(Bandit) Convex Optimization with Biased Noisy Gradient Oracles

Csaba Szepesvári
Joint work with Xiaowei Hu, L.A. Prashanth, and András György

Reinforcement Learning and Artificial Intelligence (RLAI) Group
Department of Computing Science & AICML
University of Alberta
csaba.szepesvari@ualberta.ca

November 10, 2015

Gatsby Seminar

Outline

Expectation Management

- 80% review (long history!)

Expectation Management

- 80% review (long history!)
- 20% new

Convex Optimization with Noisy Bandit Feedback

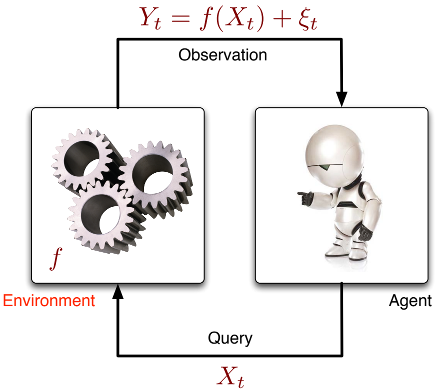

Goal

Assume f convex (smooth etc).
Find a near-minimizer of f using $n > 0$ queries!

Noisy Bandit Feedback $\equiv$ Noisy zeroth-order information

Convex Optimization with Noisy Bandit Feedback

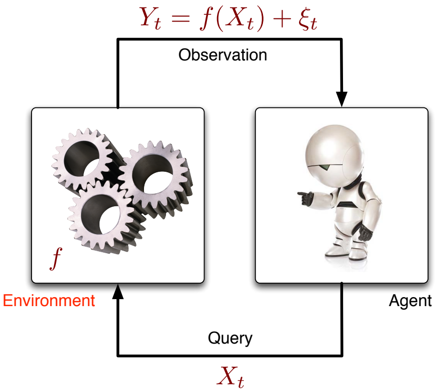

Goal

Assume f convex (smooth etc).

Find a near-minimizer of f using $n > 0$ queries!

Noisy Bandit Feedback $\equiv$ Noisy zeroth-order information

Main Question

How fast can/will/should the optimization error $\Delta_n = \mathbb{E}[f(X_n)] - \inf_{x \in \mathcal{K}} f(x)$ decrease with n?

Example: Resource Allocation

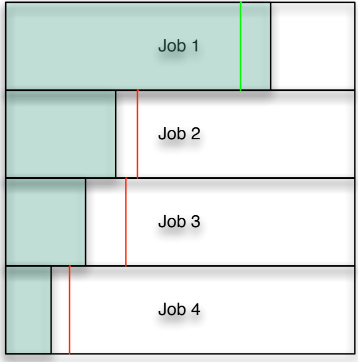

- Distribute memory in between jobs

Example: Resource Allocation

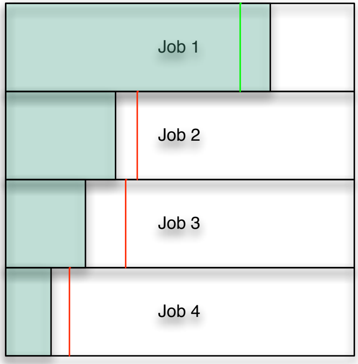

- Distribute memory in between jobs
- Maximize success: $\max_{x \in \mathcal{K}} f(x)$

Example: Resource Allocation

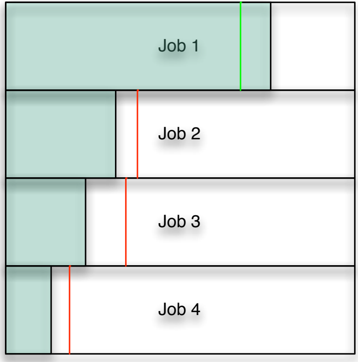

- Distribute memory in between jobs
- Maximize success: $\max_{x \in \mathcal{K}} f(x)$
- Linear constraints: $\mathcal{K} = \{ x \in [0, \infty)^d \ : \ \sum_i x_i = 1 \}$

Example: Resource Allocation

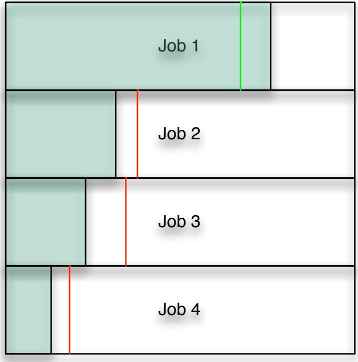

- Distribute memory in between jobs
- Maximize success: $\max_{x \in \mathcal{K}} f(x)$
- Linear constraints: $\mathcal{K} = \{\, x \in [0, \infty)^d \;:\; \sum_i x_i = 1 \,\}$
- Concave objective f: Convex combination of resources not worse than what randomization gives.

Example: Resource Allocation

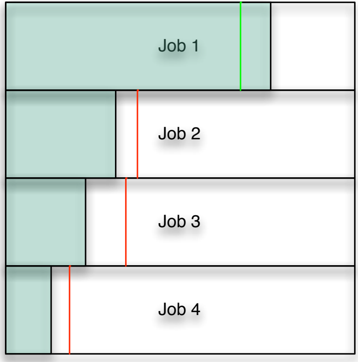

- Distribute memory in between jobs
- Maximize success: $\max_{x \in \mathcal{K}} f(x)$
- Linear constraints: $\mathcal{K} = \{ x \in [0, \infty)^d \ : \ \sum_i x_i = 1 \}$
- Concave objective f: Convex combination of resources not worse than what randomization gives.

The only way to learn about f is by trying different configurations.

Example: Resource Allocation

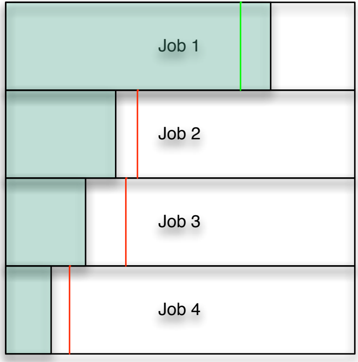

- Distribute memory in between jobs
- Maximize success: $\max_{x \in \mathcal{K}} f(x)$
- Linear constraints: $\mathcal{K} = \{\, x \in [0, \infty)^d \,:\, \sum_i x_i = 1 \,\}$
- Concave objective f: Convex combination of resources not worse than what randomization gives.

The only way to learn about f is by trying different configurations. Feedback: $f(x) + \mathrm{noise}$ for x tried. “Bandit” feedback.

Example: Resource Allocation

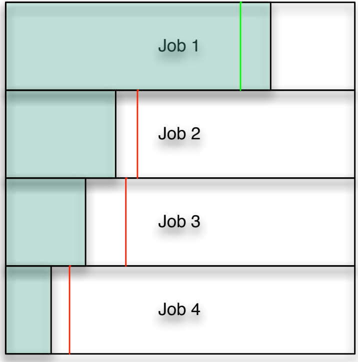

- Distribute memory in between jobs
- Maximize success: $\max_{x \in \mathcal{K}} f(x)$
- Linear constraints: $\mathcal{K} = \left\{ x \in [0, \infty)^d \ : \ \sum_i x_i = 1 \right\}$
- Concave objective f: Convex combination of resources not worse than what randomization gives.

The only way to learn about f is by trying different configurations.
Feedback: $f(x) + \mathrm{noise}$ for x tried. "Bandit" feedback.
Goal is to find the best configuration quickly.

Why Bandit Feedback?

- No other choice

Why Bandit Feedback?

- No other choice
  - Controlling an unknown system

Why Bandit Feedback?

- No other choice
  - Controlling an unknown system
  - Simulation optimization:

Why Bandit Feedback?

- No other choice
  - Controlling an unknown system
  - Simulation optimization:
    - $f(x) = \mathbb{E}\left[F(x, \xi)\right]$ — ξ: simulation noise

Why Bandit Feedback?

- No other choice
  - Controlling an unknown system
  - Simulation optimization:
    - $f(x) = \mathbb{E}[F(x, \xi)]$ — ξ: simulation noise
    - $F(x, \xi)$ is the output of a simulation;

Why Bandit Feedback?

- No other choice
  - Controlling an unknown system
  - Simulation optimization:
    - $f(x) = \mathbb{E}\left[F(x,\xi)\right]$ — ξ: simulation noise
    - $F(x,\xi)$ is the output of a simulation;
    - Gradients are unavailable

Why Bandit Feedback?

- No other choice
  - Controlling an unknown system
  - Simulation optimization:
    - $f(x) = \mathbb{E}\left[F(x, \xi)\right]$ — ξ: simulation noise
    - $F(x, \xi)$ is the output of a simulation;
    - Gradients are unavailable
- By choice:

Why Bandit Feedback?

- No other choice
  - Controlling an unknown system
  - Simulation optimization:
    - $f(x) = \mathbb{E}[F(x, \xi)]$ — ξ: simulation noise
    - $F(x, \xi)$ is the output of a simulation;
    - Gradients are unavailable
- By choice:
  - Gradient is too expensive/complicated to compute

Why Bandit Feedback?

