

CS 7641: Machine Learning - Fall 2026

Welcome

Instructor



Theodore J. LaGrow

theodore.lagrow@pe.gatech.edu
(<mailto:theodore.lagrow@pe.gatech.edu>)

Head TAs



John Mansfield

jmansfield6@gatech.edu
(<mailto:jmansfield6@gatech.edu>)



Jake Knigge



jwk@gatech.edu
(<mailto:jwk@gatech.edu>)



Jack Henderson

jhenderson88@gatech.edu
(<mailto:jhenderson88@gatech.edu>)

Creators of Recorded Material

	
Charles Isbell	Michael Littman
Chancellor, University of Illinois Urbana-Champaign	Associate Provost, Brown University

General Information

Machine Learning is a three-credit course on the study and application of the field of Machine Learning. Machine Learning is that area of Artificial Intelligence that is concerned with computational artifacts that modify and improve their performance through experience. The area is concerned with issues both theoretical and practical. This particular class is a part of a series of classes in Machine Learning and takes care to present algorithms and approaches in such a way that grounds them in larger systems. We will cover a variety of topics, including statistical supervised and unsupervised learning methods, optimization algorithms, and Bayesian learning methods. The course also covers theoretical concepts such as inductive bias, the PAC and Mistake-bound learning frameworks, minimum description length principle, and Ockham's Razor. Additionally, this term will have a focus on adding advanced topics to help bring the material closer to the state-of-the-art in ML practices. In order to ground these methods, the course includes programming and practical application with a number of projects.

Objectives

There are four primary objectives for the course:

- To provide foundational knowledge with a broad survey of approaches and techniques in ML
- To develop a deeper understanding of several major topics in ML
- To develop the design and programming skills that will help you to build intelligent, adaptive artifacts
- To develop necessary skills to communicate research and practicum in ML

The last bullet point is the core objective of this course and separates this ML offering from many others available online or even at GT. Without proper and efficient communication, no amount of coding or results will matter. You should develop enough background that you can pursue any desire you have to learn more about specific techniques in ML, either to pursue ML as a research career or to apply ML techniques in other research areas.

Prerequisites

The official prerequisite for this course is an introductory course in artificial intelligence. In particular, those of you with experience in general representational issues in AI, some AI programming, and at least some background in statistics, vector calculus, information theory, and linear algebra should be adequate. Any student who did well in an introductory AI course should be fine. You will note that most semi-modern AI courses suggest at least some tentative background in some machine learning techniques as well. Of course, having said all that, the most important prerequisite for enjoying and doing well in this class is your interest in the material. We say this in every semester and in every course, but it's true. In the end, it will be your own motivation to understand the material that gets you through it more than anything else. If you are not sure whether this class is for you, please contact the instructor or Head TAs.

Course Resources

- **Video Lectures.** All lectures are recorded and available in Ed Lessons.
- **Readings.** The textbook for the course is *Machine Learning* by Tom Mitchell. This is an older text, however the material is foundational towards understand the basics of ML and freely available online. With the rise of LLMs and GPTs, there is a LOT of misinformation afoot online. Make sure you understand your information sources. Be warned. We will follow the textbook closely for most of the semester as many of the pre-recorded lectures align week-to-week, so it is imperative that you have a copy of the book. We will also use supplemental readings, but those will be provided for you. Weekly supplemental readings will be updated to support advanced topics posted throughout the semester.
- **Computing.** Even though you will not need high-throughput dedicated resources, you will have

access to CoC clusters for your programming assignments. You can test your code on the Shuttles cluster, using your GT username and password to log in - you will not need a CoC account for this course. More info can be found here:

<https://support.cc.gatech.edu/facilities/general-access-servers> 

(<https://support.cc.gatech.edu/facilities/general-access-servers>). If needed, every assignment can be queued up and run on a free instance of Colab (barring some computational lag).

