Figures & Tables
+ Research Career Paths

CS 197 & 197C | Stanford University | Sean & **Lauren**

cs197.stanford.edu | cs197c.stanford.edu

Slides adapted from previous iterations of the course by Michael Bernstein

HCIE Lab Tour

Overview

MIT CSAIL Alliances

Figures & Tables
+ Research Career Paths

CS 197 & 197C | Stanford University | Sean & **Lauren**

cs197.stanford.edu | cs197c.stanford.edu

Slides adapted from previous iterations of the course by Michael Bernstein

Items due this Thursday

197C

Research log

that's it!

197

Progress report III

that's it!

Items due next Thursday

197C

Research log

that's it!

197

Project Milestone

+ team dynamics form

Milestone deliverables:

1 page summary of your completed milestone goal + figure / table summarizing the results of your goal

Figures & Tables

+ Research Career Paths

Figures + Tables: what's the point?

- To convey the main takeaway of your work in a nice, visual easy-to-digest format

"We find that e taken to the power of i and pi simultaneously, when added to one, equals zero."

vs.

$$e^{i\pi} + 1 = 0$$

Euler's Identity
Schuhmann et. al. 2022
Ahuja et. al. 2021

the "most beautiful equation of all of math"

Figures + Tables: what's the point?

- To convey the main takeaway of your work in a nice, visual easy-to-digest format
- To summarize numerical data succinctly

"Of the public datasets available, the MS-COCO dataset possesses 330K images paired with text while the CC3M dataset contains 3M images paired with text, the Visual Genome…

vs.

Dataset	# English Img-Txt Pairs
Public Datasets	
MS-COCO	330K
CC3M	3M
Visual Genome	5.4M

Euler's Identity
Schuhmann et. al. 2022
Ahuja et. al. 2021

Figures + Tables: what's the point?

- To convey the main takeaway of your work in a nice, visual easy-to-digest format
- To summarize numerical data succinctly
- To visualize the framework / pipeline / system you're proposing

"In the CoolMoves pipeline, a VR-enabled user interacts with a virtual world and the input from the controllers is cross-linked with a large database of motions which accentuates and …

vs.

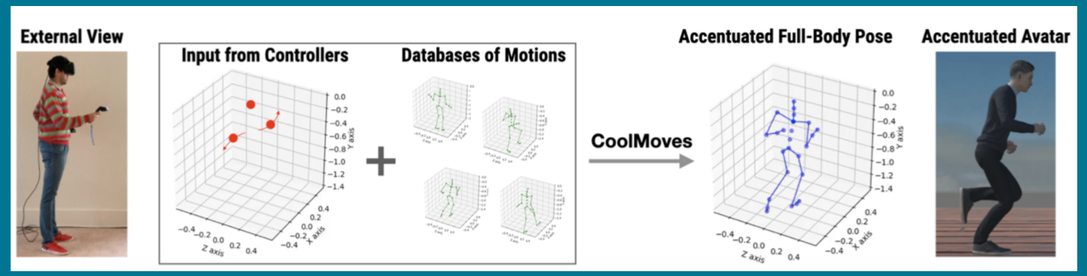

Euler's Identity
Schuhmann et. al. 2022
Ahuja et. al. 2021

Types of figures

1. Introductory
2.
3.
4.

Beery et. al. 2020
Schuhmann et. al. 2022
Kunin et. al. 2020
Ahuja et. al. 2021

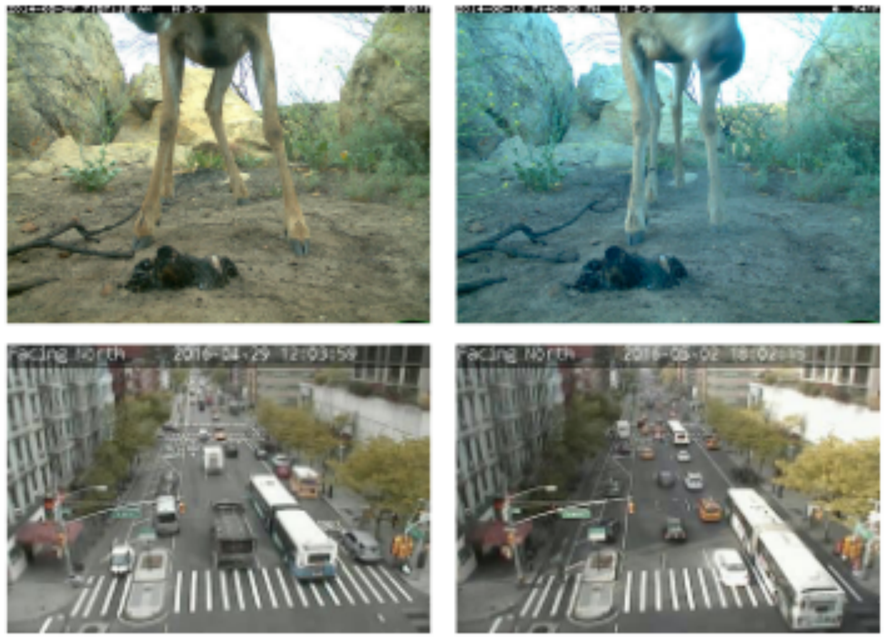

Figure 1: **Visual similarity over long time horizons.** In static cameras, there exists significantly more long term temporal consistency than in data from moving cameras. In each case above, the images were taken on separate days, yet look strikingly similar.

