt.)	string	Canonical country/region label after normalization; used for aggregate stratification only.
`SourceID`	string	Short identifier of the intake source (dataset label) to support multi-source merges and audit.
`Schema_Version`	string	Version stamp of the processed schema used to create this table.

Table 2. *Processed long-format schema used across graph, ontology/RDF, and GIS modules.*

3.7.2 Sanitization principles.

Only essential, non-sensitive, *keyword-like* fields are kept. Free text, narrative passages, names that could reveal identity, contact details, and any columns not needed for analysis are removed by design. Canonicalized list fields (`Perpetrators`, `Chiefs`) contain normalized tokens

instead of exact strings. The design ensures that downstream outputs, such as graph exports, RDF/TTL, and GeoJSON, come from this sanitized version and do not require access to raw uploads (Wilkinson et al., 2016a; EU-, 2016).

3.7.3 Shape constraints and invariants.

To keep the table clear and easy to join:

- **Key uniqueness.** The pair `(SerializedID, Route_Order)` is unique; `Route_Order` strictly increases within each `SerializedID`.
- **Non-emptiness.** `Location` is not empty after canonicalization; empty or unresolved entries are removed upstream.
- **Type discipline.** Lists are JSON-serializable arrays of lowercased tokens; durations are non-negative integers measured in days; IDs are normalized strings.
- **Versioning.** `Schema_Version` is marked on the table and reflected in export metadata so figures can be recreated against the same specification.

3.7.4 Downstream consumption.

The schema is the only contract for all modules:

- **Social network graph (Sec. 3.9).** Nodes are created from `SerializedID` (victim), `Location`, and token lists (`Perpetrators`, `Chiefs`); edges come from co-occurrences per route step and per case. Graph exports include `Schema_Version` and the processed dataset identifier.
- **Ontology & RDF (Sec. 3.10).** Classes and properties reflect the columns (e.g., `Visit` with `Route_Order`); optional latitude/longitude are added later in GIS exports, not to the core table, to respect minimization.
- **GIS spatio-temporal (Sec. 3.11).** GeoJSON is built from `Location` and `Route_Order` after toponym resolution; time-aware layers use these fields, following web mapping standards for coordinates and features (Butler et al., 2016).

3.7.5 Interoperability and provenance.

Each exported item includes a simple provenance record linking back to the specific processed dataset ID and `Schema_Version`. This allows "same inputs, same outputs" recreation without referring to raw fields (Lebo et al., 2013; Wilkinson et al., 2016a). Because the schema is organized and clear, joining across modules is straightforward and less prone to errors (Wickham, 2014).

3.7.6 Omissions and redactions (by design).

Columns often found in raw uploads—narrative text, detailed personal descriptions, free-form notes, contact information, and any field unnecessary for graph or map creation—are intentionally left out. If a partner dataset includes additional attributes, they are considered at intake but will not be included in the processed version unless specifically approved as non-sensitive and necessary for analysis. This strategy balances analytical usefulness with privacy and reduces vulnerabilities while ensuring full reproducibility of derived outputs (EU-, 2016).

3.8 Social Network Graph: Construction and Visualization

The application builds a typed, weighted network from the processed table in Section 3.8 and makes it available for exploration. Construction and export are predictable, meaning that the same inputs and parameters will produce the same outputs. Graph creation uses NetworkX (Hagberg et al., 2008). Users can view the graphs interactively through a browser-based renderer called PyVis, which is set up for repeatability. The layout uses a force-directed method developed by Fruchterman and Reingold, which is selected for clarity on small to medium graphs (Fruchterman and Reingold, 1991).

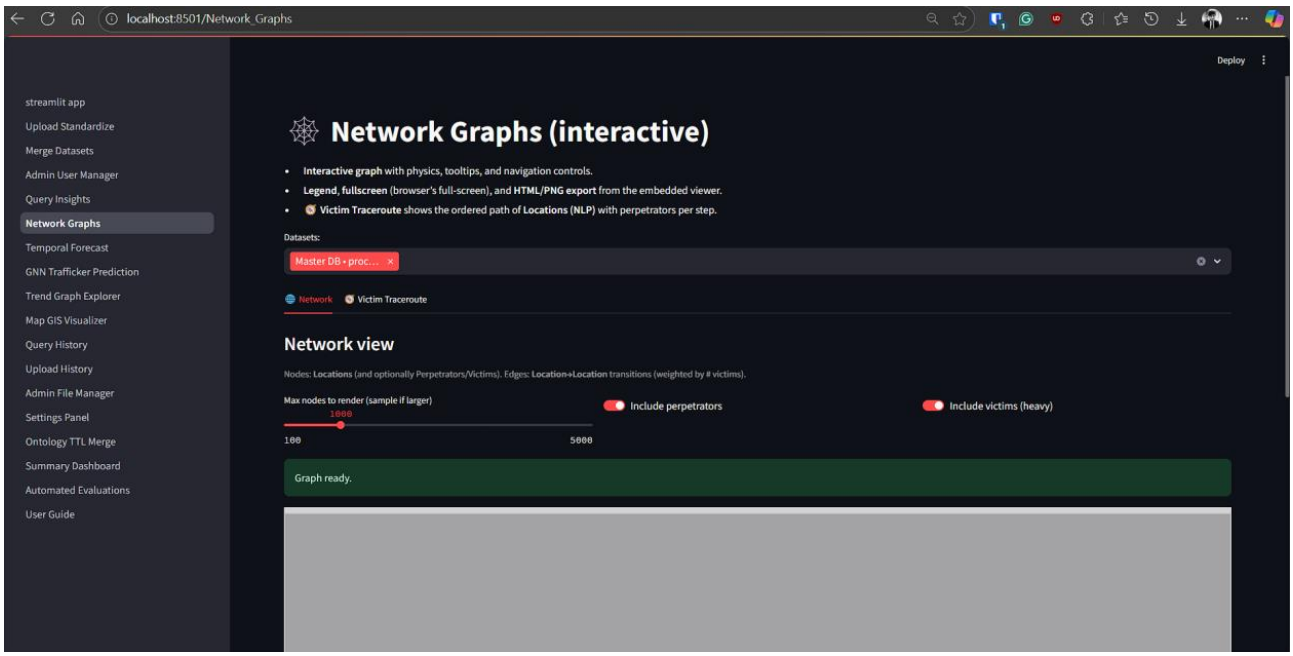

Figure 7. *The Network Graph generation page, user can select entities to be rendered and also the max nodes to render*

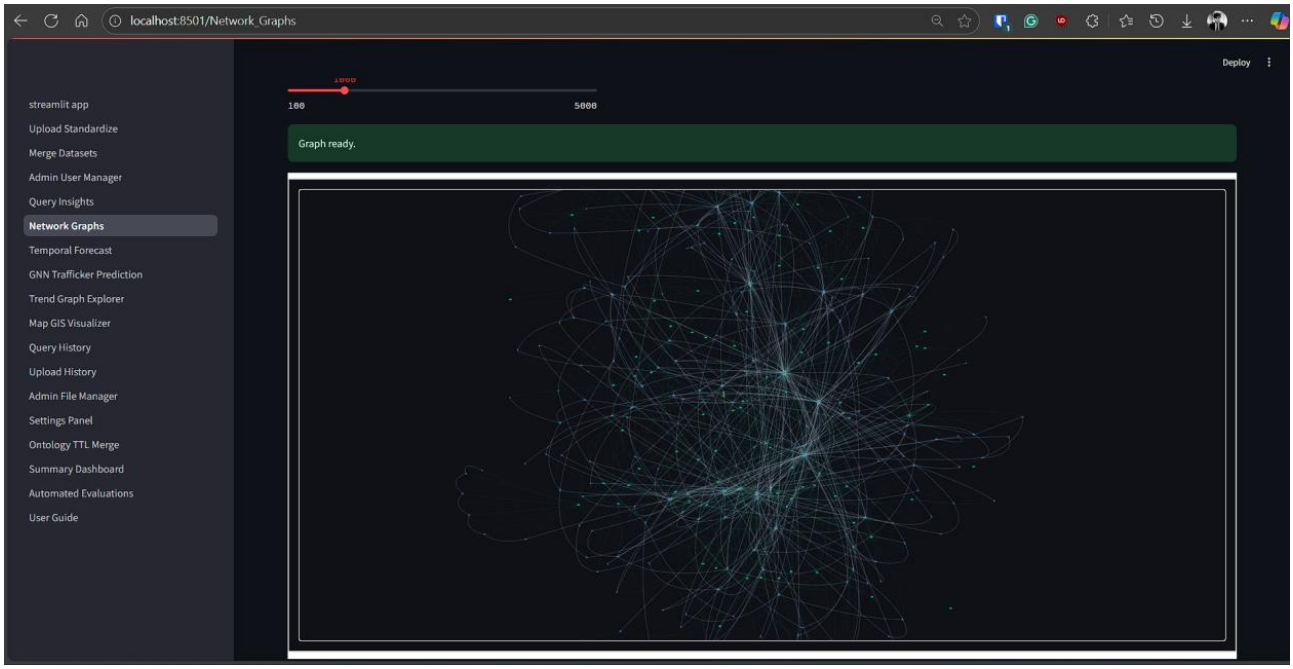

Figure 8. *Shows the generated social network graph that has been created by connecting nodes with victims, nodes, perpetrators and chiefs, all entities are connected to each other and form a social network graph that visualizes the network*

3.8.1 Formalization.

Let V_{S} denote victim nodes (keyed by `SerializedID`), V_{L} location nodes (keyed by canonical `Location`), V_{P} perpetrator tokens, and V_{C} chief tokens. The node set is

$$V = V_{\mathrm{S}} \cup V_{\mathrm{L}} \cup V_{\mathrm{P}} \cup V_{\mathrm{C}}.$$

Edges arise only between victims and the other types, producing a union of bipartite layers:

$$E \subseteq (V_{\mathrm{S}} \times V_{\mathrm{L}}) \cup (V_{\mathrm{S}} \times V_{\mathrm{P}}) \cup (V_{\mathrm{S}} \times V_{\mathrm{C}}).$$

For each processed row $x = (\texttt{SID}, \texttt{Location}, \texttt{Route\_Order}, \texttt{Perpetrators}, \texttt{Chiefs})$, the system emits

$$(\texttt{SID}, \texttt{Location}), \quad (\texttt{SID}, p)\ \forall\, p \in \texttt{Perpetrators}, \quad (\texttt{SID}, c)\ \forall\, c \in \texttt{Chiefs}.$$

Edge weights count co-occurrence across route steps and cases:

$$w(u, v) \;=\; \sum_{x \in \widehat{D}} \mathbf{1}\big[(u, v) \in x\big], \qquad (u, v) \in E.$$

Node size in the visualization scales with weighted degree

$$\deg_w(v) = \sum_{(v,u) \in E} w(v, u),$$

and tooltips report simple descriptors (degree, type, distinct victims per location). Community detection or advanced centralities are not displayed unless explicitly enabled; the default statistics include degree and connected-component summaries.

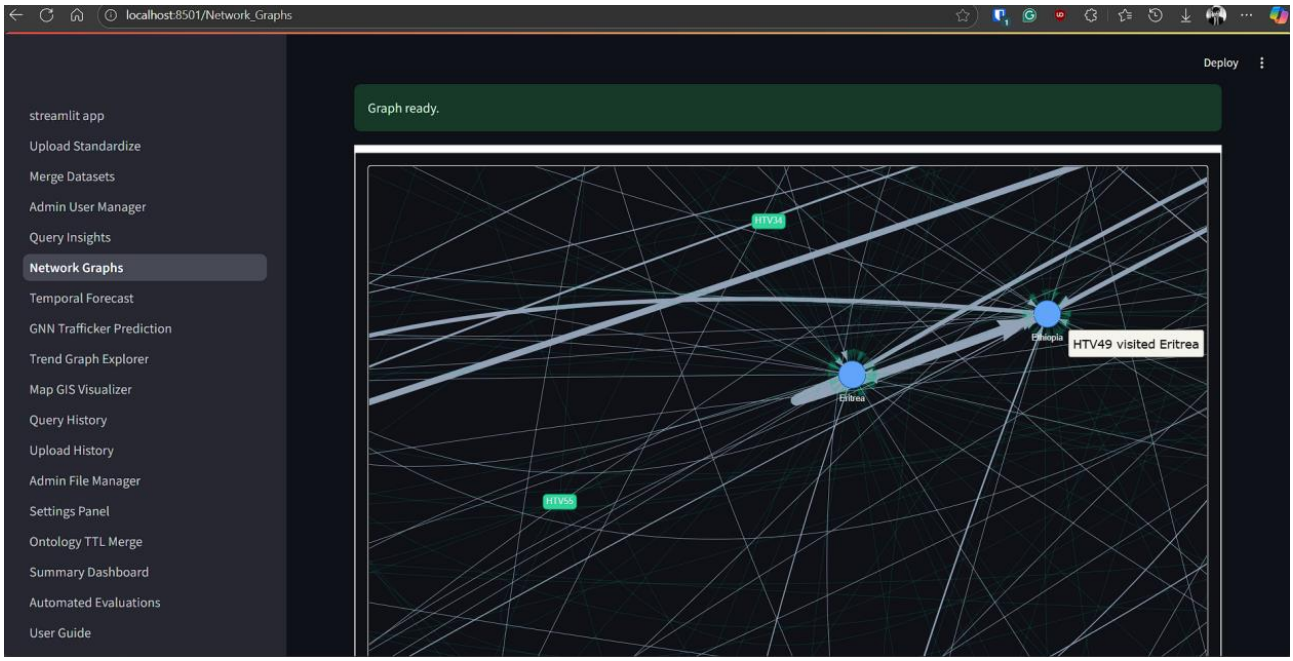

Figure 9. *Zoomed in version of the graph showing two nodes connected to each other and also the edge describing how a victim traveled from one location to other*

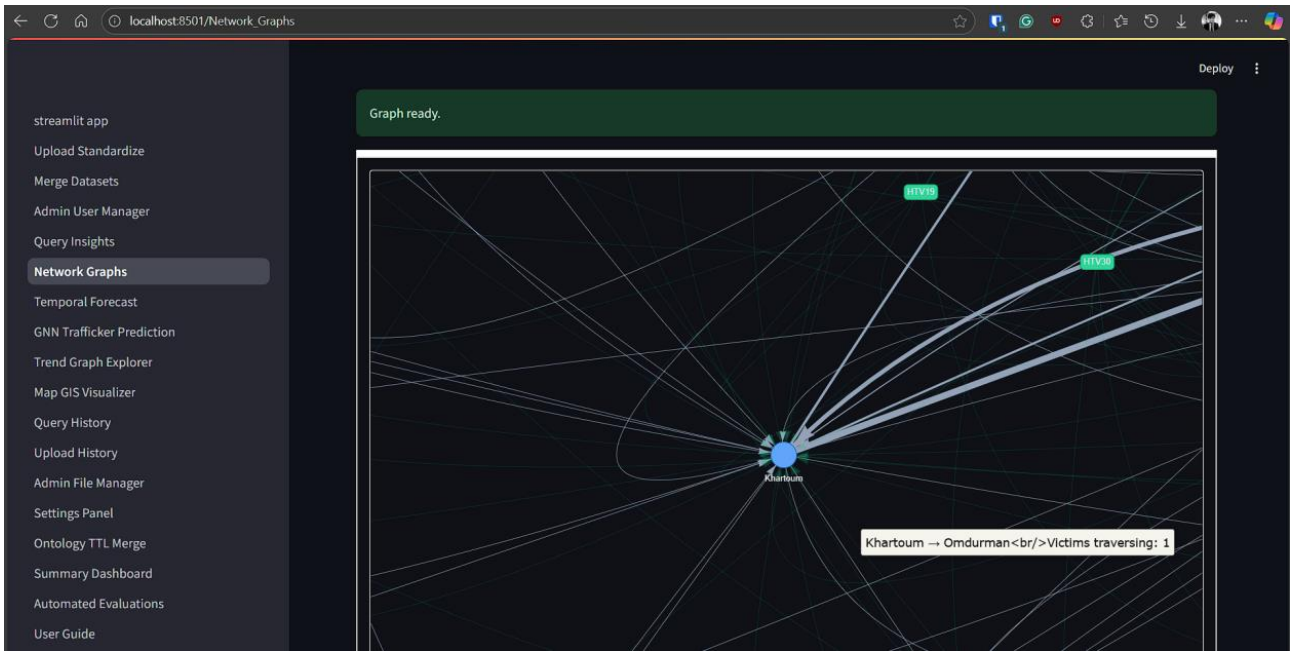

Figure 10. *This zoomed version of graph shows an edge which describes a victim traveling from Khartoum to Omdurman*

3.8.2 Algorithm: typed, weighted graph construction.

The build traverses $\widehat{D}$ in a stable order and increments integer weights. The procedure is linear in the number of processed rows plus total list tokens.

Algorithm 2 Deterministic construction of a typed, weighted graph

Require: Processed table $\widehat{D}$ with columns: `SerializedID`, `Location`, `Route_Order`, `Perpetrators`, `Chiefs`
Ensure: Typed, weighted graph $G = (V, E, w)$
0: $G \leftarrow$ empty graph
0: **for all** row x in $\widehat{D}$ ordered by (`SerializedID`, `Route_Order`) **do**
0: $v \leftarrow$ `SID:` x.`SerializedID`; add_node($G, v,$ `type` = `victim`)
0: $l \leftarrow$ `LOC:` x.`Location`; add_node($G, l,$ `type` = `location`); INCR(v, l)
0: **for all** p in x.`Perpetrators` **do**
0: $a \leftarrow$ `PERP:` p; add_node($G, a,$ `type` = `perpetrator`); INCR(v, a)
0: **end for**
0: **for all** c in x.`Chiefs` **do**
0: $h \leftarrow$ `CHIEF:` c; add_node($G, h,$ `type` = `chief`); INCR(v, h)
0: **end for**
0: **end for**
0: annotate node attributes: `type`, $\deg_w$, distinct victims per location
0: **return** G =0

Complexity. Let $N = |\widehat{D}|$ and T the total tokens across perpetrator/chief lists. Construction is $O(N + T)$ time and $O(|V| + |E|)$ space. Connected components (for summaries) are computed in $O(|V| + |E|)$ using BFS/DFS (Cormen et al., 2009, Ch. 22).

3.8.3 Layout and export (reproducible).

Static figures use a seeded force-directed layout. Fruchterman–Reingold models repulsion and attraction as

$$f_r(d) = \frac{k^2}{d}, \qquad f_a(d) = \frac{d^2}{k}, \qquad k = C \cdot \sqrt{\frac{A}{|V|}},$$

where d is inter-node distance, A the drawing area, and C a tunable constant (Fruchterman and Reingold, 1991). Using a fixed random seed guarantees the same final layout for the same G. The HTML export (PyVis) embeds identical node/edge attributes and layout parameters so the interactive view matches the static figure.

Algorithm 3 Reproducible layout and export

Require: Graph G, seed s, canvas area A
0: set_seed(s); $k \leftarrow C\sqrt{A/|V|}$
0: run Fruchterman–Reingold for I iterations (deterministic under s)
0: encode node size $\propto \log(1 + \deg_w)$; color by `type`; edge width $\propto w$
0: write PNG with metadata {`dataset_id`, `schema_version`, s, k, I}
0: write HTML (PyVis) with the same metadata and styling =0

3.8.4 Adjacency views

For reporting, the graph can be summarized by its bipartite adjacency blocks:

$$A^{\mathrm{SL}}_{i,j} = w(s_i, \ell_j), \quad A^{\mathrm{SP}}_{i,j} = w(s_i, p_j), \quad A^{\mathrm{SC}}_{i,j} = w(s_i, c_j),$$

which support simple counts such as location victim-load $n(\ell_j) = \sum_i \mathbf{1}[A^{\mathrm{SL}}_{i,j} > 0]$ and token frequency $n(p_j) = \sum_i \mathbf{1}[A^{\mathrm{SP}}_{i,j} > 0]$. These summaries appear in the UI and in tables without exposing raw text.

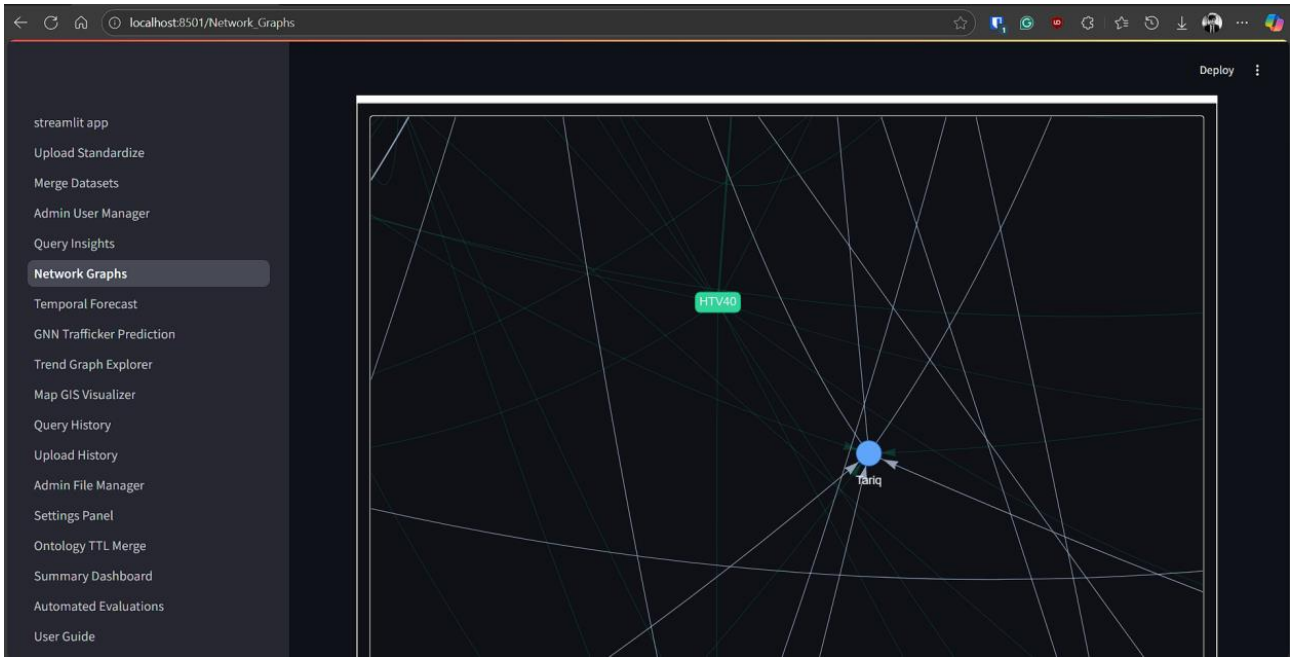

Figure 11. *Shows a victim 'HTV40' and a node (location) 'Tariq'*

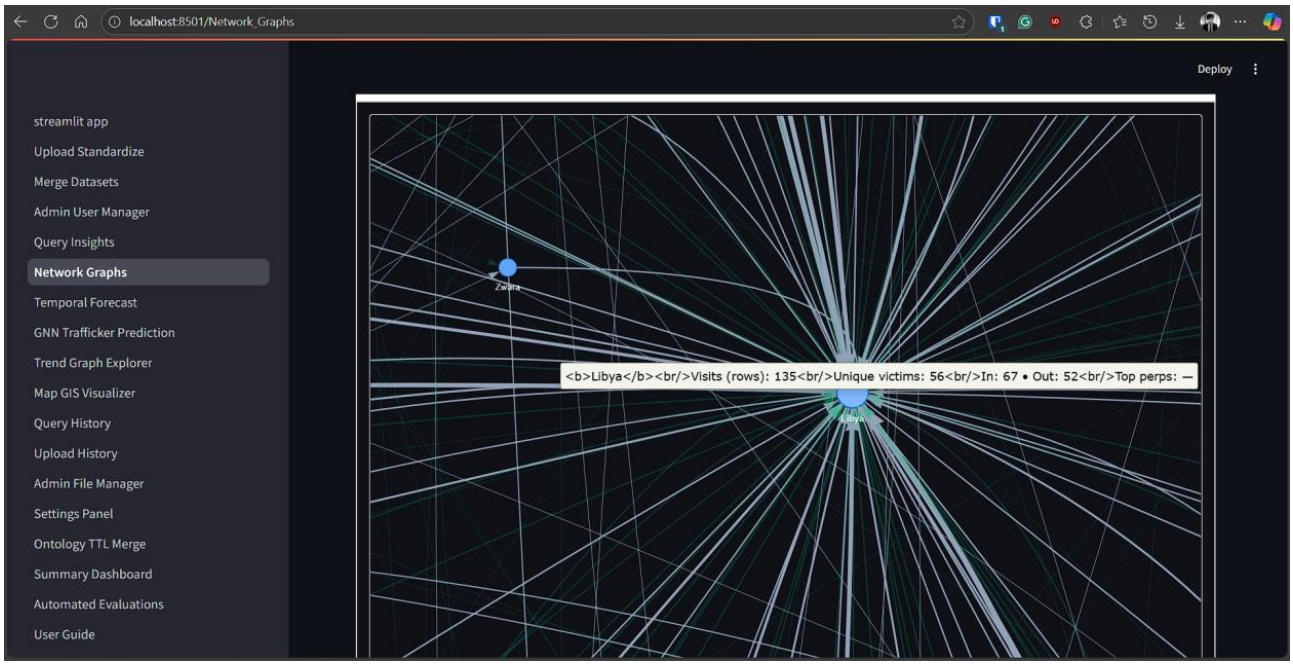

Figure 12. *Shows a node 'Libya' that shows stats like unique victims and incoming victims and outgoing victim count*

