Fundamentals

ALL FORMS OF EFFECTIVE COMMUNICATION—written, oral, or graphical—build on the same principles, addressing fundamental concerns of purpose, content, and form. What are we trying to achieve by engaging in communication? What must we then write, say, or draw, and how must we write, say, or draw it to reach our purpose? These are the key questions to analyze existing documents, presentations, or displays, and to create effective new ones.

This first part discusses the fundamentals of communication. After establishing what *effective communication* designates, it proposes a set of three laws that forms the very foundation of the further guidelines. As a more substantial dichotomy than the usual opposition between oral and written channels, it discusses the specificities of verbal and nonverbal codings. Finally, it examines the effectiveness of essential structures in terms of the number, hierarchy, and sequence of elements.

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

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

The name of the game

EFFECTIVE COMMUNICATION is getting messages across. Thus it implies someone else: it is about an audience, and it suggests that we get this audience to understand something. To ensure that they understand it, we must first get them to pay attention. In turn, getting them to understand is usually nothing but a means to an end: we may want them to remember the material communicated, be convinced of it, or, ultimately, act or at least be able to act on the basis of it.

Information

A concentration of 175 µg per m^3 has been observed in urban areas

A message

The concentration in urban areas (175 µg/m^3) is unacceptably high

A message differs from raw information in that it presents "intelligent added value", that is, something to understand about the information. A message interprets the information for a specific audience and for a specific purpose. It conveys the *so what*, whereas information merely conveys the *what*. A message is to information what conclusions are to findings. Because it makes a statement, it requires a complete sentence.

A *what* caption (a noun phrase)
Evolution of sales over the years

A *so what* caption (a statement)
Sales dropped by 40% last year

To communicate effectively, we must thus identify messages. Conveying information only is usually not enough, as it leaves the audience with the question, so what? We must moreover recognize and seize opportunities to get the messages across, for example in the captions of figures or in the titles of slides.

Often, the messages to be conveyed are numerous or complex, and the situation carries constraints. Among these are space (such as a four-page limit on a paper), time (a 15-minute limit on a presentation), and audience (background, language, etc.). Not every hindrance is a true constraint, though: for example, a suboptimal room can often be rearranged, at least to a point.

Get your audience to

- **pay attention to,**
- **understand,**
- **(be able to) act upon**

a maximum of messages, given constraints.

Effective communication is optimization under constraints: we must maximize, not what we write, say, or draw, but how much our audience gets out of our documents, presentations, and displays, in quantity or in quality—with a purpose in mind and under certain constraints. Because of these constraints, we cannot hope to be perfect. We can, however, be optimal.

Fundamentals

- The name of the game
- **The three laws of communication**
  - Adapting to the audience
  - Maximizing the signal-to-noise ratio
  - Using effective redundancy
- 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

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

The three laws of communication

THREE SIMPLE YET SOLID PRINCIPLES are all we need to optimize virtually any instance of communication. These three "laws of communication" can be derived with a simple model of one-way communication, embodying the idea of getting messages across optimally to our audience.

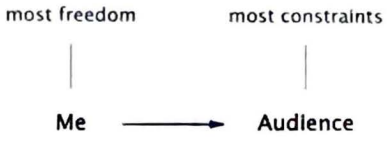

First law **Adapt to your audience**

To optimize under constraints, we must first identify what is and what is not under our control, and concentrate on what is. To a point, we cannot select our audience: we must take them as they come. Still, we can decide what to tell them and how. To optimize our communication, we must thus adapt to them.

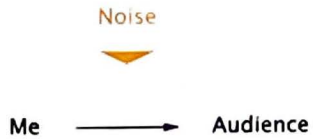

Second law **Maximize the signal-to-noise ratio**

The simple model above is ideal: it suggests that information or messages sent from one side reach the other side intact. In practice, information may suffer losses, because of noise. To prevent losses, we must filter out the noise; alternatively, we can enhance the signal so it can withstand the noise better.

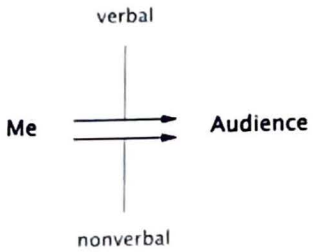

Third law **Use effective redundancy**

The second law is limited to prevention. When noise cannot be anticipated, it cannot be filtered out, so it results in losses. To compensate for the losses, we can tell things several times, by repeating the message or by replicating it across channels, preferably in different ways, such as verbally and nonverbally.

Can I not select my audience at all?

Early on, the audience is not necessarily given, indeed: a scientist deciding to which journal to send a paper is selecting his or her audience. Still, the range of options is often limited and, once a journal has been chosen, the audience becomes a given: the scientist can hardly select, within this audience, who may read the paper.

Can I not change or influence my audience?

You can most certainly influence your audience, such as increase their motivation or provide any prerequisite knowledge they might need. In doing so, you are already communicating with them and, essentially, adapting to them. In other words, the audience can be regarded as given in that they cannot be changed a priori. Influencing them requires adapting to them.

Why should I always be the one adapting?

You should be the one adapting to the extent that you are the one with a purpose—that is, that you want something from your audience. Much like being customer-minded in business or being user-friendly in software development, adapting to one's audience is really a question of effectiveness more than one of selflessness.

Is the audience never to be blamed, then?

Blaming the audience may help us feel better but seldom gets us anywhere, unless perhaps if blames can influence the audience positively. A more purpose-oriented approach is to regard their shortcomings as constraints—and adapt.

Real-world audiences know less

Adapting to the audience in a professional context is all the more difficult because practice in school usually develops the wrong communication reflexes. Real-world audiences and purposes differ markedly from those that students were long confronted with.

Students formally write and speak to demonstrate their mastery of a well-defined body of knowledge and, ultimately, to be graded. For such a purpose, their only relevant audience is the course instructor: a single, clearly identified person, who is normally more knowledgeable than they are about the topic and morally obliged to read their entire document or attend their entire presentation, however boring. Accordingly, successful strategies to good grades may involve including as much material as possible (especially when in doubt as to what the instructor will give extra credit for) or showing off with jargon.

Professionals, in contrast, formally write and speak to get their audiences to pay attention, understand, and (be able to) act. Such real-world audiences are unpredictably multiple (especially for documents), almost always less knowledgeable about the topic, and highly selective about what they read or attend. They have little patience with writers and, especially, speakers attempting to demonstrate the breadth or depth of their knowledge, often at the expense of the clarity or the conciseness of their discourse.

