

9/13

LP Duality

Primal LP:

$$\begin{aligned} \text{Variables } \vec{x} &= \langle x_1, \dots, x_n \rangle \\ \text{maximize } & \sum_{i=1}^n c_i x_i \quad \text{subject to constraints} \\ & \sum_{i=1}^n A_{ji} x_i \leq b_j, \quad j \in \{1, \dots, m\} \\ & x_i \geq 0 \quad \forall i \in \{1, \dots, n\} \end{aligned}$$

To find the dual, construct a variable w_j for each j s.t.

(1) $w_j \geq 0 \Rightarrow$ any feasible \vec{x} satisfies

$$\sum_{i=1}^n \left(\sum_{j=1}^m A_{ji} w_j \right) x_i \leq \sum_{j=1}^m b_j w_j$$

(2) $\forall i, \sum_{j=1}^m A_{ji} w_j \geq c_i \Rightarrow$ any feasible \vec{x} satisfies

$$\sum_{i=1}^n c_i x_i \leq \sum_{i=1}^n \left(\sum_{j=1}^m A_{ji} w_j \right) x_i \leq \sum_{j=1}^m b_j w_j$$

upper bound on objective of primal!

The **dual LP problem** is to find the best upper bound.
In other words,

$$\text{variables } \langle w_1, \dots, w_m \rangle = \vec{w}$$

$$\begin{aligned} \text{minimize } & \sum_{j=1}^m b_j w_j \quad \text{subject to constraints} \\ & \sum_{j=1}^m A_{ji} w_j \geq c_i \quad \forall i \in \{1, \dots, n\} \\ & w_j \geq 0 \quad \forall j \in \{1, \dots, m\} \end{aligned}$$

So, we have a primal LP and the dual LP, which is the optimization problem for the best upper bound.

Theorem Weak LP Duality

- (1) If the primal LP is unbounded ($+\infty$), the dual LP is infeasible.
- (2) If the primal LP is finite, dual LP is finite \leq primal, or infeasible.

Proof:

- (1) Suppose BWOC that the dual LP is feasible. Then, there is some upper bound on the primal LP. \times
- (2) Any feasible solution of the dual must upper bound the primal. \square

Theorem: Complementary Slackness

Consider feasible \vec{x} for the primal and feasible \vec{w} for dual. Then, the following are equivalent:

- (1) ($w_j = 0$ or $\sum_i A_{ji} x_i = b_j \forall j$) AND ($x_i = 0$ or $\sum_j A_{ji} w_j = c_i \forall i$)
(i.e. dual variable is 0 or primal bound j is tight, and vice versa)
- (1*) ($\lambda_j = 0$ or $\sum_i A_{ji} x_i = b_j \forall j \in S$) AND \vec{x} is optimal for LP²
- (2) $\sum_i c_i x_i = \sum_j w_j b_j$ (i.e. \vec{x} and \vec{w} are both optimal)
- (2*) \vec{x} is optimal for LP & \vec{w} gives best upper bound

Proof: We can say the following:

$$\underbrace{\sum_i c_i x_i - \sum_j w_j b_j}_{\text{holds because } \vec{w} \text{ is feasible}} \leq \underbrace{\sum_i (\sum_j A_{ji} w_j) x_i - \sum_j w_j b_j}_{\text{holds because } \vec{x}, \vec{w} \text{ feasible}} = \sum_i w_j (A_{ji} x_i - b_j) \leq 0$$

Condition (2) \Leftrightarrow this whole inequality being tight.

Condition (1) \Leftrightarrow no terms in the two middle sums.

So, (1) \Leftrightarrow (2). \square

Def: Begin with any primal LP. Then, the Lagrangian relaxation v.t. $\vec{\lambda}$ ($\lambda_j \geq 0 \forall j \in S$ where $S \subseteq \{1, \dots, m\}$) is

$$LP_S^2 = \text{Maximize } \sum_i c_i x_i + \sum_{j \in S} \lambda_j (b_j - \sum_i A_{ji} x_i) \text{ subject to}$$
$$\sum_i A_{ji} x_i \leq b_j \quad \forall j \notin S \quad \left(\text{make some } j\text{'s from constraint to objective} \right)$$
$$x_i \geq 0 \quad \forall i$$

Theorem: Weak Lagrangian Duality
 $\forall S \subseteq \{1, \dots, m\}$ and $\forall \vec{\lambda}$, $LP_S^1 \geq LP$

Proof: Let \vec{x} be feasible for LP. Then, \vec{x} must be feasible for LP_S^1 .

Also, since $\lambda_j \geq 0 \forall j$, $(b_j - \sum_i A_{ji} x_i) \lambda_j \geq 0$

So, we relax by allowing a larger space of feasible solutions and also by increasing the optimum.

Observe that we can search for $\vec{\lambda}$ that maximizes the objective for LP_S^1 .

Best upper bound = $\min_{\vec{\lambda} \text{ with } \lambda_j \geq 0 \forall j \in S} \left\{ \max_{\vec{x} \text{ feasible for } \mathcal{F}_S} \left\{ \sum_i c_i x_i + \sum_{j \in S} \lambda_j (b_j - \sum_i A_{ji} x_i) \right\} \right\}$

once S is fixed, every $\vec{\lambda}$ gives you a program, and every such program bounds the primal.

Note: when $S = \{1, \dots, m\}$, this best upper bound search is equivalent to the dual LP.

Theorem: Separating Hyperplane Theorem

Let P be a closed convex region in \mathbb{R}^n with $\vec{x} \notin P$.

Then, $\forall \vec{x} \notin P, \exists \vec{w} \in \mathbb{R}^n$ s.t. $\vec{x} \cdot \vec{w} > \max_{\vec{y} \in P} \{\vec{y} \cdot \vec{w}\}$



$H_{\vec{w}}$ is hyperplane $\vec{z} \in \mathbb{R}^n$

with $\vec{z} \cdot \vec{w} = \text{constant}$.

Then, there is some $\max_{\vec{y} \in P} \{\vec{y} \cdot \vec{w}\}$, and $\vec{x} \cdot \vec{w}$ is larger.

Lemma: Let \vec{x} solve the primal LP, and let $S = \{j : \sum_i A_{ij} x_i = b_j\}$

Then there exist $\{\lambda_j\}_{j \in S}$ s.t. $\lambda_j \geq 0 \forall j \in S$ and $c_i = \sum_{j \in S} \lambda_j A_{ji} \forall i$.

(i.e. for each condition j that \vec{x} tightly fits, there is a nice multiplier).

Proof: Let $X = \{\vec{y} : \exists \{\lambda_j \geq 0\}_{j \in S} \text{ s.t. } y_i = \sum_{j \in S} \lambda_j A_{ji}\}$

X is closed and convex. So, with the separating hyperplane theorem, if $\vec{c} \notin X$ we can improve our solution \vec{x} . So, $\vec{c} \in X$.

D

Theorem: Strong LP Duality

- (1) If primal is unbounded, the dual is infeasible.
- (2) If the primal is finite, the dual and primal are equal.
- (3) If the primal is infeasible, the dual is infeasible or unbounded.

Proof:

Set $w_j = 2_j \forall j \in S$, $w_j = 0 \forall j \notin S$

Because of the lemma, this is a feasible dual solution.

Now,

$$\sum_j b_j w_j = \sum_{j \in S} b_j w_j + \sum_{j \notin S} b_j w_j = \sum_{j \in S} \left(\sum_i A_{ji} x_i \right) w_j = \sum_i \left(\sum_{j \in S} A_{ji} w_j \right) x_i = \sum_i c_i x_i$$

= c_i from lemma

Lecture 9/15

LP Rounding

Motivating
Vibes: Turn LP-based problem to integer program,
solve normal LP, apply finesses to get integer solution.

ex/ Max-weight bipartite matching:

given bipartite $G(V=A \cup B, E \subseteq A \times B)$ and $w: E \rightarrow \mathbb{R}$ weights,
find matching set M of edges s.t. no node appears $>$ once
that maximizes $\max \sum_{e \in M} w_e$

Define x_e as follows:

x_e is an integer $e \in [0, 1]$

$x_e = 1 \Leftrightarrow e \in M$

$x_e = 0 \Leftrightarrow e \notin M$

The problem is to

maximize $\sum_e w_e \cdot x_e$

subject to $0 \leq x_e \leq 1 \forall e$

$\forall a \in A, \sum_{b \in B} x_{(a,b)} \leq 1$

$\forall b \in B, \sum_{a \in A} x_{(a,b)} \leq 1$

AND x_e IS AN
INTEGER

*the condition
makes the an
integer program.
If we remove
this we get
a LP that
can be solved
in poly-time,
but with
fractional solutions.*

We use the Birkhoff-Von Neuman Theorem, which states that any fractional matching to a set of convex integer matchings. Choosing any of these randomly will, in expectation, achieve the fractional expectation.

ex2 / vertex cover (LP-Mod):

Given $G=(V,E)$, weight $w:V \rightarrow \mathbb{R}$, output the set $S \subseteq V$ s.t. $\forall e \in E$, at least one endpoint of e is in S and S minimizes $\sum_{i \in S} w_i$

To convert this to an integer program, define indicator variables $x_i = \begin{cases} 1 & v_i \in S \\ 0 & v_i \notin S \end{cases}$. Then, we get the problem

variables: $x_i; \forall i \in V$
 minimize $\sum_{i \in V} w_i x_i$

~~(and x_i integer)~~

get rid of this, solve resulting LP

subject to

$$0 \leq x_i \leq 1 \quad \forall i \in V$$

$$\forall (u,v) \in E, \quad x_u + x_v \geq 1$$

at least one node is in cover

For a solution \tilde{x} to the LP, we can try to get an integer solution by rounding: place $i \in S$ iff $x_i \geq \frac{1}{2}$

(note: this is the best poly-time algo. for vertex cover)

Thm: Rounding outputs a valid vertex cover.

Proof: $\forall (u,v) \in E, \quad x_u + x_v \geq 1$. So, at least one of u,v must be in S . \square

Thm: Rounding outputs a 2-approx for the best vertex cover
 (i.e. $\sum_{i \in S} w_i \leq 2 \sum_{i \in V} w_i x_i$)

Proof: $\sum_{i \in V} w_i x_i = \sum_{i \text{ s.t. } x_i \geq \frac{1}{2}} w_i x_i + \sum_{i \text{ s.t. } x_i < \frac{1}{2}} w_i x_i \geq \sum_{i \in S} \frac{w_i}{2} + 0 = \sum_{i \in S} \frac{w_i}{2}$ \square

ex/ Distributed computing

The problem: n jobs, m machines, must assign each job processing job i on machine j takes time P_{ij}

The goal: Finish all jobs as quickly as possible.

Def indicator $x_{ij} = \begin{cases} 1 & \text{job } i \text{ put on machine } j \\ 0 & \text{else} \end{cases}$

$$\text{minimize } \max_j \left\{ \sum_i x_{ij} P_{ij} \right\}$$

(slowest machine)

A mathematical framework for this problem is to

$$\text{minimize } \max_j \left\{ \sum_i x_{ij} P_{ij} \right\} \text{ subject to } \forall i, \sum_j x_{ij} \geq 1$$

$$\forall i, j, x_{ij} \in [0, 1]$$

$$x_{ij} \text{ is integer}$$

If we define T as a variable st. $T \geq \sum_i x_{ij} P_{ij} \forall j$, minimizing T solves the program.

There is an **integrality gap**, i.e. there are instances where the best fractional solution $\leq \frac{1}{n}$ (best integral solution). ← consider case with 1 job, m machines, $P_{ij} = 1$. fractional best is $\frac{1}{m}$, best actual solution is 1.

This proves that any rounding to the relaxed LP's solution will suck.

Proof: Any rounding algorithm takes as input feasible \vec{x}^* and outputs an integral \vec{x} st. $\text{quality}(\vec{x}) \geq C \text{quality}(\vec{x}^*)$ for constant C .
Integrality gap disallows this. \square

If we add a constraint that considers the lower bound that all jobs must go somewhere, we can build

Answer: think of \rightarrow like the integral option


$$\text{LP}(\dagger): \text{minimize } T$$
$$\text{subject to } x_{ij} \in [0, 1]$$
$$\sum_j x_{ij} \geq 1 \forall i$$

$$T \geq \sum_j P_{ij} x_{ij} \forall j$$
$$x_{ij} = 0 \text{ if } P_{ij} > T$$

Key observation: If $t =$ integral optimum, the integral optimum is a feasible solution for LP(A).
 \Rightarrow int. opt. $\geq t$ & int. opt. $\geq T^k(t)$ ← soln. to LP(A)

Rounding algorithm:

\forall machines j , make $\lceil \sum_i x_{ij} \rceil$ copies of it as nodes $\in A$.
 \forall jobs j , make each one a single node $\in B$.
 So, we get a bipartite graph. In A , there are multiple nodes for each machine. In B , there is one node for each job.



Each job connects exactly x_{ij} to some copy of j & higher copies of each machine get worse jobs

Note: every job has 1 or 2 edges to machines because it will either fit inside one copy or not.

In other words, we start at earliest jobs.
 The number of copies of machine j is given by the LP solution. Each copy has capacity 1.
 We go in decreasing order of jobs, putting/splitting it in the earliest copy we can to fill capacities.
 The last copy of machine j might not be filled.

Claim 1: Let T_j^c be the smallest job assigned to copy c of machine j . Then, $T^k(t) \geq \sum_{c=2}^{\lceil \sum_i x_{ij} \rceil} T_j^c \quad \forall j$

This is a consequence of the ordering of jobs in decreasing time. As we go to lower copies, they were filled by better jobs.

* The algorithm is to find a complete matching in the graph and use it.

Claim 2: $T^*(A) \leq t + \sum_{c=2}^{\lfloor \frac{n}{2} \rfloor} T_c^j \quad \forall j$

Lecture 9/20 - Ellipsoid Algorithm

Really we refer to an LP

$$\begin{aligned} & \text{minimize} && \sum_i c_i x_i \\ & \text{subject to} && \sum_j A_{ij} x_j \leq b_i \quad \forall i \\ & && x_i \geq 0 \quad \forall i \end{aligned}$$

Sometimes we have disproportionately more constraints than variables.

Ex/Semidefinite programming

$X \in \mathbb{R}^{n \times n}$ is positive semidefinite $\iff \forall a \in \mathbb{R}^n, a^T X a \geq 0$

If we want X being pos. semi-def. to be a constraint, this is essentially infinitely many linear constraints.

This would still be an LP (linear objective, linear constraints), but you can't do anything in poly time over # of constraints.

