

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

## Online algorithms

- ▶ Optimization with incomplete information
- ▶ Guaranteed competitive ratio:  
 $r$  s.t.  $\forall I: \text{ALG}(I) \leq r \cdot \text{OPT}(I)$
- ▶ Bad performance on easy instances, overly pessimistic



## Online algorithms

- ▶ Optimization with incomplete information
- ▶ Guaranteed competitive ratio:  
 $r$  s.t.  $\forall I: \text{ALG}(I) \leq r \cdot \text{OPT}(I)$
- ▶ Bad performance on easy instances, overly pessimistic



## Online algorithms

- ▶ Optimization with incomplete information
- ▶ Guaranteed competitive ratio:  
 $r$  s.t.  $\forall I: \text{ALG}(I) \leq r \cdot \text{OPT}(I)$
- ▶ Bad performance on easy instances, overly pessimistic



## Machine learning predictions

- ▶ Often very powerful



- ▶ No guarantee, can be arbitrarily bad

## Online algorithms

- ▶ Optimization with incomplete information
- ▶ Guaranteed competitive ratio:  
 $r$  s.t.  $\forall I: \text{ALG}(I) \leq r \cdot \text{OPT}(I)$
- ▶ Bad performance on easy instances, overly pessimistic



## Machine learning predictions

- ▶ Often very powerful



- ▶ No guarantee, can be arbitrarily bad
- ▶ prediction error  $\eta$

# Motivation

## Online algorithms

- ▶ Optimization with incomplete information
- ▶ Guaranteed competitive ratio:  
 $r \text{ s.t. } \forall I: \text{ALG}(I) < r \cdot \text{OPT}(I)$
- ▶ Bad performance on easy instances, overly pessimistic



## Machine learning predictions

- ▶ Often very powerful



- ▶ No guarantee, can be arbitrarily bad
- ▶ prediction error  $\eta$

## Prediction-augmented algorithms

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

## Some previously studied problems

- ▶ Ski rental: predict #days we will ski [PurohitSK'18]
- ▶ Non-clairvoyant scheduling: predict processing times [PurohitSK'18]
- ▶ Restricted assignment: predict machine weights [LattanziLMV'20]
- ▶ Caching: predict next arrival time [LykourisV'18, Rohatgi'20, Wei'20]
- ▶ Weighted caching: predict *all* requests until next arrival [JiangPS'20]

Issue: lack of generality, predictions tailored to specific problems

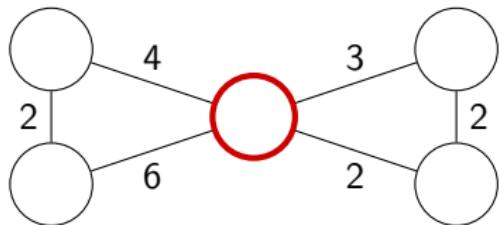
Proposition (Antoniadis, Coester, Elias, Polak, Simon; ICML'20)

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

⇒ need *more general* prediction setup

# Metrical Task Systems (MTS)

- ▶ Given: metric space  $(M, d)$ ,  $x_0 \in M$
- ▶ At time  $t$ :
  1. cost function  $\ell_t : M \mapsto \mathbb{R}_+$  revealed
  2. algo chooses  $x_t \in M$
  3. pays  $d(x_{t-1}, x_t) + \ell_t(x_t)$

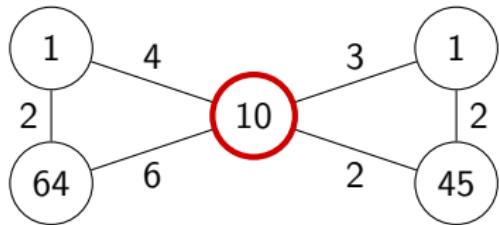


Cost: 0

MTS generalizes caching,  $k$ -server, convex body/function chasing, layered graph traversal...

# Metrical Task Systems (MTS)

- ▶ Given: metric space  $(M, d)$ ,  $x_0 \in M$
- ▶ At time  $t$ :
  1. cost function  $\ell_t : M \mapsto \mathbb{R}_+$  revealed
  2. algo chooses  $x_t \in M$
  3. pays  $d(x_{t-1}, x_t) + \ell_t(x_t)$

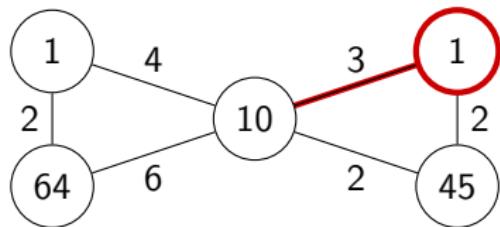


Cost: 0

MTS generalizes caching,  $k$ -server, convex body/function chasing, layered graph traversal...

