STUDY PLAN · 2026

# Three-Year Data Science Curriculum

From mathematical and computing foundations to machine learning and advanced AI, using university specializations on Coursera and reference textbooks.

| Year 1 | Year 2 | Year 3 |
| --- | --- | --- |
| Foundations | Theoretical computer science and statistics | Data, machine learning and AI |
| **398 h** | **442 h** | **453 h** |

**1,293 h** of core coursework

**6** semesters

**61** courses

**10** universities

**17** reference textbooks

Data collected from coursera.org on 4 November 2026 · Textbook availability checked on the same date

# Contents

1. About the curriculum
2. Overview and workload
3. Year 1 · Foundations (S1, S2)
4. Year 2 · Theoretical computer science and statistics (S3, S4)
5. Year 3 · Data, machine learning and AI (S5, S6)
6. Advanced options
7. Bibliography
8. Method and caveats

# 1. About the curriculum

Over three years, this curriculum trains a data scientist who understands the methods they use rather than just calling them. It is designed for a motivated adult with end-of-high-school mathematics who can devote about ten hours a week to study.

### THE PRINCIPLES BEHIND THE ORDER

- **Tools before methods.** Programming (S1) and continuous mathematics (S2) come first, because everything else uses them.
- **Discrete before probabilistic.** Proofs, combinatorics and graphs (S3) prepare probability (S4) and the complexity analysis of learning algorithms.
- **Statistics before machine learning.** Estimation, testing and regression (S4) provide the framework in which the S5 models make sense.
- **Threads running through the curriculum.** Linear algebra (S2) returns in regression, neural networks and reinforcement learning; Markov chains (S4) return in Bayesian MCMC, reinforcement learning and graphical models.
- **One project per semester.** Each semester ends with a capstone or project to consolidate what has been learned.

### HOW TO READ THIS DOCUMENT

Each semester sets out its goals, target skills, final project and things to watch, then the course-by-course programme with links, lengths, levels and Coursera ratings, followed by further reading. Links are clickable in the PDF.

# 2. Overview and workload

| YEAR | SEM. | CONTENT | HOURS |
|---|---|---|---|
| **Year 1**<br>Foundations<br>**398 h** | **S1** | **Programming and thinking with data**<br>Understanding Data: Navigating Statistics, Science, and AI (30 h)<br>Python for Everybody Specialization (79 h)<br>Fundamentals of Computing Specialization (98 h) | 207 h |
| | **S2** | **Calculus and linear algebra**<br>Calculus for Engineers (30 h)<br>Linear Algebra from Elementary to Advanced (34 h)<br>Mathematics for Engineers (127 h) | 191 h |
| **Year 2**<br>Theoretical computer science and statistics<br>**442 h** | **S3** | **Discrete mathematics and algorithms**<br>Introduction to Discrete Mathematics for Computer Science (116 h)<br>Algorithms Specialization (57 h) | 173 h |
| | **S4** | **Probability and statistical modeling**<br>Foundations of Probability and Statistics (105 h)<br>Statistical Inference and Hypothesis Testing in Data Science Applications (37 h)<br>Statistical Modeling for Data Science Applications (127 h) | 269 h |
| **Year 3**<br>Data, machine learning and AI<br>**453 h** | **S5** | **Data and machine learning**<br>Data Wrangling with Python (135 h)<br>Machine Learning: Theory and Hands-on Practice with Python (56 h) | 191 h |
| | **S6** | **Advanced AI**<br>Bayesian Statistics (94 h)<br>Foundations of Neural Networks (36 h)<br>Deep Learning (57 h)<br>Reinforcement Learning (75 h) | 262 h |
| **Core curriculum total** | | | **1,293 h** |
| Advanced options (optional) | | | + 240 h |

COURSE HOURS PER SEMESTER




SUGGESTED PACE

Based on 40 weeks of study per year, the curriculum takes about **10 h per week in year 1**, **11 h in year 2** and **11 h in year 3**. Lengths are those Coursera gives for each course; allow extra time for further reading, roughly 20 to 30% of course time.

