longer forced to make choices such as of how much detail to give. One can leave that decision to the reader. Such documents have a huge potential to improve the way mathematics is presented, and this potential will only increase as technology improves.

8.3 Letters versus Words

I will not say much about this, since most of what I have to say is very similar to what I have already said about the level of detail in which a document is written. With the kinds of electronic documents that are now possible, one can save the reader the trouble of searching through a paper to find out what a letter stands for by incorporating a reminder that appears when you click on the letter. Perhaps better still, it could appear in a little box when you hover over the letter. One could also have condensed statements involving lots of letters with the option of converting them into equivalent wordier statements. Again, the point is that there are many more options now.

8.4 Modularity

The kinds of electronic documents I have been discussing make possible a form of top-down mathematical writing that would be far less convenient in a print document. One could write a high-level account of some piece of mathematics, giving the reader the option of expanding any part of that high-level account into a lower-level account that justifies it in more detail. And there could be many levels of this, so that if you clicked on everything you would end up with a presentation of the entire argument in full gory detail.

A less ambitious possibility is one that solves the problem discussed earlier about where to place a lemma. The difficulty was that in a print document you will either put it before the proof where it is used, in which case it is not adequately motivated, or during the proof, in which case it looks ugly, or after the proof, in which case the proof itself leaves you with awkward promises to fill in gaps later. But with an electronic document, putting a lemma exactly where it is needed is no longer ugly. During the proof, one can say, "We are now going to make use of the following statement," and give the reader a button to click on that will bring up a proof of that statement.

8.5 Order of Presentation

If you do not want to decide whether to give an abstract definition first or start with motivating examples, then you can give the reader the choice. Just start with a page of headings and invite the reader to decide whether to click on "Motivating examples" first or "The formal definition" first.

To some extent, the same goes for the decision about whether to present arguments in their logical order or in a way that brings out how they were discovered. If at some point the logical order requires you to draw a rabbit out of a hat, you could at the very least introduce a slider that explains where that rabbit actually came from.

VIII.2 How to Read and Understand a Paper

Nicholas J. Higham

Whether you are a mathematician or work in another discipline and need to use mathematical results, you will need to read mathematics papers—perhaps lots of them. The purpose of this article is to give advice on how to go about reading mathematics papers and gaining understanding from them.

The advice is particularly aimed at inexperienced readers. A professional mathematician may read from tens to hundreds of papers every year, including published papers, manuscripts sent for refereeing by journals, and draft papers written by students and colleagues. To a large extent the suggestions I make here are ones that you naturally adopt after reading sufficiently many papers.

Mathematics papers fall into two main types: primary research papers and review papers. Review papers give an overview of an area and usually contain a substantial amount of background material. By design they tend to be easier to read than papers presenting new research, although they are often longer. The suggestions in this article apply to both types of papers.

1 The Anatomy of a Paper

Mathematics papers are fairly rigid in format, having some or all of the following components.

Title. The title should indicate what the paper is about and give a hint about the paper's contributions.

Abstract. The abstract describes the problem being tackled and summarizes the contributions of the paper. The length and the amount of detail both vary greatly. The abstract is meant to be able to stand alone. Often it is visible to everyone on a journal's Web site, while the paper is visible only to subscribers.

Introduction. The first section of the paper, almost always called "Introduction," sets out the context and problem being addressed in more detail than the abstract. Depending on how the paper has been written, the introduction may or may not describe the results and conclusions. Some papers lend themselves to a question being posed in the introduction but fully answered only in a conclusions section.

Conclusions. Many, but not all, papers contain a final section with a title such as "Conclusion" or "Concluding Remarks" that summarizes the main conclusions of the paper. Omission of such a section indicates that the conclusions have been stated in the introduction or perhaps at the end of a section describing experiments, or that no explicit summary has been provided. This section is often used to identify open questions and describe areas for future research, and such suggestions can be very useful if you are looking for problems to work on.

Appendix. Some papers contain one or more appendices, which contain material deemed best separated from the main paper, perhaps because it would otherwise clutter up the development or because it contains tedious details.

References. The references section contains a list of publications that are referred to in the text and that the reader might want to consult.

Supplementary materials. A relatively new concept in mathematics is the notion of additional materials that are available on the publisher's Web site along with the paper but are not actually part of the paper. These might include figures, computer programs, data, and other further material and might not have been refereed even if the paper itself has. It is not always easy to tell if a paper has supplementary materials, as different journals have different conventions for referring to them. They might be mentioned at the end of the paper or in a footnote on the first page, and they may be referred to with "see the supplementary materials" or via an item in the reference list.

