

Units: 4

Instructor: Keith M. Chugg
EEB 500A (5th floor main office)
213-740-7294 (Voice)
zoom: <https://usc.zoom.us/my/keithchugg>
chugg@usc.edu – Include 499 in subject

Office Hours: Tuesday, Thursday: 5:30 – 7:00 PM

Instructor: Antonio Ortega
EEB 436
213-740-2320 (Voice)
zoom: <https://usc.zoom.us/my/aortega>
aortega@usc.edu – Include 499 in subject

Office Hours: TBD

TA: TBD

Grader: TBD

Office: No office hours, email only for questions on grading. Grader is not a TA.

Lecture: Tuesday, Thursday, 3:30 – 5:20 PM in WPH 102

Discussion/Lab: Friday, 2:00 – 4:00 PM in WPH 102

Webpages:

- [Class Preview/Overview Page](#) for prospective students
- [Piazza Class Page](#) for most everything but grades and code
- [USC Blackboard Class Page](#) for grades
- [Github](#) for code examples used in lecture and discussion

- All HWs, handouts, solutions will be posted in PDF format
- *Student has the responsibility to stay current with webpage material*

Prereq: EE151L Introduction to Programming for Electrical Engineers (or equivalent proficiency, e.g., C++, Matlab, Python), E141L Applied Linear Algebra for Engineering (or MATH 225) , EE364 Probability and Statistics for EE/CS (or MATH 407)

Useful, but not required: EE301L

Other Requirements: Basic computer skills (e.g., plotting, Python or Matlab or other).

Grading:

- 30% Homework
- 10% Quiz 1
- 10% Quiz 2
- 20% Midterm Exam
- 30% Final Project

Note on e-mail vs. Piazza: If you have a question about the material or logistics of the class and wish to ask it electronically, please post it on the piazza page (not e-mail). You may post it anonymously if you wish. Often times, if one student has a question/comment, other also have a similar question/comment. Use e-mail with the professor, TA, graders only for issues that are specific to your individually (e.g., a scheduling issue or grade issue). You can post autonomously on piazza. Please do not use a private post on piazza.

Learning Objectives: Upon successful completion of this course a student will

- Be able to identify machine learning according to the taxonomy of supervised, unsupervised, reinforcement learning, etc.
- Apply methods of linear and nonlinear methods of regression or classification to data sets
- Apply principle component analysis to datasets and determine a reasonable amount of dimensionality reduction
- Select appropriate methods of optimization to train machine learning systems
- Apply data engineering concepts, such as cleaning, labeling, augmentation, to design and improve data sets for machine learning applications
- Be able to use Python-based machine learning tools, such as scikit-learn, Tensorflow, PyTorch to design machine learning solutions and evaluate the associated performance

Exam Dates:

- **Quiz 1:** Tuesday, February 7, 2023 (TBR, finalized Lecture 1)
- **Midterm Exam:** Thursday, March 2 (TBR, finalized Lecture 1)
- **Midterm Exam:** Tuesday, April 4 (TBR, finalized Lecture 1)
- **Final Project Presentation/Report:** Tuesday, May 9, 2:00 – 4:00 (university final exam slot)

Grading Policies:

- **Final grades** will be assigned by a combination of student score distribution (curve) and the discretion of the instructor. Final grades are nonnegotiable.
 - Final grades are non-negotiable and are assigned at the discretion of the instructor. If you cannot accept this condition, you should not enroll in this course.
- **Homework Description and Policy**
 - **Late HW** will not be accepted. A late assignment results in a zero grade. We will use electronic submission via Blackboard. Assignments will be due by midnight of the due date.
 - Homework will be assigned and collected typically every 1.5 weeks.
 - * Assignments will typically contain 3-4 analytical problems and 2-3 computational/programming problems
 - * Programming problems can typically be solved in less than 200 lines of Python.

- Show your work in your homework solution; the correct answer alone is worth only partial credit.
- Homework collaboration is encouraged. This is discussing problems and solution strategies with your classmates, the TA, and/or the instructor and is to be distinguished from copying solutions of others which is prohibited.

- **Exam Policy**

- **Make-up Exams:** No make-up exams will be given. If you cannot make the above dates due to a class schedule conflict, you must notify the instructor by the last day to add/drop. In the case of a medical emergency, please contact the instructor as soon as practical.
- Quizzes will be multiple choice
- Midterm will be closed book (with a crib sheet allowed).

- **Attendance:** Lecture attendance is encouraged but not mandatory. However, students are responsible for all material presented in lecture.

Textbooks:

- **Required Textbooks:**

- Jeremy Watt, Reza Borhani, Aggelos Katsaggelos, *Machine Learning Refined, 2nd Ed.*, Cambridge University Press.