Unsurprisingly, the most common failure to adapt to one's audience, then, is to present information that is too technical or too little relevant to them. According to a common myth in academic research, a presentation should have one third that everyone in the audience understands, one third that some understand, and one third that no one understands. What can one gain with such an approach, though? Audiences have infinitely more respect for speakers who can explain complex matters in simple ways and thus give new insights. Still, the myth endures.

Adapting to the audience

EFFECTIVE COMMUNICATION always requires motivation. If we want our audience to pay attention to, understand, and act upon our messages, we are the ones who should make the effort. That is, we must adapt to them, not expect them to adapt to us. Should they be willing and able to adapt, too, so much the better, but we have no cause to assume they will.

The first law, *adapt to your audience*, is one of empowerment: it implies that we are responsible for the success of our acts of communication. If our audiences fail to get the messages, it is our problem, not theirs, as we have not reached our goal. Blaming them makes little sense: it hardly helps us optimize. From our perspective, the degrees of freedom are on our side.

Teachers who stick to what they had planned to do regardless of whether the students pay attention, understand, or develop the required competencies are not adapting to their audience. Unsurprisingly, such teachers often blame it all on the students, too.

Adapting to our audience is normally a spontaneous attitude in our private life. For example, we do not address children the way we address adults: we recognize the need to adapt. It is far less spontaneous an attitude in our professional life, in which we tend to regard the others as similar to ourselves.

Adapting means putting ourselves in the shoes of the audience, anticipating their situation, their needs, their expectations, etc. It implies structuring the story along their line of reasoning, not ours, and recognizing the constraints they might bring: their familiarity with the topic, their mastery of the language, the time they can free for us, etc. Whenever we are not taking a certain constraint into account somehow, we fail to adapt.

Imagine that a foreigner asks you for directions and that the only language you have in common with him is English, of which he has little command. Suppose that he did not get your first explanation. Adapting to him (assuming that you are motivated) might involve making gestures, sketching a map, speaking more slowly, pronouncing more clearly, using a simpler vocabulary and a simpler syntax—or perhaps accompanying him to where he must go.

Finally, adapting to the audience suggests that, if one strategy does not work, we try a different one. If the audience failed to get the message, merely repeating it is unlikely to help: we must change the code or the channel. As the saying goes, if we do what we already did, we will get what we already got. Still, adapting to our audience does not mean losing track of our purpose. On the contrary, it means doing what it takes to get the audience to (be able to) do what we want them to do.

Is noise always bad? Can it not be used to regain the attention of the audience?

Noise is undesirable by definition. If something helps you reach the purpose you have in mind, such as by getting the attention of the audience, you can best regard it as signal, not as noise.

Still, before introducing or tolerating anything that might attract attention, consider whether the device does not distract more than it helps, such as by stealing the attention away from you or reflecting badly on your professional image. When they are handled with a touch of humor, minor mishaps in an oral presentation can help the speaker build rapport with the audience, yet the presentation would be more impressive without them. Similarly, conspicuous clothes or jewelry can easily overshadow the message. Attendees later referring to one of the speakers as *you know, the one with the dark red jacket* remember the clothes more than the person.

How can I "increase the signal" beyond merely speaking louder?

Increasing the volume in an oral presentation or perhaps the font size in a written document seem obvious applications of the second law, as is making the data lines thicker than the axes, tick marks, or grid lines in a graphical display.

More broadly, *increasing the signal* may mean conveying stronger messages. When we do not master the language, when the transmission is unusually poor, or when the audience is tired or otherwise less attentive, we may well have to be blunt and not attempt to be too subtle, as subtleties will probably not survive the noise. While not ideal, it would be optimal in this case.

Identifying sources of noise

Noise comes in many forms and from many sources. In the case of a formal presentation, for example, the noise sources that most people readily think of are the audience and the environment: attendees chatting among themselves or coming in and out, mobile phones going off, noisy air conditioning, unreliable equipment (microphone, projector, etc.), building works in progress outside the room, etc. The noise source that these people forget at first is the speaker himself or herself. Noise produced by a speaker is typically more distracting than that coming from other sources, because it is part of what the audience is supposed to pay attention to. Just because it comes from the speaker, however, it can more easily be controlled than other sources.

Noise produced by speakers in oral presentations shows more particularly in two components: slides and delivery. Busy slides compete with the speaker for the attention of the audience, and flashy slides draw attention to themselves, not to their content. Delivery noise includes imperfect pronunciation, filler words, unnecessary gestures or mannerisms, and so on, all the way to inappropriate dress code.

Noise in documents is whatever prompts readers to stop thinking about content and start thinking about form (or perhaps about irrelevant content). Examples are an unclear structure of the document, intricate sentences, unusual or superfluous words, spelling mistakes, and distracting visual elements in the figures, the typography, or the page layout.

Noise in graphical displays includes the many forms of data distortion, as with inappropriate graph types, and all "unnecessary ink", that is, visual elements that can be erased without loss of clarity or accuracy: decorative third dimensions or gradients of color and overabundant tick marks or grid lines in graphs, irrelevant backgrounds or objects in photographs (to be removed ahead of time or cropped out), etc.

Maximizing the signal-to-noise ratio

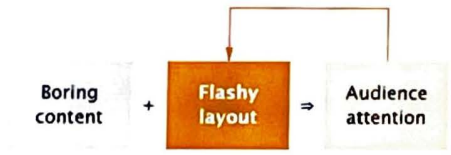

A frequent yet hopeless attempt salvaging uninteresting content with an "interesting" page layout. Will the document thus produced get the attention of the audience? Yes, it will—on the page layout, with little transfer to the content, if any. The flashy layout is noise.

NOTHING IS NEUTRAL in communication. The audience indeed sees and hears everything, so everything matters. What does not help get the message across detracts from it by needlessly mobilizing the audience's intellectual resources, even if for a short time. By definition, it is noise. Noise is thus more than unwanted sound: it is anything that can distract from the message (the signal) by drawing attention onto itself.

