
Reading a Computer Science Research Paper

Philip W. L. Fong

Department of Computer Science
University of Calgary
Calgary, Alberta, Canada T2N 1N4
pwlfg@ucalgary.ca

Abstract: This tutorial article highlights some points that a graduate or senior undergraduate student should bear in mind when reading a computer science research paper. Specifically, the reading process is divided into three tasks: comprehension, evaluation and synthesis. The genre of paper review is then introduced as a vehicle for critical reading of research papers. Lastly, guidelines on how to be initiated into the trade of conference and/or journal paper review are given. Designed to be used in a graduate course setting, this tutorial comes with a suggested marking scheme for grading paper reviews with a summary-critique-synthesis structure.

Categories and Subject Descriptors: K.3.2 [Computing Milieux]: Computer and Information Science Education --- Computer science education.

General Terms: Documentation

Keywords: Graduate Education, Reading Research Papers, Paper Review

1. COMPREHENSION

The first lesson to reading research papers is learning to understand what a paper says. A common pitfall for a beginner is to focus solely on the technicalities. Yes, technical contents are very important, but they are in no way the only focus of a careful reading. In general, you should ask yourself the following four questions when you are reading a research paper.

1. *What is the research problem the paper attempts to address?* What is the *motivation* of the research work? Is there a *crisis* in the research field that the paper attempts to resolve? Is the research work attempting to overcome the *weaknesses* of existing approaches? Is an existing *research paradigm* challenged? In short, what is the *niche* of the paper?
2. *What are the claimed contributions of the paper?* What is *new* in this paper? A new *question* is asked? A new *understanding* of the research problem? A new *methodology* for solving problems? A new *algorithm*? A new breed of software *tools* or *systems*? A new *experimental method*? A new *proof technique*? A new *formalism* or *notation*? A new *evidence* to substantiate or disprove a previously published claim? A new *research area*? In short, what is *innovative* about this paper?
3. *How do the authors substantiate their claims?* What is the *methodology* adopted to substantiate the claims? What is the *argument* of the paper? What are the major *theorems*? What *experiments* are conducted? *Data analyses*? *Simulations*? *Benchmarks*? *User*

studies? *Case studies*? *Examples*? In short, what makes the claims *scientific* (as opposed to being mere opinions¹)?

4. *What are the conclusions?* What have we *learned* from the paper? Shall the *standard practice* of the field be changed as a result of the new findings? Is the result *generalizable*? Can the result be applied to *other areas* of the field? What are the *open problems*? In short, what are the *lessons* one can learn from the paper?

Every well-written research paper contains an *abstract*, which is a summary of the paper. The role of an abstract is to outline the answers to the above questions. Look therefore, first to the abstract for answers. The paper should be an elaboration of the abstract.

Another way of looking at paper reading is that every good paper tells a *story*. Consequently, when you read a paper, ask yourself, "What is the plot?" The four questions listed above make up an archetypical plot structure for every research paper.

2. EVALUATION

An integral component of scholarship is to be critical of scientific claims. Ambitious claims are usually easy to make but difficult to substantiate. Solid scholarship involves careful validation of scientific claims. Reading research papers is therefore an exercise of critical thinking.

¹ Alternatively, what makes it a research paper rather than *science fiction*?

- *Is the research problem significant?* Is the work scratching *minor itches*? Are the authors solving *artificial problems* (aka *straw man*)? Does the work enable *practical applications*, deepen *understanding*, or explore *new design space*?
- *Are the contributions significant?* Is the paper *worth reading*? Are the authors simply *repeating* the state of the art? Are there real *surprises*? Are the authors aware of the relation of their work to *existing literature*¹? Is the paper addressing a well-known *open problem*?
- *Are the claims valid?* Have the authors been *cutting corners* (intentionally or unintentionally)? Has the right theorem been proven? Errors in proofs? Problematic experimental setup? Confounding factors? Unrealistic, artificial benchmarks? Comparing apples and oranges? Methodological misunderstanding? Do the numbers add up? Are generalizations valid? Are the claims modest enough²?

When you evaluate a research work, two caveats are worth noting:

- Consistently evaluating research works in a negative way gives a young researcher a false sense of being critical. Learn to be fair: attend to both the strengths and weaknesses of the work. If you are reading a classical paper that has been published for a while, make sure you are reading the paper in the right historical context: What seems to be obvious now might have been ground-breaking then.
- A young researcher may want to focus on point 3 (*Are the claims valid?*). Evaluating the significance of the research problem and the contributions of the paper usually requires a comprehensive understanding of the research field as a whole. Yet, do not let the above comment hinder you from disagreeing with the paper authors in matters of significance.

3. SYNTHESIS

Creativity does not arise from the void. Interacting with the scholarly community through reading research papers is one of the most effective ways for generating novel research agendas. When you read a research paper, you should see it as an opportunity for you to come up with new research projects. The following is a list of questions you can ask to help in this direction. (Of course, this list is not supposed to be exhaustive.)

- What is the crux of the research problem?

¹ Be very sceptical of work that is so “*novel*” that it bears no relation to any existing work, builds upon no existing paradigm, and yet addresses a research problem so significant that it promises to transform the world. Such are the signs that the author might not be aware of existing literature on the topic. In such a case, the authors could very well be repeating works that have already been done decades ago.

² It is very tempting for an inexperienced researcher to make overly general conclusion from limited evidence. A high quality scientific claim is always *modest* --- claiming only what can be concluded from the evidence, making explicit the limitation of the evidence, and carefully delimiting the scope of the claim.

