Proceedings of the Seventeenth International AAAI Conference on Web and Social Media (ICWSM 2023)

Who Is behind a Trend? Temporal Analysis of Interactions among Trend Participants on Twitter

John Ziegler, Michael Gertz
Heidelberg University
Institute of Computer Science
Im Neuenheimer Feld 205
69120 Heidelberg - Germany
ziegler@informatik.uni-heidelberg.de, gertz@informatik.uni-heidelberg.de

Abstract

Trends are a fundamental component of today's fast-evolving media landscape. Still, a lot of questions about who participates in such trends remain unanswered. Are trends driven by individual actors, or do interactions between actors reveal community structures? If so, do those structures change during the life cycle of a trend or between topically similar trends? In short: Who is behind a trend?

This paper contributes to a better understanding of these questions and, in general, actor networks underlying trends on social media. As a case study, we leverage a large Twitter dataset from the EURO2020 soccer competition to detect and analyze topical trends. Our novel Gaussian fitting method allows separating trend life cycles into up- and down-trend components, as well as determining the duration of trends. An event-based evaluation proves good performance results. Given separate trend stages and topically similar trends at different points in time, we conduct a temporal analysis of the actor networks during trends. Our findings not only reveal a large overlap of participants between successive trends but also indicate large variations within a trend life cycle. Furthermore, actor networks seem to be centred around a small number of dominant users and communities. Those users also show large stability across similar trends over time. In contrast, temporally stable community structures are neither found within nor across topically similar trends.

Introduction

Trends, typically described as topics that "capture the attention of a large audience for a short time" (Asur et al. 2011, p. 1), characterize the dynamics of media attention. Especially in the context of social media, they represent a central component underlying the spreading of media content and interactions among users in the form of mentions, replies, and comments. Despite their importance, trends are still not well understood in all their details, and several aspects remain unexplored. This work aims at a better understanding of the actor networks behind trends on social media, expressed by the general question of "Who is behind a trend?". In more detail, we ask the following research questions:

1. *Are trends driven by individual actors, or do interactions between those actors reveal community structures?*
2. *Do those structures change during the life cycle of a trend or between topically similar trends?*

To answer these questions using appropriate computational methods, we conduct a case study based on a large Twitter dataset of more than 16 million tweets collected over the duration of the European soccer championship 2020 (EURO2020). A Gaussian fitting method is developed to identify time spans in which selected topics become trends. Stages of a trend, meaning up- and down-trend, as well as the respective duration of a complete trend life cycle, are also determined by the detection process. To better understand the relationship among participants of a trend, we model the dataset as temporal snapshot networks with interactions representing Twitter @mentions. Changes of these networks during a trend (intra-trend) and between similar trends at different points in time (inter-trend) are analyzed. In summary, this paper makes the following contributions:

1. To conduct a comprehensive analysis, we leverage a large-scale Twitter dataset that is known to contain topically similar trends over time. Known trends are leveraged for evaluation purposes.
2. By combining change point detection and Gaussian fitting, we present a novel method for trend detection. It allows to distinguish between up- and down-trend, and it helps to determine the duration of detected trends.
3. We use identified trend durations as adaptive time windows and model the dataset as temporal snapshot-based networks accordingly. In contrast to a static window size used in similar approaches, this adaptivity overcomes the difficulty of properly aggregating snapshots, a problem related approaches are struggling with.
4. An in-depth network analysis of the interactions among actors that participate in trends is conducted. Changes within trends as well as between topically similar trends over time are taken into account.

Our findings confirm but also significantly extend previous results obtained by related work such as (Budak, Agrawal, and El Abbadi 2011), (Asur et al. 2011) and (Zhang, Zhao, and Xu 2016) as they show new methods and results of an actor-centred analysis during and across temporal trends on social media. Regarding the first research question, they show that only a few Twitter accounts take up a large portion of the mentions in tweets, as observed by high domination-ratios, and therefore form the centre of the respective trend. Additionally, only a few large communities of actors are present in respective interaction networks. Most users belong to a small group of actors. Leading to the second research question, trends are not centred around stable communities of actors, as shown by an intra-trend analysis that reveals great variability among community members across the life cycle of a trend. These communities are also not temporally stable, as shown in an inter-trend comparison. In contrast, mentioned highly dominating actors are temporally stable, i.e., re-occurring across trends. Also, trends show a large overlap of participating users for topically similar trends at different points in time. Nevertheless, actors participating in a trend are strongly changing during the trend life cycle and vary between the up- and down-trend stages.

In the following, after a discussion of related work, we present our approaches to trend detection and the derivation of temporal properties of those trends. We then present various analytical methods that are aimed at a better understanding of the interactions among trend participants. We conclude the paper with a summary and discussion of ongoing work as well as an ethical statement addressing aspects related to the analysis of Twitter and trend data.

Related Work

The methods developed in this paper lie at the intersection of trend detection and network analysis, with some focus on the study of communities. While those topics have already been studied in numerous ways individually, work that connects both aspects is rare. On the one hand, regarding trend analysis, we refer the interested reader to the survey by Sharma et al. (2016). Also, in a more general sense, Yang and Leskovec (2011) investigate temporal attention patterns of online media content. On the other hand, the books by Newman et al. (2011) and Latora et al. (2017) give an excellent overview of the field of network science in general. More specifically, Javed et al. (2018) survey different community detection approaches. Similar to that, the work by Rossetti and Cazabet (2018) provides an in-depth coverage of the field of community detection as applied to temporal networks.

