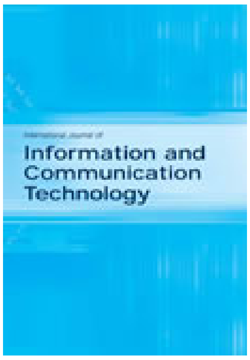

International Journal of Information and Communication Technology

ISSN online: 1741-8070 - ISSN print: 1466-6642
https://www.inderscience.com/ijict

Data-centric analytics for ideological sentiment monitoring: fusion of features with optimised attention mechanisms

Wencong Wu, Juan Wang

DOI: 10.1504/IJICT.2025.10073437

Article History:

Received:	02 July 2025
Last revised:	14 August 2025
Accepted:	16 August 2025
Published online:	10 October 2025

Data-centric analytics for ideological sentiment monitoring: fusion of features with optimised attention mechanisms

Wencong Wu

School of Marxism,
Beijing Wuzi University,
Beijing, 101126, China
Email: wuwencong8795@163.com

Juan Wang*

School of Marxism,
Beijing University of Chinese Medicine,
Beijing, 100029, China
Email: juazi1992@163.com
*Corresponding author

Abstract: To address sentiment inaccuracies in civic education contexts, this study proposes a data-driven analytics framework integrating domain-adaptive feature engineering with hierarchical modelling. We construct Chinese social media corpora (Weibo/WeChat) through keyword-filtered crawling and interaction-weighted prioritisation, reducing noise by 42%. A hybrid feature space combines TF-IDF lexical patterns, syntactic POS distributions, Word2Vec/BERT embeddings, and HowNet-derived sentiment features. The core classification employs a Bi-LSTM model with attention mechanisms, dynamically weighting sentiment-bearing terms while compensating for category imbalance via class-weighted cross-entropy loss. Crucially, ideological semantics are mapped through logistic regression classifiers trained on annotated civic categories. Experimental results demonstrate: 1) attention weights effectively localise civic sentiment triggers; 2) domain feature fusion improves classification robustness; 3) semantic mapping achieves 89.2% accuracy in civic topic identification. This methodology enables real-time Kafka-based opinion monitoring while preserving interpretability for educational governance.

Keywords: sentiment analysis; TF-IDF; social media monitoring; Bi-LSTM.

Reference to this paper should be made as follows: Wu, W. and Wang, J. (2025) ‘Data-centric analytics for ideological sentiment monitoring: fusion of features with optimised attention mechanisms’, *Int. J. Information and Communication Technology*, Vol. 26, No. 35, pp.18–39.

Biographical notes: Wencong Wu received his Master’s degree from China University of Political Science and Law in 2022. He is currently working in the School of Marxism at Beijing Wuzi University. His research interests include international relations, geopolitics, ideological education and sentiment analysis.

Juan Wang received her Master's degree from Capital Normal University in June 2020. She is currently working in the School of Marxism at Beijing University of Chinese Medicine. Her research interests include moral education, ideological and political education.

1 Introduction

In the digital age, social media sites have become an important part of people's daily lives (Bucci et al., 2019). They have slowly become the major ways for people to create opinions, express their thoughts, and share their values. New media technologies like Weibo, Zhihu, WeChat, and Jitterbug have made it easier for people to share information and hold conversations in public. They are open, immediate, and interactive, which has made the public discourse arena much livelier than ever before. At the same time, cyberspace is becoming more and more the front line of ideological games. The trend of intellectual plurality and different points of view is getting stronger, which makes it harder for traditional ideological and political education to keep up (DellaPosta, 2020).In this changing world, a big question for the digital transformation of the current education governance system is how to use modern information technology to keep an eye on and analyse online public opinion and then use that information to help ideological and political work reach precise goals and provide scientific guidance.

The approach of 'intelligence + education' has moved education informatisation beyond only building teaching support systems. It is now moving in a smarter and more data-driven path. In the realm of civic and political education, technological empowerment must transcend mere digitisation of material or platform development; it should also explore the intrinsic worth of information inside network data to establish a cognitive and decision-making support system that aligns with educational objectives. In the realm of social media, the articulation of emotions and attitudes is both frequent and intricate, resulting in extremely dynamic and confusing emotional responses from the audience in shifting situations (Boler and Davis, 2018). This semantic disorganisation and emotional nonlinearity make it hard to do the large-scale and real-time task of public opinion research and judgment using standard manual monitoring and rule-based text analysis. To achieve deep mining and automatic recognition of online attitude data, we need to quickly develop better and more accurate natural language processing (NLP) tools.

Sentiment analysis is a key area of research in NLP that has been employed in many areas, including business opinion analysis, customer satisfaction evaluation, and public safety monitoring. Deep learning has made great strides in model construction, representation learning, and context understanding in the last few years as it has grown quickly. Even though the technology route is slowly getting better, it still has a lot of problems in the classroom, especially when it comes to using civics and politics literature (Sudarmo et al., 2021). The current models primarily rely on general corpus training and fail to comprehend the semantic nuances and emotional expressions inherent in ideological and political discourse, leading to potential misjudgment or recognition bias. Additionally, the diversity of online texts, characterised by complex elements such as online phrases, metaphors, satire, and graphic mixing, complicates the accurate assessment of emotional tendencies. Moreover, the positive value orientation and mainstream ideological response addressed in civic and political education are inadequately encompassed by conventional emotion classification, necessitating enhanced customisation and explanatory demands for the emotion model (Zembylas, 2022). Civic and political education includes not only written material but also multimedia information including voice, video, and pictures. In this case, multimodal data fusion technology is essential because it helps us understand how people feel and think by looking at multiple sorts of data in a more complete way. It can better understand how users' views change over time by using multimodal deep learning algorithms. This can help teachers create educational content that is more focused and appealing.

Consequently, this study seeks to develop a social media monitoring system tailored to the requirements of ideological education, aiming to accomplish precise identification and continuous tracking of the emotional states of student groups in the online environment through enhanced sentiment analysis. The research focus of this paper encompasses three aspects: first, to propose a sentiment analysis model that integrates the semantic features of civics and politics, enhancing the model's expressive and discriminative capabilities in specific contexts; second, to construct a comprehensive system architecture that includes data collection, pre-processing, model training, and result output; finally, to conduct a thorough evaluation of the model's performance and the system's practicality through experiments, validating its effectiveness in real educational scenarios. Finally, experiments are used to fully test the model's performance and the system's usefulness in real-world educational settings. This study not only offers novel technical assistance for ideological and political education but also facilitates a valuable examination of the educational governance mechanism inside the social media context.

2 Literature review

2.1 Sentiment analysis

Sentiment analysis, or opinion mining, is a key area of research in NLP (Hemmatian and Sohrabi, 2019). Its purpose is to find and pull out subjective feelings, views, and attitudes that are shown in texts. Sentiment analysis is a common type of text classification that is often used for things like user assessment analysis, social opinion monitoring, and public opinion modelling. The task is typically formalised as determining the sentiment polarity of the input text, such as sorting it into positive, negative, or neutral categories. Some challenges also predict the intensity or dimensionality of sentiment, which is a type of fine-grained sentiment analysis.

