Online optimization with black-box advice

Nicolas Christianson, Tinashe Handina, Adam Wierman
Caltech CMS

AI/ML achieves state-of-the-art performance in many domains, but…

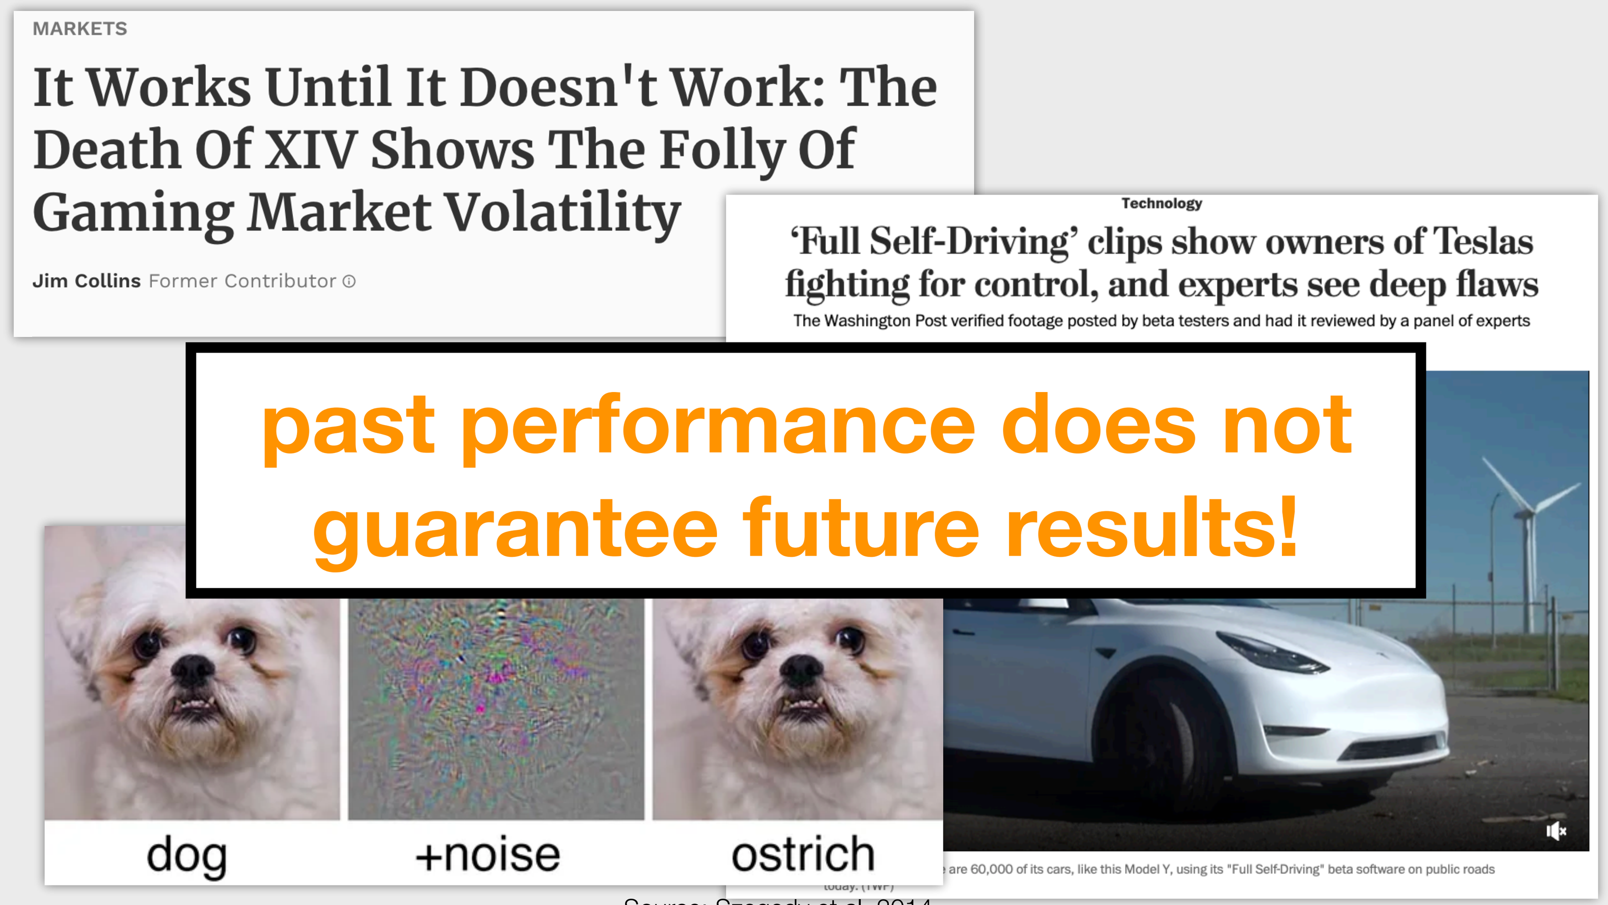

Source: Szegedy et al. 2014

To deploy ML in real-world online decision-making, we need algorithms that:

