SSR Journal of Artificial Intelligence (SSRJAI)

OPEN ACCESS

Homepage: https://ssrpublisher.com/ssrjai/

Volume 3, Issue 8, 2026 ISSN: 3049-0413

Fine-Grained Propaganda Detection for Hausa: Systematic Review of Resources, Methods, and Future Directions

Abdulhamid Abubakar[1], Gilbert I. O. Aimufua[2], Muhammed A. Abubakar[3]

[1]Centre for Cyberspace Studies, Nasarawa State University Keffi, Nigeria

Received: *05.07.2026* | ***Accepted:*** *05.08.2026* | ***Published:*** *08.08.2026*

***Corresponding author:** Abdulhamid Abubakar

DOI: *10.5281/zenodo.21848915*

Abstract **Original Research**

The proliferation of propaganda through digital platforms threatens information integrity globally. Computational tools for fine-grained persuasion technique detection remain concentrated on English and a few high-resource European languages. Hausa, a Chadic language spoken by over 120 million people across West and Central Africa, lacks systems capable of identifying specific persuasion strategies like loaded language or appeal to authority. This systematic review, conducted following PRISMA 2020 guidelines, synthesises research at the intersection of fine-grained propaganda detection, parameter-efficient transfer learning, cross-lingual knowledge transfer, and African natural language processing (NLP). Thirty-seven studies (2019-2026) were retrieved from IEEE Xplore, ScienceDirect, Scopus, Web of Science, and the ACL Anthology based on strict inclusion criteria. The evidence reveals three primary patterns: (a) fine-grained propaganda detection features mature taxonomies and transformer-based methods, but its empirical base excludes African languages; (b) parameter-efficient methods like Low-Rank Adaptation LoRA match full fine-tuning performance but remain unevaluated for span-level technique detection; and (c) Hausa NLP has largely advanced through binary and coarse classification, leaving a critical gap in fine-grained analysis. A coverage analysis confirms that no existing study addresses fine-grained propaganda detection for Hausa. This review consolidates the evidence in nine tables and seven figures, maps the resource landscape, profiles adaptation trade-offs, and proposes a five-year research roadmap. The roadmap prioritises adapter benchmarking and corpus development (2026-2028), cross-lingual transfer with cultural adaptation (2027-2029), and deployable, resource-efficient systems for African fact-checking organizations (2029-2031). The findings demonstrate that while the technical building blocks for fine-grained propaganda detection in Hausa exist, the intersection remains completely unexplored, presenting a significant opportunity for future research.

Keywords: Fine-grained propaganda detection; cross-lingual transfer learning; parameter-efficient fine-tuning; Hausa; low-resource languages; African language NLP; systematic literature review.

Copyright © 2026 The Author(s). This is an open-access article distributed under the terms of the Creative Commons Attribution-NonCommercial 4.0 International License (CC BY-NC 4.0).

1. Introduction

Propaganda is the deliberate and systematic use of persuasive communication to shape attitudes, beliefs, and behaviour in support of a particular political, ideological, or social agenda. Unlike misinformation, propaganda does not necessarily rely on falsehood; instead, it operates through rhetorical strategies such as emotional appeals, loaded language, appeals to authority, and false dilemmas (Piskorski et al., 2023). The proliferation of digital communication has amplified the reach of these strategies, enabling persuasive content to circulate rapidly across social media, online news portals, and messaging applications (Joaquim, 2025). With social media now serving as a primary channel for political information for a substantial proportion of internet users, the automated identification of manipulative content has emerged as a recognised priority within natural language processing (NLP), computational social science, and information systems research (Ausat, 2023; Saaida & Alhouseini, 2023).

The computational response to this priority has progressed along two interrelated directions. The first involves a shift from coarse binary classification which separates propagandistic from non-propagandistic text toward fine-grained technique detection, which identifies the specific persuasion strategy employed and the textual span in which it occurs (Da San Martino et al., 2019). The second entails the extension of detection capabilities beyond English into multilingual settings, exemplified by the SemEval shared tasks that have become the field's reference benchmarks (Piskorski et al., 2023). However, the benefits of these advances are unevenly distributed. The overwhelming majority of detection systems target English and a handful of European languages, with smaller efforts for Arabic, Chinese, and Russian (Refaee et al., 2022; Zhang & Zhang, 2022), while African languages despite accounting for approximately one-third of the world's languages are almost entirely absent from the fine-grained propaganda detection literature.

Hausa serves as a compelling focal point for examining this absence. As a Chadic language of the Afro-Asiatic family, it is spoken by more than 120 million first-language speakers and tens of millions of second-language speakers across Nigeria, Niger, Chad, Cameroon, and Ghana, placing it among the most widely spoken languages on the African continent and globally (Muhammad et al., 2025). It functions as a lingua franca in commerce, education, religion, and journalism, and carries a substantial media footprint through international broadcasters and domestic outlets. From a computational standpoint, however, Hausa remains under-resourced, with annotated datasets, domain-specific corpora, and high-performing models far less developed than for high-resource languages (Inuwa-Dutse, 2025).

Two recent developments make a systematic synthesis particularly timely. First, parameter-efficient fine-tuning methods now allow large models to be adapted by updating only a small fraction of their parameters a critical consideration for the computationally constrained settings in which African fact-checking organisations operate (Jibrin et al., 2026). Second, cross-lingual transfer mechanisms such as adapter composition enable knowledge from high-resource source languages to bootstrap capabilities in low-resource targets (Pfeiffer et al., 2020). Despite growth in research on Hausa NLP, cross-lingual learning, and propaganda detection, these areas have developed in isolation, and no study has investigated how parameter-efficient transfer learning can detect fine-grained propaganda techniques in Hausa.

This systematic review therefore addresses four interrelated research questions:

i. What are the predominant methodological approaches to fine-grained propaganda detection, and how have they been evaluated?

ii. To what extent do parameter-efficient and cross-lingual transfer methods extend such systems to low-resource languages, and what are the associated efficiency and performance trade-offs?

iii. What is the empirical state of Hausa and African language NLP, and where does propaganda detection sit within it?

iv. What limitations and future research directions are identifiable within the current body of work?

