Improved Bounds for Online Facility Location with Predictions

Dimitris Fotakis[1,2], **Evangelia Gergatsouli**[3,4], **Themistoklis Gouleakis**[5],
Nikolas Patris[6,2], **Thanos Tolias**[1,2]
[1]National Technical University of Athens, Greece
[2]Archimedes/Athena RC, Greece
[3]University of Wisconsin - Madison, WI, USA
[4]Georgia Institute of Technology, Atlanta, GA, USA
[5]Nanyang Technological University, Singapore
[6]University of California, Irvine, CA, USA
fotakis@cs.ntua.gr, egergatsouli3@gatech.edu, themis.gouleakis@ntu.edu.sg, npatris@uci.edu, thanostolias@mail.ntua.gr

Abstract

We consider Online Facility Location (OFL) in the framework of learning-augmented online algorithms. In OFL, demands arrive one-by-one in a metric space and must be (irrevocably) assigned to an open facility upon arrival, without any knowledge about future demands. We focus on uniform facility opening costs and present an online algorithm for OFL that exploits potentially imperfect predictions on the locations of the optimal facilities. We prove that the competitive ratio decreases from sublogarithmic in the number of demands n to constant as the total prediction error η_1, i.e., the sum of distances of the predicted locations to the optimal facility locations, decreases. E.g., our analysis implies that if for some $\varepsilon > 0$, $\eta_1 = \mathrm{OPT}/n^{\varepsilon}$, where OPT is the optimal cost, the competitive ratio becomes $O(1/\varepsilon)$. We complement our analysis with a matching lower bound establishing that the dependence of the algorithm's competitive ratio on the η_1 error is optimal, up to constant factors.

1 Introduction

Online algorithms deal with decision making in cases where the input data is not known in advance, but rather arrives sequentially. The algorithm makes irrevocable decisions, only based on the input data received so far, and incurs the corresponding irrevocable cost. Traditionally, in the analysis of online algorithms we assume, rather pessimistically, that an adversary always presents the algorithm with the worst-case input. The performance of online algorithms is evaluated by the *competitive ratio* (Borodin and El-Yaniv 1998), which is the worst-case ratio of the total algorithm's cost to the cost of a computationally unrestricted optimal algorithm aware of the request sequence in advance.

On the other hand, machine learning (ML) aims to *predict* the unknown based on historical data and to *learn* how the world looks like. A recent trend aims to use ML predictions about the future input in order to deal with the inherent uncertainty in online algorithms, while still providing worst-case performance guarantees. One might think that directly using machine learning in online problems should enhance their performance, since by predicting the input, with some error, we should be able to come up with almost optimal solutions. In reality, this turns out not to be true, since the error of the learner may vary across different parts of the input and could propagate along different phases of the algorithm.

Lykouris and Vassilvitskii (2018) proposed a framework aiming to provide formal guarantees for such *learning-augmented* online algorithms, in terms of their *consistency* and *robustness*. They require the algorithm to be near optimal, if the predictions about the future input are accurate (consistency), while for arbitrary erroneous predictions, the competitive ratio should gracefully degrade to (and not exceed by far) the worst-case one (robustness). The idea of combining online algorithms with ML advice is that in the end, we should be able to overcome the worst-case lower bounds and get the best of both worlds. In the framework of Lykouris and Vassilvitskii (2018); Purohit, Svitkina, and Kumar (2018); Antoniadis et al. (2023), the learning-augmented algorithm, given some predictions of total error η, is required to make decisions online and its competitive ratio is given as a function of η. Many online problems have been studied under this framework, such as ski rental, scheduling, metrical task systems (see e.g., the survey of Mitzenmacher and Vassilvitskii (2020)). In this work we investigate the competitive ratio of Online Facility Location in the framework of learning-augmented online algorithms, following an approach similar to (Jiang et al. 2022).

Online Facility Location (OFL) was introduced by Meyerson (2001) and considers a sequence of demands located in an underlying metric space. Each demand must be connected to an open facility upon arrival. Each facility has an opening cost, which is irrevocable in the sense that once opened, facilities cannot be closed. Every demand incurs its assignment cost, which is the distance to the closest open facility at the demand's assignment time. Our goal is to decide when and where to open the facilities, so that we minimize the total facility opening cost plus demand assignment cost. Meyerson (2001) presented an elegant randomized algorithm with competitive ratio $O(\log n/\log\log n)$, where n is the number of demands. Since then, there has been a significant volume of work on OFL and its many variants (see Section 1.2).

Online Facility Location with Predictions. In OFL with predictions (OFLpred), every demand v in the request sequence is accompanied by a prediction p on the location of the optimal facility where v is assigned. Every demand must be connected to an open facility upon arrival, and facility and assignment costs are irrevocable and are defined as before. In addition to the number n of demands, instances of OFLpred are parameterized by the *total prediction error* η_1, which is the sum, over all predictions, of their distance to the respective optimal facility, and the *maximum prediction error* $\eta_\infty \leq \eta_1$, which is the maximum such distance.

Jiang et al. (2022) studied the competitive ratio of OFLpred as a function of η_∞. They considered non-uniform facility costs, carefully adapted Meyerson's algorithm and proved a competitive ratio of $O(\log \min\{\frac{n\eta_\infty}{\mathrm{OPT}}, n\})$, where OPT is the optimal cost. They also gave an almost matching lower bound of $\Omega\left(\frac{\log \min\{\frac{n\eta_\infty}{\mathrm{OPT}}, n\}}{\log\log n}\right)$ on the competitive ratio of any randomized algorithm for OFLpred (even with uniform facility costs). The results of Jiang et al. (2022) imply that the competitive ratio of OFLpred degrades from a small constant, if all predictions are perfect (consistency), to logarithmic, if there are some inaccurate predictions and $\eta_\infty = \Omega(\mathrm{OPT})$ (robustness). Moreover, for the class of instances in their lower bound, where a small fraction of the predictions have error η_∞ and the remaining ones are colocated with the respective demand points (and thus, $\eta_1 = \Omega(\mathrm{OPT})$ as long as $\eta_\infty \geq \mathrm{OPT}/n^{1-\delta}$, for any constant $\delta > 0$), an almost logarithmic competitive ratio is unavoidable.

1.1 Motivation and Contribution

Our work is motivated from the observation that determining the dependence of OFLpred's competitive ratio on the total prediction error η_1 (in addition to its dependence on η_∞ studied by Jiang et al. (2022)) contributes to a deeper understanding about how and to which extent predictions can help in improving the performance of OFL algorithms. E.g., let us consider OFLpred instances with $\eta_\infty = \mathrm{OPT}/n^\varepsilon$, for some constant $\varepsilon > 0$. Within this class, the upper and the lower bound of Jiang et al. (2022) fail to differentiate, as far as their best possible competitive ratio is concerned, between (i) instances where all but few predictions are perfect; (ii) instances where for some $\beta \in (0, 1)$, every prediction is $1/\beta$ times closer to the respective optimal facility than the corresponding demand; and (iii) instances where every prediction is at distance η_∞ to the optimal facility. In (i), $\eta_1 \approx \mathrm{OPT}/n^\varepsilon$ and we should expect a constant competitive ratio, if ε is constant. In (ii), $\eta_1 \approx \beta\mathrm{OPT}$ and the competitive ratio must be an increasing function of β. Only in (iii), we should expect a competitive ratio fully determined by η_∞. For all these instances with very different prediction quality (and for many other instances in between), the results of Jiang et al. (2022) guarantee a competitive ratio of $O((1-\varepsilon)\log n)$, while their lower bound applies only to (iii). For another example, if $\eta_\infty = \mathrm{OPT}/(\log n)^\ell$, for some integer $\ell \geq 1$, the resulting guarantee on the competitive ratio is $O(\log n - \ell \log\log n)$ in all cases (i)-(iii).

