

About this course

This is a free self-study course outline written by Saving Planets. It is not a university course, it carries no academic credit, and no school is involved.

Its scope follows a one-semester junior numerical methods course for engineers. It is modeled on the topics and depth of MIT 18.330 Introduction to Numerical Analysis, MIT 2.086 Numerical Computation for Mechanical Engineers and MIT 10.34 (free on MIT OpenCourseWare), using free OpenStax calculus sections, which we link where they help. None of those authors, schools or agencies endorses this page.

What we add: every unit ties to real work with a cited public source (NOAA's Global Forecast System and supercomputers, the Mauna Loa CO₂ record, USGS groundwater models and NASA climate models), every number on the page is re-computed by a checking script before it goes live, and every lab has a free simulation or a spreadsheet.

Work through it in order. Each unit has goals, key ideas, equations, step-by-step examples, practice with full answers, common mistakes, flashcards and a quiz. Your progress is saved only in this browser.

How this course helps save the planet

Weather forecasts, climate projections and groundwater plans are all numerical methods running on computers: the equations cannot be solved by hand, so they are broken into grids, time steps and huge linear systems.

Knowing how those methods work, and how they fail, is how engineers judge a forecast or a model. Each unit ends with a short, sourced section. Values that are not from a source are labeled as sample or textbook values.

Before you start

Prerequisites

- Calculus I and II: derivatives, integrals and Taylor series.
- Differential Equations: first-order and second-order ODEs.
- Linear Algebra: matrices, determinants and eigenvalues.
- Intro Programming: loops, functions and arrays (a spreadsheet also works).

How to study

- Before computing, estimate the answer and decide what accuracy you need.
- Always compare against an exact answer on a test problem before trusting a method on a real one.
- Track the error as you halve the step: the ratio tells you the order of the method.
- Keep units in SI and write the step size with its units.
- Watch for subtracting nearly equal numbers; rewrite the formula if you can.
- Do every practice problem on paper or in a spreadsheet before opening the answer, and retake the quizzes a few days later.

Lab equipment (for the wet labs)

- A computer or tablet for the free PhET Least-Squares Regression, Curve Fitting, Calculus Grapher and Projectile Motion simulations.
- Spreadsheet software (any free one works), or Python if you know it.
- NOAA's Mauna Loa CO₂ data file in a web browser.
- A calculator.

Suggested 15-week schedule

About 6 to 9 hours a week: reading, practice problems and, in lab weeks, the lab.

Week	Unit	Topics	Reading
1	Errors, Floating-Point Numbers and Taylor Series	Errors, floating point, Taylor series; Lab A	MIT 18.330
2	Roots of Equations: Bisection, Newton and Secant	Bisection and fixed-point iteration	MIT 18.330
3	Roots of Equations: Bisection, Newton and Secant	Newton and secant methods; Lab B	OpenStax Calculus 1, 4.9
4	Linear Systems: Gaussian Elimination, LU and Conditioning	Gaussian elimination, LU and conditioning	MIT 2.086
5	Iterative Solvers: Jacobi, Gauss–Seidel and Convergence	Jacobi, Gauss–Seidel and sparse systems	MIT 18.330
6	Interpolation, Splines and Least-Squares Fitting	Interpolation and splines	MIT 18.330
7	Interpolation, Splines and Least-Squares Fitting	Least squares; Lab C; midterm practice	NOAA GML
8	Numerical Differentiation	Numerical differentiation	MIT 18.330

Week	Unit	Topics	Reading
9	Numerical Integration	Numerical integration; Lab D	OpenStax Calculus 2, 3.6
10	Ordinary Differential Equations: Euler, Runge–Kutta and Stability	Euler, Heun and RK4	MIT 10.34
11	Ordinary Differential Equations: Euler, Runge–Kutta and Stability	Stability, stiffness and systems; Lab E	MIT 10.34
12	Partial Differential Equations: Finite Differences	Finite differences for the heat equation	MIT 2.086
13	Partial Differential Equations: Finite Differences	Laplace, waves and CFL; Lab F	NOAA EMC
14	Optimization, Eigenvalues and Computing at Scale	Optimization and the power method	MIT 18.330
15	Optimization, Eigenvalues and Computing at Scale	Computing at scale, verification and validation; final practice	NOAA; NASA GISS

Your progress

Tick units and labs as you finish them, and take the quizzes. Progress and quiz scores are saved only in this browser, never sent to us.

Errors, Floating-Point Numbers and Taylor Series

Absolute and relative error, significant digits, IEEE 754 floating-point numbers, machine epsilon, rounding and cancellation, and truncation error from Taylor series.

Learning objectives

By the end of this unit you can:

- Compute absolute and relative error.
- Explain machine epsilon for single and double precision.
- Spot and fix catastrophic cancellation.
- Bound truncation error with the next Taylor term.

Key concepts

- Absolute error is $|\text{approximate} - \text{exact}|$; relative error divides by $|\text{exact}|$. Relative error tells you how many digits are right.
- Computers store numbers in binary floating point. Double precision has machine epsilon $2^{-52} \approx 2.22 \times 10^{-16}$ (about 16 digits); single precision has $2^{-23} \approx 1.19 \times 10^{-7}$ (about 7 digits).
- Most decimals, like 0.1, have no exact binary form, so $0.1 + 0.2 - 0.3$ comes out as about 5.55×10^{-17} instead of 0.
- Subtracting nearly equal numbers loses digits (catastrophic cancellation). Rewriting the formula algebraically often fixes it.
- Truncation error comes from cutting off an infinite process, such as a Taylor series. The first dropped term estimates the error.

Key equations

$$e_{rel} = \frac{|\tilde{x} - x|}{|x|}$$

Relative error

$$\varepsilon_{mach} = 2^{-52} \approx 2.22 \times 10^{-16}$$

Double-precision epsilon

$$f(x + h) = f(x) + hf'(x) + \frac{h^2}{2}f''(x) + \dots$$

Taylor series

$$\sqrt{x+1} - \sqrt{x} = \frac{1}{\sqrt{x+1} + \sqrt{x}}$$

Avoiding cancellation

Worked examples, step by step

Example 1.1 Find the relative error of $22/7$ as an approximation to π .

1. $|22/7 - \pi| = 0.001264$.
2. Divide by π .

Answer: Relative error = 4.02×10^{-4} , so about 3 correct digits.

Example 1.2 Compute $\sqrt{(10^8 + 1)} - \sqrt{(10^8)}$ safely.

1. Direct subtraction cancels most digits.
2. Use $1/(\sqrt{(10^8 + 1)} + \sqrt{(10^8)})$.

Answer: 5.0×10^{-5} (more exactly $4.99999999875 \times 10^{-5}$).

Example 1.3 Estimate the error of $e^{0.1} \approx 1 + 0.1$.

1. The next term is $0.1^2/2 = 0.005$.
2. The true error is $e^{0.1} - 1.1$.

Answer: Error = 0.00517, close to the 0.005 estimate.

Practice problems

Work each on paper first, then check.

1. How many correct digits does double precision give?

About 16 ($\epsilon = 2.22 \times 10^{-16}$).

2. Add the $0.1^3/6$ term to $e^{0.1} \approx 1.1$. New error?

1.71×10^{-4} .

3. Why is $0.1 + 0.2 - 0.3$ not exactly 0 on a computer?

0.1 and 0.2 have no exact binary form; the result is about 5.55×10^{-17} .

4. What is single-precision epsilon?

$$2^{-23} = 1.19 \times 10^{-7}.$$

Common misconceptions

Myth: Computers do exact arithmetic.

Actually: Every operation rounds to the nearest floating-point number.

Myth: Smaller steps always mean smaller error.

Actually: Rounding error grows as steps shrink.

Myth: Relative and absolute error are the same.

Actually: Relative error divides by the true size.

Flashcards

Absolute error

$$|\text{approx} - \text{exact}|.$$

Relative error

$$\text{Absolute error} / |\text{exact}|.$$

Machine epsilon

$$2^{-52} \approx 2.22 \times 10^{-16} \text{ in double.}$$

Cancellation

Digit loss when subtracting close numbers.

Truncation error

From cutting off an infinite process.

Single precision

About 7 digits.

Self-check quiz

1. Double-precision epsilon is about...

$$2.2 \times 10^{-16}$$

$$10^{-7}$$

$$10^{-3}$$

0

2.2×10^{-16} , 2^{-52} .

2. Subtracting nearly equal numbers causes...

cancellation

overflow

nothing

exact results

cancellation. Lost digits.

3. Truncation error comes from...

cutting off a series

binary storage

typing errors

units

cutting off a series. Definition.

4. Relative error divides by...

the true value

the step

epsilon

10

the true value. Definition.

Try this quiz again

How this helps save the planet

NOAA's Environmental Modeling Center describes a coupled global forecast run at about 3 km resolution that ran its atmospheric dynamics in 32-bit precision and everything else in 64-bit precision, a deliberate trade between speed and accuracy.

That is this unit's epsilon choice in practice: single precision (about 7 digits) is faster and uses half the memory, so modelers use it where the error it adds stays below the error the model already has.

