Overview on Political Multiclass Sentiment Analysis of Tamil X (Twitter) Comments: DravidianLangTech@NAACL 2025

Bharathi Raja Chakravarthi[1], Saranya Rajiakodi[2], Elizabeth Sherly[3], Thenmozhi Durairaj[4], Sathiyaraj Thangasamy[5], Ratnasingam Sakuntharaj[6], Prasanna Kumar Kumaresan[7], Kishore Kumar Ponnusamy[3], Arunaggiri Pandian Karunanidhi[8], Rohan R[4],

[1]University of Galway, Ireland,
[2]Central University of Tamil Nadu, India, [3]Digital University Kerala, Kerala, India,
[4]Sri Sivasubramaniya Nadar College of Engineering, Chennai,
[5]Sri Krishna Adithya College of Arts and Science, India, [6]Eastern University, Sri Lanka.
[7]Data Science Institute, University of Galway, Ireland, [8]Micron Technology, United states.

Abstract

Political multiclass detection is the task of identifying the predefined seven political classes. In this paper, we report an overview of the findings on the "Political Multiclass Sentiment Analysis of Tamil X(Twitter) Comments" shared task conducted at the workshop on DravidianLangTech@NAACL 2025. The participants were provided with annotated Twitter comments, which are split into training, development, and unlabelled test datasets. A total of 139 participants registered for this shared task, and 25 teams finally submitted their results. The performance of the submitted systems was evaluated and ranked in terms of the macro-F1 score.

1 Introduction

Online platforms are becoming the key platforms for the public conversation and the distribution of political news due to the quick development of digital and social media (Hermida et al., 2012; Kümpel et al., 2015; Tumasjan et al., 2010). Users may voice their thoughts, participate in conversations, and organize political movements with a reach and involvement previously unavailable on platforms like X (formerly Twitter) (Mustafaraj and Metaxas, 2011; Velasquez, 2012). Over the past decade, social media has fueled conversations on a wide range of divisive political topics, including climate change, gun control, abortion rights, income inequality, the death penalty, taxation policies, and LGBTQ+ rights (Rainie et al., 2012; Zhuravskaya et al., 2020). In addition to encouraging democratic participation and a range of ideas, these conversations often serve to magnify social prejudices, frequently reinforcing divisive opinions and political divisions (Blair, 2002; Devine, 1989).

As online political discourse expands, Natural language processing (NLP) models are increasingly being used to analyze public sentiment and opinion trends. However, many of these models are trained on vast datasets gathered from online sources, which inherently reflect existing societal biases. Political sentiment analysis is not solely a technological challenge but also involves issues of fairness and the ethical application of AI. (Blodgett et al., 2020; Kumar et al., 2022; Field et al., 2021). Numerous research studies have emphasized the dangers of bias in NLP models, such as incorrect sentiment categorization, unintentional reinforcement of ideological viewpoints, and distortion of minority voices (Nangia et al., 2020; Sun et al., 2019). Moreover, the subjective character of political state-of-mind labeling and differences in annotator viewpoints make attempts to create objective models much more challenging (Feng et al., 2023; Sap et al., 2019).

Sentiment analysis has advanced, but political expression poses special difficulties that need advanced strategies. Political conversations frequently contain sarcasm, coded language, and shifting rhetorical methods that are challenging for standard models to accurately interpret, unlike generic sentiment classification tasks where text is simply categorized as positive, negative, or neutral (Demszky et al., 2019). Furthermore, the framing of language is influenced by biases in political reporting and media coverage, making it significantly harder to train objective sentiment analysis models. (Joseph and Morgan, 2020).

This work presents a summary of the Political Multiclass Sentiment Analysis of Tamil X (Twitter) Comments shared task, which intends to improve multilingual and low-resource sentiment analysis research in order to overcome these issues. This work offers a chance to investigate the shortcomings of existing AI techniques for expressing sentiment in political situations by concentrating on Tamil, a linguistically rich language. The objective

Proceedings of the Fifth Workshop on Speech, Vision, and Language Technologies for Dravidian Languages, pages 104–111
May 3, 2025 ©2025 Association for Computational Linguistics

is to compare different strategies, find limitations in current techniques, and encourage improvements in the categorization of political perspective for under-resourced languages. We collected a dataset containing Tamil comments from X(Twitter) and then annotated the dataset for seven predefined classes. Then, we split it into training, development, and test sets for this task.