To address these questions, the specific objectives of this review are:

i. To systematically synthesise advances in propaganda detection, parameter-efficient learning, and cross-lingual transfer relevant to low-resource languages over the period 2019 to 2026;

ii. To benchmark parameter-efficient adapter method classes and multilingual model families against efficiency, performance, and deployment dimensions based on reported evidence;

iii. To map the empirical state and resource landscape of Hausa and African language NLP, and to precisely locate the fine-grained propaganda detection gap;

iv. To develop a forward-looking research roadmap synthesised from the gaps and recommendations identified across the reviewed literature.

The remainder of this paper is organised as follows. Section 2 details the PRISMA-compliant methodology. Section 3 presents the theoretical properties and comparative evaluation of adaptation methods. Section 4 presents the application-oriented synthesis, including coverage analysis, resource landscape, and research roadmap. Section 5 discusses the findings within the broader research discourse, and Section 6 concludes with limitations and recommendations.

2. Methodology

The review follows the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) 2020 guidelines (Page et al., 2021) to ensure transparent and reproducible reporting. PRISMA 2020 was selected for its updated treatment of modern search complexities, including preprints and grey literature, which are prevalent in fast-moving NLP research. The process comprised four phases, identification, screening, eligibility assessment, and inclusion as set out in the flow diagram in Figure 1.

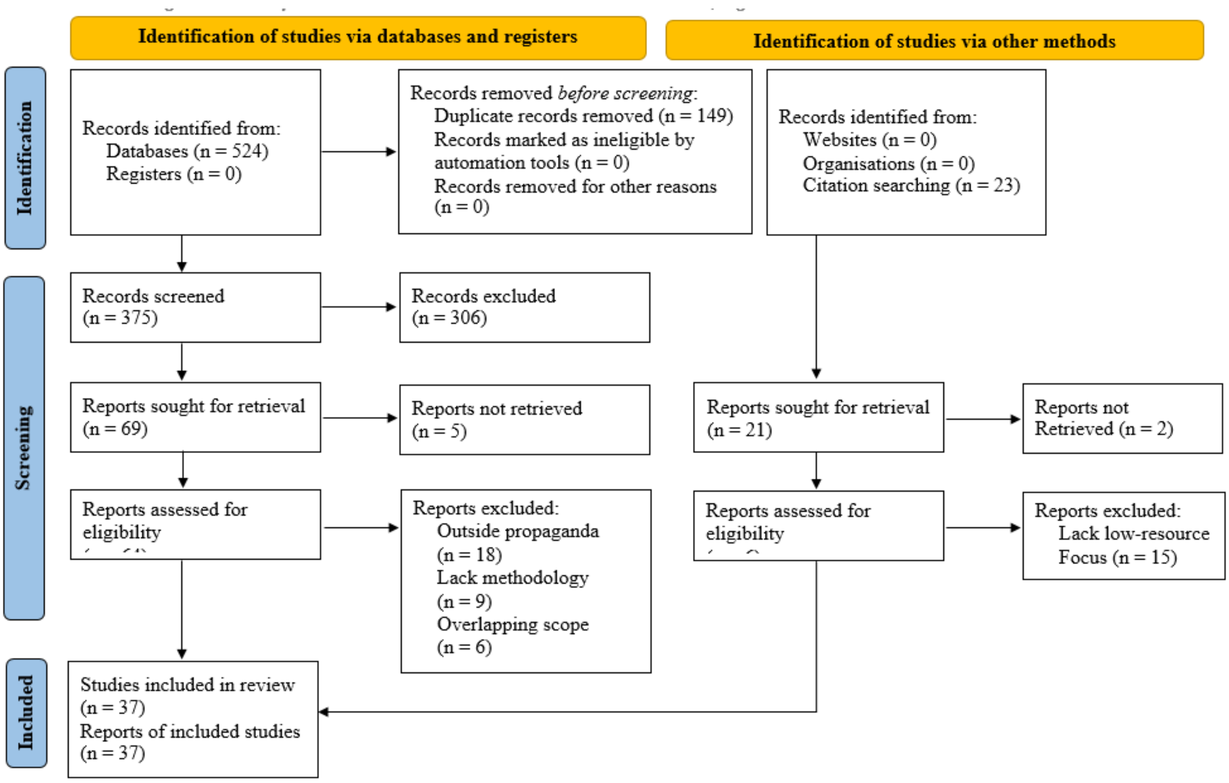

Figure 1. PRISMA 2020 Flow Diagram of the Study Selection Process.

2.1 Search Strategy and Data Collection

A structured search was conducted across IEEE Xplore, ScienceDirect, SpringerLink, Scopus, Web of Science, and the ACM Digital Library, supplemented by the ACL Anthology, arXiv, and the HausaNLP Catalog to capture domain-specific and emerging work. The search covered English-language publications from 2019 to 2026, the period during which fine-grained propaganda detection, parameter-efficient methods, and Hausa NLP matured. A Boolean strategy combined three thematic groups. The propaganda group included "propaganda detection," "persuasion techniques," "fine-grained," and "span-level." The methods group included "parameter-efficient," "adapter," "LoRA," "cross-lingual transfer," and "zero-shot." The language group included "Hausa," "African languages," "low-resource," and the names of relevant models such as AfroXLMR and AfriBERTa. Records were exported to Zotero reference manager for deduplication prior to screening.

2.2 Inclusion and Exclusion Criteria

For inclusion, a study had to meet the following criteria: (a) be a peer-reviewed journal article, conference paper, or rigorous preprint published between 2019 and 2026; (b) report empirical, theoretical, or review research bearing directly on propaganda detection, parameter-efficient or cross-lingual transfer learning, or Hausa and African language NLP; and (c) provide an explicit description of its methods together with reported performance metrics or a substantive methodological contribution. Exclusion criteria removed studies outside these domains, opinion pieces and white papers without empirical or methodological substance, studies lacking a clear model or validation description, duplicate records, and non-English publications.

