

Empirical comparison of network sampling techniques

Neli Blagus, Lovro Šubelj, Marko Bajec

University of Ljubljana, Faculty of Computer and Information Science, Ljubljana, Slovenia

Abstract

In the past few years, the storage and analysis of large-scale and fast evolving networks present a great challenge. Therefore, a number of different techniques have been proposed for sampling large networks. In general, network exploration techniques approximate the original networks more accurately than random node and link selection. Yet, link selection with additional subgraph induction step outperforms most other techniques. In this paper, we apply subgraph induction also to random walk and forest-fire sampling. We analyze different real-world networks and the changes of their properties introduced by sampling. We compare several sampling techniques based on the match between the original networks and their sampled variants. The results reveal that the techniques with subgraph induction underestimate the degree and clustering distribution, while overestimate average degree and density of the original networks. Techniques without subgraph induction step exhibit exactly the opposite behavior. Hence, the performance of the sampling techniques from random selection category compared to network exploration sampling does not differ significantly, while clear differences exist between the techniques with subgraph induction step and the ones without it.

Keywords: complex networks, network sampling, comparison of sampling techniques, subgraph induction, sampling accuracy

1. Introduction

Real-world networks are often very large and fast evolving. Therefore, not only their storage poses a problem, but their analysis and understanding present a great challenge. In the past few years, a number of different techniques have been proposed for sampling large networks to allow for their faster and more efficient analysis. However, some information about the network is lost through sampling, thus it is of key importance to understand the changes in the network structure introduced by sampling.

Several studies on network sampling analyze the match between the original networks and their sampled variants [1, 2, 3], while only few of them focus on comparing the performance of different sampling techniques [4, 5, 6]. In general, network exploration techniques like random walk and forest-fire sampling approximate the original networks

Email addresses: neli.blagus@fri.uni-lj.si (Neli Blagus), lovro.subelj@fri.uni-lj.si (Lovro Šubelj), marko.bajec@fri.uni-lj.si (Marko Bajec)

more accurately than random node and link selection [4]. However, Ahmed et al. [7] proposed link selection with additional subgraph induction step, where the sampled network consists of randomly selected links (i.e. random link selection) and any additional links between their endpoints (i.e. subgraph induction). In this way, not only the performance of random link selection is improved, the proposed technique also outperforms several other sampling techniques.

In this paper, we apply subgraph induction also to random walk and forest-fire sampling. We consider ten real-world networks and analyze the changes of their properties introduced by sampling. We compare eight sampling techniques based on the match between the original and sampled networks. The results reveal that techniques with subgraph induction step tend to underestimate the degree and clustering distribution, while other techniques tend to overestimate both properties. Moreover, the techniques with subgraph induction overestimate the average degree and density of the original networks, in contrast, techniques without subgraph induction underestimate them. It appears that the performance of the sampling techniques from random selection category compared to network exploration sampling does not differ significantly, while clear differences exist between the techniques with subgraph induction step and the ones without subgraph induction.

The rest of the paper is structured as follows. In Section 2, we first present the background on network sampling and expose the sampling techniques used in the study. The results of the empirical analysis are reported and formally discussed in Section 3, while Section 4 concludes the paper and suggests directions for further research.

2. Network sampling

Let the network be represented by a simple undirected graph $G = (V, E)$, where V denotes the set of nodes ($n = |V|$) and E is the set of links ($m = |E|$). The goal of network sampling is to create a sampled network $G' = (V', E')$, where $V' \subset V$, $E' \subset E$ and $n' = |V'| \ll n$, $m' = |E'| \ll m$. The sample G' is obtained in two steps. In the first step, nodes or links are sampled using a particular strategy like random selection and network exploration sampling. In the second step, the sampled nodes and links are retrieved from the original network. The sampled network is called a subgraph of the original network, if it consists of sampled nodes or sampled links only. Otherwise, if sampled nodes and all their mutual links are included in the sample or the sample consists of sampled links and any additional links among their endpoints, the sampled network is called an induced subgraph of the original network (i.e., subgraph induction).

However, the previous studies have shown that the size n' of the sample affects the performance of network sampling. Expectedly, larger samples approximate the original networks more accurately [6]. Still, the sample size of 10 – 30% of the original network achieves the balance between small sample and accurate approximation [4, 8, 6].

