

Effective written documents

WRITTEN DOCUMENTS are valuable references. Although documents may well remain unread, they can also be reread many times or be read at the most convenient moment for each reader. In a context where many media compete for the attention of the readers, documents must allow each reader to assimilate effortlessly just what he or she needs to know. Their essential quality is therefore to be readable, that is, clear, accurate, and concise.

An effective way to go from scratch to a complete document proceeds iteratively in five steps. First, plan your document: gather your thoughts about the writing task. Second, design it: organize your material into a clear, well-structured hierarchy. Third, draft it: turn your ideas into paragraphs and sentences. Fourth, format it: take care of the visual component. Finally, revise it: test it and improve it iteratively until it is optimal, that is, the most effective it can be, given certain constraints.

Fundamentals

- The name of the game
- The three laws of communication
- A thousand words, a thousand pictures
- Chains and magical numbers
- Trees, maps, and theorems

Effective written documents

Planning the document

- Designing the document
- Drafting the document
- Formatting the document
- Revising the document

- Defining your purpose
- Identifying your audience(s)
- Selecting your content

Effective oral presentations

- Planning the presentation
- Designing the presentation
- Creating the slides
- Delivering the presentation
- Answering questions

Effective graphical displays

- Understanding pictures
- Planning the graph
- Designing the graph
- Constructing the graph
- Drafting the caption

Applications

- Effective instructions
- Effective electronic mail
- Effective Web sites
- Effective meeting reports
- Effective scientific posters

Planning the document

PLANNING IS THINKING about what we are to create. It thus means distancing ourselves from the situation and making sure we know everything we need to know in order to write an effective document. It is best carried out looking out our window, not staring at a text-processor one.

One systematic way to think about the document is to answer a set of questions about it. An effective and mnemonic set, unsurprisingly, consists of the five basic interrogative words starting with the letter *w*: *why*, *who*, *what*, *when*, and *where*. (The sixth interrogative *how*, conveniently not starting with *w*, is not an initial question: it points to the strategy, determined later on the basis of the answers to the other five questions.)

The five planning questions

Why	Purpose
Who	Audience
What	Content
When	Time constraints
Where	Space constraints

The five planning questions are not equally hard to answer. The constraints of time and space (*when* and *where*), critical as they are, are often easy to identify. Harder to determine are the purpose (*why*), the audience (*who*), and the content needed to reach this purpose with this audience (*what*). These three are tightly linked: the purpose expresses what the audience should (be able to) do after reading the document's content.

Planning need not take long but should be done with care, because it has a far-reaching impact. Poor planning results in unnecessary iterations before the writing process converges on the optimal document—when it has time to converge at all before the resources are up—and may result in writer's block (the inability to proceed with writing). The less time we have in total, the fewer iterations we can afford, and consequently the more time we should devote, proportionally, to planning.

When unclear initial ideas about the document to be created make planning hard, a possible strategy is to start drafting at once, in order to clarify our ideas. Sooner or later, though, we will need to go through the planning and designing steps, and to adapt our draft accordingly—an unavoidable iteration.

Common shortcomings

Should the purpose always be identified first?

Because it plays a deciding role, the purpose is often best identified first. Still, sometimes you know you need to react toward someone before deciding exactly what you want to reach, in which case you have identified your audience before determining your document's purpose.

I often have difficulties identifying my purpose. Any suggestions to make this process easier?

If you are writing a document at the request of someone (client, boss, etc.), ask this person what he or she identifies as the purpose for it. Placing the document in a broader framework helps, too. Is it, for example, part of a project? If yes, how is the document contributing to it?

When the purpose seems difficult to identify, an effective approach is simply to talk about it with someone else, such as a direct colleague. Talking out loud obliges one to turn thoughts into words, making them easier to work with. Your colleague will probably ask you questions, too, which will catalyze your thought process, and he or she might actually suggest a purpose.

What if I am convinced that the document I have been asked to write serves no purpose?

If you are deeply convinced that the document will make no difference, today or in the future, take up the issue with whoever has requested it, explaining that you can hardly write effectively without a clear idea of the document's purpose. (If the document is a contractual requirement, but you suspect that no one will ever read it, do write it, but do not spend much time on it.)

A communication purpose is more than a reason for writing: it gives a direction, suggests a strategy, provides a metric for success by enabling authors to visualize the outcome of the communication act.

A first typical shortcoming, then, is simply writing a document without having a clear purpose in mind. Indeed, many authors do have a reason for writing (*because my boss told me to report my research*) but do not visualize any potential or desired action on the part of the audience. The reason for writing can be confused with personal motivation, such as *I want to make an original contribution to science*, which can of course be a worthwhile pursuit in itself but does not constitute a communication purpose, because it does not involve the audience in any way. It is perhaps a motivation for carrying out the work but not for reporting on the work thus carried out.

Rational minds who do understand that a purpose for writing involves the audience may still fall short of identifying an observable outcome as part of it. They seem content with reaching the second step (*I want my audience to understand how it works*) and do not anticipate any action for the audience. The question these people might ask themselves is, *why do I want them to understand how it works?* Or perhaps, *how will I see that they understood?* That is, *what will this understanding in fact achieve?*

Finally, some otherwise clearly identified purposes may lead to ineffective communication strategies. Such is the case of purposes focusing on what are symptoms of excellence rather than sources of it (*I want my audience to be entertained by my text*) or expressing personal and often hidden agendas (*I want my boss to be impressed with my writing*). While these may well reflect legitimate aspirations, they offer no guarantee of helping authors reach a more fundamental professional purpose, namely get the audience to (be able to) act upon messages.

Defining your purpose

PURPOSES ARE THE ONLY METRIC against which to gauge the effectiveness of documents. A document is effective to the extent that it reaches its purpose—period. Accordingly, we can hardly determine the most suitable writing strategy without first having identified a purpose for our document. This purpose simply captures the change that the document is to produce: what must be different after it has been read? All communication is purpose-driven and, unsurprisingly, this purpose is defined in relation to both audience and author.

Effective purposes focus on the audience. They identify, not what the author should achieve, but rather what the readers should (be able to) do as a result of reading the document. This potential action may be understood broadly as anything involving the audience actively, even if only intellectually. Still, expressing it in observable terms helps envision a strategy: it is thus more useful to say we want the audience to *be able to disassemble the device* or to *sign the contract* rather than to *understand* or to *agree*, which cannot be observed as such. Audience-centered purposes moreover anticipate the purpose that readers bring to the communication: why will they read?

Although audience-centered, purposes acquire full meaning only in relation to an author. When identifying your purpose, identify your role in the communication: answer the question *who am I in relation to both my purpose and my audience?* or *why am I the person writing this document?* The audience will likely want to know, so be prepared to clarify it for them.

The purpose is implicit in the idea of getting messages across, defined as getting our audience to pay attention, understand, and (be able to) act upon our messages: it is the ultimate step. The first two steps, by contrast, are mere means to this end. For example, getting our audience to read a whole document is not a purpose in itself: if they (are able to) do what we want them to do without reading every word, so much the better.

Technical terms are not jargon

Since technical terms are so useful, how about using them for nonspecialists, too, but adding a glossary as an appendix to the document?

Readers can assimilate only so many new terms at a time, so beware of introducing too many. A glossary can certainly help, but it is best used redundantly: the (few) new terms must still be introduced in the text. If you include a glossary, let readers know about it, lest they never see it.

Is every reader really either a specialist or a nonspecialist?

As with all dichotomies, dividing the readers in just two groups is simplistic, and specialism is obviously a continuum that can be divided into as many groups as called for. For example, you might want to place your direct supervisor midway between specialists and nonspecialists: probably not quite as specialized as you are while still more specialized than most readers. Thinking in terms of just two groups, however, is a useful basis for reasoning about audiences.

For specialists, must I include conclusions? They can draw these themselves, right?

Readers as specialized as you are should be able to draw conclusions as well as you do, indeed. Of course, just because they can does not mean they will: interpretation is seldom trivial. Also, few readers are in as good a position as you are to interpret your own findings. It is therefore more constructive to include your conclusion and let the readers specialized enough to have an opinion on this conclusion disagree with it if they choose to. At least, you provided a basis for their thinking—and perhaps for a dialogue.

Technical terms are not the same thing as jargon. Defined in technical dictionaries, they are meant to make the document clear, accurate, and concise for as many readers as possible (within specialists). Jargon, by contrast, always reduces the readership to a “chosen few.” It can happen purely by accident, such as when new employees unknowingly pick up company-specific terms and use these elsewhere. More often, though, it is caused by careless writing or, worse, by a desire to impress rather than help.

The 040 has a built-in 882.




The Motorola processor 68040 has a built-in floating-point unit 68882.

Technical terms are explicit, while jargon is cryptic. In the sentence *The 040 has a built-in 882*, readers who do not know what an 040 is would probably not know where to look for an explanation either. In the revised sentence, less specialized readers might of course not understand *floating-point unit*, but at least they have an idea of what the author is talking about (part of a processor) and they can also look for a definition in a technical dictionary.

Unexplained abbreviations easily become jargon: they needlessly exclude those readers who would understand the spelled-out form (including those who used to know this form and can use a reminder). Moreover, even a well-known acronym may expand to distinct meanings. Does *IP* stand for *intellectual property* or *Internet protocol*? Or *intraoperative*? Despite the presence of contextual clues, readers who readily recognize one meaningful expansion may not immediately think of the alternative ones (assuming that they are even aware of alternatives) and may therefore be misled, even if temporarily.

Identifying your audience(s)



AUDIENCES ARE MULTIPLE, for each reader is unique. Still, readers can usefully be classified in broad categories on the basis of their proximity both to the subject matter (the *content*) and to the overall writing situation (the *context*). The challenge lies not so much in addressing a given category as in dealing with a mixed audience. The strategy is twofold: identify which part of the audience matters for your purpose and, for written documents, layer the information presented, so each reader can access only what is relevant to him or her.

Although knowledge forms a continuum, readers can usually be placed in one of two groups: those who already know much about the subject matter (*specialists*) and those who know little about it (*nonspecialists*). Clearly, specialism is relative: everyone is a specialist on some subjects and a nonspecialist on others. Moreover, even a group of all specialists could be subdivided into more specialized and less specialized readers. As with any structure, the two groups are a view of the mind.

Specialists and nonspecialists differ in the type of information they are primarily interested in. Specialists want more detail: they can understand the technical aspects, can often use these in their own work, and require them anyway to be convinced. Nonspecialists need more basic information to bridge the gap between what they know and what the document discusses: more background at the beginning, to understand the need for and importance of the work; more interpretation at the end, to understand the relevance and implications of the findings.

Specialists and nonspecialists differ as well in the vocabulary they master: specialists readily understand technical terms; nonspecialists do not. Technical terms, whenever appropriate, are effective: they are both precise and concise, and they help present the author as a specialist, thus enhancing credibility. For nonspecialists, they should be replaced by paraphrases or, if they are particularly desirable, be carefully introduced.

Keeping all readers in mind

Are primary readers specialists and secondary readers nonspecialists?

The two audience typologies are independent; that is, primary readers can be nonspecialists and secondary ones specialists. For example, most executives are not (or no longer) specialists of all that they read about as primary readers. Conversely, specialists (re)reading a document in the future are definitely secondary readers.

By definition, I am writing for primary readers, so why should I worry about secondary ones?

Above all, worry about the readers who matter for you to reach your purpose. A quick reply sent by electronic mail may acceptably include little or no context, because secondary readers are generally few and probably unimportant. In contrast, a document that is worth archiving is one for which secondary readers do matter. Contextual information for secondary readers is usually a useful reminder for primary ones, too. In general, then, documents can usefully address both primary and secondary readers.

To address secondary readers, how about simply referring to a previous document?

Referring to a previous document as a sole way of providing context is hardly reader-friendly: it obliges readers to find this document first and to search through it for whatever context they might need to understand the document they actually want to read—an ill-defined task. Most readers will therefore not bother to do so, hence their understanding will be suboptimal. Such a reference is best provided redundantly, for those interested in more context or history.

Writing for a mixed audience is always a challenge: we must give to the secondary readers information that we assume the primary readers know already. How to do so, yet keep these readers interested? The solution, conceptually, is simple: just ensure that each sentence makes an interesting statement, one that is new to all readers—even if it includes information that is new to secondary readers only. This approach is illustrated in the example below, in which a few words of clarification go a long way.

We worked with IR.


