

# A Robust Cybersecurity Topic Classification Tool

Elijah Pelofske<sup>1</sup>, Lorie M. Liebrock<sup>1</sup>, and Vincent Urias<sup>2</sup>

<sup>1</sup>New Mexico Cybersecurity Center of Excellence, New Mexico Institute of Mining and Technology, Socorro, New Mexico, USA

<sup>2</sup>Sandia National Laboratories, Albuquerque, New Mexico, USA

## Abstract

In this research, we use user defined labels from three internet text sources (Reddit, StackExchange, Arxiv) to train 21 different machine learning models for the topic classification task of detecting cybersecurity discussions in natural English text. We analyze the false positive and false negative rates of each of the 21 model's in cross validation experiments. Then we present a Cybersecurity Topic Classification (CTC) tool, which takes the majority vote of the 21 trained machine learning models as the decision mechanism for detecting cybersecurity related text. We also show that the majority vote mechanism of the CTC tool provides lower false negative and false positive rates on average than any of the 21 individual models. We show that the CTC tool is scalable to the hundreds of thousands of documents with a wall clock time on the order of hours.

**Keywords**— cybersecurity, topic modeling, text classification, machine learning, neural networks, natural language processing, StackExchange, Reddit, Arxiv, social media

## 1 Introduction

Identifying cybersecurity discussions in open forums at scale is a topic of great interest for the purpose of mitigating and understanding modern cyber threats [12, 19, 22]. The challenge is that these discussions are typically quite noisy (i.e., they contain community known synonyms or acronyms or slang) and it is difficult to get labelled data in order to train resilient NLP (natural language processing) topic classifiers. Additionally, it is important that a tool that detects cybersecurity discussions in internet text sources is *scalable* and offers *low errors rates* (in particular, both low false negative rates and low false positive rates).

In order to address the challenges of finding relevant cybersecurity labelled data, we use a technique that gathers posts or articles from different internet sources that have user defined *topic labels*. We then collect and label the training text as being cybersecurity related or not based on the subset of labels that the text source offers. Thus, the labelled training data we gather is not manually labelled by researchers; instead it is labelled inherently by the system we gather the text from. This provides an additional benefit for cybersecurity related discussions in that it might be difficult for a manual labelling process to catch all of the language variation (e.g., unknown synonyms). This labelling method uses the user defined labels (which removes the need for the labelling process to identify all known cybersecurity terms used in online discussion forums) and provides a much larger and current labelled data set. Lastly, this method of gathering labelled data is *highly* scalable. The reason is that the platforms we used have publicly available data and systems to retrieve that data in very large amounts. We used three specific sources of text: Reddit, StackExchange, and Arxiv.

Using the topic classification labelled data, we train a total of 21 different machine learning models using several different algorithms and the three different text sources. We then show the validation accuracy of

the models, as well as the cross validation (i.e., validating a model on a text source upon which it was not trained) accuracy of each of the models. After that we define a confidence measure for the continuous output values of the machine learning models and present those results for the models where applicable. We next present the Cybersecurity Topic Classification (CTC) tool, which uses the majority vote consensus of all 21 trained models in order to evaluate whether a novel document is cybersecurity related or not. We show that the CTC tool has both scalability to hundreds of thousands of documents per hour and low error rates. Lastly, we provide all of the labelled data we used to train and validate the models in a Github repository [24].

This article is structured as follows. After a brief literature review in Section 1.1, we define our methods of gathering labelled data and then training and validating the machine learning models in Section 2. Section 3 describes the experiments. The investigation of how the minimum token length of the training data changes false negative and false positive rates is presented in Section 3.1. After training each of the models using the specified parameters, in Section 3.2 we show the validation accuracy rates for each of the 21 trained models. In Section 3.3 we show that the continuous outputs of some of the machine learning models can yield a simple confidence measure, where higher confidence means higher solution quality. In Section 3.4 we present the CTC tool and show its scalability (meaning real time text samples classified per hour) and high accuracy. Section 4 discusses conclusions and future work.

## 1.1 Previous work

There is significant interest surrounding the goal of being able to automate cybersecurity threat detection on social media [19, 18, 15, 11, 12, 27, 14, 2]. Twitter, Reddit, and Stackexchange are popular forums from which several previous studies have gathered cybersecurity related documents [19, 11, 14, 15, 2, 20] for the purpose of training machine learning detection systems and classifiers. In particular, [19] used tags (or other community defined mechanisms) as a document labelling method for cybersecurity topic classification related text. There is also some interest in investigating vulnerability discussions on developer sites such as Stackoverflow [18, 20].

There are several different approaches taken with which topic modelling task to use as a signal to detect cybersecurity discussions. Typically the topic classification task is related to training directly on labelled text and then perhaps developing an idea of the more relevant keywords in these discussions [19, 12]. Other researchers use sentiment analysis in conjunction with machine learning models [27, 11].

There are also interesting approaches that use social media as a signal to detect specific cyber-attacks (e.g., DDoS attacks) or vulnerabilities [15, 18, 19, 2].

For a review of text classification with deep learning models, see [22], and for a survey of gathering social media data, see [6].

## 2 Methods

This section describes the methods used to gather labelled text, preprocess and vectorize that text, train several machine learning models using the labelled data, and lastly evaluate the accuracy of these machine learning models.

All figures in this article were created using Matplotlib [13].

### 2.1 Text sources

For this research, we focus on gathering large amounts of text from the three sources Reddit, Stackexchange, and Arxiv. For gathering Reddit text we used the python modules *praw* and *psaw* [25, 1]. In order to query data from StackExchange, we used the python module *StackAPI* (which is a python wrapper for the StackExchange API [28]). For gathering Arxiv documents, we used the python module *arxiv* [5]. As a summary of the raw data collected, Table 1 shows the number of cybersecurity and non cybersecurity

Table 1: Number of documents from each text source and the labelling method used for each source

Source	cybersecurity	non cybersecurity	labelling method
Reddit	164750	4184184	sub-reddits
StackExchange	41162	4842461	post topic labels
Arxiv	12132	28996	keyword search + topic restriction

documents gathered from each source, along with the document labelling method for each source. Next we define the precise methodology for gathering and labelling the documents from each source.

### 2.1.1 Reddit

Reddit [26] is an internet discussion website that allows registered users to submit different types of content (e.g., text posts, images, links, videos) to the site. Each of these posts then get voted on (upvoted or downvoted) by other users. The entire site is logically organized into sub-reddits. Each sub-reddit has a specific scope and topic of discussion. The document labeling strategy we use is to label documents according to which sub-reddit they originate from.

For gathering cybersecurity related text, we first defined 40 cybersecurity related sub-reddits. Next, we queried each of those sub-reddits with a maximum number of posts returned of 1,000,000 (This does not mean we get 1,000,000 posts for each sub-reddit). The default sorting method for the posts returned is a Reddit defined post metric called *Hot*; for the purpose of querying with the API, this can not be changed.

