Grade 10 Science

Study Guide

Formulas · worked examples · practice problems with the answer key

5 modules · 37 formulas · 74 practice problems

The Exam Room · formula.expert · metric edition 1

September 2026 · Mathieu Leclerc, with Claude

About this study guide

Who it's for. Students working through Grade 10 Science, and the teachers, tutors and parents printing handouts for them. Print it, copy it, staple it — it's free for classroom use, whole or by the page.

How to use it. Each module teaches its formulas in order — equation, picture, and what every symbol means — then closes with a practice set. Work the problems with a pencil; the answer key waits at the back so honest work gets an honest check. The two-page formula sheet up front is the exam-day companion, and every formula carries its number (Science 10-1 onward) so your written work can cite exactly which one you used.

Do it live. Every problem in here has a living version in the Exam Room at **formula.expert/exams/grade-10-science** — new numbers every attempt, hints that pop down when you want them, stars and streaks when you don't need them. When a printed problem stumps you, the matching lesson will coach you through it, step by step, free and with no sign-up. Other courses' study guides live at **formula.expert/study-guides**.

Contents

The practice problems are edition-drawn: a later edition deals fresh numbers, so keep the key with the printing it came from. The Exam Room deals new numbers every attempt.

Figures

The formula sheet

Module 1 · Motion — Speed, Distance, Time

Science 10-1 Speed, Distance & Time

$$v = \frac{d}{t}$$

Average speed from distance travelled and elapsed time.

Science 10-2 Displacement from Average Velocity

$$d = \frac{v_0 + v}{2} t$$

Displacement as the average of initial and final velocities multiplied by the elapsed time, valid...

Module 2 · Light & Geometric Optics

Science 10-3 Wave Speed (v = fλ)

$$v = f\lambda$$

The universal wave equation: a wave's speed equals its frequency times its wavelength.

Science 10-4 Period-Frequency Relation

$$T = \frac{1}{f}$$

Period and frequency are reciprocals: seconds per cycle versus cycles per second.

Science 10-5 Index of Refraction (n = c/v)

$$n = \frac{c}{v}$$

How much a medium slows light: the ratio of c to the speed of light in the medium.

Science 10-6 Snell's Law of Refraction

$$n_1 \sin\theta_1 = n_2 \sin\theta_2$$

Light bends at an interface so that n sin θ stays the same on both sides.

Science 10-7 Critical Angle for Total Internal Reflection

$$\sin\theta_c = \frac{n_2}{n_1}$$

Beyond this angle of incidence, light in the denser medium reflects totally instead of refracting.

Science 10-8 Focal Length of a Spherical Mirror

$$f = \frac{R}{2}$$

A spherical mirror focuses parallel rays at half its radius of curvature.

Science 10-9 Thin Lens Equation

$$\frac{1}{f} = \frac{1}{d_o} + \frac{1}{d_i}$$

Relates a thin lens's focal length to its object and image distances.

Science 10-10 Lens Magnification (m = −d_i/d_o)

$$m = -\frac{d_i}{d_o}$$

Image magnification from the ratio of image to object distance; the minus sign tracks inversion.

Science 10-11 Magnification from Heights (m = h_i/h_o)

$$m = \frac{h_i}{h_o}$$

Magnification as the ratio of image height to object height.

Science 10-12 Lens Power in Diopters

$$P = \frac{1}{f}$$

The optician's unit: lens power in diopters is the reciprocal of focal length in meters.

Module 3 · Heat & Temperature

Science 10-13 Sensible Heat (Q = mcΔT)

$$Q = mc\Delta T$$

Heat needed to change a mass's temperature: mass times specific heat times the temperature change.

Science 10-14 Latent Heat

$$Q = mL$$

Heat absorbed or released when a mass changes phase at constant temperature.

Science 10-15 Thermal Linear Expansion

$$\Delta L = \alpha L_0 \Delta T$$

Length change of a solid caused by a temperature change, via the linear expansion coefficient.

Science 10-16 Heat Conduction Rate

$$P = \frac{kA\Delta T}{d}$$

Steady-state heat flow through a slab by Fourier's law of conduction.

Module 4 · Electricity — Circuits & Power

Science 10-17 Electric Charge (Q = It)

$$Q = It$$

Total charge transferred by a steady current flowing for a given time.

Science 10-18 Ohm's Law

$$V = IR$$

Relates the voltage across a conductor to the current through it and its resistance.

Science 10-19 Two Resistors in Series

$$R_t = R_1 + R_2$$

Total resistance of two resistors connected end to end is simply their sum.

Science 10-20 Two Resistors in Parallel

$$R_t = \frac{R_1 R_2}{R_1 + R_2}$$

Total resistance of two resistors connected side by side: product over sum, always less than either...

Science 10-21 Electrical Power (P = VI)

$$P = VI$$

Power delivered to a component as the product of the voltage across it and the current through it.

Science 10-22 Electrical Power (P = I²R)

$$P = I^2 R$$

Power dissipated as heat in a resistance carrying a current (Joule heating).

Science 10-23 Electrical Energy (E = Pt)

$$E = Pt$$

Energy consumed by a device drawing constant power over a period of time.

Science 10-24 Energy Cost from a Utility Rate

$$C_e = E\, p_e$$

Cost of the energy a system consumes: kilowatt-hours or fuel BTUs times the utility rate, for tower...

Module 5 · Weather, Climate & the Atmosphere

Science 10-25 Pressure (P = F/A)

$$P = \frac{F}{A}$$

Pressure as perpendicular force spread over an area.

Science 10-26 Hydrostatic Pressure (P = ρgh)

$$P = \rho g h$$

Gauge pressure at depth h in a fluid of density ρ, using g = 9.80665 m/s².

Science 10-27 Barometric Pressure with Altitude

$$P = P_0\, e^{-Mgz/RT}$$

Atmospheric pressure at a height above a reference level, for an isothermal layer.

Science 10-28 Relative Humidity from Vapour Pressure

$$\varphi = \frac{p_v}{p_{ws}}$$

Relative humidity as what it actually is: the vapour pressure present divided by the saturation...

Science 10-29 Saturation Vapour Pressure (Magnus / Alduchov–Eskridge)

$$p_{ws} = 610.94 \exp\left(\frac{17.625\, t}{t + 243.04}\right)$$

The vapour pressure of water at saturation, from temperature alone, by the Alduchov–Eskridge Magnus...

Science 10-30 Dew Point (Magnus Approximation)

$$T_d = \frac{c\,\gamma}{b-\gamma}, \quad \gamma = \ln\frac{\mathrm{RH}}{100} + \frac{b\,T}{c+T}$$

Dew point from air temperature and relative humidity by the Magnus formula with the WMO's Sonntag...

Science 10-31 Humidex (Canadian Humidity Index)

$$H = T + 0.5555\,(e-10), \quad e = 6.11\,e^{5417.753\left(\frac{1}{273.16}-\frac{1}{T_d}\right)}$$

Environment Canada's humidex: air temperature raised by the excess vapour pressure computed from...

Science 10-32 Wind Chill (2001 North American Formula)

$$T_{wc} = 13.12 + 0.6215\,T - 11.37\,V^{0.16} + 0.3965\,T\,V^{0.16}$$

The wind chill index adopted by Canada and the United States in 2001, from air temperature and wind...

Science 10-33 Heat Index (Rothfusz Regression)

$$\mathrm{HI} = c_1 + c_2 T + c_3 R + c_4 T R + c_5 T^2 + c_6 R^2 + c_7 T^2 R + c_8 T R^2 + c_9 T^2 R^2$$

The US National Weather Service heat index, Rothfusz's nine-term regression through Steadman's...

Science 10-34 Albedo and Reflected Radiation

$$G_r = \alpha\,G$$

Shortwave radiation reflected by a surface, as its albedo times the incident irradiance.

Science 10-35 Environmental Lapse Rate

$$\Gamma = \frac{T_1 - T_2}{z_2 - z_1}$$

How fast the air cools with height, from two temperatures at two altitudes.

Science 10-36 Flash-to-Bang Distance to a Lightning Strike

$$d = (331.3 + 0.606\,T_C)\,t$$

Light arrives effectively instantly and sound does not, so the gap between the flash and the...

Science 10-37 Speed of Sound in Air

$$v = 331.3 + 0.606\,T_C$$

The speed of sound in dry air grows about 0.6 m/s for every degree Celsius above freezing.

MODULE 1

Motion — Speed, Distance, Time

Reading the Quantities · The Speed Formula · Calculate the Speed · Solve for Distance · Solve for Time · Speed in Real Units · Average Speed over a Whole Trip · The Highway Run

1.1 Speed, Distance & Time

(Science 10-1) $v = \frac{d}{t}$

Where,

v = Speed (m/s)

d = Distance (m)

t = Time (s)

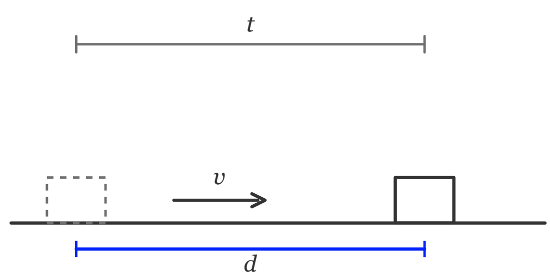

Figure 1.1 — speed, distance & time.

Speed is ground covered divided by time taken, $v = d/t$, and it is worth being clear about what kind of statement that is. It is not a law of nature that could turn out to be false — it is the *definition* of average speed. Nothing in it can be wrong; it can only be misapplied. What it gives you is the single steady speed that would have covered the same distance in the same time, which is a genuinely useful summary of a trip and tells you almost nothing about any particular moment within it.

Take a drive of 148 km that takes 1 hour 45 minutes. Convert the time to a single unit first — 1.75 h — and $v = 148/1.75 = 84.6$ km/h. In SI the same trip is 148 000 m over 6300 s, giving 23.5 m/s. Both answers describe the same drive, and the traffic light you sat at for ninety seconds is buried inside both of them.

This is the zeroth member of the kinematics family: set $a = 0$ in $d = v_0t + \frac{1}{2}at^2$ and you are left with $d = vt$. Run in the other direction, toward instantaneous speed, it becomes the derivative $v = \mathrm{d}d/\mathrm{d}t$, which is what your speedometer reads and what calculus was partly invented to handle. The relation also underwrites the modern definition of length itself: since 1983 the metre has been defined by *fixing* the speed of light at 299 792 458 m/s exactly, so distance is now measured by timing light and rearranging this formula for d.

The classic mistake is averaging the speeds instead of the trip. Drive 60 km out at 60 km/h and return at 30 km/h, and the average for the round trip is not 45 km/h. The outbound leg takes 1 h, the return 2 h, so it is 120 km in 3 h — 40 km/h. Average speed is always total distance over total time, never the mean of the individual speeds, because you spend longer at the slow one. The second trap is units: this formula does not convert anything, so 100 km and 30 minutes gives 3.33 in units of km/min, not km/h. Fix the units before the arithmetic, or use the unit selectors on this page. And note that this is *distance*, the ground covered, not displacement — a runner who finishes a 400 m lap where she started has run at a respectable speed and has an average velocity of exactly zero.

Worked example 1.1. Given d = 100 m, t = 8 s; find v (speed). *100 m in 8 s → 12.5 m/s.*

1.2 Displacement from Average Velocity

(Science 10-2) $$d = \frac{v_0 + v}{2}\, t$$

Where,

d = Displacement (m)
v_0 = Initial velocity (m/s)
v = Final velocity (m/s)
t = Time (s)

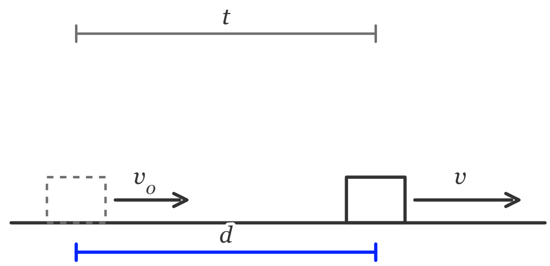

Figure 1.2 — displacement from average velocity.

Under constant acceleration, velocity changes along a straight line, so the average velocity over the interval is the plain arithmetic mean of the value at the start and the value at the end. Multiply that by the elapsed time and you have the displacement: $d = \frac{v_0+v}{2}\, t$. The reason this works is best seen on a velocity-versus-time graph, where displacement is the area underneath the curve. With constant acceleration the curve is a straight line, so the area is a trapezoid, and the area of a trapezoid is the mean of the two parallel sides times the width. The half in this formula is the half in the trapezoid rule.

A train easing from 30 m/s down to 10 m/s over 20 s covers $d = \frac{30+10}{2} \times 20 = 400$ m. Notice what you never needed: the acceleration. That is what makes this equation worth having as a separate page rather than treating it as a corollary — when the two speeds and the duration are what you have measured, it answers directly.

It is the bridge between the other kinematic equations rather than an independent fact. Substitute $v = v_0 + at$ into it and the algebra collapses to $d = v_0 t + \frac{1}{2}at^2$; substitute $v_0 = v - at$ instead and you recover $d = vt - \frac{1}{2}at^2$. All five SUVAT relations are the same two facts — velocity changes at a steady rate, displacement is the area under the velocity curve — rearranged to suit whichever variable is missing.

The failure mode is applying it when the acceleration was not constant, and it fails silently. Consider a car that sits at 10 m/s for 55 s and then accelerates hard to 30 m/s in the last 5 s. The starting velocity is 10, the final is 30, the time is 60 s, and this formula confidently reports 1200 m. The true distance is about 650 m. Nothing in the arithmetic warns you, because the shortcut assumes a straight line between the endpoints and the real motion was nothing of the kind. Whenever the acceleration varies, go back to the definition — total displacement over total time — or split the trip into segments where it genuinely is constant. One further distinction: this is average *velocity*, not average speed. If the motion reverses within the interval, the two are different numbers, and it is the velocity version that this formula computes.

Worked example 1.2. Given v_0 = 30 m/s, v = 10 m/s, t = 20 s; find d (displacement). *Train 30→10 m/s over 20 s → 400 m.*

Module 1 practice

1. Calculate the Speed — A cyclist covers 240 m of bike path at a steady pace, taking 40 s. *Calculate the average speed.*

2. Calculate the Speed — A cyclist covers 360 m of bike path at a steady pace, taking 40 s. *Calculate the average speed.*

3. Solve for Distance — A ferry travels in a straight line at a constant 4 m/s for 15 s. *Determine the distance travelled.*

4. Solve for Distance — A delivery drone travels in a straight line at a constant 10 m/s for 15 s. *Determine the distance travelled.*

5. Solve for Time — A rowing shell travels 100 m down the course at a constant 10 m/s. *Determine how long the journey takes.*

6. Solve for Time — A streetcar travels 300 m along its route at a constant 15 m/s. *Determine how long the journey takes.*

7. Speed in Real Units — A zip-line trolley hums along the cable at 25 m/s. *Express this speed in km/h.*

8. Speed in Real Units — A highway sign limits traffic to 36 km/h. *Express this speed in m/s.*

9. Average Speed over a Whole Trip — A delivery van spends 3 h at 50 km/h between towns, then 1 h at 70 km/h on rural routes. *Calculate the average speed for the whole trip.*

10. Average Speed over a Whole Trip — A family road trip runs in two legs: 3 h at 40 km/h on the highway, then 1 h at 60 km/h on smaller roads. *Calculate the average speed for the whole trip.*

11. The Highway Run — The summer highway run, mapped in three legs. Leg one: 150 km at 50 km/h. Leg two: 100 km at 100 km/h. Leg three: 60 km at 60 km/h. No calculator — every number is chosen to be friendly. Work each line; every answer feeds the next. *Determine each leg's time, then the total distance, then the whole trip's average speed.*

MODULE 2

Light & Geometric Optics

Light as a Wave · The Index of Refraction · Snell's Law · The Critical Angle · Curved Mirrors · The Thin Lens Equation · Magnification · Lens Power in Diopters · Boss — The Optics Bench

2.1 Wave Speed (v = fλ)

(Science 10-3) $v = f\lambda$

Where,

v = Wave speed (m/s)

f = Frequency (Hz)

λ = Wavelength (m)

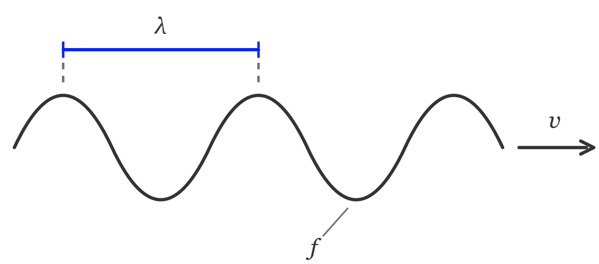

Figure 2.1 — wave speed (v = fλ).

