

Publication Culture and Peer Review

CS 197 | Stanford University | Arpita Singhal
cs197.stanford.edu

(Adapted from Prof. Michael Bernstein's slides)

Overview From Here

Home stretch!

Final paper, final talk, and team feedback due at the start of our finals slot

Reminder: iterate on any required feedback from your previous assignments for your final paper

Final presentations next Thursday at 3:30pm

Today's goals

What happens after you write a paper? And why do we always grumble about Reviewer 2?

What **peer review** is, why it matters, and how it works

How to develop a **high-quality review**

Dealing with disappointment

What are conferences, journals, arXiv, and what role do they play?

Publication culture

I finished the paper. Now what?

Now it's time for your research to take flight and enter the academic record.

...but why do we do this? Why care? And what are even the options?

Why peer review?

We often think of peer review as gating **correctness**, but it also performs another equally important function: gating contribution **impact**

There is a massive amount of research generated each year in computer science. (If you want to drink from the firehose, subscribe to daily announcements from [arXiv.org](https://arxiv.org).)

So what do you pay attention to?

An example in CS Theory

Amongst the papers written in Computer Science theory, the vast majority of them are correct proofs.

So, researchers in CS Theory are faced with a large pile of true facts about the world.

The role of the top-tier conferences is to establish **which of those true facts are the most important ones.**

(And yes, also to weed out any incorrect proofs.)

Typical gold standard: conference

Computer Science, unlike other fields, is a conference-oriented field.

There are a small set of **top-tier conferences** for each area. These are generally known to be the venues that publish the best work in the area.

There also exist a variety of second-tier and other conferences, which are less prestigious and often easier to get into.

Journals, and conference-journal hybrids, fit into this category too.

Work-in-progress venues

You can only publish a research result once. Conferences and journals are known as **archival**, meaning that they are archived permanently in the academic record.

There also exist a variety of **non-archival venues** that are intended for feedback and exposure.

