

RESEARCH

Open Access



# Uncovering complex overlapping pattern of communities in large-scale social networks

Elvis H.W. Xu<sup>1\*</sup>  and Pak Ming Hui<sup>1</sup>

\*Correspondence:

hwxu@link.cuhk.edu.hk

<sup>1</sup>Department of Physics, The Chinese University of Hong Kong, Shatin, Hong Kong, China

## Abstract

The conventional notion of community that favors a high ratio of internal edges to outbound edges becomes invalid when each vertex participates in multiple communities. Such a behavior is commonplace in social networks. The significant overlaps among communities make most existing community detection algorithms ineffective. The lack of effective and efficient tools resulted in very few empirical studies on large-scale detection and analyses of overlapping community structure in real social networks. We developed recently a scalable and accurate method called the Partial Community Merger Algorithm (PCMA) with linear complexity and demonstrated its effectiveness by analyzing two online social networks, Sina Weibo and Friendster, with 79.4 and 65.6 million vertices, respectively. Here, we report in-depth analyses of the 2.9 million communities detected by PCMA to uncover their complex overlapping structure. Each community usually overlaps with a significant number of other communities and has far more outbound edges than internal edges. Yet, the communities remain well separated from each other. Most vertices in a community are multi-membership vertices, and they can be at the core or the peripheral. Almost half of the entire network can be accounted for by an extremely dense network of communities, with the communities being the vertices and the overlaps being the edges. The empirical findings ask for rethinking the notion of community, especially the boundary of a community. Realizing that it is how the edges are organized that matters, the *f*-core is suggested as a suitable concept for overlapping community in social networks. The results shed new light on the understanding of overlapping community.

**Keywords:** Community structure, Overlapping community, Social network

## Introduction

A community in networks is conceived commonly as a group of vertices connected closely with each other but only loosely to the rest of the network. Such communities are widespread in many systems and their detection has attracted much attention in the past two decades (Fortunato 2010). This vague notion of communities is subjected to many possible interpretations. The most common one is based on the ratio of the numbers of internal edges to outbound edges, which go out of the community. The more the internal edges to outbound edges, the more definite is the community. For example, the widely used methods based on strong/weak community (Radicchi et al. 2004), LS-set (Luccio and Sami 1969), conductivity and network community profile (Leskovec et al. 2009; Jeub

et al. 2015), and fitness functions (Baumes et al. 2005; Lancichinetti et al. 2009; Goldberg et al. 2010) favor a higher internal edges to outbound edges ratio. The idea works well for disjoint communities, but it has also been adopted by algorithms for detecting overlapping communities (Xie et al. 2013). Nonetheless, the number of members, mostly at the periphery, belonging to multiple communities is still expected to be small so that an “overlapping community” remains well separated from its surrounding. However, the structure of overlapping communities in real social networks may be far more complex. It is commonplace that every individual has multiple social circles. It implies that all parts of a social community, peripheral and core, may be overlapping with a significant number of other communities and there can be far more outbound edges than internal edges. The existence of these significantly overlapped communities, as will be shown in the present work, asks for a deeper understanding of what an overlapping community really is, where their boundaries are, and how to detect them.

Analyzing big data sets of real social networks is vital in network science. An immediate problem is that most existing methods are incapable of detecting significantly overlapped groups of vertices, because these groups have too many outbound edges to be identified as well separated communities. The recently proposed methods of OSLOM (Lancichinetti et al. 2011) and BIGCLAM (Yang and Leskovec 2013) are useful to some extent in small synthetic networks, but they become inefficient for large-scale networks which readily have the size of millions to billions of vertices. Sampling small subnetworks (Maiya and Berger-Wolf 2010) would not work either due to the small-world effect (Watts and Strogatz 1998), e.g. the average distance between any two individuals on Facebook is only 4.74 (Backstrom et al. 2011; Ugander et al. 2011), while in a social group for which the size is small compared to the whole network, a member may usually need one or two hops to be connected to all the other members. A community may be considered as localized, but it is also widespread in the network. Sampling small subnetworks would preserve particular communities but decompose many others, making it inappropriate for studying the overlaps among communities. Some newly proposed algorithms (Lyu et al. 2016; Sun et al. 2017; Epasto et al. 2017) achieved linear-time complexity, but their validity and accuracy in detecting significantly overlapped communities requires further benchmarking and cross-checking. The lack of effective and efficient algorithms resulted in very few studies on detecting and analyzing overlapping community structure in large-scale social networks. An empirical study was carried out on Facebook (Ferrara 2012), but only methods for detecting disjoint communities were used. A recent study on Friendster found that about 30% of the vertices belonged to multiple communities (Epasto et al. 2017). Jebabli et al. analyzed community structure in a sampled YouTube network of 1.1 million vertices and evaluated a number of overlapping community detection algorithms (Jebabli et al. 2015; 2018). Yang and Leskovec analyzed metadata groups of some real networks and found that overlaps occur more often at the cores of communities (Yang and Leskovec 2015; 2014). This is contrary to the traditional notion that overlapping members are mostly at the periphery. Recent studies also revealed that metadata groups may not give the ground-truth of structural communities (Hric et al. 2014; Peel et al. 2017).

The present authors developed recently a scalable partial community merger algorithm (PCMA) which adopts  $f$ -core as the notion of community that a member of a community should know at least a fraction  $f$  of the other members (Xu and Hui 2018; Xu 2016). The concept of  $f$ -core imposes no constraints or implications on the fraction of overlapping

vertices in a community or the number of communities a vertex may belong to. The method is a bottom-up approach by properly reassembling partial information of communities found in ego networks of the vertices to reconstruct the complete communities. It consists of three steps:

1. Find communities in the ego network of each of the vertices. These communities are referred to as partial communities as each of them is only part of the corresponding complete community.
2. Merge partial communities that are parts of the same community to reconstruct complete communities.
3. Clean up the noise accumulated in the merged communities to sift out the real communities.

This approach is intuitive and easy to conceive. There are a number of similar algorithms such as DEMON (Coscia et al.) and EgoClustering (Rees and Gallagher 2013). The reason that PCMA achieves a far better accuracy is a novel similarity measure of communities that suppresses the amount of noise accumulated during the merging process. The present authors tested PCMA against the LFR benchmark (Lancichinetti et al. 2008) and a new benchmark designed for significantly overlapping communities, and established the accuracy and effectiveness of PCMA in detecting communities with significant overlaps, as well as slightly overlapping and disjoint ones. The linear complexity of PCMA enabled the analysis of two huge online social networks with 79.4 and 65.6 million vertices - Sina Weibo and Friendster (see Table 1) - *without* sampling small subnetworks. The  $\sim 2.9$  million communities detected by PCMA were verified to be non-duplicating and have relatively high values of internal edge density. A surprising finding is that more than 99% of them have more outbound edges than internal edges, and the outbound edges often outnumber the internal edges by many times. The communities overlap significantly, while still keeping relatively clear boundaries. These communities are strong empirical evidence against the traditional notion of an overlapping community. While we focused on developing the PCMA algorithm in Ref. (Xu and Hui 2018), we uncover the complex overlapping pattern of these communities in the present work by examining the data in detail and explain why the communities can still remain well separated from each other. After introducing the four main characteristics of the overlapping pattern, we give a macroscopic picture of the social network structure by grouping edges of the entire network into five types. The concept and possible better definitions of an overlapping community are discussed. Additional information on the data sets and the detection of communities is given in the appendix.

### Characteristics of overlapping pattern

In this section, we discuss in detail the four main characteristics of the overlapping pattern of the 2.9 million communities detected by PCMA.