The changes of network properties introduced by sampling depend also on the adopted sampling technique, since different techniques are suitable for matching different sets of network properties. Thus, sampling techniques can be roughly divided into two categories: random selection and network exploration techniques. In the first category, nodes or links are included in the sample uniformly at random or proportional to some particular characteristic like degree. In the second category, the sample is constructed




Figure 1: Random selection techniques applied to a small toy network. Highlighted nodes and links represent the samples obtained by different techniques. (a) In random node selection, the sample consists of nodes selected uniformly at random and all their mutual links. (b) In random node selection by degree, the nodes are selected to the sample with probability proportional to their degree, while all their mutual links are included in the sample. (c) In random link selection, links are selected to the sample uniformly at random. (d) In random link selection with subgraph induction, the sample consists of randomly selected links (solid lines) and also any additional links between their endpoints (dashed lines).

by retrieving a neighborhood of a randomly selected seed node using different strategies like breadth-first search, random walk and forest-fire.

2.1. Random selection

For the purpose of this study, we consider four techniques from the random selection category. We first adopt random node selection [4] (RNS), where the sample consists of nodes selected uniformly at random and all their mutual links (Fig. 1(a)). RNS accurately approximates the degree mixing [5] and preserves the relationship of transitivity and density between the original and sampled networks [6]. Moreover, it shows better performance on larger samples than on smaller [6]. Yet, RNS overestimates the degree and betweenness centrality exponent and fails to match the clustering coefficient [5], degree distribution [9] and the average path length [10] of the original network.

Furthermore, we adopt random node selection by degree [4] (RND), which improves the performance of RNS. Here, the nodes are selected randomly with probability proportional to their degrees and all their mutual links are included in the sample (Fig. 1(b)). RND matches in-degree and out-degree distributions and also spectral properties of the original network better than RNS [4]. Besides, it constructs samples with larger weakly connected component [6]. Nevertheless, despite a fully connected original network, both RNS and RND can construct a disconnected sampled network.

Next, we adopt random link selection [4] (RLS), where the sample consists of links selected uniformly at random (Fig. 1(c)). RLS matches well degree mixing [5] and the distribution of sizes of weakly connected components [4]. It constructs sparse samples and accurately approximates the average path length of the original network [11]. Yet, RLS fails to match most of other network properties [4]. Like RNS, RLS overestimates the degree and betweenness centrality exponent and underestimates the clustering coefficient [5].

We last adopt random link selection with subgraph induction [7] (RLI), which improves the performance of RLS. Here, the sample consists of links selected uniformly at random and any additional links between their endpoints (Fig. 1(d)). RLI outperforms several other techniques in matching the degree, path length and clustering coefficient distribution of the original networks [7]. It selects nodes with higher degree more likely




Figure 2: Network exploration techniques applied to a small toy network. Highlighted nodes and links represent the samples obtained by different techniques. (a) In random walk sampling, the sample consists of links, retrieved from a simulation of a random walker on the network starting at a randomly selected seed node. (b) In random walk sampling with subgraph induction, the sample consists of links selected with random walk sampling (solid lines) and also any additional links between their endpoints (dashed lines). (c) In forest-fire sampling, the broad neighborhood of a randomly selected seed node is retrieved using partial breadth-first search, where only a fraction of links is included in the sample on each step. (d) In forest-fire sampling with subgraph induction, the sample consists of links selected with forest-fire sampling (solid lines) and also any additional links between their endpoints (dashed lines).

than other random selection techniques, which increases the connectivity of the sample. Moreover, RLI is suitable for sampling large networks that can not fit into the main memory and can also be implemented as a technique for sampling streaming networks [11].

2.2. Network exploration

We consider four sampling techniques from the network exploration category (note that in the literature, this category of sampling techniques is also called topology based sampling [7], traversal based sampling [12] or link-trace sampling [13]). First, we adopt random walk sampling [4] (RWS), where the random walk is simulated on the network, starting at a randomly selected seed node (Fig. 2(a)). The sample consists of links, which are visited by a random walker and represents a connected subgraph of the original network. RWS outperforms random selection techniques in matching the transitivity [13], clustering coefficient distribution and spectral properties and also shows good performance on smaller samples [4]. Yet, RWS is biased towards selecting nodes with high degree [14] and fails to match the degree distribution [15].

