

Physics Writing Guide

How to write a Physics Journal Article

Dr. James G. McLean
Dept. of Physics & Astronomy
State University of New York at Geneseo
December, 2001 (with various later updates)

Introduction

This is a guide to the format, structure, and style of modern physics journal articles. You will note that there are significant differences between journal articles and the class lab reports that you have written in the past. There are a few matters of structure which are more rigid (e.g., how to present figures), but overall the structure is less rigid. This means that there is more freedom to put together the information in an effective way, but also more responsibility for choosing what that best way is.

Also keep in mind that journal articles exist in a different environment than lab reports. Most importantly, journal articles are **peer-reviewed**. When an article is submitted to a journal, it is sent to a few other scientists (called referees in this context) who decide whether the article is up to the standards of the journal. This is somewhat like the process of grading lab reports. The biggest difference is that the “grading” is essentially pass/fail. Luckily, unless there is a fundamental flaw in the article, it can be revised and resubmitted.

Overview

There are five major parts to a typical journal article, each handled according to its own rules. Four of these are required, and always in the same order: the Title Block, the Abstract, the Body (the written text), and the Reference List. The fifth piece consists of figures or tables, which are almost always needed to convey the message clearly and effectively. Figures and tables do not go together in a single block, but they are definitely separate from the text, and follow their own rules of presentation.

The following sections describe the structure and the content of each of these parts. Issues of structure are separated from those of content because they are conceptually very different. You will usually start your writing by worrying about content. What are you going to say, how will you order the subjects, what words or figures will best get your point across? This is the more creative step. Then, as you write, you need to follow the rules of structure.

Two other optional parts are Acknowledgements (placed after the Body and before the Reference List) and Appendices (placed after the Reference List). These have the same structure as sections of the Body, even though the Appendices are separated from it by the Reference List.

Title Block

Content

The Title Block contains at least three items, and probably a fourth.

Title: This should convey not just what physical system or effect was studied, but also what tools or techniques were used, and at least a vague idea of what kind of conclusions are reached.

Author List: This should give the names of everyone who made a major contribution to the research. Obviously, the term ‘major’ is open to interpretation. One rule of thumb: no article should ever be published without getting each and every author’s approval. Conversely, if you aren’t comfortable publishing the paper without consulting someone about its content, then they should probably go in the Author List. (The Acknowledgements section, when used, offers the opportunity to recognize people who contributed, but not enough to be an author.)

For published material, you should choose one form of your name that you will use consistently for the rest of your life. This helps others who are looking up your work, and after all, the whole reason for writing an article is to get others to read it. When you are the primary author on a paper, be kind to the other authors and make sure that you get the correct form of their name, too.

What form should you choose? Anything including your first initial and whole last name is acceptable, but I would advise using (full first name) (middle initial) (last name), especially if you have a name that is common. Some journals won’t let you use a full first name, and many abstract services only use first and middle initial; nevertheless, I advocate striving for full first name. Women sometimes prefer to only use initials, to avoid discrimination issues.

Affiliations: The school or institution of each author is indicated. There should be sufficient information to contact each author by mail, normally through the institution where they did the work. Sometimes, an author has moved between doing the work and the paper being published; in that case, the Affiliations section is written as if the author hadn’t moved, but a footnote (or endnote, depending on the journal) is added indicating the author’s current address. Often an email address or phone number for one author, known as the “corresponding author,” is also included in a footnote/endnote.

Dates: The title block will likely also include dates, such as the date of submission or the date that the article was accepted for publication. However, the dates are not the responsibility of the author(s). The publisher adds them according to their own policy.

Structure

Title: The title of the article should be informative, but it can’t be too long. If the length exceeds one line, you should begin to wonder if you could trim it down.

Author List: Typically, people are listed in order of decreasing contribution. Thus, the person who actually does the writing goes first. When using this ordering, a common exception is to put the name of a group leader (e.g., a professor) last. This is convenient because a group leader often makes a career out of focusing on one topic. Therefore, when looking for papers, you often might want a bunch of papers from Prof. X’s lab, even though the papers were written by various grad students, etc.

There are some research groups who choose a different system, however, such as alphabetical order. Therefore, you can’t reliably assume that the first author is the one who did most of the work.

