How to [read, present, review] a research paper

Perdita Stevens

School of Informatics
University of Edinburgh

How to read a research paper

Begin at the beginning...

... but don't go on to the end and then stop.

If in doubt about whether I need to read a paper, I typically do something like:

1. read abstract carefully
2. read introduction quickly
3. read conclusions quickly
4. look at references
5. skim rest of paper

and then go back and start again if I do want to read it.

Can help to articulate explicitly what questions you're trying to answer in your reading.

Problems that can arise

The main symptom in all cases is that your eyes glaze over. Consider whether you:

- don't understand the area and how it fits in
  - consider reading something else short first (Wikipedia articles can be useful, for SE fields)
- don't understand the contribution of the paper
  - for an accepted paper, it must be spelled out somewhere, but sometimes not very prominently. Can help to find the statement, underline it and then read with that concretely in mind.
- don't understand the technicalities/mathematics
  - do you really need to? If so, begin with whatever you do understand and work on from there. Try proving the opposite of what it says...
- ...?

How to present a research paper

Presentation basics

Make sure you know in advance about:

Audience: what can you assume they know?

Time: how long should you talk, and with or without leaving time for questions?

Venue, equipment

Structure of a 20 minute talk

It isn't very long... but you have to get across both the key ideas, and enough of the substance to convince the audience it's worth reading the paper.

"Failsafe" structure:

1. Title slide, with your name.
2. Summary slide, with sections of your talk (4-6)
3. Intro: hook the audience on why they should listen. Ideally, tell them what the contribution of the paper is, informally, and then go on to any technical background you need. Don't spend too long on this!
4. Main part: what does this paper do? Which bits are most important? Which are non-obvious?
   (It's OK to refer to things as being in the paper, and not give details in the talk, *but* someone who hasn't read the paper must get the jist.)
5. Related work slide: with refs
6. Conclusion: recap the contribution
7. Future work

The RTSE presentation, specifically

Aim for 20 minutes, *not* counting time for questions. Go over rather than under: I won't cut you off before 30 mins.

In some cases, the author's own slides for the original conference presentation are on the web. You may choose to use these: but

- it's much easier to talk about slides you wrote;
- writing your own can be a good way to organise your thoughts;
- it's irritating if the speaker has to say "I'm not sure what this bit means"!
- the author's slides were probably aimed at a more specialised audience, so they probably omit background that we'll need.

On the other hand, if the author's slides are good then using them might help you see how a good presentation is put together – so it's up to you.

How to review a research paper

Why is this in the course?

Obvious purpose: because if you stay in academia you will, eventually, have to write reviews. The health of the discipline depends on you learning to do it well, but it's rarely explicitly addressed.

More importantly here, though:

1. because reviewing a paper is one way to ensure you've really read it – trying to write about a paper, in whatever way, often reveals what one hasn't understood;
2. because understanding how the reviewing process works should help you to do better research and research writing. We hope.

What is reviewing?

Two purposes:

1. Quality control: publish or not?
2. Constructive criticism: how to improve?

Aim: be as efficient as possible with the first, to leave most time for the second.

Two broad kinds:

1. for competitive conferences/workshops
2. for journals.

Lots in common, but some differences.

Who reviews?

Mostly established academics – it's an eternal duty...

Typically, a second-year PhD student may co-review with supervisor; a final-year PhD student might write review independently if in own area.

The less experience you have, and the less knowledge of the field, the longer you will need to spend on a review.

Journals and conferences each have levels of responsibility. Let's look at them briefly.

Journal reviewing structure

A typical journal has:

- one or two Editors-in-Chief, who appoint the editorial board, usually allocate papers submitted to editors, and take ultimate responsibility for standards;
- an Editorial Board of established academics, who take responsibility for some papers each year, choose reviewers, collate reviews and make recommendations to EiCs;
- many people who review papers as and when asked.

