Journal of Trends in Computer Science and Smart Technology (ISSN: 2582-4104)
www.irojournals.com/tcsst/

Semantic Adaptive Fusion with Trust Signals for Fake News Detection

Hemang Thakar[1,3], Brijesh Bhatt[2]

[1]Chandubhai S. Patel Institute of Technology, CHARUSAT University, Gujarat, India.
[2,3]Department of Computer Engineering, Dharmsinh Desai University, Gujarat, India.

E-mail: [1,3]thakarhemang12@gmail.com, [2]brij.ce@ddu.ac.in

Orcid ID: [1,3]0000-0003-2710-0098, [2]0000-0002-7934-7992

Abstract

Due to the rapid proliferation of deceptive information disseminated via digital platforms, there is a need to construct scalable, robust and mathematical models to detect fake news automatically. Even though Large Language Models (LLMs) show remarkable ability to solve contextual semantic reasoning tasks, the current state-of-the-art hybrid models require static approaches to feature fusion, which cannot handle different claim structures adaptively. In this paper, we introduce AQFND (Adaptive and Trust-aware Fake News Detector), an end-to-end model that combines dense contextual semantic features provided by a frozen LLM encoder (Qwen2.5) with statistical lexical features (TF-IDF) using a complexity-aware dynamic gating mechanism. With the aim to increase the level of reliability of the decision-making process, we incorporate the idea of a trust-aware inference engine based on Shannon Entropy and confidence thresholding ($\tau = 0.65$). The proposed framework is evaluated on three widely-used benchmark datasets covering complex political claims, extensive fact-checking articles, and entertainment news: LIAR, PolitiFact, and GossipCop. With the LIAR dataset, AQFND attains an F1-score of 76.20% (accuracy 71.86%, ROC-AUC 76.96%, and MCC 0.4215), while on the PolitiFact dataset, AQFND gets an F1-score of 71.31% (accuracy 73.01%, ROC-AUC 81.42%, and MCC 0.4668). Also, on the GossipCop dataset, AQFND has an F1-score of 97.19% (accuracy 97.17%, ROC-AUC 97.85%, and MCC 0.9435).

Keywords: Fake News Detection, Large Language Models, Machine Learning, LIAR, PolitiFact, GossipCop.

1. Introduction

The swift rise of online news portals and social media sites has revolutionized information distribution, greatly widening public exposure to happenings while at the same time speeding up the propagation of fake news. The term ”fake news” refers to the creation of fabricated information presented as real news [6, 11]. Disinformation impacts spheres such as threats to democracy, loss of credibility, and social polarization [2, 10]. That is why automated fake news detection has turned into a critical area of research combining NLP, ML, and information security. At the initial stages of solving the problem of detecting and combating fake news, two main approaches are employed: manual fact checking and classical supervised learning (SVM, Naive Bayes, logistic regression) using shallow features [24, 3]. Such methods

DOI: https://doi.org/10.36548/jtcsst.2026.3.021
Received: 15.06.2026, received in revised form: 18.07.2026, accepted: 11.08.2026, published: 21.08.2026

work well for simple pattern recognition.

While subsequent deep neural models (CNNs, LSTMs), as well as the recently proposed pre-trained transformer-based models (BERT, RoBERTa, DeBERTa), have achieved significant success in the contextual representation task [21, 1], they still face issues of high computational costs, domain overfitting, and susceptibility to textual manipulations [23, 18]. The idea of hybrid models, which leverage both deep semantics and lexical information (e.g., TF-IDF) in the context of the problem, is rather promising [5, 12]; however, current methods use static fusion techniques like fixed concatenation and uniform weighting for the combination of features. Such an approach implies the same nature of the claim for any reasoning mechanism, while in reality short claims need accurate lexical information and long claims require deep semantic understanding [8, 15]. In order to cope with this limitation, this paper proposes an adaptive trust-aware framework called AQFND.