Ex/Traveling Salesman (visit every node in graph along min weight path)

Let d_{ij} = dist. from i to j ,
we can write an IP

$x_{ij} = \mathbb{1}(i,j \text{ in path})$

minimize $\sum_{i,j} d_{ij} x_{ij}$

subject to $x_{ij} \in \{0,1\} \quad \forall i,j$ (integer constraint)

$\sum_j x_{ij} = 2 \quad \forall i$ (enter + exit = in/deg)

$\sum_{i \in S} \sum_{j \in S} x_{ij} \geq 2 \quad \forall S \subseteq V, S \neq \emptyset, S \neq V$
(every cut is crossed)

To relax this into an LP, we could remove the integer constraint. However, there are 2^n cut constraints (power set), so we don't want this approach.

These examples show that sometimes we wish to do something else. We generalize.

Def **Convex Programming** Real convex means $\frac{f(i)+f(j)}{2} \geq f(k) \forall k \in (i,j)$

A **convex program** is of the form

$$\begin{aligned} & \text{minimize} && f(\vec{x}) \\ & \text{subject to} && \vec{x} \in K \\ & && f \text{ is convex} \\ & && K \text{ is convex + closed} \end{aligned}$$

A hard problem would be to only do this with a membership oracle for K and a function evaluation oracle for f .

We can ask for a stronger assumption: a separation oracle.

Def: A **separation oracle** for a closed convex region K takes as input \vec{x} and outputs

$$\begin{cases} \text{"yes"} & \vec{x} \in K \\ \text{separating hyperplane } w & \vec{x} \notin K \end{cases}$$

A separation oracle can be thought of as a constant vector, where we either return "yes" or a violated constraint.

Consider now a convex program where all we are given is a linear objective $f(\vec{x})$ and a separation oracle.

Practice: - We can make a separation oracle for the Traveling Salesman (constant vector) by solving MinCut (poly time) for the graph with X_{ij} weights and verifying that the weight of the mincut is ≥ 7 .

- We can make a separation oracle for the Semidefinite Program by returning the eigenvector with a negative eigenvalue. (Poly time)

★ Ellipsoid Algorithm

Given as input a separation oracle for $K \subseteq [-B, B]^n$, output

$\begin{cases} \text{"yes"} & \text{Vol}(K) \geq \epsilon^n \\ \text{"no"} & K \text{ is empty} \end{cases}$

(Ex let $K = \{ \vec{x} \mid A\vec{x} \leq \vec{b} \}$ and $x \in [0, 1]^n$ and A_{ij}, b_i are rational numbers of c bits.

bounded by box

finite bits need these some wiggle room for solutions, given K volume

The plan is to check $K \cap \{ \vec{x} \mid f(\vec{x}) \leq C \}$ emptiness with the ellipsoid algorithm, and run binary search on C . This is easy if f is linear, but if f is just convex we use the fact that convex functions lie above the gradient hyperplane; so, we need a gradient oracle for convex f .

Def: An **ellipsoid** is defined by a center \vec{a} and a pos. semidefinite matrix B st.

$$E_{\vec{a}, B} = \{ \vec{x} \mid (\vec{x} - \vec{a})^T B (\vec{x} - \vec{a}) \leq 1 \}$$

The algorithm follows these rules:



① Query the oracle; either it is in K and we are done, or we get a separating hyperplane and have shrunk the potential volume for K by a multiplicative factor.

② Repeat a poly # of times until $\text{vol} K < \epsilon^n$

More precisely, the algorithm works by:

vars {

$E_0 =$ smallest ellipsoid containing $[-B, B]^n$ initial bounding box
 Define $p_i = \text{center}(E_i)$

loop {

while ($\text{vol}(E_i) \geq \epsilon^n$):
 if (separator oracle(p_i)):

- return p_i

else:

- get separating hyperplane \vec{w}_i, b_i

- update $E_{i+1} =$ smallest ellipsoid containing $E_i \cap \{ \vec{x} \mid \vec{w}_i \cdot \vec{x} \leq b_i \}$

makes feasible region larger, but is easy to compute and gives you p_{i+1} for free

losing space, but smallest ellipsoid contains this space is nice

return False

Lemma 1: We can find E_{i+1} given E_i, \vec{w}_i, b_i

Lemma 2: $\text{Vol}(E_{i+1}) \leq (1 - \frac{1}{2n}) \text{Vol}(E_i)$
shrinking factor

If we define the two problems for closed, convex K

separator oracle $\rightarrow \text{separate}_K(\vec{x}) = \begin{cases} \text{yes} & \vec{x} \in K \\ \vec{w} \text{ s.t. } \vec{x} \cdot \vec{w} > \max_{\vec{y} \in K} \{ \vec{y} \cdot \vec{w} \} & \vec{x} \notin K \end{cases}$

over convex space $\rightarrow \text{optimize}_K(\vec{c}) = \text{argmax}_{\vec{x} \in K} \{ \vec{x} \cdot \vec{c} \}$

We just saw a reduction from $\text{optimize}_K \rightarrow \text{separate}_K$.

We wish to prove a reduction from $\text{separate}_K \rightarrow \text{optimize}_K$

Theorem: Separate $K \rightarrow$ optimize K

Proof: Define an LP with vars \vec{w} s.t. we

$$\text{maximize } \sum_i x_i w_i = \vec{x} \cdot \vec{w}$$

$$\text{subject to } \sum_i w_i y_i = \vec{y} \cdot \vec{w} \leq 1 \quad \forall \vec{y} \in K$$

We see that if $\vec{x} \notin K$, $\exists \vec{w}'$ s.t. $\vec{x} \cdot \vec{w}' > \max_{\vec{y} \in K} \{\vec{y} \cdot \vec{w}'\}$

Let $\vec{w} = \frac{\vec{w}'}{\max_{\vec{y} \in K} \{\vec{y} \cdot \vec{w}'\}}$. This will clearly satisfy the constraint.

We seek a separation oracle for the region $\vec{w} \cdot \vec{y} \leq 1 \quad \forall \vec{y} \in K$, which we can do by optimizing $\max_{\vec{y} \in K} \{\vec{w} \cdot \vec{y}\}$ and comparing this

to 1. With this oracle, we can then optimize the initial LP via the ellipsoid algorithm to derive a separation oracle for K .

□

Lecture 9/22 - Semidefinite Programs

Some linear algebra background:

Def: A symmetric matrix $A \in \mathbb{R}^{n \times n}$ is **positive semidefinite** if $\forall \vec{x} \in \mathbb{R}^n, \vec{x}^T A \vec{x} \geq 0$.

The following are equivalent:

- (1) A symmetric matrix is PSD
- (2) A has all nonnegative eigenvalues
- (3) A can be written as $U^T U$ for some $U \in \mathbb{R}^{n \times n}$
 $\Leftrightarrow A_{ij} = \langle \vec{u}_i, \vec{u}_j \rangle$ for n vectors $\vec{u}_1, \dots, \vec{u}_n \in \mathbb{R}^n$

Obs: The set of all PSD matrices in $\mathbb{R}^{n \times n}$ is convex. ← $\forall x, y \in S, \frac{x+y}{2} \in S$

Proof:

Let A_1, A_2 be PSD. Then, $A = \frac{A_1 + A_2}{2}$ has

$$\vec{x}^T A \vec{x} = \frac{1}{2} \vec{x}^T (A_1 + A_2) \vec{x} = \frac{1}{2} (\vec{x}^T A_1 \vec{x} + \vec{x}^T A_2 \vec{x}) \geq 0 \quad \forall \vec{x} \in \mathbb{R}^n$$

□

A **semi-definite program** is a program of the form

maximize $\sum_{ij} c_{ij} x_{ij}$ ← objective can actually be minimizing any convex fn. or maximizing any concave fn.

subject to $\sum_{ij} A_{ijk} x_{ij} \leq b_k \quad \forall k$

$$X \text{ is PSD} \Leftrightarrow \exists \vec{v}_1, \dots, \vec{v}_n \in \mathbb{R}^n \text{ s.t. } x_{ij} = \langle \vec{v}_i, \vec{v}_j \rangle \quad \forall i, j$$

Equivalently, we can write a program to search over the vectors $\{\vec{v}_1, \dots, \vec{v}_n\}$

$$\Rightarrow \text{maximize } \sum_{ij} c_{ij} \langle \vec{v}_i, \vec{v}_j \rangle$$

subject to $\sum_{ij} A_{ijk} \langle \vec{v}_i, \vec{v}_j \rangle \leq b_k \quad \forall k$

$$\vec{v}_i \in \mathbb{R}^n \quad \forall i$$

Ex/ Max-Cut

Consider the NP-Hard problem **Max-Cut**:

Given undirected, unweighted graph, find $S \subseteq V$ ($S \neq \emptyset$, $S \neq V$) maximizing # of edges between S & \bar{S} ($\sum_{u \in S} \sum_{v \in \bar{S}} I((u,v) \in E)$)

The current best approach is to do an SDP relaxation (replace # w/ vectors)

We write the integer program

$$\begin{aligned} & \text{maximize} && \sum_{(i,j) \in E} \frac{1}{4} |u_i - u_j|^2 \\ & \text{subject to} && u_i \in \{-1, 1\} \quad \forall i; \end{aligned}$$

label for which side of cut u_i is on

To make this an integer SDP, we write the u_i labels as standard basis vectors:

$$\begin{aligned} & \text{(linear function of dot products)} && \text{maximize} && \sum_{(i,j) \in E} \frac{1}{4} \|\vec{u}_i - \vec{u}_j\|^2 = \sum_{(i,j) \in E} \frac{1}{4} (\langle \vec{u}_i, \vec{u}_i \rangle + \langle \vec{u}_j, \vec{u}_j \rangle - 2\langle \vec{u}_i, \vec{u}_j \rangle) \\ & \text{(linear constraints over dot products)} && \text{subject to} && \|\vec{u}_i\|^2 = 1 \quad \forall i \iff \langle \vec{u}_i, \vec{u}_i \rangle \geq 1, \langle \vec{u}_i, \vec{u}_i \rangle \leq 1 \quad \forall i \\ & && && \cancel{u_i \in \{\hat{e}_i, -\hat{e}_i\} \quad \forall i} \end{aligned}$$


We can relax the basis element constraint to get an SDP, which is poly-time solvable. Now, we must round.

Random Hyperplane Rounding:

- (1) Choose $\vec{c} \sim \mathcal{N}(0, 1)^n$ ($c_i \sim \mathcal{N}(0, 1)$ i.i.d)
- (2) Set $u_i = \text{sign}(\langle \vec{c}, \vec{u}_i \rangle)$

The hyperplane tangent to \vec{c} at the origin splits the space and cuts the graph. If $\vec{u}_i = \vec{u}_j$, they certainly are on the same side of the hyperplane. If $\vec{u}_i = -\vec{u}_j$, they are certainly on opposite sides. So, this has the properties we want.

Consider the space spanned by \vec{u}_i, \vec{u}_j .
 We can show \vec{c} lands randomly in the space.



$$\text{So, } P\{\text{round}(\vec{u}_i) \neq \text{round}(\vec{u}_j)\} = \frac{2\theta_{ij}}{2\pi} = \frac{\theta_{ij}}{\pi}$$

Thus, the number of edges in the cut is $\sum_{(i,j) \in E} \frac{\theta_{ij}}{\pi}$ in expectation

$$\text{The LP yields a max } \sum_{(i,j) \in E} \frac{1}{4} (\|\vec{u}_i\|^2 + \|\vec{u}_j\|^2 - 2\langle \vec{u}_i, \vec{u}_j \rangle) = \sum_{(i,j) \in E} \frac{1 - \cos(\theta_{ij})}{2}$$

We can find numerically that $\forall \theta, \frac{\theta/\pi}{\frac{1 - \cos \theta}{2}} \geq 0.878$

So, the rounded solution is a 0.878-approx of the optimal solution to the relaxed SDP.

Ex/ MAX 2SAT (NP Hard)

Given n literals and m clauses w/ 2 literals each,

i.e. clause $c_l \in \{x_i \vee x_j, x_i \vee \neg x_j, \neg x_i \vee x_j, \neg x_i \vee \neg x_j\}$,
 we want to set the literals to maximize the # of satisfied clauses.

We write

$$\text{maximize } \sum_l \frac{1 - (1 - y_l^2)(1 - y_l^2)}{4}$$

$$\text{subject to } x_i^2 = 1 \quad \forall i \\ x_i \in \{-1, 1\} \quad \forall i$$

$$\left(y_{ij} = x_i \text{ if } j^{\text{th}} \text{ literal in } l \text{ is } x_i, \right. \\ \left. -x_i \text{ if } j^{\text{th}} \text{ literal in } l \text{ is } \neg x_i \right)$$

To vectorize this and relax it into an SDP, we want

$$\begin{aligned} \text{maximize} \quad & \sum_k 1 - \frac{\langle \tilde{x}_0 - \tilde{y}_k, \tilde{x}_0 - \tilde{y}_k \rangle}{4} \\ \text{subject to} \quad & \|x_i\|^2 = 1 \quad \forall i \end{aligned}$$

(filling the role of true)
 $\left(\begin{array}{l} \tilde{y}_k = \tilde{x}_i \text{ if } \dots \\ \text{else } -\tilde{x}_i \end{array} \right)$
 $\left(\tilde{x}_0 \text{ is any } L_2\text{-normalized fixed vector} \right)$

We get a solution to this SDP in poly-time.

Rounding:

(1) Pick a random direction $\tilde{z} \sim \mathcal{N}(0, I)^n$

(2) Set $x_i = \text{sign}(\langle \tilde{z}, \tilde{x}_i \rangle \cdot \langle \tilde{z}, \tilde{x}_0 \rangle)$

(the sign of this is the "true" side of \tilde{z})

Lecture 9/27 - Submodular Function Minimization

Submodular Functions

Def:

Let N be a set of n elements. A function $f: 2^N \rightarrow \mathbb{R}$ is **submodular** if

$$(1) \forall A \subseteq B \subseteq N \text{ and } \forall j \notin B, \quad f(A \cup \{j\}) - f(A) \geq f(B \cup \{j\}) - f(B)$$

or equivalently

$$(2) \forall S, T \subseteq N \quad f(S \cup T) + f(S \cap T) \leq f(S) + f(T)$$

\leftarrow power set of N

(Diminishing marginal returns)

Ex/ **Cut Function**

If $G=(V,E)$ is some graph and $N=V$, $f(S)$ is the weight of edges from S to \bar{S} , then f is submodular if all edges have nonnegative weight.

Ex/ **Bipartite Coverage Functions**

If $G=(V,E)$ bipartite, N is the set of left-hand nodes, $f(S)$ is the # of right hand nodes with an edge to something in S . Then, f is submodular.

SFM: Given submodular f , find $\operatorname{argmin}_{S \subseteq N} \{f(S)\}$

Note: Because submodular functions can be silly slow, we work in terms of **value oracle** access to $f(\cdot)$. So, we count polynomial runtime and counting the # of queries to this oracle.

