

SURFACES AND ESSENCES

ANALOGY AS THE FUEL AND FIRE OF THINKING

DOUGLAS HOFSTADTER & EMMANUEL SANDER

BASIC BOOKS

A Member of the Perseus Books Group
New York

PROLOGUE

Analogy as the Core of Cognition



Giving Analogy its Due

In this book about thinking, analogies and concepts will play the starring role, for without concepts there can be no thought, and without analogies there can be no concepts. This is the thesis that we will develop and support throughout the book.

What we mean by this thesis is that each concept in our mind owes its existence to a long succession of analogies made unconsciously over many years, initially giving birth to the concept and continuing to enrich it over the course of our lifetime. Furthermore, at every moment of our lives, our concepts are selectively triggered by analogies that our brain makes without letup, in an effort to make sense of the new and unknown in terms of the old and known. The main goal of this book, then, is simply to give analogy its due — which is to say, to show how the human ability to make analogies lies at the root of all our concepts, and how concepts are selectively evoked by analogies. In a word, we wish to show that analogy is the fuel and fire of thinking.

What Dictionaries Don't Say about Concepts

Before we can tackle this challenge, we need to paint a clear picture of the nature of concepts. It is easy — in fact, almost universal — to underestimate the subtlety and complexity of concepts, all the more so because the tendency to think of concepts in overly simple terms is reinforced by dictionaries, which claim to lay out the various different meanings of a given word by dividing the main entry into a number of subentries.

Take, for example, the noun “band”. In any reasonably-sized dictionary, there will be, in the overall entry for this word, a subentry describing a band as a piece of cloth that can be wrapped around things, another subentry describing how a band can be a colored strip or stripe on a piece of cloth or other type of surface, another subentry

describing a band as a smallish set of musicians who tend to play certain types of music or to use only certain types of instruments, another one for a group of people who work or play together, another one for a wedding ring, another one for a selection on a record or a compact disk, another one for a range of frequencies or energies or prices or ages (etc.), and perhaps a few others. The dictionary will clearly set out these various concepts, all fairly distinct from each other and all covered by the same word "band", and then it will stop, as if each of these narrow meanings had been made perfectly clear and were cleanly separable from all the others. All well and good, except that this gives the impression that each of these various narrower meanings of the word is, on its own, homogeneous and not in the least problematic, and as if there were no possible risk of confusion of any one of them with any of the others. But that's nowhere near the truth, because sub-meanings are often closely related (for instance, the colored stripe and the range of frequencies, or the wedding ring and the piece of cloth wrapped around something), and because each of these supposedly clear and separate senses of the word "band" constitutes on its own a bottomless chasm of complexity. Although dictionaries give the impression of analyzing words all the way down to their very atoms, all they do in fact is graze their surfaces.

One could spend many years compiling a huge anthology of photographs of highly diverse wedding bands, or, for that matter, an anthology of photos of headbands, or of jazz bands, or of bands of criminals — or then again, of photos of wildly different chairs or shoes or dogs or teapots or versions of the letter "A", and on and on — without ever coming close, in any such anthology, to exhausting the limitless possibilities implicitly inherent in the concept. Indeed, there are books of precisely this sort, such as *1000 Chairs*. If the concept *chair* were completely straightforward, it is hard to see what interest such a book could possibly have. To appreciate the beauty, the originality, the practicality, or the style of a particular chair requires a great deal of experience and expertise, of which dictionaries cannot convey even an iota.

One could of course make similar observations concerning the subtleties of various types of bands — thus, one could spend one's whole life studying jazz bands, or headbands, or criminal bands, and so forth. And even concepts that seem much simpler than these are actually endless swamps of complexity. Take the concept of the capital letter "A", for instance. One would need many pages of text in complex, quasi-legal language if one were trying to pin down just what it is that we recognize in common among the countless thousands of shapes that we effortlessly perceive as members of that category — something that goes way beyond the simple notion that most people have of the concept "A" — namely, that it consists of two oppositely leaning diagonal strokes connected by a horizontal crossbar.

Indeed, catalogues of typefaces are veritable gold mines for anyone interested in the richness of categories. In the facing figure, we have collected a sampler of capital "A"'s designed for use in advertising, and as is clear from a moment's observation, any *a priori* notion that one might have dreamt up of A-ness will be contradicted by one or more of these letters, and yet each of them is perfectly recognizable — if not effortlessly so when displayed all by itself, then certainly in the context of a word or sentence.



The everyday concepts *band*, *chair*, *teapot*, *mess*, and *letter 'A'* are very different from specialized notions such as *prime number* or *DNA*. The latter also have unimaginably many members, but what is shared by all their members is expressible precisely and unambiguously. By contrast, in the mental structure underpinning a word like “band”, “chair”, “mess”, or “teapot” there lurks a boundless, blurry richness that is completely passed over by dictionaries, because spelling out such subtleties is not a dictionary’s aim. And the fact is that ordinary words don’t have just two or three but an *unlimited number* of meanings, which is quite a scary thought; however, the more positive side of this thought is that each concept has a limitless potential for variety. This is a rather pleasing thought, at least for people who are curious and who are stimulated by novelty.

Zeugmas: Amusing Revealers of Conceptual Subtlety

There is a linguistic notion called “zeugma” (also sometimes called “syllepsis”) that, although it is fairly obscure, has a good deal of charm and brings out the hidden richness of words (and thus of concepts). The zeugma or syllepsis is one of the classical

figures of speech, and is often — perhaps nearly always — used to humorous effect. It is characterized by the fact that more than one meaning of a word is exploited in a sentence, although the word itself appears only once. For example:

I'll meet you in five minutes and the garden.

This sentence exploits two different meanings of the preposition “in” — one temporal and the other spatial. When one imagines meeting someone *in* a garden, one sees in one's mind's eye two relatively small entities physically surrounded by a larger entity, whereas when one imagines a meeting taking place *in* five minutes, one thinks of the period of time that separates two specific moments from each other. Everyone understands with no trouble that these are two very different concepts associated with the same word, and the fact that the preposition “in” is used only once in the sentence despite the wide gap between the two meanings that it's conveying is what makes us smile when we read the sentence.

Here are a few other somewhat humorous examples of zeugmas:

Kurt was and spoke German.

The bartender gave me a wink and a drink.

She restored my painting and my faith in humanity.

I look forward to seeing you with Patrick and much joy.

In the first, the word “German” is forced to switch rapidly, in the reader's mind, from being an adjective denoting a nationality to being a noun denoting a language.

The second zeugma involves two different aspects of the notion of *transfer* between human beings. Does one person really *give* a wink to another person? Is a wink a material object like a drink, which one person can hand another?

In the third zeugma, the speaker's faith in humanity had disappeared and was made to come back, whereas the painting had not disappeared at all. Moreover, faith in humanity is far less palpable than a painting on one's wall. What gives this zeugma its flavor of oddness is that one of the meanings of the verb “restore” that it depends on is “to return something that has been lost”, while the other meaning used is “to make something regain its former, more ideal state”, and although these two senses of the same word are clearly related, they are just as clearly not synonymous.

Finally, the last zeugma in our quartet plays on two sharply contrasting senses of the preposition “with”, one conjuring up the image of someone (Patrick) physically accompanying someone else (the speaker and the person being addressed), and the other communicating the emotional flavor (great pleasure) of a mental process (the anticipation of a reunion). As in the other cases, the zeugmatic use of “with” brings out the wide gap between two senses of one word, and to experience this distinction in such a crisp fashion is thought-provoking. We thus see that any well-designed zeugma will, by its very nature, automatically highlight certain semantic subtleties of the word (or phrase) around which it is built.

For example, what does the word “book” mean? One would at first tend to say that it designates an object made of printed sheets of paper bound together in some fashion, and having a cover (and so forth and so on). This is often correct, but the following zeugma brings out a different sense of the word:

The book was clothbound but unfortunately out of print.

This sentence reminds us that the word “book” also denotes a more abstract concept — namely, the set of all copies available in stores or warehouses. Are we thus in the presence of *one* concept, or of *two*? And when someone says, “I’m translating this book into English”, are they using a third sense of the word? How many subtly distinct concepts secretly coexist in the innocent word “book”? It would be an instructive exercise to try to construct more zeugmas based on yet other senses of the word “book”, but we have other goals here, so we will leave that challenge to our readers.

Instead, let’s look at a somewhat more complex zeugma:

When they grew up, neither of those bullies ever had to pay
for all the mean things that they did as, and to, younger kids.

Here the trickiness is in the strange, lightning-fast shifting of meaning of “younger kids” as a function of whether it is seen as part of the phrase “things that they did as younger kids” or as part of the phrase “things that they did to younger kids”, since in the first case the *younger kids* are the ex-bullies themselves (or rather, the bullies that they once were), while in the latter case the *younger kids* are their victims.

Some Revealing Zeugmas

Although the zeugmas we’ve exhibited above are mostly quite amusing, it’s not for entertainment but for enlightenment that we’ve brought up the topic. And so let’s take a look at some cases that raise more serious issues.

“You are always welcome in my home,” he said in English and all sincerity.

