

6

Lumping

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


In 1982, thousands of students in the United States had to estimate 3.04×5.3 , choosing 1.6, 16, 160, 1600, or “I don’t know.” Only 21 percent of 13-year-olds and 37 percent of 17-year-olds chose 16. As Carpenter and colleagues describe [7], the problem is not a lack of calculation skill. On questions testing exact multiplication (“multiply 2.07 by 9.3”), the 13-year-olds scored 57 percent, and the 17-year-olds scored 72 percent correct. The problem is a lack of understanding; if you earn roughly \$5 per hour for roughly 3 hours, your net worth cannot grow by \$1600. The students needed our next tool: rounding or, more generally, lumping.

choice	age 13	age 17
1.6	28%	21%
16	21	37
160	18	17
1600	23	11
don’t know	9	12

6.1 Approximate!

Fortunately, rounding is inherent in our perception of quantity: Beyond three items, our perception of “how many” comes with an inherent imprecision. This fuzziness is, for adults, 20 percent: If we briefly see two groups of dots whose totals are within 20 percent, we cannot easily judge which group has more dots. Try it by glancing at the following squares.

In the left pair, one square contains 10 percent more dots than the other square; in the right pair, one square contains 30 percent more dots than the other square. In the 10-percent pair, spotting the more numerous square is difficult. In the 30-percent pair, it is almost obvious at sight. Lumping comes naturally; we just need the courage to do it. We'll develop the courage first in rounding numbers, the most familiar kind of lumping.



6.2 Rounding on a logarithmic scale

Just as driving to visit a next-door neighbor atrophies our muscles and ability to move around the physical world, asking calculators to do simple arithmetic dulls our ability to navigate the quantitative world. We never develop an innate sense of sizes and scales in the world. The antidote is to do the computations ourselves, but approximately—by placing quantities on a logarithmic scale and rounding them to the nearest convenient value.


6.2.1 Rounding to the nearest power of ten

The simplest method of rounding is to round every number to the nearest power of ten. That simplification turns most calculations into adding and subtracting integer exponents (the exceptions come from roots, which produce fractional exponents). Here, “nearest” is judged on a logarithmic scale, where distance is measured not with differences but with ratios or factors. For example, 50—although closer to 10 than to 100 on a linear scale—is a factor of 5 greater than 10 but only a factor of 2 smaller than 100. Therefore, 50 is closer to 100 than to 10 and would get rounded to 100.


As practice, let's estimate the number of minutes in a day:

$$1 \text{ day} \times \frac{24 \text{ hours}}{\text{day}} \times \frac{60 \text{ minutes}}{\text{hour}} = 24 \times 60 \text{ minutes.} \quad (6.1)$$

Now we round each factor to the nearest power of 10. Because 24 is only a factor of 2.4 away from 10, but more than a factor of 4 away from 100, it gets rounded to 10:



In contrast, 60 is closer to 100 than to 10:



With these approximations, 1 day is approximately 1000 minutes:

$$1 \text{ day} \times \frac{\text{from } 24}{\overbrace{10}^{\text{from } 60}} \text{ hours} / \text{day} \times \frac{\overbrace{100}^{\text{from } 60}}{\text{hour}} \text{ minutes} = 1000 \text{ minutes.} \quad (6.2)$$

The exact value is 1440 minutes, so the estimate is only 30 percent too small. This error is a reasonable tradeoff to gain a method that requires almost no effort—who needs a calculator to multiply 10 by 100? Furthermore, the accuracy is enough for many calculations, where insight is needed more than accuracy.

Problem 6.1 Rounding to the nearest power of ten

Round these numbers to the nearest power of ten: 200, 0.53, 0.03, and 7.9.

Problem 6.2 Boundary between rounding up or down

We saw that 60 rounded up to 100 but that 24 rounded down to 10. What number is just at the boundary between rounding down to 10 and rounding up to 100?

Problem 6.3 Using rounding to the nearest power of ten

Round the numbers to the nearest power of ten and thereby estimate the products: (a) 27×50 , (b) 432×12 , and (c) 3.04×5.3 .

Problem 6.4 Calculating the bending of light

In Section 5.3.1, we used dimensional analysis to show that the Earth could bend starlight by approximately the angle

$$\theta \sim \frac{\overbrace{6.7 \times 10^{-11} \text{ kg}^{-1} \text{ m}^3 \text{ s}^{-2}}^G \times \overbrace{6 \times 10^{24} \text{ kg}}^{m_{\text{Earth}}}}{\underbrace{6.4 \times 10^6 \text{ m}}_{R_{\text{Earth}}} \times \underbrace{10^{17} \text{ m}^2 \text{ s}^{-2}}_{c^2}}. \quad (6.3)$$

Round each factor to the nearest power of ten in order to estimate the bending angle mentally. How long did making the estimate take?

6.2.2 Rounding to the nearest half power of ten


Rounding to the nearest power of ten gives a quick, preliminary estimate. When it is too approximate, we just round more precisely. The next increase in accuracy is to round to the nearest *half* power of ten.

As an illustration, let's estimate the number of seconds in a year. The numerical calculation (without the units) is $365 \times 24 \times 3600$. Now we round each factor by replacing it with a number of the form 10^β . In the previous method, where we rounded to the nearest power of ten, β was an integer. Now β can also be a half-integer (for example, 2.5).


► *What is a half power of ten numerically?*

Two half powers of ten multiply to make 10, so each half power is $\sqrt{10}$, or slightly more than 3 (as you found in Problem 6.2). When you need more precision, a half power of ten is 3.2 or 3.16, although that much precision is rarely needed.

In rounding the calculation for the number of seconds, 365 becomes $10^{2.5}$, and, as diagrammed below, 3600 becomes $10^{3.5}$:



The remaining factor is 24. It is closer to $10^{1.5}$ (roughly 30) than to 10^1 :



Thus, we replace 24 by $10^{1.5}$. Then the calculation simplifies to

$$\underbrace{10^{2.5}}_{365} \times \underbrace{10^{1.5}}_{24} \times \underbrace{10^{3.5}}_{3600}. \quad (6.4)$$

Now just add the powers of ten:

$$10^{2.5} \times 10^{1.5} \times 10^{3.5} = 10^{7.5}. \quad (6.5)$$

Because the 0.5 in the exponent 7.5 contributes a factor of 3, there are about 3×10^7 seconds in a year. I remember this value as $\pi \times 10^7$, which is accurate to 0.5 percent.

Problem 6.5 Earth's orbital speed

Using the estimate of $\pi \times 10^7$ seconds in a year, estimate the Earth's orbital speed around the Sun. Don't use a calculator! (The Earth–Sun distance, 1.5×10^{11} meters, is worth memorizing.)

Problem 6.6 Where does the π come from?

True or false: The π in $\pi \times 10^7$ seconds per year arises because the Earth orbits the Sun in a circle, and a circle has π in its circumference.

Problem 6.7 Only approximately pi

True or false: The π in $\pi \times 10^7$ seconds per year is not exact because the Earth orbits in a slightly noncircular ellipse.

Problem 6.8 Estimating geometric means

In Section 2.3, I talked to my gut to estimate the US annual oil imports. That discussion resulted in the geometric-mean estimate

$$\sqrt{10 \text{ million} \times 1 \text{ trillion}} \text{ barrels per year.} \quad (6.6)$$

Estimate the square root by placing the two quantities 10 million and 1 trillion on a logarithmic scale and finding their midpoint.

6.3 Typical or characteristic values

Lumping not only simplifies numbers, where it is called rounding, it also simplifies complex quantities by creating an abstraction: the typical or characteristic value. We've used this form of lumping implicitly many times; now it's time to discuss it explicitly, in order to appreciate its scope and to develop skill in applying it.

6.3.1 Estimating the population of the United States

Our first explicit example of a typical or characteristic value occurs in the following population estimate. Knowing a population is essential in many estimates about the social world, such as a country's oil imports (Section 1.4), energy consumption, or per-capita land area (Problem 1.14). Here we'll estimate the population of the United States by dividing it into two factors.

$$\text{US population} \sim N_{\text{states}} \times \text{population of a typical US state.} \quad (6.7)$$

The first factor, N_{states} , everyone in America learns in school: $N_{\text{states}} = 50$. The second factor contains the lumping approximation. Rather than using all 50 different state populations, we replace each state with a typical state.