Define a function $\hat{f}: [0,1]^N \rightarrow \mathbb{R}$ st. $\forall S \subseteq N$,
 $\hat{f}(S) = \hat{f}(\text{vector with } x_i = 1 \forall i \in S, x_i = 0 \forall i \notin S)$ and $\hat{f}(S) = f(S) \forall S$.

\hat{f} is extension of f from discrete inclusion of elements to $[0,1]^n$

We want to show that \hat{f} is convex $\Leftrightarrow f$ is submodular.
 Then, since \hat{f} and f agree over 2^N , we can minimize \hat{f} , and we want to use this to minimize f .

Claim: Given an evaluation oracle and a gradient oracle for convex \hat{f} , we can minimize \hat{f} over $[0,1]^n$ in poly time via the ellipsoid algorithm.

Proof: Recall that the ellipsoid algorithm works as follows:

Ellipsoid(K): given K convex and bounded ($\exists H$ s.t. $K \subseteq [-H, H]^n$) and a poly time separation oracle for K , determine in poly time whether K is empty.

To use ellipsoid as a subroutine, let $K_c = [0,1]^n \cap \{\vec{x} \mid \hat{f}(\vec{x}) \leq C\}$. K_c is convex because \hat{f} is convex, and it's bounded by $[0,1]^n$. We can check if $\vec{x} \in K_c$ by checking $x_i \in [0,1] \forall i$ and querying the evaluation oracle.

To find a separating hyperplane if $\vec{x} \notin K_c$, we can return

hyperplane = $\begin{cases} \text{the hyperplane } \{\vec{y} \mid y_i < x_i\} & \text{if } x_i \notin [0,1] \text{ for some } i \\ \{\vec{y} \mid \vec{\nabla} \hat{f}(\vec{x}) \cdot (\vec{y} - \vec{x}) \leq C - \hat{f}(\vec{x})\} & \end{cases}$ \square

Def: For a function $f: \{0,1\}^n \rightarrow \mathbb{R}$, the **Lovász extension** $\hat{f}: [0,1]^n \rightarrow \mathbb{R}$ is $\forall \vec{x} \in [0,1]^n \quad \hat{f}(\vec{x}) = \mathbb{E}_{\mathcal{I} \sim U(\{0,1\})} \{f(\{i \mid x_i \geq \lambda\})\}$

Sample random threshold, include all coordinates above the threshold.

We observe that there are only $n+1$ sets to query an.
 To see this, suppose wlog that \vec{x} is s.t. $x_1 \geq \dots \geq x_n$. Then, the possible sets are $\{\emptyset\}, \{x_1\}, \{x_1, x_2\}, \dots, \{x_1, \dots, x_n\}$ with thresholds $1 \geq \lambda > x_1, x_1 \geq \lambda > x_2, x_2 \geq \lambda > x_3, \dots, x_n \geq \lambda \geq 0$ that occur with \mathbb{P} 's $\mathbb{P} = 1 - x_1, \mathbb{P} = x_1 - x_2, \mathbb{P} = x_2 - x_3, \dots, \mathbb{P} = x_n - 0$

So, for such monotonically decreasing \hat{x} ,

$$\hat{f}(\hat{x}) = \sum_{i=0}^n (x_i - x_{i+1}) f(\{1, \dots, i\}) \quad (x_0=1, x_{n+1}=0)$$

Then, $\frac{\partial \hat{f}(\hat{x})}{\partial x_i} = f(\{1, \dots, i\}) - f(\{1, \dots, i-1\})$.

We can create a gradient oracle in $n \log n$ time + n evaluations of f . ← needed to sort

Theorem: \hat{f} is convex \iff f is submodular

Proof: For simplicity, suppose WOLOG that

1) $x_1 \geq \dots \geq x_n$ (we can relabel)

2) $f(\emptyset) = 0$ (shifting f doesn't change submodularity)

(submodular \Rightarrow convex)

Define $P_f = \left\{ \vec{w} \mid \forall S \subseteq \{1, \dots, n\}, \sum_{i \in S} w_i \leq f(S) \text{ and } \sum_{i=1}^n w_i = f(\{1, \dots, n\}) \right\}$ ← power set

Define $g(\vec{z}) = \max_{\vec{w} \in P_f} \left\{ \vec{z} \cdot \vec{w} \right\}$. This forms a primal LP maximized by g .

(1) g is convex because the supporting line of a linear segment will lie below (or along) the max of the next part. ✗

(2) f submodular $\Rightarrow g = \hat{f}$. To see this, we write the dual

minimize $\sum_{S \subseteq \{1, \dots, n\}} y_S f(\{1, \dots, S\})$

subject to $\sum_{i \in S} y_S = z_i \quad \forall i$

$y_S \geq 0 \quad \forall S \subseteq \{1, \dots, n\}$

We propose optimal $w_i^* = f(\{1, \dots, i\}) - f(\{1, \dots, i-1\}) = \nabla \hat{f}$

$$y_S^* = \begin{cases} z_i - z_{i-1} & S = \{1, \dots, i\} \\ z_n & S = \{1, \dots, n\} \\ 0 & \text{else} \end{cases}$$

We want to show

$$(1) \sum_{i \in S} w_i^* = \hat{f}\left(\frac{1}{2}\right)$$

$$(2) \sum_{S \subseteq \{1, \dots, n\}} y_S^* f(S) = \hat{f}\left(\frac{1}{2}\right)$$

(3) y^* is feasible in the dual

(4) w^* is feasible in the primal

} true $\forall f$

Since this will imply that there are optimal solutions, and therefore that $\hat{f} = g$. We prove (4) by induction on $|S|$.

For a given S let i be the largest index in S . Since f is submodular,

$$\begin{aligned} f(S) + f(\{1, \dots, i-1\}) &\geq f(S \cup \{1, \dots, i-1\}) + f(S \cap \{1, \dots, i-1\}) \\ &= f(\{1, \dots, i\}) + f(S \setminus \{i\}) \end{aligned}$$

$$\begin{aligned} \Rightarrow f(S) &\geq f(\{1, \dots, i\}) - f(\{1, \dots, i-1\}) + f(S \setminus \{i\}) \\ &= w_i^* + f(S \setminus \{i\}) \stackrel{\text{IH}}{=} \sum_{i \in S} w_i^* \end{aligned}$$

So, w_i^* is feasible in the primal.

Then, we can optimize \hat{f} in poly time.

Lecture 9/29 - Concentration Bounds

Vibes: what can we say about a random variable and how close it usually/always is to its expectation?

Notation: For the below notes, S is a subset of the power set $\{0,1\}^N$

Recall **Markov's Inequality**:

Let X be a nonnegative random variable. Then,

$$\mathbb{P}\{X > c \mathbb{E}\{X\}\} \leq \frac{1}{c} \quad \forall c > 0$$

and **Chebyshev's Inequality**:

Let X be a RV with mean μ and variance σ^2 . Then,

$$\mathbb{P}\{|X - \mu| \geq c\sigma\} \leq \frac{1}{c^2} \quad \forall c > 0$$

Chernoff Bounds

★ We ask what if we draw n random variables that are independent and bounded. CLT means they approach Gaussian!

Formally, what if we have random variables X_1, \dots, X_n that are independent and s.t. $X_i \in [0,1] \forall i$. What can we say about $X = \sum_{i=1}^n X_i$?

Theorem: Let X_1, \dots, X_n be independent with $X_i \in \{0,1\} \forall i$. Then,

$$\mathbb{P}\left\{\sum_{i=1}^n X_i \geq (1+\epsilon) \mathbb{E}\left\{\sum_{i=1}^n X_i\right\}\right\} \leq e^{-\frac{\epsilon^2 \mathbb{E}\left\{\sum_{i=1}^n X_i\right\}}{3+3\epsilon}}$$

← grows exponentially in expectation of the sum

Proof: Let $X = \sum_{i=1}^n X_i$. Let $p_i = \mathbb{E}\{X_i\} \forall i$:

Pick t to set later, and look at the random variable e^{tX} .
Observe that

$$\begin{aligned} \mathbb{E}\{e^{tX}\} &= \mathbb{E}\left\{\prod_{i=1}^n e^{tX_i}\right\} \stackrel{X_i \text{ independent}}{=} \prod_{i=1}^n \mathbb{E}\{e^{tX_i}\} = \prod_{i=1}^n \left((1-p_i) + p_i e^t\right) \\ &= \prod_{i=1}^n \left(1 + p_i(e^t - 1)\right) \leq \prod_{i=1}^n e^{p_i(e^t - 1)} \stackrel{\sum p_i = \mathbb{E}\{X\}}{=} e^{(e^t - 1) \mathbb{E}\{X\}} \end{aligned}$$

$1 - x \leq e^{-x}$

We can see that, since e^{tx} is monotone,
 $\mathbb{P}\{X > (1+\epsilon)\mathbb{E}\{X\}\} = \mathbb{P}\{e^{tX} > e^{t(1+\epsilon)\mathbb{E}\{X\}}\}$

By Markov's Inequality, this is bounded by

$$\mathbb{P}\{X > (1+\epsilon)\mathbb{E}\{X\}\} \leq \frac{e^{(e^t-1)\mathbb{E}\{X\}}}{e^{t(1+\epsilon)\mathbb{E}\{X\}}}$$

Letting $t = \ln(1+\epsilon)$ and noting $(1+\epsilon)\ln(1+\epsilon) > \epsilon + \frac{\epsilon^2}{3}$ for $\epsilon \in [0, 1]$,

$$\mathbb{P}\{X > (1+\epsilon)\mathbb{E}\{X\}\} \leq e^{\mathbb{E}\{X\}(\epsilon - (1+\epsilon)\ln(1+\epsilon))} \leq e^{-\frac{\epsilon^2\mathbb{E}\{X\}}{3}} \leq e^{-\frac{\epsilon^2\mathbb{E}\{X\}}{3+3\epsilon}}$$

→ to let it hold for $\epsilon=1$

READ NOTES HERE for Chernoff applications

Examples that look like sum of random variables we can use Chernoff on, but aren't!

Ex 1 / Fixed ^{unweighted} graph G . Put v in a set S independently with probability p_v to get a random cut. What is the value of $\text{cut}_G(S)$.

Ex 2 / Let F be a subset of the power set $\{0, 1\}^N$. Put v in S independently with probability p_v . We can use Chernoff bounds on the size $|S|$, but not on functions like

$$a) \max_{T \subseteq F} \{|S \cap T|\} \quad \text{or} \quad b) f(S) = \begin{cases} |S| & 0 \leq |S| < \sqrt{n} \\ \sqrt{n} & \sqrt{n} \leq |S| \leq \frac{n+\sqrt{n}}{2} \\ \sqrt{n} + |S| & \frac{n+\sqrt{n}}{2} < |S| < \frac{n+\sqrt{n}}{2} \\ 2\sqrt{n} & |S| \geq \frac{n+\sqrt{n}}{2} \end{cases}$$

Defn: A function f is **c-Lipschitz** if $\forall S \subseteq N$ and $j \in N$,
 $|f(S \cup \{j\}) - f(S)| \leq c$
bounded differences

Theorem: McDiarmid's Inequality

Let X_1, \dots, X_n be independent random variables, and let $f(\dots)$ satisfy bounded differences for c_1, \dots, c_n

$$\text{(i.e. } \forall i, \vec{x}_i, x_i, x_i' \quad |f(\vec{x}_i, x_i) - f(\vec{x}_i, x_i')| \leq c_i)$$

holding everything except X_i constant makes f c_i -Lipschitz

Then,

$$\mathbb{P}\{|f(\vec{X}) - \mathbb{E}\{f(\vec{X})\}| > \epsilon\} \leq 2e^{-\frac{2\epsilon^2}{\sum c_i^2}}$$

(Note, when f is 1-Lipschitz, $\mathbb{P}\{\dots\} \leq 2e^{-\frac{2\epsilon^2}{n}}$. So, when $\epsilon > \sqrt{n}$ this is cool.)

Better Theorem: Schoetman

Let f be subadditive and 1-Lipschitz, and let X_1, \dots, X_n be independent. Let a be the median of $f(\vec{X})$. Then,

$$\mathbb{P}\{f(\vec{X}) \geq 3a + k\} \leq 2^{-k} \quad \forall k > 0$$

Note that Example 1 above is submodular, but non-monotone and Example 2(a) above is XOS / fractionally subadditive

Def: A function f is **XOS** if there exist additive functions f_1, \dots, f_n s.t. $f(s) = \max_i \{f_i(s)\}$

Def: A function f is **(a,b)-self-bounding** if there exist f_1, \dots, f_n s.t. $0 \leq f(s) - f_i(s) \leq 1 \quad \forall i \in \{1, \dots, n\}$
 $\sum_{i=1}^n f_i(s) \leq af(s) + b \quad \forall s$

Theorem: (a,b)-self-bounded functions are Chernoff bounded.

Corollary: Since XOS functions are (1,0)-self-bounded and non-monotone submodular functions are (? , 0)-self-bounded, XOS & nonmod SM are Chernoff bounded.

Lecture 10/4 Streaming I

Streaming algorithms process large data in a small space (low memory usage).

The input stream is a sequence of inputs a_1, \dots, a_n that is processed in sequence order.

Ex Approximate Counting

Maintain a counter n initialized to 0, supporting

- $inc()$: $n \leftarrow n+1$ (no more than N $inc()$)

- $query()$: return an approximation $\tilde{n} = (1 \pm \epsilon)n$ with high probability.

We can solve this trivially with $\log(N)$ bits by maintaining exact counter. This can be quite big.

Question:

Can we represent numbers $n \in \{1, \dots, N\}$ using $\ll \log N$ bits s.t. we can recover $\tilde{n} \in [\frac{n}{2}, 2n]$ from the encoding?

We can approximate n by only storing intervals, such as the nearest power of 2 (to get a 2-approx).

For $n \in [2^x, 2^{x+1})$, we can store x in $O(\log \log N)$ bits.

