2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)

From interaction to participation: the role of the imagined audience in social media community detection and an application to political communication on Twitter

Obaida Hanteer*, Luca Rossi*, Davide Vega D'Aurelio†, Matteo Magnani†
* DssLab, IT University Of Copenhagen, Denmark
{obha, lucr}@itu.dk
† InfoLab, Department of Information Technology
Uppsala University, Sweden
{davide.vega, matteo.magnani}@it.uu.se

***Abstract*—In the context of community detection in online social media, a lot of effort has been put into the definition of sophisticated network clustering algorithms and much less on the equally crucial process of obtaining high-quality input data. User-interaction data explicitly provided by social media platforms has largely been used as the main source of data because of its easy accessibility. However, this data does not capture a fundamental and much more frequent type of participatory behavior where users do not explicitly mention others but direct their messages to an invisible audience following a common hashtag. In the context of multiplex community detection, we show how to construct an additional data layer about user participation not relying on explicit interactions between users, and how this layer can be used to find different types of communities in the context of Twitter political communication.**

I. Introduction

Community detection is one of the most studied topics in social network analysis. While effective community detection algorithms are certainly necessary to identify meaningful communities, another equally crucial aspect is the definition of which connections should form the input data. However, it is generally recognized today that online social media are complex communication systems where different types of interactions are supported, and different network datasets can be built depending on the type of interaction to be studied. If we focus on Twitter, different types of data and different combinations of them have been considered when looking for communities. A common approach is to build a network based on following/follower relations [1], that can be easily obtained from the Twitter API. Researchers have soon realized that interaction networks are also directly available from the tweets, either defined by retweets [2] or by explicit mentions indicated by the @ character [3]. More recently, advances in multiplex social network analysis have led to the application of multiplex community detection methods, motivated by the hypothesis that analyzing these three types of connections together can reveal new types of communities. More recent work [4] has also suggested to organize the interactions (e.g., @ mentions) between users into multiple layers based on the topic in the exchanged text.

IEEE/ACM ASONAM 2018, August 28-31, 2018, Barcelona, Spain
978-1-5386-6051-5/18/$31.00 © 2018 IEEE

We claim that a strong limitation of the aforementioned approaches is that they only focus on the explicit interactions among users that take place within the social media: following, retweeting, and mentioning. However, much of Twitter contemporary interaction takes place within the space of polyadic conversations defined by hashtags. By adding a specific hashtag to their tweets, users do not only label the content of the tweet declaring its general topic but also identify the imagined audience [5]. This participation in a shared discussion, taking place on this hashtag-defined topical space, is largely ignored when Twitter data is used for community detection purposes [6], [4]. In our opinion, the reason why this data has been ignored is that, differently from explicit interactions where specific users are directly mentioned in tweets, imagined audiences are not explicitly available from social media APIs being in most cases not precisely known by the users when they are tweeting [7]. Our claim is that the implicit connections among users adopting common hashtags would be a valuable and natural input to a community detection algorithm.

In this paper, we provide the following contributions. First, we discuss the different choices to model Twitter interactions for community detection tasks, claiming that the connections explicitly provided by the Twitter platform omit a fundamental portion of the complex communication patterns happening on this system, and specifically hide participation dynamics in favor of interaction dynamics. Then, we provide a way to capture this additional social layer about user participation.

II. Topical audience model

A common way to model the multiple types of relationships supported by a social media platform is to use multiplex

