

Contents lists available at [ScienceDirect](https://www.sciencedirect.com)

Information Processing and Management

journal homepage: www.elsevier.com/locate/ipm

Investigating the COVID-19 vaccine discussions on Twitter through a multilayer network-based approach

Gianluca Bonifazi ^{a,1}, Bernardo Breve ^{b,1}, Stefano Cirillo ^{b,1}, Enrico Corradini ^{a,1}, Luca Virgili ^{a,*}

^a *DII, Polytechnic University of Marche, Italy*

^b *DI, University of Salerno, Via Giovanni Paolo II 132, 84084 Fisciano (SA), Italy*

ARTICLE INFO

Keywords:

Twitter
Multilayer network
Social network analysis
Hashtag extraction
AvaxTweets dataset
COVID-19 vaccine

ABSTRACT

Modeling discussions on social networks is a challenging task, especially if we consider sensitive topics, such as politics or healthcare. However, the knowledge hidden in these debates helps to investigate trends and opinions and to identify the cohesion of users when they deal with a specific topic. To this end, we propose a general multilayer network approach to investigate discussions on a social network. In order to prove the validity of our model, we apply it on a Twitter dataset containing tweets concerning opinions on COVID-19 vaccines. We extract a set of relevant hashtags (i.e., gold-standard hashtags) for each line of thought (i.e., pro-vaxxer, neutral, and anti-vaxxer). Then, thanks to our multilayer network model, we figure out that the anti-vaxxers tend to have ego networks denser (+14.39%) and more cohesive (+64.2%) than the ones of pro-vaxxer, which leads to a higher number of interactions among anti-vaxxers than pro-vaxxers (+393.89%). Finally, we report a comparison between our approach and one based on single networks analysis. We prove the effectiveness of our model to extract influencers having ego networks with more nodes (+40.46%), edges (+39.36%), and interactions with their neighbors (+28.56%) with respect to the other approach. As a result, these influential users are much more important to analyze and can provide more valuable information.

1. Introduction

In recent years, with the growth of social networks, we have witnessed the birth of virtual public squares, where each person can express their thoughts to a considerable number of people. Due to the visibility given by the large number of users who populate these platforms, social networks have become a new communication channel. On social media, prominent personalities and newspapers can post news and updates from all over the world (Cauteruccio, Corradini, Terracina, Ursino, & Virgili, 2020; Corradini, Nocera, Ursino, & Virgili, 2021; Willnat & Weaver, 2018).

Unfortunately, the improper use of social platforms can fuel the dissemination of inaccurate, sometimes fake news (Campan, Cuzzocrea, & Truta, 2017). The heterogeneity of people surfing the socials (such as Facebook, Twitter, and Instagram Blanco & Lourenço, 2022; Burel, Farrell, & Alani, 2021), comprehend users without a sufficient level of awareness to distinguish news from reliable sources from misleading and distorted news (Cerruto, Cirillo, Desiato, Gambardella, & Polese, 2022). The latter aims at generating dissent and continuous interactions between users, making these contents bounce from one profile to another, thus

* Corresponding author.

E-mail addresses: g.bonifazi@univpm.it (G. Bonifazi), bbreve@unisa.it (B. Breve), scirillo@unisa.it (S. Cirillo), e.corradini@pm.univpm.it (E. Corradini), luca.virgili@univpm.it (L. Virgili).

¹ All authors contributed equally to the manuscript preparation.

<https://doi.org/10.1016/j.ipm.2022.103095>

Received 12 April 2022; Received in revised form 9 August 2022; Accepted 4 September 2022

Available online 12 September 2022

0306-4573/© 2022 Elsevier Ltd. All rights reserved.

feeding a dense network of disinformation. An important example that provides an idea of the impact of disinformation on social media is related to the COVID-19 pandemic. In particular, debates and controversies have been continuously initiated on the main social networks, and mostly focused on the gravity of the pandemic and the usefulness of the prevention measures adopted (Hung et al., 2020). Furthermore, the arrival of vaccines has provided new elements for discussion between people who support their importance and those who doubt both their effectiveness and safety. This debate became so heated that the users involved were divided into two categories, namely pro-vaxxer and anti-vaxxer (Furini, 2021). In this scenario, the dissemination of incorrect and/or false news is seriously likely to distort the perception of the population on a critical topic, representing a serious threat to world health, and indirectly contributing to the worsening of conditions. Therefore, it is essential to evaluate how the exchange of content on social networks impacts the conception of such situations, which is among the main goals of the social network analysis research field.

The peculiarity of social networks is the possibility to empower users to interact on multiple fronts. For instance, on social networks like Twitter, users may interact through likes, replies, retweets, and mentions, having as a common ground the tweets themselves and the corresponding topics. The information evaluated by considering the projection on different levels of interactions can open different new outcomes in the field of social network analysis, allowing us to consider such highly frequent “cross-relationships” too. For this reason, we define a generic multilayer network-based approach for user–topic analysis on social media. We exploit the flexibility of a multilayer network to map the common user interactions in a social network (e.g., like to a post, retweet, friendship, etc.) to a layer, along with a link connecting the same user over the different layers. Then, thanks to the extraction of gold-standard hashtags from posts, we create a topic layer, which allows us to project the multilayer network into a new one focused on a set of subjects. To demonstrate the effectiveness of the proposed model, we apply it on a Twitter dataset called AvaxTweets dataset (Muric, Wu, Ferrara, et al., 2021), which contains tweets regarding pro-vax and anti-vax opinions. First, we report a thorough study of the single layers composing our multilayer network model, and make a comparison between a multilayer approach and a single network one, which proved the effectiveness of the former. Then, we project the multilayer network according to the pro-vax, neutral, and anti-vax gold-standard hashtags in order to study the characteristics of the most influential users for each line of thought. The main contributions of this paper are:

- The definition and formalization of a generic multilayer network model for representing social networks. The proposed solution empowers the topic analysis in social media and can be adapted to several domains, as discussion analysis, topic analysis, and information dissemination of users. We set up the first set of layers describing the type of users’ interactions and a further layer relating users through the projection over a topic extracted from their posts.
- A comprehensive case study on people’s perception of COVID-19 vaccines through the analysis of a dataset of tweets, made possible through the application of our proposed approach. Results allowed us to deeply study the ego networks of three identified groups of users, namely pro-vaxxers, neutrals and anti-vaxxers.
- A comparison of the multilayer network approach highlights its superiority with respect to a single networks approach. In fact, the application of a multilayer network allowed us to extract influencers with more neighbors and interactions with their neighbors, which could bring more valuable information.

The outline of this paper is as follows: in Section 2, we present the Related Literature. In Section 3, we illustrate our multilayer network-based approach, define its specialization to the Twitter scenario, and introduce the concept of ego network suitable to our case. In Section 4, we report the employed dataset and extract the most relevant hashtags according to pro-vax, anti-vax, and neutral perspectives. In Section 5, we first analyze the single layers composing our multilayer network, and then highlight the differences in terms of knowledge extraction between the multilayer network and single networks approaches. In Section 6, we apply our multilayer network-based approach to study the most influential users for each topic category. In Section 7, we summarize the obtained results and discuss the advantages and limitations of our approach. Finally, in Section 8, we draw our conclusion and describe some possible future works.

2. Related literature

In this section, we survey some major works related to our approach. In particular, we will first present works focusing on the impact of COVID-19 global pandemic through social network analysis. We will then present the related literature employing multilayer networks for performing analysis on different domains.

2.1. Social network analysis on the COVID-19 pandemic

Social networks play an important role in people’s lives, and their usage has increased since the COVID-19 pandemic. In fact, the goal of social networks has evolved, since people do not consider them as simple means of communication, but as real informational platforms used by transmissions of local and global entities (Mourad, Srour, Harmanai, Jenainati, & Arafteh, 2020). Recently, several studies have evaluated the impact of the pandemic on social network (Kovacs, Caplan, Grob, & King, 2021), and e-commerce platforms (Galhotra & Dewan, 2020). For instance, Luo (2021) defines a new approach relying on Deep Recurrent Neural Network (DRNN) to predict online shopping behavior and improve e-business performance starting from the data collected during the COVID-19 pandemic. Another recent study has shown that pandemic fear slightly affects the effectiveness and credibility of e-commerce platforms (Tran, 2021). In fact, despite the economic growth of large companies, such as Amazon and eBay (Pisal, 2021), several

new e-commerce platforms have been published, especially of small companies with the aim of increasing their economy (Bhatti et al., 2020).

Even if the impact of COVID-19 on the economy represents one of the most discussed issues by citizens around the world, other critical problems have been widely debated on social network platforms, such as school closures (Hung et al., 2020), climate (Ward, Xiao, & Zhang, 2020), and vaccines (Latkin et al., 2021).

In Sharma, Zhang, and Liu (2021), the authors present an anonymized dataset of tweets on vaccine disinformation, collected during the lockdown period in 2020 by means of the Twitter streaming APIs. In particular, the study shows a preliminary analysis of the tweet contents over time and provides descriptive statistics of some general characteristics of the corresponding accounts.

A recent study (Vargas, Maier, Vallim, Banda, & Preciado, 2021) explored how COVID-19 has affected people from a psychophysical point of view. The analysis of the tweet led the authors to affirm that vaccination played a fundamental role in reducing people's negativity by promoting their psychological well-being.

The authors of Feng and Zhou (2022) propose a geo-tagged Twitter dataset that can be exploited to perform fine-grained investigations of the public reaction to the COVID-19 pandemic. The analysis of this dataset allowed the authors to perform work (and study) engagement measurements between lockdown and re-open periods. To this end, they compared the volume of tweets posted on workdays and weekends and during specific hours of the day.

In Burel et al. (2021), the authors verify the relationship between the spread of misinformation and the work that the fact-checking organizations are carrying out to stop the proliferation of false claims about the COVID-19 pandemic. The work performed an analysis on 16,521 URLs divided, more or less evenly, between URLs containing misinformation and URLs aiming to do fact-checking. By following the spread of these URLs on Twitter posts, the authors were able to analyze their impact and how they were spread across the social. This analysis showed that, although fact-checking organizations have proved more effective than previous work by the same authors (Burel, Farrell, Mensio, Khare, & Alani, 2020), they are still unable to overcome the impact of the misinformation spread.

2.2. Studies employing multilayer networks

The application of multilayer networks has been proved to be a valuable tool to represent users and their interactions in several domains.

In Türker and Sulak (2018), the authors carried out a study to evaluate the meaningfulness of hashtags within tweets and if the co-occurrence of multiple hashtags is actually linked by a semantic correlation. In fact, it can often happen that, instead of inserting hashtags that reflect the topic discussed in the tweet, the author decides to insert other hashtags completely unrelated to the actual topic. All this is an attempt to increase the visibility of the tweet, which will then be listed under different topics. The study is based on a multilayer network approach characterized by two types of interaction, i.e., the co-occurrence of hashtags and the semantic relationship between them. The results proved that the co-occurrence of hashtags is mainly present when there is also a semantic correlation. However, even a poor presence of semantically unrelated co-occurrences is sufficient for reducing node separation and network diameter in the co-occurrence network layer.

In Singh, Mitra, and Singh (2020), the authors performed a sentiment classification task by transforming tweets into a heterogeneous multilayer network composed of three layers, i.e., the hashtag layers, the keyword layer, and the mention layer. The authors then generated random walk sequences from the multilayer network to evaluate a node's prominence in the network. They did so by extending the random walk employed in the PageRank algorithm. Afterward, both tweets and sequences are embedded and trained in a neural model to output a tweet's final sentiment score. Experimental evaluation performed on a dataset of 42,422 tweets demonstrated that the proposed method outperforms its competitors in identifying the either positive, negative or neutral sentiment of Tweets.

In Pierri, Piccardi, and Ceri (2020), the authors tackle the problem of fake news identification by modeling a multilayer network that puts into correlation an article with its related discussion on Twitter. In fact, for each article, they constructed a multilayer network composed of four different layers, i.e. retweet, reply, quote, and mention. The authors then employed several global network properties for encoding each network layer in a tuple of features. Such features are then concatenated in a single feature vector and employed for training a Logistic Regression model. Experimental results show high accuracy scores proving that a multilayer network-based approach allows simple, un-tuned models, to still achieve accurate classification results.

