Decision-Theoretic Approaches for Improved Learning-Augmented Algorithms

Spyros Angelopoulos, Christoph Dürr, Georgii Melidi

The problem

Among infinitely many learning-augmented algorithms, how do we **choose** the best one?

Learning-Augmented Algorithms leverage machine-learned predictions about the input to improve performance beyond classical worst-case guarantees.

Evaluation traditionally focuses on

Consistency: performance with perfect predictions

Robustness: performance under adversarial predictions

The Challenge

Many algorithms achieve different tradeoffs between consistency and robustness.

However:

- infinitely many Pareto-optimal algorithms exist
- their behavior under prediction error differs significantly

➡ **Which algorithm should we choose?**

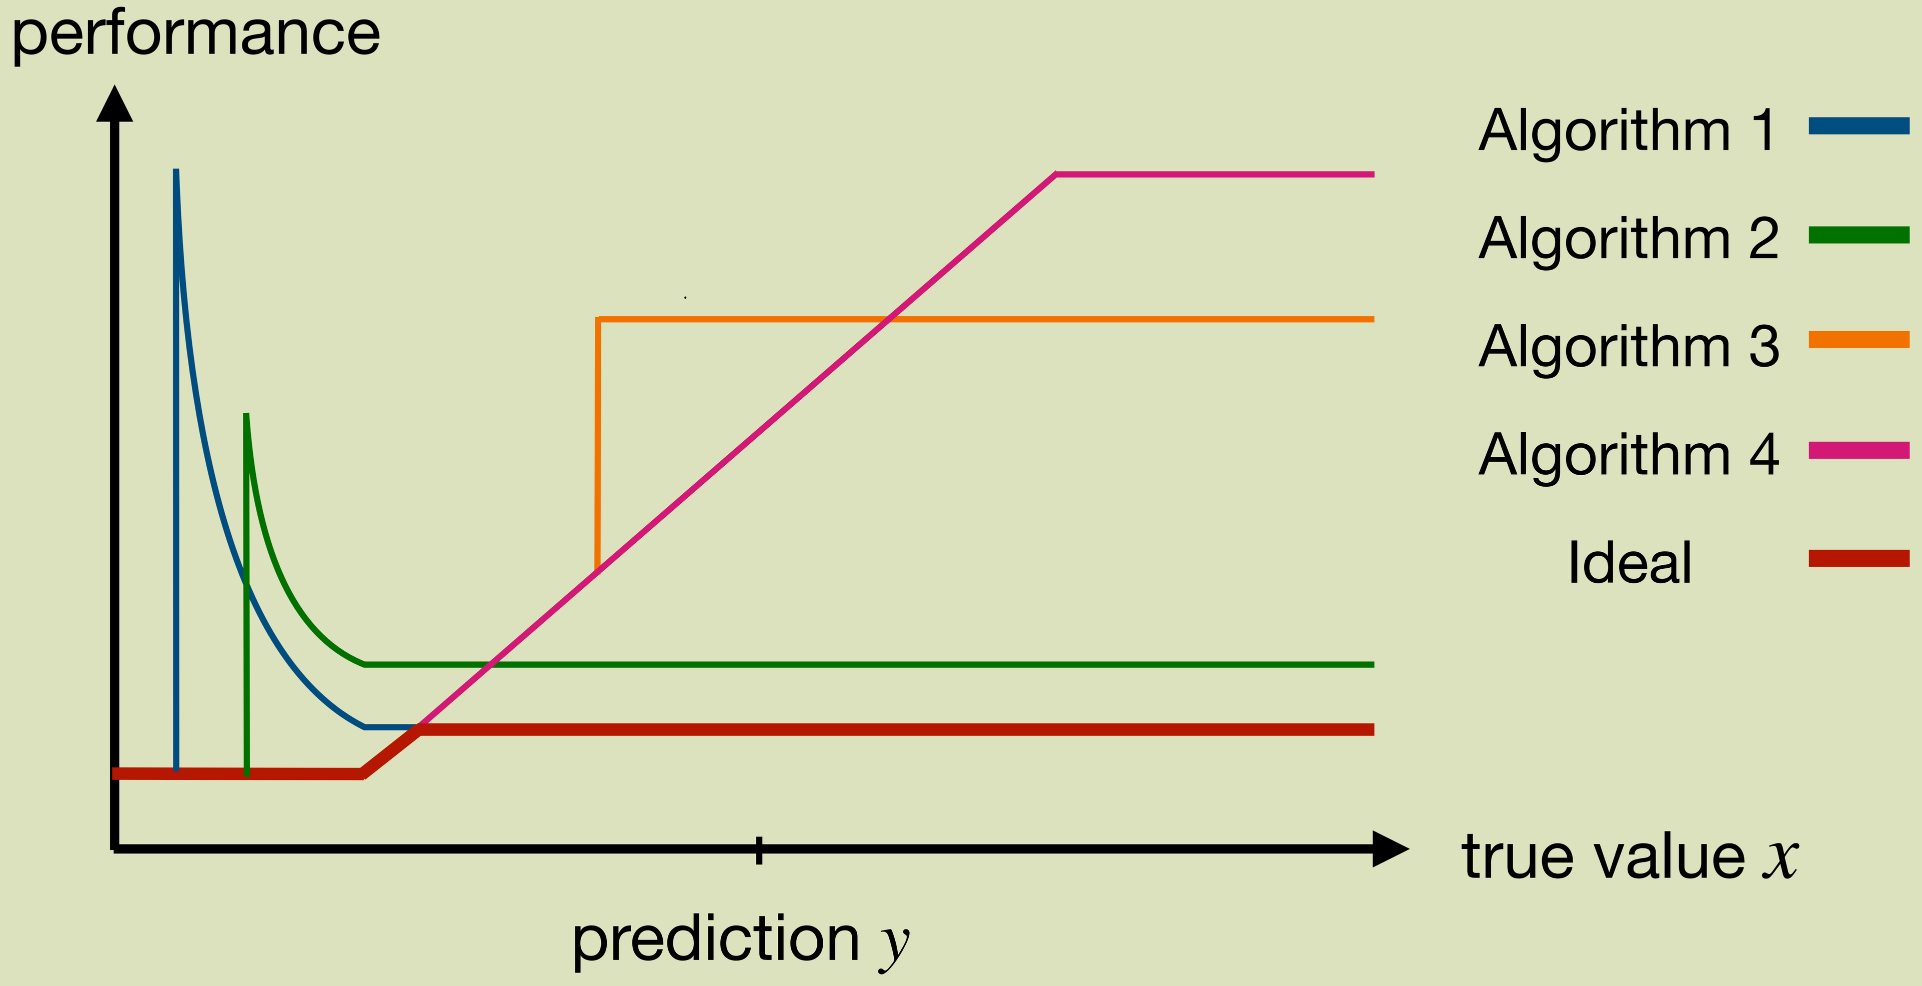

We need a principled comparison metric.

Our framework

Distance-Based View

For a fixed robustness level r, we compare an algorithm A to the ideal r-robust algorithm I_r, which achieves the best possible performance under the same robustness constraint.

We measure how far the performance $pr(A, x)$ of algorithm A is from the ideal benchmark I_r, where x is the true value and y is the prediction.

$$d_{\max}(A) = \sup_x \left(pr(A,x) - pr(I_r,x)\right)w(x)$$

