

CS-GY 6763: Lecture 9

Online and Stochastic Gradient Descent, Dimension Dependent Optimization

NYU Tandon School of Engineering, Prof. Christopher Musco

Given a function f to minimize, assume we have:

- **Function oracle:** Evaluate $f(\mathbf{x})$ for any \mathbf{x} .
- **Gradient oracle:** Evaluate $\nabla f(\mathbf{x})$ for any \mathbf{x} .

Goal: Minimize the number of oracle calls to find $\tilde{\mathbf{x}}$ such that $f(\tilde{\mathbf{x}}) \leq \min_{\mathbf{x}} f(\mathbf{x}) + \epsilon$.

In machine learning applications, $f(\mathbf{x})$ is typically a loss function for a fixed training dataset.

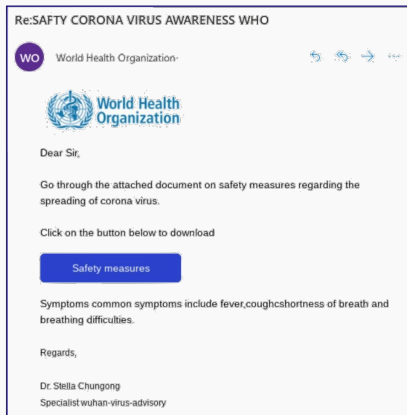
First part of class:

- Basics of an alternative setting: Online Learning + Optimization.
- Introduction to Regret Analysis.
- Application to analyzing Stochastic Gradient Descent.

Many machine learning problems are solved in an online setting with constantly changing data.

- Spam filters are incrementally updated and adapt as they see more examples of spam over time.
- Text recommendation engines (e.g. Github Copilot) need to be kept up-to-date as software libraries/APIs change.
- Content recommendation systems adapt to user behavior and clicks (which may not be a good thing...)


Machine learning based email spam filtering.



Markers for spam change overtime, so model might change.

Plant identification via iNaturalist app.

(California Academy of Science + National Geographic)



- When the app fails, image is classified via crowdsourcing (backed by huge network of amateurs and experts).
- Single model that is updated constantly, not retrained in batches.

ONLINE LEARNING FRAMEWORK

Choose some model $M_{\mathbf{x}}$ parameterized by parameters \mathbf{x} and some loss function ℓ . At time steps $1, \dots, T$, receive data vectors $\underline{\mathbf{a}}^{(1)}, \dots, \underline{\mathbf{a}}^{(T)}$.

- At each time step, we pick (“play”) a parameter vector $\underline{\mathbf{x}}^{(i)}$.
- Make prediction $\tilde{y}^{(i)} = M_{\underline{\mathbf{x}}^{(i)}}(\underline{\mathbf{a}}^{(i)})$.
- Then told true value or label $\underline{y}^{(i)}$. Possibly use this information to choose a new $\underline{\mathbf{x}}^{(i+1)}$.
- Goal is to minimize cumulative loss:

$$\left(L = \sum_{i=1}^T \ell(\underline{\mathbf{x}}^{(i)}, \underline{\mathbf{a}}^{(i)}, \underline{y}^{(i)}) \right) \quad \nearrow \quad | M_{\underline{\mathbf{x}}^{(i)}}(\underline{\mathbf{a}}^{(i)}) - \underline{y}^{(i)} |^2$$

For example, for a regression problem we might use the ℓ_2 loss:

$$\ell(\underline{\mathbf{x}}^{(i)}, \underline{\mathbf{a}}^{(i)}, \underline{y}^{(i)}) = \left| \langle \underline{\mathbf{x}}^{(i)}, \underline{\mathbf{a}}^{(i)} \rangle - \underline{y}^{(i)} \right|^2.$$

For classification, we could use logistic/cross-entropy loss.

Abstraction as optimization problem: Instead of a single objective function f , we have a single (initially unknown) function $f_1, \dots, f_T : \mathbb{R}^d \rightarrow \mathbb{R}$ for each time step.

- For time step $i \in 1, \dots, T$, select vector $\mathbf{x}^{(i)}$
- Observe f_i and pay cost $f_i(\mathbf{x}^{(i)})$
- Goal is to minimize $\sum_{i=1}^T f_i(\mathbf{x}^{(i)})$.

We make no assumptions that f_1, \dots, f_T are related to each other at all!

REGRET BOUND

In offline optimization, we wanted to find $\hat{\mathbf{x}}$ satisfying $f(\hat{\mathbf{x}}) \leq \min_{\mathbf{x}} f(\mathbf{x})$. Ask for a similar thing here.

Objective: Choose $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(T)}$ so that:

$$\frac{1}{T} \sum_{i=1}^T f_i(\mathbf{x}^{(i)}) \leq \frac{1}{T} \left[\min_{\mathbf{x}} \sum_{i=1}^T f_i(\mathbf{x}) \right] + \frac{1}{T} \epsilon.$$

Here ϵ is called the **regret** of our solution sequence $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(T)}$. Regret compares to the best fixed solution in hindsight.

We typically ϵ to be growing sublinearly in T .

Regret compares to the best fixed solution in hindsight.

$$\sum_{i=1}^T f_i(\mathbf{x}^{(i)}) \leq \left[\min_{\mathbf{x}} \sum_{i=1}^T f_i(\mathbf{x}) \right] + \epsilon.$$

It's very possible that $\left(\sum_{i=1}^T f_i(\mathbf{x}^{(i)}) \right) < \left[\min_{\mathbf{x}} \sum_{i=1}^T f_i(\mathbf{x}) \right]$. Could we hope for something stronger?

Exercise: Argue that the following is impossible to achieve:

$$\sum_{i=1}^T f_i(\mathbf{x}^{(i)}) \leq \left[\sum_{i=1}^T \min_{\mathbf{x}} f_i(\mathbf{x}) \right] + \epsilon.$$

HARD EXAMPLE FOR ONLINE OPTIMIZATION

Convex functions:

$$f_1(x) = |x - h_1|$$

\vdots

$$f_T(x) = |x - h_T|$$

$$\min_x \sum_{i=1}^T f_i(x) \approx T/2$$

where h_1, \dots, h_T are i.i.d. uniform $\{0, 1\}$.

$$\sum_{i=1}^T \min_x f_i(x) = \Theta$$

Online can't do better than

$$\sum_{i=1}^T f_i(x^{(i)}) \approx \underline{T/2}.$$

$$\sum_{i=1}^T f_i(\mathbf{x}^{(i)}) \leq \left[\min_{\mathbf{x}} \sum_{i=1}^T f_i(\mathbf{x}) \right] + \epsilon.$$

Beautiful balance:

- Either f_1, \dots, f_T are similar or changing slowly, so we can learn predict f_i from earlier functions.
- Or f_1, \dots, f_T are very different, in which case $\min_{\mathbf{x}} \sum_{i=1}^T f_i(\mathbf{x})$ is large, so regret bound is easy to achieve.
- Or we live somewhere in the middle.

