

Monitoring ‘Cyber Related’ Discussions in Online Social Platforms

Ruth Ikwu and Panos Louvieris

Brunel University London, United Kingdom

ABSTRACT

As the use of social platforms continues to evolve, in areas such as cyber-security and defence, it has become imperative to develop adaptive methods for tracking, identifying and investigating cyber-related activities on these platforms. This paper introduces a new approach for detecting “cyber-related” discussions in online social platforms using a candidate set of terms that are representative of the cyber domain. The objective of this paper is to build and evaluate these candidate terms for detection and tracking of cyber-related activities across various online platforms. The methodology presented in this paper applies natural language processing techniques to representative data from multiple social platforms to develop a representative cyber lexicon that can be applied to filter and text retrieval tasks in online social platforms. This paper also evaluates the terms’ performance in classifying discussions as ‘cyber-related’ or ‘non-cyber-related’. The results presented are most applicable to cyber threat monitoring and detection of malicious cyber activities in online social platforms.

Keywords: Cyber Lexicon, Natural Language Processing, Social platforms, Cyber-Incident Monitoring.

1 INTRODUCTION

The evolution of the World Wide Web provides new methods for communication in the hyper-connected world. These communications include sharing of texts, images and videos between geographically distributed cyber-personas. Security experts usually apply manual techniques for monitoring and detecting cyber-related activities on social platforms (Dionisio, Alves, Ferreira, & Bessani, 2019; Rowe & Saif, 2016; Varol, Ferrara, Menczer, & Flammini, 2017). However, given the ever-increasing volume of data generated on these social platforms, these manual techniques have become labour intensive, inefficient, and a significant challenge for security experts who hope to stay ahead.

The volume of text data currently generated daily online and the resultant exponential increase in the number of potential cyber-related activities requires robust techniques in analytics to identify cyber threats for early mitigation. One such technique is Natural Language processing of text data. NLP is rooted in many disciplines, such as computer science and computational linguistics. It is yet another branch of Artificial Intelligence (AI) which gives computers a better understanding of, interpretation of, as well as the ability to manipulate human language (Halper, 2017; SAS Institute, 2018). The primary aim of NLP is a step towards bridging the communication gap between human and machine.

Language processing techniques that support cyber situational awareness on the social dimension of cyberspace need more domain-specific lexicon than what is available today. For example, several studies and historical events have demonstrated the ability of individuals and groups to use social platforms as enablers for perpetrating cyber-incidents (Hernández, Sanchez, Sánchez, & Pérez, 2016; Olsen, 2013; Zetter, 2014). In July 2006, the hacker group ‘anonymous’ recruited new members, planned, organised and facilitated the famous raid on Habbo Hotel using the ‘4chan’ microblogging platform (Bernstein et al., 2011; Knuttila, 2011). Within the next decade, between 2008 and 2012, several hacker groups such as ‘Anonymous’ and its spin-off ‘LulzSec’ used social chat rooms to recruit, train new hackers, plan and coordinate cyber-attacks with hashtag trends such as #OperationPayback (Mackey, 2010; Pras et al., 2010). Furthermore, in popular and open social platforms like Twitter, the use of trending hashtags such as #OpLibtard, #OpISIS, #OpPedos, #OpIsreal are open tags for tracking discussions related to certain cyber-activities being carried out. A gap exists in automatically spotting these sorts of conversations on these platforms, given the volume of data generated.

This paper addresses this gap by providing a methodology for automatically identifying cyber-related discussions in online social platforms. The goal is to create a simple method for monitoring conversations on social platforms and detecting texts that are related to cyber activities. We aim to develop a generic lexicon of candidate terms that capture the context of analytical interest. We achieve this by creating a representative corpus of cyber-related discussions using data from popular cyber microblogging platforms such as Reddit, Twitter, stack overflow, hacker news and cyberwar news. We build a lexicon of cyber-related or context-specific terms that are known to appear frequently with a higher degree of relevance to the context of analysis, in these discussions. We rank each term based on scoring mechanisms that emphasise the frequency of term occurrence, the relevance of terms in sentences and the mutual dependence amongst terms. To evaluate the performance of the lexicon, we apply it to a ‘cyber-related’ quantification task on a set of new labelled discussions from a collection of social platforms. A quantification algorithm estimates the degree of ‘cyber-relatedness’ of random text and classifies each text into one of two classes – cyber-related or non-cyber-related -- based on an optimally selected threshold.

The main contributions of this paper are of two folds; firstly, it presents a novel approach for building a cyber domain-specific lexicon and offers a curated list of context-specific terms that are representative of the cyber domain. The lexicon contributes to already existing methods for monitoring and detecting cyber-related activities on social platforms (Khandpur et al., 2017; Lippmann et al., 2016; Sri, Yellari, & Rao, 2017).

2 MINING SOCIAL PLATFORMS FOR CYBER-RELATED DISCUSSIONS

The first challenge in building a lexicon in the context of the cyber domain is to generate a representative set of “cyber-related” discussions from these social platforms. Any strategy for gathering such information must include relevant texts from the appropriate social channels. Social media platforms are characterised by the type of content generated, user experience and policies that guide interactions on these platforms (Chen, Xu, & Whinston, 2009). To accurately identify all kinds of sources of relevant data, we define three types of social platforms based on the level of moderation of user-generated content (McKenzie et al., 2012).

2.1 The Categorisation of Social Platforms

The level of moderation of user content is crucial as it defines the manner of conversations happening on a social platform. We observe such categorisation from the degree to which users can freely express their views on these platforms (unhindered) and therefore, the extent to which they can use them for personal activities.

Pre-Moderated Social Platforms: Administrators of these platforms highly moderate user-generated content before it appears online. Pre-moderation ensures that posts generated by users but deemed “inappropriate” by site administrators never make it online. These sites do not allow users to create personalised topics for discussions but rather, contribute to topic categories created by the site administrators. Therefore, these types of social platforms have highly regulated topic forums where users can participate. Additionally, due to the need for highly regulated online communications on these platforms, discussions do not occur in real-time. While these platforms are useful for controlling inappropriate use and stamping out cybercrimes, such rules have been known to lead to the death of online communities (McKenzie et al., 2012).

Post-Moderated Social Platforms: Administrators of these sites are minimally involved in the moderation of its content. As opposed to pre-moderated content, user-generated content appears online immediately after posting but queued for moderation. This level of moderation allows for real-time communication in online communities. Eventually, user-generated content deemed inappropriate by site moderators is filtered, hidden or deleted. Post-moderation filters displayed posts generated by users, only when deemed “inappropriate” by site administrators. On these platforms, users are sometimes allowed to create their topics of discussion and contribute to topics created by other users. While these platforms offer flexibility as opposed to pre-moderated platforms, they are at risk of accommodating inappropriate content if response site moderators do not respond quickly. Examples of social networks within these categories include Stack Overflow, Reddit, Stack Exchange, Quora (McKenzie et al., 2012).