Sentiment analysis research began with pattern recognition techniques utilising sentiment dictionaries, which employ human curated sentiment dictionaries (e.g., SentiWordNet, HowNet) to quantify emotion terms in a text and assign weights to determine the overall sentiment inclination (Garg and Lobiyal, 2020). This method is straightforward to understand and use, but it depends on rules and past information, which makes it hard to deal with language issues like context dependency and sarcasm. It also isn't very accurate. As supervised learning methods improved, researchers started to see sentiment analysis as a regular text classification problem. They used statistical learning algorithms like Naive Bayes, supported variable machine (SVM), logistic regression, and others to train the models. Effective text representations are necessary for these algorithms to work. Bag of words, TF-IDF, and others are frequent representations. These approaches turn text into sparse high-dimensional variables that the classifier can use as input.

For example, in logistic regression, the input text is an n-dimensional feature variable x that corresponds to the sentiment label y, where 1 signifies positive sentiment and 0 means negative sentiment. The logistic regression model uses the Sigmoid function to turn the weight variables w and the bias b into probabilistic outputs. It does this by linearly combining the two values.

$$P(y=1|x)=\frac{1}{1+\exp^{\left(-w^{\top}x-b\right)}} \tag{1}$$

where w stands for the model's parameter variables, x stands for the input feature variable, b stands for the bias term, and exp stands for the exponential function. The result of this function is the likelihood that the text has a positive sentiment. If this number is higher than a certain level (typically 0.5), it is likely that the sentiment is positive. Traditional machine learning methods can work on small and medium-sized corpora, but they have trouble modelling contextual information and dealing with complex semantic structures like long text, sentiment transitions, and sentiment metaphors because they depend on features that are made up.

As deep learning gets better, neural network-based methods are slowly taking the place of older models and becoming the most popular way to analyse sentiment. recurrent neural network (RNN) and their derivatives, including long short-term memory network (LSTM), are very good at modelling context when it comes to sequence modelling (Moradi et al., 2021). LSTMs use a gating system to remember contextual information in input sequences for a long time. Their main computational function is to update the states of forgetting gates, input gates, and output gates. For example:

$$h_t=o_t\odot\tanh\left(c_t\right) \tag{2}$$

where h_t is the hidden state at the moment, o_t is the activation value of the output gate, c_t is the state of the memory cell, tanh is the hyperbolic tangent function, and $\odot$ is the operation for multiplying elements together. This way, LSTM can accurately capture sentiment dependencies in long texts, which makes sentiment categorisation more accurate.

Also, convolutional neural network (CNN) are used for emotion classification tasks, especially when the material is short. CNN uses sliding convolution kernels to extract local features from word variable sequences (Huan et al., 2022).They next use maximum pooling operations to choose the patterns that best represent emotions, and finally they make classification decisions. CNN is superior to LSTM in capturing local contextual data and expressing the proximity of sentiment words.

Sentiment analysis has entered a new stage of context-aware language modelling in recent years, thanks to the rise of Transformer architecture and pre-trained language models. For example, bidirectional encoder representations from transformer (BERT) can find the most common patterns of sentiment through large-scale unsupervised pre-training and learning of universal language representations. BERT is better than other deep learning models for texts that are hard to understand or have emotional shifts or semantic transitions. In real-world uses, the output text representation of the BERT model is frequently linked to a simple linear classifier that can figure out the polarity of the sentiment (Barreto et al., 2023).

To sum up, sentiment analysis technology has gone through four stages: dictionary rule approaches, classical machine learning methods, deep learning methods, and pre-trained language models. Each step has consistently enhanced its capabilities in feature expression, context modelling, sentiment modelling accuracy, and more, facilitating the extensive utilisation of sentiment computing technology across many text analysis contexts.

2.2 Social media monitoring system

As internet technology has grown quickly and social media sites have become quite popular, social media monitoring systems have slowly become an important tool for managing public opinion and analysing data. In the early days of social media monitoring, keyword matching and simple statistical approaches were used to keep track of how people were talking about certain themes or phrases. But as the amount of user-generated content (UGC) has grown and the sorts of data have become more varied, traditional approaches have had a lot of trouble dealing with vast amounts of unstructured data (Maia et al., 2025).

To solve this problem, researchers and engineers have started to use more advanced algorithms and methods. NLP has been widely used to analyse text on social media, including text preprocessing, word segmentation, named entity recognition, dependent syntax analysis, sentiment tendency determination, topic modelling, and more. For instance, Latent Dirichlet allocation (LDA) topic modelling can find hidden semantic structures in large texts without any help. It does this by assuming that documents are made up of a mix of different topics (Chauhan and Shah, 2021).This helps to look at how user interests are spread out and how public topics change on social media.

Sentiment analysis, a fundamental component of opinion monitoring, is continually advancing through various methodologies. It began with a dictionary-based approach and has slowly moved on to machine learning techniques and now the most popular deep learning models. Dictionary methods use SentiWordNet to count the weights of positive and negative sentiments. This is good because it is easy to understand and quick to compute. Machine learning methods, like SVM, can be trained on hand-made text features to make sentiment classification more flexible. SVM works well when there are clear sample boundaries and high feature dimensions, which is especially good for binary classification problems.

As deep learning has progressed, more intricate neural network architectures are extensively employed in text sentiment analysis: CNN uses convolutional kernels to get local contextual features, which works well for short texts like tweets or comments (Chen et al., 2020). RNN works well for time-series data and can model word-order information, but it has the problem of gradient vanishing. To fix this, LSTM adds a memory gating mechanism based on LSTM, which greatly improves the ability to model long-term dependencies. the ability to model. The attention mechanism model makes the system even better at picking up important information, especially when it comes to multi-round opinion tracking and sentiment mutation analysis.

Graph neural network (GNN) is getting more and more attention at the algorithmic level, along with standard supervised learning. This is critical for finding major opinion nodes and guessing how far information will spread. By modelling users or postings as nodes in a graph and using linking relationships to show how people interact with each other (retweets, comments, likes, etc.), GNN can better describe how information spreads and how users affect each other. GraphSAGE and GAT are two related versions that have also been slowly added to the system to make it better at generalising under complicated network configurations (Chang and Branco, 2024).

Along with text, social media has a lot of unstructured content including pictures, videos, and audio files. Because of this, multimodal learning methods have been a popular area of research. Multimodal models allow surveillance systems to better understand the information they collect by understanding how diverse data sources could be related to each other. For instance, the joint image recognition model (ResNet) and text encoding model (BERT) build a cross-modal emotion classifier that makes it easier to recognise emotions in mixed graphic and text content.

For large-scale data processing, the social media monitoring system uses big data platforms to do efficient calculations. Frameworks like Hadoop and Spark support offline batch processing and feature engineering. Apache Kafka and Flink are popular for collecting streaming data and processing it in real-time, which gives the system near-real-time feedback capabilities.