For gathering non cybersecurity related text, we search the top 100 most popular sub-reddits at the time of searching (this list can also be found at [24]) using the same method described above. None of the 40 cybersecurity topic focused sub-reddits are in the top 100 most popular sub-reddits.

For gathering posts in general, we perform some filtering of the data before actually labelling and storing each document. In particular, we remove all posts which are marked as **deleted** or **removed**, since those posts do not contain any post text anymore (these posts were either removed by the user who posted it or a Reddit administrator).

For each post, the title and main post are treated separately in the API. In order to construct a document out of each post, we treated the title as the first sentence and the post content as the remainder of the document (i.e., we merged the two pieces of text with a period and space in between).

Figure 1 shows the distribution of the collected tagged Reddit text in terms of token (i.e., usable words) length.

### 2.1.2 StackExchange

StackExchange [29] is a group of Q&A websites whose topics of discussion are wide ranging; but the most popular sites are developer and programming sites such as Stackoverflow. The sites are self moderating in that registered user's can upvote and downvote posts. Each site also allows the users who post to identify posts using topic tags.

We used multiple StackExchange sites as text sources (see [24] for the full list). For each of these StackExchange sites, we gathered the top (defined by most upvoted) 10,000 posts for each month since the inception of the given StackExchange site.

Next, we defined a list of cybersecurity related topic tags across different security and technology related StackExchange sites (this list can be found at our Github [24]). We labelled each post as being cybersecurity related if it used any of the tags in our list and otherwise it was labelled as not cybersecurity related.

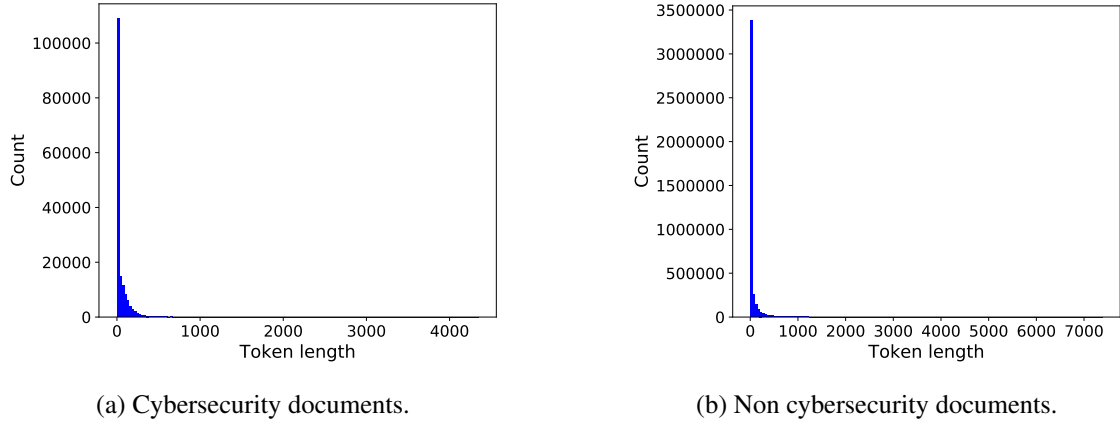


Figure 1: Reddit token data histogram




Figure 2: StackExchange token data histogram.

As with Reddit, for each post we queried, we merged the title and main post text into a piece of text.

Figure 2 shows the distribution of the collected tagged StackExchange text in terms of token (i.e., usable parsed word) length.

### 2.1.3 Arxiv

Arxiv [4] is an open access repository of e-prints (including papers before peer-review and after peer-review). The repository includes scientific papers on wide ranging topics including computer science, mathematics, physics, statistics, and economics. Each paper comes with one more topic labels and can be downloaded in the form of a PDF.

The methodology for gathering cybersecurity labelled text is as follows. We used a seed list of cybersecurity terms and topics in order to search Arxiv (the word list is provided in [24]). Of the resulting papers returned in the search, if any of the tags are **cs.CR** (which broadly is defined as computer science regarding cryptography and cybersecurity), then we download that pdf and tag the document as cybersecurity related.

For gathering non cybersecurity related documents from Arxiv, we searched all of the remaining non **cs.CR** categories and chose the top 100 (i.e., most relevant) papers from each of those categories; any of these papers with a **cs.CR** were not downloaded, since those documents would be cybersecurity related.

The last step involves some text cleaning, which is specific to Arxiv, since all of the documents are PDFs. First, we remove all non-English documents (some of the downloaded technical documents were in a va-




Figure 3: Arxiv token data histogram.

Table 2: Pre-processed text showing the number of cleaned and tokenized documents that are not empty (here by empty we mean having no detectable English words).

Source	cybersecurity	non cybersecurity
Reddit	163,739	4,129,605
Stackexchange	41,162	4,842,459
Arxiv	12,132	289,969

riety of other languages). Non-English documents were found using **langdetect** [17]. With non-English documents, the majority of the text was non-English, therefore the full document was not used. In future work, we may instead translate the non-English documents and use them as well. Next, we use **tika** to parse the PDF's into raw text. In some cases, **tika** is unable to parse the PDF's (in which case we can not use those documents). Figure 3 shows the distribution of the collected tagged Arxiv text in terms of token (i.e., usable word) length.

In Figure 3, we see that the average document length for Arxiv is in the thousands of words. In Figure 1, we saw that the Reddit average post length is less than 100 words, in contrast to Arxiv. In Figure 2, we observe that the average StackExchange post length is approximately 100. Across all three text sources, the difference in average document length between cybersecurity and non cybersecurity documents is marginal. The most significant difference between the cybersecurity and non cybersecurity documents is that there are many more non cybersecurity documents than cybersecurity documents.

## 2.2 Text preprocessing

In order to use each document in the various classification algorithms, it is necessary to vectorize each document. In order to standardize all documents so that this vectorization process is consistent, we preprocess each document in a variety of ways. Specifically we remove URL's, non ascii characters, code tags, all HTML tags, and excessive whitespace. Table 2 shows the total number of usable documents after we have cleaned the text.

## 2.3 Vectorizer

For this research, we used a term frequency - inverse document frequency (TF-IDF) vectorizer found in *scikit-learn* [23]. TF-IDF vectorization is used to show how common (or important) a given word is in the

input document. In particular, it weights words more heavily that occur at a greater frequency. The first step in creating a consistent vectorization method across all documents was compiling the more relevant and popular English words to use for this vectorizer. To this end, we used two English word lists and a custom list of cybersecurity terms. We used *google20kwords* from [10] and *30k.txt* from [9].

We merge all of these lists, remove all duplicates, and remove all entries that are not English words. To determine whether to remove or keep a word, if any of the three methods *nlk* (see [7]) words, *nlk* wordnet, or the python module *enchant* consider a word to be English, we include it in the dictionary. Pronouns and nouns, such as names and company names, are not considered English words. The resulting English dictionary has a length of 24,538 (the word list is provided in [24]). Non English words were not vectorized because the vast majority of the text we gathered was English and using this fixed English dictionary means that non English words are automatically removed. Future work may include handling foreign language words.