A statement involving an unexplained abbreviation is cryptic for secondary readers who do not know what *IR* is and are left to wonder about its nature: is it a chemical, a computer code, "infrared," etc.?

We worked with IR. IR stands for Information Resources and is a new department.

Clarifying *IR* in an independent sentence, however, may be regarded as patronizing by primary readers. These then exclaim, *Hey, I know that!* when reading the second sentence, which tells them nothing new.

We worked with the recently launched Information Resources (IR) department.

A sentence that is appropriate for all readers states something new to all, while including information that is new to some and a useful reminder to others. Secondary readers thus learn about *IR*. Presumably, primary readers know what *IR* is but may not have in the front of their mind that this department was "recently launched." While not telling them anything they do not know, these two words will put them in a certain frame of mind for their further reading.



Secondary readers

Are reading far from here or in the future from now
Require context to be able to comprehend the issue

Quoting the original information when replying to electronic mail is one example of context used as reminder for primary readers while informing secondary ones. Such quotes are best kept short (the minimum needed to remind readers of the point addressed).

Besides proximity to the subject matter, readers can be sorted on the basis of their proximity to the situation (the context), both in space and in time. Those close to the *here and now* are *primary readers*; any other ones are *secondary readers*. Clearly again, everyone is a primary reader on some matters and a secondary one on others, depending on their situation.

Primary readers are usually well-defined: they are the people we have in mind when we are writing the document. They are the ones expected to (be able to) act about the content of it and are thus listed under *to* in the header of short documents or electronic mail. Typical primary readers are our direct boss or colleagues, or our contact person at a client organization.

Secondary readers, on the other hand, are often ill-defined: they are the people who are sent a document for information (as listed under *copy to* or *cc*), but also those who obtain it via unpredictable routes, for example from a primary reader or from a library or archive. Typical secondary readers are distant colleagues, higher hierarchy, other people at a client's, or simply whoever will read the document in a few months or a few years (including, very often, the authors themselves).

To be understandable to both primary and secondary readers, documents must include context. Primary readers, who are close to the *here and now*, normally know the context but may not be thinking of it when starting to read; they benefit from a reminder, bringing it back to the front of their mind. Secondary readers, by contrast, are not (or no longer) aware of the context, so they need to find it in the document itself.

Context upfront not only broadens a document's readership, it is also essential for "getting our audience to pay attention": well phrased, it (re)establishes the importance of the subject and usefully prepares primary and secondary readers alike for the story told by putting them in a desired frame of mind.

*Why must I be selective in a written document?
Readers can select what they want to read.*

You should admittedly be even more selective for an oral presentation than for a document. Still, unnecessary material dilutes the message. Moreover, lengthy documents are intimidating enough that their reading may be postponed (*I will read this big report... when I have time*), even when only a small part of them, such as the executive summary, must in fact be read. Should you nonetheless include much material, do give your readers the means to be selective by providing adequate navigational features. Just because readers need not read everything does not mean they know what they must read.

*I am writing an article for a scientific journal:
is my audience not made of specialists only?*

Readers of scientific literature can be regarded as specialists indeed. Specialization, however, is relative, and few readers are as specialized as the authors, who usually studied the topic in great depth. Beware, therefore, of thinking that all readers already know what you know. Among less specialized readers are newcomers to the field, notably young doctoral students; people involved in multidisciplinary research, who are for example specialists of a technique but not of all the fields in which it is applied; and people having to review numerous papers or abstracts across a given field, as when serving on the program committee of a large conference.

Underestimating one's audience is not better than overestimating it, but it is less frequent. Broadening a document's readership without sacrificing accuracy is a worthwhile endeavor for scientific papers as for other documents.

Turning the story around

Deciding what to include in a document and what to leave out is difficult, especially when reporting on a large body of work. To select content effectively, we must first identify our main messages, that is, the conclusions of our work. Most people, however, draw conclusions as they are writing: expressing their thoughts into words catalyzes their thinking. Alas, this approach is all but selective: the authors are writing for themselves, not for their audience.

Typically, authors proceed in chronological fashion.

First of all, they describe almost everything they have carried out as part of the work they report on.

Next, they write down everything they have obtained as a result of carrying out the work described.

As a final step, they think about what they can conclude from it all (what readers want to know most).

Work done



Findings



Conclusions



A more selective approach to deciding what content to include in a report turns the chronology around.

First, figure out your conclusions, that is, select the main messages to be conveyed to your audience.

Then, determine which findings are both necessary and sufficient to support your main messages.

Finally, decide in how much detail to describe the part of your work that led to the findings reported.

Conclusions



Findings



Work done

Selecting your content

CONTENT SHOULD BE LIMITED to whatever material serves the purpose and should be organized in a way that suits the audience. In other words, we should be highly selective, including only the material needed, not all material available. Being selective requires having a clear vision before writing. If we do not, we may well have to write the document twice: once for ourselves, to figure out our main messages, and once for our audience, to get these messages across effectively.

Audiences are seldom, if ever, homogeneous: they combine specialists and nonspecialists, primary and secondary readers. Effective documents, therefore, elegantly address all of these: they provide enough detail for the specialists while enabling any nonspecialist to understand the motivation for the work and the outcome of it; they satisfy the needs of the readers *here and now* and of those far from *here* or years from *now*.

The key to addressing mixed audiences in a written document is structure. At their most global level (their *macrostructure*), effective documents place first the parts that most readers are interested in, then those parts that only knowledgeable and interested readers will read. At the most detailed level (the *microstructure*), effective sentences carefully interlace new material with known material, to offer enough to readers who know less without appearing patronizing to the others.

Effective documents for any audience also take into account the purpose that readers bring to the communication. Thus they anticipate and answer the questions that these readers are likely to have, not only by including the right information but also by presenting this information in the right sequence. These questions may not match the ones we have as authors. In describing work we did, for example, we may be interested in the *how* when readers want to know the *why*, *who*, and *what*. We may thus want to clarify why we did the work, including perhaps who asked us or in what capacity we decided to do it.

Fundamentals

- The name of the game
- The three laws of communication
- A thousand words, a thousand pictures
- Chains and magical numbers
- Trees, maps, and theorems

Effective written documents

- Planning the document
- Designing the document**
- Drafting the document
- Formatting the document
- Revising the document

- Breaking the chronological model
- Including a global component
- Designing fractal documents

Effective oral presentations

- Planning the presentation
- Designing the presentation
- Creating the slides
- Delivering the presentation
- Answering questions

Effective graphical displays

- Understanding pictures
- Planning the graph
- Designing the graph
- Constructing the graph
- Drafting the caption

Applications

- Effective instructions
- Effective electronic mail
- Effective Web sites
- Effective meeting reports
- Effective scientific posters


Designing the document

PLACING FIRST what readers are primarily interested in: there lies the foundation of effective document design. Readers, however, have other interests than authors: there lies the explanation for so many ineffective documents. Effective documents are audience-oriented, not self-centered.

Authors all too often report their work in a chronological way. They present the motivation for the work in an introduction (the *before*), detail this work in a body (the *during*), and report its outcome in a conclusion (the *after*). In this way, they match the work process closely, often all the way to the relative time spent on each part: many authors spend much time writing the body of their reports or articles, exactly like they spent much time carrying out the work, but devote comparatively little time to writing both the introduction and the conclusion.


Readers, in contrast, are primarily interested in the motivation for the work and in the outcome of it, not in the work itself. First of all, they need to relate the work to a broader context much more than they need to understand how it was done, all the more so if they are nonspecialists or secondary readers. Next, they want to know how they are affected by this work, that is, what the findings of the work mean in their own case. In a chronological document, they thus read the introduction, then often go straight to the conclusion, skipping the body. Specialists might later read (parts of) the details in this body, if only to convince themselves of the validity of the conclusion.

A chronological structure, in other words, is straightforward for authors but suboptimal for readers: it requires navigation. Effective document designs, therefore, break the chronology: they place first what the readers are primarily interested in. This first component also provides readers with a global view that helps them assimilate details, should they read (parts of) the body. Such a global component can usefully be included in front of more detailed material at any level in the document.



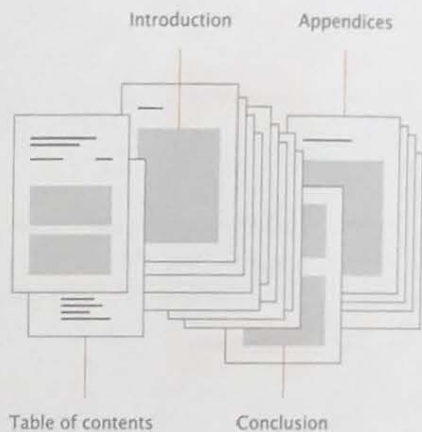
Designing short and long documents

Short and long documents alike usefully include a global component that makes sense on its own, whether or not it is redundant with any other part.



When the document is short (a few pages, perhaps) and addresses a homogeneous audience, it need not be abstracted: the foreword would not differ sufficiently from the introduction, whether in extent or in technicality, and neither would the summary from the conclusion; in each case, one of the two is enough. Still, the document can usefully present the material in the order in which readers want it, with the motivation for and outcome of the work upfront—together with the header on the first page. You can think of the resulting structure equivalently as having an abstract but neither an introduction nor a conclusion or as having, in fact, no abstract but a conclusion placed right after the introduction. Either way, what follows the first page is the body.

A multipage letter can include a global component, too, and be signed at the bottom of the first page, with more detailed material on subsequent pages, perhaps each with its own heading. You can choose to think of it as a one-page letter with attachments instead of a multipage letter, but the idea remains.



When the document is longer than just a few pages and all the more so when it is supposed to address a mixed audience of specialists and nonspecialists (two conditions that are frequently met in practice), it does benefit from including a redundant abstract. Indeed, the introduction and conclusion are likely too long to constitute a useful global component and too technical for the less specialized readers. A distinct global component, made of a foreword and a summary inspired by (but not copied from) the introduction and conclusion, will tell all readers what they want to know first of all and most of all. Ideally, it fits on the very first page, with the header.

Although documents may display the abstract as one block of text so as to be economical with paper, they should nonetheless include both motivation and outcome in it, that is, the *before* and the *after*.



Breaking the chronological model


THE READERS' PRIMARY INTEREST in the *before* and *after* suggests two strategies for effective design. First, devote sufficient time to drafting the introduction and conclusion, for these parts are typically read first of all and most of all. Second, order the material, not *chronologically*, but *logically*, that is, in the order in which readers are most likely to read it.

A chronological story can be made reader-friendly in two ways. One way is to relocate the conclusion after the introduction, before the body; this solution works best for short documents and for homogeneous audiences, such as nonspecialists only. Another way is to maintain the chronological structure as is but to precede it with a component that restates, in a concise and somewhat different way, what readers want to know first, as a form of effective redundancy most appropriate for longer or more complex documents, addressing mixed audiences.

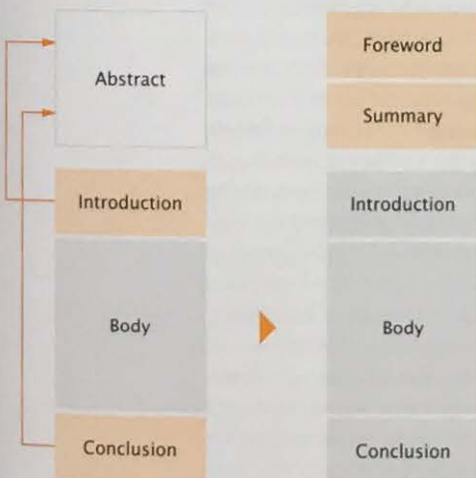
Short or long, documents can thus usefully include a global, stand-alone component, stating upfront the motivation for and outcome of the work. Accordingly, this global component, known under such names as *abstract* or *executive summary*, comprises two parts: the *before*, usually known as *foreword*, and the *after*, usefully named *summary* in a restricted sense. These two parts vary widely in absolute and relative lengths from document to document, even if they are best kept short. An abstract, in other words, is all but a condensed full story: in a sense, it reports the starting point and the end point only.

Global components serve many purposes. By telling the story in a compact, audience-oriented way, they are the only parts that most readers need to read: they thus save these readers much time and effort. For those readers who end up reading the whole document, they provide the necessary framework to understand and mentally structure the subsequent details. For all potential readers, they are a powerful selection tool: they help them decide whether to read (parts of) the document.

Short document
Relocate material



Longer document
Restate material



Equivalent terms

Is the foreword identical to the introduction?

