

Online metric algorithms with untrusted predictions*

Antonios Antoniadis[†]
University of Twente

Christian Coester[‡]
University of Oxford

Marek Eliáš[§]
Bocconi University

Adam Polak[¶]
Max Planck Institute for Informatics
Jagiellonian University

Bertrand Simon^{||}
IN2P3 Computing Center, CNRS

Abstract

Machine-learned predictors, although achieving very good results for inputs resembling training data, cannot possibly provide perfect predictions in all situations. Still, decision-making systems that are based on such predictors need not only benefit from good predictions, but should also achieve a decent performance when the predictions are inadequate.

In this paper, we propose a prediction setup for arbitrary *metrical task systems (MTS)* (e.g., *caching*, *k-server* and *convex body chasing*) and *online matching on the line*. We utilize results from the theory of online algorithms to show how to make the setup robust. Specifically for caching, we present an algorithm whose performance, as a function of the prediction error, is exponentially better than what is achievable for general MTS. Finally, we present an empirical evaluation of our methods on real world datasets, which suggests practicality.

1 Introduction

Metrical task systems (MTS), introduced by Borodin et al. (1992), are a rich class containing several fundamental problems in online optimization as special cases, including *caching*, *k-server*, *convex body chasing*, and *convex function chasing*. MTS are capable of modeling many problems arising in computing and production systems (Sleator and Tarjan, 1985; Manasse et al., 1990), movements of service vehicles (Dehghani et al., 2017; Coester and Koutsoupias, 2019), power management of embedded systems as well as data centers (Irani et al., 2003; Lin et al., 2013), and are also related to the *experts* problem in online learning, see (Daniely and Mansour, 2019; Blum and Burch, 2000).

Initially, we are given a metric space M of *states*, which can be interpreted for example as actions, investment strategies, or configurations of some production machine. We start at a predefined initial state x_0 . At each time $t = 1, 2, \dots$, we are presented with a *cost function* $\ell_t: M \rightarrow \mathbb{R}^+ \cup \{0, +\infty\}$ and our task is to decide either to stay at x_{t-1} and pay the cost $\ell_t(x_{t-1})$, or to move to some other (possibly cheaper) state x_t and pay $\text{dist}(x_{t-1}, x_t) + \ell_t(x_t)$, where $\text{dist}(x_{t-1}, x_t)$ is the cost of the transition between states x_{t-1} and x_t . The objective is to minimize the overall cost incurred over time.

Given that MTS is an online problem, one needs to make each decision without any information about the future cost functions. This makes the problem substantially difficult, as supported by

*The authors would like to thank IGAFIT for the organization of the AlgPiE workshop which made this project possible.

[†]Supported by DFG Grant AN 1262/1-1.

[‡]Supported by NWO VICI grant 639.023.812 of Nikhil Bansal.

[§]Supported by ERC Starting Grant 759471 of Michael Kapralov.

[¶]Supported by National Science Center of Poland grant 2017/27/N/ST6/01334 and by Swiss National Science Foundation within the project *Lattice Algorithms and Integer Programming* (185030).

^{||}Supported by the DFG Project Number 146371743 – TRR 89 Invasive Computing.

strong lower bounds for general MTS (Borodin et al., 1992) as well as for many special MTS problems (see e.g. Karloff et al., 1994; Fiat et al., 1998). For the recent work on MTS, see Bubeck et al. (2019); Coester and Lee (2019); Bubeck and Rabani (2020a).

In this paper, we study how to utilize predictors (possibly based on machine learning) in order to decrease the uncertainty about the future and achieve a better performance for MTS. We propose a natural prediction setup for MTS and show how to develop algorithms in this setup with the following properties of *consistency* (i), *smoothness* (ii), and *robustness* (iii).

- (i) Their performance with perfect predictions is close to optimal.
- (ii) With decreasing accuracy of the predictions, their performance deteriorates smoothly as a function of the prediction error.
- (iii) When given poor predictions, their performance is comparable to that of the best online algorithm which does not use predictions.

Caching and weighted caching problems, which are special cases of MTS, have already been studied in this context of utilizing predictors Lykouris and Vassilvitskii (2018); Rohatgi (2020); Wei (2020); Jiang et al. (2020); Bansal et al. (2022). However, the corresponding prediction setups do not seem applicable to general MTS. For example, algorithms by Lykouris and Vassilvitskii (2018) and Rohatgi (2020) provide similar guarantees by using predictions of the next reoccurrence time of the current page in the input sequence. However, as we show in this paper, such predictions are not useful for more general MTS: even for weighted caching, they do not help to improve upon the bounds achievable without predictions unless additional assumptions are made (see Bansal et al. (2022) for an example of such an assumption).

We propose a prediction setup based on *action predictions* where, at each time step, the predictor tries to predict the action that an offline algorithm would have taken. We can view these predictions as recommendations of what our algorithm should do. We show that using this prediction setup, we can achieve consistency, smoothness, and robustness for any MTS. For the (unweighted) caching problem, we develop an algorithm that obtains a better dependency on the prediction error than our general result, and whose performance in empirical tests is either better or comparable to the algorithms by Lykouris and Vassilvitskii (2018) and Rohatgi (2020). This demonstrates the flexibility of our setup. We would like to stress that specifically for the caching problem, the action predictions can be obtained by simply converting the reoccurrence time predictions used in (Lykouris and Vassilvitskii, 2018; Rohatgi, 2020; Wei, 2020), a feature that we use in order to compare our results to those previous algorithms. Nevertheless our prediction setup is applicable to the much broader context of MTS. We demonstrate this and suggest practicability of our algorithms also for MTS other than caching by providing experimental results for the *ice cream* problem (Chrobak and Larmore, 1998), a simple example of an MTS. Finally, we extend our theoretical result beyond MTS to *online matching on the line*.

Action Predictions for MTS. At each time t , the predictor proposes an *action*, i.e., a state p_t in the metric space M . We define the *prediction error* with respect to some offline algorithm OFF as

$$\eta = \sum_{t=1}^T \eta_t; \quad \eta_t = \text{dist}(p_t, o_t), \quad (1)$$

where o_t denotes the state of OFF at time t and T denotes the length of the input sequence.

The predictions could be, for instance, the output of a machine-learned model or actions of a heuristic which tends to produce good solutions in practice, but possibly without a theoretical guarantee. The offline algorithm OFF can be an optimal one, but also other options are plausible. For example, if the typical instances are composed of subpatterns known from the past and for which good solutions are known, then OFF could also be a near-optimal algorithm which composes its output from the partial solutions to the subpatterns. The task of the predictor in this case is to anticipate which subpattern is going to follow and provide the precomputed solution to that

subpattern. In the case of the caching problem, as mentioned above and explained in Section 1.3, we can actually convert the reoccurrence predictions (Lykouris and Vassilvitskii, 2018; Rohatgi, 2020; Wei, 2020) into action predictions.

Note that, even if the prediction error with respect to OFF is low, the cost of the solution composed from the predictions p_1, \dots, p_T can be much higher than the cost incurred by OFF, since $\ell_t(p_t)$ can be much larger than $\ell_t(o_t)$ even if $\text{dist}(p_t, o_t)$ is small. However, we can design algorithms which use such predictions and achieve a good performance whenever the predictions have small error with respect to any low-cost offline algorithm. We aim at expressing the performance of the prediction-based algorithms as a function of η/OFF , where (abusing notation) OFF denotes the cost of the offline algorithm. This is to avoid scaling issues: if the offline algorithm incurs movement cost 1000, predictions with total error $\eta = 1$ give us a rather precise estimate of its state, unlike when $\text{OFF} = 0.1$.

Caching Problem. In the caching problem we have a two-level computer memory, out of which the fast one (cache) can only store k pages. We need to answer a sequence of requests to pages. Such a request requires no action and incurs no cost if the page is already in the cache, but otherwise a *page fault* occurs and we have to add the page and evict some other page at a cost of 1. Caching can be seen as an MTS whose states are cache configurations¹. Therefore, also the predictions are cache configurations in our setup, but as we discuss in Section 1.3 they can be encoded much more succinctly than specifying the full cache content in each time step. The error η_t describes in this case the number of pages on which the predicted cache and the cache of OFF differ at time t .

1.1 Our results

We prove two general theorems providing robustness and consistency guarantees for any MTS.

Theorem 1. *Let A be a deterministic α -competitive online algorithm for a problem P belonging to MTS. Given action predictions for P , there is a deterministic algorithm achieving competitive ratio*

$$9 \cdot \min\{\alpha, 1 + 4\eta/\text{OFF}\}$$

against any offline algorithm OFF, where η is the prediction error with respect to OFF.

Roughly speaking, the competitive ratio (formally defined in Section 2) is the worst case ratio between the cost of two algorithms. If OFF is an optimal algorithm, then the expression in the theorem is the overall competitive ratio of the prediction-based algorithm.

Theorem 2. *Let A be a randomized α -competitive online algorithm for an MTS P with metric space diameter D . For any $\epsilon \leq 1/4$, given action predictions for P there is a randomized algorithm achieving cost*

$$(1 + \epsilon) \cdot \min\{\alpha, 1 + 4\eta/\text{OFF}\} \cdot \text{OFF} + O(D/\epsilon),$$

where η is the prediction error with respect to an offline algorithm OFF. Thus, if OFF is (near-)optimal and $\eta \ll \text{OFF}$, the competitive ratio is close to $1 + \epsilon$.

We note that the proofs of these theorems are based on the powerful results by Fiat et al. (1994) and Blum and Burch (2000). In Theorem 20, we show that the dependence on η/OFF in the preceding theorems is tight up to constant factors for some MTS instance.

For some specific MTS, however, the dependence on η/OFF can be improved, as shown in Section 3, where we present a new algorithm for caching whose competitive ratio has a logarithmic dependence on η/OFF . One of the main characteristics of our algorithm, which we call TRUST&DOUBT, compared to previous approaches, is that it is able to *gradually* adapt its level of trust in the predictor throughout the instance. Showing that our general prediction setup can be used to design such efficient algorithms for caching is the most involved result of our paper, so the following result is proved before Theorems 1 and 2.

¹A cache configuration is the set of pages in cache.

Theorem 3. *For caching with action predictions, there is a randomized algorithm with competitive ratio $\min\{f(\frac{\eta}{\text{OFF}}), g(k)\}$ against any algorithm OFF, where $f(\frac{\eta}{\text{OFF}}) \leq O(\log \frac{\eta}{\text{OFF}})$ is the smoothness with prediction error η and $g(k) \leq O(\log(k))$ is the robustness with cache size k .*

We do not attempt to optimize constant factors in the proof of Theorem 3, but we remark that f can be chosen such that $f(0) = 1 + \epsilon$, for arbitrary $\epsilon > 0$. The reason is that our algorithm in the proof of Theorem 3 can be used as algorithm A in Theorem 2.

Although we designed our prediction setup with MTS in mind, it can also be applied to problems beyond MTS. We demonstrate this in Section 5 by employing our techniques to provide an algorithm of similar flavor for *online matching on the line*, a problem not known to be an MTS.

Theorem 4. *For online matching on the line with action predictions, there is a deterministic algorithm with competitive ratio $O(\min\{\log n, 1 + \eta/\text{OFF}\})$, where η is the prediction error with respect to some offline algorithm OFF.*

We also show that Theorem 4 can be generalized to give a $O(\min\{2n - 1, \eta/\text{OFF}\})$ -competitive algorithm for *online metric bipartite matching*.

In Section C, we show that the reoccurrence time predictions introduced by Lykouris and Vassilvitskii (2018) for caching do not help for more general MTS.

Theorem 5. *The competitive ratio of any algorithm for weighted caching even if provided with precise reoccurrence time predictions is $\Omega(\log k)$.*

Note that there are $O(\log k)$ -competitive online algorithms for weighted caching which do not use any predictions (see Bansal et al., 2012). This motivates the need for a different prediction setup as introduced in this paper. This lower bound result has been obtained independently by Jiang et al. (2020) who also proved a lower bound of $\Omega(k)$ for deterministic algorithms with precise reoccurrence time predictions. However, for instances with only ℓ weight classes, Bansal et al. (2022) showed that perfect reoccurrence time predictions allow achieving a competitive ratio of $\Theta(\log \ell)$.

We round up by presenting an extensive experimental evaluation of our results that suggests practicality. We test the performance of our algorithms on public data with previously used models. With respect to caching, our algorithms outperform all previous approaches in most settings (and are always at least comparable). A very interesting use of our setup is that it allows us to employ any other online algorithm as a predictor for our algorithm. For instance, when using the Least Recently Used (LRU) algorithm – which is considered the gold standard in practice – as a predictor for our algorithm, our experiments suggest that we achieve the same practical performance as LRU, but with an exponential improvement in the theoretical worst-case guarantee ($O(\log k)$ instead of k). Finally we applied our general algorithms to a simple MTS called the ice cream problem and were able to obtain results that also suggest practicality of our setup beyond caching.

1.2 Related work

Our work is part of a larger and recent movement to prove rigorous performance guarantees for algorithms based on machine learning. The first main results have been established on both classical (see Kraska et al., 2018; Khalil et al., 2017) and online problems: Lykouris and Vassilvitskii (2018) and Rohatgi (2020) on caching, Lattanzi et al. (2020) on restricted assignment scheduling, Purohit et al. (2018) on ski rental and non-clairvoyant scheduling, Gollapudi and Panigrahi (2019) on ski rental with multiple predictors, Mitzenmacher (2020) on scheduling/queuing, and Medina and Vassilvitskii (2017) on revenue optimization.

Most of the online results are analyzed by means of *consistency* (competitive ratio in the case of perfect predictions) and *robustness* (worst-case competitive-ratio regardless of prediction quality), which was first defined in this context by Purohit et al. (2018), while Mitzenmacher (2020) uses a different measure called *price of misprediction*. It should be noted that the exact definitions of consistency and robustness are slightly inconsistent between different works in the literature, making it often difficult to directly compare results.

Results on Caching. Among the closest results to our work are the ones by Lykouris and Vassilvitskii (2018) and Rohatgi (2020), who study the caching problem (a special case of MTS) with machine learned predictions. Lykouris and Vassilvitskii (2018) introduced the following prediction setup for caching: whenever a page is requested, the algorithm receives a prediction of the time when the same page will be requested again. The prediction error is defined as the ℓ_1 -distance between the predictions and the truth, i.e., the sum – over all requests – of the absolute difference between the predicted and the real reoccurrence time of the same request. For this prediction setup, they adapted the classic Marker algorithm in order to achieve, up to constant factors, the best robustness and consistency possible. In particular, they achieved a competitive ratio of $O(1 + \min\{\sqrt{\eta/\text{OPT}}, \log k\})$ and their algorithm was shown to perform well in experiments. Later, Rohatgi (2020) achieved a better dependency on the prediction error: $O(1 + \min\{\frac{\log k}{k} \frac{\eta}{\text{OPT}}, \log k\})$. He also provides a close lower bound.

