

Identifying Topical Twitter Communities via User List Aggregation

Derek Greene, Derek O’Callaghan, Pádraig Cunningham

School of Computer Science & Informatics, University College Dublin
{derek.greene,derek.ocallaghan,padraig.cunningham}@ucd.ie

Abstract. A particular challenge in the area of social media analysis is how to find communities within a larger network of social interactions. Here a community may be a group of microblogging users who post content on a coherent topic, or who are associated with a specific event or news story. Twitter provides the ability to curate users into lists, corresponding to meaningful topics or themes. Here we describe an approach for crowdsourcing the list building efforts of many different Twitter users, in order to identify topical communities. This approach involves the use of ensemble community finding to produce stable groupings of user lists, and by extension, individual Twitter users. We examine this approach in the context of a case study surrounding the detection of communities on Twitter relating to the London 2012 Olympics.

1 Introduction

A wide variety of community finding techniques have been proposed in the literature, with recent research focusing on the challenge of identifying overlapping communities [1]. In the case of microblogging data, researchers have been interested in the identification of communities of users on Twitter, who produce tweets on a common topic, who belong to the same demographic, or who share a common ideological viewpoint [2]. These approaches have generally relied on explicit views of the Twitter network, such as follower relations or retweets.

Twitter users can organise the accounts that they follow into Twitter *user lists*. These lists are used in a variety of ways. In some cases they may correspond to personal lists of a given user’s friends and families, but frequently lists are employed to group together Twitter accounts based on a common topic or theme. In this way, every Twitter user can effectively become a community curator. Notably, journalists from news organisations such as The Telegraph and Storyful curate lists relevant to a given news story or event, as a means of monitoring breaking news. Recently, Kim *et al.* and García-Silva *et al.* [3, 4] both discussed the potential of user lists to provide latent annotations for Twitter user profiles.

Our primary goal here is to demonstrate that topical communities can be identified by harnessing the “crowd-sourced” list building efforts of a large base of Twitter users. In Section 3, we show that this can be done by constructing a graph based on the similarity of user list memberships, and then using an *ensemble community finding* approach to find robust, overlapping groups of lists

within this graph, from which user communities can be derived. We use stability information derived from the ensemble as a proxy for the reliability of the communities. In Section 4, we evaluate the proposed techniques on a case study relating to coverage of the London 2012 Olympics on Twitter.

2 Related Work

Many researchers have become interested in exploring the network structure within the Twitter network, given the potential for Twitter to facilitate the rapid spread of information. Java *et al.* [2] provided an initial analysis of the early growth of the network, and also performed a small-scale evaluation that indicated the presence of distinct Twitter user communities. Kwak *et al.* [5] performed an evaluation based on a sample of 41.7m users, studying aspects of the network such as: identifying influential users, information diffusion, and trending topics. Typically researchers have focused either on Twitter users from the perspective of the content that they produce, or in terms of explicit network representations based on follower relations or retweeting activity [5]. However, preliminary work by Kim *et al.* [3] suggested that latent groups and relations in Twitter data could be extracted by examining user list data. Wu *et al.* [6] suggested that user list memberships could be used to organise users into a pre-defined set of categories: celebrities, media, organisations, and blogs. García-Silva *et al.* [4] described approaches for extracting semantic relations from user lists, by constructing relations between co-occurring keywords taken from list names.

Many algorithms have been proposed to identify communities in graphs, based on different combinations of objective functions and search strategies [1]. A widely-employed algorithm in this area is OSLOM (Order Statistics Local Optimization Method), introduced by Lancichinetti *et al.* [7]. Kwak *et al.* [8] observed that many community detection algorithms can produce inconsistent results, due to stochastic elements in their optimisation process. Lancichinetti & Fortunato [9] demonstrated that this also applied to OSLOM, and proposed an ensemble approach to generate stable results from a set of partitions. In the more general cluster analysis literature, *ensemble clustering* methods have been developed to address similar issues. These typically involve generating a diverse set of “base clusterings”, which are aggregated to produce a consensus solution [10, 11].

3 Methods

In this section, we introduce an approach that aggregates user list information to generate communities. Firstly, we describe the construction of a graph representation of user lists, based on their membership overlaps. Then in Section 3.2 we describe an ensemble approach to identify overlapping groups of user lists. The stability of these groups is assessed as described in Section 3.3, and the selection of community labels is discussed in Section 3.4. Finally, the derivation of corresponding communities for individual users is discussed in Section 3.5.

