

Sex Trafficking Detection with Ordinal Regression Neural Networks

Longshaokan Wang^{1*}, Eric Laber², Yeng Saanchi³, Sherrie Caltagirone⁴

¹Alexa AI, Amazon, ²³Department of Statistics, North Carolina State University, ⁴Global Emancipation Network
¹longsha@amazon.com, ²³{eblaber, ysaanch}@ncsu.edu, ⁴sherrie@globalemancipation.ngo

Abstract

Sex trafficking is a global epidemic. Escort websites are a primary vehicle for selling the services of such trafficking victims and thus a major driver of trafficker revenue. Many law enforcement agencies do not have the resources to manually identify leads from the millions of escort ads posted across dozens of public websites. We propose an ordinal regression neural network to identify escort ads that are likely linked to sex trafficking. Our model uses a modified cost function to mitigate inconsistencies in predictions often associated with nonparametric ordinal regression and leverages recent advancements in deep learning to improve prediction accuracy. The proposed method significantly improves on the previous state-of-the-art on Trafficking-10K, an expert-annotated dataset of escort ads. Additionally, because traffickers use acronyms, deliberate typographical errors, and emojis to replace explicit keywords, we demonstrate how to expand the lexicon of trafficking flags through word embeddings and t-SNE.

1 Introduction

Globally, human trafficking is one of the fastest growing crimes and, with annual profits estimated to be in excess of 150 billion USD, it is also among the most lucrative (Amin 2010). Sex trafficking is a form of human trafficking which involves sexual exploitation through coercion. Recent estimates suggest that nearly 4 million adults and 1 million children are being victimized globally on any given day; furthermore, it is estimated that 99 percent of victims are female (International Labour Organization, Walk Free Foundation, and International Organization for Migration 2017). Escort websites are an increasingly popular vehicle for selling the services of trafficking victims. According to a recent survivor survey (THORN and Bouché 2018), 38% of underage trafficking victims who were enslaved prior to 2004 were advertised online, and that number rose to 75% for those enslaved after 2004. Prior to its shutdown in April 2018, the website Backpage was the most frequently used online advertising platform; other popular websites used for advertis-

ing escort service include Craigslist, Redbook, SugarDaddy, and Facebook (THORN and Bouché 2018). Despite the seizure of Backpage, there were nearly 150,000 new online sex advertisements posted *per day* in the U.S. alone in late 2018 (Tarinelli 2018); even with many of these new ads being re-posts of existing ads and traffickers often posting multiple ads for the same victims (THORN and Bouché 2018), this volume is staggering.

Because of their ubiquity and public access, escort websites are a rich resource for anti-trafficking operations. However, many law enforcement agencies do not have the resources to sift through the volume of escort ads to identify those coming from potential traffickers. One scalable and efficient solution is to build a statistical model to predict the likelihood of an ad coming from a trafficker using a dataset annotated by anti-trafficking experts. We propose an ordinal regression neural network tailored for text input. This model comprises three components: (i) a Word2Vec model (Mikolov et al. 2013b) that maps each word from the text input to a numeric vector, (ii) a gated-feedback recurrent neural network (Chung et al. 2015) that sequentially processes the word vectors, and (iii) an ordinal regression layer (Cheng, Wang, and Pollastri 2008) that produces a predicted ordinal label. We use a modified cost function to mitigate inconsistencies in predictions associated with nonparametric ordinal regression. We also leverage several regularization techniques for deep neural networks to further improve model performance, such as residual connection (He et al. 2016) and batch normalization (Ioffe and Szegedy 2015). We conduct our experiments on Trafficking-10k (Tong et al. 2017), a dataset of escort ads for which anti-trafficking experts assigned each sample one of seven ordered labels ranging from “1: Very Unlikely (to come from traffickers)” to “7: Very Likely”. Our proposed model significantly outperforms previously published models (Tong et al. 2017) on Trafficking-10k as well as a variety of baseline ordinal regression models. In addition, we analyze the emojis used in escort ads with Word2Vec and t-SNE (van der Maaten and Hinton 2008), and we show that the lexicon of trafficking-related emojis can be subsequently expanded.

The main contributions of this paper are summarized as follows: 1. We propose a neural network architecture for

*This work was done when Wang was a PhD student at North Carolina State University.
Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

text data with ordinal labels, which outperforms the previous state-of-the-art (Tong et al. 2017) on Trafficking-10k. 2. We propose a simple penalty term in the cost function to mitigate the monotonicity violation and improve the interpretability of the output of the ordinal regression layer, where the monotonicity violation was previously deemed too computationally costly to resolve (Niu et al. 2016). 3. We provide a qualitative analysis on the top escort ads flagged by our model, which offers patterns that anti-trafficking experts can potentially confirm or make use of. 4. We provide an emoji analysis that shows how to use unsupervised learning techniques on the raw data to generate leads on new trafficking key words. 5. We open source our code base and trained model to encourage further research on trafficking detection and to allow the law enforcement to make use of our research for free.

