BME 646 and ECE 60146: Deep Learning:

Theory And Practice Of Deep Neural Networks

Prof. Avi Kak and Prof. Charles Bouman
Spring 2026
Purdue University

Course Description and Grading

- Course Description:
  - The purpose of this course is to teach the theory and practice of deep neural networks from basic principles through state-of-the-art methods. The class will blend hands-on programming using a variety of state-of-the-art programming frameworks with theoretical treatment based on current literature. Implementation will emphasize the use of the Pytorch language and the use of dynamic computational graphs. Competency in multivariate calculus and linear algebra, some previous experience with optimization techniques along with some programming experience is required for success in the course
- Assignments and Points:
  - There will be regularly assigned and graded programming assignments. These programming assignments must be performed independently by each student. Violation of this rule will be considered a form of cheating. There will be two midterm exams and no final exam.
  - Final grades will use the following weighting.
    - Programming 50%

      This grade will be calculated from roughly 8 equally weighted programming assignments.
    - Theory 50%

      This grade will be calculated from two equally weighted midterms.

Academic Integrity Policies

- We expect every member of the Purdue community to practice honorable and ethical behavior both inside and outside the classroom. Any actions that might unfairly improve a student's score on homework, quizzes, labs, or examinations will be considered cheating and will not be tolerated. Examples of cheating include (but are not limited to):
  - Sharing results or other information during an examination.
  - Bringing forbidden material or devices to an examination.
  - Working on an exam before or after the official time allowed.
  - Requesting a re-grade of answers or work that has been altered.
  - Submitting a homework or laboratory report that is not your own work, or engaging in forbidden homework or laboratory report collaboration.
  - Possession of another person's laboratory solutions or report from the current or previous years.
  - Reference to, or use of another person's laboratory solutions or report from the current or previous years.
  - Allowing another person to copy your laboratory solutions or work.
  - Representing as your own work anything that is the result of the work of someone else.
- All homeworks and laboratories must be performed independently by each student. Violation of this rule will be considered a form of cheating. At the professor's discretion, cheating on an assignment, or examination will result in a failing grade for the entire course, or a reduced grade, or a zero score for the particular assignment, or exam. Occurrences of academic dishonesty will be reported to the Assistant Dean of Students and copied to the ECE Assistant Head for Education. If there is any question as to whether a given action might be construed as cheating, please see the professor or the TA before you engage in any such action.
- 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 or by calling 765-494-8778. While information may be submitted anonymously, the more information that is submitted provides the greatest opportunity for the university to investigate the concern.
- The Honor Pledge Task Force, a student organization responsible for stewarding the mission of the Honor Pledge and encouraging a culture of academic integrity, asks all instructors to prominently include the student-initiated Purdue Honor Pledge on their syllabus, as well as exams and key assignments:
  - The Purdue Honor Pledge "As a boilermaker pursuing academic excellence, I pledge to be honest and true in all that I do. Accountable together - we are Purdue"
- Also, you may refer to Purdue's student guide for academic integrity for more information.

Course Policies – Part 1

- Nondiscrimination Statement:
  - 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. Link to Purdue's nondiscrimination policy statement.
- Students with Disabilities:
  - 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 me 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.
  - Purdue has assistance available to help you make learning materials accessible. Some examples include:
    - Information on Universal Design for Learning
    - Guidance on creating accessible documents
- Emergency Preparation:
  - In the event of a major campus emergency, course requirements, deadlines and grading percentages are subject to changes that may be necessitated by a revised semester calendar or other circumstances.
  - In the event of an emergency, students can get information from the following sources:
  - The following web page:
    - https://www.purdue.edu/ehps/emergency_preparedness/flipchart/index.html

    provides resources in case of emergencies that affect the West Lafayette campus.
  - Keep your cell phone on but with the ringer and buzzer off to receive a Purdue ALERT text message.
  - Log into a Purdue computer connected to the network to receive any Desktop Popup Alerts.
  - In an emergency, students are also welcome to contact Prof. Bouman by phone at his office or home.
- 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 see the Office of the Dean of Students for drop-in hours (M-F, 8 am- 5 pm).
  - 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.

Course Policies – Part 2

- Violent Behavior Policy:
  - Purdue University is committed to providing a safe and secure campus environment for members of the university community. Purdue strives to create an educational environment for students and a work environment for employees that promote educational and career goals. Violent Behavior impedes such goals. Therefore, Violent Behavior is prohibited in or on any University Facility or while participating in any university activity.
  - See the University's full violent behavior policy for more detail.
- Diversity and Inclusion Statement:
  - In our discussions, structured and unstructured, we will explore a variety of challenging issues, which can help us enhance our understanding of different experiences and perspectives. This can be challenging, but in overcoming these challenges we find the greatest rewards. While we will design guidelines as a group, everyone should remember the following points:
    - We are all in the process of learning about others and their experiences. Please speak with me, anonymously if needed, if something has made you uncomfortable.
    - Intention and impact are not always aligned, and we should respect the impact something may have on someone even if it was not the speaker's intention.
    - We all come to the class with a variety of experiences and a range of expertise, we should respect these in others while critically examining them in ourselves.
- Course Evaluation:
  - During the last two weeks of the course, you will be provided with an opportunity to evaluate this course and your instructor. Purdue uses an online course evaluation system. You will receive an official email from evaluation administrators with a link to the online evaluation site. You will have up to two weeks to complete this evaluation. Your participation is an integral part of this course, and your feedback is vital to improving education at Purdue University. I strongly urge you to participate in the evaluation system.

