Investigating community evolutions in TikTok dangerous and non-dangerous challenges

Journal of Information Science 1–25

DOI: 10.1177/01655515221116519
journals.sagepub.com/home/jis

Investigating community evolutions in TikTok dangerous and non-dangerous challenges

Abstract

In just few years, TikTok has become a major player in the social media environment, especially with regards to teenagers. One of the key factors of this success is the idea of challenges, i.e., video competitions/emulations on a certain topic, which a user can launch and other ones can join. Most of the challenges are fun and harmless. However, there are also users who launch challenges that are dangerous, or at least suitable only for an adult audience (and TikTok is the most popular social network for teenagers). This paper focuses primarily on this kind of challenge. In particular, it investigates an aspect not yet studied in the literature, which is the different characteristics and evolutionary dynamics of the communities of users participating in non-dangerous and dangerous challenges. Its final goal is the identification of evolutionary patterns that distinguish the communities of users participating in the two types of challenges. The knowledge of these patterns could be a first step in implementing an approach to the early detection of dangerous challenges in TikTok.

Keywords: Social Network Analysis; TikTok; Challenge Lifespan; Community Evolution; Evolutionary Patterns; Challenge Classification

1 Introduction

A few years after its appearance, TikTok[1] (also known as Douyin in China) has attracted the interest of hundreds of millions users, especially, but not only, among teenagers. The strength of TikTok are videos, generally short, through which users can launch challenges. A challenge consists in a series of videos emulating the original one launching it. TikTok supplies several tools specifically designed for video editing, manages HD resolution, full screen display and provides its users with the possibility of adding a music clip to a posted video. The varied and qualified set of functionalities for video management and, above all, the possibility of launching challenges or participating in them represents the main strength of this social platform.

A challenge is identified by a hashtag; it begins when a user posts a video with that hashtag and invites other ones to replicate that video in their own way. Most of the challenges are fun and not dangerous, but some of them are dangerous or, in any case, suitable for an adult audience only, while TikTok is the most popular social network among teenagers. To give an idea of the dangerousness of

[1] www.tiktok.com

some challenges, we mention the Benadryl challenge, which encouraged users to ingest large amounts of diphenhydramine to get high and record their responses, and the Blackout challenge, which encouraged users to choke themselves until the point of losing consciousness, while uploading the results on TikTok.

TikTok has increased security controls and removed challenges judged dangerous. However, every day the authors of dangerous challenges find new tricks to bypass TikTok's controls. Taking into account the number of users of this social medium and the number of challenges launched daily on it, it is easy to understand how the definition of automatic tools able to distinguish a non-dangerous challenge from a dangerous one is a very important issue to address.

Another interesting research issue regarding this social medium concerns the study of the communities participating in a challenge and their evolution over time. In particular, some questions that can be investigated are the following: Are there differences in the evolution and dynamics of the communities related to non-dangerous and dangerous challenges? What can be said about these communities regarding the connection level of users, the configuration of friendships and all those issues typical of Social Network Analysis?

This paper is intended as an attempt to address these challenging issues. In particular, we study the characteristics of the communities participating in dangerous and non-dangerous challenges, the behavior of the corresponding users and their dynamics and evolution over time. The final goal is the possible detection of evolutionary patterns allowing the distinction of non-dangerous challenges from dangerous ones.

Regarding this fact, it must be said that TikTok has been intensively studied in the literature from multiple perspectives, especially with regards to influencers [32, 64], and their role in marketing [12, 65, 24, 52], politics [50, 39, 55], health [72, 38, 10, 31], etc. Many other studies have focused on the recommendation algorithm underlying TikTok [17, 67, 53, 71, 36, 4], privacy and security issues [42, 33, 41, 70], types of messages and contents that, directly or indirectly, are spread through this social platform [4, 66]. There are also some studies about challenges [74, 57], the principle of imitation underlying them [35] and the strategies with which the videos launching them are designed [11]. However, to the best of our knowledge, no paper specifically investigated the differences in the evolutionary dynamics of communities in dangerous and non-dangerous challenges, as well as the possibility of exploiting these differences to search for evolutionary patterns capable of distinguishing one kind of challenge from the other. Our contribution goes exactly in that direction.

To perform our analysis, we selected seven non-dangerous challenges and seven dangerous ones. For each of them, we considered the corresponding posted videos and a set of features characterizing the associated user communities (e.g., number of connected components, size of the maximum connected component, average clustering coefficient, average path length). Next, we defined a social network-based model to represent the user community associated with each TikTok challenge. Using this model, we investigated the evolutionary dynamics of the communities associated with non-dangerous and dangerous challenges. First, we focused on the characteristics of their videos and the parameters of the social networks associated with their communities. From a first analysis, taking into account the evolution of the community size during the challenge lifespans, we could observe that non-dangerous and dangerous challenges seemed to show different dynamics. Here, a clarification on the term "lifespan" is in order. By "lifespan" of a challenge, we do not mean the time period elapsing from when it is launched to when it finally disappears from TikTok. In fact, there are challenges

that never disappear from this social platform, even though they have not received new videos from months or years. Here, the lifespan of a challenge is the time period elapsing from when it is launched to when it is no longer able to elicit at least limited interactions with users.

To capture the differences on the community dynamics in the two kinds of challenge, we divided lifespans into suitable intervals. Then, we grouped these intervals into homogeneous clusters. At this point, for each cluster, we used the values of the Social Network Analysis parameters characterizing the communities corresponding to the intervals belonging to it for drawing the cluster's profile. After this, for each challenge, we identified the sequence of intervals, along with the corresponding clusters, which formed its lifespan. From examining these sequences and the characteristics of the corresponding clusters, we hypothesized that some clusters were substantially equivalent and verified the correctness of this hypothesis by means of a t-test [6].

After verifying this correctness, we could simplify the sequences related to challenges, and this allowed us to identify a main evolutionary pattern characterizing non-dangerous challenges, and two main evolutionary patterns, different from the previous one, characterizing dangerous challenges. This result provides a new way to distinguish the two types of challenges. After obtaining this result, we tested whether it was accurate and generalizable to the other TikTok challenges. To this end, we considered 300 challenges and were able to verify that our model was very accurate also for this sample, whose size was much larger than the one initially used.

We point out that the classification approach we propose in this paper is currently able to support the detection of dangerous challenges only near the end of their lifespan, or at least after that a fairly long time period has elapsed since their beginning. On the other hand, the early detection of dangerous challenges is not the goal of this paper. In fact, in it, we want to define an approach to the classification of TikTok challenges. Although our paper does not aim at early detection of dangerous challenges in TikTok, it makes its own contribution in the literature related to the classification of video content published in social media, as will be clear in Section 2. Actually, the early detection of dangerous challenges is an extremely difficult problem that cannot be solved in a single paper, but needs a multi-stage research. In fact, in the early detection of videos, we cannot rely on metadata alone because they might be deliberately falsified by malicious authors [44, 69]. Therefore, any approach based on the actual content of challenges or the behavior of people accessing them is necessarily complex and first requires a series of researches to understand the phenomenon. Only after fully understanding the latter, it is possible to think of approaches that use the knowledge gained to propose a solution. This necessarily requires lengthy multi-stage researches. Our paper is in the first stage of one such research, i.e., the stage devoted to better understand the phenomenon. In the future, in order to achieve the latter goal, we may consider reducing the granularity of the time intervals considered. In this way, we can think of identifying very soon some evolutionary patterns allowing an early detection of dangerous challenges.

This paper is structured as follows: In Section 2, we present the related literature. In Section 3, we illustrate the initial dataset storing data about the 14 challenges we used for our analysis. In Section 4, we define our Social Network Analysis based model to represent the community of users associated with a challenge. In Section 5, we propose a preliminary analysis of dangerous and non-dangerous challenges and their corresponding communities. This represents the starting point for the study of the evolution of user communities associated with the two types of challenges, which is presented in

Section 6. In Section 7, we illustrate the reasoning leading us to the discovery of evolutionary patterns for the two types of challenges. Finally, in Section 8, we draw our conclusion and take a look at possible future developments of this research.

2 Related literature

In recent years, TikTok has been the subject of analysis by researchers operating in different fields [56]. For example, it has been studied in the context of Social Network Analysis, marketing, machine learning and deep learning, politics, and so on [72, 38, 10, 50, 39, 55, 49, 34]. The opportunities and challenges posed by this social medium are clearly described in [13].

Compared to other social platforms, TikTok is characterized by a massive diffusion among teenagers [28]. This has led to the emergence of new types of influencers, suited for this social platform [32]. The study of such influencers represents the main objective of [64]. In [29], the authors propose an analysis on "personal branding". This term refers to the process of creating a brand from a person's profile. Researchers have also performed analyses to understand whether marketing strategies and influencer actions in TikTok actually lead to increased brand awareness and sales [3].

Another issue related to TikTok, on which researchers have turned their attention, concerns privacy and security [42, 33, 41]. Obviously, TikTok has largely attracted the interest of researchers working in the context of Social Network Analysis [66]. Many authors have turned their attention to the recommendation algorithm used by TikTok [17, 67, 53, 71, 36, 4]. Some authors have focused on using machine learning and deep learning approaches to understand the dynamics of this social medium [68].

As for TikTok challenges, which represent the main focus of this paper, few studies concerning them can be found in the past literature. In particular, the authors of [74] investigate the role of these challenges in fostering the imitation principle. In this analysis, they use the concept of memes and introduce the notion of "imitation publics". The author of [35] focuses on strategies that can be adopted to create a video for a challenge; to this end, he analyzes the `#distantdance` challenge in detail. The authors of [11] study the processes through which challenges can influence TikTok users. Finally, the authors of [57] analyze how TikTok challenges can be used to spread specific messages in this social medium.

The topic considered in this paper can be seen as a specialization for TikTok of a more general topic related to the discovery of communities in social media, their classification and the study of their evolution. These themes are of fundamental importance in Social Network Analysis [60]. In this context, some studies focus on static methods for community detection [20], while others analyze the activities of social network members to investigate their evolution over time [16, 47]. To perform this task, it is possible to define the concept of dynamic network, which is a special type of complex network that changes over time [45]. Changes in the network occur when new members join or leave it, when existing relationships disappear, when new relationships appear, and so on. These structural changes lead to a continuous evolution of the network, which means that the corresponding structure must be continuously recomputed.

Various approaches for studying the temporal evolution of communities have been proposed in the past literature [16, 47, 40]. In order to provide a complete overview of them, [16] proposes a taxonomy

consisting of four categories. These includes approaches for: *(i)* Independent Community Detection and Matching, *(ii)* Dependent Community Detection, *(iii)* Simultaneous Community Detection on All Snapshots, and *(iv)* Dynamic Community Detection on Temporal Networks.

Independent Community Detection and Matching approaches operate by applying the static community detection methods to the dynamic case. They consider the evolution of the network into consideration in many time steps. During each time step, the network is modeled as a set of communities. The communities of a time step can be matched with those of the previous time step based on a similarity measure. For example, the authors of [58] focus on social networks and propose an event detection algorithm to find community evolution patterns between adjacent snapshots. In this way they are able to evaluate the evolution trend of the whole network. In [73], the authors describe a framework for event reconstruction that aims to analyze the dynamic characteristics of community structure. They define a set of community attributes and reconstruct events based on the examination of these attributes. In [59], the authors propose a model and a similarity measure, called mutual transition, to track communities and to analyze significant transition events occurring in them.

Dependent Community Detection approaches leverage snapshots and detect communities for each of them. Given a certain snapshot, these methods consider the communities found in the immediately preceding snapshot or, otherwise, in the most recent ones. For example, the authors of [26] improve the Louvain algorithm by including the concept of dynamism when forming communities. They use the communities detected at time $t-1$ to identify the communities at time t. In [23], the authors associate attributes with the topology of a graph and define the topological attraction between nodes and communities. They update the current community structure based on the changes occurred in the previous time step. In [22], the authors propose an evolutionary community discover algorithm based on leader nodes (called EvoLeaders). Each community is considered as a set of follower nodes close to a potential leader. An EvoLeader represents the most central node in the corresponding community. By keeping the leader nodes over the evolution of communities, these last could show continuity with respect to the previous versions.

Simultaneous Community Detection on All Snapshots approaches take in input all the evolution stages of a social network simultaneously. They create a single network by binding together in a single graph all the snapshots of the social network. In this way, they maintain the structures aligned in time by coupling the arcs between the same nodes at different time steps. For example, the authors of [62] propose a general framework for finding communities in dynamic networks. First, they model such a task as a graph coloring problem. Then, they use a heuristic technique involving greedily matching pairs of node sets between time steps, in descending order of similarity. In [61], the authors go further and include arbitrary dynamic networks. They solve an optimization problem using a semi-definite programming relaxation and a rounding heuristic. In [30], the authors construct a single network from all snapshots by connecting similar nodes appearing in different time steps. They also create links between nodes connected to at least one common neighbor in two consecutive time steps. Finally, they use the classical Walktrap community detection algorithm.

