CS540 Machine learning
Lecture 1
Introduction

Administrivia

- Class web page

  www.cs.ubc.ca/~murphyk/Teaching/CS540-Spring10
- Join http://groups.google.ca/group/**cs540-spring10**
- Office hours: Fri 3.30-4.30 CS 187, or by appointment
- New rooms: Tuesdays Forestry 1613 (next to Tim Hortons).

  Thursdays Angus 292, a brand new room in the Sauder business school.
- Midterm: Thu Mar 4th
- Last class: Thur Apr 15th
- Final project due Tue Apr 27th

Map to Angus 292

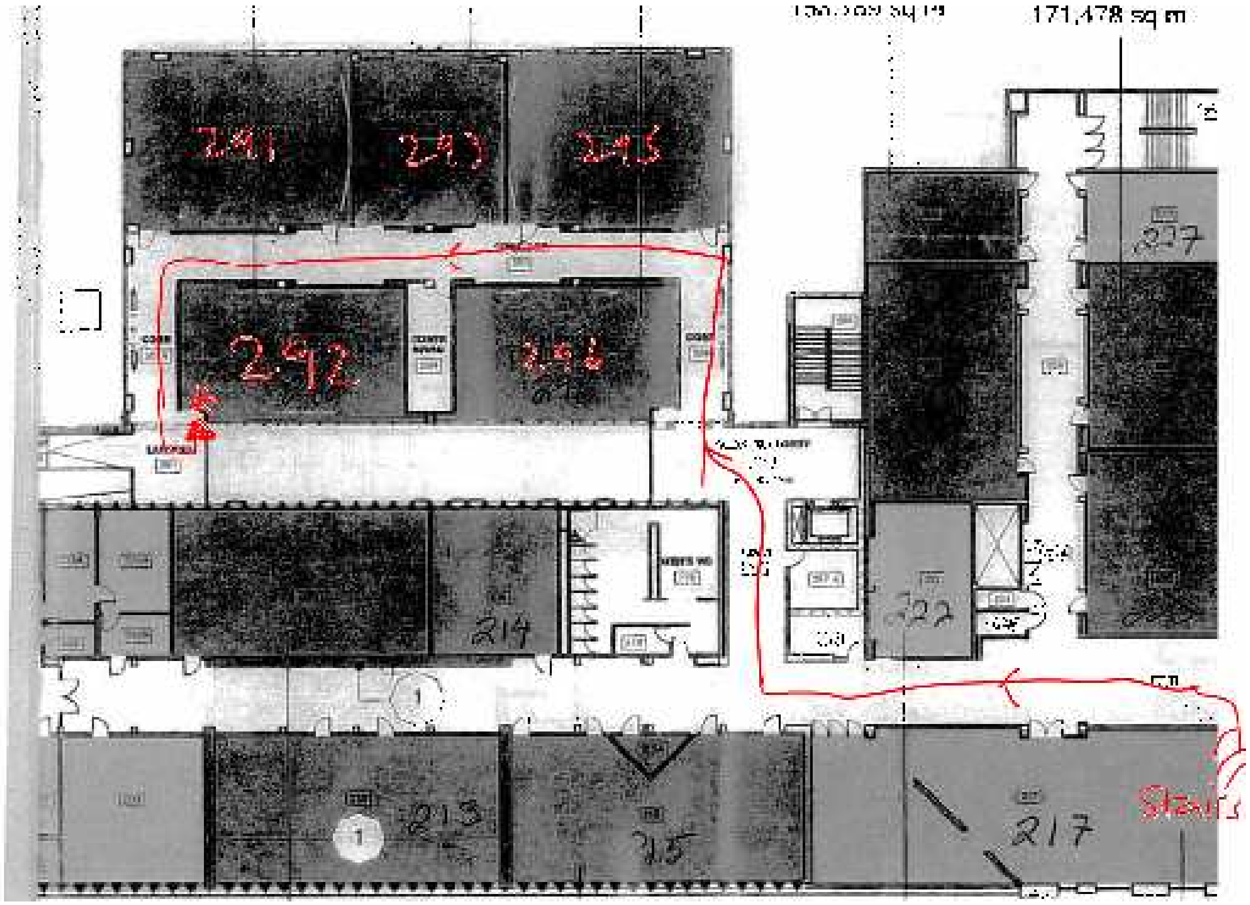

Grading

Grading

- Midterm (open-book):35%
- Final project: 40%
- Weekly Assignments: 25%

Homeworks

Weekly homeworks, out on Tue, due back on Tue

- Collaboration policy:
  - You can collaborate on homeworks if you write the name of your collaborators on what you hand in; however, you must understand everything you write, and be able to do it on your own (eg. in the exam!)
- Sickness policy:
  - If you cannot do an assignment or an exam, you must come see me in person; a doctor's note (or equivalent) will be required.