Every traveling wave advances exactly one wavelength during each cycle of its source, and it completes *f* cycles every second — so its speed is simply frequency times wavelength. The relation holds for every wave in nature: sound, light, water ripples, seismic tremors. An FM station broadcasting at 100 MHz emits radio waves that travel at the speed of light, about 3.00×10^8 m/s, so each wave is roughly 3 m long — which is why FM antennas are built around three-quarters of a metre, a quarter of a wavelength.

The common trap is thinking a higher frequency makes a wave *faster*. It doesn't: speed is set by the medium alone. Sound in room-temperature air moves at about 343 m/s whether it is a 20 Hz bass rumble (λ ≈ 17 m) or a 20 kHz whistle (λ ≈ 17 mm). Raise the frequency and the wavelength shrinks in exact proportion, leaving **v** untouched.

Worked example 2.1. Given v = 1234.8 km/h, f = 0.5 kHz; find λ (wavelength). *Sound at 1234.8 km/h, 0.5 kHz → lambda = 0.686 m.*

2.2 Period-Frequency Relation

(Science 10-4) $$T = \frac{1}{f}$$

Where,

T = Period (s)

f = Frequency (Hz)

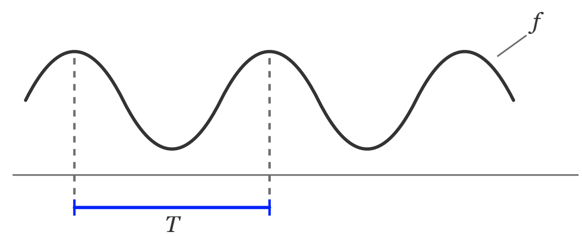

Figure 2.2 — period-frequency relation.

Period and frequency are one fact counted in opposite directions. The period T is seconds per cycle; the frequency f is cycles per second; and $T = 1/f$ is not a discovery about nature but unit algebra. If something completes four cycles in a second, each cycle takes a quarter of a second, and no experiment was required to establish that. What earns this relation a page of its own is that almost every oscillation formula you will meet returns one member of the pair while the question in front of you wants the other, so this conversion sits quietly in the middle of nearly every wave calculation.

North American mains alternates at 60 Hz, so one full cycle takes $1/60 = 16.67$ ms — the interval behind the familiar hum in audio gear. European mains at 50 Hz gives 20 ms. Concert A at 440 Hz gives 2.27 ms per cycle. Run it the other way and a resting heart beating once every 0.8 s is oscillating at 1.25 Hz, which is the same statement as 75 beats per minute.

The unit is younger than the idea. Frequency was written "cycles per second" well into the twentieth century; the International Electrotechnical Commission proposed **hertz** in 1930 and SI adopted it in 1960, honouring Heinrich Hertz, who between 1886 and 1888 generated and detected radio waves in his Karlsruhe laboratory and so turned Maxwell's equations from mathematics into an observed fact. The hertz is dimensionally just s^{-1}, and it is reserved by convention for periodic phenomena — the becquerel is also s^{-1} and counts random decays, which is precisely why the two units are kept apart despite being numerically identical. Combined with $v = f\lambda$, this page also gives the equally useful $v = \lambda/T$.

Three traps. The first is **angular** frequency: $\omega = 2\pi f$ in radians per second, and the pendulum and spring formulas carry their 2π for exactly this reason. Substituting an ω where an f belongs makes the answer wrong by a factor of 6.283, which is large enough to notice and small enough to rationalise. The second is revolutions per minute: 3600 rpm is 60 Hz, not 3600, and the conversion is a division by 60. The third catches people with pendulums — a "seconds pendulum" ticks once per second but has a period of *two* seconds, because a full cycle is out and back. Count a cycle as a return to the starting state moving in the starting direction, and the reciprocal will behave.

Worked example 2.2. Given f = 60 Hz; find T (period). *60 Hz mains → T = 1/60 s.*

2.3 Index of Refraction (n = c/v)

(Science 10-5) $$n = \frac{c}{v}$$

Where,

n = Index of refraction

v = Speed of light in the medium (m/s)

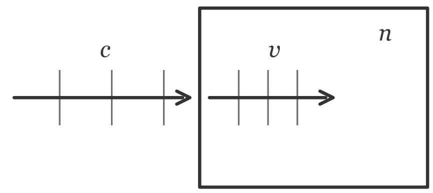

Figure 2.3 — index of refraction (n = c/v).

The refractive index is a speed ratio and nothing more: the speed of light in vacuum divided by its speed in the material, with $c = 299\,792\,458$ m/s exact by definition. Because nothing outruns light in vacuum, $n \geq 1$ for ordinary transparent materials. Every refraction effect you will meet descends from this one slowing — the bending at an interface, total internal reflection, the way a prism spreads colours, the shallowness of a pool. Snell's law, the critical angle and apparent depth are all downstream of $n = c/v$, which is a good reason to be clear about what it means before using it.

Water at $n = 1.333$ carries light at 2.25×10^8 m/s; ordinary crown glass at 1.52 gives 1.97×10^8; diamond at 2.417 slows it to 1.24×10^8 m/s. Run it forward on a number people actually use: a single-mode fibre core has $n = 1.4682$, so $v = 2.042 \times 10^8$ m/s, which is 4.90 µs per kilometre. A 1000 km link therefore cannot have a round-trip latency below about 9.8 ms no matter how good the electronics get — a hard floor set by this equation, and a number every network engineer eventually learns the hard way.

It is worth being careful about what "slowing" means, because the usual telling is misleading. Individual photons always travel at c; there is no medium in which light itself is sluggish. What happens is that the passing electromagnetic field drives the electrons in the material, those electrons re-radiate, and the superposition of the original wave with all the re-radiated wavelets is a wave whose crests advance more slowly than c. The bulk speed c/n is a property of that superposition, not of any individual photon. This also explains why n depends on wavelength: the electrons respond more strongly near their resonances, so blue is slowed more than red. That is dispersion, it is why a prism works, and it is why "$n = 1.52$ for glass" is shorthand for the value at the sodium D line at 589 nm. BK7 crown is 1.5168 there and 1.5224 in the blue at 486 nm.

Four things to watch. The n in this equation is the **phase** index, and c/n is the phase velocity. In a region of strong dispersion the phase velocity can genuinely exceed c, and engineered materials with n below 1 exist; no information travels faster than c, because signals travel at the group velocity, and this distinction is garbled in popular accounts often enough to be worth stating. Second, quoting a single index without naming a wavelength is imprecise, and it matters for anything involving colour. Third, n also varies with temperature and, for gases, with pressure — air is 1.000293 at standard conditions, which is why we quietly treat air as vacuum in most problems, but that tiny residual is exactly what produces road mirages and the twinkling of stars. Finally, Snell's law needs only the *ratio* of two indices, so when both media are given you never need c at all; this page is for the cases where the speed in the material is what you actually want.

Worked example 2.3. Given n = 150 %; find v (speed of light in the medium). *n = 150 % (1.5) → v = c/1.5 = 1.9986e8 m/s.*

2.4 Snell's Law of Refraction

(Science 10-6) $$n_1 \sin\theta_1 = n_2 \sin\theta_2$$

Where,

n_1 = Index of refraction (medium 1)
θ_1 = Angle of incidence (°)
n_2 = Index of refraction (medium 2)
θ_2 = Angle of refraction (°)

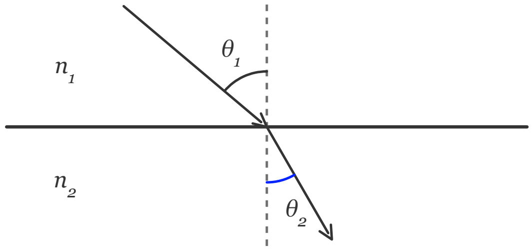

Figure 2.4 — snell's law of refraction.

Light crossing into a denser medium slows down, and to keep its wavefronts connected it must bend toward the normal — by exactly the amount that keeps n sin θ constant. A ray entering water (n = 1.333) from air at 45° refracts to arcsin(sin 45°/1.333) ≈ 32°, which is why a pool's floor looks shallower than it is and a straw appears kinked at the surface.

Run the light the other way, from dense to thin, and Snell's law eventually fails to give an answer: past the critical angle the sine would exceed 1 and the ray reflects totally instead — the trick that traps light inside optical fibers.

Worked example 2.4. Given n_1 = 1, θ_1 = 30 °, n_2 = 1.5; find θ_2 (angle of refraction). *Air to glass at 30 deg → theta2 = arcsin(1/3) = 19.4712 deg.*

2.5 Critical Angle for Total Internal Reflection

(Science 10-7) $$\sin\theta_c = \frac{n_2}{n_1}$$

Where,

θ_c = Critical angle (°)
n_1 = Index of the denser medium
n_2 = Index of the outer medium

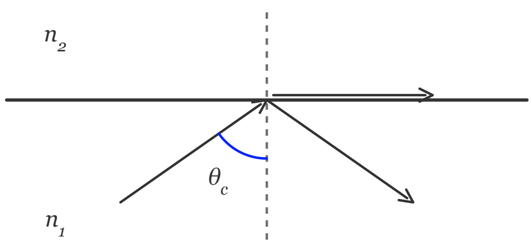

Figure 2.5 — critical angle for total internal reflection.

Take Snell's law, $n_1 \sin\theta_1 = n_2 \sin\theta_2$, with the light starting in the denser medium, and increase the angle of incidence. The refracted ray bends further and further from the normal until it is skimming along the surface at $\theta_2 = 90°$, where $\sin\theta_2 = 1$ and the relation reduces to $\sin\theta_c = n_2/n_1$. Push past that and there is simply no solution — the sine of the refracted angle would have to exceed 1 — so no light refracts at all and every bit of it reflects back inside. The word "total" is precise here in a way it never is for a mirror: a silvered surface loses a few per cent on every bounce, and total internal reflection loses nothing.

Glass to air, with $n_1 = 1.5$ and $n_2 = 1.0$, gives $\theta_c = \arcsin(0.667) = 41.8°$. That it falls below 45° is the reason a plain 45–45–90 glass prism reflects perfectly with no coating, and why binocular prisms, SLR pentaprisms and corner-cube retroreflectors are made the way they are. Water to air gives $\arcsin(1/1.333) = 48.6°$. Diamond gives $\arcsin(1/2.417) = 24.4°$, a remarkably small angle, so light entering a brilliant cut bounces repeatedly among the facets before it can escape — which, together with diamond's strong dispersion, is what a gemmologist means by fire.

Optical fibre is this equation as an industry. A step-index single-mode fibre has a core at $n = 1.4682$ and a cladding at $n = 1.4629$, a difference of about a third of one per cent, giving $\theta_c = \arcsin(1.4629/1.4682) = 85.1^\circ$. Measured from the normal, that means any ray travelling within 4.9° of the fibre axis is trapped, and it stays trapped for tens of kilometres between amplifiers. Note that the reflection happens at the core–cladding boundary, not at the glass–air surface: the cladding is what lets the fibre keep working when it is bundled, buried, bent and handled.

Three errors, in descending order of how often I see them. **Order matters: n_1 must be the denser medium, the one the light is already in.** If $n_2 > n_1$ there is no critical angle at all — light going from air into glass always refracts and never totally reflects — and asking for one is asking for the arcsine of a number greater than 1. The page refuses, and it is right to. Second, all angles here are measured *from the normal*, not from the surface, and a ray described as being "at 5° to the fibre axis" is at 85° to the normal. A good half of all critical-angle mistakes are this one substitution. Third, "no light escapes" is true of the propagating wave but not of the field: an evanescent wave extends about a wavelength beyond the surface, decaying exponentially, and it carries no energy away — unless you bring a second piece of glass within that distance, in which case light tunnels across the gap. Frustrated total internal reflection is a real effect with real products behind it, including optical fingerprint scanners and some beam splitters.

Worked example 2.5. Given n_1 = 2, n_2 = 1; find θ_c (critical angle). *n1 = 2, n2 = 1 → theta_c = 30 deg.*

2.6 Focal Length of a Spherical Mirror

(Science 10-8) $$f = \frac{R}{2}$$

Where,
f = Focal length (m)
R = Radius of curvature (m)

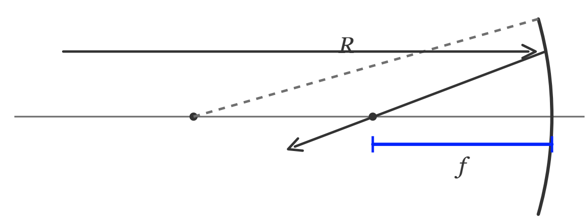

Figure 2.6 — focal length of a spherical mirror.

Rays arriving parallel to the axis of a spherical mirror cross the axis at half the radius of curvature. The proof is three lines. Take a ray parallel to the axis striking the mirror at some height; the normal at that point is the radius, pointing back to the centre of curvature C. The law of reflection sends the ray away at the same angle to that normal, and the triangle formed by the strike point, C, and the place where the reflected ray crosses the axis has two equal angles — so it is isosceles, and the crossing point sits exactly halfway between the mirror and C. That "exactly" holds only while the angles are small, and everything interesting about mirror design lives in the failure of that assumption.

Grind a telescope blank to a 4 m radius of curvature and you have a 2 m focal length. Run it the other way for a real specification: an f/5 telescope with a 200 mm aperture needs $f = 1000$ mm, so $R = 2000$ mm — and R is what the mirror maker actually measures during figuring, by finding the point where the mirror images a pinhole back onto itself. A domestic case: a concave shaving mirror with $R = 1$ m has $f = 0.5$ m, so a face held closer than 0.5 m gets an upright, enlarged, virtual image, and a face held farther away suddenly appears upside down. Most people have noticed that flip without knowing they were crossing a focal point.