In Section 2, we discuss related work on human trafficking detection and ordinal regression. In Section 3, we present our proposed model and detail its components. In Section 4, we present the experimental results, including the Trafficking-10K benchmark, a qualitative analysis of the predictions on raw data, and the emoji analysis. In Section 5, we summarize our findings and discuss future work.

2 Related Work

Trafficking detection: There have been several software products designed to aid anti-trafficking efforts. Examples include Memex¹ which focuses on search functionalities in the dark web; Spotlight² which flags suspicious ads and links images appearing in multiple ads; Traffic Jam³ which seeks to identify patterns that connect multiple ads to the same trafficking organization; and TraffickCam⁴ which aims to construct a crowd-sourced database of hotel room images to geo-locate victims. These research efforts have largely been isolated, and few research articles on machine learning for trafficking detection have been published. Closest to our work is the Human Trafficking Deep Network (HTDN) (Tong et al. 2017). HTDN has three main components: a language network that uses pretrained word embeddings and a long short-term memory network (LSTM) to process text input; a vision network that uses a convolutional network to process image input; and another convolutional network to combine the output of the previous two networks and produce a binary classification. Compared to the language network in HTDN, our model replaces LSTM with a gated-feedback recurrent neural network, adopts certain regularizations, and uses an ordinal regression layer on top. It significantly improves HTDN’s benchmark despite only using text input. As in the work of Tong et al. (2017), we pre-train word embeddings using a skip-gram model (Mikolov et al. 2013b) applied to unlabeled data from escort ads, however, we go further by analyzing the emojis’ embeddings and thereby expand the trafficking lexicon.

¹darpa.mil/program/memex

²htspotlight.com

³marinusanalytics.com/trafficjam

⁴traffickcam.com

Ordinal regression: We briefly review ordinal regression before introducing the proposed methodology. We assume that the training data are $D_{\text{train}} = \{(\mathbf{X}_i, Y_i)\}_{i=1}^n$, where $\mathbf{X}_i \in \mathcal{X}$ are the features and $Y_i \in \mathcal{Y}$ is the response; \mathcal{Y} is the set of k ordered labels $\{1, 2, \dots, k\}$ with $1 \prec 2 \dots \prec k$. Many ordinal regression methods learn a composite map $\eta = h \circ g$, where $g : \mathcal{X} \rightarrow \mathbb{R}$ and $h : \mathbb{R} \rightarrow \{1, 2, \dots, k\}$ have the interpretation that $g(\mathbf{X})$ is a latent “score” which is subsequently discretized into a category by h . η is often estimated by empirical risk minimization, i.e., by minimizing a loss function $C\{\eta(\mathbf{X}), Y\}$ averaged over the training data. Standard choices of η and C are reviewed by Rennie and Srebro (2005).

Another common approach to ordinal regression, which we adopt in our proposed method, is to transform the label prediction into a series of $k - 1$ binary classification sub-problems, wherein the i th sub-problem is to predict whether the true label exceeds i (Frank and Hall 2001; Li and Lin 2006). For example, one might use a series of logistic regression models to estimate the conditional probabilities $f_i(\mathbf{X}) = P(Y > i | \mathbf{X})$ for each $i = 1, \dots, k - 1$. Cheng, Wang, and Pollastri (2008) estimated these probabilities jointly using a neural network; this was later extended to image data (Niu et al. 2016) as well as text data (Irsay and Cardie 2015; Ruder, Ghaffari, and Breslin 2016). However, as acknowledged by Cheng, Wang, and Pollastri (2008), the estimated probabilities need not respect the ordering $f_i(\mathbf{X}) \geq f_{i+1}(\mathbf{X})$ for all i and \mathbf{X} . We force our estimator to respect this ordering through a penalty on its violation.

3 Method

Our proposed ordinal regression model consists of the following three components: Word embeddings pre-trained by a Skip-gram model, a gated-feedback recurrent neural network that constructs summary features from sentences, and a multi-labeled logistic regression layer tailored for ordinal regression. See Figure 1 for a schematic. The details of its components and their respective alternatives are discussed below.

3.1 Word Embeddings

Vector representations of words, also known as word embeddings, can be obtained through unsupervised learning on a large text corpus so that certain linguistic regularities and patterns are encoded. Compared to Latent Semantic Analysis (Dumais 2004), embedding algorithms using neural networks are particularly good at preserving linear regularities among words in addition to grouping similar words together (Mikolov et al. 2013a). Such embeddings can in turn help other algorithms achieve better performances in various natural language processing tasks (Mikolov et al. 2013b).