2 Related work

Several studies have explored sentiment analysis in Tamil, particularly focusing on social media platforms like Twitter. For instance, the study Anbukkarasi and Varadhaganapathy (2020) employed deep learning algorithms such as Long Short-Term Memory (LSTM) and Bidirectional LSTM (BiLSTM) to analyze Tamil tweets, achieving notable accuracy and F1-scores.

Another study, Thavareesan and Mahesan (2019) investigated various machine learning approaches for sentiment classification in Tamil texts, contributing to the understanding of effective methods for Tamil sentiment analysis. Additionally, Shanmugavadivel et al. (2022) addressed the challenges of analyzing sentiments in code-mixed Tamil texts, which are common in social media contexts. This study utilized machine learning techniques to classify sentiments in such code-mixed data.

Furthermore, Mahata et al. (2020) explored sentiment classification in code-mixed Tamil-English tweets using a Bi-Directional Recurrent Neural Network (RNN) approach, highlighting the complexities and solutions in handling mixed-language data. In addition to these studies, (Anish and Sumathy) proposed an SVM-based approach to analyze sentiments in Tamil political reviews, while (Devasena et al., 2022) demonstrated how sentiment analysis could be applied to predict election results based on Twitter data.

Sharmista and Ramaswami (2020), examined Tamil sentiment classification in the context of product reviews, showcasing its relevance in different domains. Lastly, (Shanmugavadivel et al., 2022) explored embedding representations for code-mixed Tamil text, addressing the challenges posed by multilingual and informal social media content.

These studies collectively contribute to the advancement of sentiment analysis methodologies for Tamil, particularly in the context of social media data. However, limited research exists on political multiclass sentiment analysis in Tamil, which involves classifying sentiments into multiple nuanced categories beyond the traditional positive, negative, and neutral classes. Our work aims to bridge this gap by introducing a detailed classification scheme tailored to Tamil political discourse.

3 Task Description

The primary goal of this task is to detect the political categories in the comments collected from X (Twitter). The participants were provided with training, development, and test datasets. The dataset is tagged using 7 classes namely, Substantiated, Sarcastic, Opinionated, Positive, Negative, Neutral and None of the above. Further information on the task is available in the Codalab site[1].

3.1 Datasets

The dataset containing Tamil text is the social media comments collected from X(Twitter). The diverse political sentiments captured in the dataset aim to reflect real-world nuances, making it well-suited for the multiclass sentiment analysis task. The dataset was divided into training, development, and testing sets. Training and validation sets are provided with class labels, and test sets are provided as unlabeled ones for evaluation. The data distribution and class distribution of training, validation, and test sets are given in Table 1

Table 1: Data Distribution

Class	Train	Development	Test	Total
Substantiated	412	52	51	515
Sarcastic	790	115	106	1,011
Opinionated	1,361	153	171	1,685
Positive	575	69	75	719
Negative	406	51	46	503
Neutral	637	84	70	791
None of the above	171	20	25	216
Tamil	4,352	544	544	5,440

4 Methodology

Totally 25 teams have actively participated in this shared task to detect the political comments in tamil. The participants have explored a variety of methodologies to classify the given comment as predefined political classes

Synapse team (KP et al., 2025) focused on preprocessing and fine-tuning to address class imbal-

[1] https://codalab.lisn.upsaclay.fr/competitions/20702

ance and optimize performance. During preprocessing, they converted emojis to text and expanded the top 160 most repeated hashtags to their full forms for better semantic understanding. For the model, they have finetuned IndicBERTv2-MLM-Back-TLM encoder based LLM model which was trained on IndicCorp v2 and Samanantar datasets, and an additional task of Translation. The fine-tuning was performed using the AutoModelForSequenceClassification architecture, incorporating class weights to rectify the class imbalance effectively. The team utilized only the train dataset for this fine-tuning process.