In Oro, Pizzuti, Procopio, and Ruffolo (2017) the authors introduce SocialAU, based on a multilayer network, to detect topic authoritative social media users by employing the greedy PARAFAC algorithm (Kolda & Bader, 2006). SocialAU, combines topological and context analysis to obtain influential users, exploiting a multilayer network composed of three layers, mapping users, items (i.e., instances of the topic), and keywords of a tweet. An extensive evaluation, performed on both Twitter and Yelp, proved the ability of SocialAU to identify influential users on several topics of interest.

The authors in Nguyen, Wang, Dai, and Dow (2021) investigated the impact of malicious Twitter accounts in a scenario where they could potentially disrupt the fairness of an election. In particular, the authors modeled the political discussion as a multilayer network for spotting the most influential users on social media as well as their communities with the application of several centrality measures. The evaluation was performed through a case study on a political discussion forum in Taiwan, proving the effectiveness of the proposed approach in the identification of influential users, suggesting that their behavior might be associated with malicious activities.

All the manuscripts included in the literature highlight the effectiveness of representing a specific problem as a multilayer network, allowing for a better exploration of the network structure, and fully enabling the analysis of users' interactions in multiple aspects. With respect to the representations presented in the literature, in this work we formalized the interactions on social media by expanding the multilayer representation with projections across set of layers, characterized by nodes of different nature. This enforces the concept of bimodality and further empowers the analysis of interaction on social media, as we will discuss later.




Fig. 1. Schematic representation of our multilayer network model.

3. Model

In this section, we define our multilayer network-based model for user–topic analysis on a social medium. Being our model extremely generic, it can be specialized to investigate how the users interact with each other on any topic and on any social media.

3.1. Definition of the multilayer network model

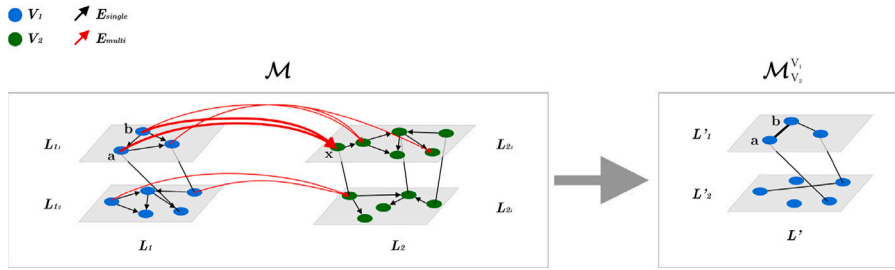
We define a multilayer network $\mathcal{M} = \langle V, E, L \rangle$ (Boccaletti et al., 2014). Here, $V = \{V_1, \dots, V_k\}$ is a set of k sets of nodes. Each set $V_i \in V$ is defined on a type of nodes different from all the other sets $V_j, \forall V_j \in V$, with $i \neq j$.

We define a set $R = \{r_1, \dots, r_h\}$ of relationships. A relationship defines a kind of interaction between nodes. We can now define $L = \{L_1, \dots, L_m\}$ as the set of layers of our multilayer network. In other words, given a set of nodes $V_x \in V$ and a set of relationships $R_y \subseteq R, L_j \in L$ is a set of layers L_{j_i} , each one related to the sets V_x and R_y . A layer $L_{j_i} = \langle V'_x, E_x \rangle$ can be identified as a network. $V'_x \subseteq V_x$ is the set of nodes. There is an edge $e = \langle n_1, n_2, w_{12} \rangle \in E_x$ between two nodes $n_1, n_2 \in V'_x$ if the corresponding nodes interact through a relationship $r \in R_y$. The edge has also a weight w_{12} , which represents the number of interactions between n_1 and n_2 through the corresponding relationship. For each layer in L_j there is one and only one relationship $r \in R_y$, so that $|L_j| = |R_y|$.

Finally, $E = \{E_{single}, E_{multi}\}$ is the set of sets of edges we can find in our multilayer network. In particular, we have a set E_{single} containing edges linking pairs of same type nodes, i.e., a subset of $V_x \times V_x \cdot |L_j|$ for each $V_x \in V$ and $L_j \in L$. So, the set E_{single} contains all edges that link nodes in each layer and the same nodes between different layers of the same type of nodes. For instance, consider a set of nodes $V_1 = \{a, b, c\}$ and 3 layers $L_1 = \{L_{1_1}, L_{1_2}, L_{1_3}\}$. All the possible edges we could find in this case are $\{a, b, c\} \times \{a, b, c\} = \{aa, ab, ac, ba, bb, bc, ca, cb, cc\}$ times 3, the number of layers. The edges between the same nodes, e.g. aa , link the same node between different layers if they exist. On the other hand, E_{multi} is a set of edges between nodes of different types, i.e., a subset of $V_x \times V_y \cdot |L_j|$ for each $V_x, V_y \in V$, with $V_x \neq V_y$, and L_j the set of layers defined for V_x . In this case, given two sets of nodes V_1 and V_2, E_{multi} contains all edges that link nodes of V_1 and V_2 , i.e., all edges that start from the layers of L_1 and end in the ones of L_2 . Note that each edge in a layer and each edge starting from the same layer represent a relationship of R in the multilayer network. Fig. 1 shows an example of the structure of our proposed model. Consider for example to have m sets of nodes $V = \{V_1, \dots, V_m\}$. For each set of nodes V_x , we have a set of layers L_x . In the figure, we can see the first set L_1 , for the nodes V_1 , and the last one L_m , for the nodes V_m . In each layer, all the nodes are linked through the corresponding relationship in that layer. In addition, we have some nodes that belong to more than one layer. So, we have an edge linking the same node in different layers, for all of those nodes. In the figure, for the first two layers L_{1_1} and L_{1_2} , these edges are aa, cc, dd . All the edges in black are the set E_{single} . On the other hand, all the edges in red are the set E_{multi} .

3.2. Knowledge extraction from the multilayer network model

As it happens with multimodal networks, working directly on \mathcal{M} is not straightforward. Indeed, we need to define and use metrics suitable to both the multilayer and bimodal natures of the model. We can work with both the single layers of the network (which represents a portion of the overall scenario) and with projections of the multilayer network.

Fig. 2. Example of a projection of \mathcal{M} .

Given two sets of nodes $V_1, V_2 \in \mathcal{V}$, we define the projection of V_1 on V_2 as a multilayer network $\mathcal{M}_{V_2}^{V_1} = \langle V_1, E', L' \rangle$. This network is defined only on the nodes of V_1 .

For each layer $L_{1j} \in L_1$, where L_1 is the set of layers defined for V_1 , there is a layer $L'_j \in L'$. Given the relationship $r \in R$ of the layer L_{1j} , which defines both the edges of L_{1j} in E_{single} and the edges starting from it in E_{multi} , two nodes a and b of L'_j are linked by an edge $e \in E'$, if both are linked through the same r to the same node of V_2 in \mathcal{M} , i.e., a and b are linked by an edge of E_{multi} defined by the relationship r to the same node x of V_2 . Fig. 2 shows an example of a projection of \mathcal{M} .

3.3. Definition of an ego network in \mathcal{M}

As we will see in the next sections, an important network structure used in our experiments is “ego network”. In the case of a single layer network, this structure is built and used to study the characteristics of a single actor (or node) (Jones & Volpe, 2011). Given a network $L = \langle V, E \rangle$, we can define an ego network of the node $n \in V$ as $\mathcal{E}_n = \langle V_n, E_n \rangle$. $V_n \subseteq V$ is the set of nodes, which contains n and all the nodes directly linked through an edge to n in L . $E_n \subseteq E$ is the set of edges of the ego network. It contains all edges linking the nodes of V_n to n , plus the edges between them.

To the best of our knowledge, in the literature, there is no formal definition of an ego network suitable to our scenario. For this reason, we propose a possible definition in the following.

Given a set of layers L_j , defined for a set of nodes V_j , we define the multilayer ego network of the node $n \in V_j$ as $\mathcal{E}_{\mathcal{M}_n} = \langle V_n, E_n \rangle$. In particular, $V_n \subseteq V$ is the set of nodes that contains n and all nodes that are connected to n in at least a layer of L_j . Two nodes $v_x, v_y \in V_n$ are linked by an edge $e \in E_n$ if there exists an edge between v_x and v_y in at least a layer of L_j .

3.4. Multilayer network model specialization for twitter

In this section, we adapt our general multilayer model to Twitter. Potentially, the model supports multiple types of nodes, $k = |\mathcal{V}|$, so $k \in [2, +\infty)$. In our scenario, we are dealing with user–topic analysis on a social medium, so, we can assume two different types of nodes in our model, $\mathcal{V} = \{V_u, V_t\}$, $k = 2$. The first, V_u is the set of user nodes, where each node represents a user in the social network. The latter, V_t is the set of topic nodes, where each node represents a discussed topic. The number of nodes in V_u depends on how many users are considered in the analysis, while the number of topics depends on the discussion modeled. In our case, we are dealing with the discussion on vaccines, with three different opinions about them, pro-vax, anti-vax, and neutral. We have a user node for each user who made at least one of the possible interactions on Twitter, and have a topic node for each hashtag used by pro-vaxxers, anti-vaxxers, and neutral users.

Accordingly, to the possible interactions on Twitter, R contains the following relationships:

- “Like” (i.e., r_l): when a user likes the tweet of another user;
- “Reply to” (i.e., r_r): when a user replies to the tweet of another user;
- “Retweet” (i.e., r_{rt}): when a user retweets the tweet of another user²;
- “Mention” (i.e., r_m): when a user mentions another one in a tweet;
- “Found together” (i.e., r_f): when a topic is found together with another topic in the same tweet.

So, $\mathbf{L} = \{L_u, L_t\}$, where L_u is the set of layers associated to V_u , while L_t is the set of layers associated to V_t . Formally speaking, L_u and L_t are defined as:

- $L_u = \{L_l, L_r, L_{rt}, L_m\}$;
- $L_t = \{L_f\}$

² Retweeted tweets may also contain quotes.

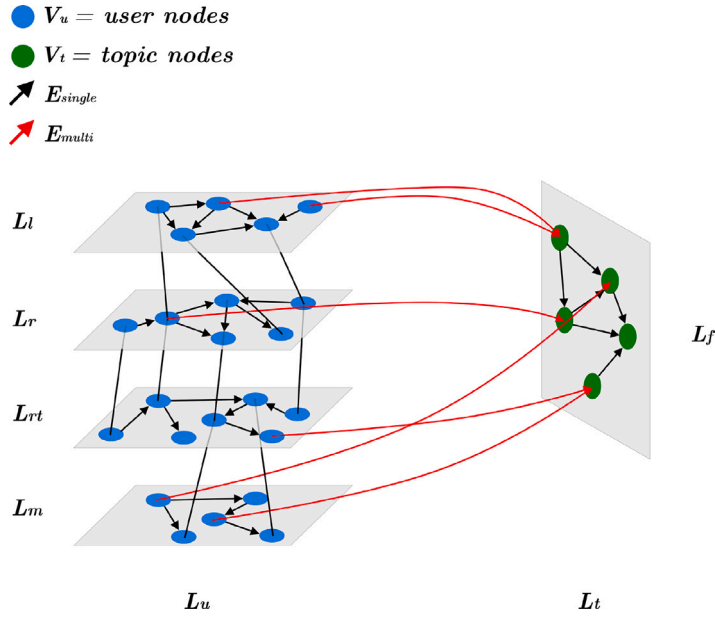


Fig. 3. Example of our specialization of \mathcal{M} .

Where $L_x = \langle V_y, E_x \rangle$ is the layer associated to the relationship x and the type of nodes y .

The set E_{single} contains all edges linking the nodes of the same type, i.e., all edges between nodes of V_u and all edges between the nodes of V_t . Plus, it contains all edges linking the same nodes in multiple layers, as we have seen in Section 3.1. In this specialization, this is true only for the nodes of V_u , as L_t has only one layer. On the other hand, E_{multi} contains the edges linking nodes of V_u to nodes of V_t .

Fig. 3 shows a graphical simplification of our specialization of \mathcal{M} to Twitter. First of all, blue nodes are user nodes. The figure shows a subset of them, and how they could be linked on each layer and between layers. The same for green nodes, i.e., topic nodes. All black edges belong to E_{single} set, i.e., all edges between blue nodes, all edges between green nodes, and edges between the same nodes in different layers of L_u . The red edges belong to E_{multi} set, which are the edges between nodes of V_u and nodes of V_t .

