A Statistical Test of Heterogeneous Subgraph Densities to Assess Clusterability

Pierre Miasnikof¹, Liudmila Prokhorenkova^{2,3}, Alexander Y. Shestopaloff⁴, and Andrei Raigorodskii^{2,3}

 ¹ University of Toronto, Dept. of Chemical Engineering and Applied Chemistry, Toronto, Canada
 ² Moscow Institute of Physics and Technology, Dolgoprudny, Russia
 ³ Yandex, Moscow, Russia

⁴ The Alan Turing Institute, London, United Kingdom

Abstract. Determining if a graph displays a clustered structure prior to subjecting it to any cluster detection technique has recently gained attention in the literature. Attempts to group graph vertices into clusters when a graph does not have a clustered structure is not only a waste of time but will also lead to misleading conclusions. To address this problem, we introduce a novel statistical test, the δ -test, which is based on comparisons of local and global densities. Our goal is to assess whether a given graph meets the necessary conditions to be meaningfully summarized by clusters of vertices. We empirically explore our test's behavior under a number of graph structures. We also compare it to other recently published tests. From a theoretical standpoint, our test is more general, versatile and transparent than recently published competing techniques. It is based on the examination of intuitive quantities, applies equally to weighted and unweighted graphs and allows comparisons across graphs. More importantly, it does not rely on any distributional assumptions, other than the universally accepted definition of a clustered graph. Empirically, our test is shown to be more responsive to graph structure than other competing tests.

1 Introduction

Graph clustering has been very well and very broadly covered in the literature (e.g., [30,8,14,2,15,29]). Nonetheless, objective measures of a graph's suitability for clustering are lacking. There are multiple graph clustering and community detection techniques in the literature and they all rely on the unmentioned assumption that the graph has a clustered (community) structure, to begin with. However, not all graphs have this type of structure. It is important to determine if a graph is a good candidate for clustering, before any vertex-grouping effort is undertaken. Applying clustering techniques to a graph that does not have a clustered structure is not only a waste of time; it will inevitably lead to misleading conclusions.

In this article, our objective is to provide a test to determine if a graph meets the prerequisite conditions for it to have a meaningful cluster structure or if it displays no significant clustered structure. Here, our goal is not to identify graph clusters and their component vertices or assess the quality of the clusters identified by an algorithm. Our goal is to assess whether a given graph meets the necessary conditions to be meaningfully summarized by grouping its vertices into clusters; whether its vertices are concentrated within densely inter-connected subgraphs. In other words, our aim is to provide a technique for accurately answering the question recently posed by Chiplunkar et al. [9], "(...) given access to a graph G = (V,E), can we quickly determine whether the graph can be partitioned into a few clusters with good inner conductance (...)?", using statistical sampling and hypothesis testing.

We want to assess whether a graph is more likely to be partition-able, like the graph shown in Figure 1a, or if it has a cluster-less structure as in Figure 1b.⁵



(a) Graph with Clustered Structure (b) Graph without Clustered Structure

Fig. 1. Graphs Displaying Clustered and Unclustered Structure

It is important to note that we are not looking to determine if all vertices are grouped within clusters, but if a significant amount of them may belong to separate clusters. We want to determine if clusters might be a good model for the graph, if clusters might help us properly partition the graph in a meaningful way. For example, a graph may have a small number of clusters that contain a small fraction of vertices while the rest of the graph is randomly connected; or a graph may have most of its vertices contained in densely connected clusters with a very small number of vertices that do not belong to clusters. In the first instance, we would not want to conclude that the graph displays a clustered structure, since clusters are uninformative features of the graph and are not useful indicators for its partitioning. Meanwhile, in the second instance, we would want to conclude that clusters are likely to offer a meaningful picture of a graph's structure.

With a view on large data sets, we develop a test that only relies on small samples from the graph and does not examine the graph in its entirety. Here, we

⁵ Also, note that in this article we assume undirected graphs with no self-loops.

follow the examples of many authors who have inferred various graph properties through sampling (e.g., [3,11,12,17,18,19]).

The rest of this article is organized as follows. We begin with a review of previous work in the field, describe and justify our methods, compare them to recent techniques from the literature and share results from our empirical tests. As a point of comparison, we also implement the techniques proposed by Gao and Lafferty [17,18] and empirically compare their performance to our method's performance.