This zeugma is clearly built around the word “in”, and the natural question here is whether we are dealing with *one* sense or *two* senses of the word. In a respectable dictionary, these two meanings would probably have distinct subentries. However, what about the following sentence?

“You are no longer welcome in my home,” he said in anger and all sincerity.

Are the two meanings of “in” here exactly the same? Perhaps — after all, they both apply to the mental states of a single person; but then again perhaps not — after all, one could replace “in anger” by “in an outburst of anger” but certainly one could not

say “in an outburst of sincerity”. So it’s rather tricky. As a matter of fact, it would be impossible to give a definitive judgment on this issue. Indeed, we chose this example precisely because it brings out certain subtle nuances of the concept *in*. How does one recognize those situations that match the English word “in”? To put it another way, how does one recognize *in*-situations? What do all *in*-situations have in common, and how do some of them differ from others, and why would it be next to impossible to make a precise and sharp classification of all the types of *in*-situation?

Let’s shift our attention from a preposition to a verb. Does the following sentence strike you as innocuous and perfectly acceptable (*i.e.*, nonzeugmatic), or does it grate on your ears (thus it would be a zeugma)?

I’m going to brush my teeth and my hair.

Are the two types of brushing really just one thing deep down, or are they worlds apart? We might gain perspective on this question by looking at a similar example in another language. In Italian, one might easily and comfortably say:

Voglio lavarmi la faccia e i denti.

(In a fairly literal translation, this says, “I want to wash my face and my teeth.”) The fact that Italian speakers say things this way sheds light on how they perceive the world — namely, it shows that they perceive the act of washing one’s face and the act of brushing one’s teeth as belonging to the same category (both are types of *washing*), and thus they are, in some sense, “the same act”.

On the other hand, to speakers of English, brushing one’s teeth is not a kind of washing (washing usually involves soap of some sort, and most people would hesitate to refer to toothpaste as “soap”, though the two have much in common), so the sentence sounds zeugmatic (that is, its double application of the same word makes us smile). As for French, although occasionally one will hear “se laver les dents” (“to wash one’s teeth”), it is more common to say (and hear) “se brosser les dents” (“to brush one’s teeth”). The latter seems more natural to French speakers than the former. And thus we see that a phrase (“to wash one’s teeth and one’s face”) can be very zeugmatic in one language (English), can have a faintly zeugmatic flavor in another language (French), and can be totally nonzeugmatic in a third language (Italian).

The preceding example shows how a zeugma can reveal a conceptual division that speakers of language A find blatantly obvious, while to speakers of language B it is difficult to spot. For instance, in English, we can say without any sense of oddness:

Sometimes I go to work by car, and other times on foot.

In German or Russian, however, these two forms of locomotion call for different verbs. When one takes a vehicle to arrive at one’s destination, then the verb “fahren” is used in German, whereas when one goes somewhere on foot, then the verb “gehen” is used.

In Russian it's trickier yet, because not only is there a distinction between *going in a vehicle* and *going on foot*, but also the choice of verb depends on whether this kind of motion is undertaken frequently or just one time. Thus a completely innocuous-seeming verb in English breaks up into several different verbs in Russian. In other words, what to English speakers seems to be a monolithic concept splits into four distinct concepts to Russian speakers.

Let's take another very simple sentence in English:

The boy and the dog were eating bread.

This sentence is nonzeugmatic in English; that is, it simply *works*, sounding neither strange nor humorous to the English-speaking ear. On the other hand, it sounds wrong in German, because different verbs apply to animal and human ingestion — “fressen” for the beasts, and “essen” for humans. In other words, German speakers split up what to us anglophones is the monolithic concept of *eating*, breaking it into two varieties, according to the type of creature that is carrying out the act.

The “Natural” Conceptual Distinctions Provided by Each Language

These examples might inspire someone to imagine a language (and culture) that has no verb that applies both to men and to women. Thus it would have one verb that would apply to eating acts by *men* and a different one that would apply to eating acts by *women* — say, “to wolf down” for men and “to fox down” for women, as in “Petunia foxed down her sandwich with relish, gusto, and pickles”. Speakers of this hypothetical language would find it jolting to learn that in English one can say, “My husband and I enjoy eating the same things” or “A girl and a boy were walking down the sidewalk.” To them, such sentences would sound nonsensical. A language like this may strike you as ludicrous, but many languages do make just such gender-based lexical distinctions.

For instance, in French there is a clear-cut distinction between enjoyment partaken of by men and enjoyment partaken of by women, which shows up in, among other venues, the standard adjective meaning “happy”: whereas a joyous man or boy will be “heureux”, a joyous woman or girl will be “heureuse”. And thus, a *curieux* French male might well wonder what it feels like to be *heureuse* — but he would do so in vain! A man simply cannot be *heureuse*! In like manner, a *curieuse* French woman might wonder what it feels like to be *heureux* — but her efforts, no matter how valiant, would be doomed to failure. A Venusian might as well try to imagine what it feels like to be Martian!

Does all this sound far-fetched to you? Well, consider that there is a famous Russian poem centered on what the poet, a man named Il'ya L'vovich Selvinsky, considered a very strange fact: namely, that every act of his lover — every single one of the mundane verbs that described her actions — was graced, when in the past tense, by a feminine ending (often the syllables or bisyllables “la”, “ala”, or “yala”). The poet describes various completely ordinary actions on her part (walking, eating, etc.), and then expresses wonderment at his own feeling of disorientation, since he, being a male,

has never once performed a single one of these “uniquely feminine” acts, nor experienced a single one of these “uniquely feminine” sensations, and, alas, will never be able to do so. In making such observations, is Selvinsky expressing something deep, or is he merely playing with words?

One can easily enough imagine a language that, with a panoply of verbs, distinguishes between a vast number of different ways of eating — the eating of a famished boy, of a high-society lady, of a pig, a horse, a rabbit, a shark, a catfish, an eagle, a hummingbird, and so forth and so on. Such a fine-grained breakup of a concept that seems to *us* completely monolithic is perfectly imaginable, because we understand that there are genuine differences between these creatures’ ways of ingesting food (indeed, if there weren’t any, we would not have written “genuine differences”). Each language has the right and the responsibility to decide where it wishes to draw distinctions in the zone of semantic space that includes all of these distinct activities. After all, there are not, on earth (and never have been, and never will be) two creatures that eat in an exactly identical fashion, nor even two different moments in which a single creature eats in exactly the same manner, down to the tiniest detail.

Every act is unique, and yet there are resemblances between certain acts, and it is precisely these resemblances that give a language the opportunity to describe them all by the same label; and when a language chooses to do so, that fact creates “families” of actions. This is a subtle challenge to which every language reacts in its own fashion, but once this has been done, each group of people who share a common native language accepts as completely natural and self-evident the specific breakdown of concepts handed to them by their language. On the other hand, the conceptual distinctions that are part and parcel of *other* languages may strike them as artificial, pointlessly finicky, even incomprehensible or stupid, unless they find some interest in the subtleties of such distinctions, which may then make them see their own set of concepts in a fresh light.

Wordplay with the Word “Play”

The verb “to play” affords us a delightful sampler of zeugmas, or else, depending on a person’s native language and on their own personal way of perceiving the actions involved, non-zeugmas. For example:

Edmond plays basketball and soccer.

This sentence, on first sight, might seem about as natural as they come, and very far from zeugmaticity, and yet the two activities involved, although they both belong to the category of *sports*, are different in numerous ways from each other. For instance, one involves a ball that is primarily in contact with the feet (and on occasion with the head), while the other involves a ball that is primarily in contact with the hands (and virtually never with the head). Certain speakers of English might therefore hear a trace of strangeness, albeit only very slight, in the application of the same verb to two rather disparate activities.

If *essen* (which is what people do when they eat food) and *fressen* (which is what, say, pigs and rabbits do with their food) are seen by German speakers as activities that belong to two different categories, then there is nothing to keep us from imagining a language in which one would say:

Edmondus snuoiqs basketballum pluss iggfruds soccerum.

The speakers of this hypothetical language would see the actions of basketball players — or rather, of basketball *snuoiqers* — as being just as different from the actions of soccer *iggfruders* as the sounds “snuoiq” and “iggfrud” are different from each other.

If this example’s zeugmaticity seems too weak, then we can try another avenue of approach to the same issue:

Sylvia plays tennis, Monopoly, and violin.

This sentence involves a musical instrument and two types of game that are much more different from each other than are basketball and soccer. If one tried to measure the distances between these three concepts by asking people to estimate them, it’s likely that most people would place *violin* quite a long ways from *tennis* and *Monopoly*, and those two games, though not extremely near each other, would be much closer than either of them is to *violin*. And finally, not too surprisingly, this matches the collective choice of Italian speakers, who would translate the above sentence as follows:

Sylvia *gioca* al tennis e a Monopoly, e *suona* il violino.

It would be unthinkable, in Italian, for anyone to *play* (in the sense of *giocare*) a musical instrument; the mere suggestion is enough to make an Italian smile. The kind of scene that such a phrase would conjure up is that of people playing catch with a Stradivarius, for instance. While it is natural for English and French speakers to see violin-playing as belonging to the same category as soccer-playing and basketball-playing, the idea would seem downright silly to Italian speakers.

