
Learning-Based Algorithms for Graph Searching Problems

Adela Frances DePavia
University of Chicago
adepavia@uchicago.edu

Erasmus Tani
University of Chicago
etani@uchicago.edu

Ali Vakilian
TTIC
vakilian@ttic.edu

Abstract

We consider the problem of graph searching with prediction recently introduced by Banerjee et al. (2023). In this problem, an agent, starting at some vertex r has to traverse a (potentially unknown) graph G to find a hidden goal node g while minimizing the total distance travelled. We study a setting in which at any node v , the agent receives a noisy estimate of the distance from v to g . We design algorithms for this search task on unknown graphs. We establish the first formal guarantees on unknown weighted graphs and provide lower bounds showing that the algorithms we propose have optimal or nearly-optimal dependence on the prediction error. Further, we perform numerical experiments demonstrating that in addition to being robust to adversarial error, our algorithms perform well in typical instances in which the error is stochastic. Finally, we provide alternative simpler performance bounds on the algorithms of Banerjee et al. (2023) for the case of searching on a known graph, and establish new lower bounds for this setting.

1 INTRODUCTION

Searching on graphs is a fundamental problem which models many real-world applications in autonomous navigation. In a *graph searching problem* instance, an agent is initialized at some vertex $r \in V$ (referred to as the root) in some (potentially weighted, directed) graph $G = (V, E)$. The agent’s task is to find a goal node $g \in V$. The agent searches for g by sequentially visiting adjacent nodes in the graph. The graph

searching problem terminates when the agent reaches the goal node, and the cost incurred by the agent is the total amount of distance they travelled.

There are two main settings of interest. In the *exploration* setting, as the agent moves through the graph, it only learns the structure of G by observing the vertices and edges adjacent to the nodes it has visited, a model sometimes referred to as the *fixed graph scenario* (Komm et al., 2015; Kalyanasundaram and Pruhs, 1994). In the strictly-easier *planning* setting, the agent is given the entire graph G ahead of time, but does not know the identity of the goal node g .

Without additional information, in the worst-case the agent must resort to visiting the entire graph, a task which amounts to finding an efficient tour in an unknown graph¹(Berman, 2005; Dobrev et al., 2012; Megow et al., 2012; Eberle et al., 2022). Recently, Banerjee et al. (2023) consider the setting when an algorithm for graph searching also receives some prediction function, $f : V \rightarrow \mathbb{R}$, representing some (noisy) estimate of the distance to the goal node g at any given node $v \in V$. This setup models applications in which the searcher receives advice from some machine-learning model designed to predict the distance to the goal; this problem fits into the broader framework of learning-based algorithms, which exploit (potentially noisy) advice from a machine learning model to enhance their performance.

In the case of planning, Banerjee et al. (2023) propose an intuitive strategy that can be deployed in weighted and unweighted graphs and analyze its performance in terms of structural properties of the instance graph, such as its maximum-degree and its doubling-dimension. They establish formal guarantees on the cost incurred by their algorithm under different notions of prediction error. In contrast, their results in the exploration setting are limited: they propose an

Proceedings of the 27th International Conference on Artificial Intelligence and Statistics (AISTATS) 2024, Valencia, Spain. PMLR: Volume 238. Copyright 2024 by the author(s).

¹We note that historically the task of finding an efficient tour in an unknown graph has been referred to as graph exploration, but following the conventions of Banerjee et al. (2023) we reserve this name for the graph search problem on an unknown graph.

algorithm for exploration in unweighted trees. Their algorithm is tailored to this restricted class of graphs and their guarantees are parameterized by the number of incorrect predictions, which does not capture the magnitude of the deviation between predictions and true distances.

This paper seeks to expand the understanding of graph exploration problems with predictions. We design algorithms which can be deployed on a variety of weighted graphs and prove worst-case guarantees on their performance. We also complement this analysis by providing lower-bounds for these problems, showing that the algorithms studied in this work are optimal or nearly optimal. In particular, we focus on two different error models. In the **absolute** error model, the magnitude of the error at a node is independent of that node’s true distance to the goal, and guarantees are given in terms of the total magnitude of the error incurred at every node. In the **relative** error model, nodes further from the goal may have larger deviation between the prediction and the truth, and algorithmic guarantees are parameterized by the maximum ratio of the error to the true distance at any vertex.

Related Works Online graph searching problems have long been used as basic models for problems in autonomous navigation Berman (2005). Searching with access to predictions, also referred to as “advice” or “heuristics” in different communities, is a commonly studied variant (Pelc, 2002; Dobrev et al., 2012; Eberle et al., 2022; Banerjee et al., 2023). The problem and prediction settings considered in this work most closely correspond to those considered by Banerjee et al. (2023). This setup models applications in which predictions are the output of some machine-learning model. Recent years have seen a marked increase in the integration of machine learning techniques to enhance traditional algorithmic challenges (Angelopoulos et al., 2020; Gupta et al., 2022; Mitzenmacher and Vassilvitskii, 2022; Antoniadis et al., 2023). More generic forms of advice and the advice-complexity of exploration tasks are long-standing subjects of study. Komm et al. (2015) study the case when the searcher receives generic advice, which can take the form of any bit string, and prove results about the advice complexity of this task. For a more detailed survey of related works, we direct the reader to Section A.

Organization In Section 1.1 we formally state the main results of the paper. In Section 2 we present technical preliminaries and define relevant notation. Sections 3 and 4 contain algorithms and analysis for exploration under absolute and relative error models respectively. In Section 5 we derive new bounds in the

planning setting via metric embeddings. Finally, in Section 6 we complement these results with numerical experiments. All missing proofs are in Section B of the supplementary material.

1.1 Summary Of Results

We begin by formally describing the exploration and planning settings:

The Exploration Problem In this setting, both the graph G and the predictions f are initially unknown to the agent: the agent is initialized with access to the root node r , the neighbors of r , and the predictions at all of these nodes. As the algorithm proceeds, on each iteration i it has access in memory to a subgraph $G_i \subseteq G$ containing the nodes it has visited, the neighbors of those nodes, and any edges between visited nodes and neighbors. The searcher can only query predictions from nodes in the subgraph G_i . This problem models exploration of an unknown environment.

The Planning Problem In this setting, both the graph G and the predictions f are fully known to the searcher upon initialization.

For any algorithm which visits an ordered sequence of vertices v_1, \dots, v_T , we denote the algorithmic cost $\text{ALG} \stackrel{\text{def}}{=} \sum_{i \in T} d_{G_i}(v_i, v_{i+1})$, where $d_{G_i}(\cdot, \cdot)$ denotes the shortest (weighted) path distance in the graph G_i described above. The guarantees in this paper compare ALG to the optimal cost a posteriori, denoted by $\text{OPT} \stackrel{\text{def}}{=} d_G(r, g)$.

Exploration Under Absolute Error A natural way of measuring the error of some predictions is the magnitude of the difference between the true distance-to-goal and prediction value at each node. In Section 3 we propose an algorithm for the exploration problem on weighted graphs and prove performance guarantees parameterized by these error measures. In particular, we prove the following theorem:

Theorem 1. *There is an algorithm (Algorithm 1) for searching arbitrary (potentially directed) graphs which finds the goal g by traveling a distance of at most $\text{OPT} + \mathcal{E}_1^- + n \cdot \mathcal{E}_\infty^+$, where $\mathcal{E}_1^- \stackrel{\text{def}}{=} \sum_{v \in V} \max\{0, d(v, g) - f(v)\}$ and $\mathcal{E}_\infty^+ \stackrel{\text{def}}{=} \max_{v \in V} \max\{0, f(v) - d(v, g)\}$.*

This algorithm enjoys several advantages over the most recent results on the exploration problem by Banerjee et al. (2023). Their work introduces an involved combinatorial algorithm for exploration on unweighted trees whose performance is parametrized by the ℓ_0 norm

(the number of non-zero entries) of the vector of errors. The guarantees of their algorithm do not apply when the graph being searched is not an unweighted tree. In contrast, the algorithm proposed in the present work is intuitive and easily implementable, and the guarantees obtained hold in a wide variety of settings, e.g. when the graph is weighted and/or directed. The performance of the algorithm is parameterized by the natural absolute deviation (ℓ_1) error measure.

We also establish that under the above parameterization, the proposed algorithm is in some sense optimal (see Theorem 6). Further our analysis highlights that the same algorithm performs particularly well when (erroneous) predictions only *underestimate* distance to the true goal. Prediction functions with this property are referred to as *admissible*. They are key objects of study in path-finding literature and are well-motivated by applications (Dechter and Pearl, 1985; Eden et al., 2022; Ferguson et al., 2005; Pohl, 1969).

Relative Errors One realistic setting for applications is one in which the error is proportionate to the magnitude of the distance to the goal. In order to capture this behavior, we consider a different error model in which the ratio of the error to true distance-to-goal at every vertex is assumed to be bounded by some value $\varepsilon \in (0, 1)$:

$$(1 - \varepsilon)d(v, g) \leq f(v) \leq (1 + \varepsilon)d(v, g). \quad (1)$$

We do not place any restriction on the total amount of error in the graph beyond this condition.

We consider two regimes of multiplicative error. In the first setting, ε is assumed to be known to the searcher a priori. We propose an algorithm and show that it achieves the following competitive ratio:

Theorem 2. *Consider the exploration problem on a weighted tree where predictions satisfy (1) with respect to $\varepsilon \in (0, 1)$, and ε is known. Then there exists an algorithm (Algorithm 2) which succeeds in finding the goal g and incurs competitive ratio at most*

$$\frac{\text{ALG}}{\text{OPT}} \leq \frac{1}{1 - \varepsilon} + n\varepsilon \cdot \frac{4}{(1 - \varepsilon)^2}.$$

In particular, if the predictions are admissible, then the same algorithm incurs competitive ratio

$$\frac{\text{ALG}}{\text{OPT}} \leq 1 + n\varepsilon \cdot \frac{2}{1 - \varepsilon}.$$

In the second regime ε is assumed to be small ($\varepsilon < 1/3$) but its exact value is not assumed to be known. For this setting, we design a different algorithm which allows us to prove the following result:

Theorem 3. *Given G a weighted tree with predictions f satisfying Equation (1) for some unknown $\varepsilon < 1/3$, then Algorithm 3 with $\beta = 2/3$ incurs competitive ratio at most*

$$\frac{\text{ALG}}{\text{OPT}} \leq 2 + O\left(n\varepsilon \frac{5 + 3\varepsilon}{(1 - 3\varepsilon)^2}\right).$$

In Section 4, we describe our algorithms for these problems, prove the above theorems, and complement these algorithmic guarantees with lower bounds that show these algorithms are nearly optimal.

Planning Problems Banerjee et al. (2023) consider problems in which both the full graph and all predictions are available to the algorithm upon initialization. This setting is referred to as the planning problem. They construct algorithms for this version of the problem under different error models and the guarantees they obtained are outlined in Table 1. In many regimes (e.g. planning on unweighted trees under error parametrized by the ℓ_0 -norm of the vector of errors, denoted \mathcal{E}_0 , and planning on graphs under error parametrized by the ℓ_1 -norm of the vector of errors, denoted \mathcal{E}_1) matching lower bounds for their algorithms can be established, as shown in the table. A notable exception is the case of planning on unweighted graphs under \mathcal{E}_0 parametrization: they establish an upper bound of $\text{OPT} + 2^{O(\alpha)}O(\mathcal{E}_0^2)$ where α is the *doubling dimension* of the graph. In particular, the lower bounds they provide fail to match the dependence on the quadratic term \mathcal{E}_0^2 in their upper bound.

We provide an alternative analysis of their algorithm based on metric properties of the instance graph which shows that in some classes of graphs one can reduce the asymptotic dependence on \mathcal{E}_0 from a quadratic to a linear factor. In particular, we consider the distortion of embedding the instance graph into a weighted path or cycle graph, and establish the following guarantee:

Theorem 4. *Consider G an unweighted graph such that G admits an embedding into a weighted path or a weighted cycle of distortion at most ρ . Then on G , the \mathcal{E}_0 planning algorithm of Banerjee et al. (2023) incurs cost at most $\text{OPT} + O(\rho\mathcal{E}_0)$.*

We note that the results present in Table 1 which were not proved by Banerjee et al. (2023) are established in this paper in Section 5 and the proved in the supplementary material.

2 TECHNICAL PRELIMINARIES AND NOTATION

Graphs are assumed to be weighted and directed, unless otherwise specified. (Weighted) shortest-path dis-

		\mathcal{E}_0 , unweighted	\mathcal{E}_1 , positive weights
Trees with maximum degree Δ	upper bound	$\text{OPT} + O(\Delta\mathcal{E}_0)$ (†)	$\text{OPT} + O(\Delta\mathcal{E}_1^2)$ (integer distances)
	lower bound	$\max\{\text{OPT}, \Omega(\Delta\mathcal{E}_0)\}$ (†)	$\max\{\text{OPT}, \Omega(\Delta\mathcal{E}_1^2)\}$
Graphs with doubling dimension α	upper bound	$\text{OPT} + 2^{O(\alpha)}O(\mathcal{E}_0^2)$ (†)	$\text{OPT} + 2^{O(\alpha)}O(\mathcal{E}_1)$ (†)
	lower bound	Unknown	$\max\{\text{OPT}, \Omega(2^\alpha\mathcal{E}_1)\}$
Graphs with path-embedding distortion ρ	upper bound	$\text{OPT} + O(\rho\mathcal{E}_0)$	$\text{OPT} + O(\rho\mathcal{E}_1)$
	lower bound	Unknown	$\max\{\text{OPT}, \Omega(\rho\mathcal{E}_1)\}$

Table 1: Known results for planning problem in different settings. (†) denotes results from Banerjee et al. (2023). Δ denotes the maximum degree of any vertex in the graph, α denotes the doubling-dimension of the graph, and ρ denotes the distortion of the path-embedding on G . Banerjee et al. (2023) in their work note the absence of a matching lower bound in the graph planning problem. This gap motivates this work’s study of planning bounds parameterized by metric embeddings, which yields a reduced asymptotic dependency on \mathcal{E}_0 . For full discussions of lower bounds for \mathcal{E}_1 parametrized by Δ , α , and ρ , see Lemmas 11 and 13, and Lemma 20 in Section B of the supplementary material.

tances in a graph G are denoted by $d_G(\cdot, \cdot)$. Given a set $S \subseteq V$, let ∂S be its external vertex boundary:

$$\partial S \stackrel{\text{def}}{=} \{v \in G \setminus S \mid \exists u \in S : v \sim u\}$$

For a vertex set $S \subseteq V$, we denote by $\text{tour}_G(S)$ the (weighted) length of the shortest walk that visits all nodes in S :

$$\text{tour}_G(S) \stackrel{\text{def}}{=} \max_{v \in S} \min_{W \in \mathcal{W}(v, S)} \text{length}_G(W), \quad (2)$$

where $\mathcal{W}(v, S)$ is the set of walks in G starting at vertex v and visiting every vertex in S .

Given a metric space (X, d_X) its doubling constant is the minimum number λ such that, for every $r > 0$, every ball of radius r can be covered by at most λ balls of radius $r/2$ (Gupta et al., 2003). The doubling constant of a graph G is the doubling constant of (G, d) where d is the shortest path distance on G . The doubling dimension of the space, denoted α , is defined as $\alpha \stackrel{\text{def}}{=} \log_2 \lambda$.

Notation For Exploration Algorithms Recall that in exploration problems, the true graph G is initially unknown to the searcher. Thus in the setting of exploration, one needs to distinguish between the shortest *known* path distance between two vertices and the true shortest path distance. To this end, let V_{i-1} be the set of vertices *visited* by iteration i and let G_i be the subgraph of G containing: all of the vertices in $V_{i-1} \cup \partial V_{i-1}$, and all of the edges adjacent to V_{i-1} . Throughout this paper, we emphasize the distinction between d_{G_i} and d_G . In general, we have $d_{G_i} \geq d_G$.

