USC Viterbi School of Engineering

EE 559
Machine Learning I: Supervised Methods (including Mathematical Pattern Recognition)
Units: 4

For Spring 2021: to get the full 4 units of credit, please also register for 1 unit of directed research (EE 590):

1. Through myviterbi.usc.edu, Directed Research, Ming Hsieh Dept. of ECE, choose EE 590 (1 unit);
2. For supervisor/instructor: Keith Jenkins; for topic/project description: "EE 559 supplement";
3. You will then be given D clearance for EE 590 (1 unit)
4. At the end of the semester, you will receive a letter grade for EE 559 and CR/NC for EE 590.

=> If you don't register for EE 590, you will only get 3 units of credit.

This syllabus is a draft (v0.8). There will be some minor changes to accommodate the shortened semester and online instruction.

Lecture: MW 3:30 – 5:20 PM Pacific Time
Discussion: Thu 5:30 – 6:20 PM Pacific Time

Location: Online, and DEN@Viterbi
* Partially in-person instruction will be included only when and if it becomes safe and reasonable to do so.

Course URL: https://courses.uscden.net/d2l/home

Instructor: B. Keith Jenkins

Office: EEB 404A (only when partial in-person instruction is included)
Office Hours: Tu Th 4:00 – 5:00 PM Pacific Time (tentative)
Contact Info:
Email: jenkins@sipi.usc.edu [Please include "EE 559" in the subject line]
Online: piazza.com

Teaching Assistants: TBA

IT Help: For help with coding machine-learning algorithms, consult piazza, online forums, the TAs, or the instructor; for help with other Python coding, working with datasets, or using library routines in Python, consult piazza, online help and documentation, reference resources given below, or the TAs; for help with USC-supplied software or on-campus networking, consult USC ITS at https://itservices.usc.edu/contact/ .

Course Description (Catalogue)

Distribution-free and probabilistic methods for supervised classification and regression; learning algorithms; optimization techniques; feature-space transformations; parametric and nonparametric methods; Bayes decision theory; artificial neural networks.

Course Description (Expanded)

Concepts and algorithms for pattern recognition and regression using machine learning are covered in depth. The course will stress an understanding of different supervised-learning algorithms at both theoretical and practical levels, as well as their advantages and disadvantages. Underlying fundamentals are emphasized, including theory and origins of learning algorithms and criterion functions. The goal is to give the student an understanding of some fundamental approaches to machine learning, to enable further study and growth on their own. The student's work will include mathematical analysis, analytical understanding, and writing and running code that learns from data. A moderately sized project in the second half of the semester will involve developing and optimizing one or more machine learning systems to perform well on real-world datasets. This course is intended for graduate students in Electrical and Computer Engineering or related fields, who wish to gain an understanding of, and some experience with, machine learning approaches, tools, and techniques. Only supervised-learning methods are covered in this course.

Learning Objectives

After successfully completing this course, the student will:

- Have a perspective of different approaches to supervised machine learning for pattern recognition and regression
- Understand the underlying math of a variety of supervised-learning methods
- Be able to use statistical and non-statistical techniques to solve machine learning problems
- Be able to code and run algorithms for learning from data
- Be able to optimize machine learning algorithms and systems, and assess their overall performance
- Know how to conceive and develop new techniques for machine learning where needed

Prerequisite(s):

Co-Requisite(s): EE 503 and EE 510

Concurrent Enrollment:

Recommended Preparation: knowledge of Python at the level of EE 541 (A Computational Introduction to Deep Learning); knowledge of multivariate calculus.

Course Notes

The Desire2Learn (D2L) system will host the course website. On the website will be posted video recordings of the lectures and discussion sessions, lecture and discussion-session notes, additional instructor-provided notes and handouts, and assignments. Your graded assignments and scores will also be visible to you on the website.