Affiliations: When different institutions are involved, there are two possible formats. One is to list all authors from place *A*, then affiliation *A*, then all authors from place *B*, then affiliation *B*, and so forth. Unfortunately, this method often conflicts with putting the author’s names in the desired order. The second method gets around this by listing all the authors with footnote-style superscripts after each name, one symbol for each affiliation. Then list the

affiliations, prefaced by the appropriate superscript. The effect is just like footnotes, except that nothing is moved to the bottom of the page.

Dates: Since these are handled by the journal, you need not worry about this.

Abstract

Content

The primary purpose of the Abstract is to summarize your results and conclusions. What new and interesting thing are you telling the world? If the whole point is that you have improved the accuracy of some measurement, then you would include the numerical result and uncertainty. Otherwise, you probably wouldn't include the uncertainty. In fact, the best experiments often lead to Abstracts with no numbers at all! A qualitative conclusion that is interesting enough to be published is usually of wider interest and application. (Note: This does not mean that you can improve an experiment by leaving numbers out of the Abstract, of course!)

Along the way, you usually give some description of your equipment or experimental technique. This need not be detailed if it is not the focus of the paper. On the other hand, if the point of your paper is that you have designed a new apparatus or made an improvement to some experimental technique, then it should be featured.

There is also an element of marketing in Abstracts. The Abstract is your advertisement enticing anyone doing an abstract search to read your full article. As a result you often refer to previous work or common assumptions of the past that contrast with your results. Or you might explain in one or two sentences why the field is vitally important.

Whatever the focus of your paper is, it should be clear by the second sentence of the Abstract. If you leave it until the end of the Abstract, it will lose its impact.

Since the Abstract is a summary, nothing should be in it that is not also in the main text. An Abstract is **not** an introduction; the paper should be complete even without the Abstract. One way to ensure this is to leave writing the Abstract until after you have written the rest of the paper.

Structure

The abstract is almost always required to be a single paragraph. It has to be self-contained (i.e., understandable without reading or looking at anything else). It can not have or refer to any figures or tables – it's just straight text. On rare occasions it is appropriate to include a citation to another published work; in that case, you put the bibliographic entry right in the text, not as a footnote.

In some sense, the abstract is not part of the paper. It is an abstract *of* the paper, which implies that the paper was there first, then a summary was *abstracted* from it. As such, the Abstract often doesn't have a section title, but instead uses a slightly different format to distinguish it from the Body (e.g., smaller font or indented margins).

Body

Content

What to include

It is very important to decide who your intended audience is. This is the controlling factor that determines what you need to explain, and what you can assume that everyone knows. When publishing, this is determined by which journal you are submitting to; for *Classical & Quantum*

Gravity you can assume that everyone knows what a ‘metric tensor’ is, but *Nature* is intended for an audience that might well not know.

Quantify everything you can. Why describe something as “small” when you could use “1 cm”? This doesn’t mean that you have to be highly accurate; if the size wasn’t terribly important, you can use “roughly 1 cm.” And the way in which you quantify something might vary; in my example maybe “10% of the total length” would be more appropriate. The general idea is to describe things in the way that is most clear to the reader, and non-quantifying words like ‘small’ require interpretation.

Give an indication of the reliability of any numerical final results. Normally, this is done by specifying an uncertainty. Note that this does not mean that every single measurement you make needs to have an uncertainty. The rules of uncertainty combination often mean that only some of the measurements contribute significantly to the uncertainty of the final result. Also, some types of graphical analysis can produce an uncertainty without having uncertainties for the individual data points.

You do not need to be very explicit about how you calculated the uncertainty. You are at a level now where we assume that you know what you are doing. There should, however, be some indication of the sources of uncertainty, and the general methods used (obtained from a graph fit, propagated using a particular equation, ...).

Equations should be treated as part of the text, extensions of the English language (more on that in Structure below). If you use equations taken from another source (with a proper citation, of course), then it is not necessary to repeat a derivation that is in the referenced work. You must, however, explain where the equation comes from; what assumptions or relations are at the origin of the derivation. This is best if you can actually identify parts of the equation, e.g., “the second term comes from the relativistic correction...”

Figures and Tables are **not** part of the text. As will be explained below, they are accessories to the text. Therefore, whenever you refer to a Figure or Table in the Body text, you should be clear about what you conclude from it and why. The Figure or Table forms a justification for what you say, but you still have to say it in the Body. A reader without access to the Figure or Table should still be able to follow the text.