Follow-the-leader algorithm: f_1, \dots, f_{i-1}

- Choose $\mathbf{x}^{(0)}$.
- For $i = 1, \dots, T$:
 - Let $\underline{\mathbf{x}}^{(i)} = \arg \min_{\mathbf{x}} \sum_{j=1}^{i-1} f_j(\mathbf{x})$.
 - Play $\mathbf{x}^{(i)}$.
 - Observe f_i and incur cost $f_i(\mathbf{x}^{(i)})$.

Simple and intuitive, but there are two issues with this approach. One is computational, one is related to the accuracy.


$$O(T^2)$$

$$f_i(\mathbf{x}) = \ell(\mathbf{M}_{\mathbf{x}}(\mathbf{a}^{(i)}) - \mathbf{y}_i)$$

FOLLOW-THE-LEADER

Hard case:

f_a f_b f_b $f_a f_a$ $f_b f_b$



Online Gradient descent:

- Choose $\mathbf{x}^{(1)}$ and $\eta = \frac{R}{G\sqrt{T}}$.
- For $i = 1, \dots, T$:
 - Play $\mathbf{x}^{(i)}$
 - Observe f_i and incur cost $f_i(\mathbf{x}^{(i)})$.
 - $\mathbf{x}^{(i+1)}$ = $\mathbf{x}^{(i)} - \eta \nabla f_i(\mathbf{x}^{(i)})$

If $f_1, \dots, f_T = f$ are all the same, this is the same as regular gradient descent. We update parameters using the gradient ∇f at each step.

ONLINE GRADIENT DESCENT (OGD)

$\mathbf{x}^* = \arg \min_{\mathbf{x}} \sum_{i=1}^T f_i(\mathbf{x})$ (the offline optimum)

Assume:

- f_1, \dots, f_T are all convex.
- Each is G -Lipschitz: for all \mathbf{x}, i , $\|\nabla f_i(\mathbf{x})\|_2 \leq G$.
- Starting radius: $\|\mathbf{x}^* - \mathbf{x}^{(1)}\|_2 \leq R$.

Online Gradient descent:

- Choose $\mathbf{x}^{(1)}$ and $\eta = \frac{R}{G\sqrt{T}}$.
- For $i = 1, \dots, T$:
 - Play $\mathbf{x}^{(i)}$.
 - Observe f_i and incur cost $f_i(\mathbf{x}^{(i)})$.
 - $\mathbf{x}^{(i+1)} = \mathbf{x}^{(i)} - \eta \nabla f_i(\mathbf{x}^{(i)})$

ONLINE GRADIENT DESCENT ANALYSIS

Let $\mathbf{x}^* = \arg \min_{\mathbf{x}} \sum_{i=1}^T f_i(\mathbf{x})$ (the offline optimum)

Theorem (OGD Regret Bound)

for convex f_1, \dots, f_T

$$\text{After } T \text{ steps, } \epsilon = \frac{1}{T} \left[\sum_{i=1}^T f_i(\mathbf{x}^{(i)}) \right] - \frac{1}{T} \left[\sum_{i=1}^T f_i(\mathbf{x}^*) \right] \leq \frac{RG\sqrt{T}}{T}$$

Average regret overtime is bounded by $\frac{\epsilon}{T} \leq \frac{RG}{\sqrt{T}}$

Goes $\rightarrow 0$ as $T \rightarrow \infty$.

$$T \geq \frac{R^2 G^2}{\delta^2}$$

All this with no assumptions on how f_1, \dots, f_T relate to each other! They could have even been chosen **adversarially** – e.g. with f_i depending on our choice of \mathbf{x}_i and all previous choices.

Theorem (OGD Regret Bound)

After T steps, $\epsilon = \left[\sum_{i=1}^T f_i(\mathbf{x}^{(i)}) \right] - \left[\sum_{i=1}^T f_i(\mathbf{x}^*) \right] \leq RG\sqrt{T}$.

Claim 1: For all $i = 1, \dots, T$,

$$f_i(\mathbf{x}^{(i)}) - f_i(\mathbf{x}^*) \leq \frac{\|\mathbf{x}^{(i)} - \mathbf{x}^*\|_2^2 - \|\mathbf{x}^{(i+1)} - \mathbf{x}^*\|_2^2}{2\eta} + \frac{\eta G^2}{2}$$

$\nearrow \mathbf{x}^{(i)} \approx \nabla f_i(\mathbf{x}^{(i)})$

(Same proof for standard GD. Only uses convexity of f_i .)

Theorem (OGD Regret Bound)

After T steps, $\epsilon = \left[\sum_{i=1}^T f_i(\mathbf{x}^{(i)}) \right] - \left[\sum_{i=1}^T f_i(\mathbf{x}^*) \right] \leq RG\sqrt{T}$.

Claim 1: For all $i = 1, \dots, T$,

$$\underbrace{f_i(\mathbf{x}^{(i)}) - f_i(\mathbf{x}^*)} \leq \frac{\|\mathbf{x}^{(i)} - \mathbf{x}^*\|_2^2 - \|\mathbf{x}^{(i+1)} - \mathbf{x}^*\|_2^2}{2\eta} + \frac{\eta G^2}{2}$$

Telescoping Sum:

$$\begin{aligned} \sum_{i=1}^T \left[\underbrace{f_i(\mathbf{x}^{(i)})} - \underbrace{f_i(\mathbf{x}^*)} \right] &\leq \frac{\|\mathbf{x}^{(1)} - \mathbf{x}^*\|_2^2 - \|\mathbf{x}^{(T)} - \mathbf{x}^*\|_2^2}{2\eta} + \frac{T\eta G^2}{2} \\ &\leq \frac{R^2}{2\eta} + \frac{T\eta G^2}{2} = \frac{R G \sqrt{T}}{2} + \frac{R G \sqrt{T}}{2} = \underline{R G \sqrt{T}} \end{aligned}$$

$\rightarrow \leq R^2$ $\rightarrow \geq 0$

$$M = \frac{R}{G\sqrt{T}}$$

STOCHASTIC GRADIENT DESCENT (SGD)

Efficient offline optimization method for functions f with finite sum structure:

$$f(\mathbf{x}) = \sum_{i=1}^n f_i(\mathbf{x}).$$

(Handwritten diagram: The equation is enclosed in large parentheses. The function $f(\mathbf{x})$ on the left and the summand $f_i(\mathbf{x})$ on the right are circled. A handwritten arrow points from the text "# of training items" to the upper limit n of the summation.)

Goal is to find $\hat{\mathbf{x}}$ such that $f(\hat{\mathbf{x}}) \leq f(\mathbf{x}^*) + \epsilon$.

