CSE 4392 / CSE 5369 NLP Research Project Specification

Course Instructor: Kenny Zhu   kenny.zhu@uta.edu
Teaching Assistant: Peter Zhang   wxz9630@mavs.uta.edu

Due Date: May 6, 2026, 11:59 PM

1   Introduction

The course project for this class is an *NLP research project*. Each team must design and develop a system or model that solves an interesting problem in Natural Language Processing (NLP). The project should be research-oriented rather than purely engineering-oriented: you are expected to formulate a clear problem, motivate its importance, implement your own method, and evaluate it empirically. The goal is to simulate a small-scale research workflow similar to an NLP conference submission.

A complete project submission must include:

- a working implementation,

- the code and data used in the project,

- and a research paper written in ACL format.

2   Project Objective

Your project must address a concrete NLP task or research question. Examples include, but are not limited to:

- Text Classification

- Named Entity Recognition (NER)

- Machine Translation

- Text Generation

- Question Answering

- Speech Recognition

- Text Summarization

- Sentiment Analysis

- Information Extraction

- Dialogue Systems

3 Core Requirements

3.1 Original Implementation

The project must be **your own implementation**. You may build on public libraries and frameworks, but the overall system, experimental pipeline, and research design must be your own work.

- Copying an existing repository and only changing configuration files is not sufficient.

- Re-implementing a baseline from a paper is allowed, provided that you also contribute your own analysis, comparison, extension, or improvement.

- You must clearly state which components are borrowed from public toolkits and which parts were implemented by you.

3.2 Research Component

Your project must contain a genuine research component. At minimum, your paper should answer the following questions:

- What problem are you solving?

- Why is this problem important or interesting?

- What data do you need and how will you obtain it?

- What method or system did you design?

- How does it compare with one or more baselines?

- How will you evaluate your method and what insights do you expect to gain?

A project that only demonstrates software usage without a research question will not receive full credit.

3.3 Experimental Evaluation

Every project must include an empirical evaluation.

- You must evaluate your method on data that is clearly described in the paper.

- You must report appropriate metrics for the task.

- You must include at least one baseline for comparison.

- You should provide analysis beyond a single result table whenever possible.

Examples of further analysis include: error analysis, ablation study, qualitative examples, comparison across data conditions, or discussion of failure cases.

4   Suggested Project Workflow

A strong project will typically include the following stages:

1. **Problem Definition**
   Define the NLP task clearly and precisely.

2. **Related Work Review**
   Summarize prior methods and identify what gap or perspective your project addresses.

3. **Data Preparation**
   Collect, clean, annotate, or preprocess the data as needed.

4. **Method Design**
   Design and implement your model, algorithm, or system.

5. **Baseline Construction**
   Implement or adapt one or more baseline methods for comparison.

6. **Evaluation**
   Run experiments, compute metrics, and analyze the results.

7. **Paper Writing**
   Present the problem, method, experiments, and findings in ACL paper format.

5   Paper Requirements

5.1   Format

You must write your final paper in **ACL format** using the official ACL style files and templates; see the official ACL paper formatting guidelines here: `https://acl-org.github.io/ACLPUB/formatting.html`

Your paper should resemble a conference-style research paper and will typically include, but is not limited to, the following sections:

- Abstract

- Introduction

- Method

- Experiments and Evaluation

- Results and Analysis

- Related Work

- Conclusion

- Limitations

- References

5.2   Length

The paper should target the ACL long paper style:

- **8 pages of main content**, in ACL format,

- plus additional pages for references,

- and optional appendix / supplementary material if needed.

Your goal is to produce a paper of approximately conference-paper quality and scope.

5.3   Writing Expectations

Your paper must:

- clearly describe the background, motivation, and contributions,

- explain the method in sufficient technical detail,

- describe the dataset(s) and preprocessing,

- report experimental settings and evaluation metrics,

- include quantitative results,

- include interpretation or analysis of the results,

- discuss the limitations of your work,

- cite prior work appropriately.