LANGUAGES AND TOOLS

Python is the main language (S1, S3, S5, S6). MATLAB or GNU Octave is used in S2 for numerical computing. R is used in S4 for statistical modeling and in S6 for Bayesian statistics. This variety is deliberate: each tool is the one its discipline uses.

YEAR 1 · 398 HOURS

# Foundations

Programming, reasoning about data, mastering calculus and linear algebra

S1

YEAR 1 · 207 HOURS · ≈ 10 H PER WEEK OVER 20 WEEKS

## Programming and thinking with data

> Write your first Python programs, work with files, the web and SQL databases, and build a critical understanding of data and AI.

### TARGET SKILLS

- Python variables, loops, functions and data structures
- Reading files, accessing the web and APIs, SQL queries
- Algorithmic thinking: interactive programs, basic combinatorics, graphs
- Critical reading of a statistic, a study or an AI prediction

### END-OF-SEMESTER PROJECT

The Python for Everybody capstone (collecting, storing and visualizing web data), then the Fundamentals of Computing final exam.

### THINGS TO WATCH

- Python for Everybody and Fundamentals of Computing overlap at the start: take Rice alongside the second half of Python for Everybody.

### PROGRAMME

1

SPECIALIZATION · 3 COURSES · 30 H

#### Understanding Data: Navigating Statistics, Science, and AI

University of Michigan · Beginner

**Stated prerequisites:** No previous knowledge of data, statistics or AI necessary

| # | COURSE | LEVEL | LENGTH | RATING (REVIEWS) |
|---|---|---|---|---|
| 1 | How to Describe Data | Beginner | 9 h | 4.7 (29) |
| 2 | How Science Turns Data Into Knowledge | Beginner | 11 h | 4.7 (18) |
| 3 | Decoding AI: A Deep Dive into AI Models and Predictions | Beginner | 10 h | 4.7 (29) |

2

SPECIALIZATION · 5 COURSES · 79 H

## Python for Everybody Specialization

University of Michigan | Beginner | Python, SQL, SQLite

**Stated prerequisites:** No prerequisites; designed as an on-ramp to programming

| # | COURSE | LEVEL | LENGTH | RATING (REVIEWS) |
|---|---|---|---|---|
| 1 | Programming for Everybody (Getting Started with Python) | Beginner | 17 h | 4.8 (233,775) |
| 2 | Python Data Structures | Beginner | 20 h | 4.9 (97,365) |
| 3 | Using Python to Access Web Data | Beginner | 19 h | 4.8 (44,591) |
| 4 | Using Databases with Python | Beginner | 14 h | 4.8 (21,659) |
| 5 | Capstone: Retrieving, Processing, and Visualizing Data with Python | Beginner | 9 h | 4.7 (14,483) |

3

SPECIALIZATION · 7 COURSES · 98 H

## Fundamentals of Computing Specialization

Rice University | Beginner | Python

**Stated prerequisites:** High school mathematics; no programming required

| # | COURSE | LEVEL | LENGTH | RATING (REVIEWS) |
|---|---|---|---|---|
| 1 | An Introduction to Interactive Programming in Python (Part 1) | — | 19 h | 4.8 (3,322) |
| 2 | An Introduction to Interactive Programming in Python (Part 2) | Beginner | 16 h | 4.9 (1,214) |
| 3 | Principles of Computing (Part 1) | Intermediate | 19 h | 4.7 (584) |
| 4 | Principles of Computing (Part 2) | Intermediate | 16 h | 4.8 (365) |
| 5 | Algorithmic Thinking (Part 1) | Intermediate | 14 h | 4.6 (385) |
| 6 | Algorithmic Thinking (Part 2) | Intermediate | 13 h | 4.7 (220) |
| 7 | The Fundamentals of Computing Capstone Exam | Intermediate | 1 h | 4.8 (116) |

### FURTHER READING

**Python for Everybody: Exploring Data in Python 3**

Charles R. Severance

The specialization's official textbook: one chapter per week of the course.

Free (PDF, HTML, EPUB), CC BY 4.0 licence · www.py4e.com/book

**Think Python, 3rd edition**

Allen B. Downey

A second, more rigorous look at functions, recursion and classes; prepares Python 3 and OOP.