In recent years, reinforcement learning (RL) has been progressively integrated into public opinion management contexts, particularly in the optimisation of public opinion guidance tactics and the recommendation of content, thereby demonstrating its potential (Den Hengst et al., 2020). RL learns the best ways to act by having a smart body interact with the environment. To help people with their opinions, the reward function can be used to specify the aim of the guidance (to lessen the spread of bad feelings), and the strategy network may be used to continuously improve the ability to intervene with information in a smarter and more flexible way.

To sum up, the history of the social media monitoring system shows how technology has changed from static keyword matching to dynamic deep modelling. The main technologies it uses include NLP, machine learning, deep neural networks, GNN, multimodal fusion, large data processing, and RL. These are all cutting-edge fields. As new technologies like pre-trained large models (e.g., the GPT and BERT families) and federated learning get better over time, the social media monitoring system will keep getting better, becoming smarter, safer, and more useful. It will also be even more important in situations involving cybersecurity, public governance, and content regulation.

3 System design and model construction

When dealing with complicated and changing information about public opinion on social media, it has become important to build a system that can accurately analyse sentiment and recognise content related to ideology and politics. This is a key technical path to achieving value guidance and content regulation in cyberspace. To facilitate effective identification and organised analysis of emotional expressions connected to ideology and politics in social media texts, this research develops a social media monitoring system grounded in optimised sentiment analysis. The system has a layered design, which makes sure that each processing link, from getting data to recognising semantics, is very scalable and stable (see Figure 1). The system has four main layers: the data gathering layer, the pre-processing and feature construction layer, the sentiment analysis and optimisation modelling layer, and the semantic recognition and correlation analysis layer for ideology and politics. Each layer opens up step by step based on the logical dependency relationship. This creates a closed-loop process that goes from capturing information to understanding feelings to making value judgments. This greatly improves the processing speed and accuracy of the sentiment content related to civic and political education in online public opinion.

Figure 1 System architecture of social media monitoring (see online version for colours)

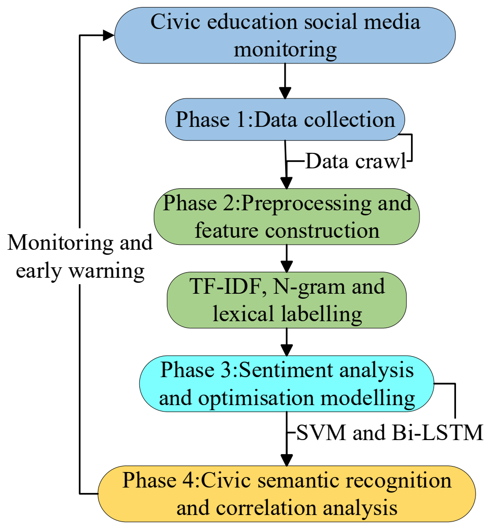

1 Data collection layer. The data collection layer is the most important part of the social media monitoring system based on optimised sentiment analysis. It is the first step in the whole process of processing information, and its quality has a direct impact on how well and quickly the next steps of pre-processing, feature construction, and modelling analysis work. As social media platforms grow quickly, UGC has a lot of different types, and it changes often (Zhuang et al., 2025).To effectively gather public opinion data on ideological and political education topics, the data collection layer must possess robust data capture capabilities, mechanisms for topic focus, and strategies for exception control to guarantee the diversity of data sources and semantic relevance.

The system mainly focuses on popular Chinese social media sites including Sina Weibo, WeChat, Zhihu, Douban, and others. It uses hybrid collecting architecture that includes a platform open interface (API) call, directed crawler scheduling, and keyword filtering (Kim et al., 2020). The collecting mechanism is based on topics and uses a time sliding window and content semantic features to crawl more material and update it in real-time. The first collection of social media corpuses should be:

$$D = \{d_1, d_2, \ldots, d_n\} \quad 3)$$

where d_i denotes the text data of article i, which includes metadata such as body content c_i, release time t_i, user identity u_i, number of comments r_i, and so on. To make it easier to process later, all the data is saved in a single data structure. All the data is stored in a single data structure, which makes it easy to process later. To make sure that the collected content stays on topic, the system creates a list of domain keywords:

$$K = \{k_1, k_2, \ldots, k_m\} \quad 4)$$

where k_j is the j^{th} term that is closely related to the meaning of civics, like young responsibility, moral norms, and so on. The selection of keywords can be changed on the fly based on social trends or research focuses. The system uses a Boolean matching function to filter each piece of input semantically:

$$f(d_i, K) = \begin{cases} 1 & if\ \exists\, k_j \in K\ and\ k_j \in d_i \\ 0 & otherwise \end{cases} \quad 5)$$

This function checks to see if the text d_i has any keywords from the keyword collection. If it does, it means that the semantics of the text are important to the study goals. This step is like a first step in filtering the dataset D to get rid of noise data that isn't useful.

The system also builds the corpus that meets the requirements for the desired text collection:

$$D' = \left\{ d_i \in D \middle| f\left(d_i, K\right) = 1 \right\} \quad (6)$$

where D' is the set of all the social texts that have the goal keywords. This is the main input data for the system to start the text pre-processing and sentiment modelling step. This collection is much smaller than the original collection D. It also makes the next computation more efficient and makes the corpus more thematically consistent at the same time. This method not only uses fewer system resources, but it also makes the model better at telling the difference between domain semantics.

To make the data collection layer more flexible and able to work in real-time, the system creates a distributed crawler scheduling mechanism to support platform-level concurrent crawling tasks. This, along with the sliding time window strategy $T = [t_s, t_e]$, lets the system dynamically get content from a certain time. For instance, during important holidays, big social events, or the publication of new policies, the system can focus on gathering data during times when sentiment changes quickly by changing the range of T to make it easier to do evolutionary analysis later.

The system not only matches keywords, but it also uses user behaviour screening logic to check for quality by looking at content length, interaction number (e.g., l_i, r_i), user activity, and other factors. It then builds the following weighting function:

$$w\left(d_i\right) = \alpha \cdot \log\left(1 + l_i\right) + \beta \cdot \log\left(1 + r_i\right) \quad (7)$$

where α and β are two of these control characteristics that help balance the impact of likes and comments. The function is used to figure out how much heat or conversation the material is getting on the platform, which is what the next step in content prioritisation is based on. The data will only be sent to the model training module if $w(d_i)$ is greater than the threshold θ specified by the system. This weighted filtering process considers the emotional relevance and social importance of the text content, which makes the system more useful and responsive.

