Studying Jihadists on Social Media: A Critique of Data Collection Methodologies

by Deven Parekh, Amarnath Amarasingam, Lorne Dawson, Derek Ruths

Abstract

In this article, we propose a general model of data collection from social media, in the context of terrorism research, focusing on recent studies of jihadists. By analyzing Twitter data collection methods in the existing research, we show that the methods used are prone to sampling biases, and that the sampled datasets are not sufficiently filtered or validated to ensure reliability of conclusions derived from them. Alternatively, we propose some best practices for the collection of data in future research on jihadist using social media (as well as other kinds of terrorist groups). Given the similarity of the methodological challenges posed by research on almost all social media platforms, in the context of terrorism studies, the critique and recommendations offered remain relevant despite the recent shift of most jihadists from Twitter to Telegram and other forms of social media.

Keywords: jihadist, terrorism, data collection, graph sampling, network sampling, dataset, Twitter, social media

Introduction

In recent years, jihadist terrorist movements have used varying social media platforms to organize, coordinate operations, and spread propaganda. Significant research has focused on understanding these online activities. Such research naturally requires the collection and analysis of social media data produced primarily by jihadist users. The findings of such studies are only as valid as the data they are based on – which is the topic of the present study. In a comprehensive survey of the largest body of research on online jihadist activity – that of jihadist activity on Twitter – we find a wide array of data collection methods. In the majority of cases, we find that these studies fail to acknowledge limitations of the data collection methods, raising serious concerns about the validity of their findings. Similar issues exist in studies that consider other social media platforms. There are known standards of practice and established methodologies for addressing these issues in the field of computer science that it would be useful for scholars of terrorism to become familiar with and apply to future work.

In order to ground our discussion, we first present a generalized framework for data collection, which can be used to understand the methods used by any study of terrorist behavior on a social media platform. We then show how decisions within this framework yield specific kinds of systemic biases in downstream analysis. In this article we extensively consider the case of Twitter – primarily, because it is the platform on which most studies have been run. However, our framework and conclusions are also useful for researchers working and doing research on other platforms, such as Facebook, Gab.ai, or Telegram. The latter, particularly, has been extensively used by a variety of jihadist groups around the world since at least 2015.[1]

In this article, we make three core contributions. First, we identify common methods of data collection found in the existing literature that studies jihadists on Twitter. As part of this exercise, we propose a general framework of data collection that consists of four phases - initialization, expansion, filtering, and validation – which impact the properties (and quality) of data produced. Our second contribution uses this four-phase formalization to analyze limitations of the data collection methods of existing terrorism research and their implications for their results/findings. Finally, based on these analyses, we recommend best practices to improve the quality of sampled social media data, and accordingly derived results, in terrorism research.

We believe that it would be helpful to clarify some of the terminology used throughout this article. First, the research we examine in this article focuses on different groups involved in terrorism, such as ISIS fighters in Syria, or ISIS foreign fighters from Europe – in this article we use general terms such as "jihadists" to include all groups or entities involved in extremist or radical activities and propaganda related to Islamism on Twitter. "Irrelevant accounts" are non-jihadist accounts, such as news reporters, researchers and ordinary Twitter users that may follow or be related to jihadist accounts on Twitter, but are not jihadists and are not directly involved in jihadist activities.

Social Media Data and Social Graph

Before delving into the substance of our study, here we provide a brief overview of the structure of social media data – with particular attention to the *social graph*.

Data on a social media platform typically consists of user profile information, each user's published content (such as text, images and other media), as well as details of relationships or interactions between the users on the platform. In the context of Twitter, the platform on which we will focus in this study, most users have public profiles (meaning anyone on the internet can see them). A single piece of published content is called a *tweet*, which can contain text, images, links, and mentions of other users. Finally, relationships are formally declared in the form of *follower-followee pairs* (discussed further below).

In many research studies, particularly in terrorism research, it is important to take into account interactions among social media users. Such interactions are typically modelled using a *social graph.* A social graph, also called a social network, is a mathematical structure consisting of a set of nodes representing social media users and a set of edges representing relationships or interactions among those users. In the context of Twitter, the set of nodes represent Twitter user accounts. The question, of course, is what "counts" as a relationship. One of the common relationships used to build a Twitter social graph - and the primary technique used by the online jihadist literature we consider here - is that of one user *following* another user's account. Figure 1(a) shows a representation of such a graph, where circles are "nodes" representing Twitter users and arrows are "edges" such that an arrow from node A to node C shows that user C is following user A on Twitter. *Neighbors* of a node A are all the nodes to which there is an edge from the node A. In case of the Twitter example,

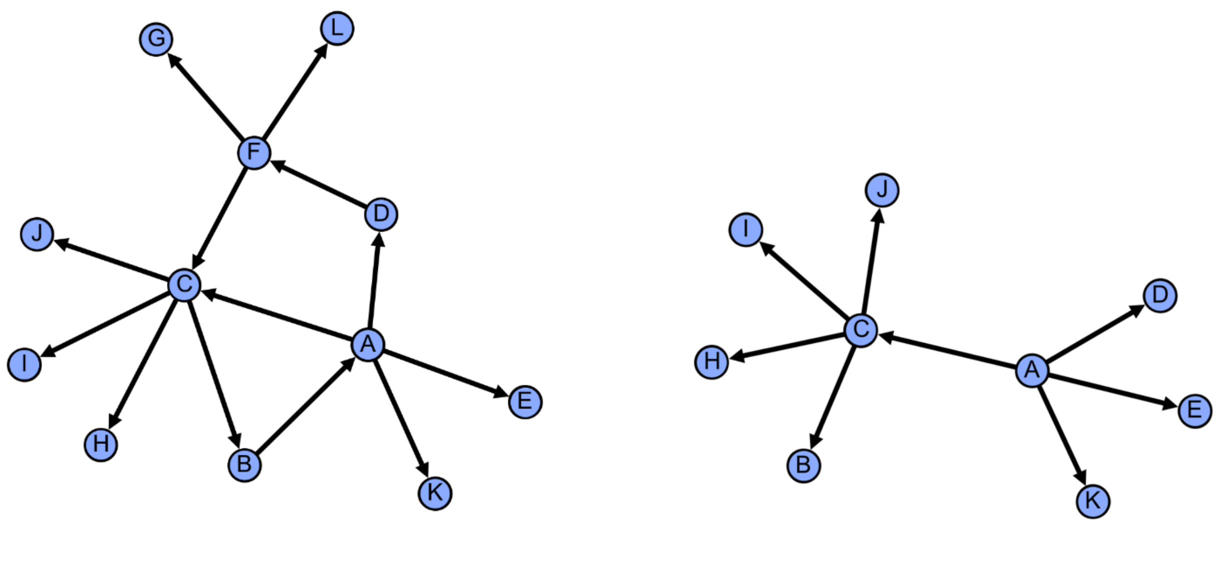

(a) (b)

Figure 1. (a) *An example of Twitter social graph created using the Follower relationship. Each node represents a Twitter user. An edge or arrow from node A to C, for example, shows that user C is following user A on Twitter.* ***(b)*** *A graph sampled from that in (a) using snowball sampling with starter nodes A and C.*

neighbors of the node A would be all the users following the user A on Twitter.

Twitter has three primary kinds of relationships that are used in terrorism research to build a social graph representing jihadist accounts and their interactions.

1. *Follower*. On Twitter a user A is able to *follow* another user B so that the user A can track latest updates and tweets by the user B. Users that follow a user B are called *followers* of the user B. Followers of a jihadist Twitter account may typically include other jihadist accounts, supporters of jihadist groups, potentially radicalizing individuals, as well as irrelevant accounts such as researchers and news reporters.

2. *Friend*. This is an inverse relationship to that of *Follower*. In other words, if user A is a follower of user B on Twitter, then user B is said to be a *Friend* of the user A. Friends of a user A are all the users that are being followed the user A. Friends of a jihadist account may include other prominent jihadist accounts as well as other irrelevant account if the jihadist accounts is attempting to appear as a normal account or posing as a researcher.

3. *Retweet and Mention*. On Twitter, a user can *retweet* (i.e., share) another user's tweet or can *mention* another user directly in their tweet. Retweets can be considered as passive engagement in jihadist activities: a user may endorse the content of a tweet by retweeting it, without explicitly communicating with other users. On the other hand, mentions can be considered as active communication between users. These relations can be used to build social graphs that highlight the flow of information among jihadist user accounts.

