JOURNAL OF
EMERGING INVESTIGATORS

Article

Uncovering the hidden trafficking trade with geographic data and natural language processing

Rusmiya Aqid[1], Clayton Greenberg[2]
[1] Lexington High School, Lexington, Massachusetts
[2] Inspirit AI, Palo Alto, California

SUMMARY

Human trafficking is a topic that is both underreported and under-prosecuted in comparison to other crimes. Often this is a result of being unable to detect the crime in the first place. To ease this detection problem, we present a data-driven map visualization and an evidence-based detection tool to improve human trafficking detection. For our tool, we hypothesized that a machine learning model, using natural language processing (NLP), could analyze text to identify socioeconomic patterns in trafficking to detect if trafficking is present. To identify potential patterns, we mapped gross domestic product (GDP) and trafficking information on separate maps. We then statistically drew connections between GDP per capita and reported trafficking rates around the world and in the US, and found a negative relationship between the two variables in the world. We created the detection tool using a logistic regression model on a manually compiled dataset to identify trafficking instances from qualitative data. In making a detection tool, we aimed to draw clearer distinctions between trafficking and other crimes or events. We anticipate that this distinction may lead to more human trafficking reports and prosecutions. Our final trafficking detection tool predicted labor trafficking cases in reports and interviews with 94% accuracy. Our f1 score, another measure of accuracy, was also 94%. We did not find evidence that the model explicitly used socioeconomic patterns to detect trafficking cases, but our analysis suggested that such patterns may have helped the model make predictions.

INTRODUCTION

Human trafficking is an under-studied crime (1). Furthermore, there are few statistical analyses, machine learning models and tools explaining or making predictions about the crime. The lack of study surrounding human trafficking does not necessarily come from an ignorance of the crime. Instead, trafficking is hard to analyze partly due to ambiguity in what qualifies as trafficking and what does not. There is typically a distinction made between sex and labor trafficking, though there tends to be more gray areas within the properties of each definition (2). For our analyses, we adopted the United States definitions of trafficking in accordance with the Trafficking Victims Protection Act of 2000. The U.S. Code defines two types of trafficking: sex trafficking and forced labor. Sex trafficking encompasses the activities involved when a perpetrator profits from sex acts involving a victim under 18 or the victim experiences force, fraud, or coercion to perform such acts (3). Forced labor encompasses the activities involved when a perpetrator uses force, fraud, or coercion to control and exploit the labor of another person (3).

The lack of analyses on human trafficking makes this an interesting topic of study for data and computer scientists, who often work with data, or a lack thereof, to form conclusions. Among the handful of research papers that currently discuss human trafficking through the lens of machine learning, current approaches to conducting machine learning-based human trafficking analyses often focus solely on sex trafficking and neglect labor trafficking (4). When labor trafficking is studied, it is often grouped together with sex trafficking and does not distinguish between the two types of trafficking. Furthermore, the lack of data available about labor trafficking makes it difficult to train AI models effectively. Thus, machine-learning based labor trafficking analysis is an underexplored application of AI that offers potential for further investigation. Natural language processing (NLP) is a branch of AI which allows computers to understand human language, whether it is presented as text or speech (5). NLP enables computers to find patterns and use those patterns to categorize data, which has been applied to human trafficking analyses in recent years. We wanted to further explore this domain of machine learning research by creating a labor trafficking detection tool which utilizes NLP.

A study named Project RESTART (The Reporting Experiences of Survivors to Analyse in Real-Time) used NLP to identify labor and sex trafficking survivor needs (6). The NLP models analyzed survivor case report data and survivor self-reflections from an app called MeL (Monitoring, Evaluation and Learning) and assigned need categories for each survivor. There were ten category labels: financial, employment, social integration, safety, education, legal, physical health, mental health, accommodation and dependents. The results showed that the NLP model's ability to categorize survivor data matched human analysts for most categories. However, it performed worse than humans for the safety and social integration labels, which were ambiguous and difficult to discern even for human analysts. The model was also trained on limited data making it harder to distinguish between patterns pertaining to each category. In another study, Trafficking in Persons reports – annually released federal trafficking reports – were analyzed using NLP to find hidden labor and sex trafficking patterns in the text that were not apparent before (7). Alongside other NLP techniques to find semantic connections between words, the study used text clustering, a technique which groups similar words together, to find associations between words in country-level data to then visualize countries as source, transit, and destination countries. Source countries are where victims are recruited, destination countries are where trafficked individuals are exploited, and transit countries are where victims are taken to on their way to the destination country. The study revealed that Thailand was often a transit country for international sex trafficking operations. These two studies presented evidence that NLP could be used to derive useful trafficking information from qualitative data. However, while the current literature provides insights into trafficking information across different regions and identifies needs for survivors, it does not discuss diagnostic or prevention tools that indicate trafficking. To the best of our knowledge, there has been no prior work to detect individual cases of trafficking using an automated machine learning approach.