Reactively-Moderated Social Platforms: Administrators of these sites are rarely involved in moderating user-generated content. As a result, content appears online immediately, and precisely as created. Reactive moderation means site moderators rely on users to report inappropriate content when they see it. User-generated content is therefore only checked if a complaint is made about them by other users. Site administrators remove reported

posts if deemed necessary by site moderators. This level of moderation ensures that users create and engage in topics they want. Users with similar ideologies can belong to the same sub-group where they share thoughts and beliefs without restrictions. Users are usually allowed to create their topics of discussion, monitor and contribute to topics created by other users. Therefore, these types of social platforms encourage highly effective real-time communication with minimal restrictions. While these platforms offer flexibility and freedom, they are known to be safe havens for cybercriminals. Examples of social platforms in this category include Twitter, Facebook, Snap Chat, Tumblr, 4chan (McKenzie et al., 2012).

Identifying text in online platforms that are relevant to a specific topic of interest is usually done with a keyword-based approach (Kaji & Kitsuregawa, 2007; Ntoulas, Pzerfos, & Cho, 2005). In a keyword-based approach, a set of terms are used to filter through text, and only texts containing any word in the set of terms are returned (Bruns & Yuxian Eugene, 2012). However, with no prior assumptions of context-related keywords, it is useful to target forums where users are actively engaged in related topics of interest.

On real-time microblogging platforms like Twitter and Facebook, a keyword-based approach may be appropriate to filter tweets on specific topics or from specific user accounts (Bian, Yang, Zhang, & Chua, 2015; Starbird, Muzny, & Palen, 2012). However, specific users and hashtags that are known to be related to events on the cyber domain are also a good source of related discussions. With content-organised platforms like Reddit and Stack Overflow, identifying a handful of context-related channels where users are actively engaged in relevant types of discussions to have shown to return better samples for context analysis (Lippmann et al., 2016).

2.2 Lexicon Building

The aim of building a lexicon is to extract a set of terms that capture the context of lingual analytical interest. From a sample of ‘domain-related’ texts, we hope to obtain a set of terms that are seen to appear frequently within these texts. In addition to the frequency of term occurrence, we aim to discriminate terms based on their level of importance and associative mutual dependence on other terms in the sentence.

Typically, when building domain-specific lexicons, two design approaches are considered and most often combined: selecting and grouping terms within some predefined categories (Kipper, Korhonen, Ryant, & Palmer,

2006) and weighting terms based on the domain of analysis (Baccianella, Esuli, & Sebastiani, 2010) and usefulness of terms to the topic of interest (Jurafsky & Martin, 2017).

The first step in most lexical creation tasks is the candidate terms generation. This step involves creating a wordlist or dictionary of terms that are known to appear frequently in discussions of interest. Keyword-based extraction of terms is most common in this step and has been instrumental to creating traditional lexicons (Ntoulas et al., 2005; Olteanu, Castillo, Diaz, & Vieweg, 2014) for natural language and information retrieval tasks such as sentiment analysis (Baccianella et al., 2010). For example, (Nielsen, 2011) creates a wordlist of positive and negative words for ranking text documents on a scaled range of -1 to 1 (-1 indicating a strongly negative text document and +1 indicating a strongly positive text document). Allahyari et al. (Allahyari et al., 2017) apply a similar approach to a multi-class text classification task. Similarly, Rose et al. (Rose, Engel, Cramer, & Cowley, 2010) keyword-based approach provide context to the terms used for lexical analysis by creating a network of lexical-semantic relations between words in a document corpus, where the meaning of each term in the lexicon is defined within the context of its relationship with other terms.

Simple keyword-based candidate term selection approach is further extended to include methods for term scoring. For each term that makes it into the candidate set, scoring techniques evaluate the importance of that term relative to other terms in the candidate set (Kaji & Kitsuregawa, 2007) — for example, quantifying the relationship between the terms ‘vulnerability’ and ‘malware’ in a document set. Popularly in research, terms are numerically scored by two main techniques: frequency-based scoring techniques (Blumenstock, 2008; Rose et al., 2010; Zhang, Jin, & Zhou, 2010) and association-based scoring techniques (Debole & Sebastiani, 2003; Jurafsky & Martin, 2017; Khan, Qamar, & Bashir, 2016). Frequency-based scoring techniques rank scores based on the number of times they occur within sentences. These methods usually address issues with stop words such as ‘I’, ‘is’, ‘then’, ‘that’, ‘have’, ‘has’. – that have no contextual meaning but have a high-frequency score due to the nature of their usage in the English language. Scoring techniques based on associative dependence further simple frequency-based scoring to address issues of term importance in the context they occur.

However, since the basis of most of these term scoring techniques is on semantic relationships between words in a text document (independent of externally perceived meanings), most of these techniques are blind to the

domain of analysis in which they occur. For example, consider these two phrases that belong to the same text corpora taken from two different subreddits: ‘MySQL Database developer needed urgently for a 2-month project in Belfast.’ and ‘New MySQL database vulnerability found on Windows operating system. Yet again!!’ Assuming a single domain of analysis, the term ‘database’ in the midst of other domain-related terms such as: [‘vulnerability’], [‘operating’ and ‘system’], [‘MySQL’], should have a higher-ranking score than the same word in the previous phrase.

This paper aims to develop a lexicon that is usable on most social platforms to automatically detect cyber-related messages. The data comprises of a standard set of discussions from various social platforms, representative of these cyber-related discussions.

3 DATA COLLECTION AND PRE-PROCESSING

The datasets used in this study are from five social platforms; Twitter (twitter.com), Reddit (reddit.com), Stack Overflow (stackoverflow.com), Cyber War News (cyberwarnews.info) and The Hacker News (thehackernews.com). These platforms were selected as they represent the types of social networks and content generated discussed in this paper.

Twitter represents a self-moderated real-time microblogging platform where users communicate uninterrupted and unedited. Twitter is a real-time reactively moderated social network platform where moderators rely on users to report content deemed inappropriate. Users populate twitter timelines with short (maximum of 280 characters) non-curated posts. Hashtags identify tweets on a single topic, ‘@’ represents users involved in a discussion thread. Cyber discussions on twitter cover a wide range of cyber event types from individuals and groups on cyber hacktivism, cyber warfare, cyber-crimes and cyber terrorism activities.

Reddit is a collection of minimally moderated sub-forums called subreddits with multiple topic threads and comment discussions in each subreddit. Users populate Reddit forums with short to medium-length posts and comments. Individuals are allowed to freely start and participate in conversations on their topic of interest. Sub-reddits dedicated to cyber discussions cover a wide range of cyber-related issues such as cyber-crime, cyber-warfare, cyber-hacktivism and cyber-terrorism.

Stack overflow is a highly-moderated question and answer community with sub-communities a wide range of technical content. Questions and answers






on stack exchange can be long, medium or sometimes short text. Discussion forums are highly moderated, and content is filtered based on what forum moderators classify ‘appropriate’ or ‘relevant’ to the forum.

Cyberwar news is an archive of articles on cyber events. Cyberwar news provides long curated, moderated and edited news articles with details on cyber events. Cyber event details include details of attackers, details of the victim and technical aspects of the attack.