Additionally, the data collection layer includes several built-in tools to help control the quality of the data. For example, the data cleaning module automatically removes HTML tags, special characters, ad content, and duplicate records. The identification of unusual user behaviour uses behavioural characteristics to find machine screen-swiping, users who are too active, and water armies, which helps the system make fewer mistakes (Rinott and Tractinsky, 2022).Finally, the multi-source redundancy comparison mechanism compares data from multiple platforms to Third, the multi-source redundancy comparison mechanism takes samples from more than one platform to make sure that the coverage and completeness are better. Setting the crawling frequency, keyword update cycle, and platform API access quota in a reasonable way makes it possible for the system to work properly, with stable, manageable data input.

To sum up, the layer's job is to get semantically significant language from the system's unstructured social data. This layer builds a high-quality, sustainably updated target corpus pool by adding keyword-based topic matching, platform heterogeneous fusion, time sliding window scheduling, and anomaly detection. This pool provides a solid data support base for the next steps in sentiment modelling and ideology analysis.

2 Pre-processing and feature construction layer. Once the text cleaning and normalisation processes are finished, the system builds multiple levels of features on the text data. This is done to meet the needs of different models for the input format and to make sentiment analysis more accurate and expressive. The feature construction technique encompasses conventional lexical level representation as well as statistical features, syntactic features, sentiment features, and multi-layer embedding representation, collectively forming the multimodal variable structure of the system input.

At the lexical level, the system uses TF-IDF, N-gram, and lexical labelling as its main features. Also, part-of-speech (POS) features are stored as sparse variable inputs that show how different types of words are distributed in the text (Maqsood et al., 2020). For example, adjective density, verb frequency, and modal proportion are all examples of this. These features have been shown to be useful in distinguishing between strong and weak emotional expressions.

Second, in terms of statistical aspects, the system pulls out indicators that are strongly related to the structure of the text and how the user interacts with it. These include the length of the text, the number of sentences, the average sentence length, the density of punctuation, the fraction of exclamations, and the amount of capital letters. For instance, a lot of exclamation marks and subjective terms could mean that the person is very emotional, while the length of the text and the depth of the remarks could mean that the person is willing to contribute (Wankhade et al., 2022). Often, these hand-built statistical features are used with depth features splicing to make the model work better when there aren't many samples.

The system uses a lot of different embedding representations to model the text with low-dimensional densities at the semantic level. Word variables are initially trained via Word2Vec, which associates each lexical item with a variable of a constant dimension in a continuous space (Yilmaz and Toklu, 2020). Subsequently, sentence representations are created through averaging or weighted aggregation:

$$v_d = \frac{1}{n}\sum_{i=1}^{n} w_i \tag{8}$$

where v_d stands for the overall variable representation of text d, w_i stands for the word variable of the i^{th} word, and n stands for the number of words in the text. This formula is based on the idea of average pooling of word variables. This is widely used as a lightweight way to represent text in the model input layer for shallow models.

Then, GloVe is used to pre-train the word variables to find global co-occurrence features in a wider range of situations. In higher-level semantic modelling, BERT is used to encode the whole text in a way that takes context into account, and this is used as a sentence semantic representation through CLS variables (Ding et al., 2023). BERT can model context in both directions, which makes it better at expressing complicated semantic structures like puns, negation, and irony than Word2Vec and GloVe. In real-world use, the system allows both static word variables and dynamic context variables as inputs, allowing both shallow and deep semantics to affect the model at the same time.

Furthermore, considering the specific characteristics of the sentiment analysis task, the system creates a distinct set of sentiment vocabulary-related features. It utilises the KnowledgeNet Hownet dictionary and the sentiment polarity dictionary to annotate the quantity and distribution ratio of positive, neutral, and negative vocabulary within each text, and it computes the sentiment score of the sentence by integrating the sentiment intensity words with the negative word collocation position. For instance, by creating emotion polarity feature variables:

$$f_s(d_i) = [pos_i, neg_i, neu_i, neg\_\text{mod}_i] \tag{9}$$

where pos_i, neg_i, and neu_i stand for the number of positive, negative, and neutral words, respectively, neg_mod_i stands for the number of negative modifiers. 'Co-occurring with emotion terms. These features in the dataset work well with the variational semantic model's capacity to make precise emotion evaluations.

The method also adds a hierarchical feature extraction mechanism for lengthy texts with complicated chapter structures. The long text is divided into sentence-level parts, and word variables and sentiment features are taken out of each segment separately (Rao et al., 2018). Then, global feature pooling is done at the paragraph level, and the Attention mechanism looks at the key sentence weights to make sentence-level variable fusion representations. This keeps the chapter-level logical structure while avoiding losing information.

In summary, this system uses a multi-dimensional, multi-granularity, multi-semantic feature extraction strategy at the feature construction layer. This accurately models text data from lexical, structural, semantic, and sentiment levels. It greatly improves the model's ability to tell the difference between and generalise complex expressions, and it gives the downstream sentiment classification task enough semantic foundation and information support.

3 Sentiment analysis and optimisation modelling layer. The layer is the main computing part of this system. Its main job is to figure out how people feel about text. This layer employs supervised learning methods to build a sentiment classification model after feature extraction in the previous layer. It also includes an optimisation process to make the model more generalisable and useful in real life. SVM is the baseline model for traditional models. It works well with high-dimensional sparse text information and is good for binary or multi-classification tasks like emotion polarity classification (Rahman-Laskar et al., 2024).The main purpose is to build the best hyperplane in the feature space to make the classification interval as wide as possible. The logic behind it is as follows:

$$\min_{w,b,\xi} J(w,\xi) = \min_{w,b,\xi}\left(\frac{1}{2}\|w\|^2 + C\sum_{i=1}^{n}\xi_i\right) \tag{10}$$

where w is the weight variable, b is the bias term, ξ_i is the slack variable, and C is the penalty parameter that balances the model's complexity with the cost for making mistakes in classification.

The system also adds a deep learning model based on bidirectional long and short-term memory networks (Bi-LSTM) to get the contextual semantic information from the text. Bi-LSTM can look at both forward and backward information in a sequence at the same time (Jian et al., 2024). This eliminates the problem of typical RNNs' gradients disappearing or inflating in long sequences. The following steps are used to calculate its forward and backward hidden states:

$$\overrightarrow{h_t} = LSTM\left(x_t, \overrightarrow{h_{t-1}}\right) \tag{11}$$

$$\overleftarrow{h_t} = LSTM\left(x_t, \overleftarrow{h_{t+1}}\right) \tag{12}$$

The hidden state at the end is a combination of the hidden state in both directions:

$$L = -\sum_{i=1}^{n}\sum_{c=1}^{C} y_{i,c}\log\left(\hat{y}_{i,c}\right) \tag{13}$$

where $y_{i,c}$ is the real label for sample i in category c and $\hat{y}_{i,c}$ is the model's prediction probability. To make the model better at generalising, we use the dropout approach and L2 regularisation during training (Wu et al., 2021).To fix the problem of social media text sentiment categories not being evenly distributed, the loss function is changed by adding category weighting. This makes it easier to recognise a few categories more accurately.

The system also includes an attention function to help people focus on important emotional phrases. The model can change how much each word contributes to the final sentiment judgment by calculating the attention weights at each time step. This makes sentiment analysis more accurate and easier to understand.

