

Metric Indexing for Graph Similarity Search^{*}

Franka Bause¹[0000-0003-4202-3692], David B. Blumenthal²[0000-0001-8651-750X],
Erich Schubert³[0000-0001-9143-4880], and Nils M. Kriege¹[0000-0003-2645-947X]

¹ University of Vienna, Faculty of Computer Science, Vienna, Austria
{franka.bause,nils.kriege}@univie.ac.at

² Friedrich-Alexander University Erlangen-Nürnberg (FAU), Department Artificial
Intelligence in Biomedical Engineering (AIBE), Erlangen, Germany
david.b.blumenthal@fau.de

³ TU Dortmund University, Department of Computer Science, Dortmund, Germany
erich.schubert@tu-dortmund.de

Abstract. Finding the graphs that are most similar to a query graph in a large database is a common task with various applications. A widely-used similarity measure is the graph edit distance, which provides an intuitive notion of similarity and naturally supports graphs with vertex and edge attributes. Since its computation is NP-hard, techniques for accelerating similarity search have been studied extensively. However, index-based approaches for this are almost exclusively designed for graphs with categorical vertex and edge labels and uniform edit costs. We propose a filter-verification framework for similarity search, which supports non-uniform edit costs for graphs with arbitrary attributes. We employ an expensive lower bound obtained by solving an optimal assignment problem. This filter distance satisfies the triangle inequality, making it suitable for acceleration by metric indexing. In subsequent stages, assignment-based upper bounds are used to avoid further exact distance computations. Our extensive experimental evaluation shows that a significant runtime advantage over both a linear scan and state-of-the-art methods is achieved.

Keywords: Graphs · Similarity search · Graph edit distance.

1 Introduction

Graph-structured data is ubiquitous in many areas such as chemo- and bioinformatics or computer vision. A common task is to search a database containing a large number of graphs for those that are most similar to a given query graph. Such queries are submitted directly by the user or occur as subproblems in downstream machine learning algorithms. A widely accepted concept of graph

^{*} This work has been supported by the Vienna Science and Technology Fund (WWTF) through project VRG19-009, and Deutsche Forschungsgemeinschaft (DFG), project number 124020371, within the Collaborative Research Center SFB 876 “Providing Information by Resource-Constrained Analysis”, SFB projects A2 and A6.

similarity is the *graph edit distance*, which is the minimum cost for transforming one graph into the other by a sequence of edit operations. A strength of this measure is that it can elegantly be applied to graphs with vertex and edge attributes by defining the costs of edit operations adequately. For example, to compare protein graphs where vertices are annotated by the amino acid sequence of the secondary structure elements they represent, the Levenshtein distance was used [20].

However, the vast majority of efficient methods for similarity search in graph databases are limited to the special case where graphs have categorical labels and the costs of edit operations are uniform (either zero or one) [28, 30, 26, 25, 29, 31, 16, 13, 10]. A fairly recent development in this domain are neural graph embeddings, e.g. [19], which do not return exact similarity search results. For the pairwise computation of the graph edit distance, several exact approaches [10, 15] and heuristics such as *bipartite graph matching* based on optimal vertex assignments [20] have been proposed, many of which support the graph edit distance in its full generality [15, 20]. Several of these yield lower and upper bounds on the graph edit distance as a byproduct, which have just recently been compared systematically [4]. However, these lower bounds for the general graph edit distance are not yet widely used for similarity search in graph databases. For the methods based on optimal vertex assignments, it has only recently been shown how to derive a distance termed *BRANCH* that is guaranteed to be a lower bound and proven to satisfy the triangle inequality [2]. *BRANCH* has been shown to provide an excellent trade-off between tightness and running time [4].

We propose a filter-verification framework for similarity search, which supports the general graph edit distance with arbitrary metric edit costs and is hence suitable for graphs with any attributes comparable with a distance measure. We employ *BRANCH* as an initial filter accelerated by metric indexing. In the next stages, we derive upper bounds from the optimal assignment and improve them via local search to reduce the candidate set further, before performing verification by exact computation of the graph edit distance. We experimentally evaluate our approach on graphs with attributes and categorical labels showing the effectivity of the filter pipeline. The results show that our approach allows scalable similarity search in attributed graphs with non-uniform edit costs. For the special case of uniform edit costs, where competing methods are available, our approach is shown to outperform the state of the art.

2 Related Work

We discuss approaches for similarity search regarding the graph edit distance and methods for its pairwise exact or approximate computation.

