

Tracing the Unseen: Uncovering Human Trafficking Patterns in Job Listings

Siyi Zhou,^{1,3} Jiankun Peng,^{2,3} Emilio Ferrara^{1,2,3}

¹ Annenberg School of Communication and Journalism, ² Viterbi School of Engineering, ³ USC Information Sciences Institute
University of Southern California
zhousiyi@usc.edu, pengrich@usc.edu, emiliofe@usc.edu

Abstract

In the shadow of the digital revolution, the insidious issue of human trafficking has found new breeding grounds within the realms of social media and online job boards. Previous research efforts have predominantly centered on identifying victims via the analysis of escort advertisements. However, our work shifts the focus towards enabling a proactive approach: pinpointing potential traffickers before they lure their preys through false job opportunities. In this study, we collect and analyze a vast dataset comprising over a quarter million job postings collected from eight relevant regions across the United States, spanning nearly two decades (2006-2024). The job boards we considered are specifically catered towards Chinese-speaking immigrants in the US. We classify the job posts into distinct groups based on the self-reported information of the posting user. Our investigation into the types of advertised opportunities, the modes of preferred contact, and the frequency of postings uncovers the patterns characterizing suspicious ads. Additionally, we highlight how external events such as health emergencies and conflicts appear to strongly correlate with increased volume of suspicious job posts: traffickers are more likely to prey upon vulnerable populations in times of crises. This research underscores the imperative for a deeper dive into how online job boards and communication platforms could be unwitting facilitators of human trafficking. More importantly, it calls for the urgent formulation of targeted strategies to dismantle these digital conduits of exploitation.

Introduction

The scourge of human trafficking presents a formidable challenge in measuring and combating its reach, often muddled by the overlapping definitions with human smuggling. This confusion was clarified in the early 2000s by the United Nations, drawing a clear line between these nefarious activities while acknowledging their frequent co-occurrence (Laczko and Gramegna 2003; UN 2003). In the United States, the narrative is particularly poignant among immigrants, who form the majority of trafficking victims, regardless of their smuggling status upon entry (ACLU 2007). Unraveling the journey of these individuals is crucial to understanding their victimization and identifying the perpetrators.

Traditionally, the nexus between traffickers and their victims was forged in person, within tightly-knit communities. However, the digital age, marked by the ascendancy of social media, has transformed this dynamic, facilitating connections between potential victims and traffickers across vast distances (Fraser 2016). Platforms like *chineseinla.com* serve as a double-edged sword for the Chinese-speaking diaspora in Los Angeles, offering vital community support alongside avenues for illicit employment and exploitation.

Highlighting the urgency of this issue, the *UNODC's 2023 Trafficking in Persons Report* reveals that a significant portion of trafficking victims between 2012 and 2022 were ensnared through online recruitment methods, with an even larger group being indirectly exploited via digital platforms, including dating apps (UNODC 2024). The US Department of State further underscores the role of online job boards as hunting grounds for traffickers aiming at vulnerable populations (Office to Monitor and Combat Trafficking in Persons 2023).

Despite these alarming trends, current anti-trafficking strategies have largely centered on mitigating risks associated with dating apps and escort services, deploying tools such as free background checks, and tools to safeguard or identify victims online (Whitney et al. 2018; Tong et al. 2017; Zhu, Li, and Jones 2019; Lee et al. 2021; Tobey et al. 2024; Khan 2023; Palmquist 2023). Yet, these initiatives stop short of addressing the root of the problem: preemptively identifying traffickers through the guise of legitimate employment offers.

This paper seeks to bridge this gap by turning the lens towards locally-targeted Chinese speaking platforms within the United States, often the first port of call for new immigrants. Through a detailed examination of dubious job advertisements on these platforms, we aim not only to shine a light on potential trafficking operations but also to pioneer detection methods capable of unmasking human traffickers lurking behind the darker corners of the Web.

Methodology

Navigating Culturally Niche Communities

Our investigation delves into the intricacies of culturally specific communities, recognizing that the fabric of human trafficking is woven with threads of demographic vulnerabil-

ities. Pioneering work by Kevin Bales (Bales 2007), which meticulously combed through a global dataset of trafficking cases alongside variables from the World Statistics Pocketbook, has illuminated the stark reality that poverty and limited economic opportunities are often harbingers of trafficking risks. Inspired by Bales’ findings, our study crafts a series of hypothetical yet realistic scenarios to simulate potential targeting by traffickers within distinct demographics, emphasizing the nuanced interplay between socioeconomic factors and vulnerability to exploitation.

