The Indonesian Journal of Computer Science
www.ijcs.net
Volume 14, Issue 6, December 2025
https://doi.org/10.33022/ijcs.v14i6.5065

Unsupervised Clustering of Vietnamese Positive and Negative News Using PhoBERT and DBSCAN

Tuan Long Dinh[1], **Ngoc An Le**[2]
dinhtuanlong@hou.edu.vn[1], anln@hou.edu.vn[2]
[1,2]Faculty of Information Technology, Hanoi Open University, Hanoi, Vietnam

Article Information

Received : 18 Nov 2025
Revised : 26 Dec 2025
Accepted : 30 Dec 2025

Keywords

Vietnamese sentiment analysis, unsupervised clustering, PhoBERT, DBSCAN, news text mining

Abstract

The proliferation of digital media has made detecting and analyzing sentiment trends in Vietnamese news content increasingly important. This paper proposes an unsupervised learning approach for clustering Vietnamese news articles into positive and negative sentiment categories. The model combines headline and content features using PhoBERT, a Vietnamese-optimized language model, with DBSCAN clustering. Text is encoded using PhoBERT-base for headlines (768 dimensions) and PhoBERT-large for content (1024 dimensions), then concatenated and reduced to 64 dimensions via UMAP before clustering. KeyPhoBERT extracts representative keywords to enhance interpretability. Evaluated on 1,180 manually annotated articles from university social media with inter-annotator agreement of Cohen's kappa 0.83, the model achieves F1-score of 94.37%, with Adjusted Rand Index of 0.87 and Normalized Mutual Information of 0.81. Comparison with BERTopic baseline demonstrates the effectiveness of our approach for Vietnamese sentiment clustering without requiring labeled training data.

A. Introduction

The rapid growth of digital media platforms has fundamentally transformed how information spreads within educational communities. University social media channels now serve as primary communication hubs where students share opinions about academic activities, institutional policies, and campus life. This information landscape includes both positive content that reinforces educational values and negative content such as complaints or controversies that may affect community perceptions. Detecting, analyzing, and monitoring these opposing information streams plays a crucial role in internal communication management, protecting the educational environment, and understanding learner psychology in a timely manner.

Automatically categorizing such content by sentiment polarity presents significant challenges for Vietnamese natural language processing. Vietnamese remains a low-resource language with limited labeled datasets, making supervised learning approaches impractical for many real-world applications. While semi-supervised methods can reduce labeling requirements, they still depend on initial label quality and may not generalize well to continuously evolving social media content. In this context, unsupervised learning methods emerge as a potential solution to effectively discover hidden information structures in data without requiring labels.

Recent research has increasingly combined semantic representation with unsupervised clustering for content analysis. Mayor et al. [1] analyzed temporal relationships between emotions in online news and Twitter using large-scale automated analysis of nearly 1 million articles and over 6 million tweets, showing that changes in general sentiment in news correspond to temporal changes in related social media content. Rybinski [2] explored news lifespan prediction integrating NLP techniques like LDA topic modeling, relative sentiment analysis, and emotion detection, finding that positive news tends to have longer engagement periods. For topic discovery in specialized domains, Alqurashi and Ahmad [3] applied BERTopic with UMAP (Uniform Manifold Approximation and Projection) [4] dimensionality reduction and HDBSCAN clustering to analyze cybersecurity literature from 15,751 academic papers and 5,831 industry papers, demonstrating the effectiveness of transformer-based embeddings for domain-specific clustering. Liu et al. [5] proposed graph-based methods for news event detection, integrating semantic and temporal embeddings into a graph structure and applying convolutional updates to capture contextual dynamics.