The overarching kinds of relationships on Twitter can also be observed on other social media platforms, albeit in a different form. For example, on Facebook a user sharing other user's post can be considered as similar to retweeting on Twitter. This allows the generalization of our framework and subsequent analysis on Twitter to similar social media platforms.

Sampling from Social Graph

To our knowledge, all data collection methods for detecting jihadist accounts on Twitter involve explicitly sampling from the Twitter social graph. When studying a particular group of users, it's natural to take only the portion of the Twitter social graph (the alternative is collecting many millions of users who have absolutely nothing to do with the study). Selecting only a subset of users and/or edges from a graph (also referred to as a network) is called *graph sampling (or network sampling).* Two commonly used sampling methods are:

- *Random Node Sampling*. This is the most basic sampling method where a random subset of nodes is selected from the node set of original graphs. Each node is typically selected independently with a uniform probability. Once the nodes are selected, the sampled graph is constructed by selecting all the edges from the original graph that connect the sampled nodes to each other.

- *Snowball Sampling*. In snowball sampling, we sample a set of *starter nodes*, either manually or randomly from the original graph. For each starter node, we then add all (or a fraction) of its neighbors to the set of sampled nodes. After the nodes are sampled, the sampled graph is constructed by adding the edges connecting the nodes. Figure 1(b) shows a graph sampled from Figure 1(a) using snowball sampling with starter nodes A and C.

A sampled graph is expected to consist of a subset of nodes and edges that are representative of the structure of the original graph. This allows the observations and results obtained from the sampled graph to generalize to the original graph. However, in practice, a crucial and unavoidable issue with sampling (shown in Figure 1(b)) is that any network sample provides a distorted view of the network. From a technical perspective, this distortion can bias various network statistics (e.g., the most highly connected nodes, the number of triangles in the network, the distance between nodes in the network) we might be interested in. It can also bias the metadata associated with the network (e.g. the topics being posted by users). As has been highlighted in the network science literature, sampling must be done with great care.[2]

The present study can be understood as a critical assessment of the impact of popular network sampling practices within the jihadist research community on the validity of research findings.

Overview of Existing Terrorism Research

We identify two distinct research objectives that, together, characterize the vast majority of existing terrorism research studies analyzing Twitter accounts: (1) qualitative and quantitative descriptions or summaries including social network analyses (SNA), and (2) characterizations of jihadist accounts. Crucially, both of these research objectives involve the collection of Twitter data. For this study, we have selected research articles from both of these categories such that they also employ different types of data collection methods. By doing this, we have assembled a representative sample of collection methods for further analysis in this article. In the rest of this section, we present a brief overview of the research in the selected articles.

Qualitative descriptions include expert analyses and commentaries on a particular jihadist or a group of jihadists, as well as the state of jihadist groups and content on Twitter. Quantitative summaries are overviews, including statistics and graphical plots, which explain the presence, activities, and interactions of the jihadist population on Twitter. Social network analyses involve the use of computational methods and tools to study interaction networks among jihadists, with typical goals such as discovering flows of information in a network or finding the most important or central accounts in a network. The following papers were chosen for this category.

- Klausen, 2015, studied the network of 59 Twitter accounts of Western-origin fighters known to be in Syria, over the period of January to March 2014, to understand the information flow and the extent to which access to, and content of, communications are controlled. The key findings point to the controlling role played by feeder accounts belonging to terrorist groups in the insurgency zones and by Europe-based organizational accounts associated with the banned British organization Al Muhajiroun.[3]
- Berger et al., 2015, created a demographic snapshot of ISIS supporters on Twitter based on data collected from September to December 2014. They proposed a methodology for discovering and characterizing relevant ISIS accounts. Furthermore, they studied ISIS supporting accounts that were suspended by Twitter and the effects of suspension in limiting the reach and scope of ISIS activities.[4]
- Berger et al., 2016, collected and analyzed the list of English-speaking ISIS supporter accounts maintained by a Twitter user "Baqiya Shoutout", that were active from June to October 2015. Using social network analysis, they found the ISIS English-language social networks are small and insular. The declining number of accounts and limited amount of pro-ISIS content in the networks suggested that suspension of jihadist Twitter accounts was having a devastating effect. In particular, individual users who repeatedly created new accounts after suspension faced a decrease in their follower counts.[5]
- Bodine-Baron et al., 2016, differentiated ISIS supporters and opponents based on whether they refer to ISIS by its full name in Arabic (The Islamic State) or by the acronym "Daesh". Lexical analysis suggested that the frequent users of Daesh had content that was highly critical of ISIS, while users of "The Islamic State" used glorifying terms. Furthermore, ISIS opponents outnumbered supporters six to one while ISIS supporters routinely out-tweeted ISIS opponents. Social network analysis of ISIS conversations on Twitter revealed identities and prominent content themes categorized as four metacommunities: Shia Muslims, Syrian mujahideen, ISIS supporters, and Sunni Muslims. The metacommunities were further studied to find central communities and their interactions with each other.[6]
- Conway et al., 2017, presented a detailed analysis of disruption (suspension/content takedown) and its effects on pro-IS Twitter accounts, in comparison to other jihadist groups including Hay'at Tahrir al-Sham (HTS), Ahrar al-Sham, the Taliban and al-Shabaa. They observed that pro-IS accounts faced significantly greater disruption than other jihadist accounts, resulting in sparser pro-IS relationship

networks. In addition, they analyzed the presence of IS propaganda on other platforms including content hosting websites by obtaining links to the websites from the jihadist tweets. Significant takedown rate was also observed on those platforms.[7]

Characterization of jihadist accounts refers to the important task of identifying the characteristic features of jihadists Twitter accounts and subsequently using the features to understand behavior of jihadist users, and to discover new potential jihadist accounts. The following are the articles chosen for this category:

- Magdy et al., 2015, classified Twitter accounts as pro- or anti-ISIS by finding whether Arabic tweets posted by the account contained the full name such as "Aldawla Alislamiya" ("Islamic State") or the acronym "da'esh." They observed a correlation between tweeting trends of such accounts that supported or opposed ISIS with major news or events around the time data was collected (Oct.-Dec. 2014). Based on tweet content, such as hashtags and temporal patterns, they found that ISIS supporters joined Twitter to show their support which, in particular, was motivated by frustration with the failure of the Arab spring revolutions. Furthermore, they built a classifier to predict if Twitter accounts support or oppose ISIS using their tweets from the pre-ISIS period.[8]

- Kaati et al., 2015, trained a machine learning model to detect Twitter accounts that support jihadist groups and disseminate propaganda content. To train the model, they used data dependent features such as most common hashtags, word bigrams and most frequent words, as well as data independent features such as frequency of word length, letters, digits, and emotion words. They showed that their model has significant accuracy for English tweets while for Arabic tweets the performance was worse.[9]

- Klausen et al., 2016, developed a behavioral model of extremist user accounts on Twitter to predict if the accounts would be suspended for extremist activity. They trained the model using Twitter data collected during the year of 2015. Using the model they could identify new extremist accounts as well as link new accounts created by the same user. Based on information about suspended users' accounts, they also propose a network search model to efficiently find new Twitter accounts created by suspended users.[10]

- Rowe et al., 2016, studied radicalization signals exhibited by European-based Twitter users by characterizing their differences in their behavior before and after they began using pro-ISIS terms and sharing pro-ISIS content. They proposed methods to identify if a user is activated (i.e., exhibiting radicalized behavior), based on content sharing patterns and pro- and anti-ISIS language used by the user. Furthermore, they studied how the behavior of users diverged before and after activation, in terms of language use, content sharing, and interactions with other users. Finally, they show that social homophily in Twitter communities has a strong influence on adoption of pro-ISIS behavior by radicalizing users.[11]

- Wright et al., 2016, proposed quantitative methods to identify resurgent jihadist accounts, which are new accounts created by the original users of accounts that have been suspended. They found that resurgent accounts grow faster (gather more followers) than naturally growing non-resurgent accounts, and that there are significant proportions (20% - 30%) of fast-growing duplicate accounts. Accordingly, they suggest terrorism researchers need to recognize and account for the biases introduced to their datasets by the large number of resurgent accounts.[12]

- Smedt et al., 2018, created a Hate corpus consisting of online jihadist hate speech from tweets posted by manually identified subversive profiles on Twitter. They also created a Safe corpus which consisted of reporters, imams and Muslims, as well as random tweets on general topics such as cooking and sports. Using Natural Language Processing, they performed quantitative analyses such as language and demographics distribution, as well as keyword analysis comparing Hate and Safe corpora. Finally, using Machine Learning techniques, they could predict jihadist hate speech with over 80% accuracy.[13]

Four-Phase Model of Data Collection from Social Graph

The overarching thesis of this study is that the way in which social media data is collected impacts on the quality of the data obtained and, by extension, the quality and validity of the insights gained from analysis of that data. Therefore, as a starting point, in this section, we describe the methods of data collection employed in the existing research on jihadism on Twitter. In order to frame this discussion, we first provide a four-phase model of data collection, which can be generalized to any social media platform similar to Twitter. Using this model, we categorize methods in the existing literature according to specific strategies they employ to implement the phases of data collection.