- No other choice
  - Controlling an unknown system
  - Simulation optimization:
    - $f(x) = \mathbb{E}[F(x,\xi)]$ — ξ: simulation noise
    - $F(x,\xi)$ is the output of a simulation;
    - Gradients are unavailable
- By choice:
  - Gradient is too expensive/complicated to compute
  - (Can this be justified?)

Subclasses of Convex Problems

Deviation from linearity, or "curviness":

$$D_f(w||v) \doteq f(w) - \{f(v) + \langle \nabla f(v), w - v \rangle\}.$$

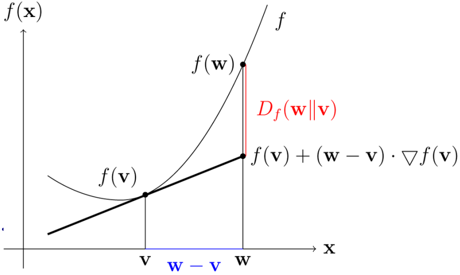

Subclasses of Convex Problems

Deviation from linearity, or "curviness":

$$D_f(w||v) \doteq f(w) - \{f(v) + \langle \nabla f(v), w - v \rangle\}.$$

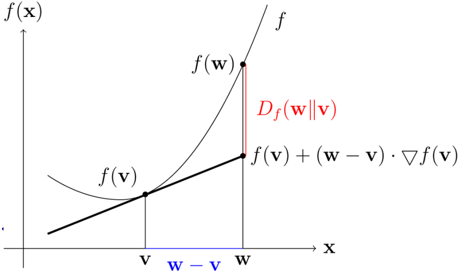

- Smoothness:

$$D_f(x||y) \leq \frac{L}{2} \|x - y\|^2 .$$

Subclasses of Convex Problems

Deviation from linearity, or "curviness":

$$D_f(w||v) \doteq f(w) - \{f(v) + \langle \nabla f(v), w - v \rangle\} .$$

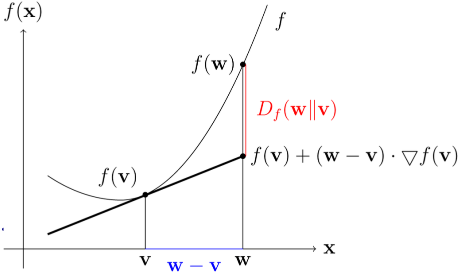

- Smoothness:
$$D_f(x||y) \leq \frac{L}{2} \|x - y\|^2 .$$
- Strong convexity:
$$D_f(x||y) \geq \frac{\mu}{2} \|x - y\|^2 .$$

Subclasses of Convex Problems

Deviation from linearity, or "curviness":

$$D_f(w\|v) \doteq f(w) - \{f(v) + \langle \nabla f(v), w - v \rangle\} .$$

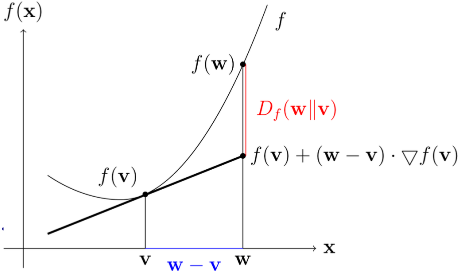

- Smoothness:
  $$D_f(x\|y) \le \frac{L}{2} \|x - y\|^2 .$$
- Strong convexity:
  $$D_f(x\|y) \ge \frac{\mu}{2} \|x - y\|^2 .$$

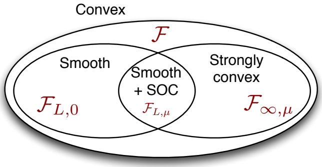

Subclasses of Convex Problems

$\mathcal{K} \subset \mathbb{R}^d$ convex, closed, non-empty.

Deviation from linearity, or "curviness":

$$D_f(w||v) \doteq f(w) - \{f(v) + \langle \nabla f(v), w - v \rangle\} .$$

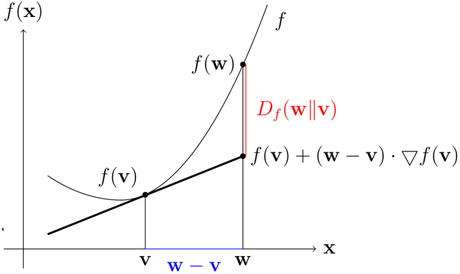

- Smoothness:
  $$D_f(x||y) \leq \frac{L}{2} \|x - y\|^2 .$$
- Strong convexity:
  $$D_f(x||y) \geq \frac{\mu}{2} \|x - y\|^2 .$$

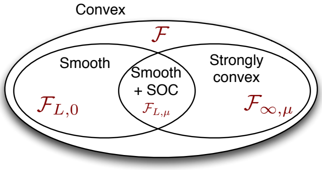

Convex Optimization 101: Ellipsoid Method

Assumption

First-order information, $(\nabla f(\cdot), f(\cdot))$, can be obtained at any point.

Convex Optimization 101: Ellipsoid Method

Assumption

First-order information, $(\nabla f(\cdot), f(\cdot))$, can be obtained at any point.

Ellipsoid Method (Shor, 1970)

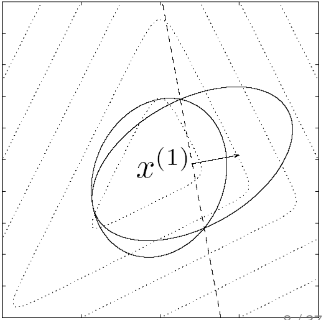

Convex Optimization 101: Ellipsoid Method

Assumption

First-order information, $(\nabla f(\cdot), f(\cdot))$, can be obtained at any point.

Ellipsoid Method (Shor, 1970)

- $\Delta_n^{\mathcal{E},d} = C\exp(-\frac{n}{d^2})$ (Nemirovskii and Yudin, 1983)

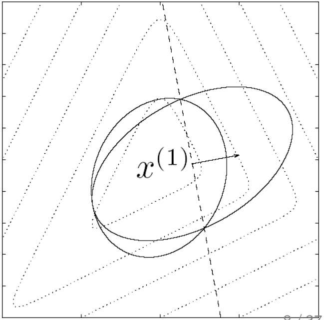

Convex Optimization 101: Ellipsoid Method

Assumption

First-order information, $(\nabla f(\cdot), f(\cdot))$, can be obtained at any point.

Ellipsoid Method (Shor, 1970)

- $\Delta_n^{\mathcal{E},d} = C \exp(-\frac{n}{d^2})$ (Nemirovskii and Yudin, 1983)
- Matching lower bound when d small and n large.

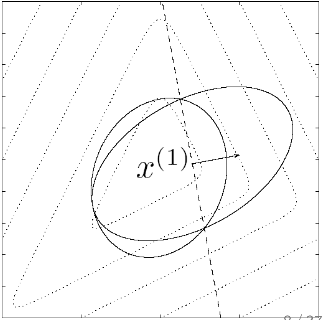

Convex Optimization 101: Ellipsoid Method

Assumption

First-order information, $(\nabla f(\cdot), f(\cdot))$, can be obtained at any point.

Ellipsoid Method (Shor, 1970)

- $\Delta_n^{\mathcal{E},d} = C \exp(-\frac{n}{d^2})$ (Nemirovskii and Yudin, 1983)
- Matching lower bound when d small and n large.

  ... but

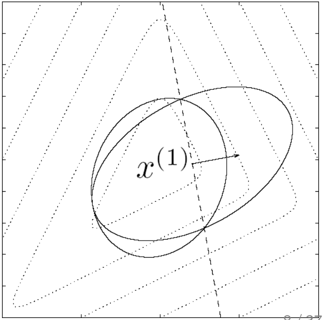

Convex Optimization 101: Ellipsoid Method

Assumption

First-order information, $(\nabla f(\cdot), f(\cdot))$, can be obtained at any point.

Ellipsoid Method (Shor, 1970)

- $\Delta_n^{\mathcal{E},d} = C \exp(-\frac{n}{d^2})$ (Nemirovskii and Yudin, 1983)
- Matching lower bound when d small and n large.

  $\ldots$ but
- Each update takes $O(d^2)$ MADDs (prohibitive when d large).

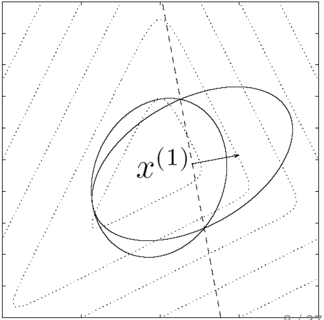

Convex Optimization 101: Ellipsoid Method

Assumption

First-order information, $(\nabla f(\cdot), f(\cdot))$, can be obtained at any point.

Ellipsoid Method (Shor, 1970)

- $\Delta_n^{\mathcal{E},d} = C \exp(-\frac{n}{d^2})$ (Nemirovskii and Yudin, 1983)
- Matching lower bound when d small and n large.
  ... but
- Each update takes $O(d^2)$ MADDs (prohibitive when d large).
- Suboptimal for $d/n^2 \gg 1$: $\lim_{d\to\infty} \Delta_n^{\mathcal{E},d} = \Omega(1)$, while $\Delta_n^* = \Theta(1/\sqrt{n})$.

