ON THE QUERY COMPLEXITY OF METRIC OPTIMIZATION UNDER WEAK-STRONG ORACLES

Anonymous authors
Paper under double-blind review

ABSTRACT

We study a Weak-Strong Oracle (WSO) framework for discrete optimization problems in a metric space $(\mathcal{X}, d)$. We assume the existence of a computationally expensive strong oracle ($\mathcal{SO}$) that outputs exact distances between points, and a computationally cheap weak oracle ($\mathcal{WO}_\alpha$) that outputs α-approximate distances that can be corrupted by an adversary with some probability. The goal is to minimize the number of queries to the strong oracle while still accurately estimating global geometric structure.

Our results reveal a structural asymmetry: while maximization problems such as diameter remain robust under adversarial corruption, minimization problems are significantly more sensitive. We give algorithms and complementary lower bounds for fundamental extremal distance problems in this model. For diameter estimation, we obtain a $(2\alpha^2 + 1)$-approximation using $O(\log^2 n)$ strong-oracle queries, and show that no algorithm using $o(n)$ strong queries can achieve an approximation factor below $2\alpha^2$. More generally, we develop a lower-bound framework based on certifying edges and their overlap, which captures the difference between metric maximization and minimization problems. While maximization problems can admit sublinear-query algorithms, we show that metric minimization can require $\Omega(n^2)$ strong-oracle queries, even for constant-factor approximation.

1 INTRODUCTION

Many fundamental machine learning problems utilize various forms of metric optimization. For example, nearest-neighbor search finds closest points to an input query (Indyk and Motwani, 1998); clustering groups similar points or selects representative centers (Gonzalez, 1985; Lloyd, 1982); and geometric covering objectives can guide the selection of examples for active learning (Sener and Savarese, 2018). In these examples, the goal is to optimize an objective function based on pairwise distances $d(x, y)$ between input points. Typically, precise distances are assumed to be available at no cost to the algorithm, but in reality, computing them can be prohibitively expensive. This is exemplified by many metrics, such as high-dimensional Euclidean distances and the Word Mover's Distance, which measures similarity between two documents through an optimal-transport computation over their word embeddings (Kusner et al., 2015). If such a computation is required for every candidate pair, even examining the input's pairwise distances can be infeasible.

A common way to mitigate this cost is to use two distance comparators rather than just one: (1) an inexpensive but coarse estimate for most pairwise distances, and (2) an accurate but expensive computation for selected pairs. In a user-modeling system deployed for Google Now, lightweight similarity scores helped restrict the number of full similarity computations used to identify related user activities. This algorithm ran about 30 times faster than a baseline using full similarity scores while producing near-optimal cluster quality (Guha et al., 2015). For text clustering, KwikBucks uses inexpensive embedding-based similarities to reduce calls to a more expensive cross-attention model. Under an assumption that the cheap signal includes all similar pairs and a bounded number of false positives, KwikBucks achieves an additive-error clustering guarantee with only $O(n)$ expensive queries, rather than examining all $\Theta(n^2)$ pairs (Silwal et al., 2023). The savings depend on correctly identifying where precise distances are required, and where cheap approximations are sufficient. Consequently, a cheap score may be wrong precisely on a pair that the objective is dependent on.

Theoretical guarantees depend heavily on how accurately cheap distances represent a target metric. Xu et al. (2024) give a bi-metric framework for nearest-neighbor search in which a data structure is

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

built using a cheap proxy metric and selectively queries an expensive target metric. Their results for DiskANN and Cover Trees obtain a $(1+\varepsilon)$-approximate nearest neighbor with respect to the target metric, for any $\varepsilon > 0$, when the proxy distance approximates the target distance within a fixed factor on every pair. However, this uniform accuracy assumption rarely holds in practice, as real-world proxy metrics frequently produce unpredictable errors on a subset of pairs. Algorithms that rely on global, uniform error bounds break down when faced with such localized, unbounded corruptions.

In the weak-strong oracle model of Bateni et al. (2024), cheap distance estimates may be adversarially corrupted. Given an n-point metric space $(\mathcal{X}, d)$ and $\alpha \geq 1$, an algorithm in the *α-weak-strong oracle* (α-WSO) model accesses two oracles:

- An expensive **strong oracle** $\mathcal{SO}$, which returns the exact distance $\mathcal{SO}(x, y) = d(x, y)$ for a queried pair.
- An inexpensive **weak oracle** $\mathcal{WO}_\alpha$. Independently for each unordered pair $\{x, y\}$, its response is uncorrupted with probability at least $\frac{2}{3}$. In that case, $\frac{1}{\alpha} d(x, y) \leq \mathcal{WO}_\alpha(x, y) \leq \alpha d(x, y)$. Otherwise, its response can be any nonnegative value chosen by an adversary.

The existing theory demonstrates that substantial savings are possible, particularly for clustering and graph objectives. Bateni et al. (2024) give constant-factor approximations for k-center, k-median, and k-means with $\tilde{O}(k^2)$ strong pair queries, as well as results for minimum spanning tree. Subsequent work studies relative-error coresets for standard and fair clustering (Braverman et al., 2025a) and improves algorithms for clustering with noisy distance oracles (Pradhan et al., 2026). Related comparison-oracle models offer further results for clustering and graph optimization (Galhotra et al., 2024; Raychaudhury et al., 2025; 2026). Despite this progress, these problem-specific results do not yet explain which features of a metric objective make strong queries avoidable.

> *What are the capabilities and limitations of the α-WSO model? Which metric objectives can be approximated using few exact distance queries, and which require inspecting a substantial fraction of all pairs?*

2 Our Contributions

In this paper, we perform a comprehensive study over the classes of metric optimization problems that admit guaranteed approximations with a "small number" of strong-oracle queries in the α-WSO model. Specifically, we give a generalized framework for proving strong-query lower bounds and algorithms for several fundamental metric objectives. Together, these results reveal metric structures that support approximations with few strong queries and structures that force many. This distinction is demonstrated by two simple extremal objectives: identifying the maximum distance $d(x, y)$ between any two points $x, y \in \mathcal{X}$, and identifying the minimum distance between any two such points. We find that minimum distance requires $\Omega(n^2)$ strong queries for every fixed approximation factor, whereas the maximum distance admits a $(2\alpha^2 + 1)$-approximation using $O(\log^2 n)$ strong queries. Our algorithms and near-matching bounds also cover farthest-pair recovery, bipartite maxima, and certified discrete 1-center. We provide additional lower bounds for nearest neighbor and exact minimum spanning tree, requiring strong queries linear in the number of candidate pairs. Formal definitions of these problems appear in Section B.

In maximization problems, the triangle inequality constrains a large distance through other distances, allowing sampled anchors to support a global approximation, while in minimization problems, an objective's optimum can be hidden at one pair of points. We refer to these different objective regimes as *Max-Rigidity* and *Min-Fluidity*.

We first examine when an exact query can detect a hidden change to a metric. Consider a baseline metric and a perturbed metric, with a distance perturbation planted at a candidate pair. For this perturbed metric, its *certifying set* consists of the pairs whose exact distances differ from the baseline. Consider a family of K such metrics, each with a distinct planted pair and let w be the maximum number of perturbed metrics in this family certified by any one pair. This is a *flagged family* if the weak oracle can use its corruption allowance to conceal the single perturbed pair and mask all other distance changes without corruption. Our lower-bound theorem follows from constructing a flagged family whose objective values differ from that of the baseline.

Theorem 2.1 (Informal meta-theorem; Theorem 4.4)**.** *Suppose a flagged family of K planted metrics has certificate overlap w, and let an α-WSO algorithm make q strong-oracle pair queries. If the set*

of acceptable outputs for each planted metric is disjoint from that for the baseline, then succeeding with probability at least $\frac{2}{3}$ requires $q = \Omega(\frac{K}{w})$. Alternatively, if the algorithm must return a list of at most L candidates containing the planted pair with probability at least $\frac{2}{3}$, then $wq + L = \Omega(K)$.

We construct a weak-oracle which adheres to the stated accuracy and corruption probability guarantees while leaving many candidates plausible, even if the algorithm makes unlimited weak queries. Under the baseline, roughly pK candidate pairs receive weak responses consistent with that of a planted pair, where p is the corruption probability. Under a uniformly planted perturbation, the true candidate is uniform among the observed candidates. We show that the two weak-transcript distributions differ in total variation by only $O(K^{-\frac{1}{2}})$ for constant p. Because one strong query can certify at most w perturbations and corruption is fixed, repeating a weak query does not remove ambiguity and the algorithm cannot identify the planted instance without $\Omega(\frac{K}{w})$ exact queries. The lower-bound applies to algorithms that return an approximately optimal pair, as well as algorithms that estimate an objective value.

Min-Fluidity arises when a perturbation can change an objective by an arbitrarily large factor while each strong query certifies only $O(1)$ planted metrics. Our construction starts with a uniform metric whose distinct pairs have distance B. For each candidate pair e, we contract only $d(e)$ to $a > 0$. Every affected triangle has side lengths (a, B, B), so the result is a valid metric. The planted pair is the only certificate of the metric, giving $w = 1$.

Theorem 2.2 (Informal Min-Fluidity result; Corollary E.5)**.** *Let $\mathcal{K}$ be a set of K candidate pairs. For every fixed $c \geq 1$, estimating $\min_{e\in\mathcal{K}} d(e)$ within factor c, or returning a pair whose distance is within factor c of this minimum, requires $\Omega(K)$ strong queries in the worst case.*

For this family, the minimum changes from B to a, and $\frac{B}{a}$ can be arbitrarily large. Taking all pairs as candidates gives an $\Omega(n^2)$ bound for closest pair, while taking pairs incident to a fixed point gives an $\Omega(n)$ bound for its nearest neighbor. These minimum-distance objectives exhibit *Min-Fluidity*, but the property applies whenever a hard perturbation family can change an objective substantially while each exact query certifies only $O(1)$ planted alternatives.

We next give algorithms for objectives whose large distances can be constrained through sampled anchors. We study two basic metric objectives: the diameter $\Delta = \max_{x,y\in\mathcal{X}} d(x, y)$ and the discrete 1-center radius $R^* = \min_{c\in\mathcal{X}} \max_{x\in\mathcal{X}} d(c, x)$. For diameter, we adapt the anchor-and-median approach of Bateni et al. (2024). Our analysis handles the additional α-distortion in uncorrupted weak responses and gives a one-sided estimate of the maximum. For discrete 1-center, a direct use of two weak responses along a path through an anchor fails due to a joint corruption probability above $\frac{1}{2}$, rendering the majority needed for the median argument impossible. We instead use one exact center-to-anchor distance and one weak target-to-anchor response in each estimate.

Theorem 2.3 (Informal Max-Rigidity algorithms; Theorems 3.2 and C.7)**.** *With probability $1 - O(n^{-3})$, an α-WSO algorithm using $O(\log^2 n)$ strong queries returns a diameter estimate $\hat{\Delta}$ satisfying $\frac{\Delta}{2\alpha^2+1} \leq \hat{\Delta} \leq \Delta$. With the same strong-query bound and success probability, an α-WSO algorithm returns a center $\tilde{c} \in \mathcal{X}$ and a certified radius $\tilde{R}$ satisfying $\max_{x\in\mathcal{X}} d(\tilde{c}, x) \leq \tilde{R} \leq (2\alpha^2+2)R^*$.*

The diameter estimator also applies to maximum distances between two sets, and a further witness-extraction step returns a pair with the same approximation factor. The discrete 1-center result shows that *Max-Rigidity* is not exclusive to maximization problems, and an objective with an outer minimum may also admit a certified approximation using $O(\log^2 n)$ strong queries.

We obtain an accuracy lower bound from a planted diameter construction using our meta-theorem. Starting from the uniform metric with distance B, for each candidate pair we set its distance to $2\alpha^2 B$, and distances on pairs incident to either endpoint to $\alpha^2 B$. The neighboring changes preserve the triangle inequality and lie at the boundary where the baseline and planted weak-oracle response intervals overlap. Thus only the planted pair requires corruption to be concealed. For $\alpha > 1$, each strong query certifies at most $\Theta(n)$ of the $\Theta(n^2)$ possible plantings, yielding the linear query barrier below factor $2\alpha^2$. Setting the neighboring dilation to 1 instead changes only the planted pair, yielding the quadratic barrier below factor 2.

Theorem 2.4 (Informal diameter lower bound; Theorem 4.8)**.** *Fix $\alpha \geq 1$. Any α-WSO algorithm that, with probability at least $\frac{2}{3}$, returns a one-sided diameter estimate satisfying $\frac{\Delta}{c} \leq \hat{\Delta} \leq \Delta$ requires $\Omega(n)$ strong-oracle queries if $c < 2\alpha^2$. If $c < 2$, it requires $\Omega(n^2)$ strong-oracle queries.*

The first bound places our $(2\alpha^2 + 1)$ guarantee within an additive 1 of the barrier for algorithms using $o(n)$ strong queries, while the second bound asserts that improving the factor below 2 requires quadratically many exact queries.

Finally, we evaluate the diameter estimator on edit-distance datasets with rare outliers. When the strong-oracle sample misses an outlier, weak-oracle medians provide information about points outside the sample; when the sample contains an outlier, exact distances sharpen the estimate. Across both datasets, combining these signals gives more accurate mean estimates than using either component alone. These experiments illustrate how the estimator can benefit from both oracles while retaining the worst-case guarantee of Theorem 3.2.

3 MAX-RIGIDITY

The Anchor-Median algorithm, an adaptation of the estimator of Bateni et al. (2024), constructs a small certificate for the global maximum. The strong queries determine an exact core among the sampled anchors, while the weak-oracle responses connect every remaining point to this core. The triangle inequality then propagates these local bounds into a global certificate for the maximum.