► *What is the population of a typical state?*

A large state is California or New York, each of which has a megacity with a population in the several millions. A tiny state is Delaware or Rhode Island. In between is Massachusetts. Because I live there, I know that it has 6 million people. Taking Massachusetts as a typical state, the US population is roughly 50×6 million or 300 million. This estimate is more accurate than we deserve: The 2012 population is 314 million.

From this example, you can also see how lumping enhances symmetry reasoning: When there is change, look for what does not change (Section 3.1). Here, each state has its own population, so there's plenty of change among the list of states. Lumping helps us find, or create, a quantity that does not change. We imagine a typical state, one that may not even exist (just as no family has the average of 2.3 children), and replace every state with this state. We've lumped away all the change—throwing away information in exchange for insight into the population of the country.

Problem 6.9 German federal states

The Federal Republic of Germany has 16 federal states. Pick one at random, multiply its population by 16, and compare that estimate with Germany's population.

6.3.2 Lumping varying physical quantities: How high can animals jump?

Using typical or characteristic values allows us to reason out seemingly impossible questions while sitting in our armchairs. For practice, we'll study how the jump height of an animal depends on its size. For example, should a person be able to jump higher than a locust?

The jump could be a running or a standing high jump. Both types are interesting, but the standing high jump teaches more about lumping. Therefore, imagine that the animal starts from rest and jumps directly upward.

Even with this assumption, the problem looks underspecified. The jump height depends at least on the animal's shape, on how much muscle the animal has, and on the muscle efficiency. It is the kind of problem where lumping, a tool for making assumptions, is most helpful. We'll use lumping and proportional reasoning to find the scaling exponent β in $h \propto m^\beta$, where h is the jump height and m is the animal's mass (our measure of its size).

Finding a scaling exponent usually requires a physical model. You can often build them by imagining an extreme, unrealistic situation, and then


asking yourself what physics prevents it from happening. Thus, why can't we jump to the Moon? Because it demands a vast amount of energy, far beyond what our muscles can supply. The point to extract from this thought experiment is that jumping demands energy, which is supplied by muscles.

The appearance of supply and demand suggests, as in the estimate of the number of taxis in Boston (Section 3.4.1), that we equate the demand to the supply. Then we estimate each piece separately—divide and conquer.

The energy demanded is the gravitational potential energy mgh . Here, g is the gravitational acceleration, and the jump height h is measured as the vertical change in the animal's center of mass (CM). Because all animals feel the same gravity, $E_{\text{demanded}} = mgh$ simplifies to the proportionality $E_{\text{demanded}} \propto mh$.

For the supplied energy, we again divide and conquer:

$$E_{\text{supplied}} \sim \text{muscle mass} \times \text{muscle energy density}, \quad (6.8)$$



where the muscle energy density is the energy per mass that muscle can supply. This product already contains a lumping assumption: that all the muscles in an animal provide the same energy density. This assumption is reflected in the single approximation sign (\sim).

Even so, the result is not simple enough. Each species, and each individual within a species, will have its own muscle mass and energy density. Therefore, let's make the further lumping assumption that all animals, even though they vary in muscle mass, share the same muscle energy density. The assumption is plausible because all muscles use similar biological technology (actin and myosin filaments). Fortunately, making this assumption is an approximation: Lumping throws away actual information, which is how it reduces complexity. With the assumption of a common energy density, the supplied energy becomes the simpler proportionality

$$E_{\text{supplied}} \propto \text{muscle mass}. \quad (6.9)$$

The simplicity of this result reminds us that an approximate result can be more useful than an exact result.

The muscle mass also varies from animal to animal. Introducing a dimensionless prefactor and making another lumping approximation will manage that complexity. The dimensionless prefactor α will be the fraction of an animal's mass that is muscle:

$$m_{\text{muscle}} = \alpha m. \quad (6.10)$$

Alas, α varies across species (compare a cheetah and a turtle), within a species, and within the lifetime of an individual—for example, my α is dropping as I sit writing this book. If we account for all these variations, we will be overwhelmed by their complexity. Lumping rescues us: It gives us permission to assume that α is the same for every animal. We replace the diversity of animals with a typical animal. This assumption is not as crazy as it might sound. It doesn't mean that all animals have the same muscle mass. Rather, it means that all animals have the same *fraction* of muscle; as an example, for people, $\alpha \sim 0.4$.

With this lumping assumption, $m_{\text{muscle}} = \alpha m$ becomes the simpler proportionality $m_{\text{muscle}} \propto m$. Because the supplied energy is proportional to the muscle mass, it is proportional to the animal's mass:

$$E_{\text{supplied}} \propto m_{\text{muscle}} \propto m. \quad (6.11)$$

This result is as simple as we can hope for, and it depends on the right quantity, the animal's mass. Now let's use it to predict how an animal's jump height depend on this mass.

Because the demanded energy and supplied energy are equal, and the demanded energy—the gravitational potential energy—is proportional to mh ,

$$mh \propto m. \quad (6.12)$$

The common factor of m^1 cancels out, leaving h independent of m :

$$h \propto m^0. \quad (6.13)$$

All (jumping) animals should be able to jump to the same height!

The result always surprises me. My intuition, before doing the calculation, cannot decide how h should behave. On the one hand, small animals seem strong: Ants can lift a mass many times their own body mass, whereas humans can lift only roughly their own body mass. However, my intuition also insists that people should be able to jump higher than locusts.

The data, from *Scaling: Why Animal Size Is So Important* [41, p. 178], contradict my intuition and confirm our lumping and scaling analysis. The mass spans more than eight orders of magnitude (eight decades, or a factor of 10^8), yet the jump height, as a change in the height of the center of mass, varies only by a factor of 3 (half a decade). On a log–log graph, the height-versus-mass line would have a slope less than $1/16$. Our predicted scaling

	m	h
flea	0.5 mg	20 cm
click beetle	40 mg	30 cm
locust	3 g	59 cm
human	70 kg	60 cm

of constant jump height ($h \propto m^0$), which corresponds to a slope of zero, is surprisingly accurate. The outliers, fleas and click beetles, are at the lighter end. (For an explanation, try Problem 6.10.) Furthermore, at the heavier end, locusts and humans, although differing by more than four orders of magnitude in mass, jump to almost exactly the same height.

A moral of this example is that lumping augments proportional reasoning. Proportional reasoning reduces complexity by showing us a notation for ignoring quantities that do not vary. For example, when all animals face the same gravitational field, then $E_{\text{demanded}} = mgh$ simplifies to $E_{\text{demanded}} \propto mh$. Alas, we live in the desert of the real, where “the same” is almost always only an approximation—for example, as with the energy density of muscle in different animals. Lumping rescues us. It gives us permission to replace these changing values with a single, constant, typical value—making the relations amenable to proportional reasoning.

Problem 6.10 Jumping fleas


The prediction of constant jump height seems to fail at small sizes: Larger animals jump about 60 centimeters, whereas fleas reach only 20 centimeters. In this problem, you evaluate whether air drag can explain the discrepancy.

- As a lumping approximation, pretend that an animal is a cube with side length l , and assume that it jumps to a height h independent of its mass m . Then find the scaling exponent β in $E_{\text{drag}}/E_{\text{demanded}} \propto l^\beta$, where E_{drag} is the energy consumed by drag and E_{demanded} is the energy needed in the absence of drag.
- Estimate $E_{\text{drag}}/E_{\text{demanded}}$ for a cubical human that can jump to 60 centimeters. Using the scaling relation, estimate it for a cubical flea. What is your judgment of drag as an explanation for the lower jump heights of fleas?

6.3.3 Period of an ideal spring

A surprising conclusion of dimensional analysis is that the period of a spring, or a small-amplitude pendulum, does not depend on its amplitude (Problem 5.14). However, the mathematical reasoning doesn't give us the why; it doesn't provide a physical insight. That insight comes from lumping, by using characteristic values. We'll try it together by finding the period of a spring; you'll practice by finding the period of a pendulum (Problem 6.11).

This kind of lumping, in exchange for providing physical insight, requires that we make a physical model. Here, a stretched or compressed spring



exerts a force on and thereby accelerates the mass. If the spring is stretched to an amplitude x_0 , then the force is kx_0 and the acceleration is kx_0/m . This acceleration varies as the mass moves, so analyzing the motion usually requires differential equations. However, this acceleration is also a characteristic acceleration. It sets the scale for the acceleration at all other times. If we replace the changing acceleration with this characteristic acceleration, the complexity vanishes. The problem becomes a constant-acceleration problem where $a \sim kx_0/m$.

