

CS-GY 9223 D: Lecture 14

Leverage Score Sampling, Spectral Sparsification, Taste of my research

NYU Tandon School of Engineering, Prof. Christopher Musco


Now IPS

CVPB

- Final project needs to be submitted by 12/18 on NYU Classes. 6 page writeup minimum. I am still available for last minute meetings if needed.
- Please fill out course feedback!
- I desperately need graders to help next year – if you will be around in the Fall 2021 semester, let me know.


Main idea: If you want to compute singular vectors or eigenvectors, multiply two matrices, solve a regression problem, etc.:

1. Compress your matrices using a randomized method.
2. Solve the problem on the smaller or sparser matrix.
 - \tilde{A} called a “sketch” or “coreset” for A .



SKETCHED REGRESSION

Randomized approximate regression using a Johnson-Lindenstrauss Matrix:



Input: $A \in \mathbb{R}^{n \times d}$, $b \in \mathbb{R}^n$.

Algorithm: Let $\tilde{x}^* = \arg \min_x \|\Pi A x - \Pi b\|_2^2$.

Goal: Want $\|\tilde{A} \tilde{x}^* - b\|_2^2 \leq (1 + \epsilon) \min_x \|A x - b\|_2^2$

Theorem (Randomized Linear Regression)

Let $\mathbf{\Pi}$ be a properly scaled JL matrix (random Gaussian, sign, sparse random, etc.) with $m = \tilde{O}\left(\frac{d}{\epsilon^2}\right)$ rows. Then with probability $(1 - \delta)$, for any $\mathbf{A} \in \mathbb{R}^{n \times d}$ and $\mathbf{b} \in \mathbb{R}^n$,

$$\|\mathbf{A}\tilde{\mathbf{x}}^* - \mathbf{b}\|_2^2 \leq (1 + \epsilon) \min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2$$

where $\tilde{\mathbf{x}}^* = \arg \min_{\mathbf{x}} \|\mathbf{\Pi A x} - \mathbf{\Pi b}\|_2^2$.

Theorem (Subspace Embedding)

Let $\mathbf{A} \in \mathbb{R}^{n \times d}$ be a matrix. If $\mathbf{\Pi} \in \mathbb{R}^{m \times n}$ is chosen from any distribution \mathcal{D} satisfying the Distributional JL Lemma, then with probability $1 - \delta$,

$$\sqrt{A} \rightarrow \mathbb{R}^m$$

$$(1 - \epsilon) \|\mathbf{Ax}\|_2^2 \leq \|\mathbf{\Pi Ax}\|_2^2 \leq (1 + \epsilon) \|\mathbf{Ax}\|_2^2$$


for all $\mathbf{x} \in \mathbb{R}^d$, as long as $m = O\left(\frac{d + \log(1/\delta)}{\epsilon^2}\right)$.

Implies regression result, and more.

Example: The any singular value $\tilde{\sigma}_i$ of $\mathbf{\Pi A}$ is a $(1 \pm \epsilon)$ approximation to the true singular value σ_i of \mathbf{B} .

SUBSAMPLING METHODS

Recurring research interest: Replace random projection methods with random sampling methods. Prove that for essentially all problems of interest, can obtain same asymptotic runtimes.




Sampling has the added benefit of preserving matrix sparsity or structure, and can be applied in a wider variety of settings where random projections are too expensive.

SUBSAMPLING METHODS


First goal: Can we use sampling to obtain subspace embeddings? I.e. for a given \mathbf{A} find $\tilde{\mathbf{A}}$ whose rows are a (weighted) subset of rows in \mathbf{A} and:

$$(1 - \epsilon)\|\mathbf{Ax}\|_2^2 \leq \|\tilde{\mathbf{A}}\mathbf{x}\|_2^2 \leq (1 + \epsilon)\|\mathbf{Ax}\|_2^2.$$



EXAMPLE WHERE STRUCTURE MATTERS

Let \mathbf{B} be the edge-vertex incidence matrix of a graph G with vertex set V , $|V| = d$. Recall that $\mathbf{B}^T \mathbf{B} = \mathbf{L}$.



