Giving Talks

CS 197 & 197C | Stanford University | **Sean Liu** & Lauren Gillespie
cs197.stanford.edu | cs197c.stanford.edu

Administrivia

197

Due Thurs 5/31, 10am:

Assignment 9: Draft Talk

5-min talk (+1 min Q&A)

Note: Section on Thurs!

197C

Due Thurs 5/31, 10am:

Assignment 6: Draft Talk

5-min talk (+1 min Q&A)

Overview From Here: 197

Section on Thurs 6/1

Practice Presentation & Peer Review

Final Project Report & Advising: Thurs 6/1—Wed 6/7

Due **Tues** 6/13, 3:30pm:

Final paper, final talk, project due.

Final exam slot: (3:30-6:30pm): we will meet in 60-109 for presentations!

Overview From Here: 197C

Small Group meeting on Thurs 6/1

    Practice Presentation

Research mentor feedback due: Tues 6/6

    Get feedback on your draft proposal from your research mentor

Due **Tues** 6/13, 3:30pm:

    Final paper, final talk, milestone due.

    Final exam slot: (3:30-6:30pm): we will meet in 60-109 for presentations!

Grading Scheme

Encourage mastery learning:

If a section of the final paper improves on the original submission, we will upgrade the original assignment grade to the higher final grade.

For example, Introduction assignment was a √, but the Introduction in the final paper was a √+ —> Update original to √+

Conversely, final submissions that do not address feedback will be penalized in the final talk and final paper grades.

If you do not submit anything for an assignment (or make a low-effort submission), you will not be able to improve your grade later for that assignment.

Anonymous Feedback!

Thank you for the feedback!

197:

Lecture topics: Project ideation, exploring topics, keeping up-to-date, etc.

Team dynamics: CAs have implemented measures and will check in regularly

Project — by now, you should have a clear list of TODOs on what to submit to get an A on your project. Please check w/ your CA!

197C:

Learning goal: iterate on your bit flip. Hit the ground running come summer!

Pre...

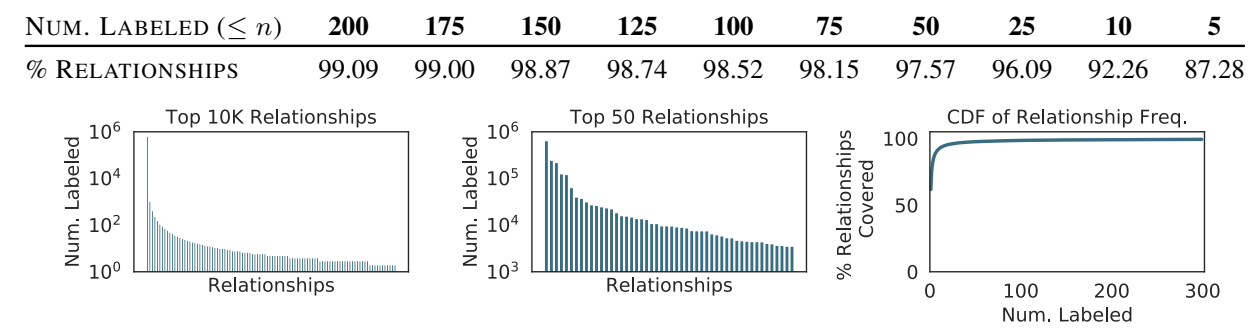| Num. Labeled ($\leq n$)	200	175	150	125	100	75	50	25	10	5
% Relationships	99.09	99.00	98.87	98.74	98.52	98.15	97.57	96.09	92.26	87.28

Figure 2. Visual relationships have a long tail (left) of infrequent relationships. Current models [49,54] only focus on the top 50 relationships (middle) in the Visual Genome dataset, which all have thousands of labeled instances. This ignores more than 98% of the relationships with few labeled instances (right, top/table).

features can characterize *some* visual relationships very well, they might fail to capture *complex relationships* with high variance. To quantify the efficacy of our image-agnostic features, we define "subtypes" that measure spatial and categorical complexity (Section 3).

Based on our analysis, we propose a semi-supervised approach that leverages image-agnostic features to label missing relationships using as few as 10 labeled instances of each relationship. We learn simple heuristics over these features and assign probabilistic labels to the unlabeled images using a generative model [39,46]. We evaluate our method's labeling efficacy using the completely-labeled VRD dataset [31] and find that it achieves an F1 score of 57.66, which is 11.84 points higher than other standard semi-supervised methods like label propagation [57]. To demonstrate the utility of our generated labels, we train a state-of-the-art scene graph model [54] (see Figure 6) and modify its loss function to support probabilistic labels. Our approach achieves 47.53 recall@100[1] for predicate classification on Visual Genome, improving over the same model trained using only labeled instances by 40.97 points. For scene graph detection, our approach achieves within 8.65 recall@100 of the same model trained on the original Visual Genome dataset with $108\times$ more labeled data. We end by comparing our approach to transfer learning, the de-facto choice for learning from limited labels. We find that our approach improves by 5.16 recall@100 for predicate classification, especially for relationships with high complexity, as it generalizes well to unlabeled subtypes.