- The most widely use optimization algorithm in modern machine learning.
- Easily analyzed as a special case of online gradient descent!


STOCHASTIC GRADIENT DESCENT

Recall the machine learning setup. In empirical risk minimization, we can typically write:


$$f(\mathbf{x}) = \sum_{i=1}^n f_i(\mathbf{x})$$

where f_i is the loss function for a particular data example $(\mathbf{a}^{(i)}, y^{(i)})$.

Example: least squares linear regression.

$$f(\mathbf{x}) = \sum_{i=1}^n (\mathbf{x}^T \mathbf{a}^{(i)} - y^{(i)})^2$$


Note that by linearity, $\nabla f(\mathbf{x}) = \sum_{i=1}^n \nabla f_i(\mathbf{x})$.



STOCHASTIC GRADIENT DESCENT

Main idea: Use random approximate gradient in place of actual gradient.

Pick random $j \in 1, \dots, n$ and update x using $\nabla f_j(x)$.

$$\mathbb{E}[\nabla f_j(x)] = \frac{1}{n} \nabla f(x).$$
$$\mathbb{E}[\nabla f_j(x)] = \sum_{j=1}^n \frac{1}{n} \nabla f_j(x) = \frac{1}{n} \underbrace{\sum_{j=1}^n \nabla f_j(x)}_{\nabla f(x)} = \frac{1}{n} \nabla f(x)$$

$\nabla f_j(x)$ is an unbiased estimate for the true gradient $\nabla f(x)$, but can typically be computed in a $1/n$ fraction of the time!

Trade slower convergence for cheaper iterations.




Stochastic first-order oracle for $f(\mathbf{x}) = \sum_{i=1}^n f_i(\mathbf{x})$.

- **Function Query:** For any chosen j, \mathbf{x} , return $f_j(\mathbf{x})$
- **Gradient Query:** For any chosen j, \mathbf{x} , return $\nabla f_j(\mathbf{x})$

Stochastic Gradient descent:

- Choose starting vector $\underline{\mathbf{x}}^{(1)}$, step size η
- For $i = 1, \dots, T$:
 - Pick random \underline{j}_i uniformly at random from $1, \dots, n$.
 - $\underline{\mathbf{x}}^{(i+1)} = \underline{\mathbf{x}}^{(i)} - \eta \nabla f_{\underline{j}_i}(\underline{\mathbf{x}}^{(i)})$
- Return $\underline{\hat{\mathbf{x}}} = \frac{1}{T} \sum_{i=1}^T \underline{\mathbf{x}}^{(i)}$

VISUALIZING SGD



STOCHASTIC GRADIENT DESCENT

Assume: $\|\nabla f(\mathbf{x})\|_2 \leq \sum_{i=1}^n \|\nabla f_i(\mathbf{x})\|_2 \leq G'$

- Finite sum structure: $f(\mathbf{x}) = \sum_{i=1}^n f_i(\mathbf{x})$, with f_1, \dots, f_n all convex.
- Lipschitz functions: for all \mathbf{x}, j , $\|\nabla f_j(\mathbf{x})\|_2 \leq \frac{G'}{n}$.
 - What does this imply about Lipschitz constant of f ?
- Starting radius: $\|\mathbf{x}^* - \mathbf{x}^{(1)}\|_2 \leq R$.

Stochastic Gradient descent:

- Choose $\mathbf{x}^{(1)}$, steps T , step size $\eta = \frac{R}{G'\sqrt{T}}$.
- For $i = 1, \dots, T$:
 - Pick random $j_i \in 1, \dots, n$.
 - $\mathbf{x}^{(i+1)} = \mathbf{x}^{(i)} - \eta \nabla f_{j_i}(\mathbf{x}^{(i)})$
- Return $\hat{\mathbf{x}} = \frac{1}{T} \sum_{i=1}^T \mathbf{x}^{(i)}$

Approach: View as online gradient descent run on function sequence f_{j_1}, \dots, f_{j_T} .

Only use the fact that step equals gradient in expectation.

JENSEN'S INEQUALITY

For a convex function f and points $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(t)}$

$$f\left(\frac{1}{t} \cdot \mathbf{x}^{(1)} + \dots + \frac{1}{t} \cdot \mathbf{x}^{(t)}\right) \leq \frac{1}{t} \cdot f(\mathbf{x}^{(1)}) + \dots + \frac{1}{t} \cdot f(\mathbf{x}^{(t)})$$

$$f(\lambda x + (1-\lambda)y) \leq \lambda f(x) + (1-\lambda)f(y)$$



STOCHASTIC GRADIENT DESCENT ANALYSIS

Claim (SGD Convergence)

After $T = \frac{R^2 G^2}{\epsilon^2}$ iterations:

$$\hat{x} = \frac{1}{T} \sum_{i=1}^T x^{(i)}$$

$$\mathbb{E} [f(\hat{x}) - f(x^*)] \leq \epsilon.$$

Claim 1:

$$f(\hat{x}) - f(x^*) \leq \frac{1}{T} \sum_{i=1}^T [f(x^{(i)}) - f(x^*)]$$

Prove using Jensen's Inequality:

$$f\left(\frac{1}{T} \sum_{i=1}^T x^{(i)}\right) - \frac{1}{T} \sum_{i=1}^T f(x^*) \leq \frac{1}{T} \sum_{i=1}^T f(x^{(i)}) - \frac{1}{T} \sum_{i=1}^T f(x^*)$$

STOCHASTIC GRADIENT DESCENT ANALYSIS

Claim (SGD Convergence)

$$n \underbrace{\mathbb{E} f_{j_i}(\gamma)} = f(\gamma)$$

After $T = \frac{R^2 G'^2}{\epsilon^2}$ iterations:

$$\mathbb{E}[f(\hat{\mathbf{x}}) - f(\mathbf{x}^*)] \leq \epsilon.$$

$$\begin{aligned} \mathbb{E}[f(\hat{\mathbf{x}}) - f(\mathbf{x}^*)] &\leq \frac{1}{T} \sum_{i=1}^T \mathbb{E}[f(\mathbf{x}^{(i)}) - f(\mathbf{x}^*)] = \mathbb{E}[f(\mathbf{x}^{(i)})] - f(\mathbf{x}^*) \\ &= \frac{1}{T} \sum_{i=1}^T n \mathbb{E}[f_{j_i}(\mathbf{x}^{(i)}) - f_{j_i}(\mathbf{x}^*)] \\ &= \frac{n}{T} \cdot \mathbb{E} \left[\sum_{i=1}^T f_{j_i}(\mathbf{x}^{(i)}) - f_{j_i}(\mathbf{x}^*) \right] \end{aligned}$$

