

RESEARCH

Open Access



# Content-based comparison of communities in social networks: Ex-Yugoslavian reactions to the Russian invasion of Ukraine

Bojan Evkoski<sup>1\*</sup>, Petra Kralj Novak<sup>2,3</sup> and Nikola Ljubešić<sup>2,4</sup>

\*Correspondence:  
bojanevkoski@outlook.com

<sup>1</sup> Institute of Contemporary History, Ljubljana, Slovenia

<sup>2</sup> Department of Knowledge Technologies, Jozef Stefan Institute, Ljubljana, Slovenia

<sup>3</sup> Department of Network and Data Science, Central European University, Vienna, Austria

<sup>4</sup> Faculty of Information and Communication Sciences, University of Ljubljana, Ljubljana, Slovenia

## Abstract

We discuss the added value of various approaches for identifying similarities in social network communities based on the content they produce. We show the limitations of observing communities using topology-only and illustrate the benefits and complementarity of including supplementary data when analyzing social networks. As a case study, we analyze the reactions of the Ex-Yugoslavian retweet communities to the Russian invasion of Ukraine, comparing topological inter-community interaction with their content-based similarity (hashtags, news sources, topics and sentiment). The findings indicate that despite the Ex-Yugoslavian countries having a common macro-language, their retweet communities exhibit diverse responses to the invasion. Certain communities exhibit a notable level of content-based similarity, although their topological similarity remains relatively low. On the other hand, there are communities that display high similarity in specific types of content, but demonstrate less similarity when considering other aspects. For example, we identify a strong echo-chamber community linked to the Serbian government that deliberately avoids the invasion topic, despite showing news source similarities with other communities highly active on the subject. In summary, our study highlights the importance of employing multifaceted approaches to analyzing community similarities, as they enable a more comprehensive understanding of social media discourse. This approach extends beyond the confines of our specific case study, presenting opportunities to gain valuable insights into complex social events across various contexts.

**Keywords:** Social network analysis, Community detection, Community comparison, Topology versus content, Large language models, Russian invasion of Ukraine, Ex-Yugoslavia

## Introduction

Community detection algorithms are powerful network analysis techniques used to identify groups of highly interconnected nodes, which can provide a large-scale map of a network and help to shed light on its underlying function. These techniques are widely used in many fields, including physics, sociology, biology, computer science, economics, politics, and neuroscience (Newman 2003; Strogatz 2001; Amaral and Ottino 2004). While the resulting communities are topological (Ding 2011) (considering interactions

between nodes), their interpretation is typically qualitative in nature, relying on additional contextual data that is not or cannot be represented in the network's topology (Morini et al. 2021; Surian et al. 2016; Cinelli et al. 2021; Grčar et al. 2017).

In social network analysis, there is a frequent need to measure the similarity of social groups, or communities. This allows one to better understand the dynamics and interactions within a network. By comparing communities, we can identify and differentiate patterns of behavior, opinions, and interests among distinct communities. This information is useful in various fields where the concept of homophily (the tendency of individuals to associate with similar others) plays a role (Bisgin et al. 2010), such as marketing (Spertus et al. 2005), politics and sociology (Evdoski et al. 2021; Durazzi et al. 2021), as it provides insights into how groups form and interact. It can also help in identifying potential sources of conflict, or areas where collaboration and cross-community engagement may be beneficial. Moreover, it can be used in studying community evolution, where similar communities have the potential to merge, or diverging communities the potential to split (Rossetti and Cazabet 2018; Cordeiro et al. 2016).

Typical topological approaches for community similarity include measuring the interaction between two communities, the node overlaps in repeated community partitions, or through the utilization of hierarchies in hierarchical community detection (Evdoski et al. 2021; Weng et al. 2014; Malin 2005). Yet, these solutions take into account only the interactions on which the network topology is based, disregarding potential similarities of the communities based on the information they produce and interact with.

To achieve a comprehensive understanding of social communities, it is crucial to consider data beyond their topology. Even when two communities are not directly connected, they might share similarities in terms of topics, opinions, or influential figures. Moreover, they can approach the same matter differently, with varying intensity, or using different keywords or hashtags (Novak et al. 2018). Therefore, it is imperative to incorporate multifaceted analysis to properly measure the similarity between communities.

Recent advances in machine learning and natural language processing allow us to perform content-based analysis with greater accuracy and efficiency than ever before. With the advent of sophisticated language representation methods such as large language models (LLMs), our ability to compute and compare content has reached new heights (Lin et al. 2022). We are now able to automatically infer the contextual meaning of words, sentences, and even documents. And these technological advances can now be readily applied to research beyond computer science, and are particularly useful in data-driven interdisciplinary areas such as complex network analysis.

In this paper, we combine the common practices of community detection in social networks with content-based analysis via state-of-the-art natural language processing, and present the following contributions:

- We discuss approaches for quantifying the similarity of social network communities based on the textual data they produce.
- We present an analysis of the limitations of measuring topology-only community similarity, as well as the benefits and complementarity of the content-based approach on three representations of social media content: hashtags, news sources, and topics with sentiment.

- We apply the discussed approaches to analyze the reaction of the Ex-Yugoslavian tweetosphere (later referred to as the BCMS tweetosphere) to the Russian invasion of Ukraine, in two scenarios:
  - Comparison of communities in one network snapshot—revealing content-based similarities not visible through the topology of the retweet network in the first four weeks of the invasion.
  - Comparison of communities between time snapshots of a network—identifying content drifts between subsequent network snapshots in community evolution, by comparing communities of the retweet network before the invasion with communities during the invasion.

### Application

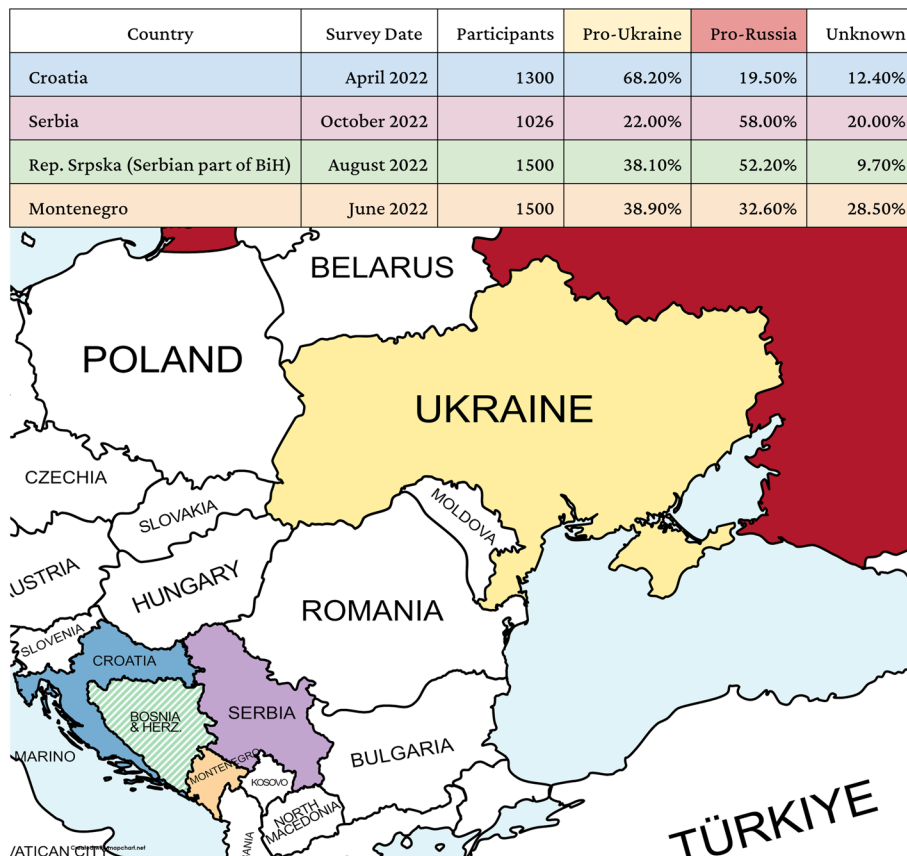
We apply content-based comparisons in the analysis of the Ex-Yugoslavian tweetosphere (Mazzucchelli 2012) and its response to the Russian invasion of Ukraine. We collect tweets posted during the first four weeks of the invasion and the four weeks prior to it, from users who predominantly tweet in one of the four official varieties of the pluricentric Serbo-Croatian macro-language (referred to as BCMS<sup>1</sup>): Bosnian (bos), Croatian (hrv), Montenegrin (cnr), or Serbian (srp). All four languages act as official languages in the respective Ex-Yugoslavian countries.

The BCMS region, despite sharing the same macro-language, is highly heterogeneous due to deep-seated cultural differences. Nationality, ethnicity, religion, past conflicts, and political affiliations are some of the factors that contribute to the region's polarization on various topics (Banac 2009; Papic 1994). Moreover, polls in 2022 indicate a divide in opinion on the Russian invasion of Ukraine, with Croats largely supporting Ukraine, Serbs and Bosnian Serbs backing Russia, and Montenegrins split between the two (see Fig. 1). Hence, we consider the BCMS tweetosphere to be a relevant case study, as it enables us to investigate complex similarities or dissimilarities that are hardly captured in a typical network topology.

