Pattern Recognition Model for Security News Media Classification in Nigeria

[1]Ipeayeda Funmilola Wumi, [2]Fadiora Babatunde Olawale, [3]Lawrence Sunday Adeolu

[1]Faculty of Computing, Ajayi Crowther University, Oyo, Nigeria

[2]Department of Computer Studies, The Polytechnic, Ibadan, Nigeria

DOI: https://doi.org/10.51584/IJRIAS.2026.11090095

Received: 24 September 2026; Accepted: 29 September 2026; Published: 15 October 2026

ABSTRACT

Security incidents in Nigeria generate substantial volumes of news reports, but the unstructured nature of these reports makes systematic classification difficult. This study developed and evaluated a pattern-recognition model for automatically classifying Nigerian security-news articles into five categories: Insurgency/Terrorism, Banditry, Communal Conflict, Kidnapping, and Armed Robbery. News articles were collected from Punch, Daily Trust, and Premium Times for the period 1 January 2022 to 31 August 2025. A hybrid labelling pipeline combined article-based rules with temporal and geographic corroboration from Armed Conflict Location & Event Data (ACLED), followed by human validation. After corpus preparation, 13,896 articles were retained and 7,345 received accepted weak labels. A manually reviewed sample of 500 articles produced a human-gold dataset of 454 articles after removal of out-of-scope cases. The classification model used TF-IDF unigram and bigram features with a Linear Support Vector Machine (SVM). To reduce the effect of class imbalance, SMOTE was applied only to the training data. Hyperparameter selection on the validation set identified $C = 0.1$ as the best configuration. On the untouched 227-article human-gold test set, the model achieved 95.59% accuracy, 95.59% macro precision, 95.72% macro recall, and 95.57% macro-F1. All five classes recorded F1-scores above 93%, although some confusion occurred between semantically related categories such as Banditry, Kidnapping, and Armed Robbery. The SVM also exceeded Multinomial Naive Bayes and a majority-class baseline evaluated on the same test set. The findings demonstrate the feasibility of a locally adapted security-news classification pipeline, while the limited size of the human-gold test set and unequal representation of news sources indicate the need for larger, independently annotated and temporally external evaluations.

Keywords: security news classification, pattern recognition, natural language processing, TF-IDF, Support Vector Machine

INTRODUCTION

Nigeria experiences diverse forms of security incidents, including insurgency, banditry, communal conflict, kidnapping, and armed robbery. These incidents generate large quantities of media reports, but the reports are predominantly unstructured text. The resulting volume makes manual identification, organisation, and analysis of incident types difficult at scale. Automated text classification offers a means of converting unstructured news into structured categories that can support subsequent analysis.

Recent text-classification research has moved from lexical and statistical approaches toward contextual language models. Transformer-based systems such as BERT can capture contextual relationships that are difficult to represent using conventional bag-of-words and TF-IDF approaches (Devlin et al., 2019; Minaee et al., 2021). However, domain adaptation remains important because models trained on general-purpose corpora may not adequately represent specialised terminology and reporting patterns in a target domain (Gururangan et al., 2020). This issue is particularly relevant to Nigerian security reporting, where locations, actors, incident terminology, and reporting conventions are locally specific.

African NLP research has expanded the availability of regional language resources, including MasakhaNEWS,

but its news-topic categories are broader than the incident taxonomy required for security-event classification (Adelani et al., 2023). Other studies have examined Nigerian Pidgin sentiment classification and Nigerian cybersecurity incident classification, but these tasks differ from classifying open-domain Nigerian news articles into security-incident categories (Kniele & Beloucif, 2023; Onyekpeze et al., 2021).

A further challenge is the cost of producing large human-labelled datasets. Weak supervision provides a mechanism for generating training labels from imperfect rules and external knowledge while retaining human review for quality control (Ratner et al., 2017). In conflict research, ACLED provides structured event information that can be used as a corroborative signal. Mohammadi et al. (2023), for example, demonstrated the use of ACLED textual records in conflict-related classification, although their task was different from article-level classification of independently collected news.

Against this background, the present study developed a Nigerian security-news classification pipeline that combined article-based weak labelling, ACLED-supported temporal and geographic corroboration, human validation, TF-IDF feature extraction, and a Linear SVM classifier. The study evaluated the resulting model using an independently held-out human-gold test set and compared it with Multinomial Naive Bayes and a majority-class baseline.