When analyzing performance, we denote the *progress* made on the i th iteration as

$$\Delta_i \stackrel{\text{def}}{=} d_G(v_i, g) - d_G(v_{i+1}, g). \quad (3)$$

Observe that, under the assumption that v_0, v_1, \dots, v_T are nodes visited by some algorithm which originates at r and terminates at g (i.e. $v_0 = r$ and $v_T = g$) we have:

$$\sum_{i=0}^{T-1} \Delta_i = d_G(r, g) - 0 = \text{OPT}.$$

Metric Embeddings When analyzing the planning algorithms of Banerjee et al. (2023), we consider metric embeddings on graphs, and parametrize results in terms of the *distortion* of the relevant embedding:

Definition 1 (Distortion). Given a function $\tau : X \rightarrow Y$ between two finite metric spaces (X, d_X) and (Y, d_Y) , we define the distortion $\text{dist}(\tau)$ of τ as the minimum value ρ satisfying the following: there exists a constant $c > 0$ such that for all $x_1, x_2 \in X$,

$$c \cdot d_X(x_1, x_2) \leq d_Y(\tau(x_1), \tau(x_2)) \leq c \cdot \rho \cdot d_X(x_1, x_2).$$

3 EXPLORATION UNDER ABSOLUTE ERROR

In this section we study the absolute error regime, in which the ℓ_1 norm of the vector of errors is bounded by some constant \mathcal{E}_1 , not necessarily known to the searcher. We consider the following natural rule: on the i th iteration, choose the next vertex to visit by

picking the node $v_i \in \partial V_{i-1}$ minimizing the sum of $d_{G_i}(v_{i-1}, v_i)$ and $f(v_i)$ (See Algorithm 1). This iterative step can be interpreted as visiting the vertex that would be on the shortest path to the goal if all the predictions $f(v)$ were correct.

Algorithm 1 ℓ_1 -Greedy_Search(G, r)

```

 $v_0 \leftarrow r$ 
 $i \leftarrow 0$ 
 $V_0 \leftarrow \{r\}$ 
while  $v_i \neq g$  do
     $i \leftarrow i + 1$ 
     $v_i \in \operatorname{argmin}_{v \in \partial V_{i-1}} d_{G_i}(v_{i-1}, v) + f(v)$ 
     $V_i \leftarrow \{v_0, \dots, v_i\}$ 
end while

```

We remark that because we are working in the exploration setting, Algorithm 1 does not have access to the true distances $d_G(\cdot, \cdot)$ and must instead make use of $d_{G_i}(\cdot, \cdot)$ the distances in the subgraph of observed vertices at iteration i .

Theorem 1 parametrizes the worst-case guarantees for Algorithm 1 in terms of the total negative error \mathcal{E}^- and the maximum-occurring positive error \mathcal{E}_∞^+ . This asymmetry corresponds to the intuition that positive errors can obstruct the search task more dramatically than negative errors by obscuring shortest paths to the goal. Indeed, an immediate corollary of Theorem 1 is the following result about the performance of Algorithm 1 in the setting in which the prediction function is admissible (i.e. error is only negative):

Corollary 5. *Consider the problem of searching a weighted (possibly directed) graph with predictions f satisfying $f(v) \leq d_G(v, g) \forall v \in V$. Then there exists an algorithm which finds the goal g with cost at most $\text{OPT} + \mathcal{E}_1$, where \mathcal{E}_1 is the total ℓ_1 error in the predictions, i.e. $\mathcal{E}_1 \stackrel{\text{def}}{=} \sum_{v \in V} |f(v) - d(v, g)|$.*

The dependency on \mathcal{E}^- and \mathcal{E}_∞^+ in Theorem 1 is optimal in the following sense:

Theorem 6. *For every $\mathcal{E}^- > 0$, there exist graph search instances with total negative error \mathcal{E}^- such that any algorithm for the exploration problem on these instances must incur cost at least $\text{OPT} + \mathcal{E}^-$ in the worst case. Additionally, for any $n > 3$ and any \mathcal{E}_∞^+ , there exist graph search instances on n nodes with maximum positive error \mathcal{E}_∞^+ such that any algorithm for the exploration problem on these instances must incur cost at least $\text{OPT} + \mathcal{E}_\infty^+(n - 2)$ in the worst case.*

We note that the lower bound in Theorem 6 also holds for the expected distance travelled of randomized search strategies, up to constant factors, as per the following result.

Proposition 7. *For every $\mathcal{E}^- > 0$ there exists a graph search instance with total negative error \mathcal{E}^- such that any randomized algorithm incurs expected costs at least $\text{OPT} + \frac{1}{2}\mathcal{E}^-$ on this instance. Moreover, for any $n > 3$ and any \mathcal{E}_∞^+ there exists a graph search instance on n nodes with maximum positive error \mathcal{E}_∞^+ such that any randomized algorithm must incur expected cost at least $\text{OPT} + (n - 2)\mathcal{E}_\infty^+/2$ on this instance.*

The full proof of Theorem 1 and the proof of Corollary 5, given in Section B, rely on a charging argument which shows that distance travelled away from the goal can be directly attributed to errors in the predictions of observed nodes. The proof of Theorem 6 and Proposition 7 are constructive and can also be found in Section B.

For completeness we now sketch the proof of Theorem 1. The proof considers the progress Δ_i as in Equation (3), which measures how much closer the agent is to the goal after the i^{th} iteration of the algorithm. The cost of the algorithm is given by $\text{ALG} = \sum_{i \in [T]} d_{G_i}(v_{i-1}, v_i)$. At each step $i \in [T]$, one can show that the distance travelled $d_{G_i}(v_{i-1}, v_i)$ in the observed subgraph G_i is bounded above by the sum of three terms:

$$d_{G_i}(v_{i-1}, v_i) \leq \Delta_i + \mathcal{E}^-(v_i) + \mathcal{E}^+(w_i), \quad (4)$$

where $\mathcal{E}^-(v_i)$ is the negative error at v_i , and $\mathcal{E}^+(w_i)$ is the positive error at some vertex w_i on a shortest path from v_{i-1} to g . In particular, all three terms in this upper bound are independent of the observed subgraph G_i . The statement of the theorem then follows by summing both sides of Equation (4) over all $i \in [T]$.

4 EXPLORATION UNDER RELATIVE ERROR

In this section, we consider the setting when the prediction function satisfies Equation (1) for every $v \in V$. We assume that $\varepsilon \in (0, 1)$: in particular, if $\varepsilon \geq 1$, then $f(v) = 0$ is a valid prediction at every vertex and no exploration algorithm can avoid visiting the entire graph in the worst case.

If ε is known to the searcher a priori then given access to the prediction at a node, the searcher can construct an upper bound on the true distance-to-goal. On trees this allows one to limit exploration to a ball of some radius R (dependent on ε and OPT) around the initial vertex, effectively “pruning” distant nodes from the vertex set. In particular, one could limit their search to the set:

$$S_{\varepsilon, r} \stackrel{\text{def}}{=} \left\{ v \in V \mid d_G(v, r) \leq \frac{1}{1 - \varepsilon} f(r) \right\}. \quad (5)$$

We couple this observation with the algorithm in the previous section and obtain the following algorithm.

Algorithm 2 ε -Known_Search(G, r, ε)

```

 $v_0 \leftarrow r$ 
 $i \leftarrow 0$ 
 $V_0 \leftarrow \{r\}$ 
while  $v_i \neq g$  do
     $i \leftarrow i + 1$ 
     $v_i \in \operatorname{argmin}_{v \in \partial V_{i-1} \cap S_{\varepsilon, r}} d_{G_i}(v_{i-1}, v) + f(v)$ 
     $V_i \leftarrow \{v_0, \dots, v_i\}$ 
end while
    
```

In Section B we leverage favorable properties of the set $S_{\varepsilon, r}$ to show that Algorithm 2 satisfies the guarantees of Theorem 2. We observe that in particular, Theorem 2 implies that for any n the competitive ratio incurred by Algorithm 2 tends to 1 as $\varepsilon \rightarrow 0$. The combination of the truncation with the shortest-path rule in Algorithm 2 was necessary to secure this property: for example, a simple scheme such as running breadth-first-search on the truncated set $S_{\varepsilon, r}$ would not enjoy such a guarantee in the worst case.

We note that Algorithm 2 crucially relies on the fact that the searcher knows the value of ε and so it cannot be deployed in the setting where ε is unknown. For the latter regime we propose an alternative algorithm, also based on Algorithm 1, which provably succeeds for unknown values of ε under the assumption that ε is small, e.g. $\varepsilon < 1/3$.

Algorithm 3 ε -Unknown_Weighted_Search(G, r, β)

```

 $v_0 \leftarrow r$ 
 $i \leftarrow 0$ 
 $V_0 \leftarrow \{r\}$ 
while  $v_i \neq g$  do
     $i \leftarrow i + 1$ 
     $v_i \in \operatorname{argmin}_{v \in \partial V_{i-1}} \beta d_{G_i}(v_{i-1}, v) + f(v)$ 
     $V_i \leftarrow \{v_0, \dots, v_i\}$ 
end while
    
```

In Section B we show that on trees this reweighting scheme ensures that Algorithm 3 never explores nodes which are far from the goal and use this property to establish the guarantees in Theorem 3 under the setting where $\beta = 2/3$.

We give a lower bound to establish that Algorithm 2 is almost optimal. Specifically, we show that even when ε is known a-priori, the asymptotic dependence on $n \cdot \varepsilon$ is tight up to factors of $1/(1 - \varepsilon)$:

Theorem 8. *For all n sufficiently large ($n \geq 6$) and for any $\varepsilon \in (0, 1)$, there exists an instance \mathcal{I} of the exploration on weighted trees G with predictions (1),*

such that any algorithm for the exploration problem on G must incur cost $\text{ALG} \geq (1 + n\varepsilon)\text{OPT}$ on \mathcal{I} .

Note that this lower bound also applies to the regime addressed by Algorithm 3. Moreover, one can prove an analogous lower bound for the case of potentially randomized exploration strategies, as per the following result.

Proposition 9. *For all n sufficiently large ($n \geq 6$) and for any $\varepsilon \in (0, 1)$, there exists an instance \mathcal{I} of the exploration on weighted trees G with predictions (1), such that any randomized algorithm for the exploration problem on G must incur cost $\text{ALG} \geq (1 + \frac{n\varepsilon}{2})\text{OPT}$ on \mathcal{I} .*

4.1 Planning With Relative Error

Under the model of relative error described by Equation (1), planning problems become either trivial or impossible, with no intermediate regimes. In the context of predictions $f(v) = (1 + \varepsilon_v)d_G(v, g)$ for some $\varepsilon_v \in [-\varepsilon, \varepsilon]$, we observe that it must be that $f(g) = 0$, independent of the value of ε . For the planning problem to be nontrivial, there must also occur vertices $v \neq g$ such that $f(v) = 0$. Thus, for nontrivial instances of the planning problem in this regime, $\varepsilon \geq 1$. However, instances with such (large) values of ε are hopeless in the worst case: such a setting allows for the prediction of *every* node to be set equal to 0, a case which forces the searcher to visit every node in the worst case.

5 PLANNING

5.1 Planning Bounds Via Metric Embeddings

In this section, we analyze the performance of the algorithm by Banerjee et al. (2023) for the planning problem, as a function of how similar the target graph is to some graph G' admitting inexpensive tours. In the planning problem, the full graph G as well as all predictions f are made available to the algorithm upon initialization. Banerjee et al. (2023) study the implied error functions $\phi_0 : V \rightarrow \mathbb{R}$ and $\phi_1 : V \rightarrow \mathbb{R}$, defined as

$$\phi_0(v) \stackrel{\text{def}}{=} |\{u \in V : f(u) \neq d_G(u, v)\}|$$

and

$$\phi_1(v) \stackrel{\text{def}}{=} \sum_{u \in V} |f(u) - d_G(u, v)|.$$

Banerjee et al. (2023) consider an algorithm that iteratively visits sublevel sets of ϕ_0 or ϕ_1 respectively, for geometrically increasing thresholds (see Algorithm 4). Their algorithm is very simple: for every threshold, it visits every node in the sublevel set before increasing

to the next threshold value. Algorithmically, Banerjee et al. (2023) accomplish this by computing a minimum length Steiner tree of the sublevel set, which is in general not computationally efficient. However, one can replace this computationally expensive procedure with a polynomial-time constant factor approximation (see e.g. (Karlín et al., 2021)), which preserves the asymptotic upper bound on the algorithms’ performance.

Algorithm 4 FullInfoX (G, r, ϕ) from Banerjee et al. (2023)

```

▷  $\forall \kappa \in \mathbb{R} : L_{\phi}^{-}(\kappa) := \{v \in G \mid \phi(v) \leq \kappa\}$ 
 $V_i \leftarrow r$ 
 $\lambda \leftarrow 0$ 
while  $g \notin V_i$  do
     $V_i \leftarrow V_i \cup L_{\phi}^{-}(2^{\lambda})$ 
    Compute a minimum length Steiner tree of  $V_i$ 
    and perform an Euler tour of the tree
     $\lambda \leftarrow \lambda + 1$ 
end while

```

This definition of Algorithm 4 motivates our analysis, which focuses on the distortion of embedding the instance graph into some graph G' which admits inexpensive tours. Recall the definition of $\text{tour}_G(S)$ as in Equation (2).

Definition 2 (Easily-tourable). A graph $G' = (V', E')$ is c -easily-tourable for some $c > 0$ if for any $S' \subseteq V'$, $\text{tour}_{G'} \leq c \cdot \text{diam}(S')$.

In Section B, we establish that in easily-tourable graphs the algorithm of Banerjee et al. enjoys good performance. We then show that if a graph G can be embedded into an easily-tourable graph G' with distortion ρ , then G itself must be easily-tourable with tour costs that scale with ρ . This culminates in the following result:

Lemma 10. *Given G an unweighted graph, if G admits an embedding $\tau : G \rightarrow G'$ of distortion ρ for G' some $c_{G'}$ -easily-tourable graph, then Algorithm 4 with objective $\phi = \phi_0$ from Banerjee et al. (2023) incurs cost at most $\text{OPT} + O(\rho \cdot c_{G'} \cdot \mathcal{E}_0)$. If G has integer-valued distances and admits an embedding of distortion ρ into G' some easily-tourable graph, then Algorithm 4 with objective $\phi = \phi_1$ incurs cost at most $\text{OPT} + O(\rho \cdot c_{G'} \cdot \mathcal{E}_1)$.*

In particular, (weighted) paths and cycles are easily-tourable with respect to constant c , resulting in the guarantees in Theorem 4. In Section C we give results suggesting that our analysis of planning problems via metric embeddings is a refinement of the analysis of Banerjee et al. (2023).




Figure 1: The lower bound construction for the proof of Lemmas 11 and 13.

5.2 Lower bounds For Planning

In this section, we provide lower bounds that extend the results in Banerjee et al. (2023). Consider the planning problem on some weighted graph G with integer weights² where prediction error is parameterized by \mathcal{E}_1 . Banerjee et al. (2023) propose an algorithm which provably incurs cost at most $\text{OPT} + O(\mathcal{E}_1 \text{poly}(\lambda))$, where λ is the doubling constant of the input graph. We provide complementary lower bounds by establishing the following result:

Lemma 11. *Let \mathcal{A} be any algorithm for the planning problem which is guaranteed to incur cost: $\text{OPT} + O(\mathcal{E}_1^a \lambda^b)$ for some $a, b \in \mathbb{R}$ when run on a weighted graph G with doubling constant λ , and with a error vector \vec{e} such that $\|\vec{e}\|_1 = \mathcal{E}_1$. Then it must be the case that $a \geq 1$. Moreover, if $a = 1$, then $b \geq 1$.*

Our lower bounds are constructive, and consider a family of graphs illustrated in Figure 1. A full proof can be found in Section B.

