Interacting with your Research Community

Pravesh Kothari (CMU)

Based on Matt Weinberg and Nicole Immorlica's Slides

Disclaimer

1. Disclaimer: Like all advice, this one comes with no warranty. And… also subjective. And… somewhat colored by TCS. And… if you reject most of it, you are probably right. ;-)

2. Meta-Advice: Always understand **why.**
   1. **Do** look to others for examples of good talks/websites/papers/etc.
   2. **Don't** try to "follow the rules" without understanding why.

3. Might be more useful if we keep it interactive!

outline.

Several aspects to your professional interactions.

1. Papers you write.
2. Reviews you write.
3. Talks/Posters you present.
4. Your website.

This is a non-exhaustive list. This talk presents some of our thoughts on each of these.

Suggestions for Writing

Main goals:

- Engage, entertain and convince the reader of the importance and coolness of your work.
- Present correct, complete, and verifiable proofs.

Everything else is secondary.

- I really mean this! View everything through this lens.
- No “formatting requirements” (although norms exist).
- Next several slides break down how we normally try to convince the reader our work is cool.
- **But these are not hard/fast rules. Ask why!**

Suggestions for Writing

Highly recommend watching Larry McEnerney's talk on academic writing!

Outline of a paper.

1. Title: phrase summarizing the work
2. Abstract: concise description of the work
3. Introduction: set context, motivate and state key

Warning: This is my preference. Not a prescription for your paper's outline.

5. Related Work: how your work fits with the literature?
6. Overview: highlight main idea of your proof, use as a roadmap for how the paper goes

Body!

7. Appendix: proofs of easy observations, proofs that are "believable" but tedious and unilluminating, minor extensions…

Writing Principles

Contents

title.

Purpose: Briefly indicate why someone might be interested in reading your paper.

Example: Is it OK to be catchy/funny?

- Sure! If it is also concise and descriptive.
- *Is it Easier to Prove Statements that are Guaranteed to be True?* Pass and Venkatasubramaniam FOCS'20.
- *The Log-Approximate-Rank Conjecture is False* Chattopadhyay, Mande, Sherif, STOC'19

title.

Good	Bad
Descriptive (but brief)	Vague

title.

Bad:

Information Aggregation in Social Networks, Feldman, Immorlica, Lucier and Weinberg, 2014, working paper.

Good:

Reaching Consensus via non-Bayesian Asynchronous Learning in Social Networks, Feldman, Immorlica, Lucier and Weinberg, APPROX 2014.

title.

Purpose: Briefly indicate why someone might be interested in reading your paper.

Example: Why is vague bad?

- *Information Aggregation in Social Networks.*
- Is it a theory paper? Empirical study?
- No idea what makes it different from the 100,000 other papers on information aggregation in social networks.

abstract.

Purpose: Normally targeted at an expert. Should they should read more? *Think about people reading ~20 abstracts in their arxiv digest at 11pm…these are experts looking to quickly see if they should read the paper.*

Example: Should I "sell" the main results?

- You should state objectively **why the main result is interesting**, so an expert knows what's the point.
- You shouldn't go overboard, you have an entire introduction for that.
- E.g. "This is the first constant-factor approximation."

abstract.

Purpose: Normally targeted at an expert. Should they should read more? *Think about people reading ~20 abstracts in their arxiv digest at 11pm…these are experts looking to quickly see if they should read the paper.*

Example: Should I “define” the problem in abstract?

- I have done it. Especially when it’s a new problem.
- It’s hard but important, IMO.

abstract.

Good	Bad
concise	wordy
fact-based	salesmanship
accurate	over-claiming
identifies keywords	inaccurate terminology

Pet Peeve: abstracts that are really introductions.

abstract (bad).

In a social learning setting, members of a society share their experiences to help others make better choices. Following the established path can boost an individual's utility but it can hurt the society as a whole since other options of higher value may never be explored. We show that when the population is diverse, this issue can be avoided as people may not be satisfied with the available choices and look for alternatives. High diversity, though, comes at a cost as past experiences become less valuable.

abstract (bad).

We model these situations in a standard setting of consumer search introduced by Weitzman and study how different diversity levels compare with each other. We … and quantify how the socially optimal diversity level changes …. Moreover, while high diversity can lead to anarchy and confusion in typical situations, we show that it can be really beneficial in settings where society may accidentally uncover a unanimously accepted hidden gem.

abstract (reasonable).

Abstract

We give the first polynomial-time algorithm for robust regression in the list-decodable setting where an adversary can corrupt a greater than 1/2 fraction of examples.

For any $\alpha < 1$, our algorithm takes as input a sample $\{(x_i, y_i)\}_{i \leqslant n}$ of n linear equations where αn of the equations satisfy $y_i = \langle x_i, \ell^* \rangle + \zeta$ for some small noise ζ and $(1-\alpha)n$ of the equations are *arbitrarily* chosen. It outputs a list L of size $O(1/\alpha)$ - a fixed constant - that contains an ℓ that is close to ℓ^*.