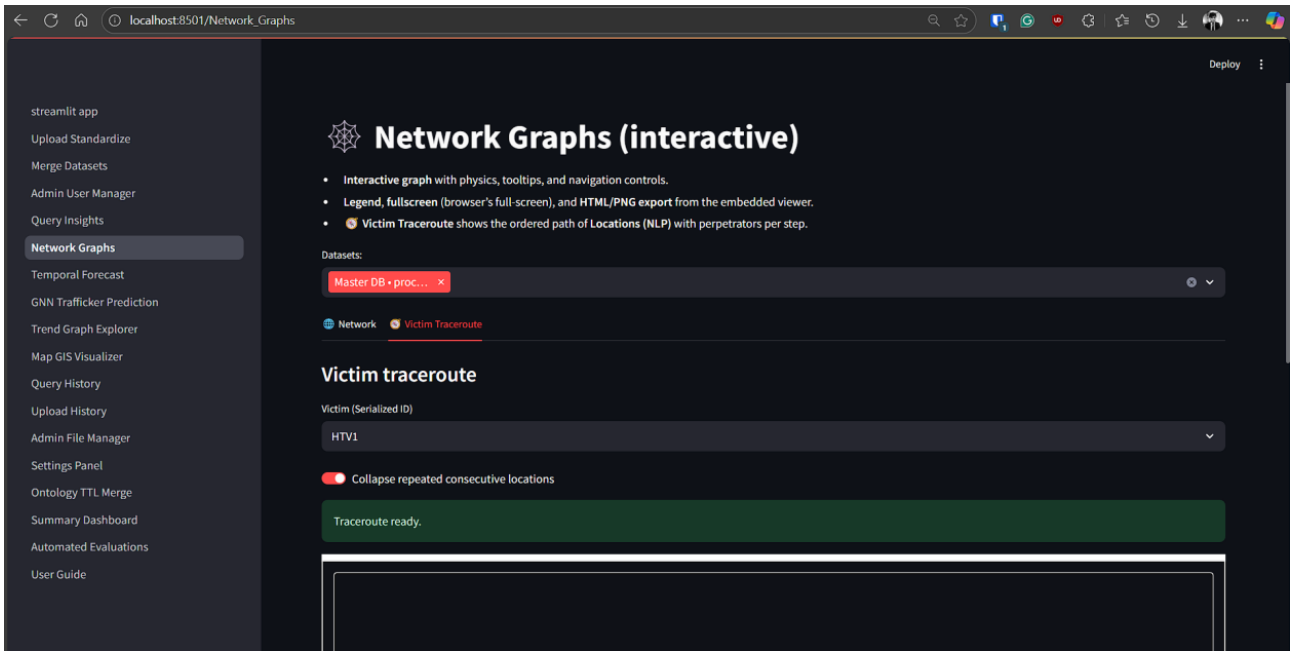

Figure 13. *There is also a subsection in the network graph page, called as 'Victim Traceroute' which generates a trajectory graph of the victims's movements*

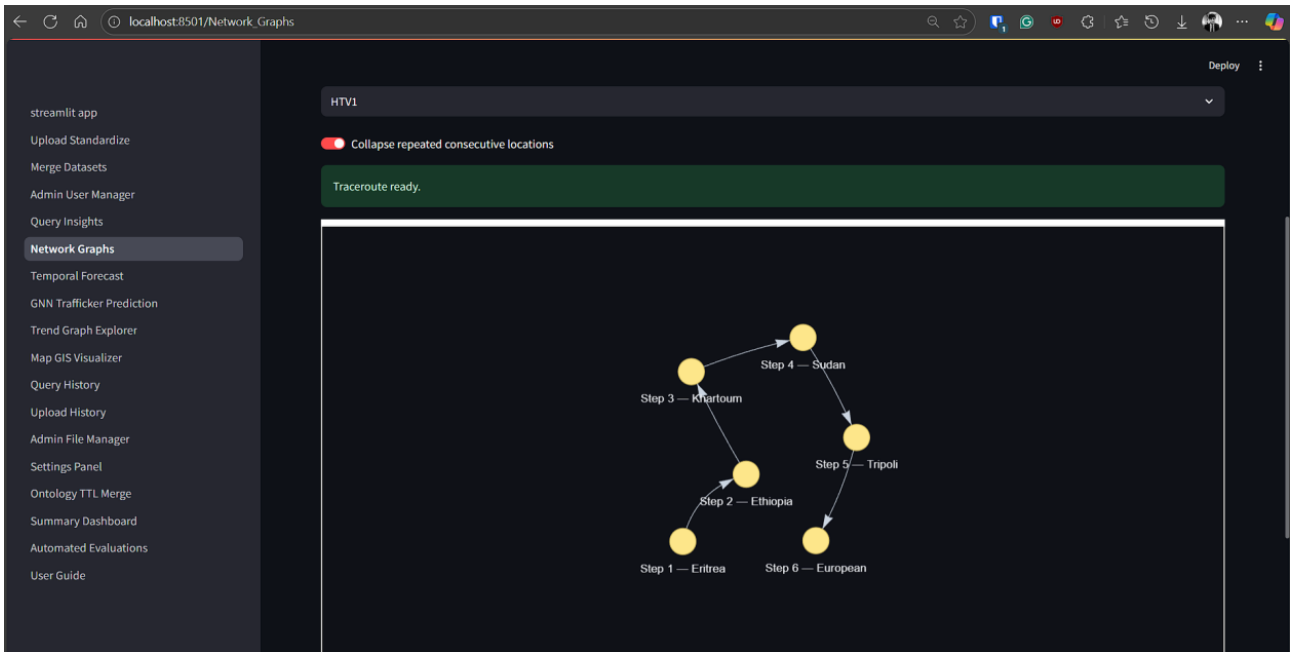

Figure 14. *Traceroute of victim HTV1*

Figure 15. *Traceroute of victim HTV11*

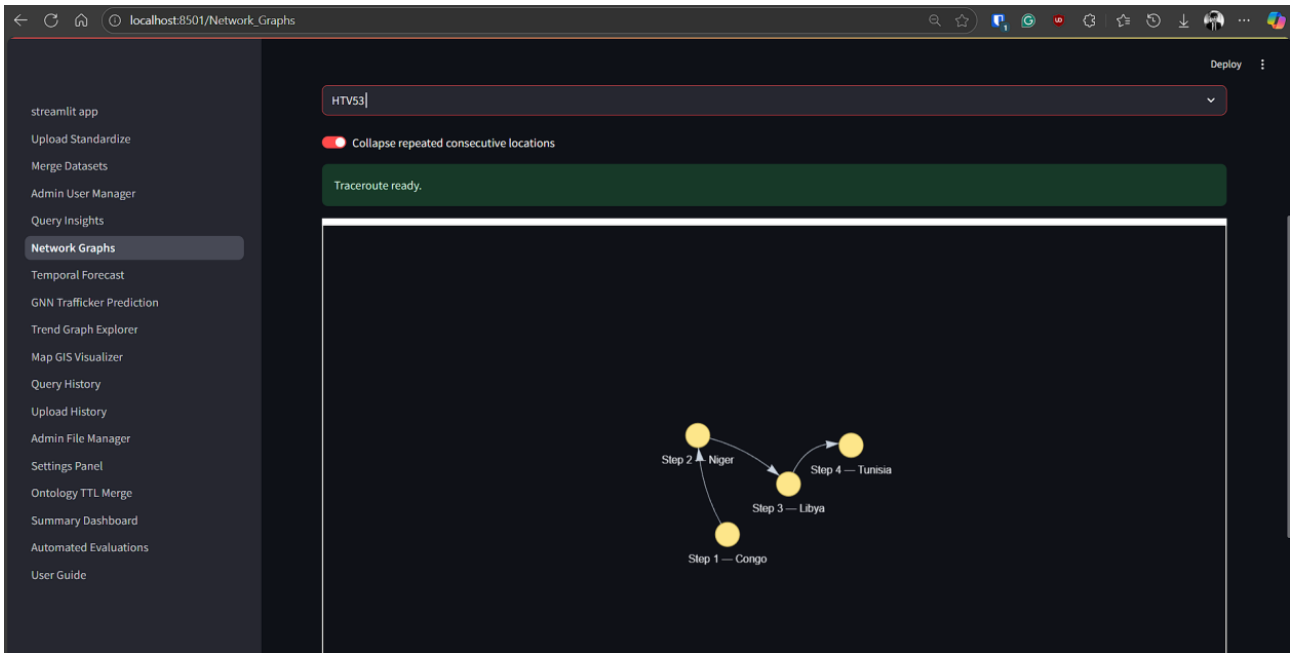

Figure 16. *Traceroute of victim HTV53*

3.8.5 Interpretation and safeguards.

Edges show co-occurrence in the processed projection. They do not imply causality. Very small counts are combined or left unlabeled in static exports to lower the risk of re-identification. Analysts should examine structural patterns with the GIS layers (Sec. 3.11) to understand hubs and routes better.

3.9 Ontology Creation & Semantic Modeling

This subsection defines the lightweight domain vocabulary and RDF serialization that the application uses to make its artifacts reusable and inspectable beyond the user interface. The ontology reflects the processed schema (Sec. 3.8) with a small set of classes and properties that keep the meaning of route steps, places, and actor tokens without showing raw text. Exports are created with `RDFlib` and written as Turtle (`.ttl`) so they are easy to compare, version, and check with common tools (Cyganiak et al., 2014; RDFLib Developers, 2025). Provenance and geospatial metadata follow established vocabularies: PROV-O for lineage and WGS84 for coordinates. This allows consumers to understand how each file was generated and, where possible, how locations relate to the earth's surface (Lebo et al., 2013; Brickley, 2003).

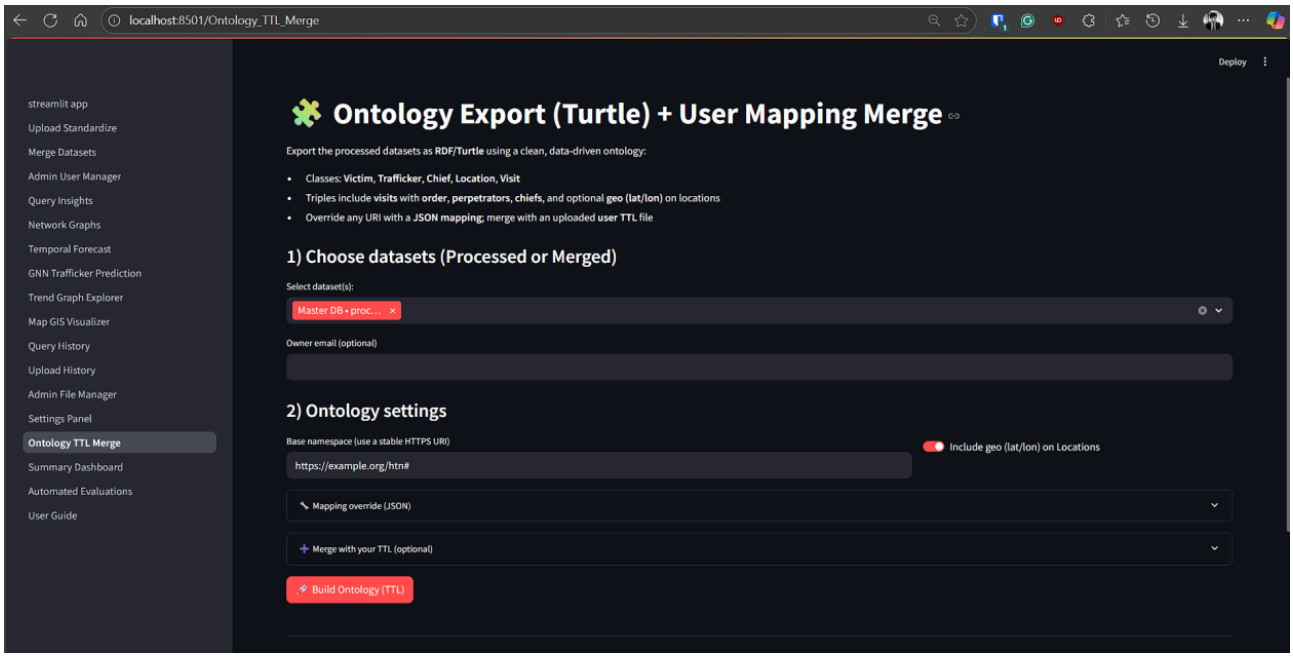

Figure 17. *This page creates the ontology export post NLP stage when we upload the data, it also support chronological merging of datasets, A user can only merge their datasets and merge them and create ontology, while the application never stores the RAW dataset it only stores the created ontology*

3.9.1 Design goals

The semantic model is intentionally minimal. It (i) matches 1:1 with the processed, sanitized projection; (ii) uses dereferenceable IRIs created from IDs and canonical tokens; (iii) records route order explicitly; and (iv) carries small, non-sensitive metadata (dataset ID, schema version, export parameters) to support reproducibility and audit, in line with FAIR principles (Wilkinson et al., 2016a). The application does not perform inference or ontology-driven reasoning; the RDF export is a direct serialization of the same information that the graph and GIS modules use.

3.9.2 Namespaces and core vocabulary.

The export declares standard prefixes plus an application namespace `ex:`. The core classes are:

- `ex:Victim` for case instances keyed by `SerializedID`.
- `ex:Location` for canonical place tokens (optionally augmented with coordinates).
- `ex:Perpetrator` and `ex:Chief` for canonicalized actor tokens.
- `ex:Visit` for a route step (a victim's presence at a location with an explicit order).

The principal object/data properties include:

- `ex:hasVisit` (`Victim` → `Visit`), `ex:atLocation` (`Visit` → `Location`), `ex:order` (`xsd:integer`) for route sequencing.

- `ex:hasPerpetrator` (`Victim` → `Perpetrator`) and `ex:hasChief` (`Victim` → `Chief`) for canonicalized tokens attached to a case.
- `ex:schemaVersion` (`xsd:string`) and `ex:sourceID` (`xsd:string`) for minimal, non-sensitive metadata.
- `wgs84:lat`, `wgs84:long` (optional) on `Location`, when coordinates are available from the GIS step (Brickley, 2003).

3.9.3 Deterministic IRI construction and mapping.

Let a processed row be $x = (\texttt{SerializedID}, \texttt{Location}, \texttt{Route\_Order}, \texttt{Perpetrators}, \texttt{Chiefs})$. IRIs are formed by stable, URL-safe encodings of canonical strings:

$$\begin{aligned} \text{Victim } s: &\ \texttt{ex:v/}\text{url\_safe}(\texttt{SerializedID}) \\ \text{Location } \ell: &\ \texttt{ex:loc/}\text{url\_safe}(\texttt{Location}) \\ \text{Visit } v_{s,j}: &\ \texttt{ex:visit/}\text{url\_safe}(\texttt{SerializedID})\texttt{/}j \\ \text{Perpetrator } p: &\ \texttt{ex:perp/}\text{url\_safe}(p) \\ \text{Chief } c: &\ \texttt{ex:chief/}\text{url\_safe}(c). \end{aligned}$$

The row-to-triples mapping μ creates the following assertions:

$$\begin{aligned} & s\ \texttt{rdf:type ex:Victim}, \quad \ell\ \texttt{rdf:type ex:Location}, \quad v_{s,j}\ \texttt{rdf:type ex:Visit}, \\ & s\ \texttt{ex:hasVisit}\ v_{s,j}, \quad v_{s,j}\ \texttt{ex:atLocation}\ \ell, \quad v_{s,j}\ \texttt{ex:order}\ j, \\ & s\ \texttt{ex:hasPerpetrator}\ p\ \ \forall p \in \texttt{Perpetrators}, \qquad s\ \texttt{ex:hasChief}\ c\ \ \forall c \in \texttt{Chiefs}. \end{aligned}$$

All literals use explicit datatypes (e.g., `xsd:integer` for `ex:order`). Because canonicalization occurs upstream (Sec. 3.6), tokens are already lowercased, trimmed, and de-duplicated, ensuring consistent identifiers in RDF.

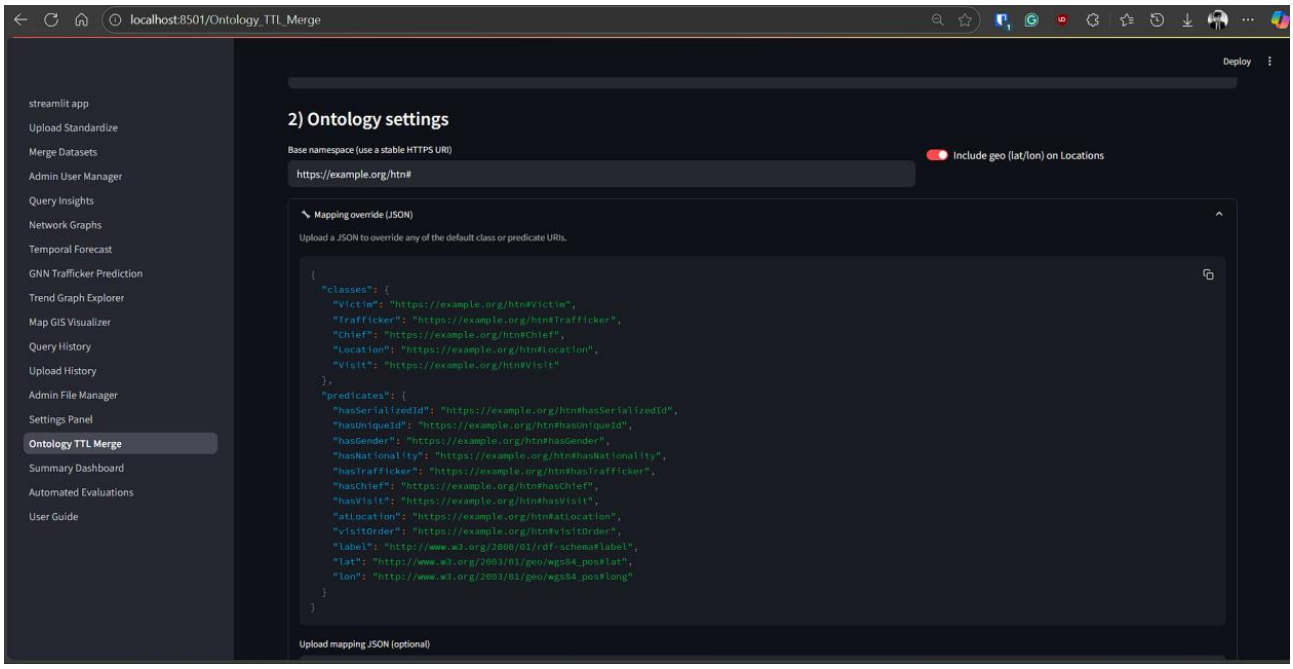

Figure 18. *The Ontology page also allows for Ontology setting, changing the base namespace and it also supports adding the geocodes (lat/lon) on locations*

3.9.4 Minimal illustrative Turtle snippet.

The following excerpt (with anonymized IDs) shows a single route step and associated tokens. It mirrors exactly the information in the processed table; there is no inclusion of raw free text.

```
@prefix ex:    <https://example.org/ont/> .
@prefix xsd:   <http://www.w3.org/2001/XMLSchema#> .
@prefix prov:  <http://www.w3.org/ns/prov#> .
@prefix wgs84: <http://www.w3.org/2003/01/geo/wgs84_pos#> .

ex:v/7f9a... a ex:Victim ;
  ex:schemaVersion "1.3" ;
  ex:sourceID "dataset_A" ;
  ex:hasVisit ex:visit/7f9a.../1 ;
  ex:hasPerpetrator ex:perp/network_op ;
  ex:hasChief ex:chief/handler_k .

ex:loc/tripoli a ex:Location ;
  wgs84:lat "32.8872"^^xsd:decimal ;
  wgs84:long "13.1913"^^xsd:decimal .

ex:visit/7f9a.../1 a ex:Visit ;
  ex:atLocation ex:loc/tripoli ;
  ex:order "1"^^xsd:integer .
```

3.9.5 Export procedure (RDFlib).

The serialization is created in a single pass over the processed table:

1. Initialize an empty `rdflib.Graph` with prefixes `rdf`, `rdfs`, `xsd`, `prov`, `wgs84`, and `ex`.
2. For each row x (ordered by `SerializedID` and `Route_Order`): mint IRIs for $s, \ell, v_{s,j}$; assert class types and `ex:hasVisit`, `ex:atLocation`, `ex:order`. For each perpetrator token p and chief token c, assert `ex:hasPerpetrator` and `ex:hasChief` links.
3. When coordinates are available from the GIS pipeline, annotate `Location` with `wgs84:lat` and `wgs84:long`.
4. Attach minimal metadata on the dataset resource (schema version, source ID). Optionally, include a small PROV header (dataset entity, export activity, software agent) for stronger lineage.
5. Serialize to Turtle; record the processed dataset ID and parameters in the export's side metadata (Sec. 3.5).

Because the export is a pure function of the processed table and parameters, re-running with the same inputs yields an identical `.ttl` (byte-for-byte when iteration order and numeric formatting are fixed), which supports figure regeneration and external validation.

3.9.6 Interoperability and downstream use.

Keeping the vocabulary close to the processed schema reduces cognitive load and maintains traceability between tables, graphs, and RDF. The presence of WGS84 properties allows GIS tools to directly use location nodes. The explicit `ex:order` on visits enables time-aware visualizations or queries to reconstruct routes without checking the original spreadsheet. External users can import the file into triple stores, run SPARQL for cross-case consolidation or hub identification, and, if desired, add their own shapes or integrity checks (for example, *Visit* must point to exactly one *Location*). While the current implementation does not include a SHACL file, the export works with SHACL validation patterns if partners want to impose additional constraints (Knublauch and Kontokostas, 2017).

3.9.7 Privacy and scope.

Only sanitized tokens and standard place names appear in RDF. There are no raw narratives, direct identifiers, or sensitive free text. Coordinates, when present, follow the same caution used in maps: unresolved or low-confidence locations are left out instead of being estimated. The ontology is intentionally compact and descriptive. It captures what the system already produces and uses while avoiding speculative classes or properties that are not supported by implementation.

3.10 GIS Spatio-Temporal Visualization

This section explains how the application converts standardized place strings from the processed table (Sec. 3.8) into geospatial layers for exploration. The design focuses on clarity

and reproducibility. Place resolution is fixed under a documented threshold. Map layers come only from the cleaned projection. Exports follow standards, which allows them to be reused outside the app. All geodata are written as GeoJSON FeatureCollections according to RFC 7946 (Butler et al., 2016) and displayed with Folium, which uses Leaflet for interactive viewing (Folium Project, 2025; Leaflet, 2025).

3.10.1 Inputs and outputs.

The input is the processed table containing the columns `Location` and `Route_Order` (along with the optional `Time_Spent`). The outputs are: (i) a point FeatureCollection for markers and clustering, (ii) a weighted point set for density heatmaps, and (iii) a sequence for each case for temporal playback when route order or timing is available. Each artifact includes metadata that links it to the processed dataset ID and schema version (Sec. 3.5).

