LLMs & Cognition 2026

Course

weeks 1 to 7: Ana Uban
weeks 8 to 14: Sergiu Nisioi
Thursdays 10:00-12:00
room 512 PBT

Labs

Claudiu Creangă

Instructor	Email
Ana Sabina Uban	auban@fmi.unibuc.ro
Sergiu Nisioi	sergiu.nisioi@unibuc.ro
Claudiu Creangă	claudiu.creanga@s.unibuc.ro

Grading

1p Free as in Beer

2p Lab Assignments

- You will receive 4 equally weighted homework assignments. Late submission is accepted but with a penalisation of -1pct for each week.
- You will receive points for seminar activity (interaction during the class).
- Attendance is not graded.

2p In-class Article Reading Participation

- Grades are given by the following criteria: Zero 0, Insufficient (IS – 0.5), Sufficient (S – 1), Good (G –1.5), Very Good (VG – 2)
- We assume everyone is very good, but we ask questions individually. If you do not know how to answer basic questions related to your work, you get awarded -0.5.
- We assign teams of people for each week and they need to be present in class and participate in the discussion at the very least. Extra points for putting together slides

or leading the discussion. It's up to the team whether they organize together or prepare separately, but each person must participate in the discussion in order to get the 2 points of your final grade (it is not possible to make up for it during the retake session).

- The team assigned for each week will be scheduled in advance at the beginning of the semester (see table below).
- The topic for each week - an article to read and get familiar with - will be decided and announced at least one or two weeks in advance, depending on the complexity of the topic.
- We recommend the following organization of the team into roles for tackling the understanding of the paper (we will pre-assign roles, but you can switch among yourselves or collaborate if you prefer). <u>All students in the team should read the entire article in order to properly understand it</u>, but each student can focus primarily on one of the following aspects:

Problem framer (PF)

- Focus on the Introduction, Related Work and Conclusions sections of the paper, describe the problem that the paper is trying to solve, and the motivation for approaching it. What previous work exists related to the topic and how the current study relates to it (what novelty it proposes to bring). Present what other works have been published in the meantime on the topic?

Methodology experts (ME)

- Focus on the Methodology sections of the paper (aka "Methods"/"Experiments"/…, including description of the data used if any), have a good grasp of the details of the methodological steps. Experts should understand perfectly what figures, plots and models mean and how they are generated.

Coder (CO)

- Investigate whether there is any codebase available associated with the study (usually linked in the paper itself), run the code and understand what it does. Present the structure of the code and prepare a running demo to show to the entire class. In case the results of the paper are not reproducible, the demo can also cover a small subsection of the paper, a smaller model or a particular component. Produce a demo, a colab, a streamlit app, a piece of code, depending on the type of paper.

Data scientist (DS)

- Find the dataset used in the paper, prepare a dataset description. Take the entire dataset or a subsample (if it's too large) and prepare a notebook, some tables and a series of plotly plots with an Exploratory Data Analysis, for example, but not limited to: number of words, type-token ratio, number of sentences, means and standard deviations, distribution plots, fertility of a multilingual tokenizer, perplexity of a small (multilingual) language model (e.g. SmolLM, gemma-270m, or a quantized version that runs locally), average

word length, average word frequency (consider wordfreq vs dataset frequency) and any other statistics or meaningful info.

Results interpreter (R):

- Focus on the Results section (“Results Analysis”/”Discussion”), understand what the study obtains in terms of experimental results - understand all results tables and plots, and be able to comment on the comparison between the results of different experiments reported in tables/plots, and what the results show in relation to the proposed task.

Challenger (CH)

- Focus on Limitations and Future Work - either in dedicated sections (“Limitations”/”Conclusions and Future Work”/”Discussion”/”Ethical Concerns”…) or throughout the paper. This role also involves understanding the paper and methodology and it requires the presenter to show details on how future work can be implemented, how much time would it take to build additional results, strategies to build future datasets, and the detailed methodological approach. Comment on what could be some problems/shortcomings of the paper, potential questions that you might ask the authors, or even possible additions to their experiments - This can be a starting point for your final project :)

- The slides, code snippets or any other artifacts should be uploaded on teams: Semester Presentations

5p Final Projects

Final projects will be solved in teams of minimum 3 and maximum 4 students.
Project topics and future deadlines will be announced on the teams platform. Where topics overlap, two teams will compete for the maximum score and the best team will qualify to receive one extra bonus point.