The foreword has exactly the same structure as the introduction, but it is normally shorter and also less specialized than the introduction. It must indeed be understandable to all readers, including less specialized ones. These readers, typically executives, may have different stakes in the need than the specialists. This difference is reflected in the foreword and introduction.

If it seems hard to write a foreword that differs from the introduction by being either shorter or less technical, you might apply the strategy for short documents: in place of an abstract, relocate the conclusion after the introduction.

How can I write an abstract for documents that are in essence detailed, such as minutes of a meeting, specifications, or procedures?

The foreword should pose no specific problem: it accounts for both the document and the work that led to it, as usual. For example, it explains why a meeting was needed, who called it, etc.

The summary can in this case best be regarded as what the audience needs to know most of all or what will help them understand the details. For a meeting, it includes decisions and actions. For specifications (and for similar documents), it might be what is new or special about them, or simply an overview of the device specified.

Is an objective the same as a need?

An objective is half a need only: it corresponds to the desired situation but does not capture the actual one. As such, it fails to justify a task.

The proposed model of global component (abstract) is applicable to almost any professional document reporting some "work done" in the broadest sense. Because it builds on the universal interest shown by professionals in the motivation for and outcome of the piece of work reported, it is just as effective for a business report as for a scientific publication. Still, it is traditionally called or described differently by different professional or even corporate cultures. Below are equivalent terms frequently encountered.

Abstract (the generic term used here) is often used in academic writing but less so in the business world, where a report's global component is usually called *executive summary*, because it enables executives (who are often less specialized) to make decisions. Scientists, engineers, and managers alike may also speak of an *extended summary*, usually to suggest that it encompasses both the motivation (foreword) and the outcome (summary in a restricted sense). A foreword is sometimes called *purpose statement*, because it justifies both the work and the document.

The items in the abstract have several names, too. You may prefer to think of the *need* as a *problem* or an *issue*, but these may be needlessly restrictive: the need may be an *opportunity*, too; furthermore, an intellectual need may be thought of as a *question*. A *task* requested by others can be called a *mission*, one focusing on the means more than on the people. It can be regarded as the *approach* or *method* used. The *object of the document* is sometimes labeled *rhetorical purpose* (although this term is also used for the foreword as a whole, resulting in confusion); when it announces the structure of the document, it can be regarded as an *overview* or as a *preview*. The *findings*, when numerical, are the *main results*. The *conclusion* is the *main message*, also known as *take-home message*, a term perhaps more used for oral presentations. Finally, the *perspectives* are also known as *future work* or perhaps *next steps*.

Including a global component

TO TELL READERS what readers primarily want to know, abstracts include two parts: a foreword and a summary. The foreword, which is similar in nature to the introduction, focuses on the situation *before* the work was done. In contrast, the summary, which is similar in nature to the conclusion, focuses on the situation *after* the work. Effective abstracts include little or no information from the document's body.

The foreword accounts for both the work and the document. First, it states the need for the work reported, as a difference between the actual and desired situations, possibly preceded by whatever context helps understand this need better. Next, it states the task carried out, without detailing what was done. (Strictly speaking, because it corresponds to the *before* part, it states what the authors had decided or been asked to do; usually, this is what they actually did, too.) Finally, switching from a content point of view to a rhetorical one, it establishes the object of the document, that is, what the document does or covers, and possibly what the readers should do with it.

Foreword *The before*

Context

Why the need is so pressing or important

Need

Why something needed to be done at all

Task

What was undertaken to address the need

Object

What the present document does or covers

Summary *The after*

Findings

What the work done yielded or revealed

Conclusion

What the findings mean for the audience

Perspectives

What the future holds, beyond this work

The summary states and, especially, interprets the outcome of the work. First, it mentions the findings or main results. Next, it clarifies what these findings mean, given the audience and the need. So doing, it may recommend a course of action for the audience to address the need. Beyond looking back at the need, it might look ahead and offer some perspectives.

Formulating tasks and objects

What exactly is the difference between task and object? Can these not be combined?

Task and object differ in focus. Schematically, the task states what the authors did, whereas the object establishes what the document does. The first thus focuses on the work carried out; the second, on the communication of this work. Should you combine the two, focus on the task (*Therefore, we developed a new method to ...*), rather than on the object (*This paper presents a new method to ...*): the task follows the need more logically and clarifies who did the work. Also, beware of mixing up places and tenses, as in *In this report, we measure ...*, which makes no sense: you do not in fact measure anything in the report; you did so, presumably, in the lab.

What does the task become for a review paper?

As author of a review, you can regard your job as reporting the work of others, in which case you sum up this work, collectively, in the task (*Over the last twenty years, researchers have ...*), or as reporting your review work, in which case you state so in the task. The need, of course, must be adapted; in the second case, it must be the need for a review of research in the field.

Should a scientific abstract include an object?

The object of the document must orient readers as to the document's contents and structure, so they can decide whether (or what) to read. When it merely repeats what the need and task already state or clearly imply, it is best omitted. Such is often the case for experimental work, for which the body is predictably structured as *Materials and methods, Results, Discussion*.

Of the many ways to formulate tasks and objects, some are more logical or more readily recognized.

The task must clarify not only the work carried out but also the agents who carried it out. If these are the authors of the document (the most usual case), it is best phrased in the active voice and first person (*we measured, we developed, we implemented, etc.*). If not, as for a review paper, it is logically phrased in the third person, preferably in the active voice (*In recent years, researchers have investigated ...*), more informative and usually more readable than an impersonal construct (*... has been investigated*).

The task can be effectively connected to the need by a phrase before the subject. A simple *therefore* is usual enough, but the desired part of the need is often an elegant option (*To increase the speed, we redesigned ...*). This initial phrase can be used to clarify any requesting party, too (*At the request of the Senior Management Team, we redesigned ...*).

The object, in contrast, focuses not on the authors but on the document itself. Thus, it can best use the document as grammatical subject of the clause (*This report summarizes, This paper presents, etc.*). Such a phrasing is equally appropriate for objects of parts of the document (*Chapter 3 compares ...*).

Task and object differ also in the tenses they use. The task, like all references to the work carried out, is best set in the past (or the present perfect) tense. The object, like all references to actual documents, is atemporal and thus best set in the present tense.

Whereas choosing a grammatical subject for a task or an object is straightforward (*we, this report, etc.*), selecting a suitable verb for it is more challenging. Did we *examine, analyze, or perhaps investigate*? Does the document *present, explain, or report on*? The verb, not the subject, carries the strength here.

Thinking in concentric layers

Abstracts can be thought of as made of four concentric layers. At the center is “communication,” a rhetorical layer focusing on the document. Around this rhetorical layer is a “work done” layer, consisting of the task (*before*) and the findings (*after*). It focuses on the authors, who usually have carried out the task and generated the findings. Around this “work done” layer is a “problem-solution” layer, comprising the need (*before*) and the conclusion (*after*). This one focuses on the readers, who, presumably, are concerned by the need and must have the findings interpreted, in view of this need. Finally, around this “problem-solution” layer is an optional “situation” layer, comprising the context (*before*) and the perspectives (*after*). It focuses neither exclusively on the readers nor exclusively on the authors (perspectives might involve either of them); in a sense, this is an “anyone” (or an “everything else”) layer.



Context	Anyone	
Need	Readers	<i>What you want ≠ what you have</i>
Task	Author(s)	<i>What I/we did to address the need</i>
Object	Document	<i>What the document does/covers</i>
Findings	Author(s)	<i>What I/we found, doing the task</i>
Conclusion	Readers	<i>What these findings mean to you</i>
Perspectives	Anyone	

Established from a chronological point of view (*before/after*), the abstract in two distinct parts can easily be reinterpreted in terms of the motivation (the *why*) and outcome (the *what*). The first part (the *before*) might be known to some readers, for these might well have pointed out the need and ordered the task, so the news appears in the second part (the *after*). Still, the first part not only informs those who did not know, it also prepares those who did to assimilate the message(s) optimally, by reminding them of the context and of the need, thus placing both the task and the document in perspective.

*The need is well known to my audience.
Should I still include it in the document?*

A need is bound to a point in space and time. It might be known to primary readers, but not to secondary ones. Moreover, primary readers seldom have it in mind as they begin to read and might not see it the way the authors see it. A carefully stated need thus puts all readers in a favorable frame of mind for their reading.

The foreword part of the abstract establishes motivation and, to some extent, relationship. It always plays an important role. Of course, it can be brief or presented as already known, especially when the audience is clearly defined (it might start with *As you may remember, ...*).

Is having the motivation in the introduction not enough? Why should it be in the abstract?

Whereas titles are the very first selection tool, abstracts are the main one. If an abstract fails to include motivation, readers are more likely to stop reading than to look for it elsewhere, such as in the introduction—all the more so when the abstracts are available free of charge, whereas the full documents must be purchased.

I must write an abstract limited to 100 words. This is very short. Can I drop any item at all?

When the number of words is severely limited, condensing or combining items is preferable to suppressing any, for all of them play a role. When needed, context, need, and task can often be written as one sentence (*In the framework of ... and in an effort to ..., we investigated ...*). At times, the object can be expressed implicitly.

Common suboptimal abstracts

While any abstract missing one or more components is suboptimal, three types are particularly frequent: *promissory, out of the blue, and self-centered ones.*

Promissory	Out of the blue	Self-centered
Context		
Need		
Task		Task
Object		
	Findings	Findings
	Conclusion	
	Perspectives	

An abstract including the motivation (the *before*) but not the outcome (the *after*) is most frustrating: it promises but does not deliver, obliging readers who would have had enough with a good abstract to dig in the details in an effort to find the message. It is perhaps acceptable as a proposal submitted to a selection committee long before a conference but far less so as the abstract of a final document.

An abstract including the outcome (the *after*) but no motivation (the *before*) comes out of the blue. The findings, appearing first, are easily confused with context. Moreover, the conclusion is usually impossible to understand properly without a need. While it may seem acceptable to primary readers, well aware of the need and task, an out-of-the-blue abstract is cryptic for secondary ones, who are not.

An abstract that is limited to a task and findings, with no real motivation (need) and no real message (conclusion), is self-centered: it ignores that which the audience most wants to know. It makes them wonder what this document has to do with them, as well as what they are supposed to do about it.

Conveying the motivation and the outcome

The two key questions, *why* and (*so*) *what*, can be associated symbolically or mnemonically to the two parts of the abstract and, within these parts, to each item, with its respective slant.

The first part of the abstract implicitly answers the question *why* at every step. The context answers the question *why now*: it explains how the recent history and current situation led to the need. If reader-oriented, the need implicitly answers the question *why you*: by reading it, readers should perceive the document's relevance to them, that is, understand why it was sent to them or decide whether to select it for reading. In turn, the task implicitly answers the question *why me/us*, clarifying what the authors have to do with the stated need. Finally, the object answers the question *why this document*, that is, *given the need and the task, what are the document's purpose or contents? And what should we, readers, do with it?*

Why? Motivation	Context	Why now	<i>My/our or your current situation</i>
	Need	Why you	<i>Why this is in fact relevant to you</i>
	Task	Why me/us	<i>What I/we have got to do with it</i>
	Object	Why this document	
What? Outcome	Findings	What	<i>What resulted from the task done</i>
	Conclusion	So what	<i>What these findings mean to you</i>
	Perspectives	What now	<i>What I/we or you should do next</i>

The second part of the abstract similarly answers the question *what* at every step. Thus the findings merely state the *what*, that is, *what have I or we found by carrying out the task?* The conclusion (the message) answers the question *so what*, typical of readers who cannot guess whether some findings mean good news or bad news, or how good or how bad it is. Closing the outermost loop of the multiply nested structure, the perspectives, like the context, focus on temporal aspects and answer the question *what now* or, equivalently, *what next*.

Creating useful headers

*What is wrong with a heavily technical abstract?
If readers do not understand it, can they not see
that they are not part of the intended audience?*

Giving up on a document because one plainly does not understand is such a negative choice. Is it not infinitely more constructive to enable all potential readers to understand our abstract and, with this understanding, decide rationally whether they can benefit from the document?

Unlike the document, the abstract will be read by virtually all readers, be it as a selection tool. Consequently, a reader-friendly abstract is one that is understandable to the least specialized of our readers (while still useful to specialists).

*Should I write the abstract first or last, that is,
before or after writing the complete document?*

Writing the abstract after the whole document usually makes for better abstracts. Writing it before the document, however, usually makes for better structured, more selective documents. Perhaps you can write abstracts first *and* last?

