



Cardinality Estimation of Subgraph Matching: A Filtering-Sampling Approach

Wonseok Shin

Seoul National University
wsshin@theory.snu.ac.kr

Siwoo Song

Seoul National University
swsong@theory.snu.ac.kr

Kunsoo Park*

Seoul National University
kpark@theory.snu.ac.kr

Wook-Shin Han

POSTECH
wshan@dbllab.postech.ac.kr

ABSTRACT

Subgraph counting is a fundamental problem in understanding and analyzing graph structured data, yet computationally challenging. This calls for an accurate and efficient algorithm for Subgraph Cardinality Estimation, which is to estimate the number of all isomorphic embeddings of a query graph in a data graph. We present FaST_{est}, a novel algorithm that combines (1) a powerful filtering technique to significantly reduce the sample space, (2) an adaptive tree sampling algorithm for accurate and efficient estimation, and (3) a worst-case optimal stratified graph sampling algorithm for hard instances. Extensive experiments on real-world datasets show that FaST_{est} outperforms state-of-the-art sampling-based methods by up to two orders of magnitude and GNN-based methods by up to three orders of magnitude in terms of accuracy.

PVLDB Reference Format:

Wonseok Shin, Siwoo Song, Kunsoo Park, and Wook-Shin Han. Cardinality Estimation of Subgraph Matching: A Filtering-Sampling Approach. PVLDB, 17(7): 1697 - 1709, 2024.
doi:10.14778/3654621.3654635

PVLDB Artifact Availability:

The source code, data, and/or other artifacts have been made available at <https://github.com/SNUCSE-CTA/FaSTest>.

1 INTRODUCTION

Subgraph matching is the fundamental problem in understanding and analyzing graph structured data [38]. Given a data graph and a query graph, subgraph matching is the problem of finding all isomorphic embeddings of the query graph in the data graph. Identifying the occurrences of specific subgraph patterns is crucial for various applications such as analyzing protein-protein interaction networks [7, 33], revealing patterns and trends of user interactions in social networks [16], and optimizing queries in relational databases [25]. Subgraph counting, the problem of counting the number of embeddings, is also of paramount importance in various applications, such as designing graph kernels [21, 39] and understanding biological networks [37, 52].

Both subgraph matching and subgraph counting are computationally challenging. Specifically, subgraph matching is NP-hard

*Corresponding author

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.

Proceedings of the VLDB Endowment, Vol. 17, No. 7 ISSN 2150-8097.
doi:10.14778/3654621.3654635

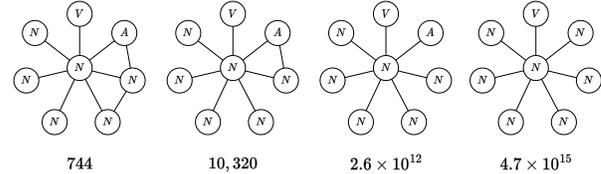


Figure 1: Query graphs and their number of embeddings in WordNet dataset ($\approx 80K$ vertices). N, V, and A represent Noun, Verb, and Adjective, respectively.

[18]. Subgraph counting is #P-hard, since subgraph counting in planar graphs (subgraph counting restricted to both data graph and query graph being planar) is #P-hard [22]. This challenge is further amplified by the exponential number of possible embeddings. For instance, Figure 1 shows some query graphs with 8 vertices and the number of embeddings for each query graph in the WordNet graph [46]. The number of embeddings varies from 744 to 4.7×10^{15} with only small changes in the query graphs. Consequently, numerous subgraph matching algorithms [6, 18–20, 24], which are primarily implemented for counting the number of embeddings rather than finding the actual embeddings, typically count at most k (usually 10^3 to 10^5) embeddings, instead of counting all embeddings. As a result, there is a pressing need for an accurate and efficient algorithm for *Subgraph Cardinality Estimation*, which is to estimate the number of all isomorphic embeddings.

Existing Approaches and Limitations. Subgraph cardinality estimation has been extensively researched due to its diverse real-world applications. In particular, the counting of small subgraphs (also known as graphlets or motifs) has been widely studied. IMPR [12] proposed a random walk-based method for counting graphlets, while MOTIVO [9] developed an adaptive graphlet sampling algorithm utilizing the color-coding technique. However, such works primarily focus on small subgraphs.

Recently, utilizing Graph Neural Networks (GNN) for subgraph cardinality estimation is gaining interest. LSS [50] and NeurSC [48] have explored GNNs for subgraph cardinality estimation. However, their accuracy for difficult instances with large queries remains unsatisfactory.

Approximate counting of homomorphic embeddings and estimating join cardinalities are also closely related problems to subgraph cardinality estimation. Existing works in these areas mainly utilize summarization or sampling-based approaches. A comprehensive survey [35] demonstrated that WanderJoin [27], which employs random walks for sampling, outperforms other summarization and sampling methods for counting homomorphic embeddings. Alley [26] extends WanderJoin by proposing random walk with intersection to improve accuracy. Despite the success of such methods, they often encounter sampling failure, particularly for instances with complex label distributions and large queries [50].

Contributions. In this paper, we present a new algorithm FaST_{est} (Filtering and Sampling Techniques for subgraph cardinality estimation), addressing the limitations of existing methods.

We propose a novel *Filtering-Sampling* approach for cardinality estimation as follows. To approximate the cardinality of a set whose elements are difficult to enumerate or count, sampling is a widely used approach. For a query graph q and a data graph G , let \mathcal{M} denote the set of all isomorphic embeddings of q in G . A sampling algorithm first defines a sample space Ω which is a superset of \mathcal{M} , and aims to approximate the ratio $\rho = |\mathcal{M}|/|\Omega|$. Assume that (1) obtaining the cardinality of Ω is easier than that of \mathcal{M} , (2) Ω is amenable to random sampling, and (3) verification of whether a random sample $x \in \Omega$ is in \mathcal{M} can be done efficiently. By obtaining uniform random samples from Ω and using the empirical ratio $\hat{\rho}$ which is the proportion of random samples lying in \mathcal{M} as an estimator to ρ , the estimate $|\Omega|\hat{\rho}$ is an unbiased consistent estimator [11] for $|\mathcal{M}|$.

Our approach reduces the size of sample space Ω greatly by using filtering, while retaining all embeddings (i.e., keeping $|\mathcal{M}|$ unchanged). That is, we increase the ratio $\rho = |\mathcal{M}|/|\Omega|$ by decreasing $|\Omega|$, which makes sampling more accurate and more efficient. (Previous approaches take samples from the whole sample space Ω , and thus the ratio ρ is very small in many cases, thereby causing sampling failures.)

- We build an auxiliary data structure *Candidate Space* (CS) which contains candidates for vertex mapping and edge mapping. By employing novel safety conditions—Triangle Safety, Four-Cycle Safety, and Edge Bipartite Safety—and the refinement order called Promising First Candidate Filtering, we reduce the candidate edges by up to 80% compared to the filterings of the state-of-the-art subgraph matching algorithms [24, 43], resulting in a compact CS.
- We develop a *candidate tree sampling* algorithm, which performs uniform sampling of spanning trees of the query graph in the compact CS. Our tree sampling employs an adaptive strategy to determine the sample size using the Clopper-Pearson confidence interval, thereby achieving efficient and accurate estimation with rigorous probabilistic guarantees.
- We devise a *stratified graph sampling* algorithm that obtains samples from diverse regions of the sample space, achieving outstanding accuracy on hard instances with a small sample size. Our graph sampling algorithm attains a worst-case optimal time complexity of $O(AGM(q))$ for the query graph q , which is a tight upper bound for $|\mathcal{M}|$ proposed by Atserias, Grohe and Marx [2, 3], guaranteeing the same time complexity even when the sampling ratio is 100%.

We demonstrate that FaST_{est} shows significant improvements in both accuracy and efficiency for subgraph cardinality estimation through extensive experiments on well-known real-world datasets. Specifically, FaST_{est} outperforms state-of-the-art sampling-based methods by up to two orders of magnitude and GNN-based methods by up to three orders of magnitude in terms of accuracy.

Organization. The rest of the paper is organized as follows. Section 2 introduces the definitions and the problem statement. Section 3 outlines an overview of FaST_{est}. Section 4 discusses the candidate filtering algorithm. Section 5 describes the candidate tree

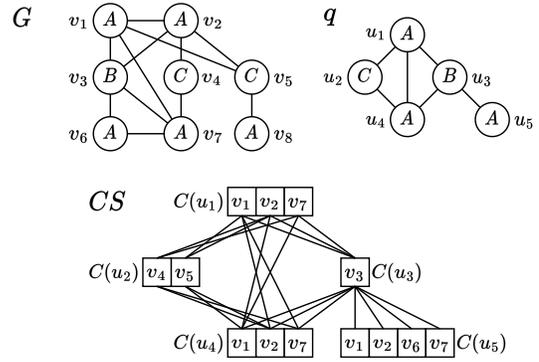


Figure 2: Data graph G , query graph q , and Candidate Space

sampling algorithm. Section 6 explains the stratified graph sampling algorithm. Section 7 presents the experimental results, and Section 8 concludes the paper. The Appendix, containing omitted proofs, is available online [40].

2 PRELIMINARIES

2.1 Problem Statement

In this paper, we mainly consider undirected and connected graphs with vertices labeled, while the techniques we propose can be extended to directed, disconnected, and edge-labeled graphs. We define a graph $g = (V_g, E_g, L_g)$ as a triplet of a set of vertices V_g , a set of edges E_g , and a labeling function $L_g : V_g \rightarrow \Sigma_V$ where Σ_V is a set of possible labels. Given query graph $q = (V_q, E_q, L_q)$ and data graph $G = (V_G, E_G, L_G)$, a vertex-mapping function $M : V_q \rightarrow V_G$ is called an *isomorphic embedding* when three conditions are met: (1) M is injective, i.e., $M(u) \neq M(u')$ for $u \neq u' \in V_q$. (2) M preserves all adjacency relationships, i.e., $(M(u), M(u')) \in E_G$ for every $(u, u') \in E_q$. (3) M preserves all vertex labels, i.e., $L_G(M(u)) = L_q(u)$ for every $u \in V_q$. We say that q is *subgraph isomorphic* to G when such an embedding exists. For example, in Figure 2, $\{(u_1, v_1), (u_2, v_5), (u_3, v_3), (u_4, v_2), (u_5, v_6)\}$ is an isomorphic embedding of q in G , as it is injective and it preserves all vertex labels and adjacency relationships. When a vertex-mapping function satisfies only (2) and (3), we call it a *homomorphic embedding*.

Table 1 lists the notations that are frequently used in the paper.