- **Web.** We will use the class Canvas page and Ed Discussions to post announcements. Please bookmark these pages and them early and often.

Readings and Lectures

The online lectures are meant to summarize the readings and stress the important points. You are expected to critically read any assigned material. Your active participation in the material, the lectures, and office hours are crucial in making the course successful. We will have office hours throughout the week with a recorded Friday office hour lead by the Head TAs for those who cannot attend live. We want to stress that the more you put into the material, the more you will get out. The full teaching staff is to assist you in learning and growing in the area of machine learning. Reach out and communicate often!

To help you to pace yourself, we have provided a nominal schedule (see the Syllabus link) that tells you when we would be covering material. We recommend you try to keep that pace.

Further, the instructor will post additional recordings of advanced topics pertaining to each module depending on the week. These will help and correspond to the course assignments.

Scoring and Grading

Your final grade is determined by how you do on several components: participation, a reading/writing quiz, hypothesis report practice, 4 unit quizzes, 4 comprehensive reports, 4 report discussions, and a cumulative final exam.

Reading/Writing Quiz. In an effort to kickstart the course, I have included a quiz to help with reading and writing formal reports. At a graduate level in a terminal program, it is not enough to be able to apply advanced algorithms but understand why you make specific design choice and comment on nuanced caveats on your results. This is a key aspect of the course you will not find in other online machine learning classes. The teaching staff prides ourselves in providing detailed feedback and we consistently hear that this skill can be immediately applied to many positions students currently hold. Therefore, I have included a lecture and quiz to help establish many of these skills you will practice in your reports over the course of the semester. You will have unlimited attempts to earn the points on the quiz until the due date. The lecture material and quiz will be available at the beginning of the first week and due at the end of the second week.

Hypothesis Report Practice. To help supplement the Reading/Writing Quiz, I have provided an exercise on hypothesis interpretation and development. There is a large disparity in understanding term-to-term, so this should be great practice with feedback before the first assignment is due. This will help with direct applications to the reports and mimic the report submission process.

Unit Quizzes. To help with understanding the course material in a more nuanced manner, there will be a unit quiz associated with each of the four units (SL, OL, UL, and RL). Since each unit will be three weeks total, you will have a shorter quiz due by the end of the unit to help in your understanding of the material. I will open each quiz during the first week of the unit and it will be open

until the end of the third week. There will be a variety of questions with both selection questions and calculations randomly picked from a test bank of questions. You will have a total of 60-minutes to complete the quiz with 4-attempts total over the window of the quiz. After each attempt, you will be able to see which questions you got correct and incorrect. You will be able to have one sheet of notes and a calculator, however this will be a closed-book, closed-internet quiz. There will be some practice questions posted on Ed Discussions before each quiz opens.

Reports. There will be four scored assignment reports, Supervised Learning (SL), Optimization & Uncertainty in Learning (OL), Unsupervised Learning (UL), and Reinforcement Learning (RL). They will be about programming and analysis. You will be required to use Overleaf and write your reports in LaTeX. You will use Private Georgia Tech Github repos to store your code and the Enterprise Georgia Tech Overleaf. *Use of a personal Github is prohibited and will not be scored if submitted.* Generally, these assignments are designed to give you deeper insight into the material and to prepare you for the exam. The programming will be in the service of allowing you to run and discuss experiments, do analysis, and write your reports.

We do not provide a rubric for any of the assignments. These will not be shared to avoid gamification, which will happen no matter how much you argue (we have tried both ways in many classes in OMSCS). This is by pedagogical design to challenge you as there is no rubric in academia or the workforce. At a graduate level in a terminal program, you will be asked to develop skills you may not have a lot of practice in. Many times, communication is a skill lacking by many engineers and practitioners, especially seasoned personnel. That being said, we have an extensive rubric behind the scenes while we grade. There is a vetted system to verify and calibrate grading between staff members which is iterated and monitored by the instructor. We understand this will be the first time for many of you with such a unique challenge to potential weak spots. There are many people who have succeeded and this aspect of the course continues to be a highlight for alumni, especially those seeking new jobs or attempt new projects at work. Our aim to to make this a manageable but challenging experience to help you grow as ML practitioner at the highest level.

