

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, luca}@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. A large number of algorithms have been defined to identify different types of communities in online social media based on the connections between users. However, 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.

In our opinion, the easy availability of large social network data from social media APIs has initially contributed to keep the focus of community detection tasks on the development of better and better algorithms, to then apply them to the large datasets directly provided by the social media platforms’ APIs. However, today it is generally recognized 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.

Despite the rich and diverse types of network data mentioned above, we claim that a strong limitation of these approaches is that they only focus on the explicit interactions between users that take place within the social media: following, retweeting, mentioning. All these actions are characterized by the presence of two actors, which can naturally be represented as connections in a network. Nevertheless, while this kind of data has been used for long time in a successful way [5], it is far from providing a complete representation of Twitter-based communication. As noticed in [6], much of Twitter contemporary interaction takes place within the space of polyadic conversations defined by hashtags. This conversational space extends beyond the range of visibility of the following/follower network (messages written using a hashtag are visible to every other user following the same hashtag, regardless if they are following the author) and constitutes what has been defined an *ad-hoc public* [7]. By adding a specific hashtag to their tweets users do not only label the content of the tweet declaring its general topic (e.g. using the #election hashtag) or providing information about how the tweet should be interpreted (e.g. with the #kidding hashtag), but also identifying the user’s imagined audience [8].

This participation in a shared discussion, taking place on this hashtag-defined topical space, is at the same time one

of the key aspects of Twitter as a communication platform and largely ignored when Twitter data is used for community detection purposes [9], [4]. 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 [10]. In summary, there are no explicit connections between a user and its imagined audience, but we claim 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 available choices to build Twitter datasets to be used as input 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. Building such a layer is not straightforward, because even if hashtags are directly available in the tweets we cannot just connect all users using the same hashtag; this would correspond to large cliques not directly suitable as input to community detection algorithms. Therefore, we show a simple procedure where tweets are grouped according to their topics and users are connected with weighted edges indicating to what extent their topical participation spaces overlap. Finally, we introduce a new dataset about the Twitter activity of all candidates to the last Danish general elections, and we test our approach on it showing how it allows us to identify new community structures and how it affects our interpretation of the data.

II. TOPICAL AUDIENCE MODEL

Beside the networks built on a single type of interaction among Twitter users (e.g. networks based on retweet activity) or built collapsing various types of interactions into a single network (e.g. networks merging retweets and replies), a common way to model the multiple types of relationships supported by a social media platform is to use multiplex networks, as illustrated in Fig. 1.

Existing work has already used multiplex networks where layers represent explicit interactions between users. Here we add an additional layer representing user 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 do so by considering the similarity in topical participation between two users as a relationship among them, 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 considered topical discussions to be included in the model. In the context of this paper we will use the explicit hashtag used by the users to self-declare the topic of the message as a proxy for the topic of the shared conversation




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

as suggested by [6]. Each topical discussion adds a layer to the multiplex and is modeled as a single clique that ties all the users who participated in the shared discussion adding the same hashtag to their messages or retweeting messages containing the hashtag. 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. Therefore, a tie among two users in a layer l in this multiplex does not necessarily imply a direct interaction among them about the topic represented by l but only implies their active participation in that discussion by using the relevant hashtag.

In the second phase we compute a single network from these topical layers by applying a weighted flattening [11]. In our model, an edge e between u_1 and u_2 in the flattened graph has a weight w_e defined using the Jaccard coefficient as:

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

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 addition of a TAM as a layer to a multiplex network that models other modes of interaction allows us, in the context of Twitter data, to include information about people participating in discussion about the same topic (or set of) but not making that co-participation explicit through retweets or replies. In other words, the addition of a topical layer allows us to include in the modelled data the social dimension of participating in a shared conversational space.

Within this perspective it is important to assess the effect that the addition of a topical layer will have on real-world data

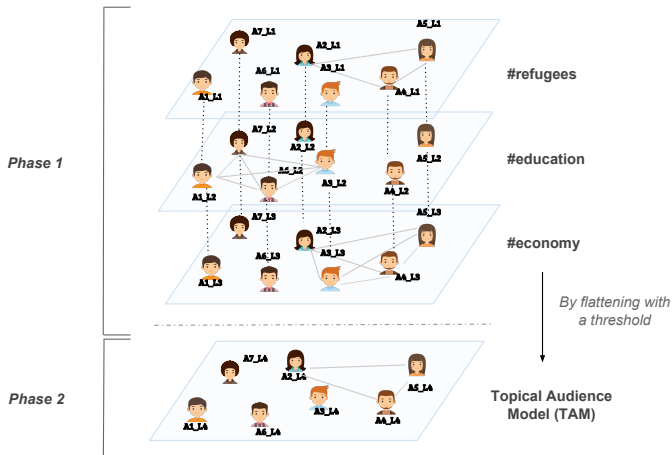


Fig. 2: The two phases in the formation of the Topical Audience Model (TAM)

to explore if and to what extent it will change the communities that can actually be observed within the data.