Conference reviewing structure

A typical conference has:

- one or two Programme Committee Chairs, who appoint the programme committee, allocate papers submitted to PC members, and take ultimate responsibility for standards;
- a Programme Committee of established academics, who review the papers they're allocated, with or without the help of subreviewers, and take part in discussion of which papers should be accepted;
- many people who review papers.

Main difference from journal: because number of slots is limited, papers are in direct competition with one another, so PC discussion very important.

Progress of a conference paper

Conference issues Call for Papers, including scope description and page limit

Authors writes papers

Papers submitted to conference by deadline

Programme committee chair selects (usually 3) reviewers

PC members and/or reviewers read paper and write reviews (independently)

PC members discuss, maybe vote

Acceptance decisions made

Reviews sent (anonymised) to authors

[Occasionally, there is an "author response" phase here]

Successful authors submit revised papers

How a journal paper differs

Papers not usually page limited: include details

Reviews are longer, take longer to do, more thorough.

Generally decide on papers independently, not in competition... question is "does this paper need to be in the literature?" rather than "is this one of the 22 most interesting submissions?"

Where conferences just decide accept/reject, journals decide accept/minor revisions/major revisions/reject: that is, if the paper is considered promising but flawed, authors may get the chance to revise.

Common, but in SE not universal, for conference paper to be followed up by a "journal version".

Should the paper be published?

Key questions:

will anybody read it?

will anybody else be spared the trouble of writing it?

Subquestions

- is the area alive/valuable/interesting/within scope?
- does the paper address a valid problem within the area?
- does it make a significant contribution?
- has it been done before?
- will readers be able to understand the paper as written?

The constructive criticism side

If the paper will be rejected, how can the authors write a better paper next time?

If it will be accepted, how can they help readers to understand it more easily?

Are there related areas/questions that the authors might want to address? Is there related work they don't know about?

Overall: how to write

Be concise, but specific.

If the paper is bad, say *why*, as specifically as possible, and try to phrase positively (The paper would have been better if...). If it's been done before, give the reference. If a statement is false, give a counterexample.

Be polite, but not bland.

Remember the authors are human and getting a bad review is a horrible experience: but also that you're part of the quality control mechanism and the PC chair/journal editor needs to be able to tell what you really think, easily.

Sections of a review

Journals often ask for completely free-text reviews.

Conferences usually use standard templates, fitting with conference management software.

The one I ask you to use for RTSE is on the web page, based on CyberChair.

Classification: overall verdict

```
A: I will champion this paper at the PC meeting
   (advocate/accept).
B: I can accept this paper, but I will not champion it
   (accept, but could reject).
C: This paper should be rejected, though I will not
   fight strongly against it (reject, but could accept).
D: Serious problems. I will argue to reject this paper
   (detractor).
(Fill in after the keyword: A, B, C or D)
CLASSIFICATION:
```

(For RTSE, it would be surprising if you used D, and mildly surprising if you used C, as the papers have been published and I tried to chose good ones!)

Expertise

```
2 - What is your overall expertise concerning the
    subject areas of this paper?
X: I am an expert;
Y: I am knowledgable in the area, though not an expert;
Z: I am not an expert. My evaluation is that of an
   informed outsider.
(Fill in after the keyword: X, Y or Z)
REVIEWER-EXPERTISE:
```

"Expert" means "have published papers in the same area", roughly.

For RTSE: you are probably Z, but say if you think Y!

Paper type

```
3 - Is the paper a research contribution or an experience
paper?
1: Research
2: Experience
3: Both research and experience
(Fill in after the keyword: 1, 2 or 3)
PAPER-TYPE:
```

RTSE: I think all the papers we're considering are Research papers.

Summary

For real: The briefer the better: shorter than the abstract.

Bring out the most important contribution(s).

For RTSE, only: do that, and then, still in this section, write about 400-500 words summarising the work reported in the paper in your own words. (Last year some people had trouble understanding their paper, which then made the rest of their review hard to interpret and mark: asking you to do this should (a) help you to be sure you've understood, and (b) if you haven't, it should help me to see where the problem was!)