On Saving Planets: [Earth Observation](#) · [Intro to Programming](#)

Sources: NOAA Environmental Modeling Center: Global Forecast System (GFS) (https://www.emc.ncep.noaa.gov/emc/pages/numerical_forecast_systems/gfs.php)

Free learning resources

- [MIT OpenCourseWare 18.330 Introduction to Numerical Analysis \(Spring 2012\)](https://ocw.mit.edu/courses/18-330-introduction-to-numerical-analysis-spring-2012/) (<https://ocw.mit.edu/courses/18-330-introduction-to-numerical-analysis-spring-2012/>): Free MIT lectures.
- [MIT OpenCourseWare 2.086 Numerical Computation for Mechanical Engineers \(Fall 2014\)](https://ocw.mit.edu/courses/2-086-numerical-computation-for-mechanical-engineers-fall-2014/) (<https://ocw.mit.edu/courses/2-086-numerical-computation-for-mechanical-engineers-fall-2014/>): Free MIT course.



I finished this unit

Roots of Equations: Bisection, Newton and Secant

Bracketing with bisection, Newton's method, the secant method, fixed-point iteration, convergence order and stopping rules.

Learning objectives

By the end of this unit you can:

- Find a root by bisection and predict the number of steps.
- Apply Newton's method and recognize quadratic convergence.
- Use the secant method when the derivative is unknown.
- Choose a stopping rule.

Key concepts

- Bisection halves a bracket $[a, b]$ where f changes sign. It always converges, but only one binary digit per step: $n \geq \log_2((b - a)/\text{tol})$.
- Newton's method follows the tangent line: $x_{n+1} = x_n - f(x_n)/f'(x_n)$. Near a simple root it converges quadratically, doubling the correct digits each step.
- The secant method replaces f' with the slope through the last two points. It needs no derivative and converges with order about 1.618.
- Newton can fail if f' is near zero or the start is far away; a bracketing fallback makes it safe.
- Stop when $|x_{n+1} - x_n|$ and $|f(x_n)|$ are both below tolerance, and always cap the number of iterations.

Key equations

$$x_{n+1} = x_n - \frac{f(x_n)}{f'(x_n)}$$

Newton's method

$$x_{n+1} = x_n - f(x_n) \frac{x_n - x_{n-1}}{f(x_n) - f(x_{n-1})}$$

Secant method

$$n \geq \log_2 \frac{b - a}{\text{tol}}$$

Bisection steps

$$|e_{n+1}| \approx C|e_n|^2$$

Worked examples, step by step

Example 2.1 Use bisection on $f(x) = x^3 - x - 2$ on $[1, 2]$. Give the first three midpoints and the root.

1. $f(1) = -2, f(2) = 4$.
2. Midpoints 1.5, 1.75, 1.625.

Answer: Root $x = 1.5214$.

Example 2.2 How many bisection steps shrink a width-1 bracket below 10^{-6} ?

1. $\log_2(10^6) = 19.93$.

Answer: 20 steps.

Example 2.3 Use Newton's method on $x^2 - 2$ from $x_0 = 1$.

1. $x_1 = 1.5, x_2 = 1.41667$.
2. $x_3 = 1.414216$.

Answer: $x_3 = 1.414216$, already 6 correct digits.

Practice problems

Work each on paper first, then check.

1. Apply Newton to $x^3 - x - 2$ from $x_0 = 1.5$. First step?

1.52174.

2. Apply the secant method to $x^2 - 2$ from 1 and 2. First three iterates?

1.3333, 1.4000, 1.4146.

3. What convergence order does the secant method have?

About 1.618.

4. Why cap the iteration count?

A method may diverge or cycle.

Common misconceptions

Myth: Newton always converges.

Actually: It can diverge from a poor start.

Myth: Bisection is fast.

Actually: It gains one bit per step.

Myth: A small $f(x)$ means x is close to the root.

Actually: Not if f is very flat.

Flashcards

Bisection

Halve a sign-change bracket.

Newton

$x - f/f'$.

Secant

Newton with a two-point slope.

Quadratic convergence

Digits double each step.

Bracket

Interval with a sign change.

Stopping rule

Step and residual below tolerance.

Self-check quiz

1. Newton converges...

quadratically

linearly

never

cubically always

quadratically. Near a simple root.

2. Bisection needs...

- a sign change
- a derivative
- a matrix
- nothing
- a sign change. Bracket.

3. Secant avoids...

- the derivative
- f
- iteration
- starting points
- the derivative. Slope from two points.

4. Steps to 10^{-6} from width 1 by bisection...

- 20
- 6
- 100
- 2
- $20 \cdot \log_2 10^6$.

[Try this quiz again](#)

How this helps save the planet

The USGS says MODFLOW is considered an international standard for simulating groundwater. One of its versions, MODFLOW-NWT, uses a Newton-Raphson formulation to improve the solution of

unconfined groundwater-flow problems.

Unconfined aquifers make the equations nonlinear, because the saturated thickness depends on the water level being solved for. Newton's method from this unit is what makes those runs converge reliably.

On Saving Planets: [Oceans](#) · [Calculus I](#)

Sources: [USGS: MODFLOW and Related Programs \(https://www.usgs.gov/mission-areas/water-resources/science/modflow-and-related-programs\)](https://www.usgs.gov/mission-areas/water-resources/science/modflow-and-related-programs)

Free learning resources

- [OpenStax Calculus Volume 1, 4.9 Newton's method \(https://openstax.org/books/calculus-volume-1/pages/4-9-newtons-method\)](https://openstax.org/books/calculus-volume-1/pages/4-9-newtons-method): Free textbook section.
- [MIT OpenCourseWare 18.330 Introduction to Numerical Analysis \(Spring 2012\) \(https://ocw.mit.edu/courses/18-330-introduction-to-numerical-analysis-spring-2012/\)](https://ocw.mit.edu/courses/18-330-introduction-to-numerical-analysis-spring-2012/): Free MIT lectures.



I finished this unit

Linear Systems: Gaussian Elimination, LU and Conditioning

Gaussian elimination with partial pivoting, LU factorization, operation counts, the condition number and how it limits accuracy.

Learning objectives

By the end of this unit you can:

- Solve a 3×3 system by elimination.
- Explain LU factorization and reuse.
- Estimate the cost of elimination.
- Use the condition number to judge accuracy.

Key concepts

- Gaussian elimination reduces $Ax = b$ to an upper-triangular system, then back-substitutes. Partial pivoting swaps rows so the largest pivot is used, which controls rounding error.
- LU factorization stores the elimination as $A = LU$. Once factored, each new right-hand side costs only two triangular solves.
- Elimination costs about $2n^3/3$ floating-point operations: doubling n multiplies the work by 8.
- The condition number $\kappa(A) = \|A\|\|A^{-1}\|$ measures sensitivity. You can lose about $\log_{10} \kappa$ digits of accuracy.
- Never compute A^{-1} to solve a system; factor and solve instead.

Key equations

$$A = LU$$

LU factorization

$$\text{cost} \approx \frac{2}{3}n^3$$

Elimination flops

$$\kappa(A) = \|A\|\|A^{-1}\|$$

Condition number

$$\frac{\|\delta x\|}{\|x\|} \leq \kappa \frac{\|\delta b\|}{\|b\|}$$

Error bound

Worked examples, step by step

Example 3.1 Solve $2x + y - z = 8$, $-3x - y + 2z = -11$, $-2x + y + 2z = -3$.

1. Eliminate x from rows 2 and 3, then y from row 3.
2. Back-substitute.

Answer: $x = 2$, $y = 3$, $z = -1$.

Example 3.2 Estimate the elimination cost for $n = 1000$.

1. $2 \times 1000^3/3$.

Answer: About 6.67×10^8 flops.

Example 3.3 Find the 2-norm condition number of $\begin{bmatrix} 1 & 1 \\ 1 & 1.0001 \end{bmatrix}$.

1. The rows are nearly parallel.
2. $\kappa = \sigma_{\max}/\sigma_{\min}$.

Answer: $\kappa \approx 40,002$, so about 4–5 digits can be lost.

Practice problems

Work each on paper first, then check.

1. How much more work is $n = 2000$ than $n = 1000$?

8 times.

2. Elimination flops for $n = 3$?

$2 \times 27/3 = 18$.

3. Why pivot?

To avoid dividing by small numbers and amplifying rounding.

4. Digits lost with $\kappa = 10^8$ in double precision?

About 8 of 16.

Common misconceptions

Myth: Use A^{-1} to solve systems.

Actually: Factor and solve; it is cheaper and more accurate.

Myth: A small determinant means ill-conditioned.

Actually: Use the condition number.

Myth: Pivoting is optional.

Actually: Without it, elimination can fail or lose accuracy.

Flashcards

Gaussian elimination

Reduce to triangular form.

LU

$A = LU$.

Pivoting

Row swaps for stability.

Condition number

$\|A\| \|A^{-1}\|$.

Back substitution

Solve triangular systems.

Cost

$\approx 2n^3/3$.

Self-check quiz

1. Elimination costs about...

$2n^3/3$

n

n^2

2^n

$2n^3/3$. Triple loop.

2. LU helps when...

many right-hand sides

$n = 1$

no solution

never

many right-hand sides. Reuse factors.

3. Large κ means...

sensitive

easy

singular always

symmetric

sensitive. Amplifies errors.

4. Partial pivoting picks...

largest pivot

smallest

zero

random

largest pivot. Stability.

