

*CS 2001: Research Topics*  
**Reviewing Papers**

Prof. Adriana Kovashka  
University of Pittsburgh  
September 15, 2021

# Plan for this lecture

- Why review papers
- How conference reviews work
- Role of a reviewer
- Structure of a paper review
- Example reviews
- How to debate merits of a paper
  - Get ready for proponent/opponent discussion!

# Why Review Papers?

## Reason 1: Testing your own comprehension

- Noting contributions, significance, strengths, and weaknesses
- Identifying promising areas for future work

## Reason 2: Group meetings / reading groups

- Similar to above, but to promote discussion within group

## Reason 3: Related work in your papers

- Can be thought of as *very* concise paper reviews
- Summarize main technical points, compare/contrast with your work

## Reason 4: Conference and journal reviews

- Peer review is used to judge the merit of scientific papers
- Reviews influence accept/reject decision and author revisions

# Purpose of a Conference Review

A conference paper review serves many purposes:

- **Synthesizes** the reviewer's understanding of the paper
- **Communicates** the reviewer's thoughts about the paper to other PC members and the PC chair
- Partially **documents** the PC's decision to accept/reject the paper
- **Provides guidance** to the authors regarding possible (or mandatory!) improvements to their work

As a result, the review is important at all stages of the process

**Bottom line:** A paper review should *not* be a book report!

# Reviewer Guideline Examples

- NeurIPS:  
<https://nips.cc/Conferences/2020/PaperInformation/ReviewerGuidelines>
- CVPR: <http://cvpr2020.thecvf.com/submission/main-conference/reviewer-guidelines>
- CHI: <https://chi2021.acm.org/for-authors/presenting/papers/guide-to-reviewing-papers>
- ISCA:  
<https://www.iscaconf.org/isca2020/submit/reviewers.php>

# Example Review Questions (ICCV'21)

1. [Summary] In 3-5 sentences, describe the key ideas, experiments, and their significance.
2. [Strengths] What are the strengths of the paper? Clearly explain why these aspects of the paper are valuable.
3. [Weaknesses] What are the weaknesses of the paper? Clearly explain why these aspects of the paper are weak. Please make the comments very concrete based on facts (e.g. list relevant citations if you feel the ideas are not novel). If applicable, please indicate key issues and questions which, if well addressed during the 1-page rebuttal, might influence you to change your rating.
4. [Overall rating] Paper rating (pre-rebuttal)
  - Strong Accept / Weak Accept / Borderline / Weak Reject / Strong Reject
5. [Detailed comments] Additional comments regarding the paper (e.g. typos, any suggestions to make the submission stronger).
6. [Reproducibility] Is the method described in this paper reproducible?
7. [Confidence] Reviewer's confidence in his/her recommendation
8. Please provide an "Overall Rating", following the rebuttal and reviewers discussions.
  - Strong Accept / Accept / Leaning to Accept / Learning to Reject / Reject / Strong Reject
9. Justification of final rating. Describe the rationale for your final rating, including notes based on the rebuttal, discussion, and other reviews.

# Scores in Reviews

- Comments that reviewers write are important; also important are scores
- CVPR reviewing scale:
  - Strong Accept, Weak Accept, Borderline, Weak Reject, Strong Reject
- NeurIPS reviewing scale:
  - 10: Top 5% of accepted NeurIPS papers. Truly groundbreaking work.
  - 9: Top 15% of accepted NeurIPS papers. An excellent submission; a strong accept.
  - 8: Top 50% of accepted NeurIPS papers. A very good submission; a clear accept.
  - 7: A good submission; accept. I vote for accepting this submission, although I would not be upset if it were rejected.
  - 6: Marginally above the acceptance threshold. I tend to vote for accepting this submission, but rejecting it would not be that bad.
  - 5: Marginally below the acceptance threshold. I tend to vote for rejecting this submission, but accepting it would not be that bad.
  - 4: An okay submission, but not good enough; a reject. I vote for rejecting this submission, although I would not be upset if it were accepted.
  - 3: A clear reject: I vote and argue for rejecting this submission.
  - 2: I'm surprised this work was submitted to NeurIPS; a strong reject.
  - 1: Trivial or wrong or already known.

