Peer Review

Frédo Durand
MIT CSAIL & EECS

Siggraph review process

- Sort: the area chairs/advisory board/chairs assign papers to 2 committee members
- The committee members assign it to 3 external/ tertiary reviewers (5 reviewers total: 2 primary, 3 external)
- After rebuttal, the primary start discussing papers by email or with the BBS
- Potentially, extra committee members are asked to review a borderline paper
- At the committee, people discuss some more, reach a decision.

http://www.siggraph.org/s2009/submissions/technical_papers/index.php

ETHICAL ISSUES

Conflicts

- Same institution
  - defined broadly: MSR Asia and MSR Redmond are conflicted. Anyone at MIT is conflicted with us, even if we don't know who they are.
- Co-authors of a paper in the last N years
  - where N is usually 3
- PhD advisors/advisee: lifetime conflict
- Common grants
  - Because the renewal of a grant could depend on publications
  - Might be waived, e.g. in Europe where everyone has joint grants
- Involvement in the work
  - A little fuzzy. You will receive credit, you're on a PhD committee
- Misc. : family, people you're too closed to, etc.

http://www.siggraph.org/s2009/submissions/technical_papers/ethics.php

Conflicts - motivations

- Fair decisions
- Perception of fairness
  - Almost as important as fairness itself.
  - A criterion for conflict is not just whether you feel you can provide a fair review, but also whether someone might think that the process is fair or not.

Reviewing competing papers

- Tough situation: often the most relevant reviewer is a researcher who does competing work
- The recent policy has been to avoid that authors of competing papers review each other's submissions
- If you feel a paper is too much in competition with yours, you can decline to review it
- If you have not submitted your work yet, decline
  - You run the risk of being too influenced
- In any case, inform the primary (and possibly write it in the confidential comments)

Reviewing ethic

- Don't share the paper, even once accepted
  - The authors might be filing for a patent
- Do not build on the ideas in the paper until it's published
- Forget you even read it ;-)
- Maybe main reason why grad students should not review papers
  - It's easier to ignore ideas when you only follow the research at a high level

The guessing game

- Do not try to guess who your reviewers are, you will be wrong way too often
- Anecdote 1: Researcher X was mad because he was convinced that researcher Y had killed his paper at the program committee. Turns out Y was defending the paper but could not save it
- Anecdote 2: Author Z was mad at researcher W because he thought that W had killed his paper at the PC meeting. Turns out that W was not on the committee that year.
- Just because a review tells you to cite papers by an author does not mean he wrote the review, especially when the citation is not fully relevant. Author understands their work, and it's often external people who suggest irrelevant citations because they just saw a talk or something.

DISCUSSING UNPUBLISHED WORK

Discussing unpublished work

- Two states:
  - submitted
  - half baked
- Various contexts:
  - A researcher is visiting the lab
  - You are invited to give a talk
- Tension between
  - Get feedback
  - Advertise
  - Respect the reviewing process
  - Protect your ideas

Protect your ideas

- Your ideas can get grabbed by even non-malicious people
  - They might forget you even talked about it
- I know a number of anecdotes where a researcher was inspired by a discussion with someone but forgot about it. It is then an awkward situation when the other person realizes you're publishing a paper based on an idea they gave you. It's even worse if you scoop them.

REVIEWING

The dual role of publications

- Communicate ideas
- Sanction someone's achievements
  - Career advancement, tenure, etc.

- Usually aligned, sometimes a little at odd
  - If two papers come up with the same idea and one is little better, goal 1 only requires the acceptance of the best paper, but goal 2 requires that we accept both
  - goal 1 might suggest that we should reject the paper, wait for a revision to get an even better paper, while goal 2 might suggest that the authors deserve that publication

What to provide in a review

- Thorough evaluation of the pros and cons
  - Importance of the problem (tricky) — The context
  - Novelty of solution
  - Quality of results
  - } The core contribution
  - Thoroughness of result analysis
  - Quality of writing
  - } The execution
- Suggestions for improvement
  - Writing
  - Analysis of results and limitations
  - Previous work
- Don't agonize over numerical scores
- The PC members will make the decision. Provide them elements to inform that decision