Claim (SGD Convergence)

$$\mathbf{x}^* = \arg \min_{\mathbf{x}} \underline{\underline{f(\mathbf{x})}}$$

After $T = \frac{R^2 G'^2}{\epsilon^2}$ iterations:

$$\mathbb{E}[f(\hat{\mathbf{x}}) - f(\mathbf{x}^*)] \leq \epsilon.$$

$$\begin{aligned} \mathbb{E}[f(\hat{\mathbf{x}}) - f(\mathbf{x}^*)] &\leq \frac{1}{T} \sum_{i=1}^T \mathbb{E}[f(\mathbf{x}^{(i)}) - f(\mathbf{x}^*)] \\ &= \frac{n}{T} \sum_{i=1}^T \mathbb{E}[f_{j_i}(\mathbf{x}^{(i)}) - f_{j_i}(\mathbf{x}^*)] \\ &\leq \frac{n}{T} \cdot \mathbb{E} \left[\sum_{i=1}^T f_{j_i}(\mathbf{x}^{(i)}) - f_{j_i}(\underline{\underline{\mathbf{x}^{offline}}}) \right], \end{aligned}$$

where $\underline{\underline{\mathbf{x}^{offline}}} = \arg \min_{\mathbf{x}} \left(\sum_{i=1}^T f_{j_i}(\mathbf{x}) \right)$.

Claim (SGD Convergence)

BG \sqrt{T} After $T = \frac{R^2 G'^2}{\epsilon^2}$ iterations:

$$\mathbb{E}[f(\hat{\mathbf{x}}) - f(\mathbf{x}^*)] \leq \epsilon.$$

$$\begin{aligned} \mathbb{E}[f(\hat{\mathbf{x}}) - f(\mathbf{x}^*)] &\leq \frac{1}{T} \sum_{i=1}^T \mathbb{E}[f(\mathbf{x}^{(i)}) - f(\mathbf{x}^*)] \\ &= \frac{1}{T} \sum_{i=1}^T n \mathbb{E}[f_j(\mathbf{x}^{(i)}) - f_j(\mathbf{x}^*)] \\ &\leq \frac{n}{T} \mathbb{E}\left[\sum_{i=1}^T \underline{f_j(\mathbf{x}^{(i)})} - \underline{f_j(\mathbf{x}^{offline})}\right] \\ &\leq \frac{n}{T} \cdot \left(R \cdot \frac{G'}{n} \cdot \sqrt{T}\right) \quad (\text{by OGD guarantee.}) \\ &= \text{BG}' / \sqrt{T} \end{aligned}$$

STOCHASTIC VS. FULL BATCH GRADIENT DESCENT

Number of iterations for error ϵ :

- Gradient Descent: $T = \frac{R^2 G^2}{\epsilon^2}$.
- Stochastic Gradient Descent: $T = \frac{R^2 G'^2}{\epsilon^2}$.

Always have $G \leq G'$: $\max_x \|\nabla f_1(x) + \nabla f_2(x) + \dots + \nabla f_n(x)\|_2$

$$\begin{aligned} G &= \max_x \|\nabla f(x)\|_2 \leq \max_x (\|\nabla f_1(x)\|_2 + \dots + \|\nabla f_n(x)\|_2) \\ &\leq \max_x (\|\nabla f_1(x)\|_2) + \dots + \max_x (\|\nabla f_n(x)\|_2) \\ &\leq n \cdot \frac{G'}{n} = G' \end{aligned}$$

So GD converges strictly faster than SGD.

But for a fair comparison:

- SGD cost = (# of iterations) $\cdot O(1)$
- GD cost = (# of iterations) $\cdot O(n)$

STOCHASTIC VS. FULL BATCH GRADIENT DESCENT

We always have $\underline{G} \leq \underline{G}'$. When it is much smaller then GD will perform better. When it is closer to this upper bound, SGD will perform better.

What is an extreme case where $G = G'$?

$$f_1 = f_2 = \dots = f_n$$

STOCHASTIC VS. FULL BATCH GRADIENT DESCENT

What if each gradient $\nabla f_i(\mathbf{x})$ looks like random vectors in \mathbb{R}^d ?
 E.g. with $\mathcal{N}(0, 1)$ entries?

$\rightarrow d$ dimensions

$$\mathbb{E} [\|\nabla f_i(\mathbf{x})\|_2^2] = d \quad \frac{G^1}{n} \approx \sqrt{d} \quad \underline{G^1 \approx \sqrt{d} \cdot \eta}$$



$$\mathbb{E} [\|\nabla f(\mathbf{x})\|_2^2] = \mathbb{E} \left[\left\| \sum_{i=1}^n \nabla f_i(\mathbf{x}) \right\|_2^2 \right] = dn$$

$$\underbrace{\int \int + \int \int + \int \int + \dots + \int \int}_{\nabla f(\mathbf{x})} = \int \int \xrightarrow{\mathcal{N}(0, \eta)}$$

$$\underline{G \approx \sqrt{dn}}$$

STOCHASTIC VS. FULL BATCH GRADIENT DESCENT


Takeaway: SGD performs better when there is more structure or repetition in the data set.



$$\nabla \ell(\underline{a}^{(j_i)}, \underline{x}^{(i)}, y^{(j_i)}) = \nabla \underline{f}_{j_i}(x^{(i)})$$

$$\left(M_{x^{(i)}} (a^{(i)} - y^{(i)}) \right)^2 \quad \begin{array}{cccc} \underline{j_1} & \underline{j_2} & \dots & \underline{j_T} \\ 4 & 7 & 21 & N \end{array}$$

PRECONDITIONING



$j_i \sim \text{vars } (1, \dots, 4)$
 $(\langle a_{j_i}, x \rangle - y_{j_i})^2$
 $= f_{j_i}(x)$

PRECONDITIONING

$$\begin{aligned} \text{Minimize } f(x) &\leq \text{Minimize } g(x) && \text{Let } x_g^* = \text{argmin } g(x) \\ \text{Minimize } f(x) &\leq f(h(x_g^*)) && = g(x_g^*) \end{aligned}$$

Main idea: Instead of minimizing $f(x)$, find another function $g(x)$ with the same minimum but which is better suited for first order optimization (e.g., has a smaller conditioner number).

Claim: Let $h(x) : \mathbb{R}^d \rightarrow \mathbb{R}^d$ be an invertible function. Let $g(x) = f(h(x))$. Then

$$\min_x f(x) = \min_x g(x) \quad \text{and} \quad \text{argmin}_x f(x) = h\left(\text{argmin}_x g(x)\right).$$

$$\text{Minimize } g(x) \leq \text{Minimize } f(x)$$

$$\begin{aligned} \text{Let } x_f^* &= \text{argmin } f(x) \\ \text{Minimize } g(x) &\leq g(h^{-1}(x_f^*)) = \end{aligned}$$

$$\begin{aligned} &= f(h(h^{-1}(x_f^*))) \\ &= f(x_f^*) \\ &= \text{Minimize } f(x) \end{aligned}$$

PRECONDITIONING

First Requirement: We need $g(\mathbf{x})$ to still be convex.

