

The background of the cover is a light blue gradient. It features several water droplets in various stages of impact. Some droplets are captured mid-fall, appearing as small spheres. Others have just hit the surface, creating a central column of water and concentric ripples that spread outwards. The lighting is soft, highlighting the textures of the water and the ripples.

**THE ART OF INSIGHT
IN SCIENCE AND ENGINEERING**

Mastering Complexity

SANJOY MAHAJAN

The Art of Insight in Science and Engineering

Mastering Complexity

Sanjoy Mahajan

The MIT Press
Cambridge, Massachusetts
London, England

© 2014 Sanjoy Mahajan

The Art of Insight in Science and Engineering: Mastering Complexity by Sanjoy Mahajan (author) and MIT Press (publisher) is licensed under the Creative Commons Attribution–Noncommercial–ShareAlike 4.0 International License. A copy of the license is available at creativecommons.org/licenses/by-nc-sa/4.0/



MIT Press books may be purchased at special quantity discounts for business or sales promotional use. For information, please email special_sales@mitpress.mit.edu.

Typeset by the author in 10.5/13.3 Palatino and Computer Modern Sans using ConT_EXt and LuaT_EX.

LIBRARY OF CONGRESS CATALOGING-IN-PUBLICATION DATA

Mahajan, Sanjoy, 1969- author.

The art of insight in science and engineering : mastering complexity / Sanjoy Mahajan.

pages cm

Includes bibliographical references and index.

ISBN 978-0-262-52654-8 (pbk. : alk. paper) 1. Statistical physics.

2. Estimation theory. 3. Hypothesis. 4. Problem solving. I. Title.

QC174.85.E88M34 2014

501'.9-dc23

2014003652

Printed and bound in the United States of America

10 9 8 7 6 5 4 3 2 1

Brief contents

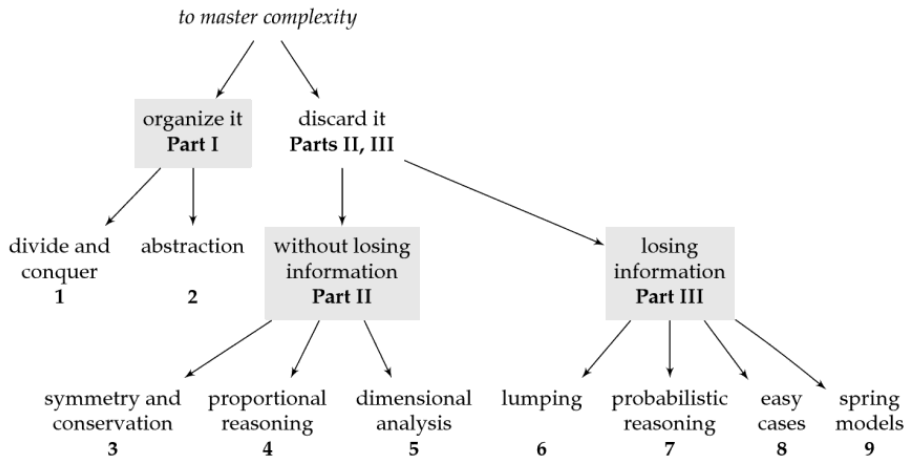
	Preface	xiii
	Values for backs of envelopes	xvii
<i>Part I</i>	<i>Organizing complexity</i>	1
	1 Divide and conquer	3
	2 Abstraction	27
<i>Part II</i>	<i>Discarding complexity without losing information</i>	55
	3 Symmetry and conservation	57
	4 Proportional reasoning	103
	5 Dimensions	137
<i>Part III</i>	<i>Discarding complexity with loss of information</i>	197
	6 Lumping	199
	7 Probabilistic reasoning	235
	8 Easy cases	279
	9 Spring models	317
<hr/>		
	Bon voyage: Long-lasting learning	357
	Bibliography	359
	Index	363

Contents

	Preface	xiii
	Values for backs of envelopes	xvii
<i>Part I</i>	<i>Organizing complexity</i>	1
	1 Divide and conquer	3
	1.1 Warming up	3
	1.2 Rails versus roads	6
	1.3 Tree representations	7
	1.4 Demand-side estimates	10
	1.5 Multiple estimates for the same quantity	16
	1.6 Talking to your gut	17
	1.7 Physical estimates	20
	1.8 Summary and further problems	25
	2 Abstraction	27
	2.1 Energy from burning hydrocarbons	28
	2.2 Coin-flip game	31
	2.3 Purpose of abstraction	34
	2.4 Analogies	36
	2.5 Summary and further problems	53
<i>Part II</i>	<i>Discarding complexity without losing information</i>	55
	3 Symmetry and conservation	57
	3.1 Invariants	57
	3.2 From invariant to symmetry operation	66
	3.3 Physical symmetry	73
	3.4 Box models and conservation	75
	3.5 Drag using conservation of energy	84
	3.6 Lift using conservation of momentum	93
	3.7 Summary and further problems	99

4	Proportional reasoning	103
4.1	Population scaling	103
4.2	Finding scaling exponents	105
4.3	Scaling exponents in fluid mechanics	117
4.4	Scaling exponents in mathematics	123
4.5	Logarithmic scales in two dimensions	126
4.6	Optimizing flight speed	128
4.7	Summary and further problems	135
5	Dimensions	137
5.1	Dimensionless groups	139
5.2	One dimensionless group	147
5.3	More dimensionless groups	152
5.4	Temperature and charge	165
5.5	Atoms, molecules, and materials	175
5.6	Summary and further problems	192
<i>Part III</i>	<i>Discarding complexity with loss of information</i>	197
6	Lumping	199
6.1	Approximate!	199
6.2	Rounding on a logarithmic scale	200
6.3	Typical or characteristic values	203
6.4	Applying lumping to shapes	212
6.5	Quantum mechanics	229
6.6	Summary and further problems	234
7	Probabilistic reasoning	235
7.1	Probability as degree of belief: Bayesian probability	235
7.2	Plausible ranges: Why divide and conquer works	239
7.3	Random walks: Viscosity and heat flow	249
7.4	Transport by random walks	263
7.5	Summary and further problems	276
8	Easy cases	279
8.1	Warming up	279
8.2	Two regimes	281
8.3	Three regimes	291
8.4	Two dimensionless quantities	308
8.5	Summary and further problems	312

9 Spring models	317
9.1 Bond springs	317
9.2 Energy reasoning	321
9.3 Generating sound, light, and gravitational radiation	331
9.4 Effect of radiation: Blue skies and red sunsets	345
9.5 Summary and further problems	353
<hr/>	
Bon voyage: Long-lasting learning	357
Bibliography	359
Index	363



SHARING THIS WORK Like my earlier *Street-Fighting Mathematics* [33], this book is licensed under a Creative Commons Attribution–Noncommercial–Share Alike license. MIT Press and I hope that you will improve and share the work noncommercially, and we would gladly receive corrections and suggestions.

INTER-SPERSED QUESTIONS The most effective teacher is a skilled tutor [2]. A tutor asks many questions, because questioning, wondering, and discussing promote learning. Questions of two types are interspersed through the book. *Questions marked with a ► in the margin*, which a tutor would pose during a tutorial, ask you to develop the next steps of an argument. They are answered in the subsequent text, where you can check your thinking. *Numbered problems*, marked with a shaded background, which a tutor would give you to take home, ask you to practice the tool, to extend an example, to use several tools, and even to resolve an occasional paradox. Merely watching workout videos produces little fitness! So, try many questions of both types.

IMPROVE OUR WORLD Through your effort, mastery will come—and with a broad benefit. As the physicist Edwin Jaynes said of teaching [25]:

[T]he goal should be, not to implant in the students’ mind every fact that the teacher knows now; but rather to implant a way of thinking that enables the student, in the future, to learn in one year what the teacher learned in two years. Only in that way can we continue to advance from one generation to the next.

May the tools in this book help you advance our world beyond the state in which my generation has left it.

Acknowledgments

In addition to the dedication, I would like to thank the following people and organizations for their generosity.

For encouragement, forbearance, and motivation: my family—Juliet Jacobsen, Else Mahajan, and Sabine Mahajan.

For a sweeping review of the manuscript and improvements to every page: Tadashi Tokieda and David MacKay. Any remaining mistakes were contributed by me subsequently!

For advice on the process of writing: Larry Cohen, Hillary Rettig, Mary Carroll Moore, and Kenneth Atchity (author of *A Writer's Time* [1]).

For editorial guidance over many years: Robert Prior.

For valuable suggestions and discussions: Dap Hartmann, Shehu Abdussalam, Matthew Rush, Jason Manuel, Robin Oswald, David Hogg, John Hopfield, Elisabeth Moyer, R. David Middlebrook, Dennis Freeman, Michael Gottlieb, Edwin Taylor, Mark Warner, and many students throughout the years.