2 Related Work

Our work is primarily inspired by the recent publications by Gao and Lafferty [17,18], who use sampling and statistical testing to determine if a network has community structure. We are also guided by the question, "(...) given access to a graph G = (V,E), can we quickly determine whether the graph can be partitioned into a few clusters with good inner conductance (...)?", which was recently posed by Chiplunkar et al. [9].

While a formal definition of "community structure" remains a topic of debate (e.g., [15,30]), virtually all authors agree a cluster (or community) is a subset of vertices that exhibit a high-level of interconnection between themselves and a low-level of connection to vertices in the rest of the graph [14,27,28,33] (we quote these authors, but their definition is very common across the literature).

The graph testing (through sampling) literature is very rich. Many authors have proposed tests for various graph properties other than the existence of clusters (e.g., [3,11,19,20,25]). Tests for clustering have also been studied in the past. For example, Arias-Castro and Verzelen [4,32] set up clustering as a hypothesis test. They use the Erdős-Rényi random (ER) graph [1,13] as a null model, but their alternative hypothesis is the existence of only one dense subgraph. Czumaj et al. [10] describe a test of clustering, but they impose a restrictive $\kappa - \phi$ model as their benchmark. Bikel and Sarkar develop tests for graph spectra, in order to determine if a graph follows the ER model [7].

As mentioned earlier, Gao and Lafferty [17,18] propose clustering tests, but their tests focus on vertex triplets and asymptotic properties. More recently, Jin et al. [23] raise the point that testing for clustering remains a non-trivial problem and propose potentially (computationally) costly tests based on the number of paths and cycles of fixed length. Also, very recently Chiplunkar et al. [9] propose a test of clustering based on a restrictive $\kappa - \phi$ model, while He et al. [22] highlight the problem as still being unresolved.

Unfortunately, most of these authors' approaches are restrictive and rely on rigid models for their hypotheses tests. Indeed, Ugander et al. [31] claim that incomplete non-empty k-node subgraphs are infrequent not just in social networks, but there exist mathematical reasons why such subgraphs are infrequent in general settings. They proved a theoretical upper bound on their probability of occurrence. Later, Yin and Leskovec discussed the need to look at more complex structures, in order to gain an understanding of complex network structures.

tures [34]. Finally, it is important to remember that closed triplet frequencies and clustering coefficients may be higher than under the ER hypothesis, while the network does not exhibit clustering [16].

The choice of null hypothesis is also a topic of debate in the literature. In addition to most of the authors above who focus on clustering, Elenberg et al. [12] approximate 3-profile counts by sampling and use the ER graph as a null model. Unfortunately, the ER random graph model has been described as overly simplistic to model real-world networks (e.g., [1,5,6,16]).

Gao and Lafferty [18] use the ER model to prove a central limit theorem for their estimators of wedges and triangles. However, as pointed out by those same authors [17], the ER model may be an insufficient null hypothesis. Meanwhile, other authors (e.g., [23]) set up a more general null hypothesis of only one single cluster and an alternative hypothesis of more than one.

3 Methods

We develop a test whose successful fulfillment forms a necessary condition for a graph (network) to exhibit a clustered (community) structure. It is founded on statistical hypothesis testing; the reliability of the conclusions it leads to is transparent and rooted in statistical theory.

Unlike many authors' recent work on the topic, our technique does not rely on any assumptions on the structure of the null graph, beyond what is implied by the definition of an unclustered graph structure. In addition, our technique detects anomalies in density, not just specific sub-structures like triangles or other fixed-length structures. It also makes no assumptions on the number of clusters, as in the case of $\kappa - \phi$ models.

3.1 Underlying Assumptions and Densities

Our test only relies on the assumption that graphs composed of more than one cluster are characterized by groups of densely connected vertices with sparse connections between vertices belonging to different clusters. This assumption is a direct consequence of the fact that a cluster can be described as a subset of vertices that are densely connected to each other but only sparsely connected to the rest of the graph. Therefore, a clustered connection pattern translates into heterogeneous local densities across the graph. A clustered graph is characterized by pockets of highly connected nodes leading to high local densities that are significantly higher than the graph's overall density.

