

A researcher who has contributed to the research must be given an opportunity to be included as an author, but authors should not be listed without their permission.

A related issue is of author order, since many readers will assume that the first author is the main contributor. A researcher who is clearly the main contributor should always be listed first—don't believe Alfred Aaby when he tells you that alphabetic ordering is the norm. Where there is no obvious first author, possible approaches to ordering include alphabetical or reverse alphabetical, perhaps with an explanatory footnote, or a reversal or rotation of the order used on a previous paper by the same authors. Many supervisors choose to put their student co-authors first.

Grammar

In this book I have avoided giving advice on grammar, because the clarity of writing largely depends on whether it conforms to accepted usage. One aspect of grammar is, however, worth considering: that some people like to use traditional grammar to criticize other people's text, based on rules such as *don't split infinitives* or *don't begin a sentence with "and" or "but"*. I dislike this attitude to writing: grammatical rules should be observed, but not at the cost of clarity or meaning. However, be aware that an overdose of grammatical errors annoys some readers.

Beauty

Authors of style guides like to apply artistic judgements to text. This does not mean that scientific writing should be judged as literary prose, indeed such prose would be quite inappropriate. But we read that text should be crystalline, transparent, and have good rhythm and cadence; and one should dislike stuffiness, softness, stodge, sludge, and sagging or soggy sentences.

How useful such judgements are to most authors is not clear. Doubtless, well-crafted text is a pleasure to read, ill-written text can be hard going, and good rhythm in text helps us to parse. But awareness of beauty in text does not, I think, help us to attain it, nor is it evident that, to a poor writer, the terminology of beauty in text is meaningful. It is sufficient to aim to achieve simplicity and clarity.

3 Writing style: specifics

Those complicated sentences seemed to him very pearls . . .

"The reason for the unreason with which you treat my reason, so weakens my reason that with reason I complain of your beauty" . . . These writings drove the poor knight out of his wits.

Cervantes
Don Quixote

Underneath the knocker there was a notice that said:

PLES RING IF AN RNSER IS REQIRD

Underneath the bell-pull there was a notice that said:

PLEZ CNOKE IF AN RNSR IS NOT REQID

These notices had been written by Christopher Robin, who was the only one in the forest who could spell.

A. A. Milne
Winnie the Pooh

Titles and headings

Titles of articles and sections should be concise and informative, use specific rather than general terms, and accurately describe the content. Complicated titles with long words are hard to swallow.

- × A New Signature File Scheme based on Multiple-Block Descriptor Files for Indexing Very Large Data Bases
- ✓ Signature File Indexes Based on Multiple-Block Descriptor Files
- × An Investigation of the Effectiveness of Extensions to Standard Ranking Techniques for Large Text Collections
- ✓ Extensions to Ranking Techniques for Large Text Collections

Don't make the title so short that it is contentless. "Limited-Memory Huffman Coding for Databases of Textual and Numeric Data" is awkward, but it is superior to "Huffman Coding for Databases", which is far too general.

Accuracy is more important than catchiness—"Strong Modes can Change the World!" is excessive, not to mention uninformative. The more interesting the title, however, the more likely that the text underneath it will be read. The title is the only part of your paper that the vast majority of people will see; if the title does not reflect the paper's contents, the paper will not be read by the right audience.

Titles and section headings do not have to be complete sentences; indeed, such titles can look rather odd.

- × Duplication of Data Leads to Reduction in Network Traffic
- ✓ Duplicating Data to Reduce Network Traffic

Section headings should reflect the article's logical structure. If a section is headed "Lists and Trees" and the first subsection is "Lists", another should be "Trees"; don't use, say, "Other Data Structures". If a section is headed "Index Organizations" the subsection heading should be "B-trees" rather than "B-tree indexes".

An article (or thesis chapter) usually consists of sections and possibly subsections. There is rarely any need to break subsections into sub-subsections. Don't break text into small blocks; three headings on a page is almost certainly too many. But beware of having too few sections,

because it is hard to continue the logical flow of a section over more than a few pages.

Headings may or may not be numbered. My preference is to use only two levels of headings, major and minor, and to only number major headings. If all headings are unnumbered, make sure that major and minor headings are clearly distinguished by font, size, or placement.

The opening paragraphs

The opening paragraphs can set the reader's attitude to the whole paper, so begin well. All of a document should be created and edited with care, but take the most care with the opening, to create the best possible impression. The abstract should be written especially well, without an unnecessary word, and the opening sentence should be direct and straightforward.

- × Trees, especially binary trees, are often applied—indeed indiscriminately applied—to management of dictionaries.
- ✓ Dictionaries are often managed by a data structure such as a tree, but trees are not always the best choice for this application.