For the free software used for typesetting: Hans Hagen, Taco Hoekwater, and the Con \TeX t user community (Con \TeX t and Lua \TeX); Donald Knuth (\TeX); Taco Hoekwater and John Hobby (MetaPost); John Bowman, Andy Hammerlindl, and Tom Prince (Asymptote); Matt Mackall (Mercurial); Richard Stallman (Emacs); and the Debian GNU/Linux project.

For the NB document-annotation system: Sacha Zyto and David Karger.

For being a wonderful place for a graduate student to think, explore, and learn: the California Institute of Technology.

For supporting my work in science and mathematics education: the Whitaker Foundation in Biomedical Engineering; the Hertz Foundation; the Gatsby Charitable Foundation; the Master and Fellows of Corpus Christi College, Cambridge; Olin College of Engineering and its Intellectual Vitality program; and the Office of Digital Learning and the Department of Electrical Engineering and Computer Science at MIT.

Values for backs of envelopes

π	pi	3	
G	Newton's constant	7×10^{-11}	$\text{kg}^{-1} \text{m}^3 \text{s}^{-2}$
c	speed of light	3×10^8	m s^{-1}
hc	h shortcut	200	eV nm
$m_e c^2$	electron rest energy	0.5	MeV
k_B	Boltzmann's constant	10^{-4}	eV K^{-1}
N_A	Avogadro's number	6×10^{23}	mol^{-1}
R	universal gas constant $k_B N_A$	8	$\text{J mol}^{-1} \text{K}^{-1}$
e	electron charge	1.6×10^{-19}	C
$e^2/4\pi\epsilon_0$	electrostatic combination	2.3×10^{-28}	$\text{kg m}^3 \text{s}^{-2}$
$(e^2/4\pi\epsilon_0)/hc$	fine-structure constant α	0.7×10^{-2}	
σ	Stefan-Boltzmann constant	6×10^{-8}	$\text{W m}^{-2} \text{K}^{-4}$
M_{Sun}	solar mass	2×10^{30}	kg
m_{Earth}	Earth's mass	6×10^{24}	kg
R_{Earth}	Earth's radius	6×10^6	m
AU	Earth-Sun distance	1.5×10^{11}	m
$\theta_{\text{Moon or Sun}}$	angular diameter of Moon or Sun	10^{-2}	rad
day	length of a day	10^5	s
year	length of a year	$\pi \times 10^7$	s
t_0	age of the universe	1.4×10^{10}	yr
F	solar constant	1.3	kW m^{-2}
p_0	atmospheric pressure at sea level	10^5	Pa
ρ_{air}	air density	1	kg m^{-3}
ρ_{rock}	rock density	2.5	g cm^{-3}
$L_{\text{vap}}^{\text{water}}$	heat of vaporization of water	2	MJ kg^{-1}
γ_{water}	surface tension of water	7×10^{-2}	N m^{-1}
P_{basal}	human basal metabolic rate	100	W
a_0	Bohr radius	0.5	Å
a	typical interatomic spacing	3	Å
E_{bond}	typical bond energy	4	eV
\mathcal{E}_{fat}	combustion energy density	9	kcal g^{-1}
ν_{air}	kinematic viscosity of air	1.5×10^{-5}	$\text{m}^2 \text{s}^{-1}$
ν_{water}	kinematic viscosity of water	10^{-6}	$\text{m}^2 \text{s}^{-1}$
K_{air}	thermal conductivity of air	2×10^{-2}	$\text{W m}^{-1} \text{K}^{-1}$
K	... of nonmetallic solids/liquids	2	$\text{W m}^{-1} \text{K}^{-1}$
K_{metal}	... of metals	2×10^2	$\text{W m}^{-1} \text{K}^{-1}$
$c_{\text{p}}^{\text{air}}$	specific heat of air	1	$\text{J g}^{-1} \text{K}^{-1}$
c_{p}	... of solids/liquids	25	$\text{J mol}^{-1} \text{K}^{-1}$

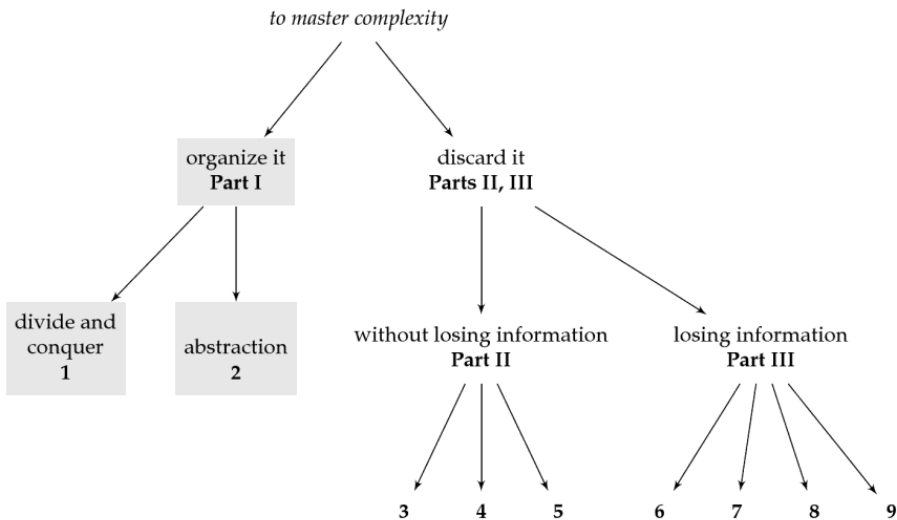
Part I

Organizing complexity

We cannot find much insight staring at a mess. We need to organize it. As an everyday example, when I look at my kitchen after a dinner party, I feel overwhelmed. It's late, I'm tired, and I dread that I will not get enough sleep. If I clean up in that scattered state of mind, I pick up a spoon here and a pot there, making little progress. However, when I remember that a large problem can be broken into smaller ones, calm and efficiency return. I begin at one corner of the kitchen, clear its mess, and move to neighboring areas until the project is done. I divide and conquer (Chapter 1).

Once the dishes are clean, I resist the temptation to dump them into one big box. I separate pots from the silverware and, within the silverware, the forks from the spoons. These groupings, or abstractions (Chapter 2), make the kitchen easy to understand and use.

In problem solving, we organize complexity by using divide-and-conquer reasoning and by making abstractions. In Part I, you'll learn how.



This range is wide, spanning a factor of 100. In contrast, the dollar bill's width probably lies between 10 and 20 centimeters—a range of only a factor of 2. The volume range is wider than the width range because we have no equivalent of a ruler for volume; thus, volumes are less familiar than lengths. Fortunately, the volume of the dollar bill is the product of lengths.

$$\text{volume} = \text{width} \times \text{height} \times \text{thickness}. \quad (1.1)$$

The harder volume estimate becomes three easier length estimates—the benefit of divide-and-conquer reasoning.

The width looks like 6 inches, which is roughly 15 centimeters. The height looks like 2 or 3 inches, which is roughly 6 centimeters. But before estimating the thickness, let's talk about unit systems.

6 cm	\$1 bill
	15 cm

► *Is it better to use metric or US customary units (such as inches, feet, and miles)?*

Your estimates will be more accurate if you use the units most familiar to you. Raised in the United States, I judge lengths more accurately in inches, feet, and miles than in centimeters, meters, or kilometers. However, for calculations requiring multiplication or division—most calculations—I convert the customary units to metric (and often convert back to customary units at the end). But you may be fortunate enough to think in metric. Then you can estimate and calculate in a single unit system.

The third piece of the divide-and-conquer estimate, the thickness, is difficult to judge. A dollar bill is thin—paper thin.

► *But how thin is “paper thin”?*

This thickness is too small to grasp and judge easily. However, a stack of several hundred bills would be graspable. Not having that much cash lying around, I'll use paper. A ream of paper, which has 500 sheets, is roughly 5 centimeters thick. Thus, one sheet of paper is roughly 0.01 centimeters thick. With this estimate for the thickness, the volume is approximately 1 cubic centimeter:

$$\text{volume} \approx \underbrace{15 \text{ cm}}_{\text{width}} \times \underbrace{6 \text{ cm}}_{\text{height}} \times \underbrace{0.01 \text{ cm}}_{\text{thickness}} \approx 1 \text{ cm}^3. \quad (1.2)$$

Although a more accurate calculation could adjust for the fiber composition of a dollar bill compared to ordinary paper and might consider the roughness of the paper, these details obscure the main result: A dollar bill is 1 cubic centimeter pounded paper thin.

To check this estimate, I folded a dollar bill until my finger strength gave out, getting a roughly cubical packet with sides of approximately 1 centimeter—making a volume of approximately 1 cubic centimeter!

In the preceding analysis, you may have noticed the $=$ and \approx symbols and their slightly different use. Throughout this book, our goal is insight over accuracy. So we'll use several kinds of equality symbols to describe the accuracy of a relation and what it omits. Here is a table of the equality symbols, in descending order of completeness and often increasing order of usefulness.