When your reports are scored, you will receive feedback explaining your errors (and your successes!) in a fair level of detail. This feedback is for your benefit, both on this assignment and for future assignments. It is considered a part of your learning goals to internalize this feedback. This is one of many learning goals for this course, such as understanding how to analyze data or the differences between each algorithm or bias/variances in each of them.

Report Discussions. After each unit report, you will complete a short Report Discussion on Canvas. The goal is to help you reflect on your results while they are still fresh and to build better habits for the next report with your peers. You will share one figure or table you are proud of and briefly explain why, describe one surprising finding from your experimentation, confirm whether you met the goal you set in the prior report, and state one concrete change you will make next time. These discussions are complete or incomplete and are designed to be brief, practical, and useful for improving assignment-to-assignment.

Reviewer Response. In an effort to learn and grow assignment-to-assignment, we will provide a mechanism to edit and respond to your feedback. We will call this the Reviewer Response. You will have one week from the assignment grade being posted to edit and provide a two-page maximum response with both edits made and reviewer feedback. You will need to reasonably respond and edit your initial paper submission to improve your paper in good faith. Both the initial submission, revised submission, and two-page response will be needed for a proper Reviewer Response. If satisfied, you will receive half of the missed points back for the assignment. For example, if the initial grade was a 70/100, if everything is satisfied for the Reviewer Response, there will be 15 points added resulting in an 85/100. Further examples will be provided when the assignment grades are posted. Reviewer Response will only apply to SL Report, OL Report, and UL Report since the RL Report's grade and feedback will be released too close to the end of the term.

Final Exam. There will be a digital, closed-book final exam at the end of the term. The final exam will also be administered via Canvas and Honorlock this term.


Final Letter Grade. Everyone has a clear path to succeed in this course. At the end of the term, I apply a curve to the overall course grade to account for term-to-term differences where specific cutoffs are set at my discretion. Individual assignments are not curved.

Due Dates

All scored assignments are due by the time and date indicated. Here "time and date" means **Eastern Time (ET)**. Canvas does not currently support Anywhere On Earth, so this is the best alternative we can offer being at Georgia Tech. Please double check your settings and assignments for the exact due dates to mark your calendars appropriately. As a good check, you should go to settings on Canvas and set your time zone.

All assignments will be due at **11:59:00 PM ET** on the date due. However, since we will not be looking at the assignments until morning, you will have officially until 7:59:00 AM ET until the assignment is marked late. I understand that there are many circumstances that you may need an additional hour or two to complete the assignment. I will be asleep through the night and see no issue in giving the extra time. However, I need to heed a stern warning. You should use the 11:59PM timestamp as your internal deadline rather than the 7:59AM official cutoff. Staying up all night is a detriment to your mental health and may not be as conducive to constructive writing and testing. I know there is a colloquialism where nothing would get done unless for the last minute, however I do hope you all manage your time wisely. Please note the exact time for the submission as many situations may incur Murphy's Law. Allow a couple of minutes for the submission upload and check as it does take a few seconds on average to upload an assignment in Canvas.

For the reports only, we accept late assignments for a 20% off the top per-day penalty, a max of 5 days, or a 0 grade. The only exceptions to late report assignment penalties will require **a submission**

to the Dean of Students ([link to form](https://cm.maxient.com/reportingform.php?GeorgiaTech=&layout_id=5)  [https://cm.maxient.com/reportingform.php ? GeorgiaTech=&layout_id=5](https://cm.maxient.com/reportingform.php?GeorgiaTech=&layout_id=5)). You will need a **note** from the appropriate authority and **immediate notification** of the problem when it arises. If you do reach out to us directly, we will direct you to the Dean of Students.