Noise can be a major impediment to effective communication. At best, it just stretches the attention span of the audience. At worst, it takes their attention entirely away from the content. As an example, typographical errors in a written document or filler words in an oral presentation can be most distracting: audience members may well find themselves on the lookout for the next typo or next *um* rather than for the next message. In graphical communication, noise easily shifts the attention from content to technology: when readers start wondering what software produced a graph instead of what experiment produced the data, they are most probably missing the point.

The second law, *maximize the signal-to-noise ratio* (or *ratios*), is all about contrast between what helps and what hinders an act of communication. The ratio between signal and noise matters more than signal or noise alone. To a point, we can thus tolerate continuous background noise, which we notice only when it varies suddenly, for example when it goes away.

Clearly, the most satisfying approach to contrast is reducing or eliminating noise: breaking the silence in a whisper is far more effective than covering the noise in a shout. Recognizing that nothing is neutral, we should thus question the relevance of anything we plan to include: words in a written document, gestures in an oral presentation, lines in a graphical display. To optimize a text or an image, we may do better to suppress, not add. By removing every unneeded drop of ink, we ensure that the audience pays attention to nothing but the message.

Identifying possible codings

Is redundancy the same as repetition?

Repetition suggests a signal at different times on one channel, as when a speaker previews his or her main points before developing them. Redundancy can be just that, but it can also be a signal across different channels at one time, as when a speaker illustrates the presentation with a slide show. Either way can be effective.

Why insist on calling it "effective" redundancy?

Not all redundancy is effective. For example, superfluous words as in *added bonus* or *oval in shape* add nothing. What is much worse still, multiple channels competing with one another, such as text-heavy slides accompanying a talk, are more harmful than helpful: each channel is indeed a source of noise for the other one(s).

Is a channel the same as a coding?

The term *channel* (or, equivalently, *medium*) refers to perception by the senses; in contrast, *coding* refers more to processing by the mind. (At times, the boundary is somewhat blurred.) For example, *paper* conveying written words or *air* conveying oral words (as sound waves) are channels, and *text* or *pictures* are codings.

What matters most for effective redundancy is codings. In first approximation, text is text, whether it is seen through the eyes or heard through the ears. A different coding, such as a picture, would be a more useful redundancy than a similar coding in a different channel, such as a second stream of text. Still, channels have their importance, in particular in terms of the nature of the noise they are subject to.

The three laws are in order of decreasing priority. To prevent losses, the first measure is to filter out the noise and, if deemed useful, increase the signal (adapting its maximum intensity to the audience). If we could filter out all noise, we would not need redundancy, at least not to compensate for losses, because there would not be any. Alas, some noise is beyond our control: attendees at a presentation may be preoccupied with something else, readers of a document may be interrupted by a phone call, etc. Using several codings is thus usually desirable.

Because nothing is neutral, most everything can be regarded as a coding, that is, as a potential source of noise if left uncontrolled and a potential signal if used well. Whether or not we like it, the clothes we wear always say something about who we are, for example. While we may decide not to worry unduly about the possible statement we thus make because it has far less impact than other codings, we may want to scrutinize our dress code for noise.

Devices we can usefully regard as distinct codings in written documents include the text itself, the set of headings in the text and in the table of contents, the page layout (revealing the structure visually), and tables or figures, all of which can be optimized.

In oral presentations, codings include most of all the verbal, vocal, and visual delivery (all three being powerful devices toward convincing an audience), possibly supported by slides or printed handouts. Just because the nonverbal ones (vocal and visual) are intuitive does not mean they cannot be managed: we can thus learn to amplify our intonation, quiet body noise, or make eye contact with the audience.

Graphical displays, too, can be seen as including more than one coding. They might convey meaning through relative lengths, positions along a scale, shapes, colors, explanatory labels, captions, etc.

Using effective redundancy

TELLING THINGS ONCE is often not enough: redundancy helps restore messages damaged by noise. It should not, however, introduce noise itself, that is, distract the audience, such as when concurrent channels compete with one another. Effective redundancy, therefore, gets a given message across several times, but coded in complementary, compatible ways.

The stop sign conveys meaning through shape, color, and label: it is the only octagonal sign, one of only two signs to be solid red (the other is the wrong-way sign), and the only one labeled "STOP" in most countries. It also comes with a white line across the lane as yet another way to mean stop.

Effective redundancy works in two ways: one is compensation, the other, collaboration. First, each coding gives the audience a chance to understand the message. Motorists, for example, can identify a stop sign in three ways: color, label, and shape. If they cannot distinguish the color, they can read the label "STOP". If this label is hidden by mud or snow (or if they see the sign from the back), they can still recognize it by shape. By giving several chances, effective redundancy helps address inhomogeneous audiences. Second, all codings work together in synergy: here, color, label, and shape, when all identified, complement one another for a faster recognition of the sign.

What makes a different coding is partly a view of the mind. Though they are both verbal codings on the same medium, the text and the headings within a document can be regarded as distinct codings, used for distinct purposes. When looking for a specific part of the document, we are thus likely to flip through the pages and look at the headings but not at the text. Conversely, when we have decided to read the full document linearly, we typically read the text but skip all the headings.

Although redundancy is a choice, the multiplicity of codings may not be: some codings are unavoidable. When speaking in public, for example, we communicate through what we say (the verbal component), how we say it (the vocal component), and everything that we let the audience see about ourselves (the visual component). Any component escaping our control can carry noise or, what is worse, convey messages that work against our intent, resulting in so-called cognitive dissonance.

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

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

A thousand words, a thousand pictures

Verbal text-like	**Nonverbal** vocal, visual
Rational abstract, learned	**Intuitive** concrete, innate
Sequential slow, exclusive	**Global** fast, nonexclusive

A PICTURE IS WORTH A THOUSAND WORDS, or so they say. In reality, however, not all pictures are created equal, and the power of visual communication is too often misunderstood, not to say misused. Pictures are no panacea; some words convey concepts better than a thousand pictures.

Intellectual processes, complex and still poorly understood, can pragmatically be modeled as either verbal or nonverbal. Verbal processes are rational, able to manipulate intangible, abstract concepts whose symbolic meaning must be learned. Nonverbal processes are intuitive, almost unconscious, tuned for concrete items with nonsymbolic, quasi innate meaning. Verbal code, such as a piece of text, is sequential; as a result, it is processed relatively slowly. In contrast, nonverbal code, such as a photograph, is global and processed in an instant. Verbal and nonverbal processes are about codes, not channels: as an example, text is verbal code, whether it is heard or read.

