

RESEARCH

Open Access



# Understanding the growth of the Fediverse through the lens of Mastodon

Lucio La Cava, Sergio Greco and Andrea Tagarelli\* 

\*Correspondence:  
andrea.tagarelli@unical.it  
Department of Computer  
Engineering, Modeling,  
Electronics, and Systems  
Engineering (DIMES),  
University of Calabria,  
87036 Rende, Italy

## Abstract

Open-source, Decentralized Online Social Networks (DOSNs) are emerging as alternatives to the popular yet centralized and profit-driven platforms like Facebook or Twitter. In DOSNs, users can set up their own server, or instance, while they can actually interact with users of other instances. Moreover, by adopting the same communication protocol, DOSNs become part of a massive social network, namely the *Fediverse*. Mastodon is the most relevant platform in the Fediverse to date, and also the one that has attracted attention from the research community. Existing studies are however limited to an analysis of a relatively outdated sample of Mastodon focusing on few aspects at a user level, while several open questions have not been answered yet, especially at the instance level. In this work, we aim at pushing forward our understanding of the Fediverse by leveraging the primary role of Mastodon therein. Our first contribution is the building of an up-to-date and highly representative dataset of Mastodon. Upon this new data, we have defined a network model over Mastodon instances and exploited it to investigate three major aspects: the structural features of the Mastodon network of instances from a macroscopic as well as a mesoscopic perspective, to unveil the distinguishing traits of the underlying federative mechanism; the backbone of the network, to discover the essential interrelations between the instances; and the growth of Mastodon, to understand how the shape of the instance network has evolved during the last few years, also when broadening the scope to account for instances belonging to other platforms. Our extensive analysis of the above aspects has provided a number of findings that reveal distinguishing features of Mastodon and that can be used as a starting point for the discovery of all the DOSN Fediverse.

**Keywords:** Decentralized online social networks, Mastodon instances, Structural network analysis, Community detection, Core decomposition, Graph pruning, Prestige ranking

## Introduction

In the last decade, we witnessed an unprecedented proliferation of Online Social Networks (OSNs). Roughly and generally speaking, OSNs aim to shrink timing and distances that characterize inter-personal relationships through the Internet. However, the extreme popularity gained by Facebook and the other worldwide available yet centralized OSN platforms (i.e., hosted and controlled by a single company) has soon led their owners to pursue a collateral social-marketing goal, which is mostly implemented

through content personalization mechanisms and advertisement strategies. As it is well-known, side-effects such as the formation of information bubbles and concerns about the protection of data and user privacy normally characterize most existing centralized OSNs.

The above aspects contributed to raise the opportunity for developing new paradigms of OSNs to become “user-centric” rather than “company-centric” platforms.

As a major consequence, privacy control, as well as spontaneous and recommendation-free communications among the users, are favored and unbiased as much as possible from the invasiveness of advertisements.

In this context, Decentralized Online Social Networks (DOSNs) are emerging as alternatives to the popular centralized platforms. DOSNs are built upon two key aspects: the availability of open-source software to allow everyone to set up their server hence avoiding centralization, and the existence of specific communication protocols to enable fluid interconnections between servers that embrace the same protocol.

These core components lead to a *federation* model, in which the servers, also called as *instances*, can communicate to each other through the same protocol. This implies that users which are signed up for a particular server can actually interact with users of other servers, analogously to what normally happens with email services. DOSNs hence become part of a massive social network, namely the *Fediverse*. As a consequence of this mechanism, users can use their accounts on a DOSN platform to follow users on other platforms, without needing an account there.

The Fediverse currently provides several services, such as *Mastodon* and *Friendica* for microblogging, *PeerTube* and *Funkwhale* for video hosting, *PixelFed* for image hosting. Among these platforms, Mastodon is the one that has encountered the greatest attention increase over the years. Mastodon provides a user experience comparable to Twitter (e.g., published contents are called *toots*, whereas the analogous of the retweet functionality is called *boost*) and Reddit (e.g., niche communities and content moderation are emphasized, however Mastodon communities are independent of each other). Moreover, Mastodon affords content warnings, i.e., synopses of toots that can preview disturbing content. Mastodon is built upon the *ActivityPub* protocol,<sup>1</sup> which implements a layer for client-to-server communications and another one for the server-to-server communications. Thanks to this protocol and a subscription-based mechanism (implicitly carried out by the instances), users can interact with each other even if they belong to different instances.

The extended followship mechanism in Mastodon also leads to an original yet remarkable timeline structure, namely *home*-timeline, which provides toots generated by followed users, *local*-timeline, which yields toots created within the instance, and *federated*-timeline, which contains all public toots from all users (either from the same instance or not) that are known to the instance where a user is registered.

Furthermore, Mastodon instances allow their users to apply rules and policies on the generated contents. Administrators can declare both the main topics of their instance and prohibited contents. Users can mark some contents as inappropriate

---

<sup>1</sup> <https://www.w3.org/TR/activitypub/>.

for a given instance by placing a *content warning* on the content itself. Along with the content warning, a spoiler (i.e., a textual component summarizing the obfuscated content) will be displayed to the user, letting her/him decide whether to view it or not.

Finally, Mastodon also allows administrators to close registrations for their instances, e.g., in the case of a “private instance” among friends, or to efficiently moderate contents. Nevertheless, this feature does not affect user interactions which, as outlined above, are guaranteed by specific protocols.

**Related work.** DOSN analysis is a relatively novel research field. Early works mainly investigated motivations, opportunities and challenges related to different solutions for the decentralized paradigm, from distributed systems like peer-to-peer networks to hybrid systems integrating external and private resources for storing user data (Guidi et al. 2018). Two surveys on these topics as well as on issues related to DOSN infrastructures, data management, privacy and information diffusion, can be found in Guidi et al. (2018), Datta et al. (2010). Focusing on more recently developed and open-source DOSNs, to the best of our knowledge, Mastodon is the only platform in the Fediverse that has received noticeable attention from the research community (Cerisara et al. 2018; Trienes et al. 2018; Zignani et al. 2018; Raman et al. 2019; Zignani et al. 2019; Zulli et al. 2020). Zulli et al. (2020) have recently performed a qualitative analysis based on an interview to a sample of instance moderators. From that study it emerges that the federative structure of Mastodon enables content variety and community autonomy, and also emphasizes horizontal growth between instances rather than growth within instances; however, any analysis of the interactions on the Mastodon instances is missing.

From a network science perspective, the studies by Zignani et al. (2018, 2019) are particularly relevant, as they were the first to analyze a portion of the Mastodon user-network, focusing on degree distribution, triadic closure, and assortativity aspects, and comparing such characteristics to those in Twitter (Zignani et al. 2018). From the analysis of the in-degree and out-degree distributions, Mastodon is found to show a more balanced behavior between followers and followees than what observed in Twitter.

Also, the 95% of users exhibit a difference between followers and followees bounded in the range  $(-250, 250)$ . Concerning social bots, the authors reported a low presence (around 5%), which is significantly lower than the 15% observed on Twitter by Varol et al. (2017). Clustering coefficient in Mastodon ranges between those of Facebook and Twitter. The degree assortativity in Mastodon was also inspected, considering source in-degree (SID), source out-degree (SOD), destination in-degree (DID), and destination out-degree (DOD). The authors observed lack of correlation between (SOD, DOD), (SOD, DID) and (SID, DOD), which indicates that users who follow many users are connected to users whose popularity may vary (DID) and who in turn follow few or many users (DOD). Moreover, the observed negative correlation  $(-0.1)$  between SID and DID implies that the higher the popularity of a user is, the less popular the users s/he follows will be. Overall, disagreement is observed between the degree assortativity in Mastodon and the ones shown by well-known social networks. Finally, Zignani et al. found that users’ hubiness is bounded within its instance and influenced by the latter. Also, in Zignani et al. (2019), the authors investigate how the decentralization process affects relationships between users, unveiling that instances

show individual footprints (based on degree distribution and clustering coefficient statistics observed on the top 10 instances in Mastodon) that influence relationships.

It should be noted that none of the above works have focused on the interactions between instances in Mastodon, and a number of related aspects are missing, such as distinction between online and offline instances, mesoscopic structure analysis, or backboning. Indeed, several open questions still remain to address on Mastodon, which inspired our study in this work.

**Contributions.** Our research stems from a twofold motivation: to provide a fresh view on Mastodon based on recently updated data, and to fill a lack of knowledge on topological features of the Mastodon network focusing at the *instance* level.

As previously discussed, early studies have primarily focused on the analysis of Mastodon users, and they captured a relatively small snapshot of Mastodon dated four years ago. Clearly, this might have overlooked salient traits that can be discovered at the instance level, as well as it raises the need for getting a timely picture of Mastodon which has presumably changed over time. To overcome these limitations, our study builds upon an up-to-date and representative network data over the instances, and utilizes it to provide insights into their relations. The goal is manifold: it includes the opportunity of enhancing our understanding of the *macroscopic* and *mesoscopic* structures of Mastodon to unveil the distinguishing traits of the underlying federative mechanism, and to discover the essential interrelations between the instances; but also we want to understand how the instance network has changed, within Mastodon as well as at the boundary of Mastodon itself.

We elaborate on the above aspects by developing an extensive analysis framework to answer the following research questions:

- Q1** *Network data and models:* How are the Mastodon instances detected and modeled as a network?
- Q2** *Structural features:* What are the salient structural features of the network of Mastodon instances, at *macroscopic* as well as *mesoscopic* level?
- Q3** *Fingerprint:* Are there any clues to the presence of notable phenomena that distinguishes Mastodon from centralized OSNs? How does a federative mechanism arise from the Mastodon instances?
- Q4** *Network backbone:* What is the backbone of the network of Mastodon instances, and does it preserve the structural features of the whole network?
- Q5** *Growth:* How has the shape of the network of Mastodon instances evolved during the last few years?

**Plan of the paper.** The remainder of the paper is organized so as to address the above stated research questions. Section "[Polite data crawling and network modeling](#)" describes our crawling methodology, the data collected and the network models we built upon this data (**Q1**). Section "[Structural analysis of the Instances network](#)" contains the structural analysis of the Mastodon instance network, from the macroscopic and mesoscopic perspectives (**Q2–Q3**). Section "[Backbone of the Instances network](#)" describes our methodology of identification of the backbone of the Mastodon instance network (**Q4**). Section "[Evolution of the network of Mastodon instances](#)"

analyzes the evolution of the Mastodon instance network (Q5) from three points of view: comparison with the earlier Mastodon network, emphasis on the online portion of the network, and an analysis of centrality of the instances. Finally, section "Conclusion and future work" concludes the paper and provides pointers for future research.

### Polite data crawling and network modeling

To answer our first research question (Q1), here we describe the crawling methodology adopted to collect public information from Mastodon, the steps carried out to build and validate our Mastodon instance dataset, and the network models we derived from the collected data.

#### Crawling methodology

The publicly available dataset on Mastodon relationships provided in Zignani et al. (2018) contains data extracted during the period between 2017 and 2018. This clearly raises concerns about the possibly partial obsolescence of those data, since social networks continuously evolve and there is no reason to assume that Mastodon and the Fediverse would represent an exception to this rule. Therefore, to satisfy the need for up-to-date data, we carried out an extensive crawling phase based on a newly designed crawler.

**Crawling requirements and design principles.** Our crawler was developed under strict and self-imposed constraints, i.e., following the *privacy by design*, *privacy by default* approach, and exclusively relying on the publicly-available Mastodon REST APIs<sup>2</sup>—using such APIs, we accessed data through *GET* and *POST* methods of the HTTP protocol, and managed the payload of requested data in a JSON format. Under these constraints, we were able to make our crawling methodology fully compliant with ethical and privacy-related principles.

Given the decentralized nature of Mastodon, it is not straightforward to detect the myriad of instances available today. Nonetheless, to get updated information on the current landscape of Mastodon instances, the *instances.social* website<sup>3</sup> is commonly used as a de-facto tracker of Mastodon instances. We exploited it to generate a list of seeds (i.e., starting points for the searching process), which correspond to currently online Mastodon instances.

Mastodon instances provide developers with *authentication tokens* to ensure control over the scope of the interactions. Moreover, by leveraging on authenticated requests, developers might achieve better interaction capabilities with instances. These conditions certainly comply with our desired privacy and ethical principles. Therefore, we submitted our seed list (i.e., the instances obtained from *instances.social*) to the authentication process, getting approved from approximately 900 instances out of about 1100. Also, being able to traverse instances timeline—via authenticated requests—we discovered about 81,000 new users to explore. Then, we carried out a *breadth-first-search* over them, detecting incoming and outgoing links and progressively increasing the number of users to explore, by discovering new ones during the link detection process.

---

<sup>2</sup> <https://docs.joinmastodon.org/api/>.

<sup>3</sup> <https://instances.social/>.

