

# BETTER AND SIMPLER LEARNING-AUGMENTED ONLINE CACHING

ALEXANDER WEI

*Harvard University*

**ABSTRACT.** Lykouris and Vassilvitskii (ICML 2018) introduce a model of online caching with machine-learned advice, where each page request additionally comes with a prediction of when that page will next be requested. In this model, a natural goal is to design algorithms that (1) perform well when the advice is accurate and (2) remain robust in the worst case a la traditional competitive analysis. Lykouris and Vassilvitskii give such an algorithm by adapting the MARKER algorithm to the learning-augmented setting. In a recent work, Rohatgi (SODA 2020) improves on their result with an approach also inspired by randomized marking. We continue the study of this problem, but with a somewhat different approach: We consider combining the BLINDORACLE algorithm, which just naïvely follows the predictions, with an optimal competitive algorithm for online caching in a black-box manner. The resulting algorithm outperforms all existing approaches while being significantly simpler. Moreover, we show that combining BLINDORACLE with LRU is in fact optimal among deterministic algorithms for this problem.

## 1. INTRODUCTION

Traditionally, the study of online algorithms focuses on *worst-case* robustness, with algorithms providing the same competitive guarantee against the offline optimal over all inputs. In recent years, however, there has been a surge of interest in studying online algorithms in the presence of structured inputs [MV17; LV18; PSK18; GP19; KPS<sup>+</sup>19; KPS<sup>+</sup>20; Roh20; LLM<sup>+</sup>20; Mit20]. A principal motivation for these works is the philosophy of beyond worst-case analysis [KP00; Rou19]: Many practical settings have inputs that follow restricted patterns, making classical worst-case competitive analysis too pessimistic to inform practice. In particular, the worst-case examples classical competitive analyses guard against often do not materialize. Furthermore, algorithms designed with the worst case in mind can be hamstrung by these considerations, ending up unnatural and losing performance on “nice” inputs.

*Learning-augmented* online algorithms, introduced by Lykouris and Vassilvitskii [LV18] and Purohit, Svitkina, and Kumar [PSK18], is a beyond worst-case framework motivated by the powerful predictive abilities of modern machine learning. The structure of the input is assumed to come in the form of a machine-learned predictor that provides predictions of future inputs. A concern with this setup may be that machine learning models typically have few worst-case guarantees. Nonetheless, with learning-augmented algorithms, we want the best of both worlds: Given a predictor, the objective is to design algorithms that (1) perform well in the optimistic scenario, where the predictor has low error, and (2) remain robust in the classical worst-case sense, when the predictor can be arbitrarily bad. That is, we would like our algorithm to be  $c(\eta)$ -competitive against the offline optimal on all inputs, where  $c$  is a function of the predictor error  $\eta$  such that  $\max_{\eta} c(\eta) \leq \gamma$  for some  $\gamma \geq 1$ . Such an algorithm is said to be  $\gamma$ -robust.

The focus of our work is learning-augmented online caching [LV18; Roh20]. In the *online caching* (a.k.a. *online paging*) problem, one seeks to maintain a cache of size  $k$  online while serving requests for pages that may or may not be in the cache. For simplicity, we assume that pages must always be served from the cache and that bringing a page into the cache has unit cost. (In particular, if the cache is full, bringing a page into the cache requires also *evicting* a page already in the cache.) Thus, we seek to minimize the number of requests for which the page is not already in the cache, i.e., the number of *cache misses*. This is

---

*E-mail address:* weia@college.harvard.edu.

a classical online problem that has been the subject of extensive study over the past several decades (see [BE98] for an overview). From the worst-case perspective, this problem is well-understood for not only the version stated above [ST85; FKL<sup>+</sup>91; ACN00], but also for weighted generalizations [BBN12a; BBN12b].

Online caching in the learning-augmented context was first considered by Lykouris and Vassilvitskii [LV18]. They introduce a model of prediction where the predictor, upon the arrival of each page, predicts the next time that this page will be requested. They show that the BLINDORACLE algorithm, which follows the predictor naïvely and evicts the page with the latest predicted arrival time, can have unbounded competitive ratio (i.e., is *non-robust*). They then give a different algorithm, PREDICTIVEMARKER, based on the MARKER algorithm of Fiat, Karp, Luby, McGeoch, Sleator, and Young [FKL<sup>+</sup>91], that achieves a competitive ratio of

$$2 + O\left(\min\left(\sqrt{\frac{\eta}{\text{OPT}}}, \log k\right)\right),$$

where  $\eta$  is the  $\ell_1$  error of the predictor and OPT is the cost of the offline optimal. (See Section 2.1 for precise definitions.) In [Roh20], Rohatgi introduces the LNONMARKER algorithm, which is also based on randomized marking (but eschews the framework somewhat), and shows that this algorithm obtains a competitive ratio of

$$O\left(1 + \min\left(\frac{\log k}{k} \frac{\eta}{\text{OPT}}, \log k\right)\right).$$

This bound is obtained by first constructing a non-robust algorithm and then using a black-box combination technique discussed in [LV18] to combine this non-robust algorithm with the MARKER algorithm. Rohatgi also provides a lower bound of

$$\Omega\left(\min\left(\log\left(\frac{1}{k \log k} \frac{\eta}{\text{OPT}}\right), \log k\right)\right)$$

for the competitive ratio of any learning-augmented online algorithm for caching in terms of  $k$ , OPT, and  $\eta$ .

**1.1. Our Contribution.** We show that the strikingly simple approach of combining BLINDORACLE with an  $O(\log k)$ -competitive online caching algorithm (e.g., MARKER) in a black-box fashion obtains a competitive ratio bound of

$$O\left(1 + \min\left(\frac{1}{k} \frac{\eta}{\text{OPT}}, \log k\right)\right),$$

improving on that of LNONMARKER. Thus, although BLINDORACLE is non-robust [LV18], we show that it should not be abandoned entirely. In fact, it has excellent performance when  $\eta/\text{OPT}$  is small. So when this algorithm is combined with a  $O(\log k)$ -competitive algorithm, we start seeing performance improvement over the robust algorithm starting at  $\eta/\text{OPT} = O(k \log k)$ .

And although our improvement in the competitive ratio is slight, previous approaches for learning-augmented online caching [LV18; Roh20] have relied on much more intricate constructions based on randomized marking. We therefore believe that our simple approach may yield better practical performance and may generalize more readily to other learning-augmented settings. (Indeed, we note that the deterministic combining algorithm is particularly simple: Just “follow” the algorithm with the better performance thus far.)

We also give precise bounds on the constant factors in the competitive ratios that we obtain. Fiat, Karp, Luby, McGeoch, Sleator, and Young [FKL<sup>+</sup>91] and Blum and Burch [BB00] provide optimal bounds for combining online algorithms online in a black-box manner, with better constants than the approach discussed in [LV18] and applied in [Roh20]. By composing a careful competitive analysis of BLINDORACLE with these “combiners,” we obtain constants in the competitive ratio that are lower than those of previous work.