Due to the gaps in identifying human trafficking instances, the goal of our study was twofold: To present useful maps characterizing human trafficking relative to socioeconomic factors and to describe a classification model made with machine learning to predict labor trafficking instances. We hypothesized that NLP can identify keywords such as socioeconomic indicators that indicate human trafficking and use those patterns to distinguish between trafficking and non-trafficking cases. To understand whether socioeconomic data is presently an identifiable indicator of trafficking, we mapped trafficking data against a socioeconomic feature: GDP. We chose to map GDP per capita specifically because it is a common metric used to measure a country's economic growth and poverty rate. We found that GDP per capita was a reasonable predictor of trafficking in most countries, but not at the US state level. For our classification tool, we used three-word vectorization approaches: TF-IDF, spaCy, and CountVectorizer. Our classification tool is not meant to be used as an all-encompassing device for trafficking. Rather, it serves as an additional tool to understand trafficking in a way that utilizes the plethora of existing qualitative trafficking data

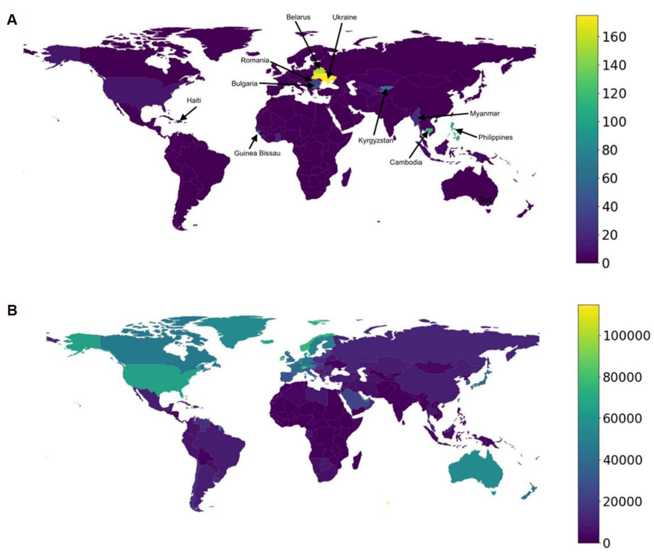

Figure 1: Global reports per million individuals and GDP per capita for countries relative to one another. A) Trafficking report proportions per million individuals in 177 countries. Countries exceeding 20 trafficking reports per million individuals are marked. **B)** GDP per Capita proportions in 177 countries. Yellow indicates the highest proportion.

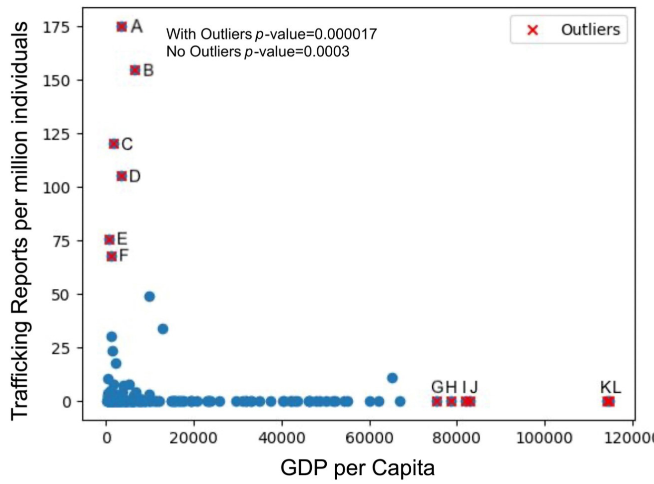

Figure 2: Global trafficking per million individuals and GDP per capita with marked outliers. Global trafficking reports per million individuals against GDP per capita is shown for 177 countries, 165 countries are non-outliers. The following are marked in red as outliers: Ukraine (A), Belarus (B), Cambodia (C), Philippines (D), Guinea-Bissau (E), Kyrgyzstan (F), Norway (G), Ireland (H), Switzerland (I), Falkland Islands (J), The French Southern and Antarctic Lands (K) and Luxembourg (L). Mahalanobis distance was used to detect outliers and Spearman's correlation was used to assess correlation between the two variables.

to predict trafficking. We anticipate that the maps of reported trafficking data and their analysis can be used in projects and research papers as a point of reference from which to extrapolate more details about the trade.

RESULTS

Before constructing our logistic regression trafficking detection model, we aimed to explore whether socioeconomic measures correlate with reported trafficking rates. We hypothesized that NLP could potentially identify socioeconomic patterns within this data, which could enhance the model's ability to detect instances of trafficking. We created two global maps to visualize trafficking and GDP information of 177 countries. In particular, we mapped trafficking reports per million individuals and GDP per capita **(Figure 1).** Trafficking reports per million individuals represents the number of calls per million individuals that the National Human Trafficking Hotline received excluding hang-ups, wrong numbers, or otherwise calls unrelated to trafficking. In our analysis of GDP per capita and trafficking reports per million individuals, we found a negative monotonic relationship between the two variables **(Figure 2, ρ = -0.318).** This negative relationship was also significant **(Figure 2, *p*-value = 0.000017)**. However, we identified twelve outliers in the data in which the Mahalanobis distance D exceeds the cut-off c (**Figure 2, c = 5.991**). Six of those outliers had abnormally high reports of trafficking, namely Ukraine, Belarus, Cambodia, Philippines, Guinea-Bissau, and Kyrgyzstan **(Figure 2, D = 55.353, 42.933, 25.295, 19.142, 9.558, and 7.577, respectively)**. The other six outliers had abnormally high GDP per capita, namely Luxembourg, The French Southern and Antarctic Lands, Falkland, Switzerland, Norway, and Ireland **(Figure 2, D = 21.226, 21.047, 9.832, 9.544, 7.755, and 8.614, respectively)**. In a second model, we omitted these outliers and determined that the correlation