\equiv	equality by definition	read as “is defined to be”
$=$	equality	“is equal to”
\approx	equality except perhaps for a purely numerical factor near 1	“is approximately equal to”
\sim	equality except perhaps for a purely numerical factor	“is roughly equal to” or “is comparable to”
\propto	equality except perhaps for a factor that may have dimensions	“is proportional to”

As examples of the kinds of equality, for the circle below, $A = \pi r^2$, and $A \approx 4r^2$, and $A \sim r^2$. For the cylinder, $V \sim hr^2$ —which implies $V \propto r^2$ and $V \propto h$. In the $V \propto h$ form, the factor hidden in the \propto symbol has dimensions of length squared.

Copyrighted image

Copyrighted image

$$A \begin{cases} = \pi r^2 \\ \approx 4r^2 \\ \sim r^2 \end{cases}$$

$$V \propto \begin{cases} r^2 \\ h \end{cases}$$

Problem 1.1 Weight of a box of books

How heavy is a small moving-box filled with books?

Problem 1.2 Mass of air in your bedroom

Estimate the mass of air in your bedroom.

Problem 1.3 Suitcase of bills

In the movies, and perhaps in reality, cocaine and elections are bought with a suitcase of \$100 bills. Estimate the dollar value in such a suitcase.

Problem 1.4 Gold or bills?

As a bank robber sitting in the vault planning your getaway, do you fill your suitcase with gold bars or \$100 bills? Assume first that how much you can carry is a fixed weight. Then redo your analysis assuming that how much you can carry is a fixed volume.

1.2 Rails versus roads

We are now warmed up and ready to use divide-and-conquer reasoning for more substantial estimates. Our next estimate, concerning traffic, comes to mind whenever I drive the congested roads to JFK Airport in New York City. The route goes on the Van Wyck Expressway, which was planned by Robert Moses. As Moses's biographer Robert Caro describes [6, pp. 904ff], when Moses was in charge of building the expressway, the traffic planners recommended that, in order to handle the expected large volume of traffic, the road include a train line to the then-new airport. Alternatively, if building the train track would be too expensive, they recommended that the city, when acquiring the land for the road, still take an extra 50 feet of width and reserve it as a median strip for a train line one day. Moses also rejected the cheaper proposal. Alas, only weeks after its opening, not long after World War Two, the rail-free highway had reached peak capacity.

Let's use our divide-and-conquer tool to compare, for rush-hour commuting, the carrying capacities of rail and road. The capacity is the rate at which passengers are transported; it is passengers per time. First we'll estimate the capacity of one lane of highway. We can use the 2-second-following rule taught in many driving courses. You are taught to leave 2 seconds of travel time between you and the car in front. When drivers follow this rule, a single lane of highway carries one car every 2 seconds. To find the carrying capacity, we also need the occupancy of each car. Even at rush hour, at least in the United States, each car carries roughly one person. (Taxis often have two people including the driver, but only one person is being transported to the destination.) Thus, the capacity is one person every 2 seconds. As an hourly rate, the capacity is 1800 people per hour:

$$\frac{1 \text{ person}}{2 \text{ s}} \times \frac{3600 \text{ s}}{1 \text{ hr}} = \frac{1800 \text{ people}}{\text{hr}}. \quad (1.3)$$

The diagonal strike-through lines help us to spot which units cancel and to check that we end up with just the units that we want (people per hour).

This rate, 1800 people per hour, is approximate, because the 2-second following rule is not a law of nature. The average gap might be 4 seconds late at night, 1 second during the day, and may vary from day to day or from highway to highway. But a 2-second gap is a reasonable compromise estimate. Replacing the complex distribution of following times with one time is an application of lumping—the tool discussed in Chapter 6. Organizing complexity almost always reduces detail. If we studied all highways at all times of day, the data, were we so unfortunate as to obtain them, would bury any insight.

- *How does the capacity of a single lane of highway compare with the capacity of a train line?*

For the other half of the comparison, we'll estimate the rush-hour capacity of a train line in an advanced train system, say the French or German system. As when we estimated the volume of a dollar bill (Section 1.1), we divide the estimate into manageable pieces: how often a train runs on the track, how many cars are in each train, and how many passengers are in each car. Here are my armchair estimates for these quantities, kept slightly conservative to avoid overestimating the train-line's capacity. A single train car, when full at rush hour, may carry 150 people. A rush-hour train may consist of 20 cars. And, on a busy train route, a train may run every 10 minutes or six times per hour. Therefore, the train line's capacity is 18 000 people per hour:

$$\frac{150 \text{ people}}{\text{car}} \times \frac{20 \text{ cars}}{\text{train}} \times \frac{6 \text{ trains}}{\text{hr}} = \frac{18\,000 \text{ people}}{\text{hr}}. \quad (1.4)$$


This capacity is ten times the capacity of a single fast-flowing highway lane. And this estimate is probably on the low side; Robert Caro [6, p. 901] gives an estimate of 40 000 to 50 000 people per hour. Using our lower rate, one train track in each direction could replace two highways even if each highway had five lanes in each direction.

1.3 Tree representations


Our estimates for the volume of a dollar bill (Section 1.1) and for the rail and highway capacities (Section 1.2) used the same method: dividing hard problems into smaller ones. However, the structure of the analysis is buried within the sentences, paragraphs, and pages. The sequential presentation hides the structure. Because the structure is hierarchical—big problems

split, or branch, into smaller problems—its most compact representation is a tree. A tree representation shows us the analysis in one glance.


Here is the tree representation for the capacity of a train line. Unlike the biological variety, our trees stand on their head. Their roots, the goals, sit at the top of the tree. Their leaves, the small problems into which we have subdivided the goal, sit at the bottom. The orientation matches the way that we divide and conquer, filling the page downward as we subdivide.




In making this first tree, we haven't estimated the quantities themselves. We have only identified the quantities. The question marks remind us of our next step: to include estimates for the three leaves. These estimates were 150 people per car, 20 cars per train, and 6 trains per hour (giving the tree in the margin).



Then we multiplied the leaf values to propagate the estimates upward from the leaves toward the root. The result was 18 000 people per hour. The completed tree shows us the entire estimate in one glance.



This train-capacity tree had the simplest possible structure with only two layers (the root layer and, as the second layer, the three leaves). The next level of complexity is a three-layer tree, which will represent our estimate for the volume of a dollar bill. It started as a two-layer tree with three leaves.




Then it grew, because, unlike the width and height, the thickness was difficult to estimate just by looking at a dollar bill. Therefore, we divided that leaf into two easier leaves.

Here, subdivide the demand—the consumption. We consume oil in so many ways; estimating the consumption in each pathway would take a long time without producing much insight. Instead, let's estimate the largest consumption—likely to be cars—then adjust for other uses and for overall consumption versus imports.

$$\text{imports} = \overline{\text{car usage}} \times \frac{\overline{\text{all usage}}}{\overline{\text{car usage}}} \times \frac{\overline{\text{imports}}}{\overline{\text{all usage}}}. \quad (1.6)$$

Here is the corresponding tree. The first factor, the most difficult of the three to estimate, will require us to sprout branches and make a subtree. The second and third factors might be possible to estimate without subdividing. Now we must decide how to continue.



- *Should we keep subdividing until we've built the entire tree and only then estimate the leaves, or should we try to estimate these leaves and then subdivide what we cannot estimate?*

It depends on one's own psychology. I feel anxious in the uncharted waters of a new estimate. Sprouting new branches before making any leaf estimates increases my anxiety. The tree might never stop sprouting branches and leaves, and I'll never estimate them all. Thus, I prefer to harvest my progress right away by estimating the leaves before sprouting new branches. You should experiment to learn your psychology. You are your best problem-solving tool, and it is helpful to know your tools.

Because of my psychology, I'll first estimate a leaf quantity:

$$\frac{\overline{\text{all usage}}}{\overline{\text{car usage}}}. \quad (1.7)$$

But don't do this estimate directly. It is more intuitive—that is, easier for our gut—to estimate first the ratio of car usage to other (noncar) usage. The ability to make such comparisons between disjoint sets, at least for physical objects, is hard wired in our brains and independent of the ability to count. Not least, it is not limited to humans. The female lions studied by Karen McComb and her colleagues [35] would judge the relative size of their troop and a group of lions intruding on their territory. The females would approach the intruders only when they outnumbered the intruders by a large-enough ratio, roughly a factor of 2.

Other uses for oil include noncar modes of transport (trucks, trains, and planes), heating and cooling, and hydrocarbon-rich products such as fertilizer, plastics, and pesticides. In judging the relative importance of other uses compared to car usage, two arguments compete: (1) Other uses are so many and so significant, so they are much more important than car usage; and (2) cars are so ubiquitous and such an inefficient mode of transport, so car usage is much larger than other uses. To my gut, both arguments feel comparably plausible. My gut is telling me that the two categories have comparable usages:

$$\frac{\text{other usage}}{\text{car usage}} \approx 1. \quad (1.8)$$

Based on this estimate, all usage (the sum of car and other usage) is roughly double the car usage:

$$\frac{\text{all usage}}{\text{car usage}} \approx 2. \quad (1.9)$$

This estimate is the first leaf. It implicitly assumes that the gasoline fraction in a barrel of oil is high enough to feed the cars. Fortunately, if this assumption were wrong, we would get warning. For if the fraction were too low, we would build our transportation infrastructure around other means of transport—such as trains powered by electricity generated by burning the nongasoline fraction in oil barrels. In this probably less-polluted world, we would estimate how much oil was used by trains.