Unfortunately, the escort ads contain a plethora of emojis, acronyms, and (sometimes deliberate) typographical errors that are not encountered in more standard text data, which suggests that it is likely better to learn word embeddings from scratch on a large collection of escort ads instead of using previously published embeddings (Tong et al. 2017).




Figure 1: Overview of the ordinal regression neural network for text input. \mathbf{H} represents a hidden state in a gated-feedback recurrent neural network.

We use 168,337 ads scraped from Backpage as our training corpus and the Skip-gram model with Negative sampling (Mikolov et al. 2013b) as our model.

3.2 Gated-Feedback Recurrent Neural Network

To process entire sentences and paragraphs after mapping the words to embeddings, we need a model to handle sequential data. Recurrent neural networks (RNNs) have recently seen great success at modeling sequential data, especially in natural language processing tasks (LeCun, Bengio, and Hinton 2015). On a high level, an RNN is a neural network that processes a sequence of inputs one at a time, taking the summary of the sequence seen so far from the previous time point as an additional input and producing a summary for the next time point. One of the most widely used variations of RNNs, a Long short-term memory network (LSTM), uses various gates to control the information flow and is able to better preserve long-term dependencies in the running summary compared to a basic RNN (see Goodfellow, Bengio, and Courville 2016 and references therein). In our implementation, we use a further refinement of multi-layered LSTMs, Gated-feedback recurrent neural networks (GF-RNNs), which tend to capture dependencies across different timescales more easily (Chung et al. 2015).

Regularization techniques for neural networks including Dropout (Srivastava et al. 2014), Residual connection (He et al. 2016), and Batch normalization (Ioffe and Szegedy 2015) are added to GF-RNN for further improvements.

After GF-RNN processes an entire escort ad, the average of the hidden states of the last layer becomes the input for the multi-labeled logistic regression layer which we discuss next.

3.3 Multi-Labeled Logistic Regression Layer

As noted previously, the ordinal regression problem can be cast into a series of binary classification problems and thereby utilize the large repository of available classification algorithms (Frank and Hall 2001; Li and Lin 2006; Niu et al. 2016). One formulation is as follows. Given k total ranks, the i -th binary classifier is trained to predict the probability that a sample \mathbf{X} has rank larger than i : $\hat{f}_i(\mathbf{X}) = \hat{\mathbb{P}}(Y > i|\mathbf{X})$. Then the predicted rank is

$$\hat{Y} = 1 + \sum_{i=1}^{k-1} \text{Round} \left\{ \hat{f}_i(\mathbf{X}) \right\}.$$

In a classification task, the final layer of a deep neural network is typically a softmax layer with dimension equal to the number of classes (Goodfellow, Bengio, and Courville 2016). Using the ordinal-regression-to-binary-classifications formulation described above, Cheng, Wang, and Pollastri (2008) replaced the softmax layer in their neural network with a $(k-1)$ -dimensional sigmoid layer, where each neuron serves as a binary classifier (see Figure 2 but without the order penalty to be discussed later).

With the sigmoid activation function, the output of the i th neuron can be viewed as the predicted probability that the sample has rank greater⁵ than i . Alternatively, the entire sigmoid layer can be viewed as performing multi-labeled logistic regression, where the i th label is the indicator of the sample's rank being greater than i . The training data are thus re-formatted accordingly so that response variable for a sample with rank i becomes $(\mathbf{1}_{i-1}^T, \mathbf{0}_{k-i}^T)^T$. The $k-1$ binary classifiers share the features constructed by the earlier layers of the neural network and can be trained jointly with mean squared error loss. A key difference between the multi-labeled logistic regression and the naive classification (ignoring the order and treating all ranks as separate classes) is that the loss for $\hat{Y} \neq Y$ is constant in the naive classification but proportional to $|\hat{Y} - Y|$ in the multi-labeled logistic regression.

Cheng, Wang, and Pollastri's (2008) final layer was preceded by a simple feed-forward network. In our case, word embeddings and GF-RNN allow us to construct a feature vector of fixed length from text input, so we can simply attach the multi-labeled logistic regression layer to the output of GF-RNN to complete an ordinal regression neural network for text input.

The violation of the monotonicity in the estimated probabilities (e.g., $\hat{f}_i(\mathbf{X}) < \hat{f}_{i+1}(\mathbf{X})$ for some \mathbf{X} and i) has remained an open issue since the original ordinal regression neural network proposal of Cheng, Wang, and Pollastri (2008). This is perhaps owed in part to the belief that correcting this issue would significantly increase training complexity (Niu et al. 2016). We propose an effective and computationally efficient solution to avoid the conflicting predictions

⁵Actually, in Cheng, Wang, and Pollastri's original formulation, the final layer is k -dimensional with the i -th neuron predicting the probability that the sample has rank greater than or equal to i . This is redundant because the first neuron should always be equal to 1. Hence we make the slight adjustment of using only $k-1$ neurons.




