SURFACES AND ESSENCES

ANALOGY AS THE FUEL AND FIRE OF THINKING

DOUGLAS HOFSTADTER
& EMMANUEL SANDER

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

Grateful acknowledgment is hereby made to the following individuals and organizations for permission to use material that they have provided or to quote from sources for which they hold the rights. Every effort has been made to locate the copyright owners of material reproduced in this book. Omissions that are brought to our attention will be corrected in subsequent editions.

Photograph of Mark Twain: © CORBIS
Photograph of Edvard Grieg: © Michael Nicholson/CORBIS
Photograph of Albert Einstein: © Philippe Halsman/Magnum Photos
Photograph of Albert Schweitzer: © Bettmann/CORBIS

We also most warmly thank Kellie Gutman and Tony Hoagland for their generous permission to publish their poems in this volume.

Published by Basic Books
A Member of the Perseus Books Group

Designed by Douglas Hofstadter
Cover by Nicole Caputo and Andrea Cardenas

Library of Congress Control Number: 2013932688

ISBN: 978-0-465-02158-1 (e-book)

10 9 8 7 6 5 4 3 2 1

TABLE OF CONTENTS

WORDS OF THANKS

A Little Background

As we look back, fondly reliving the genesis of our book, we vividly recall the first key moment, which took place in mid-July, 1998, at an academic congress, in Sofia, Bulgaria. The occasion was the first international conference on the subject of analogy. Organized by Boicho Kokinov, Keith Holyoak, and Dedre Gentner, this memorable meeting assembled researchers from many countries, who, in an easy-going and lively atmosphere, exchanged ideas about their shared passion. Chance thus brought the two of us together for the first time in Sofia, and we found we had an instant personal rapport — a joyous bright spark that gradually developed into a long-term and very strong friendship.

In 2001–2002, Douglas Hofstadter spent a sabbatical year in Bologna, Italy, and during that period he was invited by Jean-Pierre Dupuy to give a set of lectures on cognition at the École Polytechnique in Paris. At that time, Emmanuel Sander had just published his first book — an in-depth study of analogy-making and categorization — and at one of the lectures he proudly presented a copy of it to his new friend, who, upon reading it, was delighted to discover how deeply similar was the vision that its author and he had of what cognition is really about. And then some time passed, with a couple of brief get-togethers in Paris and Toulouse, supplemented by email exchanges and phone calls, mixing intellectual content and friendly feelings.

In February of 2005, Doug invited Emmanuel to Bloomington for a gala birthday party he was throwing for his many friends, as he turned 60. One day during that event, he suggested to Emmanuel that he would be very happy to come to Paris for a few weeks to work with Emmanuel in translating his book into English. Emmanuel was very pleased, but shortly after Doug arrived in Paris in July, the original goal mutated into a considerably larger one — namely, that of co-authoring a book on the fundamental role of analogy in thought, approaching the topic in a non-technical manner and from many points of view, and using a large sampling of concrete examples to justify the theoretical positions. The book would hopefully be accessible to anyone interested in thinking, yet would also have the high ambition of reaching an academic audience and of putting forth a new and original stance towards cognition. This thus was the moment of our book's conception!

Over a three-week period in Paris, many ideas were tossed about, and the result

was a forty-page document that featured snips of conversations between the two future co-authors, many notes on ideas for the book, and a very preliminary sketch of what its chapters might look like. And then, for the next four years — 2006 to 2009 — each of the authors made a month-long visit to the other one in his home town. Adding to that, Doug spent an eight-month sabbatical in Paris in 2010. During all this time, there was a constant exchange of ideas via email and via phone, allowing the book to evolve from a few cells into a viable complex organism.

As all this shows, the present book is the fruit of a long collaboration, and finally it has reached maturity. Its authors have invested in it the hope that it has a message with enduring value, even if it is clearly rooted in today's culture and style of life — in fact, it is rooted "in vibrant thought", as we fondly recall one of our friends putting it. But we hope that, despite the spatial and temporal specificity of its origins, its key ideas are universal enough that they will withstand the passage of time.

Ping-ponging between Languages and Cultures

We are quite proud of the fact that our joint book is the result of a very unusual creative process. Not just written by two people, it was written in two languages at the same time. To be more specific, this book has two originals — one in French and one in English. Each is a translation of the other, or perhaps neither of them is a translation. But however you choose to look at it, the two versions of this book have equal standing. They are two highly distinct concrete incarnations of one immaterial entity — namely, this book as it exists on the ethereal plane not of words but of ideas.

To be sure, the writing process involved countless acts of translation, but those acts took place at the very moment that the original text was being generated. Sometimes they carried ideas from English to French, and sometimes they went in the other direction, but what is key here is that these back-and-forth exchanges between the brains of the two authors were accompanied — and this is a rare thing — by back-and-forth exchanges between two languages, which led, in a convergent fashion, to many modifications of the original text, bringing it into closer alignment with its translation, and then the resulting text went once more around the bilingual, bicultural, bicerebral loop, until finally, after a good number of iterations back and forth, things reached a satisfactory equilibrium.

And thus the two versions — the English one that you are now looking at, and its counterpart in French — have gone back and forth many times through the filters of both languages. Often we would find that a high level of clarity emerged as a result of this special dynamic, as translation is nothing if not a merciless revealer of imprecision, vagueness, and lack of logical flow. Translation brings such defects out like a flashlight turned on in a dusty attic. A different metaphor is the sharpening of a knife, because our process of repeated exchanges became, for us, a constant act of sharpening of the ideas we were trying to express. And thus the fact that this book has two originals is not merely an amusing curiosity, but more importantly, it has been a guiding principle keeping us constantly focused on the goal of coherence and lucidity. At least we, the authors, see our book in this way, and we hope that our readers will see it as we do.

We encourage those of our English-speaking readers who are comfortable with French to try tackling a few passages in both versions, because each specific version

takes advantage of ideas, images, and turns of phrase that are deeply rooted in the culture at which it is aimed. This fact made for a particularly enjoyable and stimulating exercise for both authors, in that we were constantly being challenged to come up with an apt analogue for, say, a given idiomatic phrase or a given situation, or perhaps a given speech error, and the quest for optimal examples really kept us on our toes. For anyone who loves languages, then, a parallel perusal of the two texts should provide, in addition to plenty of new ideas (which was of course our primary aim), a special experience of savoring ideas fleshed out in two contrasting ways — in short, a bit of delicious icing on the cake.

“Merci” to So Many!

Two authors, two languages, two lives. While this book was being written, many people were involved with us in many diverse ways, and life predictably followed its unpredictable course. We thus have many heartfelt feelings to express.

At the top of the list are our families, whom we cherish immeasurably. For Doug this means first of all Baofen, and for Emmanuel it means Cécile. They are our muses, amusing and amazing, loving and beloved. “À B., C. – D., E.” says it all, using initials, in French. Next come our children. To Doug, his son Danny and daughter Monica mean everything. Both are rich in humor, verve, idealism, and artistic imagination, inherited largely from their loving mother Carol, who, alas, was torn from us so many years ago. On Emmanuel’s side, there are Michaël, who is protective, intense, and impetuous, and Tom, who is sensitive, social, and solid, and Talia, who is impish, witty and creative; and there is Daniela, their devoted and loving mother. Into our families have recently come, on Doug’s side, Baofen’s son David, and on Emmanuel’s side, Cécile’s son Arthur, who grace our lives with their talents and their gentle natures.

Doug expresses many thanks to his sister Laura Hofstadter, her husband Len Shar, and their two sons, Nathaniel and Jeremy, both filled with intellectual brio. Over the years, their house has been the site of innumerable “jolly evenings” marked by crack croquet competitions, wild word-wizardry, and side-splitting semantic silliness, along with the yummiest of food and the chummiest of chatting. Somewhat symmetrically, Emmanuel warmly thanks David, trusted brother and insightful colleague, David’s wife Véronique, and their daughters Hannah and Gabriela, his radiant nieces. Emmanuel also extends his deepest gratitude to his father and mother, Jean-Pierre and France Sander, for having fostered his growth, from his earliest youth, in the most generous fashion imaginable. And Doug likewise recalls with enormous thanks all the warmth and encouragement of his late parents, Robert and Nancy Hofstadter.

Over the course of these seven-plus years, each of us has experienced the grief of losing several people with whom the bonds ran very deep. Here we wish to honor the treasured memory of Raphaël Sander, Agnès Sander, Maurice Sander, Esther Sidi, Morgan Rogulski, and Lucie Cohen, on Emmanuel’s side, and of Nancy Hofstadter, Helga Keller, Steve Larson, Valentino Braitenberg, and Paolo Bozzi, on Doug’s.

It behooves us now to devote a paragraph or two to Paragraphe, the laboratory that, ever since this book was conceived, has been Emmanuel Sander’s intellectual home at the University of Paris VIII. Its director Imad Saleh leads, with ebullience, generosity, and vigor, a laboratory where human and scientific values exist side by side. Within Paragraphe, the research group CRAC (a French acronym for “Understanding,

Reasoning, and Knowledge Acquisition") is led jointly by Emmanuel Sander and Raphaële Miljkovitch. Emmanuel treasures his intellectual exchanges with Raphaële, and he is delighted to have used some of the plentiful harvest she has made of linguistic oddities issuing from the mouths of her two young sons. CRAC is a cooperative team whose members represent many diverse facets of developmental psychology and get along so well that many strong friendships have come to bloom within it.

And thus a big thank-you to Jean Baratgin (whose specialty is the study of reasoning), Christelle Bosc-Miné (problem-solving), Rémi Brissiaud (educational psychology), Sandra Bruno (conceptual development), Anne-Sophie Deborde (attachment), Corinne Demarcy (problem-solving), Sabine Guéraud (understanding), Caroline Guérini (theory of mind), Frank Jamet (naïve reasoning), Hélène Labat (learning to read), Annamaria Lammel (cultural psychology), Jean-Marc Meunier (knowledge representation), Sandra Nogry (conceptual development), and Carine Royer (learning to read). Emmanuel's doctoral students, current and former, have given much to him through their dedication and the freshness and openness of their thinking. They are a hard act to follow. In particular, we mention Valentine Chaillet, Laurence Dupuch, Sylvie Gamo, Khider Hakem, Bruno Martin, Évelyne Mengue, Lynda Taabane, and Emmanuel Trouche. We also keep in our hearts the memory of Justine Pélouard, who seemed to be headed for a wonderful scientific career, when all at once her life came to an end.

In other teams within Paragraphe, we would like to single out Anne Bationo, Ghislaine Azemard, Claude Baltz, Françoise Decortis, Hakim Hachour, Madjid Ihadjadene, Pierre Quettier, Alexandra Saemmer, Samuel Szoniecky, and Khaldoun Zreik for rich interactions with Emmanuel on numerous occasions. Over the years, some of them have become deeply appreciated friends. Emmanuel would also like to express his gratitude to a set of colleagues outside of Paragraphe, but still in the Psychology Department, for their lively ideas and their personal warmth. Above all, he thanks Marie-Carmen Castillo and Roxane Bordes, and then Aline Frey, Alain Blanchet, Samuel Demarchi, Sophie Frigout, Corinna Kohler, Michèle Montreuil, Tobie Nathan, Michael Pichat, Jean-Luc Picq, and Frédéric Rousseau.