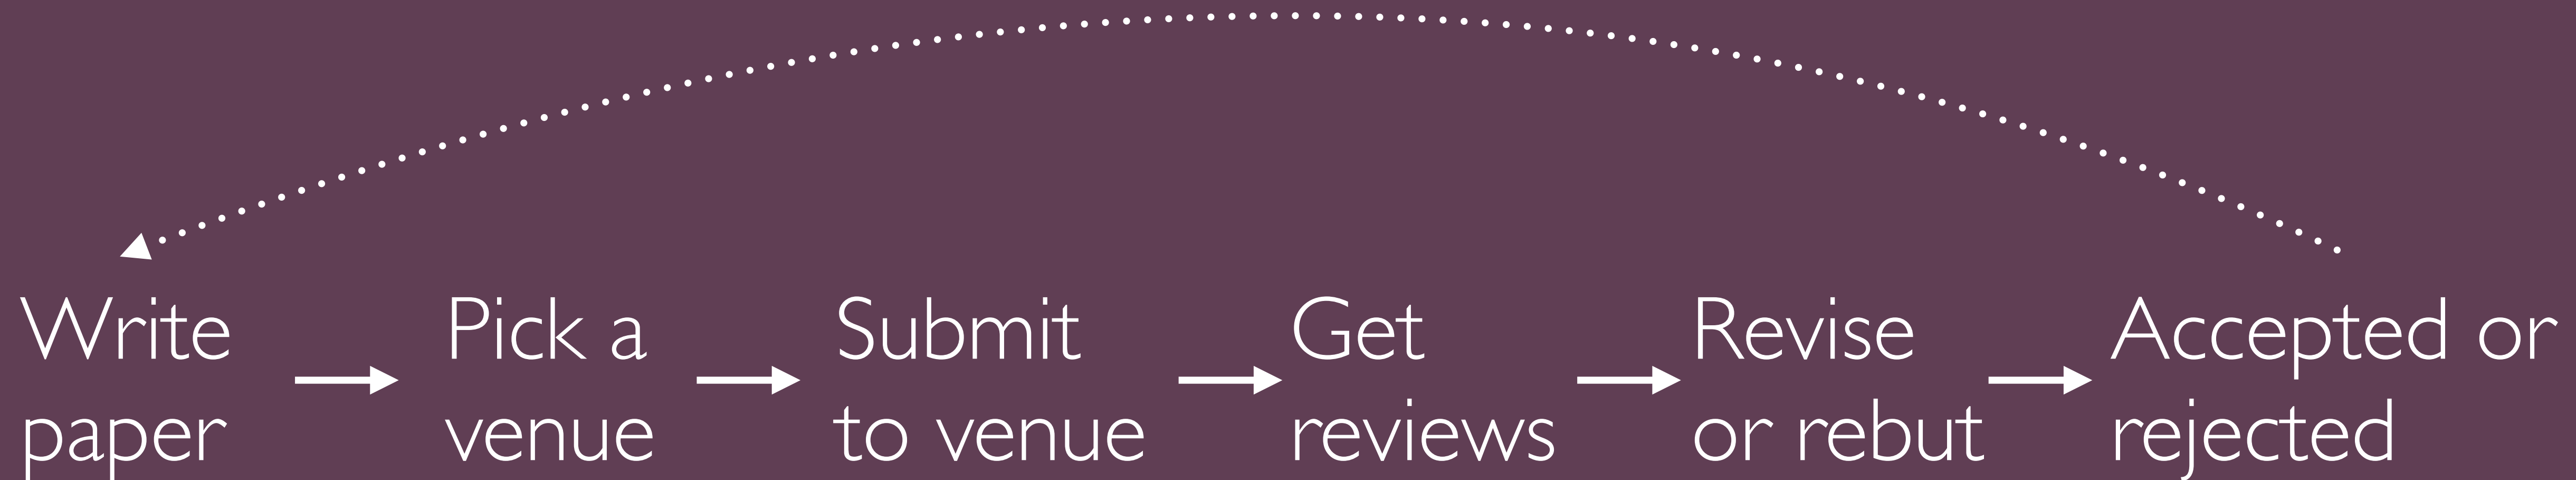
Workshops

Posters

Demos

arXiv.org

Life of a paper



“WIP venues sound fun...”

They should! VPUE provides Conference Grants for up to **\$1,500** to travel to present your research at a conference.

If you're interested, ask your TA!

They can work with you to identify a reasonable non-archival venue to submit to, and point you at the format requirements.

studentgrants.stanford.edu

Peer review

The dual role of peer review

You can always put your paper on a public report archive such as **arXiv.org**. But getting your research into a conference requires peer review.

Peer review relies on experts in the field to judge two questions:

- 1) Is this research **correct**? Does it actually achieve what it claims?
- 2) Is the **contribution** valuable enough to publish at this venue?

Who are the peers?

Ideally, your paper gets routed to people who are **experts** in the topic of your research.

People who publish in the area that you're working in

People who you cite in your submission

Anatomy of a peer review

Exact details vary, but many reviews contain the following elements:

- Overall score: e.g., 1-5

- Textual review
(~5 paragraphs)

The process

External review model

Associate Chair (AC)
Secondary Chair (2AC)



Invited reviewer 1
Invited reviewer 2
Invited reviewer 3

Internal review model

Senior Committee Member (SPC)
Senior Committee Member (SPC)



Committee member 1
Committee member 2
Committee member 3

Double-blind review


Typically, when you submit a paper to a conference, you **anonymize** yourself by not including your name or affiliation in the author block of the paper.

Goal: ensure that papers are reviewed on content, not on reputation

Likewise ACs' and reviewers' identities are hidden from the authors

Goal: avoid retaliatory behavior, focus on the institution of peer review rather than the people

What happens with reviews?



Example score distribution from a top-tier conference

Rebuttal and revision

Some conferences use **rebuttals**, where you have a short period of time (~1 week) to reply to the reviews. Reviewers read your rebuttal, adjust scores if desired, and then a final decision is made.

Machine learning conferences now often make this period interactive, with back-and-forths allowed with the reviewers

Other conferences and all journals use **revisions**, where a paper is given a specified period of time (a few weeks to a few months) to directly make changes based on the reviews. Reviewers read the revised paper, adjust scores if desired, and then a decision is made.

Who makes the final decision?

Typically, the senior members of the committee (ACs/SPCs) make a final recommendation based on the input of the reviewers.

Conference acceptance rates are often ~25%.

Sometimes you get unlucky.