Free online (Jupyter notebooks), CC BY-NC-SA 4.0 licence · allendowney.github.io/ThinkPython

S2

YEAR 1 · 191 HOURS · ≈ 10 H PER WEEK OVER 20 WEEKS

# Calculus and linear algebra

> Acquire the continuous mathematical tools on which all of statistics and machine learning depend.

### TARGET SKILLS

- Derivatives, integrals, series and applications of single-variable calculus
- Vectors, matrices, determinants, eigenvalues, orthogonality
- Differential equations and vector calculus
- Numerical methods and scientific programming in MATLAB

### END-OF-SEMESTER PROJECT

The Mathematics for Engineers capstone: numerical computation of the flow past a cylinder (Kármán vortex street).

### THINGS TO WATCH

- Calculus for Engineers comes first: Mathematics for Engineers assumes single-variable calculus.
- Install MATLAB (or GNU Octave) before the semester.

## PROGRAMME

1

COURSE · 1 COURSE · 30 H

### Calculus for Engineers

The Hong Kong University of Science and Technology | Beginner

**Stated prerequisites:** High school mathematics

| # | COURSE | LEVEL | LENGTH | RATING (REVIEWS) |
|---|---|---|---|---|
| 1 | Calculus for Engineers | Beginner | 30 h | 4.9 (54) |

2

SPECIALIZATION · 3 COURSES · 34 H

### Linear Algebra from Elementary to Advanced

Johns Hopkins University | Beginner

**Stated prerequisites:** Basic algebra; no calculus needed

| # | COURSE | LEVEL | LENGTH | RATING (REVIEWS) |
|---|---|---|---|---|
| 1 | Linear Algebra: Linear Systems and Matrix Equations | Beginner | 10 h | 4.7 (183) |
| 2 | Linear Algebra: Matrix Algebra, Determinants, & Eigenvectors | — | 15 h | 4.8 (79) |
| 3 | Linear Algebra: Orthogonality and Diagonalization | Intermediate | 9 h | 4.8 (56) |

3

SPECIALIZATION · 5 COURSES · 127 H

### Mathematics for Engineers

The Hong Kong University of Science and Technology | Beginner | MATLAB

**Stated prerequisites:** Single variable calculus and some computer programming

| # | COURSE | LEVEL | LENGTH | RATING (REVIEWS) |
|---|---|---|---|---|
| 1 | Matrix Algebra for Engineers | Beginner | 20 h | 4.9 (4,726) |
| 2 | Differential Equations for Engineers | Beginner | 27 h | 4.9 (2,251) |
| 3 | Vector Calculus for Engineers | Beginner | 30 h | 4.8 (1,425) |
| 4 | Numerical Methods for Engineers | Intermediate | 40 h | 4.8 (451) |
| 5 | Mathematics for Engineers: The Capstone Course | Intermediate | 10 h | 5.0 (45) |

## FURTHER READING

### Calculus, Volume 1

Gilbert Strang, Edwin Herman (OpenStax)

Worked exercises for Calculus for Engineers: limits, derivatives, integrals.

Free (PDF and web), OpenStax · openstax.org/details/books/calculus-volume-1

### Introduction to Applied Linear Algebra: Vectors, Matrices, and Least Squares

Stephen Boyd, Lieven Vandenberghe

Linear algebra through data: least squares, regression, clustering. Bridges to year 3.

Free (author's PDF), Python and Julia companions · web.stanford.edu/~boyd/vmls

### Mathematics for Machine Learning

Marc Peter Deisenroth, A. Aldo Faisal, Cheng Soon Ong

Part I (linear algebra, analytic geometry, vector calculus) at the end of the semester; Part II is used in S5.

Free (author's PDF) · mml-book.github.io

YEAR 2 · 442 HOURS

# Theoretical computer science and statistics

Proofs, algorithms, probability and statistical models

S3

YEAR 2 · 173 HOURS · ≈ 9 H PER WEEK OVER 20 WEEKS

## Discrete mathematics and algorithms

> Learn to prove, count and reason about graphs, then design and analyze efficient algorithms.

