CS 4713/5713 - Computational Learning Theory

Dimitris Diochnos
The University of Oklahoma,
School of Computer Science

Fall 2025

Time and Location

The course will take place at Carson Engineering Center 438. Tuesdays & Thursdays, 10:30am-11:45am.

Canvas: Canvas will be used in order to distribute homework assignments and potentially other reading material.

Personal Website: A website will also be used to keep track of our progress, links to additional resources, timetables for deliverables, and other useful information at:

https://www.diochnos.com/teaching/CS4713-5713/2025F/index.php.

Having said that, the most accurate deadlines for delivering your homework assignments and project deliverables will be provided through Canvas, since that is the submission platform for what you turn in.

Instructor

Dimitris Diochnos, 230 Devon Energy Hall (DEH), diochnos@ou.edu .

Teaching Assistants

No teaching assistants are available for this course.

Office Hours

Office hours will be held:

- on **Tuesdays** between **3:00pm-3:50pm**,
- on **Thursdays** between **3:00pm-3:50pm**,
- as well as on **Fridays** between **1:30pm-2:15pm**, or
- **by appointment**.

Please note that while anyone is welcome during the office hours, students from CS 4713-5713 will have precedence on Fridays, while students from CS 3823 will have precedence on Tuesdays and Thursdays.

Prerequisite Background

Design and analysis of algorithms, basic computational complexity theory, mathematical maturity. Tools from probability theory will be discussed (briefly) on demand as they arise.
Prerequisites: CS 4413 or DSA 4413. You can also enroll with the permission of the instructor.

Course Catalog Description

Learning using membership queries, equivalence queries, version spaces, decision trees, linear models. Probably approximately correct (PAC) learning, VC-theory, distribution-independent learning. Representation issues and intractability. Noise models, statistical queries, PAC learning under noise, poisoning attacks, adversarial examples. Distribution-specific learning and evolvability. Online learning and mistake bounds. Weak and strong learning (boosting). No student may receive credit for 4713 and 5713.

Learning Outcomes

We have the following learning outcomes in this course.

Core Concepts and Models.

- **Formal Learning Models:** Students learn about formal models of learning, such as PAC (Probably Approximately Correct) learning, which provides a framework for rigorous analysis of learning problems.
- **Computational Efficiency:** A central goal is to understand and prove the computational efficiency of learning algorithms, determining when efficient algorithms are likely to exist and when learning problems are inherently difficult.

Theoretical Analysis and Guarantees.

- **Error Bounds:** Students learn to construct and analyze probabilistic bounds on the excess risk of learning estimators to provide guarantees on generalization error.
- **Key Theoretical Tools:** Students become proficient in using tools like VC-dimension, and various concentration and tail inequalities (like Chernoff's and Hoeffding's inequalities) to analyze learning algorithms.
- **Sample Complexity:** The course examines how to determine the necessary sample sizes for learning under different scenarios.

Algorithm Design and Analysis.

- **Algorithm Analysis:** Students analyze the performance of algorithms for various learning tasks focusing on convergence rates as the dimensionality of the instances increases.
- **Robustness and Noise:** The theory explores learning with noise, poisoning attacks, and adversarial examples, leading to a deeper understanding of robust learning algorithms.

Broader Skills.

- **Research-Oriented Skills:** The course equips students with the tools to conduct research in theoretical machine learning and to contribute to understanding new research frontiers.
- **Rigorous Thinking:** Students develop the ability to think about machine learning from a rigorous, principled, and mathematically sound perspective.

Schedule of Classes

The syllabus is continuously updated and subject to change. We will cover the material at a pace that is comforable. Our **first meeting** is on **Tuesday, Auguest 26, 2025** and our **last meeting** is on **Thursday, December 11, 2025**.

A **rough outline** for the course, which is subject to change slightly depending on our pace, is shown in Table 1.

