Chapter 0

High-School Math Review

This short chapter is a compact reference, not a re-teaching of the material: it collects the facts and formulas you are expected to already know from high school, and which the rest of the course will use freely and without re-deriving. If everything here looks familiar, skip ahead. If something doesn't, this is the place to patch the gap *before* it surfaces awkwardly in the middle of a later proof.

0.1 Algebraic manipulation

Factoring and expanding

$$\begin{aligned} (a+b)^2 &= a^2 + 2ab + b^2 & (a-b)^2 &= a^2 - 2ab + b^2 \\ a^2 - b^2 &= (a-b)(a+b) & a^3 - b^3 &= (a-b)(a^2 + ab + b^2) \\ a^3 + b^3 &= (a+b)(a^2 - ab + b^2) & (a+b)^3 &= a^3 + 3a^2b + 3ab^2 + b^3 \end{aligned}$$

For a quadratic $ax^2 + bx + c$ $(a \neq 0)$ with roots x_1, x_2 (see below), $ax^2 + bx + c = a(x - x_1)(x - x_2)$.

Exponent rules

For $a, b > 0$ and $m, n \in \mathbb{R}$:

$$\begin{aligned} a^m a^n &= a^{m+n}, & \frac{a^m}{a^n} &= a^{m-n}, & (a^m)^n &= a^{mn}, \\ (ab)^n &= a^n b^n, & a^{-n} &= \frac{1}{a^n}, & a^{1/n} &= \sqrt[n]{a}. \end{aligned}$$

Absolute value

$$|x| = \begin{cases} x & x \geq 0 \\ -x & x < 0 \end{cases}$$

Key properties: $|x| \geq 0$; $|xy| = |x||y|$; $|x+y| \leq |x| + |y|$ (the **triangle inequality**); and for $r > 0$, $|x| < r \iff -r < x < r$.

Rational expressions

To add fractions, use a common denominator: $\frac{a}{b} + \frac{c}{d} = \frac{ad+bc}{bd}$. To simplify, factor the numerator and denominator and cancel common factors — never cancel terms that are only added or subtracted, only common *factors* of the whole numerator and denominator.

Exercise 0.1: Factor $x^2 - 5x + 6$, and simplify $\dfrac{x^2-9}{x^2-x-6}$.

0.2 Equations and inequalities

Solving quadratics

The equation $ax^2 + bx + c = 0$ ($a \neq 0$) has solutions

$$x = \frac{-b \pm \sqrt{b^2 - 4ac}}{2a},$$

provided the **discriminant** $\Delta = b^2 - 4ac$ is ≥ 0 (if $\Delta < 0$, there are no real solutions; if $\Delta = 0$, there is a single repeated solution). Equivalently, one can **complete the square**: $ax^2 + bx + c = a\left(x + \frac{b}{2a}\right)^2 + \left(c - \frac{b^2}{4a}\right)$.

Sign analysis for inequalities

To solve an inequality like $\dfrac{(x-1)(x+2)}{x-3} \geq 0$: find the values where each factor is zero or undefined ($x = 1, -2, 3$ here), mark them on the real line, and determine the sign of the whole expression on each resulting interval by checking one test point per interval (or by counting sign changes). Remember: multiplying or dividing an inequality by a *negative* number flips its direction.

Exercise 0.2: Solve $2x^2 - 3x - 2 = 0$, and solve the inequality $x^2 - 4 > 0$.

0.3 Functions and graphs

This is deliberately light — Chapter 4 covers functions rigorously and in full; the point here is only to jog your memory before we get there.

Lines and parabolas

A line through (x_0, y_0) with slope m is $y - y_0 = m(x - x_0)$; in slope-intercept form, $y = mx + q$. Two lines with slopes m_1, m_2 are parallel iff $m_1 = m_2$, and perpendicular iff $m_1 m_2 = -1$. A parabola $y = a(x-h)^2 + k$ has vertex (h, k), opening upward if $a > 0$ and downward if $a < 0$.

Reading a graph

For a function f, the **domain** is the set of valid inputs x, and the **range** is the set of resulting outputs $f(x)$; on a graph, these are read off as the horizontal and vertical extents of the curve, respectively. Given the graph of f, the graph of $f(x) + k$ is a vertical shift, $f(x - h)$ a horizontal shift, $-f(x)$ a reflection across the x-axis, and $cf(x)$ a vertical stretch/compression by a factor of c.

0.4 Exponentials and logarithms

Laws of logarithms

For a base $a > 0$, $a \neq 1$, and $x, y > 0$:

$$\log_a(xy) = \log_a x + \log_a y, \quad \log_a\left(\frac{x}{y}\right) = \log_a x - \log_a y,$$
$$\log_a(x^r) = r \log_a x, \qquad \log_a x = \frac{\log_b x}{\log_b a} \text{ (change to base } b > 0,\ b \neq 1).$$

Exponentials and logarithms are inverse to one another: $\log_a(a^x) = x$ and $a^{\log_a x} = x$.

