

# Publication Culture and Peer Review

CS 197 | Stanford University | Michael Bernstein

# 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?

# Announcements

We are always tuning this class to make it a better experience. A few changes:

- More 1-on-1 project feedback with TAs

- Dropping Assignment 9 (peer review) to give more time on the final project. Points are going to the final paper and the participation grade instead.

# More announcements

Second optional revision deadline for your draft paper

The staff will give feedback on whatever you submit by 11/27 @ noon

Revising the paper between the deadline and the 27th is optional; we are doing this because several teams wanted more time to add results before we look at your drafts

Peer review in section on Friday will be on the version you just submitted

Final team feedback form forthcoming

# Even more announcements

Your guest judges for your final presentations:

Cynthia Lee

Chris Piech

Phil Levis

Final presentations are in class on Week 10.

Plan on an extended class day (~**5:50pm**) to support the presentations

Final papers will be due in finals week

# “How am I doing in class?”

Our intent: give feedback on submissions to help you improve your projects, and give you a yardstick to measure your progress.

Your feedback: “I’m getting points knocked off! @\$#\*&!”

We hear you.

This isn’t a class where we’re trying to separate the good from the great.

In the future, we will be redesigning these feedback mechanisms.

Pragmatically: keep in mind the means and s.d.’s, and check Carta for CS 347 for a likely grade distribution

# 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?



# Whyyyyyyyyyyy

(Yes, the extra letters are added for emphasis.)

Class to MSB: “OK, millennial.”

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.