Phases of Data Collection

Any method for collecting data from a social graph (whether Twitter or other social media platforms) involves four main phases: *Initialization, Expansion, Filtering* and *Validation*. The first two phases consist of dataset *creation* methods, while the last two phases include methods for improvement and verification of dataset *quality*.

Initialization. This phase involves choosing or obtaining an initial set of Twitter accounts, also called "seed accounts". Seed accounts can be obtained manually by experts in terrorism research from various sources such as news. Another common way of creating a set of seed accounts is by identifying accounts that have made posts using specific keywords. For example, Bodine-Baron et al. searched for tweets using grammatical variations of "Islamic State" and "Daesh" in Arabic, and obtained a list of users who posted the tweets to form an initial seed set.[14] In Twitter and other similar post-oriented platforms, this keyword searching is done first on tweets – identifying tweets that contain the target words. The initial set of users is then obtained by identifying the authors of all these tweets.

Expansion. In this phase, the dataset is grown to include more accounts that are related to the initial seed accounts with the aim of capturing a bigger group or network of jihadists, and one that has potentially more jihadist accounts than irrelevant accounts. Without exception, related accounts are discovered by exploring a social graph – by which we mean the network of explicit relationships among social media accounts. Graph sampling, as discussed earlier, plays a crucial role in exploring the relationships. Since a social graph corresponding to a social media platform is typically huge, graph sampling allows the researcher to study a smaller part of the graph. The main idea of the expansion phase is, therefore, to form a representative dataset by exploring a social graph.

A dataset resulting from the expansion phase necessarily depends on the choice of initial seed accounts. Therefore, it is important to choose the initial set carefully, depending on research objectives. It is noteworthy that the expansion phase need not necessarily use relations, such as being a follower or a friend, typically found in social graphs. There could be different ways of relating any two users on a social media platform: a user replying to another user's tweet, retweeting another user's tweet or sharing similar content in terms of hashtags. Depending on the research objectives, some of these relationship may be more appropriate than others. Follower and Friend relationships are more commonly used in Twitter terrorism research as they suggest direct relations between jihadist accounts.

Filtering. Both the initialization and expansion phases may be followed by a *filtering* phase to improve quality of a dataset. In this phase, accounts from the sample selected using criteria that favor inclusion of jihadist accounts and exclusion of irrelevant accounts. Such criteria include removing inactive user accounts (ones that have not posted any tweets for a reasonably long time period), old accounts that were created a very long time ago, and accounts that have more than a certain number of followers or friends. The criteria for filtering accounts are often informed by domain expertise and intuitive knowledge. For example, Berger et al., 2016, removed all the accounts from their dataset with more than 9,500 followers because such accounts are unlikely to be jihadist accounts.[15]

Validation. In this phase, the final dataset is verified for its quality or reliability. If the research objective is to study the state of jihadist activities on social media, it is expected that a dataset should have a high proportion of jihadist accounts compared to that of irrelevant accounts (or at least that this ratio and bias be well characterized). If the dataset is small, it can be manually assessed. For large datasets, it is standard practice for one or more random samples to be manually assessed.[16] Below, we propose a manual annotation method that we used to validate our dataset.

Figure 2. shows the dataflow between phases of dataset creation. It is worth noting that, based on the nature of a particular study, a collection process may approach the four phases in different ways:

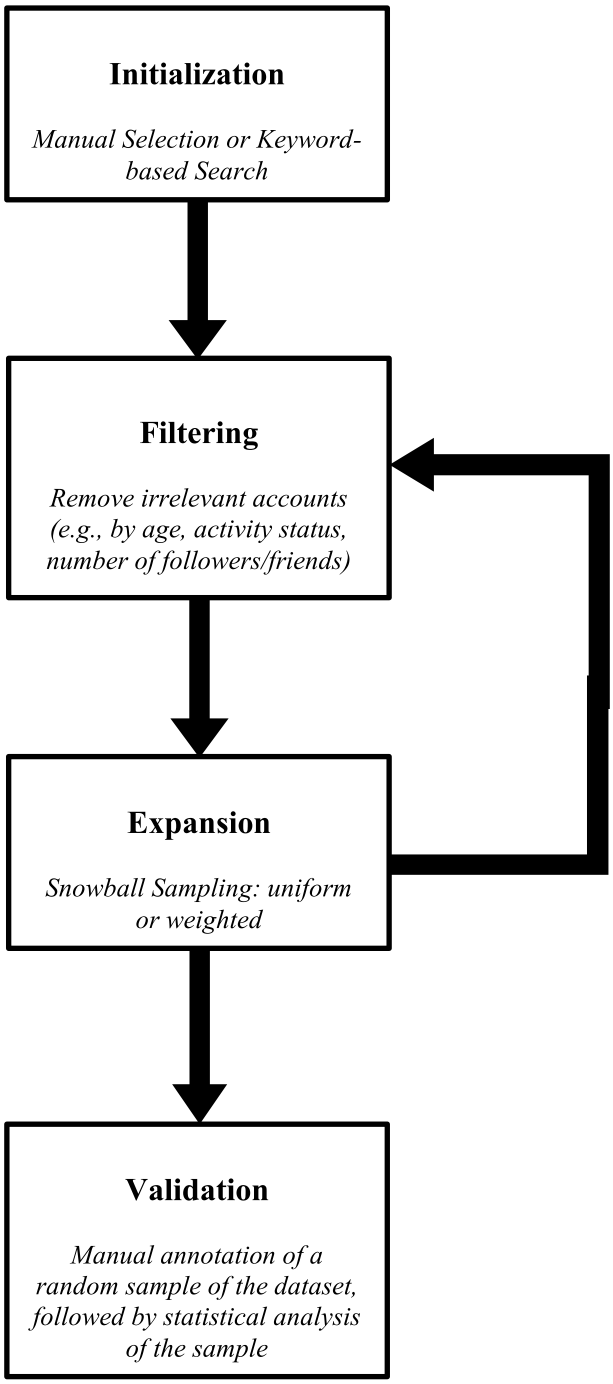

Figure 2: Four-phase Model of Data Collection

- In some cases, the Expansion phase may not be used at all. For example, Magdy et al. constructed a dataset by searching for tweets using Arabic keywords related to ISIS and the dataset was not expanded further.[17]
- In other situations, there can be more than one expansion phase where a dataset is iteratively grown to include accounts related to all the accounts collected in the previous Expansion phase.[18] This includes effectively all the second-level relations of initial seed accounts.

On the Generality of the Proposed Model

When we consider the kind of data typically used in terrorism research, we find that the structure as well as sources of such data are very similar across many social media platforms including Twitter and Telegram. There are three main types of data used in terrorism studies:

1. **User data**. This includes all user profile/account information such as profile description, demographic information and photos. In addition, any content produced by a user such as messages including text, images, videos and other media.
2. **Group data**. If a social media platform facilitates group messaging and broadcasting, it generates group data including group profile information, group member information as well as messages.
3. **Interaction data**. This is essentially network data obtained from relationships and interaction among users on a social media platform.

The data collection process for any of the above types essentially involves choosing a seed set of users or groups - initialization phase, as well as expansion phase if required. In addition, sampling from such datasets, including networks, necessitates the same kind of filtering and validation as applied to the Twitter datasets discussed in the article. As a result, our proposed model naturally extends to a wide array of social media platforms including all those that have been considered by the terrorism research community.

Strategies for Data Collection in the Existing Literature

The general model of data collection from the previous section provides a high-level framework through which we can view the data collection methods in existing literature: considering them in terms of the four phases of the model. By doing this, we can study differences between the methods and identify limitations for each of the phases. In this section, we describe in detail specific strategies, which are employed by researchers, corresponding to the four phases of the model. Table 1 summarizes these strategies for all research articles we study in this paper.