Returning to our actual world, let's estimate the second leaf:

$$\frac{\text{imports}}{\text{all usage}}. \quad (1.10)$$

This adjustment factor accounts for the fact that only a portion of the oil consumed is imported.

► *What does your gut tell you for this fraction?*


Again, don't estimate this fraction directly. Instead, to make a comparison between disjoint sets, first compare (net) imports with domestic production. In estimating this ratio, two arguments compete. On the one hand, the US media report extensively on oil production in other countries, which suggests that oil imports are large. On the other hand, there is also extensive coverage of US production and frequent comparison with countries such as Japan that have almost no domestic oil. My resulting gut feeling is that the

categories are comparable and therefore that imports are roughly one-half of all usage:

$$\frac{\text{imports}}{\text{domestic production}} \approx 1 \quad \text{so} \quad \frac{\text{imports}}{\text{all usage}} \approx \frac{1}{2}. \quad (1.11)$$


This leaf, as well as the other adjustment factor, are dimensionless numbers. Such numbers, the main topic of Chapter 5, have special value. Our perceptual system is skilled at estimating dimensionless ratios. Therefore, a leaf node that is a dimensionless ratio probably does not need to be subdivided.

The tree now has three leaves. Having plausible estimates for two of them should give us courage to subdivide the remaining leaf, the total car usage, into easier estimates. That leaf will sprout its own branches and become an internal node.



► How should we subdivide the car usage?


A reasonable subdivision is into the number of cars N_{cars} and the per-car usage. Both quantities are easier to estimate than the root. The number of cars is related to the US population—a familiar number if you live in the United States. The per-car usage is easier to estimate than is the total usage of all US cars. Our gut can more accurately judge human-scale quantities, such as the per-car usage, than it can judge vast numbers like the total usage of all US cars.



For the same reason, let's not estimate the number of cars directly. Instead, subdivide this leaf into two leaves:

1. the number of people, and
2. the number of cars per person.

The first leaf is familiar, at least to residents of the United States: $N_{\text{people}} \approx 3 \times 10^8$.



The second leaf, cars per person, is a human-sized quantity. In the United States, car ownership is widespread. Many adults own more than one car, and a cynic would say that even babies seem to own cars. Therefore, a rough and simple estimate might be one car per person—far easier to picture than the total number of cars! Then $N_{\text{cars}} \approx 3 \times 10^8$.

The per-car usage can be subdivided into three easier factors (leaves). Here are my estimates.

1. *How many miles per car year?* Used cars with 10 000 miles per year are considered low use but are not rare. Thus, for a typical year of driving, let's take a slightly longer distance: say, 20 000 miles or 30 000 kilometers.

2. *How many miles per gallon?* A typical car fuel efficiency is 30 miles per US gallon. In metric units, it is about 100 kilometers per 8 liters.

3. *How many gallons per barrel?* You might have seen barrels of asphalt along the side of the highway during road construction. Following our free-association tradition of equating the thickness of a sheet of paper and of a dollar bill, perhaps barrels of oil are like barrels of asphalt.

Their volume can be computed by divide-and-conquer reasoning. Just approximate the cylinder as a rectangular prism, estimate its three dimensions, and multiply:

$$\text{volume} \sim \frac{1 \text{ m}}{\text{height}} \times \frac{0.5 \text{ m}}{\text{width}} \times \frac{0.5 \text{ m}}{\text{depth}} = 0.25 \text{ m}^3. \quad (1.12)$$

Copyrighted image

A cubic meter is 1000 liters or, using the conversion of 4 US gallons per liter, roughly 250 gallons. Therefore, 0.25 cubic meters is roughly 60 gallons. (The official volume of a barrel of oil is not too different at 42 gallons.)


Multiplying these estimates, and not forgetting the effect of the two -1 exponents, we get approximately 10 barrels per car per year (also written as barrels per car year):

$$\frac{2 \times 10^4 \text{ miles}}{\text{car year}} \times \frac{1 \text{ gallon}}{30 \text{ miles}} \times \frac{1 \text{ barrel}}{60 \text{ gallons}} \approx \frac{10 \text{ barrels}}{\text{car year}}. \quad (1.13)$$

In doing this calculation, first evaluate the units. The gallons and miles cancel, leaving barrels per year. Then evaluate the numbers. The 30×60 in the denominator is roughly 2000. The 2×10^4 from the numerator divided by the 2000 from the denominator produces the 10.

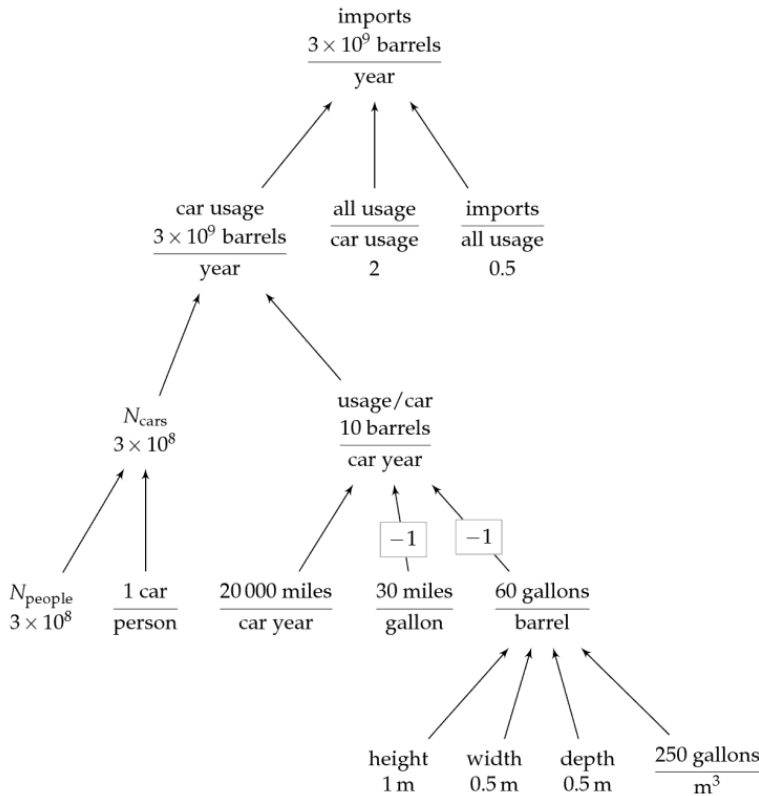

This estimate is a subtree in the tree representing total car usage. The car usage then becomes 3 billion barrels per year:

$$3 \times 10^8 \text{ cars} \times \frac{10 \text{ barrels}}{\text{car year}} = \frac{3 \times 10^9 \text{ barrels}}{\text{year}}. \quad (1.14)$$



This estimate is itself a subtree in the tree representing oil imports. Because the two adjustment factors contribute a factor of 2×0.5 , which is just 1, the oil imports are also 3 billion barrels per year.

Here is the full tree, which includes the subtree for the total car usage of oil:



Problem 1.6 Using metric units


As practice with metric units (if you grew up in a nonmetric land) or to make the results more familiar (if you grew up in a metric land), redo the calculation using the metric values for the volume of a barrel, the distance a car is driven per year, and the fuel consumption of a typical car.

► How close is our estimate to official values?

divide-and-conquer tree has two leaves. (In Section 6.3.1, you'll see a qualitatively different method, where the two leaves will be the number of US states and the population of a typical state.)

The area is the width times the height, so the area leaf itself splits into two leaves. Estimating the width and height requires only a short dialogue with the gut, at least if you live in the United States. Its width is a 6-hour plane flight at 500 miles per hour, so about 3000 miles; and the height is, as a rough estimate, two-thirds of the width, or 2000 miles. Therefore, the area is 6 million square miles:

$$3000 \text{ miles} \times 2000 \text{ miles} = 6 \times 10^6 \text{ miles}^2. \quad (1.16)$$



In metric units, it is about 16 million square kilometers.

Estimating the population density requires talking to your gut. If you are like me you have little conscious knowledge of the population density. Your gut might know, but you cannot ask it directly. The gut is connected to the right brain, which doesn't have language. Although the right brain knows a lot about the world, it cannot answer with a value, only with a feeling. To draw on its knowledge, ask it indirectly. Pick a particular population density—say, 100 people per square mile—and ask the gut for its opinion: “O, my intuitive, insightful, introverted right brain: What do you think of 100 people per square mile for the population density?” A response, a gut feeling, will come back. Keep lowering the candidate value until the gut feeling becomes, “No, that value feels way too low.”

Here is the dialogue between my left brain (LB) and right brain (RB).