Our algorithm succeeds whenever the inliers are chosen from a *certifiably* anti-concentrated distribution D. In particular, this gives a $(d/\alpha)^{O(1/\alpha^8)}$ time algorithm to find a $O(1/\alpha)$ size list when the inlier distribution is standard Gaussian. For discrete product distributions that are anti-concentrated only in *regular* directions, we give an algorithm that achieves similar guarantee under the promise that ℓ^* has all coordinates of the same magnitude. To complement our result, we prove that the anti-concentration assumption on the inliers is information-theoretically necessary.

Our algorithm is based on a new framework for list-decodable learning that strengthens the "identifiability to algorithms" paradigm based on the sum-of-squares method.

In an independent and concurrent work, Raghavendra and Yau [RY19] also used the Sum-of-Squares method to give a similar result for list-decodable regression.

introduction

1 Introduction

In this work, we design algorithms for the problem of linear regression that are robust to training sets with an overwhelming ($\gg 1/2$) fraction of adversarially chosen outliers.

Outlier-robust learning algorithms have been extensively studied (under the name *robust statistics*) in mathematical statistics [Tuk75, MMY06, Hub11, HRRS11]. However, the algorithms resulting from this line of work usually run in time exponential in the dimension of the data [Ber06]. An influential line of recent work [KLS09, ABL13, DKK+16b, LRV16, CSV17, KS17a, KS17b, HL17, DKK+17, DKS17, KKM18] has focused on designing *efficient* algorithms for outlier-robust learning.

Our work extends this line of research. Our algorithms work in the "list-decodable learning" framework. In this model, a majority of the training data (a $1-\alpha$ fraction) can be adversarially corrupted leaving only an $\alpha \ll 1/2$ fraction of "inliers". Since uniquely recovering the underlying parameters is information-theoretically *impossible* in such a setting, the goal is to output a list (with an absolute constant size) of parameters, one of which matches the ground truth. This model was introduced in [BBV08] to give a discriminative framework for clustering. More recently, beginning with [CSV17], various works [DKS18, KS17a] have considered this as a model of "untrusted" data.

There has been phenomenal progress in developing techniques for outlier-robust learning with a *small* ($\ll 1/2$)-fraction of outliers (e.g. outlier "filters" [DKK+16a, DKK+17, CDG19, DKK+18b], separation oracles for inliers [DKK+16a] or the *sum-of-squares* method [KS17b, HL17, KS17a, KKM18]). In contrast, progress on algorithms that tolerate the significantly harsher conditions in the list-decodable setting has been slower. The only prior works [CSV17, DKS18, KS17a] in this direction designed list-decodable algorithms for mean estimation via problem-specific methods.

In this paper, we develop a principled technique to give the first efficient list-decodable learning algorithm for the fundamental problem of *linear regression*. Our algorithm takes a corrupted set of linear equations with an $\alpha \ll 1/2$ fraction of inliers and outputs a $O(1/\alpha)$-size list of linear functions, one of which is guaranteed to be close to the ground truth (i.e., the linear function that correctly labels the inliers). A key conceptual insight in this result is that list-decodable regression information-theoretically requires the inlier-distribution to be "anti-concentrated". Our algorithm succeeds whenever the distribution satisfies a stronger "certifiable anti-concentration" condition that is algorithmically "usable'. This class includes the standard gaussian distribution and more generally, any spherically symmetric distribution with strictly sub-exponential tails.

Prior to our work[1], the state-of-the-art outlier-robust algorithms for linear regression [KKM18, DKS19, DKK+18a, PSBR18] could handle only a small (< 0.1)-fraction of outliers even under strong assumptions on the underlying distributions.

List-decodable regression generalizes the well-studied [DV89, JJ94, FS10, YCS13, BWY14, CYC14, ZJD16, SJA16, LL18] and *easier* problem of *mixed linear regression*: given k "clusters" of examples that are labeled by one out of k distinct unknown linear functions, find the unknown set of linear functions. All known techniques for the problem rely on faithfully estimating certain *moment tensors* from samples and thus, cannot tolerate the overwhelming fraction of outliers in the list-decodable setting. On the other hand, since we can take any cluster as inliers and treat rest as outliers, our algorithm immediately yields new efficient algorithms for mixed linear regression. Unlike all prior works, our algorithms work without any pairwise separation or bounded condition-number assumptions on the k linear functions.