We then fit a TF-IDF vectorizer to this dictionary, save it with 32 bit precision (in order to reduce memory costs), and use this vectorizer for computing the inputs to all machine learning classifiers. The TF-IDF vectorizer can be used to vectorize sentences, paragraphs, or entire documents. In our case, we vectorized each document using TF-IDF. We do not remove any stop words during this vectorization because in simple grid search experiments, removing stop words increased accuracy in some cases, while in other cases it decreased accuracy (both training and testing accuracy). Therefore, since removing stop words or not was not clearly motivated, we did not perform this extra step.

## 2.4 Machine Learning Classifiers

We test a total of 6 different types of machine learning classifiers. For all of these classifiers, we use a fixed TF-IDF vectorizer that has been fitted to the English dictionary described in the previous section. Specifically, we use *Decision Tree*, *Random Forest*, *Logistic*, *LinearSVC*, and *Multi-layer Perceptron* (MLP). Each of these five models were provided in the python module **scikit-learn** (version 0.22.0) [23, 8]. We also build a Deep Neural Network (DNN) model using **tensorflow** (version 2.3.0) [21].


For the Decision Tree classifier, we set the maximum depth to be 100, to reduce the real time computational cost. For the DNN model, we specify all of the parameters and hyperparameters used in Section 2.4.1. For all other models, we used default parameters.

### 2.4.1 Deep Neural Network


We build a simple multi-layer neural network with one hidden layer using tensor flow; we use a sequential model with three layers. The first layer has an input equal to the length of English dictionary we used (24,538 words), with 10,000 nodes. The second layer has 1,000 nodes and the last layer has 100 nodes. The output layer has two nodes. The output layer uses an activation function *softmax* and every other layer uses *relu*.

The number of epochs we use is not fixed when training the model. Instead, we train the model until a certain threshold training accuracy has been reached. In our experiments we try two different accuracy thresholds 0.95 and 0.99. The number of samples per iteration was fixed to 4,000. We use *sparse-categorical-crossentropy* for the model loss function, the model metric is *accuracy*, and *Adam* [16] as the model optimizer. To speed up the training and validation time, we also set the flags *workers* to 400 and *use-multiprocessing* to True.

Since we try two different accuracy thresholds (0.95 and 0.99), we actually train two different deep neural network models. Thus, in total we train seven different models in the following section.



(a) False negative rate as a function of token length



(b) False positive rate as a function of token length

Figure 4: False negative and false positive rates as a function of token length (i.e. usable word length) for the Reddit labelled text

### 3 Results

In this section we investigate some of the relevant parameters, both of the text and the machine learning models, that influence the training and testing accuracy.


First, we perform some simple grid-search optimizations of some of the model hyperparameters. Although not shown here for the sake of space, we found that these hyperparameters converge to reasonable performance; for example *max-iters* should be set to a reasonably large number, but the training accuracy plateaus relatively quickly (i.e., after several hundred iterations). Second, in Section 3.1 we investigate how the minimum word length of the data changes both the training and the testing accuracy. In these first two steps, we generally had to run many instances of each algorithm and therefore typically we would select a subset of the data to train and validate on. In the third step, we train on half of each of the text source's datasets and validate the models on the other half (in some cases, due to RAM limitations, we trained on 3/8 of the data and validated on 5/8). Next, in Section 3.2 we validate how each of these trained models performs on disparate text sources that they were not trained on (e.g., validating a model trained on Reddit text using text from StackExchange). In Section 3.3 we define a confidence measure from the continuous outputs of some of the machine learning models and determine how the confidence measure corresponds to error rates in the training dataset results. Finally, we combine all of these models into a unified Python tool called CTC. We show that CTC classifies large numbers of text documents with relatively low error rates.

As shown in Table 1, the labelled data we have gathered is quite skewed towards being made of mostly not cybersecurity documents. This is not unexpected, but it means that training a machine learning model on that dataset will result in unequal weighting of the importance of the classification tasks. In an attempt to correct this towards an evenly balanced dataset, for each machine learning model we use the following class weighting rule. If we have  $n_c$  cybersecurity documents and  $n_{nc}$  not cybersecurity documents, then the class weight for not cybersecurity is 1 and the class weight for cybersecurity is  $\frac{n_{nc}}{n_c}$ .


In the remainder of the article we use FN to denote false negative rate and FP to denote false positive rate.

#### 3.1 Token Length

In this section, we determine the behavior of several of the classifiers as a function of the minimum token length used on the training and validation datasets. Across all three text sources we select random subsets of the full training data to reduce the total needed computation time. The procedure we used was to eliminate all documents with usable token lengths less than  $N$  ( $N$  is plotted on the x-axis of the plots shown in the remaining subsections), train the model on the training data, and then validate (and plot) the accuracy using the unseen validation data (both the validation data and training data had at least  $N$  usable tokens in each




(a) False negative rate as a function of token length




(b) False positive rate as a function of token length

Figure 5: False negative and false positive rates as a function of token length for the StackExchange labelled text



(a) False negative rate as a function of token length



(b) False positive rate as a function of token length

Figure 6: False negative and false positive rates as a function of token length for the Arxiv labelled text

document). To save space, we show these results for four different machine learning classifiers, *Decision Tree*, *LinearSVC*, *Random Forest*, *Logistic*, on each of three text sources.

### 3.1.1 Reddit

For Reddit, we set both the training and validation datasets to have 40,000 Cybersecurity labelled documents and 50,000 non Cybersecurity labelled documents. In Figure 4, we plot the false negative and false positive rate as a function of token length for four different classifiers. Figure 4 shows on the left an increase in the false negative rate for the Random Forest Decision Tree and Logistic classifiers as token length increases, while LinearSVC's false negative rate remains relatively unchanged. Figure 4 shows on the right a decrease in the false positive rate as a function of token length across all four classifiers.