Following the original announcement of our work, we learned about further developments by Wei (2020) and Jiang et al. (2020). Wei (2020) further refined the aforementioned results for caching with reoccurrence time predictions. The paper by Jiang et al. (2020) proposes an algorithm for weighted caching in a very strong prediction setup, where the predictor reports at each time step the reoccurrence time of the currently requested page as well as *all* page requests up to that time. Jiang et al. (2020) provide a collection of lower bounds for weaker predictors (including an independent proof of Theorem 5), justifying the need for such a strong predictor. In a followup work, Bansal et al. (2022) showed, though, that the reoccurrence time predictions² are still useful for weighted caching when the number ℓ of weight classes is small, allowing to achieve a competitive ratio of $\Theta(\log \ell)$ for good predictions.

We stress that the aforementioned results use different prediction setups and they do not imply any bounds for our setup. This is due to a different way of measuring prediction errors, see Section 1.3 for details. Therefore, we cannot compare the theoretical guarantees achieved by previously published caching algorithms in their prediction setup to our new caching algorithm within our broader setup. Instead, we provide a comparison via experiments.³

Combining Worst-Case and Optimistic Algorithms. An approach in some ways similar to ours was developed by Mahdian et al. (2012), who assume the existence of an optimistic algorithm and developed a meta-algorithm that combines this algorithm with a classical one and obtains a competitive ratio that is an interpolation between the ratios of the two algorithms. They designed such algorithms for several problems including facility location and load balancing. The competitive ratios obtained depend on the performance of the optimistic algorithm and the choice of the interpolation parameter. Furthermore the meta-algorithm is designed on a problem-by-problem basis. In contrast, (i) our performance guarantees are a function of the prediction error, (ii) generally we are able to approach the performance of the best algorithm, and (iii) our way of simulating multiple algorithms can be seen as a black box and is problem independent.

Online Algorithms with Advice. Another model for augmenting online algorithms, but not directly related to the prediction setting studied in this paper, is that of *advice complexity*, where information about the future is obtained in the form of some always correct bits of advice (see Boyar et al. (2017) for a survey). Emek et al. (2011) considered MTS under advice complexity, and Angelopoulos et al. (2020) consider advice complexity with possibly adversarial advice and focus on Pareto-optimal algorithms for consistency and robustness in several similar online problems.

²Their actual algorithm only needs the relative ordering of reoccurrence times, which is also true for Lykouris and Vassilvitskii (2018); Rohatgi (2020); Wei (2020).

³One might be tempted to adapt the algorithm of Rohatgi (2020) to action predictions by replacing the page with the furthest predicted reoccurrence in the algorithm of Rohatgi (2020) by a page evicted by the predictor in our setting. However, it is not hard, following ideas similar to the first example about prediction errors in the Section 1.3, to construct an instance where this algorithm is $\Omega(\log k)$ -competitive although $\frac{\eta}{\text{OPT}} = O(1)$ in our setup.

1.3 Comparison to the setup of Lykouris and Vassilvitskii

Although the work of Lykouris and Vassilvitskii (2018) for caching served as an inspiration, our prediction setup cannot be understood as an extension or generalization of their setup. Here we list the most important connections and differences.

Conversion of Predictions for Caching. One can convert the reoccurrence time predictions of Lykouris and Vassilvitskii (2018) for caching into predictions for our setup using a natural algorithm: At each page fault, evict the page whose next request is predicted furthest in the future. Note that, if given perfect predictions, this algorithm produces an optimal solution (Belady, 1966). The states of this algorithm at each time are then interpreted as predictions in our setup. We use this conversion to compare the performance of our algorithms to those of Lykouris and Vassilvitskii (2018) and Rohatgi (2020) in empirical experiments in Section 6.

Prediction Error. The prediction error as defined by Lykouris and Vassilvitskii (2018) is not directly comparable to ours. Here are two examples.

(1) Consider a paging instance where some page p is requested at times 1 and 3, and suppose we are given reoccurrence time predictions that are almost perfect except at time 1 where it is predicted that p reoccurs at time T rather than 3, for some large T . Then the prediction error in the setting of Lykouris and Vassilvitskii (2018) is $\Omega(T)$. However, the corresponding action predictions obtained by the conversion above are wrong only at time step 2, meaning the prediction error in our setting is only 1 with respect to the offline optimum.

(2) One can create a request sequence consisting of $k + 1$ distinct pages where swapping two predicted times of next arrivals causes a different prediction to be generated by the conversion algorithm. The modified prediction in the setup of Lykouris and Vassilvitskii (2018) may only have error 2 while the error in our setup with respect to the offline optimum can be arbitrarily high (depending on how far in the future these arrivals happen). However, our results provide meaningful bounds also in this situation. Such predictions still have error 0 in our setup with respect to a near-optimal algorithm which incurs only one additional page fault compared to the offline optimum. Theorems 1–3 then provide constant-competitive algorithms with respect to this near-optimal algorithm.

The first example shows that the results of Lykouris and Vassilvitskii (2018); Rohatgi (2020); Wei (2020) do not imply any bounds in our setup. On the other hand, the recent result of Wei (2020) shows that our algorithms from Theorems 1–3, combined with the prediction-converting algorithm above, are $O(1 + \min\{\frac{1}{k} \frac{\eta}{\text{OPT}}, \log k\})$ -competitive for caching in the setup of Lykouris and Vassilvitskii (2018), thus also matching the best known competitive ratio in that setup: The output of the conversion algorithm has error 0 with respect to itself and our algorithms are constant-competitive with respect to it. Since the competitive ratio of the conversion algorithm is $O(1 + \frac{1}{k} \frac{\eta}{\text{OPT}})$ by Wei (2020), our algorithms are $O(\min\{1 + \frac{1}{k} \frac{\eta}{\text{OPT}}, \log k\})$ -competitive, where η denotes the prediction error in the setup of Lykouris and Vassilvitskii (2018).

Succinctness. In the case of caching, we can restrict ourselves to *lazy* predictors, where each predicted cache content differs from the previous predicted cache content by at most one page, and only if the previous predicted cache content did not contain the requested page. This is motivated by the fact that any algorithm can be transformed into a lazy version of itself without increasing its cost. Therefore, $O(\log k)$ bits are enough to describe each action prediction, saying which page should be evicted, compared to $\Theta(\log T)$ bits needed to encode a reoccurrence time in the setup of Lykouris and Vassilvitskii (2018). In fact, we need to receive a prediction not for all time steps but only those when the current request is not part of the previous cache content of the predictor. In cases when running an ML predictor at each of these time steps is too costly, our setup allows predictions being generated by some fast heuristic whose parameters can be recalculated by the ML algorithm only when needed.

Learnability. In order to generate the reoccurrence time predictions, Lykouris and Vassilvitskii (2018) used a simple PLECO (Anderson et al., 2014) predictor. In this paper, we introduce another simple predictor called POPU and show that the output of these predictors can be converted to action predictions.

Predictors *Hawkey* Jain and Lin (2016) and *Glider* Shi et al. (2019) use binary classifiers to identify pages in the cache which are going to be reused soon, evicting first the other ones. As shown by their empirical results, such binary information is enough to produce a very efficient cache replacement policy, i.e., action predictions. In their recent paper, Liu et al. (2020) have proposed a new predictor, called *Parrot*, that is trained using the imitation learning approach and tries to mimic the behaviour of the optimal offline algorithm (Belady, 1966). The main output of their model, implemented using a neural network, are in fact action predictions. However it also produces the reoccurrence time predictions in order to add further supervision during the training process. While at first it may seem that predicting reoccurrence times is an easier task (in particular, it has the form of a standard supervised learning task), the results of Liu et al. (2020) show that it might well be the opposite – e.g., when the input instance variance makes it impossible to predict the reoccurrence times accurately yet it is still possible to solve it (nearly) optimally online. We refer to the paper of Chłędowski et al. (2021) for an extensive evaluation of the existing learning augmented algorithms using both reoccurrence time and action predictions. Following the emergence of learning-augmented algorithms, Anand et al. (2020) even designed predictors specifically tuned to optimize the error used in the algorithms analysis. This work has been restricted so far to a much simpler online problem, ski rental.

2 Preliminaries

In MTS, we are given a metric space M of states and an initial state $x_0 \in M$. At each time $t = 1, 2, \dots$, we receive a task $\ell_t: M \rightarrow \mathbb{R}^+ \cup \{0, +\infty\}$ and we have to choose a new state x_t without knowledge of the future tasks, incurring cost $\text{dist}(x_{t-1}, x_t) + \ell_t(x_t)$. Note that $\text{dist}(x_{t-1}, x_t) = 0$ if $x_{t-1} = x_t$ by the identity property of metrics.

Although MTS share several similarities with the *experts* problem from the theory of online learning (Freund and Schapire, 1997; Chung, 1994), there are three important differences. First, there is a *switching cost*: we need to pay cost for switching between states equal to their distance in the underlying metric space. Second, an algorithm for MTS has *one-step lookahead*, i.e., it can see the task (or loss function) before choosing the new state and incurring the cost of this task. Third, there can be *unbounded costs* in MTS, which can be handled thanks to the lookahead. See Blum and Burch (2000) for more details on the relation between experts and MTS.

To assess the performance of algorithms, we use the *competitive ratio* – the classical measure used in online algorithms.

Definition 1 (Competitive ratio). *Let \mathcal{A} be an online algorithm for some cost-minimization problem P . We say that \mathcal{A} is r -competitive and call r the competitive ratio of \mathcal{A} , if for any input sequence $I \in P$, we have*

$$\mathbb{E}[\text{cost}(\mathcal{A}(I))] \leq r \cdot \text{OPT}_I + \alpha,$$

where α is a constant independent of the input sequence, $\mathcal{A}(I)$ is the solution produced by the online algorithm and OPT_I is the cost of an optimal solution computed offline with the prior knowledge of the whole input sequence. The expectation is over the randomness in the online algorithm. If OPT_I is replaced by the cost of some specific algorithm OFF , we say that \mathcal{A} is r -competitive against OFF .

Before we prove our results for general MTS, we consider in the next section the caching problem. It corresponds to the special case of MTS where the metric space is the set of cardinality- k subsets of a cardinality- n set (of pages), the distance between two sets is the number of pages in which they differ, and each cost function assigns value 0 to all sets containing some page r_t and ∞ to other sets.

3 Logarithmic Error Dependence for Caching

We describe in this section a new algorithm, which we call TRUST&DOUBT, for the (unweighted) caching problem, and prove Theorem 3. The algorithm achieves a competitive ratio logarithmic in the error (thus overcoming the lower bound of Theorem 20 that holds for general MTS even on a uniform metric), while also attaining the optimal worst-case guarantee of $O(\log k)$.

We assume that the predictor is *lazy* in the following sense. Let r_t be the page that is requested at time t and let P_t be the configuration (i.e., set of pages in the cache) of the predictor at time t . Then P_t differs from P_{t-1} only if $r_t \notin P_{t-1}$ and, in this case, $P_t = P_{t-1} \cup \{r_t\} \setminus \{q\}$ for some page $q \in P_{t-1}$. Note that any algorithm for caching can be converted into a lazy one without increasing its cost.

We partition the request sequence into *phases*, which are maximal time periods where k distinct pages are requested⁴: The first phase begins with the first request. A phase ends (and a new phase begins) after k distinct pages have been requested in the current phase and right before the next arrival of a page that is different from all these k pages. For a given point in time, we say that a page is *marked* if it has been requested at least once in the current phase. For each page p requested in a phase, we call the first request to p in that phase the *arrival* of p . This is the time when p gets marked. Many algorithms, including that of Lykouris and Vassilvitskii (2018), belong to the class of so-called *marking algorithms*, which evict a page only if it is unmarked. The classical $O(\log k)$ -competitive online algorithm of (Fiat et al., 1991) is a particularly simple marking algorithm: On a cache miss, evict a uniformly random unmarked page. In general, no marking algorithm can be better than 2-competitive even when provided with perfect predictions. As will become clear from the definition of TRUST&DOUBT later, it may follow the predictor’s advice to evict even marked pages, meaning that it is not a marking algorithm. As can be seen in our experiments in Section 6, this allows TRUST&DOUBT to outperform previous algorithms when predictions are good.⁵ We believe that one could modify the algorithm so that it is truly 1-competitive in the case of perfect predictions. However, formally proving so seems to require a significant amount of additional technical complications regarding notation and algorithm description. To keep the presentation relatively simple, we abstain from optimizing constants here.

3.1 First warm-up: A universe of $k + 1$ pages

Before we give the full-fledged TRUST&DOUBT algorithm for the general setting, we first describe an algorithm for the simpler setting when there exist only $k + 1$ different pages that can be requested. This assumption substantially simplifies both the description and the analysis of the algorithm while already showcasing some key ideas. In Sections 3.2 and 3.3, we will explain the additional ideas required to extend the algorithm to the general case.

Our assumption means that at each time, there is only one page missing from the algorithm’s cache and only one page missing from the predicted cache. Moreover, the first request in each phase is an arrival of the (unique) page that was not requested in the previous phase, and all other arrivals in a phase are requests to pages that were also requested in the previous phase.

3.1.1 Algorithm (simplified setting)

We denote by M the set of marked pages and by U the set of unmarked pages.

In each phase, we partition time into alternating *Trust intervals* and *Doubt intervals*, as follows: When a phase starts, the first Trust interval begins. Throughout each Trust interval, we ensure that the algorithm’s cache is equal to the predicted cache P_t . As soon as the page missing from P_t is requested during a Trust interval, we terminate the current Trust interval and start a new Doubt interval. In a Doubt interval, we treat page faults by evicting a uniformly random page

⁴Subdividing the input sequence into such phases is a very common technique in the analysis of caching algorithms, see for example Borodin and El-Yaniv (1998) and references therein.

⁵There exist instances where TRUST&DOUBT with perfect predictions strictly outperforms the best marking algorithm, but also vice versa, see Appendix B.

from U . As soon as there have been 2^{i-1} arrivals since the beginning of the i th Doubt interval of a phase, the Doubt interval ends and a new Trust interval begins (and we again ensure that the algorithm's cache is equal to P_t).

3.1.2 Analysis (simplified setting)

Let d_ℓ be the number of Doubt intervals in phase ℓ .

Claim 6. *The expected number of cache misses in phase ℓ is $1 + O(d_\ell)$.*

Proof. Any cache miss during a Trust interval starts a new Doubt interval, so there are d_ℓ cache misses during Trust intervals. There may be one more cache miss at the start of the phase. It remains to show that there are $O(d_\ell)$ cache misses in expectation during Doubt intervals.

