Dispersion for Data-Driven Algorithm Configuration, Online Learning, and Private Optimization

Travis Dick

Joint work with Nina Balcan and Ellen Vitercik

Carnegie Mellon University

Data-Driven Algorithm Configuration

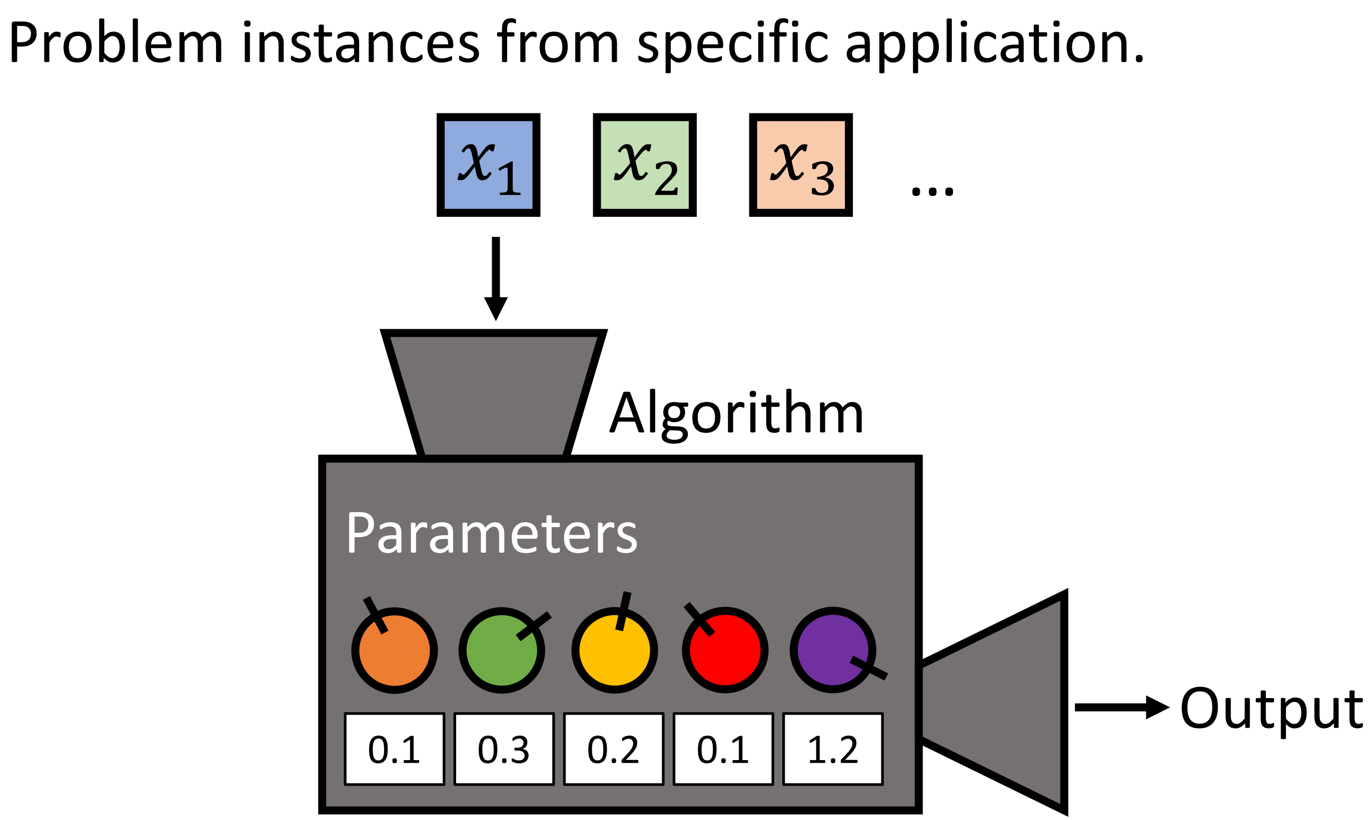

Goal:

- Automatically find the best parameters for a specific application domain.
- Algorithm is run repeatedly, historic instances are training data.
- Want provable guarantees for online and private settings.

An Overview

- Algorithm configuration often requires optimizing sums of piecewise Lipschitz functions.