In a similar way, writing the introduction early typically makes, not for better introductions, but for better research, because it helps clarify, for the authors at this point, what they have, what they want, and how they hope to find it.

Must foreword and summary be labeled thus?

The foreword and the summary do not need to be labeled; neither, in fact, does the abstract. Global components, be they abstracts or other, are mainly identified as such by their location. Above all else, they must be in the right place.

Whatever the document is, the header must allow its correct routing, filing, and selection for reading or other processing, such as destruction. Readers mostly select on the basis of the title and author(s), so these should be prominent. Equally important yet too often absent in online documents is a date and, for unambiguous reference, a unique identifier, such as a reference number. This identifier should then appear on each page next to the page number, so the page can be traced even if it becomes loose.

The title of a report or of an article, or, equivalently, the *subject* of a memorandum or electronic mail, is a challenging element to write. To remain visual, it must be kept short—a question of visual length on paper or on the screen more than of word count. The question is, what do we tell in these few words?

Titles, like documents, can be structured effectively in two parts: a global one and a more specific one, perhaps separated by a dash or a colon, or written as a title and subtitle. The two parts work together like a funnel, gradually narrowing down the topic (*Project meeting on Fri 20 Apr: revised agenda*). Used in isolation, the first part would be too vague (what about the "project meeting on Fri 20 Apr?") whereas the second one would lack its framework (the "revised agenda" for what meeting?). The titles of related documents can have the same first part and distinct second parts, making for easy sorting (*Project meeting on Fri 20 Apr: list of attendees*).

Whether the distribution list is part of the header is debatable. Clearly, it plays a role for selection, as when recipients want to know whether they are listed under *to* or under *cc*. In any case, it should not stand in the way of more important information. When it is long, it is best moved out of the header and onto a page of its own, inserted where it can easily be located without being overly prominent, such as after the first page—or after the last one.

Allowing informed decisions

The abstract, with its two parts, allows two decision moments. After reading the first part (the motivation), readers can gauge their interest for the topic and for the document. Only those who care enough about both will go on reading. After reading the second part (the message), they can normally take action about the need and, again, decide whether to quit or read on. This decision will be influenced by their confirmed interest but depends largely on their need for the details. Specialists typically want the details, either because they can use them in their own work or because they wish to assess the quality of the work done and, with it, the credibility of the message.

Header Screening	Title Author(s)		
		<i>Is this for me?</i>	If yes, go on to the foreword If no, then stop reading here
Foreword Motivation	Context Need Task Object		
		<i>Do I care?</i>	If yes, go on to the summary If no, then stop reading here
Summary Outcome	Findings Conclusion Perspectives		
		<i>Do I need more?</i>	If yes, go on to the document If no, then stop reading here

A very first selection tool, however, is the header preceding the abstract, specifically the document's title and author(s), and possibly the mention of the reader's name under *to* or *cc*. The global component of a document is thus best regarded as comprising three parts: header, foreword, and summary. For most documents, notably short to medium-length reports, the three parts together should ideally not exceed one page, as it increases the probability of their being read immediately. Having to turn a page seems to be quite a psychological barrier.

Below is the executive summary of a short report addressed by the tax officer to the CEO and CFO at Amazing Belgian Chocolates (ABC), in preparation for a meeting in which the issue will be discussed; the other members of the committee appear in cc. All of them are aware that ABC, selling its products on the Belgian market, owes Value Added Tax (VAT) to the Treasury, while sister company Choc Export, exporting ABC's products, can claim VAT refunds.

The context, limited in this case to a single sentence, is sufficient to situate the need in both time (April 2007) and space (Belgium). Next comes the need, presented as an opportunity, not a problem. All the same, it represents a gap between ABC's current situation and a more favorable alternative, which is desirable for the readers.

Given the audience of executives, the task is explicit on processes. The author's function (tax officer), normally indicated in the header, can justify that she went ahead and "prepared a rough estimate" but perhaps not that she engaged expenses for legal consultation, so the task clarifies that the CFO backed her. It also indicates who the consultant is, thus answering a possible question from the CEO.

The object of the document says not only what the report covers but also what the readers should do with it (prepare the meeting).

The summary in itself (findings, conclusion, perspectives) is set here as one solid paragraph, but, if longer, it could also be divided in two or three paragraphs, as is the case for the foreword above. Conversely, this foreword could be set as a single paragraph, too; at less than 200 words, it is still reasonably short. Which option is best (one paragraph or several) depends also on the page layout.

As of April 2007, the Belgian legislation authorizes VAT consolidation among the members of a group. Such a VAT group could be an opportunity for us at Amazing Belgian Chocolates and for our sister company Choc Export: our VAT payments could be used to offset their (notoriously slow) VAT refunds, positively affecting the cash flow of both companies. Filing a consolidated VAT return for the VAT group thus formed, however, would require that we adapt our accounting practices as well as our IT systems.

As a first step toward evaluating the benefit for us of forming a VAT group together with Choc Export, I prepared a rough estimate of the potential savings and, in agreement with our CFO, Linda Thielemans, I contacted a tax consultant at Van Belle & Partners to verify that ABC and Choc Export meet the criteria for constituting a VAT group under the new decree.

In preparation for our upcoming committee meeting on Wed 13 Jun, this note sums up my findings so far and maps out what I identify as the next three steps, to be discussed in committee and approved by you.

My rough estimate suggests consolidated savings in excess of 20 000 euros per year, and the reply from Van Belle & Partners is positive on all points. Accordingly, I recommend that we move forward with the project by carrying out these three tasks: (1) with the accounting teams of both companies, evaluate the required adjustments to the systems and the costs associated with implementing them; (2) if the anticipated savings exceed these costs, request authorization from the VAT administration; (3) once it is authorized, implement the VAT group.

Below is a possible abstract for a scientific article published in a journal (abstracts of research reports for internal use in an organization would be similar). Such an abstract addresses other researchers. Still, in an effort not to be needlessly technical, it limits acronyms to those strictly useful for conciseness: thus, it introduces *ILT* (which it uses three times), but it avoids *AAA* for *abdominal aortic aneurysm*, even though this acronym is used in the article itself.

Context—provides background for less specialized readers and, so doing, establishes or recalls the importance of the problem.

Need—motivates the audience by stating the difference between the desired and actual situations.

Task—states what the authors undertook to address the need, in the first person (*we*), past tense.

Object—clarifies what the paper covers without repeating the task, in the active voice, present tense.

Findings—state the main results in a way that both less and more specialized readers find helpful.

Conclusion—interprets findings (states the *so what*)—in this case, all the way to a recommendation.

Perspectives—broaden the view with any further needs and tasks.

An aneurysm in the abdominal aorta will rupture as soon as the wall stress exceeds the wall strength at any location, thus threatening the patient's life. Elective surgical repair, however, is costly and risky. Evaluating wall stress to predict the risk of rupture is therefore essential toward patient management, yet current models suffer from several limitations. Among others, they do not consider the presence of an IntraLuminal Thrombus (ILT), a fibrin structure present in variable degrees in 75% of aneurysms. Using computed tomography and finite elements, we investigated whether the presence of ILT alters the distribution or the magnitude of the wall stress in aneurysms of the abdominal aorta. This paper reports the wall stress distribution in four patients and discusses the impact of the ILT configuration.

In all four patients, the presence of an ILT altered the stress distribution and reduced the peak stress by 6 to 38% ($p = .067$), depending on the geometry. As a consequence, it should be taken into account in any patient-specific model of aortic aneurysms for evaluating the wall stress and the risk of rupture. Still, it may also adversely affect the wall strength and will therefore remain the focus of future work.

With just under 200 words, the above abstract can convey the motivation for and outcome of the work with some accuracy, without intimidating readers by its length. Still, and when allowed by the journal, it is best set in two paragraphs (foreword, summary). An abstract under 150 words is challenging to write; one over 250 words is seldom justified for an article.

Common shortcomings

Should every item at every level begin with a global component?

If an item constitutes a meaningful entity, surely there is something general to be said about it, before each of its constituent parts is detailed. At the very least, you can prepare the readers for the item's structure, even if this structure is already shown by a set of headings. Indeed, headings are often seen but not read. Moreover, they are not seen together when the item spans several pages, so they provide no global view.

Can the global component of a chapter or section be its first section or subsection?

The global component of a chapter or section is really a level above sections or subsections: it belongs to the chapter or section as a whole, so it is best placed directly under the heading of this chapter or section, before any heading of section or of subsection. (As a minor point, doing so results in a clearer page layout, too.)

Is an object still required when readers are used to a structure (for example, when it is identical to that of matching items in other documents)?

Primary readers might be used to a structure, but any secondary readers, almost by definition, cannot be: an object is thus a necessity for them and may be a useful reminder to other readers. However, avoid presenting as new something that is known to your readers, lest you appear patronizing. If the structure is usual, say so. For a section on *Results*, you might thus write *As usual, our results come from three areas: ...* You would list here the areas in question again, before discussing each of them in a subsection.

As a clear confirmation of the usefulness of telling the beginning of the story and the end of it together in one place, some scientific journals now require that authors announce their conclusion at the end of the introduction. Similarly, many authors start the conclusion section (at the end of the document) by restating the need, the task, or even the object of the document, which appear in the introduction. Introductions or conclusions structured in this way, while well-meant, needlessly duplicate the abstract. With a well-written abstract, readers are equipped for the complete article; the introduction should be limited to the *before*; the conclusion, to the *after*.

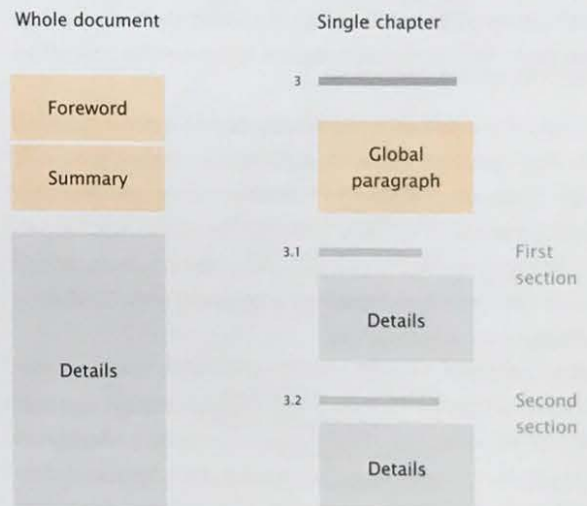
Even if deemed unnecessary as part of the abstract, the object of the document should certainly appear in the introduction, typically as a final paragraph starting with *This document is organized as follows* and then listing what each chapter or section covers, preferably in the active voice (*Chapter 2 reviews ...*). The chapters or sections, in turn, should include a global paragraph upfront, to orient the readers.

Two headings that follow each other immediately, such as those of Section 2.4 and Subsection 2.4.1, are usually an indication that a global paragraph is missing. At times, however, the first subsection is precisely what is needed for a global paragraph, so we can simply delete that subsection's heading (and, of course, renumber the subsequent headings). Usually, the heading thus deleted was of little use, as in *General description* or *Preliminary remarks*.

The presence of a paragraph between the heading of a section and that of its first subsection does not, in itself, guarantee an effective mental preparation of the readers for the structure ahead. In particular, such paragraphs often miss the *object of the section*, (the local equivalent of the object of the document), especially when they have been built by relocation of other items, rather than written for the purpose.

Designing fractal documents

EFFECTIVE DOCUMENT DESIGNS ARE FRACTAL. They exhibit an identical structure—global component, then details—for the whole document and for each of its constituent items at any level, such as the chapters, sections, and subsections, in fact all the way down to the level of individual paragraphs.



Like documents, chapters (or sections or subsections) should begin with a global component, perhaps just one paragraph, typically located before the heading of the first subbranch (for a chapter, before the first section). This global paragraph might contain any and all items of a foreword and summary. At the very least, it must include an object, preparing readers for the chapter's structure, but it can also usefully set forth a motivation for the chapter or the chapter's main messages.

Chronological chapters (or sections or ...) can often be turned into global-detailed ones in the same way that documents can. Any conclusion section, for example, can be moved upfront as part of the chapter's global component and be suppressed where it was—or for long chapters, be restated briefly upfront.