In a Doubt interval, we can have a cache miss only when a page from U arrives. The arriving page from U is the one missing from the cache with probability $1/|U|$. Moreover, when a page from U arrives, it is removed from U . The expected number of cache misses during Doubt intervals is therefore a sum of terms of the form $1/|U|$ for distinct values of $|U|$. Since the total number of arrivals during Doubt intervals is at most 2^{d_ℓ} , the expected number of cache misses during Doubt intervals is at most $\sum_{u=1}^{2^{d_\ell}} 1/u = O(d_\ell)$. \square

Due to the claim, our main remaining task is to upper bound the number of Doubt intervals.

We call a Doubt interval *error interval* if at each time t during the interval, the page missing from P_t is present in the cache of the offline algorithm. Note that each time step during an error interval contributes to the error η . Let $e_\ell \leq d_\ell$ be the number of error intervals of phase ℓ . Since for all $i < e_\ell$, the i th error interval contains at least 2^{i-1} time steps, we can bound the error as

$$\eta \geq \sum_{\ell} (2^{e_\ell-1} - 1). \quad (2)$$

Denote by OFF_ℓ the cost of the offline algorithm during phase ℓ .

Claim 7. $d_\ell \leq \text{OFF}_\ell + e_\ell + 1$.

Proof. Consider the quantity $d_\ell - e_\ell$. This is the number of Doubt intervals of phase ℓ during which P_t is equal to the offline cache at some point. Since P_t changes at the start of each Doubt interval, but P_t changes only if the page missing from P_t is requested (since we assume the predictor to be lazy), this means that the offline cache must change between any two such intervals. Thus, $d_\ell - e_\ell - 1 \leq \text{OFF}_\ell$. \square

Combining these claims, the total number of cache misses of the algorithm is at most (noting by $\#\text{phases}$ the number of phases)

$$\sum_{\ell} (1 + O(d_\ell)) \leq O \left(\text{OFF} + \#\text{phases} + \sum_{\ell} (e_\ell - 1) \right).$$

Each term $(e_\ell - 1)$ can be rewritten as $\log_2(1 + (2^{e_\ell-1} - 1))$. By concavity of $x \mapsto \log(1 + x)$, subject to the bound (2) the sum of these terms is maximized when each term $2^{e_\ell-1} - 1$ equals $\frac{\eta}{\#\text{phases}}$. Thus, the total number of cache misses of the algorithm is at most

$$O \left(\text{OFF} + \#\text{phases} + \#\text{phases} \cdot \log \left(1 + \frac{\eta}{\#\text{phases}} \right) \right) \leq \text{OFF} \cdot O \left(1 + \log \left(1 + \frac{\eta}{\text{OFF}} \right) \right),$$

where the last inequality uses that $\text{OFF} = \Omega(\#\text{phases})$ since all $k + 1$ pages are requested in any two adjacent phases, so the offline algorithm must have a cache miss in any two adjacent phases.

3.2 Second warm-up: The predictor is a marking algorithm

We now drop the assumption from the previous section and allow the number of pages in the universe to be arbitrary. However, we will assume in this section that the predictor is a marking algorithm (i.e., the predicted cache P_t always contains all marked pages). In this case, our algorithm will also be a marking algorithm.

Our algorithm is again based on phases, which are defined as before. Denote by U the set of unmarked pages that were in the cache at the beginning of the phase, and by M the set of marked pages. By our assumption that both the predictor and our algorithm are marking algorithms, at the start of a phase U is equal to both the predicted cache as well as the algorithm’s cache as it contains precisely the pages that were requested in the previous phase. An important notion in phase-based paging algorithms is that of *clean pages*. For the setting considered in this section, where the predictor is a marking algorithm, we define a page as *clean* if it is requested in the current phase but was not requested in the previous phase. We denote by C the set of clean pages that have arrived so far in the current phase. (In the general setting, we will need to define clean pages slightly differently.)

Several simultaneous interval partitions. While in the first warm-up setting with a $(k+1)$ -page universe there could be only a single clean page per phase, a main difference now is that there can be *several* clean pages in a phase. For this reason, it is no longer sufficient to partition the phase into Trust intervals and Doubt intervals that are defined “globally”. Instead, we will associate with *each* clean page a *different* subdivision of time into Trust intervals and Doubt intervals: The time from the arrival of q until the end of the phase is partitioned into alternating q -Trust intervals and q -Doubt intervals. Thus, a time t can belong to various intervals – one associated to each clean q that has arrived so far in the current phase. At any time, some of the current intervals may be Trust intervals while the rest are Doubt intervals. During a q -Trust interval, we will *not* ensure that the entire algorithm’s cache is equal to the predictor’s cache, but only that one particular page $f(q)$ that is evicted by the predictor is also predicted by the algorithm.

More precisely, we will also maintain a map $f: C \rightarrow U \setminus P_t$ that maps each clean page $q \in C_t$ (that has arrived so far) to an associated page $f(q)$ that was evicted by the predictor during the current phase (and is currently still missing from the predictor’s cache). Intuitively, we can think of $f(q)$ as the page that the predictor advises us to evict to make space for q . If it happens that the page $f(q)$ associated to some clean q is requested, the predictor has to load $f(q)$ back to its cache P_t , and we redefine $f(q)$ to be the page that the predictor evicts at this time. Observe that this ensures that the pages $f(q)$ are distinct for different q (in fact, since we assume the predictor to be a lazy marking algorithm, f is a bijection in this case).

When a clean page q arrives, the first q -Trust interval begins. Throughout each q -Trust interval, we will ensure that the page $f(q)$ is evicted from our algorithm’s cache. If during a q -Trust interval the page $f(q)$ is requested, we terminate the current q -Trust interval and start a new q -Doubt interval. In a q -Doubt interval, we ignore the advice to evict $f(q)$ and instead evict a uniformly random unmarked page when necessary. As soon as there have been 2^{i-1} arrivals since the beginning of the i th q -Doubt interval, the q -Doubt interval ends and a new q -Trust interval begins (and we again ensure that the page currently defined as $f(q)$ is evicted).

We will skip a more formal description and analysis of the algorithm for this setting as it will be contained as a special case of our algorithm in the next section.

Remark 8. *At a high level, the idea of linking evictions to individual clean pages (which is explicit for the pages $f(q)$ evicted in Trust intervals) bears some similarities to the notion of eviction chains used in Lykouris and Vassilvitskii (2018); Rohatgi (2020). However, our algorithm and charging scheme are quite different. In particular, the natural adaptations of algorithms in Lykouris and Vassilvitskii (2018); Rohatgi (2020) to our setting would only be $\Omega(\log k)$ -competitive even when $\frac{\eta}{OPT} = O(1)$, where η is the prediction error in our setting. This can happen on instances where predictions are mostly good, but occasionally very bad. To overcome this, we use Doubt intervals*

that start small and grow over time, which allows our algorithm to recover quickly from occasional very bad predictions.

3.3 Algorithm for the general case

We now describe our algorithm TRUST&DOUBT for the general case. In contrast to the previous section, we drop here the assumption that the predictor must be a marking algorithm. Thus, the predictor may evict marked pages, and since TRUST&DOUBT may trust such evictions, also TRUST&DOUBT may evict marked pages. Consequently, it is no longer true that set of pages in the algorithm’s cache at the start of a phase is equal to the set of pages requested in the previous phase. This means that some pages may be “ancient” as per the following definition.

Definition 2. *A page is called ancient if it is in TRUST&DOUBT’s cache even though it has been requested in neither the previous nor the current phase (so far).*

We partition each phase into two stages that are determined by whether ancient pages exist or not: During stage one there exists at least one ancient page, and during stage two there exist no ancient pages. We note that one of the two stages may be empty.

The algorithm for stage one is very simple: Whenever there is a page fault, evict an arbitrary ancient page. This makes sense since ancient pages have not been requested for a long time, so we treat them like a reserve of pages that are safe to evict. Once this reserve has been used up, stage two begins.

The algorithm for stage two is essentially the one already described in the previous section. Before we give a more formal description, we first fix some notation. Let U be the set of pages that were in cache at the beginning of stage two and that are currently unmarked. Let M be the set of marked pages. We call a page *clean* for a phase if it arrives in stage two and it was not in $U \cup M$ immediately before its arrival. (Pages arriving in stage one are *not* considered clean as these are easy to charge for and do not need the analysis linked to clean pages in stage two.) By C we denote the set of clean pages that have arrived so far in the current phase.




Figure 1: Illustration of definitions used to describe TRUST&DOUBT. At time t , a new phase starts and the cache contains a , c and d , where d is ancient. In the following phase, e is not clean because it arrives during stage one. When e is requested, TRUST&DOUBT evicts the ancient page d and loads e , and then stage two begins with $U = \{a, c\}$ and $M = \{e\}$ initially. The page b requested next is considered clean because, although it was also requested in the previous phase, it was not in $U \cup M$ immediately before its request. The next requested page a is not clean as it was already in $U \cup M$.

It is immediate from the definitions that the following equation is maintained during stage two:

$$|U \cup M| = k + |C| \tag{3}$$

Similarly to before, TRUST&DOUBT maintains an injective map $f: C \rightarrow (U \cup M) \setminus P_t$ that maps each clean page $q \in C$ (that has arrived so far) to a distinct page $f(q)$ that is currently missing from the predictor’s cache. Note that since the predictor may evict marked pages, it is necessary to include marked pages in the codomain of f . As before, the time from the arrival of a clean page q to the end of the phase is partitioned into alternating q -Trust intervals and q -Doubt

intervals. Depending on the type of the current interval, we will also say that q is trusted or q is doubted. Let

$$\begin{aligned} T &:= \{f(q) \mid q \in C \text{ trusted}\} \\ D &:= \{f(q) \mid q \in C \text{ doubted}\}. \end{aligned}$$

To organize the random evictions that the algorithm makes, we sort the pages of U in a uniformly random order at the beginning of stage two. We refer to the position of a page in this order as its *rank*, and we will ensure that the randomly evicted pages are those with the lowest ranks.⁶

A pseudocode of TRUST&DOUBT when a page r is requested in stage two is given in Algorithm 1.

If r is clean and the request is an arrival (the condition in line 1 is true), we first define its associated page $f(r)$ as an arbitrary⁷ page from $(U \cup M) \setminus (P_t \cup T \cup D)$. We will justify later in Lemma 10 that this set is non-empty. We then start an r -Trust interval. (Note that this adds $f(r)$ to the set T .) Since r is then trusted, we ensure that $f(r)$ is evicted from the cache. If it was already evicted, then we instead evict the page in cache with the lowest rank. Either way, there is now a free cache slot that r will be loaded to in line 6. We also initialize a variable t_r as 1. For each clean page q , we will use this variable t_q to determine the duration of the next q -Doubt interval.

Otherwise, we also ensure that r is in cache, evicting the page in cache with the lowest rank if necessary (lines 8–10).

If r is a page of the form $f(q)$, we redefine $f(q)$, and since the previous prediction to evict the old $f(q)$ was bad, we ensure that q is now doubted (lines 11–14), *i.e.*, we start a new q -Doubt interval if q was trusted.

Finally, in lines 15–20 we check for each clean page q whether it should reach the end of its Doubt interval. If q is in its i th Doubt interval, then this happens if the current arrival is the 2^{i-1} th arrival after the start of the current q -Doubt interval. For each q for which a q -Doubt interval ends, we start a new q -Trust interval and ensure that $f(q)$ is evicted from the cache. If $f(q)$ was not evicted yet, we reload the evicted page with the highest rank back to the cache so that the cache contains k pages at all times.

Remark 9. *To simplify the analysis, the algorithm is defined non-lazily here in the sense that it may load pages even when they are not requested (in line 20). An implementation should only simulate this non-lazy algorithm in the background and, whenever the actual algorithm has a page fault, it evicts an arbitrary (e.g., the least recently used) page that is present in its own cache but missing from the simulated cache.*

Correctness. It is straightforward to check that the algorithm’s cache is always a subset of $(U \cup M) \setminus T$, since any page added to T is evicted. Moreover, it is always a superset of $M \setminus T$ because pages from M are only evicted if they are in T .

The following two lemmas capture invariants that are maintained throughout the execution of the algorithm. In particular, they justify the the algorithm is well-defined.

Lemma 10. *The set $(U \cup M) \setminus (P_t \cup T \cup D)$ is non-empty before $f(r)$ or $f(q)$ is chosen from it in lines 2 and 12.*

Proof. It suffices to show that $|U \cup M| > |P_t \cup T \cup D|$ right before the respective line is executed.

Before line 2 is executed, it will be the case that $|C| = |T \cup D| + 1$ (because $r \in C$, but $f(r)$ is not defined yet). Thus, equation (3) yields $|U \cup M| = k + |T \cup D| + 1 > |P_t \cup T \cup D|$.

Before line 12 is executed, it holds that $|C| = |T \cup D|$ and $r \in (T \cup D) \cap P_t$. Again, the inequality follows from equation (3). \square

⁶Since randomly evicted pages may be reloaded even when they are not requested, maintaining such ranks leads to consistent random choices throughout a phase.

⁷e.g., the least recently used

Algorithm 1: When page r is requested in phase ℓ and no ancient pages exist

```

1 if  $r \in C$  and this is the arrival of  $r$  then // Arrival of a clean page
2   | Let  $f(r)$  be an arbitrary page from  $(U \cup M) \setminus (P_t \cup T \cup D)$ 
3   | Start an  $r$ -Trust interval
4   | if  $f(r)$  is in cache then evict  $f(r)$ 
5   | else evict the lowest ranked cached page from  $U \setminus T$ 
6   | Load  $r$  to the cache
7   |  $t_r := 1$  // Duration of next  $r$ -Doubt interval (if it exists)
8 else if  $r$  is not in cache then // Page fault, but not arrival of clean page
9   | Evict the lowest ranked cached page from  $U \setminus T$ 
10  | Load  $r$  to the cache
11 if  $r = f(q)$  for some  $q \in C$  then // Advice to evict  $f(q)$  was bad
12  | Redefine  $f(q)$  as an arbitrary page from  $(U \cup M) \setminus (P_t \cup T \cup D)$ 
13  | if  $q$  is trusted then
14  |   | End the  $q$ -Trust interval and start a  $q$ -Doubt interval
15 foreach  $q \in C$  that is doubted do // Check for end of Doubt intervals
16  | if the current request is the  $t_q$ th arrival since the start of this  $q$ -Doubt interval then
17  |   | End the  $q$ -Doubt interval and start a  $q$ -Trust interval
18  |   |  $t_q := 2 \cdot t_q$ 
19  |   | if  $f(q)$  is in cache then
20  |   |   | Evict  $f(q)$  and load the highest ranked evicted page from  $U \setminus T$  back to the
20  |   |   | cache

```

The next lemma justifies that reloading a page in line 20 will be possible, and the lemma will also be crucial for the competitive analysis later.

