TO: The Faculty of the College of Engineering

FROM: Elmore Family School of Electrical and Computer Engineering

RE: New Graduate Course, ECE 50024 Machine Learning

The faculty of the School of Electrical and Computer Engineering has approved the following new course. This action is now submitted to the Engineering Faculty with a recommendation for approval.

ECE 50024 Machine Learning
Sem. 2, Lecture 3, Cr. 3.
Prerequisite by Topic: ECE20875 Python programming, and ECE302 Probabilistic Methods in ECE, or equivalent courses

Description: An introductory course to machine learning, with a focus on supervised learning using linear models. The course will have four parts: (1) mathematical background on linear algebra, probability, and optimization. (2) classification methods including Bayesian decision, linear regression, logistic, regression,
and support vector machine. (3) robustness of classifier and adversarial examples. (4) learning theory on the feasibility of learning, VC dimension, complexity analysis, bias-variance analysis. Suitable for senior undergraduates and graduates with a background in probability, linear algebra, and programming.

Reason: This course focuses on the fundamental principles of machine learning. It is a dual level course which will prepare student to move to more specialized topic related to machine learning. The course is suitable for senior undergraduates and graduate students with a background in probability, linear algebra, and programming.

Milind Kulkarni,
Associate Head for Teaching and Learning
Elmore Family School of Electrical and Computer Engineering

ECE 595ML / STAT 598 Machine Learning

Course Information
Spring 2021

Lecture Hours: TuTh 4:30-5:45pm, Zoom

Instructor: Professor Stanley Chan
Room: MSEE 338
Email: ece595chan@gmail.com
Office Hours: By email appointment.

Teaching Assistants:
Chengxi Li: ece595chan@gmail.com

Administrative Assistants:
Camille Hamelman, and Cheryl Leuck
Email: ece595chan@gmail.com

Course Website: `https://engineering.purdue.edu/ChanGroup/ECE595/`

Course Objectives:

The objective of this course is that by the end of the semester, students will be able to

- Apply basic linear algebra, probability, and optimization tools to solve machine learning problems.
- Understand the principles of supervised learning methodologies, and can comment on their advantages and limitations.
- Explain the trade-offs in model complexity, sample complexity, bias, variance, and generalization error in the learning theory.
- Implement, debug, and execute basic machine learning algorithms on computers.

Catalog Description: Machine Learning is a dual-level (500-level) course customized for science and engineering students who are seeking to develop a solid mathematical foundation of the subject. The course focuses on the fundamental principles of machine learning instead of a scattered set of algorithmic tools. Students completing the course will be able to formulate the practical machine learning problems into mathematical frameworks, implementing algorithms to execute the statistical inference tasks, and analyzing the performance and limitations of the algorithms. The course has four parts: (1) linear regression, which covers regression models, outliers, ridge regularization, LASSO regularization, convex optimization, gradient descent algorithms, and stochastic algorithms; (2) classification, which covers separability, Bayesian classifiers, ROC curves, precision-recall curves, logistic regression, kernel methods; (3) learning theory, which covers probability inequality, the probably approximately accurate framework, generalization bound, model complexity, sample complexity, VC dimension, bias, variance, overfitting, and validation; (4) advanced topics of the state-of-the-arts, for example deep neural networks, generative models, and adversarial robustness. The course emphasizes the co-development of theory and programming. Students will have hands-on experience implementing machine learning algorithms in Python.

Required textbook:

- *Introduction to Probability for Data Science*, by Stanley Chan, Michigan Publishing, 2021.
- *Learning from Data*, by Abu-Mostafa, Magdon-Ismail and Lin, AMLBook, 2012.

Recommended textbook:

- *Pattern Classification*, by Duda, Hart and Stork, Wiley-Interscience; 2 edition, 2000.
- *Elements of Statistical Learning*, by Hastie, Tibshirani and Friedman, Springer, 2 edition, 2009.
- *Pattern Recognition and Machine Learning*, by Bishop, Springer, 2006.

Pre-requisites: ECE20875 Python programming, and ECE302 Probabilistic Methods in ECE.

In addition to these minimum requirements, the following courses are highly recommended to students who are planning to take ECE595ML.

- **Linear Algebra** (as in the materials covered by G. Strang's Linear Algebra Textbook)

  A good course at Purdue is MA 511 Linear Algebra
- **Optimization** (as in chapter 1 - chapter 4 of S. Boyd's Convex Optimization)

  A good course at Purdue is ECE 647 (Week 1 - Week 8)
- **Probability** (as in the materials covered by D. Bertseka's Intro to Probability Textbook)

  A good course at Purdue is ECE 302 Probability. Graduate probability such as ECE 600 is recommended but not required.

To help you determine if you have adequate pre-requisites, we encourage you to try homework 0 posted in the homework section. If the problems are significantly beyond your comfort level, we suggest considering taking ECE 595 at a later time.

Grades: All students will be graded by the following rubric. Graduate students and undergraduate students will be graded on two different curves.

- Homework (10%). Homework are given approximately biweekly. Please submit your homework through gradescope. Late homework will not be accepted. You are encouraged to work in smalll groups, but you have to write / type your own solution.
- We highly encourage you to type your solution using the LaTeX template provided in the course website, although we accept hand-written solutions. All programming answers should be typed.
- Quiz (40%). There will be 6 quizzes throughout the semester. The quizzes shall be taken on gradescope. Each quiz will be 30 minutes long.
- Project (50%). Please visit the course website for project instructions.