In French, the verb *jouer* is used both for musical instruments and for sports, but it is followed by different prepositions in the two cases. Thus one plays *at* a sport but one plays *of* a musical instrument. Does this syntactic convention split the concept of *jouer* into two quite clear and distinct sub-meanings? In English, there is no similar syntactic convention that would create a mental division of the verb “to play” into two separate pieces; rather, it simply feels monolithic.

Playing Music and Sports in Chinese

The distinction made in Italian between “*giocare*” (for sports) and “*suonare*” (for musical instruments) might seem a bit precious. After all, not only English but plenty of other languages are happy to use exactly the same verb for both kinds of activities —

thus French uses “jouer”, German uses “spielen”, Russian uses “играть”, and so on. What about Chinese?

It turns out that Mandarin speakers are considerably more finicky in this matter than Italian speakers: they linguistically perceive four broad types of musical instruments, each type meriting its own special verb. Thus for stringed instruments there is the verb “拉” (pronounced “lā”), meaning roughly “to pull”, while for wind instruments one says “吹” (“chuī”), which means “to blow”. Then for instruments such as the guitar, whose strings are plucked by the fingers, or the piano, whose keys are pushed by the fingers, the verb is “弹” (“tán”) — and finally, for drums, which are banged, what one says is “打” (“dǎ”).

Curiously enough, it’s possible to apply the verb that means “to play” (as in “play with a toy”) to any musical instrument (it is “玩”, pronounced “wán”); unfortunately, however, the meaning is not what an English speaker might expect: it’s essentially the idea of *fussing around* with the instrument in question, and moreover this usage of “玩” is extremely informal, indeed slangy.

One might naturally wonder how a Chinese speaker would ask a more generic question, such as “How many instruments does Baofen play?” But the best translations of this perfectly natural English sentence elegantly bypass the problem by making use of very broad verbs such as “学习” (“xuéxí”) or “会” (“huì”), which mean, respectively, “to study” and “to be able; to know”, and which have no particular connection with music. In short, there is no general verb in Mandarin that corresponds to the *musical* notion of playing, even though to us English speakers the concept seems totally logical, even inevitable; but the fact is that speakers of Chinese have no awareness of this lacuna in their lexicon, no matter how blatant it might seem to us.

Well, all right, then. But what about playing games and sports — surely there is just one verb in Chinese for this monolithic concept? To begin with, one does not, in Mandarin, play board games and sports with the same verb. For chess, one engages in the activity of “下” (“xià”), which one does not do with any kind of ball. And for a sport that uses a ball, it all depends on the kind of ball involved. For basketball, it’s “打” (“dǎ”), the verb that applies to playing a drum (the connection may seem a bit strained to a non-Chinese), whereas for soccer it’s “踢” (“tī”), which means “to kick”. Thus one might say, “I prefer kicking soccer to beating basketball.” Once again we see that in a domain that strikes an English speaker as monolithic — everything is *played*, and that’s all there is to it! — distinctions are not just rife but necessary in Chinese.

For English speakers, despite our use of the single verb “to play”, it’s not terribly hard to see that this verb conflates two activities that are quite different — namely, making rhythmic noises and having fun — and that the conceptual union thus created is not inevitable, and might even be seen as being rather arbitrary. On the other hand, *within* each of these two domains, it’s harder to see a lack of natural unity. If someone were to ask us if *playing dolls*, *playing chess*, and *playing soccer* are all really “the same activity”, we could of course point out differences, but to focus on such fine distinctions would seem quite nitpicky. And when we learn that in Mandarin, *playing soccer* and *playing basketball* require different verbs, it is likely to strike us as really overdoing things,

rather as if some exotic tongue insisted on using two different verbs to say “to drink”, depending on whether it involved drinking *white* wine or *red* wine. But then again, this is an important distinction for wine-lovers, so it’s conceivable that some of them would very much like the idea of having two such verbs.

Zeugmas and Concepts

Our brief excursion to Zeugmaland will come to a climax in the following bold prediction:

You will enjoy this zeugma as much as a piece of chocolate or of music.

This sentence has a couple of zeugmatic aspects. Firstly, it plays on two senses of the noun “piece”. In some readers recognition of this contrast will evoke a smile, even though there’s no denying that both usages of the word are completely standard. Secondly, it plays on three senses of the verb “enjoy” — one involving a gustatory experience, another involving an auditory experience, and yet another involving the savoring of a linguistic subtlety. Each reader will of course have a personal feeling for how large the distinction between these three senses of the word is.

Aside from making us smile, zeugmas offer us the chance to reflect on the hidden structure behind the scenes of a word or phrase — that is, on the concept associated with the lexical item, or more precisely, on the *set of concepts* associated with it — and since most words could potentially be used to form a zeugma (including very simple-seeming words such as “go”, as we saw above in the discussion of German and Russian), the phenomenon necessarily increases our sensitivity to the miracle of the human brain’s ability to spontaneously assign just about anything it encounters to some previously known category. After all, despite the inevitable and undefinable blurriness of the “edges” of each one of our categories, and despite the enormous number of categories, our brains manage to carry out such assignments in a tiny fraction of a second and in a manner of which we are totally unaware.

The Nature of Categorization

The spontaneous categorizations that are continually made by and in our brains, and that are deeply influenced not just by the language we are speaking but also by our era, our culture, and our current frame of mind, are quite different from the standard image, according to which categorization is the placing of various entities surrounding us into preexistent and sharply-defined mental categories, somewhat as one sorts items of clothing into the different drawers of a chest of drawers. Just as one can easily stick one’s shirts into a physical drawer labeled “shirts”, so one would easily assign dogs to the mental drawer labeled “dog”, cats to the nearby mental drawer labeled “cat”, and so forth. Every entity in the world would fit intrinsically into one specific mental “box” or “category”, and this would be the mental structure to which all the different entities

of the same type would be assigned. Thus all bridges in the world would be unambiguously assigned to the box labeled "bridge", all situations involving motion would be assigned to the box labeled "move", and all situations involving things standing still would be assigned to the box labeled "stationary". This mechanism of "boxing" everything in the world would be both automatic and completely reliable, the *raison d'être* of mental categories being to assign entities objectively to their proper conceptual label in an objective, observer-independent fashion.

Such a vision of the nature of categorization is very far from what really goes on, and in the pages to come we will do our best to show why this is so. But hopefully, already from Chapter 1 onwards, readers will feel persuaded that mental categories are anything but drawers into which clear-cut items are automatically sorted, and this idea will be reinforced ever more strongly as the book proceeds.

What, then, do we mean in this book by "category" and "categorization"? For us, a category is a mental structure that is created over time and that evolves, sometimes slowly and sometimes quickly, and that contains information in an organized form, allowing access to it under suitable conditions. The act of categorization is the tentative and gradated, gray-shaded linking of an entity or a situation to a prior category in one's mind. (Incidentally, when we use the term "category", we always mean a category in someone's mind, as opposed to mechanical labels used in computer data bases or technical labels used in scientific taxonomies, such as lists of the names of biological species.)

The tentative and non-black-and-white nature of categorization is inevitable, and yet the act of categorization often feels perfectly definite and absolute to the categorizer, since many of our most familiar categories seem on first glance to have precise and sharp boundaries, and this naïve impression is encouraged by the fact that people's everyday, run-of-the mill use of words is seldom questioned; in fact, every culture constantly, although tacitly, reinforces the impression that words are simply automatic labels that come naturally to mind and that belong intrinsically to things and entities. If a category has fringe members, they are made to seem extremely quirky and unnatural, suggesting that nature is really cut precisely at the joints by the categories that we know. The resulting illusory sense of the near-perfect certainty and clarity of categories gives rise to much confusion about categories and the mental processes that underlie categorization. The idea that category membership always comes in shades of gray rather than in just black and white runs strongly against ancient cultural conventions and is therefore disorienting and even disturbing; accordingly, it gets swept under the rug most of the time. Since the nature of mental categories is much subtler than the naïve impression suggests, it is well worth examining carefully.

A category pulls together many phenomena in a manner that benefits the creature in whose mind it resides. It allows invisible aspects of objects, actions, and situations to be "seen". Categorization gives one the feeling of understanding a situation one is in by providing a clear perspective on it, allowing hidden items and qualities to be detected (by virtue of belonging to the category *person*, an entity is known to have a *stomach* and a *sense of humor*), future events to be anticipated (the glass that my dog's tail just knocked

off the table is going to break) and the consequences of actions to be foreseen (if I press the "G" button, the elevator will go down to the ground floor). Categorization thus helps one to draw conclusions and to guess about how a situation is likely to evolve.

In short, nonstop categorization is every bit as indispensable to our survival in the world as is the nonstop beating of our hearts. Without the ceaseless pulsating heartbeat of our "categorization engine", we would understand nothing around us, could not reason in any form whatever, could not communicate with anyone else, and would have no basis on which to take any action.