4. Overview of the dataset

In this section, we provide an overview of the datasets adopted in this study and analyze the hashtags employed by users. In the first part, we describe the structure of the dataset, also highlighting occurrences and correlations between the most frequent hashtags in the set of considered tweets. In the second part, we analyze the contents of the tweets associated with the relevant hashtags, i.e., gold-standard hashtags (Di Giovanni et al., 2021), to indicate whether they represent a positive or negative perspective on the COVID-19 vaccine debate.

4.1. Dataset description

The spreading of the COVID-19 pandemic and different lockdowns imposed by the public governments have led to a strong usage of social media. Among the different social networks, Twitter proved to be a tool to rapidly communicate with citizens during public health crises aiming to inform, boost morale, and even raise awareness by encouraging active participation. As a matter of fact, during the crucial phases of the pandemic, government leaders, and virologists have continuously shared information about the treatments and law acts to fight the spread of COVID-19. In fact, 88.9% of global leaders have verified and active Twitter accounts, with more than 85 million users that have followed their “Informative”, “Morale-boosting”, and “Political” tweets (Rufai & Bunce, 2020). For our analysis, we focused on the AvaxTweets dataset (Muric et al., 2021), representing the largest dataset of tweets collected between October 2020 and April 2021 on the topics of COVID-19 vaccines. The period considered within the dataset is of timely relevance as it was the most crucial period for COVID-19 vaccine discussions. In addition, it is large enough to fully demonstrate the potential of our approach. The dataset contains two collections of tweets extracted from the historical account-level data collection of Twitter and the streaming keyword-centered data collection. In this study, we consider the streaming data collection created by using the snowballing sampling technique in DeVerna et al. (2021). This strategy initially required the definition of a small set of keywords related to strong vaccination hesitation (such as *vacciniskill* and *vaccinodamage*), which was subsequently enriched with other similar keywords extracted from the first set of tweets.

Table 1
Statistics about the AvaxTweets dataset.

	Streaming collection	
	Before hydrating	After hydrating
Number of tweets	1,832,333	1,095,621
Number of accounts	719,652	451,584
Verified accounts	9032	8736
Average tweets per account	2.5	1.75
Accounts with location	5661	1632
Most recent tweet	2021-04-21	2021-04-21

Table 1 shows statistics about the AvaxTweets dataset at the time of writing this article (early December 2021). The authors only released tweet IDs in order to comply with Twitter's Terms of Service and this required rehydrating the tweets for retrieving their contents using the Twitter API. As we can see, the resulting streaming data collection contains over 1.8 million tweets from over 700K unique accounts. However, after the operation of re-hydration of the dataset according to the strategy defined in Muric et al. (2021), the number of tweets was reduced to 1,095,621, of which 1,078,613 were tweets of unverified accounts, while 17,008 of verified accounts. This reduction in the number of tweets may be due to the fact that Twitter has blocked and/or removed many accounts that probably did not respect the platform's policies or that had been identified as fake. A tweet could also be no longer available if the author deletes it, or changes its privacy settings. For this reason, we choose to use only the tweets shared by verified accounts (i.e., 8736 accounts), aiming at performing an accurate analysis without considering bots or fake accounts. Although the operation of rehydrating the dataset allowed us to obtain information concerning retweets, mentions, and replies to tweets, it has been necessary to design a web crawler for extracting the likes of each tweet. In fact, the official Twitter APIs limit the number of likes that can be extracted for each tweet to 100, which could affect both the amount of data and the proposed experimentation.

The extended dataset has been adopted in our case study to evaluate the effectiveness of the proposed multilayer network model. Nevertheless, in the following sections, we further analyze the tweets and the corresponding hashtags in order to perform a preliminary analysis of their contents.

4.2. Preliminary analysis of hashtags

Starting from the tweets of the verified accounts, i.e., 17,008 tweets, we have performed cleaning operations of the contents by removing all special characters and/or emojis, yielding the standardization of the tweet syntax. This operation allowed us to remove any encoding errors of the characters in the tweets recovered after the re-hydration process and standardize the syntax of the hashtags, which often are syntactically different due to the use of uppercase and/or lowercase letters, such as in the case of "Covid19" and "covid19" or "vaccine" and "VACCINE". It is important to notice that, although hashtags have been standardized in their syntax, their extraction from tweets required the adoption of a specific regular expression to identify and collect them. After standardizing the content of tweets, it was possible to identify the most common hashtags used in the tweets and analyze their frequency in the dataset.

Table 2 reports the occurrences and the frequencies of the top 30 hashtags employed from verified accounts of the dataset. The occurrence values represent the total number of times a hashtag appears, while the frequency is the number of occurrences of each hashtag with respect to the tweets shared by verified accounts. It is important to notice that only 2245 of 17,008 tweets of the verified accounts contain at least one hashtag, whereas 14,763 tweets do not contain any hashtag. As shown in Table 2, we have defined three different frequency values that represent the frequency of each hashtag with respect to: (i) the number of all the hashtags used in the tweets (i.e., N_X); (ii) the number of tweets that contain at least one hashtag (i.e., N_Y), and (iii) the number of all tweets in the streaming data collection shared by verified accounts (i.e., N_Z). Let N be the number of occurrences of each hashtag in the streaming data collection, the frequencies are defined as follows:

$$F_1 = \frac{N \cdot 100}{N_X} \quad F_2 = \frac{N \cdot 100}{N_Y} \quad F_3 = \frac{N \cdot 100}{N_Z} \quad (1)$$

As we expected, there are a large number of tweets that contain general hashtags that refer to Covid or family ties, such as "#covid19", "#parenting", "#family", and so forth. However, even if several hashtags show pro-vax sentiments, such as "#vaccination", "#thisisourshot", and "#getvaccinated", there are several hashtags that show strong anti-vaccine sentiments, such as "#vaccinefraud", "#antivaccine", and "#novaccine". The latter represents about 6.5% of the top 30 hashtags used in the tweets of the verified accounts, showing that there are many people who spread anti-vaccine sentiments. To further analyze how these sentiments have affected collective thought, we show the correlation analysis between hashtags of the tweets in the streaming data collection (Fig. 4). In particular, to identify the correlation between the hashtags collected from the tweets, it was necessary to turn the text into a numeric form, by transforming the hashtags used in the text into a vector form. To do this, we define a binary vector for each tweet whose dimensions are equal to the number of hashtags involved in the evaluation. Each vector contains several elements, whose value is equal to 1 if the hashtag is contained in the text of the tweet, and 0 otherwise. For example, let us consider the hashtags "#vaccinated", "#family", "#maskup", and "#day". So the tweet:

"Protect those that cannot protect themselves #MaskUp #Vaccinated"

Table 2
Top 30 hashtags in streaming data collection.

Hashtag	N	F ₁ (%)	F ₂ (%)	F ₃ (%)	Hashtag	N	F ₁ (%)	F ₂ (%)	F ₃ (%)
#covid19	572	40.25	25.48	3.36	#vaccination	32	2.25	1.43	0.19
#mybodymychoice	133	9.36	5.92	0.78	#vaccinefraud	30	2.11	1.34	0.18
#vaccine	117	8.23	5.21	0.69	#cdc	29	2.04	1.29	0.17
#vaccines	89	6.26	3.96	0.52	#uae	28	1.97	1.25	0.16
#covidvaccine	87	6.12	3.88	0.51	#breaking	27	1.90	1.20	0.16
#vaxxed	83	5.84	3.70	0.49	#informedconsent	26	1.83	1.16	0.15
#parentalrights	72	5.07	3.21	0.42	#doctorspeakup	25	1.76	1.11	0.15
#coronavirus	71	5.00	3.16	0.42	#vaccineswork	24	1.69	1.07	0.14
#family	57	4.01	2.54	0.34	#mybodyismyown	24	1.69	1.07	0.14
#parenting	57	4.01	2.54	0.34	#thisourshot	23	1.62	1.02	0.14
#covid	49	3.45	2.18	0.29	#unvaccinated	22	1.55	0.98	0.13
#vaccinated	48	3.38	2.14	0.28	#pfizer	21	1.48	0.94	0.12
#antivaccine	38	2.67	1.69	0.22	#getvaccinated	21	1.48	0.94	0.12
#longCOVID	34	2.39	1.51	0.20	#astrazeneca	18	1.27	0.80	0.11
#israel	32	2.25	1.43	0.19	#novaccine	14	0.99	0.62	0.08




Fig. 4. Correlation matrix of all the hashtags that appear in at least 10 different tweets of the streaming data collection.

will be represented as [1, 0, 1, 0] in its vector form.

In our analysis, we consider 54 hashtags representing the hashtags that appear in at least 10 different tweets of the streaming data collection. Thus, for each of them, we have defined a binary vector of size 54 to obtain the rows of a new binary matrix. In this way, we can calculate their correlation by using the Pearson coefficient that returns a value in the range between -1.0 and 1.0 for each hashtag pair. The Pearson measure is one of the most used coefficients, since it measures the degree of the association involving

linear related variables and permits to remove the prejudices of users (Jeyasudha et al., 2021; Kalamatianos, Mallis, Symeonidis, & Arampatzis, 2015).

As we expected, there are several strong correlations between hashtags in the same domain, such as “#family” with “#parental-rights”, and “#parenting”, or “#vaccineswork” with “#covidvaccines”. However, the analysis reveals some other strong interesting correlations, such as the ones between “#vaccinefraud” and “#cybercrime”, and “#unvaccinated” and “#children”.

This could mean that the tweets contained in the streaming data collection discussing vaccines and identifying them as fraud have a strong link to cybercrime. This is probably due to the fact that many people who have joined the vaccination campaign and have used applications to track infections (such as NHS COVID-19 (Wymant et al., 2021), Immuni (Bosco & Cvajner, 2021), etc.), have often shared their personal data with institutions and/or governments. In this scenario, many people have opened several debates on personal data privacy since they were not sure of the effectiveness of these applications. Similarly, when tweets discuss unvaccinated people, they often highlight the problem of vaccinating children. In fact, during the period in which the dataset has been created, there was an intensive debate on the problem of vaccinating children, since several people have discussed the ethics of the researchers when testing vaccines on children. An example of a tweet is reported in the following:

“The reason we do not do studies to compare groups of #vaccinated vs. #unvaccinated #children is a very simple reason of ethics. When vaccines are available to prevent against diseases, it is unethical for researchers to assign kids to a study’s “control group” without vaccines.”

In fact, several people have opened debates on the ethics of the researchers when comparing the effects of the symptoms of COVID-19 in vaccinated and unvaccinated children. On the other hand, many parents have complained about the lack of effective testing to prove the vaccine’s effectiveness on their children.

These types of debates have also been intrinsically mapped in correlation analyses, where we can see that correlations are sometimes more or less strong in relation to the number of tweets that have dealt with certain types of discussions.

4.3. Preliminary analysis of the tweet contents

The creation of an effective vaccine against COVID-19 has been one of the biggest challenges of recent years. In fact, the vaccination campaign against COVID-19 has been considered a social and economic challenge. Governments had to establish a distribution plan of vaccine doses in a short time aiming at restarting the economy of each country.

Our analysis starts by considering the hashtags used within all considered tweets. Starting from these, we have defined different sets of hashtags, also known as gold-standard hashtags (Di Giovanni et al., 2021), to identify the presence of positive or negative opinions in the tweets concerning the vaccine debate. Although the use of gold-standard hashtags was defined with the aim of identifying only two types of tweets, such as for or against the vaccine (Abu-Raddad, Chemaitelly, & Butt, 2021), it was necessary to identify the third category of gold-standard hashtags, i.e., neutral. The latter represents all hashtags that do not show a clear opinion regarding vaccines and/or that do not concern the vaccine debate.

For example, the hashtags “#quarantine”, “#covid19”, and “#longhailer” are neutral gold-standard hashtags, since the first two do not express a clear opinion regarding the debate, while the others have no reference to vaccines.

It is important to notice that, the identification of tweets containing positive or negative opinions on the vaccines is not the main contribution of this study. In fact, this type of analysis aims to study the main characteristics of the dataset and investigate the peculiarities of the most influential users for each line of thought through a multilayer network.

To make our analysis, we extracted all the hashtags contained in the tweets of the verified accounts, i.e., 1421 different hashtags, and we manually identified the three sets of gold-standard hashtags. In order to obtain a proper interpretation of the hashtags, we employed a cross-checking strategy for their classification. In particular, we divided the 1421 hashtags into two sets, each of them assigned to a pair of authors, while the fifth author was in charge of solving any disagreement. In fact, each author of a pair proceeded by providing an individual classification of the hashtags assigned to him and, whether the pair disagreed, the nature of the hashtag was assigned by the fifth author. In order to obtain an objective classification of the hashtags, each one of these has been evaluated by each author of the paper. Based on the evaluation results, we have identified the three classes of gold-standard hashtags using the concept of majority voting. For the sake of clarity, Table 3 only shows the top 20 gold-standard hashtags for the pro-vax, neutral, and anti-vax categories.