We can handle increments by increasing x with probability 2^{-x} , such that we, in expectation, increment x when we should. The algorithm looks like

$init()$: $x \leftarrow 0$ after first $inc()$

$query()$: return 2^x

$inc()$: $x \leftarrow \begin{cases} x+1 & \text{w.p. } 2^{-x} \\ x & \text{w.p. } 1-2^{-x} \end{cases}$

Analysis: Let X_n be the R.V. X after n calls to $\text{mc}()$
 We WTS that $\mathbb{E}\{2^{X_n}\} = n$ and $\text{Var}\{2^{X_n}\} = O(n^2)$

Proof:

$$\begin{aligned} \mathbb{E}\{2^{X_n}\} &= \sum_x \mathbb{P}\{X_n = x\} 2^x = \sum_x (\mathbb{P}\{X_{n-1} = x\} \cdot (1-2^{-x}) + \mathbb{P}\{X_{n-1} = x-1\} \cdot 2^{-(x-1)}) 2^x \\ &= \sum_x \mathbb{P}\{X_{n-1} = x\} (2^x - 1) + \sum_x \mathbb{P}\{X_{n-1} = x-1\} \cdot 2 = \mathbb{E}\{2^{X_{n-1}}\} + 2 \\ &= \mathbb{E}\{2^{X_{n-1}}\} + 1 \Rightarrow \mathbb{E}\{2^{X_n}\} = n. \end{aligned}$$

Similar logic works for the variance. □

We can apply Chebyshev's Inequality $\mathbb{P}\{|Y - \mathbb{E}\{Y\}| > T\} \leq \frac{\text{Var}\{Y\}}{T^2}$
 to get

$$\mathbb{P}\{|2^{X_n} - n| > T\} \leq O\left(\left(\frac{n}{T}\right)^2\right)$$

Means

We can reduce the variance by averaging s independent copies.
 Let $X^{(i)}$ denote X in the i th copy. Then, letting $X^* = \frac{1}{s} \sum X^{(i)}$ be
 the average, $\mathbb{E}\{2^{X^*}\} = n$ and $\text{Var}\{2^{X^*}\} = \frac{1}{s^2} \cdot s O(n^2) = O\left(\frac{n^2}{s}\right)$

Chebyshev now gives

$$\mathbb{P}\{|2^{X^*} - n| > T\} \leq O\left(\frac{1}{s} \left(\frac{n}{T}\right)^2\right)$$

If we set $T = \varepsilon n$, $s = \frac{1}{\varepsilon^2 \delta}$, we get $\mathbb{P}\{|2^{X^*} - n| > \varepsilon n\} < \delta$

Total space used is $\Theta\left(\frac{1}{\varepsilon^2 \delta} \log \log n\right)$

Median of means

Maintain $s_1 \cdot s_2$ independent copies. On query, divide into s_1 groups
 of size s_2 . Let $X^{(i,j)}$ be the j th X of group i . For each group i ,
 compute $\tilde{n}_i = \frac{1}{s_2} \sum_j 2^{X^{(i,j)}}$. Let \tilde{n} be the median of $\tilde{n}_1, \dots, \tilde{n}_{s_1}$.

If we set s_2 to $\Theta\left(\frac{1}{\varepsilon^2}\right)$, $\mathbb{P}\{|\tilde{n} - n| > \varepsilon n\} < \frac{1}{4}$

We can find that

$$\tilde{n} > (1+\epsilon)n \Leftrightarrow \geq \frac{s_i}{2} \text{ groups have } \tilde{n}_i > (1+\epsilon)n$$

$$\tilde{n} < (1-\epsilon)n \Leftrightarrow \geq \frac{s_i}{2} \text{ groups have } \tilde{n}_i < (1-\epsilon)n$$

$$\text{Let } Y_i = \begin{cases} 1 & \text{if } \tilde{n} > (1+\epsilon)n \\ 0 & \text{else} \end{cases}$$

We know that

① All the Y_i are independent

$$\text{② } \tilde{n} > (1+\epsilon)n \Leftrightarrow \sum Y_i \geq \frac{s_i}{2}$$

$$\text{③ } \mathbb{E}\{Y_i\} < \frac{1}{n}$$

$$\text{With Chernoff, } \mathbb{P}\{\tilde{n} > (1+\epsilon)n\} = \mathbb{P}\{\sum_{i=1}^{s_i} Y_i \geq \mathbb{E}\{\sum Y_i\} \cdot 2\} \leq e^{-\frac{s_i}{2}}$$

$$\text{Similarly, } \mathbb{P}\{\tilde{n} < (1-\epsilon)n\} \leq e^{-\frac{s_i}{2}}$$

If we set s_i to $\Theta(\log(\frac{1}{\delta}))$, union bound yields

$$\mathbb{P}\{|\tilde{n} - n| \leq \epsilon n\} > 1 - \delta$$

$$\text{Total space used is } O\left(\frac{1}{\epsilon^2} \log\left(\frac{1}{\delta}\right) \log \log N\right)$$

Morris Counter

If we instead set the base to be $(1+\alpha)$ instead of 2,

we get $\mathbb{P}\{|\tilde{n} - n| \leq \epsilon n\} > 1 - \delta$ with space $O(\log(\frac{1}{\epsilon}) + \log \log N + \log \log(\frac{1}{\delta}))$

Lecture 10/6 - Streaming II

Ex/ Distinct Elements

Input: a stream a_1, \dots, a_n ($a_i \in \{1, \dots, U\}$),

Output: estimate \tilde{F} of # of distinct elements
s.t. $\tilde{F} = (1 \pm \epsilon)F$ w.p. $\geq 1 - \delta$.

Naive Solution

Store all distinct elements! $O(n \log U)$ space

Subset Sampling

Not accurate ;)

(Recall that if X_1, \dots, X_F are independent RV's with $X_i \sim U[0, 1]$ ^{uniform} and $X^{(k)}$ is the k th smallest one, then $\mathbb{E}\{X^{(k)}\} = \frac{k}{F+1}$)

We can use this in reverse: find $X^{(k)}$ for some k to estimate F .

KMV (k-minimum value)

Algorithm

Ideally, assume access to a random hash function $h: \{1, \dots, U\} \rightarrow [0, 1]$. Have a parameter $k \geq 1$ to set later.

- initialize a set S to \emptyset to store the k smallest hash values.
 - for i in $\{1, \dots, n\}$:
 $S \leftarrow S \cup \{h(a_i)\}$
 if $|S| > k$: remove $\max\{S\}$ from S
 - if $|S| = k$: return $\tilde{F} = \frac{k}{\max(S)} - 1$
 else: return $\tilde{F} = |S|$
- ← turns out to be important*

Analysis

We want two things

- ① upper bound on $\mathbb{P}\{\tilde{F} > (1+\epsilon)F\}$
- ② upper bound on $\mathbb{P}\{\tilde{F} < (1-\epsilon)F\}$

} these are v. similar, so we focus on ①

We can find

$$\mathbb{P}\{\tilde{F} > (1+\epsilon)F\} = \mathbb{P}\left\{\frac{k}{\max\{s\}} > (1+\epsilon)F\right\} = \mathbb{P}\left\{\max\{s\} < \frac{k}{(1+\epsilon)F}\right\}$$

where $\max\{s\}$ is the k th smallest hash value.

Let v_1, \dots, v_F be the hash values of the elements.

independent $\forall v_i, \mathbb{P}\left\{v_i < \frac{k}{(1+\epsilon)F}\right\} = \frac{k}{(1+\epsilon)F}$

Let X be a RV denoting the # of v_i s.t. $v_i < \frac{k}{(1+\epsilon)F}$

$$\Rightarrow \mathbb{E}\{X\} = \sum_{i=1}^F \frac{k}{(1+\epsilon)F} = \frac{k}{1+\epsilon}$$

$$\Rightarrow \text{Var}\{X\} = \sum_{i=1}^F \text{Var}\left\{v_i < \frac{k}{(1+\epsilon)F}\right\} = F \left(\frac{k}{(1+\epsilon)F} - \left(\frac{k}{(1+\epsilon)F}\right)^2 \right) < k$$

By Chebyshev,

$$\mathbb{P}\{X \geq k\} \leq \frac{\text{Var}\{X\}}{(k - \frac{k}{1+\epsilon})^2} < \frac{k(1+\epsilon)^2}{k^2 \epsilon^2} = O\left(\frac{1}{\epsilon^2 k}\right)$$

If we set $k = \frac{c}{\epsilon^2}$ $\mathbb{P}\{\tilde{F} > (1+\epsilon)F\} < O\left(\frac{1}{\epsilon}\right)$

We can apply similar logic to find that $\mathbb{P}\{\tilde{F} < (1-\epsilon)F\} < O\left(\frac{1}{\epsilon}\right)$

By Union Bound, $\mathbb{P}\{\tilde{F} \in (1\pm\epsilon)F\} > 1 - O\left(\frac{2}{\epsilon}\right)$

using space $O\left(\frac{1}{\epsilon^2}\right)$ "real numbers".

We can do better with the **median trick**: maintain T independent copies and output the median of the predictions. We saw last time that this yields

$$\mathbb{P}\{\text{median} \in (1\pm\epsilon)F\} \geq 1 - e^{-\Theta(T)}. \text{ Setting } T = O\left(\log\left(\frac{1}{\delta}\right)\right), \mathbb{P}\{\dots\} \geq 1 - \delta.$$

Note that this algorithm assumes

- ① storing real numbers in $[0, 1]$
- ② random hash function

\uparrow
space = $O\left(\frac{\log\left(\frac{1}{\delta}\right)}{\epsilon^2}\right)$ real #'s

Removing the Assumptions

① Discretize $[0, 1]$ to $\{\frac{1}{M}, \frac{2}{M}, \dots, \frac{M-1}{M}, 1\}$. We get a "rounding error" $\leq O(\frac{1}{M})$.
If we set $M=U$, things work out the same.

② Defn let \mathcal{H} be a family of hash functions

$\{1, \dots, U\} \rightarrow \{1, \dots, M\}$. \mathcal{H} is **c-wise independent** if

$\forall x_1, \dots, x_c \in \{1, \dots, U\}$ distinct, $\forall a_1, \dots, a_c \in \{1, \dots, M\}$,

$$\mathbb{P}_{h \in \mathcal{H}} \{ \forall i \in \{1, \dots, c\}, h(x_i) = a_i \} = \frac{1}{M^c}$$

Recall that there exists pairwise independent \mathcal{H} of size $\text{poly}(U, M)$.
 \Rightarrow it takes $O(\log U + \log M)$ to encode one $h \in \mathcal{H}$.

Recall also that variance is linear for pairwise independent RVs.
For KMV, the only place that we use independence of the hash values v_i is when calculating $\text{Var}\{X\}$.

So, the proof of the analysis is complete!

Total space amounts to

$$O\left(\log\left(\frac{1}{\epsilon}\right)\left(\log U + \frac{1}{\epsilon^2} \log U\right)\right) = O\left(\frac{1}{\epsilon^2} \log\left(\frac{1}{\epsilon}\right) \log(U)\right) \text{ bits}$$

There is a better result: $O\left(\frac{1}{\epsilon^2} \log\left(\frac{1}{\epsilon}\right) + \log(U)\right)$ (Blelloch 2018)

Ex/Frequency Moment

Input: a stream a_1, \dots, a_n ($a_i \in \{1, \dots, U\}$),

Denote by f_x the # of x in the stream and $F_p = \sum_{x \in \{1, \dots, U\}} (f_x)^p$

Output: We want \tilde{F} s.t. $\mathbb{P}\{\tilde{F} \in (1 \pm \epsilon)F\} > 1 - \delta$

Note that $p=0$ is # distinct, $p=1$ is count.
for $p \geq 2$, we use AMS.

AMS

Algorithm:

Assume access to a random hash $\theta: \{1, \dots, u\} \rightarrow \{-1, 1\}$

- initialize $x \leftarrow 0$
- for i in $\{1, \dots, n\}$:
 $x \leftarrow x + \theta(a_i)$
- return x^2

Correctness/Analysis

$$\begin{aligned} \text{We have } X &= \sum_{y \in \{1, \dots, u\}} f_y \cdot \theta(y) \Rightarrow X^2 = \sum_{y_1, y_2} f_{y_1} f_{y_2} \theta(y_1) \theta(y_2) \\ &= \sum_y f_y^2 \theta(y)^2 + 2 \sum_{y_1 < y_2} f_{y_1} f_{y_2} \theta(y_1) \theta(y_2) \\ &\Rightarrow \mathbb{E}\{X^2\} = \sum_y f_y^2 = F_2 \quad \checkmark \end{aligned}$$

has $\mathbb{E}=0$

Similarly, we can find (if θ u -wise independent) \leftarrow All this in w/ lecture notes

$$\text{Var}\{X^2\} = \mathbb{E}\{X^4\} - \mathbb{E}\{X^2\}^2 \leq O(F_2^2)$$

We can mention s_1, s_2 copies of AMS, divided into $g_1 = O(\log \frac{1}{\delta})$ groups of size $s_2 = O(\frac{1}{\epsilon^2})$. The median of the group means satisfies

$$\mathbb{P}\{\text{median} \in (1 \pm \epsilon) F_2\} > 1 - \delta$$


with space $O\left(\frac{1}{\epsilon^2} \log\left(\frac{1}{\delta}\right) \log u\right)$

(Note: for $p \geq 2$, space lower bound $\Omega(n^{1-2/p})$)

Lecture 10/11 - Johnson-Lindenstrauss

We focus on *dimensionality reduction*.

Given vectors $\vec{x}_1, \dots, \vec{x}_n \in \mathbb{R}^m$, and $\epsilon > 0$, find a mapping $f: \mathbb{R}^m \rightarrow \mathbb{R}^d$ (d.c.c.m) s.t. $\forall i, j \in \{1, \dots, n\}$, $\|f(x_i) - f(x_j)\|^2 \in (1 \pm \epsilon) \|x_i - x_j\|^2$



Theorem: (Johnson-Lindenstrauss)

For any $x_1, \dots, x_n \in \mathbb{R}^m$ and any $\epsilon > 0$, there exists $f: \mathbb{R}^m \rightarrow \mathbb{R}^d$ for $d = \Omega\left(\frac{1}{\epsilon^2} \log n\right)$ s.t. $\|f(x_i) - f(x_j)\|^2 \in (1 \pm \epsilon) \|x_i - x_j\|^2 \quad \forall i, j \in [n]$.
Moreover, f is linear. $f(x) = \Pi x$
 Π is $d \times m$ matrix

Proof:

The plan is as follows:

① Find a distribution \mathcal{D} over matrices in $\mathbb{R}^{d \times m}$ s.t.

$$\forall \vec{x} \in \mathbb{R}^m, \mathbb{P}_{\Pi \sim \mathcal{D}} \left[\|\Pi \vec{x}\|^2 \in (1 \pm \epsilon) \|\vec{x}\|^2 \right] > 1 - \delta \quad \text{(\epsilon, \delta) J-L Property}$$

for $d = O\left(\frac{1}{\epsilon^2} \log n\right)$, $\delta = \frac{1}{\text{poly}(n)}$.

② Union bound!

