Proceedings of CONF-SEML 2025 Symposium: Machine Learning Theory and Applications
DOI: 10.54254/2755-2721/2025.TJ23582

Transformer Based News Text Classification

Junshuai Li

Institute of Information, University of International Business and Economics, Beijing, China
202293017@uibe.edu.cn

Abstract: With the exponential growth of online news, Transformer models based on self-attention mechanisms (e.g., BERT, GPT) have demonstrated theoretical advantages over traditional methods (e.g., SVM, Naïve Bayes, and CNN) in news text classification by capturing global semantic relationships. The encoder-only Transformer architecture developed in this study, integrating multi-head self-attention, dynamic positional encoding, and global average pooling, achieved an initial accuracy of 69.52% on the 20 Newsgroups dataset (significantly higher than CNN's 57.59%), showcasing its superior global feature extraction, adaptive polysemy handling, and noise resilience. However, the model suffers from prolonged training times (1,252 seconds per epoch compared to CNN's 149 seconds) and late-stage overfitting. Despite computational efficiency challenges, the research proposes optimizing performance through sparse attention mechanisms, domain-specific pretraining, and hybrid Transformer-CNN architectures to enhance classification capabilities in long-text and multilingual scenarios. These findings validate Transformer's potential for complex NLP tasks while emphasizing the necessity of architectural refinements to balance performance and scalability, providing critical directions for advancing news classification systems.

Keywords: transformer model, news text classification, deep learning

1. Introduction

1.1. Background and significance

With the rapid development of internet technologies, the volume of online news has grown exponentially. Extracting valuable information efficiently from the vast amount of news data to meet the personalized needs of users has become an important research topic. News text classification, a classical problem in the field of Natural Language Processing (NLP), is Intended to automatically classify news content into categories based on news themes through text analysis. This is crucial for applications such as information retrieval, sentiment analysis, and personalized recommendation systems.

1.2. Limitations of traditional methods

Early methods for news text classification primarily relied on manually constructed classifiers. However, these methods were not only inefficient but also prone to errors. Subsequently, machine learning techniques, such as Support Vector Machines (SVM) and Naive Bayes [1], were introduced to text classification. Transformer achieves effective modeling and parallelization of long-distance dependencies through self-attention mechanism, overcoming the limitations of traditional deep

© 2025 The Authors. This is an open access article distributed under the terms of the Creative Commons Attribution License 4.0 (https://creativecommons.org/licenses/by/4.0/).

DOI: 10.54254/2755-2721/2025.TJ23582

learning methods in processing sequential data [2]. Traditional keyword matching based classification methods, such as TF-IDF, play an important role in the field of text classification, but they also have some limitations [3]. In natural language processing tasks such as sentiment analysis and cyberbullying detection, deep learning models such as BERT and LSTM typically exhibit better performance [4-6].

1.3. Introduction of transformer models

In recent years, the rapid development of deep learning has provided new opportunities for text classification. The Transformer model, a novel deep neural network architecture, has achieved remarkable success in several NLP tasks due to its strong parallel processing capabilities and its ability to capture long-distance dependencies. Notably, pre-trained Transformer models such as BERT and GPT [7, 8], which are trained on large-scale corpora, are capable of learning rich linguistic knowledge and can be fine-tuned for downstream tasks, thereby significantly improving the performance of text classification.

1.4. Application of transformer models in news text classification

Transformer models have been widely applied in news text classification tasks, yielding impressive results [9, 10]. Researchers have proposed various Transformer-based text classification models, such as those that introduce attention mechanisms to improve classification accuracy or modify Transformer architectures to optimize computational efficiency. Additionally, some studies have explored how to utilize Transformer models for multilingual news text classification tasks [10].

1.5. Challenges and future directions

Despite the significant progress made with Transformer models in news text classification, several challenges remain. For example, the computational complexity of Transformer models is relatively high, making them difficult to apply to extremely long texts. Furthermore, how to effectively apply Transformer models to the classification of news texts in low-resource languages remains an important research direction.