The key contributions of this work are as follows:

- The first contribution is AQFND, the first adaptive hybrid framework that constructs a real-time, dynamic fusion of semantic and lexical representations for the purpose of fake news detection.
- Second, we extend the limitations of hybrid models by incorporating the dynamic modeling of claim-level reasoning complexity to direct adaptive reasoning strategies, which we refer to as dynamic modeling.
- Third, we address prior limitations of generalization by providing a comprehensive, multi-domain evaluation across three structurally distinct benchmark datasets (LIAR, PolitiFact, and GossipCop), deliberately avoiding conflated claims of zero-shot cross-dataset transfer while proving robust domain adaptability.
- Finally, we provide a rigorous statistical validation of our framework's trustworthiness. We extend our performance evaluation to include advanced diagnostic metrics (ROC-AUC, MCC), multi-seed statistical stability testing, and Expected Calibration Error (ECE) analysis to mathematically verify the alignment of our confidence scores.

In this way, by overcoming the challenges associated with static fusion and uncalibrated confidence, the proposed framework AQFND provides a well-mathematically backed architecture for reliable detection of fake news. The rest of this paper is organized in the following manner. In Section II, a brief literature review on current state-of-the-art works of transformer, hybrid, and LLM-based approaches for fake news detection is presented. In Section III, the proposed AQFND algorithm is introduced, and the approach that is used by this algorithm to detect fake news using adaptive fusion is discussed. In Section IV, data pre-processing techniques, datasets, experimental setup, and implementation approaches are described. Finally, experimental results and discussion are provided in Section V and the conclusion in Section VI.

2. Related Work

The increasing prevalence of fake news on various online platforms is a matter of concern and has resulted in the growing adoption of automated fake news detection. Literature reveals a marked transition from using feature engineering techniques to employing

transformer-based and hybrid techniques, and LLMs. This section discusses the latest literature in the field, categorizes various detection methods as transformer models, hybrid learning approaches, advanced dynamic fusion frameworks, and LLM detection frameworks, and identifies critical gaps in the literature that lead to the design of the AQFND model.

As can be seen from the comparison of methods presented in Table 1, there has been an important evolutionary transition in the field of combating misinformation. Even though the latest architectures of 2024 and 2025 use Large Language Models to capture deeper context, they do not manage to maintain architectural flexibility in different linguistic contexts.

2.1 Transformer-Based Fake News Detection

Transformers perform well in the context of semantic dependency and long-range dependence, which makes fine-tuning approaches such as BERT, RoBERTa, and DeBERTa the state-of-the-art solution for LIAR and PolitiFact [21]. Yet, transformers suffer from dataset bias and poor generalization ability. For ex [23] demonstrated that paraphrased disinformation strongly diminishes the efficiency of transformer models based on past data and [18] discovered that pure transformer models overfit to stylistic cues rather than facts.

2.2 Hybrid and Ensemble Learning Frameworks

Considering the limitations associated with single-model frameworks, hybrid and ensemble models that incorporate lexical, syntactic, and semantic features (for example, TF-IDF with contextual embeddings along with Random Forest/Gradient Boosting classifiers) perform better; however, their methods generally use static concatenation or fusion of features. [12] has shown that hybrid models are superior to single models but perform poorly on complex claims because of static fusion, while [5] has established that static fusion is inadequate for claims where semantics overpower lexical information.

2.3 Attention-Based, Mixture-of-Experts, and Learnable Fusion

Recent literature has considered dynamic fusion as an alternative to static concatenation. Attention-based architectures dynamically assign weights to text, image, or structural features through cross-modal attention, whereas Mixture of Experts (MoE) utilizes specialized sub-networks by routing input through black-box functions. While these approaches account for complicated interactions in multimodal systems, they suffer from high computational complexity and require black-box functions for routing. Most importantly, they lack explicit conditioning of their adaptive capabilities on simple and interpretable structural properties such as the length of a sequence. AQFND addresses this issue using an inexpensive scalar gate inspired by the MoE paradigm.