② Starting from ①, assume we have done ①. Then, sample $\Pi \sim \mathcal{D}$, we get $f: \mathbb{R}^m \rightarrow \mathbb{R}^d$ s.t. $f(\vec{x}_i - \vec{x}_j) = f(\vec{x}_i) - f(\vec{x}_j)$. By (ϵ, δ) J-L,

$$\forall i, j \in [n], \mathbb{P}_{\Pi \sim \mathcal{D}} \left[\|f(\vec{x}_i) - f(\vec{x}_j)\|^2 \in (1 \pm \epsilon) \|\vec{x}_i - \vec{x}_j\|^2 \right] = \mathbb{P}_{\Pi \sim \mathcal{D}} \left[\|\Pi(\vec{x}_i - \vec{x}_j)\|^2 \in (1 \pm \epsilon) \|\vec{x}_i - \vec{x}_j\|^2 \right] > 1 - \delta$$

For $\delta = \frac{1}{n^2}$, we can union bound over all pairs to see that what we want happens with probability $1 - \frac{1}{n}$.

□

① There are two constructions of this distribution \mathcal{D} :

(a) $\pi_{ij} \sim \frac{1}{\sqrt{d}} \cdot \{-1, 1\}$ (b) $\pi_{ij} \sim \frac{1}{\sqrt{d}} \mathcal{N}(0, 1)$
Note: Give prob with this

Using scheme (b), fix $\tilde{x} \in \mathbb{R}^m$. Sample π as above, and let $\tilde{y} = \pi \tilde{x} \in \mathbb{R}^d$.
 Then, $\|\tilde{y}\|^2 = \sum_{i=1}^d y_i^2$, and $y_i = \sum_{j=1}^m \pi_{ij} x_j \quad \forall i$.

$$\begin{aligned} \Rightarrow \mathbb{E}_{\pi} [y_i^2] &= \mathbb{E}_{\pi} \left[\left(\sum_j \pi_{ij} x_j \right)^2 \right] = \mathbb{E}_{\pi} \left[\sum_{j_1, j_2=1}^m \pi_{ij_1} \pi_{ij_2} x_{j_1} x_{j_2} \right] \\ &= \mathbb{E}_{\pi} \left[\sum_{j=1}^m \pi_{ij}^2 x_j^2 \right] + \mathbb{E}_{\pi} \left[2 \sum_{j_1 < j_2} \pi_{ij_1} \pi_{ij_2} x_{j_1} x_{j_2} \right] \\ &= \frac{1}{d} \sum_{j=1}^m x_j^2 = \frac{\|\tilde{x}\|^2}{d} \end{aligned}$$

Note: 0 in expectation because independent

$\Rightarrow \mathbb{E} [\|\tilde{y}\|^2] = d \cdot \frac{\|\tilde{x}\|^2}{d} = \|\tilde{x}\|^2$. So, \mathcal{D} behaves well in expectation.
 We want to show $\sum y_i^2$ concentrates.

We can say $\mathbb{P} \left[\sum y_i^2 > (1+\epsilon) \|\tilde{x}\|^2 \right] = \mathbb{P} \left[e^{t \sum y_i^2} > e^{t(1+\epsilon) \|\tilde{x}\|^2} \right]$
 Set $t = \frac{\epsilon m}{8 \|\tilde{x}\|^2}$
 Set $d = \frac{C}{\epsilon^2} \log\left(\frac{1}{\delta}\right)$

$$\begin{aligned} &\stackrel{\text{markov}}{\leq} \frac{\mathbb{E} \left[e^{t \sum y_i^2} \right]}{e^{t(1+\epsilon) \|\tilde{x}\|^2}} = \frac{1}{e^{t(1+\epsilon) \|\tilde{x}\|^2}} \cdot \left(\frac{1}{1 - 2t \|\tilde{x}\|^2 / d} \right)^{d/2} \\ &\leq e^{-6\epsilon^2 d} \end{aligned}$$

Then $\mathbb{P} \left[\sum y_i^2 > (1+\epsilon) \|\tilde{x}\|^2 \right] \leq \delta$.

We can perform similar bounds on the lower tail. This yields

$$\mathbb{P}_{\pi \sim \mathcal{D}} \left[\|\pi \tilde{x}\|^2 \in (1 \pm \epsilon) \|\tilde{x}\|^2 \right] > 1 - \delta \quad \text{as desired.}$$

□

This reduces to dimension $d = \frac{1}{\epsilon^2} \log(\text{poly } m)$, but takes $O(mn)$ time to transform each vector $\tilde{x} \in \mathbb{R}^m$.

We can do better :-

Two strategies to speed up $\Pi \tilde{x}$:

- ① use a sparse matrix Π (sparse JL transform)
- better for sparse \tilde{x}

Consider random matrix Π . Fix parameter s .

- sample exactly s entries randomly in every column of Π to be nonzero
- fill all selected nonzero entries with random $\pm \frac{1}{\sqrt{s}}$

Theorem (KN, 2014)

$\exists c_1, c_2 > 0$ s.t., if we set $d = c_1 \cdot \frac{1}{\epsilon^2} \log(\frac{1}{\delta})$, $s = c_2 \epsilon d = \frac{c_1 c_2}{\epsilon} \log(\frac{1}{\delta})$
then $\forall \tilde{x} \in \mathbb{R}^m$, $\mathbb{P}[\|\Pi \tilde{x}\|^2 \in (1 \pm \epsilon) \|\tilde{x}\|^2] > 1 - \delta$

- ② use a structured matrix Π that allows for (fast JL transform)
fast matrix-vector multiplication.
- better for average \tilde{x}

Let Π be a product of 3 matrices, each with fast multiplication. In particular,

$$\Pi = \frac{1}{\sqrt{d}} S \cdot H \cdot D \quad (\text{assume } m \text{ is a power of } d).$$

$\begin{matrix} \swarrow & \searrow \\ d \times m & m \times m \end{matrix}$

- S is a random variable, where $S \tilde{x}$ picks d random coordinates of \tilde{x} to form a vector in \mathbb{R}^d .

Lemma: If $\|\tilde{x}\|_\infty$ is small, then $\mathbb{P}[\|S \tilde{x}\|^2 \in (1 \pm \epsilon) \|\tilde{x}\|^2]$ is large.

So, we want $H \cdot D$ to preprocess \tilde{x} to maintain the norm, but have small $\|H \cdot D \tilde{x}\|_\infty$.

- H is a deterministic Hadamard matrix $H_{2^k k} = \begin{bmatrix} H_{2^k k-1} & H_{2^k k-1} \\ H_{2^k k-1} & -H_{2^k k-1} \end{bmatrix}$, $H_0 = [1]$
 $\forall \tilde{x} \in \mathbb{R}^m$, $H \tilde{x}$ can be computed in $O(m \log m)$

- D is a randomly diagonal matrix that randomly negates coordinates.

Dx can be computed in $O(n)$ time.

$$D = \begin{bmatrix} \pm 1 & & 0 \\ & \pm 1 & \dots \\ 0 & \dots & \pm 1 \end{bmatrix}$$

We know that both $\frac{1}{\sqrt{n}}H$ and D are unitary, preserving the norm. There is a nontrivial lemma

Lemma: $\forall x \in \mathbb{R}^m$, $P[\| \frac{1}{\sqrt{n}}HDx \|_\infty \text{ is "small"}]$ is "large".

← sort of a rotation

This yields that $TH = SHD$ has the same properties, but can be multiplied in $O(m \log m)$ time.

Insert speaker notes here


Lecture 10/25 - Learning from Experts

no distributional assumptions, can be adversarial

Consider a sequence of events $E_1, \dots, E_T \in \{0, 1\}$ where each event E_t 's outcome is revealed at time t .

There are also n experts, each one predicting E_t before time t .

The goal is to predict events before they happen, minimizing # of mistakes



We will show that, without knowing which one is the best expert, we can also make about m^* mistakes as well.

Warning If $m^* = 0$, we follow the majority advice among all experts that haven't made a mistake yet at step t . There are 2 cases

- ① the majority is correct
 - ② the majority is incorrect, and so we reduce the # of experts we follow by a factor ≥ 2
- ② can only happen $\log n$ times, and so we make $\leq \log n$ mistakes.

Weighted Majority

Init: Fix parameter $\gamma \in (0, \frac{1}{2}]$, give weight $w_i = 1$ to expert i

For $t \in [T]$:

- follow the weighted majority of all experts
- for all incorrect experts, $w_i \leftarrow w_i(1-\gamma)$

Theorem: the # of mistakes M is at most $2 \cdot (1+\gamma)^{m^*} + \frac{2 \ln n}{\gamma}$

Proof: Denote by $w_i^{(t)}$ the weight of expert i at time t .
 Let $W^{(t)} = \sum_{i=1}^n w_i^{(t)}$ be the total weight.

Every time we make a mistake,

$$W^{(t+1)} = \sum_{i=1}^n w_i^{(t+1)} = \sum_{i \text{ correct}} w_i^{(t)} + (1-\gamma) \sum_{i \text{ incorrect}} w_i^{(t)} = W^{(t)} - \gamma \sum_{i \text{ incorrect}} w_i^{(t)}$$

$\geq W^{(t)} - \gamma \sum_{i \text{ incorrect}} w_i^{(t)}$ because weighted majority

$$\leq W^{(t)} - \gamma \frac{W^{(t)}}{2} = \left(1 - \frac{\gamma}{2}\right) W^{(t)}$$

The best expert i_* has $w_{i_*}^{(T)} = (1-\gamma)^{m^*}$

The final total weight is

$$W^{(T)} \leq \left(1 - \frac{\gamma}{2}\right)^M W^{(0)} = n \left(1 - \frac{\gamma}{2}\right)^M \geq (1-\gamma)^{m^*}$$

$$\Rightarrow (1-\gamma)^{m^*} \leq n \left(1 - \frac{\gamma}{2}\right)^M \Rightarrow m^* \ln\left(\frac{1}{1-\gamma}\right) \geq \ln n + M \ln\left(\frac{1}{1-\frac{\gamma}{2}}\right)$$

Since $|\gamma| \leq \frac{1}{2}$, $3 \leq \ln\left(\frac{1}{1-\gamma}\right) \leq 3 + \gamma^2$

$$\Rightarrow (3 + \gamma^2) m^* \geq \ln n + \frac{\gamma}{2} M \Rightarrow M \leq 2(1+\gamma) m^* + \frac{2 \ln n}{\gamma} \quad \square$$

Randomized Weighted Majority

The same idea, but in each round, we return $b \in \{0, 1\}$ w.p. $\sum_{i \text{ predicts } b} \frac{w_i^{(t)}}{W^{(t)}}$

Theorem: The randomized version makes at most

$$\mathbb{E}\{M\} \leq (1+\gamma) m^* + \frac{\ln n}{\gamma}$$

mistakes in expectation.

Proof: Denote by $q^{(t)}$ the probability that we make a mistake in step t . So, $q^{(t)} = \sum_{\text{incorrect}} \frac{w_i^{(t)}}{W^{(t)}}$. Then,

$$W^{(t+1)} = \sum_i w_i^{(t+1)} = \sum_{\text{correct}} w_i^{(t)} + (1-\zeta) \sum_{\text{incorrect}} w_i^{(t)} = W - \zeta \sum_{\text{incorrect}} w_i^{(t)}$$

$$= W^{(t)} - \zeta q^{(t)} W^{(t)} = (1 - \zeta q^{(t)}) W^{(t)}$$

So, the final total weight is

$$W^{(T)} = n \prod_{t=1}^T (1 - \zeta q^{(t)}) \leq n \prod_{t=1}^T e^{-\zeta q^{(t)}} = n e^{-\zeta \sum_{t=1}^T q^{(t)}} = n e^{-\zeta \mathbb{E}\{M\}}$$

Also, as before,

$$(1-\zeta)^{m^*} \leq W^{(T)} \Rightarrow (1-\zeta)^{m^*} \leq n e^{-\zeta \mathbb{E}\{M\}}$$

$$\Rightarrow m^* (\zeta + \zeta^2) \geq -\ln n + \zeta \mathbb{E}\{M\} \Rightarrow \mathbb{E}\{M\} \leq (1+\zeta) m^* + \frac{\ln n}{\zeta}$$

□

Multiplicative Weights

In the general setting, there are T rounds.

- each round has n choices $\{1, \dots, n\}$, and we choose one.
- there is a cost $m_i^{(t)} \in [1, 1]$ for choosing i in round t .
- We wish to minimize total cost.

Init: Fix parameter $\zeta \in (0, \frac{1}{2}]$ and give weight $w_i = 1$ to each choice.

For $t \in [T]$:

- return i w.p. proportional to w_i
- observe costs $\{m_i^{(t)}\}_{i \in [n]}$
- update $w_i \leftarrow w_i (1 - \zeta m_i^{(t)})$

Theorem: For every $i \in [n]$, the expected total cost is at most

$$\mathbb{E}\{M\} \leq \sum_{t=1}^T m_i^{(t)} + \zeta \sum_{t=1}^T |m_i^{(t)}| + \frac{\ln n}{\zeta}$$

\uparrow
 $\leq T \cdot \frac{1}{\zeta}$
 $m_i^{(t)} \in [1, 1]$ w.p. ζ

If we set $z = \sqrt{\frac{\ln n}{T}}$, we make $O(\sqrt{T \ln n})$ more mistakes than the expert in question.

Lecture 10/27 - Online Algorithms

Ski Rental

Every day that you ski, you can either:

- (1) use skis you already bought
- (2) rent skis for R
- (3) buy skis for B

* An important part of the model is that you don't know, until it happens, whether you plan to ski.

On day 1, you go skiing and must decide.

After day i , you may never ski again, or you go skiing on day $i+1$.

We measure the result using the **competitive ratio**: $\max_{\text{all inputs}} \left\{ \frac{\text{your cost}(\text{input})}{\text{OPT of}(\text{input})} \right\}$

↑
input = $D \in \mathbb{N}$
the number of days
I will ski:

The offline problem has an $\text{OPT}(\text{input } D) = \min\{DR, B\}$

We wish to design an online algorithm that does well under the competitive ratio metric.

Any deterministic online algorithm is fully defined by T , the number of days we are willing to rent (rent $\forall t \leq T$, buy on $t = T+1$, $t \in [D]$)

Claim: For any algorithm T , the competitive ratio is achieved at $D = T+1$.

Proof: For fixed T, D , get $\frac{R \min\{D, T\} + B \mathbb{1}_{D > T}}{\min\{DR, B\}}$ ← price we pay
← opt

(1) max cannot be achieved at $D > T+1$, since numerator doesn't change and denominator may grow.

(2) max cannot be achieved at $D < T$, since it is always ≤ 1 .

(3) $D = T$ is $\leq D = T+1$, using marginal logic.

(worst case is stopping skiing right after buying)

□

So, for any T , the competitive ratio is $\frac{RT+B}{\min\{R(T+1), B\}}$.

Claim: This is minimized at $T = \frac{B}{R} - 1$ (assuming $B/R \in \mathbb{N}$), yielding a competitive ratio of $2 - \frac{R}{B}$.