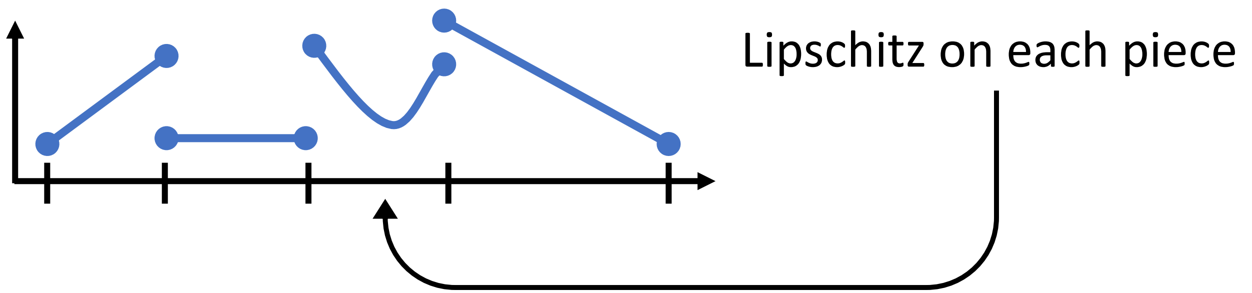

- Worst-case impossibility results for **online** and **private** optimization of PWL functions.

Our Contributions:

- Identify a general structural property called ***dispersion*** that implies
  - Good regret bounds in online optimization
  - Utility guarantees for differentially private optimization
  - Uniform convergence results in statistical settings.
- Satisfied in real problems under very mild assumptions.

Example: Greedy Knapsack Algorithm

[Gupta & Roughgarden '16]

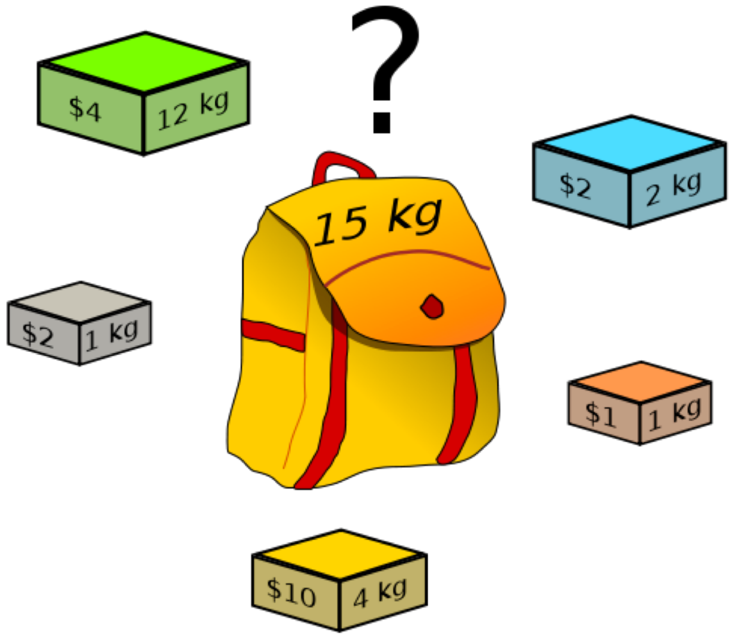

Problem Instance:

- Given n items
- item i has value v_i and size s_i
- a knapsack with capacity K

Find the most valuable subset of items that fits.

Algorithm: (Parameter $\rho \geq 0$)

Add items in decreasing order of $\text{score}_\rho(i) = v_i / s_i^\rho$.

Observation: For one instance, total value is piecewise constant in ρ.

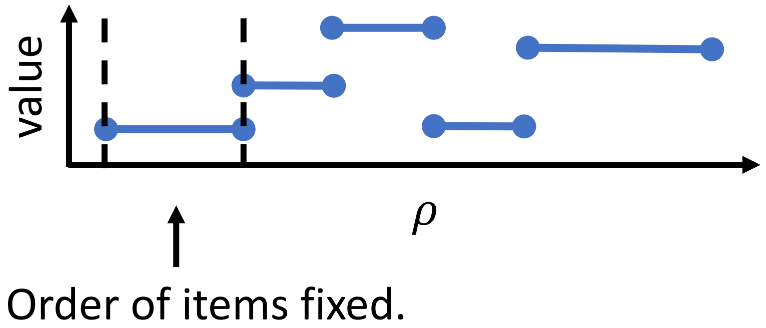

Online Piecewise Lipschitz Optimization

Learning protocol:

For each round $t = 1, \dots, T$:

1. Learner chooses point $\rho_t \in C \subset R^d$.
2. Adversary chooses piecewise L-Lipschitz function $u_t: C \to R$.
3. Learner gets reward $u_t(\rho_t)$
4. **Full information:** Learner observes entire function u_t
5. **Bandit information:** Learner only observes the scalar $u_t(\rho_t)$

Goal: Minimize regret $= \max_{\rho \in C} \sum_{t=1}^{T} u_t(\rho) - \sum_{t=1}^{T} u_t(\rho_t)$.

