Class 9

Compressed Sensing and the RIP

and a recap of what we’ve seen so far

Announcements

- HW4 due Friday!
- I won’t be here Thursday, but you will be in capable hands.
    - Prasanna will lead class!

Results from feedback form

- Main piece of feedback:

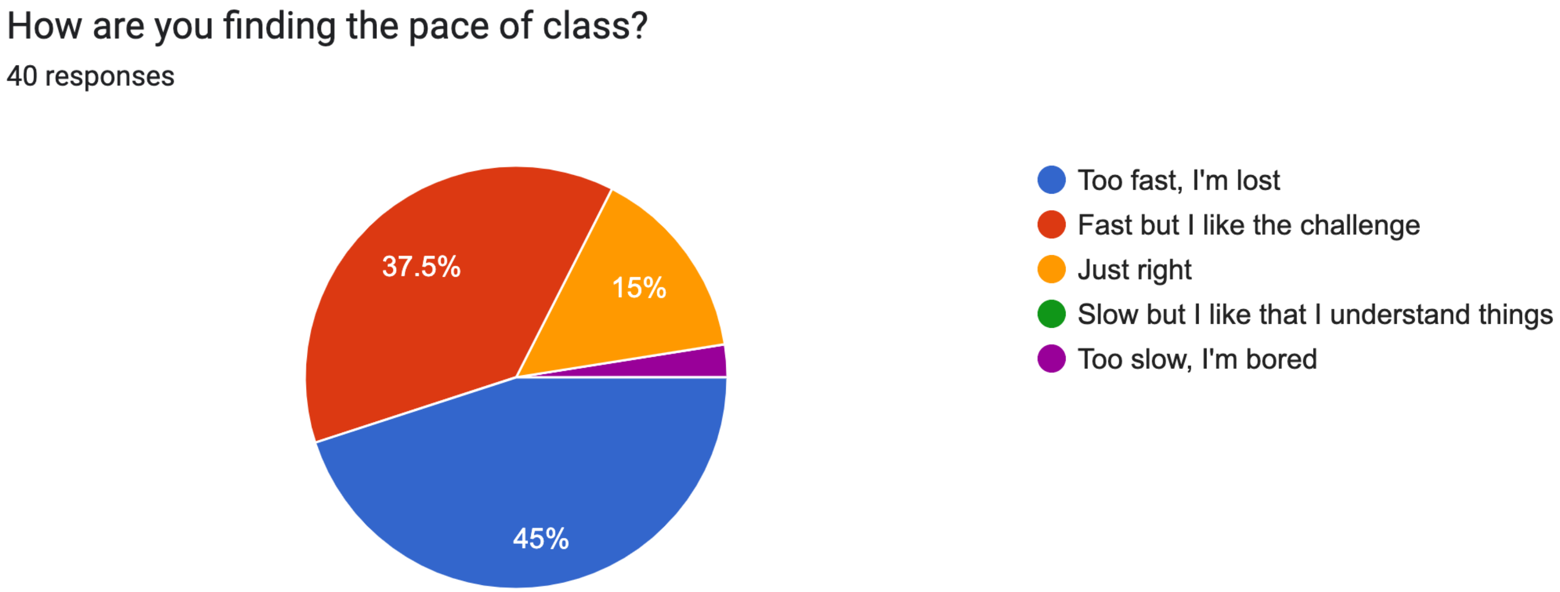

- I will try to slow down in class!
  - For today's class, we'll go back to go over some previous material folks asked for.

Other things that came up frequently

- The class/HW takes a lot of time.
  - Noted! HW5 should be a bit lighter.
- The quizzes take a lot of time.
  - That is not our intention! If you are struggling with the quizzes, PLEASE ask about them in OH or on Ed (privately). You can also discuss them with your HW group.
- Can we have more practice problems?
  - Check out the problems in “Probability and Computing” by Mitzenmacher and Upfal. The course staff are happy to chat about them with you.
  - We’ll have a few more today ☺

Note: there were several other things that came up, but not frequently. We did read them and we’ll do our best!

