CS 5713 - Computational Learning Theory

Dimitris Diochnos
The University of Oklahoma, School of Computer Science

Fall 2023

Time and Location

The course will take place at Carson Engineering Center 0119. Tuesdays & Thursdays, 10:30am-11:45am.

Website: https://www.diochnos.com/teaching/CS5713/2023F/index.php

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

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-4:00pm**,
- on **Thursdays** between **2:00pm-3:00pm**,
- as well as on **Fridays** between **10:45am-11:45am**, or
- **by appointment**.

Please note that while anyone is welcome during the office hours, students from CS 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

Topics of machine learning theory. Learning using membership queries, equivalence queries, version spaces, decision trees, perceptrons. Probably approximately correct (PAC) learning, Occam algorithms, VC-dimension, sample sizes for distribution-independent learning. Representation issues, proper learning, reductions, intractability, learning in the realizable case, agnostic learning. Noise models, statistical queries, PAC learning under noise. Adversarially robust learning against poisoning attacks and against adversarial examples. Distribution-specific learning and evolvability. Online learning and learning with expert advice in the mistake bound model. Weak and strong learning (boosting).

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, August 22, 2023** and our **last meeting** is on **Thursday, December 8, 2023**.

A **rough outline** for the course, which is subject to change slightly depending on our pace, is:

Weeks 1-2: Introduction, concept learning, membership and equivalence queries, version spaces, and search algorithms.

Weeks 3-4 Decision tree learning and learning with perceptrons.

Weeks 5-9: Probably approximately correct (PAC) learning, Occam algorithms, PAC learnability of finite concept classes. Reducibility in PAC learning. Issues on representation and computational intractability. Infinite concept classes, VC theory and sample complexity bounds.

Weeks 9-10: Noise models and statistical queries. PAC learnability under noise.

Week 11: Adversarial machine learning - poisoning attacks and adversarial examples.

Week 12: Distribution-specific learning and evolvability.

Weeks 12-13: Online learning models and learning with expert advice.

Weeks 13-14: Weak vs strong learning. Boosting.

The **midterm exam** is **in-class** and will take place near the 6th week of the classes.

There will be **no final exam** for this class. Instead, we will have a **semester-long project**.

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.

Starting from this year (2023) the course will dive a bit less into the proofs of various results and instead will try to cover additional material from **trustworthy supervised learning**. Along these lines we will touch upon the following topics:

- *Adversarial Machine Learning*,
- *Interpretability*,
- *Fairness*,

and potentially other related topics such as *class imbalance* and *distribution/dataset shift*. In order to dive deeper into such 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**.

Grading

Grading will be based on the following:

- 40% homework problems,
- 40% semester-long project,
- 20% in-class midterm exam,

Grades may also be adjusted slightly upward or downward depending on class participation. Grading will 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.

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.

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

There will be 5-6 homework assignments.

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).

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. No collaboration will be allowed on exams.

Regarding the projects, students are of course allowed to work alone. In most cases they will also be allowed to work in pairs, but they need to discuss this with me in advance and make sure that they can indeed work in pairs on a particular project that they have in mind.

General Remarks. Please note the following two.

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

Late Work Policy

You can postpone once your homework 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).

We will be using an electronic system (Canvas) for the students' submissions and therefore it is your responsibility to turn in your homework (or an exam, should this be the case) on time. You are allowed to upload multiple copies of your work, so always make sure that you have submitted something. Apart from electronic submissions (Canvas), you can turn-in homework sets also in-class, by the end of the class on the day they are due.

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.

General Policies by the University of Oklahoma

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.

Food and Drink Policy. Food is not allowed classrooms. Consuming food in learning spaces is a significant concern for transmission of illness and therefore is prohibited. Students that may need to eat for health reasons, such as blood sugar regulation, should step outside the classroom to a social distanced location. If you are in a course that extends more than an hour, please be aware that some students may need to need to eat for health and may need to step out of classes briefly.

Drinking in classes is generally discouraged. It is acceptable, specifically in longer course formats, for students to take a sip from a water bottle or cup with a lid. Faculty may also need to take sips of water while they are teaching.

Academic Misconduct. Academic misconduct hurts everyone but particularly the student who does not learn the material. All work submitted for an individual grade should be the work of that single individual and not his/her friends. It is fine to ask a fellow student for help as long as that help does not consist of copying any computer code, or solutions to other assignments. Students working on joint projects may certainly help one another and are expected to share code within the project group. However, they may not share beyond the group.

1. Collaboration is encouraged for final projects. For the projects, you will work within your groups. For the homework, you may form study groups and exchange ideas, but in the end each one of you has to submit their own work in their own words.
2. Do not show another student (or group) a copy of your projects or homework before the submission deadline. The penalties for permitting your work to be copied are the same as the penalties for copying someone else's work.
3. Make sure that your computer account is properly protected. Use a good password, and do not give your friends access to your account or your computer system. Do not leave printouts or thumb drives around a laboratory where others might access them.

Upon the first documented occurrence of academic misconduct, I will report it to the Campus Judicial Coordinator. The procedure to be followed is documented in the University of Oklahoma Academic Misconduct Code. In the unlikely event that I elect to admonish the student, the appeals process is described in http://www.ou.edu/integrity.

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.

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 TBD .

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.

Class Evaluations. The College of Engineering utilizes student ratings 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.

Religious Holidays. It is the policy of the University to excuse the absences of students that result from religious observances and to provide without penalty for the rescheduling of examinations and additional required classwork that may fall on religious holidays.

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.

Accommodation of Disabilities. The University of Oklahoma is committed to providing reasonable accommodation for all students with disabilities. Students with disabilities who require accommodations in this course are requested to speak with the professor as early in the semester as possible. Students with disabilities must be registered with the Office of Disability Services prior to receiving accommodations in this course. The Office of Disability Services is located in 730 College Ave, phone 405/325-3852 or TDD only 405/325-4173.

Adjustments for Pregnancy/Childbirth Related Issues. Should you need modifications or adjustments to your course requirements because of documented pregnancy-related or childbirth-related issues, please contact me as soon as possible to discuss. Generally, modifications will be made where medically necessary and similar in scope to accommodations based on temporary disability. For commonly asked questions, please see https://www.ou.edu/eoo/faqs/pregnancy-faqs.htm

Title IX Resources. For any concerns regarding gender-based discrimination, sexual harassment, sexual misconduct, stalking, or intimate partner violence, the University offers a variety of resources, including advocates on-call 24.7, counseling services, mutual no contact orders, scheduling adjustments and disciplinary sanctions against the perpetrator. Please contact the Institutional Equity Office 405-325- 3546 (8-5, M-F) or OU Advocates 405-615-0013 (24.7) to learn more or to report an incident.

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**

Acknowledgements

I would like to thank Professors Le Gruenwald, Dean Hougen, Amy McGovern, and Deborah Trytten, for providing valuable feedback on a preliminary version of the syllabus.

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.

[3] Michael J. Kearns and Umesh V. Vazirani. *An Introduction to Computational Learning Theory*. MIT Press, 1994.

[4] Tom M. Mitchell. *Machine Learning*. McGraw Hill Series in Computer Science. McGraw-Hill, 1997.

[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.

[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.

[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.