- *[Gupta & Roughgarden '16] have online algorithms for Max-Weight Independent Set with smoothed adversaries.*
- *[Cohen-Addad & Kanade '17] consider 1-dim. piecewise constant functions with smoothed adversaries.*

A Mean Adversary

Fact: There exists an adversary choosing piecewise constant functions from $[0,1]$ to $[0,1]$ such that **every** full information online algorithm has **linear regret**.

At round t, adversary chooses a threshold τ_t and flips a coin to choose either

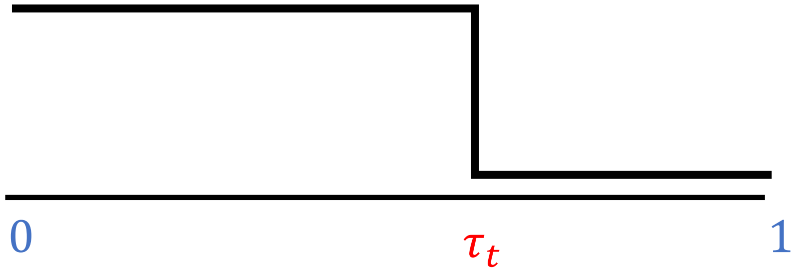

or

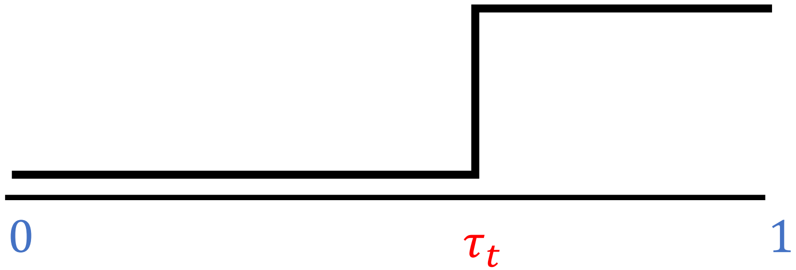

Learner has expected utility of ½ per round → expected total utility $T/2$.

Let G_t = Set of points with payout 1 in first t rounds

Set τ_t to be midpoint of G_{t-1} → G_T has utility T

Regret = $T/2$.

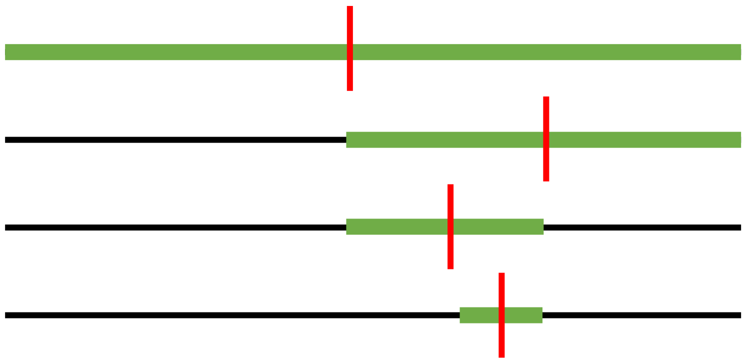

Dispersion

The mean adversary concentrated discontinuities near ρ^*. Even very near points had low utility!

Def: Functions $u_1(\cdot), \ldots, u_T(\cdot)$ are ***$(\boldsymbol{w}, \boldsymbol{k})$-dispersed at point $\boldsymbol{\rho}$*** if the ℓ_2-ball $B(\rho, w)$ contains discontinuities for at most k of $u_1 \ldots, u_T$.