Technological Proficiency and Hardware/Software Required

Python will be used throughout this class for homework assignments and the class project. All students will be responsible for installing and maintaining their own Python distribution. For students that aren't fully up to speed in Python, resources (books, online links, and discussion sessions) will be provided to help you learn during the first few weeks of class. For this, knowledge of MATLAB will be assumed and used as a springboard to Python.

Required Readings and Supplementary Materials

Required text

- C. M. Bishop, *Pattern Recognition and Machine Learning* (Springer, 2006). ISBN-13: 978-0387-31073-2. Available from USC bookstore, Amazon.com, and Springer.com.

Supplementary texts

- R. O. Duda, P. E. Hart, and D. G. Stork, *Pattern Classification*, Second Edition (Wiley-Interscience, John Wiley and Sons, Inc., New York, 2001)
- T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, Second Edition (Springer, 2009)
- Simon Haykin, *Neural Networks and Learning Machines,* 3rd Edition, (Pearson, 2009)
- Ethem Alpaydin, *Introduction to Machine Learning, Fourth Edition* (MIT Press, 2020)

Supplementary resources for Python

- Hans Fanghor, *Introduction to Python for Computational Science and Engineering* (2016) [PCSE], available for free download at:
  https://github.com/fangohr/introduction-to-python-for-computational-science-and-engineering

- Fabio Nelli, *Python Data Analytics*, (APress, 2015) [PDA], available for download from USC Library:
  http://usc.summon.serialssolutions.com/search?q=Fabio%20Nelli%2C%20Python%20Data%20Analytics#!/search?ho=t&l=en&q=Fabio%20Nelli,%20Python%20Data%20Analytics

- The Python 3 Tutorial [PT]:
  https://docs.python.org/3.8/tutorial/index.html

- EU Python 3 Tutorial [EUP]: Good for Chapters on object oriented programming, class vs. instance attributes, and inheritance
  http://www.python-course.eu/python3_course.php

Description and Assessment of Assignments

1. Homework assignments

Homework assignments will, on average, consist of approximately 50% computer problems and 50% analytical problems. Homework assignments will be posted on the course website on Friday, and due the following Friday. Each homework assignment will be worth 10 points. Starting with Homework 3, a portion of the homework score (1-2 points, depending on the amount and complexity of coding that is needed) for each homework assignment that includes coding will be based on quality of Python code that the student has written. (Guidelines for writing quality code will be given in Discission for Week 3 as well as a handout and links posted on the course website.)

Python 3 is required for all homework computer problems, with the exception of the first 2 homework assignments, for which you are also allowed to use MATLAB. Some of the computer problems will require coding machine learning algorithms yourself (without using libraries such as pandas, sk-learn, etc.); others will allow use of these and other libraries. Each assignment will state the requirements and guidelines for its computer problems.

Python 3 is open source and available for free download (*e.g.*, as a distribution like www.anaconda.com).

2. Course project

Programming languages: For the course project, you may use Python 3 and/or C/C++. If you want to use any other language for portions of your project, please check with the instructor or TA first.

Project content: You will be given a choice of 2-3 real-world datasets and associated nontrivial problems to solve. Your goal is to design and demonstrate a good-performing machine learning system, to compare some different approaches, and to understand and explain the results you obtain. To do so, you will use your choice of preprocessing, feature extraction, dimensionality adjustment, and classifiers/regressors, etc., to solve the problem; as well as cross validation, training and test datasets as appropriate. You will be given some guidelines and suggestions, as well as some requirements of items your project must include.

Deliverables: On or before the due date specified in the Project Assignment, you will turn in your project by uploading 3 files to the D2L website: (1) a written final report (as a pdf file); (2) a single, computer-readable pdf file with all of your code; and (3) a zip file that contains all your code files (e.g., .py files) in a form that can be run to verify your results. The written final report must be typewritten and include tables and graphs to show your final results as well as sample or summary intermediate results. The suggested maximum length is 15 pages. Detailed instructions for the written final report will be posted along with a template in Word and LaTeX.