Probably most similar to ours is the work by Budak et al. (2011). By connecting topical trends with the social network structure of participating users, they are able to identify structurally different types of trends. According to their study, "coordinated" trends, as opposed to "uncoordinated" ones, are mainly discussed among users that are also "friends" in the respective social network. In contrast, "uncoordinated" trends are not driven by clustered groups of actors but instead are driven by unrelated users. As validation, they also use a Twitter dataset of trends connected with information about the Twitter social graph. Although our work connects trend analysis with social interactions among participating actors, too, we specifically focus on ad-hoc interactions that are present during trends and further compare those interactions across similar trends at different points in time. With our trend detection method, we are also able to differentiate between different phases of a trend and to compare interactions during these accordingly.

Work related to socio-semantic networks also connects topical or, more broadly speaking, semantic networks with actor networks, which is similar to our approach (Arroyo-Machado, Torres-Salinas, and Robinson-Garcia 2021) (Radicioni et al. 2021) (Hellsten and Leydesdorff 2020). Nevertheless, to the best of our knowledge, no work in this field focuses on actor networks that underlie topical trends on social media.

Furthermore, our work is related to the one by Asur et al. (2011). On the basis of a Twitter dataset, they analyze the factors that influence the formation and persistence of trends. Even though they already consider factors influencing the impact of users regarding a trend, they do not investigate the network structures between those trend participants. In a similar direction, the work by Zhang et al. (2016) is concerned with the question of whether the so-called "crowd", meaning a large number of low impact users or rather "opinion leaders" with a high influence, contribute the most to trends on social media. They already highlight the importance of ordinary users as opposed to influencers and therefore underline the necessity of our work. Again their focus is not on the interactions among actors participating in a trend. They do not further analyze community structures of such "crowds", and their work lacks a temporal comparison of similar trends re-occurring over time.

Recent work by Khan et al. (2021) is solely focusing on the actual detection and ranking of trends based on Twitter data. Different from our work, they take an open domain approach and therefore detect trends related to different genres but do not check for topically similar trends over time. Also, their method is not capable of determining different phases of a trend and its duration. Marangoni-Simonsen and Xie (2015) take a different approach and use a change point methodology to detect community emergence in a sequence of networks. Finally, Huang et al. (2020) apply change point detection to a sequence of network snapshots to find temporal anomalies. In our work, we start from change point detection to find the appropriate window size for network aggregations and subsequent analyses. Somehow similar is the work by Anghinoni et al. (2019), which proposes a novel trend detection method. Trends are represented as communities of complex networks that are extracted from time series data. In contrast to our work, their work is more theoretical and does not deal with the use case of analyzing actor networks underlying trends.

Trend Detection

The first step of our analysis consists of detecting trends in a collection of social media posts. Due to the lack of a benchmark dataset that fits our use case, we rely on a Twitter dataset related to the EURO2020[1] soccer championship that we collected. We check whether our trend detection method delivers high-quality results by conducting an evaluation based on known real-world events. Detecting trends

[1]UEFA EURO 2020: https://www.uefa.com/uefaeuro-2020, accessed 31/03/23

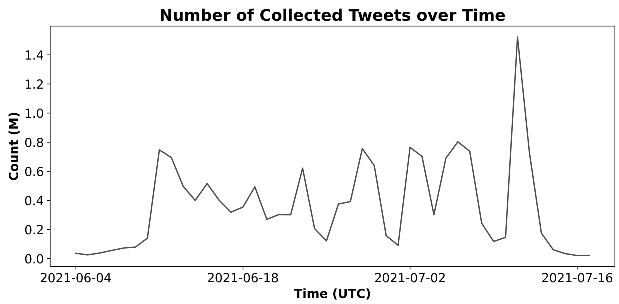

Figure 1: Number of tweets over time as part of the collected EURO2020 dataset

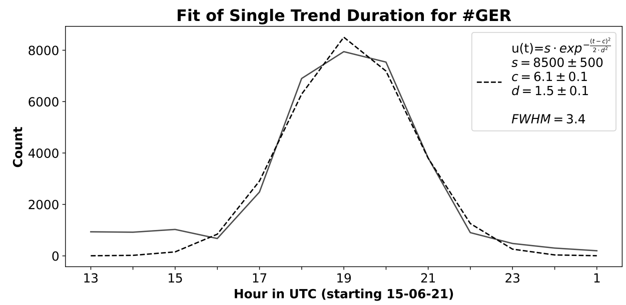

Figure 2: Example fit to identify the trend of the hashtag *GER* and its duration

and the time windows in which they are present provides the basis for subsequent analytical methods for studying interactions among actors contributing to those trends over time.

Dataset

The Twitter dataset we collected consists of the conversation related to the EURO2020 soccer competition. We rely on the Twitter search API v2[2] and gather tweets that either contain the official EURO2020 account (@*EURO2020*) or the official hashtag (*#EURO2020*). To get a complete dataset, we use a time window that starts one week before (*04/06/2021*) and ends one week after (*18/07/2021*) the competition. Time-stamped information about mentions of users in Tweets and what hashtags have been used is extracted from the raw tweets. Mentions as interactions among actors compared to retweets are used because they rather represent some kind of social tie as opposed to actions of just sharing information. The statistics of the dataset can be found in Table 1.

Description	**Count** $(in\ million)$
Tweets	16.163
Users	3.802
Hashtags	0.266
Hashtag Usages	27.594
User Mentions	19.707

Table 1: Rounded statistics of the collected Twitter dataset. For hashtag usages and user mentions, multiple occurrences in the same tweet are not considered.

As Figure 1 shows, the activity on the Twitter platform related to the EURO2020 championship varies over time. A clear peak of attention can be observed during times of the soccer championship final.

In line with the works of Asur et al. (2011) and Budak et al. (2011) we do not deal with the problem of topic extraction on its own but rather take hashtags as representatives of topics and therefore restricting ourselves to the detection of trends as determined by the temporal usage of hashtags.

For reproducibility, the source code of how the Twitter dataset was collected is made publicly available[3]. It also provides additional functionality regarding EURO2020 information, like access to match times of teams or retrieval of opponents.