We point out that Mastodon allows redirecting or moving a user's profile. Although notable, this feature could determine some inconsistencies during the crawling, such as redirects to other instances while exploring the user profiles. Therefore, we avoided generating edges for users who presented similar behaviors.

Moreover, two side yet relevant remarks arise regarding our crawler implementation. First, to efficiently handle the collected data, we used a caching mechanism (*Redis*) coupled with a NoSQL database (*MongoDB*). Also, to prevent computational bottlenecks, we avoid repeated checkings over the database during the crawling phase (e.g., for checking duplicated edges), so that we eventually refine the complete network dataset in an "off-line mode", once the crawling process has finished, by exploiting particularly efficient processing functionalities provided by suitable data and network manipulation software libraries.

Our crawling session ended up with 27,989,557 links detected. After performing basic data-cleaning steps, particularly removing duplicate links, we obtained about 1.4M and 18M unique users and links, respectively, managing to cover 16,282 instances.

It should be emphasized that, to respect privacy principles, we firmly avoided using scraping techniques or systems, i.e., we abstained crawling information from instances which did not provide us with an authentication token. Notice also that the detected links were immediately anonymized, and any information that could impact on the users' privacy was replaced with numerical data generated through a proper hashing function—as a consequence, it is not possible to trace back the original information on users from our raw dataset. Finally, we point out that our fetching of descriptive text data (e.g., toots) was *minimal*, i.e., it occurred only during the initialization of our crawling process: indeed, we produced the seed-user set by discovering them through toots available in the timelines of the seed instances, relying only on authenticated requests. Nonetheless, we never stored this data since we processed it in real-time. After this initial phase, the crawling continued via breadth-first-search, thus ignoring textual data.

**Spotting Mastodon instances.** As previously mentioned, platforms in the Fediverse utilize a shared protocol, allowing for seamless interactions among their users. A related key-aspect is that, when requesting followings or followers of a Mastodon user, the APIs return all of them, regardless of the Fediverse platform. In this regard, one question becomes how we can distinguish between instances that belong to Mastodon from other platforms' instances in the Fediverse. We answered this question through a verification process, as summarized next.

To date, some relevant websites provide up-to-date Mastodon information, namely the aforementioned *instances.social* and *fediverse.party*,<sup>4</sup> so that we exploited them to filter our data through their lists of known instances. Note however that, while *fediverse.party* does not distinguish online from offline instances, *instances.social* provides fine-grained filtering capabilities. In this regard, we focused on the setting of two main parameters provided in the *instances.social* APIs: *include\_dead* and *include\_down*. As declared in *instances.social*, an instance is considered *dead* if inactive for at least two weeks, and *down* if it is not currently online yet live within a two-week window. We set either both

---

<sup>4</sup> <https://fediverse.party/en/mastodon>.



**Table 1** Current landscape of Mastodon instances as provided by *instances.social* and *fediverse.party* websites

	Instances.social	Fediverse.party	$\cup$	$\cap$
Online	1193	Not available	1193	997
Online+Offline	7313	3396	9433	1276

Symbol  $\cup$ , resp.  $\cap$ , stands for the total of Mastodon instances calculated as the size of the set union, resp. intersection, between the instances set provided by the websites

*include\_dead* and *include\_down* to *true*, or *false*, to obtain all the known Mastodon instances, resp. online-only instances.

As reported in Table 1, we merged information retrieved from both websites, regardless of the instances’ status (i.e., online or offline), obtaining 9433 known Mastodon instances. In addition, we requested online-only ones to *instances.social*, getting 1193 instances to date. Note also that, by restricting our census of the Mastodon instances to the information shared between the two platforms, the total online and the overall total would be decreased of 16.4% (997) and 86.5% (1276).

**Validated datasets.** Based on the above information, we analyzed our crawled data to properly detect the status of the instances. Results are summarized in Fig. 1.

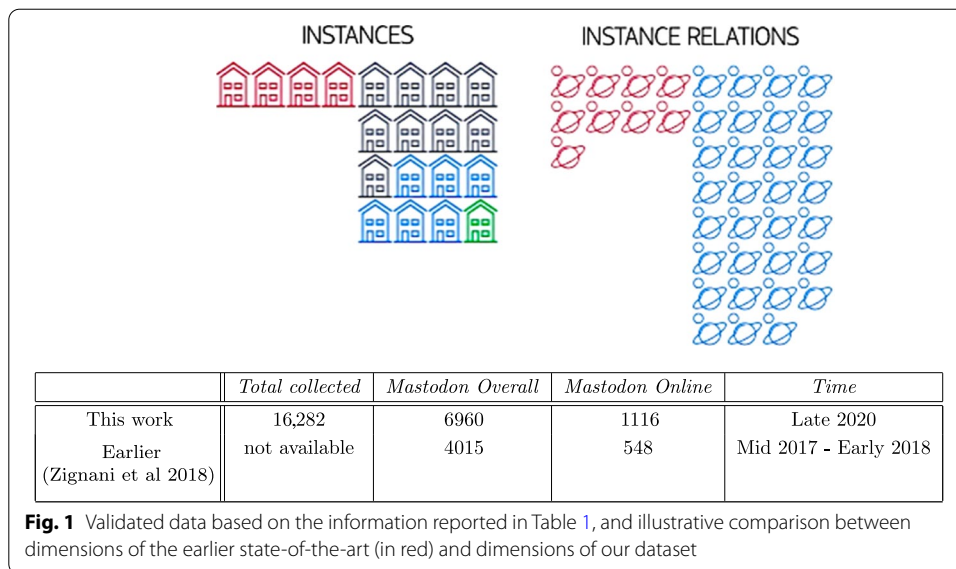
We intercepted 6960 out of 9433 Mastodon instances (both online and offline), and 1116 out of 1193 currently online instances. We point out the significance of the latter value, given the coverage of most of the online Mastodon instances to date. Moreover, our dataset doubles the earlier state-of-the-art in terms of currently online instances. Clearly, the freshness of our data (November-December 2020) influences this value.

It should be emphasized that our collected data includes a remarkable amount (9322) of non-Mastodon instances, i.e., belonging to other Fediverse platforms. This clearly strengthens the concept of Fediverse, but also opens to the discovery of the primary role taken by Mastodon within the Fediverse. In fact, although we detected non-Mastodon users and instances through Mastodon ones, and hence this knowledge of the Fediverse might be partial, our collected data offers an unprecedented opportunity for deepening our understanding of the position of Mastodon in the Fediverse (i.e., how Mastodon instances and users interact with the rest of the Fediverse), given the premises of independence yet cooperation among platforms in the Fediverse.

Further important remarks also arise regarding the growth of Mastodon. Although we are not aware of the status (i.e., online or offline) of the instances in the earlier state-of-the-art dataset (Zignani et al. 2018) at the time of their creation, we hypothesize that, after a first boost due to enthusiasm and novelty, Mastodon reached its stability as a DOSN. Indeed, the number of currently online instances is moderate compared to the number of all-time known ones and refers to non-transient yet well-rooted platforms. Overall, our dataset turns out to be significantly larger and more recent than the earlier Mastodon dataset, making it more suitable for novel and further studies.

**Network models**

Let us denote with  $\mathcal{U}$  the set of users and with  $\mathcal{I}$  the set of instances available in the extracted Mastodon data. We can define a directed network modeling the Mastodon data as  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ , where the node set  $\mathcal{V}$  contains pairs  $(u, i)$ , with  $u \in \mathcal{U}$  and  $i \in \mathcal{I}$ ,



and the edge set  $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$  corresponds to the set of following relations, such that any  $(x, y) \in \mathcal{E}$  with  $x = (u, i)$  and  $y = (v, j)$  means that user  $u$  in instance  $i$  follows user  $v$  in instance  $j$ . Note that  $u$  may coincide with  $v$  provided that  $i \neq j$ . Given  $\mathcal{G}$ , we derive three Mastodon networks at instance level, which are formally defined as follows.

*Instance network.* The first network we define is the graph modeling relations between all the instances in  $\mathcal{I}$ , hereinafter referred to as INSTANCES network, as the directed weighted network  $G_{\mathcal{I}} = \langle V, E, w \rangle$ , where  $V = \mathcal{I}$  is the set of nodes,  $E$  is the set of edges such that  $(i, j) \in E$  means that there exists at least one user in instance  $i$  that follows another user in instance  $j$ , and  $w : E \mapsto \mathcal{R}$  is an edge weighting function such that, for any  $(i, j) \in E$ ,  $w(i, j)$  stores the multiplicity of the following relation from  $i$  to  $j$  (i.e., number of users in  $i$  following users in  $j$ ).

*Online instance network.* Our second network is induced from the set of instances that are detected as online at the time of the crawling process we carried out. Therefore, by denoting with  $V^o \subseteq \mathcal{I}$  the set of online instances, the ONLINE-INSTANCES network  $G_{\mathcal{I}}^o = \langle V^o, E^o, w^o \rangle$ , with edge-set  $E^o = E \cap (V^o \times V^o)$  and edge weighting function  $w^o : E^o \mapsto \mathcal{R}$ , is defined to model the connections between the online instances only.

*Expanded network.* Our third network generalizes the first one by accounting for instances that have been recognized outside Mastodon. Actually, every link extracted during our crawling process is by definition incident with at least one instance that belongs to Mastodon. Therefore, we also define an expanded network to explore the boundary of the Mastodon network to the rest of the Fediverse. By denoting with  $V^* \supset \mathcal{I}$  such expanded set of instances, i.e., the whole set of crawled instances, we define the EXPANDED-INSTANCES network as  $G_{\mathcal{I}}^* = \langle V^*, E^*, w^* \rangle$ , where  $E^* = E \cup \{(i, j) \mid (i \in V \wedge j \in V^* \setminus V) \vee (i \in V^* \setminus V \wedge j \in V)\}$ ,

and the weighting function  $w^* : E^* \mapsto \mathcal{R}$  follows analogous definition as for the Mastodon instances network.

All the above defined networks and the one inferred from the earlier dataset are summarized in Table 2. Note that the number of nodes in the ONLINE-INSTANCES network



**Table 2** Networks created from our collected dataset, and comparison with the earlier state-of-the-art network

Network name	#Nodes	#Edges
EXPANDED-INSTANCES	16,282	318,218
INSTANCES	6960	216,504
ONLINE-INSTANCES	1115	75,046
Earlier (Zignani et al. 2018)	4015	95,221

All networks but EXPANDED refer to Mastodon-only instances

is decreased of one w.r.t. the information given in Fig. 1, since one online instance is actually disconnected from the network. Moreover, in Fig. 2, we provide an illustration of the overall network of instances created from our collected dataset, i.e., EXPANDED-INSTANCES.

Remarkably, the *Earlier* network is found to be mostly contained in our INSTANCES network; more precisely, about 80% of the instances in the *Earlier* network are also contained in our INSTANCES. As we shall discuss later in this work, this has important implications in the growth of Mastodon during the last three years.

It should also be noted that our collected data allows us in principle to build networks at the *user* level as well, e.g., we could define the network of the relations between users of each particular instance; nonetheless, this goes beyond the scope of this work, whose focus is the analysis and understanding of the relations among the instances in Mastodon. Therefore, we leave the study of instance-specific networks of users as future work (cf. section "Conclusion and future work").

### Structural analysis of the INSTANCES network


In this section, we answer our second research question (Q2) by presenting an extensive analysis of the network we built over the Mastodon instances, i.e., the previously introduced INSTANCES network. To unveil the main characteristics that define Mastodon, we will take a macroscopic as well as a mesoscopic perspective, and organize the discussion into the two next subsections.

#### Macroscopic structural analysis

We begin our investigation of the INSTANCES network at a macroscopic level. We refer to Table 3 for a summary of statistics on the main structural characteristics of the INSTANCES network, each of which is analyzed in the following.

**Degree distribution.** Figure 3 shows the boxplot, density function, and Complementary Cumulative Distribution Function (CCDF) of the INSTANCES network in-degrees, with various types of distribution fittings; results obtained for the out-degrees and the total degrees are analogous, and we report them in the Appendix (Figs. 9 and 10).

Looking at the boxplot and the histogram with associated density function, there is evidence of right-skewness of the in-degree distribution, with a small bunch of "outliers" scaling in the regime of thousands. In detail, while the first quartile, median, mean, third quartile, and non-outlier maximum degree are 91.5, 202, 331, 413, and 878, respectively, there are 21 instances having in-degree above 880, up to a maximum degree of 4685.



**Fig. 2** Illustration of the EXPANDED-INSTANCES network. Node colors denote different types of instances: white and light blue indicate online and offline Mastodon instances, respectively, whereas dark blue corresponds to non-Mastodon instances. Node size is proportional to its degree. The displayed layout is based on the force-directed drawing ForceAtlas2 model. (Produced by using the Graphistry service, available at <https://www.graphistry.com/>.)