Starting from the sets of gold-standard hashtags, we iteratively analyze the content of the tweets, keeping track of the occurrences of each hashtag. Specifically, for each of them, we computed the number of occurrences considering: (i) the entire set of tweets (i.e., N); (ii) the set of original tweets, i.e., the tweets that are not retweets (i.e., N_o), and (iii) the set of retweets (i.e., N_{RT}). Moreover, we report the frequencies of each hashtag related to tweets that are or are not retweets (N_o and N_{RT} , respectively).

From Table 3, we can see that most of the gold-standard hashtags are contained in original tweets meaning that people with a specific opinion tend to directly share it by writing original tweets. Nevertheless, for several anti-vax gold-standard hashtags, the frequency of hashtags contained in retweets is higher than the frequency of the original tweets. This could mean that people tend to express anti-vax opinions by re-posting the content of other people or answering previous tweets.

After the definition of gold-standard hashtags, we iteratively evaluated each tweet by counting the number of occurrences of the different categories of hashtags (i.e., pro-vax, anti-vax, and neutral). In particular, a tweet with a higher number of pro-vax gold-standard hashtags was considered as a tweet with a pro-vax opinion. Similarly, a tweet with a higher number of anti-vax gold-standard hashtags was considered as a tweet with an anti-vax opinion. However, concerning the tweets with a higher number

Table 3
Statistics of the tweets related to the top 20 gold-standard hashtags.

	Gold-standard hashtags	N	N_o	N_{RT}	F_o (%)	F_{RT} (%)	Gold-standard hashtags	N	N_o	N_{RT}	F_o (%)	F_{RT} (%)
pro-vax	#vaxxed	83	72	11	0.87	0.13	#protectchicago	9	5	4	0.56	0.44
	#vaccinated	44	37	7	0.84	0.16	#vaccinepassports	9	4	5	0.44	0.56
	#doctorssspeakup	25	13	12	0.52	0.48	#covidvaccines	6	4	2	0.67	0.33
	#vaccineswork	24	22	2	0.92	0.08	#vaccinationrules	6	6	0	1.0	0.0
	#thisisourshot	23	17	6	0.74	0.26	#covaxin	6	3	3	0.5	0.5
	#getvaccinated	21	18	3	0.86	0.14	#factcheck	5	3	2	0.6	0.4
	#wearamask	18	15	3	0.83	0.17	#vaccin	4	1	3	0.25	0.75
	#maskup	15	7	8	0.47	0.53	#vaccineequity	4	4	0	1.0	0.0
	#vaccinessavelives	15	13	2	0.87	0.13	#washyourhands	4	4	0	1.0	0.0
	#vaccinepassport	13	13	0	1.0	0.0	#ableg	3	3	0	1.0	0.0
anti-vax	#mybodymychoice	133	97	36	0.73	0.27	#arrestbillgates	6	4	2	0.67	0.33
	#antivaccine	38	32	6	0.84	0.16	#antivaxxers	6	4	2	0.67	0.33
	#vaccinefraud	30	22	8	0.73	0.27	#vaccinefailure	6	5	1	0.83	0.17
	#mybodyismyown	24	19	5	0.79	0.21	#novaccineforme	5	3	2	0.6	0.4
	#unvaccinated	22	20	2	0.91	0.09	#antivax	4	1	3	0.25	0.75
	#novaccine	14	10	4	0.71	0.29	#exposebillgates	4	4	0	1.0	0.0
	#cdnpoli	9	8	1	0.89	0.11	#billgatesvaccine	3	3	0	1.0	0.0
	#learntherisk	9	9	0	1.0	0.0	#researchanddestroy	3	1	2	0.33	0.67
	#medicalfreedom	8	6	2	0.75	0.25	#plandemic	3	1	2	0.33	0.67
	#billgates	8	8	0	1.0	0.0	#scamdemic	2	1	1	0.5	0.5
neutral	#covid19	561	378	183	0.67	0.33	#astrazeneca	18	4	14	0.22	0.78
	#vaccine	116	82	34	0.71	0.29	#sarscov2	13	10	3	0.77	0.23
	#vaccines	88	79	9	0.9	0.1	#covid_19	13	9	4	0.69	0.31
	#covidvaccine	86	66	20	0.77	0.23	#children	13	13	0	1.0	0.0
	#coronavirus	71	61	10	0.86	0.14	#covid19vaccine	12	11	1	0.92	0.08
	#covid	49	39	10	0.8	0.2	#onpoli	11	10	1	0.91	0.09
	#longcovid	33	13	20	0.39	0.61	#coronavaccine	10	7	3	0.7	0.3
	#vaccination	32	28	4	0.88	0.12	#healthcare	9	9	0	1.0	0.0
	#pfizer	21	17	4	0.81	0.19	#health	7	6	1	0.86	0.14
	#pandemic	18	16	2	0.89	0.11	#essentialworkers	7	0	7	0.0	1.0

of gold-standard neutral hashtags, i.e., tweets with neutral opinions, it was also necessary to analyze the number of pro-vax and anti-vax hashtags to understand the nature of the tweets. In fact, a neutral tweet can also be considered pro-vax or anti-vax according to the respective gold-standard hashtags occurrences. The results of the gold-standard hashtag analysis can be summarized as follows: 1102 tweets with pro-vax opinions, 293 with anti-vax opinions, and 850 neutral tweets.

This preliminary analysis provides a general overview of the opinions discussed by users in the tweets and allows us to evaluate the existence of posts containing both pro-vax and anti-vax opinions. As we have discussed above, there is a fairly even distribution between pro-vax and anti-vax tweets, which makes the dataset suitable for our study. Moreover, it is important to notice that this type of analysis is limited to the evaluation of a single user, without considering the impact of the shared opinions on the other users on the social network. In fact, this type of analysis does not take into account the user interactions on a topic, and so it does not consider how the content is perceived by other users and followers. For these reasons, in the following sections, we exploit the knowledge extracted so far and investigate the interactions between users thanks to our multilayer network-based approach.

5. Analysis of user interactions

In this section, we investigate the properties of our multilayer network model tailored to the Twitter scenario. To this end, in Section 5.1 we report the analysis of the single networks related to user relationships of R , i.e., “Retweet” (r_{rt}), “Reply To” (r_r), “Like” (r_l) and “Mention” (r_m). Then, in Section 5.2, we employ our model and compare it with an approach leveraging only the single networks composing our multilayer network \mathcal{M} . In order to perform this investigation, we do not consider the layer of topic L_t , but only the layers L_u , since now we point out the general characteristic of our approach regardless of the set of topics.

5.1. Analysis of single networks

As reported in Section 3.4, our model defines networks related to a set of relationships R . As for users, the considered interactions in Twitter are r_{rt} , r_r , r_l and r_m . To investigate these four networks, we started by reporting their descriptive statistics, such as the number of nodes and edges, density, clustering coefficient, number of connected components, and size of the maximum connected component in Table 4.

From the analysis of Table 4, we can observe that the number of nodes is the same for all networks, while the number of edges differs a lot. In particular, L_r is the least connected one, as we can see from both the number of edges and density. On the other hand, the L_m layer is the most connected one, with a density much higher than all the other layers. The “Mention” relationship (r_m) connects many users together and creates larger connected components w.r.t. other forms of interactions. This peculiarity helps

Table 4
Single networks descriptive analysis.

Network	Nodes	Edges	Density	Clustering Coeff.	# Conn. Comp.	Max. Conn. Comp.
L_{r_l}	8736	3429	9.2e-5	4.8e-3	5690	2128
L_r	8736	354	0.9e-5	0.6e-3	8388	70
L_l	8736	762	2.0e-5	87e-3	7884	501
L_m	8736	4885	53e-5	67e-3	3521	3047

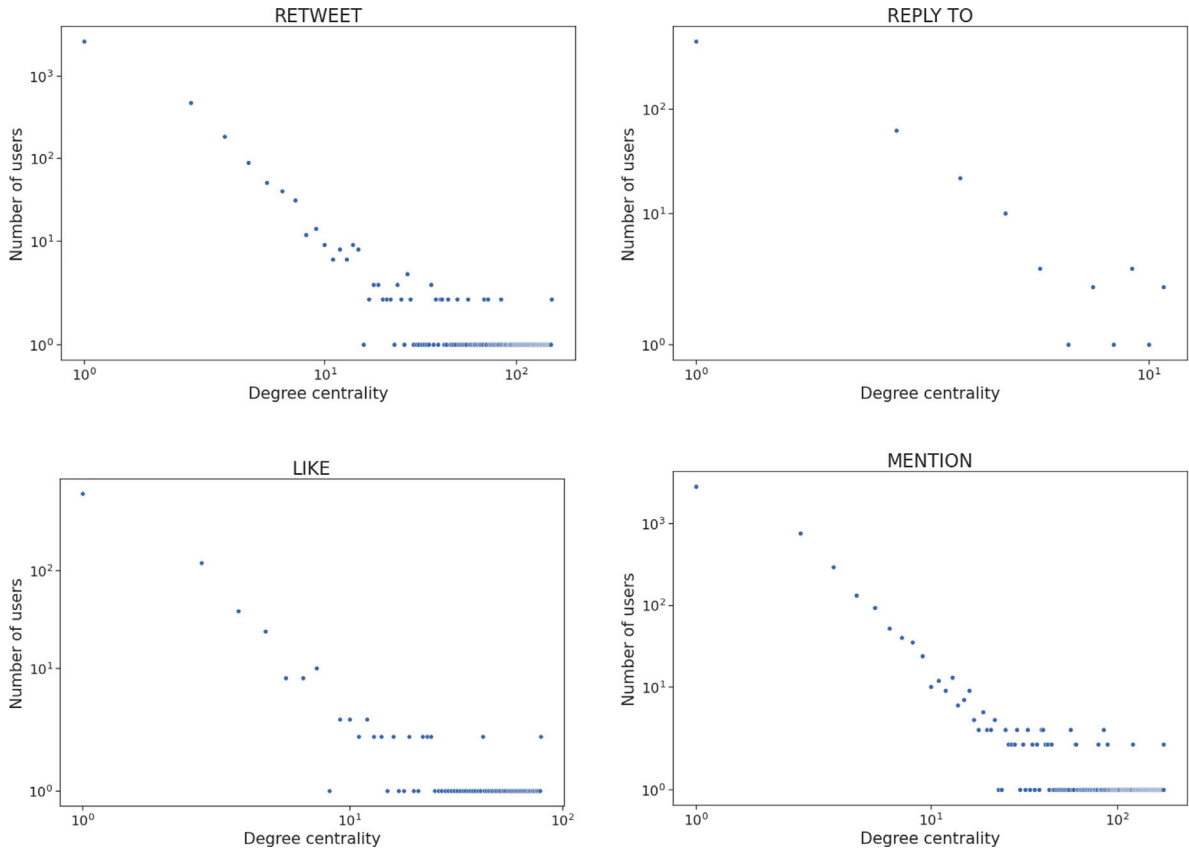


Fig. 5. Degree centrality distributions for “Retweet” (L_{r_l}), “Reply To” (L_r), “Like” (L_l) and “Mention” (L_m) networks.

users to spark many discussions (i.e., tweets) with their followers. The L_{r_l} network has similar statistics to those of L_m , but with fewer edges, smaller clustering coefficients, and smaller size of the maximum component. The behavior seems similar, even if it connects fewer users. Furthermore, it is worth noting the high clustering coefficient and low density of the L_l network. It shows that the few connected users tend to join together in closed triads. Roughly speaking, the user that likes each other posts highlight a certain level of trust in each other and mutual approval of their posts. However, it is evident that the “Like” relationship (r_l) is one of the least present ones since it seems that the verified users tend to use other communication ways to express their ideas and opinions. Finally, we can observe that the L_r network has the lowest number of edges, density, clustering coefficient, and max connected component. Indeed, it has fewer users connected and the maximum connected component is small w.r.t. the number of nodes.