Similar to Cyberwar news, the Hacker News is a widely-acknowledged cybersecurity news platform, with over 8 million active readers monthly. Readers include IT Professionals, academic researchers, hackers and technology experts. The platform features news on the latest events in cybersecurity and extensive coverage of current and future trends in information security.

3.1 Data Collection

This work aims to build a lexicon with a set of candidate terms to automatically identify texts that are related to the cyber domain. After identifying potential data sources for representative cyber-related discussions, we create a generic text corpus that is a collection of documents from all data sources.

	Data Source	Social Platform Type	Cyber Discussion Type	Number of Documents	% Corpus
1	 stackoverflow	Question and Answer Forum	Long, Medium or Short, Minimally-moderated user-generated content.	190,970	17.8%
2	 reddit	Discussion /Community Forum		333,882	31.2%
3	 twitter	Microblogging Site	Short	507,348	47.4%
4	 CYBERWAR NEWS	Blog/News Site	Long Heavily Moderated, Edited News-like Articles	10,568	1%
5	 The Hacker News	Blog/News Site	Long Heavily Moderated, Edited News-like	28,114	2.6%

			Articles		
	Total Number of Documents			1,070,882	

TABLE 1: DATA SOURCES FOR LEXICON DEVELOPMENT

From blogging platforms such as cyberwar news (cyberwarnews.info) and the hacker news (thehackernews.com), we collect a total of 1550 cyber event news articles. Each article is sentence tokenised, and each sentence is a single document in the final corpus. Cyber news articles such as on cyberwar news and hacker news, provide details of cyber events on social media such as the name of the operation, target, hacker(s) (if available), date of the event and additional social media details.

To identify relevant users and therefore construct a useful keyword-based filter for Twitter, we extract 547 twitter hashtags (189) and mentions (358) from the curated cyber news articles collected from cyberwar news and the hacker news. For example, cyber operations on Twitter are tagged with the naming convention “#Op[Name of Operation]” (#OpPayback). This list of extracted handles, hashtags and mentions are further used as filters to collect historical twitter discussions from start of topic discussions up until the 18th of June 2018. Each tweet is also added as a single document to the final corpus.

All user comments on each post from thirty-eight cyber-related subreddits were also gathered and processed. Each post and user comment are sentence-tokenised and added as a single document to the final corpus. Furthermore, questions, answers and posts from ~900 stack exchange question threads with at least a thousand (1000) votes up until the 18th of June 2018 were also collected and added to the corpus. Non-text lines were excluded from the collection process.

The search space for data collection in this paper is significantly reduced to only sources of cyber-related discussions online. Therefore, tweets, comments and posts from all data sources were collected under the assumption of being a post about cyber-related activities.

3.2 Data processing

We combined tweets, Reddit comments, stack overflow posts and blog articles from cyberwar news and hacker news into a single text corpus. We remove duplicate tweets and retweets. We ensure that each data source contributes a significant number of documents to the final corpus to achieve

an even distribution of document types. A document represents a single tweet, a question, an answer, a sentence, a post or a comment from any of our data sources. Long curated articles (text) such as from cyberwar news are split into individual sentences, where each sentence is a document in our corpus. Splitting long documents ensures that the length of each document in the final corpus remains within the same range of word count.

We use speech tagging to identify and remove entities such as names, places, people, things, events and time from each document. Speech tagging splits each document into samples of parts-of-speech identification, which act as markers to find people, places, dates, time and other related entities. Parts of speech tagging identifies the function a word plays in a sentence. Table 2 below shows an example of a document with speech-tagged words. *“Excerpt from Michael Hastings’ “The Operators” Re: Death Threats <http://bit.ly/198uVBl> #Anonymous #OpIsrael #OpPalestine #Gaza.”*

SN	TERM	TAG	DESCRIPTION
1	Excerpt	NN	Noun Singular Or Mass
2	From	IN	Preposition or subordinating conjunction
3	Michael	NNP	Proper Noun, Singular
4	Hastings	NNP	Proper Noun, Singular
5	The	DT	Determiner
6	Operators	NNPS	Proper Noun, Plural
7	death	NN	Noun, Singular Or Mass
8	threats	NNS	Noun, Plural
9	anonymous	JJ	Adjective
10	oplsreal	NNP	Proper Noun Singular
11	OpPalestine	NNP	Proper Noun Singular
12	Gaza	NNP	Proper Noun Singular

TABLE 2: SAMPLE OF SPEECH-TAGGED TWEET

“Always assume the adversary knows the method, see Kerckhoffs’ principle linked in our sidebar.”

SN	TERM	TAG	DESCRIPTION
1	Always	NNS	Noun, Plural
2	Assume	VBP	Verb, non-third person singular present
3	The	DT	Determiner
4	Adversary	NN	Noun, singular or mass
5	Knows	VBZ	Verb, third person singular present
6	The	DT	Determiner
7	Method	NN	Noun, singular or mass
8	See	VBP	Verb, non-3rd person singular present
9	Kerchoffs	NNP	Proper noun, singular
10	Principle	NN	Noun, singular or mass
11	Linked	VBD	Verb, past tense
12	in	IN	Preposition or subordinating conjunction
	Our	PRP\$	pronoun, possessive
	Sidebar	NN	Noun, singular or mass

TABLE 3: SAMPLE OF SPEECH-TAGGED REDDIT COMMENT

We use the compendium of Penn Treebank (Santorini, 1990) as a standard for speech tagging terms in our corpus. We remove all types of nouns (NN, NNS, NNP, NNPS, POS), all forms of pronouns (PRP, PRP\$, WP, WP\$), prepositions (IN, EX), conjunctions (CC), determiners (DT, PDT, WDT), articles (TO, RP) and other unwanted terms (FW, MD, SYM) leaving only verbs, adverbs and adjectives. Additionally, we also remove URLs, user mentions, hashtags and emotion encodings from each document. Finally, we remove forum specific words from corresponding documents. For example, the word “post” from twitter documents, “votes”, “upvotes” and “downvotes” from stack overflow documents, the words “subreddit”, “Reddit” and “forum” from Reddit documents.

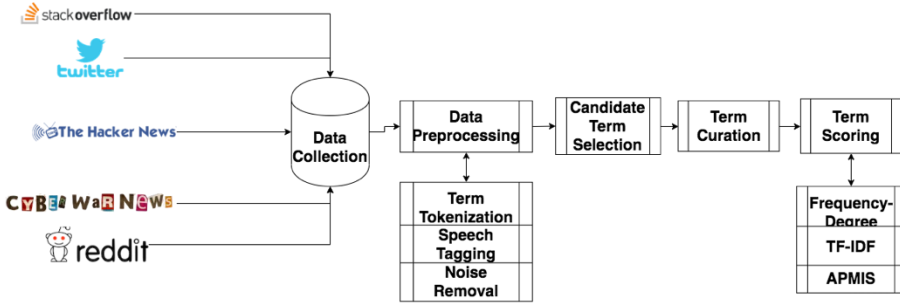


FIGURE 1: LEXICON DEVELOPMENT PROCESS

4 BUILDING THE LEXICON

Figure 1 shows the steps involved in developing our generic ‘cyber-related’ lexicon. Our process starts by selecting a set of candidate terms for our generic lexicon. Each candidate term is assigned three scores; FDR, TF-IDF and an APMIS. The FDR represents the term’s frequency of occurrence across all documents in the corpus. The TF-IDF score represents the average of the term’s importance across all documents in which it occurred. The APMIS score represents the associative dependence of the term based on its point-wise mutual information with other terms in the document. Finally, a cutoff at the Nth percentile is used to select the *Top (k)* terms from each scoring mechanism.