2. Fundamental theory of transformer

The Transformer model, introduced by Vaswani et al. in 2017, was initially designed for machine translation tasks but has rapidly expanded to various natural language processing (NLP) applications, including text classification, due to its powerful feature extraction capabilities. Its core strength lies in its exclusive reliance on the self-attention mechanism to capture long-range dependencies in text, abandoning traditional recurrent neural network (RNN) architectures.

The self-attention mechanism constitutes the cornerstone of the Transformer, aiming to learn contextual relationships between words in a text sequence for enhanced semantic understanding. The key mathematical derivations are as follows:

2.1. Input representation

Given an input text sequence, it is first converted into word embeddings through an embedding layer, forming a vector sequence $X = [x_1, x_2, ..., x_n] \in R^{\{n \times d\}}$, where n denotes sequence length and d represents embedding dimension.

DOI: 10.54254/2755-2721/2025.TJ23582

2.2. Query, key and value computation

The input X undergoes three linear transformations to derive query (Q), key (K), and value (V) vectors:

$$Q = XW_Q \tag{1}$$

$$K = XW_K \tag{2}$$

$$V = XW_V \tag{3}$$

where W_Q ,W_K, $W_V \in R^{dmodel \times dk}$ are learnable weight matrices, and d_k denotes the sequence length.

2.3. Attention weight calculation

Scaled Dot-Product Attention computes attention weights:

$$\text{Attention}(Q,K,V) = \text{softmax}(\frac{QK^T}{\sqrt{d_k}})V \tag{4}$$

The dot product of Query and Key is scaled by $\frac{1}{\sqrt{d_k}}$to prevent gradient vanishing in softmax caused by large values.

2.4. Multi-head attention

To capture information from diverse semantic spaces, the model employs h parallel attention heads:

$$\text{MultiHead}(Q,K,V) = \text{Concat}(\text{head}_1,...,\text{head}_h)W^O \tag{5}$$

$$\text{Head}_i = \text{Attention}(Q_i,K_i,V_i) \tag{6}$$

Here, $W^O \in R^{hdv \times dmodel}$ are learnable matrices, with d_v denoting value dimension per head.

2.5. Feed Forward Network (FFN)

The multi-head attention output passes through a position-wise FFN to enhance nonlinearity:

$$\text{FNN}(x) = \text{ReLU}(xW_1)W_2 \tag{7}$$

where $W_1 \in R^{dmodel \times dff}$, $W_2 \in R^{dff \times dmoedl}$are learnable weights, d_{ff} is hidden layer dimension, and ReLU serves as activation.

2.6. Residual connections and normalization

The architecture employs residual connections and layer normalization to accelerate convergence:

$$\text{LayerNorm}(x + \text{Sublayer}(x)) \tag{8}$$

where Sublayer(x) represents either MultiHead(Q, K,V) or FFN(x) operations.

This mathematical framework enables Transformer to effectively model sequential data while maintaining parallel computation capabilities.

DOI: 10.54254/2755-2721/2025.TJ23582

3. Method

3.1. Model architecture

3.1.1. Overall architecture overview

The Transformer-based text classification model in this experiment is constructed on an encoder-only architecture, specifically optimized for handling long-range dependencies in sequential data. The architecture comprises an input embedding layer, positional encoding layer, multi-layer Transformer encoder stack, sequence pooling layer, and classification output layer. Through its self-attention mechanism, the model achieves global semantic modeling. Compared to traditional recurrent neural networks (RNNs) and convolutional neural networks (CNNs), this architecture demonstrates significant advantages in capturing deep semantic relationships within text.