General Style

It is OK to use the first person (“I” and “we”) in your writing. However, if you find that you are using the first person a lot, it is almost certainly an indication that the focus of your writing is wrong. Chances are good that you are focusing on what actions you took while doing the experiment, rather on what physical relationships were revealed by what happened. Note that the problem is **not** the use of the first person itself; you could rewrite all those sentences in the passive voice, thus getting rid of the “I”s and “we”s, but the incorrect focus of the writing would still be there. Another indicator of a problem with focus is excessive use of words like “first,” “next,” “then,” “afterwards,” and the like, which again reveal a focus on the process instead of the results.

It is not acceptable to use the second person (‘you’) to refer to the reader.

It is **never** appropriate to declare that an experiment was successful or unsuccessful. The reason is that the reader has absolutely no interest in what your original goal was; they just want to know what you did. In your past, it may have been somewhat appropriate to note successful or unsuccessful, because you could compare what you did to instructions in a lab manual. (Although,

even there, I don't know why you would need to tell your instructor what the expectations were; the instructor *wrote* the expectations.) However, in real research, *there is no lab manual*. In this class, there often isn't much of a lab manual, and even when there is, the equipment might not cooperate. In both real research and this class, you just work on a problem as best you can, then report what you accomplished. You might set goals for yourself; in fact, that's a valuable way to keep motivated. But those goals are for your own benefit, and readers of your papers just don't care.

Organizing your thoughts

Often the Body is split into sections (see Structure below). Even when you don't use sections, thinking in terms of sections can help to organize the paper. How you split it is up to you, but here are a few section titles that are frequently seen: Introduction, Equipment (or Setup, Apparatus, ...), Results, Theory, Discussion, and Conclusion.

Note that I ***did not*** list a Procedure section. Back when you were writing lab reports for prepared experiments, it made sense to have separate Procedure and Results sections, because the procedure was determined well in advance. But that often results in needless repetition. (E.g., "Each mass was weighed." in Procedure and later "The masses were found to weigh 1, 2, and 3 N." in Results.) If there is a special measurement technique you used, explain it when you describe the applicable equipment. Otherwise, just include any methods when you give the results.

Note that I ***did not*** list an Uncertainty section. Any calculations or equations for uncertainties on specific quantities should appear at the point where you discuss the quantities themselves. (Although recall that including such equations at all would be rare, only when there is something unusual about them.) Any problems that you had with the equipment should be discussed in the Equipment and/or the Conclusion sections.

The order of the sections should be governed by two principles: (1) You shouldn't write something that requires knowing something else that comes later in the article; and (2) Try to get the ideas to flow naturally one into the other. For instance, if you can't describe the results without reference to the theory, then the theory had better come first. On the other hand, putting theory after results sometimes flows well into a discussion (i.e., comparison of theory and results).

The Introduction virtually always starts by putting your experiment into a context. That could be a historical context (even if the 'history' happened last year), or it could be a context of related problems or industrial needs. This is also where you put citations to background work on the problem, so that a reader with less knowledge than your intended audience could read them to get up to speed. The Introduction often concludes with an outline of the rest of the paper. Note that the Introduction is **not** about what you did for your experiment. Nor is it about the theory one needs to understand your experiment. It is about the general concept that you are exploring, and how it fits into the bigger picture.

In both the Introduction and Theory sections, make sure that you stick to things that are reasonably relevant to your experiment. Everything in the Theory section should be useful in your data analysis. In the Introduction, making connections to things that you don't use is good, but don't let them overshadow the purpose of your paper.

If some topic requires an equation, that does **not** automatically mean that it belongs in the Theory section. The Theory section should be focused on the main topic of your paper. Most often, it will include several equations that will be inter-related. In contrast, if you wish to present an equation for equipment calibration, that should go in the Equipment section.

Similarly, not everything that you figured out necessarily belongs in the Results section. The Results section should be focused on the main topic of your paper. If some preliminary work was required to understand the equipment, then that might go into the Equipment section (if it helps the reader understand the equipment), or might not belong in your paper at all.