Table 1: Tentative Course Schedule

Period	Topics
Week 1	Syllabus, expectations, mathematical background
Week 2	Concept learning using queries. Introduction to version spaces.
Week 3	Conclusion of version spaces. Introduction to decision trees.
Week 4	Conclusion of decision trees. Linear models for classification.
Week 5	Introduction to PAC learning.
Week 6	Learning in the realizable case using a finite hypothesis space. Improving explainability by learning hypotheses with few relevant variables.
Week 7	Reviewing material and preparation for the midterm. Midterm.
Week 8	Agnostic learning using finite hypotheses spaces and empirical risk minimization. Intractability of learning 3-term DNF formulae.
Week 9	VC-theory. Upper bounds and lower bounds for distribution-independent learning. Noise models. Malicious noise, induced distributions.
Week 10	Random misclassification noise. Hardness of learning conjunctions under random misclassification noise.
Week 11	PAC learning under class imbalance. Statistical queries.
Week 12	Poisoning attacks. Evasion attacks (adversarial examples). Other topics on trustworthy machine learning.
Week 13	Distribution-specific learning, evolvability.
Week 14	Online learning. Halving, randomized halving, weighted majority algorithm, randomized weighted majority algorithm, winnow.
Week 15	Weak and strong learning (boosting).
Week 16	Student presentations.

The **midterm exam** is **in-class** and will take place near the 11th week of the classes.

The university has scheduled the final exam for Thursday, December 18, 2025 between 8:00am and 10:00am. However, there will be **no final exam** for this class. Instead, we will have a **semester-long project** with presentations occuring during the last week of classes (normally, during the last Friday of the semester).

Textbook, Notes and Related Reading Material

The class will rely to a large extent on papers as well as on handouts by the instructor. A very good book for introduction to machine learning that we will be using for some of the topics that we plan to cover in this class, is the book *Machine Learning*, by Tom Mitchell [4]. The book is available online for free by the author, at:

- https://www.cs.cmu.edu/~tom/mlbook.html .

Having said that, we plan to cover material that is not available in the above book, and is instead covered by the more recent book **Understanding Machine Learning**, by Shai Shalev-Schwartz and Shai Ben-David [8], or the book **Foundations of Machine Learning**, by Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar [5]. Both of these books are available online for free by the authors at the following addresses respectively:

- http://www.cs.huji.ac.il/~shais/UnderstandingMachineLearning/ and
- https://cs.nyu.edu/~mohri/mlbook/ .

Furthermore, a classic resource for covering aspects relevant to the PAC model of learning is [3] and we will blend discussions from this book together with the more recent treatment found in [8] or [5].

It is advised that the students take notes from the material that is covered in class.

In order to dive deeper into topics related to trustworthy supervised learning we will be using references such as:

- **Interpretable Machine Learning** [6],
- **Trustworthy Machine Learning** [10],
- **Fairness and Machine Learning: Limitations and Opportunities** [1].

Other Books of Interest. Domingos in *The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World* [2], has a very nice (high-level) description, sometimes accompanied by historical anecdotes, on different aspects of machine learning. Valiant's book *Probably Approximately Correct: Nature's Algorithms for Learning and Prospering in a Complex World* [9] provides a good narrative for connections that we discuss between evolution and learning. Finally, an unfortunate side of the ever-increasing influence that machine learning algorithms have in our daily lives is discussed in *Weapons of math destruction: How big data increases inequality and threatens democracy*, by Cathy O' Neil [7]. There are many more books along these lines and will be mentioned in some preliminary slides in class. These issues motivate the urgency of developing machine learning mechanisms that are **trustworthy**.

Project

Students should form groups of 2 and work on the semester-long project. No group can be of 1 person, or of more than 2 persons. The work of the students will be presented near the end of the course and will also have a written component. Projects can be **implementation-based** or **theory-based**. In either case, the students will have to present what they read, and potentially the findings of an implementation, at the end of the semester. The **students enrolled in CS 5713** will have to do a more extensive literature review by citing and explaining **two more related papers per person** in the group compared to the load that undergraduate students have.