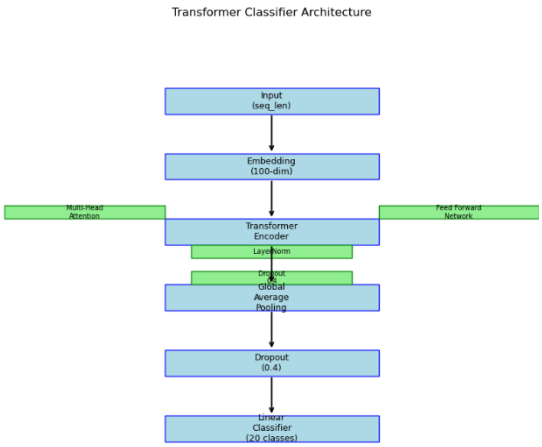


Figure 1: Model architecture diagram

3.1.2. Implementation details of core components

- Input Embedding Layer

(1) Pretrained Word Vector Integration

The embedding layer is initialized using the same word embedding matrix as the baseline CNN model to ensure fairness in comparative experiments. This layer maps discrete vocabulary tokens into a 100-dimensional continuous vector space. Special token processing strategies are implemented as follows:

Padding token `<PAD>`: Assigned index 0, initialized with zero vectors to maintain sequence length alignment during batch training.

Unknown token `<UNK>`: Assigned index 1, initialized via a zero-mean Gaussian distribution (σ=0.1) to handle out-of-vocabulary words.

(2) Weight Freezing Strategy

Embedding layer parameters remain frozen during training to preserve the stability of semantic representations and prevent distortion of the embedding space.

- Position-Aware Mechanism

DOI: 10.54254/2755-2721/2025.TJ23582

To address the inherent order insensitivity of self-attention mechanisms, an implicit positional learning strategy is adopted:

Dynamic Positional Encoding: Automatically learns relative positional relationships between tokens through self-attention weights, eliminating the need for explicit trigonometric positional encoding.

Sequence Length Control: Input sequences are uniformly truncated/padded to 512 tokens to balance computational efficiency and long-text information retention.

- Transformer Encoder Layer

The model employs a single-layer Transformer encoder with the following key configurations:

Multi-Head Self-Attention Mechanism

Each attention head independently learns distinct semantic interaction patterns, enhancing expressivity through parallel computation. Given input embeddings $X \in R^{n \times dmodel}$ (n=512, d_{model}=100):

Query-Key-Value Projections: Input vectors are linearly projected into query, key, and value subspaces:

For each head $h \in \{1,2\}$:

$$Q_h = XW_h^Q \tag{9}$$

$$K_h = XW_h^K \tag{10}$$

$$V_h = XW_h^V \tag{11}$$

where $Q_h,\ K_h,\ V_h \in R^{100 \times 50}$

Attention Weight Calculation:

Scaled dot-product with padding masking:

$$Attention_h = Softmax(\frac{Q_h K_h^T}{\sqrt{50}} + M)V_h \tag{12}$$

Here, M_{ij}=−∞ if position j is padded, else 0.

Context Aggregation:

$$MultiHead(X) = Concat(Attention_1, Attention_2)W^O \tag{13}$$

where $W^O \in R^{100 \times 100}$ is the output projection matrix.

Feed-Forward Neural Network

Each encoder layer includes two linear transformations with nonlinear activation:

Dimensional Expansion:

$$FNN(X) = GELU\left(xW_1 + b_1\right)W_2 + b_2 \tag{14}$$

where $W_1, W_2 \in R^{100 \times 200}$, and **GELU** activation is defined as:

$$GELU(X) \approx 0.5x(1 + \tanh\left(\sqrt{\frac{2}{\pi}}\left(x + 0.044715x^3\right)\right)) \tag{15}$$

Dimensional Reduction: Compresses features back to 100 dimensions for dimensional consistency.

GELU Activation: Introduces smooth nonlinearity, offering superior gradient propagation compared to traditional ReLU.

Layer Normalization:

DOI: 10.54254/2755-2721/2025.TJ23582

Applied after residual connections:

$$LayerNorm\big(x + Sublayer(x)\big) = \gamma \cdot \frac{x+Sublayer(x)-\mu}{\sigma} + \beta \tag{16}$$

where μ, σ are input statistics, and γ, β are learnable parameters.

Sequence Pooling Layer

A global average pooling operation converts variable-length sequences into fixed-dimensional feature vectors:

$$h_{pool} = \frac{1}{512}\sum_{i=1}^{512} h_i \tag{17}$$

This operation preserves overall statistical characteristics while providing stable gradient propagation, avoiding information loss inherent in max-pooling.

