EECS 183/283a:
Natural Language Processing

Instructors:
Gopala Krishna Anumanchipalli, Alane Suhr

GSIs: Akshat Gupta, Charlie Snell, Sanjay Subramanian, Kayo Yin

Instructor: Alane Suhr

- From Appalachia (southeast Ohio)
- BS from Ohio State University in Computer Science and Engineering (+ minor in Linguistics)
- PhD from Cornell University, spent most of my time in New York City at Cornell Tech
- 1 year in Seattle as a postdoc at Ai2
- Research: language use and acquisition in interaction

Instructor: Gopala Krishna Anumanchipalli

GSI: Kayo Yin

- 4th year PhD co-advised by Jacob Steinhardt and Dan Klein
- Born in Kobe (Japan) and grew up in Paris (France)
- BS from École Polytechnique in Math and Computer Science
- MS from Carnegie Mellon University in Language Technologies
- Research: AI alignment, interpretability, sign language, cognitive linguistics

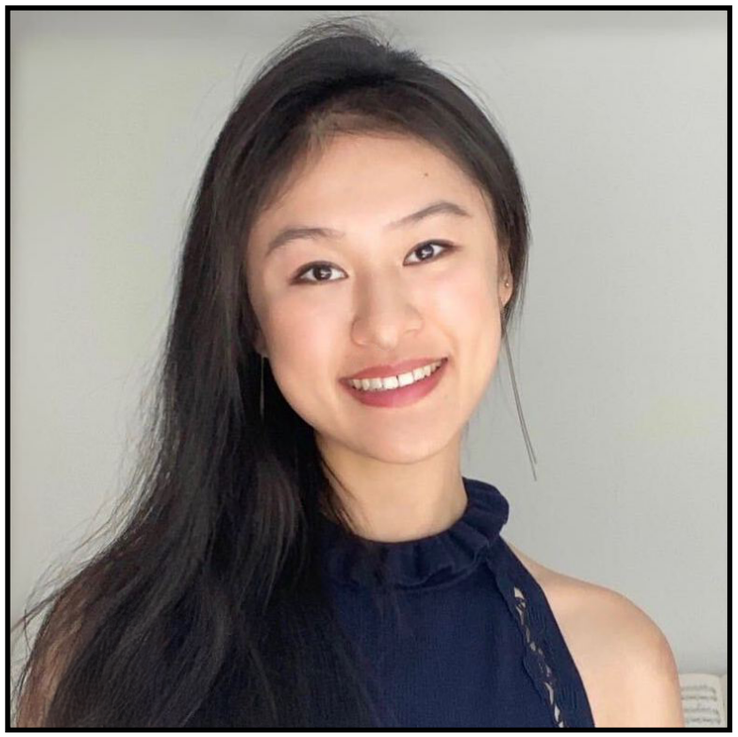

GSI: Sanjay Subramanian

GSI: Charlie Snell

GSI: Akshat Gupta

- 3rd year PhD advised by Gopala Anumanchipalli
- Born and grew up in Delhi, India
- BTech from IIT Mandi in Electrical Engineering
- MS from Carnegie Mellon University in Computer Engineering
- Research: Continual learning, Knowledge Editing, Interpretability

Questions We'll Explore

- What is "language technology"?
- How are language technologies built and deployed in practice?
- How do we know when these technologies work (or don't work)?
- How are language technologies impacting society?
- What are some open problems in natural language processing, and how can we address them?
- How can we think of language through a computational lens?

High-Level Topics

- What is natural language processing?
- Language and computation
- Language modeling: sequences, speech, and structure
- Modern LLM recipe
- Modern language technologies and applications

Course Meetings

- **Lecture:** TTh, 3:30-5pm, 2050 VLSB
- **Discussion sections will vary by week** — check Ed to see the schedule for each week (no discussion this week)
- **Exams** (except where otherwise arranged via DSP)
  - Midterm: in person during lecture on Tuesday, October 14
  - Final: Friday, Dec 19, 7-10pm, location TBD

Lectures

- Attend in person, unless there are exceptional circumstances — don’t come to class sick!
- Please arrive to lecture on time; we will begin exactly on Berkeley time
- You can bring food/drink to class, but keep noise to a minimum
- You can use any device to take notes, but please pay attention to the lecture and avoid distractions on your devices
- We will try to post slides in advance of the lectures, but cannot guarantee this
- We will record and upload lectures, but can’t guarantee they’ll be posted quickly — don’t use them as a replacement for attending lecture!

Prerequisites & Enrollment

- **You must have taken:** CS 182, CS 188, or CS 189
  - In Fall 2025: you must have received an A- or higher when taking this course
- **Will we have exceptions?** TBD
- **Required experience:** basic familiarity with neural networks, e.g., past experience with Pytorch and numpy
  - We will have no introductory lectures on programming frameworks used in the assignments
  - You can gain such experience through deep learning assignments in other ML courses in the department
- If you took CS 288, you can’t take this class for a grade
- **Undergrads must enroll in EECS 183**

Enrollment Form

- https://forms.gle/wXX6EMWDf4ntvuHp9

- Fill this out even if you are already enrolled! There are some questions relevant to everyone

Auditing

- All course materials (e.g. access to Ed forum, bCourses, Gradescope) will be made available to any currently-enrolled Berkeley student

- If you'd like to follow along with the course, you can request access by emailing us and forwarding proof current enrollment

- You can submit assignments to Gradescope and run the autograder

- Enrolled students are given priority in office hours and seating in lecture
  - For fire safety reasons, if the classroom is full and you're not officially enrolled, you'll need to leave