RELATED LITERATURE

2.1 Automated News Classification and Pattern Recognition

Automated text classification assigns predefined labels to textual documents and has become an established natural language processing task. Conventional approaches based on TF-IDF remain attractive where datasets are domain-specific and computational resources are limited. Chavan and Padatare (2024) reported the practical use of TF-IDF with traditional classifiers for multi-category news classification, while broader reviews show that deep learning can provide richer contextual representations at the cost of greater computational and data requirements (Minaee et al., 2021). Among classical classifiers, Support Vector Machines remain widely used, although their performance can be affected by imbalanced datasets, irrelevant features, and hyperparameter selection. Olusola et al. (2025), for example, addressed these issues in a cassava leaf-disease classification task by using balanced class samples and hybridising two optimisation algorithms to tune SVM hyperparameters, obtaining 96.73% accuracy against 95.51% and 94.25% for the single-algorithm variants. Although that study concerned plant-leaf images rather than text, it illustrates the importance of controlling class imbalance and hyperparameter choice in SVM-based classification.

2.2 Domain Adaptation and Nigerian NLP

Domain-specific vocabulary can reduce the effectiveness of models trained on general-purpose data. Gururangan et al. (2020) showed that continued pre-training on domain-relevant text can improve downstream performance. African NLP studies similarly demonstrate the importance of locally relevant data and language resources (Adelani et al., 2023; Kniele & Beloucif, 2023). These findings motivate a locally constructed Nigerian security-news corpus rather than direct reliance on general news datasets.

2.3 Weak Supervision and Conflict Data

Weak supervision uses noisy or heuristic labels to reduce dependence on costly manual annotation (Ratner et al., 2017). ACLED provides structured conflict-event information such as event dates, locations, event types, actors, and fatalities. In this study, ACLED information was used as corroborative evidence rather than as article-level ground truth. This distinction was maintained because a structured event record and an independently written news article are not necessarily one-to-one representations of the same event.

2.4 Research Gap

The reviewed literature indicates progress in general news classification, African NLP, weak supervision, and conflict-event analysis, but these strands do not directly provide a validated five-class Nigerian security-news

ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue IX September 2026

classification system. The present study addresses this gap by integrating Nigerian news articles, rule-based weak labels, ACLED corroboration, human validation, and a conventional high-dimensional text classifier.

METHODOLOGY

3.1 Research Design and Data Sources

The study adopted a quantitative, experimental machine-learning design. The corpus consisted of Nigerian security-related news articles from Punch, Daily Trust, and Premium Times published between 1 January 2022 and 31 August 2025. Vanguard was not included in the final corpus because historical article discovery did not produce usable historical URLs. ACLED records for Nigerian conflict events were used to corroborate candidate labels.

Table 1: Distribution of study articles by news source

News source	Articles	Percentage
Punch	7,651	55.06%
Daily Trust	5,770	41.52%
Premium Times	475	3.42%
Total	13,896	100.00%

The five target categories were Insurgency/Terrorism, Banditry, Communal Conflict, Kidnapping, and Armed Robbery. The initial retained corpus contained 13,896 articles. Accepted weak labels were assigned to 7,345 articles, while 6,551 did not receive an accepted weak label.

Table 2: Weak-label coverage and label-source outcomes

Weak-label outcome	Articles	Percentage
Accepted weak label	7,345	52.86% of corpus
No accepted weak label	6,551	47.14% of corpus
Article rule + ACLED corroboration	2,632	35.83% of weak-labelled
Article-rule-only label	4,713	64.17% of weak-labelled

3.2 Labelling and Human Validation

Candidate articles were labelled using article-based rules and, where available, temporal and geographic corroboration from aggregated ACLED records. Articles that did not meet the required confidence conditions were not forced into a class. Label propagation was used for suitable unmatched articles based on textual similarity to confidently labelled instances.

A balanced review sample of 500 unique articles was manually examined. Forty-six articles were excluded as outside the five-class scope, 39 received corrected or newly assigned labels, and 415 retained their accepted weak labels. The resulting human-gold dataset contained 454 articles. Among 417 in-scope articles with comparable weak labels, 415 agreed with the human annotation, giving an observed agreement of 99.52%.