Initialization Strategies

Strategies for the initialization phase include *manual selection* and *keyword-based search*. In manual selection, the seed dataset can be created manually in the following two ways:

1. Experts in terrorism research obtain Twitter accounts of jihadists using various sources such as news stories, blogs, reports released by law enforcement agencies, and data from other terrorism research.[19]

2. From a jihadist account that maintains a list of other jihadist supporter accounts. Berger et al., for example, used an ISIS account “Baqiya Shoutout” to get an initial set of user accounts.[20]

In a keyword-based search, there are two steps:

1. The first step is to collect a set of tweets from Twitter by searching for specific keywords such as ‘The Islamic State’ in the text of the tweets. Keywords are chosen by experts, based on practical knowledge about jihadist groups, and validated by statistical methods, or by manual annotation using human

coders.

- Bodien-Barone , used a log-likelihood based measure of distinctiveness of keywords to validate their choice of variations of "Daesh" and "The Islamic State" in Arabic.[21]
- Rowe et al. used two coders fluent in Arabic and English to label tweets with pro- and anti-ISIS terms, and selected keywords based on an inter-rater agreement statistic between the coders.[22]
- Magdy et al. hypothesized that using the full name to refer to the Islamic State indicated ISIS support as opposed to using the abbreviated name, which indicated ISIS opposition. They validated this choice of keywords (ISIS name variations) by having a human annotator judge a sample of tweets obtained using the keywords.[23]

2. In the second step, a set of users, who posted the tweets collected in the first step is obtained. This set of user accounts forms the initial seed set, which can be filtered further and expanded.

Expansion Strategies

Strategies for dataset expansion include variations of the snowball sampling that we discussed above. The user accounts in the seed set obtained from an initialization phase are used as starter nodes for the sampling of the social graph. The commonly used variations of snowball sampling are as follows:

1. *Random snowball sampling.* For each starter node, all of its neighbors are added to the set of sampled nodes and the sampled graph is constructed by taking all the edges that connect to the sampled nodes. The majority of the research papers we have studied use this sampling technique.[24]
2. *Weighted snowball sampling.* In the case of weighted sampling of any given dataset, each neighbor of a starter node is assigned a weight and the sampling process chooses neighbors with a probability that is proportional to its assigned weight. For example, in the case of *degree-weighted snowball sampling*, each neighbor in the social graph is assigned a weight equal to number of its followers or friends. Then, in the sampling process, each neighbor of a node is chosen with a probability that is proportional to its assigned weight. In other words, neighbors with higher number of connections (followers or friends), are more likely to be sampled.[25]

Filtering and Validation Strategies

As shown in Table 1, at least half of the terrorism research studies, except Berger et al. (2015), Wright et al., Conway et al., Kaati et al. and Smedt et al.,[26], either do not perform a filtering phase to improve their datasets or their criteria for filtering are too weak. The most common strategy employed was to remove inactive accounts, but this is not strict enough to ensure that the remaining accounts are mostly jihadists.

Similarly, for the validation phase, datasets are not manually annotated or assessed in many articles shown in Table 1. While Magdy et al. do validate their dataset, it is only with a small random sample size of 50.[27] Some of the research studies do not verify the proportion of jihadist user accounts in their datasets, even though they validate other aspects of their methods. Bodine-Baron et al. and Magdy et al., confirm their choice of pro- and anti-ISIS keywords or terms used to annotate user accounts as pro- or anti-ISIS. In this way, it is effectively assumed that their dataset mostly contains jihadist accounts.[28] Considering the difficulty in tracking jihadist accounts generally, and their clandestine way of operating, it is imperative that researchers do more to validate their samples.

The lack of both a proper criteria for filtering and validating in much of the existing literature is a serious problem. In the next section we describe in detail such limitations in the existing literature.

Limitations of Data Collection Strategies

Based on the analysis of different data collection strategies, we found that there are two recurring high-level problems: (1) a fundamental lack of characterizing account inclusion errors and (2) errors introduced by graph sampling and lack of filtering.

In this section, we describe these problems in the context of Twitter datasets in the literature. In all cases, our observations readily generalize to any study of jihadist use of social media.

Lack of Characterization of Account Inclusion

The majority of studies (6 out of 8) made no clear or credible attempt to characterize the extent to which their account collection process actually did collect jihadist accounts (defined as per the objective of the specific study). Given that these studies then went on to make claims about the activities, relationships among, and fates of jihadist accounts, this omission is surprising and troubling. If we do not know the proportion of accounts in the dataset that *actually* are jihadists, it is impossible to attribute observed trends to jihadist online activity.

Table 1: Data Collection Strategies in Existing Literature

Article	Dataset Initialization & Filtering	Dataset Expansion & Filtering	Validation
Berger et al., 2015	*Manual selection:* tweets of 4700 accounts manually assessed to remove non-ISIS supporters *Filtering:* remove accounts 1) with >500 followers, 2) did not tweet within past four months	*Random Snowball Sampling:* Using Friend relationship ***Filtering:*** remove accounts 1) with > 50,000 followers, and 2) that are likely to be bots *Remarks:* Dataset was further expanded to include level 2 and 3 accounts.	***For Presence of Jihadist users:*** A random sample of 1000 were manually annotated using a Data Codebook
Kaati et al., 2015	*Manual Selection:* Accounts from Shumukh al-Islam forum and their followers, manually identified *Keyword-based Search:* Tweets containing jihadist propaganda based on hashtags and network of known jihadists.	*Filtering:* Clusters of known Jihadist sympathizers were used to select tweets with hashtags	Validation of accounts: Manual coding *Validation of Tweets*: None
Klausen, 2015	*Manual selection*: 60 Western foreign fighters in Syria ***Filtering:*** 1 inactive account removed	*Random Snowball Sampling:* Using both Follower and Friend relationships ***Filtering:*** None	None
Magdy et al., 2015	*Keyword-based search:* Search Arabic tweets for full name and acronym of ISIS. ***Filtering:*** remove users: 1) suspended and deleted, 2) posted <10 tweets mentioning ISIS, 3) <70% of tweets strictly using either full name or acronym of ISIS.	None	*Validation of Keywords:* A random sample of 1000 tweets were annotated for pro-, anti-ISIS, or neutral ***For Presence of Jihadist users:*** a random sample of 50 each from pro- and anti-ISIS groups of users were annotated

Berger et al., 2016	*Manual selection:* List of accounts maintained by user "Baqiya Shoutout". ~1000 users. ***Filtering:*** remove accounts with >9,500 followers	In the article, it is mentioned that using social network analysis ~2500 English speaking ISIS-supporting accounts were found. But, details are not given.	It is not clear from the article, but it seems the list was assessed manually.
Bodine-Baron et al., 2016	*Keyword-based search:* Search tweets for grammatical variations of "Islamic State" and "Daesh" in Arabic. ~23M tweets and 771K users	Network is formed by Mentions relationship between 771K users. 3.3M mentions or edges. ***Filtering:*** None	*Validation of Keywords:* Using lexical analysis methods such as log likelihood scores and collocates. ***For Presence of Jihadist users:*** None
Klausen et al., 2016	*Manual selection:* ~5000 seed users	*Random Snowball Sampling:* Using both Follower and Friend relationships ***Filtering:*** Remove suspended accounts. ~647K accounts	*Validation of prediction model:* Using samples of active and suspended accounts, 5000 each. ***For Presence of Jihadist users:*** None
Rowe et al., 2016	*Manual selection:* 652 seed users from a prior research ***Filtering:*** Select users who are active, not deleted and their timeline visible	*Random Snowball Sampling:* Using Follower relationship ***Filtering:*** Select accounts based in Europe. ~154K users	*Validation of pro- anti-ISIS Lexicon:* Manual annotation of tweets containing pro- and anti-ISIS terms (sample of 1K each) ***For Presence of Jihadist users:*** None
Wright et al., 2016	*Manual selection:* Publicly known, official media jihadist accounts named by newspapers	*Weighted Snowball Sampling:* Using Follower relationship ***Filtering:*** As a part of sampling process. Sample accounts that: 1) have <1000 followers, 2) are followed by >10% of users in the sampled dataset, 3) are active	***For Presence of Jihadist users:*** Visual assessment of accounts to identify resurging accounts
Conway et al., 2017	*Manual selection:* 722 pro-IS and 451 other jihadist accounts obtained using three methods: 1) original seed set of accounts (27%) were manually identified from known jihadist accounts and their followers and friends, 2) second set (30%) was obtained semi-automatically and verified manually, 3) third set (43%) was obtained semi-automatically using known IS propaganda links. Seed accounts were obtained by various methods including Keyword-based search. *Filtering:* Selected only accounts with at least one recent tweet with pro-IS text/images. Manually excluded accounts maintained by journalists	None	*Validation of accounts:* Manual Coding based on content of tweet: keywords in text and images

Smedt et al., 2018	Keyword-based Search: 367 subversive profiles found by searching keywords such as "kuffar", followed by manual inspection. All tweets for the accounts were obtained	*Expansion:* None *Filtering:* Accounts were manually inspected. Tweets are not filtered but it is fine for the research task	*Validation of accounts:* Manual Coding

Sources of Error

In the case of the Twitter datasets we examined, we observed four specific and common sources of error: inclusion of irrelevant accounts, exclusion of jihadist accounts, bias in sampling networks, and bias towards jihadist communities.