Detection

Given the usage of a hashtag over time, the goal of the detection process is to find time windows in which the hashtag shows trending behaviour. More formally, let h be a hashtag and the output of the detection process $\tau_h = \{[t_1, t_2], [t_3, t_4], \dots, [t_n, t_{n+1}]\} \in \mathbb{N} \times T \times T$ the set of time windows defined as tuples of points in time that specify the duration of the found trends for that hashtag. We refer to the usage of a hashtag h over time as $u_h(t) : T \to \mathbb{N}$. In a time range $[t_{start}, t_{stop}]$ of a potential trend we model $u_h(t)$ as a Gaussian function with parameters s, c, and d:

$$\begin{aligned} u_h(t) &= s \cdot \exp\left(-\frac{(t-c)^2}{2d^2}\right) \\ \text{with} \quad & t_{start} \le t \le t_{stop}. \end{aligned} \tag{1}$$

Given that those parameters control the height, centre, and standard deviation of the function, one can think of them as the strength (s), centre (c) and duration (d) of the respective trend. In more detail, the full width at half maximum (FWHM) as characteristic value of the model is used as the actual time range $[t_1, t_2]$ defining the trend duration:

$$\begin{aligned} &\text{FWHM} = 2\sqrt{2 \ln 2}\, d, \\ &t_1 = c - \frac{\text{FWHM}}{2}, \quad t_2 = c + \frac{\text{FWHM}}{2}. \end{aligned} \tag{2}$$

Consequently, time ranges of the up- and down-trend stages of the respective trend are derived as follows:

$$\begin{aligned} &\text{Up-Trend} : [t_1, c] \in T \times T, \\ &\text{Down-Trend} : [c, t_2] \in T \times T. \end{aligned} \tag{3}$$

[2]Twitter Developer Platform: https://developer.twitter.com/en/docs/twitter-api/tweets/search/introduction, accessed 31/03/23

[3]UEFA EURO 2020 Twitter Dataset: https://github.com/jomazi/euro2020

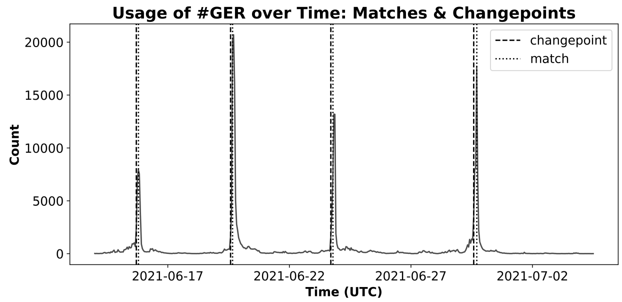

Figure 3: Detected change points for the hashtag *GER* and soccer matches of the German national team

Figure 2 shows the temporal hashtag usage for *#GER*, which stands for the German national team. It further shows how a Gaussian function is fitted to the time window that potentially contains the trend. Parameters derived by the model are then used to define the characteristics of the present trend. In this case, the trend duration is determined to be about 3.4 hours.

Although Gaussian-like popularity progressions seem to fit our use case (as they are in line with typical temporal attention patterns on social media, see clusters T1 and T2 in Yang and Leskovec (2011)), it is not yet answered how the time ranges that potentially contain a trend are determined. Constantly checking a sliding time window for a successful fit of the model is not efficient. This is why we leverage a two-step process and first detect change points of the hashtag usage via Bayesian Online Changepoint Detection (Adams and MacKay 2007). In the case a change point is detected, its point in time is taken as the centre of a $12h$ time window, and the trend detection is only then applied to this time range.

As an example, Figure 3 shows detected change points for the temporal hashtag usage of *#GER*. Interestingly, change points are in close proximity to the soccer matches of the German national team. Later on, this pattern is exploited during the evaluation, as described in the next section.

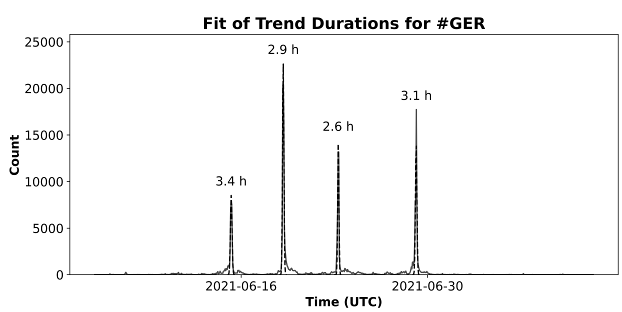

Figure 4: Trends and their durations for the hashtag *GER*

For the actual implementation of the change point detection step, we rely on the Kats[4] Python package as a toolkit for time series analysis. To detect shifts in the average hashtag usage, we use "NORMAL_KNOWN_MODEL" as the model parameter, together with the default threshold of 0.5 and a lag of 24. Given an hourly time resolution, the lag of $24h$ is inspired by a media logic of daily new trends.

Evaluation

To verify whether trends found in the present dataset do indeed reflect real-world trends, we follow an approach similar to the one used by Béres et al. (2018) and argue that real-world events characterized by their temporal limitation should also be reflected by a temporally limited shift of representative dataset statistics, in particular those of a trend.

As exemplified by Figure 3, detected change points and matches of soccer teams, as represented by certain hashtags, are close in time. This is why we resort to the official schedule[5] of the competition as ground truth and argue that during matches participating teams should be trending in the respective Twitter conversation. This follows the rationale adopted in Béres et al. (2018). To find tweets that belong to soccer teams, we rely on the official acronyms as used in the schedule. Those are often leveraged as hashtags on Twitter, e.g., *GER* referring to the German national team would often be used as *#GER*.