[Try this quiz again](#)

How this helps save the planet

The USGS says a key feature of MODFLOW 6 is its ability to solve multiple tightly coupled numerical models in a single system of equations.

Its map of depth to the water table across the continental United States came from a parallel simulation of 256 tightly coupled model domains on 256 cores of a USGS supercomputer. Each

domain is a large sparse linear system like the ones in this unit.

On Saving Planets: [Linear Algebra](#) · [Oceans](#)

Sources: [USGS: MODFLOW and Related Programs \(https://www.usgs.gov/mission-areas/water-resources/science/modflow-and-related-programs\)](https://www.usgs.gov/mission-areas/water-resources/science/modflow-and-related-programs)

Free learning resources

- [MIT OpenCourseWare 2.086 Numerical Computation for Mechanical Engineers \(Fall 2014\) \(https://ocw.mit.edu/courses/2-086-numerical-computation-for-mechanical-engineers-fall-2014/\)](https://ocw.mit.edu/courses/2-086-numerical-computation-for-mechanical-engineers-fall-2014/): Free MIT course.
- [MIT OpenCourseWare 10.34 Numerical Methods Applied to Chemical Engineering \(Fall 2015\) \(https://ocw.mit.edu/courses/10-34-numerical-methods-applied-to-chemical-engineering-fall-2015/\)](https://ocw.mit.edu/courses/10-34-numerical-methods-applied-to-chemical-engineering-fall-2015/): Free MIT course.



I finished this unit

Iterative Solvers: Jacobi, Gauss–Seidel and Convergence

Jacobi and Gauss–Seidel iteration, diagonal dominance, the iteration matrix and spectral radius, sparse systems and why iterative methods scale.

Learning objectives

By the end of this unit you can:

- Carry out Jacobi and Gauss–Seidel steps.
- Test diagonal dominance.
- Use the spectral radius to predict convergence speed.
- Explain why large sparse systems use iteration.

Key concepts

- Jacobi updates every unknown from the previous iterate; Gauss–Seidel uses new values as soon as they are computed.
- If A is strictly diagonally dominant, both methods converge from any start.
- The error shrinks each step by about the spectral radius ρ of the iteration matrix. Gauss–Seidel often has $\rho_{\text{GS}} = \rho_J^2$.
- Large grids give sparse matrices with only a few nonzeros per row. Iterative methods store only those nonzeros, while elimination fills in.
- Stop when the residual $\|b - Ax\|$ is small relative to $\|b\|$.

Key equations

$$x_i^{(k+1)} = \frac{1}{a_{ii}} \left(b_i - \sum_{j \neq i} a_{ij} x_j^{(k)} \right)$$

Jacobi

$$|a_{ii}| > \sum_{j \neq i} |a_{ij}|$$

Diagonal dominance

$$\|e_k\| \approx \rho^k \|e_0\|$$

Convergence rate

$$r = b - Ax$$

Worked examples, step by step

Example 4.1 Do two Jacobi steps on $4x + y = 1$, $x + 3y = 2$ from $(0, 0)$.

1. Step 1: $x = 1/4 = 0.25$, $y = 2/3 = 0.6667$.
2. Step 2: $x = (1 - 0.6667)/4$, $y = (2 - 0.25)/3$.

Answer: $x = 0.0833$, $y = 0.5833$ (exact 0.0909, 0.6364).

Example 4.2 Do two Gauss–Seidel steps on the same system.

1. Step 1: $x = 0.25$, $y = (2 - 0.25)/3 = 0.5833$.
2. Step 2: $x = 0.1042$, $y = 0.6319$.

Answer: $x = 0.1042$, $y = 0.6319$, closer than Jacobi.

Example 4.3 Find the Jacobi and Gauss–Seidel spectral radii for that system.

1. $\rho_J = \sqrt{1/12}$.
2. $\rho_{GS} = \rho_J^2 = 1/12$.

Answer: $\rho_J = 0.289$; $\rho_{GS} = 0.0833$.

Practice problems

Work each on paper first, then check.

1. Is $\begin{bmatrix} 4 & 1 \\ 1 & 3 \end{bmatrix}$ diagonally dominant?

Yes: $4 > 1$ and $3 > 1$.

2. Steps for Gauss–Seidel ($\rho = 0.0833$) to cut error by 10^6 ?

About 6 ($0.0833^6 = 3.35 \times 10^{-7}$).

3. Third Gauss–Seidel iterate?

$x = 0.0920$, $y = 0.6360$.

4. Why not eliminate on a million-unknown grid?

Fill-in uses too much memory and time.

Common misconceptions

Myth: Iteration always converges.

Actually: It needs $\rho < 1$.

Myth: Jacobi and Gauss–Seidel are the same.

Actually: Gauss–Seidel uses new values immediately.

Myth: A small change between steps means converged.

Actually: Check the residual too.

Flashcards

Jacobi

Update all from old values.

Gauss–Seidel

Use new values immediately.

Spectral radius

Largest $|\text{eigenvalue}|$.

Diagonal dominance

$|a_{ii}| > \text{row sum}$.

Sparse

Mostly zeros.

Residual

$b - Ax$.

Self-check quiz

1. Converges if...

$\rho < 1$

$\rho > 1$

n is even

A is singular

$\rho < 1$. Error shrinks.

2. Gauss-Seidel uses...

newest values

old values only

no values

inverses

newest values. Definition.

3. Sparse means...

mostly zeros

dense

small n

singular

mostly zeros. Few nonzeros.

4. Sufficient for convergence...

diagonal dominance

symmetry only

$n = 2$

nothing

diagonal dominance. Theorem.

Try this quiz again

How this helps save the planet

NOAA's GFS divides the atmosphere into 127 vertical layers at about 13 km horizontal spacing. Each grid point couples only to its neighbours, so the systems such models build are huge and sparse.

With a textbook Earth surface area of $5.1 \times 10^8 \text{ km}^2$, a 13 km grid has about 3.0 million columns and 383 million cells. Systems that size are handled with iterative and other sparse methods, not plain elimination.

On Saving Planets: [Earth Observation](#) · [Linear Algebra](#)

Sources: NOAA Environmental Modeling Center: Global Forecast System (GFS) (https://www.emc.ncep.noaa.gov/emc/pages/numerical_forecast_systems/gfs.php)

Free learning resources

- [MIT OpenCourseWare 18.330 Introduction to Numerical Analysis \(Spring 2012\)](https://ocw.mit.edu/courses/18-330-introduction-to-numerical-analysis-spring-2012/) (<https://ocw.mit.edu/courses/18-330-introduction-to-numerical-analysis-spring-2012/>): Free MIT lectures.
- [MIT OpenCourseWare 2.086 Numerical Computation for Mechanical Engineers \(Fall 2014\)](https://ocw.mit.edu/courses/2-086-numerical-computation-for-mechanical-engineers-fall-2014/) (<https://ocw.mit.edu/courses/2-086-numerical-computation-for-mechanical-engineers-fall-2014/>): Free MIT course.



I finished this unit

Interpolation, Splines and Least-Squares Fitting

Linear and polynomial interpolation, Lagrange form, the danger of high-degree polynomials, cubic splines, and least-squares line and polynomial fits with residuals.

Learning objectives

By the end of this unit you can:

- Interpolate linearly between data points.
- Build a Lagrange polynomial.
- Fit a least-squares line and interpret its slope.
- Choose between interpolation and fitting.

Key concepts

- Interpolation passes exactly through the data; fitting passes near noisy data and minimizes the sum of squared residuals.
- A degree- n polynomial through $n + 1$ points is unique. High degrees on equally spaced points can oscillate wildly near the ends (Runge's phenomenon).
- Cubic splines join cubic pieces with matching slope and curvature, giving smooth curves without wild swings.
- The least-squares line solves the normal equations $A^T A c = A^T y$. The slope is a rate, with units.
- Plot the residuals: a pattern in them means the model is missing something, such as curvature.

Key equations

$$y = y_0 + \frac{x - x_0}{x_1 - x_0}(y_1 - y_0)$$

Linear interpolation

$$p(x) = \sum_i y_i \prod_{j \neq i} \frac{x - x_j}{x_i - x_j}$$

Lagrange form

$$A^T A c = A^T y$$

Normal equations

$$\min \sum_i r_i^2$$

Worked examples, step by step

Example 5.1 NOAA's Mauna Loa annual means are 390.10 ppm (2010) and 414.21 ppm (2020). Interpolate 2015.

$$1. y = 390.10 + 0.5 \times 24.11.$$

Answer: 402.16 ppm; NOAA's 2015 value is 401.01, so the record curves upward.

Example 5.2 Fit a least-squares line to the decade values 1960–2020 (316.91, 325.68, 338.76, 354.45, 369.71, 390.10, 414.21 ppm).

1. Solve the normal equations for slope and intercept.

Answer: Slope 1.613 ppm per year; the line predicts 415.0 ppm for 2025.

Example 5.3 Fit a quadratic in $t = \text{year} - 1960$ instead.

1. Least squares with columns 1, t , t^2 .

Answer: $c = 316.84 + 0.810t + 0.01338t^2$; it predicts 426.0 ppm for 2025, close to NOAA's 427.35.

Practice problems

Work each on paper first, then check.

1. Average growth 2010–2020 from the decade values?

2.41 ppm per year.

2. Average growth 2000–2010?

2.04 ppm per year.