Finally, we show that combining BLINDORACLE with a  $k$ -competitive deterministic algorithm (e.g., LRU [ST85]) is the best one could hope to do among deterministic algorithms for learning-augmented online caching. In particular, we show that a linear dependence on  $\eta/(k \cdot \text{OPT})$  in the competitive ratio is necessary.

Therefore, if a logarithmic dependence on  $\eta/(k \cdot \text{OPT})$  is to be achieved, as in Rohatgi’s lower bound, then randomization is needed, (perhaps surprisingly) even in the regime where  $\eta/(k \cdot \text{OPT})$  is bounded.

Stated formally, our main result analyzing `BLINDORACLE` is the following:

**Theorem 1.1.** *For learning-augmented online caching, `BLINDORACLE` obtains a competitive ratio of*

$$\min\left(1 + 2\frac{\eta}{\text{OPT}}, 2 + \frac{4}{k-1}\frac{\eta}{\text{OPT}}\right),$$

where  $\eta$  is the  $\ell_1$  loss incurred by the predictor and  $\text{OPT}$  is the offline optimal cost. (For precise definitions, see Section 2.1.)

Plugging this bound into the results of Fiat, Karp, Luby, McGeoch, Sleator, and Young [FKL<sup>+</sup>91] and Blum and Burch [BB00] for competitively combining online algorithms online (see Section 2.2) yields the following corollaries:

**Corollary 1.2.** *There exists a deterministic algorithm for learning-augmented online caching that achieves a competitive ratio of*

$$2 \min\left(\min\left(1 + 2\frac{\eta}{\text{OPT}}, 2 + \frac{4}{k-1}\frac{\eta}{\text{OPT}}\right), k\right).$$

**Corollary 1.3.** *There exists a randomized algorithm for learning-augmented online caching that achieves a competitive ratio of*

$$(1 + \varepsilon) \min\left(\min\left(1 + 2\frac{\eta}{\text{OPT}}, 2 + \frac{4}{k-1}\frac{\eta}{\text{OPT}}\right), H_k\right)$$

for any  $\varepsilon \in (0, 1/4)$ .<sup>1</sup> (Here,  $H_k = 1 + \frac{1}{2} + \frac{1}{3} + \dots + \frac{1}{k} = \ln(k) + O(1)$  is the  $k$ -th harmonic number.)

Finally, we state our matching lower bound for Corollary 1.2 on deterministic algorithms for learning-augmented online caching:

**Theorem 1.4.** *The competitive ratio bound for any deterministic learning-augmented online caching algorithm must be at least*

$$1 + \Omega\left(\min\left(\frac{1}{k}\frac{\eta}{\text{OPT}}, k\right)\right).$$

**1.2. Related Work.** In addition to the predecessor works by Lykouris and Vassilvitskii [LV18] and Rohatgi [Roh20] on learning-augmented online caching, there have been several other recent papers in the space of learning-augmented online algorithms: Medina and Vassilvitskii [MV17] study repeated posted-price auctions, Purohit, Svitkina, and Kumar [PSK18] and Gollapudi and Panigrahi [GP19] study the ski rental problem, and Purohit, Svitkina, and Kumar [PSK18], Lattanzi, Lavastida, Moseley, and Vassilvitskii [LLM<sup>+</sup>20], and Mitzenmacher [Mit20] study online scheduling. Of these, the scheduling algorithm of Purohit, Svitkina, and Kumar [PSK18] is the most similar in spirit to this present work: Both algorithms are based on combining a naïve and optimistic algorithm with a robust algorithm.

Other threads of research falling under beyond worst-case online algorithms include work on combining multiple algorithms with different performance characteristics [FKL<sup>+</sup>91; BB00; MNS12; GP19], designing online algorithms with distributional assumptions (e.g., stochasticity) on the input [KP00; BS12; MGZ12], and semi-online algorithms, where the input is assumed to have a predictable offline component and an adversarial online component [KPS<sup>+</sup>19; KPS<sup>+</sup>20].

The idea of learning-augmentation has also been explored in many other algorithmic and data structural settings in recent years. These include learned indices [KBC<sup>+</sup>18], bloom filters [Mit18], frequency estimation in streams [HIK<sup>+</sup>19], and nearest neighbor search [DIR<sup>+</sup>19], among others.

Finally, advice for online algorithms has also been considered with a more complexity theoretic spirit through the study of advice complexity of online algorithms; see the survey of Boyar, Favrholt, Kudahl, Larsen, and Mikkelsen [BFK<sup>+</sup>17].

<sup>1</sup>The trade-off in  $\varepsilon$  and the additional cost is additive; thus, it does not factor into the competitive ratio.

**1.3. Recent Developments.** Recently, in work done independently of and concurrently with this paper, Antoniadis, Coester, Eliás, Polak, and Simon [ACE<sup>+</sup>20] also study a BLINDORACLE-like algorithm, which they term FOLLOWTHEPREDICTION, in the more general setting of learning-augmented metrical task systems; they also use the “combiner” of Blum and Burch [BB00] to make this algorithm robust. However, their prediction model, when specialized to online caching, is incomparable to that of Lykouris and Vassilvitskii [LV18] (which we follow).<sup>2</sup> Thus, the theoretical results proved in these two models do not imply each other.

## 2. PRELIMINARIES

**2.1. Setup and Notation.** In the online caching problem, we receive a sequence  $\sigma = (\sigma_1, \dots, \sigma_n)$  of page requests online, and our goal is to serve these requests using a cache of size  $k$  while minimizing cost. In this problem, pages must be served from the cache and can be served at no cost; however, evicting a page from the cache has unit cost.<sup>3</sup>

We will establish competitive bounds comparing the performance of two online caching algorithms  $\mathcal{A}$  and  $\mathcal{B}$ . More precisely, we will show bounds of the form

$$\text{ALG}_{\mathcal{B}}(\sigma) \leq \gamma \cdot \text{ALG}_{\mathcal{A}}(\sigma) + O(1),$$

where  $\text{ALG}_{\mathcal{A}}(\sigma)$  and  $\text{ALG}_{\mathcal{B}}(\sigma)$  are the costs of  $\mathcal{A}$  and  $\mathcal{B}$ , respectively, as measured in number of evictions made while serving a sequence  $\sigma$  of page requests. We will also compare our costs to the offline optimal algorithm OPT, whose cost  $\text{OPT}(\sigma)$  is the minimum possible cost of serving request sequence  $\sigma$ . We will omit the argument  $\sigma$  when the context is clear (i.e., just writing  $\text{ALG}_{\mathcal{A}}$  to represent  $\text{ALG}_{\mathcal{A}}(\sigma)$ ).