The members of Doug's research group FARG ("Fluid Analogies Research Group"), early on in Ann Arbor but mostly in Bloomington have, over three decades, shed much light on the richness of that elusive mental phenomenon called "analogy-making". We are thinking of Marsha Meredith (who developed the computer model Seek-Whence), Melanie Mitchell (Copycat), Robert French (Tabletop), Gary McGraw (Letter Spirit), John Rehling (Letter Spirit), James Marshall (Metacat), Harry Foundalis (Phaeaco), Francisco Lara-Dammer (George), Abhijit Mahabal (SeqSee), and Eric Nichols (Musicat). Standing on their shoulders and following in their footsteps are Matthew Hurley, Ben Kovitz, William York, and David Bender. Others who have brought ideas and insights to FARG over the years include Daniel Defays (Numbo), Alex Linhares (Capyblanca), David Moser (errors and humor), Donald Byrd, Gray Clossman, Steve Larson, Hamid Ekbia, David Chalmers, Wang Pei, Peter Suber, Yan Yong, Liu Haoming, Christoph Weidemann, Roy Leban, Liane Gabora, and Damien Sullivan.

Beyond FARG, Doug's life has been vitally enriched by so many good friends and sparkling colleagues in so many lands. Let's start with France, where, among the names that come to mind, are François Vannucci, Jacqueline Henry, Serge Haroche,

Daniel Kiechle, Daniel Bougnoux, André Markowicz, Jacques Pitrat, Paul Bourgine, François Récanati, Gilles Cohen, Gilles Esposito-Farèse, Alain Zalmanski, Françoise Strobbe, Jean-Pierre Strobbe, Martine Lemonnier, Anne Bourguignon, Hubert Ceram, Karine Ceram, Liana Gourdjia, Marc Coppey, Geoff Staines, Silvia Busilacchi, Michelle Brûlé, and Denis Malbos. Turning to Italy, where he has always been so warmly received, he is reminded of Benedetto Scimemi, Luisa Scimemi, Giuseppe Trautteur, Pingo Longo, Giovanni Sambin, Alberto Parmeggiani, Francesco Bianchini, Maurizio Matteuzzi, Alex Passi, Sabrina Ardizzoni, Achille Varzi, Oliviero Stock, Enrico Predazzi, Cristina Peroni, Maurizio Codogno, Enrico Laeng, Paola Turina, Patrizio Frosini, Ozalp Babaoglu, Irene Enriques, Pietro Perconti, Andrea Padova, and *la famiglia* Genco.

But we shouldn't omit his friends and colleagues on the North American continent! Doug thus takes great pleasure in saluting (and in a fairly arbitrary order) : Scott Buresh, Greg Huber, Karen Silverstein, Kellie Gutman, Richard Gutman, Caroline Strobbe, Grant Goodrich, Peter Rimbey, Scott Kim, Peter Jones, Steve Jones, Brian Jones, Iranee Zarb, Francis Zarb, David Policansky, Charles Brenner, Inga Karliner, Jon Thaler, Larry Tesler, Colleen Barton, Pentti Kanerva, Eric Hamburg, Michael Goldhaber, Rob Goldstone, Katy Börner, Rich Shiffrin, Jim Sherman, Colin Allen, John Kruschke, Mike Dunn, Breon Mitchell, Dan Friedman, George Springer, Mike Gasser, David Hertz, Willis Barnstone, Sumie Jones, Betsy Stirratt, Marc Hofstadter, Daniel Dennett, John Holland, Bob Axelrod, Dick Nisbett, Ken DeWoskin, Bill Cavnar, Gilles Fauconnier, Mark Turner, Lera Boroditsky, Mark Johnson, Bubal Wolf, Joe Becker, Donald Norman, Bernard Greenberg, Johnny Wink, Jay Curlin, Joseph Sevene, Anton Kuerti, Bill Frucht, Glen Worthey, Marilyn Stone, Jim Falen, Eve Falen, James Plath, Christopher Heinrich, Karen Bentley, Ann Trail, Sue Wunder, Julie Teague, Phoebe Wakhungu, Clark Kimberling, John Rigden, Leon Lederman, Jerry Fisher, Steve Chu, Peter Michelson, Bill Little, Paul Csonka, Sidney Nagel, Don Lichtenberg, Philip Taylor, Simone Brutlag, Doug Brutlag, Sandy Myers, Kristen Motz, and last but not least, Ollie (truly a *golden* retriever). Further afield, scattered hither and yon around the globe, are Francisco Claro, Peter Smith, Robert Boeninger, Cyril Erb, John Ellis, Alexander Rauh, Marina Eskina, Marek Karliner, Hakan Toker, and Michel Moutschen. To all of the above, Doug tips a deeply thankful hat.

Many friends and colleagues have likewise influenced Emmanuel's ideas about thought and have enriched his intellectual, professional, and personal life. Jean-François Richard occupies a special place of honor because of his constant presence, his gift of inspiring others, and his phenomenal creative drive. Emmanuel also wishes to cite the profound influence of colleagues with whom he has had long-time interactions and who have inspired him in many ways. These include Daniel Andler, Nicolas Balacheff, Jean-Marie Barbier, Claude Bastien, Luca Bonatti, Jean-François Bonnefon, Valérie Camos, Roberto Casati, Evelyne Clément, Jacques Crépault, Karine Duvignau, Michel Fayol, Jean-Paul Fischer, Bruno Gaume, Jean-Marc Labat, Jacques Lautrey, Ahn Nguyen Xuan, Jean-François Nicaud, Ira Noveck, Pierre Pastré, Sébastien Poitrenaud, Guy Politzer, Pierre Rabardel, Sandrine Rossi, Gérard Sensevy, Catherine Thevenot, Andrée Tiberghien, André Tricot, Jean-Baptiste Van der Henst, Gérard Vergnaud, Lieven Verschaffel, and Bruno Villette. On a more personal level, Emmanuel wishes to thank so many long-time friends for their loyalty and their indescribably valuable affection: Youri Beltchenko, Florence Deluca Boutrois, Patrick Grinspan, Michaël Jasmin, Audrey Norcia, Franck Lelong, Gaëlle Le Moigne, Philippe

Pétiard, Caroline Sidi, Nadine Zahoui, and Marie-Hélène Zerah.

Our publishers on both sides of the Atlantic have been extremely cordial and helpful to us. On the American side, John Sherer was, for several years, a true believer in this book, and when he left Basic Books, T. J. Kelleher and Lara Heimert kept the enthusiastic fires warmly burning. We would also like to thank Nicole Caputo, Andrea Cardenas, Tisse Takagi, Michele Jacob, Cassie Nelson, and Sue Caulfield for their top-notch contributions to this work in its English-language incarnation. On the French side, Odile Jacob and Bernard Gotlieb welcomed us on board and made us know they would support us strongly. We owe special thanks to Jean-Luc Fidel for his meticulous reading of the manuscript and for his sensitive comments. We also wish to thank Jeanne Pérou, Cécile Andrier-Taverne, and Claudine Roth-Islert for their excellent work on the production and distribution of this book in France.

We expresss deeepest gratitude to David Bender, Steven Willliams, and Jane Stewart Adams for scouring *this* book for typos and non sequiturs, and *mutatis mutandis*, we thank Christellle Bosc-Miné and Karine Duvignau for scouring *that* book. And no, certainly not to be forgotten are Greg Huber, Tom Seeber, and D. Alvin Oyzeau, all of whom tinkered most helpfully with the figures (and sometimes with the facts as well!).

It goes without saying that many of Doug's friends have become friends of Emmanuel's, and vice versa, which of course blurs the borders of all these categories. Such mingling of two worlds has been one of the great side benefits of so many years of work together. Indeed, sometimes the process seemed so long it would never end, and yet here we are, putting the finishing touches on this book. We have learned a great deal about thinking, about writing, and about language from this process, and we hope our readers will take pleasure in, and hopefully inspiration from, our joint creation.

SURFACES AND ESSENCES

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, quasilegal 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 *igg fruders* 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 means, 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 intelligencetest 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

analogymaking) 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

And the same principle holds at more g ll. Thus when one writes or utters a sentence, many of the along without being chosen one by one, since they are all that has been pre-selected. Thus, much as with letters bein words are in a sense constrained by higher-level thoughts. r upwards, we can say that the same holds when one is dev 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 standarized 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 gramma 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).* Æsop’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.

CHAPTER 1

The Evocation of Words

How do Words Pop to Mind?

At every moment we are faced with a new situation. Actually, the truth is much more complicated than that. The truth is that, at every moment, we are simultaneously faced with an indefinite number of overlapping and intermingling situations.

In the airport, we are surrounded by strangers whom we casually observe. Some seem interesting to us, others less so. We see ads everywhere. We think vaguely about the cities whose names come blaring out through loudspeakers, yet at the same time we are absorbed in our private thoughts. We wonder if there's time enough to go get a frozen yogurt, we worry about the health problems of an old friend, we are troubled by the headline we read in someone's newspaper about a terrorist attack in the Middle East, we smile to ourselves at a clever piece of wordplay in an ad on a television screen, we are puzzled as to how the little birds flying around and scavenging food survive in such a weird environment… In short, far from being faced with *one* situation, we are faced with a seething multitude of ill-defined situations, none of which comes with a sharp frame delineating it, either spatially or temporally. Our poor besieged brain is constantly grappling with this unpredictable chaos, always trying to make sense of what surrounds it and swarms into it willy-nilly.

And what does "to make sense of" mean? It means the automatic triggering, or unconscious evocation, of certain familiar categories, which, once retrieved from dormancy, help us to find some order in this chaos. To a large extent, this means the spontaneous coming to mind of all sorts of *words*. Without any effort, one finds oneself thinking, "cute little girl", "funny-looking coot", "same dumb ad as at the airport I was at yesterday", "an Amish family", "sandals", "what's she reading?", "who's whistling?", "where is their nest?", "when are we going to board?", "what an annoying ring tone", "how could I have left my cell-phone charger at home?", "and I did it last time, too", "the air-conditioning is on too high in here", and so on.

All these words! No experience is more familiar to us than this ceaseless barrage of words popping up in our mind extremely efficiently and without ever being invited. But where do these words come from, and what kind of invisible mechanism makes them bubble up? What is going on when one merely thinks silently to oneself, "a

mother and her daughter”?

It All Starts with Single-member Categories

To be able to attach the label “mother” to some entity without thinking about it, one has to be intimately familiar with the concept *mother*, which is denoted by the word. For most of us, this intimacy with the concept goes all the way back to our earliest childhood, to our first encounters with the notion. For one-year-old Tim, the core of the concept is clearly his own mother — a person who is much bigger than he is, who feeds him, comforts him when he cries, sings him lullabies, picks him up, plays with him in the park, and so forth. Once this first mental category bearing the name “Mommy” has a toehold, Tim will be able to see that in the world around him there are similar phenomena, or as we prefer to put it, analogous phenomena.