However, most existing deep learning models are primarily trained on English data, limiting their ability to understand Vietnamese semantics. Traditional techniques like n-gram representations are easy to implement but lack deep contextual representation capability, while multilingual models such as mBERT show limited performance on Vietnamese compared to language-specific models. PhoBERT [6], built on the RoBERTa architecture and trained on large-scale Vietnamese corpora including Vietnamese Wikipedia, VCCorpus, and OSCAR, has demonstrated superior performance for Vietnamese NLP tasks through its ability to capture Vietnamese language characteristics. Recent work by Nguyen et al. [7] on Vietnamese text classification and Le et al. [8] on Vietnamese sentiment analysis has shown the benefits of Vietnamese-specific pre-training. Meanwhile, clustering algorithms like DBSCAN [9] and HDBSCAN [10] have proven effective for discovering clusters of arbitrary shapes without requiring prior specification of cluster numbers, but have not been effectively combined with representations from Vietnamese-specific language models. Current approaches also lack focus on educational content analysis, particularly in understanding the sentiment dynamics unique to university environments with their colloquial language, academic terminology, and context-specific expressions.

This study addresses these challenges by proposing an unsupervised clustering approach that leverages PhoBERT combined with DBSCAN to discover natural sentiment groupings in Vietnamese news content without any supervision. The approach combines semantic representations from both headlines and article content, applies dimensionality reduction via UMAP, and uses DBSCAN to identify clusters corresponding to positive and negative sentiment categories. We also provide comparison with BERTopic [11], a widely-used topic modeling approach that combines UMAP, HDBSCAN, and c-TF-IDF, to demonstrate the effectiveness of our design choices. The main contributions of this paper are: (i) an unsupervised clustering framework combining PhoBERT representations with DBSCAN for Vietnamese sentiment classification, effectively exploiting semantic features from both headlines and content; (ii) a publicly described annotation protocol with inter-annotator agreement metrics for Vietnamese educational content from real-world university social media; and (iii) comprehensive evaluation using standard clustering metrics including Adjusted Rand Index and Normalized Mutual Information, with comparison against BERTopic baseline, confirming the model's potential for content analysis systems and public opinion monitoring without requiring any labels.

B. Research Method

B.1. Overview of Models and Processing Techniques

In recent years, deep learning models based on the Transformer architecture have achieved remarkable results in Natural Language Processing, including text classification and clustering tasks. Thanks to their ability to capture contextualized semantic representations, models such as BERT [12][13], RoBERTa [14], and PhoBERT [6] have become valuable tools in modern text mining systems.

PhoBERT is the first pre-trained language model built specifically for Vietnamese based on the RoBERTa architecture. Instead of using word-level or character-level tokenization, PhoBERT applies Byte Pair Encoding (BPE) combined with large-scale Vietnamese corpora (Vietnamese Wikipedia, VCCorpus, OSCAR, etc.) to effectively learn semantic representations from context. PhoBERT's greatest advantage is its superior ability to capture Vietnamese language characteristics compared to multilingual models like mBERT or XLM-R.

DBSCAN (Density-Based Spatial Clustering of Applications with Noise) [9] is a popular unsupervised clustering algorithm that operates based on data point density in feature space. DBSCAN can automatically determine the number of clusters without prior specification while effectively detecting and removing noise points. DBSCAN's advantages lie in its ability to handle arbitrarily shaped clusters, no distribution assumptions, and good performance on noisy or density-heterogeneous data.