In our analysis, we use  $A_t$  and  $B_t$  to denote the cache states of  $\mathcal{A}$  and  $\mathcal{B}$ , respectively, just before the  $t$ -th request. Formally,  $A_t$  and  $B_t$  are subsets of  $\{1, \dots, t-1\}$  of size at most  $k$ , containing for each cached page the index at which it was last served. That is, when serving the  $t$ -th request, we remove some old request index  $t'$  from the cache and insert  $t$ . Thus, if  $t'$  is such that  $\sigma_t = \sigma_{t'}$ , this operation is free; otherwise, it has unit cost. In the sequel, we will also refer to these indices  $t$  as *page requests*.

In the learning-augmented online caching problem, the  $t$ -th page request comes with a prediction  $h_t$  for the next time page  $\sigma_t$  is requested. That is, at the time of the  $t$ -th request, our algorithm receives the pair  $(\sigma_t, h_t)$ . Let  $h = (h_1, \dots, h_n)$  be the tuple of all  $n$  predictions. To define a notion of loss, let  $y_t$  denote for each  $t$  the next time page  $t$  is actually requested, with  $y_t = n+1$  if page  $\sigma_t$  is never requested again. The  $\ell_1$  loss is then defined to be

$$\eta(\sigma, h) = \sum_t |h_t - y_t|.$$

We will omit arguments to  $\eta$  if the context is clear. Note that if  $\eta(\sigma, h) = 0$ , then the offline optimal can be obtained, as the optimal algorithm always evicts the page that is next requested furthest into the future.

In stating our bounds, the essential quantity is often  $\eta/\text{OPT}$ . To make this clear, we take  $\varepsilon = \eta/\text{OPT}$  and state our bounds in terms of  $\varepsilon$  in the sequel.

**2.1.1. Inversions.** Call a pair  $(i, j)$  of page requests an *inversion* if  $y_i < y_j$  but  $h_i \geq h_j$ . Let  $M(\sigma, h)$  denote the total number of inversions between the pair of sequences  $\sigma$  and  $h$ . And as above, we will omit the arguments to  $M$  when the context is clear.

**2.1.2. BlindOracle.** We now formally define the BLINDORACLE algorithm as follows: For each page request, if the requested page is already in the cache, do nothing. Otherwise, evict the page request  $p$  whose predicted next arrival time  $h_p$  is furthest away among all  $p \in A_t$ , with ties broken consistently (e.g., by always evicting the least recently used page with maximal  $h_p$ ).

<sup>2</sup>Namely, their algorithms expect predictions to be in a different form: They expect predictions to be cache states (i.e., the set of pages in the cache at time  $t$ ) rather than next arrival times of pages. Moreover, there exist sequences of “corresponding” inputs for each of these two models such that the predictor error approaches infinity in one model while remaining constant in the other.

<sup>3</sup>Note that this is equivalent to the “standard” version, where each cache miss has unit cost, up to a constant of  $k$ .

**2.2. Combining Online Algorithms Competitively.** In this section, we state some classical bounds on competitively combining online algorithms, due to Fiat, Karp, Luby, McGeoch, Sleator, and Young [FKL<sup>+</sup>91] and Blum and Burch [BB00]. This type of “black-box” combination was also considered by Lykouris and Vassilvitskii [LV18], but their approach has a worst constant than that of Fiat, Karp, Luby, McGeoch, Sleator, and Young [FKL<sup>+</sup>91]. We also note that results of a similar flavor are proven by Purohit, Svitkina, and Kumar [PSK18] and Mahdian, Nazerzadeh, and Saberi [MNS12], but for other online problems.

The question of combining multiple online algorithms while remaining competitive against each was first considered in the seminal paper of Fiat, Karp, Luby, McGeoch, Sleator, and Young [FKL<sup>+</sup>91]. They consider combining  $n$  online algorithms  $\mathcal{A}_1, \dots, \mathcal{A}_n$  for the online caching problem into a single algorithm  $\mathcal{A}$  such that  $\mathcal{A}$  is  $C_i$ -competitive against  $\mathcal{A}_i$  for each  $i$ . They show that such an  $\mathcal{A}$  is achievable if and only if

$$\sum_{i=1}^n \frac{1}{C_i} \leq 1.$$

We will need only the special case of  $n = 2$  and  $C_1 = C_2 = 2$ , which we state below:

**Theorem 2.1** ([FKL<sup>+</sup>91], special case). *Given any two algorithms  $\mathcal{A}$  and  $\mathcal{B}$  for the online caching problem, there exists an algorithm  $\mathcal{C}$  such that*

$$\text{ALG}_{\mathcal{C}}(\sigma) \leq 2 \min(\text{ALG}_{\mathcal{A}}(\sigma), \text{ALG}_{\mathcal{B}}(\sigma)) + O(1).$$

Moreover, if  $\mathcal{A}$  and  $\mathcal{B}$  are deterministic, so is  $\mathcal{C}$ .

Indeed, we note that that this can be done deterministically with a “follow-the-leader” approach, in which we simulate both algorithms and at each step evict any page that is not in the cache of the better performing algorithm (as measured by total number of evictions after serving the current request).

Blum and Burch [BB00] show that one can obtain a better approximation factor using a *randomized* scheme, namely multiplicative weights.<sup>4</sup> That is, at each point in time, the probability that the combined algorithm is following one of the  $n$  algorithms is given by a probability distribution over the  $n$  algorithms governed by the multiplicative weights update rule. For  $n = 2$ , their result can be stated as follows:

**Theorem 2.2** ([BB00], special case). *Given any two algorithms  $\mathcal{A}$  and  $\mathcal{B}$  for the online caching problem and any  $\varepsilon$ ,  $0 < \varepsilon < 1/4$ , there exists an algorithm  $\mathcal{C}$  such that*

$$\text{ALG}_{\mathcal{C}}(\sigma) \leq (1 + \varepsilon) \min(\text{ALG}_{\mathcal{A}}(\sigma), \text{ALG}_{\mathcal{B}}(\sigma)) + O(\varepsilon^{-1}k).$$

*Remark.* Although we do not state the versions of these results for  $n > 2$ , one could imagine that they can be useful wishes to combine multiple machine-learned predictors.

**2.3. From  $\ell_1$  Loss to Inversions.** We now state a lemma of Rohatgi [Roh20] that relates  $\ell_1$  loss to the number of inversions, letting us lower bound the  $\ell_1$  loss  $\eta(\sigma, h)$  by lower bounding the number of inversions  $M(\sigma, h)$ . Thus, instead of reasoning in terms of  $\ell_1$  loss, we will reason in terms of inversions.

**Lemma 2.3** ([Roh20]). *For any  $\sigma$  and  $h$ ,  $\eta(\sigma, h) \geq \frac{1}{2}M(\sigma, h)$ .*

With this lemma, it suffices (up to a factor of 2) to give our competitive ratio upper bounds in terms of the number of inversions  $M$ .

