

Bachelor Thesis Proposal 26/27 winter:

LLM-Assisted Cluster Interpretation: Using LLMs to Explain and Label Text Clusters

Short Abstract

Text clustering can automatically discover groups of semantically related documents, but interpreting these clusters is often a manual and time-consuming process. After clustering, a researcher typically needs to inspect many documents and assign meaningful descriptions or labels to the discovered groups.

This thesis investigates how Large Language Models (LLMs) can assist in interpreting text clusters. Given a set of documents belonging to the same cluster, an LLM can be used to summarize the common characteristics of the cluster, identify its main topics, and propose a concise human-readable label. The automatically generated interpretations can then be compared with ground-truth labels or human judgments.

The project explores a practical application of LLMs beyond text generation: using language models as tools for understanding and interpreting the structure discovered by unsupervised learning.

Objectives

The main objectives of the thesis are:

- Create text clusters using standard clustering methods.
- Select representative documents or samples from each cluster.
- Use an LLM to generate descriptions, summaries, or labels for the clusters.
- Investigate different prompting strategies for cluster interpretation.
- Compare automatically generated interpretations with reference labels or human judgments.
- Analyze the strengths and limitations of LLM-based cluster interpretation.

Possible research questions include:

- Can LLMs generate meaningful labels for automatically discovered text clusters?
- How does the selection of representative documents affect the generated interpretation?
- Which prompting strategies produce more consistent and informative cluster descriptions?
- To what extent do LLM-generated labels agree with human or dataset-provided labels?

Requirements

- Basic knowledge of Python.
- Basic knowledge of machine learning and data analysis.
- Familiarity with scikit-learn and standard NLP concepts.
- Basic understanding of LLMs is helpful but not required.
- Experience with Hugging Face Transformers is helpful.
- No prior experience with training or fine-tuning LLMs is required.

Possible Tools

Python, pandas, scikit-learn, Hugging Face Transformers, Sentence-Transformers, and an open-source LLM suitable for local inference.

References

1. Aggarwal, C. C., & Zhai, C. (2012). *A Survey of Text Clustering Algorithms*. In **Mining Text Data**, Springer, pp. 77–128.
2. Reimers, N., & Gurevych, I. (2019). *Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks*. In **Proceedings of EMNLP-IJCNLP 2019**, pp. 3982–3992.
3. Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). *BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding*. In **Proceedings of NAACL-HLT 2019**, pp. 4171–4186.

If you are interested in this proposal, please send a short email introducing yourself to Jian Lan (lan@dbis.fh-lmu.de) as soon as possible before 10.10.2026.