Next, we ventured into data collection, pinpointing specific lexicons resonating with at-risk individuals. Our methodology involved an extensive, multilingual Web search for job advertisements, meticulously combing through the digital landscape to uncover where vulnerable demographics might seek employment. The culmination of this effort is encapsulated in Table 1, which presents three examples of constructed scenarios along with a set of keywords—translated into English for clarity—aimed at piercing the veil of where these communities converge in their quest for work.

This strategic search unveiled a constellation of eight domains, each a hub within the Chinese-speaking diaspora in the United States, teeming with job postings. These domains, on the surface ripe with opportunities, potentially mask the sinister undertones of trafficking under the guise of legitimate employment. Our approach not only highlights the importance of nuanced, culturally sensitive research in unveiling trafficking pathways, but also sets the stage for more targeted interventions within these specific communities.

Scenario	Keyword
New immigrants from China with low English proficiency looking for a job	Well-paid, Chinese language job, flexible job
Newly graduated international students looking for jobs that offer visa and green card sponsorship	Sponsor, H1B, EB-2, EB-3, OPT
Chinese speaking populations living in the US without legal work authorizations looking for under the table jobs that pay cash	Cash, under the table, temporary work, part time

Table 1: Example scenarios and corresponding keywords

Job Ads Collection Across Different Platforms

In dissecting the digital anatomy of various domains central to our investigation, a serendipitous discovery was made: a uniformity in web architecture across these platforms. This uniformity not only piqued our interest but also paved the way for a more efficient method of data collection. Figure 1, a schematic representation, sheds light on the commonalities, serving as the blueprint for our data scraping approach.

Leveraging the predictability of these structures, we adopted a breadth-first search strategy as our modus operandi. The initial phase involves a comprehensive sweep of the job listings section, aggregating every accessible link. Following this, we delve deeper, extracting the source code

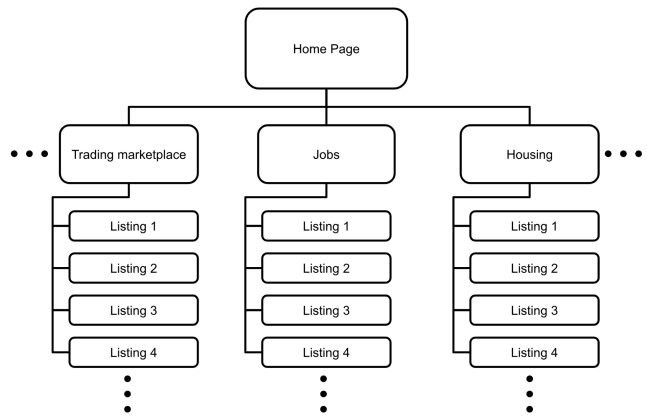


Figure 1: A simplified web structure of job ads platforms

from each link and archiving it as an HTML snapshot on our local storage system. This preemptive measure safeguards against the potential loss of data, ensuring its availability for subsequent analysis even in the event of page removal or content alteration.

The final step of our collection sees the transformation of these HTMLs into structured JSON files, ready for the analytical deep dive. This conversion process not only facilitates a streamlined analysis but also underscores our commitment to preserving the integrity and accessibility of the data collected from these digital platforms.

Extracting Information from Raw Job Ads Data

The culmination of our data collection effort—snapshots of myriad job advertisements—marks the beginning of a critical phase: distillation and organization of data gleaned from these webpages. Central to our analytical framework is the assignment of a unique identifier to each post, with the original URL serving as the pivotal attribute. This step underpins the extraction of several key details from each advertisement, including but not limited to the type of hiring, contact information, job description, listing time, number of views, and authorship information.

The significance of these data points cannot be overstated, as they collectively offer a multifaceted lens through which the complex landscape of online job postings can be scrutinized:

1. **Type of Hiring:** This is instrumental in segmenting the data across different industries, enabling us to pinpoint sectors that might be more prone to suspicious activities.
2. **Contact Information:** A gateway to unraveling the cross-domain presence of users with trafficking intents, this information aids in mapping out the digital footprint of parties potentially involved in nefarious activities.
3. **Job Descriptions:** A cornerstone for enhancing existing language models designed to detect signals of human trafficking. This information provides a deeper understanding of the modus operandi of traffickers, potentially offering novel insights into their tactics and strategies over their victims.