We check whether trends are detected on the same day as the according team played a EURO2020 soccer game. Given that for the group stage of the EURO2020 competition teams were arranged in groups of four, every team participated in at least three matches during the time window covered by the dataset. In total, 51 matches are evaluated. Per team, an average of 98% of the trends are correctly detected, and only 12% of the detected trends are false positives under the assumption that only soccer matches caused a trend. Further investigations reveal that false positive trends appear either the day before the tournament (53%) or on the last day of the group stage (47%). One can assume that the increased media coverage before the competition and the decision about which team makes it to the knockout stage, next to actual matches, also caused trends. Also, trend related discourse on Twitter is not only limited to the actual soccer matches but covers other events as well. As an example, Table 2 shows the most liked English tweets published on *12/06/21* mentioning *#DEN*. Even though the Danish team played against Finland and was trending this day, a large portion of the social media discourse is related to the cardiac arrest of the Danish soccer player Christian Eriksen[6]. These findings strengthen the assumption that the proposed dataset and methodology are not limited to a soccer-specific social media analysis but are applicable to study social media trends in general. Furthermore, the high accuracy gives reason to assume that our trend detection method is reliable

[4]Kats: https://facebookresearch.github.io/Kats, accessed 31/03/23

[5]EURO2021 Match Schedule: https://editorial.uefa.com/resources/026a-126a09addc81-6f092f1f9f89-1000/euro2021_match_schedule_-_english_-_310521_20210601103927.pdf, accessed 31/03/23

[6]CNN - Christian Eriksen suffered cardiac arrest during Euros match and 'was gone' before resuscitation, doctor says: https://edition.cnn.com/2021/06/13/football/christian-eriksen-stable-spt-intl/index.html, accessed 31/03/23

Tweet ID	Content
1403766744-624381955	YEEEES! The official word is here. #Erikson is alive! What a horrible scare. My thoughts are with his family, his friends, and all players on the pitch. #DENFIN #DEN #UEFA2020 #EURO2020 <emoji:folded hands> https://t.co/tuviGvtSVp
1403766315-681259524	Unbelievable news <emoji:clapping hands><emoji:clapping hands>. Just unbelievable <emoji:clapping hands><emoji:clapping hands>. #DEN #FIN #EURO2020 https://t.co/tozmCoft1X

Table 2: The discourse around trends on Twitter is not only related to soccer matches but also to other happenings, as in this case, the cardiac arrest by the Danish soccer player Christian Eriksen. Here, the most liked English tweets published on *12/06/21* mentioning *#DEN* are shown.

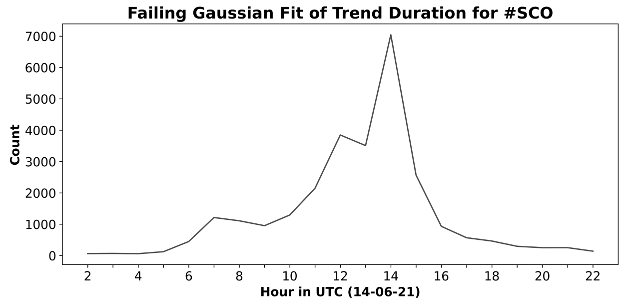

Figure 5: Example of a failing trend duration fit. In this case, the trend raises in a multi-step process as opposed to a single Gaussian-like increase.

and can be used during the next steps to analyze interactions among actors that participate in those trends.

Underlining the good performance of the approach, a median trend duration of 2.6 ± 0.5 hours is found. Given that a soccer match lasts about 1.5 hours plus a break of 0.25 hours (and often some extra time), this value seems to be reasonable. Especially one has to take into account that news coverage typically starts earlier and stops later than the actual time range of the match.

Speaking about the identification of trend durations, our proposed fitting procedure converges in all cases except for one (hashtag *SCO* (Scotland) at 14/06/21 6 am UTC), as shown in Figure 5, meaning that, in general, the Gaussian model can be seen as appropriate. It assumes a raise, peak, and downfall of the according trend, following a classical trend life cycle. Nevertheless, looking at the trend for which duration identification fails reveals another possible trend development. In this case, the trend raises in a two- or three-step process and can be seen as multiple overlapping Gaussian functions (think of a Gaussian mixture model). Compared to the findings of Yang and Leskovec (2011), the pattern might most accurately be described by their cluster T5, which also describes a two-step increase in attention. Given that this kind of up-trend pattern can be seen as rather rare, it is not studied further in this work.

Overall, the above method allows us to robustly identify trends and their durations. Applied to our dataset that also contains interactions among trend participants, it further allows us to analyze and temporally compare the social network structures underlying those trends.

Network Analysis

The network analysis part is aimed at a better understanding of the interactions among trend participants. As already discussed, we resort to mentions in tweets for this purpose. Formally, $G(V, E)$ denotes the directed mention network of users extracted from our Twitter dataset. An edge $e = (v_1, v_2, t) \in E$ contained in the network is defined as a triple of two vertices $v_1, v_2 \in V$ and a timestamp t at which the respective link occurred (i.e., account v_1 mentions account v_2 in a tweet at time t). Given that a trend of a hashtag h is present during the time window lasting from t_1 until t_2, the set of users $V_h^{t_1,t_2}$ and the set of mentions $E_h^{t_1,t_2}$ defines the interaction network $G_h^{t_1,t_2}$ among trend participants. Participants are those users that used the hashtag at least once during trend time. Formally, the set of interactions that has to be taken into account is defined as follows:

$$\begin{aligned} E \supseteq E_h^{t_1,t_2} = \{ & e = (v_1, v_2, t) \in E \mid \\ & (v_1 \in V_h^{t_1,t_2} \vee v_2 \in V_h^{t_1,t_2}) \wedge t_1 \leq t \leq t_2 \}. \end{aligned} \tag{4}$$

As a result, for a set of trends $\tau_h = \{[t_1, t_2], [t_3, t_4], \dots, [t_n, t_{n+1}]\}$ detected for a given hashtag h (see Detection), one obtains a series of snapshot networks $G_h = \{G_h^{t_1,t_2}, G_h^{t_3,t_4}, \dots, G_h^{t_n,t_{n+1}}\}$ that contain the interactions among participants during those trends. Note that according to the respective trend durations computed by our approach described above, snapshots are adaptively sized as opposed to having a fixed window size. Given the defined interaction networks, methods known from the field of network science, more specifically community detection (Javed et al. 2018), e.g., Infomap (Rosvall and Bergstrom 2008), can be applied to gain a better understanding of the interaction patterns.

