Applied Bayesian Nonparametrics

Special Topics in Machine Learning
Brown University CSCI 2950-P, Fall 2011

Instructor: *Erik Sudderth*

Machine Learning Problems

	Supervised Learning	*Unsupervised Learning*
Discrete	classification or categorization	clustering
Continuous	regression	dimensionality reduction

- Bayesian nonparametric (BNP) models lead to more flexible, data-driven methods for all of these problems
- Primary focus is on unsupervised learning

What do you want to learn about?

Course Prerequisites

- A course in modern statistical machine learning
  - Brown CS 195F: Introduction to Machine Learning
  - Brown APMA 261: Recent Applications of Probability and Statistics
  - Possibly other classes or experience...
- Programming abilities for the course project
- Readings will require “mathematical maturity”
- Insufficient background by themselves:
  - Brown CS 141: Introduction to AI
  - Traditional undergrad statistics (APMA 165/166)

Background Material

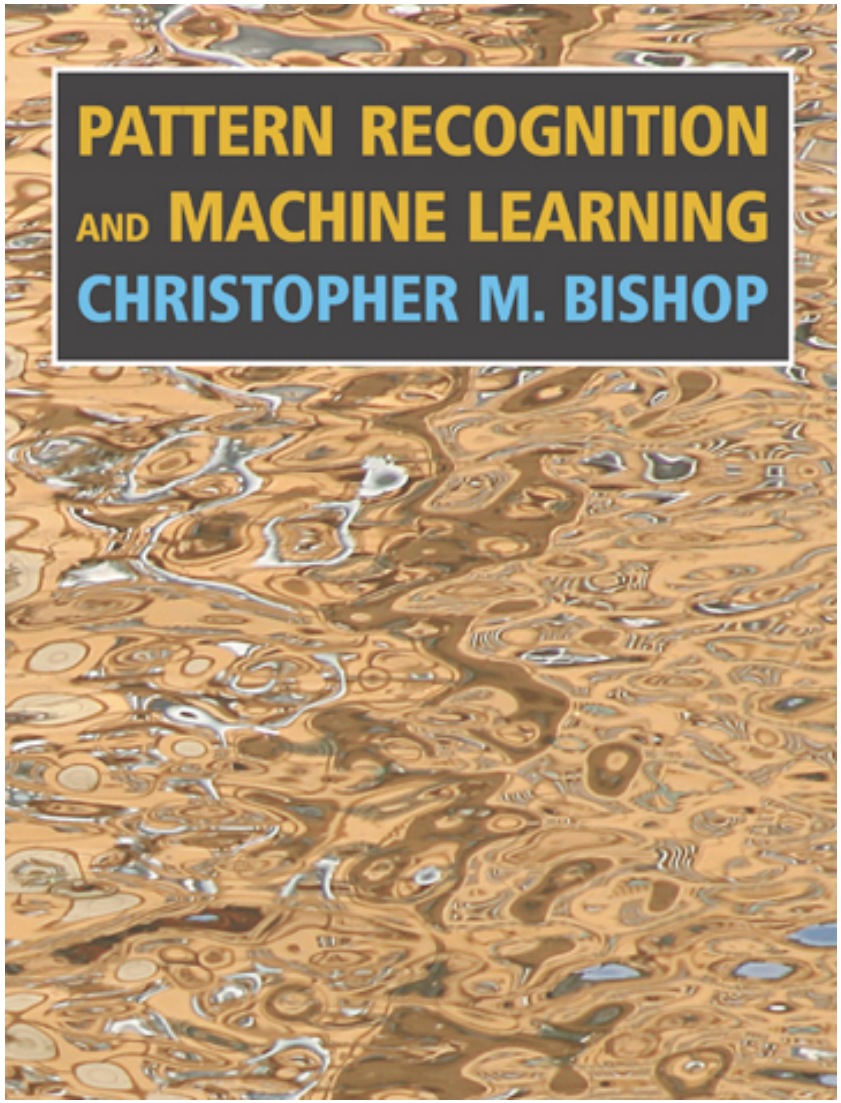

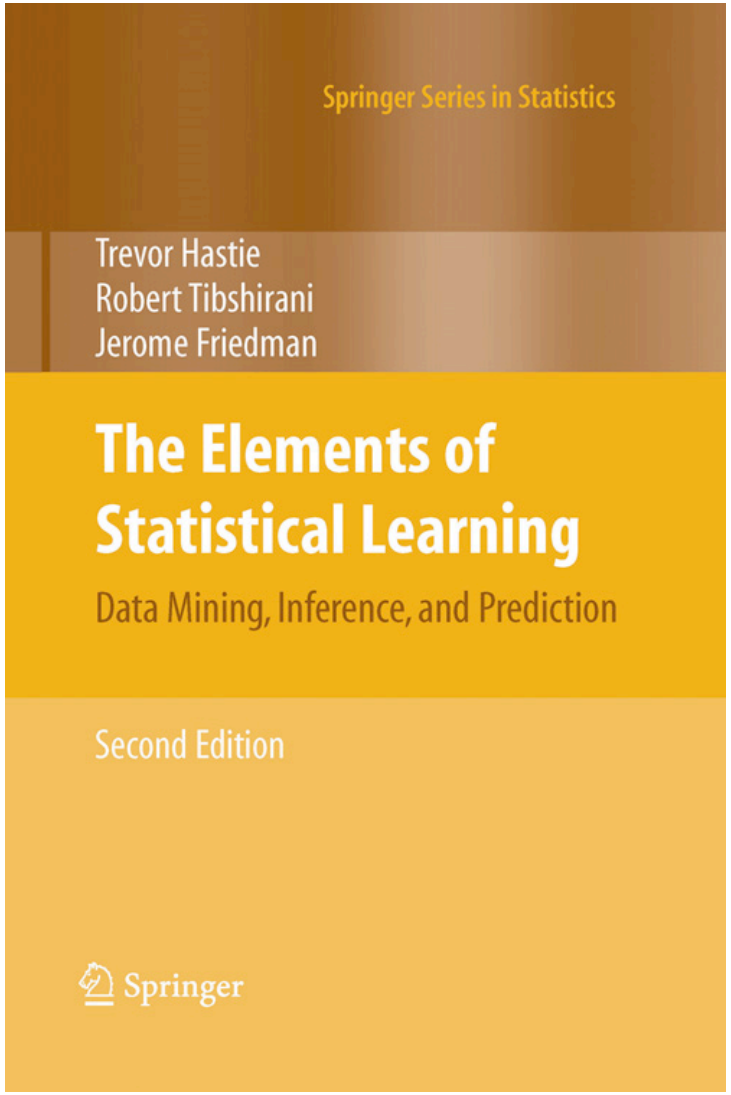

Course Format & Readings

- Two 80-minute meetings per week
  (Tuesdays & Thursdays, 2:30-3:50pm, CIT 506)
- Each day will have three 25-minute segments:
  - Average ML conference paper: 1 segment
  - Average statistics journal paper: 2 segments
  - Exceptions to every rule…
- Typical reading for a single class:
  one journal paper & one related conference paper
- Presentation & discussion of some segments will be led by instructor, others by students

Course Evaluation

Class Participation: 30%

- Attend class and participate in discussions
- Prepare summary overview presentation, and lead class discussion, for 2 segments
  - Most journal papers will be collaboratively presented
  - Prof. Sudderth will lecture for the remaining segments
- Upload brief comments about one assigned reading before each lecture (due at 8am)

Final Project: 70%

- Proposal: 1-2 pages, due in late October (10%)
