DATA 301
Syllabus at a Glance

Course Info

- **Professor**: Dr. Joanna Bieri (joanna_bieri@redlands.edu)
- **Meets**: Tuesday, Thursday - 2:40-3:55pm, Hall of Letters 101
- **Office Hours**: posted on the class website
- **Format**: flipped and individual. Watch the videos before class, and all graded work lives in your own private GitHub repository
- **Time Commitment**: minimum 4 hours outside class per hour in class (∼10 hrs/week)

Key Dates - Fall 2026

- Classes Begin: Tuesday, September 1
- Exam 1: released Thursday, October 1 - due Monday, October 5
- No Class - Fall Break: Tuesday, October 13
- Final Project Teams Assigned: Thursday, October 22
- Project Proposals Due: Tuesday, November 3
- Exam 2: released Thursday, November 5 - due Monday, November 9
- No Class - Thanksgiving: Thursday, November 26
- Final Exam, Team Project Presentations: Thursday, December 10, 3:00pm

→ Full day-by-day course schedule

Grading Breakdown

Pre-Class Preparation & Quizzes	10%
In-Class Participation	10%
Weekly Assignments	30%
Take-Home Exams	30%
Final Project	20%

→ Full grading scale and evaluation procedures

Quick Links

- Classwork Expectations
- Communication
- Academic Honesty & AI Policy
- Accommodations & Support
- Your Git Workflow

DATA 301
Introduction to Machine Learning

Professor:

Dr. Joanna Bieri

joanna_bieri@redlands.edu

Class:
Fall 2026
T, Th: 2:40-3:55
Hall of Letters 101

Office Hours:
Posted on the class website

Welcome to Machine Learning

Machine learning is the practice of building programs that improve from experience instead of from instructions you write by hand. This course takes you from that idea all the way to understanding, at a mechanical level, how a large language model works. Classes, lectures, and group interactions should be a safe learning space where each member of the community feels valued, listened to, and respected. It is the responsibility of each of us to ensure that we are lifting up our peers, breaking free from old biases, reaching across cultural or socioeconomic boundaries, and supporting each other. If at any time in the class you feel that there is something that could be done to improve your learning or your experience please let me know. I am here to lift you up and help you learn!

You come to this class from a wide range of backgrounds! Some of you are mathematics or physics majors, others come from business, biology, or data science and that is a big advantage. The course is built so that the ideas come first and the mathematics arrives when you need it. What everyone will share by December is the ability to train a model, evaluate it honestly, and explain why it does what it does.

This is a **flipped** class. Lecture videos are posted before each class meeting, and you are expected to watch them beforehand. Our time together is spent on the harder problems, on your questions, and on writing code. Attendance is expected at all class meetings unless you have notified me ahead of time.

Course Learning Objectives

By the end of the semester your work must reflect your ability to

1. Frame a real problem as a machine learning task and recognize the cases where machine learning is the wrong tool.
2. Build, train, and evaluate models in Python using scikit-learn and PyTorch.
3. Design an honest evaluation: split data without leakage, choose metrics that match the question being asked, and report what your results do *not* show.
4. Diagnose underfitting and overfitting from learning curves, and apply regularization deliberately rather than by habit.
5. Explain how a neural network learns and write a training loop yourself.
6. Explain the architecture of a transformer, and how a large language model is pretrained, fine-tuned, and used.
7. Use version control to manage, document, and share your own work.
8. Evaluate the limitations, failure modes, and ethical stakes of a model that gets deployed on real people.
9. Communicate a complete machine learning project in written, oral, and graphical form.
10. Have some fun exploring machine learning!

What are your goals for the class?

Time Commitment

This class meets 2.5 hours per week. Plan to spend a minimum of 4 hours outside of class for every hour in class, roughly 10 hours per week on videos, reading, weekly assignments, and project work, in addition to class time. Training models takes real time; start assignments early enough that a long training run isn't the thing standing between you and the deadline.

Technology

You absolutely need access to a computer that can run Python and JupyterLab. We will install everything you need together on the first day of class, including the JupyterLab Git extension, which is how you will move your work to GitHub. We will use JupyterLab for our programming and GitHub for version control.

Everything in this course is designed to run on your own laptop, on CPU. You do not need a fancy GPU. The models we build are deliberately small enough to train in minutes rather than hours. For the final project, if your team needs more compute than a laptop can give, talk to me and we will arrange access to a cloud GPU.