3. Why does the straight line underpredict 2025?

Growth is speeding up; the residuals curve.

4. Interpolate or fit noisy lab data?

Fit.

Common misconceptions

Myth: Higher-degree polynomials always fit better.

Actually: They can oscillate between points.

Myth: A good R^2 means a good model.

Actually: Check the residual pattern.

Myth: Extrapolation is as safe as interpolation.

Actually: It is much riskier.

Flashcards

Interpolation

Exactly through the data.

Least squares

Minimize squared residuals.

Spline

Smooth piecewise cubic.

Runge

Wiggles of high-degree fits.

Residual

Data minus model.

Normal equations

$A^T A c = A^T y$.

Self-check quiz

1. Least squares minimizes...

sum of squared residuals

max error

slope

degree

sum of squared residuals. Definition.

2. Runge's phenomenon is...

end oscillation

cancellation

divergence

pivoting

end oscillation. High-degree polynomials.

3. Splines are...

piecewise cubics

single lines

random

Taylor series

piecewise cubics. Definition.

4. Normal equations are...

$$A^T A c = A^T y$$

$$A x = 0$$

$$L U$$

$$f = 0$$

$$A^T A c = A^T y. \text{ Least squares.}$$

[Try this quiz again](#)

How this helps save the planet

NOAA reports that the Mauna Loa record, started by C. David Keeling of Scripps in March 1958, is the longest record of direct CO₂ measurements in the atmosphere. The monthly mean was 426.07

ppm in September 2026, up from 424.37 ppm in September 2025.

NOAA's annual means rose from 316.91 ppm in 1960 to 427.35 ppm in 2025. The upward curve shows up in this unit's residuals: a straight line misses it, and a quadratic fits.

On Saving Planets: [Solutions](#) · [Earth Observation](#)

Sources: [NOAA Global Monitoring Laboratory: Trends in atmospheric CO₂ \(Mauna Loa\)](https://gml.noaa.gov/ccgg/trends/) (<https://gml.noaa.gov/ccgg/trends/>) · [NOAA GML: Mauna Loa annual mean CO₂ data](https://gml.noaa.gov/webdata/ccgg/trends/co2/co2_annmean_mlo.txt) (https://gml.noaa.gov/webdata/ccgg/trends/co2/co2_annmean_mlo.txt)

Free learning resources

- [PhET: Least-Squares Regression](https://phet.colorado.edu/en/simulations/least-squares-regression) (<https://phet.colorado.edu/en/simulations/least-squares-regression>): Fit lines interactively.
- [PhET: Curve Fitting](https://phet.colorado.edu/en/simulations/curve-fitting) (<https://phet.colorado.edu/en/simulations/curve-fitting>): Polynomial fits interactively.



I finished this unit

Numerical Differentiation

Forward, backward and central differences, their truncation error, the balance with rounding error and the best step size, second derivatives, and Richardson extrapolation.

Learning objectives

By the end of this unit you can:

- Apply forward and central differences.
- State the order of each formula.
- Choose a step size that balances truncation and rounding.
- Estimate a second derivative.

Key concepts

- The forward difference $(f(x + h) - f(x))/h$ has error proportional to h (first order).
- The central difference $(f(x + h) - f(x - h))/(2h)$ has error proportional to h^2 (second order), so halving h cuts the error by 4.
- Rounding error grows like ϵ/h as h shrinks. For forward differences the best h is about $\sqrt{\epsilon} \approx 1.49 \times 10^{-8}$; for central differences about $\epsilon^{1/3} \approx 6.06 \times 10^{-6}$.
- The second derivative is $(f(x + h) - 2f(x) + f(x - h))/h^2$, also second order.
- Richardson extrapolation combines two step sizes to cancel the leading error term.

Key equations

$$f'(x) \approx \frac{f(x + h) - f(x)}{h}$$

Forward difference

$$f'(x) \approx \frac{f(x + h) - f(x - h)}{2h}$$

Central difference

$$f''(x) \approx \frac{f(x + h) - 2f(x) + f(x - h)}{h^2}$$

Second difference

$$h_{opt} \approx \sqrt{\epsilon}$$

Best forward step

Worked examples, step by step

Example 6.1 Estimate the derivative of $\sin x$ at $x = 1$ with a forward difference, $h = 0.1$.

1. $(\sin 1.1 - \sin 1)/0.1$.

Answer: 0.4974; exact $\cos 1 = 0.5403$; error -0.0429 .

Example 6.2 Repeat with a central difference.

1. $(\sin 1.1 - \sin 0.9)/0.2$.

Answer: 0.5394; error -0.00090 , about 47.7 times smaller.

Example 6.3 Find the best forward-difference step in double precision.

1. $h \approx \sqrt{(2.22 \times 10^{-16})}$.

Answer: $h \approx 1.49 \times 10^{-8}$.

Practice problems

Work each on paper first, then check.

1. Halve h in a central difference. Error?

About $\frac{1}{4}$ as large.

2. Halve h in a forward difference. Error?

About $\frac{1}{2}$ as large.

3. Best central-difference step in double precision?

About $\varepsilon^{(1/3)} = 6.06 \times 10^{-6}$.

4. Why not use $h = 10^{-15}$?

Rounding error ε/h dominates.

Common misconceptions

Myth: Smaller h is always better.

Actually: Below the optimum, rounding wins.

Myth: Forward and central are equally accurate.

Actually: Central is second order.

Myth: Differentiating noisy data is easy.

Actually: Differences amplify noise; smooth or fit first.

Flashcards

Forward difference

First order.

Central difference

Second order.

Second difference

$f(x+h) - 2f(x) + f(x-h)$ over h^2 .

Optimal h

$\approx \sqrt{\epsilon}$ for forward.

Richardson

Cancel the leading error.

Order

Exponent of h in the error.

Self-check quiz

1. Central difference error $\propto \dots$

h^2

h

$1/h$

h^3 always

h^2 . Symmetric.

2. Forward difference is ... order.

first

second

zero

fourth

first. Error $\propto h$.

3. Best forward h is about...

$\sqrt{\epsilon}$

ϵ

1

0.1

$\sqrt{\epsilon}$. Balance.

4. Noise in data is ... by differencing.

amplified

removed

halved

unchanged

amplified. Divides by h .

Try this quiz again

How this helps save the planet

NOAA says the GFS produces hourly forecast output for the first 120 hours, then 3-hourly output for days 5 to 16. Forecasters and downstream models difference those outputs to get rates such as pressure tendency and rainfall rate.

This unit shows why the time step of saved output matters: a rate taken from 3-hourly values with a central difference is far more accurate than a one-sided difference over the same interval.

On Saving Planets: [GOES-19](#) · [Calculus I](#)

Sources: NOAA Environmental Modeling Center: Global Forecast System (GFS) (https://www.emc.ncep.noaa.gov/emc/pages/numerical_forecast_systems/gfs.php)

Free learning resources

- MIT OpenCourseWare 18.330 Introduction to Numerical Analysis (Spring 2012) (<https://ocw.mit.edu/courses/18-330-introduction-to-numerical-analysis-spring-2012/>): Free MIT lectures.
- PhET: Calculus Grapher (<https://phet.colorado.edu/en/simulations/calculus-grapher>): Derivatives interactively.



I finished this unit

Numerical Integration

Trapezoidal and Simpson's rules, composite rules and error orders, Romberg extrapolation, Gaussian quadrature, and integrating measured data.

Learning objectives

By the end of this unit you can:

- Apply the composite trapezoidal and Simpson's rules.
- Predict how error falls with n .
- Use Romberg to improve a result.
- Apply 2-point Gaussian quadrature.

Key concepts

- The trapezoidal rule joins points with straight lines. Its error falls like h^2 : doubling n cuts it by 4.
- Simpson's rule fits parabolas through pairs of intervals (n even). Its error falls like h^4 : doubling n cuts it by 16.
- Romberg extrapolation: $T_{2n} + (T_{2n} - T_n)/3$ cancels the h^2 error of the trapezoidal rule (it equals Simpson's rule).
- Gaussian quadrature picks the best points: 2 points at $\pm 1/\sqrt{3}$ integrate cubics exactly.
- For measured data at fixed spacing, use the trapezoidal or Simpson rule directly on the samples.

Key equations

$$T_n = h \left[\frac{1}{2}f_0 + f_1 + \cdots + f_{n-1} + \frac{1}{2}f_n \right]$$

Trapezoidal rule

$$S_n = \frac{h}{3} [f_0 + 4f_1 + 2f_2 + \cdots + 4f_{n-1} + f_n]$$

Simpson's rule

$$\int_{-1}^1 f \approx f\left(-\frac{1}{\sqrt{3}}\right) + f\left(\frac{1}{\sqrt{3}}\right)$$

2-point Gauss

$$E_T \propto h^2, \quad E_S \propto h^4$$

Error orders

Worked examples, step by step

Example 7.1 Estimate $\int_0^1 e^x dx$ with the trapezoidal rule, $n = 4$.

1. $h = 0.25$; sum the weighted values.

Answer: $T_4 = 1.72722$; exact $e - 1 = 1.71828$; error 0.00894 .

Example 7.2 Repeat with Simpson's rule, $n = 4$.

1. Weights 1, 4, 2, 4, 1 times $h/3$.

Answer: $S_4 = 1.718319$; error 3.70×10^{-5} .