The links between density and clustered patterns of connectivity were shown in Miasnikof et al. [26]. Under such a pattern of connectivity, it is expected that the densities of induced subgraphs obtained by sampling vertices within a neighborhood will exhibit, on average, higher densities than the graph's global density. An example of this heterogeneous density is shown in Figure 2, where the graph's global density is K = 0.43, while each of its constituent cluster (subgraph) densities are $k_1 = 0.83, k_2 = 1$.



Fig. 2. Heterogeneous Densities

On the other hand, if a graph has no meaningful cluster structure, we should not observe a significant number of vertices falling into densely inter-connected subgraphs which are sparsely connected to other vertices on the graph, as seen in Figure 1b. A non-clustered structure implies that any two vertices have a roughly uniform probability of being connected. This homogeneity can be statistically inferred by the absence of a statistically significant difference between the mean density of a sample of randomly chosen induced subgraphs and the graph's global density.

Unlike much of the current work on this topic (e.g., [4,7,12,32]), our null model does not rely on the assumption of an ER graph and does not make any assumptions on the number of clusters (e.g., [9]). Our proposed technique does not impose any model on the null other than its lack of clustering reflected in a homogeneous density across the graph. Our null graph is simply a general graph with a near constant density, on average. This null hypothesis includes all graphs without clustered structure, not only ER graphs, but any other type of non-clustered graph. The configuration model graph and a graph which contains a small number of clusters while most of the edges are randomly connected are both examples of such general unstructured graphs. Our simulated data sets described in Section 3.6 and our empirical tests in Section 3.7 include both these examples.

3.2 (Sub)Graph Densities and the δ Statistic

Consider a graph G = (V, E), with |V| = N. First, we randomly sample a subset of S = sN vertices, where $s \in (0, 1]$ is a parameter which determines the proportion of total vertices we wish to sample. We call these vertices the root nodes.

Then, for each root node $i \in \{1, \ldots, S\}$, we take all its neighboring vertices and consider the local subgraph induced by them. Let E_i denote the set of edges in the local induced subgraph formed by the neighbors of node i and n_i denote the number of these neighbors. In the case of unweighted graphs, $|E_i|$ (or |E|when dealing with the full graph) represents the cardinality of the set (number of edges). In the case of weighted graphs, $|E_i|$ (or |E|) represents the sum of edge weights.

We compute the local induced subgraph density as

$$k_i = \frac{|E_i|}{0.5 \times n_i(n_i - 1)}$$

Then, the mean local density for S induced local subgraphs is computed as

$$\bar{K}_{\ell} = \frac{1}{S} \sum_{i=1}^{S} k_i \,.$$

At the graph level, the global density is computed as

$$K = \frac{|E|}{0.5 \times N(N-1)} \,.$$

Finally, we introduce the δ statistic, the normalized mean of k_i , as a measure of divergence between the global and mean local densities:

$$\delta = \frac{K_\ell}{K} - 1 \,.$$

With the help of Figure 3, we illustrate this sampling procedure and the computation of the δ statistic. We begin by sampling two vertices labelled v_1 and v_2 (we have S = 2) and compute the densities of the two local induced subgraphs (k_1, k_2) . Our δ statistic is then computed as follows:

$$\begin{split} k_1 &= \frac{|E_1|}{0.5 \times n_1(n_1 - 1)} = \frac{2}{0.5 \times 3 \times 2} = 0.67 \,, \\ k_2 &= \frac{|E_2|}{0.5 \times n_2(n_2 - 1)} = \frac{3}{0.5 \times 3 \times 2} = 1 \,, \\ \bar{K}_{\ell} &= \frac{1}{2}(k_1 + k_2) = \frac{1}{2}(0.67 + 1) = 0.84 \,, \\ K &= \frac{|E|}{0.5 \times N(N - 1)} = \frac{12}{0.5 \times 8 \times 7} = 0.43 \,, \\ \delta &= \frac{\bar{K}_{\ell}}{K} - 1 = \frac{0.84}{0.43} - 1 = 0.95 \,. \end{split}$$



K = 0.43

Fig. 3. Heterogeneous Densities

3.3 Statistical Hypotheses

On the basis of our assumptions on densities, we posit the following relationships between clustering and density:

- 1. If a graph does not display a clustered structure, then, on average, there should be no statistically significant difference between local densities and the graph's global density (homogeneity);
- 2. If a graph does display a clustered structure, then, on average, there should be a statistically significant difference between the densities of local induced subgraphs and the graph's global density (heterogeneity);
- 3. If a graph G_1 has a more strongly clustered structure than a graph G_2 , then its mean local density should also be significantly greater, in a statistical sense (relative heterogeneity).

