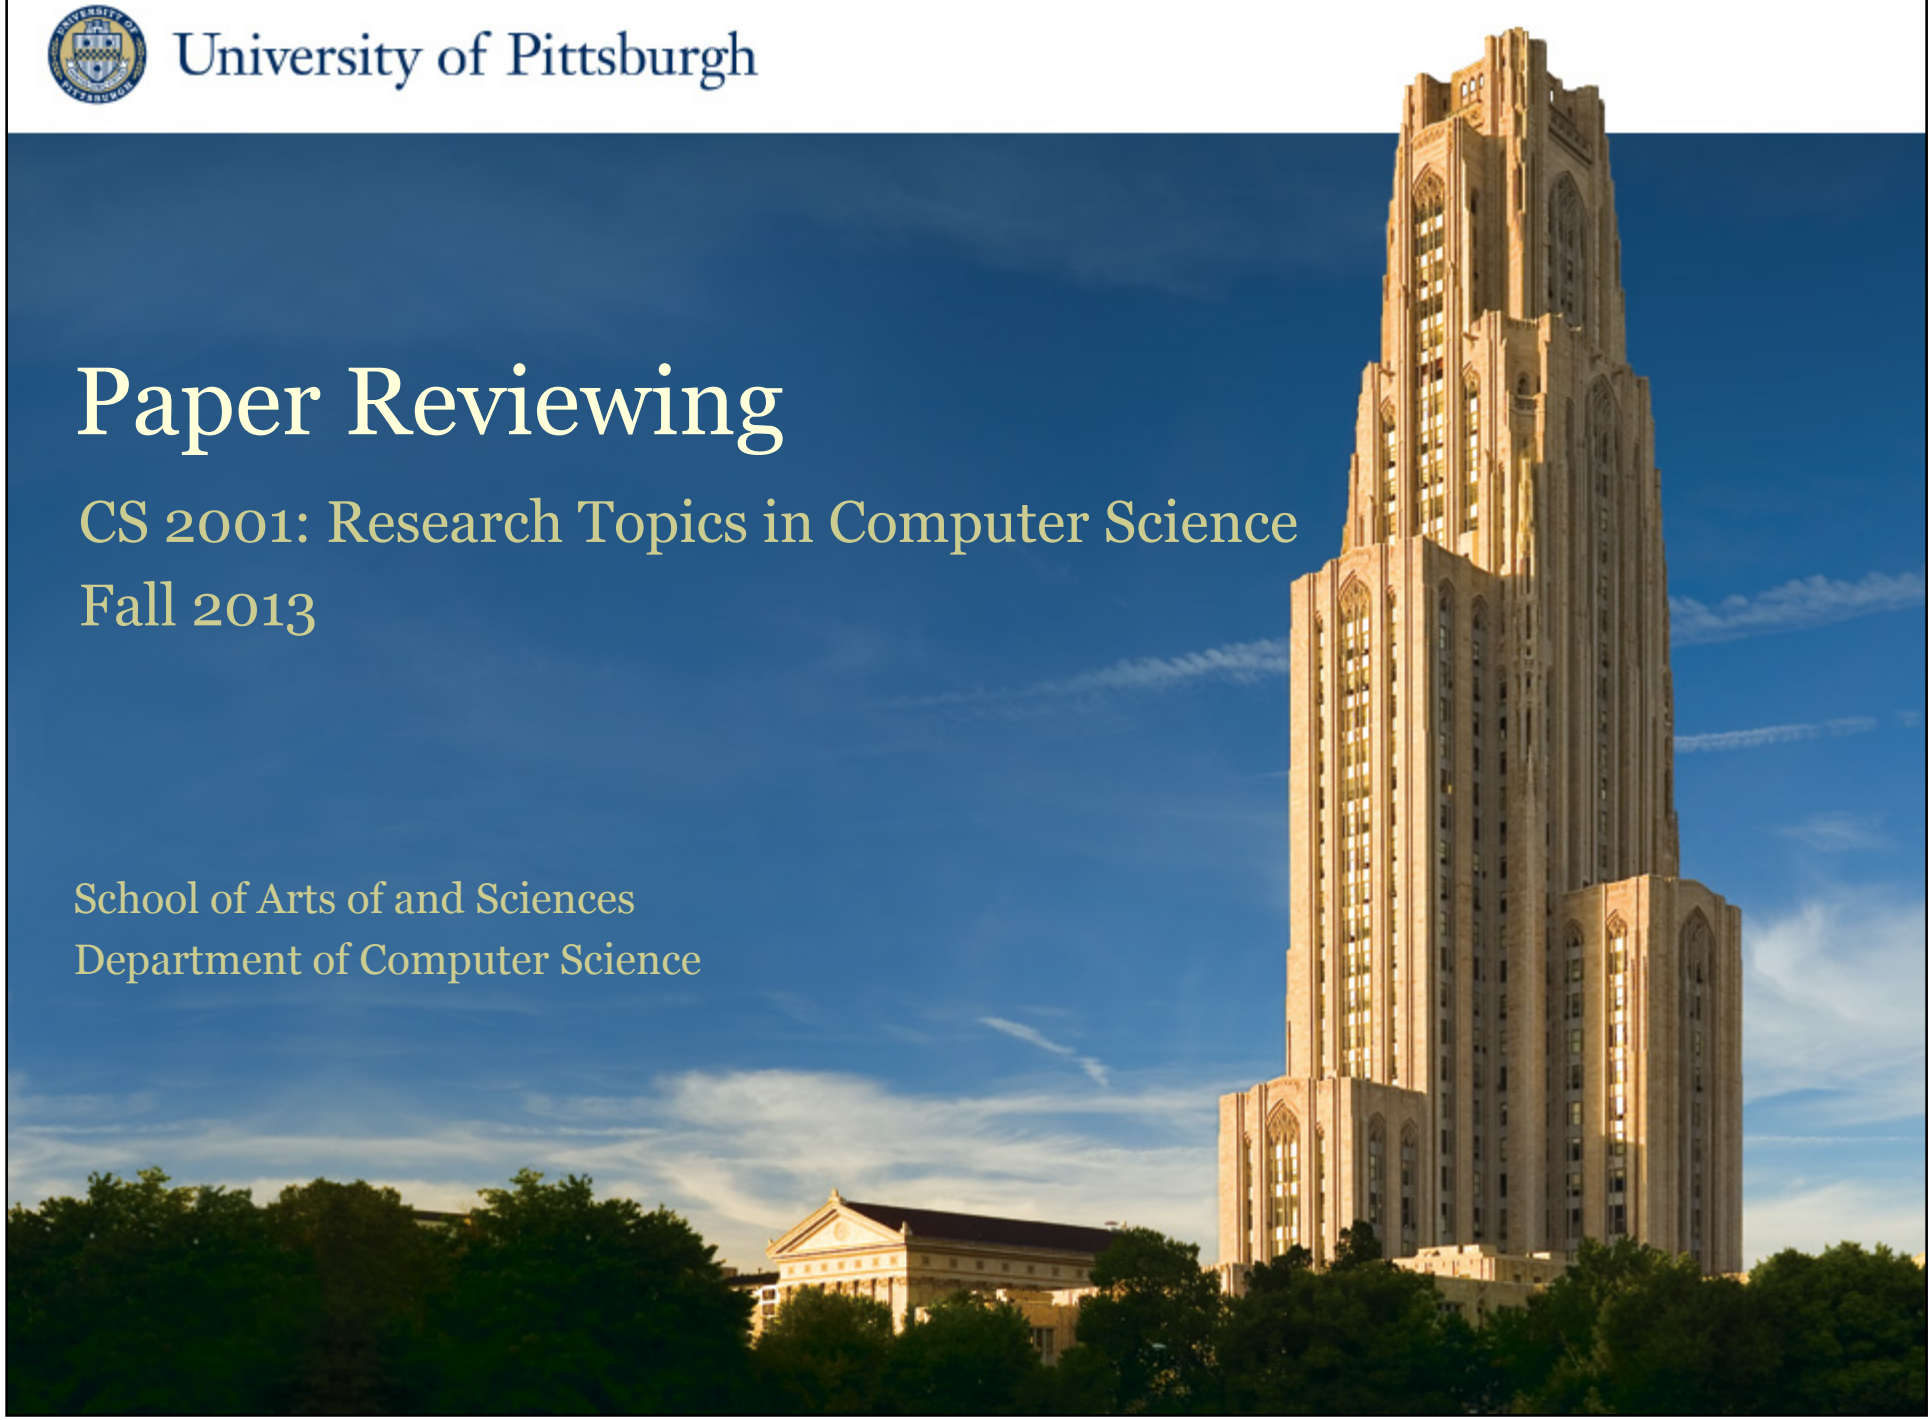
University of Pittsburgh
Paper Reviewing
CS 2001: Research Topics in Computer Science
Fall 2013
School of Arts of and Sciences
Department of Computer Science