In this work, we focus on OFLpred with *uniform* facility costs, where the cost of opening a facility at any point of the underlying metric space is f, and determine OFLpred's competitive ratio as a function of both η_∞ and η_1. As demonstrated by the discussion above, we believe that the total error η_1 is a metric at least as representative of the predictions' accuracy[1] as the maximum error η_∞.

In Section 3, we present a randomized algorithm that works similarly to Meyerson's algorithm, but with the predictions in place of the demands. By generalizing the analysis of Meyerson (2001), so that it also takes the prediction error of each demand-prediction pair into account, we show that if all predictions are perfect, our algorithm is 2-competitive (i.e., it is 2-consistent), and that for any prediction error η_1 and η_∞ with $\mathrm{OPT} \geq \eta_1 \geq \eta_\infty > \mathrm{OPT}/n$, the competitive ratio of our algorithm is $O\left(\frac{\log(\frac{n\eta_\infty}{\mathrm{OPT}})}{\log\left(\frac{\mathrm{OPT}}{\eta_1}\log(\frac{n\eta_\infty}{\mathrm{OPT}})\right)}\right)$.

In Section 4, we show how to apply cost doubling to OFLpred in order to obtain an online algorithm with thrice the minimum competitive ratio of our proposed algorithm and the algorithm of Jiang et al. (2022). For uniform facility costs, the latter works as Meyerson's algorithm, but it opens facilities in pairs, one at the demand's location and another at the predicted optimal location. Our generalized analysis framework shows an improved competitive ratio for (Jiang et al. 2022)'s algorithm with uniform facility costs. Taking the minimum of the competitive ratios, we get the following:

Theorem 1 (Upper Bound). *The competitive ratio of OFLpred is at most:*

- $3\left(3\frac{n\eta_\infty}{\mathrm{OPT}} + 2\right)$, *for all* $\eta_1 \geq \eta_\infty \geq 0$.
- $O\left(\frac{\log \min\left\{\frac{n\eta_\infty}{\mathrm{OPT}}, n\right\}}{\log\left(\max\left\{\frac{\mathrm{OPT}}{\eta_1}, 1\right\} \log \min\left\{\frac{n\eta_\infty}{\mathrm{OPT}}, n\right\}\right)}\right)$, *for all* $\eta_1 \geq \eta_\infty > \frac{\mathrm{OPT}}{n}$.

We note that η_1 and η_∞ are only used to bound the competitive ratio, since the algorithm itself does not depend on their knowledge. Moreover, our algorithm is robust in the sense that for all $\eta_1 \geq \mathrm{OPT}$ and $\eta_\infty \geq \mathrm{OPT}$, the competitive ratio is $O(\log n/\log\log n)$. It is interesting that the competitive ratio in Theorem 1 (which is best possible, as shown by Theorem 2) is achieved by combining two versions of Meyerson's algorithm: ours, which uses the predicted locations for its facilities and decisions, and (Jiang et al. 2022)'s, whose decisions are guided by the demand locations and uses the predicted locations in a supporting role (see also (Wei 2020) for a similar construction in the context of online caching with predictions).

In Section 5, we prove a lower bound on the competitive ratio of OFLpred, establishing that the dependence of the competitive ratio in Theorem 1 on η_1 and η_∞ is essentially best possible. The lower bound construction general-

[1]We should highlight that a parameterization of the competitive ratio only on η_1 is insufficient for similar reasons that a parameterization on η_∞ alone is not enough (see also footnote 2 in Section 3.3 for the technical reason behind that). To see this, we let $\eta_1 = \mathrm{OPT}$ and consider two classes of instances, one where each prediction has error OPT/n (so $\eta_\infty = \mathrm{OPT}/n$), and another where $m = \Theta(\frac{\log n}{\log\log n})$ and for each $\ell = 1, \ldots, m$, there are $m^{\ell-1}$ predictions with error OPT/m^ℓ each (so $\eta_\infty = \mathrm{OPT}/m$) presented to the algorithm in decreasing order of error. Even though $\eta_1 = \mathrm{OPT}$ in both cases, our results imply that the competitive ratio is constant in the former case and $\Omega(\frac{\log n}{\log\log n})$ in the latter case.

izes the lower bounds of (Fotakis 2008) and (Jiang et al. 2022) and requires non-trivial modifications of the sequence of demand-prediction pairs so that we can achieve virtually any allowable combination of the maximum η_∞ and the total η_1 error (see also Remark 1, Section 5).

Theorem 2 (Lower Bound – Informal)**.** *For all* $\alpha = \eta_\infty/\mathrm{OPT} < 1/3$ *and* $\beta \approx \eta_1/\mathrm{OPT} \in (3\alpha, 1]$*, there are OFL instances with* n *demand-prediction pairs where any randomized algorithm has competitive ratio*

$$\Omega\left(\frac{\log(\alpha n)}{\log(\log(\alpha n)/\beta)}\right)$$

Although our techniques are quite different from those of Jiang et al. (2022), from a conceptual viewpoint, theorems 1 and 2 significantly generalize and can be regarded as an informative refinement of their results. Specifically, for any fixed $\alpha = \eta_\infty/\mathrm{OPT}$, theorems 1 and 2 determine the best possible competitive ratio of OFLpred as a function of $\beta = \eta_1/\mathrm{OPT}$. Returning to our motivating example at the beginning of this section, we can now differentiate between cases (i)-(iii). Theorems 1 and 2 imply that the best possible competitive ratio is $\Theta(1/\varepsilon)$ in (i), $\Theta(\frac{(1-\varepsilon)\log n}{\log((1-\varepsilon)\log(n)/\beta)})$ in (ii), and $\Theta(\frac{(1-\varepsilon)\log n}{\log((1-\varepsilon)\log n)})$ in (iii). Similarly, if for some integer $\ell \geq 1$, $\eta_1 \approx \eta_\infty = \mathrm{OPT}/(\log n)^\ell$, we prove that the best possible competitive ratio for case (i) is $\Theta(\frac{\log n}{(\ell+1)\log\log n})$.

The proofs and the technical details omitted from this extended abstract can be found at (Fotakis et al. 2024).

1.2 Related Work

Learning Augmented Algorithms. Medina and Vassilvitskii (2017) and Lykouris and Vassilvitskii (2018) initiated this line of work and introduced the notions of consistency and robustness. Purohit, Svitkina, and Kumar (2018) considered ski rental and non-clairvoyant scheduling, giving consistency and robustness guarantees. Lykouris and Vassilvitskii (2018) studied the classical online caching problem, and were able to adapt the Marker algorithm (Fiat et al. 1991) to obtain a tradeoff between robustness and consistency. Rohatgi (2020) and Wei (2020) subsequently presented simpler learning-augmented caching algorithms with improved dependence of their competitive ratios on the prediction error.

