

CS5590 Foundations of Machine Learning

Instructor: J. SakethaNath

October 3, 2019

About the course

This is the basic foundations course on machine learning for CSE students. The course assumes no prior exposure to machine learning. However, it assumes the student has good understanding of basic engineering math subjects like probability & statistics, linear algebra etc.

The focus of this course is on understanding basic principles and concepts, rather than on giving a hands-on overview of state-of-the-art. Students interested in the later are encouraged to credit the “Applied ML” course.

The syllabus of this course is:

1. Overview of machine learning, challenges, terminology and definitions, Supervised Learning (6 hrs)
2. Basic probabilistic and non-probabilistic models:
 - (a) Loss composed Linear Models and parameter estimation using ERM (4 hrs)
 - (b) Generative models and parameter estimation using method of moments, maximum likelihood (3hrs)
 - (c) Discriminative models and maximum conditional likelihood estimation (2 hrs)
3. Generalized models:
 - (a) Exponential models, and kernel methods (2 hrs)
4. Model selection (1 hr)
5. Non-parametric modeling (3 hrs)
 - (a) k-NN, KDE
 - (b) variants for generative, discriminative classification and regression
6. Unsupervised Learning (3 hrs)
 - (a) 1-class SVM, SVC

- (b) Mixture models, EM algorithm, k-means
- 7. Representation Learning (6 hrs)
 - (a) PCA, probabilistic PCA, supervised PCA,
 - (b) neural networks and SGD
- 8. Online Learning (5 hrs)
- 9. Reinforcement Learning (4 hrs)
- 10. Other popular paradigms: Boosting, Bagging, trees (3 hrs)
- 11. Learning with Sequential Data (3 hrs)

References

The textbooks for this course are:

1. For probabilistic models [1].
2. For kernel methods [2].
3. In general, part2 and part 3 of [3]¹.

Evaluation

Date	Duration	Percentage
03-Sep-2019 (Tue)	90 min.	10% (Closed book/laptop; No cheat-sheets)
15-Oct-2019 (Tue)	90 min.	20% (Closed book/laptop; No cheat-sheets)
Final Exam (27-Nov to 2-Dec)	180 min.	70% (Closed book/laptop; No cheat-sheets)

Apart from exams, there will be regular assignments given. The simulation (hands-on) exercises will be given with the help of Dr. Anoop (Microsoft). The other theoretical assignment problems will serve as practice for exams. Though none of these assignments will be evaluated, 30% of exam questions will be directly from theoretical assignment problems. Hence it is important to attempt and solve them. The hands-on assignments will give you an expertise in themselves as they will be based on popular libraries like PyTorch.

¹This is freely available at <https://www.cs.huji.ac.il/~shais/UnderstandingMachineLearning/understanding-machine-learning-theory-algorithms.pdf>.

Contact

The course page is at <http://www.iith.ac.in/saketha/teaching/cs5590.html>. All official communication and technical interactions outside lecture hours will be through Google classroom. The code for regular students is “6ioh74b” and that for MDS batch is “tryl3i3”. The instructor can be contacted via email: saketha at iith. The typical working hours of the instructor are 8:30am to 12:30pm and 2:15pm to 5:15pm. His office is C-519. You are welcome to drop-by anytime for clarifying doubts during these working hours. Technical interaction (outside and in lecture hours) is highly encouraged.

References

- [1] Kevin P. Murphy. *Machine Learning: A Probabilistic Perspective*. The MIT Press, 2012.
- [2] Bernhard Scholkopf and Alexander J. Smola. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. MIT Press, Cambridge, MA, USA, 2001.
- [3] Shai Shalev-Shwartz and Shai Ben-David. *Understanding Machine Learning: From Theory to Algorithms*. 2014.