In summary, the existence of meaningful clusters implies that, on average, local clustering should be significantly stronger than the graph's overall clustering. Here, we draw the reader's attention to the fact that our technique seeks to statistically determine the presence of locally strong atypical densities. Our approach is more general than the recent work which focuses on triangles or other fixed-length cycles or restrictive $\kappa - \phi$ models (e.g., [9,17,18,23]).

On the basis of the above-mentioned relationships, we formulate the following statistical hypotheses, which can be tested using the appropriate t-tests:

- Under the null hypothesis that a graph does not have a clustered structure, i.e., it has a uniform density, we test for $E(k_i) = K \Rightarrow E(\delta) = 0$ by using a one-tailed Student t-test;
- Alternatively, if the graph does have heterogeneous density, we expect $E(k_i) > K \Rightarrow E(\delta) > 0$, since it is greater than under the null;

- Under the null hypothesis that a graph G_1 has equally strong clustered structure as a graph G_2 , we test for $E(\delta_{G_1}) = E(\delta_{G_2})$, again by using a one-tailed Student t-test;
- Alternatively, if a graph G_1 has a more heterogeneous density structure than a graph G_2 , we expect $E(\delta_{G_1}) > E(\delta_{G_2})$.

3.4 Algorithm

As mentioned in the previous section, our δ test statistic can be used to answer two related but distinct questions:

- 1. Does a graph display heterogeneity in its density?
- 2. Does a graph G_1 have a more heterogeneous density than a graph G_2 ?

In the first question, we ask whether a graph meets the necessary condition to have a clustered structure. In the second, we ask if one graph meets this condition more strongly than another. These questions can be answered by following these steps:

- Sample S induced subgraphs, each containing n_i vertices $(i \in \{1, \ldots, S\})$;
- Compute the local densities $k_i = \frac{|E_i|}{0.5 \times n_i(n_i-1)}$ for each of the S subgraphs;
- Compute graph density $K = \frac{|E|}{0.5 \times N(N-1)}$, for graph G = (V, E) with |V| = N;
- Compute the mean of the local densities: $\bar{K}_{\ell} = \frac{1}{S} \sum_{i=1}^{S} k_i$;
- Normalize the mean and obtain our test statistic: $\delta = \frac{K_{\ell}}{K} 1;$
- In the case of a single-graph test, perform a one-tailed t-test; the null hypothesis is $E(\delta) = 0$, the alternative is $E(\delta) > 0$;
- In the case of a two-graph test, perform a two-sample (unpaired with unequal variance) one-tailed t-test on the δ statistics of graphs G_1, G_2 ; the null hypothesis is $E(\delta_{G_1}) = E(\delta_{G_2})$, the alternative is $E(\delta_{G_1}) > E(\delta_{G_2})$

3.5 Methodological Comparison to Other Recently Published Techniques

In Table 1, we compare key features of methodologies employed by recently published test techniques that have inspired our own work. In the first column, we display the features associated with the methods proposed by Chiplunkar et al. [9]. In the second, we display the features of Gao and Lafferty's two recently published tests [17,18]. Finally, in the last column we display the features of our own δ test.

In Table 1, we see that our test is more general and versatile than recently published tests. We make no distributional assumptions and impose no generative model on the graph. Our test applies to weighted and unweighted graphs equally and is rooted in statistical testing methodology.

Table 1. A Comparison of Recently Published Clustering Tests

	Chip et al.	G-L Tests	δ Test
Requires Number of Clusters as Input	Y	Ν	Ν
Assumes Generative Model	Y	Ν	Ν
Provides Statistical Confidence Intervals	N	Y	Y
Imposes Distribution on Null	NA	Y	N
Applies to Weighted and Unweighted Graphs	Y	Ν	Y

3.6 Empirical Examples

We compute δ for each graph and perform single graph t-tests for the following ten graphs. We also perform two-sample t-tests (with unequal variance), to compare each pair of graphs (our weighted graph is assessed separately). With the exception of our Stanford University SNAP "real-world" graph [24,35], all graphs were generated using the Networkx library [21]. Details of each graph's characteristics are provided below, in Table 2.

