



Manufacturing & Service Operations Management

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Linking Multisite Sex Ad Data at the Individual Level to Aid Counter-Trafficking Efforts

Nickolas Freeman, Greg Bott, Burcu Keskin, Jason Parton, James Cochran

To cite this article:

Nickolas Freeman, Greg Bott, Burcu Keskin, Jason Parton, James Cochran (2026) Linking Multisite Sex Ad Data at the Individual Level to Aid Counter-Trafficking Efforts. *Manufacturing & Service Operations Management* 28(1):133-152. <https://doi.org/10.1287/msom.2024.0816>

This work is licensed under a Creative Commons Attribution 4.0 International License. You are free to copy, distribute, transmit and adapt this work, but you must attribute this work as “*Manufacturing & Service Operations Research*.” Copyright © 2025 The Author(s). <https://doi.org/10.1287/msom.2024.0816>, used under a Creative Commons Attribution License: <https://creativecommons.org/licenses/by/4.0/>.”

Copyright © 2025 The Author(s)

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.






For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Linking Multisite Sex Ad Data at the Individual Level to Aid Counter-Trafficking Efforts

Nickolas Freeman,^{a,*} Greg Bott,^a Burcu Keskin,^a Jason Parton,^a James Cochran^a

^aDepartment of Information Systems, Statistics, and Management Science, The University of Alabama, Tuscaloosa, Alabama 35478

*Corresponding author

Contact: freem028@ua.edu,  <https://orcid.org/0000-0002-3036-7335> (NF); greg.bott@ua.edu,  <https://orcid.org/0000-0003-0928-9990> (GB); bkeskin@ua.edu,  <https://orcid.org/0000-0002-2273-1002> (BK); jmparton@ua.edu,  <https://orcid.org/0000-0002-6636-1090> (JP); jjcochran@ua.edu,  <https://orcid.org/0000-0002-8632-8909> (JC)

Received: February 1, 2024

Revised: May 1, 2025; August 13, 2025


Accepted: September 29, 2025

Published Online in Articles in Advance: November 17, 2025

<https://doi.org/10.1287/msom.2024.0816>

Copyright: © 2025 The Author(s)

Abstract. *Problem definition:* The internet facilitates sex trafficking through adult service websites (ASWs) that host online advertisements for sexual services (sex ads). Since the closure of the popular site Backpage.com, the ecosystem of ASWs has expanded to include multiple competing sites that are hosted outside U.S. jurisdiction. Gaining intelligence for counter-trafficking efforts requires collecting, linking, and cleaning the data from multiple sites. However, high ad volumes, disparate data types, and the existence of generic and misappropriated data make this process challenging. We present an end-to-end process for linking sex ad data and filtering potentially erroneous links. Outputs of the developed process have been used to inform counter-trafficking operations that have helped identify more than 60 potential victims of sex trafficking, some of whom are getting help to transition out of the life. *Methodology/results:* Our process leverages concepts and techniques from network science, information systems, and artificial intelligence to link ads across sites at the level of an individual or unique posting entity. Our approach is computationally efficient, allowing millions of ads to be processed in under an hour. A key component of our process is an edge-filtering procedure that identifies and removes potentially erroneous links in a graph representation of sex ad data. A comparison of the proposed process to an existing approach shows that our process is typically more computationally efficient and yields substantial increases in the number of individuals for which we can derive actionable intelligence. *Managerial implications:* The proposed process is an efficient and effective approach for transforming the high volumes of disparate data from sex ads into intelligence that can save lives. It has been refined over years of collaboration with practitioners and represents a strong foundation upon which further counter-trafficking tools can be built.

 **Open Access Statement:** This work is licensed under a Creative Commons Attribution 4.0 International License. You are free to copy, distribute, transmit and adapt this work, but you must attribute this work as “Manufacturing & Service Operations Research. Copyright © 2025 The Author(s). <https://doi.org/10.1287/msom.2024.0816>, used under a Creative Commons Attribution License: <https://creativecommons.org/licenses/by/4.0/>.”

Funding: This research is partially supported by the National Science Foundation [Grant D-ISBN-2240299].

Keywords: sex trafficking • network science • anomaly detection • artificial intelligence

1. Introduction

Sex trafficking refers to the use of force, fraud, or coercion to compel an individual to perform commercial sex services. The internet has and continues to facilitate sex trafficking in the United States via adult service websites (ASWs) that host commercial sexual advertisements (sex ads). Recent news references highlight the ongoing role of the internet in facilitating sex trafficking via ASWs:

- The U.S. Attorney’s Office, Middle District of North Carolina (U.S. Department of Justice 2023) describes the sentencing of a Durham man to six life terms in prison

for sex trafficking. The trafficker targeted women who were homeless or suffering from substance abuse problems. Five victims testified that they were photographed and advertised on two ASWs. They also testified that the sentenced individual and a codefendant arranged commercial sex encounters and that the trafficker used violence and coercion to maintain order.

- FOX6 News Milwaukee (2022) describes the sentencing of a Milwaukee man for sex trafficking of two minors in Chicago and Wisconsin. The trafficker advertised the victims on two ASWs and arranged encounters to occur in hotel rooms.

- Ellis (2023) describes a man pleading guilty to charges of sex trafficking of a minor. The trafficker posted ads for the victim on several ASWs, interacted with potential clients, and scheduled encounters between the trafficked individual and clients.

In recent years, government agencies have engaged in a game of whack-a-mole, shuttering ASWs linked to sex trafficking activities only to have new sites quickly emerge to fill the resulting void. Zeng et al. (2022) discuss the futility of these efforts, finding no causal impact that is attributable to the shutdown of two prominent sites in 2018. Instead, they find that users quickly moved to other sites hosted in locations outside of U.S. jurisdiction. The persistence of the online commercial sex ecosystem, demonstrated by anecdotal evidence and empirical analyses, suggests it is here to stay. Consequently, organizations dedicated to combating sex trafficking continue to need methods that transform the data available on these platforms into actionable intelligence.

Not all ads posted on ASWs are associated with victims of sex trafficking, and distinguishing advertisements featuring trafficking victims from those posted voluntarily or created as scams is a significant challenge. A large body of research is devoted to detecting sex trafficking victims based solely on the attributes of sex ads, that is, keywords, emojis, and image attributes. However, a critical issue is the absence of large-scale, verifiable, ground truth data. Even when law enforcement organizations (LEOs) can identify and isolate victims, establishing definitive evidence of trafficking requires significant time and resources, and victims often deny their victimization because of trauma or fear.

In addition to the challenge of identifying trafficking victims, the process of transforming sex ad data into actionable insights requires addressing several technical obstacles. First, sex ads use various types of data, including text, images, and numbers. Thus, data transformation pipelines must handle and link this multimodal data. Second, although ASWs typically convey the same core information, the ASW ecosystem includes many sites with different formats that provide data with slight variations from others. Third, the high volume of ads created daily exceeds the ability of most LEOs and nonprofit organizations (NPOs) to harvest, store, and transform this data into meaningful data products. Fourth, because providing and consuming commercial sex are illegal in most of the continental United States, scammers target ASWs extensively, adding significant noise to the data. Finally, even legitimate ads may contribute noise in the form of erroneous connections that arise from the use of generic or misappropriated data.

Since 2019, our research has focused on developing a systematic process for collecting and analyzing sex ads to support antitrafficking efforts. Our ad-linking

process (i) uses a graph representation to link the diverse data included in ads, (ii) filters a giant component that emerges during graph construction and includes erroneous links between a large proportion of the data, and (iii) transforms the filtered graph representation into actionable data products to inform the efforts of LEO and NPO partners. These data products provide quantitative analyses of advertisement frequency within specific geographic regions, such as cities or counties, along with associated ad URLs and phone numbers. Images are available on request, which may be blurred to avoid traumatizing NPO volunteers, some of whom are survivors of sexual exploitation. These data products have supported LEO and NPO operations, resulting in the identification of more than 60 potential sex trafficking victims. We refer to these individuals as “potential victims” because the common requirement of victim testimony makes convictions difficult. Given the emotional, physical, and psychological trauma associated with exploitation, it is often difficult to secure such testimony even if trafficking has occurred.

Our ad-linking process builds on that described in Freeman et al. (2022), addressing several limitations and incorporating significant enhancements. Specifically, our process (i) incorporates recent advances in language modeling for AWS data, (ii) leverages connected components from a graph representation of the data in a novel fashion to generate training examples for a same user classifier, and (iii) implements a novel approach for filtering erroneous data links without data loss. Moreover, we disclose as much detail as possible regarding our methods and computational setup to aid reproducibility. Using a large open ASW data set, we show that the proposed approach typically offers substantial improvements in terms of output quality and computational time.

The remainder of the paper is structured as follows. Section 2 describes related research on linking data from ASWs. Section 3 summarizes key concepts and techniques from graph theory/network science, image processing, and machine learning/artificial intelligence that we incorporate in our methodology. Section 4 provides a detailed description of the various components of our ad-linking process. Section 5 discusses our approach to detecting and removing erroneous links in the giant component that arises in the graph representation of the data and presents experimental results demonstrating its efficacy. Section 6 concludes the work, acknowledges limitations, and suggests directions for future research.

2. Related Work

This paper presents an end-to-end process to help identify victims of sex trafficking by linking sex ad data at

the level of an individual or unique posting entity (e.g., escort agency, massage parlor, pimp). A significant component of our approach is a procedure for filtering the giant component that arises in graph representations of the data because of generic posts, generic images, and misappropriated images in these data sets. In this section, we provide an overview of related works that use analytical methods to link and derive insights from sex ad data. We also discuss some related research on filtering giant components in graphs. We conclude the section by discussing specific contributions of this research.

2.1. Analytical Methods for Deriving Insights from Sex Ads

Li et al. (2018) present one of the earliest studies on linking sex ad data. The authors propose an unsupervised template matching algorithm that detects organizations operating on ASWs, in which the term “organization” refers to “a single individual or a network of individuals who are trafficking a specific set of victims.” Analogously, our ad-linking process extracts connected components of a graph representation of ad data corresponding to an individual or unique posting entity. Their matching algorithm generates embeddings for sex ad texts, clusters the embeddings using hierarchical density-based spatial clustering of applications with noise (HDBSCAN) (McInnes et al. 2017), and extracts word groups or templates that exhibit a high degree of collocation in the clustered documents. Applying their algorithm to a data set of nearly 40 million ads collected from the (now defunct) site Backpage.com demonstrates that it generates (i) linkages similar to those deduced from a graph representation based on posted phone numbers and (ii) additional linkages that cannot be constructed using only phone numbers. This research represents an important first step in showing how sex ad features can be leveraged to detect similarity.