- **Optional Textbooks:**

- Ethem Alpaydin, *Introduction to Machine Learning, 3rd Ed.*, MIT Press.
- Christopher Bishop, *Pattern Recognition and Machine Learning*, Springer, 2016.
- Michael Nielsen, [Neural Networks and Deep Learning](#)
- Kevin P. Murphy, *Machine Learning: A Probabilistic Perspective*, The MIT Press, 2012.
- [Tensorflow online documentation and examples](#)

Course Outline

1. Introduction and Motivation [**Watt Chapters 1 and 2**]
 - (a) ML definitions, problem statements, and tools
 - (b) Global/Local Optimization, Curse of Dimensionality
 - (c) Comparison to inference methods based on statistical models
 - (d) Applications of ML
2. Overview of Optimization Methods [**Watt Chapters 2.1-2.5, 3.1-3.5, 4.1, 4.3**]
 - (a) Global/Local Optimization, Curse of Dimensionality
 - (b) Zero order methods
 - (c) Gradient decent
 - (d) Second order methods
 - (e) Autograd tools in Python
3. Regression [**Watt Chapters 3.1-3.5, 5.1-5.2, 5.6**]
 - (a) Linear regression
 - (b) Linear regression review from EE364 and EE141L
 - (c) Multiple-output regression
 - (d) Alternative loss functions
4. Introduction to Data Engineering [**Watt Chapters 11.1-4, 11.7-11.8, 11.10**]
 - (a) Over-fitting
 - (b) Under-fitting
 - (c) Regularization
 - (d) Cross-validation
 - (e) augmentation
 - (f) Data formats and common data sets.
5. Classification [**Watt Chapters 6.1-6.4, 7.1-7.4, 7.6**]
 - (a) Linear regression for 2-class classification
 - (b) Logistic regression and the perceptron
 - (c) Cross-entropy loss
 - (d) Multi-class classification
 - (e) Example: digital modulation classification

QUIZ 1

6. Features and Unsupervised ML [**Watt Chapters 8.1-8.3, 8.5, 9.2, 9.7**]
 - (a) Sufficient statistics and features
 - (b) Principle Component Analysis

- (c) K-mean clustering
- (d) Example: common features in audio signal processing

7. Introduction to Nonlinear Methods [**Watt Chapters 10.1-10.2, 10.4**]

- (a) Nonlinear regression
- (b) Nonlinear classifiers

8. Introduction to Kernel Machine [**Watt Chapter 12**]

- (a) Universal approximation property
- (b) Kernel trick
- (c) Optimization and learning methods

QUIZ 2

9. Introduction to Neural Networks [**Watt Chapter 13, Neilson**]

- (a) Multilayer perceptions (MLPs)
- (b) Activation functions
- (c) Backpropagation learning

10. Decision Trees [**Watt Chapter 14.1-14.4, 14.6**]

- (a) Regression trees
- (b) Classification trees
- (c) Random forests

11. Deep Learning (Neural Networks) [**Slides, Tensorflow tutorials/docs**]

- (a) Optimizers
- (b) Regularizers
- (c) Training and evaluation of MLPs in TensorFlow

MIDTERM

12. Convolutional Neural Networks [**Slides, Tensorflow tutorials/docs**]

- (a) CNN architectures and conventions
- (b) Applications to computer vision problems

13. Recurrent Neural Networks [**Slides, Tensorflow tutorials/docs**]

- (a) RNN architectures and conventions
- (b) Applications to nonlinear filtering

14. Attention mechanisms [**Slides, Tensorflow tutorials/docs**]

- (a) Attention as an alternative to state
- (b) Introduction to transformers

(c) Example capabilities of Large Language Models (LLMs)

15. Selected Advanced Topics (time allowing) [**Slides, Tensorflow tutorials/docs**]

(a) Introduction to reinforcement learning

(b) Introduction to generative models (GANs, diffusion)

(c) Example capabilities of text-to-image systems and deep fakes

MIDTERM 1

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](#).

Support Systems:

[Student Health Counseling Services](#) – (213) 740-7711, 24/7 on call. 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) 740-4900, 24/7 on call Free and confidential therapy services, workshops, and training for situations related to gender-based harm.

[Office of Equity and Diversity \(OED\)](#) and [Title IX](#) – (213) 740-5086 Information about how to get help or help a survivor of harassment or discrimination, rights of protected classes, reporting options, and additional resources for students, faculty, staff, visitors, and applicants. The university prohibits discrimination or harassment based on the following protected characteristics: race, color, national origin, ancestry, religion, sex, gender, gender identity, gender expression, sexual orientation, age, physical disability, medical condition, mental disability, marital status, pregnancy, veteran status, genetic information, and any other characteristic which may be specified in applicable laws and governmental regulations.

[Bias Assessment Response and Support](#) – (213) 740-2421 Avenue to report incidents of bias, hate crimes, and microaggressions for appropriate investigation and response.

[The Office of Disability Services and Programs](#) – (213) 740-0776 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 Support and Advocacy](#) – (213) 821-4710 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 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 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 Non-emergency assistance or information.