Homework and announcements will be posted on the class website. You submit your work as a Pull Request in your own GitHub repository; Canvas is used for the short comprehension quizzes and for exam submission.

Laptop Lending Program: If you do not have a computer, or your computer cannot handle the programming for this course, the Data Science program has a laptop lending program that can get you a computer capable of completing all the work in this class. Contact me directly and we will get you set up.

Textbook

- *Hands-On Machine Learning with Scikit-Learn and PyTorch* by Aurélien Géron. This is the new PyTorch edition. Please make sure you do not buy the older Keras/TensorFlow version, as the code will not match ours.

All of the book's code notebooks are freely available online, and the readings for each class meeting are listed in the course schedule at the end of this syllabus.

Classwork

Consistent participation and engagement are crucial to your success in the class. Please read the following carefully:
Unlike DATA 201, the graded work in this course is **individual**. Each of you has your own private GitHub repository, and you are responsible for your own code and your own writing. The one exception is the final project, which you will do in a team.

i. **Pre-Class Preparation**: This is a flipped class, so preparation is not optional, it is where the lecture actually happens. Before each class meeting, watch the posted videos and do the assigned reading. Most class meetings will begin with a short comprehension quiz on Canvas covering the video. These are low-stakes and meant to tell us both whether the ideas are making sense.

ii. **In-Class Participation**: In class you should be actively working on and discussing machine learning. Bring your computer and try out all the things we are learning about! There will also be short daily practice problems to work through during and after class, these are not collected, but the weekly assignment assumes you have done them. Missing more than 2 days of class (unexcused) will result in a 50% reduction in your In-Class Participation grade and each additional missed class will decrease your grade further.

iii. **Weekly Assignments**: Rather than a small assignment every day, you will have one substantial assignment each week, due Sunday evening. Each assignment should include working Python code along with a written discussion: what you tried, what the evaluation actually showed, and what you concluded. A model with a number attached to it is not a result until you can say what the number means and what would make it wrong. You submit by opening a Pull Request into main in your own repository; I review it and leave comments on the Pull Request, and that is your feedback. Assignments are graded pass-fail. If you fail you can resubmit the assignment once.

iv. **Exams**: There will be two take-home exams. Each is released on a Thursday afternoon after class and is due the following Monday at 11:59pm, so you have the weekend to work. Exams are individual, closed to outside help, and **AI-free**, see the policy below. You can use any code that you generated as part of the HW and daily studying to aid you on the exam. Exam 1 covers evaluation methodology and ensemble methods; Exam 2 covers neural networks and representations. There is no exam on the transformer and large language model material, that is what the final project is for.

v. **Final Project**: Working in a team of 3-4 students you will complete a machine learning project inspired by an area of interest to the group. Teams are assigned in late October and proposals are due November 3. Your team presents at our final exam, Thursday December 10 at 3:00pm. This is the one piece of collaborative work in the course, and it is where you show that you can carry a problem all the way from a question to an honestly evaluated answer.

Evaluation Procedures

There are five elements of the course which contribute to your overall grade, as follows:

Pre-Class Preparation & Quizzes	10%
In-Class Participation	10%
Weekly Assignments	30%
Take-Home Exams	30%
Final Project	20%

How to Succeed in this Class

All work for the course will be graded according to the following criteria.

- Students can demonstrate pre-class preparation by coming to class prepared. This means watching all the videos, completing all required reading, and starting on the assignments or worksheets. The short in-class quizzes are part of this.

- Students can demonstrate in-class participation by attending every class. In addition, students should be actively involved in the discussion during class time. For some this will be active speaking, for others it will be taking notes and posting them for the class, for others it will be talking in small groups or helping with collaboration. There are lots of ways to participate, let me know how you best like to contribute!

- Each weekly assignment is submitted as a Pull Request in your own GitHub repository and graded for 10 points, pass-fail, on the code and the written discussion together. See the assignment prompts on the class website for specific details.

- The final project is your chance to show mastery of the concepts for the class. You will propose your final project during the semester and work with your peers and professor to clearly define your goals.

Grading Scales:

4.0	95-100%
3.7	90-94%
3.3	87-89%
3.0	83-86%
2.7	80-82%
2.3	77-79%
2.0	73-76%
1.7	70-72%
1.3	67-69%
1.0	63-66%
0.7	60-62%

Please note: *According to the University Course Catalog (p. 13) 3.7 and 4.0 are reserved for "outstanding" work; 3.3 and 3.0 are both defined as "excellent," not mediocre.*

Communication

- The most reliable way to reach me is by email. Please note that my normal working hours are 9 a.m. to 5 p.m., Monday to Friday. I do not respond to emails after 5 p.m. or on weekends, except in an emergency.
- You can make appointments with me via email. Appointments can happen in Duke 209 or on Teams.
- It is important that you communicate throughout the semester. Let me know if there are ways I can improve your learning in the class. If you are going to miss class or need an extension on an assignment, the earlier you tell me the better!

