

Testable Algorithms for Approximately Counting Edges and Triangles in Sublinear Time and Space

Talya Eden  

Bar-Ilan University, Ramat Gan, Israel

Ronitt Rubinfeld 

Massachusetts Institute of Technology, Cambridge, MA, USA

Arsen Vasilyan  

University of Texas at Austin, TX, USA

Abstract

We consider the fundamental problems of approximately counting the numbers of edges and triangles in a graph in sublinear time. Previous algorithms for these tasks are significantly more efficient under a promise that the arboricity of the graph is bounded by some parameter $\bar{\alpha}$. However, when this promise is violated, the estimates given by these algorithms are no longer guaranteed to be correct.

For the triangle counting task, we give an algorithm that requires no promise on the input graph G , and computes a $(1 \pm \epsilon)$ -approximation for the number of triangles t in G in time $O^*\left(\frac{m \cdot \alpha(G)}{t} + \frac{m}{t^{2/3}}\right)$, where $\alpha(G)$ is the arboricity of the graph. The algorithm can be used on any graph G (no prior knowledge of the arboricity $\alpha(G)$ is required), and the algorithm adapts its run-time on the fly based on the graph G .

We accomplish this by trying a sequence of candidate values $\tilde{\alpha}$ for $\alpha(G)$ and using a novel algorithm in the framework of testable algorithms. This ensures that wrong candidates $\tilde{\alpha}$ cannot lead to wrong estimates: if the advice is incorrect, the algorithm either succeeds despite this or detects this and continues with a new candidate. Once the algorithm accepts the candidate, its output is guaranteed to be correct with high probability.

We prove that this approach preserves – up to an additive overhead – the dramatic efficiency gains obtainable when good arboricity bounds are known in advance, while ensuring robustness against misleading advice. We further complement this result with a lower bound, showing that such an overhead is unavoidable whenever the advice may be faulty.

We further demonstrate implications of our results for triangle counting in the streaming model.

2012 ACM Subject Classification Theory of computation → Streaming, sublinear and near linear time algorithms

Keywords and phrases Sublinear Algorithms, Triangle Counting, Edge Counting, Arboricity

Digital Object Identifier 10.4230/LIPIcs.ITCS.2026.54

Related Version *Full Version:* <https://arxiv.org/abs/2509.20351> [30]

Funding *Talya Eden:* Talya gratefully acknowledges the support of the MIT International Science and Technology Initiatives (MISTI) and the MIT-Israel Zuckerman STEM Fund. Part of this work was conducted while the authors were visiting the Simons Institute for the Theory of Computing.

Ronitt Rubinfeld: Supported by NSF awards DMS-2022448 and CCF-2310818, and the MIT-Israel Zuckerman STEM Fund. Part of this work was conducted while the authors were visiting the Simons Institute for the Theory of Computing.

Arsen Vasilyan: Arsen gratefully acknowledges the support of the MIT International Science and Technology Initiatives (MISTI) and the MIT-Israel Zuckerman STEM Fund. Part of this work was conducted while the authors were visiting the Simons Institute for the Theory of Computing.

Acknowledgements We thank Sepehr Assadi for discussions that influenced this work.



© Talya Eden, Ronitt Rubinfeld, and Arsen Vasilyan;
licensed under Creative Commons License CC-BY 4.0

17th Innovations in Theoretical Computer Science Conference (ITCS 2026).

Editor: Shubhangi Saraf; Article No. 54; pp. 54:1–54:24



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

Counting triangles is a fundamental graph problem with applications in network analysis, database systems, and social network theory, among others (see [1] for a comprehensive survey). A multitude of algorithms for counting triangles have been developed which work well in sequential [53, 18, 40], streaming [42, 49, 17], sublinear-time [23, 29, 12, 25, 27, 3, 31], parallel and distributed models [11, 33, 15, 19, 54, 46, 55], both from theoretical and practical perspectives. The best known *sublinear-time* algorithm for estimating the number of triangles in general graphs in the augmented model¹ runs in $O^*(\frac{m^{3/2}}{t})$ time [23], where m and t denote the number of edges and triangles in the input graph G , respectively. By the same result, this algorithm is optimal due to query complexity lower bounds.²

Given the importance of the triangle counting problem, there is much interest in even faster algorithms for typical real world graphs. These graphs often have structural properties that make it possible to perform faster computations. An important example of such a property is the graph's *arboricity*, a sparsity parameter that corresponds (up to a constant factor of 2) to the maximum average degree over all subgraphs. The arboricity is typically small in many real-world networks [34, 41, 58, 20], including planar graphs and preferential attachment networks. When an upper bound on the arboricity is known, there are fast sublinear-time triangle-counting algorithms [28, 10] that circumvent the known lower bounds for general graphs.³

The aforementioned arboricity-based algorithms require an upper bound $\bar{\alpha}$ on the actual arboricity of the graph, denoted $\alpha(G)$, as one of their input parameters. Those algorithms rely crucially on the correctness of this assumption. If the arboricity of the graph exceeds $\bar{\alpha}$, these algorithms are no longer guaranteed to give correct estimates for the number of triangles.

1.1 Our Contribution

In this work, we develop a novel algorithm in the augmented query access model that does not require any information about the arboricity of the input graph. For *all graphs* G , our algorithm gives a correct $(1 \pm \epsilon)$ -approximation to the number of triangles in G , and has a run-time that scales with the arboricity of the input graph. Specifically, the run-time of our algorithm is $O^*\left(\frac{m \cdot \alpha(G)}{t} + \frac{m}{t^{2/3}}\right)$.⁴

A naive approach would be to combine the sublinear-time arboricity estimator of [26] with the triangle counting results of [28, 10]. However, [26] runs in essentially optimal time $\Theta(n/\alpha(G))$, and so invoking it would dominate the complexity and substantially degrade the overall runtime, reflecting the inherently high query cost of arboricity estimation.

We avoid estimating arboricity directly, by first designing an algorithm in the formal framework of testable algorithms, introduced by [56] in the context of agnostic learning. A testable algorithm receives advice about the input that may or may not be accurate. The

¹ In the sublinear-time augmented query model, the algorithm can perform degree queries, neighbor queries and pair queries, as well as query for a uniformly distributed edge in the graph. See Section 2.2 for full details.

² The result of [23] is stated in the general query model, that only allows for degree, neighbor and pair queries, without uniform edge queries. Their complexity in this model is $\Theta^*\left(\frac{n}{t^{1/3}} + \frac{m^{3/2}}{t}\right)$. However, it directly follows from their work that in the augmented model the complexity is $\Theta^*\left(\frac{m^{3/2}}{t}\right)$ (that is, only the second term is needed for the upper bound, and the lower bound proof of the second term still holds for the stronger augmented model).

³ Here, the work of [28] is in the general query model, and the work of [10] is augmented model.

⁴ We use the $O^*(\cdot)$ notation to hide $\text{poly}(1/\epsilon, \log n)$ dependencies.

algorithm is allowed to either output an estimate or terminate early indicating “bad advice”. It must satisfy the following two properties: (1) *Completeness*: If the advice is correct, the algorithm outputs a correct estimate (with high probability). (2) *Soundness*: If the algorithm outputs an estimate (i.e., does not return “bad advice”), then that estimate is correct (with high probability), even if the advice was incorrect. This framework captures natural scenarios in which domain knowledge provides partial information about the input class, but the reliability of this information cannot be guaranteed.

We note that previous work on testable algorithms focused on settings in learning theory, and our central conceptual contribution is extending this framework to approximate counting and estimation tasks. Specifically, we design efficient testable algorithms for estimating the number of edges and triangles in a graph, given a potentially unreliable upper bound $\tilde{\alpha}$ on its arboricity as advice.⁵ Our algorithms return a $(1 \pm \epsilon)$ -approximation with high probability, provided it does not output “bad advice”. Moreover, if the given advice is accurate, then the algorithm is not likely to output “bad advice”.

► **Theorem 1** (Testable approximate triangle counting algorithm). *Let G be the input graph, and ϵ and δ approximation and error parameters. Let m and t be the number of edges and triangles in G , respectively. There exists an algorithm that gets as input augmented query access to G , the number⁶ of edges m , as well as an advice: a positive integer $\tilde{\alpha}$. The algorithm has expected run-time and query complexity of $O^*\left(\left(\frac{m\tilde{\alpha}}{t} + \frac{m}{t^{2/3}}\right) \log \frac{1}{\delta}\right)$. Upon each invocation, the algorithm either returns “bad advice” or a number \hat{t} , satisfying*

- (Completeness) *For every G , if $\alpha(G) \leq \tilde{\alpha}$, then the algorithm can output “bad advice” only with probability at most δ .*
- (Soundness) *For every G and $\tilde{\alpha}$, w.p. at least $1 - \delta$ the algorithm will either return “bad advice” or output $\hat{t} \in (1 \pm \epsilon)t$.*

In order to eliminate the need for advice, we use this algorithm with a sequence of guesses serving as our advice value $\tilde{\alpha}$. We start with running the algorithm with a setting of $\tilde{\alpha} = 1$, and continue invoking the algorithm while doubling the guessed value of $\tilde{\alpha}$ until the algorithm outputs an answer (rather than “bad advice”). By Soundness, the first accepting run outputs a correct estimate with high probability, and by Completeness, the number of guesses is bounded, so the overhead is only logarithmic. Consequently, we obtain an algorithm applicable to any graph G .

► **Corollary 2** (Instance-adaptive approximate triangle-counting algorithm.). *There exists an algorithm with the following guarantees. The algorithm receives augmented query access to a graph G , and the number of edges m in G together with parameters $\epsilon, \delta \in (0, 1)$. With probability at least $1 - \delta$, the algorithm produces an estimate \hat{t} that with probability at least $1 - \delta$ satisfies $\hat{t} \in (1 \pm \epsilon)t$, where t is the number of triangles in G . The expected run-time of the algorithm is $O^*\left(\left(\frac{m \cdot \alpha(G)}{t} + \frac{m}{t^{2/3}}\right) \log \left(\frac{1}{\delta}\right)\right)$, where $\alpha(G)$ is the arboricity of G .*

We refer to this algorithm as an *instance-adaptive algorithm* since it adapts “on the fly” to the structure of G , scaling with $\alpha(G)$ without requiring any a priori bound on it.

⁵ Throughout the paper we use $\alpha(G)$ to denote the actual arboricity of the graph G , $\bar{\alpha}$ to denote a guaranteed upper bound on the arboricity, and $\tilde{\alpha}$ to denote an unreliable upper bound.

⁶ We remark that the requirement on knowing m can be removed at an additive (essential) overhead of $O^*\left(\frac{n\alpha(G)}{m}\right)$, as implied by Corollary 25 below or $O^*(n^{1/4})$ due to [7]. However, for the sake of clarity, we chose to present the two results separately and focus here on the triangle estimation problem, as this is the main focus of our work.

Optimality of the bound

The query complexity of our testable algorithm for approximating the number of triangles lies between two extremes: the known lower bound in the general setting (with no advice), and the more efficient algorithms achievable under a reliable arboricity promise. We prove that this tradeoff is inherent: any algorithm that uses significantly fewer queries may produce an underestimate of the triangle count when the advice is incorrect, thus establishing the optimality of our approach. The lower bound also establishes a strict separation between the instance-adaptive estimation setting we consider here, and the trusted-advice setting of [28]. Formally, we give an information-theoretic lower-bound showing that $\Omega\left(\frac{m}{t^{2/3}}\right)$ queries are necessary, which together with the $\Omega\left(\frac{m \cdot \alpha(G)}{t}\right)$ lower bound of [28] demonstrates the optimality of our algorithm.