2.3 Screening Process

Title and abstract screening was performed against the predefined criteria, followed by full-text assessment of the records sought for retrieval. As shown in Figure 1, the database search yielded 524 records, of which 149 duplicates were removed, leaving 375 for screening. Title and abstract screening excluded 306 records, leaving 69 sought for retrieval, of which 5 could not be retrieved and 64 were assessed for eligibility. At the full-text stage, 33 reports were excluded: 18 for falling outside the propaganda or low-resource focus, 9 for lacking methodology or metrics, and 6 for overlapping scope. Citation searching identified 23 further records, of which 21 were assessed and 15 excluded, yielding 6 additional reports. In total, 37 studies met all criteria and were included in the synthesis.

2.4 Data Extraction and Analysis

A standardised extraction form was piloted on a subset of studies and refined for consistency. The following were extracted from each study: identification details; task or focus; data sources and dataset size; model and method type, including base model family and adaptation strategy; reported performance metrics such as accuracy, F1, precision, and recall; efficiency indicators such as trainable parameter share where reported; and author-acknowledged limitations and gaps. Because the included studies report heterogeneous metrics across different tasks, quantitative findings are presented as reported rather than pooled, and comparative judgements about method classes are expressed on normalised qualitative scales to enable cross-study comparison.

To support the comparative evaluation in Section 4, an Adaptation Suitability Score (ASS) was defined for method classes as a weighted combination of normalised dimensions, given in Equation (1):

$$ASS = w^1 \cdot PE + w^2 \cdot P + w^3 \cdot M + w^4 \cdot D \qquad 1.1$$

Where PE denotes parameter efficiency, P denotes reported performance, M denotes memory efficiency, and D denotes deployment readiness, with weights that sum to one and can be adjusted for context. For resource-constrained deployment, the analysis weights parameter efficiency and deployment readiness more heavily ($w_1 = 0.30$, $w_2 = 0.30$, $w_3 = 0.20$, $w_4 =$

0.20). The scores that populate this scheme are qualitative syntheses of the reviewed evidence rather than measurements from a single experiment, and they are presented to structure comparison, not to substitute for primary benchmarking.

3. Theoretical Results

Table 1 summarises the mechanism, trainable parameter share, and principal strength and limitation of each method class. Table 2 reports the comparative scores used in the qualitative profile.

Table 1: Theoretical Properties of Parameter-efficient Adaptation Method Classes

Method	Mechanism	Trainable Parameters	Key Strength	Key Limitation
LoRA	Low-rank weight-update matrices	0.1-1%	Explicit rank control and budget	Requires rank and placement tuning
Bottleneck Adapter	Down- and up-projection modules	1.0-5%	Original, well-supported method	More parameters than LoRA
Prefix Tuning	Learnable input prefixes per layer	0.1-1%	No change to core weights	Sensitive to prefix length
Compacter	Low-rank plus hypercomplex layers	< 0.1%	Extreme parameter efficiency	Implementation complexity
$(IA)^3$	Learned activation scaling vectors	< 0.01%	Smallest trainable budget	Lower performance ceiling
AdapterFusion	Attention over task adapters	Variable	Compositional reuse of adapters	Needs multiple trained adapters
Full Fine-Tuning	Update all model weights	100%	Performance ceiling reference	Costly; one copy per task

Table 2: Comparative Evaluation of Selected Adapter Methods on Normalised Qualitative Dimensions

Method	Parameter Efficiency	Memory Efficiency	Training Speed	Reported Performance	Implementation Maturity
LoRA	0.85	0.80	0.80	0.90	0.95
Bottleneck Adapter	0.70	0.70	0.75	0.85	0.90
Prefix Tuning	0.75	0.75	0.70	0.75	0.70
$(IA)^3$	0.98	0.90	0.85	0.70	0.55

3.1 Efficiency, Performance, and Deployment Trade-off

A consistent pattern emerges across Tables 1 and 2 and is visualised in Figure 2, which plots trainable parameter share against expected performance on logarithmic axes. Four methods lie on the Pareto frontier, the set that is not dominated on both criteria: (IA)³ and Compacter at very small parameter budgets, LoRA at an intermediate budget, and full fine-tuning at the performance ceiling. Bottleneck adapters and prefix tuning sit just inside the frontier, because LoRA reaches comparable or higher performance with fewer trainable parameters. For the resource-constrained settings that characterise African fact-checking organisations, the preferred region combines high efficiency with strong performance, which places LoRA and bottleneck adapters as the most suitable starting points and reserves the lighter methods such as (IA)³ and Compacter for the most constrained hardware. This conceptual trade-off should be confirmed empirically on the span-level task, since none of the reviewed studies has evaluated these methods for fine-grained propaganda detection.