**Table 1** Information on the two huge social networks analyzed

Dataset	$n$	$m$	$\langle k \rangle$	$C_{WS}$	$c$
Sina Weibo	79.4M	1046M	26.4	0.155	1.3M
Friendster	65.6M	1806M	55.1	0.205	1.6M

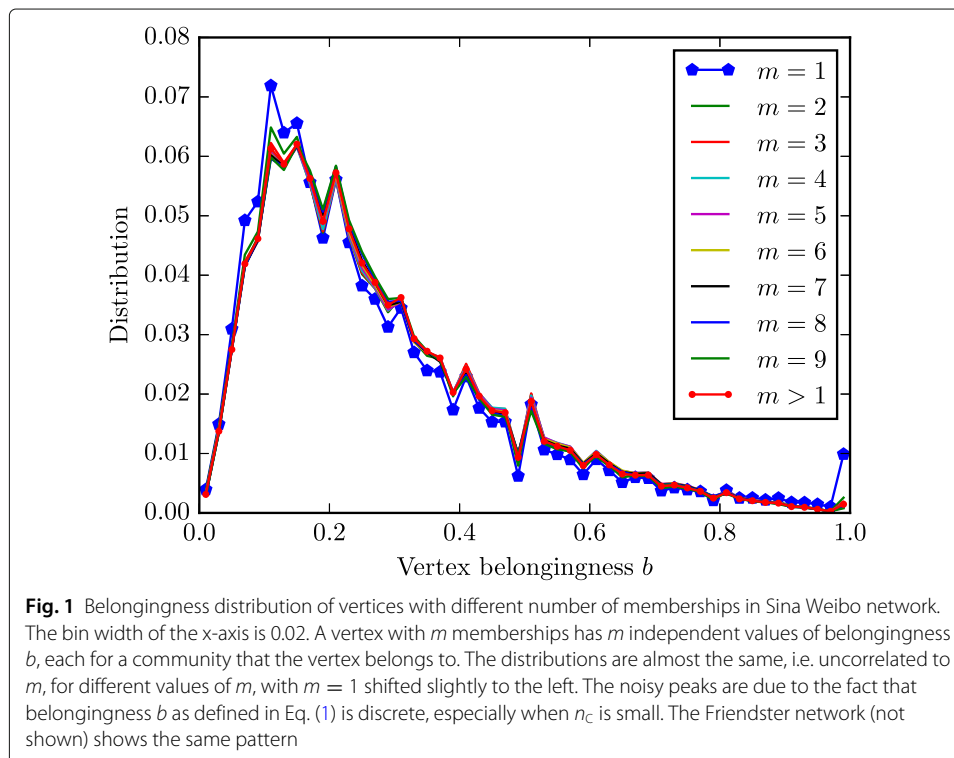
$M$  represents a million.  $n$  and  $m$  are the number of vertices and edges.  $\langle k \rangle$  is the average vertex degree.  $C_{WS}$  is the average local clustering coefficient.  $c$  is the number of communities detected by PCMA. More detailed information is given in the appendix

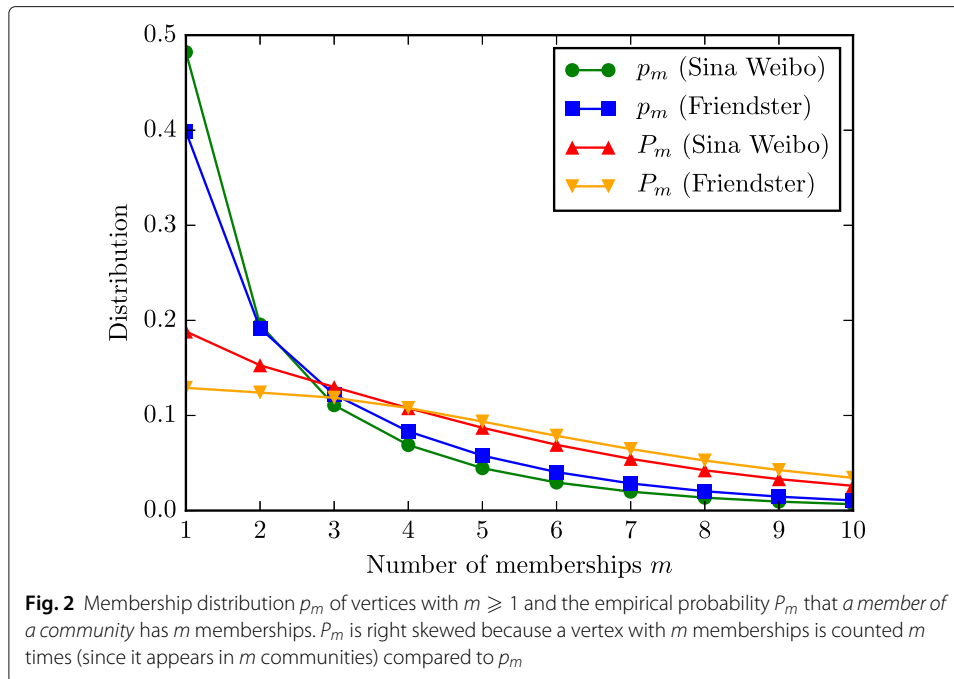
**Characteristic 1.** Multi-membership vertices or overlapping vertices account for the majority of the community, and they are everywhere. These vertices were often thought to be peripheral members. A recent study on metadata groups (Yang and Leskovec 2014) found that these vertices are more likely core members. Our analysis on the two large-scale social networks reveals that the overlapping vertices can be *anywhere*, i.e., core and periphery, in the community. In general, a vertex  $v$  may belong to  $m_v$  communities. The vertices can then be sorted by their values of  $m_v = m$  for  $m \geq 1$ . The belongingness  $b_{v,c}$  of a vertex  $v$  to a community  $C$  can be defined as

$$b_{v,c} = \frac{k_{v,c}^{int}}{n_c - 1} \tag{1}$$

where  $n_c$  is the community size and  $k_{v,c}^{int}$  is the number of other members in  $C$  that are connected to  $v$ . A high (low) value of  $b_{v,c}$  means that  $v$  is closer to the core (periphery) of  $C$ . If overlaps occur more often at the periphery (core), we would expect multi-membership vertices with  $m > 1$  to have a lower (higher) belongingness  $b$  than those with  $m = 1$ . Figure 1 shows that the belongingness distributions for vertices with different values of  $m$  are almost identical, with an insignificant tendency of multi-membership vertices having a slightly higher belongingness. The results imply that  $m_v$  is basically uncorrelated with  $b_{v,c}$ , and multi-membership vertices exist everywhere in a community with no preference towards the core or the periphery as compared with non-overlapping vertices.

Figure 2 shows the distribution  $p_m$  of the number of memberships among the vertices with  $m \geq 1$ . About 50% of the vertices, i.e., those with  $m > 1$ , have multiple memberships in Sina Weibo. For Friendster, the proportion is  $\sim 60\%$ , which is about twice of that





reported in Ref. (Epasto et al. 2017). A related quantity is

$$P_m = \frac{p_m \cdot m}{\langle m \rangle}, \tag{2}$$

which gives the empirical probability that a member of a community has  $m$  memberships. Here,  $\langle m \rangle = \sum_{m=1}^{\infty} p_m \cdot m$  is the mean value of  $m$ . Note that  $P_m$  and  $p_m$  are related but different.  $P_m$  is the expected membership distribution of the members within a community, and  $p_m$  describes the distribution in  $m$  of all vertices with  $m \geq 1$ . Referring to  $P_m$  in Fig. 2,  $P_{m=1} = 18.8\%$  and  $12.9\%$  for Sina Weibo and Friendster, respectively, implying that on average more than 80% of the members in a community are multi-membership vertices. This is in sharp contrast to the preconceived idea that only a small fraction of members in a community belong also to other communities. The results reveal that most members of a community have multiple memberships and they are everywhere in the community.