### 3. A FIRST ANALYSIS OF BLINDORACLE

In this section, we give a first analysis of BLINDORACLE, showing that it gets very good performance when the ratio  $\varepsilon = \eta/\text{OPT}$  is very small. In particular, our analysis shows that as  $\varepsilon \rightarrow 0$ , the competitive ratio achieved approaches 1.

Let  $\mathcal{A}$  be the offline optimal algorithm (i.e., such that  $\text{ALG}_{\mathcal{A}} = \text{OPT}$ ). Let  $\mathcal{B}$  be BLINDORACLE. Note that we can think of each of  $\text{ALG}_{\mathcal{A}}$ ,  $\text{ALG}_{\mathcal{B}}$ , and  $M$  as functions of the time  $t$ , i.e., they are the cost of  $\mathcal{A}$ , the cost

<sup>4</sup>The result of Blum and Burch [BB00] in fact holds for general metrical task systems.

of  $\mathcal{B}$ , and the number of inversions, respectively, on the prefix consisting of the first  $t - 1$  requests.<sup>5</sup> We use the  $\Delta$  operator to denote the change (in a function of  $t$ ) from time  $t$  to time  $t + 1$ . For example,  $\Delta \text{ALG}_{\mathcal{A}} = 1$  if  $\text{ALG}_{\mathcal{A}}$  evicts an element upon the  $t$ -th request.

In our analysis, we maintain a matching  $X_t$  between  $A_t$  and  $B_t$  at all times  $t$ . Call a matching *valid* if it consists only of pairs  $(a, b) \in A_t \times B_t$  such that the next arrival of  $b$  is no later than the next arrival of  $a$ . Indeed, our matching  $X_t \subseteq A_t \times B_t$  will be valid throughout the execution of the algorithm.

We now proceed with a potential function analysis, taking our potential  $\Phi$  (as a function of  $A_t, B_t$ , and  $X_t$ ) to be the number of unmatched pages in  $B_t$ . For notational simplicity, we will simply denote  $\Phi(A_t, B_t, X_t)$  by  $\Phi(t)$ . Given this setup, we show:

**Proposition 3.1.** *There exists a valid matching  $X_n$  such that*

$$\text{ALG}_{\mathcal{B}} + \Phi(n) \leq \text{OPT} + M.$$

*Proof.* We induct on the length  $n$  of the input and perform a case analysis to show that we can maintain a valid matching  $X_t$  such that at each time step, the right-hand side increases at least as much as the left-hand side, i.e.,  $\Delta \text{ALG}_{\mathcal{B}} + \Delta \Phi \leq \Delta \text{OPT} + \Delta M$ .

For our base case, note that  $A_1 = B_1$ , so we may take  $X_1$  to be the identity matching.

Now, upon a request at time  $t$ , we update  $X_t$  according to the following cases (and with the consequences listed for each case):

- (1) The requested page  $p$  is in both  $A_t$  and  $B_t$ .
  - (a) The cached pages are matched to each other.
    - Do nothing.
  - (b) Otherwise:
    - (i) Both cached pages are matched.
      - Remove the pairs  $(c, p)$  and  $(p, d)$  from  $X_t$ .
      - Add the pairs  $(p, p)$  and  $(c, d)$  to  $X_t$ .
      - As a result:
        - $\Delta \Phi = 0$ .
    - (ii) Otherwise:
      - Remove any pairs involving  $p$  from  $X_t$ . (There is at most one such pair.)
      - Add the pair  $(p, p)$  to  $X_t$ .
      - As a result:
        - $\Delta \Phi \leq 0$ .
- (2) The requested page  $p$  is in  $B_t$  only.
  - Remove any pairs involving the evicted page  $a$  from  $X_t$ . (There is at most one such pair.)
  - Remove any pairs involving the requested page  $p$  from  $X_t$ . (There is at most one such pair.)
  - Add the pair  $(p, p)$  to  $X_t$ .
  - As a result:
    - $\Delta \text{OPT} = 1$ .
    - $\Delta \Phi \leq 1$ .
- (3) The requested page  $p$  is in  $A_t$  only.
  - (a) The evicted page  $b \in B_t$  is unmatched.
    - Add the pair  $(p, p)$  to  $X_t$ . (The arriving page  $p \in A_t$  cannot be in any valid matching.)
    - As a result:
      - $\Delta \text{ALG}_{\mathcal{B}} = 1$ .
      - $\Delta \Phi = -1$ .
  - (b) The evicted page  $b \in B_t$  is matched.
    - (i)  $b$  arrives later than all unmatched pages in  $B_t$ .
      - Remove the pair  $(c, b)$  involving the evicted page  $b$  from  $X_t$ .

---

<sup>5</sup>This indexing is to be consistent with the definitions of  $A_t$  and  $B_t$ .

- Add the pair  $(c, b')$  to  $X_t$ , where  $b' \in B_t$  is any unmatched page.
  - Add the pair  $(p, p)$  to  $X_t$ . (The arriving page  $p \in A_t$  cannot be in any valid matching.)
  - As a result:
    - $\Delta \text{ALG}_{\mathcal{B}} = 1$ .
    - $\Delta \Phi = -1$ .
- (ii) There is an unmatched page  $b' \in B_t$  arriving later than  $b$ .
- Remove the pair  $(c, b)$  involving the evicted page  $b$  from  $X_t$ .
  - Add the pair  $(p, p)$  to  $X_t$ . (The arriving page  $p \in A_t$  cannot be in any valid matching.)
  - As a result:
    - $\Delta \text{ALG}_{\mathcal{B}} = 1$ .
    - $\Delta \Phi = 0$ .
    - $\Delta M = 1$ , as there is an inversion between  $b$  and  $b'$ . (Note that we do not count this inversion ever again, as  $b$  gets evicted.)
- (4) The requested page  $p$  is in neither  $A_t$  nor  $B_t$ .
- (a)  $\mathcal{A}$  evicts an unmatched page  $a \in A_t$ .
- (i)  $\mathcal{B}$  evicts an unmatched page  $b \in B_t$ .
- Add the pair  $(p, p)$  to  $X_t$ .
  - As a result:
    - $\Delta \text{OPT} = 1$ .
    - $\Delta \text{ALG}_{\mathcal{B}} = 1$ .
    - $\Delta \Phi = 1$ .
- (ii)  $\mathcal{B}$  evicts a matched page  $b \in B_t$ .
- Remove the pair  $(c, b)$  involving  $b$  from  $X_t$ .
  - Add the pair  $(p, p)$  to  $X_t$ .
  - As a result:
    - $\Delta \text{OPT} = 1$ .
    - $\Delta \text{ALG}_{\mathcal{B}} = 1$ .
- (b)  $\mathcal{A}$  evicts a matched page  $a \in A_t$ .
- (i)  $\mathcal{B}$  evicts an unmatched page  $b \in B_t$ .
- Remove the pair  $(a, d)$  involving  $a$  from  $X_t$ .
  - Add the pair  $(p, p)$  to  $X_t$ .
  - As a result:
    - $\Delta \text{OPT} = 1$ .
    - $\Delta \text{ALG}_{\mathcal{B}} = 1$ .
- (ii)  $\mathcal{B}$  evicts a matched page  $b \in B_t$ .
- Remove the pair  $(a, d)$  involving  $a$  from  $X_t$ .
  - Remove the pair  $(c, b)$  involving  $b$  from  $X_t$ .
  - Add the pair  $(p, p)$  to  $X_t$ .
  - As a result:
    - $\Delta \text{OPT} = 1$ .
    - $\Delta \text{ALG}_{\mathcal{B}} = 1$ .
    - Note that either  $b$  arrives after  $d$ , in which case we can add  $(c, d)$  to  $X_t$  and  $\Delta \Phi = 0$ , or the pair  $(b, d)$  forms an inversion, in which case  $\Delta \Phi = 1$  and  $\Delta M = 1$ . (As before, since  $b$  is getting evicted, we will not count this pair twice.)