3.1 User List Graph Construction

We construct a graph G of l nodes, where each node represents a distinct Twitter user list L_x . A weighted edge exists between a pair of lists if they share users in common. Rather than using the raw intersection size between a pair, we make allowance for the significance of the intersection size relative to the size of the two lists, and the total number of users assigned to lists n . For a pair (L_x, L_y) , we compute a p -value to indicate the significance of the probability of observing at least $|L_x \cap L_y|$ users from L_x within another list of size $|L_y|$:

$$PV(L_x, L_y) = 1 - \sum_{j=0}^{|L_x \cap L_y|-1} \frac{\binom{|L_x|}{j} \binom{n-|L_x|}{|L_y|-j}}{\binom{n}{|L_y|}} \quad (1)$$

To improve interpretability, we compute the associated log p -value:

$$LPV(L_x, L_y) = -\log(PV(L_x, L_y)) \quad (2)$$

where a larger value is more significant. We consider Eqn. 2 as a measure of the similarity between a pair of user lists, corrected for chance. To further increase the sparseness of the graph, we remove edges with weights $LPV < \rho$ for a weight threshold ρ . Increasing the value of ρ will result in an increasingly sparse graph.

3.2 Combining Overlapping Communities

We will naturally expect that different topical communities will potentially overlap with one another. To identify communities of lists, we apply the OSLOM algorithm which has been shown to out-perform other community finding approaches [7]. However, as noted in [9], OSLOM can produce unstable results.

Following the CSPA ensemble aggregation approach [10], and the method for combining network partitions [9], we now describe an approach for generating and combining an ensemble of overlapping community sets. Given an initial user list graph G , we construct a symmetric $l \times l$ *consensus matrix* \mathbf{M} . For the purpose of generating a collection of r *base community sets*, we apply the OSLOM algorithm [7] using a different initial random seed for each run. Motivated by the notion an ensemble of weak clusterings [11], we use the “fast” configuration of OSLOM, which uses a minimal number of optimisation iterations.

After generating a base community set, for each unique pair of nodes (L_x, L_y) in network, we compute the Jaccard similarity between the sets of community labels assigned to those nodes by OSLOM. If the pair are not both co-assigned to any community, the score is 0. If the pair are present in all communities together, the score is 1. However, unlike the binary approach of [9], if the pair are present in some but not all communities together, the Jaccard score will reflect this. See Fig. 1 for examples. In the case of non-overlapping partitions, the score will reduce to the binary scoring used in [9]. After computing all Jaccard scores, we increment the corresponding matrix entries in \mathbf{M} .

Once all r base community sets have been generated, \mathbf{M} is normalised by $1/r$ to give a matrix with entries $\in [0, 1]$. To find the *consensus communities*,

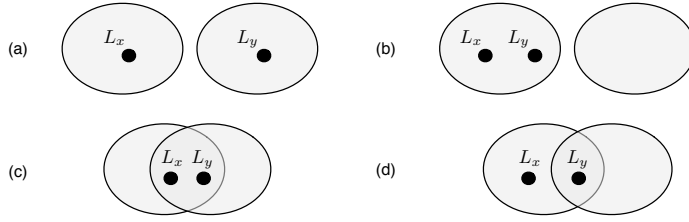


Fig. 1. Four different cases of computing the Jaccard similarity between the sets of community labels assigned to two nodes L_x and L_y . The Jaccard scores respectively are: (a) 0.0, (b) 1.0, (c) 1.0, (d) 0.5.

we follow a similar approach to that used by [9]. We construct a new undirected weighted graph such that, for every unique pair of nodes (L_x, L_y) , we create an edge with weight M_{xy} if $M_{xy} \geq \tau$. The threshold parameter $\tau \in [0, 1]$ controls the sparsity of the graph. We then apply OSLOM to this graph to produce a final grouping of user lists.

3.3 Evaluating Community Stability

When applying community detection, often we may wish to examine the most reliable or robust communities with the strongest signals in the network. Here, we rank the consensus communities generated as described in Section 3.2, based on the cohesion of their members with respect to the consensus matrix \mathbf{M} . A more stable consensus community will consist of user lists which were frequently co-assigned to one or more communities across all r base community sets.