B.2. Model Architecture

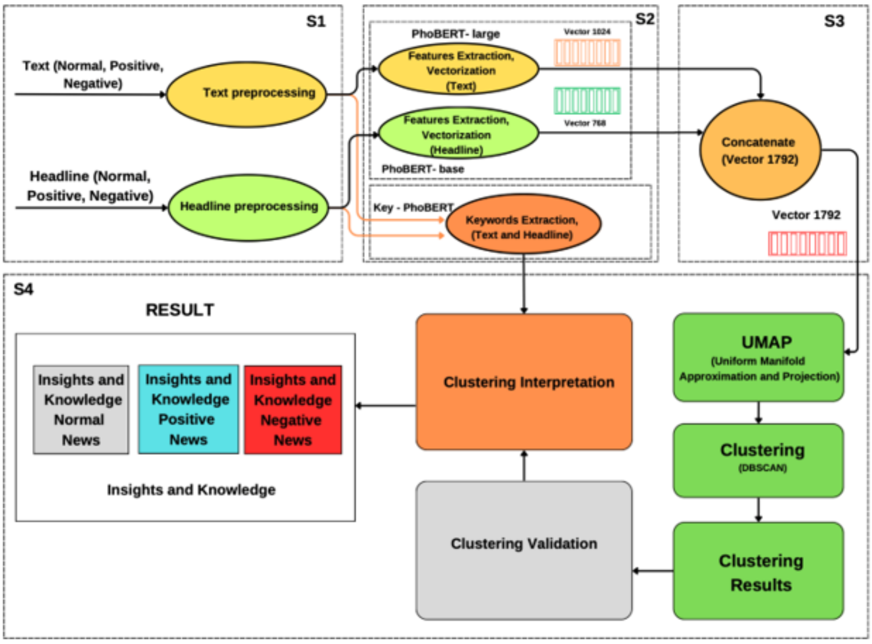

Figure 1. Architecture of the proposed model

The proposed model aims to cluster positive-negative news using unsupervised learning methods, simultaneously exploiting information from both headlines and content. Figure 1 illustrates the overall architecture of the proposed model.

The processing pipeline consists of four main stages. First, both headline and content text undergo preprocessing. Second, the model employs KeyPhoBERT to extract representative keywords for cluster labeling. Third, PhoBERT-base processes headlines to produce 768-dimensional vectors, while PhoBERT-large processes content to create 1024-dimensional vectors. These two vectors are concatenated into a unified 1792-dimensional vector representing the entire news article. Fourth, UMAP reduces dimensions from 1792 to 64, and finally DBSCAN performs clustering on the reduced vectors.

B.3. Dataset Description

Collection Period and Methodology: The dataset was collected from January to June 2024 from official fanpages and forums of Hanoi Open University. Data collection employed automated scraping tools combined with manual verification to ensure data quality and relevance. Only posts with clear distinction between headlines and body content were retained; posts with non-informative headers (e.g., emojis, single words) or lacking clear structure were filtered out during preprocessing.

Dataset Composition: The dataset comprises 1,180 Vietnamese news articles with the following distribution: 623 positive articles (52.8%) and 557 negative articles (47.2%). Positive articles include academic announcements, achievement celebrations, and informative educational content. Negative articles contain complaints, criticisms, controversial discussions, or content expressing dissatisfaction. The average headline length is 12.3 words and average content length is 89.7 words. Topics covered include academic announcements, student activities, policy changes, campus events, student feedback, and administrative communications.

Annotation Protocol: The dataset was manually annotated by a team of three trained undergraduate annotators with prior experience in processing educational communication content. A detailed annotation guideline was developed defining positive content as articles expressing satisfaction, achievement, constructive information, or positive institutional image, while negative content was defined as articles expressing complaints, criticism, controversy, or content potentially harmful to institutional reputation. Each post was independently labeled by two team members. In cases of disagreement (127 cases, 10.8%), the final label was determined through group discussion, followed by a final review and approval by a senior researcher. Inter-annotator agreement measured by Cohen's kappa was 0.83, indicating strong agreement.

Dataset Challenges: Compared to existing Vietnamese sentiment datasets, our data presents unique challenges including: (a) mixed sentiment articles containing both positive and negative elements, (b) sarcastic or ironic expressions common in student discourse, (c) domain-specific academic terminology, and (d) informal language and abbreviations typical of social media.