The small-angle proviso is the origin of spherical aberration. Rays striking farther out on a sphere cross the axis slightly closer to the mirror, so the focus is a smear rather than a point, and the smear grows as the fourth power of the aperture ratio. Newton's first reflecting telescope of 1668 used a spherical mirror and was small enough that this hardly mattered. Modern primaries are figured into paraboloids, which bring all parallel axial rays to one exact point — at the price of coma for anything off-axis, which is the trade the Schmidt and Ritchey–Chrétien designs exist to manage. Worth keeping in proportion: the deviation between the sphere and the required paraboloid on a large mirror is only a few wavelengths of light, and removing those few wavelengths is the hardest and slowest part of making one.

The sign convention is what people get wrong. A concave mirror has its centre of curvature in front of it, so R and f are positive and it can form real images. A convex mirror — the passenger-side wing mirror, the security dome in a shop — has its centre of curvature *behind* the reflecting surface, so R and f are negative, and it can only ever form upright, reduced, virtual images. That negative focal length is the whole content of "objects in mirror are closer than they appear": the reduced image reads to the eye as a more distant car. Enter a convex mirror's radius as a positive number and every answer downstream is wrong. Two further cautions: $f = R/2$ is a mirror relation only — a lens with the same surface radii has a focal length given by the lensmaker's equation and depends on the refractive index, and substituting $R/2$ there is a common and badly wrong shortcut. And R is the radius of the sphere the surface is a cap of, not the radius of the mirror's rim.

Worked example 2.6. Given R = 4 m; find f (focal length). *R = 4 m → f = 2 m.*

2.7 Thin Lens Equation

(Science 10-9) $$\frac{1}{f} = \frac{1}{d_o} + \frac{1}{d_i}$$

Where,

- f = Focal length (m)
- d_o = Object distance (m)
- d_i = Image distance (m)

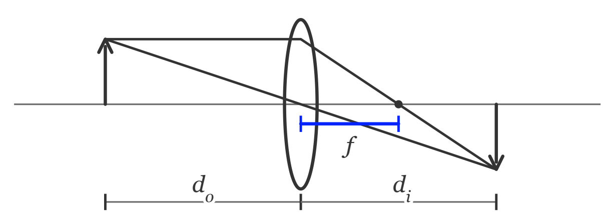

Figure 2.7 — thin lens equation.

The thin lens equation is the workhorse of geometric optics, tying together where the object sits, where its image forms, and the lens's focal length. The sign conventions carry the physics: a positive d_i is a real image on the far side of a converging lens, a negative d_i a virtual image on the near side. Place an object exactly at the focal point and $1/d_i$ must vanish — the image recedes to infinity, which is how a magnifying glass produces parallel rays.

A camera focused on a subject 2 m away with a 50 mm lens forms its image at $d_i = (1/0.05 - 1/2)^{-1} \approx 51.3$ mm behind the glass — the millimeter of travel that the focus ring provides.

Worked example 2.7. Given d_o = 30 cm, d_i = 15 cm; find f (focal length). *d_o = 30 cm, d_i = 15 cm → f = 10 cm.*

2.8 Lens Magnification (m = −d_i/d_o)

(Science 10-10) $$m = -\frac{d_i}{d_o}$$

Where,
m = Magnification
d_i = Image distance (m)
d_o = Object distance (m)

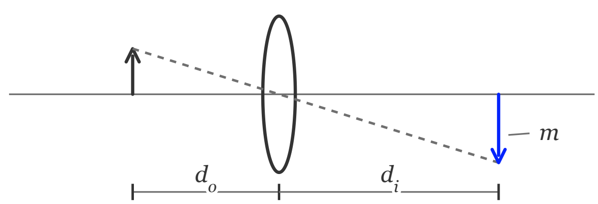

Figure 2.8 — lens magnification (m = −d_i/d_o).

This one falls out of similar triangles, and it is worth seeing why rather than accepting it. The ray that passes through the exact centre of a thin lens goes straight through undeviated, because at the centre the two surfaces are parallel. That single ray makes the same angle with the axis on both sides, so the triangle formed by the object and the triangle formed by the image are similar. Their proportions give $|h_i|/|h_o| = d_i/d_o$ — the image is as many times larger as it is farther away. The minus sign in $m = -d_i/d_o$ is a bookkeeping convention laid on top of that geometry, and it encodes orientation: under the standard convention a real image from a converging lens has both distances positive, so m comes out negative, and negative means inverted.

A camera with a 50 mm lens focused on a subject 2 m away forms its image at $d_i = 51.3$ mm, so $m = -51.3/2000 = -0.0256$. A person 1.8 m tall images 46 mm tall, which will not fit on a full-frame sensor only 24 mm high — so you step back, or fit a wider lens. A projector runs the same equation the other way. Put a 24 mm slide 102 mm in front of a 100 mm lens and the thin-lens equation gives $d_i = 5.1$ m, so $m = -50$: an image 1.2 m across and upside down. That inversion is why slides go into a carousel the wrong way up.

The two magnification pages are meant to be used together. This one gets you from distances to a ratio; $m = h_i/h_o$ gets you from that ratio to a size. A case worth memorising sits between them: when $d_o = 2f$ the thin-lens equation gives $d_i = 2f$ as well, so $m = -1$ exactly — object and image the same size, symmetric about the lens, and separated by $4f$. That is the 1:1 setting macro photographers work at, and $4f$ is the smallest object-to-sensor distance at which a given lens can form a real image at all. Try to squeeze the two closer and no solution exists.

The sign is the trap, and sign conventions are the single largest source of wrong answers in geometric optics. A negative m does not mean a negative size or a reduced image — it means *inverted*. Size lives in the magnitude, orientation lives in the sign, and $m = -0.03$ is telling you two separate things at once: strongly reduced, and upside down. This page uses the standard "real is positive" convention: d_o positive for a real object, d_i positive for a real image on the far side of the lens, negative for a virtual image on the near side. Hold a magnifying glass closer than its focal length and d_i goes negative, so m goes positive: upright and enlarged, exactly what your eye reports. Pick one convention and never borrow a formula from a textbook using another. One last distinction: this is the *linear* magnification of a real optical image. The "10×" stamped on a hand lens and the "40×" on a microscope objective are angular magnifications defined against a 25 cm reference viewing distance, a different quantity that cannot be compared with this one directly.

Worked example 2.8. Given d_i = 15 cm, d_o = 30 cm; find m (magnification). *d_i = 15 cm, d_o = 30 cm → m = -0.5.*

2.9 Magnification from Heights (m = h_i/h_o)

(Science 10-11) $$m = \frac{h_i}{h_o}$$

Where,
m = Magnification
h_i = Image height (m)
h_o = Object height (m)

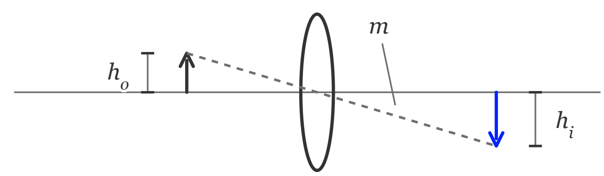

Figure 2.9 — magnification from heights (m = h_i/h_o).

This is magnification in its most direct form: how tall the image is compared with the thing it is an image of. Everything else said about magnification is derived from this ratio, including the $-d_i/d_o$ form on the companion page, which is really just this quantity re-expressed in terms of distances via similar triangles. Because it is a ratio of two lengths it is dimensionless, and the only unit discipline required is that both heights be measured in the same unit — millimetres over millimetres, micrometres over micrometres, it does not matter which so long as they match.

A microscope objective marked 40× projects a 5 µm red blood cell as a 200 µm image for the eyepiece to work on. Run it backwards for a photographic problem: a full-frame sensor is 24 mm tall, so to fit a person 1.8 m tall into the frame you need $|m| = 24/1800 = 0.0133$, and the companion page then tells you which lens and distance deliver it. Or take the Moon, 3474 km across at a distance of 384 400 km: a 500 mm lens images it $500 \times 3474/384400 = 4.5$ mm across, which is why lunar photography needs a very long lens to fill any part of a frame.

The two magnification relations are designed to be chained. Set them equal and you get $h_i/h_o = -d_i/d_o$, which is the practical heart of thin-lens problem solving: the thin-lens equation hands you d_i, the ratio form converts that into a magnification, and this page converts the magnification into a physical size you can compare against a sensor, a screen or a detector. Almost every lens problem you will meet is those three steps in some order.

The sign convention lives in h_i, and this is where it bites. Heights measured above the optical axis are positive and heights below it are negative, so an inverted image has a genuinely *negative* h_i — not a height smaller than nothing, but a height measured downward from the axis. The usual mistake is to enter the physical size of an inverted image as a positive number, which produces a magnification of the right magnitude and the wrong sign, and then to carry that error into whatever comes next. The object height h_o is conventionally taken as positive, with the object upright; making it negative inverts your own reference frame and flips everything. And keep this linear magnification separate from the other things called magnification: a telescope's or microscope's angular magnification is a ratio of apparent angles, not of heights, and the "3×" on a zoom lens is a ratio of focal lengths and is not a magnification at all.

Worked example 2.9. Given h_i = 4 cm, h_o = 2 cm; find m (magnification). *h_i = 4 cm, h_o = 2 cm → m = 2.*

2.10 Lens Power in Diopters

(Science 10-12) $$P = \frac{1}{f}$$

Where,

P = Optical power (dpt)

f = Focal length (m)

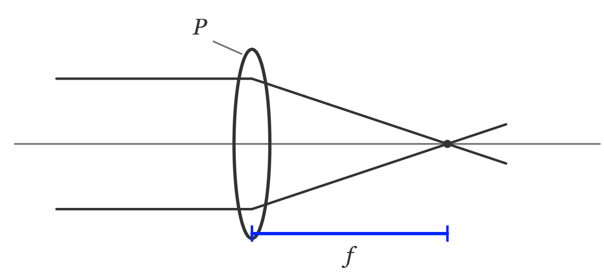

Figure 2.10 — lens power in diopters.

Optical power in dioptres is just the reciprocal of the focal length in metres, and the reason opticians prefer it to focal length is arithmetic convenience of a very practical kind. Focal lengths of thin lenses in contact combine as $1/f = 1/f_1 + 1/f_2$, which nobody wants to do at a fitting bench. Powers simply add: a $+2.0$ D reader over a $+0.5$ D correction is $+2.5$ D, and you have done it in your head. The whole of clinical refraction is built on that single property, and the unit is named for it — a dioptre is a reciprocal metre wearing a clinical hat.

A -4.0 D prescription for myopia is a diverging lens with $f = -0.25$ m, or -25 cm. A $+2.5$ D reader has $f = +40$ cm, which is roughly where you would hold a book. The relaxed human eye is about a $+60$ D system altogether, and the split is instructive: roughly $+43$ D comes from the cornea, where air meets tear film, and only about $+19$ D from the crystalline lens the muscles can adjust. That is why refractive surgery reshapes the cornea — it is where nearly three quarters of the power lives. Accommodation adds up to about $+14$ D of extra power in a child and declines steadily to nearly nothing by the sixties, which is presbyopia expressed as one number.

The reciprocal-metre form pays off a second time in what optometrists call vergence. The reciprocal of a distance in metres is the vergence of the light arriving from or heading to that point, so the thin-lens equation itself becomes plain addition in dioptres: incoming vergence plus lens power equals outgoing vergence. Light from an object 50 cm away arrives with -2 D of vergence; add a $+6$ D lens and it leaves with $+4$ D, focusing 25 cm beyond. No reciprocals, no common denominators. This is how the arithmetic is actually done in a clinic, and it is worth knowing even if you never sit in that chair.

Four traps. **Metres, always metres.** A 50 mm camera lens is $+20$ D, not $+0.02$ D, and entering millimetres puts the answer out by a factor of 1000 — the commonest error on this page by a wide margin. Second, the sign carries the whole clinical meaning: converging positive, diverging negative, so a myopia correction is negative and a hyperopia or reading correction is positive. Dropping a minus sign converts one condition into its opposite. Third, the additive rule requires the lenses to be genuinely in contact, or separated by a distance negligible compared with their focal lengths. Separate them by d and the combination becomes $P = P_1 + P_2 - dP_1P_2$, which is exactly why a spectacle lens sitting 12 mm in front of the eye is not the same prescription as a contact lens on the cornea — above about ± 4 D the vertex-distance correction is clinically significant and is routinely calculated. Fourth, the cylinder and axis figures on a prescription describe a separate astigmatic correction in a particular meridian, and this single spherical number says nothing about them.

Worked example 2.10. Given f = 0.5 m; find P (optical power). *f = 0.5 m → P = 2 dpt.*

Module 2 practice

12. Light as a Wave — A student shakes a long skipping rope at 2 Hz, sending waves of wavelength 4 m down its length. *Calculate the speed of the waves along the rope.*

13. Light as a Wave — A wave machine drives waves across its tank at 15 m/s, at a frequency of 3 Hz. *Determine the wavelength of the waves.*

14. The Index of Refraction — Inside a rectangular glass block, light is measured travelling at 2.0 ×10^8 m/s. *Calculate the index of refraction of the glass.*

15. The Index of Refraction — A clear plastic block has an index of refraction of 1.2. *Determine the speed of light inside it.*

16. Snell's Law — A ray of light in air (n = 1.00) strikes the surface of a calm pool at 40° from the normal. The index of refraction of the water is 1.33. *Determine the angle of refraction.*

17. Snell's Law — A ray of light in air (n = 1.00) strikes a glass block at 30° from the normal. The index of refraction of the glass is 1.5. *Determine the angle of refraction.*

18. The Critical Angle — A ray of light travelling inside a zircon gem (n = 1.9) heads up toward the surface, with air (n = 1.00) beyond. *Calculate the critical angle for total internal reflection.*

19. The Critical Angle — A ray of light travelling inside a glass block (n = 1.5) heads up toward the surface, with air (n = 1.00) beyond. *Calculate the critical angle for total internal reflection.*

20. Curved Mirrors — A solar-cooker mirror is a slice of a sphere with a radius of curvature of 90 cm. *Determine the focal length of the mirror.*

21. Curved Mirrors — A telescope's concave mirror has a focal length of 15 cm. *Determine the radius of curvature of the sphere it was ground from.*

22. The Thin Lens Equation — A converging lens of focal length 6 cm throws a sharp image onto a screen 15 cm behind it. *Determine how far the object stands from the lens.*

23. The Thin Lens Equation — A converging lens of focal length 10 cm throws a sharp image onto a screen 60 cm behind it. *Determine how far the object stands from the lens.*

24. Magnification — A converging lens forms a sharp image of a candle on a screen. The candle stands 25 cm from the lens, and the image forms 10 cm beyond it. *Calculate the magnification, sign included.*

25. Magnification — A document camera magnifies 5 times. A postage stamp 3 cm tall sits under it. *Determine the height of the image on the classroom screen.*

26. Lens Power in Diopters — A lens on an optometrist's bench has a focal length of 20 cm. *Calculate the power of the lens in diopters.*

27. Lens Power in Diopters — A pair of reading glasses is labelled 2.5 D. *Determine the focal length of the lenses, in centimetres.*

28. Boss — The Optics Bench — Last experiment of the term. A candle burns 30 cm from a converging lens, and its sharp image lands on a screen 60 cm past the lens. The flame is 8 cm tall. (No calculator — every reciprocal here is a kind one.) Work each line — every answer feeds the next. *Determine the focal length, then follow the candle as it moves — one line at a time.*

29. Boss — The Optics Bench — Bonus mark, for the label drawer: bench lenses are filed by power, and this one has a focal length of 25 cm. (A metre is 100 cm — still no calculator.) *Determine the power to write on the lens's sticker, in diopters.*

MODULE 3

Heat & Temperature

Heat Is Not Temperature · Specific Heat Capacity · Rearranging Q = mcΔT · Latent Heat of Phase Change · Climbing the Heating Curve · Thermal Expansion · Heat Flowing Through a Wall · The Kettle Problem

3.1 Sensible Heat (Q = mcΔT)

(Science 10-13) $$Q = mc\Delta T$$

Where,

- Q = Heat energy (J)
- m = Mass (kg)
- c_p = Specific heat capacity (J/(kg·K))
- ΔT = Temperature change (C°)

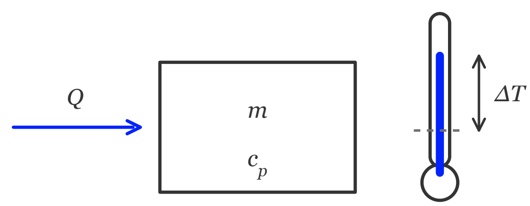

Figure 3.1 — sensible heat (q = mcδt).