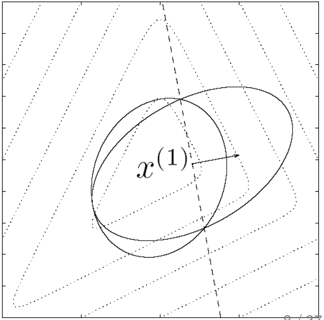

Convex Optimization 101: Gradient Methods

Think: "d large". Update complexity is $O(d)$.

- Subgradient method: Optimal for convex problems:

$$\Delta_n = O(\sqrt{1/n}).$$

No dependence on d!

Convex Optimization 101: Gradient Methods

Think: "d large". Update complexity is $O(d)$.

- Subgradient method: Optimal for convex problems:

$$\Delta_n = O(\sqrt{1/n}).$$

No dependence on d!

- Gradient method + momentum term $\equiv$
accelerated gradient method (Nesterov, 2004).

Convex Optimization 101: Gradient Methods

Think: "d large". Update complexity is $O(d)$.

- Subgradient method: Optimal for convex problems:

$$\Delta_n = O(\sqrt{1/n}).$$

No dependence on d!

- Gradient method + momentum term $\equiv$
  accelerated gradient method (Nesterov, 2004).
  - Optimal for smooth convex problems, $\mathcal{F}_{L,0}$:

$$\Delta_n = O(L/n^2).$$

Convex Optimization 101: Gradient Methods

Think: "d large". Update complexity is $O(d)$.

- Subgradient method: Optimal for convex problems:

$$\Delta_n = O(\sqrt{1/n}).$$

No dependence on d!

- Gradient method + momentum term $\equiv$
  accelerated gradient method (Nesterov, 2004).
  - Optimal for smooth convex problems, $\mathcal{F}_{L,0}$:

$$\Delta_n = O(L/n^2).$$

  - Optimal for strongly convex, smooth problems, $\mathcal{F}_{L,\mu}$:

$$\Delta_n = O\left(\exp\left(-\frac{n}{\sqrt{L/\mu}}\right)\right).$$

Outline

Results under Noisy Bandit Feedback

Recall: $f(x) = \mathbb{E}[F(x, \xi)]$.

Assumptions

(A1) $F(x, \xi)$ can be obtained at any point. "Uncontrolled noise"

Results under Noisy Bandit Feedback

Recall: $f(x) = \mathbb{E}\left[F(x, \xi)\right]$.

Assumptions

(A1) $F(x, \xi)$ can be obtained at any point. "Uncontrolled noise"

(A2) ξ can be kept fixed between queries. "Controlled noise"

Controlled Noise – Smooth Case

Simple story!

Controlled Noise – Smooth Case

Simple story!

Gradient method

Controlled Noise – Smooth Case

Simple story!

Gradient method

$\Delta_n \leq C\sqrt{d^2/n}$.

Controlled Noise – Smooth Case

Simple story!

Gradient method

$\Delta_n \leq C\sqrt{d^2/n}$.

Nesterov (2011) , Duchi et al. (2015)

Controlled Noise – Smooth Case

Simple story!

Gradient method

$\Delta_n \leq C\sqrt{d^2/n}$.

Nesterov (2011) , Duchi et al. (2015)

Optimal!

Uncontrolled Noise: Big Gaps

Ellipsoid method and relatives (Nemirovskii and Yudin (1983), Section 9.4)

Uncontrolled Noise: Big Gaps

Ellipsoid method and relatives (Nemirovskii and Yudin (1983), Section 9.4)

- Agarwal et al. (2013) – $\sqrt{d^{33}/n}$

Uncontrolled Noise: Big Gaps

Ellipsoid method and relatives (Nemirovskii and Yudin (1983), Section 9.4)

- Agarwal et al. (2013) – $\sqrt{d^{33}/n}$
- Liang et al. (2014) – $\sqrt{d^{14}/n}$

Uncontrolled Noise: Big Gaps

Ellipsoid method and relatives (Nemirovskii and Yudin (1983), Section 9.4)

- Agarwal et al. (2013) – $\sqrt{d^{33}/n}$
- Liang et al. (2014) – $\sqrt{d^{14}/n}$

Gradient methods: (Nemirovskii and Yudin (1983), Section 9.3)

Uncontrolled Noise: Big Gaps

Ellipsoid method and relatives (Nemirovskii and Yudin (1983), Section 9.4)

- Agarwal et al. (2013) – $\sqrt{d^{33}/n}$
- Liang et al. (2014) – $\sqrt{d^{14}/n}$

Gradient methods: (Nemirovskii and Yudin (1983), Section 9.3)

- Convex: – $(d^2/n)^{1/4}$
  (Nemirovskii and Yudin, 1983; Flaxman et al., 2005)

Uncontrolled Noise: Big Gaps

Ellipsoid method and relatives (Nemirovskii and Yudin (1983), Section 9.4)

- Agarwal et al. (2013) – $\sqrt{d^{33}/n}$
- Liang et al. (2014) – $\sqrt{d^{14}/n}$

Gradient methods: (Nemirovskii and Yudin (1983), Section 9.3)

- Convex: – $(d^2/n)^{1/4}$
  (Nemirovskii and Yudin, 1983; Flaxman et al., 2005)
- Smooth: – $(d^2/n)^{1/3}$
  (Nemirovskii and Yudin, 1983; Saha and Tewari, 2011)

Uncontrolled Noise: Big Gaps

Ellipsoid method and relatives (Nemirovskii and Yudin (1983), Section 9.4)

- Agarwal et al. (2013) – $\sqrt{d^{33}/n}$
- Liang et al. (2014) – $\sqrt{d^{14}/n}$

Gradient methods: (Nemirovskii and Yudin (1983), Section 9.3)

- Convex: – $(d^2/n)^{1/4}$
  (Nemirovskii and Yudin, 1983; Flaxman et al., 2005)
- Smooth: – $(d^2/n)^{1/3}$
  (Nemirovskii and Yudin, 1983; Saha and Tewari, 2011)
- Smooth + SOC: $\sqrt{d^2/n}$
  (Hazan and Levy, 2014)

Uncontrolled Noise: Big Gaps

Ellipsoid method and relatives (Nemirovskii and Yudin (1983), Section 9.4)

- Agarwal et al. (2013) – $\sqrt{d^{33}/n}$
- Liang et al. (2014) – $\sqrt{d^{14}/n}$

Gradient methods: (Nemirovskii and Yudin (1983), Section 9.3)

- Convex: – $(d^2/n)^{1/4}$
  (Nemirovskii and Yudin, 1983; Flaxman et al., 2005)
- Smooth: – $(d^2/n)^{1/3}$
  (Nemirovskii and Yudin, 1983; Saha and Tewari, 2011)
- Smooth + SOC: $\sqrt{d^2/n}$
  (Hazan and Levy, 2014)

Lower bound: $\sqrt{d^2/n}$ (Shamir, 2012).

Uncontrolled Noise: Big Gaps

Ellipsoid method and relatives (Nemirovskii and Yudin (1983), Section 9.4)

- Agarwal et al. (2013) – $\sqrt{d^{33}/n}$
- Liang et al. (2014) – $\sqrt{d^{14}/n}$

Gradient methods: (Nemirovskii and Yudin (1983), Section 9.3)

- Convex: – $(d^2/n)^{1/4}$
  (Nemirovskii and Yudin, 1983; Flaxman et al., 2005)
- Smooth: – $(d^2/n)^{1/3}$
  (Nemirovskii and Yudin, 1983; Saha and Tewari, 2011)
- Smooth + SOC: $\sqrt{d^2/n}$
  (Hazan and Levy, 2014)

Lower bound: $\sqrt{d^2/n}$ (Shamir, 2012).

Can we do better?

Uncontrolled Noise: Big Gaps

Ellipsoid method and relatives (Nemirovskii and Yudin (1983), Section 9.4)

- Agarwal et al. (2013) – $\sqrt{d^{33}/n}$
- Liang et al. (2014) – $\sqrt{d^{14}/n}$

Gradient methods: (Nemirovskii and Yudin (1983), Section 9.3)

- Convex: – $(d^2/n)^{1/4}$
  (Nemirovskii and Yudin, 1983; Flaxman et al., 2005)
- Smooth: – $(d^2/n)^{1/3}$
  (Nemirovskii and Yudin, 1983; Saha and Tewari, 2011)
- Smooth + SOC: $\sqrt{d^2/n}$
  (Hazan and Levy, 2014)

Lower bound: $\sqrt{d^2/n}$ (Shamir, 2012).

Can we do better? ... using a "clever" gradient method maybe?

Outline

Gradient Estimation – Noise-free Feedback

Finite-differences (Kiefer and Wolfowitz, 1952):

$$g_i = \frac{1}{\delta}\left(f(x+\delta e_i) - f(x)\right) , \quad i = 1, \ldots, d \,.$$

Taylor-series expansion:

$$f(x+\delta e_i) = f(x) + \delta \, \nabla f(x) e_i + \delta^2 \, e_i^\top \nabla^2 f(x) e_i + O(\delta^3).$$

Accuracy: $\|g - \nabla f(x)\|_2 = O(\sqrt{d}\,\delta)$.