Subgraph Matching and Counting. For query graph q and data graph G , let \mathcal{M} denote the set of all *embeddings* of q in G . The *subgraph matching* problem is to find \mathcal{M} exactly, while the *subgraph*

Table 1: Frequently Used Notations

Notation	Definition
q, G	Query graph and data graph
V_g, E_g, L_g	Vertices, edges, and labels of a graph g
\mathcal{M}	Set of isomorphic embeddings of q in G
$d_g(x)$	Degree of a node x in a graph g
$N_g(x)$	Neighbors of x in a graph g
$d_g(x, l)$	Number of vertices in $N_g(x)$ with label l in a graph g
δ_g, Δ_g	Degeneracy [28] and maximum degree of a graph g
T_q	A spanning tree of q
$C(u)$	Set of candidate vertices for u
$C(u' u, v)$	Set of candidate neighbors of (u, v) to u'
$E_{CS}(u, u')$	Set of candidate edges for (u, u')

counting problem is to find $|\mathcal{M}|$. It is clear that all subgraph matching algorithms can also be used for counting.

Problem Statement. Given a query graph q and a data graph G , the *subgraph cardinality estimation* problem is to approximate $|\mathcal{M}|$, the cardinality of the set of all isomorphic embeddings of q in G .

2.2 Related Works

Subgraph Matching and Counting. In the recent decade, numerous solutions have been proposed for both exact and approximate subgraph matching and subgraph counting.

A wealth of results for subgraph matching [5, 6, 8, 15, 18, 19, 24, 43] has been established based on Ullmann’s backtracking framework [45]. Recent works such as [6, 18, 24, 43] share the filtering-backtracking approach. These works first build an auxiliary data structure such as Compact Path Index (CPI) [6], Bigraph Index (BI) [43] and Candidate Space (CS) [18, 24], utilizing diverse filtering strategies such as extended DAG-graph DP using Neighbor Safety [24]. These algorithms utilize such data structures to reduce search space for backtracking, which is further accelerated by strategic selection of matching orders such as Candidate-size order [18]. Interested readers can refer to a detailed survey of the aforementioned algorithms [42].

Approximate solutions are split into three sub-classes: summarization-based, sampling-based, and machine learning-based approaches. Summarization methods, such as SumRDF [41] and Characteristic Sets [34], construct a summary data structure of the data graph. Estimates for query graph q are then processed efficiently by querying the summary structure, typically by decomposing q into smaller substructures such as stars and aggregating the results computed for each substructure. However, these methods may produce highly inaccurate results [35].

Sampling algorithms have also been widely employed for subgraph counting [4, 23, 49]. WanderJoin [27] performs random walks over relations and uses Horvitz-Thompson estimation, with the calculated weights of each walk. JSub [51] obtains independent uniform samples from join results by computing upper bounds for intermediate join sizes. Alley [26] proposes the random walk with intersection and it also uses Tangled Pattern Index (TPI), which is a summary structure obtained by mining difficult patterns.

Machine learning-based solutions, particularly those utilizing Graph Neural Networks (GNN), have recently begun to emerge. Works like LSS [50] and NeurSC [48] have explored the potential of GNNs for subgraph cardinality estimation. While machine learning enables highly efficient query processing, these solutions do not provide an unbiased estimation.

AGM Bound and Worst-Case Optimality. The AGM bound [3] is a tight upper bound for the number of embeddings of a query graph q in a data graph G . An algorithm for subgraph matching or cardinality estimation is referred to as *worst-case optimal* if it guarantees $O(AGM(q))$ time complexity [26, 31]. SSTE [2] developed an edge sampling algorithm for estimating arbitrary-sized subgraph counts in $\tilde{O}(AGM(q)/OUT(q))$ time, where $OUT(q)$ refers to the number of embeddings. While it provides strong theoretical guarantees, its performance on real-world graphs is not impressive [13].

Algorithm 1: OVERVIEW OF FaST_{est}

Input: A data graph G , a query graph q
Output: Estimated number of embeddings of q in G

- 1 CS \leftarrow BuildCompactCandidateSpace(G, q);
- 2 (estimate, #successes) \leftarrow CandidateTreeSampling(CS, q);
- 3 **if** #successes is large enough **then return** estimate;
- 4 **else return** CandidateGraphSampling(CS, q , #successes);

3 OVERVIEW OF ALGORITHM

Filtering-Sampling Approach. Since the problem of deciding whether q is subgraph isomorphic to G is NP-hard, exact enumeration or counting of embeddings is often computationally infeasible. To approximate the cardinality of a set \mathcal{M} whose elements are difficult to enumerate or count, sampling is a widely used approach. A sampling algorithm first defines a sample space Ω which is a superset of \mathcal{M} , and aims to approximate the ratio $\rho = |\mathcal{M}|/|\Omega|$.

Our approach reduces the size of sample space Ω greatly, while retaining all embeddings. We develop a filtering algorithm to reduce the sample space by removing unnecessary vertices and edges in finding all embeddings.

Filtering. We build an auxiliary data structure *Candidate Space* (CS) containing candidates for vertices and edges. To build a compact CS, we perform iterative refinements for each vertex $u \in V_q$. The refinement step involves checking a set of necessary safety conditions that candidates must satisfy, which enables us to remove invalid candidates.

- *Generic Framework.* We propose the generalized framework for candidate filtering, where different conditions and refinement strategies can be implemented. (Section 4.1)
- *Safety Conditions.* We propose novel safety conditions: *Triangle Safety*, *Four-Cycle Safety*, and *Edge Bipartite Safety*. These conditions effectively reduce the CS by better utilizing local substructures. (Section 4.2)
- *Refinement Order.* We suggest *Promising First Candidate Filtering* as the refinement order, where in each iteration a candidate set most likely to be reduced is chosen to be refined. (Section 4.3)

The combined filtering algorithm we develop is much stronger than those in [18, 24, 43] in terms of filtering power, resulting in a compact CS.

Sampling. We develop a sampling algorithm utilizing the compact CS for accurate and efficient subgraph cardinality estimation.

- *Candidate Tree Sampling.* We choose a spanning tree T_q of q , and define the sample space Ω as the set of homomorphisms (called candidate trees) of T_q in the compact CS. We count the number of candidate trees and sample them uniformly at random. (Section 5)
- *Stratified Graph Sampling.* We develop a graph sampling algorithm to handle hard cases in which obtaining accurate results via tree sampling is not possible in a reasonable time. We propose a novel sampling method inspired by stratification to sample from diverse regions of the sample space, achieving worst-case optimality. (Section 6)

The outline of our algorithm FaST_{est} is shown in Algorithm 1.

4 CANDIDATE FILTERING

Candidate Space (CS) is an auxiliary data structure built for efficient subgraph matching, proposed by DAF [18] and extended by VEQ [24]. Here we extend the definition of CS by explicitly maintaining candidate edges, which is necessary for better filtering.

4.1 Candidate Space

Definition 4.1 (Candidate Space). For a vertex $u \in V_q$ and its neighbor $u' \in N_q(u)$, we define the following terms.

- *Candidate set* $C(u)$: a subset of V_G such that if there exists an embedding M of q in G that maps u to v , v must be included in $C(u)$.
- *Candidate edges* $E_{CS}(u, u')$: a subset of E_G such that if there exists an embedding M of q in G that maps u to v and u' to v' , (v, v') must be included in $E_{CS}(u, u')$.

The *Candidate Space* is a collection of candidate sets and candidate edges for query graph q and data graph G .

By explicitly maintaining candidate edges, we can remove an edge (v, v') from candidate edges if it does not belong to any embeddings, even though vertices v and v' remain in the Candidate Space. For $v \in C(u)$ and $u' \in N_q(u)$, we define the *candidate neighbors* of (u, v) to u' as the vertex candidates of u' when u is mapped to v , i.e., $C(u' | u, v) = \{v' | (v, v') \in E_{CS}(u, u')\}$.

We build the initial CS as follows. For each vertex u , initial candidate set $C_{init}(u)$ is defined as the set of vertices $v \in V_G$ with the following two conditions: (1) v has the same label as u ; (2) for each label $l \in \Sigma$, $d_G(v, l) \geq d_q(u, l)$. We initialize candidate edges as $E_{CS}(u, u') = \{(v, v') \in E_G | v \in C(u) \text{ and } v' \in C(u')\}$. The initial candidate space can be constructed in $O(|E_q||E_G|)$ time.

Example 4.2. Figure 2 depicts an example of the initial CS given query graph q and data graph G . Among the vertex pairs with matching labels, v_6 is not in $C(u_1)$ since $d_G(v_6, C) < d_q(u_1, C)$. Other absences can also be seen analogously. In the CS, $C(u_4 | u_1, v_1)$, the candidate neighbors of (u_1, v_1) to u_4 , is $\{v_2, v_7\}$ as they are neighbors of v_1 which are in $C(u_4)$.

By definition, all embeddings of q in G are preserved in CS, as all possible vertices and edges constituting an embedding are present. We call this property of CS as *complete*. In our sampling algorithm later to be described, reducing the size of a candidate space translates to a smaller sample space. Therefore, the key is to refine the candidate space as much as possible while retaining completeness. Note that CS can contain $O(|V_q||V_G|)$ vertices and $O(|E_q||E_G|)$ edges in the worst case.

Generic Filtering Framework. We start the refinement from the initial CS. For a predefined set of safety conditions, the generic filtering framework can be defined as Algorithm 2. The framework consists of several design choices; a set of *safety conditions* to determine validity of candidate vertices or edges, a method to decide the *refinement order*, i.e., the order of query vertices whose candidate set is refined, and the *stopping criteria* by which we finish the refinement.

4.2 Safety Conditions

It is necessary that $C(u' | u, v)$ is nonempty for v to be a candidate for u for each neighbor u' of u . Together with this condition, VEQ [24] suggested the use of *safety condition* that can further filter

Algorithm 2: CANDIDATE SPACE REFINEMENT

```

1 while refinement is not finished do
2    $u \leftarrow$  a vertex from query graph;
3   foreach  $v \in C(u)$  do
4     if VertexSafety( $u, v$ ) violated then
5       Remove  $v$  from  $C(u)$ ;
6       continue;
7     foreach  $u' \in N_q(u)$  do
8       foreach  $v' \in C(u' | u, v)$  do
9         if EdgeSafety( $(u, u'), (v, v')$ ) violated then
10          Remove  $v'$  from  $C(u' | u, v)$ ;
11      if  $C(u' | u, v)$  is empty for any  $u' \in N_q(u)$  then
12        Remove  $v$  from  $C(u)$ ;

```

invalid candidates. For any necessary condition $h(u, v)$ that must be satisfied whenever $M(u) = v$ for some embedding M , in any step, we can remove candidate vertex v from $C(u)$ when $h(u, v)$ is false. VEQ considers *Neighbor Safety* for h , that can filter candidate vertices if it lacks sufficient candidate neighbors for some label. GQL [20] and VC [43] employ a stronger condition using bipartite matching.

