Summer 2025

Delivery: 100% Web-Based on Canvas, with submissions on Canvas/Gradescope **Dates course will run**: May 12th, 2025 – July 31st 2025

Instructor Information

Dr. Kira Zsolt

Email: zkira@gatech.edu

General Course Information

Description

Deep learning is a sub-field of machine learning that focuses on learning complex, hierarchical feature representations from raw data. The dominant method for achieving this, artificial neural networks, has revolutionized the processing of data (e.g. images, videos, text, and audio) as well as decision-making tasks (e.g. game-playing). Its success has enabled a tremendous amount of practical commercial applications and has had a significant impact on society.

In this course, students will learn the fundamental principles, underlying mathematics, and implementation details of deep learning. This includes the concepts and methods used to optimize these highly parameterized models (gradient descent and backpropagation, and more generally computation graphs), the modules that make them up (linear, convolution, and pooling layers, activation functions, etc.), and common neural network architectures (convolutional neural networks, recurrent neural networks, etc.). Applications ranging from computer vision to natural language processing and decision-making (reinforcement learning) will be demonstrated. Through in-depth programming assignments, students will learn how to implement these fundamental building blocks as well as how to put them together using a popular deep learning library, PyTorch. In the final project, students will apply what they have learned to real-world scenarios by exploring these concepts with a problem that they are passionate about.

Pre- &/or Co-Requisites

Students must have a strong mathematical background (linear algebra, calculus especially taking partial derivatives, and probabilities & statistics) and at least an introductory course in Machine Learning (e.g. equivalent to CS 7641). This is not a soft requirement: This should **NOT** be your first ML class, and self-study (e.g. online Coursera/Udacity courses) do not count. While you may be able to do OK on some parts of the class, other parts such as quizzes will be difficult to do if you do not have any background in ML. **Strong** programming skills (specifically Python) are necessary to complete the assignments.

Course Objectives

• Describe the major differences between deep learning and other types of machine learning algorithms.

<u>Georgia Institute of Technology</u>

Course Syllabus: CS7643 Deep Learning

- Explain the fundamental methods involved in deep learning, including the underlying optimization concepts (gradient descent and backpropagation), typical modules they consist of, and how they can be combined to solve real-world problems.
- Differentiate between the major types of neural network architectures (multi-layered perceptrons, convolutional neural networks, recurrent neural networks, etc.) and what types of problems each is appropriate for.
- Select or design neural network architectures for new data problems based on their requirements and problem characteristics and analyze their performance.
- Describe some of the latest research being conducted in the field and open problems that are yet to be solved.

Course Materials

Course Text

Deep Learning, by Ian Goodfellow and Yoshua Bengio and Aaron Courville, MIT Press. Available <u>online</u>.

Additional Materials/Resources

All additional reading materials will be available via PDF on Canvas.

Course Website and Other Classroom Management Tools

All course materials and videos are located on Canvas.

Course Requirements, Assignments & Grading

Assignment Distribution and Grading Scale

Assignments	Weight
On-Boarding Quiz	(Required to verify identity using proctoring software)
Pre-requisites problem set (PS0) and Quiz 0	1.25%
Assignments (3)	60%
Quizzes (equally weighted)	18.75%
Final Project (including proposal)	20%
Ed Participation Bonus	1% extra credit for top 3 endorsed posters

Grading Scale

Your final grade will be assigned as a letter grade, with *at least* the following grades (i.e., 90 or greater will definitely be an A).

Α	90-100%
В	80-89%
C	70-79%
D	60-69%
F	0-59%

Assignment Due Dates

All assignments are due at 08:00:00 AM Eastern Time (ET) on the day the assignment is due, unless otherwise noted. If you are located outside of ET, Canvas will display the due dates in your local time (which can be changed by editing your personal Canvas settings). We will not accept assignments submitted late due to time zone issues, so do verify your desired settings as there are no exceptions.

Every homework deliverable and project deliverable will have a 48-hour **grace period** during which no penalty will apply. This is intended to allow you time to verify that your submission has been submitted (we recommend you re-download it and look it over to make sure all questions/deliverables have been answered). Canvas will show your submission as late, but you do not have to ask for this grace period. **Deliverables after the grace period will receive a grade of 0.** There is no grace period for taking quizzes.

