
Sample Complexity of Learning Heuristic Functions for Greedy-Best-First and A* Search

Shinsaku Sakaue

The University of Tokyo
Tokyo, Japan

sakaue@mist.i.u-tokyo.ac.jp

Taihei Oki

The University of Tokyo
Tokyo, Japan

oki@mist.i.u-tokyo.ac.jp

Abstract

Greedy best-first search (GBFS) and A* search (A*) are popular algorithms for path-finding on large graphs. Both use so-called heuristic functions, which estimate how close a vertex is to the goal. While heuristic functions have been handcrafted using domain knowledge, recent studies demonstrate that learning heuristic functions from data is effective in many applications. Motivated by this emerging approach, we study the sample complexity of learning heuristic functions for GBFS and A*. We build on a recent framework called *data-driven algorithm design* and evaluate the *pseudo-dimension* of a class of utility functions that measure the performance of parameterized algorithms. Assuming that a vertex set of size n is fixed, we present $O(n \lg n)$ and $O(n^2 \lg n)$ upper bounds on the pseudo-dimensions for GBFS and A*, respectively, parameterized by heuristic function values. The upper bound for A* can be improved to $O(n^2 \lg d)$ if every vertex has a degree of at most d and to $O(n \lg n)$ if edge weights are integers bounded by $\text{poly}(n)$. We also give $\Omega(n)$ lower bounds for GBFS and A*, which imply that our bounds for GBFS and A* under the integer-weight condition are tight up to a $\lg n$ factor. Finally, we discuss a case where the performance of A* is measured by the suboptimality and show that we can sometimes obtain a better guarantee by combining a parameter-dependent worst-case bound with a sample complexity bound.

1 Introduction

Given a graph with a start vertex s , a goal vertex t , and non-negative edge weights, we consider finding an s - t path with a small total weight. The Dijkstra algorithm [16] finds an optimal path by exploring all vertices that are as close to s as t . It, however, is sometimes impractical for large graphs since exploring all such vertices is too costly. Heuristic search algorithms are used to address such situations; among them, greedy best-first search (GBFS) [17] and A* search (A*) [24] are two popular algorithms. Both GBFS and A* use so-called heuristic functions, which estimate how close an input vertex is to t . GBFS/A* attempts to avoid redundant exploration by scoring vertices based on heuristic function values and iteratively expanding vertices with the smallest score. If well-suited heuristic functions are available, GBFS/A* can run much faster than the Dijkstra algorithm. Furthermore, if A* uses an *admissible* heuristic function, i.e., it never overestimates the shortest-path distance to t , it always finds an optimal path [24]. Traditionally, heuristic functions have been made based on domain knowledge; e.g., if graphs are road networks, the Euclidean distance gives an admissible heuristic.

When applying GBFS/A* to various real-world problems, a laborious process is to handcraft heuristic functions. Learning heuristic functions from data can be a promising approach to overcoming the obstacle due to the recent development of technologies for collecting graph data. Researchers have demonstrated the effectiveness of this approach in robotics [11, 32, 28, 36], computational organic chemistry [13], and pedestrian trajectory prediction [36]. With learned heuristic functions, however,

obtaining theoretical guarantees is difficult since we can hardly understand how the search can be guided by such heuristic functions. (A recent paper [1] studies learning of admissible heuristics for A^* , but the optimality is confirmed only empirically.) Moreover, learned heuristic functions may be overfitting to problem instances at hand. That is, even if GBFS/ A^* with learned heuristic functions perform well over training instances, they may deliver poor future performance. In summary, the emerging line of work on search algorithms with learned heuristic functions is awaiting a theoretical foundation for guaranteeing their performance in a data-driven manner. Thus, a natural question is: *how many sampled instances are needed to learn heuristic functions with generalization guarantees on the performance of resulting GBFS/ A^* ?*

1.1 Our contribution

We address the above question, assuming that path-finding instances defined on a fixed vertex set of size n are drawn i.i.d. from an unknown distribution. Our analysis is based on so-called *data-driven algorithm design* [22, 4], a PAC-learning framework for bounding the sample complexity of algorithm configuration. In the analysis, the most crucial step is to evaluate the *pseudo-dimension* of a class of utility functions that measure the performance of parameterized algorithms. We study the case where GBFS/ A^* is parameterized by heuristic function values and make the following contributions:

1. Section 3 gives $O(n \lg n)$ and $O(n^2 \lg n)$ upper bounds on the pseudo-dimensions for GBFS and A^* , respectively. The bound for A^* can be improved to $O(n^2 \lg d)$ if every vertex has an at most d degree and to $O(n \lg n)$ if edge weights are non-negative integers at most $\text{poly}(n)$.
2. Section 4 presents $\Omega(n)$ lower bounds on the pseudo-dimensions for GBFS and A^* . We prove this result by constructing $\Omega(n)$ instances with unweighted graphs. Thus, our bounds for GBFS and A^* under the integer edge-weight condition are tight up to a $\lg n$ factor.
3. Section 5 studies a particular case of bounding the suboptimality of A^* . We show that we can sometimes improve the guarantee obtained in Section 3 by using an alternative $O(n \lg n)$ bound on the pseudo-dimension of a class of parameter-dependent worst-case bounds [34].

An important consequence of the above results is the tightness up to a $\lg n$ factor for GBFS and A^* under the integer-weight assumption. Note that this assumption holds in various realistic situations. For example, the Internet network and state-space graphs of games are unweighted (unit-weight) graphs, and A^* is often applied to path-finding instances on such graphs.

1.2 Related work

Data-driven algorithm design. Gupta and Roughgarden [22] proposed a PAC approach for bounding the sample complexity of algorithm configuration, which is called *data-driven algorithm design* and has been applied to a broad family of algorithms, including greedy, clustering, and sequence alignment algorithms. We refer the reader to a nice survey [4]. A recent line of work [5, 9, 10] has extensively studied the sample complexity of configuring integer-programming methods, e.g., branch-and-bound and branch-and-cut. In [9, 10], upper bounds on the pseudo-dimension for general tree search are presented, which are most closely related to our results. Our upper bounds, which are obtained by using specific properties of GBFS/ A^* , are better than the previous bounds for general tree search, as detailed in Appendix A. Balcan et al. [8] presented a general framework for evaluating the pseudo-dimension. Their idea is to suppose that performance measures form a class of functions of algorithm parameters, called *dual* functions, and characterize its complexity based on how they are piecewise structured. This idea plays a key role in the analysis of [9, 10], and our analysis of the upper bounds are also inspired by their idea. Its application to our setting, however, requires a close look at the behavior of GBFS/ A^* . Balcan et al. [7] showed that approximating dual functions with simpler ones is useful for improving sample complexity bounds, which is similar to our idea in Section 5. A difference is that while they construct simpler functions with a dynamic programming algorithm, we can use a known worst-case bound on the suboptimality of best-first search [34]. Lower bounds on the pseudo-dimension for graph-search algorithms have not been well studied.

Heuristic search with learning. Eden et al. [18] theoretically studied how the average-case running time of A^* can be affected by the dimensions or bits of learned embeddings or labels of vertex features, based on which heuristic function values and computed. The sample complexity of learning heuristic functions, however, has not been studied.

2 Preliminaries

We present the background on learning theory and our problem setting. In what follows, we let $\mathbb{I}(\cdot)$ be a boolean function that returns 1 if its argument is true and 0 otherwise. We use $\mathcal{H} \subseteq \mathcal{R}^{\mathcal{Y}}$ to denote a class of functions that map \mathcal{Y} to $\mathcal{R} \subseteq \mathbb{R}$. For any positive integer m , we let $[m] = \{1, \dots, m\}$.

2.1 Background on learning theory

The following *pseudo-dimension* [29] is a fundamental notion for quantifying the complexity of a class of real-valued functions.

Definition 1. Let $\mathcal{H} \subseteq \mathbb{R}^{\mathcal{Y}}$ be a class of functions that map some domain \mathcal{Y} to \mathbb{R} . We say a set $\{y_1, \dots, y_N\} \subseteq \mathcal{Y}$ is *shattered* by \mathcal{H} if there exist target values, $t_1, \dots, t_N \in \mathbb{R}$, such that

$$|\{(\mathbb{I}(h(y_1) \geq t_1), \dots, \mathbb{I}(h(y_N) \geq t_N)) \mid h \in \mathcal{H}\}| = 2^N.$$

The *pseudo-dimension* of \mathcal{H} , denoted by $\text{Pdim}(\mathcal{H})$, is the size of a largest set shattered by \mathcal{H} .

If \mathcal{H} is a set of binary-valued functions that map \mathcal{Y} to $\{0, 1\}$, the pseudo-dimension of \mathcal{H} coincides with the so-called *VC-dimension* [35], which is denoted by $\text{VCdim}(\mathcal{H})$.

The following proposition enables us to obtain sample complexity bounds by evaluating the pseudo-dimension (see, e.g., [2, Theorem 19.2] and [27, Theorem 11.8]).

Proposition 1. Let $H > 0$, $\mathcal{H} \subseteq [0, H]^{\mathcal{Y}}$, and \mathcal{D} be a distribution over \mathcal{Y} . For any $\delta \in (0, 1)$, with a probability of at least $1 - \delta$ over the i.i.d. draw of $\{y_1, \dots, y_N\} \sim \mathcal{D}^N$, for all $h \in \mathcal{H}$, it holds that

$$\left| \frac{1}{N} \sum_{i=1}^N h(y_i) - \mathbb{E}_{y \sim \mathcal{D}} [h(y)] \right| = O \left(H \sqrt{\frac{\text{Pdim}(\mathcal{H}) \lg \frac{N}{\text{Pdim}(\mathcal{H})} + \lg \frac{1}{\delta}}{N}} \right).$$

In other words, for any $\epsilon > 0$, $N = \Omega \left(\frac{H^2}{\epsilon^2} (\text{Pdim}(\mathcal{H}) \lg \frac{H}{\epsilon} + \lg \frac{1}{\delta}) \right)$ sampled instances are sufficient to ensure that with a probability of at least $1 - \delta$, for all $h \in \mathcal{H}$, the difference between the empirical average and the expectation over an unknown distribution \mathcal{D} is at most ϵ .

2.2 Problem formulation

We describe path-finding instances, GBFS/A* algorithm, and performance measures considered in this paper.

Path-finding instances. We consider solving randomly generated path-finding instances repetitively. Let $x = (V, E, \{w_e\}_{e \in E}, s, t)$ be a path-finding instance, where (V, E) is a simple directed graph with n vertices, $\{w_e\}_{e \in E}$ is a set of non-negative edge weights (sometimes called costs), $s \in V$ is a start vertex, and $t \in V$ is a goal vertex. We let Π be a class of possible instances. Each instance $x \in \Pi$ is drawn from an unknown distribution \mathcal{D} over Π . We impose the following assumption on Π .

Assumption 1. For all $x \in \Pi$, the vertex set V and the goal node t are identical, and there always exists at least one directed path from $s \neq t$ to t , i.e., every instance $x \in \Pi$ is feasible.