$\int \rightarrow h(\mathbf{x})$

Claim: Let \mathbf{P} be an invertible $d \times d$ matrix and let $g(\mathbf{x}) = f(\mathbf{P}\mathbf{x})$.

$g(\mathbf{x})$ is always convex.



For this choice of preconditioner, if $\tilde{\mathbf{x}} \approx \arg \min g(\mathbf{x})$, we would want to return $\mathbf{P}\tilde{\mathbf{x}}$ as a near minimizer of f .

dad

Additional Goals:

- $g(x)$ should be better conditioned (more smooth, more strongly convex, etc.) than f .
- $\{P\}$ needs to be easy to compute and we need to be able to apply Px efficiently.

Common choice: Diagonal preconditioner.

- Choose P to be a diagonal matrix D .
 (Handwritten note: $\rightarrow \|PAx - b\|_2$)
- For $f(x) = \|Ax - b\|_2^2$ common choice is $D = \text{diag}(A^T A)^{-1}$, which is known as the **Jacobi preconditioner**.
 (Handwritten note: $\rightarrow \|PAx - b\|_2$)

Often works very well in practice!

DIAGONAL PRECONDITIONER

A =

-734	1	33	9111	0
-31	-2	108	5946	-19
232	-1	101	3502	10
426	0	-65	12503	9
-373	0	26	9298	0
-236	-2	-94	2398	-1
2024	0	-132	-6904	-25
-2258	-1	92	-6516	6
2229	0	0	11921	-22
338	1	-5	-16118	-23

```
>> cond(A'*A)
```

ans =

8.4145e+07

```
>> P = sqrt(inv(diag(diag(A'*A))));
```

```
>> cond(P*A'*A*P)
```

ans =

10.3878

ADAPTIVE STEPSIZES

Another view: If $\underline{g(x)} = \underline{f(Px)}$ then $\underline{\nabla g(x)} = \underline{P^T \nabla f(Px)}$.

$\nabla g(x) = P \nabla f(Px)$ when P is symmetric.

Gradient descent on g :

- For $t = 1, \dots, T$,

$$\underline{P} \underline{x^{(t+1)}} = \underline{P} \underline{x^{(t)}} - \eta \underline{P}^2 [\underline{\nabla f(Px^{(t)})}]$$

$$\tilde{x} \quad P \tilde{x}$$

↪ e.g. when P is diagonal

Gradient descent on g :

- For $t = 1, \dots, T$,

$$\cdot \underline{y^{(t+1)}} = \underline{y^{(t)}} - \eta \underline{P}^2 [\underline{\nabla f(y^{(t)})}]$$

$$P x^{(t+1)} = y^{(t+1)}$$


When P is diagonal, this is just gradient descent with a different step size for each parameter!

ADAPTIVE STEPSIZES

Algorithms based on this idea:

- AdaGrad /
- RMSprop /
- Adam optimizer /

$$\underbrace{P^T A^+ A P}$$



(Pretty much all of the most widely used optimization methods for training neural networks.)

BREAK

FIRST ORDER CONVEX OPTIMIZATION

First Order Optimization: Given a convex function f and a convex set \mathcal{S} ,

Goal: Find $\hat{\mathbf{x}} \in \mathcal{S}$ such that $f(\hat{\mathbf{x}}) \leq \min_{\mathbf{x} \in \mathcal{S}} f(\mathbf{x}) + \epsilon$.

Assume we have:

- Function oracle: Evaluate $f(\mathbf{x})$ for any \mathbf{x} .
- Gradient oracle: Evaluate $\nabla f(\mathbf{x})$ for any \mathbf{x} .
- Projection oracle: Evaluate $P_{\mathcal{S}}(\mathbf{x})$ for any \mathbf{x} .

Gradient descent requires $O\left(\frac{R^2 G^2}{\epsilon^2}\right)$ calls to each oracle to solve the problem.

We were only able to improve the ϵ dependence by making stronger assumptions on f (strong convexity, smoothness).

DIMENSION DEPENDENT BOUND

Alternatively, we can get much better bounds if we are willing to depend on the problem dimension. I.e. on d if $f(x)$ is a function mapping d -dimensional vectors to scalars.

We already know how to do this for a few special functions:

$$f(x) = \|\mathbf{Ax} - \mathbf{b}\|_2^2$$

where

$$\mathbf{A}^T (\mathbf{Ax} - \mathbf{b})$$

$$\mathbf{A} \in \mathbb{R}^{n \times d}.$$

$$(\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \mathbf{b}$$

DIMENSION DEPENDENT BOUND

Let $f(x)$ be bounded between $[-B, B]$ on \mathcal{S} . $m: \gamma$
 $x \in \mathcal{S}$ $f(x)$

Theorem (Dimension Dependent Convex Optimization)


There is an algorithm (the Center-of-Gravity Method) which finds \hat{x} satisfying $f(\hat{x}) \leq \min_{x \in \mathcal{S}} f(x) + \epsilon$ using $O(d \log(B/\epsilon))$ calls to a function and gradient oracle for convex f .

Caveat: Assumes we have some representation of \mathcal{S} , not just a projection oracle. We will discuss this more later.

Note: For an unconstrained problem with known starting radius R , can take \mathcal{S} to be the ball of radius R around $\underline{x^{(1)}}$. If $\|\nabla f(x)\|_2 \leq G$, we always have $B = \underline{O(RG)}$.

CENTER OF GRAVITY METHOD

Natural “cutting plane” method. Developed simultaneously on opposite sides of iron curtain.



Not used in practice (we will discuss why) but the basic idea underlies many popular algorithms.


A few basic ingredients:




1. The center-of-gravity of a convex set \mathcal{S} is defined as:

$$c = \frac{\int_{x \in \mathcal{S}} x \, dx}{\text{vol}(\mathcal{S})} = \frac{\int_{x \in \mathcal{S}} x \, dx}{\int_{x \in \mathcal{S}} 1 \, dx}$$

2. For two convex sets \mathcal{A} and \mathcal{B} , $\mathcal{A} \cap \mathcal{B}$ is convex. Proof by picture:




CENTER OF GRAVITY METHOD




Natural “cutting plane” method.

- $\underline{S}_1 = S$
- For $t = 1, \dots, T$:
 - $\mathbf{c}_t =$ center of gravity of S_t .
 - Compute $\nabla f(\mathbf{c}_t)$.
 - $\mathcal{H} = \{\mathbf{x} \mid \langle \nabla f(\mathbf{c}_t), \mathbf{x} - \mathbf{c}_t \rangle \leq 0\}$.
 - $S_{t+1} = S_t \cap \mathcal{H}$
- Return $\hat{\mathbf{x}} = \arg \min_t f(\mathbf{c}_t)$




Natural “cutting plane” method.

