# A retweet network analysis of the European Parliament

Darko Cherepnalkoski, Igor Mozetič

*Jožef Stefan Institute, Jamova 39, 1000 Ljubljana, Slovenia*

*Email: {darko.cerepnalkoski, igor.mozetic}@ijs.si*

*Abstract*—**Social media analytics is increasingly used to uncover underlying real-world phenomena. The goal of this paper is to evaluate the role of Twitter in identifying communities of influence when the "ground truth" is known. We consider the European Parliament (EP) Twitter users during a period of eight months, in which they posted over 370,000 tweets. We define influence as a retweet relation between two Twitter users. We construct two networks of influence: (i) core, where both users are EP members, and (ii) extended, where one user can be outside the EP. We detect communities in both networks and compare them to the "ground truth": the political group and country of the EP members. The results show that the core network closely matches the political orientation, while the extended network reflects the country of origin. This provides empirical evidence that the formation of retweet networks and community detection are appropriate tools to reveal the actual relationships, and can therefore be used to uncover hidden properties when the "ground truth" is not known.**

*Keywords*-**European Parliament; retweet networks; community detection; networks of influence; influence spreading**

## I. INTRODUCTION

The growth of social media and user-generated contents on the web is a potentially relevant and rich source of data. This work is based on data from Twitter[1], a social networking and micro blogging platform with over 300 million monthly active users, posting over 500 million tweets per day.

There are at least two approaches to analyzing Twitter data: the (social) network analysis, and the contents analysis. In our previous research [1] we have combined both approaches. We have detected influential communities, identified discussion topics and assigned sentiment of the communities towards selected topics. However, the question whether the communities detected have corresponding real-world counterparts remained unanswered.

In this paper, we study retweet networks of the Members of the European Parliament and investigate how closely the community structure of these networks matches the political group membership and country membership. We approach these issues by employing a set of analytical tools. We use tools from the network theory, which have been applied successfully to characterize a wide variety of complex systems. We show that network theory is particularly effective at uncovering structure without prior knowledge of political orientation.

To the best of our knowledge, there has been no previous work on the analysis of retweet networks of the Members of the European Parliament. Nevertheless, there is a considerable body of literature on aspects relevant to this study.

Twitter data have been used in various contexts encompassing identification of spreading patterns of popular information, classes of dynamical collective attention, linguistic usage patterns on worldwide scale, and political activity.

Conover et al. [2] predict the political alignment of Twitter users in the run-up to the 2010 US elections based on content and network structure. They [3] analyze the polarization of retweet and mention networks for the same elections. Borondo et al. [4] analyze the user activity during the 2011 Spanish presidential elections. They [5] additionally analyze the 2012 Catalan elections focusing on the interplay between language and the community structure of the network. Most existing research, as Larsson [6] points out, focuses on the online behavior of political figures during election campaigns.

Hix et al. [7] investigate the voting cohesion of political groups in the European Parliament. Larsson [6] examines the Twitter presence of representatives outside of the election periods.

Recent research has adopted a networks science-based approach to investigate the structure of legislative work in the US Congress, including committee and subcommittee membership [8], bill cosponsoring [9], and roll-call votes [10]. In a more recent work, Dal Maso et al. [11] examine the community structure with respect to political coalitions and government structure in the Italian Parliament.

There are three different ways how users on Twitter interact: 1) a user follows posts of other users, 2) a user can respond to other user's tweets by mentioning them, and 3) a user can forward interesting tweets by retweeting them. Based on these three interaction types, one can define three measures of influence exerted by a user on Twitter: *indegree influence* (the number of followers, indicating the size of his audience), *mention influence* (the number of mentions of the user, indicating his ability to engage others in conversation), and *retweet influence* (the number of retweets, indicating the ability of the user to write content of interest to be forwarded to others).

Kwak et al. [12] compare three different network-based measures of influence on Twitter: the number of followers, page-rank, and the number of retweets—finding the ranking of the most influential users differed depending on the measure. Cha et al. [13] also compare three different measures of influence: the number of followers, the number of retweets, and the number of mentions—also finding that the most followed users did not necessarily

IEEE computer society

Table I
THE NUMBER OF TWITTER USERS BY POLITICAL GROUP.

| Group | EP seats | Twitter accounts | Core network | Extended network |
|---|---|---|---|---|
| GUE-NGL | 52 | 36 | 34 | 35 |
| S&D | 191 | 151 | 126 | 136 |
| Greens-EFA | 50 | 45 | 43 | 45 |
| ALDE | 68 | 50 | 42 | 47 |
| EPP | 218 | 152 | 118 | 136 |
| ECR | 72 | 49 | 37 | 44 |
| EFDD | 47 | 35 | 26 | 31 |
| NA | 52 | 28 | 23 | 27 |
| Total | 750 | 546 | 449 | 501 |



Figure 1. Daily volume of tweets posted by the EP members.

score highest on the other measures. Wang et al. [14] compare the number of followers and page-rank with a modified page-rank measure that accounts for topic, again finding that ranking depends on the influence measure. Suh et al. [15] investigate how different factors such as account age, use of hashtags and URLs impact the influence of the user measured by the number of retweets. Bakshy et al. [16] investigate how information spreads on a retweet network and whether there are preconditions for a user to become influential.