Lemma 11. *Before each request of stage two, there are $|D|$ pages from $U \setminus T$ missing from the cache.*

Proof. The pages in cache are a subset of $U \cup M$ of size k . By equation (3), there are $|C| = |T \cup D|$ of these pages missing from the cache. The pages in T account for $|T|$ of those missing pages. The remaining $|D|$ missing pages are all in $U \setminus T$ (because pages from M are only evicted if they are in T). \square

3.4 Competitive analysis

Let C_ℓ denote the set C at the end of phase ℓ . The next lemma and its proof are similar to a statement in Fiat et al. (1991). However, since our definition of clean pages is different, we need to reprove it in our setting.

Lemma 12. *Any offline algorithm suffers cost at least*

$$OFF \geq \Omega \left(\sum_{\ell} |C_\ell| \right).$$

Proof. We first claim that at least $k + |C_\ell|$ distinct pages are requested in phases $\ell - 1$ and ℓ together. If there is no stage two in phase ℓ , then C_ℓ is empty and the statement trivial. Otherwise, all pages that are in $U \cup M$ at the end of phase ℓ were requested in phase $\ell - 1$ or ℓ , and by equation (3) this set contains $k + |C_\ell|$ pages.

Thus, any algorithm must suffer at least cost $|C_\ell|$ during these two phases. Hence, OFF is lower bounded by the sum of $|C_\ell|$ over all even phases and, up to a constant, by the according sum over all odd phases. The lemma follows. \square

By the following lemma, it suffices to bound the cost of TRUST&DOUBT incurred during stage two.

Lemma 13. *The cost during stage one of phase ℓ is at most the cost during stage two of phase $\ell - 1$.*

Proof. The cost during stage one of phase ℓ is at most the number of ancient pages at the beginning of phase ℓ . This is at most the number of marked pages that were evicted in phase $\ell - 1$. Since a marked page can be evicted only during stage two, the lemma follows. \square

Let $d_{q,\ell}$ be the number of q -Doubt intervals in phase ℓ . The next lemma is reminiscent of Claim 6 from our first warm-up section.

Lemma 14. *The expected cost during stage two of phase ℓ is $O(|C_\ell| + \sum_{q \in C_\ell} d_{q,\ell})$.*

Proof. The cost incurred in lines 1–7 is at most $O(|C_\ell|)$. In lines 8–10, the algorithm can incur cost only if the requested page was in $U \cup T$ before the request (because the request is not an arrival of a clean page, so it was in $U \cup M$, and if it was in $M \setminus T$ then it was in cache already). If the page was in T , then a new Doubt interval will start in line 14, so the cost due to those pages is at most $\sum_{q \in C_\ell} d_{q,\ell}$. If the page was in $U \setminus T$, then by Lemma 11 and the random choice of ranks it was missing from the cache with probability $\frac{|D|}{|U \setminus T|}$. To account for this cost, we charge $\frac{1}{|U \setminus T|}$ to each clean q that is doubted at the time. Over the whole phase, the number of times we charge to each $q \in C_\ell$ in this way is at most the total number of arrivals during q -Doubt intervals, which is at most $2^{d_{q,\ell}}$. By equation (3) and since $|C| = |T| + |D|$, we have $|U| = k + |T| + |D| - |M|$, so $|U \setminus T| \geq k + |D| - |M| \geq k + 1 - |M|$. The quantity $|M|$ increases by 1 after each such request to a page in $|U \setminus T|$, so the value of $|U \setminus T|$ can be lower bounded by $1, 2, 3, \dots, 2^{d_{q,\ell}}$ during the at most $2^{d_{q,\ell}}$ arrivals when $\frac{1}{|U \setminus T|}$ is charged to q . Hence, the total cost charged to q is at most $O(d_{q,\ell})$. It follows that the overall cost incurred in lines is at most $O\left(\sum_{q \in C_\ell} d_{q,\ell}\right)$.

Finally, the only other time cost is incurred is in line 20. This also amounts to at most $\sum_{q \in C_\ell} d_{q,\ell}$ because it happens only at the end of a Doubt-interval. \square

Denote by $e_{q,\ell}$ the number of q -Doubt intervals with the property that at each time during the interval, the page currently defined as $f(q)$ is present in the offline cache. Since the current page $f(q)$ is never in the predictor's cache, and the i th doubted q -interval contains 2^{i-1} arrivals for $i < d_q$, a lower bound on the prediction error is given by

$$\eta \geq \sum_{\ell} \sum_{q \in C_\ell} (2^{e_{q,\ell}} - 1). \quad (4)$$

Denote by $\text{OFF}_{q,\ell}$ the number of times in phase ℓ when the offline algorithm incurs cost for loading the page currently defined as $f(q)$ to its cache.

The next lemma is the generalization of Claim 7.

Lemma 15. *For each $q \in C_\ell$, we have $d_{q,\ell} \leq \text{OFF}_{q,\ell} + e_{q,\ell} + 1$.*

Proof. Consider the quantity $d_{q,\ell} - e_{q,\ell}$. This is the number of q -Doubt intervals of phase ℓ during which $f(q)$ is missing from the offline cache at some point. Except for the last such interval, the page $f(q)$ will subsequently be requested during the phase, so the offline algorithm will incur cost for loading it to its cache. The lemma follows, with the “+1” term accounting for the last interval. \square

We are now ready to prove the main result of this section.

Theorem (Restated Theorem 3). *TRUST&DOUBT has competitive ratio $O(\min\{1 + \log(1 + \frac{\eta}{\text{OFF}}), \log k\})$ against any offline algorithm OFF, where η is the prediction error with respect to OFF.*

Proof. The $O(\log k)$ bound follows from Lemma 14, Lemma 12 and the fact that $d_{q,\ell} \leq O(\log k)$ for each $q \in C_\ell$. The latter fact holds because if $d_{q,\ell} \geq 2$, then the $(d_{q,\ell} - 1)$ st q -Doubt interval contains $2^{d_{q,\ell}-2}$ arrivals, but there are only k arrivals per phase.

For the main bound, combining Lemmas 13, 14 and 15 we see that the total cost of the algorithm is at most

$$O\left(\text{OFF} + \sum_{\ell} |C_{\ell}| + \sum_{\ell} \sum_{q \in C_{\ell}} e_{q,\ell}\right).$$

The summands $e_{q,\ell}$ can be rewritten as $1 + \log_2(1 + [2^{e_{q,\ell}-1} - 1])$. By concavity of $x \mapsto \log(1+x)$, while respecting the bound (4) the sum of these terms is maximized when each term in brackets equals $\frac{\eta}{\sum_{\ell} |C_{\ell}|}$, giving a bound on the cost of

$$O\left(\text{OFF} + \sum_{\ell} |C_{\ell}| \left(1 + \log\left(1 + \frac{\eta}{\sum_{\ell} |C_{\ell}|}\right)\right)\right).$$

Since this quantity is increasing in $\sum_{\ell} |C_{\ell}|$, applying Lemma 12 completes the proof of the theorem. \square

3.5 Lower bound

The $O(\log k)$ upper bound matches the known lower bound $\Omega(\log k)$ on the competitive ratio of randomized online algorithms without prediction Fiat et al. (1991). The competitive ratio of TRUST&DOUBT when expressed only as a function of the error, $O(1 + \log(1 + \frac{\eta}{\text{OFF}}))$, is also tight due to the following theorem. It should be noted, though, that for the competitive ratio as a function of both k and $\frac{\eta}{\text{OFF}}$ it is still plausible that a better bound can be achieved when $\frac{\eta}{\text{OFF}}$ is relatively small compared to k .

Theorem 16. *If an online caching algorithm achieves competitive ratio at most $f(\frac{\eta}{\text{OPT}})$ for arbitrary k when provided with action predictions with error at most η with respect to an optimal offline algorithm OPT, then $f(x) = \Omega(\log x)$ as $x \rightarrow \infty$.*

Proof. Fix some $k+1$ pages and consider the request sequence where each request is to a uniformly randomly chosen page from this set. We define phases in the same way as in the description of TRUST&DOUBT. By a standard coupon collector argument, each phase lasts $\Theta(k \log k)$ requests in expectation. An optimal offline algorithm can suffer only one page fault per page by evicting only the one page that is not requested in each phase. On the other hand, since requests are chosen uniformly at random, any online algorithm suffers a page fault with probability $1/(k+1)$ per request, giving a cost of $\Theta(\log k)$ per phase. Since $\frac{\eta}{\text{OPT}} = O(k \log k)$ due to the duration of phases, the competitive ratio of the algorithm is $\Omega(\log k) = \Omega(\log \frac{\eta}{\text{OPT}})$. \square

4 Robust Algorithms for MTS

The goal of this section is to prove Theorem 1 and Theorem 2, which deal with algorithms substantially simpler than TRUST&DOUBT, but demonstrate the usefulness of our prediction setup for the broad class of MTS problems. In Section 4.1 we will first describe a simple algorithm whose competitive ratio depends linearly on the prediction error, but the algorithm is not robust against large errors. In Section 4.2 we then robustify this algorithm based on powerful methods from the literature. Finally, in Section 4.3 we show that the linear dependency of the achieved competitive ratio on η/OPT is inevitable for some MTS.

4.1 A non-robust algorithm

We consider a simple memoryless algorithm, which we call FTP.

Algorithm Follow the Prediction (FtP). Intuitively, our algorithm follows the predictions, but still somewhat cautiously: if there exists a state “close” to the predicted one that has a much cheaper service cost, then it is to be preferred. Let us consider a metrical task system with a set of states X . We define the algorithm FtP (Follow the Prediction) as follows: at time t , after receiving task ℓ_t and prediction p_t , it moves to the state

$$x_t \leftarrow \arg \min_{x \in X} \{\ell_t(x) + 2\text{dist}(x, p_t)\}. \quad (5)$$

In other words, FtP follows the predictions except when it is beneficial to move from the predicted state to some other state, pay the service and move back to the predicted state.

Lemma 17. *For any MTS with action predictions, algorithm FtP which achieves competitive ratio $1 + \frac{4\eta}{\text{OFF}}$ against any offline algorithm OFF, where η is the prediction error with respect to OFF.*

Proof. At each time t , the FtP algorithm is located at configuration x_{t-1} and needs to choose x_t after receiving task ℓ_t and prediction p_t . Let us consider some offline algorithm OFF. We denote x_0, o_1, \dots, o_T the states of OFF, where the initial state x_0 is common for OFF and for FtP, and T denotes the length of the sequence.

We define A_t to be the algorithm which agrees with FtP in its first t configurations x_0, x_1, \dots, x_t and then agrees with the states of OFF, i.e., o_{t+1}, \dots, o_T . Note that $\text{cost}(A_0) = \text{OFF}$ and $\text{cost}(A_T) = \text{cost}(\text{FtP})$. We claim that $\text{cost}(A_t) \leq \text{cost}(A_{t-1}) + 4\eta_t$ for each t , where $\eta_t = \text{dist}(p_t, o_t)$. The algorithms A_t and A_{t-1} are in the same configuration at each time except t , when A_t is in x_t while A_{t-1} is in o_t . By the triangle inequality, we have

$$\begin{aligned} \text{cost}(A_t) &\leq \text{cost}(A_{t-1}) + 2\text{dist}(o_t, x_t) + \ell_t(x_t) - \ell_t(o_t) \\ &\leq \text{cost}(A_{t-1}) + 2\text{dist}(o_t, p_t) - \ell_t(o_t) + 2\text{dist}(p_t, x_t) + \ell_t(x_t) \\ &\leq \text{cost}(A_{t-1}) + 4\text{dist}(o_t, p_t), \end{aligned}$$

The last inequality follows from (5): we have $2\text{dist}(p_t, x_t) + \ell_t(x_t) \leq 2\text{dist}(p_t, o_t) + \ell_t(o_t)$. By summing over all times $t = 1, \dots, T$, we get

$$\text{cost}(\text{FtP}) = \text{cost}(A_T) \leq \text{cost}(A_0) + 4 \sum_{t=1}^T \eta_t,$$

which equals OFF + 4η . □

4.2 Combining online algorithms

We describe now how to make algorithm FtP robust by combining it with a classical online algorithm. Although we only need to combine two algorithms, we will formulate the combination theorems more generally for any number of algorithms.

Consider m algorithms A_0, \dots, A_{m-1} for some problem P belonging to MTS. We describe two methods to combine them into one algorithm which achieves a performance guarantee close to the best of them. Note that these methods are also applicable to problems which do not belong to MTS as long as one can simulate all the algorithms at once and bound the cost for switching between them.

Deterministic Combination. The following method was proposed by Fiat et al. (1994) for the k -server problem, but can be generalized to MTS. We note that a similar combination is also mentioned in Lykouris and Vassilvitskii (2018). We simulate the execution of A_0, \dots, A_{m-1} simultaneously. At each time, we stay in the configuration of one of them, and we switch between the algorithms in the manner of a solution for the m -lane *cow path* problem, see Algorithm 2 for details.

Theorem 18 (generalization of Theorem 1 in Fiat et al. (1994)). *Given m online algorithms A_0, \dots, A_{m-1} for a problem P in MTS, the algorithm MIN^{det} achieves cost at most $(\frac{2\gamma^m}{\gamma-1} + 1) \cdot \min_i \{\text{cost}_{A_i}(I)\}$, for any input sequence I .*

Algorithm 2: MIN^{det} (Fiat et al., 1994)

```
1 choose  $1 < \gamma \leq 2$ ; set  $\ell := 0$ 
2 repeat
3    $i := \ell \bmod m$ 
4   while  $cost(A_i) \leq \gamma^\ell$ , follow  $A_i$ 
5    $\ell := \ell + 1$ 
6 until the end of the input
```

A proof of this theorem can be found in Section A. The optimal choice of γ is $\frac{m}{m-1}$. Then $\frac{2\gamma^m}{\gamma-1} + 1$ becomes 9 for $m = 2$, and can be bounded by $2em$ for larger m . Combined with Lemma 17, we obtain Theorem 1.

Randomized Combination. Blum and Burch (2000) proposed the following way to combine online algorithms based on the WMR (Littlestone and Warmuth, 1994) (Weighted Majority Randomized) algorithm for the experts problem. At each time t , it maintains a probability distribution p^t over the m algorithms updated using WMR. Let $dist(p^t, p^{t+1}) = \sum_i \max\{0, p_i^t - p_i^{t+1}\}$ be the earth-mover distance between p^t and p^{t+1} and let $\tau_{ij} \geq 0$ be the transfer of the probability mass from p_i^t to p_j^{t+1} certifying this distance, so that $p_i^t = \sum_{j=0}^{m-1} \tau_{ij}$ and $dist(p^t, p^{t+1}) = \sum_{i \neq j} \tau_{ij}$. If we are now following algorithm A_i , we switch to A_j with probability τ_{ij}/p_i^t . See Algorithm 3 for details. The parameter D is an upper bound on the switching cost between the states of two algorithms.