- $\mathcal{S}_1 = \mathcal{S}$
- For $t = 1, \dots, T$:
 - \mathbf{c}_t = center of gravity of \mathcal{S}_t .
 - Compute $\nabla f(\mathbf{c}_t)$.
 - $\mathcal{H} = \{\mathbf{x} \mid \langle \nabla f(\mathbf{c}_t), \mathbf{x} - \mathbf{c}_t \rangle \leq 0\}$.
 - $\mathcal{S}_{t+1} = \mathcal{S}_t \cap \mathcal{H}$
- Return $\hat{\mathbf{x}} = \arg \min_t f(\mathbf{c}_t)$



Natural “cutting plane” method.


- $\mathcal{S}_1 = \mathcal{S}$
- For $t = 1, \dots, T$:
 - \mathbf{c}_t = center of gravity of \mathcal{S}_t .
 - Compute $\nabla f(\mathbf{c}_t)$.
 - $\mathcal{H} = \{\mathbf{x} \mid \langle \nabla f(\mathbf{c}_t), \mathbf{x} - \mathbf{c}_t \rangle \leq 0\}$.
 - $\mathcal{S}_{t+1} = \mathcal{S}_t \cap \mathcal{H}$
- Return $\hat{\mathbf{x}} = \arg \min_t f(\mathbf{c}_t)$



CENTER OF GRAVITY METHOD

Intuitively, why does it make sense to search in $\mathcal{S}_t \cap \mathcal{H}$ where:

$$\mathcal{H} = \{\mathbf{x} \mid \langle \nabla f(\mathbf{c}_t), \mathbf{x} - \mathbf{c}_t \rangle \leq 0\}$$



CENTER OF GRAVITY METHOD

Intuitively, why does it make sense to search in $\mathcal{S}_t \cap \mathcal{H}$ where:

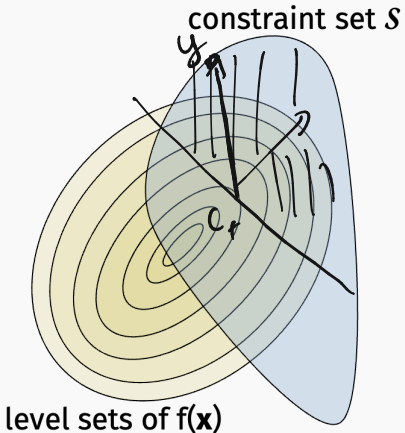
$$\mathcal{H} = \{\mathbf{x} \mid \langle \nabla f(\mathbf{c}_t), \mathbf{x} - \mathbf{c}_t \rangle \leq 0\}$$

By convexity,

$$\underline{f(\mathbf{y})} \geq \underline{f(\mathbf{c}_t)} + \underline{\langle \nabla f(\mathbf{c}_t), \mathbf{y} - \mathbf{c}_t \rangle}.$$

If $\mathbf{y} \notin \{\mathcal{S}_t \cap \mathcal{H}\}$ then

$\langle \nabla f(\mathbf{c}_t), \mathbf{y} - \mathbf{c}_t \rangle$ is ~~negative~~,
so $\underline{f(\mathbf{y})} > \underline{f(\mathbf{c}_t)}$. *positive*



CONVERGENCE THEOREM

Theorem (Center-of-Gravity Convergence)

Let f be a convex function with values in $[-B, B]$. Let \hat{x} be the output of the center-of-gravity method run for T iterations.

Then:

$$\underline{f(\hat{x})} - \underline{f(x^*)} \leq 2B \left(1 - \frac{1}{e}\right)^{T/d} \leq 2Be^{-T/3d} \leq \epsilon$$


If we set $T = 3d \log(2B/\epsilon)$, then $f(\hat{x}) - f(x^*) \leq \epsilon$.

$$T = 3d \log(2B/\epsilon)$$

. 67

KEY GEOMETRIC TOOL



Want to argue that, at every step of the algorithm, we “cut off” a large portion of the convex set we are searching over:



Theorem (Grünbaum's Theorem)

For any convex set \mathcal{S} with center-of-gravity \mathbf{c} , and any halfspace $\mathcal{Z} = \{\mathbf{x} \mid \langle \mathbf{a}, \mathbf{x} - \mathbf{c} \rangle \leq 0\}$ then:

$$\frac{\text{vol}(\mathcal{S} \cap \mathcal{Z})}{\text{vol}(\mathcal{S})} \geq \frac{1}{e} \approx \underline{\underline{.368}}$$



Want to argue that, at every step of the algorithm, we “cut off” a large portion of the convex set we are searching over.

Theorem (Grünbaum’s Theorem)

For any convex set \mathcal{S} with center-of-gravity \mathbf{c} , and any halfspace $\mathcal{Z} = \{\mathbf{x} \mid \langle \mathbf{a}, \mathbf{x} - \mathbf{c} \rangle \leq 0\}$ then:

$$\frac{\text{vol}(\mathcal{S} \cap \mathcal{Z})}{\text{vol}(\mathcal{S})} \geq \frac{1}{e} \approx .368$$


Let \mathcal{Z} be the compliment of \mathcal{H} from the algorithm. Then we cut off at least a $1/e$ fraction of the convex body on every iteration.

Corollary: After t steps, $\underline{\text{vol}(\mathcal{S}_t)} \leq \underline{\left(1 - \frac{1}{e}\right)^t \text{vol}(\mathcal{S})}$.


CONVERGENCE PROOF

Let δ be a small parameter to be chosen later.

Let $\mathcal{S}^\delta = \{(1 - \delta)\mathbf{x}^* + \delta\mathbf{x} \mid \mathbf{x} \in \mathcal{S}\}$.




Claim: Every point \mathbf{y} in \mathcal{S}^δ has good function value.



For any $\mathbf{y} \in S^\delta$:

$$\begin{aligned}
 \underline{\underline{f(\mathbf{y})}} &= \underline{\underline{f((1-\delta)\mathbf{x}^* + \delta\mathbf{x})}} \\
 &\leq (1-\delta)f(\mathbf{x}^*) + \delta f(\mathbf{x}) \\
 &\leq \underline{\underline{f(\mathbf{x}^*)}} - \underline{\underline{\delta f(\mathbf{x}^*)}} + \underline{\underline{\delta f(\mathbf{x})}} \\
 &\leq f(\mathbf{x}^*) + \underline{\underline{2B\delta}}.
 \end{aligned}$$

CONVERGENCE PROOF



We also have: $\text{vol}(\mathcal{S}^\delta) = \delta^d \text{vol}(\mathcal{S})$.