2 Deciding Whether to Read a Paper

A common scenario is that you come across a paper that, based on the title, you think you might need to read. For example, you may be signed up to receive alerts from a journal or search engine and become aware of a new paper on a topic related to your interests. How do you decide whether to read the paper? The abstract should contain enough information about the context of the work and the paper's results for you to make a decision. However, abstracts are sometimes very short and are not always well written, so it may be necessary to skim through the introduction and conclusions sections of the paper.

The reference list is worth perusing. If few of the references are familiar, this may mean that the paper presents a rather different view on the topic than you expected, perhaps because the authors are from a different field. If papers that you know are relevant are missing, this is a warning that the authors may not be fully aware of past work on the problem.

If the main results of the paper are theorems, read those to see whether it is worth spending further time on the paper. Consider also the reputation of the journal and the authors, and, unless the paper is very recent, check how often (and how, and by whom) it has been cited in order to get a feel for what other people think about it. (Citations can be checked using online tools, such as Google Scholar or one of several other services, most of which require a subscription.)

3 Getting an Overview

A paper does not have to be read linearly. You may want to make multiple passes, beginning by reading the abstract, introduction, and conclusions, as well as looking at the tables, figures, and references.

Many authors end the introduction with a paragraph that gives an overview of what appears in each part of the paper. Sometimes, though, a glance at the paper's section headings provides a more easily assimilated summary of the content and organization.

Another way in which you might get an overview of the paper is by reading the main results first: the lemmas, theorems, algorithms, and associated definitions, omitting proofs. The usefulness of this approach will depend on the topic and your familiarity with it.

4 Understanding

It is often hard to understand what you are reading. After all, research papers are meant to contain original ideas, and ideas that you have not seen before can be hard to grasp. You may want to stop and ponder an argument, perhaps playing with examples.

I strongly recommend making notes, to help you understand the text and avoid having to retrace your steps in grasping a tricky point if you come back to the paper in the future. It is also a good idea to write a summary of your overall thoughts on the paper; when you go back to the paper a few months or years later, your summary will be the first thing to look at. I recommend dating your notes and summary, as in the future it can be useful to know when they were written. Indeed, I have papers that I have read several times, and the notes show how my understanding changed on each reading. (There exist papers for which multiple readings are needed to appreciate fully the contents, perhaps because the paper is deep, because it is badly written, or both!)

As well as writing notes, it is a good idea to mark key sentences, theorems, and so on. I do this either by putting a vertical line in the margin that delineates the area of interest or by marking the relevant text with a highlighter pen.

I write my notes on a hard copy of the paper. Many programs are available that will allow you to annotate PDF files on-screen, though using mathematical notation may be problematic; one solution is to handwrite notes and then scan them in and append them to the PDF file.

A good exercise, especially if you are inexperienced at writing papers, is to write your own abstract for the paper (100–200 words, say).

Writing while you read turns you from a passive reader into an active one, and being an active reader helps you to understand and remember the contents. One useful technique is to try out special cases of results. If a theorem is stated for analytic functions, see what it says for polynomials or for the exponential. If a theorem is stated for $n \times n$ matrices, check it for $n = 1, 2, 3$. Another approach is to ask yourself what would happen if one of the conditions in a theorem were to be removed: where would the proof break down?

When you reach a point that you do not understand, it may be best to jump to the end of the argument and go back over the details later to avoid getting bogged down. Keep in mind that some ideas and techniques are so well known to researchers in the relevant field that they might not be spelled out. If you are new to the field you may at first need a bit of help from a more experienced colleague to fill in what appear to be gaps in arguments.

It is important to keep in mind that what you are reading may be badly explained or just wrong. Typographical errors are quite common, especially in preprints and in papers that have not been copy edited. Mathematical errors also occur, and even the best journals occasionally have to print corrections ("errata") to previously published articles.

In mathematical writing certain standard phrases are used that have particular meanings. "It follows that" or "it is easy to see that" mean that the next statement can be proved without using any new ideas and that giving the details would clutter the text. The detail may, however, be tedious. The shorter "hence," "therefore," or "so" imply a more straightforward conclusion. "It can be shown that" again implies that details are not felt to be worth including but is noncommittal about the difficulty of the proof.