Fixing V is necessary for evaluating the pseudo-dimension in terms of $n = |V|$. Note that we can deal with the case where some instances in Π are defined on vertex subsets $V' \subseteq V$ by removing edges adjacent to $V \setminus V'$. The feasibility assumption is needed to ensure that GBFS/A* always returns a solution, and $s \neq t$ simply rules out the trivial case where the empty set is optimal. In Appendix B, we discuss how to extend our results to the case where t can change depending on instances.

Algorithm description. We sketch algorithmic procedures that are common to both GBFS and A* (see Algorithms 1 and 2 for details, respectively). Let A_ρ be a GBFS/A* algorithm, which is parameterized by heuristic function values $\rho \in \mathbb{R}^n$. Given an instance $x \in \Pi$, A_ρ starts from s and iteratively builds a set of candidate paths. These paths are maintained by OPEN and CLOSED lists, together with pointers $p(\cdot)$ to parent vertices. The OPEN list contains vertices to be explored, and the CLOSED list consists of vertices that have been explored. In each iteration, we select a vertex v from OPEN, expand v , and move v from OPEN to CLOSED.

Heuristic function values ρ are used when selecting vertices. For each $v \in V$, the corresponding entry in ρ , denoted by ρ_v , represents an estimated shortest-path distance from v to t . (Although heuristic function values are usually denoted by $h(v)$, we here use ρ_v for convenience.) In each iteration, we select a vertex with the smallest *score*, which is defined based on ρ as detailed later. We impose the following assumption on the vertex selection step.

Assumption 2. *Define an arbitrary strict total order on V ; for example, we label elements in V by v_1, \dots, v_n and define a total order $v_1 < \dots < v_n$. When selecting a vertex with the smallest score, we break ties, if any, in favor of the smallest vertex with respect to the total order.*

If we allow A_ρ to break ties arbitrarily, its behavior becomes too complex to obtain meaningful bounds on the pseudo-dimension. Assumption 2 is a natural rule to exclude such troublesome cases.

Performance measure. Let A_ρ be GBFS/A* with parameters $\rho \in \mathbb{R}^n$. We measure performance of A_ρ on $x \in \Pi$ with a utility function u . We assume u to satisfy the following condition.

Assumption 3. *Let $H > 0$. A utility function u takes x and a series of all OPEN, CLOSED, and $\text{p}(\cdot)$ generated during the execution of A_ρ on $x \in \Pi$ as input, and returns a scalar value in $[0, H]$.*

We sometimes use A_ρ to represent the series of OPEN and CLOSED lists and pointers generated by A_ρ . Note that u meeting Assumption 3 can measure various kinds of performance. For example, since the pointers indicate an s - t path returned by A_ρ , u can represent its cost. Moreover, since the series of OPEN and CLOSED lists maintain all search states, u can represent the time and space complexity of A_ρ . We let $u_\rho : \Pi \rightarrow [0, H]$ denote the utility function that returns the performance of A_ρ on any $x \in \Pi$, and define a class of such functions as $\mathcal{U} = \{u_\rho : \Pi \rightarrow [0, H] \mid \rho \in \mathbb{R}^n\}$. The upper bound, H , is necessary to obtain sample complexity bounds with Proposition 1. Setting such an upper bound is usual in practice. For example, if u measures the running time, H represents a time-out deadline.

Generalization guarantees on performance. Given the above setting, we want to learn $\hat{\rho}$ values that attain an optimal $\mathbb{E}_{x \sim \mathcal{D}}[u_{\hat{\rho}}(x)]$ value, where available information consists of sampled instances x_1, \dots, x_N and $u_\rho(x_1), \dots, u_\rho(x_N)$ values for any $\rho \in \mathbb{R}^n$. To obtain generalization guarantees on the performance of $A_{\hat{\rho}}$, we bound $|\frac{1}{N} \sum_{i=1}^N u_\rho(x_i) - \mathbb{E}_{x \sim \mathcal{D}}[u_\rho(x)]|$ uniformly for all $\rho \in \mathbb{R}^n$. Note that the uniform bound offers performance guarantees that are independent of learning procedures, e.g., manual or automated (without being uniform, learned $\hat{\rho}$ may be overfitting sampled instances). As in Proposition 1, to bound the sample complexity of learning ρ values, we need to evaluate the pseudo-dimension of \mathcal{U} , denoted by $\text{Pdim}(\mathcal{U})$, which is the main subject of this study.

Remarks on heuristic functions. While we allow heuristic function values ρ to be any point in \mathbb{R}^n , the range of heuristic functions may be restricted to some subspace of \mathbb{R}^n . Note that our upper bounds are applicable to such situations since restricting the space of possible ρ values does not increase $\text{Pdim}(\mathcal{U})$. Meanwhile, such restriction may be useful for improving the upper bounds on $\text{Pdim}(\mathcal{U})$; exploring this direction is left for future work. Also, our setting cannot deal with heuristic functions that take some instance-dependent features as input. To study such cases, we need more analysis that is specific to heuristic function models, which goes beyond the scope of this paper. Thus, we leave this for future work. Note that our setting still includes important heuristic function models on fixed vertex sets. For example, we can set ρ using learned distances to landmarks [21], or we can let ρ be distances measured on some metric space by learning metric embeddings of vertices [37].

3 Upper bounds on the pseudo-dimension

We present details of GBFS and A* and upper bounds on the pseudo-dimensions of \mathcal{U} . In this section, we suppose that vertices in V are labeled by v_1, \dots, v_n as in Assumption 2.

3.1 Greedy best-first search

Algorithm 1 shows the details of GBFS A_ρ with heuristic function values $\rho \in \mathbb{R}^n$. When selecting vertices in Step 3, the scores are determined only by ρ . This implies an obvious but important fact.

Lemma 1. *Let $\rho, \rho' \in \mathbb{R}^n$ be a pair of heuristic function values with an identical total order up to ties on their entries, i.e., $\mathbb{I}(\rho_{v_i} \leq \rho_{v_j}) = \mathbb{I}(\rho'_{v_i} \leq \rho'_{v_j})$ for all $i, j \in [n]$ such that $i < j$. Then, we have $u_\rho(x) = u_{\rho'}(x)$ for all $x \in \Pi$.*

Algorithm 1 GBFS with heuristic function values ρ

```
1: OPEN = {s}, CLOSED = ∅, and p(s) = None.
2: while OPEN is not empty :
3:   v ← argmin{ρv' | v' ∈ OPEN}.           ▷ Break ties as in Assumption 2.
4:   for each child c of v :
5:     if c = t :
6:       return s–t path by tracing pointers p(·), where p(t) = v.
7:     if c ∉ OPEN ∪ CLOSED :
8:       p(c) ← v and OPEN ← OPEN ∪ {c}.
9:   Move v from OPEN to CLOSED.
```

Algorithm 2 A* with heuristic function values ρ

```
1: OPEN = {s}, CLOSED = ∅, p(s) = None, and gs = 0.
2: while OPEN is not empty :
3:   v ← argmin{gv' + ρv' | v' ∈ OPEN}.           ▷ Break ties as in Assumption 2.
4:   if v = t :
5:     return s–t path by tracing pointers p(·).
6:   for each child c of v :
7:     gnew ← gv + w(v,c).
8:     if c ∉ OPEN ∪ CLOSED :
9:       gc ← gnew, p(c) ← v, and OPEN ← OPEN ∪ {c}.
10:    else if c ∈ OPEN and gnew < gc :
11:      gc ← gnew and p(c) ← v.
12:    else if c ∈ CLOSED and gnew < gc :           ▷ Steps 12–14 are for reopening.
13:      gc ← gnew and p(c) ← v.
14:      Move c from CLOSED to OPEN.
15:   Move v from OPEN to CLOSED.
```

Proof. For any $x \in \Pi$, if ρ and ρ' have an identical strict total order on their entries, vertices selected in Step 3 are the same in each iteration of A_ρ and $A_{\rho'}$. Since this is the only step ρ and ρ' can affect, we have $A_\rho = A_{\rho'}$ for all $x \in \Pi$, hence $u_\rho(x) = u_{\rho'}(x)$. Moreover, this holds even if ρ and/or ρ' have ties on their entries because of Assumption 2. That is, the total order uniquely determines a vertex selected in Step 3 even in case of ties. Therefore, the statement holds. \square

From Lemma 1, the behavior of GBFS is uniquely determined once a total order on $\{\rho_v\}_{v \in V}$ is fixed. Thus, for any $x \in \Pi$, the number of distinct $u_\rho(x)$ values is at most $n!$, the number of permutations of $\{\rho_v\}_{v \in V}$. This fact enables us to obtain an $O(n \lg n)$ upper bound on the pseudo-dimension of \mathcal{U} .

Theorem 1. For GBFS A_ρ with parameters $\rho \in \mathbb{R}^n$, it holds that $\text{Pdim}(\mathcal{U}) = O(n \lg n)$.

Proof. Lemma 1 implies that we can partition \mathbb{R}^n into $n!$ regions, $\mathcal{P}_1, \mathcal{P}_2, \dots$, so that for every \mathcal{P}_i , any pair of $\rho, \rho' \in \mathcal{P}_i$ satisfies $u_\rho(x) = u_{\rho'}(x)$ for all $x \in \Pi$. Note that the construction of the regions, $\mathcal{P}_1, \mathcal{P}_2, \dots$, does not depend on x . Thus, given any N instances x_1, \dots, x_N , even if ρ moves over whole \mathbb{R}^n , the number of distinct tuples of form $(u_\rho(x_1), \dots, u_\rho(x_N))$ is at most $n!$. To shatter N instances, $n! \geq 2^N$ must hold. Solving this for the largest N yields $\text{Pdim}(\mathcal{U}) = O(n \lg n)$. \square

3.2 A* search

Algorithm 2 is the details of A*. As with GBFS, ρ only affects the vertex selection step (Step 3). However, unlike GBFS, the scores, $g_v + \rho_v$, involve not only ρ but also $\{g_v\}_{v \in V}$. Each g_v is called a g -cost and maintains a cost of some path from s to v . As in Algorithm 2, when v is expanded and a shorter path to c is found, whose cost is denoted by g_{new} , we update the g_c value. Thus, each g_v always gives an upper bound on the shortest-path distance from s to v . For each $v \in V$, there are at most $\sum_{k=0}^{n-2} k! \leq (n-1)!$ simple paths connecting s to v , and thus g_v can take at most $(n-1)!$ distinct values. We denote the set of those distinct values by \mathcal{G}_v , and define $\mathcal{G}_V = \{(v, g_v) \mid v \in V, g_v \in \mathcal{G}_v\}$ as the set of all pairs of a vertex and its possible g -cost. It holds that $|\mathcal{G}_V| \leq n \times (n-1)! = n!$.

Note that once $x \in \Pi$ is fixed, \mathcal{G}_v for $v \in V$ and \mathcal{G}_V are uniquely determined. To emphasize this fact, we sometimes use notation with references to x : $g_v(x)$, $\mathcal{G}_v(x)$, and $\mathcal{G}_V(x)$. As with the case of GBFS (Lemma 1), we can define a total order on the scores to determine the behavior of A^* uniquely.

