Reviewing and rebuttal guidelines

Jan van Gemert
Delft University of Technology
Full Research Guidelines are:
`https://jvgemert.github.io/ResearchGuidelinesInDL.pdf`

1 Reviewing and rebuttal guidelines

Why review. If you are writing peer-reviewed papers, then you are asking others to devote time and effort to your work. Thus, it's fair to return the favor and review the work of others. Reviewing is part of your academic community. It offers other advantages: improving the scientific field, learning something new, practicing your critical thinking, and helping others.

What to review. I've found the "Troubling Trends in Machine Learning Scholarship" [1] paper quite helpful, where I regularly give back a review where I state that a paper is *Confusing explanation with speculation* and/or has a *Failure to identify the sources of empirical gains*. In addition, I use Hitchens's razor: "*What can be asserted without evidence can also be dismissed without evidence*. And, reviewing involves applying the research guidelines in this document, albeit not during the process, but at a finished paper. Specifically, see if you can find the storyline as in Section **??**, look for the answers that can be asked during a research meeting in Section **??**, and the writing guidelines in Section **??**. I've labeled each guideline with a unique identifier, which could help in motivating a review by referring to this document and it's identifiers.

Goal of reviewing. The main goal of reviewing is to improve the work of others by giving feedback while preventing the publication of flaws. This might include flaws in the following. Hypothesis: Are the hypotheses sound?

Literature: Are the relation to relevant other work present, correct and properly motivated why and how the work is related?. Method: Is the method aligned with the hypotheses? Technical: Are the equations correct? Does it do what is claimed? Are there no unexplained surprises? Is it reproducible, ie: code? Experimental setup: are the hypotheses evaluated? Is the motivation evaluated? Evaluation: is it an unbiased, fair, comparison to others? Clarity: is it understandable? Figures/tables readable? it's OK if there are minor spelling/grammar mistakes, as long the paper can reasonably be well understood without too much puzzling. As a reviewer, it typically cannot be expected to rerun experiments, a review is inherently based on trust in the author's integrity. A perfect paper does not exist, all papers are limited in some sense. Thus, be critical, but appreciate the positives.

What is a good paper. It's a solid brick that others can build on: something is learned. It's well written with an intuitive motivation, for example in Fig 1. It has clearly specified hypotheses, research questions, and contributions. The method aligns well with hypotheses. The Hypotheses, research questions and contributions are backed up by empirical evidence. Comparison experiments vary only 1 variable. It has experiments on several datasets to illustrate generalization. Bold numbers are never a goal in itself, they are 'only' important to show relevance/usefulness. It's reproduceable, it has clear algorithms, or better: code.

Addressing novelty. It's easy to do something novel: merely add a layer, and it's 'novel'. Novelty is not a goal in itself: a paper about my left thumb is extremely novel, but that does not make it, nor the paper, interesting. It is up to the paper under review to explain in the introduction and related work sections how it relates to others, and what contributions it has compared to the other work. Just because the outcome is "obvious", or "trivial" is not good grounds for rejection: the paper has now confirmed this outcome; and this confirmation is a contribution in itself. A possible ground for rejection could be if there is other near-identical work but not placed in relation to the paper under review and/or if it is not experimentally compared against.

Review quality. A bad review: makes claims without giving details, citations, or motivation. Is only a few lines. Only checks the bold numbers. A good review gives constructive author feedback, so in addition to What,

it also suggests How to change a paper. Is well motivated, with detailed justification (citations / line numbers). Is well-written and self-contained: the review is readable without the paper. It makes the decision for the AC easier.

My review structure Review formats vary slightly for each conference. I always use this layout in a .txt file, which can be poured in any format. When I am reading the paper for the first time, I directly write comments per line. Once these detailed comments per line number are there, then the other points follow from them. Review structure:

- Summary. Unbiased, the authors should agree with it, introduce terms that you will build on later so that the review can be self-contained.
- List of positive points: just 1 line per point
- List of negative points: just 1 line per point
- A conclusion paragraph of score motivation and main suggestions for what the paper could address in the rebuttal. This paragraph builds on the summary and the list of positive/negatives.
- Detailed comments per line number with detailed justification.

1.1 Rebuttal guidelines

Why write a rebuttal. Several conferences, and journals, allow for a rebuttal: a factual response to the reviewers. The main goal of a rebuttal is to correct mistakes, and convince the reviewers your work is interesting. Even if there is only 1 positive reviewer it can help to write a rebuttal: there often is a discussion phase, where your "champion" can then defend the paper.

Not novel Typically, when reviewers write that there is insufficent novelty –without citing a missing paper– then what they really mean is that, unfortunately, they did not find the paper interesting. Yes; finding something interesting is subjective; and this is probably why they write 'not novel' instead of 'I didn't find it interesting'. I suspect that reviewers are afraid to be honest because it is not possible to give objective reasons for why something is not interesting (to you). This is unfortunate, because as an author this

would be valuable feedback to have. If reviewers find it not novel, then try to ask yourself why they didn't find it interesting. Try to also ask a (somewhat senior?) not co-author colleague, who is not afraid to tell you the truth to your face.

Science is done by humans Many reviewers: do their work too fast; have a rejection mentality; do not read the paper well; write a too short and unmotivated review; or base the review on the author's name if a pre-print or blog is available. Receiving a 'bad review' can be quite frustrating; especially when you spent all that effort on your paper. The only advice I can give is try to learn something from the review anyway. If the reviewer did not understand the paper: what can be improved? How can the paper be made easier to parse? How to improve "something between the lines" that they did not like?

Because science is done by humans, its also important to address the reviewers as human beings:

- Always thank the reviewers (Don't "over-thank").
- Assume they will not change their mind more than 1 point (It might happen, but is psychologically difficult)
- If technically possible: Do what they ask, even if it doesn't make sense (to you). The most convincing response is to just show it.
- If you fight/shout/insult: they will fight back in the discussion; and they will have the last word.
- Write for reviewers and area-chair/editor; having the reviewers on your side is much easier to get accepted.
- Make it easy for the reviewer to find their answer (do not 'hide' the answer somewhere in a lot of text). Thus, try to copy the comment of the reviewer verbatim
- Write the rebuttal self-contained: ideally, they should not need to go back to the paper, nor to any of the reviews.
- Do not take reviews personal (you are not your work)
- Reply positive, non-defensive and to the point

- Be polite and professional, but self-assured and firm
- Long and too dense rebuttals will scare reviewers away. Leave sufficient white space.

My approach to writing a rebuttal is as follows:

1. Copy-paste all concrete positive and negative points in a doc
2. Answer each negative point
3. Perform all requested experiments (to good approximation)
4. Group similar (positives and negative) points
5. Start by summarizing grouped (verbatim) positive comments
6. Answer grouped (verbatim) negative comments
7. Rephrase negative answers and compactly rewrite
8. Decide which answers to drop strategy (eg: Convince one reviewer, but keep others).
9. Ask someone else to read rebuttal and ask how they feel

References

[1] Zachary C. Lipton and Jacob Steinhardt. Troubling trends in machine learning scholarship. arXiv, 2018.