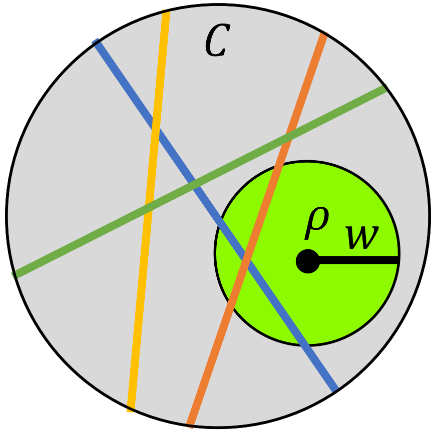

Each colored line is a discontinuity of one function.

Ball of radius w about ρ contains 2 discontinuities.
→ $(w, 2)$-dispersed.

Functions will satisfy a range of dispersion parameters:

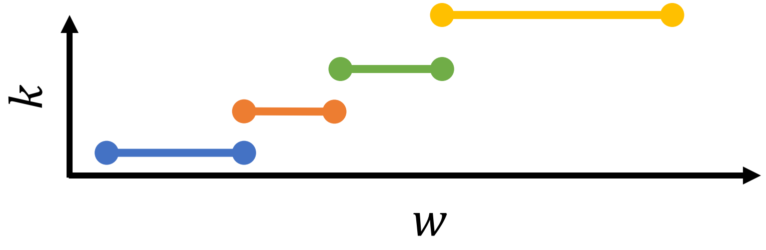

The Sum of Disperse Functions

Let $u_1, \dots, u_T$ be PWL functions and plot their sum $\sum_t u_t$

Not disperse **Disperse**

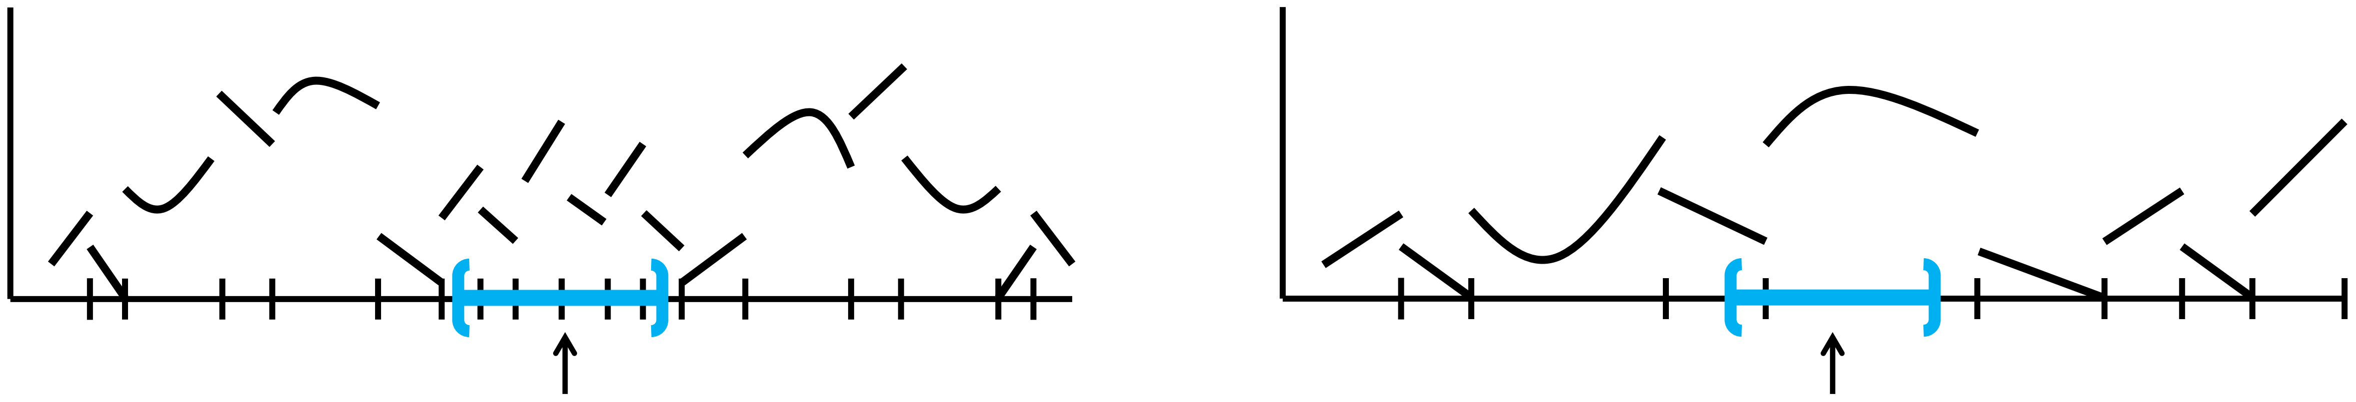

Many boundaries within interval

Few boundaries within interval

Key Property of Dispersed Functions

Lemma: Let $u_1, \ldots, u_T : C \to [0,1]$ be piecewise L-Lipschitz and (w,k)-dispersed at a maximizer ρ^*. Then every $\rho \in B(\rho^*, w)$ satisfies $\sum_{t=1}^{T} u_t(\rho) \geq OPT - TLw - k$.

"The w-neighborhood of ρ^* has high utility".
Dispersion characterizes how the utility decays with w.

Proof:

$u_1, \ldots, u_T$

Is u_t L-Lipschitz on $B(\rho^*, w)$?

no → $|u_t(\rho) - u_t(\rho^*)| \leq 1$

yes → $|u_t(\rho) - u_t(\rho^*)| \leq Lw$

At most k fn's

At most T fn's

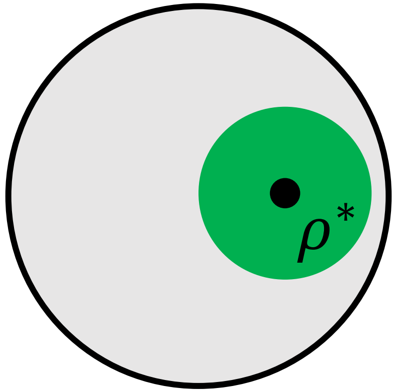

Full Information Regret Bounds

