CS206: Principles of Scientific Computing

Xiaohui Xie
Department of Computer Science
University of California, Irvine

Course information

Prerequisites: multivariate calculus, linear algebra

Textbooks:

- Numerical Linear Algebra by Trefethen and Bau
- Convex Optimization by Boyd and Vandenberghe
- **Mathematics for Machine Learning** https://mml-book.github.io/

Course website:

- https://www.ics.uci.edu/~xhx/courses/CS206/

Piazza:

- Course announcements and other communications will be carried out through piazza. Please sign up.

Grading

Grading based on:

- Midterm exam (40%)
- Final exam (50%)
- Class participation (10%)

Homework assignments:

- Not graded (no need to submit)
- Some exam questions will be taken directly from the assignments.

Scientific Computing

- What is Scientific Computing?
  - Design and analysis of algorithms for numerically solving mathematical problems in science and engineering
  - sometimes called *numerical analysis*
- What's special about Scientific Computing?
  - Deals with continuous quantities/variables
  - Considers effects of approximations
- Why Scientific Computing?
  - Simulation of natural phenomena
  - Solving real-world engineering problems

Typical topics in scientific computing

- Numerical linear algebra
- Optimization
- Numerical integration and differentiation
- Solving ODEs - initial value problems, boundary value problems
- Solving PDEs

Topics covered this quarter

- Numerical linear algebra
- Optimization
- Numerical integration and differentiation
- Solving ODEs - initial value problems, boundary value problems
- Solving PDEs

And other topics important for Machine Learning

Problems you will be able to solve by the end of this quarter

- Linear equations: $Ax = b$, where $A \in R^{m \times m}$ and $b \in R^m$
- Least square problems: Given $A \in R^{m \times m}, m \geq n, b \in R^m$, find $x \in R^n$ such that $\|b - Ax\|_2$ is minimized.
- Find eigenvalues and eigenvector of a square matrix A - $A = Q \Sigma Q^{-1}$
- Find singular value decomposition (SVD) of a matrix A - $A = U \Sigma V^T$

Solving optimization problems

Mathematical **optimization problem**:

$$\begin{aligned} \text{minimize} \quad & f_0(\mathbf{x}) \\ \text{subject to} \quad & f_i(\mathbf{x}) \leq \mathbf{b}_i, \quad i = 1, \cdots, m \end{aligned}$$

where

- $\mathbf{x} = (x_1, \cdots, x_n) \in \mathbb{R}^n$: optimization variables
- $f_0 : \mathbb{R}^n \to \mathbb{R}$: objective function
- $f_i : \mathbb{R}^n \to \mathbb{R}$: constraint function

Optimal solution $\mathbf{x}^*$ has smallest value of f_0 among all vectors that satisfy the constraints.