1. *Inclusion of Irrelevant Accounts*

This is one of the most common problems we have seen in the methods of the existing literature. It is a central critique in this article that the issue of irrelevant accounts is very prevalent and should not be ignored. Inclusion of irrelevant accounts could happen in two ways:

a) *Bias in sampling due to choice of relationships.* The bias toward irrelevant accounts occurs due to a particular choice of relationships that are explored in the process of graph sampling. Irrelevant or noisy accounts, such as news reporters and researchers, who are just "listening" to other jihadist accounts, are likely to be included when the datasets are constructed using strategies involving follower relationships in the dataset expansion phase.[29]

b) *Lack of appropriate Filtering and Validation.* Arguably, using Friend relationships might result in a lesser proportion of irrelevant accounts because jihadists are more likely to follow other jihadists or supporters than irrelevant accounts.[30] Furthermore, when using a keyword-based search, noisy accounts are likely to include those who report news and condemn or oppose jihadists content and happen to use the keywords in tweet content.[31] In these cases, when sampling methods are fixed and cannot be improved but data is still biased, proper criteria for filtering and validation of datasets can help mitigate the problem. As shown in Table 1, however, lack of proper filtering and validation phases in the literature results in the inclusion of irrelevant accounts.

2. *Exclusion of Jihadist Accounts*

Exclusion of jihadist accounts also happens, due to bias in graph sampling, where particular relationships are chosen to be explored over others, combined with research objectives. For example, on Twitter, when the Friend relationship is chosen, the resulting dataset is likely to exclude the accounts of those who are potentially "radicalizing," since they have almost no followers and are largely just "listening" to other prominent jihadists. If the research objective is to detect users who are radicalizing or tracing the process of radicalization, then this bias is problematic. Likewise, a keyword-search approach is likely to miss potential jihadist accounts, if such accounts do not post tweets containing the keywords.

Since Friend and Follower networks are generally huge, sampling is performed to get smaller networks that are easier to work with. In such cases, if the proportion of irrelevant accounts compared to that of jihadist accounts is high, there is a significant chance that the sampled dataset will miss many potential jihadist accounts. This substantially decreases the reliability of the statistical results reported for such a dataset. Therefore, as mentioned earlier, choosing appropriate steps in the filtering phase or validating the quality of the dataset using manual annotation is indeed important.

3. *Bias in Sampling Networks due to Choice of Sampling Methods*

When a large network is sampled to obtain a smaller network for social network analyses, it is important to choose the appropriate sampling method. As discussed earlier, typically in terrorism research, random node sampling or random snowball sampling is used, both of which are shown to result in a sampled network that has different properties than the original network.[32] This is likely to produce incorrect results in calculating network statistics such as highly connected nodes or distance between two nodes. Alternatives, such as random walk sampling, which sufficiently preserve network properties, should be considered.[33]

4. *Bias towards Jihadist Communities*

During an Expansion phase, relationships of seed accounts are explored to add more accounts to the dataset. Any two different sets of seed accounts are likely to result in different datasets after the Expansion phase depending on whether the sets belong to same network communities or not. Therefore, a specific choice of the initial seed accounts produces a dataset that is biased towards specific communities of jihadist accounts. This bias is acknowledged in the existing literature, and indeed, for some research tasks, it is reasonable to focus on certain jihadist communities.[34] Nonetheless, this bias is an important factor to consider when creating a dataset.

Specific Errors in the Existing Literature

The methodological limitations described above inevitably lead to biased or incorrect results from quantitative and statistical analyses. One common class of errors involves fundamental network statistics. When the presence of irrelevant accounts in a Twitter network is not accounted for, the network does not represent true relationships among jihadists. Klausen et al, 2015, reported popularity ranking of accounts in their network dataset.[35] They used the degree centrality which is a simple measure that counts how many neighbors each node or accounts has. However, if an account has significant proportion of irrelevant accounts as its neighbors, the degree centrality measure would incorrectly assign more popularity to that account.

Another popular network analysis task is to detect structured communities within a network. Bodine-Baron et al., 2016, discovered jihadist metacommunities using a community detection algorithm.[36] Their network dataset was formed using "mention" relationships extracted from 23 million tweets, which resulted in a massive network of 771K nodes and 3.3 million edges. Given the collection methods employed, many of these accounts are likely to be researchers, journalists and other non-jihadist accounts whose relationships span across the metacommunities. Not only would this confuse the community inference task, but any inferred communities would wrongly assign these non-jihadist accounts (since, by definition, they do not belong to *any* jihadist community). As a result, the community detection algorithm would inevitably yield misrepresentative metacommunities – both in terms of content and size.

Finally, we observed lack of validation of linguistic data in Rowe, 2016. In order to study differences in pro- vs. anti-ISIS linguistic characteristics, Rowe, 2016 created a pro- and anti-ISIS lexicon. They validated the lexicon by manual coding the words in tweets as pro- and anti-ISIS. However, they do not validate if the underlying Twitter data comes from jihadist accounts. As a result, it is incorrect to assume that use of pro-ISIS language implies that the account is an ISIS supporter itself.

Recommended Best Practices

Based on the issues identified in the previous section, here we provide several recommendations for standards of practice that should be adopted by the research community. To lend credence to these recommendations, we apply them to a new benchmark dataset to show how they increase the validity and decrease errors.

We identify three practices that would substantively improve the quality of social media datasets and, as a result, enhance the quality and validity of findings derived from them.

Perform manual validation on datasets. The single greatest danger and opportunity for improvement posed by current data collection practices is the lack of a validation stage in dataset collection – which is to say, a lack of characterization of error in datasets. Without knowing what kind of errors are present in a dataset, it is quite literally impossible to account for those errors in the interpretation of trends discovered in the data. As a result, findings based on data whose errors have not been systematically characterized are at extreme risk of incorrectly crediting features of the data (e.g., the behavior of users, salient structures in social networks, uses of language) to jihadist origins.

Of course, we do assume that researchers are committed to conducting thoughtful and representative analysis. This brings us to the second benefit to characterizing errors: knowing what they are will undoubtedly motivate researchers to address them. We suspect that the widespread lack of filtering and use of biased network sampling techniques are most likely due to ignorance that these practices produce serious biases. Requiring quantitative characterization of the errors in datasets would make researchers keenly aware of any biases and, therefore, more empowered to address them by recollecting data and revising their methods.

One (relatively) easy, established way of characterizing error in datasets is to manually annotate a random sub-sample of the final dataset. In the context of judging the inclusion/exclusion of jihadist accounts, this would involve randomly selecting 100 accounts from the dataset and coding their connection to a jihadist movement (e.g., "jihadist", "supporter", "irrelevant"). The distribution of labels in the random sample would provide a strong indication of the quality of the dataset and any biases or errors present in it.

Use Friend Relationships in an Expansion Phase. Berger et al. argue that irrelevant accounts are more likely to follow jihadist accounts than being followed.[37] We have independently confirmed this (see next section) and, therefore, advise that the network-based collection of jihadist accounts gather Friends of jihadist accounts in a dataset, rather than Followers.

