

Applied Natural Language Processing CSE 8803 - Fall 2026

Printable syllabus packet

Edition	Fall 2026
Course	Applied Natural Language Processing (CSE 8803)
EdStem / LMS	Canvas dashboard -> Ed Discussion
Omitted sections	Instructor/TA listings and course schedule omitted from this packet as requested

This packet compiles the course information, policies, grading structure, and resource material from the course website into a printable format.

Item	Details
Source website	https://anlp-website.github.io/fall25/
Format	Website-to-PDF syllabus packet
Content basis	Public Fall 2025 course website structure/content, organized into a printable Fall 2026 packet as requested
Layout note	Designed for printing and submission while preserving the course website's substantive content

Included sections

- Course information and success guidance
- Course overview and prerequisites
- General guidelines, office hours, and inclusion statement
- Grading categories and assessment policies
- Resource collection

Course Information

EdStem: Check your Canvas Dashboard and click on Ed Discussion.

Course description

The primary objective of this course is to introduce students to broad classes of techniques and tools for analyzing text data using Natural Language Processing (NLP) algorithms and techniques. The course emphasizes how to apply pre-processing, processing, and post-processing NLP techniques to analyze text and develop NLP models.

How to succeed in this class

Engage actively

- Participate in discussions regularly on the class Ed Discussion platform.
- Use office hours to clarify doubts, seek guidance, and deepen understanding.

Review all learning materials

- Review lectures, including videos, tutorial videos, and PDFs.
- Pay close attention to the selected readings because they complement and deepen the lecture content.

Manage assignments strategically

- Start assignments as soon as they are released.
- Use effective time management so the work supports learning, not just completion.

The site emphasizes that consistent effort, active participation, and timely engagement with course materials are crucial to success.

Course Overview

Course goals

Analyze NLP techniques and apply them to text data. The course is divided into three main categories:

- Preprocessing: demonstrate how to clean and integrate text data.
- Processing: apply NLP algorithms on pre-processed data to perform different tasks.
- Post-processing: evaluate developed NLP models.
- Solve problems with real datasets.
- Gain practical know-how through significant hands-on programming assignments.

Course pre- and/or co-requisites

Students are expected to have prior knowledge in machine learning, linear algebra, optimization, probability, and statistics. The programming language for the class is Python (Python 3.x), and students should know at minimum how to use NumPy, matrix operations, linear algebra, probability, and statistics.

Additional prerequisites

- CSE 6040
- CS 1301

Class text

- Required readings: none.
- Recommended reading: *Introduction to Natural Language Processing* by Jacob Eisenstein (draft available on GitHub).

General Guidelines, Office Hours, and Inclusion

Attendance

- This is a fully online course.
- Students should log in regularly so that work can be completed without falling behind and needing extensive review.

Class deliverables

- All class deliverables are handled through Gradescope, except quizzes, which are on Canvas.
- Deadlines will not be extended under any circumstances.
- Work submitted after the deadline receives zero credit.

EdStem

- EdStem is the main and only place for course discussions and announcements.
- If a question is suitable for the full class, students are asked to post it on Ed first so others can benefit and answers arrive faster.
- Sensitive matters may be addressed through a private post; if a more restricted conversation is needed, students can request a private chat.
- The site distinguishes good Ed questions about concepts, assignment interpretation, project questions, site issues, or suggestions from bad questions asking staff to debug code line-by-line.

Exceptional circumstances

- Requests for exceptions should be made in advance when possible and must be due to incapacitating illness, personal emergencies, or similarly serious events.
- The site directs students to work through the Dean of Students office before contacting the teaching team.

Office hours

- TAs plan to hold office hours starting in week 2, except on Georgia Tech holidays.
- Each session is one hour long and is run by at least one TA.
- Office hours will be held through live Ed Discussion threads, with details announced on Ed.
- Students are always welcome to ask questions on Ed; office hours supplement Ed and do not replace it.

