

Modeling Islamist Extremist Communications on Social Media using Contextual Dimensions: Religion, Ideology, and Hate

UGUR KURSUNCU, Kno.e.sis Center, Wright State University, USA

MANAS GAUR, Kno.e.sis Center, Wright State University, USA

CARLOS CASTILLO, Universitat Pompeu Fabra, Spain

AMANUEL ALAMBO, Kno.e.sis Center, Wright State University, USA

KRISHNAPRASAD THIRUNARAYAN, Kno.e.sis Center, Wright State University, USA

VALERIE SHALIN, Department of Psychology, Wright State University, USA

DILSHOD ACHILOV, Department of Political Science, University of Massachusetts, Dartmouth, USA

I. BUDAK ARPINAR, Department of Computer Science, The University of Georgia, USA

AMIT SHETH, AI Institute, The University of South Carolina, USA

Terror attacks have been linked in part to online extremist content. Online conversations are cloaked in religious ambiguity, with deceptive intentions, often twisted from mainstream meaning to serve a malevolent ideology. Although tens of thousands of Islamist extremism supporters consume such content, they are a small fraction relative to peaceful Muslims. The efforts to contain the ever-evolving extremism on social media platforms have remained inadequate and mostly ineffective. Divergent extremist and mainstream contexts challenge machine interpretation, with a particular threat to the precision of classification algorithms. Radicalization is a subtle *long-running persuasive* process that occurs over time. Our context-aware computational approach to the analysis of extremist content on Twitter breaks down this persuasion process into building blocks that acknowledge inherent *ambiguity* and *sparsity* that likely challenge both manual and automated classification. Based on prior empirical and qualitative research in social sciences, particularly political science, we model this process using a combination of three *contextual dimensions* – religion, ideology, and hate – each elucidating a degree of radicalization and highlighting independent features to render them computationally accessible. We utilize domain-specific knowledge resources for each of these contextual dimensions such as Qur'an for *religion*, the books of extremist ideologues and preachers for political *ideology* and a social media hate speech corpus for *hate*. The significant sensitivity of the Islamist extremist ideology and its local and global security implications require reliable algorithms for modelling such communications on Twitter. Our study makes three contributions to reliable analysis: (i) Development of a computational approach rooted in the contextual dimensions of religion, ideology, and hate, which reflects strategies employed by online Islamist extremist groups, (ii) An in-depth analysis of relevant tweet datasets with respect to these dimensions to

Authors' addresses: Ugur Kursuncu, Kno.e.sis Center, Wright State University, Dayton, OH, USA, ugur@knoesis.org; Manas Gaur, Kno.e.sis Center, Wright State University, Dayton, OH, USA, manas@knoesis.org; Carlos Castillo, Universitat Pompeu Fabra, Barcelona, Spain, chato@acm.org; Amanuel Alambo, Kno.e.sis Center, Wright State University, Dayton, OH, USA, amanuel@knoesis.org; Krishnaprasad Thirunarayan, Kno.e.sis Center, Wright State University, Dayton, OH, USA, t.k.prasad@wright.edu; Valerie Shalin, Department of Psychology, Wright State University, Dayton, OH, USA, valerie.shalin@wright.edu; Dilshod Achilov, Department of Political Science, University of Massachusetts, Dartmouth, MA, USA, dachilov@umassd.edu; I. Budak Arpinar, Department of Computer Science, The University of Georgia, Athens, GA, USA, budak@uga.edu; Amit Sheth, AI Institute, The University of South Carolina, Columbia, SC, USA, amit@sc.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2019 Association for Computing Machinery.

2573-0142/2019/11-ART151 \$15.00

<https://doi.org/10.1145/3359253>

exclude likely mislabeled users, and (iii) A framework for understanding online radicalization as a process to assist counter-programming. Given the potentially significant social impact, we evaluate the performance of our algorithms to minimize mislabeling, where our context-aware approach outperforms a competitive baseline by 10.2% in precision, thereby enhancing the potential of such tools for use in human review.

CCS Concepts: • **Information systems** → **Information systems applications**; **Information retrieval**; • **Human-centered computing** → **Collaborative and social computing theory, concepts and paradigms**; **Social media**; *HCI theory, concepts and models*; • **Applied computing** → *Psychology*.

Additional Key Words and Phrases: islamist extremism, radicalization, multi-dimensional modeling, contextual dimensions, user modeling

ACM Reference Format:

Ugur Kursuncu, Manas Gaur, Carlos Castillo, Amanuel Alambo, Krishnaprasad Thirunarayan, Valerie Shalin, Dilshod Achilov, I. Budak Arpinar, and Amit Sheth. 2019. Modeling Islamist Extremist Communications on Social Media using Contextual Dimensions: Religion, Ideology, and Hate. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 151 (November 2019), 22 pages. <https://doi.org/10.1145/3359253>

1 INTRODUCTION

In December, 2018 the United Nations Counter-Terrorism Implementation Task Force (CTITF)¹ met² in New York focusing on key action items for a global effort to understand the distortions of the narratives used by terrorists on online platforms. Their subsequent report emphasized that “terrorist organizations like Da’esh (ISIS) and Al Qaida continue to twist religion to serve their ends. The threat posed by returning and relocating fighters, as well as from individuals inspired by them, remains high and has a global reach”. Approximately one thousand Americans between 1980 and 2011 and more than five thousand individuals from Europe through 2015, *traveled* to join extremist groups abroad [17]. Since 2011, the Federal Bureau of Investigation (FBI) reported that 300 Americans attempted or traveled to Syria and Iraq to join extremist groups [41]. Since March, 2014 at least 182 individuals have been charged in the US for ISIS-related offenses³. Recent reports [29] suggest that the terror attacks were linked online extremist content, consumed by supporters as the newly radicalized recruits living in the West are active users of Twitter (e.g., ISIS supporters had higher activity than 67% of all Twitter users) [13]. Recently, a 24 year old college student from Alabama became radicalized on Twitter before moving to Syria to join ISIS⁴. Her radicalization began when she was 20 through meeting other Muslim community members on Twitter, which she refers to as the *Muslim twittersphere*. Self-taught, she read verses from the Qur’an but interpreted them with others in the twittersphere, persuaded that when *the true Islamic State* is declared, it is obligatory to do *hijrah*, which they see as the pilgrimage to ‘*the Islamic State*’. The lack of adequate knowledge about the religion combined with the extremist ideology on Twitter led her to regrettable decisions and actions.

Much has been written on how effectively violent terrorist networks, most notoriously ISIS, have utilized social media to recruit new members [58]. Nevertheless, efforts to capture systematically the ever-evolving dynamics of extremism on social media platforms have remained inadequate:

¹<https://www.un.org/counterterrorism/ctitf/en/about-task-force>

²<https://www.un.org/sg/en/content/sg/speeches/2018-12-06/un-global-counter-terrorism-compact-coordination-committee-remarks>

³GW Extremism Tracker: ISIS in America (May 2019) <https://extremism.gwu.edu/isis-america>

⁴<https://www.nytimes.com/2019/02/22/podcasts/the-daily/isis-american-women.html>

⁵<https://www.nytimes.com/2019/02/19/us/islamic-state-american-women.html>

limited in scope, opaque in approach, and mostly ineffective in practice⁶⁷⁸[6, 25, 34]. Further, as Islamist extremism is a subjective concept with serious repercussions for individuals, the analysis of this topic imposes significant social responsibility in designing reliable algorithms, to avoid discriminatory or biased classification. Merely knowing that someone is a Muslim should not label him or her as a religious extremist. Therefore, domain expertise and responsible use of knowledge provides decisive context.

In this study, we model such religious extremist communications on Twitter based on highly persuasive content, incorporating domain-specific knowledge with three distinct contextual dimensions: religion (**R**), extremist Islamist ideology (**I**) and hate (**H**), originating from domain expert's analysis of the data (see Section 4) as well as social science literature [38, 51, 57] in consultation with our domain expert co-author. The **religion** dimension refers to Muslim attitudes that range from "mainstream" through more "extreme" interpretations of Islamic scriptures. Attitudes toward political **extremist ideology** (i.e., Islamism) are another prevalent dimension of extremism. The conceptualization and measurement of variations in political and ideological attitudes toward Islamism are drawn from [1]'s (concept building) study of Political Islamism. Finally, **hate speech** or attitudinal support for violence is the third critical dimension that provides a benchmark for Islamist extremism with the potential for violent terrorist acts [32, 33]. Our approach will further enable an analysis of the radicalization process individually and collectively in a fine-granular manner, to provide a computational foundation for any form of human intervention. Our hypothesis is that the combination of these three contextual dimensions will create more coherent and distinctive representation of extremist communications, improving the performance of classification. Accordingly, we address the following research questions: **RQ1**: Does incorporation of these contextual dimensions into representation of social media communications improve extremist content classification performance? **RQ2**: Which combination of the dimensions are more effective? **RQ3**: How much does each of these dimensions contribute to the classifier performance?

To achieve these goals, we perform an in-depth analysis of datasets that contain *verified* Islamist extremist accounts (see Section 3.2 for details). We operationalize abstract models of behavior exhibited by an extremist individual under the influence of religion, extremist ideology, and hate. We generate representations for different contextual dimensions using word embeddings drawn from domain-specific resources, to render these dimensions computationally accessible. We then address the challenging problems of inherent sparsity and ambiguity in relations that are implicit in this data to obtain reliable results. Further, language and topical analyses characterize the similarity between users that we can scrutinize, and hierarchical clustering identifies outlier individuals that otherwise would mislead the analysis. Finally, we model Islamist extremist communications utilizing supervised classification algorithms operating over domain-specific representations of extremist and non-extremist users generated by incorporating the three contextual dimensions mentioned above.