Foundations of machine learning

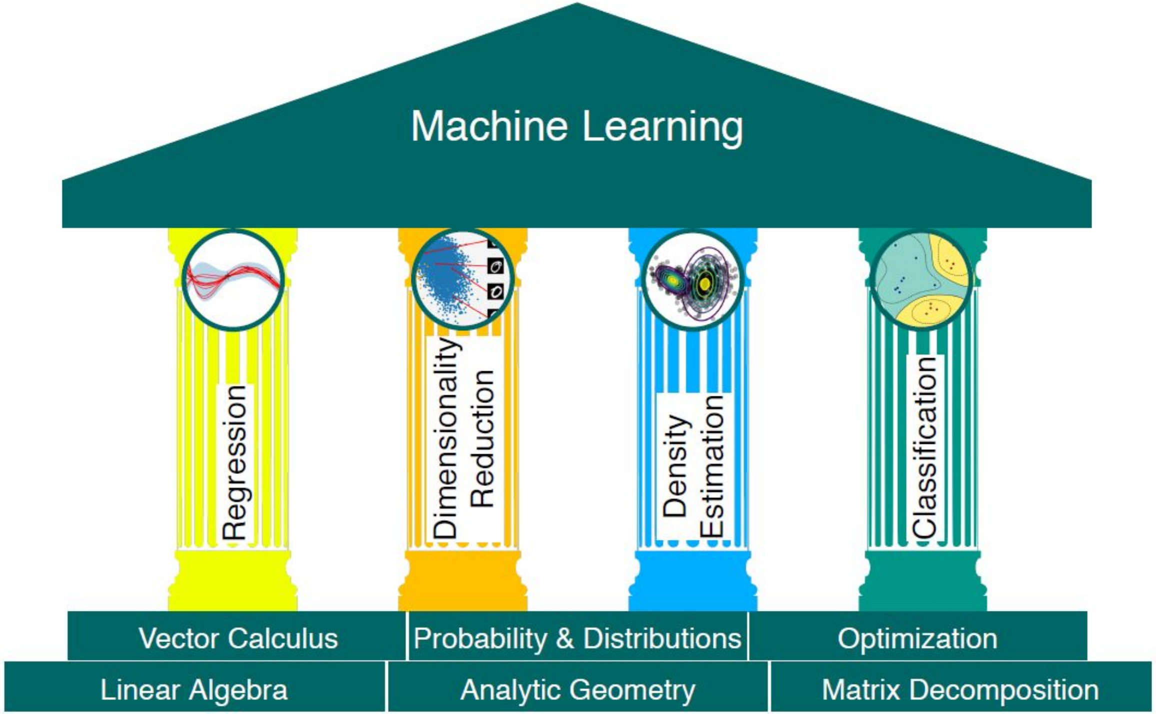

Linear algebra