$$\begin{array}{cccccccc} +1 & -1 & 0 & 0 & 0 & 0 & 0 & 0 \\ +1 & 0 & -1 & 0 & 0 & 0 & 0 & 0 \\ +1 & 0 & 0 & -1 & 0 & 0 & 0 & 0 \\ +1 & 0 & 0 & 0 & -1 & 0 & 0 & 0 \\ 0 & +1 & 0 & 0 & 0 & -1 & 0 & 0 \\ 0 & 0 & +1 & -1 & 0 & 0 & 0 & 0 \\ 0 & 0 & +1 & 0 & -1 & 0 & 0 & 0 \end{array}$$

\vdots

$$\begin{array}{cccccccc} 0 & 0 & 0 & 0 & +1 & 0 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & +1 & 0 & -1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & +1 & -1 \end{array}$$

\mathbf{B}


Recall that if $\mathbf{x} \in \{-1, 1\}^d$ is the cut indicator vector for a cut S in the graph, then $\frac{1}{4} \|\mathbf{B}\mathbf{x}\|_2^2 = \text{cut}(S, V \setminus S)$.

LINEAR ALGEBRAIC VIEW OF CUTS

$$S = \{1, 2, 3, 5\}$$

$$V \setminus S = \{4, 6, 7, 8\}$$


$$\mathbf{x} = [1, 1, 1, -1, 1, -1, -1, -1]$$



$\mathbf{x} \in \{-1, 1\}^d$ is the cut indicator vector for a cut S in the graph,
then $\frac{1}{4} \|\mathbf{B}\mathbf{x}\|_2^2 = \text{cut}(S, V \setminus S)$

WEIGHTED CUTS

Extends to weighted graphs, as long as square root of weights is included in \mathbf{B} . Still have the $\mathbf{B}^T \mathbf{B} = \mathbf{L}$.



$$\begin{bmatrix} +\sqrt{w_{12}} & -\sqrt{w_{12}} & 0 & 0 & 0 & 0 & 0 & 0 \\ +\sqrt{w_{13}} & 0 & -\sqrt{w_{13}} & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

\vdots

$$\begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & +\sqrt{w_{78}} & -\sqrt{w_{78}} \end{bmatrix}$$


\mathbf{B}

And still have that if $\mathbf{x} \in \{-1, 1\}^d$ is the cut indicator vector for a cut S in the graph, then $\frac{1}{4} \|\mathbf{B}\mathbf{x}\|_2^2 = \text{cut}(S, V \setminus S)$.

SPECTRAL SPARSIFICATION

Goal: Approximate \mathbf{B} by a weighted subsample. I.e. by $\tilde{\mathbf{B}}$ with $m \ll |E|$ rows, each of which is a scaled copy of a row from \mathbf{B} .

subsampling matrix




Natural goal: $\tilde{\mathbf{B}}$ is a subspace embedding for \mathbf{B} . In other words, $\tilde{\mathbf{B}}$ has $\approx O(d)$ rows and for all \mathbf{x} ,

$$(1 - \epsilon) \|\mathbf{B}\mathbf{x}\|_2^2 \leq \|\tilde{\mathbf{B}}\mathbf{x}\|_2^2 \leq (1 + \epsilon) \|\mathbf{B}\mathbf{x}\|_2^2.$$

HISTORY SPECTRAL SPARSIFICATION

$\tilde{\mathbf{B}}$ is itself an edge-vertex incidence matrix for some sparser graph \tilde{G} , which preserves many properties about G ! \tilde{G} is called a spectral sparsifier for G .



For example, we have that for any set S ,

$$(1 - \epsilon) \text{cut}_{\tilde{G}}(S, V \setminus S) \leq \text{cut}_G(S, V \setminus S) \leq (1 + \epsilon) \text{cut}_{\tilde{G}}(S, V \setminus S).$$

So \tilde{G} can be used in place of G in solving e.g. max/min cut problems, balanced cut problems, etc.

In contrast $\Pi \mathbf{B}$ would look nothing like an edge-vertex incidence matrix if Π is a JL matrix.

HISTORY OF SPECTRAL SPARSIFICATION

Spectral sparsifiers were introduced in 2004 by Spielman and Teng in an influential paper on faster algorithms for solving Laplacian linear systems.