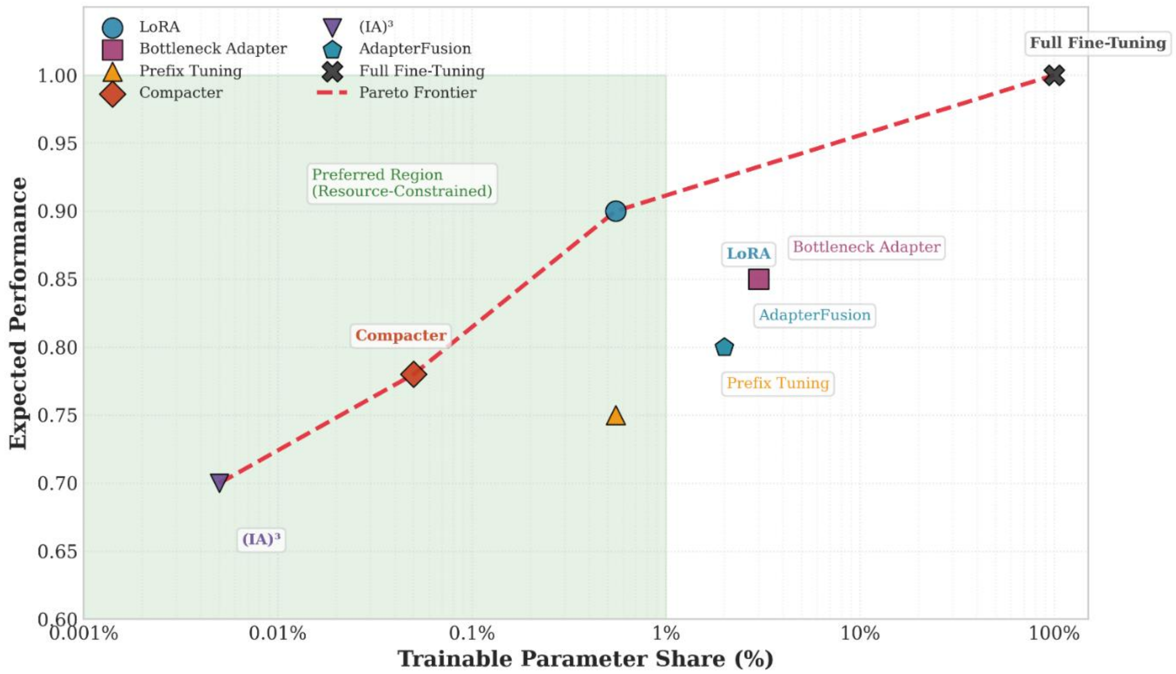

Figure 2: Efficiency Frontier of Parameter-Efficient Adaptation Methods

4. Application and Results Presentation

4.1 Reported Performance across Hausa and Related NLP Tasks

Table 3 collates the headline results reported for Hausa and closely related low-resource tasks. The metrics are heterogeneous and are therefore presented as reported rather than pooled. Two observations follow. First, strong results are achievable for Hausa classification, with reported accuracies ranging from approximately 84% for news classification (Du, 2025) to approximately 99% for machine-generated text detection (Sani et al., 2025). Second, every reported result concerns a binary or coarse classification objective; the fine-grained propaganda row is empty because no study has attempted the span-level task, which confirms the central gap of this review.

Table 3: Reported Performance across Hausa and related Low-resource NLP Tasks.

Task	Representative Study	Method or Model	Reported Metric	Note
Fake News Detection	Imam et al. (2022)	SVM and traditional ML	85% accuracy	Binary; 2,600 articles
Fake News Detection	Abdullahi Ibrahim et al. (2024)	Transfer learning	Significant improvement	6,600+ articles
News Classification	Du (2025)	Cross-lingual plus augmentation	84.1% / 83.4% macro-F1	Hierarchical fine-tuning
Machine-Generated Text	Sani et al. (2025)	AfroXLMR	99.23% accuracy	2,586 articles; binary
Sentiment Analysis	Jibrin et al. (2026)	LoRA-enhanced transformers	Competitive; 1.06% params	Document level
Cyberhate (Punjabi)	Kaur & Saini (2024)	Transliteration plus BERT	Precision 0.86 recall 0.83	Low-resource comparator
Fine-Grained Propaganda	None	Not yet attempted	Not applicable	Research gap

4.2 Coverage of Modelling Approaches across Tasks

To precisely locate the gap, the reviewed Hausa and African NLP studies were mapped onto a matrix of NLP tasks against modelling approaches. Table 4 records the count of reviewed studies for each combination, and Figure 3 shows the same evidence as a segmented bar chart, in which each task bar is divided into the approaches that address it. The pattern is clear: multilingual transformers and Afri-centric models cover most established tasks, traditional machine learning persists for the earliest fake news and offensive language work, and parameter-efficient methods appear so far only for sentiment analysis. The fine-grained propaganda task has no studies under any approach, leaving its bar empty the visual signature of an unaddressed problem.

Table 4: Coverage of Reviewed Hausa and African NLP studies

NLP Task	Traditional ML	Multilingual Transformer	Afri-Centric Model	Parameter-Efficient	Cross-Lingual Transfer
Fake News Detection	1	1	1	0	1
Sentiment Analysis	0	1	1	1	0
Offensive Language	1	1	1	0	0
News Classification	0	1	1	0	1

Machine-Generated Text	0	1	1	0	0
Named Entity Recognition	0	1	1	0	1
Fine-Grained Propaganda	0	0	0	0	0

Note: Counts reflect the reviewed corpus and are indicative of relative coverage rather than an exhaustive census of all Hausa NLP work.