We take a momentary break here to explain a typographical convention of our book. When speaking about a word, we will put it in quotation marks (“table”), whereas when speaking about a concept, we will use italics (*table)*. This is an important distinction, because whereas a word is a sequence of sounds, a set of printed letters, or a chunk of silent inner language, a concept is an abstract pattern in the brain that stands for some regular, recurrent aspect of the world, and to which any number of different words — for instance, its names in English, French, and so forth, or sometimes no word at all — can be attached. Words and concepts are different things. Although the distinction between them is crucial and often very clear, there will unavoidably be cases in our text where it will be ambiguous and blurry, and in such cases, we’ll make a choice between italics and quotation marks that might seem a bit arbitrary. Another source of ambiguity is the fact that here and there we’ll use italics for emphasis, just as we’ll use quotation marks to suggest a sense of doubt or approximation (which could sometimes be conveyed equally well by the word “so-called”), and of course we will use quotation marks when we are making a quotation. Alas, the world is simply filled with traps, but we hope that the ambiguities are more theoretical than actual. And with that said, we return to our main story.

One day in the park, Tim, aged eighteen months, sees a tot playing in the sandbox and then notices a grown-up near her who is taking care of her. In a flash, Tim makes a little mental leap and thinks to himself more or less the following (although it’s far from being fully verbalized): “That person is taking care of *her* just like Mommy takes care of me.” That key moment marks the birth of the concept *mommy* with a small “m”. The lowercase letter is because there are two members of this new category now (and of course using uppercase and lowercase letters is just *our* way of hinting at what’s going on in Tim’s head, not *his* way). From this point on, it won’t take Tim long to notice yet other instances of this concept.

At the outset, Tim’s concept of *mommy* still floats between singular and plural, and the analogies in his head will be quite concrete, a comparison always being made to the *first* mommy, which is to say, with Mommy (the one with the capital “M”), but as new instances of the concept *mommy* are superimposed and start to blur in his memory, the mental mapping that Tim will automatically carry out, each time he spots a new grown-up in the park, will start to be made not onto Mommy, but onto the nascent and growing concept of *mommy* — that is, onto a generalized, stereotyped, and even slightly abstract situation, centered on a generic grown-up (*i.e.*, stripped of specific

details) and involving a generic child who is near the grown-up and whom the grown-up talks to, smiles at, picks up, comforts, watches out for, and so on.

It's not our goal here to lay out a definitive theory of the growth of the specific concept *mommy*, as our purpose is more general than that. What we are proposing is that the birth of *any* concept takes place more or less as described above. At the outset, there is a concrete situation with concrete components, and thus it is perceived as something unique and cleanly separable from the rest of the world. After a while, though — perhaps a day later, perhaps a year — one runs into another situation that one finds to be similar, and a link is made. From that moment onward, the mental representations of the two situations begin to be connected up, to be blurred together, thus giving rise to a new mental structure that, although it is less specific than either of its two sources (*i.e.*, less detailed), is not fundamentally different from them.

And so the primordial concept *Mommy* and the slightly more sophisticated concept *mommy* act in very similar ways. In particular, both of them are easily mapped onto newly encountered situations "out there", which leads both of them to extending themselves outwards — a snowball effect that will continue all throughout life. It's this idea of concepts extending themselves forever through a long series of spontaneous analogies that we wish to spell out more carefully in the next few sections.

Passing from *Mommy* to *mommy* and then to *mother*

One day, Tim, who sadly has never met his father, is playing in the park, and he runs into a little girl accompanied by a grown-up who is encouraging the girl to play with the other children. He thinks to himself that this grown-up is the *mommy* of the little girl. That is, Tim's mind makes a link between what he's observing and his new concept of *mommy*. This is an act of categorization. Perhaps the new person is not actually the child's mother but the child's father, or perhaps it's her grandmother, or even her older brother or sister, but even so, that doesn't make Tim's mapping of this new person onto the category *mommy* irrational, because his notion of *Mommy/mommy* is wider than ours is (not richer, of course, but less discriminating, due to his lack of experience). This simple analogy Tim has made is flawless; it's just that he hasn't taken into account certain details that an adult would have used. If Sue, his mother, explains to him that this person isn't the little girl's *mommy* but her *daddy*, then Tim may well modify his concept of *mommy*, thereby coming into closer alignment with the people around him.

Gradually, as Tim uses the word "mommy" more and more, his initial image — that of his own mother — will start to recede from view, like a root being grown over ever more as time passes. He will overlay his earliest image with traits of other people whom he assigns to this mental category, and the vivid and unique features of his own Mommy will become harder and harder to find in it. Nonetheless, even when Tim is himself a grown-up, there will remain in his concept of *mommy* some residual traces of his primordial concept *Mommy*.

One day, a friendly woman who's come all the way from her home in Canada turns up and treats Tim very sweetly. He hears the word "mommy" used several times to refer to this grown-up, and so for a while he concludes that maybe he has more than one mommy. For Tim this is conceivable, since he has not yet built up a set of expectations that would rule this possibility out. Sometimes his "second mommy"

takes him to the park and she, too, chats with the other mommies. But after a week or so, Tim's second mommy vanishes, which quite understandably saddens him. The next day, one of the mommies in the park asks Tim, "Did your grandmother go back home?" Tim doesn't answer, because he doesn't yet know the concept of *grandmother*. So she reformulates her question: "Where's your mommy's mom today, Tim?" But this question makes even less sense to Tim. He knows perfectly well that *he's* the one who has the mommy (he even had two of them in the past few days!), and so his mommy (that is, the remaining one) can't have a mommy. After all, it's *children* who have mommies (and sometimes also daddies) whose purpose is to be sweet to them, to watch over them, and to help them, and Tim knows that his mommy *isn't* a child, and so she doesn't have a mommy. That's obvious! The woman doesn't push her strange question, and Tim goes back to his playing.

And time passes. A few months later, Tim starts to realize that grown-ups are sometimes accompanied by other grown-ups that they refer to as their "mother". Suddenly everything starts to be clear… What children have are *mommies*, and what grown-ups have are *mothers*. That makes sense! And into the bargain, there's even an analogical bond between *mommy* and *mother*. Of course Tim isn't aware of having made an analogy — neither this concept nor the word for it will be known to him for another ten or more years! — but he has nonetheless made one. And as is often the case with analogies, this one helps clarify things for Tim but it also misleads him a little.

We now will skip over the details, simply adding that the two concepts of *mommy* and *mother* gradually merge to create a more complex concept at whose core there is the primordial concept of *Mommy*. This doesn't mean that the primordial image of Sue springs to Tim's mind every time that he hears the word "mother" or even the word "mommy", but merely that the invisible roots are structured in that manner.

As any concept grows in generality, it also becomes more discriminating, which means that at some point it's perfectly possible that some early members of the category might be demoted from membership while new members are being welcomed on board. Thus the dad at the park whom Tim had first taken for a *mommy* is stripped of the label, and although Tim's grandmother stays on as a member of the category *mother*, she winds up in a less central zone than the *mommy* zone, which is reserved for the mothers of small children. And of course as time goes by, Tim will come to understand that his grandmother herself was once a member of the category *mommy* (just as his own Mommy was once a member of the category *small child*), but at present all of that is well beyond his grasp.

The Cloud of Concepts of *Mother*

One might think that the concept of *mother* is very precise — perhaps as precise as that of *prime number*. That would imply that to every question of the form "Is X a mother or not?", there would always be a correct, objective, black-and-white answer. But let's consider this for a moment. If a little girl is playing with two dolls, one bigger and one smaller, and she says that the big one is the small one's mother, is this an example of motherhood? Does the large doll belong to the category *mother*? Or contrariwise, could one state without risk of contradiction that she does *not* belong to that category?

And if we read a certain book in which a certain Sue is described as the mother of a certain Tim, then does this Sue, who is never anything but a made-up character in a book, truly belong to the category *mother*? Does it make any difference that Sue was modeled on a real person, and Tim on her son? Is Sue more of a mother than the doll is? What indeed is Sue? If in the book it states that she is 34 years old, that she has light brown hair, that she weighs 120 pounds, that she is five feet five inches tall, and that she's the mom of a small boy, does that mean that Sue has a body and once gave birth? A doll, at least, is a physical object, but what is Sue, when you come down to it? An abstract thought triggered by some words on a page, by some black marks on a white background. Does this thought even deserve the pronoun "she"?

When Tim gets to be six, if someone tells him that Lassie is Spot's mother, he certainly won't protest, but if he were told that the queen bee is the mother of all the bees in her hive, it's less clear what he would say, and in any case some mental effort would be needed before he could absorb this idea. And if he were told that a drop of water that he just watched dividing into two drops is the mother of the two new drops, he would almost surely find this suggestion very surprising. Everyone knows phrases that use the word "mother" in ways that go far beyond the senses that apply to Lassie, the queen bee, or even the splitting drop of water — for instance, "my motherland", "a mother cell", "the mother lode", "Mother Earth", "Greece is the mother of democracy", and "Necessity is the mother of invention". Are these true instances of the concept of *mother*, genuine cases of maternity? What is the proper way to understand such usages of the word?

Some readers may feel inclined to say that these are all "metaphorical mothers", and indeed, such a viewpoint is not without merit, but we have to point out that there is no sharp boundary that separates "true" mothers from those that are metaphorical, for categories in general don't have sharp boundaries; most of the time, metaphorical and literal meanings overlap so greatly that when one tries to draw a clear boundary, one discovers that things only get blurrier and blurrier.

When he turns seven or eight, Tim will start to be able to handle phrases in which the word "mother" is used with greater fluidity than back in nursery school. He might run into the statement "Mary is the mother of the Lord Jesus" in a religious context. This is a mild extension of the usual meaning, since Mary is imagined as a woman whereas the Lord Jesus is imagined as a divine being, magical and omnipotent in some ways, even if also, in some sense, as a baby like all others. At age seven, though, Tim probably won't have much trouble envisioning Mary giving birth to the Lord Jesus.

On the other hand, having given physical birth to a baby is not a prerequisite for attributing motherhood to an entity, since even if no one ever teaches us this explicitly, we all come to know that motherhood pulls together several different properties, such as that of *female biological parent*, that of *female nurturer*, and that of *female protector*, and these properties do not all need to be present simultaneously. For example, the familiar fact of adoption reminds us that giving birth is only one possible route for becoming a mother.

If at age nine, Tim is reading a book on Egypt or on mythology and runs into the sentence "Isis is the mother of Nature", he'll have to extend his prior conceptions of motherhood at least slightly, because this time, Isis is not a human being but a deity who, in Tim's mind, looks much like a woman but in some sense is not one, and who is capable of giving birth to some rather abstract things, such as Nature, yet without anything emerging from her body. And yet Tim will rather easily absorb this new

instance of motherhood, because she looks enough like hundreds of other members of the category *mother* that are already installed in his memory.

Moving right along, Tim will soon handle cases that are even more abstract, such as “Marie Curie is the mother of radioactivity”, “The American revolution is the mother of the French revolution”, “The American revolution is the mother of the Daughters of the American Revolution”, “Judaism is the mother of Christianity”, “Alchemy is the mother of chemistry”, “Censorship is the mother of metaphor” (Jorge Luis Borges), “Leisure is the mother of philosophy” (Thomas Hobbes), and “Death is the mother of beauty” (a quote from Wallace Stevens, and also the title of a detailed study of the role of metaphor in thought by cognitive scientist Mark Turner).