Further results in online algorithms with machine learned advice include the work of Lattanzi et al. (2020), who studied the restricted assignment scheduling problem, the work of Bamas et al. (2020), who considered energy minimization problems, and the more general framework of online primal-dual algorithms in (Bamas, Maggiori, and Svensson 2020). Moreover, Antoniadis et al. (2023) studied the following online selection problems from the viewpoint of learning-augmented algorithms: (i) the classical secretary problem; (ii) online bipartite matching with vertex arrivals; and (iii) the graphic matroid secretary problem.

More recently, Almanza et al. (2021) considered OFL with predictions in the form of different sets of suggested optimal facility locations. They present a randomized online algorithm with logarithmic competitiveness against an optimal solution restricted to facilities from the suggested sets. Azar, Panigrahi, and Touitou (2022) considered OFL and other online network design problems with the predicted demand sequence given in advance. They used cost doubling in order to combine an online algorithm applied to the actual demand sequence and an offline algorithm that computes partial solutions for the predicted input. However, their notions of prediction and error are different from ours, which makes their results and techniques incomparable to ours. Argue et al. (2022) considered OFL in a setting where the predictions are obtained by sampling an ε-fraction of the demand sequence, and Gupta et al. (2022) considered online covering and facility location problems with predictions guaranteed to be accurate with probability at least ε.

In a different research direction, there has been significant interest recently in facility location problems with predictions from the perspective of truthful mechanism design, e.g., (Agrawal et al. 2022; Barak, Gupta, and Talgam-Cohen 2024; Istrate and Bonchis 2022).

Online Facility Location. The (metric uncapacitated) Facility Location is a classical optimization problem widely studied in both Operations Research and Computer Science (see e.g., (Drezner and Hamacher 2004; Shmoys 2000)). Its online version has received significant attention since its introduction by Meyerson (2001). Fotakis (2008) established a lower bound of $\Omega(\log n/\log\log n)$ and showed that the competitive ratio of Meyerson's algorithm is asymptotically optimal. For follow up work on OFL and its variants, we refer the reader to the survey of Fotakis (2011). Recently, there has been research interest in the dynamic variant of OFL (Cohen-Addad et al. 2019; Guo et al. 2020).

2 Model and Preliminaries

Notation. We consider a metric space $(\mathcal{M}, d)$, where $d : \mathcal{M} \times \mathcal{M} \to \mathbb{R}_{\geq 0}$ is a distance function, which is non-negative, symmetric and satisfies the triangle inequality. For a point $v \in M$ and a subset $U \subseteq M$, we let $d(v, U) = \min_{u \in U} d(v, u)$. We use the convention that $d(v, \emptyset) = \infty$.

Online Facility Location. In OFL, the input consists of a demand sequence $(v_1, \ldots, v_n)$ in a metric space $(\mathcal{M}, d)$. The demands arrive one at a time and must be assigned irrevocably to an open facility upon arrival. In response to the demand sequence, the online algorithm maintains a sequence of facility configurations $(\mathcal{F}_0, \mathcal{F}_1, \ldots, \mathcal{F}_n)$.

When a new demand v_t arrives, the algorithm decides whether to assign it to an existing facility or to open a new one. If the algorithm opens a new facility at location c, the facility cost increases by f and $\mathcal{F}_t = \mathcal{F}_{t-1} \cup \{c\}$. Otherwise, $\mathcal{F}_t = \mathcal{F}_{t-1}$. Finally, v_t is assigned to the nearest facility in $\mathcal{F}_t$ and the assignment cost increases by $d(\mathcal{F}_t, v_t)$. The goal is to minimize the total assignment and facility cost:

$$f \cdot |\mathcal{F}_n| + \sum_{t=1}^{n} d(\mathcal{F}_t, v_t)$$

We let $\mathcal{F}^*$ denote an optimal set of facility locations for the corresponding offline instance of Facility Location with demand set $\{v_1, \ldots, v_n\}$ (which is fully known in advance).

Then, $\text{OPT} := f \cdot |\mathcal{F}^*| + \sum_{t=1}^{n} d(v_t, \mathcal{F}^*)$ is the optimal cost. We let $k = |\mathcal{F}^*|$ denote the number of optimal facilities. For each demand v_t, $c_t^* = \arg\min_{c^* \in \mathcal{F}^*} d(v_t, c^*)$ denotes the optimal facility where v_t is assigned. We note that in the optimal solution $\mathcal{F}^*$, for any demand v_t, $d(v_t, c_t^*) \leq f$ (we would have opened a facility at v_t, otherwise).

Predictions. We consider a learning-augmented setting, where each new demand v_t is accompanied by a *prediction* p_t of the optimal facility c_t^* where v_t is assigned. The *prediction error* of p_t, denoted as $\eta(t)$ or $\eta(p_t)$, is the distance of the predicted location p_t to the optimal location c_t^*, i.e., $\eta(t) = d(p_t, c_t^*)$. We use the *total prediction error* $\eta_1 = \sum_{t=1}^{n} \eta(t)$ and the *maximum prediction error* $\eta_\infty = \max_t \eta(t)$ in order to quantify the inaccuracy of the predictions provided to the algorithm.

Competitive Ratio. We evaluate the performance of online algorithms using the *competitive ratio* (Borodin and El-Yaniv 1998). A randomized online algorithm is *cr*-competitive if for any sequence of demands (or demand-prediction pairs), the algorithm's expected cost is at most *cr* times the optimal cost for the corresponding offline instance, where the demand sequence is fully known in advance.

The competitive ratio of an OFL algorithm with predictions may depend on the number of demands n, the total prediction error η_1 and the maximum prediction error η_∞. We pay special attention to *consistency*, which is the competitive ratio when the predictions are perfect and $\eta_1 = \eta_\infty = 0$, and *robustness*, which is the worst-case competitive ratio over all possible values of η_1 and η_∞.

Notational Conventions. Demands are typically denoted by v_t (or v) and predictions by p_t (or p). We sometimes refer to *demand-prediction pairs* (v_t, p_t) simply as *requests*, for brevity. We use the term *optimal center* (or *center*) to refer to an optimal facility in $\mathcal{F}^*$ and the term *facility* to refer to an algorithm's facility in $\mathcal{F}$. We say that a demand v_t (or the demand-prediction pair (v_t, p_t)) is mapped to the optimal center c_t^* where v_t is assigned in the optimal solution.

When the timestep t is clear from the context or not important, we omit the subscript $_t$ and simply use v for demands and p for predictions. Moreover, we usually omit the subscript in $\mathcal{F}_{t-1}$ and use $\mathcal{F}$ to denote the current set of algorithm's facilities when a new demand v_t arrives. We let $C_{(v,p)}$ denote the algorithm's cost associated with the request (v, p), and let Asg_v^* (resp. Asg^*) denote v's (resp. the total) optimal assignment cost. In general, we use * to indicate the costs and the facilities in the optimal solution.