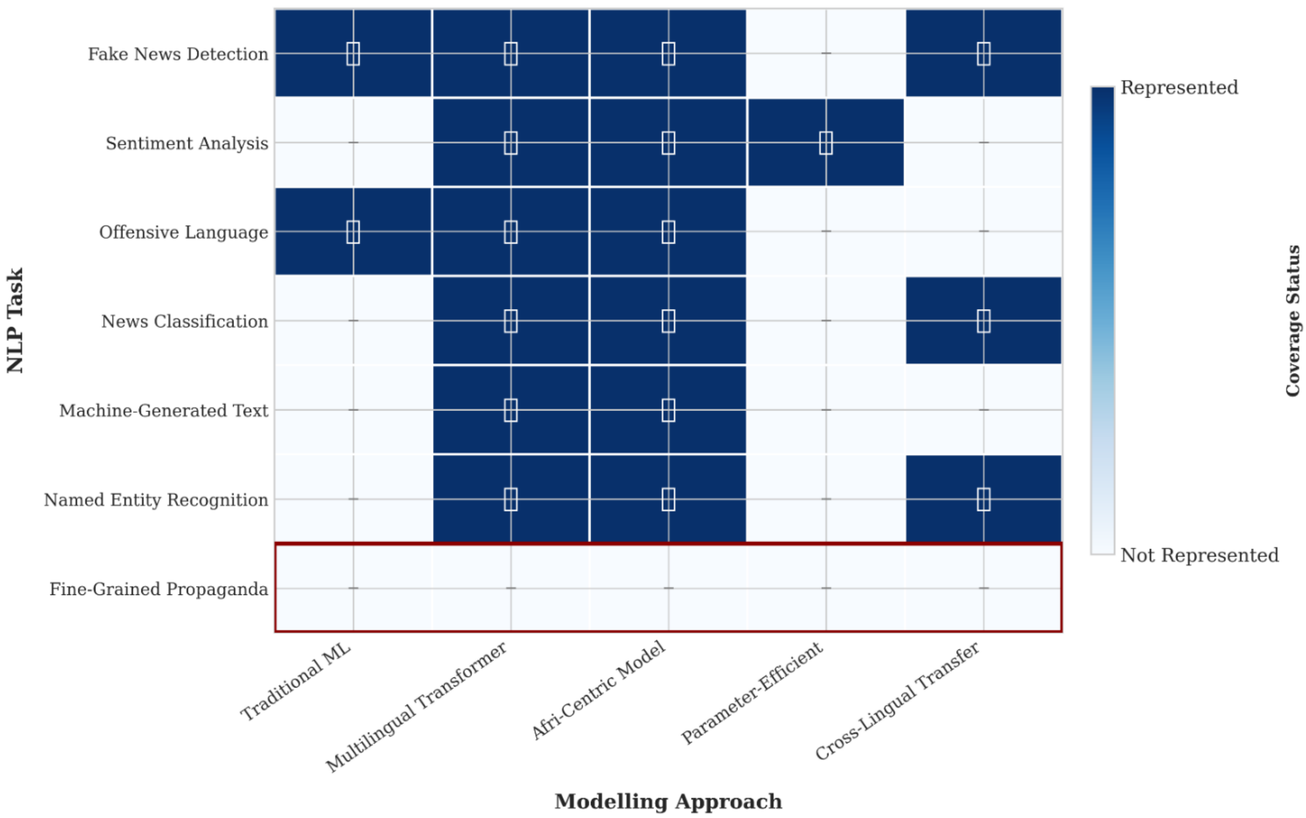

Figure 3: Heatmap Coverage of Modelling Approaches across NLP Tasks

4.3 Comparative Profile of Adaptation Methods

Figure 4 presents the comparative scores from Table 4 as a parallel-coordinates plot across five dimensions, with each method drawn as a path connecting its score on every axis. LoRA traces the most balanced path, combining high implementation maturity and reported performance with strong efficiency, consistent with its adoption in the one Hausa demonstration to date (Jibrin et al., 2026). The (IA)³ method ranks highest on parameter and memory efficiency but falls lowest on performance and maturity, producing a path that crosses the others. Bottleneck adapters and prefix tuning fall between these. For span-level propaganda detection, the profile suggests beginning with LoRA and bottleneck adapters and treating the lighter methods as efficiency-oriented comparisons rather than default choices.

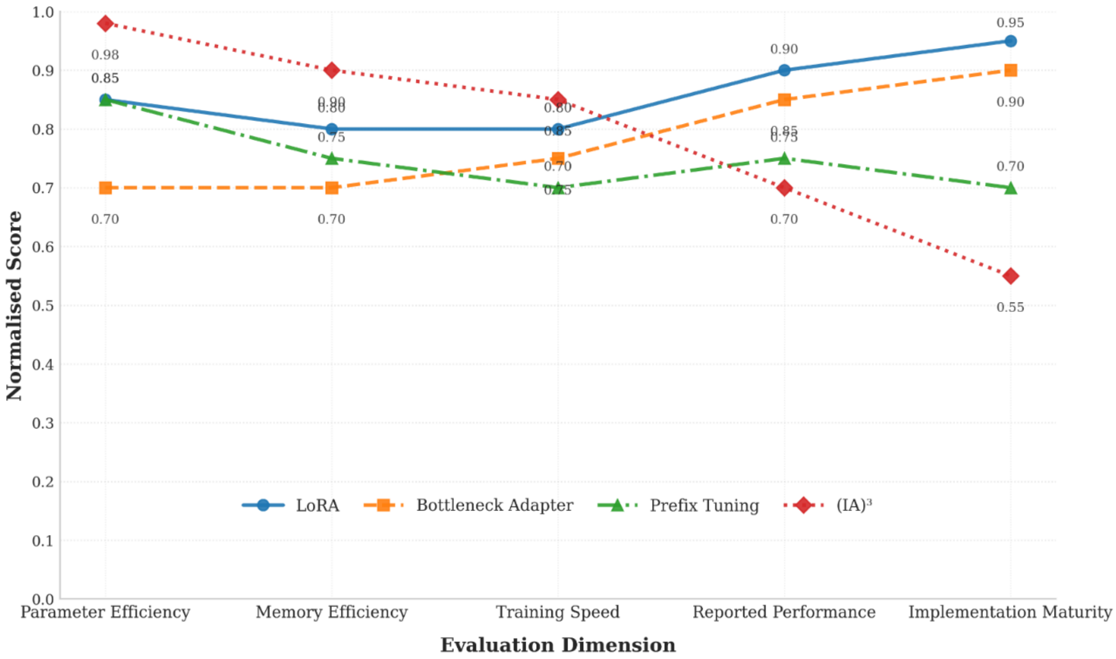

Figure 4. Parallel-coordinates Comparison of Parameter-efficient Adaptation Methods across Five Deployment-relevant Dimensions.

4.4 Evolution of Hausa NLP Capabilities

Table 5 trace the emergence of Hausa and African NLP capabilities relevant to propaganda detection. Named entity recognition, sentiment analysis, and machine translation became established from 2021 to 2022; misinformation-adjacent tasks such as fake news and offensive language detection followed from 2022 to 2023; machine-generated text detection and the first parameter-efficient Hausa work are recent, from 2025 to 2026. Fine-grained propaganda detection has no established presence and is positioned as a proposed direction. The timeline shows that the building blocks, Afri-centric models, transfer methods, and annotation practice are now in place, which makes the propaganda task tractable rather than premature.

Table 5: Evolution Timeline of Hausa and African NLP Capabilities Relevant to Propaganda Detection.

Capability or Task	Period	Representative Studies	Status
Named Entity Recognition	2021-present	African NER resources and benchmarks	Established
Sentiment Analysis	2022-present	AfriSenti; LoRA-based sentiment	Established
Machine Translation	2022-present	Hausa Visual Genome; NLLB-class models	Established
Fake News Detection	2022-present	Imam et al.; Abdullahi Ibrahim et al.	Established (binary)
Offensive Language	2023-present	Adam et al.	Established

Machine-Generated Text	2025-present	Sani et al.	Emerging
Parameter-Efficient Methods	2026-present	Jibrin et al.	Emerging
Fine-Grained Propaganda	2026 onward (proposed)	Target of this review	Not yet addressed

4.5 Resource Landscape of Hausa Datasets

Table 6 profile the principal Hausa-language datasets in the reviewed corpus by size and label granularity. The datasets show steady growth in availability but limited label granularity: most carry binary or few-class labels, and the largest resource is a multimodal translation set rather than a fine-grained classification corpus. The proposed HausaProp corpus is positioned to fill the granularity gap, pairing a moderate segment count with a twenty-technique span-level label scheme. Figure 6 ranks the datasets by size and annotates each with its label granularity, making clear that no existing Hausa resource approaches the granularity required for fine-grained propaganda detection.