Health Protocols

- In an effort to keep our classroom community healthy, please do not come to class if you are feeling ill, have a fever, or have recently been exposed to a contagious illness.
- If you need to miss class due to illness, e-mail me directly so we can make a plan for any missed class work and/or attendance issues. Please also reach out if you have any concerns about your health needs and goals for the semester.

Academic Honesty

The University of Redlands enforces strict standards as regards academic honesty, and students may be dismissed for breaches of these standards.
In light of this, please note that:

- intentional plagiarism, i.e. piecemeal or wholesale appropriation of text from one or more printed or internet source, will result in a fail grade for the course.
- plagiarism by default, i.e. uncredited adoption of ideas from source texts due to carelessness in citation, will result in a fail grade for the project.

Artificial Intelligence

This is a class where AI is the *subject*, so let's be clear about it as a *tool*. By December you will know what a large language model is actually doing when it answers you!
Machine learning asks you to do the hard, painful, and scary work of understanding ideas. I want you to do that work, not hand it to an AI. At the same time, using AI well is part of professional practice. In this class, you're expected to use it the way a professional would: to help you debug and write basic code, not to think for you.

- **Your understanding is never outsourced.** The written discussion, interpretation, and reflection you submit must always be your own thinking, in your own words, on every assignment. That means explaining what your model shows, why you made a choice, and what it means - the part no tool can do for you.
- **AI is welcome for the mechanics of code**: debugging error messages, looking up syntax, or drafting basic/boilerplate code, the same way a professional would. Treat it like a fast, knowledgeable colleague - useful, but not a substitute for understanding what your code does.
- **Be especially careful with results.** An AI will happily write you a plausible-looking evaluation that quietly leaks test data or reports a metric that doesn't answer your question. You are responsible for every number you report. "The model wrote it" is not a defense in industry and it isn't one here.
- **Disclosure scales with how much you used.** A quick debugging question or syntax lookup needs no citation. If you copy and paste code an AI wrote for you, say so: name the tool and what it helped with in a short note on the assignment. If AI helps you understand a concept, rewrite the explanation in your own words rather than pasting it in. If you draw on a large chunk of AI-generated material, disclose that too.
- **Exams are the one AI-free zone.** The take-home, no-outside-help, individual exams exist to test your own understanding without tool assistance, so no AI use there.
- **Using an AI as an object of study is always fine.** Prompting a model to see how it fails, measuring its outputs, or fine-tuning one is coursework, not a shortcut. That is different from having it do your assignment.

Using AI to bypass this - having it write your reflections or your interpretation of results, or submitting uncredited AI work as your own - is treated as an academic honesty violation, the same as plagiarism.

If you are still in any doubt about what constitutes plagiarism or acceptable AI use, please ask me before you hand in your work!

Equity and Title IX

The University of Redlands is committed to creating a safe, inclusive, and equitable learning environment. Discrimination, harassment, retaliation, and sexual misconduct, including sexual assault, dating or domestic violence, and stalking, are not tolerated. If you or someone you know experiences these behaviors, support and reporting options are available.

Faculty are considered Responsible Employees and are required to report any disclosures of such conduct to the Office of Equity and Title IX. Reporting does not require you to file a formal complaint, but ensures you are informed of your rights and resources.

For more information, support, or to file a report, visit www.redlands.edu/titleixandequity or email the Title IX Coordinator at **titleix@redlands.edu**.

Confidential support is available through the Counseling Center (909-748-8108) or TimelyCare. Additional Resources can be found at www.redlands.edu/titleixandequity.