Dynamic Community Detection on Temporal Networks approaches work directly on temporal networks. In this case, the authors do not consider snapshots but the changes occurring in the network. The idea is to search for communities and study their evolution by analyzing the addition and removal of nodes and arcs in the network. For example, the authors of [37] consider the evolution

of the network arc by arc. A node is considered belonging to the community with which it shares the largest number of arcs. Thus, the addition or removal of arcs can result in a node moving from one community to another. To avoid the continuous oscillation of a node from one community to another, if the difference between the number of arcs that a node shares with two communities is below a certain threshold, then the node will remain in the community it previously belonged to. In [48], the authors propose Tiles, an algorithm that tracks the evolution of communities over time. When a new interaction happens in the network, Tiles uses a label propagation procedure to propagate the changes to the node's neighborhood. In [27], the authors propose an algorithm to find communities based on high-connected hubs. It first searches for highly connected nodes, which will represent the hubs. Then, it assigns the non-hub nodes to the nearest hub. This assignment can evolve iteratively over time.

As specified in the Introduction, the main objective of our paper is to study the differences between non-dangerous and dangerous challenges in TikTok. This study is conducted focusing mainly on the difference in the evolution of the corresponding communities, finding different evolutionary patterns that characterize the two kinds of community. Its ultimate goal is trying to distinguish non-dangerous and dangerous challenges based on the behavior of the corresponding communities. To the best of our knowledge, no other paper in the literature investigated this issue. To achieve its goals, our approach uses a wide set of notions from Social Network Analysis [15, 14, 63], Data Mining [8, 25] and Statistics [6]. In particular, it constructs a social network for each challenge and uses several parameters typical of Social Network Analysis to characterize it. Then, it adopts Data Mining techniques (in particular, clustering) to build a first rough version of the evolutionary patterns capable of characterizing the two kinds of community. Finally, it uses the t-test [6] to test some hypotheses that allow a further refinement of the previously detected evolutionary patterns.

As anticipated in the Introduction, this paper has its own definite collocation in the scientific literature and makes its own contribution to it despite the fact that it deals only with the first stage of a research on early detection of malicious videos.

Similar to our case, several papers in the literature have dealt with the classification of videos using complex techniques, which do not take into account only the metadata that can be faked by the authors. For example, the authors of [69] propose a method of classifying inappropriate videos on YouTube that does not take metadata into account. The authors of [44] propose a classifier to identify inappropriate videos for children on YouTube, which could be even recommended by YouTube's own recommendation algorithm misled by deep fake videos. Another approach that deals with the classification of videos having within them misleading content on COVID-19 is presented in [51]. Another example of classification of child-unsafe videos is proposed in [54]. Finally, the identification of extremist videos in online video sharing sites is investigated in [21]. In the literature, our paper fits into this line of research. In this context, it makes a very different contribution from others since it is based on the behavior of user communities accessing videos.

In this line of research, our approach can be considered:

- A response to all those authors who perform classification of videos through the corresponding metadata and who argue the need to expand their approach with non-textual features [18, 21, 1]. In fact, our approach considers behavioral features.

- A response to the many papers that make classification of videos on more classical social platforms, such as YouTube, and that argue that this area of research opens up to consider more recently appeared platforms, which present video structuring and social interaction dynamics very different from classical ones (as is exactly the case with TikTok) [18, 69, 44, 1].

The motivations for various authors to propose video classification approaches in social media like ours concern:

- The possibility of reporting that malicious users posted video with deliberately wrong metadata that enable them to reach segments of users for whom they are not suitable.
- The possibility of reporting that deep-faked videos are posted as real and that these reach users with low critical sense, and thus capable of believing that those videos are true.
- The possibility of reporting that someone is posting extremist videos or, even, videos that incite terrorism and push kids and young people to enroll in terrorist organizations.
- The possibility that a deeper understanding of the phenomenon of dangerous videos is a first step toward their early detection.

3 Dataset construction

In order to perform our research, we needed a dataset recording data and metadata related to non-dangerous and dangerous challenges in TikTok. To the best of our knowledge, there was no dataset with such characteristics already available; therefore, we had to build it from scratch. In identifying the challenges to be considered in such a dataset, we focused on some of them that were very common in TikTok at the time of data extraction. Specifically, we considered seven non-dangerous challenges and seven dangerous ones. To this end, we assumed as dangerous a challenge that had received several criticisms in the media about the problems it could cause to the people participating in it. As it usually happens in TikTok, we identify each challenge through the hashtag used to post the corresponding videos. In Table 1, we report the seven non-dangerous challenges, while in Table 2 we show the seven dangerous ones. Actually, in the past, much more dangerous challenges than those shown in Table 2 have spread on TikTok. Some of them, such as the Benadryl challenge and the Blackout challenge mentioned in the Introduction, even caused the death of participants. These challenges, and other ones equally disrupting, were promptly blocked by TikTok and the access to the corresponding data was impossible.

At this point, a consideration on the number of challenges we have chosen is necessary. In fact, the classification problem we are considering is a typical "rare class problem" [6]. It arises in the presence of a strong imbalance of the two classes to be predicted with the class of greatest interest being the rare one. In such a scenario, it is better to have a model less accurate but capable of identifying as many instances of the rare class as possible [6]. To achieve this, a balancing of the two classes is done, even though in reality the rare class is much less prevalent. In practice, as mentioned above, it is very difficult to find data on dangerous challenges because they are rare and are removed from TikTok as

Challenge	*Description*
`#bussitchallenge`	Participants show themselves changing clothes.
`#copinesdancechallenge`	Participants perform a series of dance movements.
`#emojichallenge`	Participants imitate different emoji.
`#colpiditesta`	Participants virtually hit a soccer ball with their heads.
`#boredinthehouse`	Participants film a subject, often an animal, in different parts of the house.
`#itookanap`	Participants film a subject, often an animal, sleeping.
`#plankchallenge`	Participants perform dance movements based on training exercises.

Table 1: The seven non-dangerous challenges of our dataset

Challenge	*Description*
`#silhouttechallenge`	Participants expose their bodies covered by a red filter. They are often naked and the filter, being digital, can be easily removed.
`#bugsbunny`	Participants lie on their stomachs and lift their legs upward to show their feet sticking out of their heads like the ears of a rabbit. Then they begin to move their feet to the beat of a song. They often show intimate parts of their bodies.
`#strippatok`	Participants post videos related to strippers (both males and females). Clearly it regards topics not suitable for a young audience.
`#firewroks`	Participants post videos with fireworks risking their safety. The seemingly wrong hashtag is a trick to bypass TikTok's controls.
`#fightchallenge`	Participants post videos with fights that they organize. It is judged dangerous because it can lead to fighters getting injured.
`#sugarbaby`	Participants post videos about "sugar babies", i.e., young people having sex with older people for money.
`#updownchallenge`	Participants move intimate parts of their bodies to the beat of a song.

Table 2: The seven dangerous challenges of our dataset

soon as they are recognized as dangerous. Therefore, in order to have a balanced dataset, we had to undersample the non-dangerous challenges. This way of proceeding can lead to a worsening of the overall accuracy of our approach but allows us to obtain very high sensitivity values. In turn, this allows our approach to correctly classify as many dangerous challenges as possible.

Having said that, we note that in any case the number of challenges considered may seem low and, in some ways, it is. This is due in part to the rarity of the dangerous challenges and in part to the typical way of proceeding of the research investigations on TikTok. In fact, these investigations take into account few challenges, each characterized by many videos. For example, [43] analyzes 12 challenges, [2] examines 8 challenges, [7] considers 8 challenges and a total of 100 videos, [19] studies only one challenge characterized by 1,495 videos; finally, [50] and [46] each analyze two challenges. As we will see below, our 14 challenges still led us to examine 6,005 videos, which represent a significant number in the TikTok investigation scenario.

Table 3 shows the number of videos we collected for each challenge, along with the date of the first and last one.

With regard to this table, we point out that, in the period in which we carried out our tests (July 2021 - September 2021), the lifespan of all the challenges in our dataset can be considered concluded (according to the meaning we gave to the concept of lifespan conclusion in the Introduction). In fact,

Challenge	*Number of Videos*	*Date of the first video*	*Date of the last video*
Non-dangerous Challenges			
`#bussitchallenge`	803	2020-06-11	2021-03-28
`#copinesdancechallenge`	250	2020-12-10	2021-03-24
`#emojichallenge`	663	2018-09-25	2021-03-27
`#colpiditesta`	1086	2018-01-21	2021-04-08
`#boredinthehouse`	359	2019-11-12	2021-04-06
`#itookanap`	206	2018-09-16	2021-03-22
`#plankchallenge`	380	2018-06-22	2021-04-08
Dangerous Challenge			
`#silhouttechallenge`	266	2018-08-15	2021-03-24
`#bugsbunny`	252	2018-01-05	2021-04-09
`#strippatok`	756	2019-02-16	2021-04-19
`#firewroks`	118	2018-02-03	2021-04-14
`#fightchallenge`	381	2018-08-08	2021-04-20
`#sugarbaby`	174	2018-09-11	2021-04-22
`#updownchallenge`	311	2018-06-17	2021-04-25

Table 3: Number of videos and date of the first and last video for each challenge

although these challenges continued to be present in TikTok, they were no longer able to stimulate meaningful interactions with users.

After choosing the challenges, we developed a crawler to obtain public data about them and the corresponding videos. Our crawler anonymizes information about the authors of the videos. More specifically, for each challenge, it records the identifier of the video originating it and the ones of the other videos referring to it. For each of these videos, our crawler also derives its list of likes. Finally, for each like, it determines: *(i)* the user who posted it; *(ii)* her privacy policy; *(iii)* any possible video that she posted for the same challenge[2].

After downloading the data for each video and performing some pre-processing tasks, we obtained a record for it. This record contains the fields shown in Table 4.

4 Model definition

After illustrating the dataset on which performing our analyses, we want to define a model to represent a challenge. Our model is a social network-based one.

Specifically, let $\mathcal{C}'$ (resp., $\mathcal{C}''$) be the set of non-dangerous (resp., dangerous) challenges and let $\mathcal{C}$ be the union of $\mathcal{C}'$ and $\mathcal{C}''$. Let C_i be a challenge of $\mathcal{C}$; a social network $\mathcal{N}_i = \langle N_i, A_i \rangle$ can be associated with it.

N_i is the set of nodes of $\mathcal{N}_i$. There is a node n_{i_j} for each author a_{i_j} who posted at least one video for C_i. Each node n_{i_j} has associated a label l_{i_j} that registers the publication timestamp of the first video that a_{i_j} posted for C_i[3]. Since there is a biunivocal correspondence between a node $n_{i_j} \in N_i$ and the corresponding author a_{i_j}, in the following we will use these two terms interchangeably.

A_i is the set of arcs of $\mathcal{N}_i$. An arc $(n_{i_j}, n_{i_k}) \in A_i$ denotes that the author a_{i_k} liked a video published by a_{i_j} and that the timestamp recorded in l_{i_j} precedes the one recorded in l_{i_k}. Intuitively, the arc

[2]In TikTok, a user can post more videos for the same challenge.

[3]Note that a_{i_j} could post more videos for C_i over time.

Feature	*Description*
challenge_id	The hashtag of the challenge which the video belongs to.
createTime	The publication date of the video.
video_id	The video identifier.
video_duration	The video duration, expressed in seconds.
author_id	The identifier of the author of the video.
author_verified	It indicates whether the author is verified (in TikTok, a verified user denotes a notable person).
music_id	The identifier of the music track or sound used in the video.
music_title	The title of the music track or sound used in the video.
stats_diggCount	The number of likes obtained by the video.
stats_playCount	The number of views of the video.
authorStats_diggCount	The total number of likes expressed by the author of the video for other videos.
authorStats_followingCount	The number of users followed by the author of the video.
authorStats_followerCount	The number of users following the author of the video.
authorStats_heartCount	The total number of likes received by the author of the video.
originalVideo	It is set to 1 if the video began the challenge it belongs to; otherwise, it is set to 0.
likedBy_ids	The list of identifiers of the users, who put a like to the video and have their privacy policy set to "public" (our crawler can operate only with users adopting this policy; it cannot derive information from users having their privacy policy set to "private".).

Table 4: The record associated with each video of a challenge

(n_{i_j}, n_{i_k}) denotes that the challenge C_i propagated from a_{i_j} to a_{i_k}. In fact, a_{i_j} posted a video for C_i; this was liked by a_{i_k}, who, in turn, posted a video of her own for the same challenge. Accordingly, an arc from n_{i_j} to n_{i_k} indicates the joint occurrence of two facts, namely that a_{i_k} liked a video published by a_{i_j} and that, in turn, she decided to propagate the corresponding challenge by publishing of her own a video on it. Thus, the existence of an arc from n_{i_j} to n_{i_k} represents a strong adherence of a_{i_k} to the challenge. In fact, a_{i_k} not only had to like a video posted by a_{i_j} (which already denotes a form of interest in the corresponding challenge) but in turn had to actively (and not only passively) participate in the challenge by posting a video of her own related to it.