A. Data

The data we use to test the proposed approach was collected during the month leading to the 2015 Danish parliamentary election. Starting from a list of all the Danish politicians running for parliament and with a Twitter account, we collected all the tweets written during the 30 days leading to the election. The initial dataset was formed by 493 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 494 politicians. Given the complexity of the Danish multi-party system, the parties have also been grouped according to actual coalitions: Red Block, currently at the opposition, grouping Alternativet, Radikale Venstre, Enhedslisten, Socialdemokratiet and Socialistisk Folkeparti) and the Blue block, currently in government, that groups: Dansk Folkeparti, KristenDemokraterne, Liberal Alliance, Venstre, Det Konservative Folkeparti.

Table I shows the main descriptive data of the interaction networks (retweets and replies) obtained from the data, as well as for the multiplex topical network built with three different values of θ .

layer	#nodes	#edges	density	ccoef	#coms
1 Retweet	212	484	0.0007	0.011	9
2 Reply	127	169	0.0020	0.175	26
3 TAM.2	132	1594	0.0065	0.564	7
4 TAM.5	121	427	0.0017	0.738	12
5 TAM.7	68	152	0.0006	1.000	19

TABLE I: Summary table of the networks used in the analysis. **ccoef** refers to the clustering coefficient. **#coms** refers to the number of detected communities in each layer using gLouvain algorithm

B. Experimental settings

The main focus in our experiments is to execute community detection on different multiplexes constituted of different combinations of Twitter interactions as different layers and to study the nature of these resulted communities.

The main building blocks that were used as layers to build different multiplexes were the retweet interaction, the reply interaction, and the topical interaction. The retweet interaction as a layer was modeled as an undirected graph where a tie is added between two users if a retweet interaction happened between them (in any direction) and the same holds for the reply interaction. As to the topical interaction, we had to decide on two things: 1) how to choose the discussions based on which TAM can be built, 2) what threshold to use.

	Hashtag	Theme	Frequency
1	#dkgreen	climate	189
2	#sundpol	health	61
3	#talop	economy	53
4	#dkaid	economy	47
5	#arbejde	work	47
6	#nuloverdeigen	refugees	45
7	#dksocial	work	40
8	#skolechat	education	40
9	#itpol	IT	31
10	#dajegvar12	children	25

TABLE II: Top 10 most used topical hashtags

To decide on which discussions among the politicians to consider in the TAM, the hashtags used by the politicians in the DKPol dataset were listed and qualitatively analyzed. We then excluded the hashtags that were just about the election campaign as such (like #dkpol or #fv15, where fv stands for *Folketingsvalget* which is the Danish word for general election) and those referring to political TV debates (e.g. #tv2valg or #DRdinstemme). After this filtering we were left with only hashtags used to refer to specific topics. This resulted in a set of 23 hashtags. Table II shows the top 10 in terms of frequency.

Each of these hashtags was considered as a separate discussion and this resulted in 23 layers in the first phase of the formation of the TAM. As we have detailed above, to reduce these layers into a TAM we defined a threshold θ . In the context of this study, a low threshold will result into a TAM connecting politicians who interacted on only a small portion of their preferred topics, while a high threshold will produce connections only among politicians who had a large overlap in their topical interests.

To explore this dynamic and thus identify the proper threshold for the study we observed the impact of various thresholds on the assortativity of the TAM, measured on the political affiliation of the various politicians. Assortativity on nominal attributes (also known as nominal assortativity) measures the level to which the nodes of the graph are connected to similar nodes based on a given property [12], in our case the affiliation to a specific political party. Fig. 3 shows a clear dynamic between the threshold level and the political assortativity

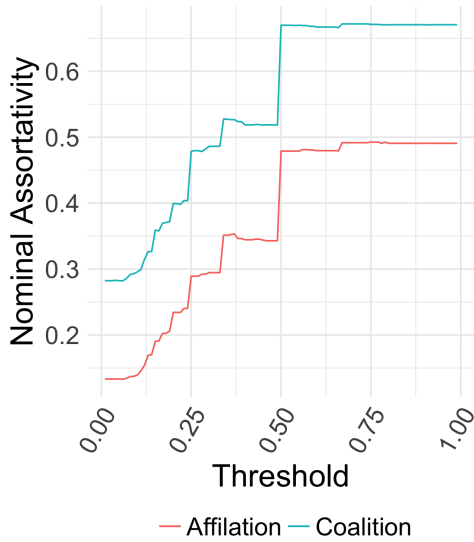


Fig. 3: Nominal assortativity of TAM with respect to the threshold θ

(assortativity based on the political affiliation of the nodes) of the topical layer.

