

Smoothed Analysis of Online Metric Matching with a Single Sample: Beyond Metric Distortion

Yingxi Li ✉ 

Department of Management Science and Engineering, Stanford University, CA, USA

Ellen Vitercik ✉ 

Department of Management Science and Engineering, Department of Computer Science,
Stanford University, CA, USA

Mingwei Yang¹ ✉ 

Department of Management Science and Engineering, Stanford University, CA, USA

Abstract

In the online metric matching problem, n servers and n requests lie in a metric space. Servers are available upfront, and requests arrive sequentially. An arriving request must be matched immediately and irrevocably to an available server, incurring a cost equal to their distance. The goal is to minimize the total matching cost.

We study this problem in $[0, 1]^d$ with the Euclidean metric, when servers are adversarial and requests are independently drawn from distinct distributions that satisfy a mild smoothness condition. Our main result is an $O(1)$ -competitive algorithm for $d \neq 2$ that requires no distributional knowledge, relying only on a single sample from each request distribution. To our knowledge, this is the first algorithm to achieve an $o(\log n)$ competitive ratio for non-trivial metrics beyond the i.i.d. setting. Our approach bypasses the $\Omega(\log n)$ barrier introduced by probabilistic metric embeddings: instead of analyzing the embedding distortion and the algorithm separately, we directly bound the cost of the algorithm on the target metric space of a simple deterministic embedding. We then combine this analysis with lower bounds on the offline optimum for Euclidean metrics, derived via majorization arguments, to obtain our guarantees.

2012 ACM Subject Classification Theory of computation → Online algorithms

Keywords and phrases Online algorithm, Metric matching, Competitive analysis, Smoothed analysis

Digital Object Identifier 10.4230/LIPIcs.ITCS.2026.94

Related Version *Full Version*: <https://arxiv.org/abs/2510.20288> [45]

Funding *Yingxi Li*: Research supported in part by NSF grant CCF-2338226.

Ellen Vitercik: Research supported in part by NSF grant CCF-2338226.

Acknowledgements We would like to thank the anonymous reviewers for their many helpful comments and suggestions.

1 Introduction

The online metric matching problem is a classic topic in the design of online algorithms and has been studied extensively for decades. In this problem, n servers are available upfront and n requests arrive sequentially, with all servers and requests lying in a common metric space. Each request must be immediately and irrevocably matched to an unmatched server, incurring a cost equal to their distance. The goal is to minimize the total matching cost.

¹ Corresponding author



© Yingxi Li, Ellen Vitercik, and Mingwei Yang;
licensed under Creative Commons License CC-BY 4.0

17th Innovations in Theoretical Computer Science Conference (ITCS 2026).

Editor: Shubhangi Saraf; Article No. 94; pp. 94:1–94:23

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

The online metric matching problem captures a variety of practical scenarios. In ride-hailing, for example, servers and requests correspond to drivers and passengers, and the cost of a match reflects the pickup distance. The Euclidean special case is also natural for applications such as kidney exchange, where patients and donors are represented by high-dimensional feature vectors and the Euclidean distance between them captures compatibility.

When servers and requests are adversarially chosen, Meyerson et al. [48] give an $O(\log^3 n)$ -competitive algorithm for general metrics, later improved to $O(\log^2 n)$ [10]. For Euclidean metrics, an $O(\log n)$ -competitive algorithm is given by Gupta and Lewi [34]. On the other hand, there exists an $\Omega(\log n)$ competitive-ratio lower bound for uniform metrics² [48], which holds even in the random-order arrival model [51]. Moreover, no algorithm can achieve an $o(\sqrt{\log n})$ competitive ratio for line metrics [50].

The fully adversarial model is often considered overly pessimistic, since worst-case scenarios rarely arise in practice. A more realistic assumption is that requests are i.i.d. sampled, while servers remain adversarially chosen. In this setting, Gupta et al. [32] give an $O((\log \log \log n)^2)$ -competitive algorithm for general metrics and distributions, and their algorithm is $O(1)$ -competitive for tree metrics. Recently, Yang and Yu [58] show that the above setting further reduces to the case where all servers and requests are i.i.d. sampled from the same distribution. Building on this, they obtain an $O(1)$ -competitive algorithm for Euclidean metrics with *smooth* distributions, where a distribution is smooth if it admits a density with respect to the uniform measure upper bounded by a constant.

The i.i.d. model yields nearly tight guarantees, but it is often unrealistic: in many applications, different requests follow markedly different distributions. A more promising approach is *smoothed analysis*, which has been widely successful in both theory and practice, and is frequently used to explain the strong empirical performance of heuristics with poor worst-case guarantees [56, 20, 3]. In this model, the adversary first selects an adversarial input, which is then randomly perturbed by nature via adding a small noise, and the algorithm is required to perform well in expectation. Since the perturbed input always follows a smooth distribution, in the more modern and general formulation of smoothed analysis, the adversary directly specifies a smooth distribution over inputs [37, 22]. In this paper, we adopt the smoothed analysis framework to study online metric matching and obtain improved competitive ratios under significantly weaker assumptions than those in the i.i.d. setting.

1.1 Our Results

Main result. We focus on $[0, 1]^d$ with the Euclidean metric, assuming that servers are adversarial and requests are independently drawn from distinct smooth distributions. As our main result, we give an algorithm which requires no knowledge of the distributions and is $O(1)$ -competitive for $d \neq 2$, given access to one sample from each of the n request distributions (Theorem 15). In particular, the algorithm does not need to know the correspondence between distributions and samples, and its competitive ratio depends polynomially on d and the smoothness parameter, which are usually treated as constants that do not grow with n [34, 42, 58]. To the best of our knowledge, this is the first algorithm to achieve an $o(\log n)$ competitive ratio for non-trivial metrics in a setting beyond the i.i.d. assumption, even when the request distributions are fully known to the algorithm (see Table 1 for a thorough comparison with prior work).

² A metric is *uniform* if every pair of distinct points has distance 1.

■ **Table 1** Comparison with prior work on online metric matching with the state-of-the-art results. Our result is the first to achieve an $o(\log n)$ competitive ratio for non-trivial metrics under distributional assumptions strictly weaker than i.i.d.

Requests	Servers	Metric Space	Distributional Knowledge	Competitive Ratio	Reference
Adversarial	Adversarial	General	None	$O(\log^2 n)$	[10]
Adversarial	Adversarial	Euclidean	None	$O(\log n)$	[34]
i.i.d.	Adversarial	General	Full knowledge	$O((\log \log \log n)^2)$	[32]
i.i.d.	Adversarial	Tree	Full knowledge	$O(1)$	[32]
i.i.d. smooth	Adversarial	Euclidean	$O(n^2)$ samples	$O(1)$ for $d \neq 2$	[58]
i.i.d. uniform	Adversarial	Euclidean	Full knowledge	$O(1)$	[58]
Independent, non-identical, smooth	Adversarial	Euclidean	One sample per distribution	$O(1)$ for $d \neq 2$	Our work

A complementary algorithm. In addition, we present a second algorithm under the same assumptions whose guarantees are not directly comparable (Theorem 15). It is $O(1)$ -competitive for $d \geq 3$ and, while still suboptimal in $d = 2$, it outperforms the first algorithm in that regime, providing a complementary strength.

Discussion. Our sample-based assumption – one sample per request distribution – is a natural and practical assumption, widely adopted in recent work on online algorithms [8, 53, 43, 27, 31]. This assumption is also necessary to some extent, as it is unclear how an algorithm could leverage the distributional properties of requests without some information, even in simpler i.i.d. settings. In comparison, our assumptions are weaker than those in prior work for the i.i.d. model: the algorithm of Gupta et al. [32] requires full knowledge of the distribution, and the algorithm of Yang and Yu [58] needs $O(n^2)$ samples.

Finally, we turn to the role of the dimension d . Two-dimensional space is a critical case for Euclidean matching: prior works highlight that the plane behaves fundamentally differently from both the line and higher dimensions [57, 42, 58]. Consistent with this, our analysis does not yield a constant-competitive bound in $d = 2$. Nevertheless, our second algorithm achieves strictly better performance in two dimensions than our main algorithm, offering partial progress. Closing the gap in two dimensions remains an important direction for future research.

1.2 Technical Overview

Our approach builds on the classical paradigm of metric embeddings but departs in a crucial way. Previous algorithms embed the input metric space into *Hierarchically well-Separated Trees* (HST) and analyze two pieces separately: the distortion of the embedding and the competitive ratio of the algorithm on HSTs [48, 34, 10]. This separation is what forces the $\Omega(\log n)$ barrier. Our key idea is to bypass the distortion step entirely: we analyze the algorithm’s cost directly in the resulting HST, using the non-contractivity of the embedding to argue about the original metric space. This avoids the logarithmic loss while retaining the algorithmic utility of the HST. With this perspective in place, we now recall the definitions of metric embeddings and HSTs.

A (*non-contractive*) *deterministic embedding* of a metric space $(\mathcal{X}, \delta)$ into another metric space $(\mathcal{X}', \delta')$ is a mapping $f : \mathcal{X} \rightarrow \mathcal{X}'$ satisfying $\delta(x, y) \leq \delta'(f(x), f(y))$ for all $x, y \in \mathcal{X}$, and the *distortion* of an embedding f is the smallest $\kappa \geq 1$ such that $\delta'(f(x), f(y)) \leq \kappa \cdot \delta(x, y)$ for all $x, y \in \mathcal{X}$. A *probabilistic embedding* is a distribution over deterministic embeddings, and the distortion of a probabilistic embedding f is the smallest $\kappa \geq 1$ such that $\mathbb{E}[\delta'(f(x), f(y))] \leq \kappa \cdot \delta(x, y)$ for all $x, y \in \mathcal{X}$. It is not hard to see that given an embedding of the input metric space to another (usually simpler) metric space with distortion κ_1 , and an algorithm for the latter metric space with competitive ratio κ_2 , we get an algorithm for the input metric space with competitive ratio $\kappa_1 \cdot \kappa_2$. In other words, a metric embedding with distortion κ reduces the problem for a complicated metric to the same problem for a simple metric with the cost of an additional factor κ in the competitive ratio.

One of the most popular simple metrics is the HST metric, which is induced by the distance function on an HST. In particular, an α -HST for $\alpha \geq 1$ is a rooted tree where all the leaf nodes are at the same depth. Its edge lengths are defined hierarchically: assuming root-adjacent edges have length 1, and for every internal node v , the edge to its parent is α times longer than each edge to its children (see Definition 6 for a formal definition and Figure 1 for an illustration). It is known that every n -point metric space can be randomly embedded into an α -HST with distortion $O(\alpha \log n)$ [29], which, when used as a reduction for online problems, contributes an $O(\log n)$ factor to the competitive ratio.

To bypass the $O(\log n)$ distortion factor, we avoid the usual two-step analysis that handles the distortion of the metric embedding and the competitive ratio for HSTs separately. Instead, we bound directly the cost the algorithm incurred on the resulting HST; by the non-contractivity of the embedding, this also upper-bounds the algorithm's cost in the original metric space. We then compare this upper bound to a lower bound for the offline optimum in the original metric space, yielding the desired competitive ratio. Notably, a simple deterministic embedding – despite having unbounded worst-case distortion – suffices for our analysis.