For a given consensus community C of size c , we compute the mean of the values M_{xy} for all unique pairs (L_x, L_y) assigned to C ; this value has the range $\in [0, 1]$. We then compute the mean expected value for a community of size c as follows: randomly select c unique nodes from G , and compute their mean pairwise score from the corresponding entries in \mathbf{M} . This process is repeated over a large number of randomised runs, yielding an approximation of the expected stability value. We then employ the widely-used adjustment technique introduced by [12] to correct stability for chance agreement:

$$\text{CorrectedStability}(C) = \frac{\text{Stability}(C) - \text{ExpectedStability}(C)}{1 - \text{ExpectedStability}(C)} \quad (3)$$

A value close to 1 indicates a highly-stable community, while a value closer to 0 is indicative of a weak community that appeared intermittently over the r runs. We rank all consensus communities based on their values for Eqn. 3.

3.4 Selecting Community Labels

To summarise the content of a consensus community, we aggregate the meta-information associated with all lists assigned to that community. Specifically, we

construct a bag-of-words model, where each user list is represented by unigrams and bigrams tokenised from the list’s name and description. Single stop-words are removed, and terms are weighted using log-based TF-IDF. For each community, we then compute the centroid vector corresponding to the mean vector of all lists assigned to that community. To generate descriptive labels for the community, we subtract the mean vector of **all user list vectors** from the community centroid vector, and rank terms in descending order based on the resulting weights. The top ranked terms are used as community labels.

3.5 Deriving User Memberships

The consensus communities generated using the method proposed in Section 3.2 can potentially provide us with an insight into the overall topics in a Twitter corpus. However, it will often be useful to assign community memberships to individual users. We can readily produce this by using the list groupings in conjunction with the original user list membership information.

For a given consensus community C of size c , we examine the memberships of all lists in the community. We consider the assignment of a user u_i to each $L_x \in C$ as being a vote with weight $1/c$ for u_i belonging to the overall community. The total membership weight for u_i is therefore given by the fraction of lists $L_x \in C$ containing u_i . Membership weights for all communities are computed in this way. We can also rank the importance of users in a given community by sorting users by weight in descending order. To produce a final set of user communities, we only include a user in a community for which the user has a membership weight $\geq \mu$, based on a membership threshold $\mu \in [0, 1]$.

4 Evaluation

4.1 Data Collection

To evaluate the proposed community finding methods, we constructed a dataset based on a list of 499 users curated by The Telegraph, which covers athletes, journalists, and organisations involved in the London 2012 Summer Olympics¹. Initially, for each user we retrieved up to their 200 most recent user list assignments. From this initial pool of lists, we then retrieved list memberships for 10,000 randomly selected lists of size ≥ 5 and containing at least 2 core list users. This yielded a dataset containing a total of 44,484 individual list membership records, where the average number of lists per user was 89. The most frequently-listed user was *@andy-murray*, assigned to 1,931 different lists.

4.2 Community Detection

Using the approach described in Section 3.1, we constructed a user list graph based on membership information for the 499 users. To limit the density of the

¹ <http://twitter.com/#!/Telegraph2012/london2012>

graph, we use a weight significance threshold of $\rho = 6$ (*i.e.* user list overlaps are considered as significant for $LP \leq 1^{-6}$). This resulted in a graph containing 4,948 nodes representing user lists, with 749,062 weighted edges between them.

To generate an ensemble of base community sets, we apply OSLOM as described in Section 3.2 for 100 random runs, selecting the lowest level of the hierarchy as the solution for each run. The average number of non-singleton communities in each run was 157. Combining the base community sets yielded a consensus matrix containing ≈ 11.8 m non-zero values. We examined a range of threshold values $\tau \in [0.1, 0.5]$, and selected a threshold $\tau = 0.2$ to generate consensus communities in order to maximise coverage over user lists, while also reducing the density of the consensus graph. Applying OSLOM to the sparse graph of ≈ 1.5 m values produced a total of 94 consensus communities, considerably lower than the average base community count. Finally, user communities were derived using a low membership threshold $\mu = 0.1$ to maximise the number of core users assigned to communities. In total, 416 core users were assigned to at least one community, with 362 users assigned to multiple communities.