B.4. Clustering Algorithm

The clustering process follows a systematic algorithm with formal notation as follows:

Let the news dataset be D = $\{x_1, x_2, ..., x_n\}$, where each x_i comprises headline_i and content_i. The key functions are defined as:

- KeyPhoBERT(x_i): Keyword extraction function for text x_i
- PhoBERT_base($headline_i$): Feature extraction from headline → 768-dim vector
- PhoBERT_large($content_i$): Feature extraction from content → 1024-dim vector
- UMAP(z_i): Dimensionality reduction function from 1792 → 64 dimensions
- DBSCAN(*eps*, *min_samples*): Unsupervised clustering algorithm

Step 1 - Data Preprocessing: Both headline and content text undergo normalization by converting to lowercase, removing special characters, URLs, emojis, and unnecessary punctuation. This step helps the model focus on core content and reduce semantic noise. Common spelling errors and abbreviations frequently encountered in internet language are standardized to full forms to ensure semantic consistency when extracting embeddings from PhoBERT.

Step 2 - Feature Extraction: After preprocessing, the model employs KeyPhoBERT to extract key semantic terms from each news article. Each input text (combination of title and body content) is processed by KeyPhoBERT to generate semantically rich keywords using n-gram techniques (with *n* ranging from 1 to 6) and semantic embeddings from PhoBERT. The extracted keywords are stored for cluster interpretation. Simultaneously, semantic feature extraction uses two PhoBERT versions: PhoBERT-base for headlines creating 768-dimensional feature vectors, while PhoBERT-large for content generates 1024-dimensional vectors.

Step 3 - Feature Combination: The 768-dimensional vector from headlines and 1024-dimensional vector from content are concatenated to form a unified 1792-dimensional vector:

$$z_i = [PhoBERT\_base(headline_i) \oplus PhoBERT\_large(content_i)]. \qquad (1)$$

This combination helps the model have a comprehensive view of news articles, with headlines providing quick overviews while content provides deeper details and context.

Step 4 - Dimensionality Reduction and Clustering: The 1792-dimensional vectors are reduced to 64 dimensions using UMAP with parameters *n_neighbors*=15 and *min_dist*=0.1, preserving the most critical semantic structure while accelerating the clustering phase. The reduced vectors are then passed into DBSCAN for clustering. DBSCAN is suitable because it does not require prior specification of the number of clusters and can detect clusters of arbitrary shapes while filtering out noise points. The key parameters include: *eps (epsilon)* = 0.8, the maximum distance between two points to be considered neighbors; and *min_samples* = 5, the minimum number of neighboring points required to form a cluster.

B.5. Evaluation Methodology

Since our approach is unsupervised, we evaluate clustering quality using both external metrics (comparing against ground truth labels) and internal metrics (assessing cluster structure).

Cluster-to-Label Matching: For each predicted cluster K_j containing at least 5 articles, we identify the five most representative articles based on their proximity to the cluster centroid. A predicted cluster K_j is considered a correct match to a ground truth category G_i if at least three out of the five representative articles belong to the same ground truth category. This enables calculation of Precision and Recall as follow [15]:

$$Precision = \sum \frac{Check(K_i, G_i)}{M} \qquad (2)$$

$$Recall = \sum \frac{Check(K_i, G_i)}{N} \qquad (3)$$

where M is number of clusters and N is number of ground truth categories.

Standard Clustering Metrics: We report Adjusted Rand Index (ARI) [16] measuring agreement between predicted and true labels adjusted for chance, Normalized Mutual Information (NMI) [16] quantifying shared information between clustering and ground truth, Silhouette Score [17] for internal cluster cohesion, and Davies-Bouldin Index [18] for cluster separation quality.

Stratified Evaluation: To assess stability, we perform 10-fold stratified evaluation where data is split into 10 folds maintaining class proportions. Each fold serves as held-out test data while the model clusters the remaining data. This stratified splitting ensures robust estimation of clustering performance across different data subsets. Importantly, no labels are used during the clustering process itself; labels are only used for post-hoc evaluation.

C. Result and Discussion

C.1. Experimental Setup

Experiments were conducted on 3 NVIDIA Tesla P40 GPUs (24GB each) using PyTorch 1.12 with the Transformers library. PhoBERT models (vinai/phobert-base and vinai/phobert-large) were used for embedding extraction. DBSCAN and UMAP implementations from scikit-learn and umap-learn libraries were employed.