Zhu et al. (2019) use machine learning to detect phrases indicative of sex trafficking in sex ad data. Specifically, they apply support vector machines on normalized ad text from the Trafficking10k data set (see Tong et al. (2017) for more information) to classify ads into various risk categories. Lee et al. (2021) describe INFOSHIELD, which detects patterns in sex ad text. The procedure uses term frequency–inverse document frequency (TF-IDF) to extract distinguishing keywords from posts and then constructs a bipartite graph that captures document–phrase relationships. The connected components of this graph represent a coarse clustering, which is refined in a fine clustering phase. The authors test the approach using open Twitter bot data (see Cresci et al. (2017) for more details) and the Trafficking10k data set.

Similar to the previously described research, we use a graph representation for the sex ad data. However, we use a transformer-based neural network model to capture and compare the semantic meanings of posts instead of TF-IDF. We also use perceptual hashes (pHashes) to incorporate images into our ad-linking procedure. Keskin et al. (2021) are the first researchers to demonstrate the use of perceptual hashing techniques for linking sex ads using text and images. They extract spatiotemporal patterns from linked ad groups and use them to simulate future movements. Freeman et al. (2022) extend the work of Keskin et al. (2021) by describing an end-to-end process for collecting and grouping sex data. They also use a graph representation of sex ad data and observe that generic text, generic images, and misappropriated images can create erroneous links, leading to the formation of large connected components in the graph. They propose a method for filtering these large components based on betweenness centrality but note that sampling techniques are required to overcome computational challenges. Although the authors are correct that erroneous links can form in the graph representation, they fail to recognize that this leads to a well-studied phenomenon in network science: the emergence of a giant component.

2.2. Filtering the Giant Component in Graphs

A significant issue we encounter when constructing the graph representation of ASW data is the presence of a giant component. In network science, a giant component is a connected component of a graph whose size scales proportional to the number of vertices in the graph (Sood and Yağan 2023). Although the existence of a giant component is often considered a useful feature of a network, there are cases in which decomposing it offers important insights into the network. Dianati (2016) discusses techniques and applications for such decomposition methods, which the author refers to as graph pruning methods. One use case, which also applies to illicit activities, is presented in Husain et al. (2020). Specifically, the authors use a graph representation to study spatiotemporal patterns in data related to terrorism. They use the Global Terrorism Database, which contains news articles on terrorism events from around the world, to construct a graph connecting terrorists and their targets. They find that the constructed graph representation includes a giant component that links approximately 92% of all vertices. They use a disparity filtering method to filter edges and extract the backbone of the graph. Similarly, Ran et al. (2013) observe a giant component in a graph representation of proteins associated with essential hypertension. They use graph centrality measures to extract the graph’s backbone, and this provides insights into important regulatory pathways for controlling the condition.

Another concept closely related to filtering edges in a giant component of a graph is known as k -core decomposition; k -core decomposition partitions a graph into hierarchical layers based on vertex connectivity, in which the core number of each vertex indicates the largest value of k for which the vertex belongs to a k -core (Batagelj and Zaversnik 2003). This concept is used extensively in the analysis of social networks. Notably, Kitsak et al. (2010) use an analogous method, which they refer to as k -shell decomposition, to identify the most efficient vertices for information spreading in a network. Interestingly, they show that the best spreaders do not always correspond to the most highly connected or the most central vertices.

A key difference between the methods discussed in this section and our approach for filtering the giant component is the features used to inform the filtering process. Specifically, most existing research relies on features of the graph, such as centrality metrics. In our case, the vertices correspond to data collected from ASW ads. Thus, instead of relying on structural features of the graph, we leverage the graph to identify data points to compare, which, in our case, are pieces of text, and then leverage custom artificial intelligence models to decide whether to remove the edge connecting the data points in the graph representation.

2.3. Contribution

The closest research to ours is Freeman et al. (2022). Although their goal is similar, there are several limitations with their proposed pipeline. First, the authors fail to recognize the emergence of a giant component in the graph representation of ASW data, which leads them to devise filtering methods that apply to other components unnecessarily. Second, the procedure proposed for identifying and filtering potentially erroneous links in Freeman et al. (2022) (i) does not leverage the data associated with vertices, (ii) can lead to data loss, and (iii) requires approximation to overcome computational challenges. Specifically, the authors use estimates of betweenness centrality (BC) to identify and remove vertices that exceed a selected percentile threshold in the graph representation. BC computation is very expensive. Thus, the authors are required to approximate the BC values for all vertices in a particular component based on shortest path computations among a subset of nodes. Regarding data loss, filtering graph vertices based on extreme BC percentile should result in the removal of only a few vertices with high connectivity. However, as we show in our experimentation, different site structures or approximation issues can lead to much more substantial data loss. A final limitation of Freeman et al. (2022) is that it lacks several details that are necessary for replication.

Our ad-linking process makes several contributions and addresses limitations of the process presented in

Freeman et al. (2022). First, we identify the emergence of a giant component, connecting the challenges of working with sex ad data to a broader phenomenon in the network science literature. Second, our approach to filtering potentially erroneous edges in the giant component is data-driven, leveraging recent advances in language modeling for ASW data to identify and remove such connections based on classifying the similarity of ad text pairs. Unlike prior methods, the present filtering procedure does not result in data loss. A major challenge of training such a classifier is the lack of ground truth. We present a novel approach to generate reliable training data for the task in a dynamic fashion that is robust to changes in the underlying data used to construct the graph representation. This approach is not specific to our application and can be applied in other contexts in which filtering the giant component provides insights into its underlying structure, for example, social network analysis. Finally, this research takes major steps toward reproducibility, including detailed descriptions for all steps and leveraging a large and open ASW data set (Freeman et al. 2025a) in experiments that demonstrate its efficacy in terms of output quality and computational time. Ultimately, the outlined process opens the door for new research directions by providing a thorough, rigorous, and well-documented procedure for extracting data corresponding to posting entities represented in ASW data.

3. Background

This section briefly describes concepts and techniques we employ from graph theory/network science, image processing, and machine learning/artificial intelligence. Our intent is not to provide an overarching review of any area, concept, or technique but to highlight aspects that aid in understanding our description of the developed ad-linking process.

3.1. Graph Theory and Network Science Fundamentals

Network science refers to the application of mathematical graph theory to problems across various disciplines, in which graph theory focuses on mathematical approaches for modeling relationships between objects (Lewis 2011). A graph $G(V, E)$ is defined by a set of vertices V that are connected by edges that form a set E . These edges can be directed or undirected, depending on the context. For example, vertices in a family tree graph may represent individuals, and directed edges can indicate who is a descendant of whom. Edges can also have weights that represent meaningful attributes of the relationship between the two connected vertices, such as distance or capacity. A path, P_{uv} , refers to an ordered collection of vertices connecting an origin or source vertex $u \in V$ to a destination or target vertex

$v \in V$. The path's length is the number of edges the path includes. A vertex is said to be reachable from another if a path connecting the two vertices exists (Börner et al. 2007).

Our ad-linking process represents sex ad data as a graph and extracts groups of vertices that are reachable from one another. The groups are called “components” in graph theory and network science terminology. A common phenomenon in the study of real-world graphs is the emergence of a giant component, which is a component whose size scales with the number of vertices in the graph (Janson et al. 1993). When a giant component exists, a significant proportion of the graph vertices belong to this component (Börner et al. 2007). In our application, the presence of generic and misappropriated data leads to the emergence of the giant component. Our ad-linking process includes an edge-filtering procedure that specifically seeks to identify and filter such erroneous edges in the giant component. Filtering the giant component allows us to decompose it into additional components, each likely to correspond to an individual or unique posting entity for which we can derive actionable intelligence from.

3.2. Hashing Techniques: Cryptographic and Perceptual

Sex ads include text, contact information (often represented as strings), and images. Individuals or posting entities tend to reuse each of these data elements. Thus, we need efficient ways to represent and match them across ads and ASWs. Achieving this is most complicated for images.

In computer science, a hashing function maps data records to a lower dimensional space such that two pieces of identical data generate the same hash value. Ideally, these functions create simple encodings that can verify more complex data or programs. Various hashing functions have been developed for different applications (Chi and Zhu 2017) with cybersecurity being one of the most well-known applications.

Cryptographic hashing refers to security-oriented methods designed to be one way and difficult or infeasible to reverse. They are also designed to ensure that small content changes result in different hash values. A common application of cryptographic hashes is verifying downloaded software. Software providers compute and publish hash values using specific algorithms, such as Message Digest (e.g., MD5) or Secure Hash Algorithm (e.g., SHA-256). Users who download the software can apply the same algorithm to verify that their hash matches the one published with different values potentially indicating corrupt software.

Perceptual hashing techniques have a similar goal of representing complex data in a simpler format but are designed to work with images. Specifically, digital images are represented as arrays of pixel values, and

perceptual hashing techniques determine whether two images are fundamentally the same based on their content rather than making strict pixel-to-pixel comparisons. This is crucial because slight image variations can arise from cropping, applying filters, compression, or other modifications. Perceptual hashing techniques are commonly used to help identify instances of copyright infringement or online child sexual abuse materials. We refer readers interested in learning more about perceptual hashing techniques and their applications to Farid (2021).

3.3. Transformer Models and Sentence Embeddings

As noted earlier, our graph representation for sex ad data is prone to the emergence of a giant component that arises from generic and misappropriated data. Our ad-linking process includes a procedure to identify and remove erroneous edges within the giant component. This filtering procedure leverages transformer neural network models tailored for ASW applications for these comparisons.

The transformer neural network architecture, introduced by Vaswani et al. (2017), utilizes an attention mechanism that allows the model to focus on important contextual elements during training. Before transformers, recurrent neural networks (RNNs) were the dominant architecture for text processing. RNNs process text sequences serially, making them slow and computationally challenging for large context windows. The transformer architecture and its attention mechanism remove significant barriers, enabling rapid advances in neural network applications for natural language and other areas (Incitti et al. 2023).