Worst-case weighted deviation from the ideal performance.

$$d_{\text{avg}}(A) = \frac{1}{|R_y|}\int_{R_y} \left(pr(A,x) - pr(I_r,x)\right)w(x)\,dx$$

Average weighted deviation across prediction errors in R_y (prediction range).

These metrics evaluate algorithms by how closely they match the best possible r-robust strategy across the entire range of prediction errors.

Risk-Based View (CVaR)

Instead of evaluating performance for each prediction error, we assume the prediction is uncertain and described by a distribution over inputs.

Conditional Value-at-Risk: $CVaR_\alpha(X) = \inf_t \left(t + \frac{1}{1-\alpha}\mathbb{E}[(X-t)^+]\right)$

CVaR measures the expected loss in the worst $1-\alpha$ fraction of outcomes.

We evaluate an algorithm A using α-consistency

$$\alpha\text{-cons}(A) = \sup_F \frac{CVaR_{\alpha,F}(A(\sigma))}{\mathbb{E}_{\sigma\sim F}[OPT(\sigma)]}$$

This measures algorithm performance when the prediction follows a distribution.

The parameter $\alpha \in [0,1)$ controls risk sensitivity:

- $\alpha = 0$ optimize expected performance
- $\alpha \to 1$ protect against worst-case outcomes

Applications

⛷📈 Threshold-Based Problems
(Ski Rental & One-Max Search)

Both problems admit infinitely many r-robust threshold strategies.