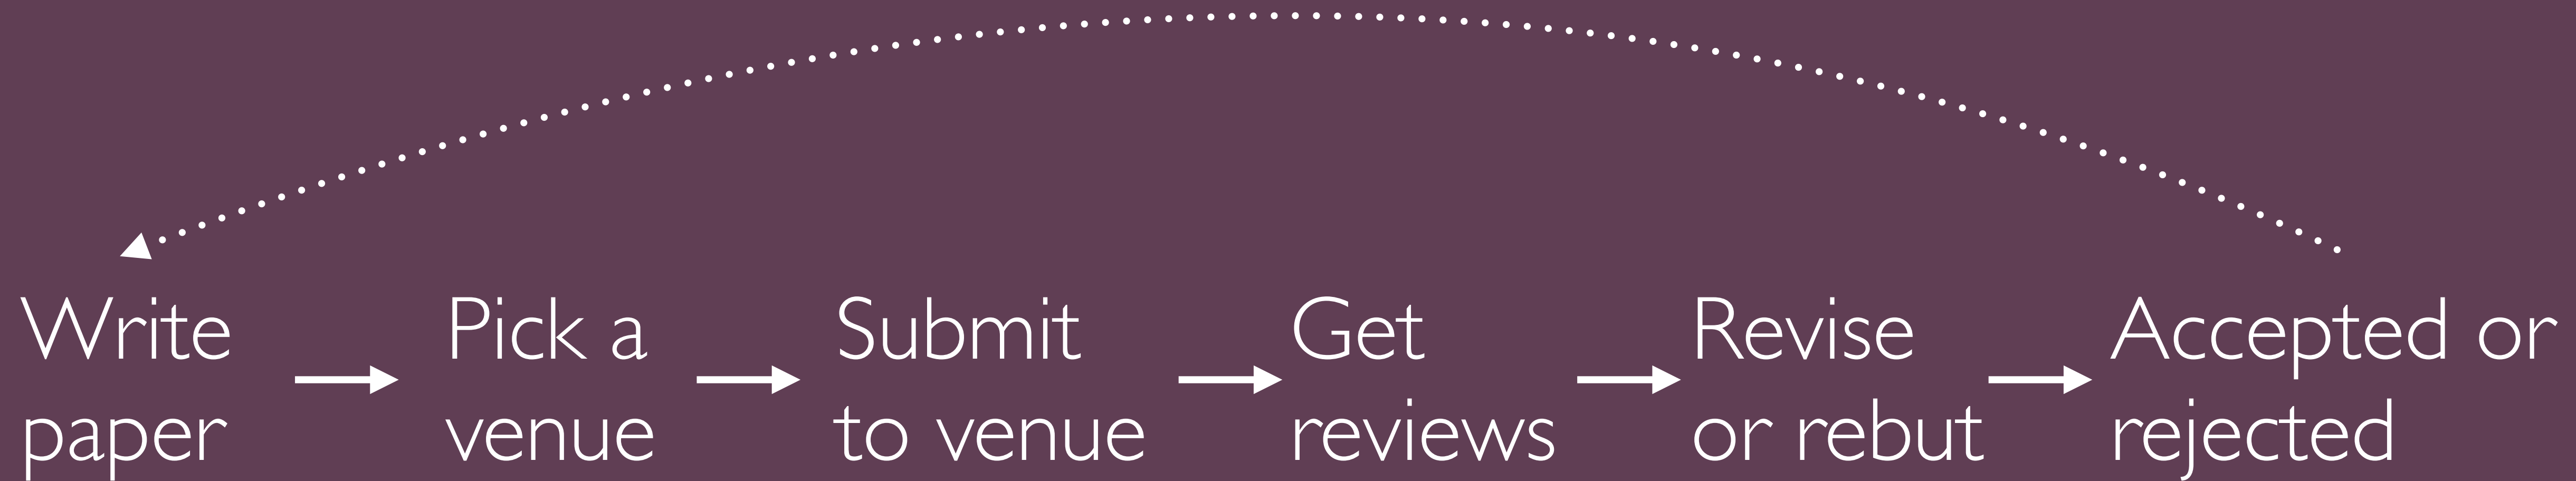
Workshops

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](http://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 most reviews contain the following elements:

- Overall score: 1-5

- Textual review  
(~5 paragraphs)

# The process

## External review model

Associate Chair (AC)  
Secondary Chair (2AC)



Think and invite

Invited reviewer 1  
Invited reviewer 2  
Invited reviewer 3

## Internal review model

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



Assign out of a  
pre-recruited pool

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.

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.

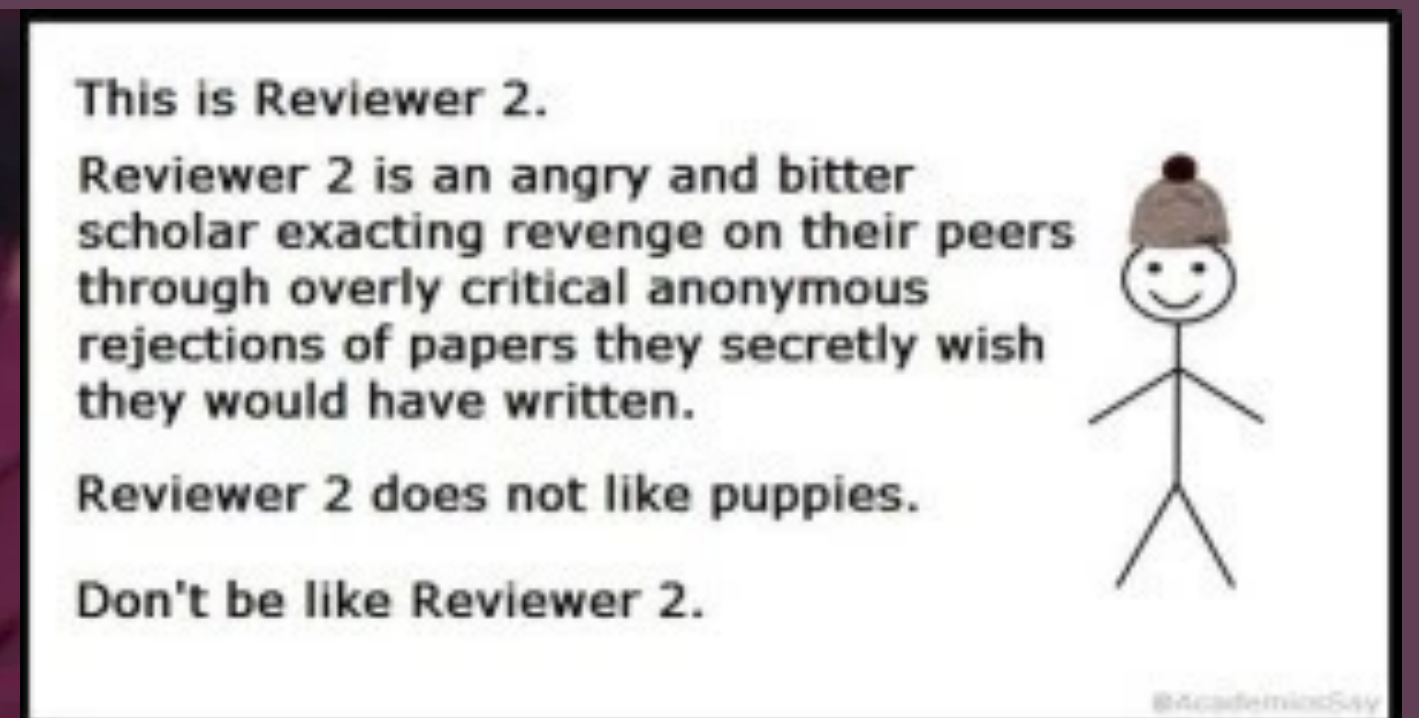


# Who makes the final decision?

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%.

# 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?

[2min]

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

# 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 CHI paper submission as a Ph.D. student 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. 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

I have, multiple times, transitioned a paper 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 is actually 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?**

# From here on out

Final presentations in our last day of class

Final papers due during finals week

Final team assessments due with the final paper

Please let me know feedback in person or via the course evaluation.  
This is the first instance of the class, and we have lots of ideas of how to improve it, and we want yours too.

# Publication Culture and Peer Review

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