As we explicitly maintain candidate edges, any necessary condition $g((u, u'), (v, v'))$ for edges which must be satisfied whenever $M(u) = v$ and $M(u') = v'$ for an embedding M can be used to remove invalid candidate edges.

In this paper, we propose novel safety conditions for stronger filtering. We propose the Triangle Safety and Four-Cycle Safety which consider cyclic substructures. We extend the ESIC condition suggested by VC [43] to Edge Bipartite Safety, increasing the filtering power while maintaining the same time complexity.

Triangle Safety. Filtering methods in [6, 18, 24, 43] did not utilize cyclic substructures, i.e., triangles and 4-cycles in the query graph. While finding all cycles is itself an NP-hard problem, using small cycles can greatly increase the filtering power without taking excessive time.

Definition 4.3. For an edge $(a, b) \in E_g$, we define $L_g^3(a, b)$ as the set of vertices $c \in V_g$ such that $(a, b), (b, c), (c, a) \in E_g$, i.e., they form a triangle with the given edge (a, b) .

We define the condition *Triangle Safety* as follows.

Definition 4.4. A data edge (v, v') for a query edge (u, u') is *Triangle Safe* when

- $|L_q^3(u, u')| \leq |L_G^3(v, v')|$, and
- for each $u^* \in L_q^3(u, u')$, there exists $v^* \in L_G^3(v, v')$ consisting of valid vertices and edges, i.e., for each u^* , there exists $v^* \in C(u^* | u, v) \cap C(u^* | u', v')$.

Four-Cycle Safety. A similar work can be done with 4-cycles in the query graph and the data graph.

Definition 4.5. For an edge $(a, b) \in E_g$, we define $L_g^4(a, b)$ as the set of edges $(c, d) \in E_g$ such that $(a, b), (b, c), (c, d), (d, a) \in E_g$, i.e., they form a 4-cycle with the given edge (a, b) .

We define the condition *Four-Cycle Safety* as follows.

Definition 4.6. A data edge (v, v') for a query edge (u, u') is *Four-Cycle Safe* when

- $|L_G^4(u, u')| \leq |L_G^4(v, v')|$, and
- for each $(u^*, \hat{u}) \in L_G^4(u, u')$, there exists $(v^*, \hat{v}) \in L_G^4(v, v')$ consisting of valid vertices and edges, i.e., for each (u^*, \hat{u}) , there exists $v^* \in C(u^* \mid u', v')$ and $\hat{v} \in C(\hat{u} \mid u, v)$ such that $(v^*, \hat{v}) \in E_{CS}(u^*, \hat{u})$.

To compute Triangle Safety and Four-Cycle Safety, we index all triangles and 4-cycles in advance. When the numbers of triangles and 4-cycles in the data graph are substantial, indexing them can result in a significant overhead. Hence, if the numbers of triangles and 4-cycles exceed some thresholds k_1 and k_2 , respectively, we disable respective safety condition to avoid excessive overhead of indexing. For the experiments, we use $k_1 = k_2 = 10^9$.

Edge Bipartite Safety. To develop a stronger safety condition, GQL [20] and VC [43] propose the use of a bipartite graph $B(u, v)$ for candidate neighbors of $v \in C(u)$, $u \in V_q$. $B(u, v)$ consists of left vertices $V_L = N_q(u)$, and right vertices $V_R = N_G(v)$. The edges of $B(u, v)$ connect u' and v' if and only if $v' \in C(u' \mid u, v)$, i.e., when a candidate edge exists between $v \in C(u)$ and $v' \in C(u')$.

It is then readily observed that for v to be a valid candidate for u , $B(u, v)$ should have a bipartite matching whose size is equal to the set of left vertices. We can determine whether $B(u, v)$ has such a bipartite matching in $O(d_q(u)^2 d_G(v))$ time utilizing Ford-Fulkerson [17] algorithm, as there are at most $d_q(u)d_G(v)$ edges and the size of the maximum matching is at most $d_q(u)$.

We further extend the use of the bipartite graph to the fine-grained filtering of candidate edges. Given the bipartite graph $B(u, v)$, the existence of a maximum bipartite matching of size $d_q(u)$ does not guarantee that for every edge there exists a maximum matching containing the edge.

Definition 4.7. An edge is called *maximally matchable* if it is included in some maximum matching.

When an edge $(u', v') \in E_{B(u, v)}$ is not maximally matchable, v' should be removed from $C(u' \mid u, v)$. For bipartite graphs, it is known that finding all maximally matchable edges can be done in linear time when a maximum matching is given [44]. Therefore, after we find the maximum bipartite matching on $B(u, v)$, we verify whether each edge is maximally matchable. This takes $O(d_q(u)d_G(v))$ time which is negligible compared to finding maximum bipartite matching.

4.3 Promising-First Candidate Filtering

Refinement Priority. In each refinement step, we select a query vertex to be refined (Algorithm 2, Line 2). Let $C_t(u)$ be the candidate set of u after the t -th refinement step finished, with $C_0(u)$ denoting the candidate set in the initial CS. Let u^1, u^2, \dots, u^t be the sequence of vertices that have already been refined, with the current refinement step being $t + 1$.

As u^{t+1} , we choose the query vertex u such that u has the lowest $penalty_t(u)$ (i.e., most promising), defined as

$$penalty_t(u) = \begin{cases} 1 & \text{if } u = u^t \\ penalty_{t-1}(u) \times \frac{|C_t(u^t)|}{|C_{t-1}(u^t)|} & \text{if } u \in N_q(u^t) \\ penalty_{t-1}(u) & \text{otherwise.} \end{cases}$$

We assign all vertices an initial penalty of $penalty_0(u) = \phi$ for some constant ϕ in $[0, 1]$. After refining u^t , $penalty_t(u^t)$ is increased to 1. For $u \in N_q(u^t)$, we reduce $penalty_t(u)$ by multiplying $|C_t(u^t)|/|C_{t-1}(u^t)|$ (the ratio by which $C(u^t)$ is reduced), as the size of $C(u)$ is also likely to decrease when the size of $C(u^t)$ decreases. If u is neither u^t itself nor a neighbor of u^t , its penalty remains unchanged. Note that the penalty always lies within the range $[0, 1]$.

Stopping Criteria. We empirically observe that the efficacy of refinements stagnates after some iterations. Therefore, we define two stopping criteria for the refinement, and stop whenever either one of the two conditions triggers.

- (1) When the minimum penalty among query vertices become higher than some predefined constant τ . This explicitly guides the algorithm to stop when the efficacy of further refinements are expected to be small and not worth the overhead.
- (2) When the sum of degrees of vertices chosen to be refined, including repeated selections, becomes higher than $R|E_q|$ for some predefined constant R . This criterion is necessary to bound the worst case complexity, and practically guides the algorithm to stop when the refinement takes too much time.

Theorem 4.8. With stated stopping criteria, various safety conditions lead to the following time complexity bounds for the filtering.

- Edge Bipartite Safety : $O(|E_q||E_G|\Delta_q)$ time
- Triangle Safety : $O(|E_q||E_G|\Delta_q\delta_G)$ time
- 4-Cycle Safety : $O(|E_q||E_G|\Delta_q^2\Delta_G\delta_G)$ time

Empirically, we choose $\tau = 0.9$ and $R = 5$ for the experiments. We observe that most of the runs were stopped by the first criterion (penalty).

5 CANDIDATE TREE SAMPLING

We present a sampling algorithm to select trees within the Candidate Space. Although there have been related works employing tree sampling for similar problems such as join cardinality estimation [51], these algorithms may face challenges in accuracy and efficiency since the number of trees is much larger than the actual number of embeddings. In contrast, our algorithm greatly benefits from the highly effective Candidate Space, which enables us to filter out numerous trees that cannot be extended to embeddings.

Our algorithm strategically chooses the spanning tree of the query graph to minimize the sample space. We prove that our method is optimal under an additional assumption, and it performs remarkably well in practice even when such an assumption is not met. Furthermore, we adaptively determine the sample size during the sampling process to provide a probabilistic bound. In most cases, our algorithm achieves the bound with a reasonably small sample size, ensuring both accuracy and efficiency in practice.

5.1 Candidate Trees

Definition 5.1 (Candidate Tree). Let T be a subgraph of q which is a tree rooted at $u_r \in V_T$. We define a *Candidate Tree* for T as a vertex mapping $s : V_T \rightarrow V_G$ such that

- for the root node u_r of T , $s(u_r) \in C(u_r)$, and
- for a node x and its parent node $p \in V_T$, $s(x) \in C(x \mid p, s(p))$.

Algorithm 3: CANDIDATETREESAMPLING(CS, q)

Input: The compact CS, a query graph q **Output:** Estimated number of embeddings and the number of successful trials in tree sampling

```
1  $T_q \leftarrow \text{GetQueryTree}(\text{CS}, q);$ 
2  $D \leftarrow \text{CountCandidateTrees}(\text{CS}, T_q);$ 
3 #successes  $\leftarrow 0;$ 
4 #trials  $\leftarrow 0;$ 
5 while sampling termination condition is not reached do
6    $s \leftarrow \text{GetSampleTree}(\text{CS}, T_q, D);$ 
7   #trials  $\leftarrow$  #trials + 1;
8   if CheckEmbedding(CS,  $s$ ,  $q$ ) then
9     #successes  $\leftarrow$  #successes + 1;
10  $u_r \leftarrow$  root of  $T_q;$ 
11 estimate  $\leftarrow \frac{\text{\#successes}}{\text{\#trials}} \sum_{v \in C(u_r)} D(u_r, v);$ 
12 return (estimate, #successes);
```

We note that a candidate tree can be seen as a homomorphism of T in the compact CS.

Example 5.2. In Figure 3, for given q and compact CS, we first build a spanning tree T_q of q by removing (u_2, u_4) and (u_3, u_4) edges. The mapping $\{(u_1, v_1), (u_2, v_3), (u_3, v_5), (u_4, v_7), (u_5, v_3)\}$ is a candidate tree for T_q . Note that the candidate tree is not necessarily an injective mapping.

Let T_q be any spanning tree of q . When needed, we consider T_q as a directed graph with edge directions assigned from root to leaf. We say that a candidate tree t is a *counterpart* of an embedding if the vertex mapping of t is the same as the vertex mapping M of the embedding. Since the CS is complete, the set of candidate trees for T_q contains the counterparts of all the embeddings of q . Therefore, we choose a spanning tree T_q and use the set of candidate trees for T_q in the compact CS as the sample space Ω . We develop an algorithm to perform exact counting of candidate trees using dynamic programming, and sample a candidate tree from Ω uniformly at random.