Exercise 0.3: Solve $2^x = 32$, and solve $\log_2(x) + \log_2(x-2) = 3$.

0.5 Trigonometry

The unit circle

An angle θ measured in **radians** is the length of the arc it subtends on the unit circle: a full turn is 2π radians ($= 360^\circ$), so $\theta_{\text{rad}} = \theta_{\text{deg}} \cdot \frac{\pi}{180}$. For a point on the unit circle at angle θ, $\cos\theta$ and $\sin\theta$ are its x- and y-coordinates, respectively, and $\tan\theta = \sin\theta / \cos\theta$.

θ	0	$\frac{\pi}{6}$	$\frac{\pi}{4}$	$\frac{\pi}{3}$	$\frac{\pi}{2}$	π
$\sin\theta$	0	$\frac{1}{2}$	$\frac{\sqrt{2}}{2}$	$\frac{\sqrt{3}}{2}$	1	0
$\cos\theta$	1	$\frac{\sqrt{3}}{2}$	$\frac{\sqrt{2}}{2}$	$\frac{1}{2}$	0	-1
$\tan\theta$	0	$\frac{\sqrt{3}}{3}$	1	$\sqrt{3}$	undefined	0

Pythagorean identity

$$\sin^2\theta + \cos^2\theta = 1.$$

0.6 Coordinate geometry

Distance, midpoint, and circles

For points (x_1, y_1) and (x_2, y_2):

$$\text{distance} = \sqrt{(x_2 - x_1)^2 + (y_2 - y_1)^2}, \qquad \text{midpoint} = \left(\frac{x_1 + x_2}{2}, \frac{y_1 + y_2}{2}\right).$$

A circle of radius r centered at (h, k) has equation $(x-h)^2 + (y-k)^2 = r^2$.

0.7 Sets, intervals, and logical notation

A quick glossary for the light notation we'll start using immediately in Chapter 2 (which develops it far more carefully): for $a < b$, the **closed interval** $[a, b] = \{x \in \mathbb{R} \mid a \le x \le b\}$, the **open interval** $(a, b) = \{x \in \mathbb{R} \mid a < x < b\}$, and half-open intervals $[a, b), (a, b]$ combine the two conventions. The symbols $\forall$ ("for all") and $\exists$ ("there exists") will appear throughout the course to state facts precisely, e.g. "$\forall x \in \mathbb{R},\ x^2 \ge 0$."

0.8 Reading and writing rigorous proofs

A quick orientation to how mathematics is written, before you see many proofs in earnest.

Direct proof vs. proof by contradiction

A **direct proof** of a statement "if P then Q" assumes P and shows, through a chain of valid deductions, that Q must follow. A **proof by contradiction** instead assumes the *negation* of what you want to prove, and derives a logical contradiction (an absurdity, such as a number being both even and odd, or $0 = 1$) — since a contradiction can never actually hold, the assumption must have been false, so the original statement is true. It's especially useful for proving something *cannot* happen (nonexistence, irrationality, uniqueness, etc.), where a direct proof has nothing to build toward.

Example 0.1: **$\sqrt{2}$ is irrational.** Suppose, for contradiction, that $\sqrt{2}$ is rational: $\sqrt{2} = p/q$ with $p, q \in \mathbb{N}_+$ sharing no common factor (the fraction is in lowest terms). Squaring gives $2q^2 = p^2$, so p^2 is even, hence p itself is even (say $p = 2k$). Substituting, $2q^2 = 4k^2$, so $q^2 = 2k^2$ is also even, hence q is even too. But then p and q are both even, contradicting the assumption that they share no common factor. This contradiction shows the original assumption — that $\sqrt{2}$ is rational — must be false, so $\sqrt{2}$ is irrational.

Theorems, propositions, lemmas, and corollaries

These labels all mark a statement that has been rigorously proved, but signal its role in the text: a **theorem** is a major, standalone result; a **proposition** is a solid result, usually less central than a theorem; a **lemma** is a smaller "helper" result, proved mainly because it's needed in the proof of something else (typically the theorem right after it); and a **corollary** is a quick consequence that follows easily from a theorem or proposition just proved, with little extra work. The distinctions are about emphasis and narrative role, not about how rigorously each one is proved — all four are proved with exactly the same standard of rigor.

Exercise 0.4: Prove, by contradiction, that there is no largest natural number.

Chapter 1

Motivational Examples

Why does any of what we will be doing this semester matter? Before we build the formal machinery of the course — sets, sequences, limits, derivatives, integrals — it's worth pausing to see, informally and without proof, why that machinery is worth building in the first place. Every example below will be made completely rigorous later in the course; for now, just enjoy the ideas.