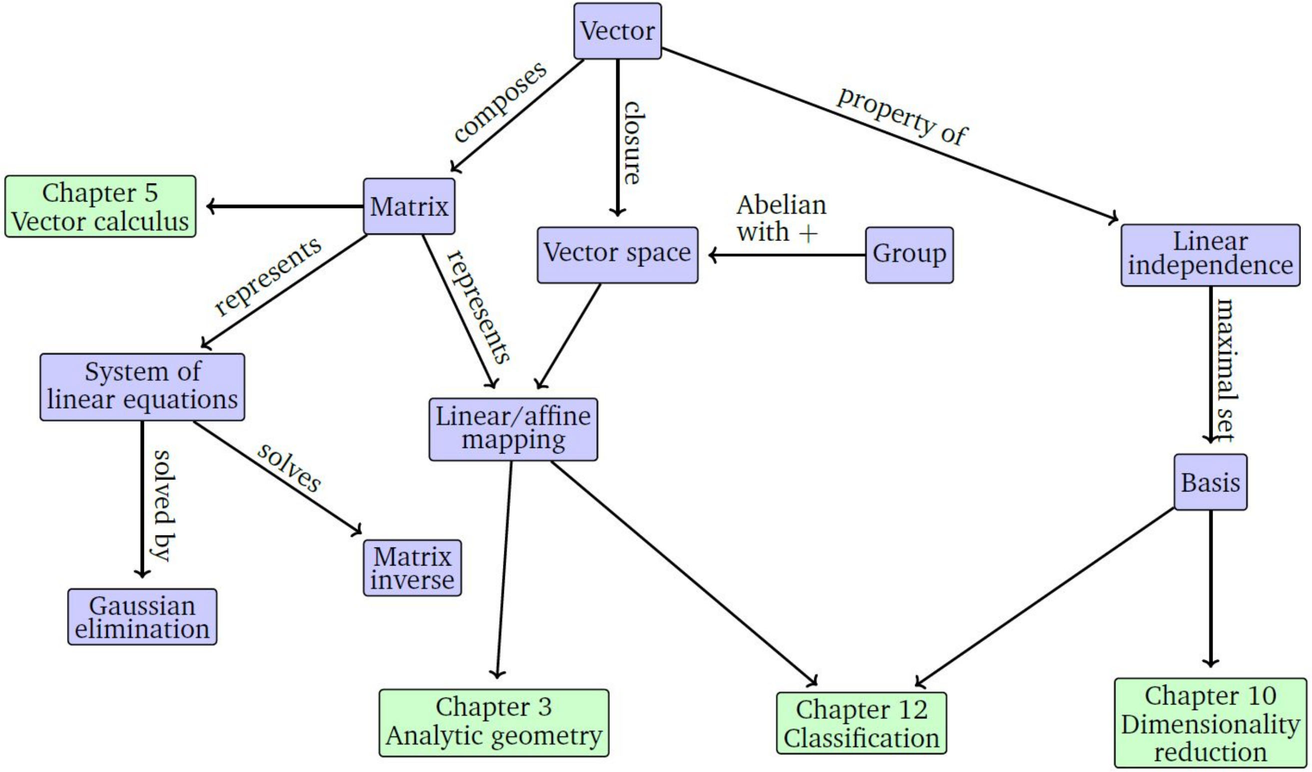

Analytic Geometry

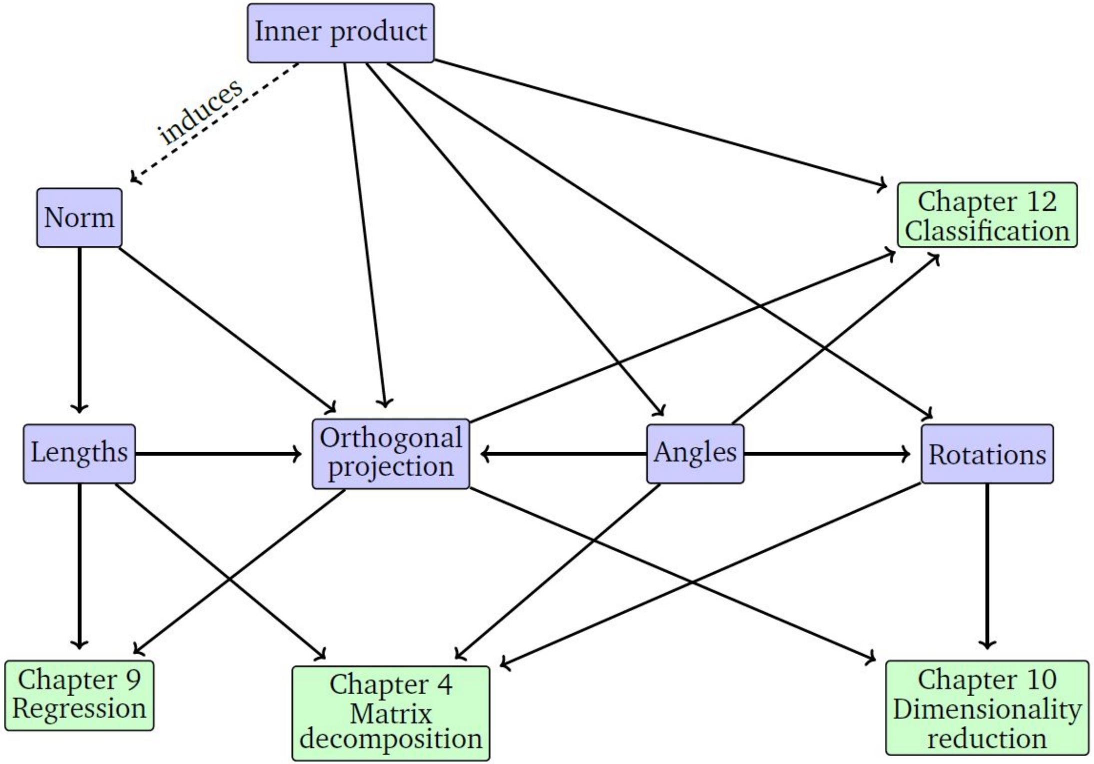

Matrix Decompositions