C.2. Main Results

Table 1 presents the primary clustering results. With optimal DBSCAN parameters (eps=0.8, min_samples=5), the model achieves precision of 98.60%, recall of 90.48%, and F1-score of 94.37%.

Table 1. Primary Clustering Results

eps	min samples	Model	Features	Precision (%)	Recall (%)	F1-score (%)
0.8	5	PhoBERT	Title + Content	98.60	90.48	94.37

Table 2 presents a parameter sensitivity analysis of the DBSCAN clustering algorithm by varying the neighborhood radius parameter (eps) while keeping min_samples fixed at 5. To ensure consistency with the overall evaluation framework, both external clustering metrics (Adjusted Rand Index – ARI, and Normalized Mutual Information – NMI) and internal validation metrics (Silhouette Score and Davies–Bouldin Index) are reported.

The results indicate that the configuration with eps = 0.8 achieves the best overall clustering performance across all evaluation criteria. Specifically, this setting yields the highest ARI (0.87) and NMI (0.81), indicating strong agreement between the clustering assignments and ground-truth labels, while simultaneously producing the highest Silhouette Score (0.82) and the lowest Davies–Bouldin Index (0.52), which reflect compact and well-separated clusters.

Table 2. DBSCAN Parameter Analysis Results

eps	min_samples	ARI	NMI	Silhouette Score	Davies-Bouldin Index
0.7	5	0.74	0.69	0,68	0,75
0.8	5	0.87	0.81	0,82	0,52
0.9	5	0.76	0.71	0,72	0,68

Based on this analysis, the DBSCAN configuration with eps = 0.8 and min_samples = 5 is selected for subsequent experiments and comparative evaluation.

C.3. Baseline Comparison

Table 3 compares our proposed DBSCAN-based approach with BERTopic, a widely used topic modeling framework that integrates UMAP, HDBSCAN, and c-TF-IDF, under the same experimental and evaluation settings. We further report results using different embedding models to analyze the impact of representation quality.

Table 3. Comparison with Baseline Methods

Method	Embedding	ARI	NMI	F1-score
Proposed (DBSCAN)	PhoBERT	0.87	0.81	94.37%
BERTopic (HDBSCAN)	PhoBERT	0.79	0.74	87.52%
Proposed (DBSCAN)	Vistral-7B [19]	0.71	0.66	84.70%
Proposed (DBSCAN)	GloVe [20]	0.62	0.58	82.70%

Our approach outperforms BERTopic by 6.85 percentage points in F1-score and achieves a notably higher ARI (0.87 vs. 0.79). The performance gains can be primarily attributed to the explicit integration of headline and content features, which is empirically validated by the ablation study in Table 4, showing that the combined configuration consistently outperforms each individual component. In addition, the use of DBSCAN with a fixed eps parameter appears to provide more stable cluster boundaries than HDBSCAN's adaptive density threshold in the binary sentiment classification setting.

C.4. Ablation Studies

Table 4 shows the impact of different feature configurations. Combining headlines and content yields 5.62% improvement over headlines alone and 3.17% over content alone, confirming that both components provide complementary information.

Table 4. Feature Configuration Analysis

Feature Configuration	F1-score (%)
Headlines Only	88.75
Content Only	91.20
Headlines + Content	94.37

Table 5 analyzes the effect of UMAP embedding dimensionality on classification performance, while DBSCAN parameters are fixed to the optimal configuration identified in Table 2.

Table 5. Impact of UMAP Dimensionality on Classification Performance

UMAP Dim	F1-score
32	92.15%
64	94.37%
128	93.52%

C.5. Stability Analysis

Table 6 reports the results of a stratified 10-fold cross-validation conducted to assess the stability of the proposed approach, DBSCAN and UMAP parameters are fixed to the optimal configuration identified in previous experiments and re-applied consistently across all folds. The model demonstrates consistent performance across folds, achieving a mean F1-score of 93.55% with a low standard deviation (±0.39%), which indicates strong robustness and generalization capability.