Field	Value
URL	https://www.chineseinla.com/f/page_viewtopic/t_1962623.html
Post Type	hiring
Hiring Type	internship
Industry	computer
Location	1455 Monterey Pass Rd #206 Monterey Park CA
Company Name	LaneCert Education Group
Title	优质转让,高新就业,小白也能轻松落地,留学生OPT, H1B 高新落地
Email	info@lanecert.com
Phone	3234880011, 6267265138, 6297265138
Post	诚意出让七年业界口碑,低收WIOA & 退役军人VA GI BILL 免费学CCIE 顶级讲师,一对一辅导,自有ISP实验室...
Date Created	2020/09/17
Date Modified	2024/01/21
Total Number of Edits	38
Views	57,020
Author	LaneCert.LA
Author Profile	https://www.chineseinla.com/user/id_2385433.html

Table 2: An example of *Job Posting* metadata information

Table 2 presents a snapshot of the organized data, showing the transformation of raw data into a structured repository ready for in-depth analysis. This process not only lays the groundwork for subsequent investigative endeavors but also enriches the toolkit available to researchers and practitioners fighting against the scourge of human trafficking.

Content of our Job Ads Dataset

This segment of our study delves into a detailed exploration and profiling of the dataset at hand. Although Table 2 exemplifies an ideal scenario—where every job listing is accompanied by a complete set of attributes—reality often tells a different story. A significant portion of the listings we encountered were bereft of one or more critical attributes, such as the industry type, hiring type, or geographical location. Our investigation, at this point, is primarily focused on unraveling the patterns associated with the methods of contact delineated within these listings, guided by the industry attribute as self-reported by the recruiters.

Moreover, our initial analysis extends to scrutinizing the temporal dynamics of job postings—how the frequency of listings by recruiters fluctuates over time across different industries. This twofold examination, encompassing both the modes of communication and the time series of posting activity, sheds light on the underlying correlations between online job advertisements and their potential role in facilitating human trafficking.

Through this analytical endeavor, we aspire to uncover the latent threads that may tie seemingly innocuous job postings to the dark underbelly of trafficking operations, thereby contributing to a more nuanced understanding of how digital platforms can be exploited for such heinous purposes.

A General Outline of the Job Ads Dataset

In an effort to pierce through the veil of online job postings within the Chinese-speaking diaspora in the United States, our collection spans a vast array of 258,619 job advertisements. These postings, collected from niche platforms

across eight major U.S. regions, trace back from the year 2006 up to 2024, offering a longitudinal view of the online job market’s evolution. Table 3 summarizes the distribution of these advertisements, laying out a diverse landscape of sources.

A characteristic of our dataset is the uniform presence of certain attributes across all postings—namely, the title, author, job description, and creation date. However, the inclusion of other attributes remains variable, painting a picture of data richness contrasted by gaps in informational completeness. Further distinguishing these job posts is their certification status as indicated by the job post’s address prefix: “j_” signifies a platform-certified posting, while “t_” denotes its uncertified counterpart. Our analytical focus sharpens on the “t_” category, navigating through the uncertified postings to unearth patterns and insights.

The dataset was methodically categorized into nine distinct job groups based on the recruiters’ self-disclosed information, as shown in Figure 2. The findings reveal a significant demand for clerical (18,369), warehouse (13,673), and driving (9,406) positions, with the clerical category encompassing a wide array of roles including secretarial jobs. A deeper dive into this category—specifically, a random examination of 15 secretary job listings—uncovered a disconcerting pattern: some listing solicited interviews at a hotel, and its provided contact number was previously listed on an escort service website, signaling a glaring sex trafficking hazard. Moreover, an analysis of the warehouse and driver job postings revealed 1,294 ads with job descriptions shorter than 200 words. This brevity, marked by a lack of detail, potentially flags an increased risk of labor trafficking, underlining the intricate challenges and dangers lurking within these online job platforms.

Contact Methods by Industry

In our exploration of job advertisement platforms, a striking pattern emerged regarding the preferred methods of contact across various industries.