We investigated about the outlier instances, focusing on the top-5 by degree, resp. in-degree and out-degree, as shown in Fig. 4. We found these correspond to *mastodon.social*, *pawoo.net*, *mastodon.xyz*, *octodon.social*, and *mstdn.io*, in all the three cases. These are clearly among the most popular instances in Mastodon and, as expected, they allow their users to discuss a large variety of topics. Interestingly, *mastodon.social* is always the top-1 instance regardless of the type of degree, whereas *pawoo.net* and *mastodon.xyz* alternate each other at the second and third rank.

Such instances are also well-recognized in the CCDF plot, where we observe a probability of 50% of having at least 200 in-degree, which already drops to 20% for an in-degree around 600, and further decreases below 4% for the outliers. The CCDF plot also displays the best fitting of power-law, lognormal, exponential and Poisson distribution to the observed data. The resulting fitting curves appear to provide indications of lognormality and, to a limited extent, of power-law fitting.

The above prompted us to assess the corresponding statistical significance, whereby we resorted to a Kolmogorov-Smirnov test. Results are summarized in Table 4. In the first subtable, the high  $p$ -values suggest that the null hypothesis that the data are from a power-law distribution cannot be rejected, although this holds on a limited regime ( $x_{\min}$ ) starting from degree values, resp. in-degree and out-degree values, of the order of hundreds. In particular, for the in-degree case, note that  $x_{\min}$  is above the mean of the distribution. The remaining subtables in Table 4 correspond to four different

**Table 3** Summary of structural characteristics of the INSTANCES network, including details on community structure and core decomposition

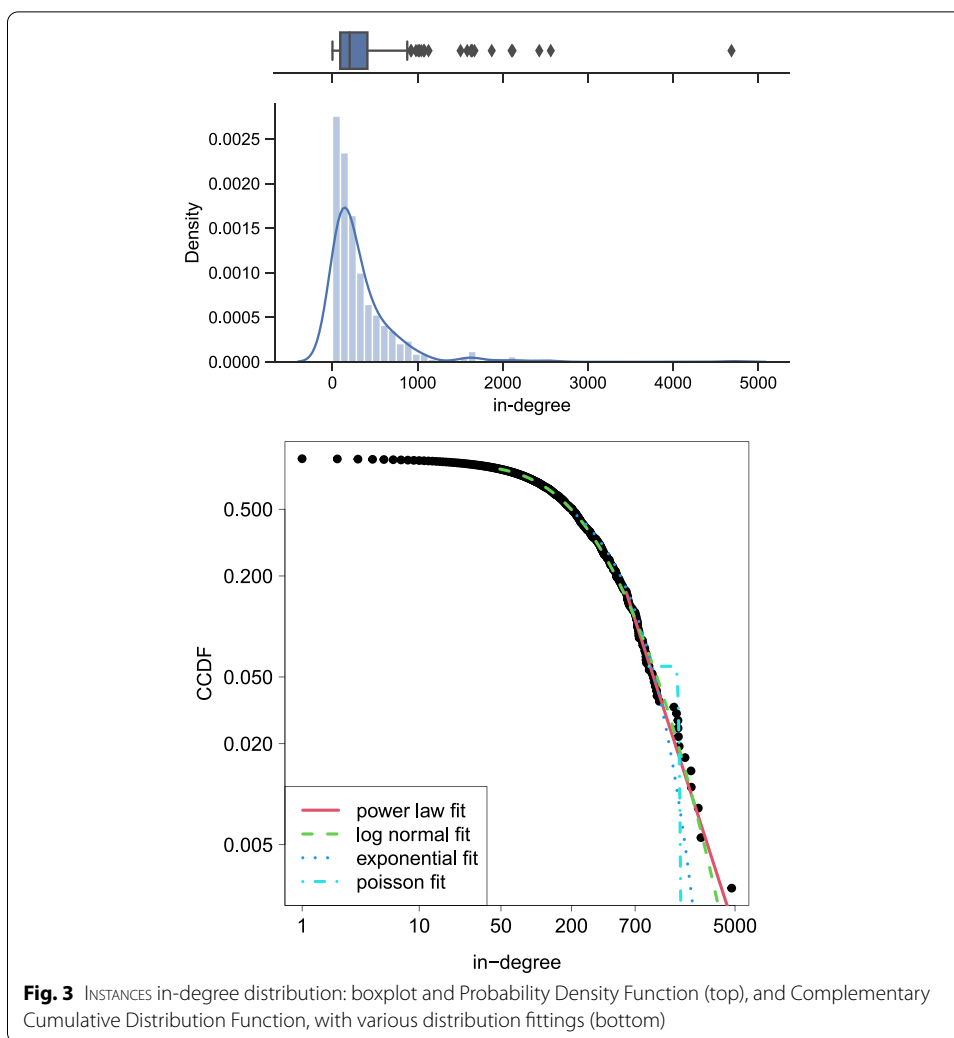
	INSTANCES	INSTANCES inner-most core		
		Degree	In-degree	Out-degree
#Nodes	6960	189	208	196
#Edges	216,504	25,790	28,690	26,463
Reciprocity	65.1%	88.4%	85.7%	88.2%
Density	0.004	0.726	0.666	0.692
Average degree*	41.966	152.328	157.702	150.98
Average in-degree	31.107	136.455	137.933	135.015
% Sources	12%	0%	0%	0%
% Sinks	6.6%	0%	0.005%	0%
Degree assortativity*	- 0.274	- 0.117	- 0.158	- 0.135
Degree assortativity	- 0.253	- 0.14	- 0.171	- 0.151
Average path length	2.330	1.270	1.330	1.310
Diameter	5	2	2	2
Transitivity*	0.128	0.832	0.798	0.807
Clustering coefficient*	0.836	0.837	0.810	0.816
Clustering coefficient (full averaging)*	0.687	0.837	0.810	0.816
#Strongly connected components	1305	1	2	1
#Weakly connected components*	1	1	1	1
Modularity by Louvain*	0.289	0.032	0.039	0.037
#Communities by Louvain*	5 (5)	3 (3)	3 (3)	3 (3)
Modularity by Louvain**	0.353	0.242	0.246	0.246
#Communities by Louvain**	6 (8)	4 (5)	3 (4)	4 (6)
#Communities by Infomap**	6 (54)	1 (3)	1 (4)	1 (3)

\*Statistic calculated by discarding the edge orientation

\*\*Statistic calculated by taking into account the edge weights

scenarios we investigated for the lognormality fitting, namely (from top to bottom in the table) full regime (i.e., whole observed data), removal of the outliers, removal of lower-degree ( $\leq 50$ ) instances, removal of both outliers and lower-degree instances. As it can be noted, the Kolmogorov-Smirnov test yielded high significance values for the lognormality fitting in all cases, except when the outliers only are discarded; particularly, the significance is maximized when the lower-degree instances are removed (i.e.,  $p$ -values from 0.687 to above 0.9). In such cases, the test informs us that we cannot reject the null hypothesis and so we conclude that the observed data are lognormally distributed.


**Sources and sinks.** We inspected the presence of instances having no incoming links (i.e., sources) as well as of instances having no outgoing links (i.e., sinks). As reported in Table 3, the percentage of both types of instances is not negligible, with the incidence of sources being nearly double than sinks. This might provide clues for the presence of users belonging to small instances (e.g., private ones) interested in contents produced by users located in other instances; indeed, since small instances would host few users, the lack of incoming links is plausible. On the opposite side, the percentage of sink instances sheds light on that they might contain well-consolidated



user groups, among which there is no need to interact with users belonging to other instances. This peculiarity might suggest sort of self-sufficiency in the federation.

**Triadic closure.** We analyze how well the triadic closure principle is met in the INSTANCES network, by looking at both transitivity (i.e., the probability that two incident edges are completed by a third one to form a triangle) and local clustering coefficient (i.e., how strongly connected are the neighbors of a node).

We observe a rather low value of transitivity (0.128), which is actually not surprising given the relatively low density of the network. By contrast, local clustering coefficient is very high (0.836), and remains as such even when accounting for sink or source instances (0.687). This evidence is remarkable as it hints at a federative structure among the instances. Note also that the dichotomy between a relatively lower transitivity and a higher local clustering coefficient, and more in general, the low correlation between the two statistics is typical of networks characterized by a skewed degree distribution. In this respect, the Mastodon INSTANCES network also keeps this feature.



As a further remark on length-2 closed loops (i.e., reciprocal edges), we observe a high fraction of reciprocal edges (above 65%). As we shall further observe through our core decomposition analysis, reciprocity tends not to be limited to users within the same instance, but involves instances that can be placed very differently, from the periphery to the internal of the network, and vice versa.

**Degree assortativity.** One key structural property of a network at macroscopic level refers to degree correlation, or degree assortativity, which measures how the probability of a link between two nodes depends on their degrees (Newman 2002, 2003). Real-world social networks are often found to have positive degree assortativity, i.e., well-connected individuals are linked to other well-connected ones. A recent study has also shown that this evidence does actually hold for those social networks built upon shared memberships of group (Fisher et al. 2017).

Remarkably, the Mastodon INSTANCES network exhibits a degree assortativity which is significantly negative (−0.253), which means that well-connected instances are connected to many instances with few other connections. This might be ascribed to the heterogeneous degrees characterizing the instances in the network. In this respect, we argue that, since instances can be bounded to specific topics, users belonging to different instances with diversified degrees tend to interact with each other to reach a broader range of contents, thus ultimately improving their experience on the platform and increasing the speed of information transfer.

The degree disassortativity, i.e., negative degree correlation, exhibited by the Mastodon INSTANCES network outlines a novelty w.r.t. well-known centralized social networks. It should be noted that, unlike centralized social networks, Mastodon users’ behavior is not impacted by recommendation mechanisms. Therefore, the followships tend to be built upon the topical interests and preferences that users have, which leads to a form of topically-induced link formation rather than a popularity-based attachment. Moreover, this further supports the strong interrelation between instances which, as we shall discuss in the next section, characterizes peripheral and inner-core locations in Mastodon.

### Mesoscopic structural analysis

We organize our presentation of the mesoscopic structural analysis of the INSTANCES network into two parts: the first one is devoted to the evaluation of the community structure that is detected over the network, whereas the second part is concerned with the core decomposition of the network.

**Table 4** Power-law and lognormal fittings through *Kolmogorov-Smirnov* test performed on the INSTANCES network

	Degree	In-degree	Out-degree
<i>Power law</i>			
$x_{min}$	890	588	457
$\alpha$	2.987	3.166	3.057
<i>p</i> -value	0.939	0.889	0.936
<i>Lognormal</i>			
Interval	[1, 9062]	[1, 4685]	[1, 4377]
$\mu, \sigma$	5.54, 1.26	5.18, 1.24	5.11, 1.19
<i>p</i> -value	0.113	0.187	0.113
Interval	[1, 1307]	[1, 878]	[1, 742]
$\mu, \sigma$	5.37, 1.15	5.05, 1.15	4.96, 1.09
<i>p</i> -value	0.01	0.054	0.024
Interval	[51, 9062]	[51, 4685]	[51, 4377]
$\mu, \sigma$	5.82, 0.95	5.53, 0.87	5.45, 0.81
<i>p</i> -value	0.963	0.687	0.932
Interval	[51, 1307]	[51, 878]	[51, 742]
$\mu, \sigma$	5.65, 0.80	5.40, 0.74	5.31, 0.67
<i>p</i> -value	0.47	0.646	0.816

### Community detection

We resorted to two well-known community detection algorithms for discovering communities in the Mastodon network of instances, i.e., the Louvain method (Blondel et al. 2008) and the Infomap method (Rosvall and Bergstrom 2008). Louvain is a two-step, hierarchical greedy optimization method that attempts to maximize the modularity of a partition of the network, whereas Infomap optimizes the Map equation, which exploits the information-theoretic duality between finding community structure in networks and minimizing the description length of a random walker's movements on a network.

It should be noted that both methods have been used with success for networks of many different types and sizes, and today, they are the most widely used methods for detecting communities in large networks.

For the Louvain algorithm, we exploited both the original, undirected implementation as well as the directed variant.<sup>5</sup> In the latter case, we also took into account the edge weights. As regards Infomap, we exploited its weighted directed implementation.<sup>6</sup>

The number of communities found by the aforementioned algorithms is shown in Table 3, for all considered scenarios (cf. notes marked with \* and \*\* below the table). Please note that we report two values for each case: the one within parenthesis corresponding to the total number of communities while the first value refers to the number of communities that contain at least ten instances. Furthermore, since the Louvain algorithm optimizes modularity, we also report the modularity values corresponding to the community structures discovered by Louvain; in this regard, we find evidence

<sup>5</sup> <https://github.com/nicolasdugue/DirectedLouvain>.

<sup>6</sup> <https://www.mapequation.org/infomap/>.



of modular structure within the INSTANCES network, with modularity from about 0.29 (undirected network) to 0.35 (weighted directed network).

