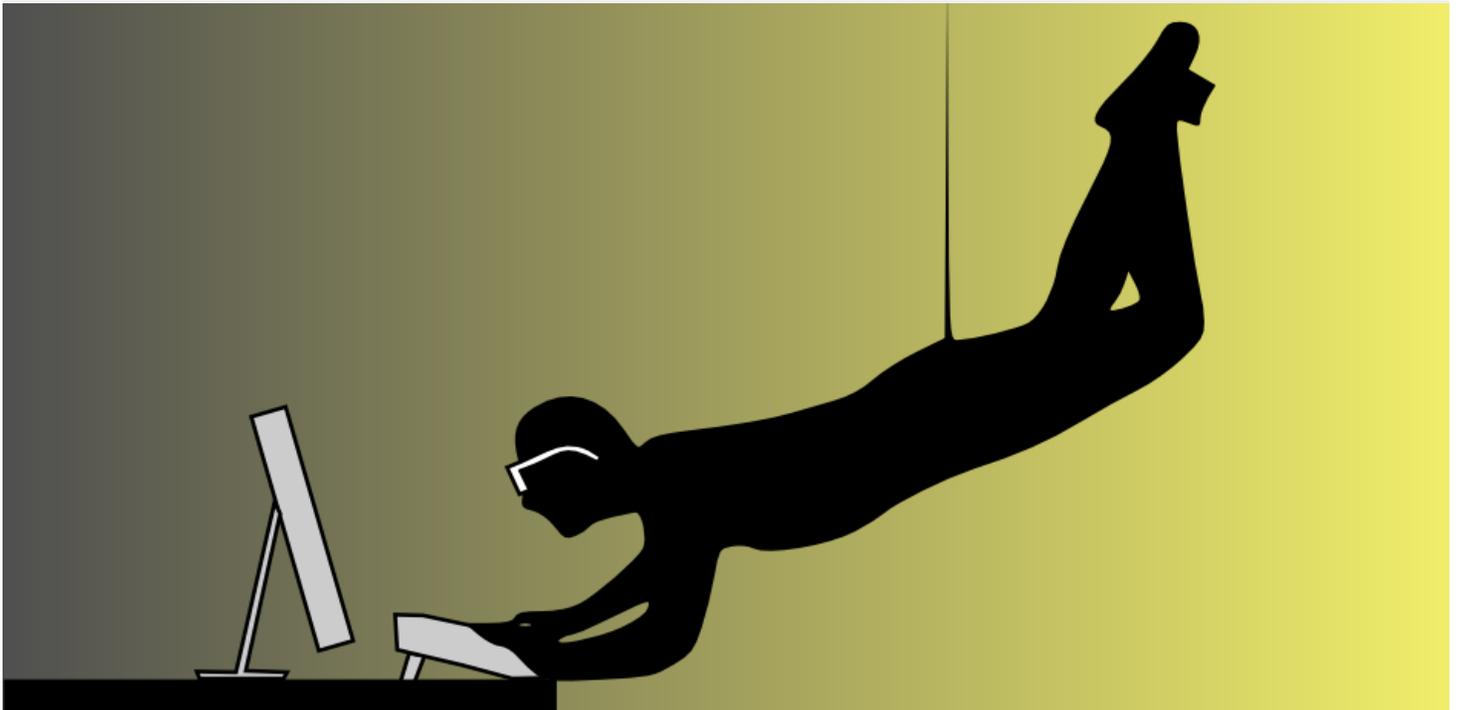


Plotting Your Paper



My guide for scientific paper writing based on starting from an outline and gradually fleshing it out into a full paper.

Plotting Your Paper: A Top-Down Approach to Scientific Writing

by Niklas Elmqvist, University of Maryland, College Park

What I will call “top-down writing” here is a method for writing scientific papers in a top-down manner. The basic idea is to start from an outline of the entire paper and then gradually add content by jumping from section to section until the paper is finished. Rather than completing each section individually (which would be a depth-first approach), you add this detail breadth first.

Top-down writing is by no means the only way to write scientific papers, but I find that it is a useful approach to avoid getting bogged down in details while keeping sight of the structure of the paper. It is particularly useful for novice writers who have a hard time knowing what needs to go into a paper, and where. I also find it particularly fitting to computer scientists, because we’re all used to top-down programming and divide-and-conquer algorithms.

Disclaimer

There are lots of guides on writing academic papers, but they often tend to focus on either macroscopic or microscopic aspects of writing. In other words, such guides will either talk in very general terms on how to write using active voice, how to avoid too long sentences, or using precise language. Or, they will instead discuss details of sentence structure, correct punctuation, or what really goes into an abstract. For this reason, it can sometimes be confusing to understand the “middle part”: how do you actually use these rules and guidelines to structure, organize, and write a full research paper from scratch? That is the goal of this guide: to propose a top-down approach that guides the entire writing process.

Note that this approach is by no means my own invention. I am merely taking standard advice on academic paper writing to its extreme. Virtually every writing guide talks about developing an outline first and then fleshing it out. In this guide, I am saying that you should slavishly use your outline as a guide for how to write your paper, gradually expanding it until you have completed the article.

