DSCI 552: Machine Learning for Data Science

Units: 4

Term—Day—Time:

- Spring 2021
- Tuesdays and Thursdays from 10.00 am to 11.50 am
- Two 110 minutes classes per week
- 28 meetings in 15 weeks (+ student hours)

Location:

- Online: https://blackboard.usc.edu/
- Zoom Link: https://usc.zoom.us/j/93676908364 (check the passcode in the "Lecture Zoom Links" section on our Blackboard page)

Instructor:

- Dr. Marcin Abram
- e-mail: mjabram@usc.edu

Teaching Assistant:

- Ninareh Mehrabi
- e-mail: ninarehm@usc.edu
- Weekly meetings with your TA: Fridays 6-7 pm (check the Blackboard page for the Zoom link)

Graders:

- Supriya Devalla, e-mail: devalla@usc.edu
- Pratik Singhavi, e-mail: psinghav@usc.edu

Students Hours (also known as Office Hours):

- Two 60 minutes slots per week
- Monday 5-6 pm and Wednesday 6-7 pm
- Zoom Link: https://usc.zoom.us/j/8934576028 (check the Blackboard page for the passcode)
- Everybody is welcome
- No prior appointment is needed

IT Help:

- Blackboard Student Help, https://studentblackboardhelp.usc.edu/
- the Viterbi Service Desk https://viterbiit.usc.edu/get-help/

Course Description

DSCI 552 is an intermediate-level course in the Data Science program. It focuses on practical applications of machine learning techniques to real-world problems. During this course, you will learn how to apply and assess various machine learning algorithms, such as linear models, k-means, support vector machines, decision trees, random forests, neural networks, hidden Markov models, etc. You will learn how to analyze real-world datasets, how to design learning algorithms, how to train and evaluate machine learning models, and how to create technical reports that describe your findings.

This is a foundational course with the primary application to data analytics but this class is intended to be accessible both to students from technical backgrounds such as computer science, computer engineering, electrical engineering, or mathematics; and to students from less technical backgrounds such as business administration, accounting, various medical specializations including preventative medicine and personalized medicine, genomics, and management information systems.

Learning Objectives

At the end of this course, you will be able to:

1. Analyze quantitatively and qualitatively real-world datasets.
2. Describe and compare standard machine learning algorithms.
3. Choose or design learning algorithms suitable for a particular task.
4. Train and evaluate machine learning models.
5. Detect and assess bias in datasets as well as in the trained machine learning models.
6. Design a full machine learning pipeline.
7. Create a technical report describing your work and presenting your results.
8. Create peer-review reports.
9. Create a longer research article.
10. Present your findings in the form of a short presentation.

Prerequisite(s)

None, but see the Recommended Preparation section below.

Co-Requisite(s) or Concurrent Enrollment

None.

Recommended Preparation

To succeed with this class, review the information from:

★ Introduction to programming in Python (you should know how to write simple functions, how to use classes, how to plot figures using matplotlib, and use pandas, numpy, scikit-learn, and scipy).

- ★ (Optionally) Introduction to programming in R (using RStudio, fitting regressions using glm, plotting figures using ggplot). In general, Python is the recommended programming language in this class, but the first few assignments can be also completed using R.
- ★ Introduction to Mathematics: Calculus (e.g., operations on hyperbolic functions), Algebra (e.g., vector calculus, matrix notation), Probability (e.g., frequentist probability, Bayes theorem), and Statistics (e.g., hypothesis testing, interpretation of p-values, etc.).

Course Notes

This course will be comprised of:

- ★ lectures,
- ★ weekly readings and quizzes (short take-home tests, where I would ask questions related to an assigned article or a book chapter),
- ★ bi-weekly problem sets (where you will be asked to either analyze a dataset, design a learning algorithm or to train and evaluate a machine learning model; you will be asked to present both a technical report that describes your findings and your code),
- ★ a final project (that consists of several parts: over the semester you will submit a work plan, a literature review, an outline, an early draft, you will prepare two peer-reviews, incorporate reviewers' suggestions, and finally record a short presentation).

The course will ordinarily be taken for a letter grade. Documents, including lecture notes, homework assignments, and additional readings, will be distributed online via the course Blackboard site. There will be no final exam.