Two Misleading Caricatures of Analogy-making

If categorization is central to thinking, then what mechanism carries it out? Analogy is the answer. But alas, analogy-making, like categorization, is also plagued by simplistic and misleading stereotypes. We therefore proceed straightaway to discuss those stereotypes, in the aim of quickly ridding ourselves of the contaminating and confusing visions that they give of the nature of the motor of cognition.

The first of these stereotypes of analogy-making takes the word "analogy" as the name of a certain very narrow class of sentences, seemingly mathematical in their precision, of the following sort:

West is to east as left is to right.

This can be made to look even more like a mathematical statement if it is written in a quasi-formal notation:

west : east :: left : right

Intelligence tests often employ puzzles expressed in this kind of notation. For example, they might pose problems of this sort: "tomato : red :: broccoli : X", or perhaps "sphere : circle :: cube : X", or "foot : sock :: hand : X", or "Saturn : rings :: Jupiter : X", or "France : Paris :: United States : X" — and so forth and so on. Statements of this form are said to constitute *proportional analogies*, a term that is itself based on an analogy between words and numbers — namely, the idea that an equation expressing the idea that one pair of numbers has the same ratio as another pair does ($A/B = C/D$) can be carried over directly to the world of words and concepts. And thus one could summarize this very analogy in its own terms:

proportionality : quantities :: analogy : concepts

There is no scarcity of people who believe that this, no more and no less, is what the phenomenon of analogy is — namely, a template always involving exactly four lexical items (in fact, usually four words), and which has the same rigorous, austere, and precise flavor as Aristotle's logical syllogisms (such as the classic "All men are mortal;

Socrates is a man; ergo, Socrates is mortal"). And indeed it was none other than Aristotle who first studied proportional analogies. For him, analogy, understood in this narrow fashion, was a type of formal reasoning belonging to the same family as deduction, induction, and abduction. The fact that many people today understand the word "analogy" in just this narrow way therefore has genuine and valid historical roots. Nonetheless, such a restrictive vision of the faculty of analogy-making leads almost ineluctably to the conclusion that it is such a precise, focused, and specialized type of mental activity that it will crop up only in very rare circumstances (particularly in intelligence tests!).

And yet analogy, as a natural form of human thought, is not by any means limited to this kind of case. Although each of the proportional analogies exhibited above was intended to have just one single correct response — the so-called *right answer* — the fact is that the world in which we live does anything but give us a long series of intelligence-test questions in the form of right-answer analogy puzzles. Thus in the case of the "Paris of the United States" puzzle given above, although we ourselves were thinking mostly of New York as "the right answer", we have collected, in informal conversations, quite a few other perfectly defensible answers, including Washington, Boston, Los Angeles, Las Vegas, Philadelphia, and — of course! — Paris, Texas.

Indeed, quite to the contrary, the world confronts us with a never-ending series of vague and ambiguous riddles, such as this one: "What disturbing experience in my life, or perhaps in the life of some friend of mine, is meaningfully similar to the sudden confiscation of my eight-year-old son's bicycle by the principal of his school?" It is by searching for strong, insight-providing analogues in our memory that we try to grasp the essences of the unfamiliar situations that we face all the time — the endless stream of curve balls that life throws at us. The quest for suitable analogues is a kind of art that certainly deserves the label "vital", and as in any other form of art, there seldom is a single right answer. For this reason, although proportional analogies may on occasion be gleaming jewels of precision and elegance, the image that they give of the nature of analogy-making is wildly misleading to anyone who would seek the crux of that mental phenomenon.

Another widely held view of analogy (and here we come to the second stereotype) is that when people make analogies, they call on sophisticated reasoning mechanisms that, through intricate machinations, somehow manage to link together far-flung domains of knowledge, sometimes in a conscious fashion; the conclusions reached thereby may be very subtle but will also be very tentative. This vision gives rise to the image of analogies as being the fruit of strokes of genius, or at least of deep and unusual insights. And there are indeed numerous famous cases of this sort that one can cite — great scientific discoveries resulting from sudden inspirations of people who found undreamt-of links between seemingly unrelated domains. Thus the mathematician Henri Poincaré wrote, "One day... the idea came to me very concisely, very suddenly, and with great certainty, that the transformations of indeterminate ternary quadratic forms were identical to those of non-Euclidean geometry." This flash of inspiration gave rise to much rich new mathematics. One can also admiringly recall various

architects, painters, and designers who, thanks to some fresh analogy, were able to transport a concept from one domain to a distant one in such a fruitful way that people were amazed. From this perspective, the making of analogies is a cognitive activity that only a small number of extremely inventive spirits engage in; it happens only when a mind dares to explore highly unlikely connections between concepts, and it reveals relationships between things that no one had ever before thought of as being related.

This stereotype of analogy-making does not presume that such acts are limited to scientists, artists, and designers; much the same vision of sophisticated reasoning that connects distant domains and leads to daring but tentative conclusions applies to people in everyday life. For example, it is universally accepted that analogy plays a major role in teaching. Most everyone can recall analogies from their schooldays, such as between the atom and the solar system, between an electrical circuit and a circuit in which water flows, between the heart and a pump, or between the benzene molecule and a snake biting its own tail. All of these cases feature connections that link rather remote domains (or, to be more precise, domains that *seem* remote when only their surface is taken into account). One can also find examples of such analogies in everyday arguments, whether someone is supporting an idea or trying to knock it down. For instance, if everyone laughs in the face of a person who dares to reveal grand ambitions, a natural retort might be, "Laugh all you want; they all laughed at Christopher Columbus!" And in political debate, analogies between far-flung situations play a key role. Thus these days, likening the leader of a foreign country to Hitler as a way to ignite patriotic fervor has become a hackneyed stratagem (for example, the elder George Bush pulled the Hitler analogy out of his hat numerous times in order to justify the first Iraq war), whereas likening a war to the Vietnam war has played precisely the opposite role in the United States (the opponents of the second Iraq war called on the Vietnam analogy over and over again). One even finds such insightful analogies, full of freshness, in childish observations, such as when the daughter of one of this book's authors, rising to the full mental height of her seven summers, proudly declared, "School is like a staircase; each new grade is one step higher!" Such a joyous moment of enlightenment is, in its humble way, an insight comparable to Poincaré's joyous insight into abstruse mathematical phenomena.

To summarize, then, the first of these two stereotypes — proportional analogy — is so formally constrained that if that were all analogy-making amounted to, it would merely be the Delaware of cognition; by contrast, the second stereotype pinpoints a far more important mental phenomenon — namely, the selective exploitation of past experiences to shed light on new and unfamiliar things belonging to another domain. And thus we will spend very little time on proportional analogies in this book; however, it's quite another matter as far as rich interdomain analogies are concerned, and we will devote a great deal of attention to them. And yet, despite its clear relevance to our central topic, this second vision of analogy-making is still impoverished, since it vastly under-represents the wide range of mental phenomena to which analogy is connected. Indeed, it completely leaves out the idea that analogy-making is the machinery behind the pulsating heartbeat of thought: categorization.

Analogy-making and Categorization

Indeed, the central thesis of our book — a simple yet nonstandard idea — is that the spotting of analogies pervades every moment of our thought, thus constituting thought's core. To put it more explicitly, analogies do not happen in our minds just once a week or once a day or once an hour or even once a minute; no, analogies spring up inside our minds numerous times every second. We swim nonstop in an ocean of small, medium-sized, and large analogies, ranging from mundane trivialities to brilliant insights. In this book, we will show how the simplest and plainest of words and phrases that we come out with in conversations (or in writing) come from rapidly, unconsciously made analogies. This incessant mental sparkling, lying somewhere below the conscious threshold, gives rise to our most basic, humdrum, low-level acts of categorization, whose purpose is to allow us to understand the situations that we encounter (or at least their most primordial elements), and to let us communicate with others about them.

A substantial fraction of the myriads of analogies constantly being born and quickly dying in our heads are made in order to allow us to find the standard words that name mundane objects and activities, but by no means all of them are dedicated to that purpose. Many are created to try to make sense of situations that we face on a much larger scale. To pinpoint, in the form of a single previously known concept, the essence of a complex situation that has just cropped up for the first time involves a much more penetrating and global understanding of a situation than one gets from simply smacking labels on its many familiar constituents. And yet this far deeper process — the retrieval of a long-buried memory by an analogy — is so central and standard in our lives that we seldom think about it or notice it at all. It is an automatic process, and virtually no one wonders why it occurs, nor how, since it is so familiar. If asked “How come that particular memory popped to your mind right after I told you what happened to me?”, a typical person might reply, with a bemused tone at being asked such a silly question, “Well, what I remembered is *very much like* what you told me. *That's* why I remembered it! How could it have been any other way?” It's as if they had been asked, “Why did you fall down?” and answered, “Because I tripped!” In other words, having X, which is in some sense very similar to Y, come to mind when Y occurs and seizes our attention seems as natural and inevitable as falling down when one is tripped — there is no mystery, hence there seems to be no need whatsoever for any explanation!