One area in which transformers are successfully applied is developing sentence embedding and classification models. These models transform sentences into numerical vector representations that can be compared with identify sentences with similar contexts. The open-source community has emerged as a primary host of pretrained transformer models. Platforms such as Hugging Face host repositories of models and publish libraries that enable easy Python integration. On Hugging Face, contributors ranging from large technology companies to individuals share models and documentation that allow users to use pretrained models directly or fine-tune them for specific applications.

Several key concepts are important when applying transformer models to natural language problems:

- Foundation models: The term “foundation model” refers to a model trained on broad, diverse data sets to provide a foundation for more specialized models (Bommasani et al. 2022). Some well-known foundation model families are BERT (Devlin et al. 2019), RoBERTa (Liu et al. 2019), GPT (Radford et al. 2018), and LLaMa (Touvron et al. 2023). Foundation models learn vector

representations for tokens, which are basic units of text that can correspond to words, subwords, characters, or bytes. Different models employ different tokenization approaches. For example, BERT uses WordPiece tokenization, whereas GPT uses Byte-Pair Encoding tokenization. It is important to note that the vocabulary of tokens that a particular language model utilizes is fixed. If a foundation model encounters a token that is not in its vocabulary, it is mapped to special tokens that represent it as unknown, for example, the [UNK] token in BERT. Understanding how unknown tokens are handled is crucial when considering ASW data, which uses emojis and slang that may not be part of the general-purpose vocabulary of mainstream foundation models.

- **Transformer components:** The transformer consists of encoder and decoder components (Vaswani et al. 2017). The encoder learns how to represent input text as vector embeddings that capture context from other tokens in an input sequence. The decoder learns to generate tokens autoregressively based on preceding tokens. Foundation models differ in the transformer components used. For example, BERT is an encoder-only model that can generate contextualized token embeddings based on all other tokens in an input sequence (Qorib et al. 2024). BERT variants are commonly fine-tuned for tasks that require understanding an entire text sequence, such as text classification and named entity recognition. GPT variants are decoder-only models that are typically fine-tuned for tasks that rely on text generation. T5 utilizes both the encoder and decoder components and is commonly fine-tuned for tasks in which text understanding and generation are important, such as question answering and text summarization.

- **Fine-tuning:** Foundation models are designed to exhibit broad text understanding for assessing text similarity or generating text. Using foundation models for a specific task, such as classification or named entity recognition, requires fine-tuning on a labeled data set (Ding et al. 2023). Some tasks require an additional step of training a model that learns to aggregate token representations for a text sequence into a single embedding that can then be used to generate features for a classifier or clustering algorithm. An example of such a model is a sentence transformer model (Reimers and Gurevych 2019).

Our research utilizes a sentence transformer model specifically fine-tuned for computing similarities between post texts found in sex ads. Freeman et al. (2025b) describes the training process, which includes custom pretraining of a BERT variant on a data set of more than 19 million ASW posts and fine-tuning a sentence transformer on a data set of four million ASW post triplets using a contrastive learning task. We use this fine-tuned sentence transformer to generate

numerical vector representations for post texts. We compute the similarity between post text pairs using cosine similarity, a standard method for comparing sentence embeddings (Reimers and Gurevych 2019). We employ a simple binary classification model to translate the similarity scores into predictions about whether the same individual created both posts.

4. Ad-Linking Process

This section describes our process for linking ads at the level of an individual or unique posting entity. The process performs five distinct tasks: (i) data collection, (ii) data cleaning and standardization, (iii) location cleaning and standardization, (iv) ad grouping, and (v) output generation. Figure 1 provides a flowchart depicting the complete process. Each task is described in detail in the following sections.

4.1. Data Collection

We use custom web-scraping applications to collect data from ASWs. Because the ecosystem of ASWs is constantly evolving and the structure and format of each site differ, we develop specific data-collection applications for each one. We add new applications for emerging sites, discontinue applications for inactive sites, and modify existing applications as sites change.

Each application collects all text and image information for newly posted ads. We use unique identifiers, such as ad IDs or unique combinations of extracted features (such as text and post date), to ensure ads are only collected once. We also utilize caching to ensure that we only collect each unique image once. All text information is stored in a JSON format, and we use cryptographic and perceptual hashing techniques to represent associated images in the JSON data. Copies of each unique image encountered are downloaded to a dedicated image store.

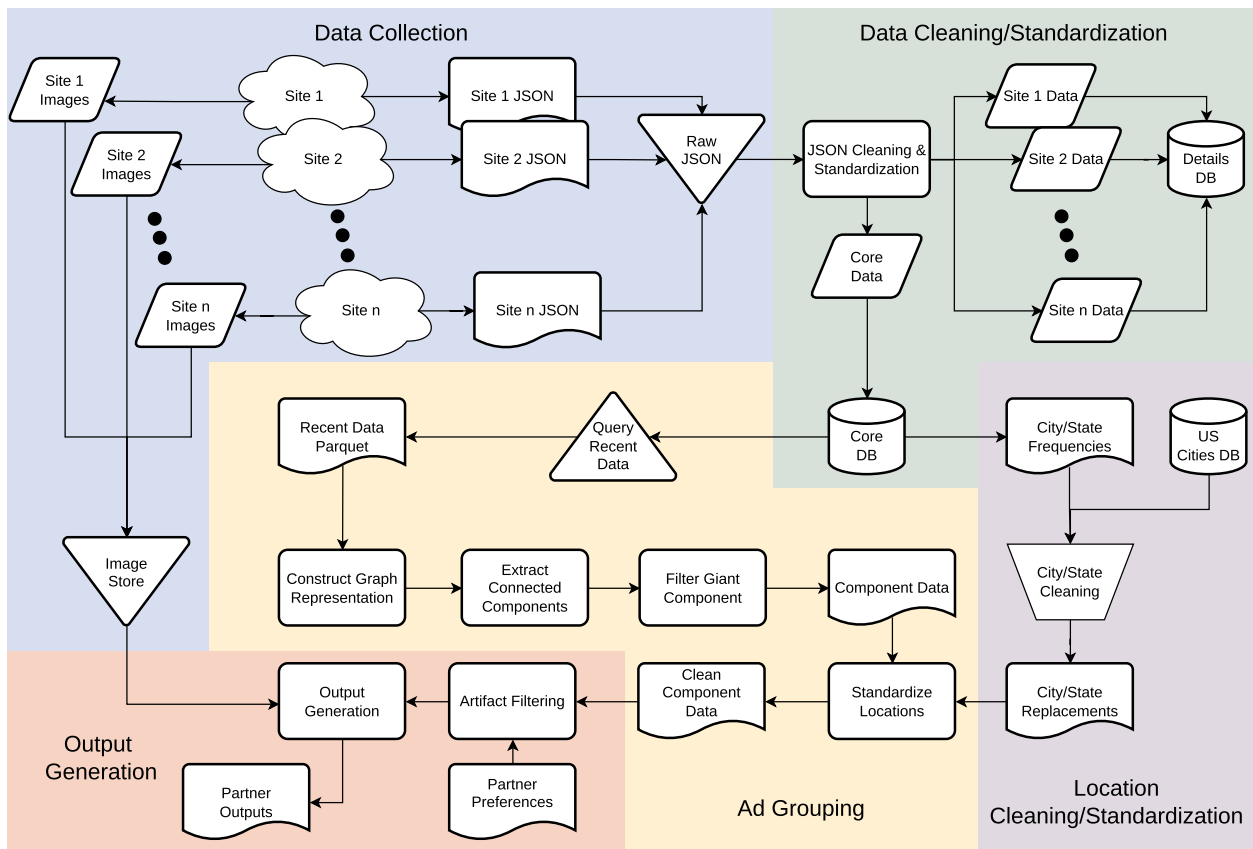
4.2. Data Cleaning and Standardization

The data cleaning and standardization task serves two purposes: (i) to extract a consistent subset of information from all sites for ad linking and (ii) to map all data to a set of standardized labels, allowing comparisons between sites. We refer to the data extracted in (i) as core data, which is stored in a relational database (core DB). The data mapped in (ii) is referred to as detail data and stored in a separate relational database (details DB).

The core data includes the following data fields for each ad, which are collected from all sites:

- Unique ID: unique identifier for ad.
- Post: heading text displayed with ad.
- Site: the ASW from which the ad was collected.
- Post date: the date the ad was posted or first observed.

Figure 1. (Color online) Ad-Linking Pipeline



- Target city: the city the ad targeted.
- Target state: the state the ad targeted.
- Location details: any freeform text providing additional information regarding the location.
 - Phash-16 values: a list of 16-bit pHash values for images posted with the ad.
 - Contact information: contact information provided with the ad.

The detail data varies significantly among sites and includes additional information about the provider, such as hair and eye color, rates, touring schedules, and even client reviews.

4.3. Location Cleaning and Standardization

The next phase of our ad-linking process focuses on cleaning and standardizing cities and states represented in the core data. Most ASWs allow posters to select from a limited set of cities and states when specifying the target location for their ad. This creates two significant challenges. First, the set of locations varies among sites. Second, the available locations may differ in their granularity and spelling of location names. Examples of these issues are one site allowing ad creators to select “Dallas, Texas” or “Fort Worth, Texas” whereas another only allows “Dallas/Fort Worth, Texas.” Moreover, “Fort Worth” can be represented as

“Fort Worth,” “Ft Worth,” or “Ft. Worth,” each interpreted as a different string by a computer.

To address these issues, our process maps the city/state pairs associated with ads to locations specified in version 1.6 of the Pro U.S. cities database made available by Simplemaps (2022). We begin by first querying the core DB to retrieve a data set showing the number of URLs collected for each unique city–state pair. We then combine the city and state strings and attempt to match the combined string to an entry in the U.S. cities database. In cases in which we successfully match the city–state string to an entry in the U.S. cities database, we join the county name and Federal Information Processing Standards (FIPS) code from the U.S. cities database to the ad data. FIPS codes are standardized numeric or alphanumeric codes issued by the U.S. government to uniquely identify geographic entities, such as states, counties, cities, and other administrative divisions. For cases in which no match is made, we sort the data by the number of affected URLs and manually specify appropriate matches, prioritizing locations that affect the most URLs.