Algorithm 3: MIN^{rand} (Blum and Burch, 2000)

```
1  $\beta := 1 - \frac{\epsilon}{2}$ ; // for parameter  $\epsilon < 1/2$ 
2  $w_i^0 := 1$  for each  $i = 0, \dots, m-1$ ;
3 foreach time  $t$  do
4    $c_i^t :=$  cost incurred by  $A_i$  at time  $t$ ;
5    $w_i^{t+1} := w_i^t \cdot \beta^{c_i^t/D}$  and  $p_i^{t+1} := \frac{w_i^{t+1}}{\sum w_i^{t+1}}$ ;
6    $\tau_{i,j} :=$  mass transferred from  $p_i^t$  to  $p_j^{t+1}$ ;
7   switch from  $A_i$  to  $A_j$  w.p.  $\tau_{ij}/p_i^t$ ;
```

Theorem 19 (Blum and Burch (2000)). *Given m on-line algorithms A_0, \dots, A_{m-1} for an MTS with diameter D and $\epsilon < 1/2$, there is a randomized algorithm MIN^{rand} such that, for any instance I , its expected cost is at most*

$$(1 + \epsilon) \cdot \min_i \{cost(A_i(I))\} + O(D/\epsilon) \ln m.$$

Combined with Lemma 17, we obtain Theorem 2.

4.3 Lower bound

We show that our upper bounds for general metrical task systems (Theorems 1 and 2) are tight up to constant factors. We show this for MTS on a uniform metric, i.e., the metric where the distance between any two points is 1.

Theorem 20. *For $\bar{\eta} \geq 0$ and $n \in \mathbb{N}$, every deterministic (or randomized) online algorithm for MTS on the n -point uniform metric with access to an action prediction oracle with error at most $\bar{\eta} \cdot OPT$ with respect to some optimal offline algorithm has competitive ratio $\Omega(\min\{\alpha_n, 1 + \bar{\eta}\})$, where $\alpha_n = \Theta(n)$ (or $\alpha_n = \Theta(\log n)$) is the optimal competitive ratio of deterministic (or randomized) algorithms without prediction.*

Proof. For deterministic algorithms, we construct an input sequence consisting of phases defined as follows. We will ensure that the online and offline algorithms are located at the same point at the beginning of a phase. The first $\min\{n - 2, \lfloor \bar{\eta} \rfloor\}$ cost functions of a phase always take value ∞ at the old position of the online algorithm and value 0 elsewhere, thus forcing the algorithm to move. Let p be a point that the online algorithm has not visited since the beginning of the phase. Only one more cost function will be issued to conclude the phase, which takes value 0 at p and ∞ elsewhere, hence forcing both the online and offline algorithms to p . The optimal offline algorithm suffers a cost of exactly 1 per phase because it can move to p already at the beginning of the phase. The error is at most $\bar{\eta}$ per phase provided that point p is predicted at the last step of the phase, simply because there are only at most $\bar{\eta}$ other steps in the phase, each of which can contribute at most 1 to the error. Thus, the total error is at most $\bar{\eta} \text{OPT}$. The online algorithm suffers a cost $\min\{n - 1, 1 + \lceil \bar{\eta} \rceil\}$ during each phase, which proves the deterministic lower bound.

For randomized algorithms, let $k := \lfloor \log_2 n \rfloor$ and fix a subset F_0 of the metric space of 2^k points. We construct again an input sequence consisting of phases: For $i = 1, \dots, \min\{k, \lfloor \bar{\eta} \rfloor\}$, the i th cost function of a phase takes value 0 on some set F_i of feasible states and ∞ outside of F_i . Here, we define $F_i \subset F_{i-1}$ to be the set consisting of the half of the points of F_{i-1} where the algorithm's probability of residing is smallest right before the i th cost function of the phase is issued (breaking ties arbitrarily). Thus, the probability of the algorithm already residing at a point from F_i when the i th cost function arrives is at most $1/2$, and hence the expected cost per step is at least $1/2$. We assume that $\bar{\eta} \geq 1$ (otherwise the theorem is trivial). Similarly to the deterministic case, the phase concludes with one more cost function that forces the online and offline algorithms to some point p in the final set F_i . Again, the optimal cost is exactly 1 per phase, the error is at most $\bar{\eta}$ in each phase provided the last prediction of the phase is correct, and the algorithm's expected cost per phase is at least $\frac{1}{2} \min\{k, \lfloor \bar{\eta} \rfloor\} = \Omega(\min(\log n, 1 + \bar{\eta}))$, concluding the proof. \square

In light of the previous theorem it may seem surprising that our algorithm TRUST&DOUBT for caching (see Section 3) achieves a competitive ratio logarithmic rather than linear in the prediction error, especially considering that the special case of caching when there are only $k + 1$ distinct pages corresponds to an MTS on the uniform metric. However, the construction of the randomized lower bound in Theorem 20 requires cost functions that take value ∞ at several points at once, whereas in caching only one page is requested per time step.

5 Beyond Metrical Task Systems

The objective of this section is to show that the prediction setup introduced in this paper is not limited to Metrical Task Systems, but can also be useful for relevant problems not known to be inside this class. This emphasizes the generality of our approach, compared to prediction setups designed for a single problem. We focus on the *online matching on the line* problem, which has been studied for three decades and has seen recent developments.

In the *online matching on the line* problem, we are given a set $S = \{s_1, s_2, \dots, s_n\}$ of server locations on the real line. A set of requests $R = \{r_1, r_2, \dots, r_n\}$ which are also locations on the real line, arrive over time. Once request r_i arrives, it has to be irrevocably matched to some previously unmatched server s_j . The cost of this edge in the matching is the distance between r_i and s_j , i.e., $|s_j - r_i|$ and the total cost is given by the sum of all such edges in the final matching, i.e., the matching that matches every request in R to some unique server in S . The objective is to minimize this total cost.

The best known lower bound on the competitive ratio of any deterministic algorithm is 9.001 (Fuchs et al., 2005) and the best known upper bound for any algorithm is $O(\log n)$, due to Raghvendra (2018).

We start by defining the notion of *distance* between two sets of servers.

Definition 3. Let P_i^1 and P_i^2 be two sets of points in a metric space, of size i each. We then say that their distance $\text{dist}(P_i^1, P_i^2)$ is equal to the cost of a minimum-cost perfect matching in the bipartite graph having P_i^1 and P_i^2 as the two sides of the bipartition.

In *online matching on the line with action predictions* we assume that, in each round i along with request r_i , we obtain a prediction $P_i \subseteq S$ with $|P_i| = i$ on the server set that the offline optimal algorithm is using for the first i many requests. We allow here even that $P_i \not\subseteq P_{i+1}$. The error in round i is given by $\eta_i := \text{dist}(P_i, \text{OFF}_i)$, where OFF_i is the server set of a (fixed) offline algorithm on the instance. The total prediction error is $\eta = \sum_{i=1}^n \eta_i$.

Since a request has to be irrevocably matched to a server, it is not straightforward that one can switch between configurations of different algorithms. Nevertheless, we are able to simulate such a switching procedure. By applying this switching procedure to the best known classic online algorithm for the problem, due to Raghvendra (2018), and designing a Follow-The-Prediction algorithm that achieves a competitive ratio of $1 + 2\eta/\text{OFF}$, we can apply the combining method of Theorem 18 to get the following result.

Theorem (Restated Theorem 4). *There exists a deterministic algorithm for the online matching on the line problem with action predictions that attains a competitive ratio of*

$$\min\{O(\log n), 9 + \frac{8\eta}{\text{OFF}}\},$$

for any offline algorithm OFF .

We note that for some instances the switching cost between these two algorithms (and therefore, in a sense, also the metric space diameter) can be as high as $\Theta(\text{OPT})$ which renders the randomized combination uninteresting for this particular problem.

5.1 A potential function

We define the *configuration* of an algorithm at some point in time as the set of servers which are currently matched to a request.

For each round of the algorithm, we define S_i as the current configuration and P_i as the predicted configuration, which verify $|S_i| = |P_i| = i$. We define a potential function after each round i to be $\Phi_i = \text{dist}(S_i, P_i)$, and let μ_i be the associated matching between S_i and P_i that realizes this distance, such that all servers in $S_i \cap P_i$ are matched to themselves for zero cost. We extend μ_i to the complete set of servers S by setting $\mu_i(q) = q$ for all $q \notin S_i \cup P_i$. The intuition behind the potential function is that after round i one can simulate being in configuration P_i instead of the actual configuration S_i , at an additional expense of Φ_i .

5.2 Distance among different configurations

The purpose of this section is to show that the distance among the configurations of two algorithms is at most the sum of their current costs. As we will see, this will imply that we can afford switching between any two algorithms.

We continue by bounding the distance between any two algorithms as a function of their costs.

Lemma 21. *Consider two algorithms A and B , and fix the set of servers S as well as the request sequence R . Let A_i and B_i be the respective configurations of the algorithms (i.e., currently matched servers) after serving the first i requests of R with servers from S . Furthermore, let OPT_i^A (resp. OPT_i^B) be the optimal matching between $\{r_1, r_2, \dots, r_i\}$ and A_i (resp. B_i), and let M_i^A (resp. M_i^B) be the corresponding matching produced by A (resp. B). Then:*

$$\begin{aligned} \text{dist}(A_i, B_i) &\leq \text{cost}(\text{OPT}_i^A) + \text{cost}(\text{OPT}_i^B) \\ &\leq \text{cost}(M_i^A) + \text{cost}(M_i^B). \end{aligned}$$

Proof. The second inequality follows by the optimality of OPT_i^A and OPT_i^B . For the first inequality let s_j^A (resp. s_j^B) be the server matched to r_j by OPT_i^A (resp. OPT_i^B), for all $j \in \{1, \dots, i\}$. Therefore, there exists a matching between A_i and B_i that matches for all $j \in \{1, \dots, i\}$, s_j^A to s_j^B which has a total cost of

$$\begin{aligned} \sum_{j=1}^i \text{dist}(s_j^A - s_j^B) &\leq \sum_{j=1}^i \text{dist}(s_j^A - r_j) + \sum_{j=1}^i \text{dist}(s_j^B - r_j) \\ &= \text{cost}(\text{OPT}_i^A) + \text{cost}(\text{OPT}_i^B), \end{aligned}$$

where the inequality follows by the triangle inequality. By the definition of distance we have that $\text{dist}(A_i, B_i) \leq \sum_{j=1}^i \text{dist}(s_j^A - s_j^B)$, which concludes the proof. \square

5.3 Follow-The-Prediction

Since *online matching on the line* is not known to be in MTS, we start by redefining the algorithm Follow-The-Prediction for this particular problem. In essence, the algorithm virtually switches from predicted configuration P_i to predicted configuration P_{i+1} .

Let S_i be the actual set of servers used by Follow-The-Prediction after round i . Follow-The-Prediction computes the optimal matching among P_{i+1} and the multiset $P_i \cup \{r_{i+1}\}$ which maps the elements of $P_{i+1} \cap P_i$ to themselves. Note that if $r_{i+1} \in P_i$, then $P_i \cup \{r_{i+1}\}$ is a multiset where r_{i+1} occurs twice. Such matching will match r_{i+1} to some server $s \in P_{i+1} \setminus P_i$. Recall that μ_i is the minimum cost bipartite matching between S_i and P_i extended by zero-cost edges to the whole set of servers. Follow-The-Prediction matches r_{i+1} to the server $\mu_i(s)$, i.e., to the server to which s is matched to under μ_i . We can show easily that $\mu(s) \notin S_i$. Since $s \notin P_i$, there are two possibilities: If $s \notin S_i$, then $\mu(s) = s \notin S_i$ by extension of μ_i to elements which do not belong to S_i nor P_i . Otherwise, $s \in S_i \setminus P_i$ and, since μ_i matches all the elements of $S_i \cap P_i$ to themselves, we have $\mu(s) \in P_i \setminus S_i$.

Theorem 22. *Follow-The-Prediction has total matching cost at most $\text{OFF} + 2\eta$ and therefore the algorithm has a competitive ratio of*

$$1 + 2\eta/\text{OFF}$$

against any offline algorithm OFF.

Proof. The idea behind the proof is that, by paying the switching cost of $\Delta\Phi_i$ at each round, we can always virtually assume that we reside in configuration P_i . So whenever a new request r_{i+1} and a new predicted configuration P_{i+1} arrive, we pay the costs for switching from P_i to P_{i+1} and for matching r_{i+1} to a server in P_{i+1} .

We first show that, for every round i , we have:

$$\begin{aligned} FtP_i + \Delta\Phi_i &\leq \text{dist}(P_{i+1}, P_i \cup \{r_{i+1}\}) \\ \Leftrightarrow \text{dist}(r_{i+1}, \mu(s)) + \Phi_{i+1} &\leq \text{dist}(P_{i+1}, P_i \cup \{r_{i+1}\}) + \Phi_i. \end{aligned}$$

Note that for all j , $\Phi_j = \text{dist}(S_j, P_j) = \text{dist}(\bar{S}_j, \bar{P}_j)$, where \bar{S}_j and \bar{P}_j denote the complements of S_j and P_j respectively.

We have in addition $\text{dist}(\bar{S}_i, \bar{P}_i) = \text{dist}(\bar{S}_i \setminus \{\mu_i(s)\}, \bar{P}_i \setminus \{s\}) + \text{dist}(s, \mu_i(s))$ as $s \notin P_i$ and $\mu_i(s) \notin S_i$, and $(s, \mu_i(s))$ is an edge in the min-cost matching between \bar{S}_i and \bar{P}_i . Note that $S_{i+1} = S_i \cup \{\mu_i(s)\}$ so $\bar{S}_i \setminus \{\mu_i(s)\} = \bar{S}_{i+1}$. Therefore, we get:

$$\Phi_i = \text{dist}(\bar{S}_i, \bar{P}_i) = \text{dist}(\bar{S}_{i+1}, \bar{P}_i \setminus \{s\}) + \text{dist}(s, \mu_i(s)) = \text{dist}(S_{i+1}, P_i \cup \{s\}) + \text{dist}(s, \mu_i(s)).$$

In addition, we have $\text{dist}(P_{i+1}, P_i \cup \{r_{i+1}\}) = \text{dist}(s, r_{i+1}) + \text{dist}(P_{i+1} \setminus \{s\}, P_i)$ because by definition of s , s is matched to r_{i+1} in a minimum cost matching between P_{i+1} and $P_i \cup \{r_{i+1}\}$.