4.1 Candidate Term Selection

Our candidate terms are based on term frequency, term unigrams and term entropy in across all documents. First, each document is term tokenised across the corpus producing a set of 2,434 initial candidate terms. After term tokenising sentences, we perform an initial curation step to remove words that are less than two characters long, greater than 15 characters and/or with term entropy of 0. The entropy of terms is estimated using the text perfection method described in (Revovna Ospanova, 2013). We estimate the entropy of each term as $TermEntropy = -(p * \log(p))$ where p is the probability of occurrence for each individual character in the term, estimated as:

$$\frac{Numberofuniquecharacters \in term}{NumberofCharacters \in term}$$

Using the term entropy eliminates terms with no substance—for example ‘aaaaa’, ‘ahhhh’, ‘hahaha’. We also remove text noise in the form of numbers, punctuation and stop words. Finally, we stem the remaining words using Porter’s Stemmer (Willett, 2006). The candidate term selection step returns 1149 candidate terms.

4.2 Term Scoring

Our term scoring strategy is robust in that it includes terms ranked by three different scoring mechanisms. In order to ensure a robust lexicon, we rank terms based on their co-occurring frequency with other terms, relevance and mutual association with other terms across all documents. For each term in our candidate set, we estimate its *Term Frequency-Inverse Document Frequency (TF-IDF)* (Ramos, 2003) to measure its relevance, *Frequency-Degree Ratio (FDR)* (Rose et al., 2010) to quantify its overall co-occurring frequency and *Aggregated Pointwise Mutual Information Score (APMIS)* to estimate its associative dependence. The frequency degree ratio scoring ensures that the lexicon includes terms, frequently occurring with other terms. Similarly, the Term-Frequency Inverse Document Frequency scoring mechanism ensures that the lexicon includes rare but relevant contextual terms. Finally, the Aggregated Pointwise Information scoring ensures that the lexicon includes terms with greater informative association with other terms across documents.

The term scoring mechanisms used in this paper starts with an initial estimation of a term-document matrix. A term-document matrix, $tf_{t,d}$, is a sparse matrix representation of the terms’ weights in each document in the corpus (Jurafsky & Martin, 2017). We start by creating the term-document matrix $tf_{t,d}$, where each row is a document d , each column is a term t and, each cell represents the number of times the term t occurs in the document d . Note that the columns in a term-document matrix represent terms across all documents in the corpus, therefore creating a highly sparse matrix representation of term weightings of zeros across documents. To address the sparsity of the term-document matrix, we remove terms that only appear in at least 90% of documents in the combined corpus.

4.2.1 Term-Frequency Inverse Document Frequency (TF-IDF)

While term frequency of a term t refers to the number of occurrences of the term t across all documents in the corpus, the document frequency of a term t , refers to the number of documents in which t occurs. Dividing the term-

frequency of each term by its corresponding document frequency produces a matrix that is directly proportional to the frequency of occurrence of each term in a document but inversely proportional to the number of documents it occurs. The inverse document frequency score diminishes the weights of frequently occurring terms and increases the weights of terms that occur rarely. The TF-IDF therefore, assigns weights that quantify relevance of terms in documents. The Inverse Document Frequency (IDF) (Spärck Jones, 1972) is estimated as:

$$idf_t = \log \left(\frac{N}{df_t} \right)$$

We, therefore, estimate the term-frequency inverse-document frequency score (TF-IDF score) for each term as a product of the term frequencies and the inverse document frequency.

$$TF_{t,d}idf_t = (TF_{t,d}) \left(\log \left(\frac{N}{df_t} \right) \right)$$

Where the total term-frequency of each term ($TF_{t,d}$) is estimated as the sum of raw counts of a term t in each document d divided by the total number of terms in d .

4.2.2 Frequency Degree Ratio (FDR)

The summed term frequency, ***freq(t)***, refers to the number of times each term appears across all documents while term degree, ***degree(t)***, is a measure of co-occurrence of each term with other terms across all documents in a given corpus (Rose et al., 2010). The term degree is a term-term first-order co-occurrence matrix (Schütze & Pedersen, 1993) $tt_{t,t}$ which represents the number of times a term t_i occurs with another term t_j within the same context across all documents in the corpus. Co-occurrence expresses links of relationships between texts, therefore acting as an indicator of cohesion between individual terms in a document (Mihalcea & Tarau, 2004). In a co-occurrence matrix or term-term matrix, each term is represented as a numeric vector of the number of occurrences with other terms. Each term in our candidate set is represented as a potential context term. Our co-occurrence matrix is an $|M|X|M|$ matrix where M = number of terms in the candidate set. The summed term degree, ***degree(t)***, is the sum of term-term co-occurrence of term t with all other candidate terms. The term degree, therefore, measures the importance of each term in a document relative to other documents in the entire corpus.

For the Frequency-degree ratio scoring, each term in the resulting candidate set is initially weighted on these two scores: the summed term frequency and the summed term degree of terms. Note that our measure of term degree, *degree(t)* subtracts the number of self-term co-occurrence (i.e. the number of times a term co-occurs with itself) from the summed total of term degree (by equating the diagonals of the term-term matrix **M** to 0). Similarly, we normalise our measure of term frequency of each term with the total number of terms in the corresponding document. The ratio of term frequency to term degree $\frac{degree(t)}{freq(t)}$ for each term in our candidate set is estimated to give the FDR term score.

4.2.3 Aggregated Pointwise Mutual Information (APMIS)

Frequency-based algorithms, however, are not the best measures for associations between terms as they are not discriminative to terms with superior level of relevance given the context of analysis. To extract a measure for the degree of context shared between individual pairs of terms, we replace the frequency with the pointwise mutual information score between the two terms. We estimate the informative dependence between two terms as their pointwise mutual information. The pointwise mutual information between two terms $t_1 \wedge t_2$ measures the amount of informative association between them (Kenneth & Hanks, 1990), i.e. the probability of observing a term t_1 with another term t_2 as opposed to observing them independently. In its simplest application, a set of terms in a candidate set is measured against a set of ‘context-terms’, usually representing a given context of analysis. It can also be described as the logged independent joint probability of occurrence between $t_1 \wedge t_2$. The PMI between two terms $t_1 \wedge t_2$ is estimated as:

$$PMI(t_1, t_2) = \log_2 \left(\frac{P(t_1, t_2)}{P(t_1)P(t_2)} \right)$$

We use our candidate terms as context terms for analysis. Therefore, our PMI Matrix is a term-term representation of the pointwise mutual information between each pair of candidate terms. The PMI estimate ranges from $-\infty$ to $+\infty$. To compute the PMI matrix from a term-term matrix **M** with n rows (terms) and n columns (context-terms; candidate terms in our case), we estimate each cell as:

$$pmi_{ij} = \log_2 \left(\frac{p_{ij}}{p_i * p_j} \right)$$

The pointwise mutual information (PMI) matrix shows the PMI score for each ‘term-term’ combination of terms in our candidate set. For example, given the term-term matrix below, the aggregated sum of co-occurrence of the term ‘attack’ with all other terms in the matrix is 3182. Similarly, the aggregated sum of co-occurrence of the term ‘target’ with all other terms in the matrix is 1398.