- (Generalize the cut sparsifiers of Benczur, Karger '96.
- Further developed in work by Spielman, Srivastava + Batson, '08.
- Have had huge influence in algorithms, and other areas of mathematics – this line of work lead to the 2013 resolution of the Kadison-Singer problem in functional analysis by Marcus, Spielman, Srivastava.

This class: Learn about an important random sampling algorithm for constructing spectral sparsifiers, and subspace embeddings for matrices more generally.

NATURAL FIRST ATTEMPT

Goal: Find \tilde{A} such that $\|\tilde{A}x\|_2^2 = (1 \pm \epsilon)\|Ax\|_2^2$ for all x .

Possible Approach: Construct \tilde{A} by uniformly sampling rows from A .



Can check that this approach fails even for the special case of a graph vertex-edge incidence matrix.

IMPORTANCE SAMPLING FRAMEWORK

Key idea: Importance sampling. Select some rows with higher probability.

Suppose \mathbf{A} has n rows $\mathbf{a}_1, \dots, \mathbf{a}_n$. Let p_1, \dots, p_n $\in [0, 1]$ be sampling probabilities. Construct $\tilde{\mathbf{A}}$ as follows:

- For $i = 1, \dots, n$
 - Select \mathbf{a}_i with probability p_i .
 - If \mathbf{a}_i is selected, add the scaled row $\frac{1}{\sqrt{p_i}} \mathbf{a}_i$ to $\tilde{\mathbf{A}}$.

Remember, ultimately want that $\|\tilde{\mathbf{A}}\mathbf{x}\|_2^2 = (1 \pm \epsilon) \|\mathbf{A}\mathbf{x}\|_2^2$ for all \mathbf{x} .

Claim 1: $\mathbb{E}[\|\tilde{\mathbf{A}}\mathbf{x}\|_2^2] = \|\mathbf{A}\mathbf{x}\|_2^2$. $\mathbf{A}\mathbf{x} = \mathbf{y}$.

$$\mathbb{E} \|\tilde{\mathbf{A}}\mathbf{x}\|_2^2 = \sum_{i=1}^n y_i^2 \cdot \left(\frac{1}{\sqrt{p_i}}\right)^2 \mathbb{E}[Z_i] \quad \text{where } Z_i = \begin{cases} 1 & \text{w/ prob } p_i \\ 0 & \text{w/ prob } 1-p_i \end{cases}$$

Claim 2: Expected number of rows in $\tilde{\mathbf{A}}$ is $\sum_{i=1}^n p_i$.

How should we choose the probabilities p_1, \dots, p_n ?

statistical

1. Introduce the idea of row **leverage scores**.
2. Motivate why these scores make for good sampling probabilities.
3. Prove (at least mostly) that sampling with probabilities proportional to these scores yields a subspace embedding (or a spectral sparsifier) with a near optimal number of rows.

MAIN RESULT

Let $\mathbf{a}_1, \dots, \mathbf{a}_n$ be \mathbf{A} 's rows. We define the **statistical leverage score** τ_i of row \mathbf{a}_i as:

$$\tau_i = \mathbf{a}_i^T (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{a}_i.$$

We will show that τ_i is a natural importance measure for each row in \mathbf{A} .

We have that $\tau_i \in [0, 1]$ and $\sum_{i=1}^n \tau_i = d$ if \mathbf{A} has d columns.

MAIN RESULT

For $i = 1, \dots, n$,

$$\tau_i = \mathbf{a}_i^T (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{a}_i.$$

fixed constant

Theorem (Subspace Embedding from Subsampling)

For each i , and fixed constant c , let $\underline{p}_i = \min \left(1, \frac{c \log d}{\epsilon^2} \cdot \underline{\tau}_i \right)$. Let

$\tilde{\mathbf{A}}$ have rows sampled from \mathbf{A} with probabilities p_1, \dots, p_n .

With probability 9/10, for any \mathbf{x}

$$(1 - \epsilon) \|\mathbf{Ax}\|_2^2 \leq \|\tilde{\mathbf{A}}\mathbf{x}\|_2^2 \leq (1 + \epsilon) \|\mathbf{Ax}\|_2^2,$$

and $\tilde{\mathbf{A}}$ has $O(d \log d / \epsilon^2)$ rows in expectation.

$$O(d/\epsilon^2)$$

How should we choose the probabilities p_1, \dots, p_n ?