You can also report to local law enforcement at 909-798-7681, ext. 1. If you are ever in immediate danger, please call 911 or email/text 911@redlandspolice.org if you cannot call. To reach Public Safety on campus, call 909-748-8888 or use the
Rave Gaudian App

Accommodations

The University of Redlands is committed to providing equal access to its programs and courses in accordance with the ADA and Section 504 of the Rehabilitation Act. If you have a disability and need reasonable accommodations in this class, please contact Academic Support and Accessibility (ASA) as early as possible to begin the registration process. ASA (located on the ground floor of the Armacost Library) can be reached at 909-748-8069 or **asa@redlands.edu**. All accommodations must be approved through ASA before they can be implemented, and accommodations cannot be applied retroactively. Once your eligibility is verified, ASA will provide you with an accommodation letter for your instructors. Faculty, students, and ASA will then work together to arrange appropriate accommodations that support your learning needs without fundamentally altering essential course requirements. You are strongly encouraged to discuss your approved accommodations with your instructors early in the term. For more information, please visit the ASA webpage (https://www.redlands.edu/student-life/student-resources/accessibility). If there are ways that simple changes to the class could improve your learning please feel free to reach out to me directly.

Counseling Center

The Counseling Center provides free and confidential mental health services, including short-term individual therapy, group therapy, single-session therapy, consultations, and urgent appointments to all students with in-person or virtual options. Our Counseling Center is committed to inclusivity and to providing a supportive space for everyone. Please call 909-748-8108 to schedule an appointment or email counseling_center@redlands.edu. If a student is in crisis, please call 909-748-8960 for the 24/7 mental health crisis line. For more information on our resources, go here. Another option for individual therapy for all students is TimelyCare, which provides virtual therapy immediately (Talk Now) or up to 12 scheduled virtual therapy sessions per year. Students can choose their therapist from a list of providers for the scheduled therapy option.

Conflict Resolution Center

Experiencing a conflict? Whether it's with a friend, roommate, another member of a student organization, or faculty or staff member, conflicts happen. Learning to navigate conflicts is important to success in virtually any field, and a vital step in being a part of a community and having healthy, meaningful relationships with others. See https://sites.redlands.edu/conflict-resolution-center/student-re for more information.

The Care Team

The University CARE Team exists to help provide support and resources to students that are overwhelmed, experiencing significant distress, or possibly present some risk to themselves or others. As a faculty member, I may reach out to students about whom I am concerned to talk individually, and/or refer them to the CARE Team. If you have concerns about a fellow student, consider sharing your concern with the CARE Team via their online form. This is part of who we are as a caring, proactive community where we all look out for one another. Additionally, if you feel that you or someone else needs immediate mental health support, the University has a 24/7 mental health crisis line at 909-748-8960, and the Timely Care app, which offers on-demand emotional care. Both services connect to a live, licensed counselor.

Additional Resources

If you are in need of additional resources, please refer to the available programs below.

Book Lending Program

The Book Lending Program is an initiative to ensure the academic success of First-Generation students (students who are the first to go to college in their families who meet a particular estimated family contribution [EFC] level). Funded through alumni donations, this program provides books and other classroom materials, when needed, for First-Generation students who could not otherwise afford to purchase them. Books are returned at the end of the course, to be used by other First-Generation students the next semester. The program works alongside the Library and faculty members to ensure the availability of books and classroom materials. For more information, click the link above or contact blp@redlands.edu.

Emergency Student Loans

Student Financial Services (SFS) administers a short-term, no-interest loan fund to assist students experiencing an emergency or cash-flow problem. Except in unusual circumstances, these loans do not exceed $200 and are billed to the student's account. Evidence of repayment ability is a prerequisite for all short-term loans made to students. Students are not eligible for more than one emergency student loan per term. Contact: SFS@redlands.edu or x8047

Student Lounges

Lounges for all students to sit, work, and eat can be found here on the University website.

Student Affairs Discretionary Fund

These endowed funds in Student Affairs can be used to support student success and remove impediments that otherwise may cause the student to stop or leave school. To utilize this fund, divisional leadership should be made aware of the student in dire need of financial support. This support can be anything from personal expenses, such as utility bills, gas money, emergency trips home due to family tragedy, off-campus counseling, and other medical costs, and occasionally mental health assessment expenses. Students receive grants based on their financial need. Contact: student_affairs@redlands.edu.

Student Food Support Pantry

The Student Food Support Pantry is a resource available to all established full and part-time University of Redlands students facing food insecurities. The Pantry is located on the north side of North Hall. This space is an open, no-questions-asked space with dried and canned goods, and non-perishable items, as well as seasonal fresh produce from our sustainable farm and limited refrigerated goods. Food for this distribution is provided in partnership with Feeding America Riverside and San Bernardino. It is also funded through private donations, ASUR, and the Office of Community Service Learning. For more information, please contact SURF@redlands.edu.