Types of figures

1. Introductory
2. Framework / pipeline / model
3.
4.

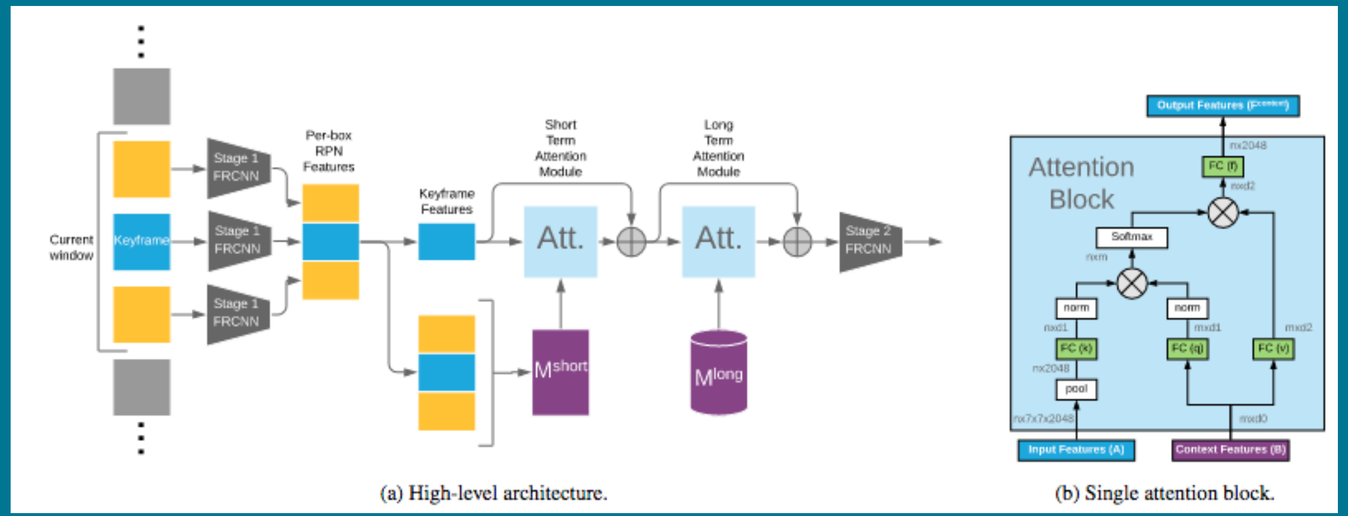

(a) High-level architecture. (b) Single attention block.