Table 6: Resource profile of reviewed Hausa-language datasets

Dataset or Study	Year	Approx. Size	Task	Label Granularity	Availability
Imam et al.	2022	2,600 articles	Fake News	Binary	Public
Hausa Visual Genome	2022	32,923 pairs	Multimodal MT	Not applicable	Public
AfriSenti (Hausa subset)	2023	~30,000 tweets	Sentiment	3 classes	Public
Abdullahi Ibrahim et al.	2024	6,600+ articles	Fake News	Binary	Reported
Sani et al.	2025	2,586 articles	Authenticity	Binary	Public
HausaProp (proposed)	2026	5,000-8,000 segments	Propaganda	20 techniques	Planned open release

4.6 Five-Year Research Roadmap

Table 7 and Figure 5 present a five-year roadmap synthesised from the gaps and recommendations identified across the reviewed literature. The roadmap is not a prediction but an aggregation of directions the literature implies. The first phase, from 2026 to 2028, prioritises benchmarking parameter-efficient adapters on an established English propaganda dataset and developing the HausaProp corpus with a culturally adapted twenty-technique scheme and a target inter-annotator agreement above 0.65. The second phase, from 2027 to 2029, tests cross-lingual transfer through MAD-X composition and few-shot adaptation and documents culturally specific persuasion through a Hausa propaganda lexicon. The third phase, from 2029 to 2031, optimises the best configuration for inference on modest hardware and evaluates it with fact-checking partners.

Table 7: Five-year research roadmap (2026 to 2031) aggregated from gaps identified across the reviewed literature

Year(s)	Phase	Focus Area	Key Objective	Representative KPIs
2026-2028	Phase 1: Foundations	Adapter benchmarking and HausaProp corpus	Build span-level 20-technique corpus; benchmark PEFT methods	Krippendorff alpha > 0.65; span Micro-F1 baseline
2027-2029	Phase 2: Transfer	Cross-lingual transfer and cultural adaptation	Apply MAD-X and few-shot transfer; build Hausa lexicon	Gains in rare-technique macro-F1 over zero-shot
2029-2031	Phase 3: Deployment	Deployment and fact-checker integration	Optimise on-device models; field evaluation with partners	Inference on 8 GB GPU or CPU; partner usability evaluation

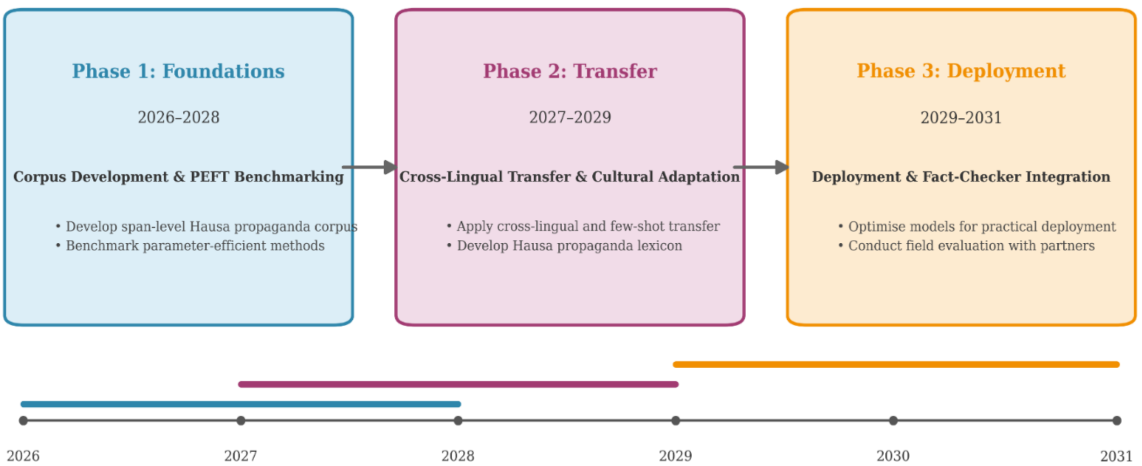

Figure 5: Three-phase Research Pipeline for Fine-grained Hausa Propaganda Detection (2026 to 2031).

5. Discussion

This systematic review of thirty-seven studies yields three principal findings. First, fine-grained propaganda detection has a mature taxonomy, established benchmarks, and robust transformer-based methods, but its empirical base is built almost entirely on English and European languages, with the SemEval-2023 multilingual task including no African language (Piskorski et al., 2023) and the only cross-lingual propaganda demonstration confined to a resource-rich English-to-Chinese pair (Zhang & Zhang, 2022). Second, parameter-efficient methods have firm theoretical grounding (Aghajanyan et al., 2021) and consistent evidence that they approach full fine-tuning at a small fraction of trainable parameters (Houlsby et al., 2019; Hu et al., 2021), including a Hausa demonstration at approximately 1.06% of parameters (Jibrin et al., 2026), yet none has been evaluated on the span-level structured prediction that propaganda detection requires. Third, Hausa NLP has advanced through a sequence of binary and coarse classification studies (Imam et al., 2022; Abdullahi Ibrahim et al., 2024; Sani et al., 2025) and several resource-building efforts (Muhammad et al., 2023; Abdulmumin et al., 2022), but it lacks both a fine-grained propaganda corpus and a systematic comparison of model families for such a task.