In more detail, we use the canonical dyadic partition of $[0, 1]^d$ to define a 2^d -ary 2-HST of height h (a tunable parameter we optimize later), where an HST is Δ -ary if every internal node has exactly Δ children. To describe the embedding, the root of the HST corresponds to $[0, 1]^d$, each node is partitioned into 2^d subcubes with half the side-length, and the leaves at depth h correspond to disjoint subcubes of side-length 2^{-h} (see Figure 2 for an illustration). Then, we map each point $x \in [0, 1]^d$ to the unique leaf node whose corresponding cube contains x .

Equipped with the metric embedding, the problem reduces to designing algorithms for HST metrics, where we adopt the *Random-Subtree* algorithm of Gupta and Lewi [34], and the algorithm of Bansal et al. [10]. Our key observation is that the expected cost analysis for both algorithms reduces to bounding the fluctuations in the number of requests in each subtree. By controlling these fluctuations (via standard deviations of Poisson–binomial counts) and applying concavity, we show that the worst case occurs when requests are uniformly distributed across children. This yields distribution-free upper bounds for the cost (Theorems 7 and 11), so no smoothness assumption is needed for the upper bounds.

Smoothness enters only in the lower bound on the offline optimum (Lemma 17). For $d \geq 2$, the bound follows directly from standard nearest-neighbor distance estimates. The proof for $d = 1$ is subtler. Any imbalance between the numbers of servers and requests in a subinterval of $[0, 1]$ forces a proportional number of matches to cross the subinterval's boundary, contributing to the total cost. Inspecting all subintervals of length $L \in (0, 1)$ yields a lower bound for the offline optimum, which we refer to as *obstacle to matching at length*

scale L . It is known that, when all servers and requests are uniformly distributed, the largest obstacle to matching for $d = 1$ occurs at the scale of a constant length [42]. To generalize this reasoning to the case where requests are drawn from distinct smooth distributions, we derive a lower bound for the obstacle to matching at a length scale proportional to the smoothness parameter. Using majorization and concavity arguments, we show that this lower bound is minimized when requests are as concentrated as possible. The desired lower bound then follows from the anti-concentration properties of smooth distributions.

Finally, we apply the result in [58] to incorporate the provided samples from the request distributions in a black-box manner (Lemma 5). Combining the above upper bounds for the algorithms with the lower bounds for the offline optimum yields our main result.

1.3 Related Work

Further related work on online metric matching. For adversarial servers and requests, Kalyanasundaram and Pruhs [40], Khuller et al. [44], and Raghvendra [51] give $(2n - 1)$ -competitive deterministic algorithms for general metrics, which is optimal. Raghvendra [51] provides a primal-dual deterministic algorithm that is $O(\log n)$ -competitive for general metrics in the random-order arrival model, which is later shown to exhibit near instance-optimality [49] and an $O(\log n)$ competitive ratio for line metrics [52].

When all servers and requests are uniformly distributed on the Euclidean space, Kanoria [42] gives an $O(1)$ -competitive algorithm, which applies the same deterministic metric embedding as ours. Nevertheless, they employ a different algorithm for the HST metric, and it is unclear how to generalize their analysis to non-identical request distributions. When all servers and requests are i.i.d. drawn from a general distribution, Chen et al. [21] present an algorithm for Euclidean metrics with nearly optimal *regret*, which is defined as the difference between the cost of the algorithm and the offline optimum. Balkanski et al. [9] show that the *Greedy* algorithm is $O(1)$ -competitive for line metrics when all servers and requests are uniformly distributed.

Akbarpour et al. [1] initiate the study of *unbalanced markets*, where servers outnumber requests by a constant factor, and they show that Greedy is $O(\log^3 n)$ -competitive when all servers and requests are uniformly distributed on a line, which is subsequently improved to $O(1)$ [9]. Kanoria [42] gives an $O(1)$ -competitive algorithm for unbalanced markets when all servers and requests are uniformly distributed on the Euclidean space. When servers are adversarial and requests are i.i.d. drawn from a smooth distribution, Yang and Yu [58] achieve competitive-ratio and regret guarantees for unbalanced markets and Euclidean metrics.

Several variants of online metric matching have also been studied, which include online transportation [38, 6], online metric matching with recourse [35, 47], online min-cost perfect matching with delays [28, 7], online metric matching with stochastic arrivals and departures [5, 2], online min-weight perfect matching [16], and online matching in geometric random graphs [55].

Smoothed analysis of online problems. For smoothed analysis of other online problems, prior studies primarily focus on online learning [41, 36, 17, 26, 37] and online discrepancy minimization [12, 13, 37], whose goal is to minimize regret. Regarding smoothed competitive analysis, Becchetti et al. [14] consider the multi-level feedback algorithm for non-clairvoyant scheduling, and Schäfer and Sivadasan [54] analyze the work function algorithm for metrical task systems. More recently, Coester and Umenberger [22] conduct smoothed analysis of classic online metric problems including k -server, k -taxi, and chasing small sets, where requests are drawn from smooth distributions, and they achieve significantly improved competitive

ratios compared to the adversarial setting. In particular, they allow the distribution followed by each request to depend on the realizations of the past requests and the decisions made by the algorithm thus far, and their algorithms require no distributional knowledge. However, their technique does not result in improved algorithms for online metric matching.

Sample complexity of online problems. A growing line of work studies online algorithms that receive samples from the underlying distributions to go beyond the strong assumption of knowing the distributions in full. Pioneered by Azar et al. [8], competitive guarantees are achieved for sample-based prophet inequalities under various combinations of the arrival model, the number of samples, and combinatorial constraints [24, 53, 19, 23, 25, 30]. This paradigm has also been adopted by literature on online resource allocation [27, 31] and online weighted matching [43].

2 Preliminaries

Let $(\mathcal{X}, \delta)$ be a metric space. There are n servers $S = \{s_1, \dots, s_n\}$ available at time $t = 0$, whose locations are known to the algorithm, and n requests $R = (r_1, \dots, r_n)$ that arrive sequentially. At each time step $t \in [n]$, the location of the request r_t is revealed, and the algorithm must immediately and irrevocably match it to an unmatched server. The cost of matching a request r and a server s is $\delta(r, s)$. We assume that servers are adversarial, and requests are independently drawn from distributions $\mathbb{D}_1, \dots, \mathbb{D}_n$ support on $\mathcal{X}$. Let $\mathbb{D} := \prod_{i=1}^n \mathbb{D}_i$ be the joint request distribution, which is not known to the algorithm.

Given a (randomized) algorithm $\mathcal{A}$, let $\text{cost}(\mathcal{A}; S, R)$ be the expected cost of $\mathcal{A}$ for server set S and request sequence R . For a distribution $\mathbb{D}$ over request sequences, define $\text{cost}(\mathcal{A}; S, \mathbb{D}) := \mathbb{E}_{R \sim \mathbb{D}}[\text{cost}(\mathcal{A}; S, R)]$. Given S and R , let $\text{OPT}(S, R)$ be the minimum cost of a perfect matching between S and R . Similarly, let $\text{OPT}(S, \mathbb{D}) := \mathbb{E}_{R \sim \mathbb{D}}[\text{OPT}(S, R)]$ be the expected optimal cost when requests are drawn from $\mathbb{D}$, and $\text{OPT}(\mathbb{D}, \mathbb{D}) := \mathbb{E}_{S \sim \mathbb{D}}[\text{OPT}(S, \mathbb{D})]$ be the expected optimal cost when servers and requests are all drawn from $\mathbb{D}$. We say that an algorithm $\mathcal{A}$ is α -competitive for $\alpha \geq 1$ if for all S and $\mathbb{D}$, $\text{cost}(\mathcal{A}; S, \mathbb{D}) \leq \alpha \cdot \text{OPT}(S, \mathbb{D})$. Finally, given a matching M between two sets S and R , we denote the element in R that is matched to $s \in S$ as $M(s)$.

► **Definition 1** (Smoothness). *We say that a measure μ over $\mathcal{X}$, which supports a uniform distribution $\mathbb{U}$, is σ -smooth for $\sigma \in (0, 1]$ if for every measurable subset $\mathcal{X}' \subseteq \mathcal{X}$, $\mu(\mathcal{X}') \leq \mathbb{U}(\mathcal{X}')/\sigma$.*

2.1 Majorization and Poisson Binomial

For a vector $\mathbf{p} \in \mathbb{R}^n$, we use $p_{[i]}$ for $i \in [n]$ to denote the i -th largest element among $p_1, \dots, p_n$.

► **Definition 2** (Majorization). *For $\mathbf{p}, \mathbf{q} \in \mathbb{R}^n$, we say that $\mathbf{p}$ majorizes $\mathbf{q}$, denoted as $\mathbf{p} \succ \mathbf{q}$, if*

- $\sum_{i=1}^k p_{[i]} \geq \sum_{i=1}^k q_{[i]}$ for $k \in [n-1]$, and
- $\sum_{i=1}^n p_{[i]} = \sum_{i=1}^n q_{[i]}$.

For $\mathbf{p} \in [0, 1]^n$, let $\text{PB}(\mathbf{p})$ denote the *Poisson binomial* random variable $\sum_{i=1}^n X_i$ with independent $X_i \sim \text{Ber}(p_i)$. For $\mathbf{p}, \mathbf{q} \in [0, 1]^n$ with $\mathbf{p} \succ \mathbf{q}$, it is known that $\text{PB}(\mathbf{p})$ is “less spread out” than $\text{PB}(\mathbf{q})$ in the sense of *convex order*, which we formally present in the following lemma.

► **Lemma 3** (Proposition 12.F.1 in [46]). *For $\mathbf{p}, \mathbf{q} \in [0, 1]^n$, if $\mathbf{p} \succ \mathbf{q}$, then $\mathbb{E}[f(\text{PB}(\mathbf{p}))] \leq \mathbb{E}[f(\text{PB}(\mathbf{q}))]$ for every convex function $f : \{0, 1, \dots, n\} \rightarrow \mathbb{R}$.*

The following corollary, which compares the mean absolute deviations of two Poisson binomial random variables, is a direct consequence of Lemma 3 since $f(t) = |t - c|$ is convex.

► **Corollary 4.** *For $\mathbf{p}, \mathbf{q} \in [0, 1]^n$ with $\mathbf{p} \succ \mathbf{q}$, $\mathbb{E}[|\text{PB}(\mathbf{p}) - \sum_{i=1}^n p_i|] \leq \mathbb{E}[|\text{PB}(\mathbf{q}) - \sum_{i=1}^n q_i|]$.*

2.2 Algorithm with Sample Access to Request Distributions

Yang and Yu [58] show that given an arbitrary algorithm $\mathcal{A}$ and one sample from each request distribution, we can construct another algorithm whose cost can be decomposed into the offline optimum and the cost of $\mathcal{A}$ when the server set is also drawn from $\mathbb{D}$.³

► **Lemma 5** (Theorem 2 in [58]). *Given an algorithm $\mathcal{A}$ and one sample from each request distribution, there exists an algorithm $\mathcal{A}'$ such that $\text{cost}(\mathcal{A}'; S, \mathbb{D}) \leq \text{OPT}(S, \mathbb{D}) + \text{cost}(\mathcal{A}; \mathbb{D}, \mathbb{D})$. In particular, $\mathcal{A}'$ does not need to know the correspondence between distributions and samples.*

2.3 Hierarchically Well-Separated Trees

In this subsection, we introduce Hierarchically well-Separated Trees.

