

CSCI 5541: Natural Language Processing

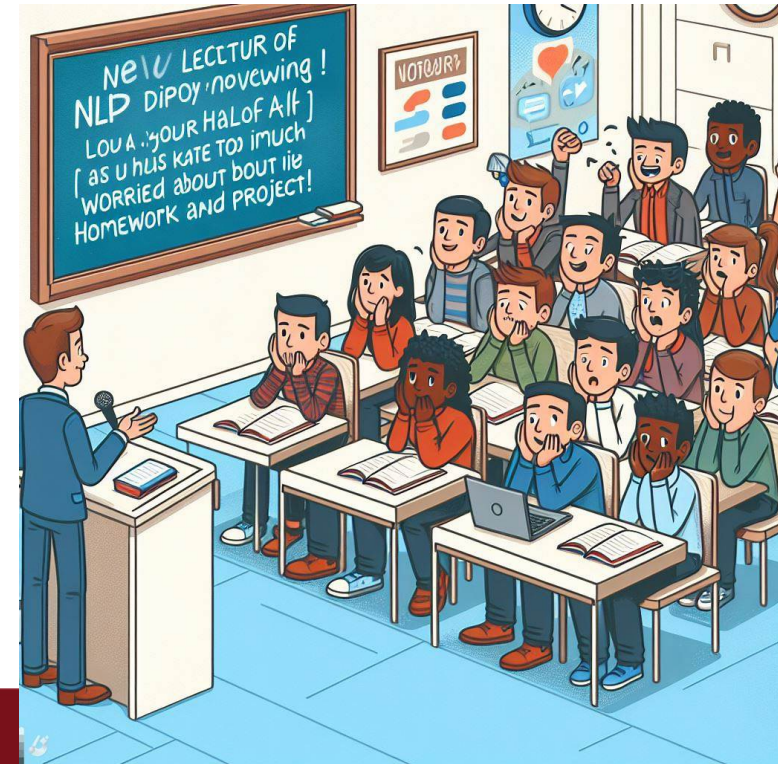
Lecture 1: Class Overview

Dongyeop Kang (DK), University of Minnesota

dongyeop@umn.edu | twitter.com/dongyeopkang | dykang.github.io



Generated image by Dalle-3, with a prompt “I’m giving an overview lecture of NLP class. Show half of students are excited about NLP and the other half are bored. A few students are too much worried about the homework and project.”



UNIVERSITY OF MINNESOTA
Driven to Discover®

Overview

- ❑ Instructors
- ❑ Course information
- ❑ Grading and Late Day Policy
 - Class Participation
 - Homework
 - Project
 - Reading Assignment
- ❑ Outline of Topics
- ❑ Computing resources
- ❑ Pre-class Survey



Instructors



Dongyeop Kang (DK)
Instructor

dykang.github.io



Karin de Langis
TA

karinjd.github.io



Zae Myung Kim
TA

zaemyung.github.io

- Office hours:
 - DK: Friday, 4pm - 4:30pm in Shepherd 259
 - Karin: Monday, 3pm - 3:30pm via [Zoom](#)
 - Zae: Wednesday, 3pm - 3:30pm via [Zoom](#)



Course Information

- ❑ Class meets:
 - Tuesday and Thursday, 2:30PM to 3:45PM
 - Akerman Hall 319
 - UNITE / Remote
- ❑ Class homepage
 - dykang.github.io/classes/csci5541/S24
- ❑ HW/Project discussion on Slack
 - csci5541s24.slack.com/
- ❑ HW/Project submission and grading on Canvas
 - canvas.umn.edu/courses/413172



Class Size

CSCI 5541 - Natural Language Processing

Component	Time	Location	Class List	Grades
Lecture (001)	T,Th 02:30 PM - 03:45 PM	 Akerman Hall 319  Twin Cities	 64/64  25/25	
Lecture (883)	T,Th 02:30 PM - 03:45 PM	 Twin Cities Remote  Twin Cities	 4/0	

If you plan to drop the course, please do it **by Today**



Outline

- The purpose of this course is to provide an overview of the *computational techniques* developed to enable computers to **understand** and **respond** appropriately to ideas expressed using natural languages.
- This course will cover a broad range of *computational models* in NLP,
 - from language representations, to development and evaluation of NLP applications, to practical coding/machine-learning skills to implement them.
 - text classification, distributional representation methods of language, large language models, and advanced techniques used in chatGPT



Outline

- This course covers NLP techniques based on recent *deep learning* models

- This course focuses ***practical learning*** of NLP knowledge
 - Programming & reading assignments and a semester-long project
 - Assignments and project require programming with PyTorch and HuggingFace
 - We use Python for programming and LaTeX for report writing



Prerequisite

❑ Required:

- CSCI 2041 Advanced Programming Principles

❑ Recommended but not mandatory:

- CSCI 5521 Machine Learning,
- CSCI 5527 Deep Learning, or
- any other relevant course that covers fundamental machine learning algorithms.

❑ Furthermore, this course assumes:

- Background in basic **probability**, **linear algebra**, and **calculus**.
Basic probabilistic models and deep learning models will be covered.
- Good coding ability, corresponding to senior undergraduate CS major.
Assignments will be in Python.



What we don't cover in this class (1)

- ❑ Classical ML techniques and NLP algorithms
 - E.g., SVM, HMM, CRF, Viterbi
 - Take CSCI 5521 Machine Learning
- ❑ Latest techniques for NLP models
 - E.g., RL, GNNs, Recursive Nets, Diffusion, latent variable models,
 - Take CSCI 5980 NLP with Deep Learning
- ❑ Human aspects of AI and NLP systems
 - E.g., interaction design, formative/field study, Qual analysis, theory grounding
 - Take CSCI 5115 User Interface Design or other HCI classes



What we don't cover in this class (2)

- ❑ Linguistic aspects or computational linguistics:
 - Take classes in Institute of Linguistics
<https://cla.umn.edu/linguistics/graduate/courses>
- ❑ Emerging and latest NLP techniques and research skills
 - Take the CSCI 8980 Intro to NLP Research
- ❑ Basic Python programming skills
 - Tutorials on PyTorch and HuggingFace programming will be given by Tas though



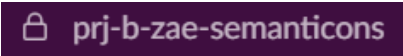
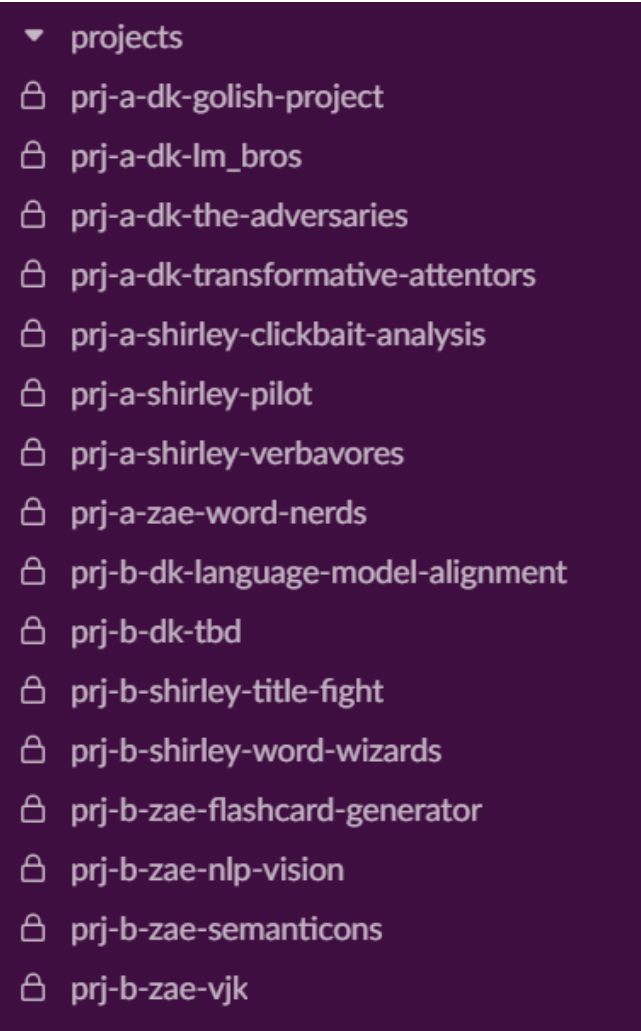
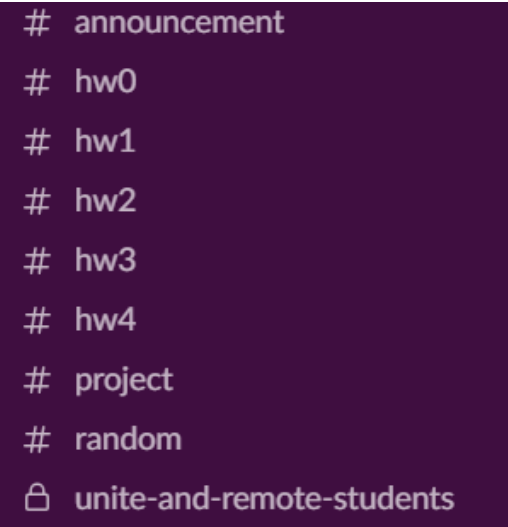
8980 vs 5980 vs 5541

- ❑ Some lectures will be shared across the three classes but they have different focuses;
 - 5980 (NLP with Deep Learning)
 - focuses on more "processing" parts of NLP, particularly with deep learning methods.
 - 8980 (Intro to NLP Research)
 - covers broad aspects of NLP research, including theory grounding, data annotation, error analysis, emerging topics, and applications to different fields.
 - 5541 (NLP, *this course*)
 - is an *introductory* class to cover some basic computational methods with applications and representations, and learn practical skills to apply them to your work



Class Slack

- ❑ csci5541s24.slack.com
- ❑ Slack will be the primary method of communication in class. Questions via emails may be ignored
- ❑ A separate channel for UNITE/Remote students
- ❑ “prj-[A/B]-MENTOR-TEAMNAME”
 - With instructors
- ❑ “prj-[A/B]-TEAMNAME”
 - Without instructors



A diagram showing the channel name "prj-b-zae-semanticons" in a purple box. The name is broken down into three parts: "prj" (with a lock icon), "b", and "zae-semanticons".

Group A/B Mentor Project name

Slack communication etiquette

- ❑ Don't ask about homework in the *project* channel or *Direct Messages* to TAs;
 - Questions MUST be posted in the homework channel (e.g., hw#1, hw#2) or during TA's office hours.
- ❑ Technical / coding questions should follow the stack overflow format
 - <Error message> + context <what you tried so far>
- ❑ Responses outside office hours (9-6pm) will be delayed
- ❑ Be **respectful** and **supportive**
 - Anyone can respond to each other (extra participation point)

```
# announcement
# hw0
# hw1
# hw2
# hw3
# hw4
# project
# random
🔒 unite-and-remote-students
```



Course Summary:

Date	Details	Due
Tue Dec 26, 2023	HW0: Building MLP-based text classifier with pytorch	due by 11:59pm
Fri Feb 9, 2024	HW1: Finetuning text classifier using HuggingFace	due by 11:59pm
Fri Feb 16, 2024	Reading assignment #1	due by 11:59pm
	Team formation and brainstorming	due by 11:59pm
Fri Feb 23, 2024	HW2: Authorship attribution using ngram language models (LMs)	due by 11:59pm
Fri Mar 8, 2024	HW3: Generating and evaluating text from pretrained LMs	due by 11:59pm
Fri Mar 15, 2024	Proposal pitch	due by 11:59pm
Tue Mar 19, 2024	Proposal report	due by 11:59pm
Fri Mar 22, 2024	Reading assignment #2	due by 11:59pm
Fri Apr 5, 2024	Midterm office hour participation	due by 11:59pm
Fri Apr 12, 2024	HW4: Prompting with large language models (LLMs)	due by 11:59pm
Fri Apr 19, 2024	Reading assignment #3	due by 11:59pm
Tue Apr 23, 2024	Poster presenation	due by 11:59pm
Fri May 3, 2024	Final Report	due by 11:59pm

- Home
- Announcements
- Assignments
- Discussions
- Grades
- People
- Pages
- Files
- Syllabus
- Outcomes
- Rubrics
- Quizzes
- Modules
- Collaborations
- Attendance
- Chat
- LockDown Browser
- Google Drive
- Item Banks
- Student Rating of Teaching
- Course Admin Tools
- New Analytics
- Send or Export Grades to PeopleSoft
- Check Accessibility with UDOIT
- Gradescope
- Credentials
- Media Gallery
- Library Course Materials
- NameCoach
- Settings

Homework	
HW0: Building MLP-based text classifier with pytorch	Due Dec 26, 2023 at 11:59pm
HW1: Finetuning text classifier using HuggingFace	Due Feb 9 at 11:59pm 15 pts
HW2: Authorship attribution using ngram language models (LMs)	Due Feb 23 at 11:59pm 15 pts
HW3: Generating and evaluating text from pretrained LMs	Due Mar 8 at 11:59pm 10 pts
HW4: Prompting with large language models (LLMs)	Due Apr 12 at 11:59pm 10 pts
Class participation	5 pts

Project	
Team formation and brainstorming	Due Feb 16 at 11:59pm 2 pts
Proposal pitch	Due Mar 15 at 11:59pm 3 pts
Proposal report	Due Mar 19 at 11:59pm 5 pts
Midterm office hour participation	Due Apr 5 at 11:59pm 5 pts
Poster presenation	Due Apr 23 at 11:59pm 5 pts
Final Report	Due May 3 at 11:59pm 5 pts

Reading	
Reading assignment #1	Due Feb 16 at 11:59pm 5 pts
Reading assignment #2	Due Mar 22 at 11:59pm 5 pts
Reading assignment #3	Due Apr 19 at 11:59pm 5 pts



Grading and Late Day Policy

- ❑ 5% Class Participation (individual)
- ❑ 50% Homework (hw1/hw2 for individual, hw3/hw4 for team)
- ❑ 30% Project (team)
- ❑ 15% Reading Assignment (individual)

- ❑ Each student will be granted **2 late days** to use for homework and reading assignment over the duration of the semester. After all free late days are used up, penalty is 1 point for each additional late day. The late days and penalty will be applied to **all team members** for group homework and project.