KCLR team (Mia et al., 2025) adopted a transformer-based deep learning architecture enhanced with multi-faceted embedding techniques. This approach combines three distinct feature extraction methods: attention-weighted representations, and CLS token embeddings from the transformer outputs. These features are concatenated to create comprehensive sentence representations before being processed through a fully connected classification layer. To address data imbalance challenges, the team implemented oversampling of minority classes using scikit-learn's resample function, ensuring robust and balanced training across all categories. This integrated approach, combining advanced feature engineering with balanced training data, enables effective multi-class classification while maintaining model robustness.

byteSizedLLM team implemented an advanced hybrid methodology combining a customized attention BiLSTM network with an XLM-RoBERTa base model, which had already been fine-tuned on the AI4Bharat dataset using Masked Language Modeling (MLM). The AI4Bharat dataset included fully and partially transliterated text, with 20–70 percentage of words randomly transliterated, enhancing transliteration-based diversity. This approach allows it to learn robust cross-lingual representations and adapt to varied transliteration patterns. The team further fine-tuned this pre-trained model and integrated BiLSTM and attention layers to capture sequential dependencies, making the model highly effective for multilingual and transliteration-heavy tasks.

Eureka-CIOL team (Eram et al., 2025) began by analyzing the dataset and identified that it consists of Tamil text with six distinct sentiment classes. Their best-performing model utilized a multilingual custom model pre-trained on general Twitter sentiment data, which allows for handling the diverse nature of social media content. To adapt the model for the specific task of sentiment classification, they applied a Multi-Layer Perceptron (MLP) on top of this pre-trained model, enabling fine-tuning. This approach leveraged the multilingual capabilities of the model and the domain-specific knowledge from general sentiment data. Finally the fine-tuned model was used to generate the predictions on the test data.

Wictory team (K et al., 2025) employed specific preprocessing techniques to prepare the data, including demojifying the text and removing unwanted characters. For their model, they converted word embeddings for generated LaBSE (Language-agnostic BERT Sentence Embedding), which were then passed into a Support Vector Machine (SVM) for classification.

MNLP team implemented the Deep Learning based model which was fine-tuned for classification. Their model achieved a 0.3026 macro F1-score and ranked 6th in the shared task.

Nova Spark developed a text classification pipeline for Tamil and English, involving text normalization, tokenization, and TF-IDF vectorization. To handle class imbalance, Borderline-SMOTE, SMOTEENN, and ADASYN were used. An optimized Support Vector Classifier (SVC) was trained using GridSearchCV for the best macro F1-score. Performance was evaluated with a classification report, and final predictions were saved as a CSV for submission.

Team_Catalysts (Shanmugavadivel et al., 2025a) implemented a robust Tamil text classification pipeline, including Unicode normalization, tokenization with Stanza, and standardization of spoken variants. Class imbalance was addressed through upsampling, followed by TF-IDF vectorization. A Random Forest Classifier was trained using stratified splitting and evaluated with accuracy and classification reports, ensuring effective sentiment analysis.

Lowes team began by preprocessing the dataset to prepare it for analysis. They then fine-tuned a BERT-based model specifically for the task. Their model achieved a 0.2908 macro F1-score and ranked 9th in the shared task.

Abhay43 team applied simple preprocessing to the dataset. They then extracted embeddings using the DeBERTa v3 model, which were subsequently fed into a two-layered LSTM model. They achieve a macro F1-score of 0.2904 and ranked tenth

GS Team explored several machine learning approaches including Logistic regression, random forest classifier, support vector machine, and XGBBoost classifier with TFIDF vectorization techniques for feature extraction techniques. Among these models, the XGBoost model outperformed the other models. Similarly, **JAS** team employed Logistic Regression as a primary approach for this classification task.

SentiTamil team utilized classical machine learning approaches, specifically support vector machine (SVM), with TFIDF vectorizer, limiting the number of features to 5,000 for efficiency. They also tried to fine-tune the tamil-llama-7b model, however the predicted value is not similar as the gold label of the training dataset.