Description and Assessment of Assignments

Weekly readings and quizzes

Each Tuesday I will publish a short quiz (worth 10 points) on Blackboard. The questions will concern some basic ideas discussed in the class and/or the topics related to the recommended readings. There will be 12 quizzes in total. You will all have approximately 7 days to complete each quiz. As long as the quiz is open, you will be able to send multiple answers (the latest submitted answer will matter). The closing time for the quizzes is always on Tuesdays at 10 am PT (Pacific Time). Specifically,

- ★ The *first* quiz closes on Tuesday, January 26, 2021, at 10 pm PST.
- ★ The *second* quiz closes on Tuesday, February 2, 2021, at 10 am PST.
- ★ …
- ★ The *twelfth* quiz closes on Tuesday, April 20, 2021, at 10 am PDT (see the full schedule below).

Note, that the quiz submission deadlines coincide with the beginning of the Tuesday lectures. At the beginning of, each Tuesday lecture, I will discuss solutions to the last quiz. Therefore, extensions will not be possible and late submissions will be worth 0 points.

Bi-Weekly problem sets

Every second Thursday I will publish a problem set. Typically, those questions will require you to analyze a dataset, design a learning algorithm, or train and evaluate a machine learning model. Each problem set will be worth 20 points. There will be 6 problem sets in total. You will all have approximately 14 days to complete each problem set. The code must be published on GitHub or on a similar platform. The technical reports must be uploaded in the pdf format (you can either write the report in LaTeX or in a WYSIWYG editor, like Word, GoogleDoc, or LibreOffice - just remember to always export your report to the pdf format). The deadline for uploading the solutions is always on Thursday at 10 am PT. Specifically,

★ The deadline for the *first* problem set is on Thursday, February 4, 2021, at 10 am PST.
★ The deadline for the *second* problem set is on Thursday, February 18, 2020, at 10 am PST.
★ ...
★ The deadline for the *sixth* problem set is on Thursday, April 15, 2021, at 10 am PDT (see the full schedule below).

Note, that in those written and programming assignments, the completeness and the clarity of your description and analysis will matter as much as the final correct answer. Sending just a single final value or presenting a single final plot (even if correct) is not enough. See the table below:

Grade Component	**Meets Expectations (75%-100%)**	**Approaches Expectations (50%-75%)**	**Needs Improvement (0%-50%)**
Completeness (50%)	All parts of the question are addressed. E.g., if the task was to a) select a machine learning algorithm, b) train, and c) validate the model - the student completed all three parts.	Most parts of the question are addressed. E.g., if the task was to a) select a machine learning algorithm, b) train, and c) validate the model - the student selected and trained the model, but the validation part is missing or is incomplete.	The main question is not addressed. The answer is irrelevant to the task. The analysis of the issues and events are either vague or incomplete.
Clarity and Support (25%)	A non-expert (e.g., a fellow student) can understand the solutions. All concepts and used techniques are defined and explained. Whenever it is applicable, the solution is accompanied by illustrative plots that are explained and interpreted. There are references to sources. Accompanied code is well commented and easy to follow. The code follows both PEP8 and PEP257.	The teacher (or other professional computer scientists) can understand the solution but a non-expert might have some trouble doing so. The solution has some minor shortcuts or some non-explained assumptions. Not every step of the analysis is explained, but it is still possible to follow the author's logic. References are missing. The code is not well commented but it is still possible to understand it.	It is hard to follow the solutions. The solution has some major shortcuts or hidden assumptions. There are no references in the texts. The analysis or evaluation of the issues and events is vague. It is either hard or impossible to understand or verify correctness of the code.
Validity (25%)	All calculations are correct. The final values (or plots) are right and the final interpretation or conclusions are correct.	Small mistake in the code and/or calculations (e.g., a wrong sign, a missing constant). The final answer is close to the correct value (e.g., it differs by a small factor - twice too large or twice too small; however, the general trend is correct).	Major mistakes in the code and/or in the analysis. The final values and/or conclusions are incorrect.

The Final Project

Your task is to:

- ★ prepare a technical report and/or a scientific article (limit of 5000 words; can be shorter, the recommended length is about 3500 words) on one of the topics below.
- ★ Peer-review two articles prepared by your colleagues.
- ★ Address the comments that you received from your peers.
- ★ Record a short summary (2-3 minutes) of your work (as a simple interview-style video presentation, business-style pitch/demo or a conference-style narrated slideshow).

