World Academy of Science, Engineering and Technology
International Journal of Humanities and Social Sciences
Vol:16, No:3, 2022

Extracting Attributes for Twitter Hashtag Communities

Ashwaq Alsulami, Jianhua Shao

Abstract—Various organisations often need to understand discussions on social media, such as what trending topics are and characteristics of the people engaged in the discussion. A number of approaches have been proposed to extract attributes that would characterise a discussion group. However, these approaches are largely based on supervised learning, and as such they require a large amount of labelled data. We propose an approach in this paper that does not require labelled data, but rely on lexical sources to detect meaningful attributes for online discussion groups. Our findings show an acceptable level of accuracy in detecting attributes for Twitter discussion groups.

Keywords—Attributed community, attribute detection, community, social network.

I. Introduction

ONLINE social network platforms have resulted in a massive amount of user interaction data that could be usefully analysed to support a range of applications. One such analysis is the discovery of communities on social networks, which attempts to find groups of users who, for instance, share the same interests [1], [2], are connected in a certain way [3], communicate regularly with each other [4] or hold the same opinions on specific topics [5]. While various ways have been proposed to determine online communities, we consider the group of users who participate in the same Twitter hashtag (a word or phrase preceded by a hash sign #) discussion, and we are interested in extracting attributes that would characterise the participants in a hashtag group.

The ability to determine the attributes of the participants who support a particular topic is highly beneficial and can be applied in various fields, for instance, to personalise advertisements posted in relation to a trending event or topic, to support marketing companies in understanding their customers' motivations or to enhance content recommendations.

In the literature, several methods [6]–[11] have been proposed to extract user attributes from social media. Studies that have inferred Twitter users' attributes can be divided into three categories: those that depend on a network, those that depend on tweet content and those that depend on participants profiles. For example, Hu et al. [10] proposed a method that extracts occupation of the participants based on tweets, and Vijayaraghavan [6] analysed users' first names to infer their gender, picture features to extract their age and gender and information about followers or who they follow to predict their political orientation and location. In this paper we use user profiles to extract attributes. This is because of the expected massive amount of data needed from tweets in obtaining attributes, especially as we consider multiple rather than individual attribute(s). For example, Georgiou et al. [12] utilised between 1K and 200K tweets to obtain user's attribute. Also, extracting attributes from networking like [13], [14] does not work in our case as they can infer attributes only of users who follow certain popular accounts.

The method that is most relevant to our work is reported in [12], but it can extract only four attributes for a topic on Twitter (location, age, gender, and political affiliation) using different techniques to infer each attribute. Unlike previous research, we suggest a novel generic approach for counting the number of twitter profiles in a certain community that support a specific attribute, we call this technique Detection-based attributes extraction for online communities.

Although the method given in [12] is applicable to social media topics and generate focused communities, i.e. communities whose users share certain characteristics (e.g. location, age, gender or political situation), it does not consider other attributes. Since there may be attributes of interest that are not studied when one only looks at these attributes, other attributes should also be examined. The proposed method focuses instead on how a given attribute is supported within a online community. For example, when given an hashtag group and an attribute-value pair of interest such as $\langle Religion, Christianity\rangle$, the method searches through the hashtag group to determine if this attribute is sufficiently shared by the participants of the group.

The detection-based approach method determines the attribute probability of all *#hashtag community* participants on the basis of semantic relationship and similarity based on user queries. First, the method finds possible attribute values of giving attribute by using Wordnet through synonyms and hyponyms. Then, it calculates similarities between words used in user profiles and possible attribute values to identify relevant value by using a Word Embedding. Both techniques are well known, but not with Twitter data, in particular with the profile data. This method attempts to address the weaknesses in previous researches as follows:

- Previous studies have focused on a few specific user attributes. For example, Sloan et al. [9] inferred age and occupation, and Messias et al. [15] extracted gender and race from data. They are however unable to identify an arbitrary attribute of interest using a single generic method. In other words, previous studies has not created

A. Alsulami is with the School of Computer Science and Informatics, Cardiff University, Cardiff, United Kingdom (e-mail: alsulamiam@cardiff.ac.uk).

J. Shao is with the School of Computer Science and Informatics, Cardiff University, Cardiff, United Kingdom (e-mail: shaoj@cardiff.ac.uk).

a general method that can be applied to any attribute. Instead of looking for specific attributes, we have proposed a new generic method that can be used to find any desired attribute.
- Previous solutions depend on a large amount of labelled data, and as such they were limited by the lack of ground truth [16]. The purposed method minimises the requirement for labelled data.

The rest of the paper is organised as follows. In Section II, we outline the work related to our work. Section III gives the necessary background and in Section IV we describe the proposed method in detail. The experimental results are reported in Section V and finally we conclude the paper in Section VI.

