

# Deep Learning for Social Sciences

## Assignment 4: Fine-tuning Large Language Models for Political Text Classification

Giordano De Marzo

26 June 2025

**To be handed in by July 16th 23:59**

### 1 Overview

In this assignment, you will explore fine-tuning techniques for large language models using real parliamentary speech data. Unlike traditional machine learning approaches that rely on hand-crafted features, modern language models can learn rich representations directly from text through fine-tuning on domain-specific data.

The assignment uses the ParlaMint parliamentary corpus, containing speeches from various European parliaments. You will predict political party affiliation and ideological orientation from speech text, comparing fine-tuned model performance against pretrained ones to understand when domain adaptation provides additional predictive power.

**Complete all tasks for receiving points. Maximum points: 30. Bonus points available but will not exceed 30 total.**

#### 1.1 Dataset Information

- **Source:** ParlaMint parliamentary corpus (CLARIN repository)
- **Country:** Several countries available, I suggest trying with Austria, but you can choose the country you prefer
- **Files:** Speech transcripts (`.txt`) and metadata (`.tsv`) files
- **Data type:** Parliamentary speeches with speaker metadata
- **Target variables:** Political party affiliation, ideological orientation
- **Features:** Speech text, party affiliation, speaker role, date, session info
- **Time period:** Focus on the last 5 years of available data for temporal consistency

Available at: <https://www.clarin.si/repository/xmlui/handle/11356/1912>

### 2 Tasks

#### 2.1 Data Import, Preprocessing and Exploratory Analysis (5 points)

- Load and parse both speech transcripts and metadata files
- Create a unified dataset matching speeches to political parties and political orientation

- **Focus on the last 5 years** of available data to ensure temporal consistency
- Implement data cleaning and preprocessing:
  - Remove speeches shorter than 50 words
  - Filter out procedural and administrative speeches
  - Handle missing or incomplete party information
  - Remove duplicate speeches
- Perform comprehensive exploratory data analysis:
  - Analyze the distribution of speech lengths and political parties
  - Compute basic text statistics (vocabulary size, average speech length, etc.)
  - Visualize party distribution and temporal patterns
  - Create word clouds or frequency distributions for different parties

## 2.2 Model Training and Comparison (10 points)

Train and evaluate fine-tuned language models for **two classification tasks**:

### 2.2.1 Task A: Political Party Prediction

- **Target:** Multi-class classification of political party affiliation
- **Hypothesis:** Different parties use distinct vocabulary and rhetorical patterns
- **Base model:** Use a small pre-trained generative language model (e.g. Qwen, Gemma, Llama ...)
- **Features:** Raw speech text (no manual feature engineering)

### 2.2.2 Task B: Ideological Orientation Prediction

- **Target:** Multi-class classification (left/center/right etc)
- **Hypothesis:** Ideological positions manifest in language use and topic focus
- **Base model:** Same architecture as Task A
- **Features:** Raw speech text

**For each task:**

- Build a suitable dataset for instruction fine-tuning
- Use the pretrained model for the task as a baseline
- Implement **fine-tuning using LoRA** (Low-Rank Adaptation)
- Use appropriate train/validation/test splits (e.g. 70%/15%/15%)
- Report comprehensive performance metrics
- Include confusion matrices and classification reports

## 2.3 Report Writing (15 points)

Write a comprehensive report (max 3 pages) that includes:

### 2.3.1 Structure:

- **Introduction:** Problem description, political science context of the country analyzed, and technical approach
- **Data Analysis:** Dataset description, preprocessing steps, exploratory findings
- **Methods:** Model architectures, fine-tuning procedures
- **Results:** Performance comparison tables, confusion matrices, error analysis
- **Discussion and Conclusions:**
  - Discuss potential applications and limitations
  - Key findings and methodological insights

## 3 Bonus Tasks

You can use up to an additional half page for each bonus task.

### 3.1 Bonus Task 1: Cross-temporal Evaluation (1 bonus points)

- Train models on historical data and test on recent speeches
- Analyze how political language evolution affects model performance
- Discuss implications for temporal generalization in political NLP

### 3.2 Bonus Task 2: Attention Analysis and Interpretability (2 bonus points)

- Implement attention visualization techniques (e.g. SHAP Values)
- Identify which words/phrases are most important for party classification
- Analyze whether the model learns politically meaningful patterns
- Compare learned patterns with political science knowledge

**Note:** This assignment reflects real-world challenges in computational political science. Focus on understanding when and why domain-specific fine-tuning improves performance over general baselines, rather than achieving perfect accuracy.