The objective of this assignment is to a) explore literature regarding data science and machine learning, b) synthesize the acquired knowledge in the form of article, c) learn how to write peer-review comments, d) learn how to respond to peer-review comments, e) be able to summarize a weeks-long project in a form of a condensed, short presentation.

Projects Propositions (choose one):

- ❏ *(For those who like open questions)* Limits of Machine Learning. A possible starting point, David J. Hand, “Classifier Technology and the Illusion of Progress” (2006). This is a relatively old paper. Are those findings still hold? What is written about this topic in more contemporary articles? Another starting point could be Anthony M. Zador, “A critique of pure learning and what artificial neural networks can learn from animal brains” (2019) if you are more into biology or Axel Seifert and Stephan Rasp, “Potential and Limitations of Machine Learning for Modeling Warm-Rain Cloud Microphysical Processes” (2020) if you are more into earth science. This is a vast topic and you would be expected to do an intensive, individual literature review. You would be also expected to have at least a small quantitative section (it could be e.g., a demonstration of how a popular model can be fooled or a demonstration, how the performance of the neural network deteriorates due to the domain shift; you can do something original or you could repeat some results from one of the paper - both approaches would be acceptable).
- ❏ *(For those who like reading)* AI Ethics. A possible starting point, Nick Bostrom, “Superintelligence: Paths, Dangers, Strategies” (2014) or Cathy O'Neil, “Weapons of Math Destruction” (2016). Those books can not be your only sources. You would need to also find some relevant (peer-reviewed) articles about your topic. Nevertheless, those books can give you a good starting point and inspire you to do further research. You would be expected to include in your paper an exhausting literature review section. Additionally, your project should be at least partly-quantitative (though, the quantitative part can be shorter than in other projects). Some possible paths: You can detect or quantify bias in various pre-trained models; you can benchmark or compare various techniques that promise to reduce bias, etc. Some possible starting points: Ian J. Goodfellow, Jonathon Shlens, Christian Szegedy, “Explaining and Harnessing Adversarial Examples” (2014) and Douglas Heaven, “Why deep-learning AIs are so easy to fool” (2019).
- ❏ (*If you care about social justice*) Machine Learning and Social Justice. A possible starting point: articles and books by Ruha Benjamin, e.g. selected chapters from “Race After

Technology: Abolitionist Tools for the New Jim Code" (2019) or "Captivating Technology: Race, Carceral Technoscience, and Liberatory Imagination in Everyday Life" (2019). You could also read Shakir Mohamed, Marie-Therese Png, and William Isaac, "Decolonial AI: Decolonial Theory as Sociotechnical Foresight in Artificial Intelligence" (2020) or Pratyusha Kalluri, "Don't ask if AI is good or fair, ask how it shifts power" (2020). Those books and articles can not be your only sources. You are expected to do your own literature review. You should use mostly peer-reviewed books and articles (this is a general remark to all those projects). The main bulk of your article could focus on the discussion, comparison, or critique of various papers and ideas. However, you are still expected to have some quantitative sections in your paper. Each project must be at least partly-quantitative (though, the quantitative part of this project can be shorter compared with other topics). Possible ideas: You could detect or quantify racial bias in various popular datasets; Possible starting points: Joy Buolamwini, "Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification" (2018), Moin Nadeem, Anna Bethke, and Siva Reddy, "StereoSet: Measuring stereotypical bias in pretrained language models" (2020), Jungseock Joo and Kimmo Kärkkäinen, "Gender Slopes: Counterfactual Fairness for Computer Vision Models by Attribute Manipulation" (2020) or Rachel Rudinger et al. "Gender Bias in Coreference Resolution" (2018).