For the quizzes, since there are multiple attempts, each quiz will be due at the exact due date.

Further, I try to maintain our preset due dates that are posted on Canvas. The only exception when due dates will be changed is a widespread, outside force preventing use of required technology (e.g. a hurricane taking out power and Wi-Fi for a week). Otherwise, I will not change due dates. If a due date needs to be changed, you will be given at minimum a 24-hour notice. Of course, anything can happen during the term and communication will be king. This will be the course policy.

As always, start early and do your best!

Numbers

Component

<u>Reading/Writing Quiz</u>	5%
<u>Hypothesis Quiz</u>	5%
<u>Unit Quizzes</u>	20%
- SL Unit Quiz	(5%)
- OL Unit Quiz	(5%)
- UL Unit Quiz	(5%)
- RL Unit Quiz	(5%)
<u>Reports</u>	45%
- SL Report	(11.25%)
- OL Report	(11.25%)

- UL Report (11.25%)

- RL Report (11.25%)

Report Discussions **5%**

- SL Report Discussion (1.25%)

- OL Report Discussion (1.25%)

- UL Report Discussion (1.25%)

- RL Report Discussion (1.25%)

Final **20%**

Extra Credit

There are several opportunities to receive extra credit in this course. We intend to provide an additional comprehensive problem set that you will be able to turn in before the Final Exam. We will provide answers but will not score the set. The problem set will help prepare you for the exam. If all the problems are attempted and turned in, we will award 1% to your overall course percentage. Remember, you need to attempt everything, including adding explanations for your answers. Everything turned in will be double-checked. These will be calculated on Canvas later in the term but before final grades are released. More details will be posted in the problem set assignment description.

Additionally, if there are significant contributions to the Ed Discussion board throughout the term, we will award 1% to your overall course percentage. "Significant contributions" will up to our discretion, however interacting with Ed Discussions will only benefit you over the course of the term. These additional points will be calculated on Canvas later in the term but before final grades are released.

Further, feedback about the course is key to help understand what is going well and how to improve when possible. The best mechanism is the end-of-term CIOS survey. If the course crosses a threshold of 70% of students completing the survey, I will award everyone in the course an additional extra credit point to the overall grade after the curve is set.


There may additional extra credit opportunities throughout the term which will be communicated on Ed Discussions.

Office Hours and Other Channels

We love the assignments in this course. As you will discover they are wonderfully open-ended, much more so than many of you will be used to. It is therefore important that in addition to watching the lectures and comprehending the required readings that you attend office hours and regularly check Ed Discussions. We will have non-recorded Office Hours Monday through Thursday with members of the TA team to help you when you have questions throughout the week. These are not required but helpful when you need to ask questions during an open hour. We will host Friday Office Hours weekly with a non-graded check-in quiz and tangential lecture (topics will vary but more focused on SOTA methods/applications). The Friday Office Hours are here for your benefit and we strongly encourage you to join live or watch the recording regularly. You should consider the information disseminated in the Friday Office Hour required.

Joining information will be shared on Ed Discussions each week.

Statement of Academic Honesty

At this point in your academic careers, we feel that it would be impolite to harp on cheating, so we won't. You are all adults and are expected to follow the university's code of academic conduct ([honor code](https://policylibrary.gatech.edu/student-life/academic-honor-code)  <https://policylibrary.gatech.edu/student-life/academic-honor-code>). Furthermore, at least some of you are researchers-in-training, and we expect that you understand proper attribution and the importance of intellectual honesty.

We should also point out that "proper attribution" does not absolve the writer of the "intellectual honesty" that comes from original writing. While it is definitely the case that copying text without attribution is considered plagiarism, it is also the case that copying too much text even with attribution betrays a lack of intellectual honesty. Too many quotes of more than, say, two sentences will be considered plagiarism and a terminal lack of academic originality. Do not overthink this issue, but do not under think it either.