CrewX team leveraged IndicBERT, a multilingual language model tailor for Indian languages, as the backbone for political multiclass sentiment analysis of Tamil Twitter comments. The dataset was preprocessed to handle challenges such as code-mixing, transliteration, and noise typical in social media text. Tokenization was performed using IndicBERT's tokenizer to preserve linguistic nuances. The team fine-tuned the pre-trained IndicBERT model on the DravidianLangTech dataset, utilizing a classification head with softmax activation to predict sentiment classes. To enhance performance, they experimented with techniques like data augmentation, stratified sampling, and weighted loss to address class imbalance. The model was trained using cross-entropy loss and optimized with AdamW, while employing early stopping to prevent overfitting. Evaluation metrics, including accuracy, F1-score, and precision-recall, were used to assess the model's effectiveness. This approach leverages IndicBERT's contextual understanding to address the intricacies of Tamil sentiment analysis in a political context.

AnalysisArchitects team (Jayaraman et al., 2025) implemented a diverse methodology by employing Naive Bayes, SVM, and LSTM models for the task of multiclass sentiment analysis. For the Naive Bayes approach, the team preprocessed the text, transformed it using CountVectorizer, and trained the model for multiclass sentiment analysis. Predictions were then generated on a test dataset, and the results were saved as a CSV file. The SVM model utilized TF-IDF features for text representation. After preprocessing the text, the team trained an SVM classifier and evaluated its performance on a test dataset. Predictions for the separate test set were also saved as a CSV file. This method tokenizes and pads Tamil text sequences, then trains an LSTM model for sentiment analysis. The model uses an embedding layer, LSTM for sequence learning, and a softmax output for classification. Input dimensions are adjusted, and sequence values are clipped to stay within valid range

Beyond_tech team (Shanmugavadivel et al., 2025b) utilized a combination of natural language processing techniques and pattern recognition to extract relevant information and generate appropriate responses. The methodology involved analyzing the task description and context, followed by segmenting the input into smaller, manageable parts. Each segment was processed to identify key concepts and relationships, facilitating the formulation of precise and coherent outputs. To ensure continuous improvement, the team applied an iterative feedback loop for refinement and alignment with task requirements. This approach allowed for efficient handling of complex queries, maintaining accuracy and clarity in response generation.

CUET_Novice team (Barua et al., 2025) utilized multiple deep learning architectures for their methodology. In the first approach (run1), they utilized a model with stacked Bidirectional GRU (BiGRU) layers, followed by normalization and a feedforward neural network for classification. In the second approach (run2), they utilized a model with multiple Bidirectional LSTM (BiLSTM) layers, similarly they applied normalization and a feedforward neural network. For the third approach (run3), they employed a transformer-based model, leveraging its advanced contextual understanding capabilities. This diverse experimentation with GRUs, LSTMs, and transformers allowed the team to explore various architectures for optimal

KSK team (M et al., 2025) implemented an incremental and continual learning for political multiclass sentiment analysis of Tamil tweets focusing on adapting models to new data while retaining prior knowledge. Algorithms like Stochastic Gradient Descent (SGD) and Online Naive Bayes dynamically update parameters for evolving sentiments. The team also utilized Incremental SVMs and Hoeffding Trees, enabling efficient updates without retraining on the entire dataset. Pretrained models like multilingual BERT are fine-tuned continually to adapt to new linguistic patterns while avoiding catastrophic forgetting. Online ensemble methods further enhance robustness, making them suitable for evolving Twitter data streams.

QuanNguyen team utilized the BERT multilingual base model (cased) to perform multiclass sentiment analysis on Tamil X (Twitter) comments. The data preprocessing involved identifying and categorizing hashtags and icons uniquely associated with each sentiment class while removing special characters and irrelevant symbols for cleaner input. The multilingual BERT model, well-suited for handling multiple languages including Tamil, was fine-tuned on the preprocessed dataset to capture contextual and semantic patterns in sentiment. While BERT formed the core of the system, the team noted the potential for exploring other deep learning models to further enhance performance.