Most of the time you end with a Conclusion section. This is **not** where you put results; i.e., the section name does not mean “what I/we conclude based on the experiment.” Rather, the section title means “wrapping up (concluding) the paper.” The Conclusion summarizes the article; in that respect, it is similar to the Abstract. However, the Conclusion should be even more focused on your results; your procedure or technique should take a back seat (unless your primary result *is* an improved procedure). This different focus is enabled by the fact that you are allowed to assume that the rest of the paper has been read, which is not the case for the Abstract. The Conclusion may also include speculation based on your results, suggestions for future experiments, or suggestions for improving your current experiment.

Words

Don’t confuse ‘uncertainty’ and ‘error.’ There is a long tradition of confusing these words in physics, but there is no need for you to continue that tradition. Reserve the word ‘error’ for a mistake, or the actual disagreement between two sources of the same quantity.

Describing a number as ‘calculated’ is usually not helpful.

- There is often a temptation to try to use ‘calculated’ to distinguish between experimental and theoretical. However, both experiment and theory require calculation.
- You may want to distinguish between a value you obtained and a value that came from somewhere else. Instead of calling your value ‘calculated,’ use an adjective telling where the other value came from.
 - ‘Nominal’ means the value used to name something. The nominal value is expected to be close, but probably not exact. For example, “For the circuit with a 100 k Ω nominal resistance, the resistance was measured to be 98.3 k Ω .”
 - Use ‘commonly accepted’ for values that you could look up in a reference.
 - A fall back is to call the other value the ‘expected’ value. However, this is always less desirable than an adjective that expresses *why* you expect that value.

When referring to graphs, do not use the generic x and y to refer to the horizontal and vertical axis. You used the graph to show the relation between two physical quantities, which almost surely have a variable to represent them; use those variables instead.

Structure

For journals with the word “Letters” in their names (e.g., *Physical Review Letters*), the Body is not allowed to be subdivided into sections. These journals also usually have strict length limits (e.g., 4 journal pages for *PRL*).

For other journals, the Body usually is divided into sections with titles. The sections may or may not be numbered. Subsections are usually allowed, but they aren’t common.

Use paragraphs. Splitting the paper into sections will only give it a large-scale structure. On a smaller scale, you need to group together concepts.

A list of equipment is almost never appropriate. The focus should be on what you did, and the equipment should come up naturally in the course of doing that. (The only exception is if the paper

presents an experiment intended for others to repeat, as in an instructional laboratory. In that case, the reader might actually find a list useful.)

Notes commenting on the text (footnotes and endnotes) are rarely used, except for citations to other work (discussed under Reference List below).

One special issue for technical writing is how to show mathematical equations. Here are some rules, which you may recognize from textbooks:

1. Variables should be in italic typeface. Units and numbers are in normal (roman) typeface. Vectors are in bold typeface or with arrows above them, and matrices are in a bold sans-serif typeface (Arial font is an example of sans-serif). Use of an equation editor is strongly encouraged. That will automatically italicize variables, fine tune the spacing, and often prevent equations from being split by line breaks.
2. Explicitly define all variables when they are first used. This is true even when you think the context makes it painfully obvious. The definition should still be in the form of a grammatically correct sentence.
3. When a variable immediately follows its description, do **not** surround it with commas. For example, use “the initial velocity v_i is near zero”, **not** “the initial velocity, v_i , is near zero”.
4. Small equations may be placed in the text just like a word.
5. Large equations, important equations, and equations to which you need to refer later should be “displayed.” That means the text breaks off and the equation gets a line of its own. Although not required, it is good practice to number every displayed equation. Use a number in parentheses aligned to the right margin. You can then refer to the equations, as in “We now combine Eqs. 4, 5, and 8 to obtain...”
6. Regardless of whether an equation is displayed or not, the equation is treated as if it were grammatically part of the text. This is why you will often see punctuation, such as a comma or period, at the end of a displayed equation; the equation is *part* of the sentence. Normally the equation plays the role of a noun. Even though relation operators are often thought of as verbs (e.g., “ $A = B$ ” translates to “ A is equal to B ”), the relation operator may not supply the verb for the main clause of the sentence. For a subordinate clause, however, you sometimes can get away with the equation supplying the verb.

Example	Comment
Thus, $p = nkT$.	Not acceptable; no main clause verb.
The pressure p is thus given by $p = nkT$.	This is fine. The equation is a noun.
The pressure p is thus given by $p = nkT \quad (24)$	As above with the equation displayed. I like to put some space between the equation and the punctuation.
Thus, we find that $p = nkT$. The pressure is given by $p = nkT$, where $n = N / V$.	Examples of equations supplying the verb for subordinate clauses.