3 Online Facility Location with Predictions

In this section, we present online algorithm PREDOFL for OFL with predictions (OFLpred) and establish its competitive ratio. PREDOFL works similarly to (Meyerson 2001)'s algorithm, but with the predictions in place of the demands. Specifically, every time a demand-prediction pair (v_t, p_t) arrives, we open a facility at the predicted location p_t with probability $d(p_t, \mathcal{F})/f$ (instead of v_t and $d(v_t, \mathcal{F})/f$ in Meyerson's algorithm). If $d(p_t, \mathcal{F}) \geq f$, we open a new facility at p_t with certainty. The remainder of this section is devoted to the analysis of PREDOFL.

Algorithm 1: PREDOFL: OFL with Predictions

Input: Demand-prediction pairs $(v_1, p_1), \ldots, (v_n, p_n)$

```
F = ∅     {set of open facilities}
for each demand-prediction pair (v_t, p_t) do
    With probability min{1, d(F, p_t)/f}:
        F = F ∪ {p_t}     {new facility opens at p_t}
    Assign v_t to nearest facility in F with cost d(F, v_t)
end for
```

3.1 Main Properties

We first prove two main properties of PREDOFL, which are repeatedly used in the analysis of its competitive ratio.

Lemma 1. *Let (v, p) be a demand-prediction pair mapped to center c^*, and let $\mathcal{F}$ be the set of algorithm's facilities when (v, p) arrives. Then, the algorithm's cost for (v, p) is*

$$\mathbb{E}\left[C_{(v,p)}\right] \leq \min\{d(\mathcal{F}, p), f\} + d(\mathcal{F}, v) \quad (1)$$

$$C_{(v,p)} \leq f + \text{Asg}_v^* + \eta(p), \qquad \text{if } d(\mathcal{F}, p) \geq f \quad (2)$$

Lemma 2. *Let $\mathcal{P} = (p_1, \ldots, p_t, \ldots)$ be a sequence of predictions, where each p_t causes a new facility to open at p_t with probability $d(\mathcal{F}_{t-1}, p_t)/f < 1$. Then, the expected value of $\sum_{\tau=1}^{t} d(\mathcal{F}_{\tau-1}, p_\tau)$ just before p_{t+1} causes the first facility at a prediction in $\mathcal{P}$ to open is at most f.*

3.2 Competitive Ratio with Good Predictions

We upper bound the competitive ratio of PREDOFL in terms of $\frac{n\eta_\infty}{\text{OPT}}$, which is useful when predictions are very accurate.

Theorem 3. *For all $\eta_\infty \geq 0$, PREDOFL's competitive ratio is at most $3\frac{n\eta_\infty}{\text{OPT}} + 2$.*

Proof. PREDOFL opens a facility at p_1 with certainty. Since $d(c_1^*, p_1) \leq \eta_\infty$, by the definition of η_∞, and using Lemma 1, we obtain that the algorithm's cost is

$$C_{(v_1,p_1)} \leq f + d(v_1, c_1^*) + d(c_1^*, p_1) \leq f + \text{Asg}_{v_1}^* + \eta_\infty \quad (3)$$

By Lemma 1, PREDOFL's expected cost for each subsequent demand-prediction pair (v_t, p_t) is:

$$\mathbb{E}\left[C_{(v_t,p_t)}\right] \leq d(\mathcal{F}_{t-1}, p_t) + d(\mathcal{F}_{t-1}, v_t) \quad (4)$$

If when (v_t, p_t) arrives, $d(\mathcal{F}_{t-1}, c_t^*) \leq \eta_\infty$, then (i) $d(\mathcal{F}_{t-1}, p_t) \leq d(\mathcal{F}_{t-1}, c_t^*) + d(c_t^*, p_t) \leq 2\eta_\infty$; and (ii) $d(\mathcal{F}_{t-1}, v_t) \leq d(v_t, c_t^*) + (\mathcal{F}_{t-1}, c_t^*) \leq \text{Asg}_{v_t}^* + \eta_\infty$. Therefore, by (4), $\mathbb{E}\left[C_{(v_t,p_t)}\right] \leq \text{Asg}_{v_t}^* + 3\eta_\infty$.

Otherwise, for each optimal center c^*, due to the facility opening rule in PREDOFL and Lemma 2, the expected value of $\sum_t d(\mathcal{F}_{t-1}, p_t)$ over all pairs (v_t, p_t) mapped to c^* and assigned to an algorithm's facility before a facility within η_∞ to c^* opens is at most f. Therefore, using (4) and that

$$\begin{aligned} d(\mathcal{F}_{t-1}, v_t) &\leq d(\mathcal{F}_{t-1}, p_t) + d(p_t, c^*) + d(c^*, v_t) \\ &\leq d(\mathcal{F}_{t-1}, p_t) + \eta_\infty + \text{Asg}_{v_t}^*, \end{aligned}$$

we obtain that for each optimal center c^*, the expected assignment cost of all demand-prediction pairs (v_t, p_t) that are mapped to c^* and are assigned to an algorithm's facility as along as $d(\mathcal{F}_{t-1}, c^*) > \eta_\infty$ is at most f plus their optimal assignment cost plus their number times η_∞. Moreover, the algorithm's total cost for the first pair (v_t, p_t) which is mapped to c^* and opens a facility at p_t (and thus, causes $d(\mathcal{F}_t, c^*) \leq \eta_\infty$ for the first time) is bounded as in (3).

Putting everything together, we obtain that the expected total cost of PREDOFLis at most $2kf + \text{Asg}^* + 3n\eta_\infty$, which divided by $\text{OPT} = kf + \text{Asg}^*$ concludes the proof. □

Theorem 3 directly implies an upper bound of 2 on the consistency of PREDOFL.

3.3 Competitive Ratio with Arbitrary Predictions

We next bound the competitive ratio of PREDOFL in case where the predictions are useful but may be far from perfect.

Theorem 4. *For all sequences of n demand-prediction pairs with $\text{OPT} \geq \eta_1 \geq \eta_\infty > \text{OPT}/n$, the competitive ratio of* PREDOFL *is at most*

$$O\left(\frac{\log\left(\frac{n\eta_\infty}{\text{OPT}}\right)}{\log\left(\frac{\text{OPT}}{\eta_1}\log\left(\frac{n\eta_\infty}{\text{OPT}}\right)\right)}\right) \tag{5}$$

For the proof, we bound the algorithm's expected cost for each optimal center c^* separately and utilize the concept of phases, a standard technique in the analysis of OFL algorithms, see e.g., (Fotakis 2011). The approach involves defining a family of concentric balls with geometrically decreasing radii centered around each optimal center c^*. Specifically, we fix a pair of integers $m, \ell \geq 2$ so that

$$m^\ell \geq \frac{n\eta_\infty}{\text{OPT}}, \tag{6}$$

and define a family of balls $B_i(c^*) = \text{Ball}(c^*, \frac{\eta_\infty}{m^i})$ around each optimal center c^*, for all $i = 0, \ldots, \ell$. We say that a demand-prediction pair (v, p) mapped to c^* belongs to a ball $B_i(c^*)$, if the prediction $p \in B_i(c^*)$. We note that the radius of each B_ℓ is at most OPT/n, due to the definition of m and ℓ in (6). The radius of each B_0 is η_∞ and $B_0(c^*)$ includes all demand-prediction pairs (v, p) mapped to c^*.

For the analysis of PREDOFL, we quantify the progress of the algorithm's facilities $\mathcal{F}$ towards converging to an optimal center c^* through phases defined using the balls $B_i(c^*)$.