### TARGET SKILLS

- Proof by induction, by contradiction, invariants
- Combinatorics, discrete probability, graph theory
- Modular arithmetic and RSA cryptography
- Divide and conquer, graphs, greedy algorithms, dynamic programming, NP-completeness

### END-OF-SEMESTER PROJECT

The UCSD Delivery Problem project: solve a travelling salesman problem and compare heuristics.

### THINGS TO WATCH

- Stanford Algorithms assumes you can read proofs: take it after Discrete Mathematics.
- For more practice: Data Structures and Algorithms (UCSD, 182 h), listed in the Computer Science file.

### PROGRAMME

1

SPECIALIZATION · 5 COURSES · 116 H

#### Introduction to Discrete Mathematics for Computer Science

University of California San Diego | Beginner | Python

**Stated prerequisites:** Basic math (fractions, squares), basic Python (functions, loops, recursion)

| # | COURSE | LEVEL | LENGTH | RATING (REVIEWS) |
|---|---|---|---|---|
| 1 | Mathematical Thinking in Computer Science | Beginner | 42 h | 4.4 (2,289) |
| 2 | Combinatorics and Probability | Beginner | 23 h | 4.6 (874) |
| 3 | Introduction to Graph Theory | Beginner | 21 h | 4.5 (1,073) |
| 4 | Number Theory and Cryptography | Beginner | 17 h | 4.5 (610) |
| 5 | Delivery Problem | Beginner | 13 h | 4.7 (376) |

2

SPECIALIZATION · 4 COURSES · 57 H

## Algorithms Specialization

Stanford University (Stanford Online) | Intermediate | C, Java, Python

**Stated prerequisites:** Programming in at least one language (C, Java or Python); familiarity with proofs (induction, contradiction); discrete probability

| # | COURSE | LEVEL | LENGTH | RATING (REVIEWS) |
|---|---|---|---|---|
| 1 | Divide and Conquer, Sorting and Searching, and Randomized Algorithms | Intermediate | 15 h | 4.8 (5,342) |
| 2 | Graph Search, Shortest Paths, and Data Structures | Intermediate | 14 h | 4.8 (2,002) |
| 3 | Greedy Algorithms, Minimum Spanning Trees, and Dynamic Programming | Intermediate | 15 h | 4.8 (1,279) |
| 4 | Shortest Paths Revisited, NP-Complete Problems and What To Do About Them | Intermediate | 13 h | 4.8 (832) |

### FURTHER READING

**Mathematics for Computer Science**

Eric Lehman, F. Thomson Leighton, Albert R. Meyer (MIT)

The reference for proofs, combinatorics and graphs; covers the whole UCSD course and more.

Free (MIT OpenCourseWare PDF), CC BY-SA licence · ocw.mit.edu/courses/6-1200j-mathematics-for-computer-science-spring-2024/pages/readings

**Algorithms**

Jeff Erickson (University of Illinois)

Companion to Stanford Algorithms: recursion, dynamic programming, graphs, reductions, with many exercises.

Free (author's PDF), CC BY 4.0 licence · jeffe.cs.illinois.edu/teaching/algorithms

S4

YEAR 2 · 269 HOURS · ≈ 13 H PER WEEK OVER 20 WEEKS

# Probability and statistical modeling

> Build a solid probabilistic foundation, then learn to estimate, test and model.

### TARGET SKILLS

- Random variables, common distributions, limit theorems
- Discrete-time Markov chains and Monte Carlo methods
- Estimation, confidence intervals, hypothesis tests
- Linear regression, ANOVA, experimental design, generalized linear models

### END-OF-SEMESTER PROJECT

A complete study in R: frame a question, fit a linear model or GLM, check the residuals and write a reproducible report.

### THINGS TO WATCH

- All three blocks are from CU Boulder and are eligible for academic credit.
- Foundations of Probability and Statistics requires calculus up to Calculus II level: S2 covers it.
- The inference course is taken on its own, because the specialization's other two courses are already in Foundations of Probability and Statistics.