Table 6. Stratified 10-Fold Evaluation Results

Fold	Precision (%)	Recall (%)	F1-score (%)
1	97.25	89.74	93.34
2	96.87	89.95	93.28
3	98.12	90.12	93.95
4	98.6	90.48	94.37
5	97.54	89.63	93.42
6	96.95	89.89	93.29
7	97.83	90.01	93.76
8	96.78	89.57	93.04
9	97.65	89.86	93.59
10	97.14	90.09	93.48
Mean	97.47	89.93	93.55
Std Dev	±0.56	±0.27	±0.39

C.6. Error Analysis and Limitations

A qualitative inspection of misclassified articles reveals several recurring error patterns. First, articles containing sarcastic or ironic expressions are occasionally misclassified, as such implicit meanings are difficult to infer from contextual embeddings alone. Second, very short articles consisting primarily of headlines tend to produce weaker semantic representations, which reduces clustering reliability. Third, articles expressing mixed sentiments which combining both positive and negative elements pose challenges for the adopted binary sentiment formulation.

The current evaluation is conducted on a single dataset collected within one university context. While the results demonstrate strong performance on Vietnamese educational news, further experiments are required to assess generalization to other domains, such as product reviews or political news. Moreover, the binary positive/negative sentiment setting does not capture neutral or fine-grained emotional categories, which limits the expressiveness of the current framework.

D. Conclusion

This paper presents an unsupervised clustering-based approach for Vietnamese sentiment analysis that integrates PhoBERT embeddings with DBSCAN clustering. Experimental results show that the proposed method achieves an F1-score of 94.37% and an ARI of 0.87 on a dataset of 1,180 Vietnamese news articles, outperforming the BERTopic baseline by 6.85 percentage points in F1-score. Ablation studies further demonstrate that combining headline and content representations through the concatenation of PhoBERT-base and PhoBERT-large embeddings significantly enhances sentiment discrimination. Notably, the proposed approach does not require labeled training data, making it well suited for low-resource scenarios.

Future work will focus on addressing current limitations by improving the handling of sarcastic and mixed-sentiment content, extending the evaluation to additional Vietnamese datasets across diverse domains, and exploring semi-supervised extensions that leverage limited labeled data. Overall, this study provides a solid foundation for unsupervised sentiment analysis in Vietnamese and offers insights that may be transferable to other low-resource languages with appropriate adaptation.

E. Acknowledgment

This work is supported by the scientific project No. MHN2025-02.22, funded by Hanoi Open University, Vietnam.

F. References

[1] E. Mayor, V. Eicher, A. Bangerter, I. Gilles, E. G. Green, and A. Clémence, "Associations between emotions expressed in internet news and subsequent emotional content on Twitter," *Heliyon*, vol. 8, no. 12, Art. no. e12421, 2022, doi: 10.1016/j.heliyon.2022.e12421.

[2] K. Rybinski, "Content still matters: A machine learning model for predicting news longevity from textual and context features," *Information Processing & Management*, vol. 60, no. 4, Art. no. 103358, 2023, doi: 10.1016/j.ipm.2023.103358.

[3] Alqurashi and A. Ahmad, "A data-driven multi-perspective approach to cybersecurity knowledge discovery through topic modeling," *Alexandria Engineering Journal*, vol. 107, pp. 374–389, 2024, doi: 10.1016/j.aej.2024.07.044.

[4] L. McInnes, J. Healy, and J. Melville, "UMAP: Uniform manifold approximation and projection for dimension reduction," *arXiv preprint*, arXiv:1802.03426, 2018. [Online]. Available: https://arxiv.org/abs/1802.03426

[5] Z. Liu, Y. Zhang, Y. Li, and Chaomurilige, "Key news event detection and event context using graph convolution, clustering, and summarizing methods," *Applied Sciences*, vol. 13, no. 9, Art. no. 5510, 2023, doi: 10.3390/app13095510.