Domain	Number of posts	Earliest	Latest
chineseinatlanta.com	3,142	2016-01-22	2024-03-12
chineseinflorida.com	2,741	2015-12-24	2024-03-12
chineseinla.com	142,325	2006-09-06	2024-03-12
dcchinaren.com	2,580	2013-03-21	2024-03-12
seattlechinaren.com	6,913	2013-03-17	2024-03-12
vegaschinaren.com	3,588	2010-12-13	2024-03-12
nychinaren.com	58,880	2011-05-13	2024-03-13
chineseinsfbay.com	38,450	2011-01-11	2024-03-14

Table 3: Composition of sources for our *Job Ads* dataset




Figure 2: Number of job ads listed for different industries (self-reported by the posting users). The count does not include those listings that does not specify their industries

Looking at Figure 3, it appears that an overwhelming 58.8% of job postings favored telephone numbers as the primary mode of communication, with 30.6% opting for email as their exclusive contact method. A smaller segment, 10.6%, offered the choice of both phone and email, ensuring a broader avenue for potential applicants to reach out. Phone communication is harder to monitor and often leads to in person contacts, amplifying the risk of harm (Khan 2023).

A deeper dive into the data reveals an industry-specific trend, particularly within the massage parlor sector, where phone numbers are universally adopted as the sole contact method. This practice is not isolated; beyond the realms of clerical and computer-related jobs, a distinct preference for phone communication pervades, spanning sectors like nanny services, restaurants, and driving. These fields demonstrate a marked predilection for phone contacts, coupled with a discernible hesitancy to diversify communication channels.

Table 4 elucidates this trend further, detailing the distribution of phone numbers and email addresses across the board for job applications. The reliance on phone numbers introduces a layer of anonymity and diminishes traceability, in stark contrast to the digital footprint left by email correspondences. This distinction has significant implications for the legitimacy of job advertisements and, by extension, the potential for exploitation. The preference for phone communication, while perhaps convenient, concurrently amplifies the vulnerabilities and risks associated with human trafficking,

Industry	Email	Phone	Both	total
Other	3,028	11,934	1,130	16,092
Sales	2,768	3,153	949	6,870
Computer-related	800	1,003	238	2,041
Driver	393	5,125	216	5,734
Restaurant	100	2,392	57	2,549
Warehouse	4,178	6,629	1,992	12,799
Nanny	11	325	3	339
Clerk	9,278	8,059	2,447	19,784
Service	900	1,748	228	2,876

Table 4: Contact method preferences for different industries




Figure 3: Percentage of preferred contact methods for different industries

underscoring the need for a critical assessment of contact methodologies in job listings.

Job Ads Posting Frequency Over Time

Our next analysis, depicted in Figures 4 and 5, closely examines the temporal dynamics of job posting frequencies, both overall and within specific industries, against the backdrop of major global events over the past five years. Intriguingly, the data suggest that geopolitical tensions, such as the Ukraine-Russia conflict and the conflict between Israel and Hamas, had moderate effects on the volume of job advertisements. However, the impact of the COVID-19 pandemic, particularly within China, paints a different picture, significantly influencing recruiter behavior.

The timeline of the pandemic reveals two distinct periods of activity: the initial outbreak between January and June 2020 saw a sharp decline in job postings, mirroring the global uncertainty and lockdown measures of the period. A similar downturn was observed during the second wave, spanning April to December 2022. Conversely, 2021 emerged as a period of recovery and heightened activity, potentially signaling an increased eagerness among the population to seek employment opportunities in the U.S., possibly driven by desires for immigration or better prospects amid the pandemic's economic fallout.

While the overarching trends in posting frequency appear consistent across industries, a closer inspection reveals