It is known that the doubling constant is always at least as large as the maximum degree but we note that in general it may be much larger even if one restricts themselves to trees. We analyze the algorithm of Banerjee et al. (2023) and prove that in trees with integer weights performance of their algorithm can be bounded in terms of maximum degree, at the cost of paying a higher asymptotic dependence on \mathcal{E}_1 :

Lemma 12. *Given G a tree with integer-valued distances and maximum degree Δ , consider the planning problem on G with predictions satisfying $\mathcal{E}_1 \geq 1$.*

²We note that the analysis of Banerjee et al. for the above setting also appears to go through for the case non-integer weights, and that the lower bound provided by Lemma 11 would then hold for that setting too.




Figure 2: Performance of Algorithms 1 and 3 against random errors. Experimental procedures are detailed in Section D. The number of nodes is fixed over all graph topologies and error settings. LEFT: Average and standard deviation of $\text{ALG} - \text{OPT}$ incurred by Algorithm 1 over 2000 independent random trials for varying values of \mathcal{E}_1 . RIGHT: Average and standard deviation of ALG/OPT incurred by Algorithm 3 over 2000 independent random trials for varying values of ε .

Then on this problem instance Algorithm 4 with objective $\phi = \phi_1$ from Banerjee et al. (2023) incurs cost at most $\text{OPT} + O(\Delta \mathcal{E}_1^2)$.

We establish corresponding lower bounds via a similar construction to that in Lemma 11:

Lemma 13. *Let \mathcal{A} be any algorithm for the planning problem on trees with integer weights which is guaranteed to incur cost: $\text{OPT} + O(\mathcal{E}_1^a \Delta^b)$ for some $a, b \in \mathbb{R}$ when run on a tree G with maximum degree Δ and with a prediction vector \vec{e} such that $\|\vec{e}\|_1 = \mathcal{E}_1$. Then it must be the case that $a \geq 1$ and $b \geq 1$. Moreover, it must be that $a + b \geq 3$.*

6 NUMERICAL EXPERIMENTS: IMPACT OF RANDOM ERRORS

Throughout this paper we have focused on worst-case theoretical guarantees. In this section, we provide numerical results exploring the effectiveness of Algorithms 1 and 3 on exploration problems beyond the worst-case setting, particularly under random errors. The experiments show that in addition to being robust to adversarial error, the algorithms considered perform well in instances with stochastic error. Moreover, we find that although the guarantees for Algorithm 3 were shown only for trees, empirically the algorithm also succeeds on general (cyclic) graphs. These results suggest that Algorithms 1 and 3 could be deployed effectively for exploration problems in practice.

We study the performance of Algorithm 1 in the absolute regime and Algorithm 3 in the relative regime. In the left subfigure of Figure 2 we plot the performance of Algorithm 1 for different graph topologies when the error is sampled at random in an absolute fashion. We

plot the performance against the total ℓ_1 -norm of the error vector. In the right subfigure of Figure 2 we explore the performance of Algorithm 3 against relative error. Once again, we find that the algorithm enjoys empirical performance superior to that predicted by the worst-case upper bound. In particular, the gap between the worst-case bound and the average empirical performance is consistent over families of graphs with very different topologies.

In Table 2 we report the average ratio of algorithmic cost incurred to value of the upper-bound in Theorem 1 (given as a percentage). We compute this percentage for different classes of graphs over 100 runs of these experiments when the error, initial node and goal node, and graph structure (when applicable) have been sampled at random.

We also empirically compare the performance of Algorithm 1 to another natural heuristic, which we call Smallest Prediction. In Smallest Prediction, the agent always travels to the vertex v in ∂V_{i-1} with the smallest value of the prediction function $f(v)$. While our theoretical results already show that Algorithm 1 achieves optimal performance in the worst-case, we show that our algorithm performs better than Smallest Prediction in the presence of random error across a variety of graph topologies. In Figure 3, we plot the average performance of Algorithm 1 against the performance of Smallest Prediction (measured as the distance travelled by the agent, minus OPT , as a fraction of OPT) in random trees with 100 vertices for a growing value of the magnitude \mathcal{E}_1 of the error vector. More details on this comparison can be found in the supplementary material.

Further details of our experiments, including details

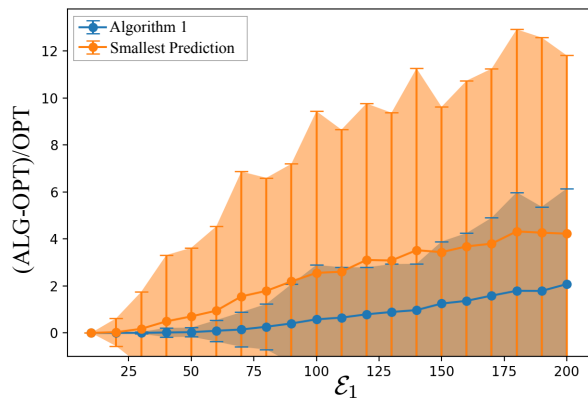


Figure 3: A comparison of the performance of Algorithm 1 with the Smallest Prediction heuristic. We plot the average and the standard deviation of the performance of Algorithm 1 and that of the Smallest Prediction heuristic against the magnitude of the error vector \mathcal{E}_1 . Experiments in this figure were conducted on random trees; for analogous results on other graph topologies, see Figure 9 in Appendix D.

GRAPH FAMILY	Random Lobster	Erdős Rényi	Random Tree	Circular Ladder
COST(%)	2.4 ± 2.5	3.1 ± 4.3	1.7 ± 2.3	0.7 ± 0.5

Table 2: Average empirical cost of running Algorithm 1 as a percentage of the upper bound in Theorem 1 for different family of graphs. Experiments performed on graphs with 300 nodes. These results demonstrate that when run with randomly-generated errors, the actual cost incurred by the algorithm is a very small fraction of the upper bound.

on the error models and the graph families being used in the experiments can be found in Section D in the supplementary material.

7 CONCLUSIONS AND FUTURE DIRECTIONS

In this work we have introduced new general algorithms for the problem of searching in an unknown graph. Under the absolute error model we design algorithms which succeed in a broad class of graphs and prove that these algorithms are optimal (Section 3). We then move beyond the absolute error regime and consider relative error; to the best of the authors' knowledge, this work is the first to address the exploration problem under this natural error model. Within this setting we propose algorithms for the exploration problem on weighted trees and show that their performance is nearly-optimal (Section 4).

We complement our advances in the exploration setting by expanding the landscape of results for the planning problem. We extend the work of Banerjee et al. (2023) by providing alternative performance guarantees which establish a linear—rather than quadratic—dependency on the error parameter \mathcal{E}_0 in some graph families, and which suggests that such a lower asymptotic dependency may be attainable in general (Theorem 4). We also complete the results of Banerjee et al. in the planning setting on integer-distance graphs by proving one cannot improve the factor of \mathcal{E}_1 in their upper bound and that achieving this linear dependence on the error requires cost linear in the doubling constant λ of the instance graph (Lemma 11).

The work in this paper directly suggests several avenues for further study. While our lower bounds demonstrate the impossibility of uniformly improving the results in Theorem 1, it is possible the bound may be overly pessimistic in certain classes of graphs; it would be interesting to consider whether making stronger structural assumptions about the instance graph would yield better guarantees. In the setting of relative error, an immediate open problem is whether the guarantees on Algorithms 2 and 3 can be extended to more general graphs. In order to improve understanding of the planning problem, a complete characterization of easily-tourable graphs would better contextualize Lemma 10. Finally, while the numerical results suggest the algorithms proposed in this paper perform well under random errors, formal guarantees studying this setting would be a valuable addition.

Acknowledgements

The authors wish to thank all the anonymous reviewers for their thoughtful feedback.

A. DePavia is supported by NSF DGE 2140001.

While working on this project, E. Tani was supported by the Institute for Data, Econometrics, Algorithms, and Learning (IDEAL) with the NSF Grant ECCS-2216912.

References

- A. Aamand, P. Indyk, and A. Vakilian. Frequency estimation algorithms under zipfian distribution. *arXiv preprint arXiv:1908.05198*, 2019.
- A. Aamand, J. Y. Chen, and P. Indyk. (Optimal) Online Bipartite Matching with Degree Information. In *Advances in Neural Information Processing Systems*, volume 35, 2022.
- S. Alpern and S. Gal. *The theory of search games and*

- rendezvous*, volume 55. Springer Science & Business Media, 2006.
- K. Anand, R. Ge, and D. Panigrahi. Customizing ml predictions for online algorithms. In *International Conference on Machine Learning*, pages 303–313, 2020.
- S. Angelopoulos, C. Dürr, S. Jin, S. Kamali, and M. Renault. Online computation with untrusted advice. In *11th Innovations in Theoretical Computer Science Conference (ITCS 2020)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2020.
- A. Antoniadis, C. Coester, M. Eliáš, A. Polak, and B. Simon. Online metric algorithms with untrusted predictions. *ACM Transactions on Algorithms*, 19(2):1–34, 2023.
- E. Bamas, A. Maggiori, and O. Svensson. The primal-dual method for learning augmented algorithms. *Advances in Neural Information Processing Systems*, 33:20083–20094, 2020.
- S. Banerjee, V. Cohen-Addad, A. Gupta, and Z. Li. Graph searching with predictions. In *14th Innovations in Theoretical Computer Science Conference, ITCS*, volume 251 of *LIPICs*, pages 12:1–12:24, 2023.
- P. Berman. On-line searching and navigation. *Online Algorithms: The State of the Art*, pages 232–241, 2005.
- A. Bhattacharya, B. Gorain, and P. S. Mandal. Treasure hunt in graph using pebbles. In *International Symposium on Stabilizing, Safety, and Security of Distributed Systems*, pages 99–113. Springer, 2022.
- S. Bouchard, Y. Dieudonne, A. Labourel, and A. Pelc. Almost-optimal deterministic treasure hunt in arbitrary graphs. In *International Colloquium on Automata, Languages and Programming (ICALP) 2021*, 2021.
- J. Y. Chen, T. Eden, P. Indyk, H. Lin, S. Narayanan, R. Rubinfeld, S. Silwal, T. Wagner, D. P. Woodruff, and M. Zhang. Triangle and four cycle counting with predictions in graph streams. In *10th International Conference on Learning Representations, ICLR, 2022*.
- E. Cohen, O. Geri, and R. Pagh. Composable sketches for functions of frequencies: Beyond the worst case. In *Proceedings of the 37th International Conference on Machine Learning*, 2020.
- R. Dechter and J. Pearl. Generalized best-first search strategies and the optimality of A*. *Journal of the ACM (JACM)*, 32(3):505–536, 1985.
- I. Diakonikolas, V. Kontonis, C. Tzamos, A. Vakilian, and N. Zarifis. Learning online algorithms with distributional advice. In *International Conference on Machine Learning*, pages 2687–2696, 2021.
- S. Dobrev, R. Kráľovič, and E. Markou. Online graph exploration with advice. In *International Colloquium on Structural Information and Communication Complexity*, pages 267–278. Springer, 2012.
- E. Du, F. Wang, and M. Mitzenmacher. Putting the “learning” into learning-augmented algorithms for frequency estimation. In *Proceedings of the 38th International Conference on Machine Learning*, pages 2860–2869, 2021.
- F. Eberle, A. Lindermayr, N. Megow, L. Nölke, and J. Schlöter. Robustification of online graph exploration methods. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 9732–9740, 2022.
- T. Eden, P. Indyk, S. Narayanan, R. Rubinfeld, S. Silwal, and T. Wagner. Learning-based support estimation in sublinear time. In *9th International Conference on Learning Representations, ICLR, 2021*.
- T. Eden, P. Indyk, and H. Xu. Embeddings and labeling schemes for a. In *13th Innovations in Theoretical Computer Science Conference (ITCS 2022)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2022.
- P. Erdős, A. Rényi, et al. On the evolution of random graphs. *Publ. math. inst. hung. acad. sci.*, 5(1):17–60, 1960.
- D. Ferguson, M. Likhachev, and A. Stentz. A guide to heuristic-based path planning. In *Proceedings of the international workshop on planning under uncertainty for autonomous systems, international conference on automated planning and scheduling (ICAPS)*, pages 9–18, 2005.
- P. Ferragina and G. Vinciguerra. Learned data structures. In *Recent Trends in Learning From Data: Tutorials from the INNS Big Data and Deep Learning Conference (INNSBDDL2019)*, pages 5–41. Springer, 2020.
- D. Foead, A. Ghifari, M. B. Kusuma, N. Hanafiah, and E. Gunawan. A systematic literature review of A* pathfinding. *Procedia Computer Science*, 179:507–514, 2021.
- B. Gorain, K. Mondal, H. Nayak, and S. Pandit. Pebble guided optimal treasure hunt in anonymous graphs. *Theoretical Computer Science*, 922:61–80, 2022.
- A. Gupta, R. Krauthgamer, and J. R. Lee. Bounded geometries, fractals, and low-distortion embeddings. In *44th Annual IEEE Symposium on Foundations of Computer Science, 2003. Proceedings.*, pages 534–543. IEEE, 2003.
- A. Gupta, D. Panigrahi, B. Subercaseaux, and K. Sun. Augmenting online algorithms with ϵ -accurate predictions. *Advances in Neural Information Processing Systems*, 35:2115–2127, 2022.

- C. Hsu, P. Indyk, D. Katabi, and A. Vakilian. Learning-based frequency estimation algorithms. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*, 2019.
- P. Indyk, A. Vakilian, and Y. Yuan. Learning-based low-rank approximations. In *Advances in Neural Information Processing Systems*, pages 7400–7410, 2019.
- T. Jiang, Y. Li, H. Lin, Y. Ruan, and D. P. Woodruff. Learning-augmented data stream algorithms. In *International Conference on Learning Representations*, 2020.
- B. Kalyanasundaram and K. R. Pruhs. Constructing competitive tours from local information. *Theoretical Computer Science*, 130(1):125–138, 1994.
- A. R. Karlin, N. Klein, and S. O. Gharam. A (slightly) improved approximation algorithm for metric tsp. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 32–45, 2021.
- D. Komm, R. Kráľovič, R. Kráľovič, and J. Smula. Treasure hunt with advice. In *Structural Information and Communication Complexity: 22nd International Colloquium, SIROCCO 2015, Montserrat, Spain, July 14-16, 2015. Post-Proceedings 22*, pages 328–341. Springer, 2015.
- T. Kraska, A. Beutel, E. H. Chi, J. Dean, and N. Polyzotis. The case for learned index structures. In *Proceedings of the 2018 International Conference on Management of Data*, pages 489–504, 2018.
- Y. Li, H. Lin, S. Liu, A. Vakilian, and D. Woodruff. Learning the positions in counts sketch. In *The Eleventh International Conference on Learning Representations*, 2023.
- H. Lin, T. Luo, and D. Woodruff. Learning augmented binary search trees. In *International Conference on Machine Learning*, pages 13431–13440. PMLR, 2022.
- T. Lykouris and S. Vassilvitskii. Competitive caching with machine learned advice. In *International Conference on Machine Learning*, pages 3302–3311, 2018.
- N. Megow, K. Mehlhorn, and P. Schweitzer. Online graph exploration: New results on old and new algorithms. *Theoretical Computer Science*, 463:62–72, 2012.
- M. Mitzenmacher. A model for learned bloom filters and optimizing by sandwiching. In *Advances in Neural Information Processing Systems*, pages 464–473, 2018.
- M. Mitzenmacher and S. Vassilvitskii. Algorithms with predictions. *Communications of the ACM*, 65(7):33–35, 2022.
- P. Paliwal. A survey of a-star algorithm family for motion planning of autonomous vehicles. In *2023 IEEE International Students’ Conference on Electrical, Electronics and Computer Science (SCEECS)*, pages 1–6. IEEE, 2023.
- A. Pelc. Searching games with errors—fifty years of coping with liars. *Theoretical Computer Science*, 270(1-2):71–109, 2002.
- I. Pohl. Bi-directional and heuristic search in path problems. Technical report, Stanford Linear Accelerator Center, Calif., 1969.
- M. Purohit, Z. Svitkina, and R. Kumar. Improving online algorithms via ml predictions. In *Advances in Neural Information Processing Systems*, pages 9661–9670, 2018.
- L. H. O. Rios and L. Chaimowicz. A survey and classification of A* based best-first heuristic search algorithms. In *Brazilian Symposium on Artificial Intelligence*, pages 253–262. Springer, 2010.
- A. Wei and F. Zhang. Optimal robustness-consistency trade-offs for learning-augmented online algorithms. *Advances in Neural Information Processing Systems*, 33:8042–8053, 2020.
- A. C.-C. Yao. Probabilistic computations: Toward a unified measure of complexity. In *18th Annual Symposium on Foundations of Computer Science (sfcs 1977)*, pages 222–227. IEEE Computer Society, 1977.

