

# Checklist for writing fair and constructive Stat-ML reviews

Aaditya Ramdas (aramdas@cmu.edu), Carnegie Mellon University

Some personal suggestions (via accumulated wisdom) for my own students.

- **Why review?** The peer-review process is central to the current academic system, and to keep our journals or conferences free/cheap and high quality, we rely on people giving up their time as a service. Carefully reading papers not directly in your area can also broaden horizons and improve writing by refining your tastes.
- **Accept an invitation to review** if a paper falls in between “no clue about the area” and “conflict of interest since I’m working on exactly the same problem”. In order to provide reasoned judgment and constructive feedback, you should be familiar with existing related work, comfortable with (or interested in) the topic.
- **How many should I review?** A rough “equilibrium-based” heuristic: say each submission causes 3 reviews. So if you submit a 3-author paper, and each author agrees to review one more paper in return, there would be some balance. Thus, prolific authors would review more often, which seems fair. In reality, high quality researchers (not necessarily prolific) will be asked to review more often, and that’s also a good thing.
- **Time spent reading the paper.** When you’re younger, you might spend more time on a review, since you may need to read related work, or are slower to follow mathematical arguments in proofs. Also, papers with new and sophisticated ideas might take much longer to process, even for experts. I rarely spend less than one hour on a paper, unless it is an obvious reject, and I rarely spend more than 24 total hours on a paper. Usually, I read a paper in a few chunks of an hour or two, with gaps of hours/days to dwell on the contents.

Once you have accepted to review the paper and read it, writing a “good” review is critical.

- **What makes a good review?** Broadly speaking, it must be *respectful*, *precise* (makes pointed praise/criticism, clear judgments that are useful for editors+authors), *balanced* (talks about pros and cons, ingenuities and shortcomings), and *constructive* (regardless of outcome, how can the paper be improved?).
- **Modular reviews.** Try to comment on each section of the paper, or at least on each of the high level splits (motivation, related work, algorithms, theory, experiments, software, appendix, etc, as relevant). Decouple the major comments (wrong arguments, missed critically related work, poor baselines) from relatively minor ones (typos, choice of notation, poorly organized supplement). Aim for 2-3 pages for most reviews.
- **Subjective or objective?** A review necessarily must contain subjective and objective judgments. The editor/chair cares both about whether the argument/algorithm/theorem is correct and works as claimed (more objective), as much as whether the problem is an important one and the contribution is novel or incremental (subjective). Both types of judgments require justification from the reviewer.
- **Focus on the field, not yourself.** You have to deem whether a paper has quality contents that should be read, considered and debated by the field. Importantly, you may recommend a well-rounded or novel paper for publication even you do not agree with its position or line of attack.
- **Be considerate.** Being critical of methods or literature review or experiments or the theory or the writing is all completely fine, but attempt to make a significant part of the criticism *constructive* in nature. And highlight things that the author did well (interesting problem, novel approach or model, etc).
- **You can be wrong.** Make it clear which criticisms you are absolutely certain about, and which criticisms are hunches to which the authors deserve a chance to rebut. Admitting your fallibility as a reviewer is a good thing. Nothing is more annoying than an extremely confident review that is completely off the mark.
- **Going above and beyond.** I have received reviews in which the reviewer did their own experiments with our method, or provided a simpler/cleaner proof, pointed out relevant connections to other parts of the literature, and asked very insightful questions that deserved more thought. These are the best kinds of reviews, and while the rest of this document described a good median to aim for, it’s worth wondering how you can help the peer review process result in the best possible publication at the end.