# Reviews Followed by Rebuttals

- A rebuttal is a short document (1-page) where authors address concerns by reviewers; typically short deadline for conferences
- Typically in question-answer format
- The goal is to rebut factual errors on the reviewers' part, and alleviate concerns for on-the-fence reviewers
- Especially important for borderline papers (WA, B, WR) or (WA/WR, B, B) or (B, B, B)

# Conference Program Committees

- CVPR, the top conference in computer vision, is probably on the huge end
- Program chairs: 4-5
- Area chairs (managing and moderating reviewing): ~250
- Reviewer count: on the order of *thousands*
- Despite being at the bottom of the hierarchy, reviewers have the most important work!

# How Conferences Work

- How much time does a reviewer spend with a paper?
- How much time does the area chair spend with it?
- Which papers does the area chair spend more/less time with?
- What does this imply for how papers, reviews, and rebuttals should be written?

# Review Score Scenarios

- What do you think are common review score distributions in CVPR?
- Which papers (with what scores) get accepted?
  - What happens to a paper with (WA, WA, WA)?
  - What happens to a paper with (WA, B, WR)?
  - What happens to a paper with (WA, WA, SR)?
  - What happens to a paper with (B, B, B)?
- Scores and decisions for one conference:
  - <https://github.com/evanzd/ICLR2021-OpenReviewData>
  - [https://docs.google.com/spreadsheets/d/1MLlgV82\\_4K1FJGSjUKm2R8cw5msn4xmrhtnS9FLAS6k/edit#gid=189496698](https://docs.google.com/spreadsheets/d/1MLlgV82_4K1FJGSjUKm2R8cw5msn4xmrhtnS9FLAS6k/edit#gid=189496698)

# Importance of a Reviewer

- Say a student will publish 5 papers (1 per year) to be granted their PhD
- Say it takes 2 attempts to publish one paper
- Say each attempt is reviewed by 3 reviews
- Thus, that PhD student will receive 30 paper reviews during their PhD
- One reviewer is responsible for 1/30-th of someone's PhD career (or more if working in related area hence common reviewer!)

# Role of a Reviewer

- A reviewer is a judge—they safeguard the quality of a publication venue
- A reviewer also gives feedback and helps steer a particular work (if suggestions that reviewer gives are sensible)
- Both of these imply reviewing should be taken very seriously!
- Important: Keep papers you review *confidential*, to protect the authors' work!

# A Reviewer Should:

- Read the paper once (1 hour)
- Read it carefully a second time (1-2 hours)
- Take notes about strengths, weaknesses, points that seem unclear, problematic, even typos (if time)
- Use notes to fill out a review template (~15-30 min)
- Be responsive to area chair requests for discussion
- Read rebuttal, others' reviews, and complete final scores (if there is a rebuttal phase)

# Note on Reviewer Anonymity

- Reviewer identity is hidden from authors, but not hidden from program and area chairs
- Thus, apart from honor and karma, you should do a good job at reviewing because these chairs should have a good opinion of you for later job opportunities
- Sometimes, fellow reviewer identities are also revealed (if reviewing on the same paper)

# Structure of a Paper Review

Technical summary

Description of contributions

Major critiques

- Strengths
- Weaknesses
- Questions

Minor points

Concluding remarks

## *Content:*

- Very short (1-2 paragraphs)
- Overview of the paper

## *Purpose:*

- It provides the **reviewer** with context for the review
- It allows the **PC chair** to get a quick synopsis of the paper
- It convinces the **author** that the reviewer read and understands the paper