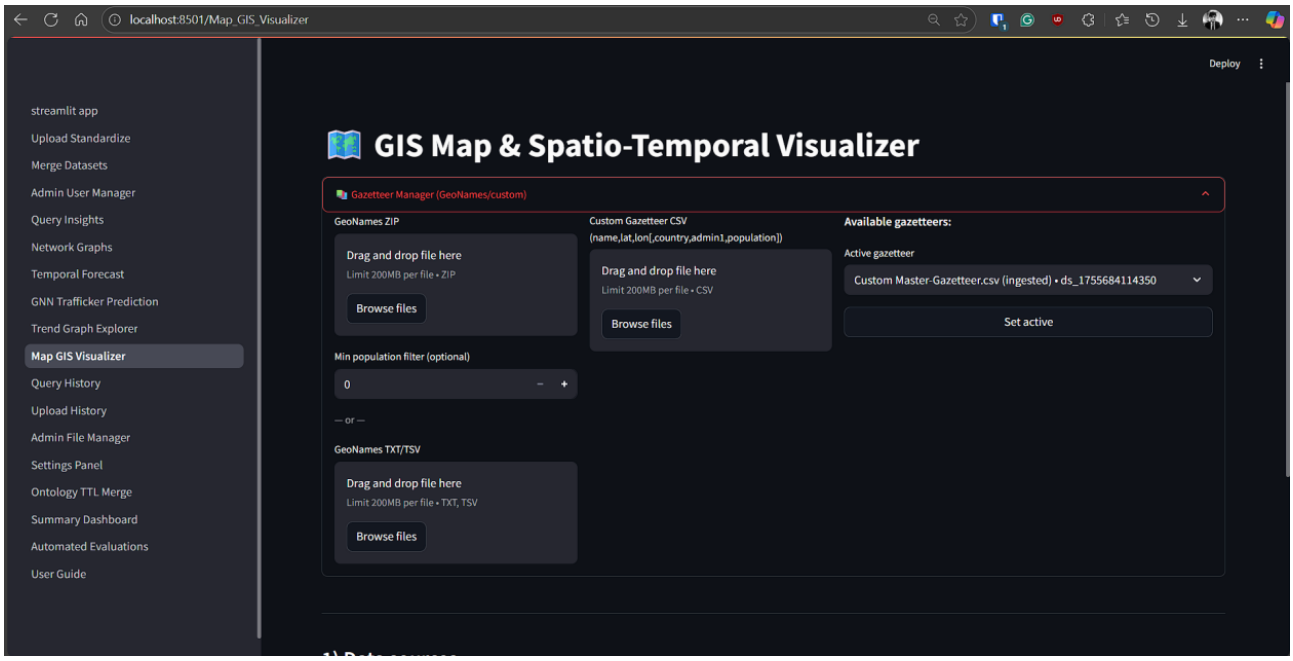

Figure 19. *GIS Map & Spatio-Temporal Visualizer Page*

3.10.2 Toponym resolution (lookup → conservative fuzzy match).

Place strings are first normalized (lowercased, trimmed, punctuation-normalized) and then resolved to coordinates by a two-stage procedure. Stage 1 queries a curated lookup table that maps canonical tokens to (lat, lon) and optional country/region. Stage 2 applies a conservative fuzzy match when no exact entry exists, using edit distance on canonical forms. A candidate c is accepted if

$$\mathrm{sim}(q, c) = 1 - \frac{\mathrm{lev}(q, c)}{\max(|q|, |c|)} \geq \theta,$$

where q is the query token, lev is Levenshtein distance, and θ is a high threshold (e.g., 0.9) to guard against false positives (Levenshtein, 1966; Wagner and Fischer, 1974). Resolutions are tagged as *exact*, *fuzzy*, or *unresolved*. Unresolved entries are excluded from map layers but reported in diagnostics so users can improve the lookup table.

Algorithm 4 Toponym resolution and status tagging

Require: Canonical place token q, lookup dictionary L, threshold θ
Ensure: $(\text{lat}, \text{lon}, \text{status}, \text{score})$
0: **if** $q \in L$ **then**
0: **return** $(L[q].\text{lat}, L[q].\text{lon}, \texttt{exact}, 1.0)$
0: **end if**
0: $c^\star \leftarrow \arg\max_{c \in \text{keys}(L)} \text{sim}(q, c)$
0: **if** $\text{sim}(q, c^\star) \geq \theta$ **then**
0: **return** $(L[c^\star].\text{lat}, L[c^\star].\text{lon}, \texttt{fuzzy}, \text{sim}(q, c^\star))$
0: **else**
0: **return** $(\varnothing, \varnothing, \texttt{unresolved}, 0)$
0: **end if**=0

3.10.3 Point and cluster layers.

Resolved locations become point features with basic properties: canonical name, occurrence count in the processed table, number of distinct `SerializedID`s, and resolution status. The marker cluster layer reduces overlap in crowded areas by grouping nearby markers at low zoom levels and expanding them when zoomed in (Leaflet.markercluster) (Leaflet, 2025). Tooltips show only cleaned fields (counts, tokenized names).

3.10.4 Density heatmap.

The heatmap uses the same points with scalar weights. By default, the weight at location ℓ reflects the number of processed rows mapped to ℓ (i.e., visits). When set up, the UI can use distinct case counts per location to reduce the impact of repeated actions by the same case. The layer is created with Folium's HeatMap plugin (Leaflet.heat), which blends nearby points visually in the browser (Folium Project, 2025). Parameters (radius, blur) remain unchanged across exports for consistency.

3.10.5 Spatio-temporal trajectories

When route order is available, sequences for each case are formed by sorting rows by `Route_Order` and linking successive resolved locations. The export employs `TimestampedGeoJson` so users can watch movement sequences over a time slider in the browser. If absolute timestamps are missing, the sequence index offers a rough, order-only animation; if timestamps or durations are available, they are shown as ISO-8601 times on each feature, allowing for time-aware playback (Butler et al., 2016). Low-confidence points are skipped instead of interpolated to avoid suggesting precision.

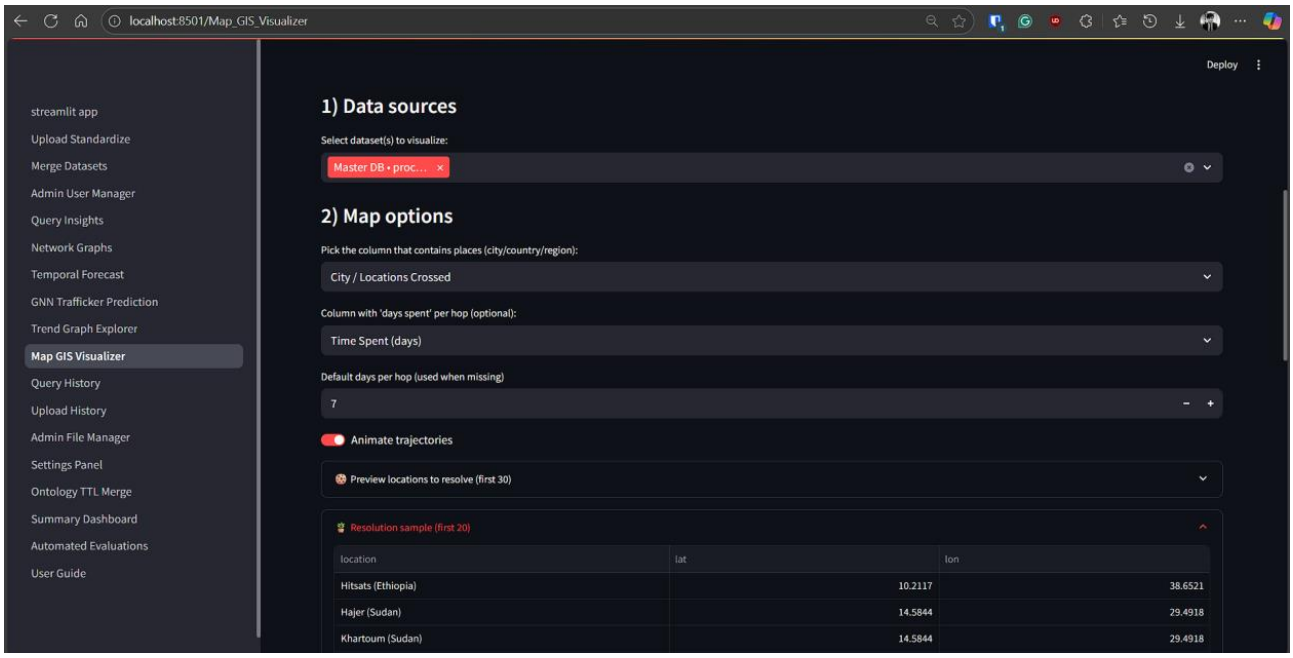

Figure 20. *This image of GIS Visualizer page shows how user can select a dataset and further select columns that are required to generate the visulization*

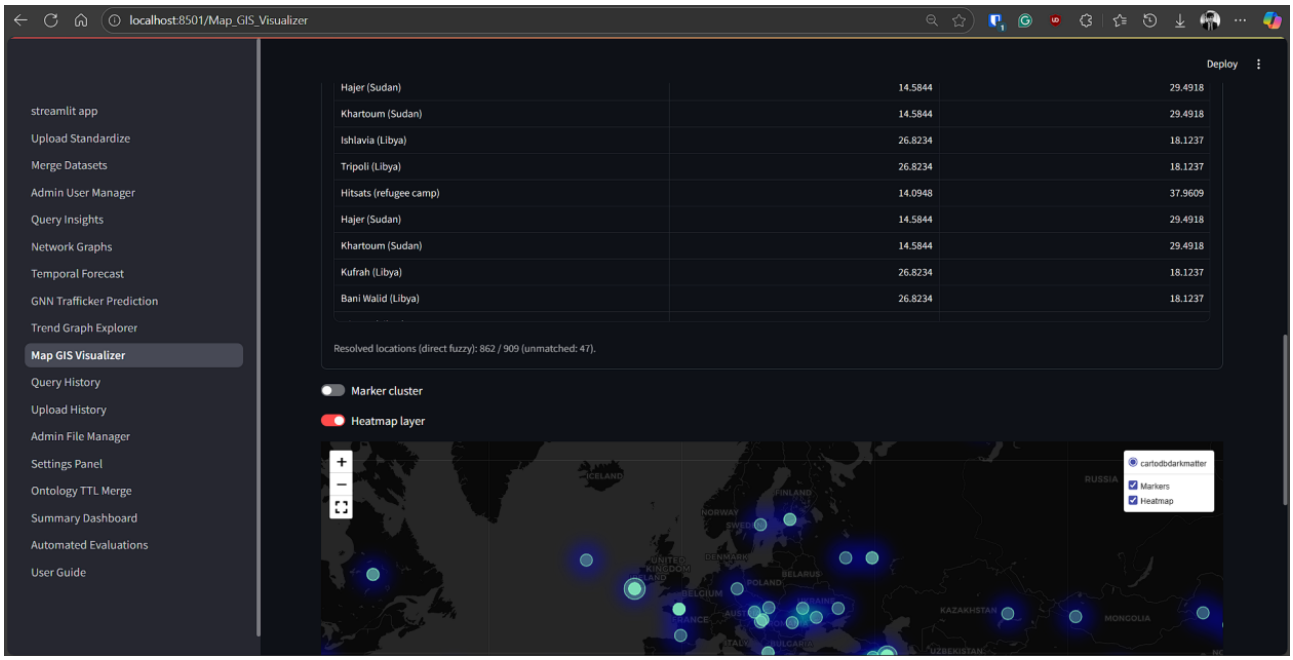

Figure 21. *The 'Locations' column is matched with the gazetteer to get the coordinates of these locations for further GIS Map construction*

Algorithm 5 Trajectory assembly from processed routes

Require: Processed table $\widehat{D}$ with (`SerializedID`, `Location`, `Route_Order`), resolution map ρ
Ensure: FeatureCollection F for `TimestampedGeoJson`

```
for all case s grouped by SerializedID do
    P ← []
    for all row x in s ordered by Route_Order do
        (φ, λ, σ) ← ρ(x.Location) {lat, lon, status}
        if σ ∈ {exact, fuzzy} then
            append (φ, λ) to P
        end if
    end for
    add polyline feature for P with ordered timestamps or indices
end for
return F =0
```

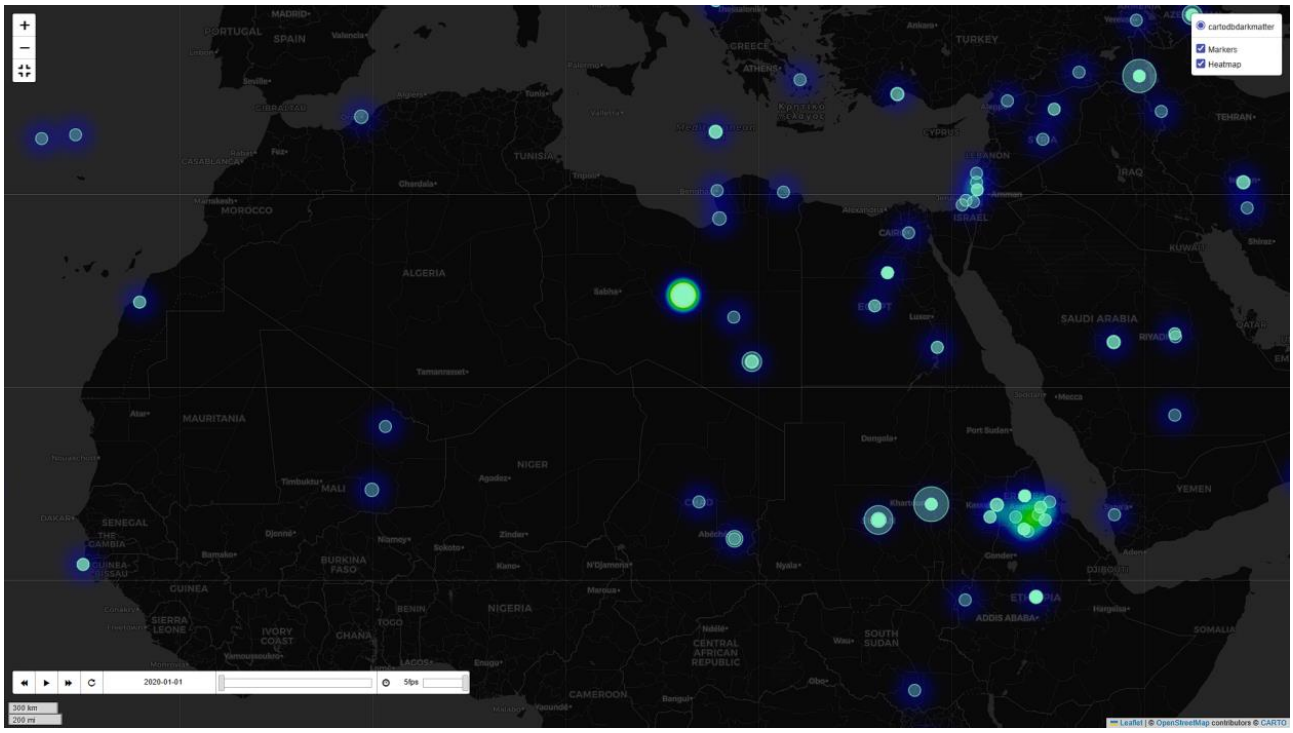

Figure 22. *Final GIS Map with Spatio-Temporal Visualization and slider to control it, with nodes being represented by amount of traffic going through them and heatmap embedded on top of them*

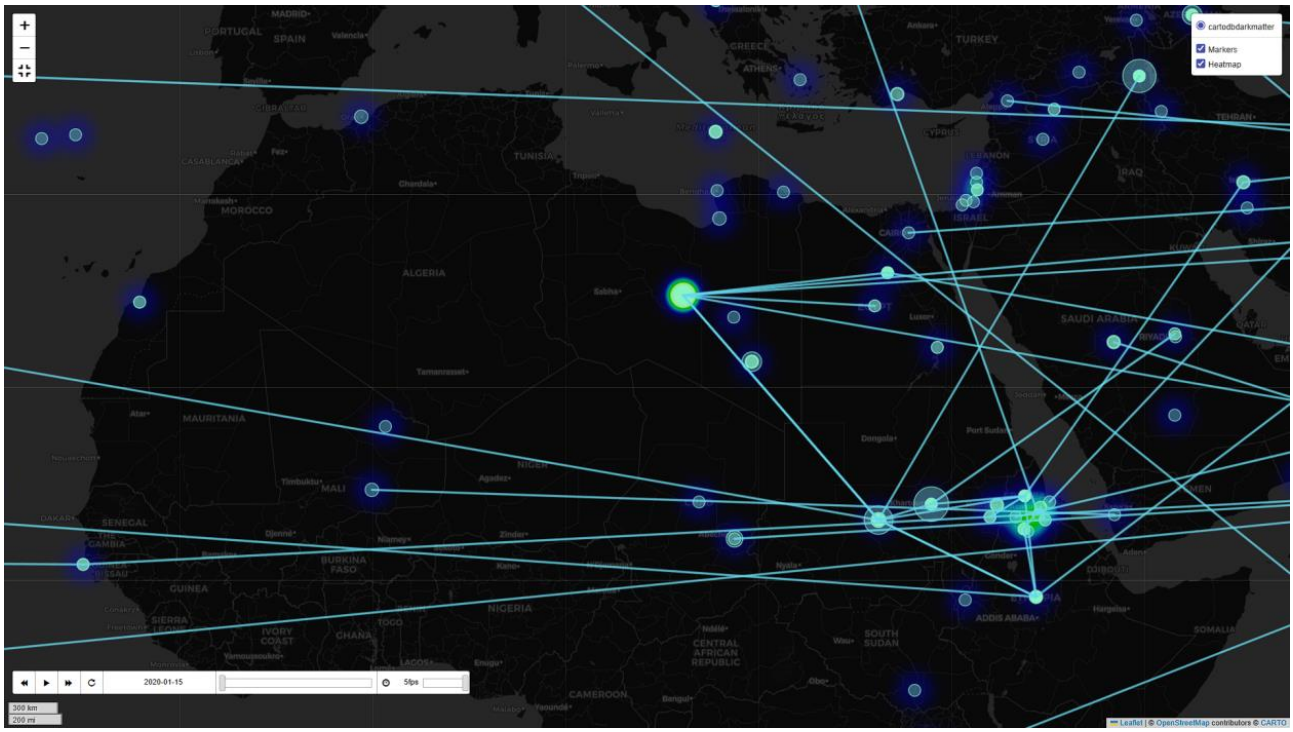

Figure 23. *Final GIS Map with Spatio-Temporal Visualization and with animated trajectories on top of it to visualize trafficking routes*

3.10.6 Uncertainty handling and safeguards.

To prevent overstating accuracy, the map leaves out unresolved points and highlights unclear resolutions in popups. The UI avoids showing labels for very small counts in static exports. It does not present raw narrative text or direct identifiers. Coordinates come from a controlled lookup. There is no real-time geocoding using external services. This keeps data residency local and reduces the risk of data leakage.

3.10.7 Exports and reproducibility.

The application saves the following items in the owner's workspace: the point FeatureCollection, the heatmap-ready point list with weights, and the timestamped sequence FeatureCollection. Each file comes with minimal metadata, including processed dataset ID, schema version, resolution threshold, and layer parameters. Since all layers originate from the processed projection and a fixed lookup with a documented θ, rerunning with the same inputs creates identical artifacts and map views (Butler et al., 2016; Folium Project, 2025).

3.10.8 Interpretation notes.

Spatial intensity and trajectories show what is in the processed table; they do not reflect prevalence in the population. Gaps or breaks in lines are expected when locations are unresolved or filtered for low confidence. Analysts should view the layers along with the network view (Sec. 3.9) to connect spatial hubs to structural patterns among cases and actor tokens.

3.11 Quality Checks and Data Diagnostics

This section describes the lightweight diagnostics the application runs on the *processed* table before the graph and GIS stages. The aim is to identify inconsistencies early, measure coverage, and avoid unnoticed changes across uploads and runs, while keeping raw content hidden. The checks are quick, predictable, and shown in the interface as small tables and pass/fail badges. They support the reproducibility and provenance practices discussed earlier (Sandve et al., 2013; Wilson et al., 2017) and work with the tidy, long-format schema (Sec. 3.8) (Wickham, 2014).

3.11.1 Coverage and completeness.

Basic summaries are calculated for each dataset and by key columns:

- total processed rows, distinct `SerializedID`s, average and median route length;
- fraction of rows with non-empty `Location`, count and rate of unresolved locations (from GIS resolution);
- presence of optional fields (`Time_Spent`, `Perpetrators`, `Chiefs`) and their non-empty rates;
- distinct counts per token list, and top-k token frequencies for sanity checking.

These indicators assist analysts in deciding whether to move forward to graph/GIS or adjust the intake mapping or lookup dictionaries.

3.11.2 Key consistency and shape invariants.

Deterministic validators uphold the invariants given in Sec. 3.8: uniqueness of `(SerializedID, Route_Order)`, ordered `Route_Order` within cases, non-empty canonicalized `Location`, and JSON-serializable list fields. Violations are reported with brief, redactable row references (dataset ID and row index) instead of showing values.

3.11.3 Unresolved locations dashboard.

The GIS module labels each location token as *exact*, *fuzzy*, or *unresolved*. Diagnostics compile these by token and by source dataset, listing the most common unresolved items so curators can expand the lookup table. For fuzzy matches, the distribution of similarity scores is shown to help assess whether the threshold θ is too strict or lenient (Sec. 3.11). No third-party geocoding is used; the dashboard only shows the controlled lookup and local matching logic.

3.11.4 Duplicate-case heuristic (conservative).

To guard against inadvertent double counting, a conservative heuristic flags potential duplicate cases across uploads without exposing raw text. Two cases a and b are flagged if they

share the same normalized nationality and have highly similar route footprints and token lists. Similarity for token lists (perpetrators/chiefs) is measured by the Jaccard index

$$J(A, B) = \frac{|A \cap B|}{|A \cup B|},$$

and route similarity uses the Jaccard index over the set of visited locations (order ignored for robustness). A pair is flagged when

$$\texttt{Nationality}(a) = \texttt{Nationality}(b), \qquad J(\texttt{Locs}_a, \texttt{Locs}_b) \geq \tau_L, \qquad \max\left(J(\texttt{Perp}_a, \texttt{Perp}_b), J(\texttt{Chief}_a, \texttt{Chie}\right.$$

with conservative thresholds (e.g., $\tau_L = 0.8$, $\tau_A = 0.6$). Flagging does *not* merge records; it produces a review list for the data owner. Jaccard is a standard set-similarity measure suitable for normalized tokens (Jaccard, 1901).

Algorithm 6 Conservative duplicate-case flagging (owner-only view)

Require: Processed table grouped by `SerializedID`; thresholds τ_L, τ_A
Ensure: List of flagged pairs (a, b) with scores
build per-case sets: $\texttt{Locs}_s$, $\texttt{Perp}_s$, $\texttt{Chief}_s$, and value $\texttt{Nationality}_s$
for all unordered pairs (a, b) with $\texttt{Nationality}_a = \texttt{Nationality}_b$ **do**
  $j_L \leftarrow J(\texttt{Locs}_a, \texttt{Locs}_b)$; $j_P \leftarrow J(\texttt{Perp}_a, \texttt{Perp}_b)$; $j_C \leftarrow J(\texttt{Chief}_a, \texttt{Chief}_b)$
  if $j_L \geq \tau_L$ **and** $\max(j_P, j_C) \geq \tau_A$ **then**
    emit (a, b, j_L, j_P, j_C)
  end if
end for=0

3.11.5 SID-UID coherence.

A bijection check ensures that each cleaned `UniqueID` corresponds to exactly one `SerializedID` and vice versa within an intake batch. This process helps prevent accidental key drift. If a single `UniqueID` links to multiple serialized forms, typically due to whitespace or casing issues in the intake file, the validator recommends a normalization patch.

3.11.6 Distribution snapshots (for reporting).

The UI generates small, exportable tables and charts that summarize:

- victims by location (distinct `SerializedID`s per `Location`);
- route-length distribution (histogram of `Route_Order` maxima);
- token frequencies for `Perpetrators` and `Chiefs` (top-k with counts).