Team_Luminaries_0227 team began by preprocessing the dataset, including cleaning text data. They utilized the TF-IDF vectorizer to convert the textual data into numerical representations. To address class imbalance in the dataset, They applied the SMOTE (Synthetic Minority Oversampling Technique) algorithm, ensuring balanced class distributions. For classification, a Random Forest classifier was trained, with performance evaluated using metrics such as precision, recall, and F1-score. The trained models were saved for later use, and predictions were generated on the test dataset, ensuring the methodology aligns with the objective of the task.

VKG_VELLORE INSTITUTE OF TECHNOLOGY team utilized classification pipeline by extracting features from a pre-trained Indic-BERT language model, and then DBOW and TF-IDF methods were applied followed by CatBoost classifier for text classification. For better performance, they performed preprocessing steps like removing special characters and converting text to lowercase. After tokenizing the text using the BERT tokenizer, Indic-BERT embeddings were created, transforming the input text into dense representations rich in contextual information. To address the class imbalance, they used SMOTE (Synthetic Minority Oversampling Technique) to balance the training dataset.Embedded data warmed-up a CatBoost classifier for the reason that it is adept at dealing with categorical nearest neighbor features and unbalanced data sets. For evaluation, the team applied a 90:10 train-validation split and macro-averaged metrics were employed to allow for a comprehensive performance appraisal This method effectively combines the advantages of pre-trained embeddings and a powerful gradient boosting model, yielding accurate multi-class classification.

CUET_NetworkSociety team (Babu et al., 2025) employed a transformer-based approach using the 'bert-base-multilingual-cased' model for text classification. The data preprocessing includes normalization and label encoding. The team utilized the Hugging Face 'Trainer' class for fine-tuning with tokenized inputs, optimized hyperparameters, and mixed precision ('fp16') was implemented to enhance computational efficiency during training.

Walter White team utilized the Indic BERT model, which is well-suited for code-mixed data and effectively handles Tamil-specific linguistic features. During the preprocessing stage, the team replaced emojis with their corresponding textual descriptions but excluded those irrelevant to the context (e.g., the kite emoji). They also removed newline characters, hashtags, and normalized spaces for consistency. For tokenization, they opted for the Trivial Tokenizer, as it is compatible with both Indic BERT and the Tamil language.

YenCS team implemented a multi-step approach to text classification. Initially, the text data is preprocessed by cleaning and tokenizing it. Then, word embeddings are generated using a pre-trained word2vec model. Three different deep learning models are trained: a Convolutional Neural Network (CNN) with a GRU layer, an LSTM model, and an LSTM model with an added GRU layer. These models are then combined using a stacking ensemble technique, where the predictions of the individual models serve as input features for a meta-model (RandomForestClassifier). Finally, the meta-model makes the final prediction, aiming to improve the overall classification accuracy compared to using any single model alone. The process is further enhanced by using early stopping and hyperparameter tuning to optimize model performance.

ARINDASCI team performed political sentiment classification using a multi-step machine learning pipeline. Initially they preprocessed the data by removing the noise like special characters, URLs, and whitespaces. Then they tokenized and used pre-trained embeddings (e.g., fastText or TamilBERT) to capture the semantic informations For classification, the team experimented with various models, including traditional machine learning algorithms like Logistic Regression and advanced deep learning models such as LSTMs and Transformer-based architectures. The system achieved a macro-F1-score of 0.0727on the test set.

5 Results and Discussion

There was a total of 139 people who registered for this shared task, and 25 teams submitted their results. The ranking for Tamil was determined based on the macro F1-score, as shown in Table 2. The Synapse team secured first place with an F1-score of 0.377 by fine-tuning the IndicBERTv2-MLM-Back-TLM encoder-based LLM, leveraging IndicCorp v2 and Samanantar datasets. The KCLR team followed closely in second place, achieving a score of 0.371 with a transformer-based deep learning model enhanced through diverse embedding techniques. The byteSizedLLM team ranked third with an F1-score of 0.349, employing a hybrid approach that integrated a customized attention BiLSTM network with a fine-tuned XLM-RoBERTa base model.