We analyze the classic Exponentially Weighted Forecaster *[Cesa-Bianchi & Lugosi '06]*

Algorithm: (Parameter $\lambda > 0$)
At round t, sample ρ_t from $p_t(\rho) \propto \exp(\lambda \sum_{s=1}^{t-1} u_s(\rho))$.

Theorem: If $u_1, \ldots, u_T : C \to [0,1]$ are piecewise L-Lipschitz and (w, k)-dispersed at ρ^*,

EWF has regret $O\left(\sqrt{Td \log \frac{1}{w}} + TLw + k\right)$.

When is this a good bound? For $w = 1/(L\sqrt{T})$ and $k = \tilde{O}(\sqrt{T})$ regret is $\tilde{O}(\sqrt{Td})$.

Intuition:

- The ball $B(\rho^*, w)$ has utility at least $OPT - TLw - k$.
- EWF can compete with $B(\rho^*, w)$ when its volume is non-negligible.

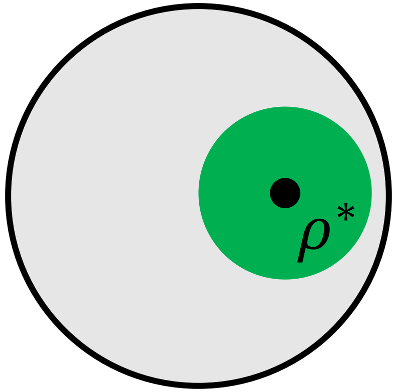

Note: Don't need to know (w, k) to run algorithm.

Matching Lower Bound

Theorem: For any algorithm A, there are piecewise constant functions $u_1, \ldots, u_T$ so that A has expected regret at least

$$\Omega\left(\inf_{(w,k)} \sqrt{Td\log\left(\frac{1}{w}\right)} + k\right)$$

Where the infimum is over all (w, k)-dispersion parameters satisfied by $u_1, \ldots, u_T$ at ρ^*.

Idea: Calculate dispersion parameters for mean adversary.