The following example of how not to begin is the first sentence of a published paper.

- × This paper does not describe a general algorithm for transactions.

Only later does the reader discover than the paper describes an algorithm for a special case.

- ✓ General-purpose transaction algorithms guarantee freedom from deadlock but can be inefficient. In this paper we describe a new transaction algorithm that is particularly efficient for a special case, the class of linear queries.

The first paragraphs should be intelligible to any likely reader; save technicalities for later on, so that readers who can't understand the details of your paper are still able to understand your results and the importance of your work. That is, describe what you have done without the details of how it was done.

Starting an abstract or introduction with "This paper concerns ..." or "In this paper ..." often means that results are going to be stated out of context.

- × In this paper we describe a new programming language with matrix manipulation operators.
- ✓ Most numerical computation is dedicated to manipulation of matrices, but matrix operations are difficult to implement efficiently in current high-level programming languages. In this paper we describe a new programming language with matrix manipulation operators.

The second version describes the context of the article's contribution.

A typical organization for the introduction of an article would be to use the first paragraphs to describe the context. It is these paragraphs that convince the reader that the article is likely to be interesting. The opening sentences should clearly indicate the topic.

- × Underutilization of main memory impairs the performance of operating systems.
- ✓ Operating systems are traditionally designed to use the least possible amount of main memory, but such design impairs their performance.

The second version is better for several reasons. It is clear; it states the context, which can be paraphrased as *operating systems don't use much memory*; and, in contrast to the first version, it is positive.

Take care to distinguish description of existing knowledge from the description of the paper's contribution.

- × Many user interfaces are confusing and poorly arranged. Interfaces are superior if developed according to rigorous principles.
- ✓ Many user interfaces are confusing and poorly arranged. We demonstrate that interfaces are superior if developed according to rigorous principles.

Don't write the introduction as if it flows on from the abstract, which is a summary of a paper rather than its opening. The paper should be complete even with the abstract removed.

Variation

Diversity—in organization, structure, length of sentences and paragraphs, and choice of words—is a useful device for keeping the reader's attention.

- × The system of rational numbers is incomplete. This was discovered 2000 years ago by the Greeks. The problem arises with squares whose sides are of unit length. The length of the diagonals of these squares is irrational. This discovery was a serious blow to the Greek mathematicians.
- ✓ The Greeks discovered 2000 years ago that the system of rational numbers is incomplete. The problem is that some quantities, such as the length of the diagonal of a square with unit sides, are irrational. This discovery was a serious blow to the Greek mathematicians.

Note how, in the second version, the final statement is more effective although it hasn't been changed.

Paragraphing

A paragraph usually consists of discussion on a single topic or issue. In a well-written paper, the gist if not the argument is often captured in the first sentence of each paragraph, with the remainder of the paragraph used for amplification or example. Every sentence in a paragraph should be related to the topic announced in the opening.

Long paragraphs can be an indication that several lines of argument have not been sufficiently disentangled by the author. Moreover, readers tend to pay more attention to the start and end of each paragraph and less to the body. If a long paragraph can be broken, break it. Variation in paragraph length makes the page less dull in appearance, however, so don't chop text into paragraphs of uniform size.

Contextual information can be forgotten between paragraphs, and references between paragraphs can be difficult to follow. For example, if a paragraph discusses a fast sorting algorithm, the next paragraph should not begin "This algorithm ..." but rather "The fast sorting algorithm ..."; if one paragraph refers to Harvey, the next should not refer to "his" but rather "Harvey's". Link paragraphs by reuse of key words or phrases, and by using expressions that connect the content of one paragraph to that of the next.

The use of formatted lists as an occasional alternative to paragraphs is common. Lists are useful for the following reasons.

- They highlight each main point clearly.

- The context remains obvious, whereas in a long list of points made in a paragraph it is hard to tell whether the later points are part of the original issue or belong to some subsequent discussion.
- An individual point can be considered in detail without confusing the main thread of narrative.
- They are easy to refer to; for example, as a checklist of the necessary properties of an algorithm.

List points can be numbered, named, or tagged. Use numbers only when ordering is important. If it is necessary to refer to an individual point, use numbers or names. Otherwise use tags, as in the list above. Acceptable tags are bullets and dashes; fancy symbols such as \leftrightarrow or graphic icons look childish.

A disadvantage of lists is that they highlight rather too well: a list of trivia can be more attention-getting than a paragraph of important information. Reserve the use of lists for material that is both important and in need of enumeration.