► **Theorem 3.** *Let \mathcal{A} be a testable algorithm for approximating the number of triangles in a graph G given arboricity advice $\tilde{\alpha}$ and augmented query access to G . Let n, m and t denote the number of nodes, edges and triangles in G . Finally, assume that the algorithm knows n and m , and that $\tilde{\alpha} \geq m/n$. Then \mathcal{A} must perform $\Omega\left(\frac{m}{t^{2/3}} + \frac{m \cdot \alpha(G)}{t}\right)$ queries in order to succeed with high probability.*

On the choice of the query model

As mentioned previously, the triangle-counting algorithm of [28] operates in the strictly less powerful *general* query model, which does not provide uniform edge samples. Under a promised arboricity bound $\bar{\alpha}$, their complexity is $\Theta^*\left(\frac{n\bar{\alpha}^2}{t} + \frac{m\bar{\alpha}}{t}\right)$, in contrast to the lower bound of $\Omega\left(\frac{n}{t^{1/3}} + \frac{m^{3/2}}{t}\right)$ for general graphs [23]. If we adopt the same (general) query model in our setting, but without any bound on the arboricity, then a lower bound of $\Omega(n/t^{1/3})$ follows immediately by hiding a clique of size $\Theta(t^{1/3})$. Thus, in order to achieve a meaningful arboricity-dependent runtime we assume access to *uniform edge samples*, as is given by the *augmented model*.⁷

Further results

Using a known reduction from sublinear-time algorithms in the augmented model to streaming algorithms in the edge arrival model (see, e.g., [26, 32]), we prove the following.

► **Theorem 4.** *There exists a 9-pass $O^*\left(\frac{m \cdot \tilde{\alpha} \log 1/\delta}{t} + \frac{m \log 1/\delta}{t^{2/3}}\right)$ -space algorithm for testable triangle-counting with arboricity advice (Definition 6) in the arbitrary-order insert-only edge arrival streaming model. As is standard and unavoidable in the streaming model, the algorithm requires a rough initial estimate $\tilde{t} = \Theta(t)$ of the triangle-count t .⁸*

⁷ For a more concrete sense on these bounds, consider a grid graph where each square is augmented with one diagonal edge. This graph has n vertices, $m = \Theta(n)$ edges, $t = \Theta(n)$ triangles, and arboricity $O(1)$. In the general query model and assuming a bound on the arboricity, the result of [28] yields an upper bound of $O\left(\frac{n\bar{\alpha}^2}{t} + \frac{m\bar{\alpha}}{t}\right) = O(1)$ queries. Without any knowledge of the arboricity and still in the general model, any algorithm requires $\Omega(n/t^{1/3}) = \Omega(n^{2/3})$ queries. By contrast, allowing uniform edge samples (but without assuming any bound on the arboricity) enables our $O^*\left(\frac{m}{t^{2/3}} + \frac{m\alpha(G)}{t}\right) = O(n^{1/3})$ algorithm.

⁸ In our implementation we assume $\tilde{t} \in [t/4, t]$, but the same argument can be adapted to work with $\tilde{t} \in [\frac{1}{c} \cdot t, c \cdot t]$ for any given known constant c .

In the guaranteed advice setting, Bera and Seshadhri [6] gave an arboricity-dependent 6-pass $O^*(m\bar{\alpha}/t)$ -space streaming algorithm for triangle counting.

We note that the instance-adaptive version of our algorithm can also be implemented in the streaming model, however with a $O(\log n)$ round complexity, due to the need to search for the correct advice $\bar{\alpha}$. For further discussion see Section 5. We also note that lower bounds do not directly transfer between the two models, and so it remains open as to whether the space bounds and round complexity of Theorem 1 are optimal.

Finally, we state our result for edge estimation in the testable setting.

► **Theorem 5.** *There exists a testable algorithm in the augmented query model (Definition 6) for $(1 \pm \epsilon)$ -approximate counting of edges on a graph G with n vertices and m edges, that has an expected run-time of $O^*\left(\frac{n \cdot \bar{\alpha} \log 1/\delta}{m}\right)$.*

Interestingly, this result does not incur any additive overhead over the $O^*\left(\frac{n\bar{\alpha}}{m}\right)$ results of [27]. It does, however, require the uniform edge samples of the stronger augmented query access model, as opposed to the algorithm of [27], which works in the general model. As in the case of triangles, this stronger query access is necessary if one wants to obtain an arboricity-dependent algorithm in the testable setting, as otherwise the presence of a small (hard to find) clique of size \sqrt{m} can fool any edge estimation algorithm which uses at most $o(n/\sqrt{m})$ vertex samples. For small-arboricity graphs, the resulting algorithm is much faster than the state-of-the-art edge-counting algorithm for general graphs in the augmented model [7]. In particular, for constant arboricity graphs, the running time of our algorithm is $O^*(1)$, whereas the algorithm of [7], that does not assume any knowledge on the arboricity, runs in time $O^*(n^{1/4})$, which is essentially optimal for general graphs.

As in the case of triangle counting, an easy consequence of our testable algorithms is an instance-adaptive algorithm for approximating the number of edges (using a similar doubling trick). Formally, the resulting instance-adaptive approximate edge counting algorithm runs in time $O^*\left(\frac{n \cdot \alpha(G)}{m}\right)$.

Our results are summarized below.

■ **Table 1** Summary of our results, compared to the state of the art in the same model for general graphs as well as bounded arboricity graphs. Our results can be seen as lying in between these two “extremes”. Here (and elsewhere) $\alpha(G)$ denotes the actual arboricity of the graph, $\bar{\alpha}$ a guaranteed upper bound on $\alpha(G)$, and $\tilde{\alpha}$ a potentially-incorrect advice on $\alpha(G)$.

	Computational model	General Graphs	Guaranteed Bound on Arboricity	Our Results (Testable Algorithms)
Triangles	sublinear-time, augmented	$\Theta^*\left(\frac{m^{3/2}}{t}\right)$	$O^*\left(\frac{m\bar{\alpha}}{t}\right)$	$\Theta^*\left(\frac{m \cdot \alpha(G)}{t} + \frac{m}{t^{2/3}}\right)$
	streaming, edge arrival	$\Theta^*\left(\min\left\{\frac{m^{3/2}}{t}, \frac{m}{\sqrt{t}}\right\}\right)$		$O^*\left(m\tilde{\alpha}/t + \frac{m}{t^{2/3}}\right)$
Edges	augmented model	$\Theta^*(n^{1/4})$	$O^*\left(\frac{n\bar{\alpha}}{m}\right)$	$O^*\left(\frac{n \cdot \alpha(G)}{m}\right)$

1.2 Overview of the algorithms and lower bounds for triangle estimation

1.2.1 Testable triangles algorithm

The ERS algorithm. Our algorithm builds on ideas from the earlier by [28], henceforth ERS, and we begin with a brief overview of how ERS works. The ERS algorithm estimates the number of triangles by classifying edges as *light* or *heavy* and assigning each triangle to

its first light edge, leaving only all-heavy triangles unassigned. It then samples a set R of uniformly random edges and attempts to find triangles incident to edges in R by sampling random neighboring edges. Whenever a triangle is found, the algorithm checks whether it is assigned to the sampled edge and uses these events to estimate the total number of assigned triangles.

ERS classifies an edge as heavy when it participates in more than τ_t triangles. This assignment rule yields a low-variance, nearly unbiased estimator. Consequently, sampling a set R of $O(m\tau_t/t)$ edges uniformly at random ensures that, with high probability, the number of triangles assigned to some edge in R is close to its expected value. To estimate this quantity, the algorithm repeatedly tries to sample triangles incident to edges of R , and whenever a triangle incident to some edge $e \in R$ is found, it checks whether the triangle is indeed assigned to e .

ERS prove that estimating the number of triangles assigned to R requires only $O(d(E)/t)$ triangle-sampling attempts in expectation, where $d(E) \stackrel{\text{def}}{=} \sum_{e \in E} d(e)$, and $d(e)$ for $e = \{u, v\}$ is defined as $\min\{d(u), d(v)\}$. Once a triangle is found, verifying whether it is assigned to an edge e can be done in $O(d(e)/\tau_t)$ queries. To bound this term, they define a *degree-heaviness threshold* τ_d beyond which edges are also classified as heavy (independent of their triangle count) and then ignored. Hence, the cost of an assignment verification step is $O(\tau_d/\tau_t)$.

The arboricity guarantee then allows them to bound each of these three terms by $O(m\bar{\alpha}/t)$, through the following arguments:

- (i) **Triangle-heaviness:** In bounded-arboricity graphs, edges with triangle load above $\tau_t = \bar{\alpha}/\epsilon$ can be ignored, since they can only contribute a small fraction of triangles. This determines the sample size of R .
- (ii) **Sum of edge degrees:** In bounded-arboricity graphs, the sum of edge degrees satisfies $d(E) = \sum_{e \in E} d(e) = O(m\bar{\alpha})$, implying that the number of triangle sampling attempts is bounded.
- (iii) **Degree-heaviness:** In bounded-arboricity graphs, only few edges exceed $\tau_d = m\bar{\alpha}^2/t$, so their contribution to the total triangles count is small and assignment verification remains efficient.

Plugging in each of these bounds in their respective terms gives the $O(m\bar{\alpha}/t)$ running time.

The testable algorithm. Our testable algorithm uses the three properties above to verify the advice $\tilde{\alpha}$. If these properties hold, then the correctness does not require that the actual arboricity be small; rather, these properties suffice to guarantee the quality of the estimate.

However, since the advice is not guaranteed, setting the thresholds τ_t and τ_d as in ERS no longer ensures that the number of unassigned heavy triangles is small. We must therefore modify the thresholds and explicitly check that the number of heavy edges remains bounded. In ERS, the set of heavy edges H satisfies $|H| = O(\epsilon t/\bar{\alpha})$, and arboricity is used to argue that H induces at most $O(\bar{\alpha}|H|) = O(\epsilon t)$ triangles. Without the arboricity guarantee, we instead rely on the general bound that any m -edge graph contains at most $O(m^{3/2})$ triangles. Hence we would like to make sure that $|H| = O((\epsilon t)^{2/3})$. Accordingly, we set

$$\gamma = c \cdot \max\{\tilde{\alpha}, t^{1/3}\}, \quad \tau_t = \gamma/\epsilon, \quad \text{and} \quad \tau_d = m\gamma^2/t,$$

for some small constant c . We show that these settings ensure that if the advice is correct then $|H|$ will be small. Therefore, we can estimate the size of H , and if $|H| > (\epsilon t)^{2/3}$ we can safely output “bad advice”. Otherwise, if $|H|$ does not exceed the bound (so we do not

reject), then it is small enough that – even in high arboricity graphs – it cannot induce too many heavy triangles, thereby preserving Properties (i) and (iii). For Property (ii), once the set R is sampled, we explicitly verify that $d(R) \stackrel{\text{def}}{=} \sum_{e \in R} d(e) = O(\tilde{\alpha}|R|)$; if this fails, we reject the advice. These modifications yield the runtime bound

$$O^*(m\tau_t/t + m\tilde{\alpha}/t + \tau_d/\tau_t) = O^*(m/t^{2/3} + m\tilde{\alpha}/t).$$

1.2.2 The $\Omega(m/t^{2/3})$ lower bound

We prove a lower bound of $\Omega(m/t^{2/3} + m\alpha(G)/t)$ for any testable triangle-counting algorithm in the augmented model. The $\Omega(m\alpha(G)/t)$ term follows from the no-advice lower bound of [28], and it remains to prove the first term, which is dominant whenever $\tilde{\alpha} < t^{1/3}$.