And we can go yet further, to the idea of Nature as the mother of all living creatures (“Mother Nature”), or the idea of *Mother Superior* in a convent, or the idea of *den mother* for a Cub Scout pack, or the idea of a company that has a *mother company* from which it sprang at some earlier point, or the idea of the *mother board* in a computer, and so forth. A mother in a park, a mother in a soap opera, an adoptive mother, a den mother, a mother doll, a mother bee, a mother cell, a mother board, a mother drop of water, a mother deity, a mother company, the mother lode… Given that some mothers, such as Tim’s mommy Sue, are certainly “real mothers”, while others, like the mother board, are just as certainly “metaphorical mothers”, the goal of drawing a sharp, objective boundary between the two distinct subcategories seems as if it might well be within reach. However, as we have shown with our list of blurry examples, such as the person in a novel, the doll mother, and the adoptive mother, that hope is but a beckoning mirage.

On the Categories and Analogies of Children

The story we’ve just told illustrates a central theme of our book — namely, that each category (in this book we use this term synonymously with the term “concept”) is the outcome of a long series of spontaneous analogies, and that the categorization of the elements in a situation takes place exclusively through analogies, however trivial they might seem to an adult. A crucial part of this thesis is that analogies created between a freshly perceived stimulus (such as the mother of a little girl in the park, as seen by Tim) and a relatively new and sparse mental category that has only one single member (such as Tim’s category *Mommy)* are no different from analogies created between a perceived stimulus (once again, take the same woman in the park) and a highly developed mental category to which thousands of analogies have already contributed (think of the very rich category *mother* in the mind of an adult).

This last statement is among the most important in our book, yet on first sight it might seem dubious. Is it really plausible that the very same mechanisms underlie the act whereby a two-year-old spots a Saint Bernard and exclaims “Sheep!” and the act whereby a physicist of great genius discovers a subtle and revelatory mapping between two highly abstract situations? Perhaps it seems implausible at first glance, but we hope to have made it convincing by the end of the book.

In the meantime, to facilitate building a pathway that will get us to this goal, we’ll set up some intermediary bridges. Toward this end, it will be useful for us to take a look at a number of statements made by children, for these statements reveal hidden analogies that underlie their word choices. And so, without further ado, here is a small

sampler of children's sentences, many of which were collected by developmental psychologist Karine Duvignau in her work with parents who were observing their children at home.

Camille, age two, proudly announces: "I undressed the banana!"

She talks about the banana as she would talk about a person or a doll, seeing the peel as an article of clothing that she has removed from it. The banana has thus been "laid bare" (a near neighbor of what Camille said).

Joane, age two, says to her mother: "Come on, Mommy, turn your eyes on!"

Here a little girl speaks to her mother as if she were dealing with an electrical device having an on–off switch.

Lenni, age two, says about a broken toy: "Gotta nurse the truck!"

Here, as in the case of Camille, we see a personification of an inanimate object. The truck is "sick" and so the child wants to help it "get well".

Talia, three years old, says: "Dentists patch people's teeth."

This represents the flip side of the coin, where the child speaks of something alive as if it were an inanimate object (as we just saw Joane do as well).

Jules, three years old, exclaims: "They turned off the rain!"

For Jules, rain is like a television set or a lamp that a person or people can turn on or off with a switch.

Danny, aged five, says to his nursery-school teacher: "I want to eat some water."

In this case, Danny was not speaking his native language but one he was just starting to learn, so he reached out and grabbed the nearest word he knew.

Talia, aged six, says to her mother, "Are you going to go scold the neighbors today?"

The night before, the upstairs neighbors had held a very noisy party, and her mother had told Talia that the next morning she would go knock on their door and complain about their noisemaking. In using the word "scold", Talia was unconsciously revealing her egalitarian value system: any person, whether it's an adult or a child, may sometimes have to be scolded.

Tom, aged eight, asks: "Dad, how long does a guinea pig last?"

While it's true that Tom talks here about his guinea pig in a most materialistic manner, the tenderness with which he treats his pet shows unmistakably that his category *entity of limited duration* is much broader in scope than that of most adults.

At the same age, Tom asks his parents, "How do you cook water?"

This question gets uttered when Tom has generously decided to fix some coffee for his parents one morning, but isn't sure how to start. The distinctions between such kitchen-bound concepts as *to heat up*, *to boil*, *to cook*, and *to fix* are not yet very clear in his mind, but since he announces to anyone who'll listen that he aspires to be a chef in a top-flight restaurant someday, it's to be hoped that this blur won't last too long.

Once again Tom, still eight, says to his uncle, "You know, your cigarette is melting."

This is stated when Tom's uncle is so involved in a conversation that he seems unaware that his cigarette is slowly being consumed in the ashtray. Although Tom knows cigarettes are not for consumption by children, here he links them with certain foods that he knows well, such as ice cream and candy, which can melt.

Tom tips over a wineglass, goes to get a sponge, and chirps, "Here, I'll erase it!"

Part of the tablecloth has just been colored dark, much as paper is colored by pencils or a blackboard is colored by chalk, and so to Tom it makes sense that the sponge will act as an eraser, eliminating all traces of the spilled liquid.

Mica, age twelve, asks his mother, "Mom, could you please roll up your hair?"

He wants to take a snapshot of her and what he means is, "Could you please put your hair up in a bun?", but his thought comes out in a more picturesque way.

Very similar examples are provided by Corentin, who says, “You can stop, Mom, your hair’s all cooked now” (meaning it’s now dry), or Ethan who observes, “I broke the book” (meaning he’s torn it), or Tiffany who declares, “I want to get my nails permed” (meaning she wants a manicure), or little Alexia, who asks, “Mom, can you glue my button back on?” (of course meaning “sew it back on”), and last but not least by Joane, who poses the classic conundrum, “Do buses eat gas?”

Impressive Heights of Abstraction by Children

In each of the cases shown above, one can ask if the child actually was making an error. The key question is, what would constitute an error? If Danny knows the word “drink” but it simply doesn’t come to mind, and if he realizes that “to eat” isn’t really what he means to say, then saying “I want to eat water” would be an error. But if he has the feeling that what he said is perfectly fine, and if he would be surprised to hear the nursery-school teacher correct him, then we’d say that his statement was correct, at least from his own point of view. Most likely Camille, who “undressed the banana”, Ethan, who “broke the book”, and Alexia, who wanted the button to be “glued back on”, had little or no idea of the existence of the verbs “to peel”, “to tear”, and “to sew”. From their viewpoint, what they were saying was correct, because their concepts of *undressing*, *breaking*, and *gluing* were more inclusive than those concepts are in the mind of an adult, and they could thus be applied to situations having a wider range of diversity. For example, Ethan could almost certainly have said, given the proper conditions, “the curtains are broken”, “I broke a loaf of bread”, or “they broke the house”.

On the other hand, it’s very unlikely, even in our society, filled as it is to the brim with technological gadgets, that Joane (“Come on, Mommy, turn your eyes on!”) would be familiar with the verb “to turn on” and yet unfamiliar with the verb “to open”. Likewise, it’s extremely unlikely that Jules (“They turned off the rain!”) knows the verb “to turn off” but is unaware of the verb “to stop”. And so we ask: are these children making errors, or not?

The line between what is and what is not an error is less precise than one might think. What these children are doing is making semantic approximations, stretching their personal concepts in a way that adults would not feel comfortable doing, because the concepts *to turn on*, *to turn off*, *to open*, and *to close* in these children’s minds have not yet reached their adult forms — no more than (to switch from verbs to nouns momentarily) the categories *horse* and *cat* had reached their relatively stable adult stages in the mind of little Abby when, at age three, she saw some greyhounds and called them “horses”, and soon thereafter saw a chihuahua and called it a “cat”. The concepts silently hidden behind these words will continue to develop in the minds of all these children, just as will the category *mother* in Tim’s mind.

The utterances made by such children are not terribly different from the semantic approximations of adults who say “I broke my DVD” instead of “I scratched it”, or “I broke my head getting out of the car” instead of “I banged it”; it’s just that adults’ concepts are a little bit more sophisticated than children’s. And then (sticking with the verb “to break” for a moment) there are many usages that are often labeled “metaphorical”, such as “to break bread”, “to break one’s fast”, “to break one’s silence”, “to break one’s brain”, “to break somebody’s neck”, “to break the ice”, “to

break wind", "to break ground", "to break the news", "to break someone's heart", "to break a habit", "to break away", "to break a code", "to break the law", "to break a world record", and on and on. Such usages are obviously built upon analogical extensions of the verb "to break" that go way beyond anything that a child does who says "the book is broken".

Our tale of children's usage of verbs hasn't "turned off" yet. Let's look at little Joane's use of the verb "Come on!" ("Come on, Mommy, turn your eyes on!"). This usage is undeniably a correct one, and it reveals a deep understanding of the situation that this two-year-old is in. What does "Come on" mean? Firstly, it's a verb that indicates that the speaker wants some change to come about, and it's directed at another person who the speaker feels would be able to make that change happen. Secondly, it's spoken as a kind of urging — stronger than and less polite than "please", almost reaching the intensity of "I insist". Thirdly, although it's an imperative based on the verb "to come", it has nothing to do with physical motion. In fact, "Come on!" is such a frozen expression that one might even argue that it is no longer a genuine verb but more of an interjection, rather like "Hey!" After all, no one would reply to the exhortation "Come on!" by saying "Okay, I'm coming on!" But grammar aside, we are dealing with a subtle word choice made by a toddler. She clearly had put her finger on the situation's essence — namely, she wanted her mother to open her eyes — and that desire led to an eager hope that she could achieve this goal by whining.

To put it in another way, already at the tender age of two, Joane had understood that there is a certain class of situations in life that match and that evoke the label "Come on!" This mental category of *Come on!* situations had gained a solid toehold in her mind. One of the situations belonging to this category was the current one, with her napping mother. To put it succinctly, then, we are saying that *Come on!* situations constitute a mental category that is every bit as real and as important as categories such as *eyes*, *truck*, and *Mommy*, which refer to physical entities in the world. The acquisition of the abstract category of *Come on!* situations by a two-year-old child is a small cognitive miracle and is thus an excellent challenge for anyone who has the goal of deeply understanding human thought.

We might equally well focus on the choice of the verb "gotta" by Lenni ("Gotta nurse the truck!"). This two-year-old boy has understood the essence of situations that are labeled with the pseudo-word "gotta" — namely, something is needed in a hurry, there's no time to lose, and so on. It's very likely that Lenni thinks that "gotta" is just one single word (which is why we didn't write "got to" or "I've got to" or "we've got to", etc.), and this would suggest that he hasn't fully understood that it is a verb, even if in different circumstances he might say, "You don't gotta do it" or "I gotted to do it", and other variants, which are clearly attempts to use it as a verb. So once again we observe a case of a high degree of abstraction carried out by a human being who belongs to the category *toddler*.

Here are a couple of other childish pearls that do not involve verbs. Six-year-old Talia announced, "Dad, we have to get some deodorant for the refrigerator!" (since it reeked of seafood), and her two-and-a-half-year-old cousin Hannah, having just licked all the chocolate off her Eskimo Pie, exclaimed with delight, "Look, now the ice cream is naked!"