Progress Reports (2p)

By **October 31 Anywhere on Earth** (AOE, strict deadline) you need to have decided on a topic and to submit a preliminary plan on how you plan to solve it (see the 5 points below for direction). No submission means automatically that you cannot present a project and you risk failing the class.

Progress reports (0.5p each) human-written report to document the progress. Strict deadlines on: **19 Nov, 03 Dec, 17 Dec, 14 Jan**, 10:00 am.
While doing your reports, consider covering the following:

1. related work and papers on the subject
2. what data is available, how large is your dataset, to make a preliminary exploratory data analysis

3. what types of modeling strategies, evaluation metrics you can use based on the compute resources that you have
4. results, qualitative evaluation, automatic evaluation, actual model comparisons
5. questions and bottlenecks

This is an opportunity to identify problems, bottlenecks, and to receive feedback and advice from us. Research projects will involve regular updates on progress and regular feedback from us.

Failing the class

If you fail the class, you must get in touch with us before **July 20 2026**, so we can schedule project plans etc. for the September session. If you do not get in touch with us to establish a project for the retake session, we will automatically consider you failed the retake session. Lab points obtained during the semester are retained until September.
If you want to increase you grade, that should be possible by getting in touch with us by the same deadline to be assigned extra project work.

Project requirements

In addition to the scheduled progress reports that amount to 2p in total, a project must have specific deliverables. For the project grades will be given using similar marks criteria: Zero 0, Sufficient (S – 1), Good (G –2), Very Good (VG – 3)

Deliverables:

- paper / technical report
- implementation code
- slide presentation

All papers/technical reports will follow the classical structure of a research article (4-8 pages). For group projects, the report must include a statement of the contribution of each team member.

The overleaf template for writing the report is available here: https://www.overleaf.com/latex/templates/association-for-computational-linguistics-acl-conference/jvxskxpnznfj (yours shouldn't be anonymous). It is mandatory to introduce at the end, after the Conclusions and before the Bibliography a short section called **Contributions** explaining the contribution of each student in the team.

Make a directory for your team and upload everything here (make a folder for your project)

Writing tips for the final report

The report should be self-contained: it should describe the models and datasets you're using and hypotheses you're testing, without assuming that the reader already knows them, from the proposal or from anywhere else. The introduction should be largely understandable by a smart undergraduate student outside the field.

Clear writing is very important. We are not going to take off points for grammar errors, but please read over your report carefully to make sure your sentences are concise, simple, and understandable by someone who is not closely familiar with your project. If your writing is difficult to follow, this may affect your grade even if the technical work you did was sound.

A typical report will consist of the following content:

1. An introduction motivating the project and describing prior work. Only cite a paper if it's relevant, for example, if you want to explain how your method or research question differs from existing work. We assume any project will have some relevant prior work, which you should cite and include in the References.
2. Methods (or Experimental Setup): your models, your datasets and how you ran your experiments. The descriptions should be detailed enough that someone reading your paper can reproduce your experiments based on it and reach the same results.
3. Results: Include figures and/or tables. Figures are usually better than tables. All figures and tables should have a caption, which should be as self-contained as possible (avoid “see text for details”). Text in the figures should be readable without zooming in.
4. Discussion: Was your hypothesis confirmed? How do your findings relate to findings from prior work? What are the limitations of your experiment, and how can they be addressed in future work?
5. Conclusion: One or two paragraphs summarizing the main takeaway.
6. Author contribution statement (for team projects) (can exceed the 8 page limit).
7. LLM disclosure section (can exceed the 8 page limit): how did you use LLMs as a tool to help in the elaboration of the project code, report etc - models and versions, prompts.

You may also want to check:

- Writing a Good Scientific Paper
- How to Get Your Paper Accepted
- Highly Opinionated Advice on How to Write ML Papers

Avoid writing your reports using LLMs, these software tend to introduce a lot of useless gibberish and you may get penalized for not being clear and concise, see also this https://hackingsemantics.xyz/2024/writing/

Course Presentations Schedule

Below you will find your institutional username and the date when you will have to present a pre-assigned paper. The papers will be announced during the semester, depending on the difficulty. If you do not find your name in the list, contact Ana or Sergiu immediately.