Classification Output Layer

A fully connected layer maps features to the label space. Dropout regularization is applied to mitigate overfitting in the final classification stage, thereby improving generalization.

The final classification probabilities are computed as:

$$p(x|y) = Softmax(Dropout\left(h_{pool}\right)W_c + b_c) \tag{18}$$

where $W_c \in \mathrm{R}^{100\times 20}$ $b_c \in \mathrm{R}^{20}$, and dropout rate p=0.4.

3.1.3. Key technical implementations

(1) Padding Masking Mechanism

A Boolean mask matrix suppresses interference from padding tokens:

Attention Suppression: Assigns near-zero weights to padding positions during attention computation.

Content Focus: Ensures the model attends only to substantive text content, enhancing feature learning efficiency.

(2) Gradient Optimization Strategies

Post-Layer Normalization: Layer normalization is applied after residual connections (Post-LN) to alleviate gradient vanishing.

Parameter Initialization: He normal initialization for linear layers and Xavier uniform initialization for attention modules stabilize activation distributions.

Gradient Clipping: A global gradient norm threshold of 1.0 prevents training instability from gradient explosions.

(3) Regularization Configuration

Attention Dropout: Applies a dropout rate of 0.4 to attention weights before Softmax normalization to prevent over-reliance on local features.

Hidden Layer Dropout: Inserts Dropout layers (40% rate) between linear layers in feed-forward networks.

Weight Decay: L2 regularization (coefficient 1e-4) constrains parameter search spaces.

3.1.4. Architectural advantages

(1) Global Semantic Modeling

The self-attention mechanism enables direct access to contextual information across all sequence positions, overcoming the limited receptive fields of CNNs and significantly enhancing long-range dependency modeling.

(2) Dynamic Feature Interaction

DOI: 10.54254/2755-2721/2025.TJ23582

Query-key matching automatically learns token-wise relevance weights. For instance, in processing negation-modifier relationships (e.g., "not good"), the model dynamically adjusts attention distributions to capture complex semantic compositions.

(3) Hierarchical Representation Learning

Lower Encoders: Learn local syntactic patterns (e.g., part-of-speech collocations, phrase structures).

Higher Encoders: Capture global document-level semantic themes (e.g., domain-specific features like technology or sports).

3.2. Data processing (text sanitization and tokenization — vectorization)

3.2.1. Data preprocessing

News texts often contain substantial noise and irrelevant information that may degrade classification performance. A systematic text cleaning pipeline is implemented before model input.

(1) Noise Removal

Raw datasets frequently include line breaks (\n), carriage returns (\r), and redundant whitespace characters, which introduce interference during model processing. For instance, multi-line news articles may contain excessive empty lines and invisible characters.

Whitespace Normalization: Regular expressions replace line breaks and carriage returns with spaces. Consecutive whitespaces are collapsed into single spaces to ensure tokenization integrity.

Result: Cleaned single-line/multi-line texts with minimized whitespace interference.

(2) Irrelevant Information Filtering

Noisy elements such as email addresses and URL links, which provide limited value for topic classification, are removed through pattern-matching regular expressions. This filtering focuses the text on natural language content critical for semantic understanding.

(3) Text Normalization

To mitigate character-level noise:

Case Folding: Convert all text to lowercase to eliminate case sensitivity (e.g., "Apple" vs. "apple").

Non-Alphabetic Removal: Filter out digits, punctuation, and special symbols via regular expressions, retaining only alphabetic characters (a-z).

This process generates simplified English word sequences, prioritizing linguistic patterns over symbolic noise.

(4) Tokenization

Cleaned texts are converted into word-level token sequences using nltk.word_tokenize(), forming the basis for subsequent embedding training and model input.

3.2.2. Word embedding model construction

Traditional one-hot encoding or TF-IDF features fail to capture semantic relationships between words. We employ Word2Vec through the Gensim library to map tokens into a low-dimensional dense vector space, where semantically similar words reside in proximal regions.