- ❏ *(For those who like math)* Private Machine Learning. You can go in two different directions. One direction is related to data privacy. Here, a good starting point would be articles describing differential privacy, e.g., Damien Desfontaines and Balázs Pejó, "SoK: Differential privacies" (2020). To see how neural networks can leak private information, see e.g., Matt Fredrikson et al. "Model Inversion Attacks that Exploit Confidence Information and Basic Countermeasures" (2015) - if you are interested in image recognition - or Nicholas Carlini et al. "Extracting Training Data from Large Language Models" (2012) and Nicholas Carlini et al. "The Secret Sharer: Evaluating and Testing Unintended Memorization in Neural Networks" (2018) - if you are interested in language models. Another direction is to read about methods of preserving the privacy of models - e.g., preventing users from reverse-engineering proprietary models. Here, the direction would be to read e.g., about homomorphic encryption and/or secure multi-party computations in the context of distributed machine learning.
- ❏ *(For those who like tinkering)* Design and deploying machine learning algorithms on a Raspberry Pi (or a similar system). You can create a weather station that gathers data about the weather and tries to predict the weather in the next hour. You can create a device that will measure (using a camera and an object detection model) traffic and inform the user if the traffic is larger or smaller than the average. There are countless possibilities. You are encouraged to propose something original. Your paper should still be structured like any other paper, with an introduction, a literature review, methods, and result sections. Just, in this case, the literature review can be shorter, and your methods and results sections will be longer. In the methods section, you would describe how you constructed and tested your system. Your presentation can be also altered; instead of narrated slides (a typical, conference-style), you can record a demonstration of how your system works (it would be more business-style pitch or demo presentation).

- ❑ *(For those who like a practical approach to machine learning)* Machine Learning at scale. You have a couple of possible directions. First can be related to massive models trained on supercomputers. Here you could explore the literature regarding e.g., large NLP models. Another direction is to focus on distributed training on the edge- and end-devices. Here, you can start by reading about federated machine learning (check for example the TensorFlow Federated library).
- ❑ *(For those who like pure machine-learning problems)* You can explore transfer learning methods. For example, you can demonstrate "catastrophic forgetting". A good starting point for you can be the article by James Kirkpatrick et al. "Overcoming catastrophic forgetting in neural networks" (2017). Another direction would be to talk about model robustness. Here, a good starting point would be "Strengthening Deep Neural Networks", a book by Katy Warr. You can also look at the works of Judy Hoffman, e.g., Judy Hoffman, Daniel A. Roberts, Sho Yaida, "Robust Learning with Jacobian Regularization" (2019) or Yogesh Balaji, Tom Goldstein, Judy Hoffman, "Instance adaptive adversarial training: Improved accuracy tradeoffs in neural nets" (2019).
- ❑ *(For those who like to know how things works)* Interpretability of Machine Learning Models. Explore various methods, that can be used to illustrate how various machine learning models (e.g., neural networks) learn patterns from data. A good starting point can be Chris Olah et al. "The Building Blocks of Interpretability" (2018). Other relevant papers can be: Aravindh Mehandran and Andrea Veldadi, "Understanding Deep Image Representations by Inverting Them" (2014) and Alexey Dosovitskiy and Thomas Brox, "Inverting Visual Representations with Convolutional Networks"(2016). If you have a physics or chemists background, you can also look at the topic of the loss function landscape, here the starting point could be Hao Li at al., "Visualizing the Loss Landscape of Neural Nets" (2018) and Sathya R Chitturi et al. "Perspective: new insights from loss function landscapes of neural networks" (2020).
- ❑ *(For those who want to be creative)* Machine Learning and Natural Language Processing models for a creative generation. For instance, people have used generative models to generate stories e.g., refer to Yao et al. "Plan-and-Write: Towards Better Automatic Storytelling" (2019), to generate poetry, cf. Ghazvininejad et al. "Generating Topical Poetry" (2016), to generate music (see e.g., Google research's Magenta project), images, cf. Dosovitskiy et al. "Generating Images with Perceptual Similarity Metrics based on Deep Networks" (2016), or videos, cf. Vondrick et al. "Generating Videos with Scene Dynamics" (2016). Of course, you are not restricted to any of these mediums. We want you to be creative and have the freedom of choosing whatever excites you.
- ❑ *(For those who want to know how things are evaluated)* Model Selection and Evaluation. Depending on the application, Machine Learning and Natural Language Processing models get evaluated differently. In addition, evaluation of some models requires extensive human annotations which is infeasible in some cases; thus, many researchers have tried to use Machine Learning models to evaluate other Machine Learning models! (Yes this sounds both weird and interesting). For instance, there is active research taking place on how to evaluate dialogue systems e.g., refer to Jan Deriu et al. "Survey on Evaluation Methods for Dialogue Systems" (2020). For this project, you can also select a specific task and explore challenges in the evaluation of that specific task - by studying what people

have done and examining what you can add-on to improve the existing evaluation metrics and methods.