- Scenario 1: generated using Erdős-Rényi model (ER);
- Scenario 2: generated using degree-corrected stochastic block model (SBM);
- Scenario 3: (connected) caveman (CC);
- Scenario 4: generated using Barabási-Albert model with an out-degree of 3 (BA3);
- Scenario 5: generated using Barabási-Albert model with an out-degree of 5 (BA5);
- Scenario 6: generated using Watts-Strogatz model (WS);
- Scenario 7: generated using the configuration model with an underlying power law with exponent of 3 (CM);
- Scenario 8: a merger of two connected caveman graphs of 75 vertices each and an ER graph of 850 vertices (MD);
- Scenario 9: a "real-world" graph. Here, we converted the SNAP "Eu-core network" into an undirected graph with no self-loops (EUC);
- Scenario 10: a weighted (connected) caveman (CCw).

3.7 Empirical Results

The tables below contain the p-values for our one-sample and two-sample δ statistics. In Table 3, the diagonal entries are for the one-sample t-tests, while the off diagonal elements are for the two-sample graph vs graph t-tests. These tests were done on graph samples consisting of a set of randomly selected root nodes of size equal to 25% of all graph nodes.

As expected, our one-sample test does not reject the null hypothesis of homogeneous density for the ER and configuration model (CM) graphs, whose edges are randomly connected. The null is also not rejected in the case of the "merged graph" in which only a small number of vertices belong to densely connected subgraphs while most of the vertices are randomly connected. On the

Graph	Graph Characteristics
ER	N = 1,000, p = 0.333
SBM	$N = 1,000, c_i = 91,21,333,555$
CC	num cliques $= 10$, size of cliques $= 100$
BA3	N = 1,000, out-degree = 3
BA5	N = 1,000, out-degree = 5
WS	N = 1,000, k = 14, p = 0.2
CM	$N = 1,000, \exp = 3$
MD	One ER and two CC, $N = 1,000 = 850 + 2 \times (3 \times 25)$
EUC	N = 1,005 in 42 known clusters
CCw	num cliques = 12, size of cliques = 100, $w \in [0, 1]$

Table 2. Test Graph Characteristics (N denotes total number of vertices, $|c_i|$ number of vertices in cluster i, p edge probability)

Table 3. P-Values for δ (Unweighted Graphs)

	\mathbf{ER}	WS	BA3	BA5	CC	SBM	EUC	CM	MD
ER	0.64	1.00	1.00	1.00	1.00	1.00	1.00	0.75	0.92
WS	0.00	0.00	1.00	1.00	1.00	0.00	1.00	0.25	0.00
BA3	0.00	0.00	0.00	0.15	1.00	0.00	1.00	0.03	0.00
BA5	0.00	0.00	0.85	0.00	1.00	0.00	1.00	0.06	0.00
CC	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00	0.00
SBM	0.00	1.00	1.00	1.00	1.00	0.00	1.00	0.69	0.00
EUC	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
CM	0.25	0.75	0.97	0.94	1.00	0.31	1.00	0.25	0.25
MD	0.08	1.00	1.00	1.00	1.00	1.00	1.00	0.75	0.08

other hand, the null is correctly rejected in cases where the graph is known to have a strongly clustered structure (SBM, CC, EUC).

Our two-sample test accurately classifies the ER graph as not having a more heterogeneous density than any of the other graphs. It also confirms the EUC and CC graphs are more clustered than all other graphs.

We also perform a one-sample test on a weighted graph, a weighted connected caveman graph. The p-value for that test is p = 0.00, which confirms the graph is indeed clustered.

Other Tests Because of their similar approach, we also perform the EZ and T^2 tests presented by Gao and Lafferty in [17,18], as a point of comparison. However, because these tests do not lend themselves to pairwise comparison or weighted graphs, we are only able to compare one-sample tests on unweighted graphs.

