How to Read a Scientific Paper

Gordon J. Pace
August 2006

Warning: *This document is written in a rather dogmatic style. You may not agree with some of the draconian measures recommended for effective reading of a scientific paper, because you can (or may believe that you can) do without some of these techniques. But first try reading a paper using this approach. Only then try others. If you find that you can be just as effective with less rigour, then go ahead and use your approach. But only then.*

Preparation 0:

0. **Is it the right paper?** Start by checking where the paper was published. Is it a journal paper, a conference paper, a workshop paper or a technical report? This usually indicates the importance of the contribution (as presented in the paper). Obviously, some technical reports eventually become published in a conference or a journal. Check whether this was the case. If so, you may be better off starting with the journal paper. Or maybe the technical report contains some details about certain proofs or algorithms not given in the conference or journal version.

Preparation 1:

1. **Print the paper:** No matter how used to reading documents off the screen you are, print the paper. You need to have a hard copy to stare at, highlight stuff and doodle on. Also, you will switch off your computer (see step 3) to effectively read a paper, so you cannot read a paper off a switched off monitor.
2. **Prepare a pile of empty sheets of paper:** You need somewhere to take notes, rough notes. I usually keep a notepad with notes about the important papers I read, but for every page I write there, usually I have scribbled dozens of pages notes on rough paper. Not using your pristine, neat notepad encourages you to scribble and write all sorts of stuff while reading the paper. You can always be selective later on and

copy bits and pieces (leaving out the doodles, poems and the shopping list).

Preparation 2:

Reading a paper is a difficult task, a very difficult task. Remember that your aim should *never* be to skim through the contents to get an idea of what the authors are saying, but to understand the contents of the paper in full detail, and be able to argue for and against different issues raised by the paper in a concrete manner. This requires concentration, and unless you focus on the paper for long stretches of time, you cannot fully take in its content. Although different people can concentrate in different environments (other people in the room, music in the background, etc), there are some things with which you cannot focus on the task at hand. So,

3. **Switch off your computer;**
4. **and your mobile;**
5. **and lock the door.**

No, you shouldn't consider the beep of an incoming email or SMS to be an alarm marking time for a short break. If people (or students for that matter) tend to drop in to your office for a discussion or chat, lock the door. Set aside time to read a paper, and consider it as though you are in the middle of an important meeting. Would you stop for a chat on MSN half way through an important meeting? Would you browse for a couple of minutes to take a break?

Skimming through the paper:

Your aim is to start by working out what the paper is all about – whether it is really relevant, and whether you should invest the time to read it thoroughly. Should you take notes at this stage? Should you highlight important parts of what you read? Since, at this stage you will be trying to obtain a correct overview of the paper by reading the less technical sections, I think that you should not. Also, these sections usually have a higher proportion of buzzwords, fluff and somewhat irrelevant waffle. Learn to ignore these, and concentrate on the actual content.

Conference/journal: Check where the paper was published. In which journal, or conference did it appear? This usually gives you a hint as to the emphasis of the paper. Some conferences concentrate more on

certain techniques than others (despite the fact that this may not necessarily appear in the conference name or description). At first you will not be able to do this, but as you go along, after reading tens of papers, you will start being able to judge this.

Title: A well-written paper usually has a carefully-chosen title with no extra, no missing words. Read the title carefully. Remember, usually every word is there for a purpose. And missing words are usually not there because they are irrelevant (although in some cases the conference or journal where the paper appeared may imply these words).

Abstract: The abstract contains a summarised version of the paper, leaving out all technical detail, but clearly identifying the motivation and contributions of the paper. This should give you an idea of what is to come, and will help you identify what is new in the paper you are reading. Don't, I repeat, don't try to judge or understand a paper after reading the abstract. No matter how much reading between the lines you are able to do, there is never enough in the abstract to conclude anything.

Introduction: So, if the abstract is a summary of the paper, what is it that usually appears in the introduction? Usually, you will get a summary of the *area* in which the contribution of the paper is set – giving you a context in which to understand the paper's result better. You will not get any technical background here – that belongs in another section, but you will still get a good overview of the knowledge required to follow the paper. If you realise that you need to learn something about a particular related topic, stop, find a relevant paper, and read that before the current one. The more you read in a field, the less you will have to do this, so do not be discouraged if at first you end up backtracking to other papers and books over and over again. Usually, the introduction also summarises the contributions of the paper to the area (more verbosely than the abstract). Read this carefully, and take some notes to remind yourself later on. Finally, at the end of the introduction, you will usually find a paragraph or two explaining the structure of the rest of the paper. You can ignore this at this stage.

