

Project Title: Scalable NLP Systems for Large-Scale Sentiment Analysis in Public Discourse

Project Description: This project develops and deploys scalable natural language processing (NLP) pipelines for sentiment analysis of social media and other online text data. Using transformer-based models (e.g., RoBERTa, DeBERTa, and domain-adapted LLMs), zero-shot sentiment classification, and fine-tuned emotion lexicons, we aim to detect and track shifts in sentiment related to a range of topics, including well-being, policy, and public trust. The focus is on engineering robust, modular systems that can operate across domains and adapt to evolving language patterns in online discourse.

The ideal candidate is a student with strong Python and NLP experience, especially in Hugging Face Transformers, PyTorch or TensorFlow, and an interest in building real-time or batch sentiment analysis tools at scale.

Supervisor: Prof. Ethan Fosse

Application: Please submit your CV, unofficial transcript, and a single paragraph describing your interest in the project in an email to Prof. Ethan Fosse (ethan.fosse@utoronto.ca) and Prof. Nicholas Spence (nicholas.spence@utoronto.ca).

Start Date: Fall 2025

Research Area: Data Science, Machine Learning, Artificial Intelligence, Information Engineering