

KUET at StanceNakba Shared Task: StanceMoE: Mixture-of-Experts Architecture for Stance Detection

Abdullah Al Shafi, Md. Milon Islam, Sk. Imran Hossain, K. M. Azharul Hasan

Department of Computer Science and Engineering, Khulna University of Engineering & Technology
abdullah@iict.kuet.ac.bd, {milonislam, imran, az}@cse.kuet.ac.bd

Abstract

Actor-level stance detection aims to determine an author's expressed position toward specific geopolitical actors mentioned or implicated in a text. Although transformer-based models have achieved relatively good performance in stance classification, they typically rely on unified representations that may not sufficiently capture heterogeneous linguistic signals, such as contrastive discourse structures, framing cues, and salient lexical indicators. This motivates the need for adaptive architectures that explicitly model diverse stance-expressive patterns. In this paper, we propose StanceMoE, a context-enhanced Mixture-of-Experts (MoE) architecture built upon a fine-tuned BERT encoder for actor-level stance detection. Our model integrates six expert modules designed to capture complementary linguistic signals, including global semantic orientation, salient lexical cues, clause-level focus, phrase-level patterns, framing indicators, and contrast-driven discourse shifts. A context-aware gating mechanism dynamically weights expert contributions, enabling adaptive routing based on input characteristics. Experiments are conducted on the StanceNakba 2026 Subtask A dataset, comprising 1,401 annotated English texts where the target actor is implicit in the text. StanceMoE achieves a macro-F1 score of 94.26%, outperforming traditional baselines, and alternative BERT-based variants.

Keywords: stance detection, mixture-of-experts, context-aware gating, adaptive weighting.

1. Introduction

Stance detection refers to the automatic identification of an author's position towards a specific topic, entity, or proposition expressed within a text (Garg and Caragea, 2024). The task aims to identify whether the presented viewpoint is supportive, opposing, or neutral toward the specific target. The study of stance detection differs from common sentiment analysis because it needs to measure emotional expressions that depend on particular targets, whereas the same text can show different stances based on which target the reader chooses to consider (Niu et al., 2024). The nature of this task requires target awareness because it generates both semantic complexity and computational difficulties, particularly when targets remain unmentioned and stances are expressed in an indirect way (Küçük and Can, 2020).

To address these challenges, we introduce StanceMoE¹, a context-enhanced Mixture-of-Experts (MoE) architecture designed for actor-level stance detection. Unlike conventional transformer-based approaches that rely on single aggregated representations, our method explicitly decomposes stance modeling into complementary expert modules that capture diverse linguistic and discourse-level phenomena. By incorporating a context-aware gating mechanism, the architecture enables adaptive and input-sensitive fusion of heteroge-

neous stance signals. Through comprehensive experimentation and detailed ablation analysis on StanceNakba 2026 Shared Task dataset (Aldous et al., 2026), we demonstrate that explicitly modeling such diverse patterns enables more robust and fine-grained actor-level stance detection with 3rd position in the competition.

2. Related Works

The field of stance detection has evolved through its transition from rule-based systems (Küçük and Can, 2020) to classical supervised systems (Alturayef et al., 2023), which rely on manually created features. Deep learning methods, such as Convolutional Neural Networks (CNNs) and Long Short-Term Memory (LSTMs) with attention mechanisms, enable models to automatically extract features and better capture target-specific stance information (Gera and Neal, 2025). Current transformer architectures use BERT (Garg and Caragea, 2024) and Large Language Models (LLMs) (Pangtey et al., 2025) to deliver excellent results through their ability to understand context and their capacity for few-shot learning.

3. Proposed StanceMoE Architecture

We propose a context-enhanced Mixture-of-Experts architecture developed upon a fine-tuned BERT encoder for stance detection. The overall framework consists of three components: (i) a contextual encoder, (ii) six parallel expert modules cap-