Filter Collected Accounts. Based on the statistics related to known jihadist accounts, it is reasonable to assume that such accounts do not usually have a very high number of followers or friends. Moreover, long-dormant or entirely inactive accounts do not provide sufficient information on the current state of jihadist activities and hence can be removed from a dataset. However, many studies do not use any of these filters. Happily, all of these criteria can be implemented using simple filters - placing a threshold on the number of followers, friends and the age of accounts – and can substantially reduce the number of irrelevant accounts.[38]

The Impact of Recommended Practices

In the previous section we provided three recommendations for best practices. In this section, we demonstrate that these do, in fact, substantially improve the quality of datasets. To do this, we conduct a study of the impact of filtering and network expansion on inclusion of irrelevant accounts in collected datasets. We also highlight just how high the number of irrelevant accounts can be, even when applying all best practices.

It would have been ideal to conduct this assessment on datasets from prior work. However, we faced two challenges: 1) datasets for most research studies could not be shared and 2) for those datasets we could obtain, many Twitter accounts are already suspended or deleted. Therefore, we created new datasets employing the standard techniques employed by the jihadist research community. We created four benchmark datasets that serve as proxies for datasets in the literature.

Dataset construction

Our datasets were initialized with 47 known jihadist accounts chosen by one of the authors of this article who is an expert in terrorism research. All data was collected during the month of June 2017.

For the Expansion phase, we used Follower and Friend relationships, each creating two separate datasets. This allows us to test the Expansion-related recommendation.

For each of the two datasets, we wanted to compare effects of the filtering phase on the quality of data (the third recommendation). Therefore, for each expansion method (Friend/Follower), we created a dataset with/without account filters. For the filtering phase, we removed accounts that satisfied two criteria: 1) accounts that were more than one year old (created one year before the time of data collection) and 2) accounts that were highly connected (with more than 1.000 followers or friends).

This process yielded the following four datasets, which we use for further analysis:

1. A dataset that was expanded, using Follower relationships, but a filtering phase was not performed.
2. The same dataset as in (1), but a filtering phase was performed.
3. A dataset that was expanded, using Friend relationships, but a filtering phase was not performed.
4. The same dataset, as in (3), but a filtering phase was performed.

In order to assess the composition of each dataset, we took a random sample of 100 accounts from each of the datasets and manually annotated them. Each Twitter account in our datasets was given one of five different labels. These labels were chosen to simplify the annotation process, so that the accounts could be labelled manually by non-experts in terrorism research. The five labels are:

1. Positive Islamic. These are the accounts that post largely positive tweets about Islam, which include messages of peace, love, equality, and condemn hatred or violence.
2. Radical Islamist. These are the accounts that clearly spread messages of hatred and violence against civilians.
3. Ambiguous accounts. These are accounts that mix both positive and radical comments, making it hard to clearly distinguish if it warrants being placed in one of the first two categories.
4. Irrelevant accounts. These are the accounts that appear to belong to normal twitter users and do not post any content related to Islam or Jihad, in either positive or negative ways.
5. Insufficient Information. This label is used when an account does not have any tweets, or otherwise lacks sufficient information to be classified using the other four categories.

Findings

Figures 3 and 4 show the proportion of account types for each of the four benchmark datasets. There are several noteworthy trends which underscore the importance of our recommendations.

Friend expansion yields more jihadist accounts than ollower expansion. Comparing between Figures 3 and 4, we see that friend expansion yields well over 10% more accounts of interest.

Filtering enriches samples for jihadist accounts. In both Figures 3 and 4, if we compare within labels, we find that filtering substantively increases the proportion of jihadist accounts in the dataset. It appears that filtering has a more dramatic effect on follower-expanded datasets – which stands to reason since follower expansion tends to include more irrelevant accounts.

The majority of accounts are irrelevant. This point cannot be overstressed. We find that in all four of our datasets, the vast majority of accounts are irrelevant. This should give all researchers in this field pause, as it suggests that the majority of accounts involved in previous studies of jihadists online may well also be irrelevant. It would be fair to point out that there are far fewer jihadists active on Twitter than in past years. However, the point still stands that *nobody knows* how many irrelevant accounts are present in past studies—efforts were not made by the researchers to estimate this problem. And, certainly, it would be unprincipled to simply assume that proportions are different without evidence. All this points to two important lessons: first, it may well be that much of what is known from past studies of online jihadist behavior is highly skewed by irrelevant accounts. Second, manual validation should be a required phase in all future research projects, and this validation should be reported and discussed in publications.

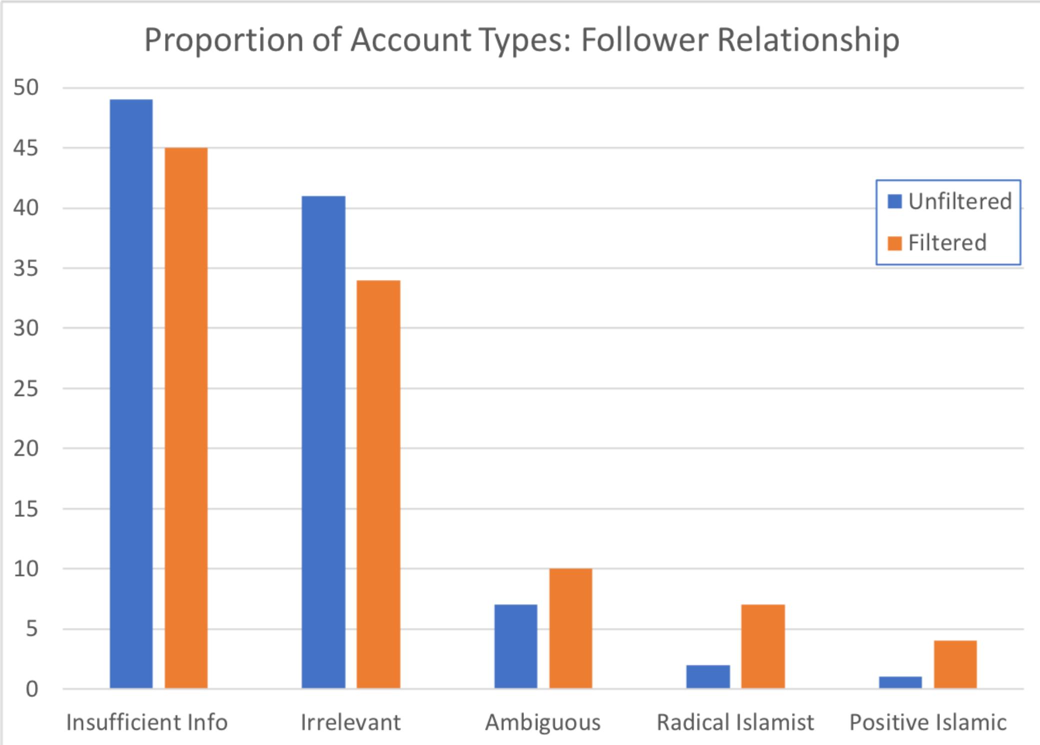

Figure 3. Proportion of different account types/labels in our annotated benchmark dataset created using Follower relationship. Size of the dataset is 100.