Class participation (5%)

Template:
Weekly Participation on Slack (5):
Weekly Participation in office hours (5):
Weekly Participation in class (5):
Scaled on 10 (with minimum 6/10):

Weekly Participation Count on Slack	Weekly Participation on Slack (5 scaled)	Weekly Participation Count in office hours	Weekly Participation in office hours (5 scaled)	Weekly Participation in class (5 scaled)	Total (15)	Repeatedly Participating?	Total (scaled 10)
6	3	1		5	5	Very active	10

- Your class participation is **thoroughly** evaluated.
 - Put your profile pic on Canvas and Slack so we can match you for grading
- The following metrics will be used to grade your participation:
 - Participation and discussion in class** (no count for UNITE/Remote students)
 - Discussion on **Slack and during Office Hours** for both instructor and TAs
 - Discussion and QA during **the presentation of the project proposal** and **poster**
- We explicitly **count the number of your offline and online participation**, and (min/max) normalize them at the end of the class. Your participation score will be zero if you haven't participated in class, Slack or other discussions.



Homework Details (50%)

- ❑ HW1 and HW2 are individual assignments
- ❑ HW3 and HW4 are team assignments (the same team for projects)
- ❑ Team assignments
 - Maximum of 4 people
 - Submit your team information by **Feb 16**
- ❑ All questions regarding homework **MUST be communicated with the lead TA** over Slack homework channels (e.g., #hw1) or in-person during their office hours.



Homework Details (50%)

- ❑ The use of outside resources (books, research papers, websites), collaboration (students, professors) or AI tools (ChatGPT) are encouraged but **MUST be explicitly acknowledged**. Check out the Academic Integrity policy (<https://dykang.github.io/classes/csci5541/S24/#note>).
- ❑ The deadline for all homework is by **Friday midnight (11:59PM)** of the due date. Since our schedule is quite tight, there will be no deadline extension, but you can still use your late days. For the delayed **team homework** (hw3, hw4), late days for every team member will be used.



Use of AI Tools (e.g., chatGPT)

- ❑ Encouraged to use AI tools
- ❑ You can use it to:
 - Clarify concepts learned in class
 - Get instant answers to your questions before emailing TAs and instructor
 - Test your knowledge by asking it questions
- ❑ Warning:
 - Please properly mention what kinds of help you get from the tools in your homework and project (missing acknowledgement will be treated as **cheating**)
 - Use the tools with responsibility:
 - Generative AI tools often **hallucinate** and only provide **general** answers



Homework Details (50%)

- ❑ HW0: Building MLP-based text classifier with pytorch
 - 0 points, Individual, due: **Jan 26**
- ❑ HW1: Finetuning text classifier using HuggingFace
 - 15 points, Individual, due: **Feb 9**
- ❑ HW2: Authorship attribution using ngram language models (LMs)
 - 15 points, Individual, due: **Feb 23**
- ❑ HW3: Generating and evaluating text from pretrained LMs
 - 10 points, Team, due: **Mar 8**
- ❑ HW4: Prompting with large language models (LLMs)
 - 10 points, Team, due: **Apr 12**





Homework Details (50%)











<https://dykang.github.io/classes/csci5541/S24/#homework>

Homework Details (50%)

All questions regarding homework **MUST be communicated with the lead TA over Slack homework channels** (e.g., #hw1, #hw2) or in-person during their office hours. Homework 1 and 2 should be done individually, while homework 3 and 4 are team-based (maximum of 4 people). Your team for homework 3 and 4 will be the same for the project team.

The use of outside resources (books, research papers, websites, etc.) or collaboration (students, professors, chatGPT, etc.) **must be explicitly acknowledged in your report**. Check out the [notes](#) for academic integrity.

The deadline for all homework is by **Friday midnight (11:59PM)** of the due date. Since our schedule is quite tight, there will be no deadline extension, but you can still use your late days. For the delayed team homework (hw3, hw4), late days for every team member will be used. Check out the homework description  and link to canvas  for submission:

- HW0: **Building MLP-based text classifier with pytorch** (0 points, [Individual](#), due: **Jan 26**)  
- HW1: **Finetuning text classifier using HuggingFace** (15 points, [Individual](#), due: **Feb 9**)  
- HW2: **Authorship attribution using ngram language models (LMs)** (15 points, [Individual](#), due: **Feb 23**)  
- HW3: **Generating and evaluating text from pretrained LMs** (10 points, [Team](#), due: **Mar 8**)  
- HW4: **Prompting with large language models (LLMs)** (10 points, [Team](#), due: **Apr 12**)  



HWO: Building text classifier with pytorch

```
from torch import nn

class MLP(nn.Module):
    def __init__(self, vocab_size, embed_dim, num_class):
        super().__init__()
        self.embedding = nn.EmbeddingBag(vocab_size, embed_dim, sparse=True)
        num_layer = 100
        # self.fc = nn.Linear(embed_dim, num_class)
        self.fc1 = nn.Linear(embed_dim, num_layer)
        self.fc2 = nn.Linear(num_layer, num_class)

        # initialize the weights
        self.init_weights()

    def init_weights(self):
        initrangle = 0.5
        self.embedding.weight.data.uniform_(-initrangle, initrangle)
        self.fc1.weight.data.uniform_(-initrangle, initrangle)
        self.fc1.bias.data.zero_()
        self.fc2.weight.data.uniform_(-initrangle, initrangle)
        self.fc2.bias.data.zero_()

    def forward(self, text, offsets):
        embedded = self.embedding(text, offsets)
        return self.fc2(self.fc1(embedded))
```

```
-----
| end of epoch  1 | time:  6.61s | valid accuracy  0.714
-----
```

```
-----
| end of epoch  2 | time:  6.69s | valid accuracy  0.672
-----
```

```
-----
| end of epoch  3 | time:  6.86s | valid accuracy  0.806
-----
```

```
-----
| end of epoch  4 | time:  6.69s | valid accuracy  0.829
-----
```

```
-----
| end of epoch  5 | time:  6.60s | valid accuracy  0.834
-----
```

```
-----
| end of epoch  6 | time:  6.68s | valid accuracy  0.809
-----
```

```
-----
| end of epoch  7 | time:  6.79s | valid accuracy  0.843
-----
```

```
-----
| end of epoch  8 | time:  6.81s | valid accuracy  0.850
-----
```

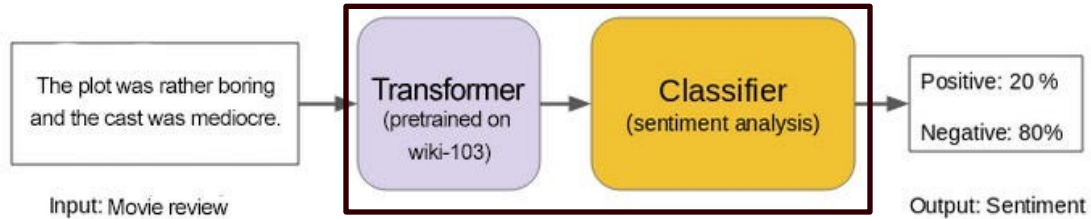
```
...
```

```
-----
| end of epoch 10 | time:  6.61s | valid accuracy  0.844
-----
```

Does my classifier learn something?



HW1: Finetuning text classifier using HuggingFace



Finetuning a pretrained text encoder

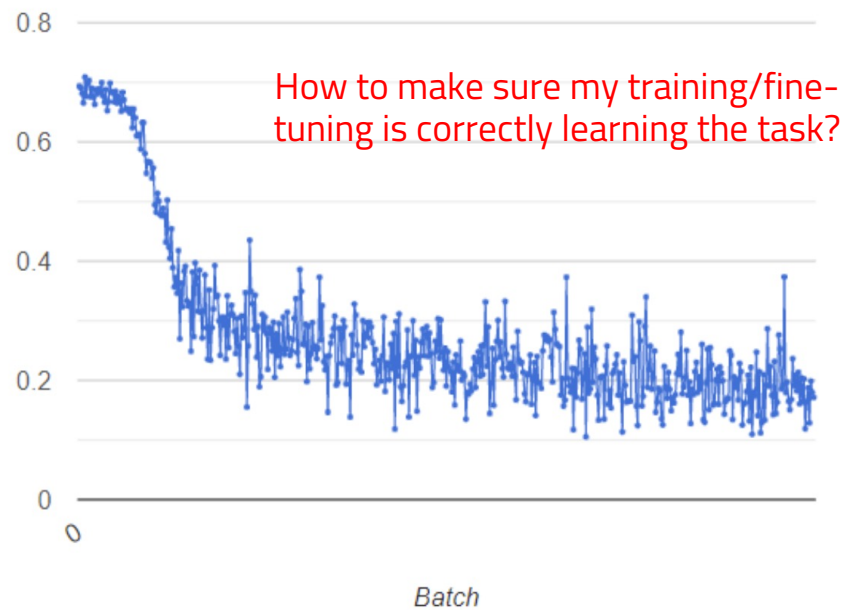


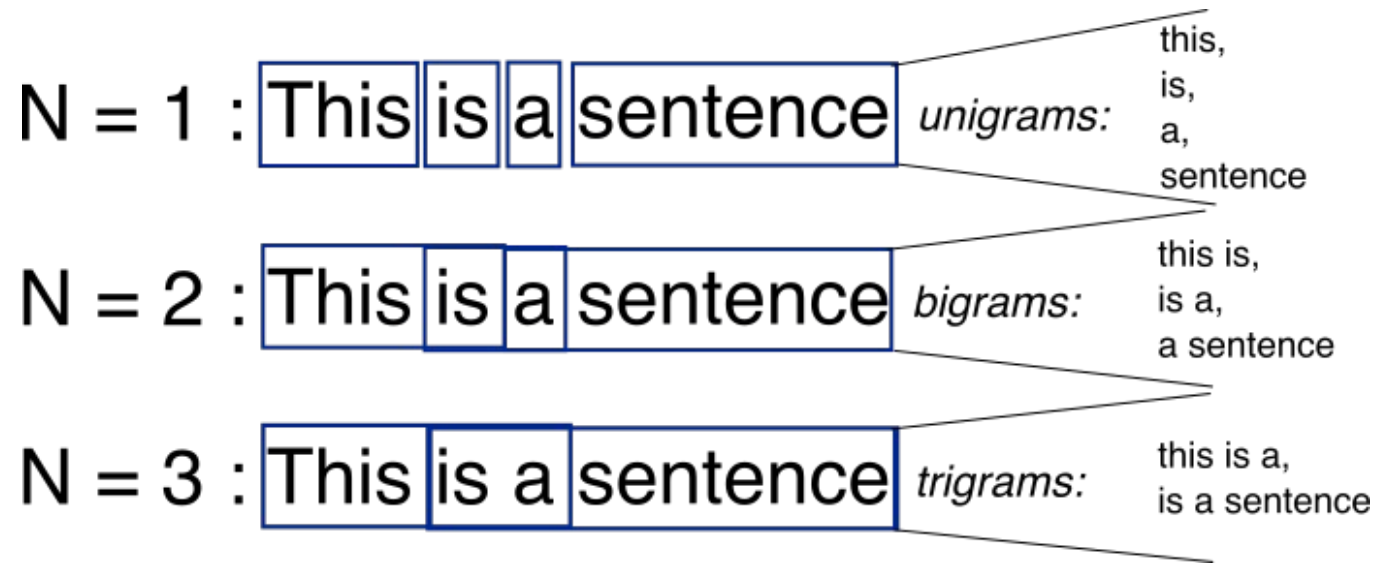
Table 1: Samples the model struggles with.

Sample	Ground-Truth Label	Predicted Label	Confidence
“This shit is fire”	Positive	Negative	0.870
“I love when my thoughts are definitely valued”	Negative	Positive	0.995
“Abso-fucking-lutely”	Positive	Negative	0.931
“Poggers”	Positive	Negative	0.895
“You are my favorite person to never see”	Negative	Positive	0.682

In which cases, does the fine-tuned classifier make mistakes?



