

**GEORGIA INSTITUTE OF TECHNOLOGY
SCHOOL OF ELECTRICAL AND COMPUTER ENGINEERING**

**ECE 8803: Generative and Geometric Deep Learning
Syllabus – Fall 2023**



Instructor: Prof. Amirali Aghazadeh | amiralia@gatech.edu | www.amirali-aghazadeh.com

TA: TBA

Course Days/Times*: 12:30-1:45 pm, Tuesday and Thursday, Klaus 2443

Office Hours: TBA

Textbook: No required textbook but the following are excellent references for this class (tentative):

1. Bronstein, M. M., Bruna, J., Cohen, T., and Veličković, P. (2021). Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *arXiv preprint arXiv:2104.13478*.
2. Salakhutdinov, Ruslan. "Learning deep generative models." *Annual Review of Statistics and Its Application* 2 (2015): 361-385.
3. Ian Goodfellow, Yoshua Bengio, and Aaron Courville, *Deep Learning*, The MIT Press (November 18, 2016), ISBN-10: 0262035618, ISBN-13: 978-0262035613
Available online: <https://www.deeplearningbook.org/>
4. Generative Deep Learning, David Foster, O'Reilly, Inc. ISBN: 9781492041948

Prerequisite:

The main pre-requisite is an advanced course in machine learning (ECE 6254, ECE 4252/8803 or equivalent). This is not an introductory course to deep learning and will require a solid background in math, statistics, machine learning, and deep learning. Students must have Python and PyTorch programming expertise to train advanced deep-learning models.

Course Objective: The design, foundation, and applications of *generative* and *geometric* deep learning models

Academic Honesty: All Georgia Tech Honor Code violations will be handled by referring the case directly to the Dean of Students for investigation and penalties.

Homework: Check the Homework link on the main course Canvas page for all assignments often. Exam problems tend to be a variation of homework problems. Students are encouraged to form **groups** to discuss homework problems but must formulate their own write-ups. Georgia Tech's Honor Code will be strictly enforced, and students must always observe the code. Check Canvas Assignment often for due dates and submission instructions.

Exams: There will be no final exam in this course. There will be a final project. The course will have two midterm exams.

Assignments: Assignments will have both an analytical part and a hands-on part. The hands-on parts are coding assignments, which will look like mini projects. Students are expected to have a background in Python and PyTorch. Students are encouraged to utilize Google's Colab for their codes in this course. We will work with coc-ice towards reserving computing resources for the students. These hands-on assignments vary between adding to existing codes, writing codes from scratch, searching the literature for codes for a specific application, or summarizing papers in the literature.

Grading:

Homework assignments 40%

Exams 20%

Final project 40%

In all cases, we will at least follow the basic traditional grading scale where: A=90-100; B=80-89; C=70-79; D=60-69; F=0-59. While in all cases a grade of 90-100 will be assigned an 'A', the boundaries/cutoffs for the other grades will be determined at the end of the semester based on the overall class performance. It is impossible to determine what the exact "cutoffs" will be for each grade but you can be assured that your assigned grade will never be lower than the "traditional" grading scale described above based on your final class average. The project will be at the level of producing a short conference paper (4~6 pages, double-column ICML style).

Programming Language: We will utilize Python throughout the course. We will also utilize a library of Jupyter notebooks, including in Colab and PyTorch.

Canvas: Course website on Canvas: when clicking the course, students will see an overall view of all the course components, including Syllabus, Lecture, Homework, Exam, Piazza, Supplements, etc. Go to <https://canvas.gatech.edu/> and if you do not see the class page, make sure you are registered for the course.

Piazza: Students are expected to utilize PIAZZA platform to post questions. If you have any problems or feedback for the developers, email team@piazza.com. Find our class page by clicking the Piazza link on the left navigation menu on Canvas. If Georgia Tech ends Piazza and moves to another platform, we will be utilizing the new platform.

Assignments Submission: All homework assignments need to be submitted on Canvas. Read the instructions of each assignment carefully.

Attendance: Your attendance and participation are strongly encouraged. Check the Institute Absence Policy at: <http://www.catalog.gatech.edu/rules/4/>.

Typically, sometime in September, we must submit a report on Verification of Participation. If this applies this year, we will use lecture attendance and/or other metrics (e.g., TurningPoint survey responses and Piazza activities) to decide if you are participating or not.

Communications: All communication is expected to be conducted on Piazza. One can utilize the private message option. Emails are strongly discouraged.

Announcements: Official announcements will be posted on Canvas and/or Piazza or announced during lectures.

Academic Honesty: All violations of the Georgia Tech Honor Code will be handled by referring the case directly to the Dean of Students for investigation and penalties. The complete honor code can be found at this link: <http://www.policylibrary.gatech.edu/student-affairs/academic-honor-code>

Available Resources:

- The Center for Academic Success has programs to help students improve their study habits and time management: <http://www.successprograms.gatech.edu/>.
- The Dean of Students Office helps students who have personal or medical issues that impact their academic performance: <http://www.deanofstudents.gatech.edu/>

Office of Disability Services: If you are a student registered with the Office of Disability Services (ODS), please make sure the appropriate forms and paperwork are completed with the instructor within the first week of classes. The instructor will abide by all accommodations required by ODS. The schedule for exams is posted in the syllabus and any potential modifications or changes will be made with at least one week's notice. It is the responsibility of the student to properly arrange test accommodations for each exam with ODS in sufficient time to guarantee space for exam administration. ALL exam accommodations must be handled through ODS. If the student does not register accommodations with ODS for the taking of an exam, then they will have to take the exam at the normally scheduled times without any additional accommodation unless the instructor is given specific directive from ODS on the student's behalf due to a mitigating circumstance. (<https://disabilityservices.gatech.edu/>)

Topical Outline (tentative):

- Generative modeling (~12 lectures)
 - o Intro to probability and statistics

- Autoregressive models
 - Autoregressive Autoencoders
 - Recurrent Neural Networks
- Variational autoencodes (VAEs)
 - Mixture models
 - Marginal Likelihood
 - Evidence lower bound (ELBO)
- Normalizing flow models
- Generative adversarial networks (GANs)
 - Likelihood-free learning
 - GAN mechanics
 - GAN training algorithm
 - Optimization challenges
- Energy-based models
- Score-based Models
 - Denoising Diffusion models
- Evaluation of generative models
- Applications of deep generative modeling
 - Computer vision (DALL-E)
 - Language modeling (ChatGPT)
 - Protein modeling (AlphaFold)
- Geometric deep learning (~12 lectures)
 - Introduction
 - Symmetries
 - Early neural networks
 - Universal approximation
 - Emergence of geometries
 - Learning in high dimensions
 - The curse of dimensionality
 - Geometric priors
 - Groups
 - Representations
 - Invariance
 - Equivariance
 - Graphs and sets
 - Permutation Invariance and Equivariance
 - Groups
 - Grids
 - Fourier in neural network
 - Convolutional operators
 - Deformation stability
 - Mathematics of CNN
 - Geodesics and Manifolds
 - Geodesics
 - Isometries

- Non-Euclidian convolution
 - Manifold Fourier transform
 - Meshes
 - Gauges
 - Graph Neural Network
 - Geometric modeling for sciences
 - Deep learning on proteins
 - Deep learning on molecules
- (Time permitted) Adversarial Attacks (3-4 lectures)
- Introduction
 - Certifiable robustness
 - Trojan attacks/defense
 - Poisoning attacks/defense