Sample project: You are given a Hand Postures dataset that consists of frames taken from videos of different hand gestures. The acquired data consists of locations of markers on a glove worn by each subject. The dataset is challenging because: (1) the markers are not labelled, so it is not known which data entry corresponds to which marker; (2) typically some of the markers are occluded in any given frame (image); and (3) the number of features varies from one data point (one frame) to another. Your goal is to design and demonstrate a pattern recognition system that can reliably recognize the hand posture, given input data from a separate test set. Your project is required to include a trial of at least 4 classifiers/learning algorithms (some of which are specified, some of which are up to you); feature dimensionality adjustment; any necessary preprocessing; proper dataset usage including cross validation and model selection; estimate of your final system's performance on unseen data; analysis, explanation, and demonstration that you understand intermediate and final results.

Participants: Each project can be done by a team of 2 students, or by an individual student. Guidelines will be provided in the Project Assignment as to as to which datasets and problems constitute a reasonable topic and workload for a team project and for an individual project.

Timeline (tentative):

- Week 10: Project Assignment, and Dataset Descriptions, posted on D2L
- Week 15: Final Report and code files due.

There are no progress reports or proposals required. However, the instructor and TAs are available for any discussions or questions that you would find helpful; and you are encouraged to seek advice or direction as you deem appropriate for your project.

Your course project will be graded by the following criteria:

- Technical approach (20%);
- Analysis, understanding, and interpretation (20%);
- Performance of your final system (15%);
- Workload (15%);
- Difficulty of the problem (10%);
- Quality of final report write-up (10%),
- Quality of written code (10%).

Grading Breakdown (tentative)

Assessment Tool (assignments)	% of Grade
Homework	30
Course project	25
Midterm exam	20
Final exam	25
TOTAL	100

Extra credit homework problems (throughout the semester) will be tallied separately from the above scores, and will determine students' grades for borderline cases (*i.e.*, students near a border between two letter grades).

Assignment Submission Policy

Homework assignments are due on Friday, one week after they are assigned. Submit your homework by uploading one pdf file of your solution, and one computer-readable pdf file of all your code (for assignments with computer problems), to the Assignment Dropbox on the D2L website. A computer-readable pdf can be generated by converting an editable text document to pdf; scanned documents and screen-shots are not computer readable.

Due date and time will be stated on each homework assignment. The late submission policy will be posted on the D2L web site. Late penalties will be uniformly applied to everyone. Your two lowest homework scores will be dropped to accommodate special circumstances, without the need for requesting an extension.

Grading Timeline

Graded assignments including comments will be available on the D2L website as soon as grading is completed and verified (typically 1-2 weeks after the due date).

Additional Policies

Policy on Collaboration and Individual Work in this Class

Collaboration on techniques for solving homework assignments and computer problems is allowed, and can be helpful; however, each student is expected to work out, code, and write up his or her own solution. Use of other solutions to homework assignments or computer problems, from any source including other students, before the assignment is turned in, is not permitted.

For class projects, general collaboration to resolve issues, or to clarify technical material, is allowed. Use of internet as well as journal and conference literature is encouraged. However, each student (or team) does their own work and writes up their own report. The author(s) of the report are presenting themselves as having done the work described in the report. Any reported work, explanations, information, or code that is obtained from others must be cited as such; instructions for doing this will be given with the project assignment. Including such work in the report or code without citing it amounts to plagiarism.

Please also see below for additional policies that apply to all USC classes.

Course Schedule: A Weekly Breakdown (to be revised to accommodate shorter-than-usual semester in Spring 2021)

Readings refer to section numbers in Bishop.