Lemma 2. *Fix any instance $x \in \Pi$. Let $\rho, \rho' \in \mathbb{R}^n$ be a pair of heuristic function values such that total orders on the sets of all possible scores, $\{g_v(x) + \rho_v \mid (v, g_v(x)) \in \mathcal{G}_V(x)\}$ and $\{g_v(x) + \rho'_v \mid (v, g_v(x)) \in \mathcal{G}_V(x)\}$, are identical up to ties. Then, it holds that $u_\rho(x) = u_{\rho'}(x)$.*

Proof. If the two sets of scores have an identical strict total order, we select the same vertex in Step 3 in each iteration of A_ρ and $A_{\rho'}$. Thus, we have $A_\rho = A_{\rho'}$ for any fixed x , implying $u_\rho(x) = u_{\rho'}(x)$. We show that this holds even in the presence of ties by using Assumption 2. First, any two scores of the same vertices, $g_v(x) + \rho_v$ and $g'_v(x) + \rho_v$, never have ties since \mathcal{G}_v consists of distinct g -costs. Next, if $g_{v_i}(x) + \rho_{v_i} = g_{v_j}(x) + \rho_{v_j}$ holds for some $i < j$, we always prefer v_i to v_j in Step 3 due to Assumption 2. Therefore, even in the presence of ties, we select a vertex in Step 3 as if the set of scores has a strict total order. Thus, if ρ and ρ' induce the same total order up to ties on the sets of possible scores, it holds that $u_\rho(x) = u_{\rho'}(x)$. \square

By using Lemma 2, we can obtain an $O(n^2 \lg n)$ upper bound on the pseudo-dimension of \mathcal{U} .

Theorem 2. *For A^* A_ρ with parameters $\rho \in \mathbb{R}^n$, it holds that $\text{Pdim}(\mathcal{U}) = O(n^2 \lg n)$.*

Proof. As with the proof of Theorem 1, we partition \mathbb{R}^n into some regions so that in each region, the behavior of A^* is unique. Unlike the case of GBFS, boundaries of such regions change over N instances. To deal with this situation, we use a geometric fact: for $m \geq n \geq 1$, m hyperplanes partition \mathbb{R}^n into $O((em)^n)$ regions.¹

Fix a tuple of any N instances (x_1, \dots, x_N) . We consider hyperplanes in \mathbb{R}^n of form $g_{v_i}(x_k) + \rho_{v_i} = g_{v_j}(x_k) + \rho_{v_j}$ for all $k \in [N]$ and all pairs of $(v_i, g_{v_i}(x_k)), (v_j, g_{v_j}(x_k)) \in \mathcal{G}_V$ such that $i \neq j$. These hyperplanes partition \mathbb{R}^n into some regions, $\mathcal{P}_1, \mathcal{P}_2, \dots$, so that the following condition holds: for every \mathcal{P}_i , any $\rho, \rho' \in \mathcal{P}_i$ have the same total order on $\{g_v(x_k) + \rho_v \mid (v, g_v(x_k)) \in \mathcal{G}_V(x_k)\}$ and $\{g_v(x_k) + \rho'_v \mid (v, g_v(x_k)) \in \mathcal{G}_V(x_k)\}$ up to ties for all $k \in [N]$, which implies $u_\rho(x_k) = u_{\rho'}(x_k)$ for all $k \in [N]$ due to Lemma 2. That is, for every $k \in [N]$, if we see $u_\rho(x_k)$ as a function of ρ , it is piecewise constant where pieces are given by $\mathcal{P}_1, \mathcal{P}_2, \dots$. Therefore, when ρ moves over whole \mathbb{R}^n , the number of distinct tuples of form $(u_\rho(x_1), \dots, u_\rho(x_N))$ is at most the number of the pieces. Note that the pieces are generated by partitioning \mathbb{R}^n with $\sum_{k \in [N]} \binom{|\mathcal{G}_V(x_k)|}{2} \leq N \binom{n!}{2}$ hyperplanes, which means there are at most $O\left(\left(eN \binom{n!}{2}\right)^n\right)$ pieces. To shatter N instances, $O\left(\left(eN \binom{n!}{2}\right)^n\right) \geq 2^N$ is necessary. Solving this for the largest N yields $\text{Pdim}(\mathcal{U}) = O(n^2 \lg n)$. \square

Compared with GBFS, the additional n factor comes from the bound of $(n-1)!$ on $|\mathcal{G}_v|$. This bound may seem too pessimistic, but it is almost tight in some cases, as implied by the following example.

Example 1. Let (V, E) be a complete graph with edges labeled as $\{e_1, \dots, e_{|E|}\}$. Set each edge weight w_{e_i} to 2^{i-1} for $i \in [|E|]$. Considering the binary representation of the edge weights, the costs of all simple s - v paths are mutually different for $v \in V$, which implies $|\mathcal{G}_v| = \sum_{k=0}^{n-2} k! \geq (n-2)!$.

This example suggests that improving the $O(n^2 \lg n)$ bound is not straightforward. Under some realistic assumptions, however, we can improve it by deriving smaller upper bounds on $|\mathcal{G}_v|$.

First, if the maximum degree of vertices is always bounded, we can obtain the following bound.

Theorem 3. *Assume that the maximum out-degrees of directed graphs (V, E) of all instances in Π are upper bounded by d . Then, it holds that $\text{Pdim}(\mathcal{U}) = O(n^2 \lg d)$.*

Proof. Under the assumption on the maximum degree, there are at most $\sum_{k=0}^{n-2} d^k \leq (n-1)d^{n-2}$ simple s - v paths, which implies $|\mathcal{G}_v| \leq (n-1)d^{n-2}$ for every $v \in V$. Therefore, we have $|\mathcal{G}_V| \leq n \times (n-1)d^{n-2}$. Following the proof of Theorem 2, we can obtain an upper bound on $\text{Pdim}(\mathcal{U})$ by solving $O\left(N^n \binom{n(n-1)d^{n-2}}{2}\right) \geq 2^N$ for the largest N , which yields $\text{Pdim}(\mathcal{U}) = O(n^2 \lg d)$. \square

¹Even if some regions degenerate, from [23, Theorem 28.1.1] and [12, Proposition A2.1], the number of all d -dimensional regions for $d = 0, \dots, n$ is $\sum_{d=0}^n \sum_{i=0}^d \binom{n-i}{d-i} \binom{m}{n-i} \leq 2(em)^n$. The fact has a close connection to Sauer's lemma [30] (see [20]). In this sense, our analysis is in a similar spirit to the general framework of [8].




Figure 1: An illustration of the instances x_1, \dots, x_{n-4} for $n = 8$. Each vertex is labeled by s, r, t , or $i \in [n-3]$, as shown nearby the vertex circles. The values in vertex circles represent ρ that makes A_ρ return suboptimal paths to x_2 and x_3 , i.e., $S = \{2, 3\}$. The thick edges indicate returned paths.

Second, if edge weights are non-negative integers bounded by W , we can obtain the following bound.

Theorem 4. *Assume that edge weights $\{w_e\}_{e \in E}$ of all instances in Π are non-negative integers bounded by a constant W from above. Then, it holds that $\text{Pdim}(\mathcal{U}) = O(n \lg(nW))$.*

Proof. Under the assumption, every g -cost g_v takes a non-negative integer value at most nW . Since \mathcal{G}_v consists of distinct g -cost values, $|\mathcal{G}_v| \leq nW$ holds, hence $|\mathcal{G}_V| \leq n^2W$. Solving $O\left(N^n \binom{n^2W}{2}^n\right) \geq 2^N$ for the largest N , we obtain $\text{Pdim}(\mathcal{U}) = O(n \lg(nW))$. \square

Note that if $W = O(\text{poly}(n))$ holds, we have $\text{Pdim}(\mathcal{U}) = O(n \lg n)$.

On reopening. A^* is usually allowed to reopen closed vertices as in Steps 12–14. This, however, causes $\Omega(2^n)$ iterations in general [26], albeit always finite [33]. A popular workaround is to simply remove Steps 12–14, and such A^* without reopening has also been extensively studied [34, 31, 14, 15]. Note that our results are applicable to A^* both with and without reopening.

4 Lower bounds on the pseudo-dimension

We present lower bounds on the pseudo-dimension for GBFS/ A^* . We prove the result by constructing $\Omega(n)$ shatterable instances with unweighted graphs. Therefore, the $O(n \lg n)$ upper bounds for GBFS (Theorem 1) and A^* under the edge-weight assumption (Theorem 4) are tight up to a $\lg n$ factor.

Theorem 5. *For GBFS/ A^* A_ρ with parameters $\rho \in \mathbb{R}^n$, it holds that $\text{Pdim}(\mathcal{U}) = \Omega(n)$.*

Proof. We construct a series of $n-4$ instances, x_1, \dots, x_{n-4} , that can be shattered by \mathcal{U} , where each u_ρ returns the length of an s - t path found by A_ρ . We label vertices in V by s, r, t , or $i \in [n-3]$. See Figure 1 for an example with $n = 8$. We define $M = V \setminus \{s, r, t\}$. For each x_i ($i \in [n-4]$), we draw edges (s, v) for $v \in M$ and (v, t) for $v \in \{v' \in M \mid v' > i\}$, which constitute optimal s - t paths of length 2. In addition, for each x_i , we draw edges (i, r) and (r, t) , where $s \rightarrow i \rightarrow r \rightarrow t$ is the only suboptimal path of length 3. Letting $t_i = 2.5$ for $i \in [n-4]$, we prove that \mathcal{U} can shatter those $n-4$ instances, i.e., A_ρ can return suboptimal solutions to any subset of $\{x_1, \dots, x_{n-4}\}$ by appropriately setting ρ .

Let $S \subseteq [n-4]$ indicate a subset of instances, to which we will make A_ρ return suboptimal solutions. We show that for any S , we can set ρ so that A_ρ returns $s \rightarrow i \rightarrow r \rightarrow t$ to x_i if and only if $i \in S$. We refer to the vertex labeled by $n-3$ as m , which we use to ensure that every instance has an optimal path $s \rightarrow m \rightarrow t$. We set ρ as follows: $\rho_s = n$ (or an arbitrary value), $\rho_r = \rho_t = 0$, $\rho_i = i + 2$ if $i \in S \cup \{m\}$, and $\rho_i = n$ (or a sufficiently large value) if $i \in [n-4] \setminus S$. If A_ρ with this ρ is applied to x_i , it iteratively selects vertices in $S \cup \{m\}$ in increasing order of their labels until a vertex that has a child is selected. Once a vertex with a child is expanded, it ends up returning $s \rightarrow i \rightarrow r \rightarrow t$ if $i \in S$ and $s \rightarrow v \rightarrow t$ for some $v > i$ if $i \notin S$. We below confirm this more precisely, separately for GBFS and A^* .