JOURNAL OF
EMERGING INVESTIGATORS
https://doi.org/10.59720/23-027

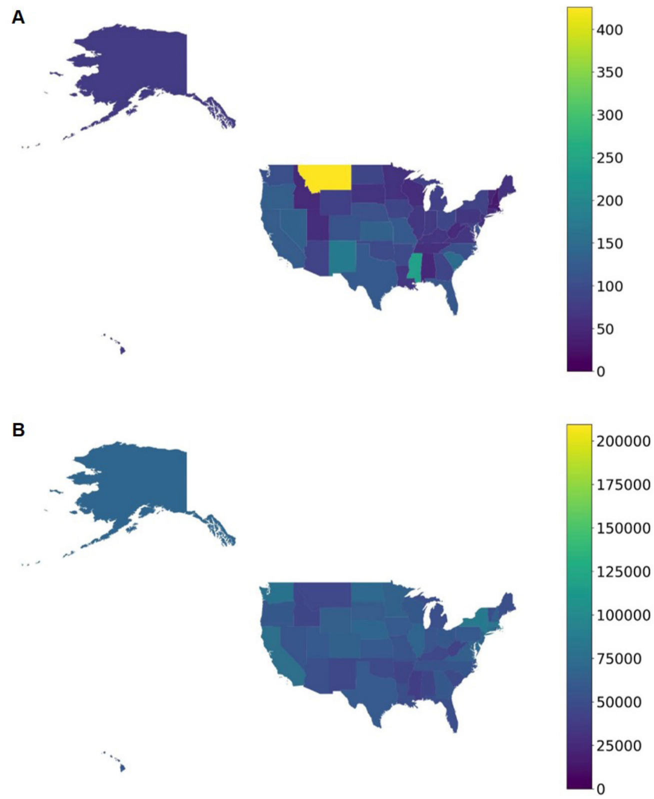

Figure 3: A comparison of the US reported trafficking cases per million individuals and the US GDP per capita. A) Trafficking reports per million individuals' proportion where states with highest trafficking proportion are marked in yellow. **B)** GDP per capita by US state, Washington D.C. marked in yellow.

between trafficking reports per million individuals and GDP per capita was also negative **(Figure 2, ρ = -0.282)**. Trafficking reports per million individuals and GDP per capita in the world without the outliers also showed a significant relationship **(Figure 2, *p*-value = 0.0003)**. Together, the global scatter plots show that global trafficking reports per million individuals and GDP per capita have a significant, negative monotonic relationship, regardless of the presence of outliers.

We then made a map of trafficking reports per million individuals in the United States **(Figure 3A).** We also made a GDP per capita map for the United States **(Figure 3B)**. In our analysis of GDP per capita and trafficking reports per million individuals in the US, we found that the two variables had a low negative monotonic correlation with a non-significant p-value **(Figure 4, *p*-value = 0.957, ρ = -0.00769)**. Notably, we had 3 outliers whose Mahalanobis distances exceeded the cut-off point **(Figure 4, *c* = 5.991).** They were Montana, which had the highest trafficking calls per million individuals, followed by the District of Columbia and Mississippi **(Figure 4, *D* = 28.951, 40.498, and 7.054, respectively).** In omitting the outliers, the relationship between US GDP per capita and trafficking calls per million individuals indicated a positive correlation with a non-significant p-value **(Figure 4, *p*-value = 0.81, ρ = 0.0364)**.

Finding a significant relationship between GDP per capita and trafficking per million individuals on a global scale supported our idea of using a machine learning model to track socioeconomic indicators of trafficking cases. We created two word clouds to find the most frequently occurring words in the news articles and reports that indicated trafficking and those that did not indicate trafficking in our dataset **(Figure 5, Table 1)**. The word cloud with trafficking articles and reports indicated "forced labor," "worker," and "victim" as the most common words **(Figure 5A)**. The word cloud without trafficking articles and reports indicated "worker," "said," and "job" as the most common words **(Figure 5B)**. After finding common words within trafficking and non-trafficking cases in our dataset, we created a logistic regression model to classify trafficking cases based on our qualitative data. We manually compiled articles from various news sources and organization reports to form the qualitative data (**Table 1**). We converted the qualitative data into input for the model using three types of word vectors and measured their accuracies using Stratified K-Fold Cross Validation by splitting the dataset into 5 folds and computing an average accuracy and f1 score **(Figure 6).** First, we used CountVectorizer, a tool in the Python scikit-learn library, which turns articles into numerical representations based on the frequency of the same words in the text. Our model was able to predict trafficking accurately in 78.67% of cases with an f1 score of 76.18% **(Figure 6).** Upon testing this model, we saw that it did an insufficient job of predicting trafficking. For the test, we used a news report of the hate crime incident at an In-n-Out as sample data (8). Our model indicated that it was an instance of labor trafficking, though the report was not a case of labor trafficking. To improve our accuracy, we utilized another tool from the scikit-learn Python library: Term Frequency-Inverse Document Frequency (TF-IDF). TF-IDF considers both the frequency of words within a document and their rarity across all documents. It assigns higher scores to words that appear frequently in a specific document while giving lower scores to words that are common across many documents. This approach helps prevent common words, which hold little informational value, from receiving disproportionately high scores. However, TF-IDF did not improve the performance of the model; the accuracy was 72.90% with an f1 score of 64.67% **(Figure 6).** Finally,