5 Documenting Your Reading

I advise keeping a record of which papers you have read, even if you have read them only partially. If you are a beginning Ph.D. student this may seem unnecessary, as at first you will be able to keep the papers in your mind. But at some point you will forget which papers you have read and having this information readily available will be very useful.

A few decades ago papers existed only as hard copies, and one would file them by author or subject. Today, most papers are obtained as PDF downloads that can be stored on our computers. Various computer programs are available for managing collections of papers. One of those, or a BibTeX database, can serve to record what you have read and provide links to the PDF files.

6 Screen or Print?

Should you read papers on a computer screen or in print form? This is a personal choice. People brought up in the digital publishing era may be happy reading on-screen, but others, such as me, may feel that they can properly read a paper only in hard copy form. There is no doubt that hard copy allows easier viewing of multiple pages at the same time, while a PDF file makes it easier to search for a particular term and can be zoomed to whatever size is most comfortable to read. It is important to try both and use whatever combination of screen and print works best for you.

If you do read on-screen, keep in mind that most PDF readers allow you to customize the colors. White or yellow text on a black background may be less strain on the eyes than the default black on white. In Adobe Acrobat the colors can be changed with the menu option Preferences–Accessibility–Document Colors Options.

7 Reading for Writing

One of the reasons to read is to become a better writer. When you read an article that you think is particularly well written, analyze it to see what techniques, words, and phrases seemed to work so well. Reading also expands your knowledge and experience, and can improve your ability to do research. Donald Knuth put it well when he said:

> In general when I'm reading a technical paper... I'm trying to get into the author's mind, trying to figure out what the concept is. The more you learn to read other people's stuff, the more able you are to invent your own in the future.

8 What Next?

Having read the paper you should ask yourself not only what the authors have achieved but also what questions remain. Can you identify open questions that you could answer? Can you see how to combine ideas from this paper with other ideas in a new way? Can you obtain stronger or more general results?

VIII.3 How to Write a General Interest Mathematics Book

Ian Stewart

> *I've always wanted to write a book.*
> *Then why don't you?*
> —Common party conversation

Popular science is a well-developed genre in its own right, and popular mathematics is an established subgenre. Several hundred popular mathematics books now appear every year, ranging from elementary introductions through school-level topics to substantial volumes about research breakthroughs. Writing about mathematics for the general public can be a rewarding experience for anyone who enjoys and values communication. Established authors include journalists, teachers, and research mathematicians; subjects are limited only by the imagination of authors and publishers' assessments of what booksellers are willing to stock. Even that is changing as the growth of e-books opens the way for less orthodox offerings. The style may be serious or lighthearted, preferably avoiding extremes of solemnity or frivolity. On the whole, most academic institutions no longer look down on "outreach" activities of this kind, and many place great value on them, both as publicity exercises and for their educational aspects. So do government funding bodies.

1 What Is Popular Mathematics?

For many people the phrase is an oxymoron. To them, mathematics is *not* popular. Never mind: popularization is the art of making things popular when they were not originally. It is also the art of presenting advanced material to people who are genuinely interested but do not have the technical background required to read professional journals. Generally speaking, most popular mathematics books address this second audience. It would be wonderful to write a book that would open up the beauty, power, and utility of mathematics to people who swore off the subject when they were five, hate it, and never want to see it or hear about it again—but, by definition, very few of them would read such a book, so you would be wasting your time.

Already we see a creative tension between the wishes of the author and the practicalities of publishing. As e-books start to take off, the whole publication model is changing. One beneficial aspect is that new kinds of book start to become publishable. If an e-book fails commercially, the main thing wasted is the author's time and energy. That may or may not be an issue—an author with a track record can use his/her time to better effect by avoiding things that are likely to fail—but it will not bankrupt the publisher.

Popular mathematics is a genre, a specific class of books with common features, attractive to fans and often repellent to everybody else. In this respect it is on a par with science fiction, detective novels, romantic fantasy, and bodice rippers. Genres have their own rules, and although these rules may not be explicit, fans notice if you break them. If you want to write a popular mathematics book, it is good preparation to read a few of them first. Many writers started that way; they began as fans and ended up as authors, motivated by the books they enjoyed reading.

Most popular mathematics books fall into a relatively small number of types. Many fall into several simultaneously. The rest are as diverse as human imagination can make them. The main classifiable types are: