

How to read a research paper

Michael Mitzenmacher

Norman Ramsey

February 3, 2000

One skill you will learn in this course (if you haven't already) is how to squeeze the juice out of a research paper. Although not every research paper is worth careful reading, let us assume that you have a good reason to read a paper carefully. Perhaps it has been assigned for this course, or you have been asked to review it, or it is relevant to your own research. We might later discuss how to *skim* a paper, so that you can decide whether it is worth a careful reading.

When you read a research paper, your goal is to understand the scientific contributions the authors are making. This task is not easy—and the low standard of scientific writing doesn't help. Most people, even seasoned professionals, have to go over some papers several times to understand the contributions. Expect to take *several hours* to read *one* paper.

Here are some guidelines.

- *Read critically:* Reading a research paper must be a critical process. You should not assume that the authors are always correct. *Be suspicious.*

You can suspect almost anything. Critical reading involves asking appropriate questions. *What kind of paper is it?* Scientific contributions can take on many forms. Some papers offer new ideas; others implement ideas, and show how they work; others bring previous ideas together and unite them under a novel framework. New ideas themselves take many forms: big ideas, little ideas, algorithms, new techniques, even (rarely) new ways of thinking.

Was this paper worth doing at all? *What is the problem the authors are trying to solve?* Is it the right problem? Are there simple solutions the authors do not seem to have considered? What are the limitations of the solution (including limitations the authors might not have noticed or clearly admitted)?

Are the assumptions the authors make reasonable? Is the logic of the paper clear and justifiable, given the assumptions, or is there a flaw in the reasoning?

If the authors present data, did they gather the right data to substantiate their argument, and did they appear to gather it in the correct manner? Did they interpret the data in a reasonable manner? Would other data be more compelling?

Some authors may try to persuade you of a particular position. What is the evidence in favor of that position?

- *Read creatively.* Reading critically is easy; it is always easier to tear down than to build up. Reading creatively involves harder, more positive thinking.

What are the good ideas in this paper? Do these ideas have other applications or extensions that the authors might not have thought of? Can they be generalized further? Are there possible improvements that might make important practical differences? If you were going to start doing research from this paper, what would be the next thing you would do?

- *Make notes as you read the paper.* Many people cover the margins of their copies of papers with notes. Use whatever style you prefer. If you have questions or criticisms, write them down so you do not forget them. Underline key points the authors make. Mark the data that is most important or that appears questionable. Such efforts help the first time you read a paper and pay big dividends when you have to re-read a paper after several months.

- After the first read-through, *try to summarize the paper* in one or two sentences.

Almost all good research papers try to provide an answer a specific question. Sometimes the question is a natural one that people specifically set out to answer; sometimes a good idea just ends up answering a worthwhile question. In our disciplines, the questions are usually *scientific questions* (e.g., what are the properties of system X or formalism Y ?) or *engineering questions* (e.g., how can I build system Z ?).

If you can succinctly describe a paper, you have probably recognized the question the authors started with and the answer they provide. Now you can go back and outline the paper, to gain insight into more specific details. If summarizing the paper in one or two sentences is easy, go back and try to deepen your outline by summarizing the three or four most important subpoints of the main idea.

- *Compare the paper to other works.* Summarizing the paper is one way to try to determine the scientific contribution of a paper. But to really gauge the scientific merit, you must compare the paper to other works in the area. Are the ideas really novel, or have they appeared before? As a beginning researcher, you may not know the answer. How would you find out?

Because scientific contributions take so many forms, knowing other work in the area can help you to determine what kind of contribution a paper is actually making.

For CS 257, you will be asked to present research papers. To supplement your presentation, I may ask you for a short, one-page review. Preparing this review will help you prepare your presentation. Although this may sound like a simple assignment, it will take a significant amount of time, especially in the beginning. (It may take several hours just to read the paper!) Keeping these guidelines in mind should make the process easier.

One-page reviews should include

- A one- or two-sentence summary of the paper.
- A deeper, more extensive outline of the main points of the paper, including, e.g., assumptions made, arguments presented, data analyzed, and conclusions drawn.
- Any limitations or extensions you see for the ideas in the paper.
- Your opinion of the paper: primarily, the *quality of the ideas* and its *potential influence*.