II. Related Work

Researchers have studied different types of community in social networks. Initial work in this field focused primarily on communities connected as graphs. A group of vertices more densely linked than other vertices on the graph indicates a community [17], [18]. Since 2008, many studies have used both graphs and attributes to detect communities [19], [20]. In these studies, the vertices of graphs are characterised by the attributes associated with them. In social networks, user attributes such as age, gender, job and interests are used as attributes. Studying attributes in addition to network structures can improve our understanding of communities. In a more recent study [12], not structural connections but only attribute values are used to detect communities. This method inferred the characteristics of participants in discussions of a particular topic. The authors then used a detection methodology that identified communities through a combination of attribute values. However, this study could characterise users based only on four attributes: location, age, gender and political affiliation. Our work is an attempt to overcome this limitation and to consider any attribute that may be of interest to us.

Several methods have been proposed to extract user attributes from social media. Studies that have inferred Twitter users' attributes can be divided into three categories: those that depend on network, those that depend on text and those that depend on images. Some studies [6], [13], [21] have inferred users' attributes based on their networks, such as their followers, people they follow, friends and retweets. Other studies [7]–[9] have extracted attributes based on textual features, such as tweet contents, screen name, location and self-description. Several works [6], [15], [22] have attempted to infer user demographics through processing their profile images or even video posts.

Culotta et al. [13] proposed a method to extract Twitter users' demographics gender, age, ethnicity, education, income and child status based on whom they follow. The authors used the Quantcast website, which infers user demographics based on the cookies of millions of websites. This approach can therefore derive user demographics without analysing the content of tweets. The authors inferred the demographics of visitors to websites and then identified the Twitter accounts of these websites.

The method used in [14] was also based on non-textual features. The authors studied users' interests to predict their demographic attributes, such as gender, age, political affiliation, education and personality. The limitation of the approaches of [13], [14] is that they can infer the attributes only of users who follow certain popular accounts.

In contrast to studies based on non-textual features [13], [14], the method introduced by Schwartz et al. [7] was based on textual features. Gender, age and personality were extracted using a language model. They detected the language features (words, phrases and topics) of many Facebook posts via open-vocabulary analysis instead of closed-vocabulary analysis. The authors found that open-vocabulary analysis produced further insights and more accurate information than closed-vocabulary analysis. However, their model identified only the age and gender of users.

A study by Hu et al. [10] followed [7]'s proposed methodology, using fixed lexicon (LIWC) and open-vocabulary approaches to investigate the relation between users' language in tweets and their occupation. The authors collected user data from Twitter and LinkedIn using an about.me search, which allows users to connect their accounts across many platforms. The authors argued people's language use varies according to their job. A disadvantage of their approach is its reliance on the availability of users' LinkedIn accounts.

In terms of specific applications, Sloan et al. [9] extract age of UK Twitter users by using pattern matching; Wood-Doughty et al. [11] proposed a neural network model that can infer gender and ethnicity from a user's name and screen name; Vicente et al. [23] predicted gender of Twitter users depending on various information such as user name and screen name, user description, content of the tweets and profile picture; Huang et al. [24] presented an approach for detecting hate speech with extracted age, gender and race/ethnicity from user's profile image, as well as geographic location by numerical location coordinates or matching a regular expression; and Mueller et al. [25] analysed tweets on Me Too Hashtag for sexual abuse and sexual harassment.

Current techniques for personal characteristics extraction can be categorized into three different categories: first, machine learning techniques like Face++ [15], [22], language model [7] and Neural model [11]. Another class of attribute extraction techniques use pattern matching. There are some methods which consider network information as well as website traffic data while deriving personal characteristics of community such as [13], [14]. All these studies tend to focus on a few specific user attributes. They are unable to identify a wider range of attributes using only one method. In contrast, we propose a generic approach, aiming to discover any attribute of interest.

III. Problem Definition

Before explaining the proposed approach, it is useful to define the problem formally first.

A. Person Characteristics

Let $T = (X, U, P)$ be a social network topic, where X is the set of posts at a specific time all about the topic; U is the set of users; P is the set of profiles, one for each user $u \in U$. Let n and k be the sizes of X and U, respectively.

Definition 1 (Community): A community $\mathcal{C}$ is represented by $\mathcal{C} = (U, P)$, where $U = \{u_1, u_2, \ldots, u_n\}$ is a set of users (people) who participate in the community and $P = \{p_1, p_2, \ldots, p_n\}$ is a set of profiles, one associated with each user.

Profiles associated with users are typically written as free text. For our work, we assume that the textual profiles have been converted into term vectors already. That is, a profile p is represented as $\langle t_1, t_2, \ldots, t_k \rangle$, where each t_i is a term (literal) extracted from p.