Figure 2: Ordinal regression layer with order penalty.

as follows: penalize such conflicts in the training phase by adding

$$P(\mathbf{X}; \lambda) = \lambda \sum_{i=1}^{k-2} \max \left\{ \hat{f}_{i+1}(\mathbf{X}) - \hat{f}_i(\mathbf{X}), 0 \right\}$$

to the loss function for a sample \mathbf{X} , where λ is a penalty parameter (Figure 2). For sufficiently large λ the estimated probabilities will respect the monotonicity condition; respecting this condition improves the interpretability of the predictions, which is vital in applications like the one we consider here as stakeholders are given the estimated probabilities. We also hypothesize that the order penalty may serve as a regularizer to improve each binary classifier (see the ablation test in Section 4.3).

All three components of our model (word embeddings, GF-RNN, and multi-labeled logistic regression layer) can be trained jointly, with word embeddings optionally held fixed or given a smaller learning rate for fine-tuning. The hyperparameters for all components are given in the Appendix. They are selected according to either literature or grid-search.

4 Experiments

We first describe the datasets we use to train and evaluate our models. Then we present a detailed comparison of our proposed model with commonly used ordinal regression models as well as the previous state-of-the-art classification model by Tong et al. (2017). To assess the effect of each component in our model, we perform an ablation test where the components are swapped by their more standard alternatives one at a time. Next, we perform a qualitative analysis on the model predictions on the raw data, which are scraped from a different escort website than the one that provides the labeled training data. Finally, we conduct an emoji analysis using the word embeddings trained on raw escort ads.

4.1 Datasets

We use raw texts scraped from Backpage and TNABoard to pre-train the word embeddings, and use the same labeled texts Tong et al. (2017) used to conduct model comparisons. The raw text dataset consists of 44,105 ads from TNABoard and 124,220 ads from Backpage. Data cleaning/preprocessing includes joining the title and the body of an ad; adding white spaces around every emoji so that it can

be tokenized properly; stripping tabs, line breaks, punctuations, and extra white spaces; removing phone numbers; and converting all letters to lower case. We have ensured that the raw dataset has no overlap with the labeled dataset to avoid bias in test accuracy. While it is possible to scrape more raw data, we did not observe significant improvements in model performances when the size of raw data increased from $\sim 70,000$ to $\sim 170,000$, hence we assume that the current raw dataset is sufficiently large.

The labeled dataset is called Trafficking-10k. It consists of 12,350 ads from Backpage labeled by experts in human trafficking detection⁶ (Tong et al. 2017). Each label is one of seven ordered levels of likelihood that the corresponding ad comes from a human trafficker. Descriptions and sample proportions of the labels are in Table 1. The original Trafficking-10K includes both texts and images, but as mentioned in Section 1, only the texts are used in our case. We apply the same preprocessing to Trafficking-10k as we do to raw data.

4.2 Comparison with Baselines

We compare our proposed ordinal regression neural network (ORNN) to Immediate-Threshold ordinal logistic regression (IT) (Rennie and Srebro 2005), All-Threshold ordinal logistic regression (AT) (Rennie and Srebro 2005), Least Absolute Deviation (LAD) (Bloomfield and Steiger 1980; Narula and Wellington 1982), and multi-class logistic regression (MC) which ignores the ordering. The primary evaluation metrics are Mean Absolute Error (MAE) and macro-averaged Mean Absolute Error (MAE^M) (Baccianella, E-suli, and Sebastiani 2009). To compare our model with the previous state-of-the-art classification model for escort ads, the Human Trafficking Deep Network (HTDN) (Tong et al. 2017), we also polarize the true and predicted labels into two classes, “1-4: Unlikely” and “5-7: Likely”; then we compute the binary classification accuracy (Acc.) as well as the weighted binary classification accuracy (Wt. Acc.) given by

$$\text{Wt. Acc.} = \frac{1}{2} \left(\frac{\text{True Positives}}{\text{Total Positives}} + \frac{\text{True Negatives}}{\text{Total Negatives}} \right).$$

Note that for applications in human trafficking detection, MAE and Acc. are of primary interest. Whereas for a more general comparison among the models, the class imbalance robust metrics, MAE^M and Wt. Acc., might be more suitable. Bootstrapping or increasing the weight of samples in smaller classes can improve MAE^M and Wt. Acc. at the cost of MAE and Acc..