List-Decodable Learning via the Sum-of-Squares Method Our algorithm relies on a strengthening of the robust-estimation framework based on the sum-of-squares (SoS) method. This paradigm has been recently used for clustering mixture models [HL17, KS17a] and obtaining algorithms for moment estimation [KS17b] and linear regression [KKM18] that are resilient to a small ($\ll 1/2$) fraction of outliers under the mildest known assumptions on the underlying distributions. At the heart of this technique is a reduction of outlier-robust algorithm design to just finding "simple" proofs of unique "identifiability" of the unknown parameter of the original distribution from a corrupted sample. However, this principled method works only in the setting with a small ($\ll 1/2$) fraction of outliers. As a consequence, the work of [KS17a] for mean estimation in the list-decodable setting relied on "supplementing" the SoS method with a somewhat problem-dependent technique.

As an important conceptual contribution, our work yields a framework for list-decodable learning that recovers some of the simplicity of the general blueprint. Central to our framework is a general method of *rounding by votes* for "pseudo-distributions" (see Section 2) in the setting with $\gg 1/2$ fraction outliers. Our rounding builds on the work of [KS19] who developed such a method to give a simpler proof of the list-decodable mean estimation result of [KS17a].

Prior results discussed above hold for any underlying distribution that has upper-bounded low-degree moments and such bounds are "captured" within the SoS system. Such conditions are called as "certified bounded moment" inequalities. An important contribution of this work is to formalize *anti-concentration* inequalities within the SoS system and prove such inequalities for natural distribution families. Unlike bounded moment inequalities, there is no canonical encoding within SoS for such statements. We choose an encoding that allows proving certified anti-concentration for a distribution by showing the existence of a certain approximating polynomial. This allows showing certified anti-concetration via a modular approach relying on a beautiful line of works that construct "weighted " polynomial approximators [Lub07].

We believe that our framework for list-decodable estimation and our formulation of certified anti-concentration condition will likely have further applications in outlier-robust learning.

1.1 Our Results

We first define our model for generating samples for list-decodable regression.

[1] There's a long line of work on robust regression algorithms (see for e.g. [BJKK17, KP19]) that can tolerate corruptions only in the *labels*. We are interested in algorithms robust against corruptions in both examples and labels.

introduction

Purpose: Serves a few purposes, tricky to balance.

- Convincingly explain the context and importance of the question you are studying
- Effectively communicate why the paper is cool.

Example: Should I “sell” the main results?

- Absolutely! Don’t be afraid to tell the reader exactly why it’s cool, new and required new insights.
- (From a selfish perspective: don’t be afraid to tell the *reviewer* exactly what they should state as the “main contributions” in their review).

introduction.

Good	Bad
motivation from practice or existing literature	flimsy stories or cartoon realities
place results in context	abstracts of related work
identify take-aways and key intuition	overly-precise statement of results and techniques

Pet Peeve: Laundry lists of results with no motivation.

Results

Purpose: Provide theorem statements of main result(s).

- Can be somewhat informal if needed.
- May want to state it in a model without many bells and whistles if needed.
- Discuss how the theorem relates to known results, briefly state (and include a pointer to detailed discussion in overview) the key idea(s).
- A non-expert will likely read your intro before all.
- An expert will likely read your statement of results before all.

Pet Peeve: Hard to find statement of main results.

Results

Purpose: Provide theorem statements of main result(s).

- Can be somewhat informal if needed.
- May want to state it in a model without many bells and whistles if needed.
- Discuss how the theorem relates to known results, briefly state (and include a pointer to detailed discussion in overview) the key idea(s).

Pet Peeve: Hard to find statement of main results.

Results

Purpose: Provide theorem statements of main result(s).

- Tricky to get the right level of informality.
- But important.

- An expert will likely read your statement of results before all. I'd open this section if an abstract catches my eye in my arxiv digest…

Pet Peeve: Hard to find statement of main results.

Related work.

Purpose: Provide context for your work.

- Most related stuff (ideally) already covered in intro.
- So this is largely to assign scientific credit for prior work.
- It might also be “works that might look related but really aren’t” at times.

related work.

Good	Bad
comprehensive	skimpy
describes connections	reads like a list
cites work from multiple fields	unaware of related literature

Pet Peeve: sections that read like a list of abstracts!

Related work.

Purpose: Provide context for your work.

- Most related stuff (ideally) already covered in intro.
- Also to assign scientific credit for prior work.

Example: How much detail should I give?

- Enough to make your point!
- Ex: “Cai and Daskalakis give a PTAS for a single unit-demand buyer with independent MHR item values, to the optimal deterministic item pricing.”
- Useful if you give a PTAS for a related problem.
- **Not** useful just because you study pricing.

Related work.

Purpose: Provide context for your work.

- Most related stuff (ideally) already covered in intro.
- Also to assign scientific credit for prior work.

Example: How much detail should I give?

- Enough to make your point!
- Ex: “Works such as [CaiD11, …] also provide approximations in different models to ours.”
- Useful if you study unrelated pricing problem.
- **Not** enough if reviewer might reasonably wonder what your work contributes over CaiD11.

Related work.

Purpose: Provide context for your work.

- Most related stuff (ideally) already covered in intro.
- Also to assign scientific credit for prior work.

Example: “concurrent/independent” work?

