

# Syllabus

## CS 795

NLP Theory [🔗](#)

### Instructors

[Antonis Anastasopoulos](#) (antonis [at] gmu [dot] edu)  
Office Hours: Email for appointments.

### TA

TBD (TBD [at] gmu [dot] edu)  
Office Hours: TBD

### Meets

Wednesdays, 4:30 pm - 7:10 pm, Enterprise Hall 173.

### Course Web Page

<https://nlp.cs.gmu.edu/course/cs795-fall25/>.  
We will use **Canvas** and **Gradescope** for course materials/assignments/grading, and **Piazza** for Q&A (sign up link: <https://piazza.com/gmu/fall2025/cs795>).

### Course Description

This PhD-level seminar course will cover seminal topics around the theory of language, learning, and the emerging theory around large language models. The class will require semester-long research projects based on the materials discussed in class.

### Prerequisites

No official prerequisites, but highly recommended are CS 678 (Advanced NLP) and CS 747 (Deep Learning). Please contact the instructor if you have questions about the necessary background.

### Class Format

- The class will be in-person. In general, each class will take the following format:
- *Reading*: Before some lectures, you will be pointed to some reading materials (*see "Reading Materials" in course schedule*) that you should read before coming to class that day.
  - *Presentations*: A group of students will be responsible for leading the discussion on each paper.

### Grading

There will be no midterm or final exam. Your final grade will be dependent on:  
**Class participation (50%)**: You are expected to provide construct feedback to the assigned groups.

**Project (50%)**: Ths will be a group research project related to LLMs and the topics we will discuss in class. **The groups will be of 2-3 people**. More details about the project will be shared in class.

Letter Grade	Points (out of 100)
A	94-100
A-	90-93
B+	86-89
B	83-85
B-	80-82
C+	76-79
C	73-75
C-	70-72
D	60-69
F	0-59

Contact the instructor in advance if you know that you will very likely not be able to attend a class due to illness or other incidents.

### Tentative Schedule

- Our discussions will be framed around the following modules:
- 1. LLMs and Language Understanding
  - 2. Language Learning in the Limit
  - 3. Learning Under Privileged Information
  - 4. Expressivity of Transformers