Please see the document **Project Description** for more information about the write-ups and the deadlines.

Examinations

The midterm exam will be a closed-book written exam in the class where we have our regular meetings.

There is no final exam for this class. Your final written report as well as your presentation serve this purpose.

Homework Assignments

Undergraduate students will have five homework assignments whereas graduate students will have six homework assignments. There will be 5-6 homework assignments; most likely six, with a tentative schedule as shown in Table 2.

Table 2: Tentative Homework Schedule

Homework	Announced	Due	Assigned To
1	Week 1 (end)	Week 3 (beginning)	All students
2	Week 3 (beginning)	Week 4 (end)	All students
3	Week 4 (end)	Week 6 (beginning)	All students
4	Week 8 (beginning)	Week 10 (beginning)	All students
5	Week 10 (beginning)	Week 12 (beginning)	All students
6	Week 12 (end)	Week 15 (end)	Graduate students only

The contribution for your grade based on homework will be computed by adding up all the points that you receive from the individual homework assignments and then dividing by the maximum amount of points that you could gather from all these assignments. I expect the assignments to be weighted roughly evenly, except apart from the first one (so that we can start working on homework problems earlier rather than later).

Grading

Grading will be based on the following:

- **40% homework assignments**,
- **40% semester-long project**,
- **20% in-class midterm exam**.

Grades may also be adjusted slightly upward or downward depending on class participation. I expect grading to be along the lines shown in the table below:

Percentage	Grade
$\geq 90\%$	A
$\geq 80\%$	B
$\geq 70\%$	C
$\geq 60\%$	D
otherwise	F

Grades may be curved at the end of the semester, but grade thresholds will never be higher than those shown above.

Additional Workload for Students Enrolled in CS 5713

The additional workload for students enrolled in CS 5713 has two components. The first component is an additional individual homework assignment near the end of the semester, which will allow these students to understand how research is being performed in machine learning theory. This is an effort of more than 20% on the homework assignments compared to students enrolled in CS 4713, which corresponds to more than 8% of the total final grade. Furthermore, the students enrolled in CS 5713 need to perform a more extensive review of the related work in their write-ups for the semester-long project and discuss such broader connections with the main paper that they are presenting during the final presentations. This combined effort of increased literature review and discussion during the final presentation corresponds to approximately 10% increase of load compared to students enrolled in CS 4713 for their projects, thus, overall, this second component corresponds to an additional 4% effort that students enrolled in CS 5713 need to accomplish for the same final grade as those enrolled in CS 4713.

As a summary, students enrolled in CS 5713 need to put about 12%-13% more effort compared to students enrolled in CS 4713, in order to achieve the same numerical (weighted) grade.

Course Policies

Collaboration Policy

Regarding homework assignments, unless otherwise specified, students may discuss problem sets with one another. However, students should afterward write the solutions on their own. Collaborators (people you speak to about an assignment) must be named at the top of the assignment (together with their university IDs). No collaboration will be allowed on exams. Students **must** form **groups of two (2) people** and work together on the semester-long project. You cannot work alone. You cannot work in a group of three (3) or more people.

General Remarks. Please note the following two.

- **If you are unsure if something is permitted, consult with me before doing it.**
- **For exams** (whether midterms or final), **students are required to work alone** and follow the stated rules exactly.

Project code. Your project code and writeups must be written exclusively by you or your group. **Use of any downloaded code or code taken from a book (whether documented or undocumented) is considered academic misconduct and will be treated as such.** Exceptions from this policy (such as a project that builds on an existing open-source project) may be granted but you MUST speak with me first. You are also **not allowed to use Generative AI (e.g., ChatGPT, etc.)** for your code or text.

