On the Complexity of Bandit and Derivative-Free Stochastic Convex Optimization

Ohad Shamir

Microsoft Research

NIPS 2012 Optimization Workshop

Bandit / Derivative-Free Stochastic Convex Optimization

Setting

- Convex domain $\mathcal{W} \subseteq \mathbb{R}^d$
- Unknown convex function $F : \mathcal{W} \mapsto \mathbb{R}$
- Can get $F(\mathbf{w})$+noise at any $\mathbf{w} \in \mathcal{W}$
- Want to optimize F with as few queries as possible

Information is *zeroth-order*: No direct access to gradients/Hessians

Bandit / Derivative-Free Stochastic Convex Optimization

Optimization Community

- Derivative-Free / Zeroth-Order SCO
- Black-box situations where gradient is hard to compute / unavailable
- Goal: Minimize *optimization error*

$$\mathbb{E}\left[F(\bar{\mathbf{w}}_T) - F(\mathbf{w}^*)\right]$$

Online Learning Community

- Bandit SCO
- Sequential decision making under uncertainty (e.g. multi-armed bandits)
- Goal: Minimize *regret*

$$\mathbb{E}\left[\frac{1}{T}\sum_{t=1}^{T} F(\mathbf{w}_t) - F(\mathbf{w}^*)\right]$$

Minimizing regret is harder than minimizing error

Attainable Performance

What is the attainable error/regret in terms of

- Number of queries T
- Dimension d

With gradient information, situation is simple:

	Error		**Regret**	
Function Type	$\mathcal{O}$	Ω	$\mathcal{O}$	Ω
Convex	$\sqrt{1/T}$			
Strongly Convex	$1/T$		$\log(T)/T$	

Attainable Performance

In Bandit/Derivative-free Setting: **???**

Attainable Performance

In Bandit/Derivative-free Setting: **???**

- Multi-armed bandits (Linear F, $\mathcal{W}$ is simplex): $\Theta(\sqrt{d/T})$

  [Auer et al. 2002],[Audibert and Bubeck 2009]

Attainable Performance

In Bandit/Derivative-free Setting: **???**

- Multi-armed bandits (Linear F, $\mathcal{W}$ is simplex): $\Theta(\sqrt{d/T})$

  [Auer et al. 2002],[Audibert and Bubeck 2009]

- Linear F, other convex domains:
  $\mathcal{O}(\sqrt{d/T})$ or $\mathcal{O}(\sqrt{d^2/T})$ bounds, depending on domain

  [Abbasi-Yadkori et al. 2011],[Audbiert et al. 2011],[Bubeck et al. 2012]

Attainable Performance

In Bandit/Derivative-free Setting: **???**

- Multi-armed bandits (Linear F, $\mathcal{W}$ is simplex): $\Theta(\sqrt{d/T})$

  [Auer et al. 2002],[Audibert and Bubeck 2009]
- Linear F, other convex domains:
  $\mathcal{O}(\sqrt{d/T})$ or $\mathcal{O}(\sqrt{d^2/T})$ bounds, depending on domain

  [Abbasi-Yadkori et al. 2011],[Audbiert et al. 2011],[Bubeck et al. 2012]
- General convex F:
  - $\mathcal{O}\left(\left(d^2/T\right)^{1/4}\right)$ [Flaxman et al. 2005]
  - $\mathcal{O}\left(\sqrt{d^{34}/T}\right)$ [Agarwal et al. 2011]

Yudin and Nemirosvki 1979

...Each of the methods suggested is in some respect unimprovable, but bad in other respects... We have not succeeded in combining their good qualities and eliminating the bad. Whether it is possible to do this... we do not know. The situation with a zeroth-order oracle is far from clear.

Our Results

We study the complexity of nonlinear bandit/derivative-free SCO - particularly strongly convex functions

1st Main Result

For strongly-convex and smooth functions, attainable error/regret is exactly $\Theta(\sqrt{d^2/T})$

Our Results

We study the complexity of nonlinear bandit/derivative-free SCO - particularly strongly convex functions

1st Main Result

For strongly-convex and smooth functions, attainable error/regret is exactly $\Theta(\sqrt{d^2/T})$

- Follows from a new lower bound, which matches upper bound in [Agarwal et al., 2010]
- First tight complexity characterization for a general nonlinear class
- T must scale quadratically with the dimension d
- Stronger lower bound for strongly convex and convex functions

Our Results

2nd Main Result

In the special case of quadratic functions, attainable **error** is exactly $\Theta(d^2/T)$

- Improvable to $\mathcal{O}(d/T)$ under additional assumptions

- Demonstrates "fast rate" is possible for a general function class, even without gradient knowledge
- "Contradicts" $\Omega(\sqrt{d/T})$ lower bound presented at this NIPS [Jamieson et al. 2012]

Our Results

2nd Main Result

In the special case of quadratic functions, attainable **error** is exactly $\Theta(d^2/T)$

- Improvable to $\mathcal{O}(d/T)$ under additional assumptions

- Demonstrates "fast rate" is possible for a general function class, even without gradient knowledge
- "Contradicts" $\Omega(\sqrt{d/T})$ lower bound presented at this NIPS [Jamieson et al. 2012]