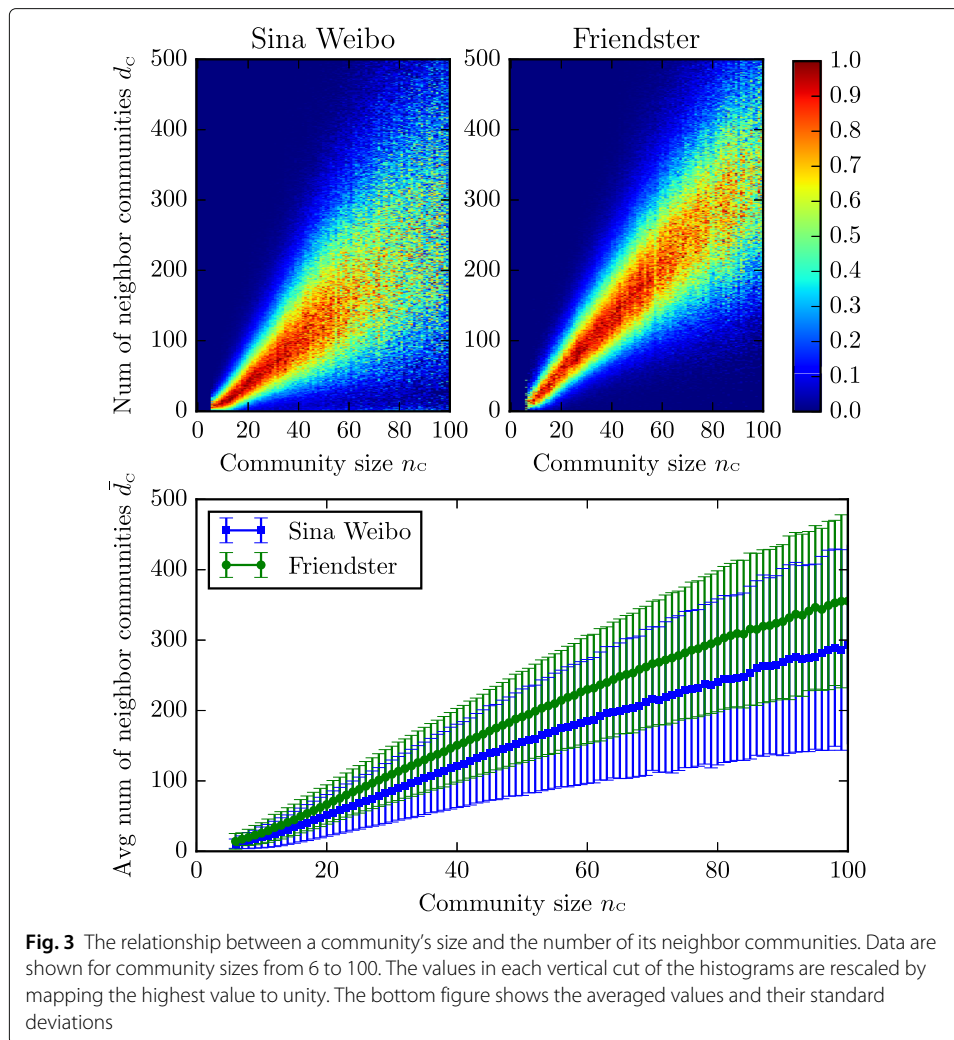
**Characteristic 2.** The multi-membership vertices lead to a community overlapping with many other communities. We refer to them as neighbor communities. Figure 3 shows the relationship between the number of neighbor communities  $d_c$  and the size  $n_c$  of a community in the two social networks. To extract information, the expected number of neighbor communities for a community of size  $n_c$  is roughly

$$\bar{d}_c(n_c) \approx (\langle m \rangle_c - 1) \cdot n_c \cdot r_{nd} \tag{3}$$

where

$$\langle m \rangle_c = \sum_{m=1}^{\infty} P_m \cdot m = \frac{\langle m^2 \rangle}{\langle m \rangle} \tag{4}$$

is the expected number of memberships of a member in the community. Although each member connects the community to  $\langle m \rangle_c - 1$  other communities of which it is also a member,  $(\langle m \rangle_c - 1) \cdot n_c$  overestimates the number of neighbor communities due to duplication, i.e., some members in the community have common neighbor communities. A



factor  $r_{nd}$  is introduced to represent the non-duplicate rate. Consider the simple case of a size- $n_c$  community with  $x$  members all in only one neighbor community. In this case,  $(\langle m \rangle_c - 1) \cdot n_c = x$  while  $\bar{d}_c = 1$ , implying  $r_{nd} = 1/x$ . Thus, the value of  $r_{nd}$  also indicates the extent of overlap between two communities. For overlaps of just 2 or 3 vertices,  $r_{nd}$  drops below 50%. The analysis in Fig. 3 confirms that  $\bar{d}_c \sim n_c$ , but with a slope gradually decreasing with increasing  $n_c$ . Thus,  $r_{nd}$  is negatively correlated with  $n_c$ . The slopes are around  $3 \sim 4$ , which are about 30% smaller than the values 4.36 and 5.25 calculated by  $(\langle m \rangle_c - 1)$  from empirical data of Sina Weibo and Friendster, respectively. Note that these slopes are very large, e.g. a community of size as small as 30 could overlap with  $\sim 100$  other communities concurrently. The resulting non-duplicate rates  $r_{nd}$  are above 70%, strongly indicating that most overlaps concern just one vertex.

**Characteristic 3.** In contrast to the generally believed notion that a community should have more internal edges than outbound edges, we found that more than 99% of the 2.9 million communities have more outbound edges than internal edges. For each community identified by PCMA, we evaluated the total number of internal edges  $k_c^{int}$  and outbound edges  $k_c^{out}$ :



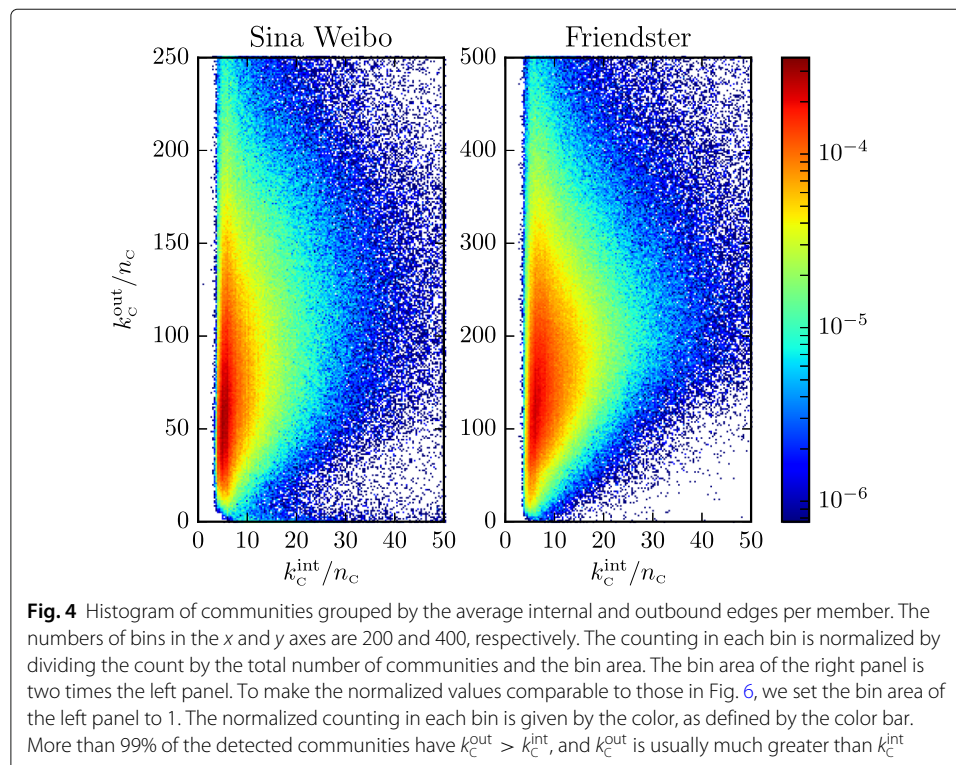
$$k_C^{\text{int}} = \sum_{v \in C} k_{v,C}^{\text{int}}, \quad k_C^{\text{out}} = \sum_{v \in C} k_{v,C}^{\text{out}} \tag{5}$$