Successive interaction networks can be compared based on their similarity. For this, we rely on the overlap coefficient as a measure of similarity (Vijaymeena and Kavitha 2016). Given two sets of vertices V_1 and V_2, the overlap coefficient α is defined as

$$\alpha(V_1, V_2) = \frac{|V_1 \cap V_2|}{min(|V_1|, |V_2|)}. \tag{5}$$

As an extension to the overlap coefficient α as defined in Equation 5, we specify the relative overlap α_r as the intersecting fraction relative to the cardinality of the first set:

$$\alpha_r(V_1, V_2) = \frac{|V_1 \cap V_2|}{|V_1|}. \tag{6}$$

In the following, we will talk about "inter-trend" comparison when it comes to comparing interaction networks for the same hashtag at different points in time. As an example, comparing the four trends of the hashtag *GER* shown in Figure 4 would follow the "inter-trend" approach. In contrast, "intra-trend" analysis refers to a comparison of the up- and down-trend phases within a single trend. It focuses on only a single trend period.

Regarding our implementation for this framework, we rely on the igraph software package (Csardi, Nepusz et al. 2006), and for community detection, we use the Infomap community detection algorithm (Rosvall and Bergstrom 2008), which is already included in the igraph package. Centrality measures are derived via PageRank (Page et al. 1999) scores, and community detection is applied per snapshot. We argue that approaches from the field of evolutionary clustering (Chakrabarti, Kumar, and Tomkins 2006) are not feasible for our community detection use case due to their inherent assumption that subsequent clusters should be similar, see the discussion about the method of "Temporal Smoothing" in Section 3.2 of Rossetti and Cazabet (2018). Compared to those methods, our approach is less restrictive, given that we do not assume to find communities or stable clusters of communities at all.

As an example, Figure 6 shows the interactions among the five largest communities during the trend on Twitter accompanying the match of Germany vs. France. Communities are labelled according to the three most central nodes within the respective cluster.

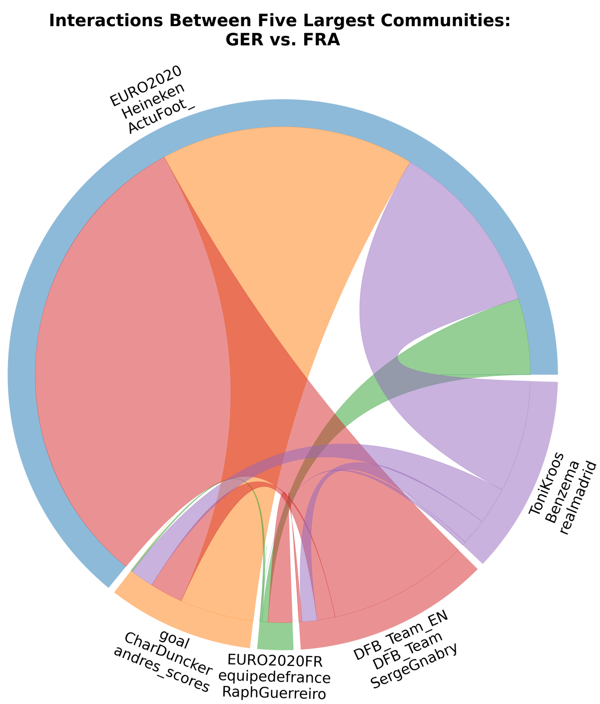

Figure 6: Exemplary mentions between the five largest communities in a trend visualized as chord diagram

Inter-Trend Results

Overall, the interaction networks derived from our dataset contain, on average, about $7 \pm 4k$ vertices and $20 \pm 10k$ edges, which corresponds to a very low density of 0.0004 ± 0.0002. This result on its own already shows that not a single highly connected group of nodes participates in trends, but instead interactions are spread across a large user base. Furthermore, the high deviations show that trends significantly vary in their spreading across the social media platform. If one takes mentions as an indicator of the social network between actors, the terminology as used in the study of Budak et al. (2011) can be used. In that sense, analyzed trends can be classified as "uncoordinated", meaning that they are discussed among distributed as opposed to clustered users.

Interestingly, a median of $21 \pm 4\%$ of nodes overlap for successive trend networks, $\alpha(V_h^{\tau_h[n]}, V_h^{\tau_h[n+1]}) = 0.21 \pm 0.04$. Taking into account that two comparable trends are caused by real-world matches against different teams, one would expect a maximum overlap of about 50%, assuming that the same amount of users represents supporters of the one and the other team. Given this assumption, the fraction of overlapping users seems to be quite high, and one can argue that a similar user base is participating across topically similar trends. From that, one could conclude that across matches, the fan base supporting a team stays similar. Alternatively, it might as well be that participants are not supporters of individual teams but instead general soccer-related actors, e.g., sports media outlets or journalists.

To determine whether single users play an outstanding role in the network, similar to Asur et al. (2011), we calculate the domination-ratio as the proportion of mentions that come from or go to the most mentioning/mentioned user. For an even more meaningful quantity, we also sum up the domination-ratios for the ten most active participants (see Table 3 "10"). Because mentions imply a direction, we differentiate between out- and ingoing mentions. As shown in Table 3, especially for ingoing mentions, the domination-ratio is quite high, which means that only a handful of user accounts take up a large portion of the overall mentions and thereby form the core of actors the trend is centred around. One might call those actors "trend-hubs" or "trend-influencer". Most interestingly, the top ten most dominating users reach a median overlap coefficient of 0.46 ± 0.04. This value is close to the expected maximum value of 0.5, as explained above.