Clearly, a high threshold ($\theta > 0.5$) in our dataset increases the political assortativity filtering out edges among differently affiliated politicians. While this is a reasonable result, although not necessarily expected, suggesting that politicians belonging to the same party show stronger topical connections, it also shows how even the highest threshold does not produce a fully assortative network.

Building on these results we have generated the topical layer for three thresholds (0.2, 0.5, 0.7). We did our experiments on 5 multiplexes for the different combinations of the layers. 1) with only the retweet layer, 2) with both the retweet and the reply layers, 3) with the retweet, reply and topical layer (one multiplex per threshold).

Given these multiplex networks we then study the resulting community structure using a modularity-maximization community detection method (generalized Louvain). As the results of the community detection might slightly vary from an execution to another based on the chosen starting node, we executed generalized Louvain 1000 times for each multiplex network.

To further investigate the social dynamics behind the observed communities beside the structural elements, the communities were evaluated against the information of the grouping of politicians in political parties through the normalized mutual information index (NMI).

This was done on the one side to evaluate how much Twitter networks of interaction replicate political affiliations and, on the other side, to evaluate the impact of the topical layer. When the community observed were compared with the political affiliation, we used, per each experiment, the execution that gave communities with the highest modularity.

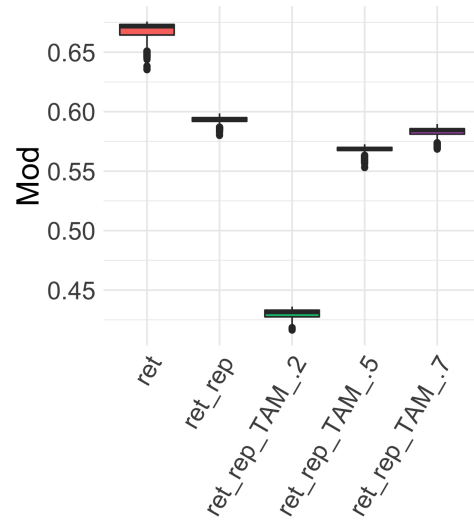


Fig. 4: Modularity value of the different combination of layers/thresholds

C. Results

In order to evaluate how the addition of a topical layer affected the resulting multiplex network we will present the analysis of both the community structure and of the community composition.

1) *Community structure*: Fig. 4 shows the modularity value observed on different single and multiplex networks, modelling the data described above. These networks are: the single layer network containing only the retweets (ret), the multiplex network containing retweets and replies (ret_rep), and the multiplex network containing retweets, replies and the topical layer. For the latter we present three versions with different thresholds used to define the topical connections (t) with the values of 0.2, 0.5, 0.7.

It can be clearly observed how the single layer retweet network shows the highest modularity value (.65) sign of a structure with clearly identifiable clusters. When we add layers representing different types of interactions this modular structure is disturbed thus the network appears to have a lower modularity. It is interesting to observe how adding the most inclusive topical layer ($\theta = .2$) results in the most conspicuous drop in the modularity value, suggesting how edges on the new topical layer where formed out of a social dynamic different from the one driving the retweet layer. While this effect gets less and less clear when we increased the threshold used in the topical layer, the addition of a topical layer results, within the context of this data, in a less modular network structure.

2) *Community composition*: While the analysis of the modularity of the multiplex networks provides an initial insight into the effect of the topical layers on the Twitter multiplex network, the full range of consequences becomes more visible when we look at the internal composition of the resulting communities. Fig. 5 shows the value of Normalized Mutual Information (NMI) when we compare the communities iden-




Fig. 5: Normalized Mutual Information (NMI) between the communities identified by generalized Louvain and the actual political affiliation

tified within the various networks with the known political affiliation of the politicians forming the network.