The text data need to be vectorized before they can be fed into the baseline models (whereas vectorization is built into ORNN). The standard practice is to tokenize the texts using n-grams and then create weighted term frequency vec-

⁶Backpage was seized by FBI in April 2018, but we have observed that escort ads across different websites are often similar, and a survivor survey shows that traffickers post their ads on multiple websites (THORN and Bouché 2018). Thus, we argue that the training data from Backpage are still useful, which is empirically supported by our qualitative analysis in Section 4.4.

Label	1	2	3	4	5	6	7
Description	Strongly Unlikely	Unlikely	Slightly Unlikely	Unsure	Weakly Likely	Likely	Strongly Likely
Count	1,977	1,904	3,619	796	3,515	457	82

Table 1: Description and distribution of labels in Trafficking-10K.

tors using the term frequency (TF)-inverse document frequency (IDF) scheme (Beel et al. 2016; Manning, Raghavan, and Schütze 2009). The specific variation we use is the recommended unigram + sublinear TF + smooth IDF (Manning, Raghavan, and Schütze 2009; Pedregosa et al. 2011). Dimension reduction techniques such as Latent Semantic Analysis (Dumais 2004) can be optionally applied to the frequency vectors, but Schuller, Mousa, and Vryniotis (2015) concluded from their experiments that dimension reduction on frequency vectors actually hurts model performance, which our preliminary experiments agree with.

All models are trained and evaluated using the same (w.r.t. data shuffle and split) 10-fold cross-validation (CV) on Trafficking-10k, except for HTDN, whose result is read from the original paper (Tong et al. 2017)⁷. During each train-test split, 2/9 of the training set is further reserved as the validation set for tuning hyperparameters such as L2-penalty in IT, AT and LAD, and learning rate in ORNN. So the overall train-validation-test ratio is 70%-20%-10%. We report the mean metrics from the CV in Table 2. As previous research has pointed out that there is no unbiased estimator of the variance of CV (Bengio and Grandvalet 2004), we report the naive standard error treating metrics across CV as independent. Recall that a 95% confidence interval is roughly the point estimate $\pm 1.96 \times$ the standard error.

We can see that ORNN has the best MAE, MAE^M and Acc. as well as a close 2nd best Wt. Acc. among all models. Its Wt. Acc. is a substantial improvement over HTDN despite the fact that the latter use both text and image data. It is important to note that HTDN is trained using binary labels, whereas the other models are trained using ordinal labels and then have their ordinal predictions converted to binary predictions. This is most likely the reason that even the baseline models except for LAD can yield better Wt. Acc. than HTDN, confirming our earlier claim that polarizing the ordinal labels during training may lead to information loss.

4.3 Ablation Test

To ensure that we do not unnecessarily complicate our ORNN model, and to assess the impact of each component on the final model performance, we perform an ablation test. Using the same CV and evaluation metrics, we make the following replacements separately and re-evaluate the model: 1. Replace word embeddings pre-trained from skip-gram model with randomly initialized word embeddings; 2. replace gated-feedback recurrent neural network with long short-term memory network (LSTM); 3. disable batch normalization; 4. disable residual connection; 5. replace the multi-

⁷The authors of HTDN used a single train-validation-test split instead of CV.

labeled logistic regression layer with a softmax layer (i.e., let the model perform classification, treating the ordinal response variable as a categorical variable with k classes); 6. replace the multi-labeled logistic regression layer with a 1-dimensional linear layer (i.e., let the model perform regression, treating the ordinal response variable as a continuous variable) and round the prediction to the nearest integer during testing; 7. set the order penalty to 0. The results are shown in Table 3.

The proposed ORNN once again has all the best metrics except for Wt. Acc. which is the 2nd best. Note that if we disregard the ordinal labels and perform classification or regression, MAE falls off by a large margin. Setting order penalty to 0 does not deteriorate the performance by much, however, the percent of conflicting binary predictions (see Section 3.3) rises from 1.4% to 5.2%. So adding an order penalty helps produce more interpretable results⁸.

4.4 Qualitative Analysis of Predictions

To qualitatively evaluate how well our model predicts on raw data and observe potential patterns in the flagged samples, we obtain predictions on the 44,105 unlabelled ads from TNABoard with the ORNN model trained on Trafficking-10k, then we examine the samples with high predicted likelihood to come from traffickers. Below are the top three samples that the model considers likely:

- “amazing reviewed crystal only here till fri book now please check our site for the services the girls provide all updates specials photos rates reviews njfantasygirls ...look who s back amazing reviewed model samantha... brand new spinner jessica special rate today 250 hr 21 5 4 120 34b total gfe total anything goes no limits...”
- “2 hot toght 18y o spinners 4 amazing providers today specials...”
- “asian college girl is visiting bellevue service type escort hair color brown eyes brown age 23 height 5 4 body type slim cup size c cup ethnicity asian service type escort i am here for you settle men i am a tiny asian girl who is waiting for a gentlemen...”