In this pursuit, our study makes the following four specific contributions: (i) Development of a computational approach rooted in the dimensions of religion, extremist ideology, and hate that are employed by online Islamist extremist groups to influence, (ii) an in-depth data analysis of relevant datasets with respect to these dimensions, demonstrating improvement in classification with ideological and hate contents, (iii) a framework for understanding online radicalization that serves as a basis for counter programming, (iv) potentially significant reduction in discriminatory

⁶<https://www.lawfareblog.com/marginalizing-violent-extremism-online>

⁷<https://www.theguardian.com/world/2017/sep/19/britain-has-large-audience-for-online-jihadist-propaganda-report-says>

⁸<https://policyexchange.org.uk/wp-content/uploads/2017/09/The-New-Netwar-2.pdf>

bias of mislabeling mainstream (non-extremist) Muslim accounts as extremist⁹. As precision in classifying the communications on Islamist extremism takes precedence over recall to interpret a response properly, we evaluate precision to emphasize minimizing mislabeled non-extremist users. Our approach outperforms a competitive baseline by 10.2% in precision and 8.5% in recall leading to an improvement of 10.7% in F1-score. The combination of all three contextual dimensions in the representation of an account outperforms other alternatives.

In Section 2, we provide details on existing research related to Islamist extremism on social media. In Section 3, we describe the preliminary concepts used in this study. Section 4 characterizes the dataset through language, topical, statistical and similarity analyses, which inform our subsequent modeling approach. We discuss the detailed modeling of Islamist extremism in Section 5, and the evaluation of results and its implications in Section 6. Finally, we present our conclusions with future directions in Section 7.

2 RELATED WORK

The state of the art approaches to detecting and analyzing Islamist extremist communications on social media are limited in their selection of features due to the sparsity and ambiguity inherent in social media data. We need novel approaches to learn coherent representations of content by making use of contextually relevant domain knowledge in a principled manner.

Previous research related to Islamist extremism on social media has focused on four problems: (i) detection of extremist content [8, 35, 47], (ii) prediction of extremist users [7, 26–28, 46, 59], (iii) detection of communities for extremist users [5, 9, 49, 50] and (iv) identification of hate promoting extremism [2–4, 21, 55]. We categorize this study as detection of extremist content and prediction of extremist users.

Ferrara et al. [28] proposed a framework to predict extremist users, their adoption of extremist content, and interaction reciprocity between extremists and regular users. They built predictive models for a binary classification of extremist users using Random Forest (RF) and Logistic Regression (LR) algorithms. To predict extremist users, they employed 52 features that include user and tweet metadata as well as information related to user network and temporal evolution of content. They performed prediction of adoption of extremist viewpoints (as a result of being influenced) based on the behavior of regular users retweeting the content from extremist users. The prediction of interactions with extremists (indicating more active involvement) was based on reply tweets. For all three prediction tasks, RF outperformed LR, with AUCs (Area Under the ROC Curve) of 0.87, 0.77 and 0.69 for the prediction of extremist users, adoption and interactions, respectively.

Rowe et al. [46] performed an analysis of 154K Twitter users in order to extract cues related to radicalization from their content, based on whether the users favor pro vs anti-extremist stances. They found that 727 of these users displayed a pro-ISIS stance, particularly when an event related to ISIS unfolded. In another study, a graph-based semantic approach for detection of radicalization in the Twitter content was proposed by [47]. They utilized knowledge graphs (e.g., DBpedia) to provide semantic relationships between the extracted entities in the content, improving robustness over prior approaches that involved lexical, sentiment, topic and network features. They applied their approach to 1132 (566 pro / 566 anti-ISIS) users, with 1.9M (0.6M pro / 1.3M anti-ISIS) tweets, achieving an F1-score of 0.92.

Fernandez et al. [27] developed an approach for detection and prediction of the influence a user is exposed to, by combining social and computational models of radicalization. They compared the radicalization level of 112 pro-ISIS v/s 112 “general” Twitter users with respect to the roots

⁹Mainstream adherents to Islam number ~1.8 billion, while only a small fraction of them adopt extremist views.

<https://www.pewresearch.org/fact-tank/2017/08/09/muslims-and-islam-key-findings-in-the-u-s-and-around-the-world/>

of radicalization at the individual (micro), community (meso) and global (macro) levels. Their approach achieved up to a 0.90 F1-score for detection and between 0.70 and 0.80 precision for prediction, utilizing vector representations of users designed based on their three level approach. In a follow-up study, Fernandez et al. [26] utilized contextual semantic features of radical content on social media employing ontologies and knowledge bases (DBpedia and Wikidata) to capture categories, topics, entities and entity types. They tested the effectiveness of extracting semantic context from radical conversations on classification of pro-ISIS and non pro-ISIS accounts. They achieved an improvement in precision, recall and F1-score of 0.04, 0.04 and 0.03, respectively.

In contrast to this literature, our work is grounded in the social science literature and incorporates domain-specific resources [30, 31, 52] in the model to better understand and detect extremist content using linguistic approaches. As Islamist extremism is a complex issue that involves different contexts, traditional approaches do not adequately capture important nuances in the language related to the multiple contextual dimensions of the problem. Our approach uncovers these nuances by decomposing social media posts along the three contextual dimensions (see Section 3.4), and provides a fine-grained basis for understanding an individual's progression towards radicalization. Moreover, this understanding will improve interpretability and serve as a crucial basis for designing and building counter extremism narratives for possible de-radicalization efforts.

3 PRELIMINARIES

3.1 Background: Islamist Extremism On Social Media

Extremist actors involved in the dissemination of persuasive content frequently disguise themselves as legitimate representatives of a religion, doctrine or ideology (e.g., extremists posing as true (mainstream) believers in Islam). From the perspective of the persuader, persuasive (propagandist) messages should resemble messages produced by common agents, but be able to perpetuate their hidden agenda by deception and foster misinformation by distorting concepts and relations. This challenges the reliable detection of radicalization content. Such persuasive content involves unconstrained doctrinal concepts and relationships inspired by religion, history and politics.

For example, the concept “jihad” commonly appears in mainstream Islamic as well as extremist communications, albeit with different context-dependent interpretations (see Table 1). The concept of “jihad” can mean (i) self-spiritual struggle, (ii) defensive war to protect lives and property from aggression, or (iii) acts of provoked or unprovoked violence, depending on its context of use [23]. Classification of the first and the second interpretation of “jihad” as extreme would cause the computational model to be gravely incorrect. Further, the degree and progression on a radicalization scale are reflected in the content. For example, users who are recruiters will have a tendency to disseminate information to influence/impress their followers, and initially utilize religious references in their narratives. As they move further in their persuasive radicalization process, they use extremist ideology propaganda by referring to resources of their ideologues. The process culminates with inciting violence by utilizing hate speech and encouraging the followers to act and commit violence. Thus, to glean reliable and comprehensive insights and to assess its intensity, it is critical to use the three contextual dimensions of Religion(R), Ideology(I) and Hate(H) to analyze all communication.

As the interpretation of an individual term depends upon its surrounding lexical context, accurate identification of the relationship between lexical features and Islamist extremism is crucial. For example (see also Table 1), when the term “jihad” co-occurs with “kill” and “attack”, it connotes *hate and violence*. In the presence of “Allah” and “Islam”, the term “jihad” stands for its original meaning denoting the *religious* concept of self-struggle. “Jihad” co-occurring with “imam_anwar_alawlaki”, who is considered [18] a prominent *ideologue* of radical Islamist groups, connotes *exhorting hate and*

No.	Extremist Content Examples	R	I	H
1.	"Here is the fragrance of Paradise , Here is the field of <i>Jihad</i> . Here is the <i>land of Islam</i> , Here is the <i>land of the Caliphate</i> "		✓	
2.	"Reportedly, a number of <i>apostates</i> were <u>killed</u> in the process. Just because they like it I guess. #SpringJihad #CountrysideCleanup"			✓
3.	"and <i>Jihad</i> means to <i>sacrifice</i> YOURSELF in <u>war</u> to save your <u>country</u> (or religion)"		✓	✓
4.	"I asked about the paths to Paradise It was said that there is no path shorter than <i>Jihad</i> "	✓	✓	
5.	" God honored us w/ <i>Jihad</i> <i>Khilafah</i> in this era of <i>Fitnah</i> "		✓	
6.	"By the Lord of Muhammad (blessings and peace be upon him) The <i>nation</i> of <i>Jihad</i> and <i>martyrdom</i> can never be <i>defeated</i> "		✓	
7.	"Anyone who prefers to raise <i>secularism</i> over Islam is a <i>kafir</i> , whether he's from Saudi, Sudan, Somalia, Mexico, Burma, Hawaii, or elsewhere."		✓	
8.	"#MyJihad is to take care of mother, then mother, then father, then other relatives in..."	✓		
9.	"Kindness is a language which the blind can see and the deaf can hear #MyJihad be kind always"	✓		
10.	"May Allah accept those who fast Monday's and Thursday's."		✓	

Table 1. Example tweets from our dataset for extremist/non-extremist social media users, annotated by our co-author domain expert for religion (R), extremist ideology (I) and hate (H) terminology. "Jihad" appears in multiple dimensions. Examples 8 and 9 contain the term "Jihad" in its mainstream meaning, whereas, in Examples 1, 2, 3, 4, 6, it refers to its meaning in the extremist context. Some of the terms are coded based on their relatedness to one of the three dimensions: Religion (Bold-faced), Ideology (Italicized) and Hate (Underlined)

violence. The word "jihad" acquires a different meaning in each of the contexts above; therefore, its representation should be semantically different as well. For this reason, we generate representations of contents and users based on the three contextual dimensions (Religion, Extremist Ideology and Hate) that we identified for the domain of Islamist extremism. The representation of contents is created through word embedding models learned from domain-specific resources, for each contextual dimension. We provide further details of these procedures in the subsequent subsections.