Contrarily, there do not seem to be any accounts dominating the actual mentioning activity, as the domination-ratios

	Complete	Up-Trend	Down-Trend
I1	0.16 ± 0.03	0.17 ± 0.04	0.14 ± 0.03
I10	$\mathbf{0.46 \pm 0.04}$	$\mathbf{0.48 \pm 0.04}$	$\mathbf{0.48 \pm 0.06}$
O1	0.008 ± 0.003	0.02 ± 0.01	0.009 ± 0.004
O10	0.05 ± 0.02	0.09 ± 0.05	0.05 ± 0.02

Table 3: Domination-ratios for different trend phases, in- (I) and outgoing (O) mentions, as well as the top one (1) and ten (10) ranked users

GER vs. FRA	GER vs. POR
EURO2020	EURO2020
goal	DFB_Team
InvictosSomos	goal
Football__Tweet	DFB_Team_EN
DFB_Team_EN	ToniKroos
DFB_Team	Cristiano
Cristiano	realmadrid
EURO2020FR	InvictosSomos
realmadrid	2010MisterChip
ToniKroos	selecaoportugal

GER vs. HUN	GER vs. ENG
EURO2020	EURO2020
goal	goal
Football__Tweet	DFB_Team
Footballogue	England
SquawkaNews	DFB_Team_EN
DFB_Team	sterling7
DFB_Team_EN	BBCSport
InvictosSomos	ChelseaFC
brfootball	ManCity
2010MisterChip	Football__Tweet

Table 4: Top ten most dominating users for trends during matches of the German national team

for outgoing mentions are low. This finding is in line with the insights of Asur et al. (2011), who highlight the link between low domination-ratios and longer trend durations. Zhang et al. (2016) as well outline the importance of the "crowd" participating in a trend for it to gain large popularity.

Most dominant user accounts regarding ingoing mentions during matches of the German national team are shown in Table 4. Most actors are related to soccer players and teams, national soccer associations or the EURO2020 championship. Obviously, those users have a large relevance for soccer-related trends. Accounts that are closely linked to the German national team, in this case *DFB_Team_EN* and *DFB_Team*, can be found across all trends. The *EURO2020* account can be seen as an artefact of the process of how the dataset was collected. On the other hand, one can also observe large variations between the different matches/trends. Depending on the opposing team accounts that are related to this one also become relevant, e.g., *England* or *sterling7*, for the match against the English team.

Besides individual actors, when it comes to community detection, on average, 700 ± 300 communities per network are found. Of these communities, most consist of only a small number of users, as shown in Figure 7. The distribution of community sizes approximately follows a power law decay. The largest communities, on average, consist of $14 \pm 4\%$ of nodes of the complete network. Again, strongly deviating values are observed, which suggests large differences in the trend dominance of single communities.

Interestingly, a pairwise comparison of the top ten largest communities between trends reveals large fluctuations. On average, even the maximum overlap coefficient with a value

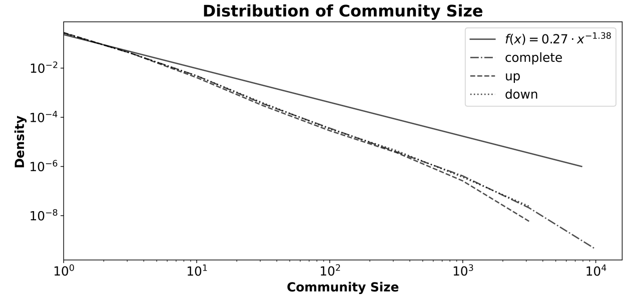

Figure 7: Community sizes differentiated by trend stages. Densities approximately follow power law distributions.

of 0.21 ± 0.04 does not surpass the similarity score between complete networks. Therefore, communities do not seem to be temporally stable across trends. This finding is underlined by a very low overlap between the largest communities of subsequent trends (0.04 ± 0.03). Also, only a small fraction of nodes of the largest community during a trend can be found during the next trend ($\alpha = 0.13 \pm 0.03$).

To verify whether certain communities of actors continue within other communities in succeeding trends, we check for the maximum fraction of users that stay together between any pair of the top ten communities in successive trends. For this, we use the relative overlap coefficient α_r as defined in Equation 6. Figure 8 gives a visual example of such a community flow in the form of an alluvial diagram. It should be noted that the sizes of the community blocks are not absolute but relative to the in- respective outgoing overlap. As community labels, the top three accounts with the highest centrality values are used. One can see that the community containing the *DFB_Team_EN* Twitter account partially continues across trends. Nevertheless, for the most part, communities get mixed quite strongly as a low maximum value of $\alpha_r = 0.10 \pm 0.03$ between successive top ten communities indicates.

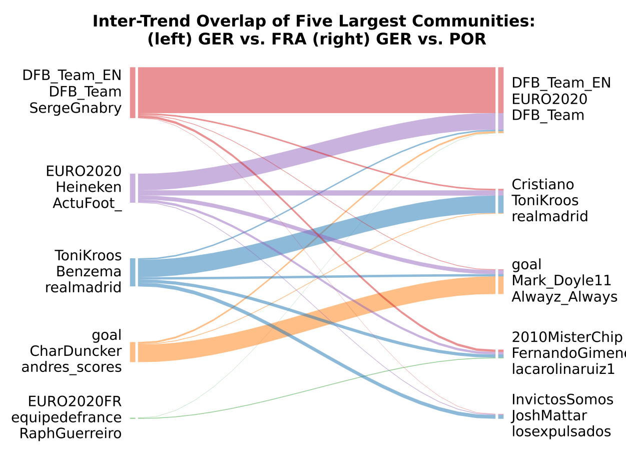

Figure 8: Exemplary inter-trend overlap of the five largest communities visualized as alluvial diagram