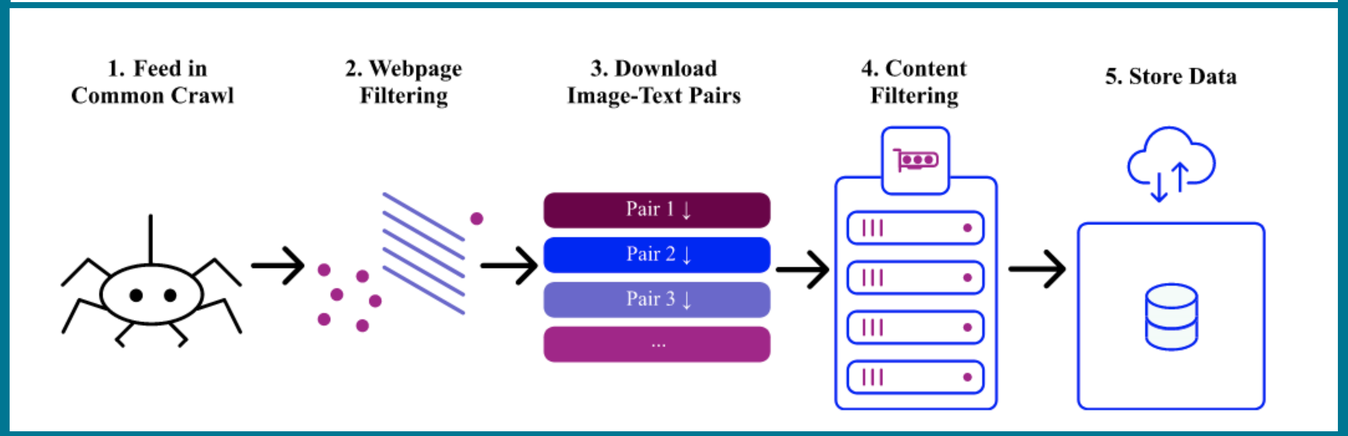

Beery et. al. 2020
Schuhmann et. al. 2022
Kunin et. al. 2020
Ahuja et. al. 2021

Types of figures

1. Introductory
2. Framework / pipeline / model
3. Experimental results
4.

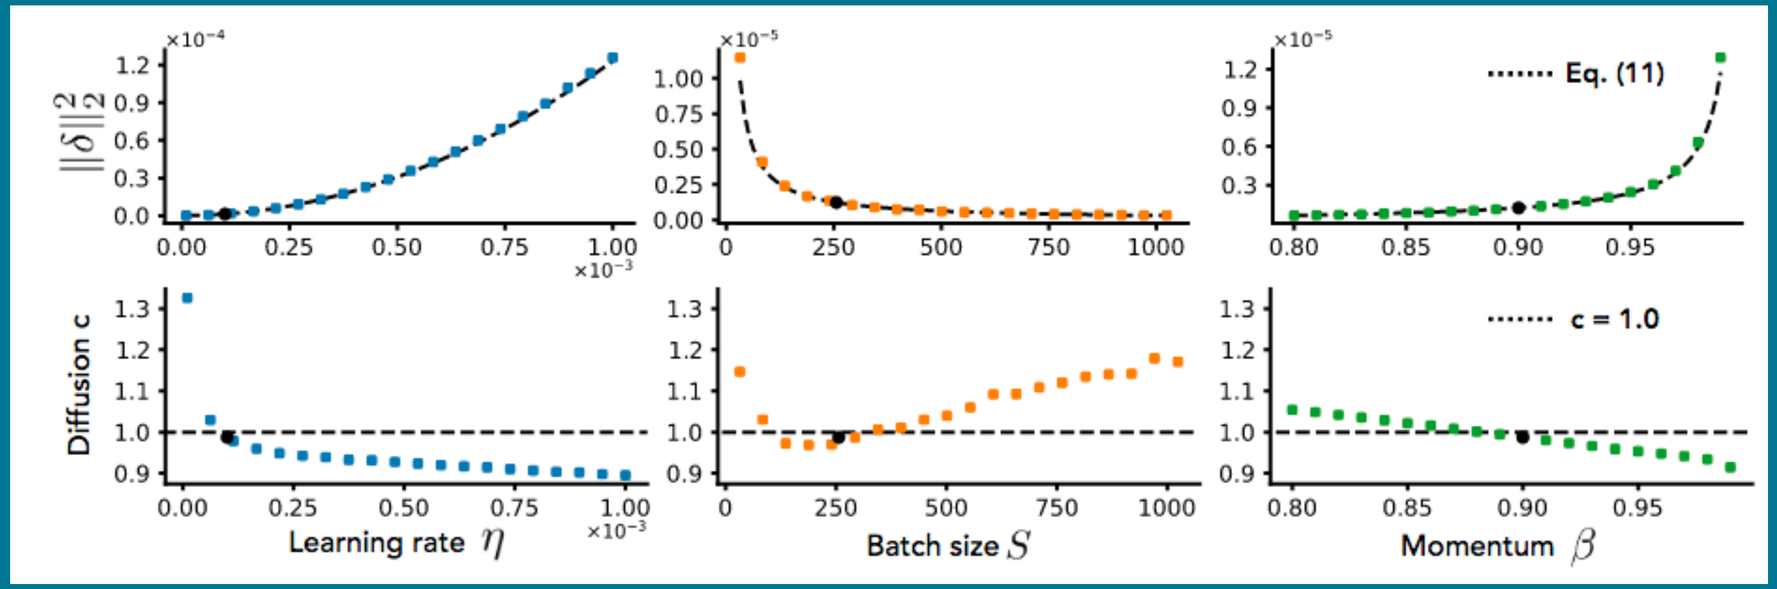

Beery et. al. 2020
Schuhmann et. al. 2022
Kunin et. al. 2020
Ahuja et. al. 2021

Types of figures

1. Introductory
2. Framework / pipeline / model
3. Experimental results
4. Downstream analysis / case studies

Beery et. al. 2020
Schuhmann et. al. 2022
Kunin et. al. 2020
Ahuja et. al. 2021

Fig. 10. CoolMoves uses limited input of the user's hands and head to recover plausible legs motions as showcased in this scene of a user climbing a ladder.

What makes a good figure?

A good figure does…

- Summarize one key takeaway at a time
- Make it immediate and easy to see your bit flip
- **Remain true to the data**

A good figure does **not**…

- Become a giant information dump of all your findings
- Bury your bit flip in a hard-to-see format
- Hide mediocre results with questionable display