Checklist

1. For all models and algorithms presented, check if you include:
 - (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model. [Yes : These are included in the main body of the paper when discussing the relevant results.]
 - (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm. [Yes : The summary results of the analysis are in the body of the paper, the formal proofs can be found in the supplementary material.]
 - (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries. [Yes: see uploaded code, along with README.md file. README.md outlines all dependencies.]

2. For any theoretical claim, check if you include:
 - (a) Statements of the full set of assumptions of all theoretical results. [Yes : These are included in the main body of the paper.]
 - (b) Complete proofs of all theoretical results. [Yes : These can be found in Section B of the supplementary material.]
 - (c) Clear explanations of any assumptions. [Yes : These are included in the main body of the paper.]
3. For all figures and tables that present empirical results, check if you include:
 - (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). [Yes: we provide both the code used to generate the figures, and data files for exactly recreating figures.]
 - (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). [Not Applicable : No model was trained for this paper.]
 - (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). [Yes: this is discussed in Section 6 and further detailed in Section D of the supplementary material.]
 - (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). [Yes: See Section D in the supplementary material.]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
 - (a) Citations of the creator If your work uses existing assets. [Not Applicable]
 - (b) The license information of the assets, if applicable. [Not Applicable]
 - (c) New assets either in the supplemental material or as a URL, if applicable. [Not Applicable]
 - (d) Information about consent from data providers/curators. [Not Applicable]
 - (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. [Not Applicable]
5. If you used crowdsourcing or conducted research with human subjects, check if you include:
 - (a) The full text of instructions given to participants and screenshots. [Not Applicable]
 - (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. [Not Applicable]
 - (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. [Not Applicable]

Supplementary Material

A SURVEY OF RELATED WORKS

Graph search with distance-to-goal predictions The problem and prediction settings considered in this work most closely correspond to those considered by Banerjee et al. (2023). Banerjee et al. (2023) study exploration and planning under absolute error models. They consider two parametrizations of prediction error: the first in terms of the number of nodes at which predictions are not equal to true distance-to-goal, denoted \mathcal{E}_0 , and the second in terms of the ℓ_1 norm of the vector of errors, denoted \mathcal{E}_1 as in this work. They develop an algorithm for exploration on unweighted trees, and prove guarantees on its performance in terms of \mathcal{E}_0 . They also develop algorithms for planning on graphs: on unweighted graphs, they establish guarantees parameterized by \mathcal{E}_0 , while in graphs with integer-valued distances their performance bounds are parameterized by \mathcal{E}_1 .

Treasure Hunt In the *treasure hunt problem*, a mobile agent must traverse some unknown environment, continuous or discrete, to locate a stationary hidden goal (Alpern and Gal, 2006). When the search environment is a graph, this problem shares many features with the exploration problem considered in this paper. Bouchard et al. (2021) study the graph treasure hunt problem when the searcher receives no additional information. They establish lower bounds on the total cost incurred by any algorithm in terms of the number of edges in the ball of radius OPT around the root node, and give algorithms with performance guarantees which asymptotically exceed these lower bounds by at most a factor of $\log(\text{OPT})$. Graph treasure hunt problems have also been considered when the agent receives help with the task. Komm et al. (2015) study the case when the searcher receives generic advice, which can take the form of any bit string. They consider the advice complexity of the treasure hunt task; they prove that there is an algorithm which achieves competitive ratio r by receiving $O(n/r)$ bits of advice along the search and moreover they establish that any algorithm achieving a competitive ratio of r must receive $\Omega(n/r)$ bits of advice (see Theorems 4 and 5 in Komm et al. (2015)). In the setting where graph vertices are anonymized, i.e. the searcher has no way to recognize whether a vertex has or has not previously been visited, recent work has studied the task of graph treasure hunt with access to advice from an omniscient oracle which marks vertices with binary labels (Bhattacharya et al., 2022; Gorain et al., 2022).

The exploration model studied in this paper can be viewed as a graph treasure hunt problem with specific kinds of advice (predictions of distance-to-goal). One main contrast with this work is that the advice considered can contain adversarial errors, and indeed the impact that different error models have on the graph exploration task is a core topic investigated in this work.

Path-Planning and A* Search Distance-to-goal predictions have been the subject of study in many path-planning and graph-traversal settings. Initialized with a root node and some known target node, *path-planning problems* seek to learn a shortest path between the root and goal and common problem models assume access to a set of distance-to-goal predictions, referred to as “heuristics” within this literature Pohl (1969); Ferguson et al. (2005). A* search is a celebrated algorithm for the problem of path-finding and graph-traversal, designed for cases when the entire graph G and all predictions f are accessible in memory, and has spawned many algorithmic variants (Rios and Chaimowicz, 2010; Foead et al., 2021; Paliwal, 2023).

Much of the theory of A* search focuses on cases when predictions have particular structural properties: a prediction function $f : V \rightarrow \mathbb{R}$ is called *admissible* if the prediction at every node v is never greater than the actual distance to the goal from v . A prediction function $f : V \rightarrow \mathbb{R}$ is *consistent* (or *monotone*) if for every node v and every neighbor u of v , $f(v) \leq f(u) + d(v, u)$ and $f(g) = 0$. Consistency is studied so heavily in part because it implies admissibility. Admissible heuristics are well-motivated and occur in various other problems. For example, Eden et al. (2022) consider access to an oracle that provides an underestimate for the probabilities of any element in some discrete probability distribution. In addition to being well-motivated by applications, admissibility is a focus of literature because A* search with admissible heuristics enjoys optimality properties (Dechter and Pearl, 1985).




Figure 4: Comparing Algorithm 1 versus A* search on a random tree with randomly generated errors. The same set of predictions is provided to both algorithms. While the set of nodes visited by the two algorithms is comparable, the computational model in A* places no penalty on traversal distance so that algorithm has a tendency to double-back on itself, leading to a more expensive tour. Nodes are colored by prediction value and labeled with the order in which they are first visited by the relevant algorithm. Tour cost is taken to be $\sum d(v_i, v_{i+1})$ for indices i ordered according to when a node is first visited: the tour cost for A* omits costs incurred by re-expanding a node within the execution of the algorithm.

While many of the problems and algorithms within path-planning may appear closely related to this work at first-glance, we emphasize that the goals and cost models differentiate the graph searching problems considered in this work from those in path-planning. In particular, the design and algorithmic guarantees on path-planning algorithms like A* search implicitly assume that the full graph and predictions are available in memory upon initialization of the algorithm. Performance guarantees and notions of optimality are proven in terms of computational procedures, rather than traversal distance. For example, when predictions are admissible A* is optimal in the sense that the set of nodes expanded (an operation analogous to visiting a node) is minimal (Dechter and Pearl, 1985). However, the sequence in which these nodes are expanded can incur high traversal distance, as illustrated in Figure 4. More generally, the goals of algorithms for path-planning differ substantially from that in the graph search problem: the path-finding problem seeks to learn and return a shortest path even if the process required to find such a path is expensive in the sense of graph traversal, whereas the aim of a graph searching problem is finding the goal node in an inexpensive manner and makes no demand that a shortest path from the root node to the goal be in the set of visited nodes upon termination of the algorithm. These differences in goals and algorithmic guarantees mean that in cases when the environment is unknown, i.e. when the graph and predictions are not available in memory but must instead be accessed by traversal, path-planning algorithms may incur much higher traversal cost than the graph search algorithms proposed in this paper.

Learning-Based Algorithms. Recent years have seen a marked increase in the integration of machine learning techniques to enhance traditional algorithmic challenges. Such algorithms have been developed for various topics including online algorithms Lykouris and Vassilvitskii (2018); Purohit et al. (2018); Angelopoulos et al. (2020); Wei and Zhang (2020); Bamas et al. (2020); Aamand et al. (2022); Antoniadis et al. (2023); Anand et al. (2020); Diakonikolas et al. (2021); Gupta et al. (2022), data structures Kraska et al. (2018); Mitzenmacher (2018); Ferragina and Vinciguerra (2020); Lin et al. (2022), and streaming models Hsu et al. (2019); Indyk et al. (2019); Aamand et al. (2019); Jiang et al. (2020); Cohen et al. (2020); Du et al. (2021); Eden et al. (2021); Chen et al. (2022); Li et al. (2023). For an extensive collection of learning-based algorithms, refer to the repository at <https://algorithms-with-predictions.github.io/>.

B MISSING PROOFS

In this section, we present detailed proofs of the results in the main body of the paper.

B.1 Proofs For Section 3

Proof of Theorem 1 and Corollary 5. Algorithm 1 follows a shortest path in G_i from v_{i-1} to v_i . A simple consequence of this is that:

$$\text{ALG} = \sum_{i \in [T]} d_{G_i}(v_{i-1}, v_i).$$

Let $\Delta_i \stackrel{\text{def}}{=} d_G(v_{i-1}, g) - d_G(v_i, g)$, and let T be the number of iterations of the while loop executed, so that $v_T = g$. We then have:

$$\sum_{i \in [T]} \Delta_i = d_G(v_0, g) - d_G(v_T, g) = d_G(v_0, g) - d_G(g, g) = \text{OPT}.$$

Consider any iteration $i \in [T]$. Let w_i be the first vertex outside of V_{i-1} encountered when traversing a shortest path from v_{i-1} to g . Note that $w_i \in \partial V_{i-1}$, and hence, by the update rule in Algorithm 1 we have:

$$f(v_i) + d_{G_i}(v_{i-1}, v_i) \leq f(w_i) + d_{G_i}(v_{i-1}, w_i). \quad (6)$$

Furthermore, by the definition of w_i , we have:

$$d_{G_i}(v_{i-1}, w_i) = d_G(v_{i-1}, w_i). \quad (7)$$

The above follows from a simple contradiction argument, for the existence of a shorter path from v_{i-1} to w_i in G which is not in G_i , would contradict the definition of w_i . We also have:

$$d_G(v_{i-1}, g) = d_G(v_{i-1}, w_i) + d_G(w_i, g). \quad (8)$$

We then have, for any $i \in [T]$:

$$\begin{aligned} d_G(v_i, g) - f(v_i) &\stackrel{(6)}{\geq} d_G(v_i, g) + d_{G_i}(v_{i-1}, v_i) - d_{G_i}(v_{i-1}, w_i) - f(w_i) \\ &\stackrel{(7)}{=} d_G(v_i, g) + d_{G_i}(v_{i-1}, v_i) - d_G(v_{i-1}, w_i) - f(w_i) \\ &= d_G(v_{i-1}, g) - \Delta_i + d_{G_i}(v_{i-1}, v_i) - d_G(v_{i-1}, w_i) - f(w_i) \\ &\stackrel{(8)}{=} d_{G_i}(v_{i-1}, v_i) - \Delta_i + d_G(w_i, g) - f(w_i). \end{aligned}$$

Note that the vertices in the sequence $\{v_i\}_{i \in [T] \cup \{0\}}$ are always distinct, while the vertices in the sequence $\{w_i\}_{i \in [T]}$ might not be. This also implies that $T \leq n$. The above then implies:

$$\begin{aligned} \text{ALG} &= \sum_{i \in [T]} d_{G_i}(v_{i-1}, v_i) \\ &\leq \sum_{i \in [T]} \Delta_i + d_G(v_i, g) - f(v_i) + f(w_i) - d_G(w_i, g) \\ &= \sum_{i \in [T]} \Delta_i + \sum_{i \in [T]} d_G(v_i, g) - f(v_i) + \sum_{i \in [T]} f(w_i) - d_G(w_i, g) \\ &\leq \text{OPT} + \mathcal{E}_1^- + T \cdot \mathcal{E}_\infty^+ \\ &\leq \text{OPT} + \mathcal{E}_1^- + n \cdot \mathcal{E}_\infty^+. \end{aligned}$$

This completes the proof of Theorem 1. When predictions are admissible, $\mathcal{E}_1^- = \mathcal{E}_1$ and $\mathcal{E}_\infty^+ = 0$ so Corollary 5 follows. \square

Proof of Theorem 6. We begin by proving the first half of the theorem. Given $\mathcal{E}_1^- > 0$ one considers the three-vertex graph path P_3 where the two edges are weighted with weight $w = \mathcal{E}^-/2$, the start vertex / root is chosen to be the middle vertex and the goal is one of the other two vertices (see left side of Figure 5). Note that the value of OPT is $d_G(r, g) = w$. When the predictions on the vertices are given by: $f(v_1) = 0$, $f(v_2) = w$ and




Figure 5: The construction in the proof of Theorem 6.

$f(v_3) = 0$, the error is equal to \mathcal{E}^- and the graph looks completely symmetric to the searcher, and hence in the worst case to find the goal, the searcher has to incur a cost of $3w = w + 2w = \text{OPT} + \mathcal{E}^-$ as needed.

For the second part of the theorem, we construct a star on n vertices, where each edge has weight $w = \mathcal{E}_\infty^+/2$, the starting vertex is at the center of the star, and the goal is chosen arbitrarily among the other vertices. The prediction at the goal is then picked to equal \mathcal{E}_∞^+ so that it equals the prediction in all other vertices, i.e. we set the predictions to $f(r) = w$ and $f(v) = 2w$ for all $w \neq r$. Every algorithm will then have to visit the entire star in the worst case, incurring a cost of $\mathcal{E}_\infty^+/2 + \mathcal{E}_\infty^+(n-2) = \text{OPT} + \mathcal{E}_\infty^+(n-2)$ as needed (See the right side of Figure 5). \square

Proof of Proposition 7. We apply Yao's minimax principle (Yao, 1977) to the same constructions used in the proof of Theorem 6. For both constructions, we consider the distribution over instances produced by choosing the goal node uniformly at random among the leaf nodes.

In particular, for the first statement, we consider the performance of any deterministic algorithm on the distribution of instances given by taking the three-node path graph on the left-hand side of Figure 5, and placing the goal g at either v_1 or v_3 with equal probability. We fix predictions $f(v_1) = f(v_3) = 0$ and $f(v_2) = w$ as in the proof of Theorem 6. The expected cost incurred by any deterministic algorithm over this distribution of instances is $2w = \text{OPT} + w = \text{OPT} + 1/2\mathcal{E}^-$. Yao's minimax principle then implies the stated lower bound for all randomized algorithms.

The proof of the second result follows analogously by considering the second construction in the proof of Theorem 6. \square

B.2 Proofs For Section 4

B.2.1 Algorithmic Guarantees Under Relative Error

We begin the analysis of Algorithm 2 by establishing the following properties of the set $S_{\varepsilon,r}$ defined in Equation (5). For G a tree, let $P_G(u,v)$ denote the (unique) shortest path between nodes u and v in G .

Lemma 14. *For $S_{\varepsilon,r}$ as defined in (5) and G a weighted tree, then the following hold:*

- (i) $\forall v \notin S_{\varepsilon,r}, d_G(v,r) > \text{OPT}$,
- (ii) $\forall v \in S_{\varepsilon,r}, d_G(v,r) \leq \frac{1+\varepsilon}{1-\varepsilon} \cdot \text{OPT}$,
- (iii) $\forall v \in S_{\varepsilon,r}, d_G(v,g) \leq \frac{2}{1-\varepsilon} \cdot \text{OPT}$,
- (iv) $P_G(r,g) \subseteq S_{\varepsilon,r}$,
- (v) For any $v \in S_{\varepsilon,r}, P_G(v,g) \subseteq S_{\varepsilon,r}$.