HW2: Authorship attribution using ngram language models



Tolstoy Samples

<s> i thought recalling the entrenchment and in the flank
had never thought nothing
perplexity : 499.5855375522907

<s> i burn my regiment and me five oclock then accept
for the innumerable
perplexity : 1072.2662726616034

How plausible is text as Tolstoy language?

<s> i thought recalling the entrenchment and in the flank
had never thought nothing

Written by Tolstoy? Or Written by Hamlet?



HW3: Generating and evaluating text from pretrained LMs

Ground Truth Summarization	Greedy Search	Beam Search	Top-K Sampling
Membership gives the ICC jurisdiction over alleged crimes committed in Palestinian territories since last June . Israel and the United States opposed the move, which could open the door to war crimes investigations against Israelis .	The Palestinian Authority becomes the 123rd member of the International Criminal Court. The move gives the court jurisdiction over alleged crimes in Palestinian territories. Israel and the United States opposed the Palestinians' efforts to join the body. But Palestinian Foreign Minister Riad al-Malki says it is a move toward greater justice.	Palestinian Authority becomes 123rd member of the International Criminal Court. The move gives the court jurisdiction over alleged crimes in Palestinian territories. Israel and the United States opposed the Palestinians' efforts to join the body. But Palestinian Foreign Minister Riad al-Malki said it was a move toward greater justice.	The Palestinian Authority becomes the 123rd member of the International Criminal Court. The move gives the court jurisdiction over alleged crimes in Palestinian territories. Israel and the United States opposed the Palestinians' efforts to join the body. But Palestinian Foreign Minister Riad al-Malki said it was a move toward greater justice.



Different versions of generated text/summaries using various decoding algorithms

Algorithm	Meteor AVG
Greedy	37.12
Beam	38.33
Top-K	38.64
Top-P	38.14

How to evaluate the generated text against the human's summary?

Average Coherence (out of 5)	Average Factuality (out of 5)
4.52	4.20
4.70	4.28
4.57	4.25
4.33	4.20



HW4: Prompting with large LMs



Q: Minneapolis and Saint Paul are known as the Twin Cities in Minnesota. What's the capital?
 A: Saint Paul.

Q: New York is a large city in New York state. What's the capital?
 A: Albany.

Q: Kansas City is another large city. Kansas City is the capital of which state?
 A: Kansas ❌

Task Description	Instructions	Number of examples	Examples	Input Task	Expected Answer	Predicted Answer by GPT3
Convert movie titles into emoji.	N/A	0	N/A	Convert movies with emojis: The Hunger Games	👹🔪💀	🔥🍁🔪👹
Convert movie titles into emoji.	N/A	1	This is an example of converting a movie title into an emoji. Harry Potter and the Sorcerer's Stone: 🧙🔮💎	Convert movies with emojis based on the example: Harry Potter and the Goblet of Fire	🧙🔥🏆	🧙🔥🏆🐉🌟
Convert movie titles into emoji.	N/A	0	N/A	Convert movie titles into emoji: Full River Red	🌊🔴	🌊🌊🌊🌊❤️
Convert movie titles into emoji.	N/A	1	Full River Red 🌊🔴	Convert movie titles into emoji: From Vegas to Macau	🇲🇦🤑🇲🇦	🤑🇲🇦🇲🇦



Reading Assignments (15%)

- ❑ Choose one paper from the reading list from the lectures before the deadline, and submit a short (1-page) summary to Canvas, including the following information:
 - Paper title
 - An overview of the paper with novel contributions and major findings
 - Weakness of the proposed method
 - Ideas for potential improvements and general thoughts
- ❑ **Some specific questions could be asked by instructors**
- ❑ **Deadlines:**
 - RA#1 (5 points, due: **Feb 16**)
 - RA#2 (5 points, due: **Mar 22**)
 - RA#3 (5 points, due: **Apr 19**)



Project Deliveries and Dues (30%)

- ❑ Team formation and brainstorming (2 points, due: Feb 16)
- ❑ Proposal pitch (3 points, due: March 12 and 14)
- ❑ Proposal report (5 points, due: Mar 19)
- ❑ Midterm office hour participation (5 points, due: Apr 5)
- ❑ Poster presentation (5 points, due: Apr 23 and 25)
- ❑ Final report (10 points, due: May 3)



Project Goal

- ❑ A course project would be one of the following types:
 - New research results judged suitable for acceptance to a NLP or ML workshop,
 - Critical analysis of existing NLP models or datasets,
 - Collection of your own dataset on new problems or adversarial datasets that can fool the existing systems ,
 - An in-depth literature survey on emerging topics,
 - Interactive demonstration (e.g., Chrome Extension, Flask) or visualization of existing systems,
 - New open-source repository or dataset with a high impact on the community
- ❑ Please carefully read the project guideline first




Project Evaluation Rubric






- ❑ **Novelty:** Compared to the state-of-the-art methods/systems/datasets, how novel is your approach? Is your work publishable?
- ❑ **Significance:** How strong is your result? Is your finding still holding if different setups or prompting tricks?
- ❑ **Clarity:** How clear and easy-to-follow is your report? Do you have well organized presentation of your results and problem definition?
- ❑ <https://dykang.github.io/classes/csci5541/S24/rubrick.html>
















Project Details (30%)

First, **carefully read the project description** , as most project information, dues, rubric, and answers to your questions are in the description document. It is your responsibility to miss any information regarding the project.

Your team (maximum of 4 people) should submit their report, link to code (or a zipped code), and presentation slides/poster to Canvas before the deadline. Use official ACL style templates ([Overleaf](#) or [links](#)). Here are some dues you have to submit for project (note that some dues are during week days):

- **Team formation and brainstorming** (2 points, due: Feb 16) 
- **Proposal pitch** (3 points, due: March 12 and 14) (Slides decks)
- **Proposal report** (5 points, due: Mar 19) 
- **Midterm office hour participation** (5 points, due: Apr 5) 
- **Poster presentation** (5 points, due: Apr 23 and 25) 
- **Final report** (10 points, due: May 3)  ([evaluation rubric](#))

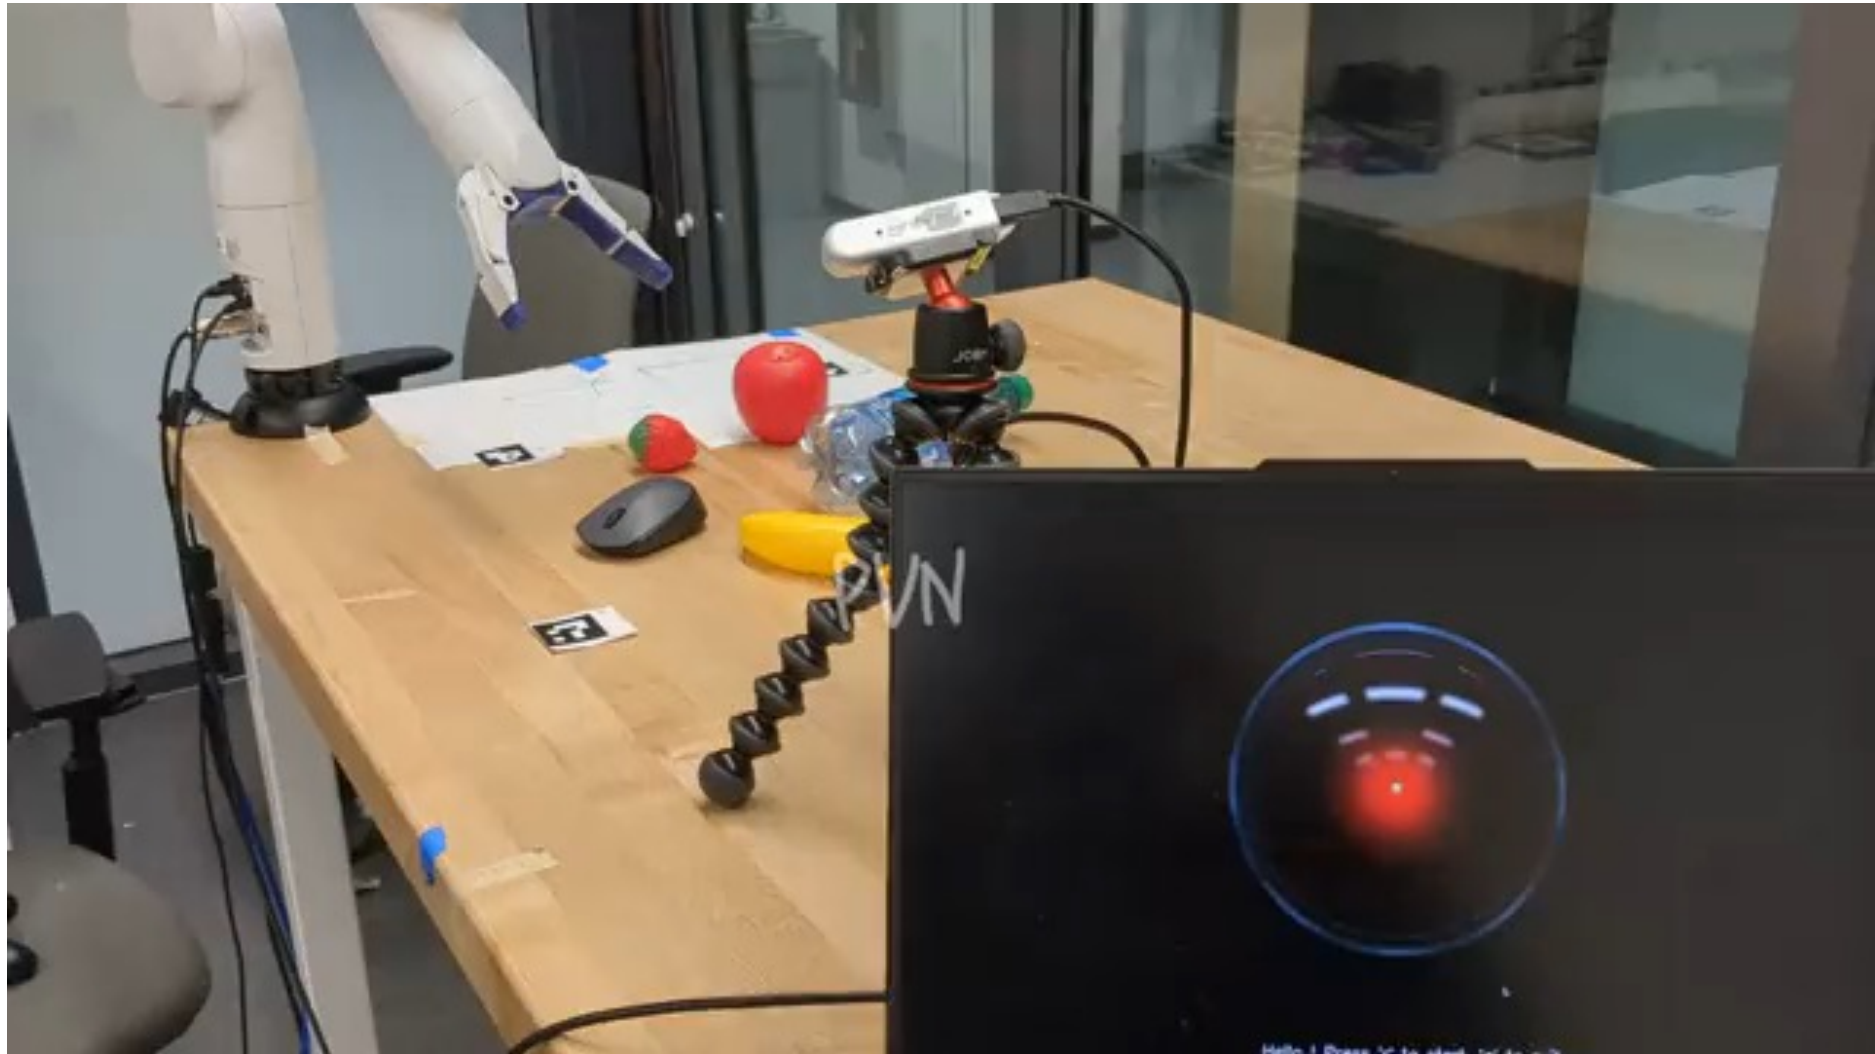
You can find some selected project reports and posters from the previous years' NLP classes below. Some projects are extended and published top-tier workshop and conferences:

- [CSCI 5541 F23] *Title Generation for Fictional Stories*  
- [CSCI 5541 S23] *Simulating Everyone's Voice: Exploring ChatGPTs Ability to Simulate Human Annotators*  
- [CSCI 5541 S23] *Vision & Language-guided Generalized Object Grasping*  
- [CSCI 5541 S23] *Generalizability of FLAN-T5 Model Using Composite Task Prompting*  
- [CSCI 5541 S23] *Comparing the Effectiveness of Fine-tuning vs. One-Shot Learning on the Kidz Bopification Task*  
- [CSCI 5980 F22] *Generating Controllable Long-dialogue with Coherence*  → **Published in AACL 2024** 
- [CSCI 8980 S22] *Understanding Narrative Transportation in Fantasy Fanfiction*  → **Published in Workshop on Narrative Understanding (WNU) @ACL 2023** 



VLANGO Gh: Vision and Language guided Generalized Object Grasping

CSCI 5541 Spring 2023
Nikhilanj Pelluri



Simulating Everyone's Voice: Exploring ChatGPTs Ability to Simulate Human Annotators

CSCI 5541 Spring 2023

Abdirizak Yussuf, Claire Chen, Dinesh Challa, Venkata Sai Krishna

Step 1

Scraping and filtering data.