Intra-Trend Results

In contrast to the inter-trend setting, the intra-trend analysis compares interaction networks between different stages of a trend (up- vs. down-trend). Overall, the mention networks during those stages contain, on average, about $3 \pm 2k$ and $5 \pm 3k$ vertices resp., as well as $6 \pm 4k$ and $11 \pm 6k$ edges resp., meaning that the mention networks reveal very low densities of 0.0008 ± 0.0005 respectively 0.0005 ± 0.0003. Similar to the findings for the inter-trend analysis, a group of loosely interacting actors seems to participate across the stages of a trend. In this case, one could as well talk about "uncoordinated" trends (Budak, Agrawal, and El Abbadi 2011). Furthermore, it has to be noted that, on average, in the later phase of a trend, more distinct users participate. This can probably be explained by an already larger popularity of the trend at that point in time.

Normally, one would not expect to have completely different networks during up- and down-trend phases. However, this is what our analysis reveals. Within a single trend, only a median of $28 \pm 4\%$ of nodes are overlapping, showing that only about a quarter of users participate across the complete trend life cycle. Also, regarding community analysis, results are similar to the ones found for the inter-trend setting. On average, 400 ± 200 and 500 ± 200 communities per network are found respectively. Figure 7 shows the distribution of community sizes over all actor networks. As can be seen, communities mostly consist of only a small number of users, a property that holds for all trend stages. The distributions of community sizes approximately follow power laws with even less large communities as expected by named distribution. The largest communities cover on average $16 \pm 6\%$ respectively $14 \pm 5\%$ of actors of the complete ad-hoc network and are thereby quite dominant among respective interactions. On average, the maximum similarity in terms of overlap coefficient between the top ten largest communities during the up- and down-trend is only 0.28 ± 0.08. As the two stages are part of the same trend, one would expect to have a more stable base of actor communities that contribute to the trend across its life cycle. By simply comparing the largest communities between the two trend stages, similar results are obtained. On average, only a fraction of $6 \pm 6\%$ of actors overlap. Also, again a low maximum value of $\alpha_r = 0.10 \pm 0.04$ is observed by comparing the top ten communities of successive trend stages. Nevertheless, in all cases, actors of the largest community found during the up-trend also participate in the down-trend, meaning that at least those actors are stable across the entire trend life cycle.

Furthermore, findings related to the domination of users, as shown in Table 3 and already discussed for the inter-trend analysis, also hold true for the intra-trend setting. Domination-ratios of 0.48 regarding the top ten ranked users can be observed for both trend stages. By comparing a maximum of ten most dominant accounts between trend stages, an overlap of $60 \pm 10\%$ is observed on average. Therefore, dominant accounts seem to be somehow stable within a trend as well. As an example, the top five dominant users during matches of the German national team regarding ingoing mentions and separated by trend stages are shown in Table 5. One might expect that changes of these dominant accounts between the two trend stages are caused by different foci in media attention, e.g., pre-match expectations vs. moderating the soccer match. Unfortunately, analysis of the actor networks does not tell much about the topical variations within a trend life cycle. More suitable semantic analyses are needed for this, which also brings up future work.

Up Stage

GER vs. FRA	GER vs. POR
DFB_Team	EURO2020
Cristiano	InvictosSomos
EURO2020	DFB_Team_EN
goal	DFB_Team
DFB_Team_EN	goal

GER vs. HUN	GER vs. ENG
goal	EURO2020
DFB_Team_EN	DFB_Team_EN
EURO2020	goal
DFB_Team	DFB_Team
InvictosSomos	England

Down Stage

GER vs. FRA	GER vs. POR
EURO2020	EURO2020
goal	DFB_Team
Football__Tweet	DFB_Team_EN
DFB_Team_EN	ToniKroos
DFB_Team	Cristiano

GER vs. HUN	GER vs. ENG
EURO2020	EURO2020
goal	goal
DFB_Team	England
DFB_Team_EN	DFB_Team_EN
InvictosSomos	BBCSport

Table 5: Top five most dominating users for different trend stages during matches of the German national team

Conclusion and Future Work

Given the dominant role of social media platforms and their influence on society, obtaining insights into who contributes to trends and how underlying actor networks are structured will continue to play an important role. In this work, we have tackled the problem of combining two core methods for such studies, trend detection and network analysis, to get a more fine-grained view of social interactions among actors during trends on social media. In the context of a large-scale Twitter dataset related to the European soccer championship 2020, we developed several novel methods that aim at analyzing and exploring trends. Our novel Gaussian-based trend detection method allows us to differentiate between up- and down-trend as well as to determine the duration of a trend. An event-based evaluation proves good performance. Furthermore, with the dataset and detected trends on hand, we are able to analyze topically similar trends across time (inter-trend) and within the trend life cycle (intra-trend). Our analysis focuses, in particular, on the actors behind those trends, modelled in the form of temporal mention networks extracted from the dataset. Trend-dependent time windows allow for an adaptive snapshot-based network aggregation.

Among other results, our findings show a large overlap of the user base in an inter-trend setting but not within a trend during its different stages. Users participating in a trend seem to vary a lot during the life cycle of a trend but not so between similar trends across time. Furthermore, trends are centred around a small set of highly influential users, as indicated by high domination-ratios. This core of actors is also stable across time. In contrast, even though large communities of actors are present, these are neither stable within nor across trends.

While these insights are primarily based on the present Twitter dataset, similar results are likely obtained for other types of social media datasets as analyzed content is not only centred around soccer matches but covers the discourse around general events as well. Also, the trend detection method captures common media attention patterns, meaning that it applies to other trend analysis scenarios as well.