Please note that unauthorized use of any previous semester course materials, such as tests, quizzes, homework, projects, videos, and any other coursework, is prohibited in this course. In particular, you are not allowed to use old exams. Using these materials will be considered a direct violation of academic policy and will be dealt with according to the GT Academic Honor Code. Furthermore, we do not allow copies of my exams outside of this course. Just as you are not to use the previous material you are not to share current material with others either now or in the future. Our policy on that is strict. If you violate the policy in any shape, form, or fashion you will be dealt with according to the GT Academic Honor Code.

Due to the size and online nature of this course, there can sometimes be students who promote behavior and language that falls outside the Student Code of Conduct, especially after assignment grades are posted. Any personal attacks or unacceptable use of language towards other students or staff on any online platform will be reported with a zero-tolerance policy. When in doubt, follow the golden rule.

Updates concerning the use of Large Language Models (LLMs) and Generative AI

We treat AI based assistance the same way we treat collaboration with people. You may discuss ideas and seek help from classmates, colleagues, and AI tools, but all submitted work must be your own. For this course a large language model is any model with more than one billion parameters. These tools can increase productivity, but they are aids, not substitutes for your skills. The goal of reports is synthesis of analysis, not merely getting an algorithm to run. Even if you locate code elsewhere or generate code with an AI tool, you must apply it thoughtfully to your data and write the analysis yourself.

Every submission must include an AI Use Statement in an appendix. List the tools used and what they assisted with and confirm that you reviewed and understood all assisted content. If an editor or platform surfaces suggestions, you may keep them only with disclosure and verification. You are not required to disable assistants. Be aware that coding editors often include assistants by default. Google Colab and Visual Studio Code commonly surface them. Also check any writing or grammar platforms you use. Overleaf through Georgia Tech is acceptable.

Allowed with disclosure. Brainstorming, outlining, grammar and clarity edits, code generation, code refactoring, and debugging.

Not allowed. Submitting AI written analysis, conclusions, or figures as your own. Fabricating results or citations. Paraphrasing AI or prior work to evade checks. Uploading course materials or private data to public tools.

Example Statement at the very end of the report before References. "AI Use Statement. I used ChatGPT and Visual Studio Code Copilot to brainstorm and outline sections of the report, generate and refactor small code snippets, debug a pandas indexing issue, and edit grammar and clarity throughout. I reviewed, verified, and understand all assisted content."


Verification will rely on provenance and reproducibility. We will not ask for live coding. We will review platform histories, including Overleaf revision timelines and copy and paste patterns, and Git commit and branch history. We may request your repository with full history, Overleaf project history, notebooks, logs, and intermediate outputs, and we may run reproducibility checks. Provide run instructions, environment details, random seeds, and the scripts used to generate figures and tables. Similarity and AI indicators, including Turnitin, are corroborating signals only and are not used in isolation.

If we suspect a violation, we will document evidence, notify you, and request artifacts. Cases may be referred to the Office of Student Integrity (OSI). Undisclosed or prohibited use is academic misconduct. This policy may be updated during the term. When in doubt, always use proper citations, and ask clarifying questions on Ed.

Diversity, Equity, and Inclusion

We commit to creating an inclusive, diverse, and equitable learning environment. Regardless of personal history, background, or identity category, every student is a valued member of the Georgia Tech community. Your experiences and ideas are valuable and essential, and you are encouraged to share them freely and respectfully. If you find aspects of the course instruction, subject matter, or online classroom environment that result in barriers to your inclusion, please reach out to me personally without fear of reprisal.

Disability Accommodation

Georgia Tech has established policies with respect to disability accommodation through the Office of Disability Services (ODS). These policies may be accessed at their website located at <https://disabilityservices.gatech.edu/>  (<https://disabilityservices.gatech.edu/>). Students seeking disability accommodation are specifically referred to the guidance and materials available from ODS. To receive accommodation, students must comply with the requirements set forth by that Office.