2.1 Similarity Search in Graph Databases

Methods for similarity search in graph databases can be divided into two categories, depending on whether they compare overlapping or non-overlapping substructures. The methods *k-AT* [25], *CStar* [28], *Segos* [26] and *GSim* [30] belong

to the first category. These techniques are inspired by the q -grams used in the computation of the string edit distance. Either q -grams based on trees [25, 28, 26] or paths [30] are used. The methods *Pars* [29], *MLIndex* [16], and *Inves* [13] partition the graphs into non-overlapping substructures. They essentially obtain lower bounds based on the observation, that if x non-overlapping substructures of a database graph are not contained in the query graph, the graph edit distance is at least x . *Pars* uses a dynamic partitioning approach to achieve this, while *MLIndex* uses a multi-layered index to manage multiple partitions for each graph. *Inves* is a method used to verify whether the graph edit distance of two graphs is below a specified threshold by first trying to generate enough mismatching non-overlapping substructures. *Mixed* [31] combines the idea of q -grams and graph partitioning. These methods only allow uniform edit costs and are therefore not suitable for graphs with continuous attributes.

The concept of a *median graph* of a set of graphs regarding the graph edit distance has been studied extensively, see [3] and references therein. An application of median graphs is their use as routing objects in hierarchical index structures [23, 3]. However, we are not aware of any concrete realization using this concept in a setting comparable to ours.

2.2 Pairwise Computation of the Graph Edit Distance

For computing the exact graph edit distance, both general-purpose algorithms [15] as well as approaches tailored to the verification step in graph databases have been proposed [9], which are usually based on depth- or breadth-first search [12, 9], or integer linear programming [15]. As the exact computation of the graph edit distance is not feasible for larger graphs, many heuristics have been proposed, e.g., [14, 20, 4, 18, 2, 11]. The properties of the dissimilarities obtained from these are in general not well investigated. For heuristics based on optimal vertex assignment [20], which are widely used in practice [24], a variant called BRANCH was recently studied thoroughly [2]. BRANCH is a lower bound on the graph edit distance, a pseudo-metric on graphs and supports arbitrary cost models (c.f., Section 4.1).

3 Preliminaries

We introduce the required basic concepts of graph theory and discuss database search with a focus on the metric space.

3.1 Graph Theory

A *graph* $G = (V, E, \mu, \nu)$ consists of a set of vertices $V(G)$, a set of edges $E(G) \subseteq V(G) \times V(G)$ between vertices of G , a labeling function for the vertices $\mu: V(G) \rightarrow L$, and a labeling function for the edges $\nu: E(G) \rightarrow L$. We discuss only undirected graphs and denote an edge between u and v by uv . The set of neighbors of a vertex $v \in V(G)$ is denoted by $N(v) = \{u \mid uv \in E(G)\}$.

Table 1: Notation for edit costs.

$c_v(u, v)$	Cost of substituting vertex u with vertex v (adjusting the label/attributes)
$c_v(u, \epsilon)$	Cost of deleting the isolated vertex u
$c_v(\epsilon, v)$	Cost of inserting the isolated vertex v
$c_e(uv, wx)$	Cost of substituting edge uv with edge wx (adjusting the label/attributes)
$c_e(uv, \epsilon)$	Cost of deleting the edge uv
$c_e(\epsilon, wx)$	Cost of inserting the edge wx

The set L can be categorical labels or arbitrary attributes including real-valued vectors and complex objects such as strings.

A measure commonly used to describe the dissimilarity of two graphs is the *graph edit distance*, which is the minimum cost for transforming one graph into the other using edit operations. An *edit operation* can be deleting or inserting an isolated vertex or an edge or relabeling any of the two. An *edit path* from graph G_1 to G_2 is a sequence of edit operations (e_1, e_2, \dots) that transforms G_1 into G_2 .

Definition 1 (Graph Edit Distance [20]). *Let c be an edit cost function assigning non-negative costs to edit operations. The graph edit distance between two graphs G_1 and G_2 is defined as*

$$d_{ged}(G_1, G_2) = \min \left\{ \sum_{i=1}^k c(e_i) \mid (e_1, \dots, e_k) \in \Upsilon(G_1, G_2) \right\},$$

where $\Upsilon(G_1, G_2)$ is the set of all possible edit paths from G_1 to G_2 .

The costs of the different edit operations can be chosen as required for the specific use case, see Table 1 for our notation. If the edit costs are symmetric, non-negative, and strictly positive for each non-identical edit operation, the graph edit distance is a metric on graphs, treating graph isomorphism as identity [3]. Note that this holds even if the edit costs do not satisfy the triangle inequality (and hence are no metric), because the graph edit distance uses the edit path with minimal cost. In this work, we nonetheless assume that the edit costs respect the triangle inequality, i.e., we assume that the following inequalities hold:⁴

$$c_v(u, w) \leq c_v(u, v) + c_v(v, w) \quad \forall (u, v, w) \in \mathcal{V}^3 \quad (1)$$

$$c_v(u, v) \leq c_v(u, \epsilon) + c_v(\epsilon, v) \quad \forall (u, v) \in \mathcal{V}^2 \quad (2)$$

$$c_e(uv, yz) \leq c_e(uv, wx) + c_e(wx, yz) \quad \forall (uv, wx, yz) \in \mathcal{E}^3 \quad (3)$$

$$c_e(uv, wx) \leq c_e(uv, \epsilon) + c_e(\epsilon, wx) \quad \forall (uv, wx) \in \mathcal{E}^2 \quad (4)$$

Equations (1), (3), and (4) can be enforced via pre-processing without changing the graph edit distance and can hence be assumed to hold w.l.o.g. [5].

