

# Online Metric Algorithms with Untrusted Predictions

Antonios Antoniadis<sup>1</sup>   Christian Coester<sup>2</sup>   Marek Elias<sup>3</sup>  
Adam Polak<sup>4</sup>   **Bertrand Simon<sup>5</sup>**

- 1: MPI, Saarbrücken (Germany).
- 2: CWI, Amsterdam (Netherlands).
- 3: EPFL, Lausanne (Switzerland).
- 4: Jagiellonian University, Kraków (Poland).
- 5: University of Bremen (Germany).

Dagstuhl – February 2020

# Motivation


## Online algorithms

- ▶ Guaranteed competitive ratio
- ▶ Bad performance on easy instances, overly pessimistic



## Machine Learning predictions

- ▶ Often relevant information
- ▶ No guarantee, can be arbitrarily bad



## Prediction-augmented algorithms

- ▶ Target competitive ratio:  $O(\min\{1 + f(\eta/OPT), ONLINE\})$


## Online algorithms

- ▶ Guaranteed competitive ratio
- ▶ Bad performance on easy instances, overly pessimistic



## Machine Learning predictions

- ▶ Often relevant information
- ▶ No guarantee, can be arbitrarily bad



## Prediction-augmented algorithms

- ▶ Target competitive ratio:  $O(\min\{1 + f(\eta/OPT), ONLINE\})$


## Online algorithms

- ▶ Guaranteed competitive ratio
- ▶ Bad performance on easy instances, overly pessimistic



## Machine Learning predictions

- ▶ Often relevant information
- ▶ No guarantee, can be arbitrarily bad



## Prediction-augmented algorithms

- ▶ Target competitive ratio:  $O(\min\{1 + f(\eta/OPT), ONLINE\})$

## Some previously studied problems

- ▶ Ski rental: predict #days we will ski [KumarPS'18]
- ▶ Non-clairvoyant scheduling: predict processing times [KumarPS'18]
- ▶ Restricted assignment: predict machine weights [LattanziMLV'20]
- ▶ Caching: predict next arrival time [LykourisV'18, Rohatgi'20]

Issue: lack of generality, predictions tailored to specific problems


### Lemma

*Previously used caching predictions are not useful for weighted caching.*

⇒ need for a *more general* prediction setup

# The Metrical Task System (MTS) problem

## Definition by picture



Round 0

Cost incurred: 0


Note: generalizes caching,  $k$ -server, convex body chasing...

## Prediction setup

- ▶ Fix OFF: offline algorithm (e.g., OPT), who goes to states  $o_1, o_2, \dots$
- ▶ At time  $t$ ,  $p_t :=$  prediction of  $o_t$ . Error:  $\eta = \sum_t d(o_t, p_t)$

# The Metrical Task System (MTS) problem

## Definition by picture



Round 1 before serving  
Cost incurred: 0


Note: generalizes caching,  $k$ -server, convex body chasing...

## Prediction setup

- ▶ Fix OFF: offline algorithm (e.g., OPT), who goes to states  $o_1, o_2, \dots$
- ▶ At time  $t$ ,  $p_t :=$  prediction of  $o_t$ . Error:  $\eta = \sum_t d(o_t, p_t)$

# The Metrical Task System (MTS) problem

## Definition by picture



Round 1 after serving  
Cost incurred:  $3+1$

Note: generalizes caching,  $k$ -server, convex body chasing...


## Prediction setup

- ▶ Fix OFF: offline algorithm (e.g., OPT), who goes to states  $o_1, o_2, \dots$
- ▶ At time  $t$ ,  $p_t :=$  prediction of  $o_t$ . Error:  $\eta = \sum_t d(o_t, p_t)$



# The Metrical Task System (MTS) problem

## Definition by picture



Round 2 before serving  
Cost incurred:  $3+1$


Note: generalizes caching,  $k$ -server, convex body chasing...

## Prediction setup

- ▶ Fix OFF: offline algorithm (e.g., OPT), who goes to states  $o_1, o_2, \dots$
- ▶ At time  $t$ ,  $p_t :=$  prediction of  $o_t$ . Error:  $\eta = \sum_t d(o_t, p_t)$

# The Metrical Task System (MTS) problem

## Definition by picture



Round 2 after serving  
Cost incurred:  $3+1+7+2$

Note: generalizes caching,  $k$ -server, convex body chasing...

## Prediction setup

- ▶ Fix OFF: offline algorithm (e.g., OPT), who goes to states  $o_1, o_2, \dots$
- ▶ At time  $t$ ,  $p_t :=$  prediction of  $o_t$ . Error:  $\eta = \sum_t d(o_t, p_t)$

## Algorithm subroutine (FTP, Follow the Prediction)

- ▶ Go to  $p_t$ , except if there is a cheap state nearby
- ▶ Proposition: this costs at most  $\text{OFF} \cdot (1 + 4\eta/\text{OFF})$

## Combining online algorithms A and B: $\text{comb}(A, B)$ [BlumB'00]

- ▶  $E(\text{cost}_{\text{comb}(A,B)}) \leq (1 + \varepsilon) \cdot \min\{E(\text{cost}_A), E(\text{cost}_B)\}$  asymptotically

Theorem ( $\text{ROBUSTFTP} := \text{comb}(\text{ONLINE}, \text{FTP})$  competitive ratio)

$\text{ROBUSTFTP}$  is  $O(\min\{1 + \eta/\text{OFF}, \text{ONLINE}\})$  - competitive.