Course Schedule

Schedule is subject to change as we progress through the semester. You will be notified of any changes in class. Readings are from Géron, Hands-On Machine Learning with Scikit-Learn and PyTorch.

DATA 301: Introduction to Machine Learning - Fall 2026						
DATE	**WEEK**	**DAY**	**Lecture Topic**	**Reading (Geron)**	**Homework / Exams Due**	**Notes**
09/01/2026	1	**Tuesday**	**Day 1: Course Setup, Git Workflow & the ML Landscape**	*ch. 1 (all)*	**Set up Python, JupyterLab, Git, and your course repo**	**Bring your laptop. Unit 1: Evaluation Done Right.**
09/03/2026	1	**Thursday**	**Day 2: Regression as a Lab Bench - Learning Curves, Under- & Overfitting**	*ch. 4: Polynomial Regression; Learning Curves*	**HW 1 due Sunday 9/6 11:59pm**	
09/08/2026	2	**Tuesday**	**Day 3: Bias-Variance and Regularization (Ridge, Lasso, Elastic Net)**	*ch. 4: Regularized Linear Models*		
09/10/2026	2	**Thursday**	**Day 4: Cross-Validation for Real - Leakage, Stratification, Nested CV**	*ch. 2: Create a Test Set; ch. 3: Measuring Accuracy Using Cross-Validation*	**HW 2 due Sunday 9/13 11:59pm**	
09/15/2026	3	**Tuesday**	**Day 5: Metrics Beyond Accuracy - ROC vs PR, Calibration, Thresholds**	*ch. 3: Performance Measures*		
09/17/2026	3	**Thursday**	**Day 6: Bagging and Random Forests**	*ch. 6: Voting Classifiers; Bagging and Pasting; Random Forests*	**HW 3 due Sunday 9/20 11:59pm**	**Unit 2: Ensemble Methods.**
09/22/2026	4	**Tuesday**	**Day 7: Boosting - Gradient Boosting and XGBoost**	*ch. 6: Boosting; Stacking*		
09/24/2026	4	**Thursday**	**Day 8: The Full Pipeline - Preprocessing, Search, Model Selection**	*ch. 2 (all)*	**HW 4 due Sunday 9/27 11:59pm**	
09/29/2026	5	**Tuesday**	**Catch-Up Day (in class)**	*review ch. 1-6*		**Catch up and review for Exam 1 (Units 1-2). Bring questions.**
10/01/2026	5	**Thursday**	**Day 9: From Logistic Regression to a Neuron - Tensors & Autograd**	*ch. 9: From Biological to Artificial Neurons; ch. 10: PyTorch Fundamentals*	**Exam 1 due Monday 10/5 11:59pm**	**Unit 3: Neural Networks. Exam 1 released after class today (take-home).**
10/06/2026	6	**Tuesday**	**Day 10: MLPs in PyTorch - Writing the Training Loop by Hand**	*ch. 10: Implementing a Regression MLP; Mini-Batch Gradient Descent*		
10/08/2026	6	**Thursday**	**Day 11: Backpropagation, Activations, and Initialization**	*App. A: Autodiff; ch. 11: Glorot and He Initialization; Better Activation Functions*	**HW 5 due Sunday 10/11 11:59pm**	