3.1 A GENERALIZED ESTIMATOR

When running the Anchor-Median algorithm, let $m := \lceil C \log n \rceil$ with $C = 800$ and assume $|A|, |B| \geq m$. For sorted values $y_1 \leq \cdots \leq y_k$, we use the upper median $\mathrm{median}(y_1, \ldots, y_k) := y_{\lfloor k/2 \rfloor + 1}$. First, we show that the median of the weak-oracle responses $\mathcal{WO}_\alpha(x, s)$ provides a reliable estimate of the true distances from x to the sampled anchors. We begin by establishing that, with high probability, more than half of the weak-oracle responses from every point in B to a sample $S_A \subseteq A$ of size $O(\log n)$ are uncorrupted. The same holds symmetrically for points in A and the sample $S_B \subseteq B$. For every point, the median of weak oracle estimates between this point and the anchors must be sandwiched between α-scaled true distances for some uncorrupted points. Note that this does not imply that the median itself is uncorrupted. Let $\mathcal{E}$ denote the event that, for every point, a strict majority of its weak-oracle responses to the sampled anchors are uncorrupted. A Chernoff bound and union bound give $\Pr[\mathcal{E}] \geq 1 - O(n^{-3})$, proven in Lemma C.1.

Lemma 3.1 (Bounded median). *For $u \in A$, define $M^B(u) := \mathrm{median}_{t \in S_B \setminus \{u\}} \mathcal{WO}_\alpha(u, t)$. Preconditioned upon $\mathcal{E}$, there exist uncorrupted anchors $t^{\min}, t^{\max} \in S_B$ satisfying $\frac{1}{\alpha} d(u, t^{\min}) \leq M^B(u) \leq \alpha d(u, t^{\max})$. Additionally, the symmetric statement holds for $v \in B$ against S_A.*

Algorithm 1 combines these median estimates with exact distances among the sampled anchors. When the maximum of all estimates and sampled strong queries is taken, it is at least a $\frac{1}{2\alpha^2+1}$ fraction of the true maximum. For problems where a witnessing pair is required, we can append the witness-extraction procedure Algorithm 2 to Algorithm 1, which uses an additional $O(\log n)$ strong-oracle queries and guarantees a returned pair (u^*, v^*) such that $d(u^*, v^*) \in [\frac{1}{2\alpha^2+1} \Delta_{\max}(A, B),\ \Delta_{\max}(A, B)]$.

Algorithm 1 Anchor-Median algorithm

Require: $A, B \subseteq \mathcal{X}$ with $|A|, |B| \geq m$, oracles $\mathcal{SO}$ and $\mathcal{WO}_\alpha$
Ensure: estimate $\hat{\Delta}$ of $\Delta_{\max}(A, B)$

$m \leftarrow \lceil C \log |\mathcal{X}| \rceil$, $S_A \leftarrow \emptyset$, $S_B \leftarrow \emptyset$
for $k = 1$ **to** m **do**
  $s \sim \mathrm{Uniform}(A \setminus S_A)$; $S_A \leftarrow S_A \cup \{s\}$
  $t \sim \mathrm{Uniform}(B \setminus S_B)$; $S_B \leftarrow S_B \cup \{t\}$
$R_S \leftarrow \max\{\mathcal{SO}(s, t) : s \in S_A,\ t \in S_B,\ s \neq t\}$
for $u \in A$ **do**
  $M^B(u) \leftarrow \mathrm{median}_{t \in S_B \setminus \{u\}} \mathcal{WO}_\alpha(u, t)$
for $v \in B$ **do**
  $M^A(v) \leftarrow \mathrm{median}_{s \in S_A \setminus \{v\}} \mathcal{WO}_\alpha(v, s)$
$M_A^{\max} \leftarrow \max_{u \in A} M^B(u)$, $M_B^{\max} \leftarrow \max_{v \in B} M^A(v)$
return $\hat{\Delta} \leftarrow \max\{R_S,\ \frac{1}{\alpha} M_A^{\max},\ \frac{1}{\alpha} M_B^{\max}\}$

The guarantee of the anchor-median algorithm exemplifies Max-Rigidity. The metric contains many overlapping sets of pairs which certify a near-maximal distance due to the triangle inequality. A fixed strong query is therefore likely to be informative under many regimes of possible optimal edges.

Theorem 3.2 (Robust supremum estimator)**.** *Suppose* $|A|, |B| \geq m$*. With probability at least* $1 - O(n^{-3})$ *and using* $O(\log^2 n)$ *strong-oracle queries, the output of Algorithm 1 satisfies* $\frac{1}{2\alpha^2+1} \cdot \Delta_{\max}(A, B) \leq \hat{\Delta} \leq \Delta_{\max}(A, B)$.

3.2 Applications

The following guarantees use the anchor-median argument. For bipartite inputs with fewer than m points on a side, exact cross queries or a one-sided median variant achieve the same bound; see Corollary C.9.

Table 1: Applications of the Anchor-Median estimator.

Problem	Inputs	Approximation	Strong queries
Scalar diameter	$A = B = X$	$2\alpha^2 + 1$	$O(\log^2 n)$
Farthest pair	$A = B = X$	$2\alpha^2 + 1$	$O(\log^2 n)$
Bipartite maximum	$A = U,\ B = V$	$2\alpha^2 + 1$	$O(\log^2 n)$
Bipartite farthest pair	$A = U,\ B = V$	$2\alpha^2 + 1$	$O(\log^2 n)$

The scalar guarantees follow directly from Theorem 3.2, while the corresponding witness guarantees follow from Corollary C.3.

Additionally, an adaptation of the Anchor-Median estimator yields an approximation for the discrete 1-center problem. The original estimator combines weak-oracle medians from a point to sampled anchors, with exact distances only among the anchors themselves. This is sufficient for estimating a global maximum, but it does not directly certify the radius of a specific candidate center. For 1-center, we instead use an interleaved estimator that, for a sampled candidate center c, combines the exact distance from c to an anchor t with a weak-oracle estimate from a target point x to the same anchor. We define $U(c, x) := \mathrm{median}_{t \in S_B \setminus \{x,c\}}(\mathcal{SO}(c, t) + \alpha \mathcal{WO}_\alpha(x, t))$. The exact first leg is necessary because using weak-oracle estimates on both legs would require both responses to be uncorrupted, which occurs with probability below $\frac{1}{2}$ at the standard corruption rate $p = \frac{1}{3}$, so the median argument does not hold. For each sampled candidate center, we take the maximum of these interleaved estimates over all target points, and then return the candidate center with the minimum such value. This gives a certified $(2\alpha^2 + 2)$-approximation to the discrete 1-center using $O(\log^2 n)$ strong-oracle queries with probability $1 - O(n^{-3})$. The full construction and proof are given in Theorem C.7. This result demonstrates that a problem that is not a naive maximization may follow the behavior of Max-Rigidity.

4 Lower Bounding Framework

Within the α-WSO model, we explain the origin of the multiplicative factor appearing in our approximation guarantees by providing a near-matching accuracy lower bound. We define a common framework of metric distortions and state a meta-theorem bounding query complexity with respect to properties of such a distortion. This framework is reused in Section 5 to establish a class of problems for which $\Omega(K)$ strong queries must be made.

4.1 Lower Bounding Meta-Theorem

Given a point set $\mathcal{X}$ and base metric d_0, if the distance of an edge is changed by a factor of at most α^2, then the weak oracle can return a common legal value under both metrics, so a weak response on that edge need not reveal the change.

Lemma 4.1 (Blind window)**.** *For metrics* d, d' *on* $\mathcal{X}$ *and* $e \in \binom{\mathcal{X}}{2}$*, if* $\alpha^{-2} \leq d'(e)/d(e) \leq \alpha^2$*, the legal weak-oracle intervals* $[d(e)/\alpha, \alpha d(e)]$ *and* $[d'(e)/\alpha, \alpha d'(e)]$ *overlap.*

For a planted perturbation d_e, the edges whose strong-oracle values differ from the baseline are precisely those that can expose the planted instance. The number and overlap of these certifying

edges determine how efficiently strong queries can identify a hidden perturbation. If many candidate perturbations share certifying edges, a single strong query can eliminate substantial uncertainty. However, if each candidate has a unique certifying edge, a strong query provides information about only a single candidate.

Definition 4.2 (Perturbation family)**.** *A* perturbation family *consists of a baseline metric d_0 on $\mathcal{X}$ and a metric d_e for each $e \in \mathcal{K}$. The* certifying set *of d_e is the set of pairs on which it differs from the baseline, $W_e := \{e' : d_e(e') \neq d_0(e')\}$, and the* certificate overlap *of the family is the maximum number of certifying sets in which a single certifying edge appears, $w := \max_{e'} \big|\{e \in \mathcal{K} : e' \in W_e\}\big|$.*

w is the maximum number of planted candidates that can be exposed by a single strong-oracle query. To allow the weak oracle to conceal the certifying set, we consider a family where under any two given metrics, all but the flagged edges fall in the blind window. Each planted metric is allowed exactly one pair whose change cannot be masked by an adversarial weak oracle, which may still be masked by corruption.

Definition 4.3 (Flagged family)**.** *A perturbation family is* flagged *if $e \in W_e$ for every $e \in \mathcal{K}$ and, for each $f \in \binom{\mathcal{X}}{2}$, the values $\{d_0(f)\} \cup \{d_e(f) : e \in \mathcal{K} \setminus \{f\}\}$ are within a factor α^2 of one another.*

Theorem 4.4 (Meta-theorem on strong queries)**.** *Let $d_1, \ldots, d_K$ be a flagged family for d_0 on $\mathcal{X}$ with certificate overlap w, $\alpha \geq 1$, and weak-oracle corruption parameter $p \in (0, \frac{1}{2})$. Consider an algorithm using at most q strong queries and unlimited weak queries.*

(i) If $\mathcal{B}_0 \cap \mathcal{B}_i = \emptyset$ for all $i \in [K]$ and, for every $i \in \{0, \ldots, K\}$ and every legal weak oracle, the algorithm outputs an element of $\mathcal{B}_i$ with probability at least $1 - \eta$, where $\eta \in [0, \frac{1}{2})$, then $wq \geq pK(1 - 2\eta - \frac{1}{2}\sqrt{(1-p)/(pK)})$.

(ii) If every execution outputs $\Lambda \subseteq [K]$ with $|\Lambda| \leq L$ and $\Pr[l \in \Lambda] \geq 1 - \eta$ on each input d_l, where $\eta \in [0, 1)$, then $wq \geq (1 - \eta)pK - L$.

We construct a weak oracle that leaves many planted candidates indistinguishable from the baseline. Let $S \subseteq \mathcal{K}$ be the set of candidates whose weak-oracle response is consistent with that of a planted flag. Under the baseline, each candidate lies in S independently with probability p, so $|S|$ is concentrated around pK. Under a uniformly random planted metric, the hidden candidate e^* lies in S and is uniform over S conditioned on the weak-oracle transcript.

A strong query to a pair e can expose e^* only if $e \in W_{e^*}$. Since each pair belongs to at most w certifying sets, q strong queries can expose at most qw candidates. Thus the weak oracle leaves roughly pK plausible candidates, while the strong oracle can expose only qw of them. Accounting for the $O((pK)^{-1/2})$ total variation distance between the baseline and planted weak-oracle transcripts gives the bounds stated in Theorem 4.4.

4.2 Approximation Factor for Sublinear Algorithms

We fix the corruption parameter at the value $p = \frac{1}{3}$ of Section B and suppress it in asymptotic statements. We now show that there is a barrier to surpassing a factor $2\alpha^2$ approximation, and thus no algorithm making sublinear strong-oracle queries can achieve a better guarantee. The argument is an application of Theorem 4.4 to a family of planted metrics built by dilation. We construct a family that raises the flagged pair's distance by a fixed factor and increases neighboring distances only as much as needed to preserve the triangle inequality. We choose the dilation parameter $\lambda = \alpha^2$, so neighboring distances increase by the largest factor the weak oracle can mask without corruption, as shown in Lemma 4.1.

Definition 4.5 (Single-pair dilation)**.** *Let d_0 be the uniform metric on $\mathcal{X}$, with $d_0(x, y) = B$ for all distinct $x, y \in \mathcal{X}$. For a pair $e = \{u, v\}$ and $\lambda \geq 1$ define, for all distinct $x, y \in \mathcal{X}$,*

$$d_e(x, y) := \begin{cases} 2\lambda B, & \{x, y\} = \{u, v\}, \\ \lambda B, & |\{x, y\} \cap \{u, v\}| = 1, \\ B, & \{x, y\} \cap \{u, v\} = \emptyset. \end{cases}$$

In order to measure the effects of a perturbation, define the separation to measure the smallest multiplicative change in F over the planted family.

Definition 4.6 (Separation). *Let F be a positive scalar objective and $\{d_e\}_{e\in\mathcal{K}}$ a flagged perturbation family for d_0. Its* separation *with respect to F is $s(F) := \min_{e\in\mathcal{K}} \max\left\{\frac{F(d_0)}{F(d_e)}, \frac{F(d_e)}{F(d_0)}\right\}$.*

Separation and certification capture distinct aspects of a problem. Separation determines how much the objective can change under a hidden perturbation, and therefore whether an approximation requires identifying that perturbation. We prove that a dilation defines a flagged family whose separation depends on the factor λ by which distances are increased.