Sentence structure

Sentences should have simple structure, which usually means that they will be no more than a line or two. Don't say too much all at once.⁴

⁴The following quote is a single sentence from a version of the standard lease agreement of the Real Estate Institute of Victoria, Australia. It is 477 words long, but the punctuation amounts to only three pairs of parentheses, one comma, and one stop. This clause is an example of “the fine print”—for example, the holder of a lease containing this clause has agreed not to take action if, in circumstances such as failure to pay rent, assaulted by the property's owner.

If the Lessee shall commit a breach or fails to observe or perform any of the covenants contained or implied in the Lease and on his part to be observed and performed or fails to pay the rent reserved as provided herein (whether expressly demanded or not) or if the Lessee or other person or persons in whom for the time being the term hereby created shall be vested, shall be found guilty of any indictable offence or felony or shall commit any act of bankruptcy or become bankrupt or make any assignment for the benefit of his her or their creditors or enter into an agreement or make any arrangement with his her or their creditors for liquidation of his her or their debts by composition or otherwise or being a company if proceedings shall be taken to wind up the same either voluntarily or compulsorily under any Act or Acts relating to Companies (except for the purposes of reconstruction or amalgamation) then

- × When the kernel process takes over, that is when in the default state, the time that is required for the kernel to deliver a message from a sending application process to another application process and to recompute the importance levels of these two application processes to determine which one has the higher priority is assumed to be randomly distributed with a constant service rate R .
- ✓ When the kernel process takes over, one of its activities is to deliver a message from a sending application process to another application process, and to then recompute the importance levels of these two application processes to determine which has the higher priority. The time required for this activity is assumed to be randomly distributed with a constant service rate R .

That the kernel process is the default state is irrelevant here, and should have been explained elsewhere.

This example also illustrates the consequence of having too many words between related phrases. The original version said that “the time that is required for *something* is assumed to be ...”, where *something*

and in any of the said cases the Lessor notwithstanding the waiver by the Lessor of any previous breach or default by the Lessee or the failure of the Lessor to have taken advantage of any previous breach or default at any time thereafter (in addition to its other power) may forthwith re-enter either by himself or by his agent upon the Premises or any part thereof in the name of the whole and the same have again repossess and enjoy as in their first and former estate and for that purpose may break open any inner or outer doorfastening or other obstruction to the Premises and forcibly eject and put out the Lessee or as permitted assigns any transferees and any other persons therefrom and any furniture property and other things found therein respectively without being liable for trespass assault or any other proceedings whatsoever for so doing but with liberty to plead the leave and licence which is hereby granted in bar of any such action or proceedings if any such be brought or otherwise and upon such re-entry this Lease and the said term shall absolutely determine but without prejudice to the right of action of the Lessor in respect of any antecedent breach of any of the Lessee's covenants herein contained provided that such right of re-entry for any breach of any covenant term agreement stipulation or condition herein contained or implied to which Section 146 of the Property Law Act 1958 extends shall not be exercisable unless and until the expiration of fourteen days after the Lessor has served on the Lessee the Notice required by Sub-section(1) of the said Section 146 specifying the particular breach complained of and if the breach is capable of remedy requiring the Lessee to remedy the breach and make reasonable compensation in money to the satisfaction of the Lessor for the breach.

was 34 words long. The main reason that the revision is clearer is that *something* has been reduced to two words; the structure of the sentence is much easier to see.

It is likewise helpful to avoid nested sentences, that is, information embedded within a sentence that is not part of its main statement.

- × In the first stage, the backtracking tokenizer with a two-element retry buffer, errors, including illegal adjacencies as well as unrecognized tokens, are stored on an error stack for collation into a complete report.

First, this is poor because crucial words are missing; the beginning should read “In the first stage, which is the backtracking tokenizer . . .”. Second, the main information—how errors are handled—is intermixed with definitions. Nested content, particularly if in parentheses, should be omitted. If it really is required then put it in a separate sentence.

- ✓ The first stage is the backtracking tokenizer with a two-element retry buffer. In this stage possible errors include illegal adjacencies as well as unrecognized tokens; when detected, errors are stored on a stack for collation into a complete report.

Watch out for fractured “if” expressions.

- × If the machine is lightly loaded then speed is acceptable whenever the data is on local disks.
- ✓ If the machine is lightly loaded and data is on local disks then speed is acceptable.
- ✓ Speed is acceptable when the machine is lightly loaded and data is on local disks.

The first version is poor because the conditions of the “if” have been separated by the consequent.

It is easy to construct long, winding sentences by, for example, stating a principle, then qualifying it—a habit that is not necessarily bad, but does often lead to poor sentence structure—then explaining the qualification, the circumstances in which it applies, and in effect allowing the sentence to continue to another topic, such as the ideas underlying the principle, cases in which the qualification does or does not apply, or material which no longer belongs in the sentence at all; a property that is arguably true of most of this sentence, which should definitely be revised.