Example 7.3 Use 2-point Gaussian quadrature on the same integral.

1. Map $[0, 1]$ to $[-1, 1]$: points $0.5 \pm 0.5/\sqrt{3}$, weight $\frac{1}{2}$ each.

Answer: 1.717896 ; error -3.85×10^{-4} with only 2 function values.

Practice problems

Work each on paper first, then check.

1. Trapezoidal error with $n = 8$?

0.00224 , about $\frac{1}{4}$ of the $n = 4$ error.

2. Romberg from T_4 and T_8 ?

1.718284 .

3. Simpson's rule with $n = 2$?

1.718861 .

4. Why must Simpson's n be even?

It uses pairs of intervals.

Common misconceptions

Myth: Simpson works for any n .

Actually: n must be even.

Myth: More points always help equally.

Actually: Higher-order rules gain much faster.

Myth: Gauss points can be anywhere.

Actually: They are fixed roots of Legendre polynomials.

Flashcards

Trapezoidal

Error $\propto h^2$.

Simpson

Error $\propto h^4$.

Romberg

Extrapolate trapezoid results.

Gauss quadrature

Optimal points and weights.

Composite rule

Apply on many subintervals.

Even n

Needed for Simpson.

Self-check quiz

1. Simpson error $\propto$...

h^4

h^2

h

1

h^4 . Fourth order.

2. Doubling n in the trapezoidal rule cuts error by...

4

2

16

1

4. h^2 .

3. 2-point Gauss is exact for...

cubics

all functions

lines only

nothing

cubics. Degree $2n - 1$.

4. Romberg cancels...

the leading error

all error

rounding

f

the leading error. Extrapolation.

Try this quiz again

How this helps save the planet

NOAA explains that every Mauna Loa monthly mean is the average of daily means, which are in turn based on hourly averages, using only hours with “background” conditions.

An average is an integral divided by the interval, so building monthly means from hourly values is numerical integration, the trapezoidal idea from this unit applied to real measurements.

On Saving Planets: [Calculus II](#) · [Solutions](#)

Sources: [NOAA Global Monitoring Laboratory: Trends in atmospheric CO₂ \(Mauna Loa\)](https://gml.noaa.gov/ccgg/trends/) (<https://gml.noaa.gov/ccgg/trends/>)

Free learning resources

- [OpenStax Calculus Volume 2, 3.6 Numerical integration](https://openstax.org/books/calculus-volume-2/pages/3-6-numerical-integration) (<https://openstax.org/books/calculus-volume-2/pages/3-6-numerical-integration>); Free textbook section.
- [PhET: Calculus Grapher](https://phet.colorado.edu/en/simulations/calculus-grapher) (<https://phet.colorado.edu/en/simulations/calculus-grapher>); Integrals interactively.



I finished this unit

Ordinary Differential Equations: Euler, Runge–Kutta and Stability

Euler's method, Heun's method, fourth-order Runge–Kutta, local and global error, stability limits, stiff equations and implicit methods, and systems of ODEs.

Learning objectives

By the end of this unit you can:

- Step an ODE with Euler and RK4.
- Relate step size to global error.
- Find the explicit stability limit for $y' = -\lambda y$.
- Recognize stiff problems.

Key concepts

- Euler's method $y_{n+1} = y_n + h f(t_n, y_n)$ has global error proportional to h (first order).
- Fourth-order Runge–Kutta (RK4) uses four slope samples per step; its global error falls like h^4 .
- For $y' = -\lambda y$, explicit Euler is stable only if $h < 2/\lambda$. Outside that, the numerical solution grows even though the true one decays.
- Stiff problems have very fast and very slow parts. Implicit methods such as backward Euler stay stable with large steps.
- Higher-order ODEs become first-order systems, which the same methods solve.

Key equations

$$y_{n+1} = y_n + h f(t_n, y_n)$$

Euler

$$y_{n+1} = y_n + \frac{h}{6}(k_1 + 2k_2 + 2k_3 + k_4)$$

RK4

$$h < \frac{2}{\lambda}$$

Euler stability for $y' = -\lambda y$

$$y_{n+1} = y_n + h f(t_{n+1}, y_{n+1})$$

Backward Euler

Worked examples, step by step

Example 8.1 Solve $y' = -y$, $y(0) = 1$ to $t = 1$ with Euler, $h = 0.1$.

1. Each step multiplies by 0.9.
2. $y(1) \approx 0.9^{10}$.

Answer: 0.34868 vs exact 0.36788; error -0.0192.

Example 8.2 Take one RK4 step of $h = 0.1$ on the same problem.

1. $k_1 = -1$, $k_2 = -0.95$, $k_3 = -0.9525$, $k_4 = -0.90475$.

Answer: $y(0.1) = 0.9048375$ vs exact 0.9048374.

Example 8.3 A coffee cools by $T' = -0.1(T - 20)$ from 90°C (sample values). Use Euler with $h = 1$ min for 10 min.

1. Each step: $T - 20$ shrinks by 0.9.

Answer: 44.41°C vs exact 45.75°C .

Practice problems

Work each on paper first, then check.

1. Euler with $h = 0.05$ to $t = 1$ on $y' = -y$?

$0.95^{20} = 0.35849$, error about half.

2. Stability limit for $y' = -50y$ with Euler?

$h < 0.04$.

3. One Heun step of $h = 0.1$ on $y' = -y$?

0.905.

4. How do you solve $y'' = -y$ with these methods?

Write it as $y' = v$, $v' = -y$.

Common misconceptions

Myth: A smaller step is always stable.

Actually: Only below the limit, and rounding grows with many steps.

Myth: RK4 is exact.

Actually: Its error is small but not zero.

Myth: Implicit methods are always slower.

Actually: For stiff problems they are much faster overall.

Flashcards

Euler

First-order explicit.

RK4

Fourth order, 4 slopes.

Stability

Errors do not grow.

Stiff

Fast and slow scales together.

Backward Euler

Implicit, very stable.

System

Higher order $\rightarrow$ first order.

Self-check quiz

1. Euler global error $\propto \dots$

☐ h

☐ h^4

☐ h^2

☐ 1

☐ h . First order.

2. RK4 uses ... slopes.

☐ 4

1

2

10

4. k_1 to k_4 .

3. Euler stable for $y' = -\lambda y$ if...

$h < 2/\lambda$

$h > 2/\lambda$

any h

$\lambda < 0$

$h < 2/\lambda$. Amplification $|1 - h\lambda| < 1$.

4. Stiff problems favour...

implicit methods

Euler

bisection

Simpson

implicit methods. Stability.

Try this quiz again

How this helps save the planet

NOAA says the GFS runs four times a day and forecasts up to 16 days ahead. The EMC page also describes a 3 km coupled research run that took about 40,000 seconds of computer time for each 24-hour forecast, using 233 nodes.

That is about 11.1 hours of computing per forecast day, so the model ran only about 2.16 times faster than real time. Each forecast is a time-stepping ODE and PDE solve, and the stable step size sets the cost.

On Saving Planets: [Differential Equations](#) · [GOES-19](#)

Sources: NOAA Environmental Modeling Center: Global Forecast System (GFS) (https://www.emc.ncep.noaa.gov/emc/pages/numerical_forecast_systems/gfs.php)

Free learning resources

- [MIT OpenCourseWare 10.34 Numerical Methods Applied to Chemical Engineering \(Fall 2015\)](https://ocw.mit.edu/courses/10-34-numerical-methods-applied-to-chemical-engineering-fall-2015/) (<https://ocw.mit.edu/courses/10-34-numerical-methods-applied-to-chemical-engineering-fall-2015/>): Free MIT course.
- [PhET: Projectile Motion](https://phet.colorado.edu/en/simulations/projectile-motion) (<https://phet.colorado.edu/en/simulations/projectile-motion>): Compare with your solver.



I finished this unit

Partial Differential Equations: Finite Differences

Grids and stencils, the explicit heat equation and its stability limit, implicit and Crank–Nicolson schemes, Laplace’s equation and iterative solution, the CFL condition for waves, and the cost of resolution.

Learning objectives

By the end of this unit you can:

- Write finite-difference stencils for the heat and Laplace equations.
- Apply the explicit stability limit $r \leq \frac{1}{2}$.
- Explain the CFL condition.
- Estimate how cost grows with resolution.

Key concepts

- Replace derivatives with differences on a grid. For $u_t = \alpha u_{xx}$, the explicit scheme is $u_i^{n+1} = u_i^n + r(u_{i+1}^n - 2u_i^n + u_{i-1}^n)$ with $r = \alpha \Delta t / \Delta x^2$.
- The explicit heat scheme is stable only if $r \leq \frac{1}{2}$, so halving Δx forces Δt down by 4.
- Implicit and Crank–Nicolson schemes are stable for any Δt but need a linear solve each step.
- Laplace’s equation $u_{xx} + u_{yy} = 0$ gives a 5-point stencil; Jacobi or Gauss–Seidel (Unit 4) solve it.
- For waves and flow, the CFL condition $c \Delta t / \Delta x \leq 1$ links time step to grid spacing. Refining a 2D-horizontal grid by a factor k costs about k^2 more columns and k more time steps.

