TRAFFICVIS: Fighting Human Trafficking through Visualization

Catalina Vajiac* Andreas Olligschlaeger† Yifei Li‡ Pratheeksha Nair‡ Meng-Chieh Lee*
Namyong Park* Reihaneh Rabbany‡ Duen Horng Chau§ Christos Faloutsos*

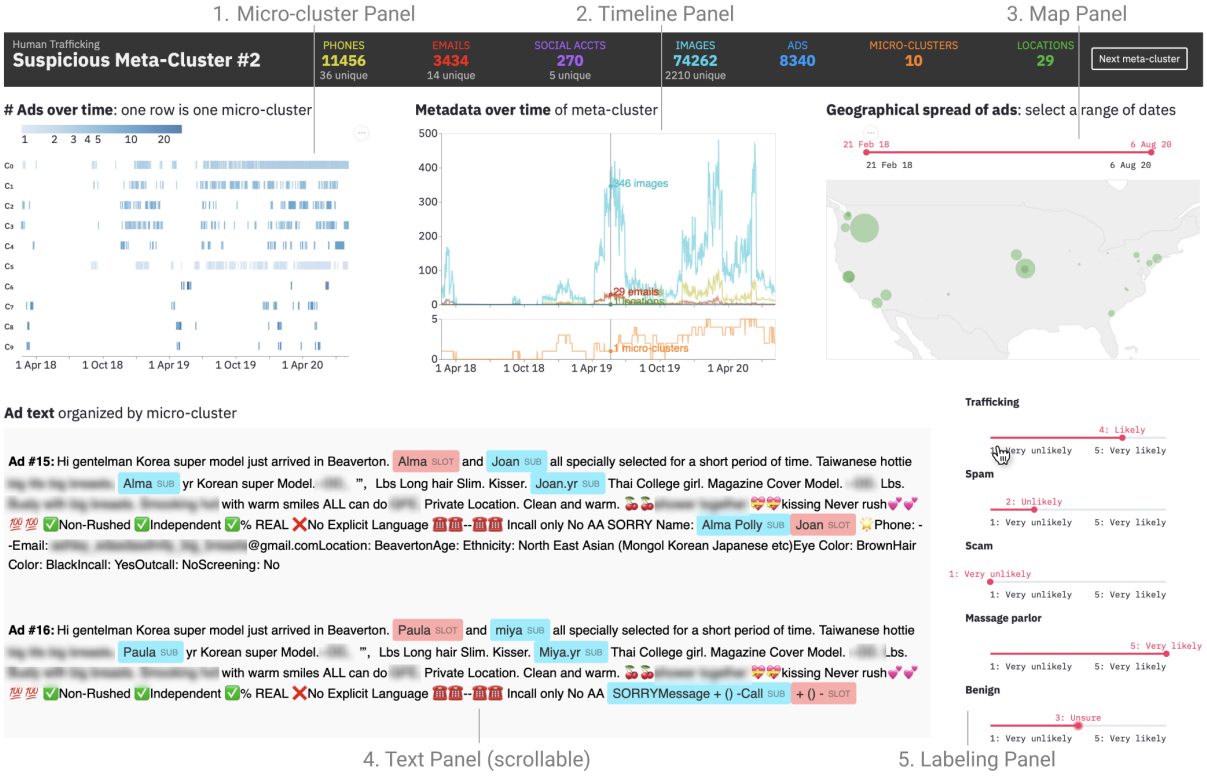

Figure 1: TRAFFICVIS showing a real meta-cluster (connected using metadata) from online escort ads (some text blurred for privacy). **1.** ***Micro-cluster panel*** shows daily activity for each micro-(text) cluster over the lifetime of the meta-cluster. **2.** ***Timeline panel*** shows metadata activity over time. **3.** ***Map panel*** shows the geographical spread. **4.** ***Text panel*** shows the escort ad text. **5.** ***Labeling panel*** allows the domain expert to label the likelihood that this meta-cluster falls under each modus operandi (M.O.).

ABSTRACT

Law enforcement can detect human trafficking (HT) in online escort websites by analyzing suspicious clusters of connected ads. Given such clusters, how can we interactively visualize potential evidence for law enforcement and domain experts? We present TRAFFICVIS, which, to our knowledge, is the first interface for cluster-level HT detection and labeling. It builds on state-of-the-art HT clustering algorithms by incorporating metadata as a signal of organized and potentially suspicious activity. Also, domain experts can label clusters as HT, spam, and more, efficiently creating labeled datasets to enable further HT research. TRAFFICVIS has been built in close collaboration with domain experts, who estimate that TRAFFICVIS provides a **median 36x speedup** over manual labeling.