Data presentation: the good, the bad, and the ugly

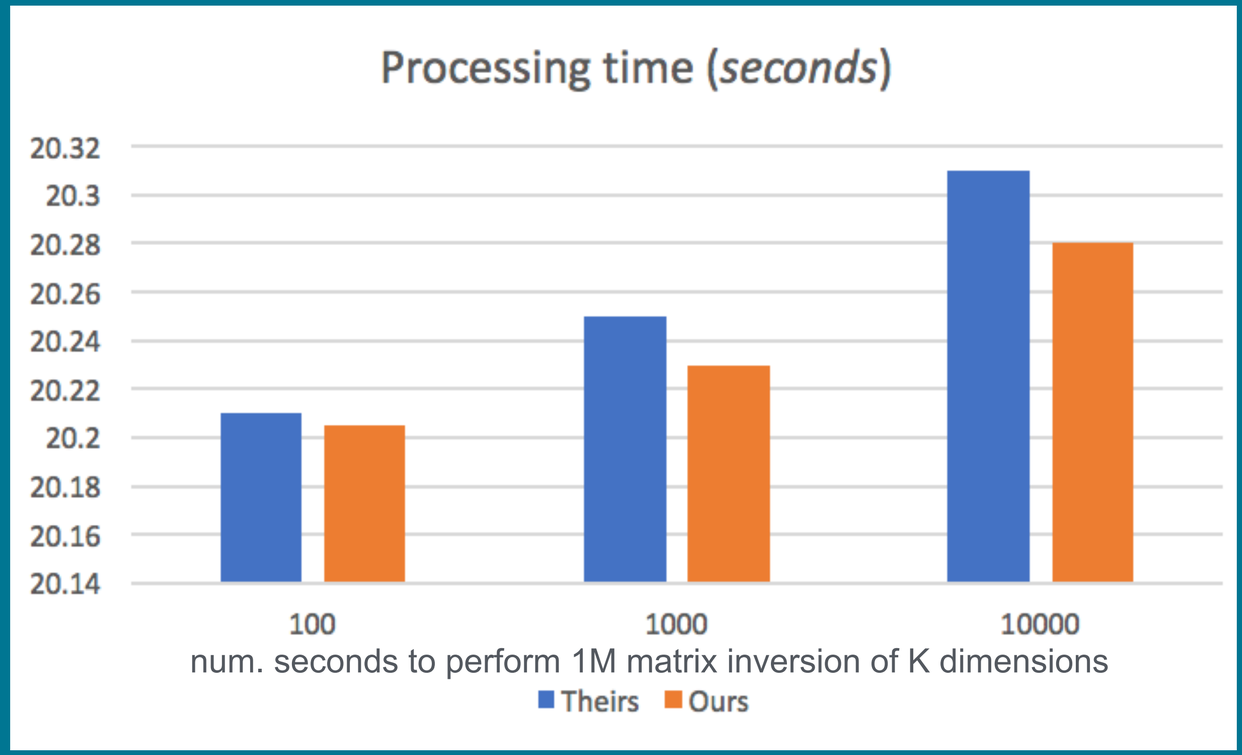

Data presentation: the good, the bad, and the ugly

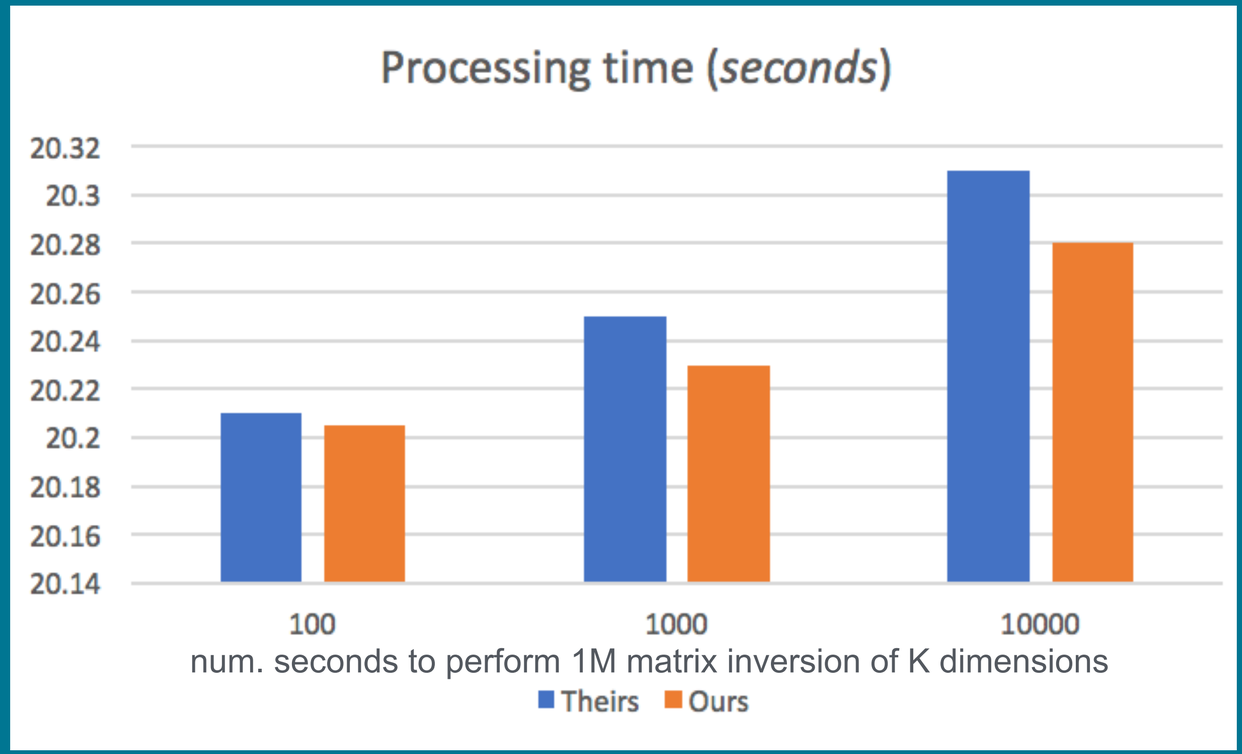

Finding figure gestalt

Rule 1: always label **all** axes and include units!!

Finding figure gestalt

Rule 1: always label **all** axes and include units!!

