Simultaneous Bounds on Competitiveness and Regret

Lachlan Andrew (Swinburne)
Siddharth Barman (Caltech)
Katrina Ligett (Caltech)
Minghong Lin (Caltech)
Adam Meyerson (Google)
Alan Roytman (UCLA)
Adam Wierman (Caltech)

Key Message: Competitive ratio and regret are incompatible, in the context of a natural class of online optimization problems.

Smoothed Online Convex Optimization (SOCO)

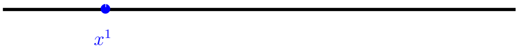

Online Sequence: x^1, c^1

Smoothed Online Convex Optimization (SOCO)

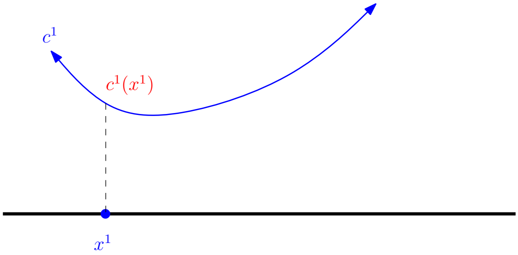

Online Sequence: x^1, c^1

Smoothed Online Convex Optimization (SOCO)

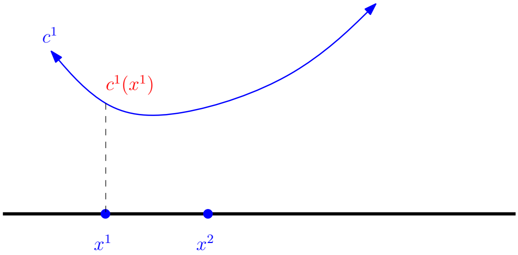

Online Sequence: x^1, c^1, x^2, c^2

Smoothed Online Convex Optimization (SOCO)

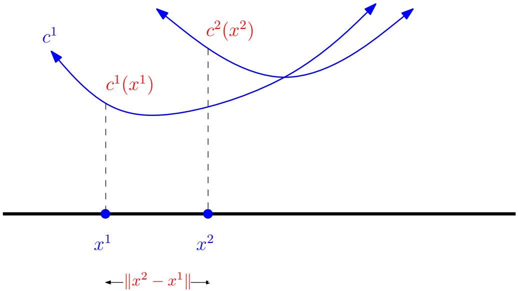

Online Sequence: x^1, c^1, x^2, c^2

Smoothed Online Convex Optimization (SOCO)

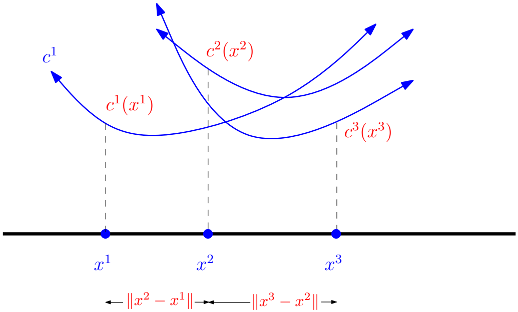

Online Sequence: $x^1, c^1, x^2, c^2, \ldots$

SOCO Formulation

Online Sequence: $x^1, c^1, x^2, c^2, \ldots$

$$\boxed{\min_{x^t \in F} \quad \Sigma_t \; c^t(x^t) + \|x^t - x^{t-1}\|}$$

Convex

Norm (smoothed)

Goal: Algorithm to minimize cost

SOCO Formulation

Online Sequence: $x^1, c^1, x^2, c^2, \ldots$

$$\min_{x^t \in F} \quad \Sigma_t \; c^t(x^t) + \|x^t - x^{t-1}\|$$

Convex

Norm (smoothed)

Applications:

- Dynamic resizing in data centers
- Geographical load balancing
- Video Streaming
- Portfolio management