Table 1 shows the top 15 communities, arranged in descending order by their stability score, as defined in Eqn. 3. The table shows the size of each community (in terms of both number of lists and users assigned), the top text labels selected for each community, and the three highest-weighted users. We observe that the most stable communities generally correspond to communities of users involved in specific, “niche” sports (*e.g.* badminton, BMX racing, fencing). In these cases, the top-weighted users correspond to either British Olympic athletes competing in these sports, or accounts of the official British organisations

Table 1. Top 15 user list communities, arranged in descending order by stability score.

Score	Lists	Users	Top Labels	Top Users
1.00	17	14	badminton, badminton players, badders	@Jennywallwork, @Nath_Robertson, @ChrisAdcock1
1.00	5	5	bmw, bmx racing, bmx atl�ti	@ShanazeReade, @liamPHILLIPS65, @bloomy181
1.00	32	11	sailing, sailors, olympic	@SkandiaTeamGBR, @AinslieBen, @matchracegirls
1.00	19	24	fencing, fencers, individuelle schermers	@britishfencing, @CBennettGBR, @LaurenceHalsted
1.00	6	5	triathlon, machines, swim run	@AliBrownleetri, @jodieswallow, @MarkCavendish
1.00	5	21	scots, red sky, 2014	@mj88live, @RobbieRenwick, @Euan_Burton
1.00	22	4	wielrennen, ciclismo, cycling	@GeraintThomas86, @UCLcycling, @MarkCavendish
0.98	5	7	track, field, track field	@allysonfelix, @TysonLGay, @tiffofili
0.98	48	19	rowing, rowers, gb rowing	@andrewthodge, @ZacPurchase, @MarkHunterGB
0.97	14	22	diving, tuffi, olympic diving	@PeterWaterfield, @matthew_mitcham, @toniacouch
0.96	36	44	hockey, hockey players, field hockey	@AlexDanson15, @RichM6, @jfair25
0.96	12	21	canoe, canoeing, canoe slalom	@GBcanoeing, @PlanetCanoe, @edmckeeper
0.93	5	5	actors athletes, internet stars, athletes tmz	@usainbolt, @ShawnJohnson, @MichaelPhelps
0.92	27	13	judo, judo clubs, judo related	@BritishJudo, @USAJudo, @IntJudoFed
0.87	6	7	runners, hardlopen, runners world	@Mo_Farah, @paulajradcliffe, @KenenisaBekele

Table 2. Validation scores achieved relative to 18 “ground truth” categories.

Category Name	Category Size	Precision	Recall	F1
<i>judo</i>	20	1.00	0.65	0.79
<i>basketball</i>	26	1.00	0.50	0.67
<i>rowing</i>	44	1.00	0.43	0.60
<i>athletics</i>	50	1.00	0.22	0.36
<i>cycling</i>	28	1.00	0.14	0.25
<i>hockey</i>	47	0.98	0.91	0.95
<i>diving</i>	23	0.95	0.91	0.93
<i>equestrianism</i>	18	0.94	0.83	0.88
<i>fencing</i>	23	0.88	0.91	0.89
<i>sailing</i>	16	0.82	0.56	0.67
<i>gymnastics</i>	24	0.77	0.42	0.54
<i>canoeing</i>	22	0.76	0.73	0.74
<i>beach-volleyball</i>	12	0.55	1.00	0.71
<i>boxing</i>	22	0.55	0.55	0.55
<i>swimming-synrho</i>	16	0.33	0.13	0.18
<i>weightlifting</i>	6	0.20	0.17	0.18
<i>archery</i>	17	0.20	0.06	0.09
<i>waterpolo</i>	22	0.05	0.05	0.05

for these sports. Interestingly, we also see some unexpected communities with high stability - a community around the Glasgow 2014 Commonwealth Games, and a community of celebrities which includes “elite” user accounts with hundreds of thousands of followers (*e.g.* *@usainbolt*, *@MichaelPhelps*). As stability decreases, we observed that communities become less homogeneous, covering highly-popular sports (*e.g.* football, basketball), or containing users and lists related to several sports. This suggests that the proposed stability provides a useful measure of the homogeneity of topical content for Twitter communities. Many of the top labels selected for communities are multi-lingual. For instance, the label for the “cycling” community in Table 1 contains terms in Dutch, Italian, and English. Unlike in certain textual analyses of tweets, the use of list membership information allows us to identify groups of users in a language-agnostic manner.