	Attack	Hack	Target	Code	Activist	zionist
attack	0	2372	731	46	63	27
hack	2372	0	611	48	72	59
target	731	611	0	18	22	16
code	46	48	18	0	4	0
activist	63	72	22	4	0	2
zionist	27	59	16	0	2	0

TABLE 4: TERM-TERM MATRIX WITH TERM-TERM CO-OCCURRENCES SCORES

If the sum of the ‘term-term’ matrix is 8182, we can estimate the PMI between the term “attack” and the context-term “target” as:

$$PMI_{attack,target} = \log_2 \left(\frac{p(attack,target)}{p(attack) * p(target)} \right)$$

Where “attack” is the term and “target” is the context term.

$$P(attack) = 3182/8182 = 0.39$$

$$P(target) = 1398/8182 = 0.17$$

$$P(attack, target) = 731/8182 = 0.09$$

$$PMI_{attack,target} = \log_2 \left(\frac{0.09}{0.39 * 0.17} \right)$$

Therefore, there is a 44% probability of observing the terms ‘attack’ and ‘target’ together in the document corpus. Estimating this value for each term-term (context) combination produces the following matrix.

	<i>Attack</i>	<i>Hack</i>	<i>Target</i>	<i>Code</i>	<i>Activist</i>	<i>zionist</i>
<i>attack</i>	0.00	0.925	0.401	-0.049	-0.069	-0.632
<i>hack</i>	0.925	0.000	0.178	0.044	0.152	0.475
<i>target</i>	0.401	0.178	0.000	-0.106	-0.300	-0.114
<i>code</i>	-0.049	0.044	-0.106	0.000	1.182	0.199
<i>activist</i>	-0.069	0.152	-0.300	1.182	0.000	0.741
<i>zionist</i>	-0.632	0.475	-0.114	0.199	0.741	0.000

TABLE 5: TERM-TERM MATRIX WITH TERM-TERM CO-OCCURRENCE PMI SCORES

The PMI scores for term-based analysis is known to be discriminative to infrequent terms where infrequent terms have very high PMI values (Jurafsky & Martin, 2017). Jurafsky & Martin proposes two solutions to this problem: a) higher probability assignments by raising context-term probabilities to $\alpha=0.75$ and b) using the Laplace [add-2] smoothed values of the co-occurrence matrix. Therefore, the APMIS for each term is the sum of PMIs of a single term with all context-terms. For example, using the table above

$$APMIS_{attack} = 0.000 + 0.925 + 0.401 - 0.049 - 0.069 - 0.632 \\ = 0.576$$

Therefore, there is a 57% probability of observing the term ‘attack’ with other terms in the matrix in table 5.

Finally, the table below shows the top 20 cyber-related terms using each of the scoring criteria.

SN	TF-IDF	FDR	APMIS
1	http	administr	actual
2	enter	agenc	allow
3	run	action	account

4	data	analyz	back
5	free	amount	activ
6	find	aim	address
7	window	breach	addit
8	actual	access	access
9	differ	aspect	app
10	file	account	basic
11	user	address	assum
12	call	autom	base
13	prize	applic	avail
14	ticket	attempt	applic
15	open	attack	appear
16	chang	capabl	attack
17	read	admin	affect
18	creat	advis	android
19	start	associ	anonym
20	key	alert	break

TABLE 6: TOP 20 TERMS (FDR, TF-IDF, APMIS SCORING).

The difference in term ranks for each scoring mechanism, in summary, the FDR word score favours co-occurring term frequency, the TF-IDF score will favour useful terms that often occur across documents while the *APMIS* will favour terms with greater association and dependence with other terms across documents.

4.3 Term Curation and Top-terms Selection

The curation step further removes unwanted terms to yield better-filtered results. We also remove terms contextually associated with specific cyber events, users, Reddit forums and hashtags. For example, terms such as the name of Reddit forums or twitter usernames and handles. Additionally, names of specific cyber events on Twitter such as tagged cyber operations,

e.g. #OpPayback or specific users such as @anonr00t were removed from the resulting word list. After candidate terms have been selected, we select terms whose term scores meet an optimal threshold K . ' K ' represents a threshold of terms scores for including corresponding terms in the lexicon.




FIGURE 2: DISTRIBUTION OF TERM SCORES

4.3.1 Term Score Cutoff Selection

We determine the cutoff ' K ' by observing the changes in the term scores at various percentiles, as shown in figure 2. Figure 2(a), 2(b) and 2(c) plot the density distribution of term scores. The red lines indicate the cutoff percentile at which there is at least an increase twice as much as the previous increase. The right-skewed distributions show that about 10% (above the 90th percentile) of all terms from each scoring mechanism meet the threshold.




FIGURE 3: TERM SCORE CUT-OFFS



Additionally, as seen in table 6 above, the terms whose scores meet the 90th percentile cutoff, returned by each scoring mechanism significantly contain terms that are representative of specific cyber event categories (Hathaway et al., 2012). For example, the *Top(K)* terms for the term frequency – term degree ratio returns terms that are characteristic of cybercrimes while the *Top(K)* terms for APMIS returns terms that are characteristic of cyber warfare and cyber espionage. To get a representative sample of terms that characterise various types of cyber incidents, we include all terms whose scores meet the optimal cutoff in the lexicon.

$$Cutoff(K)_{basicscore}, Cutoff(K)_{TF-IDF} \wedge Cutoff(K)_{APMIS}.$$

In the end, this cyber lexicon should be able to track various types of cyber-related discussions on social platforms. The application of this lexicon depends on the kind of analysis conducted. For example, the lexicon can be used in quantifying a random text to estimate its degree of cyber-relatedness or used in querying and tracking discussions on social forums. The final cyber Lexicon is a set of terms with their corresponding term scores for each scoring criteria.

4.4 Evaluating the Lexicon

To test the real-world application of the cyber lexicon, we create a new text corpus with documents collected from various social media platforms and measure the degree of ‘cyber-relatedness’ of documents in the corpus. We build the sample text corpus of cyber and non-cyber related documents (texts) obtained from the web and apply the lexicon to each document in the corpus. The test corpus is from a representative sample of all types of cyber-related discussions on various social platforms. The table below shows the data sources, amount of cyber-related and non-cyber-related texts collected for our test text corpus.