Consider a testable algorithm invoked with advice $\tilde{\alpha} < t^{1/3}$. The lower bound is shown by constructing two contrasting families of graphs.

In the first family, every graph has arboricity $\tilde{\alpha}$ and contains no triangles. The second family is obtained by taking graphs from the first family and augmenting them with a clique of size about $t^{1/3}$. (In the formal proof, the construction is refined to ensure that graphs in both families have exactly the same numbers of vertices and edges.) Consequently, graphs in the second family contain $\Theta(t)$ triangles and have arboricity $\Theta(t^{1/3})$, and so the advice is incorrect.

By definition, any testable algorithm must, with high probability, output an estimate $\hat{t} = 0$ on graphs from the first family, and either “bad advice” or an estimate $\hat{t} \in (1 \pm \epsilon)t$ on graphs from the second family. Hence, the algorithm must be able to distinguish between the two families. This requires it to hit the planted clique with high probability, which in turn forces it to perform $\Omega(m/t^{2/3})$ samples.

1.3 Additional related work

Testable learning algorithms

Testable learning is a framework introduced by Rubinfeld and Vasilyan [56] to allow a user to safely use a learning algorithm that is designed for specific distributions on labelled examples (e.g., uniform). The difficulty is that ascertaining whether examples come from the assumed distribution requires way too many samples. The new framework suggests the design of a tester and learner as a pair, where the tester efficiently tests a weaker property of the distribution, chosen such that the weaker property is both easier to test and sufficient for proving that the algorithm’s output is correct. Their work, and several followups [45, 36, 59, 37, 38, 21, 22] have demonstrated such tester-learner pairs for a number of learning tasks. The work of [35] extends the testable learning framework to handle assumptions on the label noise in classification. A recent line of work on Testable Learning with Distribution Shift (TDS learning) [45, 44, 16, 36] tests the assumption that the data encountered by a classifier during deployment comes from the same distribution encountered during training.

In concurrent work, Marcussen, Rubinfeld and Sudan introduce an analogue of testable algorithms for the setting of random graphs [47], referred to as “quality control problems”. As in this work, the algorithms in [47] have the goal of reliably approximating the number of motifs in input graphs. However, the requirements of the testable algorithm differ in our setting and theirs, particularly with respect to the completeness guarantees of the algorithms. Completeness in their setting corresponds to algorithmic correctness with high probability over random graphs, while in our setting the algorithm must be correct with high probability on any low-arboricity input graph. Our work also uses very different techniques because low-arboricity graphs have disparate structural properties from random graphs.

Testing and estimating the arboricity in sublinear-time and space

The task of estimating the arboricity was studied in a series of works, both in the sublinear-time and the streaming settings [9, 48, 43, 5, 9]. Eden, Mossel and Ron [26] improved the running times and space complexity to $O(n/\bar{\alpha})$ at the cost of increasing the approximation parameter to $O(\log^2 n)$ and the round complexity to logarithmic. Eden, Levi, and Ron [24] presented an $O(1/\epsilon)^{O(\log(1/\epsilon))}$ -time algorithm for testing whether a graph has bounded arboricity in the augmented model. However, being ϵ -close to a bounded-arboricity graph is not sufficient for applying ERS: even a small number of additional edges can still generate many triangles.

Sublinear-time and space subgraph counting algorithms

Sublinear time subgraph counting results have been extensively studied both in various query models [39, 23, 29, 25, 3, 31, 12, 2, 27, 10]. Estimating the number of edges in sublinear time, when given access to uniform edge samples as well as degree and neighbor queries, follows (implicitly) from the work of Motwani, Panigrahy and Xu [50] and from the work of Beretta and Tětek [8]. Both these results give an essentially optimal bound of $\Theta^*(n^{1/3})$. When pair queries are also allowed, implying full augmented query access, the bound was recently improved to $\Theta^*(n^{1/4})$ by Beretta, Chakrabarty and Seshadhri [7].

Sublinear arboricity-dependent bounds were studied in [27, 28, 4, 10, 6, 32].

Triangle counting with predictions

Algorithms with predictions, like testable algorithms, use auxiliary information to overcome worst-case lower bounds. However, the auxiliary information is often more informative (as it models advice given by a machine learning algorithm). In addition, such algorithms must always produce an output, even when the prediction is of low quality. The goal is that the performance improves with prediction accuracy, and, ideally, never falls below the baseline guarantee. Triangle counting in the streaming model using predictions was studied in several papers using various types of predictions [17, 13]. Common to all is that they require per-edge advice, rather than global information.

1.4 Organization

In Section 2, we formally define the notion of a testable algorithm, and the augmented query access model, and provide some useful notations and claims. In Section 3, we provide our sublinear time and space algorithms and lower bound for triangle counting, with some of the proofs deferred to Appendix B in the full version of this work [30]. The sublinear-time lower bound is proven in Section 4. The streaming result is discussed in Section 5. Finally, in Section 6, we provide the algorithm for edge counting.

2 Preliminaries

2.1 Testable algorithm definition

The following is our definition of testable algorithms for estimating the number of motifs with arboricity advice. The problems of estimating the number of edges and the number of triangles are a special case.

► **Definition 6** (Testable algorithm for motif counting with arboricity advice). *Let H be a fixed subgraph, and let t denote the number of copies of H in an input graph G . A testable arboricity advice algorithm for counting the number of copies of H in G (in the augmented query model) is an algorithm that receives: augmented query access to G , an approximation parameter ϵ , and a positive integer $\tilde{\alpha}$, which serves as a (possibly unreliable) upper bound on the arboricity of G .*

The algorithm outputs either an estimate \hat{t} or the comment “bad advice”, as follows:

- **Completeness:** *If $\alpha(G) \leq \tilde{\alpha}$, then whp \mathcal{A} outputs an estimate \hat{t} (and does not output “bad advice”).*
- **Soundness:** *For every graph G and any $\tilde{\alpha}$, if \mathcal{A} outputs a value \hat{t} , then whp $\hat{t} \in (1 \pm \epsilon)t$.*

2.2 The augmented query model

In the *augmented query model* the algorithm interacts with an unknown simple graph $G = ([n], E)$ through the following oracles: **Vertex access.** Vertices are labeled $[n]$, so the algorithm can specify any $v \in [n]$ (equivalently, it can draw a uniform random vertex by choosing v uniformly from $[n]$). **Degree query.** On input v , return the degree of v , $d(v)$. **i^{th} -neighbor query.** On input (v, i) , return the i^{th} neighbor of v if $1 \leq i \leq d(v)$ (with respect to a fixed but arbitrary ordering of v 's incident edges), and a distinguished null symbol \perp otherwise. **Pair query.** On input (u, v) , return whether $\{u, v\} \in E$. **Uniform edge sample.** Return an edge drawn uniformly at random from E , independently at each call.

If no edge sample access is allowed, then this is referred to as the *general access model*.

2.3 Notations and useful inequalities

► **Definition 7** (Arboricity of a graph). *The arboricity of a graph G is the minimum number of forests required to cover its edge set. By [51, 52], it is also equivalent to*

$$\alpha(G) = \max_{S \subseteq G} \left\lceil \frac{m(S)}{n(S) - 1} \right\rceil,$$

where $n(S)$ and $m(S)$ denote the number of vertices and edges, respectively, in the subgraph induced by the set S .

► **Theorem 8.** *For any graph G , $\alpha(G) \leq \sqrt{m}$ and $m \leq n\alpha(G)$.*

The following theorem of Chiba and Nishizeki bounds the number of triangles by the arboricity.

► **Theorem 9** (Due to [18]). *Let G be a graph with arboricity at most $\alpha(G)$. Then*

$$t(G) \leq \sum_{e \in E} d(e) \leq 2m \cdot \alpha(G) \leq 2m^{3/2}.$$

► **Definition 10** (Degree, neighbors and triangles of an edge). *The degree of an edge $e = \{u, v\}$ is the minimum degree of its endpoints, $d(e) = d(\{u, v\}) = \min\{d(u), d(v)\}$. The set of neighbors of an edge, is the set of vertices that are neighbors of its min degree endpoint.*

Also let $t(e)$ denote the number of triangles incident to the edge e .

Chernoff inequality and a useful corollary are deferred to Appendix A in the full version of this work [30].

3 The Triangles Approximation Algorithm

This section is devoted to proving the main theorem of the paper, a testable triangle counting algorithm with arboricity advice (Definition 6). We start with an algorithm that takes as an input a value \tilde{t} that serves as a rough guess of the true triangle count t . In Section B.3 of the full version of this work [30] we show how to use Theorem 16 together with the Search Theorem of [29] to search for a good guess \tilde{t} and obtain a testable triangle counting algorithm that does not require an input of \tilde{t} .

Before presenting the algorithm, we provide some useful notations and structural claims.

3.1 Preliminaries for triangle counting

The algorithm distinguishes those edges which either have high degree or too many triangles incident to them:

► **Definition 11** (Heavy and light edges). *For a given advice $\tilde{\alpha}$, and a guessed value \tilde{t} , let $\gamma = \max\{\tilde{\alpha}, \tilde{t}^{1/3}\}$. Further let $\tau_d = \frac{8m\gamma^2}{\epsilon\tilde{t}}$ be a degree threshold, and $\tau_t = 12\gamma/\epsilon$ be a triangles-degree threshold. Let H_d denote the set of edges for which $d(e) > \tau_d$, and H_t the set of edges for which $t(e) > \tau_t$. We say that an edge e is degree- (triangles-) heavy, if either $e \in H_d$ ($e \in H_t$). Otherwise, we say it is degree (triangles) light. Finally, we let $H = H_d \cup H_t$ and refer to it as the set of heavy edges, and to $E \setminus H$ as the set of light edges.*

The algorithm assigns each triangle to one of its edges as follows:

► **Definition 12** (Assigning triangles to edges). *Let $P = (E_0, E_1)$ be a partition of the graphs' edges. We assign each triangle to its first edge in E_0 if such exists (according to some arbitrary predefined order). Otherwise, if all of its edges are in E_1 , then the triangle is not assigned to any edge. For an edge e , we let $a_P(e)$ denote the number of triangles assigned to e by partition P .*

The algorithm partitions edges into two groups E_0, E_1 . E_0 contains all light edges. E_1 contains all edges that are degree-heavy and all edges that are “extremely” triangles-heavy – that is, with a higher threshold than required to just be triangles-heavy. Edges that are degree-light, and triangles-heavy with $\tau_t \leq t(e) \leq 2\tau_t$ can be in either E_0 or E_1 .