# Metrical Task Systems (MTS)

- ▶ Given: metric space  $(M, d)$ ,  $x_0 \in M$
- ▶ At time  $t$ :
  1. cost function  $\ell_t : M \mapsto \mathbb{R}_+$  revealed
  2. algo chooses  $x_t \in M$
  3. pays  $d(x_{t-1}, x_t) + \ell_t(x_t)$

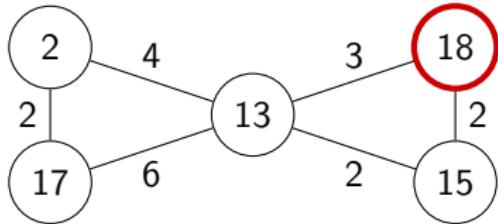


Cost: 3+1

MTS generalizes caching,  $k$ -server, convex body/function chasing, layered graph traversal...

# Metrical Task Systems (MTS)

- ▶ Given: metric space  $(M, d)$ ,  $x_0 \in M$
- ▶ At time  $t$ :
  1. cost function  $\ell_t : M \mapsto \mathbb{R}_+$  revealed
  2. algo chooses  $x_t \in M$
  3. pays  $d(x_{t-1}, x_t) + \ell_t(x_t)$

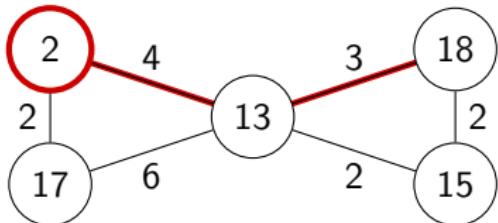


Cost: 3+1

MTS generalizes caching,  $k$ -server, convex body/function chasing, layered graph traversal...

# Metrical Task Systems (MTS)

- ▶ Given: metric space  $(M, d)$ ,  $x_0 \in M$
- ▶ At time  $t$ :
  1. cost function  $\ell_t : M \mapsto \mathbb{R}_+$  revealed
  2. algo chooses  $x_t \in M$
  3. pays  $d(x_{t-1}, x_t) + \ell_t(x_t)$

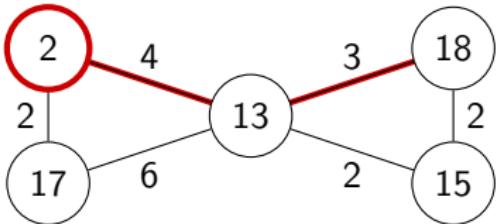


Cost:  $3+1+\textcolor{red}{7}+2$

MTS generalizes caching,  $k$ -server, convex body/function chasing, layered graph traversal...

# Metrical Task Systems (MTS)

- ▶ Given: metric space  $(M, d)$ ,  $x_0 \in M$
- ▶ At time  $t$ :
  1. cost function  $\ell_t : M \mapsto \mathbb{R}_+$  revealed
  2. algo chooses  $x_t \in M$
  3. pays  $d(x_{t-1}, x_t) + \ell_t(x_t)$



Cost:  $3+1+\textcolor{red}{7}+2$

MTS generalizes caching,  $k$ -server, convex body/function chasing, layered graph traversal...

## What to predict?

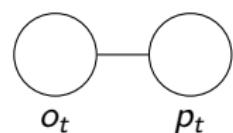
- ▶ Next costs?
  - only  $\ell_{t+1}$ ? useless with dummy rounds
  - entire future? too much information
- ▶ **Single state per round:** recommendation for  $x_t$

# Definition of the error and a simple algorithm

Definition: error  $\eta$

- ▶ 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:  $\eta := \sum_t d(o_t, p_t)$

Goal: Small cost (rel. to OFF) if  $\eta$  small



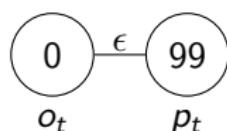
# Definition of the error and a simple algorithm

Definition: error  $\eta$

- ▶ 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:  $\eta := \sum_t d(o_t, p_t)$

Goal: Small cost (rel. to OFF) if  $\eta$  small

- ▶ Naive algo:  $x_t := p_t$



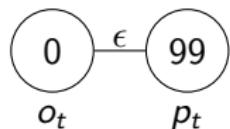
# Definition of the error and a simple algorithm

Definition: error  $\eta$

- ▶ 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:  $\eta := \sum_t d(o_t, p_t)$