Next, we adopt forest-fire sampling [4] (FFS). Here, a broad neighborhood of a randomly selected seed node is retrieved from partial breadth-first search (Fig. 2(c)). The number of links sampled on each step is selected from a geometric distribution with mean $p/(1-p)$, where p is set to 0.7 [4]. Thus, on average 2.33 links are included in the sample on each step. FFS matches well spectral properties [4] and together with RWS shows the best overall performance among several techniques [4]. However, FFS fails to match the path length and clustering coefficient of the original networks [11].

Moreover, we apply subgraph induction step to random walk and forest-fire sampling, which we term random walk sampling with subgraph induction (RWI) and forest-fire sampling with subgraph induction (FFI). Here, the sample consists of links, sampled with random walk (Fig. 2(b)) or forest-fire sampling (Fig. 2(d)), while any additional links among the endpoints of sampled links are also included in the sample. To the best of our knowledge, RWI has not been analysed in any of the previous studies. On the other hand, FFI shows worse performance than RLI in matching the path length,

degree and clustering distributions [11]. Still, the performance of FFI has not yet been compared to a larger set of sampling techniques.

3. Analysis and discussion

In the following sections, we present the adopted framework for comparison of sampling techniques (Section 3.1), report the results of the empirical analysis and formally discuss the findings (Section 3.2).

3.1. Statistical comparison

Sampling techniques are compared through four network properties (see Section 3.2) on ten real-world networks (see Table 1). For each property, we compute externally studentized residuals of the properties of sampled networks that measure the consistency of each sampling technique with the rest. We expose statistically significant inconsistencies between the techniques with two-tailed Student t -test [16] and calculate the residuals as introduced below (see also [17]).

Let x_{ij} denote some property value of j -th network for i -th sampling technique. Externally studentized residual \hat{x}_{ij} is calculated as

$$\hat{x}_{ij} = \frac{x_{ij} - \hat{\mu}_{ij}}{\hat{\sigma}_{ij} \sqrt{1 - \frac{1}{N}}}, \quad (1)$$

where N is the number of sampling techniques ($N = 8$ in our case), $\hat{\mu}_{ij}$ is the mean value and $\hat{\sigma}_{ij}$ is the standard deviation of the considered property. The mean value and standard deviation are computed for all sampling techniques, excluding the observed one

$$\hat{\mu}_{ij} = \sum_{k \neq i} \frac{x_{kj}}{(N - 1)}, \quad (2)$$

$$\hat{\sigma}_{ij}^2 = \sum_{k \neq i} \frac{(x_{kj} - \hat{\mu}_{ij})^2}{(N - 2)}. \quad (3)$$

Assuming that the errors in x are independent and normally distributed, the residuals \hat{x} have Student t -distribution with $N - 2$ degrees of freedom. Statistically significant inconsistencies between the sampling techniques are revealed by two-tailed Student t -test [16] at P -value of 0.05, rejecting the null hypothesis that the values of the considered property are consistent across the sampling techniques.

3.2. Empirical analysis

The empirical analysis is performed on ten social and information networks. Their main characteristics are presented in Table 1. Due to a large number of networks considered, the detailed description is omitted. Networks are considered to be undirected, although some of them are directed. For each network, we perform 100 realizations of each sampling technique, where we consider sample sizes of 15% of the original networks as suggested in [4, 6]. For each run of the network exploration techniques, the sample was constructed from a new randomly selected seed node.

Table 1: Real-world networks considered in the study.