Along with the small-world phenomenon and power-law degree distribution, the most salient property real-world networks exhibit is community structure, where network nodes are partitioned together in tightly knit groups, between which there are only loose connections [17]. The identification of the community structure of a network is commonly based on the optimization of its modularity [18]. Many different algorithms exist which employ various approaches [19]. In this work, we perform community detection based on the Louvain method, introduced by [20], which is among the algorithms known to perform well in a variety of domains [21].

The methodology for evaluating the degree to which the detected communities match known groups [22] used in this work is based on the $B^3$ algorithm [23]. The $B^3$ measure is the best measure according to the formal constraints for extrinsic clustering evaluation measures proposed by Amigó et al. [24].

The existing research suggests that retweets are the most suitable measure of influence on Twitter. Community detection in the retweet networks reveals the communities formed based on the spreading of influence on Twitter. Our goal is to identify the main factors along which these communities of influence are formed.

This paper is organized as follows. Section II describes the EU Parliament and the Twitter data collected. In section III we outline the Louvain community detection method, and the measures to evaluate the detected communities w.r.t. the actual groups. Sections IV and V present the results. We construct two retweet networks from the EU Parliament data, and compare them to political groups and countries of origin of the Parliament members. In section VI we discuss the results and plans for future research.

## II. THE EUROPEAN PARLIAMENT ON TWITTER

The European Union (EU) is a political and economic union which currently consists of 28 member states located in Europe. The EU operates through a system of supranational institutions which cover legislative, executive, judiciary, and monetary branches. The European Parliament, together with the Council of the European Union, is the principal legislative body.

### A. The European Parliament

The European Parliament (EP) functions analogously to national parliaments in traditional parliamentary democracies. It is elected every five years directly by the citizens of the EU. Member states are allocated a number of seats which roughly reflects the state's population. The EP members are elected on a national basis, but sit in the EP according to political groups they belong to.

Our work focuses on the period between October 1, 2014 and May 31, 2015. This period falls within the 8th EP which was elected on July 1, 2014. During this period, the EP consisted of 8 political groups:

1) European United Left–Nordic Green Left (*GUE-NGL*)—socialists and communists group,
2) Progressive Alliance of Socialists and Democrats (*S&D*)—social-democrats group,
3) The Greens-European Free Alliance (*Greens-EFA*)—greens and regionalists group,
4) Alliance of Liberals and Democrats for Europe (*ALDE*)—liberals group,
5) European People's Party (*EPP*)—christian-democrats group,
6) European Conservatives and Reformists (*ECR*)—conservatives group,
7) Europe of Freedom and Direct Democracy (*EFDD*)—euroskeptics group, and
8) the Non-Attached Members (*NA*)—independents.

### B. Collection of tweets

We acquired the list of the EP members form the official site of the EP[2]. The list consists of 750 EP members. Their

[2]http://www.europarl.europa.eu/meps/en/full-list.html (accessed June 1, 2015)
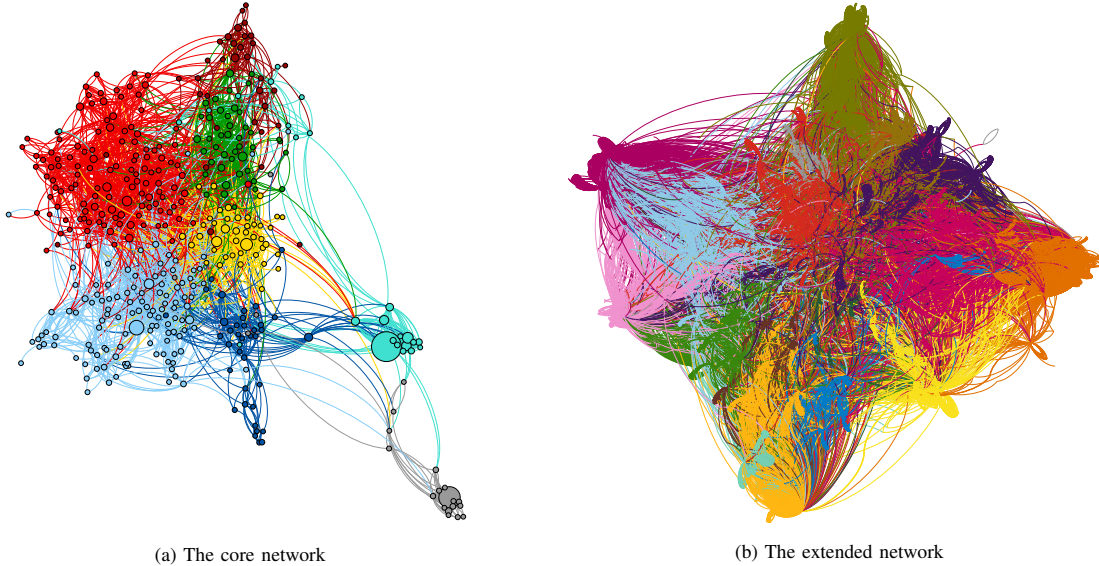
(a) The core network