### 3.1.2 StackExchange

For StackExchange, we set both the training and validation datasets to have 20,581 Cybersecurity labelled documents and 50,000 non Cybersecurity labelled documents. In Figure 5, we plot the false negative and



Table 3: Cross validate DNN with training accuracy 0.99

Training source	Validation source	Training accuracy	Validation FN	Validation FP
Reddit 1/2	Arxiv		0.0505	0.1438
Reddit 1/2	StackExchange		0.2903	0.1115
Reddit 1/2	Reddit 1/2	0.9932	0.0177	0.1801
StackExchange 1/2	Arxiv		0.0007	0.7711
StackExchange 1/2	StackExchange 1/2	0.991	0.0083	0.3081
StackExchange 1/2	Reddit		0.0666	0.6837
Arxiv 1/2	Arxiv 1/2	0.9933	0.0172	0.0236
Arxiv 1/2	StackExchange		0.2353	0.2007
Arxiv 1/2	Reddit		0.1068	0.4041

Table 4: Cross validate DNN with training accuracy 0.95

Training source	Validation source	Training accuracy	Validation FN	Validation FP
Reddit 1/2	Arxiv		0.0867	0.0886
Reddit 1/2	StackExchange		0.4485	0.0499
Reddit 1/2	Reddit 1/2	0.9573	0.0397	0.1043
StackExchange 1/2	Arxiv		0.0013	0.6222
StackExchange 1/2	StackExchange 1/2	0.9644	0.0317	0.1415
StackExchange 1/2	Reddit		0.0662	0.5342
Arxiv 1/2	Arxiv 1/2	0.9578	0.0182	0.0748
Arxiv 1/2	StackExchange		0.2052	0.2393
Arxiv 1/2	Reddit		0.0714	0.4964

false positive rate as a function of token length for four different classifiers. Figure 5 shows on the left that Random Forest, Logistic, and Decision Tree models increase in false negative rate as token length increases. Figure 5 shows on the right that the false positive rate decreases as a function of token length for DecisionTree and RandomForest, while LinearSVC and Logistic classifiers remain relatively unchanged.

### 3.1.3 Arxiv

For Arxiv, we set both the training and validation datasets to have 1,000 cybersecurity labelled documents, and 3,000 not cybersecurity labelled documents. In Figure 5 we plot the false negative and false positive rate as a function of token length for four different classifiers. Figure 5 shows on the left an increase in false negative rates as a function of usable token length for each of the four classifiers. Figure 5 shows on the right that DecisionTree and RandomForest false positive rates decrease as a function of token length, while LinearSVC and Logistic classifiers increase in false positive rates as a function of token length. Since the average usable word length of the documents from Arxiv are very large, see Figure 3, the manner in which the error rates change as a function of token length for the Arxiv source is not as important as for Reddit and StackExchange sources.

To train the machine learning models, we want to not include documents with too few tokens because the topic being discussed may be ambiguous or unclear. However, we want relatively balanced error rates between false negative and false positive. Most importantly, we want our models to be able to classify a wide range of internet text. For this reason, we would want to train on data with smaller token lengths. Given these factors, we conclude that a minimum usable token length of ten is reasonable for training the models in the remaining experiments.

Table 5: Cross validate Logistic

Training source	Validation source	Training accuracy	Validation FN	Validation FP
Reddit 1/2	Arxiv		0.27	0.0101
Reddit 1/2	StackExchange		0.5136	0.0282
Reddit 1/2	Reddit 1/2	0.9509	0.0496	0.0849
StackExchange 3/8	Arxiv		0.0255	0.2211
StackExchange 3/8	StackExchange 5/8	0.9599	0.0403	0.1037
StackExchange 3/8	Reddit		0.1162	0.4089
Arxiv 1/2	Arxiv 1/2	0.965	0.0296	0.0557
Arxiv 1/2	StackExchange		0.0347	0.6919
Arxiv 1/2	Reddit		0.0035	0.8513

Table 6: Cross validate RandomForest

Training source	Validation source	Training accuracy	Validation FN	Validation FP
Reddit 1/2	Arxiv		0.0037	0.7019
Reddit 1/2	StackExchange		0.0271	0.7481
Reddit 1/2	Reddit 1/2	0.9998	0.0018	0.6921
StackExchange 1/2	Arxiv		0.0001	0.9736
StackExchange 1/2	StackExchange 1/2	1.0	0	0.9811
StackExchange 1/2	Reddit		0.0004	0.982
Arxiv 1/2	Arxiv 1/2	1.0	0.013	0.0374
Arxiv 1/2	StackExchange		0.001	0.9705
Arxiv 1/2	Reddit		0.0001	0.9914

### 3.2 Validation

Next, we train all seven of the machine learning models described in Section 2.4 on each of the three labelled text sources (Reddit, StackExchange, Arxiv), resulting in a total of 21 distinct trained models. Following the token length experiments of Section 3.1, we use only labelled data which has word length (token length) of at least ten (this reduces the total amount of labelled data we use from the initial corpus shown in Table 2). For each model type and each text source, we train the model on one half of the labelled data, and then validate on the other half (in some cases, where training on one half of the data is too computationally costly, we train on 3/8 of the data).

Since we want these machine learning models to be applied to text that may not come from Arxiv or StackExchange or Reddit, one way to evaluate each of these models in a more robust way is to predict the topic of text from the two sources the model was not trained on; for example predicting the topic of the labelled Arxiv and StackExchange dataset given the model was trained on Reddit text. In this section, we show the validation accuracy (i.e., testing accuracy) and cross validation accuracy for each of the 21 trained models. In particular, Tables 3, 4, 5, 6, 7, 8, 9 show the cross validation results. In these tables we show the training accuracy as well as the cross validation false negative and false positive rates. Under the training accuracy column, we only get one training accuracy entry for each of the three text sources (the other rows are validation accuracy results), which means that six of the entries in that column will be empty.

As a general summary of the cross validation results, we observe the trend that the false negative and false positive rates for the models trained on a source and then validated on the same text source (not the same data, just the same text source e.g., Reddit) are usually relatively low (less than ten percent). The exceptions to this are RandomForest and Decision Tree for Reddit and StackExchange sources. We also observe varied error rates for the cross validation experiments. In particular, we usually see an asymmetry in the error rates - i.e., the false negative rate is very high and the false positive rate is very low, or the reverse.

Table 7: Cross validate LinearSVC

Training source	Validation source	Training accuracy	Validation FN	Validation FP
Reddit 1/2	Arxiv		0.1482	0.0463
Reddit 1/2	StackExchange		0.4543	0.0397
Reddit 1/2	Reddit 1/2	0.9566	0.0463	0.0992
StackExchange 3/8	Arxiv		0.0041	0.5392
StackExchange 3/8	StackExchange 5/8	0.9673	0.0335	0.1385
StackExchange 3/8	Reddit		0.0928	0.4993
Arxiv 1/2	Arxiv 1/2	0.9882	0.0147	0.0302
Arxiv 1/2	StackExchange		0.0205	0.6619
Arxiv 1/2	Reddit		0.0039	0.8224

Table 8: Cross validate DecisionTree

Training source	Validation source	Training accuracy	Validation FN	Validation FP
Reddit 1/2	Arxiv		0.2338	0.2309
Reddit 1/2	StackExchange		0.4181	0.3088
Reddit 1/2	Reddit 1/2	0.9749	0.0398	0.3233
StackExchange 1/2	Arxiv		0.0242	0.2986
StackExchange 1/2	StackExchange 1/2	0.9854	0.0188	0.4127
StackExchange 1/2	Reddit		0.0248	0.6518
Arxiv 1/2	Arxiv 1/2	1.0	0.0249	0.0641
Arxiv 1/2	StackExchange		0.0186	0.7244
Arxiv 1/2	Reddit		0.0132	0.7857

Table 4 and 3 show the direct comparison of validation and training accuracy differences when applying the 0.95 and 0.99 training accuracy cutoff. The most notable differences between these two models is that the 0.95 training accuracy model has less overall false positive error rate when predicting on the Reddit and StackExchange text, whereas the Arxiv testing false positive rates were worse. For both the 0.99 and 0.95 trained models we see that the worse performing regions are the StackExchange trained models predicting on Reddit and Arxiv text, which gives a false positive rate greater than 0.50 in all cases. The false positive rates for the Arxiv text trained models predicting on Reddit text also gives very high false positive rates (greater than 0.40).