In short, the layer combines the strength of traditional SVM with the strong semantic expressive ability of the Bi-LSTM deep model to accurately and efficiently classify sentiments. This is the basis for the social media opinion monitoring system.

4 Ideological and political semantic identification and association analysis layer. The main job in this layer is to correctly recognise and sort the text's ideological and political semantic properties. To reach this goal, the system uses deep neural networks to get the high-dimensional semantic features of the text. It then gets a raw score variable z, where each element z_j stands for the unnormalised score of the text in the j^{th} category of ideological and political topics. It is hard to use these scores directly for later probabilistic determination and decision-making since they are not probability distributions and have a range that changes. Therefore, they need to be normalised.

The softmax function is often used for multi-category classification tasks (Urazoe et al., 2021).It may turn the original score variables into probability distributions using the following formula:

$$p_j = \frac{\exp(z_j)}{\sum_{k=1}^{C} \exp(z_k)} \tag{14}$$

where p_j is the chance that the text is seen as belonging to the j^{th} category of ideological semantics, and C is the total number of all categories. The function uses the exponential operation to make the difference between larger scores bigger, and it also normalises the scores of all categories so that the output value is between 0 and 1 and adds up to 1. This fits the basic requirements of probability distribution.

From the point of view of computational efficiency, the softmax function scales well. The system can process vast amounts of social media data in real-time thanks to GPU parallel computing. This fits the rigorous time constraints of modern surveillance systems. Simultaneously, the model training phase progressively refines the parameters by maximising the softmax probability of the proper category, in conjunction with the gradient descent method, to enhance the accuracy and generalisation capability of semantic recognition in civics.

To sum up, the softmax function is the most important math tool for the semantic recognition of civics. It not only maps text multi-category probability effectively but also gives a strong probabilistic basis for the next steps in correlation analysis and decision support. The technique significantly enhances the system's capacity to comprehend intricate ideological and political semantic expressions, facilitating precise monitoring and scientific management of ideological and political public opinion (Németh, 2023).

Algorithm 1, the system's overall pseudo-code, shows the whole process from data collection to feature construction to sentiment analysis, ideological and political semantic identification, and correlation analysis. This makes it easier to see how the different layers of the system work together.

Algorithm 1 Pseudo-code for the social media sentiment and semantic monitoring system

Input: Social media platforms, pre-trained sentiment model, pre-trained semantic model, monitoring time range

Output: Sentiment classification results, Semantic classification results, Association analysis results

```
begin
Set up data source connections
Initialise storage for collected text
for each platform in the source list do
    Download text data within the time range
    Remove noise and filter invalid entries
    Store cleaned text into local storage
end for
for each text entry do
    Perform word segmentation
    Generate word embeddings and sentiment features
    Merge features into a unified representation
end for
for each representation do
    Use sentiment model to predict sentiment label
    Store sentiment result
end for
Optimise sentiment results using post-processing
for each representation do
    Use semantic model to predict ideological label
    Store semantic result
end for
Perform association analysis between sentiment and semantic labels
Output final results
end
```

4 Experimental design and analysis of results

4.1 Experimental data

This research creates a Chinese social media semantic sentiment fusion dataset to test the proposed social media monitoring system's ability to improve civic and political education through optimised sentiment analysis. There is no publicly accessible Chinese dataset that includes both sentiment tendencies and semantic labels for civic and political education. Therefore, a self-constructed methodology is employed to guarantee that the data thoroughly encompasses the content of topics pertinent to civic and political education, fulfilling the dual experimental requirements of sentiment analysis and semantic recognition.

The data collection primarily focuses on Chinese social media sites, including microblogs, WeChat public account comment sections, and prominent news comment platforms, with the specified timeframe extending from January 2023 to December 2024. We used keyword searches to find themes linked to ideological and political education, such as patriotism, civic duty, social justice, historical education, and ideological expressions. After getting roughly 28,000 text samples, we used semantic denoising and quality screening to get rid of invalid text, spam, and comments that were too similar. In the end, we kept 12,000 genuine text samples.

The dataset has a mostly balanced distribution of emotional inclination, and the five categories of semantic labels for civics and politics make up an average amount of the data. This makes it easy to keep the categories balanced when training the model. Also, the length of the text is kept to a reasonable level so that very short statements and long quoted war comments don't get in the way. Table 1 shows the main details about the dataset:

Table 1 Basic information of the custom social media dataset for ideological and emotional analysis

Item	*Description*
Data source	Weibo, WeChat public comments, news media platforms (2023.01 – 2024.12)
Crawling method	Python crawler (requests + BeautifulSoup) with keyword filtering
Example keywords	Patriotism, ideology, education policy, core socialist values
Raw data volume	Approx. 28,000 text entries
Valid samples	12,000 entries after filtering
Label types	Sentiment labels (3 classes)
	Ideological semantics (5 categories)
Semantic categories	Patriotism, civic responsibility, social justice, historical view, ideology
Average text length	42 Chinese characters (after cleaning stopwords and noise)

The dataset offers robust data support for further model training and evaluation, encompassing both conventional social media sentiment expressions and comprehensively representing the range and complexity of ideological semantic communication. It has high expandability because it can be used with a wide range of text sources and a reasonable annotation structure. It can also be used for more complicated public opinion analysis jobs in the future, like multi-task learning.

4.2 Experimental setup and evaluation

To validate the functionalities and overall efficacy of the system at each layer, this study conducts systematic tests using a self-constructed social media sentiment dataset to assess the performance of the developed model against a genuine corpus. The experiments are conducted in a consistently configured hardware and software environment, with the specific parameter combinations detailed in Table 2:

Table 2 Experimental environment and parameter settings

Item	*Description*
Hardware platform	NVIDIA RTX 3090 GPU × 1, Intel i9-13900KF CPU
Operating system	Ubuntu 22.04 LTS
Programming language and tools	Python 3.10, PyTorch 2.1, scikit-learn, pandas
Maximum epochs	30
Batch size	64
Optimiser	Adam
Activation function	ReLU (main model), Softmax (output layer)
Random seed	42 (for reproducibility)

This study presents metrics for model evaluation, aiming to provide a thorough assessment of the system's performance across multiple dimensions.

To begin with, AUC is a key measure of how well the model classifies data. It does this by computing the area under the ROC curve, which shows how well the model can tell the difference between positive and negative samples. The AUC value ranges from 0 to 1, with values closer to 1 indicating a stronger discriminative ability of the model (Carrington et al., 2022). This is particularly advantageous for emotion classification involving imbalanced categories, thereby enhancing the evaluation of the model's effectiveness in recognising the emotional tendencies of complex texts.

Second, this study introduces Cohen's Kappa coefficient to measure how consistent the model's predictions are with the real labels. This metric fixes the bias that comes from stochastic consistency and changes the typical range of values from [–1, 1] to [0, 1] through normalisation. This makes it easier to compare the results of different jobs. A higher Kappa value means that the model's prediction is more consistent and trustworthy (Chicco et al., 2021). This is especially true for social media material with a lot of tags that are clearly polarised, which is a good way to show how well the model really works.