Our contributions are three-fold. (1) We introduce the first method to complete visual knowledge bases by finding missing visual relationships (Section 5.1). (2) We show the utility of our generated labels in training existing scene graph prediction models (Section 5.2). (3) We introduce a metric to characterize the complexity of visual relationships and show it is a strong indicator ($R^2 = 0.778$) for our semi-supervised method's improvements over transfer learning (Section 5.3).

---
[1] Recall@K is a standard measure for scene graph prediction [31].

2. Related work

Textual knowledge bases were originally hand-curated by experts to structure facts [4,5,44] (e.g. <Tokyo-capital of-Japan>). To scale dataset curation efforts, recent approaches mine knowledge from the web [9] or hire non-expert annotators to manually curate knowledge [5,47]. In semi-supervised solutions, a small amount of labeled text is used to extract and exploit patterns in unlabeled sentences [2, 21,33–35,37]. Unfortunately, such approaches cannot be directly applied to visual relationships; textual relations can often be captured by external knowledge or patterns, while visual relationships are often local to an image.

Visual relationships have been studied as spatial priors [14, 16], co-occurrences [51], language statistics [28,31,53], and within entity contexts [29]. Scene graph prediction models have dealt with the difficulty of learning from incomplete knowledge, as recent methods utilize statistical motifs [54] or object-relationship dependencies [30,49,50,55]. All these methods limit their inference to the top 50 most frequently occurring predicate categories and ignore those without enough labeled examples (Figure 2).

The de-facto solution for limited label problems is **transfer learning** [15,52], which requires that the source domain used for pre-training follows a similar distribution as the target domain. In our setting, the source domain is a dataset of frequently-labeled relationships with thousands of examples [30,49,50,55], and the target domain is a set of limited label relationships. Despite similar objects in source and target domains, we find that transfer learning has difficulty generalizing to new relationships. Our method does not rely on availability of a larger, labeled set of relationships; instead, we use a small labeled set to annotate the unlabeled set of images.

To address the issue of gathering enough training labels for machine learning models, **data programming** has emerged as a popular paradigm. This approach learns to model imperfect labeling sources in order to assign training labels to unlabeled data. Imperfect labeling sources can come from crowdsourcing [10], user-defined heuristics [8,43], multi-instance learning [22,40], and distant su-
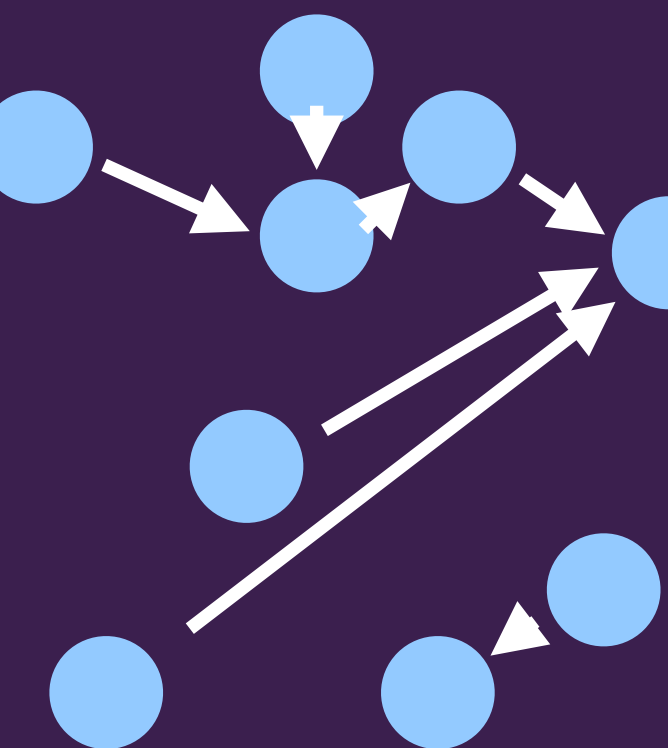
Related w...
and Lit se...

Writing a paper

You...