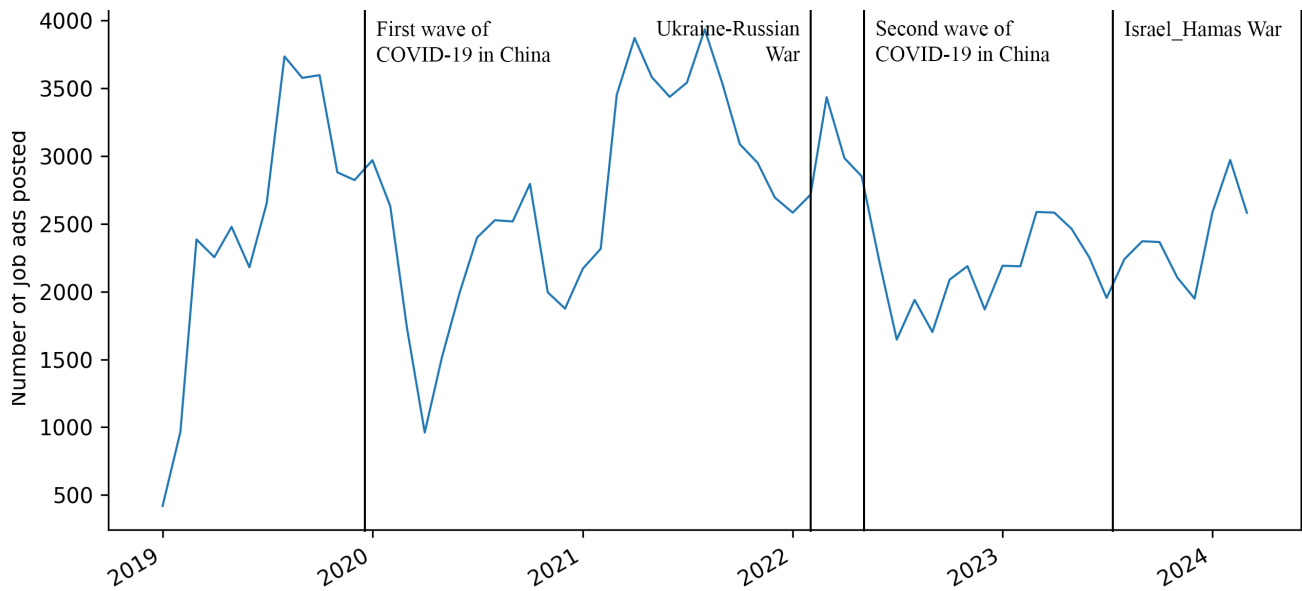


Figure 4: Posting frequencies of all job posts from 2019 to 2024

nuanced variations. Industries such as nanny services and computer-related sectors showed remarkable resilience, remaining relatively unfazed by the highlighted events. On the other hand, the service, clerical, driving, and warehouse sectors demonstrated more pronounced fluctuations, indicating a higher degree of sensitivity to external disruptions.

This disparity in response rates across industries underscores the necessity for a deeper dive into the data. It prompts us to consider whether the vulnerability of certain job sectors to external events could serve as a barometer for detecting and understanding the behavior of human traffickers online. Such insights could be invaluable in devising more targeted strategies for combating trafficking in the digital age.

Discussion and Potential Applications

The collection of our Job Ads dataset marks a significant stride towards enriching the arsenal against online human trafficking, offering new angles to examine and thwart the intricate web of traffickers. Unlike conventional datasets that pivot around escort advertisements (Szekely et al. 2015; Kejriwal and Szekely 2017a,b; Hundman et al. 2018), with a primary aim of rescuing victims post-factum, our dataset ventures into preemptive territory. It aspires not only to spotlight the perpetrators behind these heinous acts but also to stifle the propagation of their nefarious messages. Beyond its immediate utility in combating trafficking, this dataset might enable the study of immigration patterns and facilitate research within Asian American communities.

A pivotal insight from our exploratory analysis underlines a stark contrast in communication preferences on these platforms compared to more formalized channels like LinkedIn and Indeed. The predilection for phone communication in job advertisements, as revealed, is fraught with vulnerabilities. Phone numbers, with their cloak of anonymity and dis-

posability, emerge as a double-edged sword. They afford traffickers the agility to swiftly weave a rapport with unsuspecting individuals, urging decisions under duress, away from the scrutinizing eyes of law enforcement. Despite the inherent challenges in tracing phone-based exchanges, these contact numbers inadvertently weave a thread that could unravel clandestine networks, potentially bridging the gap to escort services and enabling cross-platform hunts for trafficker identities (Ferrara et al. 2014; Catanese, Ferrara, and Fiumara 2012).

Moreover, the temporal trends captured in job postings shed light on the fluid dynamics of niche job markets and their susceptibility to human trafficking, especially within minimally regulated sectors. When dovetailed with datasets chronicling pivotal events or shifts in public policy and discourse, these trends hold promise for forecasting trafficker behavior with unprecedented precision.