The present work could be extended in several ways. For example, the trend detection method could be enhanced by techniques that can deal with different trend progressions, such as the one shown in Figure 5. Also, incorporating additional information like terms and named entities into the proposed network model might complement the analysis with a better semantic understanding of given trends. This way, topical shifts within and across trends might be recognized as well. Furthermore, it might be interesting to integrate additional data sources to gather a more complete picture of who is participating in a trend on which platform. Are there cross-platform patterns that emerge synchronously, e.g., in a coordinated fashion by a small group of users? Finally, from a methodological point of view, evolutionary clustering might lead to different results when it comes to detecting temporally stable communities in the actor networks behind analyzed trends.

In general, methods and techniques described in this paper provide a solid basis for studying actor networks underlying trends on social media.

Ethics Statement

Especially relevant to this work are ethical considerations in the context of the collected dataset, as well as a potential impact of results gained from studying actor networks behind trends. After all, mainly real users are behind trends and social media postings in general. Regarding the Twitter trend dataset as used in the present work, we picked a ”harmless” topic to collect tweets about, given in the form of the EURO2020 soccer championship. Note that in our results, we do not present any sensitive information and rely on aggregated statistics whenever possible. Our work aims at a better understanding of dynamic actor networks behind trends, and in the long run, with even more techniques and methods at hand, those insights could potentially be abused in a way that topics damaging to our societies are reinforced by malicious actors or that under-represented topics are obstructed during their emergence. We do not give any hints on how gathered insights could be abused but stick to simply conducting observations. Driving for a good change, we argue that those insights might as well be or even should be used to detect undesired activities and to prevent those during their early stages.

Acknowledgements

We thank the Klaus Tschira Foundation for funding this research in the framework of the EPINetz project. For more information, please visit https://epinetz.de.

References

Adams, R. P.; and MacKay, D. J. 2007. Bayesian online changepoint detection. *arXiv preprint arXiv:0710.3742*.

Anghinoni, L.; Zhao, L.; Ji, D.; and Pan, H. 2019. Time series trend detection and forecasting using complex network topology analysis. *Neural Networks*, 117: 295–306.

Arroyo-Machado, W.; Torres-Salinas, D.; and Robinson-Garcia, N. 2021. Identifying and characterizing social media communities: a socio-semantic network approach to altmetrics. *Scientometrics*, 126(11): 9267–9289.

Asur, S.; Huberman, B. A.; Szabo, G.; and Wang, C. 2011. Trends in social media: Persistence and decay. In *Fifth international AAAI conference on weblogs and social media*.

Béres, F.; Pálovics, R.; Oláh, A.; and Benczúr, A. A. 2018. Temporal walk based centrality metric for graph streams. *Applied network science*, 3(1): 1–26.

Budak, C.; Agrawal, D.; and El Abbadi, A. 2011. Structural trend analysis for online social networks. *Proceedings of the VLDB Endowment*, 4(10): 646–656.

Chakrabarti, D.; Kumar, R.; and Tomkins, A. 2006. Evolutionary clustering. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, 554–560.

Csardi, G.; Nepusz, T.; et al. 2006. The igraph software package for complex network research. *InterJournal, complex systems*, 1695(5): 1–9.

Hellsten, I.; and Leydesdorff, L. 2020. Automated analysis of actor–topic networks on twitter: New approaches to the analysis of socio-semantic networks. *Journal of the Association for Information Science and Technology*, 71(1): 3–15.

Huang, S.; Hitti, Y.; Rabusseau, G.; and Rabbany, R. 2020. Laplacian Change Point Detection for Dynamic Graphs. *CoRR*, abs/2007.01229.

Javed, M. A.; Younis, M. S.; Latif, S.; Qadir, J.; and Baig, A. 2018. Community detection in networks: A multidisciplinary review. *Journal of Network and Computer Applications*, 108: 87–111.

Khan, H. U.; Nasir, S.; Nasim, K.; Shabbir, D.; and Mahmood, A. 2021. Twitter trends: a ranking algorithm analysis on real time data. *Expert Systems with Applications*, 164: 113990.

Latora, V.; Nicosia, V.; and Russo, G. 2017. *Complex networks: principles, methods and applications*. Cambridge University Press.

Marangoni-Simonsen, D.; and Xie, Y. 2015. Sequential Changepoint Approach for Online Community Detection. *IEEE Signal Processing Letters*, 22(8): 1035—-1039.

Newman, M.; Barabási, A.-L.; and Watts, D. J. 2011. *The Structure and Dynamics of Networks*. Princeton University Press.

Page, L.; Brin, S.; Motwani, R.; and Winograd, T. 1999. The PageRank citation ranking: Bringing order to the web. Technical report, Stanford InfoLab.

Radicioni, T.; Squartini, T.; Pavan, E.; and Saracco, F. 2021. Networked partisanship and framing: a socio-semantic network analysis of the Italian debate on migration. *PloS one*, 16(8): e0256705.

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

Rosvall, M.; and Bergstrom, C. T. 2008. Maps of random walks on complex networks reveal community structure. *Proceedings of the national academy of sciences*, 105(4): 1118–1123.

Sharma, S.; Swayne, D. A.; and Obimbo, C. 2016. Trend analysis and change point techniques: a survey. *Energy, Ecology and Environment*, 1(3): 123–130.

Vijaymeena, M.; and Kavitha, K. 2016. A survey on similarity measures in text mining. *Machine Learning and Applications: An International Journal*, 3(2): 19–28.

Yang, J.; and Leskovec, J. 2011. Patterns of temporal variation in online media. In *Proceedings of the fourth ACM international conference on Web search and data mining*, 177–186.

Zhang, L.; Zhao, J.; and Xu, K. 2016. Who creates trends in online social media: The crowd or opinion leaders? *Journal of Computer-Mediated Communication*, 21(1): 1–16.