 Table 4. p-Value Comparisons

Graph	EZ Test	T^2 Test	δ Test
ER	1.00	0.82	0.64
WS	1.00	0.00	0.00
BA3	1.00	0.00	0.00
BA5	1.00	0.00	0.00
CC	0.99	0.00	0.00
SBM	0.97	0.00	0.00
EUC	1.00	0.00	0.00
CM	1.00	0.00	0.25
MD	1.00	0.00	0.08
CCw	N.A.	N.A.	0.00

In Table 4, we observe how our test is more responsive to graph structure than either of the Gao-Lafferty tests. The p-values of the EZ test are completely unaffected by the graph's structure and the null hypothesis of an ER graph is never rejected. The T^2 test performs slightly better, since it correctly rejects the null in all but the ER case. However, unlike our δ test, it is unable to detect the lack of clustering in the configuration model (CM) and "merged" (MD) graphs. Also, neither of these competing tests is designed for graph-to-graph comparisons or defined for weighted graphs.

4 Conclusion

In this article, we presented a test for clusterability of a graph that is based on heterogeneity of local densities. We showed how variations in densities can be used as indicators of a probable clustered structure. Our tests perform as expected and compare very favorably to recently published competing techniques. Our future work will focus on sensitivity to sample sizes under various graph structures. We also intend to examine the cases of weighted graphs more closely. Currently, the topics of random and clustered weighted graphs remain subjects of debate.

Acknowledgments

Pierre Miasnikof was supported by a Mitacs-Accelerate PhD award IT05806. He also wishes to thank Lasse Leskelä of Aalto University, for the introduction to the work of Gao and Lafferty. Liudmila Prokhorenkova and Andrei Raigorodskii were supported by The Russian Science Foundation (grant number 16-11-10014).

References

- Albert, R., Barabási, A.L.: Statistical mechanics of complex networks. Reviews of Modern Physics 74, 47–97 (Jan 2002)
- Aleskerov, F., Goldengorin, B., Pardalos, P.: Clusters, Orders, and Trees: Methods and Applications. Springer Publishing Company, Incorporated (05 2014), https: //link.springer.com/book/10.1007/978-1-4939-0742-7
- 3. Alon, N., Shapira, A.: A characterization of the (natural) graph properties testable with one-sided error. SIAM J. COMPUT. 37(6), 1703–1727 (2008)
- Arias-Castro, E., Verzelen, N.: Community detection in dense random networks. Ann. Statist. 42(3), 940–969 (06 2014), https://doi.org/10.1214/14-A0S1208
- Barabási, A.L., Albert, R.: Emergence of Scaling in Random Networks. Science 286, 509–512 (Sep 1999)
- Barrat, A., Barthélemy, M., Pastor-Satorras, R., Vespignani, A.: The architecture of complex weighted networks. Proceedings of the National Academy of Science 101, 3747–3752 (Mar 2004)
- Bickel, P.J., Sarkar, P.: Hypothesis testing for automated community detection in networks. Journal of the Royal Statistical Society: Series B (Statistical Methodology) 78(1), 253-273 (2016), https://rss.onlinelibrary.wiley.com/doi/abs/ 10.1111/rssb.12117
- Butenko, S., Chaovalitwongse, W.A., Pardalos, P.M.: Clustering Challenges in Biological Networks. WORLD SCIENTIFIC (2009), https://www.worldscientific. com/doi/abs/10.1142/6602
- 9. Chiplunkar, A., Kapralov, M., Khanna, S., Mousavifar, A., Peres, Y.: Testing Graph Clusterability: Algorithms and Lower Bounds. ArXiv e-prints (Aug 2018)
- Czumaj, A., Peng, P., Sohler, C.: Testing Cluster Structure of Graphs. ArXiv eprints (Apr 2015)
- 11. Eden, T., Ron, D., Seshadhri, C.: On Approximating the Number of k-cliques in Sublinear Time. ArXiv e-prints (Mar 2018)
- Elenberg, E.R., Shanmugam, K., Borokhovich, M., Dimakis, A.G.: Beyond Triangles: A Distributed Framework for Estimating 3-profiles of Large Graphs. ArXiv e-prints (Jun 2015)
- Erdös, P., Rényi, A.: On Random Graphs I. Publicationes Mathematicae Debrecen 6, 290 (1959)
- Fortunato, S.: Community detection in graphs. Physics Reports 486, 75–174 (Feb 2010)