(Recall:  $\eta = \sum_t d(o_t, p_t)$ )

## Lemma (Lower bound)

This is tight for some MTS (i.e.,  $\eta/\text{OFF}$  - dependency).

Issue with **ROBUSTFTP**: too general, not the best for all MTS

## Focus on a specific MTS: the caching problem

- ▶ Maintain a cache of  $k$  pages, pay 1 per cache miss

Algorithm **TRUST&DOUBT**: explaining how it works its name

- ▶ “Trust” predictor: evict a page not in its cache
- ▶ If an evicted page requested: “doubt” this decision
- ▶ Regularly (depending on trustworthiness): “trust” again

Theorem (**TRUST&DOUBT** competitive ratio)


**TRUST&DOUBT** costs at most  $O(\text{OFF} \cdot \min\{1 + \log \frac{\eta}{\text{OFF}}, \log k\})$ .

# Comparison to previous caching algorithms

## How do our algorithms compare to previous ones?

- ▶ Different errors  $\rightarrow$  difficult to compare competitive ratios
- ▶ Experimentally: compute previous predictions  
deduct our predictions (evict furthest predicted)

## Results on a public dataset (BrightKite, $k = 10$ ) – (lower is better)



LRU	1.291	
Marker	1.333	
<i>Predictions</i>	PLECO	POPU
FtP	2.081	1.707
L&V [LykourisV'18]	1.340	1.262
LMarker [Rohatgi'20]	1.337	1.264
LNonMarker [Rohatgi'20]	1.339	1.292
<b>ROBUSTFtP</b>	<b>1.351</b>	<b>1.316</b>
<b>TRUST&amp;DOUBT</b>	<b>1.292</b>	<b>1.274</b>

Prediction: ground truth + lognorm error


Two predictors: simple statistics

# Comparison to previous caching algorithms

## How do our algorithms compare to previous ones?

- ▶ Different errors  $\rightarrow$  difficult to compare competitive ratios
- ▶ Experimentally: compute previous predictions  
deduct our predictions (evict furthest predicted)

## Results on a public dataset (BrightKite, $k = 10$ ) – (lower is better)



LRU	1.291	
Marker	1.333	
<i>Predictions</i>	PLECO	POPU
FtP	2.081	1.707
L&V [LykourisV'18]	1.340	1.262
LMarker [Rohatgi'20]	1.337	1.264
LNonMarker [Rohatgi'20]	1.339	1.292
<b>ROBUSTFtP</b>	<b>1.351</b>	<b>1.316</b>
<b>TRUST&amp;DOUBT</b>	<b>1.292</b>	<b>1.274</b>

Prediction: ground truth + lognorm error


Two predictors: simple statistics

# Comparison to previous caching algorithms

## How do our algorithms compare to previous ones?

- ▶ Different errors  $\rightarrow$  difficult to compare competitive ratios
- ▶ Experimentally: compute previous predictions  
deduct our predictions (evict furthest predicted)

## Results on a public dataset (BrightKite, $k = 10$ ) – (lower is better)



LRU	1.291	
Marker	1.333	
<i>Predictions</i>	PLECO	POPU
FtP	2.081	1.707
L&V [LykourisV'18]	1.340	1.262
LMarker [Rohatgi'20]	1.337	1.264
LNonMarker [Rohatgi'20]	1.339	1.292
<b>ROBUSTFtP</b>	<b>1.351</b>	<b>1.316</b>
<b>TRUST&amp;DOUBT</b>	<b>1.292</b>	<b>1.274</b>

Prediction: ground truth + lognorm error


Two predictors: simple statistics

# Comparison to previous caching algorithms

## How do our algorithms compare to previous ones?

- ▶ Different errors  $\rightarrow$  difficult to compare competitive ratios
- ▶ Experimentally: compute previous predictions  
deduct our predictions (evict furthest predicted)

## Results on a public dataset (BrightKite, $k = 10$ ) – (lower is better)



LRU	1.291	
Marker	1.333	
<i>Predictions</i>	PLECO	POPU
FtP	2.081	1.707
L&V [LykourisV'18]	1.340	1.262
LMarker [Rohatgi'20]	1.337	1.264
LNonMarker [Rohatgi'20]	1.339	1.292
<b>ROBUSTFTP</b>	<b>1.351</b>	<b>1.316</b>
<b>TRUST&amp;DOUBT</b>	<b>1.292</b>	<b>1.274</b>

Prediction: ground truth + lognorm error

Two predictors: simple statistics





is useful !

- ▶ MTS (and beyond): prediction setup and “optimal” algorithm
- ▶ Caching: specific algorithm good in theory & practice
- ▶ Perspective: other MTS (weighted caching, convex body chasing,  $k$ -server)