# CS[45]783: Machine Learning
# Fall 2019 Syllabus

Prof. Christopher Crick

**Abstract**

This course presents a probabilistic, statistical approach to automated pattern discovery applied to large datasets. What can we predict about the future, given data about the past? What kind of models can we construct with this information, and how can we assess these models' behavior and reliability? Finally, how can we represent this data and devise tools for discovering these models, efficiently and effectively? This class does not focus on tools for data analytics, nor solely on the mathematical formulations underlying statistical processing, but also on the development and analysis of learning algorithms.

## Professor

- Dr. Christopher Crick
- Office: MSCS 213
- Lab: MSCS 214
- Office hours: T 10-11:30 and 1-7, or by appointment, or walk-in
- Email: chriscrick@cs.okstate.edu

## Course Meetings

- MW 4-5:15, Human Sciences 024

## Text

- Trevor Hastie, Robert Tibshirani and Jerome Friedman, *The Elements of Statistical Learning: Data Mining, Inference and Prediction.* Available free online at `http://web.stanford.edu/~hastie/ElemStatLearn/`
- Christopher M Bishop, *Pattern Recognition and Machine Learning*

## Grading

- Assignments: 60%
- Midterm Exam: 10%
- Final Project: 30%

## Grade Breakdown

- A: 90%
- B: 80%
- C: 70%
- D: 60%
- I reserve the right to curve these percentages downwards if necessary, but they will not be curved upwards. If you score 90.0%, you will earn an A.

## Topics Covered

- $k$-nearest neighbors
- Generative models
- Decision theory
- Graphical models
- Frequentism
- Naive Bayes
- Bayesian learning
- Generative Gaussian classification
- Multivariate Gaussian distributions
- Linear and logistic regressions
- Gradient descent
- Feature induction
- Feature selection
- Neural networks
- Deep learning
- Sparsity and kernels
- Support vector machines
- Decision trees and forests
- Clustering
- $k$-means
- Expectation maximization
- Hidden Markov models
- Principal components analysis

## Policies

- Assignments will ordinarily be due on Mondays at midnight. Your code should be handed into the Canvas dropbox (canvas.okstate.edu) as a single .zip file. Extensions to the due date will only be considered in the case of legitimate academic or medical conflicts, and must be obtained in advance.

- Assignments must be demonstrated to Prof. Crick during his Tuesday office hours. You will be expected to walk through and explain your code, as well as demonstrating its behavior. Most students will be able to do this on their own laptops, but if this is a problem, you may arrange to demonstrate your code on one of the department's lab computers. If a scheduling difficulty arises that would make it difficult to attend office hours, you can schedule a meeting with Prof. Crick at some other point during the same week.
- There will be a single midterm exam which accounts for 10% of your grade. You will be permitted handwritten notes during the exam.
- Academic integrity is taken very seriously. You are permitted (and indeed encouraged) to discuss the course material with fellow students in general terms, but the programs you write must be your own. **Code copied from each other or found on the net will result in an automatic zero for the assignment**, and depending on the egregiousness of the offence may result in earning an 'F!' for the course and facing academic disciplinary measures.
- That said, you are welcome to copy code from your own previous assignments or from programming snippets that we go over in lecture. You are also expected to look for tutorials and examples from online Python, Numpy, Scipy and Tensorflow documentation. However, this should be used as a syntax reference, not as a source for large chunks of code.

## Class schedule

**Note**: Professor Crick will be at a conference on Monday, Oct. 28, and class will be cancelled. That lecture will be made up on Friday, Nov. 1, at the usual class time.

- August 19 (M): Course introduction. PRML 1.1-1.5. ESL 2.1-2.5.
- August 21 (W): Python and Numpy.
- August 26 (M): Nearest neighbor classification. PRML 8.1-8.2, ESL 8.2.
- August 28 (W): Tuning and validation.
- September 2 (M): **No class** (Labor Day). Assignment 1 due. PRML 2.1-2.2, ESL 8.3.
- September 4 (W): Directed graphical models.
- September 9 (M): Maximum likelihood estimation. PRML 2.3, 4.2.
- September 11 (W): Bayesian parameter estimation.
- September 16 (M): Correlation. PRML 3.1-3.3, ESL 3.2, 3.4.
- September 18 (W): Multivariate Gaussians.
- September 23 (M): Linear regression. Assignment 2 due. PRML 4.1, 4.3, ESL 4.1-4.4.
- September 25 (W): Discriminative models.
- September 30 (M): Gradient descent. PRML 5.1-5.3, 5.5, ESL 11.
- October 2 (W): Neural networks.
- October 7 (M): Tensorflow. PRML 6.1-6.2, 6.4. ESL 3.3.
- October 9 (W): Feature selection.
- October 14 (M): Midterm review.
- October 16 (W): Midterm.
- October 21 (M): Kernel functions. PRML 7.1, 14.4, ESL 12.2-12.3. Assignment 3 due.
- October 23 (W): Gaussian processes.

- October 28 (M): **No class**.
- October 30 (W): Support vector machines. PRML 9.3-9.4.
- November 1 (F): Unsupervised learning.
- November 4 (M): Expectation Maximization. PRML 13.1-13.2. PRML 12.1, ESL 14.5.
- November 6 (W): Hidden Markov Models.
- November 11 (M): Viterbi algorithm.
- November 13 (W): Principal component analysis.
- November 18 (M): Project presentations. Assignment 4 due.
- November 20 (W): Project presentations.
- November 25 (M): Factor analysis.
- November 27 (W): **No class** (Thanksgiving).
- December 2 (M): EM for factor analysis.
- December 4 (W): Last lecture, if necessary.
- December 13 (F): Final projects due.