GBFS. We consider applying GBFS A_ρ to x_i . A_ρ first expands s and add vertices in M to OPEN. Since vertices in $v \in [n-4] \setminus S$ have sufficiently large scores of n , they are never selected before any vertex in $S \cup \{m\}$. Thus, A_ρ selects a vertex from $S \cup \{m\}$ with the smallest score. If the selected vertex, denoted by v , satisfies $v < i$, there is no child of v ; hence, nothing is added to OPEN, and we go back to Step 3. In this way, A_ρ iteratively moves $v \in S \cup \{m\}$ that has no child from OPEN to CLOSED. Consider the first time when the selected vertex $v \in S \cup \{m\}$ has a child c (this situation is guaranteed to occur since m always has a child). If $i \notin S$, we have $v \neq i$ since v is selected from $S \cup \{m\}$. Then, since v 's child is $c = t$, A_ρ returns $s \rightarrow v \rightarrow t$ with $v \neq i$. If $i \in S$, then i has the smallest score ($\rho_i = i + 2$) among all vertices in $S \cup \{m\}$ that have a child. Thus, A_ρ selects i and opens r . Now, r has the smallest score of $\rho_r = 0$. Therefore, A_ρ selects r and reaches t , returning $s \rightarrow i \rightarrow r \rightarrow t$. Consequently, A_ρ returns the suboptimal path if and only if $i \in S$.

A*. It first expands s and add M to OPEN. Since $g_v = 1$ for all $v \in M$, only ρ values matter when comparing the scores, as with the case of GBFS. Therefore, A* iterates to move vertices in $S \cup \{m\}$ from OPEN to CLOSED until a vertex that has a child is selected. Consider the first time a selected vertex v has a child c (so far, s is not reopened since $g_s = 0$). As with the case of GBFS, we have $v \neq i$ and $c = t$ if $i \notin S$, or $v = i$ and $c = r$ if $i \in S$. Now, every $v' \in \text{OPEN} \setminus \{c\}$ has a score of at least 4 since $g_{v'} = 1$ and $\rho_{v'} \geq 3$ for $v' \in M$. Therefore, if $i \notin S$, $t \in \text{OPEN}$ has the smallest score of $g_t + \rho_t = 2 + 0 = 2$. Thus, A_ρ next selects t and returns $s \rightarrow v \rightarrow t$, where $v \neq i$. If $i \in S$, since $r \in \text{OPEN}$ has the smallest score of $g_r + \rho_r = 2 + 0 = 2$, A_ρ selects r and opens t . Then, since t has the score of $g_t + \rho_t = 3 + 0 = 3$, A_ρ selects t and returns $s \rightarrow i \rightarrow r \rightarrow t$. To conclude, A_ρ returns the suboptimal path if and only if $i \in S$. \square

5 Toward better guarantees on the suboptimality of A*

Given the results in Sections 3 and 4, a major open problem is to close the $\tilde{O}(n)$ gap² between the $O(n^2 \lg n)$ upper bound and the $\Omega(n)$ lower bound for A* in general cases. This problem seems very complicated, as we will discuss in Section 6. Instead, we here study a particular case where we want to bound the expected suboptimality of A*, which is an important performance measure since learned heuristic values are not always admissible. We show that a general bound obtained from Theorem 2 can be sometimes improved by using a ρ -dependent worst-case bound [34].

For any $x \in \Pi$, let $\text{Opt}(x)$ and $\text{Cost}_\rho(x)$ be the costs of an optimal solution and an s - t path returned by A_ρ , respectively, and let $u_\rho(x) = \text{Cost}_\rho(x) - \text{Opt}(x)$ be the suboptimality. From Theorem 2 and Proposition 1, we can obtain the following high-probability bound on the expected suboptimality:

$$\mathbb{E}_{x \sim \mathcal{D}} [\text{Cost}_\rho(x) - \text{Opt}(x)] \leq \frac{1}{N} \sum_{i=1}^N (\text{Cost}_\rho(x) - \text{Opt}(x)) + \tilde{O} \left(H \sqrt{\frac{n^2 + \lg \frac{1}{\delta}}{N}} \right). \quad (1)$$

That is, the expected suboptimality can be bounded from above by the empirical suboptimality over the N training instances (an empirical term) plus an $\tilde{O}(H \sqrt{n^2/N})$ term (a complexity term). While this bound is useful when $N \gg n^2$, we may not have enough training instances in practice. In such cases, the complexity term becomes dominant and prevents us from obtaining meaningful guarantees. In what follows, we present an alternative bound of the form ‘‘an empirical term + a complexity term’’ that can strike a better balance between the two terms when N is not large enough relative to n^2 .

To this end, we use the notion of *consistency*. We say ρ is *consistent* if $\rho_v \leq \rho_c + w_{(v,c)}$ holds for all $(v, c) \in E$. If ρ is consistent, A* without reopening returns an optimal solution. Valenzano et al. [34, Theorem 4.6] revealed that for any instance $x \in \Pi$, the suboptimality of A* can be bounded by the inconsistency accumulated along an optimal path (excluding the first edge containing s) as follows:³

$$\text{Cost}_\rho(x) - \text{Opt}(x) \leq \Delta_\rho(x) := \sum_{(v,c) \in S^*(x), v \neq s} \max\{\rho_v - \rho_c - w_{(v,c)}, 0\}, \quad (2)$$

where $S^*(x) \subseteq E$ is an optimal solution to x (if there are multiple optimal solutions, we break ties by using the lexicographical order induced from the total order defined in Assumption 2). We call $\Delta_\rho(x)$ the inconsistency (of ρ on $S^*(x)$).

²We use \tilde{O} and $\tilde{\Omega}$ to hide logarithmic factors of n and N for simplicity.

³The original theorem in [34] is applicable only to the case where A* does not reopen vertices and $\rho_t = 0$. These restrictions are unnecessary as detailed in Appendix C, and thus our result holds regardless of reopening.

Given N instances x_1, \dots, x_N , we can compute the empirical inconsistency, $\frac{1}{N} \sum_{i=1}^N \Delta_\rho(x_i)$, at the cost of solving the N instances, which we will use as an empirical term. To define the corresponding complexity term, we regard $\Delta_\rho(\cdot) : \Pi \rightarrow [0, \hat{H}]$ as an inconsistency function parameterized by ρ , where we will discuss how large $\hat{H} > 0$ can be later, and we let $\hat{\mathcal{U}} = \{\Delta_\rho : \Pi \rightarrow [0, \hat{H}] \mid \rho \in \mathbb{R}^n\}$. The following theorem says that the class $\hat{\mathcal{U}}$ of inconsistency functions has a smaller pseudo-dimension than the class \mathcal{U} of general utility functions.

Theorem 6. *For the class $\hat{\mathcal{U}}$ of inconsistency functions, it holds that $\text{Pdim}(\hat{\mathcal{U}}) = O(n \lg n)$.*

By using (2), Proposition 1, and Theorem 6, we can obtain the following high-probability bound on the expected suboptimality, whose complexity term has a better dependence on n than that of (1):

$$\mathbb{E}_{x \sim \mathcal{D}} [\text{Cost}_\rho(x) - \text{Opt}(x)] \leq \mathbb{E}_{x \sim \mathcal{D}} [\Delta_\rho(x)] \leq \frac{1}{N} \sum_{i=1}^N \Delta_\rho(x_i) + \tilde{O} \left(\hat{H} \sqrt{\frac{n + \lg \frac{1}{\delta}}{N}} \right).$$

This bound is uniform for all $\rho \in \mathbb{R}^n$, as with other bounds discussed so far. Thus, the bound holds even if we choose ρ to minimize the empirical inconsistency. Note that the empirical inconsistency is convex in ρ since $\Delta_\rho(x_i)$ consists of a maximum of a linear function of ρ and zero, hence easier to minimize than the raw empirical suboptimality in practice (and suitable for a recent online-convex-optimization framework [25]).

Before proving Theorem 6, we present a typical example to show that the inconsistency is not too large relative to the suboptimality.

Example 2. Suppose that every edge weight w_e is bounded to $[0, W]$, which ensures that the suboptimality u_ρ is at most $H = W(n-1)$ for any $\rho \in \mathbb{R}^n$. For simplicity, we consider the following natural way to compute ρ values: compute an estimate $\hat{w}_e \in [0, W]$ of w_e for each $e \in E$ and let ρ_v be the cost of a shortest v - t path with respect to $\{\hat{w}_e\}_{e \in E}$. Then, ρ enjoys the consistency with respect to $\{\hat{w}_e\}_{e \in E}$, i.e., $\rho_v \leq \rho_c + \hat{w}_{(v,c)}$ for every $(v, c) \in E$. Therefore, it holds that

$$\Delta_\rho(x) \leq \sum_{(v,c) \in S^*(x), v \neq s} \max\{\rho_v - \rho_c - w_{(v,c)}, 0\} \leq \sum_{(v,c) \in S^*(x), v \neq s} |\hat{w}_{(v,c)} - w_{(v,c)}|.$$

Hence Δ_ρ is at most $\hat{H} = W(n-2)$, implying that Δ_ρ does not largely exceed the suboptimality. If empirically accurate estimates \hat{w}_e for $e \in S^*(x)$ are available, the inconsistency becomes small.

We prove Theorem 6 by using the general analysis framework by Balcan et al. [8]. To begin with, we introduce Assouad's dual class, which provides a formal definition of the class of functions of ρ .

Definition 2 (Assouad [3]). Given a class, $\mathcal{H} \subseteq \mathbb{R}^{\mathcal{Y}}$, of functions $h : \mathcal{Y} \rightarrow \mathbb{R}$, the *dual class* of \mathcal{H} is defined as $\mathcal{H}^* = \{h_y^* : \mathcal{H} \rightarrow \mathbb{R} \mid y \in \mathcal{Y}\}$ such that $h_y^*(h) = h(y)$ for each $y \in \mathcal{Y}$.

In our case, we have $\mathcal{Y} = \Pi$ and $\mathcal{H} = \hat{\mathcal{U}}$, and each $\Delta_x^* \in \hat{\mathcal{U}}^*$ is associated with an instance $x \in \Pi$. The following definition will be used to capture the piecewise structure of the dual class $\hat{\mathcal{U}}^*$.

Definition 3 (Balcan et al. [8, Definition 3.2]). A class, $\mathcal{H} \subseteq \mathbb{R}^{\mathcal{Y}}$, of functions is $(\mathcal{F}, \mathcal{B}, K)$ -*piecewise decomposable* for a class $\mathcal{B} \subseteq \{0, 1\}^{\mathcal{Y}}$ of boundary functions and a class $\mathcal{F} \subseteq \mathbb{R}^{\mathcal{Y}}$ of piece functions if the following condition holds: for every $h \in \mathcal{H}$, there exist K boundary functions $b^{(1)}, \dots, b^{(K)} \in \mathcal{B}$ and a piece function f_b for each binary vector $\mathbf{b} \in \{0, 1\}^K$ such that for all $y \in \mathcal{Y}$, it holds that $h(y) = f_{\mathbf{b}_y}(y)$ where $\mathbf{b}_y = (b^{(1)}(y), \dots, b^{(K)}(y)) \in \{0, 1\}^K$.

The following result of [8] provides an upper bound on the pseudo-dimension of a class of functions via the piecewise structure of the dual class.