LB: What do you think of 100 people per square mile?

RB: That feels okay based on my experience growing up in the United States.

LB: I can probably support that feeling quantitatively. A square mile with 100 people means each person occupies a square whose side is 1/10th of a mile or 160 meters. Expressed in this form, does the population density feel okay?

RB: Yes, the large open spaces in the western states probably compensate for the denser regions near the coasts.

LB: Now I will lower the estimate by factors of 3 or 10 until you object strongly that the estimate feels too low. [A factor of 3 is roughly one-half of a factor of 10, because $3 \times 3 \approx 10$. A factor of 3 is the next-smallest factor by which to move when a factor of 10 is too large a jump.] In that vein, what about an average population density of 10 people per square mile?

RB: I feel uneasy. The estimate feels a bit low.

LB: I understand where you are coming from. That value may moderately overestimate the population density of farmland, but it probably greatly underestimates the population density in the cities. Because you are uneasy, let's move more slowly until you object strongly. How about 3 people per square mile?

RB: If the true value were lower than that, I'd feel fairly surprised.

LB: So, for the low end, I'll stop at 3 people per square mile. Now let's navigate to the upper end. You said that 100 people per square mile felt plausible. How do you feel about 300 people per square mile?

RB: I feel quite uneasy. That estimate feels quite high.

LB: I hear you. Your response reminds me that New Jersey and the Netherlands, both very densely populated, are at 1000 people per square mile, although I couldn't swear to this value. I cannot imagine packing the whole United States to a density comparable to New Jersey's. Therefore, let's stop here: Our upper endpoint is 300 people per square mile.

► How do you make your best guess based on these two endpoints?

A plausible guess is to use their arithmetic mean, which is roughly 150 people per square mile. However, the right method is the geometric mean:

$$\text{best guess} = \sqrt{\text{lower endpoint} \times \text{upper endpoint}}. \quad (1.17)$$

The geometric mean is the midpoint of the lower and upper bounds—but on a ratio or logarithmic scale, which is the scale built into our mental hardware. (For more about how we perceive quantity, see *The Number Sense* [9].) The geometric mean is the correct mean when combining quantities produced by our mental hardware.

Here, the geometric mean is 30 people per square mile: a factor of 10 removed from either endpoint. Using that population density,

$$\text{US population} \sim 6 \times 10^6 \frac{\text{miles}^2}{\text{miles}^2} \times \frac{30}{\text{miles}^2} \approx 2 \times 10^8. \quad (1.18)$$

The actual population is roughly 3×10^8 . The estimate based almost entirely on gut reasoning is within a factor of 1.5 of the actual population—a pleasantly surprising accuracy.

Problem 1.8 More gut estimates

By asking your gut to help you estimate the lower and upper endpoints, estimate (a) the height of a nearby tall tree that you can see, (b) the mass of a car, and (c) the number of water drops in a bathtub.

1.7 Physical estimates

Your gut understands not only the social world but also the physical world. If you trust its feelings, you can tap this vast reservoir of knowledge. For practice, we'll estimate the salinity of seawater (Section 1.7.1), human power output (Section 1.7.2), and the heat of vaporization of water (Section 1.7.3).

1.7.1 Salinity of seawater

To estimate the salinity of seawater, which will later help you estimate the conductivity of seawater (Problem 8.10), do not ask your gut directly: "How do you feel about, say, 200 millimolar?" Although that kind of question worked for estimating population density (Section 1.6), here, unless you are a chemist, the answer will be: "I have no clue. What is a millimolar anyway? I have almost no experience of that unit." Instead, offer your gut concrete data—for example, from a home experiment: adding salt to a cup of water until the mixture tastes as salty as the ocean.

This experiment can be a thought or a real experiment—another example of using multiple methods (Section 1.5). As a thought experiment, I ask my gut about various amounts of salt in a cup of water. When I propose adding 2 teaspoons, it responds, "Disgustingly salty!" At the lower end, when I propose adding 0.5 teaspoons, it responds, "Not very salty." I'll use 0.5 and 2 teaspoons as the lower and upper endpoints of the range. Their midpoint, the estimate from the thought experiment, is 1 teaspoon per cup.

I tested this prediction at the kitchen sink. With 1 teaspoon (5 milliliters) of salt, the cup of water indeed had the sharp, metallic taste of seawater that I have gulped after being knocked over by large waves. A cup of water is roughly one-fourth of a liter or 250 cubic centimeters. By mass, the resulting salt concentration is the following product:

$$\frac{1 \text{ tsp salt}}{1 \text{ cup water}} \times \frac{1 \text{ cup water}}{250 \text{ g water}} \times \frac{5 \text{ cm}^3 \text{ salt}}{1 \text{ tsp salt}} \times \frac{2 \text{ g salt}}{\underbrace{1 \text{ cm}^3 \text{ salt}}_{\rho_{\text{salt}}}} \quad (1.19)$$

The density of 2 grams per cubic centimeter comes from my gut feeling that salt is a light rock, so it should be somewhat denser than water at 1 gram per cubic centimeter, but not too much denser. (For an alternative method, more accurate but more elaborate, try Problem 1.10.) Then doing the arithmetic gives a 4 percent salt-to-water ratio (by mass).

The actual salinity of the Earth's oceans is about 3.5 percent—very close to the estimate of 4 percent. The estimate is close despite the large number of assumptions and approximations—the errors have mostly canceled. Its accuracy should give you courage to perform home experiments whenever you need data for divide-and-conquer estimates.

Problem 1.9 Density of water

Estimate the density of water by asking your gut to estimate the mass of water in a cup measure (roughly one-quarter of a liter).


Problem 1.10 Density of salt

Estimate the density of salt using the volume and mass of a typical salt container that you find in a grocery store. This value should be more accurate than my gut estimate in Section 1.7.1 (which was 2 grams per cubic centimeter).


1.7.2 Human power output

Our second example of talking to your gut is an estimate of human power output—a power that is useful in many estimates (for example, Problem 1.17). Energies and powers are good candidates for divide-and-conquer estimates, because they are connected by the subdivision shown in the following equation and represented in the tree in the margin:

$$\text{power} = \frac{\text{energy}}{\text{time}}. \quad (1.20)$$



In particular, let's estimate the power that a trained athlete can generate for an extended time (not just during a few-seconds-long, high-power burst). As a proxy for that power, I'll use my own burst power output with two adjustment factors:



Maintaining a power is harder than producing a quick burst. Therefore, the first adjustment factor, my steady power divided by my burst power, is somewhat smaller than 1—maybe 1/2 or 1/3. In contrast, an athlete's


power output will be higher than mine, perhaps by a factor of 2 or 3: Even though I am sometimes known as the street-fighting mathematician [33], I am no athlete. Then the two adjustment factors roughly cancel, so my burst power should be comparable to an athlete's steady power.

To estimate my burst power, I performed a home experiment of running up a flight of stairs as quickly as possible. Determining the power output requires estimating an energy and a time:

$$\text{power} = \frac{\text{energy}}{\text{time}}. \quad (1.21)$$

The energy, which is the change in my gravitational potential energy, itself subdivides into three factors:

$$\text{energy} = \underbrace{\text{mass}}_m \times \underbrace{\text{gravity}}_g \times \underbrace{\text{height}}_h. \quad (1.22)$$



In the academic building at my university, a building with high ceilings and staircases, I bounded up a staircase three stairs at a time. The staircase was about 12 feet or 3.5 meters high. Therefore, my mechanical energy output was roughly 2000 joules:

$$E \sim 65 \text{ kg} \times 10 \text{ m s}^{-2} \times 3.5 \text{ m} \sim 2000 \text{ J}. \quad (1.23)$$

(The units are fine: $1 \text{ J} = 1 \text{ kg m}^2 \text{ s}^{-2}$.)

The remaining leaf is the time: how long the climb took me. I made it in 6 seconds. In contrast, several students made it in 3.9 seconds—the power of youth! My mechanical power output was about 2000 joules per 6 seconds, or about 300 watts. (To check whether the estimate is reasonable, try Problem 1.12, where you estimate the typical human basal metabolism.)

This burst power output should be close to the sustained power output of a trained athlete. And it is. As an example, in the Alpe d'Huez climb in the 1989 Tour de France, the winner—Greg LeMond, a world-class athlete—put out 394 watts (over a 42.5-minute period). The cyclist Lance Armstrong, during the time-trial stage during the Tour de France in 2004, generated even more: 495 watts (roughly 7 watts per kilogram). However, he publicly admitted to blood doping to enhance performance. Indeed, because of widespread doping, many cycling power outputs of the 1990s and 2000s are suspect; 400 watts stands as a legitimate world-class sustained power output.

Copyrighted image

5 m

An electric stove requires a line voltage of 220 volts, even in the United States where most other appliances require only 110 volts. A standard fuse is about 15 amperes, which gives us an idea of a large current. If a burner corresponds to a standard fuse, a burner supplies roughly 3 kilowatts:

$$220 \text{ V} \times 15 \text{ A} \approx 3000 \text{ W}. \quad (1.26)$$


This estimate agrees with the gut estimate, so both methods gain plausibility—which should give you confidence to use both methods for your own estimates. As a check, I looked at the circuit breaker connected to my range, and it is rated for 50 amperes. The range has four burners and an oven, so 15 amperes for one burner (at least, for the large burner) is plausible.