1. exploit the **good** performance of AI/ML
2. ensure **worst-case robustness** and other desired performance guarantees

Problem focus: "smoothed" online convex optimization

At each time $t = 1, \dots, T$:

1. Adversary gives you a convex *hitting* cost $f_t : \mathbb{R}^d \to \mathbb{R}_+$
2. You choose $x_t \in \mathbb{R}^d$ and pay $f_t(x_t) + \|x_t - x_{t-1}\|$

Special case: convex body chasing (CBC):

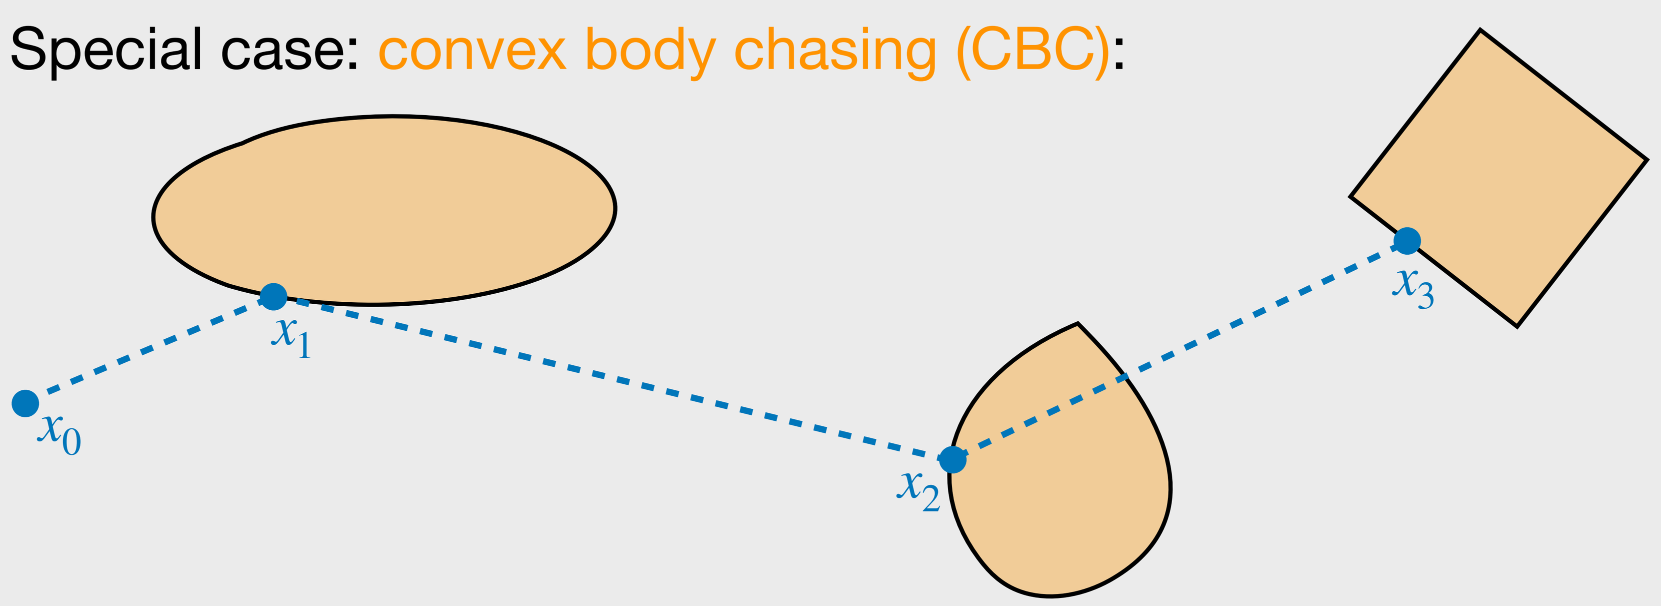

Performance metrics

Typical metric is competitive ratio (CR):

$$\mathrm{Cost}(\mathtt{ALG}) \le \mathrm{CR} \cdot \mathrm{Cost}(\mathtt{OPT}) \quad \forall \{f_t\}$$

- Worst-case metric, doesn't capture average performance - yields conservative algorithms
- $\mathcal{O}(d)$ for general, convex f_t

If data is available about typical problem instances, ML may perform better. Motivates a dual metric:

Black-box ML "advice"

Consistency: $\mathrm{Cost}(\mathtt{ALG}) \le (1+\epsilon) \cdot \mathrm{Cost}(\mathtt{ADV})$