The article is organized as follows: In [Results and discussion](#) we consider approaches for comparing social network communities based on the textual content they produce and compare the results with a topological similarity approach on data from the BCMS tweetosphere. We also critically evaluate the advantages and limitations of the proposed content-based similarity analyses and elaborate on their applicability beyond the examined use case. In [Methodology](#) we describe the building blocks for measuring content-based similarity of communities on several types of content: hashtags and news sources via ranked lists; and topics and sentiment via transformer-based language models. Finally, we wrap up with [Conclusions](#).

---

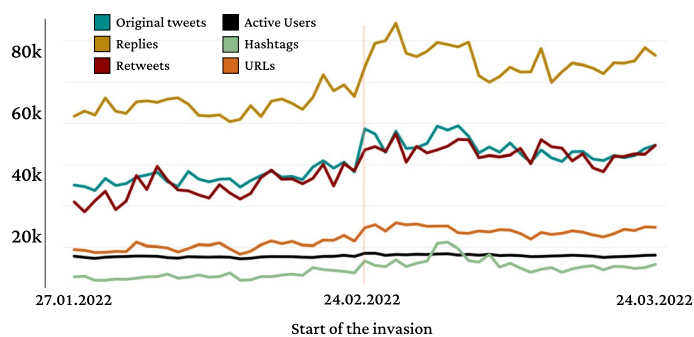
<sup>1</sup> macro-language code hbs by the ISO-639-3 standard.



**Fig. 1** Top part of the figure shows poll results in Croatia (<https://tinyurl.com/croatian-poll-riu>), Serbia (<https://tinyurl.com/serbian-poll-riu>), Republika Srpska (part of Bosnia and Herzegovina governed by Bosnian Serbs) (<https://tinyurl.com/republika-srpska-poll-riu>), and Montenegro (<https://tinyurl.com/montenegro-poll-riu>) on the following question: “Who do you support in the Ruso-Ukrainian conflict?”. The bottom part shows a map of Eastern Europe, with color marks on the BCMS-speaking countries covered in this research, Ukraine (yellow) and Russia (red)

### Results and discussion

In this section, we present results on several approaches for quantifying the similarity of social network communities based on the textual data they produce. In subsection [Data and experimental setup](#) we cover the data collection, the creation of the BCMS tweetosphere networks and the detected communities on which we are quantifying similarity. In subsection [Content-based similarity](#) we present and compare the results of the topological and the content-based similarity of the communities. In subsection [Similarity in community evolution](#) we apply the approach in a community evolution scenario, demonstrating how content-based similarity can be used in tracking communities through temporal snapshots. In [Discussion](#) we overview the benefits and complementarity of the content-based approach over the topological similarity, as well as discuss limitations and prospects.



**Fig. 2** Daily volume of original tweets, replies, retweets, active users, as well as hashtags and URLs posted. The x-axis shows the date, and the y-axis shows the volume

### Data and experimental setup

Our dataset consists of tweets spanning the last four weeks before (from January 27 to February 23, 2022), and the first four weeks of the Russian invasion of Ukraine (from February 24 to March 24, 2022). We refer to the first part of the tweets as the pre-invasion snapshot, and to the second as the invasion snapshot. We use the invasion snapshot to quantify content-based community similarity inside the network and compare the approach to the topological similarity. We use both snapshots to measure the content-based similarities between the pre-invasion and the invasion communities, getting insight into which communities evolved to take a Pro-Ukraine or Pro-Russia stance.

We detect and collect most of the tweets in the BCMS tweetosphere for the period of interest, using a collection process that is continuously running since 2017. We use the TweetCat tool (Ljubešić et al. 2014), which utilizes the Twitter search API to continuously identify new users speaking in a specific language (in our case the BCMS macro-language) by searching for tweets containing frequent words specific to that language and also by analyzing followers and followees of the already-identified users.

The final dataset contains 8.86M tweets, out of which 2.5M are original tweets, 4M are replies and 2.36M are retweets. Only 3.8M tweets belong to the pre-invasion, and 5M to the invasion snapshot, showing a massive activity increase as a reaction to the invasion. Figure 2 shows the daily volume of tweets, peaking on February 27 (the fourth day of the invasion) with 195K tweets. The total number of users in the dataset (not including retweeted users not in the original collection) is 24K, with around 13.5K active users per day. The peak of active users is 19k, marking the second day of the invasion.

### Retweet networks

Various forms of interactions on Twitter are possible between users, such as follows, mentions, replies, and retweets. The most useful indicator of social ties between Twitter users is retweeting. When a user retweets a post, it is distributed to all of its followers, just as if it were an originally authored post. Users retweet content that they find interesting or agreeable. Despite the fact that it does not always signify an endorsement, most retweets indicate links between like-minded users. Specifically, retweeting political tweets has proven to be a highly effective reflection of actual political alignments and influence (Metaxas et al. 2015; Wong et al. 2013; Cherepnalkoski and Mozetič 2016).

Practically, a retweet network is a directed graph  $G$ . The nodes are Twitter users and the edges are retweet links between the users. The edge weight is the number of times user  $A$  retweeted user  $B$ .

We build a pre-invasion and an invasion retweet network. The pre-invasion retweet network consists of 59K users, 497K unique connections, and 1 M retweets. The invasion retweet network consists of 76K users, 700K unique connections, and 1.36M retweets. The drastic increase in users is mainly due to the introduction of accounts in the network (largely international) that discuss the invasion.

### **Communities**

Informally, a community in a network is a subset of nodes more densely linked between themselves than with nodes outside the community. Community detection attracted significant attention in social network analysis, resulting in the development of numerous methods to uncover and analyze communities. These methods encompass a wide range of approaches, including modularity optimization, spectral clustering, and stochastic block models, among others (Newman 2004; Fortunato 2010). One of the most used algorithms for community detection in social networks is the modularity maximization Louvain algorithm (Blondel et al. 2008). It is computationally efficient, well-suited for large networks, and does not require ex-ante assumptions about the number or size of the communities (Lancichinetti and Fortunato 2009). However, due to greedy modularity maximization in its core (Fortunato and Hric 2016), the Louvain algorithm yields inconsistent partitions run with a different random initialization on the same network. This instability of community detection results can consequently negatively impact the consistency and the reproducibility of the downstream community similarity results.

To mitigate the instability of Louvain that could potentially negatively impact the further topological and content-based comparison of the communities, we perform robust community detection with ensembles, which is a technique of combining multiple resulting partitions into one aggregated partition (Lancichinetti and Fortunato 2012; Evkoski 2022). For our retweet networks, we apply the ECG (Ensemble Clustering for Graphs) (Poulin and Théberge 2019) implementation, combining 100 Louvain runs on the weighted retweet network into one.

Overall statistics of the communities for the pre-invasion and invasion networks are shown in Tables 1 and 2. For similarity analysis, we select communities with at least 1000 users and put the rest of the nodes in a group named `Rest`. The pre-invasion network holds 12 large communities, and the invasion network 11. We name the communities through a manual observation of the following categories: most retweeted users, most used hashtags, and most discussed topics. All communities have shown evident topical, ideological, or demographical distinctions, thus the naming procedure is unambiguous. To maintain consistency in our analysis, we utilize identical naming for communities that show continuity from the pre-invasion to the invasion period. For instance, the community named `Vučić` appears in both the pre-invasion and invasion community lists as the same set of most influential users reappears.

The statistics in Tables 1 and 2 show that the BCMS tweetsphere is dominated by political issues, with the majority of communities having politicians and news sources as the most influential users by PageRank (Gleich 2015). As the majority of the content

**Table 1** General statistics for the pre-invasion communities with more than 1000 users

Pseudonym	Users (%)	OT (%)	RY (%)	RT (%)	I-RT (%)	Negative sentiment (%)	Top topic	Top hashtags
Ex-Yu TS1	<b>24</b>	<b>27</b>	15	16	81	66	Mornings	euphoria, caosvima
RS democrats	18	18	<b>33</b>	24	82	67	Local politics	izbori2022, srbija
RS conservatives	12	10	15	11	69	<b>69</b>	Local politics	srbija, serbia
Croats	12	7	8	3	73	61	Mornings	dora2022, croatia
Ex-Yu TS2	5	8	11	4	64	62	Mornings	newprofilepic
Bosniaks	4	3	4	1	73	63	Local politics	sarajevo, bih
Vučić	4	9	1	<b>30</b>	<b>99</b>	64	Nationalism	vucic, srbija
Photography	4	<1	<1	1	63	54	Mornings	photography, nature
Zvezda	3	3	2	1	56	65	Sports	kkcz, fkcZ
RS radicals	2	<1	<1	<1	59	55	Russia/ Ukraine	serbia, srbija
Partizan	2	3	2	1	56	63	Sports	kkp, kkpartisan
Montenegro	2	3	3	2	79	67	Local politics	tbok, crnagora
Rest	7	3	3	4	70	59	Nationalism	srbija, caosvima