We now have values for all the leaf nodes. Propagating the values toward the root gives the heat of vaporization (L_{vap}) as roughly 2 megajoules per kilogram:

$$L_{\text{vap}} \sim \frac{\frac{\text{power}}{3 \text{ kW}} \times \frac{\text{time}}{600 \text{ s}}}{\frac{1 \text{ kg}}{\text{mass}}} \quad (1.27)$$

$$\approx 2 \times 10^6 \text{ J kg}^{-1}.$$

The true value is about 2.2×10^6 joules per kilogram. This value is one of the highest heats of vaporization of any liquid. As water evaporates, it carries away significant amounts of energy, making it an excellent coolant (Problem 1.17).



1.8 Summary and further problems

The main lesson that you should take away is courage: No problem is too difficult. We just use divide-and-conquer reasoning to dissolve difficult problems into smaller pieces. (For extensive practice, see the varied examples in the *Guesstimation* books [47 and 48].) This tool is a universal solvent for problems social and scientific.

Problem 1.14 Per-capita land area

Estimate the land area per person for the world, for your home country, and for your home state or province.

Problem 1.15 Mass of the Earth

Estimate the mass of the Earth. Then look it up (p. xvii) to check your estimate.

Problem 1.16 Billion

How long would it take to count to a billion (10^9)?

Problem 1.17 Sweating

Estimate how much water you need to drink to replace water lost to evaporation, if you ride a bicycle vigorously for 1 hour. Represent your estimate as a divide-and-conquer tree. Hint: Humans are only about 25 percent efficient in generating mechanical work.

Problem 1.18 Pencil line

How long a line can you write with a pencil?

Problem 1.19 Pine needles

Estimate the number of needles on a pine tree.

Problem 1.20 Hairs

How many hairs are on your head?

2

Abstraction

2.1	Energy from burning hydrocarbons	28
2.2	Coin-flip game	31
2.3	Purpose of abstraction	34
2.4	Analogies	36
2.5	Summary and further problems	53

Divide-and-conquer reasoning, the tool introduced in Chapter 1, is powerful, but it is not enough by itself to organize the complexity of the world. Try, for example, to manage the millions of files on a computer—even my laptop says that it has almost 3 million files. Without any organization, with all the files in one monster directory or folder, you could never find information that you need. However, simply using divide and conquer by dividing the files into groups—the first 100 files by date, the second 100 files by date, and so on—does not disperse the chaos. A better solution is to organize the millions of files into a hierarchy: as a tree of folders and subfolders. The elements in this hierarchy get names—for example, “photos of the children” or “files for typesetting this book”—and these names guide us to the needed information.

Naming—or, more technically, abstraction—is our other tool for organizing complexity. A name or an abstraction gets its power from its reusability. Without reusable ideas, the world would become unmanageably complicated. We might ask, “Could you, without tipping it over, move the wooden board glued to four thick sticks toward the large white plastic circle?” instead of, “Could you slide the chair toward the table?” The abstractions “chair,” “slide,” and “table” compactly represent complex ideas and physical structures. (And even the complex question itself uses abstractions.)

Similarly, without good abstractions we could hardly calculate, and modern science and technology would be impossible. As an illustration, imagine the pain of the following calculation:

$$\text{XXVII} \times \text{XXXVI}, \quad (2.1)$$


which is 27×36 in Roman numerals. The problem is not that the notation is unfamiliar, but rather that it is not based on abstractions useful for calculation. Not least, it does not lend itself to divide-and-conquer reasoning; for example, even though V (5) is a part of XXVII, $V \times \text{XXXVI}$ has no obvious answer. In contrast, our modern number system, based on the abstractions of place value and zero, makes the whole multiplication simple. Notations are abstractions, and good abstractions amplify our intelligence. In this chapter, we will practice making abstractions, discuss their high-level purpose, and continue to practice.


2.1 Energy from burning hydrocarbons

Our understanding of the world is built on layers of abstractions. Consider the idea of a fluid. At the bottom of the abstraction hierarchy are the actors of particle physics: quarks and electrons. Quarks combine to build protons and neutrons. Protons, neutrons, and electrons combine to build atoms. Atoms combine to build molecules. And large collections of molecules act, under many conditions, like a fluid.


The idea of a fluid is a new unit of thought. It helps us understand diverse phenomena, without our having to calculate or even know how quarks and electrons interact to produce fluid behavior. As one consequence, we can describe the behavior of air and water using the same equations (the Navier–Stokes equations of fluid mechanics); we need only to use different values for the density and viscosity. Then atmospheric cyclones and water vortices, although they result from widely differing sets of quarks and electrons and their interactions, can be understood as the same phenomenon.

A similarly powerful abstraction is a chemical bond. We'll use this abstraction to estimate a quantity essential to our bodies and to modern society: the energy released by burning chains made of hydrogen and carbon atoms (hydrocarbons). A hydrocarbon can be abstracted as a chain of CH_2 units:






Burning a CH₂ unit requires oxygen (O₂) and releases carbon dioxide (CO₂), water, and energy:



For a hydrocarbon with eight carbons—such as octane, a prime component of motor fuel—simply multiply this reaction by 8:



(The two additional hydrogens at the left and right ends of octane are not worth worrying about.)

► How much energy is released by burning one CH₂ unit?

To make this estimate, use the table of bond energies. It gives the energy required to break (not make) a chemical bond—for example, between carbon and hydrogen. However, there is no unique carbon–hydrogen (C–H) bond. The carbon–hydrogen bonds in methane are different from the carbon–hydrogen bonds in ethane. To make a reusable idea, we neglect those differences—placing them below our abstraction barrier—and make an abstraction called the carbon–hydrogen bond. So the table, already in its first column, is built on an abstraction.

The second gives the bond energy in kilocalories per mole of bonds. A kilocalorie is roughly 4000 joules, and a mole is Avogadro’s number (6×10^{23}) of bonds. The third column gives the energy in the SI units used by most of the world, kilojoules per mole. The final column gives the energy in electron volts (eV) per bond. An electron volt is 1.6×10^{-19} joules. An electron volt is suited for measuring atomic energies, because most bond energies have an easy-to-grasp value of a few electron volts. I wish most of the world used this unit!

	bond energy		
	$\left(\frac{\text{kcal}}{\text{mol}}\right)$	$\left(\frac{\text{kJ}}{\text{mol}}\right)$	$\left(\frac{\text{eV}}{\text{bond}}\right)$
C–H	99	414	4.3
O–H	111	464	4.8
C–C	83	347	3.6
C–O	86	360	3.7
H–H	104	435	4.5
C–N	73	305	3.2
N–H	93	389	4.0
O=O	119	498	5.2
C=O	192	803	8.3
C=C	146	611	6.3
N≡N	226	946	9.8


Playing many games might reveal a pattern to us or suggest how to compute the probability. However, playing many games by flipping a real coin becomes tedious. Instead, a computer can simulate the games, substituting pseudorandom numbers for a real coin. Here are several runs produced by a computer program. Each line begins with 1 or 2 to indicate which player won the game; the rest of the line shows the coin tosses. In these ten iterations, each player won five times. A reasonable conjecture is that each player has an equal chance to win. However, this conjecture, based on only ten games, cannot be believed too strongly.

2 TH
 2 TH
 1 H
 2 TH
 1 TTH
 2 TTTT
 2 TH
 1 H
 1 H
 1 H

Let's try 100 games. Now even counting the wins becomes tedious. My computer counted for me: 68 wins for player 1, and 32 wins for player 2. The probability of player 1's winning now seems closer to 2/3 than to 1/2.

To find the exact value, let's diagram the game as a tree reflecting the alternative endings of the game. Each layer represents one flip. The game ends at a leaf, when one player has tossed heads. The shaded leaves show the first player's wins—for example, after H, TTH, or TTTTH. The probabilities of these winning ways are 1/2 (for H), 1/8 (for TTH), and 1/32 (for TTTTH). The sum of all these winning probabilities is the probability of the first player's winning:

$$\frac{1}{2} + \frac{1}{8} + \frac{1}{32} + \dots \tag{2.5}$$



To sum this infinite series without resorting to formulas, make an abstraction: Notice that the tree contains, one level down, a near copy of itself. (In this problem, the abstraction gets reused within the same problem. In computer science, such a structure is called recursive.) For if the first player tosses tails, the second player starts the game in the position of the first player, with the same probability of winning.

To benefit from this equivalence, let's name the reusable idea, namely the probability of the first player's winning, and call it p . The second player wins the game with probability $p/2$: The factor of 1/2 is the probability that the first player tosses tails; the factor of p is the probability that the second player wins, given that the first player blew his chance by tossing tails on the first toss.