Goal: Small cost (rel. to OFF) if  $\eta$  small

- ▶ Naive algo:  $x_t := p_t$
- ▶ Better:  $x_t := \arg \min_{x \in X} \{cost_t(x) + 2d(p_t, x)\}$
- ▶ Call this algo FTP (Follow the Prediction)



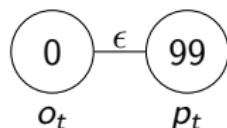
# Definition of the error and a simple algorithm

Definition: error  $\eta$

- ▶ 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:  $\eta := \sum_t d(o_t, p_t)$

Goal: Small cost (rel. to OFF) if  $\eta$  small

- ▶ Naive algo:  $x_t := p_t$
- ▶ Better:  $x_t := \arg \min_{x \in X} \{cost_t(x) + 2d(p_t, x)\}$
- ▶ Call this algo FTP (Follow the Prediction)



Lemma (Antoniadis, Coester, Elias, Polak, Simon; ICML'20)

FTP has cost  $\leq$  OFF +  $4\eta$ . So competitive ratio vs OFF is  $1 + 4\eta/\text{OFF}$ .

Guarantee holds simultaneously for all offline algos ( $\eta$  depending on it)

Issue: FTP not robust

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

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

Let ROBUSTFTP :=  $\text{comb}(\text{ONLINE}, \text{FTP})$

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

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

Let ROBUSTFTP :=  $\text{comb}(\text{ONLINE}, \text{FTP})$

Theorem (Antoniadis, Coester, Elias, Polak, Simon; ICML'20)

ROBUSTFTP has cost  $O(\min\{\text{OFF} + \eta, \text{ONLINE}\})$ .

(Recall:  $\eta = \text{prediction error wrt. OFF}$ )

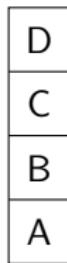
This is asymptotically optimal.

- ▶ Lower bound holds for **some** MTS
- ▶ For caching (a special case of MTS), we can do better

# Caching (aka Paging)

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

$k = 4$



E

misses: 1

# Caching (aka Paging)

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

$k = 4$



E A

misses: 1

# Caching (aka Paging)

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

$k = 4$



E A F

misses: 2

# Caching (aka Paging)

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

$k = 4$



E A F B

misses: 3

# Caching (aka Paging)

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

$k = 4$



E A F B E

misses: 3

# Caching (aka Paging)

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

$k = 4$



E A F B E B

misses: 3

# Caching (aka Paging)

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

$k = 4$



E A F B E B C

misses: 4

# Caching (aka Paging)

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

$k = 4$



E A F B E B C F

misses: 4

# Caching (aka Paging)

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



- Predicted state: set of  $k$  pages  
(intuition: recommended cache content)

## Algorithm TRUST&DOUBT: idea

- ▶ Balance 2 competing policies:
  - “Trust”: Evict what is evicted from predicted cache
  - “Doubt”: Evict according to classical policy
- ▶ In a phase, first follow “Trust”
- ▶ If this turns out bad, switch to “Doubt”
- ▶ Regularly (depending on trustworthiness): “Trust” again

Theorem (Antoniadis, Coester, Elias, Polak, Simon; ICML'20)

TRUST&DOUBT *is*  $\min \left\{ O(\log k), O(\log \frac{\eta}{O_{\text{PT}}}) \right\}$ -competitive.

## Algorithm TRUST&DOUBT: idea

- ▶ Balance 2 competing policies:
  - “Trust”: Evict what is evicted from predicted cache
  - “Doubt”: Evict according to classical policy
- ▶ In a phase, first follow “Trust”
- ▶ If this turns out bad, switch to “Doubt”
- ▶ Regularly (depending on trustworthiness): “Trust” again

Theorem (Antoniadis, Coester, Elias, Polak, Simon; ICML'20)

TRUST&DOUBT *is*  $\min \left\{ O(\log k), O(\log \frac{\eta}{\text{OPT}}) \right\}$ -competitive.

## Algorithm TRUST&DOUBT: idea

- ▶ Balance 2 competing policies:
  - “Trust”: Evict what is evicted from predicted cache
  - “Doubt”: Evict according to classical policy
- ▶ In a phase, first follow “Trust”
- ▶ If this turns out bad, switch to “Doubt”
- ▶ Regularly (depending on trustworthiness): “Trust” again

Theorem (Antoniadis, Coester, Elias, Polak, Simon; ICML'20)

TRUST&DOUBT *is*  $\min \{ O(\log k), O(\log \frac{\eta}{\text{OFF}}) \}$ -competitive vs OFF.

