

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

**ECE 2028: AI First “AI Foundations” (2-2-0-3)
Syllabus – Fall 2026**

Instructor: Prof. Ghassan AlRegib alregib@gatech.edu

Staff:

Teaching Fellows: TBA

Teaching Assistants: TBA

Undergraduate TAs: TBD

Course Days, Times, and Locations*:

Lectures on MW 09:30 – 10:20 a.m. (Scheller)

Studios: Check Oscar and Canvas, all in Klaus 2446

Office Hours:

Check Canvas

Textbook: No required textbook. On Canvas, you will find handouts, lecture notes, slides, video recordings, and links to online resources. Students will have access to materials on the AI Makerspace via the following systems StudyBuddy® and AIFirst®. The following existing books could serve as references for this class:

- Stuart Russell and Peter Norvig, Artificial Intelligence: A Modern Approach, Pearson (Fourth Edition) 2020. (<https://aima.cs.berkeley.edu/>) (ISBN-13 9780134610993)
- Chowdhary, K.R., Fundamentals of Artificial Intelligence, Springer, 2020. (ISBN 978-81-322-3972-7)
- <http://openbookproject.net/thinkcs/python/english3e/>

Prerequisite: Undergraduate Semester level [MATH 1554](#) (or MATH 1553) Minimum Grade of C **and** Undergraduate Semester level [CS 1301](#) Minimum Grade of C **OR** 1xxx-level COE Course on “Programming for Engineers.” Incoming students who have taking the equivalent to these classes while in high school can register for this class.

Course Objective: An introduction to AI Foundations.

1. To understand the basic concepts of Artificial Intelligence
2. To build on the basic concepts of computer programming in a high-level language
3. To understand the principles of data literacy
4. To understand the math principles for AI
5. To understand the process and skills necessary to effectively deal with problem solving in relation to constructing AI algorithms
6. To be able to test and evaluate AI models and systems
7. To understand the societal and ethical implications of AI systems
8. To become comfortable with common software packages in use today and on using industry-grade cluster such as the AI Makerspace

Course Description: Establish AI literacy for all students. Students will gain experience and practice with algorithm design, data structure, logical thinking, reasoning, explanations, and AI systems. A focus is on developing skills and experience in utilizing the AI Makerspace to perform certain functionalities on real data.

Course Outcomes:

1. Formulate real-life problems using data-driven approaches
2. Compute solutions using AI principles and reasoning
3. Analyze data and build AI components that fit the data and the application
4. Assess AI systems from technical as well as ethical perspectives

Students Outcomes:

1. (P) an ability to identify, formulate, and solve complex problems by applying principles of AI
2. (P) an ability to recognize ethical responsibilities of AI systems and assess their decisions in real-world contexts
3. (M) develop and conduct appropriate experimentation, analyze and interpret data, and use AI judgment to draw conclusions
4. (M) an ability to acquire and apply new knowledge as needed, using appropriate AI strategies

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. Past infractions have included cheating on tests, copying lab results, copying homework, and forging TA signatures. Additionally, you are not permitted to complete any in-class exercises or attendance checks for another student other than yourself. Doing so will be considered a violation of the Georgia Tech Honor Code and be handled accordingly. As a reminder to students, violations of the Honor Code can be met with minimum drop of one letter grade in their final course grade and potentially academic probation. The complete honor code can be found online at: <http://www.policylibrary.gatech.edu/student-affairs/academic-honor-code>

Grading:

Homework	20%	Quiz 1	15%	Quiz 2	15%
Studios	25%	Final Project	20%	Participation	5%

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 not be lower than the “traditional” grading scale described above based on your final class average.

The plan is to have 6 homework assignments and 6 studios. The grade of the lowest homework will not be counted towards the homework grade.

Exams: There will be two quizzes, amounting to **15 course points each out of 100**. Moreover, in lieu of the final exam, there will be a final portfolio project due at semester end, and it accounts for **20 course points out of 100**. There will be no make-up quizzes. If you have to miss a quiz, communicate with the instructors ahead of time with an official excuse.

Homework: Written homework will be assigned (approximately) every other week and will be due on Fridays (unless otherwise specified). For help, check the course website for office hours with instructors. Solutions will be posted for all HW assignments. **Late homework will not be accepted.** Check Canvas Assignments for submission instructions. Students will be provided by access to a number of tools and

services on the AI Makerspace such as StudyBuddy®, BuzzSync®, ChatGT®, and GTvizier®. A selection of assignments will be via one or more of these systems.

You must write up the assignments (and accordingly the code) by yourself, citing any outside references you use to arrive at your solution. **Failure to do so** will be considered a violation of the GT Academic Honor Code.

Unauthorized use of any previous semester course materials, such as tests, quizzes, homework, and any other coursework, is **prohibited** in this course. Using these materials will be considered a direct violation of the GT Academic Honor Code. For any questions involving these or any other Academic Honor Code issues, please consult me or <https://catalog.gatech.edu/policies/honor-code/>.