Tunable

Robustness: $\mathrm{Cost}(\mathtt{ALG}) \le C(\epsilon) \cdot \mathrm{Cost}(\mathtt{OPT})$

Visually:

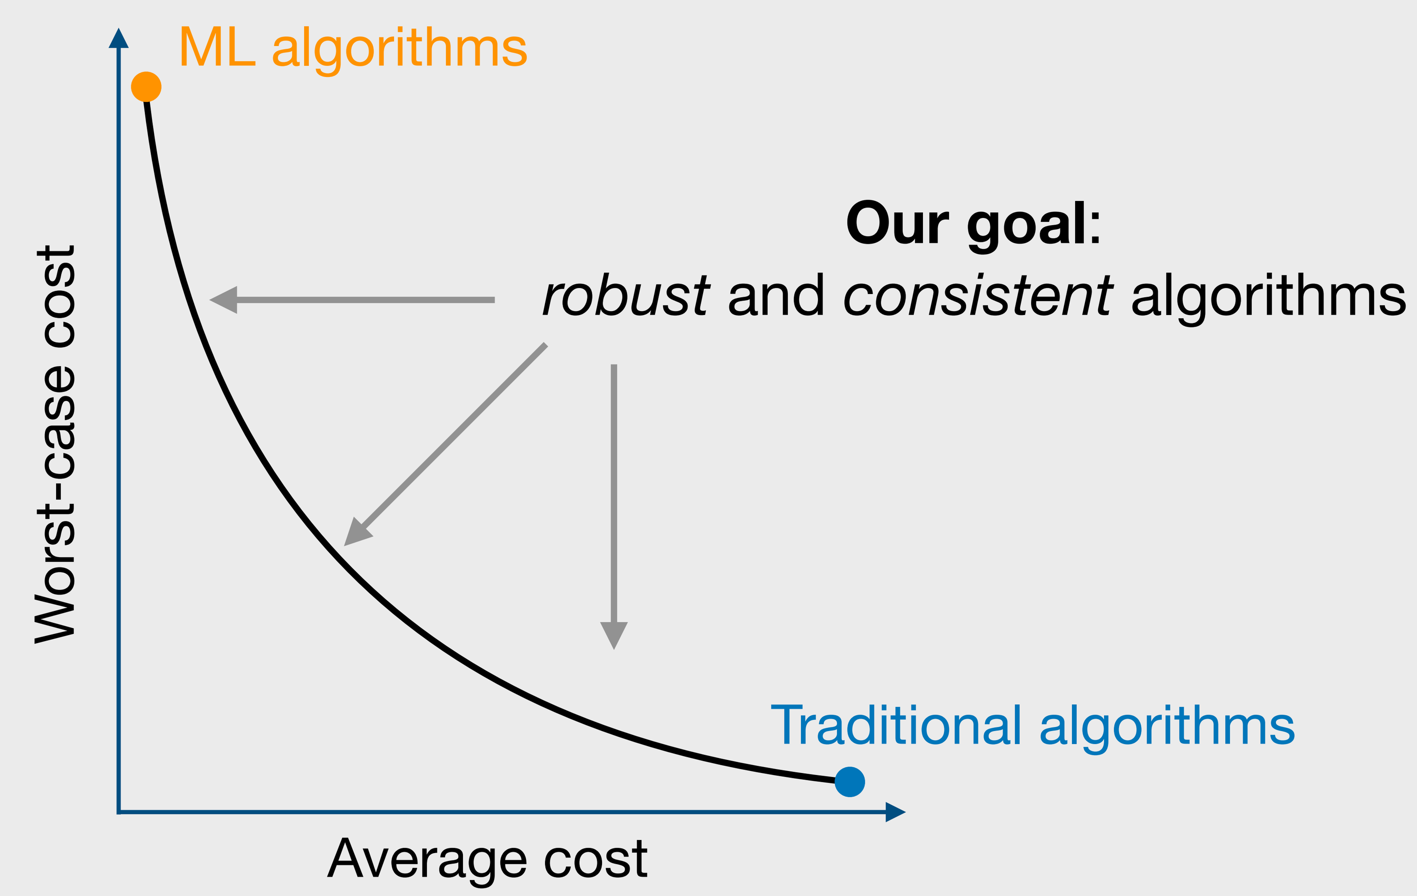

Our approach: design **meta-algorithms** to combine advice with traditional robust algorithms

First attempt: A "switching" algorithm

Basic idea: **Switch** between robust and advice algorithms based on their ongoing performance

Theorem. For any $\epsilon > 0$, **Switch** (with suitable parameters) is $(3 + \mathcal{O}(\epsilon))$-consistent and $\mathcal{O}(d\epsilon^{-2})$-robust.