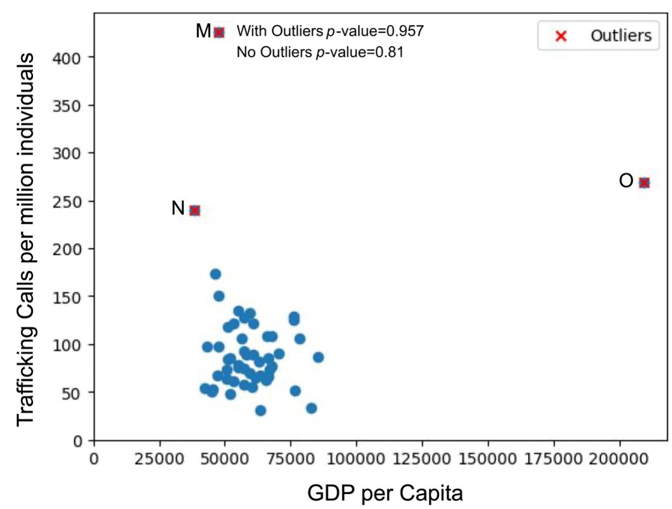

Figure 4: US trafficking per million individuals and GDP per capita with marked outliers. Outliers Montana (M), Mississippi (N), and the District of Columbia (O) are marked in red. Mahalanobis distance was used to detect outliers and Spearman's correlation was used to assess correlation between the two variables.

https://doi.org/10.59720/23-027

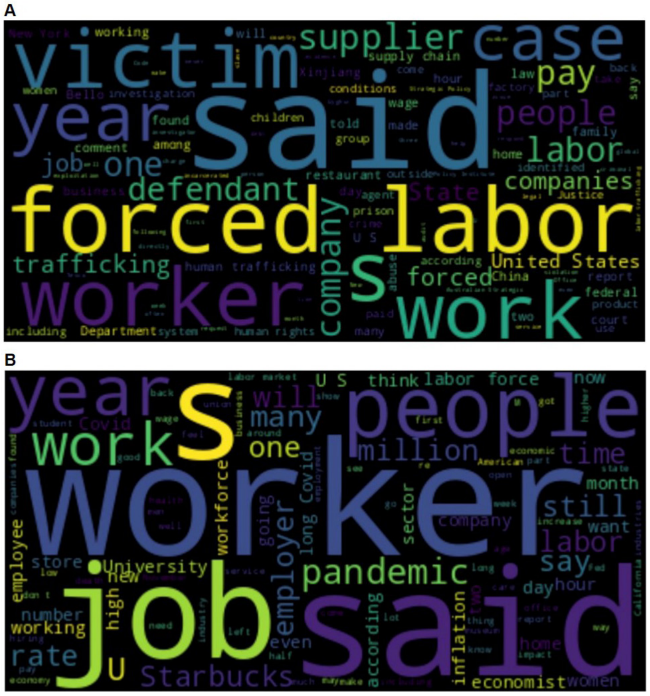

Figure 5: Word clouds representing word frequencies by word sizes. A) Displays most frequent words in cases where trafficking is present, such as "forced labor," "victim," and "worker." **B)** Displays most frequent words in cases where trafficking is not present such as "worker," "said," and "job."

we used spaCy, a Python library for NLP tasks, to obtain Global Vectors for Word Representation (GloVe) vectors. GloVe vectors are known as dense vectors and are pre-trained on a large dataset, allowing the model to learn general patterns and features from the large dataset. A dense vector is an ordered set of numbers that mathematically represent the words of each article, where each number represents a different aspect of the word's meaning. After implementing spaCy's GloVe vectors, our logistic regression model had an accuracy of 94% and an f1 score of 94% **(Figure 6).**

DISCUSSION

Human trafficking is a complex and often elusive crime. Effective detection and prevention require a comprehensive approach that goes beyond the involvement of policymakers and law enforcement. It must include advanced technologies and cross-sector collaboration. Our study considers the integration of an AI classification model and NLP to detect labor trafficking from qualitative data, and we hypothesized that NLP's capability to recognize patterns in text data can lead to the development of a successful detection tool for labor trafficking. We learned that CountVectorizer and TF-IDF were not ideal vector representations to use in the AI model as it led to 78.67% accuracy and 72.90% accuracy, respectively. SpaCy vectors performed well when we used them in our AI model, resulting in 94% accuracy, supporting our prediction.

Before making our model, we wanted to find a relationship between trafficking reports and GDP per capita because poverty-driven child labor remains prevalent. Since GDP is a basic and widely available metric that gives insight into poverty, we focused on GDP specifically. We analyzed trafficking reports per million individuals and GDP per capita and found some data that suggested a correlation between trafficking reports per million individuals and GDP per capita. The two variables had a significant negative monotonic relationship both with and without outliers in our global data. Despite finding a significant relationship, the negative correlation we found was weak. This result may not fully reflect trafficking patterns, likely due to limitations within our data. Our global trafficking reports data came from the Counter Trafficking Data Collaborative Data organization's 2020 reports (9). Since 2020, the COVID-19 pandemic may have affected the trafficking rates around the world, but due to insufficient data, we could not make our model based on 2024 data. Since our analysis is not based on 2024, our conclusions may not represent changes in human trafficking – if there were any – post-2020. For the trafficking reports per million individuals and GDP per capita in the US, there was a nonsignificant relationship between the two variables both with and without outliers. The national trafficking data faces the same dataset issue as the global data since this data comes from National Human Trafficking Hotline's report from four years ago (10). There also may be discrepancies in what is identified as trafficking in global and US reports relative to one another. Still, there were other ways we improved the validity of our results, such as by choosing an optimal correlation test for our data. Since Spearman's rank correlation coefficient test does not require normality, we met the conditions for the correlation test. Thus, our global data reliably indicates a monotonic negative relationship between trafficking per million individuals and GDP per capita and does not indicate any correlation in national data.