With a constant acceleration a for a time comparable to one period T , the mass moves a distance comparable to aT^2 , which is kx_0T^2/m . When applying lumping and characteristic values, “comparable” is the verbal translation of the single approximation sign \sim . As an equation, we would write

$$\text{distance} \sim aT^2. \quad (6.14)$$

Another useful translation is “of the order of”: The distance is of the order of aT^2 . Equivalently, the characteristic distance is aT^2 .

This characteristic distance must be comparable to the amplitude x_0 . Thus,

$$x_0 \sim \frac{kx_0T^2}{m}. \quad (6.15)$$

The amplitude x_0 cancels out! Then the period is $T \sim \sqrt{m/k}$. Lumping thus provides the following explanation for why the period is independent of amplitude: As the amplitude and therefore the travel distance increase, the force and acceleration increase just enough to compensate, leaving the period unchanged.

Problem 6.11 Period of a small-amplitude pendulum using lumping

Use characteristic values to explain why the period of a (small-amplitude) pendulum is independent of the amplitude θ_0 .

Problem 6.12 Period of a nonlinear spring

Imagine a nonlinear spring with force law $F \propto x^n$. Use lumping as follows to find how the period T varies with amplitude x_0 .

- Estimate a typical or characteristic acceleration.
- At this acceleration, roughly how far does the mass travel in one period T ?
- This distance must be comparable to the amplitude x_0 . Therefore, find the scaling exponent α in $T \propto x_0^\alpha$ (where α will be a function of the scaling exponent n in the force law). Then check your answer to Problem 5.17.

6.3.4 Lumping derivatives

The preceding analysis of the period of a spring–mass system (Section 6.3.3) illustrates a general simplification: Using characteristic values, we can replace derivatives with algebra. The algebraic expression usually provides a physical model. As an example, let’s explain, physically, an acceleration that we derived by dimensional analysis in Section 5.1.1: the inward acceleration of an object moving in a circle.

Acceleration is the derivative of velocity: $a = dv/dt$. Using the definition of derivative,

$$\frac{dv}{dt} \equiv \frac{\text{infinitesimal change in } v}{(\text{infinitesimal}) \text{ time required to make this change in } v}. \quad (6.16)$$

Infinitesimal changes and times are difficult to picture, so an analysis based on calculus often does not help us see why a result is true.

A lumping approximation, by discarding complexity, can give this insight. A way to remember the lumping approximation, first use $6 = 6$ to cancel the 6s in $16/64$:

$$\frac{16}{64} = \frac{1}{4}. \quad (6.17)$$

The result is exact! Although this particular cancellation is dubious, it suggests the analogous lumping approximation $d = d$. The resulting cancellation turns derivatives into algebra:

$$\frac{dv}{dt} \sim \frac{v}{t}. \quad (6.18)$$

► What does v/t mean?

Lumping replaces “infinitesimal” with “characteristic”:

$$\frac{v}{t} \sim \frac{\text{characteristic change in } v}{\text{time required to make this change in } v}. \quad (6.19)$$

The numerator asks us to look at the changes in v and to represent them by a characteristic or typical change. The denominator is, as an abbreviation, often called the characteristic time, or time constant, and denoted τ .

In applying this approximation to circular motion, we have to distinguish the velocity vector \mathbf{v} from its magnitude v (the speed). The speed, at least in constant-speed circular motion, never changes, so dv/dt itself is zero. We


are interested in $|d\mathbf{v}|/dt$: the magnitude of the vector's derivative, rather than the derivative of the vector's magnitude. The lumped acceleration a is

$$a \sim \frac{|\text{characteristic change in } \mathbf{v}|}{\text{time required to make this change in } \mathbf{v}}. \quad (6.20)$$

The maximum change in \mathbf{v} is reversing direction, from $+\mathbf{v}$ to $-\mathbf{v}$. A characteristic change in \mathbf{v} is \mathbf{v} itself, or any value comparable to it. This range of possibilities is captured by the single approximation sign, which stands for "is comparable to." With that notation,

$$\text{characteristic change in } \mathbf{v} \sim \mathbf{v}. \quad (6.21)$$

Here is an example of such a change, showing the velocity vectors before and after the particle has rotated partway around the circle.



When \mathbf{v} makes its significant change, from \mathbf{v}_{old} to \mathbf{v}_{new} , it produces a change in \mathbf{v} comparable in magnitude to v . The characteristic time τ —the time required to make this change—is a decent fraction of a period of revolution. Because a full period is $2\pi r/v$, the characteristic time is comparable to r/v . Using τ to represent the characteristic time, and \sim to hide dimensionless prefactors, we write $\tau \sim r/v$.

Here's another way to arrive at $\tau \sim r/v$. If the triangle of \mathbf{v}_{old} , \mathbf{v}_{new} , and $\Delta\mathbf{v}$ is an equilateral triangle, then it corresponds to the particle rotating 60° around the circle, which is almost exactly 1 radian. Because 1 radian creates an arc of length r (the same proportionality tells us that 2π radians produces the circumference $2\pi r$), the travel time is r/v , and $\tau \sim r/v$. Then the acceleration is roughly v^2/r :


$$a \sim \frac{v}{\tau} \sim \frac{v}{r/v} = \frac{v^2}{r}. \quad (6.22)$$

This equation encapsulates a physical, proportional-reasoning explanation of $a = v^2/r$. Namely, in circular motion, the velocity vector changes direction significantly in approximately 1 radian of rotation. This motion requires a time $\tau \sim r/v$. Therefore, the circular acceleration a contains two factors of v —one factor from the v itself and one factor from the time in the denominator—and it contains $1/r$, also from the time in the denominator.

6.3.5 Simplifying with characteristic values: Yield of the first atomic blast

In Section 5.2.2, we used dimensional analysis to estimate the yield of an atomic blast. We accurately predicted that the blast energy E is related to the blast radius R and the air density ρ :

$$E \sim \frac{\rho R^5}{t^2}, \quad (6.23)$$



where t is the time since the blast. However, dimensional analysis, as a mathematical argument, does not give a physical explanation for its results, which can often feel like magic. Lumping explains the magic by helping us analyze a physical model—a model whose analysis would otherwise be absurd in its complexity.

The physical model is that the blast increases the thermal energy of the air molecules, thus increasing the speed of sound—which is the speed at which the blast expands. Using this model exactly requires setting up and solving differential equations, as Taylor did [44]. Lumping, by turning calculus into algebra, simplifies the equations without discarding their physical meaning.

The first step is to estimate the thermal energy. It comes almost entirely from the hot fireball—that is, from the blast energy. This energy is spread unevenly over the blast, with a higher energy density nearer the blast center and a lower density farther away. For the lumping approximation, smear the blast energy E evenly throughout a sphere of radius R . This sphere's volume is comparable to R^3 , so the typical or characteristic energy density is $\mathcal{E} \sim E/R^3$.

The next step is to use this energy density to estimate the speed of sound c_s . As we discussed in Section 5.4.1, the speed of sound is comparable to the thermal speed, so

$$c_s \sim \sqrt{\frac{k_B T}{m}}, \quad (6.24)$$

where $k_B T$ is the approximate thermal energy of one air molecule and m is the mass of one air molecule.

To connect this speed to the energy density \mathcal{E} , let's convert energy per molecule into energy per volume, by multiplying $k_B T$ by the number density n (the number of air molecules per volume). The result is $n k_B T$, which is the thermal and therefore roughly the blast energy density $\mathcal{E} \sim E/R^3$.

To be fair, we also multiply the denominator m by n . This step converts mass per molecule into mass per volume, which gives the density ρ :

$$\underbrace{\frac{\text{mass}}{\text{molecule}}}_m \times \underbrace{\frac{\text{molecules}}{\text{volume}}}_n = \underbrace{\frac{\text{mass}}{\text{volume}}}_\rho. \quad (6.25)$$

Therefore, the speed of sound is comparable to $\sqrt{E/\rho R^3}$:

$$c_s \sim \sqrt{\frac{nk_B T}{nm}} \sim \sqrt{\frac{E/R^3}{\rho}} = \sqrt{\frac{E}{\rho R^3}}. \quad (6.26)$$

This speed is the rate at which the blast expands.