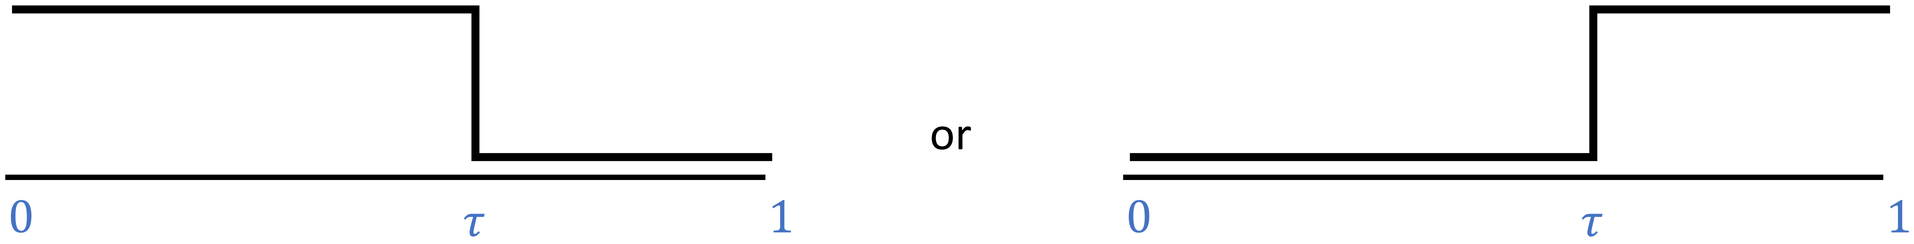

- Can make upper bound only tight for $k \geq \sqrt{T}$.
- Analysis of *[G&R '16, C-A&K '17]* both use analysis similar to $(w, 1)$-dispersion.

Smoothed Adversaries and Dispersion

Similar to *[G&R '16, C-A&K '17]* we often consider *smoothed* adversaries.

A simple example:

Consider any adversary chooses threshold functions $u_1, \ldots, u_T : [0,1] \to [0,1]$:

Location $\tau \in [0,1]$
Orientation $s \in \{\pm 1\}$

Location τ corrupted by adding $Z \sim N(0, \sigma^2)$.

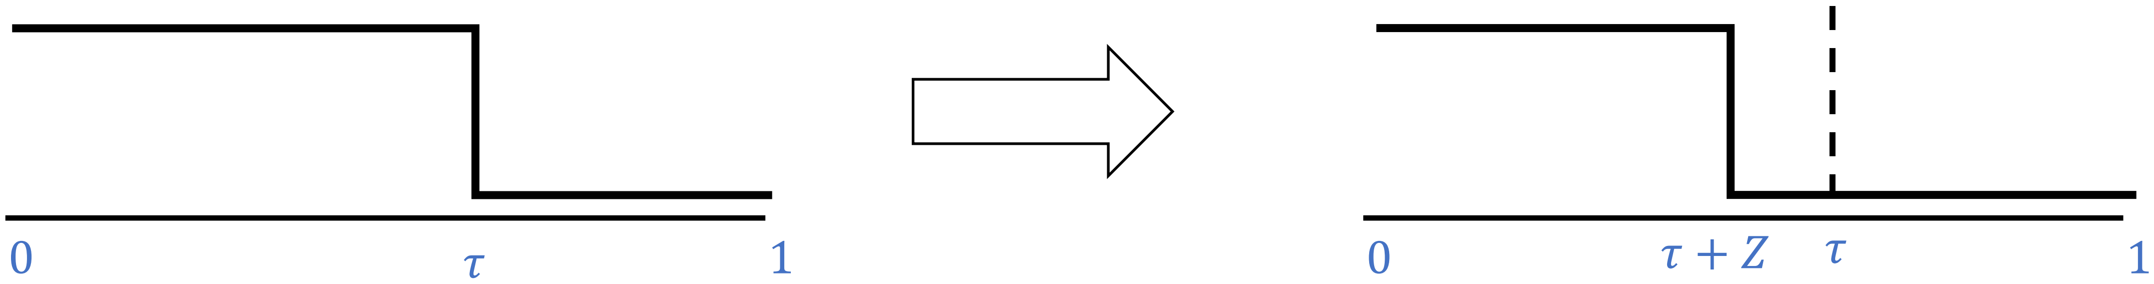

Lemma: For any $w > 0$, the fn's $u_1, \ldots, u_T$ are (w, k)-dispersed for $k = \tilde{O}(Tw/\sigma + \sqrt{T})$ w.h.p.

Density of τ_t is $O(1/\sigma)$.
For any interval I of width w:

- Expected # discontinuities is $O(Tw/\sigma)$.
- Intervals have VC-dim 2 → w.h.p. for all I, # discontinuities is $\tilde{O}(Tw/\sigma + \sqrt{T})$

Take $w = \frac{\sigma}{\sqrt{T}}$ → regret = $O\left(\sqrt{T \log \frac{T}{\sigma}}\right)$

Smoothed Adversaries and Dispersion

More generally: adversary is unable to precisely pick some problem parameters (e.g. item values in knapsack).