# Comparison with previous prediction setup for caching

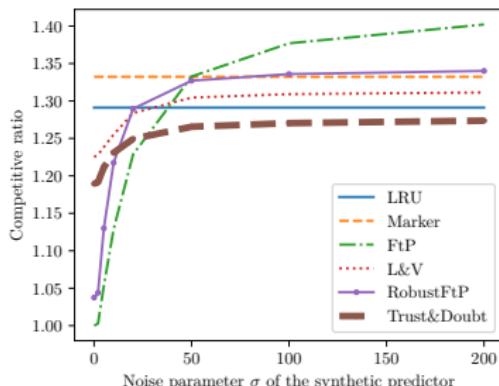
- ▶ Our setup: Prediction = recommended cache content
- ▶ In [LykourisV'18, Rohatgi'20, Wei'20]: Prediction = next time of each request

## How do setups/results compare?

- ▶ Can convert request time predictions to cache content predictions
  - Predictor evicts furthest predicted page
- ▶ Prediction error  $\eta$  incomparable
  - Nonetheless: TRUST&DOUBT also achieves guarantee of [Wei'20]
- ▶ Succinct: Instead of  $\log T$  bit predictions per request, only  $\log k$  bits when predictor has cache miss
- ▶ Our setup generalizes to MTS

# Experiments

Same datasets as [LykourisV'18]



Prediction: ground truth + lognorm error

Dataset	BK		Citi	
	PLECO	POPU	PLECO	POPU
LRU	1.29		1.85	
MARKER	1.33		1.86	
Predictions	PLECO	POPU	PLECO	POPU
	1.34	1.26	1.88	1.78
L&V [LykourisV'18]	1.34	1.26	1.88	1.78
LMarker [Rohatgi'20]	1.34	1.29	1.88	1.80
LNonMarker [Rohatgi'20]	1.34	1.26	1.88	1.78
<b>ROBUSTFtP</b>	<b>1.35</b>	<b>1.32</b>	<b>1.89</b>	<b>1.83</b>
<b>TRUST&amp;DOUBT</b>	<b>1.29</b>	<b>1.27</b>	<b>1.85</b>	<b>1.77</b>

Two predictors: simple statistics

**Definition by picture:  $n$  servers;  $n$  online requests**



## Analogous prediction setup

- ▶ Predict set of servers to be matched so far
- ▶ Error: cost of matching between these and OFF's matched servers

Theorem (Antoniadis, Coester, Elias, Polak, Simon; ICML'20)

ROBUSTFTP's analogue is  $O(\min\{1 + \frac{\eta}{\text{OFF}}, \log n\})$ -competitive vs OFF.

**Definition by picture:  $n$  servers;  $n$  online requests**



## Analogous prediction setup

- ▶ Predict set of servers to be matched so far
- ▶ Error: cost of matching between these and OFF's matched servers

Theorem (Antoniadis, Coester, Elias, Polak, Simon; ICML'20)

ROBUSTFTP's analogue is  $O(\min\{1 + \frac{\eta}{\text{OFF}}, \log n\})$ -competitive vs OFF.

**Definition by picture:  $n$  servers;  $n$  online requests**



## Analogous prediction setup

- ▶ Predict set of servers to be matched so far
- ▶ Error: cost of matching between these and OFF's matched servers

Theorem (Antoniadis, Coester, Elias, Polak, Simon; ICML'20)

ROBUSTFTP's analogue is  $O(\min\{1 + \frac{\eta}{\text{OFF}}, \log n\})$ -competitive vs OFF.

**Definition by picture:  $n$  servers;  $n$  online requests**



## Analogous prediction setup

- ▶ Predict set of servers to be matched so far
- ▶ Error: cost of matching between these and OFF's matched servers

Theorem (Antoniadis, Coester, Elias, Polak, Simon; ICML'20)

ROBUSTFTP's analogue is  $O(\min\{1 + \frac{\eta}{\text{OFF}}, \log n\})$ -competitive vs OFF.

**Definition by picture:  $n$  servers;  $n$  online requests**



## Analogous prediction setup

- ▶ Predict set of servers to be matched so far
- ▶ Error: cost of matching between these and OFF's matched servers

Theorem (Antoniadis, Coester, Elias, Polak, Simon; ICML'20)

ROBUSTFTP's analogue is  $O(\min\{1 + \frac{\eta}{\text{OFF}}, \log n\})$ -competitive vs OFF.