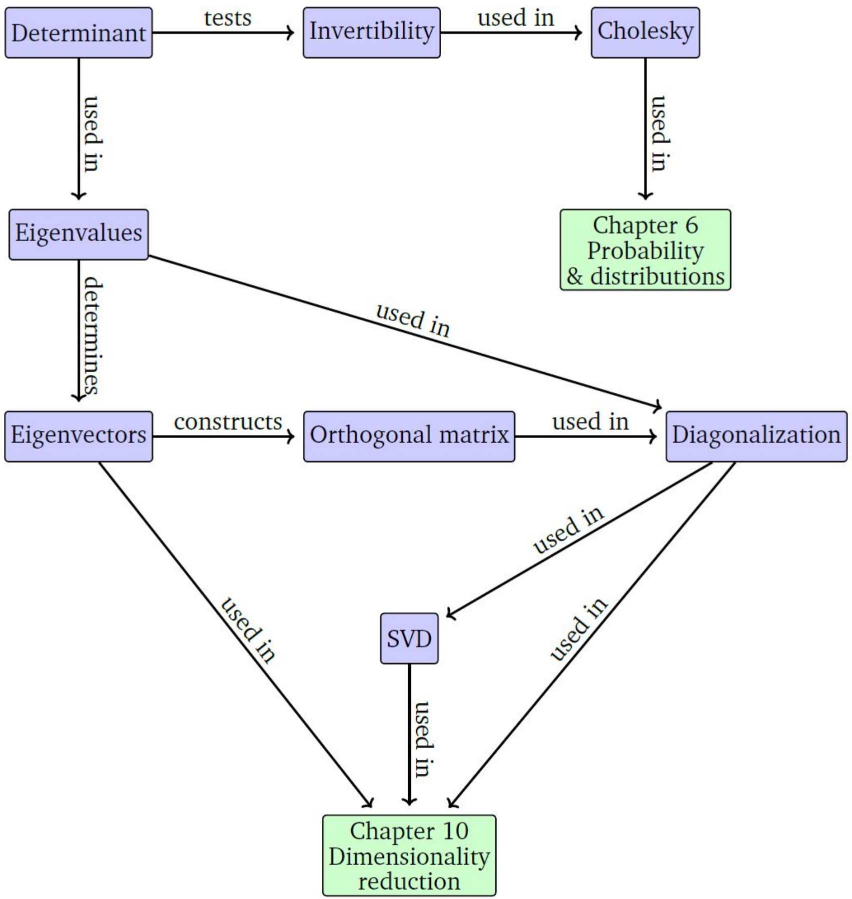

Vector Calculus

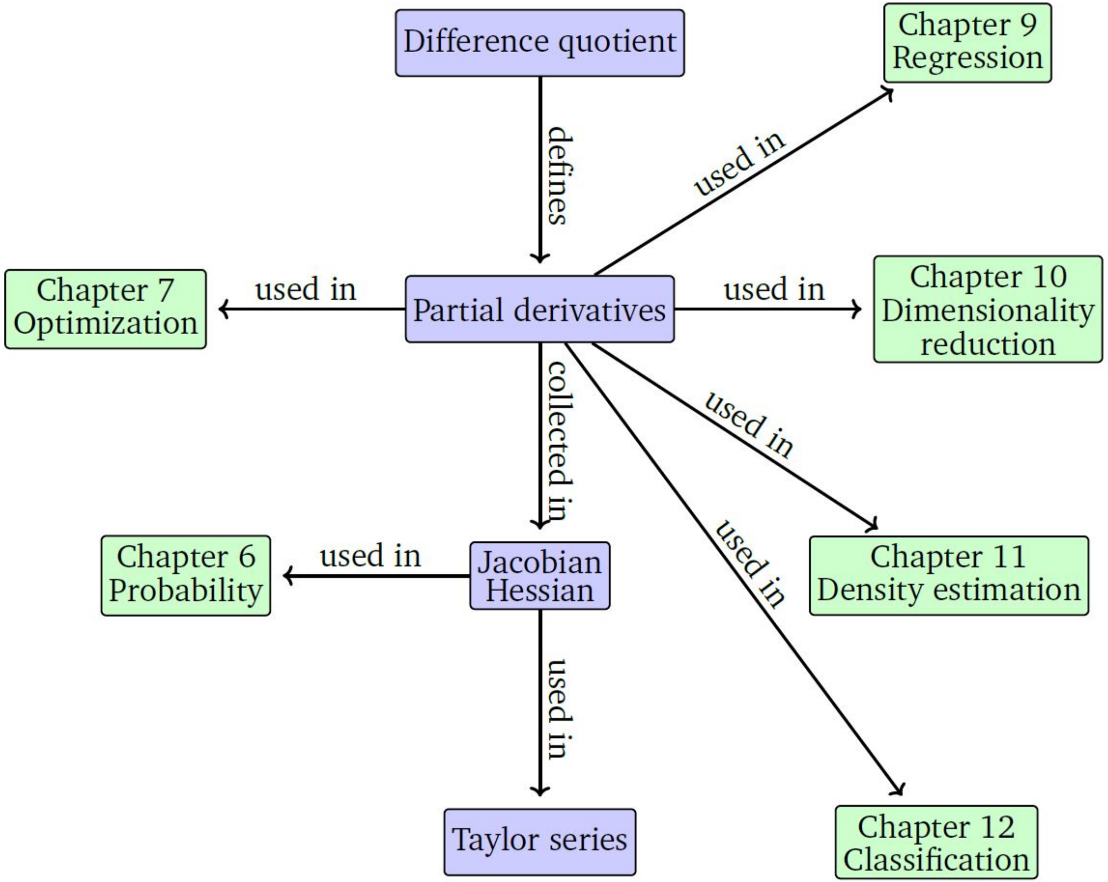