Set $\delta = \left(1 - \frac{1}{e}\right)^{T/d}$. After T steps,
 $\text{vol}(\mathcal{S}_t) \leq \left(1 - \frac{1}{e}\right)^T = \text{vol}(\mathcal{S}^\delta)$.

Either \mathcal{S}_t exactly equals \mathcal{S}^δ , in which case our next centroid gives error $\leq 2B\delta$.

Or we must have “chopped off” at least one point \mathbf{y} in \mathcal{S}^δ by the time we reach step T .

$$f^d = \left(\left(1 - \frac{1}{e}\right)^{T/d} \right)^d = \left(1 - \frac{1}{e}\right)^T \cdot \text{vol}(\mathcal{S})$$

CONVERGENCE PROOF



Claim: If we “chopped off” at least one point y in S^δ by the time we reach step T then for some centroid $c_1, \dots, c_t, \underline{f(c_t)} < 2B\delta$.

Proof:

$$f = \left(1 - \frac{1}{e}\right)^{T/d}$$

$$\underline{2B\delta} \geq f(y) \geq f(c_t) + \langle \nabla f(c_t), y - c_t \rangle > \underline{f(c_t)}.$$

Algorithm returns $\arg \min_{c_i} f(c_i)$.

CONVERGENCE THEOREM

Theorem (Center-of-Gravity Convergence)

Let f be a convex function with values in $[-B, B]$. Let $\hat{\mathbf{x}}$ be the output of the center-of-gravity method run for T iterations. Then:

$$f(\hat{\mathbf{x}}) - f(\mathbf{x}^*) \leq 2B \left(1 - \frac{1}{e}\right)^{T/d} \leq \underline{2Be^{-T/3d}}.$$

If we set $T = O(d \log(B/\epsilon))$, then $f(\hat{\mathbf{x}}) - f(\mathbf{x}^*) \leq \epsilon$.

In terms of gradient-oracle complexity, this is essentially optimal. So why isn't the algorithm used?

CENTROID COMPUTATION

In general computing the centroid is hard. #P-hard even when when \mathcal{S} is an intersection of half-spaces (a polytope).

Even if the problem isn't hard for your starting convex body \mathcal{S} , it likely will become hard for $\mathcal{S} \cap \mathcal{H}_1 \cap \mathcal{H}_2 \cap \mathcal{H}_3 \dots$

So while the oracle complexity of dimension-dependent optimization was settled in the 70s, a number of basic questions in terms of computational complexity.

We will discuss how to obtain a computationally efficient version of the center-of-gravity method called the **ellipsoid method**. This method is most famous for giving the first polynomial time algorithm for linear programming.

BREAK

Linear programs (LPs) are one of the most basic convex constrained, convex optimization problems:

Let $\mathbf{c} \in \mathbb{R}^d$, $\mathbf{b} \in \mathbb{R}^n$, $\mathbf{A} \in \mathbb{R}^{n \times d}$ be fixed vectors that define the problem, and let \mathbf{x} be our variable parameter.

$$\begin{aligned} \min f(\mathbf{x}) &= \mathbf{c}^T \mathbf{x} \\ \text{subject to } \mathbf{A}\mathbf{x} &\geq \mathbf{b}. \end{aligned}$$

Think about $\mathbf{A}\mathbf{x} \geq \mathbf{b}$ as a union of half-space constraints:

$$\{\mathbf{x} : \mathbf{a}_1^T \mathbf{x} \geq b_1\}$$


$$\{\mathbf{x} : \mathbf{a}_2^T \mathbf{x} \geq b_2\}$$

$$\vdots$$

$$\{\mathbf{x} : \mathbf{a}_n^T \mathbf{x} \geq b_n\}$$

$$\min f(x) = c^T x$$

subject to $Ax \geq b$.



- Classic optimization applications: industrial resource optimization problems were killer app in the 70s.
- Robust regression: $\min_{\mathbf{x}} \|\mathbf{Ax} - \mathbf{b}\|_1$.
- $L1$ constrained regression: $\min_{\mathbf{x}} \|\mathbf{x}\|_1$ subject to $\mathbf{Ax} = \mathbf{b}$. Lots of applications in sparse recovery/compressed sensing.
- Solve $\min_{\mathbf{x}} \|\mathbf{Ax} - \mathbf{b}\|_\infty$.
- Polynomial time algorithms for Markov Decision Processes.
- **Many combinatorial optimization problems can be solved via LP relaxations.**

Theorem (Khachiyan, 1979)

Assume $n = d$. The ellipsoid method solves any linear program with L -bit integer valued constraints in $O(n^4L)$ time.

Ellipsoid is a relatively simple center-of-gravity like method!

A Soviet Discovery Rocks World of Mathematics

By **MALCOLM W. BROWNE**

A surprise discovery by an obscure Soviet mathematician has rocked the world of mathematics and computer analysis, and experts have begun exploring its practical applications.

Mathematicians describe the discovery by L.G. Khachian as a method by which computers can find guaranteed solutions to a class of very difficult problems that have hitherto been tackled on a kind of hit-or-miss basis.

Apart from its profound theoretical interest, the discovery may be applicable

in weather prediction, complicated industrial processes, petroleum refining, the scheduling of workers at large factories, secret codes and many other things.

"I have been deluged with calls from virtually every department of government for an interpretation of the significance of this," a leading expert on computer methods, Dr. George B. Dantzig of Stanford University, said in an interview.

The solution of mathematical problems by computer must be broken down into a series of steps. One class of problem sometimes involves so many steps that it

could take billions of years to compute.

The Russian discovery offers a way by which the number of steps in a solution can be dramatically reduced. It also offers the mathematician a way of learning quickly whether a problem has a solution or not, without having to complete the entire immense computation that may be required.

According to the American journal Sci-

Continued on Page A20, Column 3

ONLY \$10.00 A MONTH!!! 24 Hr. Phone Answering Service. Totally New Concept!!! Incredible!!! 279-3870—ADVT.


Front page of New York Times, November 9, 1979.

PROBLEM SIMPLIFICATION

Simplifying the problem: Given a convex set \mathcal{K} via access to separation oracle $S_{\mathcal{K}}$ for the set, determine if \mathcal{K} is empty, or otherwise return any point $\mathbf{x} \in \mathcal{K}$.

$$S_{\mathcal{K}}(\mathbf{y}) = \begin{cases} \emptyset & \text{if } \mathbf{y} \in \mathcal{K}. \\ \text{separating hyperplane } (\mathbf{a}, c) & \text{if } \mathbf{y} \notin \mathcal{K}. \end{cases}$$


Let $\mathcal{H} = \{\mathbf{x} : \mathbf{a}^T \mathbf{x} = c\}$.



Example: How would you implement a separation oracle for a polytope $\{\mathbf{x} : \mathbf{Ax} \geq \mathbf{b}\}$.