Lemma 4.7 (Dilation flagged family). *For $\lambda = \alpha^2$ and $\mathcal{K} = \binom{\mathcal{X}}{2}$, d_e is a metric with $\Delta_{\max}(d_e) = 2\lambda B$, and the collection $\{d_e\}_{e\in\mathcal{K}}$ is a flagged family for d_0. When $\alpha > 1$, it has certifying sets $W_e = \{e' : e' \cap e \neq \emptyset\}$, certificate overlap $w = 2n - 3$, and separation $s(\Delta_{\max}) = 2\alpha^2$ for every $e \in \mathcal{K}$. When $\alpha = 1$, it has $W_e = \{e\}$ and $w = 1$.*

Applying Theorem 4.4 under a dilation flagged family yields the following sublinear approximation barrier.

Theorem 4.8 (Sublinear barrier). *For fixed $\alpha \geq 1$, let an α-WSO algorithm output $\hat{\Delta}$ satisfying $\frac{1}{c}\Delta_{\max}(\mathcal{X},\mathcal{X}) \leq \hat{\Delta} \leq \Delta_{\max}(\mathcal{X},\mathcal{X})$, with probability at least $\frac{2}{3}$. If $c < 2\alpha^2$, it must make $\Omega(n)$ strong-oracle queries, and if $c < 2$, it must make $\Omega(n^2)$ strong-oracle queries.*

Consequently, Algorithm 1 is optimal in α, as it attains approximation factor $2\alpha^2 + 1$ using $O(\log^2 n) = o(n)$ strong-oracle queries. By Theorem 4.8 no algorithm in this query regime attains a factor below $2\alpha^2$. Additionally, Lemma D.3 shows that $2\alpha^2$ is the largest separation any flagged family can achieve. One pair may exceed the blind window, which the triangle inequality caps at twice the window's maximum value. The construction of Definition 4.5 attains this bound, and it follows that no perturbation family can provide an improved bound because a metric cannot increase its maximum by more than a factor of 2 using an isolated edge perturbation. Any sufficiently large increase must propagate to many neighboring edges by the triangle inequality. Thus a large maximum cannot have a unique certifying edge, it necessarily comes with a collection of supporting certifying edges.

4.3 Meta-Theorem Applications

We apply Theorem 4.4 to the Max-Rigid problems considered in Section 3. The resulting lower bounds are summarized in Table 2; formal statements and proofs are deferred to Section D.3.

Table 2: Applications of the certificate-overlap lower-bound framework.

Problem	$c < 2\alpha^2$	$c < 2$
Bipartite maximum	$\Omega(n)$	$\Omega(n^2)$
Farthest-pair recovery	$\Omega(n)$	$\Omega(n^2)$
Discrete 1-center	$\Omega(n)$	—

For bipartite maximum and farthest-pair recovery, the same dilation families used for scalar maxima yield the two query thresholds. At dilation $\lambda = \alpha^2$, the planted perturbation has certificate overlap $w = \Theta(n)$, giving an $\Omega(n)$ lower bound for approximation factors $c < 2\alpha^2$. At $\lambda = 1$, only the planted pair changes, so $w = 1$, yielding an $\Omega(n^2)$ lower bound for $c < 2$.

The discrete 1-center application is a structurally different problem. Increasing a single distance affects only the radii of its endpoints, so substantially increasing the optimal radius requires perturbing many distances incident to a point. The resulting planted instances therefore have overlapping certifying sets, placing 1-center in the Max-Rigidity regime despite being a minimization problem. Applying Theorem 4.4 to this family gives an $\Omega(n)$ strong-query lower bound for approximation factors $c < 2\alpha^2$.

5 Min-Fluidity

Min-Fluidity arises when certifying edges are sparse. In the single-edge contraction family, planting candidate e changes only the distance $d(e)$. Thus e has exactly one certifying edge. Each strong query can therefore expose at most one candidate, while the weak oracle leaves a linear fraction of the K candidates plausible. This yields the $\Omega(K)$ strong-query lower bound.

5.1 Contraction Family Framework

Each contractive metric results from decreasing the distance between a single pair of points while keeping all other distances the same as in d_0. As a result, one distance may decrease drastically, while all others remain the same.

Definition 5.1 (Single-edge contraction family)**.** *For a set of points $\mathcal{X}$, a metric d_0 defined on $\mathcal{X}$, and a set of candidate pairs $\mathcal{K}$, construct a single-edge contraction family as follows. For each $e \in \mathcal{K}$, fix $0 < a_e < d_0(e)$, and define d_e by replacing only $d_0(e)$ with a_e. The family is admissible if each d_e is a metric.*

A contraction is admissible if it remains a metric, satisfying the triangle inequality. The uniform metric remains admissible under any single-edge contraction. We show that for a contractive family of metrics whose acceptable output set is disjoint from that of the original metric, the algorithm cannot distinguish between the two with sublinear strong-oracle queries.

Remark 5.2. *An admissible single-edge contraction family as in Definition 5.1 is a flagged family in the sense of Definition 4.3 having $w = 1$.*

For a single-edge contraction, each candidate has exactly one certifying edge, namely the contracted pair, so $w = 1$. The blind-window condition is vacuous because no other distance changes. We can apply this directly to our meta-theorem.

Corollary 5.3 (Contraction form of the meta-theorem)**.** *Any admissible single-edge contraction family $\{d_e\}_{e\in\mathcal{K}}$ has certificate overlap $w = 1$. Under the assumptions of Theorem 4.4, disjoint acceptable output sets imply $q \geq pK(1-2\eta-\frac{1}{2}\sqrt{(1-p)/(pK)})$. For a list of at most L candidates containing the planted pair with probability at least $1-\eta$, $q \geq (1-\eta)pK - L$.*

5.2 Applications

Table 3: Applications of the single-edge contraction framework.

Problem	Number of candidates	Lower bound
Closest pair	$\binom{n}{2}$	$\Omega(n^2)$
Bipartite closest pair	$\lvert U\rvert\lvert V\rvert$	$\Omega(\lvert U\rvert\lvert V\rvert)$
One nearest neighbor	$n-1$	$\Omega(n)$
Threshold emptiness	K	$\Omega(K)$
Exact minimum spanning tree	$\binom{n}{2}$	$\Omega(n^2)$

For the first three problems, the contracted pair is the unique c-approximate minimum when the baseline distance is chosen sufficiently larger than the contracted distance. Thus, estimating the minimum requires distinguishing the baseline from a planted metric, while returning an approximately minimizing pair requires identifying the planted candidate. Threshold emptiness has the same certifying-edge structure, as the baseline contains no candidate below the threshold, while each planted metric contains exactly one.

The contraction framework also applies when the output is a combinatorial structure rather than a distance. A single contracted edge belongs to every minimum spanning tree, so recovering the exact tree is at least as hard as identifying the contracted pair. This argument does not extend directly to approximate MSTs because a single contraction can change the optimal edge set while changing the optimal total weight by only a small factor.

Corollary 5.3 applies to an objective F exactly when $c < s(F)$. For $\Delta_{\min}$ it is unbounded, as contracting the minimizing pair of the uniform metric to a gives $\frac{\Delta_{\min}(d_0)}{\Delta_{\min}(d_e)} = \frac{B}{a}$. Thus, for every fixed c, choosing $a < \frac{B}{c}$ yields an $\Omega(K)$ query bound. For $\Delta_{\max}$ and R^*, the separation is bounded by 2 and 4, respectively, so this contraction argument gives no lower bound beyond those constants. The polylogarithmic strong-query guarantees for these objectives follow from the algorithms in Section 3.

Objectives that follow Max-Rigidity admit many overlapping certifying edges, so exact queries can provide information about many possible locations of a hidden optimum, while objectives that follow Min-Fluidity admit sparse certificates. Separation determines when such identification is necessary, while the structure and overlap of certifying edges determine the resulting strong-query complexity.

6 Empirical Evaluations

We empirically evaluate Algorithm 1 on edit-distance datasets designed to capture the setting where combining the strong and weak oracles is most useful. In particular, suppose that a dataset contains a small number of outliers that determine its diameter. If a uniform strong-oracle sample contains an outlier with probability close to 1, strong queries alone are sufficient. At the other extreme, if this probability is close to 0, the estimate is primarily determined by the weak-oracle terms. The more interesting regime is when the strong sample captures an outlier with nontrivial probability but is not guaranteed to do so. In this setting, the strong-oracle sample gives an accurate estimate in some trials and misses all outliers in others, while the weak-oracle terms provide information about points outside the sample.

We test diameter estimation ($A = B = \mathcal{X}$) using a shared uniform anchor set $S = S_A = S_B$ and compare α-WSO with two baselines from Algorithm 1:

1. **Uniform strong sampling:** Query all distinct pairs in S and return $R_S = \max_{s \neq t \in S} \mathcal{SO}(s, t)$.
2. **Weak-oracle median:** Return $R_W = \alpha^{-1} \max\{M_A^{\max}, M_B^{\max}\}$ using the medians from Algorithm 1.

The α-WSO estimator returns $\hat{\Delta} = \max\{R_S, R_W\}$. We measure accuracy by $\hat{\Delta}/\Delta$, where Δ is the true diameter. We evaluate the estimator on semi-synthetic datasets constructed from the Natural Language Toolkit (NLTK) word and name corpora, using Levenshtein edit distance as the metric. Most strings have length at most six, while a small number of length-20 strings act as outliers. We choose the number of outliers so that a uniformly sampled set of $C \log n$ points contains an outlier with probability approximately $\frac{1}{2}$, so the strong sample detects the outliers in roughly half of the trials. The word corpus contains 235,886 lowercase English words, while the names corpus contains 7,944 first names. For the former, we sample without replacement up to $n = 16{,}000$, and for the latter up to $n = 4{,}000$. In both experiments, the resulting diameter is approximately $\Delta = 20$. For each n, we run 10 independent trials with $C = 40$ and $\alpha = 2$ and report the mean estimate for each method. For a pair with Levenshtein distance d, we simulate an uncorrupted weak-oracle response by scaling d by a factor drawn uniformly from $[1/\alpha, \alpha]$. With probability $1/3$, we instead draw a corrupted response uniformly from $[0, 2\alpha d]$. The strong oracle returns the exact Levenshtein distance.

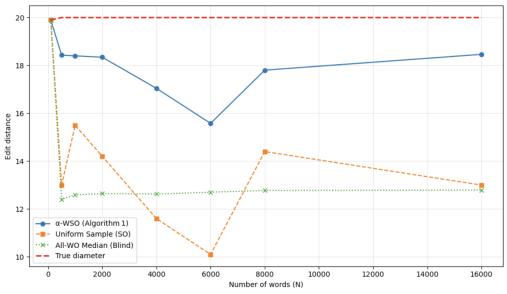


Figure 1: Diameter estimates on the NLTK word corpus. The true diameter is approximately 20. Numerical results are given in Table 7.

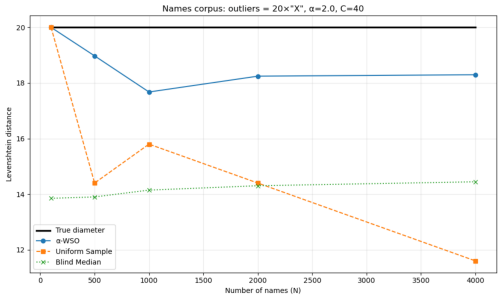


Figure 2: Diameter estimates on the NLTK names corpus without replacement. Numerical results are given in Table 8.

The results in Figures 1 and 2 demonstrate the accuracy gain of using both oracles. The strong sample encounters an outlier in roughly half of the trials, giving an accurate R_S when it does but underestimating the diameter when it does not. In contrast, $M_A^{\max}$ and $M_B^{\max}$ can capture outliers outside the strong sample, but with the multiplicative uncertainty of the weak oracle. Thus, $\hat{\Delta}$ is most accurate in this intermediate regime, where whether the strong sample captures an outlier is either roughly a toss-up or unknown.

The resulting estimates are consistent with the approximation guarantee of Theorem 3.2. While our experiments focus on a controlled outlier distribution, the distribution of real-world data may be unknown in practice. The advantage of the WSO estimator is that it provides a guarantee across all distributions while still benefiting from strong-oracle samples when they capture the extremal points.

AI USE STATEMENT

Claude was used to generate corollaries in Section C.3 and Section D.3. The authors produced distinct proofs for Theorem 4.8 and Corollary 5.3 and used Claude to assist in combining the two in Theorem 4.4. OpenAI and Claude AI models were used to verify and refine proofs in all sections. The authors reproduced all AI-generated proofs and statements independently. Code used in Section 6 was generated using OpenAI models and verified by the authors. The authors used OpenAI to search for related works, reviewed these works independently, and conducted additional searches without AI assistance. The authors wrote a draft of explanatory text in the paper and used OpenAI and Claude models to provide feedback and suggestions. AI-generated suggestions were reviewed by the authors before being incorporated.

REFERENCES

M. Bateni, P. Dharangutte, R. Jayaram, and C. Wang. Metric clustering and mst with strong and weak distance oracles. In *The thirty seventh annual conference on learning theory*, pages 498–550. PMLR, 2024.

V. Braverman, P. Dharangutte, S. H. C. Jiang, H.-A. Nguyen, C. Wang, Y. Zhang, and S. Zhou. Relative error fair clustering in the weak-strong oracle model, 2025a. URL https://arxiv.org/abs/2506.12287.

V. Braverman, J. Ergun, C. Wang, and S. Zhou. Learning-augmented hierarchical clustering. *Proceedings of Machine Learning Research*, 2025b.

M. Datar, N. Immorlica, P. Indyk, and V. S. Mirrokni. Locality-sensitive hashing scheme based on p-stable distributions. In *Proceedings of the Twentieth Annual Symposium on Computational Geometry*, pages 253–262. ACM, 2004. doi: 10.1145/997817.997857.

C. Fu, B. G. Nguyen, J. Seo, R. Zesch, and S. Zhou. Learning-augmented search data structures. In *International Conference on Learning Representations*, volume 2025, pages 99838–99865, 2025.