► **Definition 6** (HSTs). *Given $\alpha \geq 1$, an α -Hierarchically well-Separated Tree (α -HST) is a rooted tree $T = (V, E)$ along with a distance function $\delta : E \rightarrow \mathbb{R}_{\geq 0}$ on the edges that satisfies the following properties:*

1. *For each internal node v , all edges from v to its children have the same length.*
2. *For each node v , if $p(v)$ is the parent of v , and $c(v)$ is a child of v , then $\delta(p(v), v) = \alpha \cdot \delta(v, c(v))$.*
3. *For all leaf nodes v_1 and v_2 , let $p(v_1)$ and $p(v_2)$ be the parents of v_1 and v_2 , respectively. Then, $\delta(p(v_1), v_1) = \delta(p(v_2), v_2)$.*


Moreover, an HST is Δ -ary if each internal node has exactly Δ children, and the height of an HST is defined as the number of edges in the path from the root to any leaf node.

By normalization, we assume the length of each root-edge of an HST to be 1. See Figure 1 for an illustration of a 4-ary α -HST. An HST naturally induces a metric over V , where the distance between two nodes is the length of the unique tree path between them. For HST metrics, unless stated otherwise, we assume that servers and requests are located at the leaf nodes.

For each internal node v of a Δ -ary HST, let $c_i(v)$ be the i -th child of v for $i \in [\Delta]$. Define $\hat{s}(v)$ and $\hat{r}(v)$ as the number of servers and requests, respectively, in the subtree rooted at v . The height of a node v is defined as the number of edges in the path between v and any leaf node in the subtree rooted at v . For an HST with height h , we use V_j to denote the set of nodes with height j for each $j \in \{0, 1, \dots, h\}$.

When the request sequence is drawn from $\mathbb{D} = \prod_{i=1}^n \mathbb{D}_i$, for each node v , let $\mu_{\mathbb{D}_i}(v)$ be the probability that r_i is in the subtree rooted at v , and let $\mu_{\mathbb{D}}(v) := \sum_{i=1}^n \mu_{\mathbb{D}_i}(v)$ be the expected number of requests in the subtree rooted at v . Note that $\hat{r}(v) \sim \text{PB}(\mu_{\mathbb{D}_1}(v), \dots, \mu_{\mathbb{D}_n}(v))$.

³ They only prove this statement for i.i.d. requests, but the generalization to the case with non-identical request distributions follows verbatim.



■ **Figure 1** A 4-ary α -HST.

3 Random-Subtree Algorithm

The *Random-Subtree (RS)* algorithm proposed by Gupta and Lewi [34] is a natural starting point for online matching on HST metrics. In the adversarial setting, it achieves an $O(\log \Delta)$ competitive ratio for Δ -ary α -HST metrics, provided $\alpha = \Omega(\log \Delta)$. Here we revisit the RS algorithm in the stochastic setting where servers and requests are independently drawn from a distribution $\mathbb{D} = \prod_{i=1}^n \mathbb{D}_i$. Later on, we will translate the guarantees to the setting where servers are adversarial and requests are drawn from $\mathbb{D}$ by applying Lemma 5. Rather than its competitive ratio, we upper bound the cost of the RS algorithm for arbitrary $\alpha \geq 2$.

We start by describing the RS algorithm, which applies to any HST metric. For each arriving request r , let v be the lowest ancestor of r such that the subtree rooted at v contains at least one available server. Starting from v , the algorithm descends the tree toward a leaf guaranteed to contain an available server. At each internal node v , it selects uniformly at random among the children whose subtrees contain an available server, and continues this process until reaching such a leaf. The request r is then matched to a server at this leaf.

We upper bound the cost of the RS algorithm in the following theorem, where we use $H_k := 1 + 1/2 + \dots + 1/k$ to denote the k -th harmonic number.

► **Theorem 7.** *For any Δ -ary α -HST with height h and $\alpha \geq 2$, and for any $\mathbb{D} = \prod_{i=1}^n \mathbb{D}_i$,*

$$\text{cost}(\text{RS}; \mathbb{D}, \mathbb{D}) \leq 6H_\Delta \sqrt{n} \sum_{j=0}^{h-1} \left(\frac{\sqrt{\Delta^{h-j}}}{\alpha^{h-j-1}} \sum_{\ell=0}^j \xi^\ell \right),$$

where $\xi := H_\Delta / \alpha$.

The bound in Theorem 7 admits a natural interpretation. Suppose, for intuition, that every $\mathbb{D}_i$ is uniform over the leaf nodes – this turns out to be the worst case of our analysis. To facilitate interpretation, we rewrite the bound as

$$6 \sum_{j=0}^{h-1} \sqrt{n\Delta^{h-j}} \sum_{\ell=0}^j \frac{H_\Delta^{\ell+1}}{\alpha^{h-1-j+\ell}}. \quad (1)$$

The term $\sqrt{n\Delta^{h-j}} \approx \sum_{v \in V_j} \mathbb{E}[|\hat{r}(v) - \hat{s}(v)|]$ corresponds to the expected discrepancy between servers and requests at level j , i.e., in the subtrees rooted at nodes in V_j . Each discrepancy at level j contributes to at most H_Δ mismatches made by the algorithm, and each such

mismatch incurs a cost on the order of α^{j-h+1} . Moreover, a mismatch at level j can cascade downward, creating up to H_Δ additional mismatches at level $j-1$, which then trigger up to H_Δ^2 additional mismatches at level $j-2$, and so on. This propagation explains the second summation in (1).

The proof of Theorem 7 relies on the following bound for the RS algorithm, whose cost is upper bounded in terms of the number of excess requests in each subtree.

► **Lemma 8.** *For any Δ -ary α -HST with height h and $\alpha \geq 2$, and for all S and R ,*

$$\text{cost}(\text{RS}; S, R) \leq 3H_\Delta \sum_{j=0}^{h-1} \sum_{v \in V_j} \frac{(\hat{r}(v) - \hat{s}(v))^+}{\alpha^{h-j-1}} \sum_{\ell=0}^j \xi^\ell,$$

where $\xi := H_\Delta/\alpha$.

We defer the proof of Lemma 8 to Section 3.1 and proceed to finish the proof of Theorem 7.

Proof of Theorem 7. By Lemma 8, for $\xi := H_\Delta/\alpha$,

$$\text{cost}(\text{RS}; \mathbb{D}, \mathbb{D}) \leq 3H_\Delta \sum_{j=0}^{h-1} \alpha^{j-h+1} \sum_{\ell=0}^j \xi^\ell \sum_{v \in V_j} \mathbb{E}[(\hat{r}(v) - \hat{s}(v))^+]. \quad (2)$$

Next, we upper bound the expected number of excess requests for each height j .

For every node v , since $\hat{r}(v)$ and $\hat{s}(v)$ are identically distributed,

$$\begin{aligned} \mathbb{E}[(\hat{r}(v) - \hat{s}(v))^+] &\leq \mathbb{E}[|\hat{r}(v) - \mathbb{E}\hat{r}(v)|] + \mathbb{E}[|\hat{s}(v) - \mathbb{E}\hat{s}(v)|] \\ &= 2\mathbb{E}[|\hat{r}(v) - \mathbb{E}\hat{r}(v)|] \leq 2 \cdot \text{std}(\hat{r}(v)), \end{aligned} \quad (3)$$

where the last inequality holds by Jensen's inequality. Moreover,

$$\text{std}(\hat{r}(v)) = \text{std}(\text{PB}(\mu_{\mathbb{D}_1}(v), \dots, \mu_{\mathbb{D}_n}(v))) \leq \sqrt{\sum_{i=1}^n \mu_{\mathbb{D}_i}(v)} = \sqrt{\mu_{\mathbb{D}}(v)}. \quad (4)$$

As a result, for each $j \in \{0, 1, \dots, h-1\}$,

$$\sum_{v \in V_j} \mathbb{E}[(\hat{r}(v) - \hat{s}(v))^+] \leq 2 \sum_{v \in V_j} \sqrt{\mu_{\mathbb{D}}(v)} \leq 2\sqrt{n|V_j|} = 2\sqrt{n\Delta^{h-j}}, \quad (5)$$

where the first inequality follows from (3) and (4), and the second inequality holds by the concavity of $f(t) = \sqrt{t}$ and the fact that $\sum_{v \in V_j} \mu_{\mathbb{D}}(v) = n$.

Finally, combining (2) and (5) concludes the proof. ◀

3.1 Proof of Lemma 8

The proof of Lemma 8 largely follows the proof strategy of [34, Theorem 4.1], with two key differences. First, their analysis assumes $\alpha = \Omega(H_\Delta)$, ensuring the HST is sufficiently well-separated so that the costs incurred at lower levels are dominated by those at higher levels. In contrast, we do not rely on this separation, and hence the costs incurred at lower levels must be handled explicitly. Second, our proof is relatively simpler: since we only bound the absolute cost of the algorithm, we avoid the more involved step of characterizing the offline optimum, which is required in [34].

Recall that the length of each root-edge is 1, so the length of each root-leaf path is $1 + \beta$, where $\beta \leq 1/(\alpha - 1) \leq 1$. In the proof, instead of assuming all requests lie only at the leaf nodes, we allow some requests to lie at the root, and the matching process of these requests by the RS algorithm follows the same description. We also permit the number of servers to exceed the number of requests. For a server set S , leaf requests R , and root requests R' with $|S| \geq |R| + |R'|$, let $\text{cost}(\text{RS}; S, R \cup R')$ denote the expected cost of the RS algorithm. We will prove the following stronger statement:

$$\text{cost}(\text{RS}; S, R \cup R') \leq 3H_\Delta \left(|R'| \sum_{j=0}^{h-1} \xi^j + \sum_{j=0}^{h-1} \sum_{v \in V_j} \frac{(\hat{r}(v) - \hat{s}(v))^+}{\alpha^{h-j-1}} \sum_{\ell=0}^j \xi^\ell \right),$$

which gives the desired bound by setting $R' = \emptyset$.

We use v_r to denote the root, and let γ be the expected number of requests in $R \cup R'$ that traverse a root-edge in the matching produced by the RS algorithm. The following lemma upper bounds γ in terms of $|R'|$ and the number of excess requests in the subtrees of v_r .

► **Lemma 9** (Lemma 4.3 in [34]). *It holds that*

$$\gamma \leq H_\Delta \left(|R'| + \sum_{i=1}^{\Delta} (\hat{r}(c_i(v_r)) - \hat{s}(c_i(v_r)))^+ \right).$$

We analyze the cost of the RS algorithm by induction on h . We start from the base case where the HST is a star with $h = 1$. Since $\text{cost}(\text{RS}; S, R \cup R') \leq 2\gamma$, by Lemma 9,

$$\text{cost}(\text{RS}; S, R \cup R') \leq 2H_\Delta \left(|R'| + \sum_{i=1}^{\Delta} (\hat{r}(c_i(v_r)) - \hat{s}(c_i(v_r)))^+ \right),$$

concluding the proof for the base case.