- 15. Fortunato, S., Hric, D.: Community detection in networks: A user guide. ArXiv e-prints (Nov 2016)
- Fronczak, A., Hołyst, J.A., Jedynak, M., Sienkiewicz, J.: Higher order clustering coefficients in Barabási-Albert networks. Physica A Statistical Mechanics and its Applications 316, 688–694 (Dec 2002)
- 17. Gao, C., Lafferty, J.: Testing for Global Network Structure Using Small Subgraph Statistics. ArXiv e-prints (Oct 2017)
- Gao, C., Lafferty, J.: Testing Network Structure Using Relations Between Small Subgraph Probabilities. ArXiv e-prints (Apr 2017)
- Gishboliner, L., Shapira, A.: Deterministic vs Non-deterministic Graph Property Testing. ArXiv e-prints (Apr 2013)
- Goldreich, O., Ron, D.: Algorithmic aspects of property testing in the dense graphs model. SIAM J. COMPUT. 40(2), 376–445 (2011)
- Hagberg, A., Schult, D., Swart, P.: Exploring network structure, dynamics, and function using networkx. In: Varoquaux, G., Vaught, T., Millman, J. (eds.) Proceedings of the 7th Python in Science Conference. pp. 11 – 15. Pasadena, CA USA (2008)
- He, Z., Liang, H., Chen, Z., Zhao, C.: Detecting statistically significant communities. CoRR abs/1806.05602 (2018), http://arxiv.org/abs/1806.05602
- 23. Jin, J., Ke, Z.T., Luo, S.: Network Global Testing by Counting Graphlets. ArXiv e-prints (Jul 2018)
- Leskovec, J., Kleinberg, J., Faloutsos, C.: Graph evolution: Densification and shrinking diameters. ACM Transactions on Knowledge Discovery from Data 1(1) (04 2007)
- Lovász, L., Vesztergombi, K.: Nondeterministic graph property testing. Combinatorics, Probability & Computing 22, 749–762 (2013)
- Miasnikof, P., Shestopaloff, A., Bonner, A., Lawryshyn, Y.: A Statistical Performance Analysis of Graph Clustering Algorithms, chap. 11. Lecture Notes in Computer Science, Springer Nature (6 2018)
- Prokhorenkova, L.O., Prałat, P., Raigorodskii, A.: Modularity of complex networks models. In: Bonato, A., Graham., F., Prałat, P. (eds.) Algorithms and Models for the Web Graph. pp. 115–126. Springer International Publishing, Cham (2016)
- Prokhorenkova, L.O., Prałat, P., Raigorodskii, A.: Modularity in several random graph models. Electronic Notes in Discrete Mathematics 61, 947 – 953 (2017), http://www.sciencedirect.com/science/article/pii/S1571065317302238, the European Conference on Combinatorics, Graph Theory and Applications (EURO-COMB'17)
- Prokhorenkova, L., Tikhonov, A.: Community detection through likelihood optimization: in search of a sound model. Proceedings of the 2019 World Wide Web Conference (WWW '19) (2019)
- 30. Schaeffer, S.E.: Survey: Graph clustering. Comput. Sci. Rev. 1(1), 27-64 (Aug 2007), http://dx.doi.org/10.1016/j.cosrev.2007.05.001
- Ugander, J., Backstrom, L., Kleinberg, J.: Subgraph Frequencies: Mapping the Empirical and Extremal Geography of Large Graph Collections. ArXiv e-prints (Apr 2013)
- Verzelen, N., Arias-Castro, E.: Community detection in sparse random networks. Ann. Appl. Probab. 25(6), 3465-3510 (12 2015), https://doi.org/10.1214/ 14-AAP1080
- Yang, J., Leskovec, J.: Defining and Evaluating Network Communities based on Ground-truth. CoRR abs/1205.6233 (2012), http://arxiv.org/abs/1205.6233

- 34. Yin, H., Benson, A.R., Leskovec, J.: Higher-order clustering in networks. ArXiv e-prints (May 2018)
- 35. Yin, H., Benson, A., Leskovec, J.: Local higher-order graph clustering. Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. 2017 (2017)