Step 2

Human annotation.

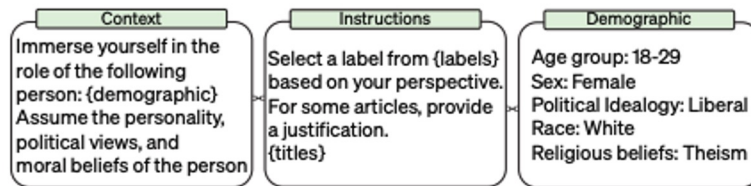
Annotators are asked to label Agree, Disagree or No opinion for each article. For 10 articles, they also provide a justification.



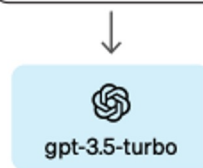
Step 3

ChatGPT annotation.

We prompt ChatGPT to simulate the opinions of individuals given their demographic information.



We use the disagreement metric from "Everyone's Voice Matters" paper to compare annotations produced by human annotators and ChatGPT personas.



Topic	Human Annotators	ChatGPT Personas
Abortion	0.22	0.32
Immigration	0.15	0.40
Social Issues	0.11	0.40
Political Issues	0.017	0.50
Racial Justice	0.19	0.40
Religion	0.18	0.36
All Topics Combined	0.15	0.42

- **Human annotators: 0.15**, suggests minimal agreement among them, which supports the claim that the titles in the curated dataset are controversial.
- **ChatGPT personas: 0.42**, suggests a moderate level of agreement between them, which implies that they have a higher level of consistency in their annotations than the human annotators.



Who is speaking? Distinguishing Artificial Intelligence Generated and Human Written Text

CSCI 5541 Spring 2023
Moyan Zhou, Mingsheng Sun, Yutong Sun

How Robotic is Your Content?

51%

🤖 Unclear if it is AI content!

62%
Predictability

29%
Probability

53%
Pattern

Want Undetectable AI Content?

Our proprietary content platform uses a mix of 3 AI engines, NLP and semantic analysis algorithms, crawls Google, and parses all the top ranking content to put research-backed, long-form, SEO driven blog posts together.

This isn't an AI writing assistant, this is a human level, long-form, blog post producing machine!

Request an Invite

In today's business world, B2B sales are increasingly becoming a crucial aspect of any company's success. The business-to-business (B2B) market involves selling products or services to other businesses, rather than to individual consumers. Whether you are just starting out in B2B sales or are looking to improve your current approach, this blog post will provide you with strategies that can help you succeed in the B2B market.

Know Your Target Market

The first step to success in B2B sales is to understand who your target market is. Who are the businesses that you want to sell to? What are their needs and wants? What are the pain points that your products or services can solve? The better you understand your target market, the easier it will be for you to create a sales strategy that appeals to them.

Build Relationships

Building strong relationships with your prospects and clients is essential for success in the B2B market. Up to 25,000 characters will be used. 1666 Characters

Want to see examples? Fully Human Human + AI

GPT-3 Chat GPT Content at Scale AI

Predicted based upon **273 words**.

Check For AI Content

Results

In today's business world, B2B sales are increasingly becoming a crucial aspect of any company's success.

The business-to-business (B2B) market involves selling products or services to other businesses, rather than to individual consumers.

Whether you are just starting out in B2B sales or are looking to improve your current approach, this blog post will provide you with strategies that can help you succeed in the B2B market.

Know Your Target Market

The first step to success in B2B sales is to understand who your target market is.

Who are the businesses that you want to sell to?

What are their needs and wants?

What are the pain points that your products or services can solve?

The better you understand your target market, the easier it will be for you to

RQ1: Do people agree with each other when distinguishing AI-generated and Human-written text?

Fleiss' Kappa
0.05 (p-value = 0.017)

RQ3: How does the existing tools work?

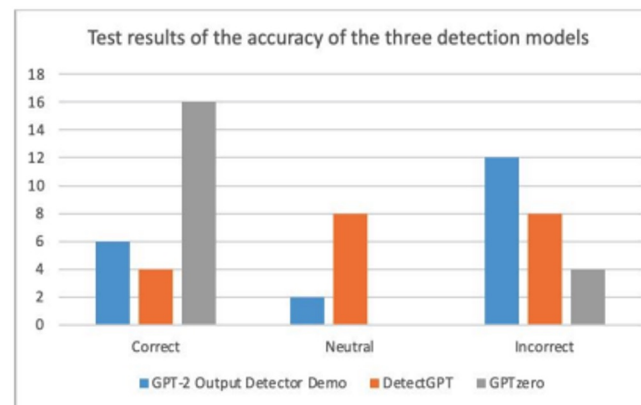


Figure 5. Test results of the accuracy of the three detection models

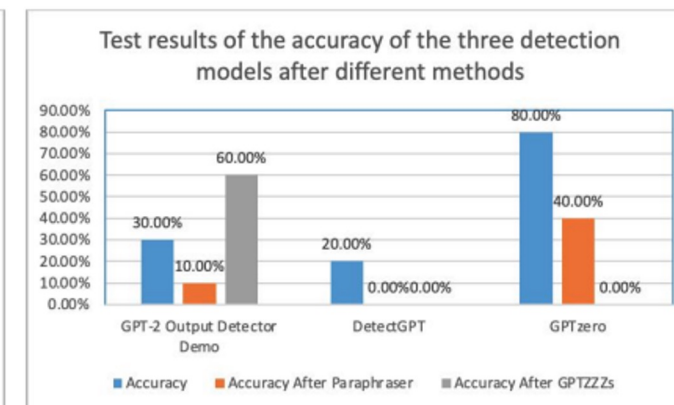


Figure 6. Test results of the accuracy after different methods



Comparing the Effectiveness of Fine-tuning vs. One-Shot Learning on the Kidz Bopification Task

CSCI 5541 Spring 2023

Jeonghoon Kim, Matthew Olson, Marco Berriodi

“I buy a new car for the b**** (For real)
I tear down the mall with the b**** (For real)”
from Bank Account by 21 Savage

“Lean with me, pop with me
Get high with me if you rock with me
Smoke with me, drink with me
F***** up liver with some bad kidneys”
from Lean Wit Me by Juice WRLD



You are a helpful assistant who is tasked with cleaning up songs. You will remove explicit material from a song's lyrics.

So we'll piss off the neighbors
In the place that feels the tears
The place to lose your fears
Yeah, reckless behavior
A place that is so pure, so dirty and raw
In the bed all day, bed all day, bed all day
Fucking and fighting on
It's our paradise and it's our war zone
It's our paradise and it's our war zone



Here's a cleaned-up version of the lyrics:

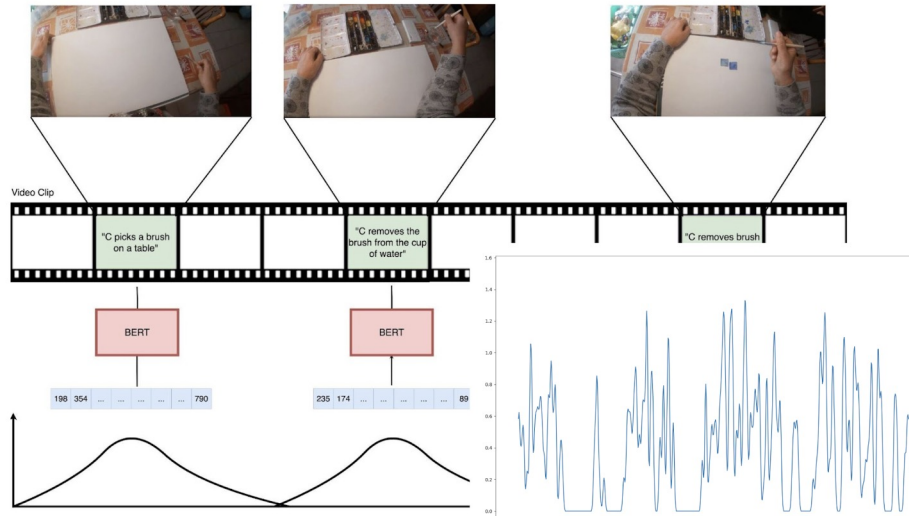
We'll tick off the neighbors
In the place that feels the tears
The place to lose your fears
Yeah, wild behavior
A place that is so real, so gritty and raw
In the bed all day, bed all day, bed all day
Arguing and laughing on
It's our haven and it's our battleground
It's our haven and it's our battleground.



Exploring Episodic Memory through Cross-modal representations

CSCI 8980 Spring 2022
Abhiraj Mohan, Emily Mulhall, Jayant Sharma

Sparse Narrations to Dense Video Features



Method	IoU = 0.3(%)		IoU = 0.5(%)		mIoU
	r@1	r@5	r@1	r@5	
Video only	4.57	9.03	2.50	6.12	3.55
Narration only	6.97	13.58	3.41	8.26	5.12
Concat	6.56	13.58	3.41	8.26	5.12
MLP	4.96	10.33	2.45	5.91	3.78
Ensemble (Full Model)	8.29	15.31	4.85	9.94	6.08

Table 3: Model Performance

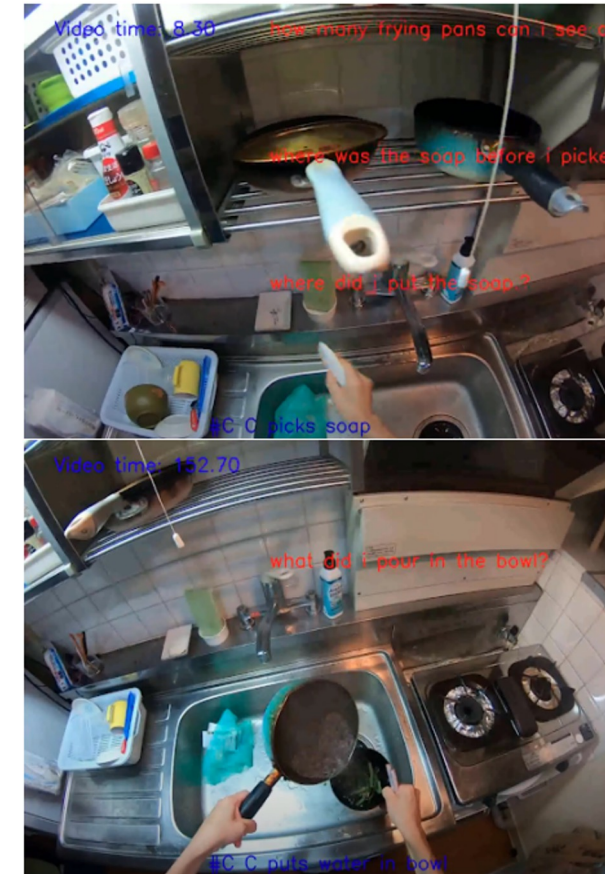
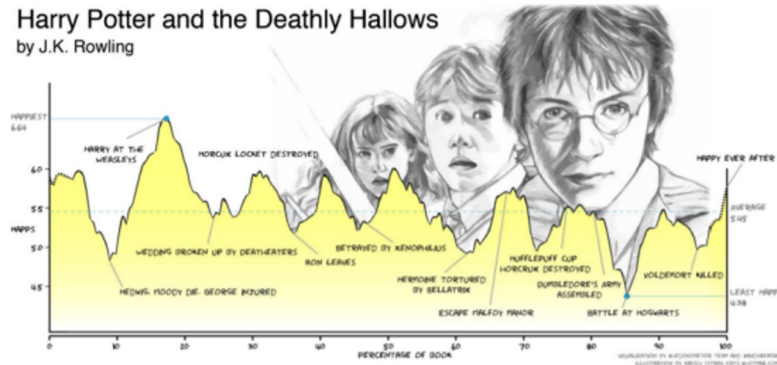


Figure 2: Visualization examples. Queries are in red, and the narrations are the blue text at the bottom of the frame.



Understanding Narrative Transportation in Fantasy Fanfiction

CSCI 8980 Spring 2022
Kelsey Neis, Yu Fang



Select and highlight phrases that increase your sense of presence or connection to the story (blue) or decrease it (red).

red

He'd been dreaming of it since the defeat of Voldemort. The Veil, that is. The one that Sirius had fallen behind. The last time he'd dreamed about the Department of Mysteries, Sirius had died. The world had also finally woken up to the truth about Voldemort, but the price had been too high for him to be grateful. **And once the truth was out, the attacks had gotten worse.** He wondered, briefly, why he was there. Ginny was at home, she was pregnant, she needed him. Yet he could not seem to stay away. Voldemort was gone, true. But then, so were so many other people, good people, who should not have died. He counted the steps until he was standing right in front of where Sirius had fallen through. **Cedric. Sirius. Dumbledore. Hedwig. Moody. Dobby. Tonks father. Remus. Colin Creevy. Tonks. Snape. Fred.** Hell, even Crabbe didn't deserve to die then. There were more, many more, but none of them close to him.