An example helps us to better understand our model. Suppose to have a challenge C_i and five users, say Alice, Bob, Mary, Peter and Kate. Alice posts a video for C_i at timestamp t_1 and another video for C_i at timestamp $t_2 > t_1$. Mary likes the video of Alice at timestamp $t_3 > t_2$ and a video for C_i at timestamp $t_4 > t_3$. Bob posts a video at timestamp $t_5 > t_4$. Peter likes the video of Mary at timestamp $t_6 > t_5$ and the video of Bob at timestamp $t_7 > t_6$. Finally, he posts a video for C_i at timestamp $t_8 > t_7$. Kate likes the video of Peter at timestamp $t_9 > t_8$ and the video of Bob at timestamp $t_{10} > t_9$. Finally, Kate posts a video for C_i at timestamp $t_{11} > t_{10}$ and Peter posts another video for C_i at timestamp $t_{12} > t_{11}$. The corresponding network $\mathcal{N}_i$ is shown in Figure 1.

Note that in $\mathcal{N}_i$ there is a node for each user. The timestamp associated with Alice is t_1 because, even though Alice posted two videos for C_i, the first of them was posted at time t_1. For the other nodes a similar reasoning applies. There is an arc between Alice and Mary because Mary first liked a video of Alice and then posted a video for C_i. There is no arc between Alice and Bob because Bob posted a video for C_i after the videos posted by Alice but he did not put a like for any video of Alice.

To give an idea of the variety of the obtained social networks (and, therefore, of the corresponding challenges), in Figure 2 (resp., 3), we show a representation of those associated with non-dangerous (resp., dangerous) challenges.

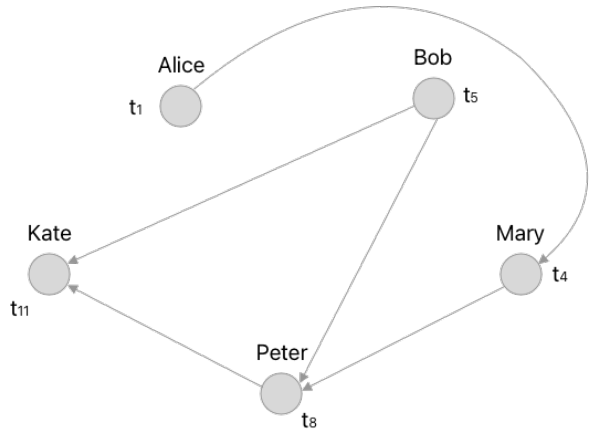

Figure 1: An example of a network corresponding to a challenge

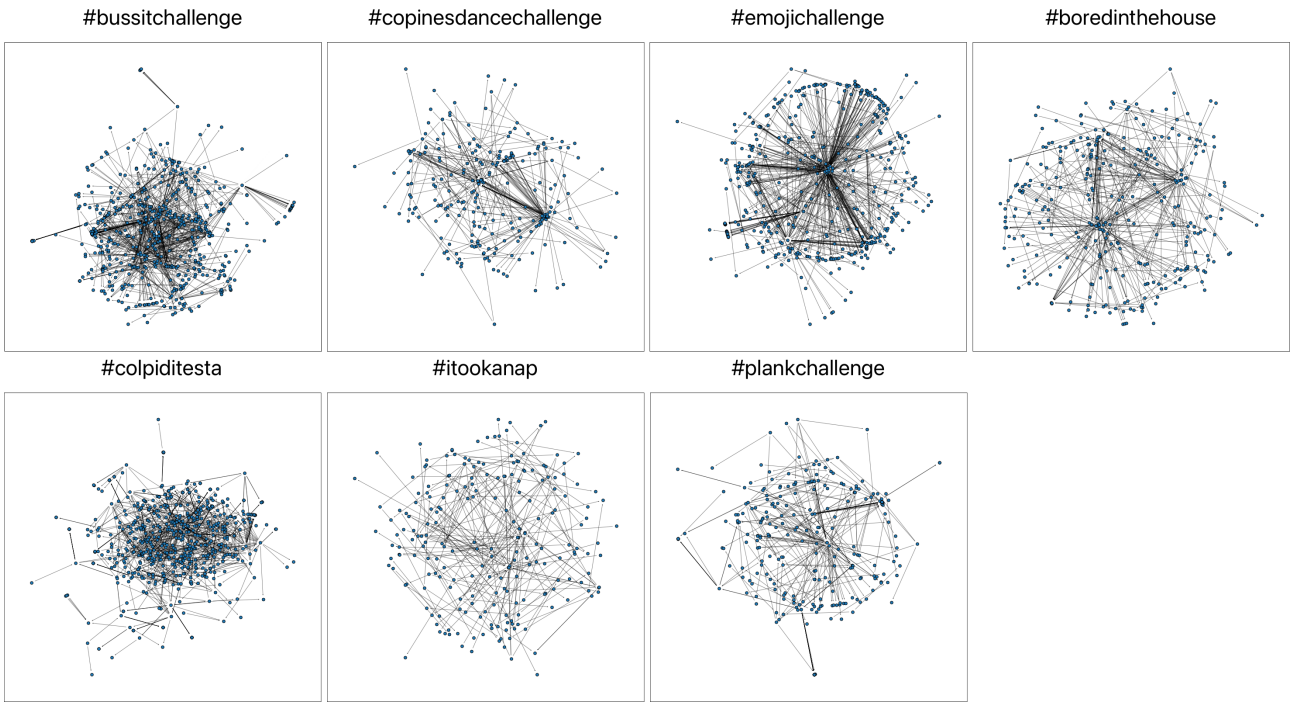

Figure 2: Structure of non-dangerous networks

5 A preliminary analysis of challenges

In this section, we begin with a preliminary analysis of the networks associated with the challenges of our dataset. It serves a dual purpose, namely: *(i)* verifying if there are structural differences between the networks associated with the two types of challenges; *(ii)* identifying interesting insights to investigate whether the user communities related to the two types of challenges have different evolutions or not, which is the core of our paper. In Tables 5 and 6, we report the values of the basic structural parameters for the two types of networks. The analysis of these tables allows us to draw the

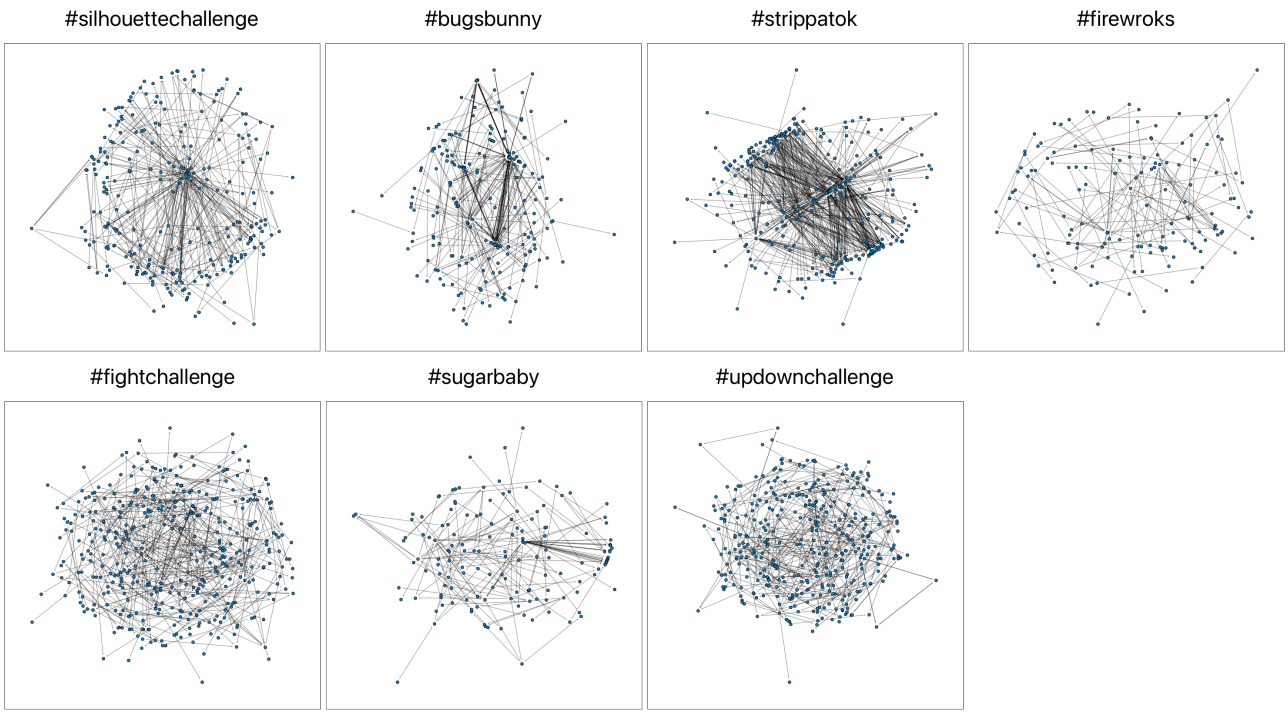

Figure 3: Structure of dangerous networks

following conclusions: *(i)* the size of the networks representing non-dangerous challenges is generally greater than that of the networks associated with dangerous challenges; *(ii)* the average degree and the average clustering coefficient of the two kinds of network are comparable; *(iii)* the density of the networks associated with dangerous challenges is higher than the one of the networks associated with non-dangerous challenges.

To assess the statistical significance of these results, we performed the appropriate t-tests and computed the corresponding p-values.

For case *(i)* the null hypotheses were: H0: "The number of nodes in the non-dangerous networks and that in the dangerous networks are equal" and H0: "The number of arcs in the non-dangerous networks and that in the dangerous networks are equal". In the first case we obtained a p-value equal to 0.012, while in the second case the p-value was equal to 0.014. In both cases the value is less than 0.05. Therefore, we can conclude that the two null hypotheses can be rejected.

For case *(ii)* the null hypotheses were: H0: "The average degree of the non-dangerous networks and that of the dangerous networks are equal" and H0: "The average clustering coefficient of the non-dangerous networks and that of the dangerous networks are equal". In the first case, we obtained a p-value equal to 0.85, while in the second case the p-value was equal to 0.91. In both cases this value is much greater than 0.05, so we can conclude that the two null hypotheses cannot be rejected.

For case *(iii)* the null hypothesis was: H0: "The density of non-dangerous networks and that of dangerous networks are equal". In this case, we obtained a p-value of 0.024, which is less than 0.05. Therefore, we can conclude that the null hypothesis can be rejected.

Finally, we point out that, in the previous tests, when the variances were not statistically different, we used the classical t-test. Instead, in the other cases, we adopted the Welch's t-test [6]. To assess whether the variances were statistically different we used the Bartlett's t-test [5].

Challenge	*Number of nodes*	*Number of arcs*	*Average degree*	*Average clustering coefficient*	*Density*
`#bussitchallenge`	618	708	1.14	0.0047	0.0019
`#copinesdancechallenge`	237	226	0.96	0	0.0040
`#emojichallenge`	440	498	1.13	0.0053	0.0026
`#colpiditesta`	691	843	1.22	0.0015	0.0018
`#boredinthehouse`	306	309	1.01	0.0018	0.0033
`#itookanap`	219	201	0.92	0	0.0042
`#plankchallenge`	271	266	0.98	0.0079	0.0036
Average Value	397.429	435.857	1.051	0.0030	0.0031

Table 5: Basic structural characteristics of non-dangerous networks

Challenge	*Number of nodes*	*Number of arcs*	*Average degree*	*Average clustering coefficient*	*Density*
`#silhouettechallenge`	262	259	0.98	0	0.0037
`#bugsbunny`	212	239	1.13	0	0.0053
`#strippatok`	297	519	1.74	0.0025	0.0059
`#firewroks`	141	111	0.79	0.0083	0.0056
`#fightchallenge`	409	339	0.83	0.0009	0.0020
`#sugarbaby`	151	143	0.94	0.0035	0.0061
`#updownchallenge`	243	199	0.81	0.010	0.0033
Average Value	245	258.429	1.031	0.0036	0.0046

Table 6: Basic structural characteristics of dangerous networks

After examining the characteristics of the networks associated with the two types of challenges, we proceeded to examine their corresponding videos. Their main characteristics are shown in Table 7. From the analysis of this table we can deduce that: *(i)* the two types of challenges have videos with similar duration; *(ii)* non-dangerous challenges have a higher average number of music tracks than dangerous challenges; *(iii)* dangerous challenges have a higher average number of likes, comments, shares and views than non-dangerous challenges. In order to assess the statistical significance of these results, we carried out the suitable t-tests and computed the corresponding values. As in the previous cases, when the variances were not statistically different, we adopted the classical t-test; otherwise, we employed the Welch's t-test. To verify whether the variances were statistically different we used the Bartlett's t-test.