Some interesting patterns in the samples with high predicted likelihood (here we only showed three) include: mentioning of multiple names or > 1 providers in a single ad; possibly intentional typos and abbreviations for the sensitive words such as “tight” \rightarrow “tought” and “18 year old” \rightarrow “18y o”; keywords that indicate traveling of the providers such as “till fri”, “look who s back”, and “visiting”; keywords that hint

⁸It is possible to increase the order penalty to further reduce or eliminate conflicting predictions, but we find that a large order penalty harms model performance.

Model	MAE	MAE ^M	Acc.	Wt. Acc.
ORNN	0.769 (0.009)	1.238 (0.016)	0.818 (0.003)	0.772 (0.004)
IT	0.807 (0.010)	1.244 (0.011)	0.801 (0.003)	0.781 (0.004)
AT	0.778 (0.009)	1.246 (0.012)	0.813 (0.003)	0.755 (0.004)
LAD	0.829 (0.008)	1.298 (0.016)	0.786 (0.004)	0.686 (0.003)
MC	0.794 (0.012)	1.286 (0.018)	0.804 (0.003)	0.767 (0.004)
HTDN	-	-	0.800	0.753

Table 2: Comparison of the proposed ordinal regression neural network (ORNN) against Immediate-Threshold ordinal logistic regression (IT), All-Threshold ordinal logistic regression (AT), Least Absolute Deviation (LAD), multi-class logistic regression (MC), and the Human Trafficking Deep Network (HTDN) in terms of Mean Absolute Error (MAE), macro-averaged Mean Absolute Error (MAE^M), binary classification accuracy (Acc.) and weighted binary classification accuracy (Wt. Acc.). The results are averaged across 10-fold CV on Trafficking-10k with naive standard errors in the parentheses. The best and second best results are highlighted.

Model	MAE	MAE ^M	Acc.	Wt. Acc.
0. Proposed ORNN	0.769 (0.009)	1.238 (0.016)	0.818 (0.003)	0.772 (0.004)
1. Random Embeddings	0.789 (0.007)	1.254 (0.013)	0.810 (0.002)	0.757 (0.003)
2. LSTM	0.778 (0.009)	1.261 (0.021)	0.815 (0.003)	0.764 (0.003)
3. No Batch Norm.	0.780 (0.009)	1.311 (0.013)	0.815 (0.003)	0.754 (0.004)
4. No Res. Connect.	0.775 (0.008)	1.271 (0.020)	0.816 (0.003)	0.766 (0.004)
5. Classification	0.785 (0.012)	1.253 (0.017)	0.812 (0.004)	0.780 (0.004)
6. Regression	0.850 (0.009)	1.279 (0.016)	0.784 (0.004)	0.686 (0.006)
7. No Order Penalty	0.769 (0.009)	1.251 (0.016)	0.818 (0.003)	0.769 (0.004)

Table 3: Ablation test. Except for models everything is the same as Table 2.

on the providers potentially being underage such as “18y o”, “college girl”, and “tiny”; and switching between third person and first person narratives.

4.5 Emoji Analysis

The fight against human traffickers is adversarial and dynamic. Traffickers often avoid using explicit keywords when advertising victims, but instead use acronyms, intentional typos, and emojis (Tong et al. 2017). Law enforcement maintains a lexicon of trafficking flags mapping certain emojis to their potential true meanings (e.g., the cherry emoji can indicate an underaged victim), but compiling such a lexicon manually is expensive, requires frequent updating, and relies on domain expertise that is hard to obtain (e.g., insider information from traffickers or their victims). To make matters worse, traffickers change their dictionaries over time and regularly switch to new emojis to replace certain keywords (Tong et al. 2017). In such a dynamic and adversarial environment, the need for a data-driven approach in updating the existing lexicon is evident.

As mentioned in Section 3.1, training a skip-gram model on a text corpus can map words (including emojis) used in similar contexts to similar numeric vectors. Besides using the vectors learned from the raw escort ads to train ORNN, we can directly visualize the vectors for the emojis to help identify their relationships, by mapping the vectors to a 2-dimensional space using t-SNE⁹ (van der Maaten and Hinton 2008) (Figure 3).

⁹t-SNE is known to produce better 2-dimensional visualizations

We can first empirically assess the quality of the emoji map by noting that similar emojis do seem clustered together: the smileys near the coordinate (2, 3), the flowers near (-6, -1), the heart shapes near (-8, 1), the phones near (-2, 4) and so on. It is worth emphasizing that the skip-gram model learns the vectors of these emojis based on their contexts in escort ads and not their visual representations, so the fact that the visually similar emojis are close to one another in the map suggests that the vectors have been learned as desired.