Abbreviations:

- ANN: artificial neural networks
- ML: machine learning
- HW: Homework

	Lecture Topics	Readings/ Preparation	Deliverables	Discussion session topics
Week 1	Introduction, basic concepts, paradigm for pattern recognition and regression, fundamental assumptions	1.0, 1.1		Course logistics (HW submission, use of D2L); ML design cycle
Week 2	**Functions and methods:** Discriminant functions for classification (2-class, multiclass); multiclass classification methods. Regression functions, regularization.	4.1.0-4.1.2	HW1 assigned.	Linear algebra summary; matrix and vector derivatives
Week 3	**Preliminaries and problem set-up for ML:** Computational complexity (big-oh, big-omega, big-theta); convex and concave functions; ML preliminaries (notation, weight space, paradigm for learning), ML problem statement and approach, criterion functions, approaches to optimization	Instructor-provided notes; handout on quality of Python code; web links for quality Python coding.	HW1 due. HW2 assigned.	Writing quality Python code (clarity, conciseness, naming conventions, ease of maintenance and debugging, refactoring, runtime)
Week 4	**Learning and optimization 1:** Optimization by gradient descent, gradient descent algorithms (batch, sequential, stochastic, mini-batch), perceptron criterion and learning, convergence proof	5.2.4, 4.1.7, instructor-provided notes	HW2 due. HW3 assigned.	Perceptron example; multiclass perceptron algorithm and its convergence
Week 5	**Learning and optimization 2:** Mean square criterion for regression, optimization by algebraic solution, linear regression, ridge regression, iterative optimization. Mean square criterion for classification; learning algorithms. Extension to nonlinear techniques.	3.0, 3.1, 4.1.3	HW3 due. HW4 assigned.	Introduction to Python's sk-learn; data normalization
Week 6	**Complexity in supervised learning 1:** Degrees of freedom and constraints. **Learning and optimization 3:** Lagrangian techniques for constrained optimization. Dual representations; kernels and kernel methods; constructing kernels; radial basis function networks	6.0-6.2, 6.3.0, App. E	HW4 due. HW5 assigned.	Examples of overfitting and underfitting caused by complexity imbalance; examples of kernels for applications to text processing, recommender systems.
Week 7	**Learning and optimization 4:** Support vector machines (classification): margin, criterion,	7.1 excluding 7.1.4	HW5 due.	Sample problems in review for midterm exam

	primal and dual formulations, linearly and nonlinearly separable cases. Kernels for support vector machines. **Complexity in supervised learning 2:** complexity (capacity) of a separating boundary, introduction to VC dimension.			
Week 8	Support vector regression. Review. Midterm exam	7.1.4	HW6 assigned.	Tips for using LIBSVM; preprocessing for data imbalance.
Week 9	**Feedforward ANN 1:** Single neuron unit; activation functions; 1 layer of neuron units; multiple layers of neuron units. Interpretations, capabilities, and limitations. Learning algorithms including perceptron, Widrow-Hoff, outer product learning, error backpropagation.	5.0-5.3	HW6 due. HW7 assigned.	Examples of ANN.
Week 10	**Feedforward ANN 2:** Regularization, radial basis function networks. Complexity, constraints, example: convolutional ANNs. **Data preprocessing techniques:** missing data, categorical features.	5.5, instructor-provided references	HW7 due. Project Assignment posted.	Data preprocessing: examples and functions.
Week 11	**Feature optimization by transformation:** Principal components analysis, Fisher's linear discriminant and generalizations, multiple discriminant analysis. **Dataset methodology:** Validation, cross validation, model selection, error estimation	12.1, 1.3	HW8 assigned.	Feature selection techniques. Mean and standard deviation of validation-error measures; confusion matrix.
Week 12	Random vectors and linear transformations of them. **Bayes decision theory (classification):** Minimum-error and minimum-risk criteria and solution, Gaussian case, Mahalanobis distance; continuous and discrete features.	2.0-2.3.3, instructor provided notes	HW8 due. HW9 assigned.	Examples of: transformations of random vectors, Mahalanobis distance, Bayes classification.
Week 13	**ML by density estimation (classification):** General approach, kernel methods, nearest-neighbor methods, convergence; generative and discriminative models; degrees of freedom, assumptions, curse of dimensionality.	2.5	HW9 due.	Additional performance metrics: receiver operating characteristic, area under curve, weighted mean-squared error. Python tips for density estimation.