(b) The extended network

Figure 2. The core network colored by the actual political groups and the extended network colored by the detected communities.

distribution according to political groups is presented in Table I (column *EP seats*). The official Twitter account of the EP, *Europarl_EN*, provides a list of Twitter accounts of the EP members[3]. We matched the EP members to the Twitter accounts and excluded Twitter accounts of former EP members; the result is a manually verified list of 546 Twitter accounts of all EP members which own one. The distribution of the EP members with Twitter accounts according to political groups is given in Table I (column *Twitter accounts*).

Through the Twitter Streaming API[4], we have monitored the activity related to the official accounts of the EP members; for each member we have collected all their tweets as well as all replies to and retweets of any tweet posted by them.

Within the period of our analysis—between October 1, 2014 and May 31, 2015—the EP members have posted 370,561 tweets, of which 195,797 (53%) are originally authored and the rest 174,764 (47%) are retweets. On average, all EP members together posted 1525 tweets per day, and each active member posted on average 3.1 tweets per day (Figure 1).

### C. Construction of retweet networks

The collected tweets described in the previous section are used to construct retweet networks. A retweet network is a directed weighted graph, where nodes represent Twitter users and edges represent the retweet relation. The direction of an edge corresponds to the direction of information spreading or influence; the weight of the edge is the number of times one user retweets the other.

[3]https://twitter.com/Europarl_EN/lists/all-meps-on-twitter/members (accessed September 30, 2014)

[4]https://dev.twitter.com/streaming/overview/request-parameters# follow

Table II
SIZE OF THE TWO RETWEET NETWORKS.

|  | Core network | Extended network |
|---|---|---|
| Nodes | 449 | 378,313 |
| Edges | 3,399 | 587,381 |
| Detected communities | 9 | 17 |

We construct two retweet networks: (i) the core network, containing as nodes only the EP members and (ii) the extended network, containing as nodes the EP members and all other users which have retweeted or have been retweeted by an EP member.

The core network consists of 449 nodes and 3,399 edges. The distribution of nodes according to political groups is in Table I (column *Core network*). The extended network consists of 378,313 nodes, of which 501 are the EP members, and 587,381 edges. The distribution of the member nodes according to political groups is also in Table I (column *Extended network*). Note that there are more EP members in the extended network (501) than in the core network (449) since the 52 EP members (the difference) were retweeted only by the non-EP members. An overview of the size of both networks is given in Table II.

### III. COMMUNITY DETECTION AND EVALUATION MEASURES

The core network consists of the EP members and the retweet relations between them. Since retweeting can be interpreted as expressing agreement on the posted tweet, it is reasonable to expect that members of the same political group will be bundled together within the network. In Figure 2a, we present the core network with a force-directed layout, where the color of the node identifies the

Figure 3. Distribution of political groups across the 9 communities in the core network. The shades of the basic color indicate different communities within a political group.



Figure 4. Composition of the 9 communities by political groups in the core network. Different colors indicate the 8 political groups in the EP.

political group of the EP member. There is an intuitive visual grouping of the EP members according to political groups. In Figure 2b, extended network with the detected communities is presented.

The goal of most community detection algorithms, implicit or explicit, is to find the best trade-off between a large intra-cluster density and a small inter-cluster density. Community detection algorithms perform maximization of modularity [25]. A good partitioning of a network in communities is one in which there are fewer than expected edges between communities. The modularity is, up to a multiplicative constant, the number of edges falling within groups minus the expected number in an equivalent network with edges placed at random. Previous work on roll-call votes suggests that the result of modularity optimization should find groups and coalitions in a parliament [11].

We perform community detection using the well established Louvain algorithm [20]. The Louvain method is a computationally very efficient algorithm that is well suited for large networks. It optimizes modularity through an iterative heuristic approach that consists of two repeating phases. In the first phase, modularity is optimized by allowing only local changes in communities; in the second, a new network is build that consists of one node for each previously found community. The algorithm repeats the iterations until the first phase can make no further improvements in modularity.