Date	Module	Details	Reading Materials
08/27	1	Introductions and Class Outline; meaning and language understanding	<b>Required</b> <ul style="list-style-type: none"><li>[1] <a href="#">Climbing towards NLU: On Meaning, Form, and Understanding in the Age of Data. Bender and Koller, 2020</a></li><li>[2] <a href="#">Eight Things to Know about Large Language Models, Bowman, 2024.</a></li></ul> <b>Optional:</b> <ul style="list-style-type: none"><li><a href="#">Is it possible for language models to achieve language understanding?, Potts, 2020.</a></li></ul>
09/03	1	The Nature of LLMs	<b>Required</b> <ul style="list-style-type: none"><li>[1] <a href="#">Dissociating language and thought in large language models. Mahowald et al., 2024</a></li></ul> <b>Optional:</b> <ul style="list-style-type: none"><li>[1] <a href="#">Do Prompt-Based Models Really Understand the Meaning of Their Prompts? Webson and Pavlick, 2022</a></li><li><a href="#">Do As I Can, Not As I Say: Grounding Language in Robotic Affordances, Ichter et al., 2023.</a></li></ul>
09/10	2	Language Learning in the Limit	<b>Required</b> <ul style="list-style-type: none"><li>[1] <a href="#">Language Identification in the Limit</a></li><li>[2] <a href="#">Inductive inference of formal languages from positive data, Angluin, 1980.</a></li></ul> <b>Optional:</b> <ul style="list-style-type: none"><li>[1] <a href="#">The Logical Problem of Language Acquisition, Cowie, 1997.</a></li><li>[2] <a href="#">Some decidability results on grammatical inference and complexity, Feldman, 1972.</a></li></ul>
09/17	2	Language Generation in the Limit	<b>Guest Lecture by Evgenios Kornaropoulos</b> <b>Strongly Recommended:</b> <ul style="list-style-type: none"><li>[1] <a href="#">Language Generation in the Limit, Kleinberg and Mullainathan, 2024.</a></li></ul> <b>Optional:</b> <ul style="list-style-type: none"><li>[1] <a href="#">Characterizations of Language Generation With Breadth, Kalavasis et al., 2025.</a></li><li>[2] <a href="#">Representative Language Generation, Peale et al., 2025.</a></li><li>[3] <a href="#">Exploring Facets of Language Generation in the Limit, Charikar and Pabbaraju, 2025.</a></li><li>[4] <a href="#">Generation from Noisy Examples, Raman and Raman, 2025.</a></li></ul>
09/24	1,2	An aside: Hallucinations and Biases	<b>Required</b> <ul style="list-style-type: none"><li>[1] <a href="#">Calibrated Language Models Must Hallucinate, Kalai end Vempala, 2023.</a></li><li>[2] <a href="#">Large Language Models Are Biased Because They Are Large Language Models, Resnik, 2025.</a></li></ul> <b>Optional:</b> <ul style="list-style-type: none"><li>[1] <a href="#">(Im)possibility of Automated Hallucination Detection in Large Language Models, Karbasi et al., 2025.</a></li><li>[2] <a href="#">LLMs Know More than they Show: on the Intrinsic Representation of LLM Hallucinations, Orqad et al., 2025.</a></li><li>[3] <a href="#">Foundational Challenges in Assuring Alignment and Safety of Large Language Models, Anwar et al., 2024. (sections 2 and 3)</a></li><li>[4] <a href="#">HALOGEN : Fantastic LLM Hallucinations and Where to Find Them, Ravichander et al., 2025.</a></li></ul>
10/01	1	The Nature of LLMs (Part II)	<b>Required</b> <ul style="list-style-type: none"><li>[1] <a href="#">Embers of Autoregression: Understanding Large Language Models Through the Problem They are Trained to Solve, McCoy et al., 2023.</a></li></ul> <b>Optional:</b> <ul style="list-style-type: none"><li>[1] <a href="#">Reasoning or Reciting? Exploring the Capabilities and Limitations of Language Models Through Counterfactual Tasks, Wu et al., 2024.</a></li><li>[2] <a href="#">Large Language Models Fail on Trivial Alterations to Theory-of-Mind Tasks, Ullman, 2023.</a></li><li>[3] <a href="#">When a Sentence Does Not Introduce a Discourse Entity, Transformer-Based Models Still Sometimes Refer to It. Schuster and Linzen, 2022.</a></li></ul>
10/08	3	Learning Under Privileged Information Part 1	<b>Required</b> <ul style="list-style-type: none"><li>[1] <a href="#">A new learning paradigm: Learning using privileged information, Vapnik and Vashisht, 2009.</a></li></ul> <b>Optional:</b> <ul style="list-style-type: none"><li>[1] <a href="#">On the Theory of Learning with Privileged Information, Pechyony and Vanpik, 2010.</a></li></ul>
10/15	3	Learning Under Privileged Information Part 2	<b>Required</b> <ul style="list-style-type: none"><li>[1] <a href="#">On the Power of Context-Enhanced Learning in LLMs, Zhu et al., 2025.</a></li><li>[2] <a href="#">Transformers as Algorithms: Generalization and Stability in In-context Learning, Li et al., 2023.</a></li></ul> <b>Optional:</b> <ul style="list-style-type: none"><li>[1] <a href="#">Deep Learning under Privileged Information Using Heteroscedastic Dropout, Lambert et al., 2019.</a></li><li>[2] <a href="#">Metadata Conditioning Accelerates Language Model Pre-training, Gao et al., 2025.</a></li></ul>
10/22	4	Representation Power of Transformers Part 1	Antonis Presenting on Preliminaries (formal languages, logic, circuits) <b>Optional (in preparation of the whole module):</b> <ul style="list-style-type: none"><li>[1] <a href="#">What Formal Languages Can Transformers Express? A Survey, Strobl et al., 2024.</a></li><li>[2] <a href="#">Overcoming a Theoretical Limitation of Self-Attention, Chiang and Cholak, 2022.</a></li></ul>