Proposition 2 (Balcan et al. [8, Theorem 3.3]). *Let $\mathcal{U} \subseteq \mathbb{R}^\Pi$ be a class of functions. If $\mathcal{U}^* \subseteq \mathbb{R}^{\mathcal{U}}$ is $(\mathcal{F}, \mathcal{B}, K)$ -piecewise decomposable with a class $\mathcal{B} \subseteq \{0, 1\}^{\mathcal{U}}$ of boundary functions and a class $\mathcal{F} \subseteq \mathbb{R}^{\mathcal{U}}$ of piece functions, the pseudo-dimension of \mathcal{U} is bounded as follows:*

$$\text{Pdim}(\mathcal{U}) = O((\text{Pdim}(\mathcal{F}^*) + \text{VCdim}(\mathcal{B}^*)) \lg(\text{Pdim}(\mathcal{F}^*) + \text{VCdim}(\mathcal{B}^*)) + \text{VCdim}(\mathcal{B}^*) \lg K).$$

Now, we are ready to prove Theorem 6.

Proof of Theorem 6. Since there is a one-to-one correspondence between $\Delta_\rho \in \hat{\mathcal{U}}$ and $\rho \in \mathbb{R}^n$, we below identify Δ_ρ with ρ for simplicity. Let $\mathcal{B} = \{\mathbb{I}(\langle \mathbf{z}, \rho \rangle + z_0) \mid (z_0, \mathbf{z}) \in \mathbb{R}^{n+1}\} \subseteq \{0, 1\}^{\hat{\mathcal{U}}}$ and $\mathcal{F} = \{\langle \mathbf{z}, \rho \rangle + z_0 \mid (z_0, \mathbf{z}) \in \mathbb{R}^{n+1}\} \subseteq \mathbb{R}^{\hat{\mathcal{U}}}$ be classes of boundary and piece functions, respectively. We show that $\hat{\mathcal{U}}^*$ is $(\mathcal{F}, \mathcal{B}, O(n^2))$ -piecewise decomposable.

Fix any $\Delta_x^* \in \hat{\mathcal{U}}^*$; this also uniquely specifies an instance $x \in \Pi$ and an optimal solution $S^*(x) \subseteq E$ (due to the tie-breaking). We consider $K = |E| = O(n^2)$ boundary functions of form $b^{(v,c)}(\rho) = \mathbb{I}(\rho_v - \rho_c - w_{(v,c)} > 0)$ for all edges $(v, c) \in E$. We below confirm that these boundary functions partition $\mathbb{R}^n \ni \rho$ into some regions so that in each region, $\Delta_x^*(\rho)$ can be written as a linear function of ρ , which belongs to \mathcal{F} . For each binary vector $\mathbf{b}_\rho = (b^{(v,c)}(\rho))_{(v,c) \in E} \in \{0, 1\}^K$, we define a subset $S_\rho(x)$ of $S^*(x)$ as $S_\rho(x) = \{(v, c) \in S^*(x) \mid b^{(v,c)}(\rho) = 1, v \neq s\}$; that is, each $(v, c) \in S_\rho(x)$ satisfies $v \neq s$ and $\rho_v - \rho_c - w_{(v,c)} > 0$. From the definition of $\Delta_\rho(x)$, we have $\Delta_x^*(\rho) = \Delta_\rho(x) = \sum_{(v,c) \in S_\rho(x)} (\rho_v - \rho_c - w_{(v,c)})$. This is a linear function of ρ , and thus we can choose a piece function $f_{\mathbf{b}_\rho} \in \mathcal{F}$ such that $\Delta_x^* = f_{\mathbf{b}_\rho}$. This relation holds for every $\mathbf{b}_\rho \in \{0, 1\}^K$, and thus we have $\Delta_x^*(\rho) = f_{\mathbf{b}_\rho}(\rho)$ for all $\rho \in \mathbb{R}^n$. Hence $\hat{\mathcal{U}}^*$ is $(\mathcal{F}, \mathcal{B}, O(n^2))$ -piecewise decomposable.

Since \mathcal{F}^* and \mathcal{B}^* can be seen as classes of linear and halfspace functions of $(z_0, \mathbf{z}) \in \mathbb{R}^{n+1}$, respectively, we have $\text{Pdim}(\mathcal{F}^*) = \text{VCdim}(\mathcal{B}^*) = n + 1$ (see also [6, Lemma 3.10], a preprint version of [8]). Therefore, from Proposition 2, we obtain $\text{Pdim}(\hat{\mathcal{U}}) = O(n \lg n)$. \square

6 Conclusion and discussion

We have studied the sample complexity of learning heuristic functions for GBFS and A* on graphs with a fixed vertex set of size n . For GBFS and A*, we have proved that the pseudo-dimensions are upper bounded by $O(n \lg n)$ and $O(n^2 \lg n)$, respectively. As for A*, we have shown that the bound can be improved to $O(n^2 \lg d)$ if every vertex has a degree of at most d and to $O(n \lg n)$ if edge weights are bounded integers. We have also presented the $\Omega(n)$ lower bounds for GBFS and A*, implying that our bounds for GBFS and A* under the integer-weight condition are tight up to a $\lg n$ factor. Finally, we have discussed bounds on the suboptimality of A* and obtained a guarantee with a better complexity term by evaluating the pseudo-dimension of the class of inconsistency functions.

An open problem is to close the gap between the upper and lower bounds regarding A* for general cases. This, however, does not seem straightforward. We here discuss the reasons for the difficulty. As regards the upper bound, the bottleneck is the bound of $(n - 2)!$ on $|\mathcal{G}_v|$, but this cannot be improved in general, as shown in Example 1. Considering this, the direct use of Sauer's lemma would not yield better upper bounds. Thus, we need to use some special structures of the hyperplanes (e.g., each has only two variables), which would require more complicated analysis. As for the lower bound, the construction of the $\Omega(n)$ instances in Theorem 5 relies on the fact that ρ has an n degree of freedom. In addition, Theorem 4 implies that we need to consider instances with non-integer edge weights (or exponentially large integer weights in n) to obtain a lower bound of $\tilde{\Omega}(n^2)$. Considering the above, we would need more involved techniques for constructing a set of $\tilde{\Omega}(n^2)$ shatterable instances.

Another interesting future direction is to improve upper bounds on the pseudo-dimension by restricting heuristic functions to some classes. Appendix D will present an illustrative example where we can achieve $\text{polylog}(n)$ upper bounds on the pseudo-dimension by assuming that heuristic functions with much fewer tunable parameters than n can be designed in an instance-specific manner.

We finally discuss limitations of our work. As mentioned in Section 2, we require every instance to be defined on (subsets of) a fixed vertex set. Also, our work does not cover the case where heuristic function values can change depending on instance-dependent features. Studying how to overcome these limitations would also constitute interesting future work.

Acknowledgements

The authors thank Siddharth Prasad for kindly telling a general result of [10] and anonymous reviewers for providing valuable suggestions, particularly on Appendix D. This work was supported by JST ERATO Grant Number JPMJER1903 and JSPS KAKENHI Grant Number JP22K17853.

References

- [1] F. Agostinelli, S. Mc Aleer, A. Shmakov, R. Fox, M. Valtorta, B. Srivastava, and P. Baldi. Obtaining approximately admissible heuristic functions through deep reinforcement learning and A* search. In *Bridging the Gap Between AI Planning and Reinforcement Learning Workshop at ICAPS 2021*, 2021. [↗ p.2](#)
- [2] M. Anthony and P. L. Bartlett. *Neural Network Learning: Theoretical Foundations*. Cambridge University Press, 1999. [↗ p.3](#)
- [3] P. Assouad. Densité et dimension. *Ann. Inst. Fourier*, 33(3):233–282, 1983. [↗ p.9](#)
- [4] M.-F. Balcan. Data-driven algorithm design. In *Beyond the Worst-Case Analysis of Algorithms*, pages 626–645. Cambridge University Press, 2021. [↗ p.2](#)
- [5] M.-F. Balcan, T. Dick, T. Sandholm, and E. Vitercik. Learning to branch. In *Proceedings of the 35th International Conference on Machine Learning (ICML 2018)*, volume 80, pages 344–353. PMLR, 2018. [↗ p.2](#)
- [6] M.-F. Balcan, D. DeBlasio, T. Dick, C. Kingsford, T. Sandholm, and E. Vitercik. How much data is sufficient to learn high-performing algorithms? Generalization guarantees for data-driven algorithm design. *arXiv:1908.02894*, 2019. [↗ p.10](#)
- [7] M.-F. Balcan, T. Sandholm, and E. Vitercik. Refined bounds for algorithm configuration: The knife-edge of dual class approximability. In *Proceedings of the 37th International Conference on Machine Learning (ICML 2020)*, volume 119, pages 580–590. PMLR, 2020. [↗ p.2](#)
- [8] M.-F. Balcan, D. DeBlasio, T. Dick, C. Kingsford, T. Sandholm, and E. Vitercik. How much data is sufficient to learn high-performing algorithms? Generalization guarantees for data-driven algorithm design. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing (STOC 2021)*, pages 919–932. ACM, 2021. [↗ p.2](#), [↗ p.6](#), [↗ p.9](#), [↗ p.10](#)
- [9] M.-F. Balcan, S. Prasad, and T. Sandholm. Sample complexity of tree search configuration: Cutting planes and beyond. In *Advances in Neural Information Processing Systems 34 (NeurIPS 2021)*, pages 4015–4027. Curran Associates, Inc., 2021. [↗ p.2](#), [↗ p.14](#)
- [10] M.-F. Balcan, S. Prasad, T. Sandholm, and E. Vitercik. Improved sample complexity bounds for branch-and-cut. In *Proceedings of the 28th International Conference on Principles and Practice of Constraint Programming (CP 2022)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022. [↗ p.2](#), [↗ p.10](#), [↗ p.14](#)
- [11] M. Bhardwaj, S. Choudhury, and S. Scherer. Learning heuristic search via imitation. In *Proceedings of the 1st Annual Conference on Robot Learning (CoRL 2017)*, volume 78, pages 271–280. PMLR, 2017. [↗ p.1](#)
- [12] A. Blumer, A. Ehrenfeucht, D. Haussler, and M. K. Warmuth. Learnability and the Vapnik–Chervonenkis dimension. *J. ACM*, 36(4):929–965, 1989. [↗ p.6](#)
- [13] B. Chen, C. Li, H. Dai, and L. Song. Retro*: Learning retrosynthetic planning with neural guided A* search. In *Proceedings of the 37th International Conference on Machine Learning (ICML 2020)*, volume 119, pages 1608–1616. PMLR, 2020. [↗ p.1](#)
- [14] J. Chen and N. R. Sturtevant. Conditions for avoiding node re-expansions in bounded suboptimal search. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI 2019)*, pages 1220–1226. IJCAI Organization, 2019. [↗ p.7](#)
- [15] J. Chen and N. R. Sturtevant. Necessary and sufficient conditions for avoiding reopenings in best first suboptimal search with general bounding functions. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI 2021)*, volume 35, pages 3688–3696. AAAI Press, 2021. [↗ p.7](#)
- [16] E. W. Dijkstra. A note on two problems in connexion with graphs. *Numer. Math.*, 1(1):269–271, 1959. [↗ p.1](#)
- [17] J. E. Doran, D. Michie, and D. G. Kendall. Experiments with the graph traverser program. *Proc. R. Soc. Lond. A Math. Phys. Sci.*, 294(1437):235–259, 1966. [↗ p.1](#)
- [18] T. Eden, P. Indyk, and H. Xu. Embeddings and labeling schemes for A*. In *Proceedings of the 13th Innovations in Theoretical Computer Science Conference (ITCS 2022)*, volume 215, pages 62:1–62:19. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022. [↗ p.2](#)