Verbal and nonverbal processes are independent of each other so they can take place concurrently: for example, an audience can watch a static picture while listening to an explanation. Concurrent verbal processes, however, are mutually exclusive: for example, an audience cannot both read text on the screen and listen to spoken text, unless perhaps if it is the same text.

Nonverbal codings, being intuitive, usually have more impact than verbal ones. To some extent, they are also more credible: we believe tone of voice and body language more than words. Dissonance between verbal and nonverbal codings can be put to good use in irony and in humor, when we let our audience know nonverbally that we do not mean what we say verbally, but is otherwise dangerous. Thus a pictorial representation of what *not* to do is misleading, even when it is accompanied by a text explanation, unless the *not* is expressed visually, too. Likewise, graphs can be intentionally or accidentally deceptive, and no amount of text can fully correct the visual deception.

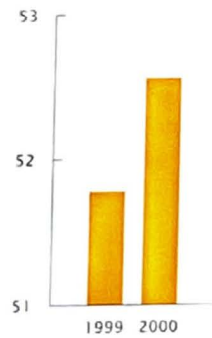

Although the left scale tells us (rationally) that the difference between 1999 and 2000 is 1.5%, the visual code prevails: we keep from the graph a strong, lasting feeling that the value for 2000 was twice that for 1999.

Why are nonverbal codings more credible?

Body language is typically more spontaneous than words are: nonverbal codings are harder to control, hence less likely to be manipulated (the body never lies, according to the proverb). As audience members, we might reason thus and decide to believe the body. More probably, however, we do not reason about the conflict: we absorb nonverbal codings unconsciously, without the analytical filter we apply to words. In other words, we believe nonverbal codings because we have no process to disbelieve them. We instantly sense that the words are untrue.

Can I never include text on presentation slides?

If your oral presentation must discuss a text, for example a novel or an article from a treaty, you might be justified in including some of it on your slides: this text is your very material and no longer a mere support for it. If you do, limit the text excerpts to those strictly needed to make your points. When showing the slide, read the text excerpt out loud for the audience, so they can read it together with you, then have their mental text processor available to handle what you have to tell them about the excerpt.

Text, here, implies sequential processing, as in a sequence of words whose order is dictated by syntax, or perhaps a sequence of sentences. Words that do not form a text and can thus be read in any order, as the labels on a diagram, are less of an issue on slides, since they do not conflict as much with concurrent spoken text. Conversely, animated visual representations, while not textual, are processed sequentially: it may be hard to watch an animated diagram and follow spoken discourse at the same time.

Removing visual noise

Pictures are powerful, and what is powerful is risky. In a verbal coding, and possible connotations aside, the conventional meaning is all that really matters: the word *apple* means "apple"—no more, no less. In contrast, in a visual coding, everything matters: the minutiae of the shape, the shades of color, etc. If such details help get a relevant message across, they are welcome indeed, otherwise, they are noise. Pictures, in other words, carry a higher potential for noise than text. It is usually easier to choose the right word than to come up with the right picture.

When details are irrelevant or otherwise undesirable, as is usually the case in technical communication, photographs can best be replaced by sober line art, less likely to carry irrelevant details. A human hand in a realistic illustration, for example of someone inserting an extension card in an electronic device, is not just any hand: it has a color, gender, and age, not to mention social status and grooming habits. A more schematic representation (line art) displays a hand with which more viewers are able to identify.

Even the most sober line art has its limits, however. The mind is so prompt to recognize visual patterns that it often interprets images in unintended ways. For example, it is as good as impossible to draw a person—or even an animal—without conveying an attitude, which may or may not be well received. Such an attitude, so conspicuous in most clip art, is noticeable even in silhouettes and in stick figures.

A picture excels at representing something intuitive, for example a real object. At the same time, it is condemned to be concrete it cannot convey abstract ideas (at any rate not unambiguously).

As an example of the ubiquitous visual ambiguity, the attitude of this little boy is read differently by different people: is he serene, sad, interested, absentminded? What do *you* see in this picture?

Visual codings, being intuitive and global, are more effective for conveying intuitive or global information. For example, maps convey relative positions more rapidly than words can, drawings describe objects more clearly than words can, and facial expressions show emotions more subtly than words can. Visual codings that mimic facial expressions such as ";-)" have thus emerged in such plain-text media as electronic mail, to convey meaning that relies on intuition—typically humor.

Visual codings, by contrast, are less effective for expressing abstract concepts. A given pictorial representation illustrates only one instance of a concept so easily expressed in words. As an example, a photograph of an apple does not say "apple": rather, it indicates a specimen of specific variety, maturity, etc. as suggested by the apple's visible shape, texture, and color. Nonsymbolic representations are condemned to be concrete, even if schematic drawings can "abstract" irrelevant details, thus broadening the drawing's suggested meaning somewhat.

Visual codings, moreover, lack the accuracy that words are endowed with through conventional association of meaning. Just like Rorschach inkblots, they are intrinsically ambiguous: being intuitive and concrete, they suggest a meaning instantly and may well suggest a different meaning to each viewer, often unable to imagine anyone else "seeing" anything else.

In a sense, a word is worth a thousand pictures, too. Indeed, verbal codings can express abstract concepts unambiguously and concisely, even if not intuitively. As an example, the word *apple* designates any apple and thus transcends all pictures, which can show specimens only. Words can convey concepts that nonsymbolic codings cannot, for example interdiction: showing it visually requires a convention, such as a red circle.

In essence, verbal and nonverbal codings are complementary. They are perhaps the essential form of effective redundancy.

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

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

Chains and magical numbers

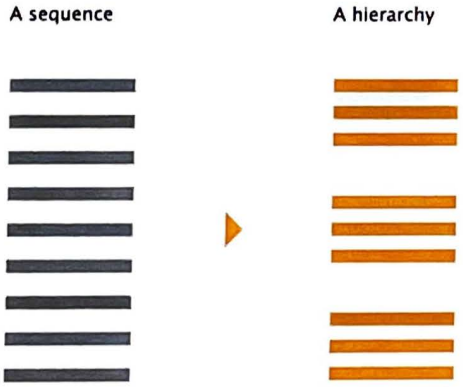

OUR CAPACITY FOR PROCESSING unstructured items of information presented together is severely limited. Series of items, or lists, tax our short-term memory, key to mental processing, and become rapidly unmanageable as the number of items grows beyond just a few. *Processing* and *short-term memory* suggest random access to the items. Longer lists can of course be committed to long-term memory by rote learning, but this memory provides sequential access only: if we forget one link in the chain, we often cannot go on. As a rule, we can process items in random fashion if we can see at a glance how many there are, without having to count.