¹<https://github.com/AbdullahRatulK/StanceMoE>

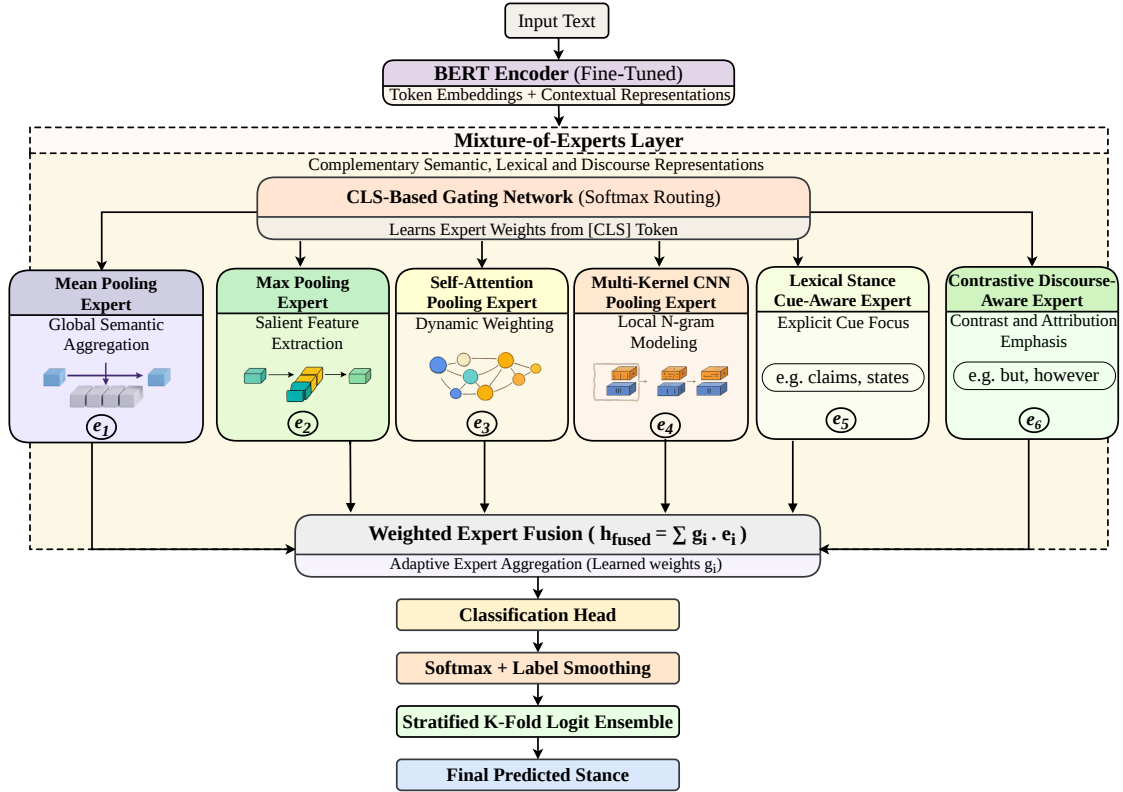


Figure 1: Proposed Mixture-of-Experts architecture for actor-level stance detection.

turing complementary linguistic signals, and (iii) a context-aware gating and fusion mechanism. The architecture is illustrated in Fig. 1.

3.1. Contextual Encoder

Given an input sequence, we employ a fine-tuned BERT encoder to obtain contextualized token representations and a special [CLS] embedding.

3.2. Expert Modules

While BERT provides promising contextual understanding, stance detection often depends on subtle lexical cues, contrast markers, and discourse patterns. Standard pooling techniques over contextual embeddings may not consistently capture these heterogeneous patterns. To address this, we introduce six complementary experts, each targeting a distinct linguistic phenomenon.

Mean Pooling Expert (Global Orientation): This expert captures the overall semantic direction of a post by averaging token representations. It is effective when stance is consistently expressed throughout the sentence. For example, “*Israel has the right to defend itself.*” (Pro-Israel), or “*Palestinians deserve statehood and equal rights.*” (Pro-Palestine).

Max Pooling Expert (Salient Polarity Tokens): Some posts contain isolated but highly indicative tokens that determine stance. For instance, the token “*occupation*” serves as the main indicator

of stance in the Pro-Palestine statement “*This ongoing occupation must end*”. Similarly, the token “*terrorist*” serves as the main lexical indicator in the Pro-Israel statement “*The terrorist attacks cannot be justified*”. Max pooling selects the most activated feature dimensions to emphasize important tokens throughout the entire sequence.

Self-Attention Pooling Expert (Context-Dependent Focus): The specific clauses of a sentence work as the place where stance exists. The sentence “*While the humanitarian situation is tragic, Israel must respond to security threats*” demonstrate an example of such configurations. In such formulations, one clause presents the main stance while the additional information is presented through the rest of the sentence. The attention mechanism gives more weight to stance components because it assesses their importance throughout the evaluation process.