Our matching process is imperfect because we cannot specify appropriate matches for all cases. However, our efforts reduce the number of URLs affected by unmatched city–state pairs from approximately 13.5%

to less than 0.5%. Figure 2 provides additional details on the impact of our manual location cleaning and standardization process. The left subplot shows the top 20 unmatched locations as a percentage of affected URLs before cleaning and standardization, whereas the right subplot shows the top 20 unmatched locations after cleaning and standardization. The output of the location cleaning and standardization process is a data set that maps the observed city–state strings to cleaned representations with county details.

4.4. Ad Grouping

Ad grouping begins by querying recent core data, which is saved to disk as a Parquet file. Parquet is an open-source file format designed to handle flat, columnar data. It is known for efficient data compression and its ability to handle various encoding types. Typically, we work with core data corresponding to ads posted or observed within 90 days of execution. This time frame is based on feedback from our partners and can be adjusted as needed.

After querying, we construct an unweighted and undirected graph representation of the data in which vertices correspond to unique artifacts, such as post text or pHash values. Graph edges identify links between artifacts that appear together in an ad. A typical ad consists of a single post and multiple images and can contain no contact information (some sites include a “message me” capability), one piece of contact information (e.g., a phone number), or multiple pieces of contact information. When constructing the graph representation, we represent each ad using the post text value as the central vertex connected to various pHash and contact information vertices. When multiple ads use the same artifact, links among ads create a larger graph structure.

Figure 3 shows a graph representation for a simple example with four ads in which ads 1–3 are linked because they share common artifacts. Specifically, ads 1 and 2 share a common contact information vertex, that is, the two ads use the same phone number. Also, ads 2 and 3 share two images. Ellipses are used to identify the shared artifacts. Because of the described overlaps, the artifacts associated with ads 1–3 form a single connected component. Ad 4 has no artifacts that overlap with other ads; thus, the artifacts associated with ad 4 form another connected component. Note that we also know the post date and target location for all ads. When an individual posts the same ad multiple times, that is, with identical post text, contact information, and images, we also have records of all post dates and target locations in the data extracted from the core DB.

Once the graph is constructed, we extract all connected components, each corresponding to a set of artifacts that can be linked via their usage across all ads. A significant issue we encounter at this point is the emergence of a giant component, which can largely be attributed to (i) generic post texts and images (e.g., “available now” or the OnlyFans logo) and (ii) misappropriated images in which a provider creates an ad using an image found elsewhere. The giant component poses a significant problem as it links artifacts associated with many providers. Without filtering the edges in the giant component and eliminating erroneous connections, the data it links is too noisy for actionable insights and must be disregarded.

Our ad-linking process includes a giant component filtering procedure that leverages a prediction model to detect and remove likely erroneous edges from the giant component. Section 5 provides a detailed explanation of the procedure. For now, we provide a high-level overview of how the procedure breaks down the

Figure 2. (Color online) County Mapping Examples


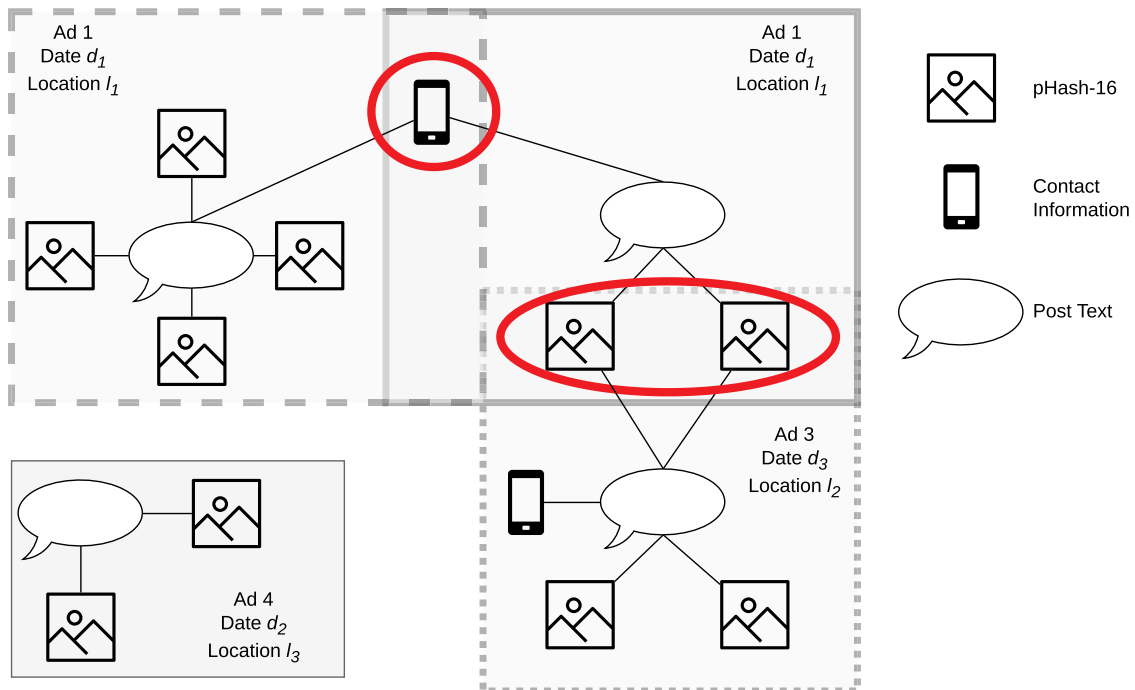


Figure 3. (Color online) Graph Representation Example



Note. The Post Text vertices correspond to the raw post text harvested from an ad.

giant component into smaller components likely to be associated with an individual or a unique posting entity. The first step is to extract the giant component from the graph representation of the data and project it so that it only includes vertices associated with posts. Edges in the projection exist between pairs of vertices corresponding to post text values that were connected by pHash or contact information vertices in the original graph representation. Figure 4 depicts this graph projection step in which images correspond to pHash values and phones correspond to contact information values.

Next, we iterate over all edges in the graph projection and use a binary classification model to assess the probability that the two posts connected by an edge were created by the same user or posting entity. We refer to this probability as the same user probability. A user-specified threshold is used to identify edges to remove from the graph projection. Specifically, if the same user probability value for two linked posts is less than the threshold, the edge is flagged for removal. Flagged edges are removed after iterating over all edges in the projection, leaving components corresponding to connected posts. Figure 5 illustrates how this procedure filters the graph projection from Figure 4 into two distinct components of connected posts when a threshold of 0.5 is applied.

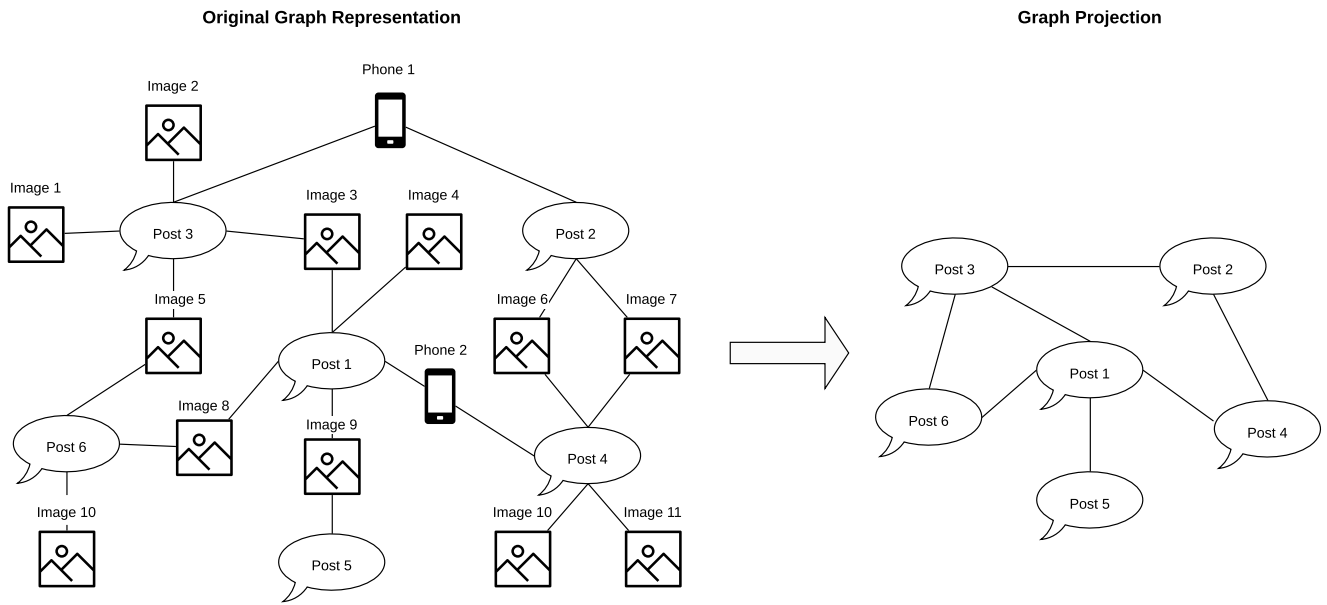
Components extracted from the filtered projection of the giant component allow us to map each post text value associated with a vertex in the giant component

to a unique component identifier. Additionally, each component not included in the giant component can also be mapped to unique component identifiers. These unique identifiers are mapped back to the original data along with county names and FIPS codes from the location cleaning/standardization phase. These data are written to a clean component data file for use in the next phase. A key thing to note about the described filtering procedure is that it results in no data loss. Instead, we are simply removing associations between post text values.

4.5. Output Generation

The final phase of our ad-linking process generates outputs for our LEO and NPO partners engaged in counter-trafficking operations. We generate two types of outputs that are tailored to the region where our partners operate. The first, summary output, provides a tabular listing of individuals or unique posting entities (corresponding to connected components in the graph representation) that have targeted the selected region in recent days. This output contains basic statistics, such as the number of observed ads, the proportion of times the individual targeted the chosen region, and recently used contact information for each individual. This allows users to quickly browse the list and attempt to make contact with individuals without visiting the sites hosting ads, minimizing exposure to explicit content. This is an important consideration for our NPO partners who rely on volunteers for outreach efforts,

Figure 4. Giant Component Filtering: Graph Projection



including volunteers who are former victims of sex trafficking. The second, detail output, provides additional details for each individual in the summary output, including recently posted images and post texts, recent URLs, and a choropleth plot showing recent movement across states by month. Our process allows specific details of the outputs to be specified on a per-partner

basis, including the blur applied to images, the locations considered, and a minimum target ad frequency for the selected locations. Figure 6 provides an example of both outputs and shows how the summary and detail outputs are linked by the unique identifier associated with connected components of the graph representation.

