Leveraging Machine Learning Techniques to Design Campaign Policies

Konstantinos Vasiladiotis
Computer Engineering and Informatics
University of Patras
Patra, Greece
vassiladiotiskostis@gmail.com

Christos Makris
Computer Engineering and Informatics
University of Patras
Patra, Greece
makri@ceid.upatras.gr

Abstract

Social media has increasingly become a platform where political trends are formed, movements amplified, and both accurate and misleading information spread. Twitter (now X) exemplifies such platforms, yet the sheer volume of data prevents traditional analysis methods. This challenge motivated the development of *Election Analysis AI*, which automates sentiment analysis, enables semantic search, and generates political content. The system supports deeper exploration of user opinions during election periods and illustrates the ethical implications of AI in political campaigns. Two public datasets of political tweets were used, which were obtained from the Kaggle platform and concern the 2020 United States presidential election (Manchunhui 2020).

Introduction

This thesis presents an advanced artificial intelligence system designed to analyze Twitter data related to an election period, providing in-depth electoral forecasts and analyses. The Election Analysis AI processes thousands of tweets from the 2020 U.S. presidential elections. By leveraging natural language processing techniques, semantic search algorithms, and fine-tuned language models, it delivers real-time political insights through an interactive conversational interface. In the first part, the system architecture and technological foundation are presented. This agentic system is based on a three-tier architecture: a Streamlit-based frontend, a FastAPI (Grinberg 2021) REST API backend, and an intermediate service layer integrating AI and data processing components. The technology stack includes Python 3.9+, FastAPI for asynchronous endpoints, MongoDB for NoSQL data storage, and Ollama for local LLM inference. The embedding infrastructure uses Sentence-Transformers (Wolf et al. 2020) to create dense vector representations, enabling semantic search across the entire dataset. For tweet and post generation, the LoRA (Low-Rank Adaptation) fine-tuned version of Llama 3.2 is utilized, trained on political datasets to produce authentic, social media–ready content (Hu et al. 2022). Additional technologies include PyTorch for deep learning, PEFT for parameter-efficient fine tuning, and Pandas/NumPy for data processing. The second part focuses on a research and experimental machine learning stage of the project, which performs sentiment analysis on election-related data and predicts likely winners at the state level. The methodology combines deep learning with sentiment analysis (Liu 2012), employing a hybrid neural network (Hochreiter & Schmidhuber 1997) that integrates textual features (via LSTM/GRU) with numeric sentiment features (polarity and subjectivity) extracted using TextBlob. The implementation employs libraries such as TensorFlow, scikit-learn, NLTK, TextBlob, and Plotly. The third part presents the core methodology, which is based on a Retrieval-Augmented Generation (RAG) architecture (Lewis et al. 2020) that combines semantic search with large language model reasoning. The semantic tweet search relies on embedding-based similarity computation using cosine similarity, enhanced with the Maximal Marginal Relevance (MMR) algorithm for result diversity (Carbonell & Goldstein 1998). Tweet embeddings are processed in batches and cached in NumPy arrays and Parquet files, using configurable similarity thresholds (min_score) and top-k selection. Additional analyses identify active user time patterns through temporal recognition algorithms. LoRA fine-tuning allows efficient adaptation of Llama 3.2, supporting eight content generation topics and integrating content filtering and spam detection mechanisms.The fourth part evaluates the system and its results. The conversational interface integrates all the aforementioned features through natural language interaction, with the RAG service handling query classification, data retrieval, and context-aware response generation. Evaluation methods include quantitative metrics such as the F1-score for model accuracy, as well as Natural Language Interface (NLI) for the evaluation of the generated text (Mersinias 2023). Additionally, human evaluation was conducted to assess the quality of the fine-tuned model's generated tweets by comparing them against real Twitter data. Finally, the thesis discusses the experimental results and analyzes the system's effectiveness while proposing future extensions, including improving prediction accuracy through larger and more representative datasets. These prospects highlight the potential of the system as a comprehensive AI-driven platform for electoral trend analysis and prediction based on social media data.

System Functionality

Election Analysis AI incorporates several key functionalities:

1 Sentiment Analysis

Tweets are preprocessed to remove noise, standardize text, and extract key linguistic features. Using TextBlob (Loria 2022), the polarity and subjectivity of each tweet are computed, which are then mapped into categorical sentiment labels: Positive, Neutral, or Negative. This structured sentiment information is used as input for downstream modeling and provides insights into public opinion trends over time. Sentiment analysis was based on the work/repository (Chandra & Saini 2021).

2 Election Prediction

The system predicts candidate support at the state level using a hybrid LSTM/GRU model (Tumasjan et al. 2010). Numerical features, such as sentiment scores and aggregated state-level tweet statistics, are combined with embedding representations of tokenized tweets. These two branches are concatenated to form a unified input to the model. To incorporate sentiment influence per state, a weighted adjustment is applied, defined as:

$$state_{wmean} = state_{mean} + A * D_{sentiment}$$

Where state_mean represents the possibility of a tweet mentioning a candidate per state, D_sentiment reflects the sentiment toward the candidate and A is a hyperparameter controlling the influence of the sentiment. This formulation enables a more secure and balanced prediction that accounts for both frequency and emotional tone in social media posts.