Even with nouns that denote the most ordinary and concrete of objects, there remain many subtleties. Lenni said, "Gotta nurse the truck!", but what truck was he speaking of? There was no truck in the apartment; there was just a broken toy. Was

that object really a truck? Well, yes and no. Lenni knows perfectly well that the trucks he sees on highways are hugely bigger than *his* truck, but for him those are distant abstractions; he's never even touched one. By contrast, his little toy truck is a physical object that drives down the invisible highways on the floor of his apartment. In that sense, this toy is, for Lenni, just as central a member of the category *truck* as are those "real" trucks that drive down the "real" highways — indeed, for him, it's probably more central than those are. Ironically, for Lenni, it's *real* trucks that are metaphorical.

Shining Light on the Moon

Earlier we suggested that there is a strong resemblance between the concrete perceptions of a small child and the abstract mental leaps made by a sophisticated physicist. We wish now to illustrate this thesis by means of a concrete example.

In 1610, Galileo Galilei, having just constructed his first telescope, turned it toward the heavens and peered at various celestial bodies. Recall that at that time, the distinction between planets and stars, which today is quite sharp, was still blurry. Certain celestial lights seemed to wander against a backdrop formed by others, but the reason for this movement was not at all clear. Galileo's choice to focus on Jupiter did not mean he knew what it was; it was probably because Jupiter was one of the brightest and thus most inviting objects in the sky to look at.

Galileo's first surprise was that in his telescope Jupiter appeared not as a mere point but as a small circle, which suggested that this "point of light" might well be a solid object with a definite size. Galileo had certainly had the experience of seeing someone with a lantern approaching him. From afar, the lantern seems to be just a dot without size, but then, little by little, the dot acquires a diameter. By analogy with this familiar phenomenon, Galileo could thus imagine that Jupiter, up till then just a dot of light, was in fact a physical object, much like the objects he knew all around him. A second surprise was that against the background of this small white circle he observed some tiny black points, and moreover — a third surprise — these tiny points moved across the circle in a straight line, some taking a few hours, others a few days. Furthermore, whenever one of these points reached the edge of the white circle, it would change color, becoming white against the backdrop of the blackness of space, and would continue moving along the same straight line, then it would slow down, stop, and reverse tracks; when it returned to the edge of the white circle, it would disappear totally, and after a while would reappear on the far side of the white circle.

We won't go into great detail about Galileo's epoch-making discovery; we want, rather, to focus on the way in which the great scientist interpreted what he was observing through his telescope. He decided that Jupiter was a roughly spherical object around which other, smaller objects were rotating perfectly periodically (with periods ranging from about two days to about fifteen days, depending on which dot he was paying attention to). Galileo knew that the Earth was round and that the Moon rotated around it in a periodic fashion, with a period of about thirty days. All these factors added up and suddenly something clicked in his mind. All at once, Galileo was "seeing" a second Earth in the sky, accompanied by several Moons. We put "seeing" in quotation marks to remind readers of the fact that the key moment of "perception" was Galileo's act of interpretation, since the light stimuli arriving at his eyes hadn't changed in the slightest. The analogy between the Moon and a spot of light (or a black

point, depending on where the dot was with respect to the white circle of Jupiter itself) was a stroke of genius — a "vision" of a visionary, so to speak.

Not everyone would have seen what Galileo saw, even if they had been given a telescope, even if they had observed the celestial lights over several weeks, and even if they had focused on Jupiter in particular. The reason is that until that moment, the word "Moon" had been applied to only one object, and the fantasy of "pluralizing" that object was well beyond the imagination of anyone alive at that time (and if someone original had the audacity to think such a thought, that was sufficient to bring about their swift demise: it suffices to recall the case of poor Giordano Bruno, who was burned at the stake in Rome in 1600 for his fantasies about worlds like our own spread throughout space). Moreover, Galileo's daring act of pluralization was the fruit of an analogy that might have seemed laughable to most people — after all, it was an analogy between the entire world, on the one hand (since for most people back then, the terms "Earth" and "world" were synonyms), and on the other hand, an infinitesimal dot of light. This analogy, which might seem far-fetched, nonetheless led to the pluralization of the Earth, since it began by taking Jupiter to be another Earth, and it was rapidly followed by the pluralization of the Moon, which naturally led to the lower-casing of the initial "M". The concept of *moon* had been born, and from that moment on it was possible to imagine one or more moons circling around any celestial body, even around moons themselves.

What Galileo envisioned, in hypothesizing that some small objects in the heavens were rotating around a larger object, was a replica, on an unknown scale, of numerous earthbound situations that were familiar to him, in which one or more objects rotated around a central object. Galileo's stroke of genius was to bank seriously on the daring heliocentric hypothesis of Copernicus and to think to himself that the sky, far from being merely a pretty two-dimensional mural whose purpose was solely to make human life more pleasant, was a genuine *place* that is completely independent of humanity, similar to the places he knew on Earth but much vaster, and as such, capable of housing entities having unknown sizes, and capable of being the site of their movement. In fact, Galileo was completely ignorant of the size of Jupiter and its moons; of course he could imagine a sphere roughly the Earth's size, but doing so would be no more than guesswork, since all he had access to was a set of tiny points. For all he knew, Jupiter might be no larger than the town of Padova, in which he was doing his stargazing, or it might be a hundred times larger than the Earth. Galileo's analogy was an analogy created (or rather, perceived) between something vast and concrete (the Earth and the Moon) and something else that was extremely tiny and immaterial (a circle and some points), but which was nonetheless imaginable as another vast and concrete thing.

Is this profound vision of Galileo's all that different from the vision of the child who sees a very small toy as being a member of the category *truck*, whose other members are so enormous that they are almost inconceivable to the child? One thing is certain — namely, that in both cases, there is a very small object that is imagined as being a very large object, and in both cases, the perceiver uses familiar phenomena in order to understand what is not familiar.

And what about the analogy that *we* are drawing between what Galileo did and what the small child does — is this, too, not just a leap between one scale of sizes and another? Isn't the small cognitive leap by the child, which links a silent, odorless plastic toy truck on the floor with a loud, smoke-belching truck on the highway, simply

a small-scale version of the sophisticated cognitive leap by Galileo, which linked the Earth under his feet with the imagined, distant Jupiter, and which linked our familiar Moon with the imagined, distant Jovian moons? Could it in fact be the case that the tiny child's act of calling an everyday object by its standard name is a close cousin of the genius's act of creating a new concept that revolutionizes human life? For the time being, we won't press the point, but we've planted the seed. To go further will require that we look more closely at the subtlety of the most ordinary categories.

Analogies in the Corridors and Behind the Scenes

Some years back, the senior author of this book went to Italy for a sabbatical year. When he arrived, he had a decent command of Italian but, like everyone in such a situation, he made plenty of mistakes — sometimes subtle, sometimes not — most of which were based on unconsciously drawn analogies to his native culture and language. The research institute where his office was located was a building in which some three hundred people worked — professors, researchers, students, writers, secretaries, administrators, technicians, cafeteria workers, and so forth. During his first few weeks, he met several dozen people, whose names he instantly forgot but whom he would continually bump into in the wide, austere corridors of the building, each time he ventured out of his small office. What to say to all these friendly folks who instantly recognized the newly arrived foreigner, the *professore americano*, and who greeted him warmly (or at least politely) whenever their paths crossed in the hallways? And what to say to the people he saw every day but whom he had never actually met?

His initial assumption, coming from his native culture, was that the proper thing to say to anyone and everyone was "Ciao!", even if it was someone that he wasn't sure he'd ever seen before. This was an innocent assumption based on the American way of saying "Hi", and perhaps it seemed charming to those who received such spontaneous greetings and who were naturally inclined to humor their sender because of his status as a foreign guest, but *il professore* soon noticed that his monosyllabic choice did not coincide with that of the majority of the native speakers of Italian whom he ran into. To be sure, there were a handful of people who said "Ciao" to him, but these were his closest colleagues whom he knew well. Otherwise, though, the people in the halls tended to say either "Salve" or "Buongiorno" to him. It took him a while to figure out the levels of formality that were linked to these two forms of greeting, but in the end he devised a fairly clear rule of thumb for himself to guide him in his hallway greetings. Basically you say "Ciao" to people with whom you are on a first-name basis; you say "Salve" to people you see from time to time and whom you recognize (or think you recognize); and finally, you say "Buongiorno" to people you're not sure you recognize, and also to people whom you would prefer to keep at arm's length.

Once he had formulated this rule of thumb and had gotten it more or less confirmed by native Italian-speaking confidants (who, in truth, had never really thought about it and who were therefore not all that sure of what they were saying), he tried to put his new insight into practice, which meant that every time he ran into someone in a corridor, he had to make an instant triage: "First-name basis? ⇒ *Ciao*. Know them a little bit? ⇒ *Salve*. Not sure who it is? ⇒ *Buongiorno.*" He rapidly discovered that this was a cognitive challenge that was not in the least trivial. Fortunately, in each of these three greeting-categories, there were one or two individuals who served as prototypes,

and using these people as starting points, he began to feel his way in the obscure corridors of acquaintanceship. “Hmm… This fellow who’s approaching me, I know him roughly as well as I know that tall curly-haired administrator” — and zing! — he whipped out a “Salve”. Around several central individuals constituting the nuclei of the three categories, there started to form mental clouds that spread out as time passed. The strategy worked pretty well, and after a few months, *il professore* was handling the challenge fluently as he strode through the corridors of what, at the outset, had been a mysterious maze.

This is a concrete example of how new categories form — in this case, those of *ciao* situations, *salve* situations, and *buongiorno* situations — thanks to the use of analogies at every step of the way. And it also allows us to stress another key point — namely, that behind the scenes of even such a simple-seeming thing as uttering an interjection, there is a complex cognitive process that depends on subtle categories.

Let’s take an example in English that has many points in common with the one just described. On certain occasions one says simply “Thanks” to convey one’s gratitude to someone; on other occasions, one says “Thank you” or “Thank you very much” or “Thanks a lot”; indeed, there is a whole range of thanking possibilities, including such familiar phrases as “Many thanks”, “Thanks ever so much”, “Thanks for everything”, “Thanks a million”, “How can I ever thank you?”, “I can’t thank you enough”, and so on. Obviously there isn’t one exact and perfect choice for each thanking occasion, but on the other hand, certain situations will very naturally evoke just one of these expressions, and some of these expressions would be wildly out of place in certain circumstances. In short, although there isn’t a one-to-one correspondence between situations and expressions, a good choice by a native speaker is far from being a random act, far from being a mere toss of dice. When one is a child, one observes thousands of occasions in which adults use one or another of these phrases without thinking about it for a split second, and pretty soon one starts to do just that oneself. Sometimes adults will smile a little, which conveys the sense that one is probably slightly off-target, while other times one can tell, watching others’ reactions, that one has hit the bull’s-eye. Thus bit by bit, one refines one’s feel for the range of applicability of each of these important and frequent phrases. However, one will probably have no memory whatsoever of the many pathways that collectively led one to one’s current status of grandmaster in the day-to-day arts of greeting and thanking.

And what holds for these seemingly trivial acts holds as well for the labels that one pins on all aspects of reality, including verbs (as we already illustrated in the case of children), adjectives, adverbs, conjunctions (as we shall shortly see), and so forth.