To asses how closely the detected communities match the political groups, we use the $B^3$ measure [23], which is considered as the most preferred measure for extrinsic evaluating of clusterings [24]. The $B^3$ measure decomposes the evaluation into calculating the precision and recall associated with each node in the network. Let $N$ be the set of all nodes in the network. For each node $n \in N$, we denote as $L(n)$ the set of nodes which have the same label as $n$, in this case, members of the same political group. With $C(n)$, we denote the set of all nodes which are members of the same community as $n$. The $B^3$ precision of a node $n$, $P(n)$, is computed as the fraction of nodes which have the same label and are in the same community as $n$, from all the nodes which are in the same community as $n$. Similarly, the recall of a node $n$, $R(n)$, is computed as the fraction of nodes with the same label

and in the same community from all nodes with the same label as $n$.

$$P(n) = \frac{|L(n) \cap C(n)|}{|C(n)|} \tag{1}$$

$$R(n) = \frac{|L(n) \cap C(n)|}{|L(n)|} \tag{2}$$

The precision and recall can be further combined into an $F_1$ score:

$$F_1(n) = 2 \cdot \frac{P(n) \cdot R(n)}{P(n) + R(n)} \tag{3}$$

The $F_1$ score is a special case of Van Rijsbergen's effectiveness measure [26], where precision and recall can be given different weights.

The precision reflects the homogeneity of a community. The more members from the same political group in the community, the higher the precision. Conversely, the recall reflects how spread-out across communities a political group is. The more members of a political group in the same community, the higher the recall. The $F_1$ score balances the precision and recall.

Furthermore, to quantify how well a political group is reflected in the community structure of the network, we calculate the mean precision, recall, and $F_1$ of the EP members of each group. Let $\{L_1, L_2, \ldots, L_k\}$ be the partitioning of the nodes according to actual labels. The precision, recall, and $F_1$ score of the set of the nodes $L_i$ are computed as:

$$P(L_i) = \frac{1}{|L_i|} \sum_{n \in L_i} P(n) \tag{4}$$

$$R(L_i) = \frac{1}{|L_i|} \sum_{n \in L_i} R(n) \tag{5}$$

$$F_1(L_i) = \frac{1}{|L_i|} \sum_{n \in L_i} F_1(n) \tag{6}$$

IV. RESULTS: THE CORE NETWORK

Community detection in the core network results in a partitioning into 9 communities. We look at how close the partitioning in communities matches the partitioning in political groups. Figure 3 shows how members of different groups are spread out across communities. Generally, most of the members of one group are located in a single community. The EP members from *S&D*, however, are

Table III
THE CORE NETWORK: A COMPARISON OF THE DETECTED AND
RANDOM COMMUNITIES WITH POLITICAL GROUPS.

| Communities | $Precision$ | $Recall$ | $F_1$ |
|---|---|---|---|
| detected | 0.785 | 0.658 | 0.684 |
| random | 0.199 | 0.127 | 0.141 |
| ratio | 3.9 | 5.2 | 4.9 |

Table IV
THE CORE NETWORK: OVERLAP MEASURES BETWEEN THE GROUPS
AND THE 9 COMMUNITIES.

| Group | $Precision$ | $Recall$ | $F_1$ |
|---|---|---|---|
| GUE-NGL | 0.471 | 0.574 | 0.518 |
| S&D | 0.902 | 0.462 | 0.611 |
| Greens-EFA | 0.783 | 0.910 | 0.842 |
| ALDE | 0.913 | 1.000 | 0.955 |
| EPP | 0.864 | 0.637 | 0.733 |
| ECR | 0.603 | 0.848 | 0.705 |
| EFDD | 0.252 | 0.479 | 0.330 |
| NA | 0.872 | 0.762 | 0.813 |
| micro avg. | 0.785 | 0.658 | 0.684 |

Table V
THE EXTENDED NETWORK: A COMPARISON OF THE DETECTED AND
RANDOM COMMUNITIES PER POLITICAL GROUPS AND COUNTRIES.

| | Communities | $Precision$ | $Recall$ | $F_1$ |
|---|---|---|---|---|
| Groups | detected | 0.389 | 0.223 | 0.249 |
| | random | 0.209 | 0.074 | 0.100 |
| | ratio | 1.9 | 3.0 | 2.5 |
| Countries | detected | 0.574 | 0.620 | 0.501 |
| | random | 0.105 | 0.112 | 0.089 |
| | ratio | 5.5 | 5.5 | 5.6 |

divided into two communities; the members from *EPP*, even though mostly contained in a single community, participate in 8 of the 9 communities. Figure 4 shows the composition of communities with respect to political groups. In general, the communities consist mostly of members of a single group. Notable exceptions are community no. 4 which contains many EP members from both *ECR* and *EFDD*, and community no. 7 which contains EP members from five different groups.

We calculate the mean precision, recall, and $F_1$ score for the network to characterize how well the community structure reflects the political group membership of the EP members. The results are in Table III. The precision is moderately high, 0.785, which reflects the fact that most of the communities, with the exception of the fourth and seventh community, are dominated by a single political group. The recall is above average, 0.658, which reflects the fact that most of the political groups, with the exception of *S&D* and *EPP*, are predominantly contained in a single community. The $F_1$ score is also above average, 0.684. In comparison, a random partitioning of the graph into 9 partitions has (on average over 1000 random partitionings) precision of 0.199, which is almost 4 times lower, recall of 0.127, which is over 5 times lower, and $F_1$ score of 0.141, which is nearly 5 times lower than the scores obtained with the partitioning into communities.