► **Definition 13** (A good partition). *We say that a partition of the graphs' edges $P = (E_0, E_1)$ is (τ_d, τ_t) -good if: (1) any light edge (according to Definition 11) is in E_0 , (2) every edge such that $d(e) > \tau_d$ or $t(e) > 2\tau_t$ is in E_1 , and (3) all other edges can be either in E_0 or E_1 .*

Our first claim shows that if the advice is good, then we can ignore triangles for which all edges are heavy, because there are not many of them.

▷ **Claim 14.** If $\tilde{\alpha} > \alpha(G)$ and $\tilde{t} \in [t/4, t]$, then $|H_d| \leq (\epsilon\tilde{t})^{2/3}$ and $|H_t| \leq (\epsilon\tilde{t})^{2/3}$.

Proof. By Theorem 9 and the assumption on $\tilde{\alpha}$, $\sum_{e \in E} d(e) \leq 2m \cdot \alpha(G) \leq 2m\tilde{\alpha}$. Therefore, there could be at most $2m\tilde{\alpha}/\tau_d$ edges with degree greater than τ_d . By the setting of $\tau_d = \frac{8m\gamma^2}{\epsilon\tilde{t}}$, we get

$$|H_d| \leq \frac{2m\tilde{\alpha}}{\tau_d} = \frac{\epsilon\tilde{\alpha}\tilde{t}}{4\gamma^2} \leq \frac{\epsilon\tilde{t}}{4\gamma} \leq \frac{1}{4} \cdot \min \left\{ (\epsilon\tilde{t})^{2/3}, \frac{\epsilon\tilde{t}}{\tilde{\alpha}} \right\},$$

where the last two inequalities are since $\gamma = \max\{\tilde{\alpha}, \tilde{t}^{1/3}\}$.

By the definition of H_t , it holds that $\sum_{e \in H_t} t(e) \geq |H_t| \cdot \tau_t$ and also, $\sum_{e \in H_t} t(e) \leq 3t$ since this holds for any set of edges, so in particular it holds for H_t . Hence, by the setting of $\tau_t = 12\gamma/\epsilon = \frac{12}{\epsilon} \cdot \max\{\tilde{\alpha}, \tilde{t}^{1/3}\}$ and since $t \leq 4\tilde{t}$, it follows that $|H_t| \leq 3t/\tau_t \leq \epsilon t / (4\tilde{t}^{1/3}) \leq (\epsilon\tilde{t})^{2/3}$. \triangleleft

\triangleright **Claim 15.** For any partition $P = (E_0, E_1)$, it holds that $\sum_{e \in E} a_P(e) \leq t$. Furthermore, if $|E_1| \leq 3(\epsilon\tilde{t})^{2/3}$ and $\tilde{t} \in [t/4, t]$, then $\sum_{e \in E_0} a_P(e) \geq (1 - 12\epsilon)t$.

Proof. By the assignment rule in Definition 12, every triangle is assigned to at most one edge, so it always holds that $\sum_{e \in E} a_P(e) \leq t$.

We now prove the second part of the claim. If there are at most $3(\epsilon\tilde{t})^{2/3}$ edges in E_1 , by Theorem 9 these edges induce at most $2 \cdot |E_1|^{3/2} \leq 12\epsilon\tilde{t} \leq 12\epsilon t$ triangles, where the last inequality is by the assumption $\tilde{t} \leq t$. That is, there are at most $12\epsilon t$ triangles with all three edges are in E_1 , and all the rest of the (at least) $(1 - 12\epsilon)t$ triangles will be assigned to one of the edges in E_0 . Hence, $\sum_{e \in E_0} t(e) > (1 - 12\epsilon)t$. \triangleleft

3.2 The algorithm

Approx-Triangles-With-Arboricity-Advice($\tilde{t}, \epsilon, \delta, \tilde{\alpha}$).

1. Set $\gamma = \max\{\tilde{\alpha}, \tilde{t}^{1/3}\}$, $\tau_d = \frac{8m\gamma^2}{\epsilon\tilde{t}}$ and $\tau_t = 12\gamma/\epsilon$.
2. Sample a set R of $r = \frac{16m \cdot \tau_t \cdot \ln(4/\delta)}{\epsilon^2 \cdot \tilde{t}}$ edges uniformly at random.
3. Query the degrees of all edges in R , and let $d(R) = \sum_{e \in R} d(e)$.
If $d(R) > |R| \cdot \tilde{\alpha} \cdot \frac{4}{\delta}$ then output “bad advice”
4. Invoke `IsHeavy` on all edges in R , and if the fraction of edges for which the algorithm returned “Heavy” is greater than $\frac{5}{2} \cdot \frac{(\epsilon\tilde{t})^{2/3}}{m}$ then output “bad advice”.
** The remaining steps of the algorithm essentially execute the algorithm of [28]. **
5. Set up a data structure to allow sampling each edge in R with probability $d(e)/d(R)$.
6. For $i = 1$ to $s = \frac{d(R)}{|R| \cdot (\tilde{t}/m)} \cdot \frac{10 \ln(8/\delta)}{\epsilon^2}$ do:
 - a. Sample an edge $e \in R$, so that each edge is sampled with probability $d(e)/d(R)$.
 - b. Pick w to be a uniform random neighbor of e .
 - c. If e, w form a triangle, then invoke `IsAssigned`(e, w).
 - d. If the triangle (e, w) is assigned to e , then let $\chi_i = 1$. Otherwise, let $\chi_i = 0$.
7. Let $\chi = \frac{1}{s} \sum_{i=1}^s \chi_i$
8. Return $\hat{t} \stackrel{\text{def}}{=} \frac{d(R) \cdot m}{|R|} \chi$

`IsAssigned`($e = (u, v), w$).

1. Invoke `IsHeavy` for each of the edges $(u, v), (u, w), (v, w)$.
2. If e is the first edge in E_0 then return YES. Otherwise, return NO.

As noted above, we begin with an algorithm that takes as input a value \tilde{t} , intended to be a rough estimate of the true triangle count t . It is a standard technique in subgraph-counting algorithms to first design a procedure that succeeds under the assumption that a coarse approximation of the target quantity is provided, and then later eliminate this assumption by iterating over a sequence of candidate estimates. Formally, to accomplish this we invoke the Search Theorem of [29]. Applying this theorem requires us to decompose the Soundness condition of our algorithm into two cases – depending on whether the guess \tilde{t} is reasonably accurate upper bound on the true triangle count t or instead exceeds t .

$\text{IsHeavy}(e = (u, v), \tau_t, \tau_d)$.

1. If $d(e) > \tau_d$ then return “Heavy”.
2. Sample $k = 18 \frac{d(e)}{\tau_t} \ln \frac{10m}{\delta}$ neighbors of e uniformly at random (from its min-degree vertex), and for each check if it forms a triangle with e .
 ** If $\text{IsHeavy}(e = (u, v), \tau_t, \tau_d)$ was called in the past, re-use the same random bits as used previously. **
3. If the number of witnessed triangles is greater than $1.5k \cdot \frac{\tau_t}{d(e)}$, then return “Heavy”. Otherwise, return “Not Heavy”.

► **Theorem 16.** *Let G be the input graph and ϵ and δ be an approximation and error parameters. Let t be the number of triangles in G . Algorithm **Approx-Triangles-With-Arboricity-Advice** receives: augmented query access to G , the number of edges m , parameters ϵ and δ , and a pair of positive integers $\tilde{\alpha}$ and \tilde{t} . The algorithm has an expected run-time and query complexity of $O\left(\frac{m \cdot \max\{\tilde{\alpha}, \tilde{t}^{1/3}\}}{\tilde{t}} \cdot \frac{\ln \frac{1}{\delta}}{\epsilon^3 \delta} \ln \frac{m}{\delta} \cdot (1 + t/\tilde{t})\right)$. Every time the algorithm is run, it can either return a number \hat{t} or it can return “bad advice”, and*

- (Completeness) *For every G and \tilde{t} , if $\alpha(G) \leq \tilde{\alpha}$, then w.p. at least $1 - \delta$ the algorithm **Approx-Triangles-With-Arboricity-Advice** will not output “bad advice”.*
- (Used for soundness) *For every G and $\tilde{\alpha}$, if $\tilde{t} \in [t/4, t]$ then w.p. at least $1 - \delta$ **Approx-Triangles-With-Arboricity-Advice** will either return “bad advice” or output $\hat{t} \in (1 \pm 20\epsilon)t$.*
- (Used for soundness) *For every G and $\tilde{\alpha}$, if $\tilde{t} > t$ then w.p. at least 5ϵ , the algorithm **Approx-Triangles-With-Arboricity-Advice** will either return “bad advice” or output a value of \hat{t} such that $\hat{t} < (1 + 20\epsilon)t$.*

The proof of the run-time bound in Theorem 16 is deferred to Appendix B.2 in the full version of this work [30] and the correctness properties in Theorem are shown in Section 3.4 (this includes the completeness property, and the two properties used for soundness). Section 3.3 focuses on the sub-routine **IsHeavy** with the proof of Claim 17 deferred to Appendix B in the full version of this work [30].

As mentioned earlier, Theorem 16 is used in Section B.3 of the full version of this work [30] together with the Search Theorem of [29] to obtain a testable algorithm (Definition 6) in the augmented query model for $(1 \pm \epsilon)$ -approximate counting of triangles (Theorem 1). The instance-adaptive algorithm (Corollary 2) is analyzed in Section 3.5.

3.3 Analyzing procedure **IsHeavy**

We rely on the following the guarantees of the procedure **IsHeavy**, and defer the proof to Section B.1 in the full version of this work [30].

▷ **Claim 17.** With probability at least $1 - \delta/10$ over all random coins, **IsHeavy** induces a (τ_d, τ_t) -good partition, as defined in Definition 13. The query complexity and the run-time of **IsHeavy** is $O\left(\frac{\min(\tau_d, d(e))}{\tau_t} \ln \frac{m}{\delta}\right)$.

3.4 Correctness of **Approx-Triangles-With-Arboricity-Advice**

In this section we prove the Completeness condition and the two Soundness conditions in Theorem 16. For the rest of the analysis we will use the following convention:

► **Remark 18.** In this section we will consider as random variables various quantities in Algorithm `Approx-Triangles-With-Arboricity-Advice` such as \hat{t} and χ_i . These random variables are well-defined even if the algorithm decides to terminate in steps 4 or 3 prior to computing these quantities. This is true because all these random variables are determined on (a) the partition defined by `IsHeavy`(b) the set R and (c) the edges sampled in Step 6.

3.4.1 Proof of Completeness

In this sub-section we prove that the algorithm `Approx-Triangles-With-Arboricity-Advice` satisfies the completeness condition in Theorem 16. Specifically, we show that for every G and \tilde{t} , if $\alpha(G) \leq \tilde{\alpha}$, then w.p. at least $1 - \delta$ the algorithm `Approx-Triangles-With-Arboricity-Advice` will not output “bad advice”. Recalling that the algorithm has the option to reject in Step 3 and Step 4, in this section we show that either is unlikely if $\alpha(G) \leq \tilde{\alpha}$. We start by considering Step 4.

▷ **Claim 19.** If the partition $P = (E_0, E_1)$ induced by the procedure `IsHeavy` satisfies $|E_1| \leq 2(\epsilon\tilde{t})^{2/3}$, then the algorithm outputs “bad advice” in Step 4 with probability only at most $\delta/4$. Moreover, $P = (E_0, E_1)$ satisfies $|E_1| \geq 3(\epsilon\tilde{t})^{2/3}$, then the algorithm will output “bad advice” in Step 4 with probability at least $1 - \delta/4$.