Two main axes for contributions

- See also the Pasteur quadrants

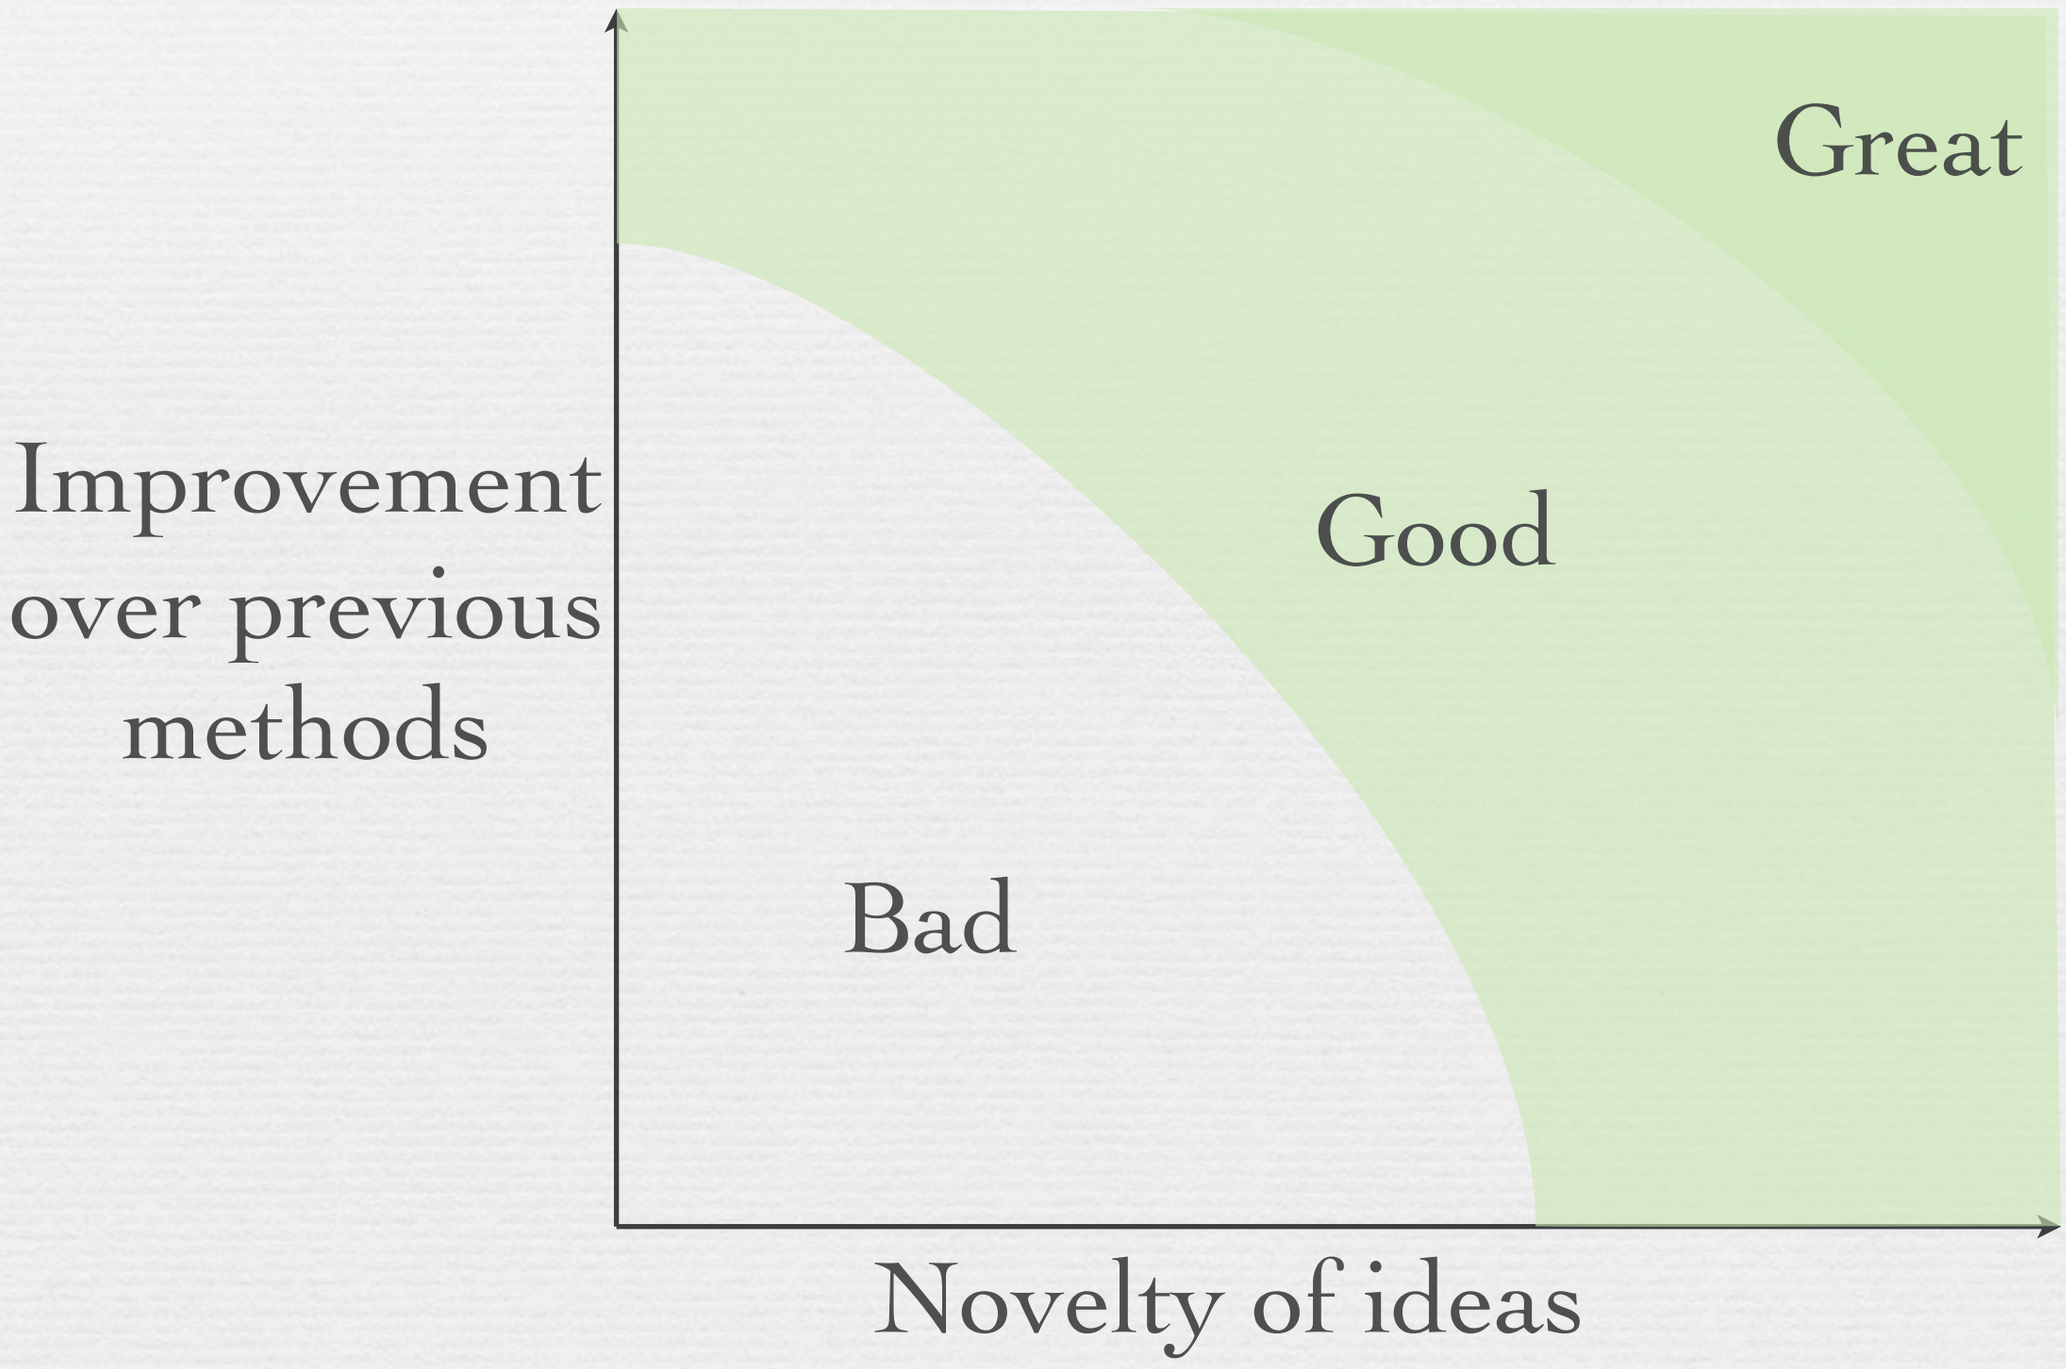

Pitfalls

- We are sometimes bad at accepting truly innovative work
- The fact that a paper opens more questions than it answers is a good thing
- Do not focus on what the paper could have been. Evaluate the submitted work. Would the field be better if this was published as is?