Bottom-Up vs. Top-Down Writing

In fiction writing, authors can roughly be divided into “plotters” and “pantsers”. Pantsers are authors who write their novel by the “seat of their pants”, i.e. they make up the story as they go without having an exact idea of where it will end. In other words, this kind of writing is done in a bottom-up fashion: while the novelist may have a rough idea where the story is going, the writing will most often progress linearly sentence-by-sentence, paragraph-by-paragraph, and chapter-by-chapter. Put differently, the author will let the story come together from the beginning to end, similar to how the events in the story unfold sequentially. Stephen King is a famous pantsler; in his book *On Writing*, he talks about how writing can be likened to archeology, where the writer is merely unearthing an existing fossil that lies buried in the ground.

A plotter, on the other hand, works top-down by first creating an outline, often following the traditional three-act structure (exposition, confrontation, resolution). The writer essentially plans out exactly what will happen in the story before committing to writing each part. Unlike pantsers, who will follow the story where the characters and events take them, the plotter leaves nothing to chance but ensures that they have captured the basic structure of the story before actually writing it. Famous plotters include J. K. Rowling, John Grisham, and Joseph Heller.

For academic papers, our focus is primarily on content, structure, and preciseness. The turn of a phrase, a particularly clever word, or believable dialogue are all aspects of a text that is of

little importance to the scientific writer. Instead, we are concerned with covering all aspects of the work we are describing, in the correct order, and with the correct amount of detail. In other words, scientific writers are mostly **plotters**. In fact, while very experienced writers can essentially compose their papers in their heads and can thus be considered pantsers, the truth is that in scientific writing, we *always* know the outcome because the paper is reporting on research that has (hopefully) already been conducted. For this reason, it makes much more sense to use a top-down approach to writing. In the next sections I will describe this idea in more detail.

The Thesis

Scientific works all have a *thesis*, which is the overall argument that the work is trying to convey. **All scientific papers should have a single main thesis: no more, no less.** If your paper has no thesis, there is no reason for your paper's existence. You should either work hard to find the thesis, or simply discard the paper. If your paper consists of two or more fundamental theses, you are probably trying to convey too much with a single paper and would be better off splitting it into several papers.

The above statement is easy to misunderstand. A paper can have several different theses at different levels, but there should only be one overarching thesis that summarizes your entire contribution. In other words, your work may have many smaller theses that all work to support the main thesis. The above restriction on only having a single thesis only applies to the main thesis. You can only have one.

Note that this is again all very different from fiction writing. A fiction author may often want to include seemingly useless details and episodes in their story simply for the sake of conveying an atmosphere, refining the personality of the main character, or provide amusement or distraction. A basic rule of fiction writing is that a passage should only be left in if it does two, or preferably three, of the following: (1) advances the plot, (2) develops character, or (3) reinforces the theme. Scientific writing is not fiction, of course, so all text in our "stories" must contribute to our own three goals of (1) clarity (2) conciseness, and (3) completeness.

The Thesis Sentence

How does your thesis manifest itself in your paper? The answer is that it should be clearly expressed as an actual sentence somewhere in your work (possibly in several places to drive the point home). In other words, this thesis sentence summarizes your contribution and therefore can be seen as the heart of your paper. If you wanted to summarize your paper with a single sentence, you should be able to use this sentence directly.

Formulating a thesis sentence can seem difficult. You may wonder how you can ever hope to summarize all of your work in a single sentence. It *is* difficult, but it must be done. If you yourself, the author of the paper, don't have a good idea of what the message of your paper is, how can you expect the reader to know? People need labels to help them think. They want to be able to think of your paper in terms such as "the paper that introduced a new way to do X using Y." They don't want to have to sift through and process an entire paper on their own in search for this label.

Hierarchies of Theses

Remember the discussion above about multiple different, minor theses? Well, it turns out that the concept of a thesis is a multiscale one. Theses exist on all levels of your paper. Just like your entire paper has a main thesis, so will individual sections, subsections, paragraphs, and even sentences have. In other words:

- Every sentence in your paper should have a *thesis word* or *thesis phrase* that you think is key to that sentence.
- Every paragraph should have one (and only one) *thesis sentence* that summarizes the main message of that paragraph.
- Every section or subsection should have one (and only one) *thesis paragraph* or sentence that summarizes the main message of that section or subsection.

The above adequacy and completeness argument also holds: if the sentence/paragraph/subsection/section has no thesis, it has no reason for existence and should be deleted. Similarly, if the sentence/paragraph/subsection/section has more than one main thesis, it should be split into smaller parts, one for each thesis. While these rules may sound confusing, it is quite the opposite: they can actually help us decide which parts of a paper to keep and which to discard (and which to split into smaller parts).

That is all very interesting, you say, but how does this translate to helping me write an actual paper? Well, it turns out that the above multiscale thesis concept can be seen as a prescription for how to write papers from the top and down. The next section describes how.