Network	Description	Nodes	Links	Average degree	Clustering coefficient	Density
<i>ca-hep</i>	<i>High E. Phys.</i> collaboration [19]	12,008	237,010	39.5	0.660	3.3×10^{-3}
<i>ca-astro</i>	<i>Astro Phys.</i> collaboration [19]	18,772	396,160	42.2	0.318	2.2×10^{-3}
<i>cit-hep</i>	<i>High E. Phys.</i> citation [20]	27,240	342,437	25.1	0.120	9.2×10^{-4}
<i>brightkite</i>	<i>Brightkite</i> friendship [21]	58,228	214,078	7.4	0.111	1.3×10^{-4}
<i>slashdot</i>	<i>Slashdot</i> friendship [22]	82,168	948,464	23.1	0.024	2.8×10^{-4}
<i>flickr</i>	<i>Flickr</i> images metadata [23]	105,938	2,316,948	43.7	0.402	4.1×10^{-4}
<i>ca-dblp</i>	DBLP collaboration [24]	317,080	1,049,866	6.6	0.306	2.1×10^{-5}
<i>nd.edu</i>	Web graph of <i>nd.edu</i> [25]	325,729	1,497,134	9.1	0.097	2.8×10^{-5}
<i>youtube</i>	<i>Youtube</i> friendship [26]	1,134,890	2,987,624	5.2	0.006	4.5×10^{-6}
<i>road-tx</i>	Texas road network [22]	1,379,917	1,921,660	5.5	0.060	3.9×10^{-6}

We analyze different properties of the original and sampled networks, including degree distribution (probability distribution of degrees of all nodes), distribution of clustering coefficient (probability distribution of the proportions of connected neighbors of each node [18]), average degree (average number of neighbors of nodes over the whole network), and density (the ratio of existing links to all possible links). We compare sampling techniques based on the match between the original and sampled networks; Section 3.2.1 reports the results for degree and clustering distributions, while in Section 3.2.2 we analyze the average degree and density.

3.2.1. Degree and clustering coefficient distributions

We first analyze the performance of sampling techniques based on the match of the degree and clustering coefficient distributions between the original and sampled networks. To compare the distributions of properties, we use Kolmogorov-Smirnov D -statistics, which is commonly used in similar studies [4, 11, 6]. Kolmogorov-Smirnov test checks the null hypothesis that the distributions of property of the original network and its sampled variant are the same, while the D -statistics measures the distance between the observed distributions. Based on the values of Kolmogorov-Smirnov D -statistics, we compare the adopted sampling techniques (see Section 3.1).

The comparison of sampling techniques based on the degree distribution is shown in Fig. 3. We observe that in most cases, the difference between the techniques with subgraph induction step (i.e., RNS, RND, RLI, RWI, FFI) and those without subgraph induction (i.e., RLS, RWS, FFS) is clear. The first group of techniques approximates the degree distribution of the original networks more accurately. In addition, the techniques with subgraph induction step improve the performance of the corresponding techniques without subgraph induction in the case of RLS and FFS. Since the subgraph induction increases the degrees of the nodes in the sample [7], this clearly contributes to better match the degree distribution between the original and sampled networks.

RNS and RND show comparable accuracy to network exploration techniques with subgraph induction, even though they construct disconnected samples. The latter indicates that the connectivity of the sample does not affect the match of the degree distribution. On the other hand, among the techniques without subgraph induction, RWS shows the best performance, which could be explained by its bias towards selecting high degree nodes and exploring densely connected parts of the network [4]. In contrast, the samples constructed by FFS are sparse, a large fraction of the nodes in the samples

has low degree, while the number of nodes with higher degree is underestimated [7]. Accordingly, FFS is the least accurate among all techniques.

Moreover, we compare the sampling techniques based on the clustering coefficient distribution. The results are presented in Fig. 4. The techniques show more comparable accuracy than in the case of the degree distribution. In general, the techniques with subgraph induction underestimate the clustering distribution, while others overestimate it. FFS and RWS prove as the best performing techniques. This indicates that in the case of matching the clustering distribution of the original networks, including all the links among sampled nodes is not the best choice. However, RNS shows the worst performance, which could be explained by its tendency to construct samples with a large number of low clustered nodes [7].

In summary, the sampling techniques with subgraph induction step underestimate the degree and clustering coefficient distributions of the original networks. On the other hand, the techniques without subgraph induction, overestimates both properties. The best overall performance is provided by RWI, which shows the most stable accuracy for matching both distributions on all considered networks. However, it appears that the properties of the original networks have relevant effect on the accuracy of the sampling techniques. For example, the techniques with subgraph induction match the degree distribution of the networks with larger average degree more accurately than when the average degree is lower and the techniques without subgraph induction also perform well. We aim to study that kind of relations in future work, since they should be observed in an even larger set of real-world networks to obtain statistically significant results.