1.1 Velocity and acceleration

Suppose an object moves along a line, and its position at time t is $s(t)$. Over a time interval from t_0 to $t_0 + \Delta t$, its **average velocity** is

$$\frac{s(t_0 + \Delta t) - s(t_0)}{\Delta t}.$$

But what is its velocity at the *exact instant* t_0? The natural idea is to shrink Δt closer and closer to 0, and see what value the average velocity approaches. This is precisely a **limit** — the central object of Chapters 5-6 — and the value it approaches is, by definition, the **instantaneous velocity**, which is exactly what the **derivative** (Chapter 10) makes rigorous. Acceleration is then simply the derivative of velocity: the rate at which velocity itself is changing. This single idea — describing an instantaneous rate of change as a limit of average rates of change — is the seed from which all of differential calculus grows.

1.2 Computers and artificial intelligence

Modern machine learning may look nothing like a calculus course, but underneath, it is calculus. Training a neural network means tuning millions of numerical parameters to make the network's predictions as accurate as possible on some data. This is done by an algorithm called **gradient descent**: at each step, the algorithm computes the derivative of an "error" function with respect to every parameter, and nudges each parameter slightly in the direction that decreases the error. Doing this over and over, many times, is — quite literally — how large language models are trained. The single-variable version of this idea, "find where the derivative is zero to locate a minimum," is something we will do explicitly in Chapter 10.

There's another, less obvious, connection: every number a computer stores is only a finite approximation of a real number, and every computation a computer performs on it accumulates a small error. Understanding *precisely* what it means for a sequence of approximations to "get arbitrarily close" to a true value — which is exactly what a limit is (Chapter 5) — is what allows us to reason rigorously about when those computer approximations can be trusted.

1.3 Compound interest

Suppose you deposit €1 in a bank account with an annual interest rate of 100%, compounded once per year. After one year, you'd have $1 \cdot (1 + 1) = 2$, i.e. €2. But what if the bank compounded the interest *twice* a year, at half the rate each time? You'd end up with

$$\left(1 + \frac{1}{2}\right)^2 = 2.25,$$

i.e. €2.25. Compounding monthly gives $\left(1 + \frac{1}{12}\right)^{12} \approx 2.61$ (€2.61); daily gives $\left(1 + \frac{1}{365}\right)^{365} \approx$ 2.7146 (€2.7146). As the compounding frequency $n \to \infty$ (continuous compounding), the amount

$$\left(1 + \frac{1}{n}\right)^n$$

approaches a specific, famous number: $e \approx 2.71828\ldots$, the base of the natural logarithm. This is one of the cleanest ways to see a limit arise naturally from an extremely concrete, everyday question — and it is exactly how the number e will be properly defined in Chapter 5.

More generally, compound interest is the simplest example of **exponential growth**: unlike linear growth, where a quantity increases by a fixed *amount* each period, exponential growth increases by a fixed *percentage* each period — and the difference between the two compounds dramatically over time. A quick rule of thumb: at a fixed annual growth rate of r%, a quantity doubles in roughly $70/r$ years — a fact that itself is a direct consequence of logarithms (Chapter 4).

1.4 Exponential growth in a pandemic

Many of you experienced this next example first-hand. In the early days of the COVID-19 pandemic, the number of new cases each day grew by roughly the same *percentage* every day — exactly the same mathematical pattern as compound interest, just in a different context. A growth rate that sounds small, say 10% per day, still means cases roughly double every week; left unchecked, this is what allows case counts to explode from a handful to millions within a few months.

This is also why public dashboards during such events so often plotted case counts on a **logarithmic scale**: on such a plot, exponential growth appears as a straight line, making it easy to spot at a glance, and easy to see when growth is slowing down (the line bending) — a visual, intuitive first encounter with logarithms, which we will define rigorously in Chapter 4. And "flattening the curve," the public-health goal you likely heard repeated constantly, was really a statement about a *derivative*: not whether cases would keep rising, but how quickly — i.e., about the *rate* of change of the case count, not the case count itself.

Remark: The full mathematical description of how an epidemic spreads — accounting for the pool of susceptible, infected, and recovered individuals simultaneously — requires a **system of differential equations**. We won't get there in this course, but Chapter 15 will give you the tools to understand the single-equation building blocks such models are made of.

1.5 What's next: accumulation

Every example above was about *rates* of change — velocities, growth rates, training adjustments. There is a natural companion question: given a rate, can we recover the total *accumulated* quantity? If we know a car's velocity at every instant, can we recover the total distance traveled? If we know the daily rate of new cases, can we recover the total case count so far? If we know the instantaneous interest rate, can we recover the total amount earned?

The answer to all three is yes, and the tool that does it is the **integral** — informally, a way of adding up infinitely many infinitesimally small contributions, which turns out to be intimately related to the derivative via the Fundamental Theorem of Calculus (Chapter 12). We'll return to all of the examples above once we get there.