The number of communities produced by the Louvain method ranges from 5 for the undirected scenario, to 8 when applying the weighted directed variant; in the latter case, only 2 out of 8 communities contain less than ten instances. Infomap, conversely, appears to detect a much higher total number of communities; however, by inspecting them, we found out that a large majority are poorly significant as they consist of less than ten instances. Indeed, as reported in the table, we point out that the two methods actually behave the same in terms of number of relatively large communities (i.e., 6 communities produced by either method considering weights and orientation of edges).

Figure 5 illustrates the communities detected by the directed Louvain method on the INSTANCES network. As it can be observed, two main communities arise in terms of size (displayed at the bottom and on the right in the figure), which together contain about 72% of the instances in the network. Other two communities (on the top in the figure) also stand out, as they contain about 21% of the instances. Remarkably, albeit less evident for the smallest communities, we can observe high connectivity between all of them, which hints at the high interrelation between instances belonging to different communities.

We also delved into the community structure obtained via the Louvain algorithm to get more insights into the community boundaries.

The largest community (i.e., the rightmost one in Fig. 5) contains the most relevant Mastodon instance in the Fediverse, namely *mastodon.social*, which represents the first instance born with the Mastodon project. Consequently, given its role as a reference point in Mastodon, the large constellation of instances observed around it is not surprising. Within the same community, we spotted *mstdn.io* and *octodon.social*. The above three instances are not topically bounded and use English as the primary language, with *mstdn.io* also embracing French and *octodon.social* extending its range of languages to Japanese and Portuguese. We point out that these instances might share the same community given their relevance and longevity (e.g., *mstdn.io* is up since early 2017) in the Fediverse.

Moving our focus to the second-largest community in size (i.e., the one at the bottom of the figure), we distinguished two relevant instances, namely *pawoo.net* and *mstdn.jp*. Their co-existence within the same community is justified by the fact that both discuss various topics and share the official language (i.e., Japanese).

Furthermore, in the third-largest community, we located the remaining instance, i.e., *mastodon.xyz*, of the previously mentioned top-5 largest ones in Mastodon. This is general-purpose and uses English and French as primary languages.

As a final remark, we point out that all the top-5 instances reported in section "Macroscopic structural analysis" are established in the largest three communities discovered by the Louvain algorithm, and hence their relevance is further strengthened due to their central role within these communities.

We replicated the same explorative analysis also for the community structure detected by Infomap. Remarkably, mostly interesting patterns are found again. Indeed, the largest community still includes the same instances as in the case of Louvain, with the addition of *mastodon.xyz*, which in the Louvain solution is included in the third largest

community. Moreover, this similarity aspect also holds for the second largest community detected by Louvain and Infomap, respectively, which comprises the two Japanese instances. These coherent findings from two different community detection methods would support an underlying logics in the organization of the instances and their interrelations.

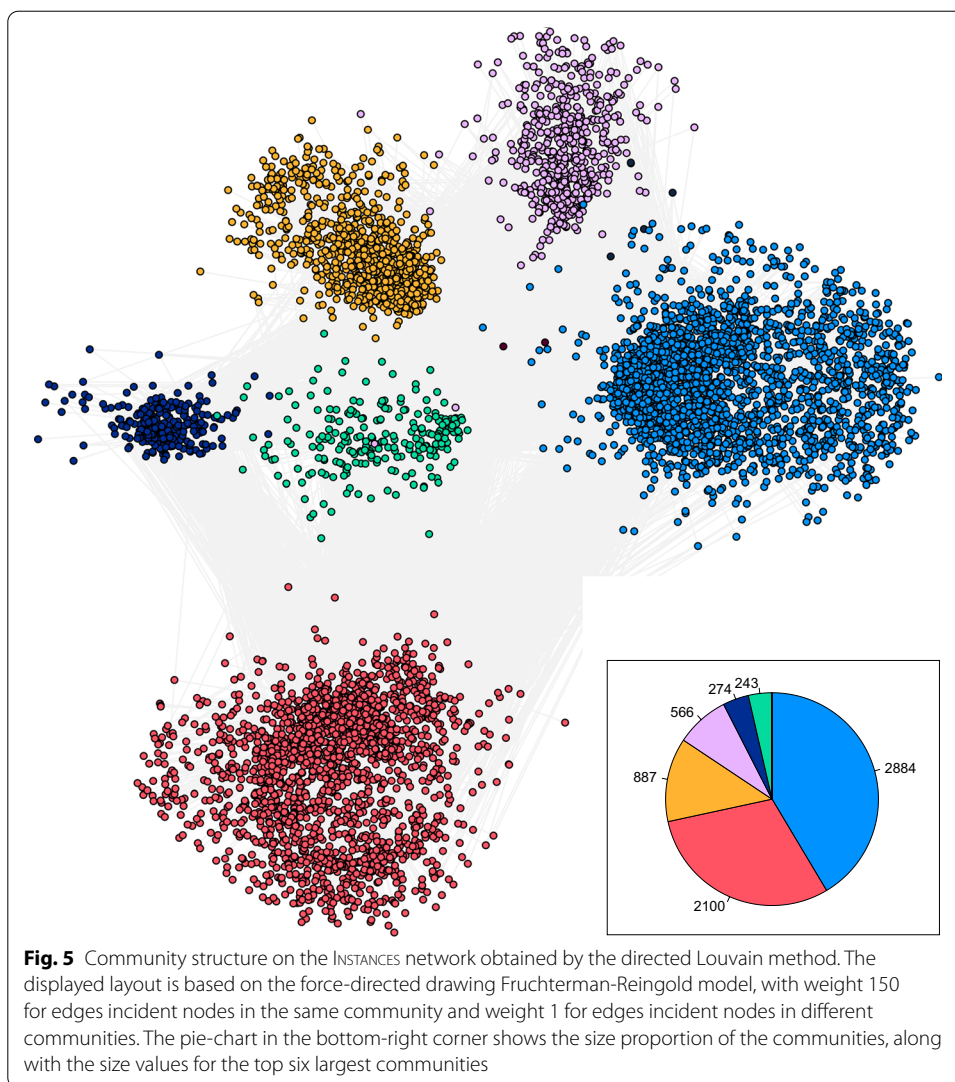
We further investigated the community structures produced by Louvain and Infomap according to the aspect of *conductance*. This is defined as the ratio between the cut size among two communities (i.e., the sum of the weights of the edges that link two sets of nodes) and the smaller of the volumes (i.e., the sum of the degrees of the nodes in a set) of the two communities; we considered both the weighted and unweighted versions of conductance, where for the latter, the edge weights are equal to one.

As reported in Table 5, conductance varies depending on the community detection approach. Both methods induce good separation among communities in the INSTANCES network, which is indicated by the low values of conductance, but in the Infomap solution this is much more evident than in the Louvain community structure. Taking a finer-grain perspective with a focus on the three largest communities produced by the two methods, respectively, the analysis of the pairwise conductance among such communities yields two main outcomes: the conductance between the largest and the third largest communities ranges from about 0.56 (0.72, for the weighted version) w.r.t. Infomap to above 0.8 when using Louvain; by contrast, regardless of the community detection method, both unweighted and weighted conductance is much lower for the pairs involving the second largest community. We tend to ascribe these differences in conductance since, on the one hand, the largest communities have significant inter-community communication flow (and hence, high cut size) due to the involvement of most relevant instances over the network; on the other hand, this inter-community connection may be limited due to different cultures and languages, as it is the case for the second largest community which indeed is centered around the Japanese-language *pawoo.net* and *mstdn.jp* instances.

The above remarks also prompted us to measure dependencies between degrees of neighbor nodes in the communities, by computing the *average nearest neighbor degree* distributions. Figure 6 reports the values of the average nearest neighbor degree as a function of the degree  $k$ , which we denote by  $knn(k)$ . As it can be noted, the decreasing trend by  $k$  is not only clear for the whole INSTANCES network—which is indeed aligned with the negative degree assortativity previously analyzed—but also for the subnetworks induced from the top three communities exhibiting disassortative traits.

### **Core decomposition**

The core decomposition of a network graph consists in assigning each node with an integer number (the core index) capturing how well the node is connected with respect to its neighbors. The result is a threshold-based hierarchical decomposition of the graph into nested subgraphs, based on a threshold ( $k$ ) which is set on the degree of nodes (Seidman 1983). The identification of such tightly-knit substructures, or cores, has long been used for understanding mesoscale structural characteristics of a network, with several



applications related to the computation of the local importance of nodes (Malliaros et al. 2020), including the estimation of the spreading potential of nodes (Kitsak et al. 2010; Caliò et al. 2020). A key advantage of core decomposition lays on theoretically grounded definition and uniqueness of its solution, which can also be computed efficiently in linear time w.r.t. the number of edges in the input graph.

Given a graph  $G = \langle V, E \rangle$  and any subset  $S \subset V$ , let us denote with  $G[S] = \langle S, E[S] \rangle$  the subgraph of  $G$  induced by  $S$ , where  $E[S] = E \cap (S \times S)$ .

For any choice of an integer value  $k \geq 0$ , the  $k$ -core of a graph is the maximal induced subgraph  $G[C_k] = \langle C_k, E[C_k] \rangle$  such that the number of neighbors of every node  $v$  in  $C_k$  is at least  $k$ . The *degeneracy*  $K$  of the graph is the highest value of  $k$  such that  $C_k \neq \emptyset$ . The core associated with the graph degeneracy is also called the *inner most core*. The set of all  $k$ -cores (i.e.,  $V = C_0 \supseteq C_1 \supseteq \dots \supseteq C_K$ ) represents the core decomposition of the graph. Moreover, the *core-index*, or *coreness*, of a node  $v$  is the largest  $k$  such that  $v \in C_k$  and  $v \notin C_{k+1}$ .

**Table 5** Unweighted and weighted *conductance* scores for the community structures obtained by Louvain and Infomap methods on the INSTANCES network. The first row refers to the average over all pairwise scores between communities, whereas the other rows refer to comparisons between the three largest communities detected by the methods. Weighted variants of conductance account for edges weight when calculating volumes and cuts

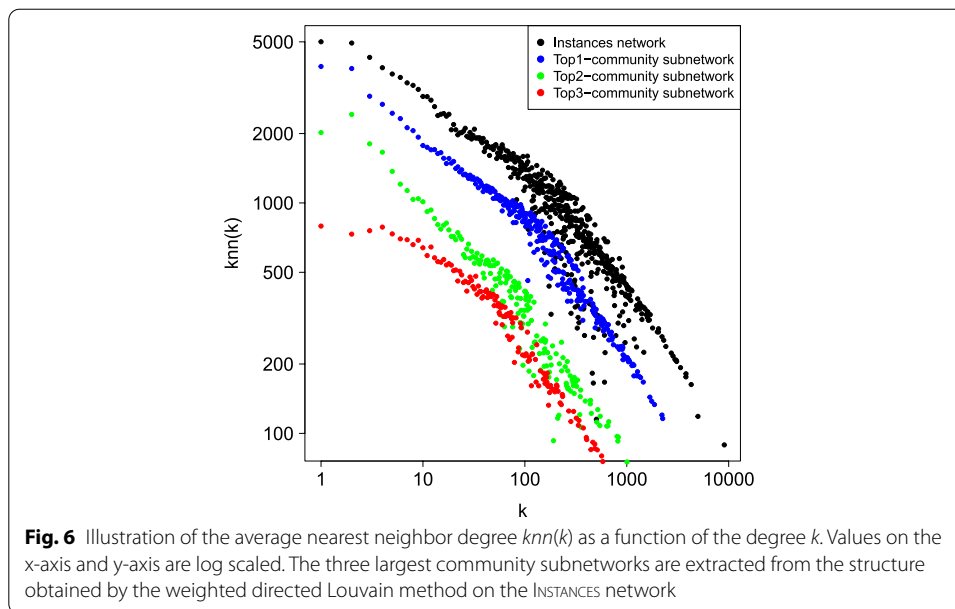
	Louvain		Infomap	
	Unweighted	Weighted	Unweighted	Weighted
INSTANCES	0.285	0.239	0.037	0.035
Top-1 vs. Top-2	0.315	0.191	0.370	0.219
Top-1 vs. Top-3	0.816	0.881	0.561	0.721
Top-2 vs. Top-3	0.100	0.059	0.231	0.335

Note also that the above definitions originally apply to undirected graphs, however they are straightforwardly adapted to directed graphs so that the degree of a node may refer to either its in-degree or out-degree.

The core decomposition of the INSTANCES network revealed further hints at the presence of a federative mechanism. As reported in Table 3, one major finding is the remarkable number of instances in the inner cores of the network, and even in the inner-most core, which ranges from 189 based on total degree to 208 based on in-degree, and also includes the top-5 instances previously discussed in section "[Macroscopic structural analysis](#)". This prompted us to explore the inner-most core of the INSTANCES network, in its three variants, i.e., via a total-degree-based, in-degree-based, and out-degree-based decomposition of the network.