Multi-Kernel CNN Expert (Phrase-Level Patterns): Geopolitical stance is mostly expressed through short slogan-like phrases which include “*Stand with Israel*” (Pro-Israel) and “*Free Palestine*” (Pro-Palestine). The compact n-grams of the expression serve as effective stance indicators which maintain their function despite minimal contextual information. Multi-kernel convolutions capture such localized phrase patterns.

Lexical Cue-Aware Expert (Framing and Reporting Signals): Neutral posts use reporting lan-

guage and distancing language to express content like the statement “*According to officials, negotiations are ongoing.*” Verbs like “*claims,*” “*reports,*” “*states,*” function as indicators that speakers use indirect methods to present information instead of showing direct support for their statements. This expert combines different ways of showing predetermined cue tokens to improve the ability to separate neutral content from opinionated postings.

Contrast-Aware Expert (Discourse Shift Modeling): Discourse contrast markers such as “*but*” or “*however*” often indicate stance refinement or emphasis. For example, “*I support peace efforts, but the blockade must end.*” or “*Civilian harm is tragic; however, security responses are necessary.*” In such cases, the clause containing the contrast marker often carries stronger stance weight. We therefore amplify representations of contrast tokens to better capture discourse-level stance shifts.

3.3. Context-Aware Gating and Fusion

We introduce a gating mechanism that dynamically weights the experts instead of treating them equally. The final representation is computed as a weighted combination of expert outputs. This design enables adaptive routing. For instance, strongly opinionated posts may rely more on max pooling, contrast-heavy statements may increase the weight of the contrast-aware expert, report-style neutral texts may emphasize cue-aware representations.

The fused representation is then passed to a linear classifier. We train the model using cross-entropy loss with label smoothing. To improve robustness, we employ stratified K -fold cross-validation during training and perform weighted logit averaging of those folds during inference. The weighting is done by giving more weight to folds having a higher F1-Score during validation.

4. Experimental Setup

4.1. Dataset

The dataset used in this paper is provided by the organizers of the StanceNakba 2026 Shared Task (Aldous et al., 2026). We focus on Subtask A, an actor-level stance detection dataset comprising 1,401 English texts labeled as Pro-Palestine, Pro-Israel, or Neutral. The stance labels denote the author’s alignment with actors, which are implicit in the text and must be inferred by the model. The official split follows a 70/15/15 split for training, development, and test sets. In our experiments, we apply stratified K -fold cross-validation on the combined training and development set (85% of the data). The official test set (15%) is kept strictly unseen during training.

4.2. Baseline Methods

We perform an extensive empirical assessment of existing stance detection techniques, which are classified into three major methodological categories: Machine Learning (ML) methods with TF-IDF features, supervised training with Deep Neural Networks (DNNs), and BERT. The ML methods include Logistic Regression (LR) (Rahman et al., 2024), Multinomial Naive Bayes (MNB) (Zannat et al., 2025), Support Vector Machine (SVM) (Rahman et al., 2024), and Random Forest (RF) (Shafi et al., 2025). We implement several popular DNN architectures as baselines, including Bi-directional Long Short-Term Memory (BiLSTM) (Rahman et al., 2024), Target specific Attentional Network (TAN) (Du et al., 2017), Gated Convolutional network with Aspect Embedding (GCAE) (Xue and Li, 2018), Cross network (CrossNet) (Du et al., 2017). Moreover, we evaluate two alternative variants developed on the BERT backbone: (i) a stacked architecture (Khan et al., 2025) where expert modules are applied sequentially, and (ii) a feature fusion model (Lee et al., 2021) where expert outputs are fused without adaptive weighting.

5. Experimental Results

5.1. Actor-level Stance Detection

Table 1 presents the comparative performance of all models for actor-level stance detection. Among traditional baselines, RF achieves the best results with an F1-Score of 84.19%, which shows that ensemble tree-based methods perform better than other methods. Among neural models, GCAE performs highest (87.62% F1), which shows that attention-based architectures perform well to model stance. The use of contextual embeddings results in better performance outcomes. BERT achieves an F1-Score of 89.86%, while expert-based architectures like Stacked (91.83%) and Fusion (91.17%) show promising results, but their improvements remain within fixed expert integration boundaries.