We note that the sample space can be further reduced if the definition of candidate trees is extended to only allow injective mappings. However, this extension is infeasible, as exact counting of isomorphic embeddings is NP-hard for tree queries [30].

5.2 Sampling Candidate Trees

Choosing the Spanning Tree. It is best to choose the spanning tree T_q that the number of candidate trees is minimized to reduce the size of the sample space. However, exact minimization by generating all possible spanning trees and counting the number of candidate trees is not viable as the number of spanning trees for q can be superexponentially large.

Thus, instead of exact minimization, we introduce a *density heuristic* to find a spanning tree with a small number of candidate trees. Since the number of candidate trees generally increases as the number of candidate edges increases, using query edges having dense candidate edges is undesirable. To avoid such cases, we

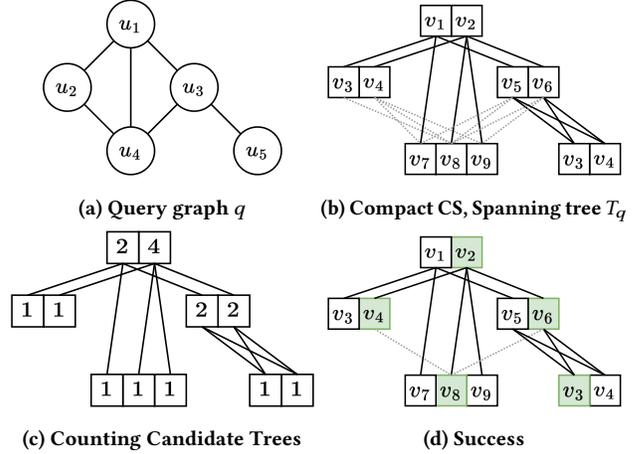


Figure 3: New Running Example - Candidate Tree Sampling

assign the weight of the edge $(u, u') \in E_q$ as

$$\text{density}(u, u') = \frac{|E_{\text{CS}}(u, u')|}{|C(u)||C(u')|} \quad (1)$$

and find a spanning tree that minimizes the sum of the logarithm of densities by Prim's Algorithm [36], which is equivalent to minimizing the product of densities of the selected edges.

This strategy of minimizing the product of density is optimal with the following assumption: for all $(v, v') \in C(u) \times C(u')$, the event that edge (v, v') is a candidate for (u, u') has the probability $\text{density}(u, u')$, and it is independent from the events that other vertex pairs are candidate edges.

Theorem 5.3. Under the above assumption, the expected number of candidate trees is minimized by the density heuristic.

Counting Candidate Trees. We develop a dynamic programming algorithm to obtain the exact count of candidate trees for a given T_q . Let T_u be a subtree of T_q rooted at $u \in V_q$ and $D(u, v)$ be the number of candidate trees for T_u with u mapped to v . The total number of candidate trees for T_u can be counted as $\sum_{v \in C(u)} D(u, v)$.

For a leaf node $u \in V_q$ and $v \in C(u)$, it is clear that $D(u, v) = 1$. For a non-leaf node $u \in V_q$ and $v \in C(u)$, a candidate tree for T_u consists of candidate trees for T_{u_c} for each child u_c of u , and the candidate for each u_c is chosen from $C(u_c | u, v)$. Hence, we have

$$D(u, v) = \prod_{u_c \in \text{children of } u} \sum_{v_c \in C(u_c | u, v)} D(u_c, v_c) \quad (2)$$

which can be computed in $O(|E_q||E_G|)$ time, employing a bottom-up dynamic programming approach similar to JSub [51] and DAF [18].

Uniform Sampling of Candidate Trees. Based on the candidate tree counts computed as above, we develop a sampling algorithm that returns each candidate tree uniformly at random. Let s denote a sample candidate tree for T_q . Recall that the candidate tree is defined as a vertex mapping. For the root vertex u_r , we sample v from $C(u_r)$ with weights proportional to $D(u_r, v)$. For $u \neq u_r$, let u^p be the parent of u in T_q . Given $v^p = s(u^p)$, we iteratively sample $s(u)$ for u with BFS traversal of T_q . We choose $s(u)$ from candidate neighbors $C(u | u^p, v^p)$ with weights proportional to $D(u, v)$. The detailed implementation is presented in Algorithm 4.

Algorithm 4: GETSAMPLETREE(CS, T_q , D)

Input: The compact CS, a spanning tree T_q of q , the computed weights D

Output: A sampled candidate tree, uniformly at random

- 1 $u_r \leftarrow$ root node of T_q ;
 - 2 Sample[u_r] \leftarrow draw v from $C(u_r)$ at random with probability proportional to $D(u_r, v)$;
 - 3 **foreach** u in BFS traversal of T_q **do**
 - 4 let u^p be the parent of u , and v^p be Sample[u^p];
 - 5 Sample[u] \leftarrow draw v from $C(u \mid u^p, v^p)$ at random with probability proportional to $D(u, v)$;
 - 6 **return** Sample
-

Example 5.4. Figure 3 demonstrates the process of counting and sampling candidate trees. For the query graph q , the compact CS is built as in Figure 3b. The edges (u_2, u_4) and (u_3, u_4) are not included to attain T_q as they have larger densities of $2/3$. Figure 3c shows the value of resulting $D(u, v)$ after counting candidate trees.

For the sampling, we start with choosing a root with weight $2 : 4$ for v_1 and v_2 , respectively. For example in Figure 3d, when v_2 is chosen, v_4 and v_6 are the only candidate neighbors of (u_1, v_2) for u_2 and u_3 respectively, thus they are sampled with probability 1. For u_4 , we take a random sample from $C(u_4 \mid u_1, v_2)$. Analogously, $s(u_5)$ is chosen among $C(u_5 \mid u_3, v_6)$. As the resulting mapping in Figure 3d is injective, we check whether non-tree edges (u_2, u_4) and (u_3, u_4) are valid. Since (v_4, v_8) and (v_6, v_8) are both in the candidate edges, Figure 3d is a success.

Theorem 5.5 (Uniform Sampling). Algorithm 4 samples *candidate tree* for T_q uniformly at random.

Sample Size. If we perform $\#t$ trials and obtain $\#s$ successes, the sample success ratio $\hat{\rho} = \#s/\#t$ is an estimate of the true proportion ρ . Let α denote the acceptable failure probability and c denote the tolerable relative error. To achieve the probability guarantee $\mathbb{P}[c^{-1}\rho \leq \hat{\rho} \leq c\rho] \geq 1 - \alpha$, we utilize the Clopper-Pearson interval [10, 14], which is the interval $(L_\alpha(\#t, \#s), U_\alpha(\#t, \#s))$ such that

$$\mathbb{P}[(L_\alpha(\#t, \#s) \leq \rho \leq U_\alpha(\#t, \#s))] \geq 1 - \alpha, \quad (3)$$

where $L_\alpha(\#t, \#s)$ and $U_\alpha(\#t, \#s)$ are the values that can be computed from α , $\#t$, and $\#s$ (e.g., using the Boost library [1]). We adaptively determine the sample size by taking samples until $c^{-1}\hat{\rho} \leq L_\alpha(\#t, \#s)$ and $U_\alpha(\#t, \#s) \leq c\hat{\rho}$ are met, leading us to:

$$\begin{aligned} \mathbb{P}[c^{-1}\rho \leq \hat{\rho} \leq c\rho] &= \mathbb{P}[c^{-1}\hat{\rho} \leq \rho \leq c\hat{\rho}] \\ &\geq \mathbb{P}[L_\alpha(\#t, \#s) \leq \rho \leq U_\alpha(\#t, \#s)] \geq 1 - \alpha. \end{aligned} \quad (4)$$

In our experiments where we had 1,000 – 1,000,000 trials, we found out that, irrespective of the number of trials, 88 successes were sufficient to satisfy the condition stated with the Clopper-Pearson interval (Equation 4) for $\alpha = 0.05$ and $c = 1.25$.

For hard cases in which achieving the condition may require an unreasonable amount of computation, we terminate the tree sampling immediately and use the graph sampling described in Section 6. For the experiments, we stop if there are no more than 10 successes for 50,000 trials. Since 88 successes are required to meet the stated condition and we only continue the tree sampling

if there are at least 11 successes in the first 50,000 trials, we expect that the number of trials won't exceed 400,000 in the tree sampling.

Theorem 5.6. Algorithm 3 is an unbiased, consistent estimator for subgraph cardinality.

6 STRATIFIED GRAPH SAMPLING

While the presented tree sampling algorithm gives accurate and efficient estimate for most cases, there are some hard instances where sampling space remains large compared to the number of embeddings even after filtering. This is due to the fact that the number of candidate trees can be asymptotically larger than $AGM(q)$, which is a tight upper bound for the number of embeddings [2].

To address such an issue, we develop a *stratified graph sampling* algorithm with the following components. (1) We consider non-tree edges together during the sampling phase, ensuring that the sample space size is bounded by $O(AGM(q))$. (2) We propose stratified sampling to select samples from diverse regions within the sample space. By combining these two components, our algorithm achieves high accuracy and efficiency, particularly on difficult instances.

6.1 Extendable Candidates

We define a vertex mapping M as a *partial embedding* when it is an embedding of an induced subgraph q' of q to G . Given a partial embedding M , a vertex $u \in V_q \setminus V_{q'}$ is an *extendable vertex* if there exists $u' \in N_q(u)$ such that $u' \in V_{q'}$. We define the set of *extendable candidates* of u regarding a partial embedding M as follows:

$$C_M(u) = \bigcap_{\{(u', v') \in M \mid u' \in N_q(u)\}} C(u \mid u', v') - \text{image of } M.$$

where image of M is the vertices of the data graph that are mapped by M , i.e., $\{v \in V_G \mid \exists u \in V_{q'} \text{ such that } M(u) = v\}$. When $M = \emptyset$, $C_M(u)$ is set to be $C(u)$. $C_M(u)$ is computed by the multi-way set intersection over candidate neighbor sets of each neighbor u' of u . To maintain injectivity, candidate vertices that are already mapped by M are removed from $C_M(u)$. For $v \in C_M(u)$, $M \cup \{(u, v)\}$ forms a partial embedding since $M \cup \{(u, v)\}$ is an embedding of an induced subgraph of q with vertex set $V_{q'} \cup \{u\}$ to G .

For example, in Figure 3b, at the time when $M = \{(u_1, v_2), (u_2, v_4)\}$, u_4 is an extendable vertex with extendable candidates $C_M(u_4) = \{v_8, v_9\}$ as they are the vertices in $C(u_4)$ which are candidate neighbors for both (u_1, v_2) and (u_2, v_4) .