Rather than concluding that there is no relationship between GDP per capita and trafficking reports per million individuals in the US, we hypothesize that there may not be enough data on the state-level to reliably indicate otherwise. Human trafficking work already suffers from a lack of data available at large, so there may not be enough data points across states. On the global scale, we had more data points with a higher range of values, which could have been helpful in observing a significant relationship. In identifying an observable correlation between country trafficking rates and

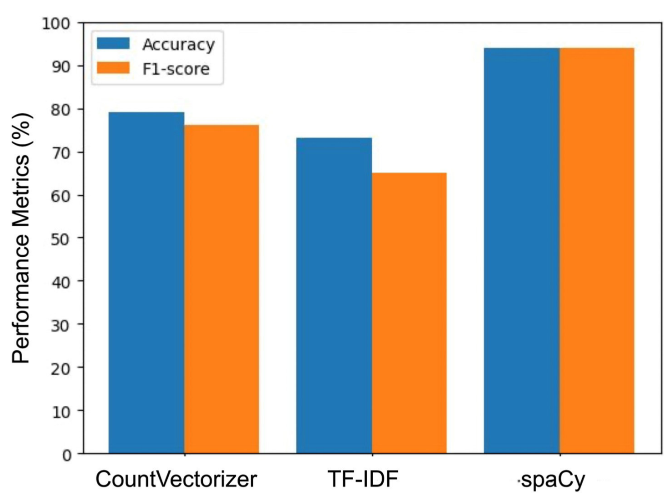

Figure 6: Performance of vectorization techniques. Cross-validated F1-scores and accuracies are shown for vectors made with CountVectorizer, TF-IDF and spaCy.