Definition 2 (Person Characteristic): A person characteristic is an attribute-value pair (A, v), where A is a literal representing a characteristic type that describes a user and v is a literal representing an instance of A.

For example, `Hobby` is a characteristic type and `swimming` is a value of this type. So (`Hobby, swimming`) represents a person characteristic. Note that a person may have multiple values for the same characteristic type. For example, a person may have (`Hobby, swimming`) and ((`Hobby, reading`). As a shorthand, we allow these to be written as a set in a person characteristic: (`Hobby,` {`swimming, reading`}).

B. Problem Definition

In this section, we define the problem of extracting attributes from profiles for a community.

Definition 3 (Attribute Extraction): Given a community $\mathcal{C} = (U, P)$, we want to extract its description

$$D_{\mathcal{C}} = \{(A_1, v_1), (A_2, v_2), \ldots, (T_m, v_m)\}$$

where each (A_i, v_i) is a person characteristic and satisfies the following condition

$$Pr[(A_i, v_i)|P] > \delta \qquad 1 \leq i \leq m$$

where δ is a user-specified threshold.

That is, when a person characteristic (A_i, v_i) appears frequently enough in a set of profiles P, then we include it as part of the description for the community, or the attributes of the community.

C. #Hashtag Communities

The definitions given in the previous sections are general enough for any type of online community, as long as each community is characterised by a set of users and a set of profiles describing the users. In this section, we consider specifically the communities formed around #hashtag in Twitter.

When a group of Twitter users tweet on the same #hashtag, we say that this group of users form a community and we call it a *#hashtag community*, denoted by $\mathcal{C}_{\#hashtag}$, and its size is the set of distinct users tweeted on the #hashtag, denoted by $|\mathcal{C}_{\#hashtag}|$. Table I gives an example of such community.[1].

[1]These four users were extracted from a much larger #hashtag community *#trump2020* which has 154 users.

TABLE I
AN EXAMPLE OF #HASHTAG COMMUNITY

User (U)	Profile (P)
u_1	prolifer Catholic Love Jesus proud mom of Iraq war veteran avid supporter of our military avid animal lover addicted to reruns of Monk he is so funny Mr. Monk.
u_2	Love God Almighty, love neighbor, love myself, Love my spouse, love playing guitars and driving fast muscle cars.
u_3	I had to #WalkAway. Proudly shadowbanned. Love President Trump. Christian. MAGA. America First. http://qmap.pub
u_4	Mother of 6 who I love more than life. Simple I am, and ladylike I try to be. My favorite colors are pink and navy, and pearls are a must. #kindnesswins #maga

From this community, it is easy to see that the words such as `Jesus`, `Catholic` and `Christian` appearing in the profiles of u_1 and u_3 should suggest that they are religious and more specifically have the faith of Christianity. If we set $\delta = 0.4$ (see Definition 3), then we can assign a person characteristic $(Religion, Christianity)$ to this community. Equally, the words such as `Mom` and `Mother of 6` appearing in the profiles of u_1 and u_4 should allow us to infer another person characteristic $(Gender, Female)$ for this community. This will then allow us to derive a description for this community: $D_{\mathcal{C}} = \{(Religion, Christianity), (Gender, Female)\}$

In the previous example, the personal characteristics or description is obtained manually. To infer them automatically from Twitter profiles, we outline two general approaches.

D. Bottom-up Derivation

The bottom-up derivation is a *search* based approach to person characteristic derivation. The main steps of this approach are given in Fig. 1. That is, we start with user profiles, and attempt to extract relevant values of person characteristics and then associate them with the types.

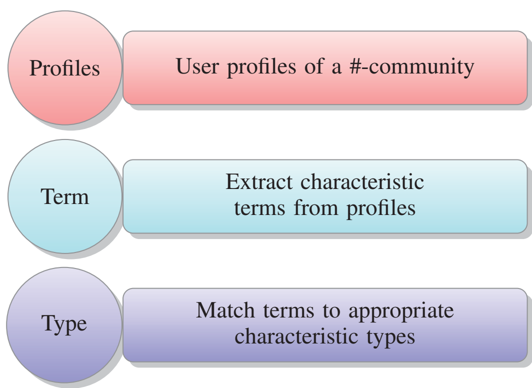

Fig. 1 Bottom-up Derivation

While this approach is desirable in that it is very general, and is not limited by what person characteristics may be extracted, it is not a trivial task and involves the following two major challenges.