The overlap measures for each group are shown in Table IV. *GUE-NGL* has an average precision (0.471) and recall (0.574), which corresponds to its members being dispersed in several groups where they are not a majority. *S&D* has a very high precision (0.902) and average recall (0.462), as a result of being almost perfectly split into two communities where its members are an overwhelming majority. *Greens-EFA* has a moderately high precision (0.783) and a very high recall (0.910) because its members are mostly contained in a single community where they are a majority. *ALDE* has the highest precision (0.913) and recall (1.000) due to the fact that all its members are contained in a single community which contains only a few other members. *EPP* has a high precision (0.864) and above average recall (0.637), reflecting the fact that it members are predominantly contained in one large community and one small community, in both of which they constitute a majority. *ECR* has above average precision (0.603) and high recall (0.848) as a consequence of it being contained predominantly in a single community which contains also quite a few members of other groups. *EFDD* has the lowest precision (0.252) and an average recall (0.479) resulting from the fact that it is spread out across several communities, in none of which its members

are a majority. And finally, *NA* has high precision (0.872) and moderately high recall (0.762) corresponding to the largest part of its members being in a single community which contains no members from other political groups.

## V. RESULTS: THE EXTENDED NETWORK

The extended network consists of the EP members as well as all other users which have retweeted or have been retweeted by them. As such, it is several orders of magnitude larger than the core network. Moreover, the edges from non-EP members to the members far outnumber the edges between the members. This network reflects the retweeting practice of the general public when it comes to political issues. In this case, we want to investigate two alternatives: Is the partitioning of the network in communities dominated by the political groups, or by the countries of origin of the EP members?

We again apply the Louvain method for community detection which results in 17 communities. A force-directed layout of the network, colored by the detected communities is in Figure 2b. For further analysis, we focus only on the EP members—for them, we know the "ground truth", i.e., the political group and country which they represent.

### A. Communities and political groups

Analogously to the core network, we analyze how close the partitioning in communities matches the partitioning in political groups. Figure 5 shows how members of different groups are spread out across communities. Generally, the EP members from all of the groups are spread out across most of the communities. Figure 5 indicates that the community structure of the extended network does not reflect the political orientation of the retweeters of the EP members.

The mean precision, recall, and $F_1$ score for the extended network, which characterize how well the commu-

Figure 5. Distribution of political groups across the 17 communities in the extended network. The shades of the basic color indicate the distribution of an individual political group between the detected communities.

Table VI
THE EXTENDED NETWORK: OVERLAP MEASURES BETWEEN THE GROUPS AND THE 17 COMMUNITIES.

| Group | $Precision$ | $Recall$ | $F_1$ |
|---|---|---|---|
| GUE-NGL | 0.364 | 0.144 | 0.184 |
| S&D | 0.450 | 0.168 | 0.234 |
| Greens-EFA | 0.191 | 0.226 | 0.175 |
| ALDE | 0.351 | 0.339 | 0.326 |
| EPP | 0.452 | 0.167 | 0.219 |
| ECR | 0.227 | 0.240 | 0.228 |
| EFDD | 0.566 | 0.505 | 0.511 |
| NA | 0.250 | 0.317 | 0.276 |
| micro avg. | 0.389 | 0.223 | 0.249 |

nity structure reflects the political group membership, are presented in Table V (rows *Groups*). Both precision and recall (and subsequently $F_1$) are low. In comparison, a random partitioning of the graph into 17 partitions has (on average over 1000 random partitionings) precision which is almost 2 times lower, recall which is 3 times lower, and $F_1$ score which is 2.5 times lower than the ones obtained with the partitioning into communities. These ratios are still substantially lower than those for the core network.

In Table VI, we present the mean precision, recall, and $F_1$ score for each political group. All groups are characterized by low scores, except *EFDD* whose scores are even higher than those for the core network. Even though the retweet behavior of its members does not facilitate the community detection algorithm to group them together, the retweet behavior of their Twitter audience allows the algorithm to do a better job of grouping them together.

### B. Communities and countries

We next investigate how the country of origin of the EP members is reflected in the community structure. To this end, we evaluate the matching of the partitioning in communities with respect to the partitioning in countries. Figure 6 illustrates how members from different countries within the EU are spread out across communities. Many countries have their members contained within only a few communities. Moreover, in the majority of countries, one community contains the prevailing number of members.