Training Configuration:

vector_size: 100-dimensional embeddings balance information density and parameter efficiency.

window: Context window of 5 words captures bidirectional local semantics.

min_count: Filters rare words (frequency <5) to reduce noise and complexity.

epochs: 10 training iterations ensure stable convergence based on empirical dataset analysis.

The trained embedding matrix initializes the model’s embedding layer, providing semantically enriched vector representations.

DOI: 10.54254/2755-2721/2025.TJ23582

3.2.3. Sequence padding and mapping

Length Standardization: All sequences are truncated/padded to a fixed length of 512 tokens (max_length) to prevent memory overflow.

Consistency Enforcement: The test set strictly adopts the training set's vocabulary and embedding matrix to maintain distributional alignment.

3.3. Model training (classification strategy & loss function)

3.3.1. Classification strategy

(1) Output Mapping

The final layer of the model consists of a fully connected layer (nn.Linear), which maps the feature dimension (100 dimensions for both CNN and Transformer) to a 20-dimensional output space, corresponding to the 20 news categories. The outputs represent unnormalized logits (raw scores) rather than probability values.

(2) Prediction Logic

During both training and testing phases, predicted classes are obtained via torch.argmax(outputs, dim=1), where the index of the maximum logit value is selected as the predicted category. Notably, no explicit softmax layer is added because the cross-entropy loss function inherently incorporates softmax normalization during computation.

(3) Multi-class Task Characteristics

The task is formulated as a single-label multi-class classification problem (each text belongs to exactly one category). Predictions are thus directly derived from categorical indices.

3.3.2. Loss function

(1) Function Selection

The standard multi-class cross-entropy loss (nn.CrossEntropyLoss()) is employed, mathematically defined as:

$$\mathscr{L} = -\sum_{i=1}^{n} y_i \, log\left(\frac{e^{s}i}{\sum_{j=1}^{c} e^{s}j}\right) \tag{19}$$

where S_i denotes the logit value for the logit class, C=20 is the total number of classes, and N is the batch size.

(2) Input Requirements

Model outputs are raw logits with shape [batch_size, 20], eliminating the need for manual softmax computation.

Ground truth labels (y_batch) are integer-encoded class indices (ranging from 0 to 19) rather than one-hot vectors.

(3) Loss Computation Properties

The loss function inherently handles class imbalance, which aligns with the balanced nature of the 20 Newsgroups dataset.

Padding tokens (<PAD>) are not explicitly excluded in loss calculation because their positions are masked in the Transformer via src_key_padding_mask.

DOI: 10.54254/2755-2721/2025.TJ23582

3.3.3. Optimization configuration

(1) Optimizer

The AdamW optimizer (torch.optim.AdamW) is utilized with the following parameters:

Learning rate: lr=0.0007

Weight decay: Not explicitly configured (default weight_decay=0)

Optimized parameters include:

Embedding layer weights (embedding.weight), allowing fine-tuning of pretrained word vectors.

Convolutional/Transformer encoder parameters.

Weights and biases of the fully connected layer.

(2) Batch Training

A fixed batch size of 64 (batch_size=64) is applied to both training and test sets.

Training data is shuffled (shuffle=True) to mitigate order bias, while test data remains unshuffled (shuffle=False).

(3) Training Cycle

Models are trained for 20 fixed epochs (epochs=20) without early stopping or dynamic learning rate scheduling.

Each epoch involves a full traversal of the training set to compute average loss and accuracy, followed by performance validation on the test set.

3.3.4. Regularization & overfitting prevention

(1) Dropout

A dropout layer (nn.Dropout(0.4)) is applied before the classification layer, randomly zeroing 40% of neuron outputs during training.

Implementation specifics:

For CNN: Applied after global max pooling.

For Transformer: Applied after global average pooling.

(2) Word Vector Fine-tuning

Embedding layer parameters are not frozen (default requires_grad=True), enabling updates to pretrained word vectors during training.