Rule 2: make sure the text on your figures is big enough to read

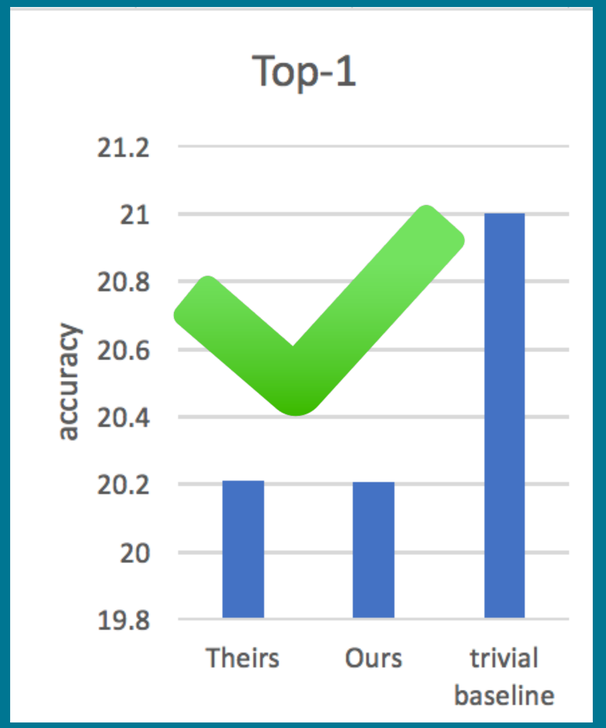

Finding figure gestalt

Rule 1: always label **all** axes and include units!!

Rule 2: make sure the text on your figures is big enough to read

Rule 3: symbology is your friend for conveying different treatments

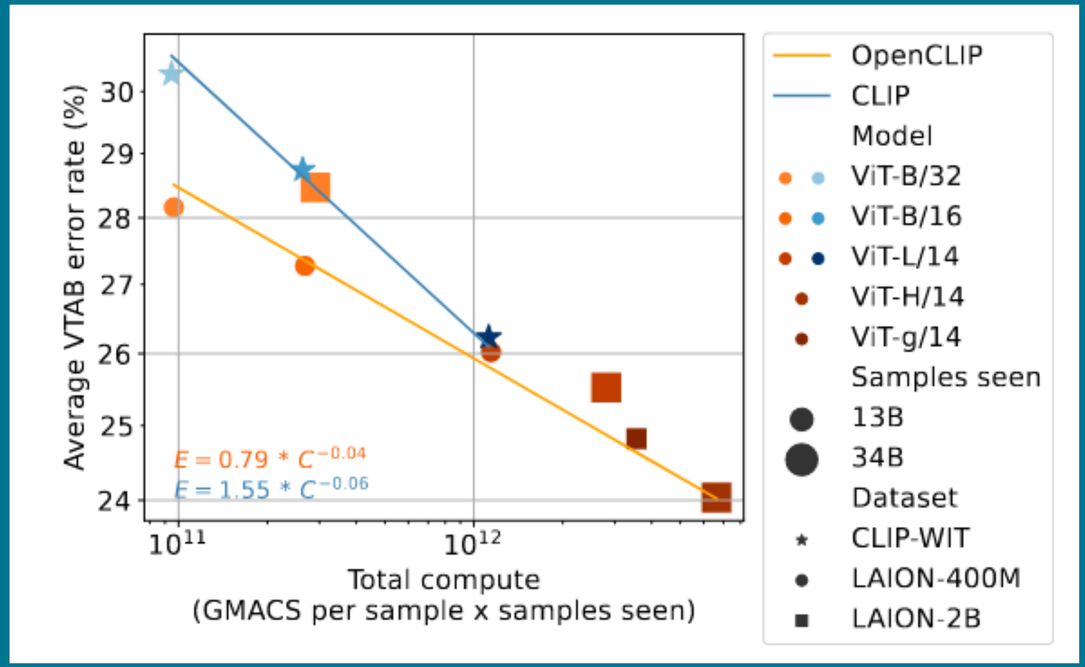

Cherti et. al. 2022

Finding figure gestalt

Rule 1: always label **all** axes and include units!!

Rule 2: make sure the text on your figures is big enough to read

Rule 3: symbology is your friend for conveying different treatments

Rule 4: aim for a clean and clear takeaway

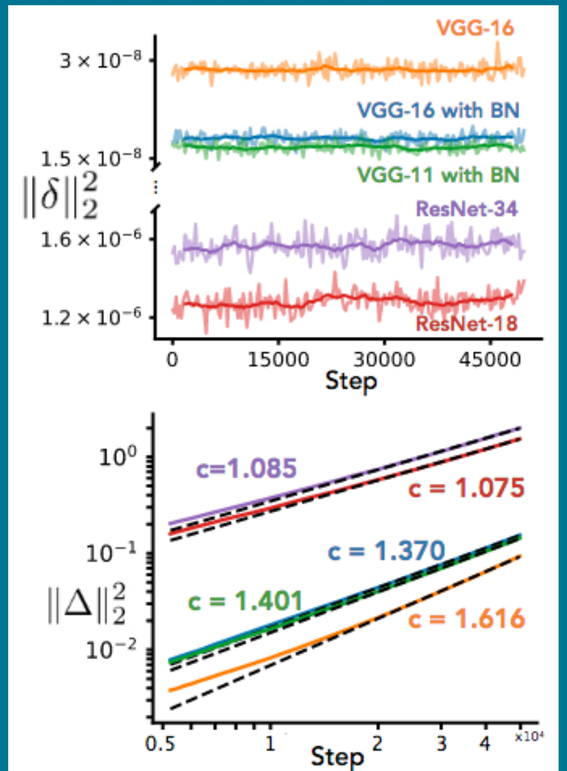

Kunin et. al. 2020

Finding figure gestalt

Rule 1: always label **all** axes and include units!!

