How to Read a Technical Paper

by Jason Eisner (2009) [https://www.cs.jhu.edu/~jason/advice/how-to-read-a-paper.html, accessed on 2/10/2019]

Multi-pass reading: Skim the paper first, skipping over anything that would take much mental effort. Just get an idea of where the paper is going, why it was written, what's old hat and what's new to you. To force yourself to keep moving, give yourself a limited time budget per page or use the autoscroll feature of your PDF reader.

Now, assuming the paper still seems worthwhile, go back and read the whole thing more carefully.

Why not practice on this webpage? Go ahead, skim it first.

S. Keshav describes three-pass reading in detail: What are you trying to do on each pass?

Write as you read: Write as you read. This keeps your attention focused and makes you engage with the paper.

Low-level notes: Often it is easiest to scribble notes on the printed-out paper itself, responding in context to the formulas, figures, and text. In that case, file or scan your annotated copy for future reference.

(Or perhaps annotate the PDF *file* directly, without printing or scanning. A free alternative to Acrobat is PDF-XChange Viewer, a Windows program that can also be run on Linux via wine. A free native Linux option is Xournal.)

You can use notes on the paper to

- restate unclear points in your own words
- fill in missing details (assumptions, algebraic steps, proofs, pseudocode)
- annotate mathematical objects with their types
- come up with examples that illustrate the author's ideas, *and* examples that would be problematic for the author
- draw connections to other methods and problems you know about
- ask questions about things that aren't stated or that don't make sense
- challenge the paper's claims or methods
- dream up followup work that you (or someone) should do

High-level notes: Low-level notes aren't enough. Also keep high-level notes about papers. You should try to distill the paper down: summarize the things that interested you, contrast with other papers, and record your own questions and ideas for future work. Writing this distillation gives you a goal while reading the paper, and the notes will be useful to you later.

Michael Mitzenmacher writes: "Read *creatively*. Reading a paper critically is easy, in that it is always easier to tear something down than to build it 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?"

At a minimum, you should re-explain the ideas in your own words: produce some text that is aimed at your future self. You should be able to reread this later and quickly reconstruct your understanding of the paper. Don't waste time repeating the parts that are easy for you. Include a URL to the original paper, and refer as needed to the paper's Figure 1, equation (2), section 3.3, etc. But do spend time writing down hard-won bits of understanding:

"They don't say this, but equation (2) is basically the same as the method of Pookie (2001), except that they add a reconfabulation step after the data purée. I was surprised at their reconfabulator, which doesn't match what I would have expected from Kachu (2004), but it does cure the exponential growth problem in this domain. To see the difference, I found it useful to think about this example: ..."

Organizing your notes: I suggest sorting your file of notes chronologically, by when you read the paper, since that may help you find vaguely remembered papers or remember what else you were reading at the time. Sometimes you'll want to search by author/title/etc., so start the notes for each paper with a rough citation. (See also How to Organize Your Files.)

If you had to put a lot of effort into really understanding some point, you can share that effort with others (and record it for your own future reference) by improving the discussion of that point on the relevant Wikipedia page.

Many people have devised software or personal systems for annotating papers and keeping track of notes. Quora users give their recommendations here and here.

When and where to read: Start early. Leave enough time that if your attention wanders, you can put the paper down and pick it up again when you're in a better reading mood. This is better than trying to force yourself through it on a deadline.

Some people find it easier to read at particular times of day, or while eating or walking or riding an exercise bike. Do you habitually pick up the closest thing to read when you're at the breakfast table or in the bathroom? Then leave papers there for yourself.

Try reading with a friend! Sit next to each other, looking at the same copy of the paper, and stay synchronized at the paragraph or sentence level. Read aloud at times. You'll keep each other moving and help each other through the hard parts. Discuss as you go along.

Set aside time: When you are starting out in a new area, it may take you hours to read a conference paper thoroughly. That's okay. It's worth spending that much time to really understand a good or foundational paper. It will pay off in your future reading and research.

I'll never find the time! Don't worry. Not all papers take that long. Many ideas are reused across papers, so you will get faster at reading. By now, in an area I know well, I can often read a paper in 30 minutes or less, because the motivation is familiar and I can recognize much of the setup as standard practice. (After all, most papers fall into an existing tradition. They extend existing work with one or two genuine new ideas, and some supporting details that may or may not be significant.)

But I'm already a third-year student. Why is this paper taking me so long? There is no shame in reading slowly. It still takes me several hours to absorb a paper on something that I genuinely don't know well. (Also, it takes me hours to *review* a paper even in my own area, because the burden is on me to spot all the problems or opportunities for improvement. 75% of submitted conference papers are rejected, and most of the remaining 25% also need improvement before publication.)