⁴ For simplicity of notation, we have defined the costs on the vertices and edges instead of their labels. Hence, the sets \mathcal{V} and \mathcal{E} are all possible vertices and edges, respectively.

E.g., if we have $c(u, v) > c(u, w) + c(w, v)$, we can simply substitute $c(u, v)$ with $c(u, w) + c(w, v)$, because a minimum cost edit path cannot contain $c(u, v)$. The only remaining constraint, Equation (2), is met (to the best of our knowledge) in all applications where the graph edit distance is used to address real-world problems [24]. Computing the graph edit distance is NP-hard [29], rendering exact computation possible for small graphs only. There are several heuristics, many of which are based on solving an assignment problem.

Definition 2 (Assignment Problem). *Let A and B be two sets with $|A| = |B| = n$ and $c: A \times B \rightarrow \mathbb{R}$ a cost function. An assignment between A and B is a bijection $f: A \rightarrow B$. The cost of an assignment f is $c(f) = \sum_{a \in A} c(a, f(a))$. The assignment problem is to find an assignment with minimum cost.*

For an assignment instance (A, B, c) , we denote the cost of an optimal assignment by $d_{\text{oa}}^c(A, B)$. The assignment problem can be solved in cubic running time using a suitable implementation of the Hungarian method [8].

3.2 Searching in Databases

Databases provide means to store data to be able to retrieve, insert or change it efficiently. In the context of data analysis, retrieval (search) is usually the crucial operation on databases, because it will be performed much more often than updates. We focus on two important types of similarity queries when searching a database DB, the first of which is the *range query* for a radius r :

Definition 3 (Range Query). *A range query $\text{range}(q, r)$, with query object q and range r , returns all objects in the database with a distance to the query object not exceeding the range, i.e., $\text{range}(q, r) = \{o \in DB \mid d(o, q) \leq r\}$.*

The second type of query considered here is the *k -nearest neighbor query*.

Definition 4 (k -Nearest Neighbor Query). *A k -nearest neighbor query (k NN query) $\text{NN}(q, k)$ with query object q and parameter k returns the smallest set $\text{NN}(q, k) \subseteq DB$, so that $|\text{NN}(q, k)| \geq k$ and $\forall o \in \text{NN}(q, k), \forall o' \in DB \setminus \text{NN}(q, k) : d(o, q) < d(o', q)$.*

In conjunction with range queries, it is preferable to return all the objects with a distance (including the query object, if part of the database), that does not exceed the distance to the k th neighbor, which may be more than k objects when tied. That yields an equivalency of the results of k NN queries and range queries, i.e., we have $\text{range}(q, r) = \text{NN}(q, |\text{range}(q, r)|)$ and $\text{NN}(q, k) = \text{range}(q, r_k)$, where r_k is the maximum distance in $\text{NN}(q, k)$. Provided that the distance used is a metric, both types of queries can be accelerated using metric indices. In our work, we use the *vantage point tree (vp-tree)* [27] as a classical method and the more recent *cover tree* [1], because they are available in the ELKI framework [21], but others could also be used. While the vp-tree is a height balanced binary tree dividing the data into *near* and *far* halves of the dataset based on the

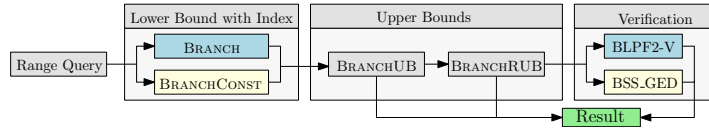


Fig. 1: Overview of the filter pipeline. For general metric edit costs the blue modules are used; yellow modules are more efficient for uniform (edge) edit costs.

median distance from the vantage point, the cover tree controls the expansion rate by reducing the maximum radius in each level of the tree, branching out if necessary into multiple branches. In both trees, queries are performed top-down by traversing all paths that cannot be dismissed using the routing objects and employing the triangle inequality.

4 Efficient Filtering for the General Graph Edit Distance

We propose a filter pipeline for range queries regarding the graph edit distance following a common paradigm for expensive distances, see e.g. [28]. Here, lower bounds allow to filter out graphs that do not satisfy the query predicate. For the remaining candidates, upper bounds are evaluated to add them immediately to the result set without exact distance computation. Finally, in the verification step, the exact distance is computed for the remaining candidates only. Our approach starts with the optimal assignment based lower bound BRANCH accelerated by metric indexing. From the same optimal assignment, an upper bound is derived (BRANCHUB) and subsequently refined by local search (BRANCHRUB) before the remaining candidates are verified. The pipeline is illustrated in Figure 1, the individual steps are described in the following.

4.1 Index-Accelerated Lower Bound Filtering

Several lower bounds on the graph edit distance have been proposed or can be derived from known heuristics, see [4]. One of the most effective lower bounds with an excellent trade-off between tightness and runtime is referred to as BRANCH.