Figure 5. (Color online) Giant Component Filtering: Edge Removal

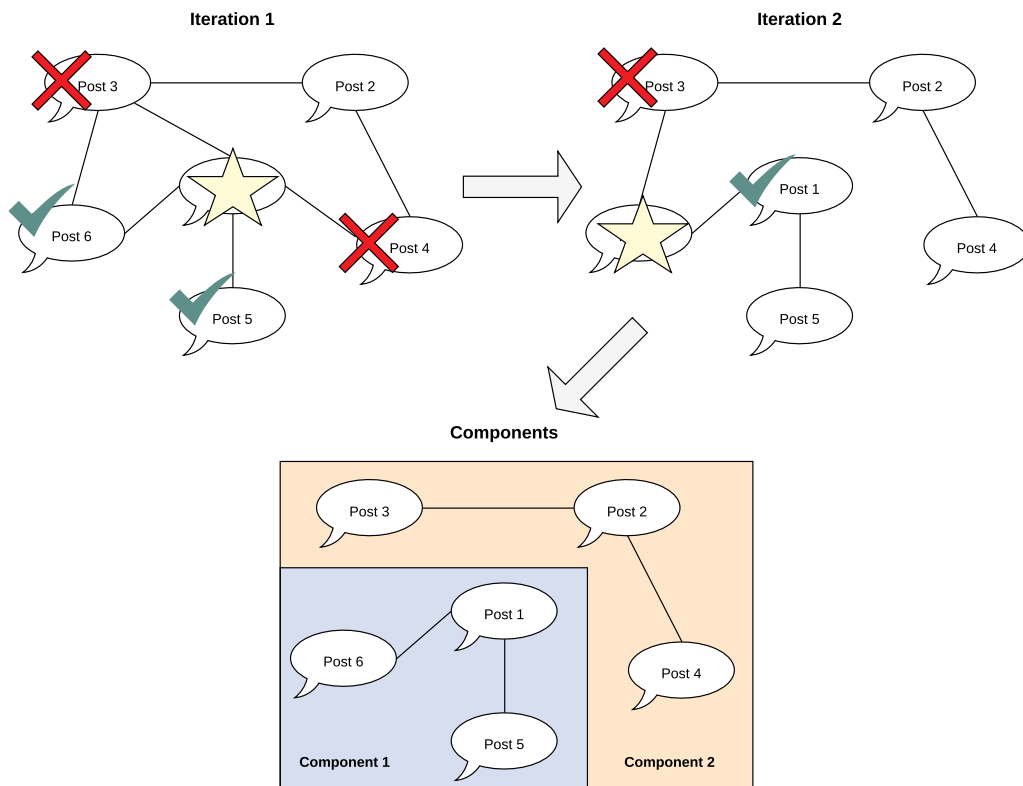
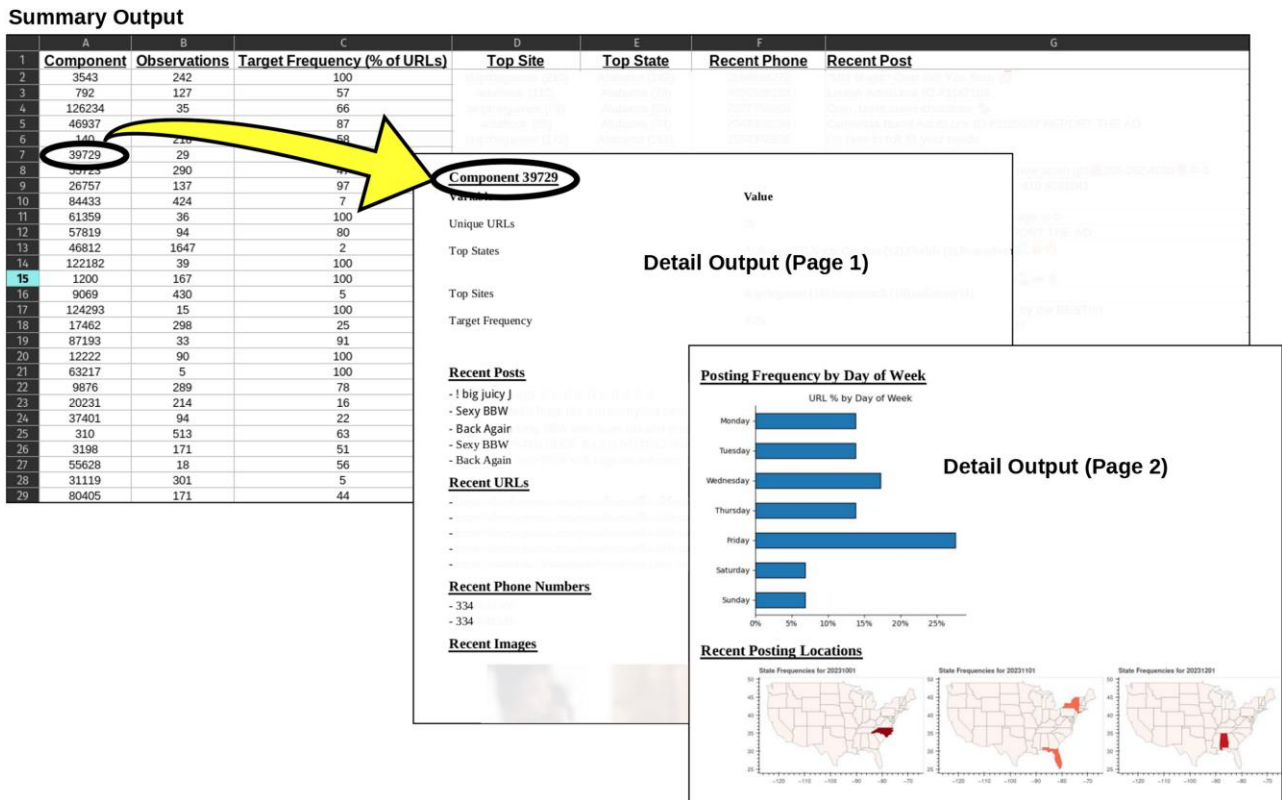


Figure 6. (Color online) Outputs Example



We do not discuss the specifics of how our partners operationalize these data products because they vary greatly depending on the type of organization (LEO versus NPO) and the resources available. A common step we take when working with a partner is filtering the data products only to include information associated with components that have been observed in areas in or near the partner’s jurisdiction in recent days. For example, if we were working with an organization in Dallas, Texas, we might filter the data only to include components observed in Dallas, Texas, or Fort Worth, Texas, over the last 14 days. Moreover, our output generation approach can be customized to filter the data based on additional criteria specified by the partner.

The key observation is that our data products distill data collected from large volumes of ads posted on one or more sites into interpretable information that can be used directly to initiate contact with a potential victim. Without these products, organization employees must try to establish links between ads by visiting sites and manually linking the information. Just the process of linking ads, excluding attempts to eliminate erroneous connections, is daunting. For example, if we observe 100 ads targeting a specific location, does that correspond to 100 individuals each posting a single ad, one individual posting 100 times, or some combination in between?

4.6. Technologies

We implement the ad-linking process using the Python programming language. Table 1 specifies essential libraries, the tasks performed by each library, and any available citations. We use the following abbreviations for the tasks: DC = data collection, DCS = data cleaning and standardization, LCS = location cleaning and standardization, AG = ad grouping, and OG = output generation.

5. Giant Component Filtering

This section provides additional details regarding our procedure for detecting and removing potentially erroneous links in the graph representation of the giant component. As described in Section 4.4, the procedure iterates over all edges in a projection of the giant component’s graph representation, in which the projection includes edges between posts that were initially connected by vertices corresponding to pHash or contact information values. During each iteration, a binary classification model assesses the probability that the same user created the posts associated with the connected vertices. Algorithm 1 provides pseudocode for the filtering procedure.

Algorithm 1 (Giant Component Filtering)

- Define same user probability threshold, δ .
- Extract connected components from graph representation of ASW data.

Table 1. Python Libraries Used by Task

Library	DC	DCS	LCS	AG	OG	Citation(s)
beautifulsoup	X					
bokeh					X	
igraph				X		Csardi and Nepusz (2006)
matplotlib					X	Hunter (2007)
numpy	X	X	X	X	X	Harris et al. (2020)
pandas		X	X	X	X	Pandas Development Team (2020), McKinney (2010)
pillow	X				X	
polars		X	X	X	X	
python-docx					X	
requests	X					
scikit-learn				X		Pedregosa et al. (2011)
scipy				X		Virtanen et al. (2020)
seaborn					X	Waskom (2021)
selenium	X					
sentence-transformers				X		Wolf et al. (2020)

- Use sentence transformer (see Section 5.3) and non-giant components to train a binary classifier for the same user classification task (see Section 5.4).

- Define GC_graph as the giant component from the original graph representation.

- Create post projection of GC_graph and store as GC_post_graph . Let E denote the set of edges in GC_post_graph

- Define $SU(x, y)$ as a procedure that computes the same user probability for a set of texts, x and y , that is, our binary classification model.

- Define an empty list, $edges_to_remove$.

for $e \in E$ **do**

- Let $source$ and $target$ represent the vertices connected by edge e in GC_post_graph .

- Let \tilde{s} represent the post text associated with the source vertex.

- Let \tilde{t} represent the post text associated with the target vertex.

if $SU(\tilde{s}, \tilde{t}) < \delta$ **then**

Add edge e to $edges_to_remove$.

end if

end for

- Remove edges in $edges_to_remove$ from GC_post_graph .

- Extract connected components from GC_post_graph . Each connected component contains a set of vertices corresponding to posts that are reachable to one another in GC_post_graph .

- Create a unique identifier for all extracted components and use these identifiers to label posts in the data.

The remainder of this section focuses on the developed classification procedure, including how it is trained and deployed. We also present experimental results demonstrating the efficacy of the edge removal procedure on an open data set that includes data from multiple ASWs. We begin with a description of the

ASW data set used throughout this section and a brief overview of the computational setup.