- Firstly, determining which term(s) or word(s) occurring in a user profile are meaningful characteristic values. For instance, in our example, it is difficult to decide that terms in the profile of u_1 such as Mom, Catholic and Christian might be relevant whereas addicted and proud might not be useful. One possible solution is to employ some machine learning techniques to identify such terms, but this would require a large amount of annotated training examples, which can be difficult to obtain.
- Secondly, assuming that we are able to obtain a list of meaningful characteristic values, determining their types is hard. For instance, in our example, linking values such as Catholic and Christian to a possible type Religion is difficult. Again, this would require a substantial knowledge base, either constructed as a dictionary or derived from machine learning, which could be difficult to obtain. Moreover, some of the terms may have multiple senses, for example author could indicate a type of occupation or a kind of hobby, and they add further complexity to this approach.

E. Top-down Derivation

The top-down derivation is a *detection* based approach to person characteristic derivation. The main steps of this approach are given in Fig. 2. That is, we start with a characteristic type given by the user, and attempt to detect values in user profiles that are relevant to the given type.

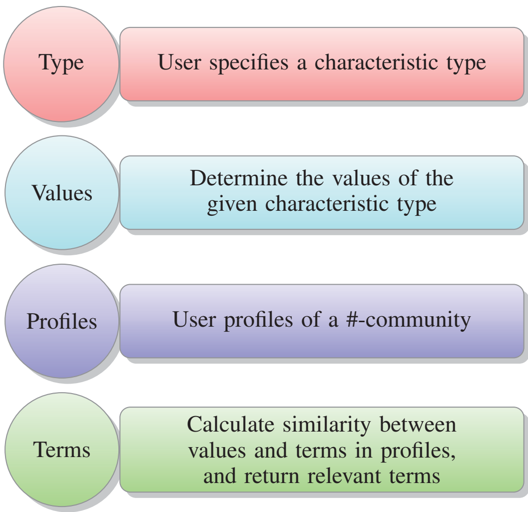

Fig. 2 Top-down Derivation

This approach is clearly less general than the bottom-up approach as it can only detect if a particular person characteristic exists amount the profiles, but is unable to find *any* person characteristic. However, the ability of detecting any given person characteristic type is still useful, as it can help monitor characteristics of a community. For example, if we wish to determine if a #-community has a type Religion, then we can search through all profiles of the users in this community to see if there are any values of Religion.

The top-down approach will not have the difficult tasks of determining relevant values from profiles and then mapping these values to the correct types like the bottom-up approach will have to deal with. But it still has some substantial challenges to address.

- Firstly, we need to have a set of values associated with a given characteristic type, so that we can use them to search through user profiles to see the given type is supported by the community. For example, when a given characteristic type is Religion, we will need the values, such as Muslim, Christian, worship and many more, to search through the user profiles to see if the community can be characterised as religious. One possible solution is to use a dictionary containing possible values for each characteristic type, for example, using Wordnet or ConceptNet. Alternatively, more dynamic semantic tools such as Word Embeddings may be used to determine association of any value to any characteristic type.
- Secondly, we need to consider how to count the support for a given characteristic type effectively and accurately. For instance, in our example, mother and mom are two different but semantically equivalent values of a characteristic type such as parent. It is clearly essential and necessary that we need to combine them when counting the support for this characteristic type. The solution to the first task above will partially address this issue, i.e. mother and mom can be both values of parent, hence will automatically be included in the counting for parent. But Twitter profiles can be written with abbreviations, spelling error, slang etc., which pose additional challenges for determining semantic equivalence between values. Again, techniques such as word embeddings or machine learning could be employed to address this issue.

IV. Proposed Approach

In this paper, we propose a detection based method for extracting attributes from hashtag communities. We explain the proposed method through an example first. Suppose that We wish to search for a characteristic type religion. From Wordnet, we discover possible values of religion as shown in Fig. 3. Then, similarities between the possible values of religion and the words appearing in the profiles are calculated using Word2vec as shown in Fig. 3. We find that words $Catholic$ in first profile and $Christian$ are similar to the possible value $Christianity$. Also, God is similar to $Faith$. So, the attribute-value pair is $(Religion, \{Catholic, Christian, god\})$ (we will discuss WordNet and Word2vec in detail in the next section).

A. Semantic Analysis

The fundamental principle behind this method is to understand a group of users by finding some common characteristics. Thus, we introduce a method for deriving

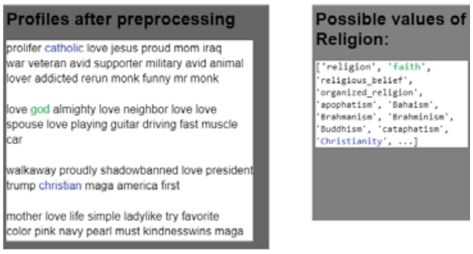

Fig. 3 Example of search religion characteristic in the example in Table I of #hashtag community

personal characteristics of a *#hashtag community* using semantic analysis. The following are definitions necessary to understand this method.