Project

The project is a group assignment comprising of 2-4 members (instructor permission is required for other group arrangements; strong justification will be needed for approval). The class project is meant for students to (1) gain experience implementing deep models and (2) try Deep Learning on problems that interest them. The amount of effort should be at the level of one homework assignment per group member. A PDF write-up describing the project in a self-contained manner will be the sole deliverable. Your final write-up will be structured like a paper from a computer vision conference (CVPR, ECCV, ICCV, etc.). We will release this template as well as rubric. Additionally, we will allow people to upload additional code, videos, and other supplementary material like code upload for assignments. While the PDF may link to supplementary material, external documents and code, such resources may or may not be used to evaluate the project. The final PDF should completely address all the points in the rubric that will be released.

Technology Requirements and Skills

Computer Hardware and Software

- High-speed Internet connection
- Laptop or desktop computer with a minimum of a 2 GHz processor and 8 GB of RAM

- CUDA compatible GPU is helpful for assignments but not necessary.
- UNIX-like OS experience is recommended (Linux/macOS)
- Windows/Linux for PC computers OR Mac iOS for Apple computers.
- Complete Microsoft Office Suite or comparable and ability to use Adobe PDF software (install, download, open and convert)
- Mozilla Firefox, Chrome browser, and/or Safari browsers (Chrome required for on-boarding quiz)

Canvas

This class will use Canvas to deliver course materials to online students. ALL course materials and quiz assessments will take place on this platform. Gradescope will be used for submission of assignments and the project.

Proctoring Information

To verify the identity of all GT online students, all online students are *required* to complete the onboarding quiz that uses Honorlock. Honorlock is utilized for student identity verification and to ensure academic integrity. Honorlock provides student identity verification via facial and ID photos. You may also be asked to scan the room around you. The onboarding quiz will be a practice quiz that will not affect your grade in the course. You can take the onboarding quiz as many times as you want. All potential violations are reviewed by a human. The Honorlock support team is available 24/7. While Honorlock will not require you to create an account, download software, or schedule an appointment in advance, you will need Google Chrome and download the Honorlock Chrome Extension. Information on how to access Honorlock and additional resources are provided below. You can also access Honorlock support at <u>https://honorlock.com/support/</u>.

Course Policies, Expectations & Guidelines

Communication Policy

You are responsible for knowing the following information:

- 1. Anything posted to this syllabus.
- 2. Anything emailed directly to you by the teaching team (including announcements via Ed Discussion), 24 hours after receiving such an email or post.

Because Ed announcements are emailed to you as well, you need only to check your Georgia Tech email once every 24 hours to remain up to date on new information during the semester. Georgia Tech generally recommends students to check their Georgia Tech email once every 24 hours. So, if an announcement or message is time sensitive, you will not be responsible for the contents of the announcement until 24 hours after it has been sent.

Late and Make-up Work Policy

There will be no make-up work provided for missed assignments. Of course, emergencies (illness, family emergencies) will happen. In those instances, please <u>contact the Dean of Students office</u>. The Dean of Students is equipped to verify emergencies and pass confirmation on to all your classes. For consistency, we ask all students to do this in the event of an emergency. Do not send any

personal/medical information to the instructor or TAs; all such information should go through the Dean of Students.

Online Student Conduct and (N)etiquette

Communicating appropriately in the online classroom can be challenging. All communication, whether by email, Ed, Canvas, or otherwise, must be professional and respectful. In order to minimize this challenge, it is important to remember several points of "**internet etiquette**" that will smooth communication for both students and instructors.