As usual, consider a single vector \mathbf{x} and understand how to sample to preserve norm of $\mathbf{y} = \mathbf{Ax}$:

$$\|\tilde{\mathbf{A}}\mathbf{x}\|_2^2 = \|\mathbf{S}\mathbf{Ax}\|_2^2 = \|\mathbf{S}\mathbf{y}\|_2^2 \approx \|\mathbf{y}\|_2^2 = \|\mathbf{Ax}\|_2^2.$$


Then we can union bound over an ϵ -net to extend to all \mathbf{x} .

VECTOR SAMPLING

As discussed a few lectures ago, uniform sampling only works well if $\mathbf{y} = \mathbf{Ax}$ is “flat”.



Instead consider sampling with probabilities at least proportional to the magnitude of \mathbf{y} 's entries:

$$\underline{p_i} \geq c \cdot \frac{y_i^2}{\|\mathbf{y}\|_2^2} \text{ for constant } c \text{ to be determined.}$$

VARIANCE ANALYSIS

Let $\tilde{\mathbf{y}}$ be the subsampled \mathbf{y} . Recall that, when sampling with probabilities p_1, \dots, p_n , for $i = 1, \dots, n$ we add y_i to $\tilde{\mathbf{y}}$ with probability p_i and reweight by $\frac{1}{\sqrt{p_i}}$.

$$\|\tilde{\mathbf{y}}\|_2^2 = \sum_{i=1}^n y_i^2 \left(\frac{1}{\sqrt{p_i}}\right)^2 \cdot Z_i \quad Z_i = \begin{cases} 1 & \text{w/ prob } p_i \\ 0 & \text{otherwise} \end{cases}$$

$$\begin{aligned} \sigma^2 = \text{Var}[\|\tilde{\mathbf{y}}\|_2^2] &= \sum_{i=1}^n \frac{y_i^4}{p_i^2} \cdot \text{Var}[Z_i] \leq \sum_{i=1}^n \frac{y_i^4}{p_i^2} \cdot p_i \\ &= \sum_{i=1}^n \frac{y_i^4}{p_i} \end{aligned}$$

$$p_i = \frac{c y_i^2}{\|\mathbf{y}\|_2^2}$$

$$\begin{aligned} \text{Var}[\|\tilde{\mathbf{y}}\|_2^2] &\leq \frac{1}{c} \sum_{i=1}^n y_i^2 \cdot \|\mathbf{y}\|_2^2 \\ &= \frac{1}{c} \|\mathbf{y}\|_2^2 \sum_{i=1}^n y_i^2 \\ &= \frac{1}{c} \|\mathbf{y}\|_2^4 \end{aligned}$$

VARIANCE ANALYSIS

Recall Chebyshev's Inequality:

$$\Pr[|\underbrace{\|\tilde{\mathbf{y}}\|_2^2}_{\geq} - \underbrace{\|\mathbf{y}\|_2^2}| \geq \underbrace{\frac{1}{\sqrt{\delta}} \cdot \sigma}] \leq \delta$$

$\frac{1}{\sqrt{c}} \propto \|\mathbf{y}\|_2^2$

We want error $|\|\tilde{\mathbf{y}}\|_2^2 - \|\mathbf{y}\|_2^2| \leq \underline{\epsilon \|\mathbf{y}\|_2^2}$.

Need set $c = \frac{1}{\delta \epsilon^2}$.¹

If we knew y_1, \dots, y_n , the number of samples we take in expectation is:

$$\underbrace{\sum_{i=1}^n p_i}_{\text{samples}} = \sum_{i=1}^n c \cdot \frac{y_i^2}{\|\mathbf{y}\|_2^2} = \underbrace{\frac{1}{\delta \epsilon^2}}_{\text{samples}}$$

¹Using the right Bernstein bound we can improve to $c = \underline{\underline{O(\log(1/\delta)/\epsilon^2)}}$.

MAXIMIZATION CHARACTERIZATION

~~Ax~~

But we of course don't know y_1, \dots, y_n , and even so these values aren't fixed. We wanted to prove a bound for $\mathbf{y} = \mathbf{A}\mathbf{x}$ for any \mathbf{x} .