Unknown Terrain

...locity

...ing x / ...suring x

Giving a talk

Today's goals

Importance of giving a good talk

Goals of a talk

Tips for giving a clear and convincing talk*

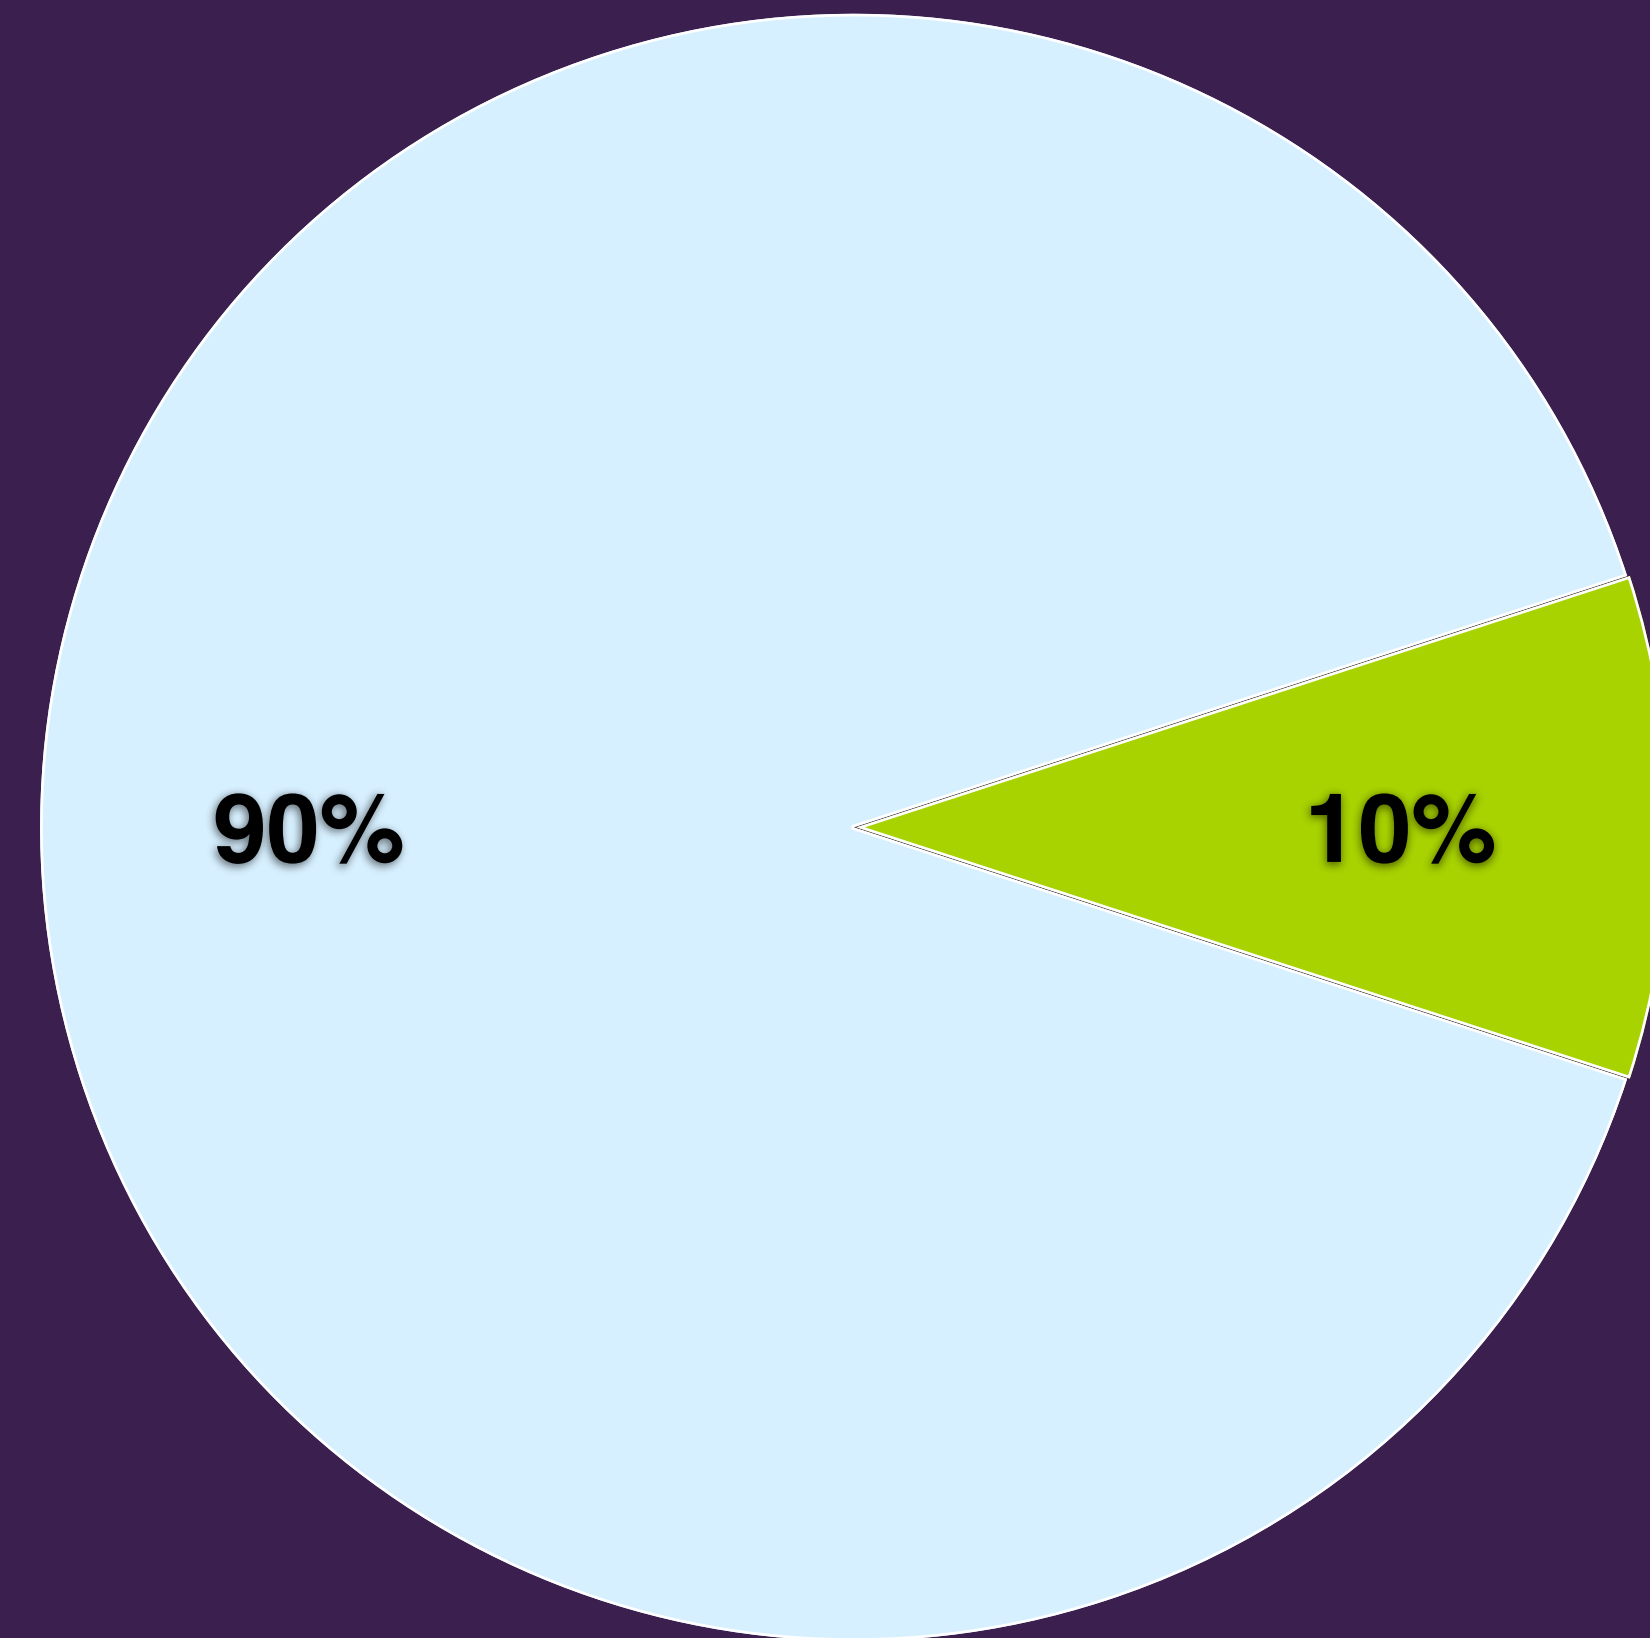*Adapted from Kayvon Fatahalian's talk: "How to give a clear talk"