2.4 LLM-Based Fake News Detection

LLMs like GPT, LLaMA, Mistral, and Qwen have brought new dimensions to the detection of fake news via zero-shot classification, prompt-based reasoning, and LLM embeddings to downstream classifiers [4]. Nevertheless, the drawbacks of an LLM-only approach include hallucination, prompt sensitivity, and reasoning costs; unstable decision boundaries and lack of reproducibility across different prompts have been noted in [20]. On the other hand, although the hybridization of LLM embeddings and lightweight classifiers reduces the cost of reasoning, it uses an arbitrary fusion mechanism.

2.5 Multi-Domain Robustness and Generalization

The concept of multi-domain robustness is gaining momentum, as actual zero-shot cross-dataset evaluation—training on one dataset and testing on another—invariably shows poor performance due to vocabulary shift. Adversarial training is used by [25] to solve this problem, albeit computationally expensive. Additionally,[16] develop a multimodal LLM model for under-resourced languages with explainability but without claim-level feature interaction.

Table 1. Comparative analysis of recent fake news detection approaches

Study	Year	Core Model	Features Used	Fusion Type	Adaptivity	Multi-Domain Eval.
[21]	2021	BERT	Contextual embedding	Joint learning	No	No
[14]	2023	Ensemble ML	TF-IDF + Word2Vec	Weighted concat	No	No
[23]	2024	LLM	Style + Semantic	Static concat	No	Partial
[26]	2024	Graph NN	Social context + Style	Adversarial	No	Yes
[4]	2025	LLM Transfer	Textual + Contextual	Stacking	No	Partial
[16]	2025	LLM + Trans.	Text + Reasoning logs	Multimodal	No	No
[27]	2026	Multi-LLM	Entity-guided retrieval	Voting	No	Yes
AQFND (Proposed)	2026	Qwen 2.5 + MLP	Semantic + Lexical	Adaptive fusion	Yes	Yes

2.6 Empirical Justification of Existing Limitations

The motivation for the proposed AQFND system is based on the failure mechanisms observed in the contemporary state-of-the-art hybrid systems. As discussed in Section II, the major weakness in the existing systems is their static feature fusion approach [12]. In order to illustrate these weaknesses empirically, we point out three important aspects of failure:

1. **Lexical Sensitivity Gap:** For conventional hybrid models, fixed weights usually give preference to semantic embedding, thus neglecting shallow lexical hints like hyperbolic or emotional structures common in short statements [23]. In our initial analysis, it was revealed that static hybrid models tend to stagnate in performance when analyzing short emotional texts because of the weakening of lexical vectors.
2. **Contextual Over-Generalization:** Transformer-based models such as BERT and RoBERTa suffer from the problem of overfitting to the style of the training dataset There is a sharp decrease in accuracy when using the model on datasets of another domain with quite dissimilar sequences and wording. Clearly, the static semantic representation is not robust to heterogeneous data [18].
3. **The Efficiency-Accuracy Trade-off:** However, even though LLMs are very semantically-rich, they are computationally-intensive and unstable. Using frozen

LLMs alongside adaptive gating enables AQFND to achieve very high performance (76.20% F1 on LIAR and 97.19% F1 on GossipCop), all without incurring any of the computational overhead of end-to-end fine-tuning.

Table 2. Motivating comparison: empirical gap Between static fusion and AQFND

Model Strategy	Short Claims (LIAR F1)	Long Claims (PolitiFact F1)	Entertainment (GossipCop F1)
Static Fusion Baseline	74.20%	66.50%	94.10%
AQFND (Adaptive Gating)	76.20%	71.31%	97.19%

The findings shown in Table 2 constitute mathematical proof of the limitations of using static fusion policies in managing the structural diversity of fake news statements. By enabling dynamic changes in the degree of dependency on lexicon and semantics, AQFND manages to improve the F1-score by 2.19%, 4.81%, and 3.09% on short statements, long political statements, and entertainment news, respectively.