Now, we assume that $h \geq 2$. For $i \in [\Delta]$, let T_i be the i -th subtree of v_r , rooted at $c_i(v_r)$. Let S_i and R_i denote the servers and requests contained in T_i , respectively. Note that $S = \bigcup_{i=1}^{\Delta} S_i$ and $R = \bigcup_{i=1}^{\Delta} R_i$. Let M_i be the set of requests outside T_i that the RS algorithm matches to servers in T_i . Then, $R_i \cup M_i$ are all the requests that are either contained in T_i or matched to servers within T_i . Let $k := \min\{|S_i|, |R_i \cup M_i|\}$. Order the requests in $R_i \cup M_i$ by their arrival time, and let X_i be the first k of them. Since T_i contains $|S_i|$ servers and the algorithm never bypasses an available server, the requests matched within T_i are exactly those in X_i . Define $\hat{R}_i := X_i \cap R_i$. Note that $M_i \subseteq X_i$, and M_i , X_i , and $\hat{R}_i$ are random variables.

We recall the following useful decomposition of $\text{cost}(\text{RS}; S, R \cup R')$ from [34].

► **Lemma 10** (Fact 4.8 in [34]). *It holds that*

$$\text{cost}(\text{RS}; S, R \cup R') = \sum_{i=1}^{\Delta} \mathbb{E}[\text{cost}(\text{RS}; S_i, \hat{R}_i \cup M_i)] + \sum_{i=1}^{\Delta} \mathbb{E}[|M_i|] \cdot (2 + \beta),$$

where the expectations are taken over the randomness of the RS algorithm.

For each $i \in [\Delta]$, let $\eta_i := (\hat{r}(c_i(v_r)) - \hat{s}(c_i(v_r)))^+$ be the number of excess requests in T_i . By Lemma 9,

$$\sum_{i=1}^{\Delta} \mathbb{E}[|M_i|] = \gamma \leq H_\Delta \left(|R'| + \sum_{i=1}^{\Delta} \eta_i \right). \quad (6)$$

By Lemma 10 and (6),

$$\begin{aligned} \text{cost}(\text{RS}; S, R \cup R') &= \sum_{i=1}^{\Delta} \mathbb{E}[\text{cost}(\text{RS}; S_i, \hat{R}_i \cup M_i)] + \sum_{i=1}^{\Delta} \mathbb{E}[|M_i|] \cdot (2 + \beta) \\ &\leq \sum_{i=1}^{\Delta} \mathbb{E}[\text{cost}(\text{RS}; S_i, \hat{R}_i \cup M_i)] + 3H_{\Delta} \left(|R'| + \sum_{i=1}^{\Delta} \eta_i \right). \end{aligned} \quad (7)$$

Next, we upper bound the first term in (7). By the inductive hypothesis, and since the length of each root-edge of each subtree T_i equals $1/\alpha$ and $\hat{R}_i \subseteq R_i$,

$$\begin{aligned} &\sum_{i=1}^{\Delta} \mathbb{E}[\text{cost}(\text{RS}; S_i, \hat{R}_i \cup M_i)] \\ &\leq \frac{1}{\alpha} \cdot 3H_{\Delta} \left(\sum_{i=1}^{\Delta} \mathbb{E}[|M_i|] \sum_{j=0}^{h-2} \xi^j + \sum_{j=0}^{h-2} \sum_{v \in V_j} \frac{(\hat{r}(v) - \hat{s}(v))^+}{\alpha^{h-j-2}} \sum_{\ell=0}^j \xi^{\ell} \right) \\ &\leq 3H_{\Delta} \left(\left(|R'| + \sum_{i=1}^{\Delta} \eta_i \right) \sum_{j=1}^{h-1} \xi^j + \sum_{j=0}^{h-2} \sum_{v \in V_j} \frac{(\hat{r}(v) - \hat{s}(v))^+}{\alpha^{h-j-1}} \sum_{\ell=0}^j \xi^{\ell} \right), \end{aligned}$$

where the second inequality holds by (6). Combining the above two displayed equations, we get

$$\begin{aligned} \text{cost}(\text{RS}; S, R \cup R') &\leq 3H_{\Delta} \left(\left(|R'| + \sum_{i=1}^{\Delta} \eta_i \right) \sum_{j=0}^{h-1} \xi^j + \sum_{j=0}^{h-2} \sum_{v \in V_j} \frac{(\hat{r}(v) - \hat{s}(v))^+}{\alpha^{h-j-1}} \sum_{\ell=0}^j \xi^{\ell} \right) \\ &= 3H_{\Delta} \left(|R'| \sum_{j=0}^{h-1} \xi^j + \sum_{j=0}^{h-1} \sum_{v \in V_j} \frac{(\hat{r}(v) - \hat{s}(v))^+}{\alpha^{h-j-1}} \sum_{\ell=0}^j \xi^{\ell} \right), \end{aligned}$$

where the equality holds by the definition of η_i 's, concluding the induction.

4 BBGN Algorithm

In this section, we turn from the RS algorithm to the algorithm proposed by Bansal et al. [10], henceforth referred to as the BBGN algorithm, which outperforms the RS algorithm in certain parameter regimes. The BBGN algorithm is $O(\log n)$ -competitive for α -HST metrics with any constant $\alpha > 1$ when both S and R are adversarial. Our interest, however, lies in analyzing its cost when both S and R are drawn from $\mathbb{D}$. As before, we will later translate these guarantees to the setting where servers are adversarial and requests are drawn from $\mathbb{D}$ by applying Lemma 5.

We now state the main theorem of this section.

► **Theorem 11.** *For any Δ -ary α -HST with height h and constant $\alpha \geq 2$ such that $\Delta^{h-1} \leq n/2$, and for any $\mathbb{D} = \prod_{i=1}^n \mathbb{D}_i$,*

$$\text{cost}(\text{BBGN}; \mathbb{D}, \mathbb{D}) \leq O \left(\sqrt{n} \sum_{j=1}^h \frac{\sqrt{\Delta^{h-j+1}}}{\alpha^{h-j}} \cdot \log \frac{n}{\Delta^{h-j}} \right).$$

The rest of this section is devoted to proving Theorem 11. For each leaf node v of the HST and $\ell \in \{0, 1, \dots, h\}$, let $p(v, \ell)$ be the ancestor of v with height ℓ . The BBGN algorithm was originally developed in a *restricted reassignment* online model, which permits limited

reassignment of previously arrived requests, and then adapted to the standard online model. We recall here a bound for the cost of the BBGN algorithm that applies to all S and R , obtained by combining [10, Lemmas 4.2, 4.3, and 4.4], and we refer to [10] for a more formal description of the algorithm.

► **Lemma 12** ([10]). *Consider an α -HST with $\alpha \geq 2$ being a constant. For all S and R , in addition to the resulting (random) matching M between S and R , the BBGN algorithm also produces a (random) min-cost perfect matching M_{OPT} between S and R such that*

$$\mathbb{E} \left[\sum_{r \in R} \delta(M(r), r) \mid M_{\text{OPT}} \right] \leq O \left(\sum_{r \in R} \sum_{\ell=1}^{L(r)} \alpha^{\ell-h} \cdot \log \hat{s}(p(r, \ell)) \right),$$

where $L(r)$ is the height of the least common ancestor of r and $M_{\text{OPT}}(r)$.

It is well known that min-cost perfect matchings in HSTs admit a useful characterization (see, e.g., [10, Lemma 4.1]): every min-cost perfect matching between S and R can be obtained by a greedy algorithm that repeatedly matches the closest request-server pair and then recurses on the remaining instance. Equivalently, the number of matched pairs whose two endpoints lie in different subtrees of a given internal node can be expressed as follows.

▷ **Claim 13.** For any HST, in any min-cost perfect matching between S and R , the number of matches whose two endpoints lie in different subtrees of an internal node v equals

$$\min \left\{ \sum_{i=1}^{\Delta} (\hat{r}(c_i(v)) - \hat{s}(c_i(v)))^+, \sum_{i=1}^{\Delta} (\hat{s}(c_i(v)) - \hat{r}(c_i(v)))^+ \right\}.$$

Combining Lemma 12 and Claim 13, we can then upper bound the expected cost of the BBGN algorithm as in the following lemma.

► **Lemma 14.** *For any α -HST with $\alpha \geq 2$ being a constant, and for all S and R ,*

$$\text{cost}(\text{BBGN}; S, R) \leq O \left(\sum_{j=1}^h \alpha^{j-h} \sum_{v \in V_j} \log(\hat{s}(v) + e) \sum_{i=1}^{\Delta} |\hat{r}(c_i(v)) - \hat{s}(c_i(v))| \right).$$

Proof. By Lemma 12, the BBGN algorithm produces an additional min-cost perfect matching M_{OPT} between S and R . For each $r \in R$, let $L(r)$ be the height of the least common ancestor of r and $M_{\text{OPT}}(r)$. For each internal node v , let $q(v)$ be the number of requests r such that $p(r, L(r)) = v$, which also equals the number of matches in M_{OPT} whose two endpoints lie in different subtrees of v . By Claim 13,

$$\begin{aligned} q(v) &= \min \left\{ \sum_{i=1}^{\Delta} (\hat{r}(c_i(v)) - \hat{s}(c_i(v)))^+, \sum_{i=1}^{\Delta} (\hat{s}(c_i(v)) - \hat{r}(c_i(v)))^+ \right\} \\ &\leq \sum_{i=1}^{\Delta} |\hat{r}(c_i(v)) - \hat{s}(c_i(v))|. \end{aligned} \tag{8}$$

If we denote M as the (random) matching resulting from the BBGN algorithm, then by Lemma 12,

$$\begin{aligned}
& \mathbb{E} \left[\sum_{r \in R} \delta(M(r), r) \mid M_{\text{OPT}} \right] \\
& \leq O \left(\sum_{r \in R} \sum_{\ell=1}^{L(r)} \alpha^{\ell-h} \cdot \log \hat{s}(p(r, \ell)) \right) \\
& \leq O \left(\sum_{r \in R} \sum_{\ell=1}^{L(r)} \alpha^{\ell-h} \cdot \log \hat{s}(p(r, L(r))) \right) \\
& \leq O \left(\sum_{r \in R} \mathbf{1}\{L(r) \geq 1\} \cdot \alpha^{L(r)-h} \cdot \log \hat{s}(p(r, L(r))) \right) \\
& \leq O \left(\sum_{j=1}^h \alpha^{j-h} \sum_{v \in V_j} \log(\hat{s}(v) + e) \cdot q(v) \right) \\
& \leq O \left(\sum_{j=1}^h \alpha^{j-h} \sum_{v \in V_j} \log(\hat{s}(v) + e) \sum_{i=1}^{\Delta} |\hat{r}(c_i(v)) - \hat{s}(c_i(v))| \right), \tag{10}
\end{aligned}$$

where the third inequality holds since $\alpha \geq 2$, and the last inequality holds by (8). Since (10) does not depend on M_{OPT} , it is also an upper bound for $\text{cost}(\text{BBGN}; S, R)$, concluding the proof. $\blacktriangleleft$

Now, we are ready to finish the proof of Theorem 11.