Needs $d+1$ queries.

Two-sided Differences

Improved estimate:

$$g_i = \frac{1}{2\delta}\left(f(x+\delta e_i) - f(x-\delta e_i)\right), \quad i = 1, \ldots, d.$$

Taylor-series expansions:

$$f(x+\delta e_i) = f(x) + \delta\,\nabla f(x) e_i + \delta^2\, e_i^\top \nabla^2 f(x) e_i + O(\delta^3).$$
$$f(x-\delta e_i) = f(x) - \delta\,\nabla f(x) e_i + \delta^2\, e_i^\top \nabla^2 f(x) e_i + O(\delta^3).$$

Accuracy: $\|g - \nabla f(x)\|_2 = O(\sqrt{d}\delta^2)$.

Needs $2d$ queries.

Two-sided Differences – Noisy Feedback

Improved estimate:

$$G_i = \frac{1}{2\delta}\left\{f(x+\delta e_i) + \xi_i^+ - \left(f(x-\delta e_i) + \xi_i^-\right)\right\}, \quad i = 1, \ldots, d.$$

Two-sided Differences – Noisy Feedback

Improved estimate:

$$G_i = \frac{1}{2\delta}\left\{f(x+\delta e_i) + \xi_i^+ - \left(f(x-\delta e_i) + \xi_i^-\right)\right\}, \quad i = 1, \ldots, d.$$

Assumption: $\mathbb{E}\left[\xi^\pm\right] = 0$, $\mathbb{E}\left[(\xi^\pm)\right] \leq \sigma^2 < +\infty$.

Two-sided Differences – Noisy Feedback

Improved estimate:

$$G_i = \frac{1}{2\delta}\left\{f(x+\delta e_i) + \xi_i^+ - \left(f(x-\delta e_i) + \xi_i^-\right)\right\}, \quad i = 1, \ldots, d.$$

Assumption: $\mathbb{E}\left[\xi^\pm\right] = 0$, $\mathbb{E}\left[(\xi^\pm)\right] \leq \sigma^2 < +\infty$.
Note: $\mathbb{E}\left[G_i\right] = g_i$.

Two-sided Differences – Noisy Feedback

Improved estimate:

$$G_i = \frac{1}{2\delta}\left\{f(x+\delta e_i) + \xi_i^+ - \left(f(x-\delta e_i) + \xi_i^-\right)\right\}, \quad i = 1, \ldots, d.$$

Assumption: $\mathbb{E}[\xi^\pm] = 0$, $\mathbb{E}[(\xi^\pm)] \le \sigma^2 < +\infty$.
Note: $\mathbb{E}[G_i] = g_i$. Hence

$$\|\mathbb{E}[G] - \nabla f(x)\|_2 = O(\sqrt{d}\delta^2)\,. \qquad \text{bias}$$

Two-sided Differences – Noisy Feedback

Improved estimate:

$$G_i = \frac{1}{2\delta}\left\{f(x+\delta e_i) + \xi_i^+ - (f(x-\delta e_i) + \xi_i^-)\right\}, \quad i = 1, \dots, d.$$

Assumption: $\mathbb{E}[\xi^\pm] = 0$, $\mathbb{E}[(\xi^\pm)] \leq \sigma^2 < +\infty$.
Note: $\mathbb{E}[G_i] = g_i$. Hence

$$\|\mathbb{E}[G] - \nabla f(x)\|_2 = O(\sqrt{d}\delta^2). \qquad \text{bias}$$

Second moment: $\mathbb{E}\left[\|G\|_2^2\right] =?$

Two-sided Differences – Noisy Feedback

Improved estimate:

$$G_i = \frac{1}{2\delta}\left\{f(x+\delta e_i) + \xi_i^+ - (f(x-\delta e_i) + \xi_i^-)\right\}, \quad i = 1, \ldots, d.$$

Assumption: $\mathbb{E}[\xi^\pm] = 0$, $\mathbb{E}[(\xi^\pm)] \leq \sigma^2 < +\infty$.
Note: $\mathbb{E}[G_i] = g_i$. Hence

$$\|\mathbb{E}[G] - \nabla f(x)\|_2 = O(\sqrt{d}\delta^2)\,. \qquad \text{bias}$$

Second moment: $\mathbb{E}\left[\|G\|_2^2\right] =?$

$G_i = g_i + \frac{\xi_i^+ - \xi_i^-}{2\delta}$

Two-sided Differences – Noisy Feedback

Improved estimate:

$$G_i = \frac{1}{2\delta}\left\{f(x+\delta e_i) + \xi_i^+ - (f(x-\delta e_i) + \xi_i^-)\right\}, \quad i = 1, \ldots, d.$$

Assumption: $\mathbb{E}[\xi^\pm] = 0$, $\mathbb{E}[(\xi^\pm)] \leq \sigma^2 < +\infty$.
Note: $\mathbb{E}[G_i] = g_i$. Hence

$$\|\mathbb{E}[G] - \nabla f(x)\|_2 = O(\sqrt{d}\delta^2). \qquad \text{bias}$$

Second moment: $\mathbb{E}\left[\|G\|_2^2\right] =?$

$G_i = g_i + \frac{\xi_i^+ - \xi_i^-}{2\delta}$, hence $\mathbb{E}[G_i^2] = g_i^2 + \frac{2\sigma^2}{4\delta^2} = g_i^2 + \frac{\sigma^2}{2\delta^2}$

Two-sided Differences – Noisy Feedback

Improved estimate:

$$G_i = \frac{1}{2\delta}\left\{f(x+\delta e_i) + \xi_i^+ - (f(x-\delta e_i) + \xi_i^-)\right\}, \quad i = 1, \ldots, d.$$

Assumption: $\mathbb{E}[\xi^\pm] = 0$, $\mathbb{E}[(\xi^\pm)] \le \sigma^2 < +\infty$.
Note: $\mathbb{E}[G_i] = g_i$. Hence

$$\|\mathbb{E}[G] - \nabla f(x)\|_2 = O(\sqrt{d}\delta^2). \qquad \text{bias}$$

Second moment: $\mathbb{E}\left[\|G\|_2^2\right] = ?$

$G_i = g_i + \frac{\xi_i^+ - \xi_i^-}{2\delta}$, hence $\mathbb{E}[G_i^2] = g_i^2 + \frac{2\sigma^2}{4\delta^2} = g_i^2 + \frac{\sigma^2}{2\delta^2}$ and

$$\mathbb{E}\left[\|G\|_2^2\right] = \|g\|_2^2 + O\left(\frac{d}{\delta^2}\right).$$

Randomization (Noise-Free Bandit Feedback)

Can we reduce the number of queries?

Randomization (Noise-Free Bandit Feedback)

Can we reduce the number of queries?
Idea: Randomize! $I \sim p(\cdot)$ a positive pmf on $\{1, \dots, d\}$.

Randomization (Noise-Free Bandit Feedback)

Can we reduce the number of queries?
Idea: Randomize! $I \sim p(\cdot)$ a positive pmf on $\{1, \ldots, d\}$.
Choose

$$G_i = \frac{1}{p(i)} \frac{f(x+\delta e_I) - f(x-\delta e_I)}{2\delta} e_{I,i}$$
$$= \begin{cases} \frac{1}{p(I)} \frac{f(x+\delta e_I) - f(x-\delta e_I)}{2\delta}, & I = i\,; \\ 0, & \text{otherwise}\,. \end{cases}$$

Randomization (Noise-Free Bandit Feedback)

Can we reduce the number of queries?
Idea: Randomize! $I \sim p(\cdot)$ a positive pmf on $\{1, \ldots, d\}$.
Choose

$$G_i = \frac{1}{p(i)} \frac{f(x + \delta e_I) - f(x - \delta e_I)}{2\delta} e_{I,i}$$
$$= \begin{cases} \frac{1}{p(I)} \frac{f(x+\delta e_I) - f(x-\delta e_I)}{2\delta}, & I = i\,; \\ 0, & \text{otherwise}\,. \end{cases}$$

Only 2-queries, regardless of d!

Randomization (Noise-Free Bandit Feedback)

Can we reduce the number of queries?
Idea: Randomize! $I \sim p(\cdot)$ a positive pmf on $\{1, \ldots, d\}$.
Choose

$$G_i = \frac{1}{p(i)} \frac{f(x+\delta e_I) - f(x-\delta e_I)}{2\delta} e_{I,i}$$
$$= \begin{cases} \frac{1}{p(I)} \frac{f(x+\delta e_I) - f(x-\delta e_I)}{2\delta}, & I = i\,; \\ 0, & \text{otherwise}\,. \end{cases}$$

Only 2-queries, regardless of d!
$\mathbb{E}[G_i] = g_i$!

Randomization (Noise-Free Bandit Feedback)

Can we reduce the number of queries?
Idea: Randomize! $I \sim p(\cdot)$ a positive pmf on $\{1, \ldots, d\}$.
Choose

$$
\begin{aligned}
G_i &= \frac{1}{p(i)} \, \frac{f(x+\delta e_I) - f(x-\delta e_I)}{2\delta} \, e_{I,i} \\
&= \begin{cases} \frac{1}{p(I)} \, \frac{f(x+\delta e_I) - f(x-\delta e_I)}{2\delta}, & I = i\,; \\ 0, & \text{otherwise}\,. \end{cases}
\end{aligned}
$$

Only 2-queries, regardless of d!
$\mathbb{E}\left[G_i\right] = g_i$! Hence, $\left\|\mathbb{E}\left[G\right] - \nabla f(x)\right\|_2 = O(\sqrt{d}\delta^2)$.