Third, this research uses normalised mean absolute error (NMAE) to figure out how accurate and stable the model is when it gives a multi-category probability output. Normalising the standard means absolute error to keep the value between 0 and 1 makes this measure easier to understand and more useful in other situations.

Lastly, the normalised information gain ratio (NIGR) is a key measure for choosing and testing features. This metric effectively stops the model from being biased toward high-base features by using the normalised information gain. It also appropriately counts how much each feature reduces uncertainty about the classification target. The addition of NIGR not only makes feature construction more scientific but also makes the model better at interpreting data (Geirhos et al., 2020).This gives a strong theoretical basis for the next steps in sentiment analysis and association mining.

These evaluation indexes not only look at how well the model classifies things, but they also look at how stable, consistent, and able to understand the meaning of the model, which is more in line with what is needed for social media text analysis. We can prevent the cosmetic improvement that comes from relying too much on accuracy rate by using a more differentiated index system. This will better show how well the system works in real-world situations.

4.3 Experimental procedure

To thoroughly evaluate the efficacy and practical utility of the social media monitoring system for civic and political education, predicated on the enhanced sentiment analysis described in this work, two principal experiments are formulated. The first experiment centres on validating the model's performance, assessing the algorithm's efficacy and improvement by juxtaposing the standard method's performance with that of the optimised model in sentiment classification and civics semantic recognition. The second experiment examines the system's functionality inside an actual social media context, evaluating its capacity to track public opinion fluctuations and risk alerts through real-time data monitoring and analysis. These two tests work together to make sure that the technology is scientific and to show how useful the system is in real life. This gives solid backing for future improvements and marketing.

To thoroughly assess the efficacy of the sentiment analysis model suggested in this research, which integrates SVM and Bi-LSTM, experiment 1 is structured for multi-model comparison. The chosen models which are Naive Bayes, random forest (RF), LSTM, CNN, SVM, and Bi-LSTM, are thoroughly evaluated in terms of their performance, pros and cons, and how well they recognise sentiment in social media writings about civics.

Figure 2 Results of performance experiments (see online version for colours)

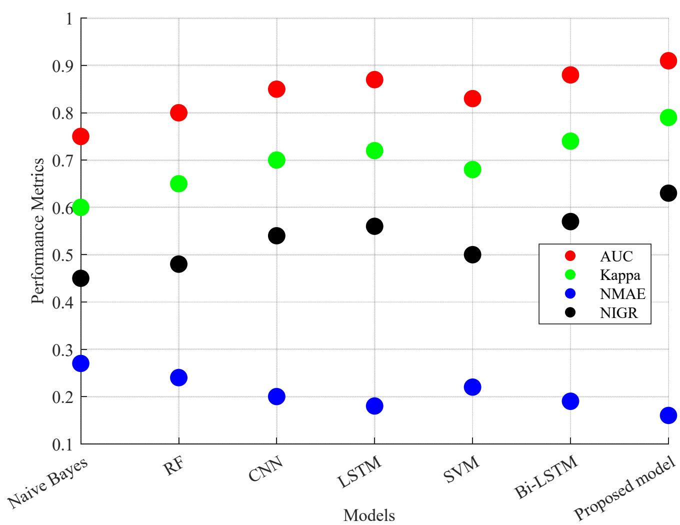

The dataset utilised for the experiment comprises extensive texts centred on civic and political themes. These texts undergo a standardised preparation regimen, which includes word segmentation, word vector embedding, and additional procedures to guarantee the integrity and uniformity of the input data. Naive Bayes and RF are two examples of traditional machine learning models that train based on data like word frequency. They work well but don't do a good job of capturing complex meanings. On the other hand, deep learning models like CNN and unidirectional LSTM use convolution and temporal structure to get local and contextual features of text. This helps them better understand the structure of the text. Bi-LSTM, on the other hand, uses pre- and post-contextual information to improve semantic comprehension even more. The fusion model described in this research integrates the discriminative strength of SVM with the sequential modelling capability of Bi-LSTM, seeking to enhance both accuracy and resilience.

The experiments employ four metrics including normalised AUC, Cohen's Kappa, NMAE, and NIGR to assess model performance. Figure 2 shows the outcomes of the experiments.

This project thoroughly assesses the efficacy of several algorithms in conducting social media sentiment analysis for civic and political education through comprehensive model comparison. Naive Bayes only gets 0.75 on the AUC measure, which is mostly because it assumes that features are conditionally independent, which makes it hard to get the context of the text. RF gets better, with an AUC of 0.80. This means that the model is better at finding non-linear characteristics by combining numerous decision trees. On the other hand, deep learning models like CNN and LSTM got AUCs of 0.85 and 0.87, which are much better than the standard models. This shows that deep networks can effectively predict local phrases and temporal semantic relations.

The classic SVM scores 0.68 on Cohen's Kappa metric, which is better than Naive Bayes and RF but still worse than Bi-LSTM's score of 0.74. The fusion model takes the best parts of both models and gets the highest Kappa value of 0.79. This means that the model is more stable when it comes to label consistency and can better detect text sentiment categories and reduce classification bias. The higher Kappa value of 0.05 shows that the fusion approach is better at handling different expressions and complex semantics.

The fusion model's NMAE error measure lowers to 0.16, which is a lot lower than the Bi-LSTM's 0.19 and the SVM's 0.22. This means that the average absolute error between the model's predicted affective tendency and the true value is roughly 16% to 27% lower. This is a big improvement in accuracy. This decrease in inaccuracy helps the algorithm make accurate predictions when it comes to real opinion monitoring and avoid making mistakes about the sentiment trend because of bias.

The fusion model has the greatest NIGR score of 0.63, which is higher than any other model. It is 10–30% better than the single model's 0.48 (RF) and 0.57 (Bi-LSTM), which shows that the model is better at choosing features and using information. Higher NIGR means that the model is better at selecting out the variables that are most helpful for sentiment categorisation. This makes the model more powerful at distinguishing between different types of messages on social media and helps it grasp more complicated texts.

In conclusion, the model presented in this study, which combines SVM and Bi-LSTM, performs best in the four main metrics of AUC, Kappa, NMAE, and NIGR. This shows how the traditional discriminative model and the deep sequence model work well together. To further validate the dynamic monitoring capability and applicability of the proposed system in real-world contexts, experiment 2 employs dynamic social media data streams to emulate the real-time formation and development of ideology and politics-related public opinion in an authentic setting. This experiment primarily assesses the model's capacity to manage incremental data, response time, and the early detection of unforeseen public opinion events, thereby providing a thorough evaluation of the system's practical application value.

First, the static dataset from experiment 1 is sorted in order by date, and then the data is loaded into the system in groups. Each batch has a particular quantity of text in it, which is meant to mimic how material is released on social media sites. This manner, the system can slowly get and process the steady flow of fresh data, much as in the real world, to keep track of and analyse public opinion on social media in real-time.