Table 2: **Task: Tamil Rank list**

Team Name	F1-score	Rank
Synapse (KP et al., 2025)	0.3773	1
KCRL (Mia et al., 2025)	0.3710	2
byteSizedLLM	0.3497	3
Eureka-CIOL (Eram et al., 2025)	0.3187	4
Wictory (K et al., 2025)	0.3115	5
MNLP	0.3026	6
Nova Spark	0.3001	7
Team_Catalysts (Shanmugavadivel et al., 2025a)	0.2933	8
Lowes	0.2908	9
abhay43	0.2904	10
GS	0.2835	11
JAS	0.2796	12
SentiTamil	0.2769	13
CrewX	0.2759	14
AnalysisArchitects (Jayaraman et al., 2025)	0.2747	15
Beyond_tech (Shanmugavadivel et al., 2025b)	0.2736	16
CUET_Novice (Barua et al., 2025)	0.2728	17
KSK (M et al., 2025)	0.2654	18
QuanNguyen	0.2613	19
Team_Luminaries_0227	0.2530	20
VKG	0.2526	21
CUET_NetworkSociety (Babu et al., 2025)	0.2178	22
WalterWhite	0.1554	23
YenCS	0.1333	24
ARINDASCI_Tamil	0.0727	25

6 Conclusion

The "Political Multiclass Sentiment Analysis of Tamil X (Twitter) Comments" shared task provided valuable insights into the classification of Tamil political comments from social media. As part of the DravidianLangTech@NAACL workshop, this task challenged participants to categorize comments into seven predefined classes using diverse machine learning, deep learning, and natural language processing approaches. With 25 participating teams, model performance was assessed using the macro-F1 score. Given the small dataset size, few-shot and zero-shot learning strategies could enhance model efficiency. Furthermore, integrating Explainable AI (XAI) techniques can improve transparency and interpretability, fostering trust in model predictions and advancing sentiment analysis for low-resource languages like Tamil.

Acknowledgments

This work was conducted with the financial support from Science Foundation Ireland (SFI) under Grant Number SFI/12/RC/2289_P2(Insight_2), supported in part of Science Foundation Ireland Centre for Research Training in Artificial Intelligence under Grant No. 18/CRT/6223.

References

S Anbukkarasi and S Varadhaganapathy. 2020. Analyzing sentiment in Tamil tweets using deep neural network. In *2020 Fourth International Conference on Computing Methodologies and Communication (ICCMC)*, pages 449–453. IEEE.

D Anish and V Sumathy. Sentiment extraction for Tamil political.

Tofayel Ahmmed Babu, MD Musa Kalimullah Ratul, Sabik Aftahee, Jawad Hossain, and Mohammed Moshiul Hoque. 2025. CUET_NetworkSociety@DravidianLangTech 2025: A Transformer-Driven Approach to Political Sentiment Analysis of Tamil X (Twitter) Comments. In *Proceedings of the Fifth Workshop on Speech, Vision, and Language Technologies for Dravidian Languages*. Association for Computational Linguistics.

Arupa Barua, Md Osama, and Ashim Dey. 2025. CUET_Novice@DravidianLangTech 2025: A Bi-GRU Approach for Multiclass Political Sentiment Analysis of Tamil Twitter (X) Comments. In *Proceedings of the Fifth Workshop on Speech, Vision, and Language Technologies for Dravidian Languages*. Association for Computational Linguistics.

Irene V Blair. 2002. The malleability of automatic stereotypes and prejudice. *Personality and social psychology review*, 6(3):242–261.

Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. 2020. Language (technology) is power: A critical survey of" bias" in nlp. *arXiv preprint arXiv:2005.14050*.

Dorottya Demszky, Nikhil Garg, Rob Voigt, James Zou, Matthew Gentzkow, Jesse Shapiro, and Dan Jurafsky. 2019. Analyzing polarization in social media: Method and application to tweets on 21 mass shootings. *arXiv preprint arXiv:1904.01596*.

