Boston University, College of Engineering
Department of Electrical and Computer Engineering

ENG EC719 (Ishwar) – Spring 2023

STATISTICAL LEARNING THEORY

Instructor: Prof. Prakash Ishwar (e-mail: pi@bu.edu)
Lectures: Mon + Wed, 4:30-6:15 pm (HAR 212)
Instructor's office hours: Fri 5:00-6:30 pm Photonics Building (tentatively in Room PHO 440)

Description:
Classical and contemporary theories of machine learning. Topics/emphasis may change based on instructor preference in different years. A project involving computer implementation of a learning or inference algorithm accompanied by or in support of theoretical analysis is required.
Prerequisites:
EC414 or EC503 or equivalent experience with machine learning. Good foundation in probability, linear algebra, multivariate calculus, and programming. Interest and appreciation of analytical work (e.g., theorems and proofs).

Syllabus:

- Statistical Decision Theory: optimum predictors and their properties.
- Frequentist and Bayesian learning with generative and discriminative probabilistic models of parametric or nonparametric kinds.
- Asymptotic consistency and limitations of ML learners.
- Universally consistent learners: generalized linear and k nearest-neighbor predictors.
- No free lunch theorem: non-existence of universal rate of convergence, error-decomposition, bias-complexity tradeoff.
- Empirical Risk Minimization: VC dimension, Rademacher complexity, sample complexity bounds, PAC learning.
- Structural Risk Minimization, nonuniform learnability, model-selection.
- Convex learning problems.
- Regularization and stability.
- Margin bounds; Analysis of some popular learning algorithms.
- *Selected Topics (as time permits or via projects)*: contrastive representation learning, domain adaptation & generalization, new perspectives on over- and under-fitting, generalization, and the role of over-paramerization.

Grading:

30%	Homeworks	~ 4 sets of predominantly analytical exercises.
65%	Project	Project involving reading, analyzing, and presenting a recent paper related to statistical learning theory or its applications (more information will follow later in the semester). Project Presentations: May 4 and May 5 from 4:00-6:00 pm (tentatively in HAR 212). Project Final report due: Sunday, May 7.
5%	Class-participation	For constructive and proactive engagement during lectures, office hours, and on Piazza.

Outcomes: As an outcome of successfully completing this course, students will:

- be able to understand and use classical and contemporary theoretical results and algorithms in machine learning;
- be able to read, understand, analyze, and implement classical and contemporary research articles related to the theoretical foundations of machine learning;
- be able to understand the pros and cons of different machine learning algorithms and their practical implications;
- be able to communicate theoretical results in written or oral forms.

Web site: `http://learn.bu.edu` will contain lecture slides, notes, links, and other useful information related to the course.

References: This course has no textbook. You will need to rely on lectures, office hours, and course materials that will be uploaded regularly to the course web site. Below is a list of reference books which we will consult.

- S. Shalev-Shwartz and S. Ben-David, *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press, 2014.
  `https://www.cs.huji.ac.il/~shais/UnderstandingMachineLearning/`
- M. Mohri, A. Rostamizadeh, and A. Talwalkar, *Foundations of Machine Learning (2nd ed.)*. The MIT Press, 2018. `https://cs.nyu.edu/~mohri/mlbook/`
- L. Devroye, L. Gyorfi, and G. Lugosi, *A Probabilistic Theory of Pattern Recognition*. Springer, 1996. `https://www.szit.bme.hu/~gyorfi/pbook.pdf`
- K. P. Murphy, *Machine Learning: A Probabilistic Perspective*. The MIT Press, 2012.
- T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning: data mining, inference, & prediction*. Springer, 2nd edition - 2009.
  `https://web.stanford.edu/~hastie/Papers/ESLII.pdf`
- C.M. Bishop, *Pattern Recognition and Machine Learning*. Springer, 2006.
- R.O. Duda, P.E. Hart, and D.G. Stork, *Pattern Classification*. Wiley-Interscience, 2nd edition - 2000.

Copyright notice: © The copyright on all ENG EC 719 course materials, including lecture slides/notes, homeworks, and solutions is asserted by Professor Prakash Ishwar and the specific individuals identified by him. All rights are reserved. Other than for personal use by *registered* students, the ENG EC 719 course materials or any portion thereof may not be reproduced or used in any manner whatsoever (including posting on a public forum like github) without the express written permission of Professor Prakash Ishwar. Acting otherwise would go against the ethical code of conduct expected of students at Boston University and from a legal standpoint it may even constitute a violation of copyright law.

Academic integrity, plagiarism: You may discuss homework problems among yourselves and with the instructor (during office hours), but any significant collaboration should be acknowledged in your solution. You must always create your own solution without copying it from your collaborator(s). In project reports, you may *paraphrase* relevant ideas from references, but not quote sentences verbatim from them.

The BU student handbook defines academic misconduct as follows: *"Academic misconduct occurs when a student intentionally misrepresents his or her academic accomplishments or impedes other students' chances of being judged fairly for their academic work. Knowingly allowing others to represent your work as theirs is as serious an offense as submitting another's work as your own."* Please see the student handbook for procedures that will follow should academic misconduct be discovered.

Inclusion: I consider the classroom to be a place of learning where all individuals are expected to contribute to provide a respectful, welcoming and inclusive environment for every member of the class irrespective of how they identify themselves.

Disability accommodations: If you are a student with a documented disability or believe you might have a disability that requires accommodations, requests for accommodations must be made in a timely fashion to Disability & Access Services, 25 Buick St, Suite 300, Boston, MA 02215; 617-353-3658 (Voice/TTY). Students seeking academic accommodations must submit appropriate medical documentation and comply with the established policies and procedures `http://www.bu.edu/disability/accommodations/`

BU health-safety policies: BU strongly encourages use of high-quality masks to reduce the risk of disease transmission in crowded settings or for individuals who are at increased risk of severe illness. Students should be compliant with all BU health-safety policies. For BU health-safety policies, please visit the university website.

Illness/Medical issues: will be resolved by case-specific discussions with the instructor. Student must inform instructor by email as soon as possible and be prepared to provide suitable documentation. May also require contacting BU's disability access services.