Idea behind leverage scores: Sample row i from \mathbf{A} using the worst case (largest necessary) sampling probability:

$$\tau_i = \mathbf{a}_i^T (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{a}_i \quad \tau_i = \max_{\mathbf{x}} \frac{y_i^2}{\|\mathbf{y}\|_2^2} \quad \text{where} \quad \mathbf{y} = \mathbf{A}\mathbf{x}.$$

If we sample with probability $p_i = \frac{1}{\epsilon^2} \cdot \tau_i$, then we will be sampling by at least $\frac{1}{\epsilon^2} \cdot \frac{y_i^2}{\|\mathbf{y}\|_2^2}$, no matter what \mathbf{y} is.

Two major concerns: 1) How to compute τ_1, \dots, τ_n , and 2) the number of samples we take will be roughly $\sum_{i=1}^n \tau_i$. How do we bound this?

MAXIMIZATION CHARACTERIZATION

$$\tau_i = \alpha_i^T (A^T A)^{-1} \alpha_i$$

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

$$A^T A = U \Sigma^2 U^T$$


$$(A^T A)^{-1} = U \Sigma^{-2} U^T$$

where

$$q_i^T x$$

$$\tau_i = \max_x \frac{y_i^2}{\|y\|_2^2}$$

$$y = Ax.$$



$$\begin{aligned} \max_x \frac{(q_i^T x)^2}{\|Ax\|_2^2} &= \max_x \frac{(q_i^T (A^T A)^{-1/2} (A^T A)^{1/2} x)^2}{x^T A^T A x} \\ &\leq \frac{q_i^T (A^T A)^{-1/2} (A^T A)^{-1/2} q_i \cdot x^T (A^T A)^{1/2} (A^T A)^{1/2} x}{x^T A^T A x} \end{aligned}$$

$$\leq q_i^T (A^T A)^{-1} q_i$$

$$x = (A^T A)^{-1} q_i$$

$$(q_i^T (A^T A)^{-1} q_i)^2$$

$$q_i^T (A^T A)^{-1} A^T A (A^T A)^{-1} q_i$$

Recall Cauchy-Schwarz inequality: $(w^T z)^2 \leq w^T w \cdot z^T z$

$$(A^T A)^{1/2} = U \Sigma^{1/2} U^T$$

$$(A^T A)^{-1/2} = U \Sigma^{-1/2} U^T$$

$$A^T A = U \Sigma^2 U^T$$

$$(A^T A)^{-1} = U \Sigma^{-2} U^T$$

EQUIVALENT MINIMIZATION CHARACTERIZATION

$$\hat{v}_i = \mathbf{a}_i^T (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{a}_i$$

$$\tau_i = \min_{\mathbf{z} \text{ such that } \mathbf{A}^T \mathbf{z} = \mathbf{a}_i} \|\mathbf{z}\|_2^2$$


$$\begin{bmatrix} \mathbf{a}_i & \mathbf{A}^T \end{bmatrix} \mathbf{z} = \mathbf{a}_i$$

EQUIVALENT MINIMIZATION CHARACTERIZATION

$$\tau_i = \min_{\substack{z \text{ such that } A^T z = a_i}} \|z\|_2^2.$$



=



=




$$\|z\|_2^2 = 1$$

$$4 \cdot \left(\frac{1}{4}\right)^2 = 1$$

Gives clearer picture of leverage score τ_i as a measure of “uniqueness” for row a_i .

LEVERAGE SCORE SAMPLING

Leverage score sampling:



- For $i = 1, \dots, n$,
 - Compute $\tau_i = \mathbf{a}_i^T (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{a}_i$
 - Set $p_i = \frac{\log(1/\delta)}{e^2} \cdot \tau_i$ *constant*
 - Add row \mathbf{a}_i to $\tilde{\mathbf{A}}$ with probability p_i and reweight by $\frac{1}{\sqrt{p_i}}$.