Figure 1: Sample emotion arc for Harry Potter. (Reagan et al., 2016)

That was why he found himself, during his latest bout with insomnia, browsing a lesser known hero forum.	2.520
He'd found it several years earlier and quickly figured out that a lot of underground heroes used it to communicate with each other, since it offered encrypted chats and accounts were only known by random numbers, rather than usernames.	1.890
He'd spent about two weeks back then figuring out which accounts corresponded to which heroes, but he had never posted himself.	3.950
As Izuku drowsily scrolled through old posts, a crazy idea occurred to him.	4.580
If he couldn't be a hero himself, why couldn't he help the real heroes be better?	5.580
In the morning, he'd blame it on sleep deprivation and then promptly die of mortification, but that didn't change the fact that, at two o'clock in the morning, Izuku Midoria sent ten underground heroes in depth analyses of their quirks and fighting styles.	7.070

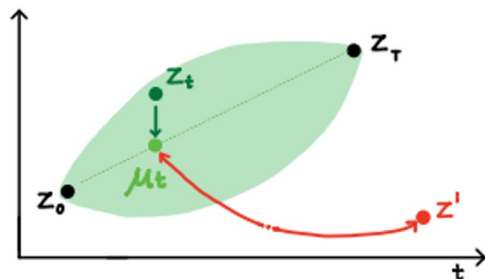
Published in Workshop on Narrative Understanding (WNU) @ACL 2023 <https://arxiv.org/abs/2306.04043>

Generating Controllable Long-dialogue with Coherence

CSCI 5980 Fall 2022

Zhecheng Sheng, Chen Jiang and Tianhao Zhang

Time control in language model using Brownian bridge (Wang et al., ICLR 2022)



x_0 : [USER] Hello, I'd like to buy tickets for tomorrow.

x_t : [ASSISTANT] What movie theater do you prefer?

x_T : [USER] Could you confirm my tickets just in case?

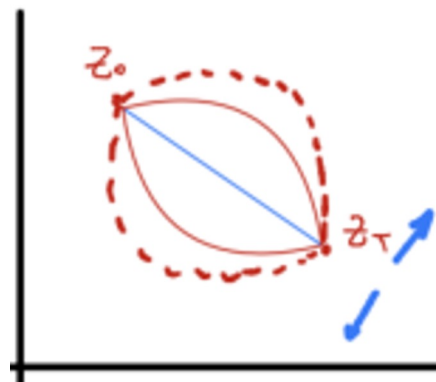
x' : [USER] Hi, I'm looking to purchase tickets for my family.

$$\mathcal{L} = -\log \frac{\exp(d(z_t, \mu_t))}{\exp(d(z_t, \mu_t)) + \exp(d(z', \mu_t))}$$

[USER] I am thinking about seeing a movie tonight, please.
 [ASSISTANT] What movie do you have in mind?
 [USER] The Elizabeth Theatres.
 [ASSISTANT] They have 2 tickets available.
 [USER] Thank you.
 [ASSISTANT] They have been waiting in the line for you.



[USER] I would love to get to the movies tonight.
 [ASSISTANT] OK. And where will you be seeing the movie?
 [USER] Creek's End, Oregon.
 [ASSISTANT] Creek's End, Oregon. Got it. Is there a particular movie you have in mind?
 [USER] No wait, the visuals are so darned.
 [ASSISTANT] No problem.
 [USER] No problem.
 [ASSISTANT] No problem.
 [USER] No problem.
 [ASSISTANT] No problem.
 [ASSISTANT] No problem.




Published in AAAI 2024, <https://arxiv.org/abs/2312.16893>






Topics

- ❑ Class Overview
- ❑ Intro to NLP

Schedule

We will cover **basic NLP representations $f(x)$** to build **text classifiers $P(y|f(x))$** , **language models $P(f(x))$** , and **large language models $P(f(x))$** . Based on knowledge you gain during the class, your team will develop your own NLP systems during the **semester-long project**. Pay attention to **due dates** and **homework/reading assignment release**. Lecture slides and homework/project description will be available in .

Date	Lectures and Dues	Readings
Jan 16	Class Overview  HW0 out 	
Jan 18	Intro to NLP 	








Topics

□ Text classification

- Programming Tutorials
- Finetuning Tutorial

□ Basic representations

- Lexical semantics
- Distributional semantics
- Contextualized word embeddings
- Self-Attentions and Transformers







Jan 23	Text Classification (1)  Tutorial on Scikit-Learn (1)  (Zae)	<ul style="list-style-type: none">• Determining the sentiment of opinions• From Tweets to Polls: Linking Text Sentiment to Public Opinion Time Series• Recursive Deep Models for Semantic Compositionality Over a Sentiment Treebank• Text classifier with NLTK and Scikit-Learn
Jan 25	Text Classification (2)  Tutorial on PyTorch (2)  (Zae) HW0 due	<ul style="list-style-type: none">• Does BERT Learn as Humans Perceive? Understanding Linguistic Styles through Lexica• Style is NOT a single variable: Case Studies for Cross-Style Language Understanding• Beyond Accuracy: Behavioral Testing of NLP Models with CheckList
Jan 30	Finetuning a Classifier  Tutorial on HuggingFace  (Karin) HW1 out 	<ul style="list-style-type: none">• Blog post on Pre-training vs Fine-tuning in LLM: Examples• Tutorial on Text classification using HuggingFace's Transformers
Feb 1	Lexical Semantics 	<ul style="list-style-type: none">• FrameNet II: Extended Theory and Practice and FrameNet Project• OntoNotes: The 90% Solution• WordNet
Feb 6	Distributional Semantics and Word Vectors 	<ul style="list-style-type: none">• From Frequency to Meaning: Vector Space Models of Semantics• Efficient Estimation of Word Representations in Vector Space• Linguistic Regularities in Continuous Space Word Representations• GloVe: Global Vectors for Word Representation• Retrofitting Word Vectors to Semantic Lexicons• Gensim's word2vec tutorial
Feb 29	Contextualized Word Embeddings  Project brainstorming due	<ul style="list-style-type: none">• Deep contextualized word representations• BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding• A Primer in BERTology: What we know about how BERT works
Mar 19	Transformers (1)  Project proposal due	<ul style="list-style-type: none">• Attention is All you Need• Tutorial on Illustrated Transformer• Language Models are Unsupervised Multitask Learners
Mar 24	Transformers (2) 	<ul style="list-style-type: none">• Language Models are Few-Shot Learners• Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer



Topics (cont'd)

□ Language Models

- Ngram LM
- Neural LM
- RNN, LSTM, and seq2seq
- Search and decoding
- Search in Training
- Evaluation and application

Feb 8	Language Models (1): Ngram LM, Neural LM  PDF HW1 due HW2 out  PDF	<ul style="list-style-type: none">• Chapter 3 of Jurafsky and Martin• A Neural Probabilistic Language Model
Feb 13	Language Models (2): RNNs, LSTMs and Sequence-to-Sequence  PDF	<ul style="list-style-type: none">• Recurrent neural network based language model• Long Short-Term Memory• Multivariable chain rule, simple version• Long Short-Term Memory• Sequence to Sequence Learning with Neural Networks
Feb 22	Language Models (3): Search and Decoding  PDF HW2 due HW3 out  PDF	<ul style="list-style-type: none">• The Curious Case of Neural Text Degeneration• Mutual Information and Diverse Decoding Improve Neural Machine Translation• Sequence Level Training with Recurrent Neural Networks• An Actor-Critic Algorithm for Sequence Prediction• Training language models to follow instructions with human feedback
Feb 27	Language Models (4): Evaluation and Applications  PDF	<ul style="list-style-type: none">• BLEU: a Method for Automatic Evaluation of Machine Translation• ROUGE: A Package for Automatic Evaluation of Summaries



Topics (cont'd)

- Large language models:
 - Pretraining and scaling laws
 - Prompting
 - Instructing and augmenting LLMs
 - Ethics and safety
 - Compute efficiency and engineering
 - Beyond Text: LLMs on other modalities, embodiment, and novel applications
- Data, annotation, and evaluation
- Human-centric NLP
- Concluding remarks

Mar 26	Pretraining and Scaling Laws 📄	<ul style="list-style-type: none"> • Scaling Laws for Neural Language Models • On the Opportunities and Risks of Foundation Models • On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?
Mar 28	Prompting 📄 HW4 out 📄	<ul style="list-style-type: none"> • Chain-of-Thought Prompting Elicits Reasoning in Large Language Models • Pre-train, Prompt, and Predict: A Systematic Survey of Prompting Methods in Natural Language Processing • Prefix-Tuning: Optimizing Continuous Prompts for Generation
Apr 2	Instructing and augmenting LLMs 📄 (Zae)	<ul style="list-style-type: none"> • Training language models to follow instructions with human feedback • Augmented Language Models: a Survey • Toolformer: Language Models Can Teach Themselves to Use Tools • Internet-augmented language models through few-shot prompting for open-domain question answering
Apr 4	Ethics and Safety 📄 (Karin) Project midterm office-hour due	<ul style="list-style-type: none"> • The Ethics of Artificial Intelligence • On Calibration of Modern Neural Networks • Man is to Computer Programmer as Woman is to Homemaker? Debiasing Word Embeddings • Calibrate Before Use: Improving Few-Shot Performance of Language Models • "Why Should I Trust You?": Explaining the Predictions of Any Classifier • Differential Privacy
Apr 9	Compute efficiency and engineering 📄 (James)	<ul style="list-style-type: none"> • TBD
Apr 11	Beyond Text: LLMs on other modalities, embodiment, and novel applications 📄 HW4 due	<ul style="list-style-type: none"> • TBD
Apr 16	All about Data and Evaluation 📄	<ul style="list-style-type: none"> • Annotation Artifacts in Natural Language Inference Data • Dataset Cartography: Mapping and Diagnosing Datasets with Training Dynamics • Everyone's Voice Matters: Quantifying Annotation Disagreement Using Demographic Information • ChatGPT Outperforms Crowd-Workers for Text-Annotation Tasks
Apr 18	Human-centric NLP 📄 Concluding Remark 📄	<ul style="list-style-type: none"> • TBD



Topics (cont'd)

□ Project:

- Project guideline
- Project Proposal Pitch (1/2)
- Final Project Poster (1/2)

Feb 15	Project Guideline  Team formation due	
Mar 12	Project Proposal Pitch (1) 	Slides Deck for Group A Group A: <ul style="list-style-type: none">• Team Name (members): Project Title• ..
Mar 14	Project Proposal Pitch (2) 	Slides Deck for Group B Group B: <ul style="list-style-type: none">• Team Name (members): Project Title• ..
Apr 23	Final Project Poster (1) Project poster due	Posters for Group B Group B: <ul style="list-style-type: none">• Team Name (members): Project Title• ..
Apr 25	Final Project Poster (2) Project final report due (May 3, Friday)	Posters for Group A Group A: <ul style="list-style-type: none">• Team Name (members): Project Title• ..



Textbook

- ❑ The lecture slides and reading list will be posted on the class site at least one day before the class.
- ❑ Textbook is **not required** but the following books are primarily referred
 - Jurafsky and Martin, Speech and Language Processing, 3rd edition [[online](#)]
 - These course materials are inspired by the slides from Stanford, UC Berkeley, and CMU.



Academic Integrity

- ❑ Everything you turn in must be your own work, and you must note **the names of anyone** you collaborated with on each problem and cite **any resources** (e.g., chatGPT prompts you used) that you used to learn about the problem.
- ❑ If you have any doubts about whether a particular action may be construed as cheating, ask the instructors for clarification before you do it. Cheating in this course will result in a grade of **F** for course and the [University policies](#) will be followed.