- Presentation: ~10 minutes, during reading period (20%)
- Conference-style technical report (40%)

Reading Comments

The Good: 1-2 sentences

- What is the most exciting or interesting model, idea, or technique described here? Why is it important?
- Don't just copy the abstract - what do *you* think?

The Bad: 1-2 sentences

- No method is perfect, and many are far from it!
- What is the biggest weakness of this model or approach?
- Problems could be a lack of empirical validation, missing theory, unacknowledged assumptions, …

The Ugly: 1-2 sentences

- Poorly written or unclear sections of the paper: terse explanations, steps you didn't follow, technical errors, etc.
- What would you like to have explained in class?

Final Projects

Best case: Application of course material to your own area of research

Key Requirements: Novelty, use of BNP models

- Identify a family of BNP models suitable for a particular application, try baseline learning algorithms
- Propose, develop, and experimentally test a new type of learning algorithm for some existing BNP model
- Experimentally compare different models or algorithms on an interesting, novel dataset
- Survey the latest advances in an area of BNP theory or application *which is not already covered by the course*
- **There will not be a list of projects to choose from.** You must propose your own (with the instructor's advice)

A Quick Poll

Administration

Registration: E-mail sudderth@cs.brown.edu with

- Your name and CS logon
- Your department, major, and year
- Your background in statistical machine learning
  - If you've taken CS195-F or APMA261, just say so
  - Otherwise, a few sentences about your background

Readings for Tuesday:

- Rasmussen & Williams, *Gaussian Processes for Machine Learning*, Chap. 2-3 (except 3.5, 3.6, 3.9)
- No comments required for Tuesday's lecture

Course webpage: Up Friday, linked from my webpage

Paper comments & coordination: Details Tuesday