Late Work Policy

You can postpone once your homework or project deliverable (checkpoint, or final write-up and code) submission by 24 hours without any penalty. After the first time that you have a late submission, a 10% (of the maximum possible grade) penalty will be applied for every day that is late – the maximum delay can be 3 days (including the first time that you have a late submission). This penalty is applied of course to every member of the group where you belong to.

We will be using an electronic system (Canvas) for the students' submissions and therefore it is your responsibility to turn in your homework (or work for the semester-long project, should this be

the case) on time. **Please coordinate within your group and make one submission per group if you are submitting a semester-long project checkpoint or the final write-up for the semester-long project.**

Chegg and Other Online Tutoring Sources

There are a wide variety of tutoring resources available through paid websites. Many of these sites have students upload assignments and solutions and surreptitiously provide these documents to other students. What appears to be a session with a tutor may be, behind the scenes, the tutor doing a search of their company database of solutions to share. By using these sites you risk being charged with academic misconduct, either by supplying other students with answers they did not author or by receiving someone else's answer that you did not author. Since these companies are not open with students about their practices, you cannot know whether a tutor is providing meaningful support (for example, identifying misunderstandings of content and explaining them) or simply feeding you someone else's solution a bit at a time. The tutor's actions can result in different students submitting answers that are identical, which may be flagged as academic misconduct during grading.

Use of Generative Artificial Intelligence

Providing solutions generated by a computer program instead of being generated from a student's mind does not demonstrate student learning. Students who learn to rely on generative AI technology may not be able to argue about computer science topics independently, as is necessary for examinations in this course and for their future careers. Amongst other problems, generative AI tools are not fact checkers and do not necessarily produce correct solutions to problems. They are sensitive to small changes in prompts. They do not quote and cite sources properly. Students who copy generated AI results into assignments are committing plagiarism, just as if they had copied from another student.

- Use of generative AI tools is only allowed when the assignment specifically permits it. **If the assignment is silent on the use of generative AI tools, they are not allowed.**
- If a generative AI tool is used in an assignment, the student must provide a summary of how the tool was used. This would include the name of the tool (including version) and a detailed description of how it was used and what work the student contributed to the assignment. Students must substantially contribute to the solution to avoid plagiarism.
- When AI tools are allowed on an assignment, students may request an alternate assignment if they have an objection to using these tools.
- If there is a question about whether a student completed work independently or with the use of these tools, I will invite the student to my office to explain the assignment in detail. **Students who choose to not come or who cannot explain the work they submitted will be charged with an academic integrity violation.**

Make-Up Midterms

In some rare cases I can offer a makeup midterm to a student (subject to my schedule and availability as well). However, if the student misses their rescheduled midterm, the percentage points of the midterm as a contribution towards their final grade will transfer to their final exam. For example, if in a class each of two midterms contributes 15% to the total final grade and the final exam contributes 25% of the total final grade, then missing one midterm would cause the final exam to contribute 40% towards the final (overall) grade that the student will receive in the class.

Furthermore, for every midterm that a student misses beyond one (e.g., a student misses both makeup midterms), then for every such midterm apart from one, the student will receive a zero.

Missing a midterm without previously informing me about this (and making arrangements ahead of time for a make-up midterm) will result in receiving a zero (0) for that particular midterm. Please read the following section for further clarification on the matter.

Not Showing Up on Midterms

Not showing up on a midterm without a reasonable justification that can explain why you have not arranged ahead of time a make-up midterm, will immediately lead to a ZERO (0) for your grade.

Claiming that you were not aware that we had a midterm on a particular day and you scheduled another appointment (e.g., doctor's appointment), or you just decided not to come to class simply because you did not want to come to class one day, is not a valid excuse. It is your responsibility to make sure that you receive all the notifications from Canvas. Additionally, a rough expectation regarding when the midterms will take place is announced in this syllabus in the tentative schedule of classes (Table 1) as well as on the webpage that we are maintaining for the course. Again, the latest and greatest information for the exam will be communicated via Canvas and also discussed in class. So, there are no cheap excuses for no-shows on the day a midterm is scheduled to take place. I will not even follow-up with an email asking you what happened. You will automatically be assigned a zero (0).