Proof:

- (1) the min is not achieved at $T > \frac{B}{R} - 1$; the denominator is constant while the numerator increases
- (2) the min is not achieved at $T < \frac{B}{R} - 1$; the denominator and numerator both grow by R , and the numerator is larger than the denominator, so the competitive ratio decreases for each additional T .

List Update -

You manage a linked list. Online, you get requests to access x . You scan the list until you hit x . You are allowed to move x up in the list however much for free after returning.

Frequency Count -

- (1) Initialize $C(x) = 0 \forall x$
- (2) If x is queried, increment $C(x)$
- (3) move x up above all y with $C(x) > C(y)$

Claim: FC has competitive ratio $\Omega(n)$

Proof: Start by querying element i times. Then, for some large K , for $i \in [K]$:
for $j \in [n]$:
query j n times. ↙ query back of the list n times

The offline optimum is, for each new query, move to front. This has total order $O(Kn^2)$

FC will pay $n + (n-1) + \dots = O(n^2)$ for each time we query j n times.
 So, FC has cost $\Omega(Kn^2) \Rightarrow C.R. = \Omega(K)$ □

Move to front

Every time you query something, move it to the front.

Theorem: MTF has $C.R. \leq 2$ ← MTF fails when querying in a cycle, which also causes offline optimum to suck.

Proof: Imagine running MTF and OPT side-by-side. ← offline opt \forall times t , denote by $\Phi(t)$ the # of pairs (x, y) s.t. $x \succ_{\text{MTF}} y$ but $y \succ_{\text{OPT}} x$.
 We can see that

(1) $\Phi(0) = 0$ (2) $\Phi(t) \geq 0$. Let $\text{MTF}(t), \text{OPT}(t)$ be the costs for query t .

Claim: $\forall t, \text{MTF}(t) + (\Phi(t) - \Phi(t-1)) \leq 2\text{OPT}(t)$

Proof: Consider accessing x @ time t . Let $\text{MTF}(x) = p$.

Suppose that k elements in front of x in MTF are also ahead of x in OPT. $\Rightarrow \text{MTF}(t) = p, \text{OPT}(t) \geq k + 1$ ← the idea is if $k=p$, we did good; if $k < p$, we fix more inversions and help Φ

The MTF operation creates k inversions, but fixes $p-k$ inversions, if we were to not change OPT. Moving x forward in OPT can only improve things, since it can only fix inversions by agreeing that x is ahead of things. So,

$$\Phi(t) - \Phi(t-1) \leq 2k - p + 1 \leq 2\text{OPT}(t) - \text{MTF}(t)$$
 □

Repeated application of the claim shows

$$\text{MTF} + \Phi \leq 2\text{OPT} \Rightarrow C.R. \leq 2.$$

□

Lecture 11/8 - Communication Complexity

Def: A **two-party communication problem** consists of a function $f: \{0,1\}^a \times \{0,1\}^b \rightarrow \{0,1\}$. Alice receives input $A \in \{0,1\}^a$ and Bob receives $B \in \{0,1\}^b$. The goal is to compute $f(A,B)$.

Def: A **deterministic communication protocol** specifies for Alice as a function of her input A and all previous messages $a_1, b_1, a_2, b_2, \dots, a_k, b_k$, what is the next message a_{k+1} Alice should send? Similarly for Bob.

Def: The **communication cost** is the maximum # of bits in all messages.

Ex/Equality $f(x,y) = 1$ iff $x=y$

Protocol: In each message i , Alice sends x_i and Bob sends y_i
 $\Rightarrow O(n)$ cost

Insert rest of lecture
here

Lecture 11/10 - Computation of Nash

Consider rock-paper-scissors

	R	P	S
R	0	1	-1
P	-1	0	1
S	1	-1	0

- This is a zero-sum game

- Consider the matrix of payoffs for the row player & the matrix for the column player. The rank of the sum of these matrices determines tractability of computation of Nash.

Recall: • Nash Equilibrium is when both players are best responding to each other.
• Can be pure or mixed NE

We generalize: A 2 player game is given by two num matrices A, B where A_{ij} denotes the payoff to A if row plays i and col plays j , and $B_{ij} \dots$

mixed strategy distribution over actions

(\vec{x}, \vec{y}) is a Nash Eq. if

$$\vec{x} A \vec{y} \geq A_{i \cdot} \vec{y} \quad \forall i \quad \& \quad \vec{x} B \vec{y} \geq \vec{x} B_{\cdot j} \quad \forall j$$

Def: (\vec{x}, \vec{y}) are ϵ -Nash Equilibrium if (Almost Nash)

$$\epsilon + \vec{x} A \vec{y} \geq A_{i \cdot} \vec{y} \quad \forall i \quad \& \quad \epsilon + \vec{x} B \vec{y} \geq \vec{x} B_{\cdot j} \quad \forall j$$

Computation of Nash

Given A, B , find a NE. This problem is PPAD-Complete.
We can find ϵ -Nash or via LP rounding, etc.

define this later

Lipton/Motakis/Mehta

set that can
have repeat
elements

Theorem: There exist two multi-sets S, T , each of size $O(\log n / \epsilon^2)$ s.t. it is an ϵ -Nash for A, B to randomly sample strategies uniformly from S, T , respectively.

This implies a brute force algorithm to exhaust all $n^{O(\frac{\log n}{\epsilon^2})}$ pairs of multisets and check if any is ϵ -Nash.

quasi-polynomial time $O(\log n)$

Side note: If $\forall \epsilon, \exists$ a $o(n \log n)$ time algorithm for finding ϵ -Nash, (Roberson) there exists a $2^{o(n)}$ algorithm for PPA.

(If you can do better than the LHM alg. above, you can do sub-exponential PPA!)

Proof of Theorem:

Let (\bar{x}, \bar{y}) be a NE (one must exist). Consider randomly sampling k strategies from \bar{x} (call it S) and k strategies from \bar{y} (call it T).

Define $x_i^* \equiv \frac{\# \text{ of times } i \in S}{k}$ $y_i^* \equiv \frac{\# \text{ times } i \in T}{k}$ (Use empirical distribution gotten by sampling from Nash)

We want to show:

(1) $\forall i, |A_{i \cdot} \cdot \bar{y} - A_{i \cdot} \cdot \bar{y}^*| < \epsilon$

(2) $\forall i, |B_{i \cdot} \cdot \bar{x} - B_{i \cdot} \cdot \bar{x}^*| < \epsilon$

(3) $|\bar{x} A_{i \cdot} - \bar{x}^* A_{i \cdot}| < \epsilon$

(4) $|\bar{x} B_{i \cdot} - \bar{x}^* B_{i \cdot}| < \epsilon$

Chernoff
City:
Population Weisberg

From this, we want to show that $\bar{x}^* A_{i \cdot} \bar{y}^* \geq A_{i \cdot} \bar{y} - 3\epsilon \forall i$

(1) gives $\bar{x}^* A_{i \cdot} \bar{y}^* \geq \bar{x}^* A_{i \cdot} \bar{y} - \epsilon$ (each row is ϵ -accurate and so is a distribution over rows)

(3) then gives $\geq \bar{x} A_{i \cdot} \bar{y} - 2\epsilon$

N.E. then gives $\geq A_{i \cdot} \bar{y} - 2\epsilon \quad \forall i$

(1) again gives $\geq A_{i \cdot} \bar{y}^* - 3\epsilon \quad \forall i$

□

Exponential Time Alg. For Exact Nash

(1) Assume WOLOG that $A = B^T$ (we can reduce anything to this by swapping players for half the actions)
This will look like playing against ourselves?

Consider the following: $(\vec{A}_{i\cdot}, \vec{x}) \leq 1 \quad \forall i$ (i doesn't give payoff more than $\frac{1}{|X_i|}$ against $\frac{\vec{x}}{|X_i|}$)
(LH polytope) $x_i \geq 0 \quad \forall i$ (i has pos. entries and is normalizable)

Being in this polytope means no strategy does better than $\frac{\vec{x}}{|X_i|}$.

(2) We call an action i **covered** if $\langle \vec{A}_{i\cdot}, \vec{x} \rangle = 0$ or $x_i = 0$ or both.

We claim: if \vec{x} satisfies LHP and has all i covered (at least one incr. is tight),
then $\frac{\vec{x}}{|X_i|}$ is Nash.

Proof: Consider using i against $\frac{\vec{x}}{|X_i|}$. If $x_i = 0$, not used and we don't care. Otherwise, i covered $\Rightarrow \langle \vec{A}_{i\cdot}, \frac{\vec{x}}{|X_i|} \rangle = \frac{1}{|X_i|} \Rightarrow i$ is a BR \square

The Algorithm: (Pivoting) (also a proof that N.E. exists)

Start from a vertex of the polytope, "walk" along the boundary (keeping all but one constraint tight) until the next vertex (pivoting).

Start at $\vec{0}$.

"relax" $x_i = 0$, see which constraints tighten to get next \vec{x} (can be done w/ LP solver)

If \vec{x} covers all i , done!

If not, $\exists i$ that is "double-covered", relax one and continue.

\square

See lecture notes
for details

Note that any decision problem (is there a Nash st. ...) is NP-Hard, but no such problem implies that no Nash exists. So, Nash = NP-Hard

PPAD-Complete Examples

Given a graph on 2^n nodes s.t. every node has indegree ≤ 1 and outdegree ≤ 1 with a source node (indegree 0), find a sink (we can nonconstructively prove a sink exists)

Given $|S| = 2^n$, $|T| = 2^{n-1}$, there exists a function $f: S \rightarrow T$ that maps s_1, s_2 to the same t .

Lecture 11/15 - Low-Rank Approx.

Let $\vec{a}_1, \dots, \vec{a}_n \in \mathbb{R}^d$ be data points. We seek $b_1, \dots, b_k \in \mathbb{R}^d$ ($k \ll d$) and $\{c_{ji}\}_{j \in [k], i \in [n]}$ s.t. $\vec{a}_i \approx \sum_{j=1}^k c_{ji} \vec{b}_j$ (approximately in low dimensional subspace)

Equivalently, let $A = \begin{pmatrix} | & & | \\ \mathbf{a}_1 & \dots & \mathbf{a}_n \\ | & & | \end{pmatrix} \in \mathbb{R}^{d \times n}$. We seek $B = \begin{pmatrix} | & & | \\ \mathbf{b}_1 & \dots & \mathbf{b}_k \\ | & & | \end{pmatrix} \in \mathbb{R}^{d \times k}$

and $C \in \mathbb{R}^{k \times n}$ s.t. $A \approx BC$.

We wish to minimize the following for a given matrix $A \in \mathbb{R}^{d \times n}$ and $k \ll d$:

minimize error,
 $B \in \mathbb{R}^{d \times k}, C \in \mathbb{R}^{k \times n}$

$$\text{error} = \sum_{i=1}^n \|\vec{a}_i - \sum_{j=1}^k c_{ji} \vec{b}_j\|_2^2 = \|A - BC\|_F^2$$

Frobenius norm
 $\|M\|_F^2 = \sum_{ij} |M_{ij}|^2$

SVD

Theorem: (SVD exists)

Let $A \in \mathbb{R}^{d \times n}$ be a matrix. Let $r = \min(d, n)$. Then, there exist matrices U, Σ, V^T

- $U \in \mathbb{R}^{d \times r}$: the columns of U (left singular vectors) are orthonormal
- $\Sigma \in \mathbb{R}^{r \times r}$; Σ is diag($\sigma_1, \dots, \sigma_r$) s.t. the singular values have $\sigma_1 \geq \dots \geq \sigma_r \geq 0$
- $V \in \mathbb{R}^{n \times r}$; the columns of V (right singular vectors) are orthonormal

$$\star U \Sigma V^T = A \star$$

This leads to some interesting properties:

- $A^T A = V \Sigma^2 V^T$
- Singular values are square roots of eigenvalues
- If \vec{v}_i column of V , $A \vec{v}_i = U \Sigma (V^T \vec{v}_i) = U \Sigma \begin{pmatrix} 1 \\ 0 \\ \vdots \end{pmatrix} = \sigma_i \vec{u}_i$

Theorem: (SVD is best)

For any $k \in [1, r]$, let U_k be the matrix of size $d \times k$ consisting of the first k columns of U . Let $V_k \in \mathbb{R}^{n \times k}$ and $\Sigma_k \in \mathbb{R}^{k \times k}$ be defined similarly. Then,

$$\|A - U_k \Sigma_k V_k^T\|_F^2 = \min_{\substack{B \in \mathbb{R}^{d \times k} \\ C \in \mathbb{R}^{k \times n}}} \|A - BC\|_F^2$$

Proof:

(k=1) Consider the case $k=1$. We seek $\vec{b} \in \mathbb{R}^d$, $\vec{c} \in \mathbb{R}^n$ s.t.


$$\|A - \vec{b} \vec{c}^T\|_F^2 = \sum_{i=1}^n \|a_i - c_i \vec{b}\|_2^2 \text{ is minimized.}$$

For any given \vec{b} , (suppose WOLOG that $\|\vec{b}\|_2 = 1$, since we can scale down \vec{b} and scale up c_i) this is minimized for the i th term $\|a_i - c_i \vec{b}\|_2^2$ when $c_i = \langle \vec{a}_i, \vec{b} \rangle$.

The minimum is then $\|a_i - c_i \vec{b}\|_2^2 = \|\vec{a}_i\|_2^2 - \|c_i \vec{b}\|_2^2 = \|\vec{a}_i\|_2^2 - c_i^2$

So, we wish to maximize

$$\max \sum c_i^2 = \sum \langle \vec{a}_i, \vec{b} \rangle^2 = \|A^T \vec{b}\|_2^2$$



For given A , we find unit vector \vec{b} maximizing $\|A^T \vec{b}\|_2^2$, and set $\vec{c} = A^T \vec{b}$.

$$\text{Let } A = U \Sigma V^T \Rightarrow \|A^T \vec{b}\|_2^2 = \|V \Sigma U^T \vec{b}\|_2^2 = \sum_{i=1}^n \sigma_i^2 \langle \vec{u}_i, \vec{b} \rangle^2$$

V orthogonal columns
 $\forall \vec{v}_i, \|\vec{v}_i\|_2 = 1$

Since $\|\vec{b}\|_2 = 1$, $\sum \langle \vec{u}_i, \vec{b} \rangle^2 = 1$ since $\{\vec{u}_i\}$ orthonormal

So, we maximize when $\vec{b} = \vec{u}_1$, ($\langle \vec{u}_1, \vec{b} \rangle = 1$ and \vec{u}_1 has largest singular value σ_1)

Therefore, the error is minimized for $\vec{b} = \vec{u}_1$, $\vec{c} = A^T \vec{u}_1 = V \Sigma U^T \vec{u}_1 = \sigma_1 \vec{v}_1$.

The claim holds for $k=1$.