Our methodological backbone is the industry classification as self-reported by the job posters, which, despite covering only a fraction (26.7%) of the total dataset, significantly surpasses the NIH's recommended threshold for epidemiological studies (Martínez-Mesa et al. 2014). This foundation not only legitimizes the robustness of our findings but also hints at the untapped potential of the unclassified posts, which, through linguistic analysis, could further refine our understanding of job market patterns and trafficking risks.

However, the dataset navigates through a sea of challenges. Temporal shifts in website content and inherent inconsistencies in job post attributes introduce a layer of complexity. While the universal presence of titles and job descriptions offers a lifeline for attribute extrapolation, the exclusion of deleted or moderator-removed postings casts a shadow on the dataset's comprehensiveness. This omission might skew the portrayal of the job market's darker underbellies, potentially glossing over the prevalence and markers

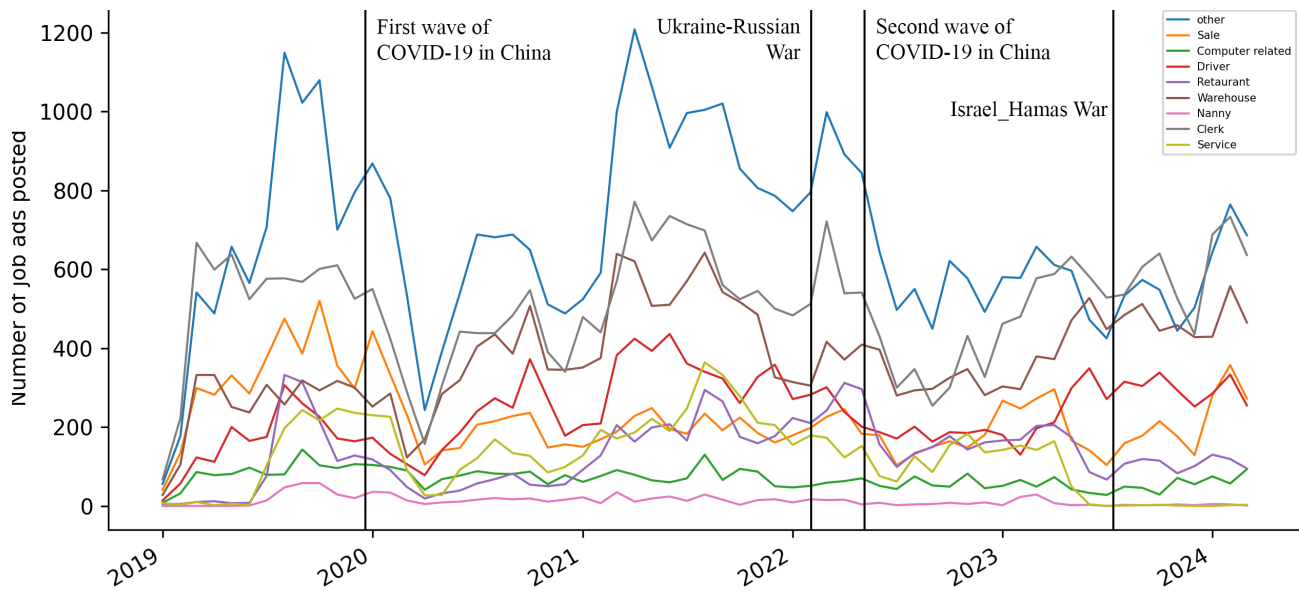


Figure 5: Posting frequencies for jobs ads in different industries from 2019 to 2024

of exploitative offers. Nonetheless, the sheer volume of data captured in real-time presents a compelling snapshot of the job market’s current state, offering valuable insights for the ongoing battle against human trafficking.

Conclusions

In this study, we compiled a dataset of 258,619 job ads from Chinese-speaking online platforms in the US, covering two decades. This dataset offers a fresh perspective on combating online human trafficking, moving beyond the focus on escort ads. Our analysis highlights a widespread use of phone communication in job ads, raising concerns due to the anonymity and untraceability of phone numbers and their potential link to trafficking activities.

We also observed how global events like the COVID-19 pandemic affected job posting patterns, revealing varying degrees of industry resilience to external shocks. Despite facing limitations such as inconsistent data and missing posts, the comprehensive nature and focus on active ads provide significant insights into the job market’s role in human trafficking, suggesting directions for future research and prevention strategies.