Proof of Theorem 11. By Lemma 14,

$$\begin{aligned}
& \text{cost}(\text{BBGN}; \mathbb{D}, \mathbb{D}) \\
& \leq O \left(\sum_{j=1}^h \alpha^{j-h} \sum_{v \in V_j} \sum_{i=1}^{\Delta} \mathbb{E}[\log(\hat{s}(v) + e) \cdot |\hat{r}(c_i(v)) - \hat{s}(c_i(v))|] \right) \\
& \leq O \left(\sum_{j=1}^h \alpha^{j-h} \sum_{v \in V_j} \sqrt{\mathbb{E}[\log^2(\hat{s}(v) + e)]} \sum_{i=1}^{\Delta} \sqrt{\mathbb{E}[(\hat{r}(c_i(v)) - \hat{s}(c_i(v)))^2]} \right), \tag{11}
\end{aligned}$$

where the second inequality holds by the Cauchy-Schwarz inequality.

We first bound the last summation in (11). For all internal node v and $i \in [\Delta]$,

$$\begin{aligned}
\sqrt{\mathbb{E}[(\hat{r}(c_i(v)) - \hat{s}(c_i(v)))^2]} & \leq \sqrt{2\mathbb{E}[(\hat{r}(c_i(v)) - \mu_{\mathbb{D}}(c_i(v)))^2 + (\hat{s}(c_i(v)) - \mu_{\mathbb{D}}(c_i(v)))^2]} \\
& = 2 \cdot \text{std}(\hat{r}(c_i(v))) \leq 2\sqrt{\mu_{\mathbb{D}}(c_i(v))},
\end{aligned}$$

where the first inequality follows since $(x + y)^2 \leq 2(x^2 + y^2)$, the equality holds since $\hat{r}(c_i(v))$ and $\hat{s}(c_i(v))$ are identically distributed with $\mathbb{E}[\hat{r}(c_i(v))] = \mathbb{E}[\hat{s}(c_i(v))] = \mu_{\mathbb{D}}(c_i(v))$, and the last inequality holds by (4). It follows that

$$\sum_{i=1}^{\Delta} \sqrt{\mathbb{E}[(\hat{r}(c_i(v)) - \hat{s}(c_i(v)))^2]} \leq 2 \sum_{i=1}^{\Delta} \sqrt{\mu_{\mathbb{D}}(c_i(v))} \leq 2\sqrt{\Delta \cdot \mu_{\mathbb{D}}(v)}, \tag{12}$$

where the last inequality holds by the concavity of $f(t) = \sqrt{t}$ and the fact that $\sum_{i=1}^{\Delta} \mu_{\mathbb{D}}(c_i(v)) = \mu_{\mathbb{D}}(v)$.

Next, we bound $\mathbb{E}[\log^2(\hat{s}(v) + e)]$ for every internal node v . Since $f(t) = \log^2(t + e)$ is concave over $[0, +\infty)$, by Jensen's inequality,

$$\mathbb{E}[\log^2(\hat{s}(v) + e)] \leq \log^2(\mathbb{E}[\hat{s}(v)] + e) = \log^2(\mu_{\mathbb{D}}(v) + e). \quad (13)$$

Finally, combining (11), (12), and (13),

$$\begin{aligned} \text{cost}(\text{BBGN}; \mathbb{D}, \mathbb{D}) &\leq O\left(\sqrt{\Delta} \sum_{j=1}^h \alpha^{j-h} \sum_{v \in V_j} \log(\mu_{\mathbb{D}}(v) + e) \sqrt{\mu_{\mathbb{D}}(v)}\right) \\ &\leq O\left(\sqrt{\Delta} \sum_{j=1}^h \alpha^{j-h} \cdot |V_j| \cdot \sqrt{\frac{n}{|V_j|}} \log\left(\frac{n}{|V_j|} + e\right)\right) \\ &\leq O\left(\sqrt{n\Delta} \sum_{j=1}^h \alpha^{j-h} \sqrt{\Delta^{h-j}} \log \frac{n}{\Delta^{h-j}}\right), \end{aligned}$$

where the second inequality holds by the concavity of $f(t) = \sqrt{t} \log(t + e)$ over $[0, +\infty)$ and the fact that $\sum_{v \in V_j} \mu_{\mathbb{D}}(v) = n$, and the third inequality holds since $|V_j| = \Delta^{h-j} \leq n/2$ for every $j \in [h]$. This concludes the proof. $\blacktriangleleft$

5 Euclidean Metrics

In this section, we turn to the Euclidean setting, where $\mathcal{X} = [0, 1]^d$ and δ is the Euclidean distance. This case is widely studied in online metric matching. We begin by stating our paper's main theorem, which characterizes the competitive ratio achievable under σ -smooth request distributions. Recall that a measure μ over $[0, 1]^d$ is σ -smooth for $\sigma \in (0, 1]$ if $\mu(\mathcal{X}) \leq \mathbb{U}(\mathcal{X})/\sigma$ for every measurable subset $\mathcal{X} \subseteq [0, 1]^d$, where $\mathbb{U}$ is the uniform distribution over $[0, 1]^d$.

► **Theorem 15 (Main Theorem).** *For $[0, 1]^d$ with the Euclidean distance, suppose that S is adversarial and R is drawn from $\mathbb{D} = \prod_{i=1}^n \mathbb{D}_i$, where each $\mathbb{D}_i$ is σ -smooth for $\sigma \in (0, 1]$. Moreover, suppose that the algorithm is given one sample from each $\mathbb{D}_i$. Then, there exists an algorithm $\mathcal{A}_{\text{RS}}$ with competitive ratio*

$$\begin{cases} O(\sigma^{-1}), & d = 1, \\ O(\sigma^{-\frac{1}{2}} n^{\frac{1}{2} - \frac{1}{2 \log(25/12)}}), & d = 2, \\ O(d^{\frac{3}{2}} \sigma^{-\frac{1}{d}}), & d \geq 3, \end{cases}$$

and an algorithm $\mathcal{A}_{\text{BBGN}}$ with competitive ratio

$$\begin{cases} O(\sigma^{-1} \log n), & d = 1, \\ O(\sigma^{-\frac{1}{2}} \log^2 n), & d = 2, \\ O(d^{\frac{3}{2}} \sigma^{-\frac{1}{d}}), & d \geq 3. \end{cases}$$

In particular, both algorithms do not need to know the correspondence between distributions and samples.

In particular, these results show that in one dimension the dependence on n is either logarithmic or absent, in two dimension the competitive ratio grows sublinearly in n , and in higher dimensions the bound is dimension-dependent but independent of n .

The proof of Theorem 15 consists of three ingredients. First, Lemma 5 shows that, given sample access to $\mathbb{D}$, the adversarial-server setting can be reduced to the stochastic case where the servers are also drawn from $\mathbb{D}$. Second, Lemma 16 provides two algorithms for Euclidean metrics along with upper bounds on their expected costs. Finally, Lemma 17 establishes lower bounds on the offline optimum. We present the latter two results next, and then combine all three ingredients to complete the proof.

► **Lemma 16.** *For $[0, 1]^d$ with the Euclidean distance, and for any $\mathbb{D} = \prod_{i=1}^n \mathbb{D}_i$, there exists an algorithm $\mathcal{A}_{\text{RS}}$ such that*

$$\text{cost}(\mathcal{A}_{\text{RS}}; \mathbb{D}, \mathbb{D}) \leq \begin{cases} O(\sqrt{n}), & d = 1, \\ O\left(n^{1 - \frac{1}{2 \log(25/12)}}\right), & d = 2, \\ O(d^{\frac{3}{2}} n^{1 - \frac{1}{d}}), & d \geq 3, \end{cases}$$

and an algorithm $\mathcal{A}_{\text{BBGN}}$ such that

$$\text{cost}(\mathcal{A}_{\text{BBGN}}; \mathbb{D}, \mathbb{D}) \leq \begin{cases} O(\sqrt{n} \log n), & d = 1, \\ O(\sqrt{n} \log^2 n), & d = 2, \\ O(d^{\frac{3}{2}} n^{1 - \frac{1}{d}}), & d \geq 3. \end{cases}$$

We defer the proof of Lemma 16 to Section 5.1 and give lower bounds for the offline optimum in the following lemma.

► **Lemma 17.** *For $[0, 1]^d$ with the Euclidean distance, and for all S and $\mathbb{D} = \prod_{i=1}^n \mathbb{D}_i$ such that each $\mathbb{D}_i$ is σ -smooth for $\sigma \in (0, 1]$,*

$$\text{OPT}(S, \mathbb{D}) \geq \begin{cases} \Omega(\sigma \sqrt{n}), & d = 1, \\ \Omega(\sigma^{\frac{1}{d}} n^{1 - \frac{1}{d}}), & d \geq 2. \end{cases}$$

We defer the proof of Lemma 17 to Section 5.2. Now, we have collected all the necessary ingredients to finish the proof of Theorem 15.

Proof of Theorem 15. By Lemma 5, given an algorithm $\mathcal{A}$, which is provided with a sample from each $\mathbb{D}_i$, there exists an algorithm $\mathcal{A}'$, which does not need to know the correspondence between distributions and samples, such that $\text{cost}(\mathcal{A}'; S, \mathbb{D}) \leq \text{OPT}(S, \mathbb{D}) + \text{cost}(\mathcal{A}; \mathbb{D}, \mathbb{D})$. Hence, it suffices to give an algorithm $\mathcal{A}$ such that $\text{cost}(\mathcal{A}; \mathbb{D}, \mathbb{D})/\text{OPT}(S, \mathbb{D})$ is upper bounded by the desired competitive ratio. The theorem then follows by applying the algorithms $\mathcal{A}_{\text{RS}}$ and $\mathcal{A}_{\text{BBGN}}$ given in Lemma 16, and the lower bounds for $\text{OPT}(S, \mathbb{D})$ given in Lemma 17. ◀


5.1 Proof of Lemma 16

Define a hierarchical decomposition $\mathcal{H}_0, \mathcal{H}_1, \dots, \mathcal{H}_h$ of $[0, 1]^d$, where h will be determined later, such that

$$\mathcal{H}_i := \left\{ \prod_{\ell=1}^d I(i, \lambda_\ell) \mid \lambda_1, \dots, \lambda_d \in [2^{h-i}] \right\},$$

where

$$I(i, \lambda) := \begin{cases} [2^{i-h}(\lambda - 1), 2^{i-h}\lambda), & \lambda < 2^{h-i}, \\ [2^{i-h}(\lambda - 1), 2^{i-h}\lambda], & \lambda = 2^{h-i}. \end{cases}$$



■ **Figure 2** Hierarchical decomposition for $[0, 1]^2$.

In other words, $\mathcal{H}_i$ is the partition of $[0, 1]^d$ into $2^{d(h-i)}$ subcubes with side-length 2^{i-h} . This gives a laminar family⁴ $\mathcal{H} := \mathcal{H}_0 \cup \mathcal{H}_1 \cup \dots \cup \mathcal{H}_h$. See Figure 2 for an illustration with $d = 2$. To prove Lemma 16, we first construct a 2^d -ary 2-HST with height h from $\mathcal{H}$, and then we show that it suffices to upper bound the cost of an algorithm for the resulting HST metric, which enables us to apply our algorithmic results for HST metrics.