3.2.2. Average degree and density

In the second part of the analysis, we study the performance of sampling techniques based on the match of the average degree and density between the original and sampled networks. We compare the properties between networks based on the actual values of properties. Furthermore, we modify Eq. (1) and (3) and use the true value of considered property of the original networks instead of the mean values $\hat{\mu}_{ij}$ (see Section 3.1) to compare the sampling techniques.

Fig. 5 shows the comparison of sampling techniques based on the average degree. We observe that the techniques with subgraph induction tend to overestimate the average degree, while others underestimate it. Among the techniques with induction step, only RNS underestimate the average degree. The samples constructed with RNS consist of a large fraction of low-degree nodes [7] and have a low average degree particularly in the smaller samples [5]. Therefore, RNS underestimates the average degree of the networks, especially compared to the techniques, which are biased in selecting nodes with higher degree (i.e., RND).

Lastly, we analyze the density of the networks. Fig. 6 shows the comparison of the sampling techniques based on the density. In general, all techniques overestimate the density, yet the techniques without subgraph induction step approximate the density more accurately than the techniques with subgraph induction. Additionally, techniques without subgraph induction improve the performance of the corresponding techniques with subgraph induction. In the previous study [27], we proved the power-law relationship between the size and density of real-world networks and their sampled variants. The results show that the network density decreases with its size. However, the techniques with subgraph induction do follow the latter relationship, while the techniques

without induction construct sparser samples than expected. Therefore, the accuracy of the sampling techniques based on the density is relative and depends on whether the samples should accurately match the density of the original networks or the relationship between the size and density of the original networks and their sampled variants should be preserved.

In addition to all above, we analyze the performance of the techniques with partial subgraph induction, where we include a different portion of mutual links between the sampled nodes in the sample (i.e., we randomly select 10% – 90% links from all possible links). The results are not described, since the techniques with partial subgraph induction did not improve the performance of the techniques with subgraph induction and they did not perform worse than techniques without subgraph induction.

In summary, the sampling techniques with subgraph induction step tend to overestimate the average degree and the density of the original networks. In contrast, the techniques without subgraph induction underestimate the average degree and approximate the density of the original networks more accurately than other techniques. Yet, the accuracy of sampling techniques does not depend only on the characteristics of the adopted technique, but also on the characteristics of the original networks. However, from exploration sampling techniques, the best performance is provided by RWI, while the best overall performance is provided by RNS, which shows very stable accuracy in matching considered properties.

4. Conclusion

In this paper, we analyze different real-world networks and study the changes of their properties introduced by network sampling. We consider eight sampling techniques with and without subgraph induction step and compare them based on the match of properties between the original networks and their sampled variants.

The results reveal that the techniques with subgraph induction underestimate the degree and clustering distribution, while they overestimate the average degree and density of the original networks. On the other hand, the techniques without subgraph induction step exhibit opposite behavior, since they overestimate the degree and clustering distribution and underestimate the average degree and density. The techniques without subgraph induction approximate the density most accurately. However, from the random selection category, random node selection shows the best overall performance, while from the network exploration category, the most stable accuracy is provided by random walk sampling with subgraph induction. Still, it appears that the performance of the sampling techniques from random selection category compared to network exploration sampling does not differ significantly, while clear differences exist between the techniques with subgraph induction step and the ones without subgraph induction. Following this, a suitable classification of network sampling techniques would be a class of techniques with subgraph induction and a class without it.

Some of the results suggest that the characteristics of the original networks affect the accuracy of the network sampling. Thus, our future work will mainly focus on analyzing relations between network properties and accuracy of the sampling techniques on a larger set of real-world networks. Besides, a prominent direction for further study is the analysis of the time and space efficiency of sampling techniques, since, for example, fitting even

a sampled network in the main memory becomes challenging with significant growth of real-world networks in the past few years.

Acknowledgment


This work has been supported in part by the Slovenian Research Agency *ARRS* within the Research Program No. P2-0359, by the Slovenian Ministry of Education, Science and Sport Grant No. 430-168/2013/91, and by the European Union, European Social Fund.