10/13/2026	7	**Tuesday**	*NO CLASS - Fall Break*			
10/15/2026	7	**Thursday**	**Day 12: Optimizers, Learning-Rate Schedules, Batch Normalization**	*ch. 11: Batch Normalization; Faster Optimizers; Learning Rate Scheduling*	**HW 6 due Sunday 10/18 11:59pm**	
10/20/2026	8	**Tuesday**	**Day 13: Deep-Net Regularization - Dropout, Early Stopping, Augmentation**	*ch. 11: Avoiding Overfitting Through Regularization*		
10/22/2026	8	**Thursday**	**Day 14: Convolutional Neural Networks - Convolution, Pooling, Architectures**	*ch. 12: Convolutional Layers; Pooling Layers; CNN Architectures*	**HW 7 due Sunday 10/25 11:59pm**	**Unit 4: Representations. Final project teams assigned.**
10/27/2026	9	**Tuesday**	**Day 15: Transfer Learning - Fine-Tuning a Pretrained Vision Model**	*ch. 12: Pretrained Models for Transfer Learning*		
10/29/2026	9	**Thursday**	**Day 16: Embeddings and Sequences - Why RNNs Fail at Long Range**	*ch. 13: LSTMs; GRUs; ch. 14: Embeddings*	**HW 8 due Sunday 11/1 11:59pm**	
11/03/2026	10	**Tuesday**	**Catch-Up Day (in class)**	*review ch. 9-14*		**Catch up and review for Exam 2 (Units 3-4). Final project proposals due.**
11/05/2026	10	**Thursday**	**Day 17: Attention from Scratch**	*ch. 14: Attention Mechanisms*	**Exam 2 due Monday 11/9 11:59pm**	**Unit 5: Transformers & LLMs. Exam 2 released after class today (take-home).**
11/10/2026	11	**Tuesday**	**Day 18: The Transformer Block**	*ch. 15: Attention Is All You Need (Positional Encodings; Multi-Head Attention)*		
11/12/2026	11	**Thursday**	**Day 19: Build and Train a Tiny GPT**	*ch. 15: Building an English-to-Spanish Transformer*	**HW 9 due Sunday 11/15 11:59pm**	
11/17/2026	12	**Tuesday**	**Day 20: Tokenization, Pretraining vs Fine-Tuning, and Scaling**	*ch. 14: Tokenization; ch. 15: Encoder-Only and Decoder-Only Transformers*		
11/19/2026	12	**Thursday**	**Day 21: Using LLMs - Retrieval-Augmented Generation and Evaluation**	*ch. 15: Fine-Tuning a Model Using the TRL Library; Off-the-Shelf Chatbot Models*	**HW 10 due Sunday 11/22 11:59pm**	
11/24/2026	13	**Tuesday**	**Day 22: Limitations, Ethics, and Where the Field Is Going**	*ch. 16: Multimodal Transformers*		**Final project work begins.**
11/26/2026	13	**Thursday**	*NO CLASS - Thanksgiving*			

12/01/2026	14	**Tuesday**	**Final Project Work (in class)**			**Team check-in with Dr. Bieri.**
12/03/2026	14	**Thursday**	**Final Project Work (in class)**			**Practice presentations.**
12/10/2026		**Thursday**	**FINAL EXAM - Team Project Presentations, 3:00pm**		**Report and code due Wed 12/9 11:59pm**	**Teams present a complete ML project. Final exam slot set by the registrar.**
Shading: alternating bands = regular weeks | pale yellow = catch-up/review | pale blue = project day | pink = no class | purple = finals. Readings: Geron, Hands-On Machine Learning with Scikit-Learn & PyTorch.						

Your Git Workflow

Two rituals to know. Do these and you're doing exactly what a working machine learning engineer does.

Weekly Sync: pull in new course material

Whenever new material is announced, run this in your own repo. New notes, notebooks, and assignments arrive this way all semester.

Check what's new
git fetch upstream

↓

Switch to your main
git checkout main

↓

Blend it in
git merge upstream/main

↓

Save it to your repo
git push origin main

↓

Your own work is never overwritten. It lives on branches until you merge it.

Branch → Pull Request → Merge: every assignment

Nobody pushes an assignment straight to main. **The Pull Request is how you turn work in.**

Start a branch
git checkout -b hw/week-03

↓

Commit, then push
git add .
git commit -m "..."
git push origin hw/week-03

↓

Open a Pull Request
into main, in your own repo
this is your submission

↓

I review it
my comments on the PR are your feedback

↓

Merge & delete the branch
main is your finished work

↓

Clean up locally
git checkout main
git pull origin main

Quick Reference

Once you know the two rituals above, this is all you'll actually need to glance at.

I want to...	Command
Get today's new content	git fetch upstream git merge upstream/main
Start this week's assignment	git checkout -b hw/week-NN
Save my progress	git add . git commit -m "..."
Turn it in	git push origin hw/week-NN then open a Pull Request
Clean up after grading	git checkout main git pull origin main git branch -d hw/week-NN

One-Time Setup

Do this once, right after I create your repo:

1. Clone your repo: git clone <your-repo-url>
2. Add the course materials remote:
   git remote add upstream https://github.com/MachineLearningUoR/FALL26.git
3. Install the packages:
   pip install -r requirements.txt
4. Turn on clean notebook diffs:
   nbstripout --install

Reminders

- Commit often, in small pieces, and commit your own work yourself. Your history is part of the record.
- **Never commit a trained model or a large dataset.** They bloat the repo permanently. Commit the code that produces them.
- You're encouraged to add a classmate as a reviewer on your Pull Requests. One rule: **reviewers comment, reviewers don't commit.** Reading someone else's code is how you learn to read code; writing it for them is not.