TO: The Faculty of the College of Engineering

FROM: Elmore Family School of Electrical and Computer Engineering

RE: New Graduate Course, ECE 60125 Optimization for Deep Learning

The faculty of the School of Electrical and Computer Engineering has approved the following new course. This action is now submitted to the Engineering Faculty with a recommendation for approval.

ECE 60125 Optimization for Deep Learning
Sem. 1, Lecture 3, Cr. 3.
Prerequisite: graduate student standing

Description:
This course discusses the optimization algorithms that have been the engine that powered the recent rise of machine learning (ML) and deep learning (DL). The "learning" in modern ML and DL tasks typically boils down to non-convex optimization problems with high-dimensional parameter spaces and objective functions involving millions of terms. Additionally, the success of DL models is reliant on finding solutions to these problems efficiently. Hence, the needs of ML are different than those of other fields in terms of optimization. This course introduces students to the theoretical principles behind stochastic, gradient-based algorithms for DL as well as practical considerations such as adaptivity, generalization, distributed learning, and non-convex loss surfaces typically present in modern DL problems.

Reason:
The goal of this 600-level course is to introduce students to modern optimization techniques (theory and applications) for machine learning and deep learning. Existing optimization courses in IE and ECE (e.g., ECE 580, ECE 647, IE 56100) discuss traditional optimization concepts including convex analysis and convex optimization with applications to wireless communication networks and control systems; there is currently no course on modern optimization theory and algorithms that underpin the recent success of ML. Hence, the proposed course addresses the increasing demand of students for an ML-tailored optimization course and equips students with knowledge and skills to conduct research (both in industry and academia) in the area of ML and deep learning optimization. The materials discussed in this course are complementary to the existing courses and aim to bridge the existing gap between the practical applications of deep learning and the theoretical understanding of why and when these deep learning models work, from an optimization perspective.

Course History: Fall 2022 – 55, Fall 2023 – 41, Fall 2024 - 51

Mithuna Thottethodi,
Associate Head for Teaching and Learning
Elmore Family School of Electrical and Computer Engineering

Optimization for Deep Learning (OPT4DL)

1. General Class Information

- Course Number and Title: ECE xxxx – Optimization for Deep Learning
- Credit hours: Three (3)
- Pre-requisites: Undergraduate probability, calculus, and linear algebra. Basic knowledge of computer vision, NLP, machine learning, and statistics is helpful but not required.

- **Meeting Times:**
- **Days:**
- **Location:**
- **Instructor: Prof. Abolfazl Hashemi**

- The course website on Brightspace: https://purdue.brightspace.com/d2l/login
- The piazza website:

This course discusses the optimization algorithms that have been the engine that powered the recent rise of machine learning (ML) and deep learning (DL). The "learning" in modern ML and DL tasks typically boils down to non-convex optimization problems with high-dimensional parameter spaces and objective functions involving millions of terms. Additionally, the success of DL models is reliant on finding solutions to these problems efficiently. Hence, the needs of ML are different than those of other fields in terms of optimization. This course introduces students to the theoretical principles behind stochastic, gradient-based algorithms for DL as well as practical considerations such as adaptivity, generalization, distributed learning, and non-convex loss surfaces typically present in modern DL problems.

2. Instructors, Teaching Assistants and Office Hours

Prof. Abolfazl Hashemi abolfazl@purdue.edu	

Please note that instructor office hours may change during the semester. The latest information for office hours can be found on the course website and on Piazza.

3. Textbook, Videos, Technology and Class Notes

Required References: None (Lecture notes provided by the instructor on course website)

Assorted Recommended References: (all freely available online)

- Sébastien Bubeck et al. Convex optimization: Algorithms and complexity. Foundations and Trends in Machine Learning, 8(3-4):231–357, 2015.

- Suh, Changho. Convex optimization for machine learning. Now Publishers, 2022.
- Lan, Guanghui. First-order and stochastic optimization methods for machine learning. Switzerland: Springer, 2020.
- Prateek Jain and Purushottam Kar. Non-convex optimization for machine learning. Foundations and Trends in Machine Learning, 2017.
- Bottou, Léon, Frank E. Curtis, and Jorge Nocedal. "Optimization methods for large-scale machine learning." Siam Review 60.2 (2018): 223-311.
- Goodfellow, Ian, Yoshua Bengio, and Aaron Courville. “Deep learning”. MIT press, 2016.
- Zhang A, Lipton ZC, Li M, Smola AJ. “Dive into Deep Learning” . Cambridge University Press, 2023
- Grohs and Kutyniok “Mathematical Aspect of Deep Learning”. Cambridge University Press, 2022
- Cutkosky, Ashok. "Lecture Notes for EC525: Optimization for Machine Learning." EC525: Optimization for Machine Learning (2022).
- Roman Vershynin. High-dimensional probability: An introduction with applications in data science, volume 47. Cambridge University Press, 2018.
- Martin J. Wainwright. High-dimensional statistics: A non-asymptotic viewpoint. Cambridge University Press, 2019.
- Shai Shalev-Shwartz and Shai Ben-David. Understanding machine learning: From theory to algorithms. Cambridge university press, 2014.