Tables 7, 5, 9, and 6 all follow similar trends to the DNN models in terms of the highest error rate validation results. In particular, the StackExchange trained models predicting on Reddit and Arxiv text, and the Arxiv trained models predicting on Reddit text all give very high false positive rates. The false negative error rates on the other hand are significantly lower than the false positive rates for most of the models with a few exceptions such as LinearSVC model trained on Reddit text predicting on StackExchange text (see Table 7).

### 3.3 Classification Confidence Measure

Here we define a confidence measure for the continuous output values of the machine learning models. Not all of the machine learning models produce continuous output values, but for the models that do we can define a simple confidence measure that quantifies how far the outputs are from either 0 or 1 in the final decision.

For each sample that is classified (if the model produces continuous output), we get an output vector of the form  $[v_0, v_1]$ , where  $v_0 + v_1 = 1$ ,  $v_0 \geq 0$ , and  $v_1 \geq 0$ .  $v_0$  corresponds to the output for the class 0 (in our

Table 9: Cross validate Multi-Layer Perceptron

Training source	Validation source	Training accuracy	Validation FN	Validation FP
Reddit 1/2	Arxiv		0.0343	0.2406
Reddit 1/2	StackExchange		0.1412	0.276
Reddit 1/2	Reddit 1/2	0.9996	0.0078	0.3183
StackExchange 3/8	Arxiv		0.0004	0.8492
StackExchange 3/8	StackExchange 5/8	0.9995	0.0027	0.5325
StackExchange 3/8	Reddit		0.0179	0.8125
Arxiv 1/2	Arxiv 1/2	0.9999	0.0147	0.0432
Arxiv 1/2	StackExchange		0.3002	0.3592
Arxiv 1/2	Reddit		0.117	0.5524




Figure 7: Histogram of confidence measures (left column) and error rates for each bin of confidence measures (right column). Logistic trained and then validated on Reddit text (top row), Arxiv text (middle row), and StackExchange text (bottom row).




Figure 8: Histogram of confidence measures (left column) and error rates for each bin of confidence measures (right column). DNN model trained to 0.95 training accuracy and then validated on Reddit text (top row), Arxiv text (middle row), and StackExchange text (bottom row).

system that means not cybersecurity related), and  $v_1$  corresponds to the output for the class 1 (in our system that means the text is cybersecurity related). The confidence measure for a particular document  $d$  (with model output  $[v_0, v_1]$ ) is then  $conf_d = \max([v_0, v_1]) - 0.5$ . Thus the  $conf_d$  measure is within  $[0, 0.5]$ , where 0 indicates no confidence (i.e., the machine learning model effectively predicted a 50-50 coin flip), and 0.5 is the highest confidence output. A natural question is whether a higher confidence measure from the model corresponds to higher solution quality (i.e., lower false negative and false positive rates). In order to evaluate this, we compute the confidence measures for labelled data in the validation datasets from Section 3.2, then bin the documents according to their confidence measure and determine the error rates for each of these bins of text samples. This confidence measure could be applied to models in general (with the exception of LinearSVC because the output from LinearSVC is not continuous). However, models such as RandomForest and DecisionTree can produce outputs that are not distributed across the confidence measure range, which can lead to potentially choppy results, where for some region of confidence measures there are no documents predicted in that region. The DNN and Logistic models give consistent continuous output, so these are the models we will demonstrate this analysis on.




Figure 9: Histogram of confidence measures (left column) and error rates for each bin of confidence measures (right column). DNN model trained to 0.99 training accuracy and then validated on Reddit text (top row), Arxiv text (middle row), and StackExchange text (bottom row).

For analysis of these confidence measures, we show two associated figures for each data set and machine learning model analyzed. In the left side graph of each figure, we show the number of documents in each bin versus the confidence measure for the bin. In the right side graph of each figure, we show the error rate for both false positives (FP) and false negatives (FN) of documents in each confidence measure bin. These two figures for each data set and machine learning model show the correlation of our confidence measure with the false positive and false negative error rates.

Figure's 7, 8, 9, 10, 11, and 12 show how the error rates (right hand side columns) change for different confidence measure values (left hand side columns). Across all of the error rate plots we observe that both the false negative and false positive rates decrease as confidence gets higher; indicating that higher confidence *does* correspond to lower error rates. Across all of the confidence value histograms, we observe that the model's predict at very high confidence; the one slight exception here is Figure 7 middle left plot, where the histogram peaks slightly before 0.5. Note that the instability observed in the Arxiv validation dataset plots is due to the small sample size of the Arxiv dataset. Conversely, on the StackExchange validation datasets we observe small overall changes to the error rates, up until very high confidence levels.




Figure 10: Histogram of confidence measures (left column) and error rates for each bin of confidence measures (right column). RandomForest model trained and then validated on Reddit text (top row), Arxiv text (middle row), and StackExchange text (bottom row).

Figures 9 and 8 show the differences between the DNN model trained to 0.99 and 0.95 accuracy respectively. The primary differences are that the 0.99 accuracy trained models have significantly higher confidence values.

### 3.4 Cybersecurity Topic Classification (CTC) tool

Now we combine all of these trained models (the validation data for these models was shown in the previous section) into a single NLP tool for cybersecurity topic modeling. In particular, for a given set of documents, we vectorize using the TF-IDF vectorizer used in all of these experiments. Then, we run all 21 models on those documents and report the results. To come to a final decision as the output of the tool, we take the majority vote on the output of the 21 ML models. This is a reasonable approach because it is simply taking the consensus of all of the trained models. This means that no outlier can change the tool's decision making process, making the system robust against outlier predictions.




Figure 11: Histogram of confidence measures (left column) and error rates for each bin of confidence measures (right column). Logistic model trained on StackExchange text and then validated on Arxiv text (top row), Reddit text (bottom row)

### 3.4.1 CTC Error rates

On this section we show that the CTC tool on average performs better than any of the individual 21 trained models.

We use the labelled dataset of [30] for this experiment. In particular, we use both the training and testing dataset from [30] in this validation test. The dataset has four different labels: 1 (World), 2 (Sports), 3 (Business), 4 (Sci/Tech). Since 4 can include cybersecurity content, we entirely remove 4, and then merge classes 1, 2, and 3 into non cybersecurity. Collectively, we call this data source **ag-news**. We also use the **philosophy** [3], which is definitely not cybersecurity discussion. Both **ag-news** and **philosophy** are used as large datasets which provide an idea of the *false positive* rate.