An *item* is whatever we can, to a point, recognize and process as a unit. To make a long list easier to process, we can group items into fewer, higher-level items, thus creating a hierarchy (a list of lists): a series of three series of three items is easier for a human brain to process than a single series of nine items. For the higher-level items to be recognized as units, however, the original items must be grouped logically, not arbitrarily. This logic must be made explicit or be readily recognizable for the audience. Visual clustering helps show the groups, too.

Balanced, multidimensional structures, in other words, match our mental patterns better than longer, unidimensional ones. Chains, structured along a single dimension, must be accessed in sequence. Trees, structured hierarchically, add a dimension. They offer, not a sequence of items, but a sequence of choices, in the form of recursive branching. They can thus organize a large number of items while offering few enough options in every choice to enable random processing of these items. With the cascade of choices kept short enough, they provide an overview of the collection of items in a way chains cannot. They are thus easier to apprehend, navigate, and remember, and constitute a more robust framework, for example to build a case: whereas a chain is only as strong as its weakest link, a weak argument in a tree does not invalidate the other ones.

Chains just do not communicate well

How about numbering the items, so readers can "see at a glance how many there are"?

"Seeing at a glance how many items there are" is a reliable sign of nonsequential perception. This perception is about *seeing* all the items; *being told* how many there are does not help.

Is sequential processing necessarily a problem?

Admittedly, linear material—for example a set of step-by-step instructions—may at first seem to require no more than sequential processing. Still, a hierarchical structure would give users an overview of the steps to be performed, thus preparing them mentally for the tasks ahead. Moreover, a long list often intimidates readers. Tree structures typically look more accessible.

Is text not processed sequentially anyway?

A verbal list, even short, must indeed be read sequentially: word after word, item after item. *Nonsequential* here refers to the initial (visual) perception of the list of items and, especially, to their manipulation in short-term memory, where they can be, in a sense, "seen" together, reviewed in any order, and finally passed on to long-term memory in a structured manner.

Must presentations always have three points?

Presentations need not have three main points: some topics are better structured in two, four, or perhaps five points. Still, because a structure in three points communicates particularly well, you might want to give it a try systematically, without forcing it on your topic if it does not fit.

While they are harder to process, chains are easier to create than trees, because they need local effort only: they can be constructed one item at a time, with little or no attention to the structure as a whole. As the examples below illustrate, chain structures are frequent in verbal code (written or spoken text) and in nonverbal items (slide shows, graphs, etc.) alike, each time taxing the intellectual capabilities of the audience—and, in one case, of the speaker.

Long series of short sentences—perhaps written in the simplistic belief that shorter sentences are easier to read—make for hard-to-read paragraphs, even if each sentence individually is very readable. The same holds for long series of short paragraphs.

Chains of premodifiers need not be long to create uncertainty as to which words are being modified. They are often found on restaurant menus (*Grilled Applewood Smoked Bacon Wrapped Mission Figs*) and in scientific publications (*Fuzzy-Logic-Controller-Based Cost-Effective Four-Switch Three-Phase Inverter-Fed IPM Synchronous Motor Drive System*).

Writing down and committing to memory the text of a presentation places the speaker at the mercy of the slightest memory lapse: the rest of the text usually cannot be recalled past any missing words (and a gap in the text would be noticeable anyway).

Showing many slides as part of an oral presentation easily creates a disorienting impression of linearity, especially if all slides have the same visual design, that is, unless there are contrasting slides meant to reveal the hierarchical structure of the material.

Graphical displays such as pie charts that include too many items not only face the spatial challenge of labeling all segments clearly but also fail to give an overview of the data. As an alternative, bar charts can display hierarchical groups of data more easily.

The question, of course, is how long a list can reasonably be, that is, how many items presented together are too many. Rather than blindly apply a single, dogmatic "magical number," let us see how small integers can usefully guide our practice of communication. As it happens, there is magic everywhere.

Zero is perfection, as in zero superfluous words on a page zero useless gestures in a talk, zero unneeded ink in a graph. Aiming for zero noise means much hard work for something the audience will *not* notice—frustrating, yes, but effective.

One is focus, as in one theme per document or presentation, one message per paragraph or slide, one idea per sentence. One is consistency and univocality, a prerequisite to meaning in verbal codings: synonyms and homonyms are suboptimal.

Two is a bit, a binary alternative. It is thus the simplest form of classification, as in specialist versus nonspecialist or verbal versus nonverbal. Two is a duality, with all its appeal and all its limitation, as in good and evil, night and day, yin and yang. Besides opposition or complementarity, two is redundancy across channels or codings—a potentially effective approach.

Three is the simplest complexity: it corresponds to a triangle (the first polygon), the number of dimensions in physical space, and the number of colors required to generate all the others. Three is of course a direct extension of two, one that breaks the duality, as by introducing gray between black and white. Interestingly, three is how we group digits in large numbers for increased readability. It is a common-sense upper limit in many cases, for example on the number of heading levels that can meaningfully be numbered together (*Section 2.4.1*). Pragmatically, three is probably the optimal number for items that must be grasped rapidly and remembered easily, such as steps in a procedure or main points in an oral presentation. Three items simply work well—for speaker and audience alike.

The difficulty is combinatorial

I thought seven was "the" magical number and a universal upper limit. Is it not so?

George Miller's now famous article published in 1956 in *Psychological Review*, "The magical number seven, plus or minus two: Some limits on our capacity for processing information," is frequently misunderstood and misquoted. It derives seven as a rough asymptotical limit from experiments that are in fact little relevant to the type of communication discussed here. Anyway, to reduce the risk of processing errors, we should limit the number of items presented together to fewer than the asymptotical value. Pragmatically, we could use the lower bound of Miller's proposed 7 ± 2 interval, namely five.

How can I group items in a table that is not otherwise structured?

You can always group the rows visually by five (or fewer), as by skipping space every five rows or, when space is at a premium, by alternating the background color for groups of five rows. Though not dictated by logic, such a grouping makes the rows of the table easier to read off, especially when the columns are set far apart.