If you have any questions, please reach out to Jack (jhenderson88@gatech.edu (<mailto:jhenderson88@gatech.edu>)) with the instructor tagged in in the email (theodore.lagrow@pe.gatech.edu (<mailto:theodore.lagrow@pe.gatech.edu>)).

Disclaimer



I reserve the right to modify any of these plans as need be during the course of the class; however, we won't do anything capriciously, anything we do change won't be too drastic, and you'll be informed as far in advance as possible. There are many, many outside factors you will not be privy to, so please do not try and make assumptions. Looking forward to an amazing learning journey together!

Reading List

Required Text

- [Tom Mitchell, Machine Learning. McGraw-Hill, 1997.](http://www.cs.cmu.edu/afs/cs.cmu.edu/user/mitchell/ftp/mlbook.html) 
(<http://www.cs.cmu.edu/afs/cs.cmu.edu/user/mitchell/ftp/mlbook.html>)

Optional Text

- [Larry Wasserman, All of Statistics. Springer, 2010](http://www.stat.cmu.edu/~larry/all-of-statistics/)  (<http://www.stat.cmu.edu/~larry/all-of-statistics/>) (Read Part 1 for an intro to Probability Theory)
- Richard Sutton and Andrew Barto, Reinforcement Learning: An introduction. (for Reinforcement Learning) ([Nov 5, 2017 version](http://incompleteideas.net/book/bookdraft2017nov5.pdf)  (<http://incompleteideas.net/book/bookdraft2017nov5.pdf>.)

Prerequisites/ Refreshers

- [Linear Algebra and Eigenproblems](https://github.com/pushkar/4641/raw/master/downloads/Eigenproblems.fm.pdf) 
(<https://github.com/pushkar/4641/raw/master/downloads/Eigenproblems.fm.pdf>)

Supervised Learning (SL1–SL8)

SL1: Decision Trees

- Mitchell Ch 1 & Ch 3

SL2: Regression & Classification

- Mitchell Ch 4

SL3: Neural Networks

- Mitchell Ch 4


SL4: Instance-Based Learning

- Mitchell Ch 8

SL5: Ensemble Learning

- [Schapire's Introduction](https://github.com/pushkar/4641/raw/master/downloads/boosting.ps)  (<https://github.com/pushkar/4641/raw/master/downloads/boosting.ps>)
- [Jiri Matas and Jan Sochman's Slides](https://github.com/pushkar/4641/raw/master/downloads/adaboost_matas.pdf) 
(https://github.com/pushkar/4641/raw/master/downloads/adaboost_matas.pdf)

SL6: Kernel Methods and SVMs

- [An introduction to SVMs for data mining](https://gatech.instructure.com/courses/499842/files/folder/Supplemental%20Readings/Week%204?preview=68842187)
(<https://gatech.instructure.com/courses/499842/files/folder/Supplemental%20Readings/Week%204?preview=68842187>)
- [Christopher Burges tutorial on SVMs for pattern recognition](https://www.microsoft.com/en-us/research/wp-content/uploads/2016/02/svmtutorial.pdf) 
(<https://www.microsoft.com/en-us/research/wp-content/uploads/2016/02/svmtutorial.pdf>)

- [Scholkopf's NIPS tutorial slides on SVMs and kernel methods](https://github.com/pushkar/4641/raw/master/downloads/svm-scholkopf.ps)  (<https://github.com/pushkar/4641/raw/master/downloads/svm-scholkopf.ps>)

SL7: Computational Learning Theory

- Mitchell Ch 7

SL8: VC Dimensions

- [High-Dimensional Probability by Vershynin \(Chapter 8.3\)](https://gatech.instructure.com/courses/499842/files/68842343?verifier=BeuZww1sXKZ6J08CaXKYost3T3tcbqG7RrTbhzim&wrap=1) (<https://gatech.instructure.com/courses/499842/files/68842343?verifier=BeuZww1sXKZ6J08CaXKYost3T3tcbqG7RrTbhzim&wrap=1>)