- ❑ *(For those who want to analyze vulnerability or robustness of Machine Learning Models)* Machine Learning and Natural Language Processing models are shown to be vulnerable towards different adversarial attacks, cf. Anthony D. Joseph et al. "Adversarial Machine Learning" (2019). Different attacks and defenses are proposed to study, understand, and improve these flaws. For this project, you can also study and quantitatively analyze the vulnerabilities of some systems and models. You can also propose solutions for robust training if possible. In case anyone is interested in adversarial Natural Language Processing or robust training, there is a lot of work in that area as well cf., e.g., Yitong Li et al. "Robust Training under Linguistic Adversity" (2017).
- ❑ *(For those who don't like the above projects)* Modify the above propositions or propose your own project. Discuss your choice with the instructor.

Structure and Formatting:

We encourage you to use the LaTeX template https://www.overleaf.com/read/crjbtrfftfhg that I prepared for you in Overleaf. If you use a WYSIWYG editor, please remember to submit your article in PDF format (not as a docx, rtf, or odt). In the paper, you must provide a link to a GitHub repository with the relevant code, scripts, or notebooks. Python 3.8+ is preferred, but in principle, you are free to use any language of your choice - as long as the code is clear and well commented (the reader should be able to understand your code even if he or she is not proficient in it; the reader should be able to go to your repository, clone your repository and run your code without getting any errors).

Steps:

1. Prepare and post a work plan **by Thursday, January 28, by 10 am**.
2. Choose your topic.
3. Find relevant literature. Read about your topic. Prepare a literature review **by Thursday, February 11, by 10 am**.
4. Make a plan for your article. Decide which aspects you are going to describe and which you will leave out. After all, you have limited space (only a couple of pages, including figures and bibliography). Submit your outline **by Thursday, February 25, by 10 am**.
5. Complete the necessary coding and calculations. Prepare plots and figures.
6. Write the first version of your article. You should have **an early draft by March 11**. You do not have to upload it yet - it is your internal deadline.
7. Proofread your article. Make sure that all key terms are defined. Make sure that the article has the right structure (abstract, introduction, the main content, discussion/summary, and bibliography). Remember, that the list of references at the end of your paper is not enough - your sources must be cited in the article (see the template).
8. Prepare a pdf of your article. Make sure that the number of words is below the maximum limit. Make sure that your name, affiliation, abstract and paper title are visible on the first page. Submit the pdf using Blackboard **by Thursday, March 25, not later than 10.00 am**.
9. Choose two articles prepared by your peers (we will coordinate this process, to make sure that each article gets an equal number of reviewers). Read the articles. Using the

Blackboard forum, give each author suggestions on how they can improve the papers. You should complete this action **by Thursday, April 8, by 10.00 am**.

10. Read the suggestions that you received from your peers. Address them (either incorporate the suggested changes or challenge them, describing why you think those changes would not improve the quality of your article).
11. Submit your final article **by Tuesday, April 27, by 10.00 am**.
12. Record a short summary of your work (2-3 minutes), either as a video-presentation or a narrated slideshow. Submit your video **by Tuesday, April 27, no later than 10.00 am**.
13. Write an academic reflection, summarizing your experience. Submit it **by Thursday, April 29, by 10 am**.

Additional Notes:

You are free to use any sources. However, you must cite all sources that you used (if not, you will violate the academic integrity standards). It might happen that you will cite non-peer-reviewed sources, like technical documentation of certain libraries or technical blog posts. This is acceptable as long as the non-peer-reviewed sources do not constitute the majority of your bibliography. If you decided to use quotes, remember to cite them correctly. Plagiarism (or using sources without proper citations) is a major violation of the university academic integrity standards and will be reported to the Office of Student Judicial Affairs and Community Standards at USC, see details at https://sjacs.usc.edu/students/academic-integrity/ and to have an overview of the sanctions and penalties, check the Appendix A: Academic Dishonesty Sanction Guidelines.

When you write your article, think about your audience. Your main audience is not the instructor, but rather your peers. Write in a way that your colleagues can understand the concepts that you describe. You can assume certain fluency in math and technology in your readers, but do not assume that your audience has any specific prior familiarity with the topic of your paper.

Technological Proficiency and Hardware/Software Required

Basic knowledge of programming is required (Python 3 recommended but familiarity with any major programming language can be sufficient). Do not use Python 2 - this version is not supported anymore. Python 3.8 or newer is recommended.