Rule 2: make sure the text on your figures is big enough to read

Rule 3: symbology is your friend for conveying different treatments

Rule 4: aim for a clean and clear takeaway

P.S. - try and use colorblind-friendly colors and avoid bad palettes!

Helpful color palette picker: https://colorbrewer2.org/

Why the rainbow palette sucks https://gorelik.net/2020/08/17/what-is-the-biggest-problem-of-the-jet-and-rainbow-color-maps-and-why-is-it-not-as-evil-as-i-thought/

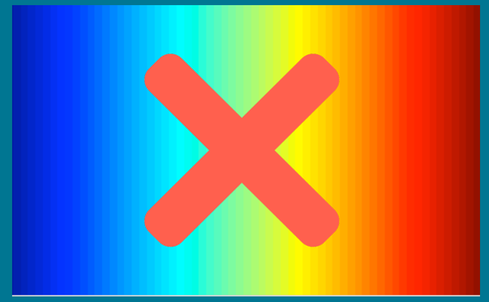

Tuning your tables

- Tables are the neutron stars of your paper
  - a.k.a. incredibly information-dense
- Ideal for summarizing 3+ treatments
- Centering your bit flip:
  - Can **bold** best model per-treatment
  - Clever alternatives: arrows, colored deltas
- Use colors, lines, bolding to distinguish treatments & improve readability

Model	Pre-training	INet	INet-v2	INet-R	INet-S	ObjNet
B/32	CLIP WIT	63.3	56.0	69.4	42.3	44.2
	LAION-400M	62.9$^{-0.4}$	55.1$^{-0.9}$	73.4$^{+4.0}$	49.4$^{+7.1}$	43.9$^{-0.3}$
	LAION-2B-en	65.7$^{+2.4}$	57.4$^{+1.4}$	75.9$^{+6.5}$	52.9$^{+10.6}$	48.7$^{+4.5}$
B/16	CLIP WIT	68.3	61.9	77.7	48.2	55.3
	LAION-400M	67.0$^{-1.3}$	59.6$^{-2.3}$	77.9$^{+0.2}$	52.4$^{+4.2}$	51.5$^{-3.8}$
B/16+	LAION-400M	69.2	61.5	80.5	54.4	53.9
L/14	CLIP WIT	75.6	69.8	87.9	59.6	69.0
	LAION-400M	72.8$^{-2.8}$	65.4$^{-4.4}$	84.7$^{-3.2}$	59.6	59.9$^{-9.1}$
	LAION-2B-en	75.2$^{-0.3}$	67.7$^{-2.0}$	87.4$^{-0.5}$	63.3$^{+3.7}$	65.5$^{-3.6}$

Schuhmann et. al. 2022

For your projects

- **CS 197**: ~2-4 figures + tables in final paper (high variance by area)
- **CS 197C:** ~1-3 figures + tables in final paper (preliminary)
- Typical top-tier CS conference: ~7-10 figures + tables
  - but again, depends on area
- **Last thing**: all figures and tables must include an accompanying text description that summarizes the figure (and helps the reader understand the bit flip / takeaway!)

Figures & Tables

+ Research Career Paths

“OK, so I took CS 197, now what?”

What can you do **at** Stanford?

What can you do **after** Stanford?

Pathways for research

I think research is interesting!

Professor

Research scientist in industry

Entrepreneur

Engineer / Engineering Lead

Professor

Work on research that you and the field find interesting.

Recruit and mentor the next generation of researchers in your field.

Teach in your area of expertise.

Typical goals:

- Do research and have impact (e.g., publications, software adoption)
- Graduate amazing students
- Inspire students to learn about your area
- Room for personalization: entrepreneurship, speaking, consulting, &etc.

Research scientist

Join a company's research division and work on research from within the company. Examples: Microsoft Research, FAIR, NVIDIA Research, OpenAI

Typical goals:

- Do research and have impact (but more focus on translation to the company's products and less on publication)
- Create innovations that transform the company you're working for (e.g., Kinect, GPT-4, DALL-E, PaLM, TPUs)

Entrepreneur

Start your own company, often based on the research you're doing, and grow it.

Typical goals:

- Scale your ideas and make them available to millions of people
- Start a new industry: your start-up is not a "me too" startup. Typically, it's pitching a dramatically new angle.
- Little focus on doing research in the short term, although the pitch might have originally been born as a research idea

Engineer / Engineering Lead

Join a company and apply your skills toward the development of product

Typical goals:

- Be the company's expert in an area, and potentially grow a team to drive product in that space
- Typically, these jobs are for types of levels of expertise and experience that cannot be acquired through a BS or MS
- Little focus on doing research in the short term, although not impossible

What's the distribution?

Michael Bernstein scraped names of all Ph.D. graduates in Computer Science from Stanford, MIT, and UC Berkeley.