The triggering of memories by analogy lies so close to what seems to be the essence of being human that it is hard to imagine what mental life would be like without it. Asking why one idea triggers another similar one would be like asking why a stone falls if one lets go of it three feet above the ground. The phenomenon of gravity is so familiar and obvious to us, striking us as so normal and so inevitable, that no one, aside from a tiny minority of physicists obsessed with explaining what others take for granted, even sees that there is anything to ask about. For most non-physicists, it's hard to see why gravity needs an explanation — and the same holds for the triggering of memories by analogy. And yet, how many scientific discoveries can hold a candle to general relativity, Albert Einstein's wildly unexpected revelation of what gravity actually is?

Categorization and Analogy-making as the Roots of Thinking

The idea that we will here defend is that a certain mental phenomenon subsumes all the aforementioned stereotypes of categorization and analogy, but is much broader than any of them are, taken in isolation. To give a foretaste of this crucial idea, we turn once again to the theme of zeugmas, because these linguistic oddities have a great deal to do with categorization through analogy. Indeed, zeugmas provide a rich wellspring of examples running the full gamut from the most mundane to the most inspired of analogies; in their own small way, then, zeugmas perfectly reflect the ubiquity and uniformity of the mechanism of categorization by analogy.

Suppose you heard someone say, "The asparagus tips and the potato dumplings were delicious." Your ever-ready zeugma detector wouldn't register a thing, because it would seem self-evident that, in this context, *asparagus tips* and *potato dumplings* simply belong to one and the same very standard category (namely, that of *scrumptious edibles*). But it would feel rather different if someone were to say, "The asparagus tips and the after-dinner witticisms were delicious", because here one senses that the adjective "delicious" has been used in two quite different senses, and so the needle on one's zeugma detector would move a bit, and as a result you would feel that a slight analogical link had been suggested between the asparagus tips and the postprandial quips. Then again, were you to hear "The asparagus tips and the expression of surprise on Anna's face were delicious", your zeugma detector would register a yet higher reading, meaning that the semantic distance (or interconceptual stretch) was yet greater; this would lead you to see and feel an analogy between the asparagus tips and a certain friend's facial expression, rather than merely thinking that they both belong to the commonplace category of *delicious things*.

In brief, it is misleading to insist on a clear-cut distinction between analogy-making and categorization, since each of them simply makes a connection between two mental entities in order to interpret new situations that we run into by giving us potentially useful points of view on them. As we will show, these mental acts cover a spectrum running from the humblest recognition of an object to the grandest contributions of the human mind. Thus analogy-making, far from being merely an occasional mental sport, is the very lifeblood of cognition, permeating it at all levels, ranging all the way from mundane perceptions ("That is a table") to subtle artistic insights and abstract scientific discoveries (such as general relativity). Between these extremes lie the mental acts that we carry out all the time every day — interpreting situations, judging the quality of various things, making decisions, learning new things — and all these acts are carried out by the same fundamental mechanism.

All of these phenomena seem quite different, but underlying them all there is just one single mechanism of nonstop categorization through analogy-making, and it operates all along the continuum we've described, which stretches from very mundane to very sophisticated acts of categorization. And it's this unified mechanism that allows us to understand sentences that run the gamut of zeugmaticity, from complete non-zeugmas (requiring only mundane categorization mediated by very basic analogy-

making) to extreme zeugmas (requiring unusually flexible categorization mediated by much more sophisticated analogy-making).

But let's take our leave of zeugmas and return to the larger picture. We claim that cognition takes place thanks to a constant flow of categorizations, and that at the base of it all is found, in contrast to *classification* (which aims to put all things into fixed and rigid mental boxes), the phenomenon of *categorization through analogy-making*, which endows human thinking with its remarkable fluidity.

Thanks to categorization through analogy-making, we have the ability to spot similarities and to exploit these similarities in order to deal with the new and strange. By connecting a freshly encountered situation to others long ago encountered, encoded, and stored in our memory, we are able to make use of our prior experiences to orient ourselves in the present. Analogy-making is the cornerstone of this faculty of our minds, allowing us to exploit the rich storehouse of wisdom rooted in our past — not only labeled concepts such as *dog*, *cat*, *joy*, *resignation*, and *contradiction*, to cite just a random sample, but also unlabeled concepts such as *that time I found myself locked outside my house in bitterly freezing weather because the door slammed shut by accident*. Such concepts, be they concrete or abstract, are selectively mobilized instant by instant, and nearly always without any awareness on our part, and it is this ceaseless activity that allows us to build up mental representations of situations we are in, to have complex feelings about them, and to have run-of-the-mill as well as more exalted thoughts. No thought can be formed that isn't informed by the past; or, more precisely, we think only thanks to analogies that link our present to our past.

The Rapid Inferences that Categories Provide

A term that will be useful to us in this context is *inference*. As is traditional in psychology, we will use the term much more broadly than it is used in the field of artificial intelligence, where it is synonymous with "formal logical deduction", as carried out by so-called "inference engines". By contrast, what we will mean by "making an inference" is simply the introduction of some new mental element into a situation that one is facing. Basically, this means that some facet of a currently active concept is lifted out of dormancy and brought to one's attention. Whether this new element is right or wrong is not the point, nor does it matter whether it follows logically from prior elements. For us, "inference" will simply mean the fact that some new element has been activated in our mind.

Thus if one sees a child crying, one infers that the child is distressed. If one sees someone shouting, one infers that the person is probably angry. If one sees that the table is set, then one infers that a meal may well soon be served. If one sees a door that is closed, one infers that it can be opened. If one sees a chair, one infers that one could sit on it. If one sees a dog, one has the ability to infer (though one does not necessarily do so) that it barks now and then, that it might bite someone, that it has a stomach, a heart, two lungs, and a brain — internal organs that one doesn't strictly perceive but that category membership allows one to infer. Inferences of this sort are a crucial

contribution to thought, and they come from categorization through analogy, for we rely ceaselessly on resemblances perceived between the present situation and ones we encountered earlier. If we did not do this at all times, we would be helpless.

Thus, it is not merely for idle fun that one calls a cat-like thing that one encounters "cat", thereby assigning it to a preexisting category in one's memory; it is principally because doing so gives one access to a great deal of extra information, such as the likely fact that it will show pleasure by purring, that it has a propensity to chase mice, that it may well scratch when threatened, tends to land on its feet, has a very autonomous character... These kinds of things, among others, can all be inferred about an entity once it has been assigned to the category *cat*, without any of them having been directly observed about the specific entity in question. Thus our categories keep us well informed at all times, allowing us to bypass the need for direct observation. If we didn't constantly extrapolate our knowledge into new situations — if we refrained from making inferences — then we would be conceptually blind. We would be unable to think or act, doomed to permanent uncertainty and to eternal groping in the dark. In short, in order to perceive the world around us, we depend just as much on categorization through analogy as we do on our eyes or our ears.

Analogy's Champions and Detractors

Some ancient philosophers, including Plato and Aristotle, were fervent defenders of analogy, seeing it as a fertile medium for thinking rather than as just a figure of speech. Nonetheless, these same thinkers felt compelled to point out its limitations. Thus Plato, using a number of analogies — among them one likening a soul to a city, in his famous work *The Republic* — warned that "likeness is a most slippery tribe". And Aristotle, although just as great an admirer of analogies, cast aspersions on many analogies made by his predecessors. Thus we see that even for some of its strongest backers, analogy has a faintly suspicious aroma, as does its cousin, metaphor. In the minds of such doubters, these two figures of speech, when used ill-advisedly, are liable to mislead both those who utter them and those who hear them.

Immanuel Kant and Friedrich Nietzsche had extremely different personalities, philosophies, and views about religion, but they were united in their unswaying belief in analogy. For Kant, analogy was the wellspring of all creativity, and Nietzsche gave a famous definition of truth as "a mobile army of metaphors". However, analogy has certainly not had such good press universally. Indeed, it's been a favorite pastime down through the centuries to berate analogy for its unreliability, its closeness to wild guessing, and the serious traps into which it leads anyone who depends on it. Some philosophers have had quite a field day denouncing analogy and metaphor, describing them as superficial, misleading, and useless forms of thought.

In particular, the empiricists in the seventeenth century and the positivists in the twentieth raked analogy and metaphor over the coals. The English philosophers Thomas Hobbes and John Locke are often quoted in this regard. Hobbes, in *Leviathan*, his best-known work, declares his love for clear words and his scorn for metaphors:

[T]he light of human minds, is perspicuous words, but by exact definitions first snuffed and purged from ambiguity; ... [M]etaphors, and senseless and ambiguous words, are like *ignes fatui*; and reasoning upon them is wandering amongst innumerable absurdities.

Hobbes leaves no doubt as to his views. Truth is light, words must be cleansed and purged of ambiguity, and metaphors are nothing but will-o'-the-wisps that would lead one to wander in a wacky world. However, if one stops to look at this passage for a moment, one is struck by a certain ironic quality — specifically, the fact that its author condemns metaphors not by using “snuffed and purged definitions” but through the repeated use of metaphors. After all, what kind of phrase is “the light of the mind”? What about “definitions that have been snuffed and purged”? And how about “wandering amongst innumerable absurdities”? What are all these phrases, if not metaphors? Does a mind really contain light? Can definitions actually be cleansed? And are metaphors in truth unpredictably flickering lights hovering above a swamp?