The coverage analysis in Section 4 makes the gap concrete. Across the reviewed tasks and approaches, the fine-grained propaganda task has no studies under any approach, while neighbouring tasks are well covered by multilingual and Afri-centric models. This is not a gap of feasibility but of attention: the building blocks are present. The trade-off analysis indicates that LoRA and bottleneck adapters offer the most favourable balance of efficiency and performance for constrained settings, while the transfer literature indicates that few-shot adaptation tends to outperform zero-shot transfer (Zhang, 2026; Providel et al., 2025), which suggests that even a moderately sized annotated Hausa corpus would yield meaningful gains. The evidence that direct translation misses culturally specific expression (Adam et al., 2023) and that Afri-centric models encode richer linguistic information than massively multilingual alternatives (Aduah & Meyer, 2025; Belay et al., 2025) argues for cultural adaptation of the taxonomy and for careful model selection rather than reliance on translation.

Several implications follow for different audiences. For researchers, the consistent value of high-quality human annotation (Maarouf et al., 2024) argues for rigorous annotation protocols and measurement of inter-annotator agreement rather than weak labelling, and the mixed results of augmentation across African tasks indicate that augmentation should be evaluated task by task rather than assumed to help. For practitioners in fact-checking organisations, the success of parameter-efficient methods supports building systems that run on modest hardware, lowering the barrier to deployment. For the wider field, extending fine-grained detection to a major African language would begin to reduce the linguistic inequality that leaves speakers of low-resource languages without tools that are routine for English.

The principal limitation of the present review is that its corpus is scoped to the literature most directly relevant to the research problem; a broader search may surface additional adjacent work, and the comparative scores presented here are qualitative syntheses that require confirmation through primary benchmarking on the span-level task. Additionally, the review does not include a quantitative meta-analysis due to the heterogeneity of reported metrics across studies. Future work should consider developing standardised evaluation frameworks for propaganda detection across African languages to enable more direct comparisons.

6. Conclusion

This review synthesised the literature at the intersection of fine-grained propaganda detection, parameter-efficient transfer learning, cross-lingual transfer, and Hausa and African language NLP, and reached a clear conclusion: the intersection itself is empty. No corpus, no lexicon, and no systematic model comparison exists for fine-grained propaganda technique detection in Hausa or any other major African language. The evidence assembled across the four areas nonetheless indicates that the task is feasible. Parameter-efficient methods work for Hausa at the document level, cross-lingual transfer works across language families for related tasks, and Afri-centric models encode useful linguistic information. What the evidence cannot yet establish is how these results combine for span-level technique detection under realistic resource constraints.

Three contributions structure the path forward. First, the review profiles the efficiency, performance, and deployment trade-offs among adaptation methods and identifies LoRA and bottleneck adapters as suitable starting points for constrained settings. Second, it maps the resource landscape and precisely locates the propaganda gap through a coverage analysis. Third, it consolidates the gaps and recommendations of the literature into a five-year roadmap that prioritises corpus development and adapter benchmarking, then cross-lingual transfer with cultural adaptation, then deployable systems evaluated with fact-checking partners. Realising this programme would extend computational propaganda detection to a major African language for the first time and contribute to more equitable access to information-integrity tools across languages.

Acknowledgments

The authors thank the HausaNLP community at Bayero University Kano for foundational resources and discussions that informed this review. The authors also acknowledge the broader African NLP research community for the datasets and models that made this synthesis possible.

Conflicts of Interest: The authors declare no conflict of interest.

Funding: This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Data Availability: This is a review article. All studies discussed are available through the cited sources.

References

Abdullahi Ibrahim, A., Umar Ali, H., Zahraddeen Yakubu, I., & Lawal, I. A. (2024). Fake news detection in Hausa language using transfer learning method. International Journal of Innovative Science and Research Technology (IJISRT), 2259-2269. https://doi.org/10.38124/ijisrt/IJISRT24OCT1050

Abdulmumin, I., et al. (2022). Hausa Visual Genome: A dataset for multi-modal English to Hausa machine translation. arXiv. https://doi.org/10.48550/ARXIV.2205.01133

Adam, F. M., Zandam, A. Y., & Inuwa-Dutse, I. (2023). Detection and analysis of offensive online content in Hausa language. arXiv. https://doi.org/10.48550/ARXIV.2311.10541

Aduah, W., & Meyer, F. (2025). Designing and contextualising probes for African languages. In Proceedings of the 6th Workshop on African Natural Language Processing (AfricaNLP 2025), 39-51. https://doi.org/10.18653/v1/2025.africanlp-1.7

Aghajanyan, A., Gupta, S., & Zettlemoyer, L. (2021). Intrinsic dimensionality explains the effectiveness of language model fine-tuning. In Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1), 7319-7328. https://doi.org/10.18653/v1/2021.acl-long.568

Ausat, A. M. A. (2023). The role of social media in shaping public opinion and its influence on economic decisions. Technology and Society Perspectives (TACIT), 1(1), 35-44. https://doi.org/10.61100/tacit.v1i1.37

Barrón-Cedeño, A., Jaradat, I., Da San Martino, G., & Nakov, P. (2019). Proppy: Organizing the news based on their propagandistic content. Information Processing & Management, 56(5), 1849-1864. https://doi.org/10.1016/j.ipm.2019.03.005

Belay, T. D., et al. (2025). AfroXLMR-Social: Adapting pre-trained language models for African languages social media text. In Findings of the Association for Computational Linguistics: EMNLP 2025, 15570-15587. https://doi.org/10.18653/v1/2025.findings-emnlp.842

Da San Martino, G., Yu, S., Barrón-Cedeño, A., Petrov, R., & Nakov, P. (2019). Fine-grained analysis of propaganda in news article. In Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), 5635-5645. https://doi.org/10.18653/v1/D19-1565

Du, C. (2025). Transformer model based on cross-language migration for low-resource Hausa news classification. In 2025 5th International Conference on Artificial Intelligence, Big Data and Algorithms (CAIBDA), 36-41. https://doi.org/10.1109/CAIBDA65784.2025.11182877

Ellul, J. (2021). Propaganda: The Formation of Men's Attitudes. Knopf Doubleday Publishing Group. https://books.google.com.ng/books/about/Propaganda.html?id=UpvZAAAAMAAJ&redir_esc=y

Houlsby, N., et al. (2019). Parameter-efficient transfer learning for NLP. arXiv. https://doi.org/10.48550/ARXIV.1902.00751