► *Based on this speed, how large should the blast be after time t ?*

Because the energy density and therefore the sound speed decreases as R increases, the blast radius is not simply the speed at time t multiplied by the time. But we can make a further lumping approximation: that the typical or characteristic speed, for the entire time t , is $\sqrt{E/\rho R^3}$. In evaluating this speed, we'll use the radius R at time t as the characteristic radius.

$$\underbrace{\text{blast radius}}_R \sim \underbrace{\text{characteristic speed}}_{c_s \sim \sqrt{E/\rho R^3}} \times \underbrace{\text{time}}_t. \quad (6.27)$$

The solution for the blast energy E is $E \sim \rho R^5/t^2$, as we found using dimensional analysis. Lumping complements that mathematical reasoning with a physical model.

6.4 Applying lumping to shapes


Lumping by replacing varying quantities with their typical, or characteristic, values is close to the next form of lumping: lumping shapes. Our first illustration is explaining a curious fact about everyday materials.

6.4.1 Densities of liquids and solids

Among books teaching the art of approximation, a classic is *Consider a Spherical Cow* [22], so named because a sphere is a much simpler shape than a cow. An even simpler shape is a cube. Thus, a powerful form of lumping is to replace complex shapes by a comparably sized cube. With this idea, we can explain why most solids and liquids have densities between 1 and 10 times the density of water.

Each atom is a complex, ill-defined shape, but pretend that it is cube. Because the atoms touch, the density of the substance is approximately the density of one approximating cube:

$$\rho \approx \frac{\text{mass of the atom}}{\text{volume of the lumped cube}}. \quad (6.28)$$



To evaluate the numerator, use A as the atom's atomic mass. Although A is called a mass, it is dimensionless: It is almost exactly the total number of protons and neutrons in the nucleus. Because protons and neutrons have almost the same mass, the proton mass m_p can represent the mass of either constituent. Then the mass of one cube is Am_p . To find the denominator, the cube's volume, make each cube's side be a typical atomic diameter of $a \sim 3$ ångströms. This size is based on our calculation in Section 5.5.1 of the diameter of the smallest atom, hydrogen (whose diameter was roughly 1 ångström). Then the density becomes

$$\rho \sim \frac{Am_p}{(3 \text{ \AA})^3}. \quad (6.29)$$

To avoid looking up the proton mass m_p , let's multiply this fraction by N_A/N_A , where N_A is Avogadro's number:

$$\rho \sim \frac{A \times m_p N_A}{(3 \text{ \AA})^3 \times N_A}. \quad (6.30)$$

The numerator is A grams per mole, because 1 mole of protons (roughly, of hydrogen) has a mass of 1 gram.

The denominator is roughly 18 cubic centimeters per mole:

$$\frac{3 \times 10^{-23} \text{ cm}^3}{(3 \text{ \AA})^3} \times \frac{6 \times 10^{23} \text{ mol}^{-1}}{N_A} = \frac{18 \text{ cm}^3}{\text{mol}}. \quad (6.31)$$

A typical solid or liquid density is then related simply to the substance's atomic mass A :

$$\rho \sim \frac{A \text{ g mol}^{-1}}{18 \text{ cm}^3 \text{ mol}^{-1}} = \frac{A}{18} \frac{\text{g}}{\text{cm}^3}. \quad (6.32)$$

Problem 6.13 Rounding to estimate atomic volume

Use rounding to the nearest half power of ten (Section 6.2.2) to show that a cube with side length 3 ångströms has a volume of approximately 3×10^{-23} cubic centimeters.

Many common elements have atomic masses between 18 and 180, so the densities of many liquids and solids should lie between 1 and 10 grams per cubic centimeter. As shown in the table, the prediction is reasonable. The table even includes a joker: water is not an element! Yet the density estimate is exact.

The table also shows why centimeters and grams are, for materials physics, more convenient than are meters and kilograms. A typical solid has a density of a few grams per cubic centimeter. Such modest numbers are easy to remember and handle. In contrast, a density of 3000 kilograms per cubic meter, although mathematically equivalent, is mentally unwieldy. On each use, you have to think, "How many powers of ten were there again?" Therefore, the table gives densities in the mentally friendly units of grams per cubic centimeter.

	A	ρ (g cm ⁻³)	
		est.	actual
Li	7	0.4	0.5
H ₂ O	18	1.0	1.0
Si	28	1.6	2.4
Fe	56	3.1	7.9
Hg	201	11.2	13.5
U	238	13.3	18.7
Au	197	10.9	19.3

Problem 6.14 Accounting for discrepancies in the density prediction

In the table of densities, iron (Fe) shows the biggest discrepancy between the actual and predicted densities, by roughly a factor of 2.5. Use that factor to make an improved estimate for the interatomic spacing in iron.

Problem 6.15 Number density of conduction electrons

Estimate the number density of free (conduction) electrons in a copper wire.

Problem 6.16 Typical drift speed in a wire


- Use your estimate of the number density of conduction electrons in a copper wire (Problem 6.15) to estimate the drift speed of electrons in the wire connecting a typical lamp to the wall socket.
- Estimate the time required for electrons to travel at this drift speed from the wall socket to the light bulb. Then how does the light bulb turn on right after you flip the switch?

6.4.2 Graph lumping: The number of undergraduate students


A particularly important shape is a graph. When applied to graphs, the idea of lumping shapes simplifies many a problem, making evident its essential features. As an example of graph lumping, we'll estimate the number of US undergraduate students. Such back-of-the-envelope market-size estimates are valuable for business planning and for making public policy.

The first step is to estimate the number of people in the United States who are 18, 19, 20, or 21 years old. This total provides, at least in the United States, the pool from which most undergraduate students come. Because not all 18-to-21-year-olds go to college, at the end we will multiply the total by the fraction of adults who are college graduates.

Finding the exact pool size requires the birth date of every person in the United States. Although these data are collected once every decade by the US Census Bureau, they would only overwhelm us. As an approximation to the voluminous data, the Census Bureau also publishes the number of people at each age. For example, the 1991 data are the wiggly line in the graph. The left side of the graph represents the number of infants and toddlers in 1991, and the right side represents the number of older people (also in 1991). The undergraduate pool size, representing all 18-, 19-, 20-, and 21-year-olds, is the shaded area. (The peak around the ages 30–35 represents the baby boomers, born in the period after World War Two.)




Unfortunately, even this graph depends on the huge resources of the Census Bureau, so it is not suited for back-of-the-envelope estimates. It also provides little insight or transfer value. Insight comes from lumping: from turning the complex, wiggly curve into a rectangle. The rectangle's dimensions can be determined without any information from the Census Bureau.



► What are the height and width of this rectangle?

The rectangle's width is a time, so it must be a characteristic time related to the population. A good guess is the life expectancy, because the age distribution varies significantly over that time. In the United States, the life expectancy is roughly 75 years, which will be the rectangle's width. In this lumping approximation, everyone lives happily until a sudden death at his or her 75th birthday. This all-or-nothing reasoning is the essential characteristic of lumping, making it such a useful approximation.



The rectangle's height can be computed from its area, which is the US population—approximately 300 million (as we estimated in Section 6.3.1). Therefore, the height is 4 million per year:


$$\text{height} = \frac{\text{area}}{\text{width}} \sim \frac{300 \text{ million}}{75 \text{ years}} = \frac{4 \text{ million}}{\text{year}}. \quad (6.33)$$

Based on this estimate, there should be 16 million people (in the United States) with the four undergraduate ages of 18, 19, 20, or 21 years.

Not everyone in this pool is an undergraduate. Therefore, the final step is to account for the fraction of adults who are college graduates. In the United States, where education has traditionally been widely spread throughout the population, this fraction (or adjustment factor) is high—say, 0.5. The number of US undergraduates should be about 8 million.

For comparison, the 2010 US census data gives 5.361 million enrolled in four-year colleges, and 4.942 million in two-year colleges, for a combined total of almost exactly 10 million. Our estimate, which lies halfway between the four-year and the combined total, is quite good!

Even had it been terrible, with a large discrepancy between the estimate and the actual number, we would still have learned useful information about a society. As an example, let's use the same method to estimate how many university students graduated in 1950 in the United Kingdom. ("College" in American English and "university" in British English are roughly equivalent.) The UK population at the time was 50 million. With a life expectancy of, say, 65 years, the rectangle's height is roughly 800 000 per year. If, as in the United States of today, 50 percent of the age-eligible population goes to university, 400 000 should graduate each year. However, in 1950, the actual number was roughly 17 000—more than a factor of 20 smaller than the estimate.