Table 3: Human-annotation outcomes for the 500-article review sample

Annotation outcome	Articles	Percentage
Accepted weak label	415	83.00%
Corrected/newly assigned label	39	7.80%
Excluded as out of scope	46	9.20%
Total reviewed	500	100.00%

ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue IX September 2026

3.3 Data Partitioning and Preprocessing

To reduce information leakage, reviewed articles and related duplicate groups were excluded from weak-label training. The remaining weak-labelled training corpus contained 6,888 articles. The human-gold dataset was divided into a 227-article validation set and a separate 227-article held-out test set. Text preprocessing included lowercasing, removal of URLs and non-lexical artefacts, tokenisation, stop-word removal using a customised list, and lemmatisation.

3.4 TF-IDF Representation and SVM

The processed texts were transformed into TF-IDF vectors using unigram and bigram features. The feature space was limited to 50,000 features, with a minimum document frequency of two and sub-linear term-frequency scaling. The vectoriser was fitted only on the training data and subsequently applied to validation and test data.

A one-versus-rest Linear Support Vector Classifier was used to learn decision boundaries between the five incident classes. Candidate regularisation values of C = 0.1, 1.0, and 10.0 were evaluated on the validation set. Macro-F1 was the principal selection criterion because it gives equal importance to each class. SMOTE was applied only to the training vectors to address class imbalance. Validation-based selection of C and the control of class imbalance are consistent with evidence that SVM-based classifiers are sensitive to imbalanced data and hyperparameter settings (Olusola et al., 2025).

Table 4: Validation results for SVM hyperparameter selection

C	Validation accuracy	Balanced accuracy	Macro-F1
0.1	96.48%	96.51%	96.44%
1.0	95.15%	95.26%	95.15%
10.0	94.27%	94.39%	94.36%

The selected configuration was C = 0.1.

3.5 Evaluation

Final performance was measured on the untouched human-gold test set using accuracy, balanced accuracy, macro precision, macro recall, macro-F1, weighted F1, and confusion-matrix analysis. A 1,000-iteration percentile bootstrap was additionally used to estimate uncertainty for accuracy and macro-F1. Multinomial Naive Bayes and a majority-class classifier were evaluated on the same test set as internal baselines.

RESULTS AND DISCUSSION

Table 5: Corpus and Label Distribution

Target class	Weak-labelled articles	Percentage
Insurgency/Terrorism	1,435	19.54%
Banditry	945	12.87%
Communal Conflict	171	2.33%
Kidnapping	4,435	60.38%
Armed Robbery	359	4.89%
Total	7,345	100.00%

The weak-labelled corpus was substantially imbalanced, with Kidnapping accounting for 60.38% of accepted labels. This imbalance justified the use of training-only SMOTE and macro-averaged evaluation measures.

ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue IX September 2026

Table 6: Overall SVM Performance

Metric	Score
Accuracy	95.59%
Balanced accuracy	95.72%
Macro precision	95.59%
Macro recall	95.72%
Macro-F1	95.57%
Weighted F1	95.58%

On the held-out 227-article human-gold test set, the selected SVM correctly classified 217 articles and misclassified 10. Accuracy was 95.59%, while macro-F1 was 95.57%. The close relationship among accuracy, balanced accuracy, macro-F1, and weighted F1 indicates that the performance was not dominated by a single class in the balanced test sample.

Table 7: Class-Level Performance

Class	Support	Precision	Recall	F1-score
Insurgency/Terrorism	43	0.9348	1.0000	0.9663
Banditry	43	0.9302	0.9302	0.9302
Communal Conflict	48	1.0000	0.9375	0.9677
Kidnapping	49	0.9783	0.9184	0.9474
Armed Robbery	44	0.9362	1.0000	0.9670

All five classes recorded F1-scores above 0.93. Insurgency/Terrorism and Armed Robbery achieved recall of 1.00 in the test set. Banditry had the lowest F1-score at 0.9302, while Kidnapping had the lowest recall at 0.9184. These values indicate generally consistent class discrimination within the pilot test, while also showing that some incident categories remain difficult to distinguish.