5.1. Data

We utilize the open data set described in Freeman et al. (2025a) to demonstrate relevant aspects of our giant component filtering approach. The data set provides a sample of more than 10 million ads collected from nine ASWs between May 1, 2022, and August 1, 2022. The authors obfuscate numerical data included in post texts, anonymize URLs and site names, and provide pHash values instead of images. These measures prevent the direct identification of individuals, preserving as much data structure as possible. Table 2 lists the specific data fields we use and describes each field exactly as provided in the original paper. We refer readers seeking an in-depth discussion of the data collection methodology and structure to the original publication.

Table 2. Relevant Data Columns

Data field	Description
url	An integer representing the URL associated with the ad.
site	An integer representing the site from which the data were collected.
post_masked	The raw heading text associated with the ad. Note: Numbers that appear in the text are obfuscated by replacing them with the character “*”.
post_int	An integer representing the raw heading text associated with the ad. Note: All occurrences of a specific text are assigned to the same integer value.
phone_int	An integer representing the phone number associated with the ad. Missing values are represented as None.
phash16	The pHash for a single image associated with the ad, computed using the imagehash library.

Table 3 summarizes the data for each of the nine sites included in the data set. We also include an overall summary (“All”) corresponding to the case in which all data are used. Even though the data were collected from May 1, 2022, to August 1, 2022, we note substantial differences in the number of URLs and unique post, pHash, and phone number values across the sites. Moreover, the summary statistics indicate substantial variation across sites concerning data duplication. For example, notice that, although site 1 has less than 10% of the number of URLs as site 9, site 1 has more unique post, pHash, and phone number values. Ultimately, the summary statistics highlight the degree of variation in data available from ASWs.

5.2. Computational Setup

The experimental results that follow show that our ad-linking process is computationally efficient and can handle practical data volumes in a time that allows for daily updates. However, processing large volumes of ad data and running neural network inference on hundreds of thousands of posts requires high-performance computational infrastructure. We run the ad-linking pipeline on a Dell Precision 7865 workstation with a 64-core AMD Ryzen Threadripper CPU, 512 GB of RAM, and an NVIDIA RTX 6000 Ada GPU with 48 GB of memory. We use the *Pop!_OS* Linux distribution and the *Pixi* package manager to build and manage reproducible Python environments with GPU support. All experimentation presented in this section was run on the described workstation.

5.3. Language Modeling

We use a binary classification model to filter edges in the giant component. Specifically, the binary classifier is trained to predict whether a pair of posts can be attributed to the same individual or posting entity. Constructing an effective classifier requires a high-quality means of assessing the similarity of ASW post texts. We leverage a sentence transformer model designed explicitly for ASW data that is proposed in Freeman et al. (2025b) for our classification task.

Table 3. Data Set Summary

Site	URLs	Unique posts	Unique pHashes	Unique phones
Site 1	461,234	283,845	1,986,408	143,226
Site 2	257,252	105,400	625,396	257,250
Site 3	1,794,625	263,084	978,388	265,552
Site 4	121,199	13,909	148,812	6,383
Site 5	27,818	16,049	65,533	5,139
Site 6	1,592,738	215,664	1,371,934	154,126
Site 7	1,070,044	320,347	1,117,073	305,370
Site 8	169,114	146,022	2,000,931	3,688
Site 9	4,784,510	273,264	501,156	17,586
All	10,278,534	1,036,868	6,219,151	497,104

Specifically, we use the sentence transformer fine-tuned from the BERT-New30522-MLM variant trained for 20 epochs.

Figure 7 demonstrates the benefit of using the ASW-tailored models in terms of tokenization. The figure shows five example posts from the data set described in Section 5.1 and how they would be tokenized by the base variant of BERT (BERT-base) and the ASW-specific model from Freeman et al. (2025b) (custom). The key observations are that the BERT-base tokenization approach (i) ignores emojis altogether and (ii) is more likely to be missing tokens in its vocabulary, which results in more occurrences of the [UNK] token.

The previous figure demonstrates that the custom foundation models introduced in Freeman et al. (2025b) are much better at retaining information from ASW post text values when tokenizing. We now demonstrate the sentence transformer model’s ability to capture the semantic similarity of posts. Specifically, we use the trained sentence transformer model to generate embeddings for each unique post text value in our data set. We then use the recently developed LocalMAP dimensionality reduction technique (Wang et al. 2024) to reduce the 512-dimensional embeddings for each post text to a two-dimensional vector. We use the default values for LocalMAP when performing the dimensionality reduction. We then use HDBSCAN (McInnes et al. 2017) to cluster the reduced-dimensional representations. Figure 8 shows examples of posts associated with three resulting clusters, demonstrating the model’s ability to generate embeddings that map similar texts nearby in the vector space.

5.4. Binary Classification

This section describes our approach to classifying whether pairs of post texts can be attributed to the same individual or posting entity. The classification task involves (i) using the sentence transformer model described in the previous section to generate embeddings for pairs of connected posts in the projected giant component representation, (ii) computing the cosine similarity of the embeddings for linked posts to measure their proximity to the embedding space, and (iii) using a simple logistic regression model to learn the classification boundary that best distinguishes posts made by the same individual or entity.

The key challenge in training such a classifier is obtaining training data. Even though some sites may provide a profile identifier with ads, it is not common, and there is no way to guarantee that an individual or posting entity does not have multiple profiles. In short, ASWs generally do not provide a means to obtain ground truth labels for individuals posting ads. We overcome this challenge by assuming that nongiant components in the initial graph representation, that is, all components that are not the largest component,

Figure 9. (Color online) Training Data Example

Post 1 str	Post 2 str	Same User i64	Cosine Similarity f32	Same User Probability f64
Hot AF content here!	Sneak away and CUM play 🍆🍆	0	-0.072021194	0.010558419947326555
🐾 ocean cat wet nd ready now	lil thick blonde 😏	0	-0.06873483	0.010893373372365713
READ FULL AD!!!! \$*** QUICKIES	Quick visit u down?	1	0.2210348	0.15122196704521143
🍆 THICKALICIOUS 🍆 OUTCALLS AVAILABLE	🍆 THICKALICIOUS 🍆 Oh So Delicious 😏😏😏	1	0.48529267	0.6928500194110849
🍆 Miss Kay 🍆	Incalls and outcalls available	0	0.051005814	0.033621517490859856
Sweet and sensual Girl Next Door	East Evansville incall 🍆🍆🍆	0	0.004015783	0.02167277936148269
Cumm n see ABQS MOST NAUGHTIEST BABE	THICKER than a SNICKER 🍆🍆🍆🍆	0	0.17620373	0.10379910435252422

specifying a same user probability threshold that removes edges between connected posts whose same user probability is less than the specified threshold value. Figure 9 shows an example of the training data set along with computed cosine similarity and same user probability values.

When generating the training data set, we typically create 10,000 training examples with 1,000 positive instances and 9,000 negative instances. Figure 10 provides additional insight regarding how the selection of the same user probability threshold affects the rate of false positives and false negatives. The figure includes subplots that depict the trade-offs between the false positive rate (FP), the false negative rate (FN), and the total error rate (total) as the threshold for the same user probability varies from 0.0 to 1.0 along the horizontal axes. The left subplot shows the trade-offs for site 5 (the smallest site by unique URLs), and the right subplot shows the trade-offs for site 9 (i.e., the largest site by unique URLs).

Note in Figure 10 that the errors are biased toward false negatives for probability thresholds greater than 0.25. In the context of our prediction task, a false negative (positive) corresponds to incorrectly leaving


(removing) an edge in the giant component projection that should be removed (retained). Regardless, the same user probability threshold provides an easy-to-interpret mechanism that allows for tailoring potential errors made by the classification model based on user preferences. In practice, we use a threshold value of 0.7, which reduces false positives but increases the number of false negatives. However, the total error is comparable to that achieved with a threshold of 0.5, and the threshold can be easily adjusted based on partner preference.

5.5. Impacts on Giant Component Size

This section demonstrates the efficacy of our giant component filtering approach. We run the procedure for each site included in our data set independently as well as for all sites combined (All). We use a same user probability threshold of 0.7 when determining edges to remove in the giant component projection. Table 4 provides relevant statistics for this experiment, including

- $|V|$: the number of vertices in the original graph representation of the data.
- $|E|$: the number of edges in the original graph representation of the data.

Figure 10. (Color online) Prediction Threshold Examples



- $|C|$: the number of components in the original graph representation of the data, including the giant component.
- GC proportion: the size of the giant component as a proportion of the number of vertices ($|V|$).
- Filtered $|C|$: the number of components after applying the giant component filtering procedure.
- Filtered GC proportion: the size of the giant component as a proportion of the number of vertices ($|V|$) after applying the giant component filtering procedure.
- $|C|$ increase: the increase in the number of components because of the application of the giant component filtering procedure, expressed as a proportion of the number of components in the original graph representation ($|C|$).

The results show that our approach reduces the size of the giant component substantially for all cases. More importantly, the filtering approach results in the extraction of significant numbers of new components that can be used to generate actionable intelligence for potential victims of sexual exploitation. For the case in which data from all sites is used to construct the initial graph representation, the giant component is reduced from a size corresponding to approximately 57% of the data artifacts to 18%, and the number of components is increased by more than 100%. Without an effective means to decompose the giant component, all of the linked data would be disregarded because the high degree of noise prohibits actionable intelligence. Thus, the proposed approach is highly effective for increasing the amount of this important and hard-to-obtain data that can be acted upon.

5.6. Benchmark Experiment

We conclude this section with a benchmarking experiment that compares our giant component filtering to that described by Freeman et al. (2022). Specifically, we compare the two approaches in terms of reducing giant component size, elapsed time, and overall component increase.

Freeman et al. (2022) use BC to identify and remove nodes with unusually high connectivity. The BC value

for a vertex v captures the number of times v lies on the shortest path between all other pairs of nodes. Thus, determining BC requires computing the shortest paths between all pairs of vertices in a graph. Noting this significant computational challenge, Freeman et al. (2022) use a random sample of approximately 40% of the vertices in a component to approximate the BC values for all vertices. They do not provide an exact BC threshold value that they use to identify vertices for deletion. Instead, they discuss how they represent BC values relative to the maximum observed value for different artifacts and examine changes in component structure at various relative cutoff thresholds to determine the appropriate settings.