It is not hard to verify that the change in the left-hand side of the bound is no more than the change in the right-hand side in each of the cases listed above, from which the proposition follows.  $\square$

**Proposition 3.2.** *The competitive ratio of algorithm  $\mathcal{B}$  is at most  $1 + 2\epsilon$ .*

*Proof.* Note that  $2\eta$  is bounded below by the number of inversions  $M$  of  $(\sigma, h)$  by Lemma 2.3. By Proposition 3.1,  $\text{ALG}_{\mathcal{A}} \leq \text{OPT} + M$ , so  $\text{ALG}_{\mathcal{A}}/\text{OPT} \leq 1 + M/\text{OPT} \leq 1 + 2\epsilon$ .  $\square$

## 4. A MORE CAREFUL ANALYSIS

In this section, we give an asymptotically better (in  $k$ ) bound for the performance of BLINDORACLE. A more careful analysis is needed to show an upper bound with a  $1/k$  coefficient on the ratio  $\varepsilon = \eta/\text{OPT}$ . We use the same high-level approach for the proof as before, but with a more complicated potential function. Again,  $\mathcal{A}$  is the offline optimal algorithm and  $\mathcal{B}$  is the BLINDORACLE algorithm, and also as before, we use  $\Delta$  to denote change (in functions of  $t$ ) from request  $t$  to request  $t + 1$ .

We maintain in this proof a matching  $X_t$  over pairs of page requests  $(a, b) \in A_t \times B_t$  such that  $h_a \geq h_b$  for each time step  $t$ . Our potential function  $\Phi$  will be a function of  $A_t$ ,  $B_t$ , and  $X_t$ . For notational simplicity, we will simply denote  $\Phi(A_t, B_t, X_t)$  by  $\Phi(t)$ .

Given  $A_t$ ,  $B_t$ , and  $X_t$  at time  $t$ , define  $\Phi_0(t)$  to be the number of  $b \in B_t$  that are unmatched. Define  $\Phi_1(t)$  to be the number of  $b \in B_t$  such that  $(b, b) \notin X_t$ . In other words,  $\Phi_1$  counts how many page requests in  $B_t$  are not matched to the *same* page request in  $A_t$ . Let  $z_a(t)$  be the number of pages in  $B_t$  predicted to appear no later than  $h_a$ , with tie-breaking done in a consistent manner (e.g., by the last time the page was requested). Next, define

$$\Phi_2(t) = \sum_{(a,b) \in X_t} (z_b(t) - z_a(t)) = \sum_{(a,b) \in X_t} (\varphi_a^A(t) + \varphi_b^B(t)),$$

where  $\varphi_a^A(t) = (k-1) - z_a(t)$  and  $\varphi_b^B(t) = z_b(t) - (k-1)$ . Finally, we take

$$\Phi(t) = (k-1)\Phi_0(t) + (k-1)\Phi_1(t) + \Phi_2(t)$$

as our overall potential function.

**Proposition 4.1.** *For any input  $(\sigma, h)$ , there exists a matching  $X_n \subseteq A_n \times B_n$  consisting only of pairs  $(a, b)$  satisfying  $h_a \geq h_b$  such that*

$$(k-1)\text{ALG}_{\mathcal{B}} + \Phi(n) \leq 2(k-1)\text{OPT} + 2M.$$

*Proof.* We again induct on the length  $n$  of the input, and we again perform a case analysis to show that we can maintain a matching  $X_t$  consisting only of pairs  $(a, b)$  satisfying  $h_a \geq h_b$  such that at each time step, the right-hand side increases at least as much as the left-hand side.

For our analysis, we split the serving of each page request into two phases:

- (1) **Matching.** Update  $X_t$  so that the page requests in  $A_t$  and  $B_t$  that are to be removed are unmatched. (Note that page requests are removed either because the corresponding page was requested again or because the corresponding page was evicted.)
- (2) **Updating.** Replace a page request from each of  $A_t$  and  $B_t$  with the new request and insert the new page request pair  $(t, t)$  into  $X_t$ .

We first analyze how updating affects the potential  $\Phi$ . This operation always decreases  $\Phi_0$  and  $\Phi_1$  each by 1, since we remove an unmatched pair. Next, for  $\Phi_2$ , observe that for a matched pair  $(a, b)$ , the difference  $z_b - z_a$  increases on the  $t$ -th request only if there exists a  $p \in B_t$  such that  $\sigma_p = \sigma_t$  and  $h_b < h_p \leq h_a$ . In this case, the pair  $(p, b)$  also forms an inversion. Any inversion  $(p, b)$  is counted at most once this way because  $p$  is evicted from  $b$ . Thus, we have  $\Delta\Phi_2 \leq \Delta M$ .

We now analyze the matching phase with a case analysis:

- (1) The requested page is in both  $A_t$  and  $B_t$ .
  - The previous page requests for  $\sigma_t$  in  $A_t$  and  $B_t$  are matched to each other, so we can just unmatch them.
  - As a result:
    - $\Delta\Phi_0 = 1$ .
    - $\Delta\Phi_1 = 1$ .
    - $\Delta\Phi_2 = 0$ , since the pages were matched to each other.
- (2) The requested page is in  $A_t$  only.