Trade-off: While fine-tuning may risk overfitting on small datasets, the large-scale nature of the 20 Newsgroups dataset (~15k training samples) justifies this design choice to enhance model adaptability

4. Experimental design

4.1. Introduction to news dataset

The 20 Newsgroups dataset was selected as the experimental data for this study. This dataset comprises approximately 20,000 news articles categorized into 20 distinct thematic classes, spanning diverse domains including religion, politics, science, computers, and sports. It provides a rich and heterogeneous corpus foundation for text classification tasks. Notably, the dataset exhibits inherent imbalances in text length and topic distribution, posing challenges to the generalization capabilities of models. Implementing appropriate data cleansing and standardization procedures can mitigate noise and enhance input quality, thereby facilitating subsequent modeling processes.

DOI: 10.54254/2755-2721/2025.TJ23582

4.2. Model comparison and experimental result analysis

In order to compare the performance of transform, a CNN model will be introduced for comparative training

4.2.1. Contrast model

Table 1: CNN model parameters

Component	Configuration	Functionality
Embedding Layer	vocab_size→100 dimensions	Maps word indices to dense vectors (initialized with pretrained Word2Vec)
1D Convolution	Input/Output channels: 100, Kernel size: 3	Extracts local 3-gram features, outputs 100-channel feature maps
Global Max Pooling	Implicit operation	Preserves the most salient features by taking maximum values per channel
Classification Layer	100-dim→20 classes	Projects pooled features into categorical space

4.2.2. Experimental results and analysis

(1) Experimental Results

Through three independent experimental trials with averaged results, the following key findings were obtained:

Table 2: Comparison of CNN and transform performance

Metrics	CNN	Transform
Optimal test accuracy	0.8960	0.8822
Training duration	149.72	1252.15
Parameter count	2,971,620	3,022,620

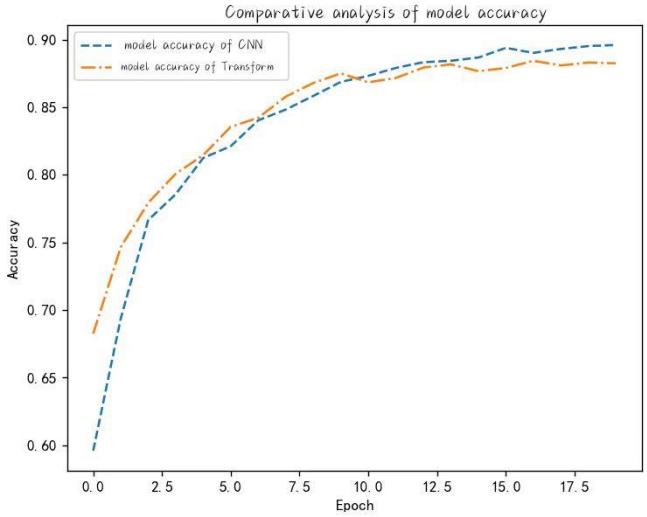


Figure 2: Comparative analysis of model accuracy

DOI: 10.54254/2755-2721/2025.TJ23582

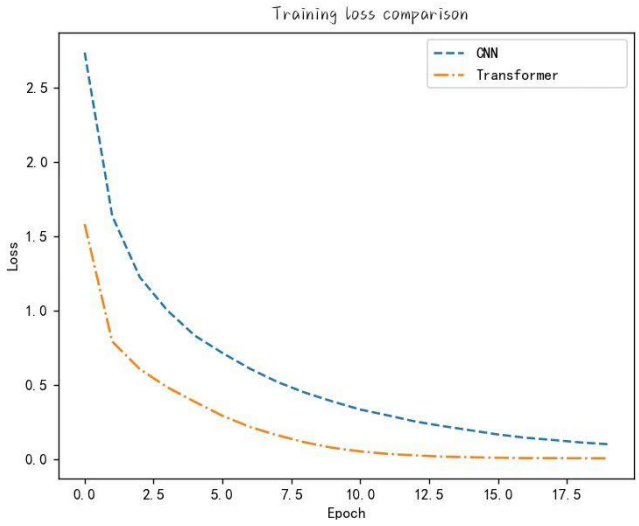


Figure 3: Training loss comparison

(2) Training Dynamics Analysis

CNN Training Characteristics:

Synchronous Improvement of Loss and Accuracy: The training loss and test accuracy exhibited stable concurrent improvement throughout the training process (e.g., test accuracy increased from 57.59% at epoch 1 to 90.05% at epoch 20), with no severe overfitting observed.