Randomization (Noise-Free Bandit Feedback)

Can we reduce the number of queries?
Idea: Randomize! $I \sim p(\cdot)$ a positive pmf on $\{1, \ldots, d\}$.
Choose

$$G_i = \frac{1}{p(i)} \frac{f(x + \delta e_I) - f(x - \delta e_I)}{2\delta} e_{I,i}$$
$$= \begin{cases} \frac{1}{p(I)} \frac{f(x+\delta e_I) - f(x-\delta e_I)}{2\delta}, & I = i\,; \\ 0, & \text{otherwise}\,. \end{cases}$$

Only 2-queries, regardless of d!
$\mathbb{E}\left[G_i\right] = g_i$! Hence, $\left\|\mathbb{E}\left[G\right] - \nabla f(x)\right\|_2 = O(\sqrt{d}\delta^2)$.

Second moment?

Second Moment from Randomization

$$G_i = \frac{1}{p(i)} \frac{f(x + \delta e_I) - f(x - \delta e_I)}{2\delta} \, e_{I,i}$$

Second Moment from Randomization

$$G_i = \frac{1}{p(i)} \frac{f(x+\delta e_I) - f(x-\delta e_I)}{2\delta} e_{I,i}$$

Taylor-series expansion:

$$f(x+\delta e_i) = f(x) + \delta \nabla f(x) e_i + \delta^2 e_i^\top \nabla^2 f(x) e_i + O(\delta^3),$$
$$f(x-\delta e_i) = f(x) - \delta \nabla f(x) e_i + \delta^2 e_i^\top \nabla^2 f(x) e_i + O(\delta^3).$$

Second Moment from Randomization

$$G_i = \frac{1}{p(i)} \frac{f(x+\delta e_I) - f(x-\delta e_I)}{2\delta} e_{I,i}$$

Taylor-series expansion:

$$f(x+\delta e_i) = f(x) + \delta \nabla f(x) e_i + \delta^2 e_i^\top \nabla^2 f(x) e_i + O(\delta^3),$$
$$f(x-\delta e_i) = f(x) - \delta \nabla f(x) e_i + \delta^2 e_i^\top \nabla^2 f(x) e_i + O(\delta^3).$$

$$\begin{aligned} G_i^2 &= \frac{\mathbb{I}_{\{I=i\}}}{p^2(i)} \frac{(\delta f_i'(x) + c_3(\delta))^2}{4\delta^2} = \frac{\mathbb{I}_{\{I=i\}}}{p^2(i)} \frac{(\delta^2 (f_i'(x))^2 + c_4(\delta))}{\delta^2} \\ &= \frac{\mathbb{I}_{\{I=i\}}}{p^2(i)} \left\{ (f_i'(x))^2 + O(\delta^2) \right\} \end{aligned}$$

Second Moment from Randomization

$$G_i = \frac{1}{p(i)} \frac{f(x+\delta e_I) - f(x-\delta e_I)}{2\delta} e_{I,i}$$

Taylor-series expansion:

$$f(x+\delta e_i) = f(x) + \delta \nabla f(x) e_i + \delta^2 e_i^\top \nabla^2 f(x) e_i + O(\delta^3),$$
$$f(x-\delta e_i) = f(x) - \delta \nabla f(x) e_i + \delta^2 e_i^\top \nabla^2 f(x) e_i + O(\delta^3).$$

$$G_i^2 = \frac{\mathbb{I}_{\{I=i\}}}{p^2(i)} \frac{(\delta f_i'(x) + c_3(\delta))^2}{4\delta^2} = \frac{\mathbb{I}_{\{I=i\}}}{p^2(i)} \frac{(\delta^2 (f_i'(x))^2 + c_4(\delta))}{\delta^2}$$
$$= \frac{\mathbb{I}_{\{I=i\}}}{p^2(i)} \left\{ (f_i'(x))^2 + O(\delta^2) \right\}$$

Hence, $\mathbb{E}[G_i^2] = O(1/p(i))$, so at best $\mathbb{E}\left[\|G\|_2^2\right] = O(d^2)$. Hmm..

Noisy Bandit Feedback

$$\tilde{G}_i = \frac{1}{p(i)} \, \frac{\left(f\left(x + \delta e_I\right) + \xi^+\right) - \left(f\left(x - \delta e_I\right) + \xi^-\right)}{2\delta} \, e_{I,i}$$

Noisy Bandit Feedback

$$\tilde{G}_i = \frac{1}{p(i)} \frac{(f(x + \delta e_I) + \xi^+) - (f(x - \delta e_I) + \xi^-)}{2\delta} e_{I,i}$$

Hence, $\tilde{G}_i = \frac{\mathbb{I}_{\{I=i\}}}{p(i)} \frac{\xi^+ - \xi^-}{2\delta} + G_i$

Noisy Bandit Feedback

$$\tilde{G}_i = \frac{1}{p(i)} \frac{(f(x+\delta e_I)+\xi^+) - (f(x-\delta e_I)+\xi^-)}{2\delta} e_{I,i}$$

Hence, $\tilde{G}_i = \frac{\mathbb{I}_{\{I=i\}}}{p(i)} \frac{\xi^+ - \xi^-}{2\delta} + G_i$ and

$$\mathbb{E}\left[\tilde{G}_i\right] = \mathbb{E}\left[G_i\right],$$

$$\mathbb{E}\left[\tilde{G}_i^2\right] = \frac{1}{p(i)} \frac{\mathbb{E}\left[(\xi^+ - \xi^-)^2\right]}{4\delta^2} + \mathbb{E}\left[G_i^2\right],$$

Noisy Bandit Feedback

$$\tilde{G}_i = \frac{1}{p(i)} \frac{(f(x+\delta e_I)+\xi^+) - (f(x-\delta e_I)+\xi^-)}{2\delta} e_{I,i}$$

Hence, $\tilde{G}_i = \frac{\mathbb{I}_{\{I=i\}}}{p(i)} \frac{\xi^+ - \xi^-}{2\delta} + G_i$ and

$$\mathbb{E}\left[\tilde{G}_i\right] = \mathbb{E}\left[G_i\right],$$

$$\mathbb{E}\left[\tilde{G}_i^2\right] = \frac{1}{p(i)} \frac{\mathbb{E}\left[(\xi^+ - \xi^-)^2\right]}{4\delta^2} + \mathbb{E}\left[G_i^2\right],$$

so

$$\left\| \mathbb{E}\left[\tilde{G}\right]_2 - \nabla f(x) \right\|_2 = O(\sqrt{d}\delta^2),$$

$$\mathbb{E}\left[\left\|\tilde{G}\right\|_2^2\right] = O(d^2(1+1/\delta^2)). \qquad \text{harsh!}$$

Summary

- Noise-free observations (or controlled noise):

Summary

- Noise-free observations (or controlled noise):
  - Bias: $O(\delta^2)$

Summary

- Noise-free observations (or controlled noise):
  - Bias: $O(\delta^2)$
  - Second-moment: $O(1)$

Summary

- Noise-free observations (or controlled noise):
  - Bias: $O(\delta^2)$
  - Second-moment: $O(1)$
- Noisy observations (a.k.a. uncontrolled noise):

Summary

- Noise-free observations (or controlled noise):
  - Bias: $O(\delta^2)$
  - Second-moment: $O(1)$
- Noisy observations (a.k.a. uncontrolled noise):
  - Bias: $O(\delta^2)$

Summary

- Noise-free observations (or controlled noise):
  - Bias: $O(\delta^2)$
  - Second-moment: $O(1)$
- Noisy observations (a.k.a. uncontrolled noise):
  - Bias: $O(\delta^2)$
  - Second moment: $O(\delta^{-2})$

Summary

- Noise-free observations (or controlled noise):
  - Bias: $O(\delta^2)$
  - Second-moment: $O(1)$
- Noisy observations (a.k.a. uncontrolled noise):
  - Bias: $O(\delta^2)$
  - Second moment: $O(\delta^{-2})$

This assumed $f \in \mathcal{C}^3$. Holds also for f convex, smooth.

Other Methods

General two-point estimate:

$$G = \frac{(f(x+U)+\xi^+) - (f(x-U)+\xi^-)}{2\delta} V \,.$$

Other Methods

General two-point estimate:

$$G = \frac{(f(x+U)+\xi^{+})-(f(x-U)+\xi^{-})}{2\delta} V\,.$$

Choose U, V such that $\mathbb{E}\left[VU^{\top}\right] = I$, $\mathbb{E}\left[V\right] = 0$.