After that, the optimised model does online sentiment recognition on each fresh batch of incoming text data, giving the appropriate sentiment classification findings and their confidence levels. The system keeps track of how long it takes to analyse each batch of data in real-time to check the accuracy of sentiment prediction. It also keeps track of the response delay to make sure the system works well in a changing environment. To further evaluate the model's capacity to react to atypical public opinion occurrences, the experiment integrates simulated unexpected event texts into the data stream, thoroughly assessing the system's functionality and stability within a complex and dynamic social context by analysing the model's sensitivity to sentiment variations and its alerting capability.

In this experiment, we will compare five models: SVM, CNN, BERT, Bi-LSTM, and the one we came up with in this study. All models are inferred online on the same data stream, and the sentiment categories and confidence levels are generated in real-time. The assessment measures include normalised AUC, Cohen's Kappa, NMAE, and NIGR. These are used to measure how well the models classify data, how accurate their predictions are, and how well they find anomalies. Figure 3 shows how well each model did in experiment 2 when it came to utility performance:

Figure 3 RESULTS of practical performance experiments (see online version for colours)

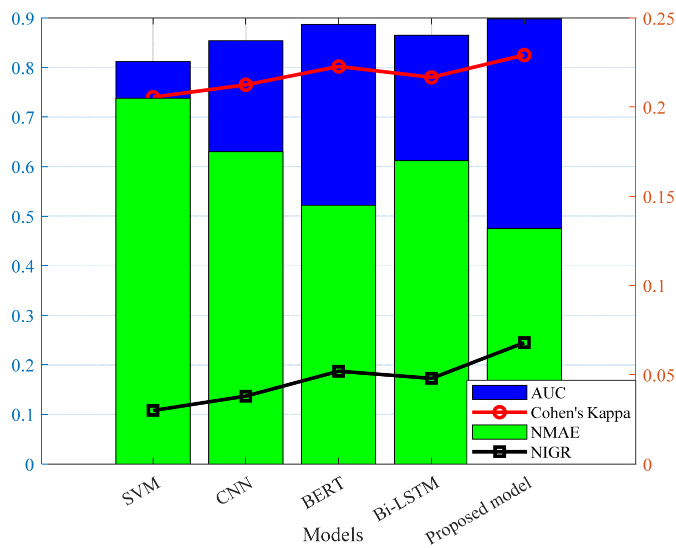

In terms of the AUC index, the model in this work gets the best score of 0.898, which is much superior to SVM (0.812), CNN (0.854), BERT (0.887), and Bi-LSTM (0.865). A higher AUC means that the model can tell the difference between positive and negative samples better. The optimised Bi-LSTM's improvement shows that it can better generalise and discriminate between different emotion categories and emotional features, and it can also more accurately identify complex emotional information in social media.

Second, this paper's model scores 0.825 on Cohen's Kappa metric, which is better than BERT (0.802) and other models. This shows that it is more consistent in its classification results. The greater Kappa value means that the model can stay stable and reliable even when the data is noisy or there are too many categories. This makes sure that the sentiment classification results are trustworthy and useful.

Lastly, the NMAE and NIGR measurements show even more how good this paper's model is. The NMAE of this model is 0.132, which is much lower than the 0.205 of the standard SVM and the 0.175 of the CNN. This means that the model's prediction error for sentiment intensity is smaller, and the prediction results are more accurate. Also, the model in this paper has a high NIGR value of 0.068, which is better than all the other models. This means that it does a great job of capturing and gaining information about unusual emotions, which helps the system find out about public opinion problems and emergencies more quickly. It also improves the early warning capability and application value of emotion monitoring.

In short, this shows that it can handle vast amounts of data rapidly and give correct analytical findings. The model can also swiftly pick up on unusual changes in sentiment and respond well when unexpected events happen in public opinion. This gives public opinion management and social monitoring significant technical assistance.

5 Summary and future research

This paper focuses on a social media monitoring system for civic and political education based on optimised sentiment analysis. It systematically designs a multi-layer architecture that includes data collection, sentiment analysis, optimised modelling, civic and political semantic recognition, and association analysis. It also proposes an optimisation scheme that combines SVM and Bi-LSTM to make sentiment recognition more accurate and faster. We built a special dataset and ran two sets of experiments to prove that the proposed model works better than others for sentiment classification and real-time public opinion monitoring. The results show that the system can better handle the complex and changing social media environment, accurately capture users' emotional changes, and has a lot of practical value and potential for use.

Nonetheless, the research presented in this paper possesses specific limitations. First, the dataset that was made is not very big and mostly contains sentences that are specialised in a certain field. This could make it harder for the model to generalise and adapt. Second, the sentiment analysis model needs to be improved when it comes to handling multimodal data. The current study just looks at text and doesn't leverage information from other types of media to its full potential. The system's response strategy to unexpected public opinion is still in its early stages. It also does not have a fuller early warning system or risk assessment approach, which makes it slower and less accurate in real-world situations when it needs to respond quickly.

Subsequent research may be conducted in several dimensions to enhance the system's efficacy and utility. First, the range of data collection can be broadened to provide larger and more varied training sets with multimodal information, which will improve the model's capacity to generalise (Gao et al., 2020).Second, by integrating with other modern technologies, we can conduct an in-depth examination of the social relationships among users and the pathways of information propagation to enhance the precision of public opinion analysis and risk forecasting. Finally, to improve early warning systems and automated intervention tactics for quick responses to catastrophes, which will help create a more effective and accurate social media monitoring platform that can aid with civic and political education and social governance.

As technology keeps getting better and the need for applications grows, the social media monitoring system based on improved sentiment analysis will become more and more important for managing public opinion, keeping people safe, and teaching. It has a lot of room to grow and be useful.

Acknowledgements

This work is supported by the Hosted the Youth Scientific Research Fund Project at Beijing Wuzi University (No. 2024XJQN28).

Declarations

All authors declare that they have no conflicts of interest.

References

Barreto, S., Moura, R., Carvalho, J., Paes, A. and Plastino, A. (2023) 'Sentiment analysis in tweets: an assessment study from classical to modern word representation models', *Data Mining and Knowledge Discovery*, Vol. 37, No. 1, pp.318–380.

Boler, M. and Davis, E. (2018) 'The affective politics of the 'post-truth' era: feeling rules and networked subjectivity', *Emotion, Space and Society*, Vol. 27, pp.75–85.

Bucci, S., Schwannauer, M. and Berry, N. (2019) 'The digital revolution and its impact on mental health care', *Psychology and Psychotherapy: Theory, Research and Practice*, Vol. 92, No. 2, pp.277–297.

Carrington, A.M., Manuel, D.G., Fieguth, P.W., Ramsay, T., Osmani, V., Wernly, B., Bennett, C., Hawken, S., Magwood, O. and Sheikh, Y. (2022) 'Deep ROC analysis and AUC as balanced average accuracy, for improved classifier selection, audit and explanation', *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 45, No. 1, pp.329–341.

