Start recording …

Admin

- **Draft due March 24th**
- Session moderators for today: **Nobody :(**
  - https://docs.google.com/spreadsheets/d/1dbmlvduupZUCDjxU4HW2_350OVrVG-g1FoEAG-uWhMk
- Speakers feel free to share your pdfs of your presentations with me

Plan

- **One talk today:**
  - Abbas Masoumzadeh
  - One 40 min slot, including 10 to 20 mins of questions / suggestions
- After that:
  - Peer review plan
  - How to write a good review

Peer review plan

- You will review 2 papers (project drafts) each
- Every project will get at least 4 peer-reviews and one instructor review
- Your job is to help your classmates
- You will be marked on your ability to write a good review

How to write a good review

Be constructive

The author is King/Queen

- **They are the ones doing the research—the hard work**
- You are merely critiquing in a state of partial information
- You want to be:
  - **Accurate**: make a fair and well calibrated assessment of quality and contribution
  - **Helpful**: the paper should improve with your advice

Most reviewer are poor

- **They violate basic principles of good reviewing**
- **They are overconfident and unmeasured**
- **Rarely constructive**
- **Follow the fashions of the field**
- **Done in a rush with little effort or thought**

This is not a competition

- It’s not you against the authors
- Many papers will be accepted
- Don’t view others working on your topic as competition:
  - It’s a sign you are working on the right research topic
  - Scoping almost never happens

You can easily become a great reviewer

- Follow the basic advice here and try and do a good job that’s it!
  - You will become a prize winning reviewer
- It’s critical to the long term health of the field
- It’s a huge part of science
- You will be rewarded!
- We are at a crisis point!

Be Kind

- It is really easy to get frustrated reading papers
  - Many are indeed incorrect, broken, or ‘student projects’
- **Remember**: everyone is actually trying to write good papers
- Think of the kind of reviews you want for your papers
- Do not start off assuming “this paper is wrong, let’s look for reasons to reject”
  - Be neutral

Don’t cheat

- We have double blind for a reason
- Yes people:
  - Put their papers online
  - Tweet about them
  - Write in a style that reveals who they are
- Don’t look up the paper online.
- Remember: famous people submit bad papers sometimes

My workflow

- Load all assigned papers into iPad (or print them)
- Read each of them slowly and carefully
- Make lots of notes
  - Including a decision of reject, middle, or accept
- On a different day start writing the reviews
  - This will require skimming the paper again and reading your notes

Two stage process

- This makes you calibrate across all the papers you review before writing the review
- This forces you to leave enough time to complete your review
- If you got mad or emotional reading the paper, that should be gone now

Never be late

- The deadline is a hard deadline
- Being late creates extra work for ACs, emergency reviewers, and slows everything down
- It’s totally asocial behavior
- I don’t care if your friends, supervisor, or Rich Sutton himself submits late reviews:
  - DON’T BE LATE

What to look for while reading

- Does the intro establish a clear problem of study? A clear hole that needs filling
- Does the intro clearly articulate measurable contributions
  - If the papers says these are our contributions, check them
- Is there a clear sense the authors are masters of the topic and cover the literature well and concisely?
  - Remember no lists of related work!

What to look for while reading (2)

- General polish: spelling, grammar, formatting, readable figures and plots, reference style and usage
- Over-claiming
- Errors in background; undefined notation
- Clear explanations of the main ideas in technical sections:
  - Don’t assume that because you don’t understand, it is your fault

What to look for in experiments

- Not enough runs
- Missing baselines
- Bad ablations
- Hyper-parameters untuned, not described, etc
- Experiments that don’t test the main idea
- Little insight or exploration of the results: Look at my numbers!
  - Over claiming … lack of significance
  - No why this happens

Make a decision

- Most papers are weak accept/reject
- Try to land on one side of this
- Prepare a list of questions that if they were answered you could decide on accept or reject
  - If the authors answer poorly that is good info also

Review structure

- Two line summary of what the paper is about. Don’t copy paste from the abstract. This is your summary based on your understanding…very helpful for AC
- Main decision. Clearly state Accept or Reject. One line listing the main reasons
- Main argument: Go through each reason. Explain it. Give evidence. Say why it matters
- Small things: these did not impact the scoring, but is a list of typos, errors and small changes to help out the author