Please go to http://graphics.stanford.edu/~kayvonf/misc/cleartalktips.pdf for full slides.

Importance of giving a good talk

Before my PhD

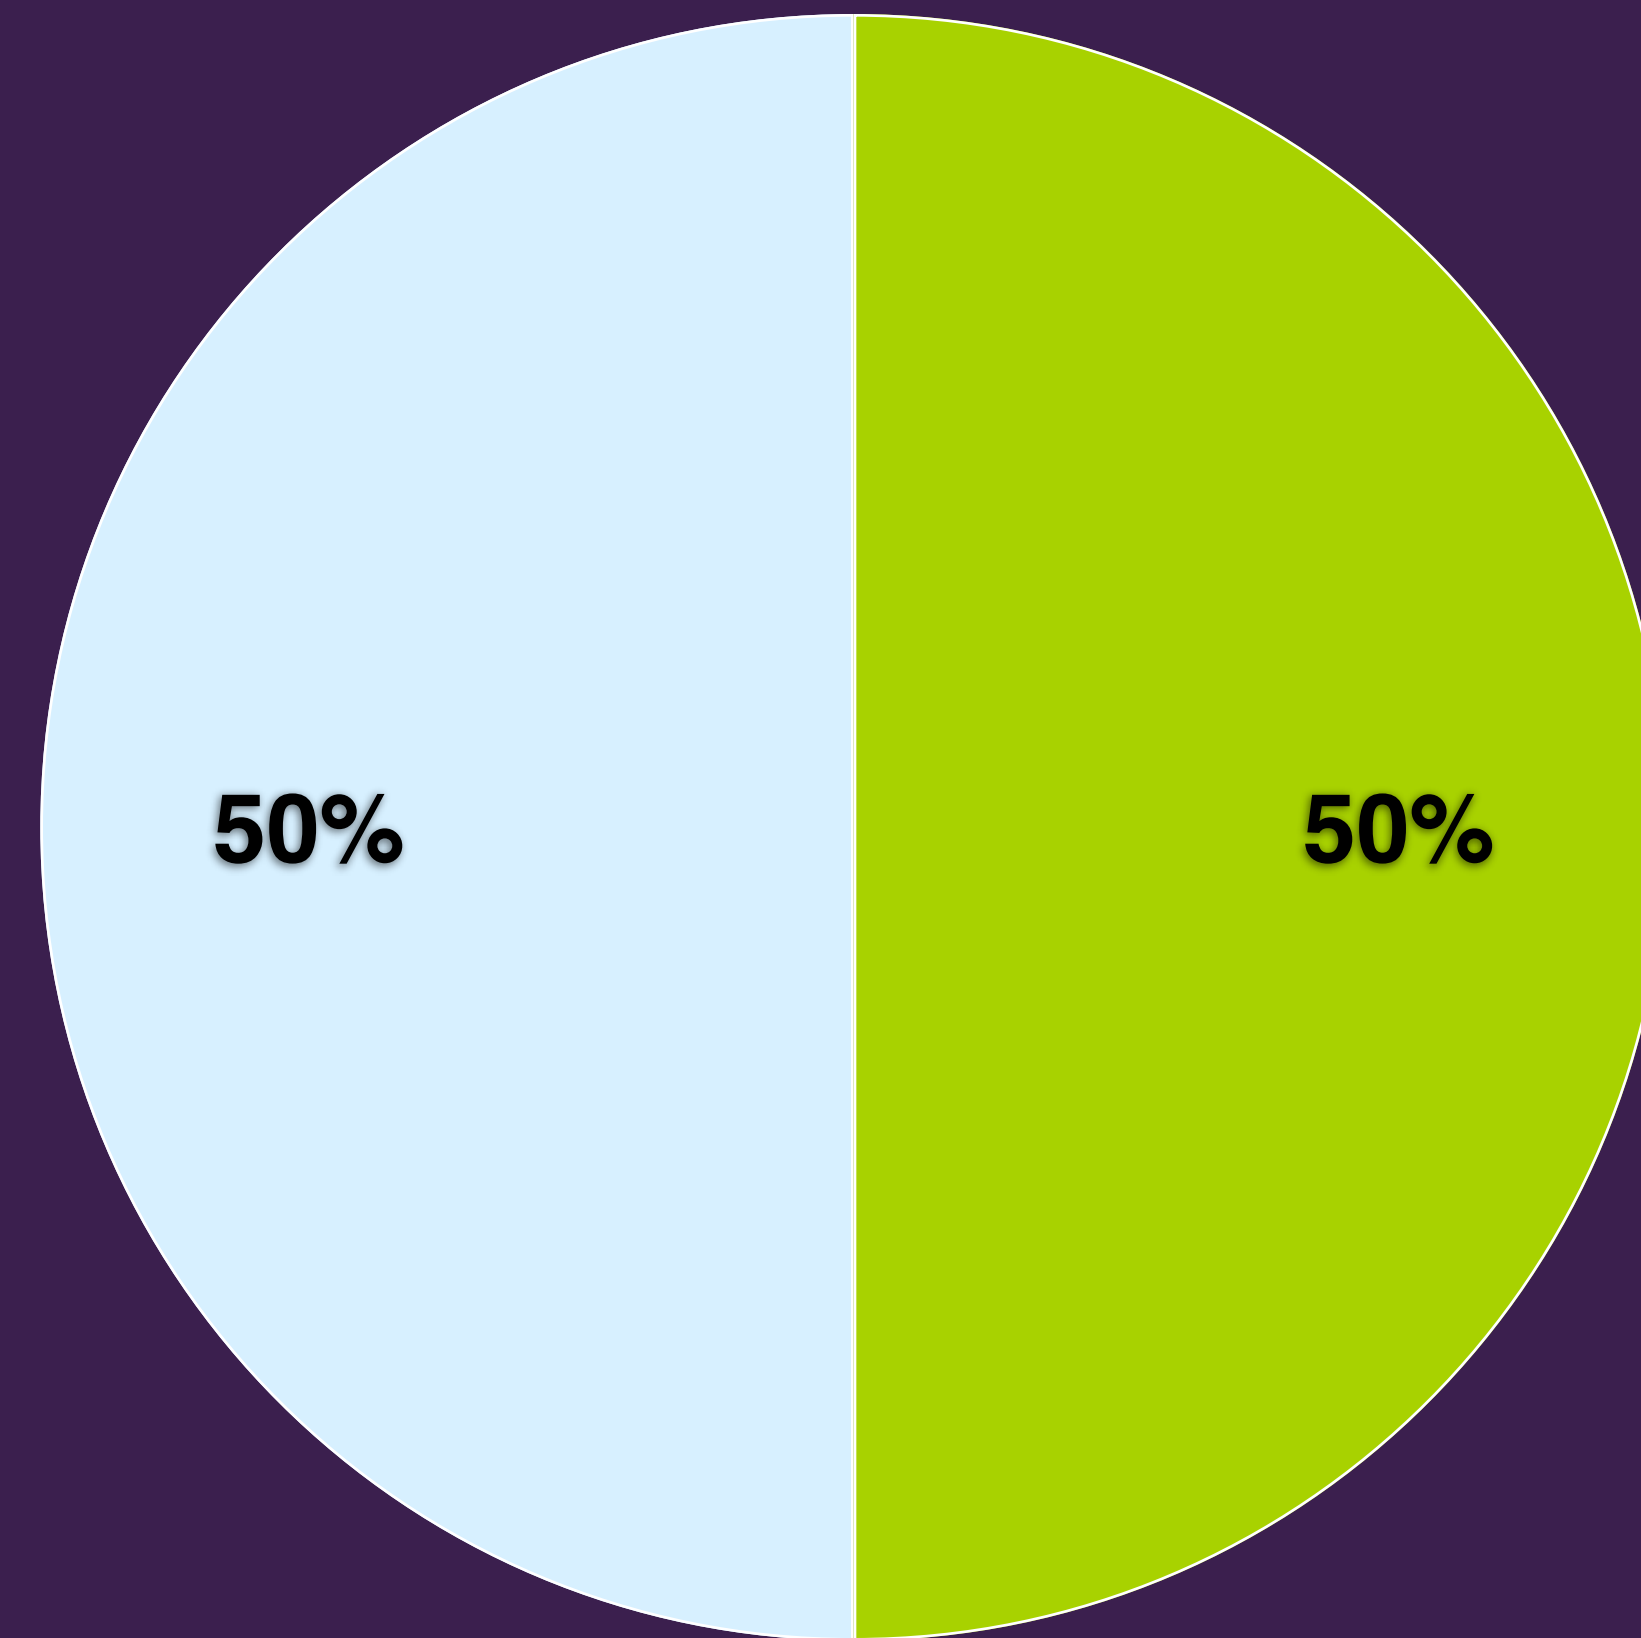## Technical
90%

Method, algorithm, system

Communication
10%

Paper writing, talks, twitter

Research Breakdown*

*Sean's opinion

After my PhD

Technical
50%

Method, algorithm, system

Communication
50%

Paper writing, talks, twitter

Research Breakdown*

*Sean's opinion

Common Misconception

If I do good work, people will automatically know, right?

No. In the real world, everyone is very busy…..

Unless they are directly working on a problem related to yours, they most likely will spend little energy in understanding your research.