Conclusions: Next skip to the concluding section. Usually, you'll find a more detailed summary of the contributions, and their context – how they performed compared to similar techniques, what impact their contribution can have (including further to that presented in the paper). Keep in mind the contribution to the field are listed in the

introduction which you should have noted somewhere.

At this stage, you will be able to judge (i) the area in which the paper is making a contribution; and (ii) the contribution of the paper. If you're still reading, the first should indicate that you can read and understand the rest of the paper, and the second will indicate how relevant the paper is to your research. At this stage, it may be useful to take a look at the bibliography to see which papers are cited and thus related to the paper you are about to read.

Reading the paper

Now that you've assessed the paper and decided that it's worth reading, it's time to actually start with the reading. Throughout the reading, make sure you take copious rough notes, which you will eventually refine into shorter ones with the main points and results presented in the paper. Obviously, as you progress in the paper, you will be getting a better overall picture, which will help you improve your notes at the end. Keeping notes avoids unnecessary highlighting, but if you are an incurable highlighter, a die-hard fluorescent colour graffiti artist, a yellow chisel-top junkie, just try not to highlight everything…

To start off your reading, reread the abstract, introduction and conclusions in more detail, and taking notes. This time round, read carefully the paragraph in the introduction which explains the structure of the paper to plan on how best to approach and read the paper.

Most papers contain a section or two with the technical background required to understand the rest of the paper. Although these sections usually contain little more than the notation which will be adopted, make sure you read it carefully and fully understand it. Go through any definitions slowly. Unless you've seen and used the definition before, understand every detail, why it is there, and why others are not. Try examples, try using it, try reproducing it. Similarly, with any algorithm and proof. Go through it step by step. Understand why it works, and why every single part is crucial. Try modifying it, and try to write it yourself from scratch. This is the process that went through the authors' head when selecting and writing the background section. Unfortunately, this is almost never documented, so it is up to you to rediscover it and document it in your notes.

Next comes the actual paper contribution and related work sections. Treat

these just like the background sections, but also try to identify the new stuff that the paper is purporting to introduce. But while reading these sections, you should not stop at understanding, but also start judging the content. Try to be sceptical, and do not accept statements made by the authors at face value, whoever the authors may be. When experimental results are presented, ask yourself whether the dataset was in some way skewed. Did they actually present a solution to the problem they identified, or just an instance of the problem? If a proof is presented, ask whether every step is justified sufficiently. Are any constraints made too restrictive, making the theorem almost universally inapplicable? If the paper presents a new algorithm, ask yourself whether they have sufficiently analysed it. If they state that the new algorithm is better when looking at the worst case, ask yourself whether it is also better in the average case. If they analyse the time complexity, ask about the space complexity. Is the algorithm obviously correct? How obvious is obvious?

Other questions to ask when reading the related work section include: Are they comparing to the state of art techniques (at the time of writing)? Have they been comprehensive in their survey of related work? Is the analysis sufficiently thorough? In some papers, it is possible to identify whether related work was studied prior to the work being done (directly or indirectly supporting and influencing the development of the contribution) or whether they simply scanned the literature after doing the work (superficially relating their results to others' techniques). In the latter case, it is usually a bad case of a list of non-committal, non-technical comparisons, and it is your job to decide how the related works (if you are familiar with them, and they are sufficiently related to your work) are actually related.

Finally, reread the concluding sections, and form your overall opinion of the paper.

After Reading the Paper

After the full reading, put the paper aside and think about it, and how it is related to your interest and work. Rewrite the notes into concise ones, with sufficient detail to allow you to recall the paper contents in a year's time. This will be invaluable when you finally start writing up your report or paper. Think about the paper, and compare to related work (even work which was not discussed in the paper), and your own work. Write notes about these comparisons.

Reread your notes in a couple of days to make sure that they still make sense, and that you have left nothing important out. Usually, the couple of days away from the paper give you new insights which you can add to your notes.

Finally

Never underestimate the time needed to understand a paper fully. Do not be discouraged if you discover that you have spend hours to understand a short proof or algorithm, or that you need days to understand a paper. Remember that the more you will read in an area, the faster you will become, but no matter how familiar you are with an area, complex papers need time to understand.

Also remember that your reading was shaped by they problem you are currently trying to solve. The next time you read the paper, you may discover that you will be looking at a completely different side to it. However, your notes will still be of great help to remember and understand again quickly the paper.

If you want to make sure you have understood the paper, you may consider explaining the results to someone else, or give a short talk about it. It is surprisingly more difficult to explain a paper's results to others than to understand it yourself. You will discover that the others' questions will identify holes in your understanding of the paper.

The process of learning how to read a technical paper is a slow and difficult one. However, it is a skill that every scientist must have. It is also a required skill before you can even dare to try to write a paper or a good report. Read, read and read, and only then attempt to write.