Hu, E. J., et al. (2021). LoRA: Low-rank adaptation of large language models. arXiv. https://doi.org/10.48550/ARXIV.2106.09685

Imam, S. H., Musa, A. A., & Choudhary, A. (2022). The first corpus for detecting fake news in Hausa language. In Emerging Technologies for Computing, Communication and Smart Cities, Vol. 875, 563-576. Springer Nature Singapore. https://doi.org/10.1007/978-981-19-0284-0_41

Inuwa-Dutse, I. (2025). NaijaNLP: A survey of Nigerian low-resource languages. arXiv. https://doi.org/10.48550/ARXIV.2502.19784

Jibrin, N., Aimufua, G., Onyedikachi, O. S., Anthony, A. A., Chukwunwike, U. S., & Aliyu, F. D. (2026). Parameter-efficient few-shot sentiment analysis using LoRA-enhanced transformers. Jurnal Masyarakat Informatika, 17(1), 146-159. https://doi.org/10.14710/jmasif.17.1.81053

Joaquim, A. F. (2025). Digital propaganda techniques: The specificities of neo-propaganda. In Advances in Wireless Technologies and Telecommunication, 39-62. IGI Global. https://doi.org/10.4018/979-8-3693-9999-6.ch003

Jowett, G., & O'Donnell, V. (2025). Propaganda & Persuasion (8th ed.). SAGE Publications.

Kaur, M., & Saini, M. (2024). Artificial intelligence inspired method for cross-lingual cyberhate detection from low resource languages. ACM Transactions on Asian and Low-Resource Language Information Processing, 23(9), 1-23. https://doi.org/10.1145/3677176

Li, X. L., & Liang, P. (2021). Prefix-tuning: Optimizing continuous prompts for generation. In Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1), 4582-4597. https://doi.org/10.18653/v1/2021.acl-long.353

Lupascu, M., Rogoz, A.-C., Stupariu, M. S., & Ionescu, R. T. (2026). Large multimodal models for low-resource languages: A survey. Information Fusion, 131, 104189. https://doi.org/10.1016/j.inffus.2026.104189

Maarouf, A., Bär, D., Geissler, D., & Feuerriegel, S. (2024). HQP: A human-annotated dataset for detecting online propaganda. In Findings of the Association for Computational Linguistics: ACL 2024, 6064-6089. https://doi.org/10.18653/v1/2024.findings-acl.363

Muhammad, S. H., et al. (2023). AfriSenti: A Twitter sentiment analysis benchmark for African languages. arXiv. https://doi.org/10.48550/ARXIV.2302.08956

Muhammad, S. H., et al. (2025). HausaNLP: Current status, challenges and future directions for Hausa natural language processing. In Proceedings of the 6th Workshop on African Natural Language Processing (AfricaNLP 2025), 176-191. https://doi.org/10.18653/v1/2025.africanlp-1.27

Nguyen, T.-P., Pham Thi Phuong, G., Do Phan, T., Le, A.-C., & Ngo Van, L. (2026). A framework for cross-lingual topic modeling using global context spaces and unified decoding with semantic knowledge distillation. Neurocomputing, 134278. https://doi.org/10.1016/j.neucom.2026.134278

Page, M. J., et al. (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. BMJ, n71. https://doi.org/10.1136/bmj.n71

Pfeiffer, J., Vulić, I., Gurevych, I., & Ruder, S. (2020). MAD-X: An adapter-based framework for multi-task cross-lingual transfer. In Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP), 7654-7673. https://doi.org/10.18653/v1/2020.emnlp-main.617

Piskorski, J., Stefanovitch, N., Da San Martino, G., & Nakov, P. (2023). SemEval-2023 Task 3: Detecting the category, the framing, and the persuasion techniques in online news in a multi-lingual setup. In Proceedings of the 17th International Workshop on Semantic Evaluation (SemEval-2023), 2343-2361. https://doi.org/10.18653/v1/2023.semeval-1.317

Providel, E., Mendoza, M., & Solar, M. (2025). Cross-lingual cross-domain transfer learning for rumor detection. Future Internet, 17(7), 287. https://doi.org/10.3390/fi17070287

Refaee, E. A., Ahmed, B., & Saad, M. (2022). AraBEM at WANLP 2022 shared task: Propaganda detection in Arabic tweets. In Proceedings of the 7th Arabic Natural Language Processing Workshop (WANLP), 524-528. https://doi.org/10.18653/v1/2022.wanlp-1.62

Ruder, S., Peters, M. E., Swayamdipta, S., & Wolf, T. (2019). Transfer learning in natural language processing. In Proceedings of the 2019 Conference of the North American Chapter of the ACL, 15-18. https://doi.org/10.18653/v1/N19-5004

Saaida, M., & Alhouseini, M. (2023). The influence of social media on contemporary global politics. https://doi.org/10.5281/ZENODO.10841255

Sani, B., et al. (2025). Who wrote this? Identifying machine vs human-generated text in Hausa. arXiv. https://doi.org/10.48550/arXiv.2503.13101

Smit Zaveri. (2024). Breaking language barriers: The role of large language models in multilingual communication. International Research Journal of Modernization in Engineering Technology and Science. https://doi.org/10.56726/IRJMETS58546

Young, K., Lee, A. M., Lee, E. B., & Strathaus, L. S. (1941). The fine art of propaganda:

A study of Father Coughlin's speeches. American Sociological Review, 6(1), 118. https://doi.org/10.2307/2086350

Zhang, W., & Zhang, X. (2022). Cross-lingual propaganda detection. In 2022 IEEE International Conference on Big Data (Big Data), 4330-4336. https://doi.org/10.1109/BigData55660.2022.10021059

Zhang, Y. (2026). Research on the application of cross-language pre-trained models in English machine translation. Franklin Open, 15, 100611. https://doi.org/10.1016/j.fraope.2026.100611

Zhang, Z., Wang, J., & Wang, X. (2023). Parameter-efficient cross-lingual transfer of vision and language models via translation-based alignment. In Findings of the Association for Computational Linguistics: EMNLP 2023, 7258-7268. https://doi.org/10.18653/v1/2023.findings-emnlp.483