Sensible heat is the energy that changes a substance's temperature without changing its phase. The specific heat capacity *c* is the price of each degree: how many joules one kilogram demands per kelvin of warming. Water's is famously steep at about 4186 J/(kg·K), which is why oceans moderate coastal climates and why a kettle takes its time. Heating 1.5 kg of water from 15 °C to 95 °C costs Q = 1.5 × 4186 × 80 ≈ 502 kJ.

The concept dates to Joseph Black's calorimetry experiments in 1760s Glasgow, which first pried apart the ideas of temperature and heat. Note that ΔT is a temperature *difference*, so a change of 80 °C equals a change of 80 K exactly — Fahrenheit differences convert by scale alone, with no offset. The formula holds as long as *c* stays roughly constant over the range and nothing melts or boils along the way.

Worked example 3.1. Given m = 2 kg, c_p = 4186 J/(kg·K), ΔT = 30 C°; find Q (heat energy). *2 kg water, c = 4186, ΔT = 30 C° → 251160 J.*

3.2 Latent Heat

(Science 10-14) $$Q = mL$$

Where,
Q = Heat absorbed or released (J)
m = Mass changing phase (kg)
L = Specific latent heat (J/kg)

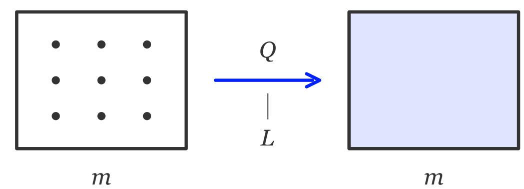

Figure 3.2 — latent heat.

Latent heat is the energy a phase change absorbs or releases while the temperature holds still. Melting 1 kg of ice at 0 °C soaks up 334 kJ — enough to heat that same water from 0 °C to 80 °C — yet the thermometer never moves until the last crystal is gone. Boiling is costlier still: vaporizing a kilogram of water takes about 2256 kJ, more than five times the energy needed to warm it from ice-cold to boiling.

Joseph Black coined the term in 1762 — *latent* means hidden, because the heat disappears into the phase change instead of the temperature reading. The physics runs everyday life: sweat cools you as it evaporates, steam scalds far worse than boiling water because it dumps its latent heat on condensing against skin, and every refrigerator moves heat by evaporating and condensing a working fluid in an endless loop.

One quantity, three notations, depending on whose book you are holding. Physics writes it L, as here, and splits it into L_f for fusion and L_v for vaporization. Engineering and every steam table print it h_{fg}, where f is saturated fluid and g is saturated gas — so $h_{fg} = h_g - h_f$ is the gap between the two columns, which is exactly why it shrinks to nothing at the critical point where those columns meet. Chemistry writes it as an enthalpy of vaporization per MOLE rather than per kilogram, so its numbers look nothing like these until you divide by the molar mass. Same energy, three addresses.

Worked example 3.2. Given m = 1 kg, L = 334 kJ/kg; find Q (heat absorbed or released). *Melting 1 kg ice at 334 kJ/kg → 334000 J.*

3.3 Thermal Linear Expansion

(Science 10-15) $$\Delta L = \alpha L_0 \Delta T$$

Where,
ΔL = Change in length (m)
α = Linear expansion coefficient (1/K)
L_0 = Original length (m)
ΔT = Temperature change (C°)

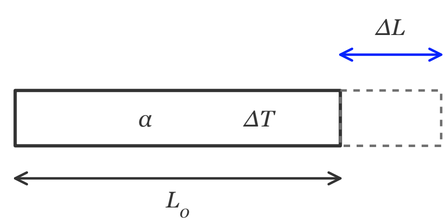

Figure 3.3 — thermal linear expansion.

Nearly every solid grows when heated, and it grows by a fixed *fraction* of whatever length it already had. That is the whole content of $\Delta L = \alpha L_0 \Delta T$: the coefficient α is the fractional growth per degree, so a long member moves more than a short one made of the same stuff. The reason solids expand at all is a detail of the interatomic bond. The potential well an atom sits in is not symmetric — pushing two atoms together costs more energy than pulling them apart by the same distance — so as the atoms vibrate harder, their *average* separation drifts outward. Expansion is that asymmetry, summed over every bond in the piece.

Take a 100 m copper riser in a hydronic building, filled at 20 °C and run at 82 °C. With α = 16.5×10^{-6} per K, $\Delta L = 16.5 \times 10^{-6} \times 100 \times 62 = 0.102\ \mathrm{m}$. The pipe wants to be **102 mm longer** than it was when it was installed — a hand's width, in a shaft where the anchors are bolted to concrete. This is why risers get expansion loops, offsets, or bellows, and why the guides that keep the pipe pointing straight matter as much as the loops themselves.

The same coefficient handles areas and volumes to good approximation: an area grows at about 2α and a volume at about 3α, because the fractional growth applies in each dimension independently. Bond two metals with different α back to back and the strip curves as it warms, which was the thermostat for most of the twentieth century. Go the other way and you get Invar, a nickel–iron alloy with α near 1.2×10^{-6} — a tenth of steel's — for which Charles Édouard Guillaume took the 1920 Nobel Prize in Physics, an award for a material rather than a discovery, given because surveying and clockmaking needed lengths that did not care about the weather.

The largest mistake is applying this equation to a member that is not free to move. A pipe anchored at both ends does not get longer; it develops stress instead. The stress is $\sigma = E\alpha\Delta T$, and for steel a 50 K rise gives $210 \times 10^9 \times 12 \times 10^{-6} \times 50 = 126\ \mathrm{MPa}$ — a serious fraction of the yield strength, and completely independent of length. That last point catches people: a short restrained member is under exactly the same stress as a long one, so shortening a run does not relieve anything. Continuous welded rail is laid under deliberate pre-tension for this reason, and a bridge without working expansion joints does not stretch, it buckles.

Two unit traps and one geometric one. ΔT is a *difference*, so it carries the same number in kelvin and in Celsius — but not in Fahrenheit, where a 50 °F change is 27.8 K, and where a coefficient tabulated "per °F" is five-ninths of the per-kelvin value. Never substitute an absolute temperature for ΔT. And α itself is not constant over a wide range; handbook values are averages over a stated interval, and they drift at cryogenic and high temperatures. The geometric one: a hole in a heated plate gets *larger*, not smaller. Every dimension scales by the same fraction, including the empty ones, which is why warming a jar lid loosens it.

Worked example 3.3. Given α = 12e-6 1/K, L_0 = 10 m, ΔT = 50 K; find ΔL (change in length). *10 m steel beam (α = 12e-6/K) heated 50 K grows 6 mm.*

3.4 Heat Conduction Rate

(Science 10-16) $$P = \frac{kA\Delta T}{d}$$

Where,

- P = Heat flow rate (W)
- k = Thermal conductivity (W/(m·K))
- A = Cross-sectional area (m²)
- ΔT = Temperature difference (C°)
- d = Thickness (mm)

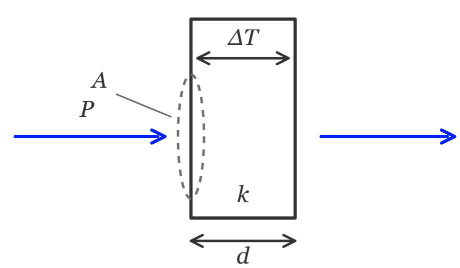

Figure 3.4 — heat conduction rate.

Fourier's law says heat flows down a temperature gradient at a rate proportional to how steep that gradient is. Written for a flat slab, the gradient is ΔT/d and the flow is $P = kA\Delta T/d$. Each term earns its place: more area gives the heat more parallel paths, a larger temperature difference drives harder, and greater thickness spreads the same difference over a longer distance and so flattens the gradient. The conductivity k is the material's own willingness to pass heat along, and it spans four orders of magnitude — copper near 400 W/(m·K), glass 0.96, mineral wool 0.04, still air 0.026.

Run the numbers on a single window pane: 1.5 m², 3 mm thick, 20 K from inside to outside. $P = 0.96 \times 1.5 \times 20/0.003 = 9600\ \mathrm{W}$. Nine and a half kilowatts through one window. That is obviously wrong, and being clear about *why* it is wrong is the most useful thing this page can teach.

The equation is fine; the model is incomplete. Heat has to reach the glass from the room air and leave it into the outdoor air, and both of those handoffs are slow. The still-air film clinging to the inside surface has a thermal resistance of about 0.12 m²·K/W and the wind-scoured outside film about 0.03, while the glass itself contributes only $d/k = 0.003/0.96 = 0.0031$. Resistances in series add, so the total is roughly 0.155 m²·K/W, giving $U = 1/R = 6.5\ \mathrm{W/(m^2{\cdot}K)}$ and a real heat flow of about **195 W** — fiftyfold less than the bare slab calculation. The glass was never the bottleneck. This series-resistance picture is exactly the electrical analogy: ΔT is voltage, P is current, and d/kA is resistance, which is why building science speaks in R-values (d/k per unit area) and U-values ($1/R_{total}$) rather than in this equation directly.

So the standard error is treating one layer as the whole assembly. Add layers by summing their resistances, never by averaging their conductivities, and never forget the two air films — in a well-insulated wall they are negligible, and in a window they dominate. It is also why a double-glazed unit works: the gain comes almost entirely from the trapped gas layer and the two extra surface films, not from doubling the glass.

Three further limits. This is a *steady-state* relation: it tells you the flow once the temperatures have settled and says nothing whatever about how long a wall takes to get there, which is a question of thermal mass and belongs to the diffusion equation. It is one-dimensional, which means it silently assumes heat goes straight through — a steel stud bridging an insulated cavity carries far more than its share of the area, and a thermal bridge can add a third to an assembly's real loss while the calculated R-value notices nothing. And ΔT is a difference, identical in kelvin and Celsius, but imperial conductivities quoted in BTU·in/(hr·ft²·°F) are a different quantity with a different thickness convention buried in them, so convert deliberately rather than by feel.

Worked example 3.4. Given k = 0.96 W/(m·K), A = 2 m², ΔT = 15 K, d = 4 mm; find P (heat flow rate). *2 m² glass pane, 4 mm, ΔT = 15 K → 7200 W.*

Module 3 practice

30. Specific Heat Capacity — An aquarium heater tops up 10 kg of water (c = 4200 J/(kg·°C)) by 6 C° overnight. *Calculate the heat the heater delivers, in kJ.*

31. Specific Heat Capacity — An aquarium heater tops up 20 kg of water (c = 4200 J/(kg·°C)) by 6 C° overnight. *Calculate the heat the heater delivers, in kJ.*

32. Rearranging Q = mcΔT — An electric kettle pushes 42 kJ into 1 kg of water (c = 4200 J/(kg·°C)) before someone flicks it off early. *Determine how much the water’s temperature rises.*

33. Rearranging Q = mcΔT — An electric kettle pushes 252 kJ into 3 kg of water (c = 4200 J/(kg·°C)) before someone flicks it off early. *Determine how much the water’s temperature rises.*

34. Latent Heat of Phase Change — A soup pot holds a rolling boil at a steady 100 °C while 0.5 kg of the water leaves as steam. (L for boiling water: 2260 kJ/kg.) *Calculate the heat that carried the steam away.*

35. Latent Heat of Phase Change — Backyard rink night: a thin flood of 1 kg of water, already chilled to 0 °C, freezes solid under a cold sky. (L for freezing water: 334 kJ/kg.) *Calculate the heat the freezing water released to the night air.*

36. Climbing the Heating Curve — A block of 1 kg of ice comes out of the deep freeze at −10 °C. A warming tray brings it up to 0 °C, melts it all, and warms the meltwater to 30 °C. (c for ice: 2100 J/(kg·°C); c for water: 4200 J/(kg·°C); L for melting: 334 kJ/kg.) *Determine the heat for the crossing at 0 °C, then for the whole journey up the curve.*

37. Climbing the Heating Curve — A camp kettle holds 2 kg of creek water at 30 °C. It is brought to a rolling boil, and then 1 kg of it boils away as steam. (c for water: 4200 J/(kg·°C); L for boiling: 2260 kJ/kg.) *Determine the heat for the climb to the boil, then for the whole job.*

38. Thermal Expansion — An aluminium power line hangs 100 m between two poles. Through a summer morning the wire warms by 25 C°. (α for aluminium: 24×10^{-6} /°C.) *Determine how much the wire lengthens, in mm.*

39. Thermal Expansion — An aluminium power line hangs 30 m between two poles. Through a summer morning the wire warms by 50 C°. (α for aluminium: 24×10^{-6} /°C.) *Determine how much the wire lengthens, in mm.*

40. Heat Flowing Through a Wall — A single-pane bedroom window has an area of 2 m² and is 4 mm thick. The glass (k = 0.96 W/(m·°C)) holds 15 C° between the warm room and the winter night. *Calculate the rate at which heat leaks out through the pane.*

41. Heat Flowing Through a Wall — A camping cooler's foam wall (k = 0.04 W/(m·°C)) has a total area of 1 m² and is 20 mm thick. The July air outside sits 15 C° above the ice inside. *Calculate the rate at which heat sneaks in.*

42. The Kettle Problem — Last question of the paper. A kitchen kettle holds 1 kg of tap water at 20 °C. It is brought to a rolling boil at 100 °C, and then 0.5 kg of it boils away as steam. (c for water: 4200 J/(kg·°C); L for boiling: 2260 kJ/kg.) Work each line — every answer feeds the next. *Determine the heat for the climb, the heat for the steam, and the total for the whole job.*

43. The Kettle Problem — Bonus mark, worked backwards. Overnight, the freezer pulls 668 kJ out of an ice-cube tray of water that was already sitting at 0 °C — and by morning every cube is solid. (L for freezing: 334 kJ/kg.) *Determine the mass of water that froze.*

MODULE 4

Electricity — Circuits & Power

Charge and Current · Ohm's Law · Rearranging Ohm's Law · Series Circuits · Parallel Circuits · Electrical Power · The Hydro Bill · The Circuit Board

4.1 Electric Charge (Q = It)

(Science 10-17) $$Q = It$$

Where,

Q = Charge (C)

I = Current (A)

t = Time (s)

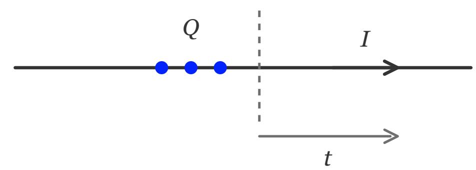

Figure 4.1 — electric charge (q = it).