He then mapped the names onto LinkedIn pages (yes, LinkedIn availability adds bias, but we found about 75% of people)

Tag their jobs on their LinkedIn:

- Faculty: job titles including words such as "faculty" or "professor"
- Entrepreneurship: triggered by titles such as "founder" or "partner"
- Research scientist: titles such as "researcher" or "scientist"
- Engineer: titles such as "programmer" or "architect"

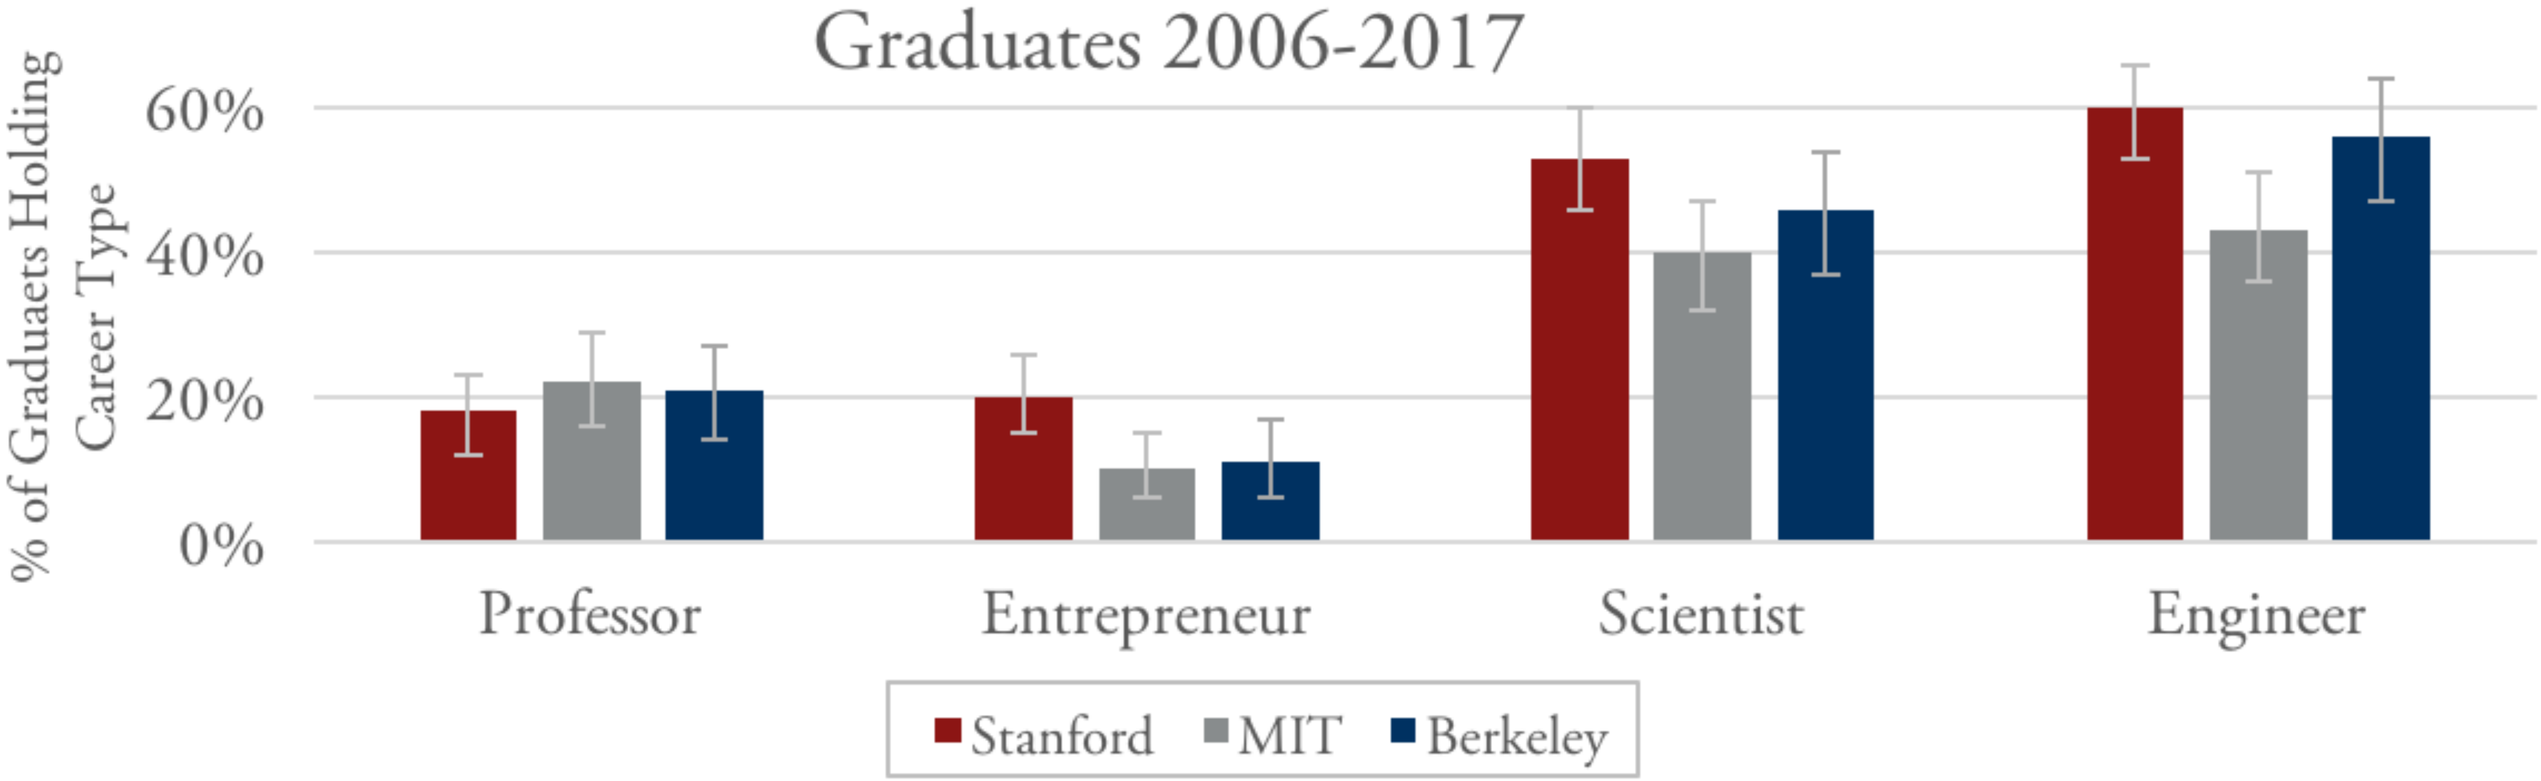

No statistically significant difference

No statistically significant difference

No statistically significant difference

Percentages add up to more than 100% because people can hold more than one position. Entrepreneurs and research scientists are a common mix. Faculty, likewise, can sometimes jump into industry research or start a company.

Pathways for research

I think research is interesting!

zooming in on the arrow

Professor

Research scientist in industry

Entrepreneur

Engineer / Engineering Lead

Pathways for research as a Stanford undergrad

I think research is interesting!

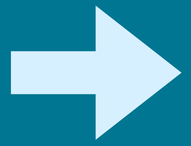

Academic year research

Summer CURIS internship

BS with honors

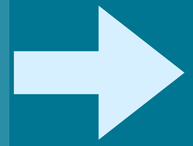

Professor

Research scientist in industry

Entrepreneur

Engineer / Engineering Lead

Academic year research

Get units for doing research with a faculty member

- Generally, start with CS 195, which fulfills the CS Senior Project requirement, then go on to CS 199
- How to get started? Talk to your CA about possible faculty to approach, and we can help facilitate an introduction.
- Typically, you'll get involved in a project ongoing in the lab

Continuing CS 197 research

The CAs are happy to keep working with you! If you'd like, we can support similar independent study courses (e.g., CS 195) to continue your CS197 project toward a workshop, work in progress, or paper

We, the staff, are also happy to help facilitate introductions to faculty you want to work with (but can't guarantee a spot or even a response)

Summer CURIS research

Apply your full effort toward a fun research project for the summer

- Get mentored by a faculty member and PhD student
- Get paid
- No need to balance the project against classes