For case *(i)* the null hypothesis was: H0: "The average video duration in the non-dangerous challenges and that in the dangerous challenges are equal". In this case, we obtained a p-value equal to 0.88, which is much greater than 0.05. Therefore, we can conclude that the null hypothesis cannot be rejected.

For cases *(ii)* and *(iii)* the null hypotheses were: *(1)* H0: "The average number of music tracks used in the non-dangerous challenges and the average number of music tracks used in the dangerous

challenges are equal", *(2)* H0: "The average number of likes in the non-dangerous challenges and that in the dangerous challenges are equal", *(3)* H0: "The average number of comments in the non-dangerous challenges and that in the dangerous challenges are equal" , *(4)* H0: "The average number of shares in the non-dangerous challenges and that in the dangerous challenges are equal", and *(5)* H0: "The average number of views in the non-dangerous challenges and that in the dangerous challenges are equal". In these five cases we obtained the following p-values: *(1)* 0.012, *(2)* 0.014, *(3)* 0.022, *(4)* 0.018, and *(5)* 0.007. All the five p-values are less than 0.05. Therefore, we can conclude that all the five null hypotheses can be rejected.

Parameter	*Non-dangerous challenges*	*Dangerous challenges*
Average video duration (seconds)	21.39	20.38
Average number of music tracks used in a challenge	208	126.20
Average number of likes	178,104.13	249,152.12
Average number of comments	1,970.03	2,559.98
Average number of shares	5,456.83	6,990.26
Average number of views	1,471,020.16	2,070,632.01

Table 7: Differences between the main basic characteristics of videos for non-dangerous and dangerous challenges

At this point, we looked at the authors of the videos posted for the two types of challenges and examined their main characteristics. These are shown in Table 8. From the analysis of this table we can deduce that: *(i)* the average number of followers is comparable for the two types of authors; *(ii)* the authors of non-dangerous challenges tend to put more likes, follow many more authors and post many more videos than the ones of dangerous challenges; *(iii)* the authors of dangerous challenges receive many more likes than the ones of non-dangerous challenges. Once again, we employed the approach already described for Tables 5 and 6 to verify the statistical significance of the results obtained.

In particular, for case *(i)* the null hypothesis was: H0: "The average number of users following authors of non-dangerous challenges and the average number of users following authors of dangerous challenges are equal". In this case, we obtained a p-value equal to 0.55, which is much greater than 0.05. Therefore, we can conclude that the null hypothesis cannot be rejected.

For cases *(ii)* and *(iii)* the null hypotheses were: *(1)* H0: "The average number of likes put by authors of non-dangerous challenges and that put by authors of dangerous challenges are equal", *(2)* H0: "The average number of likes received by authors of non-dangerous challenges and that received by authors of dangerous challenges are equal", *(3)* H0: "The average number of users followed by authors of non-dangerous challenges and the average number of users followed by authors of dangerous challenges are equal" , *(4)* H0: "The average number of videos published by authors of non-dangerous challenges and that published by authors of dangerous challenges are equal". In these four cases we obtained the following p-values: *(1)* $6.57 \cdot 10^{-4}$, *(2)* $8.46 \cdot 10^{-6}$, *(3)* 0.0042, and *(4)* 0.014. All the four p-values are less than 0.05. Therefore, we can conclude that all the four null hypotheses can be rejected.

Finally, we considered the evolution of user communities associated with non-dangerous and dangerous challenges over time. In this preliminary analysis, we focused only on the variation in the number of users. The results obtained are shown in Table 9. Examining this table, we can see im-

Parameter	*Non-dangerous challenges*	*Dangerous challenges*
Average number of likes put by an author	17,730.52	11,998.711
Average number of likes received by an author	7,033,150.71	12,080,102.18
Average number of users followed by an author	1,357.08	670.24
Average number of users following an author	400,593.58	447,762.28
Average number of videos published	384.05	263.13

Table 8: Differences between the main basic characteristics of the authors of videos for non-dangerous and dangerous challenges

portant differences between non-dangerous and dangerous challenges. First, the average lifespan of dangerous challenges is longer than the one of non-dangerous challenges. Also, the growth of the number of users in non-dangerous challenges is more gradual than the one in dangerous challenges. Indeed, as for non-dangerous challenges, when passing from 5% to 10%, 15% and 20% of the lifespan, the number of users[4] grows from 2.16% to 35.32%, 43.28% and 45.15% of the final number of users. Instead, as for dangerous challenges, when we pass from 5% to 10%, 15% and 20% of the lifespan, the number of users grows from 0.90% to 3.10%, 9.12% and 23.93% of the final number of users. For all these parameters we adopted the approach already described for Tables 5, 6 and 7 to verify the statistical significance of the results obtained. In these cases the null hypotheses were: *(1)* H0: "The average lifespan of non-dangerous challenges and that of dangerous challenges are equal", *(2)* H0: "The average number of network nodes at 5% of lifespan in non-dangerous challenges and that in dangerous challenges are equal", *(3)* H0: "The average number of network nodes at 25% of lifespan in non-dangerous challenges and that in dangerous challenges are equal" , *(4)* H0: "The average number of network nodes at 50% of lifespan in non-dangerous challenges and that in dangerous challenges are equal", *(5)* H0: "The average number of network nodes at 75% of lifespan in non-dangerous challenges and that in dangerous challenges are equal", and *(6)* H0: "The average number of network nodes in non-dangerous challenges and that in dangerous challenges are equal". In these six cases we obtained the following p-values: *(1)* 0.015, *(2)* $2.23 \cdot 10^{-6}$, *(3)* $7.54 \cdot 10^{-7}$, *(4)* $8.65 \cdot 10^{-8}$, *(5)* 0.011, and *(6)* 0.028. All the six p-values are less than 0.05. Therefore, we can conclude that all the six null hypotheses can be rejected.

This preliminary analysis seems to suggest that the communities of users associated with the two types of challenges have very different growth dynamics. Finding out whether this conjecture is true and, if so, investigating these differences in detail and finding evolutionary patterns characterizing them represent the core of this paper.

6 Analysis of the evolution of user communities for non-dangerous and dangerous challenges

In this section, we present the core of this paper, which is the identification of possible evolutionary patterns that characterize the communities of users related to TikTok challenges and allow the distinction of non-dangerous challenges from dangerous ones.

[4]Recall that there is a biunivocal correspondence between a user of a challenge and a node of the corresponding network.

Parameter	*Non-dangerous challenges*	*Dangerous challenges*
Average challenge lifespan (days)	405	550.17
Average number of network nodes at 5% of lifespan	8.6	2.2
Average number of network nodes at 25% of lifespan	140.4	7.6
Average number of network nodes at 50% of lifespan	172	22.4
Average number of network nodes at 75% of lifespan	179.4	58.8
Average number of network nodes (100% of lifespan)	397.43	245.67

Table 9: Differences between the growth of user communities associated with non-dangerous and dangerous challenges

The first step of this research consists in analyzing the temporal evolution of the 14 challenges in our dataset. In particular, we want to determine if the lifespans of the various challenges contain common typical intervals. Examples of such intervals might be: *(i)* the interval in which the challenge is born and a very first community of users begins to develop; *(ii)* the interval in which the challenge is enormously successful and becomes viral; *(iii)* the interval in which the challenge's popularity begins to decline; *(iv)* the interval in which the challenge has become obsolete and is abandoned. In addition, we want to test whether these intervals are characterized by very different behaviors from the user communities associated with challenges. Finally, behavioral differences among user communities could occur not only based on the type of intervals, but also, and perhaps most importantly, based on the type (i.e., non-dangerous and dangerous) of challenge.

To begin our research, we considered how the size of each community evolved during the lifespan of the corresponding challenge. As seen in Section 4, the community associated with each challenge can be modeled as a social network and there is a biunivocal correspondence between the users of a community and the nodes of the corresponding social network.

We now consider a plot whose x-axis represents the lifespan of a challenge and whose y-axis denotes the number of members of the community associated with it or, equivalently, the number of nodes of the corresponding social network. If we subdivide the lifespan into suitable time slots (also very small), consider the number of social network nodes in correspondence to each time slot, find the corresponding points in the diagram and join them, we obtain a broken line, which denotes the variation of the community size during the challenge lifespan. We chose a very fine granularity and, in fact, we divided the lifespan into 100 time slots. With this choice, the broken line becomes very detailed, providing a very accurate representation of how the community size varies over time. However, for reasons that will become clear later, we needed a continuous function, instead of a broken line. To obtain it, we interpolated the points of the broken line using a univariate spline.

To test whether the difference between the broken lines and the curves obtained from the interpolation is acceptable, we computed the Mean Absolute Error (MAE) by considering 100 additional equidistant points for each time slot (and, thus, 10,000 points for each lifespan). Then, we normalized the MAE value at each point to the value of the broken line at the same point. Table 10 shows the results obtained. The analysis of this table reveals that the average values of the normalized MAE are very low. This allows us to conclude that the interpolation performed by us is acceptable.

To analyze how the communities associated with challenges evolve over time, we found it useful to

Non-dangerous Challenge	*Normalized MAE*	*Dangerous Challenge*	*Normalized MAE*
#bussitchallenge	0.012	#silhouttechallenge	0.017
#copinesdancechallenge	0.015	#bugsbunny	0.017
#emojichallenge	0.021	#strippatok	0.023
#colpiditesta	0.025	#firewroks	0.026
#boredinthehouse	0.011	#fightchallenge	0.014
#itookanap	0.015	#sugarbaby	0.021
#plankchallenge	0.018	#updownchallenge	0.026

Table 10: Normalized MAE between the continuous function returned by the univariate spline interpolation and the real values for non-dangerous challenges (at left) and dangerous ones (at right)

identify the points of the lifespan where their characteristics change. Since, up to this point, the most important characteristic that we know is community size, this implies considering the points at which the broken line or the corresponding interpolation curve inverts. This is the reason why we used the interpolation curve with the univariate spline. In fact, in this way, we have a continuous function and the points where it inverts are given by the ones where it reaches a maximum or a minimum.

More formally, let C_i be a challenge, let $\mathcal{N}_i$ be the corresponding social network, and let $\nu_i(\cdot)$ be the function representing the change in the number of nodes of $\mathcal{N}_i$ during the lifespan of C_i; in other words, $\nu_i(\cdot)$ is the interpolation curve described above. To identify the points in the lifespan where $\nu_i(\cdot)$ has a maximum or a minimum, we compute the first derivative $\nu_i'(\cdot)$ of $\nu_i(\cdot)$ and check the points where it becomes null. Let $X_i = \{x_{i_1}, x_{i_2}, \cdots, x_{i_N}\}$ be the set of such points; we can split the lifespan of C_i into $N-1$ intervals (x_q, x_{q+1}), $1 \leq q \leq N-1$, such that $\nu_i(\cdot)$ is always increasing or always decreasing within each of them. As we will see in the following, these intervals represent an essential tool of our analysis because we will use them to look for the evolutionary patterns of communities capable of distinguishing non-dangerous challenges from dangerous ones.

Figures 4 and 5 show the trends of the function $\nu_i(\cdot)$ for each non-dangerous and dangerous challenge, respectively. They also show the corresponding intervals. Already from this first visual analysis, we can observe that, in the two kinds of challenge, the corresponding communities show completely different dynamics. Capturing and formalizing such dynamics represent the objective of the next sections.