Proof. In Step 4, the procedure `IsHeavy` outputs “Heavy” for an edge e if and only if $e \in E_1$, where recall that $P = (E_0, E_1)$ induced by the procedure `IsHeavy`. Therefore each edge e sampled in Step 4, has a probability of $|E_1|/m$ to be in E_1 (so `IsHeavy` outputs “Heavy” for e).

Suppose $|E_1| \leq 2(\epsilon\tilde{t})^{2/3}$, then by the multiplicative Chernoff bound, and by the setting of $|R| \geq \frac{m}{(\epsilon\tilde{t})^{2/3}} \cdot 30 \ln(4/\delta)$,

$$\Pr \left[\sum_{e \in R} \mathbb{1}_{e \in E_1} > \frac{5}{2} \cdot |R| \cdot \frac{(\epsilon\tilde{t})^{2/3}}{m} \right] < \exp \left(-\frac{|R| \cdot (\epsilon\tilde{t})^{2/3}}{24m} \right) = \delta/4,$$

proving the first part of the claim. Now, suppose that $|E_1| \geq 3(\epsilon\tilde{t})^{2/3}$. Then, again

$$\Pr \left[\sum_{e \in R} \mathbb{1}_{d(e) > \tau_d} \leq \frac{5}{2} \cdot |R| \cdot \frac{(\epsilon\tilde{t})^{2/3}}{m} \right] < \exp \left(-\frac{1}{24} \cdot \frac{2(\epsilon\tilde{t})^{2/3}}{m} \cdot |R| \right) = \delta/4.$$

◁

Building on the claim above, we now show that Step 4 is unlikely to output “bad advice” if $\tilde{\alpha} \geq \alpha(G)$.

▷ **Claim 20.** If $\tilde{\alpha} \geq \alpha(G)$ and the partition $P = (E_0, E_1)$ induced by `IsHeavy` is a (τ_d, τ_t) -good partition (see Definition 13), then the algorithm can output “bad advice” in Step 4 only with probability at most $\delta/4$.

Proof. If $P = (E_0, E_1)$ induced by `IsHeavy` is a (τ_d, τ_t) -good partition (Definition 13), then the E_1 is a subset of $H_d \cup H_t$. But if $\tilde{\alpha} \geq \alpha(G)$, then by Claim 14, we see that $|H_d| \leq (\epsilon\tilde{t})^{2/3}$ and $|H_t| \leq (\epsilon\tilde{t})^{2/3}$, and therefore

$$|E_1| \leq |H_d| + |H_t| \leq 2(\epsilon\tilde{t})^{2/3},$$

and thus by claim 19 the algorithm outputs “bad advice” in Step 4 with probability only at most $\delta/4$. ◁

54:14 Testable Algorithms for Approximately Counting Edges and Triangles

Finally, we prove that if the advice $\tilde{\alpha}$ is good, then with high probability the algorithm will not reject due to Step 3.

▷ **Claim 21.** If $\tilde{\alpha} \geq \alpha(G)$, then the algorithm can only return “bad advice” in Step 3 with probability at most $\delta/3$.

Proof. By Theorem 9, if the graph has arboricity at most $\alpha(G)$ and $\tilde{\alpha} \geq \alpha(G)$, then

$$\mathbb{E}_{e \in E}[d(e)] = \frac{1}{m} \sum_{e \in E} d(e) \leq 2\alpha(G) \leq 2\tilde{\alpha}.$$

Hence, for every i , $\mathbb{E}[d(R)] \leq |R| \cdot 2\tilde{\alpha}$, and by Markov’s inequality, $\Pr[d(R) > 4\tilde{\alpha}|R|/\delta] < \delta/4$. \triangleleft

3.4.2 Proof of Soundness Condition 1

In this sub-section we prove that the algorithm **Approx-Triangles-With-Arboricity-Advice** satisfies the first soundness condition in Theorem 16. Specifically, we prove that for every G and $\tilde{\alpha}$, if $\tilde{t} \in [t/4, t]$ then, with probability at least $1 - \delta$, **Approx-Triangles-With-Arboricity-Advice** will either return “bad advice” or output $\hat{t} \in (1 \pm 20\epsilon)t$.

Let $P = (E_0, E_1)$ be the partition induced by the procedure **IsHeavy**. By Claim 17, with probability at least $1 - \delta/10$, the partition P is a (τ_d, τ_t) -good (see Definition 13). Since the probability distribution of the partition P is (as a random variable) independent from the choice of set R and edges that **Approx-Triangles-With-Arboricity-Advice** picks in Step 6. Thus, if we condition on the event that P is a (τ_d, τ_t) -good partition, this does not change the distribution of R and edges in Step 6. Therefore, for the rest of this subsection we condition on the event that P is a (τ_d, τ_t) -good partition.

Now, suppose that $|E_1| > 3(\epsilon\tilde{t})^{2/3}$. Claim 19 tells us that in this case the algorithm **Approx-Triangles-With-Arboricity-Advice** will with probability at least $1 - \delta/4$ output “bad advice”.

Thus, it remains to consider the case that $|E_1| \leq 3(\epsilon\tilde{t})^{2/3}$. We show in Claim 23 (provided $\tilde{t} \in [t/4, t]$) in this case the random variable \hat{t} will satisfy $\hat{t} \in (1 \pm 20\epsilon)t$. This will finish the proof because, as follows from Remark 18, in the event that the *random variable* \hat{t} satisfies $\hat{t} \in (1 \pm 20\epsilon)t$ then the algorithm **Approx-Triangles-With-Arboricity-Advice** either outputs such \hat{t} or outputs “bad advice” at an earlier step.

We first need the following claim telling us that the number of triangles assigned to set R will be approximately proportional to the total number triangles t (with high probability):

▷ **Claim 22.** If $P = (E_0, E_1)$ is a (τ_d, τ_t) -good partition (Definition 13), satisfies $|E_1| \leq 3(\epsilon\tilde{t})^{2/3}$, and $\tilde{t} \in [t/4, t]$, then with probability at least $1 - \frac{\delta}{4}$, we have

$$\frac{a_P(R)}{|R|} \in \left[(1 - 13\epsilon) \cdot \frac{t}{m}, (1 + \epsilon) \cdot \frac{t}{m} \right].$$

Proof. If P is a (τ_d, τ_t) -good partition, then for any edge $e \in E$, $a_P(e) \leq 2\tau_t$. This is true since if $t(e) > 2\tau_t$, then $e \in E_1$, and so $a_P(e) = 0$, and otherwise, if $e \in E_0$, then $a_P(e) \leq t(e) \leq 2\tau_t$. Furthermore, if $\tilde{t} \leq t$, then $|R| \geq \frac{16m \cdot \tau_t \cdot \ln(4/\delta)}{\epsilon^2 \cdot \tilde{t}} \geq \frac{m \cdot (2\tau_t)}{t} \cdot \frac{3 \ln(4/\delta)}{\epsilon^2}$. Therefore, by the Chernoff bound, with probability at least $1 - \frac{\delta}{4}$,

$$\frac{1}{r} \sum_{e \in R} a_P(e) \in (1 \pm \epsilon) \mathbb{E}_{e \in E}[a_P(e)] \in (1 \pm \epsilon) \cdot \frac{\sum_{e \in E_0} a_P(e)}{m}.$$

Claim 15 tells us that if $|E_1| \leq 3(\epsilon\tilde{t})^{2/3}$ and $\tilde{t} \in [t/4, t]$, then $t \geq \sum_{e \in E_0} a_P(e) \geq (1 - 12\epsilon)t$, which combined with the bound above yields the conclusion of the claim. \triangleleft

Overall, recall that the partition $P = (E_0, E_1)$ induced by the algorithm `IsHeavy` is a (τ_d, τ_t) -good partition with probability at least $1 - \frac{\delta}{4}$ (Claim 17), this subsection has the premise that $\tilde{t} \in [t/4, t]$ and we are focusing on the case that $|E_1| > 3(\epsilon\tilde{t})^{2/3}$. With all this in mind, the following claim finishes the proof in this subsection, showing that the estimate \hat{t} will indeed be a good approximation to the true triangle count t .

▷ **Claim 23.** Suppose $\epsilon \leq 1/20$ and the partition $P = (E_0, E_1)$ induced by the algorithm `IsHeavy` is a (τ_d, τ_t) -good partition (Definition 13), and also suppose $|E_1| \leq 3(\epsilon\tilde{t})^{2/3}$ and $\tilde{t} \in [t/4, t]$, then with probability at least $1 - \delta/4$ the algorithm outputs a value \hat{t} such that $\hat{t} \in (1 \pm 20\epsilon)t$.

Proof. Observe that conditioned on an edge $e \in R$ being sampled in Step 6a, $\chi_i = 1$ iff the triangle (e, w) is assigned to e , which happens with probability $\Pr[\chi_i = 1 \mid e] = \frac{a_P(e)}{d(e)}$. Therefore,

$$\mathbb{E}[\chi_i] = \sum_{e \in R} \frac{d(e)}{d(R)} \cdot \frac{a_P(e)}{d(e)} = \frac{a_P(R)}{d(R)}. \quad (1)$$

By Claim 22, and since $\frac{a_P(R)}{R} \in [(1 - 12\epsilon)\frac{t}{m}, (1 + \epsilon)\frac{t}{m}]$, it is the case that

$$\mathbb{E}[\chi_i] \in \left[(1 - 12\epsilon) \frac{|R| \cdot t}{d(R) \cdot m}, (1 + \epsilon) \frac{|R| \cdot t}{d(R) \cdot m} \right].$$

This, together with the premise that $\tilde{t} \in [t/4, t]$ and $\epsilon \leq 1/20$ implies that

$$s \stackrel{\text{def}}{=} \frac{d(R)}{|R| \cdot \frac{\tilde{t}}{m}} \cdot \frac{10 \ln(8/\delta)}{\epsilon^2} \geq \frac{1}{\mathbb{E}[\chi_i]} \cdot \frac{3 \ln(8/\delta)}{\epsilon^2},$$

and therefore the Chernoff bound allows us to conclude that

$$\Pr \left[\left| \frac{1}{s} \cdot \chi - \mathbb{E}[\chi_i] \right| > \epsilon \mathbb{E}[\chi_i] \right] < \frac{\delta}{4}.$$

If this holds, then we have

$$\hat{t} \stackrel{\text{def}}{=} \frac{d(R) \cdot m}{|R|} \chi = (1 \pm 13\epsilon)t,$$

proving the claim. ◁

3.4.3 Proof of Soundness Condition 2

In this sub-section we prove that the algorithm `Approx-Triangles-With-Arboricity-Advice` satisfies the second soundness condition in Theorem 16. This is accomplished by the following claim:

▷ **Claim 24.** For every G and $\tilde{\alpha}$, if $\tilde{t} > t$ then w.p. at least $\epsilon/4$, the algorithm `Approx-Triangles-With-Arboricity-Advice` will either return “bad advice” or output a value of \hat{t} such that $\hat{t} < (1 + 20\epsilon)t$.