Figures and Tables

Structure

Unlike the other sections of this guide, I will describe the structure of Figures and Tables before the content. That is because the structure of Figures and Tables is very likely much different from what you are used to, and the structure controls the content to some extent.

The overriding fact about Figures and Tables is that they must “float.” When you submit an article to a journal, you control the content, but *they* control the layout. This means that you can’t count on a figure appearing immediately after a particular sentence. Maybe that sentence will be near the bottom of a page, and there won’t be enough room for the figure there.

The solution is to make figures and tables somewhat self-contained objects that can move, or “float,” to the best place on the page. For this reason, Figures and Tables are also referred to as “floating material.” Usually they float to the top of the page, but sometimes to the bottom. They must appear on the same page as the first time they are mentioned in the text, or as soon after that as possible. Note that in the above example, with the first reference to the figure coming near the bottom of the page, it is impossible for both the floating item and the first reference to it to appear on the same page.

When submitting a paper to a journal for review, a second organization is acceptable. Since Figures and Tables are going to be floating about anyhow, it is also accepted practice to put them all at the end of the article, one per page. You then let the journal worry about where they will actually appear.

(In fact, you will still occasionally see a structure where the rest of the paper goes first, followed by a separate page with the captions for all the Figures and Tables, followed by the Figures and Tables themselves. This format is handy for typesetters, but it is a pain to read because you have to keep flipping between the caption page and the figure itself.)

A Table is information arranged in columns or on a grid, with column labels and possibly row labels. Anything that is not a table is a Figure. This includes graphs, drawings, pictures, diagrams, photographs (rarely), etc.

Both Figures and Tables must have captions, which explain their content. Tables have their captions above the actual table, while Figures put their captions at the bottom. The captions always start with the word ‘Table’ or ‘Figure’ followed by a number, e.g., “Table 3: ...” or “Figure 4: ...”. This is, of course, how you refer to the Figure or Table in the text (e.g., “As we see from the entries in Table 3, ...”). Figures and Tables are numbered by separate sequences.

If someone else has already published just the figure you want to use, you may do so as long as you get the consent of the original author (and/or the journal it was published in, depending on the specific journal). If this is done, your caption must contain a citation to a reference for the original publication (see the Reference List section for details).

Because Figures have captions, it is rare to put titles on graphs. The caption serves that purpose. In fact, you should avoid having words in the figure at all. Instead use labels (A, B, C... for instance) that are explained in the caption. An exception is that you should still put labels on the axes of graphs, and words in graph legends when appropriate.

Often, you will want to have several pictures or graphs that are very closely related. It is perfectly acceptable to group them together into a single Figure with a single caption. Use (a), (b),

(c)... to label the parts. You can then use those labels to refer to the parts in the caption, and also in the Body text, e.g., “Comparing Fig. 8a to Fig. 8b, we see that...”.

Color is rarely used in Figures. It is costly to print in color, and usually the journal will require you to pay for the extra expense. We’re talking thousands of dollars here. So unless color is absolutely crucial to get the point across, don’t use it.

Content

Figures and Tables are supporting items. The reader should be able to understand the Body even if they can't see them. This doesn't mean, however, that the Body needs to be completely convincing without reference to them.

Floating material should never be included in the article unless they are referred to in the Body. If you don’t need to refer to it in the Body, what good is it doing?

Tables are used much less than you might think. The reason is that there is very often a graphical way to present the results that is easier to understand and/or more persuasive. If you show your results in a graphical way, then there is no need whatsoever to also give them in a table of data. At this level, the reader will trust you to have made the graph correctly.

The caption (for either Table or Figure) should very briefly (a few sentences) state what is being shown. It is OK to refer to things that are described in the Body. However, somebody who read your paper three weeks ago should be able to have their memory jogged *without* looking in the Body. The caption should also explain or define any abbreviations or symbols used in the Figure or Table itself.

Reference List

Content

You will always need to refer to other people’s work. (If you think you don’t, that means that you haven’t done enough background library research.) There are five possible sources for anything you write in your paper:

1. You observed it directly yourself.
2. Everyone in your intended audience knows it from previous experience.
3. It follows logically or mathematically from other things you have said.
4. You made it up (e.g., you came up with a hypothesis to explain the observations).
5. You learned it from someone/somewhere else.