Sometimes longer sentences can be divided by, say, simply replacing an “and” or a semi-colon with a period. If there is no particular reason to join two sentences, keep them separate.

Beware of misplaced modifiers.

- × We collated the responses from the users, which were usually short, into the following table.
- ✓ The users’ responses, most of which were short, were collated into the following table.

Double negatives are difficult to parse and are often ambiguous.

- × There do not seem to be any reasons not to adopt the new approach.

The impression here is of condemnation—*we don’t like the new approach but we’re not sure why*—but praise was intended; the quote is from a paper advocating the new approach. This is another example of the academic tendency to overqualification. The revision “There is no reason not to adopt the new approach” is punchier, but still negative. It is difficult to suggest further improvement with the same meaning, because the meaning was probably unintended; the following better reflects the authors’ aims.

- ✓ The new approach is at least as good as the old and should be adopted.

Sing-song phrases are distracting, as are rhymes and alliteration.

- × We propose that the principal procedure of proof be use of primary predicates.
- × Semantics and phonetics are combined by heuristics to give a mix that is new for computational linguistics.

Repetition and parallelism

Text that consists of the same form of sentence used again and again is monotonous. Watch out for sequences of sentences beginning with “however”, “moreover”, “therefore”, “hence”, “thus”, “and”, “but”, “then”,

“so”, “nevertheless”, or “nonetheless”. Likewise, don’t overuse the pattern “First, ... Second, ... Last, ...”

Complementary concepts should be explained as parallels, or the reader will have difficulty seeing how the concepts relate.

- × In SIMD, the same instructions are applied simultaneously to multiple data sets, whereas in MIMD different data sets are processed with different instructions.
- ✓ In SIMD, multiple data sets are processed simultaneously by the same instructions, whereas in MIMD multiple data sets are processed simultaneously by different instructions.

Parallels can be based on antonyms.

- × Access is fast, but at the expense of slow update.
- ✓ Access is fast but update is slow.

Lack of parallel structure can result in ambiguity.

- × The performance gains are the result of tuning the low-level code used for data access and improved interface design.
- ✓ The performance gains are the result of tuning the low-level code used for data access and of improved interface design.

This can be improved again: it is kinder to the reader to move the longer clauses in a list to the end.

- ✓ The performance gains are the result of improved interface design and of tuning the low-level code used for data access.

There are some standard forms of parallel. The phrase “on the one hand” should have a matching “on the other hand”. A sentence beginning “One ...” suggests that a sentence beginning “Another ...” is imminent. If you flag a point with “First” then every following point should have a similar flag, such as “Second”, “Next”, or “Last”.

Parallel structures should be used in lists.

- × To achieve good performance there should be sufficient memory, parallel disk arrays should be used, and caching.

The syntax can be fixed by adding “should be used” at the end but the result is clumsy. A complete revision is preferable.

- ✓ Achievement of good performance requires sufficient memory, parallel disk arrays, and caching.

Direct statements

Avoid excessive use of indirect statements (also known as passive voice), particularly descriptions of actions that don’t include any indication of who or what performs the actions.

- × The following theorem can now be proved.
- ✓ We can now prove the following theorem.

The direct style (or active voice) is often less stilted and easier to read.

Another unpleasant indirect style is the artificial use of verbs like “perform” or “utilize”, in the false belief that such writing is more precise or scientific.

- × Tree structures can be utilized for dynamic storage of terms.
- ✓ Terms can be stored in dynamic tree structures.
- × Local packet transmission was performed to test error rates.
- ✓ Error rates were tested by local packet transmission.

Other words often used in this way include “achieved”, “carried out”, “conducted”, “done”, “occurred”, and “effected”.

Change of voice sometimes changes meaning and often changes emphasis. If passive voice is necessary, use it. Complete absence of active voice is unpleasant, but that does not mean that all use of passive voice is poor.

Use of “we” is valuable when trying to distinguish between the contribution made in an article and existing results in a field, especially in an abstract or introduction. For example, in “it is shown that stable graphs are closed” the reader may have difficulty deciding who is doing the showing, and in “it was hypothesized that ...” the reader will be unsure whether the hypothesis was posed in this article or another one. Use of “we” can also allow some kinds of statements to be made more simply—consider “we show” versus “in this paper it is shown that”. And “we” is preferable to pretentious expressions such as “the authors”.

Some authors use phrases such as “this paper shows ...” and “this section argues ...” These phrases, with their implication that the paper is sentient, should not be used.

In some cases the use of “we” is wrong.