References


References

- [1] C. Hübler, H. P. Kriegel, K. Borgwardt, Z. Ghahramani, Metropolis algorithms for representative subgraph sampling, in: Proceedings of the 8th International Conference on Data Mining, IEEE, 2008, pp. 283–292.
- [2] H. Sethu, X. Chu, A new algorithm for extracting a small representative subgraph from a very large graph, e-print arXiv:1207.4825.
- [3] N. Blagus, L. Šubelj, G. Weiss, M. Bajec, Sampling promotes community structure in social and information networks, *Physica A* 432 (2015) 206–215.
- [4] J. Leskovec, C. Faloutsos, Sampling from large graphs, in: Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, ACM, 2006, pp. 631–636.
- [5] S. H. Lee, P. J. Kim, H. Jeong, Statistical properties of sampled networks, *Phys. Rev. E* 73 (1) (2006) 016102.
- [6] N. Blagus, L. Šubelj, M. Bajec, Assessing the effectiveness of real-world network simplification, *Physica A* 413 (2014) 134–146.
- [7] N. Ahmed, J. Neville, R. R. Kompella, Network sampling via edge-based node selection with graph induction, Tech. rep., Purdue University (2011).
- [8] C. Doerr, N. Blenn, Metric convergence in social network sampling, in: Proceedings of the 5th ACM workshop on HotPlanet, ACM, 2013, pp. 45–50.
- [9] M. P. H. Stumpf, C. Wiuf, R. M. May, Subnets of scale-free networks are not scale-free: sampling properties of networks, *P. Natl. Acad. Sci. USA* 102 (12) (2005) 4221–4224.
- [10] N. K. Ahmed, J. Neville, R. Kompella, Reconsidering the foundations of network sampling, in: Proceedings of the 2nd Workshop on Information in Networks, 2010.
- [11] N. K. Ahmed, J. Neville, R. Kompella, Network sampling: from static to streaming graphs, e-print arXiv:11211.3412.
- [12] P. Hu, W. C. Lau, A survey and taxonomy of graph sampling, e-print arXiv:1308.5865.
- [13] A. S. Maiya, T. Y. Berger-Wolf, Benefits of bias: towards better characterization of network sampling, in: Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, ACM, 2011, pp. 105–113.
- [14] L. Lovász, Random walks on graphs: A survey, *Combinatorics: Paul Erdős is eighty 2* (1) (1993) 1–46.
- [15] M. Kurant, A. Markopoulou, P. Thiran, On the bias of BFS, in: Proceedings of the 22nd International Teletraffic Congress, IEEE, 2010, pp. 1–8.
- [16] R. D. Cook, S. Weisberg, Residuals and influence in regression, New York: Chapman and Hall, 1982.
- [17] L. Šubelj, D. Fiala, M. Bajec, Network-based statistical comparison of citation topology of bibliographic databases, *Sci. Rep.* 4 (2014) 6496.
- [18] D. J. Watts, S. H. Strogatz, Collective dynamics of "small-world" networks, *Nature* 393 (6684) (1998) 440–442.
- [19] J. Leskovec, J. Kleinberg, C. Faloutsos, Graph evolution: Densification and shrinking diameters, *ACM Trans. Knowl. Discov. Data* 1 (1) (2007) 1–40.
- [20] KDD Cup '03, <http://www.cs.cornell.edu/projects/kddcup/> (2013).
- [21] E. Cho, S. A. Myers, J. Leskovec, Friendship and mobility: user movement in location-based social networks, in: Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, ACM, 2011, pp. 1082–1090.

- [22] J. Leskovec, K. J. Lang, A. Dasgupta, M. W. Mahoney, Community structure in large networks: Natural cluster sizes and the absence of large well-defined clusters, *Internet Math.* 6 (1) (2009) 29–123.
- [23] J. McAuley, J. Leskovec, Learning to discover social circles in ego networks, in: *Advances in Neural Information Processing Systems* 25, 2012, pp. 548–556.
- [24] J. Yang, J. Leskovec, Community-affiliation graph model for overlapping network community detection, in: *Proceedings of the 12th International Conference on Data Mining, IEEE*, 2012, pp. 1170–1175.
- [25] R. Albert, H. Jeong, A.-L. Barabási, Internet: Diameter of the world-wide web, *Nature* 401 (6749) (1999) 130–131.
- [26] J. Yang, J. Leskovec, Defining and evaluating network communities based on ground-truth, *Knowledge and Information Systems* 42 (1) (2015) 181–213.
- [27] N. Blagus, L. Šubelj, M. Bajec, Self-similar scaling of density in complex real-world networks, *Physica A* 391 (8) (2012) 2794–2802.