For any fixed \mathbf{x} , we will have that

$$(1 - \epsilon) \|\mathbf{Ax}\|_2^2 \leq \|\tilde{\mathbf{A}}\mathbf{x}\|_2^2 \leq (1 + \epsilon) \|\mathbf{Ax}\|_2^2 \text{ with probability } (1 - \delta).$$

(How many rows do we sample in expectation?)

$$\sum_{i=1}^n \frac{\log(1/\delta)}{e^2} \cdot \tau_i = \frac{\log(1/\delta)}{e^2} \cdot \underbrace{\sum_{i=1}^n \tau_i}_{\text{down}}$$

SUM OF LEVERAGE SCORES

Claim: No matter how large n is, $\sum_{i=1}^n \tau_i = d$ a matrix $A \in \mathbb{R}^{n \times d}$.

$$\begin{aligned}
 & \sum_{i=1}^n \mathbf{a}_i^T (A^T A)^{-1} \mathbf{a}_i && \text{tr}(I) \\
 &= \sum_{i=1}^n \text{tr}(\mathbf{a}_i^T (A^T A)^{-1} \mathbf{a}_i) && n \\
 &= \sum_{i=1}^n \text{tr}((A^T A)^{-1} \mathbf{a}_i \mathbf{a}_i^T) \\
 &= \text{tr}((A^T A)^{-1} \sum_{i=1}^n \mathbf{a}_i \mathbf{a}_i^T) = \text{tr}((A^T A)^{-1} A^T A) = \text{tr}(I)
 \end{aligned}$$

"Zero-sum" law for the importance of matrix rows.



$d \times d$

LEVERAGE SCORE SAMPLING

Leverage score sampling:

- For $i = 1, \dots, n$,
 - Compute $\tau_i = \mathbf{a}_i^T (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{a}_i$.
 - Set $p_i = \frac{c \log(1/\delta)}{\epsilon^2} \cdot \tau_i$.
 - Add row \mathbf{a}_i to $\tilde{\mathbf{A}}$ with probability p_i and reweight by $\frac{1}{\sqrt{p_i}}$.

For any fixed \mathbf{x} , we will have that

$(1 - \epsilon) \|\mathbf{Ax}\|_2^2 \leq \|\tilde{\mathbf{A}}\mathbf{x}\|_2^2 \leq (1 + \epsilon) \|\mathbf{Ax}\|_2^2$ with high probability.

And since $\sum_{i=1}^n p_i = \frac{c \log(1/\delta)}{\epsilon^2} \cdot \sum_{i=1}^n \tau_i$, $\tilde{\mathbf{A}}$ contains $O\left(\frac{d \log(1/\delta)}{\epsilon^2}\right)$ rows in expectation.

↓
d

Last step: need to extend to all \mathbf{x} .

MAIN RESULT

$$6 \left(d + \frac{d \log(1/\delta)}{\epsilon^2} \right) \rightarrow d$$

Naive ϵ -net argument leads to d^2 dependence since we need to set $\delta = c^d$. Getting the right $d \log d$ dependence below requires a standard “matrix Chernoff bound” (see e.g. Tropp 2015).

Theorem (Subspace Embedding from Subsampling)

For each i , and fixed constant c , let $p_i = \min \left(1, \frac{c \log d}{\epsilon^2} \cdot \tau_i \right)$. Let $\tilde{\mathbf{A}}$ have rows sampled from \mathbf{A} with probabilities p_1, \dots, p_n . With probability $9/10$,


$$(1 - \epsilon) \|\mathbf{Ax}\|_2^2 \leq \|\tilde{\mathbf{A}}\mathbf{x}\|_2^2 \leq (1 + \epsilon) \|\mathbf{Ax}\|_2^2,$$

and $\tilde{\mathbf{A}}$ has $O(d \log d / \epsilon^2)$ rows in expectation.

$$d^r \quad \frac{d \log(1/\delta)}{\epsilon^2} \rightarrow d \quad \delta = c^d$$

SPECTRAL SPARSIFICATION COROLLARY

For any graph G with d nodes, there exists a graph \tilde{G} with $O(\underline{d \log d / \epsilon^2})$ edges such that, for all \mathbf{x} , $\|\tilde{\mathbf{B}}\mathbf{x}\|_2^2 = (1 \pm \epsilon)\|\mathbf{B}\mathbf{x}\|_2^2$.



As a result, the value of any cut in \tilde{G} is within a $(1 \pm \epsilon)$ factor of the value in G , the Laplacian eigenvalues are with a $(1 \pm \epsilon)$ factors, etc.


