

COSC-450: ANALYTICAL PERFORMANCE MODELING

GUIDE TO READING PAPERS

1 Reading the Paper

Allow enough time. Reading a research paper can take a very long time, even if you're an expert in the field. You'll probably make several passes through the paper, reading at different levels of detail and looking for different information each time.

Read the abstract. The abstract of a paper is a short summary at the very beginning of the paper. It is meant to give you a sense of what the paper is about: the problem the paper addresses, the general approach to solving the problem, and the important conclusions. This can give you a good idea about what to look for in the rest of the paper, but often the abstract is quite dense so it should not be the only thing you look at before diving into the text of the paper.

Make an initial scan through the paper. Before reading the paper from start to finish, page through it quickly to get a sense of the big picture. Note the section headers: often this will give you some idea of what approach the authors took (if you see a section header called "analysis," that suggests you're looking at a theory paper; a section header "experimental results" suggests that the authors simulated or implemented a system). Look at any graphs to see what kind of results the authors obtained. You may want to read the introduction and the conclusion at this point, but you'll likely want to skip the details of any mathematical analysis or experiments.

Read the paper in more depth. After you have a clear understanding of the main points of the paper, read through each section carefully (and possibly several times). In this reading, think about the following questions:

1. What is the research problem the paper attempts to address? What is the motivation of the work? Is there a real-world problem to which the paper proposes a solution? Does the work attempt to overcome the weaknesses of existing approaches? Is an existing research paradigm challenged? In short, what is the specific goal of the paper?
2. What contributions does the paper claim to make? What is new versus work that was published earlier? Is this a brand new research question? A new algorithm or methodology for solving a problem? A new proof technique? Something else original? Often the authors will explicitly state "the contributions of this paper are..." or something similar, so this can be easy to find!
3. How do the authors substantiate their claims? That is, how do they justify that they have successfully achieved their stated contributions? In a biology paper or a chemistry paper, there often is a section titled "Methods" that will lay this out for you. Computer science papers often do not have a "methods" section, but there's always some sort of approach used to justify the claims. Are there major theorems that the authors prove? Simulations? Other types of experiments?

4. What are the conclusions? What have we learned from the paper, and what are the big ideas that we should take away from it?

It's important, when reading a paper, not to accept as fact everything that you read. Papers that are submitted for publication go through a rigorous review process during which researchers who are experts in the field read the submitted paper, evaluate the quality of the work, and make suggestions for improvement. It's important to think critically about the content of a paper any time you're reading a paper, not only when reading it as a peer reviewer. Here are some more questions to think about:

1. Is the research problem significant? Is it solving a big problem, or simply addressing some low-hanging fruit?
2. Are the contributions truly new and important? Are there surprises, or are the results fairly self-evident? Are the authors aware of how their work relates to existing literature?
3. Are the claims valid? Have the claims been sufficiently supported by theorems, experimental results, etc? Are there any problems with the proofs, experimental setup, benchmarks, methodology? Are the generalizations valid?

Finally, one of our goals in reading research papers is to learn about the state of the art and then to think creatively about possibly new research ideas. Here are some questions that might help you come up with further research problems to explore:

1. Are there other approaches to address this research problem?
2. What is a better way to substantiate the authors' claims?
3. Is there a compelling argument against the authors' claims?
4. Can the results, theorems, or methodological approach be applied to a different context?
5. What are the open problems raised by this work?

2 Summarizing the Paper

In advance of the class meeting in which we'll discuss a paper, I ask that you submit a summary and a list of questions that you have about the paper (you can submit this through Moodle). Here's some advice for writing your summary.

Avoid plagiarism. Plagiarism is the act of taking ideas or text produced by someone else and presenting it as your own. When you're summarizing someone else's work, it can be easy to inadvertently plagiarize because you might feel that the authors' own words are the clearest way to express their point. To avoid this, take notes in your own words as you're reading. Using short notes or summarizing key points in your own words forces you to rewrite the ideas into your own words later. If you find yourself sticking closely to the original language and making only minor changes to the wording, then you probably don't understand the study well enough. Keep reading!

Draft your summary. Your summary should follow the same order as the paper itself. Begin by explaining the research question and why it is important. You'll then describe the methods (theoretical analysis, simulation, experiment, etc.), and explain the major results. What were the important theorems? Were the experimental results significant? What were the most important findings in the paper, and what are the implications? Keep in mind that the results and interpretation should directly relate to the research question. Focus on the content of your summary, not on the length; you can always add more detail or condense as needed later on.

Revise! First, make sure that your summary includes enough relevant information for your reader to understand what the paper was about. On the other hand, this is a summary, so you don't want to get too bogged down in technical details. An important question to keep in mind is: who is the audience for your summary? In this case you can imagine your audience as being your peers: other upper-level computer science students who might want to learn a little bit about the research paper that you're reading. Your audience is intelligent, but doesn't want to have to work hard to understand the ideas you're explaining. After you've revised your summary for content, take another pass through it to clean up the language. Don't be overly wordy, be precise in both language and scientific meaning, and paraphrase rather than quoting directly from the paper. Any time you're submitting written work, be sure to read through it a final time to make sure you haven't missed anything before you submit it!

List your questions. At the end of your summary, list any questions that the paper has raised for you. Your questions could be about aspects of the paper that you didn't fully understand, or they could be about ideas for extending or generalizing the study. We will use these questions to kick off our in-class discussion of the paper.

Parts of this guide were adapted from Prof. Trapani's "How to Summarize a Research Article," itself adapted from: http://web2.uconn.edu/ahking/How_to_Summarize_a_Research_Article.pdf, which in turn was adapted from "Summarizing a Research Article" 1997-2006, University of Washington.

Other parts of this guide were adapted from: <http://www2.cs.uregina.ca/pwlfong/CS499/reading-paper.pdf>.