S. Galhotra, R. Raychaudhury, and S. Sintos. k-clustering with comparison and distance oracles. *Proceedings of the ACM on Management of Data*, 2(5):1–26, 2024. doi: 10.1145/3695830.

T. F. Gonzalez. Clustering to minimize the maximum intercluster distance. *Theoretical Computer Science*, 38:293–306, 1985. doi: 10.1016/0304-3975(85)90224-5.

R. Guha, V. Gupta, V. Raghunathan, and R. Srikant. User modeling for a personal assistant. In *Proceedings of the Eighth ACM International Conference on Web Search and Data Mining*, pages 275–284. ACM, 2015. doi: 10.1145/2684822.2685309.

P. Indyk and R. Motwani. Approximate nearest neighbors: Towards removing the curse of dimensionality. In *Proceedings of the 30th Annual ACM Symposium on Theory of Computing*, pages 604–613. ACM, 1998. doi: 10.1145/276698.276876.

P. Indyk and T. Wagner. Approximate nearest neighbors in limited space. In *Proceedings of the 31st Conference On Learning Theory*, volume 75 of *Proceedings of Machine Learning Research*, pages 2012–2036. PMLR, 2018.

M. Kusner, Y. Sun, N. Kolkin, and K. Weinberger. From word embeddings to document distances. In *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 957–966. PMLR, 2015. URL https://proceedings.mlr.press/v37/kusnerb15.html.

S. P. Lloyd. Least squares quantization in PCM. *IEEE Transactions on Information Theory*, 28(2): 129–137, 1982. doi: 10.1109/TIT.1982.1056489.

P. Pradhan, A. Bhattacharya, and R. Jaiswal. Improved algorithms for clustering with noisy distance oracles. In *Proceedings of the 29th International Conference on Artificial Intelligence and Statistics*, volume 300 of *Proceedings of Machine Learning Research*, pages 1261–1269. PMLR, 2026.

R. Raychaudhury, W.-Z. Li, S. Das, S. Galhotra, and S. Sintos. Metric clustering and graph optimization problems using weak comparison oracles. In *Proceedings of Thirty Eighth Conference on Learning Theory*, volume 291 of *Proceedings of Machine Learning Research*, pages 4777–4830. PMLR, 2025.

R. Raychaudhury, A. Esmailpour, S. Galhotra, and S. Sintos. Metric k-clustering using only weak comparison oracles. In *International Conference on Learning Representations*, 2026.

O. Sener and S. Savarese. Active learning for convolutional neural networks: A core-set approach. In *International Conference on Learning Representations*, 2018. URL https://openreview.net/forum?id=H1aIuk-RW.

S. Silwal, S. Ahmadian, A. Nystrom, A. McCallum, D. Ramachandran, and S. M. Kazemi. Kwikbucks: Correlation clustering with cheap-weak and expensive-strong signals. In *The Eleventh International Conference on Learning Representations*, 2023.

H. Xu, S. Silwal, and P. Indyk. A bi-metric framework for fast similarity search, 2024. URL https://arxiv.org/abs/2406.02891.

A Related Work

Approximate distance computation in high-dimensional spaces is well studied. Approximate nearest-neighbor data structures trade exactness for faster search, beginning with locality-sensitive hashing (Indyk and Motwani, 1998; Datar et al., 2004; Indyk and Wagner, 2018). These methods motivate the α-WSO framework, where approximate distances are cheaply available while exact queries are inefficient. Bateni et al. (2024) introduced the weak-strong oracle model for metric clustering and related problems, including k-center, k-median, k-means, and minimum spanning tree. Related work considers correlation clustering with cheap weak and expensive strong similarity signals (Silwal et al., 2023), as well as fair clustering and strong-query-efficient coresets (Braverman et al., 2025a).

A related line of work considers comparison oracles rather than approximate distance values. Galhotra et al. (2024) study k-clustering using quadruplet comparison and distance oracles, while Raychaudhury et al. (2025; 2026) study clustering and graph optimization using weak or noisy comparison oracles. These works demonstrate that approximate or noisy information can substantially reduce reliance on exact distance computations.

Our setting is also related to learning-augmented algorithms, which use auxiliary predictions to reduce computational cost (Fu et al., 2025; Braverman et al., 2025b). The main distinction is that learning-augmented models typically characterize performance as a function of prediction error, whereas our weak oracle allows a constant fraction of responses to be arbitrarily corrupted. We instead study how many expensive exact distance queries are necessary in the presence of this noisy auxiliary information.

B Mathematical Setting and Definitions

Metric objectives. Let $(\mathcal{X}, d)$ be a finite metric space with $n := |\mathcal{X}| \geq 3$. We write $e = \{x, y\}$ for an unordered pair of distinct points and $d(e) := d(x, y)$. For $A, B \subseteq \mathcal{X}$ containing at least one distinct pair $u \in A$, $v \in B$, define

$$\Delta_{\max}(A, B; d) := \max_{\substack{u \in A,\ v \in B \\ u \neq v}} d(u, v), \qquad \Delta_{\min}(A, B; d) := \min_{\substack{u \in A,\ v \in B \\ u \neq v}} d(u, v).$$

We omit the metric argument when it is d, writing $\Delta_{\max}(A, B)$ and $\Delta_{\min}(A, B)$. When the set arguments are omitted, they are understood to be $(\mathcal{X}, \mathcal{X})$. The diameter is $\Delta := \Delta_{\max}(\mathcal{X}, \mathcal{X})$, and the closest-pair distance is $\Delta_{\min}(\mathcal{X}, \mathcal{X})$. For disjoint, nonempty $U, V \subseteq \mathcal{X}$, the bipartite maximum and minimum are $\Delta_{\max}(U, V)$ and $\Delta_{\min}(U, V)$. A pair attaining the required approximation is a *witness*, as opposed to an estimate of the extremal distance. The nearest-neighbor problem for a fixed $x \in \mathcal{X}$ minimizes $d(x, y)$ over $y \in \mathcal{X} \setminus \{x\}$.

For a center $c \in \mathcal{X}$, let $R(c) := \max_{x \in \mathcal{X}} d(c, x)$. The optimal discrete 1-center radius is $R^* := \min_{c \in \mathcal{X}} R(c)$. A reported center $\tilde{c}$ and radius $\tilde{R}$ are *certified* when $R(\tilde{c}) \leq \tilde{R}$. An exact minimum spanning tree is a minimum-total-weight spanning tree of the complete graph on $\mathcal{X}$ with edge weights $d(e)$.

Distance access and query cost. Fix $\alpha \geq 1$ and a corruption probability $p \in (0, \frac{1}{2})$; unless stated otherwise, $p = \frac{1}{3}$. An algorithm has access to a strong oracle $\mathcal{SO}$ returning $\mathcal{SO}(e) = d(e)$ and a weak oracle $\mathcal{WO}_\alpha$. Independently for each unordered pair e, the weak response is uncorrupted with probability at least $1 - p$, in which case

$$\mathcal{WO}_\alpha(e) \in \left[\frac{d(e)}{\alpha},\ \alpha d(e)\right].$$

On a corrupted pair, an adversary may return any value in $\mathbb{R}_{\geq 0}$. The adversary may choose values within the uncorrupted interval and may adapt its responses to the algorithm's past actions. The response assigned to a pair is persistent: repeated weak queries to that pair do not resample its corruption status or its value. Weak responses need not satisfy the triangle inequality. We count strong-oracle *pair queries*; weak queries are unrestricted. Success probabilities include the algorithm's randomness and the weak oracle's per-pair corruption.

Lower-bound notation. In our lower bounds, $\mathcal{K} \subseteq \binom{\mathcal{X}}{2}$ indexes candidate perturbations, and $K := |\mathcal{K}|$. For example, we use $\mathcal{K} = \binom{\mathcal{X}}{2}$ for all pairs, or

$$\mathcal{K} = \big\{\{u, v\} : u \in U,\ v \in V\big\}$$

for disjoint $U, V \subseteq \mathcal{X}$. Given a metric d_i, let $\mathcal{B}_i$ denote its set of acceptable outputs. For a positive scalar objective F, a one-sided c-approximation has acceptable interval $[F(d_i)/c, F(d_i)]$ when estimating a maximum from below, and $[F(d_i), cF(d_i)]$ when estimating a minimum from above. For a solution-returning problem, $\mathcal{B}_i$ instead contains the acceptable reported points or pairs.

C Additional Details for Max-Rigidity

C.1 Additional Proofs for the General Estimator

Lemma C.1 (Uncorrupted majority). *Let $m = \lceil C \log n \rceil$ and $|A|, |B| \geq m$, and let $\mathcal{E}^+$ be the event that for every $u \in A$ more than $\frac{m+1}{2}$ of $\{\mathcal{WO}_\alpha(u, t) : t \in S_B \setminus \{u\}\}$ are uncorrupted, and symmetrically for every $v \in B$ against S_A. Then $\Pr[\mathcal{E}^+] \geq 1 - O(n^{-3})$. In particular, $\mathcal{E}^+$ implies the event $\mathcal{E}$ used in the main text.*

Proof. Fix $u \in A$. Let Z_t indicate that $\mathcal{WO}_\alpha(u, t)$ is uncorrupted for $t \in S_B \setminus \{u\}$. These indicators are independent with $\mathbb{E}[Z_t] \geq 2/3$, so $Z := \sum_{t \in S_B \setminus \{u\}} Z_t$ satisfies $\mathbb{E}[Z] \geq \frac{2}{3}(m-1)$. For $m \geq 31$, a multiplicative Chernoff bound gives

$$\Pr\left[Z \leq \frac{m+1}{2}\right] \leq \Pr\left[Z \leq \frac{4}{5}\mathbb{E}[Z]\right] \leq \exp\left(-\frac{m-1}{75}\right) = O(n^{-4}),$$

where the last bound uses $C = 800$. The symmetric argument holds for every $v \in B$. A union bound over at most $|A| + |B| \leq 2n$ points proves the claim. Moreover, deleting one additional anchor removes at most one uncorrupted response, leaving a strict uncorrupted majority. □

Lemma C.2 (Proof of Lemma 3.1). *For $u \in A$, define $M^B(u) := \mathrm{median}_{t \in S_B \setminus \{u\}} \mathcal{WO}_\alpha(u, t)$. Preconditioned upon $\mathcal{E}^+$, there exist uncorrupted anchors $t^{\min}, t^{\max} \in S_B$ satisfying*

$$\frac{1}{\alpha} d(u, t^{\min}) \leq M^B(u) \leq \alpha d(u, t^{\max}).$$

Additionally, the symmetric statement holds for $v \in B$ against S_A.

Proof. Let $k = |S_B \setminus \{u\}|$ and let U be the set of uncorrupted anchors, such that for every $v \in U$, $\mathcal{WO}_\alpha(u, v) \in [\frac{1}{\alpha} d(u, v), \alpha d(u, v)]$. On $\mathcal{E}$, $|U| > k/2$; in particular, this holds on $\mathcal{E}^+ \subseteq \mathcal{E}$. The median is the $(\lfloor \frac{k}{2} \rfloor + 1)$-th smallest value, there are at least $\lfloor \frac{k}{2} \rfloor + 1 > \frac{k}{2}$ values no greater than the median, which we define as set U_0 and at least $k - \lfloor \frac{k}{2} \rfloor \geq \frac{k}{2}$ no smaller than the median, which we define as set U_1. It follows that $U_0 \cap U \neq \emptyset$ and $U_1 \cap U \neq \emptyset$, since otherwise would lead to a contradiction in the form of implying $|S_B \setminus \{u\}| > k$. Thus, there is some uncorrupted $t^{\min} \in U_0 \cap U$ satisfying $M(u) \geq \mathcal{WO}_\alpha(u, t^{\min}) \geq \frac{1}{\alpha} d(u, t^{\min})$, and some uncorrupted $t^{\max} \in U_1 \cap U$ satisfying $M(u) \leq \mathcal{WO}_\alpha(u, t^{\max}) \leq \alpha d(u, t^{\max})$. □

Proof of Theorem 3.2. Let $\Delta_{A,B} := \Delta_{\max}(A, B)$. Precondition on the event $\mathcal{E}^+$, which occurs with probability $1 - O(n^{-3})$. For the upper bound, Lemma 3.1 gives $M^B(u) \leq \alpha d(u, t^{\max}) \leq \alpha \Delta_{A,B}$ for each $u \in A$, because $t^{\max} \in S_B \subseteq B$. Hence, $\frac{1}{\alpha} M_B^{\max} \leq \Delta_{A,B}$, and symmetrically $\frac{1}{\alpha} M_A^{\max} \leq \Delta_{A,B}$. R_S is a true distance between a point in A and a point in B, so $R_S \leq \Delta_{A,B}$. As these bound every term of the returned maximum, we get $\hat{\Delta} \leq \Delta_{A,B}$.

For the lower bound, let $(u^*, v^*) \in A \times B$ attain $\Delta_{A,B}$. By Lemma 3.1, there is an uncorrupted $s^{\min} \in S_A$ such that $d(v^*, s^{\min}) \leq \alpha M^A(v^*) \leq \alpha M_B^{\max}$, and a $t^{\min} \in S_B$ such that $d(u^*, t^{\min}) \leq \alpha M^B(u^*) \leq \alpha M_A^{\max}$. The triangle inequality on the path $u^* \to t^{\min} \to s^{\min} \to v^*$ gives

$$\Delta_{A,B} = d(u^*, v^*) \leq d(u^*, t^{\min}) + d(t^{\min}, s^{\min}) + d(s^{\min}, v^*) \leq \alpha M_A^{\max} + R_S + \alpha M_B^{\max}.$$

The middle leg is bounded by R_S because $t^{\min} \in S_B$ and $s^{\min} \in S_A$. As $\hat{\Delta}$ is the maximum of the three terms, $M_A^{\max}, M_B^{\max} \le \alpha\hat{\Delta}$ and $R_S \le \hat{\Delta}$, so $\Delta_{A,B} \le (2\alpha^2+1)\hat{\Delta}$.

