Mixing Predictions for Online Metric Algorithms

Antonios Antoniadis   Christian Coester   Marek Eliáš   Adam Polak   Bertrand Simon

You're solving your favorite online problem…

that can be formulated as a Metrical Task System

e.g., k-server, caching, convex body chasing

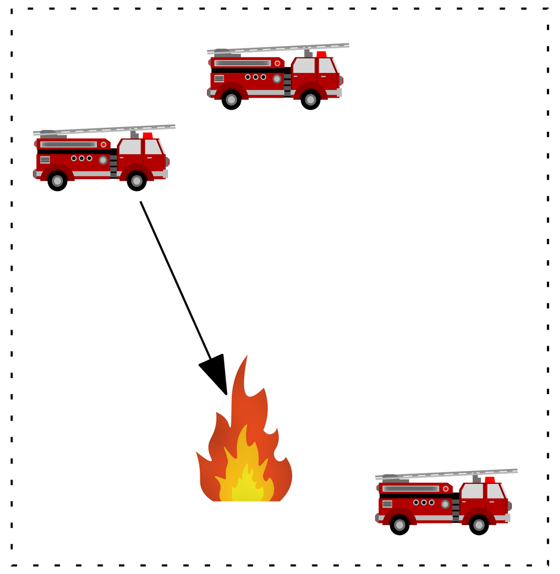

…and there are ℓ predictors…

telling you in each round what action each of them would take.

Think of, e.g., trained machine-learning models,
or a classical algorithm with worst-case guarantees.

Who do you follow to optimize your outcome?

Blum&Burch [COLT 1997]

You can achieve cost

$\leqslant (1+\varepsilon) \cdot$ **single best predictor in hindsight** + const

(const: depends on ε and the diameter of the metric space)

Main idea: reduction to online learning with experts

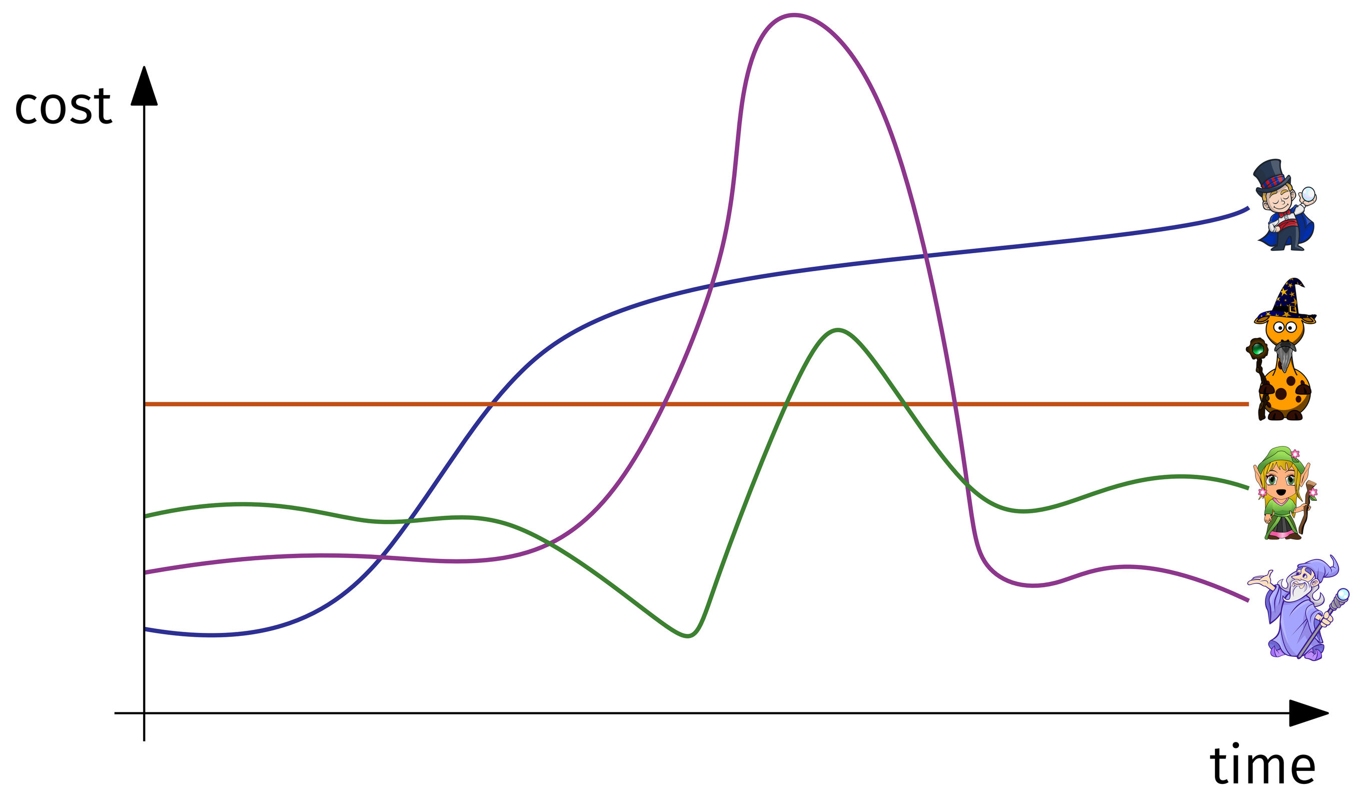

…but what if the best predictor changes over time …

Our result #1

You can achieve cost $\leqslant O(\ell^2) \cdot$ **best combination of predictors**

(best combination of predictors: optimal in hindsight choice of a possibly different predictor in each round)

Main idea: reduction to Layered Graph Traversal
+ Bubeck-Coester-Rabani LGT algorithm [FOCS 2022]

Our result #2

You can achieve cost

$\leqslant (1+\varepsilon)^2 \cdot$ **best up-to-M-switches combination** + const

$M \approx \frac{\varepsilon^2}{\log \ell} \cdot \frac{\text{OPT}}{\text{diam}}$

Main idea: reduction to unfair MTS on uniform metric
+ r-unfair competitive algorithm of Bartal et al. [STOC 1997]

Our result #3

With bandit access to predictors, you can achieve cost

$\leqslant (1+\varepsilon)^3 \cdot$ **best up-to-M-switches combination** + const

$M \approx \frac{\varepsilon^3}{\ell \log \ell \log(2+\varepsilon^{-1})} \cdot \frac{\text{OPT}}{\text{diam}}$

(bandit access: in each round you have to chose only one predictor to hear from)

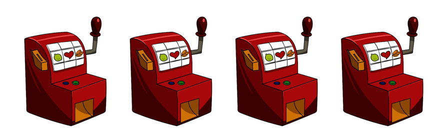

Main idea: $P(\text{explore}) = \varepsilon,\ P(\text{exploit}) = 1-\varepsilon$

UNIVERSITY OF TWENTE.   UNIVERSITY OF OXFORD   Bocconi   max planck institut informatik   CCIN2P3