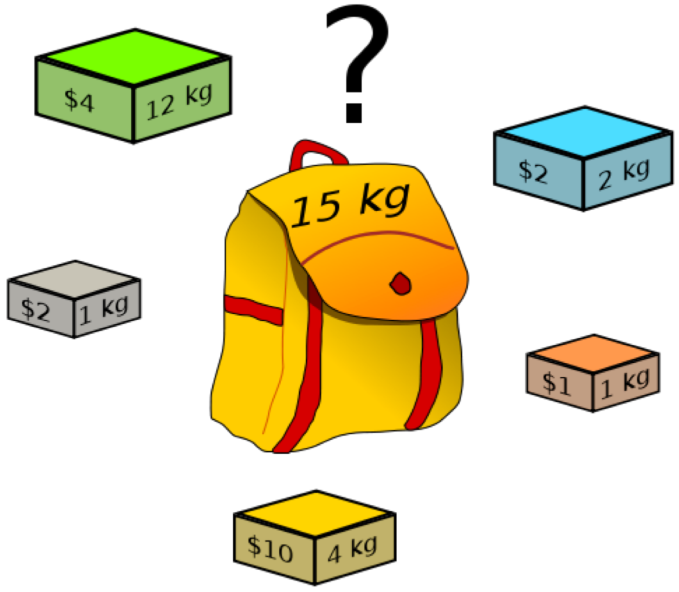

Challenges:

- Each utility function has multiple dependent discontinuities.
- Distribution of discontinuity location depends on setting.
- How do we generalize to multiple dimensions?

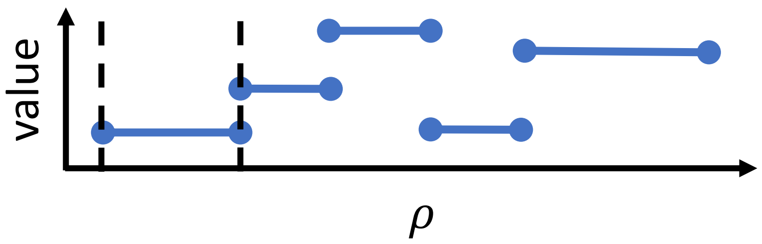

$$\rho = \frac{\ln(v_i/v_j)}{\ln(s_i/s_j)}$$

Dispersion decouples problem-specific smoothness arguments from regret bounds, differential privacy utility guarantees, and uniform convergence.

Dispersion in Algorithm Configuration and Pricing

We show (w, k)-dispersion with $w \approx 1/\sqrt{T}$ and $k \approx \sqrt{T}$ under smoothness assumptions for

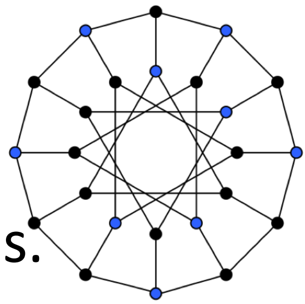

- Greedy algorithms for knapsack and max-weight independent set.
- n-bidder, m-item posted price mechanisms and second price auctions w/ reserves.
  - Maximizing revenue or social welfare.
  - Additive, unit-demand, and general valuations.

Under ***no assumption*** we show dispersion for

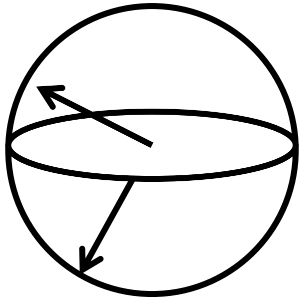

- Semidefinite rounding algorithms for integer quadratic programming
  - s-linear rounding *[Feige & Langberg '06]*
  - Outward rotations *[Zwick '99]*
  - Both are generalizations of the Goemans & Williamson max cut algorithm *['95]*.