**Definition by picture:  $n$  servers;  $n$  online requests**



## Analogous prediction setup

- ▶ Predict set of servers to be matched so far
- ▶ Error: cost of matching between these and OFF's matched servers

Theorem (Antoniadis, Coester, Elias, Polak, Simon; ICML'20)

ROBUSTFTP's analogue is  $O(\min\{1 + \frac{\eta}{\text{OFF}}, \log n\})$ -competitive vs OFF.

**Definition by picture:  $n$  servers;  $n$  online requests**



## Analogous prediction setup

- ▶ Predict set of servers to be matched so far
- ▶ Error: cost of matching between these and OFF's matched servers

Theorem (Antoniadis, Coester, Elias, Polak, Simon; ICML'20)

ROBUSTFTP's analogue is  $O(\min\{1 + \frac{\eta}{\text{OFF}}, \log n\})$ -competitive vs OFF.

**Definition by picture:  $n$  servers;  $n$  online requests**



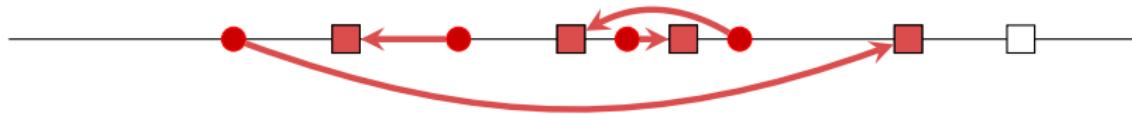
## Analogous prediction setup

- ▶ Predict set of servers to be matched so far
- ▶ Error: cost of matching between these and OFF's matched servers

Theorem (Antoniadis, Coester, Elias, Polak, Simon; ICML'20)

ROBUSTFTP's analogue is  $O(\min\{1 + \frac{\eta}{\text{OFF}}, \log n\})$ -competitive vs OFF.

**Definition by picture:  $n$  servers;  $n$  online requests**



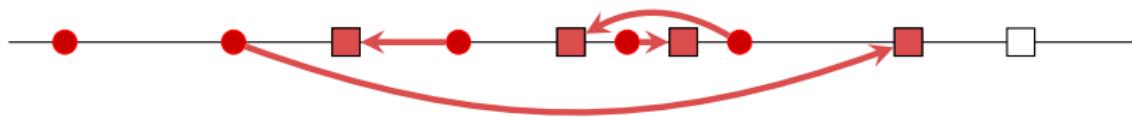
## Analogous prediction setup

- ▶ Predict set of servers to be matched so far
- ▶ Error: cost of matching between these and OFF's matched servers

Theorem (Antoniadis, Coester, Elias, Polak, Simon; ICML'20)

ROBUSTFTP's analogue is  $O(\min\{1 + \frac{\eta}{\text{OFF}}, \log n\})$ -competitive vs OFF.

Definition by picture:  $n$  servers;  $n$  online requests



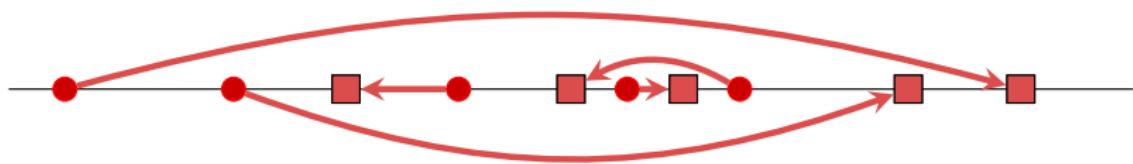
## Analogous prediction setup

- ▶ Predict set of servers to be matched so far
- ▶ Error: cost of matching between these and OFF's matched servers

Theorem (Antoniadis, Coester, Elias, Polak, Simon; ICML'20)

ROBUSTFTP's analogue is  $O(\min\{1 + \frac{\eta}{\text{OFF}}, \log n\})$ -competitive vs OFF.

Definition by picture:  $n$  servers;  $n$  online requests



## Analogous prediction setup

- ▶ Predict set of servers to be matched so far
- ▶ Error: cost of matching between these and OFF's matched servers

Theorem (Antoniadis, Coester, Elias, Polak, Simon; ICML'20)

ROBUSTFTP's analogue is  $O(\min\{1 + \frac{\eta}{\text{OFF}}, \log n\})$ -competitive vs OFF.



## MTS (and beyond)

- ▶ General prediction setup
- ▶ “Optimal” algorithm: ROBUSTFTP

## Caching

- ▶ Better algorithm: TRUST&DOUBT
- ▶ Less predicted information and better empirical performance than caching-specific setup
- ▶ Better than LRU with simple predictor

## Perspectives

- ▶ Better algorithms for other MTS?  
(e.g., weighted caching,  $k$ -server, convex body chasing, graph traversal)