	Data Source	Social Type	Platform	Cyber-related Texts	Non-Cyber-related Texts	Total Documents	% of Test Corpus
1		News Site (Long edited articles)		15	15	30	13.7%
2		Microblogging Comments (Short –		15	13	28	12.8%



		Medium minimally moderated user comments)				
3		Blogging/News Site (Long edited articles)	14	0	14	6.4%
4		Microblogging Comments (Short – Medium minimally moderated user comments)	15	14	29	13.3%
5		Discussion Forum	15	15	30	13.7%
6		Discussion Forum	15	15	30	13.7%
7		Question and Answer Forum	14	14	28	12.8%
8		Blogging Comments (Short – Medium minimally moderated user comments)	14	15	29	13.3%
	TOTAL		117	101	218	≈ 100%

TABLE 7: DATA SOURCES FOR EVALUATION DATA

4.4.1 Evaluation Strategy

For each data source in the table above, we collect a set of ‘cyber-related’ texts and non-cyber-related texts. We collect cyber-related texts from comments in forums or sections tagged as being related to cyber-attacks or incidents and non-cyber-related texts from random sections of forums such as fashion, entertainment, economy and finance. We label posts under cyber-related topics as ‘cyber’ (or 1), and posts on other topics as ‘non-cyber’ (or 0). Finally, we curate 117 cyber-related and 101 non-cyber-related texts from 8 social platforms-as shown in table 7- to create the test corpus. Note that there is an approximately equal number of documents assigned to each category.

To calculate the degree of cyber-relatedness of a document using the lexicon, we rely on a bag-of-words-based algorithm that is a function of the

document's terms and terms in the cyber lexicon. For this evaluation, the estimate of 'cyber-relatedness' of a given text is a measure of the total term scores of individuals words in the text matched to terms in our generic lexicon.

Given a random text and the lexicon with respective term scores, the 'cyber-relatedness' is measured by the process as presented below:

Pseudo-Code: Estimating the Cyber-Relatedness (CR) of a Random Text		
Input: L < Cyber Lexicon [Term: Score] >, S < A Random Text or Sentence >		
Output: CR < Numeric Quantity for the Cyber-relatedness of S >		
1: set match = 0;		<< (a)
2: set sum_scores = 0;		<< (b)
3: set words = SentenceTokenise(S);		
4: set wordcount = length(words);		<< (c)
5: For each word in words:		
6: If word in L.Terms:		
7: set sum_scores = sum_scores + L.Term.Score		
8: set match = match+1		
9: End		
10: End		
11: set scalar = match ÷ wordcount		<< (d)
12: set CR = sum_scores * scalar		
13: return CR		

TABLE 8: PSEUDOCODE-ESTIMATING CYBER-RELATEDNESS OF A RANDOM TEXT

The pseudocode above demonstrates the steps taken to quantify the degree of cyber-relatedness for a plain text. The expected output is a numeric quantity of how 'cyber-related' the sentence is. The Pseudocode above has two inputs: a) a cyber lexicon with cyber-related terms and respective term scores and b) a sentence of which is to be quantified. The pseudocode has four case points:

- match: this tracks the number of words in sentence matched to terms in cyber lexicon,
- sum_scores: this tracks the sum of term scores for matched terms,
- wordcount: this is the number of words in the given sentence,
- scalar: this scales the total matches found by the total number of words in the sentence.

The process splits the sentence into single words and searches the lexicon for a match on each word. It sums up the term scores of each matched term

and scales it by a scalar quantity. The final CR score is a sum of the term scores multiplied by the scalar quantity.

4.4.2 Classifying Documents

Document classification is the task of grouping documents in our test corpus as ‘cyber-related’ or ‘non-cyber-related’ based on their CR score. After pre-processing estimating the CR score of each text in the test corpus with the process described in Table 8, we re-scale the CR scores on a scale of 0 through 100. Classifying each document given its re-scaled CR score is treated here as a two-class classification (binary classification) task in a supervised learning environment. The actual class for each document is the assigned classes (‘cyber-related’ or ‘non-cyber-related’) from the previous document labelling phase.




FIGURE 4: RECEIVER OPERATING CURVE

Therefore, we determine a threshold to specify a classifier boundary (a threshold) between the two categories. To select an optimal classification threshold, we use the ‘Receiver Operating Characteristic’ curve (ROC)

analysis (Hanley & McNeil, 1982). Typically, the ROC curve is created by plotting the recall or sensitivity against the false positive rate given a set of estimated and actual values.

We select a threshold value that maximises the ‘area under the ROC curve’ (Hanley & McNeil, 1982) as the optimal boundary for our classification task. The intersection of the two red lines in the figure above indicates the optimal value (0.19) for a threshold. A threshold of 0.19 maximises the probability (0.77) that the lexicon will score a randomly selected cyber-related document higher than a randomly chosen non-cyber-related document.

We create a confusion matrix from comparing the classes produced by applying the lexicon and the actual classes of each document. The confusion matrix in the table below shows the number of cyber-related documents our generic lexicon correctly classifies as ‘cyber-related’ TP_s , the number of cyber-related documents it incorrectly classifies as ‘non-cyber-related’ FN_s , the number of ‘non-cyber-related’ documents it correctly classifies as ‘non-cyber-related’ TN_s , and the number of ‘non-cyber-related’ documents it incorrectly classifies as ‘cyber-related’ FP_s . We run this process with scores from the three different scoring mechanisms and compare their performance.

FDR Scoring		
Cyber	31 (TP)	80 (FP)
Non-Cyber	5 (FN)	96 (TN)
TF-IDF Scoring		
Cyber	79 (TP)	32 (FP)
Non-Cyber	48 (FN)	53 (TN)
APMIS Scoring		
Cyber	86 (TP)	25 (FP)
Non-Cyber	63 (FN)	38 (TN)

TABLE 9: SCORE PERFORMANCE COMPARISON

Table 9 above shows the performance of term weights for each scoring mechanism in correctly identifying cyber-related discussions. Each confusion matrix above is an $n \times n$ matrix tool used for performance evaluation. The diagonals of the confusion TP_s , and TN_s , represent the total number of documents the lexicon correctly classifies. On the other

hand, the off-diagonals of the confusion matrix, FN_s , and FP_s , represents the total number of documents our generic lexicon incorrectly classifies.

The term scores of the frequency-degree ratio scoring are seen to maximise the ability to correctly identify non-cyber related documents while also minimising the probability of classifying a cyber-related document as non-cyber related. On the other hand, the term scores from the APMIS scoring are seen to maximise the ability of the lexicon to identify all cyber-related documents correctly but minimally identifies all non-cyber related documents.

We determine various performance evaluation metrics from the confusion matrices above. The set of metrics we estimate are the Error rate, Performance accuracy, Precision, Recall and the F1 Score.