More Results from Dispersion

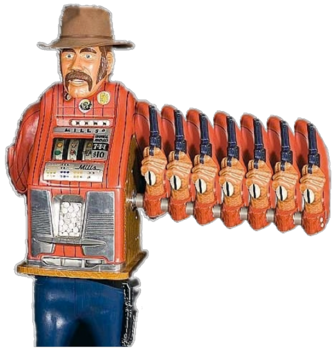

Bandit online optimization of piecewise Lipschitz functions:

- Learner only observes the scalar $u_t(\rho_t)$ each round.
- Dispersion with $w \approx 1/T^{1/(d+2)}$ and $k \approx T^{\frac{d+1}{d+2}}$ implies regret $\tilde{O}\left(T^{\frac{d+1}{d+2}}\left(\sqrt{d3^d} + L\right)\right)$.

Differentially Private Batch Optimization

- Given $u_1, \ldots, u_T$ up-front, estimate maximizer of $\frac{1}{T}\sum_t u_t$.
- Satisfy (ϵ, δ)-differential privacy (w.r.t. changing any one function).
- Exponential mechanism has suboptimality $\tilde{O}\left(\frac{1}{T\epsilon} d \log \frac{1}{w} + Lw + \frac{k}{T}\right)$.
- Matching lower bounds.

Uniform Convergence via Empirical Rademacher Bounds

- If $u_1, \ldots, u_T$ are drawn i.i.d. from a distribution P and are *globally* (w, k)-dispersed, then
- w.h.p., for every $\rho \in C$, we have $\left|\frac{1}{T}\sum_t u_t(\rho) - \mathbb{E}_{u\sim P}[u(\rho)]\right| = \tilde{O}\left(\sqrt{\frac{d \log \frac{1}{w}}{T}} + Lw + \frac{k}{T}\right)$.

Conclusions & Open Questions

- Introduced **dispersion**, which measures concentration of discontinuities of PLW fns.
  - Dispersion-based regret bounds for online optimization of PWL functions.
  - Utility guarantees for differentially private optimization.
  - Uniform convergence in statistical settings.
- Examples of dispersion in real problems.

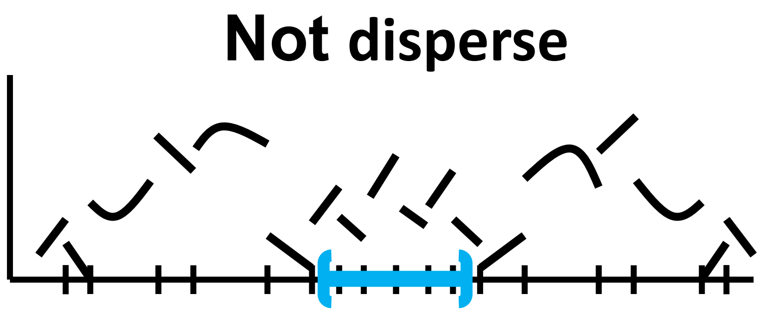

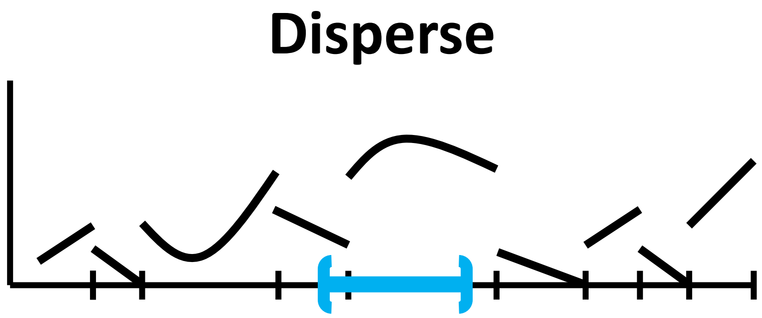

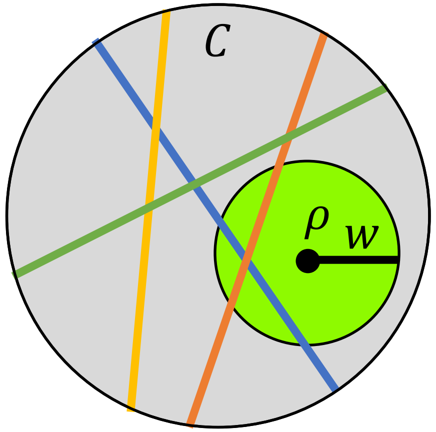

Open Questions:

- Dispersion-like definitions for other types of bad properties beyond discontinuities.
- Algorithm configuration is between full-info and bandit feedback. Can we take advantage?
- More refined algorithms for bandit online optimization.

Thanks!