Definition 1 (Phases)**.** *An optimal center c^* is in* phase $i \in \{0, \ldots, \ell - 1\}$ *during the execution of* PREDOFL *as long as*

$$\frac{\eta_\infty}{m^{i+1}} < d(\mathcal{F}, c^*) \leq \frac{\eta_\infty}{m^i},$$

is in phase -1 *as along as $d(\mathcal{F}, c^*) > \eta_\infty$, and is in* phase ℓ *as along as $d(\mathcal{F}, c^*) \leq \frac{\eta_\infty}{m^\ell} \leq \frac{\text{OPT}}{n}$.*

For the analysis of the algorithm's cost, we partition the demand-prediction pairs that are mapped to an optimal center c^* and arrive in phase i of c^* into *close pairs* (or *close requests*) and *far pairs* (or *far requests*) based on the distance of the predictions to c^*.

Definition 2 (Far and Close Requests)**.** *A demand-prediction pair (v, p) that is mapped to an optimal center c^* and arrives in phase i of c^* is a* close pair *(or a* close request*), if $d(c^*, p) < \eta_\infty/m^{i+1}$, and a* far pair *(or a* far request*) otherwise. All demand-prediction pairs arriving in phase -1 (resp. ℓ) of c^* are* close *(resp.* far*) pairs.*

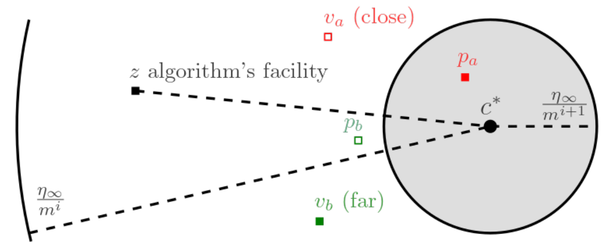

Figure 1: Phase i begins when PREDOFL opens a facility at distance at most η_∞/m^i to the optimal center c^*. Demand-prediction pairs arriving in phase i with predictions inside $\text{Ball}(c^*, \eta_\infty/m^{i+1})$ are *close pairs* (or *close requests*), while the remaining demand-prediction pairs are *far pairs* (or *far requests*). Hence, the pair (v_a, p_a) is a close one, because the prediction $p_a \in B_{i+1}(c^*)$, while the pair (v_b, p_b) is a far one, because $p_b \notin B_{i+1}(c^*)$. Phase i ends as soon as a close demand-prediction pair causes a new facility to open.

We next focus on the demand-prediction pairs mapped to an optimal center c^*. We bound the expected cost of PREDOFL separately for the far and the close demand-prediction pairs. Lemma 3 exploits the definition of phases in order to bound the expected cost of far and close requests.

For an optimal center c^*, we let $C_{\text{far}}(c^*)$ and $C_{\text{close}}(c^*)$ (respectively, $\text{Asg}^*_{\text{far}}(c^*)$ and $\text{Asg}^*_{\text{close}}(c^*)$, and $\eta_{1,\text{far}}(c^*)$ and $\eta_{1,\text{close}}(c^*)$) denote the algorithm's expected cost (respectively, optimal assignment cost and total prediction error) for far and close pairs mapped to c^*, respectively. We also let $\text{Far}(c^*)$ denote the set of far requests mapped to c^*.

Lemma 3 (Cost of Far and Close Requests)**.** *For any pair of integers m, ℓ that satisfy* (6) *and any optimal center c^*, the expected total cost incurred by* PREDOFL *for all far and close demand-prediction pairs mapped to c^* is*

$$\begin{aligned}\mathbb{E}[C_{\text{far}}(c^*)] &\leq \text{Asg}^*_{\text{far}}(c^*) + 2\,|\text{Far}(c^*)|\,\tfrac{\text{OPT}}{n} \\ &\quad + (2m+1)\eta_{1,\text{far}}(c^*) \\ \mathbb{E}[C_{\text{close}}(c^*)] &\leq \text{Asg}^*_{\text{close}}(c^*) + \eta_{1,\text{close}}(c^*) + 2(\ell+1)f\end{aligned}$$

To prove Lemma 3, we observe that for every far request (v, p) that arrives in phase $i < \ell$ and is mapped to an optimal center c^*, the definitions 1 and 2 imply that:

$$d(\mathcal{F}, c^*) \leq \frac{\eta_\infty}{m^i} \leq m\, d(p, c^*) = m\, \eta_1(p)\,. \tag{7}$$

Then, using Lemma 1, we obtain that the algorithm's expected cost due to such a far request (v, p) is:

$$\begin{aligned}\mathbb{E}\left[C_{(v,p)}\right] &\leq \big(d(\mathcal{F}, c^*) + d(c^*, p)\big) + \big(d(\mathcal{F}, c^*) + d(c^*, v)\big) \\ &\leq (2m+1)\eta_1(p) + \text{Asg}^*_v\,.\end{aligned}$$

For far requests arriving in phase ℓ, we use that $d(\mathcal{F}, c^*) \leq \text{OPT}/n$ (instead of (7)) and work similarly.

As for the expected cost of close requests, we use Lemma 1 and bound their cost for each phase separately. We show that due to Lemma 2, their expected cost up to the point where the first facility due to a close prediction opens, which causes a new phase to begin, is at most f plus additional terms due to the error of the predictions and the optimal assignment cost of the demands (see also the second case in the proof of Theorem 3).

Putting far and close demand-prediction pairs together, summing up over all k optimal centers and using that the total number of requests is n, we conclude[2] that:

Corollary 1. *For any pair of integers m, ℓ that satisfy* (6)*, the expected total cost incurred by* PREDOFL*is at most*

$$2(\ell+1)kf + \mathrm{Asg}^* + 2\,\mathrm{OPT} + (2m+1)\eta_1 \tag{8}$$

Since $\mathrm{OPT} = kf + \mathrm{Asg}^*$, the competitive ratio is $O(\ell + m\eta_1/\mathrm{OPT})$. To optimize the upper bound on the competitive ratio and prove Theorem 4, we consider the case where $\eta_\infty > \mathrm{OPT}/n$ (if $\eta_\infty \leq \mathrm{OPT}/n$, the competitive ratio is at most 5, by Theorem 3) and $\eta_1 \leq \mathrm{OPT}$ (if $\eta_1 > \mathrm{OPT}$, we should guide algorithm's decisions by the demand points). We let $\ell = m\eta_1/\mathrm{OPT}$ and select m so that (6) is satisfied:

$$m^{\frac{m\eta_1}{\mathrm{OPT}}} \geq \frac{n\eta_\infty}{\mathrm{OPT}} \Rightarrow m\log m \geq \frac{\mathrm{OPT}}{\eta_1}\log\Big(\frac{n\eta_\infty}{\mathrm{OPT}}\Big)\,. \tag{9}$$

We note that (9) gives $m = \exp(W_0(B))$, where $B = \frac{\mathrm{OPT}}{\eta_1}\log(\frac{n\eta_\infty}{\mathrm{OPT}})$ and $W_0(x)$ is the Lambert W function. Using the bound on $W_0(x)$ in (Hoorfar and Hassani 2008), for $x > e$ (which imposes an upper bound on η_1 with respect to OPT and a lower bound on η_∞ with respect to OPT/n, so that $\frac{\mathrm{OPT}}{\eta_1}\log(\frac{n\eta_\infty}{\mathrm{OPT}}) > e$), we get that $m = \Theta(\frac{B}{\log B})$, which concludes the proof of Theorem 4. $\square$

4 Combining Different OFL Algorithms

We next show how to combine two OFL algorithms, A_0 and A_1, and obtain $\mathrm{MIN}(A_0, A_1)$, which for every sequence of demand-prediction pairs achieves a total cost within thrice total cost of the best of A_0 and A_1 on the same sequence.