⚙ Contract Scheduling

Infinitely many schedules achieve optimal robustness.

We prove the optimal threshold / schedule under our distance-based and CVaR metrics.

Results

Experimental Results (One-Max Search)

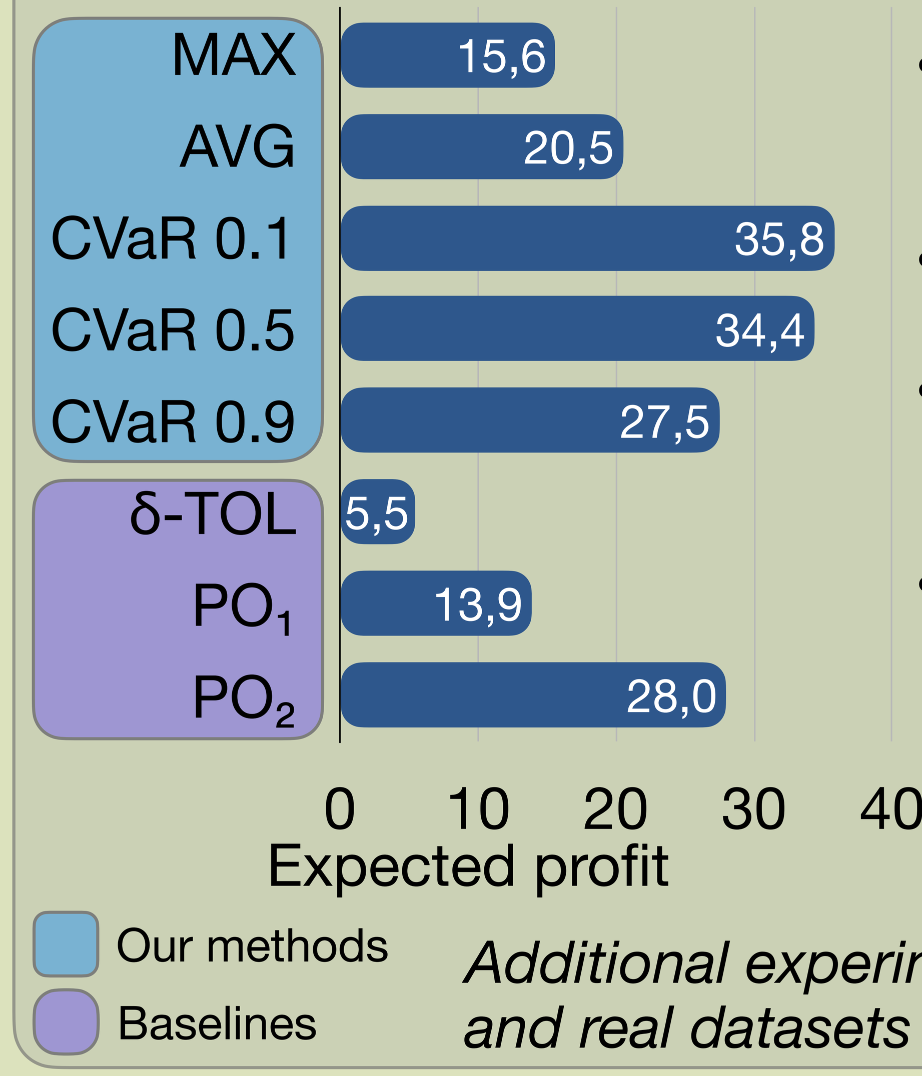

- **Distance-based algorithms (MAX, AVG)** achieve strong profit and competitive performance ratios.
- CVaR algorithms provide a risk–performance tradeoff.
- Compared to **baselines** (δ-TOL, PO_1, PO_2), our methods achieve **higher expected profit**.
- **PO_2** exhibits very poor performance ratio, illustrating the **brittleness of purely Pareto-optimal strategies**.

Additional experiments across all applications and real datasets are available in the paper →

Our decision-theoretic framework systematically selects the best algorithm within large classes of robust strategies.

- We introduce *NEW* decision-theoretic metrics for learning-augmented algorithms.
- For several classic problems, we prove the optimal algorithm within large classes of r-robust strategies.
- Distance-based and CVaR metrics provide complementary views of performance and risk.
- Experiments confirm significant improvements over existing baselines.