“Office” or “Study”?

If one pays attention to the words that are spontaneously uttered in the most mundane of conversations, one will run into many surprises that reveal something of the processes underlying these choices (if indeed “choice” is the *mot juste* here, since words generally bubble up so automatically that they do not feel like choices one has made). Here we’ll take an example involving Kellie and Dick, two friends who came from Boston to the house of the above-mentioned *professore* a number of years after he had returned to the United States, and who visited for a few days. As it happened, Kellie and Dick both used the term “your office” to designate the standard workplace

of their host, while he himself would always call it “my study”. After he had put up with this cognitive dissonance for a couple of days, it occurred to him to ask them, “How come the two of you always go around talking about my ‘office’ when you both know perfectly well that *I* always call it my ‘study’?”

This question caught the Bostonians by surprise, but they quickly hit upon an answer to it, and it was almost surely *the* answer. They said, “In our Boston house, the place where we work [they had a small public-relations firm that they ran from their house] is on the third floor — our house’s top floor — and we always call it our ‘office’. It’s the place where we have our computer, printer, and photocopy machine, all our filing cabinets, and all the slides and videos we’ve made over the course of the three decades we’ve been doing this. And for you it’s the same thing: your work area is on the second floor — the top floor of your house — and it’s where you have all the stuff that you rely on for your work: your computer, printer, and photocopy machine, your filing cabinets, your books, and so forth. To us the analogy is blatant, crystal-clear. It just jumps out at us, no need to think at all. So to us, your workplace is your *office*, clear as clear can be. That’s the whole story.”

After some reflection on the matter, their host answered, “Aha! I think I see what’s going on here. When I was a kid in California, my father had what he called his ‘study’, which was on the second floor — once again the top floor — of our house. It was the spot where he had lots of papers, books, slide rules, filing cabinets, a mechanical calculator, and so forth. Every day I would see him working there, and it left a vivid impression on me. And also, at the university, on the campus, he had an *office*, where he had many more books, and he often worked down there as well, but the difference between his *study* and his *office* was crystal-clear for me. And today, I too have both a *study* at home and an *office* on the campus here in Indiana. But I would never confuse the two of them. So that’s how I see things.”

And on this note the exchange between friends closed, but there are important lessons that can be drawn from it. First of all, what’s clear is that all parties concerned had depended unconsciously on analogies they had made to very familiar situations. These analogies involved slight “slippages” (third floor instead of second floor; slides and videos instead of books; public-relations work instead of academic work; calculator instead of computer; etc.), but at the same time they respected and preserved a more important essence — namely, both sides of each analogy involved the standard daily workplace, which was separated from the rest of the house and was the storage area for professional material, and so forth. In each case, one sees how the choice of the word to apply to the workplace came from an analogy made to one single familiar situation, rather than what one might have thought *a priori*, which is that assigning an entity to a general category like *office* would depend on the fact that the rich and abstract category *office* had been built up from thousands of different examples encountered over the course of a lifetime. And yet no connection to such a general category took place in this case. Each of the three people, although they all had rich and abstract concepts at their disposal, completely ignored them and instead made a concrete and down-to-earth analogy to a single familiar situation. The numberless prototypical instances of the concept *office*, such as executives’ offices, dentists’ offices, doctors’ offices, lawyers’ offices, and so forth, had nothing to do with what went on in Kellie’s and Dick’s minds. All that mattered was that primordial image of *office* from their own house. This is reminiscent of little Tim’s primordial concept of *Mommy*. Even though the concept *mother* has been enormously enriched for Tim as an

adult, there's no doubt that his own mother has remained over the decades a potential source for analogies; she never got melted down and lost in the abstract concept of *mother*.

As a postscript to this episode, we might add that the Bostonians, during a visit to their friend's home a year later, occasionally used the term "your attic" when referring to his study. Surprised once again by this word choice, he asked them about it, and they explained that they often used the term "attic" in talking about their office in Boston. For them, in this context, the word "attic" had nothing whatsoever to do with a typical messy and dusty attic in a typical house; quite to the contrary, they were thinking of a room at the opposite end of the messiness spectrum — a very clean space in their house, constantly in use on a daily basis. And thus, once again, we see an extremely down-to-earth analogy linking the new place to just one single familiar place rather than to a generic category in which many places are blurred together.

If Kellie and Dick had discovered a truly prototypical attic in their friend's house, full of cobwebs, ancient checkbooks, huge old wooden trunks shipped from abroad, discarded amateurish paintings, and such things, the word "attic" would certainly have sprung to mind because each of them has in their memory not only the concept *our own attic* but also the concept *typical attic*, which allows them to envisage a standard attic, if need be. For example, if Kellie were reading a mystery novel and she came across the sentence, "Trembling, the aged aunt slowly groped her way up the steep and narrow stairway towards the attic to look for the golden statuette, but after three quarters of an hour she hadn't yet come down", the chance is next to zero that this description would evoke in Kellie's mind an image of her Boston house's attic.

This example of the host's *study*, designated first by his visitors as "your office" and later as "your attic", shows how we are guided by unconscious analogies towards labels that seem to pinpoint just what we want to say. It illustrates why there is no boundary line — indeed, no distinction — between categorization and the making of analogies.

The Structure of Categories and of Conceptual Space

The anecdote we've just related shows that a concept (such as those designated by the terms "attic", "truck", "to open", "to melt", "to nurse", "come on!", "ciao", and so forth) can have specific and very distinct instances. Indeed, if we ask you to think of a golfer, you *might* conjure up the image of an anonymous middle-aged lady riding, on a Sunday morning, down some anonymous fairway on a golf cart. But it's more likely that you would conjure up the image of a famous golfer such as Tiger Woods swinging a five-iron, or perhaps you would recall a golf pro you once took lessons from. Instances of the category *golfer* abound, and around each of these specific and concrete instances there is a halo that extends far out. For example, around Tiger Woods, one can imagine seeing him not only making (or missing) a long putt on a tricky green, but also teeing off with great power, hitting out of the rough, and getting out of a sand trap, not to mention appearing in various airport ads and on television, and so forth. Moreover, in this cloud surrounding Tiger Woods, any golf aficionado will surely find a number of Woods' famous predecessors, such as Jack Nicklaus, Arnold Palmer, Sam Snead, Ben Hogan, and others. Anyone familiar with golf will evoke such images without any trouble at all. And so, what does the concept of *golfer* amount to?

It might seem, *a priori*, that asking about the nature of the concept *golfer* is a minuscule question in comparison with the huge question, "What is human thought all about?", but the fact is that it is no smaller. In any case, our modest musings about the concept *golfer* are bringing to the fore the obvious fact that concepts are densely stitched together through relationships of similarity and context. The concept of *golfer* is quite closely linked to that of *minigolfer*, and less closely to such concepts as *tennis player*, *runner*, *bicycle racer*, and so forth. Among these connections, some are very close while others are so distant that they barely exist (for example, there will be practically no relation at all between the notions of *golfer* and *sumo wrestler* in anyone's mind, aside from the fact that both are types of athletes).

The concept of *golfer* is also connected (at various conceptual distances) to a multitude of other concepts, such as *golf course*, *hole*, *fairway*, *tee*, *wood*, *iron*, *putt*, *green*, *par*, *birdie*, *eagle*, *bogey*, *double-bogey*, *hole-in-one*, *hook*, *slice*, *golf cart*, *caddie*, and *tournament*, and also, of course, to a large number of specific people (or more precisely, to the concepts one has formed that represent these people). Despite the large number of golfers of whom any fan has certainly heard the name and has very likely stored it in memory, it's far more probable that a fan will think of Tiger Woods than of some middling player from the 1960's. Thus the distance from the "center" of the concept *golfer* to the concept *Tiger Woods* is quite small, whereas the distance from the center out to the middling player from decades ago is very great, unless, perhaps, the particular player happens to be one's mother or one's uncle or something along those lines.

And thus we come to the idea of a multidimensional space in which concepts exist, somewhat like separate points; however, around each such point there is a halo that accounts for the vague, blurry, and flexible quality of the concept, and this halo becomes ever more tenuous as one moves further out from the core.

The Endless Chunking of Concepts in a Human Mind

We could not make an analogy between one concept and another if those concepts had no internal structure in our mind. The very essence of an analogy is that it maps some mental structure onto another mental structure. We can only understand how a hand is analogous to a foot if we mentally recall the fingers and the toes, for instance, as well as the way the hand is physically attached to an arm and the foot to a leg. These kinds of facts are part of what "hand" means; they are integral to the concept *hand*, and they make it what it is. But how many such facts are there "inside" the concept *hand*? How detailed are the internal structures of our concepts? This is the question to which we wish to turn for a moment.

Consider a rather complex memory in the mind of a certain professor — say, the sabbatical year she spent in Aix-en-Provence. When she recalls that year, of course she doesn't replay its 300-plus days like a movie; rather, she sees just the tiniest part of it, in its barest outlines. It's as if she were looking down at a mountain range from an airplane, but an extensive cloud layer allowed only a handful of the chain's highest peaks to peek through.

If someone asks her about details of the city of Aix, or about some major event that happened during the year, or about the most interesting people she met there, or about the schools that her children went to, and so on, then any of these aspects will become

available upon request, but until that happens, they are all hidden under the “cloud cover”. And if she decides to shift her focus to the school that her children attended that year, then still just a handful of the school’s most salient aspects will come into view. If her focus shifts still further down onto a particular teacher, then a handful of that person’s most salient features come into view — and on it goes. The overarching memory — the sabbatical year in Aix — is never seen in its full glory; rather, just a tiny (but very salient) fraction of it is ever made available. However, pieces of it can be focused in on, and in this way, the large memory can be unpacked into its component pieces, and the same can be done to those pieces, in turn.

All our concepts, from the grandest to the humblest, have the same quality of being largely hidden from view but partially unpackable on request, and the unpacking process is repeatable, several levels down. One might at first think that concepts named by simple words, in contrast to a vast and complex event like a sabbatical year that one is recalling, don’t have much inner structure, but that’s not the case.

Consider the concept of *foot*. When you first think about a foot, you don’t think of cuticles or sweat glands or hairs on it or the fancy swirls making up its five toeprints; you think about toes and an ankle and a large vague central mass, and perhaps a sole and a heel. If you then wish to, you can mentally focus in on a toe and “see” bones and joints inside it, as well as the toenail on top and the toeprint on the bottom. And then, if you wish, you can mentally focus in on the toenail, and so it goes.

So far, our discussion might suggest that concepts are structured according to the physical parts that make them up, with unpacking always moving towards smaller and smaller pieces. Of course, that wouldn’t make sense for concepts of events or other sorts of abstractions, but even when a concept is of something physical, this needn’t be the case. We’ll now give an example that makes this very clear; it is the contemporary concept of a *hub* for a given airline. We chose this concept because the word “hub” is monosyllabic, just three letters long, and sounds very down-to-earth, at the opposite end of the spectrum from fancy technical concepts like *photon*, *ketone*, *entropy*, *mitochondrion*, *autocatalysis*, or *diffeomorphism*. And yet when one looks “inside” this concept, one finds that it, too, is complex — indeed, it has much in common with technical terms. To be concrete, what comes to your mind if we say, “Denver is a hub for Frontier Airlines”? Most people will picture in their mind a map of the United States, with a set of black lines radiating into (or out of) a dot representing Denver, as is shown below.