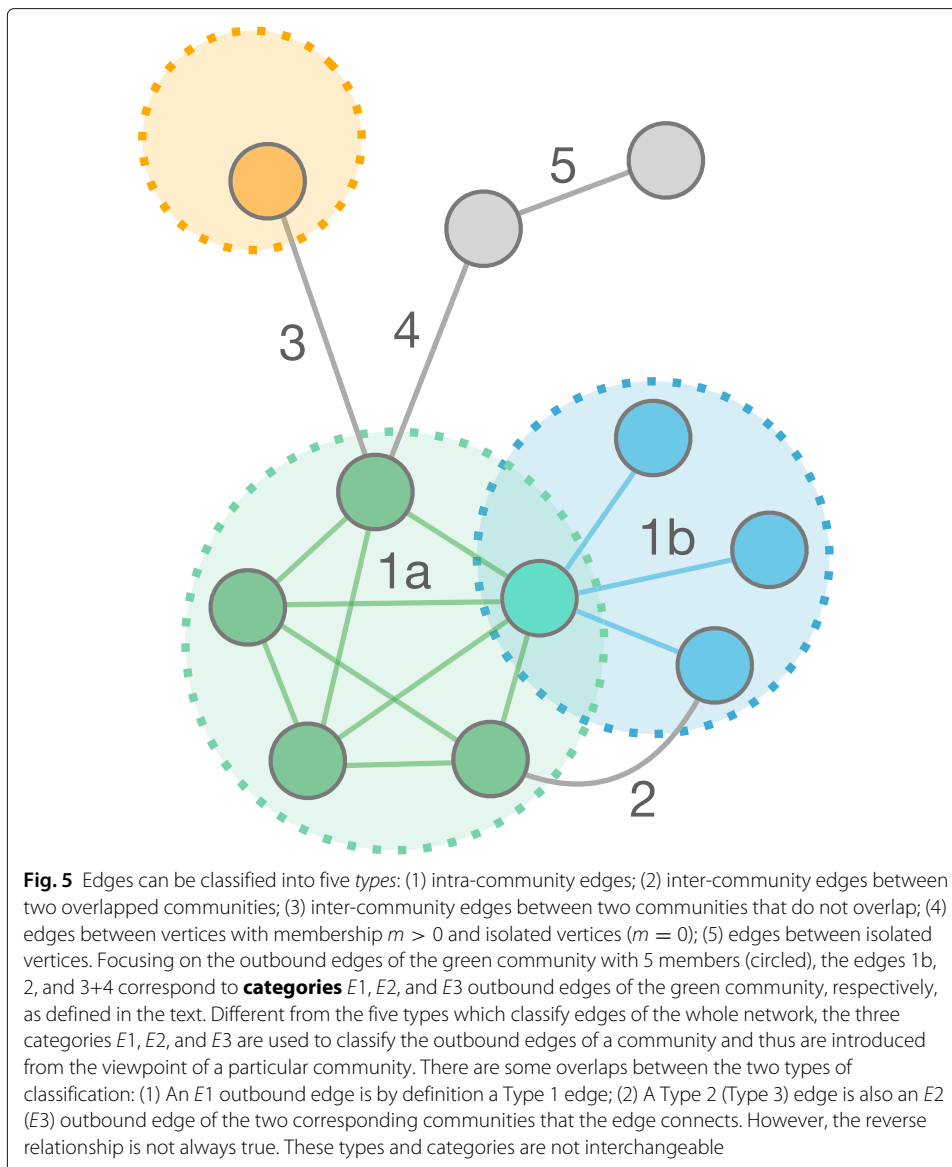
where  $k_{v,C}^{\text{int}}$  ( $k_{v,C}^{\text{out}}$ ) denotes the number of a vertex  $v$ 's edges that go inside (outside) the community  $C$ . The summations are over all  $n_C$  vertices in the community. Note that each internal edge is counted twice as both ends are within the community and each outbound edge is included only once. Figure 4 shows that the number of outbound edges of a community is not only greater, but often many times greater than the number of internal edges. More than 99% of the 2.9 million communities have more outbound edges than internal edges, in contrast to the traditional notion.

To investigate into the network structure, we focused on the outbound edges and classified them into 3 categories (see Fig. 5) as

- E1: outbound edges from a member to a neighbor community to which the member also belongs;
- E2: outbound edges from a member to a neighbor community that the member does not belong to;
- E3: outbound edges not to a neighbor community.

Their proportions  $e_1, e_2, e_3$ , with  $e_1 + e_2 + e_3 = 1$ , are calculated for each community. Figure 6 shows the histograms. Typically, the edges to a neighbor community are usually through the common member(s) of the two communities as  $e_1$  is much greater than  $e_2$ . In addition, a significant proportion of outbound edges go to neighbor communities. In Sina Weibo, most communities (red region) have  $e_1 + e_2 \approx 0.5$ . It means that  $\sim 50\%$  outbound edges are due to the vertices' multi-membership and communities

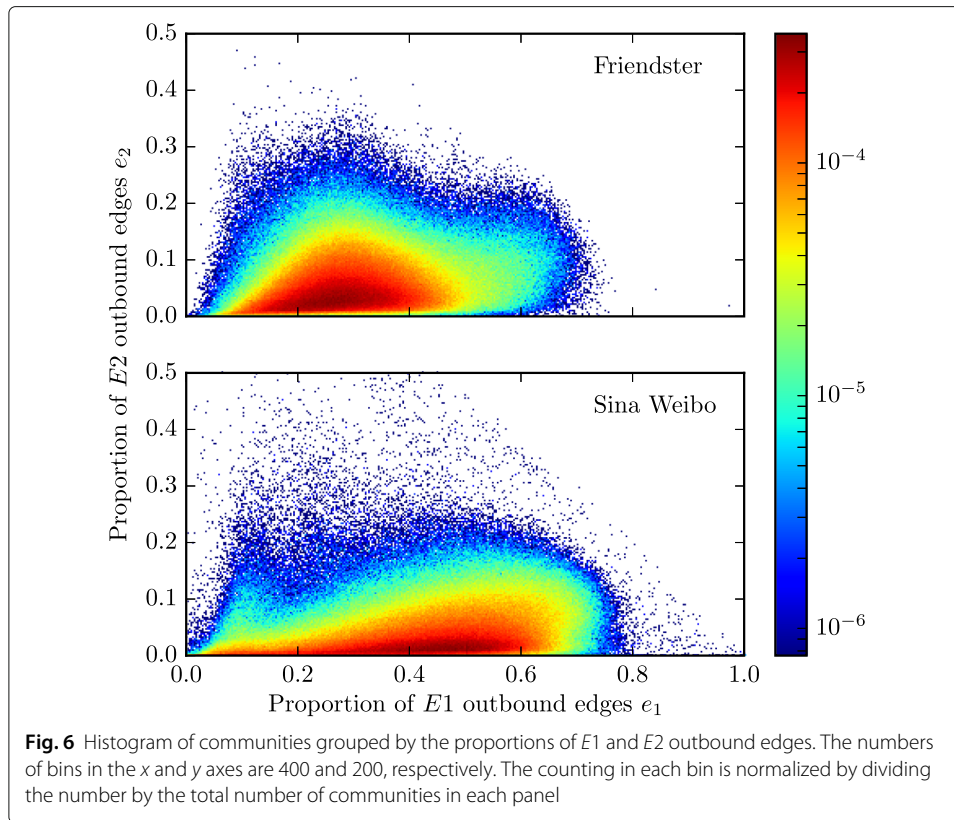




are densely connected to their neighbor communities. Note that if a community's outbound edges were randomly connected to vertices in the network, most edges would be of category  $E_3$ .