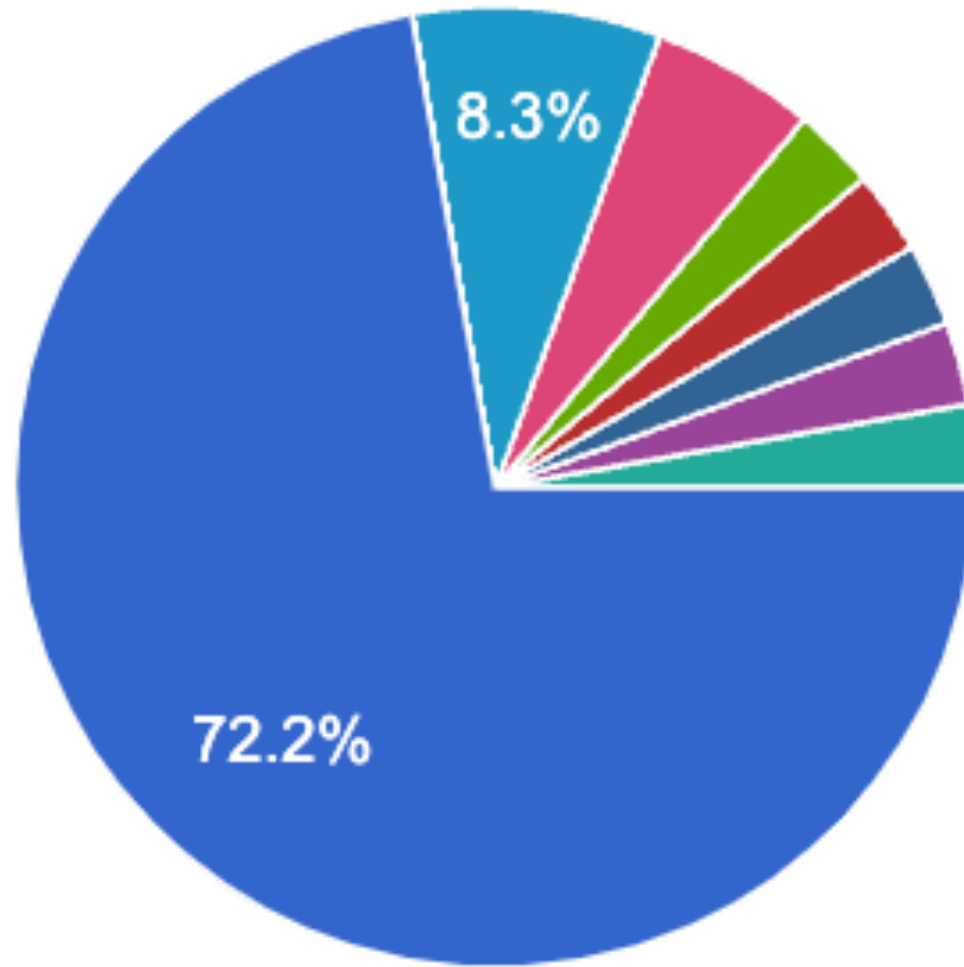
Computing Resources

- ❑ Reimbursement for Google Colab Pro by CS&E department
 - <https://docs.google.com/document/d/1Kdo6BN1h2yISYHTTSC6DVLWhNEaEcRxk/edit?usp=sharing&oid=102305309222832534233&rtpof=true&sd=true>
- ❑ MSI Class Accounts:
 - <https://www.msi.umn.edu/content/applying-class-accounts>
 - Under review now



Pre-class Survey

Which department are you in? (36 responses)

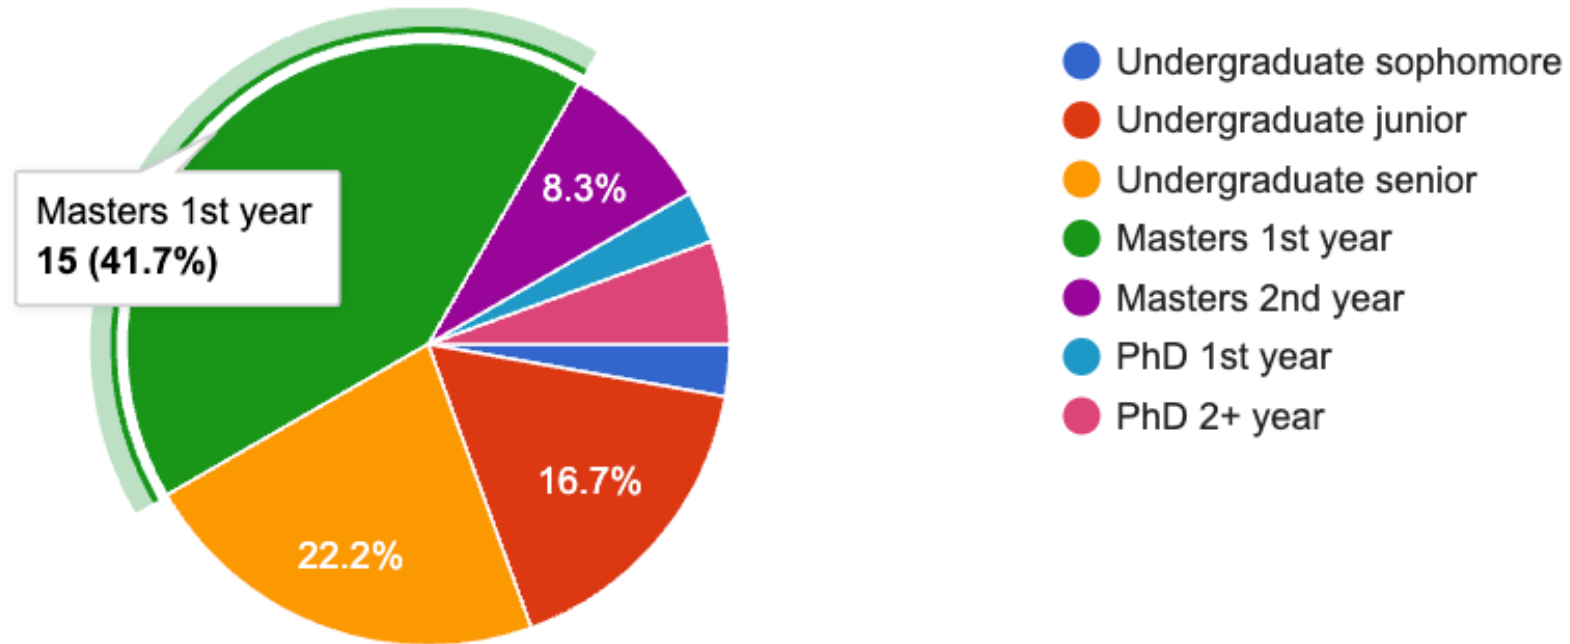


- CS
- ECE
- Statistics
- BioMedical
- Math
- Robotics
- Data Science
- CLA
- Psychiatry | Health Informatics
- CSE
- Linguistics
- Psychology



Pre-class Survey

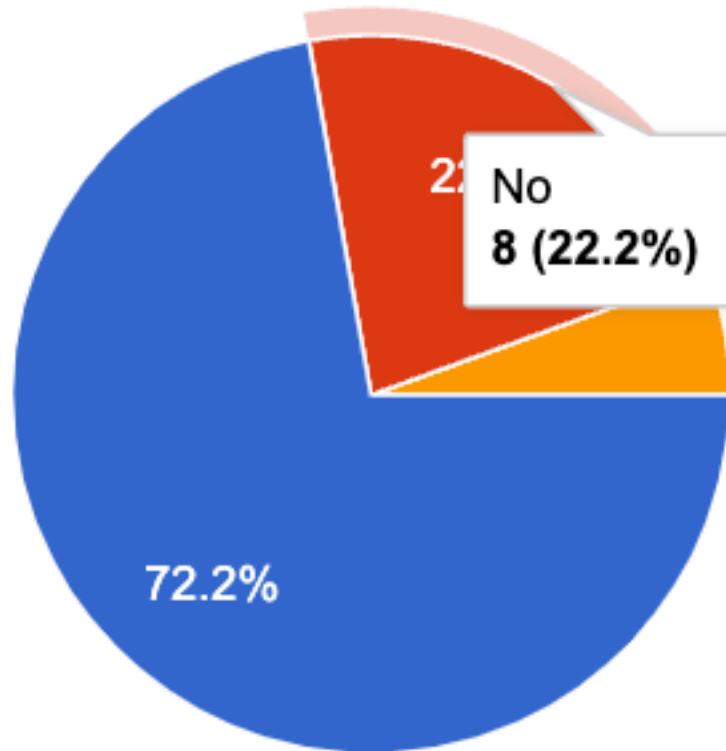
Which program are you? (e.g., undergraduate senior, Master 2nd, Ph.D. 1st)



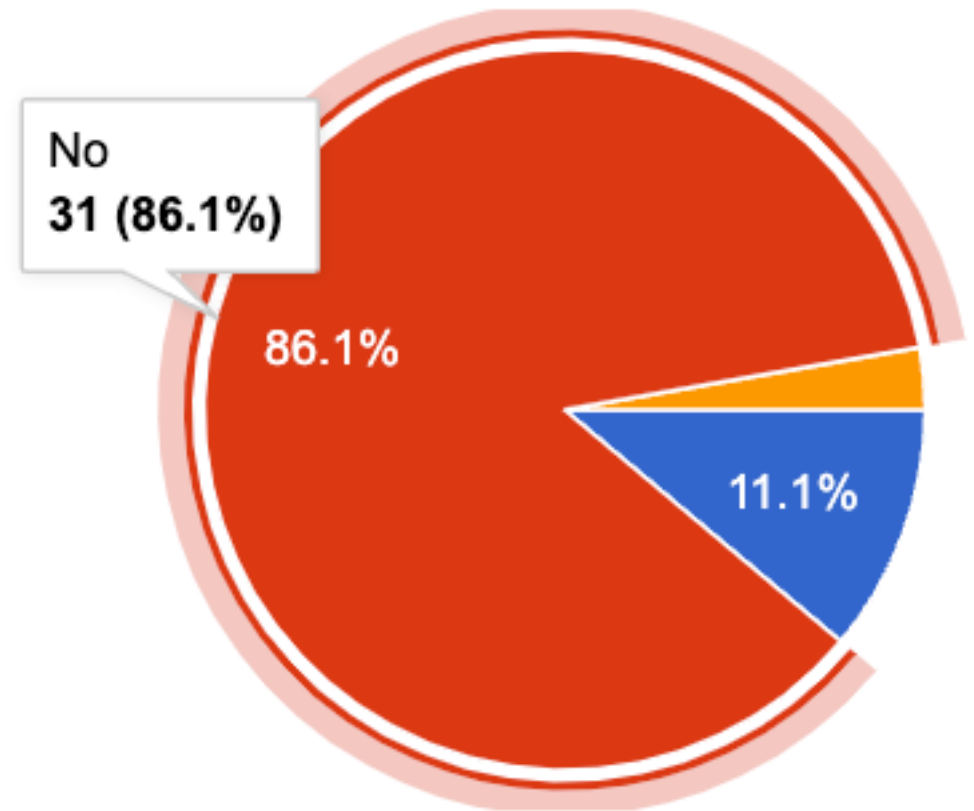
Pre-class Survey

Have you ever taken any machine learning or deep learning courses before?