NMI has been largely used, even if in slightly different variants, as a viable strategy to evaluate the result of a community detection effort [13]. Within the context of this paper, we will interpret NMI not as a measure of “quality” of the proposed community structure but as a measure of similarity between the proposed community structure and how the politicians forming our network are clustered in political parties. In other words, we are interpreting NMI as a measurement of how much the network structures we observe in the networks are based on shared political affiliation between the nodes.

Fig. 5 shows how the NMI behaves very similarly to the modularity value. 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. It is worth noticing that the number of communities identified in the various multiplex networks is very similar for retweets, retweets + replies and full multiplex networks and it increases only when the threshold θ to build the topical layer is at high values. These results are summarized in Table I.

This data shows that while the retweet layer contains communities that reflect the political affiliation of the nodes, 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. Before discussing more in details the consequences of this in the following section, it is interesting to observe in more details the actual composition of the




Fig. 6: Distribution of the number of parties that constitute the communities detected on both the retweet network and the multiplex network including retweets, replies and topical layer $\theta = .2$

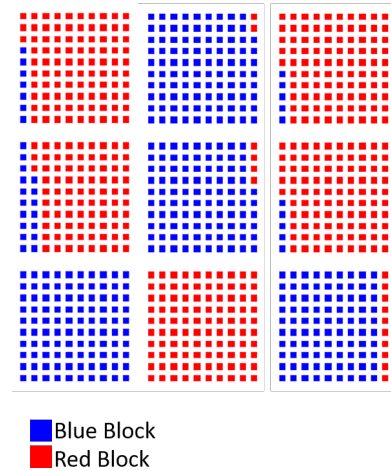


Fig. 7: Proportion of politicians belonging to the two coalitions (Red Block and Blue Block) within the communities detected on the retweet network (color figure)

communities that are detected on the retweet network and on the full multiplex.

Fig. 6 shows the distribution of how many parties were present within each community that was identified in the single-layer retweet network and the multiplex network constituted of retweets, replies and TAM with a threshold θ set to 0.2. In a network where communities were only based on party affiliation we could expect a number of communities containing a single party equal to the number of parties. On the opposite, in a network with external forces we would expect a larger number of communities bringing together different parties. Partially supporting this hypothesis, Fig. 6 shows how the multiplex constituted of retweets, replies and TAM has a bimodal distribution. There is a number of communities that gather 8 distinct parties, not present if only the retweet network is used. Nevertheless, the multiplex also has an equal high number (3) of single-party communities, more than those identified on the retweet network.

While these results might appear contradictory, it is worth remembering the political context in which the data was collected, that is characterized by political parties grouped within coalitions. The analysis of the data at a coalition level provides a better understanding of it. Fig. 7 and 8




Fig. 8: Proportion of politicians belonging to the two coalitions (Red Block and Blue Block) within the communities detected on the multiplex network including retweets, replies and the TAM $\theta = .2$ (color figure)

show 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 the retweet network (9 communities) and in the multiplex topical network with threshold θ set at $.2$ (10 communities). Looking at these figures it appears evident that while in the case of the retweet network communities are largely politically homogeneous, the multiplex topical network shows a significant number of communities that are actually formed by the members of both coalitions. This stark difference suggests how 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. These interactions, as we will discuss further in the following section, took place in the shared space of Twitter hashtag-based conversations and they undoubtedly contributed to the political debate.

IV. DISCUSSION AND CONCLUSIONS

In this paper we have introduced a novel approach to model the participation to 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 to the general election. We organize the discussion of the results as follows: first we will discuss the interpretation of the data observed within the context of the Danish election; then we will discuss these results in the larger context of the study of Twitter-based participation and social media participation in general, and propose future improvements.

A. Political community of conversation

In the context of the Twitter data collected during the 2015 Danish general election we have observed that the community structure emerging from the single layer retweet network, or from the multiplex retweet + replies network is significantly different from the TAM multiplex that includes topical interactions. More precisely, the edges on the networks based on explicit interactions seem to be largely driven by political homophily, connecting nodes that are either part of the same political party or of the same political coalition. Adding the topical layer to the multiplex model allows us to retrieve different communities that seem to bring together politicians from different political positions. 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 [7]. 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 [14].

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 raises the question if the levels of polarization that have been previously observed in political social media data [15], [2] were actual social dynamics or the result of the inherently biased data available that were unable to observe non-explicit interactions among users.

B. Participation and imagined audience in other social media

Many-to-many polyadic conversations, where users address an unknown imagined audience, are common in a growing number of contemporary social media. While originally introduced by Twitter, the idea of using hashtags to gather communication of users that are not otherwise connected (e.g. not following each other) have been adopted in various platforms such as Facebook and Instagram. This has generally been implemented through hashtags or similar solutions. 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 [16]. 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).