- Be as transparent as possible.
- Mention independent work as early as possible (I do it in the abstract itself).
- Run your comments/discussion on the concurrent work by their authors before posting publicly if possible.

Model.

If your model details take a para, include it before results.
If not, include it in a separate section.

Purpose: Should be formal.

- Most intuition (ideally) already given in intro.
- Need to be precise, but also clear.

model.

Good	Bad
notation consistent with existing norms	overloaded or excessive notation
covers limited prelim results/background	contains major theorem statements and proofs
rigorous, yet clear	overly formal, or imprecise

Model.

Purpose: Should be formal.

- Most intuition (ideally) already given in intro.
- Need to be precise, but also clear.

Example: Should I give an example?

- Sure! If it serves a purpose.
- If complicated definitions, to illustrate definitions.
  - If simple, no point.
- To illustrate subtle counter-intuitive properties.
- To illustrate special case of main proof ideas.

Overview.

Purpose: present an intuitive roadmap for your proof with several steps and ideas.

- Can be slightly informal (but not too much).
- Instead, simplify by perhaps work with a special case or discussing a proof plan and focusing only on one or two most novel/insightful steps.
- In theory papers, this section is supposed to be the main “long-term” take-away of a close-to-expert reader.
- Overview $\approx$ 35 mins of your technical talk.

Overview.

Purpose: present an intuitive roadmap for your proof with several steps and ideas.

Tip 1: Present an essentially complete proof of some insightful new result-lite!

Not always possible. But great when it is!

Tip 2: ***Don't*** include a long overview if your proof is simple enough to follow as is.

overview.

Good	Bad
Discusses a couple of your “aha” moments. points out natural wrong turns – “you may think X, but it fails because…” Explains your innovation in techniques and compares to older ones.	List of buzzwords without clearly explained conn. to your proof. Includes only boring proofs or proof ideas without intuition.

Body

Purpose: Present complete proofs of your theorem.

- Present proof top down – Main theorem follows from Steps 1, 2, 3 – presented in sections 4,5,6. Step 1 follows from lemmas….
- Delegate believable but tedious proofs to appendix.
- Always have the main goal of a section (thm,lem,…) stated right at the beginning.

Example: Include a tedious/subtle proof in the body?

- Could it be surprising to an expert? Then, yes.
- Mark that subsection and suggest a reader to “skip” in the first reading. Point out the subtlety if you can.

Body

Purpose: Present complete proofs of your theorem.

- Present proof top down – Main theorem follows from Steps 1, 2, 3 – presented in sections 4,5,6. Step 1 follows from lemmas….
- Delegate believable but tedious proofs to appendix.
- Always have the main goal of a section (thm,lem,…) stated right at the beginning.

Example: Why?

We are mathematical researchers. Our mathematics should always be complete, verifiable and convincing.

Results/Proofs.

Guiding Principle: **The entire body should “flow”.** It should be organized in a way that a reader can put the map of your proof in their mind. This is often *not* the way you came up with the proof/ideas.

If ideas too complex, distill main digestible aspects at the beginning.

Appendix.

Purpose: Believable but tedious proofs, extensions, proofs of observations used to make a point in the intro

- Yes, very few people will read the appendix, but…
- Important…I have cited several papers because I needed lemmas from their appendices.

Style suggestions:

- Appendix should be **easy to read.**
- It does **not** need to be engaging/stellarly written.
- Appendix should be **easy to navigate**.

Huge pet peeve: appendices with serious errors which were obviously never proofread.

appendix.

Good	Bad
clean to follow, but perhaps not engaging	unreadable
Relevant statements being proved are available in the appendix itself.	disorganized

Suggestions for Reviewing

impact.

You will be judged by the quality of your review.

People who read your review: Basically everyone you're meeting through this workshop (through the PC).

Every visibility opportunity counts! For some PC members, this may be the first time they see your name. Make a good first impression!

purpose.

Your job is not to directly decide whether to accept/reject the paper.

Your job is to give arguments/evidence/information to the PC so they can decide whether to accept/reject.

Parse remaining advice in this context.

content.

Briefly describe main results. Should contain enough context to explain why authors think its exciting.

Thought experiment: Would the authors agree?
This is not the place to disagree with authors.

Excellent: So clear that PC doesn't need re-read intro.

Bad: PC needs to read entire intro, ignore your summary.

content.

List major concerns. Is there an error in a proof, is the result a trivial generalization of existing work? (rare)

Excellent: "I might be misunderstanding something, but as stated, it seems that Theorem 2 is false. Here is a sketch of a counterexample. Is it possible that the authors meant to place additional assumptions?"

Point: You may be about to kill the paper. If it needs to be done, it needs to be done. But be thoughtful.

content.

List major concerns. Is there an error in a proof, is the result a trivial generalization of existing work? (rare)

Bad: “Theorem 2 is false, integral might diverge.” (Often resolved by: “It is easy to verify that if integral diverges, results still hold with notational updates.”).