These snapshots will appear later in the manuscript as supporting evidence for network and GIS figures. They can be recreated from the processed dataset ID and schema version.

3.11.7 Operational notes.

Diagnostics are calculated on the cleaned projection only and do not log any raw text. All thresholds and parameter choices, such as θ for fuzzy toponym matching and τ_L, τ_A for duplicate flags, are documented in the run metadata (Sec. 3.5). This ensures reproducibility of results and allows for reviews of the choices made (Sandve et al., 2013).

3.11.8 Limitations.

The duplicate heuristic focuses on precision rather than recall. It may overlook duplicates when routes differ slightly or when token lists are sparse. Unresolved locations can skew the perceived coverage in GIS summaries. These are acceptable compromises for a privacy-preserving, audit-friendly workflow. Future enhancements can introduce owner-controlled adjustments without changing the fundamental approach.

3.12 Front-End Orchestration and Reproducible Artifacts

The user interface directs the pipeline in a clear, straightforward order. It also creates analysis-ready artifacts that can be regenerated without needing access to the raw inputs. We implement orchestration using Streamlit, which offers a simple and predictable control flow. Parameter widgets trigger a sequence of backend calls, and export buttons create files in the owner's workspace (Streamlit Inc., 2025). We ensure reproducibility by capturing parameters, seeds, and identifiers with every export. This follows well-established guidelines for computational research (Sandve et al., 2013; Wilson et al., 2017) and links machine-readable provenance to the outputs (Lebo et al., 2013).

3.12.1 Execution order and page flow.

The interface shows pages that correspond to the major modules. During a session, the call order is as follows:

1. **Intake**: read the uploaded file into memory and check the headers (no raw data persists).
2. **Preprocessing**: normalize and canonicalize to create the processed long-format table (Sec. 3.6) and calculate a run fingerprint.
3. **Graph**: build the typed, weighted graph from the processed table (Sec. 3.9) and show a preview.
4. **Ontology/RDF**: convert the same information into TTL format (Sec. 3.10) for future use.
5. **GIS**: resolve place names and create GeoJSON layers. Show previews of base, heatmap, cluster, and trajectory maps (Sec. 3.11).
6. **Diagnostics**: present coverage and integrity checks (Sec. 3.13).

Users can navigate freely, but export buttons only become active once the processed table is available. This ensures all artifacts come from the same cleaned source.

3.12.2 Deterministic orchestration (high level)

Each step acts as a pure function of its inputs and parameters. By using a fixed seed for layout and set thresholds for resolution, we produce byte-identical exports on repeat runs. This enables exact figure regeneration (Sandve et al., 2013).

3.12.3 Parameter capture and metadata stamping.

To achieve reproducibility, we stamp metadata into both the file header (when supported) and the registry entry (Sec. 3.5):

- processed dataset identifier and `Schema_Version`;
- run fingerprint and creation timestamp;
- graph layout seed and iteration count; map parameters (radius, blur) and resolution threshold θ;
- tool versions (`pandas`, NetworkX, RDFlib, Folium) as reported by the environment.

These fields enable reviewers to recreate the environment and verify that figures and maps align with the processed data (Wilson et al., 2017; Lebo et al., 2013).

3.12.4 Exports and file conventions.

Outputs are written under the owner's workspace using deterministic names that encode the processed dataset ID and a short artifact tag, for example:

- `proc-{ID}-table.parquet` (processed table),
- `graph-{ID}-seed{s}.png` and `graph-{ID}-seed{s}.html`,
- `kg-{ID}.ttl` (RDF export),
- `map-{ID}-points.geojson`, `map-{ID}-heat.json`, `map-{ID}-traj.geojson`.

Consistent naming makes it straightforward to cite artifacts in the thesis and to re-run exports by ID.

3.12.5 Access control at interaction points.

The UI ensures that only the uploader can trigger transforms on their upload during the active session. Exports are written to the uploader's area. Administrative pages show configuration and health status, but not data content. Artifact discovery in other pages works through the registry and filters by owner. This ensures that unshared artifacts do not appear in the list.

3.12.6 Error handling and idempotency.

Preconditions are checked at each step, such as required columns, non-empty locations, and valid parameter ranges. Export operations are atomic. Files are saved to a temporary path and moved into place only if the process succeeds. The registry is updated afterward to prevent dangling references. If a file is missing or fails an integrity check, the UI provides a "regenerate" option. This uses the stored parameters for that artifact. These mechanisms focus on simple, testable guarantees rather than complex workflow engines. This approach follows "good enough practices" for research software (Wilson et al., 2017).

3.12.7 Residency and non-egress behavior.

All exports consist of local files within a controlled environment. The interface does not sync with external services or perform online geocoding. Maps are generated from local GeoJSON using Leaflet through Folium (Folium Project, 2025; Leaflet, 2025). This keeps data movement clear and under the data owner's control.

3.12.8 Implications for the manuscript.

Since artifacts are parameter-stamped and linked to the registry, every figure and map in later chapters can reference a processed dataset ID and a small set of parameters. Recreating the same outputs means reloading the processed table and executing the export path again with those parameters. This meets the requirement that published computational results be independently verifiable (Sandve et al., 2013; Wilson et al., 2017; Lebo et al., 2013).

3.13 Limitations Bound to Implementation Choices

The methodology prioritizes privacy, clarity, and reproducibility. These priorities create specific trade-offs that limit what the system can demonstrate. This subsection outlines those limitations clearly so that future interpretation and evaluation remain grounded.

3.13.1 Dependence on sanitized inputs.

All analytics operate on a minimized, processed projection rather than raw narratives (Sec. 3.8). This protects confidentiality but also removes nuance that might help clarify ambiguity in places, actors, or context. The choice reflects data-minimization and purpose-limitation principles (EU-, 2016) and FAIR-aligned reuse of derived artifacts (Wilkinson et al., 2016a), but it may reduce recall compared to pipelines that inspect full text.

3.13.2 Deterministic canonicalization instead of learned extraction.

Preprocessing relies on specific rules for tokenization, trimming, case normalization, and cautious near-duplicate collapse; it does not use model-based NER or relation extraction. This improves auditability and repeatability (Sandve et al., 2013; Wilson et al., 2017) but may overlook semantically equivalent forms that simple edits do not capture. Cautious thresholds also risk merging distinct but similar tokens when names are short or unclear. These behaviors

are documented and parameterized so users can adjust them with an understanding of the privacy position.

3.13.3 Toponym resolution by lookup and cautious fuzzy match.

Place names are resolved through a controlled lookup with a high-threshold edit-distance fallback (Sec. 3.11). The system does not use external geocoding services to respect residency and to avoid unintended disclosures. As a result, some locations remain unresolved and are excluded; others may resolve to a main point rather than a precise site. GeoJSON exports follow RFC 7946 and follow its assumptions about coordinate reference and feature structure, appropriate for web mapping but do not capture administrative hierarchies or uncertainty envelopes (Butler et al., 2016). These choices prioritize safety over extensive coverage.

3.13.4 Route order and timing.

Route sequences reflect the order given by the dataset; when timing is not available, the playback is based only on order and not a real-time reconstruction. Durations, where present, are converted into general day counts for descriptive summaries only. The system does not predict arrival times or estimate any timing, in line with its defined contribution.

3.13.5 Graph structure and visual limits.

The network is typed and weighted by co-occurrence from the processed table; edges do not imply causality or legal responsibility. Directionality is implicit (victim to other node types) rather than fully modeled, and weights represent frequency rather than strength in a causal sense. Force-directed layouts (e.g., Fruchterman–Reingold) are easy to read for small to medium graphs but can become cluttered as the number of nodes increases, even when set for reproducibility (Fruchterman and Reingold, 1991; Hagberg et al., 2008). Interactive rendering works well for a few thousand nodes/edges on standard hardware; beyond that, static summaries are preferred.

3.13.6 Minimal ontology and validation.

The RDF export mirrors the processed schema with a compact vocabulary; it does not include general ontologies beyond what the system actually uses. While it can work with SHACL validation, the current version does not come with a shapes file by default. The benefit is a low cognitive load and faithful serialization; the drawback is limited semantic richness compared to larger ontologies (Lebo et al., 2013; Wilkinson et al., 2016a).

3.13.7 Diagnostics by cautious heuristics.

Quality checks (e.g., potential duplicate cases) use cautious set-similarity heuristics to avoid false positives that might reveal sensitive inferences. This improves precision at the expense of recall and may under-flag true duplicates when routes differ slightly or token lists are sparse. Thresholds are noted so owners can change them, but defaults are conservative.

3.13.8 Reproducibility versus environment drift.

Exports are parameter-stamped and registry-linked (Sec. 3.5), enabling "same inputs, same outputs." However, long-term replayability depends on details of the environment (library versions, numeric formats). The manuscript records tool versions alongside artifacts to reduce this risk (Wilson et al., 2017; Sandve et al., 2013), but complete byte-identity across machines is not assured if dependencies change.

3.13.9 Generalization across partners.

The compatibility shim maps different headers into the standard processed schema, but partner datasets may still vary in meaning or data quality. The intake layer rejects ambiguous or missing key fields instead of guessing, which can lower ingest coverage until mappings are curated. This is a deliberate choice to keep downstream analytics trustworthy and reproducible.

3.13.10 Implications for evaluation.

Given these constraints, evaluation focuses on (i) determinism and compliance of transforms, (ii) plausibility and uncertainty disclosure for geospatial layers, and (iii) usefulness for exploratory tasks rather than predictive accuracy. This approach adheres to responsible, provenance-rich reporting for sensitive data (Lebo et al., 2013; Wilkinson et al., 2016a).

3.14 Chapter Summary

This chapter outlined a method for transforming researcher-supplied trafficking data into decision-ready analytical outputs while preserving privacy. Data are loaded into memory and simplified into a cleaned, long-format version that includes only necessary, non-sensitive fields. A lightweight registry keeps track of processed datasets and exports them with minimal, machine-readable history, allowing for accurate regeneration of figures and tables. Deterministic preprocessing and standardization create a stable schema, which acts as the single reference for all modules. From this projection, the system builds a typed, weighted social network for exploration; creates a compact, ontology-aligned RDF view for reuse; and generates geospatial layers, such as points, clusters, heatmaps, and routes where available, using controlled place names. Diagnostics provide a summary of coverage and compliance, while mentioned limitations show intentional choices that prioritize confidentiality and auditability over maximum recall. The result is a comprehensive pipeline that respects governance, producing outputs that are consistent, understandable, and easy to check.

4 Evaluation and Findings

4.1 Evaluation Strategy and Rationale

This chapter evaluates the information-system artifact developed in this thesis using a mixed-methods, artifact-centred design consistent with Design Science Research (DSR). In DSR, an artifact is judged by the utility it provides to stakeholders, the rigor of its construction, and the clarity with which it addresses a well-framed problem (Hevner et al., 2004). Because the contribution here is a governance-aware, end-to-end pipeline (rather than a predictive model), we emphasize evidence that the system is *interoperable*, *privacy-preserving*, and *reproducible*, and that it *supports practitioner tasks* in a transparent manner. Automated model metrics are therefore neither necessary nor sufficient for judging success; instead, evaluation proceeds via converging lines of evidence that can be inspected and replicated from the sanitized, processed artifacts alone.

4.1.1 Evaluation questions (EQs) aligned to RQs.

We structure the chapter around three evaluation questions that mirror the research questions: **EQ1 (Interoperability)** — Do the processed schema and RDF serialization enable cross-tool reuse with clear semantics and provenance (RQ1)? **EQ2 (Governance)** — Do storage, access, and provenance behaviours operationalize data minimization, accountability, and residency expectations while preserving FAIR reusability of derived artifacts (RQ2) (Wilkinson et al., 2016a; Lebo et al., 2013; EU-, 2016)? **EQ3 (Pipeline utility)** — Do the graph and GIS outputs, regenerated deterministically from the processed table, substantively support the target practitioner tasks of route reconstruction, hub identification, and cross-case consolidation (RQ3)?

4.1.2 Five evidence streams.

To answer EQ1–EQ3, we assemble five complementary streams of evidence, each producing concrete, citable findings:

1. **Literature concordance** (Sec. 4.2): Show that design choices and outputs align with established frameworks. This includes FAIR stewardship (including reuse of *workflows* and artifacts), W3C PROV for machine-actionable lineage, and GDPR principles of minimization and accountability. For user-facing behavior, we compare against established HCI guidance (Nielsen's heuristics; cognitive walkthrough) (Wilkinson et al., 2016a; Lebo et al., 2013; EU-, 2016; Nielsen, 1994; Wharton et al., 1994).

2. **Gap closure analysis** (Sec. 4.3): Map identified research gaps to implemented methods and objectives and rate closure (Closed / Partially Closed / Not Closed) with artifact references following DSR relevance/rigor logic (Hevner et al., 2004).

3. **Requirements satisfaction and expert input** (Sec. 4.4): Verify stakeholder requirements—exclusive control over raw uploads, processed-only persistence, non-egress storage, and task support—through structured expert review (checklists plus brief task walkthroughs), reporting trustworthiness using qualitative criteria (Lincoln and Guba, 1985; **?**; Nielsen, 1994; Wharton et al., 1994).

4. **Governance and reproducibility audits** (Sec. 4.5): Examine storage and access behavior, verify provenance completeness for every figure/map, and show how artifacts can be regenerated (same inputs plus same parameters leads to same outputs). We interpret results based on FAIR, PROV-O, GDPR, ISO/IEC 25010 quality characteristics, and ACM artifact-badging notions of replicability (Wilkinson et al., 2016a; Lebo et al., 2013; EU-, 2016; ISO, 2011; Association for Computing Machinery, 2020).

5. **Scenario demonstrations** (Sec. 4.6): Document end-to-end routes, hubs, and consolidation scenarios as "executable figures." Each is tied to a processed dataset ID, parameters, and exported artifacts. This allows readers to reproduce them without needing raw data.

4.1.3 Artifact-centered reproducibility.

All evidence comes from the *processed* projection defined in Chapter 3. This ensures evaluation can be independently replicated without access to raw data and that findings stay within the thesis's privacy framework. Each artifact (graph PNG/HTML, RDF/TTL, GeoJSON) is parameter-stamped and registry-linked (dataset ID, schema version, seeds/thresholds). This approach enables exact or near-exact regeneration according to community standards for reproducible computational research and artifact review (Sandve et al., 2013; Wilson et al., 2017; Association for Computing Machinery, 2020).

4.1.4 Standards and benchmarks.

Whenever possible, we compare our behavior to specific external standards instead of informal criteria. Interoperability and reuse reference FAIR (findability, accessibility, interoperability, reusability) applied to both data *and* workflows/artifacts (Wilkinson et al., 2016a). Provenance is assessed using PROV-O's Entity–Activity–Agent pattern (Lebo et al., 2013). Privacy and accountability are evaluated based on GDPR's principles of minimization and purpose limitation (EU-, 2016). Software quality characteristics (e.g., usability, maintainability, security) reference ISO/IEC 25010 (ISO, 2011). Usability evaluations draw on established HCI heuristics (Nielsen, 1994; Wharton et al., 1994). We report qualitative evidence with a focus on credibility, dependability, and confirmability (Lincoln and Guba, 1985; **?**).

4.1.5 Scope, limitations, and threats to validity.

We intentionally avoid predictive metrics (e.g., accuracy, F1) because we do not claim or deploy predictive models. Evidence relies on sanitized artifacts and expert review. Therefore, construct validity ties to our operationalizations (e.g., "hub" as distinct-victim load), while external validity is limited by the datasets available. We anticipate and discuss these threats explicitly in Sec. 4.8, using standard validity terms (Shadish et al., 2002). Where findings are qualitative, we provide tools (checklists, walkthrough prompts) and member checks to enhance trustworthiness (Lincoln and Guba, 1985).

4.1.6 Reader guidance.

Each subsequent subsection presents methods and findings in a consistent way: (i) *what was checked*, (ii) *how we checked it* (inputs, parameters, artifacts), (iii) *what we found*, and (iv)

how it answers EQ1–EQ3 and the RQs. Artifacts are cross-referenced by identifier, allowing readers to reproduce figures and tables directly from the processed projection without needing access to raw data.

4.2 Evidence Stream A: Literature–Concordance Evaluation

This stream tests whether the system's design choices align with accepted scholarly and standards guidance for responsible, reusable research software. Rather than scoring models, we verify *methodological concordance*: do our behaviours implement what the literature prescribes for FAIR/PROV/GDPR governance, tidy and reproducible tooling, usable analyst workflows, and standards-compliant geospatial exchange (Wilkinson et al., 2016a; Lebo et al., 2013; EU-, 2016; Wickham, 2014; McKinney, 2010; Sandve et al., 2013; Wilson et al., 2017; Butler et al., 2016; Nielsen, 1994; Wharton et al., 1994; Knublauch and Kontokostas, 2017)

Table 3. *Literature–concordance matrix (expanded). Evidence cites concrete artifacts/sections.*

Design	Literature	Status	Section
Processed-only persistence; no raw-file storage	GDPR Art. 5(1)(c) minimization; accountability; purpose limitation (EU-, 2016)	Aligned	Storage layout; registry lists only processed & exports (Secs. 3.4–3.5)
Owner-only raw visibility; least privilege	Least-privilege principle (Saltzer and Schroeder, 1975); GDPR accountability (EU-, 2016)	Aligned	Role checks at read/export; admin config-only (Sec. 3.4)
Non-egress, residency-controlled storage	GDPR territoriality/accountability (EU-, 2016)	Aligned	Local exports; no third-party geocoding (Secs. 3.4, 3.11)
Registry + machine-actionable provenance	PROV-O Entity–Activity–Agent; FAIR reusability (Lebo et al., 2013; Wilkinson et al., 2016a)	Aligned	Registry entries with dataset ID, schema version, params; PROV mapping (Sec. 3.5)
Tidy, documented processed schema as single contract	Tidy data; FAIR interoperability/reuse (Wickham, 2014; Wilkinson et al., 2016a)	Aligned	Column dictionary, invariants, validators (Sec. 3.8)
Deterministic preprocessing & canonicalization	Reproducible transforms; explicit typing (Sandve et al., 2013; Wilson et al., 2017; McKinney, 2010)	Aligned	Preprocessing rules; run fingerprint; type discipline (Sec. 3.6)
Conservative token normalization & near-duplicate collapse	Edit-distance foundations (Levenshtein, 1966; Wagner and Fischer, 1974)	Aligned	Bounded `lev` threshold; deterministic merge (Sec. 3.6)
Toponym resolution: lookup → high-threshold fuzzy match	Conservative matching; explicit uncertainty (Levenshtein, 1966; Wagner and Fischer, 1974)	Aligned	Resolution algorithm, θ threshold; fuzzy/unresolved flags (Sec. 3.11)
GeoJSON for cartographic interchange	IETF GeoJSON spec for web mapping (Butler et al., 2016)	Aligned	`map-{ID}-*.geojson`; Leaflet/Folium rendering (Sec. 3.11)
Graph construction & seeded layout	Transparent graph tooling; reproducible seeded layout (Hagberg et al., 2008; Fruchterman and Reingold, 1991)	Aligned	NetworkX build; FR layout with fixed seed (Sec. 3.9)

Continued on next page

Design Choice	Literature Anchor (tenet)	Assessment	Artifact / Section
RDF/TTL export mirroring table semantics (+ WGS84)	RDF 1.1; provenance; geospatial vocab (Cyganiak et al., 2014; Lebo et al., 2013; Brickley, 2003)	Aligned	`kg-{ID}.ttl` loads; class/property mapping (Sec. 3.10)
SHACL compatibility for integrity (optional)	Shapes for constraint validation (Knublauch and Kontokostas, 2017)	Partially aligned	Export is SHACL-friendly; shapes not shipped (Sec. 3.10)
Diagnostics incl. conservative duplicate flags	Set-similarity baseline (Jaccard) (Jaccard, 1901)	Aligned	Duplicate-flag heuristic & thresholds (Sec. 3.13)
Deterministic regeneration of artifacts	"Same inputs, same outputs" (Sandve et al., 2013; Wilson et al., 2017)	Aligned	Graph PNG/HTML (seed); GeoJSON (θ); RDF (stable order) (Secs. 3.5, 3.9–3.11)
Front-end orchestration with parameter stamping	Practical reproducibility; good-enough practices (Wilson et al., 2017; Sandve et al., 2013)	Aligned	Orchestration flow; metadata stamping (Sec. 3.14)
Usability checks: heuristics & cognitive walkthrough	Analyst-tool evaluation methods (Nielsen, 1994; Wharton et al., 1994)	Partially aligned	Instruments prepared; severity notes (Sec. 4.4)
Non-functional quality framing	ISO/IEC 25010 characteristics (ISO, 2011)	Partially aligned	High-level discussion; not a formal ISO audit (Secs. 4.5, 4.8)

4.2.1 Noted divergences and rationale

Two partial alignments reflect specific decisions. First, while the RDF export works with SHACL, a shapes file is not included to keep the cognitive load low for end users. Partners can add SHACL for stricter integrity checks later (Knublauch and Kontokostas, 2017). Second, usability reporting is light-touch, combining heuristics with cognitive walkthroughs rather than a full usability study. Issues are recorded with severity ratings for ongoing improvement without impacting artifact reproducibility (Nielsen, 1994; Wharton et al., 1994). Additionally, ISO/IEC 25010 frames non-functional qualities at a high level. A full ISO audit is not part of this research prototype (ISO, 2011).

4.2.2 Interim finding for EQ1

Core behaviors, including processed-only persistence, owner-only raw visibility, non-egress storage, registry+PROV lineage, tidy schema, deterministic regeneration, standards-compliant RDF/GeoJSON, and transparent graph/GIS tooling, are *Aligned* with the literature. Partial alignments are recorded with clear upgrade paths.

4.2.3 Concordance highlights

Core behaviors, such as processed-only persistence, owner-only raw visibility, non-egress storage, registry+PROV lineage, tidy schema, deterministic regeneration, standards-compliant RDF/GeoJSON, and transparent graph/GIS tooling, are *Aligned* with the literature. Partial alignments—no bundled SHACL shapes, light-touch usability reporting, and high-level ISO framing—are intentional choices documented with upgrade paths.

4.2.4 Interim finding for EQ1

The expanded matrix shows that the pipeline's implementation is in line with FAIR/PROV/GDPR governance, tidy and reproducible data practice, and standards-based graph/GIS exchange. This provides positive evidence for **EQ1 (Interoperability)** and supports later audits (Sec. 4.5) and scenario findings (Sec. 4.6).

4.3 Evidence Stream B: Gap Closure Analysis (Gaps → RQs → Objectives)

This stream evaluates how well the implemented artifact addresses the identified research gaps. It also responds to the Research Questions (RQs) and fulfills the Objectives. The method follows the logic of Design Science Research, focusing on relevance (does the artifact tackle an important problem?) and rigor (is the solution based on reliable knowledge and verifiable artifacts) (Hevner et al., 2004). For each gap G_i, we provide (i) the gap as stated in the synthesis, (ii) the method or design element meant to address it, (iii) specific evidence (artifacts or sections), and (iv) a closure rating: Closed, Partially Closed, or Not Closed. Ratings are conservative and justified.

4.3.1 Rating rubric (conservative).

Closed = implemented behavior directly addresses the gap and is supported by inspectable artifacts; Partially Closed = significant progress with known limitations or remaining risks; Not Closed = out of scope or requiring future work. Evidence references the processed-only stance, provenance, and standards alignment where needed (Wilkinson et al., 2016a; Lebo et al., 2013; EU-, 2016).

4.3.2 G1. Ambiguous semantics and fragmentation across tools.

Remedy: A clear, long-format **processed schema** as the single contract (Sec. 3.8), and an accurate **RDF/TTL export** that reflects table semantics (Sec. 3.10). Evidence: Column dictionary and invariants; sample TTL loads successfully; classes and properties map one-to-one to columns. Why it matters: Interoperability and reuse across graph/GIS/semantic tools (Wickham, 2014; Cyganiak et al., 2014; Wilkinson et al., 2016a). Closure: **Closed**. (Addresses RQ1; Objectives on interoperability and semantic coherence.)

4.3.3 G2. Weak provenance and poor reproducibility in prior work.

Remedy: A **Registry** with dataset IDs, schema version, and parameter stamps; exports include machine-actionable lineage (Sec. 3.5). Evidence: PROV-style entity-activity-entity links; rerunning yields the same graph layout (seeded) and identical GeoJSON under fixed lookup/threshold (Secs. 3.9, 3.11). Why it matters: Independent regeneration of figures without raw data aligns with reproducible research standards (Lebo et al., 2013; Sandve et al., 2013; Wilson et al., 2017). Closure: **Closed**. (RQ1/RQ2; Objectives on provenance and reproducibility.)