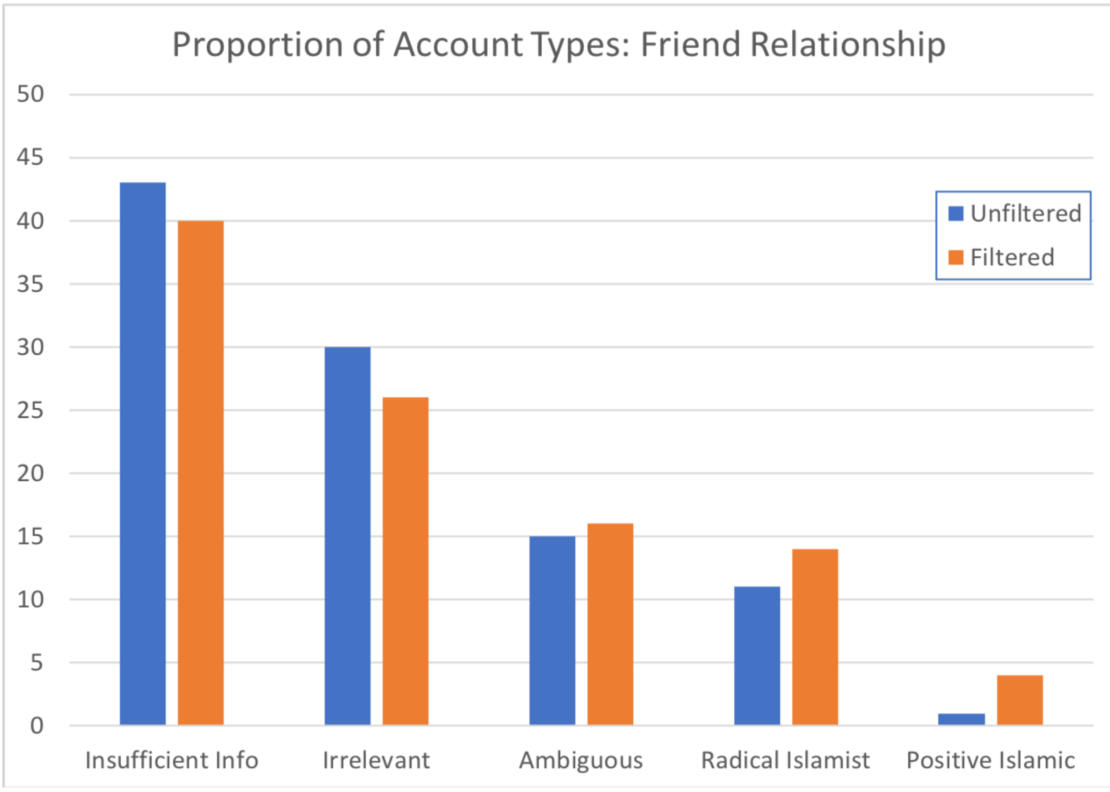

Figure 4. Proportion of different account types/labels in our annotated benchmark dataset created using Friend relationship. Size of the dataset is 100.

Conclusion

In this article, we set out to provide a critical analysis of existing data collection methods on Twitter. In the process of doing so, we have made three distinct contributions.

First, we presented an overview of the structure of social media data and proposed a general model of the data collection processes. We showed how users of social media platforms and their relationships or interactions are modeled using social graph. We highlighted the significance of graph sampling for the process of data collection from social graphs. By analyzing existing data collection methods in Twitter terrorism research, and generalizing them, using social graphs and graph sampling, we devised a four-phase model of data collection from any social media platform similar in structure to Twitter.

Second, applying our proposed model to existing data collection methods, we categorized existing approaches according to the different strategies they implement. To our knowledge, this article is the first to compare data collection practices in the context of social media terrorism research. The comparison enabled us to observe best practices in the literature for different phases, which we justified through the analysis of a set of benchmark datasets.

Finally, as part of the analysis of our datasets, we discovered that current data collection methods (even when using our recommended best practices) have exceptionally high rates of irrelevant account inclusions. Quite alarmingly, we have, at present, no basis to assume that rates are not similarly high in datasets used by prior research studies – raising serious questions about the validity of findings and trends reported in past work.

Our hope is that this work has highlighted key ways in which research on online jihadist behavior can be placed on more solid methodological grounds – and, in doing so, render analyses and findings that propel forward our understanding of the phenomena of jihadism in the modern online world.

About the Authors: **Deven Perekh** *is a PhD student at the School of Computer Science at McGill University.* **Amarnath Amarasingam** *is a Senior Fellow at the Institute for Strategic Dialogue and a postdoctoral fellow at the University of Waterloo.* **Lorne Dawson** *is Full Professor in the Department of Sociology and Legal Studies and the Department of Religious Studies at the University of Waterloo.* **Derek Ruths** *is an Associate Professor in the School of Computer Science* at *McGill University, where he runs the Network Dynamics Lab.*

Notes

[1] Bloom, M., Tiflati, H. and Horgan, J., 2017. Navigating ISIS's Preferred Platform: Telegram. *Terrorism and Political Violence,* Vol. 29, pp.1-13.

[2] See: Ahmed, N.K., Neville, J. and Kompella, R., 2014. Network sampling: From static to streaming graphs. *ACM Transactions on Knowledge Discovery from Data* (TKDD), 8(2), p.7. URL: https://arxiv.org/abs/1211.3412 as well as Wagner, C., Singer, P., Karimi, F., Pfeffer, J. and Strohmaier, M., 2017 (April). Sampling from Social Networks with Attributes. In *Proceedings of the 26th International Conference on World Wide Web* (pp. 1181-1190). International World Wide Web Conferences Steering Committee. URL: https://arxiv.org/abs/1702.05427.

[3] Klausen, J., 2015. Tweeting the Jihad: Social media networks of Western foreign fighters in Syria and Iraq. *Studies in Conflict & Terrorism*, 38(1), pp.1-22.

[4] Berger, J.M. and Morgan, J., 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, Analysis Paper No. 20.* URL: https://www.brookings.edu/wp-content/uploads/2016/06/isis_twitter_census_berger_morgan.pdf .

[5] Berger, J.M. and Perez, H., 2016. The Islamic State's Diminishing Returns on Twitter: How Suspensions are Limiting the Social Networks of English-speaking ISIS Supporters. George Washington University. Program on Extremism. URL: https://cchs.gwu.edu/sites/g/files/zaxdzs2371/f/downloads/ Berger_Occasional%20Paper.pdf .

[6] Bodine-Baron, E., Helmus, T.C., Magnuson, M. and Winkelman, Z., 2016. *Examining ISIS Support and Opposition Networks on Twitter.* RAND Corporation. URL: https://www.rand.org/pubs/resea rch_reports /RR1328.html .

[7] Conway, M., Khawaja, M., Lakhani, S., Reffin, J., Robertson, A. and Weir, D., 2017. Disrupting Daesh: measuring takedown of online terrorist material and its impacts. Vox-Pol. URL: http://www.voxpol.eu/wpfb-file/dcuj5528-disrupting-daesh-1706-web-v2-pdf .

[8] Magdy, W., Darwish, K. and Weber, I., 2015. #Failed Revolutions: Using Twitter to study the antecedents of ISIS support. First Monday. URL: http://firstmonday.org/ojs/index.php/fm/article /view/6372/5194 .

[9] Kaati, L., Omer, E., Prucha, N. and Shrestha, A., 2015, November. Detecting multipliers of jihadism on twitter. In *Data Mining Workshop (ICDMW), 2015 IEEE International Conference on* (pp. 954-960). IEEE.

[10] Klausen, J., Marks, C. and Zaman, T., 2016. Finding Online Extremists in Social Networks. URL: https://arxiv.org/abs/1610.06242 .

[11] Rowe, M. and Saif, H., 2016 (May). Mining Pro-ISIS Radicalisation Signals from Social Media Users. *Proceedings of the Tenth International AAAI Conference on Web and Social Media* (ICWSM), pp. 329-338.

[12] Wright, S., Denney, D., Pinkerton, A., Jansen, V. and Bryden, J., 2016. Resurgent insurgents: Quantitative research into Jihadists who get suspended but return on Twitter. *Journal of Terrorism Research,* 7(2). URL: https://cvir.st-andrews.ac.uk/articles/10.15664/jtr.1213/print/ .

[13] De Smedt, T., De Pauw, G. and Van Ostaeyen, P., 2018. Automatic Detection of Online Jihadist Hate Speech. *arXiv preprint arXiv:1803.04596.*

[14] Bodine-Baron, E., Helmus, T.C., Magnuson, M. and Winkelman, Z., 2016. *Examining ISIS Support and Opposition Networks on Twitter.* RAND Corporation. URL: https://www.rand.org/pubs/ research_reports/RR1328.html .

[15] Berger, J.M. and Perez, H., 2016. *The Islamic State's Diminishing Returns on Twitter: How Suspensions are Limiting the Social Networks of English-speaking ISIS Supporters*. George Washington University. Program on Extremism. URL: https://cchs.gwu.edu/sites/g/files/zaxdzs2371/f/ downloads/Berger_Occasional%20Paper.pdf .

[16] Tufekci, Z. 2014. Big Questions for Social Media Big Data: Representativeness, Validity and Other Methodological Pitfalls. In *Proceedings of the 8th International Conference on Web and Social Media* (ICWSM) (pp. 505-514); Ruths, D. and Pfeffer, J. 2014. Social Media for large studies of behavior. Science, 346(6213), pp. 1063-1064.

[17] Magdy, W., Darwish, K. and Weber, I., 2015. #Failed Revolutions: Using Twitter to study the antecedents of ISIS support. First Monday. URL: http://firstmonday.org/ojs/index.php/fm/article/ view/6372/5194 .

[18] Berger, J.M. and Morgan, J., 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, Analysis Paper No. 20.* URL: https://www.brookings.edu/wp-content/uploads/2016/06/isis_twitter_census_ berger_morgan.pdf .