As we have previously observed, these descriptive statistics highlight the different communication ways between users on Twitter. “Retweet” (r_{r_l}) and “Mention” (r_m) relationships are more widespread than “Reply To” (r_r) and “Like” (r_l), and so they are more employed by the verified users when they post a tweet. On the other hand, r_r and r_l have a low impact on verified user interactions, which points out that these users tend to like posts and reply to them very rarely.

Then, we focused our attention on the ability of the users to drive discussion for each of the considered relationships. To this end, in Fig. 5, we report the distributions of the degree centrality of the users for each network.

From Fig. 5, we can observe that these distributions follow a power law (Alstott, Bullmore, & Plenz, 2014). Generally speaking, this type of distribution is described through the formula: $p(x) = c \cdot x^{-\alpha}$, where c is a constant and α defines the steepness of the curve (the higher α , the steeper the power law). Power laws are heavy-tailed distributions, hence they usually contain few very large values compared to the lowest ones of the distributions. Power laws are endemic in some settings, such as social networks, web

Table 5
Power law distribution parameters of the degree centrality distributions.

Network	α	δ
L_{rt}	1.655	0.109
L_r	1.348	0.102
L_l	1.633	0.123
L_m	1.549	0.060

Table 6
Intersection of Top-800 users according to the degree centrality of the L_{rt} , L_r , L_l and L_m networks.

Top-800 users	Number of users	Percentage of users
$Retweet \cap Reply To \cap Mention$	143	17.87%
$Retweet \cap Reply To \cap Mention \cap Like$	67	8.37%

graphs, computer networks, collaboration networks, and so on. Twitter is no different since we can observe power law distributions characterizing user engagements on this social network, as reported in Fig. 5. We can observe that the majority of users have a low degree centrality value, while very few of them have a high degree centrality. In our scenario, this phenomenon implies that very few verified users are very active on Twitter (in terms of retweets, replies, likes, and mentions), while many of them are less way present.

In order to quantitatively evaluate the power law distributions of L_{rt} , L_r , L_l and L_m , we compute the α and δ parameters (Alstott et al., 2014) and report them in Table 5. δ is the lowest Kolmogorov–Smirnov distance between the original distribution and the best model that fits it. The lower δ , the more accurate the power law fit is Alstott et al. (2014).

From Table 5, we can observe that all the distributions are power laws, since $\alpha > 1$ and δ is low. These degree centrality distributions point out that the networks are scale-free networks (Barabási & Bonabeau, 2003). Considering the steepness of the power law distributions and the Pareto principle, which states that 80% of outcomes are due to 20% of causes (80–20 or 90–10), we decided to consider the first 800 (approximately 10%) users with the highest degree centrality as the most influential ones, and consequently derive the most important characteristics of the networks.

Following this reasoning, we select the Top-800 users according to the degree centrality of the considered relationships. The next step is to verify if these sets of users are overlapping or not. Possible overlap between two or more sets means that the same user is influential in different relationships, which highlights a higher ability to drive discussion. We compute the intersections of the top users of each network with and without considering the “Like” relationship due to the fact that it was not present in the AvaxTweets dataset but was added by us afterward. The computation of these two intersections is useful to observe the impact of the “Like” perspective. The results are reported in Table 6.

From Table 6, we can observe that the sizes of the intersections are not high. Indeed, starting from the four sets of Top-800 users, we obtain an intersection of 143 users for $Retweet \cap Reply To \cap Mention$ (i.e., 17.87% of the initial set), and 67 for $Retweet \cap Reply To \cap Mention \cap Like$ (i.e., 8.37% of the initial set). This highlights that few users are influential in all these communication ways, which could also mean that each verified user chooses one way over another to convey his/her tweets.

Once we have the most influential users from all the perspectives, we study their behavior and their neighborhood on Twitter. To this end, we compute the ego networks of the $Retweet \cap Reply To \cap Mention$ users and extract the corresponding density and clustering coefficients. In Fig. 6, we report the obtained results.

From Fig. 6, we obtain interesting insights. For instance, it is worth noting that ego networks of “Like” and “Reply To” tend to have a low clustering coefficient, even if in some cases the density is high. The high density and low clustering coefficient mean that the interactions between users are only one way since they reply or like a tweet but do not involve further users. This behavior does not foster a discussion since there are no more than two people involved in these interactions (e.g. users reply to a tweet but do not receive an answer from a third user to close the triad). However, this could highlight that verified users do not want to get involved in already opened discussions, agreeing with the expressed opinion or not, and so they prefer to convey their ideas in other ways. On the other hand, “Retweet” and “Mention” have a different distribution of density and clustering coefficient values of the ego networks compared to “Like” and “Reply To”. In the “Retweet” and “Mention” cases, we can observe that many ego networks are dense and have a medium-high clustering coefficient. This leads us to think that these two relationships are able to connect users well on Twitter since there are many connections in an ego network and these connections tend to form many triads. The capability to close triads fosters the discussions since we can observe mentions and retweets to a user and the corresponding answers back.

Surely, in this context, the analysis of users’ behaviors could be useful to find and analyze different information diffusion patterns, in order to classify them. However, this would require extra base knowledge, like the temporal evolution of the discussion, which leads to a possible future expansion of our approach.

5.2. Analysis of our multilayer network model

In this section, we leverage our multilayer network approach to extract further and different knowledge w.r.t. the single networks analysis. Recall that, the multilayer network \mathcal{M} is composed of five layers, four for the relationships considered in the previous

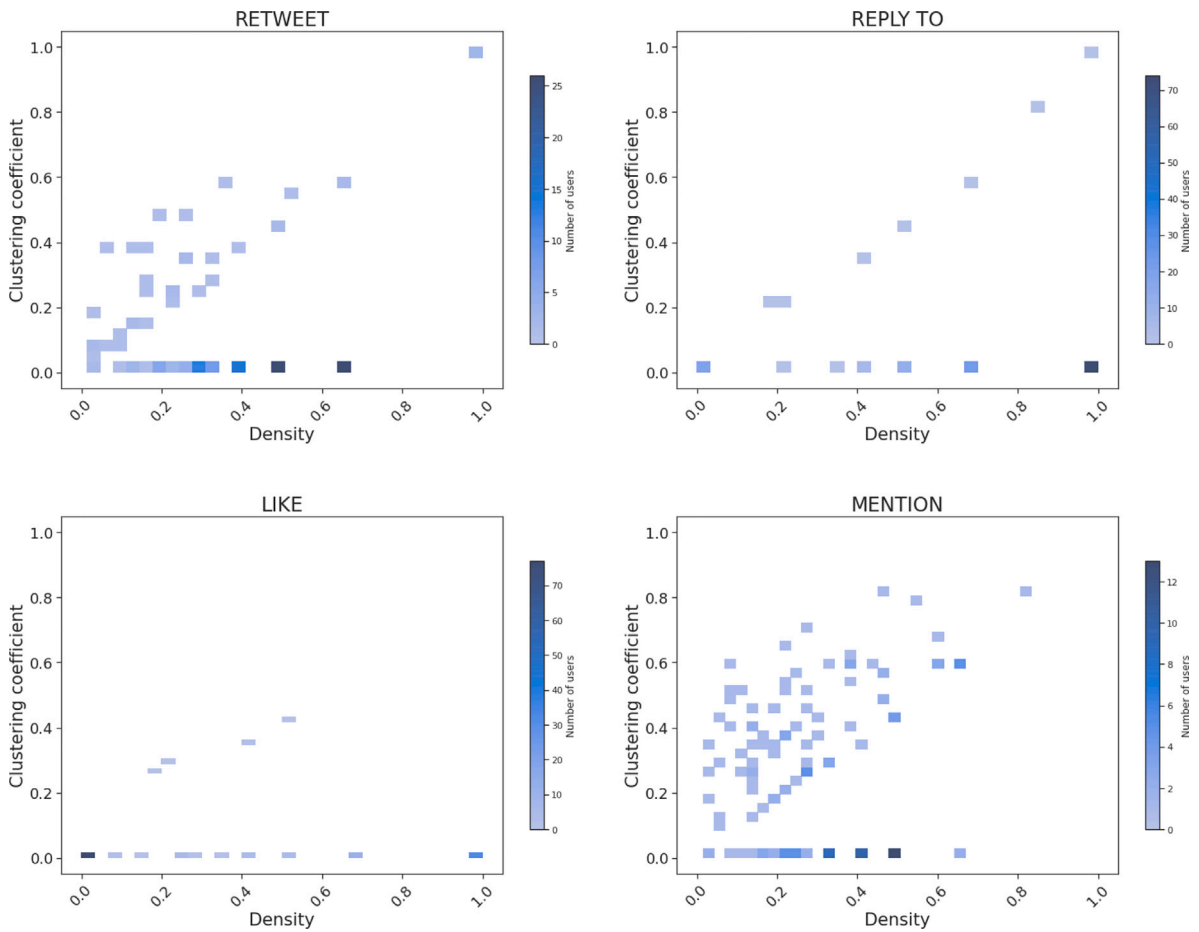


Fig. 6. Density and clustering coefficient of the ego networks of the $Retweet \cap Reply To \cap Mention$ users for each considered relationship.

sections, and one for the topics. This architecture allows us to consider all the perspectives at once and it allows us to project both users and topics on each other. This is useful to highlight the interactions between users starting from a topic and/or investigate the co-occurrences of the topics starting from the users and their relationships. Since not considering the topics so far, the number of nodes and edges of the multilayer network for each layer are the same as presented in Table 4.

As a first analysis, we compute the degree centrality of the multilayer network users. For each user, the degree centrality considers the contributions of each layer, and so it consists of the mean of the degree centrality across the L_{rt} , L_r , L_l , and L_m networks (Boccaletti et al., 2014). The results are reported in Fig. 7.

From Fig. 7, we can observe that the distribution of the degree centrality follows a power law distribution, which is also confirmed from the $\alpha = 2.518$ and $\delta = 0.038$ parameters. This power law is far steeper than the previous ones, since $\alpha > 2$. Similar to the previous case, we can study the most influential users to observe the important characteristics of the overall network.

Hence, we first collect the top users of the multilayer network according to the degree centrality, and then verify if there is an overlapping with the most influential users extracted in the single networks analysis. Following the same reasoning of Table 6, we compute the intersections with and without the “Like” relationship since it was added to the dataset afterward. In order to make a fair comparison in terms of the dimension of the sets of users, we extract the same number of top users from the multilayer network as the $Retweet \cap Reply To \cap Mention$ and $Retweet \cap Reply To \cap Mention \cap Like$ intersections reported in Table 6. For this reason, when comparing with the $Retweet \cap Reply To \cap Mention$ case, we extract the first 143 top users according to the multilayer network degree centrality (i.e., $MultilayerTop - 143$), while when comparing with the $Retweet \cap Reply To \cap Mention \cap Like$ case, we consider the first 67 top users according to the multilayer network degree centrality (i.e., $MultilayerTop - 67$).

The results are reported in Table 7.

From the analysis of Table 7, we can observe that the two approaches do not consider the same set of users as the most influential ones. Indeed, in the first case ($Retweet \cap Reply To \cap Mention \cap MultilayerTop - 143$), we have an overlap of 24 users (i.e., 16.78% of the initial users), while in the latter case ($Retweet \cap Reply To \cap Mention \cap Like \cap MultilayerTop - 67$), we can observe that 55 users are overlapping (i.e., 82.09%). The reason could be related to the fact that the multilayer network approach considers the contribution of each layer to the resulting degree centrality of a node, which is different from the procedure of the single layer analysis.

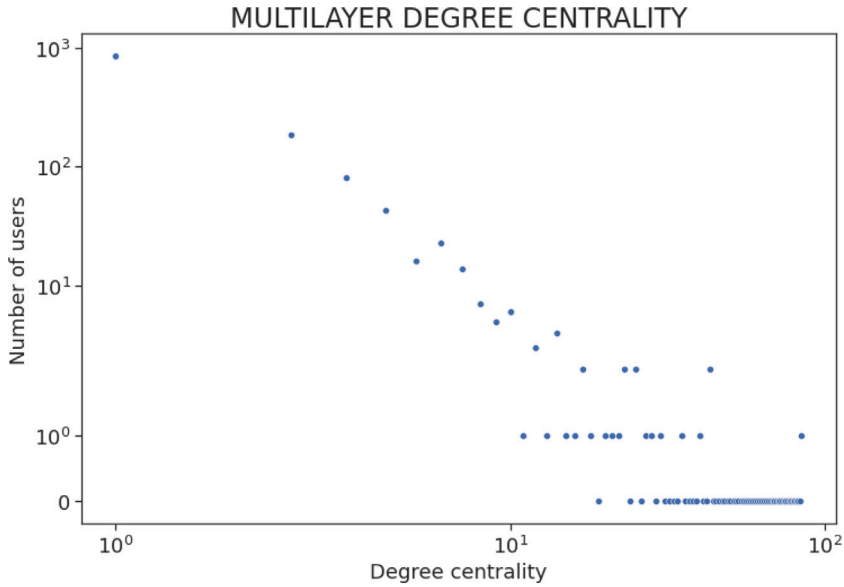


Fig. 7. Multilayer network degree centrality.