Diversity and inclusion

- The course aims to create an environment where all voices are valued and where diversity of gender, sexuality, age, socioeconomic status, ability, ethnicity, race, and culture is respected.
- Students are encouraged to share suggestions that support that goal.
- Conflicts with religious events should be communicated to the instructional team.
- Disability accommodations are handled through Georgia Tech's Office of Disability Services.

Grading Categories and Assessment Policies

Assignments and assessment weights

Component	Weight / Notes
HW1	10% - week 1 to week 2 topics
HW2	15% - week 3 to week 4 topics
HW3	20% - week 5 to week 6 topics
HW4	20% - week 7 to week 11 topics
HW5	15% - week 12 to week 13 topics
Proctored Coding Sessions	5% - three short Honorlock / Vocareum sessions tied to the homework assignments
Quizzes	15% - 12 graded quizzes, lowest grade dropped, all quizzes mandatory

Proctored coding sessions

- Three short proctored coding sessions use Vocareum and Honorlock.
- Each session lasts up to one hour and is designed to take less than the full hour.
- They are intended to assess application of course understanding without relying on generative AI or online lookup tools.

Quizzes

- There are 12 graded quizzes on Canvas, including the knowledge-based quiz (Quiz 0).
- The lowest quiz grade is dropped so that 11 quizzes count toward the grade.
- Each quiz is worth 5 points, and missing a quiz results in a grade of -5 out of 5 rather than 0 out of 5.
- Quizzes are seven minutes long, have five multiple-choice questions, and are mostly conceptual.
- Quiz questions are drawn randomly from a question bank.
- Honorlock is used for quizzes; quizzes are open book and open notes, but additional electronic devices are not permitted and browser activity is restricted to Canvas.

Grading

- Students must achieve an overall weighted average of 60% to pass the course.
- All deliverables are graded by TAs / Gradescope.
- Standard grade thresholds (90, 80, etc.) may be lowered, but not raised.
- Some assignments may include bonus points.

Plagiarism, collaboration, and honor code

- All participants are expected to know and abide by the Georgia Tech Academic Honor Code.
- Plagiarism is a serious offense; submitted work must be the student's own.
- Students may discuss high-level ideas with peers, but each student must write and submit their own answers.

- Students must not place course code in public repositories because future students could copy it.
- Suspected dishonesty is handled through Georgia Tech's academic integrity procedures and can result in severe consequences.

Late policy and due dates

- All homework and quiz deliverables are due at the times shown in the course schedule.
- The course offers no late policy for homework and quizzes.
- Any deliverable submitted after the deadline receives zero credit.
- Students should submit early and verify that every required file is included.

Timing, attendance, and netiquette

- Course videos follow a logical sequence supporting both knowledge-building and experience-building.
- Students are expected to complete assignments by their due dates.
- This is a fully online course, and students should log in regularly.
- The site asks students to follow online discussion ground rules and to communicate respectfully, with correct spelling, punctuation, and grammar consistent with the academic environment.

Resource Collection

Recommended reading

- All content and course materials can be accessed online. The site states that there is no required textbook.
- Georgia Tech students have free access to O'Reilly resources using their official GT email address.

Software engineering / becoming a better programmer

- Debugging
- Clean Code
- Refactoring
- Design Patterns: Elements of Reusable Object-Oriented Software
- The Pragmatic Programmer: From Journeyman to Master

Python

- Python Bootcamp for campus MS Analytics students (by Chris Simpkins).

Data science, machine learning, data mining

- Data Science for Business by Foster Provost and Tom Fawcett.
- The Elements of Statistical Learning: Data Mining, Inference, and Prediction by Trevor Hastie, Robert Tibshirani, and Jerome Friedman.

Probability

- Free probability book by Prof. Guy Lebanon.

Human computation

- Human Computation by Edith Law and Luis von Ahn.

Managing multiple Python versions

- Recommended article: Which Python package manager should you use?
- Managing Multiple Python Versions With pyenv.
- The right and wrong way to set Python 3 as default on a Mac.
- Poetry is highlighted as an increasingly common alternative to conda environments and even setuptools for package workflows.