Keywords: visualization, social good, anti human trafficking

*Carnegie Mellon: {cvajiac, mengchil, namyongp, christos}@cmu.edu
†Marinus Analytics: olli@marinusanalytics.com
‡McGill MILA: {yifei.li, pratheeksha.nair, rrabba}@mail.mcgill.ca
§Georgia Institute of Technology: polo@gatech.edu

1 INTRODUCTION

Human trafficking (HT) for forced sexual exploitation is a pervasive societal problem that affects over 4.8 million people world-wide [1], and the majority of HT victims are advertised on online escort websites [6]. However, legitimate sex workers also post on these sites, and law enforcement agencies are focused on separating HT rings from legitimate sex workers. There is one critical insight to detecting HT; since traffickers entirely control the ad content for their victims [6], ads posted by the same trafficker tend to be similar. This insight has already been exploited algorithmically in the state-of-the-art HT detection algorithm, InfoShield [2], as well as a few other methods [3, 5]. However, the problem is more complex; domain experts have recently discovered additional modus operandi (M.O.s) in the escort ad space. For example, there are *spam* ads with fake contact information that are posted to flood the ad space, and *scam* ads that ask for prepayment to trick customers out of their money. *Massage parlor* ads are also commonly found, which can be, but are not necessarily, hubs for HT. Occasionally, a cluster can be *benign*, not signifying any suspicious behavior.

The presence of these M.O.s have made HT detection more difficult for law enforcement. In particular, spam ads flood escort websites and often look like leads at first glance. Features of these M.O.s are not yet well known, as it is currently prohibitively time-

consuming for domain experts to label clusters. However, having good cluster labels would enable the development and evaluation of M.O. classification algorithms.

Previous work has focused on finding *micro-clusters* of ads based on text. To our knowledge, no published work has been previously done to leverage advertisement metadata by linking micro-clusters into larger organized activities or to visualize them for domain experts. As a first step, we propose TRAFFICVIS, an interactive application for domain experts to visually inspect suspicious meta-clusters and label their likelihood to be a particular M.O. In particular, this ongoing work makes the following contributions:

1. **First of its kind:** TRAFFICVIS is the first interactive application for HT detection and labeling. It will be used to curate labels for future use in developing M.O. classification algorithms. However, without *Labeling panel*, TRAFFICVIS can also be used as a standalone application to enable domain experts or law enforcement to inspect suspicious meta-clusters.
2. **Meta-clustering:** We improve upon the state-of-the-art by connecting text-based micro-clusters into *meta-clusters* based on metadata, such as contact information and images used.
3. **Time-saving:** Through regularly incorporating domain expert's suggestions, TRAFFICVIS will give an estimated **36x speedup** in ad labeling, finally making labeling scalable.

2 DESIGNING TRAFFICVIS

2.1 Meta-Clustering

Since micro-clusters are constructed only using text features, multiple micro-clusters can actually be part of the same activity. Therefore, we connect micro-clusters (c_i) into *meta-clusters* (M_j) based on extracted metadata, such as images, emails, phone numbers, and social media accounts. We consider two micro-clusters c_1, c_2 to be part of the same meta-cluster M_j if any ads $a_m \in c_1, a_n \in c_2$ share at least one metadata field. Figure 2 shows an example of how six micro-clusters can be connected into three meta-clusters.

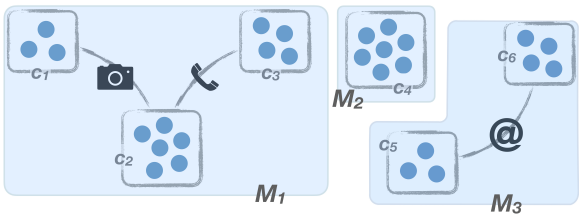

Figure 2: *From micro-clusters* (c_i) *to meta-clusters* (M_i)*:* By incorporating metadata such as images, phone numbers, and social media accounts, we combine 6 micro-clusters into 3 meta-clusters, each of which are part of the same group.