4. Grading (tentative)

- Homework: 30% (6 biweekly sets)
- Midterm: 35%
- Final Project: 35%
- Some bonus points for participation

Homework policy: We will have biweekly homeworks. Homeworks are to be submitted at the beginning of the class when they are due. Your homework must be submitted to GradeScope by the beginning of class. You may discuss homework problems with other students but must submit your own independent solution. Late homework assignments will not be accepted.

Remark: For coding problems of each HW, you may use

- Google Colab: https://colab.research.google.com/
- Scholar cluster at Purdue RCAC: https://www.rcac.purdue.edu/compute/scholar

Midterm policy: The midterm will be a take-home exam. No collaboration is allowed. Late submissions will not be accepted. Using Piazza is not allowed.

Final project policy: A 6-page report (excluding references) in the format of a journal or a conference paper that provides a survey of a chosen topic, described in a one-page abstract due mid-semester. Additionally,

the team must make slides and record a 10-to-12-minute video on the chosen topic. The project should have a detailed survey component and a simulation component for full credit. Projects should be unbiased and well written. You may do the projects individually or in groups of two. However, the effort expected from group projects is more than the combined effort of two individual projects. More details will be released on around midsemester about the requirements of the final project. *We will use iThenticate to detect plagiarism and the use of LLMs such as ChatGPT both of which may be instances of academic dishonesty.*

Extra credit (bonus) policy: There is some bonus for participation. It is measured by two mechanisms: 1) Successfully filling out both the midsemester course and the final semester course. You will be asked to provide the proof of submission which will typically be in the form of taking a screen shot of submission confirmation. The proofs must be submitted to the links that will be provided on Brightspace when they are due. Late submission will not be accepted.

2) Scribing one lecture. The scribing should be done in a comprehensive and professional manner. You must use LaTeX for scribing. A template is available on Brightspace. You can refer to the instructor's notes. By scribing, you agree to share your notes with other students as well as students taking this class in future offerings. Scribing students must attend the lecture. In case the enrollment across all sections is more than the number of lectures, we use complementing mechanisms. No fractional point will be given.

Grading cutoffs: Your final class letter grade will be determined based on the following *approximate* cut-offs (range includes minus and plus grades as well):

Grade cut-offs	%
A range	85-100
B range	75-84.99
C range	60-74.99
D range	50-59.99
F range	0-50

The plus/minus grading scheme will be used within each grade range. Depending on the average class performance, these cutoffs may be modified up to, typically, 5% in either direction.

5. **Course outline (tentative)**

- **Introduction and foundations**
    - Relevant concept review
    - ML and optimization basics
    - Stochastic gradient descent
    - Smooth, nonconvex problems
    - Expected and high probability results
    - Gradient-free SGD

    - Beyond first-order results
    - Stochastic lower-bounds
- **Deep learning training techniques and insights**
    - Deep learning architectures
    - Backpropagation
    - Automatic differentiation and computation graphs
    - Initialization and normalization methods
    - Learning rate tuning methods
    - Regularization
- **Deep learning training algorithms**
    - Adaptive methods
    - Momentum
    - Variance reduction
    - Distributed DL
    - Decentralized SGD and Consensus
- **Special topics I**
    - Compression
    - Privacy-preserving ML
    - Overparameterized models and Interpolation
    - Neural Tangent Kernel
    - Implicit bias of SGD
- **Special topics II**
    - Robustness
    - Generalization error in DL
    - Double descent
    - Min-max optimization and GANs

6. **Learning Outcomes**

- Familiarize the students with modern non-convex, stochastic optimization concepts and algorithms
- Students gain a theoretical understanding of why and when deep learning models work.
- Students learn how to design new algorithms with theoretical guarantees.
- Students use the techniques from this course to tackle research problems in their theses.
- Students gain familiarity with recent research papers in this field.
- Students' final research project lead to publications in prestigious journal and conferences.