- (a) The previous request  $p \in A_t$  for the requested page is matched as  $(p, d)$  and the page request  $b \in B_t$  evicted by  $\mathcal{B}$  is matched as  $(c, b)$ .
- Unmatch  $(p, d)$  and  $(c, b)$  and then match  $(c, d)$ . The latter is okay since  $h_c \geq h_b \geq h_d$ . Note that  $p \neq d$ .
  - As a result:
    - $\Delta \text{ALG}_{\mathcal{B}} = 1$ .
    - $\Delta \Phi_0 = 1$ .
    - $\Delta \Phi_1 \leq 1$ .
    - $\Delta \Phi_2 = z_p - (k - 1)$ , since  $\varphi_b^B = 0$  and  $\varphi_p^A = (k - 1) - z_p$ .
    - $\Delta M \geq z_p$ , since the arrival of  $\sigma_p$  also generates  $z_p$  inversions of the form  $(p, b')$  for all  $b' \in B_t$  such that  $h_p \geq h_{b'}$ .
- (b) The previous request  $p \in A_t$  for the requested page is matched as  $(p, d)$  and the page request  $b \in B_t$  evicted by  $\mathcal{B}$  is unmatched.
- Unmatch  $(p, d)$ . Note that  $p \neq d$ .
  - As a result:
    - $\Delta \text{ALG}_{\mathcal{B}} = 1$ .
    - $\Delta \Phi_0 = 1$ .
    - $\Delta \Phi_1 = 0$ .
    - $\Delta \Phi_2 \leq z_p$ , since  $\varphi_p^A = (k - 1) - z_p$  and  $\varphi_d^B = z_d - (k - 1) \geq -(k - 1)$ .
    - $\Delta M \geq z_p$ , since the arrival of  $\sigma_p$  also generates  $z_p$  inversions of the form  $(p, b')$  for all  $b' \in B_t$  such that  $h_p \geq h_{b'}$ .
- (c) The previous request  $p \in A_t$  for the requested page is unmatched and the page request  $b \in B_t$  evicted by  $\mathcal{B}$  is matched as  $(c, b)$ .
- Unmatch  $(c, b)$  and match  $(c, d)$  for an arbitrary unmatched  $d \in B_t \setminus \{b\}$ . Doing so is okay because  $h_c \geq h_b \geq h_d$ .
  - As a result:
    - $\Delta \text{ALG}_{\mathcal{B}} = 1$ .
    - $\Delta \Phi_0 = 0$ .
    - $\Delta \Phi_1 \leq 1$ .
    - $\Delta \Phi_2 \leq 0$ , since  $\varphi_b^B = 0$  and  $-\varphi_d^B = (k - 1) - z_d \geq 0$ .
- (d) The previous request  $p \in A_t$  for the requested page is unmatched and the page request  $b \in B_t$  evicted by  $\mathcal{B}$  is unmatched.
- Do nothing.
  - As a result:
    - $\Delta \text{ALG}_{\mathcal{B}} = 1$ .
- (3) The requested page is in  $B_t$  only.
- (a) The previous request  $p \in B_t$  for the requested page is matched as  $(c, p)$  and the page request  $a \in A_t$  evicted by  $\mathcal{A}$  is matched as  $(a, d)$ .
- Unmatch  $(c, p)$  and  $(a, d)$ . Note that  $c \neq p$ .
  - As a result:
    - $\Delta \text{OPT} = 1$ .
    - $\Delta \Phi_0 = 2$ .
    - If  $a = d$ , then  $\Delta \Phi_1 = 1$  and  $\varphi_a^A + \varphi_d^B = 0$ ; otherwise,  $a \neq d$ , in which case  $\Delta \Phi_1 = 0$  and  $\varphi_a^A + \varphi_d^B \geq -(k - 1)$ .
    - Moreover,  $\varphi_c^A + \varphi_p^B \geq -(k - 1)$ .
- (b) The previous request  $p \in B_t$  for the requested page is matched as  $(c, p)$  and the page request  $a \in A_t$  evicted by  $\mathcal{A}$  is unmatched.
- Unmatch  $(c, p)$ . Note that  $c \neq p$ .
  - As a result:
    - $\Delta \text{OPT} = 1$ .

- $\Delta\Phi_0 = 1$ .
  - $\Delta\Phi_1 = 0$ .
  - $\Delta\Phi_2 \leq k - 1$ , since  $\varphi_c^A + \varphi_p^B \geq -(k - 1)$ .
- (c) The previous request  $p \in B_t$  for the requested page is unmatched and the page request  $a \in A_t$  evicted by  $\mathcal{A}$  is matched as  $(a, d)$ .
- Unmatch  $(a, d)$ .
  - As a result:
    - $\Delta\text{OPT} = 1$ .
    - $\Delta\Phi_0 = 1$ .
    - If  $a = d$ , then  $\Delta\Phi_1 = 1$  and  $\varphi_a^A + \varphi_d^B = 0$ ; otherwise,  $a \neq d$ , in which case  $\Delta\Phi_1 = 0$  and  $\varphi_a^A + \varphi_d^B \geq -(k - 1)$ .
- (d) The previous request  $p \in B_t$  for the requested page is unmatched and the page request  $a \in A_t$  evicted by  $\mathcal{A}$  is unmatched.
- Do nothing.
  - As a result:
    - $\Delta\text{OPT} = 1$ .
- (4) The requested page is in neither  $A_t$  nor  $B_t$ .
- (a) The previous request  $a \in A_t$  evicted by  $\mathcal{A}$  is matched as  $(a, d)$  and the page request  $b \in B_t$  evicted by  $\mathcal{B}$  is matched as  $(c, b)$ .
- Unmatch  $(a, d)$  and  $(c, b)$  and then match  $(c, d)$ . The latter is okay because  $h_c \geq h_b \geq h_d$ .
  - As a result:
    - $\Delta\text{ALG}_{\mathcal{B}} = 1$ .
    - $\Delta\text{OPT} = 1$ .
    - $\Delta\Phi_0 = 1$ .
    - $\Delta\Phi_1 \leq 2$ .
    - $\Delta\Phi_2 \leq 0$ , since  $\varphi_a^A \geq 0$  and  $\varphi_b^B = 0$ .
- (b) The previous request  $a \in A_t$  evicted by  $\mathcal{A}$  is matched as  $(a, d)$  and the page request  $b \in B_t$  evicted by  $\mathcal{B}$  is unmatched.
- Unmatch  $(a, d)$ .
  - As a result:
    - $\Delta\text{ALG}_{\mathcal{B}} = 1$ .
    - $\Delta\text{OPT} = 1$ .
    - $\Delta\Phi_0 = 1$ .
    - If  $a = d$ , then  $\Delta\Phi_1 = 1$  and  $\varphi_a^A + \varphi_d^B = 0$ ; otherwise,  $a \neq d$ , in which case  $\Delta\Phi_1 = 0$  and  $\varphi_a^A + \varphi_d^B \geq -(k - 1)$ .
- (c) The previous request  $a \in A_t$  evicted by  $\mathcal{A}$  is unmatched and the page request  $b \in B_t$  evicted by  $\mathcal{B}$  is matched as  $(c, b)$ .
- Unmatch  $(c, b)$  and match  $(c, d)$ .
  - As a result:
    - $\Delta\text{ALG}_{\mathcal{B}} = 1$ .
    - $\Delta\text{OPT} = 1$ .
    - $\Delta\Phi_0 = 0$ .
    - $\Delta\Phi_1 \leq 1$ .
    - $\Delta\Phi_2 \leq 0$ , since  $\varphi_b^B = 0$  and  $-\varphi_d^B = (k - 1) - z_d \geq 0$ .
- (d) The previous request  $a \in A_t$  evicted by  $\mathcal{A}$  is unmatched and the page request  $b \in B_t$  evicted by  $\mathcal{B}$  is unmatched.
- Do nothing.
  - As a result:
    - $\Delta\text{ALG}_{\mathcal{B}} = 1$ .
    - $\Delta\text{OPT} = 1$ .