Next, we want to develop a labelled dataset for cybersecurity related text to validate the performance of CTC. To this end, we pull a random subset of 100 documents from seven internet forums/discussion blogs and then hand label the topic's of each of these 700 documents. Note that several of these sources are heavily cybersecurity related, which is valuable as we want to get a reasonable sample of cybersecurity text with which to validate the models. In several cases, the random documents we accessed did not have any English words or enough English words (i.e., some documents had only one or two words), so those documents were discarded. In total, from the seven sources, we have 698 labelled documents.

Table 10 shows the number of incorrectly labelled documents when using the CTC tool (the majority vote of the 21 individual models) and when using each of the individual 21 models applied to the new validation data described above. On average, the CTC tool outperforms the individual models across different text sources. Figure 14 shows a more concise version of Table 10, where we aggregate the results across the *cybersecurity* and *non cybersecurity* labelled validation text. Figure 14 shows that while there are some individual models that have very low *false positive* or low *false negative* rates, the majority vote of the 21 models has the lowest overall *false positive* and low *false negative*. For example, we observe that *Arxiv-RandomForest* has a very low false positive rate, but then has a very high false negative rate. Thus, the



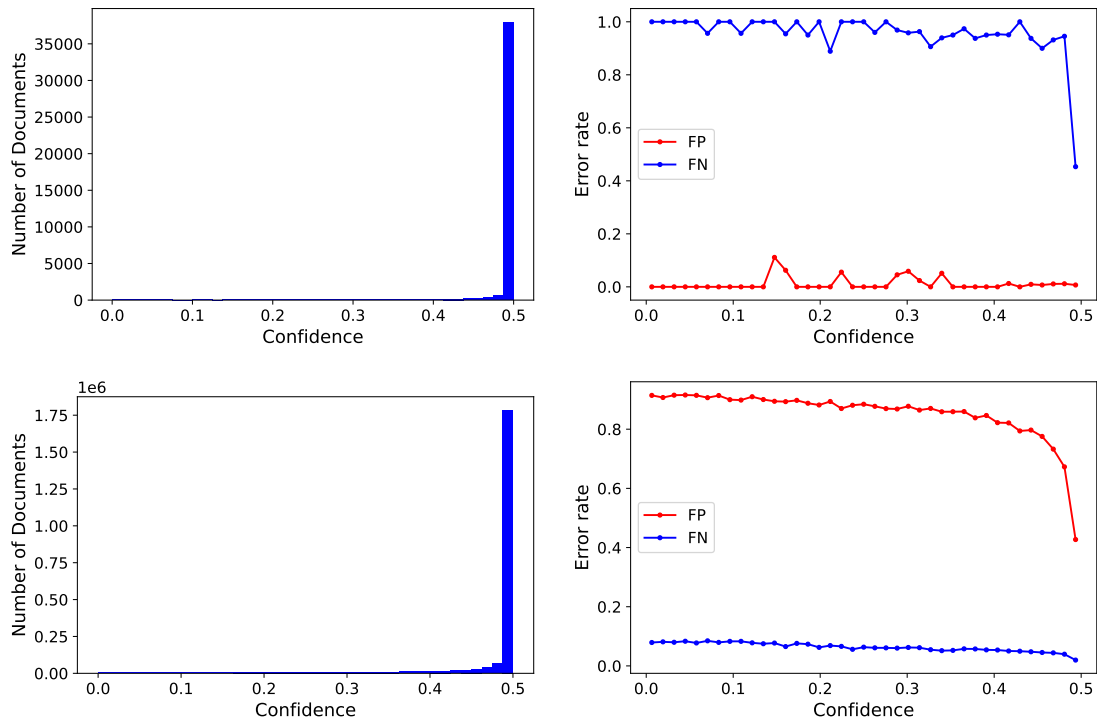


Figure 12: Histogram of confidence measures (left column) and error rates for each bin of confidence measures (right column). DNN 0.95 model trained on StackExchange text and then validated on Arxiv text (top row), Reddit text (bottom row)




Figure 13: Number of input documents vs Wall clock time for the CTC tool

majority vote mechanism used in the CTC tool is more robust compared to the individual models.

### 3.4.2 CTC Confidence Metric

The models that make up the CTC tool can still provide confidence measures for the overall decision made by the tool. As an example, Figure 15 shows the distribution of confidence values for the three Logistic models in the CTC tool when applied to the **philosophy** (right column) dataset and the **ag-news** dataset (left

column). We see that the distribution of confidence values changes depending on the data the model was trained on, as well as the data the model is predicting on. In particular, the StackExchange trained models (bottom row) have less confident outputs compared to the Arxiv trained models (middle row). The Reddit trained models (top row) have more confident outputs when classifying the **ag-news** dataset compared to the **philosophy** dataset. The behavior of the Reddit text trained model is different than the other two models in this respect; this could be because the Reddit text is predominantly very short text length (see Figure 1), whereas the other two source had higher diversity of text length. As with the previous confidence measure plots, we observe that model's predict at very high confidence (meaning very close to the maximum confidence measure of 0.5) on average.

### 3.4.3 CTC Timing

Lastly, we characterize how the CTC tool scales in terms of documents analyzed over time. We pull random subsets of some internet discussion posts of varying length and content and then measure the total wall clock time needed to classify that set of documents. We repeat this process for increasing numbers of input documents.

Figure 13 shows this scaling of documents analyzed over time. We observe that the wall clock time needed to classify  $N$  input documents has a consistently linear scaling. This scaling is largely due to using multiprocessing for the Tensorflow DNN models. The limiting factor of how many documents CTC can ingest is the available RAM on the host computer. For the result shown in Figure 13, the device had 500 GB of RAM, but we do not characterize exactly where the RAM limit is.

## 4 Conclusion and Future Work

This article proposed a methodology for gathering labelled English text for the purpose of topic modelling cybersecurity discussions. We then trained multiple machine learning models using this labelled data, and showed that combining these models into a consensus majority voting tool results in both low error rates and good scalability in terms of text samples classified per hour.

There remain many possible future research avenues:

1. Considering unsupervised clustering methods, which can determine how similar each source of training data is, and therefore also how similar a new input document is to each of these sources can yield lower error rates and lower computation times by using only a single machine learning classifier that has the highest accuracy for that type of document (based on the validation experiments done previously).
2. Using clustering methods to first cluster the corpus of labelled training data into a large number of clusters and then separately training machine learning models on each of those clusters could yield higher accuracy when using all of the models together.
3. Using other vectorization algorithms that make use of a larger dictionary of words, as well as the sequence in which words are used, may improve performance over using a bag of words model.
4. Gathering more labelled data from other online sources, for example from Twitter, Quora, or Medium, would increase the size of the training set. These sites also offer methods of mining labelled natural language posts, for example [19] and [15] use data from Twitter.
5. Analyzing the explainability of the machine learning algorithms could generate new insight. For example, the decision tree algorithm might provide a more compact and explainable logical decision path for this particular topic modeling task.
6. Using other machine learning algorithms that may be better suited to topic classification tasks, for example Recurrent Neural Networks.