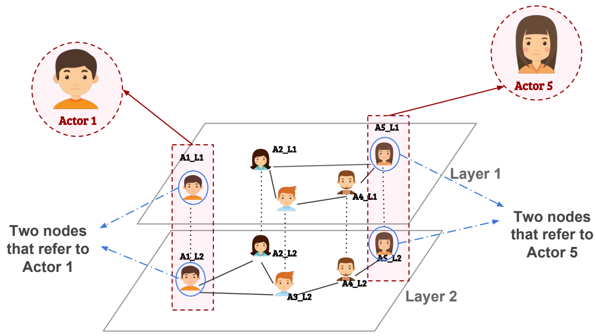

Fig. 1: An example of a multiplex network modeling two modes of interaction among five actors. This is modeled as five replicated nodes in two layers. A node and its replica are linked by a dotted line, to denote that they refer to the same actor, e.g., the same Twitter user

networks (Fig. 1). Existing work has already used multiplex networks where layers represent explicit interactions between users. Here we add a layer representing users participation. This layer, that we call topical audience model (TAM), aims at modeling the shared interests among users based on their participation in public discussions. We build the TAM layer in two phases. In the first phase, the discussions among the users of interest are modeled as a multiplex of n layers where n is the number of the topical discussions to be considered in the model. In the context of this paper, we use the explicit hashtag as a proxy for the topic of the shared conversation as suggested by [8]. Each discussion adds a layer to the multiplex and is modeled as a single clique that ties all the users who were part of the discussion by including the same hashtag in their messages. The intuition behind this is that the hashtag functions as a shared channel for discussions about a specific topic where one aims at broadcasting his/her views and opinions to everyone else in the channel.

In the second phase we compute a single network from these topical layers by applying a weighted flattening [9]. In our model, an edge e between $u1$ and $u2$ in the flattened graph has a weight w_e defined using the Jaccard coefficient as:

$$w_e = \frac{N(u_1, u_2)}{N(u_1) + N(u_2) - N(u_1, u_2)} \tag{1}$$

where $N(u_1)$ refers to the number of topical layers user u_1 has been part of and $N(u_1, u_2)$ refers to the number of topical layers users u_1 and u_2 have been both part of.

Once the weights have been computed, we can either keep them if we want to apply a weighted community detection algorithm, or we can use a threshold θ to create a TAM which considers only edges with weights exceeding θ, as we do in the case study described in the next section where we also show the effect of using different thresholds.

III. A CASE STUDY

The data we use in our case study was collected during the month leading to the 2015 Danish parliamentary election.

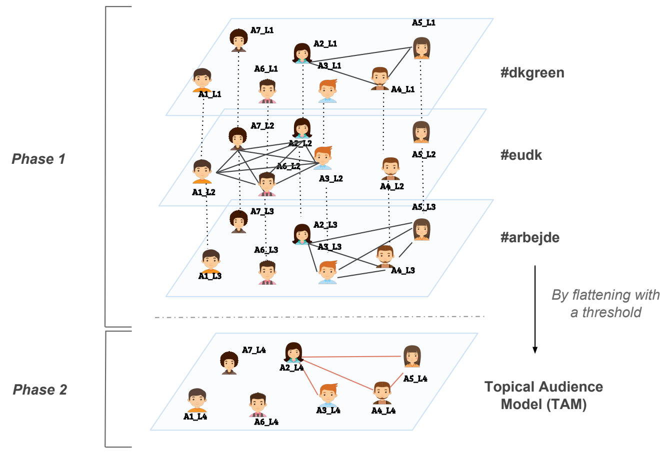

Fig. 2: The two phases in the formation of the Topical Audience Model (TAM). In this example, two actors are connected in Phase 2 if they are connected at least twice in Phase 1

Starting from a list of all the Danish politicians running for parliament and with a Twitter account, we collected all their Twitter content produced during the 30 days leading to the election. The initial dataset was formed by 490 politicians distributed across 10 parties, 5985 original tweets, 633 replies and 3993 retweets. Together with their Twitter activity, we registered also the political affiliation of the 490 politicians. Given the complexity of the Danish multi-party system, the parties have also been grouped according to actual coalitions: Red Block and Blue block.

	layer	#nodes	#edges	density	ccoef
1	Retweet	212	484	0.0007	0.011
2	Reply	127	169	0.0020	0.175
3	TAM.2	132	1594	0.0065	0.564
4	TAM.5	121	427	0.0017	0.738
5	TAM.7	68	152	0.0006	1.000