Definition 5 (Branch Distance). For two graphs G_1 and G_2 the branch distance is defined as $d_{branch}(G_1, G_2) = d_{\text{oa}}^c(V(G_1) \cup \varepsilon_1, V(G_2) \cup \varepsilon_2)$, where ε_i denotes a multiset of ϵ elements, so that $|V(G_i) \cup \varepsilon_i| = |V(G_1) \cup V(G_2)|$ for $i \in \{1, 2\}$, and

$$c(u, v) = \begin{cases} 0 & \text{if } u = v = \epsilon \\ c_v(u, v) + d_e(u, v) & \text{if } u \neq \epsilon \text{ and } v \neq \epsilon \\ c_v(\epsilon, v) + 1/2 \cdot \sum_{n \in N(v)} c_e(\epsilon, vn) & \text{if } u = \epsilon \text{ and } v \neq \epsilon \\ c_v(u, \epsilon) + 1/2 \cdot \sum_{n \in N(u)} c_e(un, \epsilon) & \text{if } u \neq \epsilon \text{ and } v = \epsilon \end{cases}$$

with $d_e(u, v) = d_{\text{oa}}^{c'}(N(u) \cup \varepsilon_u, N(v) \cup \varepsilon_v)$, where $c'(w, x) = 1/2 \cdot c_e(uw, vx)$.

Remark 1. Note that, by using a customized version of the Hungarian algorithm, BRANCH can also be implemented in a slightly more efficient way, where only one dummy vertex ϵ is added to the vertex sets $V(G_1)$ and $V(G_2)$ (see [7] for details). In this paper, we use the classical implementation employed in [20, 2], which corresponds to the characterization provided in Definition 5.

BRANCH has its origin in one of the most successful heuristics for the graph edit distance proposed by Riesen and Bunke [20]. However, in contrast to the original approach, it is guaranteed to underestimate the graph edit distance by dividing all edge costs by two to avoid that the cost of a single edge edit operation is counted twice, once for each endpoint [2]. Since an instance of the assignment problem on the vertices of the two graphs has to be solved, and for each vertex pair an assignment on their edges, BRANCH can be computed in $O(n^2\Delta^3 + n^3)$ time for graphs with n vertices and maximum degree Δ . In the case of uniform edge edit costs, d_e can be computed by multiset intersection of edge labels and the running time reduces to $O(n^3)$. This special case is referred to as BRANCHCONST [2]. It has been shown that, if the edit costs are metric, the branch distance is a pseudo-metric on graphs [2]. This allows to accelerate computing the candidate set w. r. t. this lower bound by employing metric indexing.

4.2 Upper Bound Filtering and Verification

From the solution of the assignment problem of BRANCH, an upper bound can be obtained by deriving the corresponding edit path [20], denoted BRANCHUB here. By definition of the graph edit distance, the cost of every edit path is an upper bound of the graph edit distance. Following [28], we refine the assignment by local search to gain a tighter upper bound. Starting with the assignment obtained for the lower bound, the mapping of two vertex pairs is iteratively swapped, and kept whenever it induces a cheaper edit path, until there is no improvement. We refer to the refined upper bound obtained from the BRANCH assignment as BRANCHRUB.

Eventually, the graphs that were neither filtered out by the lower bound nor approved by the upper bounds are verified by exact graph edit distance computation. We use *BSS_GED* [10] for datasets with discrete labels and uniform costs and *BLPF2-V* otherwise. The latter is based on the integer programming formulation F2 of [15] with the additional constraint that the objective function does not exceed the threshold to allow for early termination.

4.3 Nearest-Neighbor Queries

For k NN queries it is not possible to separate the different steps of the filter pipeline as clearly as shown in Figure 1. We realize k NN queries using the *optimal multi-step k-nearest neighbor search* algorithm [22]. The database graphs are scanned in ascending order according the lower bound BRANCH regarding the query graph. For each graph, the exact graph edit distance is computed and the k

Table 2: Datasets and their statistics [17]. Some datasets contain graphs with labeled or attributed vertices and edges, as can be seen in the last two columns.

Name	Graphs	avg Vertices	avg Edges	Labels (V/E)	Attributes (V/E)
<i>Cuneiform</i>	267	21.27	44.80	+/+	+/+
<i>Fingerprint</i>	2800	5.42	4.42	-/-	+/+
<i>Letter-high</i>	2250	4.67	4.50	-/-	+/-
<i>Letter-low</i>	2250	4.68	3.13	-/-	+/-
<i>MUTAG</i>	188	17.93	19.79	+/+	-/-
<i>PTC_FM</i>	349	14.11	14.48	+/+	-/-
<i>QM9</i>	129433	18.03	18.63	-/-	+/+

graphs with the smallest exact graph edit distance are maintained. Once we have found at least k objects with an exact distance smaller than the lower bound of all remaining objects, the search can be terminated. This is optimal in the sense that none of the exact distance computations could have been avoided [22]. Accessing the graphs ordered regarding the BRANCH lower bound can be achieved naively by sorting all graphs, or by using suitable metric index structures.