The table shows the percentage of retweet network users belonging to a community (Users), the percentage of original tweets (OT), replies (RY), and retweets (RT) posted by the community. I-RT is the percentage of the retweets of a community that retweet the same community (intra-retweets). Negative sentiment is the percentage of tweets with negative sentiment for a community. Finally, it shows the topic with the most tweets and the two most used hashtags. Bold values indicate the highest value for each column

**Table 2** General statistics for the invasion communities with more than 1000 users

Pseudonym	Users (%)	OT (%)	RY (%)	RT (%)	I-RT (%)	Negative sentiment (%)	Top topic	Top hashtags
Ex-Yu TS1	<b>19</b>	<b>25</b>	14	14	76	67	Truth/lies	pze, caosvima
Pro-Ukraine	19	15	17	11	80	68	Russia/ Ukraine	ukraine, russia
RS conservatives	18	16	21	22	78	<b>71</b>	Russia/ Ukraine	ukraine, nato
RS democrats	15	17	<b>27</b>	20	74	66	Russia/ Ukraine	ukraine, izbori
Ex-Yu TS2	5	8	11	4	65	63	Truth/Lies	pze, newprofilepic
Photography	4	<1	<1	1	50	55	Russia/ Ukraine	ukraine, nature
Vučić	3	8	1	<b>22</b>	<b>99</b>	54	Nationalism	vucic, srbija
Zvezda	3	2	1	1	47	59	Sports	fkcz, kkcZ
Partizan	2	2	1	1	56	65	Sports	kkpartisan, partizan
Ex-Yu communists	2	<1	<1	<1	49	69	Russia/ Ukraine	ukraine, russia
Ex-Yu TS3	2	<1	<1	<1	57	58	Russia/ Ukraine	ukraine, eu
Rest	9	3	3	3	60	59	Russia/ Ukraine	ukraine, srbija

See Table 1 for a description of the measured values



and users are Serbian, there are more peculiar political fractions in the “Serbian” part of the network. We observe communities based on political ideology in Serbia, such as RS Conservatives, RS Democrats, and RS Radicals. The non-Serbian political communities are more general and based on nationality, such as Bosniaks, Croats, and Montenegro.

An interesting case of a Serbian political community is the extreme echo chamber of the Vučić community. This community is formed around the profile of the Serbian president and his political party SNS (Srpska Napredna Stranka). With less than 4% of the users (3.76% for the pre-invasion, and 2.84% for the invasion snapshot), it produces even 29.73% and 21.92% of all the retweets for the pre-invasion and invasion snapshots respectively. Yet, 99% of these retweets are intra-community.

### **Content extraction**

We extract three types of data from the content the retweet communities produce: hashtags, news sources, and topics with sentiment.

Since the BCMS macro-language is interchangeably written in Cyrillic and Latin, we standardize the tweets by transliterating all Cyrillic tweets into Latin. This enables a more accurate comparison of topics through the utilized language models, while also unifying hashtags that are written in different scripts.

*Hashtags.* Out of all the tweets in the dataset, 3.3% contain a hashtag. We observe 58K lowercase unique hashtags, out of 19.3% are used both in the pre-invasion and invasion snapshot. Their popularity distribution follows a steep power-law curve, with the top 10 covering 20.7% of all hashtag examples. Four Serbian domestic hashtags (#fkcz, #sns, #srbija, #vucic) appear in the top 10 in both pre-invasion and invasion snapshots. Expectedly, #ukraine and #russia appear in the top 10 for the invasion snapshot, but not in the top 10 of the pre-invasion snapshot.

*News Sources.* URLs are present in 12.3% of the tweets. More than 1.1M URLs were shared through the tweets, out of 600K unique. These URLs originate from 18K different domains. In order to robustly measure news source similarities, we filter out the domains that belong to social media platforms, such as t.co (Twitter), fb.watch (Facebook),youtu.be (Youtube), etc., as well as URL shorteners, such as bit.ly and ow.ly. After the filtration, 356K URLs remain, with more than 99% pointing to news sources.

Top-level domains uncover the following: 38.4% of the URLs originate from Serbia (.rs), 8.9% from Croatia (.hr), 6.7% from Bosnia and Herzegovina (.ba), and finally 4.5% from Montenegro (.me). We notice an increase of international top-level domains in the invasion snapshot, going from 38.5% to 43.5%, suggesting a shift towards international issues, with domains such as dw.com (the German “Deutsche Welle”) and sputniknews.com (the Russian “Sputnik News”) gaining more visibility in the region by reporting on the invasion.


*Topics.* To identify the topic of tweets, we apply BERTopic, a clustering technique that utilizes embeddings from a multilingual LLM sentence transformer. Details on the technique and the hyperparameters used in this scenario can be found in subsection [Topic similarity with sentiment](#).



Before applying topic modeling to the tweets, we removed URLs and mentions, as well as stripped the hashtags from the “#” character to adapt the tweets to the model trained on more structured text. Finally, we excluded tweets with less than 20 characters to reduce noise and increase the quality of the embeddings.

The topic modeling resulted in 46 distinct topics, which we named accordingly by manually analyzing the 20 most representative keywords of each topic. We observe five topics tightly related to the invasion: Russia/Ukraine, Putin, Sanctions to Russia, War and Drones/Aircraft.

To enrich the information on topics, we also detect the sentiment of tweets using a transformer model with BERT architecture trained on BCMS data, called BERTić (Ljubešić and Lauc 2021). The model is fine-tuned on the BCS Political Sentiment dataset (Mochtak et al. 2022) which contains sentence-level labeled data from the Bosnian, Croatian, and Serbian parliaments (Mochtak et al. 2022). The model classifies text as “Negative” or “Non-Negative”, the latter covering positive or neutral. The sentiment results imply that the BCMS tweetsphere is dominated by negative sentiment, with 65.1% “Negative” tweets for the pre-invasion snapshot, and 65.7% for the invasion snapshot.




**Fig. 3** Volumes and sentiment ratio of the topics for the pre-invasion (left) and invasion (right) snapshots produced by BERTopic on the invasion snapshot of the tweets. Purple color highlights topics related to politics. The red color highlights the part of tweets with negative sentiment. The yellow color highlights the non-negative tweets

Figure 3 shows the distribution of topics and their sentiment for the pre-invasion and invasion snapshot. Expectedly, the biggest difference in the two distributions lies in the drastic increase of invasion-related topics, as they become the most discussed topics with over 15% of all tweets and almost 80% negative tweets.

**Topological similarity of communities**

There are multiple ways of measuring community similarity based on the topological proximity between the communities. Here, we consider the interaction between the communities as our topological baseline. Similarity of community *A* to community *B* is represented as the number of retweets where community *A* is retweeting community *B*, normalized by the total retweets of community *A*. Note that this similarity is asymmetric, with typically smaller communities having more interaction with large communities than vice versa.

Figure 4 presents a chord diagram illustrating the topological similarities, or interactions, among the invasion communities. Upon examination, we observe a general trend of low similarities between the communities, with smaller communities displaying greater interaction with larger communities, while the largest communities tend to be more internally cohesive. Notably, there are interesting exceptions to this pattern, such as the heightened interaction between RS Democrats and RS Conservatives, as well as between RS Democrats and Pro-Ukraine. The interactions suggest that the RS Democrats foster inter-community discussions.



**Fig. 4** Chord diagram on the topological similarity between the invasion communities. A value represents the number of retweets where community **A** is retweeting community **B**. The edge colors correspond to the retweeting community. Community sizes correspond to their total number of retweets. The diagram shows relatively low interaction in the network with the larger communities retweeting mostly themselves in an echo-chamber manner

However, these results alone fall short when analyzing communities within a specific context, such as the time of the invasion of Ukraine. To delve deeper, we need to address questions regarding the content-based similarity of communities, their alignment or conflict in viewpoints, their information sources, and whether they possess divergent perspectives despite topical similarities. Given the limitations of relying solely on the retweet network for answers, we leverage content-based analysis to compare communities and quantify their similarity.

### Content-based similarity

In our analysis, we explore content-based similarity by examining hashtags, news sources, and topic distributions with sentiment. However, it's important to note that the same approach can be effectively applied to other types of textual data as well.

For hashtags and news sources, we extract the top  $N$  (in our case we set  $N = 100$ ) most used hashtags and news sources by each community, sorted by frequency. Then we evaluate the similarity of the rankings using the Rank-Biased Overlap (RBO), as described in subsection [Hyperlink ranking similarities in social media](#), with  $p = 0.9$ . A maximum value of 1 would mean an identical list, while 0 would mean a completely disjoint list.