We construct a 2^d -ary 2-HST with height h , denoted as $\mathcal{T}$, from $\mathcal{H}$ as follows: each cube in $\mathcal{H}$ corresponds to a node, and the children of a node corresponding to $H \in \mathcal{H}$ are the nodes corresponding to maximal subsets of H in $\mathcal{H}$. For every $x \in [0, 1]^d$, let $\mathcal{T}(x)$ be the leaf node corresponding to the (unique) cube in $\mathcal{H}_0$ that contains x . For $S \subseteq [0, 1]^d$, define $\mathcal{T}(S) := \{\mathcal{T}(s) \mid s \in S\}$. We show in the following lemma that a cost upper bound for an algorithm on $\mathcal{T}$ gives rise to a cost upper bound for the corresponding algorithm on $[0, 1]^d$.

► **Lemma 18.** *Given an algorithm $\mathcal{A}$ on $\mathcal{T}$, there exists an algorithm $\mathcal{A}'$ on $[0, 1]^d$ such that for all S and R , $\text{cost}(\mathcal{A}'; S, R) \leq \sqrt{d} \cdot (\text{cost}(\mathcal{A}; \mathcal{T}(S), \mathcal{T}(R))/4 + n \cdot 2^{-h})$.*

Proof. Given the server set $S = \{s_1, \dots, s_n\}$, the algorithm $\mathcal{A}'$ initializes the given algorithm $\mathcal{A}$ with the server set being $\mathcal{T}(S)$. For each arriving request $r \in [0, 1]^d$, if $\mathcal{A}$ matches request $\mathcal{T}(r)$ to server $\mathcal{T}(s)$ for $s \in S$, then $\mathcal{A}'$ matches r to s . To establish the desired upper bound for $\text{cost}(\mathcal{A}'; S, R)$, it suffices to show that $\|s - r\|_2 \leq \sqrt{d} \cdot (\delta(\mathcal{T}(s), \mathcal{T}(r))/4 + 2^{-h})$ for all $s, r \in [0, 1]^d$, where $\delta(v, v')$ denotes the distance between v and v' on $\mathcal{T}$.

Fix $s, r \in [0, 1]^d$. Let k be the smallest integer in $\{0, 1, \dots, h\}$ such that there exists $H \in \mathcal{H}_k$ that contains both s and r . In other words, H is the (unique) smallest cube in $\mathcal{H}$ that contains both s and r , which implies $\|s - r\|_2 \leq \text{diam}(H) = \sqrt{d} \cdot 2^{k-h}$. Note that the node corresponding to H is the least common ancestor of $\mathcal{T}(s)$ and $\mathcal{T}(r)$, and the height of this node is k . Therefore,

$$\delta(\mathcal{T}(s), \mathcal{T}(r)) = 2 \sum_{i=h-k}^{h-1} 2^{-i} = 2^{k-h+2} - 2^{2-h} \geq \frac{4\|s - r\|_2}{\sqrt{d}} - 2^{2-h},$$

concluding the proof. ◀

⁴ Recall that a *laminar family* $\mathcal{F}$ is a family of subsets such that for any $X, Y \in \mathcal{F}$, one of the three following cases holds: (1) $X \subseteq Y$, (2) $Y \subseteq X$, or (3) $X \cap Y = \emptyset$.

By Lemma 18, it suffices to provide algorithms for any 2^d -ary 2-HST with a certain height h , for which we apply Theorems 7 and 11, where the height h is chosen to minimize the resulting cost for $[0, 1]^d$. We establish upper bounds for the expected cost of the RS and BBGN algorithms for the specified HST in the following two corollaries, whose proofs only involve mechanical calculation and are deferred to the full version of this paper [45].

► **Corollary 19.** *For the 2^d -ary 2-HST with height*

$$h := \begin{cases} \left\lfloor \frac{\log n}{2 \log(25/12)} \right\rfloor, & d = 2, \\ \lfloor \log(n)/d \rfloor, & d \neq 2, \end{cases}$$

and for any $\mathbb{D} = \prod_{i=1}^n \mathbb{D}_i$,

$$\text{cost}(\text{RS}; \mathbb{D}, \mathbb{D}) \leq \begin{cases} O(\sqrt{n}), & d = 1, \\ O\left(n^{1 - \frac{1}{2 \log(25/12)}}\right), & d = 2, \\ O(dn^{1 - \frac{1}{d}}), & d \geq 3. \end{cases}$$

► **Corollary 20.** *For the 2^d -ary 2-HST with height $h := \lfloor \log(n)/d \rfloor$, and for any $\mathbb{D} = \prod_{i=1}^n \mathbb{D}_i$,*

$$\text{cost}(\text{BBGN}; \mathbb{D}, \mathbb{D}) \leq \begin{cases} O(\sqrt{n} \log n), & d = 1, \\ O(\sqrt{n} \log^2 n), & d = 2, \\ O(dn^{1 - \frac{1}{d}}), & d \geq 3. \end{cases}$$

For $\mathbb{D} = \prod_{i=1}^n \mathbb{D}_i$, where each $\mathbb{D}_i$ is supported on $[0, 1]^d$, denote $\mathcal{T}(\mathbb{D})$ as the distribution followed by $\mathcal{T}(S)$, where $S \sim \mathbb{D}$. Denote the algorithm for $[0, 1]^d$ obtained by applying Lemma 18 to the RS algorithm as $\mathcal{A}_{\text{RS}}$, which implies $\text{cost}(\mathcal{A}_{\text{RS}}; \mathbb{D}, \mathbb{D}) \leq \sqrt{d} \cdot (\text{cost}(\text{RS}; \mathcal{T}(\mathbb{D}), \mathcal{T}(\mathbb{D}))/4 + n \cdot 2^{-h})$ for any $\mathbb{D} = \prod_{i=1}^n \mathbb{D}_i$. Next, we upper bound the cost of $\mathcal{A}_{\text{RS}}$ by Corollary 19. For $d = 1$,

$$\text{cost}(\mathcal{A}_{\text{RS}}; \mathbb{D}, \mathbb{D}) \leq O(\text{cost}(\text{RS}; \mathcal{T}(\mathbb{D}), \mathcal{T}(\mathbb{D})) + n \cdot 2^{-\lfloor \log n \rfloor}) \leq O(\sqrt{n}).$$

For $d = 2$,

$$\text{cost}(\mathcal{A}_{\text{RS}}; \mathbb{D}, \mathbb{D}) \leq O\left(\text{cost}(\text{RS}; \mathcal{T}(\mathbb{D}), \mathcal{T}(\mathbb{D})) + n \cdot 2^{-\left\lfloor \frac{\log n}{2 \log(25/12)} \right\rfloor}\right) \leq O\left(n^{1 - \frac{1}{2 \log(25/12)}}\right).$$

For $d \geq 3$,

$$\text{cost}(\mathcal{A}_{\text{RS}}; \mathbb{D}, \mathbb{D}) \leq O(\sqrt{d} \cdot (\text{cost}(\text{RS}; \mathcal{T}(\mathbb{D}), \mathcal{T}(\mathbb{D})) + n \cdot 2^{-\lfloor \log(n)/d \rfloor})) \leq O(d^{\frac{3}{2}} n^{1 - \frac{1}{d}}).$$

Denote the algorithm for $[0, 1]^d$ obtained by applying Lemma 18 to the BBGN algorithm as $\mathcal{A}_{\text{BBGN}}$, which implies $\text{cost}(\mathcal{A}_{\text{BBGN}}; \mathbb{D}, \mathbb{D}) \leq \sqrt{d} \cdot (\text{cost}(\text{BBGN}; \mathcal{T}(\mathbb{D}), \mathcal{T}(\mathbb{D}))/4 + n \cdot 2^{-h})$ for any $\mathbb{D} = \prod_{i=1}^n \mathbb{D}_i$. Next, we upper bound the cost of $\mathcal{A}_{\text{BBGN}}$ by Corollary 20. For $d = 1$,

$$\text{cost}(\mathcal{A}_{\text{BBGN}}; \mathbb{D}, \mathbb{D}) \leq O(\text{cost}(\text{BBGN}; \mathcal{T}(\mathbb{D}), \mathcal{T}(\mathbb{D})) + n \cdot 2^{-\lfloor \log n \rfloor}) \leq O(\sqrt{n} \log n).$$

For $d = 2$,

$$\text{cost}(\mathcal{A}_{\text{BBGN}}; \mathbb{D}, \mathbb{D}) \leq O(\text{cost}(\text{BBGN}; \mathcal{T}(\mathbb{D}), \mathcal{T}(\mathbb{D})) + n \cdot 2^{-\lfloor \log(n)/2 \rfloor}) \leq O(\sqrt{n} \log^2 n).$$

For $d \geq 3$,

$$\begin{aligned} \text{cost}(\mathcal{A}_{\text{BBGN}}; \mathbb{D}, \mathbb{D}) &\leq O(\sqrt{d} \cdot (\text{cost}(\text{BBGN}; \mathcal{T}(\mathbb{D}), \mathcal{T}(\mathbb{D})) + n \cdot 2^{-\lfloor \log(n)/d \rfloor})) \\ &\leq O(d^{\frac{3}{2}} n^{1 - \frac{1}{d}}), \end{aligned}$$

concluding the proof of Lemma 16.

5.2 Proof of Lemma 17

The proof for $d \geq 2$ relies on the following lower bound for the minimum cost of matching each request, which is established by employing nearest-neighbor-distance.

► **Lemma 21** (Lemma 14 in the arXiv version of [58]). *Let $\mathbb{D}_i$ be a σ -smooth distribution over $[0, 1]^d$ with $\sigma \in (0, 1]$. For all $d \geq 1$ and finite $S \subseteq [0, 1]^d$,*

$$\mathbb{E}_{r_i \sim \mathbb{D}_i} \left[\min_{s \in S} \|s - r_i\|_2 \right] \geq \Omega \left(\sigma^{\frac{1}{d}} |S|^{-\frac{1}{d}} \right).$$

To see that Lemma 21 implies the lower bound for $d \geq 2$, note that

$$\text{OPT}(S, \mathbb{D}) \geq \sum_{i=1}^n \mathbb{E}_{r_i \sim \mathbb{D}_i} \left[\min_{s \in S} \|s - r_i\|_2 \right] \geq \Omega \left(\sigma^{\frac{1}{d}} n^{1-\frac{1}{d}} \right),$$

as desired.