5 Experimental Evaluation

In this section, we experimentally address the following research questions:

- Q1 What speed-up in range queries can be achieved when using metric indices compared to a linear scan of the database?
- Q2 How effective are the individual lower and upper bounds in our pipeline?
- Q3 Can the proposed filter pipeline compete with state-of-the-art methods for uniform edit costs?
- Q4 What speed-up in k NN queries can be achieved when using metric indices?
- Q5 Does the proposed filter pipeline scale to a very large dataset?

5.1 Setup

As metric index we chose the *vp-tree* as a classical method and the *cover tree* as a state-of-the-art approach. For both we used the implementation provided by ELKI [21] with a sample size of 5 for the *vp-tree* and an expansion rate of 1.2 for the *cover tree*.

For a comparison in databases containing graphs with categorical labels, we used *MLIndex* [16] and *GSim* [30], since the former is considered state-of-the-art, while the latter provided much better results in our experiments. For *MLIndex* the number of partitions was set to $\text{threshold} + 1$ and in *GSim* all provided filters were used. We used the implementations provided by the authors. In addition, we used *CStar* [28], which follows a filter-verification approach related to ours. For verification we used *BSS_GED* [10] and *BLPF2-V* [15] with the Gurobi solver.

We conducted experiments on a wide range of real-world datasets with different characteristics, see Table 2. The costs of inserting, deleting or relabeling

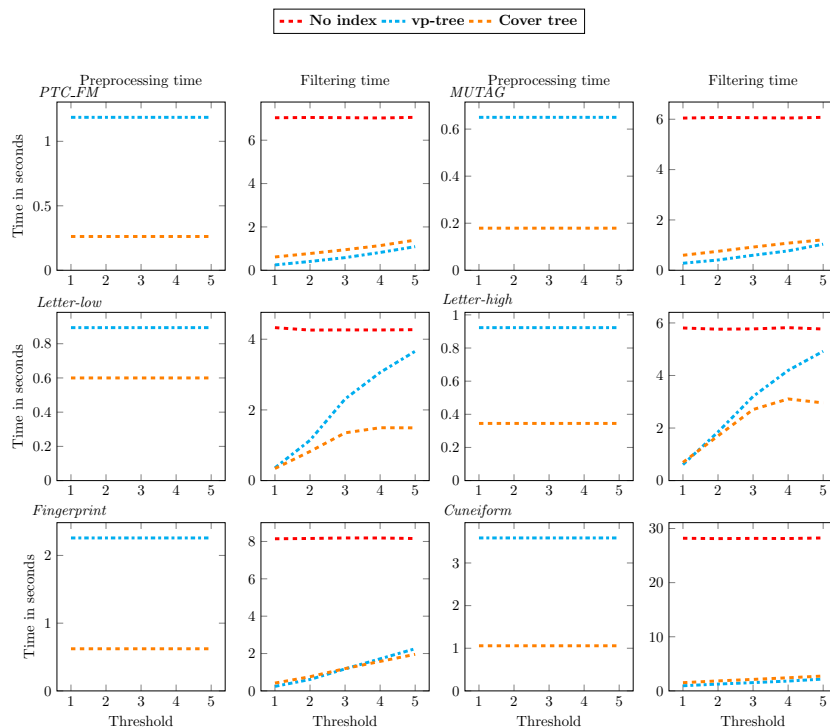


Fig. 2: Runtime comparison for filtering 100 range queries using BRANCH with thresholds 1 to 5 and preprocessing time for constructing the index.

a vertex or edge with a categorical label were set to 1, which is equivalent to the fixed setting in *MLIndex*, *GSim*, and *CStar*. For continuous attributes, the Euclidean distance was used to define the relabeling cost. For simplicity, we did not use domain-specific distances. Continuous attributes were normalized to the range $[0, 1]$ (separately for each dimension), to make distances roughly comparable between different datasets.

5.2 Results

We report on our findings regarding the above research questions.

Q1: Speed-up of range queries through metric indices. We first investigate how much of a speed-up can be achieved by using an index structure when filtering candidates for a range query by a lower bound. We randomly sampled 100 graphs from the dataset to be queries and then performed lower bound filtering without an index, using the cover tree, and the vp-tree.