In his protest, Hobbes is a bit like someone who screams in order to praise silence, or like evangelistic television preachers who whip up the masses by speaking of the sins that lead straight to hell while themselves engaging in the very acts of debauchery that they decry. His protest is also reminiscent of a paradoxical phrase that encapsulated the tragedy of the Vietnam war: “We destroyed the village in order to save it.” In short, Hobbes undermines his anti-metaphor credo by expressing it metaphorically.

The eleventh-century Benedictine monk Alberic of Monte Cassino never knew anything approaching the fame of Hobbes, but he too wrote a virulent diatribe against the use of metaphors in his book *The Flowers of Rhetoric*. Here is an excerpt:

Expressing oneself with metaphors has the quality of distracting a person's attention from the specific qualities of the object being described; in one manner or another, this distracting of attention makes the object resemble something different; it dresses it, if one may put it thusly, in a new wedding dress, and in so dressing it, it suggests that a new kind of nobility has been accorded to the object... Were a meal were served in this fashion, it would disgust and nauseate us, and we would discard it... Consider that in one's enthusiasm for giving pleasure through delicious novelty, it is unwise to begin by serving up flapdoodle. Be careful, I repeat, when you invite someone in the hopes of giving pleasure, that you not afflict him with so much malaise that he will vomit from it.

As we glide from “dressing an object in a new wedding dress” to “serving up flapdoodle” and “vomiting from it”, we are treated to one metaphor after another in a passage written for no other purpose than to criticize the use of metaphors.

Eight centuries later, Gaston Bachelard, a highly respected French philosopher of science, did not completely avoid the same trap when he wrote: “A science that accepts images is, more than any other, a victim of metaphors. Consequently, the scientific mind must never cease to fight against images, against analogies, against metaphors.” But how can science become a “victim”, and how can a mind, scientific or otherwise, “constantly fight” against anything, unless they do so metaphorically?

Are Analogies Seductive and Dangerous Sirens?

And so, are analogies like seductive and dangerous siren songs, likely to lead us astray, or are they more like indispensable searchlights, without which we would be plunged in total darkness? If one never trusted a single analogy, how could one understand anything in this world? What, other than one's past, can one rely on in grounding decisions that one makes when facing a new situation? And of course all situations *are* in fact new, from the largest and most abstract ones down to the tiniest and most concrete ones. There isn't a single thought that isn't deeply and multiply anchored in the past.

To use the elevator in an apartment building that one has never been in before, does one not tacitly depend on the analogy with countless elevators that one has used before? And when one examines this analogy, one sees that, despite its seeming blandness, it depends on numerous others. For example, once you've entered the elevator, you have to choose a small button you've never seen before, and you have to press it with a certain finger and a certain force, and you do that without thinking about it whatsoever (or more accurately, without noticing that you are thinking about it). This means that you are unconsciously depending on your prior experiences with thousands of buttons in hundreds of elevators (and also buttons on keyboards, stereo systems, dashboards, etc.), and that you are working out the best way to deal with this new button by relying on an analogy between it and your personal category *button*.

And when, after you've stepped out of the elevator and are just setting foot in the sixth-floor apartment, you see a big dog coming towards you, how do you deal with this situation if not on the basis of your prior experience with dogs, particularly large dogs? And much the same could be said for when you wash your hands in the sink that you've never seen before with soap that you've never touched before — not to mention the bathroom door, the doorknob, the electric switch, the faucet, the towel, all never before seen or touched.

And if you go into a grocery store that you've never seen before and are looking for the sugar or the olives or the paper towels, where do you go? Which aisle, which shelf, and how high up on the shelf? Without any conscious effort, you recall "the" spot where these articles are found in other familiar stores. Of course you're not thinking of just *one* place, but of a collection of various places that you mentally superimpose. You think, "The sugar should be around *here*", where the word "here" refers simultaneously to a collection of small areas in various familiar grocery stores and also to a small area in the new store, and it's "right there" that one looks first of all.

How mundane is the scene of an employee who, requesting an extra day of vacation, says to her boss, "Last year you offered an extra weekend to Katyanna, so I was wondering if you would be able to give me just one extra day next month..." How could one do anything in life if one felt that it was crucial to be constantly on the alert in order to mercilessly squelch any resemblance that came to mind at any level of abstraction or concreteness? And worse yet, once we'd squelched them all, what would we then do? On what basis would we make even the tiniest decision?

Might there be a rigorous proof that all analogies are dubious? Obviously not, because, as we just saw, everyone depends, without thinking, on a dense avalanche of mini-analogies between everyday things, and these mini-analogies follow on the heels of one another all day long, day in, day out — and seldom do such mundane analogies mislead anyone. Indeed, if they did, we would not be here to tell the tale.

Giant Electronic Dunces

How can computers be so terribly stupid, despite being so blindingly fast and having such huge and infallible memories? Contrariwise, how can human beings be so insightful despite being so limited in speed and having such small and fallible memories? Though perhaps hackneyed, these are reasonable and important questions, focusing as they do on the nearly paradoxical quality of human thought.

Indeed, the human mind, next to a computer, appears fraught with defects of every sort, coming off as hopelessly inferior along most dimensions of comparison. For instance, in carrying out pure reasoning tasks, well-polished computer algorithms reach logically valid conclusions virtually instantly, while people tend to fail most of the time. Much the same can be said about large amounts of knowledge. Where people's minds are saturated after only a few pieces of information are presented, a computer can take into account a virtually unlimited amount of information. And of course human memory is notoriously unreliable; whereas computers never forget and never distort, those are activities at which we human beings excel, for better or for worse. Three days, three weeks, three months, or three years after we've seen a movie or read a book, what details of it remain accessible in our minds? And how distorted are they? We might also mention the speed at which processing takes place in computers as opposed to human brains. What might take us minutes, hours, or far longer can be done by a computer in an infinitesimal eyeblink. Just consider simple arithmetical calculations such as "3 + 5" (a bit under a second for a person), or "27 + 92" (perhaps five or ten seconds), or "27 x 92", a calculation that most people could not carry out in their heads. Counting the number of words in a selected passage of text and correcting a multiple-choice exam are activities that we humans can carry out, but only with pathetic slowness compared to computers.

Overall, the comparison is extremely lopsided in favor of computers, for, as we just noted, computers carry out flawless reasoning and calculation way beyond human reach, handle unimaginably larger amounts of information than people can handle, do not forget things over short or long time scales, do not distort what they memorize, and carry out their processing at speeds incomparably greater than that of the human mind. In terms of rationality, size, reliability, and speed, the machines we have designed and built beat us hands down. If we then add to the human side of the ledger our easily distractable attention, the fatigue that often seriously interferes with our capacities, and the imprecision of our sensory organs, we are left straggling in the dust. If one were to draw up a table of numerical specifications, as is standardly done in comparing one computer with another, *Homo sapiens sapiens* would wind up in the recycling bin.

Given all this, how can we explain the fact that, in terms of serious thought, machines lag woefully behind us? Why is machine translation so often inept and awkward? Why are robots so primitive? Why is computer vision restricted to the simplest kinds of tasks? Why is it that today's search engines can instantly search billions of Web sites for passages containing the phrase "in good faith", yet are incapable of spotting Web sites in which the *idea* of good faith (as opposed to the string of alphanumeric characters) is the central theme?

Readers will of course have anticipated the answer — namely, that our advantage is intimately linked to categorization through analogy, a mental mechanism that lies at the very center of human thought but at the furthest fringes of most attempts to realize artificial cognition. It is only thanks to this mental mechanism that human thoughts, despite their slowness and vagueness, are generally reliable, relevant, and insight-giving, whereas computer "thoughts" (if the word even applies at all) are extremely fragile, brittle, and limited, despite their enormous rapidity and precision.

As soon as categorization enters the scene, the competition with computers takes on a new kind of lopsidedness — but this time greatly in favor of humans. The primordial importance of categorization through analogy in helping living organisms survive becomes obvious if one tries to imagine what it would be like to "perceive" the world in a manner entirely devoid of categories — something like how the world must appear to a newborn, for whom each new concept has to be acquired from scratch and with great difficulty. By contrast, seeing the new in terms of the old and familiar allows one to benefit, and at only a slight cognitive cost, from knowledge previously acquired. Thus, if there were two creatures, one of which (an adult human being) perceived the world using categorization through analogy while the other (a computer) had no such mechanism to help it out, their competition in understanding the world around them would be comparable to a race between a person and a robot to climb up to a high roof, with the human allowed to use a preexistent staircase but with the robot required to construct its own staircase from scratch.

Analogy Operating at All Levels

Categorization through analogy drives thinking at all levels, from the smallest to the largest. Consider a conversation in which several hierarchical linguistic levels are continually interacting. First of all, the choice of a specific word will of course determine the sounds that make it up; similarly, when one is typing at one's keyboard, each word chosen determines the letters composing it, so that they come along automatically rather than being chosen one by one. Analogously, words are often determined by larger structures of which they are but pieces. This happens most clearly whenever one uses a stock phrase (such as "so to speak" or "cut to the chase" or "down to the wire" or "when push comes to shove" or "as easy as stealing candy from a baby"), but it also often happens when no such expression is involved, because one is always working under the constraints of the syntactic and semantic patterns of the language one is speaking, as well as those of one's own habitual speech patterns.