Such a huge error probably does not come from the population estimate. Instead, the fraction of 50 percent of university participation must be far too high. Indeed, rather than 50 percent, the actual figure for 1950 is 3.4 percent. The difference between 50- and 3-percent university participation makes the United Kingdom of 1950 a very different society from the United States of 2010 or even from the United Kingdom of 2010, where the fraction going to university is roughly 40 percent.


6.4.3 Cone free-fall time and distance

For our next illustration of graph lumping, let's turn from the social to the physical world, to our falling cones. For the single cone of Section 3.5.2 and the cone race of Section 4.3.1, we assumed that the cones made their entire journey at their terminal speed. But this assumption cannot be exact: Just after the cones are released, their fall speed is zero. Graph lumping will help us evaluate and refine this assumption.


► *How far does the cone fall before it reaches its terminal speed?*

Taken literally, the answer is infinity, because no object reaches its terminal speed. As it approaches its terminal speed, the drag and weight more nearly balance, the force and acceleration get closer to zero, and the speed changes ever more slowly. So an object can approach its terminal speed only ever more closely. Let's therefore rephrase the question to ask how far the cone falls before it has reached a significant fraction of its terminal speed. In the "significant," you can see the door through which we will bring in lumping.


Here is a sketch of the actual fall speed versus time. At first, the speed increases rapidly. As the speed increases, so does the drag. The net force and the acceleration fall, so the speed increases more slowly. As a lumping approximation, we'll replace the smooth curve by a slanted and a horizontal segment.



A partly triangular lumping approximation might look different from the population lumping rectangles we constructed in Section 6.4.2. However, it is merely the integral of a rectangle: It is equivalent to replacing the actual, complicated acceleration by a rectangle representing the period of free fall. Then the acceleration drops abruptly to zero.



Just as we labeled the population rectangle with its height and width, here we label the velocity graph with speeds, times, and slopes. The graph has two segments. The second, horizontal segment shows the cone's falling at its terminal speed v_{term} . This speed, as we found experimentally in Section 3.5.2, is roughly 1 meter per second.




The first, slanted segment represents free fall, as if there were no air resistance. The free-fall acceleration is g , so the free-fall velocity has slope g : In 1 second, the speed increases by 10 meters per second. The slanted and horizontal segments meet where $gt = v_{\text{term}}$. Thus, they meet at approximately 0.1 seconds. Based on the lumping approximation, we expect the cone to reach a significant fraction of its terminal speed by 0.1 seconds—which is only 5 percent of the total fall time of 2 seconds.

► *Roughly how far does the cone fall in this time?*

Use lumping again! The distance is the area of the shaded triangle. The triangle's base is 0.1 seconds, and its height is the terminal speed, approximately 1 meter per second. So its area is 5 centimeters:

$$\frac{1}{2} \times 0.1 \text{ s} \times 1 \text{ m s}^{-1} = 5 \text{ cm.} \quad (6.34)$$



After only 2.5 percent of its 2-meter journey, the cone has reached a significant fraction of the terminal speed. The “always at terminal speed” approximation is quite accurate. Thanks to the lumping, we could judge the approximation without setting up or solving differential equations.

Problem 6.17 Raindrop terminal speed

Sketch the fall speed of a large raindrop versus time. Roughly how long and how far does it fall before it reaches (a significant fraction of) its terminal speed?

Problem 6.18 Actual cone fall speed

Set up and solve the differential equations for the fall of a cone with drag proportional to v^2 and terminal speed of 1 meter per second. What is its speed as a fraction of v_{term} after the cone (a) has fallen for 0.1 seconds, or (b) has fallen 5 centimeters?


6.4.4 How viscosity burns up energy

Lumping is particularly useful for gaining insight into fluid flow, a subject where the governing equations, the Navier–Stokes equations, are so complex that almost no problem has an exact solution. These equations were introduced in Section 3.5 to encourage you to use conservation reasoning. Here their specter is invoked to encourage you to use lumping.


In everyday life, an important feature of fluid flow is drag. As we discussed in Section 5.3.2, drag (in steady-state flow) results from viscosity: Without viscosity, there can be no drag. Somehow, viscosity dissipates energy.

Using graph lumping, we can understand the mechanism of energy dissipation without having to understand the detailed physics of viscosity. The essential physical idea is that the viscous force, a force from a neighboring region of fluid, slows down fast pieces of fluid and speeds up slow pieces.

As an example, the arrows in the velocity-profile diagram show a fluid's horizontal speed above a flat boundary—for example, the air speed above a frozen lake. Farther above the boundary, the arrows are longer, indicating that the fluid moves faster. In a high-viscosity liquid, such as honey, the viscous forces are large, and they force nearby regions move at nearly the same speed: The flow oozes.



Here is the lumped version of the velocity profile. The varying velocity has been replaced by two rectangles corresponding to two chunks of fluid. The top chunk is moving faster than the bottom chunk; because of the velocity difference, each chunk exerts a viscous force on the other.



Let's see what consequence this pair of forces has for the total energy of the chunks. To avoid cluttering the essential idea in the analysis with unit algebra, let's give the chunks concrete velocities and masses in a simple unit system. In this unit system, both chunks will have unit mass. The top chunk will move at speed 6 and the bottom chunk at speed 4.

► *After a long time, how will viscosity have affected the velocities of the two slabs?*

Because momentum is conserved, the final speeds sum to 10, as they do at the start. Because of the viscous force between the chunks, the top chunk slows down, and the bottom chunk speeds up. The final speeds, once viscosity has done its work, literally and figuratively, are 5 and 5.

The total kinetic energy of the two chunks starts at 26:

$$0.5 \times 1 \times (6^2 + 4^2) = 26. \quad (6.35)$$

However, their final kinetic energy is only 25:

$$0.5 \times 1 \times (5^2 + 5^2) = 25. \quad (6.36)$$


Simply because the velocities equalized, the kinetic energy fell! This reduction happens no matter what the initial velocities are (Problem 6.19). The energy difference turns into heat. Our simple physical model, based on lumping, is that viscosity burns up energy by equalizing velocities.

Problem 6.19 Generalizing the argument to other initial velocities

Let the initial velocities of the two chunks be v_1 and v_2 , where $v_1 \neq v_2$. Show that the initial kinetic energy is greater than the final kinetic energy.

Problem 6.20 Shortest-time path

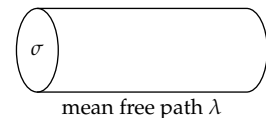
The classic problem in the calculus of variations is to find the shortest-time path for a mass sliding without friction between two points. This path is called the brachistochrone. A surprising conclusion of the full analysis is that this path need not be straight. You can make this result plausible by using shape lumping: Rather than considering all possible paths, consider only paths with one corner. Under what conditions is such a path faster than the straight path (zero corners)?

**6.4.5 Mean free path**

Lumping smooths out variation. In the famous words of Isaiah 40:4, *the crooked shall be made straight, and the rough places plain*. We've used this idea in Section 6.4.2 to simplify a function of time (population versus age) and in Section 6.4.4 to simplify a function of space (a velocity profile). Now we'll combine both kinds of lumping to understand and estimate a mean free path. The mean free path is the average distance that a gas molecule travels before colliding with another molecule. As we will find in Section 7.3.1 using probabilistic reasoning, mean free paths determine important material properties, including viscosity and thermal conductivity.

Let's first estimate the mean free path in the simplest model: a spherical molecule moving in a gas of point molecules. The sphere will have radius r and the gas molecules a number density n . After analyzing this simplified situation, we'll replace the point molecules with more realistic spherical molecules.

The moving molecule has a cross-sectional area $\sigma = \pi r^2$, and it sweeps out a tube with the same cross-sectional area (analogous to the tube in our analysis of drag via conservation of energy in Section 3.5.1).



► *How far does the spherical molecule travel down the tube before it hits a point molecule?*

This distance is the mean free path λ . However, finding it is complicated because the point molecules, always in motion, keep exiting and entering the tube. Let's simplify. First, lump in time: Freeze the motion, and pretend

that the point molecules hold those positions. Then lump in space: Rather than spreading the molecules throughout the tube, place them at the far end, at a distance λ from the near end.