**Characteristic 4.** How can communities ever be distinguished when each community overlaps with a significant number of others? The answer is that the overlap size between two communities is usually small, and the connection between them is mostly through the overlap. Table 2 lists the frequency of occurrence of the most common overlap sizes. Out of 232M (millions) overlaps among the 2.9M detected communities, more than 80% are of just a single vertex. Figure 7 shows the actual structure of two detected communities. The outbound edges from community A (left) to its neighbor community B are highly organized through the overlap. Members of B usually only know the overlapping part of A, and vice versa. The overlapped vertex serves as the sole bridge and plays a unique role





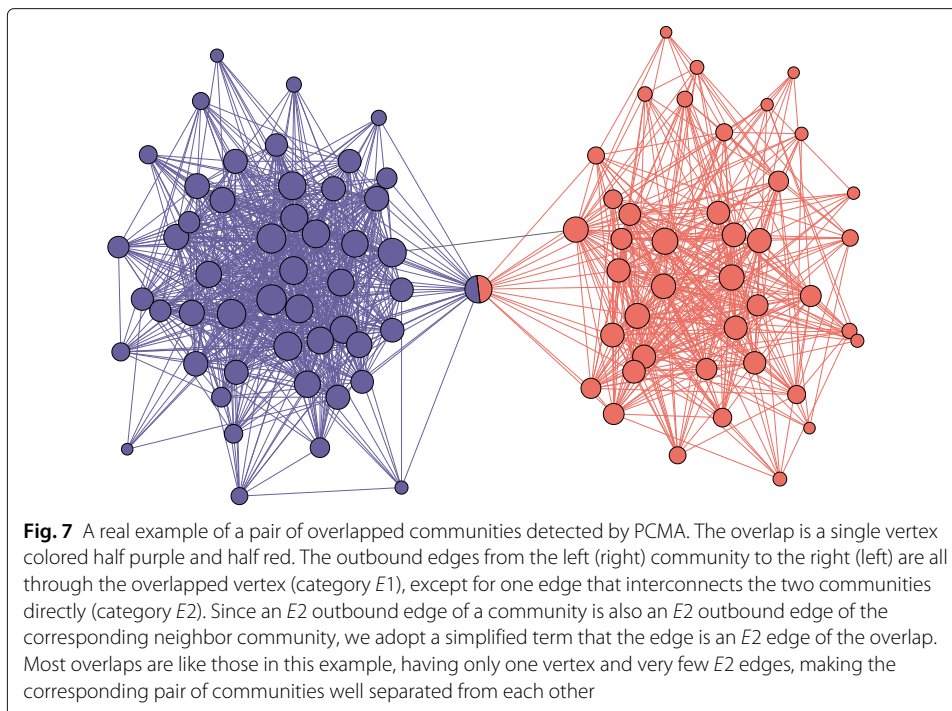
in passing information between the communities. Yet, there may exist some  $E2$  edges between the communities. In social networks, they are possibly due to the common member introducing members of the two communities to know each other. In Fig. 6,  $e_2$  is below 10% or even 5% for most communities and far less than  $e_1$ . It is the small proportion of  $E2$  edges that facilitates the easy separation of communities. The proportion  $e_2$  is thus an indicator of the clearness of the boundary between a community and its neighbor communities. We checked every pair of overlapped communities on  $E2$  edges. Results are listed in Table 2. For 37.8% (Sina Weibo) and 30.1% (Friendster) of them, there is not even a single  $E2$  edge. The communities maintain a good separation from their surrounding despite each overlaps with a significant number of neighbor communities.

### Mesoscopic view of social network structure

For the 2.9 million detected communities, we can classify all the edges in the two social networks into 5 types (see the caption of Fig. 5). The results are given in Table 3. The number of Type 1 edges suggests that the communities account for 30 ~ 35% of the entire network in terms of edges. These communities, connected together by the huge

**Table 2** Distribution of overlaps as a function of size of overlaps and the number of  $E2$  edges between a pair of communities

Dataset	No. of overlaps	Overlap size (no. of vertices)				No. of $E2$ edges of an overlap			
		1	2	3	4	0	1	2	$\leq 5$
Sina Weibo	77 million	84.5%	8.3%	2.6%	1.3%	37.8%	10.5%	6.0%	64.0%
Friendster	155 million	86.1%	7.7%	2.4%	1.1%	30.1%	11.0%	6.7%	59.4%



number of overlaps, form an extremely dense and tight network by themselves. There are 10 ~ 20% of the edges further connecting the overlapped communities (Type 2). The total number of them is comparable to that of Type 1, but since they are distributed among the huge number of overlaps, each overlap shares only a very few such edges. For example, in Sina Weibo, the 117M Type 2 edges are distributed among 73M overlaps, on average only 1.6 per overlap. The numbers further confirm the structure shown in Fig. 7. The Types 1 and 2 edges, together occupying half of the entire network, form an immense network of communities that can be regarded as a hidden skeleton of the social network in the mesoscopic scale. The remaining half of the edges are outside the skeleton, mostly Type 3 or Type 4. The former are “long-range” weak ties connecting different parts of the skeleton, thus making the skeleton an even smaller world. Although the majority of the vertices are outside the skeleton, i.e., vertices with  $m = 0$ , the edges among them (Type 5) account for less than 10%. For Friendster it is only 1.5%. These vertices are possibly the inactive users in the two online social network services.

**Table 3** Classification of edges and vertices in social networks. M represents a million. The types of edges are defined as in Fig. 5

Dataset	Sina Weibo		Friendster	
Vertices ( $m > 0$ )	21.0M	26.5%	28.0M	42.7%
Vertices ( $m = 0$ )	58.3M	73.5%	37.6M	57.3%
Edge Type 1	363M	34.7%	531M	29.4%
Edge Type 2	117M	11.2%	333M	18.4%
Edge Type 3	204M	19.5%	644M	35.7%
Edge Type 4	273M	26.1%	271M	15.0%
Edge Type 5	91M	8.7%	28M	1.5%

The edge classification helps decompose the entire network and reveals a remarkably high proportion of the significantly overlapped communities. The proportion could be even higher if less tightly connected vertices are also accepted as communities. The immense size of the network of communities confirms its important role in social networks and invites in-depth analyses on the properties of the huge and dense skeleton of social networks.

### Rethinking the concept of overlapping community

The strong empirical evidence from the analyses of the two social networks contradicts what we usually think a community is and asks for a reconsideration of the concept of community. Despite a wide variety of definitions, most of them, if not all, share an intuitive idea: members of a community should have some sort of internal cohesion and good separation from the rest of the network. The problem is how the idea should be interpreted, especially what a good separation and the boundary of a community are about.