The rest of this subsection is devoted to proving the lower bound for $d = 1$. Fix S and $\mathbb{D} = \prod_{i=1}^n \mathbb{D}_i$, and let $R \sim \mathbb{D}$. Fix an arbitrary min-cost perfect matching M between S and R . For each $x \in [0, 1]$, define $\hat{m}(x)$ as the number of matches in M that “cross” x , i.e., one endpoint of the match is in $[0, x]$, and the other endpoint is in $[x, 1]$. Note that

$$\text{OPT}(S, \mathbb{D}) = \mathbb{E} \left[\int_0^1 \hat{m}(x) dx \right].$$

Let $L := \sigma/4$. For each $x \in [0, 1 - L]$, define $\hat{s}(x) := |S \cap [x, x + L]|$ as the number of servers in $[x, x + L]$ and $\hat{r}(x) := |R \cap [x, x + L]|$ as the (random) number of requests in $[x, x + L]$. Note that if $\hat{s}(x) > \hat{r}(x)$, then at least $\hat{s}(x) - \hat{r}(x)$ servers in $[x, x + L]$ have to be matched to requests outside of $[x, x + L]$; similarly, if $\hat{s}(x) < \hat{r}(x)$, then at least $\hat{r}(x) - \hat{s}(x)$ requests in $[x, x + L]$ have to be matched to servers outside of $[x, x + L]$. Hence, for each $x \in [0, 1 - L]$, $\hat{m}(x) + \hat{m}(x + L) \geq |\hat{s}(x) - \hat{r}(x)|$. As a result,

$$\int_0^{1-L} |\hat{s}(x) - \hat{r}(x)| dx \leq \int_0^{1-L} (\hat{m}(x) + \hat{m}(x + L)) dx \leq 2 \int_0^1 \hat{m}(x) dx.$$

Hence, it suffices to show that

$$\mathbb{E} \left[\int_0^{1-L} |\hat{s}(x) - \hat{r}(x)| dx \right] \geq \Omega(\sigma \sqrt{n}). \quad (14)$$

Fix $x \in [0, 1 - L]$, and we analyze $\mathbb{E}[|\hat{s}(x) - \hat{r}(x)|]$. For each $i \in [n]$, we use $\mu_{\mathbb{D}_i}$ to denote the density of $\mathbb{D}_i$ with respect to the uniform distribution over $[0, 1]$. Note that $\hat{r}(x) \sim \text{PB}(\mathbf{w}(x))$, where $w_i(x) := \int_x^{x+L} \mu_{\mathbb{D}_i}(y) dy$ for each $i \in [n]$. Define $W(x) := \|\mathbf{w}(x)\|_1$, which implies $\mathbb{E}[\hat{r}(x)] = W(x)$. The following lemma lower bounds $\mathbb{E}[|\hat{s}(x) - \hat{r}(x)|]$.

► **Lemma 22.** $\mathbb{E}[|\hat{s}(x) - \hat{r}(x)|] \geq \Omega(\sqrt{W(x)}) - O(1)$.

Proof. By Jensen’s inequality, $\mathbb{E}[|\hat{s}(x) - \hat{r}(x)|] \geq \mathbb{E}[|\hat{r}(x) - W(x)|]$, and it suffices to show that

$$\mathbb{E}[|\hat{r}(x) - W(x)|] = \mathbb{E}[|\text{PB}(\mathbf{w}(x)) - W(x)|] \geq \Omega(\sqrt{W(x)}) - O(1). \quad (15)$$

For each $i \in [n]$, by the σ -smoothness of $\mathbb{D}_i$,

$$w_i(x) = \int_x^{x+L} \mu_{\mathbb{D}_i}(y) dy \leq \frac{L}{\sigma} \leq \frac{1}{2}.$$

Let $\mathbf{w}' \in \mathbb{R}_{\geq 0}^n$ satisfy

$$w'_i = \begin{cases} 1/2, & i \leq \lfloor 2W(x) \rfloor, \\ W(x) - \lfloor 2W(x) \rfloor / 2, & i = \lfloor 2W(x) \rfloor + 1, \\ 0, & \text{otherwise,} \end{cases}$$

which gives $\|\mathbf{w}'\|_1 = W(x)$ and $\|\mathbf{w}'\|_\infty \leq 1/2$. Since $\mathbf{w}' \succ \mathbf{w}(x)$, by Corollary 4,

$$\mathbb{E}[|\text{PB}(\mathbf{w}(x)) - W(x)|] \geq \mathbb{E}[|\text{PB}(\mathbf{w}') - W(x)|].$$

This would imply (15) since

$$\begin{aligned} \mathbb{E}[|\text{PB}(\mathbf{w}') - W(x)|] &\geq \mathbb{E} \left[\left| \text{Bin} \left(\lfloor 2W(x) \rfloor, \frac{1}{2} \right) - \frac{\lfloor 2W(x) \rfloor}{2} \right| \right] - 1 \\ &\geq \frac{\sqrt{\lfloor 2W(x) \rfloor}}{2\sqrt{2}} - 1 \geq \Omega(\sqrt{W(x)}) - O(1), \end{aligned}$$

where the second inequality holds by the following probabilistic bound.

▷ **Claim 23 ([15]).** Let $Z \sim \text{Bin}(n, p)$, with $n \geq 2$ and $p \in [1/n, 1 - 1/n]$. Then, we have

$$\mathbb{E}[|Z - \mathbb{E}Z|] \geq \text{std}(Z)/\sqrt{2}. \quad \blacktriangleleft$$

By Lemma 22,

$$\mathbb{E} \left[\int_0^{1-L} |\hat{s}(x) - \hat{r}(x)| dx \right] = \int_0^{1-L} \mathbb{E}[|\hat{s}(x) - \hat{r}(x)|] dx \geq \Omega \left(\int_0^{1-L} \sqrt{W(x)} dx \right) - O(1),$$

and (14) follows from the following lemma.

► **Lemma 24.** *It holds that*

$$\int_0^{1-L} \sqrt{W(x)} dx \geq \Omega(\sigma\sqrt{n}).$$

Proof. By the definition of $W(x)$,

$$\begin{aligned} \int_0^{1-L} W(x) dx &= \int_0^{1-L} \sum_{i=1}^n \int_x^{x+L} \mu_{\mathbb{D}_i}(y) dy dx \\ &= \sum_{i=1}^n \int_0^1 \mu_{\mathbb{D}_i}(y) \int_0^{1-L} \mathbf{1}\{y \in [x, x+L]\} dx dy \\ &= \sum_{i=1}^n \int_0^1 \mu_{\mathbb{D}_i}(y) \cdot \min\{L, y, 1-y\} dy \\ &\geq L \sum_{i=1}^n \int_L^{1-L} \mu_{\mathbb{D}_i}(y) dy. \end{aligned}$$

For each $i \in [n]$, by the σ -smoothness of $\mathbb{D}_i$,

$$\int_L^{1-L} \mu_{\mathbb{D}_i}(y) dy = 1 - \int_0^L \mu_{\mathbb{D}_i}(y) dy - \int_{1-L}^1 \mu_{\mathbb{D}_i}(y) dy \geq 1 - \frac{2L}{\sigma} = \frac{1}{2}.$$

Combining the above two displayed equations,

$$\int_0^{1-L} W(x) dx \geq \frac{Ln}{2} = \Omega(\sigma n). \quad (16)$$

For every $x \in [0, 1 - L]$, since $W(x) \in [0, n]$, we have $\sqrt{W(x)} \geq W(x)/\sqrt{n}$. It follows that

$$\int_0^{1-L} \sqrt{W(x)} dx \geq \frac{1}{\sqrt{n}} \int_0^{1-L} W(x) dx \geq \Omega(\sigma \sqrt{n}),$$

where the second inequality holds by (16), concluding the proof. $\blacktriangleleft$

6 Discussion and Future Directions

In this paper, we study the online metric matching problem for the Euclidean space $[0, 1]^d$ when servers are adversarial and requests are independently drawn from distinct smooth distributions. We present an $O(1)$ -competitive algorithm for $[0, 1]^d$ with $d \neq 2$, given a single sample from each request distribution. A key feature of our approach is that, by directly upper-bounding the algorithm's cost after a simple deterministic metric embedding, we bypass the $\Omega(\log n)$ competitive-ratio barrier that arises in the adversarial setting due to metric distortion. Since metric embeddings into HSTs have already been proven extremely effective for related online problems such as k -server [11, 18], k -taxi [33], and several variants of online metric matching [28, 16], a natural and exciting future direction is to extend our techniques to these problems.

Our guarantees rely on requests being independently sampled. An intriguing direction would be to see what forms of correlation among requests might still permit an $o(\log n)$ competitive ratio. As a starting point, recent breakthroughs in smoothed analysis of online learning [17, 37] allow each arrival's distribution – while required to be smooth – to depend on the realized history of arrivals and algorithmic decisions, and their techniques may extend to our setting. In addition, the correlation models studied in online stochastic matching [4] and prophet inequalities [39] present additional promising avenues.

References

- 1 Mohammad Akbarpour, Yeganeh Alimohammadi, Shengwu Li, and Amin Saberi. The value of excess supply in spatial matching markets. In *EC*, page 62. ACM, 2022. doi:10.1145/3490486.3538375.
- 2 Alireza AmaniHamedani, Ali Aouad, and Amin Saberi. Adaptive approximation schemes for matching queues. In *STOC*, pages 1454–1464. ACM, 2025. doi:10.1145/3717823.3718317.
- 3 Michael Anastos, Matthew Kwan, and Benjamin Moore. Smoothed analysis for graph isomorphism. In *STOC*, pages 2098–2106. ACM, 2025. doi:10.1145/3717823.3718173.
- 4 Ali Aouad and Will Ma. A nonparametric framework for online stochastic matching with correlated arrivals. In *EC*, page 114. ACM, 2023. doi:10.1145/3580507.3597773.
- 5 Ali Aouad and Ömer Saritaç. Dynamic stochastic matching under limited time. *Oper. Res.*, 70(4):2349–2383, 2022. doi:10.1287/OPRE.2022.2293.
- 6 Stephen Arndt, Benjamin Moseley, Kirk Pruhs, and Marc Uetz. Competitive online transportation simplified. *arXiv preprint arXiv:2508.08381*, 2025. doi:10.48550/arXiv.2508.08381.
- 7 Itai Ashlagi, Yossi Azar, Moses Charikar, Ashish Chiplunkar, Ofir Geri, Haim Kaplan, Rahul Makhijani, Yuyi Wang, and Roger Wattenhofer. Min-cost bipartite perfect matching with delays. In *APPROX-RANDOM*, volume 81 of *LIPICs*, pages 1:1–1:20. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2017. doi:10.4230/LIPICs.APPROX-RANDOM.2017.1.