Spectral, Sparsification 2008

ANOTHER APPLICATION: ACTIVE REGRESSION

In many applications, computational costs are second order to data collection costs. We have a huge range of possible data points $\mathbf{a}_1, \dots, \mathbf{a}_n$ that we can collect labels/values b_1, \dots, b_n for. Goal is to learn \mathbf{x} such that:


$$\mathbf{a}_i^T \mathbf{x} \approx b_i.$$

Want to do so after observing as few b_1, \dots, b_n as possible. Applications include healthcare, environmental science, etc.



ANOTHER APPLICATION: ACTIVE REGRESSION

Can be solved via random sampling for linear models.



$$\underline{a_i^T (A^T A)^{-1} a_i}$$

Claim: Let $\tilde{\mathbf{A}}$ is an $O(1)$ -factor subspace embedding for \mathbf{A} (obtained via leverage score sampling). Then $\tilde{\mathbf{x}} = \arg \min \|\tilde{\mathbf{A}}\mathbf{x} - \tilde{\mathbf{b}}\|_2^2$ satisfies:

$$\|\mathbf{A}\tilde{\mathbf{x}} - \mathbf{b}\|_2^2 \leq O(1)\|\mathbf{A}\mathbf{x}^* - \mathbf{b}\|_2^2,$$

where $\mathbf{x}^* = \arg \min \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2$. Computing $\tilde{\mathbf{x}}$ only requires collecting $O(d \log d)$ labels (independent of n).

Lots of applications:

- Robust bandlimited, multiband, and polynomial interpolation [STOC 2019].
- Robust active learning for Gaussian process regression [NeurIPS 2020].

Claim: $\tilde{\mathbf{x}} = \arg \min \|\tilde{\mathbf{A}}\mathbf{x} - \tilde{\mathbf{b}}\|_2^2$ satisfies:

$$\|\mathbf{A}\tilde{\mathbf{x}} - \mathbf{b}\|_2^2 \leq O(1)\|\mathbf{A}\mathbf{x}^* - \mathbf{b}\|_2^2,$$

where $\mathbf{x}^* = \arg \min \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2$. Computing $\tilde{\mathbf{x}}$ only requires collecting $O(d \log d)$ labels (independent of n).

Proof:


SOME OTHER THINGS I HAVE WORKED ON

Problem: Computing leverage scores $\tau_i = \mathbf{a}_i^T \overset{O(nd^2)}{\underline{\underline{(\mathbf{A}^T \mathbf{A})}^{-1}}} \mathbf{a}_i$ is expensive.

Main algorithmic idea: Bootstrap leverage score sampling from uniform sampling (ITCS 2015).

SOME OTHER THINGS I HAVE WORKED ON

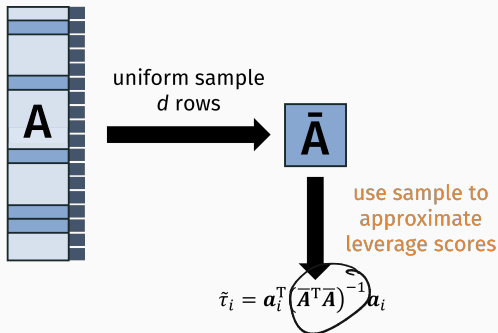
Problem: Computing leverage scores $\tau_i = \mathbf{a}_i^T (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{a}_i$ is expensive.



Main algorithmic idea: Bootstrap leverage score sampling from uniform sampling (ITCS 2015).

SOME OTHER THINGS I HAVE WORKED ON

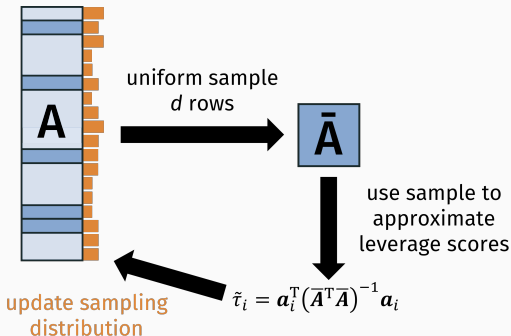
Problem: Computing leverage scores $\tau_i = \mathbf{a}_i^T (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{a}_i$ is expensive.



Main algorithmic idea: Bootstrap leverage score sampling from uniform sampling (ITCS 2015).