3.2 Dataset

Our ground truth dataset includes 538 extremist users and their 47,376 tweets in the positive class spanning nearly seven years between October 2010 and August 2017. We have used two datasets (i) tweets of Pro-ISIS users¹⁰, and (ii) tweets of users reported by the Lucky Troll Club¹¹ that have been *verified and suspended by Twitter* because of ISIS-related supportive activity¹² [12]. The data and labels were manually curated by annotators who are experts in the Arabic language and verified by Twitter's anti-abuse team. The dataset has also been used in recent studies for modeling radicalization [27, 28]. From this dataset, we selected only English tweets. In our positive samples, the prevalent concepts, topics and terms usually refer to key domain-specific entities such as people (e.g., ideologues, historic person), locations (region, city) and verbs (fight, kill, join). In the rest of this paper, we refer to the positive examples as "extremist users" and to the negative examples as "non-extremist users".

Creation of Negative Class Samples: For development and testing, we use a dataset of 6040 non-extremist mainstream Muslim religious users and 7000 of their tweets created by Chen et al. [22] to constitute the negative class. Note that Islamist extremist content and mainstream/non-extremist content have overlapping vocabulary terms (e.g., jihad) though used in different senses. This word-sense disambiguation challenges the accurate detection of extremist cues in the content, reducing

¹⁰<https://www.kaggle.com/fifthtribe/how-isis-uses-twitter>

¹¹<http://archive.is/V24aS>

¹²<https://www.technologyreview.com/s/603626/data-mining-reveals-the-rise-of-isis-propaganda-on-twitter/>

precision. To evaluate the effectiveness of our approach to disambiguate content, we create the negative class dataset from a dataset of Muslim religious users. We employ Hierarchical Dirichlet Processing (HDP) (non-parametric) clustering [56] on the tweets from users to hierarchically organize users probabilistically based on their content's topical similarity. The application of HDP over the 7K tweets resulted in 600 coherent clusters based on 20 topics with 30 sub-topics each [54]. We treat these clusters as standing for non-extremist users, and randomly select 538 from the 600-user clusters to create our non-extremist user dataset. This approach allows us to deal with data sparsity associated with users without sacrificing their coherence from their normal usage.

3.3 Word Embeddings

Generating embeddings of content provides a numerical vector representation that captures the context of a word/phrase in a corpus. Embedding algorithms including Word2Vec [43], GloVe [44] and FastText [10] have proven to be effective for creating rich representations tuned to a specific domain. Word (or phrase, sentence, document) embedding models generate numerical vector representations of words (or phrases, sentences, documents) that can be used to represent the content [37]. We can use vector operations such as addition, multiplication, and concatenation to aggregate word representations into representations of phrases, sentences, or short documents such as tweets. Domain-specific words can have frequencies, neighboring words, and usages that differ significantly. Hence, it is important to learn embeddings based on domain-specific corpora, e.g., related to Islamist extremism.

3.4 Contextual Dimension Modeling

Highly persuasive religious extremist communications on Twitter frequently employ language and topical cues related to the religion of Islam (R), extremist Islamist ideology (I) and hate (H) for effectiveness (see Section 4). The distribution of prevalent terms (i.e., words, phrases, concepts) in the content of both extremist and non-extremist users reflects different contextual dimensions of the Islamist extremism problem (see Tables 2 and 3). For example, extremist users often make references to concepts and terminologies related to extremist ideology and hate language, while they share content containing relatively fewer Islamic concepts. On the other hand, in the content of non-extremist users, they more often share content related to the religion of Islam unlike extremist users. Moreover, the ambiguity of diagnostic terms (e.g., jihad) also mandates representation of terms in different contexts. Therefore, to better reflect these differences, we create multiple models that will represent the three contextual dimensions for a reliable analysis. We were guided by authoritative sources to ground our hypotheses and look to operationalize the approach that was not performed for social media communications by social scientists. Specifically, [11, 19, 38, 45, 51, 57] show that extremism and subsequent acts of violence are linked to each of these contexts, describing that extremist groups create different interpretations of *religion* that serves their political *extremist ideological* interests, inciting *hate* and finally leading individuals to commit acts of violence. While [57] found, with context theory, significant positive correlations between cognitive complexity and extremist ideology, [38] argued that extremists teach their young followers to hate "the West" forming an *ideological* premise that they should follow *the 'true' Islam* and it should be considered above everything. Further, [51] found that there are connections between web sites operated by *extremist* and *hate* organizations and select episodes of *violence*. In the lights of these findings from the literature and our observations from our dataset, we identify the three contextual dimensions by carefully scrutinizing the Islamist extremist communications on Twitter as well as social science literature in consultation with our domain expert co-author.

We create three word embedding models for the three contextual dimensions, employing *domain specific resources*¹³ as follows: (i) For Religion: The Qur'an English translation¹⁴ by Muhammad Taqi-Din Al-Hilali and Muhammad Muhsin Khan, and Hadith (collection of Prophetic Narrations) resources referenced by ISIS the most¹⁵: Sahih Al-Bukhari¹⁶ and Sahih Muslim¹⁷ well-known authentic Hadith collection, (ii) For Extremist Ideology: Magazines published by ISIS (e.g., Dabiq¹⁸, Rumiyah¹⁹), books and transcribed lectures of extremist ideologues identified by our domain expert co-author (e.g., Anwar Al-Awlaki, Hassan Al-Banna, Said Qutb, Yusuf al-Qaradawi, Abul A'la Maududi), and (iii) For Hate: Hate speech corpus [24]. Figure 1 illustrates the overall flow of creation of contextual dimension models and representations of a user for each contextual dimension.

In this study, we use Word2Vec [42] with skip-grams to generate contextual dimension models. As we aggregate the tweets for each user, we take *union* of unigrams (U), bigrams (B), and trigrams (T), from tweets of a user, and "average" their word embeddings to generate an embedding vector of the user, following the commonly-used method. We formally define our contextual dimension modeling procedure as follows:

Let us represent a set of words/phrases in tweets of a user (u) as G which is generated through the following procedure: $(U \setminus \{B, T\}) \cup (B \setminus \{T\}) \cup T$. Then, we generate representations of a user (u) along the three contextual dimensions of Religion (R), Ideology (I), and Hate (H) as follows:

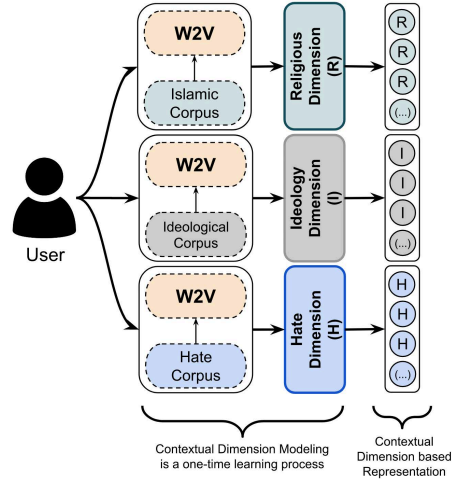


Figure 1. Creation of representations for a user using contextual dimension models using Word2Vec (W2V).

$$\vec{\mathcal{V}}(u^d) = \frac{\sum_{w \in (G \cap W^d)} \vec{v}_w}{|G \cap W^d|}, d \in \Psi \quad (1)$$

where, Ψ denotes contextual dimensions, $\vec{\mathcal{V}}(u^d)$ represents the embedding vector of a user (u) generated from a word embedding model for a contextual dimension d having vocabulary W^d , and w is a word in a tweet of a user u in the vocabulary W^d . The denominator of the Equation 1 is the cardinality of the set of words in the intersection between tweets of a user (G) and vocabulary of a dimension (W^d). We generate the embedding of a user along three dimensions as follows:

$$\vec{\mathcal{V}}(u) = SVD(\mathcal{V}(u^R) \oplus \mathcal{V}(u^I) \oplus \mathcal{V}(u^H)) \quad (2)$$

where (\oplus) concatenates the vector representations of a user (u) along the three contextual dimensions: R : $\mathcal{V}(u^R)$, I : $\mathcal{V}(u^I)$, and H : $\mathcal{V}(u^H)$. Then, we utilize singular value decomposition(SVD) to perform dimensionality reduction [53] which reduces the dimensions to 300, which is standard in the word embeddings literature [15, 42, 44].

¹³These resources are available upon request to reproduce our experiments.