Table 7

Top users intersection between the single networks and multilayer network approaches.

Top users intersection	Number of users	Percentage of users
$Retweet \cap Reply \cap To \cap Mention \cap MultilayerTop - 143$	24	16.78%
$Retweet \cap Reply \cap To \cap Mention \cap Like \cap MultilayerTop - 67$	55	82.09%

Table 8

Ego networks statistics for the top users extracted from Multilayer and single Layers Networks.

	Multilayer network			Single networks		
	Nodes	Edges	Weights	Nodes	Edges	Weights
Mean	22.32	27.44	70.08	15.89	19.69	54.51
Std	23.86	34.73	136.95	22.19	32.19	125.99
Min	5	4	4	3	2	2
25%	10	9	10.5	6	5	6
50%	15	17	21	9	9	12
75%	23	28.5	39.5	15.5	18.5	24.5
Max	178	245	824	178	245	824

Starting from this consideration, we investigated the possible differences between the most influential users extracted from the two approaches. To this end, we first compute the ego networks of these users for each layer L . Recall from Section 3.3 that an ego network for a single layer L and a node n in that layer is defined as $\mathcal{E}_n = \langle V_n, E_n \rangle$. $V_n \subseteq V$ is the set of nodes, containing n and all the nodes directly linked through an edge to n in L , while $E_n \subseteq E$ is the set of the edges linking those nodes to n and between them. Then, for each of these users, we computed the multilayer ego networks $\mathcal{E}_{\mathcal{M}_n}$. Recall from Section 3.3 that, given a set of layers L_j , defined for a set of nodes V_i and a node $n \in V_i$, $\mathcal{E} = \langle V_n, E_n \rangle$. $V_n \subseteq V$ is the set of nodes containing n and all the nodes that are connected to n in at least a layer of L_j . Two nodes $v_x, v_y \in V_n$ are linked by an edge $e \in E_n$ if there exists an edge between v_x and v_y in at least a layer of L_j . The obtained results are reported in Table 8.

From the analysis of Table 8, we can note several differences between the two approaches. Indeed, in the multilayer network, the users have larger ego networks w.r.t. number of nodes and edges. This means that these last users interact a lot on Twitter and engage in many more discussions compared to the ones extracted from the single networks analysis. Furthermore, we can observe an important difference between the multilayer and single networks perspectives in the Weights column, which reports the statistics of the edges' weights in the ego networks. Indeed, according to the former approach, the top users not only have a well-connected neighborhood, but they tend to interact much more with their neighbors than with the top users, according to the latter approach. Again, this feature highlights the peculiarity of the most influential users in the multilayer network to start and foster many discussions that do not end with a single tweet but continue for some time.

Finally, we want to investigate if both the approaches considered so far can highlight the topic discussions between the users having different perspectives. Following this reason, we compute the hashtags occurrences from the tweets corresponding to the

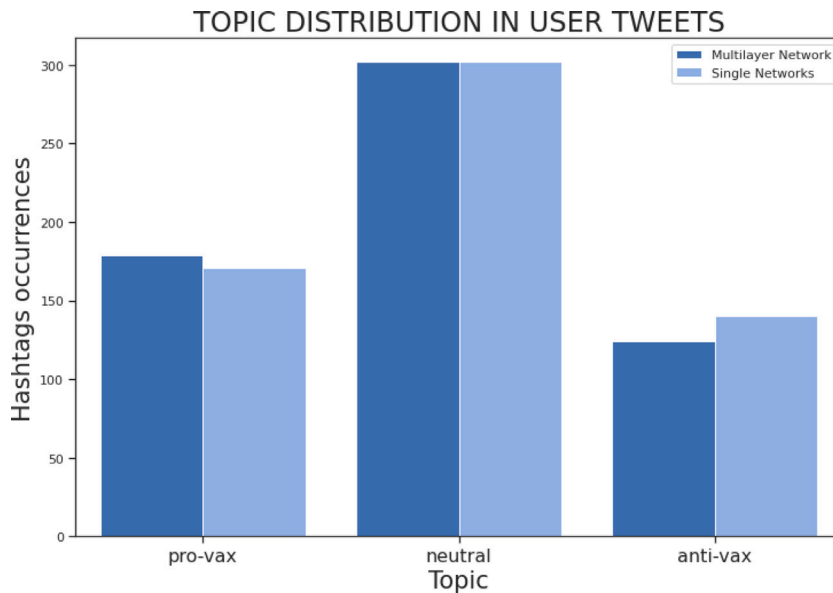


Fig. 8. Gold-standard hashtags occurrences extracted from the tweets of the most influential users according to the single networks and multilayer network approaches.

most influential users, according to the three gold-standard hashtag categories introduced in Section 4, i.e., pro-vax, neutral, and anti-vax. The results are reported in Fig. 8.

From the analysis of Fig. 8, we can observe that the users tend to use neutral hashtags in their tweets. Since they have many followers watching their tweets, they probably have to think carefully about what they are publishing. We can also note that these tweets report many hashtags coming from the pro-vax and anti-vax set, with a slight preference for the former one. Even if the sets of influential users are not the same, we cannot see a meaningful difference.

However, in order to describe the discussion on these three topics, we need to take advantage of our model and project the topic layer to the user ones. In this way, we can analyze the level of interaction for each topic, study the corresponding most influential users and observe the differences between them.

6. Analysis of topic–user projection

One of the strengths of our multilayer network approach is the possibility to project the topic layer L_t on the user layers L_u , which creates a new multilayer network $\mathcal{M}_{V_u}^V$ focused on a specific topic or set of topics. In the specialization of our multilayer network to the Twitter dataset, L_t contains all the hashtags found in tweets. As pointed out in the analysis of tweet contents in Section 4.3 and in Di Giovanni et al. (2021), a discussion (such as pro-vax or anti-vax opinions) could be represented by a set of gold-standard hashtags. In our case, we want to study three different sets of gold-standard hashtag categories: (i) pro-vax, (ii) neutral, and (iii) anti-vax. The first group contains all the hashtags from tweets that somehow express a positive opinion about vaccines. The second group is made up of hashtags from tweets that talk about vaccines but do not show any positive or negative opinions. Finally, the third group contains all the hashtags from tweets that express a negative opinion about vaccines. Each group (anti-vax, neutral and pro-vax) is represented by those hashtags extracted as specified in Section 4.

In our case, all the hashtags are modeled in the L_t layer. Thanks to this layer, it is possible to project our multilayer network in order to understand the polarization of users. Indeed, the projection let us build a new multilayer network where users are linked together if they used at least one common hashtag. Formally, two users of L_u are linked together in the projection if they are linked to at least a hashtag node of L_t through an edge of E_{multi} . Once we obtained the projection $\mathcal{M}_{V_u}^V$, we split this in three different multilayer networks, one for each gold-standard hashtag category, i.e., pro-vax, neutral, and anti-vax. This step will help us to better understand the different dynamics in those three categories of users. As a first analysis, we computed the number of nodes and edges present in each layer after the projection. We removed the isolated nodes and reported the results in Fig. 9.

From Fig. 9, we can observe that the neutral projection has more nodes and edges w.r.t. the pro-vax and anti-vax projections. Probably it is an expected result since the neutral set of topics contains hashtags that do not create controversies and do not offend anyone’s opinions. Moreover, we can see that the anti-vax projection has fewer nodes and edges than the pro-vax one. It seems that the anti-vax topics do not connect many users and tend to establish fewer edges (which means fewer interactions) between users. Finally, it is interesting noting that the number of nodes and edges of neutral and pro-vax projections are close only in the “Reply To” network.

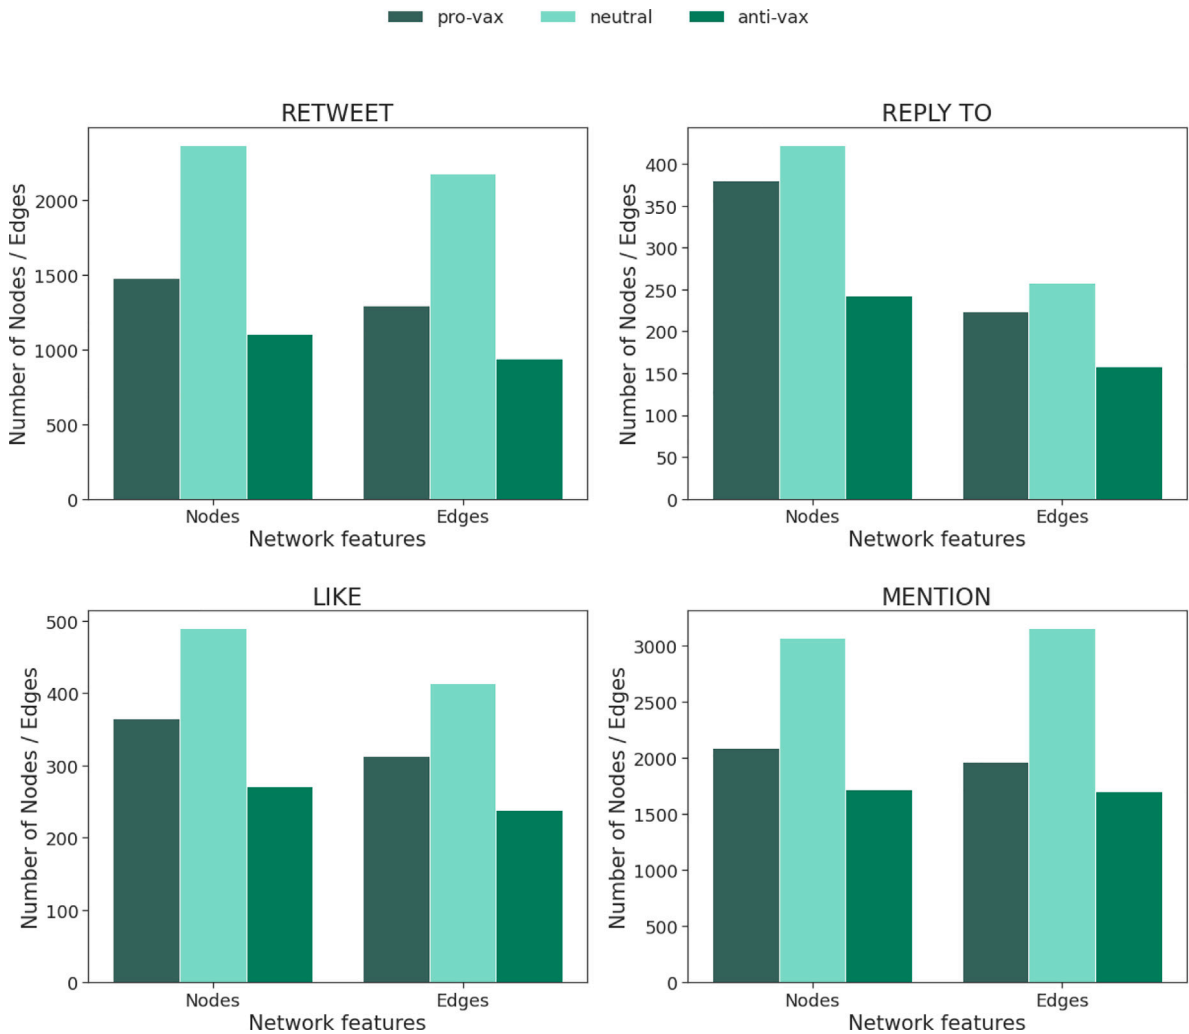


Fig. 9. Number of nodes and edges of the L_{r1} , L_r , L_l and L_m layers of \mathcal{M} after the pro-vax, neutral, and anti-vax projections.

After that, we investigate the most influential users in terms of degree centrality for each projection. Previously, we observed that the considered networks are scale-free networks. As a matter of fact, the multilayer networks of the pro-vax, neutral, and anti-vax projections are scale-free too, which means that the degree centrality distributions of the users are power laws. Following the previous reasoning, we can study the most influential users in order to extract the fundamental characteristics of the multilayer network.