C. Temporal-topical networks

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. RELATED WORKS

Twitter provides researchers with a large corpus of complex social interactions, offering the opportunity to investigate social phenomena at various scales, from conversations and localized debates to information cascades and global communities. A key issue within this context is how to build a data model that represents the original source so that qualitative analysis methods can find coherent and consistent results. So far, modeling methods have been developed in two different directions. On one side, thread-based models have focused on mapping the structure of interactions into simple networks based on the *polyadic conversations* [17] found in the data set. On the other side, *network multiplicity* [18], [19] has been used to represent multiple types of interactions (e.g., following, mentions and retweets) simultaneously as multi-relational data sets with attributed nodes and/or edges. Compared with thread-based strategies, multiplicity produces more flexible and richer models at the expense of increasing their complexity. In the following we describe recent research efforts using both modelling methods.

A. Thread-based models

Thread-based models have been explored in a wide range of analysis from the response of online communities to natural disasters [20] to the analysis of political alignment [2]. Different thread-based models have been suggested in the past based on the properties of the data source and the type of analysis intended [21]. All these modeling methods follow the same generative process, which can be summarized in two steps: (i) independently capturing the discussion threads (usually as polyadic conversations), and (ii) merging them into a single network. The main objective of this process is to produce models with none or small noise, while accurately capture the main properties of the original data.

For example, in [22] and [23] the authors generate, for each of the collected tweets, a tree of polyadic conversations by inversely following the chain of Twitter users' interactions (replies, mentions and retweets). Then, it is assumed that actual conversations have a limited duration and are more likely to involve a small number of participants; afterwards, the trees are pruned both in depth, based on the time elapsed since the root tweet; and in breadth, based on the number of unique participants. During the second step, all the polyadic conversations are merged to produce a single network between actors, where the weight of the edges represent the sum of messages exchanged on all conversations.

While thread-based models produce simple network structures easy to analyze, the possibilities of the model can be limited by the decisions made early during the modelling process. What actually constitutes a conversation or thread or what are the boundaries (in time, number of actors and messages) that limit them is still an open issue. For example, based on the type interactions considered, Twitter conversations can be formed using replies instead of mentions, which will generate sparser but larger polyadic conversations [24].

Some recent studies have focused on expanding the conversation trees used by the thread-based models by including participatory information such as the interconnections among specific users of a related hashtag, for example its early adopters [25]. As a result, the polyadic conversations become a *focused graph* instead, with mixed information about the users' interactions and their participation in some hashtag-related community. While this type of network is relevant to study diffusion processes or to contextualize the conversations inside a larger community, once the different focused graphs are merged the final model can be highly noisy.

B. Multiplexity

Multilayer, multiplex, multi-relational [18] and heterogeneous information networks [26] have been extensively used to represent relations between different types of objects and/or different types of relations. Compared with the thread-based models, these models are usually more complete and flexible, as less information is ignored during the modelling process.

In several microblogging sites (e.g. Twitter) users employ platform-specific conventions (e.g., mentions, retweets) to endorse previous messages, increase their visibility or just indicate a direct reply [27] [28]. Several works enforce such difference by creating complex networks where actors represent users and the edges on each layer represent a different type of interactions. Temporal networks [29] can also be represented using multilayer models using (i) a set of nodes, to represent the actors interacting in the social graph, (ii) a set of directed edges, to represent the direction of such interactions and (iii) a set of layers to group together the interactions happening during the period time slice.

Heterogeneous networks, instead, are typically used in socio-semantic models that combine structural and semantic data simultaneously [30]. An illustrative example can be found in [31] where the authors represent users, words and topics in

the messages as three different types of vertices in the network. The edges, then, represent relations between actors and words or topics and words. This type of network is useful to cluster elements of the data set based on categorical properties, but it makes difficult to detect conversational and participatory patterns.

Finally, multiplexity has been also used in combination with topic modelling to relate information from multiple sources [30]. For example, in [23] the authors used thread-based modelling to generate a communication layer, and then applied several semantic analysis procedures such as natural language processes and semantic relatedness to generate supporting layers describing multiple text features. One of the key differences with the present work is their use of thread-based models to represent the communication network rather than the TAM which, as we discussed in the paper, allows us to capture an additional type of communication practice.