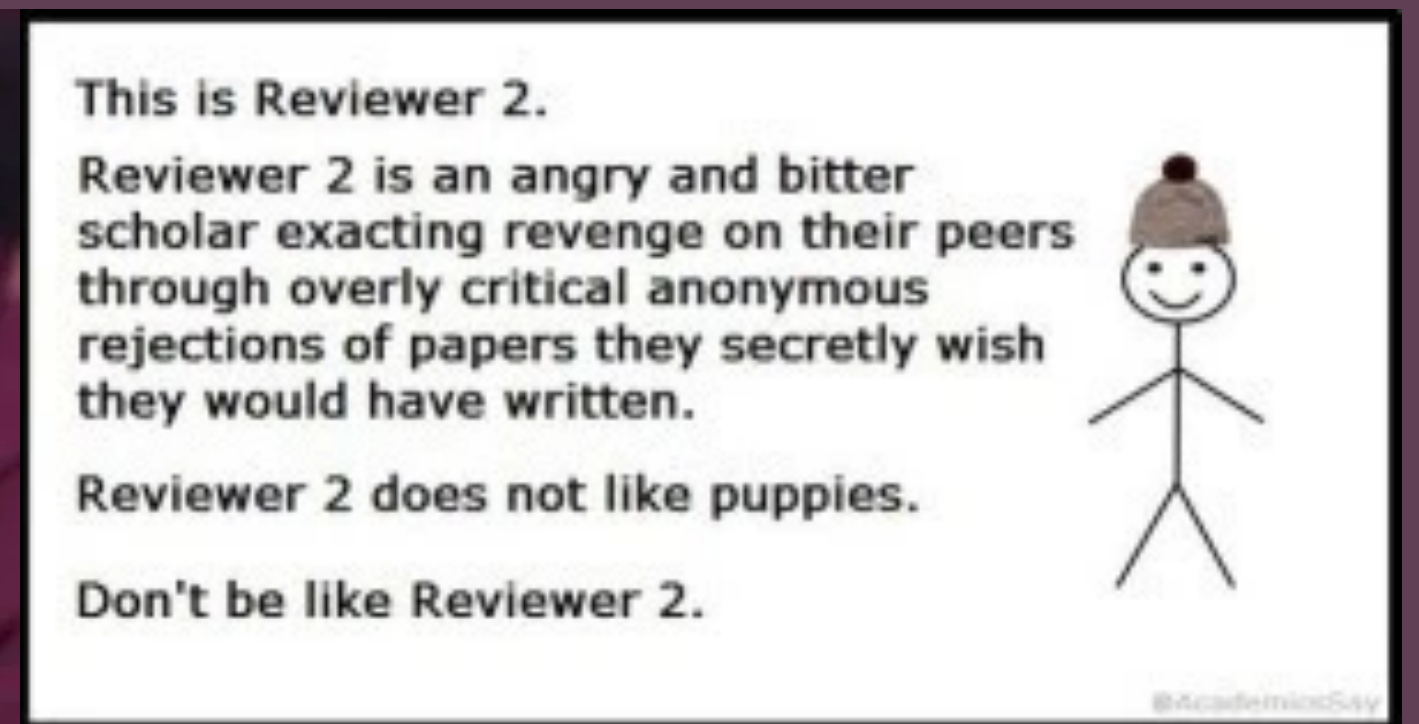


At NeurIPS 2021, 10% of the papers got redundantly reviewed by two parallel subcommittees.

Original \ Copy	Oral	Spotlight	Poster	Reject	Withdrawn
Oral	0	0	4	0	0
Spotlight	0	3	9	13	0
Poster	2	7	74	94	0
Reject	0	13	83	462	0
Withdrawn	0	0	0	0	118

51% of papers accepted by one committee were rejected by the second

Why do we shake our fist at R2?

Reviews can be quite harsh to read. Researchers refer half-jokingly to Reviewer 2 as the one who always has some bone to pick with your research and is unreasonably negative, trying to sink the paper.



How to write an effective review

The tempting behavior

- 1) Read the paper
- 2) Keep track of objections you have as you read the paper
- 3) Collate those objections into a review
- 4) Decide what score to give based on your objections

Why is that behavior problematic?

This winds up with **nitpicky** reviews: here's what's wrong, without placing those issues in context of the broader contribution.

Peer review is not an invitation to demonstrate your skill at identifying small flaws.

Writing a good review

Step one: ask yourself, **what goal is the paper trying to achieve?**

This may not be super clear from the paper. As a reviewer, your goal is to figure out what the bit flip is that they are arguing for, even if the authors aren't great at articulating it themselves.

Step two: **how well did the paper achieve that goal?**

Did they follow through on what their goal was? Did they demonstrate their thesis well?

Step three: **how could it have better achieved that goal?**

This is where you offer constructive critiques.

Writing a good review

Once you've taken those three steps, you can translate the result into a review. Essentially (but in your own words):

This paper sets out to [goal]. [Goal] is...

An important goal and well executed...

Making an implicit assumption that I disagree with...

(If relevant:) the execution...