6   Implementation Requirements

6.1   Programming Language and Tools

You may use any reasonable programming language or framework for implementation, such as Python, Java, Py-Torch, TensorFlow, or other NLP toolkits.

6.2   Reproducibility

Your submission must be organized so that the grader can understand and reproduce your work.

At minimum, your code package should include:

- source code,

- `README.md` file,

- instructions for the full pipeline (e.g., data preparation, training, inference, and evaluation),

- dependency information,

- and a description of where the data comes from.

If your experiments require substantial computation, provide:

- preprocessed data if allowed,

- saved checkpoints if helpful,

- and enough scripts so that the pipeline is understandable.

6.3 Data

You must submit the data used in the project whenever legally and practically possible.

- If the dataset is public, provide download instructions and the exact version used.

- If the dataset is too large to submit directly, provide scripts and instructions for obtaining it.

- If you create your own data, clearly document how it was collected and labeled.

7 Deliverables

Each team should submit a single `.zip` file named `LastName_StudentID_Project.zip`; all team members' names and student IDs must be listed in the top-level `README.md`.

Your final submission must contain the following:

- **Paper:** a PDF paper in ACL format, along with the LaTeX source files and any required style/template files.

- **Source Code:** all code needed to implement and run the project.

- **Data:** the dataset, processed data, or instructions/scripts to obtain the data.

- **README:** clear instructions for setup and execution, as well as the names and student IDs of all team members.

Before submission, make sure:

- the paper compiles correctly,

- the code is organized,

- the instructions are complete,

- and all files are named clearly.

A good submission folder may look like this:

```
LastName_StudentID_Project/
|-- paper/
|   |-- project_paper.pdf
|   '-- source_tex_files/
|-- code/
|   |-- src/
|   |-- scripts/
|   |-- train.py
|   |-- eval.py
|   |-- model.py
|   |-- requirements.txt
|   '-- ...
|-- data/
|   |-- README.md
|   '-- ...
|-- results/
|   |-- tables/
|   '-- figures/
'-- README.md
```

8   Project Grading Criteria

The project will be graded based on the following dimensions:

1. **Problem Quality and Motivation**
   Is the problem meaningful, well-defined, and appropriate for NLP research?

2. **Technical Depth**
   Does the project demonstrate real modeling, system design, or algorithmic effort?

3. **Implementation Quality**
   Is the system complete, organized, and reproducible?

4. **Experimental Design**
   Are the baselines, metrics, and comparisons appropriate?

5. **Analysis and Insight**
   Does the paper go beyond reporting numbers and provide interpretation?

6. **Paper Quality**
   Is the paper well-structured, clearly written, and formatted in ACL style?

9   Project Milestones

The project will involve the following key milestones:

- **Milestone 1: Proposal (Jan. 28)**
  Project topic, task definition, motivation, and initial related work.

- **Milestone 2: Progress Check (Mar. 23)**
  Current progress on data preparation, method development, and preliminary experiments.

- **Milestone 3: Final Submission (May. 6)**
  Final paper, code, data, and documentation.

The instructor or TA may announce concrete deadlines separately.

10   Academic Integrity

All submitted work must be the original work of your team.

- You must properly cite papers, codebases, datasets, and external resources.

- Plagiarism, fabricated results, or undisclosed copying of code/text will result in a score of zero and may trigger additional academic penalties.

11   Use of External Models and LLMs

You may use pretrained models, APIs, or large language models as part of your project, but you must clearly document their role.

For example, you should specify:

- whether the model is used as a baseline, feature extractor, or core method,

- whether prompting, fine-tuning, or retrieval augmentation is involved,

- and what part of the implementation is still your own contribution.

Projects that rely entirely on black-box models with little methodological contribution may receive limited credit unless accompanied by strong experimental analysis or a novel research question.

12   Reference Information

You are strongly encouraged to consult the ACL 2026 main conference submission information for the expected paper style, length, and research presentation standard:

https://2026.aclweb.org/calls/main_conference_papers/#paper-submission-information

This course project does not necessarily require an actual ACL submission, but your report should be written to resemble a serious ACL-style research paper.