- Yes
- No
- Maybe

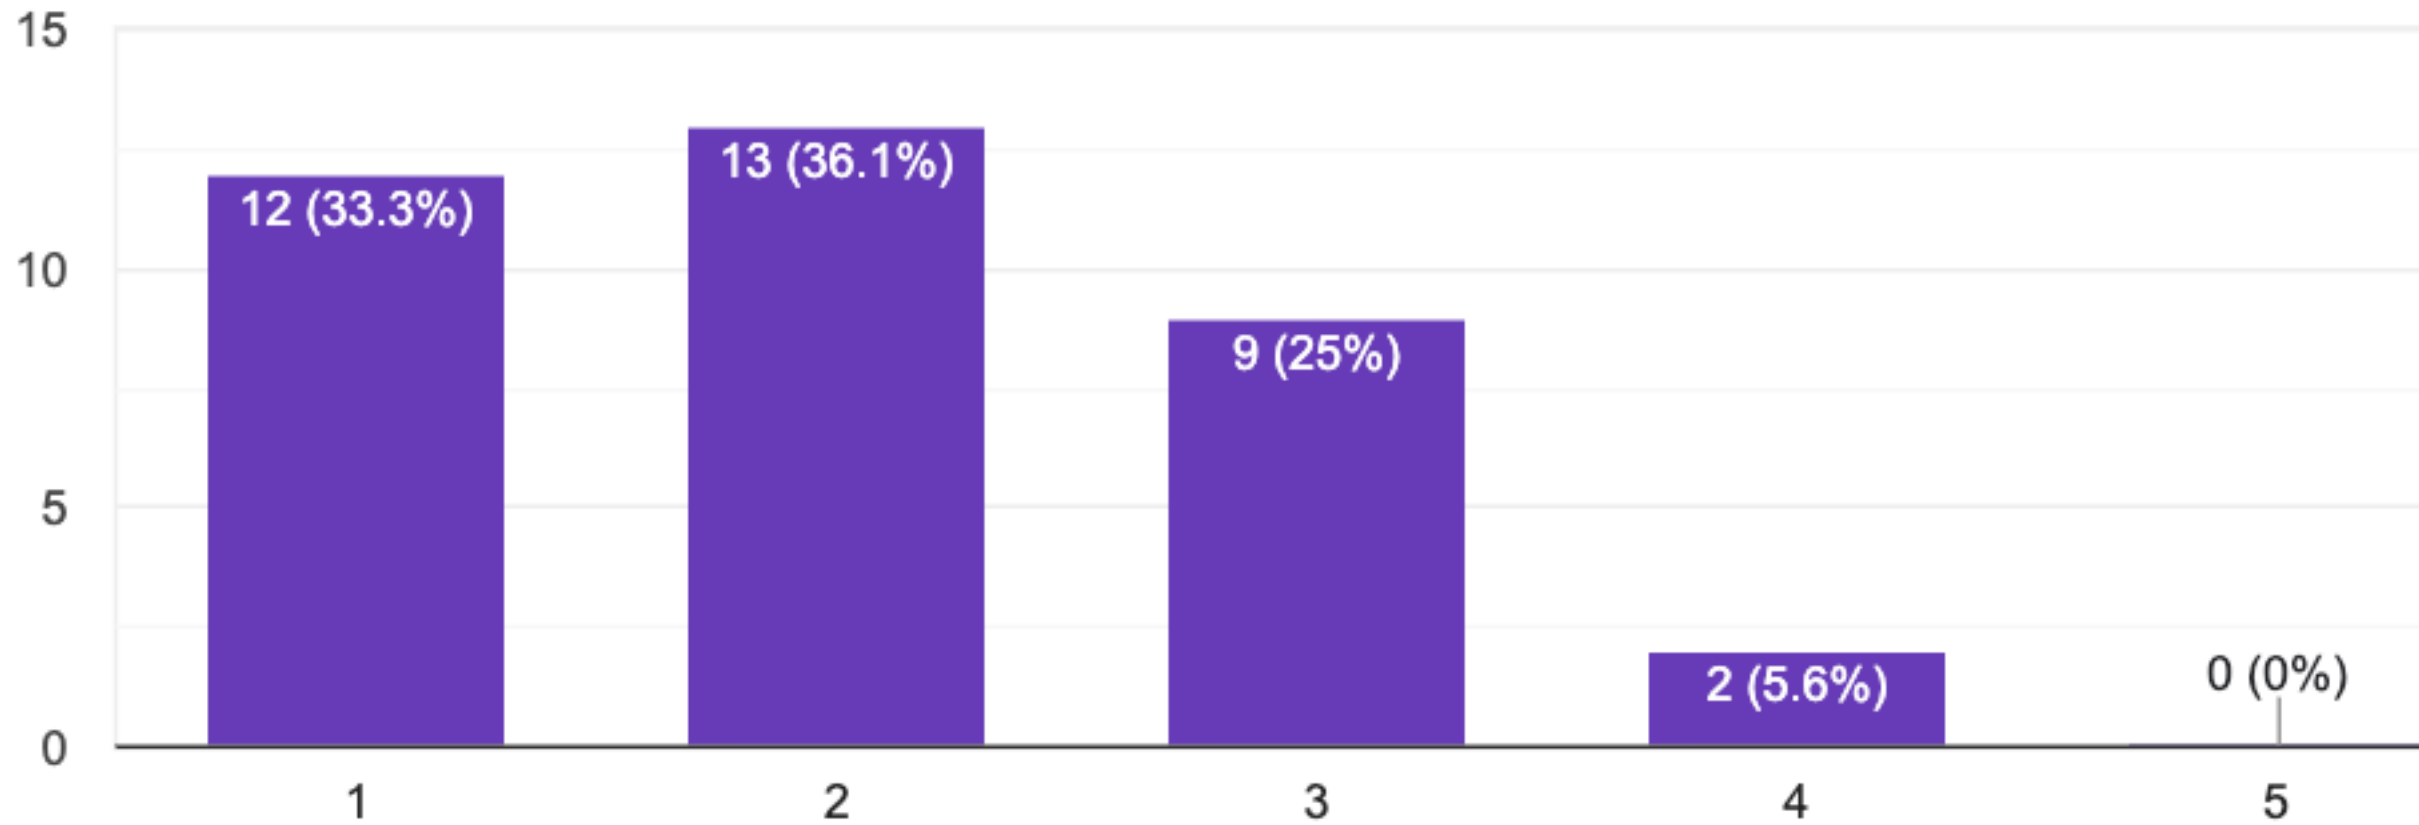


Have you ever taken any NLP courses before?



Pre-class Survey

Can you rate your current knowledge about NLP from 1 to 5?



Pre-class Survey

What is the most interesting topic you like to learn from this class?

Have you had a chance to look at the current schedule in our class homepage?

(<https://dykang.github.io/classes/csci5541/S24/>) What is the most interesting topic you like to learn from this class?

36 responses

Lexical Semantics

Language Models

Language models

BERT; there's a COVID-Twitter-BERT I want to use in my research, so I'm excited to learn about BERT.

Deep learning topics/Transformer/Attention

Hugging face transformer

Yes. Language Models

I am interested in the LLMs part.



Pre-class Survey

What is the most interesting topic you like to learn from this class?



You

Here are topics of interests from students. Can you summarize them by count?



ChatGPT

Here is a summary of the topics of interest from the students, along with the count of mentions:

1. Lexical Semantics: 3 mentions
2. Language Models: 8 mentions
3. BERT: 1 mention
4. COVID-Twitter-BERT: 1 mention
5. Deep Learning topics/Transformer/Attention: 2 mentions
6. Hugging Face Transformer: 2 mentions
7. PyTorch: 1 mention
8. Creating our own dataset: 1 mention
9. Mamba: 1 mention
10. Transformers in the context of NLP: 2 mentions
11. Introduction to LLMs: 2 mentions
12. AI-related concepts: 1 mention
13. Machine Translation: 1 mention
14. Probabilistic Language Model: 1 mention
15. Distributional Semantics: 1 mention
16. Computation efficiency: 1 mention
17. Text encoding: 1 mention
18. Encoder-Decoder: 1 mention

Note: Some topics may have variations in their naming, and I have consolidated similar topics under a common term (e.g., Language Models, LLMs).



Pre-class Survey

What is the most interesting topic you like to learn from this class?



You

Here are topics of interests from students. Can you summarize them by count?

Lexical Semantics

I'm interested in learning about **lexical** semantics and using that to build language models.



ChatGPT

Here is a summary of the topics of interest from the students, along with the count of mentions:

1. Lexical Semantics: 3 mentions
2. Language Models: 8 mentions
3. BERT: 1 mention
4. COVID-Twitter-BERT: 1 mention
5. Deep Learning topics/Transformer/Attention: 2 mentions
6. Hugging Face Transformer: 2 mentions
7. PyTorch: 1 mention
8. Creating our own dataset: 1 mention
9. Mamba: 1 mention
10. Transformers in the context of NLP: 2 mentions
11. Introduction to LLMs: 2 mentions
12. AI-related concepts: 1 mention

1 mention

tion

ion

18. Encoder-Decoder: 1 mention

Note: Some topics may have variations in their naming, and I have consolidated similar topics under a common term (e.g., Language Models, LLMs).



Pre-class Survey

What do you want to get out of this course? What is the most important thing you hope to learn from this course? What are you most excited about in this course?

23 responses

This is a very topical subject, and so I hope to get a solid foundational understanding of it at least

I'm hoping to gain enough competence to be an informed collaborator as I try to automatically identify anti-vax misinformation on Twitter accounts of my research participants.

I have been working as a Software Engineer in search/IR space for over 15 year, but I have little experience with NLP. I am hoping to gain an appericiation of the theory as well as advanced applications of NLP.

Get a good grasp of theory behind the NLP. Also working



Pre-class Survey

What do you want to get out of this course? What is the most important thing you hope to learn from this course? What are you most excited about in this course?

ChatGPT

Here is a summary of the opinions and goals expressed by the students:

1. Gain a solid foundational understanding of NLP.
2. Identify anti-vax misinformation on Twitter accounts of research participants using NLP.
3. Gain appreciation for the theory and advanced applications of NLP, especially for someone with over 15 years of experience in search/IR space.
4. Grasp the theory behind NLP and work on projects to gain hands-on experience.
5. Acquire practical and useful knowledge to improve research in a project involving NLP techniques.
6. Build an NLP project to deepen understanding.
7. Explore the potential of large language models in parsing human commands for robotics and reinforcement learning.
8. Build a language model for specific tasks through fine-tuning on a provided dataset.
9. Learn the intricate nuances of how machines identify, process, and extract meaning from input.
10. Explore Attention-based models and see them in action.
11. Conduct literature review on NLP, complete a non-trivial project, and learn from Prof. DK and team.
12. Gain a more in-depth perspective on NLP compared to previous linguistics classes.
13. Understand the mechanisms enabling computers to distinguish the context of word usage in sentences.
14. Learn the basics of NLP and identify topics of interest for further study.
15. Contribute to projects involving LLMs in a part-time UNITE student role.
16. Understand the math behind turning language into interpretable information for models.
17. Gain a better understanding of how NLP is currently used and its future applications.
18. Deepen practical knowledge of using ML with NLP.
19. Seek hands-on experience in NLP for current robotics research to make robots smarter.



Pre-class Survey

(optional) What are you most worried about? (22 responses)



ChatGPT

Here is a summary of concerns expressed by students, along with the count of mentions:

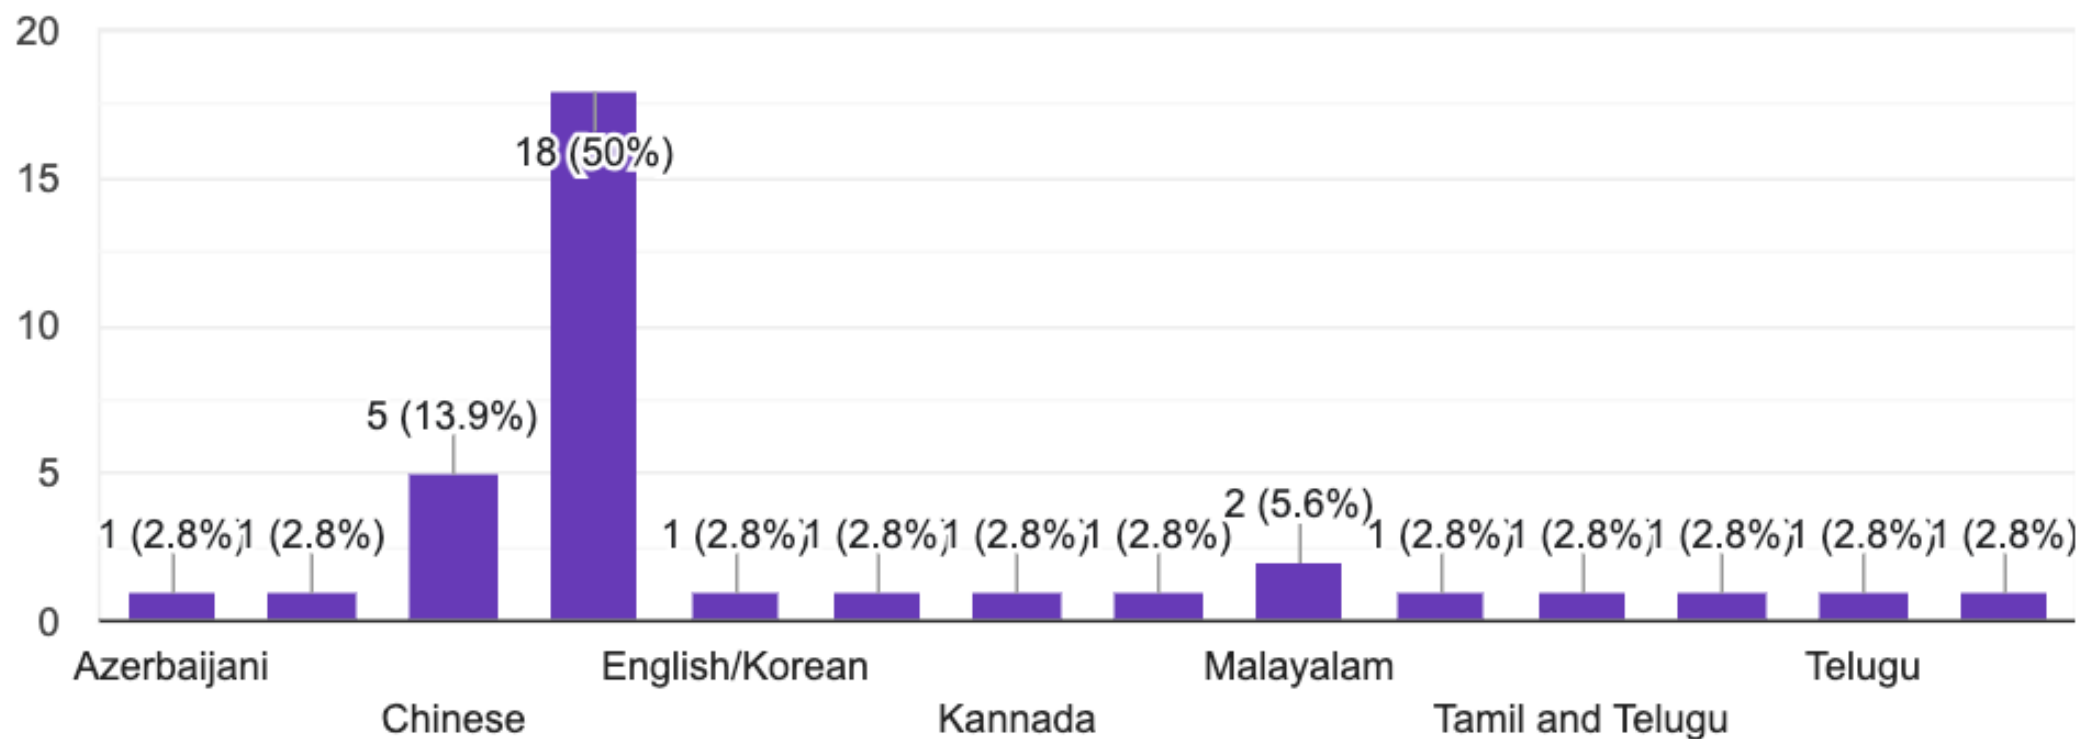
1. Assignment scheduling and procrastination habit: 1 mention
2. Not a fan of group projects due to unreliability: 1 mention
3. Time constraints due to full-time work and family responsibilities: 1 mention
4. Concerns about forming groups on their own: 1 mention
5. Waitlisted for the class and hopes to be enrolled: 1 mention
6. Stress: 1 mention
7. Concurrently taking CSCI 5521, worried about understanding deep learning concepts: 1 mention
8. Not getting into the course: 2 mentions
9. Concerns about team collaboration: 1 mention
10. Worries about not understanding obscure concepts: 1 mention
11. Timely completion of homework and projects: 1 mention
12. Concerns about the current tech job market in the USA: 1 mention
13. Worries about mathematical aspects, especially linear algebra: 1 mention
14. Zero experience with Python frameworks: 1 mention
15. Rusty on linear algebra, calculus, and ML algorithms: 1 mention
16. Concerns about the course load: 1 mention
17. Unable to grasp the content due to being in a new domain: 1 mention
18. Difficulty in successfully enrolling in the class: 1 mention