Point: Totally valid minor concern to be rigorous with divergence (and this should be raised). But don’t kill papers for oversights which can be easily resolved.

content.

Form subjective evaluation. Quality of results, innovation in techniques, quality of presentation. This is the time to disagree with the authors!

Results: (In your opinion)

- What makes the results significant (or not)?
- Is there any context that the PC needs to appreciate?

content.

Form subjective evaluation. Quality of results, innovation in techniques, quality of presentation. This is the time to disagree with the authors!

Results: Rough scale to have in mind:

- So strong it doesn’t matter how it’s proved.
- Strong.
- Motivated enough if the techniques are awesome.
- Extremely specialized or toy.

content.

Form subjective evaluation. Quality of results, innovation in techniques, quality of presentation.
This is the time to disagree with the authors!

Techniques: (In your opinion)

- Will they help you/others solve problems?
- Did you find them engaging/illuminating?
- **Don't ask:** "How much hard work?"
- Trivial proofs are bad because aren't engaging, don't help others, **not** because they're not hard enough.
- But simple, engaging, thoughtful proofs are great!

content.

Form subjective evaluation. Quality of results, innovation in techniques, quality of presentation.
This is the time to disagree with the authors!

Techniques: Rough scale to have in mind:

- Super interesting/insightful, loved reading it even if results meh.
- A strength.
- Enjoyed reading, but not a strength.
- Trivial or entirely (hard but) tedious calculations.

content.

Form subjective evaluation. Quality of results, innovation in techniques, quality of presentation.
This is the time to disagree with the authors!

Presentation: (In your opinion)

- Were the stated results completely proved?
- Do you understand everything you want to?
- Do you understand what the authors want you to?

content.

Form subjective evaluation. Quality of results, innovation in techniques, quality of presentation.
This is the time to disagree with the authors!

Presentation: Rough scale to have in mind

- Truly a pleasure to read.
- Fine.
- I was truly miserable reading. (could kill paper).

content.

Briefly describe main results. Should contain enough context to explain why authors think its exciting.

List major concerns. Is there an error in a proof, is the result a trivial generalization of existing work? (rare)

Form subjective evaluation. Quality of results, innovation in techniques, quality of presentation.

Recommendation: Accept? Enjoy it? **Why or why not**?

Suggestions for Talks

giving talks

"There is only one rule for giving talks: *know your audience.*"

Avi Wigderson

additional disclaimer.

People disagree a LOT on ”principles” for a good talk.

- **But**, goals are universal:
  - **Your goal is to engage the audience**.
  - Everything else is just subjective tips to achieve this.
  - Constantly remind yourself of this during prep.
  - (And prep a lot!).

I have given complete talks to myself.
Sometimes writing what you want to say helps.
I have rehearsed even the “jokes” at times.

typical outline of a long tech talk.

1. Problem (and rarely – if you get lucky -- results!): may want to simplify!
2. Context/Motivation: context for your work
3. Results: what you did differently, and why
4. Key Takeaways: technical innovation? Connections?--

   -----End Overview Section of your talk ---
5. Model: the setting you consider
6. Proof Plan: Remember your overview section?
7. Conclusion: restate main take-aways, use “hindsight”
8. (Future Directions): open questions

content.

Identify at most three main take-aways!

content.

Everything should have a point.

Time is short, and brain capacity is scarce.

- Identify at most 3 main take-aways, e.g., a new objective, model, technique, practical insight, etc.
- State take-aways explicitly and repeat them often.
- Think of your technical slides as teaching the coolest bit or two in the paper! Keep it engaging, surprising and if possible, get an “aha” moment.

content.

Everything should have a point.

Thought experiment: for every slide, “what is the purpose of this slide?” should have quick answer.

Bad answer: “to cover related work.”
Bad answer: “to list open questions.”

Good answer: “convey that related work covers n=2.”
Good answer: “convey that model has rich questions.”

design.

Keep it simple!

design.

Fonts: Use a single font.

Fonts should come in three sizes.

Titles, Main Text, Figures

Choosing a minimum font size helps prevent you from trying to stuff too much stuff on a single slide, which overwhelms your audience and makes them squint.

design.

Colors: Use colors wisely and sparingly.

Each color should have a meaning, e.g.,

- Title of slide
- Type of statement: theorem, proof, etc.
- Emphasis: word being defined, main point

design (Nicole).

Text: White space is an undervalued asset.

Break paragraphs into small sensible chunks, use line breaks between chunks instead of bullets. Don’t sacrifice grammar if you don’t have to. Complete sentences are appreciated. So is spelling. Don’t be too wordy. Say only what’s essential. Remember, you can say things in words, not everything needs to be written down.

design (Matt, Pravesh agrees).

Text: White space is an undervalued asset.