It is likely that you have been assigned with colleagues whom you do not know. This is exactly the point. Try to meet and get to know each other. Print the paper and read it together, go for a beer afterwards.

In class, you will have to prepare a joint presentation or discussion panel around that paper. Presenters are graded individually. Each one of you will be assigned a separate role to disentangle a paper (engineering part, theoretical part, the motivation for the work, the methods and results of the main experiments of the paper, and the paper's takeaways etc.). For each presentation, prepare a short slide deck. The goal of the presentation is to serve as a forum for discussion in class.

People who do not participate in this lose 20% of their final grade and it's not possible to recover these points during the retake/reexamination session. This is a once-in-a-lifetime opportunity.

Date	**Team (names and roles)**	**Paper**
October 08		Introductory class
October 15	GUSEV ROMAN (R) GHEORGHE BOGDAN-ALEXANDRU (CO) CUCLEA LUCA-NICOLAE (CO) LEAHU TEODOR FILIP (ME) HERMENEANU IONUT SILVIU (PF) SZOCS LUCA (ME)	Transformers introduction: **Paper: Attention is all you need** Additional materials: https://jalammar.github.io/illustrated-transformer/ https://magazine.sebastianraschka.com/p/understanding-encoder-and-decoder https://poloclub.github.io/transformer-explainer/ Code suggestion - attention visualization: https://colab.research.google.com/drive/1gDgnVaWQ33Wq82dzPd2tmCwIfZVao68m?usp=shari

		ng We welcome other tutorials/tools that illustrate the architecture of the transformer or newer versions of attention: https://arxiv.org/pdf/2608.28444
October 22	Alcaziu Robert-Constantin (PF) DINCA MARIUS STEFAN (DS) HUS LUCIAN-ANTONIO (ME) ROBERT-VLADUT MOROSANU (CO) POPA TEODORA-IOANA (R) ROMAN BIANCA-GABRIELA (ME/DS)	**The LLM Language Network: A Neuroscientific Approach for Identifying Causally Task-Relevant Units**
October 29	-	**Conference EMNLP (online?)**
November 5	NEACSU-DALU ANAIS MARIA ANDREI-COSMIN SOARE DANIEL-ANDREI VELCEA MIHNEA-ANDREI MOCANU ANDREI IONUT TOMA ANDREEA	**TBD**
November 12	VIERIU ELENA LARISA AXINTE CALIN-CRISTIAN Dragota Andrei PISLARU LUCA-ANDREI Gherghina Bogdan PANOIU SILVIU CRISTIAN	**TBD**
November 19	AELENEI ALEX IOAN LELIUC ROBERT CONSTANTIN TRUSCA FABIAN MOLDOVAN SARA-MARIA Roscan Rares-Alexandru IOANA ANDREI-MADALIN	**TBD**
November 26	TANASA ANA-MARIA LUMINARARU IONUT-ANDREI angelos fountanas CUCU MIHNEA COTIGA DAVID Neacsu Eric Marian	**TBD**
December 3	DUMITRU ALEXANDRU DANIEL DINU ANDREI-LUCA GALERIU VLAD PUSCASU TUDOR Faliboga Dimitrie-Aron DODITA ALEXANDRU TOMI	**TBD**
December 10	JILAVU MARIA-IZABELA PINZARU EVELINA PORT DRAGOS	**TBD**

	Glod Cosmin-Stefan DUDCEAC OCTAVIAN IOAN NEACSU MIHNEA-VALENTIN	
December 17	TULCEANU CEZAR-ALEXANDRU Gheorghe Andrei-Bogdan BAGHICI GABRIEL-CRISTIAN MORARIU ALEXANDRU GABRIEL UNGUREANU VLAD MADALIN	**TBD**
January 14	CRACIUN CATALIN Iaconi Andreea RARES ALEXANDRU MEDELET Ivan Emil-Ovidiu Moraru Otilia-Gabriela	**TBD**
January 21	Project discussions	