

How to Conduct a Computer Science Literary Review

Amanda Wulf
University of Texas at San Antonio
1 UTSA Circle
San Antonio, TX 78249
amandakwulf@gmail.com

ABSTRACT

The purpose of this document is to provide a structural overview of how to go about conducting a computer science literature review. It provides a detailed account of how to read a computer science paper and how to perform a systematic review of many papers.

Categories and Subject Descriptors

A.1 [Introductory and Survey]: Miscellaneous

General Terms

Documentation

1. INTRODUCTION

Learning how to effectively conduct a literature review in one's field is vitally important. A typical researcher spends hundreds of hours per year reading research papers [3] in an effort to keep up with the latest advances in his or her field, as well as in order to come up with new ideas in his or her own research.

Becoming well-read is about quality over quantity [1], so it's better to read ten research papers well than one hundred papers poorly. One might wonder how one goes about reading research papers well in the first place. This document serves to help researchers achieve that goal.

2. READING GUIDELINES

There are two goals of reading a computer science paper: to understand the scientific contributions the authors are making [4], and to help you come up with new research ideas [2].

There are three main goals people pursue when reading in general: entertainment, information, and understanding [1]. This document will not discuss reading for entertainment, as people reading computer science papers generally read them for information and understanding.

Reading for information involves learning new facts that you didn't know before. For example, after reading a paper about databases you might learn about a new type of databases you've never heard of. On the other hand, reading for understanding involves enlightenment of the mind. It involves understanding why things are the way they are and making connections between many disparate pieces of knowledge [1].

Becoming informed is a prerequisite for becoming enlightened, but in general it is better to not be satisfied solely with learning new facts when reading papers, and to try to understand the field of computer science more deeply with each paper one reads. Therefore, this document will touch briefly on reading for information and go into reading for understanding with more depth.

The definition of reading for understanding is "the process whereby a mind, with nothing to operate on but this symbols of the readable matter, and with no help from outside, elevates itself by the power of its own observations" [1]. This is a difficult goal for anyone to achieve, particularly if they are used to reading for entertainment only, because it involves analyzing your own thoughts [1].

The goal of reading for understanding fits in nicely with the first goal of computer science paper reading (understanding the scientific contributions of the authors), so it will be discussed in detail.

In order to begin reading for understanding, one must understand the four levels of reading: elementary, inspectional, analytical, and literature review [1]. Each level encapsulates the others, so this paper will touch on each briefly in order.

2.1 Elementary Reading

Elementary reading is the level at which anyone who has learned to read can read. It involves reading every word in a document from the beginning to the end. If you are able to read this document, you are capable of elementary reading.

Performing an elementary reading alone is fine for reading for entertainment, but it will not work well for reading computer science papers, so this document will not cover it in detail.

2.2 Inspectional Reading

As a prerequisite to performing inspectional reading, one must become an active reader. Active reading is defined as *asking questions while reading*. [1]. Active reading forces the mind to engage with the text and prevents it from falling asleep or drifting off onto another train of thought. Active reading is especially important for technical reading such as computer science reading, because the papers require more concentration to understand than those of some other disciplines.

Here are some tips to keep reading active.

- Print papers to read rather than reading them on your computer, to prevent distractions from websites or games and to allow for easier paper annotation.
- Do your reading at your desk (rather than the couch or bed) to avoid dozing off.
- Keep a highlighter and pen handy while reading so you can highlight the important sentences in the paper and take notes in the margins of the paper or in a notebook.

Active reading is necessary but not sufficient for inspectional reading. Inspectional reading involves two passes through the paper. The first pass is a brief skim of the paper, then the second pass is a superficial reading [1].

On the first pass, read the abstract, introduction, headings of the paper's body, conclusion, and references. The goal of the first pass through a paper is to get a good high-level idea of what the paper is about, so you know what to look for while doing the next pass, as well as to help determine if you want to do another pass at all. The entire process should take about five to ten minutes [3].

On the second pass, read the entire paper all the way through without getting too bogged down by the details. Highlight and take a few notes as you go, but don't stop to look up or ponder things you don't understand, and try not to labor over each premise [1]. This helps you get a bird's-eye view for a later deeper reading, and again helps you determine whether this paper is worth an in-depth reading. The second pass is also sufficient for papers you're interested in but that are not specifically related to your research. It should take you approximately one hour to complete [3].

During and after the second pass, you should take notes on the structure of the paper. Write down what kind of paper it is, a summary of it in your own words, and a sketch of its structure. This may or may not be a one-to-one correlation with the headings and hierarchies the authors have created.

2.3 Analytical Reading

Analytical reading takes inspectional reading one step further. It involves not only reading the words and phrases, but also analyzing the thoughts behind them [1].

A prerequisite to an analytical reading is an inspectional reading. Once you have completed an inspectional reading, it's time to give the paper its third pass.