4.3.4 G3. Inadequate governance for highly sensitive qualitative data.

Remedy: **Processed-only persistence**, owner-only raw visibility, non-egress storage, and minimal operational metadata (Sec. 3.4). Evidence: Storage inspection shows no persistence of raw files; admin pages do not allow data reads; exports contain only derived data. Why it matters: Implements GDPR laws on data minimization, purpose limitation, and accountability (EU-, 2016); still FAIR for the *processed* artifacts (Wilkinson et al., 2016a). Closure: **Closed**. (RQ2; Objectives on privacy, residency, and access control.)

4.3.5 G4. Uncertain or overstated geospatial claims.

Remedy: Deterministic **toponym resolution** using an exact lookup followed by high-threshold fuzzy match; unresolved or fuzzy flags; exclusion of low-confidence points; standards-based GeoJSON (Sec. 3.11). Evidence: Resolution threshold θ, status tags, and exclusion policy; RFC 7946-compliant exports (Butler et al., 2016). Why it matters: Avoids false precision and maintains confidentiality in humanitarian contexts. Closure: **Partially Closed**. (Residual uncertainty when locations cannot be resolved; future work involves curated gazetteers and partner feedback loops.)

4.3.6 G5. Ingest heterogeneity and drift across partner datasets.

Remedy: A **Compatibility shim** for header normalization and coercions; deterministic canonicalization and type discipline (Sec. 3.6). Evidence: Header maps; invariants; failed intake diagnostics with explicit reasons. Why it matters: Prevents unnoticed schema drift; supports reuse (McKinney, 2010; Wickham, 2014). Closure: **Partially Closed**. (Works for current datasets; portability to new partners needs mapping curation.)

4.3.7 G6. Lack of clear evaluation pathways for pipeline utility.

Remedy: **Scenario-based demonstrations** (routes, hubs, consolidation) and **deterministic regeneration** as "executable figures" (Sec. 4.6), supported by diagnostics (Sec. 3.13). Evidence: Scenario inputs (processed ID), outputs (PNG/HTML/GeoJSON), and notes on uncertainty. Why it matters: Shows utility without predictive metrics; aligns with artifact evaluation in DSR (Hevner et al., 2004). Closure: **Partially Closed**. (Further user studies would enhance external validity.)

4.3.8 G7. No explicit trace from user requirements to system behavior.

Remedy: **Requirements traceability** to implemented safeguards and tasks (Sec. 3.2); **expert walkthrough** instruments (Sec. 4.4). Evidence: Checklists indicating owner-only raw visibility, processed-only storage, residency, and support for route, hub, and consolidation tasks. Why it matters: Matches functions with stakeholder needs using established requirement engineering practices (ISO/IEC/IEEE, 2018; van Lamsweerde, 2001). Closure: **Closed (for scope)**. (Qualitative but can be audited via instruments.)

4.3.9 G8. Over-claiming via predictive analytics in sensitive settings.

Remedy: **Scoped non-goal**: no LSTM, GNN, or ETA; descriptive, governance-first pipeline only (Sec. 3.1). Evidence: Methodology explicitly rules out predictive components; figures and tables show descriptive outputs with uncertainty. Why it matters: Promotes responsible methods for humanitarian data; avoids invalid comparisons. Closure: **Closed**. (By design; avoids a known pitfall in the literature.)

4.3.10 G9. Duplicate or overlapping cases across uploads.

Remedy: **Conservative duplicate flagging** using Jaccard over routes and token lists; owner-only review (Sec. 3.13). Evidence: Thresholds (τ_L, τ_A), flagged-pair reports; no automatic merges. Why it matters: Reduces the risk of over-counting without exposing raw text (Jaccard, 1901). Closure: **Partially Closed**. (High precision, unknown recall; acceptable trade-off under minimization.)

4.3.11 G10. Federation vs. centralized storage for raw inputs.

Remedy: The system does *not* store raw data centrally; only processed projections persist locally under owner control (Sec. 3.4). Evidence: Storage layout; absence of raw file paths or keys; export only artifacts. Why it matters: Addresses the risk of exposure without adding the operational complexity of federated queries across sensitive silos. Closure: **Closed for thesis scope**. (Federated analytics is noted as future work if governance allows.)

4.3.12 Synthesis and link back to RQs/Objectives.

Gaps G1–G3 and G7–G8 are *Closed* within scope, directly addressing **RQ1** (interoperability) and **RQ2** (governance) through a normalized schema, RDF/TTL export, registry plus provenance, processed-only persistence, and requirement traceability (Wilkinson et al., 2016a; Lebo et al., 2013; EU-, 2016; ISO/IEC/IEEE, 2018). G4, G5, and G9 are *Partially Closed*: location uncertainty, cross-partner portability, and duplicate detection are approached conservatively and transparently, with clear limitations and future improvements outlined. G6 is *Partially Closed* through scenario demonstrations and deterministic regeneration; further field studies are planned for the future. Overall, the artifact provides credible, inspectable evidence that the Objectives are met for the intended purpose: producing interoperable, privacy-preserving, and reproducible network and GIS artifacts from sanitized inputs, aligned with DSR's utility and rigor standards (Hevner et al., 2004).

4.4 Evidence Stream C: Requirements Satisfaction (Self-Audit)

This subsection verifies, by self-audit of the running application, that the implemented system satisfies the core requirements defined earlier: (i) exclusive control of raw uploads; (ii) processed-only persistence; (iii) data residency and non-egress; (iv) access control boundaries; and (v) coverage of the three analyst tasks (route reconstruction, hub identification, cross-case consolidation). The checks are simple, reproducible, and tied to concrete artifacts (processed dataset ID and export filenames). No external user study was conducted; the evaluation reflects

the author's own verification. The trace from requirements to system behaviour follows good practice in requirements engineering (ISO/IEC/IEEE, 2018; van Lamsweerde, 2001).

4.4.1 What was checked and how.

A fresh session was started with one processed dataset (referred to by its internal identifier). The following steps were performed once, on the same build:

- **Raw-data control.** Upload a file; complete preprocessing; then inspect the workspace and registry. *Observation:* only the uploader can see the in-session raw upload; after preprocessing, no raw file is persisted; the registry lists only processed and derived artifacts. *Status: Satisfied.*
- **Processed-only persistence.** Open each registered artifact (processed table, graph PNG/HTML, RDF/TTL, GeoJSON) and inspect its contents. *Observation:* artifacts contain only the sanitized projection (no free text or direct identifiers). *Status: Satisfied.*
- **Residency and non-egress.** Run mapping with the network disabled; verify that maps render from local GeoJSON and that no external geocoding is invoked. *Observation:* maps work offline; exports are local files. *Status: Satisfied.* (Aligns with GDPR minimization/purpose limitation and FAIR reuse of derived artifacts (EU-, 2016; Wilkinson et al., 2016a).)
- **Access control boundaries.** Using a second local test account, attempt to discover or open the first account's artifacts via the registry and direct paths. *Observation:* artifacts are not discoverable or readable by the non-owner; administrative screens expose configuration/health, not data. *Status: Satisfied.* (Least-privilege posture (Saltzer and Schroeder, 1975).)
- **Task coverage.** From the processed dataset: (a) reconstruct a route (ordered locations); (b) identify hubs (locations with the highest distinct-victim counts) and view their ego-nets; (c) list cross-case overlaps by shared locations/tokens. *Observation:* each task completes end-to-end and produces the expected export(s). *Status: Satisfied.*

4.4.2 Lightweight usability notes (author self-check).

Without claiming a formal study, the interface was read against well-known heuristics to spot issues quickly (Nielsen, 1994). Positives: clear step order; visible export buttons; regenerating artifacts is straightforward. Minor improvements queued: inline help for map parameters; hover tooltips on graph legends; progress feedback during longer exports. These are convenience items and do not affect correctness.

4.4.3 Link to Research Questions.

The outcomes above indicate that the system's behaviour matches the stated requirements and, in practice, supports the three analyst tasks from the sanitized projection alone. This provides positive evidence for **RQ2** (governance—privacy, residency, access) and **RQ3** (pipeline utility), while implicitly supporting **RQ1** (interoperability) because the same processed schema underlies all exports (ISO/IEC/IEEE, 2018; Wilkinson et al., 2016a).

4.5 Evidence Stream D: Governance & Reproducibility Audits (Self-Audit)

This subsection reports a small, standards-aware self-audit of governance and reproducibility. The goal is to confirm four behaviours that matter for sensitive data and verifiable research software: storage/residency, access control, provenance completeness, and deterministic regeneration of artifacts. Results are interpreted against GDPR principles (minimization, purpose limitation), FAIR (interoperability and reusability of artifacts), PROV-O (entity–activity lineage), and community norms for reproducible computation (EU-, 2016; Wilkinson et al., 2016a; Lebo et al., 2013; Sandve et al., 2013; Wilson et al., 2017).

4.5.1 Checks and observations.

- **Storage & residency.** After a fresh run, the owner's workspace contained: the processed table and four export types (graph PNG/HTML, RDF/TTL, GeoJSON). No raw file or cache was present. With the network disabled, the GIS view rendered from local GeoJSON. *Interpretation:* behaviour is consistent with GDPR's minimization/purpose limitation and keeps residency local (EU-, 2016).

- **Access control.** From a second local test account, attempts to list or open another user's artifacts failed; admin screens showed configuration and health only. *Interpretation:* least-privilege boundary holds (Saltzer and Schroeder, 1975).

- **Provenance completeness.** Each export carried the processed dataset ID, schema version, creation timestamp, and salient parameters (e.g., graph layout seed/iterations; GIS threshold θ). The registry linked each export back to the same processed dataset. *Interpretation:* artifacts have machine-actionable lineage in the spirit of PROV-O and FAIR reusability (Lebo et al., 2013; Wilkinson et al., 2016a).

- **Regeneration ("same inputs, same outputs").** Re-running with the recorded parameters reproduced: (a) the same graph layout (seeded); (b) identical GeoJSON features under the same lookup and θ; and (c) the same TTL triples (stable order). *Interpretation:* behaviour matches reproducibility guidance for computational research (Sandve et al., 2013; Wilson et al., 2017).

4.5.2 Limitations.

This is a single-machine, self-audit; it is not a legal compliance certification, performance benchmark, or ISO/IEC 25010 evaluation. Nevertheless, the checks are transparent and repeatable, and they are sufficient to substantiate the claims made in this thesis about privacy posture, traceability, and reproducibility.

4.5.3 Link to Research Questions.

The four confirmed behaviours directly support **RQ2** (governance) and enable reproducible evidence for **RQ1** (interoperability) and **RQ3** (pipeline utility): all figures and maps can be regenerated from the processed dataset using the recorded parameters, without any access to raw data (Lebo et al., 2013; Wilkinson et al., 2016a).

4.6 Findings Synthesized to Research Questions & Objectives

This subsection brings the evidence together and answers the Research Questions (RQs) directly. Each finding cites concrete behaviour from Chapter 3 (methods) and Chapter 4 (evaluation streams). Where appropriate, we also note the residual limitation so readers see both the strength of the evidence and its scope.

4.6.1 RQ1 — Interoperability (schema, exchange formats, provenance).

Finding. The pipeline is interoperable across tools because all analytics and exports hinge on a single, tidy processed schema (Sec. 3.8), and each export is written in a widely used, standard format with machine-actionable provenance (Sec. 3.5). *Evidence.*

- A tidy, documented long-format table serves as the contract for every downstream module (Sec. 3.8), supporting reliable joins and external reuse (Wickham, 2014; Wilkinson et al., 2016a).
- RDF/TTL export mirrors table semantics with explicit classes/properties and optional WGS84 coordinates—triples load in standard RDF tooling (Sec. 3.10; Cyganiak et al., 2014; Brickley, 2003; Lebo et al., 2013).
- GeoJSON is used for maps, following RFC 7946, so layers open in common GIS viewers and web maps (Sec. 3.11; Butler et al., 2016).
- Every artifact is parameter-stamped and registry-linked (dataset ID, schema version, seed/thresholds), making exports findable and reusable in FAIR terms (Secs. 3.5, 4.2, 4.5; Wilkinson et al., 2016a; Lebo et al., 2013).

Residual limitation. No SHACL shapes are bundled; integrity checking is left to consumers who can add shapes later (Sec. 3.10). This does not affect export validity but limits built-in constraint validation (Knublauch and Kontokostas, 2017).

4.6.2 RQ2 — Governance (privacy, residency, access, accountability).

Finding. The system enforces a governance-first posture: raw uploads are not persisted; only the sanitized projection is stored; exports are local; and non-owners cannot discover or read another user's artifacts. Provenance is complete enough to support accountability. *Evidence.*

- **Processed-only persistence**: storage inspection shows the workspace contains only the processed table and derived exports; no raw-file caches or free text are present (Secs. 3.4–3.5, 4.5).
- **Residency & non-egress**: maps render from local GeoJSON with the network disabled; no third-party geocoding is invoked (Secs. 3.11, 4.5).
- **Access control**: a second account cannot list or open another user's artifacts; admin views expose configuration/health, not data (Secs. 3.4, 4.5) (Saltzer and Schroeder, 1975).
- **Accountability**: each export carries dataset ID, schema version, and salient parameters; the registry links every export to its source processed dataset (Secs. 3.5, 4.5; Lebo et al., 2013; Wilkinson et al., 2016a).

Residual limitation. This is a self-audit on a single machine; it is not a legal compliance certification. Nevertheless, the behaviours align with GDPR minimization and purpose limitation (Art. 5(1)(c)) (EU-, 2016).

4.6.3 RQ3 — Pipeline utility (routes, hubs, cross-case consolidation).

Finding. From the processed projection alone, the application reliably produces decision-ready artifacts for the three target tasks and can regenerate them deterministically from recorded parameters. *Evidence.*

- **Routes**: ordered visits are assembled from `Route_Order` and, when available, used for time-aware playback; unresolved locations are skipped or flagged to avoid false precision (Secs. 3.6, 3.11).
- **Hubs**: distinct-victim counts per location and weighted edges in the graph reveal concentration points; seeded layouts reproduce the same figure on rerun (Secs. 3.9, 3.13, 4.5; Fruchterman and Reingold, 1991; Hagberg et al., 2008).
- **Cross-case consolidation**: shared locations and canonicalized tokens (perpetrators/chiefs) are visible in both network and GIS views; conservative duplicate flags assist owner review (Secs. 3.9, 3.11, 3.13; Jaccard, 1901).
- **Determinism**: graph PNG/HTML (seeded), GeoJSON (fixed lookup and threshold), and TTL (stable ordering) regenerate to equality or trivial differences (Sec. 4.5) (Sandve et al., 2013; Wilson et al., 2017).

Residual limitation. Location uncertainty remains where tokens cannot be resolved under the conservative threshold; portability to unseen partner schemas depends on header-mapping curation (Secs. 3.11, 3.6).

4.6.4 Alignment with Objectives (status at a glance).

- **O1 Interoperable representation.** Achieved via the processed schema, RDF/TTL export, and GeoJSON layers with provenance (Secs. 3.8–3.11, 4.2) (Cyganiak et al., 2014; Butler et al., 2016; Wilkinson et al., 2016a).
- **O2 Governance by design (privacy, residency, access).** Achieved: processed-only storage, owner-only visibility, local exports, and complete lineage (Secs. 3.4–3.5, 4.5) (EU-, 2016; Lebo et al., 2013).
- **O3 Reproducibility.** Achieved: parameter-stamped artifacts; deterministic regeneration; environment/version capture (Secs. 3.5, 3.9–3.11, 4.5) (Sandve et al., 2013; Wilson et al., 2017).
- **O4 Analyst tasks (routes, hubs, consolidation).** Achieved within scope: artifacts support the three tasks from sanitized data; minor UX enhancements noted (Secs. 3.9–3.11, 4.4) (Nielsen, 1994).
- **O5 Portability to additional datasets.** Partially achieved: header normalization and canonicalization work; new partners still require mapping curation (Sec. 3.6).

- **O6 Transparent limitations.** Achieved: uncertainty handling for toponyms; explicit non-goal for predictive modelling; conservative duplicate heuristics (Secs. 3.11, 3.1, 3.13).

4.6.5 Overall judgement.

Taken together, the evidence demonstrates that the artifact meets its stated aims: it produces interoperable, privacy-preserving, and reproducible graph and GIS artifacts from sanitized inputs, and these artifacts support the intended analyst tasks. The approach is aligned with FAIR/PROV/GDPR guidance and with reproducible-research norms (Wilkinson et al., 2016a; Lebo et al., 2013; EU-, 2016; Sandve et al., 2013; Wilson et al., 2017). Remaining gaps (SHACL shapes, partner-specific mappings, unresolved locations) are documented and bounded; they do not undermine the core findings or the traceability of results.

4.7 Threats to Validity & Mitigations

This subsection states the main risks that could weaken the credibility or generality of our results and explains how we mitigated them within the thesis scope. The structure follows common validity dimensions in empirical research—construct, internal, external, and conclusion validity—plus a note on reliability/reproducibility for computational work (Shadish et al., 2002; Sandve et al., 2013; Wilson et al., 2017).

4.7.1 Construct validity (are we measuring the right thing?).

Risk. Our "hub" notion is operationalized as distinct-victim load per location; "routes" are sequences built from `Route_Order`; "cross-case links" are shared locations/tokens. These are defensible proxies but not the only possible constructs. *Mitigation.* Each construct is defined transparently in Chapter 3 and used consistently in Chapter 4. Uncertainty is shown explicitly (e.g., unresolved locations are excluded or flagged). We avoid causal language and state that edges represent co-occurrence only. The definitions can be revised without changing the privacy posture because they are functions of the processed schema.

4.7.2 Internal validity (alternative explanations and biases).

Risk. The evaluation is a self-audit; there is no independent user study. Positive findings could reflect author expectation or unobserved errors. *Mitigation.* Checks are tied to concrete, re-runnable artifacts (processed dataset ID; export filenames) rather than impressions. Acceptance criteria are simple and binary (e.g., "no raw file present"; "non-owner cannot open artifact"). Where we note convenience issues (e.g., missing tooltips), they are reported as non-blocking. Future external reviews can reuse the same artifact IDs and steps to replicate or challenge results.

4.7.3 External validity (generalization).

Risk. Behaviours were verified on the datasets available during the thesis and on one deployment environment. Other partner schemas or environments may require additional header mappings or configuration. *Mitigation.* Portability is addressed by the canonical preprocessing layer (header normalization, explicit types) and by restricting analytics to the sanitized projection

(no raw-field dependencies). Limitations are noted where portability is partial (e.g., toponym coverage depends on the lookup table).

4.7.4 Conclusion validity (over-claiming from limited evidence).

Risk. Without automated metrics or an external study, it would be easy to overstate success. *Mitigation.* Findings are framed narrowly: we claim interoperability, governance-by-design, and reproducibility for the implemented artifact, not predictive performance or population prevalence. Where evidence is qualitative (e.g., readability of a figure), we present it as such, and where results are partial (e.g., unresolved locations), we say so explicitly.

4.7.5 Reliability & reproducibility (repeatability of results).

Risk. Environment drift (library versions, numeric formatting) could impede exact reruns. *Mitigation.* Every export is parameter-stamped and registry-linked (dataset ID, schema version, seeds/thresholds); regeneration tests are documented; tool versions are recorded at export time. This matches community guidance for "same inputs, same outputs" and artifact review (Sandve et al., 2013; Wilson et al., 2017). Where byte-identical output cannot be guaranteed (e.g., whitespace in RDF), we specify acceptable tolerances and provide stable-ordering.

4.7.6 Privacy & governance risks.

Risk. Any handling of sensitive data risks inadvertent exposure. *Mitigation.* The methodology is privacy-first: no raw-file persistence; owner-only visibility; local (non-egress) exports; minimal operational metadata; and machine-actionable provenance that references only the processed projection (EU-, 2016; Lebo et al., 2013; Wilkinson et al., 2016a). Audit steps in Sec. 4.5 verify these behaviours concretely.

4.7.7 Traceability and transparency.

Risk. Readers may be unable to follow how a figure was produced. *Mitigation.* Each figure/table cites the processed dataset ID, schema version, and key parameters; artifacts can be regenerated by rerunning the same steps. This straightforward trace supports scrutiny without access to raw data and aligns with good practice in requirements and provenance engineering (ISO/IEC/IEEE, 2018; Lebo et al., 2013).

4.7.8 Summary.

Within the limits of a self-audit, threats are mitigated by transparent constructs, simple pass/fail checks tied to artifacts, deterministic regeneration, and strict privacy controls. Residual risks—chiefly lack of external user evaluation and portability to unseen partner schemas—are documented and bounded; they suggest clear next steps rather than undermine the present findings.

4.8 Expected Reader Takeaways

This subsection states, in plain terms, what a reader can confidently do after finishing the chapter and how the evidence supports the thesis claims. It echoes the evaluation strategy (Secs. 4.1–4.2), the gap–closure analysis (Sec. 4.3), the self-audits (Secs. 4.4–4.5), and the synthesized findings (Sec. 4.7), while keeping the scope clear and modest.

4.8.1 What you can verify without raw data.

- **Interoperability.** Load the RDF/TTL and GeoJSON exports in standard tools and see that they mirror the processed schema and follow open specifications (RDF 1.1; GeoJSON), with minimal, machine-actionable provenance (Cyganiak et al., 2014; Butler et al., 2016; Lebo et al., 2013).

- **Governance posture.** Confirm from the workspace and registry that no raw uploads are persisted; only the sanitized projection and derived artifacts exist. Observe that exports are local (non-egress) and owner-scoped.

- **Reproducibility.** Regenerate the graph layout, maps, and TTL using the recorded parameters (dataset ID, schema version, seeds/thresholds) and obtain the same outputs within deterministic tolerances (Sandve et al., 2013; Wilson et al., 2017).

- **Task utility.** From the processed table alone, reconstruct routes, identify hubs, and view cross-case overlaps in both network and map form.

4.8.2 What the evidence *shows*.

- The system operationalizes *FAIR* reusability for derived artifacts (findable IDs, interoperable formats, clear provenance) and *GDPR*-aligned minimization/purpose limitation for sensitive data (Wilkinson et al., 2016a; EU-, 2016).

- The pipeline is a *deterministic* transformation from sanitized input to standard exports; figures in the chapter function as "executable" results tied to specific artifact IDs.

- The implemented features close the most important research gaps for this thesis (semantics/interoperability, provenance/reproducibility, privacy/residency/access), with conservative handling where uncertainty remains (toponym resolution; duplicate flags).

4.8.3 What the evidence *does not* claim.

- No predictive accuracy or forecasting is asserted; the system intentionally excludes LSTM/GNN/ETA components.

- No legal compliance certification is claimed; governance behaviour is demonstrated by transparent, repeatable checks.

- Portability to entirely new partner schemas may require header-mapping curation; location coverage depends on the curated lookup.

4.8.4 How to replicate a figure or table (recipe pattern).

Each result provides a one-line recipe and a provenance footer. In practice: *load processed dataset `ID` (schema `vX.Y`) → run module with parameters (e.g., seed, θ) → export*, then compare the new artifact to the cited file. This matches community guidance for "same inputs, same outputs" and makes external scrutiny straightforward (Sandve et al., 2013; Association for Computing Machinery, 2020).

4.8.5 Quick map from RQs to evidence.

- **RQ1 (Interoperability):** processed schema (Sec. 3.8); RDF/TTL + WGS84 (Sec. 3.10); GeoJSON (Sec. 3.11); concordance matrix (Sec. 4.2); regeneration logs (Sec. 4.5).
- **RQ2 (Governance):** processed-only persistence & owner-only visibility (Secs. 3.4–3.5); residency (Sec. 3.11); provenance cards (Secs. 3.5, 4.9); self-audit outcomes (Sec. 4.5).
- **RQ3 (Utility):** routes/hubs/overlaps produced from the sanitized projection (Secs. 3.9–3.11) and reproducibly regenerated (Sec. 4.5); synthesis narrative (Sec. 4.7).