2.2 Design and Usage Scenario

Figure 1 demonstrates how an analyst (e.g., law enforcement agent, HT domain expert) may use TRAFFICVIS to investigate meta-clusters. First, the analyst sees basic statistics on the top banner, observing that for over 8000 ads, there are relatively few unique metadata fields. She then moves to the *Micro-cluster panel* to inspect the top 10 micro-clusters with the most activity (chosen due to space constraints) per day, throughout the lifetime of the meta-cluster. By hovering over a particular cell, a tooltip displays the number of ads per day, which the analyst can use to investigate the consistent volume of ads in c_0 during the second half of the meta-cluster.

Next, the analyst can look at metadata usage and the number of micro-clusters per day in the *Timeline panel*. By hovering over any date, the time-series values will be displayed. Since the number of micro-clusters is a feature derived from InfoShield, display it separately. She may notice that there are few unique locations per day and want to know what they are. In the *Map panel*, she sees the geographical spread of the meta-cluster. The analyst can also use the range slider to only look at the spread over a couple of days.

Next, she inspects the text of each ad in the scrollable *Text panel*, noticing that most of the ad text is the same, but many different people are advertised, as highlighted by InfoShield, which is a possible sign of HT and massage parlors.

Finally, after inspecting the information in each panel, the analyst uses the *Labeling panel* to label the meta-cluster on a scale of 1 (very unlikely) to 5 (very likely) for each possible M.O. Upon clicking the 'Next meta-cluster' button on the top right corner, these labels are saved and a new meta-cluster is displayed.

2.3 Preliminary results: Feedback & Speedup

In designing TRAFFICVIS, we met weekly with domain experts and incorporated their feedback in order to ensure we developed a tool they would actually use. For example, TRAFFICVIS is a widescreen application because domain experts prefer seeing all panels at once and will not use this interface on mobile devices.

According to domain experts, it takes one person 8 hours to manually label 500 ads [4]. Therefore, just to label the meta-cluster from Figure 1 with 8340 ads, it would take one person 133 hours. Using TRAFFICVIS, we estimate a domain expert could label the same meta-cluster in 5 minutes. Therefore, TRAFFICVIS would provide a 1600x speedup for this meta-cluster. Furthermore, we have applied TRAFFICVIS on data containing suspected spam and trafficking ads. Resulting meta-clusters range from 38 to 47,786 ads, with a median of 227 ads. It would take one person from 36 minutes to 764 hours, median 3 hours, to label these clusters. Thus, we give a **median 36x speedup**, but speedup range from 7.3x to over 9000x.

3 CONCLUSION AND ONGOING WORK

TRAFFICVIS makes three major contributions: **1. First of its kind** interactive application for HT detection and labeling; **2. Meta-clustering** to connect ads under the same organized activity by building on previous results; and **3. Time-saving**, granting a *median 36x speedup* over manual labeling. So far, we have collaborated with domain experts with technical experience. Next, we will conduct formal user studies with a diverse set of users to evaluate efficacy and usability. More specifically, we will compare manual labelers and TRAFFICVIS labelers. Each group will label the same ads, and we will study elapsed time and labeling agreement between them.

Furthermore, this user study will naturally curate a dataset labeled by TRAFFICVIS that will enable further algorithm development towards M.O. detection, allowing law enforcement to quickly find meta-clusters of ads that actually represent real HT cases. *Marinus Analytics* has expressed interest in incorporating TRAFFICVIS in their pipeline. This will enable researchers to continually develop and evaluate M.O. detection algorithms as we see emerging trends in escort ads over the years to come.

REFERENCES

[1] Ilo global estimate of forced labour. Available: Link, 2012.

[2] M.-C. Lee, C. Vajiac, A. Kulshrestha, S. Levy, N. Park, C. Jones, R. Rabbany, and C. Faloutsos. InfoShield: generalizable information-theoretic human-trafficking detection. In *IEEE ICDE*. IEEE, 2021.

[3] L. Li, O. Simek, A. Lai, M. P. Daggett, C. K. Dagli, and C. Jones. Detection and characterization of human trafficking networks using unsupervised scalable text template matching. In *IEEE BigData*, pages 3111–3120. IEEE, 2018.

[4] A. Olligschlager and C. Jones. Personal communication.

[5] R. Rabbany, D. Bayani, and A. Dubrawski. Active search of connections for case building and combating human trafficking. In *KDD*, pages 2120–2129. ACM, 2018.

[6] Thorn. Report on the use of technology to recruit, groom and sell domestic minor sex trafficking victims. Available: Link, 2015.