Sacrifice grammar if saves space.
Complete sentences unnecessary (but spelling important).
Don't be too wordy. Say only what's essential.
Helps to give each idea its own line, when possible.

Match everything you say out loud to **something** on slide.
- Could be short bullet. Could be figure.
- Some listeners distracted for 30 seconds and want to jump back in. Can't rewind speech, but text remains.

design (Others).

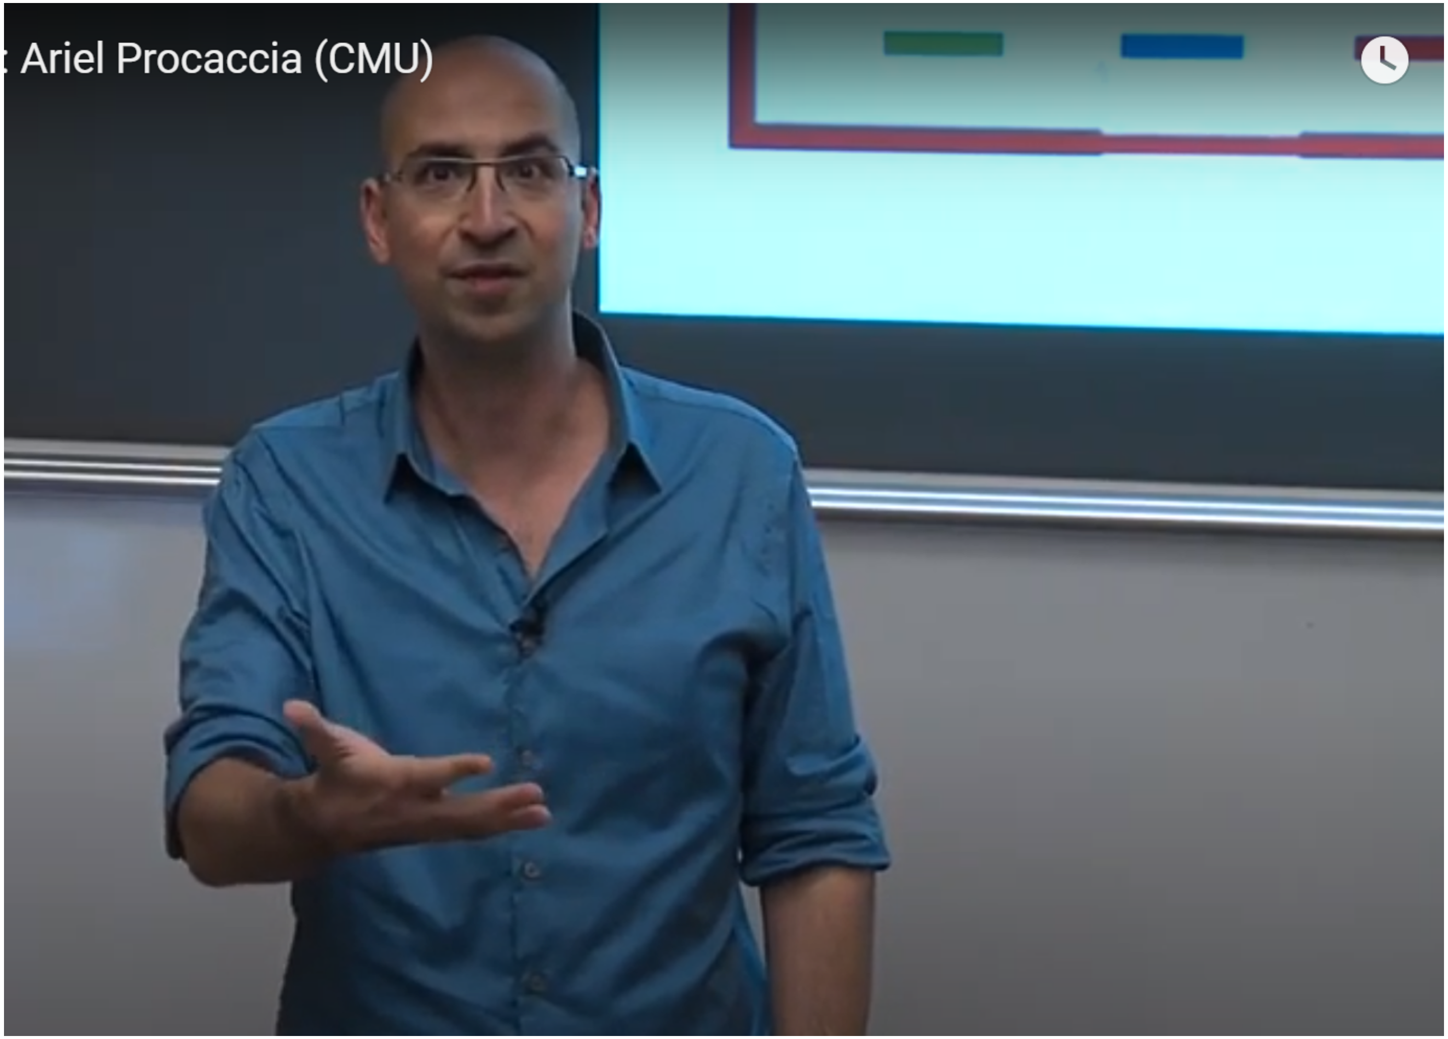

design.