7. **Class Participation - Piazza.com**

This term we will be using Piazza for class discussion. The system is highly catered to getting you help fast and efficiently from classmates and the instructors. The class on piazza is structured so you can discuss each homework and topics on each exam. Please post to the relevant thread to ensure a proper response.

To honor and respect one another, it is important to adhere to the netiquette posted on our Piazza site.

You will receive an invitation to join piazza. If you register late for the course or if you don't receive the email, please go to piazza.com and register for the class yourself.

8. **Alternate Plan for Lecture delivery**

In case of health safety if Purdue decides to go fully online again in the middle of the semester, we may flip the lectures into Zoom lectures. Everything else should remain the same.

9. **Academic Guidance in the Event a Student is Quarantined/Isolated**

If you become quarantined or isolated at any point in time during the semester, in addition to support from the Protect Purdue Health Center, you will also have access to an Academic Case Manager who can provide you academic support during this time. Your Academic Case Manager can be reached at acmq@purdue.edu and will provide you with general guidelines/resources around communicating with your instructors, be available for academic support, and offer suggestions for how to be successful when learning remotely. Importantly, if you find yourself too sick to progress in the course, notify your academic case manager and notify the instructor via email. We will make arrangements based on your particular situation. The Office of the Dean of Students (odos@purdue.edu) is also available to support you should this situation occur.

10. **Attendance Policy**

Students are normally expected to attend all in-person classes and online students are expected to watch the recordings. Students should stay home and contact the Protect Purdue Health Center (496-INFO) if they feel ill, have any symptoms associated with COVID-19, or suspect they have been exposed to the virus. In the current context of COVID-19, in-person attendance will not be a factor in the final grades, but the student still needs to inform the TA of any conflict that can be anticipated and will affect the submission of an assignment or the ability to take an exam. Only the

TA and the instructor can excuse a student from a course requirement or responsibility. When conflicts can be anticipated, such as for many University-sponsored activities and religious observations, the student should inform the TA of the situation as far in advance as possible. For unanticipated or emergency conflict, when advance notification to the TA is not possible, the student should contact the TA as soon as possible by email. When the student is unable to make direct contact with the TA and is unable to leave word with the instructor's department because of circumstances beyond the student's control, and in cases of bereavement, quarantine, or isolation, the student or the student's representative should contact the Office of the Dean of Students via email or phone at 765-494-1747. Our course Brightspace includes a link on Attendance and Grief Absence policies under the University Policies menu.

11. Classroom Guidance Regarding Protect Purdue

The Protect Purdue Plan, which includes the Protect Purdue Pledge, is campus policy and as such all members of the Purdue community must comply with the required health and safety guidelines. Required behaviors in this class include: staying home and contacting the Protect Purdue Health Center (496-INFO) if you feel ill or know you have been exposed to the virus, adhering to mask requirements at all times (e.g., mask covers nose and mouth, no eating/drinking in the classroom), disinfecting desk/workspace prior to and after use, maintaining appropriate social distancing with peers and instructors (including when entering/exiting classrooms), refraining from moving furniture, avoiding shared use of personal items, maintaining robust hygiene (e.g., handwashing, disposal of tissues) prior to, during and after class, and following all safety directions from the instructor.

Students who are not engaging in these behaviors (e.g., wearing a mask) will be offered the opportunity to comply. If non-compliance continues, possible results include instructors asking the student to leave class and instructors dismissing the whole class. Students who do not comply with the required health behaviors are violating the University Code of Conduct and will be reported to the Dean of Students Office with sanctions ranging from educational requirements to dismissal from the university.

Any student who has substantial reason to believe that another person in a campus room (e.g., classroom) is threatening the safety of others by not complying (e.g., not wearing a mask) may leave the room without consequence. The student is encouraged to report the behavior to and discuss next steps with their instructor. Students also have the option of reporting the behavior to the Office of the Student Rights and Responsibilities. See also Purdue University Bill of Student Rights.

Related Considerations:

1. *A listing of recommended safe practices for the specific class or laboratory setting (other PPE or safety behavior) can be found at the links below.*
    - Overarching SOP for Classrooms, Instructional Laboratories, and Experiential Courses
2. *References Supporting Protect Purdue Compliance:*
    - Office of the Dean of Students Protect Purdue Compliance Plan: Ask, Offer, Leave, Report

Office of the Dean of Students Managing Classroom Behavior and Expectations

12. Accessibility

Purdue University strives to make learning experiences as accessible as possible. If you anticipate or experience physical or academic barriers based on disability, you are welcome to let the lead TA know so that we can discuss options. You are also encouraged to contact the Disability Resource Center at: drc@purdue.edu or by phone: 765-494-1247. More details are available on our course Brightspace under Accessibility Information.