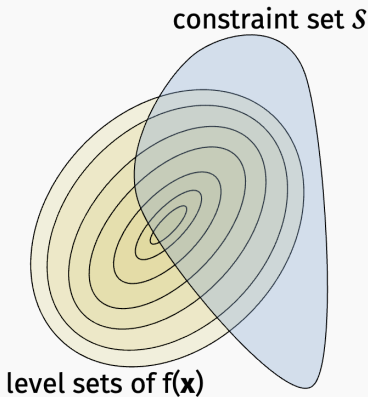
Original problem: $\min_{\mathbf{x} \in \mathcal{S}} f(\mathbf{x})$.

How to reduce to determining if a convex set \mathcal{K} is empty or not?



Original problem: $\min_{\mathbf{x} \in \mathcal{S}} f(\mathbf{x})$.

How to reduce to determining if a convex set \mathcal{K} is empty or not?

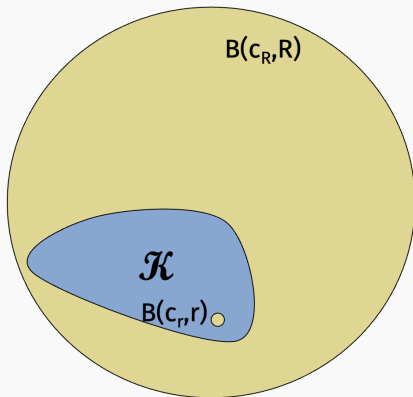


Binary search! For a convex function $f(\mathbf{x})$, $\{\mathbf{x} : f(\mathbf{x}) \leq c\}$ is convex, and you can get a separation oracle via the gradient.

ELLIPSOID METHOD SKETCH

Goal of ellipsoid algorithm: Solve “Is \mathcal{K} empty or not?” given a separation oracle for \mathcal{K} under the assumptions that:


1. $\mathcal{K} \subset B(\mathbf{c}_R, R)$.
2. If non-empty, \mathcal{K} contains $B(\mathbf{c}_r, r)$ for some $r < R$.



ELLIPSOID METHOD SKETCH


Iterative method similar to center-of-gravity:

1. Check if center \mathbf{c}_R of $B(\mathbf{c}_R, R)$ is in \mathcal{K} .
2. If it is, we are done.
3. If not, cut search space in half, using separating hyperplane.



ELLIPSOID METHOD SKETCH

Key insight: Before moving on, approximate new search region by something that we can easily compute the centroid of. Specifically an ellipse!!






Produce a sequence of ellipses that always contain \mathcal{K} and decrease in volume: $B(\mathbf{c}_R, R) = E_1, E_2, \dots$. Once we get to an ellipse with volume $\leq B(\mathbf{c}_r, r)$, we know that \mathcal{K} must be empty.

ELLIPSE

An ellipse is a convex set of the form: $\{\mathbf{x} : \|\mathbf{A}(\mathbf{x} - \mathbf{c})\|_2^2 \leq \alpha\}$ for some constant α and matrix \mathbf{A} . The center-of-mass is \mathbf{c} .

$$\{\mathbf{x} : \|\mathbf{I}(\mathbf{x} - \mathbf{c})\| < \alpha\} \quad \{\mathbf{x} : \|\mathbf{D}(\mathbf{x} - \mathbf{c})\| < \alpha\} \quad \{\mathbf{x} : \|\mathbf{A}(\mathbf{x} - \mathbf{c})\| < \alpha\}$$



Often re-parameterized to say that the ellipse is all \mathbf{x} with $\{\mathbf{x} : (\mathbf{x} - \mathbf{c})^T \mathbf{Q}^{-1} (\mathbf{x} - \mathbf{c}) \leq 1\}$

There is a closed form solution for the equation of the smallest ellipse containing a given half-ellipse. I.e. let E_j have parameters Q_j, c_j and consider the half-ellipse:

$$E_j \cap \{x : a_j^T x \leq a_j^T c_j\}.$$

Then E_{j+1} is the ellipse with parameters:


$$Q_{j+1} = \frac{d^2}{d^2 - 1} \left(Q_j - \frac{2}{d+1} h h^T \right) \quad c_{j+1} = c_j - \frac{1}{n+1} h,$$

where $h = \sqrt{a_j^T Q_j a_j} \cdot a_j$.

GEOMETRIC OBSERVATION

Claim: $\text{vol}(E_{i+1}) \leq (1 - \frac{1}{2d}) \text{vol}(E_i)$.


Proof: Via reduction to the “isotropic case”. I will post a proof on the course website if you are interested.



Not as good as the $(1 - \frac{1}{e})$ constant-factor volume reduction we got from center-of-gravity, but still very good!

GEOMETRIC OBSERVATION

Claim: $\text{vol}(E_{i+1}) \leq (1 - \frac{1}{2d}) \text{vol}(E_i)$



After $O(d)$ iterations, we reduce the volume by a constant.

In total require $O(d^2 \log(R/r))$ iterations to solve the problem.

Theorem (Khachiyan, 1979)

Assume $n = d$. The ellipsoid method solves any linear program with L -bit integer valued constraints in $O(n^4L)$ time. I.e. linear programming is in (weakly) polynomial time!

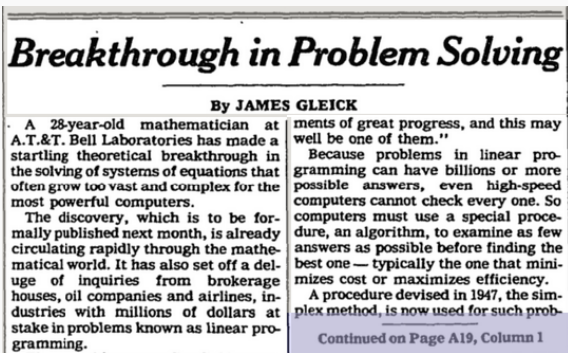
The method works for any convex program.

For LPs, we have an $O(nd)$ time separation oracle, and ellipsoid update take $O(d^2)$ time.

Careful analysis of the binary search method, how to set B_r and B_R , etc. leads to the final runtime bound.


Theorem (Karmarkar, 1984)

Assume $n = d$. The interior point method solves any linear program with L -bit integer valued constraints in $O(n^{3.5}L)$ time.



Front page of New York Times, November 19, 1984.


Will post lecture notes on the website (optional reading).



Projected Gradient Descent Optimization Path

INTERIOR POINT METHODS

Will post lecture notes on the website (optional reading).



Ideal Interior Point Optimization Path

Both results had a huge impact on the theory of optimization, although at the time neither the ellipsoid method or interior point method were faster than a heuristic known as the Simplex Method.

These days, improved interior point methods compete with and often outperform simplex.