6.2 Stratified Graph Sampling

We say that a partial embedding *extends* M if it contains M as a subset. Let w_M denote the number of embeddings of q to G extending the partial embedding M . It is clear that $w_M = 1$ when $|M| = |V_q|$, and $w_M = 0$ when $C_M(u) = \emptyset$. For other cases, we employ the method of *stratified sampling* to obtain an unbiased consistent estimator for w_M .

Stratified Sampling. *Stratification* is a strategy of dividing the population into multiple subpopulations that are mutually exclusive and jointly exhaustive, namely *strata*, prior to sampling in order to ensure diversity of resultant samples [29]. In our instance, to estimate w_M , the population is the set of partial embeddings extending M . For an extendable vertex u and candidates $C_M(u)$ for M , the population can be then stratified into subpopulations

Algorithm 5: Stratified Graph Sampling

Input: The compact CS, a query graph q , the number of successful tree samples $\#successes$

Output: Estimated number of embeddings

```

1 Function CandidateGraphSampling(CS,  $q$ ,  $\#successes$ ):
2    $ub_0 \leftarrow \text{GetSampleSize}(q, \#successes)$ ;
3   (estimate,  $\#trials$ )  $\leftarrow \text{EstimateW}(\text{CS}, q, \emptyset, ub_0)$ ;
4   return estimate;

```

Input: The compact CS, a query graph q , a current partial embedding M , an upper bound for sample size ub_M

Output: Estimated number of embeddings extending M , number of samples used

```

5 Function EstimateW(CS,  $q$ ,  $M$ ,  $ub_M$ ):
6   if  $|M| = |V_q|$  then return (1, 1);
7    $u \leftarrow \text{ChooseExtendableVertex}(q, M)$ ;
8    $C_M(u) \leftarrow \text{GetExtendableCandidates}(\text{CS}, q, u)$ ;
9   if  $C_M(u)$  is empty then return (0, 1);
10   $sz \leftarrow \min(C_M(u)/k, ub_M)$ ;
11   $S \leftarrow \text{GetRandomSubset}(C_M(u), sz)$ ;
12   $(\hat{w}_M, \#trials) \leftarrow (0, 0)$ ;
13  for  $i = 1 \dots |S|$  do
14     $M' \leftarrow M \cup \{(u, S_i)\}$ ;
15     $(\hat{w}_{M'}, n) \leftarrow \text{EstimateW}(\text{CS}, q, M', \frac{ub_M - \#trials}{|S| - i + 1})$ ;
16     $(\hat{w}_M, \#trials) \leftarrow (\hat{w}_M + \hat{w}_{M'}, \#trials + n)$ ;
17   $\hat{w}_M \leftarrow \frac{|C_M(u)|}{|S|} \times \hat{w}_M$ ;
18  return  $(\hat{w}_M, \#trials)$ ;

```

(called *groups*) of partial embeddings extending $M \cup \{(u, v)\}$ for each $v \in C_M(u)$. Let \hat{w}_M be an estimate for w_M . By definition,

$$w_M = \sum_{v \in C_M(u)} w_{M \cup \{(u, v)\}}. \quad (5)$$

Therefore, by taking a random subset S of candidates from $C_M(u)$, and estimating the values of $w_{M \cup \{(u, v)\}}$ for $v \in S$, we obtain an unbiased estimator for w_M . Formally, consider a random subset S of $C_M(u)$, where the probability of each $v \in C_M(u)$ being included in S is uniform as $|S|/|C_M(u)|$. We define \hat{w}_M as

$$\hat{w}_M = \frac{|C_M(u)|}{|S|} \sum_{v \in S} \hat{w}_{M \cup \{(u, v)\}}. \quad (6)$$

If $|M| = |V_q|$, \hat{w}_M is set to 1, and if $C_M(u) = \emptyset$, \hat{w}_M is set to 0. Each $\hat{w}_{M \cup \{(u, v)\}}$ are computed recursively.

The pseudocode of the algorithm is presented in Algorithm 5.

Sampling Order. In each step, we choose an extendable vertex $u \in V_q$, and compute extendable candidates $C_M(u)$ given partial embedding M . To choose the extendable vertex with the number of extendable candidates as small as possible, we choose the vertex u which is not in M and has the highest number of neighbors in M in each step. At first, we choose the vertex with smallest $|C(u)|$. We note that for given q and the compact CS, the order of extendable vertices is deterministic.

Adaptive Allocation. Examining a fixed proportion of extendable candidates in each recursive step by taking $|S| = |C_M(u)|/k$ for some constant k may result in an exponentially large number of

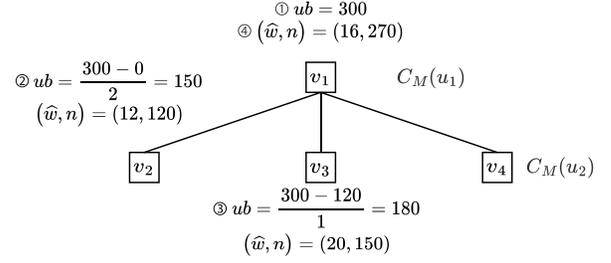


Figure 4: Example of Stratified Sampling

samples. To circumvent this issue, we limit the maximum number of samples used for estimating w_M as ub_M for each M . We choose the subset S to be a random subset with a size of $\min(|C_M(u)|/k, ub_M)$ (Algorithm 5, Line 10). The value of ub_0 is determined based on the difficulty of the query, with further details to be discussed later. We then set the upper bound $ub_{M \cup \{(u, v)\}}$ for each $v \in S$ in such a way that the sum of upper bounds does not exceed ub_M .

For the first $v \in S$, we assign $ub_{M \cup \{(u, v)\}} = \frac{ub_M}{|S|}$. The recursive call on $M \cup \{(u, v)\}$ returns an estimate for $w_{M \cup \{(u, v)\}}$ and the actual number of samples taken throughout the recursion. The recursion on M may terminate before the number of samples reaches ub_M (e.g., if $|M| = |V_q|$ or $C_M(u) = \emptyset$, only one sample is encountered regardless of how large ub_M is). In such cases, we adaptively increase the sample size for subsequent calls equally using this information to enhance the overall accuracy. The upper bound for the i -th $v \in S$ is determined as the remaining upper bound (i.e., ub_M minus the number of samples seen from the first to the $(i-1)$ -th) divided by the number of remaining calls ($|S| - i + 1$), as in Line 15. As this can only increase the upper bound from $ub_M/|S|$, we ensure that each $M \cup \{(u, v)\}$ receives at least $ub_M/|S|$ for the upper bound, thereby obtaining diverse samples even with limited sample sizes.

Example 6.1. An example with $ub_M = 300$ is depicted in Figure 4. For $C_M(u_2) = \{v_2, v_3, v_4\}$, suppose $S = \{v_2, v_3\}$ is the randomly chosen subset. For the group with v_2 , the upper bound $ub_{M \cup \{(u_2, v_2)\}}$ is set to 150 as ub_M is expected to split into two groups. Suppose that for the first group, the returned estimate is 12 with 120 samples used. The upper bound for the second group is then determined as 180, since only 120 samples were used. Assume that the group returned 20 with 150 samples used. Aggregating the results, the initial call on v_1 returns 16 (average of estimates returned by two subcalls) and reports that 270 samples are actually used.

Theorem 6.2. Algorithm 5 is an unbiased consistent estimator for subgraph cardinality, and the worst-case time complexity is bounded by $O(AGM(q))$.

Sample Size. It is favorable to use more samples on difficult instances for accuracy, and use a small sample size on easier instances for efficient estimation. However, measuring the difficulty is itself challenging as it depends on the ratio between the number of embeddings and the size of the sample space.

The query size is commonly used in prior works on graph homomorphism counting, as a larger query requires more samples in general [26, 35]. Empirically, while the difficulty of an instance increases with the query size, the size of the data graph does not seem to be strongly relevant. Since we perform the graph sampling

only when the tree sampling has not found enough successes after some number of samples, the number of successes from the tree sampling can act as a proxy measure of the difficulty of an instance. Hence we take larger ub_θ when the number of successes from the tree sampling is small. However, this association is not linear empirically, and therefore we take the square root on the number of successes, setting ub_θ as

$$ub_\theta = \frac{|V_q| \times K}{\sqrt{\# \text{ successes} + 1}} \quad (7)$$

for some constant K . Using $K = 100,000$, we achieve a reasonable accuracy over datasets of wildly different statistics.

7 PERFORMANCE EVALUATION

7.1 Experimental Setup

We conduct experiments to evaluate the performance of our algorithm, FaST_{est}, by answering the following questions.

- **Accuracy.** Compared to state-of-the-art competitors, how accurate is FaST_{est} on various real-world datasets and difficult query instances with diverse characteristics? (Section 7.2)
- **Efficiency.** Compared to state-of-the-art competitors, how efficient is FaST_{est} in terms of execution time and memory usage? (Section 7.3)
- **Evaluation of Techniques.** How do our safety conditions and tree/graph sampling algorithms contribute to accurate and efficient estimation? (Section 7.4)

We mainly compare our algorithm with Alley [26], since it has significantly outperformed the synopses-based algorithms and sampling-based algorithms such as Characteristic Sets [34], SumRDF [41], Correlated Sampling [47], JSub [51], and WanderJoin [27]. Since Alley was originally developed for estimating the number of homomorphic embeddings, we modified it for isomorphic embeddings. We list the variants of Alley we compared to as follows.

- **Alley:** We modified Alley so that for each random walk sampled, whether the walk is injective is additionally checked.
- **Alley+:** We modified Alley to not choose an already visited data vertex twice during each random walk, ensuring that only injective walks are sampled.
- **Alley+TPI:** We modified the Tangled Pattern Index (TPI) structure proposed in [26] for isomorphism and built it on top of Alley+. Following [26], we set the maximum pattern size to 4 for vertex labeled datasets, and 5 if both vertex labels and edge labels are present.

To compare our algorithm with GNN-based methods, we consider two recent state-of-the-art works, LSS [50] and NeurSC [48] as two additional baselines.

Environment. FaST_{est} is implemented in C++. The source code of Alley and LSS were publicly available, and NeurSC was obtained from the authors. Hyperparameters were set to their suggested default values. Experiments were conducted on a machine with two Intel Xeon Silver 4114 2.20GHz CPUs, an NVIDIA RTX 3090 Ti GPU, and 256GB memory.

Datasets. We conduct experiments on real-world large-scale graph datasets used in previous works [18, 24, 26, 42, 48]. For fair comparison, query graphs were obtained from a comparative study on

Table 2: Statistics of Data Graphs. Σ_V and Σ_E denote the sets of vertex labels and edge labels, respectively.