	FDR Scores	TF-IDF Scores	APMIS Scores
Precision	85%	61%	58%
Recall	31%	72%	78%
Accuracy	61%	62%	60%
F1 Score	46%	66%	67%
Error Rate	38%	38%	40%

TABLE 10: PERFORMANCE EVALUATION

Table 10 above is a cross evaluation table of the lexicon using each scoring criteria. The TF-IDF scores maximises the lexicon's accuracy (the lexicon's general ability to identify both cyber-related and non-cyber-related discussions); the FDR scores maximise the lexicon's precision (the lexicon's performance in identifying only cyber-related discussions); while the APMIS scores maximise the lexicon's recall (the lexicon's ability to identify all cyber-related discussions). A naïve approach to performance evaluation returns an optimum accuracy for monitoring cyber-related discussions (TF-IDF Scoring). However, there are several issues to consider when selecting an optimal metric for evaluation. For example, given the increased cost of deploying cyber mitigation strategies, an analyst may wish to reduce the lexicon's tolerance for false positives (FDR Scoring), by maximising its ability to identify non-cyber-related discussions correctly. On the other hand, an incident monitoring analyst in a fragile operating environment may choose to maximise the lexicon's ability to identify all cyber-related document with an acceptable tolerance for false positives (APMIS Scoring).

5 CONCLUSION

This paper describes a methodology for creating a lexicon that captures the analytical context of the cyber domain. The lexicon can be applied for monitoring and tracking cyber-related incidents in online social platforms. Additionally, the methodology presented for creating the lexicon can be applied to create custom lexicons for tracking cyber-incidents of interest. Our results are based on a balanced representation of all types of cyber-related discussions from multiple online social platforms. Our method produces a useful cyber lexicon with terms evaluated based on their *TF-IDF*, *FDR* and *APMIS* scores. The *TF-IDF*, *FDR* and *APMIS* scores, together, are shown to optimise the lexicon's overall accuracy in monitoring cyber discussions, low tolerance for false positives and high sensitivity to true positives. In its most straightforward application, this study provides researchers and cyber analysts with an integrated approach for detecting 'cyber-related' discussions in online platforms. However, there are extensive theoretical and practical impacts of this study. Firstly, the theoretical methodology outlined in this study can be applied to develop custom cyber-lexicons based on different evidence sources or social platforms. The methodology can be used in a combination of various social platforms to extract keywords that capture the context of cyber discussions taking place on these platforms. Likewise, the lexicon provided in this study can be applied to some information retrieval tasks on social platforms. The terms in the lexicon can be used as filters for sampling cyber-related documents or discussions. Lastly, the results of this study are useful for cyber analysts to detect malicious cyber activities and detect early warning signs of cyber threats on social platforms. To further this study, cyber analysts are usually interested in identifying keywords on social platforms that characterise the proliferation of various types of cyber events. We are currently working on classifying filtered 'cyber-related' discussions from online platforms based on the classifications of cyber incidents presented by Hathaway (Hathaway et al., 2012).

6 REFERENCES

- Allahyari, M., Pouriyeh, S., Assefi, M., Safaei, S., Trippe, E. D., Gutierrez, J. B., & Kochut, K. (2017). *A Brief Survey of Text Mining: Classification, Clustering and Extraction Techniques*. Retrieved from <http://arxiv.org/abs/1707.02919>
- Baccianella, S., Esuli, A., & Sebastiani, F. (2010). SentiWordNet 3.0: An Enhanced Lexical Resource for Sentiment Analysis and Opinion Mining SentiWordNet. *Analysis*, 10(January 2010), 1–12. <https://doi.org/10.1.1.61.7217>
- Bernstein, M., Monroy-Hernández, A., Harry, D., André, P., Panovich, K., & Vargas, G. (2011). 4chan and /b/: An Analysis of Anonymity and Ephemerality in a Large Online Community. *Proceedings of the Fifth International AAAI Conference on Weblogs and Social Media*, (Coleman), 50–57. <https://doi.org/10.1.1.207.9761>

- Bian, J., Yang, Y., Zhang, H., & Chua, T. S. (2015). Multimedia summarization for social events in microblog stream. *IEEE Transactions on Multimedia*, 17(2), 216–228. <https://doi.org/10.1109/TMM.2014.2384912>
- Blumenstock, J. E. (2008). Size matters: word count as a measure of quality on wikipedia. *Proceedings of the 17th International Conference on World Wide Web*, 1095–1096. <https://doi.org/10.1145/1367497.1367673>
- Bruns, A., & Yuxian Eugene, L. (2012). Tools and methods for capturing Twitter data during natural disasters. *First Monday*, 17(4). <https://doi.org/http://dx.doi.org/10.5210/fm.v17i4.3937>
- Chen, J., Xu, H., & Whinston, A. B. (2009). Moderated Online Communities and Quality of User-Generated Content. *Ssrn*, 1222. <https://doi.org/10.2139/ssrn.1481772>
- Debole, F., & Sebastiani, F. (2003). Supervised term weighting for automated text categorization. *Proceedings of the 2003 ACM Symposium on Applied Computing - SAC '03*, 784. <https://doi.org/10.1145/952686.952688>
- Dionisio, N., Alves, F., Ferreira, P. M., & Bessani, A. (2019). *Cyberthreat Detection from Twitter using Deep Neural Networks*. 1–8. <https://doi.org/10.1109/ijcnn.2019.8852475>
- Halper, F. (2017). Advanced Analytics: Moving Toward AI, Machine Learning, and Natural Language Processing. *TDWI Best Practices Report*.
- Hanley, J., & McNeil, B. (1982). The Meaning and Use of the Area Under The a Receiver Operating Characteristic Curve. *Radiology*, 143(1), 29–36.
- Hathaway, O. A., Crotoof, R., Levitz, P., Nix, H., Nowlan, A., Perdue, W., & Spiegel, J. (2012). The law of cyber-attack. *California Law Review*, 100(4), 817–885. <https://doi.org/10.15779/Z38CR6N>
- Hernández, A., Sanchez, V., Sánchez, G., & Pérez, H. (2016). Security Attack Prediction Based on User Sentiment Analysis of Twitter Data. *Proceedings of the 2016 IEEE International Conference on Industrial Technology (ICIT)*, 610–617. <https://doi.org/10.1109/ICIT.2016.7474819>
- Jurafsky, D., & Martin, J. (2017). Vector Semantics. In *Speech and Language Processing* (2nd ed., pp. 99–124). Prentice Hall.
- Kaji, N., & Kitsuregawa, M. (2007). Building Lexicon for Sentiment Analysis from Massive Collection of HTML Documents. *EMNLP-CoNLL*, 43(June), 1075–1083.
- Kenneth, C., & Hanks, P. (1990). Word Association Norms, Mutual Information, And Lexicography. *Computational Linguistics*, 16(1).
- Khan, F. H., Qamar, U., & Bashir, S. (2016). SentiMI: Introducing point-wise mutual information with SentiWordNet to improve sentiment polarity detection. *Applied Soft Computing Journal*, 39, 140–153. <https://doi.org/10.1016/j.asoc.2015.11.016>
- Khandpur, R. P., Ji, T., Jan, S., Wang, G., Lu, C.-T., & Ramakrishnan, N. (2017). *Crowdsourcing Cybersecurity: Cyber Attack Detection using Social Media*. <https://doi.org/10.1145/3132847.3132866>
- Kipper, K., Korhonen, A., Ryant, N., & Palmer, M. (2006). Extending VerbNet with novel verb classes. *Proceedings of LREC*, 2006(2.2), 1.
- Knuttila, L. (2011). User unknown: 4chan, anonymity and contingency. *First Monday*, 16(10). <https://doi.org/10.5210/fm.v16i10.3665>
- Lippmann, R. P., Campbell, W. M., Weller-Fahy, D. J., Mensch, A. C., Zeno, G. M., & Campbell, J. P. (2016). Finding Malicious Cyber Discussions in Social Media. *Lincoln Laboratory Journal*, 22(1), 203–209.
- Mackey, R. (2010). "Operation Payback" Attacks Target MasterCard and PayPal