4.4 Confusion Matrix and Error Analysis

All 43 Insurgency/Terrorism and all 44 Armed Robbery articles were correctly classified. Three Banditry articles were misclassified: two as Insurgency/Terrorism and one as Armed Robbery. Three Communal Conflict articles were misclassified, one each as Insurgency/Terrorism, Banditry, and Kidnapping. Four Kidnapping articles were misclassified, two as Banditry and two as Armed Robbery.

The observed errors are consistent with overlap in the language used to describe related security incidents. Banditry, kidnapping, and armed robbery reports may share vocabulary concerning armed attacks, raids, abductions, weapons, and perpetrators. The error pattern therefore suggests that further gains may require richer contextual representations or more detailed event-level annotation.

Table 8: Uncertainty and Baseline Comparison

Metric	Point estimate	95% bootstrap percentile interval
Accuracy	95.59%	92.51%–97.81%
Macro-F1	95.57%	92.48%–98.01%

The bootstrap intervals show that the point estimates have non-trivial sampling variation because the held-out test set contains only 227 articles. The intervals describe uncertainty from resampling the observed predictions and do not account for selection bias, annotation subjectivity, unequal source representation, or future changes in Nigerian security terminology.

Table 9: Comparison of the SVM with baseline classifiers on the held-out test set

Approach	Accuracy	Balanced accuracy	Macro-F1
TF-IDF + Linear SVM	95.59%	95.72%	95.57%
Multinomial Naive Bayes	37.44%	36.74%	28.90%
Majority class (Kidnapping)	21.59%	20.00%	7.10%

On the same held-out test set, the SVM produced higher accuracy, balanced accuracy, and macro-F1 than the two internal baselines. The majority-class classifier failed to identify four of the five categories, while Multinomial Naive Bayes did not retrieve Communal Conflict or Armed Robbery examples in this test sample. These results support the suitability of the evaluated Linear SVM configuration for the TF-IDF representation used in this study. They should not, however, be interpreted as a universal ranking of classifiers because the comparison is specific to this dataset, feature representation, and test protocol.

IMPLICATIONS AND LIMITATIONS

The study demonstrates a practical route for converting a large Nigerian security-news corpus into a structured incident-classification resource. The combination of conservative weak labelling, external corroboration, human validation, and leakage-controlled evaluation provides a framework that can be extended to larger corpora. The results also show that a conventional linear classifier can achieve strong performance when paired with an appropriate domain-specific dataset and TF-IDF representation.

Several limitations constrain generalisation. First, the final human-gold test set contained only 227 articles. Second, the manually reviewed sample was deliberately balanced and therefore does not represent the natural distribution of the complete corpus. Third, the three retained news outlets were unequally represented, with Punch and Daily Trust accounting for most articles. Fourth, the study focused on English-language reporting and five predefined incident classes. Finally, security terminology and reporting practices can change over time, creating potential concept drift.

Accordingly, the reported 95.59% accuracy should be interpreted as performance on the defined held-out pilot test rather than as a universal estimate of performance on all Nigerian security news.

CONCLUSION

This study developed and evaluated a pattern-recognition pipeline for classifying Nigerian security-news articles into Insurgency/Terrorism, Banditry, Communal Conflict, Kidnapping, and Armed Robbery. The pipeline integrated weak supervision, ACLED-supported corroboration, human validation, TF-IDF feature extraction, SMOTE-based training-set balancing, and a one-versus-rest Linear SVM. On the held-out human-gold test set, the model achieved 95.59% accuracy and 95.57% macro-F1, with all five classes recording F1-scores above 93%. The model also exceeded the Multinomial Naive Bayes and majority-class baselines under the same evaluation conditions.

The main contribution is therefore a locally focused and empirically evaluated framework for Nigerian security-news classification. Future work should enlarge the independently annotated corpus, include more news sources, employ multiple annotators and inter-annotator agreement measures, introduce an uncertain/out-of-scope class, and evaluate the model on chronologically later and externally sourced articles. Such extensions would provide stronger evidence about the model's robustness and generalisability.