Current is not a thing that flows. It is a *rate* — the amount of charge passing a chosen cross-section of the conductor each second — and one ampere means one coulomb per second. Once that is clear, $Q = It$ needs no proof, because it is the definition read backwards: if charge crosses at a steady rate, the total that crossed is the rate multiplied by how long it kept up. The only condition the equation imposes is the word **steady**. A current that varies has to be integrated, $Q = \int I\,dt$, and this page is the special case where the integral collapses to a rectangle.

A 2 A charger running for one hour moves $2 \times 3600 = 7200\ \mathrm{C}$. Battery ratings are the same arithmetic wearing different units: a phone cell marked 3000 mAh holds 3 Ah, and $3 \times 3600 = 10\,800\ \mathrm{C}$ of deliverable charge. Divide by the elementary charge, 1.602×10^{-19} C, and that is about 6.7×10^{22} electrons — a number that only sounds absurd until you remember a gram of copper contains ten times as many free ones already sitting in the metal, drifting at well under a millimetre per second.

Since the 2019 redefinition of the SI, this relation is closer to the foundation than it used to be. The ampere is now fixed by declaring the elementary charge to be exactly $1.602176634 \times 10^{-19}$ C, which makes the coulomb a *count* of charges and the ampere a count per second.

Michael Faraday got there experimentally in the 1830s: his laws of electrolysis measure the charge needed to plate out a mole of a substance, and that constant — 96 485 C per mole — is nothing but $Q = It$ run on a plating tank. Electroplating, anodising and battery capacity testing all still bill in ampere-hours for exactly this reason.

Two errors are worth naming, and one convention deserves an apology. The first error is treating milliamp-hours as energy. They are charge; a 3000 mAh cell at 3.7 V holds $3 \times 3.7 = 11.1\ \mathrm{Wh}$, and the same 3000 mAh at 1.2 V holds a third of that, so comparing two batteries by mAh alone tells you very little. The second is applying the equation to a current that is not constant — a motor's inrush, a switching supply's chopped input, or anything on AC, where over a full cycle the net charge transferred is zero even though the current is real all along. As for the convention: current is drawn flowing from plus to minus, while in a metal the electrons actually travel the other way. Benjamin Franklin guessed the sign in the 1750s, a century before anyone knew a charge carrier existed, and he guessed wrong. Nothing in the physics breaks — a deficit of negatives moving left is indistinguishable from positives moving right — but it is a historical accident, not a discovery, and it is worth knowing that it is one.

Worked example 4.1. Given I = 2 A, t = 30 s; find Q (charge). *2 A for 30 s → 60 C.*

4.2 Ohm's Law

(Science 10-18) $V = IR$

Where,
V = Voltage (V)
I = Current (A)
R = Resistance (Ω)

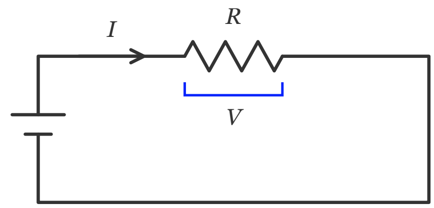

Figure 4.2 — ohm's law.

Georg Ohm published this relation in 1827 after painstaking experiments with wires of different lengths and thicknesses — and was initially ridiculed for reducing electricity to arithmetic. The idea is simple: voltage is the electrical *push*, resistance is the *opposition*, and current is what results. Double the push and you double the flow; double the opposition and you halve it. A 12 V car battery connected across a 6 Ω lamp drives 12/6 = 2 A through it; swap in a 3 Ω lamp and the current doubles to 4 A.

The law holds for **ohmic** conductors at constant temperature — metals, resistors, most wiring. Components like diodes, filament bulbs, and thermistors bend the rule because their resistance shifts as they heat up or as voltage changes. The classic V–I–R triangle mnemonic works because every rearrangement here is a single multiplication or division.

Worked example 4.2. Given I = 2 A, R = 6 Ω; find V (voltage). *2 A through 6 Ω → 12 V.*

4.3 Two Resistors in Series

(Science 10-19) $R_t = R_1 + R_2$

Where,
R_t = Total resistance (Ω)
R_1 = Resistance 1 (Ω)
R_2 = Resistance 2 (Ω)

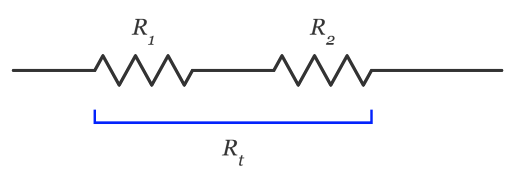

Figure 4.3 — two resistors in series.

Two resistors wired end to end sit on one path, and that single fact settles everything else. Charge has nowhere to go but forward, so the **same current passes through both** — it is not divided between them and it is not shared out according to size. Each resistor then takes its own voltage drop, $V_1 = IR_1$ and $V_2 = IR_2$, and the two drops must add up to whatever the supply provides. Divide that sum by the current they have in common and the resistances add: $R_t = R_1 + R_2$. This is not a rule to memorise. It is Ohm's law applied twice in a circuit that has only one loop.

Put a 47 Ω resistor in series with a 220 Ω resistor across a 12 V supply. The total is 267 Ω, so the current is $12/267 = 45\ \mathrm{mA}$, and it is 45 mA at every point in the loop — before the first resistor, between them, and after the second. The 47 Ω part drops $0.045 \times 47 = 2.1\ \mathrm{V}$; the 220 Ω part drops

$0.045 \times 220 = 9.9$ V; and 2.1 + 9.9 gives back the 12 V we started with. Adding the drops as a check costs nothing and catches most arithmetic errors on the spot.

The bookkeeping behind that check is Kirchhoff's voltage law, published by Gustav Kirchhoff in 1845 while he was still a student: go once around any closed loop and the voltage rises equal the voltage falls, because the loop returns you to the potential you started at. Every series result descends from it. The voltage divider is the same equation rearranged — each resistor claims the fraction $R_1/(R_t)$ of the supply — and that is how a potentiometer, a thermistor bridge and a sensor's biasing network all work. Solving this page backwards for one resistor is subtraction, $R_1 = R_t - R_2$, which is why the total must exceed the branch you already know.

The mistakes cluster in three places. The first is reaching for the wrong combination rule: series resistors add, parallel resistors combine as product over sum, and *capacitors do exactly the reverse* — series capacitors are the reciprocal case. Whenever you find yourself using product-over-sum on a series string, stop and ask which component you are holding. The second is assuming the larger resistor gets the larger current; it gets the larger *voltage drop* at the same current, and confusing those two makes a mess of any divider. The third only appears on AC: a coil or a capacitor in series with a resistor cannot be added arithmetically to it. Reactance is 90° out of phase with resistance, so a 30 Ω resistor in series with 40 Ω of reactance presents $\sqrt{30^2 + 40^2} = 50\ \Omega$, not 70. This page adds resistances, and resistances are what it will add — the quadrature sum belongs on the impedance pages.

Worked example 4.3. Given R_1 = 220 Ω, R_2 = 330 Ω; find R_t (total resistance). *220 Ω + 330 Ω in series → 550 Ω.*

4.4 Two Resistors in Parallel

(Science 10-20) $$R_t = \frac{R_1 R_2}{R_1 + R_2}$$

Where,
R_t = Total resistance (Ω)
R_1 = Resistance 1 (Ω)
R_2 = Resistance 2 (Ω)

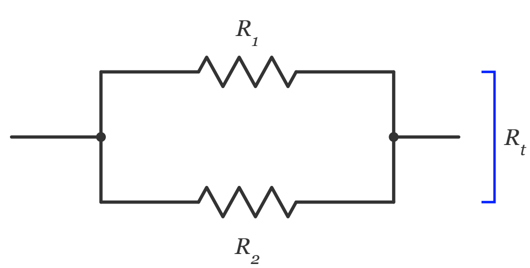

Figure 4.4 — two resistors in parallel.

Wired side by side, two resistors give the current two paths at once, so more current flows for the same voltage and the combination resists *less* than either branch alone. The tidy product-over-sum form is just $1/R_t = 1/R_1 + 1/R_2$ rearranged — conductances, not resistances, are what add in parallel. Worked example: 100 Ω in parallel with 25 Ω gives (100 × 25)/(100 + 25) = 2500/125 = 20 Ω, comfortably below the smaller branch.

Two handy special cases: equal resistors in parallel halve (two 100 Ω resistors make 50 Ω), and a much smaller resistor dominates — 10 Ω in parallel with 10 kΩ is essentially 10 Ω. Household outlets are wired in parallel so every appliance sees full mains voltage. When solving for a branch, the other resistance must exceed the total, since the total is always the smallest value in the circuit.

Worked example 4.4. Given R_1 = 4 Ω, R_2 = 12 Ω; find R_t (total resistance). *4 Ω ∥ 12 Ω → 3 Ω.*

4.5 Electrical Power (P = VI)

(Science 10-21) $$P = VI$$

Where,
P = Power (W)
V = Voltage (V)
I = Current (A)

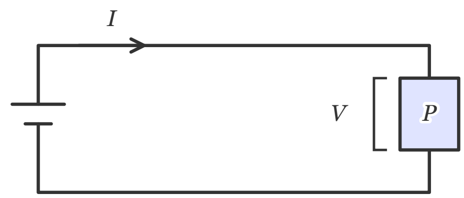

Figure 4.5 — electrical power (p = vi).

This formula falls straight out of the definitions. A volt is a joule per coulomb — the energy each unit of charge carries — and an ampere is a coulomb per second — how many units of charge arrive each second. Multiply them and the coulombs cancel, leaving joules per second: watts. A 1500 W space heater on a 120 V household circuit draws 1500/120 = 12.5 A, which is why it crowds a standard 15 A breaker and shouldn't share the circuit with much else.

Unlike the resistor-specific forms $P = I^2R$ and $P = V^2/R$, this version works for **any** component — ohmic or not — including motors, LEDs, and batteries, because it comes from the definitions of voltage and current rather than from Ohm's law. Electric utilities meter exactly this product, accumulated over time, when they bill you for energy.

Worked example 4.5. Given V = 120 V, I = 0.5 A; find P (power). *120 V at 0.5 A → 60 W.*

4.6 Electrical Power (P = I²R)

(Science 10-22) $$P = I^2R$$

Where,
P = Power (W)
I = Current (A)
R = Resistance (Ω)

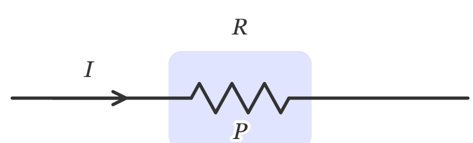

Figure 4.6 — electrical power (p = i²r).

This is the heat form of electrical power, and the square is the whole point of it. Start from $P = VI$, substitute Ohm's law for the voltage across the resistance, $V = IR$, and you get $P = I^2R$. Because the current appears twice, the heat does not track the load — it tracks the *square* of the load.

Double the current and a conductor dissipates four times the heat; triple it and nine times. Nothing else in ordinary wiring punishes a modest overload so hard, and it is the reason a circuit that runs warm at rated current runs dangerously hot at 150% of it.

Take a 30 m branch circuit run in 12 AWG copper. The conductor is about 3.31 mm², copper's resistivity is 1.68×10^{-8} Ω·m, so each metre is roughly 5.1 mΩ — and the current has to go out and come back, so the loop is 60 m and about 0.30 Ω. At 15 A the copper dissipates $15^2 \times 0.30 = 68\ \mathrm{W}$, spread along the run, and drops $15 \times 0.30 = 4.6\ \mathrm{V}$ of the supply before it ever reaches the load. Drop the current to 5 A and the loss falls to 7.5 W, not to a third: that is the square doing its work.

James Joule established the law in 1841 by immersing coils in water and measuring the temperature rise, which was also his route to the mechanical equivalent of heat and thence to the first law of thermodynamics. The same relation is the reason the grid transmits at hundreds of kilovolts. Delivered power is VI, so raising the voltage a hundredfold cuts the current a hundredfold for the same power, and cuts the line loss by ten thousand. Every transformer between a generating station and a house exists to move a fixed quantity of watts into a lower current, purely so that this equation returns a smaller number.

Three traps, in rising order of consequence. First, R is the resistance *of the thing dissipating the heat*, and I is the current through that same thing — mixing the conductor's resistance with the load's current is fine only because they are in series, and on a branched circuit it is not fine at all. Second, on AC the current must be RMS. A peak reading gives twice the power for a sine wave, and using it is one of the most common ways to double an answer without noticing. Third, and the one that bites in the field: this equation uses *resistance*, and on an AC circuit the opposition to current is impedance. A long run of steel-armoured cable, a coil, or a motor feeder has reactance as well as resistance, and the voltage drop computed from DC resistance alone will understate the real drop. The heat, though, still comes only from the resistive part — reactance stores energy and hands it back, so it moves voltage around without ever warming the copper.

Worked example 4.6. Given I = 3 A, R = 10 Ω; find P (power). *3 A through 10 Ω → 90 W.*

4.7 Electrical Energy (E = Pt)

(Science 10-23) $$E = Pt$$

Where,

E = Energy (J)

P = Power (W)

t = Time (s)

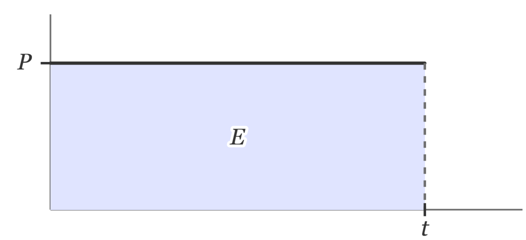

Figure 4.7 — electrical energy (e = pt).

Power is the rate at which energy is delivered, so energy is power kept up for a while. That is all $E = Pt$ says, and like every rate-times-time relation it is true by definition rather than by discovery — a watt *is* a joule per second, so watts multiplied by seconds give joules back. The distinction it enforces is the one people most often lose: **power is not energy**. A 2000 W heater is not consuming 2000 of anything; it is consuming at a rate of 2000 joules every second, and what it costs depends entirely on how long you leave it on.

Utilities meter in kilowatt-hours because the joule is inconveniently small: 1 kWh is 1000 W sustained for 3600 s, or exactly 3.6×10^6 J. A 1500 W baseboard heater running six hours a day for a thirty-day month uses $1.5 \times 6 \times 30 = 270\ \mathrm{kWh}$; at ten cents a kilowatt-hour that is \$27 on the bill. The same 270 kWh would run a 15 W LED lamp continuously for about two years. Energy comparisons like that are the only honest way to judge where a bill actually goes, and they almost always show that the heating and hot water dwarf everything with a screen on it.

The kilowatt-hour is a compound unit of the sort engineers usually avoid, and it survives because it matches how people buy electricity: a rate you can read off a nameplate multiplied by hours you can read off a clock. Watt-hours, ampere-hours and joules all measure the same physical stock of energy in different currencies: $1\ \mathrm{Wh} = 3600\ \mathrm{J}$, and an ampere-hour becomes watt-hours only after you multiply by the voltage. On the site's other pages this same relation appears as work over time in mechanics; there is no separate electrical version of it, only a separate unit.