6.1 Capturing community evolution during a challenge lifespan

In order to capture the evolution of communities during a challenge lifespan, it is first necessary to identify features capable of representing this evolution in detail and from multiple perspectives. To this end, we are helped by the social network-based model that we introduced in Section 4. Thanks to this model, given a challenge C_i and the social network $\mathcal{N}_i$ representing its community at a given interval $\mathcal{I}$, during which the trend of $\nu_i(\cdot)$ is always increasing or always decreasing, it is possible to identify 18 features of interest. These are:

- `node_number`: number of nodes of $\mathcal{N}_i$;
- `arc_number`: number of arcs of $\mathcal{N}_i$;
- `density`: density of $\mathcal{N}_i$;

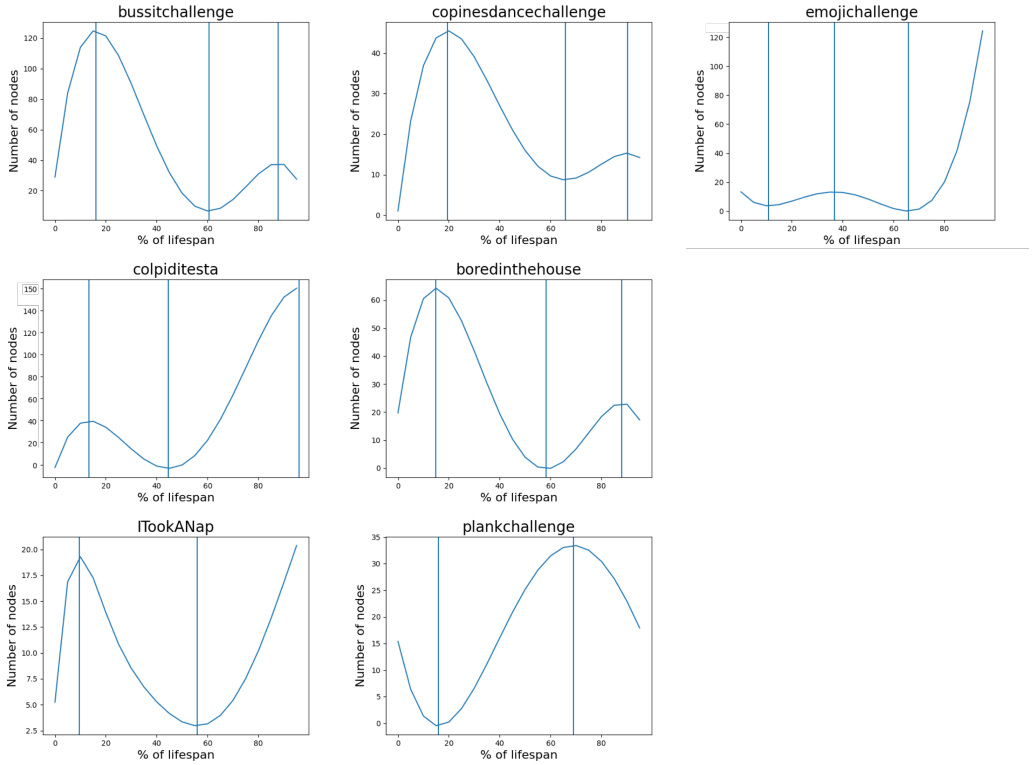

Figure 4: Trends and intervals of $\nu_i(\cdot)$ for non-dangerous challenges

- `conn_components_number`: number of connected components of $\mathcal{N}_i$;
- `max_conn_comp_node_number`: number of nodes of the maximum connected component of $\mathcal{N}_i$;
- `avg_indegree_centrality`: average indegree centrality of the nodes of $\mathcal{N}_i$;
- `avg_outdegree_centrality`: average outdegree centrality of the nodes of $\mathcal{N}_i$;
- `avg_eigenvector_centrality`: average eigenvector centrality of the nodes of $\mathcal{N}_i$;
- `avg_pagerank`: average PageRank of the nodes of $\mathcal{N}_i$;
- `avg_closeness_centrality`: average closeness centrality of the nodes of $\mathcal{N}_i$;
- `avg_clustering_coefficient`: average clustering coefficient of the nodes of $\mathcal{N}_i$.
- `radius_max_conn_comp`: radius of the maximum connected component of $\mathcal{N}_i$;
- `diameter_max_conn_comp`: diameter of the maximum connected component of $\mathcal{N}_i$;
- `perc_nodes_in_max_conn_comp`: percentage of nodes of $\mathcal{N}_i$ belonging to its maximum connected component;
- `avg_eccentricity`: average eccentricity of the nodes of $\mathcal{N}_i$;

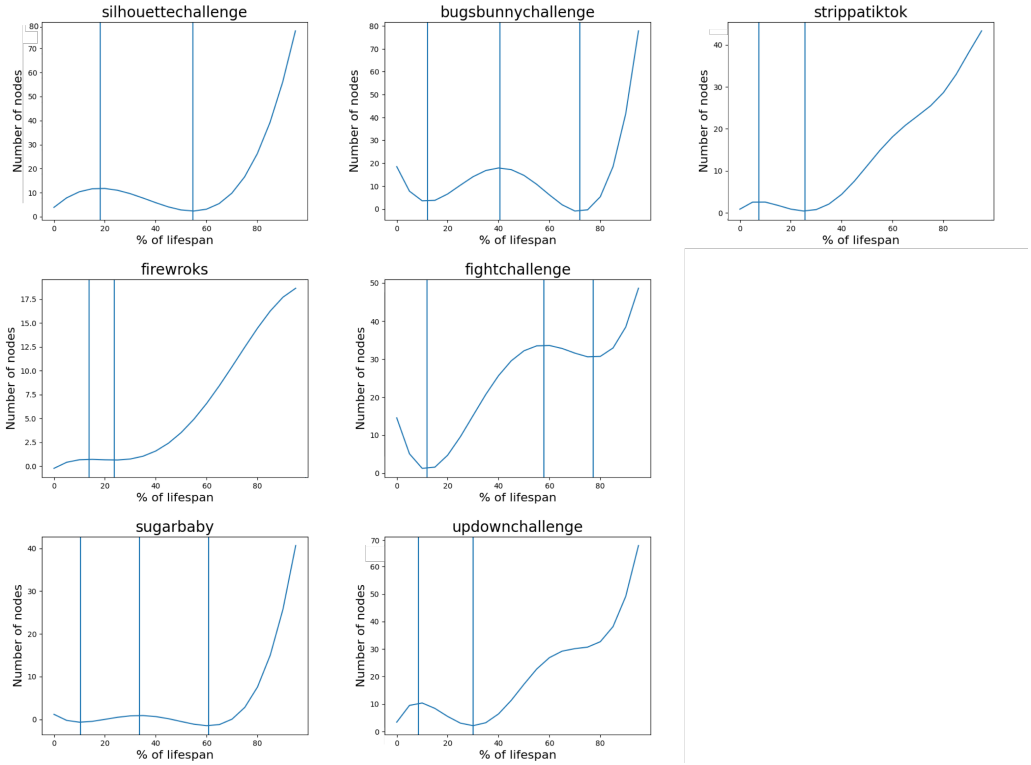

Figure 5: Trends and intervals of $\nu_i(\cdot)$ for dangerous challenges

- `avg_path_length`: average length of the paths of $\mathcal{N}_i$;
- `max_ego_network_node_number`: number of nodes present in the ego-network with the maximum size in $\mathcal{N}_i$;
- `avg_ego_network_node_number`: average number of nodes in the ego-networks of $\mathcal{N}_i$.

As we can see, we have a lot of available features, and managing all of them can be complex. Therefore, we decided to check for possible correlations between them. In fact, if a group of features is correlated, we can keep only one of them and filter out the others. Figure 6 shows the correlation matrix we obtained by applying the Pearson's correlation coefficient [6] to the pairs of features identified above.

Considering the various groups of correlated features and choosing one for each group, we identified the following features to keep for the next analyses:

- `conn_components_number`;
- `avg_indegree_centrality`;
- `avg_outdegree_centrality`;
- `avg_clustering_coefficient`;
- `perc_nodes_in_max_conn_comp`;

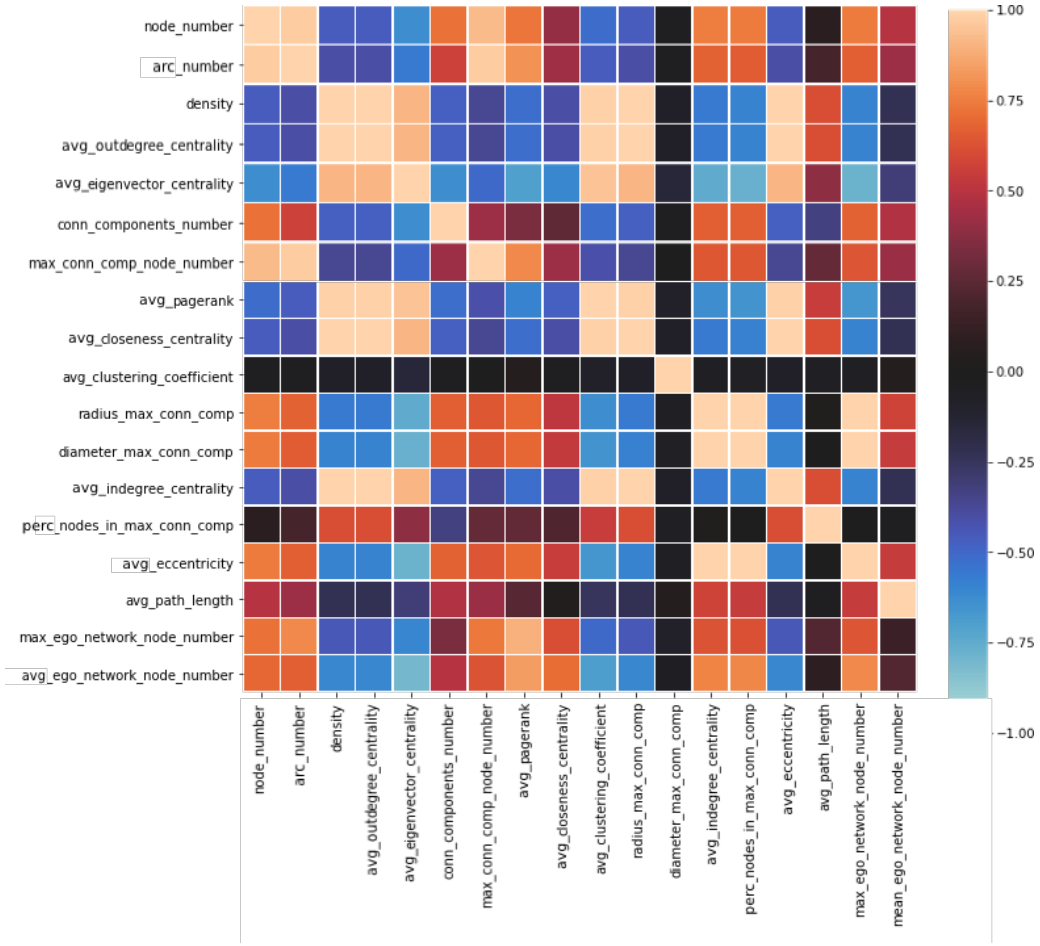

Figure 6: Correlation matrix for the 18 features representing the behavior of the communities during a challenge

- `avg_path_length`;
- `avg_ego_network_node_number`.

6.2 Detecting the similarities and differences of the evolutionary dynamics of communities

In the previous section, we have identified a list of features that can describe the behavior of the community of users associated with a challenge during a time interval. In this section, we want to use these features to group the intervals related to the lifespan of the 14 challenges of our dataset into clusters that are homogeneous from the perspective of the evolutionary dynamics of the communities involved.

First of all, we considered a new dataset formed by a single table whose rows represent the intervals of the 14 challenges under consideration and whose columns are associated to the 7 selected features. The element (h, k) of this table indicates the value assumed by the k^{th} feature in the h^{th} interval.

Afterwards, we applied a clustering technique to group the intervals into homogeneous clusters from the user community behavior perspective. Specifically, we chose the Autoclass [9] clustering algorithm. The reason for this choice lies in the fact that this algorithm, among the various positive properties characterizing it, also has that of being able to automatically determine the number of clusters. This property was particularly important in our case because it was not possible to make any a priori conjecture on this number, and the application of the elbow method carried out with k-means returned no results. Applying Autoclass to our dataset we obtained four clusters. In order to visualize them, we applied the Principal Component Analysis (hereafter, PCA) [25] to the dataset. In this way, we reduced the number of dimensions from 7 to 2, which allowed us to visualize data into a bidimensional plane whose axes correspond to the two dimensions returned by PCA. This visualization improved the interpretation of the clusters obtained. We adopted linear PCA for dimensionality reduction. Actually, we also considered other approaches to perform this task, such as t-SNE and several forms of kernel PCA. However, linear PCA is the one that provided the best tradeoff between the needs of visualization, interpretability and determinism of result.