[6] D. Q. Nguyen and A. T. Nguyen, "PhoBERT: Pre-trained language models for Vietnamese," *in Findings of the Association for Computational Linguistics: EMNLP 2020*, 2020, pp. 1037–1042, doi: 10.18653/v1/2020.findings-emnlp.92.

[7] P. H. Nguyen and M. A. Nguyen, "Vietnamese text classification using long short-term memory and Word2Vec," *Informatics and Automation*, vol. 19, pp. 1255–1279, 2020, doi: 10.15622/ia.2020.19.6.5.

[8] B. H. Le, H. M. Nguyen, N. K.-P. Nguyen, and B. T. Nguyen, "A new approach for Vietnamese aspect-based sentiment analysis," *in Proc. 14th Int. Conf. on Knowledge and Systems Engineering (KSE)*, Nha Trang, Vietnam, 2022, pp. 1–6, doi: 10.1109/KSE56063.2022.9953759.

[9] M. Ester, H. P. Kriegel, J. Sander, and X. Xu, "A density-based algorithm for discovering clusters in large spatial databases with noise," *in Proc. 2nd Int. Conf. on Knowledge Discovery and Data Mining (KDD)*, 1996, pp. 226–231.

[10] M. F. Rahman, W. Liu, S. B. Suhaim, S. Thirumuruganathan, N. Zhang, and G. Das, "HDBSCAN: Density-based clustering over location-based services," *arXiv preprint*, arXiv:1602.03730, 2016, doi: 10.48550/arXiv.1602.03730.

[11] M. Grootendorst, "BERTopic: Neural topic modeling with a class-based TF-IDF procedure," *arXiv preprint*, arXiv:2203.05794, 2022, doi: 10.48550/arXiv.2203.05794.

[12] N. M. Gardazi, A. Daud, M. K. Malik, et al., "BERT applications in natural language processing: A review," *Artificial Intelligence Review*, vol. 58, Art. no. 166, 2025, doi: 10.1007/s10462-025-11162-5.
[13] J. Devlin, M. W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," *in Proc. NAACL-HLT*, 2019, pp. 4171–4186, doi: 10.18653/v1/N19-1423.
[14] Y. Liu et al., "RoBERTa: A robustly optimized BERT pretraining approach," *arXiv preprint*, arXiv:1907.11692, 2019, doi: 10.48550/arXiv.1907.11692.
[15] C. J. van Rijsbergen, *Information Retrieval*, 2nd ed. London, UK: Butterworths, 1979.
[16] S. Shilpa, K. P. Shailaja, S. G. Nischitha, H. N. Vishwanath, and K. T. Kondoju, "External clustering validation using ARI, NMI and FMI," *ITM Web Conference,* vol. 79, Art. no. 01004, 2025, doi: 10.1051/itmconf/20257901004.
[17] P. J. Rousseeuw, "Silhouettes: A graphical aid to the interpretation and validation of cluster analysis," *Journal of Computational and Applied Mathematics,* vol. 20, pp. 53–65, 1987, doi: 10.1016/0377-0427(87)90125-7.
[18] D. L. Davies and D. W. Bouldin, "A cluster separation measure," *IEEE Transactions on Pattern Analysis and Machine Intelligence,* vol. PAMI-1, no. 2, pp. 224–227, 1979, doi: 10.1109/TPAMI.1979.4766909.
[19] C. V. Nguyen et al., "Vistral-7B-Chat: Towards a state-of-the-art large language model for Vietnamese," 2023. [Online]. Available: https://huggingface.co/Viet-Mistral/Vistral-7B-Chat, Retrieved Jun 15, 2025
[20] J. Pennington, R. Socher, and C. D. Manning, "GloVe: Global vectors for word representation," *in Proc. EMNLP*, 2014, pp. 1532–1543, doi: 10.3115/v1/D14-1162.