A few remarks arise from the observation of the inner-most cores, such as the high degeneracy under all the considered scenarios: indeed, we found out the values 201, 96, and 97 of degeneracy corresponding to a core decomposition based on total degree, in-degree, and out-degree, respectively. As expected, density is high, and so is also the reciprocity—note that statistics still refer to directed subgraphs, although the core-decomposition may have been generated based on the undirected (i.e., total degree-based) definition. Although less evident, the disassortative trait is still present, denoting the coexistence in the inner-most core of instances exhibiting diversified yet high degrees. Given the high density, average path length and diameter almost halve compared to their original values. According to a more cohesive structure, transitivity significantly increases (nearly seven times the value in the whole network) and aligns with the local clustering coefficient. Also, the number of strongly connected components (i.e., #SCCs) heavily shrinks to 1, resp. 2, for the indegree-based, resp. degree- and out-degree-based decomposition.

We also evaluated the presence of a community structure in the three variants of inner-most core. The (undirected) Louvain method determines an exceptional reduction in modularity compared to the value found for the whole network (i.e., from nearly 0.3 to 0.03-0.04), along with a small decrease in the number of significant communities (from 5 to 3). By contrast, when considering edge orientation and weights, we observe a slight decrease in modularity (from about 0.35 to 0.25) and community structure size (from 8 to 4-6). Interestingly, the number of communities detected by Infomap strongly



decreases compared to the value observed for the whole network, with one giant community and the remaining 3-4 containing less than ten instances.

Our major focus in this analysis was about whether there exists a significant amount of connections between the periphery and the inner parts of the INSTANCES network. To this purpose, we started with a visual exploration of the cores of the network, whose results are shown in Fig. 7 for the indegree-based decomposition. Looking at the figure, the display of the cores unveils evidence of directed links with a relatively high weight (i.e., solid lines in the chart), which correspond to followships from users of instances in the inner-most core to users of more peripheral instances, but also followships in the opposite direction, even coming from the outer-most cores towards the inner-most core. This certainly relates to the negative degree assortativity previously emphasized, and it represents quite a novel pattern in social networks, which usually do not show direct links between core and peripheral nodes.

We further delved into this trait by examining the number of incoming links and outgoing links for each core-index, as depicted in Fig. 8. Lower core-index values (up to about 20) correspond to a number of outgoing links that is significantly higher than the number of incoming links, which is expected (i.e., instances within these cores tend to interact with more relevant ones, while the opposite occurs less, or not at all). However, as the core-index value increases, unveiling an interesting feature of bidirectionality in the followships behavior for mid and internal cores. Further interestingly, some instances with high core-index (i.e., from about 70) show a prevalence of outgoing links. A final, remarkable trait of the INSTANCES network is that a major portion of connections (i.e., about 190,000, which is more than 87% of the total edge set size) depart from and arrive to the inner-most core.

## Discussion

Our analysis of the structural characteristics of the INSTANCES network has revealed noteworthy features, which might set Mastodon apart from well-known centralized social networks. Here, we summarize such features, thus answering our third research question (Q3) on the “fingerprint” of the network of Mastodon instances.

The evaluation at the macroscopic level has unveiled various unique traits of Mastodon. Interesting aspects are already present in the degree distribution of the instances: indeed, as opposed to well-known centralized social networks that tends to fit a power-law distribution, Mastodon reveals a better and extensive fitting with a lognormal distribution, along with the presence of few instances that show a degree up to one order of magnitude higher than the average degree in the network.

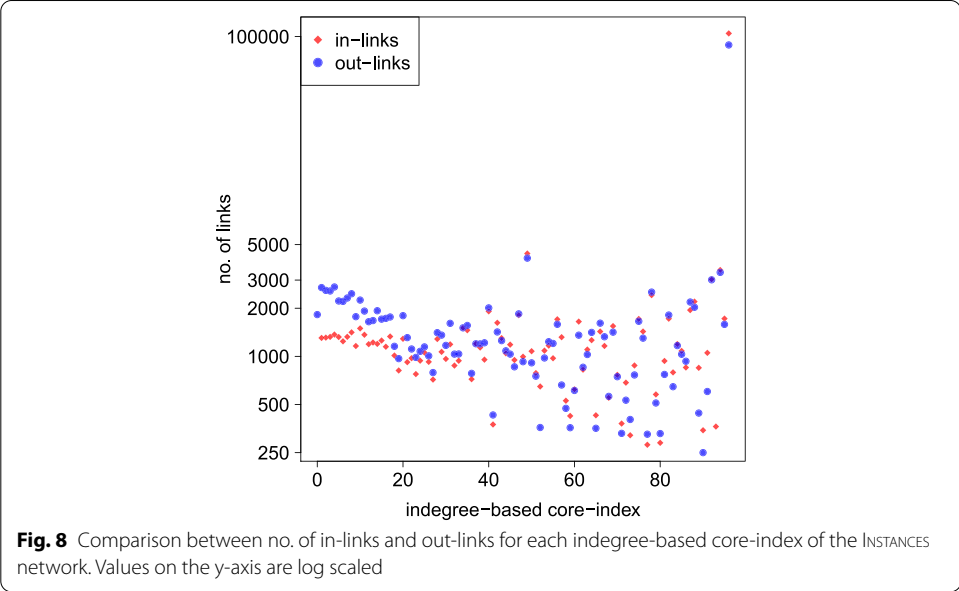
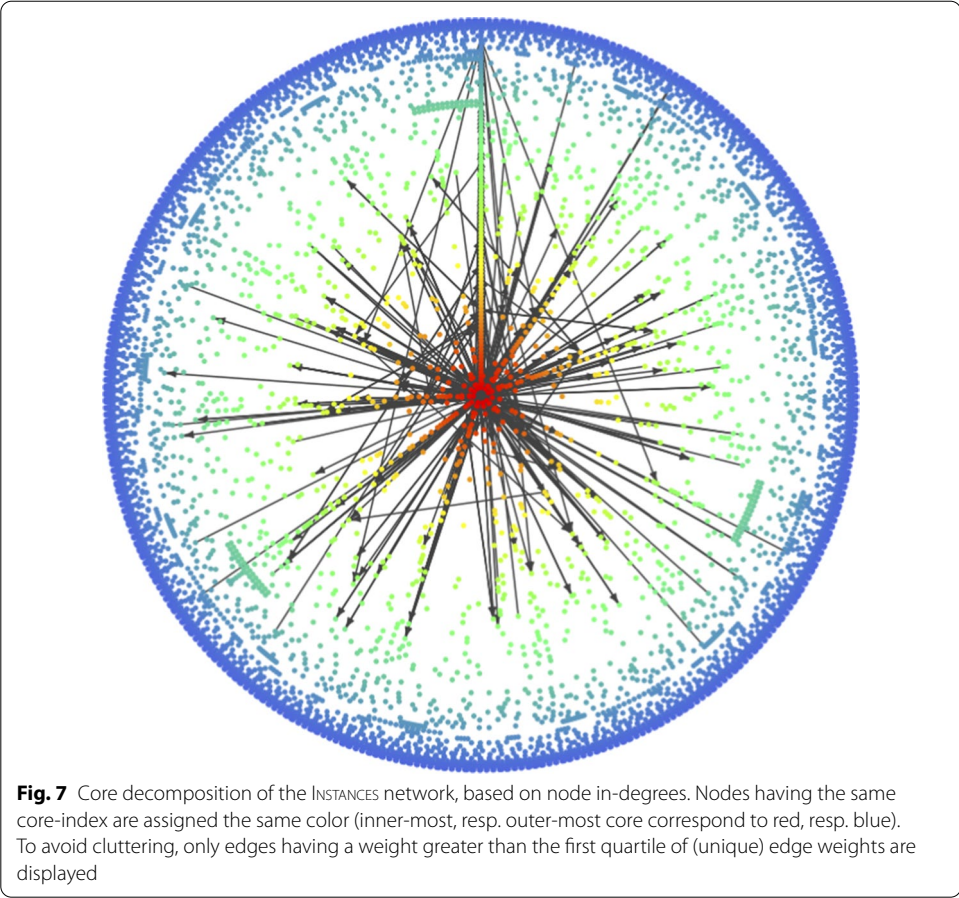
Referring to the federation concept as a set of independent yet cooperating instances, we have found a number of aspects (e.g., high clustering coefficient values and percentage of reciprocal edges) indicating that Mastodon adopts a federative mechanism. We mark this mechanism as a *mutual reinforcement* to reduce the sectorization bias that can characterize the individual instances: indeed, they might be topically-bounded due to the decentralized nature of the platform, however, their users generally look for a broader spectrum of topics, and hence they interact across different instances. We point out that, besides distinguishing Mastodon from other social networks, this trait is central to the platform itself, thanks to the shared protocol (i.e., ActivityPub) among instances.

Related to this mutual reinforcement is the observed negative degree correlation (i.e., degree disassortativity), which represents another distinctive Mastodon feature. This indicates that users belonging to different instances with heterogeneous degrees tend to interact with each other, aiming at achieving a better user experience and increasing the speed of information transfer.

The evaluation at the mesoscopic level has also highlighted traits that contribute to determine the fingerprint of Mastodon. Through community detection, we shed light on the modular structures within instances, which provide further clues to the existence of a cohesive and federative framework among them. Further investigations have revealed how these modules might be composed and influenced. We observed topics, languages and temporal processes (e.g., their creation time) as main influence factors. Moreover, through core decomposition, we spotted an unusual and conspicuous number of connections from the inner cores to the peripheral ones (Fig. 7), with a peculiar bidirectional balance between links observed starting from intermediate core-index values. Additionally, we observed that the majority of links between instances involve the inner-most core.

In conclusion, we can state that Mastodon reveals clear distinctive signs compared to what is commonly observed for other social networks, making it unique in multiple aspects. Among these traits, we detected the logical emergence of a federative mechanism, which allows independent instances to cooperate with each other, thus connecting their respective users.





### Backbone of the INSTANCES network

We focus here on a *network simplification* task which is designed to “prune” the network graph, i.e., to detect and remove irrelevant or spurious edges with the purpose of making it easier unveiling hidden substructures of a network. One simple solution to the above problem is to exploit information on the edge weights so to remove all edges having weight below a pre-determined, global threshold. Unfortunately, besides the difficulty of choosing a proper threshold for any input network, this approach tends to remove all ties that are weak at network level, thus discarding local properties at node level.

By contrast, a theoretically well-founded approach to filtering out noisy edges from a network is based on *generative null models*. The general idea is to define a null model based on node distribution properties, use it to compute a  $p$ -value for every edge (i.e., to determine the statistical significance of properties assigned to edges from a given distribution), and finally filter out all edges having  $p$ -value above a chosen significance level. In other terms, this allows to maintain only those edges that are least likely to have occurred due to random chance, hereinafter referred to as *backbone* of the network. Clearly, imposing lower significance level will yield to more restrictive substructures, thus giving place to a potential hierarchy of backbones.

Identifying the backbone of a network graph allows us to isolate the latent structure of the network under analysis, based on the removal of statistically insignificant edges. In this regard, statistical models for graph pruning have been conceived to deal with weighted networks, so that the node degree and/or the node strength are used to generate a model that defines a random set of graphs resembling the observed network.

One of the earliest methods is the *disparity* filter (Serrano et al. 2009), which evaluates the strength and degree of each node locally. The null hypothesis is that the strength of a node is redistributed uniformly at random over the node’s incident edges.

Unlike disparity, the null model proposed by Dianati (2016) is maximum-entropy based and hence unbiased. Upon it, two models are defined: the *marginal likelihood filter (MLF)*, which is a linear-cost method that assigns a significance score to each edge based on the marginal distribution of edge weights, and the *global likelihood filter*, which accounts for the correlations among edges. While performing similarly, the latter is more costly than the former, therefore we will consider the MLF model in our analysis.

### Details on the disparity and MLF models for weighted directed networks

Both disparity and MLF were originally conceived for undirected networks, but they can easily be extended to weighted directed networks as well. In the following, we provide a brief review of the two models, and refer the interested reader to the original works for further details (Serrano et al. 2009; Dianati 2016).

In the weighted directed network scenario, let us denote with  $k_i^{in}$  and  $k_i^{out}$  the in-degree and out-degree, respectively, for any node  $v_i$ . Also, the total strength  $s_i$  associated to the node has two contributions, namely the incoming strength  $s_i^{in}$  and the outgoing strength  $s_i^{out}$ , which are obtained by summing up all the weights of the incoming or outgoing links, respectively.

The disparity model aims to preserve the edges carrying a weight that represents a local significant deviation with respect to a statistical null model for the local assignment of weights by using the disparity function. Moreover, since in-degree and out-degree

distributions are assumed not to be necessarily correlated, incoming and outgoing links associated to a node are considered separately.