Proof. The properties follow immediately from the definition of the relative error model in Equation (1) and the definition of $S_{\varepsilon,r}$ in Equation (5).

- (i) $\forall v \notin S_{\varepsilon,r}, f(v) > \frac{1}{1-\varepsilon}f(r)$ and by Equation (1), $f(r) \geq (1-\varepsilon)d_G(r,g) = (1-\varepsilon)\text{OPT}$.
- (ii) $\forall v \in S_{\varepsilon,r}, d_G(v,r) \leq \frac{1}{1-\varepsilon}f(r)$, and by Equation (1), $f(r) \leq (1+\varepsilon)\text{OPT}$.
- (iii) By triangle inequality, for any $v \in G$

$$d_G(v,g) \leq d_G(v,r) + d_G(r,g) = d_G(v,r) + \text{OPT}$$

and so by property (ii) above, for any $v \in S_{\varepsilon,r}$ the result follows.

- (iv) $\forall v \in P_G(r,g), d_G(r,v) \leq d_G(r,g) = \text{OPT}$. Thus property (i) above implies $\forall v \in P_G(r,g), v \in S_{\varepsilon,r}$.
- (v) Consider $v \in S_{\varepsilon,r}$, and let $u \stackrel{\text{def}}{=} \text{argmin}\{d_G(u,g) \mid u \in P_G(v,r)\}$. Because G is a tree,

$$P_G(v,g) = P_G(v,u) \cup P_G(u,g).$$

In particular, $P_G(v,u) \subseteq P_G(v,r)$; observe that because $v \in S_{\varepsilon,r}, \forall w \in P_G(v,r), d_G(w,r) \leq d_G(v,r) \leq \frac{1}{1-\varepsilon}f(r)$ and thus by the definition of $S_{\varepsilon,r}, P_G(v,r) \subseteq S_{\varepsilon,r}$. We've thus concluded that $P_G(v,u) \subseteq S_{\varepsilon,r}$.

For the second portion of the path, $P_G(u,g)$, observe that by the definition of $u = \text{argmin}\{d_G(u,g) \mid u \in P_G(v,r)\}$, it must be that $u \in P_G(r,g)$. Thus $P_G(u,g) \subseteq P_G(r,g)$ and so by property (iv), $P_G(u,g) \subseteq S_{\varepsilon,r}$.

□

With these properties, we now bound the distance travelled on the i th step of the algorithm:

Lemma 15. *Let G be a weighted tree, with predictions satisfying Equation (1) with respect to parameter $\varepsilon < 1$. Then, the distance travelled by Algorithm 2 on the i th iteration satisfies*

$$(1-\varepsilon)d_{G_i}(v_{i-1}, v_i) \leq \Delta_i + 2\varepsilon d_G(v_i, g). \quad (9)$$

Additionally, if the predictions on G are multiplicative and decremental,

$$d_{G_i}(v_{i-1}, v_i) \leq \Delta_i + \varepsilon d_G(v_i, g). \quad (10)$$

Proof of Lemma 15. We observe that, by definition of Algorithm 2, for all iterations $i, V_{i-1} \subset S_{\varepsilon,r}$. Consider w_i the first vertex outside of V_{i-1} encountered when traversing $P(v_{i-1}, g)$, the shortest path from v_{i-1} to g . Note that $w_i \in \partial V_{i-1}$, and that by property (v) of Lemma 14, $w_i \in S_{\varepsilon,r}$.

Thus, on iteration $i, \exists w_i \in \partial V_{i-1} \cap S_{\varepsilon,r} \cap P(v_{i-1}, g)$, so by the definition of Algorithm 2 such w_i satisfies

$$f(v_i) + d_{G_i}(v_{i-1}, v_i) \leq f(w_i) + d_{G_i}(v_{i-1}, w_i).$$

In particular, because $w_i \in P(v_{i-1}, g) \cap \partial V_{i-1}$ and by the tree properties of G , we have

$$d_{G_i}(v_{i-1}, w_i) = d_G(v_{i-1}, w_i) = d_G(v_{i-1}, g) - d_G(w_i, g).$$

We can thus upper bound

$$\begin{aligned} d_{G_i}(v_{i-1}, v_i) &\leq d_G(v_{i-1}, g) - d_G(w_i, g) + f(w_i) - f(v_i) \\ &= d_G(v_{i-1}, g) - d_G(v_i, g) + (d_G(v_i, g) - f(v_i)) + (f(w_i) - d_G(w_i, g)) \\ &= \Delta_i - \varepsilon_{v_i} d_G(v_i, g) + \varepsilon_{w_i} d_G(w_i, g) \end{aligned}$$

where $\varepsilon_{v_i}, \varepsilon_{w_i} \in [-\varepsilon, \varepsilon]$ are the constants whose existence is implied by Equation (1) such that

$$f(v) = (1 + \varepsilon_v) d_G(v, g). \quad (11)$$

Consider two cases: first, consider the case when $\varepsilon_{w_i} > 0$. Because $w_i \in P(v_{i-1}, g)$,

$$d_G(w_i, g) \leq d_G(v_{i-1}, g) \leq d_G(v_{i-1}, v_i) + d_G(v_i, g).$$

Combining this bound and the fact that $\varepsilon_{v_i}, \varepsilon_{w_i} \in [-\varepsilon, \varepsilon]$ yields:

$$-\varepsilon_{v_i} d_G(v_i, g) + \varepsilon_{w_i} d_G(w_i, g) \leq 2\varepsilon d_G(v_i, g) + \varepsilon d_G(v_{i-1}, v_i).$$

Thus in this case, the desired bound in (9) is satisfied.

In the second case, when $\varepsilon_{w_i} \leq 0$,

$$-\varepsilon_{v_i} d_G(v_i, g) + \varepsilon_{w_i} d_G(w_i, g) \leq -\varepsilon_{v_i} d_G(v_i, g) \leq \varepsilon d_G(v_i, g)$$

which is trivially upper bounded by $2\varepsilon d_G(v_i, g) + \varepsilon d_G(v_{i-1}, v_i)$. Thus, in both cases, the desired bound in (9) is satisfied.

In the case of decremental errors, $\varepsilon_v \leq 0 \forall v \in V$. Thus the latter case always applies, so we can bound

$$d_{G_i}(v_{i-1}, v_i) \leq \Delta_i - \varepsilon_{v_i} d_G(v_i, g) + \varepsilon_{w_i} d_G(w_i, g) \leq \Delta_i + \varepsilon d_G(v_i, g)$$

thus establishing the bound in (10). \square

We are now equipped to prove Theorem 2.

Proof of Theorem 2. We will show that Algorithm 2 achieves the desired competitive ratio. We begin by establishing that the algorithm will terminate at the goal node g : by Property (iv) in Lemma 14, $P(r, g) \subseteq S_{\varepsilon, r}$, and further by definition $P(r, g)$ is connected, so Algorithm 2 initialized at r will explore a connected subgraph of G that contains g , and will thus terminate at g .

We now bound the total distance travelled by Algorithm 2. Consider the general case, when errors can be incremental or decremental. Then, by Lemma 15,

$$\begin{aligned} \text{ALG} &= \sum_{i \in [T]} d_{G_i}(v_{i-1}, v_i) = \sum_{i \in [T]} (1 - \varepsilon) d_{G_i}(v_{i-1}, v_i) + \varepsilon d_{G_i}(v_{i-1}, v_i) \\ &\leq \sum_{i \in [T]} \Delta_i + \sum_{i \in [T]} 2\varepsilon d_G(v_i, g) + \sum_{i \in [T]} \varepsilon d_G(v_{i-1}, v_i) = \text{OPT} + 2\varepsilon \left(\sum_{i \in [T]} d_G(v_i, g) \right) + \varepsilon \text{ALG}. \end{aligned}$$

Because $V_i \subseteq S_{\varepsilon, r}$ for all iterations i , and by property (iii) in Lemma 14,

$$2\varepsilon \sum_{i \in [T]} d_G(v_i, g) \leq 2\varepsilon |S_{\varepsilon, r}| \cdot \frac{2}{1 - \varepsilon} \text{OPT}.$$

Thus, re-arranging,

$$(1 - \varepsilon) \text{ALG} \leq \text{OPT} \left(1 + |S_{\varepsilon, r}| \varepsilon \cdot \frac{4}{1 - \varepsilon} \right),$$

yielding the claimed competitive ratio from the trivial upper bound $|S_{\varepsilon, r}| \leq n$.

In the case of decremental errors, Lemma 15 and a similar argument give

$$\text{ALG} = \sum_{i \in [T]} d_{G_i}(v_{i-1}, v_i) \leq \sum_{i \in [T]} \Delta_i + \sum_{i \in [T]} \varepsilon d_G(v_i, g) \leq \text{OPT} \left(1 + |S_{\varepsilon, r}| \varepsilon \cdot \frac{2}{1 - \varepsilon} \right),$$

yielding the claimed competitive ratio. \square

To prove Theorem 3, we'll use the following lemmas: the first (Lemma 16) establishes that certain algorithms never explore nodes too far from g . We emphasize that the below lemma makes use of distances in the full G , not G_i . In the case of weighted trees, these two distances are always identical: on a tree, for all iterations i and for any $u, v \in V_i \cup \partial V_i$, $d_{G_i}(u, v) = d_G(u, v)$. The second lemma (Lemma 17) bounds the distance travelled by these algorithms on any given iteration.

Lemma 16. Consider the exploration problem on G a weighted, undirected graph with predictions f satisfying Equation (1). Consider the update rule used in Algorithm 3:

$$v_i = \operatorname{argmin}_{v \in \partial V_{i-1}} \{\beta d_G(v_{i-1}, v) + f(v)\}, \quad (12)$$

and assume $\beta > 0$ satisfies $\beta < 1 - \varepsilon$. Then, for every iteration $i \in [T]$, the node v_i visited by the algorithm on the i^{th} iteration satisfies

$$d_G(v_i, g) \leq \frac{1 + \varepsilon + \beta}{1 - (\varepsilon + \beta)} \text{OPT}.$$

Proof of Lemma 16. Let r_i be the first vertex outside of V_{i-1} encountered when traversing $P_G(r, g)$ a shortest path from the root r to g . Note that $r_i \in \partial V_{i-1}$, and thus by Equation (12),

$$\beta d_G(v_{i-1}, v_i) + f(v_i) \leq \beta d_G(v_{i-1}, r_i) + f(r_i).$$

In particular, by the triangle inequality we can bound

$$\begin{aligned} f(v_i) &\leq \beta (d_G(v_{i-1}, r_i) - d_G(v_{i-1}, v_i)) + f(r_i) \\ &\leq \beta d_G(r_i, v_i) + f(r_i) \\ &\leq \beta (d_G(r_i, g) + d_G(v_i, g)) + f(r_i). \end{aligned}$$

Given $r_i \in P_G(r, g)$, $d_G(r_i, g) \leq d_G(r, g) = \text{OPT}$. Moreover, recalling that the predictions f must satisfy Equation (1), let ε_{v_i} and ε_{r_i} be the relative errors at vertex v_i and r_i respectively (defined as in Equation (11)). Then we can rewrite the bound above as

$$(1 + \varepsilon_{v_i})d_G(v_i, g) \leq \beta (d_G(r_i, g) + d_G(v_i, g)) + (1 + \varepsilon_{r_i})d_G(r_i, g),$$

and hence:

$$(1 + \varepsilon_{v_i} - \beta)d_G(v_i, g) \leq (\beta + 1 + \varepsilon_{r_i})\text{OPT}.$$

Using the fact that $\varepsilon_{v_i}, \varepsilon_{r_i} \in [-\varepsilon, \varepsilon]$ and the assumption that β satisfies $\beta < 1 - \varepsilon$, we can use the above bound to conclude

$$d_G(v_i, g) \leq \frac{1 + \varepsilon + \beta}{1 - (\varepsilon + \beta)} \text{OPT}.$$

In particular, this holds on any iteration independently of i , so we obtain the desired result. \square

Lemma 17. Consider the exploration problem on G a weighted, undirected graph with predictions f satisfying Equation (1). Assume $\beta > 0$ satisfies

$$\frac{1 + \varepsilon}{2} < \beta < 1 - \varepsilon. \quad (13)$$

Then the distance traversed by Algorithm 3 on the i^{th} iteration is bounded by

$$d_{G_i}(v_{i-1}, v_i) \leq \frac{\beta}{2\beta - 1 - \varepsilon} \Delta_i + \frac{2\varepsilon}{2\beta - 1 - \varepsilon} d_G(v_i, g).$$

Proof of Lemma 17. Let $P_G(u, v)$ denote an (arbitrary) shortest path between nodes u and v in G . Algorithm 3 follows a shortest path in G_i from v_{i-1} to v_i . Let w_i be the first vertex outside of V_{i-1} encountered when traversing $P_G(v_{i-1}, g)$. Note that as a consequence, $w_i \in \partial V_{i-1}$, and hence by (12) we have

$$\beta d_{G_i}(v_{i-1}, v_i) + f(v_i) \leq \beta d_{G_i}(v_{i-1}, w_i) + f(w_i).$$

Since $w_i \in P_G(v_{i-1}, g)$,

$$d_{G_i}(v_{i-1}, w_i) = d_G(v_{i-1}, w_i) = d_G(v_{i-1}, g) - d_G(w_i, g).$$

Re-arranging and using the above fact, we obtain

$$\begin{aligned} \beta d_{G_i}(v_{i-1}, v_i) &\leq \beta d_{G_i}(v_{i-1}, w_i) + f(w_i) - f(v_i) \\ &= \beta (d_G(v_{i-1}, g) - d_G(w_i, g)) + f(w_i) - f(v_i) \end{aligned}$$

$$= \beta(d_G(v_{i-1}, g) - d_G(v_i, g)) + (\beta d_G(v_i, g) - f(v_i)) + (f(w_i) - \beta d_G(w_i, g)).$$

Let $\varepsilon_{v_i}, \varepsilon_{w_i}$ be the relative prediction errors at v_i and w_i respectively (defined as in Equation (11)), and recall the definition of Δ_i in (3). Then we can rewrite the above as

$$\beta d_{G_i}(v_{i-1}, v_i) \leq \beta \Delta_i + (\beta - 1 - \varepsilon_{v_i})d_G(v_i, g) + (1 + \varepsilon_{w_i} - \beta)d_G(w_i, g).$$

Because $w_i \in P_G(v_{i-1}, g)$, we have that

$$d_G(v_{i-1}, w_i) + d_G(w_i, g) = d_G(v_{i-1}, g) \leq d_G(v_{i-1}, v_i) + d_G(v_i, g).$$

Moreover, by upper bound on β in (13), $(1 + \varepsilon_{w_i} - \beta) \geq 0$, so we can revise our upper bound:

$$\beta d_{G_i}(v_{i-1}, v_i) \leq \beta \Delta_i + (\beta - 1 - \varepsilon_{v_i})d_G(v_i, g) + (1 + \varepsilon_{w_i} - \beta)(d_G(v_{i-1}, v_i) + d_G(v_i, g)).$$

Re-arranging and recalling $\varepsilon_{v_{i-1}}, \varepsilon_{w_i} \in [-\varepsilon, \varepsilon]$ yields

$$(2\beta - 1 - \varepsilon)d_{G_i}(v_{i-1}, v_i) \leq \beta \Delta_i + 2\varepsilon d_G(v_i, g)$$

Leveraging the lower bound on β in (13), we can divide to obtain the desired result. \square

Theorem 3 follows immediately from Lemmas 16 and 17:

Proof of Theorem 3. Observe that for $\varepsilon \in (0, 1/3)$, $\beta = 2/3$ always satisfies Equation (13), independently of the value of ε . Thus for this setting, Lemma 17 implies that the update cost on a single iteration of Algorithm 3 is bounded as

$$d_{G_i}(v_{i-1}, v_i) \leq \frac{2}{1-3\varepsilon} \Delta_i + \frac{6\varepsilon}{1-3\varepsilon} d_G(v_i, g). \quad (14)$$