Academic Dishonesty:

You are in college, not high school. We respect you as adults, and we expect you behave as adults. Therefore, we ask you to be honest and ethical in the course. In that respect, any action that might give a student unfair advantage on homework or exams will be considered dishonest. Examples include, but are not limited to:

- Sharing information during exam;
- Using forbidden material or device during exam;
- Viewing and/or working on an exam before or after the official time allowed;
- Requesting a re-grade of work that has been altered;
- Submitting work that is not your own. (You can discuss problems with your classmates. But you must write your own solution.)

All cases of academic dishonesty will be reported to the Office of Student Rights and Responsibilities, and will result in punishment. Possible punishments include, but are not limited to, a score of zero on work related to the cheating incident, a failing grade for the course, and, in severe cases, expulsion from the university.

Copyright of Course Material:
All ECE 595 course material, including lecture, homework, project, solutions and exams are protected by copyright law. Without Prof Chan's permission, you are not allowed to distribute through any media including online sources. Below is an excerpt from `http://www.purdue.edu/studentregulations/student_conduct/misc.html`

... Students enrolled in, and authorized visitors to, Purdue University courses are permitted to take notes, which they may use for individual/group study or for other non-commercial purposes reasonably arising from enrollment in the course or the University generally. Notes taken in class are, however, generally considered to be "derivative works" of the instructor's presentations and materials, and they are thus subject to the instructor's copyright in such presentations and materials. No individual is permitted to sell or otherwise barter notes, either to other students or to any commercial concern, for a course without the express written permission of the course instructor...

Emergency Procedure:
Purdue University is a very safe campus and there is a low probability that a serious incident will occur here at Purdue. However, just as we receive a "safety briefing" each time we get on an aircraft, we want to emphasize our emergency procedures for evacuation and shelter in place incidents. Our preparedness will be critical *if* an unexpected event occurs.

Purdue prepares for natural disasters or human-caused incidents with the ultimate goal of maintaining a safe and secure campus, but in the end, emergency preparedness is your personal responsibility. Let's quickly review the following procedures:

- To report an emergency, call 911. To obtain updates regarding an ongoing emergency, sign up for Purdue Alert text messages, view `www.purdue.edu/ea`.
- There are nearly 300 Emergency Telephones outdoors across campus and in parking garages that connect directly to the PUPD. If you feel threatened or need help, push the button and you will be connected immediately.

- If we hear a fire alarm during class we will immediately suspend class, evacuate the building, and proceed outdoors. Do not use the elevator.
- If we are notified during class of a Shelter in Place requirement for a tornado warning, we will suspend class and shelter in [the basement].
- If we are notified during class of a Shelter in Place requirement for a hazardous materials release, or a civil disturbance, including a shooting or other use of weapons, we will suspend class and shelter in the classroom, shutting the door and turning off the lights.
- Please review the Emergency Preparedness website for additional information. `http://www.purdue.edu/ehps/emergency_preparedness/index.html`

In the event of a major campus emergency, course requirements, deadlines and grading percentages are subject to changes that may be necessitated by a revised semester calendar or other circumstances. In such an event, information will be provided through the course website and through emails.

Schedule:

	Part 1: Linear regression	
Jan 19, 2021.	Lecture 0.	Course overview. Mathematics review.
Jan 21, 2021.	Lecture 1.	Linear regression
Jan 26, 2021.	Lecture 2.	Outliers in linear regression
Jan 28, 2021.	Lecture 3.	Ridge and LASSO Regularization
Feb 2, 2021.	Lecture 4.	Optimization: Concepts
Feb 4, 2021.	Lecture 5.	Optimization: Algorithms
	Part 2: Classification	
Feb 9, 2021.	Lecture 6.	Linearly Separable
Feb 11, 2021.	Lecture 7.	Bayesian Classifier 1
Feb 16, 2021.	Lecture 8.	Bayesian Classifier 2
Feb 18, 2021.	Lecture 9.	Classification Error / ROC Curve
Feb 23, 2021.	Lecture 10.	Estimating parameters
Feb 25, 2021.	Lecture 11.	Logistic regression 1
Mar 2, 2021.	Lecture 12.	Logistic regression 2
Mar 4, 2021.	Lecture 13.	Kernel trick
Mar 9, 2021.	Lecture 14.	Kernel trick + Probability Inequality
	Part 3: Learning Theory	
Mar 11, 2021.	Lecture 15.	Probability Inequality
Mar 16, 2021.	Lecture 16.	Is learning feasible?
Mar 18, 2021.	No class.	University reading day
Mar 23, 2021.	Lecture 17.	Probably-Approximately Correct
Mar 25, 2021.	Lecture 18.	Generalization bound
Mar 30, 2021.	Lecture 19.	Growth function
Apr 1, 2021.	Lecture 20.	Growth function examples
Apr 6, 2021.	Lecture 21.	VC dimension
Apr 8, 2021.	Lecture 22.	Bias and Variance
Apr 13, 2021.	No class	University reading day
Apr 15, 2021.	Lecture 23.	Overfitting
	Part 4: Advanced Topics	
Apr 20, 2021.	Lecture 24.	Into to neural network
Apr 22, 2021.	Lecture 25.	Convolutional structures and back propagation
Apr 27, 2021.	Lecture 25.	Generative adversarial networks
Apr 29, 2021.	Lecture 26.	Conclusion

Addendum: EFD 14-22

Grade Mapping:

Homework: 10%

Quiz: 40%

Project: 50%

Mapping:

A: > 85

B: 75-84

C: 60-74

D: 50-59

F: < 50

Previous Course Offerings:

ECE 59500 Machine Learning, Spring 21 (80), Spring 20 (105), Spring 19 (180)