- 1. <u>Read first, Write later</u>. Read the ENTIRE set of posts/comments on a discussion board before posting your reply, to prevent repeating commentary or asking questions that have already been answered.
- <u>Avoid language that may come across as strong or offensive.</u> Language can be easily
 misinterpreted in written electronic communication. Review email and discussion board posts
 BEFORE submitting. Humor and sarcasm may be easily misinterpreted by your reader(s).
 Try to be as matter of fact and as professional as possible.
- 3. <u>Follow the language rules of the Internet.</u> Do not write using all capital letters, because it will appear as shouting. Also, the use of emoticons can be helpful when used to convey nonverbal feelings. ☺
- 4. <u>Consider the privacy of others</u>. Ask permission prior to giving out a classmate's email address or other information.
- 5. <u>Keep attachments small</u>. If it is necessary to send pictures, change the size to an acceptable 250kb or less (one free, web-based tool to try is picresize.com).
- 6. <u>No inappropriate material.</u> Do not forward virus warnings, chain letters, jokes, etc. to classmates or instructors. The sharing of pornographic material is forbidden.

NOTE: The instructor reserves the right to remove posts that are not collegial in nature and/or do not meet the Online Student Conduct and Etiquette guidelines listed above.

University Use of Electronic Email

A university-assigned student e-mail account is the official university means of communication with all students at Georgia Institute of Technology. Students are responsible for all information sent to them via their university-assigned e-mail account. If a student chooses to forward information in their university e-mail account, he or she is responsible for all information, including attachments, sent to any other e-mail account. To stay current with university information, students are expected to check their official university e-mail account and other electronic communications on a frequent and consistent basis. Recognizing that some communications may be time-critical, the university recommends that electronic communications be checked minimally twice a week.

Plagiarism & Academic Integrity

Georgia Tech aims to cultivate a community based on trust, academic integrity, and honor. Students are expected to act according to the highest ethical standards. All students enrolled at Georgia Tech, and all its campuses, are to perform their academic work according to standards set by faculty members, departments, schools and colleges of the university; and cheating and plagiarism constitute fraudulent misrepresentation for which no credit can be given and for which appropriate sanctions are warranted and will be applied. For information on Georgia Tech's Academic Honor

Code, please visit http://www.catalog.gatech.edu/policies/honor-code/ or http://www.catalog.gatech.edu/rules/18/.

You are encouraged to discuss problems and papers with others as long as this does not involve the copying of code or solutions. After discussions, all materials that are part of a submission should be wholly your own. Any public material that you use to gain an understanding of the materials (opensource software, help from a textbook, or substantial help from a friend, etc.) should be acknowledged explicitly in anything you submit to us. To re-emphasize, no matter what the source *you cannot copy any existing code*, from other students, online, or otherwise, and all code must be wholly your own code that you wrote by yourself. Our current policy is that you **CANNOT** use any Al-assisted (e.g. based on large language models such as Github Copilot) text or coding responses in part or whole; again, every piece of text or code you submit must be wholly your own. If you have any doubts about whether something is legal or not, please do check with the class Instructor or the TA.

We will actively check for cheating, and any act of dishonesty will result in a Fail grade. Any student suspected of cheating or plagiarizing on a quiz, exam, or assignment will be reported to the Office of Student Integrity, who will investigate the incident and identify the appropriate penalty for violations.

AI-based assistance

We will use the AI-based assistance policy developed by <u>David Joyner</u>. We treat AI-based assistance, such as ChatGPT and Copilot, the same way we treat collaboration with other people: you are welcome to talk about your ideas and work with other people, both inside and outside the class, as well as with AI-based assistants.

However, all work you submit must be your own. You should never include in your assignment anything that was not written directly by you without proper citation (including quotation marks and in-line citation for direct quotes).

Including anything you did not write in your assignment without proper citation will be treated as an academic misconduct case.

If you are unsure where the line is between collaborating with AI and copying from AI, we recommend the following heuristics:

Heuristic 1: Never hit "Copy" within your conversation with an AI assistant. You can copy your own work into your conversation, but do not copy anything from the conversation back into your assignment.

Instead, use your interaction with the AI assistant as a learning experience, then let your assignment reflect your improved understanding.

Heuristic 2: Do not have your assignment and the AI agent open at the same time. Similar to above, use your conversation with the AI as a learning experience, then close the interaction down, open your assignment, and let your assignment reflect your revised knowledge.

This heuristic includes avoiding using AI directly integrated into your composition environment: just as you should not let a classmate write content or code directly into your submission, so also you should avoid using tools that directly add content to your submission.