In particular, for G a weighted tree, on all iterations i , for all $u, v \in V_i \cup \partial V_i$,

$$d_{G_i}(u, v) = d_G(u, v).$$

This, in combination with the choice of $\beta = 2/3$ implies that Lemma 16 applies, so for all vertices v_i visited by the algorithm, we have

$$d_G(v_i, g) \leq \frac{5+3\varepsilon}{1-3\varepsilon} \text{OPT}. \quad (15)$$

We can then upper bound the last term in Equation (14) and obtain:

$$d_{G_i}(v_{i-1}, v_i) \leq \frac{2}{1-3\varepsilon} \Delta_i + \frac{6\varepsilon}{1-3\varepsilon} \cdot \frac{5+3\varepsilon}{1-3\varepsilon} \text{OPT}$$

Thus, letting T denote the total number of iterations made by Algorithm 3, summing over all iterations yields

$$\text{ALG} = \sum_{i \in [T]} d_{G_i}(v_{i-1}, v_i) \leq \frac{2}{1-3\varepsilon} \sum_{i \in [T]} \Delta_i + \frac{6\varepsilon}{1-3\varepsilon} \cdot \frac{5+3\varepsilon}{1-3\varepsilon} \text{OPT} \cdot T.$$

In particular, $\sum_{i \in [T]} \Delta_i = \text{OPT}$, and as every iteration i must end at some distinct v_i satisfying (15),

$$T \leq \left| B \left(g, \frac{5+3\varepsilon}{1-3\varepsilon} \text{OPT} \right) \right| \leq n.$$

We then have:

$$\begin{aligned} \text{ALG} &\leq \frac{2}{1-3\varepsilon} \text{OPT} + \frac{6\varepsilon}{1-3\varepsilon} \cdot \frac{5+3\varepsilon}{1-3\varepsilon} \text{OPT} \cdot T \\ &\leq \text{OPT} \left(2 + \frac{6\varepsilon}{1-3\varepsilon} + \frac{6\varepsilon}{1-3\varepsilon} \cdot \frac{5+3\varepsilon}{1-3\varepsilon} \cdot n \right). \end{aligned}$$

Giving the result in the statement of the theorem. \square

B.2.2 Lower bounds For Exploration With Relative Error

Proof of Theorem 8. We consider a star in which every edge has weight w_1 , to this, we add a new vertex g connected to one of the outside vertices by an edge of weight w_2 . We consider the case when the initial position r is the central node of the star. In this case, the optimal algorithmic cost is

$$\text{OPT} = d_G(r, g) = w_1 + w_2.$$

For the given $\varepsilon \in (0, 1)$, consider choice of w_1 and w_2 such that

$$\varepsilon = \frac{w_1}{w_1 + w_2}.$$

Note that in particular, such a setting of weights allows for the following predictions: every node except for the root and the goal can have prediction

$$f(v) = (1 - \varepsilon)(2w_1 + w_2).$$

For the above setting of w_1 and w_2 , this satisfies Equation 1 with respect to ε . In particular, for a searcher starting at the root, all neighbors of the root appear identical.

In this error regime, every algorithm for the exploration problem has to explore all branches of the star before finding g in the worst-case. Thus any algorithm has to incur cost at least

$$\text{ALG} = 2w_1 \cdot (n - 3) + (w_1 + w_2) = w_1(2(n - 3) + 1) + w_2.$$

The competitive ratio in this instance is thus

$$\frac{\text{ALG}}{\text{OPT}} = (2(n - 3) + 1) \frac{w_1}{w_1 + w_2} + \frac{w_2}{w_1 + w_2} = (2(n - 3) + 1)\varepsilon + (1 - \varepsilon) = \Theta(1 + n\varepsilon).$$

□

Proof of Proposition 9. Consider a distribution over instances obtained by taking the instance defined in the proof of Theorem 8, selecting a neighbor v of the root vertex r uniformly at random, and replacing the edge incident to the goal vertex g with an edge gv of weight w_2 .

Any deterministic algorithm for the exploration problem running on an instance sampled from this distributions incurs expected cost at least:

$$\text{ALG} = (n - 3) \cdot \frac{1}{2}2w_2 + \text{OPT} = ((n - 3)\varepsilon + 1)\text{OPT} \geq \left(1 + \frac{n\varepsilon}{2}\right)\text{OPT}.$$

The lower bound for randomized algorithms then follows from applying Yao's minimax principle (Yao, 1977). □

B.3 Proofs For Section 5

Throughout this subsection, we will use the following properties of ϕ_0 and ϕ_1 established by Banerjee et al. (2023):

Lemma 18. *Corollary 5.4 and Lemma 5.10 in Banerjee et al. (2023).* Given G an unweighted graph, for any $u, v \in G$,

$$\phi_0(u) + \phi_0(v) \geq d_G(u, v).$$

For G weighted,

$$\phi_1(u) + \phi_1(v) \geq 2d_G(u, v).$$

B.3.1 Planning Bounds Via Metric Embeddings

The distortion of an embedding can be related to its Lipschitz constant and that of its inverse: the Lipschitz constant of τ is defined as:

$$\|\tau\|_{\text{Lip}} \stackrel{\text{def}}{=} \max_{x_1, x_2 \in X} \frac{d_Y(\tau(x_1), \tau(x_2))}{d_X(x_1, x_2)}.$$

Note that any map with non-trivial distortion must be injective, and thus considering $\tau^{-1} : Y \rightarrow X$,

$$\text{dist}(\tau) = \|\tau\|_{\text{Lip}} \cdot \|\tau^{-1}\|_{\text{Lip}}.$$

To prove Lemma 10, we'll use the following fact to relate tours in G to tours in some embedding.

Lemma 19. *Consider an embedding $\tau : G \rightarrow G'$ for $G = (V, E)$ and $G' = (V', E')$. Then for any $S \subseteq V$,*

$$\text{tour}_G(S) \leq \|\tau^{-1}\|_{\text{Lip}} \cdot \text{tour}_{G'}(\tau(S)).$$

Proof of Lemma 19. Recall the definition of $\text{tour}_G(S)$ given in Equation (2):

$$\text{tour}_G(S) \stackrel{\text{def}}{=} \max_{v \in S} \min_{W \in \mathcal{W}(v, S)} \text{length}_G(W),$$

where $\mathcal{W}(v, S)$ is the set of walks in G starting at vertex v and visiting every vertex in S . Consider any walk $W = (u_1, \dots, u_k)$ in G' starting at some $u_1 \in \tau(S)$ and visiting all of $\tau(S)$. Let $W' = (u'_1, \dots, u'_{k'})$ be the subsequence of W containing only the points in $\tau(S)$. Note that W' contains all of the points in $\tau(S)$. We have:

$$\text{length}_{G'}(W) = \sum_{i=1}^{k-1} d_{G'}(u_i, u_{i+1}) \geq \sum_{i=1}^{k'-1} d_{G'}(u'_i, u'_{i+1}) \quad (16)$$

$$\geq \sum_{i=1}^{k'-1} \frac{1}{\|\tau^{-1}\|_{\text{Lip}}} \cdot d_G(\tau^{-1}(u'_i), \tau^{-1}(u'_{i+1})). \quad (17)$$

So, letting W'' be the walk visiting the vertices $(\tau^{-1}(u'_i))_{i=1}^{k'}$ in order while walking the shortest path in G between them. We then have:

$$\text{length}_{G'}(W) \geq \sum_{i=1}^{k'-1} \frac{1}{\|\tau^{-1}\|_{\text{Lip}}} \cdot d_G(\tau^{-1}(u'_i), \tau^{-1}(u'_{i+1})) = \frac{1}{\|\tau^{-1}\|_{\text{Lip}}} \cdot \text{length}_G(W'').$$

In particular, for any starting point $v \in S$ and any walk in $W \in \mathcal{W}(\tau(v), \tau(S))$ there exists some walk $W'' \in \mathcal{W}(v, S)$ such that:

$$\text{length}_G(W'') \leq \|\tau^{-1}\|_{\text{Lip}} \cdot \text{length}_{G'}(W),$$

so that, for every $v \in S$:

$$\min_{W \in \mathcal{W}(v, S)} \text{length}_G(W) \leq \|\tau^{-1}\|_{\text{Lip}} \min_{W \in \mathcal{W}(\tau(v), \tau(S))} \text{length}_{G'}(W) \leq \|\tau^{-1}\|_{\text{Lip}} \max_{u \in \tau(S)} \min_{W \in \mathcal{W}(u, \tau(S))} \text{length}_{G'}(W),$$

and hence:

$$\max_{v \in S} \min_{W \in \mathcal{W}(v, S)} \text{length}_G(W) \leq \|\tau^{-1}\|_{\text{Lip}} \max_{u \in \tau(S)} \min_{W \in \mathcal{W}(u, \tau(S))} \text{length}_{G'}(W),$$

completing the proof. \square

We now use Lemma 19 to establish Lemma 10.

Proof of Lemma 10. Given a real-valued function $f : V \rightarrow \mathbb{R}$, we denote the sublevel set of f about threshold c as

$$L_f^-(c) \stackrel{\text{def}}{=} \{v \in V : f(v) \leq c\}.$$

For the first part of the result, let G be an unweighted graph and consider the sublevel set $L_{\phi_0}^-(\lambda)$. By definition of the Lipschitz constant of $\tau : G \rightarrow G'$, for all $u, v \in G$

$$d_{G'}(\tau(u), \tau(v)) \leq \|\tau\|_{\text{Lip}} d_G(u, v).$$

Thus by Lemma 18, the embedding of the sublevel set has bounded diameter: let $u, v \in L_{\phi_0}^-(\lambda)$ such that $\text{diam}(\tau(L_{\phi_0}^-(\lambda))) = d_{G'}(\tau(u), \tau(v))$. Then

$$\text{diam}(\tau(L_{\phi_0}^-(\lambda))) = d_{G'}(\tau(u), \tau(v)) \leq \|\tau\|_{\text{Lip}} d_G(u, v) \leq \|\tau\|_{\text{Lip}} (\phi_0(u) + \phi_0(v)) \leq \|\tau\|_{\text{Lip}} \cdot 2\lambda.$$

Using this result along with the bound from Lemma 19 and the assumption that G' is $c_{G'}$ easily-tourable, we can bound

$$\begin{aligned} \text{tour}_G(L_{\phi_0}^-(\lambda)) &\leq \|\tau^{-1}\|_{\text{Lip}} \cdot \text{tour}_{G'}(\tau(L_{\phi_0}^-(\lambda))) \\ &\leq \|\tau^{-1}\|_{\text{Lip}} \cdot c_{G'} \text{diam}(\tau(L_{\phi_0}^-(\lambda))) \\ &\leq \|\tau^{-1}\|_{\text{Lip}} \cdot c_{G'} \cdot 2\lambda \|\tau\|_{\text{Lip}} \\ &= 2\rho c_{G'} \lambda. \end{aligned}$$

We now use this bound to analyze the cost of Algorithm 4. Algorithm 4 sequentially visits sublevel sets of ϕ_0 . On an iteration k corresponding to threshold λ_k , the algorithm visits each node in $L_{\phi_0}^-(\lambda_k) \subseteq V$ by computing a constant-factor approximation to the following problem: for $\{v_1, \dots, v_{|L_{\phi_0}^-(\lambda_k)|}\}$ the nodes of $L_{\phi_0}^-(\lambda_k)$ and $\Pi(n)$ the set of permutations on integers $1, \dots, n$, the algorithm computes

$$\sigma_{\lambda_k}^* \stackrel{\text{def}}{=} \underset{\sigma \in \Pi(|L_{\phi_0}^-(\lambda_k)|)}{\text{argmin}} \sum_{i=1}^{|L_{\phi_0}^-(\lambda_k)|} d_G(v_{\sigma(i)}, v_{\sigma(i+1)}).$$

The total distance travelled on the iteration k is thus

$$\sum_{i=1}^{|L_{\phi_0}^-(\lambda_k)|} d_G(v_{\sigma^*(i)}, v_{\sigma^*(i+1)}) \leq \max_{v \in L_{\phi_0}^-(\lambda_k)} \min_{W \in \mathcal{W}(v, L_{\phi_0}^-(\lambda_k))} \text{length}_G(W) = \text{tour}_G(L_{\phi_0}^-(\lambda_k)).$$

Algorithm 4 begins with $\lambda_0 = 1$ and doubles the threshold on each iteration, such that $\lambda_k = 2^k$. In particular, $\phi_0(g) = \mathcal{E}_0$, so the algorithm is guaranteed to terminate by the time it has visited every node of $L_{\phi_0}^-(\lambda_k)$ for the first sufficiently large threshold $\lambda_k \geq \mathcal{E}_0$. The algorithmic cost can thus be bounded as

$$\begin{aligned} \text{ALG} &\leq d_G(r, L_{\phi_0}^-(1)) + \sum_{k=0}^{\lceil \log_2(\mathcal{E}_0) \rceil} \text{tour}_G(L_{\phi_0}^-(2^k)) \\ &\leq d_G(r, L_{\phi_0}^-(1)) + 2\rho c_{G'} \sum_{k=0}^{\lceil \log_2(\mathcal{E}_0) \rceil} 2^k \\ &\leq d_G(r, L_{\phi_0}^-(1)) + 2\rho c_{G'} (4\mathcal{E}_0 - 1). \end{aligned}$$

Using Lemma 18, we can bound the transition cost from r to the first sublevel set as

$$d_G(r, L_{\phi_0}^-(1)) \leq d_G(r, g) + d_G(g, L_{\phi_0}^-(1)) \leq \text{OPT} + (\phi_0(g) + 1) = \text{OPT} + \mathcal{E}_0 + 1.$$

Combining these yields

$$\text{ALG} \leq \text{OPT} + \mathcal{E}_0(8\rho c_{G'} + 1) - 2\rho c_{G'} + 1 = \text{OPT} + O(\rho c_{G'} \mathcal{E}_0),$$

as desired.

For the second part of the result, let G be an graph with integer-valued distances consider the sublevel set $L_{\phi_1}^-(\lambda)$, and note that Lemma 19 holds for both weighted and unweighted graphs. The proof then follows analogously to the above argument, using the appropriate bound relating ϕ_1 to distances in G from Lemma 18. \square

B.3.2 Lower Bounds For Planning Problems




Figure 6: Reproduction of Figure 1. The lower bound construction for the proof of Lemma 11

Proof of Lemma 11. We construct a family of graphs with uniform predictions and analyze the worst-case cost incurred by any algorithm for the planning problem.

For a given Δ and $W \in \mathbb{N}$ let $G_{\Delta, W}$ be following graph: consider a root node r with Δ child nodes v_1, \dots, v_Δ . Let every edge (r, v_i) have edge weight 1. Each child node v_i then has $\Delta - 1$ descendants $u_1^i, \dots, u_{\Delta-1}^i$ with edge weights 1 for each edge (v_i, u_j^i) . Each of these descendants has a single child node w_j^i to which u_j^i is attached with edge weight W . This construction is illustrated in Figure 1.