### PROGRAMME

1

SPECIALIZATION · 3 COURSES · 105 H

## Foundations of Probability and Statistics

University of Colorado Boulder | Intermediate | R, Python, Jupyter Notebooks | academic credit

**Stated prerequisites:** Calculus through Calculus II (multivariate preferred) and some R or Python experience

| # | COURSE | LEVEL | LENGTH | RATING (REVIEWS) |
|---|---|---|---|---|
| 1 | Probability Foundations for Data Science and AI | Intermediate | 43 h | 4.5 (300) |
| 2 | Discrete-Time Markov Chains and Monte Carlo Methods | Intermediate | 33 h | 4.5 (22) |
| 3 | Statistical Estimation for Data Science and AI | Intermediate | 29 h | 4.0 (99) |

2

COURSE · 1 COURSE · 37 H

## Statistical Inference and Hypothesis Testing in Data Science Applications

University of Colorado Boulder | Intermediate | R, Jupyter Notebooks | academic credit

single course from Data Science Foundations: Statistical Inference

**Stated prerequisites:** Calculus through Calculus II (multivariate preferred) and some R programming experience

| # | COURSE | LEVEL | LENGTH | RATING (REVIEWS) |
|---|---|---|---|---|
| 1 | Statistical Inference and Hypothesis Testing in Data Science Applications | Intermediate | 37 h | 4.6 (55) |

3

SPECIALIZATION · 3 COURSES · 127 H

## Statistical Modeling for Data Science Applications

University of Colorado Boulder | Intermediate | R, Jupyter Notebooks | academic credit

**Stated prerequisites:** Calculus, basic linear algebra and probability theory

| # | COURSE | LEVEL | LENGTH | RATING (REVIEWS) |
|---|---|---|---|---|
| 1 | Modern Regression Analysis in R | Intermediate | 45 h | 4.3 (40) |
| 2 | ANOVA and Experimental Design | Intermediate | 40 h | 4.0 (23) |
| 3 | Generalized Linear Models and Nonparametric Regression | Intermediate | 42 h | 4.2 (26) |

### FURTHER READING

**Introduction to Probability, 2nd edition**

Joseph K. Blitzstein, Jessica Hwang (Harvard)

The most teachable probability textbook: distributions, conditional expectation, Markov chains, Monte Carlo.

Free online version (probabilitybook.net), with the Stat 110 course · stat110.hsites.harvard.edu

**R for Data Science, 2nd edition**

Hadley Wickham, Mine Çetinkaya-Rundel, Garrett Grolemund

Getting started with R and the tidyverse before Statistical Modeling.

Free online, CC BY-NC-ND 3.0 licence · r4ds.hadley.nz

YEAR 3 · 453 HOURS

# Data, machine learning and AI

Preparing data, learning models, going on to deep, Bayesian and reinforcement learning

S5

YEAR 3 · 191 HOURS · ≈ 10 H PER WEEK OVER 20 WEEKS

## Data and machine learning

> Prepare real-world data at scale in Python, then train and evaluate supervised and unsupervised models.

### TARGET SKILLS

- Collecting, integrating, cleaning and transforming data with pandas
- Exploratory visualization with Matplotlib and Seaborn
- Trees, regressions, SVMs, ensemble methods, cross-validation
- Clustering, dimensionality reduction, introduction to deep learning

### END-OF-SEMESTER PROJECT

The Data Wrangling with Python final project, then a complete predictive model on the same data.

### THINGS TO WATCH

- Machine Learning: Theory and Hands-on Practice has only 84 reviews: use ISLP if anything remains unclear.

### PROGRAMME

1

SPECIALIZATION · 5 COURSES · 135 H

#### Data Wrangling with Python

University of Colorado Boulder | Beginner | Python, pandas, NumPy, Seaborn

**Stated prerequisites:** Python fundamentals and statistics

| # | COURSE | LEVEL | LENGTH | RATING (REVIEWS) |
|---|---|---|---|---|
| 1 | Fundamental Tools of Data Wrangling | Beginner | 27 h | 4.3 (22) |
| 2 | Data Collection and Integration | Intermediate | 28 h | 4.5 (16) |
| 3 | Data Understanding and Visualization | Intermediate | 25 h | — |
| 4 | Data Processing and Manipulation | Intermediate | 29 h | — |
| 5 | Data Wrangling with Python Project | Intermediate | 26 h | — |