Optimization & Uncertainty in Learning (OL1–OL5)




OL1: Randomized Optimization

- Mitchell Ch 9

OL2: Deconstructing AdamW

- Loshchilov & Hutter ICLR 2019
- Kingma Ba ICLR 2015
- From SGD to AdamW LaGrow 2025
- Linear Algebra Refresher and Context OL 2 Deconstructing AdamW LaGrow (Optional)

OL3: Information Theory

- [No Free Lunch Theorem](https://ml-cs7641.s3.us-east-1.amazonaws.com/nfl-optimization-explanation.pdf)  (<https://ml-cs7641.s3.us-east-1.amazonaws.com/nfl-optimization-explanation.pdf>)
- [Charles Isbell's Note on Information Theory](https://www.cc.gatech.edu/~isbell/tutorials/InfoTheory.fm.pdf)  (<https://www.cc.gatech.edu/~isbell/tutorials/InfoTheory.fm.pdf>)
- [An Introduction to Information Theory and Entropy](https://github.com/pushkar/4641/raw/master/downloads/gentle_intro_to_information_theory.pdf)  (https://github.com/pushkar/4641/raw/master/downloads/gentle_intro_to_information_theory.pdf)

OL4: Bayesian Learning

- Mitchell Ch 6

OL5: Bayesian Inference

- Mitchell Ch 6

Unsupervised Learning (UL1–UL4)

UL1: Clustering

- [Intuitive Explanation of EM](https://ml-cs7641.s3.us-east-1.amazonaws.com/em-intuitive-explanation.pdf)  (<https://ml-cs7641.s3.us-east-1.amazonaws.com/em-intuitive-explanation.pdf>)

- [Statal View of EM](https://github.com/pushkar/4641/raw/master/downloads/em.pdf) ↗ (<https://github.com/pushkar/4641/raw/master/downloads/em.pdf>)
- [Jon Kleinberg's Impossibility Theorem for Clustering](https://www.cs.cornell.edu/home/kleinber/nips15.pdf) ↗ (<https://www.cs.cornell.edu/home/kleinber/nips15.pdf>)

UL2: Feature Selection

- (Optional Primer: Bishop 12.1 and 12.4)

UL3: Feature Transformation

- Bishop Ch 12
- [ICA: Algorithms and Applications](https://ml-cs7641.s3.us-east-1.amazonaws.com/ica-algorithms-and-applications.pdf) ↗ (<https://ml-cs7641.s3.us-east-1.amazonaws.com/ica-algorithms-and-applications.pdf>)
- [Restructuring High Dimensional Data by Charles and Paul Viola](https://www.cc.gatech.edu/~isbell/papers/isbell-ica-nips-1999.pdf) ↗ (<https://www.cc.gatech.edu/~isbell/papers/isbell-ica-nips-1999.pdf>)

UL4: Manifold Learning

- [No Straight Lines Here](https://sites.gatech.edu/omscs7641/2024/03/10/no-straight-lines-here-the-wacky-world-of-non-linear-manifold-learning/) ↗ (<https://sites.gatech.edu/omscs7641/2024/03/10/no-straight-lines-here-the-wacky-world-of-non-linear-manifold-learning/>) (Blog by Aviral and LaGrow)
- Visualizing Data with t-SNE by L.J.P van der Maaten and Geoffrey E. Hinton
- Troy Whitfield Dimensionality Reduction (slides)
- (Extensive math, optional) Burges 2009 posted to Canvas

Reinforcement Learning (RL1–RL5)