# Structure of a Paper Review

Technical summary

Description of contributions

Major critiques

- Strengths
- Weaknesses
- Questions

Minor points

Concluding remarks

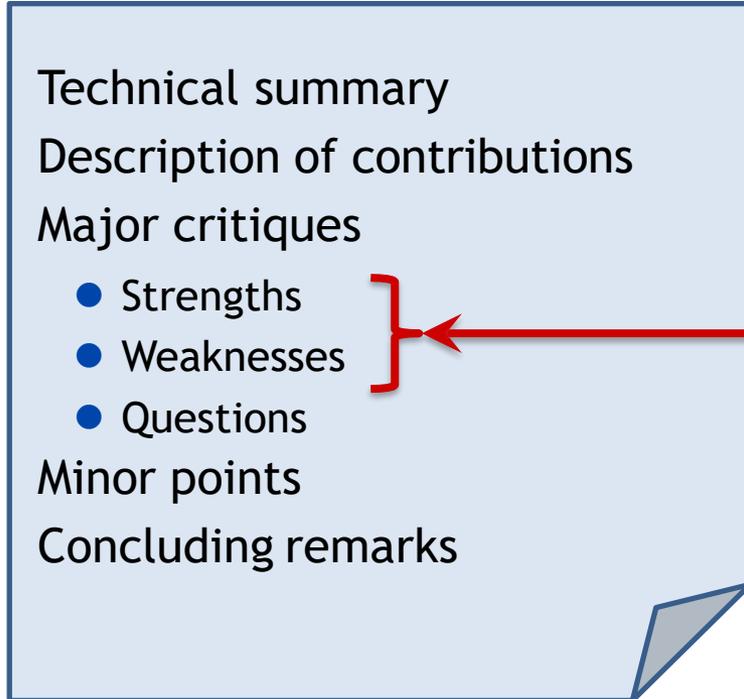
## *Content:*

- Very short (1-2 paragraphs)
- Quick summary of the **novel aspects** of the paper

## *Purpose:*

- Novelty is **paramount!** This provides evidence for the final accept/reject decision
- Again, it convinces the author that the reviewer understands the novelty of their contribution
- It sets the stage for detailed critiques

# Structure of a Paper Review



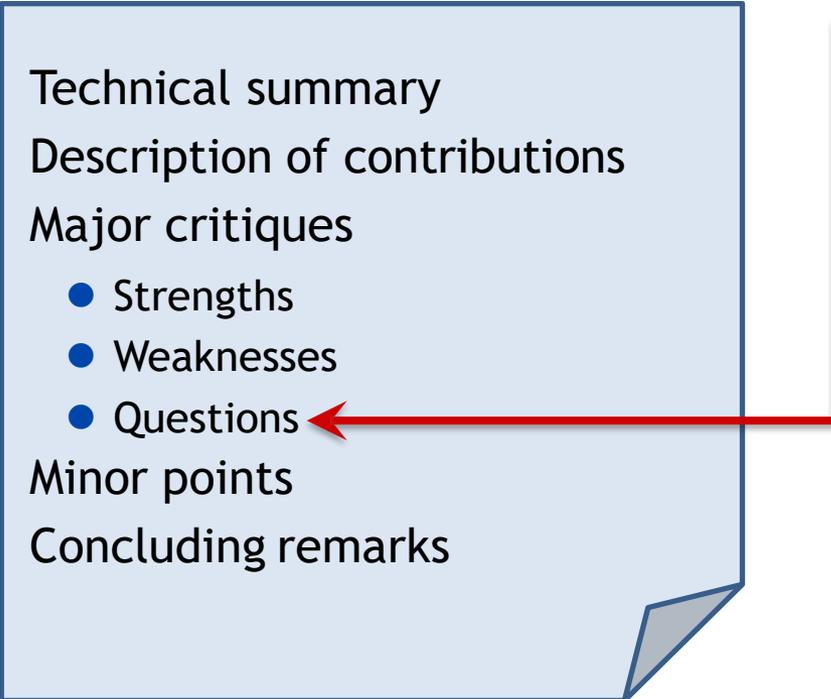
## Content:

- **Technical** and/or **methodological** strengths and weaknesses
- *Examples:*
  - How interesting is the problem?
  - Novel proof techniques or solutions
  - Missing related work
  - Assessment of the (in)completeness of the evaluation
  - ...

## Purpose:

- Primary assessment of the paper
  - Do the authors bring something of intellectual value to the table?
  - Is the paper somewhat **incremental**, but well executed?
  - Does the paper have **fatal flaws**?
- Typically, this provides **fodder for discussion** at the PC meeting<sup>18</sup>