The assumption doing the quiet work here is constant power, and most real loads are not. A thermostatted heater is either fully on or fully off, so its *average* power over an hour is the rated power times its duty cycle — a 1500 W baseboard cycling a third of the time is a 500 W load as far as the meter is concerned, and using the nameplate figure triples the estimate. A refrigerator, a well pump and a furnace blower all behave the same way. The other trap is a billing one worth knowing if you read a commercial invoice: those bills carry both an energy charge in kilowatt-hours and a *demand* charge in kilowatts, set by the highest fifteen-minute average draw in the period. The demand charge is a power charge, this equation does not produce it, and no amount of shortening run times will reduce it — only flattening the peak will.

Worked example 4.7. Given P = 60 W, t = 120 s; find E (energy). *60 W for 120 s → 7200 J.*

4.8 Energy Cost from a Utility Rate

(Science 10-24) $$C_e = E\,p_e$$

Where,

C_e = Energy cost ($)

E = Energy consumed (kWh)

p_e = Energy rate ($/kWh)

Water is not the only meter a cooling system spins. Tower fans, condenser-water pumps and the compressor itself all draw power, and boilers burn gas — so the same product, energy times a rate, prices both. A tower's fans and pumps drawing 250,000 kWh a year at $0.11/kWh cost $27,500; a boiler burning 20,000 MMBTU of gas at $8.00/MMBTU costs $160,000. North American electricity sits around $0.08–0.15/kWh commercial and natural gas around $6–10/MMBTU, but demand charges, ratchets and time-of-use blocks mean the effective rate on a bill is often well above the headline commodity rate — take it from the bill, dividing total dollars by total kilowatt-hours, rather than from the tariff sheet.

This calculation is what makes the water-treatment argument financial rather than technical. Scale is an insulator: a 0.6 mm (1/64 in) carbonate film on condenser tubes lifts compressor power by roughly 20%, and on a plant with a six-figure electricity bill that dwarfs the entire chemical budget. The same arithmetic prices the other direction too — boiler blowdown leaves at saturation temperature, so every percent of continuous blowdown costs a fraction of a percent of fuel, and a blowdown heat exchanger's payback is nothing more than this equation applied to recovered energy. Energy is entered and answered in kilowatt-hours and the rate in dollars per kilowatt-hour regardless of the metric/imperial toggle, because that is how every electricity meter on earth reads; and as with every money answer here, the currency is whatever currency you typed the rate in.

Worked example 4.8. Given E = 2e10 BTU, p_e = 8.00 $/MMBTU; find C_e (energy cost). *20,000 MMBTU of gas at $8.00/MMBTU → $160,000.*

Module 4 practice

44. Charge and Current — A doorbell circuit carries 5 A for the 50 s someone leans on the button. *Determine the charge that flows through the bell.*

45. Ohm's Law — A toaster element with a resistance of 20 Ω is plugged into the 120 V wall outlet. *Calculate the current through the element.*

46. Ohm's Law — A toaster element with a resistance of 10 Ω is plugged into the 120 V wall outlet. *Calculate the current through the element.*

47. Rearranging Ohm's Law — A 24 V battery pack pushes current through a single 8 Ω resistor. *Determine the current in the circuit.*

48. Rearranging Ohm's Law — A small heater in the science room draws 2 A through its 7 Ω element. *Determine the voltage across the element.*

49. Series Circuits — Two resistors, 1 Ω and 4 Ω, are connected end to end in a single loop with a 10 V battery. *Determine the total resistance, then the current the battery drives.*

50. Series Circuits — Two resistors, 2 Ω and 4 Ω, are connected end to end in a single loop with a 24 V battery. *Determine the total resistance, then the current the battery drives.*

51. Parallel Circuits — Two resistors, 18 Ω and 9 Ω, are connected side by side across the same battery — each on its own branch. *Determine the total resistance of the parallel pair.*

52. Parallel Circuits — Two resistors, 6 Ω and 12 Ω, are connected side by side across the same battery — each on its own branch. *Determine the total resistance of the parallel pair.*

53. Electrical Power — The label has worn off a space heater's element, but the meter shows 2 A flowing through its 10 Ω of resistance. *Calculate the power the element gives off as heat.*

54. Electrical Power — A flat-screen TV on the 120 V outlet draws 3 A. *Calculate the device's power.*

55. The Hydro Bill — A 1000 W window air conditioner hums along for 8 h on a hot afternoon. Electricity costs 10 ¢ for every kilowatt-hour. *Determine the energy used, and what it cost.*

56. The Hydro Bill — A 3000 W hot tub heater keeps the tub warm for 3 h. Electricity costs 12 ¢ for every kilowatt-hour. *Determine the energy used, and what it cost.*

57. The Circuit Board — Last bench of the exam. A circuit board holds two resistors in series — 4 Ω and 8 Ω — across a 24 V battery, and the board is left running for 2 h. Work each line — every answer feeds the next. *Determine the energy the board uses in 2 h, one line at a time.*

MODULE 5

Weather, Climate & the Atmosphere

Pressure Is Force over Area · Pressure Under Water · Thin Air — Pressure with Altitude · Relative Humidity · The Dew Point · Feels-Like Temperatures · Sunlight, Albedo and the Lapse Rate · Counting the Storm · Boss — The Weather Station

5.1 Pressure (P = F/A)

(Science 10-25) $$P = \frac{F}{A}$$

Where,
P = Pressure (kPa)
F = Force (N)
A = Area (m²)

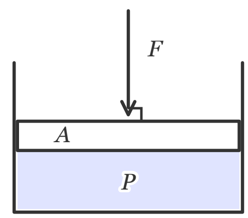

Figure 5.1 — pressure (p = f/a).

Pressure is force divided by the area it is spread over, $P = F/A$. The same force can be gentle or destructive depending entirely on how much surface it acts through, and that is the whole content of the idea. One pascal is one newton per square metre, which is a very small pressure — a sheet of paper lying on a table exerts about 1 Pa — so real numbers are almost always in kilopascals or higher. Standard atmospheric pressure is 101.325 kPa.

An 80 kg person weighs about 785 N. Standing in winter boots with maybe 350 cm² of sole in contact — 0.035 m² — the pressure under them is $785/0.035 \approx 22$ kPa, enough to break through crusted snow. Strap on snowshoes with 0.30 m² of bearing surface and the same 785 N becomes 2.6 kPa, and the snow holds. Run it the other way: lean on a thumbtack with 20 N through a point of about 0.01 mm², which is 1×10^{-8} m², and the pressure at the tip is around 2 GPa — well past what wood fibre can resist.

The same relation is behind hydraulics. Pascal's principle says pressure applied to a confined fluid is transmitted undiminished throughout it, so a small force on a small piston becomes a large force on a large one: same P, bigger A, bigger F. A hydraulic jack with a 200:1 area ratio multiplies force 200-fold, at the price of moving 200 times as far. It also underlies the hydrostatic pressure page, where the force is simply the weight of the fluid standing above.

The area to use is the area actually in contact, and it is usually smaller than it looks. A tyre bears on its contact patch, not on the outline of the tread; a bolted flange bears on the annulus under the washer, not on the whole plate. Getting that wrong understates the real pressure, sometimes badly. The unit conversion is the other reliable trap: a square metre is 10 000 square centimetres and a million square millimetres, so an area entered in cm² against a force in newtons is out by four orders of magnitude. And F must be the component *perpendicular* to the surface — a force at an angle contributes only its normal component to pressure, with the rest showing up as shear. Finally, be clear whether a quoted pressure is gauge or absolute. Gauge pressure is measured against the surrounding atmosphere, so a tyre reading "zero" is not empty; it holds about 101 kPa absolute. The suffixes psig and psia exist precisely because this goes wrong so often.

Worked example 5.1. Given F = 500 N, A = 0.25 m²; find P (pressure). *500 N on 0.25 m² → 2 kPa.*

5.2 Hydrostatic Pressure (P = ρgh)

(Science 10-26) $P = \rho g h$

Where,

P = Gauge pressure (kPa)

ρ = Fluid density (kg/m³)

h = Depth (m)

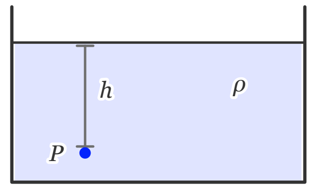

Figure 5.2 — hydrostatic pressure (p = ρgh).

Stand a column of fluid up and it presses down with its own weight. The pressure at depth h is $P = \rho g h$ — density times gravity times depth — and the remarkable thing about it is what is *not* in the formula. There is no area, no volume, no mention of the shape of the container. Pressure at a given depth depends only on how far down you are and what the fluid is. A litre of water in a narrow tube 3 m tall produces exactly the same pressure at its base as a swimming pool 3 m deep.

Three metres down in fresh water: $P = 1000 \times 9.80665 \times 3 \approx 29.4$ kPa. The useful number to carry is that every metre of water is about 9.81 kPa, so every 10 m of water adds roughly one atmosphere — which is why divers count depth in atmospheres and why your ears complain at the deep end of a pool. Mercury, at 13 546 kg/m³, does the same job in 760 mm, which is where that famous barometric height comes from.

Simon Stevin worked this out in the 1580s, and Pascal is said to have demonstrated it by fixing a long thin tube into the top of a sealed barrel and bursting the barrel with a few cups of water poured down the tube. Whether or not the barrel story is literally true, the point stands and is still called the hydrostatic paradox. In practice this equation is mostly used as a translator: divide a pressure by ρg and you get *head* in metres, which is the language pump curves are written in.

Three things go wrong, and the density one costs the most money. Head in metres is fluid-specific: a pump rated for 30 m of head delivers $1000 \times 9.80665 \times 30 \approx 294$ kPa on water, but on a 1.20 SG brine the same 30 m of head is 353 kPa, and sizing the system as though a metre were a fixed pressure will have you short. Second, h is *vertical* depth, measured straight down from the free surface — not the length of pipe, not the run along a sloping hose, and not the distance around a bend. A hundred metres of hose lying flat on the ground develops no static head at all. Third, this is **gauge** pressure, the amount by which the fluid exceeds the atmosphere above it. For absolute pressure, add about 101 kPa; at 3 m depth in a pool the absolute pressure is roughly 131 kPa, not 29.4.

Worked example 5.2. Given ρ = 1000 kg/m³, h = 10 m; find P (gauge pressure). *10 m of water → 98.0665 kPa.*

5.3 Barometric Pressure with Altitude

(Science 10-27) $P = P_0\, e^{-Mgz/RT}$

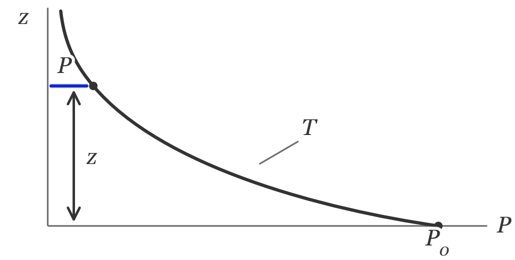

Figure 5.3 — barometric pressure with altitude.

Where,

P = Pressure at height (kPa)
P_0 = Reference pressure (kPa)
z = Height above reference (m)
T = Layer temperature (°C)

Stack a column of air on itself and each layer has to carry the weight of everything above it. Write that as a hydrostatic balance, substitute the ideal gas law for the density, hold the temperature constant, and the integration gives an exponential: $P = P_0 e^{-Mgz/RT}$. The group RT/Mg has units of length and is called the scale height, the climb over which pressure falls by a factor of e. At 15 °C it works out to $8.31446 \times 288.15/(0.0289644 \times 9.80665) = 2395.8/0.28404 = 8435$ m. So a climb of 1000 m from sea level is 0.1186 scale heights, the exponential factor is $e^{-0.1186} = 0.8882$, and the pressure falls from 101.325 kPa to 89.997 kPa.

The molar mass in that expression, 0.0289644 kg/mol, is the value the US Standard Atmosphere assigns to dry air, and it is a weighted average over nitrogen, oxygen, argon and carbon dioxide rather than a property of any single substance. Humid air is lighter than dry air, because a water molecule at 18 g/mol displaces a nitrogen molecule at 28, so a saturated tropical atmosphere has a slightly larger scale height than this constant admits. The effect is under one percent and is usually swamped by the temperature assumption, which is the real weakness here.

That assumption is worth being blunt about. The atmosphere is not isothermal, and this equation knows nothing about the lapse rate. Compare it against the standard atmosphere, which integrates a 6.5 °C/km gradient properly and gives 89.875 kPa at 1000 m: the isothermal answer of 89.997 kPa is high by 0.12 kPa, or about a tenth of a percent, which nobody cares about. Go to 5000 m and the isothermal form using the sea-level temperature returns 56.0 kPa against the standard atmosphere's 54.05 kPa, an error near four percent. The fix is free and is the whole reason the variable here is called the LAYER temperature: feed it the mean temperature of the layer rather than the temperature at the bottom. Using the average of 15 °C and the 17.5 °C below zero found at 5000 m gives 54.06 kPa, which is right to a few hundredths of a percent.

In dispersion work this equation earns its place three ways. It supplies the ambient pressure that Holland's plume-rise equation needs. It sets the air density that converts a stack's volumetric flow to a mass flow, and a plant at 1500 m elevation is moving air about fifteen percent less dense than a plant at sea level, which changes both the exit velocity and the buoyancy flux. And it is the correction that puts a measured emission rate onto the standard conditions a permit is written against. The recurring errors are the obvious two: temperatures in Celsius rather than kelvin, which makes the exponent nonsense, and using the equation above the tropopause, where the real atmosphere switches to an isothermal and then a warming regime and a single-layer model has nothing left to say.

Worked example 5.3. Given P_0 = 101.325 kPa, z = 1000 m, T = 15 °C; find P (pressure at height). *101.325 kPa at 15 °C → 89.997 kPa at 1000 m.*

5.4 Relative Humidity from Vapour Pressure

(Science 10-28) $$\varphi = \frac{p_v}{p_{ws}}$$

Where,

φ = Relative humidity (%)

p_v = Water vapour partial pressure (kPa)

p_{ws} = Saturation vapour pressure (kPa)

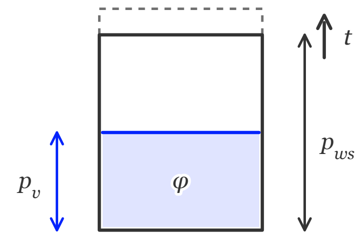

Figure 5.4 — relative humidity from vapour pressure.

Relative humidity is the most quoted and least understood number in building science. Its definition is simple enough: the vapour pressure actually present divided by the saturation vapour pressure at the same temperature, $\varphi = p_v / p_{ws}$. The trouble is entirely in that denominator, because the denominator moves.

Relative humidity is a ratio to a moving target. The saturation pressure roughly doubles for every 11 °C of warming, so heating air without adding one molecule of water sends its RH down while the actual moisture content is completely untouched. Cool the same air and RH climbs, again with no moisture change at all, until at the dew point it reaches 100 % and water starts falling out. This means a relative humidity reading tells you almost nothing about how much water is present unless you also know the temperature — and it is quoted constantly as though it did.