The algorithm makes at most $|S_A||S_B| = m^2 = O(\log^2 n)$ strong-oracle queries. □

Algorithm 2 Witness extraction, appended to Algorithm 1

$u_A^* \leftarrow \arg\max_{u\in A} M^B(u)$
$v_B^* \leftarrow \arg\max_{v\in B} M^A(v)$
if $M^B(u_A^*) \ge M^A(v_B^*)$ **then**
    $x^* \leftarrow u_A^*$; $S^* \leftarrow S_B$
else
    $x^* \leftarrow v_B^*$; $S^* \leftarrow S_A$
for $s \in S^* \setminus \{x^*\}$ **do**
    query $\mathcal{SO}(x^*, s)$
$s^* \leftarrow \arg\max_{s\in S^*\setminus\{x^*\}} \mathcal{SO}(x^*, s)$
$(u_S, v_S) \leftarrow$ the pair of $S_A \times S_B$ attaining R_S
return the pair among (u_S, v_S) and (x^*, s^*) of larger exact distance

Corollary C.3. *With probability at least $1 - O(n^{-3})$, the output of Algorithm 2 appended to Algorithm 1 returns a pair of points $(u^*, v^*) \in A \times B$ such that*

$$d(u^*, v^*) \in \left[\frac{1}{2\alpha^2+1}\Delta_{\max}(A,B), \Delta_{\max}(A,B)\right]$$

using $O(\log^2 n)$ strong-oracle queries.

Proof. Let $M^* := \max\{M_A^{\max}, M_B^{\max}\}$ and $\Delta_{A,B} := \Delta_{\max}(A,B)$. By construction, x^* is chosen with the anchor set S^* corresponding to the statistic attaining M^*. Suppose without loss of generality that $M^* = M_A^{\max} = M^B(x^*)$, so that $x^* \in A$ and $S^* = S_B$.

By Lemma 3.1, there exists an uncorrupted $t^{\max} \in S_B$ such that

$$M^B(x^*) \le \alpha d(x^*, t^{\max}).$$

Since s^* maximizes the exact distance from x^* over $S_B \setminus \{x^*\}$,

$$d(x^*, s^*) \ge d(x^*, t^{\max}) \ge \frac{1}{\alpha}M^*.$$

The symmetric case is identical. From the proof of Theorem 3.2, $\Delta_{A,B} \le \alpha M_A^{\max} + R_S + \alpha M_B^{\max}$. Since both $M_A^{\max}$ and $M_B^{\max}$ are at most M^*, $\Delta_{A,B} \le 2\alpha M^* + R_S \le 2\alpha^2 d(x^*, s^*) + R_S$. Let (u_S, v_S) be the sampled pair attaining R_S. Therefore, $\Delta_{A,B} \le 2\alpha^2 d(x^*, s^*) + d(u_S, v_S)$ and $\max\{d(x^*, s^*), d(u_S, v_S)\} \ge \frac{\Delta_{A,B}}{2\alpha^2+1}$. Both candidate pairs lie in $A \times B$, and their distances are at most $\Delta_{A,B}$, proving the claimed approximation. The witness-extraction step uses at most $|S^*| = m = O(\log n)$ additional strong-oracle queries, so the total remains $O(\log^2 n)$. □

C.2 Certified 1-Center

Definition C.4 (Interleaved estimator)**.** *For a candidate center $c \in S_A$, and a target $x \in \mathcal{X} \setminus \{c\}$, let*

$$U(c, x) := \mathrm{median}_{t\in S_B\setminus\{x,c\}} \left(\mathcal{SO}(c,t) + \alpha\mathcal{WO}_\alpha(x,t)\right).$$

To obtain a certified center, take $A = B = \mathcal{X}$ and run the sampling and strong-query steps of Algorithm 1, obtaining S_A, S_B and the exact distances $\mathcal{SO}(c,t)$ for all $c \in S_A$ and $t \in S_B \setminus \{c\}$. For each $c \in S_A$, define

$$\tilde{R}(c) := \max_{x\in\mathcal{X}\setminus\{c\}} U(c,x).$$

Return

$$\tilde{c} \in \arg\min_{c\in S_A} \tilde{R}(c), \qquad \tilde{R} := \tilde{R}(\tilde{c}).$$

Every strong-oracle value used to compute $U(c, x)$ is among the sampled cross-pair distances.

This estimator provides an approximation for the distance of some length-two path from c to x by a similar argument as in Lemma 3.1. We must use this hybrid estimator rather than relying on a weak oracle estimate as in Algorithm 1 because not only must we compute a valid distance approximation, but also ensure the endpoints of this approximation are precisely c and x. Using a median of the sum of two weak oracle estimates breaks the Chernoff bound on the median, as the probability that both legs are uncorrupted can be as low as $\frac{4}{9} < \frac{1}{2}$. More details can be found in Remark C.6.

Lemma C.5 (Bounded interleaved estimator)**.** *For a candidate center $c \in S_A$, and a target $x \in \mathcal{X} \setminus \{c\}$, let $U(c, x)$ be as in Definition C.4. Preconditioned upon $\mathcal{E}^+$, there exist uncorrupted anchors $t^{\min}, t^{\max} \in S_B \setminus \{x, c\}$ satisfying*

$$d(c, t^{\min}) + d(x, t^{\min}) \le U(c, x) \le d(c, t^{\max}) + \alpha^2 d(x, t^{\max}).$$

Proof. By $\mathcal{E}^+$, deleting the additional anchor c still leaves a strict uncorrupted majority in $S_B \setminus \{x, c\}$. As in Lemma 3.1, this forces an uncorrupted anchor $t^{\max}$ whose term is at least the median and an uncorrupted $t^{\min}$ whose term is at most the median. The exact term is clearly unaffected by corruption, and does not contribute to the corruption probability of the total sum. For $t^{\max}$,

$$U(c, x) \le \mathcal{SO}(c, t^{\max}) + \alpha \mathcal{WO}_\alpha(x, t^{\max}) \le d(c, t^{\max}) + \alpha^2 d(x, t^{\max}),$$

and for $t^{\min}$,

$$U(c, x) \ge \mathcal{SO}(c, t^{\min}) + \alpha \mathcal{WO}_\alpha(x, t^{\min}) \ge d(c, t^{\min}) + d(x, t^{\min}).$$

□

Remark C.6 (Strong-oracle queries are necessary)**.** *The interleaved estimator of Definition C.4 cannot, in general, be replaced by an estimator using only weak-oracle queries. In particular, replacing $\mathcal{SO}(c, t)$ by $\alpha\mathcal{WO}_\alpha(c, t)$ does not preserve the median guarantee of Lemma 3.1 for all $p < \frac{1}{2}$.*

For each anchor t, the estimator combines the exact distance $d(c, t)$ with a weak-oracle estimate of $d(x, t)$. Thus, the corresponding term satisfies the desired bounds whenever the single weak-oracle response is uncorrupted, which occurs with probability $1 - p > \frac{1}{2}$. If $\mathcal{SO}(c, t)$ is instead replaced by $\alpha\mathcal{WO}_\alpha(c, t)$, both weak-oracle responses must be uncorrupted for the same guarantee to hold. This occurs with probability $(1 - p)^2$, which exceeds $\frac{1}{2}$ only when $p < 1 - 2^{-1/2} \approx 0.293$. In particular, for $p = \frac{1}{3}$, this probability is $\frac{4}{9} < \frac{1}{2}$, so an uncorrupted majority is no longer guaranteed and the median argument of Lemma 3.1 does not apply.

Theorem C.7 (Certified discrete 1-center)**.** *Suppose $n \ge m$. With probability at least $1 - O(n^{-3})$, the construction above uses $O(\log^2 n)$ strong-oracle queries and returns $(\tilde{c}, \tilde{R})$ satisfying*

$$R(\tilde{c}) \le \tilde{R} \le (2\alpha^2 + 2)R^*.$$

Proof. Precondition on the event $\mathcal{E}^+$. Let $x_0 \in \arg\max_{x \in \mathcal{X}} d(\tilde{c}, x)$, so that $R(\tilde{c}) = d(\tilde{c}, x_0)$. By the lower bound of Lemma C.5, there exists an uncorrupted $t^{\min} \in S_B$ such that $U(\tilde{c}, x_0) \ge d(\tilde{c}, t^{\min}) + d(x_0, t^{\min}) \ge d(\tilde{c}, x_0) = R(\tilde{c})$, where the second inequality follows from the triangle inequality. Since $\tilde{R} = \max_{x \in \mathcal{X} \setminus \{\tilde{c}\}} U(\tilde{c}, x)$, we obtain $\tilde{R} \ge U(\tilde{c}, x_0) \ge R(\tilde{c}) \ge R^*$.

For the upper bound, fix any $x \in \mathcal{X} \setminus \{\tilde{c}\}$. By the upper bound of Lemma C.5, there exists an uncorrupted $t^{\max} \in S_B$ such that $U(\tilde{c}, x) \le d(\tilde{c}, t^{\max}) + \alpha^2 d(x, t^{\max})$. For any $y, z \in \mathcal{X}$, the optimal center c^* gives $d(y, z) \le d(y, c^*) + d(c^*, z) \le 2R^*$. Then $U(\tilde{c}, x) \le 2R^* + 2\alpha^2 R^* = (2\alpha^2 + 2)R^*$. Since this holds for every x, $\tilde{R} = \max_x U(\tilde{c}, x) \le (2\alpha^2 + 2)R^*$. Combining the two bounds proves $R^* \le \tilde{R} \le (2\alpha^2 + 2)R^*$. The sampled cross pairs require at most $|S_A||S_B| = m^2 = O(\log^2 n)$ strong-oracle queries. □

C.3 Max-Rigidity Applications

We instantiate the estimator for various metric extremal problems. In the scalar diameter problem, we aim to approximate the farthest pairwise distance in the metric, $\Delta = \max_{u,v \in \mathcal{X}} d(u, v)$. We take $A = B = \mathcal{X}$ as inputs, and attain the following guarantee.

Corollary C.8. *Algorithm 1 run with $A = B = \mathcal{X}$ computes a conservative $(2\alpha^2+1)$-approximation of the scalar diameter Δ of the input metric with probability $1 - O(n^{-3})$ using $O(\log^2 n)$ queries. Combining with Algorithm 2 produces a $(2\alpha^2 + 1)$-approximation of the farthest pair.*

We can also consider the input to be a bipartite weighted graph such that $U \cup V = \mathcal{X}$ and $U \cap V = \emptyset$. Taking $A = U, B = V$ we can compute the approximate maximum edge length and corresponding pair of points.

Corollary C.9. *Algorithm 1 run with $A = U, B = V$ computes a conservative $(2\alpha^2 + 1)$-approximation of $\Delta_{\max}(U, V)$ with probability $1 - O(n^{-3})$ using $O(\log^2 n)$ queries. Combining with Algorithm 2 produces a $(2\alpha^2 + 1)$-approximation of the farthest pair.*

Proof. If $|U|, |V| \geq m$, apply Theorem 3.2 and corollary C.3. If both sizes are less than m, query all $|U||V| < m^2$ cross pairs and return the exact maximum and its pair.

It remains to consider $|U| < m \leq |V|$; the other case is symmetric. Sample $S_V \subseteq V$ of size m and query every pair in $U \times S_V$. Let

$$R := \max_{u \in U,\ s \in S_V} d(u, s), \qquad M(v) := \operatorname{median}_{s \in S_V \setminus \{v\}} \mathcal{WO}_\alpha(v, s).$$

Choose $\hat{v} \in \arg\max_{v \in V} M(v)$ and any $u_0 \in U$, and query $W := d(u_0, \hat{v})$. Return $D := \max\{R, W\}$ and a cross pair attaining it.

With probability $1 - O(n^{-3})$, the same uncorrupted-majority and bounded-median arguments as in Lemmas C.1 and 3.1 hold for every $v \in V$ against S_V. Let (u^*, v^*) attain $\Delta := \Delta_{\max}(U, V)$. An uncorrupted anchor below the median for v^* gives

$$\Delta \leq R + \alpha M(v^*) \leq R + \alpha M(\hat{v}).$$

An uncorrupted anchor $s^+ \in S_V$ above the median for $\hat{v}$ gives

$$M(\hat{v}) \leq \alpha d(\hat{v}, s^+) \leq \alpha\big(W + d(u_0, s^+)\big) \leq \alpha(W + R).$$

Consequently,

$$\Delta \leq (1 + \alpha^2)R + \alpha^2 W \leq (1 + 2\alpha^2)D.$$

Both candidate pairs are cross pairs, so $D \leq \Delta$. The number of strong queries is at most $|U|m + 1 = O(m^2) = O(\log^2 n)$. □

D Additional Details for Lower Bounds

D.1 Proofs for the Meta-Theorem

Proof of Blind Window Lemma 4.1. Assume without loss of generality that $d'(e) \geq d(e)$. The two weak-oracle intervals overlap whenever $\frac{d'(e)}{\alpha} \leq \alpha d(e)$. This is equivalent to $\frac{d'(e)}{d(e)} \leq \alpha^2$, which holds by assumption. The case $d'(e) < d(e)$ follows symmetrically. Therefore, whenever the two true distances differ by at most a multiplicative factor α^2, a common legal uncorrupted weak-oracle response exists under both metrics. □

When bounding the approximation constant obtainable by a sublinear α-WSO algorithm, we rely on a vanishing total variation distance between two distributions of metric spaces with characteristically different true global geometric extrema. For distributions P and P' on a finite set S,

$$\mathrm{TV}(P, P') := \frac{1}{2} \sum_{x \in S} \big|P(x) - P'(x)\big|.$$

Proof of meta-theorem Theorem 4.4. Recall that e ranges over unordered pairs of $\mathcal{X}$. By Definition 4.3 and Lemma 4.1, for each pair e, there is a common point μ_e such that μ_e is a legal uncorrupted response at e under d_0 and under every d_l for which e is not the flag. Let ν_e be a legal uncorrupted response for $\mathcal{WO}_\alpha(e)$ under metric d_e and assume that $\nu_e \neq \mu_e$. Construct a legal weak oracle as follows.