Proof. Let $P = (E_0, E_1)$ again denote the partition induced by the algorithm `IsHeavy`. Conditioned on a specific choice of R , the expectation of each χ_i still satisfies Equation 1, and therefore averaging over R we have

$$\mathbb{E}[\hat{t}] = \mathbb{E} \left[\frac{d(R) \cdot m}{|R|} \chi \right] = \mathbb{E} \left[\frac{d(R) \cdot m}{|R|} \cdot \frac{a_P(R)}{d(R)} \right] = \frac{m}{|R|} \cdot \mathbb{E}[a_P(R)].$$

54:16 Testable Algorithms for Approximately Counting Edges and Triangles

By the assignment rule in Definition 12, and by the first part of Claim 15, we have $\mathbb{E}_{e \sim E}[a_P(e)] = \frac{1}{m} \sum_{e \in E} a_P(e) \leq \frac{t}{m}$, and, by linearity of expectation, $\mathbb{E}[a_P(R)] \leq |R| \cdot \frac{t}{m}$. Combining with the inequality above, we get $EX[\hat{t}] \leq t$. Therefore, by Markov's inequality, we have

$$\Pr[\hat{t} > (1 + 20\epsilon)t] < \frac{1}{1 + \epsilon} < 1 - 5\epsilon.$$

Hence, with probability at least 5ϵ , the random variable \hat{t} satisfies $\hat{t} \leq (1 + 20\epsilon)t$. Conditioned on this event, either the algorithm outputs such value of \hat{t} satisfying $\hat{t} \leq (1 + \epsilon)t$, or it terminates and returns “bad advice” at an earlier step. \triangleleft

3.5 From testable triangle-counting to instance-adaptive counting

In this section, we use Theorem 1 to obtain an instance-adaptive algorithm for approximate triangle counting (Corollary 2).

Proof of Corollary 2. We run the algorithm in Theorem 1 with a parameter $\bar{\delta} = \delta/(10 \log m)$ and an increasing sequence of values $\tilde{\alpha} = 2^j$ where $j \in \{1, 2, 3, \dots, \log m\}$. Once for some j the algorithm outputs value \hat{t}_j (Rather than “bad advice”), we output \hat{t}_j and terminate.

The soundness condition in (Definition 6), together with a union bound over j , implies that once we receive an estimate \hat{t}_j at one of the steps j , this estimate will satisfy $\hat{t}_j \in (1 \pm \epsilon)t$.

The true arboricity $\alpha(G)$ is at most $2\sqrt{m}$. For one of the iterations j_0 , we have $\tilde{\alpha}_{j_0} = 2^{j_0} \in [\alpha(G), 2\alpha(G)]$. The completeness condition tells us that for each $j \geq j_0$ with probability at least $1 - \delta$ the algorithm from Theorem 1 terminates. Thus, the expected run-time of the algorithm is upper-bounded by

$$\begin{aligned} j_0 \cdot O^* \left(\frac{m \cdot \tilde{\alpha} \log 1/\delta}{t} + \frac{m \log 1/\delta}{t^{2/3}} \right) + \sum_{j > j_0} \delta^{j-j_0} O^* \left(\frac{m \cdot \tilde{\alpha} \cdot 2^{j-j_0} \log 1/\delta}{t} + \frac{m \log 1/\delta}{t^{2/3}} \right) &\leq \\ j_0 \cdot O^* \left(\frac{m \tilde{\alpha} \log 1/\delta}{t} + \frac{m \log 1/\delta}{t^{2/3}} \right) + \sum_{j > j_0} 0.1^{j-j_0} O^* \left(\frac{m \tilde{\alpha} \cdot 2^{j-j_0} \log 1/\delta}{t} + \frac{m \log 1/\delta}{t^{2/3}} \right) &\leq \\ O^* \left(\frac{m \tilde{\alpha} \cdot 2^{j-j_0} \log 1/\delta}{t} + \frac{m \log 1/\delta}{t^{2/3}} \right), & \end{aligned}$$

finishing the proof. \triangleleft

4 A Lower Bound for Testable Triangle Counting Algorithms

Proof of Theorem 3. Note that if $\tilde{\alpha} < m/n$, then by Theorem 8, any testable algorithm can simply reject, so that proving any non-trivial lower bound is impossible. Hence, the lower bound only holds for the case where $\tilde{\alpha} \geq m/n$.

Since any testable algorithm for triangle counting also implies a triangle-counting algorithm in the case where the graph is guaranteed to be low-arboricity, the lower bound of $\Omega^* \left(\frac{m \cdot \alpha(G)}{t} \right)$ by [28, 10] applies to testable algorithms, and so it remains to prove a lower bound of $\Omega(m/t^{2/3})$ for the regime $\tilde{\alpha} < t^{1/3}$.

Consider the following two families of graphs. In the first family, \mathcal{G}_1 , all graphs have $\Theta(n)$ nodes, $\Theta(m)$ edges, arboricity $\tilde{\alpha}$ and zero triangles. In the second family \mathcal{G}_2 , all graphs have $\Theta(n)$ nodes, $\Theta(m)$ edges, arboricity $t^{1/3}$ and $\Theta(t)$ triangles.

Observe that any testable algorithm must (whp) return $\hat{t} = 0$ for graphs from the first family, and (whp) either output “bad advice” or a correct estimate of t for graphs from the second family. In particular, this implies that the algorithm must be able to distinguish between the two families with high probability.

More concretely, in the first family the graphs all consist of two disjoint sets. The first, A_1 , is a bipartite graph over n vertices where each vertex has degree $2m/n$, and the second, A_2 is a bipartite set over $r = \Theta(t^{2/3}/\tilde{\alpha})$ vertices so that each vertex has degree $\tilde{\alpha}$ and in total there are $\binom{t^{1/3}}{2} = \Theta(t^{2/3})$ edges. The arboricity of these graphs is $\tilde{\alpha}$ due to the set A_2 .

In the second family the graphs are very similar, but the set A_2 is a clique over $t^{1/3}$ vertices, so that here too it contributes $\Theta(t^{2/3})$ edges, however, here it also contributes $\Theta(t)$ triangles. In addition, the graphs in this family have an additional independent set, A_3 over $r - t^{1/3}$ vertices, so that both graph families have the same number of vertices and edges. The arboricity of these graphs is $t^{1/3}$ due to the set A_2 . Within each family the graphs only differ on the labelings of the vertices.

In order to distinguish the two families with high probability, any testable algorithm must hit at least one vertex or edge from $V \setminus A_1$. Hence, it must make

$$\Omega\left(\min\left\{\frac{n}{r}, \frac{m}{t^{2/3}}\right\}\right) = \Omega\left(\min\left\{\frac{n\tilde{\alpha}}{t^{2/3}}, \frac{m}{t^{2/3}}\right\}\right) = \Omega\left(\frac{m}{t^{2/3}}\right)$$

many uniform vertex or edge samples, where the last equality is by the assumption that $\frac{m}{n} < \tilde{\alpha}$. ◀

5 Testable Streaming Algorithm for Counting Triangles

In this section we prove that our testable estimation algorithm can be adapted to the streaming model, achieving 9 rounds and space complexity $O^*\left(\frac{m}{t^{2/3}} + \frac{m\tilde{\alpha}}{t}\right)$. The state-of-the-art arboricity-dependent bound is due to Bera and Seshadhri [6], which uses only 6 rounds and $O^*(m\tilde{\alpha}/t)$ space, but requires a trusted bound on $\alpha(G)$.

As shown in [26, 32], any sublinear-time algorithm in the augmented query model with $O(q)$ queries and adaptivity depth k can be simulated in the arbitrary-order insert-only edge arrival streaming model using $O(q) \cdot \text{poly log } n$ space and k rounds.⁹ That is, query complexity translates into space complexity, while adaptivity depth translates into the number of rounds. This is true since neighbor and edge queries can be implemented with ℓ_0 -samplers, degree queries with counters, and pair queries with a one-bit verifier. Hence, a set of non-adaptive queries can be simulated in a single pass. Inspecting the algorithm `Approx-Triangles-With-Arboricity-Advice`, one can verify that it has an adaptivity depth of 9.

Recall that `Approx-Triangles-With-Arboricity-Advice` receives both an advice value $\tilde{\alpha}$ for $\alpha(G)$ and a guessed value \tilde{t} for the number of triangles in the graph. In the sublinear setting, the latter can be removed using a search procedure. However, this search cannot be implemented in the streaming model with a constant number of passes; it would either introduce a polylogarithmic round overhead or, if executed in parallel, require linear space. This limitation is common in streaming counting algorithms. Consequently, the streaming literature typically assumes access to a coarse constant-factor estimate of the triangle count

⁹ An algorithm \mathcal{A} in the augmented query model has adaptivity depth k if the following holds. For every execution of \mathcal{A} , the set of queries Q performed by it can be partitioned into k sets Q_1, \dots, Q_k , such that for each $j \in [k]$, the queries in Q_j can be determined solely from the responses to Q_1, \dots, Q_{j-1} .

(see, e.g., [14, 49] for a detailed discussion on this assumption). Given such an estimate, `Approx-Triangles-With-Arboricity-Advice` yields a constant-round streaming testable algorithm for approximating the number of triangles, as stated in Corollary 4.

A similar polylogarithmic round overhead would also arise if the algorithm were not provided with the advice value $\tilde{\alpha}$ in advance, and instead had to adapt its space complexity to the true arboricity $\alpha(G)$ of the graph (cf. Corollary 2)

6 Approximating the Number of Edges

In this section we prove Theorem 5 and its following corollary.

► **Corollary 25.** *There exists an algorithm with the following guarantees. The algorithm receives access to a graph G , and is given the number of vertices n in G together with parameters ϵ and $\delta \in (0, 1)$. With probability at least $1 - \delta$, the algorithm produces an estimate \hat{m} that with probability at least $1 - \delta$ satisfies $\hat{m} \in (1 \pm \epsilon)m$, where m is the number of edges in G . The expected run-time of the algorithm is $O^*\left(\frac{n \cdot \alpha(G) \log 1/\delta}{\epsilon}\right)$, where $\alpha(G)$ is the arboricity of G .*

Recall that in this section too the algorithm relies on access to uniform vertex samples, degree queries and neighbor queries, as well as uniform edge samples. However, we remove the assumption on having prior knowledge of m , as this is the task at hand. Instead, the algorithm is given a guess \tilde{m} , and we use the search theorem in order to output a good estimate of m (see Section 6.3).

We now present the algorithm which is a simple adaptation of [57], and the main theorem of this section.

`Approx-Edges-With-Arboricity-Advice`($\tilde{m}, \epsilon, \tilde{\alpha}, \delta$).

1. Let $\delta = \delta/2$ and $\epsilon = \epsilon/6$.
2. Sample a set R of $r = \frac{12 \ln(1/\delta)}{\epsilon^2}$ edges uniformly at random.
3. For every $e \in R$, let $x_e = 1$ if $d(e) > 2\tilde{\alpha}/\epsilon$.
4. If $\frac{1}{r} \sum_{e \in R} x_e > 2\epsilon$ then output “bad advice”.
5. (Otherwise, continue with the bounded-arboricity algorithm of [57])
6. Let $q = \frac{n\tilde{\alpha}}{m} \cdot \frac{12 \ln(2/\delta)}{\epsilon^3}$.
7. If $q \geq n$, query $d(v)$ for all $v \in V$ and **return** $\sum_{v \in V} d(v)$.
8. For $i = 1$ to q :
 - a. Sample $u \in V$ uniformly at random.
 - b. Sample $v \in \Gamma(u)$ uniformly at random.
 - c. If $d(u) \leq 2\tilde{\alpha}/\epsilon$ and $u \prec v$, then let $\chi_i = d(u)$. Otherwise, let $\chi_i = 0$.
9. **Return** $\hat{m} = \frac{n}{q} \sum_i \chi_i$.