To assess the effect of inhomogeneities in the weights at the local level, the following functions are defined for any node  $v_i$ :

$$\begin{aligned} \Upsilon^{out}(v_i) &= k_i^{out} \sum_j (p_{ij}^{out})^2, \\ \Upsilon^{in}(v_i) &= k_i^{in} \sum_j (p_{ji}^{in})^2, \end{aligned}$$

where the summation term in the first, resp. second, expression characterizes the level of local heterogeneity in the outgoing edge weights, resp. incoming edge weights.

The null model for the disparity filter requires for each incoming, resp. outgoing, link of  $v_i$  with weight  $w$ , the calculation of the  $p$ -value  $\alpha_{ji}^{in}$ , resp.  $\alpha_{ij}^{out}$ , that its normalized weight  $p_{ji}^{in} = w/s_j^{in}$ , resp.  $p_{ij}^{out} = w/s_i^{out}$ , is compatible with the null hypothesis. The incoming and outgoing links that carry weights that can be considered not compatible with a random distribution, can be filtered out according to a chosen significance level  $\alpha$ ; more specifically, all the incoming, resp. outgoing, links with  $\alpha_{ji}^{in} < \alpha$ , resp.  $\alpha_{ij}^{out} < \alpha$  reject the null hypothesis and hence can be considered as significant heterogeneities.

Let  $T$  denote the number of unit edges, where each weighted edge is seen as multiple edges of unit weight. In the MLF model, the null model assigns each edge to nodes  $v_i$  and  $v_j$  which are selected independently and randomly with probabilities proportional to their outgoing and incoming strengths, respectively. More in detail, the probability that an edge from  $v_i$  to  $v_j$  is associated to them is given by  $p_{ij} = \frac{s_i^{out} s_j^{in}}{T^2}$ .

The probability that a weight  $w$  out of  $T$  unit edges will choose nodes  $v_i$  and  $v_j$  as their endpoints is given by a binomial distribution with probability of success  $w$  and number of trials  $T$ . Therefore, the null model defines for every couple of nodes a probability for their edge weight,  $\sigma_{ij}$ , depending on their strengths.

For each directed edge  $(v_i, v_j)$  with weight  $w_{ij}$ , its  $p$ -value, denoted as  $\gamma_{ij}$ , is computed as:

$$\gamma_{ij} = \sum_{w \geq w_{ij}} \Pr[\sigma_{ij} = w | s_i^{out}, s_j^{in}, T] = \sum_{w \geq w_{ij}} \binom{T}{w} p_{ij}^w (1 - p_{ij})^{T-w}.$$

According to this null model, the higher the strengths of two nodes, the higher the weight of an edge connecting them is to be in order to be considered statistically significant. Conversely, the lower the strengths of the two nodes, the lower the weight of a linking edge must be in order to be retained by the filter.

### Results on the pruned networks

Table 6 reports on main structural characteristics of the INSTANCES network after graph pruning processes through the application of the MLF and disparity models (for the sake of presentation, values in the original, i.e., unpruned network are also reported).

Using MLF, 55.1% and 68.5% of the edges are removed with significance level of 0.05 and 0.01, respectively. The pruning effect on the network size is much more evident when

using the disparity model, with at least 43.3% of nodes and 91.1% of edges removed. By contrast, the reciprocity percentage in the MLF pruned networks is about 7% less than in the original network, whereas the disparity pruned networks show a small increase of reciprocal edges (about 3%).

Interestingly, while it can be observed a significant decrease in the average in-degree—from around 31 in the original, unpruned network up to 11–15, resp. about 4–5, using MLF and disparity, respectively—the degree assortativity remains negative in all cases, with a small decrease in module under MLF, and comparable or increased in module for disparity with significance level 0.05 and 0.01, respectively. Average path length is also comparable, yet with some fluctuations, to its original value after applying disparity, while it marginally increases under MLF. The latter determines an increase in the network diameter by one and two units for significance level 0.05 and 0.01, respectively, while disparity raises it by one unit only under 0.01.

We also observe a general reduction of the transitivity, which becomes nearly one-half by MLF and even lower by disparity model, compared to its original value. The local clustering coefficient remains substantially unchanged under disparity, while it is nearly halved under MLF; however, when excluding source and sink nodes, the effect is similar for both pruning models and the gap from the original value is smaller.

Both pruning approaches lead to a relatively small decrease in the number of strongly connected components, which is slightly more evident under the disparity pruning. Looking at the community structures, the impact of the two pruning models varies depending on the specific community detection method. As concerns Louvain, the number of communities has little variations, while the pruning is always beneficial according to an improved modularity. On the other hand, the number of communities produced by Infomap almost doubles in the MLF pruned network, which suggests that the pruning effect due to the MLF model would favor the detection of many small communities by Infomap.

## Discussion

To answer our fourth research question (Q4), we discussed how the backbone of the Mastodon network of instances is unveiled by leveraging on graph pruning approaches based on generative null models. These approaches allow us to discard noisy edges in the network and shed light on its well-rooted underlying structures.

MLF and disparity models showed to differ in pruning intensity, with the latter being more severe. To explain this, recall that the disparity filter accounts for both the degree and strength of a node locally, whereas MLF relies on a maximum-entropy based, unbiased null model: with this in mind, the disparity filter is likely to be more heavily conditioned on in-degree and out-degree distributions that were observed in the INSTANCES network to best fit with a lognormal distribution and, to a less extent, with a power-law distribution, which is the best scenario of application of disparity (Serrano et al. 2009).

A major remark that stands out is nonetheless a certain consistency with the structural properties of the INSTANCES network; that is, apart from a significant decrease in transitivity as an expected effect of the pruning, the observed structural characteristics in

**Table 6** Summary of structural characteristics of the INSTANCES network before and after graph pruning. For each of the two graph pruning models, we tested under significance levels at 1% and 5%, which are indicated within brackets

	Unpruned	MLF (0.05)	MLF (0.01)	Disparity (0.05)	Disparity (0.01)
#Nodes	6960	6454	6081	3949	3260
#Edges	216,504	97,141	68,100	19,208	12,849
Reciprocity	65.1%	57.9%	58.4%	67.8%	68.5%
Density	0.004	0.002	0.002	0.001	0.001
Average degree*	41.966	21.390	15.860	6.429	5.182
Average in-degree	31.107	15.051	11.199	4.864	3.941
% Sources	12%	13.1%	14.1%	24.1%	26.5%
% Sinks	6.6%	6.5%	6.7%	6.9%	7.9%
Degree assortativity*	- 0.274	- 0.249	- 0.224	- 0.244	- 0.272
Degree assortativity	- 0.253	- 0.231	- 0.214	- 0.253	- 0.283
Average path length	2.330	3.080	3.270	2.530	2.520
Diameter	5	6	7	5	6
Transitivity*	0.128	0.070	0.066	0.024	0.018
Clustering coefficient*	0.836	0.502	0.493	0.846	0.839
Clustering coefficient (full averaging)*	0.687	0.411	0.390	0.447	0.370
#Strongly connected components	1305	1276	1278	1228	1121
#Weakly connected components*	1	1	1	1	1
Modularity by Louvain*	0.289	0.495	0.540	0.469	0.495
#Communities by Louvain*	5 (5)	7 (7)	7 (7)	8 (8)	8 (9)
Modularity by Louvain**	0.353	0.442	0.442	0.369	0.375
#Communities by Louvain**	6 (8)	6 (8)	7 (9)	7 (9)	7 (11)
#Communities by Infomap**	6 (54)	10 (103)	12 (82)	6 (30)	5 (38)

\*Statistic calculated by discarding edge orientation

\*\*Statistic calculated by taking into account the edge weights

the pruned networks remain comparable to those of the original network, or even more emphasized as for the modularity. In particular, the disassortative trait is still present in the backbone, which certainly strengthens our previous finding so to distinguish Mastodon from many other social networks in terms of negative degree correlation.

It is also worth noticing that the top-5 instances discussed in section "Macroscopic structural analysis", regardless of the particular community detection method are still found after all the considered pruning scenarios, which indicates that these instances are important constituent of the backbone of the INSTANCES network. Overall, even the pruned scenarios exhibit traits related to a tight structure among instances. We recall that pruning approaches aim to filter out noise from networks, and therefore these traits are further strengthened following a cleaner perspective.

### Evolution of the network of Mastodon instances

In this section we answer our fifth research question (Q5), namely "how has Mastodon evolved during the last few years?". To this aim, we organize our presentation into four parts: the first one is devoted to a comparison between the main network we introduced into this work, i.e., INSTANCES, and the network of instances previously

studied in Zignani et al. (2018); the second part focuses on the subnetwork composed of online instances in our updated network; the third part describes our analysis of centrality of the instances based on the PageRank method; finally, we report a summary of the main lessons learned from our study of the Mastodon growth.

### Comparison with the earlier Mastodon network

In this section we discuss a comparative evaluation between the Mastodon instances network introduced in this work and the earlier network presented in Zignani et al. (2018). By replicating on the latter the same structural analysis as we have carried out on the INSTANCES network, we gain insights into the evolution of Mastodon during the last 3 years, focusing on a set of macroscopic and mesoscopic characteristics that may confirm some traits or highlight novel trends. Table 7 shows the structural properties that we observed on earlier Mastodon network along with the percentage increase values obtained by the corresponding characteristics in our network (cf. Table 3).

Besides the already noticed larger size of the current network (with a +73% of instances and a +127% of links) compared to the earlier one, a general remark that arises is that the two networks of Mastodon instances show differences in several characteristics, though with two major exceptions: the one relating to the average path length and the diameter, which are both unchanged, and the other one referring to both global and local clustering coefficient, which are slightly lower (i.e., from -5% to 1%) in the current network. This would hint at a consolidated small-world behavior by the Mastodon instances.

As for the remaining properties, we observe in our updated network a small decrease in the percentage of reciprocal edges, while the percentage of sources and sinks is respectively more than doubled and decreased of less than one third. One reasonable explanation for that relates to an increased presence of offline instances in the INSTANCES network whose further neighborhood however cannot be explored via API. It should however be noted that our updated INSTANCES network still outperforms in size the earlier network even when removing their respective sources and sinks, with a percentage increase of 66% in the number of instances.

Also, in the INSTANCES network, the number of strongly connected components is more than doubled, while the decrease in density (about one third) should be related to a high increase in the number of sources, which is in turn an effect of our deep crawling.

A major aspect of interest we found out in the INSTANCES network, that is, degree disassortativity, holds similarly for the earlier network as well. This is noteworthy as it unveils consistency of this property in Mastodon over time.

Our updated network also shows more communities but a slightly less modular structure than the earlier one. Again, this should be ascribed to the impact due to the increase in the number of source instances. As concerns core decomposition, we notice that a high degeneracy also characterizes the earlier network, and the values according to three degree variants are actually close to those observed in the INSTANCES network in proportion of the respective node-set sizes, with only 9%, resp. 2%, of change when considering degree-based decomposition, resp. in-degree- or out-degree-based decomposition.



**Table 7** Evolution of the INSTANCES network through a comparison of structural characteristics with the earlier state-of-the-art network. Percentage changes refer to the increase/decrease of a statistic from the value observed in the earlier network. For the community statistics, percentage values refer to the number of communities containing at least ten instances

	Earlier (Zignani et al. 2018)	% Change
#Nodes	4015	+ 73%
#Edges	95,221	+ 127%
Reciprocity	70.9%	− 8%
Density	0.006	− 33%
Average degree*	30.612	+ 37%
Average in-degree	23.716	+ 31%
% Sources	5.63%	+ 113%
% Sinks	9.39%	− 30%
Degree assortativity*	− 0.287	− 5%
Degree assortativity	− 0.291	− 13%
Average path length	2.340	0%
Diameter	5	0%
Transitivity*	0.135	− 5%
Clustering coefficient*	0.848	− 1%
Clustering coefficient ( <i>full averaging</i> )*	0.710	− 3%
#Strongly connected components	604	+ 116%
#Weakly connected components*	1	0%
Modularity by <i>Louvain</i> *	0.356	− 19%
#Communities by <i>Louvain</i> *	4 (4)	+ 25%
Modularity by <i>Louvain</i> **	0.397	− 11%
#Communities by <i>Louvain</i> **	3 (5)	+ 100%
#Communities by <i>Infomap</i> **	5 (63)	+ 20%
Degree-based degeneracy	141	+ 43%
Degree-based inner-most-core #nodes	120	+ 58%
Degree-based inner-most-core #edges	11,227	+ 130%
In-degree-based degeneracy	69	+ 39%
In-degree-based inner-most-core #nodes	123	+ 69%
In-degree-based inner-most-core #edges	11,401	+ 152%
Out-degree-based degeneracy	70	+ 39%
Out-degree-based inner-most-core #nodes	115	+ 70%
Out-degree-based inner-most-core #edges	10,385	+ 155%

\*Statistic calculated by discarding edge orientation

\*\*Statistic calculated by taking into account the edge weights

### Narrowing the focus on online Mastodon instances

Our further step towards a comprehensive understanding of the current landscape of Mastodon instances is an analysis of the subnetwork corresponding to the online instances only, i.e., the ONLINE-INSTANCES network, which was previously introduced in section "Network models". The goal here is to check whether our findings on the overall INSTANCES network reflect on its subnetwork of online instances as well, or on the

contrary there are significant deviations on some traits. Table 8 summarizes the comparison between the two networks.