Figure 2 shows the time needed for filtering 100 range queries (each with thresholds 1 to 5) and additionally the preprocessing time for index construction. The runtime does not depend on the given threshold for a linear scan, but

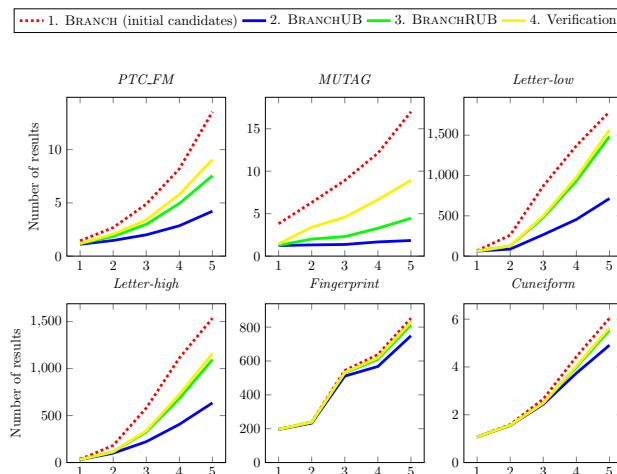


Fig. 3: Average number of initial candidates (dashed) and hits identified in the different stages of the filter pipeline for each threshold.

increases for the metric indices with the threshold, in particular for the *Letter*-datasets. It can be seen that, while on most datasets both index methods provide the same runtime benefit for filtering, the cover tree is much faster in preprocessing than the vp-tree. The runtime advantage on the *Letter*-datasets is quite small for larger thresholds. The runtime of the index structures directly corresponds to the number of BRANCH distance computations. Compared to the cover tree, the vp-tree requires many more distance computations in the preprocessing due to the chosen sample size. In general, the runtime corresponds to the number of candidates, which we investigate in the following.

Q2: Filter pipeline. In this experiment we investigate how the candidate and result set are updated during filtering. Figure 3 shows the average number of candidates for BRANCH and the number of results after each step for 100 range queries. When comparing the size of the candidate sets with the results of the previous experiment, it can be seen, that the runtime for filtering highly depends on the number of candidates. For some datasets almost all candidates remaining after the upper bound filtering are not results. This indicates that improvement is possible with tighter lower bounds. In general, BRANCHRUB manages to report almost all results, except in dataset *MUTAG*.

Q3: Comparison with state-of-the-art methods. Many methods for similarity search in graph databases limited to uniform edit costs have been proposed. We compare to *MLIndex* [16], *CStar* [28] and *GSim* [30]. We used *BSS-GED* [10] for a fast verification in our filter pipeline, as well as in *CStar*. The implementations of *MLIndex* and *GSim* contain their own verification algorithm. Figure 4 shows the runtime for preprocessing and filtering as well as the total query time

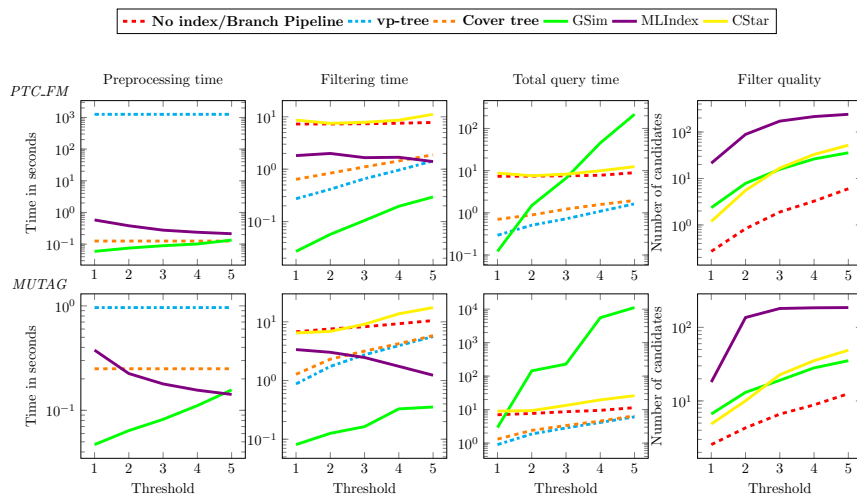


Fig. 4: Runtime for answering 100 range queries and average number of candidates remaining after applying all filters for thresholds 1 to 5. Our approaches are shown with dashed lines and are marked bold in the legend. For *MIndex* no verification time is given, since it did not finish within the time limit of 2 days.

including filtering and verification for 100 range queries. The average number of candidates remaining after application of all filters in the different methods is also shown. Only these need to be verified by exact graph edit distance computation. For *BRANCH* only one line is shown, since linear scan, *vp-tree* and the *cover tree* variant apply the same filters and generate the same candidates.

MIndex produces the largest candidate set, and did not finish the verification process in the time limit. It can be seen that, while *GSim* is quite fast in preprocessing and filtering, the verification step takes a long time. This is due to a combination of a slower verification algorithm and a higher number of candidates that have to be verified. The results indicate that, even when using *BSS_GED* for verification, the approach would not be competitive with the *cover tree* due to the high number of candidates. *CStar* needs much more time for filtering and cannot filter out as many candidates leading also to a higher verification time. Interestingly, the time for verification does not increase proportionally to the number of candidates, which might indicate, that the verification algorithm needs more time to verify certain *difficult* graphs.