Should rating scales not have some kind of neutral between positive and negative?

The middle point provided by an odd number of options may be desirable, but it can become an easy noncommittal retreat (though perhaps mitigated by an out-of-scale *no opinion* option). Still, three options (+/o/−) provide no degree of positive or negative appreciation, while five are already enough to drive some respondents into avoiding the two extremes systematically.

The abrupt saturation, beyond five, of our capacity for processing a set of items presented together may come as a surprise. Going from five to six items means adding only 20% to the sequence, whereas going from two to three is adding 50%—a lot more. We might thus expect processing six items instead of five to be 2 5 times easier than processing three instead of two, yet experience suggests otherwise.

One plausible explanation is combinatorial analysis. Because the key to apprehending sequences fully is the possibility to process them nonsequentially, we should reason, not in terms of sequence length, but in terms of nonsequential combinations. A set of n items can be combined in $n!$ (factorial n) ways, a function that grows much more sharply than n. Going from two to three items means multiplying the possibilities by three (from $2! = 2$ to $3! = 6$); going from five to six items, by contrast, means multiplying them by six (from $5! = 120$ to $6! = 720$). This model would explain not only why six items are so much harder to handle than five but also why key items benefit from being fewer than five (three items being thus 20 times less demanding).

The sequential process required beyond five items applies to visual codes, too, as soon as their details have to be processed one by one. As an example, how easily can you identify below the differences between upper and lower drawing, besides rotation? First global, the comparison becomes sequential as the number of items increases, unless (part of) the figure becomes meaningful to you as a whole.

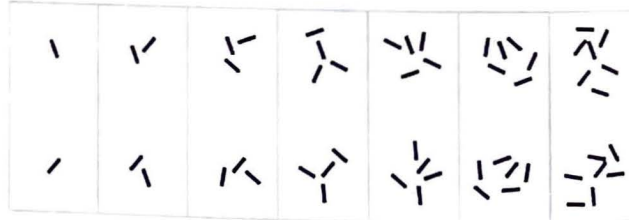

Four is a square (2^2): it is a combination of two binary options. Four is therefore a useful number of answers for rating scales (++/+/−/−−), as it embodies a cascade of two binary choices: first, is it rather positive or negative; next, is it a little or a lot. Four is also a direct extension of three: whatever works well in threes might also work, though not nearly as well, in fours. While large numbers are usually set in groups of three digits, a year such as 1984 is set as a single group of all four digits, with no apparent readability problem (yet it is read in English as *nineteen eighty-four*, or 2 × 2 digits). As another example, whereas *Section 2.4.1* is reasonably easy to situate mentally or to remember, *Section 2.4.1.3* is immediately much harder.

Five is a handful: it is the number of fingers on a human hand but also the limit above which we must count items to know their number—unless they are organized visually in groups of five or less, as can usefully be done with rows in long tables. It is thus a useful upper limit on the number of items in a list.

Six is... just after five, the same way that four is just after three. Consequently, if five is a useful maximum number for a group of items not otherwise structured, then six is just past the limit. In other words, six might work for some people, in some cases.

Seven is many: it is usually too many for the communication to be effective. In a sense, seven is the smallest numerousness, in the same way three is the simplest complexity: seven items presented together are just too numerous to be manageable. Apart perhaps for overwhelming the audience economically, seven is not a particularly useful number for communication.

From the eight numbers above, and beyond the obvious zero and one, you might remember the first three prime numbers: two for effective redundancy, as with verbal versus nonverbal; three as an optimal number—fast to grasp, easy to remember; five as the maximum number that ensures global processing.

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**
  - Balancing the structure
  - Allowing easy navigation
  - Stating messages first

Effective written documents

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

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

Trees, maps, and theorems

STRUCTURE IS A VIEW OF THE MIND. While one scheme may well seem more logical than alternative ones, none is inexorably prescribed by the material to be presented. Organizing material involves choice, so it allows optimization. The optimal structure is the one that makes the most sense for the audience. In other words, it is easy for the audience to recognize and remember, it can be navigated effortlessly even if not memorized, and it limits the need for navigation.

Organizing knowledge effectively requires a hierarchy: a tree, not a chain. At the same time, any instance of communication is trapped in time and forms a sequence: a chain, not a tree. At an elementary level, verbal discourse is indeed sequential: words, sentences, etc., are read or heard one after the other, and their order largely determines the meaning of the text. At a higher level, even reasonably self-sufficient components such as independent chapters, slides, or graphical displays are presented in a certain sequence. In an oral presentation, this sequence is imposed by the speaker; in a document, it is proposed by the writer but ultimately chosen by each reader, who elects both what to read and in what order to read it. Still, sequence there is: even highly selective readers cannot read two different chapters of a document at the same time.

Engineering our communication, then, is a triple challenge: we must organize our material into a well-balanced hierarchy, reveal this hierarchy through what is unavoidably a sequence in time, and ensure that the sequence we propose or impose suits the logic of the audience—and all of this, at all levels, from an entire document or presentation all the way down to a single sentence, which embodies the structure of an idea. Here again, effective redundancy is sure to help. A structure already revealed by verbal discourse can be visualized ideally by nonverbal code, processed globally. The layout of pages or slides (and, especially, of the table of contents or preview) thus plays a key role in revealing a structure to the audience.

How should I number the sections?

As a rule, use the so-called decimal numbering, as in *2.4.1* to indicate Subsection 1 of Section 4 of Chapter 2: it makes the hierarchy apparent. In contrast, selecting a single number or letter for the subsection, as in *IV* or *D*, fails to reveal its place within the overall structure: readers may not remember what it is the fourth part of, especially if they must interrupt their reading or when they are browsing through a document.

To remain readable, decimal numbering is best limited to three levels. Fourth-level headings, if any, can be set without a number. Similarly, when a book includes parts, the part number can probably be omitted from the numbering of the part's chapters, sections, and subsections.

Why a different number of levels and of items for written documents and oral presentations?

Listeners are in a far less favorable situation than readers to process the material presented: for example, they cannot choose their rhythm, cannot reread a part they did not understand, and have fewer visual clues about the structure (as offered to readers by a document's layout). They therefore cannot assimilate as complex a structure as readers can and should thus be presented, when possible, with simpler ones.