Week 14	**Bias-variance decomposition. Parameter estimation:** Maximum likelihood, maximum a posteriori, Bayesian inference, conjugate priors. Example: Gaussian case.	2.3.4-2.3.6, 3.2	HW10 assigned	Bias-variance examples from synthetic data; Python tips for parameter estimation techniques.
Week 15	**Regression based on parameter estimation:** Discriminative models. **Classification based on parameter estimation:** Generative models.	3.1.1, 3.3	Final project report due (Monday). HW10 due (Friday).	Sample problems in review for midterm exam
FINAL	Final exam.		Refer to the final exam schedule in the USC *Schedule of Classes* at classes.usc.edu.	

Statement on Academic Conduct and Support Systems

Academic Conduct:

Plagiarism – presenting someone else's ideas as your own, either verbatim or recast in your own words – is a serious academic offense with serious consequences. Please familiarize yourself with the discussion of plagiarism in SCampus in Part B, Section 11, "Behavior Violating University Standards" policy.usc.edu/scampus-part-b. Other forms of academic dishonesty are equally unacceptable. See additional information in SCampus and university policies on scientific misconduct, policy.usc.edu/scientific-misconduct.

Support Systems:

Counseling and Mental Health - (213) 740-9355 – 24/7 on call
studenthealth.usc.edu/counseling
Free and confidential mental health treatment for students, including short-term psychotherapy, group counseling, stress fitness workshops, and crisis intervention.

National Suicide Prevention Lifeline - 1 (800) 273-8255 – 24/7 on call
suicidepreventionlifeline.org
Free and confidential emotional support to people in suicidal crisis or emotional distress 24 hours a day, 7 days a week.

Relationship and Sexual Violence Prevention Services (RSVP) - (213) 740-9355(WELL), press "0" after hours – 24/7 on call
studenthealth.usc.edu/sexual-assault
Free and confidential therapy services, workshops, and training for situations related to gender-based harm.

Office of Equity and Diversity (OED) - (213) 740-5086 | Title IX – (213) 821-8298
equity.usc.edu, titleix.usc.edu
Information about how to get help or help someone affected by harassment or discrimination, rights of protected classes, reporting options, and additional resources for students, faculty, staff, visitors, and applicants.

Reporting Incidents of Bias or Harassment - (213) 740-5086 or (213) 821-8298
usc-advocate.symplicity.com/care_report
Avenue to report incidents of bias, hate crimes, and microaggressions to the Office of Equity and Diversity |Title IX for appropriate investigation, supportive measures, and response.

The Office of Disability Services and Programs - (213) 740-0776
dsp.usc.edu
Support and accommodations for students with disabilities. Services include assistance in providing readers/notetakers/interpreters, special accommodations for test taking needs, assistance with architectural barriers, assistive technology, and support for individual needs.

USC Campus Support and Intervention - (213) 821-4710
campussupport.usc.edu
Assists students and families in resolving complex personal, financial, and academic issues adversely affecting their success as a student.

Diversity at USC - (213) 740-2101
diversity.usc.edu
Information on events, programs and training, the Provost's Diversity and Inclusion Council, Diversity Liaisons for each academic school, chronology, participation, and various resources for students.

USC Emergency - UPC: (213) 740-4321, HSC: (323) 442-1000 – 24/7 on call
dps.usc.edu, emergency.usc.edu
Emergency assistance and avenue to report a crime. Latest updates regarding safety, including ways in which instruction will be continued if an officially declared emergency makes travel to campus infeasible.

USC Department of Public Safety - UPC: (213) 740-6000, HSC: (323) 442-120 – 24/7 on call
dps.usc.edu
Non-emergency assistance or information.