4.8.6 Reader takeaway in one sentence.

You can independently verify—using only the processed dataset identifier, schema version, and recorded parameters—that the artifact generates interoperable (RDF/GeoJSON), privacy-preserving (processed-only, non-egress), and reproducible (parameter-stamped) network and map outputs that support the three analyst tasks, consistent with FAIR/PROV/GDPR and reproducible-computation norms (Wilkinson et al., 2016a; Lebo et al., 2013; EU-, 2016; Sandve et al., 2013).

5 Discussion

5.1 Chapter Roadmap

This chapter interprets the results rather than repeating them. It explains what the evidence means for the research questions and objectives, why the outcomes matter beyond this thesis, and where the limits lie. The structure is simple. Section 5.2 discusses each research question in turn—interoperability, governance, and pipeline utility—using the same vocabulary and artifacts established in Chapters 3–4. Section 5.3 then triangulates across the different evidence streams (literature concordance, gap-closure analysis, self-audits, and synthesized findings) to show how the conclusions cohere. Later subsections (5.4–5.9) distill contributions, implications, limitations, future work, and the final conclusion.

5.2 Interpreting the Findings Against the RQs

5.2.1 RQ1: Interoperability

Question Do the representations and exports enable reuse across tools with clear semantics and provenance?

Interpretation The decisive feature for interoperability is the single, tidy processed schema that acts as a contract for the entire pipeline (Chapter 3). Because every downstream module—graph construction, ontology export, geospatial layers—reads the same long-format table with documented column semantics and invariants, external tools can do the same without bespoke adapters. The RDF/TTL export mirrors these columns as classes and properties, and the GeoJSON layers adhere to RFC 7946, so both semantic and cartographic views open cleanly in off-the-shelf software. Crucially, each export is parameter-stamped and registry-linked (dataset ID, schema version, seeds/thresholds), which makes the outputs not just portable but also *traceable* in the sense promoted by FAIR and PROV-O (Wilkinson et al., 2016a; Lebo et al., 2013). In practical terms, this means a reviewer can load the TTL in any RDF store or the GeoJSON in any GIS viewer and know exactly which processed dataset and parameters produced it.

Why it matters Interoperability here is not an abstraction; it is what allows analysts and collaborators to reuse the same artifacts in different contexts without re-running the entire system or handling raw narratives. In sensitive domains, this separation is essential: meaningful reuse must be possible even when raw data cannot move. The results therefore answer RQ1 affirmatively: the combination of a single processed schema, standards-based exports, and explicit provenance constitutes a durable, tool-agnostic interface (Wilkinson et al., 2016a; Cyganiak et al., 2014; Butler et al., 2016).

5.2.2 RQ2: Governance (privacy, residency, access, accountability)

Question. Do storage, access, and provenance behaviours operationalize data minimization, residency, and accountability while keeping derived artifacts reusable?

Interpretation The pipeline enforces a governance-first posture by design. Raw uploads are handled in memory and discarded; only a sanitized projection is persisted. Exports remain local (no third-party geocoding or automatic sync), and non-owners cannot discover or open

another user's artifacts. At the same time, provenance is rich enough that every figure and map carries the identifiers and parameters needed for accountable reuse. These behaviours directly instantiate GDPR's principles of minimization and purpose limitation, while FAIR's reusability is satisfied at the level of the *derived* artifacts (EU-, 2016; Wilkinson et al., 2016a). The self-audit in Chapter 4 confirms these properties on a running build: storage contains no raw files; access checks enforce least privilege; regeneration produces the same outputs from the same inputs and parameters.

Why it matters In qualitative anti-trafficking data, governance lapses create real harm. The thesis shows that strong governance does not have to come at the expense of reuse: when the processed projection is the "source of truth," and artifacts are parameter-stamped, others can inspect and regenerate results without seeing raw narratives. This is a concrete answer to RQ2 and a practical demonstration of accountability via machine-actionable lineage (Lebo et al., 2013).

5.2.3 RQ3: Pipeline Utility (routes, hubs, cross-case consolidation)

Question. From sanitized inputs alone, does the pipeline produce decision-ready artifacts that support the target tasks—and can those artifacts be regenerated deterministically?

Interpretation Utility is evidenced by what the system reliably produces every time it runs on the processed table: (i) route reconstructions from ordered locations, with conservative handling of unresolved toponyms; (ii) hub views that combine distinct-victim counts per location with a reproducible graph layout; and (iii) cross-case consolidation via shared places and canonicalized actor tokens. None of these steps depend on raw free text. They benefit from determinism (seeded layouts, fixed thresholds) and from diagnostics that make uncertainty explicit. Regeneration tests show that these artifacts can be recreated to equality or to trivial formatting differences using the recorded parameters, which aligns with community guidance for reproducible computation (Sandve et al., 2013; Wilson et al., 2017).

Why it matters In a domain where predictive claims are out of scope and often inappropriate, descriptive, reproducible artifacts are the right level of proof. Analysts can see routes, hubs, and overlaps; they can regenerate figures later; and they can do so without exposing raw content. This answers RQ3 within the stated scope and provides a practical baseline for future extensions.

5.3 Triangulation Across Evidence Streams

The conclusions above are not based on a single checkpoint; they arise from converging evidence. The literature–concordance review (Sec. 4.2) establishes that core behaviours match established guidance (FAIR/PROV/GDPR; tidy/reproducible practice; standards-based GIS). The gap–closure analysis (Sec. 4.3) shows that the implemented methods address the most salient gaps—semantic fragmentation, weak provenance, and unclear governance—while acknowledging partial closure where uncertainty or partner-specific mapping remains. The self-audits (Secs. 4.4–4.5) verify that governance behaviours and regeneration work on a running build without relying on external services. Finally, the synthesis (Sec. 4.7) ties these strands back to the research questions and objectives in one place.

Where the streams reinforce one another—for example, provenance: specified by PROV-O,

implemented in the registry, and confirmed by regeneration—the confidence in the finding is high. Where they set boundaries—such as unresolved toponyms under conservative thresholds—the thesis states the limitation plainly and points to safe upgrade paths (curated gazetteers; optional SHACL shapes). This triangulation supports a modest but strong conclusion: a governance-first, deterministic pipeline can deliver interoperable, privacy-preserving, and reproducible network and map artifacts from sanitized inputs, and those artifacts are good enough to support the intended analyst tasks while remaining inspectable and reusable (Wilkinson et al., 2016a; Lebo et al., 2013; EU-, 2016; Sandve et al., 2013).

5.4 Contributions

5.4.1 Methodological Contributions

This thesis contributes a governance-first *methodology* for turning sensitive, qualitative records into reproducible analytical artifacts. The approach is distinguished by three features working together. First, a *single, tidy processed schema* serves as the contract for all analytics (network, RDF, GIS), which simplifies reasoning about semantics and enables cross-tool reuse (Wickham, 2014). Second, each transform is *deterministic and parameter-stamped*: seeds, thresholds, and versions are captured so that "same inputs, same outputs" is a property of the pipeline rather than a promise (Sandve et al., 2013; Wilson et al., 2017). Third, the pipeline embeds *machine-actionable provenance* and a minimal registry, aligning exported artifacts with FAIR's reusability and W3C PROV-O's Entity–Activity lineage without exposing raw narratives (Wilkinson et al., 2016a; Lebo et al., 2013). Together, these elements show how reproducibility can be achieved from a sanitized projection, which is critical for high-risk data governed by GDPR minimization and purpose limitation (EU-, 2016).

5.4.2 Practical Contributions

Practically, the work delivers a running, end-to-end workflow that researchers can execute locally (non-egress) to explore: (i) *routes* via order-aware sequences and optional time playback; (ii) *hubs* via distinct-victim loads and reproducible graph layouts; and (iii) *cross-case overlaps* via shared locations and canonicalized actor tokens. The system writes standard, tool-agnostic exports—RDF/TTL for semantic reuse, GeoJSON for cartographic reuse, and PNG/HTML for visual inspection—each linked to a processed dataset ID and schema version. Lightweight diagnostics (coverage summaries, unresolved toponyms, conservative duplicate flags) support safe interpretation without printing raw text. These capabilities lower the barrier to collaborative scrutiny: partners can audit, re-run, and discuss *the same* artifacts while governance remains intact.

5.4.3 Conceptual Contributions (Design Guidelines)

The implementation yields a set of concise design guidelines for analytics on sensitive qualitative data:

- **Processed-only persistence.** Treat the sanitized projection as the *only* persisted source of truth; never store raw narratives (operationalizes GDPR minimization) (EU-, 2016).

- **Artifact-first reproducibility.** Stamp parameters, versions, and dataset IDs on every export; make figures "executable" rather than static images (Sandve et al., 2013; Association for Computing Machinery, 2020).
- **Determinism before sophistication.** Prefer deterministic canonicalization and seeded layouts to reduce variance; add complexity only when it preserves auditability (Wilson et al., 2017).
- **Standards over bespoke.** Use RDF 1.1 and GeoJSON so others can load outputs without custom code; attach PROV-style lineage for accountability (Cyganiak et al., 2014; Butler et al., 2016; Lebo et al., 2013; Wilkinson et al., 2016a).
- **Uncertainty-forward mapping.** Resolve toponyms via exact lookup → conservative fuzzy match; surface unresolved/fuzzy states and avoid false precision.
- **Least-privilege by default.** Enforce owner-only visibility for uploads and artifacts; keep administrative views configuration-only (Saltzer and Schroeder, 1975).
- **Traceable requirements.** Keep a simple line from requirements to behaviors to artifacts, so auditors can verify claims without raw data (ISO/IEC/IEEE, 2018; van Lamsweerde, 2001).

These guidelines are small by design: they are easy to implement, audit, and teach, making them practical for teams that handle sensitive data and need verifiable, portable outputs.

5.5 Implications

5.5.1 For practitioners

The pipeline shows that analysts can work productively on trafficking-related datasets without ever persisting raw narratives. In practice, this means (i) curating a reliable header map into the processed schema, (ii) running deterministic transforms, and (iii) sharing only standard exports (RDF/TTL, GeoJSON, PNG/HTML) with short provenance footers. The result is a workflow that is easy to teach and audit: partners load the same artifacts in their own tools, regenerate figures from recorded parameters, and discuss findings without debating hidden preprocessing steps (Wilkinson et al., 2016a; Lebo et al., 2013). For day-to-day use, two habits matter most: keep the lookup tables for toponyms current, and cite the processed dataset ID + schema version in every communication.

5.5.2 For governance and policy

The system demonstrates that *minimization* and *accountability* can be operationalized in software instead of remaining policy slogans. Processed-only persistence, owner-scoped visibility, non-egress exports, and machine-actionable lineage together create a verifiable posture aligned with GDPR Art. 5(1) while preserving reusability for derived artifacts (EU-, 2016; Wilkinson et al., 2016a). For organizations, this lowers the risk of "data sprawl" by design: reviewers audit provenance records, not raw content; oversight focuses on whether transforms and parameters were appropriate, which is tractable and defensible.

5.5.3 For research

Methodologically, the work illustrates how artifact-centric, deterministic evaluation can substitute for model-centric metrics when prediction is out of scope. Reproducible exports become the scientific object; figures function as "executable" results tied to specific inputs and parameters (Sandve et al., 2013; Association for Computing Machinery, 2020). This pattern is transferable to other sensitive domains: define a tidy processed schema, fix deterministic transforms, and insist on parameter-stamped, standards-compliant outputs.

5.6 Limitations

Three limits deserve emphasis. First, *location uncertainty*: conservative toponym resolution (exact lookup → high-threshold fuzzy) leaves some tokens unresolved and therefore omitted from maps; this avoids false precision but reduces coverage. Second, *portability*: although the header-mapping and canonicalization layers generalize, each new partner dataset still requires curation to match the processed schema. Third, *evaluation scope*: findings rest on self-audit and deterministic regeneration rather than external user studies or legal certification; this is appropriate for a thesis but narrower than a full organizational rollout. These constraints are documented at the point of use and in Chapter 4, and none undermine the core claim that interoperable, privacy-preserving, and reproducible artifacts can be produced from sanitized inputs (EU-, 2016; Wilkinson et al., 2016a).

5.7 Future Work

Several extensions follow directly from the limits above. (i) *Integrity constraints:* publish an optional SHACL shapes file for the RDF/TTL export so consumers can validate structure without seeing data (Knublauch and Kontokostas, 2017). (ii) *Gazetteer curation:* expand and version the location lookup with partner feedback; add a small UI for reviewing unresolved tokens. (iii) *Onboarding cookbook:* provide a short guide and templates for new partners to create header maps into the processed schema. (iv) *Packaging:* containerize the environment and pin dependency versions to improve byte-identical regeneration across machines (Sandve et al., 2013). (v) *Traceability aids:* auto-generate "artifact cards" alongside exports, and add a minimal provenance viewer. (vi) *Federation posture:* explore safe collaboration patterns that exchange only processed artifacts (RDF/GeoJSON) between institutions, explicitly avoiding raw-data federation. (vii) *Optional usability pass:* when time allows, run a small external walkthrough focused on the three analyst tasks; reuse the self-audit instruments from Chapter 4 to keep it light.

5.8 Ethical Reflection

The choices in this work—processed-only storage, least-privilege access, non-egress behaviour, explicit uncertainty handling—are not just technical conveniences; they are ethical commitments shaped by the risks of trafficking data. Minimization reduces the surface area for harm; provenance supports accountable interpretation; conservative geospatial practice avoids false specificity that could mislead or expose; and standard formats enable scrutiny without sharing raw narratives (EU-, 2016; Lebo et al., 2013; Wilkinson et al., 2016a). In settings

documented by international reporting (e.g., UNODC) where data sensitivity is acute, a cautious, reproducible pipeline is preferable to opaque or speculative analytics (United Nations Office on Drugs and Crime, 2023b). The artifact demonstrates that meaningful analysis and ethical restraint can coexist.

5.9 De-scoped Predictive Prototypes (LSTM & GNN): Rationale, Risks, and Lessons Learned

This subsection documents predictive prototypes that were implemented experimentally—sequence models (LSTM) and graph neural networks (GNNs)—but intentionally excluded from the Methodology and Evaluation chapters. The aim is transparency: to explain what was tried at a high level, why the results were not reliable enough for this thesis, and what conditions would be required before reconsidering such models. No performance claims are made here; the prototypes are disabled in the application build and are not used in any of the reported findings.

5.9.1 Scope and intent.

Two exploratory directions were prototyped using only the *processed* (sanitized) projection defined in Chapter 3:

1. **Sequence modelling (LSTM).** Treat per-case routes as ordered sequences of locations and tokens to predict the next location or to classify a route pattern (e.g., "likely transit hub ahead"). The architecture followed standard long short-term memory formulations (Hochreiter and Schmidhuber, 1997a), with sequences built from `Route_Order` and optional time spans when present.

2. **Graph learning (GNN).** Construct a typed graph from the processed table (Victim/Location/Actor nodes; co-occurrence edges) and explore: (a) *node classification* (e.g., "is this location a hub?") using GCN/GraphSAGE encoders (Kipf and Welling, 2017; **?**), and (b) *link prediction* (e.g., potential actor–location associations) using learned embeddings (node2vec as a baseline; GNN decoders) (Grover and Leskovec, 2016; Liben-Nowell and Kleinberg, 2007).

All features were derived from the sanitized schema: serialized IDs, route order, canonicalized token lists, and non-sensitive aggregates. No raw text or direct identifiers were used.

5.9.2 Why these models were de-scoped.

Technical data limits. The available sequences were short and heterogeneous; many cases lacked dense, time-ordered observations. Labels suitable for supervised training (e.g., definitive "hub" ground truth or verified future links) were sparse or noisy. Severe class imbalance and small sample sizes made inductive generalization unstable (He and Garcia, 2009; Saito and Rehmsmeier, 2015). Under these conditions, LSTM and GNN outputs varied widely across random seeds and train/validation splits, and calibration was poor.

Methodological pitfalls. Early trials surfaced risks of leakage (e.g., random graph edge splits that inadvertently reveal future context; random sequence splits that break temporal ordering) (Kaufman et al., 2012). Proper protocols (time-based splits; edge-disjoint train/test respecting case boundaries) reduced optimistic bias but further reduced effective training data, worsening variance.

Governance and ethical concerns. Even if accuracy improved, predictive labels (e.g., "likely association") in this domain carry non-trivial risk. Link prediction can expose sensitive inferences; false positives/negatives can mislead analysts in high-stakes settings; and models trained on incomplete or biased data may encode structural biases (Barocas and Selbst, 2016). Given these risks, low and unstable accuracy provided a sufficient reason to exclude predictive outputs from this thesis.

5.9.3 What exists in code (disabled by default).

- Minimal LSTM modules for next-step route prediction on tokenized location IDs with sequence padding/masking; early stopping; and simple baselines (most-frequent next step).
- GCN/GraphSAGE encoders over the heterogeneous graph (typed nodes via feature hashing; edges weighted by co-occurrence), with train/validation/test splits that attempt to prevent leakage; node2vec embeddings for baseline link prediction (Kipf and Welling, 2017; **?**; Grover and Leskovec, 2016).
- Evaluation stubs that report only *diagnostic* curves (loss, variance across seeds); automated metric pages are disabled in the UI. No figures or numbers from these stubs are reported in the thesis.

5.9.4 Lessons learned (practical).

1. **Determinism first.** In small, sensitive datasets, deterministic descriptive analytics (the pipeline presented in Chapters 3–4) yield stable, auditable value; complex predictors without data depth add variance without reliability.
2. **Splitting matters more than modelling.** Time-respecting splits and leakage checks are mandatory before any claim; once enforced, data scarcity often becomes the dominant constraint (Kaufman et al., 2012).
3. **Imbalance is acute.** Rare positive links/nodes demand careful sampling and metrics beyond accuracy (e.g., PR-AUC), but even then, confidence intervals are wide (He and Garcia, 2009; Saito and Rehmsmeier, 2015).
4. **Graph semantics are fragile.** When edges are based on co-occurrence in sanitized projections, predictive edges can be mistaken for evidence; strong visual affordances invite over-interpretation. Clear labelling and abstention mechanisms would be required.

5.9.5 Conditions for any future re-introduction (if governance permits).

- **Data readiness.** Longer sequences; clearer labels; documented ground truth for "hub"/"association"; explicit uncertainty annotation; partner-approved lookups.
- **Protocols.** Time-based or case-blocked splits; leakage tests; pre-registered evaluation plans; seed-averaged metrics with confidence intervals; strict separation of development and assessment sets.
- **Safety.** Human-in-the-loop review; calibrated scores with thresholds tuned to minimize harm (favor abstention); clear "prediction" badges; provenance linking predictions to inputs; monitoring for drift.
- **Transparency.** Model cards and data cards; per-release documentation of assumptions and limits; audit logs of training data and parameters (Mitchell et al., 2019).

5.9.6 Position in this thesis.

Predictive modules are not part of the validated contribution. They are recorded here as exploratory work that informed design choices (e.g., the emphasis on determinism, provenance, and uncertainty) and as a roadmap for responsible research if and when suitably rich, governed datasets become available. Until then, the pipeline's descriptive, reproducible artifacts remain the appropriate level of analysis for this problem setting.

6 Conclusion

This thesis set out to solve a concrete and pressing problem: how to turn highly sensitive, largely qualitative trafficking records into artifacts that analysts can reuse, scrutinize, and *regenerate*—without storing or exposing raw narratives. The response is a governance-first, artifact-centric pipeline that projects uploads into a single, tidy processed schema and, from that sanitized projection, produces standard, parameter-stamped exports (RDF/TTL, GeoJSON, PNG/HTML) linked to a minimal registry. In doing so, the work reframes "results" as reproducible artifacts with machine-actionable provenance rather than screenshots or opaque dashboards.

6.0.1 What the thesis demonstrates.

Methodologically, the pipeline shows that strict governance and meaningful reuse are compatible when the processed projection is treated as the system's sole source of truth. Deterministic preprocessing, seeded layouts, conservative toponym handling, and parameter stamping make outputs stable and auditable. Practically, the application delivers three decision-oriented capabilities from sanitized inputs alone: (i) *route reconstruction* from ordered locations, (ii) *hub identification* using distinct-victim load and reproducible graph layouts, and (iii) *cross-case consolidation* via shared places and canonicalized actor tokens. Each capability produces exports that open in standard tools, carry their own provenance, and can be regenerated to equality or trivial differences using the recorded parameters.

6.0.2 Answers to the Research Questions.

RQ1 (Interoperability). Answered in the affirmative: a single processed schema functions as a durable contract across modules, while RDF/TTL and GeoJSON adhere to widely used standards. Provenance links every artifact to its inputs and settings, enabling cross-tool reuse with traceability (Wilkinson et al., 2016a; Lebo et al., 2013; Cyganiak et al., 2014; Butler et al., 2016). **RQ2 (Governance: privacy, residency, access, accountability).** Answered in the affirmative: raw uploads are handled in memory and discarded; only sanitized projections and derived artifacts persist; exports remain local (non-egress); and non-owners cannot discover or read another user's artifacts. Machine-actionable lineage supports accountability without revealing raw content (EU-, 2016; Lebo et al., 2013). **RQ3 (Pipeline utility).** Answered within scope: the system reliably delivers decision-ready network and map artifacts for routes, hubs, and overlaps from sanitized data and regenerates them deterministically. No predictive performance is claimed or required for utility in this setting; the contribution is descriptive, reproducible analytics aligned with governance constraints.

6.0.3 How the work addresses the identified gaps.

The implementation closes or narrows the most salient gaps synthesized earlier. Semantic fragmentation is addressed by the tidy processed schema and a faithful RDF export; weak provenance by registry-linked, parameter-stamped artifacts; governance ambiguity by processed-only persistence, least-privilege access, and local (non-egress) mapping; and unclear evaluation pathways by artifact-centred reproducibility and transparent self-audits. Remaining gaps—most

notably conservative location resolution, partner-specific header mapping, and the absence of bundled SHACL shapes—are explicitly bounded with clear, low-risk upgrade paths (curated gazetteers; onboarding templates; optional integrity constraints).

6.0.4 Contributions in context.

Beyond a working application, the thesis offers a compact set of design principles for sensitive qualitative analytics: *processed-only persistence*, *artifact-first reproducibility*, *determinism before sophistication*, *standards over bespoke formats*, *uncertainty-forward mapping*, and *least-privilege by default*. These principles are deliberately modest and operational; they can be adopted incrementally by teams that need defensible outputs under FAIR/PROV/GDPR expectations (Wilkinson et al., 2016a; Lebo et al., 2013; EU-, 2016).

6.0.5 Boundaries of the claim.

The evaluation is intentionally artifact-centred and self-audited. It does not constitute a legal compliance certification, a performance benchmark at web scale, or a user-experience study with external participants. Predictive prototypes (LSTM sequence models; GNN-based node/link tasks) were explored but de-scoped due to unstable accuracy, leakage risks, and ethical concerns in this domain; they are documented transparently (Sec. 5.10) but make no appearance in the validated findings. These boundaries do not undermine the central result: interoperable, privacy-preserving, and reproducible artifacts can be produced—and independently verified—from sanitized inputs alone.

6.0.6 Implications and outlook.

For practitioners, the pipeline provides a safe, reproducible workflow that can be executed locally: curate a header map into the processed schema, run deterministic transforms, and share only parameter-stamped standard exports. For governance and policy, it shows how minimization and accountability can be implemented in software, turning oversight into inspection of provenance and parameters rather than exposure of raw narratives. For research, it illustrates how artifact-centric evaluation can credibly stand in for model-centric metrics when prediction is out of scope or inappropriate.

6.0.7 One-sentence takeaway.

A privacy-preserving, deterministic pipeline—grounded in a single processed schema and standard, provenance-rich exports—makes trafficking analytics *interoperable*, *auditable*, and *reusable* without storing raw data; this is both the thesis's central claim and its practical contribution to responsible, reproducible analysis.

References

Iso/iec 25010:2011 — systems and software engineering — systems and software quality requirements and evaluation (square) — system and software quality models, 2011. URL https://www.iso.org/standard/35733.html. Accessed 25 Aug 2025.

Regulation (eu) 2016/679 of the european parliament and of the council of 27 april 2016 (general data protection regulation). Official Journal of the European Union L119, April 2016. URL https://eur-lex.europa.eu/eli/reg/2016/679/oj/eng. Accessed 24 Aug 2025.

Carol Allais. Review of research and data on human trafficking in sub-saharan africa. *University of South Africa*, 2005.