The proposed StanceMoE achieves the best performance with an F1-Score of 94.26%, which surpasses all other cases of comparison. Additionally, our team has stood 3rd in the competition. The system demonstrates performance enhancement because of its adaptive expert weighting system, which uses a gating mechanism. The proposed system also shows stable performance with low cross-validation variance ($\approx \pm 1.1\%$).

5.2. Ablation Study

Table 2 presents an ablation analysis to examine the contribution of each expert in the proposed StanceMoE architecture. Removing any expert

Methods	K-fold (mean±std)				Weighted Logit Ensemble			
	Acc	Pre	Rec	F1	Acc	Pre	Rec	F1
LR	80.91±1.79	80.90±1.78	80.93±1.78	80.86±1.78	81.28	81.25	81.33	81.22
MNB	77.25±1.64	75.88±2.02	78.49±1.65	77.20±1.65	77.38	76.02	78.63	77.33
SVM	83.25±1.71	84.73±1.68	83.26±1.71	83.27±1.72	83.37	84.94	83.35	83.43
RF	84.05±1.83	84.73±1.73	84.05±1.82	84.05±1.90	84.16	84.83	84.14	84.19
BiLSTM	85.63±2.99	85.87±3.00	85.63±2.99	85.51±3.07	85.74	85.98	85.74	85.72
TAN	85.99±3.10	86.01±2.81	86.36±2.08	85.91±2.15	86.13	86.10	86.43	86.03
GCAE	87.69±2.85	87.98±2.92	87.48±2.16	87.49±2.78	87.82	88.10	87.53	87.62
CrossNet	85.13±2.28	85.55±1.89	84.90±1.90	84.77±2.25	85.29	85.65	84.97	84.88
BERT	89.77±2.35	90.04±2.30	89.77±2.31	89.61±2.29	90.05	90.28	90.03	89.86
Stacked	91.61±2.46	91.73±2.43	91.62±2.44	91.60±2.47	91.94	92.14	91.93	91.83
Fusion	91.03±2.26	91.20±2.29	91.03±2.24	91.02±2.26	91.18	91.45	91.17	91.17
StanceMoE	94.09±1.11	94.18±1.12	94.08±1.12	94.03±1.12	94.31	94.45	94.31	94.26

Table 1: Performance comparison of baseline models and the proposed StanceMoE on the held-out test set. “*K-fold (mean ± std)*” reports the mean and standard deviation of test performance across fold-specific models, while “*weighted logit ensemble*” is the final ensemble prediction obtained via logit averaging. Here, Acc = Accuracy, Pre = Precision, Rec = Recall, and F1 = F1-Score.

Methods	Acc	Pro-Palestine			Pro-Israel			Neutral		
		Pre	Rec	F1	Pre	Rec	F1	Pre	Rec	F1
w/o Mean	92.89	89.47	95.75	92.52	94.52	98.57	96.50	95.16	84.29	89.39
w/o Max	93.36	91.89	95.77	93.79	92.00	98.57	95.17	96.77	85.71	90.91
w/o Self-Attention	91.94	89.33	94.37	91.78	93.24	98.57	95.83	93.55	82.86	87.88
w/o CNN	93.36	92.96	92.96	92.96	93.33	99.26	96.55	93.85	87.14	90.37
w/o Lexical-cue	93.36	90.54	94.37	92.41	95.83	98.57	97.18	95.08	82.86	88.55
w/o Contrastive	91.94	90.41	92.96	91.67	90.91	99.32	95.24	93.85	87.14	90.37
StanceMoE	94.31	94.37	94.37	94.37	92.11	100	95.89	96.88	88.57	92.54

Table 2: Classwise ablation study of StanceMoE using weighted logit ensemble, showing the effect of removing individual expert modules.

generally degrades overall performance, confirming that the experts capture complementary stance signals. The largest drops occur when removing the self-attention or contrastive experts, indicating their importance in modeling contextual dependencies and discourse-level stance shifts. Excluding the mean or max pooling experts also reduces performance, suggesting the importance of salient lexical cues and global contextual signals for stance identification.

The lexical-cue expert particularly benefits the Pro-Palestine class, while the contrastive expert improves recognition of stance transitions and argumentative structures. Although some variants maintain competitive scores for individual classes, they generally exhibit reduced balance across classes, especially for the neutral category. Overall, the full StanceMoE model achieves the best and most balanced performance across all classes, demonstrating the effectiveness of adaptive expert integration.