1. Under d_0, for each $e \in \mathcal{K}$, return ν_e with probability p, and μ_e otherwise. For $e \notin \mathcal{K}$, return μ_e always.

2. Under d_{e^*}, for each $e \in \mathcal{K} \setminus \{e^*\}$, return ν_e with probability p, and μ_e otherwise. For e^*, deterministically return ν_{e^*}. For $e \notin \mathcal{K}$, return μ_e.

Note that when comparing the outputs under d_0 and d_{e^*} for a single edge $e \neq e^*$, the distributions are identical. We show this oracle is legal. Under d_0 each $e \in \mathcal{K}$ admits μ_e as a legal uncorrupted answer, while returning ν_e with legal corruption probability p. Under d_e the flag e admits ν_e as a legal uncorrupted answer, so returning ν_e is again legal. The adversary spends no corruption budget on the flagged pair. Since every uncorrupted answer lies in the legal interval for every member of the family, the construction is legal for every $\alpha \geq 1$, and since no answer depends on the algorithm's past actions, the adversary is oblivious.

Define indicator random variables for the divergence from the baseline value of the output at a flag $e \in \mathcal{K}$ as $Z_e = 1$ if $\mathcal{WO}_\alpha(e) = \nu_e$ and $Z_e = 0$ if $\mathcal{WO}_\alpha(e) = \mu_e$. All responses $\mathcal{WO}_\alpha(e)$ for $e \notin \mathcal{K}$ are deterministically fixed and identical under every member of the family. Under d_0, $Z_e \sim \mathrm{Bern}(p)$. Under d_{e^*}, where the flagged perturbation occurs, $Z_{e^*} = 1$, while other Z_e follow the same distribution. As such, information gained by the algorithm from the weak oracle querying the K flags can be conveyed fully by the vector $Z \in \{0,1\}^{\mathcal{K}}$.

Let $P_0(z)$ denote the probability that the weak oracle outputs $z \in \{0,1\}^{\mathcal{K}}$ under the baseline metric d_0, and let $P_1(z)$ denote the probability that the weak oracle outputs z under a regime where a metric d_l is selected uniformly at random. Clearly $P_0(z) = p^{\|z\|_1}(1-p)^{K-\|z\|_1}$. The expected value of $\|Z\|_1$ under d_0 is $\mathbb{E}_0[\|Z\|_1] = Kp$, and variance is $\mathrm{Var}_0[\|Z\|_1] = Kp(1-p)$ as this sum follows a binomial distribution. Summing over possible l values gives,

$$P_1(z) = \frac{1}{K} \sum_{\substack{f \in \mathcal{K} \\ z_f = 1}} p^{\|z\|_1 - 1}(1-p)^{K-\|z\|_1} = \frac{\|z\|_1}{pK} P_0(z).$$

The same computation identifies the conditional law of the hidden candidate. Writing e^* for the hidden index, for any $e \in \mathcal{K}$,

$$\Pr\big[e = e^* \mid Z = z\big] = \frac{\frac{1}{K} p^{\|z\|_1 - 1}(1-p)^{K-\|z\|_1} \mathbf{1}\{z_e = 1\}}{P_1(z)} = \frac{\mathbf{1}\{z_e = 1\}}{\|z\|_1}.$$

That is, conditioned on the observed vector, the hidden index is uniformly distributed over the candidates whose flag reports the perturbed value, and the algorithm has no reason to prefer one over another.

We can bound the total variation distance using the computed variance,

$$\begin{aligned}
\mathrm{TV}(P_0, P_1) &= \frac{1}{2} \sum_{z \in \{0,1\}^{\mathcal{K}}} |P_1(z) - P_0(z)| \\
&= \frac{1}{2} \sum_{z \in \{0,1\}^{\mathcal{K}}} P_0(z) \left| \frac{\|z\|_1}{pK} - 1 \right| \\
&= \frac{1}{2} \mathbb{E}_0 \left[\left| \frac{\|Z\|_1}{pK} - 1 \right| \right] \\
&\leq \frac{1}{2} \sqrt{\mathrm{Var}_0 \left[\frac{\|Z\|_1}{pK} - 1 \right]} \\
&\leq \frac{1}{2pK} \sqrt{\mathrm{Var}_0 \left[\|Z\|_1 \right]} = \frac{1}{2} \sqrt{\frac{1-p}{pK}}
\end{aligned}$$

Consider granting the algorithm the weak vector Z in advance, increasing information. Fix this $Z = z$ and random coins. The algorithm may be adaptive, as such let $Q(z)$ be the fixed set of pairs the algorithm chooses to query under d_0. We must have $|Q(z)| \leq q$. Now consider an execution under d_l, where l is hidden from the algorithm, and the same random coins and z. Until the algorithm

queries a pair of the certifying set W_l, the strong oracle answers are identical to those produced under d_0. There are $\|z\|_1$ flagged candidates, one of which is the perturbed one. Querying q pairs, each of which lies in at most w certifying sets, the probability of identifying the perturbed candidate by diverging is at most

$$\Pr[\exists e \in Q(z) \cap W_l] \leq \frac{qw}{\|z\|_1},$$

given l is uniformly random. Taking the expectation of this value over the distribution of Z yields,

$$\mathbb{E}_1\left[\frac{qw}{\|Z\|_1}\right] = qw \sum_{\substack{z\in\{0,1\}^\kappa \\ \|z\|_1\geq 1}} \frac{P_1(z)}{\|z\|_1} = qw \sum_{\substack{z\in\{0,1\}^\kappa \\ \|z\|_1\geq 1}} \frac{P_0(z)}{pK} \leq \frac{qw}{pK}.$$

Let T_0 denote the distribution of the algorithm's output under d_0 with $z \sim P_0$, let T_1 denote its distribution under a uniformly selected d_l with $z \sim P_1$, and let T' denote its distribution when $z \sim P_1$ but every strong query $e \in Q(z)$ is answered with $d_0(e)$. Under d_0 every strong query is already answered with d_0, so T_0 and T' are produced by one and the same procedure and differ only in whether z is drawn from P_0 or from P_1. Processing cannot increase total variation distance between two input laws, so $\mathrm{TV}(T_0, T') \leq \mathrm{TV}(P_0, P_1)$. We can bound the total variation distance,

$$\begin{aligned}\mathrm{TV}(T_0, T_1) &\leq \mathrm{TV}(T_0, T') + \mathrm{TV}(T', T_1) \\ &\leq \mathrm{TV}(P_0, P_1) + \Pr[\exists f \in Q(z) \text{ with } f \in W_l] \\ &\leq \frac{1}{2}\sqrt{\frac{1-p}{pK}} + \frac{qw}{pK}\end{aligned}$$

For part (i), the event that the output belongs to $\mathcal{B}_0$ has probability at least $1-\eta$ under T_0 and at most η under T_1, since $\mathcal{B}_0$ is disjoint from every planted acceptable set. Hence $\mathrm{TV}(T_0, T_1) \geq 1 - 2\eta$. Combining and solving for qw in the two inequalities yields the desired result.

For part (ii), let $\Lambda_0(z)$ be the output set produced by an execution on fixed z in a run where for each f queried via strong oracle, the result is $d_0(f)$. This will match any execution unless some $e \in Q(z)$ lies in W_l. Successfully outputting the perturbed candidate requires $l \in \{j : \exists e \in Q(z) \cap W_j\} \cup \Lambda_0(z)$, as success can either occur if a certifying pair is queried via strong oracle and the candidate thus identified, or if the algorithm would have blindly output it under d_0. This set has size at most $qw + L$, so

$$\Pr\big[l \in \{j : \exists f \in Q(z) \cap W_j\} \cup \Lambda_0(z)\big] \leq \mathbb{E}_1\left[\frac{qw+L}{\|Z\|_1}\right] \leq \frac{qw+L}{pK}.$$

Using the fact that this probability is at least $1-\eta$ yields the desired result. □

D.2 Proofs of Lower Bounds

Proof of dilation flagged family Lemma 4.7. If a triangle contains both u and v, then its side lengths are $2\lambda B$, λB, and λB, so the triangle inequality holds. If it contains exactly one of u or v, then its side lengths are λB, λB, and B, which also satisfy the triangle inequality since $\lambda \geq 1$. All remaining triangles have side lengths B, B, B. Hence d_e is a metric, and its maximum distance is $2\lambda B$.

Now fix $\lambda = \alpha^2$. Since d_e differs from d_0 exactly at those pairs meeting e, its certifying set is $W_e = \{e' : e' \cap e \neq \emptyset\}$. A pair e' lies in W_e precisely when e meets e', and the number of pairs meeting a fixed pair is $2(n-2)+1 = 2n-3$, so the certificate overlap is $w = 2n-3$. The flag condition holds because $d_e(e) = 2\alpha^2 B \neq B = d_0(e)$, so $e \in W_e$.

For the blind window, fix a pair e' and any $e_1, e_2 \in \mathcal{K} \cup \{0\}$ with $e' \notin \{e_1, e_2\}$. As e' is the flag of neither, each of $d_{e_1}(e')$ and $d_{e_2}(e')$ equals $\alpha^2 B$ if that candidate meets e' and B otherwise, while $d_0(e') = B$. Both values lie in $\{B, \alpha^2 B\}$, so their ratio is at most α^2 and the collection is a flagged family. Finally $\Delta_{\max}(d_e) = 2\alpha^2 B = 2\alpha^2 \Delta_{\max}(d_0)$. □

Theorem D.1 (Sublinear $2\alpha^2$ barrier)**.** *For fixed $\alpha \geq 1$, let an α-WSO algorithm output $\hat{\Delta}$ satisfying*

$$\tfrac{1}{c}\,\Delta_{\max}(\mathcal{X}, \mathcal{X}) \;\leq\; \hat{\Delta} \;\leq\; \Delta_{\max}(\mathcal{X}, \mathcal{X})$$

with probability at least $\frac{2}{3}$. If $c < 2\alpha^2$, it must make $\Omega(n)$ strong-oracle queries.

Proof. If $\alpha > 1$, apply Lemma 4.7, so $K = \binom{n}{2}$ and $w = 2n - 3$. The acceptable output sets under d_0 and d_e are disjoint whenever $c < 2\alpha^2$. Applying Theorem 4.4 gives $q = \Omega(K/w) = \Omega(n)$. If $\alpha = 1$, then $\lambda = 1$ and Lemma 4.7 gives $w = 1$. Since $c < 2\alpha^2 = 2$, the same argument and Theorem 4.4 yield $q = \Omega(K) = \Omega(n^2)$, which implies the claimed $\Omega(n)$ bound. □

Theorem D.2 (Sublinear 2 barrier)**.** *For fixed $\alpha \geq 1$, any α-WSO algorithm that outputs $\hat{\Delta}$ satisfying*

$$\tfrac{1}{c}\,\Delta_{\max}(\mathcal{X}, \mathcal{X}) \;\leq\; \hat{\Delta} \;\leq\; \Delta_{\max}(\mathcal{X}, \mathcal{X})$$

with probability at least $\frac{2}{3}$, for some $c < 2$, must make $\Omega(n^2)$ strong-oracle queries. Equivalently, an algorithm making $o(n^2)$ strong-oracle queries has approximation factor $c \geq 2$. The bound holds already at $\alpha = 1$.

Proof. Apply Definition 4.5 with $\lambda = 1$ and $\mathcal{K} = \binom{\mathcal{X}}{2}$, so $K = \binom{n}{2}$. Each planted metric raises a single distance to $2B$ and leaves every other distance at B, so $W_e = \{e\}$, the flag condition $e \in W_e$ holds, the blind-window condition is vacuous, and $w = 1$. The acceptable outputs are $\mathcal{B}_0 = [B/c, B]$ under the baseline and $\mathcal{B}_e = [2B/c,\, 2B]$ under each planted metric. These are disjoint precisely when $2B/c > B$, that is when $c < 2$, which holds by hypothesis. Theorem 4.4(i) with $\eta = \frac{1}{3}$ therefore gives

$$wq \;\geq\; pK\left(\frac{1}{3} - \frac{1}{2}\sqrt{\frac{1-p}{pK}}\right),$$

and substituting $K = \binom{n}{2}$ and $w = 1$ yields $q = \Omega(pn^2)$. For constant p this yields the claimed bound. □

Lemma D.3 ($2\alpha^2$ ceiling)**.** *Let d_0 and d_1 be metrics on $\mathcal{X}$ and suppose there is a pair e such that $d_1(e') \leq \alpha^2 d_0(e')$ for every $e' \neq e$. Then $\Delta_{\max}(d_1) \leq 2\alpha^2 \Delta_{\max}(d_0)$.*

Proof. For $e' \neq e$ we have $d_1(e') \leq \alpha^2 d_0(e') \leq \alpha^2 \Delta_{\max}(d_0)$. For $e = \{u, v\}$, choose any $z \notin \{u, v\}$; neither $\{u, z\}$ nor $\{z, v\}$ is the distinguished pair, so by the triangle inequality $d_1(u, v) \leq d_1(u, z) + d_1(z, v) \leq 2\alpha^2 \Delta_{\max}(d_0)$. □

Proposition D.4 (Dilation cost)**.** *Let d_0, d_1 be metrics on n points and let $W := \{e' : d_1(e') \neq d_0(e')\}$. If $\Delta_{\max}(d_1) > 2\Delta_{\max}(d_0)$ then $|W| \geq n - 2$.*

Proof. Let $e = \{u, v\}$ attain $\Delta_{\max}(d_1)$ and fix $z \in \mathcal{X} \setminus \{u, v\}$. If neither $\{u, z\}$ nor $\{z, v\}$ lies in W then $d_1(u, v) \leq d_0(u, z) + d_0(z, v) \leq 2\Delta_{\max}(d_0)$, a contradiction. Each of the $n - 2$ choices of z therefore contributes a distinct pair of W. □

D.3 Lower Bound Applications

We apply the framework to the maximization problems considered in Section 3, observing a query-approximation tradeoff. In particular, the dilation construction and certifying set meta-theorem identify two distinct thresholds: achieving an approximation factor strictly below 2 requires quadratically many strong-oracle queries, while achieving a factor strictly below $2\alpha^2$ requires linearly many strong-oracle queries. In conjunction with our $O(\log^2 n)$ upper bound and the classical farthest-point estimator, this gives a robust characterization of this tradeoff. Again, throughout this subsection we take the corruption parameter to be the constant $p = \frac{1}{3}$ fixed in Section B, and suppress it in the asymptotic bounds.