- 8 Pablo Daniel Azar, Robert Kleinberg, and S. Matthew Weinberg. Prophet inequalities with limited information. In *SODA*, pages 1358–1377. SIAM, 2014. doi:10.1137/1.9781611973402.100.
- 9 Eric Balkanski, Yuri Faenza, and Noémie Périer. The power of greedy for online minimum cost matching on the line. In *EC*, pages 185–205. ACM, 2023. doi:10.1145/3580507.3597794.
- 10 Nikhil Bansal, Niv Buchbinder, Anupam Gupta, and Joseph Naor. A randomized $o(\log^2 k)$ -competitive algorithm for metric bipartite matching. *Algorithmica*, 68(2):390–403, 2014. doi:10.1007/S00453-012-9676-9.
- 11 Nikhil Bansal, Niv Buchbinder, Aleksander Madry, and Joseph Naor. A polylogarithmic-competitive algorithm for the k -server problem. *J. ACM*, 62(5):40:1–40:49, 2015. doi:10.1145/2783434.
- 12 Nikhil Bansal, Haotian Jiang, Raghu Meka, Sahil Singla, and Makrand Sinha. Prefix discrepancy, smoothed analysis, and combinatorial vector balancing. In *ITCS*, volume 215 of *LIPIcs*, pages 13:1–13:22. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPIcs.ITCS.2022.13.
- 13 Nikhil Bansal, Haotian Jiang, Raghu Meka, Sahil Singla, and Makrand Sinha. Smoothed analysis of the komlós conjecture. In *ICALP*, volume 229 of *LIPIcs*, pages 14:1–14:12. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPIcs.ICALP.2022.14.
- 14 Luca Becchetti, Stefano Leonardi, Alberto Marchetti-Spaccamela, Guido Schäfer, and Tjark Vredeveld. Average-case and smoothed competitive analysis of the multilevel feedback algorithm. *Math. Oper. Res.*, 31(1):85–108, 2006. doi:10.1287/MOOR.1050.0170.
- 15 Daniel Berend and Aryeh Kontorovich. A sharp estimate of the binomial mean absolute deviation with applications. *Statistics & Probability Letters*, 83(4):1254–1259, 2013.
- 16 Sujoy Bhore, Arnold Filtser, and Csaba D. Tóth. Online duet between metric embeddings and minimum-weight perfect matchings. In *SODA*, pages 4564–4579. SIAM, 2024. doi:10.1137/1.9781611977912.162.
- 17 Adam Block, Yuval Dagan, Noah Golowich, and Alexander Rakhlin. Smoothed online learning is as easy as statistical learning. In *COLT*, volume 178 of *Proceedings of Machine Learning Research*, pages 1716–1786. PMLR, 2022. URL: <https://proceedings.mlr.press/v178/block22a.html>.
- 18 Sébastien Bubeck, Michael B. Cohen, Yin Tat Lee, James R. Lee, and Aleksander Madry. k -server via multiscale entropic regularization. In *STOC*, pages 3–16. ACM, 2018. doi:10.1145/3188745.3188798.
- 19 Constantine Caramanis, Paul Dütting, Matthew Faw, Federico Fusco, Philip Lazos, Stefano Leonardi, Orestis Papadigenopoulos, Emmanouil Pountourakis, and Rebecca Reiffenhäuser. Single-sample prophet inequalities via greedy-ordered selection. In *SODA*, pages 1298–1325. SIAM, 2022. doi:10.1137/1.9781611977073.54.
- 20 Xi Chen, Chenghao Guo, Emmanouil V. Vlatakis-Gkaragkounis, and Mihalis Yannakakis. Smoothed complexity of SWAP in local graph partitioning. In *SODA*, pages 5057–5083. SIAM, 2024. doi:10.1137/1.9781611977912.182.
- 21 Yilun Chen, Yash Kanoria, Akshit Kumar, and Wenxin Zhang. Feature based dynamic matching. In *EC*, page 451. ACM, 2023. doi:10.1145/3580507.3597797.
- 22 Christian Coester and Jack Umenberger. Smoothed analysis of online metric problems. In *ESA*, volume 351 of *LIPIcs*, pages 115:1–115:14. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2025. doi:10.4230/LIPIcs.ESA.2025.115.
- 23 José Correa, Andrés Cristi, Boris Epstein, and José A. Soto. Sample-driven optimal stopping: From the secretary problem to the i.i.d. prophet inequality. *Math. Oper. Res.*, 49(1):441–475, 2024. doi:10.1287/MOOR.2023.1363.
- 24 José Correa, Paul Dütting, Felix A. Fischer, and Kevin Schewior. Prophet inequalities for I.I.D. random variables from an unknown distribution. In *EC*, pages 3–17. ACM, 2019. doi:10.1145/3328526.3329627.

- 25 Andrés Cristi and Bruno Ziliotto. Prophet inequalities require only a constant number of samples. In *STOC*, pages 491–502. ACM, 2024. doi:10.1145/3618260.3649773.
- 26 Naveen Durvasula, Nika Haghtalab, and Manolis Zampetakis. Smoothed analysis of online non-parametric auctions. In *EC*, pages 540–560. ACM, 2023. doi:10.1145/3580507.3597787.
- 27 Paul Dütting, Thomas Kesselheim, Brendan Lucier, Rebecca Reiffenhäuser, and Sahil Singla. Online combinatorial allocations and auctions with few samples. In *Focs*, pages 1231–1250. IEEE, 2024. doi:10.1109/Focs61266.2024.00081.
- 28 Yuval Emek, Shay Kuten, and Roger Wattenhofer. Online matching: haste makes waste! In *STOC*, pages 333–344. ACM, 2016. doi:10.1145/2897518.2897557.
- 29 Jittat Fakcharoenphol, Satish Rao, and Kunal Talwar. A tight bound on approximating arbitrary metrics by tree metrics. *J. Comput. Syst. Sci.*, 69(3):485–497, 2004. doi:10.1016/J.JCSS.2004.04.011.
- 30 Hu Fu, Pinyan Lu, Zhihao Gavin Tang, Hongxun Wu, Jinzhao Wu, and Qianfan Zhang. Sample-based matroid prophet inequalities. In *EC*, page 781. ACM, 2024. doi:10.1145/3670865.3673506.
- 31 Rohan Ghuge, Sahil Singla, and Yifan Wang. Single-sample and robust online resource allocation. In *STOC*, pages 1442–1453. ACM, 2025. doi:10.1145/3717823.3718246.
- 32 Anupam Gupta, Guru Guruganesh, Binghui Peng, and David Wajc. Stochastic online metric matching. In *ICALP*, volume 132 of *LIPIcs*, pages 67:1–67:14. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPIcs.ICALP.2019.67.
- 33 Anupam Gupta, Amit Kumar, and Debmalya Panigrahi. Poly-logarithmic competitiveness for the k -taxi problem. In *SODA*, pages 4220–4246. SIAM, 2024. doi:10.1137/1.9781611977912.146.
- 34 Anupam Gupta and Kevin Lewi. The online metric matching problem for doubling metrics. In *ICALP (1)*, volume 7391 of *Lecture Notes in Computer Science*, pages 424–435. Springer, 2012. doi:10.1007/978-3-642-31594-7_36.
- 35 Varun Gupta, Ravishankar Krishnaswamy, and Sai Sandeep. Permutation strikes back: The power of recourse in online metric matching. In *APPROX-RANDOM*, volume 176 of *LIPIcs*, pages 40:1–40:20. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPIcs.APPROX/RANDOM.2020.40.
- 36 Nika Haghtalab, Yanjun Han, Abhishek Shetty, and Kunhe Yang. Oracle-efficient online learning for smoothed adversaries. In *NeurIPS*, 2022.
- 37 Nika Haghtalab, Tim Roughgarden, and Abhishek Shetty. Smoothed analysis with adaptive adversaries. *J. ACM*, 71(3):19, 2024. doi:10.1145/3656638.
- 38 Tsubasa Harada and Toshiya Itoh. A nearly optimal deterministic algorithm for online transportation problem. In *ICALP*, volume 334 of *LIPIcs*, pages 94:1–94:17. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2025. doi:10.4230/LIPIcs.ICALP.2025.94.
- 39 Nicole Immorlica, Sahil Singla, and Bo Waggoner. Prophet inequalities with linear correlations and augmentations. *ACM Trans. Economics and Comput.*, 11(3-4):1–29, 2023. doi:10.1145/3623273.
- 40 Bala Kalyanasundaram and Kirk Pruhs. Online weighted matching. *J. Algorithms*, 14(3):478–488, 1993. doi:10.1006/JAGM.1993.1026.
- 41 Sampath Kannan, Jamie Morgenstern, Aaron Roth, Bo Waggoner, and Zhiwei Steven Wu. A smoothed analysis of the greedy algorithm for the linear contextual bandit problem. In *NeurIPS*, pages 2231–2241, 2018. URL: <https://proceedings.neurips.cc/paper/2018/hash/2cfd4560539f887a5e420412b370b361-Abstract.html>.
- 42 Yash Kanoria. Dynamic spatial matching. In *EC*, pages 63–64. ACM, 2022. doi:10.1145/3490486.3538278.
- 43 Haim Kaplan, David Naori, and Danny Raz. Online weighted matching with a sample. In *SODA*, pages 1247–1272. SIAM, 2022. doi:10.1137/1.9781611977073.52.

- 44 Samir Khuller, Stephen G. Mitchell, and Vijay V. Vazirani. On-line algorithms for weighted bipartite matching and stable marriages. *Theor. Comput. Sci.*, 127(2):255–267, 1994. doi:10.1016/0304-3975(94)90042-6.
- 45 Yingxi Li, Ellen Vitercik, and Mingwei Yang. Smoothed analysis of online metric matching with a single sample: Beyond metric distortion, 2025. doi:10.48550/arXiv.2510.20288.
- 46 Albert W. Marshall, Ingram Olkin, and Barry C. Arnold. *Inequalities: Theory of Majorization and its Applications*, volume 143. Springer, second edition, 2011. doi:10.1007/978-0-387-68276-1.
- 47 Nicole Megow and Lukas Nölke. Online metric matching on the line with recourse. *Algorithmica*, 87(6):813–841, 2025. doi:10.1007/S00453-025-01299-8.
- 48 Adam Meyerson, Akash Nanavati, and Laura J. Poplawski. Randomized online algorithms for minimum metric bipartite matching. In *SODA*, pages 954–959. ACM Press, 2006. URL: <http://dl.acm.org/citation.cfm?id=1109557.1109662>.
- 49 Krati Nayyar and Sharath Raghvendra. An input sensitive online algorithm for the metric bipartite matching problem. In *FOCS*, pages 505–515. IEEE Computer Society, 2017. doi:10.1109/FOCS.2017.53.
- 50 Enoch Peserico and Michele Scquizzato. Matching on the line admits no $o(\sqrt{\log n})$ -competitive algorithm. *ACM Trans. Algorithms*, 19(3):28:1–28:4, 2023. doi:10.1145/3594873.
- 51 Sharath Raghvendra. A robust and optimal online algorithm for minimum metric bipartite matching. In *APPROX-RANDOM*, volume 60 of *LIPIcs*, pages 18:1–18:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2016. doi:10.4230/LIPIcs.APPROX-RANDOM.2016.18.
- 52 Sharath Raghvendra. Optimal analysis of an online algorithm for the bipartite matching problem on a line. In *SoCG*, volume 99 of *LIPIcs*, pages 67:1–67:14. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2018. doi:10.4230/LIPIcs.SOCG.2018.67.
- 53 Aviad Rubinstein, Jack Z. Wang, and S. Matthew Weinberg. Optimal single-choice prophet inequalities from samples. In *ITCS*, volume 151 of *LIPIcs*, pages 60:1–60:10. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPIcs.ITCS.2020.60.
- 54 Guido Schäfer and Naveen Sivadasan. Topology matters: Smoothed competitiveness of metrical task systems. *Theor. Comput. Sci.*, 341(1-3):216–246, 2005. doi:10.1016/J.TCS.2005.04.006.
- 55 Flore Sentenac, Nathan Noiry, Matthieu Lerasle, Laurent Ménard, and Vianney Perchet. Online matching in geometric random graphs. *CoRR*, abs/2306.07891, 2023. doi:10.48550/arXiv.2306.07891.
- 56 Daniel A. Spielman and Shang-Hua Teng. Smoothed analysis: an attempt to explain the behavior of algorithms in practice. *Commun. ACM*, 52(10):76–84, 2009. doi:10.1145/1562764.1562785.
- 57 Michel Talagrand. *Upper and lower bounds for stochastic processes: decomposition theorems*, volume 60. Springer Nature, 2022.
- 58 Mingwei Yang and Sophie H Yu. Online metric matching: Beyond the worst case. *Operations Research*, 2025.