TABLE I: Layers used in the analysis. **ccoef** is clustering coefficient

The main focus in our experiments is to execute community detection on different multiplexes constituted of different combinations of Twitter interactions (retweet, reply, and topical interactions represented by TAM) so we can study the nature of the resulted communities on each multiplex. Table I shows the main descriptive data about the layers used to constitute these multiplexes. To build the TAM, the hashtags used by the politicians in the DKPol dataset were listed and qualitatively analyzed. We then excluded the hashtags used for the election campaign and those referring to political TV debates. After this filtering, we were left with only 23 hashtags used to refer to specific topics. The TAM was then constructed in two phases. In the first phase, a multiplex of 23 layers (layer per hashtag) was built as detailed in section II. In the second phase, the multiplex has been flattened into a TAM using a threshold

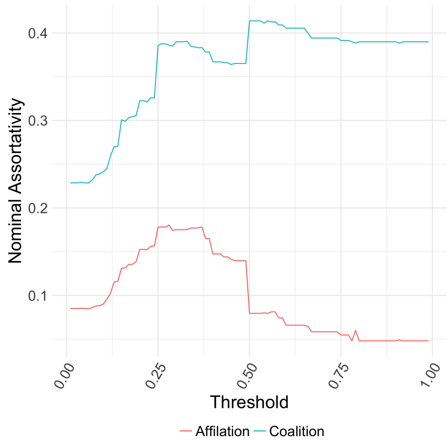

Fig. 3: Nominal assortativity of TAM with respect to the threshold θ

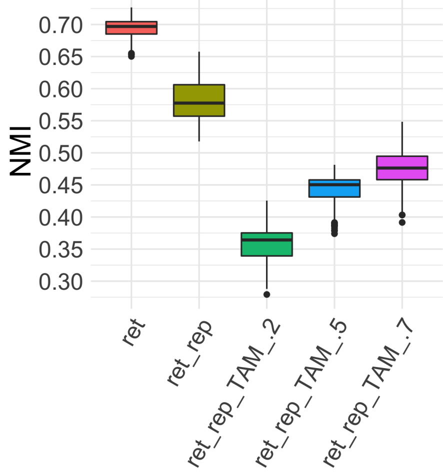

Fig. 4: (NMI) index of clusterigns identified by gLouvain and the political affiliation of ploticians

θ to filter out all edges with a weight less than θ. We observed the impact of various thresholds on the nominal assortativity [10] of the TAM layer, measured on the political affiliation and the political coalition of the politicians (Fig 3). We built the TAM layer for 3 different values of θ (0.2, 0.5, and 0.7)

We executed community detection using *Generalized Louvain* on 1) only the retweet layer, 2) the multiplex constituted of both the retweet and the reply layers, 3) the multiplex constituted of retweet, reply and TAM layer (one multiplex per threshold). Compared with other community detection methods, *Generalized Louvain* detected communities that are the closest to the groupings of politicians into political parties. While we do not use this as an evaluation criterion for how good a community detection method is, it is a good starting point to observe how the addition of other layers might affect the resulted communities. As the results of community detection using this method on the same multiplex slightly varied from an execution to another based on the order in which the nodes are scanned by the algorithm, we have run the algorithm 1000 times for each experiment. To investigate the social dynamics behind the observed communities beside the structural elements, the communities were evaluated against the groupings of politicians in political parties using the normalized mutual information metric (NMI) [11]. Within the context of this paper, we do not interpret NMI as a "quality" measure of the proposed community structure but as a measure of similarity between the proposed community structure and how the politicians are grouped into political parties.