Topics for Today

Why review papers?

Structuring a good technical paper review

Common pitfalls and how to avoid them

Examples & Exercises

There are many reasons to write paper reviews!

Reason 1: Testing your own comprehension

- Noting contributions, significance, strengths, and weaknesses
- Identifying promising areas for future work

Reason 2: Group meetings / reading groups

- Similar to above, but to promote discussion within group

Reason 3: Related work in your papers

- Can be thought of as *very* concise paper reviews
- Summarize main technical points, compare/contrast with your work

Reason 4: Conference and journal reviews

- Peer review is used to judge the merit of scientific papers
- Reviews influence accept/reject decision and author revisions

How are conference programs decided?

The program committee chair appoints a program committee

- A collection of experts in the field
- Typically consists of 15-50 people, depending on conference size

Once papers are submitted, the PC members bid on papers within their specific areas of interest and expertise

Each paper is typically assigned to at least 3 reviewers who

- Read the paper carefully
- Draft a review of the paper
- Discuss the paper with other PC members

The final program is decided upon after discussion at the PC meeting

What is the purpose of a conference review?

A conference paper review serves many purposes:

- Synthesizes the reviewer's understanding of the paper
- Communicates the reviewer's thoughts about the paper to other PC members and the PC chair
- Partially documents the PC's decision to accept/reject the paper
- Provides guidance to the authors regarding possible (or mandatory!) improvements to their work

As a result, the review is important at all stages of the process

Bottom line: A paper review should *not* be a book report!

What is the structure of a good paper review?

Technical summary

Description of contributions

Major critiques

- Strengths
- Weaknesses
- Questions

Minor points

Concluding remarks

Content:

- Very short (1-2 paragraphs)
- Overview of the paper

Purpose:

- As the reviewer, this provides you with context for the review
- Allows the PC chair to get a quick synopsis of the paper
- Convinces the author that you, as the reviewer, actually read and understand the paper

What is the structure of a good paper review?

Technical summary

Description of contributions

Major critiques

- Strengths
- Weaknesses
- Questions

Minor points

Concluding remarks

Content:

- Very short (1-2 paragraphs)
- Quick summary of the novel aspects of the paper

Purpose:

- Novelty is paramount! This provides evidence for the final accept/reject decision
- Again, convinces the author that you understand the novelty of their contribution
- Sets the stage for detailed critiques

What is the structure of a good paper review?

Technical summary

Description of contributions

Major critiques

- Strengths
- Weaknesses
- Questions

Minor points

Concluding remarks

Content:

- Technical and/or methodological strengths and weaknesses
- *Examples*:
  - How interesting is the problem?
  - Novel proof techniques or solutions
  - Missing related work
  - Assessment of the (in)completeness of the evaluation
  - …

Purpose:

- Primary assessment of the paper
  - Do the authors bring something really cool to the table?
  - Is the paper somewhat incremental, but well executed?
  - Does the paper have fatal flaws?
- Typically, this provides fodder for discussion at the PC meeting

What is the structure of a good paper review?