- [19] R. E. Fikes and N. J. Nilsson. STRIPS: A new approach to the application of theorem proving to problem solving. *Artif. Intell.*, 2(3):189–208, 1971. ↗ p.18
- [20] B. Gärtner and E. Welzl. Vapnik–Chervonenkis dimension and (pseudo-)hyperplane arrangements. *Discrete Comput. Geom.*, 12(4):399–432, 1994. ↗ p.6
- [21] A. V. Goldberg and C. Harrelson. Computing the shortest path: A* search meets graph theory. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2005)*, pages 156–165. SIAM, 2005. ↗ p.4
- [22] R. Gupta and T. Roughgarden. A PAC approach to application-specific algorithm selection. *SIAM J. Comput.*, pages 123–134, 2017. ↗ p.2
- [23] D. Halperin and M. Sharir. Arrangements. In *Handbook of Discrete and Computational Geometry*, pages 723–762. CRC press, third edition, 2017. ↗ p.6
- [24] P. E. Hart, N. J. Nilsson, and B. Raphael. A formal basis for the heuristic determination of minimum cost paths. *IEEE Trans. Syst. Man Cybern. Syst.*, 4(2):100–107, 1968. ↗ p.1
- [25] M. Khodak, M.-F. Balcan, A. Talwalkar, and S. Vassilvitskii. Learning predictions for algorithms with predictions. *arXiv:2202.09312*, 2022. ↗ p.9
- [26] A. Martelli. On the complexity of admissible search algorithms. *Artif. Intell.*, 8(1):1–13, 1977. ↗ p.7
- [27] M. Mohri, A. Rostamizadeh, and A. Talwalkar. *Foundations of Machine Learning*. MIT Press, second edition, 2018. ↗ p.3
- [28] M. Pándy, R. Ying, G. Corso, P. Velickovic, J. Leskovec, and P. Liò. Learning graph search heuristics. In *Physical Reasoning and Inductive Biases for the Real World Workshop at NeurIPS 2021*, 2021. ↗ p.1
- [29] D. Pollard. *Convergence of Stochastic Processes*. Springer, first edition, 1984. ↗ p.3
- [30] N. Sauer. On the density of families of sets. *J. Combin. Theory Ser. A*, 13(1):145–147, 1972. ↗ p.6
- [31] V. Sepetnitsky, A. Felner, and R. Stern. Repair policies for not reopening nodes in different search settings. In *Proceedings of the 9th International Symposium on Combinatorial Search (SoCS 2016)*, volume 7, pages 81–88. AAAI Press, 2016. ↗ p.7, ↗ p.14
- [32] T. Takahashi, H. Sun, D. Tian, and Y. Wang. Learning heuristic functions for mobile robot path planning using deep neural networks. In *Proceedings of the 29th International Conference on Automated Planning and Scheduling (ICAPS 2019)*, volume 29, pages 764–772. AAAI Press, 2019. ↗ p.1
- [33] R. Valenzano and F. Xie. On the completeness of best-first search variants that use random exploration. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI 2016)*, volume 30. AAAI Press, 2016. ↗ p.7
- [34] R. Valenzano, N. Sturtevant, and J. Schaeffer. Worst-case solution quality analysis when not re-expanding nodes in best-first search. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI 2014)*, volume 28, pages 885–892. AAAI Press, 2014. ↗ p.2, ↗ p.7, ↗ p.8, ↗ p.14, ↗ p.15
- [35] V. N. Vapnik and A. Y. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory Probab. Appl.*, 16(2):264–280, 1971. ↗ p.3
- [36] R. Yonetani, T. Taniai, M. Barekatin, M. Nishimura, and A. Kanazaki. Path planning using neural A* search. In *Proceedings of the 38th International Conference on Machine Learning (ICML 2021)*, volume 139, pages 12029–12039. PMLR, 2021. ↗ p.1
- [37] J. You, R. Ying, and J. Leskovec. Position-aware graph neural networks. In *Proceedings of the 36th International Conference on Machine Learning (ICML 2019)*, volume 97, pages 7134–7143. PMLR, 2019. ↗ p.4

Checklist

1. For all authors...

- (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [Yes] See Section 1.1 for details of our main claims.
 - (b) Did you describe the limitations of your work? [Yes] See Section 6.
 - (c) Did you discuss any potential negative societal impacts of your work? [N/A] This is a theoretical paper and no negative societal impacts are expected.
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
- (a) Did you state the full set of assumptions of all theoretical results? [Yes] See Section 2.2.
 - (b) Did you include complete proofs of all theoretical results? [Yes] Some complete proofs are shown in the supplementary due to the page limitation.
3. If you ran experiments...
- (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [N/A]
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [N/A]
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A]
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [N/A]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
- (a) If your work uses existing assets, did you cite the creators? [N/A]
 - (b) Did you mention the license of the assets? [N/A]
 - (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
 - (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? [N/A]
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
- (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

Appendix

A Comparisons with previous results on tree search

We compare our upper bounds with those of existing results on general tree search [9, 10].

Balcan et al. [9] studied the pseudo-dimension for tree-search algorithms in the following situation; a tree-search algorithm with d configurable parameters builds a search tree of size at most κ by iteratively choosing an action from a set of at most T possible actions. In this setting, they obtained an $O(d\kappa \log T + d \log d)$ upper bound. Balcan et al. [10] removed the dependence on κ assuming scores governing the tree search to be defined by *path-wise* functions. Their bound is $O(d\Delta^2 \log k + d\Delta \log T)$, where Δ and k are the maximum depth and the number of children, respectively, of search trees. Since κ can be exponential in the depth Δ , this is a considerable improvement in this context.

In our setting, since there are n configurable parameters, $\rho \in \mathbb{R}^n$, we have $d = n$. If we apply the bound of [9] to GBFS/A* regarded as a tree-search algorithm that iteratively builds search states, κ and T values can be as large as $\Omega(n)$. This is because GBFS/A* can perform $\Omega(n)$ iterations, where each iteration increases the tree size, and the number of possible actions is equal to the size of OPEN, which is $\Omega(n)$ in general. Moreover, for A* with reopening, the number of iterations can be $\Omega(2^n)$ as mentioned in Section 3.2, implying $\kappa = \Omega(2^n)$. Thus, only in the case of A* without reopening, the previous bound matches our $O(n^2 \lg n)$ bound (Theorem 2). As for GBFS and A* with reopening, our Theorems 1 and 2 provide $O(n \lg n)$ and $O(n^2 \lg n)$ bounds, respectively, which improve the $O(n^2 \lg n)$ and $O(n2^n \lg n)$ bounds, respectively, implied by the previous result.

As for the result of [10], seeing GBFS/A* as a tree-search algorithm again, the maximum tree depth is as large as the number of vertices in general, i.e., $\Delta = \Omega(n)$. Also, the number of children can be as large as the size of OPEN, hence $k = \Omega(n)$. Thus, the result of [10] leads to an $O(n^3 \lg n)$ bound, which is larger than any of our upper bounds.

B How to deal with varying t

As mentioned in Section 2.2, given a fixed $t \in V$, each entry ρ_v in ρ indicates an estimated cost of the shortest v - t path. Therefore, if t changes over instances, we need to define ρ for each t , which we here denote by $\rho^t \in \mathbb{R}^n$. If t changes, the structure of path-finding instances also greatly changes. Thus, it is natural to evaluate the performance of GBFS/A* separately for each t . Specifically, we take \mathcal{D} to be a conditional distribution, from which path-finding instances with fixed t are drawn, and we analyze the sample complexity for each possible $t \in V$. In practice, $\{\rho^t\}_{t \in V}$ may be obtained by, e.g., learning a function that embeds vertices into a metric space and measuring distances on the space, as mentioned in Section 2.2. In this case, an embedding function with tunable parameters governs all the $\{\rho^t\}_{t \in V}$ values. Considering such situations, we need to bound the sample complexity of learning heuristic functions for all possible $t \in V$. This can be done at a slight cost of increasing the bound in Proposition 1 by taking a union bound over all possible $t \in V$. (Note that the upper bounds on the pseudo-dimension in Theorems 1, 2 and 6 hold separately for each $t \in V$ by regarding ρ as ρ^t .) Since there are at most n possible choices of t , this replaces δ in Proposition 1 with δ/n , yielding only a $\lg n$ additive factor. This effect is small relative to that of the pseudo-dimension term.

C Worst-case analysis of A* regardless of reopening

We show that the existing bound on the suboptimality of A* by Valenzano et al. [34] holds regardless of whether we allow A* to reopen vertices or not, and we also remove a minor assumption of $\rho_t = 0$. Note that the result of [34], which focuses on the case without reopening, does not immediately imply the same bound for A* with reopening since reopening sometimes degrades the solution quality [31].

We fix an instance and define the inconsistency of an edge $(v, v') \in E$ as

$$\text{Inc}(v, v') = \max\{\rho_v - \rho_{v'} - w_{(v, v')}, 0\}. \quad (\text{A1})$$

Fix an optimal s - t path and let $P_{\text{opt}} = v_0, v_1, \dots, v_k$ be a sequence of vertices on the optimal path, where $v_0 = s$, $v_k = t$, and the optimal s - t path visits the vertices in this order. Suppose that $v_k = t$ is first selected at the $(T + 1)$ st iteration, at which the algorithm terminates.

Theorem 7. Let Cost be the cost of an s - t path returned by A^* (Algorithm 2) with/without reopening, and let Opt be the cost of $P_{\text{opt}} = v_0, v_1, \dots, v_k$. It holds that

$$\text{Cost} \leq \text{Opt} + \sum_{j=1}^{k-1} \text{Inc}(v_j, v_{j+1}).$$

The theorem was proved by Valenzano et al. [34] for Algorithm 2 without reopening. Their proof uses the fact that once a vertex is added to CLOSED, it never gets out of CLOSED. If we allow A^* to reopen vertices, the fact is not always true. Therefore, we need to prove the theorem without using the property of CLOSED. To this end, we define lists of selected vertices, which play a similar role to CLOSED in our proof. Formally, for $\tau = 0, 1, \dots, T + 1$, we define SELECTED_τ as a list of vertices that have been selected in Step 3 in the first τ iterations. Note that even with reopening, once a vertex is added to SELECTED_τ , it never gets out of the list.

As in [34], we derive the bound in Theorem 7 by decomposing P_{opt} into two subpaths, which are defined based on the following shallowest vertex.