Many definitions and quality measures of a community interpret “good separation” as the less the  $k_c^{\text{out}}$  (or  $k_c^{\text{out}}/k_c^{\text{int}}$ ), the more definite is the community. Examples include the widely used *weak community*  $k_c^{\text{int}} > k_c^{\text{out}}$  (Radicchi et al. 2004), fitness function  $k_c^{\text{int}} / (k_c^{\text{int}} + k_c^{\text{out}})^\alpha$  (Lancichinetti et al. 2009; Goldberg et al. 2010), *conductance*  $k_c^{\text{out}} / (k_c^{\text{int}} + k_c^{\text{out}})$  and *network community profile* (Leskovec et al. 2009; Jeub et al. 2015), dynamic-based definitions such as random walk (Rosvall and Bergstrom 2008) and label propagation (Raghavan et al. 2007). We argue that comparing  $k_c^{\text{out}}/k_c^{\text{int}}$  is ineffective in large-scale networks, no matter for overlapping or disjoint communities. As shown in Figs. 4 and 6, there are more outbound edges than internal edges, even if we ignore the neighbor community edges produced by the multi-membership vertices. The point is that simply a larger value of  $k_c^{\text{out}}$  does not necessarily mean the community is less definite. Consider the case that an arbitrary large number of outbound edges of a community are randomly distributed in the whole network, the community is not really strongly connected to any part of the network as long as the network size  $n \gg k_c^{\text{out}}$ . This point has also been discussed in a recent review by Fortunato and Hric (2016). They suggested using edge probabilities instead of the number of edges. A member of a community should have a higher probability  $p_{\text{in}}$  to form edges with the other members than  $p_{\text{out}}$  with vertices outside the community. Recent studies on detectability transitions in the Stochastic Block Model (Decelle et al. 2011; Nadakuditi and Newman 2012; Radicchi 2013; 2014; 2018) found that  $p_{\text{in}} > p_{\text{out}}$  is insufficient to guarantee that the community is detectable. There exists a region  $0 < p_{\text{in}} - p_{\text{out}} < \Delta$  that, although the community structure exists, no algorithms are able to detect. It is generally difficult to infer the edge probability between each pair of vertices. A simplified way is to assume the edge probabilities within a community (to the outside) are the same and equal to the internal (outbound) edge density  $\delta_c^{\text{int}} = k_c^{\text{int}} / [n_c(n_c - 1)]$ , ( $\delta_c^{\text{out}} = k_c^{\text{out}} / [n_c(n - n_c)]$ ), where  $n$  and  $n_c$  are the network and community sizes, respectively. However for large networks  $n \gg n_c$ , usually  $\delta_c^{\text{out}} \rightarrow 0$ , making the definition  $\delta_c^{\text{int}} > \delta_c^{\text{out}}$  not useful.

The problem of  $k_c^{\text{out}}$  (and so of  $\delta_c^{\text{out}}$ ) is that it counts the outbound edges to the whole network and reports only a summed quantity. What really matters is not the number  $k_c^{\text{out}}$ , but where the  $k_c^{\text{out}}$  outbound edges are distributed. As discussed under Characteristic 4 of the overlapping pattern, a multi-membership vertex may contribute much to  $k_c^{\text{out}}$  without messing up the boundary between the community and its neighbors. On the contrary,

adding a number of outbound edges to a particular vertex outside is sufficient to change the boundary of the community. These two cases are due to the different distribution patterns of outbound edges:

- Outbound edges from the same member to vertices outside the community
- Outbound edges from different members to a particular vertex outside the community

For the first case, it does not matter how many outbound edges there are. For the second case, however, the fewer the better. A good definition of overlapping community should be able to distinguish between the two cases. A useful concept here, as discussed in Ref. (Xu and Hui 2018), is the  $f$ -core – a maximal connected subgraph in which each vertex is connected to equal to or more than a fraction  $f$  of the other vertices in the subgraph:

$$b_{v,c} \geq f, \quad \forall v \in C \tag{6}$$

with  $b_{v,c}$  being the belongingness of  $v$  to  $C$  as defined in Eq. (1). A vertex is acknowledged as a member of an  $f$ -core as long as the vertex has sufficient connections to the other members of the  $f$ -core. It is irrelevant whether it is connected to a large number of vertices outside the  $f$ -core. This property of  $f$ -core distinguishes the two cases of outbound edges successfully and allows a vertex to belong to multiple  $f$ -cores naturally. In contrast, the number-based counterpart called  $k$ -core, which requires each vertex to be a neighbor to at least  $k$  other vertices in the subgraph, is non-overlapping by definition. The “maximal connected subgraph” in the definition ensures all vertices outside the  $f$ -core having belongingness less than  $f$ , as defined in Eq. (6), except for the case that there does exist one vertex outside, but including it will result in some other member(s) of the  $f$ -core to be kicked out. The fraction  $f$  defines the boundary of the community. A problem is that there is no standard way to determine what value of  $f$  should be used. Communities in social networks often show core-periphery structures (Csermerly et al. 2013; Rombach et al. 2014; Zhang et al. 2015) and have no definite boundaries. A large value of  $f$  extracts the core members of communities, and a small value results in more peripheral vertices being accepted as members. We are of the opinion that the belongingness  $b_{v,c}$  is a better way to describe vertex memberships instead of forcing a vertex to be either inside or outside of a community.

While the  $f$ -core is a good candidate, better definitions of overlapping community may still be possible. The key point is that the definition should take into account of the possibility of ubiquitous presence of multi-membership vertices:

- The proportion of multi-membership vertices may range from 0 ~ 100%,
- A vertex may belong to an arbitrary number of communities,

as revealed by data analysis. These are the causes of the significant overlaps among communities and a much greater number of outbound edges than internal edges.

### Summary and outlook

We studied the overlapping structure of 2.9 million communities detected by PCMA in the two huge online social networks. We found four main characteristics:

- Most members of a community have multiple memberships. They are everywhere, at the periphery or in the core.

- A community usually overlaps with a significant number of other communities, the number typically is several times its size.
- The number of outbound edges of a community is many times greater than the number of internal edges.
- Although communities overlap significantly, they remain relatively in good separation from each other. Most overlaps concern just one or sometimes two vertices.

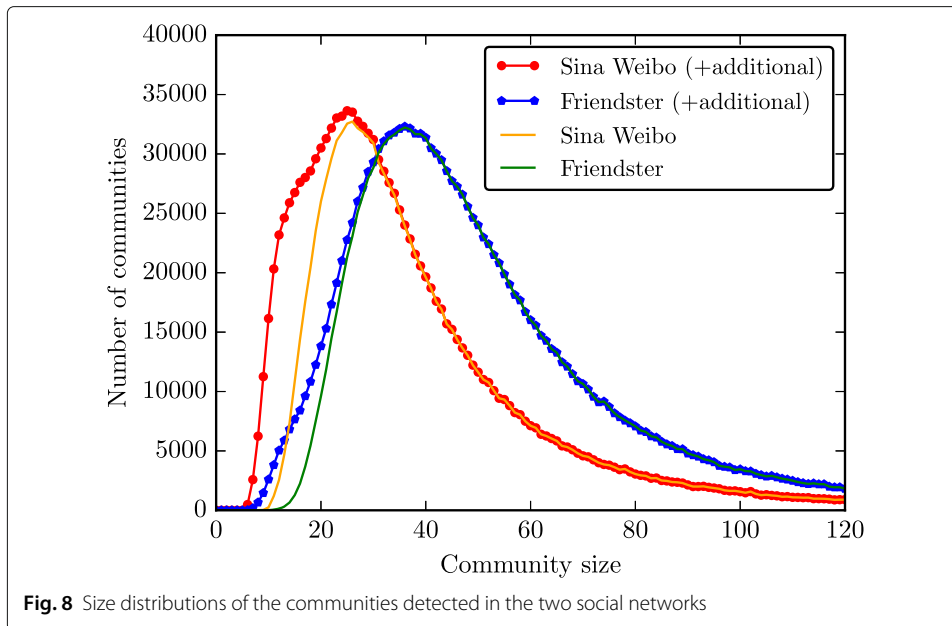
Note that PCMA does not impose any constraint or implications on the fraction of overlapping vertices in a community or the number of communities a vertex may have. It is also capable of detecting non-overlapping or slightly overlapping communities, as verified in the LFR benchmarking test (Lancichinetti et al. 2008) in Ref. (Xu and Hui 2018). The significant overlapping pattern found in the two empirical social networks asks for a rethinking of what the boundary of a community really is. We discussed several traditional interpretations and related issues, and suggested the  $f$ -core as a possible definition for overlapping community. Our study also showed a dense and tight network of communities, with the communities taking the role of vertices and the overlaps being the edges. Most overlaps are just of a single vertex. Each of these vertices plays a unique role in passing on information between the communities that it belongs to. This network of communities accounts for almost half of the entire network. It serves more studies on how its structural properties would couple to many phenomena in social dynamics.