6. Conclusion

In this paper, we introduced StanceMoE, a context-enhanced MoE architecture for stance detection. Motivated by the limitations of unified transformer representations capturing heterogeneous linguistic phenomena, our approach explicitly models complementary stance-indicative signals through expert modules and integrates them via a context-aware gating mechanism. This design enables adaptive modeling of lexical, semantic, and discourse-level patterns that characterize stance expression in sensitive contexts. Extensive experiments demonstrate that StanceMoE substantially outperforms traditional baseline models. Ablation analysis further confirms the complementary contributions of individual experts and the importance of adaptive expert fusion for robust detection. Overall, our findings highlight the effectiveness of explicitly disentangling and dynamically integrating diverse linguistic signals for stance detection.

Bibliographical References

- Kholoud Khalil Aldous, Md Rafiul Biswas, Mabrouka Bessghaier, Shimaa Ibrahim, Kais Attia, and Wajdi Zaghouani. 2026. StanceNakba shared task: Actor and topic-aware stance detection in public discourse. In *Proceedings of the 15th International Conference on Language Resources and Evaluation (LREC'26)*, Palma, Spain.
- Izzat Alsmadi, Iyad Alazzam, Mohammad Al-Ramahi, and Mohammad Zarour. 2024. Stance detection in the context of fake news—a new approach. *Future Internet*, 16(10):364.
- Nora Alturayef, Hamzah Luqman, and Moataz Ahmed. 2023. A systematic review of machine learning techniques for stance detection and its applications. *Neural Computing and Applications*, 35(7):5113–5144.
- Michael Burnham. 2025. Stance detection: a practical guide to classifying political beliefs in text. *Political Science Research and Methods*, 13(3):611–628.
- Jiachen Du, Ruifeng Xu, Yulan He, and Lin Gui. 2017. Stance classification with target-specific neural attention networks. In *2017 26th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 3988–3994.
- Krishna Garg and Cornelia Caragea. 2024. Stanceformer: Target-aware transformer for stance detection. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 4969–4984.
- Parush Gera and Tempestt Neal. 2025. Deep learning in stance detection: A survey. *ACM Computing Surveys*, 58(1):1–37.
- Zhijiang Guo, Michael Schlichtkrull, and Andreas Vlachos. 2022. A survey on automated fact-checking. *Transactions of the Association for Computational Linguistics*, 10:178–206.
- Lal Khan, Atika Qazi, Hsien-Tsung Chang, Mousa Alhajlah, and Awais Mahmood. 2025. Empowering urdu sentiment analysis: an attention-based stacked cnn-bi-lstm dnn with multilingual bert. *Complex & Intelligent Systems*, 11(1):10.
- Dilek Küçük and Fazli Can. 2020. Stance detection: A survey. *ACM Computing Surveys*, 53(1):1–37.
- Sanghyun Lee, David K Han, and Hanseok Ko. 2021. Multimodal emotion recognition fusion analysis adapting bert with heterogeneous feature unification. *IEEE Access*, 9:94557–94572.
- Junxia Ma, Changjiang Wang, Lu Rong, Bo Wang, and Yaoli Xu. 2025. Exploring multi-agent debate for zero-shot stance detection: A novel approach. *Applied Sciences*, 15(9):4612.
- Fuqiang Niu, Min Yang, Ang Li, Baoquan Zhang, Xiaojiang Peng, and Bowen Zhang. 2024. A challenge dataset and effective models for conversational stance detection. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 122–132.
- Lata Pangtey, Anukriti Bhatnagar, Shubhi Bansal, Shahid Shafi Dar, and Nagendra Kumar. 2025. Large language models meet stance detection: A survey of tasks, methods, applications, challenges and future directions. *arXiv:2505.08464*.
- Arifur Rahman, Shahriar Parvej, Kazi Saeed Alam, and HM Abdul Fattah. 2024. Optimizing sms spam detection: comparative analysis of hybrid voting ensembles and bi-lstm networks with stratified cross-validation. In *2024 5th International Conference on Data Intelligence and Cognitive Informatics (ICDICI)*, pages 1030–1035. IEEE.
- Abdullah Al Shafi, Rowzatul Zannat, Abdul Muntakim, and Mahmudul Hasan. 2025. A structured dataset of disease-symptom associations to improve diagnostic accuracy. *arXiv:2506.13610*.
- Wei Xue and Tao Li. 2018. Aspect based sentiment analysis with gated convolutional networks. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2514–2523.
- Ruichao Yang, Jing Ma, Wei Gao, and Hongzhan Lin. 2025. Llm-enhanced multiple instance learning for joint rumor and stance detection with social context information. *ACM Transactions on Intelligent Systems and Technology*, 16(3):1–27.
- Rowzatul Zannat, Abdullah Al Shafi, and Abdul Muntakim. 2025. Bridging the gap in bangla healthcare: Machine learning based disease prediction using a symptoms-disease dataset. In *2025 International Conference on Electrical, Computer and Communication Engineering (ECCE)*, pages 1–6. IEEE.
- Bowen Zhang, Genan Dai, Fuqiang Niu, Nan Yin, Xiaomao Fan, Senzhang Wang, Xiaochun Cao, and Hu Huang. 2024. A survey of stance detection on social media: New directions and perspectives. *arXiv:2409.15690*.
- Zhengyuan Zhu, Zeyu Zhang, Haiqi Zhang, and Chengkai Li. 2025. Ratsd: Retrieval augmented