Dataset	$ V $	$ E $	$ \Sigma_V $	$ \Sigma_E $	#Query	$ V_q $
Yeast (Ye)	3.1K	12.5K	71	-	1,707	4 to 32
WordNet (Wo)	76.9K	120.4K	5	-	1,164	4 to 20
DBLP (Db)	317.1K	1.0M	15	-	911	4 to 20
Youtube (Yo)	1.1M	3.0M	25	-	922	4 to 32
Patents (Pa)	3.8M	16.5M	20	-	1,485	4 to 32
YAGO (Ya)	12.8M	15.8M	185K	91	960	4 to 20

subgraph matching algorithms [42] if possible. Large query graphs for the DBLP dataset ($|V_q| = 12, 20$) and query graphs for the YAGO dataset were generated via random walk. Since YAGO has edge labels, we extended our algorithm to handle edge-labeled graphs. We used an exact subgraph matching algorithm DAF [18] for obtaining the ground truth results. Since the computational cost for exact counting is extremely high, we only used the query graphs for which the exact count could be computed within 2 hours. The data graphs are listed in Table 2.

Performance Measure. Accuracy is measured by q -error [32], defined as $\max\left(\frac{\max(1, w)}{\max(1, \hat{w})}, \frac{\max(1, \hat{w})}{\max(1, w)}\right)$, where \hat{w} is the estimate and w is the ground truth number of embeddings, as done in previous works [26, 35, 48, 50]. Efficiency is measured by average elapsed time per query. For LSS and NeurSC, we evaluated with five-fold cross-validation.

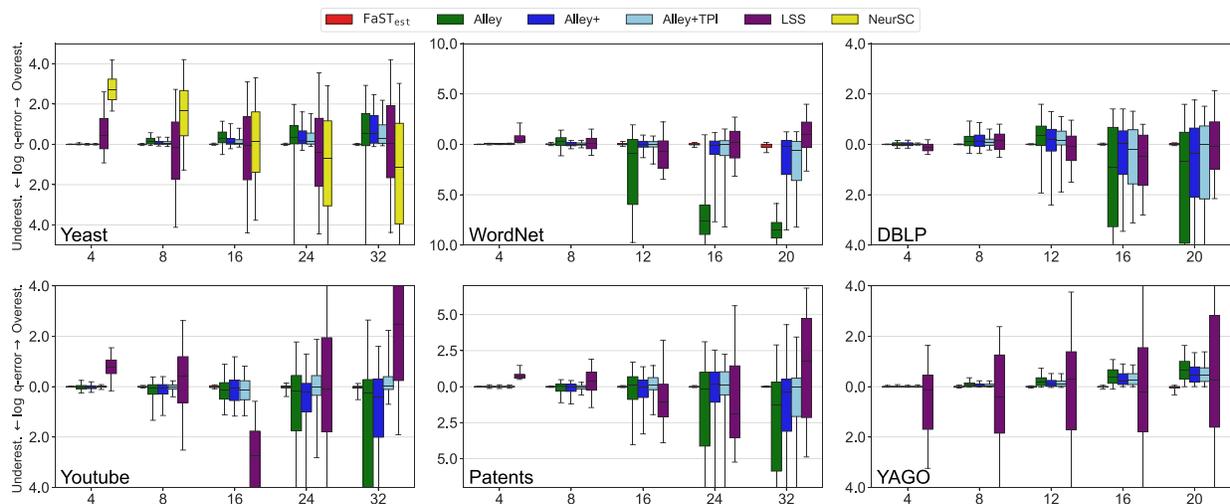
For the tree sampling, we set the acceptable failure probability to $\alpha = 0.05$ and the tolerable relative error to $c = 1.25$. Using these parameters, the sample size for the tree sampling is adaptively determined by the Clopper-Pearson interval (Equation 4).

7.2 Accuracy

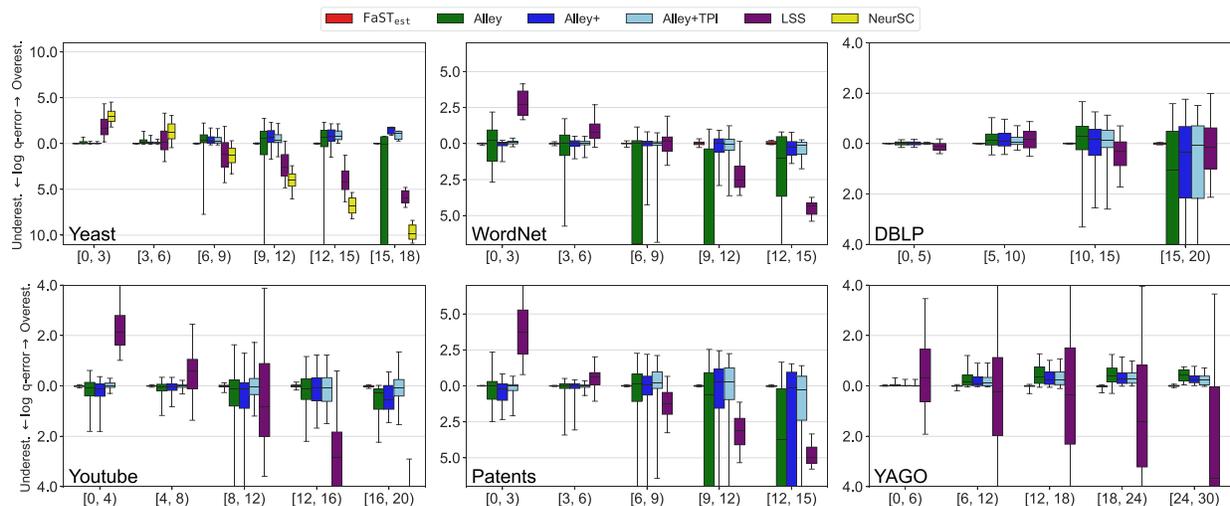
Figure 5 presents the accuracy of FaST_{est} compared to state-of-the-art competitors on real-world graph datasets, varying the size of the query graph (Figure 5a) and the number of embeddings (Figure 5b). In Figure 5, *difficult* instances for subgraph cardinality estimation are large queries or queries with many embeddings. To differentiate overestimation and underestimation, if $\hat{w} \geq w$, $\log_{10} \frac{\max(1, \hat{w})}{\max(1, w)}$

is shown above $y = \log_{10} 1 = 0$; if $\hat{w} < w$, $\log_{10} \frac{\max(1, w)}{\max(1, \hat{w})}$ is shown below $y = 0$, as in [26, 35, 48, 50]. To ensure a fair comparison, we present the results with $K = 100,000$ (Equation 7) for our graph sampling, and choose the sample size of Alley so that its overall execution time is not less than ours. Except for Yeast, we were unable to train NeurSC due to time and memory limits, and therefore we omit the results. Across data graphs with diverse statistical properties and query sizes ranging from small (4) to large (32), FaST_{est} significantly outperforms both sampling methods and GNN methods in terms of accuracy. Note that a log q -error of 1 represents tenfold over/underestimation.

Figure 5a shows that the q -error of GNNs increases as the size of the query graph grows. For example, on Patents-32 (queries with size 32 on the Patents dataset), FaST_{est} shows an average q -error of 1.04 while LSS shows 4328.3 (log q -error of 0.02 and 3.64, respectively), representing an improvement of more than three orders



(a) \log_{10} q-error with the x-axis representing the number of query vertices.



(b) \log_{10} q-error with the x-axis representing \log_{10} (#embeddings) of the query graph.

Figure 5: q-error evaluation on real world datasets (closer to 0 on the y-axis indicates better performance). The boxplot presents values of $\log q$ -error of a set of queries. The lower and upper whisker denotes the 5% and 95% quantile, respectively. The box itself represents the interquartile range, covering values between the 25% and 75% quantiles. The median of each set of queries is represented by the black line within the box.

of magnitude. Compared to NeurSC, FaST_{est} shows an average q-error of 1.07 on Yeast-32, where NeurSC shows 811.5 (log q-error of 0.03 and 2.91, respectively).

Additionally, GNNs tend to overestimate when the true cardinality is small, while they underestimate when the true cardinality is large, as observed in Figure 5b. While GNN-based methods can make fast inference for each query, the accuracy on difficult instances is substantially lower than that of sampling-based algorithms. Therefore, for further evaluation, we focus on comparison with the state-of-the-art sampling algorithm, Alley.

Sampling algorithms are often prone to sampling failure, where no successful sample is found and the algorithm consequently reports zero as the estimate. Alley also exhibits sampling failure on

difficult instances, such as large queries on Patents, Youtube, and WordNet datasets. However, FaST_{est} significantly reduces the sample space by strong filtering, thereby achieving high accuracy on such instances. Especially, on WordNet-20, FaST_{est} shows an average q-error of 1.81 (log q-error of 0.26), where Alley+TPI exhibits an average q-error of 418.5 (log q-error of 2.62), which represents an improvement of more than two orders of magnitude.

The accuracy of Alley tends to degrade as the true cardinality increases. Yet, as seen in Figure 5b, FaST_{est} demonstrates impressive accuracy even in cases where the true cardinality is exceptionally large. For the Yeast dataset, on queries with more than 10^{15} embeddings (represented by [15, 18]), FaST_{est} exhibits an average q-error of 1.10, which corresponds to a relative error of about 10%.

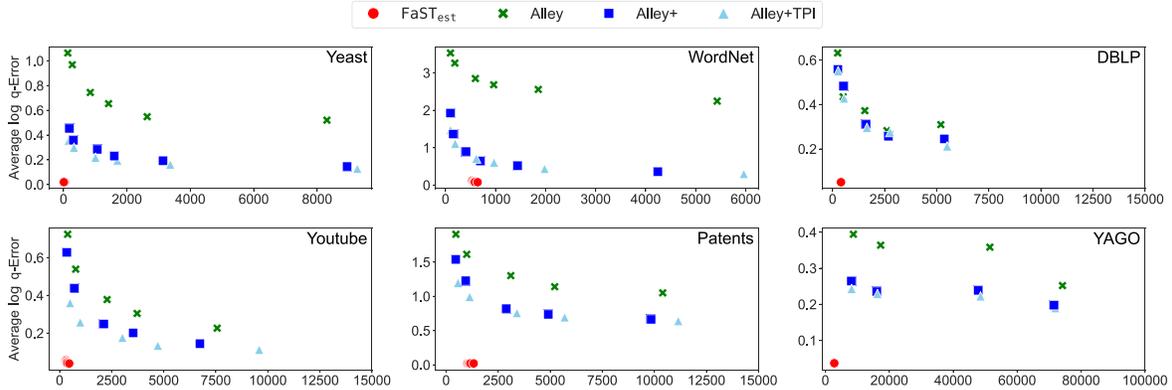


Figure 6: Accuracy-Time Tradeoff with different sample sizes. (Lower = Better accuracy, Left = Faster query processing)
The x-axis represents the average elapsed time per query (ms).