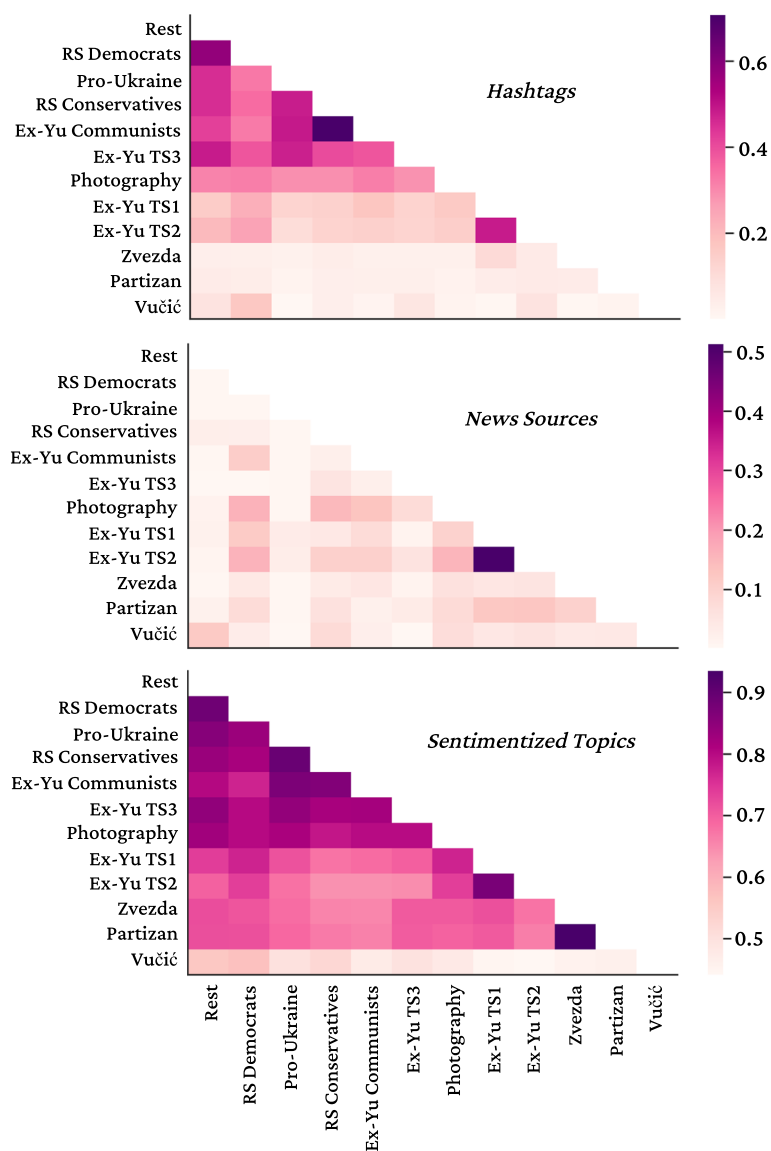
For topics, the more sophisticated technique based on the topics and their sentiment produced by LLMs, we measure community similarity by measuring the Jensen-Shannon distance of the topic probability distributions. The methodology is described in subsection [Topic similarity with sentiment](#).

Figure 5 shows the similarity of communities based on the three textual aspects, each represented by a heatmap of similarities, sorted via blockmodeling on the hashtag results. As each aspect reveals novel insights into the similarities between communities, we follow with an interpretation of the results for each of the layers.

*Hashtag Similarity.* The first heatmap of Fig. 5 shows the hashtag RBO similarity matrix of the communities in the invasion snapshot. The results reveal two strong clusters of communities not visible in the topological similarities and four non-clustered communities. The first big cluster (consisting of the communities: RS Democrats, Ex-Yu TS3, Pro-Ukraine, Ex-Yu Communists, and RS Conservatives) contains communities that predominantly discuss the invasion and use hashtags that refer to it (#russia, #ukraine, #nato, etc.). The second cluster (Ex-Yu TS1 and Ex-Yu TS2) belongs to the non-political Ex-Yu tweetsphere communities, with hashtags and content not related to politics.

Despite being purely political, the Vučić community shows very low similarity to any of the political communities. In fact, Vučić is the least similar community to the Pro-Ukraine community (even less similar than the sports-oriented Partizan and Zvezda), as the users from the first do not use hashtags related to the invasion. Rather, they use hashtags such as #srbija, #vucic, and #sns.

*News Sources.* News sources bring a different perspective to the analysis, previously unseen by the topology and hashtags. While the hashtags have more intuitive interpretability, understanding similarity in news sources requires additional manual or automatic exploration of the domains for their background and content. In our results, news



**Fig. 5** Hashtag, news source, and topic similarity of the invasion communities. Both hashtag and news source similarities are calculated using the Rank-biased overlap on the top 100 most popular hyperlinks. The topic similarity is calculated using Jensen-Shannon distance on the probability distributions of the 46 BERTopic topics for each community

source similarities contain a weaker signal than the hashtags due to their lower cross-community overlaps.

In Serbia, there are two media clusters, each catering to a specific political group—one supporting the opposition (anti-government) and the other aligned with the government (Djordjević 2020). The first is mostly shared by the RS Democrats community, with news sources such as N1 ([rs.n1info.com](http://rs.n1info.com)), Nova ([nova.rs](http://nova.rs)), as well as Western news sources reporting on the invasion, such as Reuters ([reuters.com](http://reuters.com)), and The Guardian ([theguardian.com](http://theguardian.com)). The second media cluster is shared mostly among the Vučić community and the RS Conservatives. The news sources specific to this cluster are the pro-government Informer ([informer.rs](http://informer.rs))

and Kurir (`kurir.rs`), as well as the Russian Sputnik (`rs-lat.sputniknews.com`), a domain suspended in the EU. These results are also visible through the heatmap, where we observe a particularly low similarity between `RS Democrats` and `RS Conservatives` despite their relatively high similarity in topology and hashtags, suggesting a news source polarization.

Finally, the `Pro-Ukraine` community does not have major similarities with any of the other communities, as it is the only community that is not dominated by Serbian domains, but it consists of mostly Bosnian (`radiosarajevo.ba`, `bhrt.ba`, `slobodna-bosna.ba`), Croatian (`dalmatinskiportal.hr`, `zagreb.info`) and Montenegrin media (`vijesti.me`, `bankar.me`).

*Topics.* Next, we discuss the similarity based on the topics with sentiment. The last heatmap in Fig. 5 holds comparable results to the hashtag similarity results, with a few differences. The first difference is that the `Zvezda` and `Partizan` communities now form the strongest cluster despite their low topological, hashtag and news source similarity. The topics similarity approach finally shows an evident fact for us as readers: these two communities actually have one main common interest in sports, yet they are also the two main opposing factions in Serbian sports culture. Hence, exceptional similarities in topic distribution and low similarities in topology or other types of content could point to potential polarization between two communities.

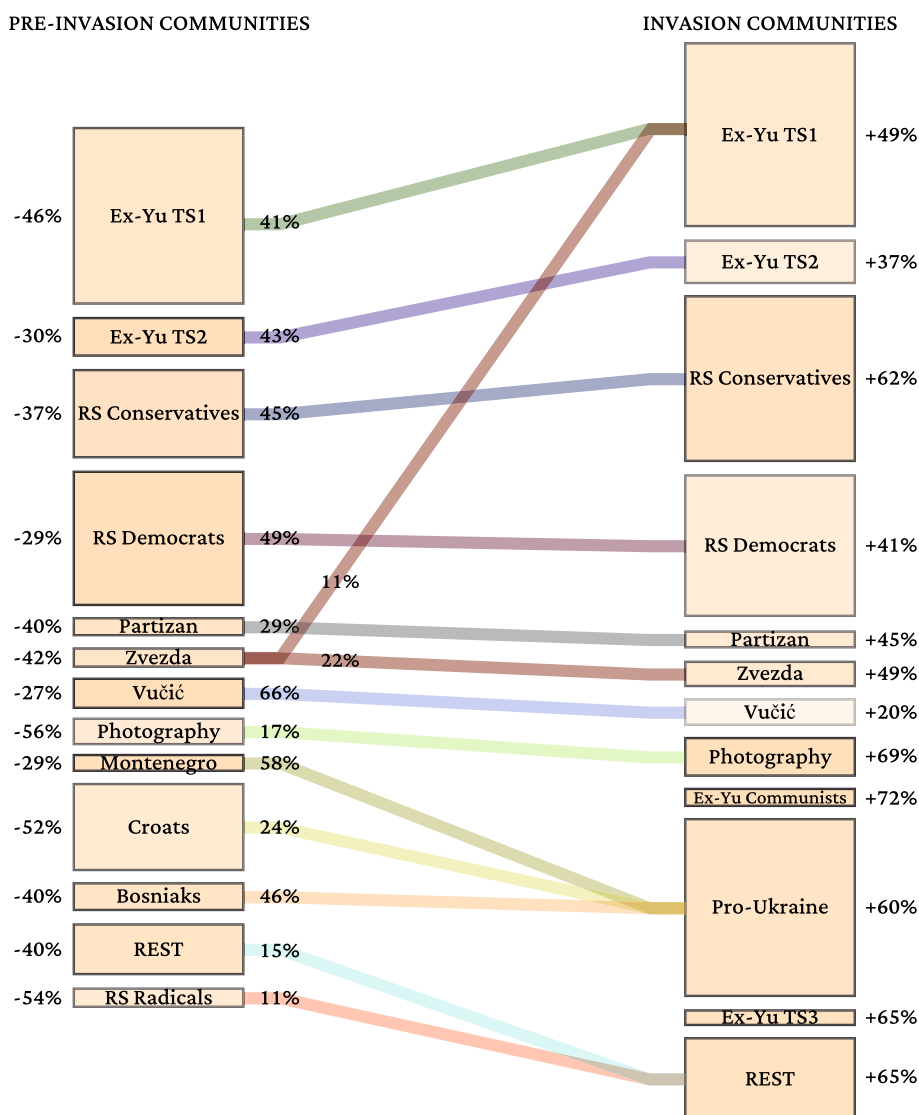
Another case where a similar phenomenon is visible is the `Pro-Ukraine` and `RS Conservatives`. Although the topology and news source results show low similarity between the two communities, the topics show that the `Pro-Ukraine` has the highest topical similarity with `RS Conservatives` and vice versa. Politically, these results are expected, since the `RS Conservatives` community is heavily interested in the conflict while supporting Russia.