truthfulness stance detection from social media posts toward factual claims. In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 3366–3381.

A. Need of Stance Detection

The ability to detect stance has become increasingly important due to the growth of online discourse. Applications span multiple domains, including electoral trend analysis (Burnham, 2025), examination of argumentative interactions in online debates (Ma et al., 2025), monitoring of social media discussions (Zhang et al., 2024), rumor assessment and verification (Yang et al., 2025), automated fact-checking systems (Guo et al., 2022), fake news identification (Alsmadi et al., 2024), and stance-aware information retrieval (Zhu et al., 2025). In these contexts, understanding directional opinion toward a claim or actor is more informative than simply measuring sentiment (Garg and Caragea, 2024).

B. Operational Details of Expert Modules

Given an input sequence $X = \{x_1, \dots, x_T\}$, pre-trained BERT encoder obtains contextualized token representations: $H = \{h_1, h_2, \dots, h_T\}$, $h_i \in \mathbb{R}^d$ where d denotes the hidden dimension and T is the length of the tokenized input sequence.

Mean Pooling Expert: This expert captures global semantic information as shown in (1).

$$e_1 = W_1 \left(\frac{1}{T} \sum_{i=1}^T h_i \right) \quad (1)$$

Max Pooling Expert: To extract salient token-level features, we utilize max pooling expert as in (2).

$$e_2 = W_2 \left(\max_{i=1}^T h_i \right) \quad (2)$$

Self-Attention Pooling Expert: We introduce a trainable attention vector v to compute token importance as mentioned in (3).

$$\alpha_i = \frac{\exp(\tanh(h_i^\top v))}{\sum_{j=1}^T \exp(\tanh(h_j^\top v))} \quad (3)$$

$$e_3 = W_3 \left(\sum_{i=1}^T \alpha_i h_i \right) \quad (4)$$

Multi-Kernel CNN Expert: To model local n-gram patterns, we apply one-dimensional convolutions with kernel sizes $k \in \{2, 3, 4, 5\}$ as demonstrated in (5).

$$e_4 = W_4 \left(\text{Concat}(\text{MeanPool}(\text{ReLU}(\text{Conv}_k(H)))) \right) \quad (5)$$

Lexical Cue-Aware Expert: We define a set of stance-indicative lexical cues and generate a binary mask over their positions C . The cue-aware representation is computed as (6).

$$e_5 = W_5 \left(\frac{\sum_{i \in C} h_i}{|C| + \epsilon} \right) \quad (6)$$

where, ϵ avoids division by zero.

Contrast-Aware Expert: To model discourse contrast markers (e.g., “but”, “however”), we enhance their contextual influence. Let, D denote the set of contrast token positions. The overall process is shown in (8).

$$\tilde{h}_i = \begin{cases} 3h_i & \text{if } i \in D, \\ h_i & \text{otherwise,} \end{cases} \quad (7)$$

$$e_6 = W_6 \left(\frac{\sum_{i=1}^T \tilde{h}_i}{|D| + \epsilon} \right) \quad (8)$$