Required Readings and Supplementary Materials

★ *(our main textbook; theory)* Ethem Alpaydin, *Introduction to Machine Learning*, 3rd Edition, MIT Press (2014).

★ *(our main textbook; practice)* Aurélien Géron, *Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems*, 2nd Edition, O'Reilly (2019).

★ *(our supplementary textbook; statistical learning perspective; theory and practice)* Gareth James, Daniela Witten, Trevor Hastie, Robert Tibshirani, *An Introduction to Statistical Learning*, Springer (2013, corrected at 8th printing 2017).

★ *(our supplementary textbook; more mathematically-heavy than Gareth James et al.; theory)* Trevor Hastie, Robert Tibshirani, and Jerome Friedman, *The Elements of Statistical Learning*, 2nd Edition, Springer (2009, corrected at 12th printing 2017).

★ *(our supplementary textbook in topics regarding neural networks; theory)* Ian Goodfellow, Yoshua Bengio, and Aaron Courville, *Deep Learning*, MIT Press (2016).
★ *(our supplementary textbook in topics regarding neural networks; practice)* François Chollet, *Deep Learning with Python*, Manning (2017).
★ *(our supplementary textbook in topics regarding unsupervised learning; practice)* Ankur A. Patel, *Hands-On Unsupervised Learning Using Python*, O'Reilly (2019).
★ *(our supplementary textbook in topics regarding bias and fairness; practice and theory)* Aileen Nielsen, *Practical Fairness*, O'Reilly (2020).
★ *(our supplementary textbook in topics regarding robustness and adversarial attacks; practice and theory)* Katy Warr, *Strengthening Deep Neural Networks*, O'Reilly (2019).
★ *(our supplementary textbook in topics regarding generative models; practice and theory)* David Foster, *Generative Deep Learning*, O'Reilly (2019).
★ *(your supplementary textbook if you need a review from statistics)* Peter Bruce, Andrew Bruce, Peter Gedeck, *Practical Statistics for Data Scientists*, 2nd Edition, O'Reilly (2020).

Note that the digital versions of those books are available for *free* via the USC libraries. You can search for books using https://libraries.usc.edu/. Note that USC students have also free access to the O'Reilly books. To register, go to https://learning.oreilly.com/home/.

You will also be asked to read various research articles. Those articles will be distributed via Blackboard.

Grading Breakdown

Course Element	Points
Weekly Quizzes(12)	120 (=12x10)
Bi-Weekly Problem Sets (6)	120 (=6x20)
Work Plan	5
Literature Review	15
Project Outline	10
Project Draft	20
Peer Reviews	20
Student Project	100
Final Presentation	20
Academic Reflection	10
TOTAL	**440**

Grading Scale

Course final grades will be determined using the following scale.

Final Grade	% of Total Points	Number of Total Points (rounded down)
A	[92% - 100%]	404-440
A-	[89% - 92%)	391-403.9
B+	[86% - 89%)	378-390.9
B	[81% - 86%)	356-377.9
B-	[78% - 81%)	343-355.9
C+	[75% - 78%)	330-342.9
C	[70% - 75%)	308-329.9
C-	[67% - 70%)	294-307.9
D+	[64% - 67%)	281-293.9
D	[59% - 64%)	259-280.9
D-	[55% - 59%)	242-258.9
F	[0% - 55%)	0-241.9

Assignment Submission Policy

Late solutions to quizzes or problem set solutions will not be accepted. Note, that the grading brackets in the table above are lower than in other courses (for example C is from 70% not from 73%). It means that you can be late (or omit) one quizz (10 points) and one problem set (20 points) and still be able to collect up to 410 points and receive A.

Note, that all the submission deadlines coincide with the beginning of the lectures. At the beginning of, each lecture, I will discuss solutions to the last quiz or problem set. Therefore, any extensions will not be possible and late submissions will be worth 0 points.

Grading Timeline

I will make every effort to grade and return homework within one week after it is received. Homework solutions will be either described during the lectures or posted on Blackboard.

Additional Policies

Names, Gender:
If you have a name and/or pronouns that differ from those in your official USC records, please let me know. If I am mispronouncing your name, please correct me. I am highly empathetic on this point because my given name (Marcin) is pronounced [ˈmart͡ɕin] using the International Phonetic Alphabet and often mispronounced in English-speaking countries.