REFERENCES

1. Adelani, D. I., Masiak, M., Azime, I. A., Alabi, J. O., Tonja, A. L., Mwase, C., Ogundepo, O., Dossou, B. F. P., Oladipo, A., Nixdorf, D., Emezue, C. C., Al-Azzawi, S. S., Sibanda, B. K., David, D., Ndolela, L., Mukiibi, J., Ajayi, T. O., Ngoli, T. M., Odhiambo, B., … Stenetorp, P. (2023). MasakhaNEWS: News

topic classification for African languages. In Proceedings of the 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (Vol. 1, pp. 144–159). Association for Computational Linguistics. https://doi.org/10.18653/v1/2023.ijcnlp-main.10

2. Allam, H., Makubvure, L., Gyamfi, B., Graham, K. N., & Akinwolere, K. (2025). Text classification: How machine learning is revolutionizing text categorization. Information, 16(2), 130. https://doi.org/10.3390/info16020130
3. Armed Conflict Location & Event Data Project. (2025). Nigeria: Conflict events [Data set]. Humanitarian Data Exchange. https://data.humdata.org/dataset/nigeria-acled-conflict-data
4. Atkinson, M., Piskorski, J., Tanev, H., & Zavarella, V. (2017). On the creation of a security-related event corpus. In Proceedings of the Events and Stories in the News Workshop (pp. 59–65). Association for Computational Linguistics. https://doi.org/10.18653/v1/W17-2709
5. Chavan, D., & Padatare, S. (2024). Explainable AI for news classification. International Journal for Research in Applied Science and Engineering Technology, 12(11), 2400–2408. https://doi.org/10.22214/ijraset.2024.65670
6. Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In Proceedings of NAACL-HLT 2019 (Vol. 1, pp. 4171–4186). Association for Computational Linguistics. https://doi.org/10.18653/v1/N19-1423
7. Gururangan, S., Marasović, A., Swayamdipta, S., Lo, K., Beltagy, I., Downey, D., & Smith, N. A. (2020). Don't stop pretraining: Adapt language models to domains and tasks. In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (pp. 8342–8360). Association for Computational Linguistics. https://doi.org/10.18653/v1/2020.acl-main.740
8. Henning, S., Beluch, W., Fraser, A., & Friedrich, A. (2023). A survey of methods for addressing class imbalance in deep-learning based natural language processing. In Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics (pp. 523–540). Association for Computational Linguistics. https://doi.org/10.18653/v1/2023.eacl-main.38
9. Kniele, A., & Beloucif, M. (2023). Uppsala University at SemEval-2023 Task 12: Zero-shot sentiment classification for Nigerian Pidgin tweets. In Proceedings of the 17th International Workshop on Semantic Evaluation (pp. 1491–1497). Association for Computational Linguistics. https://doi.org/10.18653/v1/2023.semeval-1.205
10. Minaee, S., Kalchbrenner, N., Cambria, E., Nikzad, N., Chenaghlu, M., & Gao, J. (2021). Deep learning-based text classification: A comprehensive review. ACM Computing Surveys, 54(3), Article 62. https://doi.org/10.1145/3439726
11. Mohammadi, H., Momand, Z., Habibi, P., Ramaki, N., Fazli, B. S., Rohany, S. Z., & Samsoor, I. (2023). Analyzing textual data for fatality classification in Afghanistan's armed conflicts: A BERT approach [Preprint]. arXiv. https://doi.org/10.48550/arXiv.2310.08653
12. Olusola, B. A., Oyediran, M., Ipeayeda, F., & Akindele, A. A. (2025). Enhance Support Vector Machine models for Cassava leaf disease variants identification and classification. Journal of Applied Science Information and Computing, 6(1), 87–104. https://doi.org/10.59568/JASIC-2025-6-1-10
13. Onyekpeze, O., Owolabi, O., & Ibrahim, B. H. (2021). Classification of cybersecurity incidents in Nigeria using machine learning methods. Covenant Journal of Informatics and Communication Technology, 9(2), 1–17. https://journals.covenantuniversity.edu.ng/index.php/cjict/article/view/2796
14. Ratner, A., Bach, S. H., Ehrenberg, H., Fries, J., Wu, S., & Ré, C. (2017). Snorkel: Rapid training data creation with weak supervision. Proceedings of the VLDB Endowment, 11(3), 269–282. https://doi.org/10.14778/3157794.3157797