COS 302 / SML 305: Mathematics for Numerical Computing and Machine Learning

Szymon Rusinkiewicz

Fall 2020

PRINCETON UNIVERSITY

COS 302

Instructor: Szymon Rusinkiewicz

TAs: Joshua Aduol, Yaniv Ovaida, Jad Rahme, Xinguan Sun, Fangyin Wei

Web page:

`http://www.cs.princeton.edu/courses/archive/fall20/cos302`

Why COS 302?

- COS historically centered on *discrete* math
  - Graph theory, boolean logic, number theory
- But *many* applications of computer science, data analysis, and AI these days require *continuous* math
  - Linear algebra, vector calculus, probability, optimization
  - Used in machine learning, computer vision, natural language processing, computational biology, theoretical neuroscience, computer graphics, …

Why COS 302?

- Suppose you get a summer internship helping self-driving cars estimate how far a car is behind you.
- You might collect data with sensor features and ground-truth distances…

… then use sensor feature vectors as basis functions for your model

… then perform a linear regression assuming Gaussian noise

… then set up an optimization problem to maximize the log likelihood

… then take the gradient of the resulting loss to find the zero

… and find the weights by solving the overdetermined linear system.

Why COS 302?

- Suppose you get a summer internship helping self-driving cars estimate how far a car is behind you.
- You might collect data … truth distances…

This is all *continuous* math!

Is This a Machine Learning Class?

We consider simple applications of machine learning, but in general, no.

This course is about the math you need for machine learning and other areas of computer science that rely on continuous mathematics.

It gathers material that would normally be taught in several different classes, and presents it in a focused way with examples from ML.

COS 302 Topics

- Linear algebra (weeks 1-5)
  - abstractions for reasoning about vector spaces, which are a good model for collections of data

"… by solving the overdetermined linear system …"

COS 302 Topics

- Linear algebra (weeks 1-5)
  - Linear systems
  - Vector spaces
  - Bases
  - Norms and inner products
  - Orthogonality
  - Projections
  - Eigenstuff
  - Cholesky factorization
  - Singular Value Decomposition

COS 302 Topics

- Linear algebra (weeks 1-5)
- Vector calculus (week 7)
  - understanding the geometry of multi-dimensional functions

  "… took the gradient of the resulting loss …"

COS 302 Topics

- Linear algebra (weeks 1-5)
- Vector calculus (week 7)
  - Differentiating functions of vectors

COS 302 Topics

- Linear algebra (weeks 1-5)
- Vector calculus (week 7)
- Probability (weeks 8-10)
  - reasoning about noise and uncertainty

  "… assuming Gaussian noise …"

COS 302 Topics

- Linear algebra (weeks 1-5)
- Vector calculus (week 7)
- Probability (weeks 8-10)
  - Random variables
  - Sampling
  - Independence and dependence
  - Monte Carlo
  - Gaussian distributions
  - Information theory

COS 302 Topics

- Linear algebra (weeks 1-5)
- Vector calculus (week 7)
- Probability (weeks 8-10)
- Optimization (weeks 11-12)
  - finding the best fit to data

"… maximizing log likelihood …"

COS 302 Topics

- Linear algebra (weeks 1-5)
- Vector calculus (week 7)
- Probability (weeks 8-10)
- Optimization (weeks 11-12)
  - Constrained optimization
  - Convex optimization
  - Gradient descent
  - Conjugate gradients

COS 302 Topics

- Linear algebra (weeks 1-5)
- Vector calculus (week 7)
- Probability (weeks 8-10)
- Optimization (weeks 11-12)
- Programming in Python

Python

- Widely used in scientific computing and especially machine learning because of its flexibility and large number of available libraries
  - NumPy, SciPy, Matplotlib, TensorFlow, PyTorch, JAX, etc.
  - Many numerical algorithms happen in efficient Fortran/C/C++ compiled code

- Big differences from Java / C:
  - Dynamically typed (no declarations, variables can change types)
  - Compiled to bytecode at run time (no separate compiler)
  - Uses indentation instead of {braces} for grouping

Python

- We're using "Jupyter" notebooks – in-browser execution with interleaved rich text and figures
- We ask you to run all code on the cloud in Google Colaboratory, to avoid the hassle of installation on your own computer

https://colab.research.google.com

Course Mechanics: Lectures and Precepts

- With the exception of today, lectures will be pre-recorded (typically in several chunks)
- Monday and Wednesday "class" meetings will be used for Q&A
  - Ask for explanations of lectures, readings
  - Material that is suitable for broad discussion among everyone – no debugging!
  - Proposing to move the Wednesday class to 4 PM or 9 PM EDT to accommodate students in various time zones – please fill out questionnaire on website!
- Thursday and Friday precepts are a combination of new material and help on homework assignments – please attend!

Course Mechanics: Book

Mathematics for Machine Learning

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

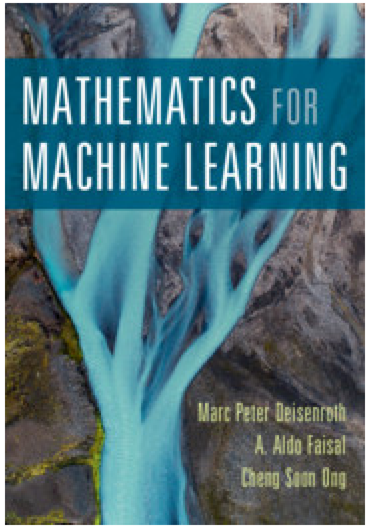

Available in hardcopy and online:

https://mml-book.github.io/

Course Mechanics: Homework

- Due weekly on Mondays
  (12:00 noon Eastern time – before the start of lecture)
- Count for 60% of grade (remainder: two exams)
- Can be submitted 1 day late for 90% credit. No other late work.
- Lowest grade dropped
- Written up in LaTeX and submitted via Gradescope

Course Mechanics: Collaboration

- You ***must not*** share code, writeups
- You ***must not*** discuss solutions, or search Internet for solutions
- You ***may*** discuss high-level concepts with others
- You ***should*** use the official online documentation for Python, NumPy, SciPy, Matplotlib, LaTeX, etc.
- You ***must*** state in your writeup who you discussed problems with

Course Mechanics: Q&A

- We encourage asking questions during “class” time, but will be using ed discussion for additional Q&A. Please direct all questions there.

- Feel free to answer each others’ questions (we will monitor and endorse students’ answers) but keep in mind collaboration policy.

- You can also mark questions private, if necessary.

Course Mechanics: Questionnaire

- Give us an idea of your background, and where you are physically
- As mentioned before, we might change future Wednesday class times (though not for *this* Wednesday)

- Available from course website – please fill it out today!

COS 302

Instructor: Szymon Rusinkiewicz

TAs: Joshua Aduol, Yaniv Ovaida, Jad Rahme, Xinguan Sun, Fangyin Wei

Web page:

`http://www.cs.princeton.edu/courses/archive/fall20/cos302`