¹⁴<https://www.noblequran.com/translation/>

¹⁵<https://www.kaggle.com/fifthtribe/isis-religious-texts>

¹⁶https://en.wikipedia.org/wiki/Sahih_al-Bukhari

¹⁷https://en.wikipedia.org/wiki/Sahih_Muslim

¹⁸[https://en.wikipedia.org/wiki/Dabiq_\(magazine\)](https://en.wikipedia.org/wiki/Dabiq_(magazine))

¹⁹[https://en.wikipedia.org/wiki/Rumiyah_\(magazine\)](https://en.wikipedia.org/wiki/Rumiyah_(magazine))

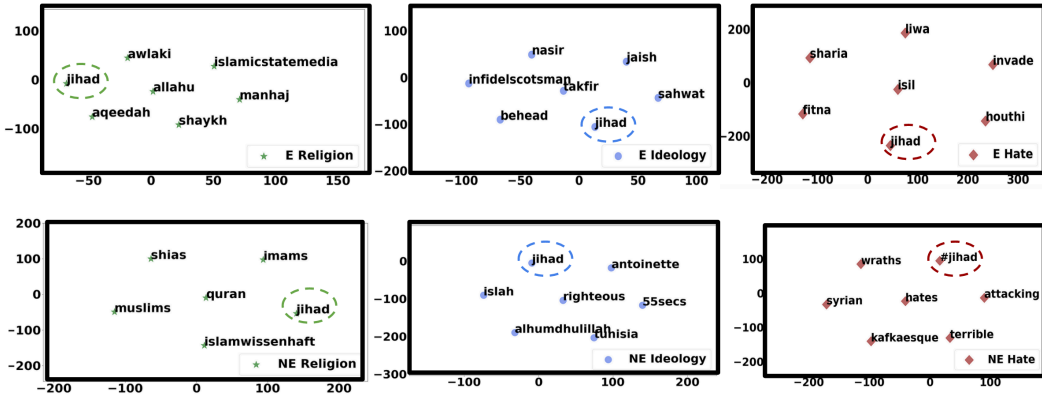


Figure 2. The representations of the word “jihad” with different meanings based on different contextual dimensions. Upper figure: The closest terms to “jihad” in the contexts: religion, ideology and hate for **extremist (E)** users. Lower figure: The closest terms to “jihad” in the three contexts: religion, ideology and hate for **non-extremist (NE)** users.

As we argue that such representations of a user involving the three contextual dimensions will be more coherent and representative, it will disambiguate diagnostic domain-specific words/phrases providing different representations. Thus, the very same terms with different meanings will be represented differently as well. To illustrate how the meaning changes in the extremist and non-extremist content with respect to each contextual dimension, we show the representations of the term “jihad” and its closest terms in Figure 2, for extremist and non-extremist users. In (a), the closest terms to “jihad”, in the content of extremist users, are related to the concepts that extremist groups use to justify their ideology. These closest terms include “infidels”, “scotsman”, “behead”, “nasir”, “takfir” for ideology; “isil”, “fitna”, “houthi”, “invade” for hate; and “aqeedah”, “awlaki”, “shaykh”, “allahu”, for religion. On the other hand, (b) displays the terms closest to “jihad” that are mostly related to the mainstream (non-extremist) Islamic terminologies; such as “alhumdullillah”, “islah”, “righteous” for ideology, “quran”, “muslims”, “imams” for religion, and “terrible”, “attacking”, “hates” for hate. As the meaning of the word “jihad” changes depending on its context, its numerical representation changes as well.

4 EXPLORATORY DATA ANALYSIS

As noted, Islamist extremism on social media has security implications, and requires careful judgment and reliable labeling of content and individuals. Hence, before attempting to use our dataset for modeling Islamist extremism, we examine our dataset carefully, identifying patterns and potential anomalies, checking our assumptions, and testing our hypothesis. We use a multi-pronged approach involving lexical, topical, statistical, and user similarity as discussed below.

4.1 N-Gram Analysis

The language characteristics of extremist and non-extremist contents differ with respect to different contextual dimensions. To determine the language characteristics, we extract n-grams ($n=1$ to 3) from tweets of users in extremist and non-extremist datasets using the skip n-gram model [16, 42]. In our experiments with n-grams, we have observed that 2- and 3-grams were particularly informative because multi-word concepts and entities are prevalent in these communications. For example, “*imam anwar al awlaki*” occurs often because he is a prominent and popular extremist ideologue. Moreover, Islamist extremist groups including “Ahrar al-Sham”, “Jabhat al-Nusra”, “Islamic State”

N-grams	Extremist Users	Non-Extremist Users
Unigrams	<i>isis, syria, <u>kill</u>, iraq, muslim, allah, <u>attack</u>, <u>break</u>, aleppo, assad, islamicstate, army, soldier, cynthias-truth, islam, support, mosul, libya, rebel, <u>destroy</u>, airstrike</i>	person, majd, opening, belief , follower, knowing, khazarrose, al-beltagy's, forza, rally, smyrna, togethernabilahasya, cyrus, islam , okuyamadigimi, dibawain, waalaikumsalam
Bigrams	<i>Caliphate news, islamic state, iraq army, soldier kill, iraqi army, syria isis, syria iraq, assad army, terror group, shia militia, isis attack, aleppo syria, martyrdom operation, ahrar sham, assad regime, follow support, lead coalition, turkey army, isis claim, <u>kill isis</u></i>	sleep controversial, time activist, ali alham-duillah , pape gratefulness, out violence, inii riots, anti-muslim fukushima's, afternoon commit, agree regimes, #patientsafety personality, mahdi muslims, movie muslimap, ahmad worried, biblical festival, jummah soldier, mixe masjid, masmil-waukee reminder, mubarak title, imams koplok
Trigrams	<i>Imam anwar awlaki, video message islamic-state, <u>fight islamic state</u>, <u>isisclaim responsibility attack</u>, <u>muwahideen powerful middleeast</u>, isis tikrit tikritop, amaqagency islamicstate fighter, sinai explosion target, alone state fighter, intelligence reportedly kill, <u>khilafahnew islamic state</u>, yemanqaida commander kill, isis militant hasakah, breakingnew assad army, <u>isis explode middle</u>, hater trier haleemah, trust <u>isis tighten</u>, <u>qamishlus isis fighting</u>, <u>defeat enemy allah</u>, <u>kill terrorist baby</u>, ahrar sham leader</i>	allah bowtie raised, holars killll studios, muhammad lingkkan fdraiser, homefeed wajib akal, israeli paid fajr, eradicating nations project, 2500 muslims homicides, suicide espinoza excess, flow producin shiekh, non-muslims defend reality, masalah taft makan, beneficial right knalan, push serious idea, jahannam philosophy prostration, brotherhood tranquility korean, saturday defile astagfirullah , quick taught america, bbe quran goal , alhamdulillah sat week , touching kids killed, fodation islamic state, islamicate sama-jhten defined

Table 2. Most prevalent unigrams, bigrams, trigrams in the content of extremist and non-extremist users. The n-grams related to the religion of Islam, extremist Islamist ideology and hate appear in bold, italics, and underlined, respectively.

(IS), and locations such as “Deir ez-Zor” once held by ISIS in the Syria and Iraq region are mentioned frequently.

Among the users in the extremist dataset, terms such as “*allah*”, “*fear allah*”, “*jannah*”, “*muslim*”, “*attack*”, “*kill*”, “*isis*”, and “*islamic state*”, are the most frequent, where the first four terms are related to the religion of Islam. The terms *attack* and *kill* are related to hate, and the last two terms refer to the most prominent Islamist extremist group (see Table 2). In contrast, among the users in the non-extremist dataset, the most frequent n-grams are “*amendment yourself*”, “*#truthmonkey*”, “*booth volunteering*”, “*time activist*”, “*drink upon*”, which reflect their contrasting social/political communication.

4.2 Topical Analysis

Topics in the content of extremist and non-extremist users can play a critical role in determining intra-class similarity and inter-class differences among users. We used Latent Dirichlet Allocation (LDA) over n-grams ($n=1-3$) to assess topical similarity of content of *extremist* users to characteristics of Islamist extremism. As LDA is a parametric probabilistic approach for retrieving topics, the number of topics should be carefully selected to capture the themes *optimally*. Hence, we used *perplexity measure* to obtain the optimal number of topics that best represents the content. A higher perplexity score implies higher representativeness and semantic integrity among the topics. Researchers [31] have applied LDA over various combinations of unigrams (U), bigrams (B), and trigrams (T) to obtain informative topics. We apply this procedure creating three different topic models covering the following combinations: (i) U, (ii) U+B, (iii) U+B+T (see Section 3.2). Using the perplexity score depicted in Figure 3, we identified 90 as the threshold for the optimal number of topics for each user for each of the four topic models. We have also used the same number (90) of

topics for non-extremist users. That is, we did not compute perplexity scores for non-extremist users separately as they were created using Hierarchical Dirichlet Process (HDP), which is a non-parametric form of LDA.