Perhaps they will also think, “Most of Frontier’s flights go in and out of Denver”, or else “Lots of Frontier planes and gates are found at the Denver airport”. This small set of “highest peaks” (*i.e.*, most salient facts) is pretty much all that one needs in most cases where a hub is being talked about. But in fact it leaves out nearly all of what makes up the concept *hub*, and of which most adults in our culture are perfectly aware. The stipulation “in our culture” is crucial, because hundreds of concepts that we take for granted are not part of other cultures or eras. For instance, imagine trying to explain the concept of *hub* to Johann Sebastian Bach, or to Joan of Arc, or to Archimedes, or to Nebuchadnezzar. These were all remarkable individuals in their respective cultures — but how would you go about trying to get across this “simple” concept to any one of them? It would be a rather long story.

To begin with, the word “hub” is the name of a very concrete, visual concept that we learn when we first ride a bicycle and we see the many spokes radiating out of (or into) the wheel’s center — its hub. Indeed, it’s because of wheels with spokes that airline hubs sport the name “hub”, and the concept of *bicycle wheel* is certainly more “primitive” or “elementary” than is that of *airline hub*, not only being learned far earlier in life but also being far simpler to grasp. Let’s list some other concepts that are more primitive than *hub* and that are likewise prerequisites to it. There’s *airline*, for instance, and *route* and *schedule* and *route map*. And in order to understand the concept of *airline*, you first have to be familiar with the concepts of *airplane* and *company*. And the concept of *route* depends on the concepts of *starting point*, *destination*, *leg*, and *connection*. We won’t go on forever, but let’s not forget that the *raison d’être* of hubs is economic efficiency — the relentless pressure to cut costs and

to reduce the number of different flights — and thus one has to know about the concepts of *trade*, *gain*, *loss*, *profits*, *competition*...

We have only scratched the surface of what goes into the concept *hub*. All of those ingredients are "in there", and they could all, if and when the need arose, be unpacked and revealed. Such unpackings carry one back towards more and more basic, elementary notions — concepts involving motion, vehicles, roundness, acquisition, trading, winning and losing, large and small numbers, and on and on. And note that none of what we have so far spoken of has anything to do with the fact that airports are associated with large cities, or that airports are more than just black dots on maps — indeed, we completely skipped the internal *physical* structure of an airport, with its runways, tarmacs, concourses, gates, jetways, food courts, etc.

The image we've just given of a chain of concepts that depend on other concepts, moving ever downwards in complexity, is reminiscent of the nesting of Russian dolls, and might give the impression that concepts are in fact structured in this boxlike fashion. In truth, however, the phenomenon of concept-building is much subtler and more fluid than that. Concepts are not like nested boxes, with any given concept being rigidly defined in terms of a precise set of previously-acquired concepts, and with concepts always being acquired in a fixed order. Instead, when new concepts are acquired, their arrival often exerts a major impact upon the "more primitive" concepts on which they are based, a bit as if the construction of a house affected the very nature of the bricks with which it was built. Although houses that modify the nature of the bricks of which they are made do not exactly grow on trees, we are all nonetheless familiar with this basic idea, since, for example, children are dependent upon parents in order to exist, but at the same time their existence radically transforms the lives of their parents.

This is true also for concepts. Thus the concept of *hub* depends, without any doubt, on many others, such as *airport*, but at the same time, the concept of *airport* is itself modified by the concept of *hub*. For instance, familiarity with the *hub* idea inevitably brings out the fact that airports are entities that can help airlines to become more streamlined and thus to save money; this notion is certainly not the most obvious fact about airports. Similarly, recalling that airports tend to be transit areas for travelers somewhat reduces the saliency of airports as final destinations. Even if such effects do not cause radical modifications of the concept *airport*, they are undeniably real and demonstrate that the original concept doesn't remain unaffected by the newer one. One can imagine more radical effects of the *hub* concept on that of *airport*, such as novel kinds of architecture aimed at optimizing the design of airports to function as hubs, or the design of new kinds of airport shopping malls specially designed to serve passengers who are making rapid plane-changes and who have only twenty or thirty minutes. And the existence of hubs can change the seemingly obvious correlation between city size and airport size; that is, with hubs, it becomes perfectly conceivable for a relatively small city (such as Charlotte, North Carolina) to have an airport with an enormous volume of air traffic but very few passengers who actually disembark there. We thus see that although there cannot be a "child" concept of *hub* without the prior "parent" concept of *airport*, the child nonetheless changes the identity of the parent.

There are countless examples of this general sort. It happens particularly often in science, where a new idea depends intrinsically on previous ones, but at the same time it casts the old ones in a fresh new light, and often a deeper light. For example, non-Euclidean geometry not only came historically out of Euclidean geometry, but it also

allowed a much deeper understanding of Euclidean geometry to emerge. In physics, much the same could be said for relativistic mechanics and quantum mechanics, both of which are “children” of classical mechanics, and together have yielded a far deeper understanding of it.

The same is true for concepts in everyday life. Thus, the relatively new notions of *surrogate mother*, *adoptive mother*, and *single mother* all come out of the concept of *mother*, as does that of a homosexual couple that adopt a child, and each of these new notions modifies the concept of *mother*, showing how a mother need not give birth to a child, need not raise a child, need not be part of a couple, and may even not be a female. In like manner, the concept of *divorce* depends on that of *marriage*, and yet it also has reverse effects on the nature of *marriage* itself (think, for instance, of the effect of prenuptial contracts, and of the fact that today everyone knows, when going into a marriage, that half of all marriages finish in divorce). The notion of *homosexual marriage* clearly depends on the prior concept of *marriage*, and the intensity of the debate over homosexual marriage is in large part due to the fact that opponents claim that the idea not only extends the concept of *marriage* but in fact does the concept serious harm. The concept of death both depends on and modifies the concept of *life*. The concept of *fast food* both depends on and modifies the concept of *restaurant*. The concept of *credit card* both depends on and modifies the concept of *money*. The concept of *cell phone* depends on and changes that of *phone*. The concept of *traffic accident* depends on and changes that of *car*. The concept of *airplane* depends on and changes that of *distance*. The concept of *recycling* depends on and changes the concept of *garbage*. The concepts of *rape*, *slavery*, *genocide*, *serial killer*, and others not only depend on but change that of *human being*.

Although the repertoire of human concepts is in a sense hierarchical, in that some concepts are prerequisites to other ones, thus implying a rough temporal order in which various concepts generally are acquired, it is nonetheless extremely different in nature from the precise and rigid way that concepts are built up systematically and strictly hierarchically in mathematics or computer science. In the latter contexts, formal definitions are introduced that make each new concept depend explicitly and in an ironclad fashion on a well-defined set of prior concepts. Ordinary concepts have none of this rigidity or precise dependence. True, a person probably needs some familiarity with such concepts as *wheel*, *spoke*, *takeoff*, *landing*, *leg of a trip*, *jetway*, *concourse*, and *transit area*, for instance, before they can acquire the concept of hub, but it’s by no means clear what precise role such concepts play in any specific person’s notion of what a hub is, nor how deeply such concepts have to have been internalized by someone who feels perfectly comfortable with the sentence “Denver is a hub for Frontier Airlines.”

Over the course of our lives, we humans build up concept after concept after concept. This process continues incessantly until we die. This is not the case for many animals, whose conceptual repertoires seem fixed from an early age, and in some cases very limited (think of the conceptual inventory of a frog or a cockroach). And each new concept depends on a number (often very large, as we’ve just seen in the case of hub) of previously existent concepts. But each of those old concepts depended, in its turn, on previous and more primitive concepts. The regress all the way back to babyhood is an extremely long one, indeed. And as we stated earlier, this buildup of concepts over time does not in any way establish a strict and rigid hierarchy. The dependencies are blurry and shaded rather than precise, and there is no strict sense of

"higher" or "lower" in the hierarchy, since, as we've shown, dependencies can be reciprocal. New concepts transform the concepts that existed prior to them and that enabled them to come into being; in this way, newer concepts are incorporated inside their "parents", as well as the reverse. Moreover, this continual process of conceptual chunking goes hand in hand with a continual process of conceptual refinement.

Classical Concepts

Until quite recently, philosophers believed that the physical world was divided into natural categories — that is, that each and every thing, by its very nature, belonged eternally to an objective category. These philosophers focused primarily on categories such as *bird*, *table*, *planet*, and so on, whose members were visible entities. In part as a result of these conjectures from long ago, there remains a tendency, even among most contemporary thinkers, to link the notion of *category* with the idea of classifying physical objects, especially objects that we can perceive visually. The idea that situations of someone being *nursed* back to health, for example, or situations of *hoping* for an outcome or of *changing one's mind*, might constitute categories with just as much legitimacy as *table* or *bird* was far from such philosophers' beliefs, let alone the even further-out idea that words such as "and", "but", "so", "nevertheless", "probably" (and so forth) are the names of important categories. If you find it difficult to imagine that a word like "but", which seems so general and perhaps even bland, denotes a category, don't worry; we will come soon enough to this matter, but for the time being we would like to make some observations on the more classical types of categories, since over the millennia certain ideas have become so entrenched in our culture that it is very difficult to overcome them and to start afresh down new pathways. It will thus be helpful for us to make some elementary observations that will paint a picture of concepts that is markedly different from the classical one.

We might begin by asking what a bird is. According to classical philosophers, whose view went essentially unchallenged in philosophy for centuries, until the studies of philosopher Ludwig Wittgenstein, published in the 1950's, and which also reigned supreme in psychology until the pioneering research of Eleanor Rosch two decades later, the category *bird* should have a precise definition consisting of necessary and sufficient conditions for an entity's membership in the category, such as "possesses two feet", "has skin covered with feathers", "has a beak", "lays eggs". (Obviously one could add further or more refined membership criteria for the category *bird*; these few simply constitute a gesture towards the idea.) The set of membership criteria (the defining properties) is said to be the *intension* of the category, while the set of actual entities that meet the criteria (the *members)* is said to be the *extension* of the category. The notions of intension and extension, borrowed from mathematical logic, are thought of as being just as precise and rigorous as that discipline itself, and the use of these terms reveals the ardent desire to render crystal-clear that which at first seems utterly elusive — namely, the abstract essence of all the highly variegated objects that surround us.

A source of problems, however, is the fact that the words used to express the membership criteria are not any more precise than the concept that one is trying to pin down — in this case, *bird*. What, for instance, is a *foot*? And what does "to possess" mean? What does "covered with" mean? And of course, everyone knows that there are

all sorts of birds that don't have two feet (perhaps because of an injury or a genetic defect) or that are not covered with feathers (ducklings and chicks, for example). And turning things around, we human beings have two feet, but if we hold a spray of feathers in our hand, this "possession" does not suffice to turn us into birds. And the famous *plume de ma tante* — my ancient aunt's quill pen, which she loved to use to make beautiful calligraphy — would that count as a feather? And if so, would possession thereof make my bipedal old aunt a bird?