Anything that falls into category 5 **must** be referenced: you need to indicate the source. Note that the boundary between 2 and 5 is not very well defined. As discussed in the Body section, you must determine who your audience is, and that will imply where the boundary falls.

Why should the reader believe what you write? The first four sources are justified either within the paper (1 and 4), or by a vast amount of common experience (2 and 3). When you reference a source, the reader needs some reason to believe that the source is reliable. In a journal article this is almost always accomplished by referencing *previously published* work. This justifies the source in two ways. First, the previously published work has been through a peer review process, so several knowledgeable people besides the original author have given it the stamp of approval. Second, if your reader doesn't believe it, they can go back to the previous paper to get more details. (There is also the added benefit that if the reader doesn't understand it, they can go back to the previous paper to learn more.)

On rare occasion, you will see a citation something like “Private communication with Dr. I. M. Smart.” This lacks the first justification, but at least the reader can contact Dr. Smart to ask questions. Obviously, you would need to get Dr. Smart’s permission before you used a reference like this. This reference type is only appropriate when you are writing about other people’s experimental observations, but those observations haven’t been published. This is **not** appropriate when you learned something from someone else; if you learned it, then you should be able to explain it in your paper. In general, the person must be the sole available source of the information.

I have never seen a World Wide Web page used as a supporting reference in a journal article. In general, not only are web pages not peer reviewed, but they may disappear at any time (we’ve all had the experience of following broken links), leaving the reader with nothing to fall back on. So, while the web is a great place to find information to get started on a topic, it’s not authoritative and it’s not permanent.

On the other hand, it has become quite common to provide information linking to where information can be found on the web. The distinction here main seem subtle. Being on the internet is not a bad thing itself; the problem occurs when there is nothing to back up that web presence.

Structure

References to other work are handled by endnotes. (Endnotes are like footnotes, except that they are collected at the very end of the paper.) A number (called the “citation” according to the Physical Review Style and Notation Guide) is put in the Body text where you make the reference, and in the Reference List section you associate the number with something previously published (with a “bibliographic entry,” also called a “reference”). The citations must be numbered in the order that they first occur in the Body, and the references are listed in order of citation number.

Note that a Reference List is different from a Bibliography, which is frequently found at the end of papers in the humanities. A Bibliography is also a list of supporting works, but it makes no attempt to link each supporting work to specific points in your paper. There are no citation numbers in a Bibliography. Bibliographies may be appropriate in long scientific writing (books or review articles), but they are never used in journal articles.

Different journals require different formats for the citation numbers. Currently, the most fashionable method is to enclose the numbers in square brackets. Usually, the citation numbers will appear in such a way that the sentence would make perfect sense even if the citations were removed, e.g., “This effect has been studied thousands of times [2, 4-10].” Or “As pointed out by Smith [8], the deely has a bopper.” Occasionally, the number will be a necessary part of the sentence, in which case you prefix it with “Ref.”, e.g., “It was shown in Ref. [14] that this is utterly negligible.”

There are occasions where the placement of the citation is inconvenient. For instance, what if you want to give a citation for a displayed equation? I believe that the best solution is to put the citation at the end of the text that immediately precedes the equation. If you want to put a citation in a caption, or maybe on a specific entry in a table; what number should you use? Journals have various arcane rules for this situation.

In the Reference List, we must worry about the format of the references. They always start with the citation number in the same format that is used in the Body. After that, the format varies between journals, but here are some typical formats. For a reference to an article in a journal, give the authors’ name(s), the name of the journal, the number of the volume that the article is in, the page number on which the article starts, and in parentheses the publication year. For example,

[5] James G. McLean, Peter Kruse, and Andrew C. Kummel, *Surface Science* v. 424, p. 206 (1999).

Note that the article title most often is **not** included. For a book, give the authors' name(s), the book title, in parentheses the publisher, the publisher's address, and the publication year, and after the parentheses the page range of interest. For example,

[8] N. W. Ashcroft and N. D. Mermin, *Solid State Physics* (Holt, Rinehart, and Winston, New York, 1976), p. 389.

For other references, a similar amount of info should be given. It should be enough for an interested reader to locate the cited material without reading a lot of extraneous material.