Finally, the topical similarity suggests even more clearly that the `Vučić` community is the only one that is especially different from the `Pro-Ukraine` community. Despite the `Russia/Ukraine` topic being the most dominant topic in the invasion snapshot (see Fig. 3), it is only the 19th most discussed topic by the `Vučić` community. On another note, what makes the `Vučić` community different is that it is the least negative out of all communities, drastically going from 64% in the pre-invasion to 54% negative tweets in the invasion (see sentiment results in Table 2). This is largely due to the positive campaign regarding the then-upcoming snap elections in Serbia held on 3 April 2022.

Here, considering the results for the `Vučić` community, it is important to emphasize the downside of analyzing major events on social media solely relying on event-specific keywords. In such a scenario, one could see what is present but fail to recognize the importance of the absence of certain communities, such as the main governing power in Serbia avoiding the discussion on the invasion.

### Similarity in community evolution

Community evolution refers to the concept of tracking and analyzing changes in communities over time (Dakiche et al. 2019; Rossetti and Cazabet 2018; Cordeiro et al. 2016). One standard approach is to observe a given timeline as a list of (overlapping or non-overlapping) network snapshots, and then apply community detection on each of



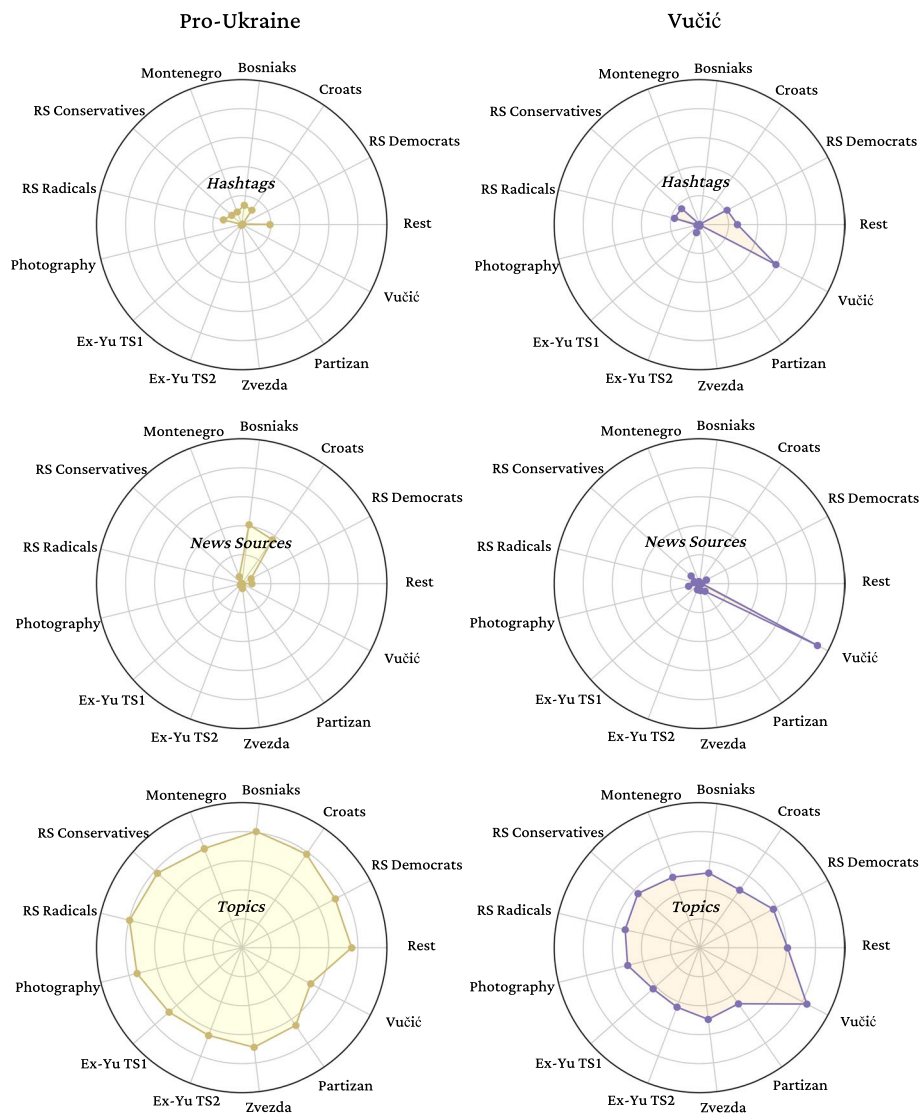
**Fig. 6** A Sankey diagram showing the transitions of users from the pre-invasion network communities (left) to the invasion network communities (right). Rectangle height is proportional to the community size. Percentages on the right-hand side near the pre-invasion communities show the portion of users found in the corresponding invasion communities. Percentages on the left-hand side near the pre-invasion communities show the portion of users of the pre-invasion communities that are inactive in the invasion period. Percentages on the right-hand side of the invasion communities show the portion of users not previously present in the large communities of the pre-invasion network.

the snapshots. This is followed by a matching phase where communities from snapshot  $t$  are identified in snapshot  $t + 1$  in a *many-to-many* manner. The matching is commonly approached by observing the user transitions using metrics such as Jaccard coefficient or BCubed F1 (Oliveira et al. 2014; Evkoski et al. 2021), which is conceptually similar to the topological similarity where only the node structure is observed.

We follow the same idea to compare the pre-invasion communities to the invasion communities. Figure 6 shows a Sankey diagram of user transitions from the first to the second network, where an edge represents the overlap of users between a pre-invasion

and an invasion community. An edge exists if at least 10% of the users from a pre-invasion community transition into an invasion one.

Most of the results are one-to-one relations between the pre-invasion and the invasion communities. A prime example is the Vučić community as the most stable, retaining 66% from the pre-invasion, gaining 20% new users, and losing only 27% of its pre-invasion users. An exception to the one-to-one relations is the merging of the Croats, Bosniaks, and Montenegro communities, which led to the formation of the Pro-Ukraine community. Although matching according to the user overlaps would be relatively accurate, what are the content-based similarities between the pre-invasion and the invasion communities? And do we observe any major contextual shifts in the communities?



**Fig. 7** Similarities of the Pro-Ukraine and the Vučić community to the pre-invasion communities according to hashtags, news sources, and topic distributions



Here, we showcase how content comparison can be beneficial in a community evolution scenario by getting insight into the origin of the invasion-related communities via their content-based similarity with pre-invasion communities. To illustrate the results, we examine two intriguing examples: the invasion communities `Vučić` and `Pro-Ukraine`. In Fig. 7, we visualize their hashtag, news source, and topic similarity to the pre-invasion communities.

The results for the `Vučić` community show a strong similarity with a single pre-invasion community, confirming the stable continuation of the community not only in terms of users but also in content. Meanwhile, the news source similarities of the `Pro-Ukraine` confirm that it is a continuation of the `Bosniaks` and `Croats` communities, but almost no similarity to any community according to hashtags and no outstanding similarity on the level of topics with any particular community. While the user transition results from Fig. 6 point that the `Pro-Ukraine` community formed as a result of merging three pre-invasion communities, the low content similarity results show that the main homophily factor (the invasion) of this newly merged community was not previously present in the network.

In conclusion, content-based similarity analysis of evolving communities can help in understanding the reasons behind merging and dissolving communities, topical shifts of stable communities, or the disassociation phenomena in dying communities. Additionally, it can help in more accurate tracking of evolving communities in network snapshots, where matching communities from snapshot  $t$  with communities from snapshot  $t + 1$  becomes a challenge.

## Discussion

The discussed approaches for multifaceted content-based comparison of communities in social networks allow the analyst to go beyond the topological structure of communities and investigate their similarity in terms of the textual content they produce. In our application of this methodology to the Ex-Yugoslavian retweet network, we substantiate this claim: little or no interaction between two communities does not mean that their content is divergent. We also propose numerical quantification of content similarity based on three types of textual data: hashtags, topics and their sentiment, and news sources.

Hashtags are widely used in social network analysis, but it is important to note that they represent a small fraction of tweets in the Ex-Yugoslavian dataset, accounting for just 3%. As a result, it is crucial to determine if hashtag users accurately reflect the communities under investigation. We supplement the hashtag analysis with a topic comparison based on transformer-based, context-aware large language models to generate a topic distribution that covers the entire Twitter corpus (with the exception of very short tweets). However, both hashtags and topics can conceal significant content-based polarization in cases where third-party news sources drastically differ. This is where the comparison of communities based on news sources comes into play, adding another layer of complexity that complements the first two. In our application of this approach, we show how multiple layers of content-based comparison can provide additional insight into a complex political landscape. The same ideas can be applied in other contexts, data sources, or even data types with the goal of obtaining a multidimensional insight into the phenomenon of interest

### **Summary of the application results**

Our analysis of the retweet network communities reveals how the Ex-Yugoslavian tweetsphere reacted in the first weeks of the Russian Invasion of Ukraine.

The Ex-Yugoslavian retweet network is heavily political, and there are mainly two types of communities: political with a moderate to high reaction to the topic (e.g., RS Democrats for moderate, and RS Conservatives for high); and non-political with a low reaction to the topic (e.g., Ex-Yu TS1 and Zvezda). An exception is the Vučić community which represents the governing power in Serbia; it is the only political community that actively avoids the discussion about the invasion.

The response from the predominantly Serbian communities is sporadic, with the retweet network reflecting Serbian society's ambivalence toward the conflict, with the exception of Serbian conservatives and radicals who openly support Russia. Meanwhile, the three pro-Ukraine communities, although differing in their interests and news sources in the pre-invasion period, are responding to the conflict by coalescing into one large community that openly supports Ukraine in the conflict.

On another note, we emphasize the importance of analyzing major events on Twitter by considering the entire tweetsphere instead of solely relying on event-specific keywords. While selecting a dataset based on a pool of invasion-related keywords may offer clearer results on the polarization surrounding the invasion, it fails to provide a comprehensive view of the situation and can result in an incomplete and potentially biased analysis. Through the analysis of the complete network, we discovered that certain strong political factors deliberately avoid the topic and maintain a neutral stance, exemplified by the extreme political echo chamber of the Serbian president Vučić. Therefore, it is important to analyze the entire tweetsphere to obtain a more nuanced and accurate understanding of major events on Twitter.

### **Limitations**

Although we used sophisticated methods such as large language models, the final similarity score generated by these methods is general and limited in its ability to capture the nuances of the phenomenon being analyzed. Thus, we caution that it is important to carefully study the content of the communities themselves, in addition to the similarity scores, before drawing any conclusions. Ultimately, the comparison only serves as a means of quantifying a general similarity, and further analysis is needed to fully understand the similarities and differences between the communities.

Additionally, since most types of content have limitations in terms of their coverage (e.g., not all tweets contain URLs or hashtags), one should also account for the representability of the content types for the communities and drop certain comparison perspectives if necessary. In our case, for example, assessing community similarity through hashtags is insufficient, as tweets with hashtags represent just 3% of all tweets.

A limitation regarding the application results is the usage of Twitter data as a source of information. Although it is the most researched social network due to its accessibility, Twitter comes with inherent limitations, primarily due to its potential lack of representativeness of the broader population. Twitter users constitute a self-selected sample, with demographics and characteristics that may not mirror those of the general population. This selectivity can introduce biases and skew the findings obtained from analyzing

Twitter data. Consequently, generalizing insights derived solely from Twitter to the broader population should be done cautiously. While acknowledging its limitations, leveraging Twitter data as a preliminary tool can still facilitate valuable avenues of inquiry.

### **Prospects**

There are several directions for further development in community comparison. We can use the proposed approach to compare communities from any two social networks. In subsection [Similarity in community evolution](#), we showed that we can use the content-based comparison of communities to investigate their content evolution. The application of the two network snapshots shows that the content-based comparison can also help in differentiating causes for community evolution events (merging, splitting, shrinking, growth, or death of communities). Automatic tools could also be developed on top of this, easing the processes of community-level network comparison and evolution analysis.

While we covered hashtags, topics, and news sources as the three main layers of content-based comparison, one can use a similar approach to compare emojis, user metadata, user descriptions, etc. Furthermore, one could compare images and videos by comparing their automatically generated descriptions. Moreover, with the rise of transformer-based machine translation and multilingual transformer language models that map text from different languages into a common latent space, content-based comparison of communities is also not limited to one language and one can potentially compare the content of multilingual social networks.

### **Methodology**

In this section, we cover the methodology for measuring the content-based similarity of social network communities through hyperlinks (such as hashtags and URLs), and topic distributions with sentiment.

#### **Hyperlink ranking similarities in social media**

A hyperlink is a reference to data that a user can follow or be guided to by clicking or tapping. Hyperlinks are crucial in social media, as they allow access to information from third-party sources and threading. There are three main types of hyperlinks on Twitter (and most of the other social media platforms): hashtags, URLs, and mentions (Chang and Iyer 2012). For Twitter, the metadata collected from the Twitter API contains a list of hashtags, URLs, and mentions for each tweet.

Typically, the frequency of unique hyperlinks in social networks follows a power law distribution with a long and low signal tail (Muchnik et al. 2013). In other words, the most popular hyperlinks in a community hold most of the hyperlink information. Hence, we first compute ranked lists of hyperlinks for each community, sorted by frequency. The ranked lists can contain a different number of elements and have a different popularity scale for each community, depending on the size of the community, activity levels, and the variety of content they produce.

A fitting metric to evaluate the similarity of ranked lists, especially hyperlinks in social media content, is Rank-Biased Overlap (RBO) (Webber et al. 2010). RBO is designed for measuring the similarity of rankings where the significance of each rank is not equal

and incorporates a probabilistic approach that models a user's persistence. It applies top-weighting, which means that the similarity of the top elements in the list is more impactful than the similarity of the elements from the bottom. This makes RBO suitable for our comparison due to the power-law distribution of the hyperlinks. It is defined as follows:

$$RBO(S, T, p) = (1 - p) \sum_{d=1}^{\infty} p^{(d-1)} \cdot A_d$$

$$A_d = \frac{|S_d \cap T_d|}{d}$$

where  $S$  and  $T$  are the two indefinite rankings of hashtags,  $S_d$  and  $T_d$  refer to the element with rank  $d$  in the lists, and  $p$  determines how steep the decline in weights in the top-weighting is. The smaller  $p$ , the more top-weighted the metric is. In the limit, when  $p = 0$ , only the top-ranked item is considered, and the RBO score is either zero or one. On the other hand, as  $p$  approaches arbitrarily close to 1, the weights become arbitrarily flat, and the evaluation becomes arbitrarily deep. In practice,  $p = 0.9$  means that the first 10 ranks have 86% of the weight of the evaluation; and  $p = 0.98$  would give 86% of the weight to the first 50 ranks.

Rank-Biased overlap falls in the range  $[0, 1]$ , where 0 means disjoint, and 1 means identical rankings, or in our case, hashtag-identical communities.

#### Topic similarity with sentiment

The introduction of transformer-based language models led to significant advancements in deep learning and major improvements in natural language tasks, even outperforming humans on specific benchmarks such as SuperGLUE (Wang et al. 2019) and BIG-BENCH (Srivastava et al. 2022). The transformers adopt the mechanism of self-attention, weighing the significance of each part of the input data. It provides context for all elements in the input and it allows the models to grasp semantics like never before.

Two of the tasks in which transformer-based language models had a major improvement are sentiment identification (Mochtak et al. 2022) and topic modeling (Grootendorst 2022). Sentiment identification, in its simplest form, classifies a text by its sentiment load into two or three classes: negative and non-negative; or negative, neutral, and positive. Sentiment identification with transformer-based language models is performed by fine-tuning the language model to the specific task of sentiment identification via manually annotated training data. During fine-tuning, the language model's parameters are adapted in such a manner that the document representation produced by the language model is most informative for the final classification task of placing a piece of text into the corresponding sentiment category (Wolf et al. 2019).

Topic modeling with large language models works by first applying transformers [known as sentence transformers (Reimers and Gurevych 2019)] to embed documents into vectors that have high cosine similarity if they are contextually similar. Then these documents are clustered using already-established numerical clustering techniques. The state-of-the-art technique we apply for topic modeling is BERTopic (Grootendorst 2022), which first extracts a representation of each document via BERT-based text embedders (Reimers and Gurevych 2019). It then reduces the number of dimensions of the vector

representations using a dimensionality reduction technique, such as UMAP or PCA (Jia et al. 2022). Then it applies a clustering algorithm (such as HDBSCAN or KMeans) to group the documents into topics (Ezugwu et al. 2022). The main practical difference between HDBSCAN and KMeans clustering is that HDBSCAN allows tweets to be categorized as outliers (not belonging to any topic), while the KMeans inherently forces all tweets into a cluster and allows potential outliers to distort the pre-selected centroids. Finally, after the clustering, BERTopic calculates a bag-of-words (c-Tf-Idf) representation for each topic, in order to provide human-readable representations of the topics consisting of the words with the highest Tf-Idf score for each topic.