Q4: Speed-up of k NN queries through metric indices. We investigate how much of a speed-up can be achieved by using an index structure compared to not using one, when answering k NN queries using the optimal multi-step k -nearest neighbor search, cf. Section 4.3. We randomly sampled 20 graphs from the dataset to be queries and then used the cover tree as well as the *vp-tree* as the underlying metric index to compare them. Figure 5 shows the time needed for answering 20 k NN queries, each with $k \in \{1, \dots, 5\}$ (excluding preprocessing). Since in the

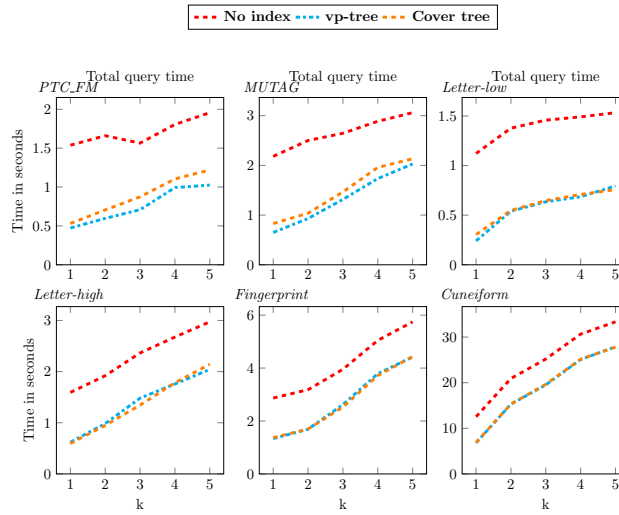


Fig. 5: Runtime comparison for answering 20 k NN queries using BRANCH and $k \in \{1, \dots, 5\}$.

optimal multi-step k -nearest neighbor search, the candidates have to be verified during search, before further candidates are explored, the runtime also includes the time needed for verification. It can be seen that, again both index structures provide the same runtime benefit. Taking into account the preprocessing time however, the cover tree has a clear advantage over the vp-tree.

Q5: Similarity search in a large dataset. We investigate the scalability of our approach on the dataset *QM9* with 129 433 graphs with attributed vertices and edges. The results shown in Figure 6 confirm the high preprocessing time of the vp-tree compared to the cover tree. Both index methods achieve a significant advantage over a linear scan in filtering by reducing the running time by several orders of magnitude depending on the selectivity of the query.

6 Conclusions

We have shown that the recently studied lower and upper bounds on the graph edit distance can be employed to realize scalable graph similarity search in a filter-verification framework accelerated by metric indexing. Our approach supports attributed graphs without restrictions of edit costs. For the extensively studied special case of graphs with discrete labels and uniform edit costs, our approach was shown experimentally to outperform the state-of-the-art methods.

There are several directions of future work to improve the filter-verification pipeline further. Our tightest upper bound was obtained via local search using a straightforward approach. More sophisticated techniques have been proposed recently [6] and can be incorporated to reduce verification. For the verification step,

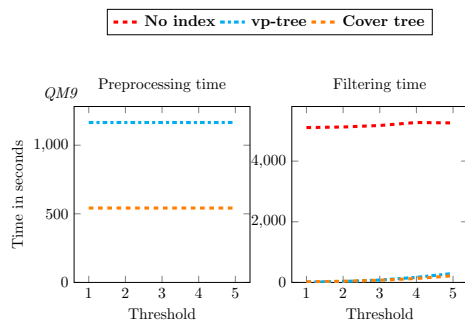


Fig. 6: Runtime comparison for preprocessing and filtering 100 range queries in the dataset *QM9* using BRANCH and thresholds 1 to 5.

tailored methods that benefit from the already obtained assignment or the upper and lower bound can be developed. A well-known phenomenon of metric trees is that their effectivity decreases with increasing intrinsic dimensionality of the data/distance. Therefore, a suitable lower bound should not only be efficiently computed and tight, but ideally also have a low intrinsic dimensionality. Studying this property for the available lower bounds remains future work. Finally, recent advances in median graph computation [3] suggest to compute routing objects instead of using database graphs. An experimental comparison to such orthogonal approaches remains future work.

References

1. Beygelzimer, A., Kakade, S.M., Langford, J.: Cover trees for nearest neighbor. In: Int. Conf. Machine Learning, ICML. vol. 148, pp. 97–104 (2006)
2. Blumenthal, D.B., Gamper, J.: Improved lower bounds for graph edit distance. *IEEE Transactions on Knowledge and Data Engineering* **30**(3), 503–516 (2018)
3. Blumenthal, D.B., Boria, N., Bougleux, S., Brun, L., Gamper, J., Gaüzère, B.: Scalable generalized median graph estimation and its manifold use in bioinformatics, clustering, classification, and indexing. *Inf. Syst.* **100**, 101766 (2021)
4. Blumenthal, D.B., Boria, N., Gamper, J., Bougleux, S., Brun, L.: Comparing heuristics for graph edit distance computation. *VLDB J.* **29**(1), 419–458 (2020)
5. Blumenthal, D.B., Gamper, J.: On the exact computation of the graph edit distance. *Pattern Recognit. Lett.* **134**, 46–57 (2020)
6. Boria, N., Blumenthal, D.B., Bougleux, S., Brun, L.: Improved local search for graph edit distance. *Pattern Recognit. Lett.* **129**, 19–25 (2020)
7. Bougleux, S., Gaüzère, B., Blumenthal, D.B., Brun, L.: Fast linear sum assignment with error-correction and no cost constraints. *Pattern Recognit. Lett.* **134**, 37–45 (2020)
8. Burkard, R.E., Dell’Amico, M., Martello, S.: *Assignment Problems*. SIAM (2012)
9. Chang, L., Feng, X., Lin, X., Qin, L., Zhang, W., Ouyang, D.: Speeding up GED verification for graph similarity search. In: Int. Conf. Data Engineering, ICDE. pp. 793–804 (2020)