Since  $\Phi_1$  and  $\Phi_2$  each decrease by 1 in the updating phase, we have  $2(k - 1)$  in extra potential that we can use to pay for costs in the matching phase. Indeed, one can verify that this is sufficient for all of the cases described above—the tight cases are 2(a), 2(b), 2(c), 3(a), and 4(a). Thus, the proposition follows.  $\square$

**Proposition 4.2.** *The competitive ratio of algorithm  $\mathcal{B}$  is at most  $2 + 4\epsilon/(k - 1)$ .*

*Proof.* Compose Proposition 4.1 with Lemma 2.3.  $\square$

*Remark.* This analysis of BLINDORACLE is tight in the constant term—we can make  $\epsilon/(k - 1)$  arbitrarily small while having a competitive ratio of 2. (However, for very small  $\epsilon$ , the bound of the previous section is better—in particular, if  $\epsilon < \frac{1}{2} + \frac{1}{k-3}$ .)

## 5. PROOFS OF UPPER BOUNDS

We are now ready to prove the results stated in Section 1:

**Theorem 1.1.** *For learning-augmented online caching, BLINDORACLE obtains a competitive ratio of*

$$\min\left(1 + 2\frac{\eta}{\text{OPT}}, 2 + \frac{4}{k-1}\frac{\eta}{\text{OPT}}\right),$$

where  $\eta$  is the  $\ell_1$  loss incurred by the predictor and OPT is the offline optimal cost. (For precise definitions, see Section 2.1.)

*Proof.* From the analysis of the previous two sections, the desired bound immediately follows from taking the minimum of the bounds in Propositions 3.2 and 4.2.  $\square$

**Corollary 1.2.** *There exists a deterministic algorithm for learning-augmented online caching that achieves a competitive ratio of*

$$2 \min\left(\min\left(1 + 2\frac{\eta}{\text{OPT}}, 2 + \frac{4}{k-1}\frac{\eta}{\text{OPT}}\right), k\right).$$

*Proof.* Combine BLINDORACLE with LRU using the “combiner” from Theorem 2.1, with the performance of BLINDORACLE being bounded by Theorem 1.1.  $\square$

**Corollary 1.3.** *There exists a randomized algorithm for learning-augmented online caching that achieves a competitive ratio of*

$$(1 + \epsilon) \min\left(\min\left(1 + 2\frac{\eta}{\text{OPT}}, 2 + \frac{4}{k-1}\frac{\eta}{\text{OPT}}\right), H_k\right)$$

for any  $\epsilon \in (0, 1/4)$ .<sup>6</sup> (Here,  $H_k = 1 + \frac{1}{2} + \frac{1}{3} + \dots + \frac{1}{k} = \ln(k) + O(1)$  is the  $k$ -th harmonic number.)

*Proof.* Like in the proof above, combine BLINDORACLE with algorithm EQUITABLE of Achlioptas, Chrobak, and Noga [ACN00]<sup>7</sup>, this time using the “combiner” from Theorem 2.2.  $\square$

## 6. DETERMINISTIC LOWER BOUND

We now show that combining BLINDORACLE with LRU gets an optimal competitive ratio bound (in terms of  $\eta$ , OPT, and  $k$ ) among all deterministic algorithms for learning-augmented online caching by proving Theorem 1.4:

**Theorem 1.4.** *The competitive ratio bound for any deterministic learning-augmented online caching algorithm must be at least*

$$1 + \Omega\left(\min\left(\frac{1}{k}\frac{\eta}{\text{OPT}}, k\right)\right).$$

<sup>6</sup>The trade-off in  $\epsilon$  and the additional cost is additive; thus, it does not factor into the competitive ratio.

<sup>7</sup>We use EQUITABLE because it achieves the optimal worst-case competitive ratio of  $H_k$  for online caching; MARKER has a competitive ratio of  $2H_k - 1$  [ACN00].

*Proof.* Let  $\mathcal{A}$  be any deterministic algorithm for learning-augmented online caching.

We show there exists a family of inputs  $(\sigma, h)$  with  $\varepsilon/k$  ranging from 0 to  $k$  and  $\text{OPT}$  arbitrarily large such  $\text{ALG}_{\mathcal{A}} \geq \text{OPT} + C\eta/k$ , for some constant  $C > 0$ . That is, for  $\varepsilon/k$  ranging from 0 to  $k$ , we will show that we can make this inequality hold as  $\text{OPT} \rightarrow \infty$  with  $\varepsilon/k$  fixed. Dividing through by  $\text{OPT}$  then implies the theorem.

We now construct such inputs  $(\sigma, h)$ . First, fix  $j < k$ . Let  $P_1, \dots, P_k, Q_0$  be  $k + 1$  distinct pages. We make the following sequence of requests, which we call a *phase*:

- (1) Repeat  $k$  times the following:
  - (a) Make requests to  $P_1, \dots, P_k$  in order, predicting each page to next appear  $k$  requests from now *except* during the last iteration, where we predict each page to next appear  $k + j + 1$  requests from now.
- (2) Make a request to  $Q_0$  and predict that it will next appear  $k^2 + j + 1$  pages from now.
- (3) For  $i = 1, \dots, j$ :
  - (a) Request the page evicted by in  $\mathcal{A}$  during the previous request, if it exists. Otherwise, request an arbitrary page. For each page, provide the same prediction as the last time this page was requested.

We repeat the above as many times as needed.

In a single phase, observe that  $\text{OPT}$  makes at most two evictions—once to evict  $Q_0$  and once upon the arrival of  $Q_0$ . On the other hand, I claim  $\mathcal{A}$  makes at least  $j + 1$  evictions. First, if the cache of  $\mathcal{A}$  after (1) does not consist of  $P_1, \dots, P_k$ , then  $\mathcal{A}$  must have incurred cost at least  $k \geq j + 1$  during (1). Thus, we may assume that  $\mathcal{A}$ 's cache consists of  $P_1, \dots, P_k$  after (1). If so,  $\mathcal{A}$  has to evict a page for each of the remaining  $j + 1$  requests in the phase, as the arrival of  $Q_0$  forces an eviction and by induction, each arrival of (3) forces an eviction. Finally, observe that all the predictions are accurate except those for pages arriving in (3), in which case they are off by at most  $k + j + 1 \leq 2k$ . Thus, over a single phase,  $\eta \leq 2jk$ . Putting all of these observations together, we get  $\text{ALG}_{\mathcal{A}} \geq j + 1 \geq \text{OPT} + j - 1$ , with  $j - 1 = \Omega(\eta/k)$ .

