How to write a review

Ulle Endriss
Institute for Logic, Language and Computation
University of Amsterdam

[Invited talk at AAMAS-2025 Doctoral Consortium
Detroit, Michigan, 19 May 2025]

Plan for Today

- Organisation of Peer Review for AI Conferences and Journals
- Best Practices for Peer Review
- Challenges for Peer Review Policy

Peer Review at Conferences

Roles at large AI conferences (some omitted for smaller conferences):

- Programme Committee Chair(s)
- Area Chairs (AC)
- Senior Programme Committee Members (SPC)
- (Regular) Programme Committee Members (PC)
- Auxiliary Reviewers

Some variation in terminology across subcommunities.

Process: Bidding $\rightarrow$ Assignment $\rightarrow$ Reviews $\rightarrow$
(Rebuttal) $\rightarrow$ Discussion $\rightarrow$ (Metareview) $\rightarrow$ Decision

Outcomes: Accept | Reject | (Short)

Peer Review at Journals

Roles at most journals:

- Editor(s)-in-Chief
- Associate Editors (AE)
- Reviewers (possibly from Editorial Board)

Process: Desk Reject? $\rightarrow$ AE Assignment $\rightarrow$ Review Invitations $\rightarrow$ Reviews $\rightarrow$ AE Recommendation $\rightarrow$ Decision

Outcomes: Accept | Minor Revisions | Major Revisions | Reject

Who can see who?

single-blind | double-blind | triple-blind

Anonymity Levels

- single-blind reviewing: authors don't know reviewer identity
- double-blind reviewing: reviewers also don't know author identity
- triple-blind reviewing: reviewers also don't know of each other

Most AI (but not, e.g., Philosophy) journals use *single-blind* reviewing.
Reason: Reviewers anonymous to avoid peer pressure. Authors not because most submissions are based on earlier conference publications.

Most AI conferences use *double-blind* reviewing.
Reason: Reviewers should not favour well-known authors/institutions.
Concerns: Never perfect. Also: knowing author identity can help.

Some conferences [unfortunately] switching to *triple-blind* reviewing.
Reason: Avoid peer pressure. Empower junior reviewers to speak up.
Concerns: Makes discussion harder/unpleasant. Less accountability.

Why do we review?

Why do we review?

To help *select* good and formally correct work for publication, so others can focus their attention on these (hopefully) high-quality papers.

To help authors *improve* their work before it gets published.

And of course there are (perfectly honourable) *selfish* motives:

- Find out about hot topics
- Learn what 'works' for your own papers
- Learn from fellow reviewers how to evaluate research
- Build a professional profile
- Sometimes it's just good fun

Review Criteria

Relevance | Clarity | Significance | Originality
Soundness | Reproducibility | Scholarship | Presentation

Review Forms

Always present:

- Comments for the authors
- Confidential comments for editors / PC members
- Recommendation (journals) or score (conferences)

Sometimes present (for conferences):

- Summary of paper
- Strengths and weaknesses
- Detailed comments
- Specific questions (for rebuttal)
- Scores (or even fields) for specific criteria

Simple review forms generally work best.

Review Structure

Summarise paper *contribution* in your own words (1 paragraph). *Why?*

- to signal that you've made a sincere effort to understand the paper
- to help chairs/AE quickly grasp what the paper is about
- to enable authors to clarify possible misunderstandings
- maybe: to suggest better ways to authors to explain significance

Then organise the points you want to make *by importance:*

- first: criticism aimed at the overall approach taken
- then: discussion of specific technical issues
- finally: remarks on presentation quality
- optional service: (partial) list of typos etc.

This tends to work better than structuring *by criterion* (also possible).

And: *Don't forget about the positive feedback!*

Review Length

Of course, there are no rules about this, nor should there be. Simply provide enough detail to be helpful (to authors and decision makers).

Very *short reviews* are almost always unprofessional.

But using an LLM to generate *lots of extra text* is even worse.

Anecdote: For AAMAS-2021, the average review length was just over 600 words/review, with clear correlation between length and quality.

Scientific Integrity & Professionalism

What are some of the things to look out for? What can go wrong?

Scientific Integrity

Avoid conflicts of interest: Don't review the work of friends.

Respect confidentiality: Submission content is confidential. Obviously, don't reject the papers of rivals to then scoop them.

Be open-minded: Holding a negative personal opinion about a specific subarea is not a valid reason to reject papers belonging to that subarea.

Don't abuse power: Never ask authors to cite your own work.
(To be safe, simply avoid discussing your own work in your reviews.)

Don't auto-generate: Your review should reflect your own professional judgment, so you must not (even partly) delegate this task to an LLM.

Respect author anonymity: Often you can guess who the authors are, but don't actively try to break anonymity by searching for preprints.

Respect reviewer anonymity: Don't undermine the system by (broadly) revealing to others that you reviewed (accepted or rejected) paper X.

Professional Conduct

Respond *quickly* to invitations to review (certainly if you will decline).

Return reviews *on time* (an absolute must for conferences!).

Carefully *edit* your review to make sure it's understandable to others.
But: no need to write perfect prose (it's not a paper!).

Avoid adding new criticism *after the rebuttal* (sometimes unavoidable).

Always express your criticism in *polite* and professional language.

Keep in mind that the *final decision* is a collective one. Don't put your fist down too early. Avoid words *accept* or *reject* in author comments.

If a claimed result is not really new, give *full bibliographic references*.
Just writing "this is a known result" is clearly unprofessional.

Reward *accessible writing:* easy to understand $\neq$ trivial

Help authors who are *non-native speakers* of English (within limits).

Who should write the reviews?

Duties and Rights

Who is *obliged* to review?

- Everyone publishing (so receiving reviews) must contribute.
- Exceptions are possible: other service; personal issues; in training

Who should be *permitted* to review?

- Experience *receiving* reviews *before writing* them yourself
- First stage: review *individual papers* someone specifically selected for you (auxiliary reviewer) and discuss your review with a mentor
- Second stage: review a whole *bunch of papers* only roughly matching your expertise (PC member)

Having completed the PhD is usually the right moment for stage two. Indeed, many conferences (very reasonably!) insist on the PhD.

Anecdote: Some of the best reviews tend to come from PhD students (as auxiliary reviewers). Note how this is consistent with the above.

Should you get paid for reviewing?

Should you get paid for reviewing?

Clearly no. Any employer expecting you to publish should also expect you to contribute to service. So you're being paid for reviewing already.

So should your employer get paid for your reviews? Maybe, but your employer must also pay for the infrastructure allowing you to publish. Better to limit such circular flow of money.

But what about them evil commercial publishers? They should charge our employers much less. Asking for payment might be a valid threat.

Should reviewers consult the supplementary material?

Should reviewers consult the supplementary material?

Authors cannot expect reviewers to go through all of it systematically.

Instead: if a reviewer wants to verify correctness of a specific claim, the supplementary material should provide clear and direct answers.

Papers that require more to be understandable should go to journals.

How far should we go to ensure anonymity?

How far should we go to ensure anonymity?

Perfect anonymity is impossible. But let's aim for decent anonymity!

Some specific suggestions (some of which are controversial):

- Reviewers should not actively try to discover author identities.
- Authors should refrain from aggressively promoting papers on social media while they are under (double-blind) review.
- ArXiv is a wonderful tool. But authors should refrain from posting preprints close to reviewing period. ($\hookrightarrow$ *ACL anonymity period)

Conclusion: How to write a review

Topics discussed:

Organisation | Best Practices | Policy

[http://www.illc.uva.nl/~ulle/how-to/write-a-review/]