Dataset (User types)	Prevalent Topics
Extremist Users	<i>islamic state, syria, isis, <u>kill</u>, allah, video, minute propaganda video scenes, jaish islam release, restock missile, kaffir, join isis, aftermath, mercy, martyrdom operation syrian opposition, punish libya isis, syria assad, islam sunni, swat, lose head, wilayatalfurat, somali, <u>child kill</u>, takfir, jaish fateh, baghdad, iraq, kashmir muslim, capture, damascus, report rebel, british, qala moon, jannat, isis capture, border cross, aleppo, iranian soldier, tikrit tikrittop, <u>lead shia military kill, saleh abdeslam refuse cooperate</u></i>
Non-Extremist Users	<i>masjid job kill, smelled valentines day myanmar, black pada newspaper, quran, tarek necessary like lost, radioactive bande khuda delivered, kaiciid united nations between sky, movement, mustafa human reference, dislodge fatir, kids cruise islamophobia language, active people justice party, hati tiba jihad, abdel lawful farrakhan, adha suhaib, hiasan racist, darinya alhamdulillah, order u.s. iran strike, light headed narcissist stuff, truth monkey, protest jihad controversial, moon accept boycott states, arabi fornicate expiration, al-beltagy rose, khuda jannat, brotherhood maaf, sunni islam, wasidiyah allahumma, muhammad laws onward walking, desperation rather hugo, okurs show rotnhell, american smurf, abraham killed, shifters controversy military, allah, prophet muhammad, rest peace, iraq asks bible jerebu</i>

Table 3. Topics extracted from the content of extremist and non-extremist users using LDA and HDP. The topics related to religion, ideology and hate are bold-faced, italicized and underlined respectively. Remaining topics did not fall under any particular dimension.

In the content of extremist users, topics related to the Islamist ideology are prevalent compared to topics related to hate and religion. For instance, Islamist extremist users frequently make use of ideology-related words/phrases to promote their organization (e.g., “islamic state”, “isis”, “join isis”), its activities (e.g., “martyrdom operation syrian opposition”) or attacks on non-muslim people (e.g., “kaffir”) (see Table 3). On the other hand, topics related to religion are prevalent in the content of non-extremist users. For instance, non-extremist users invoke religious concepts such as “allah”, “prophet muhammad”, “quran”, “alhamdulillah” and “jannat”, whereas there is only one reference to hate, “masjid job kill”. We observe that the prevalence of these topics related to different contextual dimensions in the content of extremist and non-extremist users varies.

4.3 User Similarity

As observed in Sections 4.1 and 4.2, the content of extremist and non-extremist users show strong dissimilarities in the use of language and the topics of conversations based on extremist Islamist ideology and hate, whereas they are relatively similar based on religion. Assessment of similarity between extremist and non-extremist users reveals the contrast between these users with respect




Figure 3. Identification of optimal number of relevant topics based on perplexity scores. From these graphs, we identify 90 as an optimal number of topics that best represent the content of users.

to the three contextual dimensions. Further, assessing similarity between users in each group (extremist/extremist, non-extremist/non-extremist), shows the coherence of the extremist and non-extremist datasets, and further allows us to observe potential anomalies if any exist. Hence, we provide a similarity analysis of users through comparison between pairs of users. We utilize the embedding representations of such users created through three contextual dimension models (R, I, H) (see Section 3.4) and measure the distance between them using cosine similarity. In Figures 4, 5, and 6, the heat maps depict user similarity, where similarity values range from 0.0 to 1.0, represented using shades of red, with white being 0 and dark red being 1; therefore, the darker areas correspond to more similar users.




Figure 4. Similarity between extremist (x-axis) and non-extremist users (y-axis) based on religion, ideology and hate dimensions. User id appears on the x axis for extremist and the y axis for non-extremist users. Extremist and non-extremist users show strong similarity for religion, and weak similarity for hate. They show stronger similarity based on ideology compared to hate, but weaker compared to religion.

Figure 4 shows the similarity between extremist and non-extremist users. The pairs of extremist/non-extremist users (E-N) show strong similarity for religion in Figure 4 (left), while the similarity is relatively weak for hate in Figure 4 (right). Based on ideology, these pairs of users display significant similarity in Figure 4 (middle) compared to hate. On the other hand, it is noteworthy that a small set of pairs for both hate and ideology have stronger similarity compared to other users.




Figure 5. Similarity between extremist users based on a single dimension (religion, ideology and hate). User id appears on the x and y axes for extremist users. Extremist users display strong similarity among themselves based on religion, while a small set of extremist users on the x-axis do now show this similarity with other extremist users on the y-axis. A cluster of users between 250 and 350 on the x-axis for ideology in (middle) and users between 0 and 50 for hate in (right) show strong similarity with other users. On the other hand, a set of extremist users on the y-axis of (left) triangle-figure and x-axis of (right) triangle figure do not show any similarity (white cells) with the majority of other extremist users.

Figure 5 depicts similarity among extremist users for religion, ideology and hate contextual dimensions. In Figure 5 (left), extremist users generally show strong similarity based on religion. On ideology, while a significant number of extremist users are not similar with each other, a

collection of extremist users between 250 and 350 on the x-axis shows stronger similarity with other extremist users between 240 and 538 on the y-axis in Figure 5 (middle). In Figure 5 (right), only a small collection of extremist users between 0 and 50 on the x-axis shows strong similarity with the majority of other extremist users on the y-axis based on hate. We believe that this is due to extremist users employing different hate tactics for different targets. When we consider tweets from users that are “distant” for the hate dimension, their content looks different. For example, considering the example tweets from Table 1, one extremist user’s tweets might be emboldening hatred against “apostates” (i.e., Muslims in other countries who, in the eyes of Islamist extremist groups, have deserted Muslim ideals, ideology and religion), while another might be about hatred against “the West”. The context in each of these conversations would be different because of their target, despite the hatred being incited. It is noteworthy that a spectrum of extremist users between 0 and 100 on the x-axis representing 19% of these extremist users of Figure 5 (left) for religion, and a collection of disparate users on the y-axis of Figure 5 (right) for hate, do not show any similarity (white cells) with a significant number of other extremist users on the y-axis. *This implies that the extremist user dataset might contain outliers or mislabeled users as extremist.*




Figure 6. Similarity between non-extremist users only based on religion(left), ideology(middle) and hate(right) dimensions. User id appears on the x and y axes for non-extremist users. Non-extremist users are strongly similar to each other based on religion as well as ideology, while they do not display similarity based on hate.

Figure 6 depicts similarity among non-extremist users for religion and ideology dimensions. In Figure 6 (left) and (middle), non-extremist users show strong similarity based on religion and ideology respectively. Note that the darker shade of red does represent the similarity of two users, but not relatedness of their content to ideology. Hence, strongly similar non-extremist users based on ideology might still have low relatedness to specific ideological content. In Figure 6 (right), non-extremist users do not show similarity with each other based on hate.

Observations: Through this exploratory analysis, we make the following observations that guide our modeling approach: (i) The content of extremist users heavily contains language, topical and contextual features from Islamist extremist ideology, religion of Islam, and hate speech. (ii) While extremist and non-extremist users are similar in their appeal to religion, they differ in their appeal to Islamist extremist ideology and hate. This might be because they employ different ideological and hate tactics for their targets. (iii) A small subset of extremist users are dissimilar to other extremist users for religion (see Figure 5), implying that the extremist user dataset may contain *likely outlier users*. Hence, we pursue a set of experiments to examine the presence of likely outliers and their identification in the extremist dataset, as described in Section 5.1

5 METHOD FOR MODELING EXTREMIST AND NON-EXTREMIST USERS

In this section, we explain our modeling approach informed by our observations in Section 4. We first examine the existence of likely outliers using a chain of techniques. After we identify and

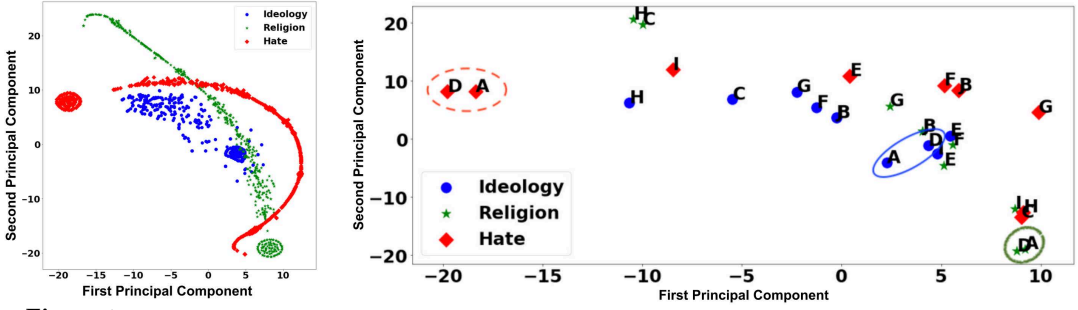


Figure 7. Placement of representations of users in contextual dimensions of religion, ideology and hate in 2-D space using t-SNE (Best seen in color). In (left), which overlays three representations in different coordinates, users show similarity on their spread over the space based on hate and religion, while a small cluster of users fall far from others. (right) provides a closer look over random samples of 20 users on the 2-d space. Users A and D (circled) are close to each other for all contextual dimensions, while they potentially form an outlier cluster of users for hate and religion.

remove likely outlier users in the extremist user dataset, we perform imputation to deal with sparse representations of users for the three contextual dimensions. Finally, we create and evaluate our models employing various combinations of contextual dimensions.

5.1 Identification of Outliers

In Section 4.3, we have seen that the extremist dataset might contain anomalous users (potentially non-extremist), whom we called *likely outliers*. As described in Section 3.2, sole expertise in the Arabic language or the problem of online abusive behavior would not suffice for reliable labeling process, as such a complex problem requires deep domain knowledge and expertise. Therefore, we suspect that the likely outliers that we have observed in the extremist dataset result from the lack of knowledge and expertise in the problem of Islamist extremism. Especially, given the immense sensitivity and related security implications of the problem, we undertake a further in-depth analysis that includes hierarchical clustering and statistical analysis, followed by validation by our co-author domain expert in the field of religious extremism.