As implied by the no-free lunch theorem for community detection that there can be no algorithm which is optimal for all possible community detection tasks (Peel et al. 2017), methods based on different approaches may reveal different aspects of the community structure. In fact, there is no standard answer to what a community really is, and it is largely unnecessary to enforce only one definition. This is especially the case for empirical networks. The important thing is that the communities detected satisfy the general notion of a community that they have internal cohesion and relatively clear boundaries. We verified that the 2.9 million communities analyzed in the present work have good separation among each other, and high values of the intra-community edge density, as shown in the Appendix.

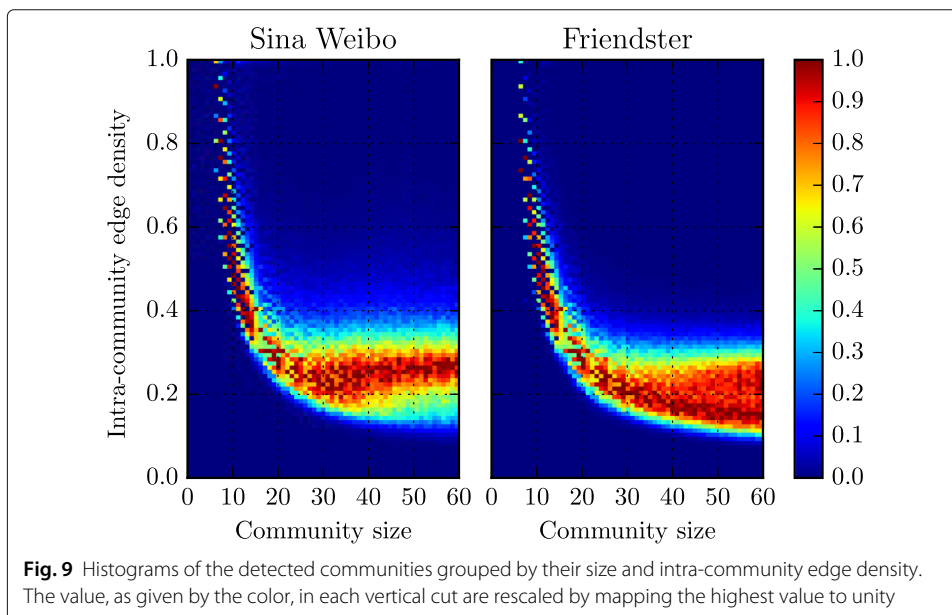
In conclusion, our empirical study unfolded new aspects of overlapping community. The results provided researchers with clues for designing effective detection algorithms, generative models, and benchmarks for overlapping communities, especially in social networks. We look forward to more empirical studies powered by new tools, to cross-check the present work and explore areas not covered by PCMA.

### **Appendix: Datasets**

For completeness, we describe the two social networks we analyzed. Table 1 gives the basic information. Sina Weibo is a directed network akin to Twitter. We focused on the embedded friendship network in which two connected individuals are following each other. Instead of sampling small subnetworks, we collected almost the whole giant component of the network, because the structural completeness of the sampled network is vital to the preservation of community structure, especially the overlapping pattern among communities. The network data of Friendster was downloaded from SNAP Datasets (Leskovec and Krevl 2014).



We detected about 1.3 and 1.6 million communities in the two networks with PCMA (Xu and Hui 2018). The algorithm is especially suitable for detecting communities in which the vertices have multiple memberships. Detailed information on the detection was reported in Ref. (Xu and Hui 2018). Specifically, the three steps of PCMA were discussed in Section 2 and the choice of the parameters for detecting the communities within PCMA were discussed in Appendix B of the paper. Using the symbols introduced in Subsection 2.2 in Ref. (Xu and Hui 2018), we used a harsh threshold  $l \geq 10$  to ensure that the detected communities are reliable (the larger the  $l$ , the more reliable the community). A drawback is that many small size communities were not included. In the present work, we add additional communities of which  $6 \leq l \leq 9$





and  $g > 3.0/l$ . The results are shown in Fig. 8. The latter condition ensures relatively high intra-community edge density of these communities, especially for those with low  $l$ . Figure 9 shows that all communities, including the newly added ones, have high values of intra-community edge density.

The large values of the proportion of intra-community and  $E2$  edges, as shown in Table 3, indicate that the number of communities we detected is close to the possible total number of communities in the two networks. However, it should be noted that there is no standard answer as to how many communities there are in a real network. We found that adding or removing the additional communities in the analyses only produces minor changes to the statistics. In particular, it does not change the characteristics of the overlapping pattern we discussed. The 2.9 million detected communities are believed to be adequate and representative.

#### Abbreviations

PCMA: Partial community merger algorithm

#### Acknowledgements

One of us (EHWX) gratefully acknowledges the support from the Research Grants Council of Hong Kong SAR Government through a Hong Kong PhD Fellowship.

#### Funding

The authors declare that no funding was received for the research reported.

#### Availability of data and materials

The network data of Friendster is available in the SNAP Datasets repository, <http://snap.stanford.edu/data>. The network data of Sina Weibo was collected from [weibo.com](http://weibo.com). Due to the current terms and conditions of the platform, the authors cannot distribute the collected data. The source code of PCMA is available at <https://github.com/hwxu/pcma>.

#### Authors' contributions

EHWX designed the research and performed the analyses. EHWX and PMH wrote the manuscript. Both authors read and approved the final manuscript.

#### Competing interests

The authors declare that they have no competing interests.

#### Publisher's Note

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

Received: 29 October 2018 Accepted: 24 April 2019

Published online: 31 May 2019

#### References

- Backstrom L, Boldi P, Rosa M, Ugander J, Vigna S (2011) Four Degrees of Separation. arXiv. 1111.4570v3
- Baumes J, Goldberg M, Krishnamoorthy M, Magdon-Ismael M, Preston N (2005) Finding communities by clustering a graph into overlapping subgraphs. In: Proceedings of the IADIS International Conference on Applied Computing. IADIS, Algarve. pp 97–104. <http://www.iadisportal.org/digital-library/cover-ac2005>
- Coscia M, Rossetti G, Giannotti F, Pedreschi D DEMON: a Local-First Discovery Method for Overlapping Communities. In: 2012 (ed). Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. ACM Press, New York. pp 615–623. <https://doi.10.1145/2339530.2339630>
- Csermely P, London A, Wu LY, Uzzi B (2013) Structure and dynamics of core/periphery networks. *J Complex Netw* 1(2):93–123. <https://doi.10.1093/comnet/cnt016>
- Decelle A, Krzakala F, Moore C, Zdeborová L (2011) Inference and Phase Transitions in the Detection of Modules in Sparse Networks. *Phys Rev Lett* 6:065701. <https://doi.10.1103/PhysRevLett.107.065701>
- Epasto A, Lattanzi S, Paes Leme R (2017) Ego-Splitting Framework: from Non-Overlapping to Overlapping Clusters. In: Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '17. ACM, New York. pp 145–154. <https://doi.10.1145/3097983.3098054>
- Ferrara E (2012) A large-scale community structure analysis in Facebook. *EPJ Data Sci* 1(1):9. <https://doi.10.1140/epjds9>
- Fortunato S (2010) Community detection in graphs. *Phys Rep* 486(3–5):75–174. <https://doi.10.1016/j.physrep.2009.11.002>
- Fortunato S, Hric D (2016) Community detection in networks: A user guide. *Phys Rep* 659:1–44. <https://doi.10.1016/j.physrep.2016.09.002>
- Goldberg M, Kelley S, Magdon-Ismael M, Mertsalov K, Wallace A (2010) Finding Overlapping Communities in Social Networks. In: 2010 IEEE Second International Conference on Social Computing (SocialCom). IEEE. pp 104–113. <https://doi.10.1109/SocialCom.2010.24>
- Hric D, Darst RK, Fortunato S (2014) Community detection in networks: Structural communities versus ground truth. *Phys Rev E* 90(6):062805. <https://doi.10.1103/PhysRevE.90.062805>