- What are some alternative way to substantiate the claim of the authors?
- Is there an alternative way to substantiate the claim of the authors?
- What is a good argument against the case made by the authors?
- Can the research results be strengthened?
- Can the research results be applied to another context?
- What are the open problems raised by this work?

Bottom line: If you were to do the research, how would you do differently?

4. PAPER REVIEW

A paper review is a short essay (5 pages, single-space, 1-inch margins, 12 point font) reporting what you have learned from reading a research paper. Writing reviews for the papers you have read is a great way to sharpen your paper reading skills. Such a review is typically structured in three sections — *summary*, *critique* and *synthesis*.

1. *Summary.* Give a brief summary of the work *in your own words*. This section demonstrates your understanding of the paper, and as such it should answer the four questions outlined in Section 1. The summary section should be structured as follows: (1) motivation, (2) contribution, (3) methodology and/or argument, and (4) conclusion. It is imperative that you use your own words to summarize the paper. Failing to adhere to this guideline not only constitutes plagiarism, but also demonstrates that you probably do not quite understand the work. You can be sure that you understand something only when you are capable of explaining it in your own words.
2. *Critique.* Pick two to three points³ you want to argue with the authors⁴. Use the questions outlined in Section 2 to help you come up with meaningful critiques. Do not repeat the *Limitations* section of the paper. Doing so means that you *agree* with the authors! Pick points of *disagreement*, and launch an intellectual debate with the authors. Carefully articulate and substantiate your case. Do not just say, “I don’t like this point.” Instead, give technical reasons to substantiate your critiques. Be specific in your choice of words. Avoid generic adjectives such as “bad”, “poor”, “lame”, “stupid”, etc, and their synonyms and antonyms. You can go a lot further by replacing such vague words with more specific

³ This restriction on the number of critiques is intentional. Firstly, the restriction forces you to pinpoint the major weaknesses of the paper, rather than to spend efforts debating issues of peripheral importance. Secondly, such a restriction allows you to enjoy the mental room necessary for developing a substantial case against the authors.

⁴ Notice that the Critique section presents only negative evaluations of the paper. Have we forgotten about being fair to a research work? No, positive evaluations are omitted for a good reason. Experience tells us that students tend to give positive evaluation in the following form: “I agree with the authors. They did this and that, and they did a good job.” The end result is usually a repetition of the authors’ claims. I find that focusing on critiques offers a more substantial learning experience to the students, forcing them to think rather than to parrot.

ones: “inelegant”, “inefficient”, “memory-intensive”, “ill-defined”, etc.

3. **Synthesis.** Propose one to two ways in which the research work can be further developed. Do not repeat the *Future Work* section of the paper. Be original. Consult the list of questions in Section 3 if you run out of ideas.

I use this format when I ask students in my graduate classes to review a paper. Consult the Appendix for a sample outline, a marking scheme, and page length suggestions.

5. AN ALTERNATIVE PAPER REVIEW FORMAT

The format of paper review outlined in the previous section is the one I adopt for my graduate classes. It works for me, but it is definitely not the only way to structure a paper review. I outline here an alternative format I learned from a friend of mine.

1. *What is the purpose of the work?*
2. *How do the authors achieve this purpose? Why is this particular approach adopted?*
3. *Do you think the purpose has been achieved?*
4. *What insights have you gained from reading this work?*

Notice the parallel between this alternative structure and the summary-critique-synthesis structure in the previous section.

6. READING RESEARCH PAPERS LIKE A PRO

When a research paper is submitted to a conference or a journal, it will undergo a *peer review* process, in which the paper is subject to the intense scrutiny of peer researchers. The *referees* who review the submitted paper will read the paper in more or less the same way as we outlined in Sections 1 and 2, and then they will write up a *referee report* in a style similar to the paper review discussed in Section 4, except for the synthesis section. Based on the referee reports, the program chair of a conference or the editor of a journal will then make the decision of whether to accept the paper. It is therefore instructional to understand how a referee goes about reviewing a paper, and learn to read research papers like a professional. A very good introduction to the subject can be found in an article by Smith [1]. The paper is slanted towards experimental computer science. For a perspective focusing on theoretical computer science, consult the article by Parberry [2]. After reading these papers, I highly recommend graduate students to find opportunities to practice professional paper reviewing. Your thesis supervisor will likely be involved in the program committees of conferences, or asked to review papers for conferences. Approach your supervisors, and volunteer to help out with paper reviews. By actually writing up a professional review report, and discussing your review with your supervisor, you will gain tremendous insight into the paper publishing process, as well as the implicit value system of the academic world. This kind of training is hard to acquire through other means.

BIBLIOGRAPHY

- [1] Alan Jay Smith, The task of the referee, *IEEE Computer*, 23(4):65-71, April 1990.
 [2] Ian Parberry, A guide for new referees in theoretical computer science, *Information and Computation*, 112(1):96-116, July 1995.

APPENDIX: MARKING SCHEME AND OUTLINE OF A PAPER REVIEW

1. SUMMARY (40%: 2.5 pages)
 - a. Motivation (8%)
 - b. Contribution (8%)
 - c. Methodology (16%)
 - d. Conclusion (8%)
2. CRITIQUE (30%: 1.5 pages)
 - a. 1st Critique (15%)
 - b. 2nd Critique (15%)
 - c. Optional: 3rd Idea

If a 3rd critique is given, then each critique is worth 10%. Students are encouraged to focus their efforts in two rather than three critiques.

3. SYNTHESIS (30%: 1 page)
 - a. 1st Idea (30%)
 - b. Optional: 2nd Idea

If a 2nd idea is presented, then each idea is worth 15%. Students are encouraged to focus their efforts on one rather than two ideas.