Pre-class Survey

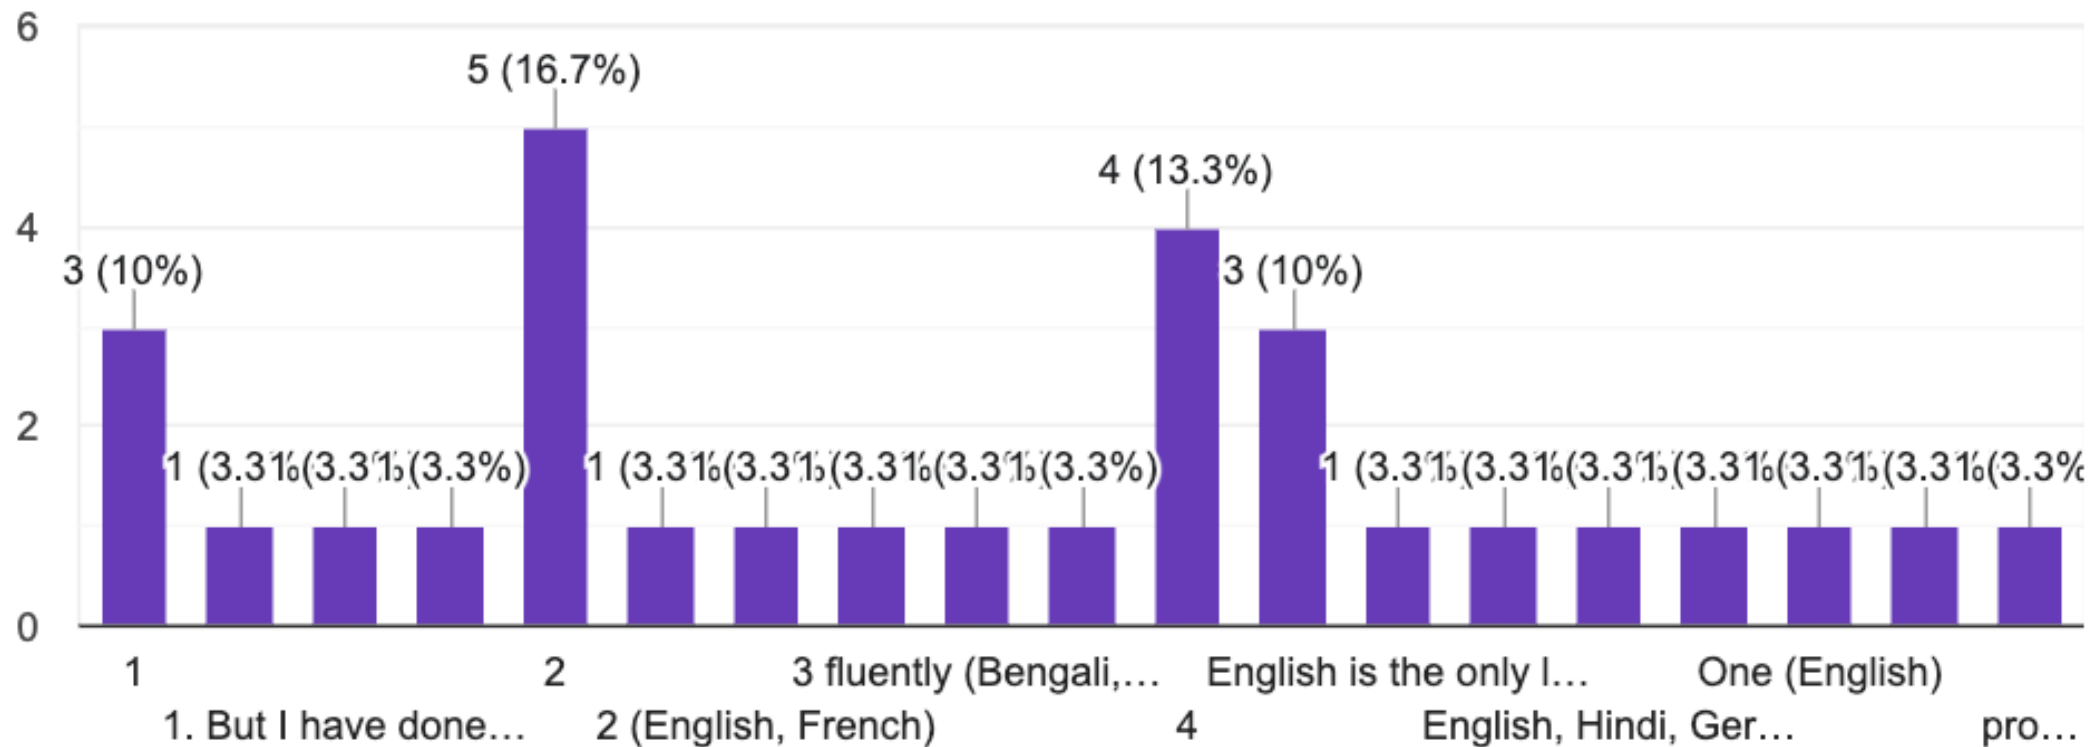
What's your mother tongue? (e.g., English, Chinese, Spanish)

36 responses



Pre-class Survey

(optional) How many language can you speak?



For those who audit or are on waitlist

- ❑ If you plan to drop the course, please do it **by Today** so I can let some students on the waitlist into the class.
- ❑ I will be teaching the same course in next semester, so please consider to take next one if you didn't make it this time



Next

- ❑ Recap: If you plan to drop the course, please do it **by Today**
- ❑ You will be added to the class Slack and Canvas after the class
- ❑ Next lecture
 - Introduction to NLP



Generate an image to the following situation: I'm giving an overview lecture of my NLP class. Show half of students are excited about NLP and the other half are bored. A few students are too much worried about the homework and project

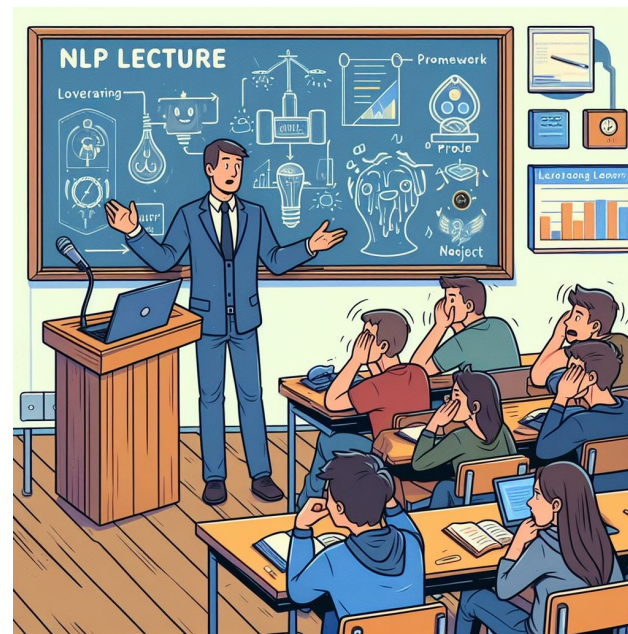
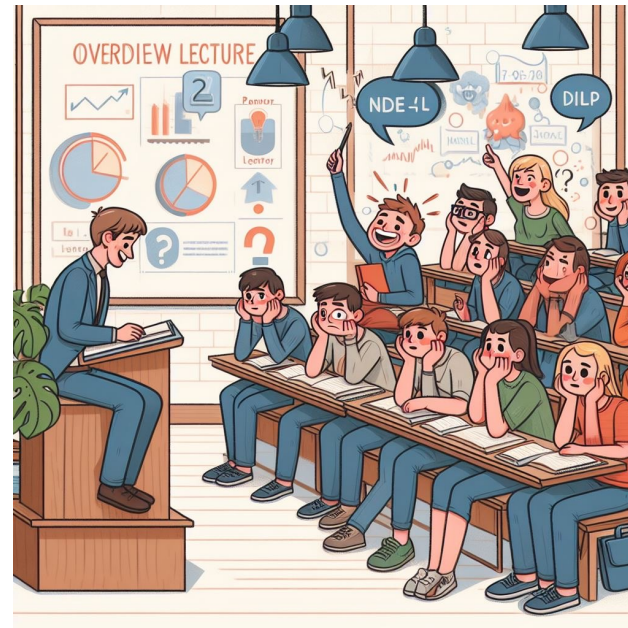
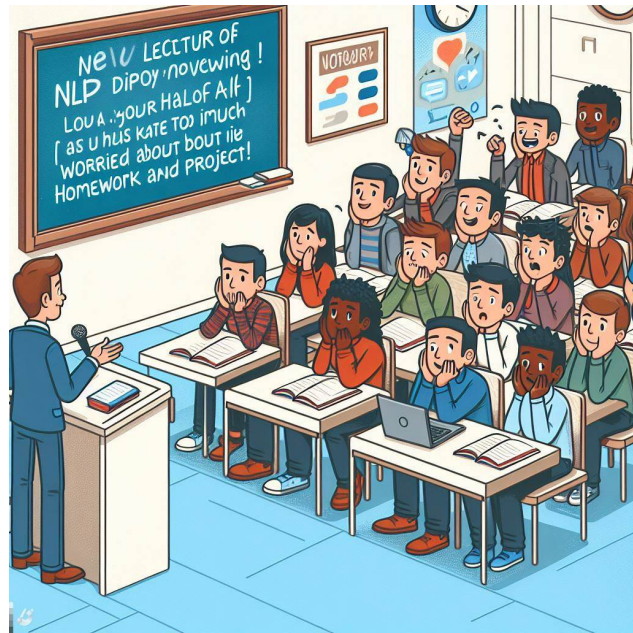
I'll try to create that.

👍 👎 📄 1 of 5 ●

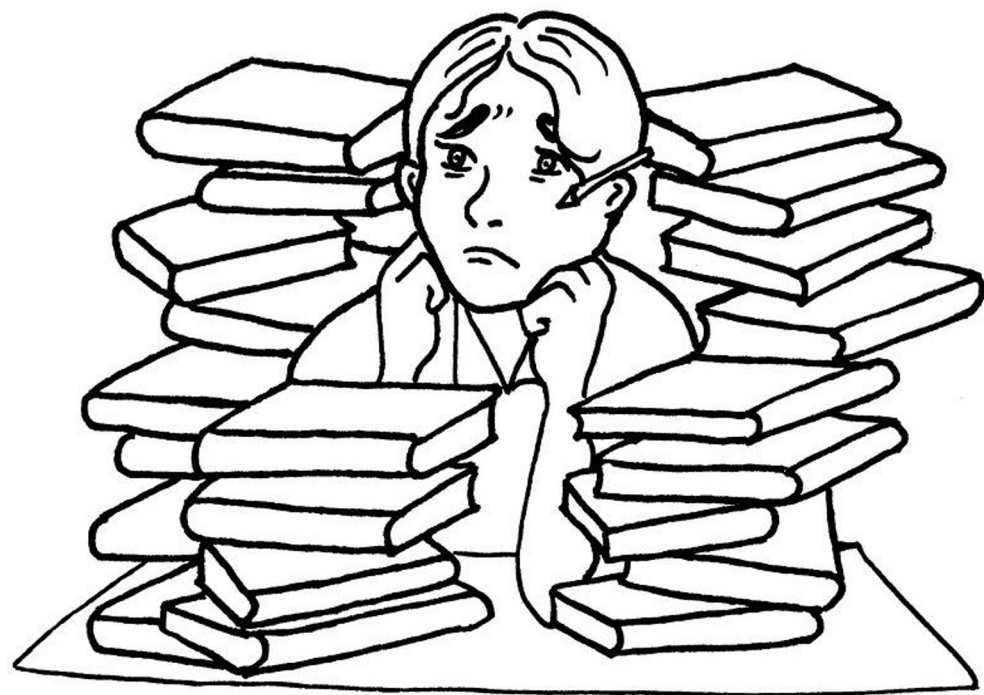


New topic

Ask me anything... 0/2000



Questions?



Overwhelmed?



Bored?