

How to ML Rebuttal – A Brief Guide

[How to ML](#) is a series of non-technical ML research advice posts by [@rocket](#) and [@j_foerst](#). If you like this, check out “[How to ML Paper – A Brief Guide](#)”, follow [@HowTo_ML](#) for future posts, and [share](#) on Twitter.

Before the reviews come in

- 1) The rebuttal process begins right after you *submit* the paper. Once you have submitted your paper, start a document with *additional experiments* and other fixes that you plan to have ready during the rebuttal period. Likely you had new insights in the lead up to the deadline, but didn’t have the time to follow up on them. Write them down now, otherwise you forget them.
- 2) Next step is to *anticipate reviewer comments* and start running *additional experiments* to address them. Common (simple) examples are *additional seeds* for the experiments, missing ablations, or additional qualitative analyses.
- 3) During this period it’s easy to start thinking about the next project or constantly checking whether reviews for your paper are out. Instead, focus on *preparation* for the rebuttal period.
- 4) We also recommend *clearing your calendar* for the rebuttal period as much as possible ahead of time to avoid stress. Often, the rebuttal phase is quite short and ideally you should post rebuttals in the beginning of this period to give reviewers enough time to engage in discussions with you to give you the chance to get reviewers to support your paper stronger.

Once you have the reviews (i.e. *planning* the rebuttal)

- 1) *Maintain perspective!* Maybe the reviews were “good”, maybe they were “bad” - this does not (yet) say much about the quality of your work. Remember the process has a *lot of noise* which can work both ways. Your task now is to *separate the signal from the noise* and maintain a positive, productive outlook. Also take low-quality reviews seriously and write a strong rebuttal whose audience in this case are the other reviewers and the AC.
- 2) Immediately start a Google Doc (or similar) and copy-paste over the reviews. Share it with all co-authors.
- 3) Carefully read the reviews and *colour code* them. Suggestion: **Red** – requires extra experiments, **Orange** – requires writing changes, **Blue** – requires clarification in the rebuttal, **Green** – noteworthy praise of the paper. We also suggest **bolding** any “make-or-break” statements in the text ([example](#) document).
- 4) At this point you should have a pretty good idea what the overall situation is and what work is required / would be most impactful. Even if the paper looks unlikely to be accepted, focus on addressing the reviewers and improving the paper. We have seen reviewers really change their mind (JF had [an ICLR paper](#) that went from 3/3/8 to 8/8/8).
- 5) If this is a collaborative paper, this is a good time to have a *meeting* and plan/split the tasks. We recommend adding a “todo” with a name assigned to every highlighted part of the reviews.

- 6) Start structuring your rebuttal according to your target audience. In most cases your rebuttal will have one of the following purposes:
 - a) to get a particular reviewer to increase their support for the paper
 - b) to convince other reviewers that a particular reviewer's concerns are *unjustified* or can easily be addressed in the camera ready copy
 - c) to equip a supportive reviewer (i.e. your champion) with further evidence for the internal discussion phase to convince other reviewers that your paper is worthy of acceptance
 - d) to convince the AC that the paper is worthy of acceptance despite some reviewer's criticism

Tips / checklist for writing the rebuttal

- Check "[How to ML paper](#)" to remind yourself of common *writing issues*. Remember to be *clear and concise* above all.
- Initially drafting of the rebuttal can happen in the same (Google) doc but should ultimately use the *most expressive* format available. E.g. if a pdf is possible use overleaf, and if markdown is possible (e.g. openreview), use [hackmd.io](#) ([example](#)).
- If you can update the paper as part of the review process (e.g. NeurIPS), **never** say "we will change the paper" for writing changes. Make those changes *immediately* instead and, in the rebuttal, point the reviewers to the changes in the updated paper.
- If paper revisions can be uploaded (e.g. OpenReview), you can make it much easier for the reviewer to see your changes—and thus maximise the impact of your changes—by also uploading a red-line markup version of your paper (credit to [@MingjiJiang](#)).
- If there is an author-reviewer discussion period (e.g. ICLR & NeurIPS), post your rebuttals as *soon as possible* to give reviewers enough time to ask additional questions (and raise their score!) and for you to answer any remaining questions.
- New results should never come out of the blue (e.g. new experiments you were planning to do anyway but no reviewer asked about them)—they need to be contextualised with respect to reviewer questions or demands.
- Ask reviewers explicitly that, unless they still have any arguments against the paper or rebuttal, they improve their support for the paper (i.e. increase their score!), or to point out additional issues.
- Reemphasize *positive points* made by the reviewers. Especially if reviewers *disagree* regarding the merits, point this out. That is, make it easy for the AC to grasp the positive aspects raised by reviewers.
- Consistently use the notation and nomenclature from the paper.
- Be polite, even when you strongly disagree "We respectfully disagree with R2's point about..."
- Colour-code reviewers if you can only submit a PDF ([example](#)). On OpenReview, make use of Markdown features (boldface reviewer ids, use heading and LaTeX support). Make it easy for reviewers to visually spot which responses are addressed to them.
- Pool *common* concerns into a *common* response to all reviewers.
- Structure the rebuttal around the reviewers' main concerns, followed by responses for their minor concerns and questions ([example rebuttal](#)).

- Use present tense “We argue that” instead of “We will argue that.”
- Pro-actively* ask reviewers for ways to further improve your paper. That is, make them accomplices in arriving at a strong camera ready paper.
- Remember any changes that you promise to make, and make sure to execute on them for the Camera-Ready Copy (see below).
- It goes without saying that despite your efforts to convince reviewers and the AC to support your paper, you need to stay *scientifically honest*. Accept valid criticism for what it is—a way to improve your paper and help the scientific community to advance their understanding of the field.
- Control-f and search for your name or co-authors. Make sure you didn’t accidentally copy-paste it over.

Discussion Period

1. Keep engaging reviewers and the AC. It’s ok to send reminders and to ask for engagement. It’s *not ok* for reviewers to ignore your rebuttals and it’s not okay for the AC to be inactive during this period in case the reviewers do not engage with your rebuttal. Be specific in your requests, though.
2. Contact AC sparingly, only for severe issues: complete garbage reviews, lack of discussion, and lack of engagement by reviewers, break of code of conduct by reviewers, heavily biased & unfair reviews etc.
3. Remain optimistic! Whatever the scores, hopefully there was something you have learned about your paper and the process you just went through has helped you improve it.

After the Discussion Period

1. Update your list of future experiments and (crucially) todos for a potential camera-ready-copy (CRC).
2. Plan for a potential resubmission. Deadlines for ICML, NeurIPS, ICLR usually work out such that a rejected paper can be *resubmitted* to the next conference.
3. Keep improving your paper! Whether for the CRC or for a resubmission, most likely you now have new ideas for experiments and writing improvements. Work on this now, ahead of the next deadline.
4. Make sure to include any changes in the CRC that you promised to the reviewers and AC during discussions.

Acknowledgements

We would like to thank Edward Grefenstette, Minqi Jiang, Shimon Whiteson, and other paper co-authors for their help and advice over the years that went in distilled form into this post.

Licence

[CC BY-NC](#)