10. Chen, X., Huo, H., Huan, J., Vitter, J.S.: An efficient algorithm for graph edit distance computation. *Knowledge-Based Systems* **163**, 762–775 (2019)
11. Gouda, K., Arafa, M., Calders, T.: BFST_ED: A novel upper bound computation framework for the graph edit distance. In: *SISAP*. pp. 3–19 (2016)
12. Gouda, K., Hassaan, M.: CSLGED: An efficient approach for graph edit similarity computation. In: *Int. Conf. Data Engineering, ICDE*. pp. 265–276 (2016)
13. Kim, J., Choi, D., Li, C.: Inves: Incremental partitioning-based verification for graph similarity search. In: *Extending Database Technology, EDBT*. pp. 229–240 (2019)
14. Kriege, N.M., Giscard, P., Bause, F., Wilson, R.C.: Computing optimal assignments in linear time for approximate graph matching. In: *ICDM*. pp. 349–358 (2019)
15. Lerouge, J., Abu-Aisheh, Z., Raveaux, R., Héroux, P., Adam, S.: New binary linear programming formulation to compute the graph edit distance. *Pattern Recognit.* **72**, 254–265 (2017)
16. Liang, Y., Zhao, P.: Similarity search in graph databases: A multi-layered indexing approach. In: *Int. Conf. Data Engineering, ICDE*. pp. 783–794 (2017)
17. Morris, C., Kriege, N.M., Bause, F., Kersting, K., Mutzel, P., Neumann, M.: TU-Dataset: A collection of benchmark datasets for learning with graphs. In: *ICML 2020 Workshop on Graph Representation Learning and Beyond, GRL+* (2020)
18. Neuhaus, M., Riesen, K., Bunke, H.: Fast suboptimal algorithms for the computation of graph edit distance. In: *Structural, Syntactic, and Statistical Pattern Recognition*. pp. 163–172 (08 2006)
19. Qin, Z., Bai, Y., Sun, Y.: Ghashing: Semantic graph hashing for approximate similarity search in graph databases. In: *ACM SIGKDD*. pp. 2062–2072 (2020)
20. Riesen, K., Bunke, H.: Approximate graph edit distance computation by means of bipartite graph matching. *Image Vision Comput.* **27**(7), 950–959 (2009)
21. Schubert, E., Zimek, A.: ELKI: A large open-source library for data analysis - ELKI release 0.7.5 "Heidelberg". *CoRR* **abs/1902.03616** (2019)
22. Seidl, T., Kriegel, H.: Optimal multi-step k-nearest neighbor search. In: *SIGMOD Int. Conf. Management of Data*. pp. 154–165 (1998)
23. Serratos, F., Cortés, X., Solé-Ribalta, A.: Graph database retrieval based on metric-trees. In: *SSPR*. pp. 437–447 (2012)
24. Stauffer, M., Tschachtli, T., Fischer, A., Riesen, K.: A survey on applications of bipartite graph edit distance. In: *GbRPR*. pp. 242–252 (2017)
25. Wang, G., Wang, B., Yang, X., Yu, G.: Efficiently indexing large sparse graphs for similarity search. *IEEE Trans. Knowl. Data Eng.* **24**(3), 440–451 (2012)
26. Wang, X., Ding, X., Tung, A., Ying, S., Jin, H.: An efficient graph indexing method. In: *Int. Conf. Data Engineering, ICDE* (2012)
27. Yianilos, P.N.: Data structures and algorithms for nearest neighbor search in general metric spaces. In: *SODA*. pp. 311–321 (1993)
28. Zeng, Z., Tung, A.K.H., Wang, J., Feng, J., Zhou, L.: Comparing stars: On approximating graph edit distance. *Proc. VLDB Endow.* **2**(1), 25–36 (2009)
29. Zhao, X., Xiao, C., Lin, X., Liu, Q., Zhang, W.: A partition-based approach to structure similarity search. *Proc. VLDB Endow.* **7**(3), 169–180 (2013)
30. Zhao, X., Xiao, C., Lin, X., Wang, W.: Efficient graph similarity joins with edit distance constraints. In: *Int. Conf. Data Engineering, ICDE* (2012)
31. Zheng, W., Zou, L., Lian, X., Wang, D., Zhao, D.: Efficient graph similarity search over large graph databases. *IEEE Trans. Knowl. Data Eng.* **27**(4), 964–978 (2015)