Also, presentations and documents normally differ in their purpose. Oral presentations are for convincing an audience of the key messages, while written documents more often attempt to convey a large or complex body of material. Presentations can thus typically accommodate a simpler structure, whereas (long) documents often require a somewhat more elaborate one.

Common ineffective structures

Long documents (reports, theses, procedures, etc.) tend to include too many levels in their hierarchy, perhaps with few items at each level—a structure that extends in depth, often with heading numbers such as 2.3.2.1.1.2 that no longer allow readers to visualize the hierarchy. Here, "too many levels" mean more than three or just more than necessary, as when there are almost more headings than text.

At the same time, there is of course nothing wrong with deeply structured thinking. What is suboptimal is turning each item of a mental tree into a heading. Perhaps each lowest-level item in one's mental tree can be written as one paragraph or one sentence, with fewer (levels of) headings in the tree structure.

Deep structure

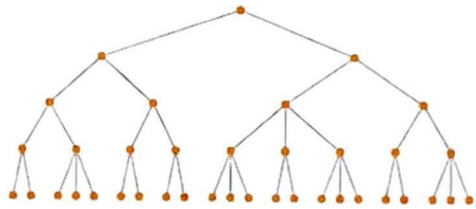

Flat structure

Shorter documents for less specialized audiences, such as magazines and newsletters, often simplify the hierarchy to an extreme, keeping only one level of heading. In the process, however, they end up with many headings at this level—a flat structure, by no means easier to assimilate than a deep one.

For a well-balanced structure, generate for yourself a complete table of contents. Are there many levels? Many headings on few pages? Single subbranches?

Balancing the structure

AN EFFECTIVE STRUCTURE IS HIERARCHICAL, not sequential. Furthermore, it consists of a limited number of levels and a limited number of items at every level. Each such item must form a meaningful entity—one comparable in scope to other entities at the same level, within and across branches. Items within a branch should be preceded by a component that gives a motivation for the branch, previews its structure implicitly or explicitly, and perhaps states its main messages. (The guidelines below apply to the levels of the tree revealed through headings and perhaps numbering, such as chapters and sections, not to levels such as paragraphs and sentences.)

As a rule, use fewer hierarchical levels than items per level, for we handle recursion with even more difficulty than lists. In written documents, endeavor to limit the number of levels to three—for example, chapters, sections, and subsections. If you must group paragraphs within a subsection, consider unnumbered headings, which would not appear in the table of contents. In oral presentations, limit yourself to one level for a short presentation, perhaps two levels for a longer one.

Limit the number of items per level, too, like you do for lists. In written documents, aim for no more than five subbranches for each branch, to afford readers a global view of the branch. Should you seem to need more, group closely related ones and substructure the entity thus obtained with paragraphs. If you have too many chapters, try grouping them into parts. In oral presentations, consider a body in exactly three points.

Before dividing a branch in subbranches, provide a global view. In written documents, include a paragraph (or more) between the heading of, say, a section and that of the first subsection. Among other things, this paragraph must let readers know what the subsections are, as a form of effective redundancy with the set of headings in the text and in the table of contents. In oral presentations, include a preview just before the body.

Well-balanced structure

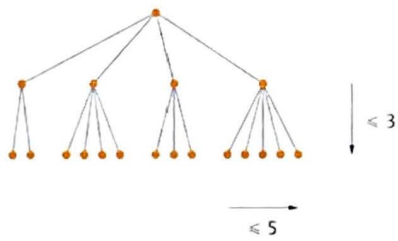

With up to three levels in the tree (chapters, sections, subsections) and up to five items at each level, a balanced document could thus have up to $5^3 = 125$ subsections, which is plenty for most material. Exceptions to this limit of 5^3 are of course possible when justified.

What is wrong with a detailed table of contents on several pages? Does it not enable readers to locate precisely what they are looking for?

Readers may or may not know what to look for in a document (or what the document calls it). When they do, they are best helped by an index, not a table of contents, no matter how detailed. When they do not, or more generally to form an overview of the material, they will normally go through the table of contents hierarchically, not sequentially, identifying first the chapter most likely to be of interest for their purpose, then the section within this chapter, and so on. If they cannot see the whole structure at once, they have to process it sequentially; they miss an overview of the major entry point (chapters).

An alternative to limiting the number of levels presented is to provide two tables of contents: a global one, limited to the top level (chapters), then a fully detailed one. In a sense, the first is a table of contents of the table of contents.

Why must links use the wording of the map?

If links use a different wording from the map to express destinations, readers may be able to find whatever information they are seeking but not to visualize their itinerary, as the links point to places they cannot put on the map. Thus they cannot easily know when and why they might have been there before, which is a major factor in deciding whether to go there.

Using exactly the same wording on the map and in links is a simple issue of consistency: always calling a given thing by the same name helps the audience recognize this thing easily and avoids ambiguity—in links and elsewhere.

Hypertext, in a sense, predates the electronic age. While the term suggests clickable links (hyperlinks), the idea of linking a piece of text to another piece of text is not a recent one. And although new media open new possibilities, effective paper documents have provided for centuries both the motivation and the means for readers to jump to other parts of the page, to other pages, or to other documents.

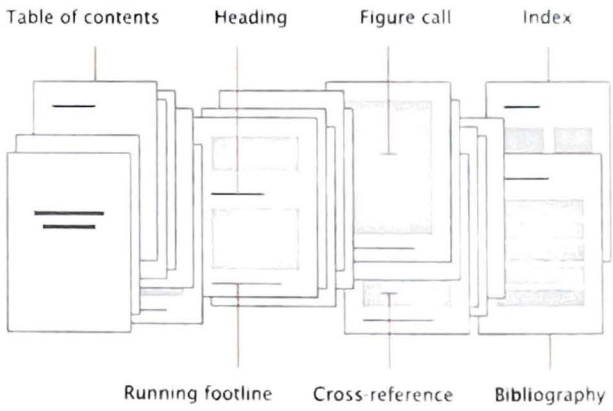

Effective paper documents, in particular the longer and strongly structured ones such as books, carry a surprising amount of navigational information, both by nature and by contents. As physical sets of pages, they provide clues to where readers are, such as near the beginning. Numbered headings and running headlines or footlines can further help readers locate their current position in a structure already made explicit through a table of contents. This table, perhaps together with a text overview near the end of the introduction, enables readers to make informed decisions regarding where to go. So does an index, along a different logic. Finally, the references to bibliographical entries, to tables and figures, and to other parts of the document are as many "hyperlinks"—just not clickable ones.

