

Reviewing a Paper for a Top Data Science Conference

Background

ICML, SIGIR, IJCAI, and UAI are all top tier conferences in computer science in general and artificial intelligence in particular. Acceptance rate for these conferences is typically around 15-30%, which is relatively low, in particular considering that most submissions are high-quality state-of-the-art papers from top institutes in the US and Europe. Just as an example, people like Stuart Russell, one of the authors of the leading textbook on artificial intelligence, and Judea Pearl, winner of the Turing award, regularly submit papers to UAI, without having any guarantee that these will get in. Some papers can be quite technical and properly reviewing them can be quite a challenge. Just reading the paper will not be enough: you will have to collect and read related literature to be able to understand the paper and to judge whether the paper is really novel and presents a relevant advancement.

Proposed plan of attack

1) Read the paper, make a list of relevant literature that you feel you should read as well, and meet with your coach to discuss/ask questions.

At first, you may have a hard time reading and understanding your paper. Focus on the main concepts and search for earlier papers that explain these. Your coach may be able to help and give you pointers if you can't find these yourself. Often the papers cited by your paper form a good starting point. But you also have to search for other papers, e.g., by entering keywords into Google scholar <https://scholar.google.com/> to try and find similar papers that you then may also need to judge the novelty of your paper, but also to find, for example, review papers on the same topic. Make sure to write down any questions that you have and discuss them with your coach. Get in touch right away when you feel that you're stuck. Note that submissions may contain supplementary materials, e.g., giving a proof of a theorem in the main paper. For most conferences, you are under no obligations to read the supplementary materials, and you should base your review primarily on the main paper.

2) Collect and read relevant literature Meet and discuss.

There is a vast literature and grasping the main concepts that the paper depends upon may be quite challenging. Again, write down the things that you do not understand and ask your coach. By the end of this phase, you should be able to read your paper and put it into the context of previous literature.

3) Write a preliminary review. Meet and discuss.

Start forming an opinion of your paper and start filling in your answer to the various questions in the review form. Discuss this with your coach: he or she will not tell you whether (s)he agrees with your judgement, but may be able to help you to substantiate your opinion.

4) Finalize review. Meet and discuss.

Go over the paper once more. Make sure that your English is flawless and your criticism on the paper is constructive. Give concrete hints on how (you feel) the authors can improve their paper. At this point we will freeze your review: this is your review that later will be graded. It is up to your coach to decide whether this review is sufficiently strong and independent to be incorporated as an additional review for the conference (when applicable). If there are valid points in your review, but the review as a whole is not there yet (which, for the record, does not mean that it is not good enough for the course!), you and your coach may discuss if and how your coach can integrate some of your points into his own review.

5) Author feedback and reviewer discussion.

Typically, authors can reply to the reviews. Reviewers and meta-reviewer then discuss potential discrepancies, in particular for borderline papers. Your coach may involve you in the discussion, in particular when you contributed to the review. In any case, ask your coach to send you a record of the discussion (if any), which you can then incorporate in your reflection.

6) Final decision and reflection.

At some point, the accepted papers will be announced and you can check whether the paper that you reviewed got accepted. After that, you can write a reflection on the review procedure. How did your review relate to the other reviews on your paper? What's your view on the discussion (if any) between the reviewers? Do you agree with the final decision? This reflection (no more than half a page) together with your review should be submitted through Blackboard and will be graded.

Note that you do not necessarily have to meet the official review deadline for the conference: if you need more time, that's perfectly fine, but your review can then not be incorporated in the official review procedure. In any case, you should submit your review by June 1, the latest, to your coach. Only after you submitted your review, you will receive the reviews of the other reviewers as well as a record of the discussions, so that you can incorporate these in your reflection. Final deadline is then June 17 for your review and discussion, which can be submitted through Blackboard. In any case, do keep in touch with your coach and regularly inform him or her of your progress.

You will probably notice that serious reviewing is not an easy task, also because the papers are quite difficult and you are probably not (yet) an expert. Try and consider this a challenge and do not let this frustrate you. Your coach, who is an expert (otherwise he or she wouldn't get the paper to review...) should be able to help, but it is up to you to ask the right questions. If you're stuck nevertheless, do not hesitate to call on me.

Literature on writing reviews

How to review a paper: <http://www.inf.ethz.ch/personal/troscoe/pubs/review-writing.pdf> or <https://www.cs.princeton.edu/~jrex/teaching/spring2005/fft/reviewing.html>.

How not to review a paper:

<http://www.sigmod.org/publications/sigmod-record/0812/p100.open.cormode.pdf>.

The NIPS consistency experiment, which received a lot of attention from (computer) scientists, described on <http://blog.mrtz.org/2014/12/15/the-nips-experiment.html>, may offer inspiration for your reflection on the final decision.

Confidentiality

Paper reviewing for top conferences is typically double blind: the authors do not know who the reviewers are and the reviewers should not know who the authors are. In some cases, it may be relatively easy to identify the authors of your paper, but you should still not make any attempts to do so. You are also not supposed to share the paper with someone else and cannot make use of any of the ideas presented in the paper for your own research, unless after the paper has been published. The author feedback and reviewer discussion are also confidential, but not always blind among the reviewers: you may hear who the other reviewers and the meta-reviewer are.

Good luck!!!