To visualize potential dissimilarity and the presence of likely outliers among the extremist users with respect to contextual dimensions, we plot representations of extremist users in a two dimensional space in Figure 7 (left), using T-distributed Stochastic Neighbor Embedding (t-SNE) [39]. It demonstrates such dissimilarity, the spread of users over the space, and existence of likely outlier users for the three contextual dimensions. When we stretch the patterns of the spread of users for hate and religion, it provides strong similarity, while a circling set of users for both contextual dimensions forming the cluster of likely outliers. To ensure that the users in these small clusters are same users for both contextual dimensions, we have picked 10 random users from the extremist user dataset and placed their representations in each contextual dimension on the 2-D space. As shown in Figure 7 (right), users **A and D** fall far from other users forming an outlier cluster for religion and hate contextual dimensions. We have repeated this procedure with different sets of random 10 users multiple times, and found that the users in these small clusters are same, which confirms our observations from Figure 7 (left). Therefore, we must identify these likely outliers before creating our models.

Hierarchical Clustering: To identify the likely outliers in the extremist user dataset, we perform an unsupervised hierarchical density based clustering (HDBC) [20] over the 538 users in the extremist dataset for each contextual dimension. HDBC forms clusters based on the euclidean distance between the users with respect to their representations for each of the contextual dimensions of religion, ideology and hate. HDBC clustering reveals two main clusters, one of which forms the majority of users, the *Likely Extremist* users, while the small cluster of users is the *Likely Outliers*.

Figure 8 shows the distribution of users over the two clusters, where the y axis represents the percentage of users in each cluster for the dimensions of religion, ideology and hate.

Statistical Analysis: From the HDBC clustering, we identified 99 (18%), 48 (9%) and 141 (26%) users in the extremist dataset, clustered as *likely outliers* for *religion, ideology and hate contextual dimensions*, respectively. To confirm a clear separation between these two clusters along with its statistical significance, we performed a non-parametric Mann-Whitney U-test for each contextual dimension. Table 4 shows the U-statistics (U-stats) and their p-values along with the effect size. While all analyses reveal a significant difference, all effects are significant with moderate effect sizes.

The effect size for ideology is slightly higher than that for the hate and the religion, implying that ideology is more effective in clustering. This outcome suggests that the variance between the content of the two clusters of users based on each dimension (especially ideology) is high; hence, the users in these two clusters are significantly different from each other.




Figure 8. Based on the HDBC clustering algorithm, we obtain two main clusters, and we call the majority cluster as “Likely Extremist” and the small cluster as “Likely Outliers”. The y-axis represents the percentage of users for each cluster. 141 hate, 99 religion, 48 ideology.

Dimension	U -stats	z-score	p-value	Effect Size
Religion	5049	12.08	0.0027	0.53
Ideology	9566	13.95	0.001	0.61
Hate	8178	12.4	0.0016	0.54

Table 4. Non-parametric Mann-Whitney U-test between extremist users using their content representations from different dimensions. While outliers and likely extremist users differ on all dimensions, the did not differ quite as much on religion.

Validation: We created a random sample of 76 users comprising 15% of the extremist dataset, to validate the two clusters for identified likely outliers and likely extremists. Our co-author domain expert annotated these users as *likely outliers*, and *likely extremist*. We obtained a kappa score [40] of 0.82, (69 correct and 7 incorrect matches).

Lastly, upon the validation of outliers by our co-author domain expert, we obtained the set of 49 outlier users in the extremist dataset. Content of the outlier users contains the following frequent terms: *marriage, Allah, bonded, silence, Islam leaders, Berjaya hilarious, cake, miss mit, kemaren, Quran, Khuda, prophet, Muhammad, Ahmad*. We found that these outlier users different from other extremist users, and they are most likely non-extremist users. Keeping these outlier users in the extremist user dataset will cause the model to *unfairly* classify non-extremist users as extremist, which will create serious implications in a real world scenario. Hence, we remove this set of outlier users from the dataset to be used in our modeling phase, yielding an extremist dataset with 489 users.

5.2 Imputation for Sparse Representations

Users on social media often use slang terms and informal language rather than the archaic language used in religious and ideological resources. Moreover, some users use hateful language mixed with religious terms, concepts, and topics in their content, while they do not share information related to Islamist extremist ideology. On the other hand, some users mostly share ideological content while they neither share religious content nor use hate speech. This situation creates sparse content of users in any single dimension, translating to sparse embedding vectors (in some cases zero vectors).

We have identified 148 users who had relatively sparse contextual content for at least one of the three dimensions.

Missing values in a dataset is a common problem, and statistical approaches have been developed to approximate such missing components for numerical data. Similarly, natural language processing addresses this problem through, such as out of vocabulary (OOV) [14] words found in domain-specific applications [48].

5.3 Modeling

We develop models employing different combinations (uni-bi-tri-dimensional) of the three contextual dimensions (religion, ideology, hate) to identify the best possible representation of users for classification [36], and determine the effectiveness and contribution of each dimension. We generate vector representations of users (489 extremist users after removal of outliers and 538 non-extremists users) for each dimension, and concatenate them to create uni-dimensional, bi-dimensional and tri-dimensional models. Uni-dimensional models include only the representation for one contextual dimension, bi-dimensional models include two dimensions and tri-dimensional model include all three dimensions. Then, for the bi-dimensional and tri-dimensional models, we perform SVD to reduce the dimensionality of the vector after concatenation, down to 300. Further, we perform our experiments developing models with and without the imputed representations to assess the effectiveness of imputation for sparse representation of users (see Section 5.2). For models without imputation, we eliminate users with sparse representation, which correspond to removing 148 users from our dataset.

Our hypothesis is that a model with three dimensions will create more coherent representations of users leading to improvements in the performance of classification. To test our hypothesis, we create and compare models that include uni-dimensional (R, I, H), bi-dimensional (IH, RI, RH) and tri-dimensional (RIH) models with and without imputation, apart from the baseline model (see below).

Since the existing related work [27, 28] has utilized Random Forest (RF) and Naive Bayes (NB) algorithms, we employ the same algorithms for a fair comparison. As we identified and removed 49 outliers from the extremist dataset (see Section 5.1), we start with 1027 users with imputation and 879 users without imputation, and then create a hold-out dataset of 300 users. We perform training using stratified 6-fold cross-validation. In Section 6, we report the results for our modeling approach, discuss possible implications, and provide comparison with our baseline. We have chosen the state-of-the-art baseline model defined in [27] which is grounded in social science models of radicalization. Note that this is a modeling comparison using our dataset, described in Section 3.2. They use a frequency-based weighting scheme with a NB model for dichotomous classification over two levels (micro and meso). As we were unable to secure their proprietary resources (i.e., lexicon) used by [27], we made our best effort at replicating their approach on our dataset, for a fair comparison.

Algorithm 1 Imputation for Extremist Users (E)

Input: $\Psi, \hat{U}_E, U_E, \Gamma_E$ ▷ Set of dimensions, Users with sparse vectors, All users, Topic model, respectively.

```

1: function IMPUTATION( $\Psi, \hat{U}_E, U_E, \Gamma_E$ )
2:   for  $d$  in  $\Psi$  do
3:     for  $\hat{u}^d$  in  $\hat{U}_E$  do
4:        $\tilde{u}^d \leftarrow \max_{u^d \in U_E^d} \left\{ \frac{\Gamma_E(\hat{u}^d) \cap \Gamma_E(u^d)}{\Gamma_E(\hat{u}^d) \cup \Gamma_E(u^d)} \right\}$ 
5:        $\vec{\mathcal{V}}(\hat{u}^d) \leftarrow \vec{\mathcal{V}}(\tilde{u}^d)$  ▷ Equation 1
6:     end for
7:   end for
8: end function

```

6 RESULTS

In this section, we report the results on performance of the models we created using precision, recall, F1-score and AUC metrics. From Table 5 and Figure 9, the baseline model has a precision of 0.88, recall of 0.82 and F1-score of 0.84 using a feature size of 23K based on the frequency of unigrams [27].

As shown in Table 5, *RF models with imputation* outperform others in precision, recall and F1-score. Uni-dimensional models with imputation achieve a precision of 0.90 using only ideology, recall of 0.86 using only hate and F1-score of 0.87 with ideology and hate each. We observe that the bi-dimensional model with the combination of religion and hate provides better precision, recall and F1-score of 0.91, 0.90 and 0.91, respectively. The combination of ideology dimension individually with the other two dimensions achieves a modest improvement (2.2%) in precision, and a greater improvement (8.5%) in recall. The inclusion of ideology improves recall reducing false negatives, and improving the identification of extremist users.




Figure 9. Precision and Recall of the models using Random Forest (RF) with and without imputation, based on different combinations of contextual dimensions of religion (R), ideology (I) and hate (H). In general, the models with imputation outperform other models without imputation. In (a), the tri-dimensional (RIH) model provides the best performance with 0.97 precision, and in (b) the bi-dimensional (RH) model provides best recall of 0.9, followed by the RIH model with 0.89.