Mental Health:
If you feel that experiences outside of class are impacting your course performance, please come and talk to me. If you would rather consult someone outside the classroom, USC Counseling and

Mental Health (https://studenthealth.usc.edu/counseling/) and Academic Counseling (https://undergrad.usc.edu/services/counseling/) are great resources.

Equity and Diversity and Title IX:

The Office of Equity and Diversity (OED) and the USC Title IX Office works with faculty, staff, visitors, applicants, and students around issues of protected class: https://eeotix.usc.edu/. Incidents of bias, hate crimes and microaggressions can be confidentially reported to: https://studentaffairs.usc.edu/bias-assessment-response-support/.

Accommodations:

Any student requesting academic accommodations based on a disability is required to register with Disability Services and Programs (DSP) each semester. A letter of verification for approved accommodations can be obtained from DSP. Please be sure the letter is delivered to me as early in the semester as possible. If you have registered accommodations with the Disability Services and Programs Office (https://dsp.usc.edu/), please communicate those to me at your earliest convenience so we can discuss your needs in this course. For those on or near campus, DSP is located in STU 301 and is open 8:30 a.m. - 5:00 p.m, Monday through Friday. They can be contacted online or by phone at (213) 740-0776 (Phone), (213) 740-6948 (TDD only), or via email, ability@usc.edu.

Statement for the observance of religious holidays

USC's policy grants students excused absences from class to observe religious holidays: http://orl.usc.edu/life/calendar/absences/. In this case, please contact your instructor in advance to agree on alternative course requirements.

Zoom Classroom Policies

The pandemic has upended our collective and individual lives. Logistically speaking, we are spread across multiple time zones, and I can only expect attendance for students for whom our course time falls within reasonable learning hours in their time zone, i.e., between 7:00 AM and 10:00 PM, see https://www.provost.usc.edu/policy-and-guidelines-for-asynchronous-learning/ (the section about the Class Participation and Attendance in Synchronous Sessions). If you are in a timezone, that prevents you from attending the classes (or if you have other situations like family responsibilities, e.g. taking care of children or dependents, that prevent you from attending the synchronous sessions), please let me know as soon as possible.

Camera Policy

The official Camera Policy can be found at
https://www.provost.usc.edu/policy-and-guidelines-for-asynchronous-learning/.

Seeing your faces can help me to gauge if the tempo of the lectures is adequate. Therefore, it would be a great help if you keep your cameras turned on. However, I acknowledge that there might be many reasons why you might wish to keep your privacy. You might also face bandwidth limitations that prevent you from using the camera. I encourage the use of virtual backgrounds and earphones/headsets whenever it is possible to mitigate privacy concerns.

