

How to Read a CS Research Paper?

Philip W. L. Fong

July 15, 2004

This article highlights some points a young researcher should bear in mind when reading a CS research paper.

1 Comprehension

The first lesson to reading research paper 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 *original* 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

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

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. Fancy claims are usually easy to make but difficult to substantiate. Solid scholarship involves careful validation of scientific claims. Reading research paper is therefore an exercise of critical thinking.

1. **Is the research problem significant?** Is the work scratching *minor itches*? Are the authors solving *artificial problems* (aka *strawman*)? Does the work enable *practical applications*, deepen *understanding*, or explore *new design space*?
2. **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*²? Is the paper addressing a well-known *open problem*?
3. **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 the generalizations valid? Are the claims modest enough?

3 Synthesis

Creativity does not arise from the void. Interacting with the scholarly community through reading research papers is one of the most effective way 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.)

²Be 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 simply repeating works that have already been done decades ago.

- What is the crux of the research problem?
- What are some alternative approaches to address the research problem?
- What is a better way to substantiate the claim of the authors?
- What is a good argument against the case made by the authors?
- How can the research results be improved?
- Can the research results be applied to another context?
- What are the open problems raised by this work?
- Bottomline: Can we do better than the authors?

4 Paper Review

A paper review is a short essay (3–4 pages) 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*, *evaluation*, 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. It is imperative that you use your own words to summarize the paper. Another way to think of it is that you are writing an alternative, elaborate *abstract* for the paper.
2. **Evaluation.** Evaluate the work by answering the questions outlined in Section 2. Learn to be fair: point out 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.
3. **Synthesis.** Generate any interesting thoughts you have on the work by consulting the list of questions in Section 3.

5 Related Work

The classic by Adler and van Doren [1] provides lots of wisdom on how to read a book. The guide by Murphy and Griswold also provides a helpful introduction to reading an engineering research paper [3].

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. 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 go 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 [5]. The paper is slanted towards experimental computer science. For a perspective focusing on theoretical computer science, consult the article by Parberry [4]. See also [2].

References

- [1] Mortimer J. Adler and Charles van Doren. *How to Read a Book*. Simon and Schuster Trade, 1980.
- [2] Allen S. Lee. Reviewing a manuscript for publication. <http://www.people.vcu.edu/~aslee/referee.html>.
- [3] Gail Murphy and Bill Griswold. How to read an engineering research paper. <http://www.cs.ubc.ca/~murphy/cpsc507/winter02/documents/reading-eval.htm>.
- [4] Ian Parberry. A guide for new referees in theoretical computer science. *Information and Computation*, 112(1):96–116, 1994.
- [5] Alan Jay Smith. The task of the referee. *IEEE Computer*, 23(4):65–71, April 1990.