When we combine the three contextual dimensions, our tri-dimensional model achieves best performance with a precision of 0.97, a recall of 0.89 and an F1-score of 0.93, with imputation. Our tri-dimension model improves the precision, recall and F1-score over the baseline model by 9.3%, 7.9%, 10.7%, respectively. On the other hand, the tri-dimensional model (RIH) improves precision and F1-score over the bi-dimensional model (RH) by 6.6% and 2.2% respectively, at the expense of a decrease of 1.1% in recall. Reducing misclassification of non-extremist users creates a minor decrease in recall, which can be considered a trade-off for a significant improvement in precision, with respect to a large set of non-extremist users. This trade-off between the tri-dimensional (RIH) and bi-dimensional (RH) models in a real world application, translates as follows: while more (~1%) extremist users are mislabeled as non-extremist (false negatives), less (~6%) non-extremist users are mislabeled as extremist (false positives). *As misclassification of non-extremist users can have significant implications in a large-scale application where non-extremists vastly outnumber extremists, the higher precision reduces potential social discrimination.*

To better illustrate the diagnostic ability of these models, we plot ROC curves and compute AUC scores where we can examine the performance of the models at different thresholds with respect to true positive rate (TPR) and false positive rate (FPR). Higher TPR with lower FPR indicates better performance. When a model approaches 1.0 of TPR, the corresponding lower FPR at this point compared to other models signifies better performance.

Figure 10 shows ROC curves for RF and NB models with different dimensions. Using RF, our tri-dimensional model along with the bi-dimensional model RH, achieves the best performance with

Dimension	Algorithm	Precision		Recall		F1-Score	
		w/o Imp	w/ Imp	w/o Imp	w/ Imp	w/o Imp	w/ Imp
Baseline	NB	0.88	–	0.82	–	0.84	–
Ideology (I)	RF	0.89	0.90	0.82	0.85	0.85	0.87
Religion (R)	RF	0.79	0.81	0.80	0.82	0.80	0.81
Hate (H)	RF	0.84	0.88	0.85	0.86	0.85	0.87
I	NB	0.80	0.88	0.71	0.75	0.75	0.81
R	NB	0.70	0.79	0.71	0.74	0.75	0.76
H	NB	0.79	0.80	0.81	0.85	0.80	0.83
I+H	RF	0.88	0.90	0.85	0.87	0.86	0.89
I+R	RF	0.84	0.90	0.87	0.89	0.86	0.89
R+H	RF	0.85	0.91	0.87	0.90	0.86	0.91
I+H	NB	0.88	0.88	0.83	0.85	0.85	0.86
I+R	NB	0.86	0.89	0.81	0.87	0.83	0.88
R+H	NB	0.81	0.92	0.80	0.84	0.80	0.88
R+I+H	RF	0.95	0.97	0.86	0.89	0.91	0.93
R+I+H	NB	0.90	0.91	0.82	0.82	0.86	0.87

Table 5. Results of the uni-bi-tri-dimensional models with and without imputation (Imp). The models without imputation were created based on 879 users after the removal of 49 identified outlier users and 148 users with sparse representations as described in section 5.1. The models with imputation were created based on the 1027 users after the removal of the 49 identified outlier users.




Figure 10. ROC curves and AUC scores of uni-bi-tri-dimensional models with RF and NB (Best seen in color). The RIH tri-dimensional and RH bi-dimensional RF models outperform other models with an AUC score of 0.93. The ROC curve for RIH converges to 1.0 at TPR earlier than the RH model, providing better precision. Note: The baseline approach involves only NB and was not tested with RF.

AUC of 0.93, improving upon the baseline by 16.3%. Further, the ROC curve of the tri-dimensional RF model converges to 1.0 of true positive rate (TPR) at the false positive rate (FPR) of 0.65, whereas the tri-dimensional NB model reaches 1.0 of TPR at FPR of 0.82, implying a 10.9% gain in precision.

Moreover, using NB, our models with multiple contextual dimensions, except the uni-dimensional model with religion, are outperforming the baseline in AUC (80%) by up to 12.5%. This shows the effectiveness of our approach with contextual dimensions of Religion, Ideology and Hate, using the same classifier, namely Naive Bayes. However, the bi-dimensional model RH provides the best performance, with AUC of 0.90, followed by the two bi-dimensional models with RI and IH. While the bi-dimensional model outperforms the tri-dimensional model, the ROC curve of the tri-dimensional NB model converges to 1.0 of TPR at FPR of 0.83, whereas the bi-dimensional model reaches the closest point to 1.0 of TPR at FPR of 0.97, implying a 9.8% gain in precision.

Key Insights: (i) Ideology and hate dimensions are often coupled with religious concepts in the content of extremist users. The inclusion of all three contextual dimensions displays the best performance compared to other models in terms of precision. Specifically, this improvement is important because it significantly reduces potential security implications of a possible deployment of this model; thus, reducing the likelihood of an unfair mistreatment towards non-extremist individuals, in a real world application. (ii) Considering that all three dimensions performing well, different extremist users employ diverse strategies to effectively cover broader set of followers at different levels of radicalization. Specifically, given that each of the three contextual dimensions plays different roles in different levels of radicalization, tri-dimensional model captures nuances as well as linguistic and semantic cues better with respect to different density of these contexts throughout the radicalization process. (iii) The contextual dimensions of religion and hate had more power in classification, suggesting that the extremist users were often using religious content along with hate language. This maybe because they use religious concepts, events, places and historic figures, to justify their hatred towards their targets, such as "apostates" or "the West", as in the examples 2 and 3 in Table 1; they could encourage their followers to commit acts of *violence*.

7 CONCLUSION

Using a principled, multi-dimensional approach to the analysis of Islamist extremist content as defined in social science, we excluded likely outlier non-extremist users to develop a robust classifier with improved precision. The success of our method highlights the limitations of more superficial, manual approaches to the identification of extremist users. Furthermore, it suggests a radicalization process over time through a careful contextual metering approach involving religion, Islamist extremist ideology and hate conversations, with fine granularity.

We improved upon the state-of-the-art in automated classification using the three contextual dimensions of Islamist extremism on social media and learned three domain-specific embedding models for interpreting content shared by the users. Overall, our comprehensive approach achieved 10.2%, 8.5% and 10.7% improvement in precision, recall and F1-score, respectively, over a competing baseline. We make the dataset and domain specific corpora for the three dimensions available upon request for research and reproducibility purposes.

Limitations and Future work: The limited number of labeled instances available for training may fail to track the changing nature of concepts and relationships. We plan to address the dynamic nature of this problem in our future work. Specifically, as the past data may not be representative of the future dynamic changes, we will explore the use of domain-specific knowledge of the radicalization process and progression to make our analysis less fragile.

ACKNOWLEDGEMENT

We acknowledge partial support from the National Science Foundation (NSF) award CNS-1513721: "Context-Aware Harassment Detection on Social Media". C. Castillo was partially supported by La Caixa (LCF) project (LCF/PR/PR16/11110009). D. Achilov was supported in part by the University

of Notre Dame (UND) Global Religion Research Initiative (GRRI) grant through Templeton Religion Trust (Grant ID: TRT0118). Any opinions, conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the NSF, LCF, or UND.

We also thank reviewers of the CSCW 2019 for their constructive feedback that greatly improved the presentation of this work.