$\mathrm{MIN}(A_0, A_1)$ applies to OFL the binary search approach, also used for the Cow Path problem (Kao, Reif, and Tate 1996). $\mathrm{MIN}(A_0, A_1)$ has access to OFL algorithms A_0 and A_1, and receives online a sequence $(v_1, p_1), \ldots, (v_n, p_n)$ of demand-prediction pairs. $\mathrm{MIN}(A_0, A_1)$ simulates both A_0 and A_1 on $(v_1, p_1), \ldots, (v_n, p_n)$, proceeds in phases guided by cost doubling, and aims to *follow* (i.e., to adopt the set of facilities of) the less costly of them. Every time the total cost of the algorithm currently followed exceeds the next power of 2, a new phase begins. While in present phase, $\mathrm{MIN}(A_0, A_1)$ follows the less costly of the two algorithms at the beginning of the phase (see also Algorithm 2).

We let $C_n(A_i)$ (resp. $C_n^{\min}$) denote the total cost of algorithm A_i (resp. of $\mathrm{MIN}(A_0, A_1)$) for the entire request sequence. The following is the main result of this section:

[2]Intuitively, η_∞ gives an upper bound on the maximum distance of the algorithm's initial facility to the optimal center c^* (this is quantified in (6)), while η_1 imposes an upper bound on the length of each phase in the convergence process (note that η_1 is multiplied by m in (8)). Balancing the number of phases, denoted by ℓ, against the phase length, denoted by m, results in an optimal competitive ratio. This explains why both η_∞ and η_1 are necessary in order to accurately express the competitive ratio of OFL with predictions.

Algorithm 2: $\mathrm{MIN}(A_0, A_1)$ combines algorithms A_0, A_1

Input: OFL algorithms A_0 and A_1,
demand-prediction pairs $(v_1, p_1), \ldots, (v_n, p_n)$

$\ell = 0, i = 0$ {phase index ℓ, algorithm index i}
$C(A_0) = C(A_1) = 0$ {$C(A_i)$ is A_i's cost so far}
$\mathcal{F}_0 = \mathcal{F}_1 = \emptyset$ {$\mathcal{F}_i$ is A_i's current set of facilities}
$\mathcal{F} = \emptyset$ {$\mathcal{F}$ is $\mathrm{MIN}(A_0, A_1)$'s set of facilities}
for each demand-prediction pair (v_t, p_t) **do**
  Serve (v_t, p_t) using A_0 and update $\mathcal{F}_0$ and $C(A_0)$
  Serve (v_t, p_t) using A_1 and update $\mathcal{F}_1$ and $C(A_1)$
  if $C(A_i) > 2^\ell$ **then**
    $\ell = \ell + 1$ {proceed to next phase}
    if $C(A_i) > C(A_{1-i})$ **then**
      $i = 1 - i$ {switch from A_i to less costly A_{1-i}}
    end if
    $\mathcal{F} = \mathcal{F} \cup \mathcal{F}_i$ {update $\mathcal{F}$ with current $\mathcal{F}_i$}
  end if
  Assign v_t to nearest facility in $\mathcal{F}$ with cost $d(\mathcal{F}, v_t)$
end for

Theorem 5. *Let A_0 and A_1 be algorithms for OFLpred. Then, for every sequence of n demand-prediction pairs, the total cost of* $\mathrm{MIN}(A_0, A_1)$*, described in Algorithm 2, is*

$$C_n^{\min} \leq 3\min\{C_n(A_0), C_n(A_1)\}\,.$$

Proofsketch. At any point in time, the total cost of $\mathrm{MIN}(A_0, A_1)$ is at most the total cost of the algorithm A_i currently followed plus the total cost of A_{1-i} up to the last pair where $\mathrm{MIN}(A_0, A_1)$ followed A_{1-i}. Then, due to the use of cost doubling in the definition of phases, the most costly of the two algorithms A_0 and A_1 cannot contribute more than twice the total cost of $\min\{C_n(A_0), C_n(A_1)\}$ to the final total cost of $\mathrm{MIN}(A_0, A_1)$. $\square$

For randomized algorithms A_0 and A_1, Theorem 5 holds for the algorithms' realized cost on any request sequence. Hence, the expected cost of $\mathrm{MIN}(A_0, A_1)$ is at most 3 times the minimum of the expected costs of A_0 and A_1.

The Proof of Theorem 1. To obtain Theorem 1, we apply Theorem 5 to PREDOFL and to the algorithm of Jiang et al. (2022) for uniform facility costs. The latter algorithm, for each request (v_t, p_t), opens a pair of new facilities, one at v_t and another at p_t, with probability $\min\{1, d(\mathcal{F}, v_t)/f\}$.

For the latter algorithm, $\mathbb{E}\big[C_{(v,p)}\big] \leq \min\{2f, 3d(\mathcal{F}, v)\}$ for any pair (v, p) (the proof is similar to Lemma 1). By defining (6) and phases wrt. $\frac{n}{\mathrm{OPT}}\min\{\eta_\infty, f\}$, using that $\mathrm{Asg}_v^* \leq f$ for any demand v, and adapting the proof of Lemma 3, we can show that for any $m, \ell \geq 2$ that satisfy $m^\ell \geq \frac{n}{\mathrm{OPT}}\min\{\eta_\infty, f\}$, the expected total cost is at most

$$3(\ell+1)kf + 3(m+1)\mathrm{Asg}^* + 3\,\mathrm{OPT}$$

Then, working as in the last paragraph of Section 3 and using that $\mathrm{OPT} \geq f$, we obtain the following upper bound on the competitive ratio (which is independent of the total prediction error η_1):

$$O\left(\frac{\log\min\{\frac{n\eta_\infty}{\mathrm{OPT}}, n\}}{\log\log\min\{\frac{n\eta_\infty}{\mathrm{OPT}}, n\}}\right)\,. \tag{10}$$

Theorem 5 shows how to get an online algorithm with thrice the minimum of the competitive ratio of PREDOFL (as established in Theorem 3 for small η_∞ and in Theorem 4 for $\text{OPT} \geq \eta_1 \geq \eta_\infty > \text{OPT}/n$) and the competitive ratio in (10), which is achieved by (Jiang et al. 2022)'s algorithm for uniform facility costs and holds for all $\eta_1 \geq 0$.

5 Lower Bound

We next show that the dependence of PREDOFL's competitive ratio on η_1 (and also on η_∞ and n) is essentially optimal.