C. Operational Details of Context-Aware Gating

The gating network takes the ‘[CLS]’ representation $h_{\text{cls}} \in \mathbb{R}^d$ from the BERT encoder as input. A learnable linear layer $W_g \in \mathbb{R}^{N \times d}$ with bias $b_g \in \mathbb{R}^N$ computes logits for $N = 6$ experts, which are normalized via a softmax to generate the gating weights as in (9).

$$g = \text{Softmax}(W_g h_{\text{cls}} + b_g) \quad (9)$$

where $g \in \mathbb{R}^K$, $g_i \geq 0$ for all i , and $\sum_{i=1}^K g_i = 1$. Each gating weight g_i determines the contribution of the corresponding expert output e_i to the final representation, which is obtained via weighted aggregation using (10).

$$h_{\text{moe}} = \sum_{i=1}^K g_i e_i \quad (10)$$

The fused representation h_{moe} is then passed through a task-specific linear layer with learnable weights W_o and bias b_o to generate the model predictions through (11).

$$\hat{y} = \text{softmax}(W_o h_{\text{moe}} + b_o) \quad (11)$$

where \hat{y} represents the predicted class probabilities. The parameters W_g , b_g , W_o , and b_o are all learnable and optimized jointly with the rest of the network via backpropagation, allowing the model to learn both how to combine experts and how to make accurate predictions.

D. Hyperparameters

Table 3 shows the hyperparameters used in the experiment.

Hyperparameter	Value
Max sequence length	128
Batch size	16
Number of epochs	10
Learning rate	5×10^{-5}
Number of splits (k)	10
Label smoothing factor	0.25
Random seed	42

Table 3: Training hyperparameters used in StanceMoE.

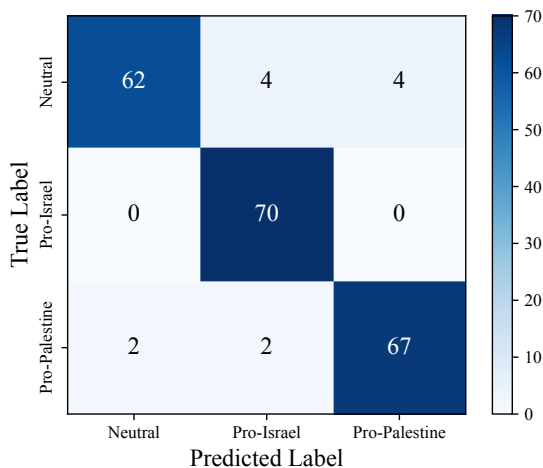


Figure 2: Confusion matrix using the proposed StanceMoE architecture.

E. Baseline ML methods

We utilize four popular ML models as our baseline systems: (1) Logistic Regression (LR): A linear model that predicts class probabilities using weighted features, effective for high-dimensional text data. (2) Multinomial Naive Bayes (MNB): A probabilistic model based on word frequency distributions, assuming feature independence, well-suited for text classification. (3) Support Vector Machine (SVM): A margin-based classifier that finds the optimal hyperplane to separate classes, robust in sparse text feature spaces. (4) Random Forest (RF): An ensemble of decision trees that improves classification performance by reducing overfitting through bagging and feature randomness.

F. Baseline DNN

We use four popular deep neural network models as our baseline systems: (1) BiLSTM (Rahman et al., 2024) is used to create bi-directional context models which predict stances without requiring to use target information. (2) GCAE (Xue and Li, 2018) provides a convolutional system which employs a gating mechanism to block non-target fea-

tures. (3) TAN (Du et al., 2017) uses an attention-boosted BiLSTM to identify important contextual details which help to determine stances. (4) CrossNet (Du et al., 2017) enhances attention processing through an aspect-based attention layer which operates before the classification process to enhance target-based feature extraction.

G. Evaluation Metrics

We evaluate model performance using macro-F1 score as the primary metric. Additionally, accuracy, macro precision, and macro recall are reported.

H. Overall Ablation Study

Table 4 shows the overall ablation study on the test set, rather than class-specific one.

I. Confusion matrix

Fig. 2 shows the confusion matrix using proposed StanceMoE architecture.

J. Error analysis

The qualitative error analysis in Table 5 reveals several systematic patterns.

(1) The model occasionally exhibits an excessive dependence on strong lexical polarity cues. The first and fifth examples show religious praise as well as humanitarian language which are mistakenly interpreted as explicit political alignment.