2

SPECIALIZATION · 3 COURSES · 56 H

#### Machine Learning: Theory and Hands-on Practice with Python

University of Colorado Boulder | Intermediate | Python

| # | COURSE | LEVEL | LENGTH | RATING (REVIEWS) |
|---|---|---|---|---|
| 1 | Introduction to Machine Learning: Supervised Learning | Intermediate | 23 h | 4.5 (67) |
| 2 | Introduction to Machine Learning: Unsupervised Learning | Intermediate | 15 h | 4.6 (24) |
| 3 | Introduction to Deep Learning | Intermediate | 18 h | 4.5 (24) |

## FURTHER READING

### Python for Data Analysis, 3rd edition

Wes McKinney

Written by the creator of pandas: the reference for Data Wrangling with Python.

Open access online (HTML), code under MIT licence · wesmckinney.com/book

### An Introduction to Statistical Learning, with Applications in Python

G. James, D. Witten, T. Hastie, R. Tibshirani, J. Taylor

The reference textbook for statistical machine learning; each chapter has a Python lab.

Free PDF on the authors' website · www.statlearning.com

### Mathematics for Machine Learning (Part II)

M. P. Deisenroth, A. A. Faisal, C. S. Ong

Linear regression, PCA, Gaussian mixtures and SVMs, using the mathematics from year 1.

Free (author's PDF) · mml-book.github.io

S6

YEAR 3 · 262 HOURS · ≈ 13 H PER WEEK OVER 20 WEEKS

# Advanced AI

> Master the three main pillars of modern AI: Bayesian inference, deep neural networks and reinforcement learning.

### TARGET SKILLS

- Bayesian inference, MCMC, mixture models, time series
- Convolutional and recurrent networks, autoencoders, GANs
- Transfer learning and generative AI
- Markov decision processes, TD learning, Q-learning, policy gradient

### END-OF-SEMESTER PROJECT

Two capstones: Bayesian Statistics (a complete Bayesian analysis of a dataset) and A Complete Reinforcement Learning System (an end-to-end agent).

### THINGS TO WATCH

- Foundations of Neural Networks and the Illinois Tech Deep Learning course have few reviews (24 and 37); rely on Bishop or Goodfellow for theory.
- The Deep Learning course is taken on its own from Illinois Tech's Data Analytics and Deep Learning specialization.

### PROGRAMME

1

SPECIALIZATION · 5 COURSES · 94 H

## Bayesian Statistics

University of California, Santa Cruz | Intermediate | R, Microsoft Excel

| # | COURSE | LEVEL | LENGTH | RATING (REVIEWS) |
|---|---|---|---|---|
| 1 | Bayesian Statistics: From Concept to Data Analysis | Intermediate | 11 h | 4.6 (3,233) |
| 2 | Bayesian Statistics: Techniques and Models | Intermediate | 30 h | 4.8 (498) |
| 3 | Bayesian Statistics: Mixture Models | Intermediate | 22 h | 4.5 (63) |
| 4 | Bayesian Statistics: Time Series Analysis | Intermediate | 19 h | 4.3 (18) |
| 5 | Bayesian Statistics: Capstone Project | Advanced | 12 h | 4.5 (10) |

2

SPECIALIZATION · 3 COURSES · 36 H

## Foundations of Neural Networks

Johns Hopkins University | Intermediate | Python, TensorFlow, PyTorch

| # | COURSE | LEVEL | LENGTH | RATING (REVIEWS) |
|---|---|---|---|---|
| 1 | Introduction to Neural Networks | Intermediate | 18 h | 4.2 (10) |
| 2 | Advanced Neural Network Techniques | Intermediate | 11 h | — |
| 3 | Practical Methodology and Ethics in AI | Intermediate | 7 h | 4.5 (11) |

3

COURSE · 1 COURSE · 57 H

## Deep Learning

Illinois Tech | Intermediate | Apache Hadoop, Apache Spark, NoSQL | single course from Data Analytics and Deep Learning

| # | COURSE | LEVEL | LENGTH | RATING (REVIEWS) |
|---|---|---|---|---|
| 1 | Deep Learning | Beginner | 57 h | 4.5 (37) |

4

SPECIALIZATION · 4 COURSES · 75 H

## Reinforcement Learning

University of Alberta (with Alberta Machine Intelligence Institute) | Intermediate | Python