Because either the first or the second player wins, the two winning probabilities add to 1:

$$\underbrace{p}_{P(\text{first player wins})} + \underbrace{p/2}_{P(\text{second player wins})} = 1. \quad (2.6)$$

The solution is $p = 2/3$, as suggested by the 100-game simulation. The benefit of the abstraction solution, compared to calculating the infinite probability sum explicitly, is insight. In the abstraction solution, the answer has to be what it is. It leaves almost nothing to remember. An amusing illustration of the same benefit comes from the problem of the fly that zooms back and forth between two approaching trains.

- *If the fly starts when the trains are 60 miles apart, each train travels at 20 miles per hour, and the fly travels at 30 miles per hour, how far does the fly travel, in total, before meeting its maker when the trains collide? (Apologies that physics problems are often so violent.)*

Right after hearing the problem, John von Neumann, inventor of game theory and the modern computer, gave the correct distance. “That was quick,” said a colleague. “Everyone else tries to sum the infinite series.” “What’s wrong with that?” said von Neumann. “That’s how I did it.” In Problem 2.7, you get to work out the infinite-series and the insightful solutions.

Problem 2.4 Summing a geometric series using abstraction

Use abstraction to find the sum of the infinite geometric series

$$1 + r + r^2 + r^3 + \dots \quad (2.7)$$

Problem 2.5 Using the geometric-series sum

Use Problem 2.4 to check that the probability of the first player’s winning is $2/3$:

$$p = \frac{1}{2} + \frac{1}{8} + \frac{1}{32} + \dots = \frac{2}{3}. \quad (2.8)$$

Problem 2.6 Nested square roots

Evaluate these infinite mixes of arithmetic and square roots:

$$\sqrt{3 \times \sqrt{3 \times \sqrt{3 \times \sqrt{3 \times \dots}}}} \quad (2.9)$$

$$\sqrt{2 + \sqrt{2 + \sqrt{2 + \sqrt{2 + \dots}}}} \quad (2.10)$$

Problem 2.7 Two trains and a fly

Find the insightful and the infinite-series solution to the problem of the fly and the approaching trains (Section 2.2). Check that they give the same answer for the distance that the fly travels!

Problem 2.8 Resistive ladder

In the following infinite ladder of 1-ohm resistors, what is the resistance between points A and B? This measurement is indicated by the ohmmeter connected between these points.


**2.3 Purpose of abstraction**

The coin game (Section 2.2), like the geometric series (Problem 2.4) or the resistive ladder (Problem 2.8), contained a copy of itself. Noticing this reuse greatly simplified the analysis. Abstraction has a second benefit: giving us a high-level view of a problem or situation. Abstractions then show us structural similarities between seemingly disparate situations.

As an example, let's revisit the geometric mean, introduced in Section 1.6 to make gut estimates. The geometric mean of two nonnegative quantities a and b is defined as

$$\text{geometric mean} \equiv \sqrt{ab}. \quad (2.11)$$

This mean is called the geometric mean because it has a pleasing geometric construction. Divide the diameter of a circle into two lengths, a and b , and inscribe a right triangle whose hypotenuse is the diameter. The triangle's altitude is the geometric mean of a and b .



This mean reappears in surprising places, including the beach. When you stand at the shore and look at the horizon, you are seeing a geometric mean. The distance to the horizon is the geometric mean of two important lengths in the problem (Problem 2.9).

For me, its most surprising appearance was in the "Programming and Problem-Solving Seminar" course taught by Donald Knuth [40] (who also created \TeX , the typesetting system for this book). The course, taught as a series of two-week problems, helped first-year PhD students transition from undergraduate homework problems to PhD research problems. A homework problem requires perhaps 1 hour. A research problem requires, say,

1000 hours: roughly a year of work, allowing for other projects. (A few problems stapled together become a PhD.) In the course, each 2-week module required about 30 hours—approximately the geometric mean of the two endpoints. The modules were just the right length to help us cross the bridge from homework to research.

Problem 2.9 Horizon distance

How far is the horizon when you are standing at the shore? Hint: It's farther for an adult than for a child.

Problem 2.10 Distance to a ship

Standing at the shore, you see a ship (drawn to scale) with a 10-meter mast sail into the distance and disappear from view. How far away was it when it disappeared?



As further evidence that the geometric mean is a useful abstraction, the idea appears even when there is no geometric construction to produce it, such as in making gut estimates. We used this method in Section 1.6 to estimate the population density and then the population of the United States. Let's practice by estimating the oil imports of the United States in barrels per year—without the divide-and-conquer reasoning of Section 1.4.

The method requires that the gut supply a lower and an upper bound. My gut reports back that it would feel fairly surprised if the imports were less than 10 million barrels per year. On the upper end, my gut would be fairly surprised if the imports were higher than 1 trillion barrels per year—a barrel is a lot of oil, and a trillion is a large number!

You might wonder how your gut too can come up with such large numbers and how you can have any confidence in them. Admittedly, I have practiced a lot. But you can practice too. The key is the practice effectively. First, have the courage to guess even when you feel anxious about it (I feel this anxiety still, so I practice this courage often). Second, compare your guess to values in which you can place more confidence—for example, to your own more careful estimates or to official values. The comparison helps calibrate your gut (your right brain) to these large magnitudes. You will find a growing and justified confidence in your judgment of magnitude.

My best guess for the amount is the geometric mean of the lower and upper estimates:

$$\sqrt{10 \text{ million} \times 1 \text{ trillion}} \frac{\text{barrels}}{\text{year}}. \quad (2.12)$$

The result is roughly 3 billion barrels per year—close to our estimate using divide and conquer and close to the true value. In contrast, the arithmetic mean would have produced an estimate of 500 billion barrels per year, which is far too high.

Problem 2.11 Arithmetic-mean–geometric-mean inequality

Use the geometric construction for the geometric mean to show that the arithmetic mean of a and b (assumed to be nonnegative) is always greater than or equal to their geometric mean. When are the means equal?

Problem 2.12 Weighted geometric mean


A generalization of the arithmetic mean of a and b as $(a + b)/2$ is to give a and b unequal weights. What is the analogous generalization for a geometric mean? (The weighted geometric mean shows up in Problem 6.29 when you estimate the contact time of a ball bouncing from a table.)

2.4 Analogies

Because abstractions are so useful, it is helpful to have methods for making them. One way is to construct an analogy between two systems. Each common feature leads to an abstraction; each abstraction connects our knowledge in one system to our knowledge in the other system. One piece of knowledge does double duty. Like a mental lever, analogy and, more generally, abstraction are intelligence amplifiers.

2.4.1 Electrical–mechanical analogies

An illustration with many abstractions on which we can practice is the analogy between a spring–mass system and an inductor–capacitor (LC) circuit.



In the circuit, the voltage source—the V_{in} on its left side—supplies a current that flows through the inductor (a wire wrapped around an iron rod) and capacitor (two metal plates separated by air). As current flows through the capacitor, it alters the charge on the capacitor. This “charge” is confusingly named, because the net charge on the capacitor remains zero. Instead,

2.4.2 Energy density in the gravitational field

With the electrical–mechanical analogy as practice, let’s try a less familiar analogy: between the electric and the gravitational field. In particular, we’ll connect the energy densities (energy per volume) in the corresponding fields. An electric field E represents an energy density of $\epsilon_0 E^2/2$, where ϵ_0 is the permittivity of free space appearing in the electrostatic force between two charges q_1 and q_2 :

$$F = \frac{q_1 q_2}{4\pi\epsilon_0 r^2}. \quad (2.19)$$

Because electrostatic and gravitational forces are both inverse-square forces (the force is proportional to $1/r^2$), the energy densities should be analogous. Not least, there should be a gravitational energy density. But how is it related to the gravitational field?

To answer that question, our first step is to find the gravitational analog of the electric field. Rather than thinking of the electric field only as something electric, focus on the common idea of a field. In that sense, the electric field is the object that, when multiplied by the charge, gives the force:

$$\text{force} = \text{charge} \times \text{field}. \quad (2.20)$$

We use words rather than the normal symbols, such as E for field or q for charge, because the symbols might bind our thinking to particular cases and prevent us from climbing the abstraction ladder.

This verbal form prompts us to ask: What is gravitational charge? In electrostatics, charge is the source of the field. In gravitation, the source of the field is mass. Therefore, gravitational charge is mass. Because field is force per charge, the gravitational field strength is an acceleration:

$$\text{gravitational field} = \frac{\text{force}}{\text{charge}} = \frac{\text{force}}{\text{mass}} = \text{acceleration}. \quad (2.21)$$

Indeed, at the surface of the Earth, the field strength is g , also called the acceleration due to gravity.

The definition of gravitational field is the first half of the puzzle (we are using divide-and-conquer reasoning again). For the second half, we’ll use the field to compute the energy density. To do so, let’s revisit the route from electric field to electrostatic energy density:

$$E \rightarrow \frac{1}{2}\epsilon_0 E^2. \quad (2.22)$$

With g as the gravitational field, the analogous route is

$$g \rightarrow \frac{1}{2} \times \text{something} \times g^2, \quad (2.23)$$

where the “something” represents our ignorance of what to do about ϵ_0 .