(a) Degree distribution




(b) Degree distribution

Figure 3: Statistical comparison of the degree distribution of the original networks and their sampled variants obtained by different sampling techniques. We show externally studentized residuals that measure the consistency of each sampling technique with the rest and expose statistically significant inconsistencies between the techniques with two tailed Student t -test at P -value of 0.05 (see (a); shaded region corresponds to 95% confidence intervals). Merely for complete understanding, we present average externally studentized residuals over all networks (see (b); the star marker indicate the residuals for the original networks, while error bars display standard deviations). Notice that techniques with subgraph induction step (full markers) approximate the degree distribution more accurately than the techniques without subgraph induction (empty markers).

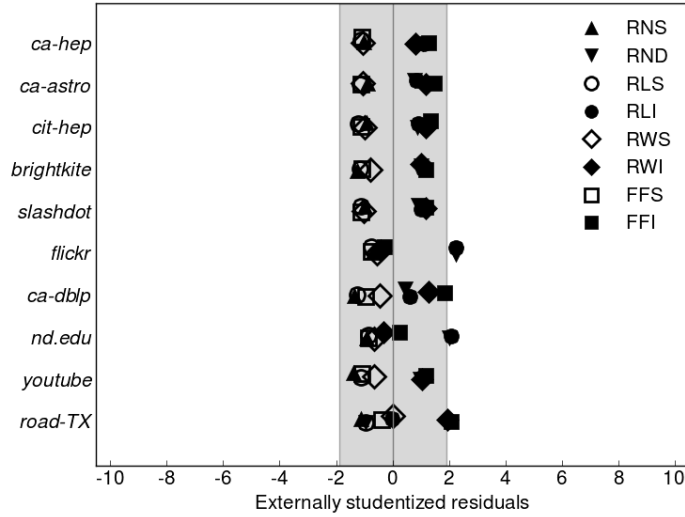


(a) Clustering coefficient distribution




(b) Clustering coefficient distribution

Figure 4: Statistical comparison of the clustering coefficient distribution of the original networks and their sampled variants obtained by different sampling techniques. We show externally studentized residuals that measure the consistency of each sampling technique with the rest and expose statistically significant inconsistencies between the techniques with two tailed Student t -test at P -value of 0.05 (see (a); shaded region corresponds to 95% confidence intervals). Merely for complete understanding, we present average externally studentized residuals over all networks (see (b); the star marker indicate the residuals for the original networks, while error bars display standard deviations). Notice that techniques with subgraph induction step (full markers) improve the performance of the corresponding techniques without subgraph induction (empty markers).




(a) Average degree




(b) Average degree

Figure 5: Statistical comparison of the average degree of the original networks and their sampled variants obtained by different sampling techniques. We show externally studentized residuals that measure the consistency of each sampling technique with the rest and expose statistically significant inconsistencies between the techniques with two tailed Student t -test at P -value of 0.05 (see (a); shaded region corresponds to 95% confidence intervals). Merely for complete understanding, we present average externally studentized residuals over all networks (see (b); the star marker indicates the residuals for the original networks, while error bars display standard deviations). Notice that techniques with subgraph induction step (full markers) tend to overestimate the average degree, while the techniques without subgraph induction (empty markers) underestimate it.



(a) Density



(b) Density

Figure 6: Statistical comparison of the density of the original networks and their sampled variants obtained by different sampling techniques. We show externally studentized residuals that measure the consistency of each sampling technique with the rest and expose statistically significant inconsistencies between the techniques with two tailed Student t -test at P -value of 0.05 (see (a); shaded region corresponds to 95% confidence intervals). Merely for complete understanding, we present average externally studentized residuals over all networks (see (b); the star marker indicates the residuals for the original networks, while error bars display standard deviations). Notice that the techniques without subgraph induction step (empty markers) approximate the density more accurately than the corresponding techniques with subgraph induction (full markers). 14