The first two applications use the $\lambda = 1$ case of Definition 4.5, in which a planted metric raises a single distance to $2B$ and leaves every other distance at B. Such a family is flagged and the certifying set of d_e is $W_e = \{e\}$, so $e \in W_e$ and the certificate overlap is $w = 1$. In this case the blind-window condition is vacuous because no pair other than the flag is modified. We begin with a lower bound on the bipartite supremum contrasting Algorithm 1's $(2\alpha^2 + 1)$-approximation using $O(\log^2 n)$ strong-oracle queries.

Corollary D.5 (Bipartite supremum)**.** *Let $U, V \subseteq \mathcal{X}$ be disjoint with $|U|, |V| = \Theta(n)$, and let $\Delta_{\max}(U, V) := \max_{u \in U,\, v \in V} d(u, v)$. Any algorithm that computes a conservative c-approximation of $\Delta_{\max}(U, V)$ with constant success probability requires $\Omega(n)$ strong-oracle queries if $c < 2\alpha^2$, and $\Omega(n^2)$ strong-oracle queries if $c < 2$.*

Proof. Take $\mathcal{K} = \{\{u, v\} : u \in U,\ v \in V\}$, so that $K = |U||V| = \Theta(n^2)$, and apply Definition 4.5 with a planted pair $e \in \mathcal{K}$. In either case the planted metric attains $\Delta_{\max}(U, V) = 2\lambda B$ at e, against a baseline value of B, so the acceptable outputs are $[B/c, B]$ and $[2\lambda B/c,\ 2\lambda B]$, which are disjoint precisely when $c < 2\lambda$.

For $\lambda = \alpha^2$, the certifying set W_e consists of the pairs meeting e. A pair $e' = \{x, y\}$ lies in W_e for $e = \{u, v\}$ exactly when $u \in e'$ or $v \in e'$, and there are at most $2|V|$ candidates of the first kind and $2|U|$ of the second, so $w \le 2(|U| + |V|) = \Theta(n)$. Theorem 4.4 gives $q = \Omega(K/w) = \Omega(n)$ whenever $c < 2\alpha^2$.

For $\lambda = 1$ only the planted pair changes, so $W_e = \{e\}$ and $w = 1$, and Theorem 4.4 gives $q = \Omega(K) = \Omega(n^2)$ whenever $c < 2$. □

In Section 3.2, we provided a bound on returning the metric diameter. We extend this bound to the application of returning an approximate farthest pair which corresponds to an output diameter estimate. Again, our lower bound on approximation factor matches the algorithm's upper bound up to a constant factor.

Corollary D.6 (Farthest-pair recovery)**.** *Suppose an algorithm must return a pair (x, y) satisfying $d(x, y) \ge \frac{1}{c}\Delta_{\max}$. With constant success probability, any such algorithm requires $\Omega(n)$ strong-oracle queries if $c < 2\alpha^2$, and $\Omega(n^2)$ strong-oracle queries if $c < 2$.*

Proof. Take $\mathcal{K} = \binom{\mathcal{X}}{2}$, so that $K = \binom{n}{2}$, and apply Definition 4.5 with a planted pair e. For $\lambda = 1$ the baseline has all distances equal to B, while each planted metric has its own pair at distance $d_e(e) = 2B$ and all remaining pairs at distance B. Hence $W_e = \{e\}$ and $w = 1$. When $c < 2$ we have $\frac{2B}{c} > B$, so every other pair fails the requirement and the only admissible output is the planted pair itself. The list-recovery part of Theorem 4.4 with $L = 1$ gives $q = \Omega(K) = \Omega(n^2)$.

For $\lambda = \alpha^2$ the planted metric takes the value $2\alpha^2 B$ at e, the value $\alpha^2 B$ at the pairs meeting e, and B elsewhere. Thus W_e is the set of pairs meeting e and $w = 2n - 3$. When $c < 2\alpha^2$ we have $\frac{1}{c} \cdot 2\alpha^2 B > B$, so an admissible output e' must meet e, though it need no longer equal e. Taking Λ to be the set of candidates meeting e', we have $|\Lambda| \le 2n - 3$ and $e \in \Lambda$ on every successful execution, so Theorem 4.4 with $L = 2n - 3$ gives $(2n - 3)q \ge (1 - \eta)K - (2n - 3)$ and hence $q = \Omega(n)$. □

The final application differs from the preceding two in the metric family that attains the largest separation. For the supremum, a single distance determines the objective, so increasing one pair suffices to change its value. In contrast, the 1-center radius is a minimum over candidate centers, and increasing a single distance affects only the radii of its two endpoints, leaving other centers unchanged. Thus, to substantially increase the optimal radius, the planted metric must dilate an entire point rather than a single pair. Its certifying set therefore consists of all pairs incident to that point, while the candidate instances correspond to pairs joining each point to a distinguished center.

For $\alpha > 1$, each planted metric changes $n - 1$ pairs, but any fixed pair changes under at most two plantings, so the certificate overlap is $w = 2$. When $\alpha = 1$, only the flagged pair changes and $w = 1$. Since $K = n - 1$, Theorem 4.4 yields an $\Omega(n)$ strong-query bound below factor $2\alpha^2$. The reason is that R^* cannot be substantially changed by a single-edge contraction. In particular, if d' is obtained from d by contracting one edge, then

$$R^*(d) \le \Delta_{\max}(d) \le 2\Delta_{\max}(d') \le 4R^*(d').$$

The outer inequalities follow from $\frac{1}{2}\Delta_{\max} \le R^* \le \Delta_{\max}$ for every metric, while the middle inequality follows from Lemma D.3. Hence every single-edge contraction family has separation $s(R^*) \le 4$, and by Definition 4.6 cannot rule out approximation factors above 4. Thus, although 1-center is a minimization problem, its large, overlapping certificates place it on the Max-Rigidity side of our dichotomy.

Corollary D.7 (Discrete 1-center)**.** *Let $R^* := \min_{c \in \mathcal{X}} \max_{x \in \mathcal{X}} d(c, x)$, and let $\eta \in [0, \frac{1}{2})$. Any α-WSO algorithm that returns a certified radius $\widehat{R}$ with $R^* \le \widehat{R} \le c\, R^*$ with probability at least $1 - \eta$, for some $c < 2\alpha^2$, must make $q = \Omega(n)$ strong-oracle queries.*

Proof. Fix $c^* \in \mathcal{X}$ and $B > 0$, and take as the candidate set the $n-1$ pairs $\{v, c^*\}$ at c^*,

$$\mathcal{K} \;=\; \big\{\, \{v, c^*\} : v \in \mathcal{X} \setminus \{c^*\} \,\big\}, \qquad K = n-1,$$

so that each candidate is itself the flag of its planted metric, as Definition 4.3 requires. Let d_0 be the star metric $d_0(c^*, x) = B$ and $d_0(x, y) = 2B$ for $x, y \neq c^*$, and for $e = \{v, c^*\} \in \mathcal{K}$ let d_e be the spiked metric with outlier v:

$$d_e(v, c^*) = (2\alpha^2 + 1)B, \qquad d_e(v, x) = 2\alpha^2 B, \qquad d_e = d_0 \text{ elsewhere.}$$

Both are metrics, the tightest triples being $(2B, B, B)$ and $\big((2\alpha^2+1)B,\, 2\alpha^2 B,\, B\big)$, each an equality. Computing radii, $R^*(d_0) = B$ at c^*, while under d_e both c^* and v have radius $(2\alpha^2+1)B$ and every other point has radius $2\alpha^2 B$, so $R^*(d_e) = 2\alpha^2 B$.

We begin by identifying the certificate overlap. Assume $\alpha > 1$, so that d_e moves exactly the pairs incident to v and $W_e = \{\{v, y\} : y \neq v\}$; in particular $e \in W_e$. A pair e' lies in $W_{\{v,c^*\}}$ iff $v \in e'$ and $v \neq c^*$, which holds for two candidates when $c^* \notin e'$ and for one when $c^* \in e'$; hence $w = 2$. When $\alpha = 1$ only the flag itself moves, giving $w = 1$, and the bound below only improves.

We verify that the family is flagged. Let e' be any pair and let $e'' \in \mathcal{K} \setminus \{e'\}$. If $e' = \{v, c^*\}$ is a candidate, then v is unperturbed under $d_{e''}$, so $d_{e''}(e') = B = d_0(e')$ and the ratio condition of Definition 4.3 is immediate. Otherwise $e' = \{x, y\}$ with $x, y \neq c^*$, and $d_{e''}(e') = 2\alpha^2 B$ when $e'' \in \{\{x, c^*\}, \{y, c^*\}\}$ and $d_{e''}(e') = 2B$ otherwise, a ratio of exactly α^2. By Lemma 4.1 the responses $\mu_{e'} = \alpha B$ in the first case and $\mu_{e'} = 2\alpha B$ in the second are legal and uncorrupted throughout. Taking $\nu_e = (2\alpha^2+1)B/\alpha$, the infimum of the legal interval for $d_e(e)$, the flag is answered without corruption under d_e; and since $\nu_e > \alpha B = \mu_e$, the two responses differ.

With the stated convention the acceptable sets are $\mathcal{B}_0 = [B, cB]$ and $\mathcal{B}_e = [2\alpha^2 B,\, 2\alpha^2 cB]$, which are disjoint precisely when $c < 2\alpha^2$.

Theorem 4.4 with $K = n-1$ and $w = 2$ now gives

$$wq \;\geq\; p(n-1)\left(1 - 2\eta - \tfrac{1}{2}\sqrt{\tfrac{1-p}{p(n-1)}}\right),$$

hence $q = \Omega(n)$ for n sufficiently large. □

E Additional Details for Min-Fluidity

E.1 Min-Fluidity Proofs

Lemma E.1 (Admissibility criteria)**.** *For a candidate pair $e = \{u, v\} \in \mathcal{K}$, a contraction d_e is admissible if and only if*

$$a_e \geq \max_{x \in \mathcal{X} \setminus e} |d_0(u, x) - d_0(v, x)|$$

Proof. Assume for all $x \in \mathcal{X} \setminus \{u, v\}$, $a_e \geq |d_0(u, x) - d_0(v, x)|$. Distances outside of $d_e(e)$ are the same as in d_0, so the triangle inequality holds for triples of points not including the altered pair. Consider any x outside of this pair. $a_e + d_0(v, x) \geq d_0(u, x)$, which implies $d_e(u, v) + d_e(v, x) \geq d_e(u, x)$, which holds for swapped v, u. Additionally, $d_e(v, u) \leq d_0(v, u)$, so d_e is admissible.

Now, assume there is an x such that $a_e < |d_0(u, x) - d_0(v, x)|$, and without loss of generality that $d_0(u, x) > d_0(v, x)$. Then, $d_e(u, v) + d_e(v, x) < d_e(u, x)$, implying that d_e is not admissible. □

Corollary E.2 (Uniform metric contraction)**.** *Let $B > 0$, candidate set $\mathcal{K}$ be fixed, and $a_e = a < B$ for every $e \in \mathcal{K}$. Under a uniform metric having $d_0(x, y) = B$ for every distinct pair $x, y \in \mathcal{X}$, the single-edge contraction family induced is admissible.*

Proof. Every difference appearing in Lemma E.1 is $|B - B| = 0$, so the criterion $a_e = a > 0$ is satisfied for every $e \in \mathcal{K}$. Equivalently, after contracting a single candidate every triangle has side lengths (B, B, B) or (a, B, B), and both satisfy all three triangle inequalities. □

Corollary E.3 (Corollary 5.3 restated)**.** *Let $\{d_e : e \in \mathcal{K}\}$ be an admissible single-edge contraction family for d_0 on $\mathcal{X}$, let $\alpha \geq 1$, and let $p \in (0, \frac{1}{2})$. Suppose an algorithm makes at most q strong-oracle queries together with arbitrarily many weak-oracle queries.*

(i) If $\mathcal{B}_0 \cap \mathcal{B}_e = \emptyset$ for every $e \in \mathcal{K}$ and the algorithm returns an output in $\mathcal{B}_e$ with probability at least $1 - \eta$ for every metric in the family and every legal weak oracle, where $\eta \in [0, \frac{1}{2})$, then

$$q \;\geq\; pK\left(1 - 2\eta - \frac{1}{2}\sqrt{\frac{1-p}{pK}}\right).$$

(ii) If instead the algorithm outputs on every execution a set $\Lambda \subseteq \mathcal{K}$ of candidate pairs with $|\Lambda| \leq L$, containing e on input d_e with probability at least $1 - \eta$ for $\eta \in [0, 1)$, then

$$q \;\geq\; (1-\eta)pK - L.$$

Proof. Immediate from Theorem 4.4 and Remark 5.2, substituting $w = 1$. □

Remark E.4 (Constants)**.** *At the corruption probability $p = \frac{1}{3}$ of Section B and success probability $\frac{2}{3}$, so $\eta = \frac{1}{3}$, Corollary 5.3 reads*

$$q \;\geq\; \frac{K}{9} - \frac{\sqrt{2K}}{6} \;=\; \Omega(K), \qquad q + L \;\geq\; \frac{2K}{9} \;=\; \Omega(K),$$

the second giving $q = \Omega(K)$ whenever the output list has size $L = o(K)$.