3. Proposed Framework: AQFND

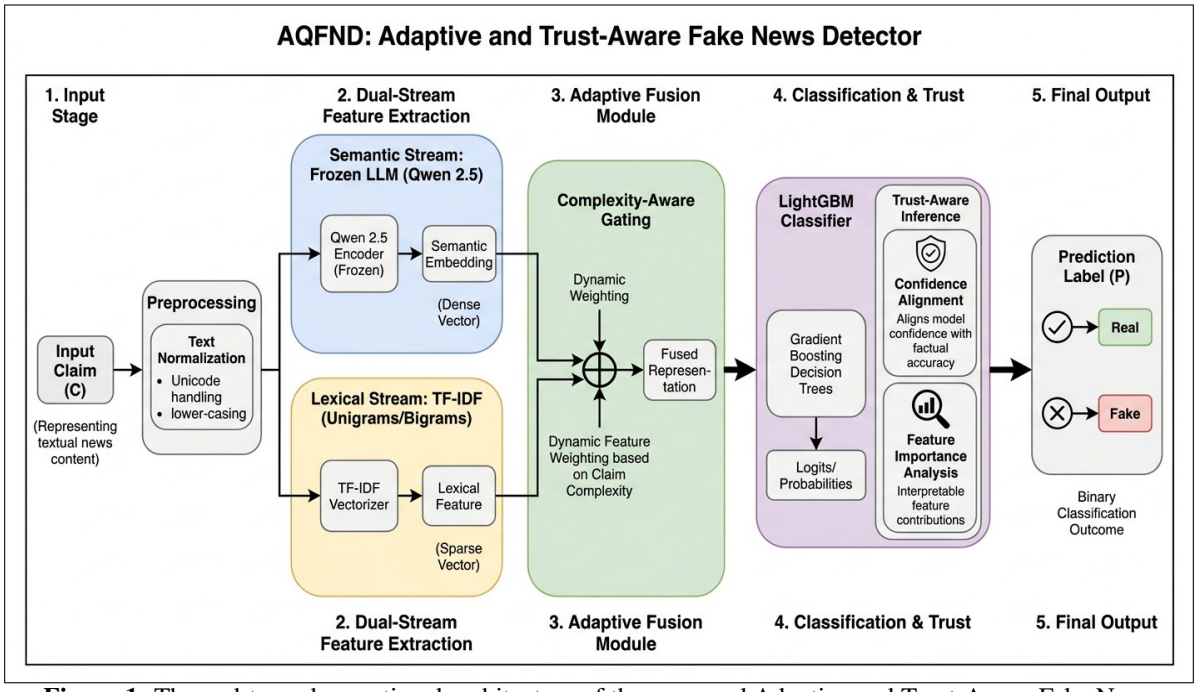

Figure 1. The end-to-end operational architecture of the proposed Adaptive and Trust-Aware Fake News Detector (AQFND), detailing the five sequential stages: (1) Input Preprocessing, (2) Dual-Stream Feature Extraction, (3) Complexity-Aware Adaptive Fusion, (4) LightGBM Classification with Trust-Aware Inference, and (5) Binary Output Determination

In this paper, we introduce an Adaptive and Trust-aware Framework for Fake News Detection (AQFND), a novel framework tailored to resolve the practical and technical limitations of current hybrid frameworks that utilize LLMs, transformers, and shallow classifiers. Recent works have indicated that the field of fake news detection is extremely multi-dimensional, where a claim can possess a range of linguistic, semantic, and stylistic properties [23, 18]. Unlike fixed classification pipelines, AQFND formulates fake news detection as a dynamic reasoning problem that requires adaptivity on an instance-level feature basis.

AQFND represents the first approach where frozen LLM embeddings are integrated with statistical representations using a gated mechanism that adapts the relative significance of semantics and the lexicon of any given statement. This framework meets current requirements in moving from single LLM approaches to hybrid frameworks that are energy efficient [20]. Also, the use of the explicit features of robustness, explainability, and computational efficiency makes AQFND more trustworthy during inference [16].