Please note that any DRC letter may take up to 2-5 business days to process. Student must send DRC letters at least a week in advance before you want your accommodation. For very serious time sensitive issues contact instructor, Prof. Abolfazl Hashemi at abolfazl@purdue.edu, to discuss.

13. Nondiscrimination Policy

Purdue University is committed to maintaining a community which recognizes and values the inherent worth and dignity of every person; fosters tolerance, sensitivity, understanding, and mutual respect among its members; and encourages each individual to strive to reach his or her own potential. In pursuit of its goal of academic excellence, the University seeks to develop and nurture diversity. The University believes that diversity among its many members strengthens the institution, stimulates creativity, promotes the exchange of ideas, and enriches campus life. More details are available on our course Brightspace table of contents, under University Policies.

14. Academic Honesty Policy

Academic integrity is one of the highest values that Purdue University holds. Individuals are encouraged to alert university officials to potential breaches of this value by either emailing integrity@purdue.edu or by calling 765-494-8778. Check https://www.purdue.edu/odos/osrr/academic-integrity/index.html for more information.

In addition, the Purdue Honors Pledge applies to this course. The statement as written by our own students is **"As a boilermaker pursuing academic excellence, I pledge to be honest and true in all that I do. Accountable together - we are Purdue."**

Based on the above, we expect every member of the Purdue community to practice honorable and ethical behavior in and outside of the classroom. Any actions which might unfairly improve a student's score on homework or examinations will be considered cheating, and will not be tolerated.

A few examples of cheating include:

- Submitting homework that is not your own work or work conducted while working with your team.
- Sharing results or notes during exams. This includes accessing any websites including the ones like Chegg during exams, even if no aid is obtained.
- Continuing work on your exam after the exam time is over.
- Hacking or attempting to hack the class electronic tools/websites.
- Requesting a regrade on an exam that has been altered.

Cheating on homework, quizzes or exams will result in a zero score for the assignment/exam, or a failing grade for the course, at the instructors' discretion. In addition, we will report the case to the Dean of Students and request additional disciplinary action to be taken.

15. Mental Health Statement

If you find yourself beginning to feel some stress, anxiety, and/or feeling slightly overwhelmed, try WellTrack. Sign in and find information and tools at your fingertips, available to you at any time.

If you need support and information about options and resources, please contact or see the Office of the Dean of Students. Call 765-494-1747. The hours of operation are M-F, 8 am- 5 pm.

If you find yourself struggling to find a healthy balance between academics, social life, stress, etc. sign up for free one-on-one virtual or in-person sessions with a Purdue Wellness Coach at RecWell. Student coaches can help you navigate through barriers and challenges toward your goals throughout the semester. Sign up is completely free and can be done on BoilerConnect. If you have any questions, please contact Purdue Wellness at evans240@purdue.edu.

If you're struggling and need mental health services: Purdue University is committed to advancing the mental health and well-being of its students. If you or someone you know is feeling overwhelmed, depressed, and/or in need of mental health support, services are available. For help, such individuals should contact Counseling and Psychological Services (CAPS) at 765-494-6995 during and after hours, on weekends and holidays, or by going to the CAPS office of the second floor of the Purdue University Student Health Center (PUSH) during business hours.

16. Additional Wellness Resources

TaskHuman offers 1-on-1 live video calls with coaches who help you focus on wellness topics such as anxiety, mindfulness, reducing stress, clean eating, time management, in-home workouts, relationship tensions, and nearly a thousand more topics. You can log on at any time to access experiences as diverse as working through heightened anxiety to a personalized yoga session with carefully vetted providers. Using this link gets you access to all the perks: https://taskhuman.com/referral/purdue. Learn more here: https://engineering.purdue.edu/ECE/TaskHuman.

Don't see a topic you want or have other questions? Contact Brooke Parks, Lead Instructional Specialist in ECE, at brookeparks@purdue.edu.

17. Basic Needs Security:

Any student who faces challenges securing their food or housing and believes this may affect their performance in the course is urged to contact the Dean of Students for support. There is no appointment needed and Student Support Services is available to serve students 8 a.m.-5 p.m. Monday through Friday. Considering the significant disruptions caused by the current global crisis as it related to COVID-19, students may submit requests for emergency assistance from the Critical Needs Fund.