- Jebabli M, Cherifi H, Cherifi C, Hamouda A (2015) User and group networks on YouTube: A comparative analysis. In: 2015 IEEE/ACS 12th International Conference of Computer Systems and Applications (AICCSA). IEEE, pp 1–8. <https://doi.org/10.1109/AICCSA.2015.7507126>
- Jebabli M, Cherifi H, Cherifi C, Hamouda A (2018) Community detection algorithm evaluation with ground-truth data. *Physica A* 492:651–706. <https://doi.org/10.1016/j.physa.2017.10.018>
- Jeub LGS, Balachandran P, Porter MA, Mucha PJ, Mahoney MW (2015) Think locally, act locally: Detection of small, medium-sized, and large communities in large networks. *Phys Rev E* 91(1):012821. <https://doi.org/10.1103/PhysRevE.91.012821>
- Lancichinetti A, Fortunato S, Radicchi F (2008) Benchmark graphs for testing community detection algorithms. *Phys Rev E* 78(4):046110. <https://doi.org/10.1103/PhysRevE.78.046110>
- Lancichinetti A, Fortunato S, Kertész J (2009) Detecting the overlapping and hierarchical community structure in complex networks. *New J Phys* 11(3):033015. <https://doi.org/10.1088/1367-2630/11/3/033015>
- Lancichinetti A, Radicchi F, Ramasco JJ, Fortunato S (2011) Finding Statistically Significant Communities in Networks. *PLoS ONE* 6(4):e18961. <https://doi.org/10.1371/journal.pone.0018961>
- Leskovec J, Krevl A (2014) SNAP Datasets: Stanford large network dataset collection. <http://snap.stanford.edu/data>
- Leskovec J, Lang KJ, Dasgupta A, Mahoney MW (2009) Community Structure in Large Networks: Natural Cluster Sizes and the Absence of Large Well-Defined Clusters. *Internet Math* 6(1):29–123. <https://doi.org/10.1080/15427951.2009.10129177>
- Luccio F, Sami M (1969) On the Decomposition of Networks in Minimally Interconnected Subnetworks. *IEEE Trans Circuit Theory* 16(2):184–188. <https://doi.org/10.1109/TCT.1969.1082924>
- Lyu T, Bing L, Zhang Z, Zhang Y (2016) Efficient and Scalable Detection of Overlapping Communities in Big Networks. In: 2016 IEEE 16th International Conference on Data Mining (ICDM). IEEE, pp 1071–1076. <https://doi.org/10.1109/ICDM.2016.0138>
- Maiya AS, Berger-Wolf TY (2010) Sampling community structure. In: Proceedings of the 19th International Conference on World Wide Web, WWW '10. ACM, New York, pp 701–710. <https://doi.org/10.1145/1772690.1772762>
- Nadakuditi RR, Newman MEJ (2012) Graph Spectra and the Detectability of Community Structure in Networks. *Physical Review Letters* 108:188701. <https://doi.org/10.1103/PhysRevLett.108.188701>
- Peel L, Larremore DB, Clauset A (2017) The ground truth about metadata and community detection in networks. *Sci Adv* 3(5):e1602548. <https://doi.org/10.1126/sciadv.1602548>
- Radicchi F (2013) Detectability of communities in heterogeneous networks. *Physical Review E* 88(1):010801. <https://doi.org/10.1103/PhysRevE.88.010801>
- Radicchi F (2014) A paradox in community detection. *Europhysics Letters* 106(3):38001. <https://doi.org/10.1209/0295-5075/106/38001>
- Radicchi F (2018) Decoding communities in networks. *Phys Rev E* 97(2):022316. <https://doi.org/10.1103/PhysRevE.97.022316>
- Radicchi F, Castellano C, Cecconi F, Loreto V, Parisi D (2004) Defining and identifying communities in networks. *Proc Natl Acad Sci USA* 101(9):2658–2663. <https://doi.org/10.1073/pnas.0400054101>
- Raghavan UN, Albert R, Kumara S (2007) Near linear time algorithm to detect community structures in large-scale networks. *Phys Rev E* 76(3):036106. <https://doi.org/10.1103/PhysRevE.76.036106>
- Rees BS, Gallagher KB (2013) EgoClustering: Overlapping Community Detection via Merged Friendship-Groups. In: The Influence of Technology on Social Network Analysis and Mining. Springer Vienna, Vienna, pp 1–20. [https://doi.org/10.1007/978-3-7091-1346-2\\_1](https://doi.org/10.1007/978-3-7091-1346-2_1)
- Rombach MP, Porter MA, Fowler JH, Mucha PJ (2014) Core-Periphery Structure in Networks. *SIAM J Appl Math* 74(1):167–190. <https://doi.org/10.1137/120881683>
- Rosvall M, Bergstrom CT (2008) Maps of random walks on complex networks reveal community structure. *Proc Natl Acad Sci USA* 105(4):1118–1123. <https://doi.org/10.1073/pnas.0706851105>
- Sun H, Jie W, Sauer C, Ma S, Han G, Wang Z, Xing K (2017) A Parallel Self-Organizing Community Detection Algorithm Based on Swarm Intelligence for Large Scale Complex Networks. In: 2017 IEEE 41st Annual Computer Software and Applications Conference (COMPSAC). IEEE, pp 806–815. <https://doi.org/10.1109/COMPSAC.2017.31>
- Ugander J, Karrer B, Backstrom L, Marlow C (2011) The Anatomy of the Facebook Social Graph. arXiv. 1111.4503v1
- Watts DJ, Strogatz SH (1998) Collective dynamics of 'small-world' networks. *Nature* 393(6684):440–442. <https://doi.org/10.1038/30918>
- Xie J, Kelley S, Szymanski BK (2013) Overlapping community detection in networks: The state-of-the-art and comparative study. *ACM Comput Surv* 45(4):43:1–43:35. <https://doi.org/10.1145/2501654.2501657>
- Xu EHW (2016) Partial Community Merger Algorithm. <https://github.com/hwxu/pcma>
- Xu EHW, Hui PM (2018) Efficient detection of communities with significant overlaps in networks: Partial community merger algorithm. *Netw Sci* 6(1):71–96. <https://doi.org/10.1017/nws.2017.32>
- Yang J, Leskovec J (2013) Overlapping Community Detection at Scale: A Nonnegative Matrix Factorization Approach. In: Proceedings of the 6th ACM International Conference on Web Search and Data Mining, WSDM '13. ACM, New York, pp 587–596. <https://doi.org/10.1145/2433396.2433471>
- Yang J, Leskovec J (2014) Structure and Overlaps of Ground-Truth Communities in Networks. *ACM Trans Intell Syst Technol* 5(2):26:1–26:35. <https://doi.org/10.1145/2594454>
- Yang J, Leskovec J (2015) Defining and evaluating network communities based on ground-truth. *Knowl Inf Syst* 42(1):181–213. <https://doi.org/10.1007/s10115-013-0693-z>
- Zhang X, Martin T, Newman MEJ (2015) Identification of core-periphery structure in networks. *Phys Rev E* 91(3):032803. <https://doi.org/10.1103/PhysRevE.91.032803>