Which parts to focus on: So do you really have to read the *whole* paper carefully on your second pass? Sometimes, but not always. It depends on why you're reading the paper.

I do think that when you are learning a new area, you should read at least *some* papers extremely thoroughly. That means knowing what every sentence and every subscript is doing, so that you really learn all of the techniques used in the paper. And understanding why things were done as they were: ask yourself dumb questions and answer them. Practice the *ability* to decode the *entire* paper—as if you were reviewing it critically and trying to catch any errors, sloppy thinking, or incompleteness. This will sharpen your critical thinking. You will want to turn this practiced critical eye on yourself as you plan, execute, and write up your own research.

However, there will also be occasional papers where it is not worth reading all the details right now. Maybe the details are of limited interest, or you simply don't feel equipped to understand them yet. Consider the parts of a typical paper:

- **Motivation.** You'll want to understand this fairly well, or there's no point in reading the paper at all. But part of the motivation may depend on things you don't know (mathematical background or past work). If you don't want to chase those references down now, you could just raise their priority on your reading list.
- **Mathematics and algorithms.** These parts are the technical heart of the paper. So don't make a habit of skimming them. (You can learn a lot from how the authors solved *their* problems.) Nonetheless, you might skim a technical section if
    - It seems like an explanation of something you already know. In that case, just check that it really says what you think.
    - While you probably would benefit right away from knowing the method in detail, this paper is just not a good place to learn it, or it is too advanced for you right now. Understand what you reasonably can, and then put it on your list of things to learn for real. Perhaps ask someone else to explain it to you or to recommend a reading.
    - It seems like an ugly *ad hoc* solution that no one would ever want to use anyway. The only reason to understand it fully would be if you wanted to criticize it or improve upon it. (Still, even if you skip the ugly details, understand what the authors' intuitions were. Think about how to capture those intuitions more elegantly.)
    - It's enough to know for now that the method exists. It seems specialized, so you might never need it. You'll come back to the paper if you do.

But you should still achieve clarity now about what the method accomplishes (its interface). Also try to glean when it is applicable, how hard it would be to use, and what determines its runtime and accuracy. Then you'll remember the method when you need it.

What you might skip for now are the hard parts: the internal workings of the method (its implementation) and any proofs of correctness or efficiency.

- **Experiments.** Many papers test their methods empirically. When you're new to a field, you should examine carefully how this is typically done (and whether you approve!). It can also be helpful to notice what datasets and code were used—as you may want to use them yourself in future.

But once you've learned the ropes, you may not always care so much about a paper's experiments. After all, sometimes you're only reading the paper to stoke your creativity with some new problems or techniques. I confess that I often pay less attention to the experimental details—though examples or error analysis do catch my attention because they often shed light.

If you do care about the conclusions of the paper ("did the method work?" "should I use it?"), then you should go back and carefully examine the experimental design, including the choice of data. Were the experiments fair? Do they support the claims? What's really going on? Are the conclusions likely to generalize beyond this experimental scenario?

In short, invest your time wisely. Focus on what is valuable to take away. *If you can't figure out which parts of the paper are most "interesting" or "important," do ask someone who should know!* If you don't know who to ask, find other papers that cite this one (via Google Scholar) and see what they say about this paper.

Delip Rao suggests: "Never read the original paper on X first. Instead read several later papers on what they say about X, get an idea of X and then read the original paper. Somehow the research community is much better in explaining ideas clearly than the original authors themselves."

What to read:

- do creative web search
    - experiment with several searches
    - put yourself in an author's shoes; what phrases might they have used?
    - become a power searcher! (read the help pages for your search engine)
    - specifically search at the ACL Anthology, Google Scholar, etc.

- track down related work (once you've got a relevant paper)
    - **backward references**: follow the bibliography to earlier papers
    - **forward references**: see who else has cited the work (via an interface such as Google Scholar)

- has someone else already listed the right papers for you?
    - survey papers in journals (also called "review articles")
    - course syllabi
    - reading group webpages
    - chapters in textbooks
    - online tutorials
    - literature review chapters from dissertations
    - direct recommendations from friends or professors (perhaps at other institutions)

- breadth-first exploration
    - read a lot of abstracts (and skim the papers as needed) before deciding which papers are best to read
    - it's okay to read multiple related papers at once, flipping back and forth so that they clarify one another
    - to get a feel for the research landscape in an area, flip through the proceedings of a relevant recent workshop, conference, or special-theme journal issue

- when the going gets tough, switch to background reading
    - textbooks or tutorials
    - review articles
    - introductions and lit review chapters from dissertations
    - early papers that are heavily cited
    - sometimes Wikipedia