The evaluation measures for the partitioning in countries are presented in Table V (rows *Countries*). In comparison to the partitioning in political groups, they are substantially higher. We also evaluated the average random partitioning,

Table VII
THE EXTENDED NETWORK: OVERLAP MEASURES BETWEEN THE COUNTRIES AND THE 17 COMMUNITIES.

| Country | $Precision$ | $Recall$ | $F_1$ |
|---|---|---|---|
| Austria | 0.124 | 1.000 | 0.220 |
| Belgium | 0.127 | 0.459 | 0.198 |
| Bulgaria | 0.028 | 0.333 | 0.050 |
| Croatia | 0.049 | 0.406 | 0.088 |
| Cyprus | 0.208 | 0.625 | 0.313 |
| Czech Republic | 0.037 | 0.167 | 0.058 |
| Denmark | 0.297 | 0.802 | 0.433 |
| Estonia | 0.021 | 0.500 | 0.041 |
| Finland | 0.821 | 0.686 | 0.741 |
| France | 0.852 | 0.415 | 0.553 |
| Germany | 0.460 | 0.893 | 0.607 |
| Greece | 0.750 | 1.000 | 0.857 |
| Hungary | 0.034 | 0.680 | 0.065 |
| Ireland | 0.833 | 1.000 | 0.909 |
| Italy | 0.884 | 0.605 | 0.685 |
| Latvia | 0.140 | 1.000 | 0.245 |
| Lithuania | 0.037 | 0.556 | 0.070 |
| Luxembourg | 0.019 | 0.556 | 0.037 |
| Malta | 0.042 | 0.500 | 0.076 |
| Netherlands | 0.370 | 0.837 | 0.513 |
| Poland | 0.676 | 0.938 | 0.785 |
| Portugal | 0.052 | 0.440 | 0.092 |
| Romania | 0.061 | 0.420 | 0.104 |
| Slovakia | 0.029 | 0.278 | 0.051 |
| Slovenia | 0.067 | 1.000 | 0.125 |
| Spain | 0.913 | 0.319 | 0.450 |
| Sweden | 0.444 | 0.680 | 0.534 |
| United Kingdom | 0.769 | 0.517 | 0.591 |
| micro avg. | 0.574 | 0.620 | 0.501 |

Table VIII
A SUMMARY OF THE $F_1$ SCORES.

| | | Communities | Groups (8) | Countries (28) |
|---|---|---|---|---|
| **Core** | | detected (9) | 0.684 | / |
| **network** | | random (9) | 0.141 | / |
| **Extended** | | detected (17) | 0.249 | 0.501 |
| **network** | | random (17) | 0.100 | 0.089 |

which has precision, recall, and $F_1$ score that are around 5.5 times lower than the ones obtained with the partitioning into communities. These ratios are comparable with those for the partitioning in political groups of the core network.

Table VII shows the mean precision, recall, and $F_1$ score for each country represented in the EP. The $F_1$ scores for the different countries vary substantially, ranging from 0.037 for Luxembourg to 0.909 for Ireland. The complete matrix of overlaps between the countries and communities is in Figure 7. A summary of the $F_1$ scores for the core and extended network, in comparison to the political groups and countries, for the detected and random communities, is in Table VIII.

## VI. CONCLUSIONS

In this paper we investigate the retweeting behavior of the EP members in a period of eight months. We have used the Twitter data to identify communities of influence and evaluated the detected communities with respect to the known "ground truth". The analysis reproduces the actual political groups and countries of origin of the EP members, without prior assumptions.

Figure 6. Distribution of countries across the 17 communities in the extended network. Different colors correspond to the detected communities.



| | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Austria | | | | | 13 | | | | | | | | | | | | |
| Belgium | | | | 1 | | | 9 | 2 | | | | 2 | | | | | |
| Bulgaria | | | 3 | | | | | 1 | 1 | | | | 1 | | | | |
| Croatia | | | 3 | | | 4 | 1 | | | | | | | | | | |
| Cyprus | | | | | | | | | | | | | | | 1 | 3 | |
| Czech Republic | 1 | | | | 3 | | 2 | 2 | 1 | 1 | | 2 | | | | | |
| Denmark | | | | | 1 | | | | | | | | 8 | | | | |
| Estonia | | | | | | | | | 1 | | | 1 | | | | | |
| Finland | | | | | 1 | | | | 1 | | | | | | | | 9 |
| France | | 31 | | | | 22 | | 7 | | | | | | | | | |
| Germany | | | | | 51 | | | 1 | 1 | | 1 | | | | | | |
| Greece | | | | | | | | | | | | | | | | 9 | |
| Hungary | | | | | 4 | | | 1 | | | | | | | | | |
| Ireland | | | | | | | | | | | | | | 10 | | | |
| Italy | | | | | 15 | 1 | | | 45 | | | | | | | | |
| Latvia | | | | | | | | | | | | 6 | | | | | |
| Lithuania | | | | | | | | 1 | | | | 2 | | | | | |
| Luxembourg | | | | | 2 | | | 1 | | | | | | | | | |
| Malta | | | | | 3 | | 3 | | | | | | | | | | |
| Netherlands | | | | | 1 | | | 1 | 21 | | | | | | | | |
| Poland | | | | | | | | 1 | | | 30 | | | | | | |
| Portugal | | | | | | | | 3 | | 1 | | 1 | | | | | |
| Romania | | | 2 | | 8 | | | 1 | 1 | 1 | | | | | | | |
| Slovakia | | | | | 2 | | 2 | 1 | | | | | 1 | | | | |
| Slovenia | | | | | 7 | | | | | | | | | | | | |
| Spain | 5 | | | | | | | 1 | | 20 | 6 | | 10 | | | | |
| Sweden | | | | | 2 | | | 1 | | | | | 13 | | | | |
| United Kingdom | | 20 | | | | | 39 | | | | 1 | | | 1 | | | |

Figure 7. The overlaps of countries and communities. The body of the table consists of the numbers of the EP members from different countries, belonging to the detected communities, with the majority of a community highlighted.