Short reviews are usually bad

- They typically don’t give reasons for accept, reject
- They typically appeal to unclear things like:
  - I wanted more experiments
  - Method was not complex enough
  - More theory
  - Idea was unknown, but simple in retrospect
- Reviews should have substance; reviews are typically not short
- Don’t you wish the reviewers of your paper would take more time and give more thoughtful feedback

Decision

- Example: *This paper should be rejected because: (1) the experiments do not provide clear evidence of a contribution, and (2) the paper has major notational problems. I have posed a series of questions below that will help me refine my score after author response*

- Be clear. Say the most important things first

- Know that you could be totally wrong: mentally prepare to change your score later

- These reasons are a contract with the authors: if they explain away these concerns **you should accept**

  - **Don’t move the goal posts!**

Main argument

- This is the most important part of the review
- It should be multiple paragraphs
- At least one paragraph per reason listed in the main decision
- This is where you give the evidence and understanding for why you accept of reject
- This allows the authors to point out:
  - How you misunderstood parts of the paper, algorithm, theory, experiments
  - How you misunderstood the area (not all papers will be in your area of expertise)

Main argument (2)

- Finish with or include throughout a clear set of questions
  - Clearly ask the authors to respond
- This allows calibration later
- Shows humility
- Directly communicates: *I could be wrong and I am willing to change my score*
- Sometimes there is a special section for this

Small things

- A list of things to make the paper better
- Tell the authors these things did not impact the score
- This is showing you read the paper in detail
- This also shows you are committed to helping the authors make the paper better

Take opportunities to be positive

- If you think the problem of study is interesting: say so
- If its a reject but the writing was good: say so
- Want them to keep working on this topic: say so
- Getting feedback is painful. Seeing our mistakes pointed out is painful
- We are all in this together, so encourage the authors

Common reviewer mistakes

- Not valuing research areas, approaches, or topics you would not or do not work on
- Making assumptions about what a paper looks like: “every ICML paper should …”
- Chasing fashions: don’t ask for things just because you saw them in other papers
- Stating folk knowledge from the community
- Not valuing firsts

Common reviewer mistakes (2)

- Asking for too much: open problem, new algorithm, Atari experiments, and convergence theory…in 8 pages
- Asking for something and not providing evidence it possible:
  - Bound this term (theory); make the algorithm do X
- Related: asking for things that would be another paper all by itself
- Rejecting because you think not enough people will be interested in this

Your main job is correctness

- IF the paper tackles an interesting open problem
- IF the paper covers the relevant literature
- IF the paper looks like a conference paper (polish, writing)
- Then your main job is two things:
  - Ensure it is correct
  - Ensure the contributions as stated are demonstrated

Missing citations do not always matter

- If the experiments, theory and main contributions would not change with knowledge of the missing citations, then the authors can add it in later

- Else its a big problem

- More generally: you are trying to decide if the paper as submitted would be acceptable with minor changes!!!

  - Nobody checks the papers after accept!

- No paper is perfect: don’t expect that!

What is a contribution?

- New knowledge, New understanding (including empirical)
- New or improved algorithm
- New theory result or proof technique
- Putting old things together in an interesting way
- Experiments are not contributions
  - They provide evidence of contributions
  - They help you evaluate the contributions

Author response and discussion

- Read response, other reviews, and responses to those
- Think: did they answer my questions? Did they rebuke my main concerns? Contract remember
- Did the other reviewers bring up positive and negative things I missed?
- Engage in discussion:
  - Don’t be silent, don’t agree to disagree. FIGHT!
  - Easiest way to gain respect from senior people in the field

CHANGE YOUR RATING

- Consider the following …
- Since you reviewed the paper:
  - You have read 3 or more reviews from others that are different from yours
  - The authors attempted to give additional info / explanations
  - You discussed with other reviewers and the AC (some of them senior researchers)
- Ask yourself: how likely is it that I correctly evaluated the paper?
- UPDATE YOUR REVIEW to reflect all the above!

Links to resources

- **MUST READ:**
  - https://sites.umiacs.umd.edu/elm/2016/02/01/mistakes-reviewers-make/
- https://iclr.cc/Conferences/2020/ReviewerGuide
  - Really good advice and sample reviews