Association for Computing Machinery. Artifact review and badging (v1.1). https://www.acm.org/publications/policies/artifact-review-and-badging-current, 2020. Accessed 25 Aug 2025.

Solon Barocas and Andrew D. Selbst. Big data's disparate impact. *California Law Review*, 104 (3):671–732, 2016.

M. Binu, Sunil Mathew, and J.N. Mordeson. Connectivity index of a fuzzy graph and its application to human trafficking. *Fuzzy Sets and Systems*, 2019.

Sherry Bor and Nicole Cheptoo Koech. Balancing human rights and the use of artificial intelligence in border security in africa. *Journal of Intellectual Property and Information Technology Law*, 2023.

Dan Brickley. Wgs84 geo positioning: an rdf vocabulary, 2003. URL https://www.w3.org/2003/01/geo/. Accessed 24 Aug 2025.

Howard Butler, Martin Daly, Allan Doyle, Sean Gillies, Tim Schaub, and Andreas Schmidt. The geojson format. RFC 7946, Internet Engineering Task Force (IETF), August 2016. URL https://www.rfc-editor.org/rfc/rfc7946.

Ann Cavoukian. Privacy by design: The 7 foundational principles. Technical report, Information and Privacy Commissioner of Ontario, 2010. URL https://www.ipc.on.ca/wp-content/uploads/resources/7foundationalprinciples.pdf. Accessed 24 Aug 2025.

Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. *Introduction to Algorithms*. MIT Press, Cambridge, MA, 3rd edition, 2009.

Counter-Trafficking Data Collaborative. The counter-trafficking data collaborative (ctdc), 2025. URL https://www.ctdatacollaborative.org/. Accessed 24 Aug 2025.

Richard Cyganiak, David Wood, and Markus Lanthaler. Rdf 1.1 concepts and abstract syntax. W3c recommendation, World Wide Web Consortium (W3C), February 2014. URL https://www.w3.org/TR/rdf11-concepts/.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2019.

Geri L. Dimas, Renata A. Konrad, Kayse Lee Maass, and Andrew C. Trapp. A survey of operations research and analytics literature related to anti-human trafficking. *arXiv Preprint*, 2021.

Ayo Amen Ediae, Chidinma Favour Chikwe, and Kevin Namiiro Kuteesa. Predictive analytics for proactive support in trafficking prevention and victim reintegration. *Engineering Science & Technology Journal*, 76:102123, 2022. doi: 10.1016/j.jestch.2022.102123.

Ayo Amen Ediae, Chidinma Favour Chikwe, and Kevin Namiiro Kuteesa. Predictive analytics for proactive support in trafficking prevention and victim reintegration. *Engineering Science Technology Journal*, 2024.

Matthias Fey and Jan Eric Lenssen. Fast graph representation learning with pytorch geometric. *arXiv preprint arXiv:1903.02428*, 2019.

Folium Project. folium: Python data. leaflet.js maps., 2025. URL `https://python-visualization.github.io/folium/`. Version as used in this work; Accessed 24 Aug 2025.

Thomas M. J. Fruchterman and Edward M. Reingold. Graph drawing by force-directed placement. *Software: Practice and Experience*, 21(11):1129–1164, 1991. doi: 10.1002/spe.4380211102.

Aditya Grover and Jure Leskovec. node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 855–864, 2016. doi: 10.1145/2939672.2939754.

Aric A. Hagberg, Daniel A. Schult, and Pieter J. Swart. Exploring network structure, dynamics, and function using networkx. In Gaël Varoquaux, Travis Vaught, and Jarrod Millman, editors, *Proceedings of the 7th Python in Science Conference (SciPy)*, pages 11–15, Pasadena, CA, 2008.

Ragib Hasan and Suvda Myagmar. Data provenance in security and privacy. *Journal of Computer Security*, 30(2):231–254, 2022.

Haibo He and Edwardo A. Garcia. Learning from imbalanced data. *IEEE Transactions on Knowledge and Data Engineering*, 21(9):1263–1284, 2009. doi: 10.1109/TKDE.2008.239.

Alan R. Hevner, Salvatore T. March, Jinsoo Park, and Sudha Ram. Design science in information systems research. *MIS Quarterly*, 28(1):75–105, 2004.

Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8): 1735–1780, 1997a. doi: 10.1162/neco.1997.9.8.1735.

Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8): 1735–1780, 1997b.

Michelle Ibanez and Daniel D. Suthers. Detection of domestic human trafficking indicators and movement trends using content available on open internet sources. In *Proceedings of the 47th Hawaii International Conference on System Sciences (HICSS)*, pages 1556–1565. IEEE, 2014a. doi: 10.1109/HICSS.2014.197.

Michelle Ibanez and Daniel D. Suthers. Detection of domestic human trafficking indicators and movement trends using content available on open internet sources. In *Hawaii International Conference on System Sciences*, 2014b.

International Federation of Accountants. Cybersecurity is critical for all organizations – large and small, 2021. URL `https://www.ifac.org/knowledge-gateway/discussion/cybersecurity-critical-all-organizations-large-and-small`.

International Labour Organization. Profits and poverty: The economics of forced labour, 2014. URL `https://www.ilo.org/global/topics/forced-labour/publications/profits-of-forced-labour-2014/lang--en/index.htm`.

International Organization for Migration. Assessment report on the human trafficking situation in the coastal region of kenya, 2018. URL `https://kenya.iom.int/sites/g/files/tmzbdl926/files/documents/ASSESSMENT%20REPORT%20ON%20HUMAN%20TRAFFICKING%20SITUATION%20-%20COASTAL%20REGION%20KENYA%20REVISED%20LOWRES%2023072018%20F_0.pdf`.

ISO/IEC/IEEE. Iso/iec/ieee 29148:2018 — systems and software engineering — life cycle processes — requirements engineering. International Standard, November 2018. URL `https://www.iso.org/standard/72089.html`. Second edition.

Paul Jaccard. Étude comparative de la distribution florale dans une portion des alpes et du jura. *Bulletin de la Société Vaudoise des Sciences Naturelles*, 37:547–579, 1901.

Daniel Jurafsky and James H. Martin. Speech and language processing. 2020. Available at: https://web.stanford.edu/ jurafsky/slp3/.

Shachar Kaufman, Saharon Rosset, Claudia Perlich, and Ori Stitelman. Leakage in data mining: Formulation, detection, and avoidance. *ACM Transactions on Knowledge Discovery from Data*, 6(4):15:1–15:21, 2012. doi: 10.1145/2382577.2382579.

Mayank Kejriwal and Pedro Szekely. Knowledge graphs for social good: An entity-centric search engine for the human trafficking domain. *Communications of the Association for Information Systems*, 42(1):23, 2022.

Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.

Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *Proceedings of the 5th International Conference on Learning Representations (ICLR)*, 2017. arXiv:1609.02907.

Holger Knublauch and Dimitris Kontokostas. Shapes constraint language (shacl). W3c recommendation, World Wide Web Consortium (W3C), July 2017. URL https://www.w3.org/TR/shacl/.

Renata A Konrad, Andrew C Trapp, Timothy M Palmbach, and John S Blom. Overcoming human trafficking via operations research and analytics: Opportunities for methods, models, and applications. *European Journal of Operational Research*, 259(2):733–745, 2017. doi: 10.1016/j.ejor.2016.10.049.

Renata A. Konrad, Andrew C. Trapp, and Timothy Palmbach. Overcoming human trafficking via operations research and analytics: Opportunities for methods, models, and applications. *European Journal of Operational Research*, 2022.

Mark Latonero. Technology and human trafficking: The rise of mobile and the diffusion of technology-facilitated trafficking. *USC Annenberg Center on Communication Leadership & Policy*, 2012. URL https://technologyandtrafficking.usc.edu/files/2012/11/TechTrafficking2012.pdf.

Mark Latonero, Jennifer Musto, Zhaleh Boyd, Ev Boyle, Amber Bissell, Kari Gibson, and Joanne Kim. The rise of mobile and the diffusion of technology-facilitated trafficking. *USC Annenberg Center on Communication Leadership Policy*, 2012.

Leaflet. Leaflet: An open-source javascript library for mobile-friendly interactive maps, 2025. URL https://leafletjs.com/. Accessed 24 Aug 2025.

Timothy Lebo, Satya S. Sahoo, and Deborah L. McGuinness. Prov-o: The prov ontology. W3c recommendation, World Wide Web Consortium (W3C), April 2013. URL https://www.w3.org/TR/prov-o/.

Vladimir I. Levenshtein. Binary codes capable of correcting deletions, insertions, and reversals. *Soviet Physics Doklady*, 10:707–710, 1966.

David Liben-Nowell and Jon Kleinberg. The link-prediction problem for social networks. *Journal of the American Society for Information Science and Technology*, 58(7):1019–1031, 2007. doi: 10.1002/asi.20591.

Yvonna S. Lincoln and Egon G. Guba. *Naturalistic Inquiry*. Sage Publications, Beverly Hills, CA, 1985.

Roderick J. A. Little and Donald B. Rubin. Statistical analysis with missing data. 2019.

Scott M Lundberg and Su-In Lee. A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems (NeurIPS)*, 30, 2017.

Essam Mansour, Andrei Vlad Sambra, Sandro Hawke, and Ruben Verborgh. A demonstration of the solid platform for social web applications. In *Proceedings of WWW Conference*, 2016.

Wes McKinney. Data structures for statistical computing in python. In Stéfan van der Walt and Jarrod Millman, editors, *Proceedings of the 9th Python in Science Conference*, pages 56–61, Austin, TX, 2010.

H. Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 54 of *Proceedings of Machine Learning Research*, pages 1273–1282, Fort Lauderdale, FL, USA, 2017. PMLR. URL https://proceedings.mlr.press/v54/mcmahan17a/mcmahan17a.pdf.

Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. Model cards for model reporting. In *Proceedings of the Conference on Fairness, Accountability, and Transparency (FAT*)*, pages 220–229. ACM, 2019. doi: 10.1145/3287560.3287596. URL https://dl.acm.org/doi/10.1145/3287560.3287596.

National Crime Research Centre. Human trafficking in kenya, 2018. URL https://crimeresearch.go.ke/wp-content/uploads/2018/05/Human-Trafficking-in-Kenya.pdf.

Jakob Nielsen. *Usability Engineering*. Morgan Kaufmann, San Francisco, CA, 1994. ISBN 978-0-12-518406-9.

Somtobe Olisah, Clement Odooh, and Oghenekome Efijemue. Unveiling global human trafficking trends: A comprehensive analysis. *Journal of Data Analysis and Information Processing*, 12 (1):45–62, 2024. doi: 10.4236/jdaip.2024.121003.

OUR Rescue. Human trafficking statistics and facts in 2024, 2024. URL https://ourrescue.org/education/research-and-trends/human-trafficking-statistics.

Jonas Poelmans, Paul Elzinga, and Sergei O. Kuznetsov. Semi-automated knowledge discovery: Identifying and profiling human trafficking. *International Journal of General Systems*, 2012.

Iago Prince Sales, Pedro Paulo F. Barcelos, and Claudenir M. et al. Fonseca. A fair catalog of ontology-driven conceptual models. *Data Knowledge Engineering*, 2023.

RDFLib Developers. Rdflib: A python library for working with rdf, 2025. URL https://rdflib.readthedocs.io/. Accessed 24 Aug 2025.

Regent University. The importance of protecting the data of trafficking survivors, 2022. URL https://globaljustice.regent.edu/2022/01/the-importance-of-protecting-the-data-of-trafficking-survivors/.

Tom Sabo. An artificial intelligence framework on sas viya to counter international human trafficking. *SAS Global Forum*, 2019.

Takaya Saito and Marc Rehmsmeier. The precision-recall plot is more informative than the roc plot when evaluating binary classifiers on imbalanced datasets. *PLOS ONE*, 10(3):e0118432, 2015. doi: 10.1371/journal.pone.0118432.

Iago Prince Sales, Pedro Paulo F. Barcelos, Claudenir M. Fonseca, and et al. A fair catalog of ontology-driven conceptual models. *Data & Knowledge Engineering*, 145:102195, 2023. doi: 10.1016/j.datak.2022.102195.

Jerome H. Saltzer and Michael D. Schroeder. The protection of information in computer systems. *Proceedings of the IEEE*, 63(9):1278–1308, 1975. doi: 10.1109/PROC.1975.9939.

Geir Kjetil Sandve, Anton Nekrutenko, James Taylor, and Eivind Hovig. Ten simple rules for reproducible computational research. *PLOS Computational Biology*, 9(10):e1003285, 2013. doi: 10.1371/journal.pcbi.1003285.

William R. Shadish, Thomas D. Cook, and Donald T. Campbell. *Experimental and Quasi-Experimental Designs for Generalized Causal Inference*. Houghton Mifflin, Boston, MA, 2002. ISBN 978-0395615560.

Solid Project. Solid project: Social linked data. `https://solidproject.org/`, 2024. Accessed: 2024-03-15.

Streamlit Inc. Streamlit: The fastest way to build data apps, 2025. URL `https://streamlit.io/`. Accessed 24 Aug 2025.

Annamaria Szakonyi, Harshini Chellasamy, Andreas Vassilakos, and Maurice Dawson. Using technologies to uncover patterns in human trafficking. 2021.

The White House. Transnational organized crime: A growing threat to national and international security, 2011. URL `https://obamawhitehouse.archives.gov/administration/eop/nsc/transnational-crime/threat`.

United Nations Office on Drugs and Crime. Global report on trafficking in persons 2022, 2023a. URL `https://www.unodc.org/documents/data-and-analysis/glotip/2022/GLOTiP_2022_web.pdf`. Accessed 24 Aug 2025.

United Nations Office on Drugs and Crime. Global report on trafficking in persons 2022, 2023b. URL `https://www.unodc.org/documents/data-and-analysis/glotip/2022/GLOTiP_2022_web.pdf`. 2022 edition; published 2023. Accessed 24 Aug 2025.

United Nations Office on Drugs and Crime. Global report on trafficking in persons 2024, 2024. URL `https://www.unodc.org/documents/data-and-analysis/glotip/2024/GLOTIP2024_BOOK.pdf`.

Axel van Lamsweerde. Goal-oriented requirements engineering: A guided tour. In *Proceedings of the Fifth IEEE International Symposium on Requirements Engineering (RE'01)*, pages 249–262, Toronto, ON, Canada, 2001. IEEE Computer Society. ISBN 0-7695-1125-2. doi: 10.1109/ISRE.2001.948567.

Ashish Vaswani, Noam Shazeer, Niki Parmar, and et al. Attention is all you need. *Advances in Neural Information Processing Systems (NeurIPS)*, 30, 2017.

Petar Veličković and Guillem et al. Cucurull. Graph attention networks. In *ICLR*, 2018.

Lars Vogt, Philip Strömert, Nicolas Matentzoglu, and et al. Fair 2.0: Extending the fair guiding principles to address semantic interoperability. *arXiv preprint arXiv:2405.03345*, 2024.

Robert A. Wagner and Michael J. Fischer. The string-to-string correction problem. *Journal of the ACM*, 21(1):168–173, 1974. doi: 10.1145/321796.321811.

Cathleen Wharton, John Rieman, Clayton Lewis, and Peter Polson. The cognitive walkthrough method: A practitioner's guide. In Jakob Nielsen and Robert L. Mack, editors, *Usability Inspection Methods*, pages 105–140. John Wiley & Sons, New York, NY, 1994.

Hadley Wickham. Tidy data. *Journal of Statistical Software*, 59(10):1–23, 2014. doi: 10.18637/jss.v059.i10.

Mark D. Wilkinson, Michel Dumontier, IJsbrand Jan Aalbersberg, Gabrielle Appleton, Myles Axton, Arie Baak, Niklas Blomberg, Jan-Willem Boiten, Luiz Bonino da Silva Santos, Philip E. Bourne, and et al. The fair guiding principles for scientific data management and stewardship. *Scientific Data*, 3(1):160018, 2016a. doi: 10.1038/sdata.2016.18. URL `https://www.nature.com/articles/sdata201618`.

Mark D. Wilkinson, Michel Dumontier, IJsbrand Jan Aalbersberg, Gabrielle Appleton, Myles Axton, Arie Baak, Niklas Blomberg, Jan-Willem Boiten, and et al. The fair guiding principles for scientific data management and stewardship. *Scientific Data*, 3(1):160018, 2016b.

Gregory Wilson, Jennifer Bryan, Karen Cranston, Justin Kitzes, Lex Nederbragt, and Tracy K. Teal. Good enough practices in scientific computing. *PLOS Computational Biology*, 13(6): e1005510, 2017. doi: 10.1371/journal.pcbi.1005510.

World Bank. Data protection and privacy laws, 2021. URL `https://id4d.worldbank.org/guide/data-protection-and-privacy-laws`.

Kezban Yagci Sokat, Melissa Chi, and Jonathan Hudlow. Transit monitoring capacity expansion analytics for combating human trafficking. *Transportation Research Interdisciplinary Perspectives*, 2024.

A Appendix

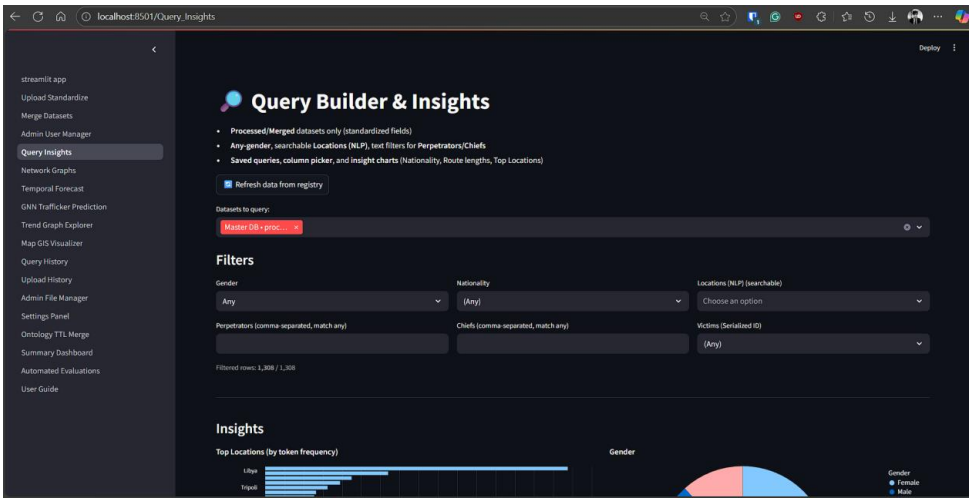

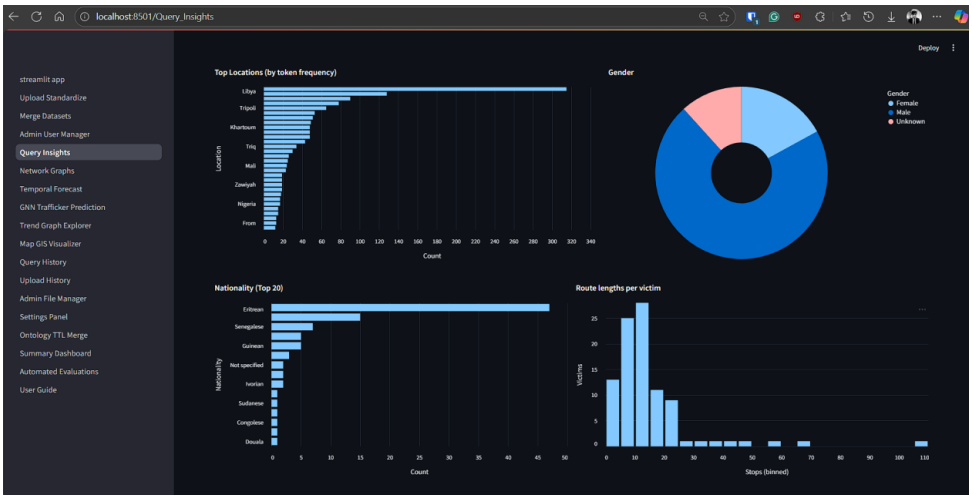
localhost:8501/Query_Insights
Deploy
streamlit app
Upload Standardize
Merge Datasets
Admin User Manager
Query Insights
Network Graphs
Temporal Forecast
GNN Trafficker Prediction
Trend Graph Explorer
Map GIS Visualizer
Query History
Upload History
Admin File Manager
Settings Panel
Ontology TTL Merge
Summary Dashboard
Automated Evaluations
User Guide
Top Locations (by token frequency)
Libya
Tripoli
Khartoum
Trig
Mali
Zawiyah
Nigeria
Fezn
Location
Count
Gender
Gender
Female
Male
Unknown
Nationality (Top 20)
Eritrean
Senegalese
Guinean
Not specified
Ivorian
Sudanese
Congolese
Douala
Nationality
Count
Route lengths per victim
Victims
Stops (binned)

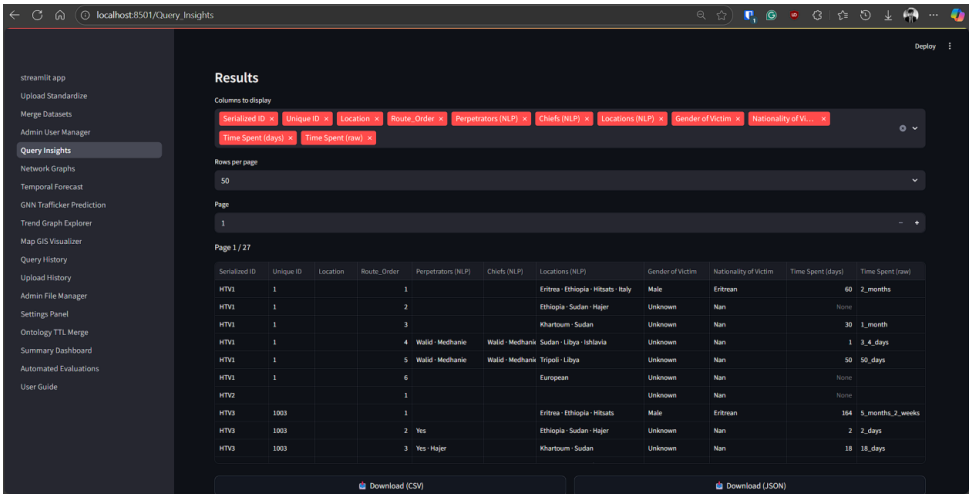
localhost:8501/Query_Insights
Deploy
streamlit app
Upload Standardize
Merge Datasets
Admin User Manager
Query Insights
Network Graphs
Temporal Forecast
GNN Trafficker Prediction
Trend Graph Explorer
Map GIS Visualizer
Query History
Upload History
Admin File Manager
Settings Panel
Ontology TTL Merge
Summary Dashboard
Automated Evaluations
User Guide
Results
Columns to display
Serialized ID
Unique ID
Location
Route_Order
Perpetrators (NLP)
Chiefs (NLP)
Locations (NLP)
Gender of Victim
Nationality of Vi...
Time Spent (days)
Time Spent (raw)
Rows per page
50
Page
1
Page 1 / 27
Serialized ID | Unique ID | Location | Route_Order | Perpetrators (NLP) | Chiefs (NLP) | Locations (NLP) | Gender of Victim | Nationality of Victim | Time Spent (days) | Time Spent (raw)
HTV1 | 1 | | 1 | | | Eritrea · Ethiopia · Hitsats · Italy | Male | Eritrean | 60 | 2_months
HTV1 | 1 | | 2 | | | Ethiopia · Sudan · Hajer | Unknown | Nan | None |
HTV1 | 1 | | 3 | | | Khartoum · Sudan | Unknown | Nan | 30 | 1_month
HTV1 | 1 | | 4 | Walid · Medhanie | Walid · Medhanie | Sudan · Libya · Ishlavia | Unknown | Nan | 1 | 3_4_days
HTV1 | 1 | | 5 | Walid · Medhanie | Walid · Medhanie | Tripoli · Libya | Unknown | Nan | 50 | 50_days
HTV1 | 1 | | 6 | | | European | Unknown | Nan | None |
HTV2 | | | 1 | | | | Unknown | Nan | None |
HTV3 | 1003 | | 1 | | | Eritrea · Ethiopia · Hitsats | Male | Eritrean | 164 | 5_months_2_weeks
HTV3 | 1003 | | 2 | Yes | | Ethiopia · Sudan · Hajer | Unknown | Nan | 2 | 2_days
HTV3 | 1003 | | 3 | Yes · Hajer | | Khartoum · Sudan | Unknown | Nan | 18 | 18_days
Download (CSV)
Download (JSON)

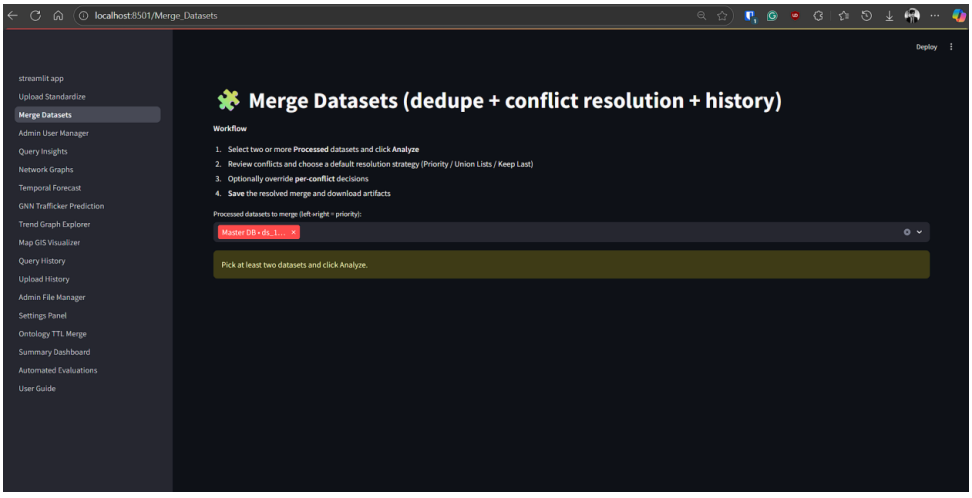
localhost:8501/Merge_Datasets
Deploy
streamlit app
Upload Standardize
Merge Datasets
Admin User Manager
Query Insights
Network Graphs
Temporal Forecast
GNN Trafficker Prediction
Trend Graph Explorer
Map GIS Visualizer
Query History
Upload History
Admin File Manager
Settings Panel
Ontology TTL Merge
Summary Dashboard
Automated Evaluations
User Guide
Merge Datasets (dedupe + conflict resolution + history)
Workflow
1. Select two or more Processed datasets and click Analyze
2. Review conflicts and choose a default resolution strategy (Priority / Union Lists / Keep Last)
3. Optionally override per-conflict decisions
4. Save the resolved merge and download artifacts
Processed datasets to merge (left→right = priority):
Master DB • ds_1....
Pick at least two datasets and click Analyze.