If you want, you are free to explain your case within 48 hours as to why you did not show up in class for the midterm and provide some valid justification for some extreme event that happened to you. However, I will not initiate this communication asking you why you did not show up in class for the midterm. It is your responsibility to do so.

Classroom Conduct

Disruptions of class will not be permitted. Examples of disruptive behavior include:

- Allowing a cell phone or pager to repeatedly beep audibly.
- Playing music or computer games during class in such a way that they are visible or audible to other class members.
- Exhibiting erratic or irrational behavior.
- Behavior that distracts the class from the subject matter or discussion.
- Making physical or verbal threats to a faculty member, teaching assistant, or class member.
- Refusal to comply with faculty direction.

In the case of disruptive behavior, I may ask that you leave the classroom and may charge you with a violation of the Student Code of Responsibilities and Conduct.

Class Web Page

The main web page for the class is

https://www.diochnos.com/teaching/CS4713-5713/2025F/index.php

Login to the Canvas website using your 4+4 (first four letters of your last name followed by the last four digits of your student number), using your standard OU password. If you have difficulty logging in, call 325-HELP. This software provides a number of useful features, including a list of assignments and announcements, an electronic mailing list, newsgroups, and grade book. All handouts are available from Canvas. You should check the site daily. When I update the site, I will post an announcement telling you what has been added and where it is located. You are responsible for things posted on the site with a 24 hour delay.

Student's Feedback for the Course

The College of Engineering utilizes students' feedback as one of the bases for evaluating the teaching effectiveness of each of its faculty members. The results of these forms are important data used in the process of awarding tenure, making promotions, and giving salary increases. In addition, the faculty uses these forms to improve their own teaching effectiveness. The original request for the use of these forms came from students, and it is students who eventually benefit most from their use. Please take this task seriously and respond as honestly and precisely as possible, both to the machine-scored items and to the open-ended questions.

Class Email Alias

Urgent announcements will be sent through email. It is your responsibility to:

- Have your university supplied email account properly forwarded to the location where you read email.
- Make sure that your email address in Canvas is correct, and forwards email to the place where you read it. I'll send out a test message during the first week of class. If you do not receive this message, it is your responsibility to get the problem resolved immediately.
- Have your email program set up properly so that replying to your email will work correctly the first time. You can send email to yourself and reply to yourself to test this.

If you need assistance in accomplishing any of these tasks, contact 325-HELP.

Newsgroups and Email

The newsgroup on Canvas should be the primary method of communication, outside of class. This allows everyone in the class to benefit from the answer to your question. If you email me a question of general interest, I may post your question and my answer to the newsgroup. Matters of personal interest should be directed to email instead of to the newsgroup, e.g. informing me of an extended personal illness. Posting guidelines for the newsgroup are available on Canvas.

Incompletes

The grade of I is intended for the rare circumstance when a student who has been successful in a class has an unexpected event occur shortly before the end of the class. I will not consider giving a student a grade of I unless the following three conditions have been met.

1. It is within two weeks of the end of the semester.
2. The student has a grade of C or better in the class.
3. The reason that the student cannot complete the class is properly documented and compelling.

Add/Drop/Withdrawal Deadlines. Please consult the OU academic calendar (as well as the policies of the School of Engineering) for the following deadlines:

- **Add a course**
- **Drop a course without penalty (course removed from transcript)**
- **Drop a course with a W on transcript**

University Policies

The instructor reserves the right to add, remove, or change any element of class policy at any time and for any reason, within the limits of University policy.

OU is committed to creating a learning environment that meets the needs of its diverse student body. If you anticipate or experience any barriers to learning in this course, please feel welcome to discuss your concerns with me.