Timeliness and long term

- Beware on incremental improvement
- Beware of fads
- Be open: maybe the current context is not right for this work, but it will be relevant five years from now
  - e.g. don't say: this work is crap because current generations of GPUs can't do X

Be specific, be positive

- When complaining about an issue, suggest improvements
- Avoid vague “the paper is poorly written”. Point out possible improvements.

Be positive

- Extract useful ideas
- Again, don't focus on what's missing in the paper but on what the paper contributes

Bad papers

- Be polite, be professional
- If the paper is impossible to read, that can be a reason for rejection
- Find a balance
  - Provide feedback and suggestions for improvement
  - It is not your job to do the author's work

Tricky case: flawed submission

- Sometimes actual dishonesty, sometimes a bug
- Always err on the side of a bug:
  - if you're wrong, it's always better to be on the generous side
  - if the authors are being dishonest, they'll get the message
  - it's difficult to prove dishonesty
- Be courteous, don't be inflammatory, it will only strengthen your feedback

Great papers

- Write a long review even when there is no criticism
  - Risk: a short review might be downweighted compared to a long rant.

Prior work

- Separate different cases:
  - There is prior work that already does this
  - There is work that is kind of related and would deserve to be cited
- Usual criterion: published article that can be purchased after an event is over (conference proceedings)
  - But the web is making this complicated
- Sketches, posters, short papers: not publications
- Patents: usually considered publications, but inconsistent. (but does not preclude the author from publishing it as an article)
- Theses: previous publication, but does not preclude the author from publishing it as an article

http://www.siggraph.org/s2008/submissions/juried/papers/policy.php

Prior work & contribution scope

- If the paper improves on specific papers, be clear about the differences and their magnitude.
  - Siggraph suggests to put this in the first question, but it's equally relevant in the evaluation question

Self contained & reproducible

- Depends on paper (impossible for systems papers)
- Two different issues
- Self contained: will the readers need to read tons of references to understand the paper?
  - Judgement call, depends what is considered well known in a field, whether standard implementations of a tool exist, etc.
- Reproducible
  - Do the authors provide the necessary information to reproduce their work?
  - Often a problem with industry papers

Readability

- Unreadable paper
  - Be specific about what you don't understand (at least a few examples)
  - Suggest improvements (better notations, define vocabulary, high-level overview, system diagram, etc.)
  - If you don't understand, then it is unclear
    - after all, you have been selected for your expertise in that area
- Tone, excessive claims
  - Don't hesitate to tell authors to reduce their claims

Scores

- 5 you're pissed if the paper does not get in
- 4 is already well in favor of acceptance
- 3 is non-committal "useless" review
- 2 you want to reject
- 1 I usually advise against using 1 unless the authors are misled or you think there is something dishonest. A 2 already sends a clear message to the primaries and is not as depressing for the authors.
- Accepted papers typically have an average above 3.3

Private comments

- Avoid using that field: the authors deserve to know your evaluation. Do not use this field as a summary of your evaluation
- Do not discuss other papers you are reviewing
  - You could reveal yourself to the authors

How long should you spend?

- Between two hours and a day
- I usually like to do it in two steps to digest the paper
- Recursive reviewing
  - Sometimes you will need to check a few other papers for reference

Why should you review (well)?

- Community service
- The senior person who asked you to review will think more highly of you if you do a good job
- Springboard to program committees
- Teaches you how to write

- As a graduate student, you should not decline a review
- Later, take numbers into account

Don't get a bad reputation

- Be polite
- Be positive
- Be punctual

- I know a number of people who were not invited on the Siggraph Program committee because of one of these issues
- The weight of your review will be reduced if you're known as a paper killer

Refs

- Greg Turk's advice on reviewing
  - http://www.siggraph.org/s2008/submissions/juried/papers/review_writing.php
- A Guide for New Referees in Theoretical Computer Science
  - http://www.eng.unt.edu/ian/pubs/referee.pdf
- Rules for Referees by Bernard K. Forscher
  - http://www.sciencemag.org/cgi/content/citation/150/3694/319
- The task of the referee by Alan Jay Smith
  - http://www.cs.uiuc.edu/homes/caesar/classes/CS598.F08/readings/reviewing.html
- David Gifford. How to referee a research paper