Definition 4 (Possible Values): Possible values

$$P_v = \{p.v_1, p.v_2, \ldots, p.v_f\}$$

are the instances of a characteristic type A, where $p.v_i$ is a possible value used to find v in a person characteristic (A_i, v_i).

Definition 5 (Relevant Terms): Relevant term is a term r_i extracted from the profile p, and is similar to one or more possible values. It is possible to have one or more relevant terms for the same characteristic type. This term is represented as v in a person characteristic (A_i, v_i).

Definition 6 (MSDA): MSDA denotes Minimum Similarity Distance Accepted. Given possible values $P_v = \{p.v_1, p.v_2, \ldots, p.v_f\}$ of a characteristic type giving A and profiles, the measure of similarity distance between one possible value and a word in profile can be given by Word2vec [26], one of word embedding techniques. Where word in profile is accepted as value of v in a person characteristic (A_i, v_i), if the similarity is equal to or more than α, where α is a user-specified threshold.

The proposed method has two main steps as shown in Fig. 4.

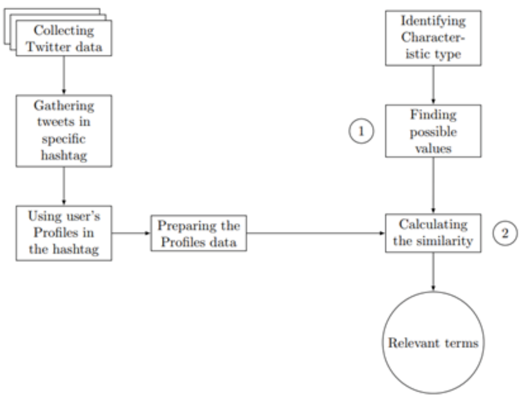

Fig. 4 The proposed method, with two key steps

1. Possible Values of the characteristic type: The method searches for the attribute of interest to the user, this is what distinguish our work than others. In the first step, we try to find all possible value of the attribute. To obtain these values, WordNet [27] is used. Like a traditional dictionary, WordNet provides definitions of terms as well as their relationships. WordNet, however, differs from a normal dictionary in that it is organised conceptually rather than alphabetically. A synonym set, or synset, is the basic unit in WordNet, and it represents a lexicalized concept. Synset instances are a collection of synonyms that communicate the same idea. Some words have only one Synset, while others have multiple. Words in WordNet are organised in a hierarchical tree structure on the basis of hypernym/hyponymy. Hypernym (more broad terms) and hyponym (more specific terms) are semantic relations between Synsets that are transitive. The semantics of concepts in the upper layers of the hierarchy are more generic, with less resemblance between them, whereas concepts at the lower layers or within the same layer are more concrete, with higher similarity [28].

In reality, along WordNet, we can recursively acquire a word's hypernyms relations up to the top root and hyponyms relations down to the bottom. Gong et al. [29] studied ways to improve Internet searches by using hypernyms and hyponyms in WordNet to broaden the query. They investigated how many levels along hypernyms and hyponyms to broaden the inquiries is interesting. In both circumstances, their results suggest that one level of hypernyms and hyponyms yields better results. Multiple hyponyms relations, in reality, stretch too many words, causing the basic keyword meanings to diverge. That is, they will introduce a lot of noise into the results, lowering the average precision values.

The pseudo-code given in Algorithm 1 is used to find possible values. Given a characteristic type, the task first finds the synonyms of it, for example, if the characteristic type is occupation, then the synonyms would be job, career, profession . . . etc. Then, the task finds hyponyms for each synonym. Hyponyms give abstract concepts of the word that are much more specific. For example, hyponyms of occupation would be teacher, farming, trade . . . etc.

Algorithm 1 Finding possible values of attribute

Requirements: a characteristic type "A"
Results : Possible values of A
Method : Using WordNet

1) Look up a Attribute using synsets() function, (result is a set of synsets (synonyms) that can all refer to the same concept)
2) Find lemma for each synsets (Each synset contains one or more lemmas, which represent a specific sense of a specific word).
3) Find hyponyms for each synsets (result is words that is more specific than a given word)
4) Find lemma for each hyponyms
5) Return lemma of synsets and hyponyms

2. Attribute-value pair (A, v): This step finds a value v of characteristic type A. Simplifying the search in the profile by using possible values would miss the majority of existing attribute terms. People are likely to use terms that are similar to possible values rather than being identical to them. As a result, to identify v, the distance between possible values and the words in profiles needs to be measured. The similarity is determined based on Word2vec technique [26]. Word2Vec is one of the most popular technique for natural language processing to learn word embeddings. It learns word associations from a large text by using a neural network model. We used a pre-trained word2vec model, which is trained from Google News. The model includes word vectors for 3 million words and phrases, which were trained on around 100 billion words.