K Devasena, M Sarika, and J Shana. 2022. Predicting Tamil nadu election 2021 results using sentimental analysis before counting. In *Proceedings of the International Conference on Computational Intelligence and Sustainable Technologies: ICoCIST 2021*, pages 279–289. Springer.

Patricia G Devine. 1989. Stereotypes and prejudice: Their automatic and controlled components. *Journal of personality and social psychology*, 56(1):5.

Enjamamul Haque Eram, Anisha Ahmed, Sabrina Afroz Mitu, and Azmine Toushik Wasi. 2025. Eureka-CIOL@DravidianLangTech 2025: Using Customized BERTs for Sentiment Analysis of Tamil Political Comments. In *Proceedings of the Fifth Workshop on Speech, Vision, and Language Technologies for Dravidian Languages*. Association for Computational Linguistics.

Shangbin Feng, Chan Young Park, Yuhan Liu, and Yulia Tsvetkov. 2023. From pretraining data to language models to downstream tasks: Tracking the trails of political biases leading to unfair NLP models. *arXiv preprint arXiv:2305.08283*.

Anjalie Field, Su Lin Blodgett, Zeerak Waseem, and Yulia Tsvetkov. 2021. A survey of race, racism, and anti-racism in NLP. *arXiv preprint arXiv:2106.11410*.

Alfred Hermida, Fred Fletcher, Darryl Korell, and Donna Logan. 2012. Share, like, recommend: Decoding the social media news consumer. *Journalism studies*, 13(5-6):815–824.

Abirami Jayaraman, Aruna Devi Shanmugam, Dharunika Sasikumar, and Bharathi B. 2025. AnalysisArchitects@DravidianLangTech 2025: Machine Learning Approach to Political Multiclass Sentiment Analysis of Tamil. In *Proceedings of the Fifth Workshop on Speech, Vision, and Language Technologies for Dravidian Languages*. Association for Computational Linguistics.

Kenneth Joseph and Jonathan H Morgan. 2020. When do word embeddings accurately reflect surveys on our beliefs about people? *arXiv preprint arXiv:2004.12043*.

Nithish Ariyha K, Eshwanth Karti T R, Yeshwanth Balaji A P, Vikash J, and Sachin Kumar S. 2025. Wictory@DravidianLangTech 2025: Political Sentiment Analysis of Tamil X(Twitter) Comments using LaBSE and SVM. In *Proceedings of the Fifth Workshop on Speech, Vision, and Language Technologies for Dravidian Languages*. Association for Computational Linguistics.

Suriya KP, Durai Singh K, Vishal A S, Kishor S, and Sachin Kumar S. 2025. Synapse@DravidianLangTech 2025: Multiclass Political Sentiment Analysis in Tamil X (twitter) Comments: Leveraging Feature Fusion of IndicBERTv2 and Lexical Representations. In *Proceedings of the Fifth Workshop on Speech, Vision, and Language Technologies for Dravidian Languages*. Association for Computational Linguistics.

S Kumar, V Balachandran, L Njoo, A Anastasopoulos, and Y Tsvetkov. 2022. Language generation models can cause harm: so what can we do about it. *An actionable survey. CoRR abs/2210.07700*.

Anna Sophie Kümpel, Veronika Karnowski, and Till Keyling. 2015. News sharing in social media: A review of current research on news sharing users, content, and networks. *Social media+ society*, 1(2):2056305115610141.

Thissyakkanna S M, Kalaivani K S, Sanjay R, and NIRENJHANRAM S K. 2025. KEC_AI_KSK@DravidianLangTech 2025: Political Multiclass Sentiment Analysis of Tamil X (Twitter) Comments Using Incremental Learning. In *Proceedings of the Fifth Workshop on Speech, Vision, and Language Technologies for Dravidian Languages*. Association for Computational Linguistics.

Sainik Kumar Mahata, Dipankar Das, and Sivaji Bandyopadhyay. 2020. JUNLP@ Dravidian-CodeMix-FIRE2020: Sentiment classification of code-mixed tweets using bi-directional RNN and language tags. *arXiv preprint arXiv:2010.10111*.