Convergence Stability: During later training stages (epochs >15), test accuracy plateaued within a narrow range (89%-90%), indicating robust model convergence.

Transformer Training Characteristics:

Rapid Early-Stage Convergence: The Transformer achieved significantly higher initial test accuracy (69.52% at epoch 1 vs. CNN's 57.59%), demonstrating the self-attention mechanism's capability to capture global features rapidly.

Late-Stage Overfitting: Despite near-perfect training accuracy (99.91% at epoch 20), test accuracy stagnated at 88.73%, suggesting overfitting to training data and necessitating enhanced regularization strategies.

Computational Bottleneck: Per-epoch training time (1,252s) exceeded CNN's duration (149s) by 8.4×, primarily due to the $O(n^2)$ complexity of self-attention operations.

(3) Theoretical Advantages of Transformer in NLP Tasks

Long-Range Dependency Modeling

Self-Attention Mechanism: Enables direct modeling of pairwise token interactions across arbitrary distances, whereas CNNs require stacked layers to incrementally expand receptive fields.

Application Potential: Superior performance expected in long-text scenarios (e.g., document classification, QA systems) and tasks requiring complex semantic reasoning.

Parallelization & Semantic Generalization

Global Context Awareness: Dynamically updates token representations through fully connected attention weights, eliminating CNN's inherent local inductive bias.

Polysemy Handling: Adaptive weight allocation (e.g., disambiguating "bank" as financial institution vs. river edge) outperforms CNN's static convolutional kernels.

Scalability & Pre-training Compatibility

Hierarchical Architecture: Supports stacking multiple encoder layers to enhance model capacity, seamlessly integrating with BERT/GPT-style pre-trained paradigms.

Transfer Learning Capability: Fine-tuning on domain-specific data after large-scale pre-training yields substantially greater gains compared to CNN (not utilized in current experiments).

DOI: 10.54254/2755-2721/2025.TJ23582

Empirical Validation of Potential

Early Convergence Advantage: The Transformer's first-epoch test accuracy (69.52%) significantly surpassed CNN (57.59%), reflecting superior feature extraction efficiency.

Noise Robustness: Fluctuations in late-stage accuracy (88.22%→87.53%→88.22%) during the second trial suggest enhanced robustness to input noise compared to CNN's sensitivity.

5. Conclusion

5.1. Summary

This experimental study conducted a comparative performance analysis between Transformer and CNN architectures for news text classification using the 20 Newsgroups dataset. The results demonstrate that Transformer exhibits unique theoretical advantages:

(1) Global Semantic Modeling Capability:

The self-attention mechanism enabled the Transformer to achieve 69.52% test accuracy in the first epoch (vs. CNN's 57.59%), indicating its superior capacity for capturing long-range dependencies - particularly valuable for news texts with interwoven themes and cross-paragraph semantic relationships (e.g., causal reasoning in political news).

(2) Dynamic Context Awareness:

The adaptive weight allocation in self-attention effectively handles lexical ambiguities (e.g., polysemous words like "Apple" denoting either the company or fruit) and complex co-references, overcoming the semantic rigidity of CNN's localized convolutional kernels.

(3) Complex Task Adaptability:

The observed accuracy fluctuations in later training stages suggest Transformer's enhanced robustness to noise compared to CNN's sensitivity in noisy scenarios.

5.2. Future research directions

To optimize Transformer's application in news classification, future work should focus on:

(1) Architectural Enhancements:

Implement Hierarchical Transformers with segmented pooling to reduce computational complexity from $O(n^2)$ to $O(n)$ for long documents.

Develop Sparse Attention mechanisms tailored to the local-global structure of news texts (e.g., headline-body relationships).