Chang, L. and Branco, P. (2024) 'Embedding residuals in graph-based solutions: the E-ResSAGE and E-ResGAT algorithms. A case study in intrusion detection', *Applied Intelligence*, Vol. 54, No. 8, pp.6025–6040.

Chauhan, U. and Shah, A. (2021) 'Topic modeling using latent Dirichlet allocation: a survey', *ACM Computing Surveys* (*CSUR*), Vol. 54, No. 7, pp.1–35.

Chen, J., Yan, S. and Wong, K-C. (2020) 'Verbal aggression detection on Twitter comments: convolutional neural network for short-text sentiment analysis', *Neural Computing and Applications*, Vol. 32, pp.10809–10818.

Chicco, D., Warrens, M.J. and Jurman, G. (2021) 'The Matthews correlation coefficient (MCC) is more informative than Cohen's Kappa and Brier score in binary classification assessment', *IEEE Access*, Vol. 9, pp.78368–78381.

DellaPosta, D. (2020) 'Pluralistic collapse: The 'oil spill' model of mass opinion polarization', *American Sociological Review*, Vol. 85, No. 3, pp.507–536.

Den Hengst, F., Grua, E.M., el Hassouni, A. and Hoogendoorn, M. (2020) 'Reinforcement learning for personalization: a systematic literature review', *Data Science*, Vol. 3, No. 2, pp.107–147.

Ding, Y., Guo, Y-H., Lu, W., Li, H-X., Zhang, M.-H., Li, H., Pan, A-Q. and Du, X-Y. (2023) 'Context-aware semantic type identification for relational attributes', *Journal of Computer Science and Technology*, Vol. 38, No. 4, pp.927–946.

Gao, J., Li, P., Chen, Z. and Zhang, J. (2020) 'A survey on deep learning for multimodal data fusion', *Neural Computation*, Vol. 32, No. 5, pp.829–864.

Garg, K. and Lobiyal, D. (2020) 'Hindi EmotionNet: A scalable emotion lexicon for sentiment classification of Hindi text', *ACM Transactions on Asian and Low-Resource Language Information Processing* (*TALLIP*), Vol. 19, No. 4, pp.1–35.

Geirhos, R., Jacobsen, J-H., Michaelis, C., Zemel, R., Brendel, W., Bethge, M. and Wichmann, F.A. (2020) 'Shortcut learning in deep neural networks', *Nature Machine Intelligence*, Vol. 2, No. 11, pp.665–673.

Hemmatian, F. and Sohrabi, M.K. (2019) 'A survey on classification techniques for opinion mining and sentiment analysis', *Artificial Intelligence Review*, Vol. 52, No. 3, pp.1495–1545.

Huan, H., Guo, Z., Cai, T. and He, Z. (2022) 'A text classification method based on a convolutional and bidirectional long short-term memory model', *Connection Science*, Vol. 34, No. 1, pp.2108–2124.

Jian, W., Li, J.P., Akbar, M.A., Haq, A.U., Khan, S., Alotaibi, R.M. and Alajlan, S.A. (2024) 'SA-Bi-LSTM: self attention with bi-directional LSTM-based intelligent model for accurate fake news detection to ensured information integrity on social media platforms', *IEEE Access*, Vol. 12, pp.48436–48452.

Kim, Y-Y., Kim, Y-K., Kim, D-S. and Kim, M-H. (2020) 'Implementation of hybrid P2P networking distributed web crawler using AWS for smart work news big data', *Peer-to-peer Networking and Applications*, Vol. 13, pp.659–670.

Maia, S., Teixeira Domingues, J.P., Rocha Varela, M.L.R. and Fonseca, L.M. (2025) 'Exploring the user-generated content data to improve quality management', *The TQM Journal*, Vol. 37, No. 3, pp.877–901.

Maqsood, S., Shahid, A., Nazar, F., Asif, M., Ahmad, M. and Mazzara, M. (2020) 'C-POS: a context-aware adaptive part-of-speech language learning framework', IEEE Access, Vol. 8, pp. 30720-30733.

Moradi, M., Sadrossadat, S.A. and Derhami, V. (2021) 'Long short-term memory neural networks for modeling nonlinear electronic components', IEEE Transactions on Components, Packaging and Manufacturing Technology, Vol. 11, No. 5, pp. 840-847.

Németh, R. (2023) 'A scoping review on the use of natural language processing in research on political polarization: trends and research prospects', *Journal of Computational Social Science*, Vol. 6, No. 1, pp.289–313.

Rahman-Laskar, S., Gupta, G., Badhani, R. and Pinto-Avendaño, D.E. (2024) 'Cyberbullying detection in a multi-classification codemixed dataset', *Computación y Sistemas*, Vol. 28, No. 3, pp.1091–1113.

Rao, G., Huang, W., Feng, Z. and Cong, Q. (2018) 'LSTM with sentence representations for document-level sentiment classification', *Neurocomputing*, Vol. 308, pp.49–57.

Rinott, M. and Tractinsky, N. (2022) 'Designing for interpersonal motor synchronization', Human–Computer Interaction, Vol. 37, No. 1, pp. 69-116.

Sudarmo, S., Arifin, A., Pattiasina, P.J., Wirawan, V. and Aslan, A. (2021) 'The future of instruction media in Indonesian education: systematic review', *Al-Ishlah: Jurnal Pendidikan*, Vol. 13, No. 2, pp.1302–1311.

Urazoe, K., Kuroki, N., Kato, Y., Ohtani, S., Hirose, T. and Numa, M. (2021) 'Multi-category image super-resolution with convolutional neural network and multi-task learning', *IEICE TRANSACTIONS on Information and Systems*, Vol. 104, No. 1, pp.183–193.

Wankhade, M., Rao, A.C.S. and Kulkarni, C. (2022) 'A survey on sentiment analysis methods, applications, and challenges', *Artificial Intelligence Review*, Vol. 55, No. 7, pp.5731–5780.

Wu, L., Li, J., Wang, Y., Meng, Q., Qin, T., Chen, W., Zhang, M. and Liu, T-Y. (2021) 'R-drop: regularized dropout for neural networks', *Advances in Neural Information Processing Systems*, Vol. 34, pp.10890–10905.

Yilmaz, S. and Toklu, S. (2020) 'A deep learning analysis on question classification task using Word2vec representations', *Neural Computing and Applications*, Vol. 32, No. 7, pp.2909–2928.

Zembylas, M. (2022) 'Affective ideology and education policy: implications for critical policy research and practice', *Journal of Education Policy*, Vol. 37, No. 4, pp.511–526.

Zhuang, W., Zeng, Q., Zhang, Y., Lin, D. and Fan, W. (2025) 'What makes UGC more popular on social media platforms? Insights from information adoption theory', *Behaviour and Information Technology*, Vol. 44, No. 8, pp.1521–1538.