# Structure of a Paper Review



Technical summary  
Description of contributions  
Major critiques

- Strengths
- Weaknesses
- Questions

Minor points  
Concluding remarks

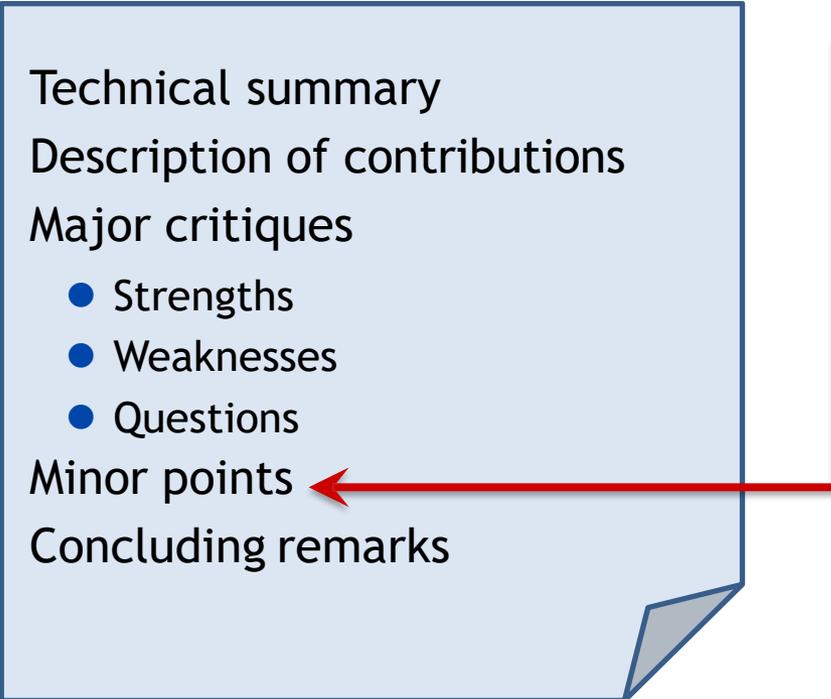
## *Content:*

- Remarks on any thing that was unclear in the paper

## *Purpose:*

- To stimulate discussion with **other reviewers**
- To inform the **author** of questions that remain unanswered after reading the paper
- May be addressed by authors in rebuttal

# Structure of a Paper Review



Technical summary  
Description of contributions  
Major critiques

- Strengths
- Weaknesses
- Questions

Minor points  
Concluding remarks

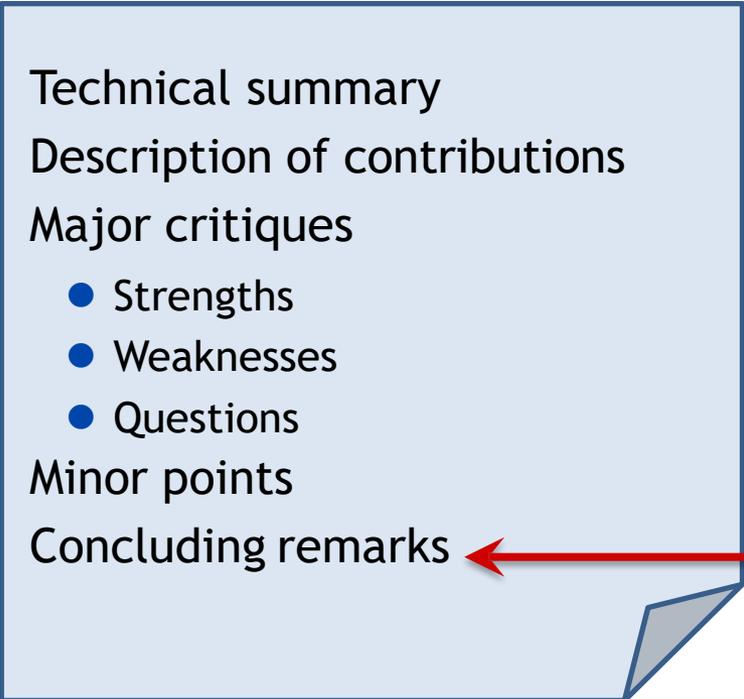
## *Content:*

- Aspects of the paper that don't influence the novelty of the contribution, but do impact the quality of the paper overall
- *Examples:*
  - Typos and grammar errors
  - Suggestions for better examples
  - Corrections to minor logical flaws
  - ...

## *Purpose:*

- Helpful for planning revisions

# Structure of a Paper Review



Technical summary  
Description of contributions  
Major critiques

- Strengths
- Weaknesses
- Questions

Minor points  
Concluding remarks

## *Content:*

- Very short (1-2 paragraphs)
- Final assessment of paper, with justification

## *Purpose:*

- Provide final suggestions
- Communicate your views on the paper to others
  - You might love the paper, yet make many negative critiques
  - You might hate the paper, yet say some positive things about it
  - This is where you clarify

# Being a Good Reviewer

Acceptance rates at good conferences are very low

- < 15% is not unusual at competitive venues
- < 10% not unheard of!