V. Experimental and Results

A. Experiment Setup

The steps taken for data collection and preparation are summarised in Fig. 5. We used Twitters streaming API to search for tweets containing a specific symbol (#) and to filter non-English tweets. We downloaded tweet content and their hashtags, author name, location and profile.

To extract useful information, we start by pre-processing the description using the following steps: we first remove emails, new line characters, single and doubles quotes, links, commas, full stops, and punctuation; we then tokenise the descriptions with genism; following that, we remove stop words with NLTK; lemmatisation with spacy; and convert all uppercase letters to lowercase letters.

One level of hyponyms method in Wordnet was used in this experiment, as discussed in the previous section, since one level produced better results. Different MSDA were evaluated to determine how similarity between a possible value and the profile word is to be measured, and we used 0.45 in our experiments.

All the experiments are executed on a computer with 8GB main memory and an Intel(R) Core(TM) i5-8265U CPU @1.80 GHz running an Windows 10 operating system. Python 3.8.8 is used to implement the algorithms.

B. Results of Human Validation

To establish how well a particular method can recognise attributes from Twitter users' profile description fields, human validation is needed. Also, there is a lack of benchmarking datasets to compare and evaluate the performance of different approaches. We therefore evaluate the performance of our attribute extracting technique against human judgment, and we use the F-measure in our study, which is commonly used to evaluate the performance of attribute extraction techniques [9], [13].

To conduct this research, we use human judgment to provide ground truth for these person characteristics (Belief, Religion), (Religion, Christianity), (Gender, Female) and (Gender, Male). We then used our proposed approach to automatically find these person characteristics. Table II shows the F-measure values for deriving several person characteristics over a range of hashtag topics. As can be seen, the proposed method worked quite well with the exception of (Gender, Male).

TABLE II
F1 Results for Twitter User Attributes Extraction on Manually Annotated Data

Person Characteristics	Number of profiles	F1-measure
(Belief, Religion)	116	83.78%
(Religion, Christianity)	130	97.96%
(Gender, Female)	116	75.95%
(Gender, Male)	116	38.71 %

C. Results of Attributes Extraction

Our experiments are split into two parts. The first section comprises experiments on smaller communities with less than 1000 users, while the second section contains experiments on two large communities with 4650 and 5992 users. The results are given Table III, which shows the percentage of person characteristics (Belief, Religion), (Religion, Christianity), (Gender, Female), (Gender, Male) for each hashtag community. The results suggest that the method proposed in this paper can be used to discover attributes for communities by issuing queries such as "find a community that has at least 30% participants who are religious".

It is worth noting that as our method is designed to extract any attribute from user profiles, we observed that in our experiments it does not work equally well for all attributes. For example, when we searched for republican and democrat as an attribute, because we have only utilised their general English meaning, which is quite different when they are used to refer to respective political parties, our extraction based on lexical sources such as Wordnet appears to be limited, and much future work is needed to take this work forward.

VI. Conclusions

The user characteristics associated with a particular topic may provide a greater understanding of that topic. In this paper, we proposed a method for obtaining social network user attributes from their profiles. The experimental results have shown that user profiles are a promising source for attribute extraction, as it helps reduce the amount of labelled data required to derive attributes from data.