IV. Results and Discussion

Fig. 4 shows that the highest level of NMI is observed when the communities are detected from the single layer network containing the retweets. Both the multiplex network including the replies, as well as those including the topical layer, score a lower value of NMI when communities are detected. This means that while the retweet layer contains communities that reflect the political affiliation of the politicians, this is no longer clearly visible when communities are detected including the other relations. This suggests the existence of two different dynamics behind the connections existing on the various layers of the multiplex structure: of political homophily in the case of the retweet layer and of different nature for the other layers. Fig. 5 shows the proportion of members belonging to each one of the two coalitions (Blue Block and Red Block) assigned to each one of the communities identified in (a) the retweet network (9 communities) and in the multiplex constituted of the retweet, the reply and the TAM with θ= 0.2 (10 communities). Looking at this figure, it appears evident that while in the case of the retweet network communities are largely politically homogeneous, the multiplex network including the TAM shows a significant number of communities that are actually formed by the members of both coalitions. This suggests that adding the TAM to the multiplex network allows us to observe interactions between political members that not only belong to different parties but also to different coalitions. While the users on the topical layer were connected because they used the same hashtag to refer to discussion topic during the same political campaign, it is hard to claim that they were not participating in the same conversation. On the opposite, we claim that even if they were not explicitly referring to each other, they were very aware of each other's presence as they were debating in the public topical space defined by the hashtags [12]. While this interaction is not easily captured since it is not readily available through the Twitter API, the proposed approach quantitatively captures the idea of users dealing with their imagined audience as repeatedly observed in qualitative studies of Twitter use [13]. From a political point of view, these results show how Twitter works as a public sphere and how topical debates gathered politicians from opposite parties. This rises the question if the levels of polarization that have been previously observed in political social media data [14], [2] were actual social dynamics or the result of the inherently biased data available that was unable to observe non-explicit interactions among users.

While originally introduced by Twitter, the idea of using hashtags to gather communication of users that are not otherwise connected has been adopted in various platforms. These platforms have thus evolved into a form of digital public space where discussions about the news, casual conversations but also political participation take place [15]. While the study of these participatory processes is more and more relevant to understand contemporary society, network approaches have only looked at direct and explicit interactions. Introducing the topical model to study hashtag-based interaction, we propose to extend the range of phenomena that can be fruitfully studied with a network approach. Moreover we suggest that this model should not be limited to Twitter data and that it could easily be applied to other hashtag-based communicative contexts (e.g. Instagram) as well as to other conceptually similar digital contexts (e.g. participation in Facebook pages).

A future extension of the proposed topical model should include the temporal aspects of interaction into the multiplex network model. While the current implementation assumes a topical stability, it is obvious that topics, as well as the association between actors and topics, change over time. Users might want to discuss a specific issue when it is highly relevant in society and then switch to another topic a few days or hours later. Twitter itself acknowledges this dynamic though the identification of ever changing trending topics that describe what is being discussed in a specific moment in a specific geographical context. Recent contributions in multiplex networks [4] have proposed to model the temporal dimension as layers of a multiplex structure to be subsequently used for community detection approaches that include temporal information. Such an approach, combined with the topical model we have introduced, could address more of the complexity we encounter in social media, where groups of users discuss within topical spaces constantly moving from one topic to the next one, in an ever evolving network of actors, moments and themes.

V. Conclusions

In this paper we have introduced a novel approach to model the participation in hashtag-based Twitter conversations. We have done this by modelling the participation into a hashtag-based discussion as a layer of a multiplex network where users are connected if their shared participation is above a given threshold θ. We have also applied this approach in the context of Twitter data collected in 2015 in Denmark during the month leading the the general election.