Probability and Distributions

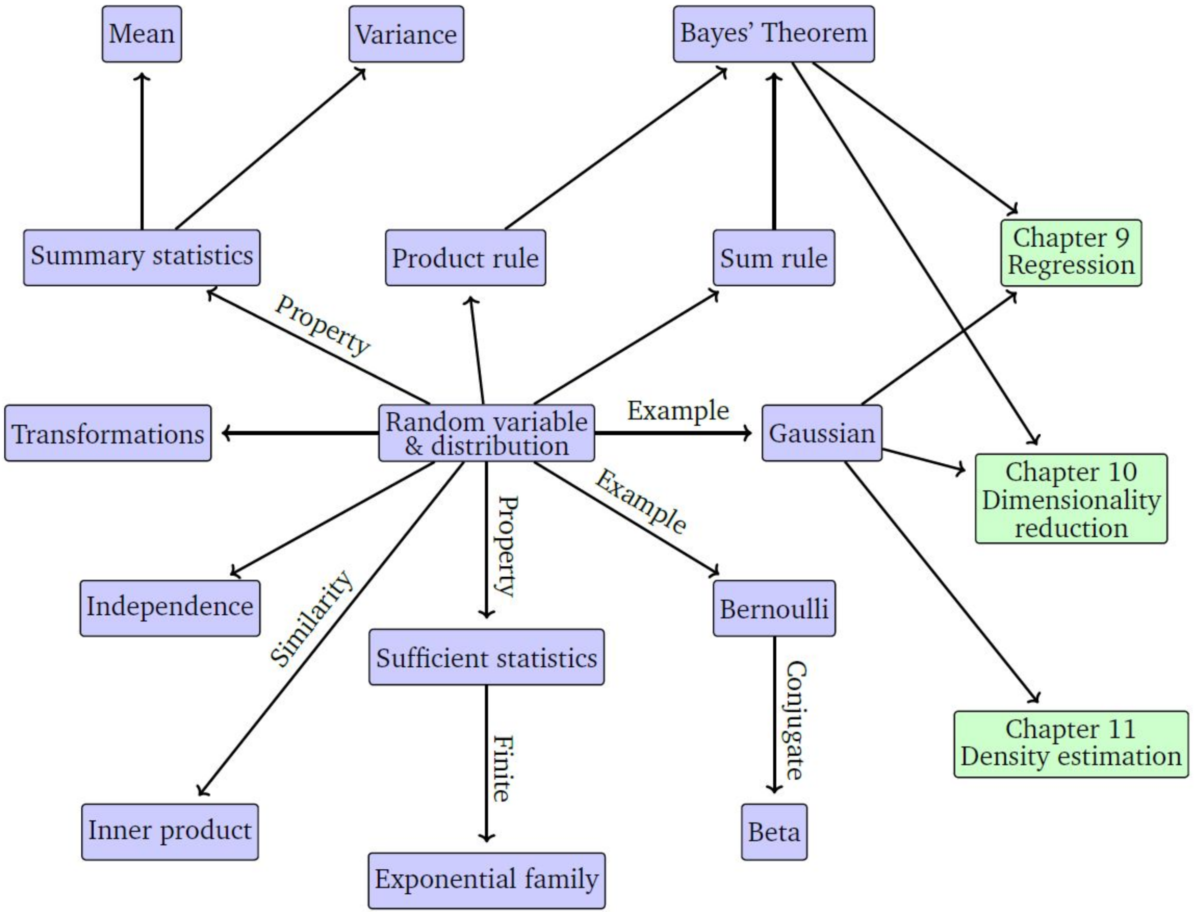

Continuous Optimizations