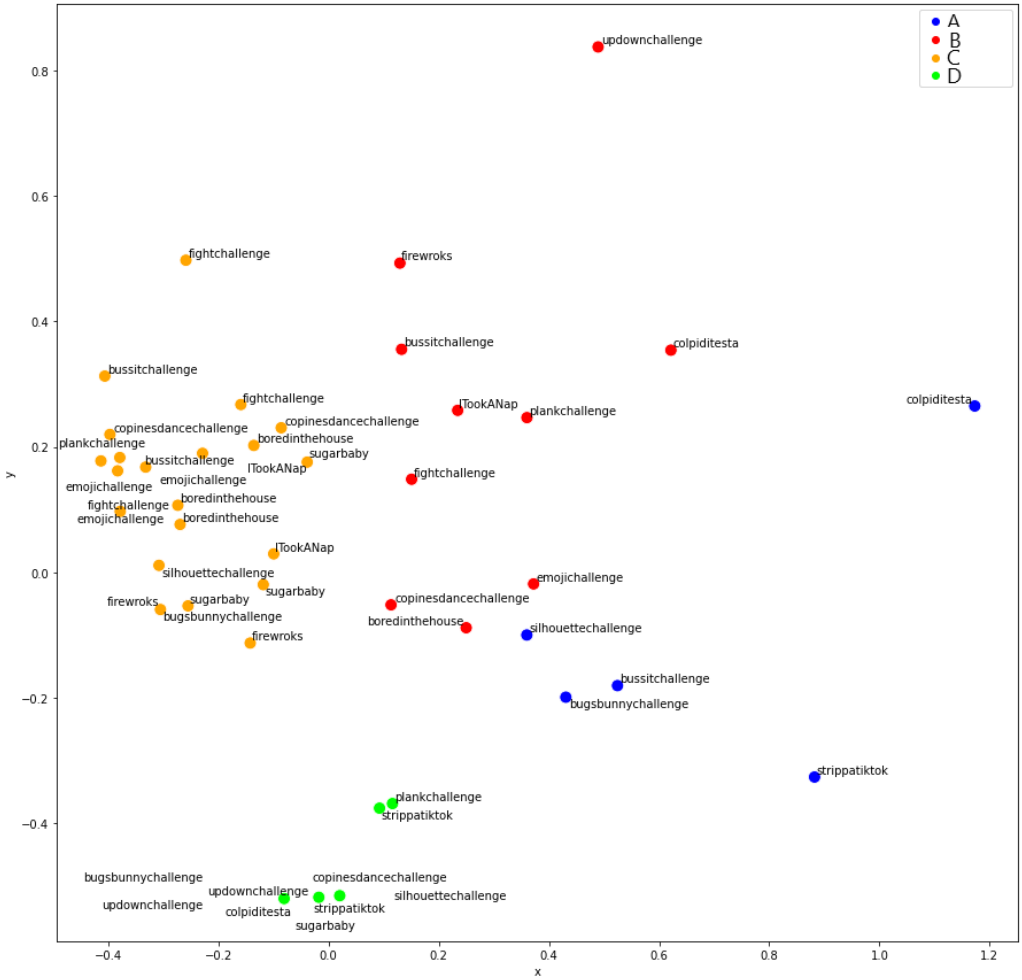

Figure 7: The four clusters of intervals returned by Expectation Maximization

After identifying clusters and representing them in a bidimensional plane, we tried to understand what each of them denoted in terms of the behavior and the dynamics of the challenge communities during the time intervals belonging to it. At the end of this activity, we drew the following characterizations:

- *Cluster A*: during the intervals belonging to this cluster, networks are characterized by a quite high number of nodes and a high number of connected components. The nodes of each connected component have a high average indegree and average outdegree. This implies that the corresponding communities consist of highly connected users. As a confirmation of the latter property, the average size of the ego networks is large and the average clustering coefficient is high.

- *Cluster B*: during the intervals belonging to this cluster, networks are characterized by a very high number of nodes and a rather high number of connected components (although less than in Cluster A). The maximum connected component includes most of the nodes, while the other ones are all made up of few nodes, albeit their number is still high. The average clustering coefficient and the average size of the ego networks remain very high, even if this is mainly due to the contribution of the nodes of the maximum connected component.

- *Cluster C*: during the intervals belonging to this cluster, networks are characterized by a limited number of nodes and a certain number of connected components. The nodes of each connected component have a small-medium average indegree and average outdegree. The average size of the ego networks is small and the average clustering coefficient is medium-small.

- *Cluster D*: during the intervals belonging to this cluster, networks have a high number of nodes and a high number of connected components. The nodes of each connected component have a medium average indegree and a medium average outdegree. Both the average size of the ego networks and the average clustering coefficient are medium-high.

In Figure 8, we show an example of the structure of a user community associated with a challenge for each cluster.

To give a quantitative idea of the characteristics of clusters, in Table 11 we show the average values taken in each cluster by the seven features we selected to represent the lifespan intervals.

Feature	*Cluster A*	*Cluster B*	*Cluster C*	*Cluster D*
`conn_components_number`	86	92	12	65
`avg_indegree_centrality`	68	74	37	55
`avg_outdegree_centrality`	152	164	11	84
`avg_clustering_coefficient`	0.0021	0.0025	0.00009	0.00072
`perc_nodes_in_max_conn_comp`	38.02%	79.74%	41.54%	56.89%
`avg_path_length`	21	23	3	18
`avg_ego_network_node_number`	301	312	24	68

Table 11: Average value taken in each cluster by the features selected to represent the lifespan intervals

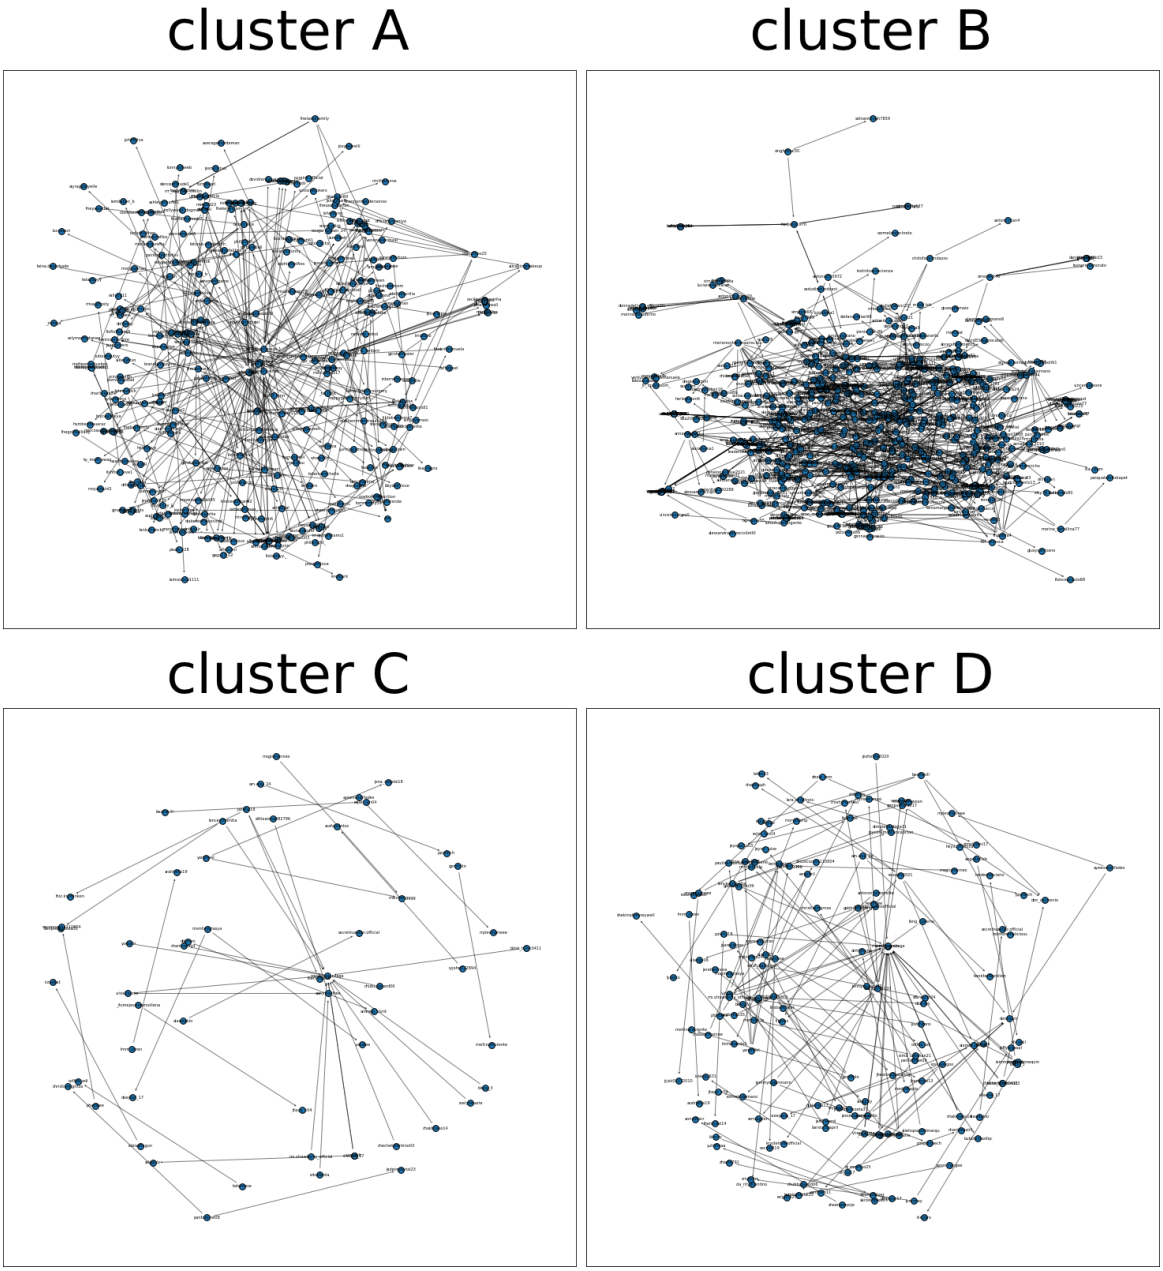

Figure 8: Example of the structure of a user community associated with a challenge for each cluster

7 Searching for evolutionary patterns in the challenge lifespans

After grouping the intervals into clusters, and after identifying the characteristics of each cluster, we tested whether there were evolutionary patterns characterizing the communities of non-dangerous and dangerous challenges while also providing the capability of distinguishing them. To this end, we considered the lifespans of the 14 challenges of the dataset and, for each of the corresponding intervals, we recorded the cluster to which it belonged. If two consecutive intervals belonged to the same cluster,

we recorded them only once. At the end of this process, we obtained the sequences of intervals shown in Table 12.

Non-dangerous Challenge	*Evolutionary Paths*	*Dangerous Challenge*	*Evolutionary Paths*
#bussitchallenge	C, B, D	#silhouttechallenge	C, A
#copinesdancechallenge	C, A, B, D	#bugsbunny	C, D
#emojichallenge	A, B, D	#strippatok	C, D
#colpiditesta	C, A, D	#firewroks	C, A, B
#boredinthehouse	A, D	#fightchallenge	C, A
#itookanap	C, A, D	#sugarbaby	C, A, D
#plankchallenge	C, B, D	#updownchallenge	C, B

Table 12: Sequences of intervals for non-dangerous and dangerous challenges

Examining these sequences, we can draw some observations. In particular:

- In the non-dangerous challenges, there is no dominant pattern although intervals of type C and D are frequent. Specifically, an interval of type D is present in each non-dangerous challenge.
- Dangerous challenges always begin with an interval of type C, whereas they end with intervals of type A, B or D.

Examining the description of clusters in Section 6.2, we can note that the user communities during the intervals belonging to clusters A and B have similar features. Also observing Figure 7 we can see that cluster B can be seen as an extension of cluster A. Therefore, we decided to analyze the data corresponding to the intervals of these clusters in more detail. We have previously seen that:

- The intervals of cluster A are characterized by networks with a high number of connected components. The average indegree and outdegree of the network nodes are high. As a result, during these intervals, there are many connections between users. This is also witnessed by the average clustering coefficient that is very high.
- The intervals of type B are characterized by networks with a rather high number of connected components and high average indegree and outdegree of the network nodes. The main difference with the intervals of type A is that, in this case, the maximum connected component contains most of the network nodes. In fact, the other connected components generally consist of pairs of nodes.

Despite the main difference mentioned above, and other small existing ones, we can hypothesize that the two clusters of intervals A and B represent the same reality. In particular, given the high average indegree, average outdegree, average clustering coefficient and the large size of ego networks, we can hypothesize that these intervals represent the peak of the evolution of a challenge.

In order to test our hypothesis, we performed a t-test [6], based on the following null hypothesis: H0: "The means of the samples for the intervals of clusters A and B are equal". Prior to performing it, we had to test whether the items in the two samples had comparable variances or not. In fact, this step is necessary to choose whether to perform the classical t-test (used when the two samples

have comparable variances) or the Welch's t-test (used otherwise) [6]. In order to decide on the comparability of the variances of the intervals of the clusters A and B, we performed the Bartlett's t-test [5]. It allows us to determine whether two samples with different numbers of items have the same variance or not. More formally, we applied the Bartlett's t-test with the following null hypothesis: H0: "The variances of the samples for the intervals of clusters A and B are equal". We computed the corresponding p-value and saw that it was equal to 0.52, which is much higher than the classical threshold of 0.05 generally considered for this parameter. Therefore, the null hypothesis cannot be rejected. As a consequence of this fact, in order to test whether the difference between the intervals of clusters A and B was statistically significant, we had to adopt the classic t-test and not the Welch's one.

Applying the classic t-test on the null hypothesis H0: "The means of the samples for the intervals of clusters A and B are equal", we obtained a p-value of 0.63. This is much greater than 0.05 and allowed us to conclude that the null hypothesis cannot be rejected. In turn, this implied that the clusters A and B were statistically equivalent and represented two very similar scenarios, despite the previously highlighted differences.

Thanks to this result, it was possible to substitute A for B in all the interval sequences of the challenges under consideration.