The tube, whose length is the mean free path λ , should be just long enough so that the tube contained, before lumping, one point molecule. Then, in the lumped model with that one molecule sitting at the end of the tube, the spherical molecule will hit one molecule after traveling a distance λ .

The number of molecules in the tube is $n\sigma\lambda$:

$$\text{number} = \underbrace{\text{number density}}_n \times \underbrace{\text{volume}}_{\sigma\lambda} = n\sigma\lambda. \quad (6.37)$$


Therefore, λ is determined by the requirement that

$$n\sigma\lambda \sim 1. \quad (6.38)$$

For motion through the gas of point molecules, σ is πr^2 , so

$$\lambda \sim \frac{1}{n\pi r^2}. \quad (6.39)$$

To make the model more realistic, let's replace the point molecules with spherical molecules. For simplicity, and to model the most frequent case, these molecules will also have radius r .



► How does this change affect the mean free path?

Now a collision happens if the centers of two molecules approach within a distance $d = 2r$, the diameter of the sphere. Thus, σ —called the scattering cross section—becomes $\pi(2r)^2$ or πd^2 . The mean free path is then a factor of 4 smaller and is

$$\lambda \sim \frac{1}{n\pi d^2}. \quad (6.40)$$

Let's evaluate λ for an air molecule traveling in air. For the diameter d , use the typical atomic diameter of 3 ångströms or 3×10^{-10} meters, which also works for small molecules (like air and water). To estimate the number density n , use the ideal gas law in the form that, at standard temperature and pressure, 1 mole occupies 22 liters. Therefore,

$$\frac{1}{n} = \frac{22 \ell}{6 \times 10^{23} \text{ molecules}}, \quad (6.41)$$

and the mean free path is

$$\lambda \sim \frac{2.2 \times 10^{-2} \text{ m}^3}{6 \times 10^{23}} \times \frac{1}{\pi \times (3 \times 10^{-10} \text{ m})^2}. \quad (6.42)$$

To evaluate this expression in your head, divide the calculation into three steps—doing the most important first.

1. *Units.* The numerator contains cubic meters; the denominator contains square meters. Their quotient is meters to the first power—as it should be for a mean free path.
2. *Powers of ten.* The numerator contains -2 powers of ten. The denominator contains 3 powers of ten: 23 in Avogadro's number and -20 in $(10^{-10})^2$. The quotient is -5 powers of ten. Along with the units, the expression so far is 10^{-5} meters.
3. *Everything else.* The remaining factors are

$$\frac{2.2}{6 \times \pi \times 3 \times 3}. \quad (6.43)$$


The $2.2/6$ is roughly $10^{-0.5}$. The three factors of 3 (one contributed by π) are $10^{1.5}$. Therefore, the remaining factors contribute 10^{-2} .

Putting the pieces together, the mean free path becomes 10^{-7} meters, which is 100 nanometers—quite close to the true value of 68 nanometers.

Problem 6.21 Electric field due to a uniform sheet of charge

In this problem, you investigate a physical model to estimate the electric field above a uniform sheet of charge with areal charge density σ .

- a. Explain why the electric field must be vertical.
- b. At a height z above the sheet, what region of the sheet contributes significantly to this field?
- c. By lumping the sheet into a point charge with the same charge as the significant region, estimate the electric field at a height z , and give the scaling exponent n in $E \propto z^n$.



Confirm your result by comparing it to the prediction from dimensional analysis (Problem 5.34).

Problem 6.22 Electric field inside a spherical shell

Another puzzling electrostatic phenomenon is that the electric field from a uniform shell of charge is zero *everywhere* inside the shell. At the center, the field must be zero by symmetry. However, even away from the center, the field is still zero. Explain this phenomenon using a physical model and lumping.

6.4.6 Lumping the path of light bent by gravity



Lumping, as we have seen, replaces a complex, changing process with a simpler, constant process. In the next example, we'll use that simplification to build and analyze a physical model for the bending of starlight by the Sun. In Section 5.3.1, using dimensional analysis and educated guessing, we concluded that the bending angle is roughly Gm/rc^2 , where m is the mass of the Sun and r is the distance of closest approach (for a ray grazing the Sun, it is just the radius of the Sun). Lumping provides a physical model for this result; this model will, in Section 8.2.2.2, allow us to predict the effect of extremely strong gravitational fields.


Imagine again a beam (or photon) of light that leaves a distant star. In its journey, it grazes the surface of the Sun and reaches our eye. To estimate the deflection angle using lumping, first identify the changing process—the source of the complexity. Here, the light beam deflects from its original, straight path and the gravitational force from the Sun changes in magnitude and direction as the photon travels. Therefore, calculating the deflection angle requires setting up and evaluating an integral—and carefully checking its trigonometric factors, such as the cosines and secants.

The antidote to complicated integrals is lumping. The lumping approximation simply pretends that the beam bends only near the Sun. In this approximation, only near the Sun does gravity operate. We further assume that, while the photon is near the Sun, its downward acceleration (the acceleration perpendicular to the path) is constant, rather than varying rapidly with position.

The problem then simplifies to estimating the deflection while the beam is near the Sun. As a further lumping approximation, let's define "near" to mean "within r on either side of the location of closest approach." The justification is dimensional: The only length in the problem is the distance of closest approach, which is r ; therefore, "near" and "far" are defined relative to the characteristic distance r .

The geometry is simplest at the point of closest approach. Therefore, let's make the further lumping approximation that, although the light beam tracks how much downward deflection should happen, the total deflection happens only at the point of closest approach.






This lumped path, instead of changing direction smoothly, has a kink (a corner) at the point of closest approach.

The deflection angle, in the small-angle approximation, is the ratio of velocity components:


$$\theta \approx \frac{v_{\perp}}{c}, \quad (6.44)$$

where c is the speed of light, which is the forward velocity, and v_{\perp} is the accumulated downward velocity. This downward velocity comes from the downward acceleration due to the Sun's gravity.

In the exact analysis, the downward acceleration varies along the path. As a function of time, it looks like a bell curve centered at the point of closest approach (labeled as $t = 0$, for symmetry). The downward velocity is the integral of (the area under) the entire $a_{\downarrow}(t)$ curve.



Our lumping analysis replaces the entire area by a rectangle centered around the peak. This rectangle has height Gm/r^2 : the characteristic (and peak) downward acceleration. It has width comparable to r/c : the time that the beam spends near the Sun. Therefore, its area, which is the lumping approximation to v_{\perp} , is roughly Gm/rc :



$$\underbrace{v_{\perp}}_{Gm/rc} \sim \underbrace{\text{characteristic downward acceleration}}_{Gm/r^2} \times \underbrace{\text{deflection time}}_{\sim r/c}. \quad (6.45)$$

The deflection angle is therefore comparable to Gm/rc^2 :

$$\theta \sim \frac{\overbrace{v_{\perp}}^{Gm/rc}}{c} = \frac{Gm}{rc^2}. \quad (6.46)$$

The lumping argument explains, with a physical model, the deflection angle that we predicted using dimensional analysis and educated guessing (Section 5.3.1). Lumping once again complements dimensional analysis.

$$\text{force} \sim \text{typical stress} \times \text{contact area.} \quad (6.48)$$

We can estimate the stress based on the strain (the fractional compression). The stress and strain are related by the Young's modulus Y :

$$\text{typical stress} \sim Y \times \text{typical strain.} \quad (6.49)$$

The strain, like the stress, varies throughout the ball. In the lumping or all-or-nothing approximation, the strain becomes a characteristic or typical strain in a region around the contact surface, and zero outside this region. The typical strain is a fractional length change:

$$\text{strain} = \frac{\text{length change}}{\text{length}}. \quad (6.50)$$

The numerator is δ : the amount by which the tip of the ball is compressed. The denominator is the size of the mysterious compressed region.

► *How large is the compressed region?*


Because the ball's radius or diameter changes, the compressed region might be the whole ball. This tempting reasoning turns out to be incorrect. To see why, imagine an extreme case: compressing a huge, 10-meter cube of rubber (rubber, because one can imagine compressing it). By pressing your finger onto the center of one face, you'll indent the face by, say, 1 millimeter over an area of 1 square centimeter. In the notation that we use for the ball, $\delta \sim 1$ millimeter, $r \sim 1$ centimeter, and $R \sim 10$ meters.

The strained region is not the whole cube, nor any significant fraction of it! Rather, its radius is comparable to the radius of the contact region—a fingertip ($r \sim 1$ centimeter). From this thought experiment, we learn that when the object is large enough ($R \gg r$), the strained volume is related not to the size of the object, but rather to the size of contact region (r).