A first remark clearly arises from the reduction in both the number of instances (−84%) and links (−65%) compared to the INSTANCES network. This is clearly expected due to the deep API-based crawling capabilities, yet at the same time it captures the currently active snapshot of Mastodon in the Fediverse. We conjecture that, after a few years of novelty and curiosity that Mastodon attracted, a steady-state has been reached as composed by those instances that might be recognized as most established for the Mastodon users. The subnetwork of online instances also turns out to be more tightly knit than the INSTANCES network, as indicated by the increase in reciprocity (+ 8%), density (one order of magnitude higher), increase in average degree and in-degree (more than doubled), decrease in average path length and diameter (−13% and −20%, respectively), the almost tripled transitivity, halved number of strongly connected components, and decrease in the percentage of sources and sinks (about 80%). The latter aspect suggests that the high percentages (particularly for sources) found in the INSTANCES network is likely to be ascribed to the different status (i.e., online or offline) of the instances, which affects the search outcomes. Nonetheless, since the ONLINE-INSTANCES network contains currently active and online instances, their users might be more engaged in developing interrelations.

Notably, the disassortative trait is still present in the ONLINE-INSTANCES network, thus confirming it as a well-rooted feature in the network of Mastodon instances.

As concerns the community structure, the ONLINE-INSTANCES network exhibits similar characteristics to the overall INSTANCES, with a relatively small reduction in undirected modularity—which becomes negligible for the weighted directed case—and in the number of communities only when accounting for edge weight and orientation. Interestingly, we point out that the number of significant communities (i.e., those containing at least ten instances) in the ONLINE-INSTANCES is the same for both Louvain and Infomap method, which is consistent with the situation already observed in the INSTANCES network. Analogous considerations can be made for the core decomposition results: indeed, not only the degeneracy remains still high in all three variants—with a decrease of about 15% w.r.t. INSTANCES—but also the sizes of the corresponding inner-most cores are proportionally higher than those in INSTANCES according to the respective node-set sizes, with 405%, resp. 419 and 403%, of increase when considering degree-based decomposition, resp. in-degree- and out-degree-based decomposition. This is remarkable as it provides evidence of a well-established trend in Mastodon as to have a conspicuous concentration of instances in the inner-most central region of the network.

### Instance centrality

To further investigate the growth of Mastodon, we evaluated how the *importance* of the instances developed during the years. To this aim, we resorted to the well-known *PageRank* method (Brin and Page 1998), which assigns prestige scores to the nodes in a network. Using this tool, we comparatively evaluated results obtained on all the networks derived from our crawling data, i.e., the EXPANDED-INSTANCES, the INSTANCES, and

**Table 8** Evolution of the INSTANCES network through a comparison of structural characteristics with the ONLINE-INSTANCES network. Percentage changes refer to the increase/decrease of a statistic from the value observed in the INSTANCES network. For the community statistics, percentage values refer to the number of communities containing at least ten instances

	ONLINE-INSTANCES	% Change
#Nodes	1115	− 84%
#Edges	75,046	− 65%
Reciprocity	70.6%	+ 8%
Density	0.06	+ 1400%
Average degree*	87.074	+ 107%
Average in-degree	67.306	+ 116%
% Sources	2.1%	− 83%
% Sinks	1.4%	− 79%
Degree assortativity*	− 0.290	+ 6%
Degree assortativity	− 0.284	+ 12%
Average path length	2.020	− 13%
Diameter	4	− 20%
Transitivity*	0.369	+ 188%
Clustering coefficient*	0.712	− 15%
Clustering coefficient ( <i>full averaging</i> )*	0.689	0%
#Strongly connected components	41	− 97%
#Weakly connected components*	1	0%
Modularity by <i>Louvain</i> *	0.229	− 21%
#Communities by <i>Louvain</i> *	5 (5)	0%
Modularity by <i>Louvain</i> **	0.336	− 5%
#Communities by <i>Louvain</i> **	4 (7)	− 33%
#Communities by <i>Infomap</i> **	4 (15)	− 33%
Degree-based degeneracy	168	− 16%
Degree-based inner-most-core #nodes	153	− 19%
Degree-based inner-most-core #edges	17,449	− 32%
In-degree-based degeneracy	80	− 17%
In-degree-based inner-most-core #nodes	173	− 17%
In-degree-based inner-most-core #edges	20,320	− 29%
Out-degree-based degeneracy	82	− 15%
Out-degree-based inner-most-core #nodes	158	− 19%
Out-degree-based inner-most-core #edges	17,987	− 32%

\*Statistic calculated by discarding edge orientation

\*\*Statistic calculated by taking into account the edge weights

the ONLINE-INSTANCES network, as well as the earlier state-of-the-art network. This approach allowed us to assess whether and to what extent the centrality of instances was affected by the introduction of new instances and/or the dismissal of instances (e.g., instances went to offline mode for a relatively long period, cf. section "[Polite data crawling and network modeling](#)").

We assessed the strength of relatedness between the PageRank solutions obtained on the above networks by means of two standard rank correlation methods, namely *Kendall correlation coefficient* (Abdi 2007) and *Fagin's intersection metric* (Fagin et al. 2003).

The Kendall correlation  $\tau$  evaluates the similarity between rankings represented through set of ordered pairs assigned to the same set of nodes, relying on the number of inversions of pairs needed to transform one ranking into the other. Given two rankings  $\mathcal{R}'$  and  $\mathcal{R}''$  obtained on the same set of  $N$  items, this non-parametric correlation is formally expressed as follows:

$$\tau(\mathcal{R}', \mathcal{R}'') = 1 - \frac{2\Delta(\mathcal{P}(\mathcal{R}'), \mathcal{P}(\mathcal{R}''))}{N(N-1)},$$

where  $\Delta(\mathcal{P}(\mathcal{R}'), \mathcal{P}(\mathcal{R}''))$  denotes the symmetric difference between  $\mathcal{R}'$  and  $\mathcal{R}''$ , i.e., the number of unshared pairs between the two rankings. The values returned by  $\tau$  are within the range  $[-1, 1]$ , where 1 indicates that the two rankings are identical and -1 means that one rank is the reverse order of the other.

The Kendall correlation assigns the same importance to all items, regardless of their position. Therefore, we also employed Fagin's intersection metric  $F$ , which considers partial rankings and assigns higher weights to items at the top of the lists. Given two ranking lists  $\mathcal{R}'$  and  $\mathcal{R}''$ , Fagin's metric is formalized as follows:

$$F(\mathcal{R}', \mathcal{R}'', k) = \frac{1}{k} \sum_{q=1}^k \frac{|\mathcal{R}'_{:q} \cap \mathcal{R}''_{:q}|}{q},$$

where  $k$  is a parameter to determine the number of items to consider from the top of both rankings, and  $\mathcal{R}_{:q}$  indicates the set of nodes from the 1st to the  $q$ th position in the ranking. Hence, we refer to  $F$  as the average over the sum of the weighted overlaps calculated considering the first  $k$  nodes in both rankings. Its values are within the range  $[0, 1]$ , where the highest the value, the better the score.

The results of our evaluations are reported in Table 9. Since Kendall correlation requires a comparison of rankings of nodes from the same set, the results obtained on any pair of networks correspond to the shared set of nodes. Note also that in all cases the Kendall correlation values are high, and associated with  $p$ -values equal to zero, which means rejection of the hypothesis of absence of correlation.

Let us first consider the comparison between EXPANDED-INSTANCES and INSTANCES networks (i.e., first column of Table 9). The Fagin's intersection values observed for various  $k$  values are very high, as expected. Indeed, as described in section "Polite data crawling and network modeling", non-Mastodon instances represent the boundary of our EXPANDED-INSTANCES network, and we have only partial knowledge of them. Hence, the more central role taken by the Mastodon instances is reasonable. However, while these remarkably high values denote that non-Mastodon instances only slightly influence the Mastodon ones, the relatively smaller value obtained for  $F_{10}$  indicates that this influence actually affects the PageRank scores of the top-10 instances.

The comparison between ONLINE-INSTANCES and INSTANCES (i.e., second column of the table) allows us to inspect the role of online vs. non-online instances. Paying

**Table 9** Ranking analysis performed via Kendall’s tau ( $\tau$ ) and Fagin’s intersection metric ( $F$ ) with various  $k$  values (indicated as subscripts) on the PageRank solutions obtained on the Mastodon networks

	EXPANDED-INSTANCES vs. INSTANCES	ONLINE-INSTANCES vs. INSTANCES	EXPANDED-INSTANCES vs. ONLINE-INSTANCES	INSTANCES vs. <i>Earlier</i>
$\tau$	0.953	0.899	0.883	0.626
$F_{10}$	0.860	0.935	0.886	0.569
$F_{50}$	0.911	0.814	0.813	0.563
$F_{100}$	0.918	0.832	0.835	0.545
$F_{500}$	0.938	0.821	0.828	0.487
$F_{1000}$	0.915	0.748	0.738	0.473

attention to  $F_{10}$ , the noticeable value indicates how the most central Mastodon instances are well-rooted in their role, regardless of potential variations, even evident, in the number of the online ones at any given time. Although a bit less evident, this trend remains valid in the range [50, 100], while we observed a further decrease with  $F_{1000}$ ; the latter might indicate the existence of relevant instances among the offline or temporarily inactive ones.

As concerns EXPANDED-INSTANCES vs. ONLINE-INSTANCES, we observe in general slightly lower values w.r.t. the previous comparison, which might be ascribed to the absence of some offline yet relevant instances as well as to the lack of contributions given by non-Mastodon instances. However, the still high correlation values observed further support the centrality of the currently online Mastodon instances.

The last comparison we consider is between the INSTANCES and the *Earlier* network. It is not surprising to observe a general decrease in the correlation values, given the temporal difference between the two networks. Nonetheless, the corresponding ranking lists over the shared set of instances show a good Kendall correlation (above 0.6); in addition, the Fagin’s intersection metric values are also quite high, for various  $k$  values. Overall, despite the introduction of new instances after about three years, the results indicate that a significant portion of those instances that are shared between the two networks have not drastically changed their strategic location in their respective networks.

**Discussion**

To investigate the evolution of Mastodon from an instance perspective, we compared the various networks presented in section "Polite data crawling and network modeling", keeping our focus on the INSTANCES network. First, we performed a comparison with the *Earlier* network, which refers to about three years ago, as previously discussed. This allowed us to highlight the structural development of Mastodon. Although the different sizes have some effect on the measures, we noticed some well-established features. In particular, we report the same average path length and diameter in both networks, whose consolidation suggests that Mastodon instances act in a small-world fashion. The *Earlier* network also yields the same disassortative trait spotted for the more recent INSTANCES network. Thus, we can establish the existence of this property over time. This consolidation allows us to consider the negative degree correlation as a distinctive trait of Mastodon w.r.t. centralized social networks and paves the way for further

studies concerning interactions among instances through the federative approach. From a mesoscopic perspective, the evaluation of the core decomposition brings us to confirm another important trait. Indeed, as initially shown by the *Earlier* network and subsequently confirmed by the *INSTANCES* one, Mastodon instances exhibit high degeneracy values, comparable when accounting for the different sizes.

Along with the temporal development, we analyzed the active status of Mastodon, focusing on the subnetwork induced from its online instances (i.e., the *ONLINE-INSTANCES* network). This narrowing produces an exceptional reduction in size (-84% instances and -65% links among them), which might indicate a stationary state in the number of online Mastodon instances. Indeed, the spotted online ones are found to be no longer transient (e.g., testing instances) or driven by the feeling of novelty, but they are rather well-defined and established. Remarkably, the *ONLINE-INSTANCES* network still exhibits the disassortative trait. The core decomposition of the *ONLINE-INSTANCES* network confirms the high degeneracy found in the other Mastodon networks, setting it as a well-established trait of Mastodon, along with a remarkable concentration of instances in the inner-most core.

Finally, the study of the instance centrality based on PageRank allowed us to assert that the most prestigious instances are settled in their roles, also considering the temporal development, which indicates an actual evolution of the platform towards a stable state. The instance-rankings computed over different pairs of instance networks have generally shown good or very high correlation, with relatively small fluctuations mainly due to the centrality contributions from offline or boundary instances, but also as an effect of the growth of the platform over time.

Overall, and in light of all the above findings, we can state that Mastodon has reached sort of structural stability. This well-established structure couples with a solid federative mechanism among instances, which acts as mutual-reinforcement towards the sectorization bias induced by the decentralized scenario.