Topic modeling with large language models works by applying transformers to embed the documents into vectors and then clustering the vectors using some typical numerical clustering techniques. The state-of-the-art technique we apply for topic modeling is BERTopic (Grootendorst 2022). First, it extracts a representation of each document via BERT-based text embedders (Reimers and Gurevych 2019) that optimize embeddings in a way that semantically similar texts have high cosine similarity. Second, it reduces the number of dimensions of the vector representations using a dimensionality reduction technique, such as UMAP or PCA (Jia et al. 2022), so that Euclidean similarity measures can be applied to a dense low-dimensional space. Third, it applies a clustering algorithm (such as HDBSCAN or KMeans) to group the documents into topics (Ezugwu et al. 2022). Finally, after the clustering, BERTopic calculates a bag-of-words (c-Tf-Idf) representation for each topic, in order to provide human-readable representations of the topics consisting of the words with the highest Tf-Idf score for each topic.

For this paper, we use a multilingual text transformer with 368 dimensions as an embedder. For dimensionality reduction, we opted for UMAP with 100 neighboring points for manifold approximation, and reduce the embeddings to 10 dimensions. We cluster the tweets into topics using HDBSCAN and Euclidean distance, limiting the minimum cluster size to 10K. The reason we chose HDBSCAN over KMeans is that HDBSCAN allows tweets to be categorized as outliers (not belonging to any topic), while the KMeans inherently forces all tweets into  $K$  clusters and allows potential outliers to distort the pre-selected centroids.

Finally, we extract unigrams and bigrams for the Tf-Idf human-readable representation with a minimum n-gram frequency of 10, excluding the typical BCMS stopwords in the process. We used guidelines by the BERTopic authors and applied an iterative process to define the BERTopic hyper-parameters, inspecting a few preliminary results before attaining coherent clusters with a low number of outliers.

We explored various approaches for applying BERTopic to the data. Our goal was to address the challenges posed by the interdependence of topic distribution between the pre-invasion and invasion snapshots while ensuring comparability and accurate representation of topics. More specifically, we considered three possible approaches for training the model:

- *Option 1:* Train and apply one model on both snapshots—the main issue with this option is the inter-dependence of the topic distribution between the pre-invasion and invasion snapshots. Here, the topic space of the invasion tweets is distorted by the dependence on the pre-invasion tweets, and more importantly, the topic space of

the pre-invasion snapshot is dependent on major events that have not yet occurred. We opted out of this option due to its potential to damage the topic distribution comparison for the second application scenario.

- *Option 2:* Train and apply a model for each snapshot—with this approach, we eliminate the topical interdependence between the two snapshots, but introduce another comparability issue. Here, the two snapshots would have a different embedding distribution, resulting in different topics and different human-readable topic representations. Thus, with  $M$  topics on one side, and different  $N$  topics on the other, the comparison of the invasion topic distribution with the pre-invasion topic distribution becomes ambiguous.
- *Option 3:* Train a model on the invasion snapshot and apply it to both snapshots—this solution is the one we applied, as it mitigates the issues of the first two options, by projecting the pre-invasion tweets into the topic distribution of the invasion snapshot.

To enrich the information on topics, we also detected the sentiment of tweets using a transformer model with BERT architecture trained on BCMS data, called BERTiC (Ljubešić and Lauc 2021). Once we inferred the topic and sentiment of the posts, we computed a frequency distribution of topics and their sentiment for each of the communities. In other words, we count the number of negative and non-negative posts for each topic, in each community. To compare the communities irrespective of their size, we transform the frequencies into probability distributions.

A suitable measure of the similarity between two probability distributions (in our case the topic distributions with sentiment for two communities),  $P$  and  $Q$ , is the Jensen-Shannon divergence ( $JSD$ ) (Lin 1991):

$$JSD(P \parallel Q) = \frac{1}{2}KLD(P \parallel M) + \frac{1}{2}KLD(Q \parallel M),$$

where  $M$  is the average of the two distributions:

$$M = \frac{1}{2}(P + Q).$$

$JSD$  is defined in terms of the Kullback–Leibler divergence ( $KLD$ ) (Kullback and Leibler 1951):

$$KLD(P \parallel Q) = \sum_x P(x) \cdot \log_2 \left( \frac{P(x)}{Q(x)} \right)$$

The square root of  $JSD$ , which makes the measure a metric, is known as Jensen-Shannon distance ( $JS$ ) (Endres and Schindelin 2003):

$$JS(P \parallel Q) = \sqrt{JSD(P \parallel Q)}, \quad 0 \leq JS(P \parallel Q) \leq 1.$$

$JS(P \parallel Q)$  of 0 indicates that  $P$  and  $Q$  are identical distributions, while values close to 1 indicate very different distributions.

Finally, to show the similarity between two communities (which we use for the visualizations in the results), we use the Jensen-Shannon similarity ( $JSS$ ) as:

$$JSS(P \parallel Q) = 1 - JS(P \parallel Q)$$

indicating high similarity of two communities for values close to 1 and low similarity for values close to 0.

## Conclusions

We discussed various approaches to measure the similarity of social network communities by analyzing the textual content they generate. One approach we employed leverages recent advancements in natural language processing tools, specifically transformer-based large language models, to compare content based on topics. We emphasize the significance of conducting a multifaceted analysis of community-generated content to gain a comprehensive understanding of their distinctions, especially when examining social networks in the context of specific societal events. In cases of polarization, relying solely on network topology fails to reveal content similarity, which is crucial for comprehending the nature of polarization itself. Given the growing influence of social networks in shaping public opinion and discourse, the accurate and effective analysis of community similarities and differences becomes increasingly vital. Content-based comparisons offer valuable insights for community analysis and have far-reaching implications across diverse domains, such as marketing, politics, and social policy.

To demonstrate the effectiveness of content-based comparisons, we analyzed similarities and differences across groups in the reaction of the Ex-Yugoslavian tweetosphere to the Russian invasion of Ukraine. We built an exhaustive collection of tweets in four<sup>5</sup> official varieties of the Serbo-Croatian (BCMS) macro-language and applied community detection on the retweet network. We investigated content-based community similarity by analyzing hashtags, news sources, and transformer-based topics with sentiment. This multifaceted comparison unveiled important social phenomena by revealing community similarities and dissimilarities that were not apparent through the retweet network's topology. Notably, our study represents the first exploration of the BCMS tweetosphere as a whole.

The versatility of content-based comparisons extends beyond analyzing a single social network or language. By utilizing appropriate tools like multilingual language models or machine translation, this approach can effectively quantify content-based similarity between any two groups, regardless of the networks they belong to.

### Author contributions

BE wrote the initial draft of the paper and did the experiments. NL collected the Twitter data and was involved in the topic modeling. PKN was involved in formulating the research idea and supervised the work. All the authors analyzed the results, and contributed to, read, and approved the final manuscript.

### Funding

BE received funding from the Slovenian Research Agency research programme, research infrastructure P6-0436: Digital Humanities: resources, tools, and methods (2022–2027). NL received funding from the Slovenian Research Agency within the Slovenian-Flemish bilateral basic research project "Linguistic landscape of hate speech on social media" (N06-0099 and FWO-G070619N, 2019–2023) and the research programme "Language resources and technologies for Slovene" (P6-0411).

### Availability of data and materials

The datasets generated and/or analyzed during the current study are not publicly available due to Twitter's terms of service but are available from the corresponding author upon reasonable request.



**Code availability**

The BERTopic implementation we applied can be found on the Github repository at <https://github.com/MaartenGr/BERTopic>. The Rank-Biased Overlap implementation we used is available at <https://github.com/changyaochen/rbo>. For Jensen-Shannon similarity, we used the Scipy library (Virtanen et al. 2020).

**Declarations****Competing interests**

The authors declare that they have no competing interests.