7. Exploring the use of subset's of the 21 machine learning models (with the majority voting mechanism) could improve results and reduce computational time. In particular, some subset may perform better than all 21 models and individual models.
8. Including non English text may provide additional valuable data once translated.

## Acknowledgements

Sandia National Laboratories (SNL) is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000. The New Mexico Cybersecurity Center of Excellence (NMCCoE) is a statewide Research and Public Service Project supported center for economic development, education, and research. The authors would like to thank both SNL and NMCCoE for funding and computing system access/support.

## References

- [1] 2021. URL: <https://pypi.org/project/psaw/> (visited on 06/11/2021).
- [2] Rasim Alguliyev, Ramiz Aliguliyev, and Fargana Abdullayeva. "Deep Learning Method for Prediction of DDoS Attacks on Social Media". In: *Advances in Data Science and Adaptive Analysis* 11 (Feb. 2019). DOI: 10.1142/S2424922X19500025.
- [3] Kourosh Alizadeh. *Text-Based Ideological Classification*. [https://github.com/kcalizadeh/phil\\_nlp](https://github.com/kcalizadeh/phil_nlp). 2021.
- [4] Arxiv. 2020. URL: <https://arxiv.org/> (visited on 06/11/2021).
- [5] Arxiv. 2021. URL: <https://pypi.org/project/arxiv/> (visited on 06/11/2021).
- [6] Bogdan Batrinca and Philip C. Treleaven. "Social media analytics: a survey of techniques, tools and platforms". In: *AI & SOCIETY* 30.1 (Feb. 2015), pp. 89–116. ISSN: 1435-5655. DOI: 10.1007/s00146-014-0549-4. URL: <https://doi.org/10.1007/s00146-014-0549-4>.
- [7] Steven Bird, Ewan Klein, and Edward Loper. *Natural language processing with Python: analyzing text with the natural language toolkit*. "O'Reilly Media, Inc.", 2009.
- [8] Lars Buitinck et al. "API design for machine learning software: experiences from the scikit-learn project". In: *ECML PKDD Workshop: Languages for Data Mining and Machine Learning*. 2013, pp. 108–122.
- [9] Derek Chuank. *high-frequency-vocabulary*. <https://github.com/derekchuank/high-frequency-vocabulary>. 2020.
- [10] furiousapathy. *most-common-english-words*. <https://github.com/furiousapathy/most-common-english-words>. 2020.
- [11] Aldo Hernandez-Suarez et al. "Social Sentiment Sensor in Twitter for Predicting Cyber-Attacks Using l1 Regularization". In: *Sensors* 18 (Apr. 2018), p. 1380. DOI: 10.3390/s18051380.
- [12] Jack Hughes et al. "Detecting Trending Terms in Cybersecurity Forum Discussions". In: *Proceedings of the Sixth Workshop on Noisy User-generated Text (W-NUT 2020)*. Online: Association for Computational Linguistics, Nov. 2020, pp. 107–115. DOI: 10.18653/v1/2020.wnut-1.15. URL: <https://www.aclweb.org/anthology/2020.wnut-1.15>.

- [13] J. D. Hunter. “Matplotlib: A 2D graphics environment”. In: *Computing in Science & Engineering* 9.3 (2007), pp. 90–95. DOI: 10.1109/MCSE.2007.55.
- [14] Ruth Ikwu and Panos Louvieris. “Monitoring “Cyber Related” Discussions in Online Social Platforms”. In: *International Journal on Cyber Situational Awareness* 4.1 (Dec. 2019), pp. 69–98. ISSN: 2633-495X. DOI: 10.22619/ijcsa.2019.100126. URL: <http://dx.doi.org/10.22619/IJCSA.2019.100126>.
- [15] Rupinder Paul Khandpur et al. “Crowdsourcing Cybersecurity: Cyber Attack Detection Using Social Media”. In: *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*. CIKM ’17. Singapore, Singapore: Association for Computing Machinery, 2017, pp. 1049–1057. ISBN: 9781450349185. DOI: 10.1145/3132847.3132866. URL: <https://doi.org/10.1145/3132847.3132866>.
- [16] Diederik P. Kingma and Jimmy Ba. *Adam: A Method for Stochastic Optimization*. 2017. arXiv: 1412.6980 [cs.LG].
- [17] *langdetect*. 2021. URL: <https://pypi.org/project/langdetect/> (visited on 06/30/2021).
- [18] Triet Huynh Minh Le et al. “Demystifying the Mysteries of Security Vulnerability Discussions on Developer Q&A Sites”. In: *CoRR* abs/2008.04176 (2020). arXiv: 2008.04176. URL: <https://arxiv.org/abs/2008.04176>.
- [19] Richard P. Lippman et al. “Toward Finding Malicious Cyber Discussions in Social Media”. In: *AAAI Workshops*. 2017. URL: <http://aaai.org/ocs/index.php/WS/AAAIW17/paper/view/15201>.
- [20] Tamara Lopez et al. “An Anatomy of Security Conversations in Stack Overflow”. In: *41st ACM/IEEE International Conference on Software Engineering*. Aug. 2019, pp. 31–40. URL: <http://oro.open.ac.uk/59243/>.
- [21] Martin Abadi et al. *TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems*. Software available from tensorflow.org. 2015. URL: <https://www.tensorflow.org/>.
- [22] Shervin Minaee et al. *Deep Learning Based Text Classification: A Comprehensive Review*. 2021. arXiv: 2004.03705 [cs.CL].
- [23] F. Pedregosa et al. “Scikit-learn: Machine Learning in Python”. In: *Journal of Machine Learning Research* 12 (2011), pp. 2825–2830.
- [24] Elijah Pelofske. *CTC*. <https://github.com/epelofske-student/CTC>. 2021.
- [25] *PRAW*. 2021. URL: <https://praw.readthedocs.io/en/latest/> (visited on 06/11/2021).
- [26] *Reddit*. 2020. URL: <https://www.reddit.com/> (visited on 06/11/2021).
- [27] Kai Shu et al. “Understanding Cyber Attack Behaviors with Sentiment Information on Social Media”. In: *Social, Cultural, and Behavioral Modeling*. Ed. by Robert Thomson et al. Cham: Springer International Publishing, 2018, pp. 377–388. ISBN: 978-3-319-93372-6.
- [28] *StackAPI*. 2021. URL: <https://stackapi.readthedocs.io/en/latest/> (visited on 06/11/2021).
- [29] *Stackexchange*. 2020. URL: <https://stackexchange.com/sites#> (visited on 06/11/2021).
- [30] Xiang Zhang, Junbo Jake Zhao, and Yann LeCun. “Character-level Convolutional Networks for Text Classification”. In: *NIPS*. 2015.