In this pass, you should be taking conceptual notes that address the whys of the paper, not just the whats [1]. The analytical stage involves first breaking a paper down to its bare-bones structure, then building up from there to the level of all arguments made in a paper. If you perform this stage well, you will understand the paper inside and out [1].

There are two stages of analytical reading which can be completed either in subsequent readings or at the same time.

The first stage of analytical reading involves taking the following notes: what kind of paper it is, a 1-2 sentence summary of the entire paper, a sketch of the paper's structure, and the problems and questions the authors set out to answer with this work [1].

The second stage of analytical reading of a paper involves finding and understanding the important terms, understanding the author's claims, understanding the arguments the authors make based on those claims, and finally understanding the solutions the authors came to [1]. It's also crucial to understand the mathematics and proofs in the paper.

After you have completed these two stages, the next step is to evaluate whether you have truly completed a good analytical reading. Do you understand the paper, inside and out? If the answer is no, you should go back through and do another in-depth analytical reading. If you have completed a second in-depth analysis and you still don't understand the paper, it's time to turn to some kind of aid to help your understanding. The aids you should resort to, in order, are books on the subject or similar papers, then other people's writings about this paper if they exist, then Google and Wikipedia, and finally your advisor if you're really stuck [1].

Once you have completed this, you completely understand the paper. Then, and only then, are you qualified to criticize the paper. Begin this by listing out the authors' assumptions, arguments, and conclusions. When constructing the authors' arguments, look for reasons the authors give for assumptions and conclusions. Discriminate between inductive and deductive reasoning, and discriminate between axiomatic assumptions and propositions with evidence behind them. Computer science papers often involve more inductive than deductive reasoning, so make sure you understand the premises [1].

Once you have this list, you are ready to determine whether you agree or disagree with the paper. There are four ways in which a paper can be wrong: either the authors are uninformed, the authors are misinformed, the authors' reasoning is not cogent, or the authors' analysis is incomplete. If you believe one of these is the case, write a summary detailing why. Make sure your disagreement is not due to emotion or disagreement on the assumptions that the authors have chosen, and try to take the authors' point of view to make sure you're not disagreeing with them unfairly [1].

If you completely understand the paper and have not found any of the four cases of disagreement to be the case, then you must agree with the paper, even if you don't like the conclusions [1].

Whether you agree or disagree with the paper, you must also evaluate whether the problem and contributions of the paper are significant [2] and determine whether you think this paper helps give you a new research direction to traverse. Add this to your notes as well [2].

This analytical stage of reading a paper should take approximately four to five hours for beginners, and around one hour for experienced computer science paper readers [3].

3. DOING A FULL LITERATURE SURVEY

Performing a literature survey is more than reading just one paper. It involves compiling a list of papers to read in order to increase your overall understanding of the computer science field or subfield, or to learn more about a particular subtopic within computer science. It's impossible to read all the papers on one subtopic as thousands of computer science papers are published yearly, so doing a literature survey is a good way to make sure the ones you do read are good papers [5].

The first step in a literature survey is to compile a list of papers to read. In order to do this, first decide the subtopic you're going to read about, then go to Google Scholar and find the three to five most recent papers about the subtopic. Look in the bibliographies of these papers and write down the authors you see cited in multiple ones. Then, go to websites where they have recently published and add the recent papers there to your list of to-read papers [3]. Make sure you prioritize your to-read papers by citation count and whether the authors are known researchers [5].

Once you have this list, you should do an inspectional reading of each paper on your list. This gives you an overarching idea of what the field and what each paper is about without getting too bogged down in the details. It also gives you the chance to decide whether to cull some of the papers on your list that are irrelevant, and it helps you find the most relevant parts of each paper to make the later process of reading them analytically easier [1].

Next, begin an analytical reading of the papers remaining on your list. Keep a central journal in which you take your notes and write down the summaries of each paper [5]. When multiple papers disagree on one topic, analyze the discussions to determine your own opinion as objectively as possible [1].

4. CONCLUSIONS

This document has laid out a framework for reading a computer science paper, as well as a system for performing a literature review of many papers. There are many rules and guidelines listed here which might take some practice before they become natural, but if you put in the time and effort to get good at performing literature reviews, you will have a worthwhile habit [1].

5. REFERENCES

- [1] M. J. Adler and C. V. Doren. *How to Read a Book: The Classic Guide to Intelligent Reading*. Simon and Schuster, New York, New York, revised edition, 1972.
- [2] P. W. L. Fong. How to read a cs research paper?

- [3] S. Keshav. How to read a paper. *ACM SIGCOMM Computer Communication Review*, 37(3):83–84, July 2007.
- [4] M. Mitzenmacher. How to read a research paper.
- [5] A. Stent. How to read a computer science research paper.