

CS 2490-01: Advanced Machine Learning (Spring 2019-2020)

Class Hours: M, W – 10:10 am – 11:40 pm in AC02-TR008

Instruction: Dr. Ravi Kothari ([ravi.kothari@ashoka.edu.in](mailto:ravi.kothari@ashoka.edu.in))

Office Hours: M, W – 12:00 noon – 1:00 pm in AC03-316

TA: Ms. Chanda Grover ([chanda\\_grover\\_ga@ashoka.edu.in](mailto:chanda_grover_ga@ashoka.edu.in))

Office Hours: TBD

---

We shall not cease from exploration, and the end of all our exploring will be to arrive where we started and know the place for the first time - T. S. Eliot

## 1 Introduction

We stand at the threshold of a fundamental discontinuity brought about by the confluence of *at least* two parallel developments *viz.* (i) an increased ability to sense and produce enormous amounts of data (the digital universe continues to double every 2 years), and (ii) the ability to create extremely large clusters ( 200,000 cores) driven by parallel runtimes. These developments, the first of which has been occurring across domains, have enabled *empirical* approaches to become central and indispensable to the creation of new systems, and indeed, to the creation of new knowledge.

It is not that empirical approaches are new. Indeed, a lot of the theory behind empirical approaches has been steadily developing though varyingly labeled (e.g. pattern recognition, machine learning, AI). What has changed is that the above two developments have enabled the theory to be applied to create spectacular demonstrations (e.g. self driven cars, a machine that can write a novel or create a composition). Indeed, these developments and the substantial theory that has been developed emboldens us to ponder *if we can create a mind!*

This *advanced* course goes deep into machine learning - relying on *mathematical rigor*, *programming and system building skills*, *self-study*, and *team learning* to take a step towards principled understanding and application of machine learning. This is **not** a course for everyone. Please examine the italicized words in this paragraph and the prerequisites listed to assess if this course is for you.

If you have continued reading, welcome to [Advanced Machine Learning](#). At the end of the course, you should have a good understanding of advanced neural architectures and learning algorithms, you should be able to understand factors that affect the performance, and you should be able to systematically apply advanced approaches to solving real world problems.

## 2 Prerequisites

CS 303 (Introduction to Machine Learning), Strong programming skills, Mathematics (Probability, Calculus)

## 3 Required Reading

- Class notes and handouts.

## 4 Suggested Reading

- I. A. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*, MIT Press, 2016.
- K. P. Murphy, *Machine Learning: A Probabilistic Perspective*, MIT Press, 2012.
- V. M. Vapnik, *The Nature of Statistical Learning Theory*, Springer, 1998.
- T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning Theory*, Springer, 2008.
- C. M. Bishop, *Pattern Recognition and Machine Learning*, Springer Verlag, 2006.
- R. O. Duda, and P. E. Hart, *Pattern Classification and Scene Analysis*, John Wiley, 1973.
- R. Kurzweil, *How to Create a Mind: The Secret of Human Thought Revealed*, Viking Penguin, 2012.

- T. M. Mitchell, *Machine Learning*, McGraw Hill, 1997.
- (Stretch Topic) R. S. Sutton, and A. G. Barto, *Reinforcement Learning: An Introduction*, MIT Press, 2017.

## 5 Topics and Schedule

This is the approximate schedule of classes. My intent is to follow this schedule modulo any unavoidable adjustments.

Date	Topic	Deadline/Remarks
Jan. 20	Introduction; Approximation of functions	
Jan. 22	...; Approximation by FFNNs	
Jan. 27	...; Prediction error	HW1 Given; Due Feb. 4
Jan. 29	...; Motivation for Deep Networks	
Feb. 3	CNN	
Feb. 5	CNN	HW2 Given; Due Feb. 13
Feb. 10	Bayesian Neural Networks	
Feb. 12	Bayesian Neural Networks	HW3 Given; Due Feb. 20
Feb. 17	Approximate Inference	
Feb. 19	Variational Auto-encoders	
Feb. 24	Variational Auto-encoders	
Feb. 26	Test 1	Upto Approximate Inference
Mar. 2	Generative Models	
Mar. 4	Generative Models	HW4 Given; Due March 20
Mar. 16	LSTM	Projects Given – Select by Mar. 20
Mar. 18	LSTM	
Mar. 23	Test 2	Upto LSTM
Mar. 25	Caption Generation	
Mar. 30	Transformers/BERT	
Apr. 1	Transformers/BERT	
Apr. 6	Evolutionary Algorithms	
Apr. 8	Evolutionary Algorithms	
Apr. 13	Learning rules	
Apr. 15	FAT in ML	

Apr. 20	Group discussion on how to create a mind	Read prior to this date
Apr. 22	Final thoughts	
May. 1		Final project reports due

## 6 Grading

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

- **Home-Work (35%):** 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 (15%):** Date given in the previous section
- **Test 2 (15%):** Date given in the previous section
- **Project (35%):** 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 individually (or in groups - I will decide later based on total enrollment and other factors) to develop and implement the project you choose. Project reports are due by midnight (IST) on the day it is due. Late submission carries a penalty of 50%/day.
- 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. Please do not do it**