Key equations

$$r = \frac{\alpha \Delta t}{\Delta x^2} \leq \frac{1}{2}$$

Explicit heat stability

$$u_{i,j} = \frac{1}{4}(u_{i+1,j} + u_{i-1,j} + u_{i,j+1} + u_{i,j-1})$$

Laplace 5-point

$$\frac{c \Delta t}{\Delta x} \leq 1$$

CFL condition

$$\text{cost} \propto k^3$$

Horizontal refinement by k

Worked examples, step by step

Example 9.1 Find the largest stable explicit time step for $\alpha = 1.0 \times 10^{-4} \text{ m}^2/\text{s}$ and $\Delta x = 0.01 \text{ m}$ (sample values).

1. $\Delta t \leq 0.5\Delta x^2/\alpha$.

Answer: $\Delta t \leq 0.5 \text{ s}$.

Example 9.2 NOAA plans to refine the GFS from 13 km to 9 km. Estimate the change in horizontal columns.

1. $(13/9)^2$.

Answer: About 2.09 times as many columns.

Example 9.3 Including a CFL-limited time step, estimate the total cost increase.

1. $(13/9)^3$.

Answer: About 3.01 times the work.

Practice problems

Work each on paper first, then check.

1. Halve Δx in the explicit heat scheme. New Δt limit?

$\frac{1}{4}$ as large.

2. Columns on a 13 km global grid (Earth $5.1 \times 10^8 \text{ km}^2$, textbook)?

About 3.02 million.

3. Cells with 127 layers?

About 3.83×10^8 .

4. Cells at 9 km with 127 layers?

About 8.0×10^8 .

Common misconceptions

Myth: Any time step works if the grid is fine.

Actually: Finer grids need smaller explicit steps.

Myth: Implicit schemes are always more accurate.

Actually: They are more stable, not automatically more accurate.

Myth: Doubling resolution doubles cost.

Actually: In 2D plus time it is about 8×.

Flashcards

Stencil

Pattern of neighbour points.

r

$\alpha \Delta t / \Delta x^2$.

CFL

$c \Delta t / \Delta x \leq 1$.

Crank–Nicolson

Implicit, second order.

5-point Laplacian

Average of neighbours.

Resolution cost

$\approx k^3$ for k-fold refinement.

Self-check quiz

1. Explicit heat stability...

$r \leq \frac{1}{2}$

$r \leq 2$

any r

$r \geq 1$

$r \leq \frac{1}{2}$. Von Neumann.

2. CFL links...

Δt and Δx

α and u

n and ε

x and y

Δt and Δx . Information speed.

3. Laplace stencil uses ... points.

5

2

9 always

1

5. Centre + 4.

4. Crank–Nicolson is...

unconditionally stable

explicit

unstable

first order

unconditionally stable. Implicit average.

Try this quiz again

How this helps save the planet

NOAA's 2023 supercomputer upgrade was meant to let the GFS move to a horizontal resolution of nine kilometres, compared with 13 in the current version, so smaller-scale features can be seen and simulated.

This unit's scaling shows why that needed new computers: 2.09 times as many columns and, with a CFL-limited time step, about 3 times the work for every forecast.

On Saving Planets: [Fluid Mechanics](#) · [Thermodynamics](#)

Sources: NOAA: [NOAA completes upgrade to weather and climate supercomputer system \(Aug. 10, 2023\)](https://www.noaa.gov/news-release/noaa-completes-upgrade-to-weather-and-climate-supercomputer-system) (<https://www.noaa.gov/news-release/noaa-completes-upgrade-to-weather-and-climate-supercomputer-system>) · NOAA Environmental Modeling Center: [Global Forecast System \(GFS\)](https://www.emc.ncep.noaa.gov/emc/pages/numerical_forecast_systems/gfs.php) (https://www.emc.ncep.noaa.gov/emc/pages/numerical_forecast_systems/gfs.php)

Free learning resources

- [MIT OpenCourseWare 2.086 Numerical Computation for Mechanical Engineers \(Fall 2014\)](https://ocw.mit.edu/courses/2-086-numerical-computation-for-mechanical-engineers-fall-2014/) (<https://ocw.mit.edu/courses/2-086-numerical-computation-for-mechanical-engineers-fall-2014/>): Free MIT course.
- [MIT OpenCourseWare 18.330 Introduction to Numerical Analysis \(Spring 2012\)](https://ocw.mit.edu/courses/18-330-introduction-to-numerical-analysis-spring-2012/) (<https://ocw.mit.edu/courses/18-330-introduction-to-numerical-analysis-spring-2012/>): Free MIT lectures.



I finished this unit

Optimization, Eigenvalues and Computing at Scale

Golden-section search, gradient descent and Newton's method for minimization, the power method for eigenvalues, flop counts, parallel computing, and verifying and validating models.

Learning objectives

By the end of this unit you can:

- Minimize a 1D function by golden-section search.
- Apply gradient descent and choose a step.
- Find a dominant eigenvalue by the power method.
- Estimate run time from flop counts and machine speed.

Key concepts

- Golden-section search shrinks a bracket by 0.618 each step without derivatives.
- Gradient descent steps downhill: $x \leftarrow x - \alpha \nabla f$. Too large an α diverges; too small crawls.
- Newton's method for minimization solves $\nabla f = 0$ using the Hessian; it converges quadratically near the minimum.
- The power method repeatedly multiplies by A ; the vector lines up with the dominant eigenvector, and the Rayleigh quotient gives its eigenvalue.
- Verification asks whether the code solves the equations right; validation asks whether they are the right equations, checked against observations.

Key equations

$$\tau = \frac{\sqrt{5} - 1}{2} \approx 0.618$$

Golden ratio factor

$$x_{k+1} = x_k - \alpha \nabla f(x_k)$$

Gradient descent

$$\lambda \approx \frac{v^T A v}{v^T v}$$

Rayleigh quotient

$$t = \frac{\text{flops}}{\text{flop rate}}$$

Worked examples, step by step

Example 10.1 How much does golden-section search shrink a bracket in 10 steps?

1. 0.618^{10} .

Answer: To 0.0081 (0.81%) of its width.

Example 10.2 Minimize $f = (x - 3)^2$ by gradient descent with $\alpha = 0.1$ from $x = 0$. Where are you after 5 steps?

1. Each step: $x - 3$ shrinks by 0.8.
2. $x_5 = 3(1 - 0.8^5)$.

Answer: $x_5 = 2.017$.

Example 10.3 Apply the power method to $\begin{bmatrix} 2 & 1 \\ 1 & 2 \end{bmatrix}$ from $(1, 0)$, 5 steps.

1. $v = (122, 121)$.
2. Rayleigh quotient.

Answer: $\lambda \approx 2.99997$ (exact 3).

Practice problems

Work each on paper first, then check.

1. How long does a 10,000-unknown dense solve ($2n^3/3$ flops) take at 14.5 petaflops?

$$6.67 \times 10^{11} / 1.45 \times 10^{16} = 4.6 \times 10^{-5} \text{ s.}$$

2. Combined speed of NOAA's two 14.5-petaflop computers?

29 petaflops.

3. Operational plus research capacity (20 petaflops)?

49 petaflops.

4. Verification or validation: comparing a model with Mauna Loa data?

Validation.

Common misconceptions

Myth: A bigger step always speeds gradient descent.

Actually: It can overshoot and diverge.

Myth: The power method finds every eigenvalue.

Actually: Only the dominant one.

Myth: A converged code is a correct model.

Actually: It still needs validation.

Flashcards

Golden section

Bracket shrinks by 0.618.

Gradient descent

$x - \alpha \nabla f$.

Power method

Dominant eigenvalue.

Rayleigh quotient

$v^T A v / v^T v$.

Verification

Solving the equations right.

Validation

Solving the right equations.

Self-check quiz

1. Golden section shrinks by...

0.618

0.5

0.1

2

0.618. Golden ratio.

2. Power method finds...

dominant eigenvalue

all roots

the inverse

the determinant

dominant eigenvalue. Repeated multiplication.

3. Validation compares with...

observations

the code

nothing

epsilon

observations. Reality check.

4. Gradient descent moves...

downhill

uphill

randomly

sideways

downhill. $-\nabla f$.

[Try this quiz again](#)

How this helps save the planet

NOAA reports that its twin weather and climate supercomputers, Dogwood in Manassas, Virginia, and Cactus in Phoenix, Arizona, were upgraded in August 2023 to 14.5 petaflops each, so together

the forecast system can process 29 quadrillion calculations per second. With research machines totalling 20 petaflops, NOAA's capacity is 49 petaflops, and over 20 operational numerical weather prediction models run on the system.

NASA GISS makes the code of its ModelE Earth system model freely available, and its GISTEMP surface-temperature analysis is updated about the 10th of every month. Open code and data let anyone verify and validate the numerical methods behind climate science.

On Saving Planets: [Solutions](#) · [Intro to Programming](#)

Sources: NOAA: NOAA completes upgrade to weather and climate supercomputer system (Aug. 10, 2023) (<https://www.noaa.gov/news-release/noaa-completes-upgrade-to-weather-and-climate-supercomputer-system>) · NASA GISS: GISS Earth System Model, ModelE (<https://www.giss.nasa.gov/tools/modelE/>) · NASA GISS: GISS Surface Temperature Analysis (GISTEMP v4) (<https://data.giss.nasa.gov/gistemp/>)

Free learning resources

- MIT OpenCourseWare 18.330 Introduction to Numerical Analysis (Spring 2012) (<https://ocw.mit.edu/courses/18-330-introduction-to-numerical-analysis-spring-2012/>): Free MIT lectures.
- MIT OpenCourseWare 10.34 Numerical Methods Applied to Chemical Engineering (Fall 2015) (<https://ocw.mit.edu/courses/10-34-numerical-methods-applied-to-chemical-engineering-fall-2015/>): Free MIT course.