3rd Main Result

Even for quadratic functions, attainable **regret** is exactly $\Theta(\sqrt{d^2/T})$

- Large gap between bandit and derivative-free SCO for nonlinear F

Our Results

	Error		Regret	
Function Type	$\mathcal{O}$	Ω	$\mathcal{O}$	Ω
Quadratic	$\mathbf{\frac{d^2}{T}}$		$\sqrt{\mathbf{\frac{d^2}{T}}}$	
St. Convex + Smooth	$\sqrt{\mathbf{\frac{d^2}{T}}}$			
Str. Convex	$\min\left\{\sqrt[3]{\frac{d^2}{T}}, \sqrt{\frac{d^{34}}{T}}\right\}$	$\sqrt{\mathbf{\frac{d^2}{T}}}$	$\min\left\{\sqrt[3]{\frac{d^2}{T}}, \sqrt{\frac{d^{34}}{T}}\right\}$	$\sqrt{\mathbf{\frac{d^2}{T}}}$
Convex	$\min\left\{\sqrt[4]{\frac{d^2}{T}}, \sqrt{\frac{d^{34}}{T}}\right\}$	$\sqrt{\mathbf{\frac{d^2}{T}}}$	$\min\left\{\sqrt[4]{\frac{d^2}{T}}, \sqrt{\frac{d^{34}}{T}}\right\}$	$\sqrt{\mathbf{\frac{d^2}{T}}}$

Quadratic Functions: Upper Bounds

Special case of strongly-convex and smooth functions

$$F(\mathbf{w}) = \mathbf{w}^\top A\mathbf{w} + \mathbf{b}^\top \mathbf{w} + c$$

- A is positive-definite (with minimal eigenvalue at least $\lambda > 0$)
- Scaled so that $\|A\|, \|\mathbf{b}\|, |c| \leq 1$

Assumption: we can query in a ball of fixed radius ϵ around optimum $\mathbf{w}^*$

Quadratic Functions: Upper Bounds

Theorem

If $\mathbf{w}^*$ has constant norm, then

$$\mathbb{E}\left[F(\bar{\mathbf{w}}_T) - F(\mathbf{w}^*)\right] \leq \mathcal{O}\left(\frac{1}{\epsilon^2}\frac{d^2}{\lambda T}\right)$$

- Jamieson et al. (NIPS 2012) show $\Omega(\sqrt{d/T})$ lower bound for such quadratic functions
- However, their domain shrinks with T, implying $\epsilon \to 0$, while we assume ϵ is fixed

Quadratic Functions: Upper Bounds

Algorithm

Input: $\lambda, \epsilon > 0$

Initialize $\mathbf{w}_1 = \mathbf{0}$.

for $t = 1, \ldots, T-1$ **do**

Pick $\mathbf{r} \in \{-1, +1\}^d$ uniformly at random

Query noisy function value v at point $\mathbf{w}_t + \frac{\epsilon}{\sqrt{d}}\mathbf{r}$

Let $\tilde{\mathbf{g}} = \frac{\sqrt{d}v}{\epsilon}\mathbf{r}$ // unbiased estimate of $\nabla F(\mathbf{w}_t)$

Let $\mathbf{w}_{t+1} = \Pi_{\bar{\mathcal{W}}}\left(\mathbf{w}_t - \frac{1}{\lambda t}\tilde{\mathbf{g}}\right)$

end for

Return $\bar{\mathbf{w}}_T = \frac{2}{T}\sum_{t=T/2}^{T}\mathbf{w}_t$.

Key Obervation: For quadratic functions, 1-point Gradient estimate is unbiased even if query far from $\mathbf{w}_t$

Quadratic Functions: Upper Bounds

Aside: Result is improvable in some cases:

Theorem

Suppose $F(\mathbf{w}) = R(\mathbf{w}) + \mathbb{E}\left[\mathbf{w}^\top \hat{A}\mathbf{w} + \hat{\mathbf{b}}^\top \mathbf{w} + \hat{c}\right]$, where

- $R(\mathbf{w})$ is a known strongly convex function
- Queries are based on noisy realizations of $\hat{A}, \hat{\mathbf{b}}, \hat{c}$, and can be anywhere in $\mathbb{R}^d$

Then $\exists$ algorithm such that $\mathbb{E}\left[F(\bar{\mathbf{w}}_T) - F(\mathbf{w}^*)\right] \leq \mathcal{O}\left(\frac{d\ \mathbb{E}\left[\|\hat{A}\|_F^2\right]}{\lambda T}\right)$

Example (Ridge Regression)

$$F(\mathbf{w}) = \frac{\lambda}{2}\|\mathbf{w}\|^2 + \mathbb{E}\left[(\langle \mathbf{w}, \mathbf{x}\rangle - y)^2\right]$$

$$\hat{A} = \mathbf{x}\mathbf{x}^\top \quad \Rightarrow \quad \mathbb{E}\left[\|\hat{A}\|_F^2\right] \leq \mathcal{O}(1) \quad \Rightarrow \quad \text{error } \mathcal{O}\left(\frac{d}{\lambda T}\right)$$