The emoji map can assist anti-trafficking experts in expanding the existing lexicon of trafficking flags. For example, according to the lexicon we obtained from Global Emancipation Network¹⁰, the cherry emoji and the lollipop emoji are both flags for underaged victims. Near (-3, -4) in the map, right next to these two emojis are the porcelain dolls emoji, the grapes emoji, the strawberry emoji, the candy emoji, the ice cream emojis, and maybe the 18-slash emoji, indicating that they are all used in similar contexts and perhaps should all be flags for underaged victims in the updated lexicon.

If we re-train the skip-gram model and update the emoji map periodically on new escort ads, when traffickers

than other dimension reduction techniques such as Principal Component Analysis, Multi-dimensional Scaling, and Local Linear Embedding (van der Maaten and Hinton 2008).

¹⁰Global Emancipation Network is a non-profit organization dedicated to combating human trafficking. For more information see <https://www.globalemancipation.ngo>.




Figure 3: Emoji map produced by applying t-SNE to the emojis’ vectors learned from escort ads using skip-gram model. For visual clarity, only the emojis that appeared most frequently in the escort ads we scraped are shown out of the total 968 emojis that appeared.

switch to new emojis, the map can link the new emojis to the old ones, assisting anti-trafficking experts in expanding the lexicon of trafficking flags. This approach also works for acronyms and deliberate typos.

5 Discussion

Human trafficking is a form of modern day slavery that victimizes millions of people. It has become the norm for sex traffickers to use escort websites to openly advertise their victims. We designed an ordinal regression neural network (ORNN) to predict the likelihood that an escort ad comes from a trafficker, which can drastically narrow down the set of possible leads for law enforcement. Our ORNN achieved the state-of-the-art performance on Trafficking-10K (Tong et al. 2017), outperforming all baseline ordinal regression models as well as improving the classification accuracy over the Human Trafficking Deep Network (Tong et al. 2017). We also conducted an emoji analysis and showed how to use word embeddings learned from raw text data to help expand the lexicon of trafficking flags.

Since our experiments, there have been considerable advancements in language representation models, such as BERT (Devlin et al. 2018). The new language representation models can be combined with our ordinal regression layer, replacing the skip-gram model and GF-RNN, to potentially further improve our results. However, our contributions of improving the cost function for ordinal regression neu-

ral networks, qualitatively analyzing patterns in the predicted samples, and expanding the trafficking lexicon through a data-driven approach are not dependent on a particular choice of language representation model.

As for future work in trafficking detection, we can design multi-modal ordinal regression networks that utilize both image and text data. But given the time and resources required to label escort ads, we may explore more unsupervised learning or transfer learning algorithms, such as using object detection (Ren et al. 2015) and matching algorithms to match hotel rooms in the images.

Acknowledgments

We thank Cara Jones and Marinus Analytics LLC for sharing the Trafficking-10K dataset. We thank Praveen Bodigutla for his suggestions on Natural Language Processing literature.

Supplemental Materials

Hyperparameters of the Proposed Ordinal Regression Neural Network

Word Embeddings: speedup method: negative sampling; number of negative samples: 100; noise distribution: unigram distribution raised to 3/4rd; batch size: 16; window size: 5; minimum word count: 5; number of epochs: 50; embedding size: 128; pretraining learning rate: 0.2; fine-tuning

learning rate scale: 1.0.

GF-RNN: hidden size: 128; dropout: 0.2; number of layers: 3; gradient clipping norm: 0.25; L2 penalty: 0.00001; learning rate decay factor: 2.0; learning rate decay patience: 3; early stop patience: 9; batch size: 200; output layer type: mean-pooling; minimum word count: 5; maximum input length: 120.

Multi-labeled Logistic Regression Layer: task weight scheme: uniform; conflict penalty: 0.5.

Access to the Source Materials

The fight against human trafficking is adversarial, hence the access to the source materials in anti-trafficking research is typically not available to the general public by choice, but granted to researchers and law enforcement individually upon request.