RL1: Markov Decision Processes

- Mitchell Ch 13

RL2: Reinforcement Learning

- [Richard Sutton and Andrew Barto, Reinforcement Learning: An introduction. MIT Press, 1998.](http://incompleteideas.net/book/bookdraft2017nov5.pdf) ↗ (<http://incompleteideas.net/book/bookdraft2017nov5.pdf>)
 - This is a full textbook, however chapters 2, 4, and 6 are helpful but not required.
- [Reinforcement Learning: A Survey](https://github.com/pushkar/4641/raw/master/downloads/kaelbling96reinforcement.pdf) ↗ (<https://github.com/pushkar/4641/raw/master/downloads/kaelbling96reinforcement.pdf>)

RL3: Game Theory

- [Andrew Moore's slides](http://www.cs.cmu.edu/~awm/tutorials.html) ↗ (<http://www.cs.cmu.edu/~awm/tutorials.html>)

RL4: Game Theory Continued

- [Andrew Moore's slides](http://www.cs.cmu.edu/~awm/tutorials.html) ↗ (<http://www.cs.cmu.edu/~awm/tutorials.html>)

RL5: Horizon of Variants

- Rainbow DQN Silver

Software

Core ML (Python)

- **scikit-learn** | the course workhorse for SL/UL; pipelines, CV, calibration, metrics.
- **imbalanced-learn** | resampling and class-weight tools for imbalanced targets.
- **XGBoost / LightGBM / CatBoost** | gradient-boosted trees for strong tabular baselines.
- **scikit-learn-extra** | extra algorithms (k-medoids, etc.) that sometimes help.

Deep Learning (for NNs)

- **PyTorch** (+ **Lightning** if you want cleaner training loops) | recommended for MLPs on tabular data.
- **Keras / TensorFlow 2** | solid, high-level API for quick NN prototypes.
- **JAX / Flax** | only if you need research-grade autodiff; not required for this course.

Optimization, Tuning, and Experiment Tracking

- **Optuna** or **Ray Tune** | hyperparameter search with pruning; integrates with sklearn/PyTorch.
- **scikit-optimize** | simple Bayesian optimization for small searches.
- **MLflow** or **Weights & Biases (W&B)** | track runs, metrics, parameters, artifacts.

Interpretability, Diagnostics, and Calibration

- **SHAP** (tabular feature attributions), **LIME** (local explanations), **ELI5** (quick checks).
- **scikit-learn calibration** or **netcal** | reliability curves, probability calibration.
- **fairlearn** | fairness metrics/mitigations (use when appropriate).

Time Series (if you venture there)

- **statsmodels**, **pmdarima**, **darts** (batteries-included forecasting), **prophet** (additive models).
(Not required for the main assignments; listed for completeness.)

Reinforcement Learning & MDPs

- **Gymnasium** (the maintained fork of Gym) | environments.
- **Stable-Baselines3** | reliable RL baselines (DQN, PPO, A2C, etc.).
- **RLlib** | scalable RL on Ray (only if you need distributed runs).
- **MDPtoolbox (Python port)** | value/policy iteration for finite MDPs.
- **Murphy's MDP Toolbox (MATLAB)** | classic finite-MDP routines.

Java / MATLAB Ecosystem (if you prefer those stacks)

- **WEKA** | classic Java ML suite (good for quick comparisons).
- **SMILE (Java/Scala)** | modern ML for the JVM.
- **Deeplearning4j** | DL on the JVM (only if you must stay in Java).
- **ABAGAIL (Java)** | still available, but community momentum is limited.
- **MATLAB**: Statistics & Machine Learning Toolbox; **Dellaert's Clustering** (legacy but useful); MDP toolbox as above.

Data & Visualization

- **pandas / polars** | tabular data wrangling (polars is fast, pandas is ubiquitous).
- **NumPy** | arrays, linear algebra.
- **matplotlib / plotly / altair** | plotting; use plotly/altair when you need interactive PR/ROC/residuals.

Environments & Notebooks

- **Anaconda/conda-forge** or **mamba/uv** | reproducible environments.
- **JupyterLab / VS Code / Colab** | notebooks and IDEs.