Received: 28 February 2023 Accepted: 10 June 2023

Published: 28 June 2023

**References**

- Amaral LA, Ottino JM (2004) Complex networks. *Eur Phys J B* 38(2):147–162
- Banac I (2009) What happened in the Balkans (or rather Ex-Yugoslavia)? *East Eur Polit Soc* 23(4):461–478
- Bisgin H, Agarwal N, Xu X (2010) Investigating homophily in online social networks. In: 2010 IEEE/WIC/ACM international conference on web intelligence and intelligent agent technology, vol 1, pp 533–536. IEEE
- Blondel VD, Guillaume J-L, Lambiotte R, Lefebvre E (2008) Fast unfolding of communities in large networks. *J Stat Mech Theor Exp* 2008(10):10008
- Chang H-C, Iyer H (2012) Trends in twitter hashtag applications: design features for value-added dimensions to future library catalogues. *Libr Trends* 61(1):248–258
- Cherepnalkoski D, Mozetič I (2016) Retweet networks of the European parliament: evaluation of the community structure. *Appl Netw Sci* 1:1–20
- Cinelli M, De Francisci Morales G, Galeazzi A, Quattrociocchi W, Starnini M (2021) The echo chamber effect on social media. *Proc Natl Acad Sci* 118(9):2023301118
- Cordeiro M, Sarmiento RP, Gama J (2016) Dynamic community detection in evolving networks using locality modularity optimization. *Soc Netw Anal Min* 6:1–20
- Dakiche N, Tayeb FB-S, Slimani Y, Benatchba K (2019) Tracking community evolution in social networks: a survey. *Inf Process Manag* 56(3):1084–1102
- Ding Y (2011) Community detection: topological vs. topical. *J Inf* 5(4):498–514
- Djordjević J (2020) Discourse in Serbian online media: Global expectations vs. local reality. *Digitalne medijske tehnologije i društveno-obrazovne promene* 9:87
- Durazzi F, Müller M, Salathé M, Remondini D (2021) Clusters of science and health related twitter users become more isolated during the Covid-19 pandemic. *Sci Rep* 11(1):19655
- Endres DM, Schindelin JE (2003) A new metric for probability distributions. *IEEE Trans Inf Theory* 49(7):1858–1860. <https://doi.org/10.1109/TIT.2003.813506>
- Evkoski B (2022) Master's thesis: community evolution analysis with Ensemble Louvain. <https://doi.org/10.13140/RG.2.2.21595.64806>
- Evkoski B, Ljubešić N, Pelicon A, Mozetič I, Kralj Novak P (2021) Evolution of topics and hate speech in retweet network communities. *Appl Netw Sci* 6(1):1–20
- Evkoski B, Mozetič I, Ljubešić N, Kralj Novak P (2021) Community evolution in retweet networks. *PLoS One* 16(9):0256175
- Ezugwu AE, Ikotun AM, Oyelade OO, Abualigah L, Agushaka JO, Eke CI, Akinyelu AA (2022) A comprehensive survey of clustering algorithms: state-of-the-art machine learning applications, taxonomy, challenges, and future research prospects. *Eng Appl Artif Intel* 110:104743
- Fortunato S (2010) Community detection in graphs. *Phys Rep* 486(3–5):75–174
- Fortunato S, Hric D (2016) Community detection in networks: A user guide. *Phys Rep* 659:1–44. <https://doi.org/10.1016/j.physrep.2016.09.002>
- Gleich DF (2015) Pagerank beyond the web. *Siam REV* 57(3):321–363
- Grčar M, Cherepnalkoski D, Mozetič I, Kralj Novak P (2017) Stance and influence of twitter users regarding the brexit referendum. *Comput Soc Netw* 4:1–25
- Grootendorst M (2022) Bertopic: neural topic modeling with a class-based TF-IDF procedure. arXiv preprint [arXiv:2203.05794](https://arxiv.org/abs/2203.05794)
- Jia W, Sun M, Lian J, Hou S (2022) Feature dimensionality reduction: a review. *Compl Intel Syst* 8(3):2663–2693
- Kullback S, Leibler RA (1951) On information and sufficiency. *Ann Math Stat* 22(1):79–86. <https://doi.org/10.1214/aoms/1177729694>
- Lancichinetti A, Fortunato S (2009) Community detection algorithms: a comparative analysis. *Phys Rev E* 80(5):056117
- Lancichinetti A, Fortunato S (2012) Consensus clustering in complex networks. *Sci Rep* 2(1):1–7
- Lin J (1991) Divergence measures based on the Shannon entropy. *IEEE Trans Inf Theory* 37(1):145–151. <https://doi.org/10.1109/18.61115>
- Lin T, Wang Y, Liu X, Qiu X (2022) A survey of transformers. *AI Open*
- Ljubešić N, Fišer D, Erjavec T (2014) Tweet-cat: a tool for building twitter corpora of smaller languages. In: *Proceedings of LREC*, pp 2279–2283
- Ljubešić N, Lauc D (2021) BERTiC - the transformer language model for Bosnian, Croatian, Montenegrin and Serbian. In: *Proceedings of the 8th Workshop on Balto-Slavic Natural Language Processing*, pp. 37–42. Association for Computational Linguistics, Kiyv, Ukraine. <https://aclanthology.org/2021.bsnlp-1.5>
- Malin B (2005) Unsupervised name disambiguation via social network similarity. *Workshop Link Anal Count Secur* 1401:93–102

- Mazzucchelli F (2012) What remains of Yugoslavia? From the geopolitical space of Yugoslavia to the virtual space of the web Yugosphere. *Soc Sci Inf* 51(4):631–648
- Metaxas P, Mustafaraj E, Wong K, Zeng L, O’Keefe M, Finn S (2015) What do retweets indicate? results from user survey and meta-review of research. *Proc Int AAAI Conf Web Soc Med* 9:658–661
- Mochtak M, Rupnik P, Ljubešić N (2022) The ParlaSent-BCS dataset of sentiment-annotated parliamentary debates from Bosnia-Herzegovina, Croatia, and Serbia. *arXiv preprint* [arXiv:2206.00929](https://arxiv.org/abs/2206.00929)
- Mochtak M, Rupnik P, Ljubešić N (2022) The Sentiment Corpus of Parliamentary Debates ParlaSent-BCS v1.0. Slovenian language resource repository CLARIN.SI. <http://hdl.handle.net/11356/1585>
- Morini V, Pollacci L, Rossetti G (2021) Toward a standard approach for echo chamber detection: reddit case study. *Appl Sci* 11(12):5390
- Muchnik L, Pei S, Parra LC, Reis SD, Andrade JS Jr, Havlin S, Makse HA (2013) Origins of power-law degree distribution in the heterogeneity of human activity in social networks. *Sci Rep* 3(1):1783
- Newman ME (2003) The structure and function of complex networks. *SIAM Rev* 45(2):167–256
- Newman ME (2004) Detecting community structure in networks. *Eur Phys J B* 38:321–330
- Novak PK, Amicis LD, Mozetič I (2018) Impact investing market on twitter: influential users and communities. *Appl Netw Sci* 3:1–20
- Oliveira M, Guerreiro A, Gama J (2014) Dynamic communities in evolving customer networks: an analysis using landmark and sliding windows. *Soc Netw Anal Min* 4:1–19
- Papic Z (1994) Nationalism, patriarchy and war in Ex-Yugoslavia. *Women’s Hist Rev* 3(1):115–117
- Poulin V, Théberge F (2019) Ensemble clustering for graphs: comparisons and applications. *Appl Netw Sci* 4(1):1–13
- Reimers N, Gurevych I (2019) Sentence-bert: sentence embeddings using siamese bert-networks. In: *Proceedings of the 2019 conference on empirical methods in natural language processing*. Association for computational linguistics. [arXiv:1908.10084](https://arxiv.org/abs/1908.10084)
- Rossetti G, Cazabet R (2018) Community discovery in dynamic networks: a survey. *ACM Comput Surv* 51(2):1–37
- Spertus E, Sahami M, Buyukkokten O (2005) Evaluating similarity measures: a large-scale study in the orkut social network. In: *Proceedings of the eleventh ACM SIGKDD international conference on knowledge discovery in data mining*, pp 678–684
- Srivastava A, Rastogi A, Rao A, Shoeb AAM, Abid A, Fisch A, Brown AR, Santoro A, Gupta A, Garriga-Alonso A, et al (2022) Beyond the imitation game: quantifying and extrapolating the capabilities of language models. *arXiv preprint* [arXiv:2206.04615](https://arxiv.org/abs/2206.04615)
- Strogatz SH (2001) Exploring complex networks. *Nature* 410(6825):268–276
- Surian D, Nguyen DQ, Kennedy G, Johnson M, Coiera E, Dunn AG (2016) Characterizing twitter discussions about HPV vaccines using topic modeling and community detection. *J Med Internet Res* 18(8):232
- Virtanen P, Gommers R, Oliphant TE, Haberland M, Reddy T, Cournapeau D, Burovski E, Peterson P, Weckesser W, Bright J, van der Walt SJ, Brett M, Wilson J, Millman KJ, Mayorov N, Nelson ARJ, Jones E, Kern R, Larson E, Carey CJ, Polat I, Feng Y, Moore EW, VanderPlas J, Laxalde D, Perktold J, Cimrman R, Henriksen I, Quintero EA, Harris CR, Archibald AM, Ribeiro AH, Pedregosa F, van Mulbregt P (2020) SciPy 1.0 contributors: SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nat Methods* 17:261–272. <https://doi.org/10.1038/s41592-019-0686-2>
- Wang A, Pruschatkun Y, Nangia N, Singh A, Michael J, Hill F, Levy O, Bowman S (2019) Superglue: a stickier benchmark for general-purpose language understanding systems. *Adv Neural Inf Process Syst* 32
- Webber W, Moffat A, Zobel J (2010) A similarity measure for indefinite rankings. *ACM Trans Inf Syst* 28(4):1–38
- Weng W, Zhu S, Xu H (2014) Hierarchical community detection algorithm based on local similarity. *J Digit Inf Manag* 12(4):274–280
- Wolf T, Debut L, Sanh V, Chaumond J, Delangue C, Moi A, Cistac P, Rault T, Louf R, Funtowicz M, et al. (2019) Huggingface’s transformers: state-of-the-art natural language processing. *arXiv preprint* [arXiv:1910.03771](https://arxiv.org/abs/1910.03771)
- Wong FMF, Tan CW, Sen S, Chiang M (2013) Quantifying political leaning from tweets and retweets. *Proc Int AAAI Conf Web Soc Med* 7:640–649

## Publisher’s Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.