### **Conclusion and future work**

Nowadays, DOSNs express an alternative, user-centric approach to support online interactions among users, with attention on self-hosted networking services, and ownerships in terms of code of conduct and moderation policies for each operating server. Mastodon represents the most widely adopted and recognized platform in the Fediverse of DOSNs. In this work, we provided the first in-depth analysis of the network of Mastodon instances based on the largest and most up-to-date Mastodon relational data existing so far, which was originally built in this work.

We analyzed our instance network models from different perspectives (i.e., macroscopic, mesoscopic, backbone) to highlight the main structural characteristics that describe Mastodon. Among these, we spotted remarkable traits that define the fingerprint of Mastodon, setting it apart from well-known OSNs. Finally, we investigated the

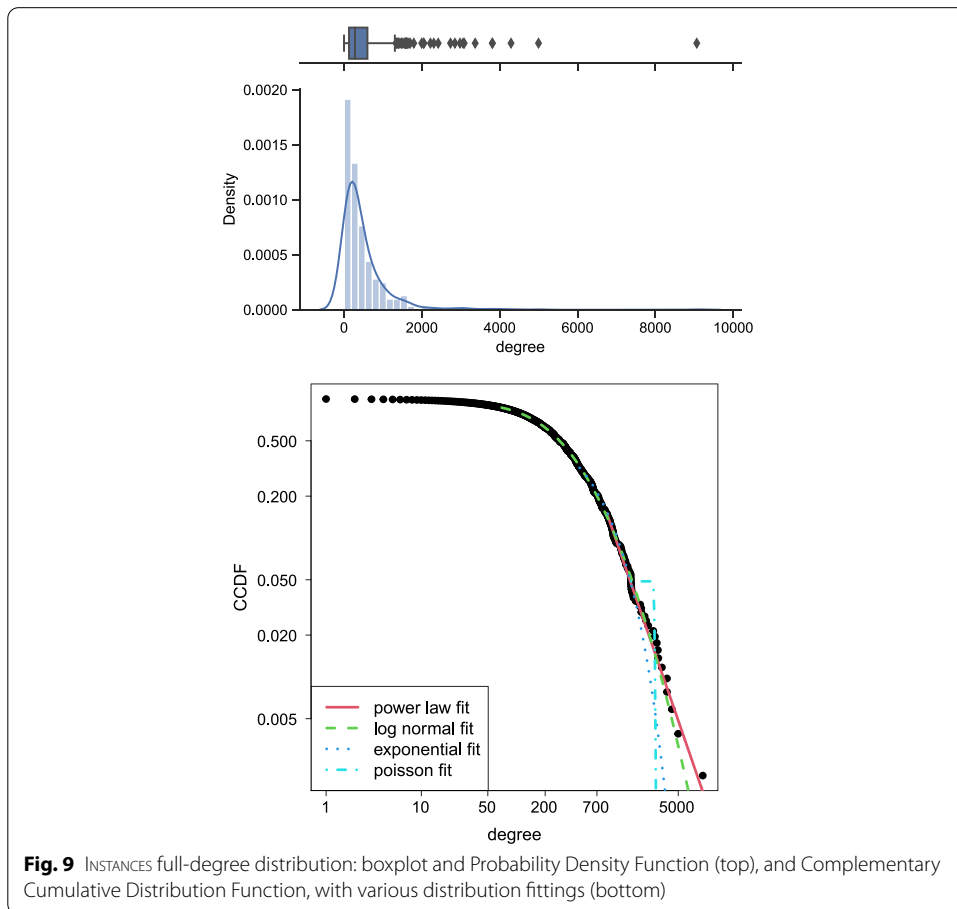


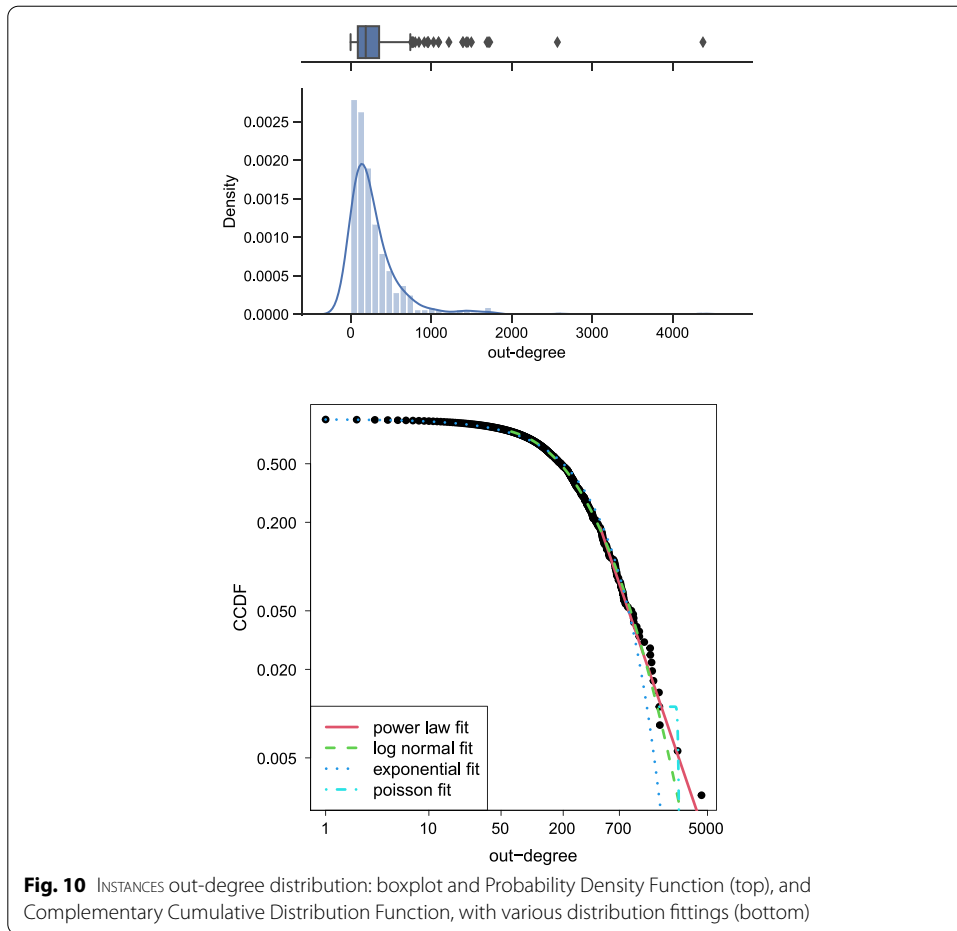
evolution of Mastodon, by identifying features that have changed over time and those representing supporting pillars.

Our next research stage will be focused on the underlying network of Mastodon *users*. As discussed in this work, our collected data are fine-grained at a user level, although the stored information concern user memberships and relations only, i.e., social contents are discarded. In this regard, by still leveraging the network of links formed by users, we plan to study the behavior of users not only within their home instances, but also their relations across instances. In particular, we believe it would be interesting to investigate the roles that users may take in their instances, at various levels of activity, and possibly alternate or opposite roles that the same users may exhibit when interacting with users of other instances. This would allow us to also analyze the dichotomy between contribution and consumption of information induced by the user behavior relations over the instances of Mastodon and other DOSNs in the Fediverse.

### Appendix

See Figs. 9 and 10.





**Abbreviations**

OSN: Online social network; DOSN: Decentralized online social network.

**Acknowledgements**

An abridged version of this work appeared in the *Proc. of the 9th Int. Conf. on Complex Networks and their Applications, 2020* (La Cava et al. 2020).

**Authors' contributions**

AT and LL conceived the study and performed the analysis; LL carried out the data collection and the experiments; LL and AT prepared the result presentation; AT, LL and SG wrote the manuscript. All authors read and approved the final manuscript.

**Funding**

The authors received no specific funding for this work.

**Availability of data and materials**

Data available on request from the authors.

**Declarations**

**Ethics approval and consent to participate**

Not applicable.

**Consent for publication**

Yes. No figure in this manuscript is under copyright or has been adapted from copyrighted material.

**Competing interests**

The authors declare that they have no competing interests.

Received: 29 March 2021 Accepted: 12 June 2021

Published online: 01 September 2021

**References**

- Abdi H (2007) The Kendall rank correlation coefficient. In: Encyclopedia of measurement and statistics
- Blondel VD, Guillaume J-L, Lambiotte R, Lefebvre E (2008) Fast unfolding of communities in large networks. *J Stat Mech Theory Exp* 10:10008
- Brin S, Page L (1998) The anatomy of a large-scale hypertextual Web search engine. *Comput Netw ISDN Syst* 30(1–7):107–117
- Cerisara C, Jafaritazehjani S, Oluokun A, Le HT (2018) Multi-task dialog act and sentiment recognition on Mastodon. In: Proceedings of the COLING conference, pp 745–754. <https://www.aclweb.org/anthology/C18-1063/>
- Caliò A, Tagarelli A, Bonchi F (2020) Cores matter? An analysis of graph decomposition effects on influence maximization problems. In: Proceedings of the 12th ACM conference on web science (WebSci). Association for Computing Machinery, New York, NY, USA, pp 184–193. <https://doi.org/10.1145/3394231.3397908>
- Datta A, Buchegger S, Vu L-H, Strufe T, Rzacca K (2010) Decentralized online social networks. In: Furht B (ed) Handbook of social network technologies and applications. Springer, Boston, pp 349–378. [https://doi.org/10.1007/978-1-4419-7142-5\\_17](https://doi.org/10.1007/978-1-4419-7142-5_17)
- Dianati N (2016) Unwinding the hairball graph: pruning algorithms for weighted complex networks. *Phys Rev E* 93:012304
- Fagin R, Kumar R, Sivakumar D (2003) Comparing top k lists. *SIAM J Discrete Math* 17(1):134–160
- Fisher DN, Silk MJ, Franks DW (2017) The perceived assortativity of social networks: methodological problems and solutions. *CoRR arXiv:1701.08671*
- Guidi B, Conti M, Passarella A, Ricci L (2018) Managing social contents in decentralized online social networks: a survey. *Online Soc Netw Media* 7:12–29. <https://doi.org/10.1016/j.osnem.2018.07.001>
- Kitsak M, Gallos LK, Havlin S, Liljeros F, Muchnik L, Stanley HE, Makse H (2010) Identification of influential spreaders in complex networks. *Nat Phys* 6(11):888–893
- La Cava L, Ruffo LE, Tagarelli A (2020) Towards mesoscopic structural analysis of the Fediverse of decentralized social networks. In: Proceedings of the 9th international conference on complex networks and their applications—book of abstracts
- Malliaros FD, Giatsidis C, Papadopoulos AN, Vazirgiannis M (2020) The core decomposition of networks: theory, algorithms and applications. *VLDB J* 29(1):61–92
- Newman MEJ (2002) Assortative mixing in networks. *Phys Rev Lett*. <https://doi.org/10.1103/physrevlett.89.208701>
- Newman MEJ (2003) Mixing patterns in networks. *Phys Rev E*. <https://doi.org/10.1103/physreve.67.026126>
- Raman A, Joglekar S, Cristofaro ED, Sastry N, Tyson G (2019) Challenges in the decentralised web: the Mastodon case. In: Proceedings of the ACM IMC conference, pp 217–229. <https://doi.org/10.1145/3355369.3355572>
- Rosvall M, Bergstrom CT (2008) Maps of information flow reveal community structure in complex networks. *Proc Natl Acad Sci (PNAS)* 105:1118
- Seidman S (1983) Network structure and minimum degree. *Soc Netw* 5:269–287. [https://doi.org/10.1016/0378-8733\(83\)90028-X](https://doi.org/10.1016/0378-8733(83)90028-X)
- Serrano MÁ, Boguñá M, Vespignani A (2009) Extracting the multiscale backbone of complex weighted networks. *Proc Natl Acad Sci* 106(16):6483–6488. <https://doi.org/10.1073/pnas.0808904106>
- Trienes J, Cano AT, Hiemstra D (2018) Recommending users: whom to follow on federated social networks. *CoRR arXiv:1811.09292*
- Varol O, Ferrara E, Davis CA, Menczer F, Flammini A (2017) Online human–bot interactions: detection, estimation, and characterization. In: Proceedings of the international conference on web and social media (ICWSM), pp 280–289
- Zignani M, Gaito S, Rossi GP (2018) Follow the “Mastodon”: structure and evolution of a decentralized online social network. In: Proceedings of the international conference on web and social media (ICWSM), pp 541–551
- Zignani M, Quadri C, Gaito S, Cherifi H, Rossi GP (2019) The footprints of a “Mastodon”: how a decentralized architecture influences online social relationships. In: Proceedings of the IEEE INFOCOM workshops, pp 472–477. <https://doi.org/10.1109/INFOCOMW.2019.8845221>
- Zulli D, Liu M, Gehl R (2020) Rethinking the “social” in “social media”: insights into topology, abstraction, and scale on the mastodon social network. *New Media Soc* 22(7):1188–1205

**Publisher’s Note**

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