Deviating from these heuristics does not automatically qualify as academic misconduct; however, following these heuristics essentially guarantees your collaboration will not cross the line into misconduct.

Students with Disabilities

If you are a student with learning needs that require special accommodation, contact the Office of Disability Services at 404.894.2563 or <u>http://disabilityservices.gatech.edu/</u>, as soon as possible, to make an appointment to discuss your special needs and to obtain an accommodations letter. Please also e-mail me as soon as possible to set up a time to discuss your learning needs.

Illness and Other Ailments

If you are a student that is negatively impacted by a health-related matter, please contact the Office of Disability Services or the Office of Dean of Students at 404.894.6367 or https://studentlife.gatech.edu/request-assistance. Do NOT send us any personal health information. They will provide you with an accommodation letter that will allow us to try to find a suitable schedule for completing all assignments. You MUST submit this and inform us that you did so on Ed before the due date for the deliverable.

Mental Health Resources and Support Services

Campus Resources

Georgia Tech Police Department Emergency: Call 911 | 404-894-2500

Dean of Students Office

404-894-2565 | studentlife.gatech.edu Afterhours Assistance Line & Dean on Call: 404-894-2204

Center for Assessment, Referral and Education (CARE)

404-894-3498 | care.gatech.edu

Collegiate Recovery Program 404-894-2575 |

counseling.gatech.edu

Counseling Center 404-894-2575 | counseling.gatech.edu

Health Initiatives 404-894-9980 healthinitiatives.gatech. edu

LGBTQIA Resource

Center

404-385-4780 | Igtbqia.gatech.edu Stamps Psychiatry Center 404-894-1420

VOICE 404-385-4464 | 404-385-4451 24/7 Info Line: 404-894-9000 | voice.gatech.edu

Women's Resource Center 404-385-0230 | womenscenter.gatech.edu

Veterans Resource Center 404-894-4953 | veterans.gatech.edu

Community Resources

Georgia Crisis and Access Line

1-800-715-4225 social workers and counselors 24 hours per day, every day, to assist those with urgent and emergency needs.

Trevor Project

1-866-488-7386 Trained counselors are available to support anyone in need.

National Suicide Prevention Hotline 1-800-273-8255

The crisis line is staffed with professional A national network of local crisis centers that provides free and confidential emotional support to people in suicidal crisis or emotional distress 24/7.

Georgia State Psychology Clinic

404-413-2500 The clinic offers high quality and affordable psychological services to adults, children, adolescents, families and couples from the greater Atlanta area.

Student-Faculty Expectations Agreement

At Georgia Tech we believe that it is important to strive for an atmosphere of mutual respect, acknowledgement, and responsibility between faculty members and the student body. See http://www.catalog.gatech.edu/rules/22/ for an articulation of some basic expectation that you can have of me and that I have of you. In the end, simple respect for knowledge, hard work, and cordial interactions will help build the environment we seek. Therefore, I encourage you to remain committed to the ideals of Georgia Tech while in this class.

Subject to Change Statement

The syllabus and course schedule may be subject to change. Changes will be communicated via the Canvas announcement tool. It is the responsibility of students to check Ed Discussions, email messages, and course announcements to stay current in their online courses.