Other Methods

General two-point estimate:

$$G = \frac{(f(x+U)+\xi^{+}) - (f(x-U)+\xi^{-})}{2\delta} V \,.$$

Choose U, V such that $\mathbb{E}\left[VU^{\top}\right] = I$, $\mathbb{E}\left[V\right] = 0$.
One-point estimate!

$$G = \frac{(f(x+U)+\xi^{+})}{\delta} V \,.$$

Other Methods

General two-point estimate:

$$G = \frac{(f(x+U)+\xi^+) - (f(x-U)+\xi^-)}{2\delta} V \,.$$

Choose U, V such that $\mathbb{E}\left[VU^\top\right] = I$, $\mathbb{E}\left[V\right] = 0$.

One-point estimate!

$$G = \frac{(f(x+U)+\xi^+)}{\delta} V \,.$$

Choose U, V such that $\mathbb{E}\left[VU^\top\right] = I$, $\mathbb{E}\left[V\right] = 0$.

Other Methods

General two-point estimate:

$$G = \frac{(f(x+U)+\xi^{+}) - (f(x-U)+\xi^{-})}{2\delta} V \,.$$

Choose U, V such that $\mathbb{E}\left[VU^{\top}\right] = I$, $\mathbb{E}\left[V\right] = 0$.
One-point estimate!

$$G = \frac{(f(x+U)+\xi^{+})}{\delta} V \,.$$

Choose U, V such that $\mathbb{E}\left[VU^{\top}\right] = I$, $\mathbb{E}\left[V\right] = 0$. Works??

Other Methods

General two-point estimate:

$$G = \frac{(f(x+U)+\xi^{+}) - (f(x-U)+\xi^{-})}{2\delta} V\,.$$

Choose U, V such that $\mathbb{E}\left[VU^{\top}\right] = I$, $\mathbb{E}\left[V\right] = 0$.
One-point estimate!

$$G = \frac{(f(x+U)+\xi^{+})}{\delta} V\,.$$

Choose U, V such that $\mathbb{E}\left[VU^{\top}\right] = I$, $\mathbb{E}\left[V\right] = 0$. Works??
$\mathbb{E}\left[G\right] = \mathbb{E}\left[G - \frac{f(x)}{\delta} V\right] = \mathbb{E}\left[\frac{(f(x+U)+\xi^{+}) - f(x)}{\delta} V\right]$.

“Zoo” of Gradient Estimation Methods

- $U \sim \delta\,\mathcal{N}(0, I)$, $V = \delta^{-1}\,U$

"Zoo" of Gradient Estimation Methods

- $U \sim \delta \mathcal{N}(0, I)$, $V = \delta^{-1} U$
  - Smoothed functional scheme by Katkovnik and Kulchitsky (1972);

“Zoo” of Gradient Estimation Methods

- $U \sim \delta\,\mathcal{N}(0, I)$, $V = \delta^{-1}\,U$
  - Smoothed functional scheme by Katkovnik and Kulchitsky (1972);
  - Refined by Polyak and Tsybakov (1990);

“Zoo” of Gradient Estimation Methods

- $U \sim \delta\,\mathcal{N}(0, I)$, $V = \delta^{-1}\,U$
  - Smoothed functional scheme by Katkovnik and Kulchitsky (1972);
  - Refined by Polyak and Tsybakov (1990);
  - Further studied by Dippon (2003); Nesterov (2011).

“Zoo” of Gradient Estimation Methods

- $U \sim \delta\,\mathcal{N}(0, I)$, $V = \delta^{-1}\,U$
  - Smoothed functional scheme by Katkovnik and Kulchitsky (1972);
  - Refined by Polyak and Tsybakov (1990);
  - Further studied by Dippon (2003); Nesterov (2011).
- $U \sim \delta\,\mathrm{Unif}(\mathbb{S}_d)$, $V = d\delta^{-1}\,U$

"Zoo" of Gradient Estimation Methods

- $U \sim \delta\,\mathcal{N}(0, I)$, $V = \delta^{-1}\,U$
  - Smoothed functional scheme by Katkovnik and Kulchitsky (1972);
  - Refined by Polyak and Tsybakov (1990);
  - Further studied by Dippon (2003); Nesterov (2011).
- $U \sim \delta\,\mathrm{Unif}(\mathbb{S}_d)$, $V = d\delta^{-1}\,U$
  - RDSA by Kushner and Clark (1978);

“Zoo” of Gradient Estimation Methods

- $U \sim \delta\,\mathcal{N}(0, I)$, $V = \delta^{-1}\,U$
  - Smoothed functional scheme by Katkovnik and Kulchitsky (1972);
  - Refined by Polyak and Tsybakov (1990);
  - Further studied by Dippon (2003); Nesterov (2011).
- $U \sim \delta\,\mathrm{Unif}(\mathbb{S}_d)$, $V = d\delta^{-1}\,U$
  - RDSA by Kushner and Clark (1978);
  - Rediscovered by Flaxman et al. (2005)

"Zoo" of Gradient Estimation Methods

- $U \sim \delta\,\mathcal{N}(0, I)$, $V = \delta^{-1}\,U$
  - Smoothed functional scheme by Katkovnik and Kulchitsky (1972);
  - Refined by Polyak and Tsybakov (1990);
  - Further studied by Dippon (2003); Nesterov (2011).
- $U \sim \delta\,\mathrm{Unif}(\mathbb{S}_d)$, $V = d\delta^{-1}\,U$
  - RDSA by Kushner and Clark (1978);
  - Rediscovered by Flaxman et al. (2005)
- $U_i \sim \delta\,\mathrm{Rademacher}(\pm 1)$, $V = \delta^{-1}\,U$

"Zoo" of Gradient Estimation Methods

- $U \sim \delta\,\mathcal{N}(0, I)$, $V = \delta^{-1}\,U$
  - Smoothed functional scheme by Katkovnik and Kulchitsky (1972);
  - Refined by Polyak and Tsybakov (1990);
  - Further studied by Dippon (2003); Nesterov (2011).
- $U \sim \delta\,\mathrm{Unif}(\mathbb{S}_d)$, $V = d\delta^{-1}\,U$
  - RDSA by Kushner and Clark (1978);
  - Rediscovered by Flaxman et al. (2005)
- $U_i \sim \delta\,\mathrm{Rademacher}(\pm 1)$, $V = \delta^{-1}\,U$
  - SPSA by Spall (1992).

"Zoo" of Gradient Estimation Methods

- $U \sim \delta\,\mathcal{N}(0, I)$, $V = \delta^{-1}\,U$
  - Smoothed functional scheme by Katkovnik and Kulchitsky (1972);
  - Refined by Polyak and Tsybakov (1990);
  - Further studied by Dippon (2003); Nesterov (2011).
- $U \sim \delta\,\mathrm{Unif}(\mathbb{S}_d)$, $V = d\delta^{-1}\,U$
  - RDSA by Kushner and Clark (1978);
  - Rediscovered by Flaxman et al. (2005)
- $U_i \sim \delta\,\mathrm{Rademacher}(\pm 1)$, $V = \delta^{-1}\,U$
  - SPSA by Spall (1992).
- $\ldots$

“Zoo” of Gradient Estimation Methods

- $U \sim \delta\,\mathcal{N}(0, I)$, $V = \delta^{-1}\,U$
  - Smoothed functional scheme by Katkovnik and Kulchitsky (1972);
  - Refined by Polyak and Tsybakov (1990);
  - Further studied by Dippon (2003); Nesterov (2011).
- $U \sim \delta\,\mathrm{Unif}(\mathbb{S}_d)$, $V = d\delta^{-1}\,U$
  - RDSA by Kushner and Clark (1978);
  - Rediscovered by Flaxman et al. (2005)
- $U_i \sim \delta\,\mathrm{Rademacher}(\pm 1)$, $V = \delta^{-1}\,U$
  - SPSA by Spall (1992).
- $\ldots$

Does it matter which of these we select? Not really:
Bias: $O(\delta^2)$, second moment: $O(1)$ or $O(\delta^{-2})$.

Outline

Gradient Estimation Oracles

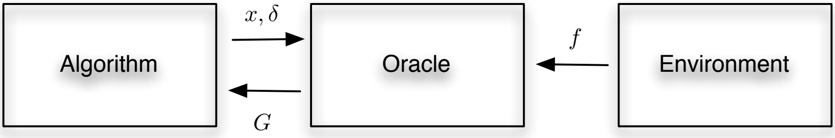

1. Bias: $\|\mathbb{E}[G] - \nabla f(x)\|_* \leq c_1(\delta)$; and

Gradient Estimation Oracles

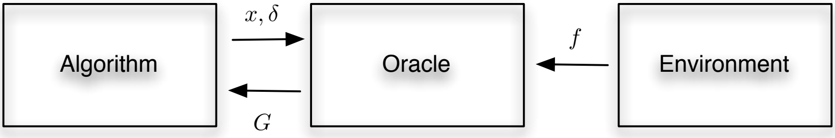

1. Bias: $\left\|\mathbb{E}\left[G\right] - \nabla f(x)\right\|_* \leq c_1(\delta)$; and
2. Second moment: $\mathbb{E}\left[\|G\|_*^2\right] \leq c_2(\delta)$.

