CS 303-01: Introduction to Machine Learning
Monsoon '19-20
Class Hours: M, W – 10:10 to 11:40 (AC 01, LR 103)

Instruction:	Dr. Ravi Kothari (`ravi.kothari@ashoka.edu.in`)
Office Hours:	M, W - 15:00 to 16:00 (AC-03, Room 316)
TA:	Dhruv Agarwal (`dhruv.agarwal_asp20@ashoka.edu.in`)
Office Hours:	T, Th - 10:00 to 12:00
TA:	Chanda Grover (`chanda.grover_phd19@ashoka.edu.in`)
Office Hours:	W - 16:00 to 18:00, Fr - 19:00 to 21:00

---

1 Introduction

An "intelligent system" may be characterized as one that has,

- The ability to store, represent, and retrieve knowledge;
- The ability to extract pertinent information from irrelevant details;
- The ability to learn (adapt);
- The ability to draw inferences from incomplete information.

In many cases, such systems can be constructed by "learning from examples" rather than from first principles or domain specific rules (both of which may be incompletely known or be too cumbersome to apply).

An impressive array of applications have already been realized using this paradigm of "learning from examples" and you will be able to construct your "intelligent systems" on completion of this class.

Sounds Promising? *Welcome to Introduction to Machine Learning*.

2 Prerequisites

Mathematics (Probability, Calculus), Programming in a high level language (Python, Java, . . .)

3 Required Reading

- Class notes and handouts. A copy of the notes will be on the course web page. *Do not use notes from previous years as I have revised them considerably.* The new notes are dated June 1, 2019 or later.

4 Suggested Reading

- C. M. Bishop, *Pattern Recognition and Machine Learning*, Springer, 2006.
- T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning Theory*, Springer, 2008.
- R. O. Duda, and P. E. Hart, *Pattern Classification and Scene Analysis*, John Wiley, 1973.
- I. A. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*, MIT Press, 2016.

5 Topics and Schedule

Date	Topic	Sub-Topic	Deadline/Remarks
Aug. 26	Adaptation and Empirical Models		
Aug. 28		Loss Functions	
Sept. 2		Numerical Optimization	HW # 1 given
Sept. 4		Performance of Empirical Models	
Sep. 9		Performance of Empirical Models	HW # 2 given
Sep. 11	Linear Models		
Sep. 11		Regression	
Sep. 16		Classification	HW # 3 given
Sep. 18		Perceptrons	
Sep. 18		Perceptrons	
Sep. 23		Unifying Classification and Regression	
Sep. 25		Applications and Review	

Sep. 30	Test # 1		Up to Perceptrons
Oct. 14	Non-linear Models		
Oct. 14		Multi-Layered Perceptrons (MLP)	
Oct. 16		MLP (contd.)	
Oct. 21		Decision Trees	HW # 4 given
Oct. 23		Random Forests	
Oct. 30		Bayes Classifier	
Nov. 4		Bayes Classifier (contd.); NN	HW # 5 given
Nov. 6	Project List Given		Select by Nov. 10
Nov. 6	Deep Networks		
Nov. 11		Convolutional Neural Networks	
Nov. 13	Test # 2		Up to Bayes Classifier
Nov. 18		Long Short Term Memory	
Nov. 20	Unsupervised Learning		
Nov. 20		Clustering	
Nov. 25	Applications		
Nov. 27	The Future		
Dec. 1	Final Project Reports Due		No extensions

6 Grading

Percentage in parentheses indicate the contribution to the final score used to determine grade in the class.

- **Home-Work (**30%**):** Home-work will be assigned as indicated in the previous section and is due by midnight (IST) on the day it is due. Late home-work carries a penalty of 50%/day. Home-work may involve building a system, constructing proofs, thought experiments, reading/presenting (in class)/critiquing a paper, and other such activities
- **Test 1 (**25%**):** Date given in the previous section

- **Test 2 (25%):** Date given in the previous section
- **Project (20%):** A set of candidate topics will be provided (feel free to propose and discuss any specific ideas you wish to pursue) and you will work in groups to develop and implement the project you choose. Project reports are due by midnight (IST) on the day it is due. No extensions.
- No makeup examinations unless it is truly an exceptional circumstance that is supported by documentary evidence
- Use of any unfair means or copying will result in an *F* for the course for everyone involved (the individual(s) who copied and the individual(s) who allowed the copying to occur). Please do not do it.