Solving homework assignments will take time. Plan accordingly. This is where most of the learning takes place. Effectively, homework is worth much more than its assigned weight of your grade. In teaching this and other classes, exam success has correlated with homework effort; students who do not engage with assignments invariably struggle on exams.

Programming Language: We will utilize Python on the AI Makerspace.

Laptops: We will spend a portion of some lectures or studios solving problems in class. There will be a set of exercises that must be completed in class and/or during studio. These tasks will be **online**, and the website will only be open for the class/studio period and possibly for a short time afterwards. Thus, **you will need to bring your laptop to every class and every studio.**

Studios (times vary, check your registered schedule): The Studio sections will meet once per week unless otherwise announced. See the Studio schedule on Canvas. A Studio has two parts: **Concepts Drills** and **Practicum**. In the first part, the purpose is to answer questions and facilitate a deeper and more personal interaction with an instructor on the course concepts that are covered during lectures. Ideally, this part is student led in that the instructor functions to assist students in the areas they feel most challenging. The instructors will use part of the studio time to perform **Concepts Drills**. The second part, Studio **Practicum**, is to explore hands-on applications of the course concepts using Python and the AI Makerspace. It is a critical component of the course. There will be approximately seven written Studio Practicum reports, including corresponding codes and plots, compiled into a single pdf file. Check Canvas Assignments for submission instructions. Attendance is expected at each studio. If you must miss a studio session for any reason, you **MUST** contact your instructor before the session with a valid excuse, or you will get an absence mark for the missed session. You can also ask the instructor's permit to attend another session in the same week. The instructors will be responsible for 5% of your grade. Do discuss grading expectations with your studio instructors. **Failing to submit one or more studio reports will result in grade F.**

Course Project: The final project will require additional requirements to be outlined in class at the middle of the semester. The project will be at the level of producing a short conference paper (max. 4 pages, double-column IEEE style).

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 and engage into online discussions. Make sure you enroll into the course site on Piazza. Check for sign-up link at LINK HERE. 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.

Attendance: Your attendance and participation are strongly encouraged. Check the Institute Absence Policy at: <https://studentlife.gatech.edu/resources/class-attendance>. Typically, in the first two weeks, 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, Canvas activities, or Piazza activities) to decide your participation. **You must attend the studio section you signed up for on the Registrar's record.**

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

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

Available Resources:

- The Center for Academic Success has programs to help students improve their study habits and time management: <https://www.success.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/>)

Student Collaboration: Students are encouraged to study together for homework, studio problems and exams to openly discuss course topics. However, each assignment that is turned in must reflect the work of each individual student. No copying of work from other students in (or out) of this class is allowed and such activity would represent a violation of the Academic Honor Code. If you are not certain of the nature of a student collaboration you are involved in, please feel free to contact Prof. AlRegib, or your recitation instructor for their expectations.

Teacher Commitment: All of the instructors and graduate teaching assistants commit to dedicating our time and energy to ensure that you have a productive learning environment for this course.

Student Commitment: As the student, you agree to commit your time and energy to learn the material by completing all assignments in a timely manner, attending all class sessions, and seeking help when you require it.

Topical Outline:

- What is AI?
 - AI Problem Solving; Real-world AI Systems and Applications; Hardware and GPUs; AI capabilities: narrow, general, and super AI; AI functionalities: from reactive to self-aware AI; AI Techniques: symbolic, machine learning, deep learning ... etc.; Predictive AI versus Generative AI
- Data:
 - Sensing and acquisition; Data Structures; Data Attributes; Data Exploration with basic statistics; Representations; Basis Functions; Orthogonality; PCA; Data Modalities; Entropy; Greedy Algorithms; Similarity and Distance; Mutual Information; KL Divergence; Cross Entropy
- Ethics and AI:
 - Case studies; Issues and Challenges; Policy making and governance; bias; fairness and trust in AI; uncertainty; adversarial challenges; transparency and explainability; implications on society and individuals; UN resolution; Engineering professional society ethics; ethical and professional responsibilities in real-world context; ethical principles versus deliberative processes
- Knowledge, Logic, and Reasoning:
 - Rules of Inference; Reasoning Patterns; Interpretations and Inferences; causality; Deductive, inductive, and abductive reasoning
- Learning:
 - Decision boundary; SVM; Inductive, Abductive and Deductive Learning; Learning by Analogy; supervised/unsupervised learning, backpropagation; gradient descent; activation foundations; regression; artificial neural networks; deep learning; finetuning; active learning; robustness; autoencoders; validation, cross-validation, and testing; overfitting and underfitting; regularization
- Gen AI:
 - Transformers; self-supervised learning; RAG; finetuning; LLMs
- Search, Decision and Planning
 - mini-max; Adversarial search; mathematical programming; resource allocation using LP; risk/reward optimization using QP; sequential decision making; dynamic programming; reinforcement learning
- Applications:
 - Anomaly detection; classification; clustering; segmentation; prediction; adversarial attacks; out-of-distribution detection; transfer learning