Mental Health Support Services

Support is available for any student experiencing mental health issues that are impacting their academic success. Students can either been seen at the University Counseling Center (UCC) located on the second floor of Goddard Health Center or receive 24/7/365 crisis support from a licensed mental health provider through TimelyCare. To schedule an appointment or receive more information about mental health resources at OU please call the UCC at 405-325-2911 or visit University Counseling Center. The UCC is located at 620 Elm Ave., Room 201, Norman, OK 73019.

Title IX Resources and Reporting Requirement

The University of Oklahoma faculty are committed to creating a safe learning environment for all members of our community, free from gender and sex-based discrimination, including sexual harassment, domestic and dating violence, sexual assault, and stalking, in accordance with Title IX. There are resources available to those impacted, including: speaking with someone confidentially about your options, medical attention, counseling, reporting, academic support, and safety plans. If you have (or someone you know has) experienced any form of sex or gender-based discrimination or violence and wish to speak with someone confidentially, please contact OU Advocates (available 24/7 at 405-615-0013) or University Counseling Center (M-F 8 a.m. to 5 p.m. at 405-325-2911).

Because the University of Oklahoma is committed to the safety of you and other students, and because of our Title IX obligations, I, as well as other faculty, Graduate Assistants, and Teaching Assistants, are mandatory reporters. This means that we are obligated to report gender-based violence that has been disclosed to us to the Institutional Equity Office. This means that we are obligated to report gender-based violence that has been disclosed to us to the Institutional Equity Office. This includes disclosures that occur in: class discussion, writing assignments, discussion boards, emails and during Student/Office Hours. You may also choose to report directly to the Institutional Equity Office. After a report is filed, the Title IX Coordinator will reach out to provide resources, support, and information and the reported information will remain private. For more information regarding the University's Title IX Grievance procedures, reporting, or support measures, please visit Institutional Equity Office at 405-325-3546.

Adjustments for Pregnancy and Related Issues

Should you need modifications or adjustments to your course requirements because of pregnancy or a pregnancy-related condition, please request modifications via the Institutional Equity Office website or call the Institutional Equity Office at 405/325-3546 as soon as possible. Also, see the Institutional Equity Office FAQ on Pregnant and Parenting Students' Rights for answers to commonly asked questions.

Reasonable Accommodation Policy

The University of Oklahoma (OU) is committed to the goal of achieving equal educational opportunity and full educational participation for students with disabilities. If you have already established

reasonable accommodations with the Accessibility and Disability Resource Center (ADRC), please log into iAdvise to request your semester accommodations as soon as possible and contact me privately, so that we have adequate time to arrange your approved academic accommodations.

If you have not yet established services through ADRC, but have a documented disability and require accommodations, please complete ADRC's pre-registration form to begin the registration process. ADRC facilitates the interactive process that establishes reasonable accommodations for students at OU. For more information on ADRC registration procedures, please review their website. You may also contact them at (405)325-3852 or adrc@ou.edu, or visit www.ou.edu/adrc for more information.

Note: disabilities may include, but are not limited to, mental health, chronic health, physical, vision, hearing, learning and attention disabilities, pregnancy-related. ADRC can also support students experiencing temporary medical conditions.

Religious Observance

It is the policy of the University to excuse the absences of students that result from religious observances and to reschedule examinations and additional required classwork that may fall on religious holidays, without penalty. [See Faculty Handbook 3.15.2]

Final Exam Preparation Period

Pre-finals week will be defined as the seven calendar days before the first day of finals. Faculty may cover new course material throughout this week. For specific provisions of the policy please refer to OU's Final Exam Preparation Period policy.

Emergency Protocol

During an emergency, there are official university procedures that will maximize your safety.

Severe Weather: If you receive an OU Alert to seek refuge or hear a tornado siren that signals severe weather.

1. Look for severe weather refuge location maps located inside most OU buildings near the entrances.
2. Seek refuge inside a building. Do not leave one building to seek shelter in another building that you deem safer. If outside, get into the nearest building.
3. Go to the building's severe weather refuge location. If you do not know where that is, go to the lowest level possible and seek refuge in an innermost room. Avoid outside doors and windows.
4. Get in, Get Down, Cover Up
5. Wait for official notice to resume normal activities.