Theorem 6. *For every large enough n, any $\alpha \in (27/n, 1/3)$ and any $\beta \in (3\alpha, 1]$ (α and β may depend on n), there are instances of OFLpred with n demand-prediction pairs, $\eta_\infty/\text{OPT} = \alpha$ and $\beta/3 \leq \eta_1/\text{OPT} \leq \beta$ where any randomized algorithm has competitive ratio at least*

$$\Omega\left(\frac{\log\left(\frac{n\eta_\infty}{\text{OPT}}\right)}{\log\left(\frac{\text{OPT}}{\eta_1}\log\left(\frac{n\eta_\infty}{\text{OPT}}\right)\right)}\right) = \Omega\left(\frac{\log(\alpha n)}{\log\left(\frac{\log(\alpha n)}{\beta}\right)}\right)$$

Proofsketch. We generalize the lower bound of Fotakis (2008). Using (Yao 1977)'s principle, we obtain the lower bound by considering the expected cost of any deterministic algorithm against an appropriately constructed probability distribution on sequences of n demand-prediction pairs.

The metric is a binary Hierarchically Well-Separated Tree T. Given n, α and β, we select $m \geq \max\{1/\beta, 4\}$ and $\ell = \lceil \beta m \rceil$ so that m is the least integer that satisfies $m^\ell \geq \alpha n$. T has $\ell + 1 \geq 2$ levels. We consider T's root as being at level 0 and T's leaves as being at level ℓ. For each vertex v, we let T_v denote the subtree of T rooted at v.

The distance of the root to its children is m^ℓ/β. The edge lengths along each path from the root to a leaf decrease by a factor of m at every level. Hence, the distance of any vertex at level $i \in \{0, \dots, \ell - 1\}$ to its children is $m^{\ell-i}/\beta$. The following hold for any level-i vertex v_i: (i) the distance of v_i to any vertex in T_{v_i} is at most $(m^{\ell-i}/\beta)\frac{m}{m-1}$; and (ii) the distance of v_i to any vertex not in T_{v_i} is a least $m^{\ell+1-i}/\beta$.

Demand Sequence. The demand sequence is divided into $\ell + 1$ phases. Phase 0 consists of $1/\alpha$ demands located at the root v_0 of T (if necessary, we round up the number of demands to the nearest integer). After the end of phase $i-1$, $i \in \{1, \dots, \ell\}$, the adversary selects v_i uniformly at random from the two children of v_{i-1}. In the next phase i, m^i/α demands arrive at v_i. The total number of demands is:

$$\sum_{i=0}^{\ell} \frac{m^i}{\alpha} = \frac{1}{\alpha}\frac{m^{\ell+1}-1}{m-1} \geq \frac{m^\ell}{\alpha}\,. \tag{11}$$

By removing demands from the last phase (if necessary), we ensure that the total number of demands is $n \leq m^\ell/\alpha$.

Facility Opening Cost. We set $f = \frac{m-2}{m-1} \cdot \frac{m^{\ell+1}}{\alpha}$.

Optimal Cost. We can show that $\text{OPT} \leq 2m^{\ell+1}/\alpha \leq 3f$, and that the optimal assignment cost Asg^* is at least $m^{\ell+1}/\alpha$, which implies that $\text{Asg}^* \geq \text{OPT}/2$, and at most $\frac{m^{\ell+1}}{\alpha}\frac{m}{m-1}$, which implies that $\text{Asg}^* \leq 2\text{OPT}/3$.

Prediction Sequence. Similarly to the demand sequence, the prediction sequence is divided into $\ell + 1$ phases. For the first demand of phase 0 located at v_0, the corresponding prediction is at distance $\eta_\infty = \alpha\text{OPT} \leq 2m^{\ell+1}$ to v_ℓ (which is the leaf where the optimal facility is located). If $\alpha\text{OPT} > d(v_0, v_\ell)$, we add a new root v_0' to T at distance η_∞ to v_ℓ and place the corresponding prediction at v_0'.

For each demand of phase $i = 0, \dots, \ell - 1$ located at v_i (with the exception of the first demand of phase 0), the prediction is at distance $\beta d(v_i, v_\ell)$ to v_ℓ. The predictions for the demands of phase ℓ are located at v_ℓ. Since $\beta \geq 1/m$, the predictions for the demands located at v_i are located along the edge connecting v_i to the location v_{i+1} of the demands of phase $i + 1$. Using $\beta \geq 3\alpha$ and $\alpha \leq 1/3$, we can show that the total prediction error is $\beta\text{OPT}/3 \leq \eta_1 \leq \beta\text{OPT}$.

Algorithm's Cost. We can show that the expected cost of any deterministic algorithm on this probability distribution over request sequences is at least $\ell f/4$.

Since $\text{OPT} \leq 3f$, the competitive ratio is $\Omega(\ell) = \Omega(\beta m)$. We recall that $\ell = \beta m$ and that m is chosen so that $m^{\beta m} \geq \alpha n$. Therefore, using that $\alpha = \eta_\infty/\text{OPT}$ and that $\beta \leq 3\eta_1/\text{OPT}$, we obtain that

$$m^{\beta m} \geq \alpha n \Rightarrow m \log m \geq \frac{3\text{OPT}}{\eta_1}\log\left(\frac{n\eta_\infty}{\text{OPT}}\right) \tag{12}$$

We note that (12), which gives a lower bound on the competitive ratio of any randomized algorithm for OFLpred as a function of n, $\beta \approx \frac{\eta_1}{\text{OPT}} < 1$ and $\alpha n = \frac{n\eta_\infty}{\text{OPT}} < n/3$, is essentially identical to (9), which determines the competitive ratio of PREDOFL. Working as in Section 3, we obtain that for $B = \frac{3\text{OPT}}{\eta_1}\log(\frac{n\eta_\infty}{\text{OPT}}) > e$, (12) is satisfied by $m = \Omega(\frac{B}{\log B})$, which implies a lower bound of $\Omega(\beta m) = \Omega(\beta\frac{B}{\log B})$. Using $\beta \geq \eta_1/\text{OPT}$ and $B = \frac{3\text{OPT}}{\eta_1}\log(\frac{n\eta_\infty}{\text{OPT}})$ in $\Omega(\beta\frac{B}{\log B})$, we obtain the desired lower bound. □

Remark 1. *Theorem 6 can be regarded as a refined version, also parameterized by η_1/OPT, of (Jiang et al. 2022)'s lower bound. Specifically, in the lower bound of (Jiang et al. 2022, Theorem F.1), which is based on a similar metric space and demand sequence, for any fixed $\alpha = \eta_\infty/\text{OPT} \in (0, 1]$, the predictions corresponding to the demands at the first levels of T (those closer to the root) are located at distance η_∞ to v_ℓ, so that the maximum prediction error is η_∞. For the remaining requests, the predictions are at the same locations as the corresponding demand points, which results in a total prediction error $\eta_1 = \Theta(\text{OPT})$, even if $\alpha = \eta_\infty/\text{OPT}$ is very small, e.g., even if $\eta_\infty = \text{OPT}/n^{(1-\delta)}$, for any constant $\delta > 0$. Hence, the lower bound of (Jiang et al. 2022) does not quantify how fast the competitive ratio of OFLpred can improve as η_1/OPT decreases (assuming a fixed value of $\alpha = \eta_\infty/\text{OPT}$). Thus, it fails to differentiate, as far as their best possible competitive ratio is concerned, between instances described in cases (i)-(iii) in Section 1.1. To close this gap, Theorem 6 establishes a lower bound on the best possible competitive ratio of OFLpred which for every fixed $\alpha = \eta_\infty/\text{OPT}$, is also parameterized by $\beta \approx \frac{\eta_1}{\text{OPT}}$ and can be applied to tell such instances apart as far as their best possible competitive ratio is concerned.* □

Acknowledgements

This work has been partially supported by project MIS 5154714 of the National Recovery and Resilience Plan Greece 2.0 funded by the European Union under the NextGenerationEU Program. Nikolas Patris was supported by an ICS research award from UC Irvine.