Now, $s \notin P_i$, so $\text{dist}(P_{i+1} \setminus \{s\}, P_i) = \text{dist}(P_{i+1}, P_i \cup \{s\})$ as this is equivalent to adding a zero-length edge from s to itself to the associated matching. Therefore, we get:

$$\text{dist}(P_{i+1}, P_i \cup \{r_{i+1}\}) = \text{dist}(s, r_{i+1}) + \text{dist}(P_{i+1}, P_i \cup \{s\}).$$

Combining the results above, we obtain:

$$\begin{aligned} FtP_i + \Delta\Phi_i &\leq \text{dist}(P_{i+1}, P_i \cup \{r_{i+1}\}) \\ \Leftrightarrow \text{dist}(r_{i+1}, \mu_i(s)) + \Phi_{i+1} &\leq \text{dist}(P_{i+1}, P_i \cup \{r_{i+1}\}) + \Phi_i \\ \Leftrightarrow \text{dist}(r_{i+1}, \mu_i(s)) + \text{dist}(S_{i+1}, P_{i+1}) & \\ &\leq \text{dist}(S_{i+1}, P_i \cup \{s\}) + \text{dist}(s, \mu_i(s)) + \text{dist}(s, r_{i+1}) + \text{dist}(P_{i+1}, P_i \cup \{s\}) \end{aligned}$$

The last equation holds by the triangle inequality.

Finally, we bound $\text{dist}(P_{i+1}, P_i \cup \{r_{i+1}\})$ using the triangle inequality. In the following OFF_i refers to the configuration of offline algorithm OFF after the first i requests have been served.

$$\begin{aligned} &\text{dist}(P_{i+1}, P_i \cup \{r_{i+1}\}) \\ &\leq \text{dist}(P_i \cup \{r_{i+1}\}, \text{OFF}_i \cup \{r_{i+1}\}) + \text{dist}(\text{OFF}_i \cup \{r_{i+1}\}, \text{OFF}_{i+1}) + \text{dist}(\text{OFF}_{i+1}, P_{i+1}) \\ &\leq \eta_i + |\text{OFF}_i| + \eta_{i+1}. \end{aligned}$$

Summing up over all rounds, and using that $\Phi_1 = \Phi_n = 0$ completes the proof of the theorem. \square

5.4 The main theorem

The goal of this subsection is to prove Theorem 4.

Proof of Theorem 4. The main idea behind the proof is to show that we can apply Theorem 18 and virtually simulate the two algorithms (Follow-The-Prediction and the online algorithm of Raghvendra (2018)).

We need to show that we can assume that we are in some configuration and executing the respective algorithm, and that the switching cost between these configurations is upper bounded by the cost of the two algorithms. Similarly to the analysis of Follow-The-Prediction, we can virtually be in any configuration as long as we pay for the distance between any two consecutive configurations. When we currently simulate an algorithm A , the distance between the two consecutive configurations is exactly the cost of the edge that A introduces in this round. When we switch from the configuration of some algorithm A to the configuration of some algorithm B , then by Lemma 21, the distance between the two configurations is at most the total current cost of A and B .

This along with Theorem 24 (which is generalizing Theorem 18 beyond MTS and can be found in Appendix A) concludes the proof. \square

5.5 Bipartite metric matching

Bipartite metric matching is the generalization of online matching on the line where the servers and requests can be points of any metric space. The problem is known to have a tight $(2n - 1)$ -competitive algorithm, due to Kalyanasundaram and Pruhs (1993) as well as Khuller et al. (1994).

We note that our arguments in this section are not line-specific and apply to that problem as well. This gives the following result:

Theorem 23. *There exists a deterministic algorithm for the online metric bipartite matching problem with action predictions that attains a competitive ratio of*

$$\min\{2n - 1, 9 + \frac{8e\eta}{OFF}\},$$

against any offline algorithm OFF.

6 Experiments

We evaluate the practicality of our approach on real-world datasets for two MTS: *caching* and *ice cream* problem. The source code and datasets are available at GitHub⁸. Each experiment was run 10 times and we report the mean competitive ratios. The maximum standard deviation we observed was of the order of 0.001.

6.1 The caching problem

Datasets. For the sake of comparability, we used the same two datasets as Lykouris and Vassilvitskii (2018).

- BK dataset comes from a former social network BrightKite (Cho et al., 2011). It contains checkins with user IDs and locations. We treat the sequence of checkin locations of each users as a separate instance of caching problem. We filter users with the maximum sequence length (2100) who require at least 50 evictions in an optimum cache policy. Out of those we take the first 100 instances. We set the cache size to $k = 10$.
- Citi dataset comes from a bike sharing platform CitiBike (2017). For each month of 2017, we consider the first 25 000 bike trips and build an instance where a request corresponds to the starting station of a trip. We set the cache size to $k = 100$.

Predictions. We first generate the reoccurrence time predictions, these predictions being used by previous prediction-augmented algorithms. To this purpose, we use the same two predictors as Lykouris and Vassilvitskii (2018). Additionally we also consider a simple predictor, which we call POPU (from *popularity*), and the LRU heuristic adapted to serve as a predictor.

- Synthetic predictions: we first compute the exact reoccurrence time for each request, setting it to the end of the instance if it does not reappear. We then add some noise drawn from a lognormal distribution, with the mean parameter 0 and the standard deviation σ , in order to model rare but large failures.
- PLECO predictions: we use the PLECO model described in Anderson et al. (2014), with the same parameters as Lykouris and Vassilvitskii (2018), which were fitted for the BK dataset (but not refitted for Citi). This model estimates that a page requested x steps earlier will be the next request with a probability proportional to $(x + 10)^{-1.8}e^{-x/670}$. We sum the weights corresponding to all the earlier appearances of the current request to obtain the probability p that this request is also the next one. We then estimate that such a request will reappear $1/p$ steps later.
- POPU predictions: if the current request has been seen in a fraction p of the past requests, we predict it will be repeated $1/p$ steps later.
- LRU predictions: Lykouris and Vassilvitskii (2018) already remarked on (but did not evaluate experimentally) a predictor that emulates the behavior of the LRU heuristic. A page requested at time t is predicted to appear at time $-t$. Note that the algorithms only consider the order of predicted times among pages, and not their values, so the negative predictions pointing to the past are not an issue.

⁸<https://github.com/adampolak/mts-with-predictions>

Algorithm	Competitive ratio	Property	Reference
LRU	k		(Sleator and Tarjan, 1985)
Marker	$O(\log k)$	Robust	(Fiat et al., 1991)
FtP	$1 + 4\frac{\eta}{\text{OPT}}$	C+S	Lemma 17
L&V	$2 + O(\min\{\sqrt{\frac{\eta'}{\text{OPT}}}, \log k\})$	C+S+R	(Lykouris and Vassilvitskii, 2018)
RobustFtP	$(1 + \epsilon) \min\{1 + 4\frac{\eta}{\text{OPT}}, O(\log k)\}$	C+S+R	Theorem 2
LMarker	$O(1 + \min\{\log \frac{\eta'}{\text{OPT}}, \log k\})$	C+S+R	(Rohatgi, 2020)
LNonMarker	$O(1 + \min\{1, \frac{\eta'}{k \cdot \text{OPT}}\} \log k)$	C+S+R	(Rohatgi, 2020)
TRUST&DOUBT	$O(\min\{1 + \log(1 + \frac{\eta}{\text{OPT}}), \log k\})$	C+S+R	Theorem 3

Table 1: Summary of caching algorithms evaluated in experiments. Note that η and η' are different measures of prediction error, so their functions should not be compared directly. Properties C+S+R mean Consistency, Smoothness, and Robustness respectively.

We then transform the reoccurrence time predictions to action predictions by simulating the algorithm that evicts the element predicted to appear the furthest in the future. In each step the prediction to our algorithm is the configuration of this algorithm. Note that in the case of LRU predictions, the predicted configuration is precisely the configuration of the LRU algorithm.

Algorithms. We considered the following algorithms, whose competitive ratios are reported in Table 1. Two online algorithms: the heuristic LRU, which is considered the gold standard for caching, and the $O(\log k)$ -competitive Marker (Fiat et al., 1994). Three robust algorithms from the literature using the “next-arrival time” predictions: L&V (Lykouris and Vassilvitskii, 2018), LMarker (Rohatgi, 2020), and LNonMarker (Rohatgi, 2020). Three algorithms using the prediction setup which is the focus of this paper: FtP, which naively follows the predicted state, RobustFtP, which is defined as $MIN^{rand}(\text{FtP}, \text{Marker})$, and is an instance of the general MTS algorithm described in Section 4, and TRUST&DOUBT, the caching algorithm described in Section 3.

We implemented the deterministic and randomized combination schemes described in Section 4.2 with a subtlety for the caching problem: we do not flush the whole cache when switching algorithms, but perform only a single eviction per page fault in the same way as described in Remark 9. We set the parameters to $\gamma = 1 + 0.01$ and $\epsilon = 0.5$. These values, chosen from $\{0.001, 0.01, 0.1, 0.5\}$, happen to be consistently the best choice in all our experimental settings.

<i>Dataset</i>	BK			Citi		
LRU	1.291			1.848		
Marker	1.333			1.861		
<i>Predictions</i>	PLECO	POPU	LRU	PLECO	POPU	LRU
FtP	2.081	1.707	1.291	2.277	1.734	1.848
L&V	1.340	1.262	1.291	1.877	1.776	1.848
LMarker	1.337	1.264	1.291	1.876	1.780	1.848
LNonMarker	1.333	1.292	1.299	1.862	1.771	1.855
RobustFtP	1.338	1.316	1.297	1.862	1.831	1.849
Trust&Doubt	1.292	1.276	1.291	1.847	1.775	1.849

Table 2: Competitive ratios of caching algorithms using PLECO, POPU, and LRU predictions on both datasets.

Results. For both datasets, for each algorithm and each prediction considered, we computed the total number of page faults over all the instances and divided it by the optimal number in order to obtain a *competitive ratio*. Figure 2 presents the performance of a selection of the




Figure 2: Comparison of caching algorithms augmented with synthetic predictions on the BK dataset.

algorithms depending on the noise of synthetic predictions for the BK dataset. We omit LMarker and LNonMarker for readability since they perform no better than L&V. This experiment shows that our algorithm TRUST&DOUBT outperforms previous prediction-based algorithms as well as LRU on the BK dataset with such predictions. Figures 3 and 4 present the performance of all algorithms on the BK and Citi datasets, respectively. On the Citi dataset (Figure 4), FTP achieves very good results even with a noisy synthetic predictor, and therefore RobustFtP surpasses other guaranteed algorithms. LNonMarker presents better performance for noisy predictions than the other algorithms.

In Table 2 we provide the results obtained on both datasets using PLECO, POPU, and LRU predictions. We observe that PLECO predictions are not accurate enough to allow previously known algorithms to improve over the Marker algorithm. This may be due to the sensitivity of this predictor to consecutive identical requests, which are irrelevant for the caching problem. However, using the simple POPU predictions enables the prediction-augmented algorithms to significantly improve their performance compared to the classical online algorithms. Using TRUST&DOUBT with either of the predictions is however sufficient to get a performance similar or better than LRU (and than all other alternatives, excepted for POPU predictions on the BK dataset). RobustFtP, although being a very generic algorithm with worse theoretical guarantees than TRUST&DOUBT, achieves a performance which is not that far from previously known algorithms. Note that we did not use a prediction model tailored to our setup, which suggests that even better results can be achieved. When we use the LRU heuristic as a predictor, all the prediction-augmented algorithms perform comparably to the bare LRU algorithm. For TRUST&DOUBT and RobustFtP, there is a theoretical guarantee that this must be the case: Since the prediction error with respect to LRU is 0, these algorithms are $O(1)$ -competitive against LRU. Thus, TRUST&DOUBT achieves both the practical performance of LRU with an exponentially better worst-case guarantee than LRU. Note that Lykouris and Vassilvitskii (2018) also discuss how their algorithm framework performs when using LRU predictions, but did not provide both of these theoretical guarantees simultaneously.

6.2 A simple MTS: the *ice cream* problem

We consider a simple MTS example from Chrobak and Larmore (1998), named *ice cream* problem. It is an MTS with two states, named v and c , at distance 1 from each other, and two types of requests, V and C . Serving a request while being in the matching state costs 1 for V and 2 for C ,




Figure 3: Comparison of caching algorithms augmented with synthetic predictions on the BK dataset.




Figure 4: Comparison of caching algorithms augmented with synthetic predictions on the Citi dataset.

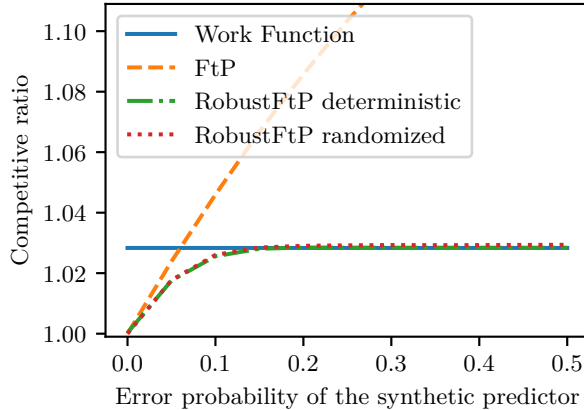


Figure 5: Performance on the ice cream problem with synthetic predictions.

and the costs are doubled for the mismatched state. The problem is motivated by an ice cream machine which operates in two modes (states) – vanilla or chocolate – each facilitating a cheaper production of a type of ice cream (requests).

We use the BrightKite dataset to prepare test instances for the problem. We extract the same 100 users as for caching. For each user we look at the geographic coordinates of the checkins, and we issue a V request for each checkin in the northmost half, and a C request for each checkin in the southmost half.

In order to obtain synthetic predictions, we first compute the optimal offline policy, using dynamic programming. Then, for an error parameter p , for each request we follow the policy with probability $1 - p$, and do the opposite with probability p .

We consider the following algorithms: the Work Function algorithm (Borodin et al., 1992; Borodin and El-Yaniv, 1998), of competitive ratio of 3 in this setting ($2n - 1$ in general); FtP, defined in Section 4 (in case of ties in Equation 5, we follow the prediction); and the deterministic and randomized combination of the two above algorithms (with the same ϵ and γ as previously) as proposed in Section 4.

Figure 5 presents the competitive ratios we obtained. We can see that the general MTS algorithms we propose in Section 4 allow to benefit from good predictions while providing the worst-case guarantee of the classical online algorithm. The deterministic and randomized combinations are comparable to the best of the algorithms combined, and improve upon them when both algorithms have a similar performance.

7 Conclusion

In this paper, we proposed a prediction setup that allowed us to design a general prediction-augmented algorithm for a large class of problems encompassing MTS. For the MTS problem of caching in particular, the setup requires less information from the predictor than previously studied ones (since previous predictions can be converted to ours). Despite the more general setup, we can design a specific algorithm for the caching problem in our setup which offers guarantees of a similar flavor to previous algorithms and even performs better in most of our experiments.