Fundamentals

- The name of the game
- The three laws of communication
- A thousand words, a thousand pictures
- Chains and magical numbers
- Trees, maps, and theorems

Effective written documents

- Planning the document
- Designing the document
- Drafting the document**
- Formatting the document
- Revising the document

- Stating and developing messages
- Expressing ideas
- Being clear, accurate, and concise

Effective oral presentations

- Planning the presentation
- Designing the presentation
- Creating the slides
- Delivering the presentation
- Answering questions

Effective graphical displays

- Understanding pictures
- Planning the graph
- Designing the graph
- Constructing the graph
- Drafting the caption

Applications

- Effective instructions
- Effective electronic mail
- Effective Web sites
- Effective meeting reports
- Effective scientific posters

Drafting the document

CAREFULLY WRITTEN DOCUMENTS are reader-friendly, just like carefully designed software is user-friendly. They are simple and direct, trying to help the readers, not to impress or confuse them; in a word, they are readable. They not only contain the right information in the right place but also phrase this information in a way that is easy to read. Hence, they allow readers to spend as little time as possible on the text and as much of it as possible on the ideas expressed.

Readable written documents are clear, accurate, and concise. They carry a single, readily understandable meaning. They tell “the truth, the whole truth, and nothing but the truth.” And they achieve clarity and accuracy in as few words as possible. Clarity, accuracy, conciseness (in that order) are the priorities at the level of the paragraphs, sentences, and individual words.

Text entities at different levels have different functions, hence they may require different approaches for optimal readability. Paragraphs convey messages, ideally in a stand-alone way: each paragraph thus states its message first, then develops it in theorem-proof fashion. Sentences state ideas, in one-to-one correspondence: short sentences can express simple ideas, while longer sentences are required for more complex ideas. Well-chosen words contribute to a clear, accurate, concise text.

Writing simply does not suggest simplifying the material but conveying the complexity of this material in the simplest way, something many of us have learned to do about mathematics. We thus have been taught to “simplify fractions”; for example to replace $679/194$ by the more readable, yet equivalent $7/2$; we have seen the advantage of a change of coordinate system to describe an object or a phenomenon with simpler equations; we have been rewarded for deriving the “most elegant” proof of a theorem. Still, when useful, we know how to approximate, that is, to simplify the material to the desired level of accuracy, as when writing $\pi \approx 3.14$, which is in fact a correct statement.

Stating the message upfront

Must every paragraph convey a message?

Every paragraph must have a unifying theme (a reason for its sentences to be set together) and must serve the purpose of the document. Often, this theme is a message, but not always. An introductory paragraph, for example, may not hold a true message, expressing a *so what*. Conversely, a one-paragraph abstract typically comprises several messages, at various levels. Whether or not its theme can be considered a true message, each paragraph should orient the readers as to its content and its structure.

I find it difficult to ensure that each paragraph states a message upfront. Any suggestion?

As much as possible, determine what you want to say with each paragraph before you write it. You could even write all the first sentences first, that is, before writing the rest of any paragraph, as a way to envision the structure of a section. Together, these sentences answer the question, *what ideas must my audience remember here?*

After writing, you can inspect each paragraph. If the message is at the end, move it to the front, adjusting the paragraph as needed. If there is no message, question the paragraph's function.

Is there a maximum length for paragraphs?

Each paragraph should have whatever length it needs to develop its message appropriately. Some might be short—as short as one sentence, perhaps—while others are comparatively long. Long paragraphs can be visually intimidating, though; paragraphs that must get the attention of the audience are therefore best kept short.

Paragraphs that fail to clarify their topic upfront are frequently misleading. In the example below, the first two sentences suggest that the paragraph discusses "single-use, disposable medical devices," then the third sentence reveals that what it does, actually, is compare two types of medical devices.

Single-use, disposable medical devices are packaged and sterilized by the manufacturer. Their packaging must provide protection, facilitate sterilization, maintain sterility, ... Reusable devices, by contrast, must be ...

Medical devices may be broadly divided into two categories, disposable and reusable, having different sterilization requirements. Single-use, disposable devices are packaged and sterilized by the manufacturer. Their ... Reusable devices, by contrast, must be ...

More common than a missing topic is a message that appears too late and therefore lacks visibility. Such is typically the case in a paragraph structured chronologically with its conclusion at the very end, thus giving the proof before stating the theorem, or one postponing the *so what* until after the *what*, as when describing a figure before interpreting it.

Figure 2 shows the evolution of the Ge content in the SiGe layer. Obviously there is a nearly linear decrease of the Ge content with increasing fluence. Knowing the ...

The germanium content decreases linearly with increasing fluence (Figure 2).

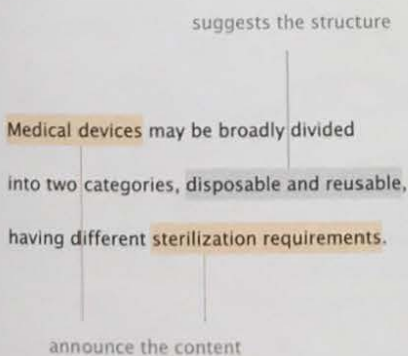
Stating and developing messages

PARAGRAPHS ARE THE ESSENTIAL structuring components of the document: they appear even in short documents that are not otherwise divided into sections, such as letters. They retain document characteristics that sentences do not. First, paragraphs—like documents—tell a story of their own: effective paragraphs remain understandable even in isolation, whereas many otherwise effective sentences do not. Second, paragraphs convey messages: each states and develops one.

Paragraphs, like full documents, should remain meaningful (to a point, that is) even out of context. Each paragraph must have its reason for being, namely to convey a certain message. Conversely, each message usually deserves its own paragraph: one paragraph per message and one message per paragraph. Paragraph size then simply depends on message complexity: more complex messages typically require longer paragraphs. As with sections, the optimal granularity is also a question of balance between paragraph length and paragraph number: few long paragraphs are hard to read; so are many short ones. In most cases, a sequence of one-sentence paragraphs makes little structural sense: messages are stated but not developed.

Effective paragraphs, like effective documents (or chapters), state their message early, ideally as early as the first sentence. A well-written first sentence, not unlike a well-written abstract, enables the readers to decide whether to read the paragraph, by announcing its content and by suggesting its structure. With prominent key words, first sentences help the readers locate quickly the paragraphs that are most relevant to them or find again some information they remember having read. By reading the first sentence of every paragraph, the readers should be able to form a good idea of the document's content.

After stating its message, an effective paragraph develops it along an appropriate structure, typically revealed by the way in which the various sentences are connected to one another.



Drafting effective lists

Need paragraphs always be parallel or serial?

Paragraphs need not always be entirely parallel or entirely serial. They may use a combination of the two structures, or be "pseudo-parallel" (lining up comparable yet not identical subjects). To be readable, however, they should not miss opportunities for a parallel or serial structure, such as introducing a switch in subject ($A \rightarrow C$) that does not reflect a switch in topic (yielding $A-B \ C-A$, instead of the parallel link $A-B \ A-C$) or positioning a new item (C) before an item (B) mentioned in the previous sentence (yielding $A-B \ C-B$, rather than the serial link $A-B \ B-C$).

Is the parallel structure not boring to read?

Parallelism may seem to encourage repetition. Not so, however: unpleasant repetitions must of course be removed lest they become noise, but not by uncalled-for variations in structure. When attempting to "parallelize" a paragraph, you can remove resulting repetitions by using pronouns and by combining related sentences, not unlike rewriting $3ax + 5ay$ as $a(3x + 5y)$.

How should I punctuate a displayed list?

The rules for punctuating displayed lists vary from book to book (and language to language). Whichever you decide to apply, be consistent. For written documents, consistency suggests using in lists the general rules of punctuation and capitalization: thus, capitals and periods for full sentences, and commas or semicolons to separate phrases or clauses within a sentence. For oral presentations, the desire to be visual may suggest dropping the punctuation marks in lists (and perhaps in some other text items).

Common not only in written documents but also on oral presentation slides, lists too often exhibit shortcomings that render them plainly ineffective. Lists are for displaying comparable items in a way that encourages their comparison or memorization, not for making a loose set of items look organized. Whether they are displayed (with or without bullets) or typeset as part of a solid paragraph, lists should

- comprise few items (in other words, five or fewer), to allow their nonsequential, visual processing;
- introduce the items by a clause (or part of one), to let the readers know what the list is about;
- phrase all items in a grammatically similar way, to reflect in the form the parallelism of content.

The manner in which the items are phrased should obviously be a grammatically correct continuation of the introductory component. The use of bullets to reveal items does not alter the rules of grammar.

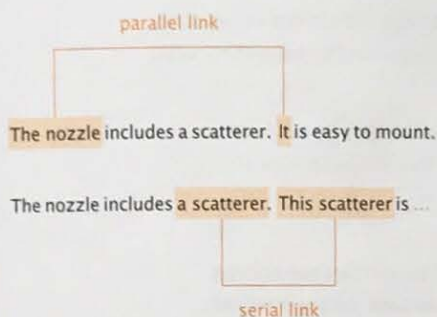
- To prepare a meeting, define its purpose
- You must also prepare an agenda
- Everyone should receive this agenda
- Does everyone know who the others are?
- The chairperson should not be secretary
- Ground rules may be appropriate, too
- Always review the purpose and agenda

When preparing a meeting,

- define the purpose and agenda,
- send the agenda to all participants.

As you start the meeting,

- welcome and introduce participants,
- clarify the roles (chair, secretary, etc.),
- set up ground rules if appropriate,
- review the purpose and the agenda.



Above all else, sentences within a paragraph are connected by content: one element in a sentence, normally its subject, points to an element in the previous sentence. A reference to the previous subject is best done with a personal pronoun (*it, they, we, etc.*), whereas a reference to the previous object or other item appearing at the end of the previous sentence is best indicated by a demonstrative adjective (*this, such, etc.*). By analogy with elementary electrical circuits, we might call the first case a parallel link and the second case a serial link.

Parallel or serial links can be repeated for several sentences: a parallel structure lines up sentences with the same subject,

Recent years have seen an increased popularity of codes based on the Diabolo algorithm. Speed is a main advantage of these codes, compared to the traditional Demon ones. Also, one can implement them reasonably easily, and it is possible to extend them so they can handle hybrid transforms. On the other hand, they require about 45% more memory, but this is less critical with today's architectures. Typical applications are ...

In recent years, codes based on the Diabolo algorithm have become increasingly popular. Compared to the traditional Demon codes, they are about twice as fast, are reasonably easy to implement, and can be extended to handle hybrid transforms. As a drawback, they require about 45% more memory, a less critical limitation with today's architectures. Typically, they are applied to ...

whereas a serial structure chains sentences by using what is introduced in one sentence as the subject of the next sentence.

All current implementations of the Diabolo algorithm are based on the so-called Angel transform. F. Angel first described this transform in [2]. The idea is to separate the data into high and low values before proceeding with generation. The implementation then stores the high and low values separately ...

All current implementations of the Diabolo algorithm are based on the so-called Angel transform. This transform, first described by F. Angel [2], separates the data into high and low values before proceeding with generation. The high and low values are then stored separately ...

Both structures can be used to construct entire paragraphs. The parallel structure, using the paragraph's topic as subject of all its sentences, is the more readily applicable of the two. The serial structure is well suited to introductory paragraphs, organized from general to particular, and to substructures, for example within a more complex parallel-serial paragraph.

*I have been told to write in the active voice.
How can I then use the topic as the subject?*

Do not confuse *active voice* and *first person*. The subject of a verb in the active voice need not be a person; it can be an inanimate object, as in *This paper presents* or *The results show*. Even if you choose to write in the first person, the topic is largely expressed in the verb, too. If you write *we decided*, readers will understand that the sentence is about a *decision* you made, although this word is not the sentence's subject. If all the sentences use *we* as subject, however, readers may feel you are talking about yourself.

I can recognize instances of main information in a subordinate clause, but how can I fix them?

To improve an instance of main information in a subordinate clause, proceed in two steps, as follows. First of all, delete the main clause, so the subordinate clause is now the main one. Then, identify what you have lost, if anything, and add it back, keeping your new main clause.

As an example, if your original sentence reads *Figure 3 shows that the simulation worked well*, first delete *Figure 3 shows that*: your sentence now simply reads *The simulation worked well*. You have, however, lost the link with Figure 3; just put it back, either as a subordinate clause (*As Figure 3 shows, ...*) or simply in parentheses at the end: *The simulation worked well (Fig. 3)*.

In contrast, if your sentence is something like *It was observed that the simulation worked well*, you lose little by deleting *It was observed that*. Quite on the contrary, the simpler statement *The simulation worked well* has more impact, especially as the first sentence of a paragraph.

Common shortcomings

To avoid a passive voice, authors may end up using a subject that is in fact not the paragraph's topic.