Therefore, in the estimate for the typical strain, the length in the denominator is r . The typical strain ϵ is then δ/r .

Because the compression δ , in contrast to the contact radius r and the ball's radius R , is not easily visible, let's rewrite δ in terms of r and R . Amazingly, they are related by a geometric mean, because their geometry reproduces the geometry of the horizon distance (Section 2.3). The compression δ is analogous to one's height above sea level. The contact radius r is analogous to the horizon distance. And the ball's radius R is analogous to the Earth's radius.

Just as the horizon distance is roughly the geometric mean of the two surrounding lengths, so should the contact radius be roughly the geometric mean of the compression length and the ball's radius. On a logarithmic scale, r is roughly halfway between δ and R . (Like the horizon distance, r is actually halfway between δ and the diameter $2R$. However, the factor of 2 does not matter to this lumping analysis.)



Therefore, δ/r , which is the typical strain, is also roughly r/R .

The restoring force is therefore comparable to Yr^3/R :

$$\underbrace{\text{restoring force}}_{Yr^3/R} \sim \overbrace{Y \times \text{typical strain}}^{\text{typical stress}} \times \underbrace{\text{contact area}}_{r^2}. \quad (6.51)$$

This force balances the weight, which is comparable to ρgR^3 :

$$\frac{Yr^3}{R} \sim \rho gR^3. \quad (6.52)$$

Therefore, in dimensionless form, the contact radius is given by

$$\frac{r}{R} \sim \left(\frac{\rho g R}{Y} \right)^{1/3}, \quad (6.53)$$

and the dimensionless function f in

$$\frac{r}{R} = f \left(\frac{\rho g R}{Y} \right) \quad (6.54)$$

is a cube root (multiplied by a dimensionless constant).

► *How large is the contact radius r in practice?*

Let's plug in numbers for a superball (a small, highly elastic rubber ball) resting on the ground. The density of rubber is roughly the density of water. A superball is small—say, $R \sim 1$ centimeter. Its elastic modulus is roughly $Y \sim 3 \times 10^7$ pascals. (This elastic modulus is a factor of 300 smaller than oak's and almost a factor of 10^4 smaller than steel's.) Then

$$\frac{r}{R} \sim \left(\frac{\overbrace{10^3 \text{ kg/m}^3}^{\rho} \times \overbrace{10 \text{ m s}^{-2}}^g \times \overbrace{10^{-2} \text{ m}}^R}{\underbrace{3 \times 10^7 \text{ Pa}}_Y} \right)^{1/3}. \quad (6.55)$$

The quotient inside the parentheses is $10^{-5.5}$. Its cube root is roughly 10^{-2} , so r/R is roughly 10^{-2} , and r is roughly 0.1 millimeters (a sheet of paper).

From this example, we see how lumping builds on dimensional analysis. Dimensional analysis gave the form of the relation between r/R and $\rho g R / Y$. Lumping provided a physical model that specified the particular form.

Problem 6.24 Compression of the tip

For the superball, estimate δ , the compression of its tip when it is resting on the ground.

Problem 6.25 Contact radius of a marble

Estimate the contact radius of a small glass marble sitting on a (very) hard table.

Problem 6.26 Energy argument

The potential-energy density due to deforming a solid is comparable to $Y\epsilon^2$. Using this relation, find a second explanation for why r/R is comparable to $(\rho g R / Y)^{1/3}$.

Problem 6.27 Mountain heights

On Earth, the tallest mountain (Everest) is 9 kilometers high. On Mars, the tallest mountain (Olympus) is 27 kilometers high. Make a lumped model of a mountain to explain the factor-of-3 difference. If the same reasoning applied to the Moon, how high would its tallest mountain be? Why is it so much shorter?

Problem 6.28 Asteroid shapes

The tallest mountain on Earth is much smaller than the Earth. Using the mountain-height data of Problem 6.27, estimate the maximum radius of a planetary body made of rock (like the Earth or Mars) that has mountains comparable in size to the body. What consequence does this size have for the shapes of asteroids?

Problem 6.29 Contact time

Imagine a ball dropped from a height, hitting a hard table with impact speed v . Find the scaling exponent β in

$$\frac{\tau c_s}{R} \sim \left(\frac{v}{c_s} \right)^\beta, \quad (6.56)$$

where τ is the contact time and c_s is the speed of sound in the ball. Then write τ in the form $\tau \sim R/v_{\text{effective}}$, where $v_{\text{effective}}$ is a weighted geometric mean of the impact speed v and the sound speed c_s .

Problem 6.30 Contact force

For a small steel ball bouncing from a steel table with impact speed 1 meter per second (Problem 6.29), how large is the contact force compared to the ball's weight?


6.5 Quantum mechanics

When we introduced quantum mechanics to estimate the size of hydrogen (Section 5.5.1), it was an actor in dimensional analysis. There, quantum mechanics merely contributed a new constant of nature \hbar . This contribution was mathematical. By using lumping to introduce quantum mechanics, we will gain a physical intuition for the effect of quantum mechanics.

6.5.1 Particle in a box: Size of neutron stars

In mechanics, the simplest useful model is motion in a straight line at constant acceleration (which includes constant velocity). This model underlies numerous analyses—for example, the lumping analysis of a pendulum period (Section 6.3.3). In quantum mechanics, the simplest useful model is a particle confined to a box. Let's give the particle a mass m and the box a width a .

- What is the lowest possible energy of the particle—that is, its ground-state energy?



Because this box can be a lumping model for more complex problems, this energy will help us explain the binding energy of hydrogen. A lumping analysis starts with Heisenberg's uncertainty principle:

$$\Delta p \Delta x \sim \hbar. \quad (6.57)$$

Here, Δp is the spread (the uncertainty) in the particle's momentum, Δx is the spread in its position, and \hbar is the quantum constant. If we know the position precisely (that is, if the uncertainty Δx is tiny), then the momentum uncertainty Δp must be large in order to make the product $\Delta p \Delta x$ comparable to \hbar . In contrast, if we hardly know the position (Δx is large), then the momentum uncertainty Δp can be small. This relation is the physical contribution of quantum mechanics.

Let's apply it to the particle in the box. The particle could be anywhere in the box, so its position uncertainty Δx is comparable to the box width a :

$$\Delta x \sim a. \quad (6.58)$$

Because of the uncertainty principle, the consequence of confining the particle to the box implies a momentum uncertainty \hbar/a :

$$\Delta p \sim \frac{\hbar}{\Delta x} \sim \frac{\hbar}{a}. \quad (6.59)$$

This momentum corresponds to a kinetic energy $E \sim (\Delta p)^2/m$. As a consequence, the particle acquires an energy, called the confinement energy:

$$E \sim \frac{\hbar^2}{ma^2}. \quad (6.60)$$

This simple result enables us to estimate the radius of a neutron star: a star without any nuclear fusion (which normally resists gravity) and heavy enough that gravity has collapsed the protons and electrons into neutrons. The atomic structure is gone, and the star is one giant, neutral nucleus (except in the crust, where gravity is not strong enough to make neutrons).

Problem 6.31 **Dimensional analysis for the size of a neutron star**

Use dimensional analysis to find the radius R of a neutron star of mass M .

Two physical effects compete: Gravity tries to crush the star; and quantum mechanics, armed with the uncertainty principle, resists by penalizing confinement. To begin quantifying these effects, let the star have radius R and contain N neutrons, so it has mass $M = Nm_n$.

Gravity's contribution to crushing the star is reflected in the star's gravitational potential energy

$$E_{\text{potential}} \sim \frac{GM^2}{R}. \quad (6.61)$$

This energy is negative: Gravity is happier (the energy is lower) when R is smaller. Here, the minus sign is incorporated into the single approximation sign \sim .

For the other side of the competition, we confine each neutron to a box. To find the box width a , imagine the star as a three-dimensional cubic lattice of neutrons. Because the star contains N neutrons, the lattice is $N^{1/3} \times N^{1/3} \times N^{1/3}$. Each side has length comparable to R , so $aN^{1/3} \sim R$ and $a \sim R/N^{1/3}$.