https://doi.org/10.59720/23-027

Trafficking Category	Report/Article	Source
1	https://humantraffickingsearch.org/resource/remulla-nbi-probing-myanmar-human-trafficking-scheme/	Human Trafficking Search
1	https://www.wric.com/news/virginia-news/midlothian-man-three-others-charged-in-labor-trafficking-at-williamsburg-laundry-facility/	WRIC
1	https://www.cbsnews.com/sanfrancisco/news/san-francisco-couple-charged-with-labor-trafficking-foreign-nanny/	CBS News
1	https://www.ksat.com/news/local/2022/11/17/woman-forced-six-migrants-into-manual-labor-for-slave-wages-sheriff-says/	KSAT
1	https://www.justice.gov/opa/pr/micronesian-couple-pleads-guilty-withholding-passports-labor-trafficking	U.S. Department of Justice Office of Public Affairs
1	https://www.justice.gov/opa/pr/georgia-woman-sentenced-140-months-prison-human-trafficking-two-young-women-nigeria	U.S. Department of Justice Office of Public Affairs
1	https://www.justice.gov/opa/pr/family-convicted-conspiring-force-pakistani-woman-labor-their-virginia-home-12-years	U.S. Department of Justice Office of Public Affairs
1	https://www.justice.gov/opa/pr/illinois-family-charged-kidnapping-forced-labor-and-conspiracy-coercing-two-minors-and-third	U.S. Department of Justice Office of Public Affairs
0	https://www.cnbc.com/2022/12/08/long-covid-is-distorting-the-labor-market-hurting-the-us-economy.html	CNBC
0	https://www.cnbc.com/2022/12/10/us-soccer-journalist-grant-wahl-dies-while-covering-world-cup-in-qatar.html	CNBC
0	https://www.theguardian.com/business/2022/nov/17/starbucks-union-fight-delays	The Guardian
1	https://polarisproject.org/labor-trafficking-examples/	The Polaris Project
1	https://polarisproject.org/labor-trafficking-examples	The Polaris Project
1	https://polarisproject.org/labor-trafficking-examples/	The Polaris Project
1	https://polarisproject.org/labor-trafficking-examples/	The Polaris Project
1	https://polarisproject.org/labor-trafficking-examples/	The Polaris Project
1	https://polarisproject.org/labor-trafficking-examples/	The Polaris Project
0	https://www.cnn.com/2022/12/15/economy/men-missing-from-the-labor-force/index.html	CNN
0	https://www.cnn.com/2022/09/09/economy/overlooked-workers-labor-shortage/index.html	CNN
0	https://www.cnn.com/2022/08/19/economy/worker-shortage-small-business/index.html	CNN
0	https://www.cnn.com/2021/12/25/economy/labor-shortage-early-retirement-charts/index.html	CNN
0	https://www.cnn.com/2022/06/13/perspectives/labor-market-economy-jobs/index.html	CNN
0	https://www.reuters.com/markets/us/us-retail-sales-fall-more-than-expected-november-weekly-jobless-claims-decrease-2022-12-15/	Reuters
0	https://www.businessinsider.com/labor-shortages-immigration-retirement-covid-2022-12	Business Insider
0	https://www.bostonglobe.com/2022/12/17/business/tewksbury-nursing-home-labor-shortage-reveals-an-industry-distress/	Boston Globe
1	https://www.npr.org/2022/10/25/1131279303/advocates-say-the-number-of-labor-trafficking-victims-is-vastly-undercounted	NPR
1	https://www.state.gov/countering-labortrafficking-two-case-studies/	US Department of State
0	https://www.businessinsider.com/stack-ranking-examples-facebook-salesforce-google-2022-12	Business Insider
0	https://gazette.com/business/why-are-labor-participation-rates-among-young-men-low-tatiana-bailey/article_9dddff96-7bbf-11ed-bccd-ef08e60ce212.html	Gazette
1	https://www.businessinsider.com/companies-brands-china-supply-chains-illegal-forced-labor-2022-12	Business Insider
0	https://www.bbc.com/worklife/article/20220908-the-jobs-employers-just-cant-fill	BBC
0	https://www.bbc.com/news/business-62767748	BBC
0	https://www.wabe.org/starbucks-workers-plan-3-day-walkout-at-100-us-stores/	WABE
0	https://www.npr.org/2022/12/06/1141063784/why-there-are-some-contradictory-narratives-on-the-job-market	NPR
0	https://www.npr.org/2022/10/05/1127047818/why-worker-productivity-has-fallen-in-the-u-s	NPR
0	https://www.npr.org/2022/09/05/1121145005/more-than-2-years-into-the-pandemic-covid-19-continues-to-roil-the-labor-market	NPR
0	https://www.nytimes.com/2022/12/16/us/university-of-california-strike-deal.html	NY Times
0	https://www.uschamber.com/workforce/understanding-americas-labor-shortage-the-most-impacted-industries	US Chamber of Commerce
1	https://humantraffickingsearch.org/resource/incarcerated-people-forced-to-do-dangerous-work-for-slave-wages/	Human Trafficking Search
1	https://humantraffickingsearch.org/resource/workers-who-made-jeans-for-tesco-trapped-in-effective-forced-labour/	Human Trafficking Search
1	https://humantraffickingsearch.org/fighting-human-slavery-why-the-private-sector-should-care/	Human Trafficking Search
1	https://humantraffickingsearch.org/resource/senegal-the-state-must-move-from-commitment-to-strong-action-to-protect-talibe-children/	Human Trafficking Search
0	https://www.cnn.com/2023/01/02/entertainment/jeremy-renner-snow-plow-hospital/index.html	CNN
0	https://www.cnn.com/2023/01/01/us/bryan-kohberger-university-of-idaho-killings-suspect-sunday/index.html	CNN
0	https://www.cnn.com/style/article/yayoi-kusama-hong-kong-m-plus-exhibition/index.html	CNN
0	https://www.cnn.com/2023/01/02/europe/pope-benedict-xvi-lying-in-state-intl	CNN
0	https://www.cnn.com/2022/12/31/asia/china-covid-surge-taiwan-help-intl-hnk	CNN
1	https://www.wgbh.org/news/local-news/2023/01/17/working-like-a-slave-why-human-trafficking-in-restaurants-is-underreported	WGBH
1	https://humantraffickingsearch.org/resource/kashmiri-women-lured-to-uae-describe-labor-trafficking-hellholes/	Human Trafficking Search

Table 1: All articles and reports from our novel dataset used to create Figures 5 and 6. The table displays all 49 rows of our training data, including links to the report or article and the sources of each report or article. For the 'Trafficking Cateogry' column,‘1’ indicates trafficking and ‘0’ indicates no trafficking. Of the columns here, only the ‘Trafficking Category’ column and the article/report text (not included in table) were used in training.

https://doi.org/10.59720/23-027

GDP around the world, we had more confidence that machine learning could pick up on GDP and related factors that indicate labor trafficking. That being said, GDP is just one of the many variables that give insight into a country's economic standing. As such, it could be a confounding variable with many other indicators such as the human development index (HDI), genuine progress indicator (GPI), and the better life index (BLI). Due to lack of access to these metrics, we were not able to compare them to trafficking reports. This was both a caution and an opportunity as we prepared to make our logistic regression classifier. We may have been limited in information for statistical modeling but given the complexity and pattern-finding abilities of NLP, we might draw clearer connections that are missing in today's trafficking analyses. We also anticipate that the statistical analyses we performed on GDP and trafficking and their relationship will be particularly useful for economists and public health/security researchers.

For our logistic regression classifier, we faced challenges when trying to compile a training dataset: there was not enough data available on the public web. Signs of trafficking that have commonly been recognized in nonprofit work and academia beyond GDP, such as belonging to a minority group or lacking access to medical care, were not yet sufficiently reported in quantitative data. Such details were sometimes reported in individual stories and reports, so we focused on qualitative data instead. We resorted to a somewhat arbitrary method for choosing qualitative data to train our model on. On the search engine, we used search terms like "labor" and "work" to find cases that did not indicate labor trafficking, and terms like "labor trafficking" and "forced labor" to find cases that did indicate labor trafficking. This method may have introduced bias into the model – the model perceived cases that we, the authors, found as indicating trafficking to indicate trafficking. The model perceived cases that we found as not indicating trafficking to not indicate trafficking. We may have also left out important articles that did not contain the keywords we specifically searched for when compiling articles. It is important to note that the model only aims to provide insight into trafficking when some sort of documentation is present. It cannot draw conclusions about the presence of trafficking in unreported cases. Though machine learning and NLP algorithms connect and assign importance to words that it sees as most valuable in making classifications, we cannot confirm that the model can identify socioeconomic factors specifically. Furthermore, while we anticipate that our model is assigning more importance to certain words in the input than others to help it classify cases, we do not explicitly know which words are most important in making the distinction between trafficking and non-trafficking cases. Our word clouds give us an indication of words that appear most frequently in relation to trafficking and non-trafficking cases, but we cannot pinpoint specific keywords that are most relevant in making the distinction. Despite being limited to data from the general news wire, we built an efficient and accurate classifier for detecting labor trafficking. We were able to make a successful model that could, for the most part, classify articles as instances of labor trafficking.