Animations: Can be helpful, don't go overboard.

A full page at once can be daunting.
Hard for audience to focus.
Good to break up arrival into connected chunks.

But per-line animation gets annoying.
Especially if there's any clicker lag.
Forces all audience to go exactly at your speed.

design.

Math: Minimize mathematical notation.

Mechanism ϕ inputs ordinal preferences $\succ$ and outputs a matching $\phi(\succ)$. Let $r_x\big(\phi(\succ)\big)$ be the rank of the match of agent x under preferences $\succ_x$.

Definition. Mechanism ϕ' stochastically dominates mechanism ϕ for agent x if for all ranks k,

$$\Pr[r_x(\phi'(\succ)(x)) \le k] \ge \Pr[r_x\big(\phi(\succ)(x)\big) \le k].$$

design.

Math: Minimize mathematical notation.

Mechanism ϕ inputs ordinal preferences $\succ$ and outputs matching $\phi(\succ)$.

Definition. Mechanism ϕ' stochastically dominates mechanism ϕ for agent x if for all ranks k and all $\succ$,

$$\Pr[x \text{ gets top k in } \phi'(\succ)] \geq \Pr[x \text{ gets top k in } \phi(\succ)].$$

design.

Use images when appropriate. Example:

Construction: S← random set of size m/2.

- $T \leftarrow$ random set of size m/2, **conditioned on** $|\boldsymbol{S} \cap \boldsymbol{T}| = \boldsymbol{m}/\boldsymbol{3}$.
- Alice gets one special set A_0. Bob gets B_0.
- $|A_0| = |B_0| = m/2$.
- $|A_0 \cap S| = |B_0 \cap T| = m/3$.
- Case 1: $A_0 \cap B_0 = \emptyset$. Case 2: $A_0 = B_0$.

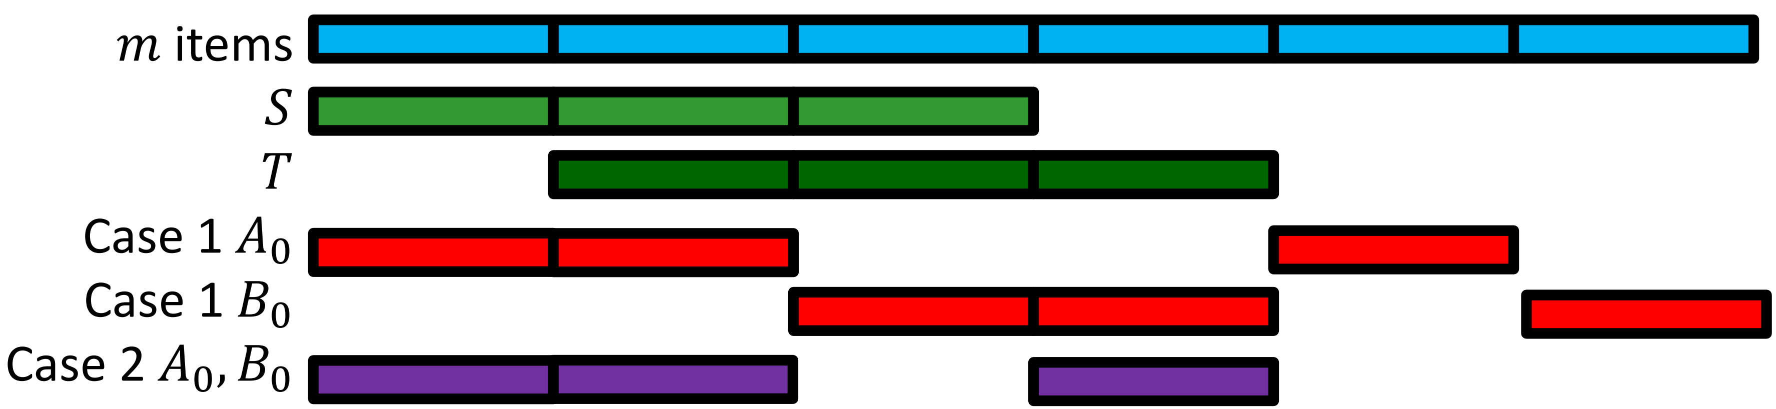

speaking.

Entertain your audience!

A talk IS performance art.

speaking.

Working the crowd: Your job is to please the audience!

- Always better to under-estimate your audience.
- Ask and answer questions, when appropriate.
- Adapt your pace to your read of the audience.

Voice: Be confident!