The behavior of the most influential users could be studied thanks to their ego networks. As in the previous sections, the ego network of a user is defined thanks to the composition of his/her ego networks in each layer. As a first step, we reported the average number of nodes, edges, and interactions in Fig. 10, which confirms the previous conclusion about the three projections. In order to prove the effectiveness of the holistic perspective of the multilayer network, we computed the average number of nodes, edges, and interactions of the ego networks of the most influential users according to the single networks approach.

From Fig. 10, we can see that the neutral projection has more nodes and edges w.r.t. the other topics. Moreover, there are no significant differences between pro-vax and anti-vax, which means that the most communicative users of both parties tend to attract few users. However, there is an interesting peculiarity of the anti-vax projections. Indeed, the weights of the ego network edges are much higher than in the other cases. This means that, even if the most influential anti-vax users tend to attract fewer users, they make a lot of interactions (in terms of retweets, replies, mentions, and likes) with the same follower/following. This could represent behavior in which the anti-vax users tend to nurture each other opinions without involving too many users. It is worth noting that the ego networks of the most influential users according to the single networks approach are slightly smaller in terms of nodes and edges compared to the ones extracted from the multilayer network approach. The most evident difference regards the mean number of interactions. Indeed, in all the projections, the top users according to the multilayer network approach have more interactions compared to the single networks case (especially in the anti-vax projection), which means that they communicate more with their neighborhood.

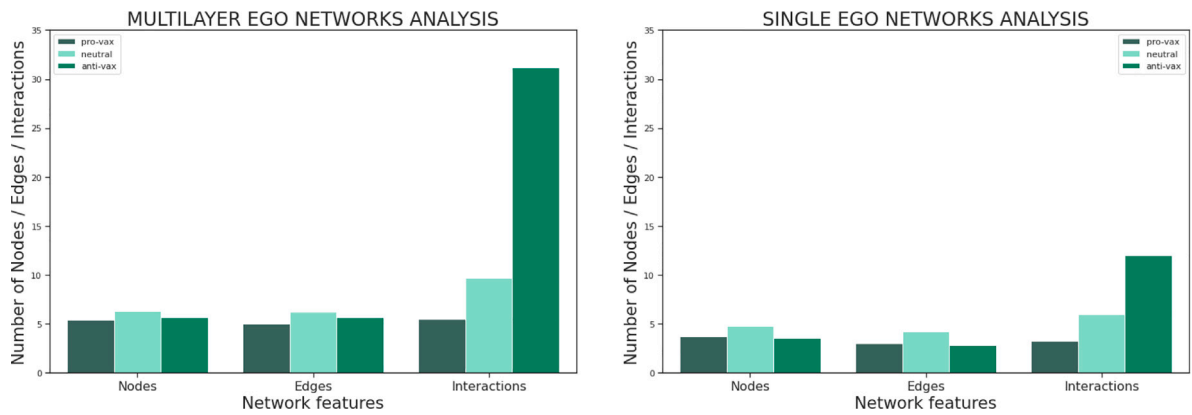


Fig. 10. Average number of nodes, edges, and interactions of the ego network of the most influential users for pro-vax, neutral and anti-vax projections for both multilayer network and single networks approaches.

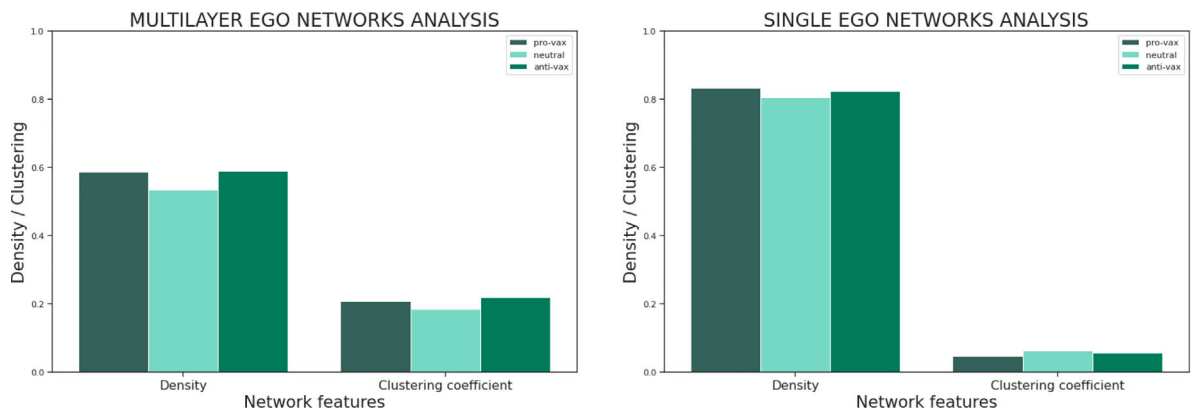


Fig. 11. Average density and clustering coefficient of the ego networks of the most influential users for pro-vax, neutral, and anti-vax projections for both multilayer network and single networks approaches.

In order to deepen our analysis and confirm the previous hypothesis, we computed the average density and clustering coefficient of the ego networks of the most influential users in Fig. 11. As in the previous case, we compute the same statistics for the ego networks of the top users according to the single networks approach.

Fig. 11 contains interesting insights that are in line with the conclusions derived from Fig. 10. In the multilayer case, even if the ego networks of anti-vax users have slightly fewer nodes and edges (which we point out is a negligible difference), the anti-vax ego networks are denser and have a higher clustering coefficient. This highlights the fact that if a user joins an ego network of another one, he/she tends to be well-connected therein. The higher clustering coefficient of the anti-vax projection w.r.t. pro-vax and neutral projections means that the users present in an ego network tend to trust each other since they retweet, reply, mention, and/or like reciprocally, nurturing each other opinions (and also supporting it). Moreover, the pro-vax projection has the lowest clustering coefficient value, which means a low level of trust among users in the ego networks. Also in this case, the comparison between the multilayer network and single networks perspectives is interesting. Indeed, the ego networks of the latter approach are denser than the ones of the former. As reported in Fig. 10, the ego networks of the single networks approach have fewer nodes leading to a higher density even if there are less edges. However, we can see a high difference in the clustering coefficient results. Indeed, in the multilayer case, the ego networks tend to have more triads than in the case of the single networks, which means that in the former perspective the users are much closer to their neighborhood. This result is consistent with the high number of interactions of the top users according to the multilayer network.

Finally, we verified if the most influential users of a topic could be influential in another one. Roughly speaking, we wanted to observe if the most communicative pro-vax users are still communicative in the anti-vax case. This describes possible contamination of topics, which surely convey discussions and (probably, but not so often) could lead to a change of mind in one direction or another. For this reason, we computed the intersections of the Top-100, Top-200, Top-500, and Top-800 users of the pro-vax, neutral, and anti-vax projections, and reported the results in Fig. 12.

From the analysis of Fig. 12, we can derive interesting insights. Indeed, we can see that in all the cases, there are few users that are influential in both pro-vax and anti-vax projections (35%). This leads us to conclude that the contamination between the two

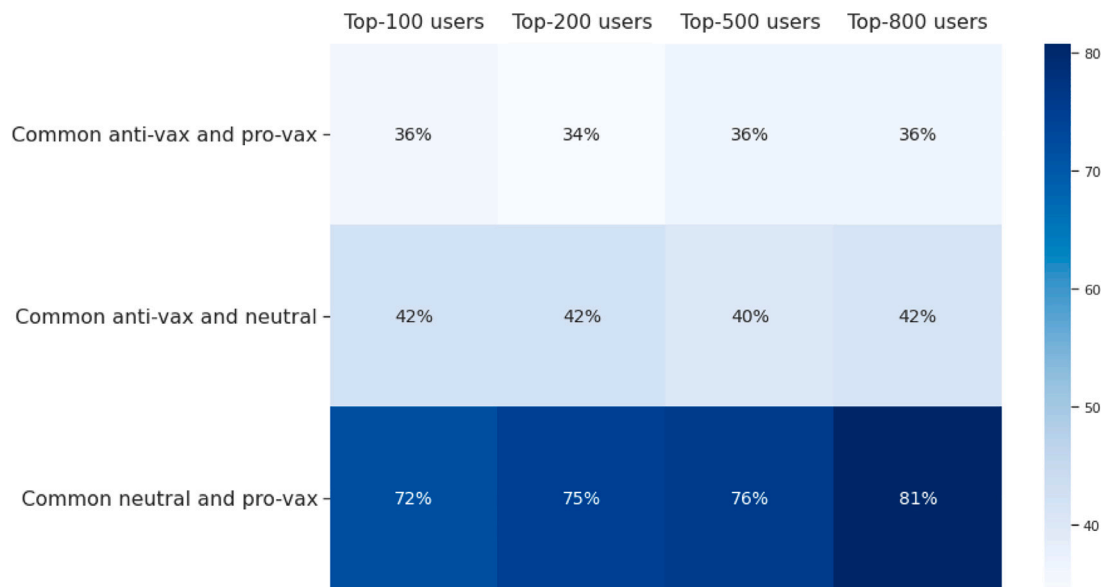


Fig. 12. Percentage of common users among the most influential ones extracted from pro-vax, neutral, and anti-vax projections.

Table 9

Summary of our findings thanks to the multilayer network approach.

	Pro-vax	Neutral	Anti-vax
# Nodes	Medium	High	Low
# Edges	Medium	High	Low
# Interactions	Low	Medium	High
Density	Medium	Low	High
Clustering Coeff.	Low	Medium	High
Findings	Low level of trust and interactions	Big ego networks but low trust among users	Small ego networks but high trust and high number of interactions
	High overlap with the neutral topic	High overlap with the pro-vax topic	Low overlap with the other topics
	Mostly support pro-vax and neutral topics	Mostly support neutral and pro-vax topics	Strongly support their topic

extreme parties is low. Moreover, we can observe a low (but still higher than the previous case) intersection between the neutral and anti-vax projections (41%), which is not the case between neutral and pro-vax projections (75%). This means that, while the pro-vax influential users tend to express themselves through both pro-vax and neutral hashtags, the anti-vaxxer ones mainly leverage the typical hashtags of their field. Obviously, this behavior does not help to foster discussions, nor causes common users to change their minds.

7. Discussion

In the previous sections, we highlighted the differences between the representation of Twitter through a multilayer network and single networks. We figured out that the multilayer network approach is able to extract top users with more nodes (+40.46%), edges (+38.36%), and interactions with their neighbors (+28.56%) than the ones retrieved from the single network approach. The multilayer top users according to the degree centrality carry a lot more information than the single networks ones and allow us to unveil their communication patterns when dealing with specific topics.

Thanks to our multilayer network approach, we identified that the anti-vaxxers tend to have ego networks denser (+14.39%) and more cohesive (+64.2%) than the ones of pro-vaxxer, which leads to a higher number of interactions among anti-vaxxers than among pro-vaxxers (+393.89%). These findings point out the different users' behavior in the three lines of thought. For instance, we figured out that the top anti-vaxxers tend to attract fewer users but they make a lot of interactions (in terms of retweets, replies, mentions, and likes) with their followers, which highlights the nurturing of each other opinions. Moreover, we identified a high clustering coefficient of the anti-vax projection w.r.t. pro-vax and neutral projections. This means that anti-vaxxers tend to show mutual trust and support for their opinions. On the other hand, the pro-vax projection has the lowest clustering coefficient value, which means a low level of trust among the corresponding users. In order to summarize our findings, we report in Table 9 the most important differences in the ego networks of pro-vax, anti-vax, and neutral influential users.

Besides the results we obtained from the Twitter dataset, our multilayer network approach has other interesting advantages. For instance, our approach is not tied to any social network. Indeed, we can think to adapt this approach to any setting in which there are posts made by users and possible interactions between them (such as reposting, liking, tagging, etc.). We can represent any relationship in any social network and derive the corresponding topics from the texts of the posts through a content analysis on hashtags or specific keywords defined by the domain experts. Moreover, we can also think to extend our approach by including the user metadata as node attributes and investigate possible patterns according to the users' interactions. These two observations will be the subject of our next future efforts.

It is also worth pointing out that our multilayer approach is not strictly related to a language or a context. This is evident from the generic content analysis we performed as the preliminary step of the overall approach. Starting from a set of posts (not necessarily tweets, it depends on the considered platform) in any language, we can extract the most important hashtags or keywords according to their frequencies which will identify the topics discussed on that social network. Then, we can create the multilayer network of users and topics and investigate it as we did in this paper. It will be our interest to study how people deal with the same topics in different languages since there will probably be some cultural differences.