Table 5 summarizes the results of our tests.
These results show the superiority of ...

The results of our tests are summarized
in Table 5. They show the superiority of ...

or better, with the details in a subordinate clause,

The results of our tests, summarized
in Table 5, show the superiority of ...

Focused on their work, scientists too often describe an observed phenomenon in a subordinate clause and state in the main clause that it was observed.

Figure 2 shows that, in most experiments,
the rate was lower at higher temperature.

In most experiments, the rate was lower
at higher temperature (Figure 2).

Impersonal main clauses of the form *it is ... that ...* also relegate the message to a subordinate clause, weakening the point they try to draw attention to.

It is clear that ...	Clearly, ...
It is obvious that ...	Obviously, ...
It is a surprise to us that ...	Surprisingly, ...
It can be concluded that ...	In conclusion, ...

Expressing ideas

SENTENCES ARE FOR CONVEYING IDEAS. Well-constructed sentences, therefore, stem from well-constructed ideas or, conversely, suggest clear and accurate ideas to the readers. Effective sentences reveal the structure and sequence of ideas logically, yet without taxing the readers' short-term memory.

With each sentence, strive to convey one idea and one only. How long the sentence should be thus depends on the idea expressed. Simple ideas can be conveyed in simple sentences consisting of an independent clause; complex ideas require complex sentences, then consisting of a main clause and one or more subordinate clauses. What constitutes a single idea is of course a view of the mind. As recommended for sections and paragraphs, an optimal division into single ideas strikes a balance between few long sentences and many short ones.

A sentence (or a clause) consists of a subject and a predicate, that is, a verb phrase asserting something about the subject. Because subject and verb are central in conveying meaning, they should be selected with care. As a rule, use the "subject" of the expressed idea as grammatical subject of the sentence. If all sentences in a paragraph cover the same topic, this topic can consistently be the subject, yielding a parallel structure.

When writing complex sentences, state the main information in the main clause and relegate the subordinate information to subordinate clauses. Such sentences prioritize the ideas, giving maximum impact to what is most important. They can effectively address mixed audiences, by placing the *new to all* in the main clause and the *new to some* in a subordinate one.

As we agreed last week, I contacted Van Belle & Partners about VAT consolidation between ABC and Choc Export. In their opinion, these two companies satisfy the criteria set forth in the April decree for constituting a VAT group. Consequently, we can now move forward with the idea ...

We agreed last week that I would contact Van Belle & Partners about VAT consolidation between ABC and Choc Export. I just did. They believe that the criteria set forth in the April decree for constituting a VAT group are satisfied by these two companies. This means that we can now move forward with ...

Keeping together what goes together

Is there a maximum length for sentences?

There is no absolute maximum for the length of a sentence, especially when this sentence comprises several clauses: readability stems from structure, not length. The average length of sentences is nonetheless best harmonized with that of paragraphs, to avoid an imbalance.

In languages that locate the verb, or part of it, at the end of a clause, the length of this clause is a direct measure of distance between subject and verb, and should then be kept reasonable.

How can I fix an overly long subject, while keeping it in subject position?

Admittedly, turning a sentence around makes the subject become the object, so it may well shift the sentence's focus in an unwanted way. As an alternative, especially for long subjects that are lists of items, try a forward reference. Instead of writing ..., ..., and ... were carried out, put the list beyond the sentence (after a colon), and simply replace it by a phrase pointing to it, as in *The following three tasks were carried out.*

Should I rely on readability formulas?

Formulas assessing the readability of a piece of text on the basis of its average word length, average sentence length, and similar measures have little applicability to educated audiences. Most were developed for school-age children reading in English as their first language, hence they must be extrapolated with caution, if at all. For example, educated non-native speakers are often more familiar with long technical terms than with short words having several meanings.

A long subject with a short predicate may be logical, but it is difficult to process. It is a common issue with tasks or objects expressed in the passive voice.

A finite-element simulation of the principal component's critical subparts was performed.

We performed a finite-element simulation of the principal component's critical subparts.

Equally difficult to process is a new item positioned before an item that refers to the previous sentence.

.... Chapter 4 of the *Safety handbook for the laboratory* discusses this issue.

.... This issue is discussed in Chapter 4 of the *Safety handbook for the laboratory.*

Finally, long parenthetical information in the middle of a sentence is best relocated at the very end of it.

Two hours later, a standard breakfast, consisting of 4 slices of buttered bread, 1 slice of cheese, 1 slice of ham, jelly, and 2 cups of coffee or tea, was served to the subjects.

Two hours later, the subjects were served a standard breakfast, consisting of four slices of buttered bread, a slice of cheese, a slice of ham, jelly, and two cups of coffee or tea.

Although logically constructed, some sentences prove hard to read because they strain the readers' short-term memory. When two related sentence parts, such as a subject and a verb, are far away from each other, readers must store the first part in their short-term memory (not knowing what to do with it) while processing the sentence further; only when they reach the second part can they attribute meaning to the two together. A similar difficulty arises when the two related sentence parts are close to each other, but the first one happens to be long.

What makes a sentence hard to read, then, is the necessity to store items before they can acquire meaning. In contrast, sentence length, frequently decried as hindering readability, is but a symptom: the longer the sentence, the higher the risk of taxing the short-term memory of the readers. Short or long, however, sentences with two or more levels of nested clauses will likely be hard to read. In contrast, carefully constructed long sentences can be almost as easy to read as short ones, while conveying relationships that short sentences could not.

In other words, look out for distance issues in your writing: keep together what goes together, within and across sentences, and place short items before long ones as much as possible. Moreover, to help readers figure out quickly how a sentence relates to previous ones (so they need not store it to do so), show links with connection words: conjunctions, adverbs, etc.

Some information is intrinsically hard to convey in a sentence, no matter how well written. It might then be better conveyed more visually, for example as a formula, a table, or a diagram.

Dick, who, when Jane arrived, left, returns.



Although short, the above sentence is hard to read because of its embedded structure (a subordinate clause within another one).

The sensors were built on a silicon substrate prepared with 370 nm SiO_2 and 160 nm Si_3N_4 , and consist of a 215-nm-thick Ti/Pt heating element, 165-nm-thick TiW/Au electric contacts, and a 286-nm-thick SnO_x gas-sensitive layer.

The sensors were built on a silicon substrate made of SiO_2 (370 nm) and Si_3N_4 (160 nm). They consist of three layers:

- | | | |
|------------------------------|----------------|--------|
| ■ a heating element | Ti/Pt | 215 nm |
| ■ a set of electric contacts | TiW/Au | 165 nm |
| ■ a gas-sensitive layer | SnO_x | 286 nm |

Rational minds never do anything

Is the passive voice always suboptimal?

The passive voice can be effective, for example to use a certain topic as the sentence's subject, as part of a parallel or serial structure. Its use must not be judged in isolation but in context. (Used as a default, however, it is suboptimal.)

Can I write "the authors" instead of "we"?

Writing *the authors* is as accurate as writing *we* (except near a reference call, where it might seem to designate the authors of the document referred to). However, it is not equally concise.

Can I clarify the agent with a reference call?

Writing *It is believed [5]* in an attempt to mean *The authors of [5] believe* remains ambiguous. All the reference call indicates to the readers is that they can find more on this belief in [5].

Are there ineffective uses of the first person?

The first person is best used exclusively to refer to the authors for tasks, decisions, beliefs, etc. If used too often or in reference to other people, it loses impact through dilution. For example, it is unnecessary for objects of the document. Instead of writing *In this chapter, we present*, write about the chapter: *This chapter presents*. Also unnecessary in most documents is a *we* designating the reader and the author together, as in *We see in Equation 4 that ...*, a first person that ends up relegating the main information to a subordinate clause. (Such a collective *we* might be useful in tutorials, but it should then not be mixed with a *we* that means *the authors*.)

In an attempt to sound objective, many documents seem to imply that their authors never do anything. A task as simple as *We analyzed the data* typically sees its agent removed (*The data were analyzed*) and often its action verb nominalized (turned into a noun): *An analysis of the data was carried out*. Authors might have been taught in school, usually without rationale, never to write in the first person, or have figured it out from the documents they read. Powerful cultural influences quickly lead to myths: many authors now regard the first person as taboo.

More than a question of style, the widespread use of the passive voice raises an issue of accuracy—the concern here is the missing agent, not the voice or grammatical person, even if the three are linked. Admittedly, the agent may not matter or might be understood from context; in such cases, however, the outcome likely matters more than the action: *The analysis of the data indicated ...*. Often, though, knowing *who did it* makes a difference to readers, so a task almost always requires an explicit agent. As another example, the main clause *It is believed* is ambiguous: it may be interpreted equally well as *The authors believe* or as *The community believes*.

Overuse of the first person unpleasantly suggests that the authors are in fact writing about themselves, not about their topic. As an alternative, authors can focus on outcome, not action, and clarify their role with a first-person subject in a subordinate clause (*The option that we selected is ...*) or a possessive (*Our analysis suggested a new course of action: ...*), in particular as part of a parallel or serial structure. A collective is another option (*The team decided ...*).

In documents having a single author, a first person plural (*we*) not otherwise clarified lacks accuracy (the author and who else?). It is thus best replaced the first time it appears by an explicit expression: *My supervisor and I* or *My department*, for example.

Being clear, accurate, and concise

WORD CHOICE AFFECTS READABILITY: unusual words make a text unclear; uncalled-for passive voices make it inaccurate; wordy expressions make it longer than necessary. Typing or language mistakes are a major source of noise, too.

For clarity, use words with a known, unambiguous meaning for your audience. Thus, for specialists, use technical terms, but avoid jargon. For nonspecialists, prefer common words, used consistently; avoid synonyms. For non-native readers, beware of so-called false friends (*faux amis*) and of words with multiple meanings; avoid idioms and cultural references.

For accuracy, select each sentence's subject and verb carefully. Identify the agent whenever this agent matters to the readers, as with verbs that imply a human judgment or responsibility, such as *decide*, *believe*, or *recommend*. If you are the agent, consider using a first person. Express the action with a verb, not with a noun. By default, cast this verb in the active voice, even with an inanimate subject (*with this method, the volume is underestimated* → *this method underestimates the volume*).

For conciseness, use the shortest phrase allowing the desired level of clarity and accuracy. Avoid ineffective redundancies (*fewer in number*), wordy expressions (*in the event that* → *if*), and nominalizations (*perform an examination of* → *examine*). Replace frequently used phrases by acronyms, typeset these and these only in capitals, and introduce them systematically. To avoid repetitions, try combining closely related sentences.

The number of triangles that are used to describe a mesh can be scaled up or down. This can be achieved by performing edge collapses and vertex splits, respectively, as shown in Figure 6.

The number of triangles used to describe a mesh can be scaled up or down through edge collapses or vertex splits, respectively (Figure 6).

Conciseness is a second-draft optimization. On a first draft, concentrate on expressing messages clearly and accurately. Then look for opportunities to say the same in fewer words.

Fundamentals

- The name of the game
- The three laws of communication
- A thousand words, a thousand pictures
- Chains and magical numbers
- Trees, maps, and theorems

Effective written documents

- Planning the document
- Designing the document
- Drafting the document
- Formatting the document**
- Revising the document

Designing intuitive page layouts
Achieving simplicity and harmony

Effective oral presentations

- Planning the presentation
- Designing the presentation
- Creating the slides
- Delivering the presentation
- Answering questions

Effective graphical displays

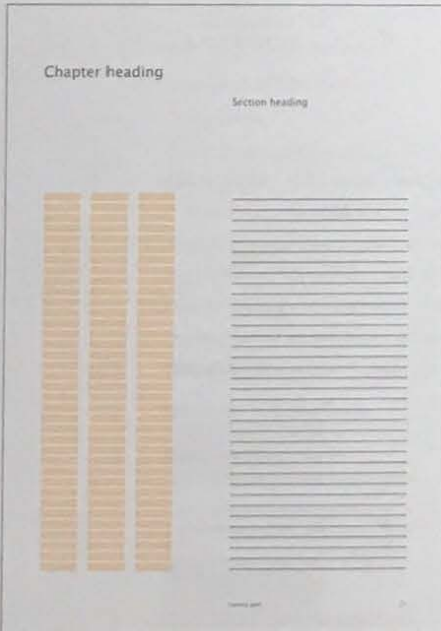
- Understanding pictures
- Planning the graph
- Designing the graph
- Constructing the graph
- Drafting the caption

Applications

- Effective instructions
- Effective electronic mail
- Effective Web sites
- Effective meeting reports
- Effective scientific posters

Formatting the document

The grid used for the right pages in this book, showing the column of text (right) and the grid subset for aligning the illustrations (left).