► **Theorem 26.** *Let G be the input graph, and ϵ and δ the approximation and failure parameters. Let m be the number of edges in G . In addition to augmented query access to G the algorithm takes as an input a pair of positive integers $\tilde{\alpha}$ and \tilde{m} . Then, the following holds.*

- *For every G and \tilde{m} , if $\alpha(G) \leq \tilde{\alpha}$, then w.p. at least $1 - \delta$ the algorithm `Approx-Edges-With-Arboricity-Advice` will not output “bad advice”.*

- (used for soundness) For every G and $\tilde{\alpha}$, if $\tilde{m} < m$ then with probability at least $1 - \delta$ Approx-Edges-With-Arboricity-Advice will either return “bad advice” or output $\hat{m} \in (1 \pm 10\epsilon)m$.
- (used for soundness) For every G and $\tilde{\alpha}$, if $\tilde{m} > m$ then w.p. at least $\epsilon/4$, the algorithm Approx-Edges-With-Arboricity-Advice will either return “bad advice” or output a value of \hat{m} such that $\hat{m} < (1 + \epsilon)m$.
- The query complexity and running time are the minimum between $O\left(\frac{n\bar{\alpha}}{m}\right) \cdot \text{poly}(1/\epsilon, \ln(1/\delta))$ and $O(n)$.

We start with some preliminary definitions and claims.

6.1 Preliminary notations and claims

► **Definition 27** (Heavy vertices and edges.). Let H denote the set of all vertices with $d(v) > 2\bar{\alpha}/\epsilon$. Let $m(H)$ denote the set of edges in the subgraph induced by H . I.e., edges with both endpoints having degree greater than $2\bar{\alpha}/\epsilon$ (so that $d(e) > 2\bar{\alpha}/\epsilon$). Let $L = V \setminus H$.

► **Notation 28** (An ordering and outgoing edges). We assume an ordering on the graphs’ vertices such that $u \prec v$ if $d(u) < d(v)$ or if $d(u) = d(v)$ and $id(u) < id(v)$. We let $d^+(u)$ denote the number of neighbors of v such that $u \prec v$ and refer to this set as the set of outgoing neighbors of u . For a set S we let $d^+(S) = \sum_{v \in S} d^+(v)$.

Observe that

$$\sum_{v \in V} d^+(v) = d^+(L) + d^+(H) = m.$$

► **Observation 29.** Suppose $\alpha > \bar{\alpha}$. Observe that since every $v \in H$ contributes at least $2\alpha/\epsilon$ edges to m , it holds that $|H| \leq \epsilon m / (2\alpha)$. Thus, by the Nash-Williams theorem, $m(H) \leq 2|H| \cdot \alpha \leq \epsilon m$. Therefore,

$$m(H) = \sum_{v \in H} d^+(v) < \epsilon \cdot m, \text{ so that } d^+(L) \stackrel{\text{def}}{=} \sum_{v \in L} d^+(v) \in [(1 - \epsilon)m, m]$$

6.2 Proof of Theorem 26

Proof of Theorem 26. We prove each of the items separately.

Proof of item 1. By Observation 29, if $\bar{\alpha} > \alpha(G)$, then $m(H) < \epsilon m$. Hence, for every $e \in E$, $\Pr_{e \sim E}[d(e) > 2\bar{\alpha}/\epsilon] < \epsilon$. By Chernoff’s inequality, for $r = 12 \ln(1/\delta)/\epsilon^2$,

$$\Pr \left[\frac{1}{r} \sum_{e \in R} x_e > 2\epsilon \right] < \delta.$$

Hence, the probability of rejecting in Step 4 is at most δ .

Proof of item 2. If $m(H) > 4\epsilon m$, then by Chernoff’s inequality,

$$\Pr \left[\frac{1}{r} \sum_{e \in R} x_e < 2\epsilon \right] < \delta.$$

Therefore, with probability at least $1 - \delta$, the algorithm returns “Bad Advice” in Step 4 and the claim holds.

Now consider the case where $m(H) < 4\epsilon m$. Then $d^+(L) = m - d^+(H) \in [(1 - 4\epsilon)m, m]$.

Let u denote the vertex sampled in Step 8a. If $u \in L$ then $\mathbb{E}_{v \sim \Gamma(u)}[\chi_i \mid u, u \in L] = \frac{d^+(u)}{d(u)} \cdot d(u) = d^+(u)$. Also, for any $u \in H$, $\mathbb{E}[\chi_i \mid u, u \in H] = 0$. Hence, by linearity of expectation,

$$\mathbb{E}_{u \in V, v \in \Gamma(u)}[\chi_i] = \frac{1}{n} \left(\sum_{u \in L} d^+(u) + \sum_{u \in H} 0 \right) = \frac{d^+(L)}{n}. \quad (2)$$

Furthermore, for any u , $\chi_i \leq 2\bar{\alpha}/\epsilon$. Therefore, the χ_i 's are $[0, B]$ random variables with $B = 2\bar{\alpha}/\epsilon$ and $\mu = \mathbb{E}[\chi_i] = \frac{d^+(L)}{n} \geq \frac{\frac{1}{2}m}{n}$. Since $r = \frac{n \cdot (2\bar{\alpha}/\epsilon)}{\frac{1}{2}m} \cdot \frac{3 \ln(2/\epsilon)}{\epsilon^2} \geq \frac{B}{\mu} \cdot \frac{3 \ln(2/\epsilon)}{\epsilon^2}$, by the Chernoff bound, $\hat{\mu} \stackrel{\text{def}}{=} \frac{1}{r} \sum_{i=1}^r \chi_i \in (1 \pm \epsilon) \cdot \frac{d^+(L)}{n} \in (1 \pm 10\epsilon) \frac{m}{n}$. Hence, the returned value $\hat{m} = \frac{n}{r} \sum_{i=1}^r \chi_i$ in Step 9 is with probability at least $1 - \delta$ in $(1 \pm 10\epsilon)m$.

Proof of item 3. By Equation 2 and linearity of expectation,

$$\mathbb{E} \left[\frac{1}{q} \sum_i \chi_i \right] = \mathbb{E}[\chi_1] = \frac{d^+(L)}{n} \leq \frac{m}{n}.$$

By Markov's inequality,

$$\Pr \left[\frac{1}{q} \sum_i \chi_i > (1 + \epsilon)m \right] < \frac{1}{1 + \epsilon} < 1 - \epsilon/4.$$

Proof of item 4. The query complexity and running time are

$$r + \min n, q = O(\min\{n, q\}) = \min \left\{ n, \frac{n\bar{\alpha}}{m} \cdot \frac{\ln(1/\delta)}{\epsilon^3} \right\}. \quad \blacktriangleleft$$

6.3 Searching for the “right” \tilde{m} value

We refer the reader to the full version of this work [30] for details on searching for the right \tilde{m} value serving as a rough estimate of the true edge count m .

References

- 1 Mohammad Al Hasan and Vachik S Dave. Triangle counting in large networks: a review. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 8(2):e1226, 2018. doi:10.1002/WIDM.1226.
- 2 Maryam Aliakbarpour, Amartya Shankha Biswas, Themis Gouleakis, John Peebles, Ronitt Rubinfeld, and Anak Yodpinyanee. Sublinear-time algorithms for counting star subgraphs via edge sampling. *Algorithmica*, 80(2):668–697, 2018. doi:10.1007/S00453-017-0287-3.
- 3 Sepehr Assadi, Michael Kapralov, and Sanjeev Khanna. A Simple Sublinear-Time Algorithm for Counting Arbitrary Subgraphs via Edge Sampling. In *Innovations in Theoretical Computer Science Conference (ITCS)*, pages 6:1–6:20, 2019. doi:10.4230/LIPIcs.ITCS.2019.6.
- 4 Sepehr Assadi and Hoai-an Nguyen. Asymptotically optimal bounds for estimating h-index in sublinear time with applications to subgraph counting. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2022, September 19-21, 2022, University of Illinois, Urbana-Champaign, USA (Virtual Conference)*, 2022. doi:10.4230/LIPIcs.APPROX/RANDOM.2022.48.
- 5 Bahman Bahmani, Ravi Kumar, and Sergei Vassilvitskii. Densest subgraph in streaming and mapreduce. *Proceedings of the VLDB Endowment*, 5(5), 2012. doi:10.14778/2140436.2140442.

- 6 Suman K. Bera and C. Seshadhri. How the degeneracy helps for triangle counting in graph streams. In *Principles of Database Systems (PODS)*, pages 457–467, 2020. doi:10.1145/3375395.3387665.
- 7 Lorenzo Beretta, Deeparnab Chakrabarty, and C Seshadhri. Faster estimation of the average degree of a graph using random edges and structural queries. *arXiv preprint arXiv:2507.06925*, 2025. doi:10.48550/arXiv.2507.06925.
- 8 Lorenzo Beretta and Jakub Tětek. Better sum estimation via weighted sampling. *ACM Transactions on Algorithms*, 20(3):1–33, 2024.
- 9 Sayan Bhattacharya, Monika Henzinger, Danupon Nanongkai, and Charalampos E Tsourakakis. Space-and time-efficient algorithm for maintaining dense subgraphs on one-pass dynamic streams. In *Forty-Seventh Annual ACM Symposium on Theory of Computing*, pages 173–182. ACM, 2015. doi:10.1145/2746539.2746592.
- 10 Arijit Bishnu, Debarshi Chanda, and Gopinath Mishra. Arboricity and random edge queries matter for triangle counting using sublinear queries. *arXiv preprint arXiv:2502.15379*, 2025. doi:10.48550/arXiv.2502.15379.
- 11 Amartya Shankha Biswas, Talya Eden, Quanquan C Liu, Ronitt Rubinfeld, and Slobodan Mitrović. Massively parallel algorithms for small subgraph counting. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2022)*, pages 39:1–39:28. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPIcs.APPROX/RANDOM.2022.39.
- 12 Amartya Shankha Biswas, Talya Eden, and Ronitt Rubinfeld. Towards a decomposition-optimal algorithm for counting and sampling arbitrary motifs in sublinear time. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPIcs.APPROX/RANDOM.2021.55.
- 13 Cristian Boldrin and Fabio Vandin. Fast and accurate triangle counting in graph streams using predictions. In *2024 IEEE International Conference on Data Mining (ICDM)*, pages 31–40. IEEE, 2024. doi:10.1109/ICDM59182.2024.00010.
- 14 Vladimir Braverman, Rafail Ostrovsky, and Dan Vilenchik. How hard is counting triangles in the streaming model? In *International Colloquium on Automata, Languages, and Programming*, pages 244–254. Springer, 2013. doi:10.1007/978-3-642-39206-1_21.
- 15 Keren Censor-Hillel, Dean Leitersdorf, and David Vulakh. Deterministic near-optimal distributed listing of cliques. In *Proceedings of the 2022 ACM Symposium on Principles of Distributed Computing*, pages 271–280, 2022. doi:10.1145/3519270.3538434.
- 16 Gautam Chandrasekaran, Adam R Klivans, Vasilis Kontonis, Konstantinos Stavropoulos, and Arsen Vasilyan. Efficient discrepancy testing for learning with distribution shift. *Advances in Neural Information Processing Systems*, 37, 2024.
- 17 Justin Y Chen, Talya Eden, Piotr Indyk, Honghao Lin, Shyam Narayanan, Ronitt Rubinfeld, Sandeep Silwal, Tal Wagner, David P Woodruff, and Michael Zhang. Triangle and four cycle counting with predictions in graph streams. In *International Conference on Learning Representations (ICLR)*, 2022.
- 18 Norishige Chiba and Takao Nishizeki. Arboricity and subgraph listing algorithms. *SIAM Journal on Computing*, 14(1):210–223, 1985. doi:10.1137/0214017.
- 19 Artur Czumaj and Christian Konrad. Detecting cliques in congest networks. *Distributed Computing*, 33(6):533–543, 2020. doi:10.1007/S00446-019-00368-W.
- 20 Maximilien Danisch, Oana Denisa Balalau, and Mauro Sozio. Listing k -cliques in sparse real-world graphs. In *Proceedings of the 2018 World Wide Web Conference on World Wide Web, WWW 2018*, pages 589–598, 2018.
- 21 Ilias Diakonikolas, Daniel Kane, Vasilis Kontonis, Sihan Liu, and Nikos Zarifis. Efficient testable learning of halfspaces with adversarial label noise. *Advances in Neural Information Processing Systems*, 36, 2023.