 I finished this unit

The lab (Computation lab)

The lab is the hands-on part of the course. Here it is free and for practice only, with no credit.

Every lab runs in a free simulation or a spreadsheet, so there are no physical hazards. Take screen breaks, sit with good posture, and keep cables out of walkways.

Safety reference: OSHA: Walking-working surfaces (slip, trip and fall hazards) (<https://www.osha.gov/walking-working-surfaces>)

Lab A: Machine epsilon and Taylor series in a spreadsheet

WEEK 1

Goal: Find machine epsilon by experiment and watch a Taylor series converge to e.

$$e \approx \sum_{k=0}^{n-1} \frac{1}{k!}$$

Taylor series for e

Procedure

1. In a column, start at 1 and halve each row.
2. In the next column, test whether $1 + \text{value} > 1$.
3. The last value that passes is machine epsilon.
4. Sum $1/k!$ for 5 and 10 terms and compare with e.

Data table

Quantity	Value
Machine epsilon (double)	2.22×10^{-16}
e, 5 terms	2.70833
e, 10 terms	2.7182815
e	2.7182818

Computed values; most spreadsheets use double precision.

Sample calculation

1. Halving 52 times gives $2^{-52} = 2.22 \times 10^{-16}$.
2. $1 + 1 + \frac{1}{2} + \frac{1}{6} + \frac{1}{24} = 2.70833$.

Safety

- Computer only.

Home or virtual alternative

- Any spreadsheet works.

MIT OpenCourseWare 18.330 Introduction to Numerical Analysis (Spring 2012) (<https://ocw.mit.edu/course/s/18-330-introduction-to-numerical-analysis-spring-2012/>)

Planet link: NOAA describes a forecast run with dynamics in 32-bit and everything else in 64-bit precision. NOAA Environmental Modeling Center: Global Forecast System (GFS) (https://www.emc.ncep.noaa.gov/emc/pages/numerical_forecast_systems/gfs.php)



I completed this lab (wet or virtual)

Lab B: Bisection versus Newton

WEEK 3

Goal: Solve $\cos x = x$ by bisection and by Newton's method and count the correct digits each step.

$$x_{n+1} = x_n - \frac{\cos x_n - x_n}{-\sin x_n - 1}$$

Newton for $\cos x - x = 0$

Procedure

1. Bracket the root in $[0, 1]$.
2. Run 20 bisection steps in a spreadsheet.
3. Run Newton from $x_0 = 1$.
4. Plot $\log_{10}$ of the error against the step number for both.

Data table

Newton step	x
1	0.750364
2	0.739113
3	0.739085133
4	0.739085133

Computed values. Bisection needs about 20 steps for 6 digits; Newton needs 3.

Sample calculation

1. $x_1 = 1 - (\cos 1 - 1)/(-\sin 1 - 1) = 0.750364$.
2. Root 0.7390851.

Safety

- Computer only.

Home or virtual alternative

- Any spreadsheet works.

OpenStax Calculus Volume 1, 4.9 Newton's method (<https://openstax.org/books/calculus-volume-1/pages/4-9-newtons-method>)

Planet link: The USGS says MODFLOW-NWT uses a Newton-Raphson formulation for unconfined groundwater flow. USGS: MODFLOW and Related Programs (<https://www.usgs.gov/mission-areas/water-resources/science/modflow-and-related-programs>)



I completed this lab (wet or virtual)

Lab C: Fitting the Mauna Loa CO₂ record

WEEK 7

Goal: Fit straight-line and quadratic models to NOAA's Mauna Loa annual means and judge them by their residuals.

$$A^T A c = A^T y$$

Least squares

Procedure

1. Explore PhET Least-Squares Regression and Curve Fitting to see how residuals change with the model.
2. Open NOAA's annual-mean file and copy the 1960–2020 decade values.
3. Fit a line and a quadratic (spreadsheet LINEST or trendline).
4. Predict 2025 and compare with NOAA's 427.35 ppm.

Data table

Model	Key coefficient	2025 prediction (ppm)
Line	slope 1.613 ppm/yr	415.0
Quadratic	0.01338 ppm/yr ²	426.0
NOAA 2025	—	427.35

Data from NOAA GML; fits computed.

Sample calculation

1. The line underpredicts 2025 by 12.3 ppm.
2. The quadratic is within 1.3 ppm.

Safety

- Computer only.

Home or virtual alternative

- Any spreadsheet works.

PhET: Least-Squares Regression (<https://phet.colorado.edu/en/simulations/least-squares-regression>) · PhET: Curve Fitting (<https://phet.colorado.edu/en/simulations/curve-fitting>) · NOAA GML: Mauna Loa annual mean CO₂ data (https://gml.noaa.gov/webdata/ccgg/trends/co2/co2_annmean_mlo.txt)

Planet link: NOAA reports the Mauna Loa monthly mean was 426.07 ppm in September 2026.

NOAA Global Monitoring Laboratory: Trends in atmospheric CO₂ (Mauna Loa) (<https://gml.noaa.gov/ccgg/trends/>)



I completed this lab (wet or virtual)

Lab D: Trapezoid versus Simpson

WEEK 9

Goal: Integrate $\sin x$ from 0 to π with both rules, halve the step and measure the error order.

$$\int_0^{\pi} \sin x \, dx = 2$$

Test integral

Procedure

1. Use PhET Calculus Grapher to see the area under a curve.
2. Build columns for x , $\sin x$ and the rule weights.
3. Compute T and S for $n = 4$, then $n = 8$.
4. Find the error ratios.

Data table

Rule, $n = 4$	Estimate	Error
Trapezoidal	1.89612	-0.10388
Simpson	2.00456	+0.00456

Computed values; doubling to $n = 8$ cuts the errors by 4.03 and 16.9, close to the theoretical 4 and 16.

Sample calculation

1. $T_4 = (\pi/4)(\sin \pi/4 + \sin \pi/2 + \sin 3\pi/4) = 1.89612$.
2. $S_4 = (\pi/12)(4 \sin \pi/4 + 2 \sin \pi/2 + 4 \sin 3\pi/4) = 2.00456$.

Safety

- Computer only.

Home or virtual alternative

- Any spreadsheet works.

PhET: Calculus Grapher (<https://phet.colorado.edu/en/simulations/calculus-grapher>) • OpenStax Calculus Volume 2, 3.6 Numerical integration (<https://openstax.org/books/calculus-volume-2/pages/3-6-numerical-integration>)

Planet link: NOAA builds each Mauna Loa monthly mean from daily means based on hourly averages. NOAA Global Monitoring Laboratory: Trends in atmospheric CO₂ (Mauna Loa) (https://gml.noaa.gov/records/mauna_loa/)



I completed this lab (wet or virtual)

Lab E: Euler and RK4 error orders

WEEK 11

Goal: Solve $y' = -y$ to $t = 1$ with several step sizes and confirm first- and fourth-order convergence.

$$y(1) = e^{-1} = 0.367879$$

Exact answer

Procedure

1. Build an Euler column for $h = 0.2, 0.1$ and 0.05 .
2. Build an RK4 block for $h = 0.2$ and 0.1 .
3. Record the error at $t = 1$.
4. Compare the error ratios when h halves; then try the same solver on PhET Projectile Motion's no-drag case.

Data table

Method, h	Error at $t = 1$
Euler, 0.2	-0.0402
Euler, 0.1	-0.0192
Euler, 0.05	-0.00939
RK4, 0.2	5.80×10^{-6}
RK4, 0.1	3.33×10^{-7}

Computed values. Euler's error halves with h ; RK4's falls by about 17 (near $2^4 = 16$).

Sample calculation

1. Euler $h = 0.1$: $0.9^{10} - e^{-1} = -0.0192$.
2. RK4 ratio $5.80 \times 10^{-6} / 3.33 \times 10^{-7} = 17.4$.

Safety

- Computer only.

Home or virtual alternative

- Any spreadsheet works.

PhET: Projectile Motion (<https://phet.colorado.edu/en/simulations/projectile-motion>) · MIT OpenCourseWare
10.34 Numerical Methods Applied to Chemical Engineering (Fall 2015) (<https://ocw.mit.edu/courses/10-34-numerical-methods-applied-to-chemical-engineering-fall-2015/>)

Planet link: NOAA's GFS runs four times a day, forecasting up to 16 days ahead. NOAA
Environmental Modeling Center: Global Forecast System (GFS) (https://www.emc.ncep.noaa.gov/emc/pages/numerical_forecast_systems/gfs.php)



I completed this lab (wet or virtual)

Lab F: The heat equation and its stability limit

WEEK 13

Goal: Run the explicit heat scheme on a rod with a hot spot and see what happens above $r = \frac{1}{2}$.

$$u_i^{n+1} = u_i^n + r(u_{i+1}^n - 2u_i^n + u_{i-1}^n)$$

Explicit scheme

Procedure

1. Set 7 nodes with ends at 0 and the middle node at 100 (sample values).
2. Step with $r = 0.25$ for 10 steps.
3. Repeat with $r = 0.6$.
4. Describe both results.