► *What is the gravitational equivalent of ϵ_0 ?*

To find its equivalent, compare the simplest case in both worlds: the field of a point charge. A point electric charge q produces a field

$$E = \frac{1}{4\pi\epsilon_0} \frac{q}{r^2}. \quad (2.24)$$

A point gravitational charge m (a point mass) produces a gravitational field (an acceleration)

$$g = \frac{Gm}{r^2}, \quad (2.25)$$

where G is Newton’s constant.

The gravitational field has a similar structure to the electric field. Both are inverse-square forces, as expected. Both are proportional to the charge. The difference is the constant of proportionality. For the electric field, it is $1/4\pi\epsilon_0$. For the gravitational field, it is simply G . Therefore, G is analogous to $1/4\pi\epsilon_0$; equivalently, ϵ_0 is analogous to $1/4\pi G$.

Then the gravitational energy density becomes

$$\frac{1}{2} \times \frac{1}{4\pi G} \times g^2 = \frac{g^2}{8\pi G}. \quad (2.26)$$

We will use this analogy in Section 9.3.3 when we transfer our hard-won knowledge of electromagnetic radiation to understand the even more subtle physics of gravitational radiation.

Problem 2.13 Gravitational energy of the Sun

What is the energy in the gravitational field of the Sun? (Just consider the field outside the Sun.)

Problem 2.14 Pendulum period including buoyancy


The period of a pendulum in vacuum is (for small amplitudes) $T = 2\pi\sqrt{l/g}$, where l is the bob length and g is the gravitational field strength. Now imagine the pendulum swinging in a fluid (say, air). By replacing g with a modified value, include the effect of buoyancy in the formula for the pendulum period.

Problem 2.15 Comparing field energies

Find the ratio of electrical to gravitational field energies in the fields produced by a proton.

2.4.3 Parallel combination

Analogies not only reuse work, they help us rewrite expressions in compact, insightful forms. An example is the idea of parallel combination. It appears in the analysis of the infinite resistive ladder of Problem 2.8.



To find the resistance R across the ladder (in other words, what the ohmmeter measures between the nodes A and B), you represent the entire ladder as a single resistor R . Then the whole ladder is 1 ohm in series with the parallel combination of 1 ohm and R :

$$\begin{array}{c} | \\ \text{---} \\ | \\ \text{---} \\ | \\ \text{---} \\ | \\ \text{---} \\ | \\ \text{---} \\ | \\ \text{---} \\ | \\ \text{---} \\ | \end{array} \quad = \quad \begin{array}{c} | \\ \text{---} \\ | \\ \text{---} \\ | \\ \text{---} \\ | \\ \text{---} \\ | \end{array} \quad (2.27)$$

The next step in finding R usually invokes the parallel-resistance formula: that the resistance of R_1 and R_2 in parallel is

$$\frac{R_1 R_2}{R_1 + R_2}. \quad (2.28)$$

For our resistive ladder, the parallel combination of 1 ohm with the ladder is $1 \text{ ohm} \times R / (1 \text{ ohm} + R)$. Placing this combination in series with 1 ohm gives a resistance

$$1 \Omega + \frac{1 \Omega \times R}{1 \Omega + R}. \quad (2.29)$$

This recursive construction reproduces the ladder, only one unit longer. We therefore get an equation for R :

$$R = 1 \Omega + \frac{1 \Omega \times R}{1 \Omega + R}. \quad (2.30)$$

The (positive) solution is $R = (1 + \sqrt{5})/2$ ohms. The numerical part is the golden ratio ϕ (approximately 1.618). Thus, the ladder, when built with 1-ohm resistors, offers a resistance of ϕ ohms.

Although the solution is correct, it skips over a reusable idea: the parallel combination. To facilitate its reuse, let's name the idea with a notation:

$$R_1 \parallel R_2. \quad (2.31)$$

This notation is self-documenting, as long as you recognize the symbol \parallel to mean “parallel,” a recognition promoted by the parallel bars. A good notation should help thinking, not hinder it by requiring us to remember how the notation works. With this notation, the equation for the ladder resistance R is

$$R = 1 \Omega + 1 \Omega \parallel R \quad (2.32)$$

(the parallel-combination operator has higher priority than—is computed before—the addition). This expression more plainly reflects the structure of the system, and our reasoning about it, than does the version

$$R = 1 \Omega + \frac{1 \Omega \times R}{1 \Omega + R}. \quad (2.33)$$

The \parallel notation organizes the complexity.


Once you name an idea, you find it everywhere. As a child, after my family bought a Volvo, I saw Volvos on every street. Similarly, we'll now look at examples of parallel combination far beyond the original appearance of the idea in circuits. For example, it gives the spring constant of two connected springs (Problem 2.16):

$$\underbrace{\text{Spring } k_1} \parallel \underbrace{\text{Spring } k_2} = \underbrace{\text{Spring } k_1 \parallel k_2} \quad (2.34)$$

Problem 2.16 Springs as capacitors

Using the analogy between springs and capacitors (discussed in Section 2.4.1), explain why springs in series combine using the parallel combination of their spring constants.

Another surprising example is the following spring–mass system with two masses:




The natural frequency ω , expressed without our \parallel abstraction, is

$$\omega = \frac{k(m + M)}{mM}. \quad (2.35)$$

This form looks complicated until we use the \parallel abstraction:

$$\omega = \frac{k}{m \parallel M}. \quad (2.36)$$

Now the frequency makes more sense. The two masses act like their parallel combination $m \parallel M$:



The replacement mass $m \parallel M$ is so useful that it has a special name: the reduced mass. Our abstraction organizes complexity by turning a three-component system (a spring and two masses) into a simpler two-component system.

In the spirit of notation that promotes insight, use lowercase (“small”) m for the mass that is probably smaller, and uppercase (“big”) M for the mass that is probably larger. Then write $m \parallel M$ rather than $M \parallel m$. These two forms produce the same result, but the $m \parallel M$ order minimizes surprise: The parallel combination of m and M is smaller than either mass (Problem 2.17), so it is closer to m , the smaller mass, than to M . Writing $m \parallel M$, rather than $M \parallel m$, places the most salient information first.

Problem 2.17 Using the resistance analogy

By using the analogy with parallel resistances, explain why $m \parallel M$ is smaller than m and M .

► *Why do the two masses combine like resistors in parallel?*

The answer lies in the analogy between mass and resistance. Resistance appears in Ohm’s law:

$$\text{voltage} = \text{resistance} \times \text{current}. \quad (2.37)$$

Voltage is an effort. Current, which results from the effort, is a flow. Therefore, the more general form—one step higher on the abstraction ladder—is

$$\text{effort} = \text{resistance} \times \text{flow}. \quad (2.38)$$

In this form, Newton’s second law,

With this changing voltage, the capacitor equation,

$$\text{current} = \text{capacitance} \times \frac{d(\text{voltage})}{dt}, \quad (2.49)$$

becomes

$$\text{current} = \text{capacitance} \times j\omega \times \text{voltage}. \quad (2.50)$$

Let's compare this form to its analog for a resistor (Ohm's law):

$$\text{current} = \frac{1}{\text{resistance}} \times \text{voltage}. \quad (2.51)$$


Matching up the pieces, we find that a capacitor offers a resistance

$$Z_C = \frac{1}{j\omega C}. \quad (2.52)$$


This more general resistance, which depends on the frequency, is called impedance and denoted Z . (In the analogy of Section 2.4.1 between capacitors and springs, we found that capacitor offered a resistance to being charged of $1/C$. Impedance, the result of an analogy between capacitors and resistors, contains $1/C$ as well, but also contains the frequency in the $1/j\omega$ factor.)

Using impedance, we can describe what happens to any sinusoidal signal in a circuit containing capacitors. Our thinking is aided by the compact notation—the capacitive impedance Z_C (or even R_C). The notation hides the details of the capacitor differential equation and allows us to transfer our intuition about resistance and flow to a broader class of circuits.

The simplest circuit with resistors and capacitors is the so-called low-pass RC circuit. Not only is it the simplest interesting circuit, it will also be, thanks to further analogies, a model for heat flow. Let's apply the impedance analogy to this circuit.



To help us make and use abstractions, let's imagine defocusing our eyes. Under blurry vision, the capacitor looks like a resistor that just happens to have a funny resistance $R_C = 1/j\omega C$. Now the entire circuit looks just like a pure-resistance circuit. Indeed, it is the simplest such circuit, a voltage divider. Its behavior is described by one number: the gain, which is the ratio of output to input voltage V_{out}/V_{in} .



In the RC circuit, thought of as a voltage divider,

$$\text{gain} = \frac{\text{capacitive resistance}}{\text{total resistance from } V_{in} \text{ to ground}} = \frac{R_C}{R + R_C}. \quad (2.53)$$