Allowing easy navigation

AN EFFECTIVE STRUCTURE CAN BE NAVIGATED effortlessly if made visible. To this end, give the audience a map, tell them at all times where they are on the map, and tell them (when appropriate) where they can go with respect to the map.

Effective maps provide an immediate overview of the territory they chart. They usually include a strong visual component. In written documents, readers should see the tree structure at a glance in a table of contents before reading the first word. In Web sites, this table of contents, often called a site map, can be a two-dimensional diagram instead of a vertical list. In oral presentations, the preview can be shown on a slide, besides being spoken and perhaps underlined with gestures.

No matter how many pages it appears on in the end, too detailed a table of contents is hard for readers to process globally and, especially, to remember in a form that helps them navigate the document. Two levels is probably all they can absorb at first, even though a document can have up to three levels. For an oral presentation, subtract one everywhere: up to two levels in the tree, and one in the preview.

To provide an immediate overview, a map should be visible as a whole. As a counterexample, a table of contents running on several (double) pages offers partial views but no overview. Maps, however, need not display the complete depth: they can be limited to the top two levels, such as chapters and sections; each chapter can then include a local map, listing its sections and its subsections—not unlike road maps at different scales.

Audience members need to be reassured about where they are. Because they may wonder about it at any time, let them know at all times, for example in a running footline in a document. When it is impractical to do so at all times, as in oral discourse, let them know often, so they never have to wait long for help, should they feel lost at any time. To tell them where they are, use the same wording (or identifying picture) as on the map.

protected by copyright

protected by copyright (see Terms of use for information on authorized reproduction).

A cryptic hyperlink: where does it lead?

A better hyperlink: it refers to the map and allows a better informed decision.

In written documents, audience members must also be told where they can go, as in a cross-reference in a paper document or a hyperlink on a Web site. So they can orient themselves and visualize their itinerary, indicate possible destinations like locations, with the same wording or picture as on the map. Provide whatever information will allow informed decisions.

Is it not better to have the message as a climax, so as to maintain the interest of the audience?

Professional audiences listen to presentations or read documents chiefly because they hope to find in this way the information they need. They easily lose patience when what they read or listen to does not seem relevant to them, all the more so when they are pressed by time. Hence, they have little tolerance for suspense, at least whenever it lasts longer than an instant. Presenting last what an audience wants to know first or most of all is more likely to make them *lose* interest. Details acquire their full relevance and significance in the light of the conclusion.

Keeping the audience interested is not a purpose in itself, really; getting the message across is, and this does not require forcing the audience to read or listen to everything. If they "buy it" early, they need not go on reading or listening to the supporting details for us to be successful.

A message placed early has a practical benefit, too: if audience members must stop reading or listening at any time, they will more likely have read or have heard our message already.

Must I not detail everything I have carried out if only so others can reproduce my experiment?

In scientific publications, whatever experiments help support the conclusions must indeed be described in enough detail to allow replication. Still, *enough detail* does not mean every detail: it means every *relevant* detail. These details need moreover not stand in the way. Because they seldom, in themselves, convey messages, they can often be relegated to an appendix or, for a presentation, to a companion document.

Building from the top down

As a rule, identify your message(s) early, not only for your audience within a given written document or oral presentation but also, first of all, for yourself when designing this document or this presentation. If you do not quite know what you are heading for, you have no relevant way to select what to include and what to leave out: you will likely include more than what is necessary to get your message across. Identifying the message at the end of the process and relocating it upfront is better than leaving it at the end (or, worse, including no message at all), but it does not promote careful content selection.

Top-down approach applied to an oral presentation

1 **Main message** What you want your audience to remember

2 **Main points** What supports this message

3 **Subpoints**

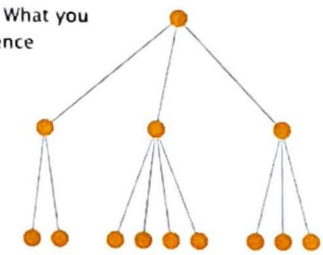

Ask yourself first, *If my audience is to remember only one sentence from my entire oral presentation, what must it be?* The answer is your main message. Repeat the question at each level, broadening out as you go down (a few sentences, instead of one).

The top-down strategy works for entire documents or presentations and also for individual paragraphs or slides. Identify the message that the paragraph or the slide must convey, then—and only then—select the information and format that will convey this message optimally (instead of first developing the exact contents of the paragraph or slide, then searching for a suitable message to tack on to it).

Stating messages first

AN EFFECTIVE STRUCTURE PRESENTS MATERIAL in the order in which the audience is most likely to want to learn it, thus minimizing the navigation required—or the impatience of an audience who cannot navigate, as in oral presentations. It presents first what the audience is primarily interested in and afterwards what is less important or less urgent to them. It presents separately what fewer of them will want to know (perhaps in a companion document for an oral presentation).

Motivation	Make the audience receptive to the topic of the document
Message	Once you have their attention, tell them your main message
Details	Next, support this message: tell them how you got there
Appendix	Last of all, present separately what fewer will want to know

Professional audiences want to be told the message (that is, the *so what*) early, though not without proper motivation first, for messages make little sense out of context. They usually have little patience for "detective stories" that do not reveal the bottom line until the end. The details of an investigation are indeed largely irrelevant until we know the outcome of it. In other words, a chronological structure reporting work done in the order in which it was done, with conclusions at the end, works poorly: it focuses on the authors, not on the audience. An effective sequence presents first the motivation (the *why*), then the main message(s), and finally the details of the work.

A useful model, one that breaks the chronological paradigm, is that of mathematicians. When reporting their work, these normally present their conclusion first, calling it a theorem, then detail their hard work for those interested. In doing so, they strive to limit the details to whatever is strictly needed to prove the theorem, giving preference to the most elegant (that is, the simplest) proof of all. Thus a report is not a story of the work: it needs neither narrate every detail of the work nor report events in the sequence in which these took place.

The theorem-proof sequence is a useful model at all levels, not only for entire written documents or oral presentations, but also for single paragraphs or slides: these can usefully state a message upfront, then develop it verbally or visually, respectively—a prototypical way of getting messages across.