Two practical consequences follow. First, dry winter indoor air is not caused by heating removing moisture; heating removes nothing. Outdoor air at −10 °C and 80 % RH carries about 1.6 g/kg, and once warmed to 21 °C that identical air reads near 11 % RH. Second, RH is nonetheless the right variable for a great many questions, because most of the things we care about respond to relative humidity rather than absolute content: mould germinates above roughly 80 % surface RH, wood and paper equilibrate their moisture content against RH, static electricity becomes a nuisance below about 30 %, and human comfort tracks it. Materials do not count grams; they respond to how close the air is to saturation at their own temperature.

The chronic field error is comparing an RH reading in one place with an RH reading in another at a different temperature and concluding something about moisture migration. Two rooms at 50 % RH and 18 °C and 24 °C hold quite different amounts of water, and the difference will drive vapour from one to the other. If you want to reason about where moisture is going, convert both to humidity ratio or dew point first — those are the quantities that compare directly across temperatures.

Readings above 100 % deserve a note rather than a refusal. Supersaturation is genuinely real: it is what fog and mist and the visible plume from a cooling tower are, and cloud physics runs on it. In a duct, though, it almost always means one of two things — either p_{ws} was evaluated at the wrong temperature, which must be the same dry-bulb the vapour is sitting at, or the air really has hit its dew point and the excess is already condensing on the nearest cold surface.

One further subtlety in the definition. Strictly, relative humidity is defined against the saturation pressure of pure water, while the vapour in real air is very slightly more soluble than that idealisation allows; ASHRAE carries an enhancement factor of about 1.004 at ordinary conditions to account for it. Almost nobody applies it, because 0.4 % is well inside the accuracy of any field instrument — but it is one more reason to expect small disagreements between references and to stop hunting for the source of a half-percent gap.

Worked example 5.4. Given p_v = 1.2 kPa, p_{ws} = 2.339 kPa; find φ (relative humidity). *1.2 kPa against 2.339 kPa saturation → 51.30% RH.*

5.5 Saturation Vapour Pressure (Magnus / Alduchov–Eskridge)

(Science 10-29) $$p_{ws} = 610.94 \exp\left(\frac{17.625\, t}{t + 243.04}\right)$$

Where,

p_{ws} = Saturation vapour pressure (kPa)

t = Temperature (°C)

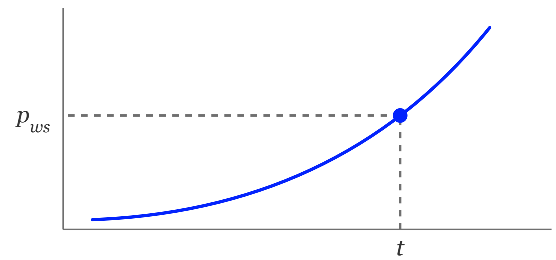

Figure 5.5 — saturation vapour pressure (magnus / alduchov–eskridge).

Air does not "hold" water the way a sponge holds it, and the sponge picture is the source of most confusion about humidity. Water vapour is a gas sharing a container with nitrogen and oxygen, and it exerts its own partial pressure quite independently of them. What temperature sets is the maximum partial pressure that vapour can sustain before it starts condensing back to liquid faster than it evaporates. That ceiling is the saturation vapour pressure, and it is a property of WATER ALONE. The air is not involved. Water saturating into a vacuum at 24 °C reaches the same pressure it reaches into a room.

The relationship is violently non-linear. From 0 °C to 24 °C the saturation pressure climbs from 0.611 kPa to 2.978 kPa — nearly five times — and it roughly doubles for every 11 °C. That single curve explains an enormous amount: why tropical air carries so much more moisture than arctic air, why a small drop in surface temperature makes a window run with condensation, and why the latent load on a coil in Houston dwarfs the one in Calgary at the same relative humidity.

Clausius and Clapeyron give the exact shape thermodynamically, but their equation has no closed-form solution, so everyone uses a fitted approximation. This page implements Alduchov and Eskridge's 1996 refinement of the Magnus form, $p_{ws} = 610.94\exp(17.625t/(t + 243.04))$, valid from −40 °C to +50 °C with a stated maximum error of 0.384 %.

Expect the third digit to disagree with your reference, and do not treat that as an error. At least four coefficient sets are in wide circulation: the original Magnus/Tetens (6.1078 hPa, 17.27, 237.3), Buck's 1981 set, the WMO/Sonntag set that this catalog's dew-point page uses, and the Alduchov–Eskridge set here. At 24 °C they give roughly 2978 to 2985 Pa. A psychrometric chart drawn from the IAPWS reference formulation will read 2.985 kPa where this page reads 2.978 — a gap of 0.23 %, which is far smaller than the error in reading a sling psychrometer and far smaller than the difference between two thermometers on the same wall. What matters is knowing which fit you are quoting, and this one says so.

One trap worth naming. Below 0 °C this curve is saturation over SUPERCOOLED LIQUID WATER, which is what psychrometric charts tabulate. Saturation over ice is lower — about 4 % lower at −10 °C — and that gap is exactly why frost grows on a cold surface while surrounding droplets stay liquid: the ice is a lower-pressure sink, so vapour migrates to it. If your problem is a freezer coil, a frost line or an outdoor coil in a defrost cycle, you want the sublimation curve and different coefficients, not this one.

Worked example 5.5. Given t = 24 °C; find p_{ws} (saturation vapour pressure). *24 °C → saturation vapour pressure 2.978 kPa.*

5.6 Dew Point (Magnus Approximation)

(Science 10-30) $$T_d = \frac{c\,\gamma}{b-\gamma}, \quad \gamma = \ln\frac{\mathrm{RH}}{100} + \frac{b\,T}{c+T}$$

Where,
T_d = Dew point (°C)
T = Air temperature (°C)
RH = Relative humidity (%)

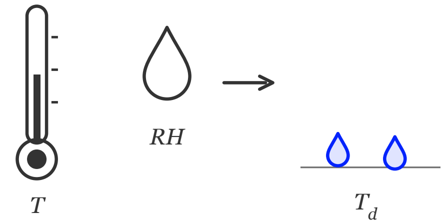

Figure 5.6 — dew point (magnus approximation).

The dew point is the temperature air must be cooled to before its water vapour starts condensing. Unlike relative humidity, which is a ratio that swings all day as the temperature moves, the dew point is close to an absolute measure of how much water is actually in the air. That is why forecasters and HVAC technicians reach for it: 70 % humidity means something entirely different in November than in July, while a dew point of 20 °C means sticky everywhere on earth.

The formula is a Magnus fit, an empirical curve for saturation vapour pressure with coefficients chosen to match laboratory measurements. This page uses the Sonntag values $b = 17.62$ and $c = 243.12\,°\mathrm{C}$ recommended by the World Meteorological Organization, good to about 0.1 °C between −45 and 60 °C. Other coefficient sets published by Tetens, Buck and Alduchov appear in textbooks and give answers a tenth of a degree apart, which is a fair statement of how well anyone knows this curve.

Two things fall straight out of the algebra. The dew point can never exceed the air temperature, since that would require more than 100 % humidity, and when the two are equal the air is saturated and fog or dew is imminent. Cooling a surface below the dew point is precisely how a cold glass sweats, why ductwork needs insulation, and how a dehumidifier works.

Worked example 5.6. Given T = 20 °C, RH = 50 %; find T_d (dew point). *20 C at 50 % RH → dew point 9.255 C.*

5.7 Humidex (Canadian Humidity Index)

(Science 10-31) $$H = T + 0.5555\,(e-10), \quad e = 6.11\,e^{5417.753\left(\frac{1}{273.16}-\frac{1}{T_d}\right)}$$

Where,
H = Humidex (humidex)
T = Air temperature (°C)
T_d = Dew point (°C)

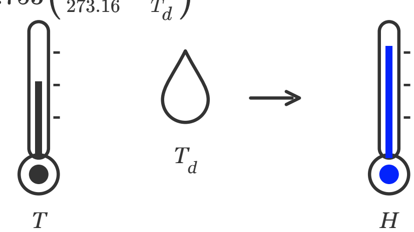

Figure 5.7 — humidex (canadian humidity index).

The humidex is Canada's answer to the question of how hot it feels when the air is wet. J. M. Masterton and F. A. Richardson devised it in 1979 at the Atmospheric Environment Service by adding to the air temperature a term proportional to the excess of the actual vapour pressure over 10 hPa, a value chosen to represent comfortably dry air. When the air is that dry the correction is zero and the humidex is simply the temperature.

Environment Canada is careful to call the result a unitless index rather than a temperature, and this page follows that convention. It sits on the Celsius scale by construction, so a humidex of 40 feels roughly like 40 °C in dry air, but it is a comfort scale and not a reading you could take with a thermometer. The published guidance runs: below 29 no discomfort, 30 to 39 some discomfort, 40 to 45 great discomfort with outdoor exertion to be curtailed, and above 45 dangerous.

The vapour pressure comes from the dew point through a Clausius-Clapeyron style exponential, which is why the humidex climbs so steeply on muggy days. A rise of a few degrees in dew point moves the vapour pressure far more than the same rise in air temperature moves anything, and the index reflects a real physical fact: sweat is your cooling system, and humid air will not accept it.

Worked example 5.7. Given T = 30 °C, T_d = 15 °C; find H (humidex). *30 C, dew point 15 C → humidex 33.97 (EC quotes 34).*

5.8 Wind Chill (2001 North American Formula)

(Science 10-32) $$T_{wc} = 13.12 + 0.6215\,T - 11.37\,V^{0.16} + 0.3965\,T\,V^{0.16}$$

Where,

T_{wc} = Wind chill index (°C)

T = Air temperature (°C)

V = Wind speed (km/h)

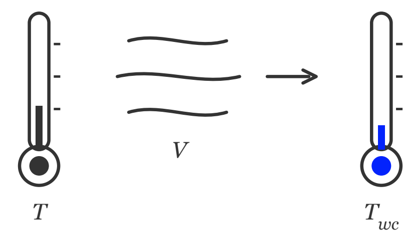

Figure 5.8 — wind chill (2001 north american formula).

Wind chill answers one question: how cold does bare skin lose heat, expressed as the still-air temperature that would strip it at the same rate. The 2001 North American formula replaced the 1945 Siple and Passel index, which had been calibrated by measuring how fast water froze in plastic cylinders hung from a pole in Antarctica and which produced numbers most people found implausibly severe. The replacement came from a joint Canadian and American working group that modelled heat loss from a human face, then validated it on volunteers walking in a chilled wind tunnel at the Defence and Civil Institute of Environmental Medicine in Toronto.

The details in the coefficients are worth knowing. Wind speed enters as $V^{0.16}$, a strongly flattening power, which is why the first 20 km/h of wind does most of the damage and the next 20 does much less. The speed is measured at the standard 10 m anemometer height and internally scaled down to face height, so the number already accounts for the fact that your face is not at the top of a mast. And the index is defined only at or below 10 °C with winds above 4.8 km/h, which is why this page refuses to compute outside that window rather than returning a confident-looking answer from a fit that was never tested there.

One thing it does not tell you is how cold an *object* gets. Wind chill cannot push anything below the true air temperature; a windy −10 °C night will not freeze your pipes any harder than a calm one, it will merely get them there faster. What the index does track well is frostbite risk on exposed skin, which is the reason it is broadcast.

Worked example 5.8. Given T = -10 °C, V = 30 km/h; find T_{wc} (wind chill index). *-10 C with 30 km/h → -19.520 C.*

5.9 Heat Index (Rothfusz Regression)

(Science 10-33) $$\mathrm{HI} = c_1 + c_2 T + c_3 R + c_4 T R + c_5 T^2 + c_6 R^2 + c_7 T^2 R + c_8 T R^2 + c_9 T^2 R^2$$

Where,

HI = Heat index (°C)

T = Air temperature (°C)

R = Relative humidity (%)

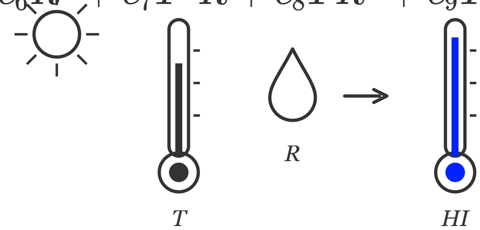

Figure 5.9 — heat index (rothfusz regression).

The heat index is the American cousin of the humidex, and it comes from a much more elaborate model. In 1979 Robert Steadman published a multi-page thermoregulation calculation involving clothing, activity, skin resistance and evaporation, producing tables of the apparent temperature. Nobody wanted to run that on air, so in 1990 Lans Rothfusz fit a nine-term polynomial regression through Steadman's tables. That regression is what the National Weather Service broadcasts and what this page computes.

Since it is a fit to a model, it inherits every assumption underneath: a person of average build walking at about 5 km/h in the shade, with a light breeze and normal clothing. Full sun can add up to 8 °C to the effective value. The regression is published for temperatures of 80 °F (26.7 °C) and above with humidity of 40 % or more, which is why this page declines outside that box, and the NWS applies small corrections at very low and very high humidity that are not included here.

Solving it backwards is left out on purpose. Taken as a function of temperature it is a quadratic with two roots, and there is no honest way to tell a user which branch the weather is on, so the page offers the forward calculation only. A regression that gives one confident answer in the direction it was fit, and an ambiguous one in reverse, should say so rather than pick.

Worked example 5.9. Given T = 35 °C, R = 50 %; find HI (heat index). *35 C at 50 % RH → 40.675 C.*

5.10 Albedo and Reflected Radiation

(Science 10-34) $$G_r = \alpha\, G$$

Where,

G_r = Reflected irradiance (W/m²)

α = Albedo

G = Incident irradiance (W/m²)

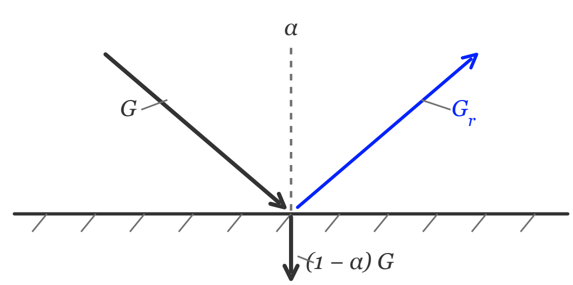

Figure 5.10 — albedo and reflected radiation.

Albedo is the plainest formula on this shard and the one with the largest consequences. A surface reflects the fraction α of the shortwave radiation that lands on it and absorbs the rest, so reflected is αG and absorbed is (1 − α)G. Under 800 W/m² of sun, grass at α = 0.25 bounces 200 W/m² back and keeps 600. Fresh snow at 0.85 bounces 850 W/m² of a full 1,000 and keeps only 150, which is why deep snow can persist in bright sunshine and why snow blindness is a real injury rather than an exaggeration. The word is Latin for whiteness, and Johann Heinrich Lambert brought it into optics in 1760.

Typical values are worth carrying in your head, because they span more than an order of magnitude. Fresh asphalt sits near 0.05 and aged asphalt near 0.10 to 0.15. Bare soil runs 0.10 to 0.25 depending on how wet it is; wet soil is markedly darker than dry. Grass and most crops fall between 0.20 and 0.26, deciduous forest near 0.15 to 0.18, and conifer forest lower still at 0.08 to 0.15 — a dark forest canopy absorbs almost as greedily as tarmac. Water is the awkward one: near 0.05 with the sun overhead, but rising steeply toward 0.6 and beyond at grazing incidence, which is why a lake dazzles at sunset and not at noon. Fresh snow reaches 0.80 to 0.90 and decays to 0.4 or lower as it ages and dirties. The planetary average is about 0.30.