Observe that, after determining the equivalence between the intervals of A and B, we have three kinds of interval, namely: *(i)* intervals of type A, whose characteristics described above suggest that they correspond to the peak of a challenge; *(ii)* intervals of type C, whose characteristics suggest that they are the initial ones in a challenge; *(iii)* intervals of type D, whose characteristics suggest that they are the ones relating to the end of a challenge.

Now, after the substitution of B with A, and recalling that our evolutionary pattern model states that two consecutive intervals of the same cluster are represented only once, the sequences of intervals that characterize non-dangerous and dangerous challenges are shown in Table 13.

Non-dangerous Challenge	*Evolutionary Paths*	*Dangerous Challenge*	*Evolutionary Paths*
#bussitchallenge	C, A, D	#silhouttechallenge	C, A
#copinesdancechallenge	C, A, D	#bugsbunny	C, D
#emojichallenge	A, D	#strippatok	C, D
#colpiditesta	C, A, D	#firewroks	C, A
#boredinthehouse	A, D	#fightchallenge	C, A
#itookanap	C, A, D	#sugarbaby	C, A, D
#plankchallenge	C, A, D	#updownchallenge	C, A

Table 13: Sequences of intervals for non-dangerous and dangerous challenges after the verification of the hypothesis that A and B are equivalent

Thanks to this result, we were able to identify some evolutionary patterns characterizing non-dangerous and dangerous challenges. Furthermore, since these patterns are different in the two cases, they also allow the distinction of non-dangerous challenges from dangerous ones.

Let us first examine non-dangerous challenges. In this case, we always have the presence of a sequence of intervals of type A, D. This sequence is very often preceded by an interval of type C, so that we have an evolutionary pattern of type C, A, D. We argued that the typical evolutionary pattern of a non-dangerous sequence is C, A, D. In fact, the challenges showing a pattern of type A, D already

existed when our research on them began, although the interactions with users that they were able to elicit were almost negligible.

Let us now examine dangerous challenges. In this case, unlike the previous one, there is no single sequence of intervals characterizing most of them. Instead, we identified two dominant sequences that correspond to two different "fates" generally characterizing the challenges of this type. They are:

- C, A: these challenges had a standard initial phase with an interval of type C; then, they reached a peak phase. Finally, they almost suddenly ceased to have meaningful interactions with users.
- C, D: these challenges had an initial phase, which was followed by a decay one. In other words, they never reached the peak. They were born, survived for a certain period on the network, and then died.

In order to verify the suitability of our approach, we decided to test it on a new dataset, larger than the previous one. It consists of 300 challenges (150 non-dangerous ones and 150 dangerous ones). As dangerous challenges are very rare, the 150 dangerous challenges of our dataset were obtained from 25 real challenges using the oversampling technique implemented through bootstrap [6]. Due to space limitations, we cannot present in detail the 175 real challenges we used. However, in Table 14, we report the aggregate values of some fields referring to them.

Parameter	*Non-dangerous challenges*	*Dangerous challenges*
Publication month of the first video	From 2018-01 to 2019-12	From 2017-01 to 2020-12
Publication month of the last video	From 2018-03 to 2021-02	From 2017-02 to 2021-04
Average lifespan in days	523.45	364.73
Average number of videos	542.54	366.55
Average number of likes received	184,234.52	247,325.48
Average number of comments received	1,984.05	2,654.03
Average number of shares	5,548.72	7,002.44
Average number of views	1,475,042.16	2,084,544.06

Table 14: Aggregate values of some fields referring to non-dangerous and dangerous challenges

The results obtained are the following:

- As for non-dangerous challenges:
  - 132 (i.e., 88.00% of them) followed the evolutionary pattern C, A, D. This is the only significant one we identified for this type of challenges.
  - 18 (i.e., 12.00% of them) followed a variety of other sequences of intervals.
- As for dangerous challenges:
  - 65 (i.e., 43.33% of them) followed the evolutionary pattern C, A;
  - 69 (i.e., 46.00% of them) followed the evolutionary pattern C, D;
  - 7 (i.e., 4.67% of them) followed the evolutionary pattern C, A, D;

– 9 (i.e., 6.00% of them) followed a variety of other sequences of intervals.

The results obtained represent a confirmation that the evolutionary patterns we detected actually exist for the two types of challenges into consideration and are capable of discriminating them. In addition, these results show that the patterns we found are really able to capture almost all the behaviors of the communities of TikTok challenges.

7.1 Discussion

In the previous section, we have seen that non-dangerous challenges generally follow the evolutionary pattern C, A, D, while dangerous challenges generally follow the evolutionary patterns C, A or C, D. The pattern C, A, D is regular while the patterns C, A and C, D are both irregular, even if for different reasons. In fact, the pattern C, A, D represents a context in which there is the appearance of a new challenge, its growth to a peak and, finally, its decrease more or less slow, but regular. The pattern C, A is typical of a context in which there has been an almost sudden end of user interactions. This may happen because the challenge ran out of steam very quickly or it was recognized by TikTok as dangerous and was stopped or removed from the social network. The pattern C, D is representative of a challenge that had an initial phase, survived for a certain period without never reaching a success, and then decayed.

The knowledge derived from the analyses described in Section 5 tells us that dangerous challenges have fewer authors than non-dangerous ones and that these authors are more connected to each other. This tends to set up a more closed scenario, where authors are mutually self-supportive. This is also evidenced by the fact that dangerous challenges have a higher average number of likes, comments, shares and views than non-dangerous ones, as well as by the fact that the authors of dangerous challenges receive many more likes than the ones of non-dangerous challenges. The greater openness of non-dangerous challenges is evidenced by the fact that their authors tend to follow more authors than the ones of dangerous challenges.

As shown in Table 9, the evolution of the two types of challenges, is very different. The number of authors of non-dangerous challenges grows in a much more regular way than the number of authors of dangerous challenges. The latter grows very little up to 50% - 75% of the lifespan. At this point, in the challenges following the behavioral pattern C, D, it decays without ever having achieved success. Instead, in the challenges following the behavioral pattern C, A, it shows an exponential growth. This suddenly stops and decays either because the challenges are recognized as dangerous by TikTok, and therefore are suppressed, or because they lose their appeal to users. This loss happens quickly and, once again, in a much more irregular way than non-dangerous challenges. In fact, the dangerous challenges having a regular decrease are those following the behavioral pattern C, A, D, which, as we have seen above, are a strict minority of the overall dangerous challenges (i.e., 4.67% in the test described in Section 7).

8 Conclusion

In this paper, we have studied the different characteristics and evolutionary dynamics of the user communities participating in non-dangerous and dangerous TikTok challenges. This study led us to the identification of evolutionary patterns allowing us to discriminate the communities of users participating in the two types of challenges. We point out again that the approach proposed in this paper should be considered a first step in our overall research. Indeed, in its current version, it is able to classify a challenge only near the end of its lifespan, or at least after a rather long period of time since its beginning. However, as we have seen above, defining a mechanism for the early detection of dangerous challenges in TikTok is an important issue, which many researchers are focusing on. In fact, the early detection of dangerous challenges is critical to prevent the latter from being too successful and achieving an exponential growth rate. The early detection of dangerous challenges starting from the evolutionary dynamics of the reference communities can be seen as the final goal of our research, of which the approach proposed in this paper can be considered the first step. In fact, we believe that if we were able to reduce the granularity of the time intervals considered, making it much finer, we could verify the possible extension of our approach to identify behavioral patterns characterizing communities. These patterns would allow the distinction of the dangerous challenges from the non-dangerous ones already at the beginning of their lifespan.

Our approach, based on the analysis of the behavior of hundreds or thousands of users participating in a challenge, is robust to the classical tricks used to bypass the current TikTok's controls. The importance of the detection of dangerous challenges is also motivated by another relevant result we obtained in the paper, namely the fact that when these challenges begin to succeed, they tend to have an exponential growth of the number of their users, even much greater than that of the communities associated with non-dangerous challenges.

In the future, besides investigating the possibility of an early detection of dangerous challenges, we plan to further analyze the evolutionary dynamics of the communities associated with challenges using additional features and concepts derived from Social Network Analysis. Moreover, we plan to further study the distinction between dangerous and non-dangerous challenges by identifying additional criteria allowing the detection of a dangerous challenge as soon as possible and in the most robust possible way. Last, but not the least, we could extend our analysis from TikTok challenges to TikTok trends. Indeed, these last ones have certainly several analogies with challenges, but, at the same time, present also several differences. Consequently, we can assume that many of the results found for challenges can be extended to trends by making suitable modifications, which consider the peculiarities of trends with respect to challenges.

References

[1] N. Aggarwal, S. Agrawal, and A. Sureka. Mining YouTube metadata for detecting privacy invading harassment and misdemeanor videos. In *Proc. of the International Conference on Privacy, Security and Trust (PST'14)*, pages 84–93, Auckland, New Zealand, 2014. IEEE.

[2] N. Alonso-López, P. Sidorenko-Bautistal, and F. Giacomelli. Beyond challenges and viral dance moves: TikTok as a vehicle for disinformation and fact-checking in Spain, Portugal, Brazil, and the USA. *Anàlisi*, 64:65–84, 2021.

[3] J. Azpeitia. *Social Media Marketing and its Effects on TikTok Users*. 2021. PhD Thesis.

[4] J. Bandy and N. Diakopoulos. # TulsaFlop: A Case Study of Algorithmically-Influenced Collective Action on TikTok. *arXiv preprint arXiv:2012.07716*, 2020.

[5] M.S. Bartlett. The effect of non-normality on the t distribution. *Mathematical Proceedings of the Cambridge Philosophical Society*, 31(2):223–231, 1935. Cambridge University Press.

[6] P. Bruce, A. Bruce, and P. Gedeck. *Practical Statistics for Data Scientist, Second Edition.* O'Reilly, Sebastopol, CA, USA, 2020.

[7] C.M. Bruno. A Content Analysis of How Healthcare Workers Use TikTok. *Elon Journal of Undergraduate Research in Communications*, 11(2):5–16, 2020.

[8] N. Cassavia, E. Masciari, C. Pulice, and D. Saccà. Discovering User Behavioral Features to Enhance Information Search on Big Data. *ACM Transactions on Interactive Intelligent Systems*, 7(2), 2017. ACM.

[9] P.C. Cheeseman and J.C. Stutz. Bayesian classification (AutoClass): theory and results. *Advances in knowledge discovery and data mining*, 180:153–180, 1996. Philadelphia, PA, USA.

[10] Q. Chen, C. Min, W. Zhang, X. Ma, R. Evans, et al. Factors driving citizen engagement with government TikTok accounts during the COVID-19 pandemic: Model development and analysis. *Journal of Medical Internet Research*, 23(2):e21463, 2021. JMIR Publications.

[11] X. Chen, K.D.B. Valdovinos, and J. Zeng. # PositiveEnergy Douyin: constructing "playful patriotism" in a Chinese short-video application. *Chinese Journal of Communication*, 14(1):97–117, 2021. Taylor & Francis.

[12] Z. Chen and Q. Zhang. A Survey Study on Successful Marketing Factors for Douyin (Tik-Tok). In *Proc. of the International Conference on Human-Computer Interaction (HCI'21)*, pages 22–42, Washington DC, USA, 2021. Springer.

[13] N. Choudhary, C. Gautam, and V. Arya. Digital marketing challenge and opportunity with reference to TikTok - A new rising social media platform. *International Journal of Multidisciplinary Educational Research*, 9(10), 2020.

[14] E. Corradini, A. Nocera, D. Ursino, and L. Virgili. Defining and detecting k-bridges in a social network: the Yelp case, and more. *Knowledge-Based Systems*, 187:104820, 2020. Elsevier.

[15] E. Corradini, A. Nocera, D. Ursino, and L. Virgili. Investigating negative reviews and detecting negative influencers in Yelp through a multi-dimensional social network based model. *International Journal of Information Management*, 60:102377, 2021. Elsevier.

[16] N. Dakiche, F.B. Tayeb, Y. Slimani, and K. Benatchba. Tracking community evolution in social networks: A survey. *Information Processing & Management*, 56(3):1084–1102, 2019. Elsevier.

[17] M. Davis. "This is For You": An Anthropological Approach to Relationships to TikTok and its Algorithm. Technical report, University of Chicago, 2021.

[18] C. Eickhoff and A.P. de Vries. Identifying suitable YouTube videos for children. In *Networked and electronic media summit (NEM'10)*, Barcelona, Spain, 2010.

[19] A. Fiallos, C. Fiallos, and S. Figueroa. Tiktok and Education: Discovering Knowledge through Learning Videos. In *Proc. of the International Conference on eDemocracy & eGovernment (ICEDEG'21)*, pages 172–176, Quito, Ecuador, 2021. IEEE.

[20] S. Fortunato. Community detection in graphs. *Physics reports*, 486(3-5):75–174, 2010. Elsevier.