- 22 Ilias Diakonikolas, Daniel Kane, Sihan Liu, and Nikos Zarifis. Testable learning of general halfspaces with adversarial label noise. In *The Thirty Seventh Annual Conference on Learning Theory*, pages 1308–1335. PMLR, 2024. URL: <https://proceedings.mlr.press/v247/diakonikolas24a.html>.
- 23 Talya Eden, Amit Levi, Dana Ron, and C. Seshadhri. Approximately counting triangles in sublinear time. *SIAM Journal on Computing*, 46(5):1603–1646, 2017. doi:10.1137/15M1054389.
- 24 Talya Eden, Reut Levi, and Dana Ron. Testing bounded arboricity. *ACM Transactions on Algorithms (TALG)*, 16(2):1–22, 2020. doi:10.1145/3381418.
- 25 Talya Eden, Reut Levi, Dana Ron, and Ronitt Rubinfeld. Approximately counting and sampling hamiltonian motifs in sublinear time. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*, pages 1043–1054, 2025. doi:10.1145/3717823.3718160.
- 26 Talya Eden, Saleet Mossel, and Dana Ron. Approximating the arboricity in sublinear time. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2404–2425. SIAM, 2022. doi:10.1137/1.9781611977073.96.
- 27 Talya Eden, Dana Ron, and C. Seshadhri. Sublinear time estimation of degree distribution moments: The arboricity connection. *SIAM Journal on Discrete Math*, 33(4):2267–2285, 2019. doi:10.1137/17M1159014.
- 28 Talya Eden, Dana Ron, and C. Seshadhri. Faster sublinear approximation of the number of k -cliques in low-arboricity graphs. In *Symposium on Discrete Algorithms (SODA)*, pages 1467–1478. SIAM, 2020. doi:10.1137/1.9781611975994.89.
- 29 Talya Eden, Dana Ron, and C. Seshadhri. On approximating the number of k -cliques in sublinear time. *SIAM Journal on Computing*, 49(4):747–771, 2020. doi:10.1137/18M1176701.
- 30 Talya Eden, Ronitt Rubinfeld, and Arsen Vasilyan. Testable algorithms for approximately counting edges and triangles in sublinear time and space. *Preprint ArXiv:2509.20351*, 2025.
- 31 Hendrik Fichtenberger, Mingze Gao, and Pan Peng. Sampling arbitrary subgraphs exactly uniformly in sublinear time. In *47th International Colloquium on Automata, Languages, and Programming, ICALP*, 2020.
- 32 Hendrik Fichtenberger and Pan Peng. Approximately counting subgraphs in data streams. In *Proceedings of the 41st ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*, pages 413–425, 2022. doi:10.1145/3517804.3524145.
- 33 Orr Fischer, Tzlil Gonen, Fabian Kuhn, and Rotem Oshman. Possibilities and impossibilities for distributed subgraph detection. In *Proceedings of the 30th on Symposium on Parallelism in Algorithms and Architectures*, pages 153–162, 2018. doi:10.1145/3210377.3210401.
- 34 Gaurav Goel and Jens Gustedt. Bounded arboricity to determine the local structure of sparse graphs. In *International Workshop on Graph-Theoretic Concepts in Computer Science*, pages 159–167. Springer, 2006. doi:10.1007/11917496_15.
- 35 Surbhi Goel, Adam R Klivans, Konstantinos Stavropoulos, and Arsen Vasilyan. Testing noise assumptions of learning algorithms. *arXiv preprint arXiv:2501.09189*, 2025. doi:10.48550/arXiv.2501.09189.
- 36 Surbhi Goel, Abhishek Shetty, Konstantinos Stavropoulos, and Arsen Vasilyan. Tolerant algorithms for learning with arbitrary covariate shift. *Advances in Neural Information Processing Systems*, 37:124979–125018, 2024.
- 37 Aravind Gollakota, Parikshit Gopalan, Adam Klivans, and Konstantinos Stavropoulos. Agnostically learning single-index models using omnipredictors. *Advances in Neural Information Processing Systems*, 36, 2024.
- 38 Aravind Gollakota, Adam Klivans, Konstantinos Stavropoulos, and Arsen Vasilyan. Tester-learners for halfspaces: Universal algorithms. *Advances in Neural Information Processing Systems*, 36, 2023.
- 39 Mira Gonen, Dana Ron, and Yuval Shavitt. Counting stars and other small subgraphs in sublinear-time. *SIAM Journal on Discrete Mathematics*, 25(3):1365–1411, 2011. doi:10.1137/100783066.

- 40 Alon Itai and Michael Rodeh. Finding a minimum circuit in a graph. In *Proceedings of the ninth annual ACM symposium on Theory of computing*, pages 1–10, 1977.
- 41 Shweta Jain and C. Seshadhri. A fast and provable method for estimating clique counts using Turán’s theorem. In *Proceedings of the 26th international conference on world wide web*, pages 441–449, 2017. doi:10.1145/3038912.3052636.
- 42 John Kallaughar, Andrew McGregor, Eric Price, and Sofya Vorotnikova. The complexity of counting cycles in the adjacency list streaming model. In *Proceedings of the 38th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*, pages 119–133, 2019. doi:10.1145/3294052.3319706.
- 43 Valerie King, Alex Thomo, and Quinton Yong. Computing $(1+\epsilon)$ -approximate degeneracy in sublinear time. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, pages 2160–2168. International Joint Conferences on Artificial Intelligence Organization, 2023. doi:10.24963/IJCAI.2023/240.
- 44 Adam Klivans, Konstantinos Stavropoulos, and Arsen Vasilyan. Learning intersections of halfspaces with distribution shift: Improved algorithms and sq lower bounds. In Shipra Agrawal and Aaron Roth, editors, *Proceedings of Thirty Seventh Conference on Learning Theory*, volume 247 of *Proceedings of Machine Learning Research*, pages 2944–2978. PMLR, 30 June–03 July 2024. URL: <https://proceedings.mlr.press/v247/klivans24b.html>.
- 45 Adam Klivans, Konstantinos Stavropoulos, and Arsen Vasilyan. Testable learning with distribution shift. In Shipra Agrawal and Aaron Roth, editors, *Proceedings of Thirty Seventh Conference on Learning Theory*, volume 247 of *Proceedings of Machine Learning Research*, pages 2887–2943. PMLR, 30 June–03 July 2024. URL: <https://proceedings.mlr.press/v247/klivans24a.html>.
- 46 Tamara G Kolda, Ali Pinar, Todd Plantenga, C Seshadhri, and Christine Task. Counting triangles in massive graphs with mapreduce. *SIAM Journal on Scientific Computing*, 36(5):S48–S77, 2014. doi:10.1137/13090729X.
- 47 Cassandra Marcussen, Ronitt Rubinfeld, and Madhu Sudan. Quality control in sublinear time: a case study via random graphs. *arXiv:2508.16531*, 2025. doi:10.48550/arXiv.2508.16531.
- 48 Andrew McGregor, David Tench, Sofya Vorotnikova, and Hoa T Vu. Densest subgraph in dynamic graph streams. In *International Symposium on Mathematical Foundations of Computer Science*, pages 472–482. Springer, 2015. doi:10.1007/978-3-662-48054-0_39.
- 49 Andrew McGregor, Sofya Vorotnikova, and Hoa T Vu. Better algorithms for counting triangles in data streams. In *Proceedings of the 35th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*, pages 401–411, 2016. doi:10.1145/2902251.2902283.
- 50 Rajeev Motwani, Rina Panigrahy, and Ying Xu. Estimating sum by weighted sampling. In *International Colloquium on Automata, Languages, and Programming*, pages 53–64. Springer, 2007. doi:10.1007/978-3-540-73420-8_7.
- 51 C St JA Nash-Williams. Edge-disjoint spanning trees of finite graphs. *Journal of the London Mathematical Society*, 1(1):445–450, 1961.
- 52 C St JA Nash-Williams. Decomposition of finite graphs into forests. *Journal of the London Mathematical Society*, 1(1):12–12, 1964.
- 53 Jaroslav Nešetřil and Svatopluk Poljak. On the complexity of the subgraph problem. *Comment. Math. Univ. Carol.*, 26(2):415–419, 1985.
- 54 Rasmus Pagh and Charalampos E Tsourakakis. Colorful triangle counting and a mapreduce implementation. *Information Processing Letters*, 112(7):277–281, 2012. doi:10.1016/J.IPL.2011.12.007.
- 55 Ha-Myung Park and Chin-Wan Chung. An efficient mapreduce algorithm for counting triangles in a very large graph. In *Proceedings of the 22nd ACM international conference on Information & Knowledge Management*, pages 539–548, 2013. doi:10.1145/2505515.2505563.
- 56 Ronitt Rubinfeld and Arsen Vasilyan. Testing distributional assumptions of learning algorithms. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing, STOC 2023, Orlando, FL, USA, June 20-23, 2023*, pages 1643–1656. ACM, 2023. doi:10.1145/3564246.3585117.

54:24 Testable Algorithms for Approximately Counting Edges and Triangles

- 57 C. Seshadhri. An extremely simple algorithm for approximate edge counting in bounded arboricity graphs. Unpublished manuscript, 2025.
- 58 Kijung Shin, Tina Eliassi-Rad, and Christos Faloutsos. Patterns and anomalies in k-cores of real-world graphs with applications. *Knowledge and Information Systems*, 54(3):677–710, 2018. doi:10.1007/S10115-017-1077-6.
- 59 Lucas Slot, Stefan Tiegel, and Manuel Wiedmer. Testably learning polynomial threshold functions. *Advances in Neural Information Processing Systems*, 37, 2024.