SOCO: Two views

$$\min_{x^t \in F} \sum_t c^t(x^t) + \|x^t - x^{t-1}\|$$

Online Algorithms

- Metrical Task System (MTS) with extra convexity
- Metric: Competitive ratio

Online Learning

- Online Convex Optimization (OCO) with switching costs
- Metric: Regret

Metrical Task Systems (MTS)

$$\min_{x^t} \sum_t f^t(x^{t+1}) + d(x^{t+1}, x^t)$$

Non convex

n point metric

Action chosen after cost is known

Performance Metric: Competitive Ratio

$$\mathrm{CR}(\mathrm{ALG}) := \max_{\text{input}} \frac{\mathrm{Cost}_1(\mathrm{ALG})}{\mathrm{OPT}_{\text{dynamic}}}$$

Typical competitive ratio results:

- $O(n)$ Work-function algorithms [KP95,...]
- $\Omega(\sqrt{\log n / \log\log n})$ [BKRS92]

SOCO: Two views

$$\boxed{\min_{x^t \in F} \quad \sum_t c^t(x^t) + \|x^t - x^{t-1}\|}$$

Online Algorithms

- Metrical Task System (MTS) with extra convexity
- 3 competitive for one-dimensional SOCO [LWAT11]

Online Learning

- Online Convex Optimization (OCO) with switching costs

Online Convex Optimization (OCO)

$$\boxed{\min_{x^t} \quad \sum_{t=1}^{T} \; c^t(x^t)}$$

Convex

Action chosen before cost is known

Performance Metric: Regret

$$\text{Rg(ALG)} := \max_{\text{input}} \quad \text{Cost}_0(\text{ALG}) - \text{OPT}_{\text{static}}$$

Typical regret results

- $O(\sqrt{T})$ Online gradient descent [Z03]
- $O(\log T)$ for bounded curvature [HAK07]

SOCO: Two views

$$\min_{x^t \in F} \sum_t c^t(x^t) + \|x^t - x^{t-1}\|$$

Online Algorithms

- Metrical Task System (MTS) with extra convexity
- 3 competitive for one-dimensional SOCO [LWAT11]

Online Learning

- Online Convex Optimization (OCO) with switching costs
- Online gradient descent achieves sub-linear regret for SOCO

SOCO: Fundamental Incompatibility

Natural Question[1]:

Can an algorithm maintain sub-linear regret *and* a constant competitive ratio?

[1]Similar tradeoffs also considered in [BB00] and [BCNS12]

SOCO: Fundamental Incompatibility

Natural Question[1]**:**

Can an algorithm maintain sub-linear regret *and* a constant competitive ratio?

✗ No

[1] Similar tradeoffs also considered in [BB00] and [BCNS12]

SOCO: Fundamental Incompatibility

For any algorithm (deterministic or randomized), either regret is $\Omega(T)$ or competitive ratio is unbounded in T.

Online Sequence: $x^1, c^1, \ldots, c^{t-1}, x^t, c^t, \ldots$

$$\mathrm{CR(ALG)} := \max_{\text{input}} \frac{\mathrm{Cost}_1(\mathrm{ALG})}{\mathrm{OPT}_{\text{dynamic}}}$$
$$\mathrm{Rg(ALG)} := \max_{\text{input}} \mathrm{Cost}_0(\mathrm{ALG}) - \mathrm{OPT}_{\text{static}}$$

Key Argument: x^t can't be expected to do well simultaneously for both c^{t-1} and c^t.

SOCO: Fundamental Incompatibility

For any algorithm (deterministic or randomized), either regret is $\Omega(T)$ or competitive ratio is unbounded in T.

Online Sequence: $x^1, c^1, \ldots, c^{t-1}, x^t, c^t, \ldots$

$$\mathrm{CR}(\mathrm{ALG}) := \max_{\text{input}} \frac{\mathrm{Cost}_1(\mathrm{ALG})}{\mathrm{OPT}_{\text{dynamic}}}$$

$$\mathrm{Rg}(\mathrm{ALG}) := \max_{\text{input}} \mathrm{Cost}_0(\mathrm{ALG}) - \mathrm{OPT}_{\text{static}}$$

$$\mathrm{Rg}'(\mathrm{ALG}) := \max_{\text{input}} \mathrm{Cost}_1(\mathrm{ALG}) - \mathrm{OPT}_{\text{static}}$$

SOCO: Fundamental Incompatibility

Theorem 2

For any algorithm (deterministic or randomized), either $\text{Rg}' = \Omega(T)$ or competitive ratio is unbounded in T.