Acknowledgment

This research is supported by the VIRT-EU project funded by the European Union's Horizon 2020 research and innovation program under grant agreement No 727040

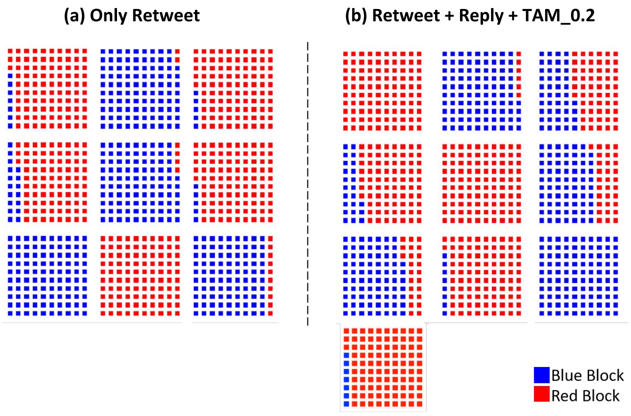

Fig. 5: Proportion of political coalitions (Red Block and Blue Block) within the communities detected on both **a)** only the retweet network **b)** the multiplex network including retweets, replies and the TAM $\theta = .2$ (color figure)

References

[1] H. Kwak, C. Lee, H. Park, and S. Moon, "What is twitter, a social network or a news media?" in *Proceedings of the 19th international conference on World wide web*. ACM, 2010, pp. 591–600.

[2] M. Conover, J. Ratkiewicz, M. R. Francisco, B. Gonçalves, F. Menczer, and A. Flammini, "Political polarization on twitter," in *Proc. of the fifth International Conference on Web and Social Media*. AAAI Press, 2011, pp. 89–96.

[3] J. Yang and S. Counts, "Predicting the speed, scale, and range of information diffusion in twitter." *Icwsm*, vol. 10, no. 2010, pp. 355–358, 2010.

[4] D. Vega and M. Magnani, "Foundations of temporal text networks," *Applied Network Science*, 3 2018.

[5] E. Litt, "Knock, knock. who's there? the imagined audience," *Journal of Broadcasting & Electronic Media*, vol. 56, no. 3, pp. 330–345, 2012.

[6] T. Highfield, L. Kirchhoff, and T. Nicolai, "Challenges of tracking topical discussion networks online," *Social Science Computer Review*, vol. 29, no. 3, pp. 340–353, 2011.

[7] M. S. Bernstein, E. Bakshy, M. Burke, and B. Karrer, "Quantifying the invisible audience in social networks," in *Proc. of SIGCHI Conference on Human Factors in Computing Systems*. ACM, 2013, pp. 21–30.

[8] A. Bruns and J. E. Burgess, "The use of twitter hashtags in the formation of ad hoc publics," in *Proc. of 6th European Consortium for Political Research General Conference*, 2011, pp. 1–9.

[9] M. Berlingerio, M. Coscia, and F. Giannotti, "Finding redundant and complementary communities in multidimensional networks," in *International Conference on Information and Knowledge Management (CIKM)*, 2011, pp. 2181–2184.

[10] M. Newman, "Mixing patterns in networks," *Physical Review E*, vol. 67, no. 2, pp. 261–287, 2011.

[11] A. Lancichinetti and S. Fortunato, "Community detection algorithms: a comparative analysis," *Physical review E*, vol. 80, no. 5, pp. 561–577, 10 2009.

[12] A. Bruns and J. Burgess, "Twitter hashtags from ad hoc to calculated publics," *Hashtag publics: The power and politics of discursive networks*, 2015.

[13] A. E. Marwick and D. Boyd, "I tweet honestly, i tweet passionately: Twitter users, context collapse, and the imagined audience," *New media & society*, vol. 13, no. 1, pp. 114–133, 2011.

[14] A. Bessi, F. Zollo, M. D. Vicario, M. Puliga, A. Scala, G. Caldarelli, B. Uzzi, and W. Quattrociocchi, "Users polarization on facebook and youtube," *PloS ONE*, vol. 11, no. 8, pp. 17–37, 8 2016.

[15] Y. Bonilla and J. Rosa, "Ferguson: Digital protest, hashtag ethnography, and the racial politics of social media in the united states," *American Ethnologist*, vol. 42, no. 11, pp. 4–17, 1 2015.