A Generative Thesis-Driven Model for Writing

A corollary of the multiscale thesis structure in a paper is that the theses in a scientific paper form a tree (not a graph since there should only be one main thesis, i.e., the root of the tree). Another corollary is that there should be no part of the paper that is also not represented by a

thesis (node) in this tree. The logical conclusion of these two corollaries is that it is possible to write a paper by constructing a “thesis tree” in a top-down manner, where new tree nodes are added to provide more support for the parent node. This “tree” is obviously your outline, organized hierarchically. This progressive refinement is repeated until all aspects of the work have been described in sufficient detail (remember: clear, concise, complete). The paper can then be filled out by adding supplemental content that complements these sentences with further details, references, or examples.

Here is this generative thesis-driven model for writing expressed as an algorithm:

1. Define an overarching main thesis for the entire paper. Write it as a sentence. This sentence becomes the root of your thesis tree.
2. For every thesis node at all levels of your thesis tree, determine if it is explained in sufficient depth.
 - If it is, move on to the next thesis node (sibling, or next level if all nodes on this level have been explored).
 - If it is not, add one or more (usually three or more) sub-theses as children to the thesis node that needs further explanation.
 - Use your knowledge of your intended audience to determine whether you need to provide more detail. Err on the side of providing too much detail rather than too little.
3. Repeat Step 2 until no more thesis nodes need to be added.
4. All of your thesis sentences in your thesis tree now become a skeleton of one-sentence paragraphs in your research paper.
5. Start turning this “skeleton” of these sentences into a proper research paper by adding supporting sentences to each paragraph. Supporting sentences are often one of the following types:
 - Clarification or additional details for a statement in the thesis sentence.
 - Description or review of one or several existing research that supports the thesis sentence (with references).
 - Concrete example for a statement in the thesis sentence.

If you follow this algorithm, you will essentially be writing your paper in a breadth-first manner, whereas writing is traditionally taught to proceed depth-first. In other words, using top-down writing, you will first define the main chapters, then the sections for each chapter, then the subsections for each section, and then the paragraphs, and so on. Remember all those

guides that tell you to write the outline of your paper first? This method is a generalization of this method, but applied to all parts at all levels of abstraction in your paper.

Benefits and Consequences

To some people, the above process might seem a little extreme. It takes a diametrically opposite approach to writing compared to what most people do. Breadth-first is certainly not the way you are taught to write in school, and it can be difficult to do at first. However, there are several important benefits that make this approach very powerful:

- **Completeness.** Completeness of content is vital for scientific work. This approach forces you to think about the big picture of your work and will ensure that you do not lose track of every piece of the message you are trying to convey. Furthermore, because of this focus on completeness of content, you will likely have less incentive to revise and polish text too early in the process. You are more likely to concentrate on describing the correct things before you concentrate on the actual grammar and vocabulary.
- **Structure.** Scientific papers live and die with their structure. A logical structure helps the reader find the information they are looking for and gives a professional impression, whereas a bad structure will do the exact opposite. Top-down writing forces you to explicitly think about the high-level and mid-level structure before you start worrying about exact formulations on a sentence level.
- **Composition.** Many novice writers struggle with basic composition: how to write your sentences, how to organize them into paragraphs, and which paragraph comes first? The answer given by the top-down model is simple: a paragraph is a single thesis sentence and zero or more supporting sentences that clarify, refine, reference, or exemplify that thesis sentence. Using the top-down writing model, you never need to decide when to stop writing a paragraph because your paragraphs will be growing in parallel, starting with a thesis sentence followed by supporting sentences that get tacked onto them, one at a time, until the paragraph is (1) clear, (2) concise, and (3) complete.

Hybrid Top-Down and Bottom-Up Writing

Having said all this, the top-down writing model described here is indeed extreme, and most experienced writers likely write papers with a combination of both top-down and bottom-up writing. This is fine, because when you have reached the stage when you can comfortably switch between the two depending on the situation, you no longer need to follow this methodology slavishly. You will automatically be writing with these sentences in mind.

As a case in point, it is common for an experienced writer to first define a general outline of section headings in the paper, then focus on each of these sections in turn. For example, it sometimes is a good idea to finish writing the introduction before the rest of the paper. This does not hinder us from taking a breadth-first approach to writing the introduction section, however. When I am in the middle of writing a section, I often find myself adding small phrases on lines by themselves after the paragraph I am currently writing. While these are not complete theses sentences, they help me shape my thinking and not to forget important things I need to add to the section later on.

However, as a novice writer, it is probably a good idea to test the top-down writing idea fully for your next writing project. If you have taken this approach to its extreme at least once, you will have a better awareness of the importance of structure in academic writing, and will hopefully write better papers in the future.

Summary

Top-down writing turns the scientific writer into a plotter who will be constantly working on an outline until the outline eventually metamorphizes into the final paper. In a future update, I will provide a concrete example of one of my papers that was written using this approach as an illustration of this method in practice.

Niklas

Niklas Elmqvist is an associate professor in the iSchool (College of Information Studies) at University of Maryland, College Park.

July 25, 2017

Guides