We consider the planning problem when the searcher is initialized at the root node r described above, and the goal is the leaf node w_1^1 . We consider the case when error in the predictions is such that each subtree rooted at v_i appears to have the same predictions. In this construction, predictions are equal to true distance-to-goal for all nodes which are descendants of v_1 for $i \neq 1$, and error is allocated only over descendants of v_1 . We first calculate the total error in such predictions:

$$\mathcal{E}_1 = |f(v_1) - d_G(v_1, g)| + |f(g)| + \sum_{j=2}^{\Delta-1} |f(u_j^1) - d_G(u_j^1, g)| + |f(w_j^1) - d_G(w_j^1, g)| = 2W + 6\Delta + 4.$$

Under these predictions, all nodes on a given level from the root appear identical to the searcher. As a result, in the worst-case any algorithm for this problem must visit every node and incur cost ALG at least as large as the shortest tour of the graph starting at r and ending at g . Hence:

$$\text{ALG} \geq W(2\Delta^2 - 2\Delta - 1) + 2(\Delta^2 - 1).$$

On the other hand, we have $\text{OPT} = d_G(r, g) = W + 2$. This gives:

$$\text{ALG} - \text{OPT} \geq 2W(\Delta^2 - \Delta - 1) + 2(\Delta^2 - 2).$$

In order to understand how this cost scales with our parameters of interest, we now establish bounds on the doubling constant of G that show that $\lambda = \Theta(\Delta^2)$. Recall that the doubling constant is defined to be the

minimum value of λ such that for any radius R , any ball of radius R can be covered with at most λ many balls of radius $R/2$. Observe that the number of nodes n always upper bounds the doubling constant, so $\lambda \leq 2\Delta^2 + \Delta + 1$. To lower bound the doubling constant, consider the ball of radius $W + 2$ centered at the root node r , and note that this ball contains the entire graph. For W large ($W \geq 4$), $\Delta^2 + 1$ many balls of radius $W/2 + 1$ are required to cover the nodes of the graph, hence $\Delta^2 - 1 \leq \lambda$.

The fact that $a \geq 1$ then follows from observing that λ is independent of W while $\text{ALG} - \text{OPT} = \Omega(W\Delta^2)$. The fact that if $a = 1$, $b \geq 1$ similarly follows from the above along with $\mathcal{E}_1 = \Theta(W + \Delta)$. \square

Proof of Lemma 13. To establish that $a \geq 1$ and that $a + b \geq 3$, we consider the same construction outlined in the proof of Lemma 11 above. The same argument implies $a \geq 1$. Setting $W = \Delta$ yields a family of problem instances on integer-weighted trees for which $\text{ALG} - \text{OPT} = \Omega(\Delta^3)$ where $\mathcal{E}_1 = \Theta(\Delta)$. Thus on this family of instances any algorithm which is guaranteed to incur cost $\text{ALG} - \text{OPT} = O(\mathcal{E}_1^a \Delta^b)$ must have $a + b \geq 3$.

To establish that $b \geq 1$, we consider a construction in which \mathcal{E}_1 scales independently of Δ . Consider the weighted star with edge weights w , and assume the searcher is initialized at the central root node r and the goal node g is a leaf node as illustrated on the right side of Figure 5. We consider the same predictions constructed in the proof of Theorem 6: all nodes except for g have $f(v) = d_G(v, g)$, and $f(g) = 2w$ so that predictions at all leaf nodes appear uniform. Then in the worst case, the searcher must visit every node in the graph, incurring traversal cost

$$\text{ALG} = w(2\Delta - 1).$$

However, the total ℓ_1 norm of the vector of errors is $\mathcal{E}_1 = 2w$ independent of Δ , and $\text{OPT} = w$, so the result follows. \square

Lemma 20. *Let \mathcal{A} be any algorithm for the planning problem on weighted trees which is guaranteed to incur cost: $\text{OPT} + O(\mathcal{E}_1^a \rho^b)$ on graph searching instance \mathcal{I} , where ρ is the minimum distortion of embedding the instance graph into the path. Then $a \geq 1$ and $b \geq 1$.*

Proof. The proof is entirely analogous to that of the second part of Lemma 13. For any value of \mathcal{E}_1 , one can make use of the same construction (the weighted star, with r the central node and g a leaf) with weights $w = \mathcal{E}_1/2$ on each edge.

The result then follows from observing that for the family of constructed graphs, $\rho = \Delta$ and the analogous calculations. \square

B.4 Planning On Trees With Integer-Valued Distances

In this section, we prove Lemma 12. The result follows from arguments analogous to those outlined in Banerjee et al. (2023) Section 5.1: we first state and prove three necessary lemmas.

Lemma 21 (Analogous to Lemma 5.3 in Banerjee et al. (2023)). *Given G with positive integer distances $d : V \times V \rightarrow \mathbb{Z}_{\geq 0}$, for any $U \subseteq V$*

$$|S \setminus M(U)| \leq \sum_{u \in U} \varphi_1(u),$$

where

$$M(U) \stackrel{\text{def}}{=} \{v \in V : d(v, u) = d(v, u') \ \forall u, u' \in U\}.$$

Proof Lemma 21. Given $d : V \times V \rightarrow \mathbb{Z}_{\geq 0}$, for any $U \subseteq V$, $\forall w \notin M(U)$ let u_w, v_w denote elements of U such that $d_G(w, u_w) \neq d_G(w, v_w)$. In particular, because $d(\cdot, \cdot)$ is integer-valued, $\forall w \notin M(U)$

$$|d_G(w, u_w) - d_G(w, v_w)| \geq 1.$$

In particular, for any $S \subseteq V$

$$\begin{aligned} \sum_{u \in U} \varphi_1(u) &= \sum_{u \in U} \sum_{w \in V} |f(w) - d_G(w, u)| \\ &\geq \sum_{u \in U} \sum_{w \in S \setminus M(U)} |f(w) - d_G(w, u)| \end{aligned}$$

$$\begin{aligned}
 &\geq \sum_{w \in S \setminus M(U)} |f(w) - d_G(w, u_w)| + |f(w) - d_G(w, v_w)| \\
 &\geq \sum_{w \in S \setminus M(U)} |d_G(w, u_w) - d_G(w, v_w)| \\
 &\geq |S \setminus M(U)|.
 \end{aligned}$$

□

Lemma 22 (Generalization of Lemma 5.10 in Banerjee et al. (2023)). *For any $u, v \in V$, we have:*

$$\varphi_1(u) + \varphi_1(v) \geq 2d(u, v) + \sum_{w \in V \setminus \{u, v\}} |d(u, w) - d(v, w)|.$$

Proof of Lemma 22. The proof is a straight-forward application of the triangle inequality:

$$\begin{aligned}
 \varphi_1(u) + \varphi_1(v) &= \sum_{w \in V} |d(u, w) - f(w)| + |d(v, w) - f(w)| \\
 &\geq \sum_{w \in V} |d(u, w) - f(w) - d(v, w) + f(w)| \\
 &= \sum_{w \in V} |d(u, w) - d(v, w)| \\
 &= 2d(u, v) + \sum_{w \in V \setminus \{u, v\}} |d(u, w) - d(v, w)|.
 \end{aligned}$$

□

We also utilize the following bound on the size of the minimum Steiner tree of any sublevel set of ϕ_1 : recall that for a real-valued function $f : V \rightarrow \mathbb{R}$, we denote the sublevel set of f about threshold c as

$$L_f^-(c) \stackrel{\text{def}}{=} \{v \in V : f(v) \leq c\}.$$

Lemma 23. *For G a connected tree with at least three nodes, integer edge weights, and maximum degree Δ , let C_λ denote the set of vertices in the minimum Steiner tree containing all vertices in $L_{\phi_1}^-(\lambda)$. Then*

$$|C_\lambda| \leq \lambda \Delta.$$

Proof of Lemma 23. By definition, $L_{\phi_1}^-(\lambda) \subseteq C_\lambda$. Let $u_1, u_2 \in L_{\phi_1}^-(\lambda)$ such that

$$d(u_1, u_2) = \text{diam}(L_{\phi_1}^-(\lambda))$$

If $\nexists w \in C_\lambda$ such that $d(u_1, w) = d(u_2, w)$, then Lemma 21 implies

$$|C_\lambda| = |C_\lambda \setminus M(u_1, u_2)| \leq \varphi_1(u_1) + \varphi_1(u_2) \leq 2\lambda.$$

In particular, for G connected with at least three nodes, $\Delta \geq 2$, so the desired bound holds.

We now consider the case when $\exists w \in C_\lambda$ such that $d(u_1, w) = d(u_2, w)$. Let q_1, \dots, q_k denote the neighbors of w and let $T_i \subseteq C_\lambda$ denote the subtree of descendants of w that contains q_i . Assume without loss of generality that $T_1 \ni u_1$ and $T_2 \ni u_2$. Note that, because C_λ is defined to be minimal, $\forall i \in [k] L_{\phi_1}^-(\lambda) \cap T_i \neq \emptyset$. Let u_3, \dots, u_k be points such that $u_i \in L_{\phi_1}^-(\lambda) \cap T_i$. Consider any $x \in C_\lambda \setminus \{w\}$. Then $\exists j \in [k]$ such that $x \in T_j$. This case is illustrated in Figure 7.

Assume without loss of generality that $j \neq 1$ (this can be assumed WLOG because if $x \in T_1$, then the below argument can be carried out with respect to u_2). Because G is a tree and $x \notin T_1$,

$$d(u_1, x) = d(u_1, w) + d(w, x)$$




Figure 7: Visual aid for proof of Lemma 23, for the case when $\exists w \in C_\lambda$ such that $d(u_1, w) = d(u_2, w)$.

By choice of u_1, u_2 , $\text{diam}(L_{\phi_1}^-(\lambda)) = 2d(u_1, w)$ so in particular $d(u_1, w) \geq d(u_j, w) \forall u_j \in L_{\phi_1}^-(\lambda)$. Thus

$$d(u_1, x) = d(u_1, w) + d(w, x) \geq d(u_j, w) + d(w, x).$$

Moreover, for all $v \in T_j$, $d(v, w) = d(v, q_j) + d(q_j, w)$ by definition of subtree T_j , so for non-zero weights,

$$d(u_j, w) + d(w, x) > d(u_j, q_j) + d(x, q_j) \geq d(u_j, x)$$

We thus conclude $d(u_1, x) > d(u_j, x)$, which in particular implies $x \notin M(\{u_1, \dots, u_k\})$.

Thus for all $(C_\lambda \setminus \{w\}) \cap M(\{u_1, \dots, u_k\}) = \emptyset$, so

$$|C_\lambda| - 1 = |C_\lambda \setminus M(\{u_1, \dots, u_k\})| \leq \sum_{i=1}^k \varphi_1(u_i) \leq \lambda \Delta.$$

We have thus established the result in both cases (i.e. when $C_\lambda \cap M(\{u_1, \dots, u_k\}) = \emptyset$ and when $C_\lambda \cap M(\{u_1, \dots, u_k\}) \neq \emptyset$). \square

Proof of Lemma 12. Consider the cost of visiting every node in C_λ , a minimum Steiner tree containing the sublevel set $L_{\phi_1}^-(\lambda)$. Because C_λ is minimal,

$$\text{diam}(C_\lambda) = \text{diam}(L_{\phi_1}^-(\lambda)) \leq \lambda$$

where the last inequality follows from Lemma 22. In particular, traversing C_λ to visit every node incurs travel cost at most $\text{diam}(C_\lambda) \cdot |C_\lambda|$. Combining the above bound and Lemma 23 implies $\text{diam}(C_\lambda) \cdot |C_\lambda| \leq \lambda^2 \Delta$.

Algorithm 4 with objective ϕ_1 proceeds by iteratively visiting every node in the sublevel set $L_{\phi_1}^-(\lambda)$ by computing and traversing a minimum Steiner tree C_λ that contains the sublevel set. Algorithm 4 begins with $\lambda_0 = 1$ and doubles the threshold on each iteration, such that $\lambda_k = 2^k$. In particular, $\phi_1(g) = \mathcal{E}_1$, so the algorithm is guaranteed to terminate by the time it has visited every node of $L_{\phi_1}^-(\lambda_k)$ for the first sufficiently large threshold $\lambda_k \geq \mathcal{E}_1$. The algorithmic cost of Algorithm 4 with objective ϕ_1 is thus bounded by

$$\begin{aligned} \text{ALG} &= d(r, L_{\phi_1}^-(1)) + \sum_{k=0}^{\lceil \log_2(\mathcal{E}_1) \rceil} \text{diam}(C_{2^k}) \cdot |C_{2^k}| \\ &\leq d(r, L_{\phi_1}^-(1)) + \Delta \sum_{k=0}^{\lceil \log_2(\mathcal{E}_1) \rceil} (2^k)^2 \\ &\leq d(r, L_{\phi_1}^-(1)) + \frac{\Delta}{3} (16\mathcal{E}_1^2 - 1) \end{aligned}$$

Additionally, leveraging Lemma 22 and the fact that $\phi_1(g) \leq c\mathcal{E}_1$,

$$d(r, L_{\phi_1}^-(1)) \leq d(r, g) + d(g, L_{\phi_1}^-(1)) \leq \text{OPT} + \frac{1}{2}(\mathcal{E}_1 + 1).$$

Combining these bounds yields the desired result in the regime $\mathcal{E}_1 \geq 1$. \square

C RELATION BETWEEN EMBEDDING DISTORTION AND DOUBLING DIMENSION

We show that every graph with a low-distortion embedding into the path also has small doubling dimension / doubling constant, as per the following lemma. Recall that the doubling constant of a metric space is the smallest value λ such that, for any choice of radius $R \in \mathbb{R}$, every ball of radius R can be covered with the union of λ balls of radius $R/2$, and that the doubling dimension is given by $\log_2 \lambda$.

Lemma 24. *Let G be an undirected graph on n vertices admitting an embedding into a path on n vertices with distortion ρ and let λ be the doubling constant of G . Then:*

$$\lambda \leq \lceil 8\rho \rceil.$$

In contrast there exist graphs with constant doubling dimension that admit no embeddings into the unweighted path of distortion independent of n . For example, the 2D planar grid graph on n vertices has constant doubling dimension / doubling constant, but a simple argument shows that every embedding of the 2D planar grid into the path has distortion $\Omega(\sqrt{n})$.

Proof of Lemma 24. Let $G = (V, E)$ be an undirected graph which embeds into $[n]$ with distortion ρ . Let τ be an embedding which achieves this distortion. For any $R > 0$, let $B_G(u, R) \subseteq V$ denote the ball in G centered at u of radius R , and let $B_{[n]}(\tau(u), R)$ denote the ball of radius R in $[n]$ centered at $\tau(u)$. Let $\tau(S) \subseteq [n]$ denote the image of $S \subseteq V$ under τ . For any radius R and any $u \in V$, by the definition of the Lipschitz constant we can bound

$$d_{[n]}(\tau(u), \tau(v)) \leq \|\tau\|_{\text{Lip}} d_G(u, v) \leq \|\tau\|_{\text{Lip}} R \quad \forall v \in B_G(u, R)$$

so $\tau(B_G(u, R)) \subseteq B_{[n]}(\tau(u), \|\tau\|_{\text{Lip}} R)$. In particular, for $S_1 \subseteq V$, $\tau(S_1) \subseteq S_2$ implies $S_1 \subseteq \tau^{-1}(S_2)$, so we conclude $B_G(u, R) \subseteq \tau^{-1}(B_{[n]}(\tau(u), \|\tau\|_{\text{Lip}} R))$.