The mistakes are mostly bookkeeping. Albedo is a SHORTWAVE property, not the longwave emissivity used one page over, and the two are genuinely independent: white paint has a high albedo and a high longwave emissivity at the same time, which is exactly the combination a cool roof wants. Albedo also varies with sun angle, with cloudiness, with wetness and with season, so a single tabulated number is a daily average at best. And it is a proportion, never a percentage in disguise: 25 and 0.25 are not the same input unless you pick the % unit.

The absorbed complement is the term that matters for everything downstream, so this page computes and reports it alongside the reflection. That absorbed flux is what warms the ground, drives evaporation, and sets the daytime side of a microclimate — and it is why an urban surface of dark roofs and asphalt runs hotter than the countryside beside it, before you have said a word about waste heat.

The feedback that makes albedo interesting rather than merely arithmetic is the one snow provides. Snow reflects, so the ground beneath it stays cold, so the snow lasts; melt it and you expose dark ground with an albedo four times lower, which absorbs far more, which melts the snow around it faster. That runaway works in both directions and it is the reason polar regions warm and cool faster than anywhere else. The same loop operates on a garden scale in spring: a patch of bare dark soil among snow will clear itself in a day, while the shaded white beside it sits there for a week.

Worked example 5.10. Given α = 0.25, G = 800 W/m²; find G_r (reflected irradiance). *Grass at albedo 0.25 under 800 W/m2 → 200 W/m2 reflected.*

5.11 Environmental Lapse Rate

(Science 10-35) $$\Gamma = \frac{T_1 - T_2}{z_2 - z_1}$$

Where,

- Γ = Lapse rate (°C/km) (°C/km)
- T_1 = Temperature below (°C)
- T_2 = Temperature above (°C)
- Δz = Height difference (m)

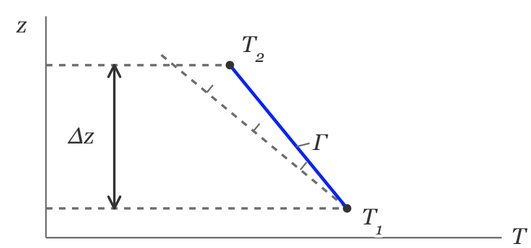

Figure 5.11 — environmental lapse rate.

A lapse rate is one subtraction and one division, and the only trick in it is the sign convention. The rate is written as the temperature BELOW minus the temperature ABOVE, so that ordinary air, which cools as you climb, comes out positive. Two thermometers reading 20 °C at the surface and 13.5 °C on a mast 1000 m higher give $\Gamma = (20 - 13.5)/1000 = 0.0065$ °C per metre, which the whole meteorological literature writes as 6.5 °C/km. That number is not an accident of the example. ICAO fixed 6.5 °C/km as the troposphere of the standard atmosphere in 1952, and every pressure altimeter in the world is calibrated against it. Note that only the DIFFERENCE of the two temperatures enters, so a Celsius interval and a kelvin interval are the same thing and the conversion cancels itself.

The rate becomes useful only when it is compared against the rate a parcel of air would cool at if it were lifted with no heat exchange at all. That one is not measured, it is derived. A rising parcel expands and does work against the surrounding pressure, and the energy comes out of its own heat content, so $\Gamma_d = g/c_p = 9.80665/1005 = 0.00976$ K/m, or 9.8 °C/km. Below that value the atmosphere is stable: a lifted parcel cools faster than its surroundings, becomes denser, and sinks back where it came from. Above it the parcel stays warmer than the air around it and keeps climbing, which is instability. Saturated air releases latent heat as it rises and so cools more slowly, roughly 5.4 °C/km in warm air and nearer 8 in cold air, which is why the band between 5.4 and 9.8 is called conditionally unstable.

Everything a plume does follows from that comparison. In unstable air the plume loops, dragged up and down by convective eddies, and a loop that touches down produces a brief ground-level concentration far above anything an hourly average would suggest. In neutral air it cones, spreading symmetrically, which is the case the Gaussian model describes best because it is the case the Gaussian model was fitted to. Under an inversion it fans, flattening into a thin ribbon that can travel many kilometres almost undiluted while the ground underneath sees essentially nothing. The dangerous part is what happens next. When morning sun heats the surface and erodes the inversion from below, the entire night's ribbon is mixed down at once. That is fumigation, and it produces the highest short-term ground-level concentrations most sources ever cause.

Two mistakes are worth naming. The first is measuring the gradient across too thin a layer near the ground, where on a sunny afternoon the bottom few metres can be superadiabatic by tens of degrees per kilometre while the air a hundred metres up is perfectly neutral. A lapse rate is only meaningful over the layer the plume actually occupies. The second is treating the lapse rate as the whole of stability. The Pasquill class that supplies the dispersion coefficients is set by wind speed and solar radiation as well, and a strong wind mixes mechanically no matter what the temperature profile says. Use the lapse rate to know which regime you are in, then take the class from the standard insolation and wind table rather than from the gradient alone.

Worked example 5.11. Given T_1 = 20 °C, T_2 = 13.5 °C, Δz = 1000 m; find Γ (lapse rate (°c/km)). *6.5 °C drop over 1 km → 6.5 °C/km (the standard atmosphere).*

5.12 Flash-to-Bang Distance to a Lightning Strike

(Science 10-36) $$d = (331.3 + 0.606\,T_C)\,t$$

Where,
d = Distance to the strike (km)
t = Seconds from flash to thunder (s)
T = Air temperature (°C)

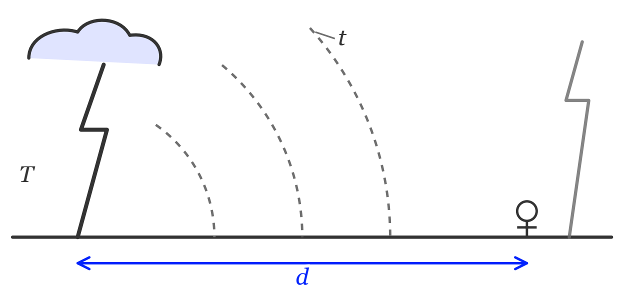

Figure 5.12 — flash-to-bang distance to a lightning strike.

Light takes about eleven microseconds to cover three kilometres and sound takes nearly nine seconds, so for any purpose you have standing in a field, the flash is instantaneous and the thunder is not. Count the gap, multiply by the speed of sound, and you have the distance to the part of the channel nearest you. That is the whole of the physics, and it is genuinely reliable — which is exactly why the page has to be careful about what it does *not* tell you.

Start with the thing this page cannot do, because everything else is arithmetic. A distance is not a safety margin. Lightning reaches ground more than ten kilometres from the storm that produced it — the phenomenon has a name, the bolt from the blue — striking ahead of the visible cloud, behind it, and out of a patch of sky that looks perfectly clear overhead. A measured eight kilometres tells you where *that* flash went. It says nothing whatsoever about where the next one goes. The published guidance belongs to the meteorological services and is quoted here rather than invented: NOAA's National Weather Service and Environment and Climate Change Canada both publish the 30-30 rule — if the gap between flash and thunder is under thirty seconds, be inside a substantial building or a hard-topped vehicle already, and wait thirty minutes after the *last* thunder before going back out. NOAA's shorter public version of the first half is "When Thunder Roars, Go Indoors".

And silence proves nothing at all. Sound in the lower atmosphere is refracted upward by the ordinary decrease of temperature with height and by wind shear, so beyond roughly twenty-five kilometres — often far less, and less again in wind, in hills, or in traffic noise — the thunder from a completely normal flash simply passes over your head and never arrives. A flash you can see but cannot hear is not a distant flash. It is a flash whose sound went somewhere else.

Now the folk rules, which is why the temperature is an input here. "Five seconds per mile" and "three seconds per kilometre" are both rounded from one calculation done at one temperature. At 20 °C sound covers a mile in 4.69 s and a kilometre in 2.91 s, so both rules round up, and rounding up reports the storm as *closer* than it is — about 6 % closer, which is the conservative direction and is why nobody ever minds. What the rules cannot do is follow the air. The speed of sound in dry air is very nearly linear in temperature, $v = 331.3 + 0.606\,T_C$, running from about 319 m/s at −20 °C to about 353 m/s at +35 °C. That is a 10 % spread, and an identical five-count means 1.067 mi on a summer afternoon and 0.992 mi in a −20 °C snow squall.

Two practical cautions on the count itself. Thunder from a distant flash arrives as a long roll, and the distance corresponds to the *first* sound — timing to the loudest part or to the end overstates the range, sometimes badly. And in a busy storm it is very easy to time one flash against another flash's thunder, which produces a confident number about nothing. If the counts from successive flashes are wildly inconsistent, that is usually what has happened.

Humidity, incidentally, barely matters: moist air is slightly less dense than dry air at the same pressure, so sound travels marginally faster in it, but the effect is a fraction of a per cent and is swamped by the temperature term. Wind matters more, since it adds vectorially to the propagation speed along the path, but for a storm at a few kilometres it is still small compared with the uncertainty in your count.

Worked example 5.12. Given t = 5 s, T = 20 °C; find d (distance to the strike). *Five seconds at 20 °C → 1.717 km.*

5.13 Speed of Sound in Air

(Science 10-37) $$v = 331.3 + 0.606\,T_C$$

Where,
v = Speed of sound (m/s)
T = Air temperature (°C)

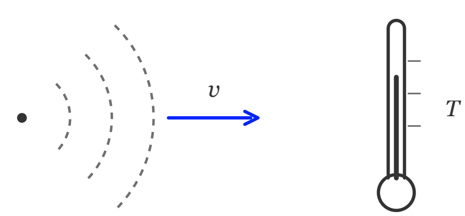

Figure 5.13 — speed of sound in air.

Sound travels by molecular collisions, and warmer molecules move faster, passing the disturbance along more quickly. In dry air the speed starts at 331.3 m/s at 0 °C and gains about 0.606 m/s per degree: 343 m/s in a 20 °C classroom, 349 m/s on a 30 °C summer day, but only 325 m/s in −10 °C winter air. The linear formula is an excellent approximation of the true √T dependence across everyday temperatures.

You use this physics every thunderstorm: light arrives essentially instantly, so counting seconds between flash and rumble and dividing by three gives the storm's distance in kilometres (sound covers about 1 km every 2.9 s at 20 °C). The temperature dependence has audible consequences too — sound bends toward where it travels slower, so on a cold, still morning with warm air above a chilled ground layer, wavefronts curve back down toward the earth and a distant train sounds uncannily close. Orchestras feel it as well: a flute's pitch rides on v, drifting sharp as the concert hall warms through the evening.

Worked example 5.13. Given T = 20 °C; find v (speed of sound). *20 C air → v = 343.42 m/s.*

Module 5 practice

58. Pressure Is Force over Area — A shipping crate presses on the warehouse floor with 600 N over a base of 0.25 m². *Calculate the pressure on the floor.*

59. Pressure Is Force over Area — A shipping crate presses on the warehouse floor with 600 N over a base of 0.25 m². *Calculate the pressure on the floor.*

60. Pressure Under Water — A diving ring rests on a pool floor 6 m deep. *Determine the water pressure on the ring, in kilopascals.*

61. Pressure Under Water — A diving ring rests on a pool floor 5 m deep. *Determine the water pressure on the ring, in kilopascals.*

62. Thin Air — Pressure with Altitude — A research aircraft climbs from sea level — 100 kPa on the ground today — to a cruising height of 11,000 m. *Estimate the air pressure up there, using the halving rule.*

63. Thin Air — Pressure with Altitude — A hiking trail gains 500 m from a trailhead sitting at 100 kPa. *Estimate the air pressure at the top, using the 12-kPa-per-kilometre rule.*

64. Relative Humidity — On a 29 °C afternoon, air can hold water vapour up to a saturation pressure of 4 kPa. Today's actual vapour pressure measures 1 kPa. *Calculate the relative humidity.*

65. Relative Humidity — On a 21 °C afternoon, air can hold water vapour up to a saturation pressure of 2.5 kPa. Today's actual vapour pressure measures 1.5 kPa. *Calculate the relative humidity.*

66. The Dew Point — An evening weather report gives 25 °C with 70 % relative humidity. *Determine the dew point, then read what the night will do.*

67. The Dew Point — An evening weather report gives 20 °C with 50 % relative humidity. *Determine the dew point, then read what the night will do.*

68. Feels-Like Temperatures — A sticky July afternoon in Windsor sits at 28 °C with a dew point of 21 °C, the air thick and still. *Determine how hot the afternoon actually feels.*

69. Feels-Like Temperatures — A January morning in Timmins reads -5 °C with a steady 30 km/h north wind. *Determine what the morning feels like on bare skin.*

70. Sunlight, Albedo and the Lapse Rate — A soccer pitch of summer grass has an albedo of 0.25, with 700 W/m² of sun landing on it. *Calculate the reflected sunlight.*

71. Sunlight, Albedo and the Lapse Rate — Open lake water — albedo 0.06 — catches 800 W/m² of afternoon sun. *Determine how much of that sunshine the lake reflects.*

72. Counting the Storm — Watching a 20 °C storm from a classroom window, a student sees lightning and counts 3 s before the thunder. *Determine how far away the lightning struck.*

73. Counting the Storm — Watching a 25 °C storm from a classroom window, a student sees lightning and counts 6 s before the thunder. *Determine how far away the lightning struck.*

74. Boss — The Weather Station — Storm duty at the school weather station, 28 °C and darkening. The gusts press on the station's 0.25 m^2 instrument plate with 600 N. The hygrometer reads 1.6 kPa of vapour against today's 4 kPa saturation ceiling. Then the sky flashes, and you count 6 s to the thunder. Work each line — every answer feeds the next. *Determine the plate pressure, the humidity, the dew point and the storm's distance — one line at a time.*

Answer key

1. 6 m/s
2. 9 m/s
3. 60 m
4. 150 m
5. 10 s
6. 20 s
7. 90 km/h
8. 10 m/s
9. 220 km
10. 180 km
11. 3 h
12. 8 m/s
13. 5 m
14. 1.5 (no unit)
15. 2.5 $\times 10^8$ m/s
16. 28.9 °
17. 19.5 °
18. 31.8 °
19. 41.8 °
20. 45 cm
21. 30 cm
22. 10 cm
23. 12 cm
24. -0.4 (no unit)
25. 15 cm
26. 5 D
27. 40 cm
28. 20 cm
29. 4 D
30. 252 kJ
31. 504 kJ
32. 10 C°
33. 20 C°
34. 1,130 kJ
35. 334 kJ
36. 334 kJ
37. 588 kJ
38. 60 mm
39. 36 mm
40. 7,200 W
41. 30 W
42. 336 kJ
43. 2 kg
44. 250 C
45. 6 A
46. 12 A
47. 3 A
48. 14 V
49. 5 Ω
50. 6 Ω
51. 6 Ω
52. 4 Ω
53. 40 W
54. 360 W
55. 8 kWh
56. 9 kWh
57. 12 Ω
58. 2,400 Pa
59. 2,400 Pa
60. 58.8 kPa
61. 49 kPa
62. 25 kPa
63. 94 kPa
64. 25 %
65. 60 %
66. 19.1 °C
67. 9.3 °C
68. 36.4 (humidex)
69. -13 °C
70. 175 W/m^2
71. 48 W/m^2
72. 343.4 m/s
73. 346.5 m/s
74. 2,400 Pa