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

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." *Icwsn*, 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] W. Deitrick and W. Hu, "Mutually enhancing community detection and sentiment analysis on twitter networks," *Journal of Data Analysis and Information Processing*, vol. 1, no. 3, pp. 1–19, 2013.
- [6] 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.
- [7] A. Bruns and J. Burgess, "Twitter hashtags from ad hoc to calculated publics," *Hashtag publics: The power and politics of discursive networks*, 2015.
- [8] E. Litt, "Knock, knock. who's there? the imagined audience," *Journal of Broadcasting & Electronic Media*, vol. 56, no. 3, pp. 330–345, 2012.
- [9] 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.
- [10] 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.
- [11] 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.
- [12] M. Newman, "Mixing patterns in networks," *Physical Review E*, vol. 67, no. 2, pp. 261–287, 2011.
- [13] A. Lancichinetti and S. Fortunato, "Community detection algorithms: a comparative analysis," *Physical review E*, vol. 80, no. 5, pp. 561–577, 10 2009.
- [14] 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.
- [15] 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.
- [16] 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.
- [17] M. Magnani and L. Rossi, "Conversation retrieval for microblogging sites," *Inf. Retr.*, vol. 15, no. 3, pp. 354–372, 6 2012.
- [18] M. Kivelä, A. Arenas, M. Barthelemy, J. P. Gleeson, Y. Moreno, and M. A. Porter, "Multilayer networks," *Journal of Complex Networks*, vol. 2, no. 3, pp. 203–271, 9 2014.
- [19] M. E. Dickison, M. Magnani, and L. Rossi, *Multilayer Social Networks*. Cambridge University Press, 2016.
- [20] Y. Qu, P. F. Wu, and X. Wang, "Online community response to major disaster: A study of tianya forum in the 2008 sichuan earthquake," in *Proc. of 42nd Hawaii International Conference on System Sciences*. IEEE, 2009, pp. 1–11.
- [21] P. Aragón, V. Gómez, D. García, and A. Kaltenbrunner, "Generative models of online discussion threads: state of the art and research challenges," *Journal of Internet Services and Applications*, vol. 8, no. 1, pp. 1–17, 10 2017.
- [22] L. Tamine, L. Soulier, L. B. Jabeur, F. Amblard, C. Hanachi, G. Hubert, and C. Roth, "Social media-based collaborative information access: Analysis of online crisis-related twitter conversations," in *Proc. of 27th ACM Conference on Hypertext and Social Media*. ACM, 2016, pp. 159–168.
- [23] S. Šćepanović, I. Mishkovski, B. Gonçalves, T. H. Nguyen, and P. Hui, "Semantic homophily in online communication: Evidence from twitter," *Online Social Networks and Media*, vol. 2, no. 1, pp. 1–18, 2017.
- [24] P. Cogan, M. Andrews, M. Bradonjic, W. S. Kennedy, A. Sala, and G. Tucci, "Reconstruction and analysis of twitter conversation graphs," in *Proc. of 1st ACM International Workshop on Hot Topics on Interdisciplinary Social Networks Research*. ACM, 2012, pp. 25–31.
- [25] D. M. Romero, B. Meeder, and J. Kleinberg, "Differences in the mechanics of information diffusion across topics: idioms, political hashtags, and complex contagion on twitter," in *Proc. of 20th International Conference on World Wide Web*. ACM, 2004, pp. 695–704.
- [26] C. Shi, Y. Li, J. Zhang, Y. Sun, and S. Y. Philip, "A survey of heterogeneous information network analysis," *IEEE Transactions on Knowledge and Data Engineering*, vol. 29, no. 1, pp. 17–37, 8 2017.
- [27] D. Boyd, S. Golder, and G. Lotan, "Tweet, tweet, retweet: Conversational aspects of retweeting on twitter," in *Proc. of 43th Hawaii International Conference on System Sciences*. IEEE, 2010, pp. 1–10.
- [28] A. Bruns and S. Stieglitz, "Towards more systematic twitter analysis: metrics for tweeting activities," *International Journal of Social Research Methodology*, vol. 16, no. 2, pp. 91–108, 10 2013.
- [29] P. Holme and J. Saramäki, "Temporal networks," *Physics Reports*, vol. 519, no. 3, pp. 97–125, 2012.
- [30] C. Roth and J.-P. Cointet, "Social and semantic coevolution in knowledge networks," *Social Networks*, vol. 32, no. 1, pp. 16–29, 2010.
- [31] I. Hellsten and L. Leydesdorff, "Automated analysis of topic-actor networks on twitter: New approach to the analysis of socio-semantic networks," *CoRR*, pp. 16–29, 2017.