Data availability: Open data sharing has revolutionized computational social science research by fostering collaboration, enhancing the quality of research, and enabling the application of advanced computational methods (Davis et al. 2016; Chen, Deb, and Ferrara 2021; Chen and Ferrara 2023). Data access democratization allows researchers to delve into human behavior, social trends, and economic patterns more deeply and accurately and at scale (Chen et al. 2022; Benjamin et al. 2023). It supports interdisciplinary studies, combining fields like sociology and computer science, to tackle complex societal issues affecting vulnerable groups (Rao et al. 2021; Zhou, Luceri, and Ferrara 2023). Making diverse datasets widely available also enables researchers to

apply machine learning and data mining techniques to identify societal threats rapidly and at scale (Sapienza et al. 2018; Tavabi et al. 2019; Luceri et al. 2024).

In that spirit, we decided to make our Job Ads dataset openly available to the research community. We will also continue to maintain and expand it in the future. Looking ahead, we hope other scholars will leverage our Job Ads dataset, alongside other available data, to dive into the complexities of human trafficking within the digital age, to not only enrich our comprehension of this global malaise but also to forge robust, data-informed frameworks for prevention and intervention.

Dataset URL: <https://github.com/serenalyoko/humantrafficking>

Acknowledgements. The authors are grateful to the Many Hopes foundation, a 501(c)3 organization that supports research to combat human trafficking.

References

- ACLU. 2007. Human Trafficking: Modern Enslavement of Immigrant Women in the United States. *American Civil Liberties Union*.
- Bales, K. 2007. What Predicts Human Trafficking? *International Journal of Comparative and Applied Criminal Justice*, 31(2): 269–279.
- Benjamin, D. M.; Morstatter, F.; Abbas, A. E.; Abeliuk, A.; Atanasov, P.; Bennett, S.; Beger, A.; Birari, S.; Budescu, D. V.; Catasta, M.; et al. 2023. Hybrid Forecasting of Geopolitical Events. *AI Magazine*.
- Catanese, S.; Ferrara, E.; and Fiumara, G. 2012. Forensic analysis of phone call networks. *Social Network Analysis and Mining*, 3(1): 15–33.
- Chen, E.; Deb, A.; and Ferrara, E. 2021. #Election2020: The First Public Twitter Dataset on the 2020 US Presidential Election. *Journal of Computational Social Science*, 5: 1–18.