- Speak up, slow down/pause, and use intonation.
- Trick: take a breath after each slide.
- Have a concrete plan for important talking points (I sometimes write my speech for some slides).

Timing: NEVER EVER go over your allotted time!

- Plan what to skip if you are running short on time.

summary.

Important elements.

1. Identify three main take-aways (at most).

2. Keep it simple (design and content).

3. Entertain your audience.

Suggestions for your website

Why should you make a website?

- Visibility is important! Say you give a good talk/poster presentation and Prof. X notices. She wants to learn what year you are, where you are, or check out your other work. Make it feasible!

- Good to match norms. In CS, most senior PhD students have websites.

1. Picture

James R. Wright

Hello, I'm James Wright. I am a postdoctoral researcher at Microsoft Research in New York City. In July 2018, I will start as an Assistant Professor at the University of Alberta. I completed my Ph.D. at UBC in 2016, advised by Kevin Leyton-Brown.

2. Bio/Affiliation

Research

…rch interest is in using data-driven…
…; that is, behavior in interactions w…
…tions of other participants. My lon…
…ly designing algorithms for media…
…led agents rather than idealized, perfectly rational game theoretic agents.

Curriculum Vitae

My academic CV is available as both an HTML page and a PDF document. I also have a public Google Scholar citations page.

Contact

james.wright@ualberta.ca

3. Contact

Publications

1. Predicting Human Behavior in Unrepeated, Simultaneous-Move Games.
   James R. Wright and Kevin Leyton-Brown.
   Games and Economic Behavior, Volume 106, pages 16–37, November 2017.
   (supersedes Wright & Leyton-Brown [2010, 2012])
2. Learning in the Repeated Secretary Problem.
   Daniel G. Goldstein, R. Preston McAfee, Siddarth Suri, and James R. Wright.
   ACM Conference on Economics and Computation (ACM-EC), 2017.
3. Deep Learning for Predicting Human Strategic Behavior.
   Jason Hartford, James R. Wright, and Kevin Leyton-Brown.
   NIPS 20…
   Oral p…
4. Incenti…
   Worse. …
   Xi Alice …
   Worksh…
   Economics and Computation, 2016.

4. Research/Publications

Suggestions for your website

Purpose: introduce yourself. Use this to guide design.

Example: What if I don't have any publications yet?

- OK to omit. OK to put "coming soon!"
- People still want to know you without publications.

Example: Should I list unfinished manuscripts?

- One reason: excited about not-yet-published work.
- One reason: make research activity look larger.
    - Still: please don't *mislead* readers.

Suggestions for your website

Purpose: introduce yourself. Use this to guide design.

Example: Should I include personal information?

- Depends on what you want others to see!

Example: What if I'm bad at HTML? **(I (Pravesh) am.)**

- OK to copy a friend's source code (but ask them to avoid awkwardness…).
- OK to use generic (but clean) free layouts.
- People still want to know you if you're bad at HTML!

Suggestions for your website

Purpose: introduce yourself. Use this to guide design.

Example: Should I list awards?

- Sure! If you want others to know about them.

Last suggestion: Be transparent.

- It’s OK to brag a bit.
- OK to list unpublished work to appear more active.
- Just use same social norms: don’t be coy about it.

Appendix: How to Parse Feedback

Context: you'll constantly get feedback forever.

Some feedback is amazing and easy to parse.

- "I think you should do XYZ because ABC."
- And you completely agree immediately.

Most feedback is not. Especially when it's non-interactive from reviewers, course evals, etc.

- "The paper is OK I guess."
- But you can still get something out of it!

feedback examples.

Theme of advice: **Why** did this person give this feedback.

"I can't follow the proof of Theorem 1, why does Ax = b?"

- If you clearly stated why Ax = b, OK to complain! (I do).
- But don't **only** complain. Make it clearer.

"The results are fine, but incremental compared to [ABC]."

- If [ABC] is completely different, OK to complain! (I do).
- But adjust your next draft to better distinguish.

personal anecdote (Pravesh).

- I got my job talk reviewed by ~5 senior researchers.
  - ~50 minutes on my work on Sum-of-Squares.
- Drastically different takes: ”focus on algorithms”, “focus on only one result”, “display breadth”, plenty of comments on adding/removing details…
- Great advice: Getting ~5 opinions was like getting a sampling of what feedback might look like from people attending my talk. Made me aware of how it was perceived. And of course, I ended up making choices that agreed with some and disagreed with others.
- Point: Advice can be varying. But important to have an idea of how your talk is being perceived.

Disclaimer

1. Disclaimer: Like all advice, this one comes with no warranty. And… also subjective. And… somewhat colored by TCS. And… if you reject most of it, you are probably right. ;-)

2. Meta-Advice: Always understand **why.**
   1. **Do** look to others for examples of good talks/websites/papers/etc.
   2. **Don't** try to "follow the rules" without understanding why.