At best, they will likely just come to your conference talk.
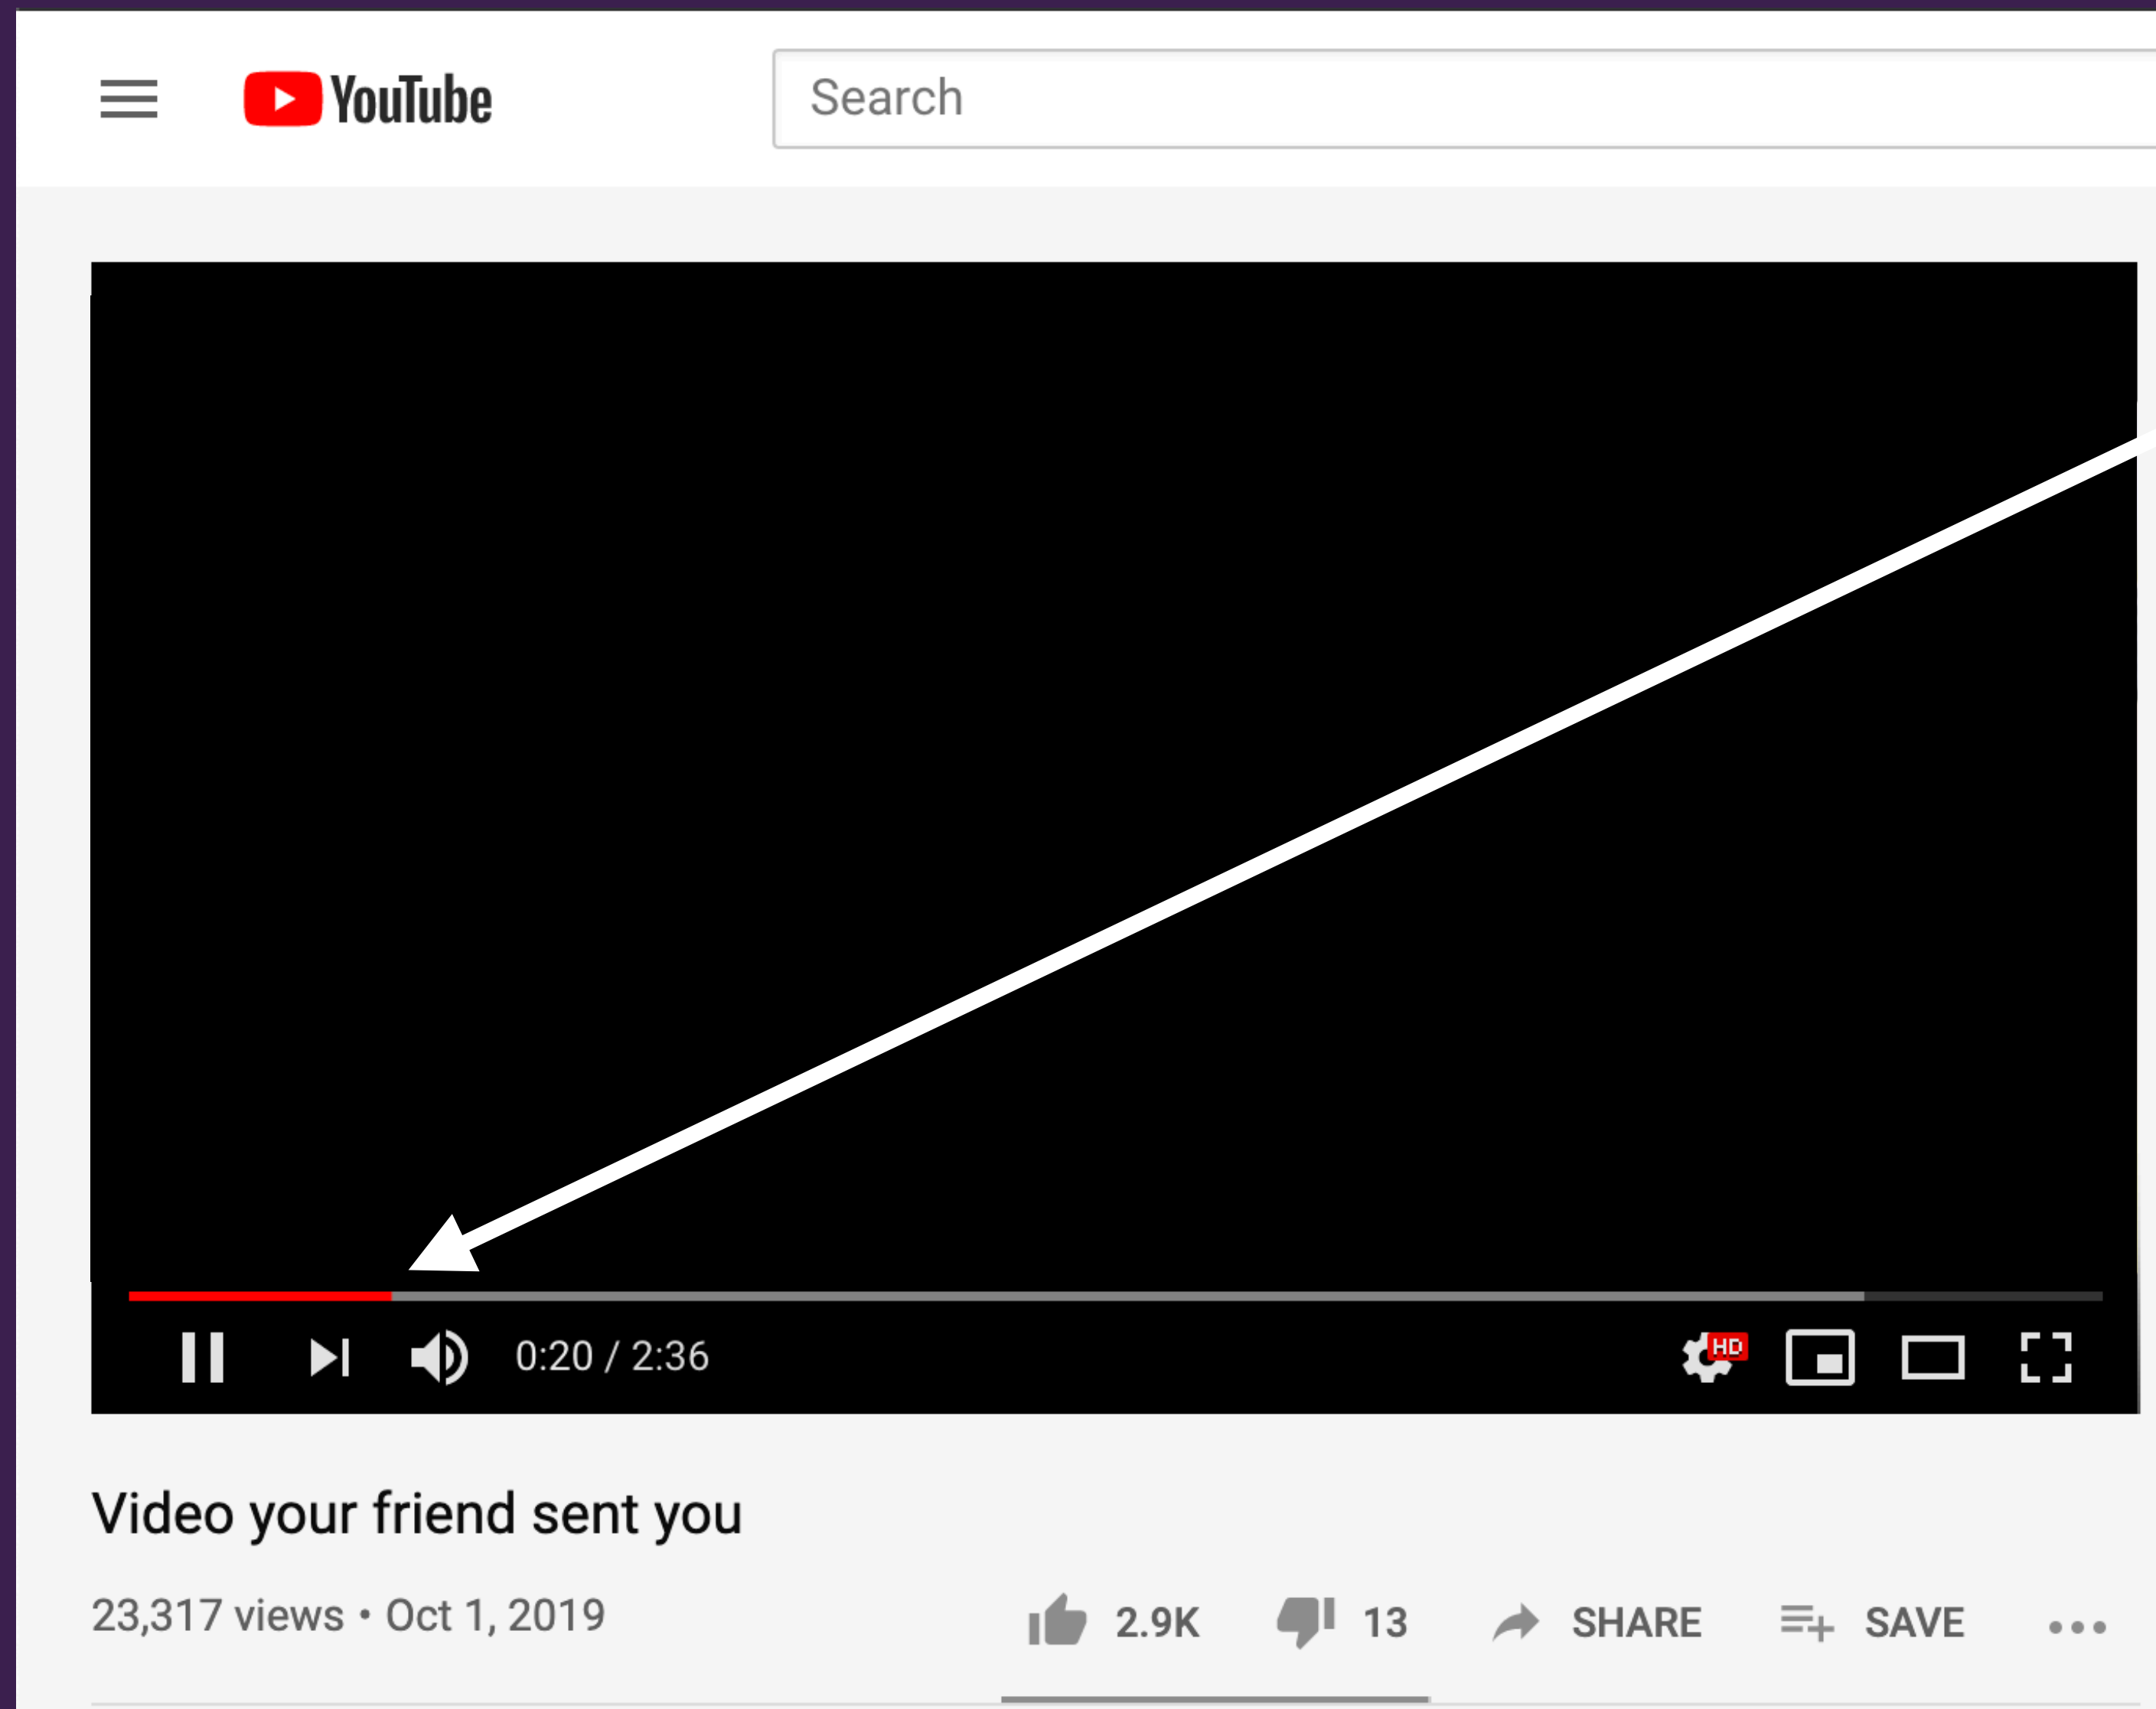
By this point, you've made a decision about whether to invest the time & energy to keep listening.

Goal of a Talk

1. Your Research Project

    Importance of the problem

    Why your solution is the right one

2. Introduce a new framing or way of thinking

    What the research community can take away and apply to their own research

3. Show that you are a good researcher — your brand!

    The way you pick problems, think about problems, and approach them

14

"Content Creators"

Catch people's attention: Good presentation

Content: Good research

Style / Brand: Your way of picking and thinking about problems

Tips for giving a good talk*

*Abridged version of Kayvon Fatahalian's talk: "How to give a clear talk"

Please go to http://graphics.stanford.edu/~kayvonf/misc/cleartalktips.pdf for full slides.

Tip 13: Present confidently!

Project your voice. Let the back of the room hear you.

Grab people's attention

Practice makes perfect :)

Reminder:

Submit your attendance on Canvas!

Giving Talks

Slide content shareable under a Creative Commons Attribution-NonCommercial 4.0 International License.