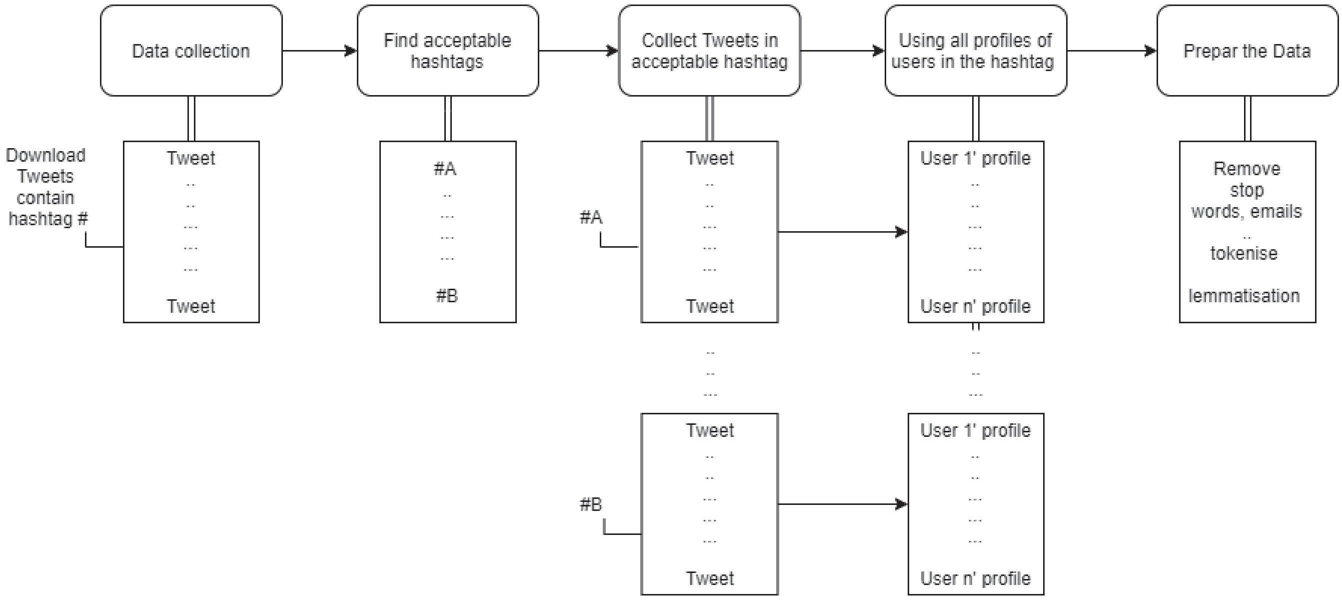

Fig. 5 Data collection and preparation

TABLE III
RESULTS

(a) EXPERIMENTS ON SMALL COMMUNITIES

Hashtag	Number of users	(Belief, Religion)	(Religion, Christianity)	(Gender, Female)	(Gender, Male)
Humantrafficking	193	31.09%	14.51%	20.73%	23.83%
qanon	962	29.94%	15.80%	19.54%	24.01%
Strengthinunity	47	25.53%	8.51%	25.53%	29.79%
Thegreatawakening	94	31.91%	14.89%	21.27%	24.46%
trump2020	116	31.03%	17.24%	24.14%	26.72%
qdrop	98	42.86%	22.45%	22.45%	29.59%
Maga	365	31.51%	16.71%	20.55%	25.75%
Pureevil	64	29.69%	9.38%	9.38%	15.63%
netflix	434	5.30%	2.30%	15.44%	21.89%
rachelchandler	713	35.20%	18.79%	23.14%	28.33%
Greatawakening	99	25.25%	15.15%	21.21%	23.23%
anons	140	24.29%	12.86%	21.43%	23.57%
Magaveteran	34	23.53%	11.76%	29.41%	35.29%
Mariajohnsen	41	4.88%	2.44%	2.44%	0.00%
Muellerreport	100	29%	15%	24%	28%
WWG1WGA	614	32.41%	17.26%	20.84%	24.59%

(b) EXPERIMENTS ON TWO LARGE COMMUNITIES

Hashtag	Number of users	(Belief, Religion)	(Religion, Christianity)	(Gender, Female)	(Gender, Male)
Brexit	4650	7.27%	2.26%	5.46%	10.26%
coronavirus	5992	5.72%	0.82%	4.09%	7.76%

REFERENCES

[1] K. H. Lim and A. Datta, "Finding twitter communities with common interests using following links of celebrities," in *Proceedings of the 3rd international workshop on Modeling social media*, 2012, pp. 25–32.

[2] ——, "Following the follower: Detecting communities with common interests on twitter," in *Proceedings of the 23rd ACM conference on Hypertext and social media*, 2012, pp. 317–318.

[3] M. Bakillah, R.-Y. Li, and S. H. Liang, "Geo-located community detection in twitter with enhanced fast-greedy optimization of modularity: the case study of typhoon haiyan," *International Journal of Geographical Information Science*, vol. 29, no. 2, pp. 258–279, 2015.

[4] B. R. Amor, S. I. Vuik, R. Callahan, A. Darzi, S. N. Yaliraki, and M. Barahona, "Community detection and role identification in directed networks: understanding the twitter network of the care. data debate," in *Dynamic networks and cyber-security*. World Scientific, 2016, pp. 111–136.

[5] N. Cao, L. Lu, Y.-R. Lin, F. Wang, and Z. Wen, "Socialhelix: visual analysis of sentiment divergence in social media," *Journal of visualization*, vol. 18, no. 2, pp. 221–235, 2015.

[6] P. Vijayaraghavan, S. Vosoughi, and D. Roy, "Twitter demographic classification using deep multi-modal multi-task learning," in *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, 2017, pp. 478–483.

[7] H. A. Schwartz, J. C. Eichstaedt, M. L. Kern, L. Dziurzynski, S. M. Ramones, M. Agrawal, A. Shah, M. Kosinski, D. Stillwell, M. E. Seligman *et al.*, "Personality, gender, and age in the language of social media: The open-vocabulary approach," *PloS one*, vol. 8, no. 9, p. e73791, 2013.