Our benchmark experiment applies the giant component filtering approaches to our data set. We apply the two approaches to each site individually and to all sites simultaneously (All). When implementing the approach from Freeman et al. (2022), which we refer to as the BC-based method, we (i) construct the graph representation of the data, (ii) extract the giant component, (iii) approximate the BC values for all vertices in the giant component using a random sample of approximately 40% of the giant component's vertices, (iv) express the BC values relative to the maximum value observed by artifact type, and (v) apply a percentile threshold to identify vertices for removal. For example, applying a threshold of 0.90 removes all vertices with a relative BC value that equals or exceeds the 90th percentile for all artifact types. Noting that BC computation is expensive, we limit the time allowed to one day (86,400 seconds). Also, we use the *nx-cugraph* library for Python to accelerate BC computation using the GPU available on the testing workstation (see Section 5.2).

Table 5 presents the benchmark experiment results for giant component size reduction and elapsed time. We consider four different threshold values of 0.75, 0.9, 0.95, and 0.99 for the BC-based filtering procedure. The results associated with these thresholds are given in the columns BC - 0.75, BC - 0.9, BC - 0.95, and BC - 0.99, respectively. The reduction in giant component size is expressed as a percentage of the number of vertices in

Table 4. Giant Component Filtering Performance Summary

Site	$ V $	$ E $	$ C $	GC proportion	Filtered $ C $	Filtered GC proportion	$ C $ increase
Site 1	2,413,478	3,740,639	69,990	0.23	93,370	0.02	0.33
Site 2	988,045	1,158,752	61,061	0.47	79,009	0.25	0.29
Site 3	1,507,023	2,329,791	56,715	0.50	100,682	0.21	0.78
Site 4	169,103	320,596	1,833	0.73	5,854	0.10	2.19
Site 5	86,720	194,188	3,738	0.31	6,294	0.02	0.68
Site 6	1,741,723	4,744,453	44,914	0.45	79,385	0.16	0.77
Site 7	1,742,789	2,708,611	62,769	0.53	119,987	0.22	0.91
Site 8	2,150,640	2,736,256	78,225	0.05	81,867	0.00	0.05
Site 9	792,005	1,007,027	82,841	0.08	95,751	0.02	0.16
All	7,753,122	15,667,074	143,489	0.57	291,530	0.18	1.03

Table 5. Benchmark Experiment Results: Giant Component Reduction and Elapsed Time

Site	Giant component size reduction, %					Elapsed time, seconds	
	Proposed	BC - 0.75	BC - 0.90	BC - 0.95	BC - 0.99	Proposed	BC
Site 1	20.90	0.21	0.20	0.20	0.20	402	6,773
Site 2	22.90	16.77	16.65	16.65	16.65	275	5,802
Site 3	29.55	50.39	50.36	50.29	49.62	1,070	38,326
Site 4	62.55	10.51	10.35	10.35	10.35	82	3,257
Site 5	28.65	5.05	5.03	4.40	4.40	249	141
Site 6	28.83	0.08	0.08	0.08	0.08	1,039	7,863
Site 7	30.33	0.34	0.16	0.16	0.16	1,374	10,829
Site 8	4.85	0.04	0.03	0.03	0.03	74	803
Site 9	5.60	0.63	0.63	0.63	0.63	215	242
All	39.16	—	—	—	—	9,300	—

the original graph representation. The elapsed time is expressed in seconds. Bold font indicates the best performance for each site with larger reductions in giant component size and shorter elapsed time values being preferred.

Table 5 shows that our giant component filtering procedure results in significantly larger size reductions when compared with that of Freeman et al. (2022). Moreover, our procedure tends to outperform in terms of elapsed time with the BC-based procedure failing to complete within the one-day time limit for the All case. There are two exceptions: the BC-based method (i) yields a larger reduction in giant component size for site 3 and (ii) requires less computational time to filter the giant component for site 5. We see that the larger giant component reduction for site 3 does not necessarily translate into a more effective decomposition. Regarding the elapsed time result for site 5, this site has the smallest graph representation (see Table 4), suggesting that the computational overhead associated with computing post embeddings in our process can result in it being outperformed on small graphs. Table 6 compares the increase in components associated with the two approaches. Bold text indicates the method that yields the largest increase in the number of components.

Table 6. Benchmark Experiment Results: Component Increase

Site	Component increase, %				
	Proposed	BC - 0.75	BC - 0.9	BC - 0.95	BC - 0.99
Site 1	33.41	2.95	2.94	2.94	2.89
Site 2	29.40	3.19	1.70	1.70	1.70
Site 3	77.52	0.08	0.44	1.35	7.07
Site 4	219.42	602.45	602.40	602.40	602.40
Site 5	68.41	1.69	1.55	1.47	1.44
Site 6	76.75	0.94	0.94	0.94	0.94
Site 7	91.16	5.63	2.09	2.09	2.09
Site 8	4.66	0.10	0.05	0.04	0.01
Site 9	15.59	0.67	0.67	0.67	0.66
All	103.17	—	—	—	—

Table 6 shows that our procedure typically results in substantial increases in the number of components. Recall that each component represents data associated with an individual or unique posting entity. Thus, the increase in components directly translates into a greater potential to identify potential victims of sexual exploitation. One exception is site 4, on which both methods yield substantial increases in components, but the BC-based approach increases the number of components by more than 600%. Table 4 shows that the giant component in the graph representation for site 4 includes approximately 75% of all vertices, explaining the high potential for additional components. Moreover, the performance of the BC-based approach suggests it may be better for some site characteristics.

Another interesting case is site 3. Whereas Table 5 shows that the BC-based method provides a larger reduction in giant component size for this site, our procedure is more effective at increasing the number of components. This observation highlights the potential for the BC-based approach to result in significant data loss. In other words, the substantial reduction in giant component size seems to be a result of deleting a significant proportion of vertices, and this limits the increase in components. Our procedure does not delete data; instead, it removes data associations. Together, the results presented in Tables 5 and 6 show that our approach generally improves upon the performance of that described in Freeman et al. (2022) in terms of giant component size reduction, computational time, and the degree of component increase. Moreover, it provides these benefits without data loss.

6. Conclusion

We present an end-to-end process for collecting and linking sex ad data, identifying and removing erroneous links, and transforming this data into intelligence for counter-trafficking operations. Our methodology is refined through ongoing collaborations with LEO and NPO partners, directly supporting operations that have identified and connected with more than 60 potential

sex trafficking victims. Our ad-linking process addresses significant limitations in previous approaches, particularly that described by Freeman et al. (2022). A critical contribution is recognizing the emergence of a giant component in the graph representation of sex ad data, establishing an important connection between counter-trafficking efforts and network science. This recognition enables us to develop a tailored approach for detecting and removing potentially erroneous edges in the giant component. This procedure

- Is data-driven, leveraging recent advances in language models specifically tailored for ASW data.
- Results in no data loss, preserving all collected information and removing potentially erroneous associations.
- Includes a novel approach to create training data based on observations from components not connected in the giant component, which allows our approach to evolve based on the available data.

To aid reproducibility and validate our approach, we demonstrate its efficacy using an open data set that contains a diverse sample of data collected from nine different ASWs. Our experimental results show substantial improvements over previous methods in terms of giant component size reduction, computational efficiency, and component increase with our approach yielding an increase in the number of components exceeding 100% when applied to the combined data set. Moreover, we take care to explicitly describe our tools, techniques, and the reasoning behind the design choices embedded in the process. We will consider requests for Python scripts that demonstrate the process of transforming the data set described in Freeman et al. (2025a) into a graph representation and filtering the resulting giant component. All requests should be made to the corresponding author. Given the sensitive nature of the application area, all requests will first be vetted with our network of law enforcement partners.

Regarding limitations, first and foremost, we want to emphasize that not all ads posted on ASWs are associated with victims of sex trafficking, and our process does not attempt to predict whether an individual is being trafficked. Based on our fieldwork experience, we believe such a determination is impossible based solely on ASW data. Instead, our ultimate goal is to develop techniques that transform the large volumes of ASW data created daily into data products that enhance the efforts of subject matter experts in the field. We are mindful of the risk that our data products could have unintended consequences and have made substantial efforts to build a network of subject matter experts that we trust and can consult with. Specifically, our team has established relationships with law enforcement agencies, nonprofit organizations, and survivor advocates with which we actively consult to ensure our work is ethically sound, practical, and minimizes the risk of unintended consequences. Moreover, we only

provide data products to organizations focused on the identification or potential sex trafficking victims rather than general law enforcement operations.

Second, ASWs are only one area of the internet that facilitates sex trafficking. Others include social media and sites on the Dark Web. However, given the well-documented history of sex trafficking being facilitated by ASWs and the general availability of these sites, we believe this work represents a significant advance in the fight against sexual exploitation.

Third, a key feature of our process is that it results in no data loss. However, we know that some proportion of ads posted on ASWs are scams (Freeman et al. 2022). Instead of trying to detect them explicitly in our process, we allow them to remain. Table 4 shows that, even though our procedure is very effective at extracting components from a giant component, it is often the case that some residual (but much smaller) giant component remains. It may be that the majority of scam ad data are contained in this residual giant component, or it may not. However, instead of making assumptions about how scam ads manifest in the graph representation, we choose to allow them to persist. A better understanding of this manifestation is a fruitful area for future research. Note that, when working with partners in the field, we apply filters to limit the components that are used to construct data products. One filter we commonly apply only includes components that have posted ads targeting the partner's location at a specified targeting rate or higher. Our experience suggests that scammers tend to target many locations simultaneously with their ads, which makes the targeting rate for any specific location rather low. Thus, our experience suggests the risk of scam ads making it to data products is low.

In addition to a better understanding of how scam ads manifest in the graph representation of ad data, two additional research directions come to mind. The first is a better understanding of the online commercial sex ecosystem and tools to quickly identify new sites and shifts in volume. As noted in the introduction, this ecosystem is highly fluid and subject to significant disruptions in response to policy changes. During our years working in this area, we have observed dramatic shifts in site usage following changes such as the implementation of more stringent account creation guidelines or the introduction of posting fees. We have also identified cases in which posting entities in certain regions use specific sites or boards that are not widely used in other areas. Improving discovery mechanisms for these emerging patterns remains a challenge. Second, although we use images for linking advertisements, we do not currently use them to detect erroneous links. Although we have experimented with tools to utilize data embedded in images (such as facial recognition), many of these tools perform poorly with the

low-resolution images prevalent on ASWs. Given the volume of advertisements hosted by these sites, the majority of images are of low resolution, necessitating the development of specialized tools and techniques that can accommodate this reality.