Committee comments

Usually left blank.

Can use it for comments that would reveal your identity.

Can use it for comments about how you read, e.g. "I only skimmed Section 3 but can go back to it if it's controversial".

Sometimes used to be blunter than you would to the author, e.g. "nothing exactly wrong with it but who cares?" but probably better practice not to.

Author comments 1

This is the main body of the review.

1. General comments: typically two or three paragraphs (sometimes more), explaining what you think is good and bad about the paper. This part should make it clear to the author *why* you are recommending acceptance/rejection. E.g.

I found your paper clear, practical and ingenious: I could imagine using your work when [...]. Especially, I was impressed by your finding that [...] which is surprising given that [...]. However, your claim that [...] is weak, given that you do not support the claim. How do you know it isn't the case that [...]? It is also crucial that you compare your approach with [...] because [...] ...

This part should make sense to your fellow PC members and to the authors.

Author comments 2

2. Detailed comments: e.g.

p4 first full bullet point: not understandable. What kind of scenario do you have in mind, where this question might be asked?

It's OK if these don't make sense without consulting the paper

Some will be typos, some will be contentful remarks that are specific to a particular place in the paper. Some reviewers separate the two; I generally mix them.

Commenting on minor issues

i.e. typos/English errors/minor mistakes – should you given an exhaustive list, or just examples?

Depends:

- how many there are - if zillions, don't list them all;
- what your recommendation is - if the paper will be accepted, it's more worthwhile to suggest improvements than if it'll be rejected anyway;
- what you think the origin of the errors is. E.g., if I know one of the authors is a native speaker of English, or if they are in an English-speaking country, but the paper's English is poor, I won't bother correcting English, I'll just say "needs proofreading by a native English speaker". If I think the authors may have difficulty getting help with the English, I'm more likely to spend the time trying to help.

Points in favour and against

Pithy. E.g.

For: a paper in a neglected area

Against: far too vague: this paper would not be useful to anyone, academic or practitioner.

Ethics, for reviewers

Maintain confidentiality of the paper.

Be fair.

Don't steal ideas!

Decline to review if you have a conflict of interests, e.g., if an author is a current or recent colleague, supervisor, student, coauthor or close personal friend: if in doubt, consult the PC member/PC chair/journal editor. This can be a difficult area, and different fields/people have different standards.

How thorough?

There is a school of thought that a referee is responsible for finding every error in the paper (Donald Knuth)!

(Corresponding to this view is the suggestion that, even if reviewers remain anonymous during the refereeing process, the names of those who ultimately recommend acceptance should appear with the paper, so that they declare that they share responsibility for the results.)

In SE, this is not how the refereeing task is regarded. Ideally reviews do go into detail and do find all errors, but this is the exception.

Anonymity: of reviewers

Typically in SE reviews are anonymous, i.e., authors are not told who the reviewers are.

Pro: it allows free criticism, especially by junior reviewers of influential authors.

Con: lack of accountability (especially for *accept* decisions!)

However, where the author and reviewer know one another, the author can often guess who the reviewer was.

Don't write anything you couldn't bear to see with your name attached

Anonymity: of authors

A few venues keep authors anonymous: they ask for submissions without author names on, and (to some extent) they ask authors to write the paper in a way which avoids them being identified. E.g., to write "in [paper] X and Y showed" rather than "in our earlier work [paper] we showed".

Pro: avoids "argument by authority"

Con: in practice, if you know the field you can often tell who wrote the paper anyway, so is it worthwhile?

RTSE review evaluation

Each of you will write a review of the paper you present: due in at the beginning of your presentation.

I will be looking for:

1. Evidence that you have read and understood the paper.
2. Evidence that you have understood how the paper fits into its field: for top marks, you should have read around the subject, not just read the paper itself.
3. Sensible identification of strengths, weaknesses, areas that could [have] be[en] improved, suggestions for future work.
4. Comprehensible review with professional tone.