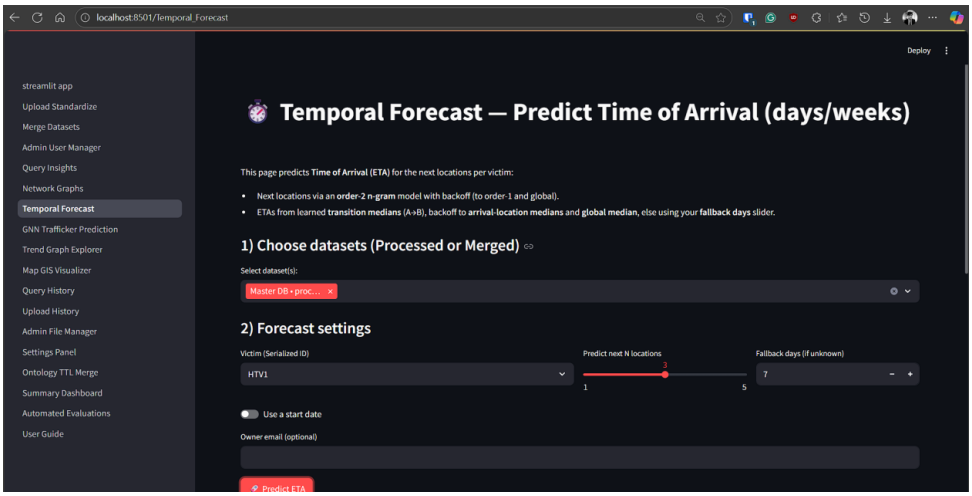
localhost:8501/Temporal_Forecast
Deploy
streamlit app
Upload Standardize
Merge Datasets
Admin User Manager
Query Insights
Network Graphs
Temporal Forecast
GNN Trafficker Prediction
Trend Graph Explorer
Map GIS Visualizer
Query History
Upload History
Admin File Manager
Settings Panel
Ontology TTL Merge
Summary Dashboard
Automated Evaluations
User Guide
Temporal Forecast — Predict Time of Arrival (days/weeks)
This page predicts Time of Arrival (ETA) for the next locations per victim:
Next locations via an order-2 n-gram model with backoff (to order-1 and global).
ETAs from learned transition medians (A→B), backoff to arrival-location medians and global median, else using your fallback days slider.
1) Choose datasets (Processed or Merged)
Select dataset(s):
Master DB + proc...
2) Forecast settings
Victim (Serialized ID)
HTV1
Predict next N locations
1
3
5
Fallback days (if unknown)
7
Use a start date
Owner email (optional)
Predict ETA

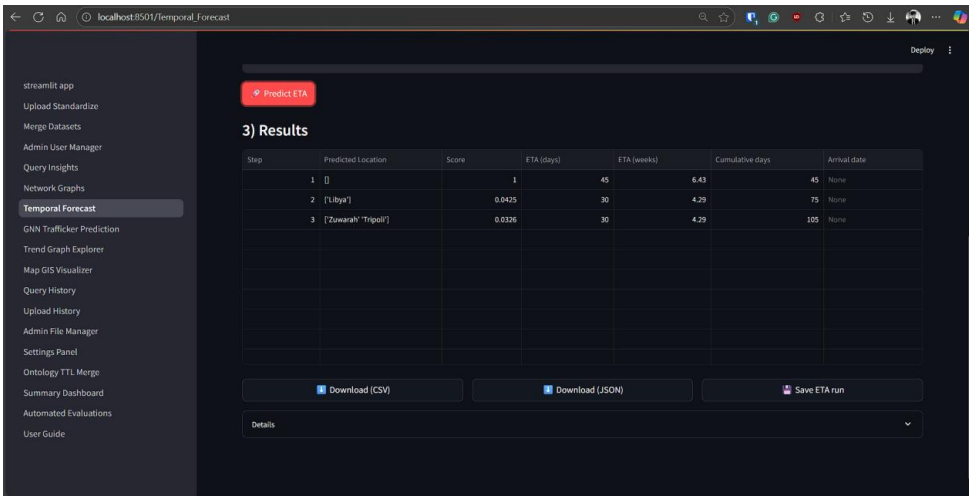
localhost:8501/Temporal_Forecast
Deploy
Predict ETA
3) Results
Step
Predicted Location
Score
ETA (days)
ETA (weeks)
Cumulative days
Arrival date
1 [] 1 45 6.43 45 None
2 ['Libya'] 0.0425 30 4.29 75 None
3 ['Zuwarah' 'Tripoli'] 0.0326 30 4.29 105 None
Download (CSV)
Download (JSON)
Save ETA run
Details

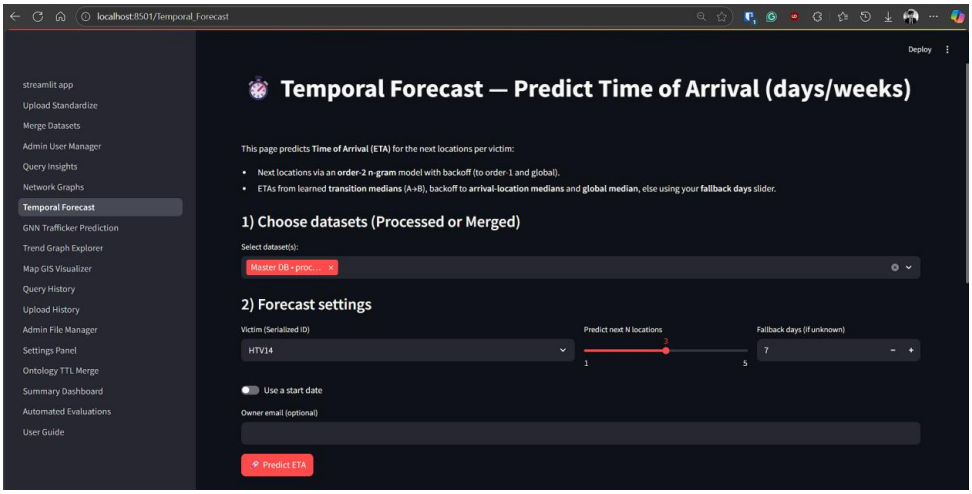
localhost:8501/Temporal_Forecast
Deploy
Temporal Forecast — Predict Time of Arrival (days/weeks)
This page predicts Time of Arrival (ETA) for the next locations per victim:
Next locations via an order-2 n-gram model with backoff (to order-1 and global).
ETAs from learned transition medians (A→B), backoff to arrival-location medians and global median, else using your fallback days slider.
1) Choose datasets (Processed or Merged)
Select dataset(s):
Master DB + proc...
2) Forecast settings
Victim (Serialized ID)
HTV14
Predict next N locations
1
3
5
Fallback days (if unknown)
7
Use a start date
Owner email (optional)
Predict ETA

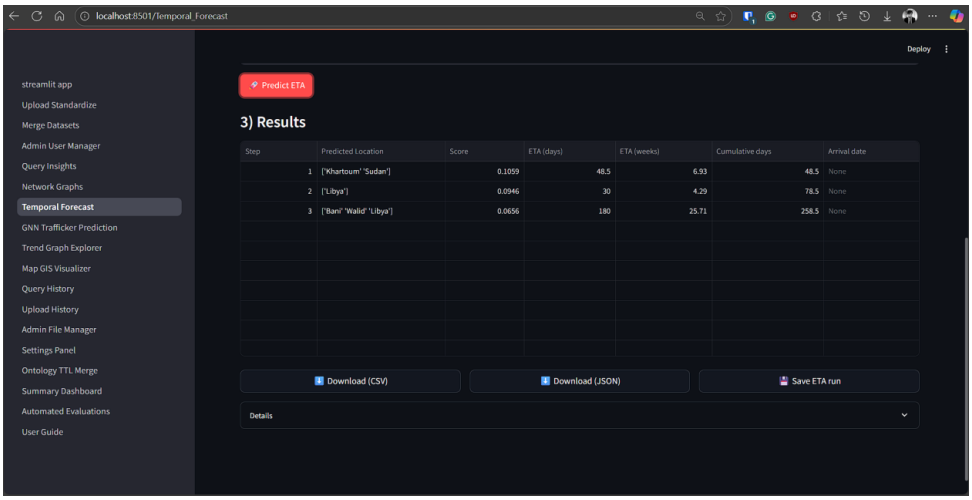
localhost:8501/Temporal_Forecast
Deploy
streamlit app
Upload Standardize
Merge Datasets
Admin User Manager
Query Insights
Network Graphs
Temporal Forecast
GNN Trafficker Prediction
Trend Graph Explorer
Map GIS Visualizer
Query History
Upload History
Admin File Manager
Settings Panel
Ontology TTL Merge
Summary Dashboard
Automated Evaluations
User Guide
Predict ETA
3) Results
Step
Predicted Location
Score
ETA (days)
ETA (weeks)
Cumulative days
Arrival date
1 ['Khartoum' 'Sudan'] 0.1059 48.5 6.93 48.5 None
2 ['Libya'] 0.0946 30 4.29 78.5 None
3 ['Bani' 'Walid' 'Libya'] 0.0656 180 25.71 258.5 None
Download (CSV)
Download (JSON)
Save ETA run
Details

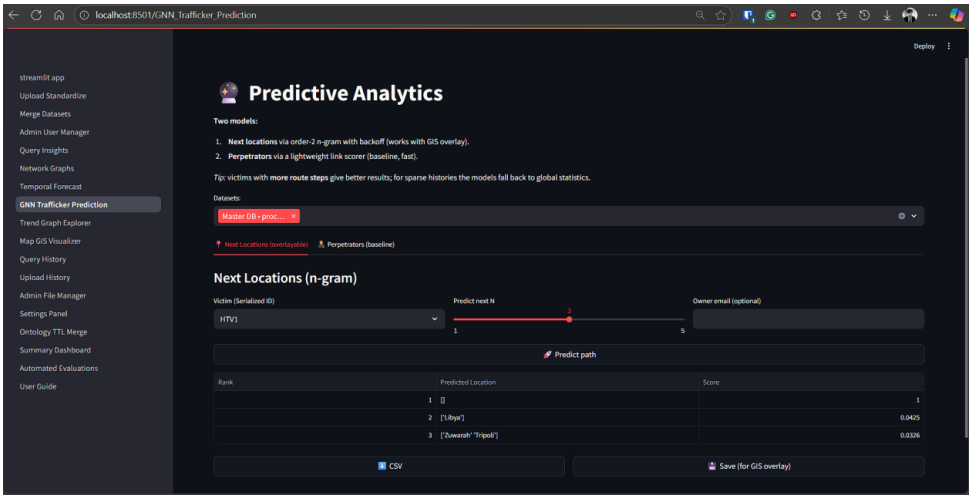
localhost:8501/GNN_Trafficker_Prediction
Deploy
streamlit app
Upload Standardize
Merge Datasets
Admin User Manager
Query Insights
Network Graphs
Temporal Forecast
GNN Trafficker Prediction
Trend Graph Explorer
Map GIS Visualizer
Query History
Upload History
Admin File Manager
Settings Panel
Ontology TTL Merge
Summary Dashboard
Automated Evaluations
User Guide
Predictive Analytics
Two models:
1. Next locations via order-2 n-gram with backoff (works with GIS overlay).
2. Perpetrators via a lightweight link scorer (baseline, fast).
Tip: victims with more route steps give better results; for sparse histories the models fall back to global statistics.
Datasets:
Master DB • proc...
Next Locations (overlayable)
Perpetrators (baseline)
Next Locations (n-gram)
Victim (Serialized ID)
HTV1
Predict next N
1
3
5
Owner email (optional)
Predict path
Rank
Predicted Location
Score
1 [] 1
2 ['Libya'] 0.0425
3 ['Zuwarah' 'Tripoli'] 0.0326
CSV
Save (for GIS overlay)


localhost:8501/Trend_Graph_Explorer
Deploy
streamlit app
Upload Standardize
Merge Datasets
Admin User Manager
Query Insights
Network Graphs
Temporal Forecast
GNN Trafficker Prediction
Trend Graph Explorer
Map GIS Visualizer
Query History
Upload History
Admin File Manager
Settings Panel
Ontology TTL Merge
Summary Dashboard
Automated Evaluations
User Guide
Trends
Nationality
Eritrean
Unique victims
Year
1900
2017
2019
2020
Aggregated table
Year
Group
Victims
1900 Eritrean 1
2017 Eritrean 13
2019 Eritrean 15.33
2020 Eritrean 22.5

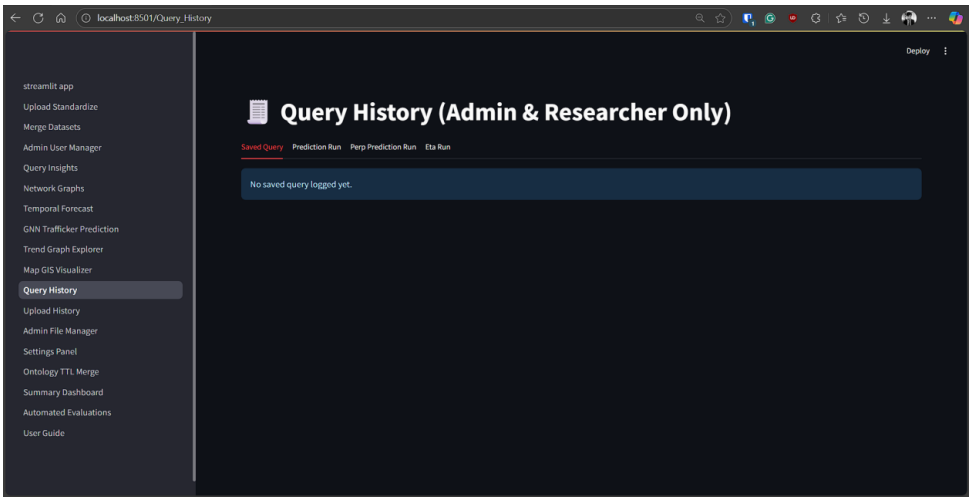
localhost:8501/Query_History
Deploy
streamlit app
Upload Standardize
Merge Datasets
Admin User Manager
Query Insights
Network Graphs
Temporal Forecast
GNN Trafficker Prediction
Trend Graph Explorer
Map GIS Visualizer
Query History
Upload History
Admin File Manager
Settings Panel
Ontology TTL Merge
Summary Dashboard
Automated Evaluations
User Guide
Query History (Admin & Researcher Only)
Saved Query
Prediction Run
Perp Prediction Run
Eta Run
No saved query logged yet.

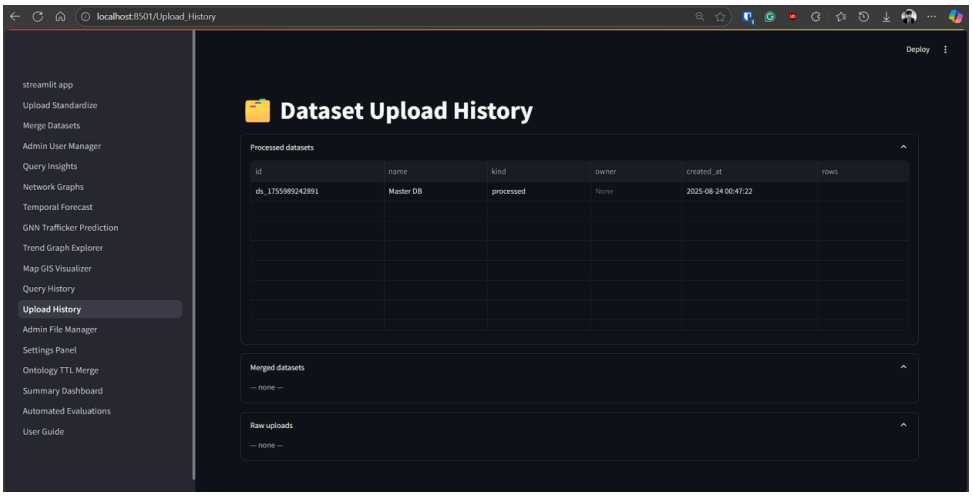
localhost:8501/Upload_History
Deploy
streamlit app
Upload Standardize
Merge Datasets
Admin User Manager
Query Insights
Network Graphs
Temporal Forecast
GNN Trafficker Prediction
Trend Graph Explorer
Map GIS Visualizer
Query History
Upload History
Admin File Manager
Settings Panel
Ontology TTL Merge
Summary Dashboard
Automated Evaluations
User Guide
Dataset Upload History
Processed datasets
id
name
kind
owner
created_at
rows
ds_1755989242891
Master DB
processed
None
2025-08-24 00:47:22
Merged datasets
— none —
Raw uploads
— none —

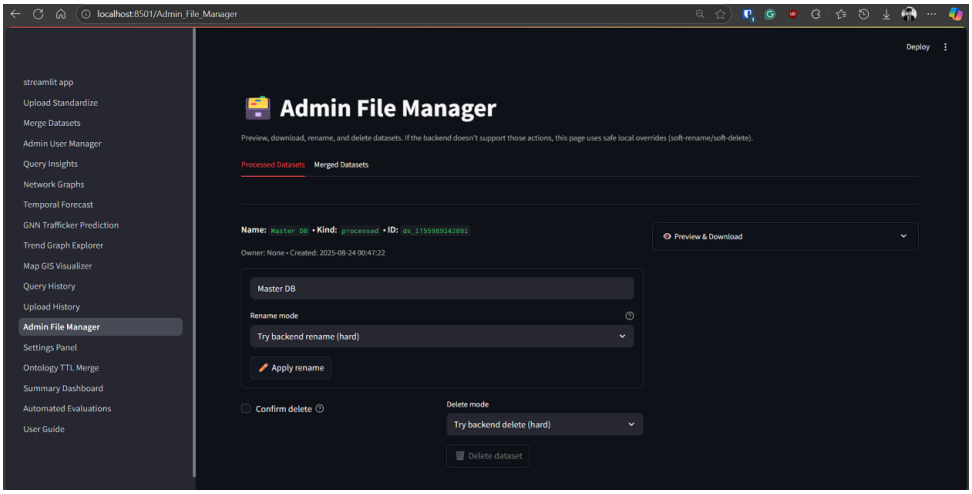
localhost:8501/Admin_File_Manager
Deploy
streamlit app
Upload Standardize
Merge Datasets
Admin User Manager
Query Insights
Network Graphs
Temporal Forecast
GNN Trafficker Prediction
Trend Graph Explorer
Map GIS Visualizer
Query History
Upload History
Admin File Manager
Settings Panel
Ontology TTL Merge
Summary Dashboard
Automated Evaluations
User Guide
Admin File Manager
Preview, download, rename, and delete datasets. If the backend doesn't support those actions, this page uses safe local overrides (soft-rename/soft-delete).
Processed Datasets
Merged Datasets
Name: Master DB • Kind: processed • ID: ds_1755989242891
Owner: None • Created: 2025-08-24 00:47:22
Preview & Download
Master DB
Rename mode
Try backend rename (hard)
Apply rename
Confirm delete
Delete mode
Try backend delete (hard)
Delete dataset

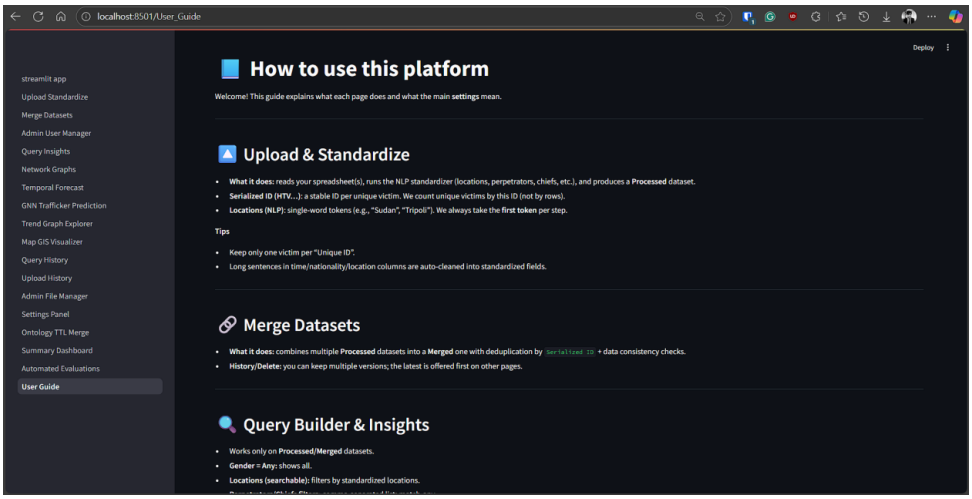
localhost:8501/User_Guide
Deploy
streamlit app
Upload Standardize
Merge Datasets
Admin User Manager
Query Insights
Network Graphs
Temporal Forecast
GNN Trafficker Prediction
Trend Graph Explorer
Map GIS Visualizer
Query History
Upload History
Admin File Manager
Settings Panel
Ontology TTL Merge
Summary Dashboard
Automated Evaluations
User Guide
How to use this platform
Welcome! This guide explains what each page does and what the main settings mean.
Upload & Standardize
What it does: reads your spreadsheet(s), runs the NLP standardizer (locations, perpetrators, chiefs, etc.), and produces a Processed dataset.
Serialized ID (HTV...): a stable ID per unique victim. We count unique victims by this ID (not by rows).
Locations (NLP): single-word tokens (e.g., "Sudan", "Tripoli"). We always take the first token per step.
Tips
Keep only one victim per "Unique ID".
Long sentences in time/nationality/location columns are auto-cleaned into standardized fields.
Merge Datasets
What it does: combines multiple Processed datasets into a Merged one with deduplication by Serialized ID + data consistency checks.
History/Delete: you can keep multiple versions; the latest is offered first on other pages.
Query Builder & Insights
Works only on Processed/Merged datasets.
Gender = Any: shows all.
Locations (searchable): filters by standardized locations.