In conclusion, our ad-linking process represents a significant methodological advancement in the fight against sex trafficking, providing LEO and NPO partners with actionable intelligence derived from the vast amount of data available on ASWs. By connecting network science principles with counter-trafficking efforts and developing specialized tools for this domain, we establish a process that supports the efficient and effective identification of potential victims and provides a foundation for future research in this critical area.

References

- Batagelj V, Zaversnik M (2003) An O(m) algorithm for cores decomposition of networks. Preprint, submitted October 25, <https://arxiv.org/abs/cs/0310049>.
- Bommasani R, Hudson DA, Adeli E, Altman R, Arora S, von Arx S, Bernstein MS, et al. (2022) On the opportunities and risks of foundation models. Preprint, submitted August 16, <https://arxiv.org/abs/2108.07258>.
- Börner K, Sanyal S, Vespignani A (2007) Network science. *Annual Rev. Inform. Sci. Tech.* 41(1):537–607.
- Chi L, Zhu X (2017) Hashing techniques: A survey and taxonomy. *ACM Comput. Surveys* 50(1):1–36.
- Cresci S, Di Pietro R, Petrocchi M, Spognardi A, Tesconi M (2017) The paradigm-shift of social spambots: Evidence, theories, and tools for the arms race. Barrett R, Cummings R eds. *Proc. 26th Internat. Conf. World Wide Web Companion* (International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva), 963–972.
- Csardi G, Nepusz T (2006) The igraph software package for complex network research. *InterJournal Complex Systems*:1695. Accessed October 29, 2025, <https://igraph.org>.
- Devlin J, Chang MW, Lee K, Toutanova K (2019) Bert: Pre-training of deep bidirectional transformers for language understanding. Burstein J, Doran C, Solorio T, eds. *Proc. 2019 Conf. North Amer. Chapter Assoc. Comput. Linguistics Human Language Tech.*, vol. 1 (Association for Computational Linguistics, Stroudsburg, PA), 4171–4186.
- Dianati N (2016) Unwinding the hairball graph: Pruning algorithms for weighted complex networks. *Phys. Rev. E* 93(1):012304.
- Ding N, Qin Y, Yang G, Wei F, Yang Z, Su Y, Hu S, et al. (2023) Parameter-efficient fine-tuning of large-scale pre-trained language models. *Nature Machine Intelligence* 5(3):220–235.
- Ellis D (2023) Pulaski County man pleads guilty in federal court to sex trafficking of a minor. *Northwest Arkansas Democrat Gazette Online* (March 29), <https://www.nwaonline.com/news/2023/mar/29/pulaski-county-man-pleads-guilty-in-federal-court/>.
- Farid H (2021) An overview of perceptual hashing. *J. Online Trust Safety* 1(1).
- FOX6 News Milwaukee (2022) Child sex trafficking case: Milwaukee man gets 25 years in federal prison. Accessed October 29, 2025, <https://www.fox6now.com/news/child-sex-trafficking-case-milwaukee-man-federal-prison>.
- Freeman N, Keskin BB, Bott GJ (2022) Collaborating with local and federal law enforcement for disrupting sex trafficking networks. *INFORMS J. Appl. Anal.* 52(5):446–459.
- Freeman N, Bott GJ, Keskin BB, Marcantonio TL (2025a) A multisite data sample for analyzing the online commercial sex ecosystem. *Sci. Data* 12:243.
- Freeman N, Nguyen T, Bott G, Parton J, Francel C (2025b) Language models for adult service website text analysis. Preprint, submitted July 14, <https://arxiv.org/abs/2507.10743>.
- Harris CR, Millman KJ, van der Walt SJ, Gommers R, Virtanen P, Cournapeau D, Wieser E, et al. (2020) Array programming with NumPy. *Nature* 585(7825):357–362.
- Hunter JD (2007) Matplotlib: A 2D graphics environment. *Comput. Sci. Engrg.* 9(3):90–95.
- Husain SS, Sharma K, Kukreti V, Chakraborti A (2020) Identifying the global terror hubs and vulnerable motifs using complex network dynamics. *Physica A Statist. Mech. Appl.* 540:123113.
- Incitti F, Urli F, Snidaro L (2023) Beyond word embeddings: A survey. *Inform. Fusion* 89:418–436.
- Janson S, Knuth DE, Luczak T, Pittel B (1993) The birth of the giant component. *Random Structures Algorithms* 4(3):233–358.
- Keskin BB, Bott GJ, Freeman NK (2021) Cracking sex trafficking: Data analysis, pattern recognition, and path prediction. *Production Oper. Management* 30(4):1110–1135.
- Kitsak M, Gallos LK, Havlin S, Liljeros F, Muchnik L, Stanley HE, Makse HA (2010) Identification of influential spreaders in complex networks. *Nature Phys.* 6(11):888–893.
- Lee MC, Vajiac C, Kulshrestha A, Levy S, Park N, Jones C, Rabbany R, Faloutsos C (2021) InfoShield: Generalizable information-theoretic human-trafficking detection. Ailamaki A, Garofalakis M, eds. *IEEE 37th Internat. Conf. Data Engrg.* (IEEE Computer Society, Washington, DC), 1116–1127.
- Lewis TG (2011) *Network Science: Theory and Applications* (John Wiley & Sons, Hoboken, NJ).
- Li L, Simek O, Lai A, Daggett M, Dagli CK, Jones C (2018) Detection and characterization of human trafficking networks using unsupervised scalable text template matching. Abe N, Liu H, Pu C, Hu X, Ahmed N, Qiao M, Song Y, et al., eds. *IEEE Internat. Conf. Big Data* (IEEE, Piscataway, NJ), 3111–3120.
- Liu Y, Ott M, Goyal N, Du J, Joshi M, Chen D, Levy O, Lewis M, Zettlemoyer L, Stoyanov V (2019) RoBERTa: A robustly optimized BERT pretraining approach. Preprint, submitted July 26, <https://arxiv.org/abs/1907.11692>.
- McInnes L, Healy J, Astels S (2017) HDBSCAN: Hierarchical density based clustering. *J. Open Source Software* 2(11):205.
- McKinney W (2010) Data structures for statistical computing in Python. vander Walt S, Millman J, eds. *Proc. 9th Python Sci. Conf. (SciPy)* (Austin, TX), 56–61.
- Pandas Development Team (2020) pandas-dev/pandas: Pandas. Accessed October 29, 2025, <http://dx.doi.org/10.5281/zenodo.3509134>.
- Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, Blondel M, et al. (2011) Scikit-learn: Machine learning in Python. *J. Machine Learn. Res.* 12(85):2825–2830.
- Qorib M, Moon G, Ng HT (2024) Are decoder-only language models better than encoder-only language models in understanding word meaning? Ku LW, Martins A, Srikumar V, eds. *Findings Assoc. Comput. Linguistics* (Association for Computational Linguistics, Stroudsburg, PA), 16339–16347.
- Radford A, Narasimhan K, Salimans T, Sutskever I (2018) Improving language understanding by generative pre-training. Accessed October 29, 2025, https://cdn.openai.com/research-covers/language-unsupervised/language_understanding_paper.pdf.
- Ran J, Li H, Fu J, Liu L, Xing Y, Li X, Shen H, et al. (2013) Construction and analysis of the protein-protein interaction network related to essential hypertension. *BMC Systems Biol.* 7(1):32.
- Reimers N, Gurevych I (2019) Sentence-BERT: Sentence embeddings using Siamese BERT-networks. Preprint, submitted August 27, <https://arxiv.org/abs/1908.10084>.
- Simplemaps (2022) version 1.6. Accessed October 29, 2025, <https://simplemaps.com/data/us-cities>.
- Sood M, Yağın O (2023) Existence and size of the giant component in inhomogeneous random k-out graphs. *IEEE Trans. Inform. Theory* 69(12):8081–8096.

- Tong E, Zadeh A, Jones C, Morency LP (2017) Combating human trafficking with deep multimodal models. Preprint, submitted May 8, <https://arxiv.org/abs/1705.02735>.
- Touvron H, Lavril T, Izacard G, Martinet X, Lachaux MA, Lacroix T, Rozière B, et al. (2023) Llama: Open and efficient foundation language models. Preprint, submitted February 27, <https://arxiv.org/abs/2302.13971>.
- U.S. Department of Justice (2023) Attorney's Office, Middle District of North Carolina. Durham man sentenced to six life terms in prison for sex trafficking. Accessed October 29, 2025, <https://www.justice.gov/usao-mdnc/pr/durham-man-sentenced-six-life-terms-prison-sex-trafficking>.
- Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, Kaiser L, Polosukhin I (2017) Attention is all you need. Guyon I, Von Luxburg U, Bengio S, Wallach H, Fergus R, Vishwanathan S, Garnett R eds. *Adv. Neural Inform. Processing Systems*, vol. 30 (Curran Associates Inc., Red Hook, NY).
- Virtanen P, Gommers R, Oliphant TE, Haberland M, Reddy T, Cournapeau D, Burovski E, et al. (2020) SciPy 1.0: Fundamental algorithms for scientific computing in Python. *Nature Methods* 17(3):261–272.
- Wang Y, Sun Y, Huang H, Rudin C (2024) Dimension reduction with locally adjusted graphs. Preprint, submitted December 19, <https://arxiv.org/abs/2412.15426>.
- Waskom ML (2021) Seaborn: Statistical data visualization. *J. Open Source Software* 6(60):3021.
- Wolf T, Debut L, Sanh V, Chaumond J, Delangue C, Moi A, Cistac P, et al. (2020) Huggingface's transformers: State-of-the-art natural language processing. Preprint, submitted July 14, <https://arxiv.org/abs/1910.03771>.
- Zeng HS, Danaher B, Smith MD (2022) Internet governance through site shutdowns: The impact of shutting down two major commercial sex advertising sites. *Management Sci.* 68(11):8234–8248.
- Zhu J, Li L, Jones C (2019) Identification and detection of human trafficking using language models. Brynielsson J, ed. *2019 Eur. Intelligence Security Informatics Conf.* (IEEE, Piscataway, NJ), 24–31.