FORMATTING IS ABOUT STRUCTURE, not about aesthetics. An effective page layout conveys the structure visually, as a form of effective redundancy, yet without drawing attention onto itself, lest it create noise. Like careful phrasing, it goes unnoticed: readers take note of content, not of form.

Visual structure is primarily a matter of spatial arrangement (where items are) and secondarily a matter of visual aspect (how items look). Much can be achieved with spacing alone: the margins and the possible indentations, the line spacing, the space between paragraphs or between headings and text, and juxtaposition, clustering, or alignment of items on a page. Type variation, such as a change in face, size, weight, or color, should merely reinforce what is already conveyed by position.

In contrast to text, a page has two dimensions. A page layout can do more than break a text into a vertical sequence of lines: it can use the horizontal dimension of the page, too, to enrich the visual structure and thus provide additional entry points to the readers. For example, a narrower main column of text (as on this page) leaves space for illustrations and other items directly next to the text, thus opening up a second dimension.

Creating a page layout is more than tinkering with a text file. It requires coming up with a global picture of the final page. This picture can be described as a design grid: a set of points, lines, and areas that guide the positioning and dimensioning of all items consistently and harmoniously across the pages.

Aesthetics, of course, does matter—ugliness distracts. Still, it is best regarded as a consequence of effective formatting, not as a purpose in itself. Even beautiful pages should reveal the structure and thus invite the readers to process the texts, not dazzle the readers to the point that they not wish to read. More than beauty, formatted pages should possess elegance, a quality implying appropriateness, simplicity, and harmony.

Should drafting always precede formatting?

Logically, a text must have been drafted before it can be formatted, so *drafting* appears before *formatting* in the methodology proposed here. Still, the format might pre-exist or be designed before the text is drafted. Layout constraints, as on the length of texts, should be identified early, so the texts can be optimized accordingly.

Should I prefer boldface or italics for emphasis?

Boldface and italics have different properties. One can notice bold words merely by looking at the page, without actually reading any text. As a consequence, boldface should be reserved for that which readers must notice at a glance, that is, almost solely navigational information: headings, "pseudoheadings" such as *Warning* next to a warning (but not the warning itself), and identifiers such as *Figure 4* under a figure (but, again, not the caption itself). Bold words within a sentence make little sense: typically, they cannot be assigned meaning in isolation. Italicized words, by contrast, do not stand out at a glance but still provide enough emphasis to stand out when the sentence is being read.

How can I use color optimally?

Color in page layout is a dangerous temptation: use it whenever it helps convey the structure, not whenever you can. Use it redundantly, too: not everyone distinguishes colors equally well. Unless you master color design, use few colors, perhaps just one (besides black) in a few tints. Design the page in black and white first, then apply color in touches wherever it adds value. Of course, use this color scheme consistently.

Common shortcomings

As a common violation of the proximity principle, many reports formatted with direct-manipulation text processors show headings as close (or closer) to the previous paragraph as to the next paragraph, probably as a typing error, not as a design choice.



Ideally, any piece of text should be typeset in a way that keeps together what goes together. In practice, however, few can. Among these few are headings on more than one line, which can be cut by hand.



Finally, many documents leave little empty space on the page and thus afford little design freedom. Often, they attempt to make headings prominent by setting them in a much larger or heavier type, which looks cramped between paragraphs of text. In most cases, a heading need not be (much) larger than the body text; instead, it needs enough space.



Designing intuitive page layouts

PAGE LAYOUTS, like all visual displays, should be intuitive: they should use no arbitrary codes; rather, they must rely on natural principles of proximity, similarity, prominence, and visual sequence. An effective layout reveals the structure of a page at a glance, before the readers even start reading.

When designing a page or spread, group related items visually. Leave more distance between unrelated items than between related ones. For example, leave more space above a heading (between the heading and material preceding it) than below it. Display a caption closer to the corresponding figure or table than to other items on the page, such as other figures or tables. Avoid breaking a line of text between tightly connected items, as in cross-references (Figure 9) and values with units (4 km).

To be intuitive, be visually consistent: format identical items identically, similar items similarly, different items differently. For example, format all same-level headings in the same way (same typeface, type size, relative position, etc.). Conversely, format headings at different levels in visibly different ways, not only in the document, but also in the table of contents. Strive for visual consistency within each document separately but also across comparable documents, as within a collection.

To indicate hierarchy, display more prominently those items that rank higher or that are more important. For example, leave more space around the higher-level headings, set them in a larger or heavier type, or combine several such features. Prominence stems from contrast, so beware of visual inflation: the more items you emphasize, the less they will stand out.

Finally, make sure that each page guides the readers visually along a useful reading sequence or, alternatively, that it gives a clear picture of possible entry points. Where will readers likely look first? Why? Is this where you want them to start? Where do you want them to go next? Is the transition clear?

- 1 Introduction
- 2 Materials and methods
- 3 Experimental results
- 3.1 Influence of the temperature
- 3.2 Influence of the pressure
- 3.3 Transient phenomena
- 4 Discussion
- 5 Conclusion and perspectives



- 1 Introduction
- 2 Materials and methods
- 3 Experimental results
- 3.1 Influence of the temperature
- 3.2 Influence of the pressure
- 3.3 Transient phenomena
- 4 Discussion
- 5 Conclusion and perspectives

What kind of typeface should I use?

The legibility of text typefaces has more to do with their familiarity than with their features. Select a design that is familiar to the audience or one that is part of your corporate identity. Optimize text legibility by adjusting type size and, accordingly, line length and line spacing.

May I mix several typefaces?

When mixing typefaces, strive for consistency and harmony. Give each face a function, one that will be readily understood by the readers. Beware of creating a visual cacophony, though; when possible, select faces within one family, as these were designed to work well together.

Is there an optimal text width?

As a first approximation, and unless the page design dictates otherwise, make the text width twice the length of a lowercase Latin alphabet set in the chosen typeface at the chosen size. (Narrower text allows a slightly faster reading but requires more line breaks, some of which may hinder readability, defeating the purpose.)

Should I justify text?

Justifying a text results in sharper text blocks but less natural word spacing and line breaks. Advantage and drawback should be balanced. If markedly rectangular paragraphs contribute to a clearer overall page structure, for example because they align nicely with the other items on the page, you may prefer to justify the text. For a less formal look, leave the text unjustified.

Plain text is most legible

To be recognized rapidly, words must look familiar. They are then identified globally, that is, as a whole rather than letter by letter, as evidenced by the fact we can still largely make out sloppy handwriting. Any visual effect modifying a word's overall shape unavoidably renders this word harder to recognize. The five effects below are common shortcomings.

physics

Avoid unusual typefaces, especially for body text. Such faces may help create a desired atmosphere, as for a commercial ad or for a wedding invitation, but they are not designed for legibility. (Of course, what is an unusual face depends on the audience.)

physics

Avoid setting text in uppercase (except acronyms). By giving them a more similar shape, all uppercase makes the words harder to differentiate. Moreover, by suggesting an acronym, it may mislead readers. In e-mail, it is considered the equivalent of shouting.

PHYSICS

Avoid underlining text or crossing the descenders (the part of the letter extending below the baseline). Besides affecting the silhouette of the whole word, underlining renders the individual letters harder to recognize (an underlined y resembles a v, etc.).

physics

Avoid setting text on top of a background picture or any nonuniform background (a solid color is OK). A background may make the page more attractive, but it always reduces the legibility of overlaid text. Text and illustrations are best set next to each other.

physics

Avoid packing the lines tightly. Tight lines interfere with one another visually. When long, they make it hard for readers to find the start of the next line. For its overall shape to be readily perceived, a word must be surrounded by enough unoccupied space. Accordingly, plan for sufficient margins around text, too, in particular inside a box: let your text breathe.

physics
biology

physics

Achieving simplicity and harmony

<i>Train</i>	<i>Paris</i>	<i>London</i>
9005	06:43	07:58
9007	07:13	08:28
9009	07:43	08:59
9011	08:07	09:34
9015	09:07	10:36
9019	10:13	11:28
9027	12:13	13:28
9031	13:01	14:34
9039	15:13	16:36
9043	16:13	17:34
9047	17:13	18:34
9049	17:43	18:59
9051	18:13	19:34
⋮	⋮	⋮



<i>Train</i>	<i>Paris</i>	<i>London</i>
9005	06:43	07:58
9007	07:13	08:28
9009	07:43	08:59
9011	08:07	09:34
9015	09:07	10:36
9019	10:13	11:28
9027	12:13	13:28
9031	13:01	14:34
9039	15:13	16:36
9043	16:13	17:34
9047	17:13	18:34
9049	17:43	18:59
9051	18:13	19:34
⋮	⋮	⋮

THE SEEMINGLY ENDLESS formatting possibilities offered by software applications are a curse as much as a blessing: used indiscriminately (too often because they are available rather than because they are useful), they introduce noise. In contrast, optimal formats are visually concise: they reveal the structure by adding the least amount of ink to the page. Prevention is better than cure. In visual design, a healthy dose of self-restraint favors simplicity, consistency, and harmony.

Consider developing a first design under strict, self-imposed constraints, on the basis of relative position on the page only, thus using a single typeface at one or perhaps two type sizes, no variation such as boldface or italic, no color besides black. Once this design satisfactorily conveys the structure visually, you can relax any of the constraints when doing so adds value. If you do, carry the change consistently across the document.

When setting text, use typographical devices sparingly, if only to maintain their impact when used; use space above all else. To make a piece of text stand out, just set it apart: increase its distance from other items, thus surrounding it with space. For example, leave enough space above and below headings. For more contrast, consider increasing the type size or weight slightly. To emphasize words within a paragraph, use italics. As a rule, avoid other typesetting effects, such as underlining.

For harmony, coordinate both the dimension and the position of all the items on the page (paragraphs, figures, tables, etc.), for example by adjusting them to an underlying design grid. For languages that read from left to right, prefer left-aligned blocks of text (not necessarily at the left margin of the page) and consider aligning the other items on their left edge, too, perhaps on the same alignment axis as the paragraphs of text.

Finally, remember that the format should work redundantly: structure and emphasis must already be conveyed by the text.

Fundamentals

- The name of the game
- The three laws of communication
- A thousand words, a thousand pictures
- Chains and magical numbers
- Trees, maps, and theorems

Effective written documents

- Planning the document
- Designing the document
- Drafting the document
- Formatting the document
- Revising the document**

- Testing the document
- Improving the document
- Reviewing documents of others

Effective oral presentations

- Planning the presentation
- Designing the presentation
- Creating the slides
- Delivering the presentation
- Answering questions


Effective graphical displays

- Understanding pictures
- Planning the graph
- Designing the graph
- Constructing the graph
- Drafting the caption

Applications

- Effective instructions
- Effective electronic mail
- Effective Web sites
- Effective meeting reports
- Effective scientific posters

Revising the document



DOCUMENTS ARE SELDOM OPTIMAL as a very first draft. Hence they must be revised more or less intensively, according to the required or desired level of quality. While a routine letter may need no more than proofreading by its author, a more important document usually benefits from a second opinion or perhaps from a more formal test.

Key to optimizing the document, revising is an integral part of the writing process, not "something you do if you have time." Accordingly, it must be taken into account in the time budget. It is not a last step as much as a decision to iterate and refine the other steps: drafting, formatting, and if needed designing or even planning. Each step will need all the fewer iterations when the preceding steps have been completed with care: if the planning must be reconsidered, so must, most probably, the designing, the drafting, and the formatting, in that order.

The need to iterate (part of) the writing process always rests on the answer to the question, is the document good enough? Past a certain stage, this question is a difficult one to answer for the authors themselves: how can they assess the clarity of what they have written unless they test it on someone else? Feedback from such test readers is an invaluable contribution to the optimization process: we must treasure it, not fear it.

After testing it in one way or another, we can start improving the document. More than a specific set of skills, this step requires a positive attitude, including the willingness to listen to test readers (even if discerningly), to change the document accordingly, and eventually to proofread everything carefully. Proofreading, a way to filter out noise, should not be neglected, even for such quick documents as short electronic messages.

Finally, knowing what feedback is useful to us should help us provide more useful feedback when we review the documents of others—a more difficult task than it might seem at first.

Common shortcomings

Can I test the document myself?