(k>1) We do the same thing. For any given $B \in \mathbb{R}^{d \times k}$ with orthogonal columns, the best C is $C = B^T A$. We seek the B maximizing

$$\|B^T A\|_F^2 = \|B^T U \Sigma V^T\|_F^2 = \|B^T U \Sigma\|_F^2 = \sum_{i=1}^k \sigma_i^2 \|B^T \vec{u}_i\|_2^2 \leq \sum_{i=1}^k \sigma_i^2$$

Since B has orthogonal columns, $\sum_{i=1}^k \|B^T \vec{u}_i\|_2^2 \leq k$, $\|B^T \vec{u}_i\|_2 \leq 1$.

So, we want $\|B^T \vec{u}_i\|_2 = 1 \quad \forall i \Rightarrow B = \begin{pmatrix} | & & | \\ \vec{u}_1 & \dots & \vec{u}_k \\ | & & | \end{pmatrix} = U_k$. Thus, $C = B^T A = B^T U \Sigma V^T = \Sigma_k V_k^T$

□

Algorithm (SVD Solver)

- Initialize $A^{(0)} = A$
- For $i=1, \dots, r$:
 - compute the optimal rank 1 approximation of $A^{(i)}$ (Find $\vec{b}^{(i)} \in \mathbb{R}^d$, $\vec{c}^{(i)} \in \mathbb{R}^n$ s.t. $\|A^{(i)} - \vec{b}^{(i)} \vec{c}^{(i)T}\|_F^2$ minimized)
 - update $A^{(i+1)} \leftarrow A^{(i)} - \vec{b}^{(i)} \vec{c}^{(i)T}$, $\vec{u}_i = \frac{\vec{b}^{(i)}}{\|\vec{b}^{(i)}\|_2}$, $\vec{v}_i = \frac{\vec{c}^{(i)}}{\|\vec{c}^{(i)}\|_2}$, $\sigma_i = \|\vec{b}^{(i)}\|_2 \|\vec{c}^{(i)}\|_2$
- Set $U = (\vec{u}_1, \dots, \vec{u}_r)$, $V = (\vec{v}_1, \dots, \vec{v}_r)$, $\Sigma = \text{diag}(\sigma_1, \dots, \sigma_r)$.

We need to show that $A^{(r+1)} = 0$ and that $\{\vec{u}_i\}_i$ and $\{\vec{v}_i\}_i$ are both orthonormal.

Claim: For any round i with $\vec{b} = \vec{b}^{(i)}$, $\vec{c} = \vec{c}^{(i)}$, $A = A^{(i)}$,

- ① $\vec{b} \in \text{column space of } A$
- ② $\vec{b} \perp \text{column space of } (A - \vec{b} \vec{c}^T)$

This claim (if we were to prove it) shows

$$\vec{b}_i \in \text{span}\{A_i\} \text{ and } \vec{b}_i \perp \text{span}\{A_{i+1}\} \Rightarrow \vec{b}_i \perp \vec{b}_j \text{ for } i \neq j.$$

and $\text{span}\{A_{i+1}\} \subseteq \text{span}\{A_i\} \Rightarrow \text{dim reduces by 1 each round.}$

These two results show U is orthonormal and $A^{(r+1)} = 0$.

□

We need to fill in one piece: finding the best rank 1 approximation for A . We will use the **Power Method**. The idea is we wish to find the top eigenvalue of $A^T A$. We keep multiplying a vector by $A^T A$, which will push it more in the direction of the top eigenvector of $A^T A$ (or $A A^T$, same spectrum).

Power Method

- initialize \vec{z} to random vector with i.i.d $\mathcal{N}(0,1)$ entries. Set $\vec{z}_1 = \frac{\vec{z}}{\|\vec{z}\|_2}$.
- For $t=1, \dots, T$:
 - set $\vec{z}_{t+1} \leftarrow A A^T \vec{z}_t$. Normalize $\vec{z}_{t+1} \leftarrow \frac{\vec{z}_{t+1}}{\|\vec{z}_{t+1}\|_2}$
- Return \vec{z}_T as \vec{b} . Return $\vec{c} = A^T \vec{b}$.

This works because

$$\vec{z}_1 = \sum_i \alpha_i \vec{u}_i \Rightarrow \vec{z}_2 \propto A A^T \vec{z}_1 = \sum_i \sigma_i^2 \alpha_i \vec{u}_i \Rightarrow \dots \Rightarrow \vec{z}_r = \sum_i (\sigma_i^2)^T \alpha_i \vec{u}_i$$

↑ largest eigenvalue will dominate

If we set $T = O\left(\frac{\log d}{\epsilon}\right)$, we have

$$\|A - \hat{b} (A^T \hat{b})^T\|_F^2 \leq (1 + \epsilon) \|A - \sigma_1 \vec{u}_1 \vec{v}_1^T\|_F^2$$

top-1 SVD

The total time to find \hat{b}, \hat{c} is $O\left(\frac{\log d \cdot nd}{\epsilon}\right)$

Therefore the total time to find the K -rank SVD approximation

is

$$\boxed{O\left(\frac{K \cdot nd \cdot \log d}{\epsilon}\right)}$$

Lecture 11/29 - Static Data Structure Lower Bound

Polynomial Evaluation

Given a polynomial $P \in \mathbb{F}_q[x]$ of degree n , where q is prime, we would like to preprocess it into a data structure of size S s.t. given a query $x \in \mathbb{F}_q$, $P(x)$ can be computed in time T

We focus on minimizing S and T .

There are two trivial structures:


- ① Store coefficients of P $O(n)$ space + $O(n)$ query time
- ② Precompute $P(x)$ for \mathbb{F}_q $O(q)$ space + $O(1)$ query time

A nontrivial result from [Kellaya, Umrigar '08] is that $\forall \delta > 0$, we can achieve $S = O(n^{1+\delta} \text{polylog } q)$, $T = O(\text{polylog}(n, q))$

Today we will prove lower bounds on (S, T) !

Def: [Yao '81] The **Cell-Probe Model** for data structure analysis is

Memory of size S :



- ① Cells are indexed by $[S]$
- ② Can read/write a cell in unit time *read/writes*
- ③ Computation is free $\Rightarrow T := \#$ of cell probes the algorithm makes

Since this model is stronger than actual computers, lower bounds here apply everywhere!

We make the usual assumption $w \geq \Omega(\log n + \log q)$ (can store pointers and elements of \mathbb{F}_a)

An interesting setting we will focus on is when $w = \Theta(\log n)$ and $q = \text{poly } n$.

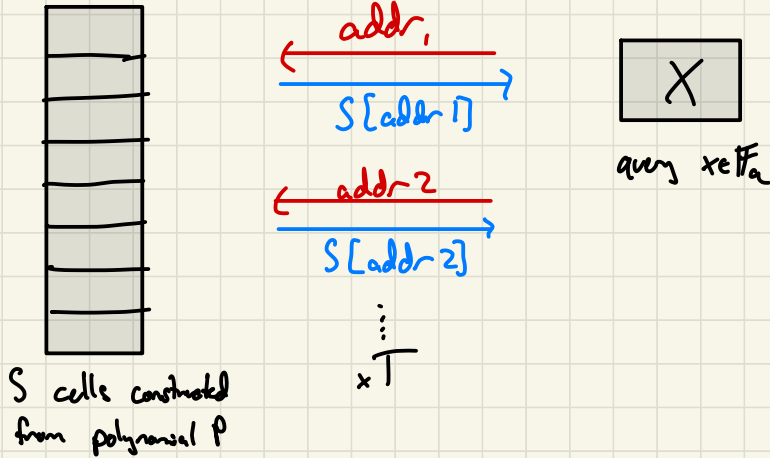
Reduction from Communication Complexity Problem [Miltersen, Nisan, ... '94]

Recall: Alice gets input X and Bob gets input Y , and the goal is to compute a function $f(X, Y)$ using minimal communication.

We can view polynomial evaluation as a communication game as follows:

Alice knows the polynomial

Bob knows the query



Memory accesses are a 2-way communication! Formally

Lemma: Suppose that \exists a data structure for polynomial evaluation w/ space S and query time T . Then, there is a protocol for the following communication protocol:

- Alice gets P . Bob gets x . The goal is to send $P(x)$ s.t. Bob sends $T \cdot \log S$ and Alice sends $T \log$ bits.

Proof: Alice preprocesses P into a data structure of size S locally. Then, Bob & Alice simulate the query alg:

For $t \in [T]$:

- Bob sends an address in $\log S$ bits
- Alice responds with the cell contents in w bits

□

So, a lower bound in this communication problem is a lower bound for PE. Note that this reasoning works for any data structure in the cell-probe model!

Claim: To compute $P(x)$, for any $c \in [0, \min\{\log n, \log q\}]$,

- Bob sends $\geq 2^c \log q$ bits
- Alice sends $\geq \log q - c$ bits

← Communication lower bound

Proof: omitted ☺

This yields that

$$\forall c \text{ s.t. } Tw < 2^c \log q, \quad T \cdot \log S \geq \log q - c$$
$$\Rightarrow S^T \geq \frac{q}{2^c} \quad \forall c \text{ s.t. } 2^c > \frac{Tw}{\log q}$$

$$\Rightarrow S^T \geq \frac{q \log q}{Tw} \Rightarrow S \geq \left(\frac{q \log q}{Tw} \right)^{1/T}$$

When $T=1$, this means $S \geq \frac{q \log q}{w}$ ← # bits to store q elements of \mathbb{F}_q
(which is the second) ← # bits/cell
trivial solution!


Cell-Sampling

Suppose now that T is a large constant s.t. $\left(\frac{q \log q}{Tw} \right)^{1/T} < n$

We focus on when $w = \log q$ and $q = \text{polyn}$ for convenience.

each cell stores one element of \mathbb{F}_q

The idea is to find a small set of cells C s.t. too many queries can be answered by accessing C .



We want to produce contradictions by being able to reconstruct the polynomial too easily.

For $T=1$, \exists one cell that is accessed by q_s different queries.
In general,

Lemma: Let $\epsilon > 0$. Suppose that \exists a data structure for polynomial evaluation w. space S and time $T(\epsilon n)$. Then, there exists a set C of ϵn cells st. $\geq q \left(\frac{\epsilon n}{S}\right)^{O(T)}$ different queries can be answered by only accessing the cells in C .

Proof: We will do this by a probabilistic argument.

1. Sample a random set C of ϵn cells.
2. Fix a query $x \in \mathbb{F}^n$.

$$\begin{aligned} \text{Then, } \mathbb{P}_C \left\{ x \text{ can be answered by accesses of only } C \right\} &= \frac{\binom{S-T}{|C|-T}}{\binom{S}{|C|}} = \frac{S-T!}{S!} \cdot \frac{|C|!}{(|C|-T)!} \\ &= \frac{|C| \cdot (|C|-1) \cdot \dots \cdot (|C|-T+1)}{S \cdot (S-1) \cdot \dots \cdot (S-T+1)} \geq \left(\frac{|C|-T}{S}\right)^T \geq \left(\frac{\epsilon n}{S}\right)^{O(T)} \quad (T \ll \epsilon n) \end{aligned}$$

3. By linearity, $\mathbb{E} \{ \#x \text{ that can be answered within } C \} \geq q \left(\frac{\epsilon n}{S}\right)^{O(T)}$.
So, there exist this many queries for some set C .

□

So, this setup would allow us to answer enough queries to reconstruct an n -degree polynomial with ϵn cells if $q \left(\frac{\epsilon n}{S}\right)^{O(T)} > n$. Therefore,

Theorem: We must have $q \left(\frac{\epsilon n}{S}\right)^{O(T)} \leq n \Leftrightarrow T \geq \Omega \left(\frac{\log(a/n)}{\log(S/n)} \right)$

if $S = O(n)$, $T \geq \Omega \left(\frac{\log(a/n)}{\log(S/n)} \right) = \Omega(\log n)$
if $T = O(1)$, $S \geq n^{1/S \cdot O(T)}$

Proof: Suppose BWOC that $q \left(\frac{\epsilon n}{S}\right)^{O(T)} \geq n+1$.

encode

1. Construct the data structure and find the set C with the claimed property (from the lemma).

2. Write down the (address, content) pairs for all cells in C .

This is the encoding.

3. To decode,

(a) Recover C from reading the encoding

(b) Query the algorithm $\forall x \in \mathbb{F}_q$. Collect the answers for all queries

that can be answered within C . The lemma condition implies

that we will have $\geq q \left(\frac{q-1}{s}\right)^{OCT} \geq n+1$ different $(x, P(x))$ pairs.

This uniquely determines the polynomial by interpolation.

So, we get a procedure that can encode the whole polynomial in

$$s |C| (\log S + w) = en (\log S + \log q) < (n+1) \log q$$

Therefore, $\exists P_1, P_2$ with the same encoding. Contradiction!

□

Lecture 12/1 - Fine-Grained Complexity

We focus on **K-SAT**: given n variables x_1, \dots, x_n and a K -CNF formula $C_1 \wedge C_2 \wedge \dots \wedge C_m$, where each C_i is of the form $y_1 \vee y_2 \vee \dots \vee y_k$ and y_j is either x_t or $\neg x_t$ for some $t \in [n]$, compute if \exists an assignment $\vec{x} \in \{0, 1\}^n$ that satisfies all C_i .

Brute force: try all 2^n assignments.

Best known: $O(2^{n(1-\frac{c}{k})})$ for constants $c, d > 0$

There is a hypothesis that this is the best we can do.

Strong Exponential Time Hypothesis - (implies $P \neq NP$)

$\forall \epsilon > 0, \exists k \geq 3$ st. K -SAT cannot be solved in $O(2^{(1-\epsilon)n} \text{poly } n)$

Theorem [Impagliazzo, Paturi, 2001]

$SETH \Leftrightarrow SETH$ with $m = O(n)$

Consider the following problem:

Orthogonal Vectors (OV)

Input: a set of N vectors in $\{0, 1\}^d$ ($d = O(\log N)$)

Output: if $\exists u, v$ st. $\langle u, v \rangle = 0 \Leftrightarrow u \wedge v = \emptyset$

Brute force: $O(N^2 d)$ time, compute $\langle u, v \rangle$ for u, v

Theorem [Williams '04]

$SETH \Rightarrow \forall \delta > 0$ OV cannot be solved in $O(n^{2-\delta} \text{poly } d)$ time


Proof: We want to prove the contrapositive: with an $O(n^{2-\delta} \text{poly } d)$ time OV alg., we can construct a $O(2^{(1-\frac{\delta}{2})n} \text{poly } n)$ K -SAT alg.

Consider a K-SAT instance C_1, \dots, C_m , $m = O(n)$

- divide the variables $\{x_1, \dots, x_n\}$ into V_1, V_2 of size $\frac{n}{2}$
- for $j=1, 2$ consider all possible $2^{\frac{n}{2}}$ partial assignments \emptyset to V_j

construct a vector of dimension $m+2$ for each (j, \emptyset)

The vectors are constructed below



- there are $2^{\frac{n}{2}+1}$ vectors in total.