Lemma 3. We say a vertex $v_i \in P_{\text{opt}} = v_0, v_1, \dots, v_k$ is the shallowest vertex at $\tau \in \{0, 1, \dots, T\}$ if it satisfies the following conditions after the τ th iteration:

$$v_i \in \text{OPEN} \setminus \text{SELECTED}_\tau \quad \text{and} \quad \{v_j \in P_{\text{opt}} \mid j < i\} \subseteq \text{SELECTED}_\tau.$$

For every $\tau = 0, 1, \dots, T$, a shallowest vertex always exists.

Proof. We prove the claim by induction. If $\tau = 0$, v_0 is the shallowest since we have $\text{SELECTED}_0 = \emptyset$ and $\text{OPEN} = \{v_0\}$. If $\tau = 1$, v_1 is the shallowest vertex since we have $\text{SELECTED}_1 = \{v_0\}$ and $v_1 \in \text{OPEN} \setminus \text{SELECTED}_1$. Assume that the claim is true for $\tau' < \tau$ and let $v_{i'} \in \text{OPEN} \setminus \text{SELECTED}_{\tau-1}$ be the shallowest vertex at $\tau - 1$. We consider two cases: $v_{i'}$ or $v \neq v_{i'}$ is selected at the τ th iteration. If $v \neq v_{i'}$ is selected, we have $\text{SELECTED}_\tau = \text{SELECTED}_{\tau-1} \cup \{v\}$ and $v_{i'} \in \text{OPEN} \setminus \text{SELECTED}_\tau$; thus $v_{i'}$ remains the shallowest at τ . If $v_{i'}$ is selected, take the longest subpath of P_{opt} that starts from $v_{i'}$ and is contained in SELECTED_τ . We denote such a subpath by $v_{i'}, \dots, v_{i''}$, where $i'' < k$ holds since v_k is never selected in the first T iterations. From the definition of the subpath, we have $v_{i''+1} \notin \text{SELECTED}_\tau$. Moreover, $v_{i''+1}$ must have been opened since its parent $v_{i''}$ has been selected. Thus, $v_{i''+1} \in \text{OPEN} \setminus \text{SELECTED}_\tau$ holds. Furthermore, we have $\{v_0, \dots, v_{i'-1}\} \subseteq \text{SELECTED}_\tau$ due to the induction hypothesis and $\{v_{i'}, \dots, v_{i''}\} \subseteq \text{SELECTED}_\tau$ from the definition of the subpath. Thus, $v_{i''+1}$ is the shallowest vertex at τ . To conclude, the shallowest vertex at τ exists in any case. The proof is completed by induction. \square

C.1 Decomposing the suboptimality term with the shallowest vertex

For every $v \in V$, we let g_v^* and ρ_v^* denote the costs of optimal s - v and v - t paths, respectively. We use $g_{v,v'}^*$ to denote the cost of an optimal v - v' path for any pair of $v, v' \in V$; it holds that $g_v^* = g_{s,v}^*$. For each $v \in V$, let $g_v^{(\tau)}$ be the g_v value after the τ th iteration, where $g_s^{(0)} = 0$ and $g_v^{(0)} = \infty$ for $v \neq s$. If g_v is updated in the $(\tau + 1)$ st iteration, we have $g_v^{(\tau+1)} < g_v^{(\tau)}$; otherwise we have $g_v^{(\tau+1)} = g_v^{(\tau)}$. Thus, $g_v^{(\tau)}$ is non-increasing in τ . We define $\delta g_v^{(\tau)} = g_v^{(\tau)} - g_v^*$ as the g -cost error of v after the τ th iteration, which is also non-increasing in τ .

The following lemma states that the suboptimality can be decomposed into two parts: a g -cost error of the shallowest vertex v_i and the inadmissibility of ρ_{v_i} (subtracted by ρ_t).

Lemma 4. If v_i is the shallowest vertex at T , it holds that

$$\text{Cost} \leq \text{Opt} + \delta g_{v_i}^{(T)} + \rho_{v_i} - \rho_{v_i}^* - \rho_t.$$

Proof. After the T th iteration, the score of v_i is

$$\begin{aligned} g_{v_i}^{(T)} + \rho_{v_i} &= g_{v_i}^* + \delta g_{v_i}^{(T)} + \rho_{v_i} \\ &= g_{v_i}^* + \rho_{v_i}^* + \delta g_{v_i}^{(T)} + \rho_{v_i} - \rho_{v_i}^* \\ &= \text{Opt} + \delta g_{v_i}^{(T)} + \rho_{v_i} - \rho_{v_i}^*. \end{aligned}$$

Since $v_k = t$ is selected at the $(T + 1)$ st iteration instead of v_i , it holds that $g_t^{(T)} + \rho_t \leq g_{v_i}^{(T)} + \rho_{v_i}$. Since we have $\text{Cost} \leq g_t^{(T)}$, we obtain the statement by rearranging the terms. \square

C.2 Bounding $\rho_{v_i} - \rho_{v_i}^* - \rho_t$

We prove a general lemma for later use, which implies an upper bound on $\rho_{v_i} - \rho_{v_i}^* - \rho_t$.

Lemma 5. *Let $P = v_0, v_1, \dots, v_k$ be any optimal v_0 - v_k path. It holds that*

$$\rho_{v_0} - \rho_{v_k} - g_{v_0, v_k}^* \leq \sum_{i=0}^{k-1} \text{Inc}(v_i, v_{i+1}).$$

Proof. From the definition (A1), $\text{Inc}(v_i, v_{i+1}) \geq \rho_{v_i} - \rho_{v_{i+1}} - w(v_i, v_{i+1})$ holds. Therefore, we have

$$\sum_{i=0}^{k-1} (\rho_{v_i} - \rho_{v_{i+1}} - w(v_i, v_{i+1})) \leq \sum_{i=0}^{k-1} \text{Inc}(v_i, v_{i+1}).$$

Using a telescoping sum argument, we obtain

$$\rho_{v_0} - \rho_{v_k} - \sum_{i=0}^{k-1} w(v_i, v_{i+1}) \leq \sum_{i=0}^{k-1} \text{Inc}(v_i, v_{i+1}).$$

Since P is optimal, we have $\sum_{i=0}^{k-1} w(v_i, v_{i+1}) = g_{v_0, v_k}^*$, thus completing the proof. \square

Consider applying Lemma 5 to the subpath $P = v_i, \dots, v_k$ of P_{opt} , which is an optimal v_i - v_k path. Since $g_{v_i, v_k}^* = \rho_{v_i}^*$ and $\rho_{v_k} = \rho_t$, it holds that

$$\rho_{v_i} - \rho_{v_i}^* - \rho_t \leq \sum_{j=i}^{k-1} \text{Inc}(v_j, v_{j+1}). \quad (\text{A2})$$

Thus, we can obtain an upper bound on $\rho_{v_i} - \rho_{v_i}^* - \rho_t$.

C.3 Bounding $\delta g_{v_i}^{(T)}$

Our goal is to prove the following lemma.

Lemma 6. *Let $P_{\text{opt}} = v_0, \dots, v_k$ be the optimal s - t path in the statement of Theorem 7. Then, the shallowest vertex $v_i \in P_{\text{opt}}$ at T satisfies*

$$\delta g_{v_i}^{(T)} \leq \sum_{j=1}^{i-1} \text{Inc}(v_j, v_{j+1}).$$

To prove Lemma 6, we need the following two lemmas.

Lemma 7. *For P_{opt} in Lemma 6 and $i \geq 1$, if $v_{i-1} \in P_{\text{opt}}$ is first selected at the τ' th iteration and satisfies $\delta g_{v_{i-1}}^{(\tau')} \leq \sum_{j=1}^{i-2} \text{Inc}(v_j, v_{j+1})$, then $v_i \in P_{\text{opt}}$ satisfies $\delta g_{v_i}^{(\tau)} \leq \sum_{j=1}^{i-1} \text{Inc}(v_j, v_{j+1})$ for all $\tau = \tau', \dots, T$.*

Proof. In the τ' th iteration, we update g_{v_i} if it is larger than $g_{\text{new}} = g_{v_{i-1}}^{(\tau')} + w(v_{i-1}, v_i)$, hence $g_{v_i}^{(\tau')} \leq g_{v_{i-1}}^{(\tau')} + w(v_{i-1}, v_i)$. Since (v_{i-1}, v_i) is an edge on P_{opt} , $g_{v_i}^* = g_{v_{i-1}}^* + w(v_{i-1}, v_i)$ holds. Therefore, it holds that $\delta g_{v_i}^{(\tau')} \leq \delta g_{v_{i-1}}^{(\tau')}$. Since $\delta g_{v_i}^{(\tau)}$ is non-increasing in τ , we have $\delta g_{v_i}^{(\tau)} \leq \delta g_{v_i}^{(\tau')}$. If $i = 1$, since $v_0 = s$, we obtain

$$\delta g_{v_1}^{(\tau)} \leq \delta g_{v_1}^{(\tau')} \leq \delta g_{v_0}^{(\tau')} = g_{v_0}^{(\tau')} - g_{v_0}^* = 0 - 0 = 0.$$

If $i > 1$, we have

$$\delta g_{v_i}^{(\tau)} \leq \delta g_{v_i}^{(\tau')} \leq \delta g_{v_{i-1}}^{(\tau')} \leq \sum_{j=1}^{i-2} \text{Inc}(v_j, v_{j+1}) \leq \sum_{j=1}^{i-1} \text{Inc}(v_j, v_{j+1}),$$

where we used the assumption on $\delta g_{v_{i-1}}^{(\tau')}$ and $\text{Inc}(v_{i-1}, v_i) \geq 0$. \square

Lemma 8. For P_{opt} in Lemma 6 and any $\tau = 0, 1, \dots, T$, every $v_i \in P_{\text{opt}} \cap \text{SELECTED}_\tau$ satisfies $\delta g_{v_i}^{(\tau)} \leq \sum_{j=1}^{i-1} \text{Inc}(v_j, v_{j+1})$.

Proof. The proof is by induction on τ . If $\tau = 0$, the claim is vacuously true since $\text{SELECTED}_\tau = \emptyset$. If $\tau = 1$, only v_0 is in SELECTED_1 . Since $g_{v_0}^{(1)} = g_{v_0}^* = 0$, we have $\delta g_{v_0}^{(1)} = 0$. Thus, the claim is true.

Assume that the claim is true after the first $\tau \geq 1$ iterations; in other words, for any $\tau' \leq \tau$, every $v_{i'} \in P_{\text{opt}} \cap \text{SELECTED}_{\tau'}$ satisfies $\delta g_{v_{i'}}^{(\tau')} \leq \sum_{j=1}^{i'-1} \text{Inc}(v_j, v_{j+1})$. Since $\delta g_{v_{i'}}^{(\tau)}$ is non-increasing in τ , from the induction hypothesis, vertices in $P_{\text{opt}} \cap \text{SELECTED}_\tau$ remain to satisfy the inequality after the $(\tau + 1)$ st iteration. Therefore, we focus on the vertex selected at the $(\tau + 1)$ st iteration, which is the only new vertex in $\text{SELECTED}_{\tau+1}$. If the selected vertex is not in P_{opt} , the statement is true after the $(\tau + 1)$ st iteration. We below discuss the case where the selected vertex is in P_{opt} . We let $v_i \in P_{\text{opt}}$ be the selected vertex, where $i \geq 1$, and discuss two cases: v_i 's parent, v_{i-1} , is in SELECTED_τ or not.