And the same principle holds at more global levels of speech as well. Thus when one writes or utters a sentence, many of the words comprising it come along without being chosen one by one, since they are all serving a higher-level goal that has been pre-selected. Thus, much as with letters being constrained by a word, the words are in a sense constrained by higher-level thoughts. And then, moving yet further upwards, we can say that the same holds when one is developing an idea; that is, the sentences one produces to express this idea are once again constrained by a yet higher-level structure, even if there is more freedom at this level than at the letter-choice level. And the same holds at the level of the conversation itself, because its overall topic, its tone, the particular people involved in it, and so forth, all constrain the ideas that will be thought of. Of course at this level, there is much more flexibility than at the level of letters composing words. And so, in summary, a conversation constrains the ideas in it, the ideas constrain the sentences, the sentences constrain the phrases, the phrases constrain the words, and finally, the words constrain their letters.

Our claim that choices on each of these levels are carried out by categorization by analogy runs against the naïve image of categories as corresponding, more or less, to single words. To be sure, some categories are indeed named by words, but others are far larger, residing essentially at the level of an entire conversation.

For example, consider arguments about the size of the military budget. Those who advocate a large budget frequently trot out the same old arguments over and over again, based on the vital need to protect our nation against unnamed threats of all sorts, the intense pressure to develop ever newer technologies, the idea that advances in military technology help to drive the civilian marketplace, and so forth. Such a line of reasoning can be spun out over a long time, while always depending on a well-known, even hackneyed, conceptual skeleton that has been "seasoned to taste", depending on the context, the occasion, and so forth. But whatever the variations on the theme are, it's always the same conceptual skeleton centered on the need for national defense and for advances in technology. The high-level category determining the overall flow of one's argument is defined by this conceptual skeleton.

Conversely, advocates of trimming the military budget will almost invariably cite the enormous importance of other sectors of the economy and the great inefficiencies in the military. Here once again, such arguments can be spun out at great length, but however they run, they will always be centered on the bloatedness of the current military budget and the crying need for funding other sectors of the economy. This is the familiar conceptual skeleton that will guide the overall flow of the argument.

And thus we see that at the top level of the conversation, we are dealing with the very high-level categories *need for a bigger military budget* and *need for a smaller military budget*, and the activation of either category in an advocate's mind will trigger, with a bit of variability, yet also a considerable degree of predictability, the auxiliary ideas that will pop to mind, and these will promptly enlist appropriate stock phrases and well-worn grammatical patterns, which will in their turn call up the standardized words that comprise them — and these words, in the end, will call up, with essentially no maneuvering room, the letters or sounds that make them up.

One can examine any conversation, whether it's a deep or a shallow one, in this fashion, and one will see how analogies, at all the different levels, are in the driver's seat. Here's a rather light-hearted example based on a real event.

One Saturday evening, Glen and Marina Bayh had a few friends over for dinner. The food was savory, the wine and witticisms flowed copiously, and at last, around midnight, people started rising to get their coats. As they were filing out the front door, Larry Miller, one of the guests, said warmly to Glen and Marina, "It was a terrific evening. Haven't had so much fun in a long time. Thanks a lot. Hope to see you again soon. Bye-bye!" On hearing this innocent remark, Jennifer, another guest, commented, "I always have a hard time saying good-bye to them." Larry, puzzled, replied to her, "But all good things come to an end. We had a great time, it finished, and now we're taking off. What's the big deal about saying bye?" Jennifer answered, "Yeah, you're right, but still it sounds weird, because 'Bayh' is their last name. I mean, for them to hear 'Bye' or 'Bye-bye' all the time must be a bit like it would be for you to hear people saying 'Miller, Miller' all the time, no? That would come across as strange to you, wouldn't it?" Larry burst out laughing and said, "I guess I'm just dense! I'd never thought of that!" Right then, Larry's wife Colleen chimed in, saying, "It reminds me of when I was a teen-ager and every time my parents took me and my brother to our grandmother's house, he and I would always whisper to each other, 'Now don't forget — you have to mind your grandma around here!' We always called her 'Gramma', and she was always correcting our English, so this was our private little way of getting back at her, though she never knew it." Everyone could easily relate to Colleen's story, and of course they all understood why her comment wasn't a *non sequitur* coming out of left field but a perfectly apt segue.

What are the key analogies behind the scenes of this down-to-earth interchange? It was launched by an analogy between the sounds of the words "Bayh" and "bye", which spurred Jennifer to invent an analogy between the last names "Bayh" and "Miller", after which the surface-level topic veered off to visits paid decades earlier to the home of a relative of Colleen's, but still driven by the momentum of analogy — this time an analogy between the name "Gramma" and the word "grammar". At a higher level, however, the trigger that sparked the retrieval of this dormant memory in Colleen's mind was the similarity between a fresh episode and a long-ago episode, both of which involved humorous phonetic resemblances between a normal word and someone's name — in one case, that of "bye" to "Bayh" and in the other case, that of "grammar" to "Gramma". So here we're dealing with a similarity between resemblances — which is to say, with an analogy between analogies.

There is nothing unusual about this conversation or this type of analogy-spotting behind the scenes. It's all par for the course. We exhibited it simply to show how a conversation as a whole mobilizes one or two brigades at a very high conceptual level, how those high-level concepts mobilize a few lower-level conceptual regiments, how these in turn mobilize a larger number of conceptual platoons or squadrons more or less at the level of stock phrases, and finally, how these many "smaller" concepts mobilize hundreds of individual soldiers way down at the word level.

Abstract or Concrete?

What lies behind this universality of analogy-making? In order to survive, humans rely upon comparing what's happening to them *now* with what happened to them in the past. They exploit the similarity of past experiences to new situations, letting it guide them at all times in this world. This incessant flow of analogies, made in broad brushstrokes, forms the crux of our thoughts, and our utterances reflect them, although our specific word choices are usually fast forgotten. The concrete meets the abstract when a down-to-earth phrase is applied to describe a down-to-earth situation but where the concepts that the phrase is built from are distant, on a literal level, from the situation. For instance, in an idiomatic utterance such as "Marie is off her rocker" or "Their love affair went down the drain", the thought is at such a high level of abstraction that one seldom will consciously envision someone falling off a rocking chair or water flowing out of a sink or bathtub.

Much the same happens when a new situation reminds you of another situation (or family of situations) that you previously encountered and that is, on its surface, totally different, but that shares an abstract essence with the new one. Thus if one day your child is not allowed to register for a crucial event at school because the relevant Web site's deadline was at 4 o'clock sharp and you logged onto the site at 4:01, this may summon up a long-buried memory from some fifteen years earlier of a time when you missed a plane because after dashing to the gate, you arrived just seconds after the doors had been closed, and no matter what you said, they wouldn't let you board.

Our daily talk is filled with this kind of meeting of concrete and abstract, but we are unaware of it most of the time. Thus if a professor says, "Only a handful of students dropped in on me in my office yesterday", we aren't likely to envision students tumbling out of a giant hand in the sky and landing in the professor's office. And if someone says, "There hasn't been any snow to speak of today", we don't feel inclined to protest, "What do you mean? You just *spoke* of it!" The concreteness of the words and phrases that we constantly use to formulate our thoughts on all sorts of topics is at one and the same time a sign of the great concreteness of our way of thinking and also a sign of our extraordinary propensity to carry out abstraction, allowing us to cast a situation using words that would seem, on their surface, to refer to totally unrelated things.

Thus a Japanese stockbroker, commenting on the unstoppable collapse of the stock market, said, "One should never try to catch a falling knife." Most people understand this image effortlessly, as well as its relevance to the circumstances. And yet if there is a falling knife here, it's certainly not a knife in its most ordinary sense, and the way this "knife" is falling is invisible, comparatively slow, and not spatially localized. It took a considerable act of abstraction for that individual to use that phrase in that context — but that's nowhere close to the end of the story, for the same phrase might just as easily be applied to a politician in the throes of a corruption scandal and whose once-ardent supporters are suddenly nowhere to be seen, or to a skyscraper on fire that it would be folly to enter, or to someone so deeply depressed that even their best friends are caught up in the atmosphere of bleak pessimism, or to someone who has fallen overboard in a

storm so violent that no one dares to go to their aid, or to an approaching hurricane that has destroyed a nearby town and from which everyone has been ordered to evacuate immediately, or perhaps, who knows, even to a person who was injured in their kitchen because they tried (of all things) to catch a falling knife. In short, we see that here we are dealing with a full-fledged category, rich and multifaceted.