Course Schedule: A Weekly Breakdown

	Topics	Readings	Deliverables
Week 1 January 19 January 21	Introduction to machine learning. Machine learning work pipeline. Datasets. General overview: linear vs. non-linear models; supervised, semi-supervised, and unsupervised training; model validation; underfitting and overfitting. Structure of the class and students' project overview.	Alpaydin Ch. 1. James Ch. 1-2. Hastie Ch. 1-2. Géron Chs 1-2.	
Week 2 January 26 January 28	Linear Models. Linear regression. Regularization. Rigid regression. LASSO. Elastic Net. Kernel methods. Support vector machine (optional).	Alpaydin Ch. 2. James Chs. 3, 6. Hastie Ch. 3.	Quiz 1 (Jan 26) Work Plan (Jan 28)
Week 3 February 2 February 4	Classification. Logistic regression. Performance measures: confusion matrix, F1 score, AUC. K-Nearest Neighbours. Nested models (optional). Bayesian theory. Parametric Models. Multivariate methods.	Alpaydin Chs. 3-5 James Chs. 4. Géron Ch 3.	Quiz 2 (Feb 2) Problem Set 1 (Feb 4)
Week 4 February 9 February 11	Unsupervised machine learning. Dimensionality reduction. Clustering. K-Means. PCA. t-SNE. Non-parametric models. Gaussian Mixture Model.	Alpaydin Chs. 6-8 Géron Ch 8.	Quiz 3 (Feb 9) Literature Review (Feb 11)
Week 5 February 16 February 18	Introduction to Neural Networks. Feedforward neural networks. Universality theorem for neural networks. Deep neural networks vs. wide neural networks. Regularization. Dropout. Overfitting. Early stop.	Alpaydin Ch. 11. Géron Chs 10-12. Chollet Chs. 1-3. Goodfellow Ch. 1.	Quiz 4 (Feb 16) Problem Set 2 (Feb 18)
Week 6 February 23 February 25	Modern machine learning practices. The bias-variance trade-off revisited. Massively overtrained neural networks and deep double and triple descent concept.	ArXiv 1812.11118 ArXiv 1912.02292 ArXiv 2006.10455 ArXiv 2006.03509 Goodfellow Chs. 6-9.	Quiz 5 (Feb 23) Project Outline (Feb 25)
Week 7 March 2 March 4	Symmetries and Invariance. Convolutional neural networks (CNN). Deep learning. Model architectures: LeNet, AlexNet, VGG, ResNet. Variational autoencoders. Combined learners (ensemble of models).	Articles TBA. Géron Chs 15-16. Goodfellow Ch. 9.	Quiz 6 (Mar 2) Problem Set 3 (Mar 4)
Week 8 March 9 March 11	Introduction to Natural language processing. Naive Bayes. Word embeddings. Recurrent neural networks: RNN, GRU, LSTM. Sentiment analysis. Attention (optional). *(Midterm Grading Period begins)*	Articles TBA. Géron Chs 14. Goodfellow Ch. 10.	Quiz 7 (Mar 9) Early Draft (Mar 11)
Week 9 March 16 March 18	Model Robustness. Adversarial attack. Fairness in machine learning. Bias in data and models. AI ethics.	Articles TBA. Nielsen Chs. 1-6.	Quiz 8 (Mar 16) Problem Set 4 (Mar 18)
Week 10 --- March 25	Data privacy. Model inversion. *(March 23 is designated as a Wellness Day)*	Articles TBA.	Draft (Mar 25)
Week 11 March 30 April 1	Decision trees. Bagging and bootstrapping. Random forest Ensemble. Gradient boosting machines (optional). Generalized random forests (optional).	Articles TBA. Alpaydin Chs. 9, 17. Géron Chs 6-7.	Quiz 9 (Mar 30) Problem Set 5 (Apr 1)
Week 12 April 6 April 8	Generative models (GAN). Variational autoencoder. (Optional) Special architectures: U-Net, object segmentation, deep-wide neural networks, working on heterogeneous datasets. *(Midterm Grading Period ends)*	Géron Ch 17. Nielsen Ch. 20.	Quiz 10 (Apr 6) Peer-Reviews (Apr 8)
Week 13 April 13 April 15	Causal inference.	TBA	Quiz 11 (Apr 13) Problem Set 6 (Apr 15)
Week 14 April 20 ---	Buffer week or a Special Topic I, e.g.: Online-machine learning, multi-armed bandit, algorithmic auctions, elements of reinforcement learning. (April 22 is designated as a Wellness Day)	Various Articles.	Quiz 12 (Apr 20)

Week 15 April 27 April 29	Special Topics II, e.g.: Adversarial attacks, data leakage, transfer learning, fine-tuning, machine learning at scale, distributed machine learning, deploying and maintaining a machine learning pipeline. Bonus: Job perspectives and job market for data scientists and machine learning engineers, working as a machine learning scientist in startups (a personal perspective).	Various Articles.	Final Project (Apr 27) Final Presentation (Apr 27) Academic Reflection (Apr 29)
FINAL May 11	There is no exam on May 11. You have a longer project instead. *(Grading Period ends May 18)*		

Student Hours (also known as Office Hours)

I will host two 60 minute meetings per week, on Monday at 5-6 pm and on Wednesday at 6-7 pm. You can access them via Zoom: https://usc.zoom.us/j/8934576028 (you will find the passcode on Blackboard). Those Student Hours (also known as Office Hours) are a dedicated time-slots when you can come to ask questions and resolve confusion about course material, as well as discuss career and educational goals as they relate to this course.

No appointment needed, however, if you sent me an email a day earlier, announcing a type of question you have, I might be able to prepare a better answer for you in advance.

If you have any sensitive questions, you can also contact me via mail, mjabram@usc.edu and we can schedule a 1-on-1 appointment via Zoom outside the student hours period.

Contact for Support Systems

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 and 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. 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. The university also prohibits sexual assault, non-consensual sexual contact, sexual misconduct, intimate partner violence, stalking, malicious dissuasion, retaliation, and violation of interim measures.

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 Support and Advocacy - (213) 821-4710
uscsa.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.