REFERENCES

- [1] Dilshod Achilov and Sedat Sen. 2017. Got political Islam? Are politically moderate Muslims really different from radicals? *International Political Science Review* (2017).
- [2] Swati Agarwal and Ashish Sureka. 2014. A focused crawler for mining hate and extremism promoting videos on YouTube.. In *Proceedings of the 25th ACM conference on Hypertext and social media*.
- [3] Swati Agarwal and Ashish Sureka. 2015. Using knn and svm based one-class classifier for detecting online radicalization on twitter. In *International Conference on Distributed Computing and Internet Technology*.
- [4] Swati Agarwal and Ashish Sureka. 2016. Spider and the flies: Focused crawling on tumblr to detect hate promoting communities. *arXiv preprint arXiv:1603.09164* (2016).
- [5] Swati Agarwal, Ashish Sureka, and Vikram Goyal. 2015. Open source social media analytics for intelligence and security informatics applications. In *International Conference on Big Data Analytics*.
- [6] Séraphin Alava, Divina Frau-Meigs, and Ghayda Hassan. 2017. *Youth and violent extremism on social media: mapping the research*.
- [7] Tarique Anwar and Muhammad Abulaish. 2015. Ranking radically influential web forum users. *IEEE Transactions on Information Forensics and Security* (2015).
- [8] I Budak Arpinar, Ugur Kursuncu, and Dilshod Achilov. 2016. Social media analytics to identify and counter islamist extremism: Systematic detection, evaluation, and challenging of extremist narratives online. In *2016 International Conference on Collaboration Technologies and Systems (CTS)*. IEEE, 611–612.
- [9] Michael Ashcroft, Ali Fisher, Lisa Kaati, Enghin Omer, and Nico Prucha. 2015. Detecting jihadist messages on twitter. In *2015 European Intelligence and Security Informatics Conference*.
- [10] Ben Athiwaratkun, Andrew Gordon Wilson, and Anima Anandkumar. 2018. Probabilistic fasttext for multi-sense word embeddings. *arXiv preprint arXiv:1806.02901* (2018).
- [11] Imran Awan. 2017. Cyber-extremism: Isis and the power of social media. *Society* 54, 2 (2017), 138–149.
- [12] Adam Badawy and Emilio Ferrara. 2018. The rise of jihadist propaganda on social networks. *Journal of Computational Social Science* (2018).
- [13] Jonathon M Berger and Jonathon Morgan. 2015. The ISIS Twitter Census: Defining and describing the population of ISIS supporters on Twitter. *The Brookings Project on US Relations with the Islamic World* (2015).
- [14] Phil Blunsom, Kyunghyun Cho, Shay Cohen, Edward Grefenstette, Karl Moritz Hermann, Laura Rimell, Jason Weston, and Scott Wen-tau Yih. 2016. Proceedings of the 1st Workshop on Representation Learning for NLP. In *Proceedings of the 1st Workshop on Representation Learning for NLP*.
- [15] Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics* 5 (2017), 135–146.
- [16] Gerlof Bouma. 2009. Normalized (pointwise) mutual information in collocation extraction. *Proceedings of GSCL* (2009), 31–40.
- [17] Bérénice Boutin, Grégory Chauzal, Jessica Dorsey, Marjolein Jegerings, Christophe Paulussen, Johanna Pohl, Alastair Reed, and Sofia Zavagli. 2016. The foreign fighters phenomenon in the European Union. *Profiles, threats & policies. The International Centre for Counter-Terrorism—The Hague*. doi (2016).
- [18] Lorraine Bowman-Grieve and Maura Conway. 2012. Exploring the form and function of dissident Irish Republican online discourses. *Media, War & Conflict* (2012).
- [19] Gary R Bunt. 2003. *Islam in the digital age: E-jihad, online fatwas and cyber Islamic environments*. Pluto Press.
- [20] Ricardo JGB Campello, Davoud Moulavi, and Jörg Sander. 2013. Density-based clustering based on hierarchical density estimates. In *Pacific-Asia conference on knowledge discovery and data mining*.
- [21] Amparo Elizabeth Cano Basave, Yulan He, Kang Liu, and Jun Zhao. 2013. A weakly supervised bayesian model for violence detection in social media. (2013).
- [22] Lu Chen, Ingmar Weber, and Adam Okulicz-Kozaryn. 2014. US religious landscape on Twitter. In *International Conference on Social Informatics*.
- [23] David Cook. 2015. *Understanding jihad*. Univ of California Press.
- [24] Thomas Davidson, Dana Warmusley, Michael Macy, and Ingmar Weber. 2017. Automated Hate Speech Detection and the Problem of Offensive Language. In *Proceedings of the 11th International AAAI Conference on Web and Social Media (ICWSM '17)*. 512–515.

- [25] Seran De Leede, Renate Haupfleisch, Katja Korolkova, and Monika Natter. 2017. *Radicalisation and Violent Extremism-Focus on Women: How Women Become Radicalised, and how to Empower Them to Prevent Radicalisation*.
- [26] Miriam Fernandez and Harith Alani. 2018. Contextual Semantics for Radicalisation Detection on Twitter. (2018).
- [27] Miriam Fernandez, Moizzah Asif, and Harith Alani. 2018. Understanding the roots of radicalisation on twitter. In *Proceedings of the 10th ACM Conference on Web Science*.
- [28] Emilio Ferrara, Wen-Qiang Wang, Onur Varol, Alessandro Flammini, and Aram Galstyan. 2016. Predicting online extremism, content adopters, and interaction reciprocity. In *International conference on social informatics*.
- [29] Martyn Frampton, Ali Fisher, Nico Prucha, and David H Petraeus. 2017. *The new Netwar: Countering extremism online*. Policy Exchange.
- [30] Manas Gaur, Amanuel Alambo, Joy Prakash Sain, Ugur Kursuncu, Krishnaprasad Thirunarayan, Ramakanth Kavuluru, Amit Sheth, Randy Welton, and Jyotishman Pathak. 2019. Knowledge-aware assessment of severity of suicide risk for early intervention. In *The World Wide Web Conference*. ACM, 514–525.
- [31] Manas Gaur, Ugur Kursuncu, Amanuel Alambo, Amit Sheth, Raminta Daniulaityte, Krishnaprasad Thirunarayan, and Jyotishman Pathak. 2018. Let Me Tell You About Your Mental Health!: Contextualized Classification of Reddit Posts to DSM-5 for Web-based Intervention. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*. ACM, 753–762.
- [32] Mohammed Hafez and Creighton Mullins. 2015. The radicalization puzzle: A theoretical synthesis of empirical approaches to homegrown extremism. *Studies in Conflict & Terrorism* (2015).
- [33] Scott Helfstein. 2012. *Edges of radicalization: Ideas, individuals and networks in violent extremism*. Technical Report. MILITARY ACADEMY WEST POINT NY COMBATING TERRORISM CENTER.
- [34] Ghaffar Hussain and Erin Marie Saltman. 2014. *Jihad trending: A comprehensive analysis of online extremism and how to counter it*.
- [35] Lisa Kaati, Enghin Omer, Nico Prucha, and Amendra Shrestha. 2015. Detecting multipliers of jihadism on twitter. In *2015 IEEE International Conference on Data Mining Workshop (ICDMW)*.
- [36] Ugur Kursuncu, Manas Gaur, Usha Lokala, Anurag Illendula, Krishnaprasad Thirunarayan, Raminta Daniulaityte, Amit Sheth, and I Budak Arpinar. 2018. What's ur Type? Contextualized Classification of User Types in Marijuana-Related Communications Using Compositional Multiview Embedding. In *2018 IEEE/WIC/ACM International Conference on Web Intelligence (WI)*. IEEE, 474–479.
- [37] Ugur Kursuncu, Manas Gaur, Usha Lokala, Krishnaprasad Thirunarayan, Amit Sheth, and I Budak Arpinar. 2019. Predictive Analysis on Twitter: Techniques and Applications. In *Emerging Research Challenges and Opportunities in Computational Social Network Analysis and Mining*. Springer, 67–104.
- [38] Wagdy Loza. 2007. The psychology of extremism and terrorism: A Middle-Eastern perspective. *Aggression and Violent Behavior* (2007).
- [39] Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. *Journal of machine learning research* (2008).
- [40] Mary L McHugh. 2012. Interrater reliability: the kappa statistic. *Biochemia medica: Biochemia medica* (2012).
- [41] Alexander Meleagrou-Hitchens, Seamus Hughes, and Bennett Clifford. 2018. *The travelers: American jihadists in Syria and Iraq*.
- [42] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781* (2013).
- [43] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*.
- [44] Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*. 1532–1543.
- [45] Jialun Qin, Yilu Zhou, Edna Reid, Guanpi Lai, and Hsinchun Chen. 2007. Analyzing terror campaigns on the internet: Technical sophistication, content richness, and Web interactivity. *International Journal of Human-Computer Studies* 65, 1 (2007), 71–84.
- [46] Matthew Rowe and Hassan Saif. 2016. Mining pro-ISIS radicalisation signals from social media users. In *Tenth International AAAI Conference on Web and Social Media*.
- [47] Hassan Saif, Thomas Dickinson, Leon Kastler, Miriam Fernandez, and Harith Alani. 2017. A semantic graph-based approach for radicalisation detection on social media. In *European semantic web conference*.
- [48] Prathusha K Sarma, Yingyu Liang, and William A Sethares. 2018. Domain Adapted Word Embeddings for Improved Sentiment Classification. *arXiv preprint arXiv:1805.04576* (2018).
- [49] Jacob R Scanlon and Matthew S Gerber. 2014. Automatic detection of cyber-recruitment by violent extremists. *Security Informatics* (2014).
- [50] Jacob R Scanlon and Matthew S Gerber. 2015. Forecasting violent extremist cyber recruitment. *IEEE Transactions on Information Forensics and Security* (2015).

- [51] Joseph A Schafer. 2002. Spinning the web of hate: Web-based hate propagation by extremist organizations. *Journal of Criminal Justice and Popular Culture* (2002).
- [52] Amit Sheth, Sujan Perera, Sanjaya Wijeratne, and Krishnaprasad Thirunarayan. 2017. Knowledge will propel machine understanding of content: extrapolating from current examples. In *Proceedings of the International Conference on Web Intelligence*. ACM, 1–9.
- [53] Jamin Shin, Andrea Madotto, and Pascale Fung. 2018. Interpreting Word Embeddings with Eigenvector Analysis. (2018).
- [54] PK Srijith, Mark Hepple, Kalina Bontcheva, and Daniel Preotiuc-Pietro. 2017. Sub-story detection in Twitter with hierarchical Dirichlet processes. *Information Processing & Management* (2017).
- [55] Ashish Sureka and Swati Agarwal. 2014. Learning to classify hate and extremism promoting tweets. In *2014 IEEE Joint Intelligence and Security Informatics Conference*.
- [56] Yee W Teh, Michael I Jordan, Matthew J Beal, and David M Blei. 2005. Sharing clusters among related groups: Hierarchical Dirichlet processes. In *Advances in neural information processing systems*.
- [57] Alain Van Hiel and Ivan Mervielde. 2003. The measurement of cognitive complexity and its relationship with political extremism. *Political Psychology* (2003).
- [58] Lorenzo Vidino and Seamus Hughes. 2015. *ISIS in America: From retweets to Raqqa*.
- [59] Pooja Wadhwa and MPS Bhatia. 2013. Tracking on-line radicalization using investigative data mining. In *2013 National Conference on Communications (NCC)*.

Received April 2019; revised June 2019; accepted August 2019