Workload

- This class will be quite time consuming.
- Attending lectures: 3h.
- Weekly homeworks: about 6h.
- Weekly reading: about 6h.
- Total: 15h/week.

- If the pace is too fast, why not take Stat406 with me this spring instead?

Pre-requisites

- You should know
  - Basic multivariate calculus e.g.,

$$\nabla_{\mathbf{x}} \mathbf{x}^T \mathbf{a} = \mathbf{a}$$

  - Basic linear algebra e.g.,

$$A\vec{u}_i = \lambda_i \vec{u}_i$$

  - Basic probability/ statistics e.g.

$$Cov(X,Y) = E\left[(X - EX)(Y - EY)\right] = E[XY] - E[X]E[Y]$$

  - Basic data structures and algorithms (e.g., trees, lists, sorting, dynamic programming, etc)

Textbook

- “Machine learning: a probabilistic approach”
- Draft copies available from Copiesmart in the UBC Village (next to Macdonald’s) for $56.50
- Extra credit (up to 5% of your grade) for finding errors (5 points) or typos (1 point) – more details later
- Please bring your book to every class (and the exam).
- Ch1 is online for free, so you can see if you want to take or drop the course

Other good books

If you want a book that is already “debugged”, see one of these

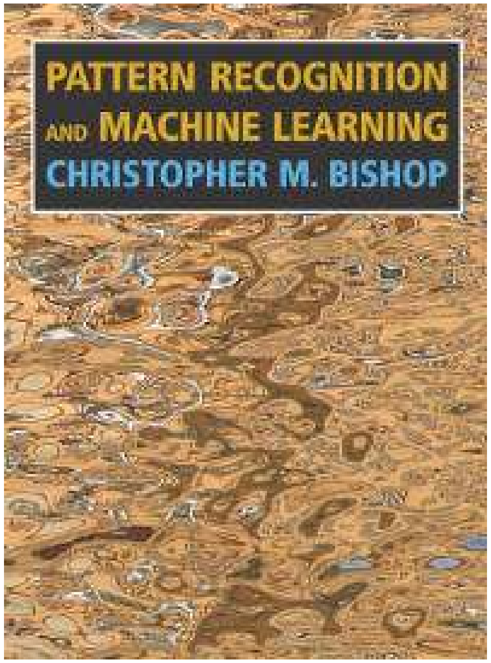

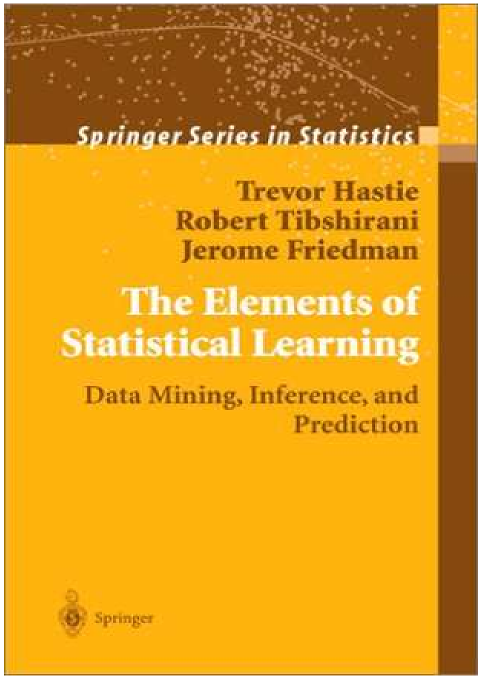

Free online

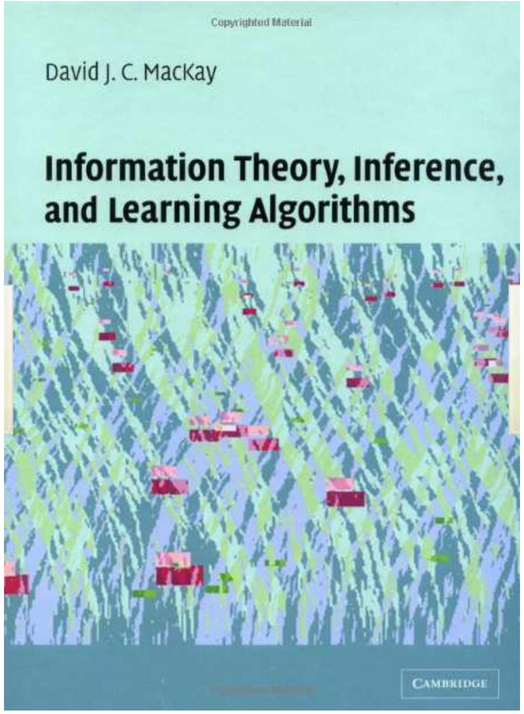

Free online

If you plan to “major” in machine learning, you should buy and read all of these!