**Stated prerequisites:** 1 year undergraduate CS or 2-3 years software development; Python; basic statistics, linear algebra and calculus

| # | COURSE | LEVEL | LENGTH | RATING (REVIEWS) |
|---|---|---|---|---|
| 1 | Fundamentals of Reinforcement Learning | Intermediate | 15 h | 4.8 (2,911) |
| 2 | Sample-based Learning Methods | Intermediate | 22 h | 4.8 (1,259) |
| 3 | Prediction and Control with Function Approximation | Intermediate | 22 h | 4.8 (851) |
| 4 | A Complete Reinforcement Learning System (Capstone) | Intermediate | 16 h | 4.7 (654) |

### FURTHER READING

**Bayesian Data Analysis, 3rd edition**

Andrew Gelman et al.

The reference for applied Bayesian statistics; Parts I to III alongside UC Santa Cruz.

Free PDF for non-commercial use (author's website) · sites.stat.columbia.edu/gelman/book

**Deep Learning: Foundations and Concepts**

Christopher M. Bishop, Hugh Bishop

Recent textbook (2024): from simple networks to transformers and diffusion models.

Free online version (bishopbook.com) · www.bishopbook.com

**Deep Learning**

Ian Goodfellow, Yoshua Bengio, Aaron Courville

The classic; Part I reviews the mathematics, Part II covers CNNs and RNNs.

Free online (HTML) · www.deeplearningbook.org

**Reinforcement Learning: An Introduction, 2nd edition**

Richard S. Sutton, Andrew G. Barto

The official textbook of the Alberta specialization: chapters follow the order of the four courses.

Free PDF on the authors' website · incompleteideas.net/book/the-book-2nd.html

# 6. Advanced options

These two specializations are not counted in the total. They suit the end of year 3 or after the curriculum.

The theoretical summit of the curriculum: Bayesian networks, Markov random fields, inference and learning. It brings together Bayesian statistics, Markov chains and algorithms.

SPECIALIZATION · 3 COURSES · 171 H

## Probabilistic Graphical Models

Stanford University (Stanford Online) | Advanced | MATLAB, Octave

**Stated prerequisites:** Abstract thinking and math skills, programming in at least one language, basic discrete probability (independence, conditional independence, Bayes' rule)

| # | COURSE | LEVEL | LENGTH | RATING (REVIEWS) |
|---|---|---|---|---|
| 1 | Probabilistic Graphical Models 1: Representation | Advanced | 67 h | 4.6 (1,445) |
| 2 | Probabilistic Graphical Models 2: Inference | Advanced | 38 h | 4.6 (491) |
| 3 | Probabilistic Graphical Models 3: Learning | Advanced | 66 h | 4.6 (305) |

READING

**Probabilistic Graphical Models: Principles and Techniques**

Daphne Koller, Nir Friedman (MIT Press)

The specialization's textbook, written by the instructor.

Paid book · mitpress.mit.edu/9780262013192/probabilistic-graphical-models

Ethics, cybersecurity and visualization: professional practice skills, useful at the end of the curriculum and eligible for CU Boulder academic credit.

SPECIALIZATION · 4 COURSES · 69 H

## Vital Skills for Data Science

University of Colorado Boulder | Intermediate | R | academic credit

**Stated prerequisites:** Some R programming

| # | COURSE | LEVEL | LENGTH | RATING (REVIEWS) |
|---|---|---|---|---|
| 1 | Data Science as a Field | Intermediate | 11 h | 4.3 (41) |
| 2 | Ethical Issues in Data Science | Beginner | 24 h | 4.8 (53) |
| 3 | Cybersecurity for Data Science | Beginner | 19 h | 3.6 (36) |
| 4 | Fundamentals of Data Visualization | Intermediate | 15 h | 4.6 (44) |

# 7. Bibliography

All textbooks except the last can be read free online, under the conditions stated.