Gradient Estimation Oracles

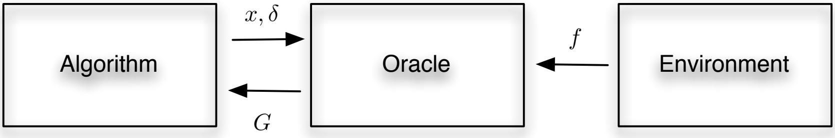

1. Bias: $\left\|\mathbb{E}\left[G\right] - \nabla f(x)\right\|_* \leq c_1(\delta)$; and
2. Second moment: $\mathbb{E}\left[\|G\|_*^2\right] \leq c_2(\delta)$.

Controlled noise: $c_1(\delta) = C_1\delta^2$, $c_2(\delta) = C_2$.

Gradient Estimation Oracles

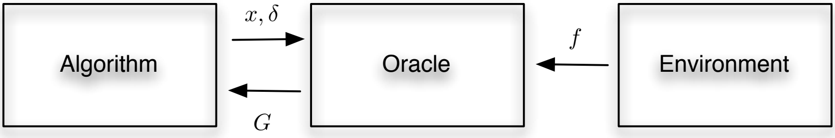

1. Bias: $\|\mathbb{E}[G] - \nabla f(x)\|_* \leq c_1(\delta)$; and
2. Second moment: $\mathbb{E}\left[\|G\|_*^2\right] \leq c_2(\delta)$.

Controlled noise: $c_1(\delta) = C_1\delta^2$, $c_2(\delta) = C_2$.
Uncontrolled noise: $c_1(\delta) = C_1\delta^2$, $c_2(\delta) = C_2\delta^{-2}$.

Gradient Estimation Oracles

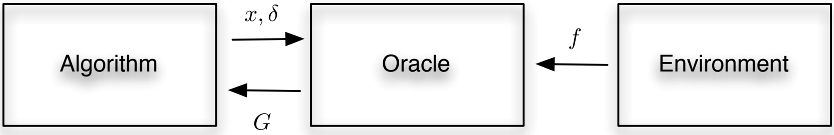

1. Bias: $\|\mathbb{E}[G] - \nabla f(x)\|_* \leq c_1(\delta)$; and
2. Second moment: $\mathbb{E}\left[\|G\|_*^2\right] \leq c_2(\delta)$.

Controlled noise: $c_1(\delta) = C_1\delta^2$, $c_2(\delta) = C_2$.
Uncontrolled noise: $c_1(\delta) = C_1\delta^2$, $c_2(\delta) = C_2\delta^{-2}$.

Polynomial oracle: $c_1(\delta) = C_1\delta^p$, $c_2(\delta) = C_2\delta^{-q}$, $p, q > 0$.

Outline

Upper Bound: Algorithm

Mirror descent (Nemirovskii and Yudin, 1983)

Input: Closed convex set $\mathcal{K}$, regularization function $\mathcal{R} : \mathbb{R}^d \to \mathbb{R}$, tolerance parameter δ, learning rates $\{\eta_t\}_{t=1}^{n-1}$.

Initialize $X_1 \in \mathcal{K}$ arbitrarily.

for $t = 1, 2, \cdots, n-1$ **do**

Query the oracle at X_t.

Receive G_t.

Update

$$X_{t+1} = \operatorname*{argmin}_{x \in \mathcal{K}} \left[\eta_t \langle G_t, x \rangle + D_{\mathcal{R}}(x, X_t)\right] .$$

end for

Return: $\hat{X}_n = \frac{1}{n} \sum_{t=1}^{n} X_t$.

Upper Bound

Theorem (Upper bound)

Consider MD with a (p, q)-order polynomial oracle, α SOC regularizer $\mathcal{R}$. Then:

$$\Delta_n(\mathcal{F}_{L,0}, \mathrm{MD}, c_1, c_2) = O(n^{-\frac{p}{2p+q}})$$
$$\Delta_n(\mathcal{F}_{L,\mu}, \mathrm{MD}, c_1, c_2) = O(n^{-\frac{p}{p+q}})\,.$$

Is This Any Good?

Recall:

$$\Delta_n(\mathcal{F}_{L,0}, \mathrm{MD}, c_1, c_2) = O(n^{-\frac{p}{2p+q}})$$
$$\Delta_n(\mathcal{F}_{L,\mu}, \mathrm{MD}, c_1, c_2) = O(n^{-\frac{p}{p+q}})\,.$$

Is This Any Good?

Recall:
$$\Delta_n(\mathcal{F}_{L,0}, \mathrm{MD}, c_1, c_2) = O(n^{-\frac{p}{2p+q}})$$
$$\Delta_n(\mathcal{F}_{L,\mu}, \mathrm{MD}, c_1, c_2) = O(n^{-\frac{p}{p+q}})\,.$$

Can get an $n^{-1/2}$ rate?

Is This Any Good?

Recall:

$$\Delta_n(\mathcal{F}_{L,0}, \mathrm{MD}, c_1, c_2) = O(n^{-\frac{p}{2p+q}})$$
$$\Delta_n(\mathcal{F}_{L,\mu}, \mathrm{MD}, c_1, c_2) = O(n^{-\frac{p}{p+q}})\,.$$

Can get an $n^{-1/2}$ rate?
Yes, if $p/(2p+q) \geq 1/2$ vs. $p/(p+q) \geq 1/2$.

Is This Any Good?

Recall:

$$\Delta_n(\mathcal{F}_{L,0}, \mathrm{MD}, c_1, c_2) = O(n^{-\frac{p}{2p+q}})$$
$$\Delta_n(\mathcal{F}_{L,\mu}, \mathrm{MD}, c_1, c_2) = O(n^{-\frac{p}{p+q}})\,.$$

Can get an $n^{-1/2}$ rate?
Yes, if $p/(2p+q) \geq 1/2$ vs. $p/(p+q) \geq 1/2$.
First holds iff $q = 0$. Second holds iff $p \geq q$.

Is This Any Good?

Recall:

$$\Delta_n(\mathcal{F}_{L,0}, \mathrm{MD}, c_1, c_2) = O(n^{-\frac{p}{2p+q}})$$
$$\Delta_n(\mathcal{F}_{L,\mu}, \mathrm{MD}, c_1, c_2) = O(n^{-\frac{p}{p+q}})\,.$$

Can get an $n^{-1/2}$ rate?
Yes, if $p/(2p+q) \geq 1/2$ vs. $p/(p+q) \geq 1/2$.
First holds iff $q = 0$. Second holds iff $p \geq q$.

Uncontrolled noise; under smoothness, $p = q = 2$.

Is This Any Good?

Recall:

$$\Delta_n(\mathcal{F}_{L,0}, \mathrm{MD}, c_1, c_2) = O(n^{-\frac{p}{2p+q}})$$
$$\Delta_n(\mathcal{F}_{L,\mu}, \mathrm{MD}, c_1, c_2) = O(n^{-\frac{p}{p+q}})\,.$$

Can get an $n^{-1/2}$ rate?
Yes, if $p/(2p+q) \geq 1/2$ vs. $p/(p+q) \geq 1/2$.
First holds iff $q = 0$. Second holds iff $p \geq q$.

Uncontrolled noise; under smoothness, $p = q = 2$.
For $\mathcal{F}_{L,\mu}$ we get $O(n^{-1/2})$ as Hazan and Levy (2014).

Is This Any Good?

Recall:

$$\Delta_n(\mathcal{F}_{L,0}, \mathrm{MD}, c_1, c_2) = O(n^{-\frac{p}{2p+q}})$$
$$\Delta_n(\mathcal{F}_{L,\mu}, \mathrm{MD}, c_1, c_2) = O(n^{-\frac{p}{p+q}})\,.$$

Can get an $n^{-1/2}$ rate?
Yes, if $p/(2p+q) \geq 1/2$ vs. $p/(p+q) \geq 1/2$.
First holds iff $q = 0$. Second holds iff $p \geq q$.

Uncontrolled noise; under smoothness, $p = q = 2$.
For $\mathcal{F}_{L,\mu}$ we get $O(n^{-1/2})$ as Hazan and Levy (2014).
For $\mathcal{F}_{L,0}$ we get $O(n^{-1/3})$ as Saha and Tewari (2011).

Is This Any Good?

Recall:
$$\Delta_n(\mathcal{F}_{L,0}, \mathrm{MD}, c_1, c_2) = O(n^{-\frac{p}{2p+q}})$$
$$\Delta_n(\mathcal{F}_{L,\mu}, \mathrm{MD}, c_1, c_2) = O(n^{-\frac{p}{p+q}})\,.$$

Can get an $n^{-1/2}$ rate?
Yes, if $p/(2p+q) \geq 1/2$ vs. $p/(p+q) \geq 1/2$.
First holds iff $q = 0$. Second holds iff $p \geq q$.

Uncontrolled noise; under smoothness, $p = q = 2$.
For $\mathcal{F}_{L,\mu}$ we get $O(n^{-1/2})$ as Hazan and Levy (2014).
For $\mathcal{F}_{L,0}$ we get $O(n^{-1/3})$ as Saha and Tewari (2011).

Controlled noise: under smoothness, $p = 2$, $q = 0$.

Is This Any Good?

Recall:

$$\Delta_n(\mathcal{F}_{L,0}, \mathrm{MD}, c_1, c_2) = O(n^{-\frac{p}{2p+q}})$$
$$\Delta_n(\mathcal{F}_{L,\mu}, \mathrm{MD}, c_1, c_2) = O(n^{-\frac{p}{p+q}})\,.$$

Can get an $n^{-1/2}$ rate?