- Live on campus

Typically, you join a project that's ongoing in the faculty member's lab

Apply early in winter quarter at curis.stanford.edu

BS with honors

Receive a special designation on your diploma (“BS with honors”)

Engage in a yearlong research project your senior year

- Takes the place of the senior project
- Typically, you do this with faculty who you’ve already been working with

Apply in the spring of your junior year

Pathways for research

I think research is interesting!

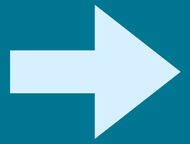

Academic year research

Summer CURIS internship

BS with honors

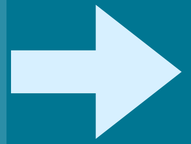

Professor

Research scientist in industry

Entrepreneur

Engineer / Engineering Lead

Pathways for research

I think research is interesting!

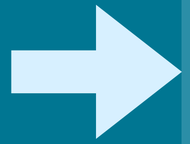

Academic year research

Summer CURIS internship

BS with honors

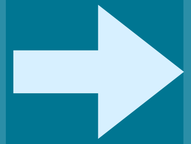

Ph.D.

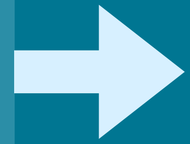

Professor

Research scientist in industry

Entrepreneur

Engineer / Engineering Lead

All of you can succeed at a PhD!

A Ph.D. is a grown-up version of the research you do as an undergraduate or master's student. You get much more control over the projects you are working on, and become first author on the resulting publication.

It's challenging because we doubt ourselves constantly. But you also earn the ability to tackle any complex problem.

Cool side benefit: become Dr. [Lastname]

How do I get in to a Ph.D.?

The most important criteria for getting into a Ph.D. program is **demonstrated interest and ability** to do research.

"How do I demonstrate interest and ability?" **Do research!**

How do I get in to a Ph.D. program?

In your statement, talk about research you did and the impact you had on the project. (You can include your CS 197 class project in it!)

You will want three recommendation letters from people with Ph.D.s to support your case.

> Typically, one is the faculty you worked most closely with on research. The other two can be supporting letters, or other research mentors.

Research career panel

Lauren

Akshat

Yujie

What questions do you have for us?

Figures & Tables
+ Research Career Paths