At times one gets the impression that the actual goal of ancient philosophers was not to classify specific entities from the material world, such as individual birds, whose variety is bewildering, but rather to characterize the relationships that hold between generic, immaterial abstractions, such as the categories *bee*, *bat*, *egg*, *chick*, *ostrich*, *pigeon*, *dragonfly*, *swallow*, *flying fish*, and so forth. If this is one's goal, then the crucial question would be "Which of these *classes* of entities are birds?" It's clear that one has moved far from the specific and concrete, and has replaced it by an intellectual activity where everything is generic and abstract. This rarefied universe of Platonic concepts, since it lacks annoying exceptions like the plucked or the injured bird, not to mention the old aunt who keeps a quill in her drawer, might appear to be as pure, immutable, and objective as the universes of Euclidean geometry or chess, and this could suggest that in this universe there are a vast number of eternal verities lying in wait to be discovered, much like theorems in geometry. But appearances are deceptive. Even if one considers only abstract categories and pays no attention to their annoyingly problematic instances, one still faces enormous obstacles.

Would a chick's lack of plumage make it lose its membership in the category *bird*? That seems unlikely. Or is there a specific instant, for each chick, when it passes over from the category *chick* to that of *bird*? Would that switchover in status take place at the instant when its skin becomes "covered" with feathers? How many feathers does it take for a chick to be "covered" with them? Or what percentage of the skin's area must be covered for it to count as "covered"? And how does one measure the surface area of a chick, if that is needed in order to decide if we are dealing with a bird or not?

The closer one looks, the more such questions one will find, and the more they are going to seem absurd. And we have only scratched the surface of the issues. Consider the generic idea of a bird that has just died. Is it still a bird? And if so, for how long will this entity remain a member of the category *bird*? Will there be a sharp transitional moment at which the category membership no longer obtains? And let's go backwards in time by a few million years. Where is the boundary line between birds and their predecessors (certain flying dinosaurs)? And to push matters in yet another direction, what about questions such as, "Is a plucked chicken still a bird?" The moment one has created the expression "plucked chicken", the question we posed becomes a legitimate question in the hypothetical formal algebra that governs abstract categories. And with this, we have opened a Pandora's box of questions: "Is a robin whose feet have been cut off still a bird?" (since the first noun phrase is the valid name of a category of entities), or "Is a snake onto which one has grafted some feathers and two eagle's feet a bird?", and so on, without any end in sight.

Even without imagining such radical transformations, one can ask whether sandals are *shoes*, whether olives are *fruits*, whether Big Ben is a *clock*, whether a stereo set is a *piece of furniture*, whether a calendar hanging on one's wall is a *book*, whether a wig is an *article of clothing*, and so forth. People turn out to have highly divided opinions on such questions. In an experiment conducted by the psychologist James Hampton,

sinks turned out to be just barely included in the category *kitchen utensils*, while sponges were just barely excluded. Since these close calls are the result of averaging over many subjects in a large experiment, one might imagine that if one were to ask individuals instead, one would find clear-cut and fixed boundaries for each person (even if they would vary from individual to individual). However, even that idea, which runs considerably against the idea of Platonic concepts (which are supposed to be objective, not subjective), turns out to be quite wrong. Many people change their mind if they are asked whether pillows and night-table lamps are articles of furniture and then are asked the same question a few days later. Are these individuals suffering from a pathological state of permanent vacillation, never able to make up their mind about anything? It seems more likely that they are quite ordinary individuals whose categories simply grow blurry toward their edges; if these people were asked about more typical cases, such as whether dogs are *animals*, they would be extremely stable in their judgments about category membership.

Anyone who has taken an interest in the letters of the alphabet will have savored the dazzling richness of a "simple" category like the letter "A", whether capital or small. What geometric shapes belong to the category "A", and what shapes do not? All that one needs to do is take a look at a few handwritten postcards or a collection of typefaces employed in advertising, or for that matter, the figure in the Prologue, in order to see why the boundaries of the twenty-six categories *a*, *b*, *c*, *d*, and so on are impossible to specify exactly. And, to be sure, what holds for the letters of the alphabet holds just as much for other familiar categories, such as *bird*, *bill*, *boss*, *box*, and *brag*.

Summing up, then, the ancient hope of making the categories describing physical objects in the world into precise and rigorous theoretical entities is a vain hope. Such categories are as fleeting and elusive, as blurry and as vague, as clouds. Where are the boundaries of a cloud? How many clouds are there in the sky today? Sometimes, when looking at the sky, one has the impression that such questions have clear and exact answers, and perhaps that's the case on some particular day; however, the next day, the sky will have a radically more complex appearance, and the idea of applying such notions to it as *how many* and *boundary* will simply be a source of smiles.

Concepts Seen in a More Contemporary Fashion

Since the classical view of categories is now generally perceived as a dead end, some contemporary psychologists have tackled the challenge of making the very blurriness and vagueness of categories into a precise science. That is, their goal is to explore those mental nebulas that are our concepts. This has led them to formulating theories of categorization that reject the role of precise membership criteria and instead invoke either the notion of a *prototype* (a generic mental entity found in long-term memory, which summarizes all one's life's experiences with the given category) or else the notion of *the complete set of exemplars* of the given category that one has encountered over one's lifetime. Another influential view involves stored "mental simulators" of experiences one has undergone, which, in response to a fresh stimulus, reactivate certain regions of the brain that were once stimulated by the closest experiences to the current stimulus.

Behind all these efforts lies the appealing idea of non-homogeneous categories — that is, categories having stronger and weaker members — which amounts to

distinguishing between more central and less central members. For example, if one times the responses of experimental subjects when they are asked questions of the form "Is an X a Y?", or if one asks them to write down a list of members of a certain category, or if one gives them a list and asks them to indicate, for each item, its degree of typicality as a member of a specific category, one finds that some very striking trends emerge, and these trends turn out to be stable across all these different ways of testing. Certain members of the category turn out to belong *more* to the given category than others do (recalling how some animals in Orwell's *Animal Farm* were "more equal" than others). For instance, ostriches and penguins turn out to lie close to the outer fringes of the category *bird*, whereas sparrows and pigeons are near its core.

This phenomenon can affect the difficulty one has in understanding a sentence inside a passage that one has been asked to read. Thus, it turns out that the time taken to read and understand a sentence such as "The bird was now just a few yards away" depends on whether, earlier in the passage, there was a reference to an ostrich (an atypical bird) or to a pigeon (a typical bird), in preparatory sentences such as "The ostrich was approaching" or "The pigeon was approaching". The link in memory between *ostrich* and *bird* turns out to be less strong than that between *pigeon* and *bird*, and this tends to impair the understanding of the passage in the first case.

It's important to point out that categorization goes well beyond the intellectual realm of connections among words, which is to say, the names of various categories (such as "sparrow", "ostrich", and "bird"). If, for example, someone were to ask Eleanor "Is a spider an insect?", she might well reply, on the basis of her knowledge from books, "No", and yet if she were to espy a dark blob hanging from the ceiling of her bedroom, it is likely that she would cry out, "Yikes! Get it out of here! I hate insects in my room — they're scary!" If someone were to object to her word choice, Eleanor would say that she knows very well that the "insect" was in fact not an insect but a spider.

Generally speaking, context has a great influence on categorization. The spider in this anecdote was seen as an *insect* in the bedroom, but it would not have been seen as such in the context of a biology test, for instance. And much the same holds in general: a single item in the world belongs to thousands of categories, which can be extremely different from each other, and a good fraction of our mental life consists in placing entities in one category and then in reassigning them to another category. During a basketball game, everyone is aware of the fact that basketballs roll, but it has been experimentally shown that only situations that involve water (such as the loading of a bunch of basketballs on board a ship) evoke the notion that basketballs float.

Context thus changes categorization and can modify how we perceive even the most familiar of items. For example, an object can slip in the blink of an eye from the category *chair* to that of *stool* when a light bulb has just burned out and one needs something to stand on in order to change it. Usually one is unaware of these category shifts because one is mentally immersed in a specific context and such shifts are carried out in a totally unconscious manner. In a given context, just one categorization seems possible to most people. Their lack of awareness of the contextual blinders that they are wearing reinforces the widespread belief in a world in which every object belongs to one and only one Platonic category — its "true" category.

On the other hand, one cannot help but recognize how complex category membership is if one considers the fact that a single entity can easily belong to many diverse categories, such as, for instance:

> 60-kilogram mass, mirror-symmetric object, living entity, biped, mammal, primate, mosquito attractor, arachnophobe, human being, forty-something, book-lover, nature-lover, non-compromiser, non-speaker of Portuguese, romantic, Iowan, blood-type A+, possessor of excellent long-distance vision, insomniac, idealist, vegetarian, member of the bar, mother, mother hen, beloved daughter, sister, big sister, little sister, best friend, sworn enemy, blonde, woman, pedestrian, car driver, cyclist, feminist, wife, twice-married woman, divorcée, neighbor, Dalmatian owner, intermediate-level salsa dancer, breast-cancer survivor, parent of a third-grader, parents' representative…

To be sure, this is but a small excerpt from a much longer list one could draw up, a list having essentially no end, and whose entries would all be terms that anyone and everyone would, without any trouble, recognize as designating various categories.

When Ann had to be hospitalized on an emergency basis and a transfusion was needed, her membership in the category *blood-type A+* dominated all her other category memberships, but in a restaurant she is above all a *vegetarian*, while at work she is a *lawyer*, at home a *mother*, in a PTA meeting a *parents' representative*, and so forth. It may seem useless to point out such obvious facts, but such simple observations carry one well outside the realm of classical categories.

When I Imitate Tweety, Am I a Bird?

Let's come back to the one-word category *bird*, which still has some lessons to teach us. Consider the following candidates for membership in the category:

- a bat;
- an airplane;
- a bronze seagull;
- an eagle in a photograph;
- the shadow of a vulture in the sky;
- Tweety the (cartoon-inhabiting) canary;
- an entire avian species, such as *eagle* or *robin*;
- a chick inside an egg two hours before it hatches;
- a flying dinosaur (or rather, a dinosaur that once flew);
- a pigeon on the screen in a showing of Hitchcock's film *The Birds*;
- the song of a nightingale recorded and played back fifty years after it died;
- a rubber-band-powered wing-flapping plastic object that swoops about in the air.

If you are like the vast majority of humans, you probably felt a keen desire to say "yes" or "no" to each of the candidates in the list above, as if you were taking an exam in school and had to demonstrate the precision of your knowledge, and as if, in each of these cases, there really were a *correct* answer to the question. A sparrow — is it a bird? *Yes!* When you spot a black spot moving unpredictably through the air against a light cloudy background, are you seeing a bird? *Of course!* And when one sees the shadow of a vulture on the ground, is one seeing a bird? *Of course not!* When one hears a loud hooting during the night, is one hearing a bird? *Yes!* And if one hears a recorded hooting (perhaps without being aware that it is recorded)? And what about the case where some person imitates hooting extremely well? And if one dreams about an owl, is there a bird involved? And if one reads a comic book featuring Tweety?

No one ever taught us the boundaries of categories. Our spontaneous sense for their boundaries is an outcome of what we often call "common sense", and no one teaches that in any school. There are no courses on category membership, and even if there were, there would be endless arguments among the students as well as between teachers and students, not to mention the passionate debates that would take place