References

Agrawal, P.; Balkanski, E.; Gkatzelis, V.; Ou, T.; and Tan, X. 2022. Learning-Augmented Mechanism Design: Leveraging Predictions for Facility Location. In *Proc. of the 23rd ACM Conference on Economics and Computation (EC 2022)*, 497–528. ACM.

Almanza, M.; Chierichetti, F.; Lattanzi, S.; Panconesi, A.; and Re, G. 2021. Online Facility Location with Multiple Advice. In *Proc. of the 34th Conference on Neural Information Processing Systems (NeurIPS 2021)*, 4661–4673.

Antoniadis, A.; Gouleakis, T.; Kleer, P.; and Kolev, P. 2023. Secretary and online matching problems with machine learned advice. *Discrete Optimization*, 48(Part 2): 100778.

Argue, C.; Frieze, A. M.; Gupta, A.; and Seiler, C. 2022. Learning from a Sample in Online Algorithms. In *Proc. of the 35th Conference on Neural Information Processing Systems (NeurIPS 2022)*.

Azar, Y.; Panigrahi, D.; and Touitou, N. 2022. Online Graph Algorithms with Predictions. In *Proc. of the 2022 ACM-SIAM Symposium on Discrete Algorithms (SODA 2022)*, 35–66. SIAM.

Bamas, É.; Maggiori, A.; Rohwedder, L.; and Svensson, O. 2020. Learning Augmented Energy Minimization via Speed Scaling. In *Proc. of the 33rd Conference on Neural Information Processing Systems 2020 (NeurIPS 2020)*.

Bamas, É.; Maggiori, A.; and Svensson, O. 2020. The Primal-Dual method for Learning Augmented Algorithms. In *Proc. of the 33rd Conference on Neural Information Processing Systems 2020 (NeurIPS 2020)*.

Barak, Z.; Gupta, A.; and Talgam-Cohen, I. 2024. MAC Advice for Facility Location Mechanism Design. *CoRR*, abs/2403.12181.

Borodin, A.; and El-Yaniv, R. 1998. *Online Computation and Competitive Analysis*. Cambridge University Press. ISBN 978-0-521-56392-5.

Cohen-Addad, V.; Hjuler, N.; Parotsidis, N.; Saulpic, D.; and Schwiegelshohn, C. 2019. Fully Dynamic Consistent Facility Location. In *Proc. of 32nd Conference on Neural Information Processing Systems (NeurIPS 2019)*, 3250–3260.

Drezner, Z.; and Hamacher, H. 2004. *Facility Location: Applications and Theory*. Springer.

Fiat, A.; Karp, R. M.; Luby, M.; McGeoch, L. A.; Sleator, D. D.; and Young, N. E. 1991. Competitive Paging Algorithms. *J. Algorithms*, 12(4): 685–699.

Fotakis, D. 2008. On the Competitive Ratio for Online Facility Location. *Algorithmica*, 50(1): 1–57.

Fotakis, D. 2011. Online and incremental algorithms for facility location. *SIGACT News*, 42(1): 97–131.

Fotakis, D.; Gergatsouli, E.; Gouleakis, T.; Patris, N.; and Tolias, T. 2024. Improved Bounds for Online Facility Location with Predictions. arXiv:2107.08277.

Guo, X.; Kulkarni, J.; Li, S.; and Xian, J. 2020. On the Facility Location Problem in Online and Dynamic Models. In *Proc. of the 23rd Conference on Approximation, Randomization, and Combinatorial Optimization (APPROX/RANDOM 2020)*, volume 176 of *LIPIcs*, 42:1–42:23. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.

Gupta, A.; Panigrahi, D.; Subercaseaux, B.; and Sun, K. 2022. Augmenting Online Algorithms with ε-Accurate Predictions. In *Proc. of the 35th Conference on Neural Information Processing Systems (NeurIPS 2022)*.

Hoorfar, A.; and Hassani, M. 2008. Inequalities on the Lambert W function and hyperpower function. *J. Inequal. Pure and Appl. Math*, 9(2): 5–9.

Istrate, G.; and Bonchis, C. 2022. Mechanism Design With Predictions for Obnoxious Facility Location. *CoRR*, abs/2212.09521.

Jiang, S. H.; Liu, E.; Lyu, Y.; Tang, Z. G.; and Zhang, Y. 2022. Online Facility Location with Predictions. In *Proc. of the 10th International Conference on Learning Representations (ICLR 2022)*.

Kao, M.; Reif, J.; and Tate, S. 1996. Searching in an Unknown Environment: An Optimal Randomized Algorithm for the Cow-Path Problem. *Information and Computation*, 131(1): 63–79.

Lattanzi, S.; Lavastida, T.; Moseley, B.; and Vassilvitskii, S. 2020. Online Scheduling via Learned Weights. In *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms (SODA 2020)*, 1859–1877.

Lykouris, T.; and Vassilvitskii, S. 2018. Competitive Caching with Machine Learned Advice. In *Proc. of the 35th International Conference on Machine Learning (ICML 2018)*, 3302–3311.

Medina, A.; and Vassilvitskii, S. 2017. Revenue Optimization with Approximate Bid Predictions. In *Proc. of 30th Conference on Neural Information Processing Systems (NeurIPS 2017)*, 1858–1866.

Meyerson, A. 2001. Online Facility Location. In *Proc. of the 42nd Symposium on Foundations of Computer Science (FOCS 2001)*, 426–431. IEEE.

Mitzenmacher, M.; and Vassilvitskii, S. 2020. Algorithms with Predictions. In Roughgarden, T., ed., *Beyond the Worst-Case Analysis of Algorithms*, 646–662. Cambridge University Press.

Purohit, M.; Svitkina, Z.; and Kumar, R. 2018. Improving Online Algorithms via ML Predictions. In *Proc. of 31st Conference on Neural Information Processing Systems (NeurIPS 2018)*, 9684–9693.

Rohatgi, . 2020. Near-Optimal Bounds for Online Caching with Machine Learned Advice. In *Proc. of the 2020 ACM-SIAM Symposium on Discrete Algorithms, (SODA 2020)*, 1834–1845.

Shmoys, D. 2000. Approximation Algorithms for Facility Location Problems. In *3rd Workshop on Approximation Algorithms for Combinatorial Optimization*, volume 1913 of *LNCS*, 27–33.

Wei, A. 2020. Better and Simpler Learning-Augmented Online Caching. In *Proc. of the 23rd Conference on Approximation, Randomization, and Combinatorial Optimization (APPROX/RANDOM 2020)*, volume 176 of *LIPIcs*, 60:1–60:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.

Yao, A. C. 1977. Probabilistic Computations: Toward a Unified Measure of Complexity (Extended Abstract). In *Proc. of the 18th Symposium on Foundations of Computer Science (FOCS 1977)*, 222–227. IEEE Computer Society.