It may be considered somewhat surprising that a better bound is attainable for caching than for general MTS, given that our lower bound instance for MTS uses a uniform metric (and caching with a $(k + 1)$ -point universe also corresponds to a uniform metric). We conjecture logarithmic smoothness guarantees are also attainable for other MTS problems with a request structure similar to caching, like weighted caching and the k -server problem. Further special cases of MTS can be obtained by restricting the number of possible distinct requests (for example an MTS with

two different possible requests can model an important power management problem Irani et al. (2003)), or requiring a specific structure from the metric space. Several such parametrizations of MTS were considered by Bubeck and Rabani (2020b) and it would be interesting to study whether an improved dependence on the prediction error can be obtained in such settings.

With respect to matching problems, there have been recent investigations through the lens of learning augmentation under specific matroid constraints in (Antoniadis et al., 2020), but this territory is still largely unexplored. It would also be interesting to evaluate our resource augmented algorithm for online metric matchings in real-life situations. As an example online matching algorithms are employed in several cities in order to match cars to parking spots (for example SFpark in San Francisco or ParkPlus in Calgary). Not only have such matching problems been studied from an algorithmic point of view (see, e.g., Bender et al. (2020), Bender et al. (2021)), but arguably it should be doable to generate high-quality predictions from historical data making our approach very promising.

Another research direction is to identify more sophisticated predictors for caching and other problems that will further enhance the performance of prediction-augmented algorithms.

References

- K. Anand, R. Ge, and D. Panigrahi. Customizing ml predictions for online algorithms. In *International Conference on Machine Learning*, pages 303–313. PMLR, 2020.
- A. Anderson, R. Kumar, A. Tomkins, and S. Vassilvitskii. The dynamics of repeat consumption. In *Proceedings of conference World Wide Web '14*, pages 419–430, 2014. doi: 10.1145/2566486.2568018.
- S. Angelopoulos, C. Dürr, S. Jin, S. Kamali, and M. Renault. Online Computation with Untrusted Advice. In *Proceedings of ITCS'20*, volume 151, pages 52:1–52:15, 2020. doi: 10.4230/LIPIcs.ITCS.2020.52.
- A. Antoniadis, T. Gouleakis, P. Kleer, and P. Koley. Secretary and online matching problems with machine learned advice. In *Proceedings of NeurIPS'20*, 2020.
- N. Bansal, N. Buchbinder, and J. Naor. A primal-dual randomized algorithm for weighted paging. *J. ACM*, 59(4):19:1–19:24, 2012. doi: 10.1145/2339123.2339126.
- N. Bansal, C. Coester, R. Kumar, M. Purohit, and E. Vee. Learning-augmented weighted paging. In *Proceedings of the Thirty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, SODA'22*, 2022.
- L. A. Belady. A study of replacement algorithms for a virtual-storage computer. *IBM Syst. J.*, 5(2):78–101, 1966. doi: 10.1147/sj.52.0078.
- M. Bender, J. Gilbert, A. Krishnan, and K. Pruhs. Competitively pricing parking in a tree. In *WINE*, volume 12495 of *Lecture Notes in Computer Science*, pages 220–233. Springer, 2020.
- M. Bender, J. Gilbert, and K. Pruhs. A poly-log competitive posted-price algorithm for online metrical matching on a spider. In *FCT*, volume 12867 of *Lecture Notes in Computer Science*, pages 67–84. Springer, 2021.
- A. Blum and C. Burch. On-line learning and the metrical task system problem. *Machine Learning*, 39(1):35–58, 2000. doi: 10.1023/A:1007621832648.
- A. Borodin and R. El-Yaniv. *Online computation and competitive analysis*. Cambridge University Press, 1998.
- A. Borodin, N. Linial, and M. E. Saks. An optimal on-line algorithm for metrical task system. *J. ACM*, 39(4):745–763, 1992. doi: 10.1145/146585.146588.

- J. Boyar, L. M. Favrholdt, C. Kudahl, K. S. Larsen, and J. W. Mikkelsen. Online algorithms with advice: A survey. *ACM Comput. Surv.*, 50(2):19:1–19:34, 2017. doi: 10.1145/3056461.
- S. Bubeck and Y. Rabani. Parametrized metrical task systems. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2020)*, 2020a.
- S. Bubeck and Y. Rabani. Parametrized metrical task systems. In *APPROX-RANDOM*, volume 176 of *LIPICs*, pages 54:1–54:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020b.
- S. Bubeck, M. B. Cohen, J. R. Lee, and Y. T. Lee. Metrical task systems on trees via mirror descent and unfair gluing. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pages 89–97, 2019. doi: 10.1137/1.9781611975482.6.
- J. Chłędowski, A. Polak, B. Szabucki, and K. T. Żoła. Robust learning-augmented caching: An experimental study. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 1920–1930. PMLR, 2021. URL <https://proceedings.mlr.press/v139/chledowski21a.html>.
- E. Cho, S. A. Myers, and J. Leskovec. Friendship and mobility: user movement in location-based social networks. In *Proceedings of SIGKDD’11*, pages 1082–1090, 2011. doi: 10.1145/2020408.2020579. URL <https://snap.stanford.edu/data/loc-brightkite.html>.
- M. Chrobak and L. L. Larmore. Metrical task systems, the server problem and the work function algorithm. In *Online Algorithms*, pages 74–96. Springer, 1998. doi: 10.1007/BFb0029565.
- T. H. Chung. Approximate methods for sequential decision making using expert advice. In *Proceedings of COLT’94*, pages 183–189. Association for Computing Machinery, 1994. doi: 10.1145/180139.181097.
- CitiBike. Citi bike trip histories. <https://www.citibikenyc.com/system-data>, 2017. Accessed: 02/02/2020.
- C. Coester and E. Koutsoupias. The online k -taxi problem. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019*, pages 1136–1147, 2019. doi: 10.1145/3313276.3316370.
- C. Coester and J. R. Lee. Pure entropic regularization for metrical task systems. In *Conference on Learning Theory, COLT 2019*, pages 835–848, 2019.
- A. Daniely and Y. Mansour. Competitive ratio vs regret minimization: achieving the best of both worlds. In *Proceedings of ALT 2019*, pages 333–368, 2019. URL <http://proceedings.mlr.press/v98/daniely19a.html>.
- S. Dehghani, S. Ehsani, M. Hajiaghayi, V. Liaghat, and S. Seddighin. Stochastic k -Server: How Should Uber Work? In *44th International Colloquium on Automata, Languages, and Programming (ICALP 2017)*, volume 80, pages 126:1–126:14, 2017. doi: 10.4230/LIPICs.ICALP.2017.126.
- Y. Emek, P. Fraigniaud, A. Korman, and A. Rosén. Online computation with advice. *Theor. Comput. Sci.*, 412(24):2642–2656, 2011. doi: 10.1016/j.tcs.2010.08.007.
- A. Fiat, R. M. Karp, M. Luby, L. A. McGeoch, D. D. Sleator, and N. E. Young. Competitive paging algorithms. *J. Algorithms*, 12(4):685–699, 1991. doi: 10.1016/0196-6774(91)90041-V.
- A. Fiat, Y. Rabani, and Y. Ravid. Competitive k -server algorithms. *J. Comput. Syst. Sci.*, 48(3): 410–428, 1994. doi: 10.1016/S0022-0000(05)80060-1.

- A. Fiat, D. P. Foster, H. J. Karloff, Y. Rabani, Y. Ravid, and S. Vishwanathan. Competitive algorithms for layered graph traversal. *SIAM J. Comput.*, 28(2):447–462, 1998. doi: 10.1137/S0097539795279943.
- Y. Freund and R. E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences*, 55(1):119–139, 1997. doi: <https://doi.org/10.1006/jcss.1997.1504>.
- B. Fuchs, W. Hochstättler, and W. Kern. Online matching on a line. *Theor. Comput. Sci.*, 332(1-3):251–264, 2005. doi: 10.1016/j.tcs.2004.10.028.
- S. Gollapudi and D. Panigrahi. Online algorithms for rent-or-buy with expert advice. In *Proceedings of ICML’19*, pages 2319–2327, 2019. URL <http://proceedings.mlr.press/v97/gollapudi19a.html>.
- S. Irani, S. Shukla, and R. Gupta. Online strategies for dynamic power management in systems with multiple power-saving states. *ACM Trans. Embed. Comput. Syst.*, 2(3):325–346, 2003. doi: 10.1145/860176.860180.
- A. Jain and C. Lin. Back to the future: Leveraging belady’s algorithm for improved cache replacement. *SIGARCH Comput. Archit. News*, 44(3):78–89, June 2016. ISSN 0163-5964. doi: 10.1145/3007787.3001146. URL <https://doi.org/10.1145/3007787.3001146>.
- Z. Jiang, D. Panigrahi, and K. Su. Online algorithms for weighted paging with predictions. In *47th International Colloquium on Automata, Languages, and Programming (ICALP 2020)*, 2020.
- B. Kalyanasundaram and K. Pruhs. Online weighted matching. *J. Algorithms*, 14(3):478–488, 1993. doi: 10.1006/jagm.1993.1026.
- H. J. Karloff, Y. Rabani, and Y. Ravid. Lower bounds for randomized k-server and motion-planning algorithms. *SIAM J. Comput.*, 23(2):293–312, 1994. doi: 10.1137/S0097539792224838.
- E. B. Khalil, B. Dilkina, G. L. Nemhauser, S. Ahmed, and Y. Shao. Learning to run heuristics in tree search. In *Proceedings of IJCAI’17*, pages 659–666, 2017. doi: 10.24963/ijcai.2017/92.
- S. Khuller, S. G. Mitchell, and V. V. Vazirani. On-line algorithms for weighted bipartite matching and stable marriages. *Theor. Comput. Sci.*, 127(2):255–267, 1994. doi: 10.1016/0304-3975(94)90042-6.
- T. Kraska, A. Beutel, E. H. Chi, J. Dean, and N. Polyzotis. The case for learned index structures. In *Proceedings of SIGMOD’18*, pages 489–504, 2018. doi: 10.1145/3183713.3196909.
- S. Lattanzi, T. Lavastida, B. Moseley, and S. Vassilvitskii. Online scheduling via learned weights. In *Proceedings of the Thirty-First Annual ACM-SIAM Symposium on Discrete Algorithms, SODA’20*, pages 1859–1877, 2020.
- J. R. Lee. Lower bounds for MTS. Lecture notes, 2018. URL <https://tcsmath.github.io/online/2018/04/20/mts-lower-bounds/>. Accessed: 02/02/2020.
- M. Lin, A. Wierman, L. L. H. Andrew, and E. Thereska. Dynamic right-sizing for power-proportional data centers. *IEEE/ACM Trans. Netw.*, 21(5):1378–1391, 2013. doi: 10.1109/TNET.2012.2226216.
- N. Littlestone and M. Warmuth. The weighted majority algorithm. *Information and Computation*, 108(2):212–261, Feb. 1994. doi: 10.1006/inco.1994.1009.

- E. Liu, M. Hashemi, K. Swersky, P. Ranganathan, and J. Ahn. An imitation learning approach for cache replacement. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 6237–6247. PMLR, 13–18 Jul 2020. URL <https://proceedings.mlr.press/v119/liu20f.html>.
- T. Lykouris and S. Vassilvitskii. Competitive caching with machine learned advice. In *Proceedings of ICML'18*, pages 3302–3311, 2018. URL <http://proceedings.mlr.press/v80/lykouris18a.html>.
- M. Mahdian, H. Nazerzadeh, and A. Saberi. Online optimization with uncertain information. *ACM Trans. Algorithms*, 8(1):2:1–2:29, 2012. doi: 10.1145/2071379.2071381.
- M. S. Manasse, L. A. McGeoch, and D. D. Sleator. Competitive algorithms for server problems. *J. ACM*, 11(2):208–230, 1990. doi: 10.1016/0196-6774(90)90003-W.
- A. M. Medina and S. Vassilvitskii. Revenue optimization with approximate bid predictions. In *Proceedings of NeurIPS'17*, pages 1858–1866, 2017.
- M. Mitzenmacher. Scheduling with predictions and the price of misprediction. In *Proceedings of ITCS'20*, pages 14:1–14:18, 2020. doi: 10.4230/LIPIcs.ITCS.2020.14.
- M. Purohit, Z. Svitkina, and R. Kumar. Improving online algorithms via ML predictions. In *Proceedings of NeurIPS'18*, pages 9684–9693, 2018.
- S. Raghvendra. Optimal analysis of an online algorithm for the bipartite matching problem on a line. In *Proceedings of SoCG'18*, pages 67:1–67:14, 2018. doi: 10.4230/LIPIcs.SoCG.2018.67.
- D. Rohatgi. Near-optimal bounds for online caching with machine learned advice. In *Proceedings of the Thirty-First Annual ACM-SIAM Symposium on Discrete Algorithms, SODA'20*, pages 1834–1845, 2020.
- Z. Shi, X. Huang, A. Jain, and C. Lin. Applying deep learning to the cache replacement problem. In *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture, MICRO '52*, page 413–425, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450369381. doi: 10.1145/3352460.3358319. URL <https://doi.org/10.1145/3352460.3358319>.
- D. D. Sleator and R. E. Tarjan. Amortized efficiency of list update and paging rules. *Commun. ACM*, 28(2):202–208, 1985. doi: 10.1145/2786.2793.
- A. Wei. Better and simpler learning-augmented online caching. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2020)*, 2020.

A Deterministic combination of a collection of algorithms

We consider a problem P and m algorithms A_0, A_1, \dots, A_{m-1} for this problem which fulfill the following requirements.

- A_0, \dots, A_{m-1} start at the same state and we are able to simulate the run of all of them simultaneously
- for two algorithms A_i and A_j , the cost of switching between their states is bounded by $cost(A_i) + cost(A_j)$.

Theorem 24 (Restated Theorem 18; generalization of Theorem 1 in Fiat et al. (1994)). *Given m on-line algorithms A_0, \dots, A_{m-1} for a problem P which satisfy the requirements above, the algorithm MIN^{det} with parameter $1 < \gamma \leq 2$ incurs cost at most*

$$\left(\frac{2\gamma^m}{\gamma-1} + 1 \right) \cdot \min_i \{ \text{cost}(A_i(I)) \},$$

on any input instance I such that $OPT_I \geq 1$. If we choose $\gamma = \frac{m}{m-1}$, the coefficient $\frac{2\gamma^m}{\gamma-1} + 1$ equals 9 if $m = 2$ and can be bounded by $2em$.

Note that assumption on $OPT_I \geq 1$ is just to take care of the corner-case instances with very small costs. If we can only assume $OPT_I \geq c$ for some $0 < c < 1$, then we scale all the costs fed to MIN^{det} by $1/c$ and instances with $OPT_I = 0$ are usually not very interesting. The value of c is usually clear from the particular problem in hand, e.g., for caching we only care about instances which need at least one page fault, i.e., $OPT_I \geq 1$.