Grading

- **60% assignments**
  - 6 total; each comprises 10% of your total grade
  - For EECS 183 only:
    - The 10% of your grade coming from your lowest-scoring assignment is given “for free”
    - You only need to complete 5 of the 6 to receive full credit for this portion of your total grade

- **40% exams**
  - 20% midterm
  - 20% final

Exam Policies

- If you are not on a DSP plan and miss an exam, we cannot guarantee a makeup exam will be available
- If you must miss an exam, arrange this with us as soon as you are aware of the conflict
- For unanticipated emergencies causing you to miss an exam, contact us as soon as possible
- Do not discuss exams with anyone until grades have been announced

Grading

- Letter grades assignments will be no stricter than this
- A+ only on case-by-case basis
  - Substantial course engagement, e.g. on Ed
  - In EECS 183, you need to have completed all 6 assignments
- Course is 4 credits

A	[93, 100]
A-	[90, 93)
B+	[88, 90)
B	[83, 88)
B-	[80, 83)
C+	[78, 80)
C	[73, 77)
C-	[70, 73)
D+	[68, 70)
D	[63, 68)
D-	[60, 62)
F	[0, 60)

Late Work

- The following is the general policy; if you have a DSP plan we will work out a different policy with you

- Assignments are due **by 10pm PT** on the day of their deadline

- You will have a total of **14 late days** to use during the semester
  - This is meant to cover all possible circumstances of unanticipated delays: illness, unexpected travel, emergency
  - If you suddenly require additional late days, please arrange via DSP what would work best for your needs

- **Each late day used past the 14 will result in a 1% penalty on your final course grade**

Late Work

- You may not use more than **3 late days** per assignment without approval from us (granted in exceptional circumstances) — we will automatically close the Gradescope submission window 72 hours after the original deadline
- What we track
  - We won’t keep track of partial late days — one late day is spent per 24 hour window after the deadline, including if the submission occurs within (at the beginning of) the window
  - If you’re enrolled in EECS 183, you will not accrue late days for your lowest-scoring assignment
  - You can email us anytime to check what we have tracked, but please also keep track of your own use of late days

Multiple Submissions

Your submissions to Gradescope were...	We will grade...	Any late penalty?
before the deadline	your most recent submission	no
after the deadline	your most recent submission within 3 days after the dealdine	based on your most recent submission
both before and after the deadline	your most recent submission before the deadline, unless you tell us not to	no, unless you ask us to grade something after the deadline

Multiple Submissions

- If you'd like us to grade something else (an earlier or later version of the assignment), you must inform us by email within 3 days after the original deadline

- If you would like us to grade a version that was submitted after the original deadline, we will apply a late penalty according to when it was submitted

Academic Integrity

- We encourage study groups!

- But all assignments must be completed independently — you can discuss and work out high-level points, but don't share code with one another (we will check using code similarity tools)

- Use of generative AI
  - Generative AI tools, e.g. for code completion or for clarifying class material, can be useful
  - But we don't want you to rely on these tools to do the thinking for you, both for coding and writing!
  - **You must use proper attribution** including reporting the use of generative AI in completing assignments

Contacting Course Staff

- Course staff email, e.g. for enrollment questions, checking late day usage, etc: eecs183@lists.berkeley.edu

- For inquiries you want to direct only to instructors, precede email subjects with [NLP CLASS]:
  - Alane's email: suhr@berkeley.edu
  - Gopala's email: gopala@berkeley.edu

- Office hours might change at any time, temporarily or permanently, but we'll announce any changes ahead of time

- Course webpage: https://cal-nlp-class.github.io/fa25/

Instructor Office Hours

- In person, during the half hour after each lecture, right after class
- **Purpose:**
  - Clarification on course materials, discussion of course and research topics
  - We can’t debug your code or provide advice on assignment solutions

GSI Office Hours

- On Zoom, open without a waiting room, screen sharing disabled
- **Purpose:**
  - In addition to instructor office hours, GSIs can provide advice on working through assignments
  - GSIs will not provide debugging support and will not read your code
  - However, you can discuss implementation

GSI Office Hours

- Monday 4-5pm
- Tuesday 9-10am
- Wednesday 9-10am
- Subject to change! Zoom link will be available via Ed

Ed Forum

- Ed forum: https://edstem.org/us/join/nDkqFH
  - We'll send announcements here, and expect you to be subscribed if you are enrolled
  - Also includes links to Gradescope
  - Use it for asking questions to the GSIs or your peers!
- Ed policies
  - Ensure content you post is accessible (best practices)
  - Cite any sources you use (academic integrity applies here too)
  - No minimum engagement requirement, but will be taken into account when assigning A+

Ed Forum

- Ed policies (cont.)
  - Please make sure your question hasn’t been asked before
  - Please use the right tags for your questions
  - Don’t post any answers, solutions, or unreasonably long code snippets (links to documentation and very short code snippets used as examples are fine)
  - Don’t use Ed for administrative matters; instead, contact the course staff

Course Feedback

- This is a new class and we want to know how we are doing!
- We'll set up anonymous feedback forms 2-3 times over the semester so you can provide course staff feedback
- Please don't wait until the standard course evals to bring up any concerns that you have!

Accommodations

- Disability Support
  - If you need a disability-related accommodation in this class, contact Disabled Students Program (DSP / (510) 642-0518 / website: http://dsp.berkeley.edu) to prepare an accommodation letter
  - Accommodation letters are necessary for us to grant an accommodation

- If you need to reschedule an exam for religious reasons, inform us by the second week of the course