Is a tour de force exploration of [goal]

Doesn't follow through on [goal] in the following way: [...]

(The execution may be a secondary matter if the goal is ill-formed!)

**What questions
do you have?**

Try it

Think back to your nearest neighbor paper. Take five minutes with your group to construct a review of that paper.

What goal is the paper trying to achieve?

How well does it achieve that goal?

How could the paper have better achieved that goal?

Dealing with rejection



Rejection is a fact of life in research.

My first NeurIPS paper submission as an Undergrad researcher got flatly rejected.

...

I've gotten rejected a lot. It hurts.

JOHANNES HAUSHOFER
CV OF FAILURES

Most of what I try fails, but these failures are often invisible, while the successes are visible. I have noticed that this sometimes gives others the impression that most things work out for me. As a result, they are more likely to attribute their own failures to themselves, rather than the fact that the world is stochastic, applications are crapshoots, and selection committees and referees have bad days. This CV of Failures is an attempt to balance the record and provide some perspective.

This idea is not mine, but due to a wonderful article in *Nature* by **Melanie I. Stefan**, who is a Lecturer in the School of Biomedical Sciences at the University of Edinburgh. You can find her original article here, her website here, her publications here, and follow her on Twitter under *@MelanieIStefan*.

I am also not the first academic to post their CV of failures. Earlier examples are here, here, here, and here.

This CV is unlikely to be complete – it was written from memory and probably omits a lot of stuff. So if it's shorter than yours, it's likely because you have better memory, or because you're better at trying things than me.

Degree programs I did not get into

2008	PhD Program in Economics, Stockholm School of Economics
2003	Graduate Course in Medicine, Cambridge University Graduate Course in Medicine, UCL PhD Program in Psychology, Harvard University PhD Program in Neuroscience and Psychology, Stanford University
1999	BA in International Relations, London School of Economics

Academic positions and fellowships I did not get

We are pleased to inform you that your paper has been accepted



As a grad student




As junior faculty



As tenured faculty

We regret to inform you that your paper has not been accepted



As a grad student



As junior faculty



As tenured faculty

How to handle bad reviews

First, take the time you need to **emotionally** process it. My process basically follows the Kübler-Ross model:

1. Denial and isolation
2. Anger
3. Bargaining
4. Depression
5. Acceptance

This is a very natural human reaction, and not one we directly have control over, so just let it happen.

Making the most of it

I see two common clusters of bad reviews:

- 1) People who **don't get the paper**. These reviews don't engage with the core idea, or engage with the wrong aspects of the idea, and their critiques come across as surface-level as a result.
- 2) People who **get the paper and still didn't like it**. These reviews are often really incisive and take down core assumptions or approaches you're taking.

Each of these clusters has something to tell us about our paper.

“They don’t get it”

These reviews suggest one of two things:



Your paper didn’t get in front of the right kind of reviewer, like it didn’t hit someone who works in the right area.

(Then: what are you signaling in your title or abstract that is attracting the wrong kind of reviewer?)

Your paper got in front of the right kind of reviewer, but they didn’t connect with your idea

(Let’s talk about Plato’s Cave...)


Plato's Cave




Your brilliant idea



The shadow cast by the paper you actually wrote





What reviewers thought you were saying




The shadow cast by their reaction in the review you read


Plato's Cave




Your brilliant idea



The shadow cast by the paper you actually wrote



What reviewers thought you were saying



The shadow cast by their reaction in the review you read

Your goal: invert the transformation to understand what really needs to change about your idea or its presentation.

Corollary: don't take the feedback literally.

“They get it”

These reviews are the really good kind of burn. It hurts because they're right.

You can shortcut the Plato's Cave process here, and take their advice more at face value.

Possible outcomes

Non-exclusive options

Reframe the paper: reconsider your bit flip (“what is the goal?”)

Perform additional engineering or evaluation work (“how well did the paper achieve the goal?”)

Revise and resubmit

Dear Michael S Bernstein,
Later published at NeurIPS
We regret to inform you that your paper 6901 (AAAI-22). This year we received a record 9,251 acceptance rate of 15%.

Papers have the potential to transition **from a flat-out reject to a Best Paper winner.**

Did those papers get in front of more sympathetic reviewers? Maybe.

Did those papers benefit from a more refined vision, execution, and articulation? Absolutely.

In some cases, **rejection can actually be the best outcome.** I'd rather have a paper rejected, iterate, and then win an award, than barely get a paper accepted and never have the impact it could have had.

**What questions
do you have?**

Publication Culture and Peer Review

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