As soon as one starts thinking about situations to which the phrase "One should never try to catch a falling knife" could apply, they start to pop up from under nearly every stone in sight. At least for a while, one has the impression that one could paint a large fraction of the world in terms of that phrase alone, since the world is rife with huge, irrepressible forces against which one has no power and which would carry one off to one's doom if one were so rash as to try to stop them. Thus we will find this image jumping to our mind willy-nilly, imposing itself on us whether we like it or not, and unless we can somehow stop thinking altogether, we will simply have to let this aggressive metaphor have its way with us until it has had its fill; after all, one shouldn't try to catch a falling knife.

Synopsis of This Book

In the chapters that follow, we don't aim to speak about the brain at a biological level, but about cognition as a psychological phenomenon. We will not speculate about the cerebral or neural processes that underlie the psychological processes we describe, because our goal is not to explain cognition in terms of its biological substrate but to present an unconventional viewpoint concerning what thought itself is. Our discussion will thus take place at this rather abstract level, but even at that level there will be plenty of grist for the mill.

The book's first three chapters constitute our attempt to provide an account of what categories and analogies are. Chapter 1 focuses on categories associated with single words, and it also puts forth some of the key theses of the book. We show how concepts designated by a single word are constantly having their boundaries extended by analogies. We take a careful look at the development of concepts by observing the long progression that starts with the concept of a child's *Mommy* (a specific adult human) and that gradually leads to all sorts of metaphorical uses such as *motherland*, passing en route through such concepts as *birth mother* and *surrogate mother*. We also show that less concrete words, such as "thanks", "much", "to fix", "to open", "but", "and" (and so on) are, no less than nouns, the names of mental categories that are the outcome of a lifelong series of analogies.

In Chapter 2, we study concepts having lexical labels that are longer than single words. We show that hidden behind the scenes of multi-word stock phrases, even ones as long as a proverb or a fable, there lie concepts that are very similar to those designated by isolated words. Thus a phrase such as "Achilles' heel" is the linguistic "hat" worn by a particular category (namely, the category of *serious weaknesses that may lead to someone's undoing*). Aesop's famous fable in which a fox tries to reach some luscious-looking grapes and when he fails, he declares that he didn't want them anyway because

they are sour, is a linguistic embodiment of the abstract mental category of situations featuring something that is the object of someone's ardor, but that, having turned out to be out of reach, is subsequently deprecated by the person who desired it. This abstract quality, often concisely called "sour grapes", is potentially recognizable in thousands of situations, and this phrase could thus be used as the verbal label of any such situation, in just the same way as there are myriads of objects meriting the label "bottlecap" and myriads of actions meriting the label "retrieve". And the same can be said of more abstract categories, some of which have to do with the act of communication taking place at the moment, and which are labeled by adverbial phrases, such as "after all", "on the other hand", "as a matter of fact", "that having been said", and so on. In other words, there are situations in our everyday interchanges that call for the label "after all", and when such situations arise, we recognize them (almost always unconsciously) as such, and we apply that label, deftly inserting it into our real-time speech stream. The chapter concludes with a discussion of intelligence as the ability to put one's finger on what counts in any given situation, and how the repertoire of categories that is handed to one by one's native language and culture tailors one's way of doing this.

Chapter 3 deals with categories for which there is no standard linguistic label; people manufacture such categories spontaneously on their own as they deal with new situations in their complex personal worlds. Later on, such categories often give rise to "reminding episodes", where one event recalls another from another time and place, possibly very distant. As an example, when D. noticed an old friend leaning down to pick up a bottlecap in Egypt's renowned Karnak Temple, he was suddenly and spontaneously reminded of a time, some fifteen years earlier, when his one-year-old son was sitting near the edge of the Grand Canyon and, completely oblivious to the spectacular scenery, was intently focused on some ants and leaves on the ground. Despite all the superficial differences that can be found separating any two situations from each other, when such a reminding incident takes place, it reveals that the two situations in question share a conceptual skeleton at a deeper level, and it shows how extremely rich and subtle is our storehouse of non-lexicalized concepts. By analyzing a series of sentences containing such high-frequency phrases as "me too", "next time", and "like that", we show that lurking behind any such phrase, no matter how casual and simple it may seem, there is a non-verbalized category, sometimes simple and sometimes subtle, based on an implicit perception of sameness, which is to say, based on an analogy.

Chapter 4 deals with the way in which, in our interactions with the world around us, we constantly and fluently move about in our repertoire of categories, and yet nearly always without the least awareness of doing so. The chapter focuses particularly on inter-category leaps that involve shifts between levels of abstraction. The flexibility of human cognition relies profoundly on our ability to move up or down the ladder of abstraction, for the simple reason that sometimes it is crucial to make fine distinctions but other times it is crucial to ignore differences and to blur things together in order to find commonalities. For instance, while one is dining, one will take care to distinguish between one's own glass and that of one's neighbor, but afterwards, when one is placing

them in the dishwasher, that distinction will be irrelevant. As another example, parents will try to assure that their children get involved in “activities”, whether this means acting in plays, doing judo, or playing a musical instrument. *Activity for my child* is a highly abstract category. The most humble of our acts conceals choices of abstraction that are hard for us to recognize accurately, because such acts are central to cognition. We have a very hard time “seeing” our cognitive activity because it is the medium in which we swim. The attempt to put our finger on what counts in any given situation leads us at times to making connections between situations that are enormously different on their surface and at other times to distinguishing between situations that on first glance seem nearly identical. Our constant jockeying back and forth among our categories runs the gamut from the most routine behaviors to the most creative ones.

Chapter 5 is devoted to the role of analogy in very ordinary, everyday situations. It deals with analogies that, because they are essentially invisible, manipulate us. We are unaware of being taken over by an analogical interpretation of a situation. In this sense, the invisible analogy manipulates us because it has simply imposed itself on us, willy-nilly. And it manipulates us also in another sense — namely, it foists new ideas on us, pushing us around. Unsatisfied with being merely an agent that enriches our comprehension of a situation we are facing, the analogy rushes in and structures our entire view of the situation, trying to make us align the newly encountered situation with the familiar old one. For instance, when a small private plane crashed into a building in Manhattan on October 11, 2006, the analogy with the events of September 11, 2001 was irrepressible, leading instantly to speculations of terrorism, even though the building was not seriously damaged; the Dow Jones average even took a noticeable nosedive for a short while. And thus analogies just jump in uninvitedly, thinking and making decisions for us, without our being aware of what is going on.

In Chapter 6, by contrast, we deal with analogies that, in some sense, we ourselves manipulate — analogies that we freshly and deliberately construct when we run into a situation that arouses our interest, sometimes in order to explain it to ourselves or to others, sometimes to argue for our own point of view. This is especially the case for what we have dubbed *caricature analogies*. These are analogies that one dreams up on the spur of the moment in order to convince someone else of an idea in which one believes. They transpose a situation into a new domain while exaggerating it. For instance, a scientist seeking a job abroad wrote to a colleague: “I love my native land, but trying to get research done here is like trying to play soccer with a bowling ball!” Also discussed in this chapter is the way in which political decisions at all levels flow from analogies perceived by decision-makers between current situations and historical events, our main case study being some of the key analogies that shaped the Vietnam war. A few studies of inter-language translation conclude this chapter, focusing on the analogies used by skilled people in order to create coordinated parallelisms between two languages and two cultures on many scales, ranging from the very small to the very large.

Chapters 7 and 8 deal with analogy in scientific thinking. Chapter 7 is concerned with what we call “naïve analogies” — in particular, the kinds of analogies on which nonspecialists tend to base their notions of scientific concepts. We show that notions

that one picks up in school, whether in mathematics, physics, or biology, are acquired thanks to appealing and helpful but often overly simple analogies to concepts with which one is already familiar. Thus an elementary arithmetical operation such as division, which supposedly is totally under one's belt by the time one starts middle school, is generally still rooted (even in the case of most university students) in a naïve analogy with the down-to-earth operation of *sharing* (as in the act of distributing 24 candies evenly to 3 children). To be sure, sharing is quite often a perfectly good way to look at division, but the view that it affords of the phenomenon is overly narrow. For example, this naïve view of division makes it very hard for people to devise a word problem that involves a division whose answer is *larger*, rather than *smaller*, than the quantity being divided. This chapter analyzes the implications, both positive and negative, of naïve analogies for education.

Chapter 8 then looks at the extreme other end of the spectrum — namely, how great discoveries are made by insightful scientists. We show how the history of mathematics and physics consists of a series of snowballing analogies. By examining from up close certain great moments in the history of these disciplines, we reveal the crucial role played over and over again by analogies — sometimes very obvious ones, sometimes very hidden ones. In particular, the deep analogies of Albert Einstein play a starring role in this chapter, including a little-known analogy that led to his hypothesis in 1905 that light consists of particles, an idea that was mightily resisted by the entire physics community for nearly two decades. The most carefully examined historical episode is that of Einstein's own slow and gradual process of coming to understand the various levels of meaning of his celebrated equation " $E = mc^2$ ".

The epilogue to our book is a dialogue — thus it is entitled "Epidialogue" — in which categorization and analogy-making are compared and contrasted along many dimensions, and although at first the two processes may seem very different, at the end of this careful comparison, the spirited debaters conclude that there is no difference between them, and they realize that in fact they are one and the same.