Can switching algorithms give better consistency? No! (lower bound)

Beyond switching algorithms

We propose an algorithm **INTERP** that exploits convexity to bypass the limits of switching algorithms

Theorem. In Euclidean setting, for any $\epsilon > 0$, **INTERP** is $(\sqrt{2} + \epsilon)$-consistent and $\mathcal{O}(d\epsilon^{-2})$-robust.

Empirical performance - energy dispatch

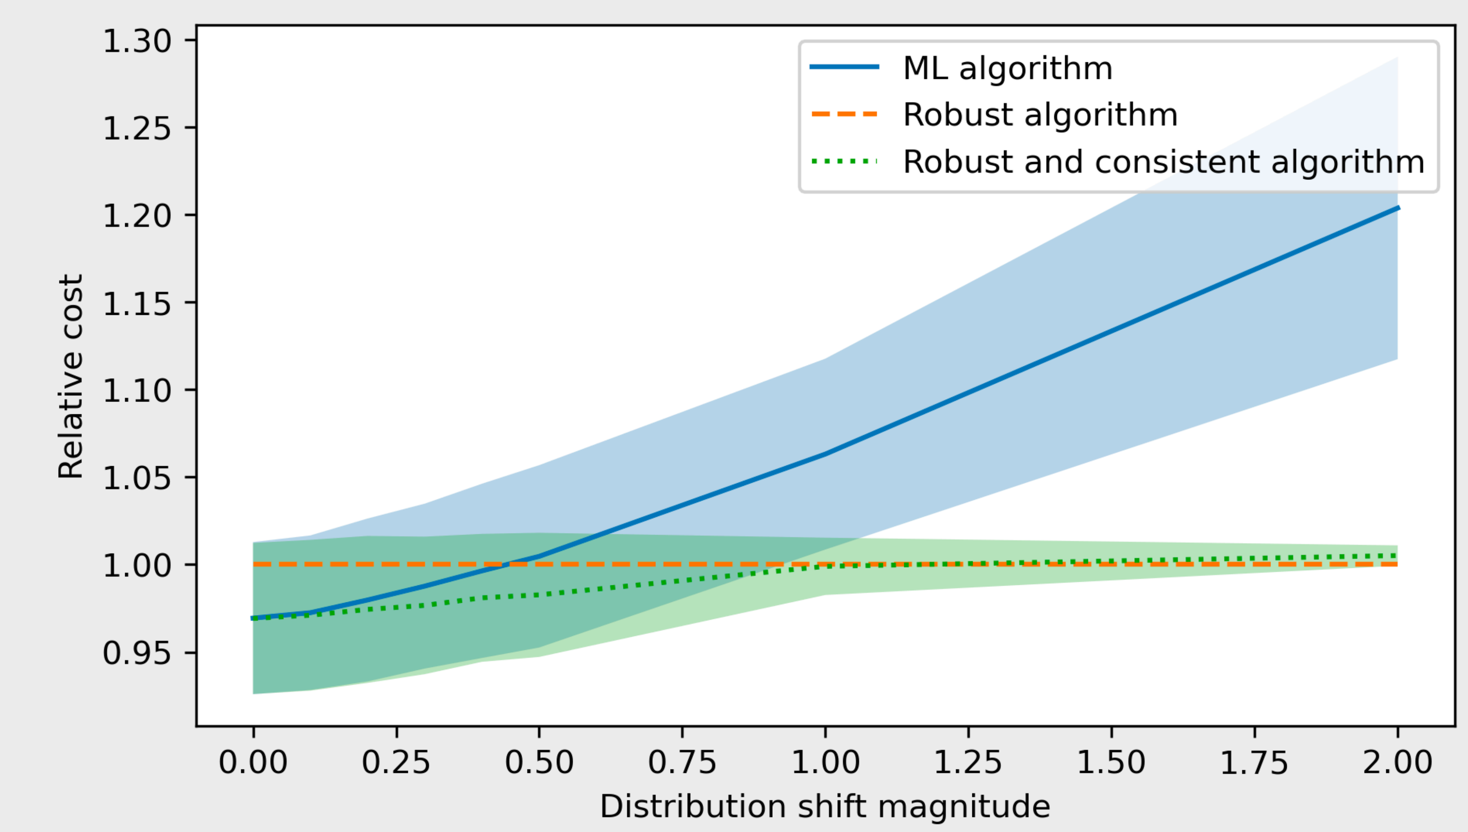

Ongoing/future directions

- Can the paradigm be extended to endow "black-box" algorithms with other sorts of guarantees (e.g., fairness, safety)?
- Other online problems (e.g., mechanism design, following Agrawal et al. '22)

N. Christianson, T. Handina, and A. Wierman. *Chasing convex bodies and functions with black-box advice*. COLT 2022.