Proof. Let us consider the ℓ -th cycle of the algorithm and denote $i = \ell \bmod m$ and $i' = (\ell - 1) \bmod m$. We are switching from algorithm $A_{i'}$, whose current cost we denote $\text{cost}'(A_{i'}) = \gamma^{\ell-1}$ to A_i , whose current cost we denote $\text{cost}'(A_i)$, and its cost at the end of this cycle will become $\text{cost}(A_i) = \gamma^\ell$. Our cost during this cycle, i.e., for switching and for execution of A_i , is at most

$$\text{cost}'(A_{i'}) + \text{cost}'(A_i) + (\text{cost}(A_i) - \text{cost}'(A_i)) = \text{cost}'(A_{i'}) + \text{cost}(A_i) = \gamma^{\ell-1} + \gamma^\ell.$$

Now, let us consider the last cycle L , when we run the algorithm number $i = L \bmod m$. By the preceding equation, the total cost of MIN^{det} can be bounded as

$$\text{cost}(MIN^{det}) \leq 2 \cdot \sum_{\ell}^{L-1} \gamma^\ell + \text{cost}(A_i) = 2 \frac{\gamma^L - 1}{\gamma - 1} + \text{cost}(A_i) \leq 2 \frac{\gamma^L}{\gamma - 1} + \text{cost}(A_i).$$

If $L < m$, we use the fact that $OPT \geq 1$ and therefore the cost of each algorithm processing the whole instance would be at least one. Therefore, we have

$$\text{cost}(MIN^{det}) \leq 2 \frac{\gamma^L}{\gamma - 1} + \gamma^L \leq 2 \frac{\gamma^m}{\gamma - 1} \cdot \min_i \{ \text{cost}(A_i) \},$$

because $\frac{\gamma^L}{\gamma-1} + \gamma^L = \frac{\gamma^{L+1}}{\gamma-1}$ and $L+1 \leq m$.

Now, we have $L \geq m$, denoting $i = L \bmod m$, and we distinguish two cases.

(1) If $\min_j \{ \text{cost}(A_j) \} = \text{cost}(A_i)$, then $\text{cost}(A_i) \geq \gamma^{L-m}$ for each i , and therefore

$$\frac{\text{cost}(MIN^{det})}{\min_j \{ \text{cost}(A_j) \}} \leq \frac{2 \frac{\gamma^L}{\gamma-1} + \text{cost}(A_i)}{\min_j \{ \text{cost}(A_j) \}}$$

Note that $\text{cost}(A_i) \geq \gamma^{L-m}$, its cost from the previous usage. Since $\min_j \{ \text{cost}(A_j) \} = \text{cost}(A_i)$, we get

$$\frac{\text{cost}(MIN^{det})}{\min_j \{ \text{cost}(A_j) \}} \leq 2 \frac{\gamma^m}{\gamma - 1} + 1.$$

(2) Otherwise, we have $\min_j \{ \text{cost}(A_j) \} \geq \gamma^{L-m+1}$ and $\text{cost}(A_j) \leq \gamma^L$ and therefore

$$\frac{\text{cost}(MIN^{det})}{\min_j \{ \text{cost}(A_j) \}} \leq 2 \frac{\gamma^{m-1}}{\gamma - 1} + \gamma^{m-1} \leq 2 \frac{\gamma^m}{\gamma - 1}.$$

For $\gamma = \frac{m}{m-1}$ we have

$$2 \frac{\gamma^m}{\gamma - 1} + 1 = 2(m-1) \left(\frac{m}{m-1} \right)^m + 1,$$

which equals 9 for $m = 2$ and can be bounded by $2em$. □

B Comparison between Trust&Doubt and the best marking algorithm

The algorithm TRUST&DOUBT does not belong to the broad class of *marking* algorithms. We notice in this section that, given perfect predictions, this property allows it to outperform all marking algorithms on some instances, but, at the same time, it does not always perform as well as the best marking algorithm even when given perfect predictions.

Remark 25. *With perfectly accurate predictions, there exist both a caching instance on which TRUST&DOUBT performs better than the best marking algorithm, and another caching instance on which TRUST&DOUBT is outperformed by a marking algorithm.*

Proof. We first build an instance where TRUST&DOUBT, given predictions corresponding to the optimal algorithm evicting the page arriving the furthest in the future, outperforms the best marking algorithm. Consider a cache of size 3 and the request sequence 1, 2, 3; 4, 5, 6; 1, 2, 3, composed of three phases of length three (separated by semicolons). TRUST&DOUBT keeps the pages 1 and 2 in cache during the second phase so suffers seven cache misses. The best marking algorithm is not able to keep such old pages in cache so suffers nine cache misses.

Now, we build an instance where TRUST&DOUBT, given again predictions corresponding to the optimal algorithm evicting the page arriving the furthest in the future, suffers more cache misses than the best marking algorithm. Consider a cache of size 3, and the request sequence 1, 2, 3; 4, 5, 6, 5, 6; 7, 1, 4, composed of three phases of length three, five and three. The best marking algorithm suffers eight cache misses, the page 4 being present in the cache for the last request. The cache of TRUST&DOUBT after the second phase contains 1, 5, 6, as the page 1 is given priority over the page 4, and, at the start of the last phase, the now ancient page 1 is evicted, so the algorithm suffers nine cache misses. \square

C Limitations of the reoccurrence time predictions

In this section, we prove Theorem 5. In previous works on caching (Lykouris and Vassilvitskii, 2018; Rohatgi, 2020; Wei, 2020), the predictions are the time of the next reoccurrence to each page. It is natural to try extending this type of predictions to other problems, such as weighted caching. In weighted caching each page has a weight/cost that is paid each time the page enters the cache. However, it turns out that even with perfect predictions of this type for weighted caching, one cannot improve upon the competitive ratio $\Theta(\log k)$, which can already be attained without predictions (Bansal et al., 2012). Our proof is based on a known lower bound for MTS on a so-called “superincreasing” metric (Karloff et al., 1994). Following a presentation of this lower bound by Lee (2018), we modify the lower bound so that the perfect predictions provide no additional information.

We call an algorithm for weighted caching *semi-online* if it is online except that it receives in each time step, as an additional input, the reoccurrence time of the currently requested page (guaranteed to be without error). We prove the following result:

Theorem (Restated Theorem 5). *Every randomized semi-online algorithm for weighted caching is $\Omega(\log k)$ -competitive.*

Proof. Let $\tau > 0$ be some large constant. Consider an instance of weighted caching with cache size k and $k + 1$ pages, denoted by the numbers $0, \dots, k$, and such that the weight of page i is $2\tau^i$. It is somewhat easier to think of the following equivalent *evader* problem: Let S_k be the weighted star with leaves $0, 1, \dots, k$ and such that leaf i is at distance τ^i from the root. A single evader is located in the metric space. Whenever there is a request to page i , the evader must be located at some leaf of S_k other than i . The cost is the distance traveled by the evader. Any weighted caching algorithm gives rise to the evader algorithm that keeps its evader at the one leaf that is *not* currently in the algorithm’s cache. The cost between the two models differs only by an additive constant (depending on k and τ).

For $h = 1, \dots, k$ and a non-empty time interval (a, b) , we will define inductively a random sequence $\sigma_h = \sigma_h(a, b)$ of requests to the leaves $0, \dots, h$, such that each request arrives in the time interval (a, b) and

$$A_h \geq 4\alpha_{h-1}\tau^h \geq \alpha_h \cdot \text{OPT}_h, \quad (6)$$

where A_h denotes the expected cost of an arbitrary semi-online algorithm to serve the random sequence σ_h while staying among the leaves $0, \dots, h$, OPT_h denotes the expected optimal offline cost of doing so with an offline evader that starts and ends at leaf 0, $\alpha_0 = \frac{1+\tau}{4\tau}$, and $\alpha_h = 1/2 + \beta \log h$ for $h \geq 1$, where $\beta > 0$ is a constant determined later. The inequality between the first and last term in (6) implies the theorem. We will also ensure that $(0, 1, \dots, h)$ is both a prefix and a suffix of the sequence of requests in σ_h .

For the base case $h = 1$, the inequality is satisfied by the request sequence σ_1 that requests first 0 and then 1 at arbitrary times within the interval (a, b) .

For $h \geq 2$, the request sequence σ_h consists of subsequences (iterations) of the following two types (we will only describe the sequence of request locations for now and later how to choose the exact arrival times of these requests): A *type 1 iteration* is the sequence $(0, 1, \dots, h)$. A *type 2 iteration* is the concatenation of $\lceil \frac{\tau^h}{\alpha_{h-1} \text{OPT}_{h-1}} \rceil$ independent samples of a random request sequence of the form σ_{h-1} . The request sequence σ_h is formed by concatenating $\lceil 8\alpha_{h-1} \rceil$ iterations, where each iteration is chosen uniformly at random to be of type 1 or type 2. If the last iteration is of type 2, an additional final request at h is issued. Thus, by induction, $(0, \dots, h)$ is both a prefix and a suffix of σ_h .

We next show (6) under the assumption that at the start of each iteration, the iteration is of type 1 or 2 each with probability $1/2$ even when conditioned on the knowledge of the semi-online algorithm at that time. We will later show how to design the arrival times of individual requests so that this assumption is satisfied. We begin by proving the first inequality of (6). We claim that in each iteration of σ_h , the expected cost of any semi-online algorithm (restricted to staying at the leaves $0, \dots, h$) is at least $\tau^h/2$. Indeed, if the evader starts the iteration at leaf h , then with probability $1/2$ we have a type 1 iteration forcing the evader to vacate leaf h for cost τ^h , giving an expected cost of $\tau^h/2$. If the evader is at one of the leaves $0, \dots, h-1$, then with probability $1/2$ we have a type 2 iteration. In this case, it must either move to h for cost at least τ^h , or $\lceil \frac{\tau^h}{\alpha_{h-1} \text{OPT}_{h-1}} \rceil$ times it suffers expected cost at least $\alpha_{h-1} \text{OPT}_{h-1}$ by the induction hypothesis. So again, the expected cost is at least $\tau^h/2$. Since σ_h consists of $\lceil 8\alpha_{h-1} \rceil$ iterations, we have

$$A_h \geq 4\alpha_{h-1}\tau^h,$$

giving the first inequality of (6).

To show the second inequality of (6), we describe an offline strategy. With probability $2^{-\lceil 8\alpha_{h-1} \rceil}$, all iterations of σ_h are of type 2. In this case, the offline evader moves to leaf h at the beginning of σ_h and back to leaf 0 upon the one request to h at the end of σ_h , for total cost $2(1 + \tau^h)$. With the remaining probability, there is at least one type 1 iteration. Conditioned on this being the case, the expected number of type 1 iterations is $\lceil 8\alpha_{h-1} + 1 \rceil / 2$, and the expected number of type 2 iterations is $\lceil 8\alpha_{h-1} - 1 \rceil / 2$. The offline evader can serve each type 1 iteration for cost $2(1 + \tau)$ and each type 2 iteration for expected cost $\lceil \frac{\tau^h}{\alpha_{h-1} \text{OPT}_{h-1}} \rceil \text{OPT}_{h-1}$, and it finishes each iteration at leaf 0. (Thus, if the last iteration is of type 2, then the final request to h incurs no additional cost.) By the induction hypothesis, $\text{OPT}_{h-1} \leq O(\tau^{h-1})$. Hence, we can rewrite the expected cost of a type 2 iteration as

$$\left\lceil \frac{\tau^h}{\alpha_{h-1} \text{OPT}_{h-1}} \right\rceil \text{OPT}_{h-1} = (1 + o(1)) \frac{\tau^h}{\alpha_{h-1}},$$

as $\tau \rightarrow \infty$. Since $h \geq 2$, the expected cost of all type 1 iterations is only an $o(1)$ fraction of the expected cost of the type 2 iterations. Overall, we get

$$\begin{aligned}
\text{OPT}_h &\leq 2^{-\lceil 8\alpha_{h-1} \rceil} 2(1 + \tau^h) + \\
&\quad (1 + o(1)) \left(1 - 2^{-\lceil 8\alpha_{h-1} \rceil}\right) \frac{\lceil 8\alpha_{h-1} - 1 \rceil}{2} \frac{\tau^h}{\alpha_{h-1}} \\
&\leq (1 + o(1)) \left[2^{-\lceil 8\alpha_{h-1} \rceil} 2\tau^h + \left(1 - 2^{-\lceil 8\alpha_{h-1} \rceil}\right) 4\tau^h \right] \\
&= (1 + o(1)) \left(1 - 2^{-\lceil 8\alpha_{h-1} \rceil - 1}\right) 4\tau^h \\
&\leq \frac{4\tau^h}{1 + 2^{-\lceil 8\alpha_{h-1} \rceil - 1}}.
\end{aligned}$$

We obtain the second inequality in (6) by

$$\begin{aligned}
\frac{4\alpha_{h-1}\tau^h}{\text{OPT}_h} &\geq \alpha_{h-1} \left(1 + 2^{-\lceil 8\alpha_{h-1} \rceil - 1}\right) \\
&\geq \frac{1}{2} + \beta \log(h-1) + 2^{-8\beta \log(h-1) - 7} \\
&\geq \frac{1}{2} + \beta \log(h-1) + \frac{\beta}{h-1} \\
&\geq \frac{1}{2} + \beta \log h \\
&= \alpha_h,
\end{aligned}$$

where the third inequality holds for $\beta = 2^{-7}$.

It remains to define the arrival times for the requests of sequence σ_h within the interval (a, b) . We do this as follows: Let $m \geq 1$ be the number of requests to leaf h in σ_h . These requests to h will be issued at times $a + (b-a) \sum_{i=1}^j 2^{-i}$ for $j = 1, \dots, m$.

To define the arrival times of the other requests, we will maintain a time variable $c \in [a, b]$ indicating the *current* time, and a variable $n > c$ indicating the time of the *next* request to leaf h after time c . Initially, $c := a$ and $n := (a+b)/2$. Consider the first iteration for which the arrival times have not been defined yet. If the iteration is of type 2, we choose the arrival times according to the induction hypothesis so that all subsequences σ_{h-1} within the iteration fit into the time window $(c, (c+n)/2)$, and we update $c := (c+n)/2$. If the iteration is of type 1, sample a type 2 iteration I and let t_1, \dots, t_{h-1} be such that t_i would be the time of the next request to page i if the next iteration were this iteration I of type 2 instead of a type 1 iteration. We define the arrival times of the (single) request to leaf $i < h$ in this type 1 iteration to be t_i . If this was not the last iteration, we update $c := n$ and increase n to the time of the next request to h (as defined above).

Notice that at the beginning of each iteration within σ_h , ordering the pages by the time of their next request always yields the sequence $0, 1, \dots, k$, and the time of the next request to each page is independent of whether the next iteration is of type 1 or type 2. Thus, as promised, whether the next iteration is of type 1 or type 2 is independent of the knowledge of the semi-online algorithm. \square