4.3 External Validation

To validate the consensus user communities that were identified by aggregating list information, we use a set of fine-grained Olympics lists also produced by The Telegraph², consisting of Twitter users associated with individual sports (*e.g.* “archery”, “equestrianism”). This provided us with 18 external “ground truth” categories, covering 423 of the 499 users in the dataset.

We computed *precision*, *recall*, and *F1* scores for all communities, and subsequently matched categories to communities based on precision. Table 2 shows the resulting scores for all categories, arranged in descending order by precision. Communities produced by user list aggregation allowed us to identify eight categories with precision ≥ 0.9 , while generally maintaining high recall. Only in the case of four categories did the proposed approach lead to precision and recall scores of both ≤ 0.5 . Subsequent examination of the data suggests that list information was relatively sparse for these categories, and that the users were generally assigned to more generic lists (*e.g.* “aquatics” for “waterpolo”).

² <http://twitter.com/#!/Telegraph2012/lists>

5 Conclusions

In this paper, we have presented initial work on the idea of identifying topical communities on Twitter by aggregating the “wisdom of the crowds”, as encoded in the form of user lists. We show that this information can be mined to identify and label coherent overlapping clusters of both lists and users. We intend to expand this study by examining the identification of other types of user groups within Twitter (*e.g.* academic communities), and by comparing these groups with those mined from other views of the data (*e.g.* follower networks, tweet content).

While the evaluation in this paper used a fixed network of users, a similar approach could be applied to identify topical sub-communities around trending terms or hashtags. Also, in some cases a user will not have been assigned to any user lists. We suggest that a classification process, using an alternative network view (*e.g.* follower links) could be used to assign such users to communities.

Acknowledgements. This research was supported by Science Foundation Ireland Grant 08/SRC/I1407 (Cliques: Graph and Network Analysis Cluster).

References

1. Fortunato, S.: Community detection in graphs. *Phy. Rep.* **486**(3-5) (2010) 75–174
2. Java, A., Song, X., Finin, T., Tseng, B.: Why we Twitter: understanding microblogging usage and communities. In: *Proc. 9th WebKDD and 1st SNA-KDD Workshop on Web Mining and Social Network Analysis.* (2007) 56–65
3. Kim, D., Jo, Y., Moon, I.C., Oh, A.: Analysis of Twitter lists as a potential source for discovering latent characteristics of users. In: *Workshop on Microblogging at the ACM Conference on Human Factors in Computer Systems (CHI’10).* (2010)
4. García-Silva, A., Kang, J., Lerman, K., Corcho, O.: Characterising emergent semantics in Twitter lists. *The Semantic Web: Research and Applications* (2012) 530–544
5. Kwak, H., Lee, C., Park, H., Moon, S.: What is Twitter, a social network or a news media? In: *Proc. 19th Int. Conf. on World Wide Web.* (2010) 591–600
6. Wu, S., Hofman, J., Mason, W., Watts, D.: Who says what to whom on Twitter. In: *Proc. 20th International Conference on World Wide Web, ACM* (2011) 705–714
7. Lancichinetti, A., Radicchi, F., Ramasco, J., Fortunato, S., Ben-Jacob, E.: Finding statistically significant communities in networks. *PLoS ONE* **6**(4) (2011) e18961
8. Kwak, H., Eom, Y.H., Choi, Y., Jeong, H., Moon, S.: Consistent Community Identification in Complex Networks. *ArXiv e-prints* (October 2009)
9. Lancichinetti, A., Fortunato, S.: Consensus clustering in complex networks. *Sci. Rep.* **2** (03 2012)
10. Strehl, A., Ghosh, J.: Cluster ensembles - a knowledge reuse framework for combining partitionings. In: *Proc. Conference on Artificial Intelligence (AAAI’02), AAAI/MIT Press* (2002) 93–98
11. Topchy, A., Jain, A., Punch, W.: Combining multiple weak clusterings. In: *Proc. 3rd IEEE International Conference on Data Mining (ICDM’03).* (2003) 331–338
12. Hubert, L., Arabie, P.: Comparing partitions. *J. Classification* (1985) 193–218