Moving forward, we could make changes to our data collection methods and the AI tools we use. In terms of data collection, an affiliation to trafficking prevention organizations or to the Department of State would allow us to get case summaries and reports of acquitted and convicted labor trafficking cases. Thus, the model could predict verdicts based on affidavits and more detailed evidence. With more data to train the model, the model will have more information to identify subtle patterns. It may also perform better on new data, improving its generalizability. To better inform the model, analyzing other socioeconomic factors such as the HDI, GPI, and BLI would also be useful. It would be interesting to compare those factors to sex trafficking data as well. In terms of AI, we could explore other classification models and word embedding techniques to improve our accuracy even more. Support Vector Machines (SVMs) are suitable models for binary classification, generalizable, and robust against overfitting – a problem in which a model performs well on the training set but poorly on unseen data. For the word vectors, we could use BERT, a technique which creates vectors by analyzing the contextual relationships between words in a sentence. We could also consider customizing our classifier to sex trafficking as opposed to labor trafficking. We anticipate that our work and future adaptations to this work will help those in nonprofit sectors, law enforcement, or individuals who may feel that they are experiencing trafficking to detect communications about trafficking automatically, and to help prevent revictimization. While our work cannot replace social workers, it explores the possibility of using automation tools as a layer of protection against victimization.

MATERIALS AND METHODS

Geographic data statistics

To visualize and analyze our geographic data, we used the following Python libraries specializing in data visualization and analysis: Pandas, GeoPandas, Matplotlib, Seaborn, and SciPy. We used Pandas to read and merge our desired datasets together. To make the trafficking reports and GDP maps, we used GeoPandas, which specializes in geospatial data and mapping. To visualize the maps, we used Matplotlib. We made this visualization with global trafficking counts, population, GDP, and their respective country names. The trafficking counts and country names came from the Counter Trafficking Data Collaborative (CTDC) organization (8). CTDC's raw dataset captured the type of trafficking, country location, the victim's relationship to the recruiter, the industry, the means of control and the "source" for each reported trafficking case. "Source" indicated whether the report came from hotline calls or case management. We could not find additional information about the specifics of the case management data collection. Most of these trafficking variables did not have data for many of the countries, so we decided to count the total instances of trafficking by totaling the "source" values for each country without considering the type of source or trafficking. In short, we summed the instances of trafficking from "source" for each country to create a trafficking calls variable and a "country" variable. The frequency of trafficking calls was counted under the column "Reported trafficking incidents." A 2020 population count for each country was added to the dataset which came from the World Population Prospects (11). We also added GDP to the dataset which was adopted by Stanford University from the World Bank's World Development Indicators database (12). From there, we found a GDP per capita count using USD as the currency. To make our two US geographic data visualizations, we took data from the 2020 National Hotline

Annual Report to construct two variables for our datasets: the frequency of trafficking calls for each state and the names of those states (10). A 2020 US state population count was added to our dataset from a report by the US Department of Agriculture (13). We used the population and trafficking calls to calculate trafficking reports per million individuals. We also added the 2020 state GDP from a report by the US Department of Commerce's Bureau of Economic Analysis to our dataset (14). We used the population and GDP data to calculate GDP per capita.

After making the maps, we used Seaborn to construct a histogram to check the distribution of the data. Then, for the global data we plotted GDP per capita against trafficking reports per million individuals. To statistically analyze its pattern, we imported "spearmanr" from the SciPy library, which allowed us to use the Spearman rank correlation coefficient and find a *p*-value for the data. Spearman's ρ measures the strength and direction of monotonic relationships between two variables. After finding a Spearman's ρ value, we suspected that some points on the scatter plot may be outliers, which could affect the validity of the results. In particular, there were some data points with either very high GDP per capita or very high trafficking reports per capita. By importing 'mahalanobis' from SciPy we used the Mahalanobis Distance – a test for finding outliers in multivariate data – to statistically find outliers in our bivariate data. For the US data, we replicated the process we used for the global data. We made a scatter plot of GDP per capita and trafficking reports per million individuals in the US. We reported a correlation and a p-value. We hypothesized that the District of Columbia might be an outlier given its seemingly high GDP per capita. To statistically check for an outlier, we used the Mahalanobis Distance test again. Using the test, we found three outliers including the District of Columbia.

Logistic regression prediction model

Before creating our logistic regression model, we first created word clouds for our input data to see if there were keywords that were common among the trafficking and non-trafficking cases, respectively. We then built the logistic regression model to classify trafficking cases. For our logistic regression model, our data came from Human Trafficking Search, a global resource and research database, and national news channels (15).