We have already successfully applied the Louvain method for community detection to uncover influential communities in the retweet networks, albeit in the context of climate and energy issues [1]. The results of the present study reinforce the suitability of the Louvain method for uncovering communities in the retweet networks. On the same data, we have also performed preliminary community detection by hierarchical stochastic block modeling [27]. The first experiment resulted in 12 communities for the core network, with a noticeably lower $F_1$ score (0.601). The extended network was partitioned into 77 communities, also with substantially lower $F_1$ scores: 0.156 in comparison to the political groups, and 0.270 in comparison to the countries.

Comparison of different community detection algorithms is interesting, but we plan to focus our future research in the following three key areas.

The presence and activities of the EP members on Twitter can be coupled with their actions in the Parliament. We plan to investigate the relations between the retweet networks and the roll-call vote networks. One of the findings of this study is that community detection can recreate the structure of different political groups with different degrees of effectiveness. Different political groups, also, manifest different levels of coherency in their voting behavior. Investigating whether these two phenomena are related will contribute to the overarching theme of the social media engagement by elected representatives.

So far, we have disregarded the contents of the tweets posted, and focused on the aggregated retweet behavior. The spreading of influence on Twitter is, however, dependant on the discussion topics. Different topics are accompanied by different levels of agreement and controversy, and may bring two political groups closer together or move them further apart. We plan to implement topic detection on Twitter data, and investigate how different topics influence the community structure of the retweet network of the EP members.

Different topics convey different sentiment. Sentiment analysis can be applied to uncover the attitude of different communities toward various issues. We have already applied the sentiment analysis to various domains, such as: (i) to compare the sentiment leaning of different network communities towards various environmental topics [1], (ii) to study the emotional dynamics of Facebook comments on conspiracy theories [28], (iii) to analyze the effects of Twitter sentiment on stock prices [29], (iv) to monitor the sentiment about political parties before and after the elections [30], and (v) to rank the widely used emojis by sentiment [31]. In the future we plan to employ sentiment analysis to characterize the sentiment of the EP political groups towards different policy and regulation issues.

REFERENCES

[1] B. Sluban, J. Smailović, S. Battiston, and I. Mozetič, "Sentiment leaning of influential communities in social

networks," *Computational Social Networks*, vol. 2, no. 9, pp. 1–21, 2015. [Online]. Available: http://dx.doi.org/10.1186/s40649-015-0016-5

[2] M. Conover, B. Gonçalves, J. Ratkiewicz, A. Flammini, and F. Menczer, "Predicting the political alignment of twitter users," in *Proc. 3rd Conf. Social Computing*, 2011.

[3] M. Conover, J. Ratkiewicz, M. Francisco, B. Gonçalves, A. Flammini, and F. Menczer, "Political polarization on twitter," in *Proc. 5th Intl. Conf. Weblogs & Soc. Med.*, 2011.

[4] J. Borondo, A. J. Morales, J. C. Losada, and R. M. Benito, "Characterizing and modeling an electoral campaign in the context of twitter: 2011 spanish presidential election as a case study," *Chaos*, vol. 22, no. 2, 2012.

[5] J. Borondo, A. Morales, R. Benito, and J. Losada, "Mapping the online communication patterns of political conversations," *Physica A: Statistical Mechanics and its Applications*, vol. 414, no. 0, pp. 403–413, 2014.

[6] A. O. Larsson, "The eu parliament on twitter—assessing the permanent online practices of parliamentarians," *Journal of Information Technology & Politics*, pp. 1–18, 2014.

[7] S. Hix, A. Noury, and G. Roland, "Voting patterns and alliance formation in the european parliament," *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, vol. 364, no. 1518, pp. 821–831, 2009.

[8] M. A. Porter, P. J. Mucha, M. E. J. Newman, and C. M. Warmbrand, "A network analysis of committees in the u.s. house of representatives," *Proc. Natl. Academy of Sciences*, vol. 102, no. 20, pp. 7057–7062, 2005.

[9] Y. Zhang, A. Friend, A. L. Traud, M. A. Porter, J. H. Fowler, and P. J. Mucha, "Community structure in congressional cosponsorship networks," *Physica A: Statistical Mechanics and its Applications*, vol. 387, no. 7, 2008.