## Authors



**Elijah Pelofske** is working on his B.S. in Computer Science and Mathematics at New Mexico Institute of Mining and Technology. He is interested in the fields of machine learning, quantum computing, cryptography, algorithm design, data science, tensor networks, and novel computing architectures. He has published several journal and conference papers in the fields of quantum computing and machine learning. He has written several novel quantum annealing open source python tools. He won the Los Alamos National Laboratory Distinguished Student award in 2019.




**Dr. Lorie M. Liebrock** is the Director of the New Mexico Cybersecurity Center of Excellence (NMC-CoE) and is a Professor of Computer Science and Engineering. Her work with the NMCCoE leads projects including a Summer Institute with Sandia National Laboratories, an expansion project for Codebreaker Challenge, and an economic development project to help New Mexico companies prepare for Cybersecurity Maturity Model Certification. In addition, she is the Chair for the Transdisciplinary Cybersecurity graduate programs (Master of Science, Professional Masters, and PhD). She is also a Professor of Computer Science and Engineering. Her background in graduate education comes from years of experience as NMT's Graduate Dean. Dr. Liebrock has extensive research and real-world experience in cybersecurity, as well as parallel and high performance computing. She has published twenty-six journal articles, forty conference and workshop papers, and holds two US patents. Her research includes enterprise-wide cybersecurity, foundations of computer science, information assurance, parallel processing, and visualization with a focus on complex problems that require the integration of many aspects of computer science. This provides many student research opportunities, as she integrates students in her research - from freshmen to Ph.D. candidates. Dr. Liebrock holds both M.S. and Ph.D. in Computer Science from Rice University, B.S. and M.S. in Computer Science from Michigan Technological University, and an Associates degree from Delta Community College.




**Vincent Urias** is a computer engineer, and Senior Member of Technical Staff in Sandia's Cyber Analysis Research Development Department continuing to make major contributions to Sandia's cyber defense programs, especially in the simulation of complex networks, in developing innovative cyber security methods, and in designing exercise scenarios that test the limits of current network security. This work is helping Sandia's customers anticipate current and emerging security threats and make critical decisions about their investments. Vince and his team use technologies to conduct cyber defense exercises in partnership with

the U.S. Department of Defense, and to support national security in collaboration with colleagues at other U.S. Department of Energy national laboratories, Department of Defense national laboratories, and the U.S. military. Vince gives back to the community in a variety of ways, providing guidance and inspiration to college interns in the lab's Center for Cyber Defenders, he supports building computer labs for local organizations and is also helping to create an Urban Wildlife Refuge in Albuquerque's South Valley among other things. Vincent is currently pursuing his Ph.D. in computer science, at New Mexico Tech. He was honored by GMiS with a HENAAC Luminary Award in October of 2016.



(a) False negative rate comparison across models.



(b) False positive rate comparison across models.

Figure 14: CTC error rate comparison to single ML models

Text source and label	Num. of docs.	21 models incor. lbd.	individual models incorrectly labelled vector
Google security blog Cybersec	98	3	[0, 3, 5, 8, 63, 22, 4, 3, 0, 0, 0, 73, 11, 22, 47, 20, 3, 10, 98, 37, 54]
Darkreading Cybersec	98	4	[2, 9, 6, 8, 41, 23, 11, 6, 2, 0, 1, 83, 21, 29, 27, 17, 1, 11, 97, 30, 34]
Fireeye Cybersec	94	6	[2, 11, 5, 6, 56, 27, 3, 3, 3, 2, 2, 71, 10, 18, 51, 20, 3, 19, 94, 30, 52]
Ycombinator Cybersec	2	2	[1, 1, 2, 2, 2, 1, 1, 2, 2, 0, 2, 2, 1, 2, 2, 2, 1, 2, 2, 1, 2]
schneier Cybersec	35	5	[2, 7, 14, 12, 29, 7, 5, 4, 2, 1, 1, 20, 4, 7, 19, 13, 1, 7, 34, 14, 24]
toms-forum Cybersec	1	0	[0, 0, 0, 0, 1, 0, 0, 1, 0, 0, 0, 1, 0, 1, 0, 0, 0, 0, 1, 1, 1]
thehackernews Cybersec	92	19	[8, 20, 32, 30, 60, 47, 28, 13, 7, 2, 3, 69, 23, 35, 32, 16, 2, 10, 92, 34, 32]
Google security blog Not Cybersec	2	2	[2, 2, 2, 2, 0, 1, 2, 1, 2, 2, 2, 0, 1, 0, 0, 0, 2, 1, 0, 0, 1]
Darkreading Not Cybersec	1	0	[0, 0, 0, 0, 0, 0, 0, 0, 1, 1, 1, 0, 1, 0, 0, 0, 0, 0, 0, 0, 0]
Fireeye Not Cybersec	5	0	[2, 1, 0, 0, 0, 3, 1, 0, 0, 2, 1, 0, 2, 0, 1, 0, 3, 1, 0, 2, 0]
Ycombinator Not Cybersec	98	0	[22, 3, 2, 0, 0, 2, 14, 7, 14, 22, 15, 0, 19, 0, 1, 2, 22, 6, 0, 5, 0]
schneier Not Cybersec	65	15	[30, 19, 20, 21, 0, 41, 27, 11, 17, 31, 24, 4, 36, 6, 5, 15, 37, 19, 0, 38, 5]
toms-forum Not Cybersec	99	0	[14, 5, 0, 0, 0, 0, 12, 10, 16, 20, 19, 0, 12, 4, 5, 6, 6, 0, 0, 1, 0]
thehackernews Not Cybersec	8	0	[4, 1, 2, 0, 0, 0, 3, 1, 2, 6, 5, 0, 1, 0, 0, 0, 4, 2, 0, 0, 0]
ag-news Not Cybersec	95698	521	[13957, 8143, 1119, 1670, 0, 3832, 12603, 285, 862, 6435, 4565, 104, 4920, 345, 939, 1638, 14158, 9228, 112, 5929, 1319]
philosophy Not Cybersec	360808	434	[37654, 23525, 284, 454, 0, 1074, 35416, 34027, 52652, 13267, 23344, 59, 4831, 3913, 79093, 68829, 23065, 27708, 12, 2631, 4992]

Table 10: Number of incorrectly labelled documents across multiple text sources. The first column shows the labelled text source. Second column shows the total number of documents from that text source. The third column shows the total number of documents that the CTC tool (i.e., the majority vote of the 21 models) incorrectly labelled. The fourth column shows the number of incorrectly labelled documents by each of the 21 individual models.






Figure 15: Histogram of confidence measures from the Logistic model trained on Reddit text (top row), Arxiv text (middle row), and StackExchange text (bottom row) for the **ag-news** (left column) and **philosophy** (right column) datasets.