3 Semantic Search

Users can submit natural language queries to explore political discussions across states and time periods. The system retrieves relevant tweets based on semantic similarity rather than keyword matching, allowing for more context – aware insights. Responses are evaluated using Natural Language Inference (NLI) metrics: Entailment, Neutral and Contradiction. Entailment measures whether the answer is directly supported by retrieval data, Neutral indicates partial or indirect alignment and Contradiction highlights conflicts with the source data. These metrics ensure that reliability and credibility of the responses provided to campaign analysts.

4 Political Content Generation

The system employs a LoRA fine-tuned Llama 3.2 model to generate realistic political tweets. It reproduces common stylistic elements such as slogans, topical references and emotional tone observed on Twitter (Ji et al. 2023). For experimental purposes, the model generates content from a selected stance while avoiding promotion of any real political agenda.

Human Evaluation Experiment

Human evaluation confirms that participants often struggle to distinguish AI-generated tweets from authentic ones, demonstrating the system's capacity to produce highly realistic content while raising ethical considerations regarding AI's use in political concepts. In our research 30 volunteers attempted to distinguish seven randomly presented tweets as AI-generated or real. Results revealed that users struggled to identify fake content (Lazer et al. 2018).

5 Trend and Activity Analysis

Temporal and engagement patterns analyzed to identify peak user activity hours, providing insight into when audiences are most likely to interact with content. The system also tracks frequently used hashtags and mentions, allowing for the identification of dominant discussion topics, influential accounts, and regional trends. By combining activity patterns with trend analysis, the system enables campaigns to optimize the timing and targeting of communications while monitoring emerging issues in near real time.

System Preview

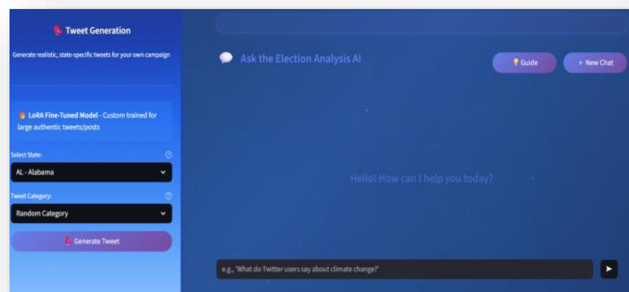

Conclusion

This work demonstrates that modern NLP (Bird et al. 2009) and data analysis techniques can transform raw social media data into actionable insights for political campaigns. *Election Analysis AI* integrates sentiment analysis, semantic search, and content generation to provide a comprehensive political analysis platform. At the same time, it highlights ethical challenges inherent in AI applications within politics, emphasizing the need for responsible use to avoid manipulation, misinformation, or undue influence on public opinion.

References

[1] Manchunhui. 2020. US Election 2020 Tweets. Kaggle.

[2] Chandra, S. & Saini, R. 2021. Sentiment Analysis of US Presidential Election using Deep Learning.

[3] Lewis, P., E. Perez, A. Piktus, et al. 2020. "Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks." Advances in Neural Information Processing Systems.

[4] Hu, E.J., Y. Shen, P. Wallis, et al. 2022. "LoRA: Low-Rank Adaptation of Large Language Models.

[5] Liu, B. 2012. *Sentiment Analysis and Opinion Mining.* San Rafael, CA: Morgan & Claypool Publishers.

[6] Ji, Z., N. Lee, R. Frieske, et al. 2023. “Survey of Hallucination in Natural Language Generation.” *ACM Computing Surveys* 55 (12): 1–38.

[7] Hochreiter, S., και J. Schmidhuber. 1997. “Long Short-Term Memory.” *Neural Computation* 9 (8): 1735–1780.

[8] Mersinias, Michail. 2023. For Generated Text, Is NLI-Neutral Text the Best Text?

[9] Bird, S., Klein, E., & Loper, E. 2009. *Natural Language Processing with Python.* O’Reilly Media.

[10] Loria, S. (2022). TextBlob: Simplified Text Processing.

[11] Wolf, T., L. Debut, V. Sanh, et al. 2020. Transformers: State-of-the-Art Natural Language Processing. Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations, 38–45.

[12] Carbonell, J., and J. Goldstein. 1998. *The Use of MMR, Diversity-Based Reranking for Reordering Documents and Producing Summaries.* Proceedings of the 21st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, 335–336.

[13] Lazer, D.M.J., M.A. Baum, Y. Benkler, et al. 2018. “The Science of Fake News.” *Science* 359 (6380): 1094–1096.

[14] Grinberg, M. 2021. *Building Web Applications with FastAPI.* Sebastopol, CA: O’Reilly Media.

[15] Tumasjan, A., T.O. Sprenger, P.G. Sandner, and I.M. Welpe. 2010. “Predicting Elections with Twitter: What 140 Characters Reveal about Political Sentiment.” *Proceedings of the Fourth International AAAI Conference on Weblogs and Social Media*: 178–185.