Recap

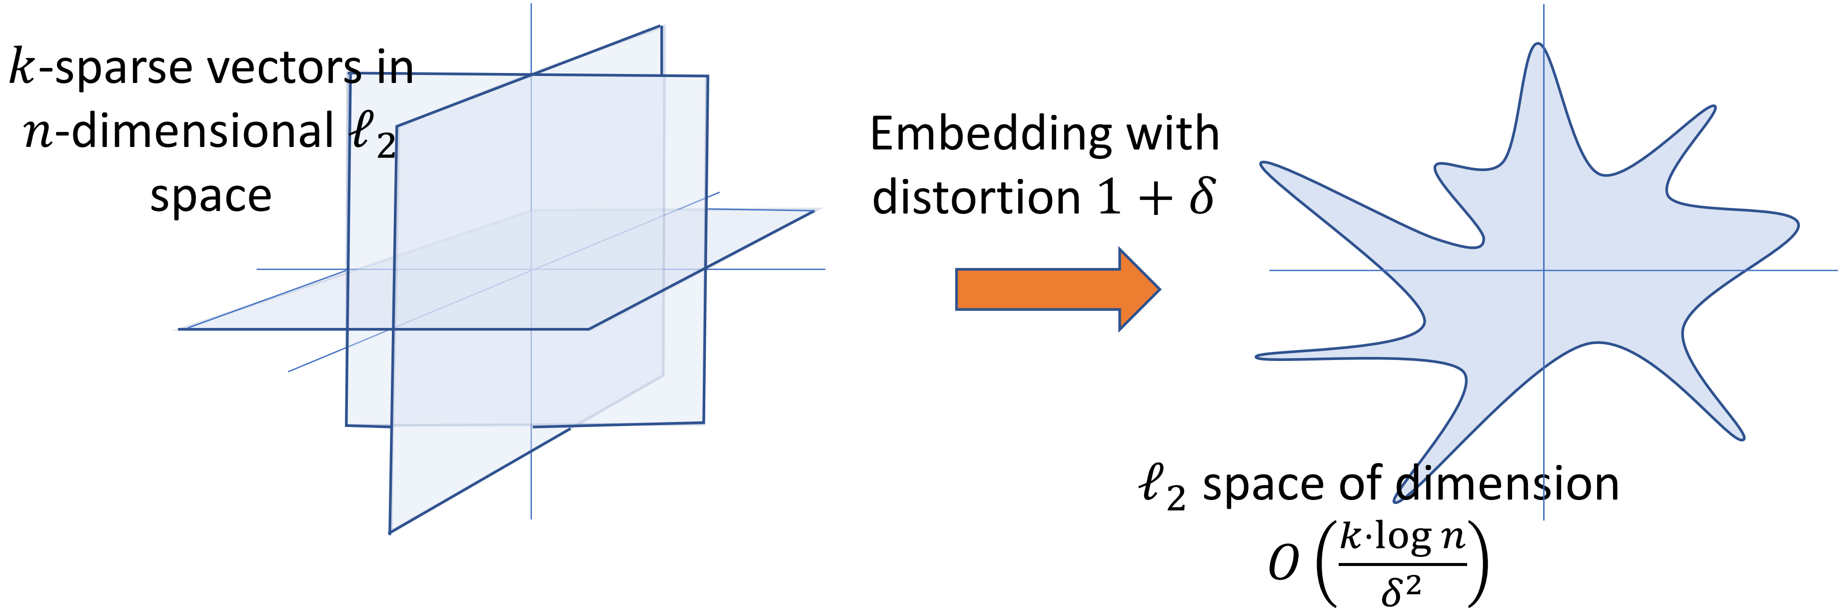

- Restricted Isometry Property: Embedding sparse vectors in low-dimensional space.
- A matrix $A \in \mathbb{R}^{m \times n}$ has the **Restricted Isometry Property** (RIP) with parameters k, δ if for any $x \in \mathbb{R}^n$ with $||x||_0 \leq k$,

$$(1 - \delta)||x||_2 \leq ||Ax||_2 \leq (1+\delta)||x||_2$$

- Random Gaussian Matrices have the RIP whp! ← Application: "Compressed Sensing"

Questions?

Minilectures, quiz?

Q1 Does it have the RIP?

6 Points

Which of the following matrices have the RIP with high probability (say, at least 0.9)?

Below, if you check "yes," the matrix should have the (k, ϵ)-RIP with high probability for all sufficiently large n; for all k with $k = \omega(1)$ and $k \leq \sqrt{n}$; and for $\epsilon = 0.1$.

Q1.1

2 Points

The identity matrix $I \in \mathbb{R}^{n \times n}$.

- (selected) Yes, this has the RIP with high probability.
- ( ) No, this does not have the RIP with high probability.

Q1.2

2 Points

A matrix $A \in \mathbb{R}^{m \times n}$ so that each entry of A is an independent $N(0, 1/m)$ random variable, where $m \geq Ck \log n$, for an appropriately large constant C.

- (selected) Yes, this has the RIP with high probability. (For large enough C).
- ( ) No, this does not have the RIP with high probability. (No matter how large C is).

Q1.3

1 Point

The matrix A from part 1.2, except that one of the columns (say, the first one) has been replaced with all zeros.

- ○ Yes, this has the RIP with high probability.
- ◉ No, this does not have the RIP with high probability.

Q1.4

1 Point

The matrix A from part 1.2, except that one of the rows (say, the first one) has been replaced with all zeros.

- ◉ Yes, this has the RIP with high probability.
- ○ No, this does not have the RIP with high probability.

Q2

6 Points

Let $B_p^n(\rho) = \{x \in \mathbb{R}^n \, : \, \|x\|_p \leq \rho\}$ denote the ℓ_p ball of radius ρ in $\mathbb{R}^n$.

Suppose that $0 < \epsilon < 1/4$.

Which of the following is a true statement?

Q2.1

3 Points

You can cover $B_2^n(1)$ with $(3/\epsilon)^n$ copies (appropriately shifted) of $B_2^n(\epsilon)$.