The data for the labor trafficking logistic regression predictor came from a variety of sources. In total, there were 49 articles used, of which 23 indicated trafficking and 26 did not indicate trafficking. The articles which indicated labor trafficking came primarily from the Human Trafficking Search organization (15). Cases where labor trafficking was not present came primarily from news sources such as CNN, Aljazeera, and a handful of other news channels. Since Human Trafficking Search only reported on cases of labor trafficking, and not more broadly about labor, outside sources were used for cases that did not indicate trafficking. To clarify, articles were manually classified as containing trafficking content or not. The choice of which articles to include did not follow a strict criterion. On the search engine, we used search terms like "labor" and "work" to find cases that did not indicate labor trafficking, and terms like "labor trafficking" and "forced labor" to find cases that did indicate labor trafficking.

To prepare the data for the regression model, the text was cleaned by turning uppercase letters to lowercase, splitting the text into separate words through tokenization, and getting rid of stop words with the spaCy library. Stop words consist of common words that add little value to a machine learning algorithm and are taken out so that the algorithm can focus on more meaningful words. Examples of words that are commonly taken out include "and," "or," "am," and "has." This process then counted the frequency of words in the input text with CountVectorizer and would later be used by the algorithm to make classification decisions. After preparing the data, the logistic regression model was made by identifying the independent variable (the text itself) and the dependent variable (whether the text would be identified as trafficking, 1, or not trafficking, 0). Since we received subpar accuracy for our first model, which used CountVectorizer to turn the articles into vectors, we then used TF-IDF and spaCy to create vector representations of the articles. As with our first model, we inputted the vectors into the logistic regression model using the same independent and dependent variables.

Source code for trafficking statistical analyses and classification tool

Our source code is available from the GitHub repository, JEI-NLP-Trafficking-Detection-Tool, Aqid, Rusmiya (2024): https://github.com/rusmiyaaqid/JEI-NLP-Trafficking-Detection-Tool.git.

Received: February 13, 2023
Accepted: December 13, 2023
Published: October 14, 2024

REFERENCES

1. "Gaps in Reporting Human Trafficking Incidents Result in Significant Undercounting." *National Institute of Justice*. www.nij.ojp.gov/topics/articles/gaps-reporting-human-trafficking-incidents-result-significant-undercounting. Accessed 28 Jan. 2023.
2. "Migrant Farm Labor Trafficking in North Carolina: Pinning Down Elusive Data." *National Institute of Justice*. www.nij.ojp.gov/topics/articles/migrant-farm-labor-trafficking-north-carolina-pinning-down-elusive-data. Accessed 30 Jan. 2023.
3. United States, Congress, 22. The Government Printing Office, 2006. www.Uscode.house.gov/View.xhtml?Path=/Prelim@title22/chapter78&Edition=Prelim Congress.
4. Deeb-Swihart, Julia, et al. "Ethical Tensions in Applications of AI for Addressing Human Trafficking: A Human Rights Perspective." *Proceedings of the ACM on Human Computer Interaction*, vol. 6, no. 295, 11 November 2022, https://doi.org/10.1145/3555186
5. Holdsworth, Jim. "What Is NLP (Natural Language Processing)?" *IBM*, IBM, www.ibm.com/topics/natural-language-processing. Accessed 4 Sept. 2024.
6. Piotrowicz, Ryszard, et al. "Utilising AI technology to improve modern slavery survivor support." *Modern Slavery and Human Rights Policy and Evidence Centre*. www.modernslaverypec.org/resources/ai-modern-slavery-support. Accessed 1 Aug. 2024.
7. "Analytics tackles the scourge of human trafficking." *SAS*. www.sas.com/en_sg/insights/articles/analytics-tackles-human-trafficking.html
8. "Man Arrested after Hateful Rant Goes Viral on TikTok:

https://doi.org/10.59720/23-027

‘He Just Kept Going on’.” *Gray Television*. www.fox8live.com/2022/12/28/man-arrested-after-hateful-rant-goes-viral-tiktok-he-just-kept-going/. Accessed 20 Jan. 2023.
9. “Counter-Trafficking Data Collaborative.” *CTDC*. www.ctdatacollaborative.org/global-dataset-0. Accessed 1 Jan. 2023.
10. “2020 National Hotline Annual Report.” *National Human Trafficking Hotline*. www.humantraffickinghotline.org/en/resources/2020-national-hotline-annual-report. Accessed 1 Jan. 2023.
11. “World Population Prospects 2020”. *United Nations*. www.population.un.org/wpp/Download/Standard/MostUsed/. Accessed 1 Jan. 2023.
12. “Gross Domestic Product 2019.” *PH241 - Stanford University*. www.large.stanford.edu/courses/2021/ph241/. Accessed 1 Jan. 2023.
13. “Population.” *Economic Research Service U.S. Department of Agriculture*. www.data.ers.usda.gov/reports.aspx?ID=17827. Accessed 1 Jan. 2023.
14. “Gross Domestic Product by State, 2nd Quarter 2021.” *Gross Domestic Product U.S. Bureau of Economic Analysis.* *www.*bea.gov/news/2021/gross-domestic-product-state-2nd-quarter-2021. Accessed 1 Jan. 2023.
15. “Global Resource Database.” *Human Trafficking Search*. www.humantraffickingsearch.org/traffickingresources/. Accessed 5 Jan. 2023.