To make  $\text{OPT}$  arbitrarily large, note that we can simply repeat the above phase multiple times in sequence; the same analysis holds, with all the values scaling linearly. Hence we have  $\text{ALG}_{\mathcal{A}} \geq \text{OPT} + \Omega(\eta/k)$  for arbitrarily large  $\text{OPT}$  over the desired range of  $\varepsilon/k$ , so the theorem follows.  $\square$

#### ACKNOWLEDGMENTS

I would like to thank Jelani Nelson for advising this project and Bailey Flanigan for providing many helpful references.

#### REFERENCES

- [ACE<sup>+</sup>20] Antonios Antoniadis, Christian Coester, Marek Eliás, Adam Polak, and Bertrand Simon. Online metric algorithms with untrusted predictions. *CoRR*, abs/2003.02144, 2020.
- [ACN00] Dimitris Achlioptas, Marek Chrobak, and John Noga. Competitive analysis of randomized paging algorithms. *Theor. Comput. Sci.*, 234(1-2):203–218, 2000.
- [BB00] Avrim Blum and Carl Burch. On-line learning and the metrical task system problem. *Mach. Learn.*, 39(1):35–58, 2000.
- [BBN12a] Nikhil Bansal, Niv Buchbinder, and Joseph Naor. A primal-dual randomized algorithm for weighted paging. *J. ACM*, 59(4):19:1–19:24, 2012.
- [BBN12b] Nikhil Bansal, Niv Buchbinder, and Joseph Naor. Randomized competitive algorithms for generalized caching. *SIAM J. Comput.*, 41(2):391–414, 2012.
- [BE98] Allan Borodin and Ran El-Yaniv. *Online computation and competitive analysis*. Cambridge University Press, 1998.
- [BFK<sup>+</sup>17] Joan Boyar, Lene M. Favrholdt, Christian Kudahl, Kim S. Larsen, and Jesper W. Mikkelsen. Online algorithms with advice: A survey. *ACM Comput. Surv.*, 50(2):19:1–19:34, 2017.

- [BS12] Sébastien Bubeck and Aleksandrs Slivkins. The best of both worlds: stochastic and adversarial bandits. In *COLT 2012 - The 25th Annual Conference on Learning Theory, June 25-27, 2012, Edinburgh, Scotland*, pages 42.1–42.23, 2012.
- [DIR<sup>+</sup>19] Yihe Dong, Piotr Indyk, Ilya P. Razenshteyn, and Tal Wagner. Learning sublinear-time indexing for nearest neighbor search. *CoRR*, abs/1901.08544, 2019.
- [FKL<sup>+</sup>91] Amos Fiat, Richard M. Karp, Michael Luby, Lyle A. McGeoch, Daniel Dominic Sleator, and Neal E. Young. Competitive paging algorithms. *J. Algorithms*, 12(4):685–699, 1991.
- [GP19] Sreenivas Gollapudi and Debmalya Panigrahi. Online algorithms for rent-or-buy with expert advice. In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, pages 2319–2327, 2019.
- [HIK<sup>+</sup>19] Chen-Yu Hsu, Piotr Indyk, Dina Katabi, and Ali Vakilian. Learning-based frequency estimation algorithms. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*, 2019.
- [KBC<sup>+</sup>18] Tim Kraska, Alex Beutel, Ed H. Chi, Jeffrey Dean, and Neoklis Polyzotis. The case for learned index structures. In *Proceedings of the 2018 International Conference on Management of Data, SIGMOD Conference 2018, Houston, TX, USA, June 10-15, 2018*, pages 489–504, 2018.
- [KP00] Elias Koutsoupias and Christos H. Papadimitriou. Beyond competitive analysis. *SIAM J. Comput.*, 30(1):300–317, 2000.
- [KPS<sup>+</sup>19] Ravi Kumar, Manish Purohit, Aaron Schild, Zoya Svitkina, and Erik Vee. Semi-online bipartite matching. In *10th Innovations in Theoretical Computer Science Conference, ITCS 2019, January 10-12, 2019, San Diego, California, USA*, 50:1–50:20, 2019.
- [KPS<sup>+</sup>20] Ravi Kumar, Manish Purohit, Zoya Svitkina, and Erik Vee. Interleaved caching with access graphs. In *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, January 5-8, 2020*, pages 1846–1858, 2020.
- [LLM<sup>+</sup>20] Silvio Lattanzi, Thomas Lavastida, Benjamin Moseley, and Sergei Vassilvitskii. Online scheduling via learned weights. In *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, January 5-8, 2020*, pages 1859–1877, 2020.
- [LV18] Thodoris Lykouris and Sergei Vassilvitskii. Competitive caching with machine learned advice. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, pages 3302–3311, 2018.
- [MGZ12] Vahab S. Mirrokni, Shayan Oveis Gharan, and Morteza Zadimoghaddam. Simultaneous approximations for adversarial and stochastic online budgeted allocation. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2012, Kyoto, Japan, January 17-19, 2012*, pages 1690–1701, 2012.
- [Mit18] Michael Mitzenmacher. A model for learned Bloom filters and optimizing by sandwiching. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, 3-8 December 2018, Montréal, Canada*, pages 462–471, 2018.
- [Mit20] Michael Mitzenmacher. Scheduling with predictions and the price of misprediction. In *11th Innovations in Theoretical Computer Science Conference, ITCS 2020, January 12-14, 2020, Seattle, Washington, USA*, 14:1–14:18, 2020.
- [MNS12] Mohammad Mahdian, Hamid Nazerzadeh, and Amin Saberi. Online optimization with uncertain information. *ACM Trans. Algorithms*, 8(1):2:1–2:29, 2012.
- [MV17] Andres Muñoz Medina and Sergei Vassilvitskii. Revenue optimization with approximate bid predictions. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, pages 1858–1866, 2017.

- [PSK18] Manish Purohit, Zoya Svitkina, and Ravi Kumar. Improving online algorithms via ML predictions. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, 3-8 December 2018, Montréal, Canada*, pages 9684–9693, 2018.
- [Roh20] Dhruv Rohatgi. Near-optimal bounds for online caching with machine learned advice. In *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, January 5-8, 2020*, pages 1834–1845, 2020.
- [Rou19] Tim Roughgarden. Beyond worst-case analysis. *Commun. ACM*, 62(3):88–96, 2019.
- [ST85] Daniel Dominic Sleator and Robert Endre Tarjan. Amortized efficiency of list update and paging rules. *Commun. ACM*, 28(2):202–208, 1985.