[19] See: Klausen, J., 2015. Tweeting the Jihad: Social media networks of Western foreign fighters in Syria and Iraq. *Studies in Conflict & Terrorism,* 38(1), pp.1-22; Klausen, J., Marks, C. and Zaman, T., 2016. Finding Online Extremists in Social Networks. URL: https://arxiv.org/abs/1610.06242; Rowe, M. and Saif, H., 2016 (May).
Mining Pro-ISIS Radicalisation Signals from Social Media Users. Proceedings of the Tenth International AAAI Conference on Web and Social Media (ICWSM), pp. 329-338; Wright, S., Denney, D., Pinkerton, A., Jansen, V. and Bryden, J., 2016. Resurgent insurgents: Quantitative research into Jihadists who get suspended but return on Twitter. *Journal of Terrorism Research,* 7(2). URL: https://cvir.st-andrews.ac.uk/articles/10.15664/jtr.1213/print/; Berger, J.M. and Morgan, J., 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, Analysis Paper No. 20.* URL: https://www.brookings.edu/wp-content/uploads/2016/06/isis_twitter_census_berger_morgan.pdf; Conway et al., D., 2017. Disrupting Daesh: measuring takedown of online terrorist material and its impacts.

[20] Berger, J.M. and Perez, H., 2016. *The Islamic State's Diminishing Returns on Twitter: How Suspensions are Limiting the Social Networks of English-speaking ISIS Supporters*. George Washington University. Program on Extremism. URL: https://cchs.gwu.edu/sites/g/files/zaxdzs2371/f/ downloads/Berger_Occasional%20Paper.pdf .

[21] Bodine-Baron, E., Helmus, T.C., Magnuson, M. and Winkelman, Z., 2016. *Examining ISIS Support and Opposition Networks on Twitter.* RAND Corporation. URL: https://www.rand.org/pubs/resea rch_reports /RR1328.html .

[22] Rowe, M. and Saif, H., 2016 (May). Mining Pro-ISIS Radicalisation Signals from Social Media Users. *Proceedings of the Tenth International AAAI Conference on Web and Social Media* (ICWSM), pp. 329-338.

[23] Magdy, W., Darwish, K. and Weber, I., 2015. #Failed Revolutions: Using Twitter to study the antecedents of ISIS support. First Monday. URL: http://firstmonday.org/ojs/index.php/fm/article /view/6372/5194 .

[24] See: Berger, J.M. and Morgan, J., 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, Analysis Paper No. 20.* URL: https://www.brookings.edu/wp-

content/uploads/2016/06/isis_twitter__census_berger_morgan.pdf; Klausen, J., 2015. Tweeting the Jihad: Social media networks of Western foreign fighters in Syria and Iraq. *Studies in Conflict & Terrorism,* 38(1), pp.1-22; Klausen, J., Marks, C. and Zaman, T., 2016. Finding Online Extremists in Social Networks. URL: https://arxiv.org/abs/ 1610.06242; Rowe, M. and Saif, H., 2016 (May). Mining Pro-ISIS Radicalisation Signals from Social Media Users. *Proceedings of the Tenth International AAAI Conference on Web and Social Media* (ICWSM), pp. 329-338.

[25] See: Wright, S., Denney, D., Pinkerton, A., Jansen, V. and Bryden, J., 2016. Resurgent insurgents: Quantitative research into Jihadists who get suspended but return on Twitter.

[26] Berger, J.M. and Morgan, J., 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, Analysis Paper No. 20.* URL: https://www.brookings.edu/wp-content/uploads/2016/06/isis_twitter_census__berger_morgan.pdf; Wright, S., Denney, D., Pinkerton, A., Jansen, V. and Bryden, J., 2016. Resurgent insurgents: Quantitative research into Jihadists who get suspended but return on Twitter. *Journal of Terrorism Research,* 7(2). URL: https://cvir.st-andrews.ac.uk/articles/10.15664/jtr.1213/print/.

[27] Magdy, W., Darwish, K. and Weber, I., 2015. #Failed Revolutions: Using Twitter to study the antecedents of ISIS support. First Monday. URL: http://firstmonday.org/ojs/index.php/fm/article/ view/6372/5194 .

[28] See: Bodine-Baron, E., Helmus, T.C., Magnuson, M. and Winkelman, Z., 2016. *Examining ISIS Support and Opposition Networks on Twitter.* RAND Corporation. URL: https://www.rand.org/pubs/research__reports/RR1328.html; Magdy, W., Darwish, K. and Weber, I., 2015. #Failed Revolutions: Using Twitter to study the antecedents of ISIS support. First Monday. URL: http://firstmonday.org/ojs/index. php/fm/article/view/6372/5194.

[29] See: Klausen, J., 2015. Tweeting the Jihad: Social media networks of Western foreign fighters in Syria and Iraq. *Studies in Conflict & Terrorism,* 38(1), pp.1-22; Klausen, J., Marks, C. and Zaman, T., 2016. Finding Online Extremists in Social Networks. URL: https://arxiv.org/abs/1610.06242; Rowe, M. and Saif, H., 2016 (May). Mining Pro-ISIS Radicalisation Signals from Social Media Users. Proceedings of the Tenth International AAAI Conference on Web and Social Media (ICWSM), pp. 329-338.

[30] Berger, J.M. and Morgan, J., 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, Analysis Paper No. 20.* URL: https://www.brookings.edu/wp-content/uploads/2016/06/isis_twitter_census__berger_morgan.pdf .

[31] Magdy, W., Darwish, K. and Weber, I., 2015. #Failed Revolutions: Using Twitter to study the antecedents of ISIS support. First Monday. URL:: http://firstmonday.org/ojs/index.php/fm/article/ view/6372/5194; Bodine-Baron, E., Helmus, T.C., Magnuson, M. and Winkelman, Z., 2016. *Examining ISIS Support and Opposition Networks on Twitter.* RAND Corporation. URL: https://www.rand.org/pubs/research__reports/RR1328.html .

[32] See: Ahmed, N.K., Neville, J. and Kompella, R., 2014. Network sampling: From static to streaming graphs. ACM Transactions on Knowledge Discovery from Data (TKDD), 8(2), p.7. URL: https://arxiv.org/abs/1211.3412; Wagner, C., Singer, P., Karimi, F., Pfeffer, J. and Strohmaier, M., 2017 (April). Sampling from Social Networks with Attributes. In *Proceedings of the 26th International Conference on World Wide Web* (pp. 1181-1190). International World Wide Web Conferences Steering Committee. URL: https://arxiv.org/abs/1702.05427 .

[33] See Ibid. for further references on network sampling.

[34] See: Klausen, J., 2015. Tweeting the Jihad: Social media networks of Western foreign fighters in Syria and Iraq. *Studies in Conflict & Terrorism,* 38(1), pp.1-22; Magdy, W., Darwish, K. and Weber, I., 2015. #Failed Revolutions: Using Twitter to study the antecedents of ISIS support. First Monday. URL: http://firstmonday.org/ojs/index.php/fm/article/view/6372/5194; Wright, S., Denney, D., Pinkerton, A., Jansen, V. and Bryden, J., 2016. Resurgent insurgents: Quantitative research into Jihadists who get suspended but return on Twitter. *Journal of Terrorism Research,* 7(2). URL: https://cvir.st-andrews.ac.uk/articles/10.15664/jtr.1213/print/ .

[35] See: Klausen, J., 2015. Tweeting the Jihad: Social media networks of Western foreign fighters in Syria and Iraq, 7.

[36] See: Bodine-Baron et al., 2016. Examining ISIS Support and Opposition Networks on Twitter, pp. 13-27.

[37] Berger, J.M. and Morgan, J., 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, Analysis Paper No. 20.* URL: https://www.brookings.edu/wp-content/uploads/2016/06/isis_twitter_census__berger_morgan.pdf .

[38] See: Berger, J.M. and Morgan, J., 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, Analysis Paper No. 20.* URL: https://www.brookings.edu/wp-content/uploads/2016/06/isis_twitter_._census_berger_morgan.pdf; Wright, S., Denney, D., Pinkerton, A., Jansen, V. and Bryden, J., 2016. Resurgent insurgents: Quantitative research into Jihadists who get suspended but return on Twitter. *Journal of Terrorism Research,* 7(2). URL: https://cvir.st-andrews.ac.uk/articles/10.15664/jtr.1213/print/