[8] T. Yo and K. Sasahara, "Inference of personal attributes from tweets using machine learning," in *2017 IEEE International Conference on Big Data (Big Data)*. IEEE, 2017, pp. 3168–3174.
[9] L. Sloan, J. Morgan, P. Burnap, and M. Williams, "Who tweets? deriving the demographic characteristics of age, occupation and social class from twitter user meta-data," *PloS one*, vol. 10, no. 3, p. e0115545, 2015.
[10] T. Hu, H. Xiao, J. Luo, and T.-v. T. Nguyen, "What the language you tweet says about your occupation," in *Tenth International AAAI Conference on Web and Social Media*, 2016.
[11] Z. Wood-Doughty, N. Andrews, R. Marvin, and M. Dredze, "Predicting twitter user demographics from names alone," in *Proceedings of the Second Workshop on Computational Modeling of People's Opinions, Personality, and Emotions in Social Media*, 2018, pp. 105–111.
[12] T. Georgiou, A. El Abbadi, and X. Yan, "Extracting topics with focused communities for social content recommendation," in *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*. ACM, 2017, pp. 1432–1443.
[13] A. Culotta, N. R. Kumar, and J. Cutler, "Predicting the demographics of twitter users from website traffic data," in *Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015.
[14] S. Volkova, Y. Bachrach, and B. Van Durme, "Mining user interests to predict perceived psycho-demographic traits on twitter," in *2016 IEEE Second International Conference on Big Data Computing Service and Applications (BigDataService)*. IEEE, 2016, pp. 36–43.
[15] J. Messias, P. Vikatos, and F. Benevenuto, "White, man, and highly followed: Gender and race inequalities in twitter," in *Proceedings of the International Conference on Web Intelligence*. ACM, 2017, pp. 266–274.
[16] N. Cesare, C. Grant, and E. O. Nsoesie, "Detection of user demographics on social media: A review of methods and recommendations for best practices," *arXiv preprint arXiv:1702.01807*, 2017.
[17] M. E. Newman and M. Girvan, "Finding and evaluating community structure in networks," *Physical review E*, vol. 69, no. 2, p. 026113, 2004.
[18] S. Fortunato, "Community detection in graphs," *Physics reports*, vol. 486, no. 3-5, pp. 75–174, 2010.
[19] Z. Xu, Y. Ke, Y. Wang, H. Cheng, and J. Cheng, "A model-based approach to attributed graph clustering," in *Proceedings of the 2012 ACM SIGMOD international conference on management of data*. ACM, 2012, pp. 505–516.
[20] Y. Ruan, D. Fuhry, and S. Parthasarathy, "Efficient community detection in large networks using content and links," in *Proceedings of the 22nd international conference on World Wide Web*. ACM, 2013, pp. 1089–1098.
[21] J. Pan, R. Bhardwaj, W. Lu, H. L. Chieu, X. Pan, and N. Y. Puay, "Twitter homophily: Network based prediction of user's occupation," in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 2019, pp. 2633–2638.
[22] A. Chakraborty, J. Messias, F. Benevenuto, S. Ghosh, N. Ganguly, and K. P. Gummadi, "Who makes trends? understanding demographic biases in crowdsourced recommendations," in *Eleventh International AAAI Conference on Web and Social Media*, 2017.
[23] M. Vicente, F. Batista, and J. P. Carvalho, "Gender detection of twitter users based on multiple information sources," in *Interactions Between Computational Intelligence and Mathematics Part 2*. Springer, 2019, pp. 39–54.
[24] X. Huang, L. Xing, F. Dernoncourt, and M. J. Paul, "Multilingual twitter corpus and baselines for evaluating demographic bias in hate speech recognition," *arXiv preprint arXiv:2002.10361*, 2020.
[25] A. Mueller, Z. Wood-Doughty, S. Amir, M. Dredze, and A. L. Nobles, "Demographic representation and collective storytelling in the me too twitter hashtag activism movement," *arXiv preprint arXiv:2010.06472*, 2020.
[26] T. Mikolov, K. Chen, G. Corrado, and J. Dean, "Efficient estimation of word representations in vector space," *arXiv preprint arXiv:1301.3781*, 2013.
[27] G. A. Miller, "Wordnet: a lexical database for english," *Communications of the ACM*, vol. 38, no. 11, pp. 39–41, 1995.
[28] Y. Li, Z. A. Bandar, and D. McLean, "An approach for measuring semantic similarity between words using multiple information sources," *IEEE Transactions on knowledge and data engineering*, vol. 15, no. 4, pp. 871–882, 2003.
[29] Z. Gong, M. Muyeba, and J. Guo, "Business information query expansion through semantic network," *Enterprise Information Systems*, vol. 4, no. 1, pp. 1–22, 2010.