This leads to the following situation

- **Problem:** Writing a good review takes time
- **Problem:** PC members often must review many papers
- **Problem:** Very few papers can be accepted anyway
- **“Solution”:** Look for reasons to reject a paper and be done with it

This is extremely counterproductive, and not good for science (**Why?**)

Hill and McKinley\* offer suggestions on avoiding this type of pitfalls

# Avoiding Pitfalls

	Pitfall	Recommendation
1	Seek to find all flaws in the paper, in part to show your expertise as a reviewer.	Look for reasons to accept a paper. Despite its flaws, does it point in new directions or expose promising insights? The community can benefit from imperfect, insightful papers.
2	Since the review process is anonymous, it is appropriate to criticize the paper as if the authors did not have feelings.	Your tone should be the same as if you are giving comments to a colleague face-to-face. It is always possible to be constructive, focus on the work, and do not attack the researchers behind it. The purpose of a review is not only for selecting papers, but to improve the quality of all the work in our area.
3	Reject papers that build on recently-published new directions but accept those that build on the established norm.	While truly new papers are best (and rare), consider accepting papers that follow-up on recently-published promising directions. These papers allow the community to explore ideas that can not be fully-developed in one paper.
4	Advocate rejecting a paper with little comment, because it is obvious that all will agree with you. Ditto for accept.	Explain why you advocate a rejection or acceptance, because people will often disagree with you. Your explanations will make you a more effective advocate or detractor for the paper.
5	Advocate rejecting (almost) all papers to show how tough you are.	Your job is to decide what is best which is not usually accomplished by rejecting every submission.
6	Advocate rejecting a paper because you seem to remember it being the same as (or similar to) unidentified prior work.	In this situation, the professional should reference important prior work after refreshing one's memory regarding what it contains. One missing reference is usually not a reason to reject a paper.

# Common Reviewing Pitfalls

- Saying something has been done before, but not citing papers
- Citing your own work—don't want to reveal your identity, and chances are, other work is relevant and can be mentioned too, if you think the authors made major omissions in citing relevant work
- Being unreasonably picky about writing, method and experiments—you should be picky, but pay attention to contributions too, and be willing to forgive small mistakes
- Not being critical enough, or being too critical

# Common Paper Strengths

- Good writing
- Methods is sufficiently novel
- Method outperforms chosen baselines (by a large margin?)
- Experiments are extensive
- Application is compelling

# Common Paper Weaknesses

- Not novel (this or very similar techniques have been proposed before)
- Missed relevant baselines (competitor methods)
- Hard to tell where improvement over baselines comes from—need ablation experiments
- Doesn't improve over baselines by a lot, very small gain over the most similar methods
- Not well written, method not clear, motivation not clear

# Discussion of a Paper

- AC may call for a discussion (on reviewing platform)
- Reviewers bring up especially important points
- Other reviewers agree/disagree with whether an issue raised is truly problematic, whether it is a critical problem, etc
- Sometimes/often a review will change their mind, based on rebuttal, or based on other reviewer's comments
- It's ok to change your mind, if you have a good reason—don't be afraid to
- Also don't be afraid to be the only person being very positive about a paper—and provide arguments

# Common Reviewing Consoles (in AI)

- CMT
- OpenReview
- EasyChair
- SoftConf

# Example Reviews (ICLR)

<https://github.com/evanzd/ICLR2021-OpenReviewData>

Look at some of the course readings (Gestalt, Referee)

# Activity 1 (20+20 min)

- Choose a paper from our list of iconic papers
- Do a review using the ICCV 2021 review format (20 min)
- A few of you will share your review with the class (5 min for each of 2-3 papers)
- Class will ask further questions, present opposing viewpoints (3 min per paper)

# Activity 2: Proponent/Opponent

- Proponent presents strengths of chosen paper
- Opponent presents weaknesses
- Proponent attempts to defend paper by addressing each weakness, opponent attacks each strength (if possible)
- Class asks further questions
- At the end, class chooses who did a better job, proponent or opponent (there may be a bias, it's ok)

What did you learn from yours,  
others' presentations?