Source code: <https://gitlab.com/BlazingBlade/TrafficKill>

Trafficking-10k: cara@marinusanalytics.com

Trafficking lexicon: sherrie@globalemancipation.ngo

References

- Amin, S. 2010. A step towards modeling and destabilizing human trafficking networks using machine learning methods. Conference: Artificial intelligence for development, papers from the 2010 AAAI Spring Symposium, Technical Report SS10-01 (pp. 2-7), Stanford.
- Baccianella, S.; Esuli, A.; and Sebastiani, F. 2009. Evaluation measures for ordinal regression. *9th International Conference on Intelligent Systems Design and Applications*.
- Beel, J.; Gipp, B.; Langer, S.; and Breitingner, C. 2016. Research-paper recommender systems: a literature survey. *International Journal on Digital Libraries* 17(4):305–338.
- Bengio, Y., and Grandvalet, Y. 2004. No unbiased estimator of the variance of k-fold cross-validation. *Journal of Machine Learning Research* 5:1089–1105.
- Bloomfield, P., and Steiger, W. 1980. Least absolute deviations curve-fitting. *SIAM Journal on Scientific and Statistical Computing* 1(2):290–301.
- Cheng, J.; Wang, Z.; and Pollastri, G. 2008. A neural network approach to ordinal regression. *2008 IEEE International Joint Conference on Neural Networks (IEEE World Congress on Computational Intelligence)* 1279–1284.
- Chung, J.; Gulcehre, C.; Cho, K.; and Bengio, Y. 2015. Gated feedback recurrent neural networks. *ICML-15*.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *CoRR* abs/1810.04805.
- Dumais, S. 2004. Latent semantic analysis. *Annual Review of Information Science and Technology* 38(1):188–230.
- Frank, E., and Hall, M. 2001. A simple approach to ordinal classification. *Lecture Notes in Artificial Intelligence* 145–156.
- Goodfellow, I.; Bengio, Y.; and Courville, A. 2016. *Deep Learning*. MIT Press.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. *CVPR*.
- International Labour Organization; Walk Free Foundation; and International Organization for Migration. 2017. *Global estimates of modern slavery: forced labour and forced marriage*. Geneva: International Labour Organization. ISBN: 978-92-2-130131-8.
- Ioffe, S., and Szegedy, C. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *ICML*.
- Irsoy, O., and Cardie, C. 2015. Modeling compositionality with multiplicative recurrent neural networks. *ICLR*.
- LeCun, Y.; Bengio, Y.; and Hinton, G. 2015. Deep learning. *Nature* 521:436–444.
- Li, L., and Lin, H. 2006. Ordinal regression by extended binary classification. *NIPS* 865–872.
- Manning, C.; Raghavan, P.; and Schütze, H. 2009. *An Introduction to Information Retrieval*. Cambridge University Press.
- Mikolov, T.; Chen, K.; Corrado, G.; and Dean, J. 2013a. Efficient estimation of word representations in vector space. *ICLR Workshop Papers*.
- Mikolov, T.; Sutskever, I.; Chen, K.; Corrado, G.; and Dean, J. 2013b. Distributed representations of words and phrases and their compositionality. *NIPS* 3111–3119.
- Narula, S., and Wellington, J. 1982. The minimum sum of absolute errors regression: A state of the art survey. *International Statistical Review* 317–326.
- Niu, Z.; Zhou, M.; Wang, L.; Gao, X.; and Hua, G. 2016. Ordinal regression with multiple output cnn for age estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 4920–4928.
- Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Weiss, R.; Dubourg, V.; Vanderplas, J.; Passos, A.; Cournapeau, D.; Brucher, M.; Perrot, M.; and Duchesnay, E. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* 12:2825–2830.
- Ren, S.; He, K.; Girshick, R.; and Sun, J. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. *NIPS*.
- Rennie, J., and Srebro, N. 2005. Loss functions for preference levels: Regression with discrete ordered labels. In *Proc. Int’l Joint Conf. Artificial Intelligence Multidisciplinary Workshop Advances in Preference Handling*.
- Ruder, S.; Ghaffari, P.; and Breslin, J. 2016. Insight-1 at semeval-2016 task 4: Convolutional neural networks for sentiment classification and quantification. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval 2016)*.
- Schuller, B.; Mousa, A.; and Vryniotis, V. 2015. Sentiment analysis and opinion mining: on optimal parameters and performances. *WIREs Data Mining Knowl. Discov.* 5:255–263.
- Srivastava, N.; Hinton, G.; Krizhevsky, A.; Sutskever, I.; and Salakhutdinov, R. 2014. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research* 15:1929–1958.
- Tarinelli, R. 2018. Online sex ads rebound, months after shutdown of backpage. *The Associated Press*.
- THORN, and Bouché, V. 2018. *Survivor insights: The role of technology in domestic minor sex trafficking*. THORN.
- Tong, E.; Zadeh, A.; Jones, C.; and Morency, L. 2017. Combating human trafficking with deep multimodal models. *Association for Computational Linguistics*.
- van der Maaten, L., and Hinton, G. 2008. Visualizing data using t-sne. *Journal of Machine Learning Research* 9:2431–2456.