Online Sequence: $x^1, c^1, \ldots, c^{t-1}, x^t, c^t \ldots$

$$\text{CR}(\text{ALG}) := \max_{\text{input}} \frac{\text{Cost}_1(\text{ALG})}{\text{OPT}_{\text{dynamic}}}$$

$$\text{Rg}'(\text{ALG}) := \max_{\text{input}} \text{Cost}_1(\text{ALG}) - \text{OPT}_{\text{static}}$$

Incompatibility: Proof Sketch

Theorem 2

For any algorithm (deterministic or randomized), either $\mathrm{Rg}' = \Omega(T)$ or competitive ratio is unbounded in T.

Construct an adversarial instance using linear functions in batches.

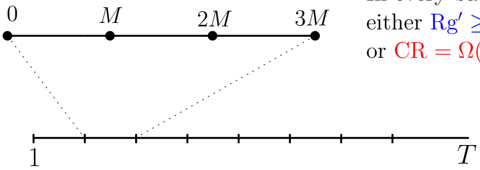

In every batch
either $\mathrm{Rg}' \geq$ const.
or $\mathrm{CR} = \Omega(M)$

Incompatibility: Proof Sketch

Theorem 2

For any algorithm (deterministic or randomized), either $\mathrm{Rg}' = \Omega(T)$ or competitive ratio is unbounded in T.

Construct an adversarial instance using linear functions in batches.

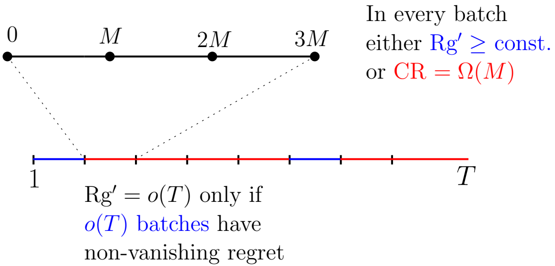

Incompatibility: Proof Sketch

Abundance of f_0s ensures $x^*_{\text{static}} = 0$.

If $x^t > \delta$ then $\mathrm{Rg}' \geq$ const.

Else if $x^t \leq \delta$ then CR$= \Omega(M)$

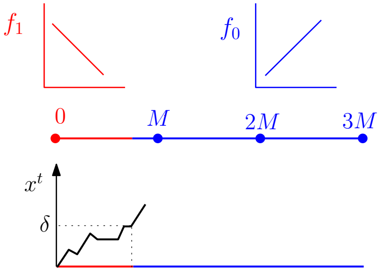

In every batch
either $\mathrm{Rg}' \geq$ const.
or $\mathrm{CR} = \Omega(M)$

SOCO: Balancing the metrics

Can an algorithm maintain sub-linear regret *and* a constant competitive ratio?

✗ No

It is possible to "nearly" achieve this goal in *one-dimensional* action space?

✓ Yes

SOCO: Balancing the metrics

Algorithm Randomized Biased Greed: RBG(θ)

- Pick a uniformly random number $r \in (-1, 1)$.
- In each time step t, we consider a work function w^t: $w^t(x) = \min_y\{w^{t-1}(y) + \|y - x\|\}$.
- Go to point a x that minimizes $w^t(x) + r\theta\|x\|$.

Theorem

For one-dimensional SOCO, using RBG we can achieve a competitive ratio of θ while maintaining $\mathrm{Rg}' = O(T/\theta)$.

Summary

Can an algorithm maintain sub-linear regret *a*nd a constant competitive ratio?

✗ No

It is possible to "nearly" achieve this goal in *one-dimensional* action space?

✓ Yes

Summary

Can an algorithm maintain sub-linear regret *a*nd a constant competitive ratio?

✗ No

It is possible to "nearly" achieve this goal in *one-dimensional* action space?

✓ Yes

Thank you!