Each neutron gets a confinement (kinetic) energy $\hbar^2/m_n a^2$. For all the N neutrons together, and using $a \sim R/N^{1/3}$,

$$E_{\text{kinetic}} \sim N \frac{\hbar^2}{m_n a^2} \sim N^{5/3} \frac{\hbar^2}{m_n R^2}. \quad (6.62)$$

In terms of the known masses M and m_n (instead of N), the total confinement energy is

$$E_{\text{kinetic}} \sim \frac{M^{5/3}}{m_n^{8/3}} \frac{\hbar^2}{R^2}. \quad (6.63)$$

The kinetic energy is proportional to R^{-2} , whereas the potential energy is proportional to R^{-1} . On log–log axes of energy versus radius, the kinetic energy is a straight line with a -2 slope, and the potential energy is a straight line with a -1 slope. Because the slopes differ, the lines cross. The radius where they cross, which is roughly the radius that minimizes the total energy, is also roughly the radius of the neutron star. (We used the same reasoning in Section 4.6.1, to find the speed that minimized the energy required to fly.)



With all the constants included, the crossing happens when R satisfies

$$\frac{GM^2}{R} \sim \frac{M^{5/3} \hbar^2}{m_n^{8/3} R^2}. \quad (6.64)$$

The radius of the neutron star is then

$$R \sim \frac{\hbar^2}{GM^{1/3} m_n^{8/3}}. \quad (6.65)$$

If the Sun, with a mass of 2×10^{30} kilograms, became a neutron star, it would have a radius of roughly 3 kilometers:

$$R \sim \frac{(10^{-34} \text{ kg m}^2 \text{ s}^{-1})^2}{7 \times 10^{-11} \text{ kg}^{-1} \text{ m}^3 \text{ s}^{-2} \times (2 \times 10^{30} \text{ kg})^{1/3} \times (1.6 \times 10^{-27} \text{ kg})^{8/3}} \quad (6.66)$$

$$\sim 3 \text{ km}.$$

The true neutron-star radius for the Sun, after including the complicated physics of a varying density and pressure, is approximately 10 kilometers. (In reality, the Sun's gravity wouldn't be strong enough to crush the electrons and protons into neutrons, and it would end up as a white dwarf rather than a neutron star. However, Sirius, the subject of Problem 6.32, is massive enough.)

Here is a friendly numerical form that incorporates the 10-kilometer information. It makes explicit the scaling relation between the radius and the mass, and it measures mass in the convenient units of the solar mass.

$$R \approx 10 \text{ km} \times \left(\frac{M_{\text{Sun}}}{M} \right)^{1/3}. \quad (6.67)$$

Our analysis, which used several lumping approximations and neglected many dimensionless constants throughout, predicted a prefactor of 3 kilometers instead of 10 kilometers. An error of a factor of 3 is a worthwhile

tradeoff. We avoid the complicated analysis of a quantum fluid of electrons, protons, and neutrons, and still gain insight into the essential physical idea: The size of a neutron star results from a competition between gravity and quantum mechanics. (You will find many insightful astrophysical scalings in an article on the “Astronomical reach of fundamental physics” [5].)

Problem 6.32 Sirius as a neutron star

Sirius, the brightest star in the night sky, has a mass of 4×10^{30} kilograms. What would be its radius as a neutron star?


6.5.2 Hydrogen by lumping

Having practiced quantum mechanics and lumping in the large (a neutron star), let’s go to the small: We’ll use lumping to complement the dimensional-analysis estimate of the size of hydrogen in Section 5.5.1.

Its size is the radius of the orbit with the lowest energy. Because the energy is a sum of the potential energy from the electrostatic attraction and the kinetic energy from quantum mechanics, hydrogen is analogous to a neutron star, where quantum mechanics competes instead against gravitation. Partly because electrostatics is much stronger than gravity (Problem 6.34), hydrogen is very tiny.

For now, let the hydrogen atom have an unknown radius r . As a lumping approximation, the complicated electrostatic potential becomes a box of width r . Then the electron’s momentum uncertainty is $\Delta p \sim \hbar/r$ and its confinement energy is $\hbar^2/m_e r^2$:

$$E_{\text{kinetic}} \sim \frac{(\Delta p)^2}{m_e} \sim \frac{\hbar^2}{m_e r^2}. \quad (6.68)$$



This energy competes against the electrostatic energy, whose estimate also requires a lumping approximation. For in quantum mechanics, an electron is not at any definite location. Depending on your interpretation of quantum mechanics, and speaking roughly, it is either smeared over the box, or it has a probability of being anywhere in the box. On either interpretation, calculating the potential energy requires an integral. However, we can approximate it by using lumping: The electron’s typical, or characteristic, distance from the proton is simply r , the radius of hydrogen (which is the size of the lumping box).

Then the potential energy is the electrostatic energy of a proton and electron separated by r :

$$E_{\text{potential}} \sim -\frac{e^2}{4\pi\epsilon_0 r}. \quad (6.69)$$

The total energy is their sum

$$E = E_{\text{potential}} + E_{\text{kinetic}} \sim -\frac{e^2}{4\pi\epsilon_0 r} + \frac{\hbar^2}{m_e r^2}. \quad (6.70)$$


The electron adjusts r to minimize this energy. As we learned in the analysis of lift (Section 4.6.1), when two terms with different scaling exponents compete, the minimum occurs where the terms are comparable. (In Section 8.3.2.2, we'll return to this example with the additional tool of easy cases and also sketch the energies on log-log axes.) Using the Bohr radius a_0 as the separation of minimum energy, comparability means

$$\frac{e^2}{4\pi\epsilon_0 a_0} \sim \frac{\hbar^2}{m_e a_0^2}. \quad (6.71)$$

The Bohr radius is therefore

$$a_0 \sim \frac{\hbar^2}{m_e (e^2/4\pi\epsilon_0)}, \quad (6.72)$$

as we found in Section 5.5.1 using dimensional analysis. Now we also get a physical model: The size of hydrogen results from competition between electrostatics and quantum mechanics.



Problem 6.33 Minimizing the energy in hydrogen using symmetry

Use symmetry to find the minimum-energy separation in hydrogen, where the total energy has the form

$$\frac{A}{r^2} - \frac{B}{r}. \quad (6.73)$$

(See Section 3.2.3 for related examples.)

Problem 6.34 Electrostatics versus gravitation


Estimate the ratio of the gravitational to the electrostatic force in hydrogen. What are the similarities and differences between this estimate and the estimate that you made in Problem 2.15?

6.6 Summary and further problems

Lumping is our first tool for discarding complexity with loss of information. By doing so, it simplifies complicated problems where our previous set of tools could not. Curves become straight lines, calculus becomes algebra, and even quantum mechanics becomes comprehensible.

Problem 6.35 Precession of the equinoxes

Because the Earth is oblate (Problem 5.46), and thus has a bit of fat near the equator, and because its rotation axis is tilted relative to the Earth's orbital plane, the Sun and the Moon each exert a slight torque on the Earth. This torque slowly rotates (precesses) the Earth's axis of rotation (so the north star will not always be the north star).



- Explain why F_1 and F_2 , the gravitational forces on the two bulges, are almost exactly equal in direction but not equal in magnitude.
- Use lumping to find the scaling exponents x and y in

$$\text{torque} \propto m^x l^y, \quad (6.74)$$

where m is the mass of the object (either the Sun or Moon) and l is its distance from the Earth. Then estimate the ratio

$$\frac{\text{torque from the Sun}}{\text{torque from the Moon}}. \quad (6.75)$$

- By including the constants of proportionality, estimate the total torque and the precession rate.

Problem 6.36 Graph lumping in reverse

In a constant-temperature (isothermal) atmosphere, which isn't a terrible approximation to the actual atmosphere, the air density falls exponentially with height:

$$\rho = \rho_0 e^{-z/H}, \quad (6.76)$$

where z is the height above sea level, ρ_0 is the density at sea level, and H is the scale height of the atmosphere. We estimated H in Section 5.4.1 using dimensional analysis; in this problem, you'll estimate it by using lumping in reverse.

- Sketch $\rho(z)$ versus z as given above. On the same graph, sketch a lumped $\rho(z)$ that represents a sea-level-density atmosphere for $z < H$ and zero density for $z \geq H$. Ensure that your lumping rectangle has the same area as the exponentially decaying $\rho(z)$ curve.
- Use your lumping rectangle and the sea-level pressure p_0 to estimate H . Then estimate the relative density at the top of Mount Everest (roughly 9 kilometers). Check your estimate by looking up the actual air density on Mount Everest.