[10] A. S. Waugh, L. Pei, J. H. Fowler, P. J. Mucha, and M. A. Porter, "Party polarization in congress: A network science approach," 2009.

[11] C. Dal Maso, G. Pompa, M. Puliga, G. Riotta, and A. Chessa, "Voting behavior, coalitions and government strength through a complex network analysis," *PLoS ONE*, vol. 9, no. 12, p. e116046, 2014.

[12] H. Kwak, C. Lee, H. Park, and S. Moon, "What is twitter, a social network or a news media?" in *Proc. 19th Intl. Conf. WWW*, 2010, pp. 591–600.

[13] M. Cha, H. Haddadi, F. Benevenuto, and P. K. Gummadi, "Measuring user influence in twitter: The million follower fallacy," in *Proc. Intl. AAAI Conf. on Weblogs and Social Media*, vol. 10, 2010, pp. 10–17.

[14] J. Weng, E.-P. Lim, J. Jiang, and Q. He, "Twitterrank: Finding topic-sensitive influential twitterers," in *Proc. 3rd Intl. Conf. Web Search & Data Mining*, 2010, pp. 261–270.

[15] B. Suh, L. Hong, P. Pirolli, and E. H. Chi, "Want to be retweeted? large scale analytics on factors impacting retweet in twitter network," in *2nd Intl. Conf. Social Computing*, 2010, pp. 177–184.

[16] E. Bakshy, J. M. Hofman, W. A. Mason, and D. J. Watts, "Everyone's an influencer: Quantifying influence on twitter," in *Proc. 4th Intl. Conf. on Web Search and Data Mining*, New York, NY, USA, 2011, pp. 65–74.

[17] M. Girvan and M. E. J. Newman, "Community structure in social and biological networks," *Proc. Natl. Academy of Sciences*, vol. 99, no. 12, pp. 7821–7826, 2002.

[18] M. E. J. Newman and M. Girvan, "Finding and evaluating community structure in networks," *Phys. Rev. E*, vol. 69, p. 026113, Feb 2004.

[19] S. Fortunato, "Community detection in graphs," *Physics Reports*, vol. 486, no. 3–5, pp. 75–174, 2010.

[20] V. D. Blondel, J.-L. Guillaume, R. Lambiotte, and E. Lefebvre, "Fast unfolding of communities in large networks," *Jou. Statistical Mechanics*, no. 10, 2008.

[21] S. Harenberg, G. Bello, L. Gjeltema, S. Ranshous, J. Harlalka, R. Seay, K. Padmanabhan, and N. Samatova, "Community detection in large-scale networks: a survey and empirical evaluation," *Wiley Interdisciplinary Reviews: Computational Statistics*, vol. 6, no. 6, 2014.

[22] J. Yang and J. Leskovec, "Defining and evaluating network communities based on ground-truth," *Knowledge and Information Systems*, vol. 42, no. 1, pp. 181–213, 2015.

[23] A. Bagga and B. Baldwin, "Entity-based cross-document coreferencing using the vector space model," in *Proc. 17th Intl. Conf. Computational Linguistics*, 1998, pp. 79–85.

[24] E. Amigó, J. Gonzalo, J. Artiles, and F. Verdejo, "A comparison of extrinsic clustering evaluation metrics based on formal constraints," *Information Retrieval*, vol. 12, no. 4, pp. 461–486, 2009.

[25] M. E. J. Newman, "Modularity and community structure in networks," *Proceedings of the National Academy of Sciences*, vol. 103, no. 23, pp. 8577–8582, 2006.

[26] C. Van Rijsbergen, *Information Retrieval*, 1979.

[27] T. P. Peixoto, "Hierarchical block structures and high-resolution model selection in large networks," *Phys. Rev. X*, vol. 4, p. 011047, Mar 2014.

[28] F. Zollo, P. Kralj Novak, M. Del Vicario, A. Bessi, I. Mozetič, A. Scala, G. Caldarelli, and W. Quattrociocchi, "Emotional dynamics in the age of misinformation," *PLoS ONE*, vol. 10, no. 9, p. e138740, 2015. [Online]. Available: http://dx.doi.org/10.1371/journal.pone.0138740

[29] G. Ranco, A. Aleksovski, G. Caldarelli, M. Grčar, and I. Mozetič, "The effects of Twitter sentiment on stock price returns," *PLoS ONE*, vol. 10, no. 9, p. e138441, 2015. [Online]. Available: http://dx.doi.org/10.1371/journal.pone.0138441

[30] J. Smailović, J. Kranjc, M. Grčar, M. Žnidaršič, and I. Mozetič, "Monitoring the Twitter sentiment during the Bulgarian elections," in *Proc. IEEE Intl. Conf. on Data Science and Advanced Analytics*. IEEE, 2015.

[31] P. Kralj Novak, J. Smailović, B. Sluban, and I. Mozetič, "Sentiment of emojis," *arXiv preprint*, 2015. [Online]. Available: http://arxiv.org/abs/1509.07761