However, we identify some limitations to our approach. First of all, we extracted the information about the tweet likes thanks to the free Twitter API, which limits the number of requests to get data. We are restricted to adding further information to the original dataset, and hence to model new concepts.

A further limitation regards the fact that we have considered only the verified users. On the one hand, this helped us to remove Twitter bots, which hide the interesting patterns present in the multilayer network, due to their overwhelming number of tweets, and their focus on a specific topic and/or action (e.g. like, retweet). On the other hand, we can observe that some networks (such as "Reply to" and "Like" networks) have low density and a low level of interactions, probably due to the missing engagements with unverified accounts. Following our reasoning, we are only describing the interactions between verified users over specific topics. However, it is straightforward to investigate the interactions of all Twitter users (including bots and malicious users) by considering all the tweets present in the dataset and then performing the same analysis we did in our case.

Another limitation regards the way we investigated the discussion on Twitter. We have assumed that a discussion is made up of several tweets on the same topics, but we did not consider the conversation perspective. Indeed, a conversation defines new relationships between tweets, which could highlight new ways to observe the evolution of the discussion. However, our multilayer approach is generic enough to model a conversation in a social network considering the chains of posts of the discussion, and then study the characteristics of the involved users.

Finally, another limitation regards the user–topic projection of our multilayer network. Indeed, we found that this particular projection requires a high number of tweets to describe the topic interactions present in the users' tweets. The reasoning behind this phenomenon could regard the low frequency of the users writing tweets on the same set of topics.

8. Conclusion

In this paper, we have presented a multilayer network-based approach to investigate discussions on a social network and prove that, in the anti-vax community there is a strong interaction and trust among the various users we analyzed. Our approach is general and not tailored to a specific scenario, which gives the freedom to study any debate carried out on any social medium. In order to prove the validity of our approach, we have tested it on the Twitter scenario. We have mapped the most important user interactions of Twitter to the layers of the multilayer network and created the corresponding relationships (i.e., "Retweet", "Reply To", "Mention", and "Like"). Moreover, we have added a further layer representing the topic detected on the user tweets, which has then been used to project the multilayer network and obtain a new one focused on specific subjects. To investigate the COVID-19 vaccine discussions on Twitter, we have employed the AvaxTweets dataset. Then, we have shown the differences between multilayer and single networks approaches, which proved the strengths of the former. Finally, we have analyzed the pro-vax, neutral, and anti-vax discussions based on the extracted gold-standard hashtags, and shown the high level of interaction and trust among users employing anti-vax topics.

However, this paper should not be considered as an ending point. Indeed, we plan to apply our multilayer network approach to other Twitter debates, such as political campaigns, climate changes, and so forth. We want to study the most influential verified users and observe their polarization over time. Moreover, we would like to analyze the projection of the layer L_u of users on the layer L_t of the topics. This will require many more tweets, but will highlight the subject co-occurrences according to the type of user interactions, and so the logical connections between hashtags from different lines of thought. In addition, we could think of the analysis of information dissemination patterns of users through our multilayer network, thanks to the presence of multiple types of relationships. For instance, this could be done by including a temporal factor in our model to represent the evolutions of discussion. Finally, we plan to identify the user communities on our multilayer network model thanks to suitable algorithms. Starting from specific topics, it would be interesting to analyze the community structures, their behavior over time, and the evolution of opinions inside the communities.

Data availability

Data will be made available on request.

References

- Abu-Raddad, L. J., Chemaitelly, H., & Butt, A. A. (2021). Effectiveness of the BNT162b2 Covid-19 vaccine against the B.1.1.7 and B.1.351 variants. *New England Journal of Medicine*, (2), 187–189.
- Alstott, J., Bullmore, E., & Plenz, D. (2014). Powerlaw: a Python package for analysis of heavy-tailed distributions. *PLoS One*, (1), Article e85777.
- Barabási, A., & Bonabeau, E. (2003). Scale-free networks. *Scientific American*, (5), 60–69.
- Bhatti, A., Akram, H., Basit, H. M., Khan, A. U., Raza, S. M., & Naqvi, M. B. (2020). E-commerce trends during COVID-19 pandemic. *International Journal of Future Generation Communication and Networking*, (2), 1449–1452.
- Blanco, G., & Lourenço, A. (2022). Optimism and pessimism analysis using deep learning on COVID-19 related Twitter conversations. *Information Processing & Management*, Article 102918.
- Boccaletti, S., Bianconi, G., Criado, R., Del Genio, C. I., Gómez-Gardenes, J., Romance, M., et al. (2014). The structure and dynamics of multilayer networks. *Physics Reports*, 544(1), 1–122.
- Bosco, C., & Cvajner, M. (2021). Investigating Italian citizens' attitudes towards immuni, the Italian contact tracing app. In *Proceedings of the conference on human-computer interaction* (pp. 34–42).
- Burel, G., Farrell, T., & Alani, H. (2021). Demographics and topics impact on the co-spread of COVID-19 misinformation and fact-checks on Twitter. *Information Processing & Management*, (6), Article 102732.
- Burel, G., Farrell, T., Mensio, M., Khare, P., & Alani, H. (2020). Co-spread of misinformation and fact-checking content during the COVID-19 pandemic. In *Proceedings of the international conference on social informatics* (pp. 28–42).
- Campan, A., Cuzzocrea, A., & Truta, T. M. (2017). Fighting fake news spread in online social networks: Actual trends and future research directions. In *Proceedings of the IEEE international conference on big data* (pp. 4453–4457).
- Cauteruccio, F., Corradini, E., Terracina, G., Ursino, D., & Virgili, L. (2020). Investigating Reddit to detect subreddit and author stereotypes and to evaluate author assortativity. *Journal of Information Science*, Article 0165551520979869.
- Cerruto, F., Cirillo, S., Desiato, D., Gambardella, S. M., & Polese, G. (2022). Social network data analysis to highlight privacy threats in sharing data. *Journal of Big Data*, (1), 1–26.
- Corradini, E., Nocera, A., Ursino, D., & Virgili, L. (2021). Investigating the phenomenon of NSFW posts in Reddit. *Information Sciences*, 140–164.
- DeVerina, M. R., Pierri, F., Truong, B. T., Bollenbacher, J., Axelrod, D., Loynes, N., et al. (2021). CoVaxxy: A collection of English-language Twitter posts about COVID-19 vaccines. In *Proceedings of the fifteenth international AAAI conference on web and social media* (pp. 992–999).
- Di Giovannini, M., Corti, L., Pavanetto, S., Pierri, F., Tocchetti, A., & Brambilla, M. G. (2021). A content-based approach for the analysis and classification of vaccine-related stances on Twitter: the Italian scenario. In *Proceedings of the information credibility and alternative realities in troubled democracies* (pp. 1–6).
- Feng, Y., & Zhou, W. (2022). Work from home during the COVID-19 pandemic: An observational study based on a large geo-tagged COVID-19 Twitter dataset (UsaGeoCov19). *Information Processing & Management*, (2), Article 102820.
- Furini, M. (2021). Identifying the features of ProVax and NoVax groups from social media conversations. *Computers in Human Behavior*, Article 106751.
- Galhotra, B., & Dewan, A. (2020). Impact of COVID-19 on digital platforms and change in E-commerce shopping trends. In *Proceedings of the fourth international conference on IoT in social, mobile, analytics and cloud* (pp. 861–866).
- Hung, M., Lauren, E., Hon, E. S., Birmingham, W. C., Xu, J., Su, S., et al. (2020). Social network analysis of COVID-19 sentiments: Application of artificial intelligence. *Journal of Medical Internet Research*, (8), Article e22590.
- Jeyasudha, J., et al. (2021). Topic modeling and sentimental analysis of tweets on covid19 to find the weightage of the popular hashtag. *Turkish Journal of Computer and Mathematics Education (TURCOMAT)*, (9), 1856–1861.
- Jones, C., & Volpe, E. H. (2011). Organizational identification: Extending our understanding of social identities through social networks. *Journal of Organizational Behavior*, (3), 413–434.
- Kalamatianos, G., Mallis, D., Symeonidis, S., & Arampatzis, A. (2015). Sentiment analysis of greek tweets and hashtags using a sentiment lexicon. In *Proceedings of the 19th Panhellenic conference on informatics* (pp. 63–68).
- Kolda, T., & Bader, B. (2006). The TOPHITS model for higher-order web link analysis. In *Workshop on link analysis, counterterrorism and security, Vol. 7* (pp. 26–29).
- Kovacs, B., Caplan, N., Grob, S., & King, M. (2021). Social networks and loneliness during the COVID-19 pandemic. *Socius*, Article 2378023120985254.
- Latkin, C., Dayton, L. A., Yi, G., Konstantopoulos, A., Park, J., Maulsby, C., et al. (2021). COVID-19 vaccine intentions in the United States, a social-ecological framework. *Vaccine*, (16), 2288–2294.
- Luo, C. (2021). Analyzing the impact of social networks and social behavior on electronic business during COVID-19 pandemic. *Information Processing & Management*, (5), Article 102667.
- Mourad, A., Srour, A., Harmanai, H., Jenainati, C., & Arafeh, M. (2020). Critical impact of social networks infodemic on defeating coronavirus COVID-19 pandemic: Twitter-based study and research directions. *IEEE Transactions on Network and Service Management*, (4), 2145–2155.
- Muric, G., Wu, Y., Ferrara, E., et al. (2021). COVID-19 vaccine hesitancy on social media: building a public twitter data set of antivaccine content, vaccine misinformation, and conspiracies. *JMIR Public Health and Surveillance*, (11), Article e30642.
- Nguyen, N.-L., Wang, M.-H., Dai, Y.-C., & Dow, C.-R. (2021). Understanding malicious accounts in online political discussions: A multilayer network approach. *Sensors*, 21(6), 2183.
- Oro, E., Pizzuti, C., Procopio, N., & Ruffolo, M. (2017). Detecting topic authoritative social media users: a multilayer network approach. *IEEE Transactions on Multimedia*, 20(5), 1195–1208.
- Pierri, F., Piccardi, C., & Ceri, S. (2020). A multi-layer approach to disinformation detection in US and Italian news spreading on Twitter. *EPJ Data Science*, 9(1), 35.
- Pisal, S. (2021). Rise of facebook, amazon, apple, netflix, google during COVID-19 pandemic. *Electronic Theses, Projects, and Dissertations*, 1311.
- Rufai, S. R., & Bunce, C. (2020). World leaders' usage of Twitter in response to the COVID-19 pandemic: a content analysis. *Journal of Public Health*, (3), 510–516.
- Sharma, K., Zhang, Y., & Liu, Y. (2021). COVID-19 vaccines: Characterizing misinformation campaigns and vaccine hesitancy on Twitter. arXiv preprint arXiv:2106.08423.
- Singh, L. G., Mitra, A., & Singh, S. R. (2020). Sentiment analysis of tweets using heterogeneous multi-layer network representation and embedding. In *Proceedings of the 2020 conference on empirical methods in natural language processing EMNLP*, (pp. 8932–8946).
- Tran, L. T. T. (2021). Managing the effectiveness of e-commerce platforms in a pandemic. *Journal of Retailing and Consumer Services*, Article 102287.
- Türker, İ., & Sulak, E. E. (2018). A multilayer network analysis of hashtags in twitter via co-occurrence and semantic links. *International Journal of Modern Physics B. Condensed Matter Physics. Statistical Physics. Applied Physics.*, (04), Article 1850029.
- Vargas, A. N., Maier, A., Vallim, M. B., Banda, J. M., & Preciado, V. M. (2021). Negative perception of the COVID-19 pandemic is dropping: Evidence from Twitter posts. *Frontiers in Psychology*, 4067.
- Ward, M. P., Xiao, S., & Zhang, Z. (2020). The role of climate during the COVID-19 epidemic in New South Wales, Australia. *Transboundary and Emerging Diseases*, (6), 2313–2317.
- Willnat, L., & Weaver, D. H. (2018). Social media and US journalists: Uses and perceived effects on perceived norms and values. *Digital Journalism*, (7), 889–909.
- Wymant, C., Ferretti, L., Tsallis, D., Charalambides, M., Abeler-Dörner, L., Bonsall, D., et al. (2021). The epidemiological impact of the NHS COVID-19 app. *Nature*, (7863), 408–412.