Textbook References

- Andrew W. Trask, Grokking Deep Learning, January 2019 ISBN 9781617293702
  - https://www.manning.com/books/grokking-deep-learning
- François Chollet, Deep Learning with Python, Manning Publications, ISBN 9781617294433, Nov. 2017.
  - https://www.manning.com/books/deep-learning-with-python
- Ian Goodfellow, Yoshua Bengio, and Aaron Courville, Deep Learning, MIT Press, 2016.
  - http://www.deeplearningbook.org
- Michael Nielsen, Neural Networks and Deep Learning, Dec. 2019.
  - http://neuralnetworksanddeeplearning.com/index.html
- Joel Grus, Data Science from Scratch, 2nd Edition, O'Reilly Media, Inc., 2019, ISBN: 9781492041122.
  - https://www.oreilly.com/library/view/data-science-from/9781492041122
- Sebastian Raschka, Python Machine Learning, ISBN 9781783555130, Sept. 2015.
  - https://www.packtpub.com/big-data-and-business-intelligence/python-machine-learning

Some Online Resources

- Stanford CS231n: Convolutional Neural Networks for Visual Recognition
  - http://cs231n.github.io
- CMU 11-785: Introduction to Deep Learning
  - http://deeplearning.cs.cmu.edu
- Oxford: Machine Learning
  - https://www.cs.ox.ac.uk/people/nando.defreitas/machinelearning

TA Support

- Akshita Kamsali
  - Email: akamsali@purdue.edu
- Sowmya Seeram
  - Email: sseeram@purdue.edu

Intellectual Ground Rules

- Thinking is required
  - Deep learning causes shallow thinking
  - Abstraction is both powerful and practical
- Theory is practical
  - I like math, so you will be expected to like it too.
  - Equations are a very compact way of expressing powerful **abstract** ideas
  - Pseudo-code is a type of mathematical language
- Think with both your Left and Right brain:
  - Emotional equations
  - Use both analytical tools, intuition, and **common sense**
- I don't like mushy thinking
  - Gobbledygook; Gibberish; Nonsense B.S.:
    - Sequences of words with no meaning and therefore no truth value.
    - ChatGPT, Transformative Administrator Speak on iPhone
  - $f(x) = g(y)$ or $z = \frac{A}{B}x$ where A and B are matrices
  - You can't prove something true or false if it has no meaning.

Big Ideas

- Signals are points in a high dimensional space.
  - Almost anything can be thought of as a point in a space.
  - But high dimensional spaces are different than you think…
- Images, text, anything, can be thought of as a random signal.
  - A random signal is a cloud of points in a high dimensional space.
  - Machine learning is a form of estimation
- Loss function minimization can be used to design systems.
  - Loss function can be design to achieve various goals, MLE, MMSE, …
  - Reliable minimization of non-convex loss functions is tricky.
- Random vector generation is changing the world.
  - Generative Models: Variational Encoders, GANs, Generative Diffusion
  - Goal: Learn to generate random variables from empirical distributions
- Attention/Self-similarity is extraordinarily powerful
  - The basis of transformers and LLMs

What is the Deep Learning Breakthrough?

- Around 2012, many things came together that were maturing over many years, and a paradigm change occurred.
- Computing resources
  - High performance CPU and GPU computers
  - Enormous databases of training data (e.g. Google)
- Optimization
  - Automatic differentiation for optimization of complex and deep networks
- Software
  - Large scale open source software systems such as python, numpy, pytorch, TensorFlow, etc.
- Architecture
  - Convolutional Neural Networks (CNN), Recurrent Neural Networks, Generative Neural Networks
- Frictionless Reproducibility[1]
  - The combination of Data Sharing, Code Sharing, and common Challenges has dramatically accelerated research in the AI field.

[1] David Donoho, "Data Science at the Singularity", *Harvard Data Science Review*, 6(1), Jan. 2024.

Theory Outline

Fall 2026 Syllabus

- Week 1 – Jan 15
  - What is Machine Learning?
  - Single Layer NN
  - The Loss Function
- Week 2 – Jan 22
  - Gradient Descent Optimization
  - Mathematical calculation of gradient
  - Matrix interpretation of gradient computation
- Week 3 – Jan 29
  - Tensors
  - Gradient Descent for Single Layer NN
  - Local and Global Minima
- Week 4 – Feb 5
  - Optimization of Deep Functions
  - Gradient Descent on Acyclic Graphs
  - General Loss Functions
- Week 5 – Feb 12
  - Convolutional Neural Networks
  - Gradients for CNNs
- Week 6 – Feb 19
  - Probability and Estimation
  - Frequentist versus Bayesian Estimation
  - Bias and Variance
- Week 7 – Feb 26
  - Training and Generalization
  - Regularization and Dropout Methods
- Week 8 – March 5
  - Exam #1
- Week 9 – March 12
  - Stochastic Gradient Descent
  - Batches and Epochs
  - Learning Rate and Momentum
- Week 10
  - Spring break
- Week 11 – March 26
  - Widely Used DL Techniques
  - Vanishing Gradients; Skip Connections
  - Batch Normalization
  - Transfer Learning and Data Augmentation
- Week 12 – April 2
  - Multi-Head Attention Layers
  - Transformer Layers
  - The Visual Transformer
- Week 13 – April 9
  - Recurrent Neural Networks
  - LSTM and GRU Networks
  - Unsupervised Learning and Autoencoders
- Week 14 – April 16
  - Adversarial learning
  - Generative Adversarial Networks (GAN)
  - Nash equilibrium
  - GAN convergence, theory and practice
  - Gibbs Distributions, and Score
- Week 15 – April 23
  - Exam #2
- Week 16 – April 30
  - Denoising Score Matching, and Langevin dynamics
  - Generative Diffusion Models
  - DALL-E Generative Diffusion Model

Exams in Weeks 8 and 15