(2) Meta-discursive commentary is sometimes confused with stance. The text in the third example describes antisemitism together with social norms without showing any specific support. The system uses target-related keywords to make stance predictions about the content. This is difficult because it requires the model to differentiate between entity mentioned and evaluative positioning.

(3) Complex multi-clause and contrastive reasoning remains challenging. The second and sixth examples contain intricate discourse structures that display stance through their multiple levels of argumentation. The system design includes experts to handle both attention and contrast decoding tasks, but in some cases the gating system fails to deliver proper evaluation for these specialized discourse elements.

(4) Very short and context-poor statements like the fourth example are inherently ambiguous. In such cases, the model appears to associate the polarity with dominant training patterns rather than grounding the prediction in the explicit evaluation of the target.

Methods	K-fold (mean±std)				Weighted Logit Ensemble			
	Acc	Pre	Rec	F1	Acc	Pre	Rec	F1
w/o Mean	91.75±1.76	92.02±1.54	91.74±1.77	91.65±1.86	92.89	93.05	92.88	92.80
w/o Max	91.52±1.23	91.83±1.06	91.51±1.23	91.43±1.30	93.36	93.56	93.35	93.29
w/o Self-att	91.47±1.04	91.65±0.99	91.46±1.04	91.38±1.08	91.94	92.04	91.93	91.83
w/o CNN	92.09±1.04	92.20±0.97	92.08±1.04	92.00±1.11	93.36	93.38	93.37	93.29
w/o Lexical-cue	91.70±1.84	91.83±1.84	91.70±1.84	91.67±1.83	93.36	93.41	93.36	93.32
w/o Contrastive	91.75±1.81	91.99±1.67	91.74±1.82	91.64±1.87	91.94	92.13	91.94	91.82
StanceMoE	94.09±1.11	94.18±1.12	94.08±1.12	94.03±1.12	94.31	94.45	94.31	94.26

Table 4: Overall K-fold and weighted ensemble ablation study on the test set.

Text	Actual	Pre-dicted	Primary Error Source
This is shameful. And the Jewish people, God’s chosen people. May God bless and protect you.	Neutral	Pro-Israel	Religious praise misinterpreted as political stance due to strong positive lexical cues.
Most comments are clothed with sycophancy. Nigerians have a right of association, but you can stand with Gaza. The people of Gaza started this aggressive behaviour on 7th October when they launched a coordinated attack against the state of Israel.	Pro-Palestine	Neutral	Complex multi-clause reasoning; insufficient emphasis on dominant stance clause.
If you are talking about the doctor not talking about Palestine. He is German living in Germany where any anti-Jewish sentiment is taken very seriously and to many people anti-Israel equals antisemitism, so it is in his better interest not to speak about it at all, not necessarily meaning he supports anyone.	Neutral	Pro-Israel	Confusion between topic discussion and stance expression due to keyword activation.
Against the Islamist terrorists who hate Jews and the West.	Neutral	Pro-Palestine	Short ambiguous input; polarity incorrectly associated with a stance.
As I stood up with BLM, I stand against oppression in all forms. Release Palestinian hostages before talking about anything else.	Neutral	Pro-Palestine	Humanitarian framing interpreted as explicit alignment due to lexical cues.
I was raised Muslim, and I know intimately how oppressive the religion can be in various ways. That being said, because I understand that Palestinian gay people are being oppressed for more than their sexuality right now, and that matters too.	Pro-Palestine	Neutral	Failure to capture nuanced contrastive stance; insufficient weighting of discourse signals.

Table 5: Representative genuine misclassification cases of StanceMoE. Different background colors are used for visual distinction.

Overall, these errors suggest that while StanceMoE is effective for explicit and lexically polarized stance expressions, nuanced and implicitly framed discourse remains a challenging direction for future improvement.

K. Future Work

Future research may extend this framework toward explicitly target-aware stance modeling or explore cross-topic and cross-domain stance detection to further evaluate the robustness and generalization capacity of the proposed architecture. Investigating multilingual adaptation may also enhance its applicability to broader geopolitical discourse contexts.

L. Acknowledgement

We would like to express our sincere gratitude to the organizer of StanceNakba 2026 shared task, and anonymous reviewers for their support. The authors acknowledge the use of ChatGPT (OpenAI) as an assistive tool for language refinement, code suggestions, and conceptual structuring. All generated content was critically reviewed, validated, and appropriately adapted by the authors, who take full responsibility for the accuracy and integrity of the work.