Md Ayon Mia, Fariha Haq, Md. Tanvir Ahammed Shawon, Golam Sarwar Md. Mursalin, and MUHAMMAD IBRAHIM KHAN. 2025. KCRL@DravidianLangTech 2025: Multi-View Feature Fusion with XLM-R for Tamil Political Sentiment Analysis. In *Proceedings of the Fifth Workshop on Speech, Vision, and Language Technologies for Dravidian Languages*. Association for Computational Linguistics.

Eni Mustafaraj and Panagiotis Takis Metaxas. 2011. What edited retweets reveal about online political discourse. In *Workshops at the Twenty-Fifth AAAI Conference on Artificial Intelligence*.

Nikita Nangia, Clara Vania, Rasika Bhalerao, and Samuel R Bowman. 2020. Crows-pairs: A challenge dataset for measuring social biases in masked language models. *arXiv preprint arXiv:2010.00133*.

Lee Rainie, Aaron Smith, Kay Lehman Schlozman, Henry Brady, Sidney Verba, et al. 2012. Social media and political engagement. *Pew Internet & American Life Project*, 19(1):2–13.

Maarten Sap, Dallas Card, Saadia Gabriel, Yejin Choi, and Noah A Smith. 2019. The risk of racial bias in hate speech detection. In *Proceedings of the 57th annual meeting of the association for computational linguistics*, pages 1668–1678.

Kogilavani Shanmugavadivel, Sai Haritha Sampath, Pramod Nandhakumar, Prasath Mahalingam, Malliga Subramanian, Prasanna Kumar Kumaresan, and Ruba Priyadharshini. 2022. An analysis of machine learning models for sentiment analysis of Tamil code-mixed data. *Computer Speech & Language*, 76:101407.

Kogilavani Shanmugavadivel, Malliga Subramanian, Subhadevi K, Sowbharanika Janani Sivakumar, and Rahul K. 2025a. Team_Catalysts@DravidianLangTech-NAACL 2025: Leveraging Political Sentiment Analysis using Machine Learning Techniques for Classifying Tamil Tweets. In *Proceedings of the Fifth Workshop on Speech, Vision, and Language Technologies for Dravidian Languages*. Association for Computational Linguistics.

Kogilavani Shanmugavadivel, Malliga Subramanian, Sanjai R, Mohammed sameer, and Motheeswaran K. 2025b. Beyond_Tech@DravidianLangTech 2025: Political Multiclass Sentiment Analysis using Machine Learning and Neural Network. In *Proceedings of the Fifth Workshop on Speech, Vision, and Language Technologies for Dravidian Languages*. Association for Computational Linguistics.

A Sharmista and Dr M Ramaswami. 2020. Sentiment analysis on Tamil reviews as products in social media using machine learning techniques: A novel study. *Madurai Kamaraj University Madurai-625*, 21.

Tony Sun, Andrew Gaut, Shirlyn Tang, Yuxin Huang, Mai ElSherief, Jieyu Zhao, Diba Mirza, Elizabeth Belding, Kai-Wei Chang, and William Yang Wang. 2019. Mitigating gender bias in natural language processing: Literature review. *arXiv preprint arXiv:1906.08976*.

Sajeetha Thavareesan and Sinnathamby Mahesan. 2019. Sentiment analysis in Tamil texts: A study on machine learning techniques and feature representation. In *2019 14th Conference on industrial and information systems (ICIIS)*, pages 320–325. IEEE.

Andranik Tumasjan, Timm Sprenger, Philipp Sandner, and Isabell Welpe. 2010. Predicting elections with Twitter: What 140 characters reveal about political sentiment. In *Proceedings of the international AAAI conference on web and social media*, volume 4, pages 178–185.

Alcides Velasquez. 2012. Social media and online political discussion: The effect of cues and informational cascades on participation in online political communities. *New Media & Society*, 14(8):1286–1303.

Ekaterina Zhuravskaya, Maria Petrova, and Ruben Enikolopov. 2020. Political effects of the internet and social media. *Annual review of economics*, 12(1):415–438.