Other good books

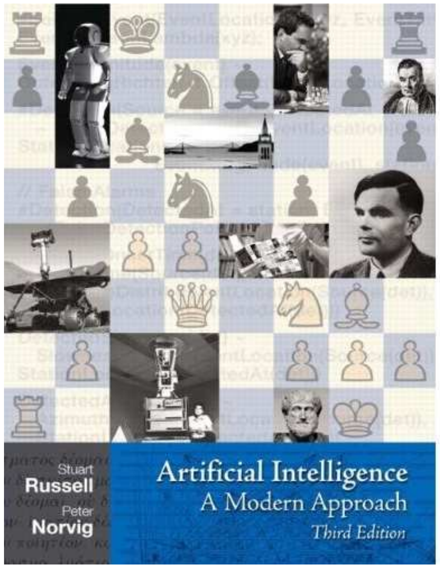

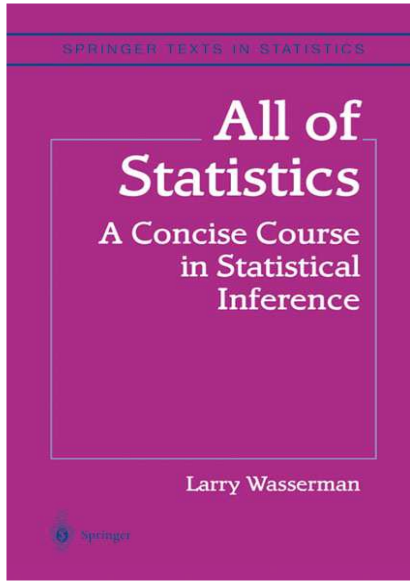

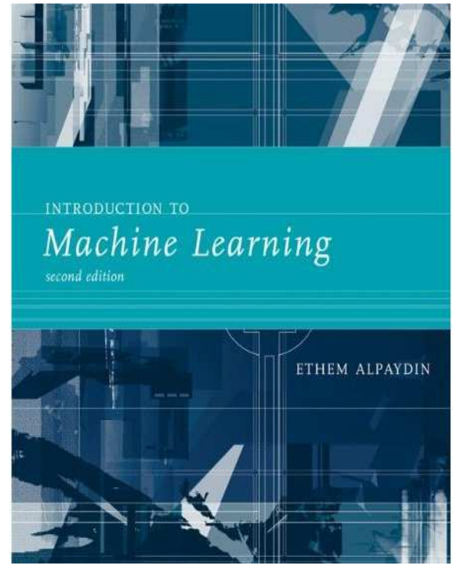

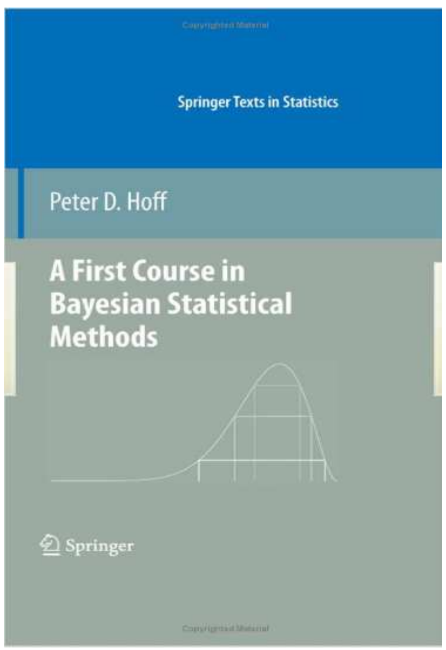

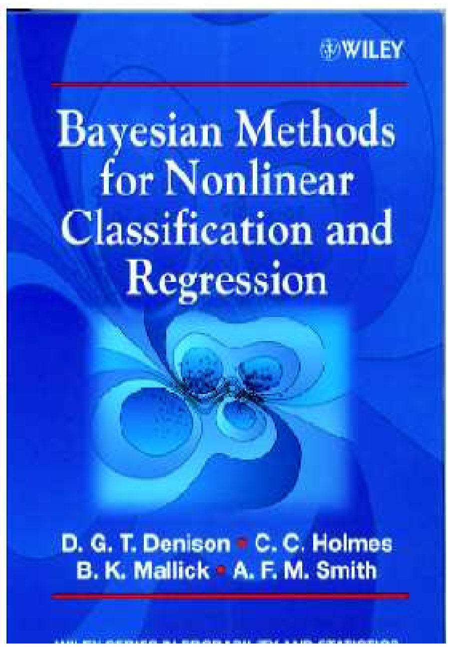

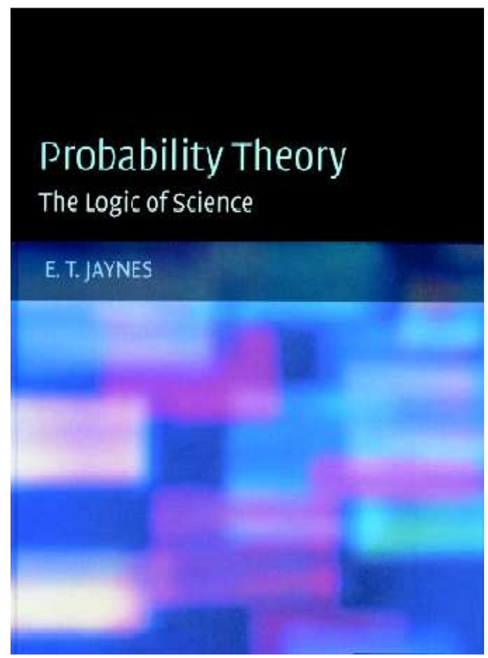

Matlab

- Matlab is a mathematical scripting language widely used for machine learning (and engineering and numerical computation in general).
- Everyone should have access to Matlab. If not, ask for a CS guest account.
- You can buy a student version for $170 from the UBC bookstore. Please make sure it has the Stats toolbox.
- Matt Dunham has written an excellent Matlab tutorial which is on the class web site – please study it carefully!
- On Wed 6th, 4-5pm, LSK 302, I will offer a brief Matlab tutorial as part of Stat406; all welcome

Software

- My textbook is accompanied by Matlab code, called PMTK3 (probabilistic modeling toolkit), which is on the book’s webpage. It changes frequently.
- Currently there is not 100% consistency between the book and the code; this will be fixed during the semester.
- I also have 2 older code projects associated with the book (PMTK and PMTK2) which you can ignore, since they are deprecated
- We will also use several other Matlab packages such as Netlab, GPML, etc.

Learning objectives I

- By the end of this class, you should be able to
- 1. For any given data analysis problem, be able to identify a suitable probabilistic model from those that are currently available, or if none exists, to devise one of your own.