7.3 Efficiency

Query Processing Time. Table 3 presents both indexing time and query processing time of FaST_{est}. Specifically, the query processing time is broken down into filtering time, tree counting time, tree sampling time, and graph sampling time. The values for Filtering, Tree Counting, Tree Sampling, and Graph Sampling (A) represent the average time spent per query (i.e., total graph sampling time divided by the number of queries). Graph Sampling (H) shows the average graph sampling time per hard query (i.e., total graph sampling time divided by the number of hard queries), where a *hard* query means a query that proceeds to the graph sampling because no more than 10 successes were observed in 50,000 trials of the tree sampling. For example, in Yeast, 25 out of 1,707 queries are hard queries. While the graph sampling usually takes much longer compared to the tree sampling, its overall impact on the query processing time remains limited because the majority of queries are resolved through the tree sampling. In general, filtering takes the most time, especially for large datasets.

Accuracy-Time Tradeoff. Sampling algorithms have a natural tradeoff between accuracy and query processing time, as the estimation can be more accurate by taking a larger sample size. In Figure 6, we demonstrate the performance of our algorithm in this tradeoff by conducting an experiment varying sample sizes.

For the graph sampling, we vary the sample size by changing value of K in Equation 7 from 5×10^3 to 3×10^5 . We determine the sample sizes for Alley as $K \times |V_q|$, using the same range of K . Each mark in Figure 6 represents the elapsed time and the log q-error averaged over all queries at different sample sizes. As the sample

Table 3: Breakdown of the running time of FaST_{est} (ms). Hard queries denote the queries that required graph sampling.

Dataset	Yeast	WordNet	DBLP	Youtube	Patents	YAGO
Indexing	658	509	49093	3419	16066	13512
Query Processing	28.06	555.53	403.94	382.08	1122.13	2870.08
Filtering	4.70	443.96	400.67	293.83	980.89	2844.65
Tree Counting	0.21	23.75	2.04	1.45	2.72	7.20
Tree Sampling	4.47	42.48	1.23	21.27	11.97	6.09
Graph Sampling (A)	18.68	45.54	0	65.53	89.55	12.14
Graph Sampling (H)	1275.47	84.01	0	663.94	929.94	3884.80
# Hard Queries	25/1707	631/1164	0/911	91/922	143/1485	3/960

size increases, the sampling algorithms tend to require more time (moving from left to right on the x-axis), while the average log q-error decreases, indicating improved accuracy (moving from top to bottom on the y-axis). For large datasets, the results of Alley variants taking an excessive amount of time were omitted (more than 15,000 ms for DBLP, Youtube, Patents, and 100,000 ms for YAGO per query on average).

Figure 6 highlights the accuracy-to-time tradeoff of FaST_{est} when compared to Alley. As seen in Table 3, most queries on Yeast, DBLP, and YAGO are handled by the tree sampling, whose sample size is determined by Equation 4, independent of K . This explains the result appearing almost as a single dot for Yeast, DBLP, and YAGO. In these datasets, the tree sampling demonstrates outstanding accuracy and efficiency. For WordNet, Youtube, and Patents, we observe that increasing the sample size improves the accuracy for hard queries, resulting in a reduction of the average log q-error.

Although FaST_{est} incurs some overhead due to filtering for large datasets, it remains significantly more efficient in terms of the required execution time to achieve a specific accuracy level. In WordNet, Alley+TPI, the top competitor, takes about 6,000 milliseconds per query to attain an average log q-error of 0.3 (q-error of about 2). On the other hand, FaST_{est} reaches an average log q-error of 0.1 (q-error of about 1.26) with just about 600 milliseconds per query.

The efficiency of FaST_{est} is achieved by the synergy of our three key components. Firstly, our strong filtering substantially reduces the sample space before sampling, which reduces the number of required samples. Filtering also reduces the cost of obtaining a sample. Secondly, the tree sampling generates accurate results without requiring costly intersection, addressing most cases with high efficiency. Lastly, the graph sampling tackles the remaining hard cases.

Indexing Time. Table 4 shows the indexing time for triangles and 4-cycles for our datasets compared to the TPI (Tangled Pattern Index) from Alley [26]. In our experiments, all six datasets utilize Triangle Safety, and three datasets (Yeast, WordNet, and DBLP) use

Table 4: Indexing time of Triangles, 4-Cycles and TPI (sec)

	Ye	Wo	Db	Yo	Pa	Ya
Triangles	0.006	0.04	1.3	3.4	16.1	13.5
Four-Cycles	0.65	0.47	47.8	-	-	-
TPI	278.0	41.5	5467.1	6069.8	24837.6	4732.7

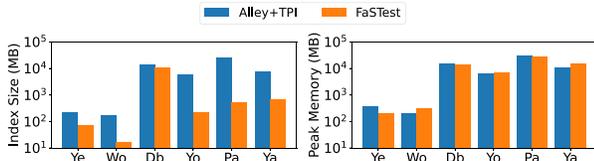


Figure 7: Index size and peak memory consumption.

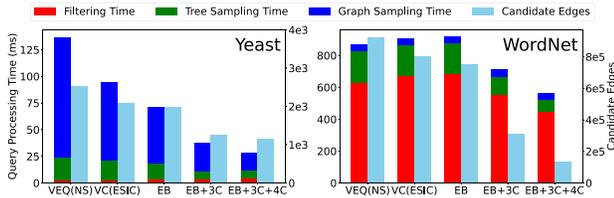


Figure 8: Query processing time and size of the Candidate Space after filtering, varying safety conditions.

Four-Cycle Safety. Compared to TPI which takes exponential time, our indexing can be done much more efficiently in polynomial time. We note that since this indexing is required only once per data graph, it can be done in advance to process many queries.

Memory Usage. Figure 7 shows the size of the index and peak memory consumption (measured as the highest memory usage observed across all queries, including both the index and the Candidate Space) of FaST_{est} compared to Alley. While the Candidate Space is built for each query, the index is built only once per data graph. FaST_{est} builds a smaller index compared to Alley’s TPI structure, and its peak memory usage remains comparable to that of Alley.

7.4 Evaluation of Techniques

Safety Conditions. Figure 8 shows the breakdown of the query processing time and the number of candidate edges after filtering for Yeast and WordNet, varying only the safety conditions used. The red, green, and blue portion of the left bar indicates the time spent on filtering, tree sampling, and graph sampling, respectively. The light blue bars show the number of candidate edges after filtering.

Together with Edge Bipartite Safety (EB), Triangle Safety (3C) and Four-Cycle Safety (4C), we implemented the safety conditions from two state-of-the-art subgraph matching algorithms: Neighbor Safety (NS) proposed by VEQ [24] and Exact Star Isomorphism Constraint (ESIC) by VC [43] as the baselines.

On Yeast, a stronger safety condition significantly reduces the number of candidate edges, at the expense of slight increase in the filtering time. This reduction in candidate edges accelerates both tree sampling and graph sampling, thereby improving overall query processing time. For WordNet, using EB increases the filtering time compared to ESIC. However, combining EB and substructure conditions (3C, 4C) reduces the filtering time as well, since the candidate edges removed by strong safety conditions are not considered in subsequent filtering steps, reducing the overall cost of filtering.

Specifically, our filtering method combining all three safety conditions reduces the number of candidate edges in WordNet by 80% compared to NS and by 75% compared to ESIC. On both datasets, each individual safety condition we propose has a clear effect in reducing query processing time and candidate edges.

Comparison of Tree and Graph Sampling. As shown in Table 3, the average tree sampling time per query is significantly

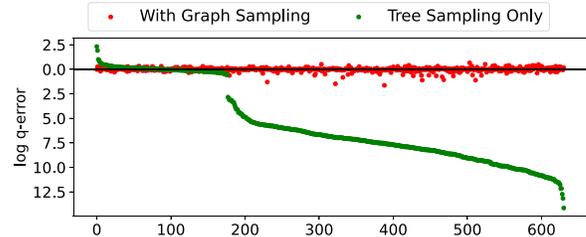


Figure 9: Comparative analysis of log q-error on WordNet with (red) and without (green) the graph sampling. The x-axis indicates queries sorted by q-error obtained from the tree sampling. Values closer to 0 is better, with dots above (resp. below) black line representing overestimation (resp. underestimation).

lower than the average graph sampling time per hard query. However, graph sampling is necessary for hard queries as tree sampling alone may not yield accurate estimates within a reasonable number of samples. In the extreme case, there is a query which has 1.5×10^6 embeddings in WordNet and for which 1.2×10^{26} candidate trees exist, necessitating about 7×10^{21} samples to get 88 successes, which is required for reaching the desired accuracy. In the experiment, FaST_{est} proceeded to the graph sampling after 50,000 trials of the tree sampling as it failed to find any successes.

Figure 9 shows the log q-errors with and without the graph sampling for 631 hard queries on WordNet that required the graph sampling. The x-axis indicates queries from those where the tree sampling overestimated to those where it underestimated. Each query is represented by a red dot and a green dot, denoting the log q-errors with and without the graph sampling, respectively. The results indicate that the tree sampling can lead to severe overestimation when one or two extremely unlikely successes occur, or severe underestimation when there are no successes. For these hard queries, the graph sampling provides reasonably accurate estimates. Even for the aforementioned extreme case, the graph sampling achieves a q-error of about 1.6.

8 CONCLUSION

In this paper, we have proposed FaST_{est} for subgraph cardinality estimation. Our novel filtering-sampling approach synergistically combines (1) a strong filtering method for drastically reducing the sample space, (2) an efficient and accurate estimation by candidate tree sampling, and (3) a worst-case optimal stratified graph sampling with outstanding accuracy on hard instances. Extensive experiments on real-world datasets show that FaST_{est} significantly outperforms state-of-the-art algorithms in terms of accuracy and time, while maintaining comparable memory usage.

ACKNOWLEDGEMENTS

W. Shin, S. Song, and K. Park were supported by Institute of Information communications Technology Planning Evaluation (IITP) grant funded by the Korea government (MSIT) (No. 2018-0-00551, Framework of Practical Algorithms for NP-hard Graph Problems). W.-S. Han was supported by Institute of Information communications Technology Planning Evaluation (IITP) grant funded by the Korea government (MSIT) (No. 2021-0-00859, Development of a Distributed Graph DBMS for Intelligent Processing of Big Graph).