10/29	4	Representation Power of Transformers Part 2	<div>Required</div> <ul style="list-style-type: none"><li>[1] <a href="#">Theoretical limitations of self-attention in neural sequence models, Hahn, 2020.</a></li><li>[2] <a href="#">Masked hard-attention transformers and Boolean RASP recognize exactly the star-free languages, Yang et al., 2023.</a></li></ul> <div>Optional:</div> <ul style="list-style-type: none"><li>[1] <a href="#">Formal language recognition by hard attention Transformers: Perspectives from circuit complexity, Hao et al., 2022.</a></li><li>[2] <a href="#">Characterizing the Expressivity of Transformer Language Models, Li and Cotterell, 2025.</a></li><li><a href="#">Why are Sensitive Functions Hard for Transformers? Hahn and Rofin, 2024.</a></li></ul>
11/05	4	Representation Power of Transformers Part 3	<div>Required</div> <ul style="list-style-type: none"><li>[1] <a href="#">Attention is Turing Complete, Pérez et al., 2021.</a></li><li>[2] <a href="#">Transformers as transducers, Strobl et al., 2024.</a></li></ul> <div>Optional:</div> <ul style="list-style-type: none"><li>[1] <a href="#">Logical languages accepted by transformer encoders with hard attention, Barceló et al., 2024.</a></li></ul>
11/12			Guest Lecture - TBD
11/19	4	Transformers and Impossible Languages	<div>Required</div> <ul style="list-style-type: none"><li>[1] <a href="#">Mission: Impossible Language Models, Kallini et al., 2024.</a></li><li>[2] <a href="#">Can Language Models Learn Typologically Implausible Languages?, Xu et al., 2025.</a></li></ul> <div>Optional:</div> <ul style="list-style-type: none"><li>[1] <a href="#">Lexinvariant Language Models, Huang et al., 2023.</a></li><li>[2] <a href="#">Priorless Recurrent Networks Learn Curiously, Mitchell and Bowers, 2020.</a></li></ul>
11/26			NO CLASS (Thanksgiving)
12/03			Final Project Presentations

Honor Code

The class enforces the [GMU Honor Code](#), and the [more specific honor code policy](#) special to the Department of Computer Science. You will be expected to adhere to this code and policy.

- IMPORTANT NOTE about the use of ChatGPT and similar technologies:** Unless permission to do so is granted by the instructor, you (or your group, if a group assignment) may not use assistive technology, artificial intelligence, or other tools to complete assignments which can generate, translate, or otherwise create/correct code or answers (many types of assistive technology may be permitted, but you must ask permission).

Note to Students

Take care of yourself! As a student, you may experience a range of challenges that can interfere with learning, such as strained relationships, increased anxiety, substance use, global pandemics, feeling down, difficulty concentrating and/or lack of motivation. All of us benefit from support during times of struggle. There are many helpful resources available on campus and an important part of having a healthy life is learning how to ask for help. Asking for support sooner rather than later is almost always helpful. GMU services are available, and treatment does work. You can learn more about confidential mental health services available on campus at: <https://caps.gmu.edu/>. Support is always available (24/7) from Counseling and Psychological Services: 703-527-4077.

Disabilities

If you have a documented learning disability or other condition which may affect academic performance, make sure this documentation is on file with the [Office of Disability Services](#) and come talk to me about accommodations. I will work with you to ensure that accommodations are provided as appropriate. If you suspect that you may have a disability and would benefit from accommodations but are not yet registered with the Office of Disability Services, I encourage you to contact them at [ods@gmu.edu](mailto:ods@gmu.edu).

Last updated on Jan 1, 0001