CMPT 726: Machine Learning

Instructor: Greg Mori

Fall 2019 SYLLABUS

Overview

Machine learning is the study of computer algorithms that improve automatically through experience. Machine learning algorithms play an important role in industrial applications and commercial data analysis. The goal of this course is to present students with both the theoretical justification for, and practical application of, machine learning algorithms. Students will gain hands-on experience with major machine learning tools and their applications to real-world data sets. This course will cover techniques in supervised and unsupervised learning, the graphical model formalism, and algorithms for combining models. This course is intended for graduate students who are interested in machine learning or who conduct research in fields which use machine learning, such as computer vision, natural language processing, data mining, bioinformatics, and robotics. No previous knowledge of pattern recognition or machine learning concepts is assumed, but students are expected to have, or obtain, background knowledge in mathematics and statistics.

Administrivia

Lectures: Monday 4:30-7:20 in B9200

TA office hours: see website

Course website:

`http://www.cs.sfu.ca/~mori/courses/cmpt726`

Lecture Schedule (subject to change)

Sept. 9: Introduction (Ch. 1)

Sept. 16: Linear Models for Regression (Ch. 3)

Sept. 23: Linear Models for Classification (Ch. 4)

Sep. 30, Oct. 7: Deep Learning (Neural Networks) (Ch. 5)

Oct. 21, 28: Graphical Models (Ch. 8)

Nov. 4, 18: Sequential Data (Ch. 13), Recurrent Neural Networks

Nov. 25: Variational Inference, Variational Auto-Encoders, Generative Adversarial Networks

Dec. 2: Exam

Dec. 8: Poster Session **4pm-7pm downtown Vancouver** (tentative)

Grading

Evaluation will be based on individual programming and written assignments, an exam, and a project (in groups of up to 5).

- 30% Assignments
- 30% Exam
- 40% Final project

Assignments

There will be three assignments, worth 10% each. Assignment dates:

- A1: Regression (out Sept. 16, due Oct. 4)
- A2: Classification / deep learning (out Oct. 7, due Oct. 25)
- A3: Graphical models / recurrent neural networks (out Oct. 28, due Nov. 15)
- FINAL PROJECT: poster session Dec. 8, report due Dec. 13 11:59pm

All assignments are to be done individually.

Late policy

Students will be permitted 3 grace days to use at their discretion over the trimester. Late days are counted from the time an assignment is due, rounded up to the nearest whole day. For example, if an assignment is due on Friday at 3:30am, and is submitted on Saturday at 5pm, 2 grace days will be used.

IMPORTANT: Other than the 3 grace days, late assignments will not be accepted, and will receive zero marks. If you have extenuating medical circumstances, please consult SFU's policies and discuss with the instructor.

IMPORTANT: Grace days may not be used for the final project.

Textbooks

REQUIRED:

- Pattern Recognition and Machine Learning. Christopher M. Bishop, Springer (2006)

REFERENCE:

- Deep Learning. Ian Goodfellow, Yoshua Bengio and Aaron Courville (`http://www.deeplearningbook.org`)
- The Elements of Statistical Learning. Trevor Hastie, Robert Tibshirani, and Jerome Friedman, Springer-Verlag (2001)
- Information Theory, Inference, and Learning Algorithms. David MacKay, Cambridge University Press (2003) (available online at `http://www.inference.phy.cam.ac.uk/mackay/itila/`)
- Machine Learning: a Probabilistic Perspective. Kevin Murphy, MIT Press (2012) `http://www.cs.ubc.ca/~murphyk/MLbook/index.html`

Academic Honesty

Academic Honesty plays a key role in our efforts to maintain a high standard of academic excellence and integrity. Students are advised that ALL acts of intellectual dishonesty are subject to disciplinary action by the School; serious infractions are dealt with in accordance with the Code of Academic Integrity and Good Conduct (S 10.01) (`http://www.sfu.ca/policies/gazette/student/s10-01.html`).