- (•) True
- ( ) False

Q2.2

3 Points

You can cover $B_1^n(1)$ with $(3/\epsilon)^n$ copies (appropriately shifted) of $B_1^n(\epsilon)$.

The difference between this one and Q2.1 is that it's ℓ_1 balls instead of ℓ_2 balls. In case it matters, we have $\mathrm{Vol}(B_1^n(\rho)) = C_n' \rho^n$ for some constant C_n', just like for ℓ_2 balls.

- (•) True
- ( ) False

Plan for today

- Slow down and recap!

More comprehensive recap

OMG it's week 5 already -- what have we seen so far? ← LOTS!

ALGORITHMS / APPLICATIONS

- Polynomial Identity Testing
- Karger's alg for min cuts
- Primality testing
- Sampling-based Median algorithm
- Randomized routing on the hypercube
- Load balancing / power of 2 choices
- Finding (appx) sparsest cuts
- Approximate Nearest Neighbors

TECHNIQUES

- Linearity of $\mathbb{E}$
- Writing things as sums of indicator r.v.'s
- Markov's inequality
- Chebyshev's inequality
- Chernoff bounds + MGFs
- "Poissonization"

METRIC EMBEDDINGS (both "technique" and "application")

- Bourgain's embedding
- Johnson-Lindenstrauss lemma
- Locality-Sensitive Hashing
- Compressed Sensing / RIP

Let's practice tail bounds!

When to use what bound?

$$\zeta(k) := \sum_{i=1}^{\infty} i^{-k}$$

$$\mathbb{P}[X=i] = \frac{i^{-3}}{\zeta(3)} \quad \text{for } i=1,2,\ldots$$

Markov?

Chebyshev?

Other?

Let's practice tail bounds!

When to use what bound?

$\zeta(k) := \sum_{i=1}^{\infty} i^{-k}$

$$\mathbb{P}[X = i] = \frac{i^{-4}}{\zeta(4)}$$

Markov?

Chebyshev?

Other?

$\zeta(k) := \sum_{i=1}^{\infty} i^{-k}$

Let's practice tail bounds!

When to use what bound?

$$\mathbb{P}[X=i] = i^{-k}/\zeta(k)$$

Markov?

Chebyshev?

Other?

Let's practice tail bounds!

When to use what bound?

$X_i \sim Ber(p)$ iid. $X = \sum_i X_i$.

How big does ε need to be to show...

	Chernoff (mult)	Chernoff (add.)	Chebyshev
$p = 1/4$, $\mathbb{P}[X \geq \mu(1+\varepsilon)] \leq 0.001$?			
$p = 1/\sqrt{n}$, $\mathbb{P}[X \geq \mu(1+\varepsilon)] \leq 0.001$?			
$p = 1/4$ $\mathbb{P}[X \geq \mu(1+\varepsilon)] \leq \delta$?			

Let's practice tail bounds!

When to use what bound?

$X_i \sim Ber(p)$ iid. $X = \sum_i X_i$.

How big does ε need to be to show…

	Chernoff (mult)	Chernoff (add.)	Chebyshev
$p = 1/4$, $\mathbb{P}[X \geq \mu(1+\varepsilon)] \leq 0.001$?	$\varepsilon = \Theta(1/\sqrt{n})$	$\varepsilon = \Theta(1/\sqrt{n})$	$\varepsilon = \Theta(1/\sqrt{n})$
$p = 1/4$ $\mathbb{P}[X \geq \mu(1+\varepsilon)] \leq \delta$?	$\varepsilon = \Theta\left(\sqrt{\frac{\log(1/\delta)}{n}}\right)$	$\varepsilon = \Theta\left(\sqrt{\frac{\log(1/\delta)}{n}}\right)$	$\varepsilon = \Theta\left(\frac{1}{\sqrt{\delta n}}\right)$ (much worse!)
$p = 1/\sqrt{n}$, $\mathbb{P}[X \geq \mu(1+\varepsilon)] \leq 0.001$?	$\varepsilon = \Theta\left(1/n^{1/4}\right)$	$\varepsilon =$ some constant? [if it works at all…]	$\varepsilon = \Theta\left(1/n^{1/4}\right)$

Let's practice Poissonization!

How to set up the argument structure

n balls into n bins

$X_i = \mathbb{1}\{\text{bin } i \text{ has} \le 2 \text{ balls}\}$

$X = \sum_{i=1}^{n} X_i$

$\mathbb{P}[X \ge 0.999n] \le$ ______

should this be $>n$ or $<n$?

Use Poissonization! Say $k \sim \text{Poi}($______$)$, drop k balls into n bins.

Let's practice Poissonization!

How to set up the argument structure

n balls into n bins

$Y_i = \mathbb{1}\{\text{bin } i \text{ has } > 2 \text{ balls}\}$

$Y = \sum_i Y_i$

$\mathbb{P}[Y > 0.3n] \leq$ ____________

should this be $> n$ or $< n$?

Use Poissonization! Say $k \sim \text{Poi}($______$)$, drop k balls into n bins.