Testing a document suggests “discovering” it (looking at it and reading it for the first time), so you cannot quite test your own document as someone else can. You can, however, take a fresh look at it by setting it aside for a while. Obviously, the longer you can wait, the better. In any case, let it rest at least a night: sleeping helps you distance yourself from the situation.

Why should I not defend my text when test readers attack it?

What seems like an attack should not be taken as one. For example, a reader who points out *this is unclear* really means *I found this unclear*; this is an incontrovertible statement—arguing back that *this is perfectly clear* makes no sense. Feel free, of course, to think so and to decide not to modify that part of the text, but arguing will only make readers less willing to provide feedback in the future. Moreover, if one reader found a piece of text unclear, so may others. Therefore, finding out what exactly this reader found unclear is a more constructive approach.

What if the test reader is my boss?

Saying that the author has final responsibility on the document is a first approximation only: the author's organization is responsible, too, so disregarding comments you disagree with may be harder when they come from the boss. If you feel you must argue, argue about ideas, not about words. If your supervisor dislikes a phrase and you dislike his or her rewording, see if you disagree on content. If not, search for a third wording that you can both agree on.

Revising one's document is usually a difficult task for everyone: hurried writers regard it as a waste of time; careful writers put much pride in their draft and hence feel an emotional barrier to amending it. As a result, this fifth step is prone to imperfections.

For documents not tested on or reviewed by others, the most common (and most severe) shortcoming is skipping the revising step entirely and regarding as a finished document what is in fact still a draft. This attitude is all too frequent for electronic mail, sent as soon as it is drafted, without even as much as undergoing proofreading. Besides carelessness, a lack of revision may come from a tight deadline, which might be caused by poor time management. Writing at the last minute leaves little time not only to go through several iterations but also to leave the draft aside for a while and take a fresh look at it.

Too often, documents that get tested or reviewed go to the wrong people, typically direct colleagues, who know much about the content and the context. True, these readers can usefully check the accuracy of the content; however, they cannot easily detect more fundamental issues of selection and structure: is it the right content, presented in the right order? Anxious authors tend to tell test readers too much about the document, for example by pointing out expected difficulties to them, thus biasing the test (*I have structured the text in [such and such] way. Is this apparent as you read?—well, it will be now*).

Of course, the most frequent revising shortcoming is a defensive inclination on the part of the authors, who then explain or justify what they have written, instead of listening to what test readers have to say. This understandable but unproductive resistance to change also makes authors reluctant to question their document as a whole and hence to consider making in-depth changes in structure or in writing, instead of quick fixes. Revising requires dedication.

Testing the document

DIFFERENT QUALITIES OF A DOCUMENT (clarity, accuracy, correctness, etc.) are best tested on different audiences. To test the clarity of your document, select nonspecialized, secondary readers, that is, people representative of the least specialized, most secondary among your anticipated readers. To test the accuracy of your document, by contrast, prefer specialists. To improve language correctness, select readers who are proficient in the language of your document—ideally native speakers, though being a native speaker is not enough.

When entrusting a copy of your document to readers for a test, clarify what you expect from them and by when, but refrain from telling them anything that might bias the test. Do say whether you are looking for comments on the form (clarity), on the content (accuracy), or on the language (correctness). Specify whether you prefer high-level feedback (for example, because you are still likely to revise the document in depth) or a more detailed analysis. Agree on a deadline, whether it is dictated by schedule (*This must go out on Tuesday, so I need your feedback by Monday*) or negotiated with the test readers (*By when do you think you can get back to me about this?*).

When receiving your commented copy from the test readers, be receptive, not defensive. Whenever practical, sit together with each test reader to review his or her feedback. Make sure you understand each annotation's meaning and importance. Ask questions on the parts that received few or no comments: the readers may have an opinion that they did not write down. If the "comment" is in fact a rewording, ask what was wrong with the original wording. Should you disagree with a remark or rewording, focus on understanding the reader's viewpoint, not on defending your own. You, as author, can decide later whether to change the sentence as suggested by this reader or to keep it as is. Defending your text defeats the purpose of the test, which is to collect readers' reactions accurately, so you can make informed decisions about what to change.

Getting help to improve our writing

Can I proofread the document myself?

Proofreading one's own document is difficult. When expecting a certain word in a certain place, one uses in fact very little visual information to recognize the word—too little, most often, to ascertain that the word is spelled correctly. As with testing the document, setting it aside for a while certainly helps. So does changing its visual appearance, such as setting the text in a different typeface, at a different type size, or with a different line spacing or line width.

Should I use a grammar checker?

Grammar checkers are much less reliable than spell-checkers are: they are limited to flagging potential problems on the basis of symptoms and letting the author decide whether there is indeed a grammatical error and how to fix it. While they may be useful to those who master the language, they tend to be confusing, if not downright misleading, for non-native writers, who may be unable to recognize false alarms and who thus tend to trust the checker blindly.

Why should I proofread a manuscript I submit to a journal? Surely they have a proofreader.


Texts that will be further edited, as is usually the case for scientific papers being submitted to a journal, will most likely be proofread later by a qualified person. If the article must first be refereed, however, errors might needlessly antagonize the referees and also prevent them from paying due attention to your messages. Polished documents always reflect positively on the authors and, indirectly, on their work. Careless documents suggest careless research.

Having a skilled editor improve our documents is certainly comfortable but not necessarily optimal. Such an editor is unlikely to know as well as we do what messages or ideas we are trying to get across. If we hand him or her an unpolished draft (which is tempting when we know it will still be improved), we risk misunderstandings and inappropriate edits. Moreover, and unless we take time with the editor to review the changes and the reasons behind them, we will not learn much from the editing process and thus not improve our writing for the situations for which we will not have the luxury of an editor.

A way to systematize how we improve documents and to sharpen our writing skills by the same token is to maintain a list of the errors we tend to make, as identified in the feedback we collect from others. We can then look for each of these errors in turn, using the search capabilities of our text processor. Such a search typically helps us identify in isolation errors we understand but nonetheless overlooked. If we tend to write *insure* when we mean *ensure*, for example, we can search for instances of *insure* and verify whether they are what we meant to write. This process will not only improve our documents but also help automatize the correct use of the term: after a while, the documents we write will probably no longer contain any instance of a misused *insure*.

Besides misused words, easy-to-search-for errors include the following ones: inaccurate expressions such as *a number of* (to be replaced by *the number*), verbose expressions such as *at this point in time* to mean *now*, misspellings that pass spell-checkers such as *it's* instead of *its*, false friends such as *actual*, having a different meaning in English and Spanish, and consistency of terminology, such as choosing between *user guide*, *user's guide*, and *user manual*. With some ingenuity, you can even devise patterns to search for suboptimal sentences as with *it must be pointed out that*, *was carried out*, or *is believed*.

Improving the document



IMPROVING A DOCUMENT proceeds in two steps: editing it iteratively, among other things on the basis of the feedback provided by the test readers; then, proofreading it carefully. The first step suggests looking for ways to improve further what is still unfinished work. In contrast, the second means inspecting what is otherwise considered a finished document. While editing is normally an activity for the authors, if only for a consistent style, proofreading can be entrusted to others.

When editing a written document, make major changes first and minor ones afterwards. Indeed, major changes may well get rid of minor problems, whereas the reverse is seldom true. To identify possible major changes, work on a paper printout, not at your computer screen. Printed pages provide not only a better overview of the document but also a flexible medium for fast marking, such as crossing out words or sentences, indicating moves with arrows, or making notes in the margin. Text processors on computer screens, in contrast, encourage local revisions in the phrasing and the formatting of the text.

After editing the document iteratively until it seems optimal, proofread it in several passes. In each pass, check one aspect of the document: content (such as numerical data), spelling, formatting, cross-references, etc. Remember to proofread not only the full paragraphs but also the isolated text items, such as headings or captions. Use any means available to you, such as a checklist, a spell-checker, or help from colleagues, but do not trust technology blindly: there are spelling errors that a spell-checker, however smart, will be unable to spot.

What fraction of a writing project's time budget should be reserved for editing and proofreading is difficult to determine: it depends of course on the quality of the draft to be revised. Anyway, it should not be underestimated: it may take as long, if not several times as long, as developing a solid first draft, depending on the quality level required of the final document.

Common shortcomings

How can I clarify different types of feedback?

Colors may help qualify feedback. Still, black and blue, while neutral, typically lack contrast with the printed text. Red is better contrasted but often connotes error or perhaps judgment and might then needlessly antagonize authors; green is usually a less threatening alternative. As a useful color scheme, you might thus use red for language errors or mandatory changes, green for suggested changes, and black pencil (perhaps with text highlighter) for comments.

Some authors are defensive or easily offended. How can I provide tactful feedback to them?

Strive not to seem threatening or judgmental: show that your comments are meant to help. To suggest that you are "on the authors' side," try phrasing suggestions with a *we*, not a *you*. For example, instead of writing *Could you turn this passive voice into an active one?*, consider writing *Could we turn ...*. Similarly, and to avoid antagonizing authors who do not take orders, try pointing out possible issues as a question, not as a statement (or, worse, as a command). When spotting an apparent lack of consistency, for example, instead of ordering *Be consistent*, you may ask them, *Is the difference intentional?* They probably prefer to conclude themselves as to the inconsistency than to be told by you.

Is reviewing the same as refereeing?

Refereeing an article implies a judgment about whether to accept it for publication in a journal. It thus involves more than reviewing the article. Referees ascertain its suitability for the journal and—ideally—also help its author(s) improve it.

To contribute effectively to improving a document and, ideally, to help authors become better writers, reviewers must do more than indicate corrections: they must strive to provide tactful, helpful feedback, that is, feedback that will be accepted by authors and from which authors can draw valuable lessons. An oral discussion helps get a constructive attitude across, but it is not always practical. In any case, oral feedback typically builds on written feedback, so this feedback should be phrased most carefully.

Reviewers, unfortunately, easily confuse reviewing with rewriting, much to the frustration of authors. Unless they have a poor command of the language, authors indeed expect comments about their text, not replacement text, which they feel is not theirs. When they see a sentence crossed out and replaced by another, they are left to think, *what was wrong with my original sentence?* They expect reviewers to point out any issues rather than offer solutions to unidentified problems. (As it happens, however, reviewers often feel something is wrong but lack the analytical framework to put the issue in words.)

When they leave it to authors to correct a problem, reviewers are sometimes unclear in their comments. Many authors have been puzzled by a squiggly line next to a paragraph, without further clarification—or with a question mark or a vague *To be rewritten*. Authors can rewrite the paragraph a dozen times, but how could they improve it if they are not told what exactly the reviewer finds suboptimal about it?

How blunt reviewers can be in providing feedback depends largely on national and corporate culture. Still, three types of criticism are needlessly negative: judgment (*this sounds silly*), reproach (*you should have used the active voice here*), and interpretation (*obviously, you are trying to bluff your audience*). They are best replaced by constructive alternatives, more likely to elicit learning, no matter the culture.

Reviewing documents of others

REVIEWING A DOCUMENT so as to provide helpful feedback to its authors is no easy task. It requires careful reading to identify problematic passages, sharp analytical thinking to pinpoint the exact shortcoming in each of these passages, and a fair amount of tact to communicate these shortcomings to the authors in a constructive way, for an optimal process.

As usual, refrain from rushing into the task at hand: plan first. When you are requested to review someone else's document, agree on the deliverable: make sure both you and the authors or review coordinator are clear on the type of review needed, the state of the document (partial or complete, draft level), and the deadline. Next, choose an appropriate place and time, and arrange not to be disturbed during a long enough period. As for revising your own text, work preferably on a paper copy.

When reviewing the document, center yourself on the purpose that was agreed upon, such as clarity, accuracy, or correctness. Should this purpose be multiple, review one aspect at a time, focusing on content first. Typos are usually more conspicuous than reasoning flaws but also less important. When reviewing for content, try to ignore any other potential sources of noise.

In your comments to the authors, strive to help, not to judge. First, provide a global assessment, to place further comments in proper perspective. As a rule, point out the weaknesses, to prompt improvements, but also the strengths, to increase the authors' willingness to revise the document and to learn. Make comments explicit in both their object and their nature: content versus form, incorrect use versus inelegant wording, and (if you have authority over the authors) required changes versus mere suggestions. Moreover, and whenever practical, consider an oral discussion to clarify your written comments and, if needed, to look with the authors for suitable solutions. As a last step, make sure everyone is clear on what to do next with the document (revision, second review, approval, etc.).