E.2 Min-Fluidity Applications

Corollary E.5 (Graph minima)**.** *For any candidate set $\mathcal{K}$, both a c-multiplicative approximation and outputting $e \in \mathcal{K}$ with $d(e) \leq c \min_{e' \in \mathcal{K}} d(e')$ require $\Omega(K)$ strong-oracle queries.*

Proof. Use the uniform baseline metric with distance $B > c^2$ and contract one candidate pair to $a = 1$. The scalar objective changes from B to 1 and even the two-sided acceptable intervals $[B/c, cB]$ and $[1/c, c]$ are disjoint, so Corollary 5.3 applies. On a planted metric, any c-approximate returned pair must equal the contracted pair, since every other candidate has distance $B > c$. Using Corollary 5.3 with $L = 1$ demonstrates the query requirement for returning an approximate pair. □

Table 4: Specializations of the graph minima corollary. Each row fixes a candidate set $\mathcal{K}$ of size K; the bound $\Omega(K)$ applies both to scalar c-approximation and to returning a pair $e \in \mathcal{K}$ with $d(e) \leq c \min_{e' \in \mathcal{K}} d(e')$, for any fixed $c \geq 1$.

Task	Candidate set $\mathcal{K}$	Lower bound
Closest pair	all unordered pairs	$\Omega(n^2)$
Bipartite closest pair	$U \times V, U \cap V = \emptyset$	$\Omega(\lvert U\rvert\lvert V\rvert)$
One nearest neighbor	$\{\{x, y\} : y \in \mathcal{X} \setminus \{x\}\}$, x fixed	$\Omega(n)$

The meta-theorem also extends to threshold emptiness, showing that deciding whether any candidate distance is below a threshold requires linear queries to candidate edges.

Corollary E.6 (Threshold emptiness)**.** *For any threshold $\tau > 0$, deciding whether there exists $e \in \mathcal{K}$ with $d(e) \leq \tau$ requires $\Omega(K)$ strong queries. Likewise, reporting all candidate pairs within the threshold requires $\Omega(K)$ strong-oracle queries.*

Proof. Take $B = 2\tau$ and $a = \tau/2$ in the uniform contraction family. The correct decision is **NO** on the baseline and **YES** on every planted metric, so the acceptable answer sets are disjoint and Theorem 4.4 applies with no approximation parameter. A correct reporting algorithm decides emptiness by testing whether its output is empty, at no additional query cost, which gives the final claim. □

The contraction framework also applies when the output is a combinatorial structure rather than a distance. A single contracted edge belongs to every minimum spanning tree, so recovering the exact tree is at least as hard as identifying the contracted pair. We note that this argument does not extend to approximate MSTs. Contracting one edge from B to a changes the optimal weight from $(n-1)B$

to $(n-2)B + a$, while forcing that edge into every minimum spanning tree. Thus, sensitivity must be measured with respect to the required output, not just its scalar objective value. In certification terms, the contracted edge is essential for certifying the exact MST even though it has little effect on the total weight.

Corollary E.7 (Exact minimum spanning tree)**.** *Returning an exact minimum spanning tree of the complete metric graph requires $\Omega(n^2)$ strong-oracle queries.*

Proof. Take the uniform baseline with distance $B > 1$ and contract one of the $K = \binom{n}{2}$ pairs, say e, to distance 1. Every spanning tree omitting e has weight $(n-1)B$. Adding e to such a tree and deleting any other edge of the resulting cycle yields a spanning tree of strictly smaller weight, since the deleted edge has weight $B > 1$. Hence every exact minimum spanning tree contains e, and its edge set has size $n-1$. Theorem 4.4 with $L = n-1$ gives $q + n - 1 \geq \frac{2}{3}p\binom{n}{2}$, so $q = \Omega(n^2)$. □

Corollary 5.3 applies to an objective F exactly when $c < s(F)$. It follows that the separation of a contraction family is the largest lower bound on approximation factor that can be produced by our framework. For $\Delta_{\min}$ it is unbounded, as contracting the minimizing pair of the uniform metric to a gives $\frac{\Delta_{\min}(d_0)}{\Delta_{\min}(d_e)} = \frac{B}{a}$, so every fixed c is unattainable and $\Omega(K)$ queries are required at every approximation factor. For $\Delta_{\max}$ and R^* the separation is bounded by 2 and 4 respectively, so no contraction family constrains them beyond those constants and both admit polylogarithmic certificates.

Max-Rigid objectives admit many overlapping certifying edges, so exact queries can provide information about many possible locations of a hidden optimum. Min-Fluid objectives admit sparse certificates, often with essentially one certifying edge per candidate. Separation determines when such identification is necessary, while the structure and overlap of certifying edges determine the resulting strong-query complexity.

F Additional Empirical Analysis

F.1 Synthetic Datasets

Intuitively, when the experiment is performed on a heavy-tailed distribution with a small number of far outliers, the $M_S^{\max}$ results take over in the Algorithm 1 estimates. As such, the $\mathcal{WO}_\alpha$ median baseline results are similar, but the $\mathcal{SO}$ uniform sample baseline performs poorly, as there is a very high probability of only sampling points from the center of the distribution. As a consequence, the WSO algorithm's advantage is revealed on datasets where the distribution is unknown or can vary.

This advantage over both baselines is revealed in Figure 3, which shows the average of 20 trials for each pair of algorithm and n value. In each trial, we compute a $k \sim 1 + \text{Exponential}(3)$, and generate n points $x \in \mathbb{R}^2$ which make up the input set, such that each x is sampled from the Student's t-distribution, $x \sim t_k^{(2)}$. As a result, each n will be tested on a variety of distributions, skewing more tail-heavy, with some more normally distributed datasets included as well. Due to the robustness of the WSO algorithm estimator, it consistently outperforms both baselines in this experiment.

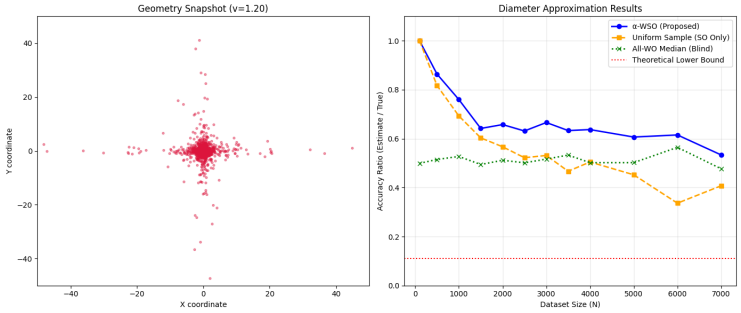


Figure 3: Mean approximation ratios on inputs of randomized t-distributions

As discussed, our first synthetic dataset experiment demonstrated the algorithm's superior performance on datasets of unknown distributions. However, the WSO algorithm also produces superior results on datasets where the probability that a $C \log n$ uniform sample picks up a pair of relatively distant points is a toss-up. As previously discussed, if this probability is close to 1, the uniform sample baseline will perform as well as the WSO algorithm, whereas if it is close to 0, it will perform poorly, and the $\mathcal{WO}_\alpha$ median baseline will perform as well as WSO algorithm. In order to demonstrate this advantage, we tested on a naive dataset where the number of outliers was fixed so that the probability of the uniform sample baseline sampling an outlier was fixed to approximately $\frac{1}{2}$.

For each value of n, our dataset consisted of points in $\mathbb{R}$, where a total of x outliers have value 1, and the remainder of points have value 0. The x value was set to ensure the probability that the $C \log n$ sample contained no 1-valued points was $\frac{1}{2}$, approximately satisfying

$$\binom{n-x}{C\log n}\binom{n}{C\log n}^{-1} = \frac{1}{2}.$$

By taking a Stirling approximation of the permutations, the value of x is obtained:

$$x \approx \frac{n \log 2}{C \log n}.$$

Figure 4 shows the result of running the experiment on such datasets of n values of 50 through 8000. For each value of n, the dataset was generated, and 20 trials of each algorithm were performed. The mean of all trials was computed and plotted for the WSO algorithm and each baseline. Figure 4 shows the uniform sample baseline's approximation ratio is consistently close to $\frac{1}{2}$, as expected. The approximation ratio of the $\mathcal{WO}_\alpha$ median baseline hovers close to 0.6 for $n \geq 1000$. This indicates that in around $\frac{1}{2}$ of trials, the R_S value found by Algorithm 1 is precisely 1, and the WSO algorithm produces the exact result. In all other cases, $R_S = 0$, and Algorithm 1 utilizes $M_S^{\max}$, the value of which is approximated by the $\mathcal{WO}_\alpha$ median baseline. The WSO algorithm's mean approximation ratio remains consistently above 0.85 for $n \geq 500$. On this dataset, the WSO algorithm outperforms both baselines by over 50% for $n \geq 5000$.

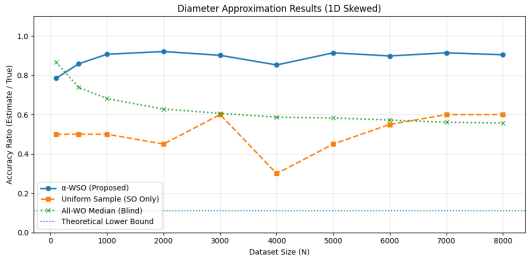


Figure 4: Mean approximation ratios on inputs of binary datasets with fixed probability of randomly sampling a 1-valued element

F.2 Data Tables

Table 5: Mean approximation ratios by input dataset size, as plotted in Figure 4

n	WO Median Baseline	Uniform Sample Baseline	WSO Algorithm
100	0.866	0.500	0.784
500	0.738	0.500	0.859
1000	0.681	0.500	0.907
2000	0.628	0.450	0.921
3000	0.607	0.600	0.901
4000	0.587	0.300	0.852
5000	0.583	0.450	0.914
6000	0.572	0.550	0.898
7000	0.561	0.600	0.914
8000	0.557	0.600	0.904

Table 7: Mean absolute diameter estimates for $\alpha = 2$ on the word-length outlier dataset. The true diameter is 20. Ten trials per n. Visualized by Figure 1

n	True Metric Diameter	α-WSO	Uniform (SO)	WO-Median (Blind)
100	19.5	19.5	18.1	11.9
500	20.0	15.0	11.6	12.7
1000	20.0	18.3	15.5	12.9
2000	20.0	18.4	17.1	13.1
4000	20.0	17.9	10.2	13.3
6000	20.0	16.6	14.4	13.1
8000	20.0	18.0	12.8	13.3
16000	20.0	17.9	11.6	13.4

Table 6: Mean approximation ratios by input dataset size, as plotted in Figure 3

n	WO Median Baseline	Uniform Sample Baseline	WSO Algorithm
100	0.499	1.000	1.000
500	0.515	0.817	0.864
1000	0.527	0.693	0.760
1500	0.494	0.603	0.642
2000	0.512	0.567	0.658
2500	0.501	0.522	0.631
3000	0.517	0.531	0.666
3500	0.533	0.467	0.633
4000	0.502	0.505	0.637
5000	0.502	0.453	0.607
6000	0.565	0.337	0.615
7000	0.477	0.408	0.534

Table 8: Mean absolute diameter estimates for $\alpha = 1$ on the NLTK word dataset without replacements. Ten trials per n. Visualized by Figure 2

n	True Metric Diameter	α-WSO	Uniform (SO)	WO-Median (Blind)
100	20	20	20	14
500	20	19	14	14
1000	20	18	16	14
2000	20	18	14	14
4000	20	18	12	14

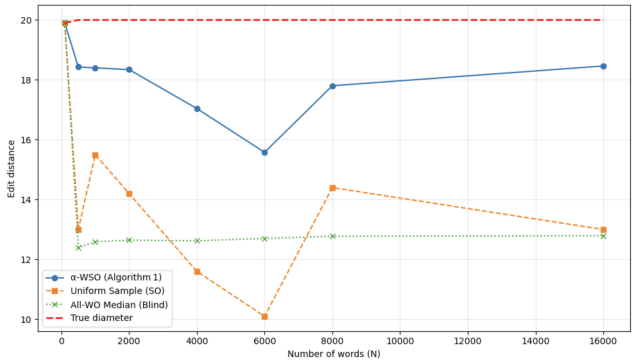


Figure 5: Comparisons of estimated distance on nltk.names corpus with replacements

Table 9: Comparison of metric diameter approximation methods across different dataset sizes for nltk.names corpus. Visualized by Figure 5

N	True Metric Diameter	α-WSO	Uniform (SO)	WO-Median(Blind)
100	20.0	20.0	20.0	13.8
500	20.0	18.4	18.6	14.4
1000	20.0	17.0	10.2	14.1
2000	20.0	18.4	13.0	14.2
4000	20.0	18.3	13.0	14.5
6000	20.0	17.7	17.2	14.6
8000	20.0	17.7	15.8	14.6
16000	20.0	17.9	14.4	14.6