Yes, if $p/(2p+q) \geq 1/2$ vs. $p/(p+q) \geq 1/2$.

First holds iff $q = 0$. Second holds iff $p \geq q$.

Uncontrolled noise; under smoothness, $p = q = 2$.

For $\mathcal{F}_{L,\mu}$ we get $O(n^{-1/2})$ as Hazan and Levy (2014).

For $\mathcal{F}_{L,0}$ we get $O(n^{-1/3})$ as Saha and Tewari (2011).

Controlled noise: under smoothness, $p = 2$, $q = 0$.

For $\mathcal{F}_{L,\mu}$ we get $O(n^{-1})$ as Nesterov (2011) .

Is This Any Good?

Recall:

$$\Delta_n(\mathcal{F}_{L,0}, \mathrm{MD}, c_1, c_2) = O(n^{-\frac{p}{2p+q}})$$
$$\Delta_n(\mathcal{F}_{L,\mu}, \mathrm{MD}, c_1, c_2) = O(n^{-\frac{p}{p+q}})\,.$$

Can get an $n^{-1/2}$ rate?
Yes, if $p/(2p+q) \geq 1/2$ vs. $p/(p+q) \geq 1/2$.
First holds iff $q = 0$. Second holds iff $p \geq q$.

Uncontrolled noise; under smoothness, $p = q = 2$.
For $\mathcal{F}_{L,\mu}$ we get $O(n^{-1/2})$ as Hazan and Levy (2014).
For $\mathcal{F}_{L,0}$ we get $O(n^{-1/3})$ as Saha and Tewari (2011).

Controlled noise: under smoothness, $p = 2$, $q = 0$.
For $\mathcal{F}_{L,\mu}$ we get $O(n^{-1})$ as Nesterov (2011) .
For $\mathcal{F}_{L,0}$ we get $O(n^{-1/2})$ as Duchi et al. (2015).

Could We Have Done Better?

Theorem (**Lower bound**)

$\mathcal{K} \subset \mathbb{R}^d$ convex, closed, with $\{+1, -1\}^d \subset \mathcal{K}$, n large enough. For any algorithm A that observes n random elements from a (p, q) polynomial oracle, we have

$$\Delta_n(\mathcal{F}_{L,0}, \mathrm{A}, c_1, c_2) = \Omega(n^{-\frac{p}{2p+q}}),$$

$$\Delta_n(\mathcal{F}_{L,1}, \mathrm{A}, c_1, c_2) = \Omega(n^{-\frac{2p}{2p+q}}).$$

Could We Have Done Better?

Theorem (Lower bound)

$\mathcal{K} \subset \mathbb{R}^d$ convex, closed, with $\{+1,-1\}^d \subset \mathcal{K}$, n large enough. For any algorithm A that observes n random elements from a (p,q) polynomial oracle, we have

$$\Delta_n(\mathcal{F}_{L,0}, \mathrm{A}, c_1, c_2) = \Omega(n^{-\frac{p}{2p+q}}),$$
$$\Delta_n(\mathcal{F}_{L,1}, \mathrm{A}, c_1, c_2) = \Omega(n^{-\frac{2p}{2p+q}}).$$

Compare with

$$\Delta_n(\mathcal{F}_{L,0}, \mathrm{MD}, c_1, c_2) = O(n^{-\frac{p}{2p+q}})$$
$$\Delta_n(\mathcal{F}_{L,1}, \mathrm{MD}, c_1, c_2) = O(n^{-\frac{p}{p+q}}) = O(n^{-\frac{2p}{2p+2q}}).$$

(The lower bound for $\mathcal{F}_{L,0}$ is tight, for $\mathcal{F}_{L,1}$ it is weak.)

Lower Bound Idea

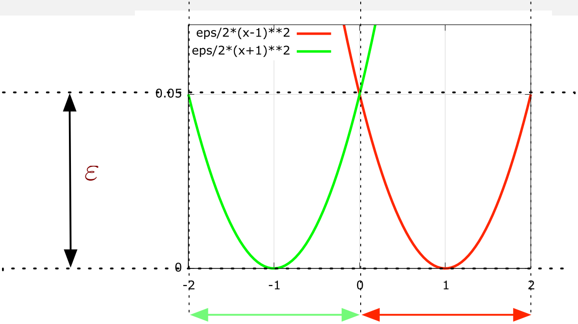

$$\max_{\delta} \frac{c_2(\delta)}{(2\varepsilon - 2c_1(\delta))^2} < n$$

$$\Delta = 2\varepsilon - 2c_1(\delta)$$

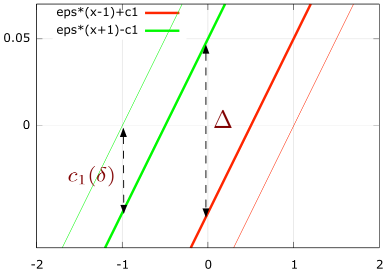

Outline

Conclusion

Corollary

To get the optimal $O(n^{-1/2})$ rate for $\mathcal{F}_{L,0}$ with uncontrolled noise, with a low-complexity algorithm, one of the following must be done:

1. *An oracle with $q = 0$ (constant second moment bound) must be designed.*
2. ~~*An algorithm that makes better use of the gradient estimates must be designed.*~~
3. *Some extra properties of gradient estimates must be exploited beyond bias/variance.*
4. *Design a non-gradient algorithm.*

Big Picture

- Noisy Bandit Convex Optimization

Big Picture

- Noisy Bandit Convex Optimization
- Controlled noise: Pretty well understood

Big Picture

- Noisy Bandit Convex Optimization
- Controlled noise: Pretty well understood
- Linear case: Pretty well understood

Big Picture

- Noisy Bandit Convex Optimization
- Controlled noise: Pretty well understood
- Linear case: Pretty well understood
- Uncontrolled noise: Not much is known about low complexity, optimal algorithms!

Big Picture

- Noisy Bandit Convex Optimization
- Controlled noise: Pretty well understood
- Linear case: Pretty well understood
- Uncontrolled noise: Not much is known about low complexity, optimal algorithms!
- Fascinating results on regret minimization in the online setting

Thanks! Questions?

References I

Agarwal, A., Foster, D. P., Hsu, D., Kakade, S. M., and Rakhlin, A. (2013). Stochastic convex optimization with bandit feedback. *SIAM Journal on Optimization*, 23(1):213–240.

Dippon, J. (2003). Accelerated randomized stochastic optimization. *The Annals of Statistics*, 31(4):1260–1281.

Duchi, J. C., Jordan, M., Wainwright, M. J., and Wibisono, A. (2015). Optimal rates for zero-order convex optimization: The power of two function evaluations. *IEEE Transactions on Information Theory*, 61(5):2788–2806.

Flaxman, A. D., Kalai, A. T., and McMahan, H. B. (2005). Online convex optimization in the bandit setting: gradient descent without a gradient. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 385–394. Society for Industrial and Applied Mathematics.

Hazan, E. and Levy, K. (2014). Bandit convex optimization: Towards tight bounds. In *NIPS*, pages 784–792.

Katkovnik, V. Y. and Kulchitsky, Y. (1972). Convergence of a class of random search algorithms. *Automation Remote Control*, 8:1321–1326.

Kiefer, J. and Wolfowitz, J. (1952). Stochastic estimation of the maximum of a regression function. *The Annals of Mathematical Statistics*, 23(3):462–466.

Kushner, H. J. and Clark, D. S. (1978). *Stochastic Approximation Methods for Constrained and Unconstrained Systems*. Springer Verlag, New York.

References II

Liang, T., Narayanan, H., and Rakhlin, A. (2014). On zeroth-order stochastic convex optimization via random walks. *arXiv preprint arXiv:1402.2667*.

Nemirovskii, A. and Yudin, D. B. (1983). *Problem complexity and method efficiency in optimization*. Wiley-Interscience series in discrete mathematics. Wiley.

Nesterov, Y. (2004). *Introductory lectures on convex optimization*, volume 87. Springer Science & Business Media.

Nesterov, Y. (2011). Random gradient-free minimization of convex functions. CORE Discussion Papers 2011001, Université catholique de Louvain, Center for Operations Research and Econometrics (CORE).

Polyak, B. and Tsybakov, A. (1990). Optimal orders of accuracy for search algorithms of stochastic optimization. *Problems in Information Transmission*, pages 126–133.

Saha, A. and Tewari, A. (2011). Improved regret guarantees for online smooth convex optimization with bandit feedback. In *International Conference on Artificial Intelligence and Statistics*, pages 636–642.

Shamir, O. (2012). On the complexity of bandit and derivative-free stochastic convex optimization. In *Proceedings of the Twenty Sixth Annual Conference on Computational Learning Theory*.

Shor, N. Z. (1970). Convergence rate of the gradient descent method with dilatation of the space. *Kibernetika*, 6(2):102–108.

Spall, J. C. (1992). Multivariate stochastic approximation using a simultaneous perturbation gradient approximation. *Automatic Control, IEEE Transactions on*, 37(3):332–341.