Data table

Case	Middle value after 1 step	After 10 steps
$r = 0.25$	50.0	16.7
$r = 0.6$	-20.0	290.5

Computed values. With $r = 0.6$ the values oscillate and grow: the scheme is unstable.

Sample calculation

1. $r = 0.25$: $100 + 0.25(0 - 200 + 0) = 50$.
2. $r = 0.6$: $100 + 0.6(-200) = -20$.

Safety

- Computer only.

Home or virtual alternative

- Any spreadsheet works.

MIT OpenCourseWare 2.086 Numerical Computation for Mechanical Engineers (Fall 2014) (<https://ocw.mit.edu/courses/2-086-numerical-computation-for-mechanical-engineers-fall-2014/>)

Planet link: NOAA planned to move the GFS from 13 km to 9 km resolution after its 2023 upgrade.
NOAA: NOAA completes upgrade to weather and climate supercomputer system (Aug. 10, 2023) (<https://www.noaa.gov/news-release/noaa-completes-upgrade-to-weather-and-climate-supercomputer-system>)



I completed this lab (wet or virtual)

Exam prep practice sets

Time yourself: about 10 minutes per problem. Then compare with the full solutions.

Midterm practice set (Units 1–5)

1. Find the relative error of 3.14 as an approximation to π .

1. $|3.14 - \pi|/\pi$.

Answer: 5.07×10^{-4} .

2. Take two Newton steps on $x^2 - 5$ from $x_0 = 2$.

1. $x_1 = 2 - (-1)/4$.

2. $x_2 = 2.25 - 0.0625/4.5$.

Answer: $x_1 = 2.25$; $x_2 = 2.23611$.

3. How many bisection steps shrink a width-4 bracket below 10^{-3} ?

1. $\log_2(4000) = 11.97$.

Answer: 12 steps.

4. How much more elimination work is $n = 500$ than $n = 100$?

1. $(500/100)^3$.

Answer: 125 times.

5. Interpolate Mauna Loa CO_2 for 2023 from NOAA's 2020 (414.21) and 2025 (427.35) annual means.

1. $414.21 + 0.6 \times 13.14$.

Answer: 422.09 ppm (NOAA's 2023 value is 421.08).

Final practice set (all units)

1. Estimate $\int_0^2 x^2 dx$ with the trapezoidal rule, $n = 2$.

1. $h = 1$: $0.5 \times 0 + 1 + 0.5 \times 4$.

Answer: 3 (exact 2.667).

2. Repeat with Simpson's rule, $n = 2$.

1. $(1/3)(0 + 4 + 4)$.

Answer: 2.667, exact for a quadratic.

3. Take two Euler steps of $h = 0.1$ on $y' = -2y$, $y(0) = 1$.

1. 0.8^2 .

Answer: 0.64 (exact 0.670).

4. Find the largest explicit heat time step for $\alpha = 2 \times 10^{-5} \text{ m}^2/\text{s}$, $\Delta x = 0.005 \text{ m}$.

1. $0.5 \times 0.005^2 / 2 \times 10^{-5}$.

Answer: 0.625 s.

5. Estimate the derivative of e^x at 0 by central difference, $h = 0.01$.

1. $(e^{0.01} - e^{-0.01})/0.02$.

Answer: 1.0000167.

6. How much does golden-section search shrink a bracket in 20 steps?

1. 0.618^{20} .

Answer: To 6.6×10^{-5} of its width.

Formula sheet

$$e_{rel} = |\tilde{x} - x|/|x|, \varepsilon_{mach} = 2^{-52}$$

$$x_{n+1} = x_n - f(x_n)/f'(x_n), n \geq \log_2((b-a)/tol)$$

$$A = LU, \text{ cost} \approx \frac{2}{3}n^3, \kappa = \|A\| \|A^{-1}\|$$

$$\|e_k\| \approx \rho^k \|e_0\|$$

$$A^T A c = A^T y$$

$$f' \approx (f(x+h) - f(x-h))/(2h), h_{opt} \approx \sqrt{\varepsilon}$$

$$T_n : E \propto h^2, S_n : E \propto h^4$$

$$y_{n+1} = y_n + hf, h < 2/\lambda$$

$$r = \alpha \Delta t / \Delta x^2 \leq \frac{1}{2}, c \Delta t / \Delta x \leq 1$$

$$x_{k+1} = x_k - \alpha \nabla f, \lambda \approx v^T A v / v^T v$$

Constants

Constant	Value
Double precision (IEEE 754)	machine epsilon $2^{-52} = 2.22 \times 10^{-16}$, about 16 significant digits

Constant	Value
Single precision (IEEE 754)	machine epsilon $2^{-23} = 1.19 \times 10^{-7}$, about 7 significant digits
e	2.718281828
π	3.141592654
Earth's surface area (textbook value)	$5.1 \times 10^8 \text{ km}^2$
1 petaflop	10^{15} floating-point operations per second

Source: [NIST Guide to the SI, Appendix B.9: conversion factors](https://www.nist.gov/pml/special-publication-811/nist-guide-si-appendix-b-conversion-factors/nist-guide-si-appendix-b9) (<https://www.nist.gov/pml/special-publication-811/nist-guide-si-appendix-b-conversion-factors/nist-guide-si-appendix-b9>)

Free texts and simulations

Use these free lectures, textbook sections and simulations alongside the course.

- [MIT OpenCourseWare 18.330 Introduction to Numerical Analysis \(Spring 2012\)](https://ocw.mit.edu/courses/18-330-introduction-to-numerical-analysis-spring-2012/) (<https://ocw.mit.edu/courses/18-330-introduction-to-numerical-analysis-spring-2012/>): Free MIT numerical analysis lectures and problem sets.
- [MIT OpenCourseWare 2.086 Numerical Computation for Mechanical Engineers \(Fall 2014\)](https://ocw.mit.edu/courses/2-086-numerical-computation-for-mechanical-engineers-fall-2014/) (<https://ocw.mit.edu/courses/2-086-numerical-computation-for-mechanical-engineers-fall-2014/>): Free MIT engineering computation course.
- [MIT OpenCourseWare 10.34 Numerical Methods Applied to Chemical Engineering \(Fall 2015\)](https://ocw.mit.edu/courses/10-34-numerical-methods-applied-to-chemical-engineering-fall-2015/) (<https://ocw.mit.edu/courses/10-34-numerical-methods-applied-to-chemical-engineering-fall-2015/>): Free MIT applied numerical methods course.
- [OpenStax Calculus Volume 2, 3.6 Numerical integration](https://openstax.org/books/calculus-volume-2/pages/3-6-numerical-integration) (<https://openstax.org/books/calculus-volume-2/pages/3-6-numerical-integration>): Free textbook section.
- [OpenStax Calculus Volume 1, 4.9 Newton's method](https://openstax.org/books/calculus-volume-1/pages/4-9-newtons-method) (<https://openstax.org/books/calculus-volume-1/pages/4-9-newtons-method>): Free textbook section.
- [PhET: Least-Squares Regression](https://phet.colorado.edu/en/simulations/least-squares-regression) (<https://phet.colorado.edu/en/simulations/least-squares-regression>): Fit lines interactively.

Glossary

Absolute error

Size of the difference from the true value.

Bisection

Root finding by halving a sign-change bracket.

CFL condition

Time-step limit $c\Delta t/\Delta x \leq 1$.

Catastrophic cancellation

Digit loss when subtracting close numbers.

Condition number

How much a problem amplifies input errors.

Crank–Nicolson

Implicit second-order time scheme.

Diagonal dominance

Each diagonal entry outweighs the rest of its row.

Gaussian quadrature

Integration with optimal points and weights.

Gauss–Seidel

Iterative solver using new values at once.

LU factorization

$A = LU$, reusable elimination.

Least squares

Fit minimizing the sum of squared residuals.

Machine epsilon

Gap between 1 and the next float.

Newton's method

Root finding along tangent lines.

Power method

Repeated multiplication to find the dominant eigenvalue.

Runge–Kutta (RK4)

Fourth-order ODE solver.

Simpson's rule

Fourth-order integration with parabolas.

Spline

Smooth piecewise cubic curve.

Stiff equation

ODE with very fast and very slow parts.

Truncation error

Error from cutting off an infinite process.

Validation

Checking a model against observations.

Self-assessment rubric

No grade or credit is given. Use this to judge honestly where you are in each unit.

Level	What it looks like
4 · Mastery	Chooses a method for the accuracy needed, tests it on a problem with a known answer, measures its order, and explains its stability limits.
3 · Proficient	Solves the practice problems with occasional slips in step size or units.
2 · Developing	Follows worked examples but cannot tell truncation from rounding error.
1 · Beginning	Runs formulas without checking error or stability; needs to revisit the unit.

Suggested weighting if you want a score

Part	Weight
Practice problems	20%
Labs	25%
Midterm practice set	25%
Final practice set	30%

Aim for level 3 or better in every unit before moving to the next course.

Related courses

- [Orbital Mechanics](#) AERO 340
- [Aerodynamics I](#) AERO 310
- [Thermodynamics](#) ENGR 320
- [All courses](#)