Quadratic Functions: Lower Bounds

Theorem

$\forall$ querying strategy, $\exists$ quadratic F (1-strongly convex, Lipschitz, minimized within unit Euclidean ball) such that

$$\mathbb{E}\left[F(\bar{\mathbf{w}}_T) - F(\mathbf{w}^*)\right] \geq \Omega\left(\frac{d^2}{T}\right)$$

Note: Result holds even if we can query anywhere in $\mathbb{R}^d$ (under reasonable noise assumptions)

Proof Idea

Adversary Strategy:

- Pick $\mathbf{e}$ uniformly at random from $\Theta\left(\sqrt{\frac{d}{T}}\right) \times \{-1,+1\}^d$
- $F(\mathbf{w}) = \frac{1}{2}\|\mathbf{w}\|^2 - \langle \mathbf{e}, \mathbf{w} \rangle$
  - Minimized at $\mathbf{e}$
- Given query point $\mathbf{w}$, return $F(\mathbf{w}) + \xi$

Proof idea:

- Due to strong convexity, suboptimality reduced to a **sum of** d hypothesis testing problems:

$$\mathbb{E}\left[F(\bar{\mathbf{w}}_T) - F(\mathbf{w}^*)\right] \geq \mathbb{E}\left[\frac{1}{2}\|\bar{\mathbf{w}}_T - \mathbf{e}\|^2\right] \geq \mathbb{E}\left[\Theta\left(\frac{d}{T}\right) \times \sum_{i=1}^{d} \mathbf{1}_{\bar{w}_i e_i < 0}\right]$$

- Result derived from a relative-entropy argument
  - # samples needed to distinguish $\text{sign}(e_i)$

Quadratic Functions: Lower Bounds

Theorem

Under same conditions as above, $\forall$ querying strategy, $\exists$ quadratic F such that $\mathbb{E}\left[\frac{1}{T}\sum_{t=1}^{T} F(\mathbf{w}_t) - F(\mathbf{w}^*)\right] \geq \Omega\left(\sqrt{\frac{d^2}{T}}\right)$

Proof Idea

- With more careful analysis, relative entropy terms actually depend on $\|\mathbf{w}_t\|$
- For small regret, $\mathbf{w}_t$ must be close to optimum $\mathbf{e}$
- If $\|\mathbf{e}\|$ small $\Rightarrow$ $\|\mathbf{w}_t\|$ must be small $\Rightarrow$ Larger lower bound

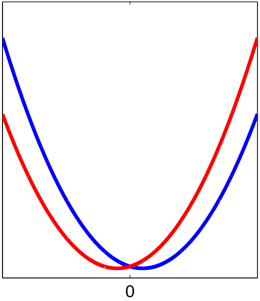

Strongly Convex and Smooth Functions

Theorem

$\forall$ querying strategy, $\exists$ strongly-convex and smooth F (minimized within unit Euclidean ball) such that

$$\mathbb{E}\left[F(\bar{\mathbf{w}}_T) - F(\mathbf{w}^*)\right] \geq \Omega\left(\sqrt{\frac{d^2}{T}}\right)$$

Key Proof Idea

- $F(\mathbf{w}) = \|\mathbf{w}\|^2 - \sum_{i=1}^{d} \frac{e_i w_i}{1+(w_i/e_i)^2}$
  - $\mathbf{e}$ again selected at random
- $F(\mathbf{w}) \approx \|\mathbf{w}\|^2 - 0.9\langle \mathbf{e}, \mathbf{w}\rangle$ near optimum, but $F(\mathbf{w}) \approx \|\mathbf{w}\|^2$ further away

  $\Rightarrow$ Querying far from optimum doesn't give information on $\mathbf{e}$
- $\Rightarrow$ Same $\sqrt{d^2/T}$ lower bound as for *regret* in quadratic case

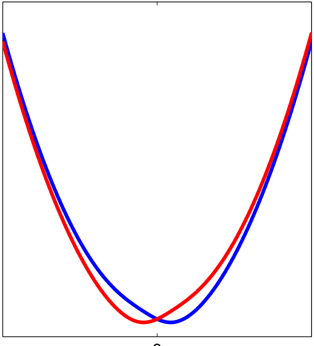

Summary

- Studied the complexity of bandit and derivative-free stochastic convex optimization
- **Exact characterization for strongly-convex and smooth functions**
  - Implies new lower bounds for more general settings
  - Quadratic dependence on the dimension is inevitable
- **"Fast" error rate achievable even without gradients**, for quadratic functions
- **Substantial gaps** between optimization error and regret
- **Open questions:** Complexity of strongly convex (non-smooth) problems? General convex problems?
  - Huge gap: $\Omega(\sqrt{d^2/T})$ vs. $\mathcal{O}\left(\min\left\{\sqrt{\frac{d^{34}}{T}}, \left(\frac{d^2}{T}\right)^{1/4}\right\}\right)$
  - Conjecture: $\Theta\left(\sqrt{d^2/T}\right)$, but **need new algorithms!**