Review of Related Literature

Sam Silvestro, University of Texas at San Antonio

Provides a brief introduction to basic best practices for reading technical papers in the field of computer science. These methods are primarily concerned with critically examining the contents and conclusions of a paper, as well as the significance of the problem and the contribution.

1. INTRODUCTION

Researchers spend many long hours reading papers. This is essential to both exploring the current state of a particular field, as well as finding related work concerning problems they wish to address themselves. As such, knowing how to efficiently and critically examine a research paper is of high importance, both with respect to learning how to dismiss papers that are not worthy of repeated reference, as well as to uncovering flaws and weaknesses in their assumptions, methods, or conclusions.

[Fong 2009] recommends to readers an approach based on dividing the reading process into three parts: comprehension, evaluation, and synthesis. During the comprehension phase, the researcher learns to read what the paper says, without focusing exclusively on the technical details. The reader should ask themselves (and a high-quality paper should readily communicate to them), "What is the research problem the paper attempts to address?"

1.1. What makes a good paper

As eluded to, we should first define what constitutes a high-quality paper. A high-quality paper should clearly state the research problem it sets out to address in both the abstract and the introduction. As a general rule, a paper should ideally be an elaboration of its abstract. A good paper should also discuss the broader impact and technical importance of its contribution [Stent]. Furthermore, the author should also provide a clear description of their research methodology (that is, their experiment, system, theory, simulation, etc.). The paper should describe and analyze the results of the work, as well as include some ideas for non-trivial future work to be performed on this or similar open problems. Lastly, related work should be described and cited clearly and correctly. After quickly skimming a paper, the reader should be able to provide the type of paper (theoretical, engineering, or empirical), the area of computer science to which it applies and, most importantly, the problem the paper addresses.

2. READING COMPREHENSION

Returning to the discussion of methods, the reader should focus on reading what the paper actually "'says"'; that is, he should not simply focus one-sidedly on the technicalities. To this end, readers should ask of themselves the following four questions while reading. We will later critically examine the answers to these and other questions.

(1) **What is the research problem the paper attempts to address?** This is perhaps the most important single question the reader can ask. By knowing the answer to this question, the reader may choose to skim or altogether omit reading this paper. This question is also critically important to evaluating the paper's methodology, as well as analyzing its conclusions (questions related to these items follow next). To this point, a high-quality paper should clearly define its research problem, making the reader's task here easier.
(2) **What are the claimed contributions of the paper?** (That is, what new or innovative algorithms, methods, tools, systems, evidence, etc. can the author claim?)

(3) **How do the authors substantiate their claims?** What methodology was used to do so? What arguments, theorems, experiments, data analyses, simulations, benchmarks, case studies, etc. were used? In short, what makes their claims scientific rather than mere unsubstantiated statements?
(4) **What are the conclusions?** Is the result generalizable? Can the result be applied to other areas? What are the open problems?

3. EVALUATION

After having read the paper and answered the questions above, the reader should then move to examine the statements, methods, and conclusions of the paper. First, was the stated research problem significant? Is the problem trivial or artifical? Work that does not enable practical applications or broaden the knowledge of the field should be evaluated as having less importance than papers to which we can answer these questions affirmatively.

Next, even if the research problem is significant, we should examine whether the contributions of this paper were themselves significant. Is the author simply repeating the state of the art? Additionally, pay close attention to any assumptions or limitations (stated or otherwise). What may sound like a contribution that addresses a well-known open problem may in fact only address a very niche subset of that problem.

Moving on, even presuming the paper represents a significant contribution to the problem, we need to verify the validity of any claims or results. When answering this question, we seek to uncover any way in which the paper or its conclusions are technically incorrect or unsubstantiated. We should also examine whether the author has proven exactly what they claim to have proved. Here, we should carefully check for flaws in logic, math, or experimental design. Did the author miss any confounding variables? Were their simulations or benchmarks too artificial or idealized to draw the conclusions the author wishes to make? In other words, are their claims modest enough? [Stent] reminds us, "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."

Lastly, [Keshav 2007] suggests we should first skim a document with a mindset toward being capable of answering "'The Five C's"' after doing so:

(1) **Category** What type of paper is this?
(2) **Context** Which other papers are related? Which theories underlie the problem?
(3) **Correctness** Are the assumptions, methods, and conclusions valid?
(4) **Contributions** In what ways does this paper contribute to the scientific community?
(5) **Clarity** Is the paper well-written?

4. CONCLUSION

More than other readers, the researcher needs to take an active role in the reading process. He should read suspiciously, remaining alert to unstated assumptions, unsupported claims, inflated statements of importance, and sweeping conclusions. To do so, the reader should strive to find answers to some of the key questions presented in this review. While this list was by no means exhaustive, it provides a starting point from which one can adapt and develop additional questions as they see fit. The most important idea, however, is that the reader should be asking these questions from the outset.

REFERENCES

Philip W.L. Fong. 2009. Reading a Computer Science Research Paper. *SIGCSE Bull.* 41, 2 (June 2009), 138–140. DOI:http://dx.doi.org/10.1145/1595453.1595493

S. Keshav. 2007. How to Read a Paper. *SIGCOMM Comput. Commun. Rev.* 37, 3 (July 2007), 83–84. DOI:http://dx.doi.org/10.1145/1273445.1273458

Amanda Stent. How to Read a Computer Science Research Paper. (????). Retrieved September 7, 2015 from http://uet.vnu.edu.vn/~chauttm/cs-english/reading-guides/howtoreadacspaper.pdf