Additional Weather Safety Information is available through the Department of Campus Safety.

The University of Oklahoma Active Threat Guidance

The University of Oklahoma embraces a Run, Hide, Fight strategy for active threats on campus. This strategy is well known, widely accepted, and proven to save lives. To receive emergency campus alerts, be sure to update your contact information and preferences in the account settings section at one.ou.edu.

RUN: Running away from the threat is usually the best option. If it is safe to run, run as far away from the threat as possible. Call 911 when you are in a safe location and let them know from which OU campus you're calling from and location of active threat.

HIDE: If running is not practical, the next best option is to hide. Lock and barricade all doors; turn of all lights; turn down your phone's volume; search for improvised weapons; hide behind solid objects and walls; and hide yourself completely and stay quiet. Remain in place until law enforcement arrives. Be patient and remain hidden.

FIGHT: If you are unable to run or hide, the last best option is to fight. Have one or more improvised weapons with you and be prepared to attack. Attack them when they are least expecting it and hit them where it hurts most: the face (specifically eyes, nose, and ears), the throat, the diaphragm (solar plexus), and the groin.

Please save OUPD's contact information in your phone.

NORMAN campus: *For non-emergencies call (405) 325-1717. For emergencies call (405) 325-1911 or dial 911.*

TULSA campus: *For non-emergencies call (918) 660-3900. For emergencies call (918) 660-3333 or dial 911.*

Fire Alarm/General Emergency

If you receive an OU Alert that there is danger inside or near the building, or the fire alarm inside the building activates:

1. *LEAVE* the building. Do not use the elevators.
2. *KNOW* at least two building exits
3. *ASSIST* those that may need help
4. *PROCEED* to the emergency assembly area
5. *ONCE* safely outside, *NOTIFY* first responders of anyone that may still be inside building due to mobility issues.
6. *WAIT* for official notice before attempting to re-enter the building.

OU Fire Safety on Campus

Acknowledgements

I would like to thank, in alphabetical order, Professors Le Gruenwald, Dean Hougen, Amy McGovern, and Deborah Trytten, for providing valuable feedback on a preliminary version of the syllabus when I first arrived at the University of Oklahoma and proposed this new course.

References

[1] Solon Barocas, Moritz Hardt, and Arvind Narayanan. *Fairness and Machine Learning: Limitations and Opportunities*. fairmlbook.org, 2019. http://www.fairmlbook.org.

[2] Pedro Domingos. *The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World*. Basic Books, Inc., USA, 2018. ISBN 0465094279.

[3] Michael J. Kearns and Umesh V. Vazirani. *An Introduction to Computational Learning Theory*. MIT Press, 1994. ISBN 0-262-11193-4.

[4] Tom M. Mitchell. *Machine Learning*. McGraw Hill Series in Computer Science. McGraw-Hill, 1997. ISBN 978-0-07-042807-2.

[5] Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations of Machine Learning*. Adaptive computation and machine learning. MIT Press, 2012.

[6] Christoph Molnar. *Interpretable Machine Learning*. Independently Published, 2 edition, 2022. URL https://christophm.github.io/interpretable-ml-book.

[7] Cathy O'Neil. *Weapons of math destruction: How big data increases inequality and threatens democracy*. Broadway Books, 2017.

[8] Shai Shalev-Shwartz and Shai Ben-David. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press, New York, NY, USA, 2014. ISBN 1107057132, 9781107057135.

[9] Leslie Valiant. *Probably Approximately Correct: Nature's Algorithms for Learning and Prospering in a Complex World*. Basic Books, Inc., New York, NY, USA, 2013.

[10] Kush R. Varshney. *Trustworthy Machine Learning*. Independently Published, Chappaqua, NY, USA, 2022.