- Sites to Avenge WikiLeaks. *The New York Times*. Retrieved from http://thelede.blogs.nytimes.com/2010/12/08/operation-payback-targets-mastercard-and-paypal-sites-to-avenge-wikileaks/?_r=0
- McKenzie, P., Burkell, J., Wong, L., Whippey, C., Trosow, S. E., & McNally, M. B. (2012). User-generated online content 1: Overview, current state and context. *First Monday*, 17(4). <https://doi.org/http://dx.doi.org/10.5210/fm.v17i6.3912>
- Mihalcea, R., & Tarau, P. (2004). TextRank: Bringing Order into Texts. *Proceedings of EMNLP*, 404–411.
- Nielsen, F. Å. (2011). A new ANEW: Evaluation of a word list for sentiment analysis in microblogs. *CEUR Workshop Proceedings*, 718, 93–98.
- Ntoulas, a., Pzerfos, P., & Cho, J. C. J. (2005). Downloading textual hidden web content through keyword queries. *Proceedings of the 5th ACM/IEEE-CS Joint Conference on Digital Libraries (JCDL '05)*, 100–109. <https://doi.org/10.1145/1065385.1065407>
- Olsen, P. (2013). *We Are Anonymous*. London: William Heinemann.
- Olteanu, A., Castillo, C., Diaz, F., & Vieweg, S. (2014). CrisisLex: A Lexicon for Collecting and Filtering Microblogged Communications in Crises. *Proc. of the 8th International Conference on Weblogs and Social Media*, 376. <https://doi.org/10.1.1.452.7691>
- Pras, A., Sperotto, A., Moura, G. C. M., Drago, I., Barbosa, R., Sadre, R., ... Hofstede, R. (2010). Attacks by "Anonymous" WikiLeaks Proponents not Anonymous. *CTIT Technical Report*, (10.41), 1–10. <https://doi.org/10.41>
- Ramos, J. (2003). Using TF-IDF to Determine Word Relevance in Document Queries. *Proceedings of the First Instructional Conference on Machine Learning*, 1–4. <https://doi.org/10.1.1.121.1424>
- Revovna Ospanova, B. (2013). Calculating Information Entropy of Language Texts. *World Applied Sciences Journal*, 22(1), 41–45. <https://doi.org/10.5829/idosi.wasj.2013.22.01.2964>
- Rose, S., Engel, D., Cramer, N., & Cowley, W. (2010). Automatic keyword extraction. *Text Mining: Applications and Theory*, 1–277. <https://doi.org/10.1002/9780470689646.ch1>
- Rowe, M., & Saif, H. (2016). Mining pro-ISIS radicalisation signals from social media users. *Proceedings of the 10th International Conference on Web and Social Media, ICWSM 2016, (Icws)*, 329–338.
- Santorini, B. (1990). Part-of-Speech Tagging Guidelines for the Penn Treebank Project (3rd Revision). *University of Pennsylvania 3rd Revision 2nd Printing*, 53(MS-CIS-90-47), 33. <https://doi.org/10.1017/CBO9781107415324.004>
- SAS Institute. (2018). Natural Language Processing - What it is and why it matters. Retrieved from https://www.sas.com/en_us/insights/analytics/what-is-natural-language-processing-nlp.html
- Schütze, H., & Pedersen, J. (1993). A Vector Model for Syntagmatic and Paradigmatic Relatedness. *Making Sense of Words: Proceedings of the Conference*, 104–113.
- Spärck Jones, K. (1972). A statistical interpretation of term specificity and its application in retrieval. *Journal of Documentation*, 28(1), 11–21.
- Sri, M. S., Yellari, L., & Rao, M. S. (2017). Identifying Malicious Data in Social Media. *International Research Journal of Engineering and Technology (IRJET)*, 4(3). Retrieved from <https://irjet.net/archives/V4/i3/IRJET-V4I3479.pdf>
- Starbird, K., Muzny, G., & Palen, L. (2012). Learning from the crowd: Collaborative filtering techniques for identifying on-the-ground Twitterers during mass disruptions. *Proceedings of 9th International Conference on Information*

- Systems for Crisis Response and Management, ISCRAM, 2011*(April), 1–10.
- Varol, O., Ferrara, E., Menczer, F., & Flammini, A. (2017). Early detection of promoted campaigns on social media. *EPJ Data Science*, 6(1), 563–566. <https://doi.org/10.1140/epjds/s13688-017-0111-y>
- Willett, P. (2006). The Porter stemming algorithm: Then and now. *Program*, 40(3), 219–223. <https://doi.org/10.1108/00330330610681295>
- Zetter, K. (2014). *Countdown to Zero Day*. Retrieved from <http://www.randomhouse.com/book/219931/countdown-to-zero-day-by-kim-zetter>
- Zhang, Y., Jin, R., & Zhou, Z. H. (2010). Understanding bag-of-words model: A statistical framework. *International Journal of Machine Learning and Cybernetics*, 1(1–4), 43–52. <https://doi.org/10.1007/s13042-010-0001-0>

KEY TERMS

Lexicon: A lexicon is a set of vocabulary used to capture the lingual characteristics of a domain of analytical interest.

Social Platforms: Open online communities used for sharing, collaboration and discussion of various topics.

Cyber-related Discussion: A text or sentence written with reference to any topic within the domain of cyber security, its applications, tools and techniques.

BIOGRAPHICAL NOTES

Ruth Ikwu is a research student with the Defence & Cyber Security (DCS) research group in the Department of Computer Science at Brunel University London. She received her PhD in Computer Science and MSc in Information Systems Management from Brunel University London. Her current research areas include cyber security, machine learning, and deep learning.

Panos Louvieris is a Professor of Information Systems and leads the Defence & Cyber Security (DCS) research group in the Department of Computer Science at Brunel University London. He is a committee member of the EPSRC Digital Personhood Network. In addition, he is a member of EPSRC ITaaU+ Network and NEMODE+ Network

REFERENCE

Reference to this paper should be made as follows: Ikwu, R. & Louvieris, P. (2019). Monitoring ‘Cyber-Related’ Discussions in Online Social Platforms. *International Journal on Cyber Situational Awareness*, Vol. 4, No. 1, pp69-98.