Technical summary

Description of contributions

Major critiques

- Strengths
- Weaknesses
- Questions

Minor points

Concluding remarks

Content:

- Remarks on any thing that was unclear in the paper

Purpose:

- Stimulate discussion with other reviewers
- Inform the author of questions still remaining after reading

What is the structure of a good paper review?

Technical summary

Description of contributions

Major critiques

- Strengths
- Weaknesses
- Questions

Minor points

Concluding remarks

Content:

- Aspects of the paper that don't influence the novelty of the contribution, but do impact the quality of the paper overall
- *Examples:*
  - Typos and grammar errors
  - Suggestions for better examples
  - Corrections to minor logical flaws
  - …

Purpose:

- Helpful for planning revisions

What is the structure of a good paper review?

Technical summary

Description of contributions

Major critiques

- Strengths
- Weaknesses
- Questions

Minor points

Concluding remarks

Content:

- Very short (1-2 paragraphs)
- Final assessment of paper, with justification

Purpose:

- Communicate your views on the paper to others
  - ➢ You might love the paper, yet make many negative critiques
  - ➢ You might hate the paper, yet say some positive things about it
  - ➢ This is where you clarify
- Provide final suggestions

Tips on being a good reviewer

Acceptance rates at good conferences are very low

- < 15% is not unusual at competitive venues
- < 10% not unheard of!

This leads to the following situation

- Problem: Writing a good review takes time
- Problem: PC members often must review many papers
- Problem: Very few papers can be taken anyway
- "Solution": Look for reasons to reject a paper and be done with it

This is extremely counterproductive, and not good for science (Why?)

Hill and McKinley* offer suggestions on avoiding this type of pitfalls

* http://www.cs.utexas.edu/users/mckinley/notes/reviewing.html

Avoiding Pitfalls

	Pitfalls	Recommendation
1	Seek to find all flaws in the paper, in part to show your expertise as a reviewer	Look for reasons to accept a paper. Despite its flaws, does it point in new directions or expose promising insights? The community can benefit from imperfect, insightful papers.
2	Since the review process is anonymous, it is appropriate to criticize the paper as if the authors did not have feelings.	Your tone should be the same as if you are giving comments to a colleague face-to-face. It is always possible to be constructive, focus on the work, and do not attack the researchers behind it. The purpose of a review is not only for selecting papers, but to improve the quality of all the work in our area.
3	Reject papers that build on recently-published new directions, but accept those that build on the established norm.	While truly new papers are best (and rare), consider accepting papers that follow-up on recently-published promising directions. These papers allow the community to explore ideas that can not be fully-developed in one paper.
4	Advocate rejecting a paper with little comment, because it is obvious that all with agree with you. Ditto for accept.	Explain why you advocate a rejection or acceptance, because people will often disagree with you. Your explanations will make you a more effective advocate or detractor for the paper.
5	Advocate rejecting (almost) all papers to show about tough you are.	Your job is decide what is best which is not usually accomplished by rejecting every submission.
6	Advocate rejecting a paper because you seem to remember it being the same as (or similar to) unidentified prior work.	In this situation, the professional should reference important prior work after refreshing one's memory regarding what it contains. One missing reference is usually not a reason to reject a paper.

John Kubiatowicz, David Bindel, Yan Chen, Steven E. Czerwinski, Patrick R. Eaton, Dennis Geels, Ramakrishna Gummadi, Sean C. Rhea, Hakim Weatherspoon, Westley Weimer, Chris Wells, Ben Y. Zhao: OceanStore: An Architecture for Global-Scale Persistent Storage. ASPLOS 2000: 190-201

Technical summary

Description of contributions

Major critiques

- Strengths
- Weaknesses
- Questions

~~Minor points~~

Concluding remarks

Concluding Thoughts

Paper reviews serve many purposes

- Developing your own comprehension
- Preparing for group meetings
- Critiquing the work of students in your lab
- Evaluating conference or journal submissions

Writing a good review is not hard, it just takes time and practice

- Right now, time is a resource that you do have
- You will get practice in this class, and as you advance

Try to avoid common pitfalls and focus on

- Recognizing the strengths of a paper
- Preparing useful feedback for the authors

In short: Write the paper review that you want to receive!