SOME OTHER THINGS I HAVE WORKED ON

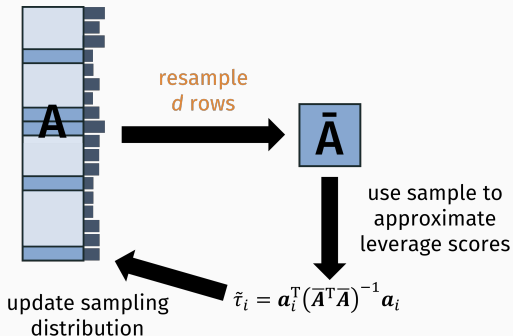
Problem: Computing leverage scores $\tau_i = \mathbf{a}_i^T (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{a}_i$ is expensive.



Main algorithmic idea: Bootstrap leverage score sampling from uniform sampling (ITCS 2015).

SOME OTHER THINGS I HAVE WORKED ON

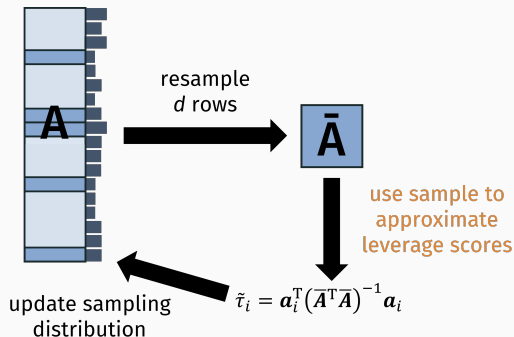
Problem: Computing leverage scores $\tau_i = \mathbf{a}_i^T (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{a}_i$ is expensive.



Main algorithmic idea: Bootstrap leverage score sampling from uniform sampling (ITCS 2015).

SOME OTHER THINGS I HAVE WORKED ON


Problem: Computing leverage scores $\tau_i = \mathbf{a}_i^T (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{a}_i$ is expensive.



Main algorithmic idea: Bootstrap leverage score sampling from uniform sampling (ITCS 2015).

SOME OTHER THINGS I HAVE WORKED ON

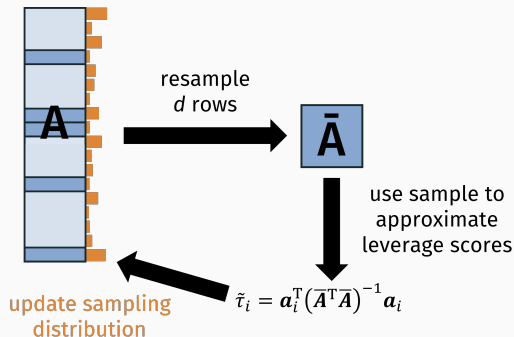
Problem: Computing leverage scores $\tau_i = \mathbf{a}_i^T (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{a}_i$ is expensive.



Main algorithmic idea: Bootstrap leverage score sampling from uniform sampling (ITCS 2015).

SOME OTHER THINGS I HAVE WORKED ON

Problem: Computing leverage scores $\tau_i = \mathbf{a}_i^T (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{a}_i$ is expensive.



After $O(\log n)$ rounds, $\tilde{\tau}_i \approx \tau_i$ for all i .


Main algorithmic idea: Bootstrap leverage score sampling from uniform sampling (ITCS 2015).

SOME THINGS I HAVE WORKED ON

Problem: Sometimes we want to compress down to $\ll d$ rows or columns. E.g. we don't need a full subspace embedding, but just want to find a near optimal rank k approximation.

Approach: Use “regularized” version of the leverage scores:


$$\bar{\tau}_i = \mathbf{a}_i^T (\mathbf{A}^T \mathbf{A} + \lambda \mathbf{I})^{-1} \mathbf{a}_i$$



Result: Sample $O(k \log k / \epsilon)$ columns whose span contains a near-optimal ~~low~~ k -approximation to \mathbf{A} (SODA 2017).

EXAMPLE RESULT: SUBLINEAR TIME KERNEL APPROXIMATION

The first $O(\underline{nk^2/\epsilon^2})$ time algorithm² for near optimal rank- k approximation of any $n \times n$ positive semidefinite kernel matrix:



Based on the classic Nyström method. Importantly, does not even require constructing \mathbf{K} explicitly, which takes $O(n^2)$ time.

²NeurIPS 2017.