Claim: two vectors are orthogonal iff they combine to satisfy.
i.e. $\langle \vec{v}_{j_1, \emptyset_1}, \vec{v}_{j_2, \emptyset_2} \rangle = 0$ iff $j_1 \neq j_2$ and $(\emptyset_1, \emptyset_2)$ satisfies.

So, $\exists \emptyset_1, \emptyset_2$ s.t. $\langle \vec{v}_{j_1, \emptyset_1}, \vec{v}_{j_2, \emptyset_2} \rangle = 0 \Leftrightarrow C_1, C_2, \dots, C_m$ is satisfiable

\Rightarrow alg for OV in $O(N^{2-\delta} \text{poly } d) \Rightarrow$ K-SAT in time

$$O\left(\left(2^{\frac{n}{2}}\right)^{2-\delta} \text{poly } n\right) = O\left(2^{n(1-\frac{\delta}{2})} \text{poly } n\right)$$

□

Graph Diameter

Given an undirected, unweighted graph $G=(V, E)$ w/ $|V|=n$, $|E|=m$
compute $D = \max_{u, v \in V} d_G(u, v)$

Brute force: breadth-first search in $O(mn)$ time

[RV '13] $\frac{3}{2}$ -approx in time $O(m^{1.5} \text{polylog } m)$
 $\left(\frac{2D}{3} \leq D' \leq D\right)$

Theorem:

SETH $\Rightarrow \forall \epsilon > 0$, $(\frac{3}{2} - \epsilon)$ -approx must take $m \cdot n^{1-o(1)}$ time

Proof: Reduce from OV on N vectors. **Check lecture notes.**

This shows that if \exists alg. for diameter in $m^{1-\delta}$ time,
 \Rightarrow OV is solved in $O(N^{1-\delta}Nd) = O(N^{2-\delta}d)$

□

Now let us look at this through **3-SUM**:

given a set S of n numbers, output whether
 $\exists a, b, c \in S$ s.t. $a+b=c$

There is also 3-SUM convolution:

given $A[1, \dots, n]$, output whether $\exists x, y \in [n]$ s.t. $A[x] + A[y] = A[x+y]$

Naive: $O(n^2)$ for both

3-Sum Conjecture - $\forall \delta > 0$, no alg. solves 3-SUM in $O(n^{2-\delta})$ time.

\iff
 $\forall \delta > 0$, no alg. solves 3-SUM-conv in $O(n^{2-\delta})$ time

Exact triangle (ET)

Given a weighted, undirected graph G , output if $\exists a, b, c \in V$ s.t.
 $w(a, b) + w(b, c) + w(c, a) = 0$

Theorem: If $\exists O(n^{2-\delta})$ alg. for ET, then \exists 3-sum-conv alg. with $O(n^{2-\frac{\delta}{2}})$ time

Proof: Consider an input A to 3-sum-conv. We will construct $O(\sqrt{n})$ graphs $G_1, \dots, G_{\sqrt{n}}$ of size $O(\sqrt{n})$.

A



$T = \Theta(\sqrt{n})$

- G_i is tripartite U_i, V_i, W_i
- U_i has vertices $j \in [n_i^U]$
 - V_i has vertices $s \in [T]$
 - W_i has vertices $t \in [2T]$

- for $j \in U_i, s \in V_i$ add edge (j,s) with weight $A[jT+s]$
- for $j \in U_i, t \in W_i$ add edge (j,t) with weight $-A[(i+j)T+t]$
- for $s \in V_i, t \in W_i$ add edge (s,t) with weight $A[iT+(t-s)]$

Claim: G_i has a zero- Δ iff $\exists x$ in block i s.t. $A[x] + A[y] = A[x+y]$

Since $|G_i| = O(T + \frac{n_i^U}{T})$, $i \in [n^U]$, if $\exists ET$ alg in time $O(n^{2-\delta})$, total time is $O(\sqrt{n} \cdot (\sqrt{n})^{3-\delta}) = O(n^{2-\frac{\delta}{2}})$.

□

Lecture 12/6-

Differential Privacy

Def: A **database** D is a tuple $(\vec{x}_1, \dots, \vec{x}_n)$.

Def: A **counting query** q is a predicate that takes input \vec{x} and outputs $q(\vec{x}) \in \{0, 1\}$. Over a whole DB, $q(D) = \sum_{i=1}^n \frac{q(\vec{x}_i)}{n}$

In the worst case, if the whole world were out to get you or an attacker had all the possible outside information, even a large-scale survey where you answer honestly is not private (even when n large, $0 \ll q(D) \ll 1$)

examples

- ① all other respondents know what they put and can find your answer
- ② Netflix de-anonymization via pattern-matching with external IMDB DB.

We would like machinery to robustly prove that no matter what an attacker knows, they can't break your privacy.

Def: A randomized algorithm M is **α -accurate for q** if, w.h.p.,
 $|M(D) - q(D)| \leq \alpha \quad \forall D$

Def: A randomized algorithm M is **ϵ -differentially private** if $\forall i$, all D, D' s.t. $D_{-i} = D'_{-i}$, \forall sets S of possible outputs,

\forall pairs of databases differing by at most 1 respondent

$$\mathbb{P}\{M(D) \in S\} \leq e^\epsilon \mathbb{P}\{M(D') \in S\}$$

\iff

$$\forall \text{ outputs } r, \quad \left| \ln \left(\frac{\mathbb{P}\{M(D) = r\}}{\mathbb{P}\{M(D') = r\}} \right) \right| \leq \epsilon$$

We want to ensure that ϵ -DP ensures that your participation cannot affect anyone else's (insurance, Man, etc.) Bayesian prior about you.

Prop:

Formally, suppose that someone has a Bayesian prior P about the database state that they will update to \bar{P} after seeing M on the database.
 ϵ -DP guarantees

$$P_{\bar{P}} \{D | M(D)=r\} \leq e^{\pm 2\epsilon} P_{\bar{P}'} \{D | M(D)=r\}$$

Proof:

$$\begin{aligned} P_{\bar{P}} \{D | M(D)=r\} &= \frac{P \{D \in P \wedge M(D)=r\}}{P \{M(D)=r\}} = \frac{P \{D \in P\} P \{M(D)=r\}}{\sum_{\hat{D}} P \{M(\hat{D})=r\} P \{\hat{D} \in P\}} \\ &= \frac{P \{D \in P\} \cdot e^{\pm \epsilon} P \{M(D)=r\}}{\sum_{\hat{D}'} P \{M(\hat{D}')=r\} e^{\pm \epsilon} P \{\hat{D}' \in P\}} = e^{\pm 2\epsilon} P_{\bar{P}'} \{D | M(D)=r\} \end{aligned}$$

← \bar{P} if they see $M(D)=r$ instead of $M(D)=r$

D

Idea 1: Randomized response

With probability p give correct answer $q(\vec{x}_i)$, w.p. $1-p$ flip it v.i. Your response will be more private without worrying about the total dataset or the algorithm. The output vector is $\vec{r} = (r_1, \dots, r_n)$.

$$\Rightarrow \frac{P \{M(D)=\vec{r}\}}{P \{M(D)=\vec{r}'\}} = \frac{P \{M_{\cdot, i}(D)=\vec{r}_{\cdot, i}\} P \{M_{\cdot, \neq i}(D)=\vec{r}_{\cdot, \neq i}\}}{P \{M_{\cdot, i}(D)=\vec{r}'_{\cdot, i}\} P \{M_{\cdot, \neq i}(D)=\vec{r}_{\cdot, \neq i}\}} = \frac{p}{1-p} \approx \epsilon$$

$$\text{if } p = \frac{1}{2} + \frac{\epsilon}{2}$$

The estimate should then be $\frac{1}{2p-1} \left(\sum_i r_i - (1-p) \right)$

which is correct in expectation with variance $\frac{p(1-p)}{n(2p-1)^2}$

Idea 2: Add noise

Add noise to each response r : drawn from $\text{Lap}(\frac{1}{\epsilon n})$. So, the PDF of the noise is $f(x) = \frac{1}{2(\frac{1}{\epsilon n})} e^{-|x|/(\frac{1}{\epsilon n})}$

We are concerned with the density ratios between D and D'

$$\frac{f_{M(D)}(r)}{f_{M(D')}(r)} = \frac{e^{-\epsilon n |r - q(D)|}}{e^{-\epsilon n |r - q(D')|}} \leq e^{-\epsilon}$$

differs by at most ϵ because cost changes by $\leq \epsilon$ when changing one response

It is correct in expectation with variance $\frac{2}{(\epsilon n)^2}$.

Def: A randomized algorithm M is (ϵ, δ) -differentially private if \forall all D, D' s.t. $D_i = D'_i$, \forall sets S of possible outputs,

\forall pairs of databases differing by at most 1 respondent

$$\forall \text{ outputs } r, \frac{\mathbb{P}\{M(D) = r\}}{\mathbb{P}\{M(D') = r\}} \leq e^{\epsilon} + \delta \mathbb{P}\{M(D') = r\}$$

check this

Theorem: If M_1, \dots, M_k all ϵ -DP, then the algorithm that answers all queries (M_1, \dots, M_k) is $k\epsilon$ -DP.

If M_1, \dots, M_k all (ϵ, δ) -DP, then (M_1, \dots, M_k) is $(k\epsilon/2 + \sqrt{2k \ln(1/\delta)} \epsilon, \delta)$ -DP.

Proof: Lol no \therefore

□

ϵ -DP is also robust to groups of individuals!

ϵ -DP is also robust to postprocessing:

\forall algorithms A , M is ϵ -DP $\Rightarrow A \circ M$ is ϵ -DP

Lecture 12/8 - Smoothed Analysis

- ① given a worst-case input \vec{x} (adversarial)
- ② randomly smooth \vec{x} to \vec{y} using some distribution of magnitude σ
(maintain adversarial big picture, but randomize lower-order bits)
- ③ \vec{y} is the true input
- ④ $\forall \vec{x}$ (even adversarial), $\mathbb{E}_{\vec{y} \in \text{smooth}_{\sigma}(\vec{x})} \{ \text{runtime}_A(\vec{y}) \} = \text{poly}(|\vec{x}|, \frac{1}{\sigma})$

Super cool result we won't prove:

Theorem: (Spielman, Teng '01)


Let \vec{c}, A, \vec{b} define a LP.
objective \vec{c} , *constraint matrix* A , \vec{b}

Let $\text{smooth}(\vec{c}, A, \vec{b})$ add i.i.d. $\mathcal{U}(0, \sigma^2)$ to $A_{ij}, b_j \forall i, j$.
Then, the simplex algorithm is smoothed poly-time in this model.
exponential worst-case





check Tim's notes for discussion about what this means about simplex in practice

Metric TSP

Consider a metric space $[0, 1] \times [0, 1] \subset \mathbb{R}^2$ with L₁ metric.



TSP is to find the lowest cost Hamiltonian cycle.

There is an algorithm called **2-OPT** that performs local searches. Basically, for any current tour , consider replacing each pair of nonadjacent edges one-by-one as follows:  \rightarrow  \rightarrow 

Keep the improvements and terminate when no pair replacement helps.

2-OPT has worst case exponential time to terminate but smoothed polynomial runtime, as we will see below.

Theorem: If node x_i is smoothed to y_i independently according to any distribution w/ PDF f_i s.t. $f_i(y_i) \leq \frac{1}{\sigma}$ $\forall y_i$ (bounded density), then 2-OPT is smoothed poly-time.

Def: A swap $(u,v), (x,y)$ is **ϵ -bad** if

$$\| \|u-v\| + \|x-y\|, -(\|u-x\| + \|v-y\|) \in (0, \epsilon)$$

Swapping makes $c\epsilon$ progress, so alg. continues with very little progress

Lemma: $\forall \epsilon, \mathbb{P}_{y_i \text{ smooth}_{\sigma}(\cdot)} \left\{ \text{any swap in } \tilde{g} \text{ is } \epsilon\text{-bad} \right\} \leq \frac{\epsilon n^4}{\sigma}$

Proof of Lemma: First, observe that there are n^4 possible choices of $((u,v), (x,y))$ and so if all of these are good, there can be no graph with an ϵ -bad swap. For each $(u,v), (x,y)$, let

$$\star = \|u-v\| + \|x-y\| - (\|u-x\| + \|v-y\|)$$

If we fix the relative ordering of $\{u, v, x, y\}$ and $\{u_2, v_2, x_2, y_2\}$, then \star is linear in all vars and all coeffs are in $\{-2, 0, 2\}$. So, $\star \in \{\text{linear fns with coeffs } \in \{-2, 0, 2\}\}$.

There are $(4!)^2 n^4$ possible linear functions. Now, for any linear function in this set, in our smoothed model, we want to show that it is $\in (0, \epsilon)$. If all the coefficients are 0, it holds trivially.

Now, suppose WLOG that u_i has a ± 2 coefficient. Sample the smoothed versions of all other variables except u_i .

The function is $\pm 2u_i + \alpha_2 u_2 + \alpha_3 u_3 + \dots$

So, the function is in $(0, \epsilon)$ iff $u_i \in \left(\frac{-c}{\pm 2}, \frac{\epsilon - c}{\pm 2} \right)$

of width $\frac{\epsilon}{2}$. The max probability that u_i can lie in this range is $\leq \frac{\epsilon}{2\sigma}$ because of the bounded densities.

For each possible swap, a union bound over the $(n!)^2$ possible functions yields that $\mathbb{P}\{\epsilon\text{-bad swap}\} \leq \frac{(n!)^2}{2} \frac{\epsilon}{\sigma}$

Now, a union bound over the n^n possible swaps proves the Lemma. \square

Proof of Theorem:

Note first that since each edge weight ≤ 1 , the initial tour is $\leq 2n$. So, if no ϵ -bad swaps, there will be $\leq \frac{2n}{\epsilon}$ iterations.

The Lemma gives that

$$\mathbb{P}\{\text{more than } M \text{ iterations}\} \leq O\left(\frac{2n}{\epsilon} \frac{n^4}{\sigma}\right)$$

w.h.p. poly(n, m)

The expected # of iterations, since there must be $\leq n!$ possible tours, is

$$\begin{aligned} \mathbb{E}\{\#\text{ites}\} &= \sum_{m=1}^{n!} \mathbb{P}\{\text{more than } m \text{ iters}\} \leq \sum_{m=1}^{n!} O\left(\frac{n^5}{\sigma}\right) \cdot \frac{1}{m} \\ &= O\left(\frac{n^5}{\sigma}\right) n \log n \end{aligned}$$

poly in expectation

Both together give smoothed poly runtime. \square