1. **Charles R. Severance**. *Python for Everybody: Exploring Data in Python 3*. Free (PDF, HTML, EPUB), CC BY 4.0 licence. www.py4e.com/book
2. **Allen B. Downey**. *Think Python, 3rd edition*. Free online (Jupyter notebooks), CC BY-NC-SA 4.0 licence. allendowney.github.io/ThinkPython
3. **Gilbert Strang, Edwin Herman (OpenStax)**. *Calculus, Volume 1*. Free (PDF and web), OpenStax. openstax.org/details/books/calculus-volume-1
4. **Stephen Boyd, Lieven Vandenberghe**. *Introduction to Applied Linear Algebra: Vectors, Matrices, and Least Squares*. Free (author's PDF), Python and Julia companions. web.stanford.edu/~boyd/vmls
5. **Marc Peter Deisenroth, A. Aldo Faisal, Cheng Soon Ong**. *Mathematics for Machine Learning*. Free (author's PDF). mml-book.github.io
6. **Eric Lehman, F. Thomson Leighton, Albert R. Meyer (MIT)**. *Mathematics for Computer Science*. Free (MIT OpenCourseWare PDF), CC BY-SA licence. ocw.mit.edu/courses/6-1200j-mathematics-for-computer-science-spring-2024/pages/readings
7. **Jeff Erickson (University of Illinois)**. *Algorithms*. Free (author's PDF), CC BY 4.0 licence. jeffe.cs.illinois.edu/teaching/algorithms
8. **Joseph K. Blitzstein, Jessica Hwang (Harvard)**. *Introduction to Probability, 2nd edition*. Free online version (probabilitybook.net), with the Stat 110 course. stat110.hsites.harvard.edu
9. **Hadley Wickham, Mine Çetinkaya-Rundel, Garrett Grolemund**. *R for Data Science, 2nd edition*. Free online, CC BY-NC-ND 3.0 licence. r4ds.hadley.nz
10. **Wes McKinney**. *Python for Data Analysis, 3rd edition*. Open access online (HTML), code under MIT licence. wesmckinney.com/book
11. **G. James, D. Witten, T. Hastie, R. Tibshirani, J. Taylor**. *An Introduction to Statistical Learning, with Applications in Python*. Free PDF on the authors' website. www.statlearning.com
12. **M. P. Deisenroth, A. A. Faisal, C. S. Ong**. *Mathematics for Machine Learning (Part II)*. Free (author's PDF). mml-book.github.io
13. **Andrew Gelman et al.**. *Bayesian Data Analysis, 3rd edition*. Free PDF for non-commercial use (author's website). sites.stat.columbia.edu/gelman/book
14. **Christopher M. Bishop, Hugh Bishop**. *Deep Learning: Foundations and Concepts*. Free online version (bishopbook.com). www.bishopbook.com
15. **Ian Goodfellow, Yoshua Bengio, Aaron Courville**. *Deep Learning*. Free online (HTML). www.deeplearningbook.org
16. **Richard S. Sutton, Andrew G. Barto**. *Reinforcement Learning: An Introduction, 2nd edition*. Free PDF on the authors' website. incompleteideas.net/book/the-book-2nd.html
17. **Daphne Koller, Nir Friedman (MIT Press)**. *Probabilistic Graphical Models: Principles and Techniques*. Paid book. mitpress.mit.edu/9780262013192/probabilistic-graphical-models

# 8. Method and caveats

- **Data source.** Specializations, courses, lengths, levels and ratings come from coursera.org pages, collected on 4 November 2026, and from the “Data Science Specializations Coursera” and “Computer Science Specializations” files.
- **How hours are counted.** Totals add up the length stated for each course. A course that appears in two specializations is counted once. Coursera's header estimates (“3 months at 10 hours a week”) often differ from the real sum of the courses and are not used.
- **Single courses.** Three items are courses taken on their own from a specialization: Calculus for Engineers, Statistical Inference and Hypothesis Testing, and Deep Learning (Illinois Tech). All the others are complete specializations.
- **Academic credit.** The CU Boulder specializations in S4 and the Vital Skills option are marked eligible by Coursera; check recognition conditions with the university.
- **Catalogue changes.** Coursera regularly changes its courses, lengths and ratings. Check each page before enrolling.
- **Textbooks.** Free availability was checked on the authors' or publishers' websites on the same date; licences differ (some are non-commercial only).