Consider $B_{[n]}(\tau(u), \|\tau\|_{\text{Lip}} R)$. Fix ϵ and let k_ϵ denote the cardinality of an ϵ -covering of $B_{[n]}(\tau(u), \|\tau\|_{\text{Lip}} R)$. Observe that for any $c, \epsilon > 0$ and any $v \in [n]$, $B_{[n]}(v, c)$ admits an ϵ -covering of cardinality at most $\lceil 2c/\epsilon \rceil$, so $k_\epsilon \leq \lceil 2\|\tau\|_{\text{Lip}} R/\epsilon \rceil$. Let $\{x_1, \dots, x_{k_\epsilon}\} \subseteq [n]$ denote the centers of the covering balls. Given $B_G(u, R) \subseteq \tau^{-1}(B_{[n]}(\tau(u), \|\tau\|_{\text{Lip}} R))$, we observe that

$$B_G(u, R) \subseteq \bigcup_{i=1}^{k_\epsilon} \tau^{-1}(B_{[n]}(x_i, \epsilon)).$$

In particular, $\forall i \in [k_\epsilon]$, for any $x, y \in B_{[n]}(x_i, \epsilon)$, the definition of the Lipschitz constant implies

$$d_G(\tau^{-1}(x), \tau^{-1}(y)) \leq \|\tau^{-1}\|_{\text{Lip}} d_{[n]}(x, y) \leq \|\tau^{-1}\|_{\text{Lip}} \cdot 2\epsilon.$$

Thus $\text{diam}(\tau^{-1}(B_{[n]}(x_i, \epsilon))) \leq 2\epsilon\|\tau^{-1}\|_{\text{Lip}}$. In particular, this implies that $\forall i \in [k_\epsilon]$ such that $\tau^{-1}(B_{[n]}(x_i, \epsilon)) \neq \emptyset$, $\exists v_i \in V$ such that $\tau^{-1}(B_{[n]}(x_i, \epsilon)) \subseteq B_G(v_i, 2\epsilon\|\tau^{-1}\|_{\text{Lip}})$. Thus

$$B_G(u, R) \subseteq \bigcup_{i=1}^{k_\epsilon} \tau^{-1}(B_{[n]}(x_i, \epsilon)) \subseteq \bigcup_{i=1}^{k_\epsilon} B_G(v_i, 2\epsilon\|\tau^{-1}\|_{\text{Lip}}).$$

We have thus produced a covering of $B_G(u, R)$ using k_ϵ balls of radius $2\epsilon\|\tau^{-1}\|_{\text{Lip}}$. We now choose ϵ so that $2\epsilon\|\tau^{-1}\|_{\text{Lip}} = R/2$, namely let $\epsilon = R/(4\|\tau^{-1}\|_{\text{Lip}})$. The cardinality of the covering is then

$$k_\epsilon \leq \left\lceil \frac{2\|\tau\|_{\text{Lip}} R}{\epsilon} \right\rceil = \left\lceil 2\|\tau\|_{\text{Lip}} R \cdot \frac{4\|\tau^{-1}\|_{\text{Lip}}}{R} \right\rceil = \lceil 8\|\tau\|_{\text{Lip}}\|\tau^{-1}\|_{\text{Lip}} \rceil$$

Using the fact that $\rho = \|\tau\|_{\text{Lip}}\|\tau^{-1}\|_{\text{Lip}}$ we conclude that the doubling constant of G is at most $\lceil 8\rho \rceil$. \square

D EXPERIMENTAL DETAILS

In this section, we provide a detailed descriptions of the experiments discussed in Section 6 of the main body of the paper. We outline two sets of experiments: the first set of experiments is used to evaluate the performance of Algorithm 1 in the presence of absolute error, the second set is used to evaluate the performance of Algorithm 2 in the presence of relative error. Both these sets of experiments focus on stochastic error.

All experiments in this section were run on a 2019 MacBook Pro with a 1.4 GHz Quad-Core Intel Core i5 Processor with 16 GB of RAM. No GPUs were used for this experiment.




Figure 8: A larger rendering of the left subfigure in Figure 2 in the main body of the paper.

Absolute Error The first set of experiments corresponds to the left side of Figure 2 (replicated above as Figure 8). Here, we generate an error vector $\vec{e} \in \mathbb{R}^n$ according to the following procedure: we fixed a value \mathcal{E}_1 representing the total desired ℓ_1 -norm of the vector of errors, and then we sampled a vector $\vec{e}_{unsigned}$ uniformly at random from the scaled simplex with ℓ_1 -norm equal to \mathcal{E}_1 , i.e.:

$$\vec{e}_{unsigned} \sim \mathcal{E}_1 \cdot \Delta_n \stackrel{\text{def}}{=} \{\vec{x} \in \mathbb{R}^n \mid \vec{x} \geq 0, \|\vec{x}\|_1 = \mathcal{E}_1\}.$$

We then assign a random sign to each entry of $\vec{e}_{unsigned}$ to obtain \vec{e} , this is done by multiplying each $\vec{e}_{unsigned}[v]$ by a Rademacher random variable σ_v . Fixing a graph G , the predictions at each vertex $v \in V$ are then given by $f(v) = d_G(v, g) + \sigma_v \cdot \vec{e}_{unsigned}[v]$. This is repeated over many instance graphs selected from four classes: Random Tree, Random Lobster, Erdos-Rényi and Circular Ladder (See paragraph **Graph Families** below). Whenever the family of graphs chosen is stochastic, as it is the case for all classes except for Circular Ladder, the graph is also resampled from its family at each iteration, so that the expectation is taken over the sampling of the graph topology as well as the random error. For each of these problem instances, we run Algorithm 1 and record the difference between the total distance ALG travelled by the algorithm to find g , and the true shortest-path distance OPT from the starting point to g , and we plot \mathcal{E}_1 against it. We report mean and standard deviation of $\text{ALG} - \text{OPT}$ over 2000 independent trials.

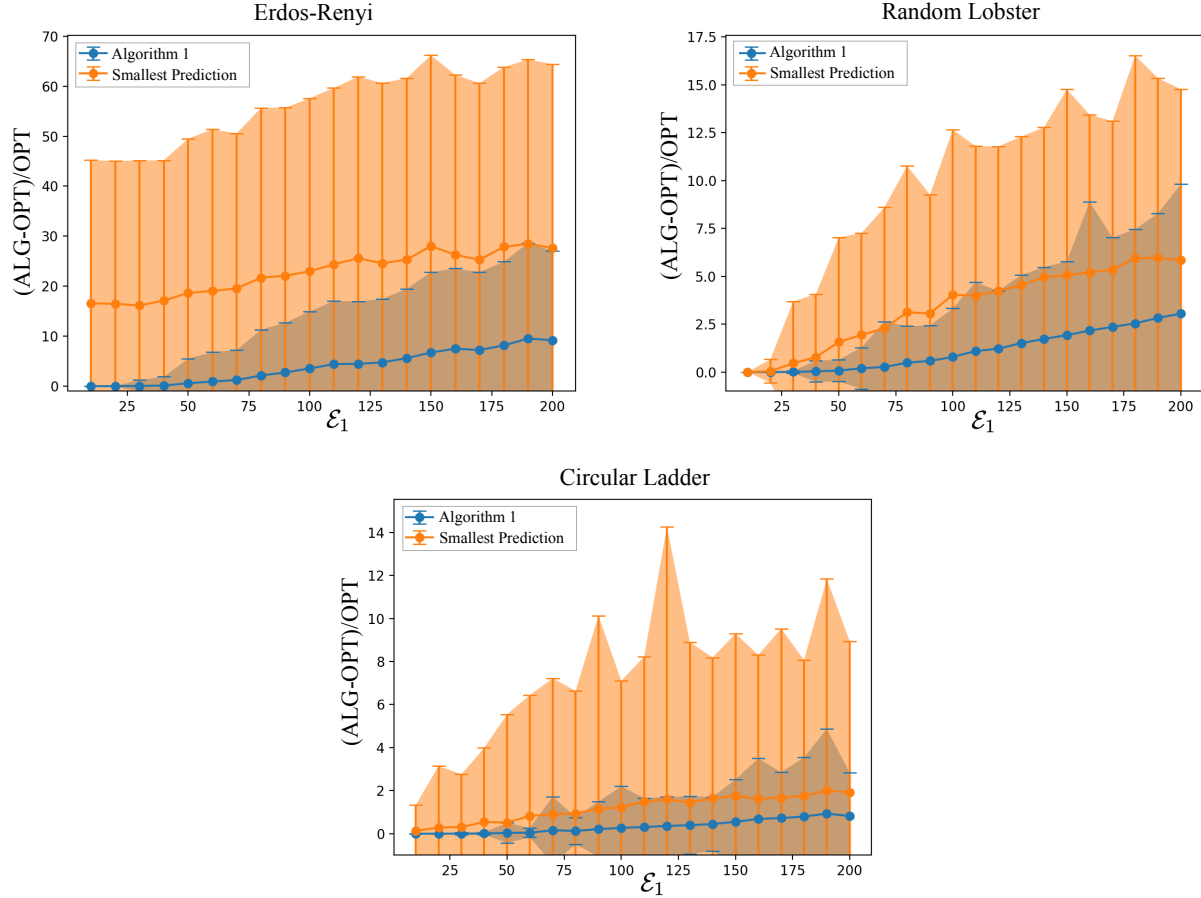


Figure 9: A comparison of the performance of Algorithm 1 with the Smallest Prediction heuristic. Each sub-figure represents one family of graphs: the top-left corresponds to random Erdős-Rényi graphs, the top-right corresponds to Random Lobster, and the middle one at the bottom corresponds to circular ladder (See **Graph Families** at the end of this section). In each figure, we plot the average and the standard deviation of the performance of Algorithm 1 and that of the Smallest Prediction heuristic against the magnitude of the error vector \mathcal{E}_1 .

Comparison to Smallest Prediction Heuristic We then compare the performance of Algorithm 1 to the Smallest Prediction heuristic defined in Section 6 (Figure 9). Recall that in Smallest Prediction, at each iteration i , the agent travels to the an arbitrary vertex $v_i \in \text{argmin}_{v \in V_{i-1}} f(v)$. We consider the same families of graphs as in the previous section. For each family, we compare the performance of our algorithm with that of Smallest Prediction for different values of the error magnitude \mathcal{E}_1 . The instances, including the errors, are generated like in the previous set of experiments. Performance is measured as the total distance travelled by the agent, minus the true distance OPT from r to g , as a fraction of OPT . Just like in the previous experiments, we run 2000 trials for every value of \mathcal{E}_1 and report the average and standard deviation of the performance across those trials.

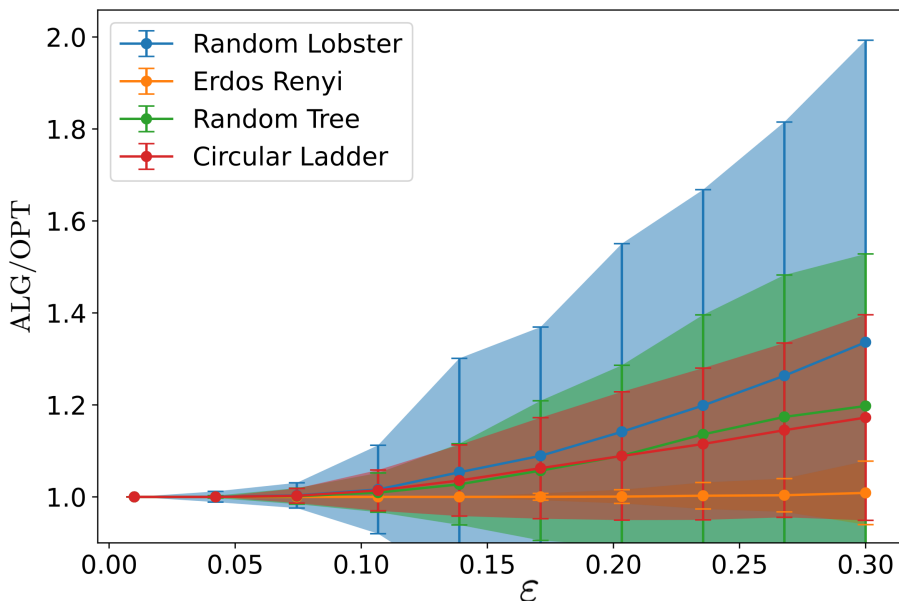


Figure 10: A larger rendering of the right subfigure of Figure 2 in the main body of the paper.

Relative Error In the right subfigure of Figure 2, for each value of ε the predictions are generated by setting $f(v) = (1 + \varepsilon_v) \cdot d_G(v, g)$ where ε_v is sampled from a Gaussian distribution with mean 0 and standard deviation $\varepsilon/2$ conditioned on the event: $\varepsilon_v \in [-\varepsilon, \varepsilon]$. We run Algorithm 3 and plot the value of the competitive ratio ALG/OPT against the value of ε for $\varepsilon \in [0, 0.3]$, and report the mean and standard deviation incurred over 2000 independent trials.

Scaling with Number of Nodes Finally, we plot the performance of Algorithm 3 as a function of n (Figure 11). For this experiment we generate instances with different numbers of vertices and plot report the average and standard deviation of the respective empirical competitive ratios (ALG/OPT). We run 2000 trials for each family of graphs and for each number $n \in \{50, 100, 500, 1000\}$ of nodes. The error is generated as in the previous section with $\varepsilon = 0.2$.

Graph Families In the above experiments we consider the four graph families described below. All the graphs considered are undirected and unweighted. In Figure 2, we sample the below graphs on $n = 100$ nodes, and in Table 2, we sample them on $n = 300$ nodes.

- **Erdős-Renyi Random Graphs:** Erdős-Rényi $G_{n,p}$ graphs (Erdős et al., 1960) are a popular random graph model in the literature. We sample from the distribution of Erdős-Rényi graphs with n nodes and edge probability $p = 0.1$, conditioned on the graph being connected;
- **Random Trees:** Trees are just connected acyclic graphs. We sample trees on n vertices uniformly at random;
- **Random Lobster Graphs:** A lobster graph is a tree which becomes a caterpillar graph when its leaves are removed, we sample random lobster graphs on n vertices;
- **Circular Ladder Graphs:** A circular ladder graph is a graph obtained by gluing the endpoints of a ladder graph, i.e. it's a graph on vertices $\{1, \dots, 2k\}$ where the edges are of the form $(i, i+1)$ for $i = 1, \dots, k-1$ and $i = k+1, \dots, 2k$, and $(i, k+i)$ for $i = 1, \dots, k$ as well as $(1, k)$, $(k+1, 2k)$. We consider the circular ladder graph on n vertices.

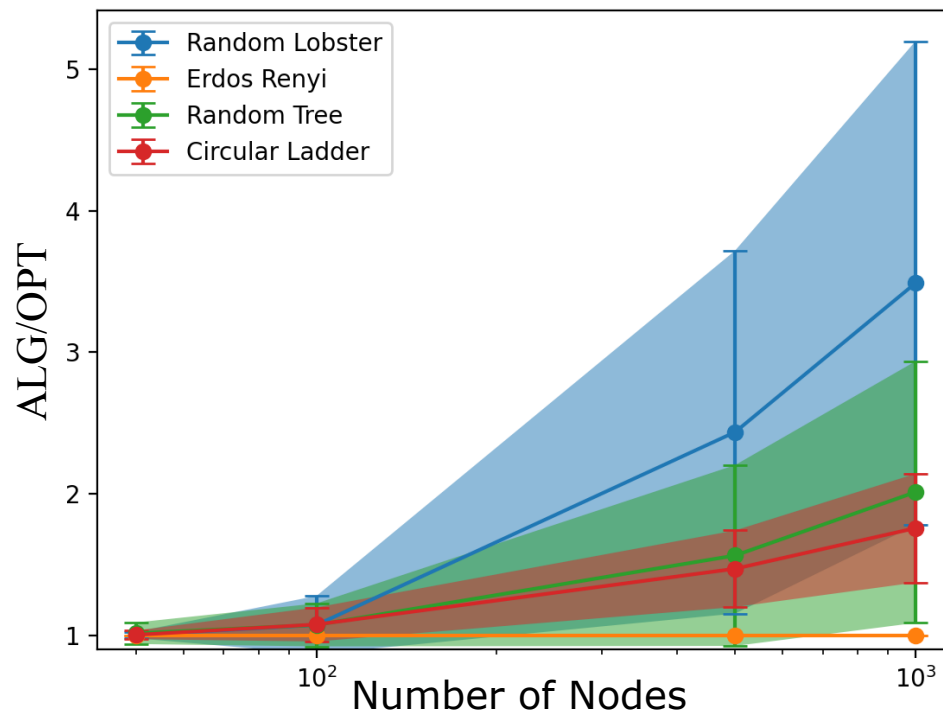


Figure 11: The empirical competitive ratio of Algorithm 3 for different graph families and for different values of n . On the x -axis: the number of vertices n in the instance graphs considered. We consider $n = 50, 100, 500, 1000$. On the y -axis: the ratio between the distance travelled by the agent, and the true distance OPT from r to g in G .