- 2. If the model is novel, be able to derive a suitable algorithm. This includes deriving the update equations for fitting the parameters (eg gradient expressions for the MLE/MAP) and inferring the hidden states (eg variational message passing updates, or full conditionals for Gibbs sampling).
- 3. If the model is novel, be able to implement the above equations in clear and efficient code (preferably Matlab).
- 4. Be able to evaluate the performance of your model/ algorithm in an objective way, and compare to other models/ algorithms.

Learning objectives II

- 5. Be able to write a clear and concise report summarizing the model, algorithm and results, with a suitable mix of English, Math and (properly labeled) figures. The report should be accompanied by code which enables the reader to easily reproduce all the results and figures.
- 6. Be able to give a clear and concise oral presentation summarizing your work, and answer oral questions on it, including defending your choice of model, algorithm or evaluation method, and discussion of possible alternatives.

Lecture style

- This year, rather than spend time making lecture slides, I will just project my book onto the screen, and use the blackboard.
- Each “lecture” then becomes just an executive summary of relevant parts of the book.
- I will try to include some in-class demos and/or interactive activities (in the spirit of the Carl Weiman initiative)
- You are expected to carefully read the book either before and/or after each lecture to learn all the details – the relevant sections will be listed on the class web page.
- Homeworks are an essential aid to learning the material; assessing performance is a relatively minor concern