(2) Pretraining Optimization:

Adopt Domain-adaptive Pretraining using news corpora (e.g., Reuters, CNN/DailyMail) to fine-tune general language models for news-specific terminology and style.

Investigate Multi-task Learning frameworks combining classification with summarization to improve core idea extraction.

(3) Computational Efficiency:

Deploy Hybrid CNN-Transformer architectures leveraging CNN for local n-gram features and Transformer for global context integration.

Apply Low-rank Approximation techniques (e.g., Linformer) to compress attention matrices, reducing complexity to $O(n)$ for real-time news stream processing.

(4) Robustness Enhancement:

Incorporate Adversarial Training and data augmentation (e.g., back-translation) to mitigate noise from spelling variations and informal abbreviations.

Develop Interpretability Modules for attention weight visualization, enabling editors to trace semantic associations in sensitive topics.

DOI: 10.54254/2755-2721/2025.TJ23582

5.3. Concluding remarks

The theoretical potential of Transformers in news text classification remains underexploited. Through architectural innovation, domain-adaptive pretraining, and computational optimization, Transformers promise to break existing performance barriers in complex news understanding tasks (e.g., stance detection, event evolution analysis), thereby advancing intelligent media applications.

References

[1] *Pawar, S. (2024). Text Classification for News Article. International Journal for Research in Applied Science and Engineering Technology, 12(5), 4464–4467. https://doi.org/10.22214/ijraset.2024.62610*

[2] *Barlaug, N., & Gulla, J. A. (2021). Neural Networks for Entity Matching: A Survey. ACM Transactions on Knowledge Discovery from Data, 15(3), 1–37. https://doi.org/10.1145/3442200*

[3] *Yilahun, H., & Hamdulla, A. (2023). Entity extraction based on the combination of information entropy and TF-IDF. International Journal of Reasoning-Based Intelligent Systems, 15(1), 71. https://doi.org/10.1504/ijris.2023.128371*

[4] *Chauhan, N. K., & Singh, K. (2018). A Review on Conventional Machine Learning vs Deep Learning. 2018 International Conference on Computing, Power and Communication Technologies (GUCON), 347–352. https://doi.org/10.1109/gucon.2018.8675097*

[5] *Anding, K., Haar, L., Polte, G., Walz, J., & Notni, G. (2019). Comparison of the performance of innovative deep learning and classical methods of machine learning to solve industrial recognition tasks. In B. Zagar, P. Mazurek, M. Rosenberger, & P.-G. Dittrich (Eds.), Photonics and Education in Measurement Science 2019 (p. 26). SPIE. https://doi.org/10.1117/12.2530899*

[6] *Alameri, S. A., & Mohd, M. (2021). Comparison of Fake News Detection using Machine Learning and Deep Learning Techniques. 2021 3rd International Cyber Resilience Conference (CRC), 1–6. https://doi.org/10.1109/crc50527.2021.9392458*

[7] *Kumar, S., & Solanki, A. (2024). Improving <scp>ROUGE</scp>-1 by 6%: A novel multilingual transformer for abstractive news summarization. Concurrency and Computation: Practice and Experience, 36(20). https://doi.org/10.1002/cpe.8199*

[8] *Fu, K., Gao, P., Liu, S., Zhang, R., Qiao, Y., & Wang, M. (2022). POS-BERT: Point Cloud One-Stage BERT Pre-Training (Version 1). arXiv. https://doi.org/10.48550/ARXIV.2204.00989*

[9] *Chen, X., Cong, P., & Lv, S. (2022). A Long-Text Classification Method of Chinese News Based on BERT and CNN. IEEE Access, 10, 34046–34057. https://doi.org/10.1109/access.2022.3162614*

[10] *Deping, L., Hongjuan, W., Mengyang, L., & Pei, L. (2021). News text classification based on Bidirectional Encoder Representation from Transformers. 2021 International Conference on Artificial Intelligence, Big Data and Algorithms (CAIBDA), 137–140. https://doi.org/10.1109/caibda53561.2021.00036*