REFERENCES

- [1] 2024. *Boost C++ Libraries*. Retrieved 2024-03-16 from <http://www.boost.org/>
- [2] Sepehr Assadi, Michael Kapralov, and Sanjeev Khanna. 2018. A Simple Sublinear-Time Algorithm for Counting Arbitrary Subgraphs via Edge Sampling. In *Proceedings of ITCS*. 6:1–6:20.
- [3] Albert Atserias, Martin Grohe, and Dániel Marx. 2008. Size bounds and query plans for relational joins. In *Proceedings of FOCS*. 739–748.
- [4] Suman K. Bera and C. Seshadhri. 2020. How to Count Triangles, without Seeing the Whole Graph. In *Proceedings of ACM SIGKDD*. 306–316.
- [5] Bibek Bhattarai, Hang Liu, and H. Howie Huang. 2019. CECI: Compact Embedding Cluster Index for Scalable Subgraph Matching. In *Proceedings of ACM SIGMOD*. 1447–1462.
- [6] Fei Bi, Lijun Chang, Xuemin Lin, Lu Qin, and Wenjie Zhang. 2016. Efficient Subgraph Matching by Postponing Cartesian Products. In *Proceedings of ACM SIGMOD*. 1199–1214.
- [7] Vincenzo Bonnici and Rosalba Giugno. 2017. On the Variable Ordering in Subgraph Isomorphism Algorithms. *IEEE/ACM Transactions on Computational Biology and Bioinformatics* 14 (2017), 193–203. Issue 1.
- [8] Vincenzo Bonnici, Rosalba Giugno, Alfredo Pulvirenti, Dennis Shasha, and Alfredo Ferro. 2013. A subgraph isomorphism algorithm and its application to biochemical data. *BMC Bioinformatics* 14 (2013).
- [9] Marco Bressan, Stefano Leucci, and Alessandro Panconesi. 2019. Motivo: Fast Motif Counting via Succinct Color Coding and Adaptive Sampling. In *Proceedings of the VLDB Endowment*. 1651–1663.
- [10] Lawrence D. Brown, T. Tony Cai, and Anirban DasGupta. 2001. Interval Estimation for a Binomial Proportion. *Statist. Sci.* 16, 2 (2001), 101–133.
- [11] George Casella and Roger L. Berger. 2021. *Statistical inference*. Cengage Learning.
- [12] Xiaowei Chen and John C. S. Lui. 2016. Mining Graphlet Counts in Online Social Networks. In *Proceedings of IEEE ICDM*. 71–80.
- [13] Yu Chen and Ke Yi. 2020. Random Sampling and Size Estimation Over Cyclic Joins. In *Proceedings of ICDT*. 7:1–7:18.
- [14] C. J. Clopper and E. S. Pearson. 1934. The Use of Confidence or Fiducial Limits Illustrated in the Case of the Binomial. *Biometrika* 26, 4 (1934), 404–413.
- [15] L.P. Cordella, P. Foggia, C. Sansone, and M. Vento. 2004. A (sub)graph isomorphism algorithm for matching large graphs. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 26, 10 (2004), 1367–1372.
- [16] Wenfei Fan. 2012. Graph Pattern Matching Revised for Social Network Analysis. In *Proceedings of ICDT*. 8–21.
- [17] L. R. Ford and D. R. Fulkerson. 1956. Maximal Flow Through a Network. *Canadian Journal of Mathematics* 8 (1956), 399–404.
- [18] Myoungji Han, Hyunjoon Kim, Geonmo Gu, Kunsoo Park, and Wook-Shin Han. 2019. Efficient Subgraph Matching: Harmonizing Dynamic Programming, Adaptive Matching Order, and Failing Set Together. In *Proceedings of ACM SIGMOD*. 1429–1446.
- [19] Wook-Shin Han, Jinsoo Lee, and Jeong-Hoon Lee. 2013. Turboiso: towards ultrafast and robust subgraph isomorphism search in large graph databases. In *Proceedings of ACM SIGMOD*. 337–348.
- [20] Huahai He and Ambuj K. Singh. 2008. Graphs-at-a-Time: Query Language and Access Methods for Graph Databases. In *Proceedings of ACM SIGMOD*. 405–418.
- [21] Tamás Horváth, Thomas Gärtner, and Stefan Wrobel. 2004. Cyclic Pattern Kernels for Predictive Graph Mining. In *Proceedings of ACM SIGKDD*. 158–167.
- [22] Hunt, Harry B. and Marathe, Madhav V. and Radhakrishnan, Venkatesh and Stearns, Richard E. 1998. The Complexity of Planar Counting Problems. *Siam Journal on Computing* 27, 4 (1998), 1142–1167.
- [23] Shweta Jain and C. Seshadhri. 2017. A Fast and Provable Method for Estimating Clique Counts Using Turán’s Theorem. In *Proceedings of ACM WWW*. 441–449.
- [24] Hyunjoon Kim, Yunyoung Choi, Kunsoo Park, Xuemin Lin, Seok Hee Hong, and Wook Shin Han. 2021. Versatile Equivalences: Speeding up Subgraph Query Processing and Subgraph Matching. In *Proceedings of ACM SIGMOD*. 925–937.
- [25] Jinha Kim, Hyungyu Shin, Wook-Shin Han, Sungpack Hong, and Hassan Chafi. 2015. Taming subgraph isomorphism for RDF query processing. In *Proceedings of the VLDB Endowment*. 1238–1249.
- [26] Kyoungmin Kim, Hyeonji Kim, George Fletcher, and Wook Shin Han. 2021. Combining Sampling and Synopses with Worst-Case Optimal Runtime and Quality Guarantees for Graph Pattern Cardinality Estimation. In *Proceedings of ACM SIGMOD*. 964–976.
- [27] Feifei Li, Bin Wu, Ke Yi, and Zhuoyue Zhao. 2016. Wander Join: Online Aggregation via Random Walks. In *Proceedings of ACM SIGMOD*. 615–629.
- [28] Don R. Lick and Arthur T. White. 1970. k -Degenerate Graphs. *Canadian Journal of Mathematics* 22, 5 (1970).
- [29] Sharon L. Lohr. 2021. *Sampling: design and analysis*. CRC press.
- [30] Dániel Marx and Michal Pilipczuk. 2014. Everything you always wanted to know about the parameterized complexity of Subgraph Isomorphism (but were afraid to ask). In *Proceedings of STACS*. 542–553.
- [31] Amine Mhedhbi and Semih Salihoglu. 2018. Optimizing subgraph queries by combining binary and worstcase optimal joins. In *Proceedings of the VLDB Endowment*. 1692–1704.
- [32] Guido Moerkotte, Thomas Neumann, and Gabriele Steidl. 2009. Preventing Bad Plans by Bounding the Impact of Cardinality Estimation Errors. In *Proceedings of the VLDB Endowment*. 982–993.
- [33] Raffaele Di Natale, Alfredo Ferro, Rosalba Giugno, Misael Mongiovi, Alfredo Pulvirenti, and Dennis Shasha. 2010. SING: Subgraph search In Non-homogeneous Graphs. *BMC Bioinformatics* 11 (2010), 96. Issue 1.
- [34] Thomas Neumann and Guido Moerkotte. 2011. Characteristic sets: Accurate cardinality estimation for RDF queries with multiple joins. In *Proceedings of IEEE ICDE*. 984–994.
- [35] Yeonsu Park, Seongyun Ko, Sourav S. Bhowmick, Kyoungmin Kim, Kijae Hong, and Wook Shin Han. 2020. G-CARE: A Framework for Performance Benchmarking of Cardinality Estimation Techniques for Subgraph Matching. In *Proceedings of ACM SIGMOD*. 1099–1114.
- [36] R. C. Prim. 1957. Shortest connection networks and some generalizations. *The Bell System Technical Journal* 36, 6 (1957), 1389–1401.
- [37] Nataša Pržulj. 2007. Biological network comparison using graphlet degree distribution. *Bioinformatics* 23, 2 (2007), e177–e183.
- [38] Siddhartha Sahu, Amine Mhedhbi, Semih Salihoglu, Jimmy Lin, and M. Tamer Özsu. 2017. The ubiquity of large graphs and surprising challenges of graph processing. In *Proceedings of the VLDB Endowment*. 420–431.
- [39] Nino Shervashidze, SVN Vishwanathan, Tobias Petri, Kurt Mehlhorn, and Karsten Borgwardt. 2009. Efficient graphlet kernels for large graph comparison. In *Proceedings of the AISTATS*. 488–495.
- [40] Wonseok Shin, Siwoo Song, Kunsoo Park, and Wook-Shin Han. 2023. Cardinality Estimation of Subgraph Matching: A Filtering-Sampling Approach. arXiv:2309.15433
- [41] Giorgio Stefanoni, Boris Motik, and Egor V. Kostylev. 2018. Estimating the Cardinality of Conjunctive Queries over RDF Data Using Graph Summarisation. *Proceedings of ACM WWW*, 1043–1052.
- [42] Shixuan Sun and Qiong Luo. 2020. In-Memory Subgraph Matching: An In-depth Study. In *Proceedings of ACM SIGMOD*. 1083–1098.
- [43] Shixuan Sun and Qiong Luo. 2022. Subgraph Matching With Effective Matching Order and Indexing. *IEEE Transactions on Knowledge and Data Engineering* 34, 1 (2022), 491–505.
- [44] Tamir Tassa. 2012. Finding all maximally-matchable edges in a bipartite graph. *Theoretical Computer Science* 423 (2012), 50–58.
- [45] J. R. Ullmann. 1976. An Algorithm for Subgraph Isomorphism. *J. ACM* (1976), 31–42.
- [46] Princeton University. 2010. About WordNet. Retrieved 2023-04-07 from <https://wordnet.princeton.edu/>.
- [47] David Vengerov, Andre Cavalheiro Menck, Mohamed Zait, and Sunil P. Chakkapen. 2015. Join Size Estimation Subject to Filter Conditions. In *Proceedings of the VLDB Endowment*. 1530–1541.
- [48] Hanchen Wang, Rong Hu, Ying Zhang, Lu Qin, Wei Wang, and Wenjie Zhang. 2022. Neural Subgraph Counting with Wasserstein Estimator. In *Proceedings of ACM SIGMOD*. 160–175.
- [49] Xiaowei Ye, Rong Hua Li, Qiangqiang Dai, Hongzhi Chen, and Guoren Wang. 2022. Lightning Fast and Space Efficient k -clique Counting. In *Proceedings of ACM WWW*. 1191–1202.
- [50] Kangfei Zhao, Jeffrey Xu Yu, Hao Zhang, Qiyang Li, and Yu Rong. 2021. A Learned Sketch for Subgraph Counting. In *Proceedings of ACM SIGMOD*. 2142–2155.
- [51] Zhuoyue Zhao, Robert Christensen, Feifei Li, Xiao Hu, and Ke Yi. 2018. Random Sampling over Joins Revisited. In *Proceedings of ACM SIGMOD*. 1525–1539.
- [52] Dongxiao Zhu and Zhaohui S Qin. 2005. Structural comparison of metabolic networks in selected single cell organisms. *BMC Bioinformatics* 6 (2005), 8. Issue 1.