- Chen, E.; and Ferrara, E. 2023. Tweets in Time of Conflict: A Public Dataset Tracking the Twitter Discourse on the War between Ukraine and Russia. In *ICWSM 2023 - 17th International AAAI Conference on Web and Social Media*. arXiv preprint arXiv:2203.07488.
- Chen, E.; Jiang, J.; Chang, H.-C. H.; Muric, G.; and Ferrara, E. 2022. Charting the information and misinformation landscape to characterize misinfodemics on social media: COVID-19 infodemiology study at a planetary scale. *JMIR Infodemiology*, 2(1): e32378.
- Davis, C. A.; Ciampaglia, G. L.; Aiello, L. M.; Chung, K.; Conover, M. D.; Ferrara, E.; Flammini, A.; Fox, G. C.; Gao, X.; Gonçalves, B.; et al. 2016. OSoMe: the IUNI observatory on social media. *PeerJ Computer Science*, 2: e87.
- Ferrara, E.; De Meo, P.; Catanese, S.; and Fiumara, G. 2014. Detecting criminal organizations in mobile phone networks. *Expert Systems with Applications*, 41(13): 5733–5750.
- Fraser, C. 2016. An analysis of the emerging role of social media in human trafficking: Examples from labour and human organ trading. 15(2): 98–112.
- Hundman, K.; Gowda, T.; Kejriwal, M.; and Boecking, B. 2018. Always lurking: Understanding and mitigating bias in online human trafficking detection. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, 137–143.
- Kejriwal, M.; and Szekely, P. 2017a. Information extraction in illicit web domains. In *Proceedings of the 26th international conference on world wide web*, 997–1006.
- Kejriwal, M.; and Szekely, P. 2017b. Knowledge graphs for social good: An entity-centric search engine for the human trafficking domain. *IEEE Transactions on Big Data*, 8(3): 592–606.
- Khan, M. I. 2023. Job Scams in LinkedIn Posts: How to Spot and Avoid Them.
- Laczko, F.; and Gramegna, M. A. 2003. Developing Better Indicators of Human Trafficking. *The Brown Journal of World Affairs*, 10(1): 179–194.
- Lee, M.-C.; Vajiac, C.; Kulshrestha, A.; Levy, S.; Park, N.; Jones, C.; Rabbany, R.; and Faloutsos, C. 2021. INFOSHIELD: Generalizable Information-Theoretic Human-Trafficking Detection. In *2021 IEEE 37th International Conference on Data Engineering (ICDE)*, 1116–1127.
- Luceri, L.; Pantè, V.; Burghardt, K.; and Ferrara, E. 2024. Unmasking the Web of Deceit: Uncovering Coordinated Activity to Expose Information Operations on Twitter. In *WWW'24*.
- Martínez-Mesa, J.; González-Chica, D. A.; Bastos, J. L.; Bonamigo, R. R.; and Duquia, R. P. 2014. Sample size: how many participants do I need in my research? *Anais brasileiros de dermatologia*, 89: 609–615.
- Office to Monitor and Combat Trafficking in Persons. 2023. Trafficking in person report. In *Department of State annual report*.
- Palmquist, K. 2023. 17 Common Job Scams and How To Protect Yourself. *Indeed*.
- Rao, A.; Morstatter, F.; Hu, M.; Chen, E.; Burghardt, K.; Ferrara, E.; and Lerman, K. 2021. Political partisanship and antisense attitudes in online discussions about COVID-19: Twitter content analysis. *Journal of medical Internet research*, 23(6): e26692.
- Sapienza, A.; Ernala, S. K.; Bessi, A.; Lerman, K.; and Ferrara, E. 2018. DISCOVER: Mining Online Chatter for Emerging Cyber Threats. In *Companion of the The Web Conference 2018*, 983–990. International World Wide Web Conferences Steering Committee.
- Szekely, P.; Knoblock, C. A.; Slepicka, J.; Philpot, A.; Singh, A.; Yin, C.; Kapoor, D.; Natarajan, P.; Marcu, D.; Knight, K.; et al. 2015. Building and using a knowledge graph to combat human trafficking. In *The Semantic Web-ISWC 2015: 14th International Semantic Web Conference, Bethlehem, PA, USA, October 11-15, 2015, Proceedings, Part II 14*, 205–221. Springer.
- Tavabi, N.; Bartley, N.; Abeliuk, A.; Soni, S.; Ferrara, E.; and Lerman, K. 2019. Characterizing Activity on the Deep and Dark Web. In *Companion Proceedings of the 2019 World Wide Web Conference*, 206–213.
- Tobey, M.; Li, R.; Özalın, O. Y.; Mayorga, M. E.; and Caltagirone, S. 2024. Interpretable models for the automated detection of human trafficking in illicit massage businesses. *IJSE Transactions*, 56(3): 311–324.
- Tong, E.; Zadeh, A.; Jones, C.; and Morency, L.-P. 2017. Combating Human Trafficking with Multimodal Deep Models. In Barzilay, R.; and Kan, M.-Y., eds., *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 1547–1556. Vancouver, Canada: Association for Computational Linguistics.
- UN. 2003. Report of the Ad Hoc Committee on the Elaboration of a Convention against Transnational Organized Crime on the work of its first to eleventh sessions. *United Nation*.
- UNODC. 2024. Global Report on Trafficking in Persons. *United Nation*.
- Whitney, J.; Jennex, M.; Elkins, A.; and Frost, E. 2018. Don't want to get caught? don't say it: The use of emojis in online human sex trafficking ads. *Hawaii International Conference on System Sciences*.
- Zhou, S.; Luceri, L.; and Ferrara, E. 2023. Unveiling the Dynamics of Censorship, COVID-19 Regulations, and Protest: An Empirical Study of Chinese Subreddit r/china.irl. In *ICWSM 2023 Companion Proceedings*. arXiv preprint arXiv:2304.02800.
- Zhu, J.; Li, L.; and Jones, C. 2019. Identification and Detection of Human Trafficking Using Language Models. In *2019 European Intelligence and Security Informatics Conference (EISIC)*, 24–31.