Course Schedule

Week/Dates	Modules/Lessons	Deliverables
1 May 12 th	 Module 1: Introduction to Neural Networks Go through Welcome/Getting Started Lesson 1: Linear Classifiers and Gradient Descent Readings: DL book: Linear Algebra background DL book: Probability background DL book: ML Background LeCun et al., Nature '15 Shannon, 1956 	Quiz #0: Course prereqs and policies due May 19 th 8:00 AM ET (NO GRACE PERIOD) Pre-req Problem set (PS0) May 19 th 8:00 AM ET (NO GRACE PERIOD)
2 May 19 th	Lesson 2: Neural Networks Readings: <u>DL book: Deep Feedforward Nets</u> <u>Matrix calculus for deep learning</u> <u>Automatic Differentiation Survey, Baydin et al.</u> 	Quiz #1: Linear Classifiers and Gradient Descent (Lesson 1) and Neural Networks (Lesson 2) due May 26 th 8:00 AM ET (NO GRACE PERIOD)
3 May 26 th	Lesson 3: Optimization of Deep Neural Networks Readings: <u>DL book: Regularization for DL</u> <u>DL book: Optimization for Training Deep</u> <u>Models</u>	Assignment 1 due Jun 2nd 8:00 AM ET (grace period until Jun 4 th)
4 Jun 2 nd	 Module 2: Convolutional Neural Networks (OPTIONAL) Lesson 6: Data Wrangling Lesson 5: Convolution and Pooling Layers Lesson 6: Convolutional Neural Network Architectures Readings: Preprocessing for deep learning: from covariance matrix to image whitening cs231n on preprocessing DL book: Convolutional Networks Optional: Khetarpal, Khimya, et al. <u>Reevaluate: Reproducibility in evaluating</u> 	Quiz #2: Optimization of Deep Neural Networks (Lesson 3), and Convolution and Pooling Layers (Lesson 5) due Jun 9th 8:00 AM ET (NO GRACE PERIOD)

Week/Dates	Modules/Lessons	Deliverables
	reinforcement learning algorithms." (2018). See related <u>blog post</u>	
5	Lesson 7: Visualization	
Jun 9 th	Lesson 8: PyTorch and Scalable Training	
	Lesson 9: Advanced Computer Vision Architectures	
	Lesson 10: Bias and Fairness	
	Readings:	
	 <u>Understanding Neural Networks Through</u> <u>Deep Visualization</u> <u>Grad-CAM: Visual Explanations from Deep</u> <u>Networks via Gradient-based Localization</u> <u>Fully Convolutional Network for Semantic</u> <u>Segmentation</u> 	
6 Jun 16 th	Module 3: Structured Neural Representations Lesson 11: Introduction to Structured Representations	Assignment 2 due Jun 23rd 8:00 AM ET (grace period until Jun 25 th)
	Lesson 12: Language Models	
	Lesson 13: Embeddings	Project Proposal due Jun 20th 8:00 AM ET (grace period until June 22nd)
	Readings:	
	 <u>DL Book: Sequential Modeling and</u> <u>Recurrent Neural Networks (RNNs)</u> 	Quiz #3: Convolutional Neural Network Architectures (Lesson
	Word2vec tutorial	6), Visualization (Lesson 7), Advanced Computer Vision Architectures (Lesson 9) due Jun 23 rd 8:00 AM ET (NO GRACE PERIOD)
	Word2vec paper	
	<u>StarSpace paper</u>	
7	Lesson 14: Neural Attention Models	(Jun 28 th withdrawal deadline)
Jun 23 rd	Readings:	
	 <u>Attention is all you need</u> <u>BERT Paper</u> The Illustrated Transformer 	

Week/Dates	Modules/Lessons	Deliverables
8	Lesson 15: Neural Machine Translation	Assignment 3 due Jul 14 th 8:00 AM ET (grace period until Jul 16 th)
Jun 30 th	Lesson 16: Automated Speech Recognition (ASR	
		Quiz #4: Module 3 due Jul 7th 8:00 AM ET (NO GRACE PERIOD)
9	Module 4: Advanced Topics	
Jul 7 th Lesson 17: Deep Reinforce Readings: • <u>MDP Notes (courtesy E</u> • <u>Notes on Q-learning (courtesy E</u> • <u>Policy iteration notes (courtesy E</u>)	Lesson 17: Deep Reinforcement Learning	
	Readings:	
	 MDP Notes (courtesy Byron Boots) Notes on Q-learning (courtesy Byron Boots) Policy iteration notes (courtesy Byron Boots) 	
10 Jul 14 th	Lesson 18: Unsupervised and Semi-Supervised Learning	Quiz #5: Module 4 due Jul 21st 8:00 AM ET (NO GRACE PERIOD)
11	Lesson 19: Generative Models	
Jul 21 st	Readings:	
	 <u>Tutorial on Variational Autoencoder</u> <u>NIPS 2016 Tutorial: Generative Adversarial</u> <u>Networks</u> 	
12 Jul 28 th	Wrap-Up	Final Project due Jul 29th 8:00 AM ET (grace period until Jul 31 st)