Case 1: $v_{i-1} \in \text{SELECTED}_\tau$. Suppose that v_{i-1} has been first selected at the τ' th iteration ($\tau' \leq \tau$). The induction hypothesis implies $\delta g_{v_{i-1}}^{(\tau')} \leq \sum_{j=1}^{i-2} \text{Inc}(v_j, v_{j+1})$. Thus, from Lemma 7, we obtain $\delta g_{v_i}^{(\tau+1)} \leq \sum_{j=1}^{i-1} \text{Inc}(v_j, v_{j+1})$.

Case 2: $v_{i-1} \notin \text{SELECTED}_\tau$. In this case, we have $i > 1$ since v_0 is selected at the first iteration. Let $v_{i'}$ be the shallowest vertex at τ . Since $v_j \in \text{SELECTED}_\tau$ must hold for all $j < i'$, we have $i' < i$. Since v_i is selected instead of $v_{i'}$, it holds that $g_{v_i}^{(\tau)} + \rho_{v_i} \leq g_{v_{i'}}^{(\tau)} + \rho_{v_{i'}}$. Therefore, we have

$$\delta g_{v_i}^{(\tau)} \leq g_{v_{i'}}^{(\tau)} - g_{v_i}^* + \rho_{v_{i'}} - \rho_{v_i} = \delta g_{v_{i'}}^{(\tau)} + g_{v_{i'}}^* - g_{v_i}^* + \rho_{v_{i'}} - \rho_{v_i} = \delta g_{v_{i'}}^{(\tau)} + \rho_{v_{i'}} - \rho_{v_i} - g_{v_{i'}, v_i}^*.$$

We below consider bounding the right-hand side. First, we discuss a bound on $\delta g_{v_{i'}}^{(\tau)}$. Suppose that $v_{i'-1} \in \text{SELECTED}_\tau$ is first selected at the τ' th iteration, where $\tau' \leq \tau$. From the induction hypothesis, we have $\delta g_{v_{i'-1}}^{(\tau')} \leq \sum_{j=1}^{i'-2} \text{Inc}(v_j, v_{j+1})$. Therefore, Lemma 7 implies $\delta g_{v_{i'}}^{(\tau)} \leq \sum_{j=1}^{i'-1} \text{Inc}(v_j, v_{j+1})$. Next, from Lemma 5, we have $\rho_{v_{i'}} - \rho_{v_i} - g_{v_{i'}, v_i}^* \leq \sum_{j=i'}^{i-1} \text{Inc}(v_j, v_{j+1})$. Consequently, we obtain

$$\delta g_{v_i}^{(\tau+1)} \leq \delta g_{v_{i'}}^{(\tau)} \leq \sum_{j=1}^{i'-1} \text{Inc}(v_j, v_{j+1}) + \sum_{j=i'}^{i-1} \text{Inc}(v_j, v_{j+1}) = \sum_{j=1}^{i-1} \text{Inc}(v_j, v_{j+1}).$$

To conclude, every $v_i \in P_{\text{opt}} \cap \text{SELECTED}_{\tau+1}$ satisfies $\delta g_{v_i}^{(\tau+1)} \leq \sum_{j=1}^{i-1} \text{Inc}(v_j, v_{j+1})$. Therefore, the statement is true by induction. \square

Now, we are ready to prove Lemma 6.

Proof of Lemma 6. Since v_i is the shallowest at T , v_{i-1} has been first selected at some τ th iteration, where $\tau \leq T$, i.e., $v_{i-1} \in P_{\text{opt}} \cap \text{SELECTED}_\tau$. Thus, Lemma 8 implies $\delta g_{v_{i-1}}^{(\tau)} \leq \sum_{j=1}^{i-2} \text{Inc}(v_j, v_{j+1})$. Therefore, from Lemma 7, we obtain $\delta g_{v_i}^{(T)} \leq \sum_{j=1}^{i-1} \text{Inc}(v_j, v_{j+1})$. \square

C.4 Proof of Theorem 7

By summing up the above lemmas and equations, we prove Theorem 7.

Proof of Theorem 7. By using Lemmas 4 and 6 and eq. (A2), we obtain

$$\begin{aligned} \text{Cost} &\leq \text{Opt} + \delta g_{v_i}^{(T)} + \rho_{v_i} - \rho_{v_i}^* - \rho_t \\ &\leq \text{Opt} + \sum_{j=1}^{i-1} \text{Inc}(v_j, v_{j+1}) + \sum_{j=i}^{k-1} \text{Inc}(v_j, v_{j+1}) \\ &= \text{Opt} + \sum_{j=1}^{k-1} \text{Inc}(v_j, v_{j+1}). \end{aligned} \quad \square$$

D An example of improving upper bounds with instance-specific heuristics

GBFS/A* is often applied to path-finding instances that have extremely many vertices, for which our bounds on the pseudo-dimension depending on n or n^2 imply somewhat pessimistic sample complexity bounds. To exhibit more practical results for such cases, we study an example situation where we can exponentially improve the upper bounds by using instance-specific heuristic functions.

We assume the vertex set V to be fixed as in Assumption 1 and that each vertex $v \in V$ corresponds to a *state* represented by a vector $\mathbf{q}_v \in \Sigma^L$, where Σ is a finite set of cardinality B . For example, if $\Sigma = \{0, 1\}$, we have $B = 2$ and each $v \in V$ corresponds to a state represented by a bit vector of length L . Such situations arise, e.g., when applying GBFS/A* to planning instances represented by the STRIPS model [19]. As for heuristic functions, we assume that each ρ_v value is computed as $\rho_v = \mathbf{q}_v^\top \boldsymbol{\theta} + \eta$, where $(\boldsymbol{\theta}, \eta) \in \mathbb{R}^{L+1}$ is a vector of tunable parameters. (Although η is not essential in the following analysis, we can use it to make every ρ_v non-negative in practice.) Here, an important observation is that the number of tunable parameters is $L + 1$, while V has up to $n = B^L$ vertices. Using this exponential decrease in the number of tunable parameters, we below obtain $\text{poly}(L \lg B) \simeq \text{polylog}(n)$ upper bounds on the pseudo-dimension for GBFS and A* under the integer-weight assumption. That is, while the number of vertices is exponential in L , the upper bounds can scale polynomially with L .

GBFS. As discussed in the proof of Lemma 1, if two heuristic function values $\rho, \rho' \in \mathbb{R}^n$ have the same total order on V , it holds $u_\rho(x) = u_{\rho'}(x)$ for all instances $x \in \Pi$. Also, note that the total order is uniquely determined by comparing ρ_v and $\rho_{v'}$ for all pairs of $v, v' \in V$. Therefore, if we partition the space \mathbb{R}^{L+1} of tunable parameters into some regions $\mathcal{P}_1, \mathcal{P}_2, \dots \subseteq \mathbb{R}^{L+1}$ by $\binom{n}{2}$ hyperplanes of form

$$\rho_v = \rho_{v'} \iff \mathbf{q}_v^\top \boldsymbol{\theta} + \eta = \mathbf{q}_{v'}^\top \boldsymbol{\theta} + \eta \iff (\mathbf{q}_v - \mathbf{q}_{v'})^\top \boldsymbol{\theta} = 0,$$

then all $(\boldsymbol{\theta}, \eta)$ values belonging to the same region \mathcal{P}_i result in the same $u_\rho(x)$ value for all $x \in \Pi$. Hence, when $(\boldsymbol{\theta}, \eta)$ is allowed to take any value in \mathbb{R}^{L+1} , the number of distinct tuples of form $(u_\rho(x_1), \dots, u_\rho(x_N))$ is bounded by the number of the regions $\mathcal{P}_1, \mathcal{P}_2, \dots$. As in the proof of Theorem 2, the number of such regions is $O\left(\left(e \binom{n}{2}\right)^{L+1}\right)$ due to Sauer's lemma. On the other hand, to shatter N instances x_1, \dots, x_N , we need to make 2^N distinct tuples of form $(u_\rho(x_1), \dots, u_\rho(x_N))$ by varying $(\boldsymbol{\theta}, \eta) \in \mathbb{R}^{L+1}$. Since $n = O(B^L)$, solving $O\left(\left(e \binom{n}{2}\right)^{L+1}\right) \geq 2^N$ for the largest N yields an $O(L \lg n) \simeq O(L^2 \lg B)$ bound on $\text{Pdim}(\mathcal{U})$.

A* under the integer-weight assumption. We can also obtain a $\text{polylog}(n)$ upper bound for A* under the condition of Theorem 4, i.e., all edge weights take non-negative integer values at most W . The proof begins with the same discussion as that of Theorem 2. For a fixed instance $x_k \in \Pi$, the number of possible $u_\rho(x_k)$ values is at most the number of regions created by hyperplanes of form $g_v(x_k) + \rho_v = g_{v'}(x_k) + \rho_{v'}$ for all pairs of $v, v' \in V$. Since $g_v(x_k)$ can take nW distinct values, the number of those hyperplanes is $\binom{n^2 W}{2}$, as discussed in the proof of Theorem 4. Therefore, given N instances x_1, \dots, x_N , the number of distinct tuples of form $(u_\rho(x_1), \dots, u_\rho(x_N))$ is bounded by the number of regions created by $N \binom{n^2 W}{2}$ hyperplanes. Next, as with the above GBFS case, we bound the number of those regions using Sauer's lemma. With the tunable parameters $(\boldsymbol{\theta}, \eta) \in \mathbb{R}^{L+1}$, each hyperplane can be written as

$$g_v(x_k) + \rho_v = g_{v'}(x_k) + \rho_{v'} \iff (\mathbf{q}_v - \mathbf{q}_{v'})^\top \boldsymbol{\theta} + g_v(x_k) - g_{v'}(x_k) = 0.$$

Sauer's lemma implies that $N \binom{n^2 W}{2}$ such hyperplanes partition \mathbb{R}^{L+1} into $O\left(\left(e N \binom{n^2 W}{2}\right)^{L+1}\right)$ regions. To shatter the N instances, the number of the regions must be at least 2^N . Since $n = O(B^L)$, solving $O\left(\left(e N \binom{n^2 W}{2}\right)^{L+1}\right) \geq 2^N$ for the largest N yields an $O(L \lg nW) \simeq O(L^2 \lg B + L \lg W)$ upper bound on $\text{Pdim}(\mathcal{U})$.

As regards A* for general cases, it is open whether a similar $\text{polylog}(n)$ upper bound can be achieved. The obstacle is the additional n factor, which remains even if the number of tunable parameters decreases. Thus, this problem would be as difficult as closing the gap between the $\Omega(n)$ lower bound and $\tilde{O}(n^2)$ upper bound.