[21] T. Fu, C.N. Huang, and H. Chen. Identification of extremist videos in online video sharing sites. In *Proc. of the International Conference on Intelligence and Security Informatics*, pages 179–181, Richardson, TX, USA, 2009. IEEE.

[22] W. Gao, W. Lu, and C. Bu. Evolutionary community discovery in dynamic networks based on leader nodes. In *Proc. of the International Conference on Big Data and Smart Computing (BigComp'16)*, pages 53–60, Hong Kong, China, 2016. IEEE.

[23] C. Guo, J. Wang, and Z. Zhang. Evolutionary community structure discovery in dynamic weighted networks. *Physica A: Statistical Mechanics and its Applications*, 413:565–576, 2014. Elsevier.

[24] M. Haenlein, E. Anadol, T. Farnsworth, H. Hugo, J. Hunichen, and D. Welte. Navigating the New Era of Influencer Marketing: How to be Successful on Instagram, TikTok, & Co. *California Management Review*, 63(1):5–25, 2020.

[25] J. Han, M. Kamber, and J. Pei. *Data Mining: Concepts and Techniques - Third Edition.* 2011. Morgan Kaufmann notes.

[26] J. He and D. Chen. A fast algorithm for community detection in temporal network. *Physica A: Statistical Mechanics and its Applications*, 429:87–94, 2015. Elsevier.

[27] P. Held and R. Kruse. Detecting overlapping community hierarchies in dynamic graphs. In *Proc. of the International Conference on Advances in Social Networks Analysis and Mining (ASONAM'16)*, pages 1063–1070, San Francisco, CA, USA, 2016. IEEE.

[28] J. Herrman. How TikTok is rewriting the world. *The New York Times*, 10, 2019.

[29] Y.Y.U. Ishihara and R. Oktavianti. Personal Branding Influencer di Media Sosial TikTok. *Koneksi*, 5(1):76–82, 2020.

[30] M.B. Jdidia, C. Robardet, and E. Fleury. Communities detection and analysis of their dynamics in collaborative networks. In *Proc. of the International Conference on Digital Information Management (ICDIM'07)*, volume 2, pages 744–749, Lyon, France, 2007. IEEE.

[31] X. Ji, S.A. Chun, P. Cappellari, and J. Geller. Linking and using social media data for enhancing public health analytics. *Journal of Information Science*, 43(2):221–245, 2017. SAGE Publications Sage UK: London, England.

[32] M. Kennedy. 'If the rise of the TikTok dance and e-girl aesthetic has taught us anything, it's that teenage girls rule the internet right now': TikTok celebrity, girls and the Coronavirus crisis. *European Journal of Cultural Studies*, 23(6):1069–1076, 2020. SAGE.

[33] N.H. Khoa, P.T. Duy, H.D. Hoang, D.T.T. Hien, and V.H. Pham. Forensic analysis of TikTok application to seek digital artifacts on Android smartphone. In *Proc. of the International Conference on Computing and Communication Technologies (RIVF'20)*, pages 1–5, Ho Chi Minh City, Vietnam, 2020. IEEE.

[34] Y.H. Kim, D. Lee, N.G. Han, and M. Song. Exploring characteristics of video consuming behaviour in different social media using k-pop videos. *Journal of Information Science*, 40(6):806–822, 2014. Sage Publications Sage UK: London, England.

[35] D. Klug. "It took me almost 30 minutes to practice this". Performance and Production Practices in Dance Challenge Videos on TikTok. *arXiv preprint arXiv:2008.13040*, 2020.

[36] D. Klug, Y. Qin, M. Evans, and G. Kaufman. Trick and Please. A Mixed-Method Study On User Assumptions About the TikTok Algorithm. In *Proc. of the International Web Science Conference (WebSci'21)*, pages 84–92, Southampton, England, UK, 2021.

[37] J. Li, L. Huang, T. Bai, Z. Wang, and H. Chen. CDBIA: A dynamic community detection method based on incremental analysis. In *Proc. of the International Conference on Systems and Informatics (ICSAI'12)*, pages 2224–2228, Yantai, China, 2012. IEEE.

[38] Y. Li, M. Guan, P. Hammond, and L.E. Berrey. Communicating COVID-19 information on TikTok: a content analysis of TikTok videos from official accounts featured in the COVID-19 information hub. *Health Education Research*, 2021. Oxford University Press.

[39] A. Lujain, H. Alhamarna, Y. AlWawi, Y. ElSayed, and H. Harb. Analysis of the representation of the 2019 Lebanese protests and the 2020 Beirut explosion on TikTok. *KIU Interdisciplinary Journal of Humanities and Social Sciences*, 1(3):53–72, 2020.

[40] P. De Meo, A. Nocera, G. Quattrone, and D. Ursino. A conceptual framework for community detection, characterization and membership in a Social Internetworking Scenario. *International Journal of Data Mining, Modelling and Management*, 6(1):22–48, 2014. Inderscience notes.

[41] K.Z. Meral. Social Media Short Video-Sharing TikTok Application and Ethics: Data Privacy and Addiction Issues. In *Multidisciplinary Approaches to Ethics in the Digital Era*, pages 147–165. IGI Global, 2021.

[42] A. Neyaz, A. Kumar, S. Krishnan, J. Placker, and Q. Liu. Security, privacy and steganographic analysis of FaceApp and TikTok. *International Journal of Computer Science and Security*, 14(2):38–59, 2020.

[43] L.H.X. Ng, J.Y.H. Tan, J.H. Darryl, and R.K.W. Lee. Will you dance to the challenge? predicting user participation of TikTok challenges. In *Proc. of the IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM'21)*, pages 356–360, Virtual event, The Netherlands, 2021.

[44] K. Papadamou, A. Papasavva, S. Zannettou, J. Blackburn, N. Kourtellis, I. Leontiadis, G. Stringhini, and M. Sirivianos. Disturbed Youtube for kids: Characterizing and detecting inappropriate videos targeting young children. In *Proceedings of the International Conference on Web and Social Media (ICWSM'20)*, volume 14, pages 522–533, Atlanta, GA, USA, 2020. Association for the Advancement of Artificial Intelligence.

[45] B. Qiu, K. Ivanova, J. Yen, and P. Liu. Behavior evolution and event-driven growth dynamics in social networks. In *Proc. of the IEEE International Conference on Social Computing (SocialCom'10)*, pages 217–224, Minneapolis, MN, USA, 2010. IEEE.

[46] Z. Qiyang and H. Jung. Learning and sharing creative skills with short videos: A case study of user behavior in tiktok and bilibili. In *Proc. of the International Association of Societies of Design Research Conference (IASDR'19)*, Manchester, UK, 2019.

[47] G. Rossetti and R. Cazabet. Community discovery in dynamic networks: a survey. *ACM Computing Surveys*, 51(2):1–37, 2018. ACM.

[48] G. Rossetti, L. Pappalardo, D. Pedreschi, and F. Giannotti. Tiles: an online algorithm for community discovery in dynamic social networks. *Machine Learning*, 106(8):1213–1241, 2017. Springer.

[49] A.R. Naghsh-Nilchi S. Nemati. Incorporating social media comments in affective video retrieval. *Journal of Information Science*, 42(4):524–538, 2016. SAGE Publications Sage UK: London, England.

[50] J.C. Medina Serrano, O. Papakyriakopoulos, and S. Hegelich. Dancing to the partisan beat: a first analysis of political communication on TikTok. In *Proc. of the International Web Science Conference (WebSci'20)*, pages 257–266, Southampton, England, UK, 2020.

[51] L. Shang, Z. Kou, Y. Zhang, and D. Wang. A Multimodal Misinformation Detector for COVID-19 Short Videos on TikTok, 2021. IEEE.

[52] A. Sheikhahmadi and M.A. Nematbakhsh. Identification of multi-spreader users in social networks for viral marketing. *Journal of Information Science*, 43(3):412–423, 2017. SAGE Publications Sage UK: London, England.

[53] E. Simpson and B. Semaan. For You, or For"You"? Everyday LGBTQ+ Encounters with TikTok. *Proc. of the International Conference on Human-Computer Interaction (HCI'21)*, 4(CSCW3):1–34, 2021. ACM.

[54] S. Singh, R. Kaushal, A.B. Buduru, and P. Kumaraguru. KidsGUARD: fine grained approach for child unsafe video representation and detection. In *Proc. of the International ACM/SIGAPP Symposium on Applied Computing (SAC'19)*, pages 2104–2111, Limassol, Cyprus, 2019.

[55] T. Sodani and S. Mendenhall. Binge-Swiping Through Politics: TikTok's Emerging Role in American Government. *Journal of Student Research*, 10(2), 2021.

[56] C. Stokel-Walker. TikTok's global surge. *New Scientist*, 245(3273):31, 2020. Elsevier.

[57] Y. Su, B.J. Baker, J.P. Doyle, and M. Yan. Fan engagement in 15 seconds: Athletes' relationship marketing during a pandemic via TikTok. *International Journal of Sport Communication*, 13(3):436–446, 2020.

[58] Y. Sun, J. Tang, L. Pan, and J. Li. Matrix based community evolution events detection in online social networks. In *Proc. of the IEEE International Conference on Smart City/SocialCom/SustainCom (SmartCity'15)*, pages 465–470, Chengdu, China, 2015. IEEE.

[59] E.G. Tajeuna, M. Bouguessa, and S. Wang. Tracking the evolution of community structures in time-evolving social networks. In *Proc. of the IEEE International Conference on Data Science and Advanced Analytics (DSAA'15)*, pages 1–10, Paris, France, 2015. IEEE.

[60] L. Tang and H. Liu. Community detection and mining in social media. *Synthesis lectures on data mining and knowledge discovery*, 2(1):1–137, 2010. Morgan & Claypool Publishers.

[61] C. Tantipathananandh and T. Berger-Wolf. Finding communities in dynamic social networks. In *Proc. of the International Conference on Data Mining (ICDM'11)*, pages 1236–1241, Vancouver, British Columbia, Canada, 2011. IEEE.

[62] C. Tantipathananandh, T. Berger-Wolf, and D. Kempe. A framework for community identification in dynamic social networks. In *Proc. of the International Conference on Knowledge Discovery and Data Mining (KDD'07)*, pages 717–726, New York, NY, United States, 2007. ACM.

[63] M. Tsvetovat and A. Kouznetsov. *Social Network Analysis for Startups: Finding connections on the social web*. Sebastopol, CA, USA, 2011. O'Reilly Media, Inc.

[64] M. De Veirman, S. De Jans, E. Van den Abeele, and L. Hudders. Unravelling the power of social media influencers: a qualitative study on teenage influencers as commercial content creators on social media. In *The regulation of social media influencers*. 2020. Edward Elgar Publishing.

[65] B. Wei and L. Chenxi. Study on the Win-Win Strategy of Douyin and Its Users. In *Proc. of the International Conference on Information Systems and Computer Aided Education (ICISCAE'20)*, pages 183–186, Dalian, China, 2020. IEEE.

[66] G. Weimann and N. Masri. Research note: spreading hate on TikTok. *Studies in Conflict & Terrorism*, pages 1–14, 2020. Taylor & Francis.

[67] L. Xu, X. Yan, and Z. Zhang. Research on the causes of the "Tik Tok" app becoming popular and the existing problems. *Journal of Advanced Management Science*, 7(2), 2019.

[68] J. Yang, J. Zhang, and Y. Zhang. First Law of Motion: Influencer Video Advertising on TikTok. *Available at SSRN 3815124*, 2021.

[69] K. Yousaf and T. Nawaz. A deep learning-based approach for inappropriate content detection and classification of youtube videos. *IEEE Access*, 10:16283–16298, 2022. IEEE.

[70] Z. Zhang and K. Wang. A trust model for multimedia social networks. *Social Network Analysis and Mining*, 3(4):969–979, 2013. Springer.

[71] Z. Zhao. Analysis on the "Douyin (Tiktok) Mania" Phenomenon Based on Recommendation Algorithms. In *Proc. of the International Conference on New Energy Technology and Industrial Development (NETID'20)*, volume 235, page 03029, Dali, China, 2021. EDP Sciences.

[72] C. Zhu, X. Xu, W. Zhang, J. Chen, and R. Evans. How health communication via Tik Tok makes a difference: a content analysis of Tik Tok accounts run by Chinese Provincial Health Committees. *International Journal of Environmental Research and Public Health*, 17(1):192, 2020. Multidisciplinary Digital Publishing Institute.

[73] J. Zhu, J. Liu, X. Zhang, and Y. Zhao. A reconstructed event-based framework for analyzing community evolution. In *Proc. of the IEEE International Conference on Big Data Analysis (ICBDA'16)*, pages 1–4, Hang Zhou, China, 2016. IEEE.

[74] D. Zulli and D.J. Zulli. Extending the Internet meme: Conceptualizing technological mimesis and imitation publics on the TikTok platform. *New Media & Society*, page 1461444820983603, 2020. SAGE.