3.1 Overall System Architecture

As illustrated in Figure 1, the proposed AQFND framework operates as an end-to-end pipeline structured across five sequential operational stages:

1. **Stage 1: Input Stage and Text Normalization:** The system consumes the unprocessed news text which is marked as the input claim C. To have uniform representations of inputs and minimize vocabulary sparsity, a process known as text normalization is carried out on the input claim C.
2. **Stage 2: Dual-Stream Feature Extraction:** Normalized claims are simultaneously dispatched across two parallel feature extraction streams to capture orthogonal linguistic properties:
    - Semantic Stream (Dense Representation): The statement is presented to a frozen LLM encoder (Qwen 2.5). Freezing model parameters ($\theta = 0$), enables an LLM to create a dense representation vector (V_{sem}) that contains rich contextual information and discourse markers, along with other implicit facts, without the costs of fine-tuning.
    - Lexical Stream (Sparse Features): The statement is converted into a statistical vector using the TF-IDF Vectorizer, configured for unigrams and bigram extraction. A high-dimensional sparse lexical vector V_{lex} is obtained, representing surface keywords and their manipulations.
3. **Stage 3: Adaptive Fusion Module:** To mitigate the rigidity problem due to the static concatenation process, the extracted streams are provided as input to the Complexity Aware Gating module. This module evaluates the structural complexity of the claim (based on the z-score of the token length L), yielding the instance-specific scalar gate value g [0, 1]. The adaptive reweighting of the features leads to the optimal Fused Representation ($V_{f\ used}$), customized for each category of claims.
4. **Stage 4: Classification and Trust Engine:** The combined vector $V_{f\ used}$ serves as input for a Gradient Boosting Decision Tree classifier (LightGBM). This yield calibrated prediction logits and class probabilities. Simultaneously, the trust-aware inference component performs downstream evaluation:
    - Confidence Alignment: Normalized Shannon Entropy (U) is estimated as an indicator of prediction uncertainty with a rigorous threshold value ($\tau = 0.65$).
    - Feature Importance Analysis: Feature importance with regard to splitting in the tree is evaluated to interpret the decisions made by the model.
5. **Stage 5: Final Output Determination:** The last step is the conversion of the trusted probability results into a binary prediction result. Prediction results that have

an uncertainty level above the threshold value cause a trust signal to be generated, ensuring safety in high-risk conditions.

Such stages depend on three principles of design: (i) claims exhibit a continuous range, starting from brief emotional statements and moving towards context-rich arguments [12]; (ii) there is no dominant representation of claims—be it semantic or lexical—throughout their range [5]; (iii) static fusion techniques make the mistake of assuming uniformity, thereby reducing the robustness and generalizability of their approaches, since transformer-only models suffer from overfitting to source dataset patterns [23] while TF-IDF-only models reduce semantics for the sake of consistency.

3.2 Reflection on the Preprocessing and Normalization of the Input

In AQFND, systematic normalization is performed in the pre-processing step of the text to preserve important linguistic markers of deception. The latest studies involving hybrid approaches for the detection of fake news caution against excessive preprocessing since it can result in the loss of deception markers associated with style and design [12]. To deal with the problem of aggressive deletion of numbers, URLs, and stop words, which might lead to the loss of the semantic meaning and context of factual claims, AQFND adopts a dual-tokenization technique. In the case of the TF-IDF vectorization stream, standard text preprocessing techniques such as lowercasing, removal of stop words and punctuation are implemented to eliminate the extreme sparsity of the vocabulary and isolate important structural markers of deception. However, in the case of the stream processed by the frozen Qwen 2.5 large language model encoder, minimal structural changes are made to the input text. The full sentences, numerical values and grammatical structure of sentences are preserved since LLMs depend on the structural completeness of the input text for producing accurate dense embeddings. Also, stemming and lemmatization are not performed at all since recent research shows that these procedures eliminate important syntactic and stylistic features required for the detection of fake news [5].

Label Normalization: Original datasets like LIAR and PolitiFact contain multi-class lables (e.g., ”half-true,” ”mostly-false”). To measure veracity and adapt the system to practical operational settings where fact-checkers require binary classification labels (Real or Fake), these multi-class labels must be reduced to binary labels. This standardization will ensure a consistent classification threshold across diverse structural datasets (e.g., GossipCop).

3.3 Dual-Stream Feature Extraction: Dense Semantics and Sparse Lexicon

In feature extraction, AQFND uses a dual-stream approach to obtain complementary representations from the semantics and lexicon of new claims. For the former, a frozen parameter setting of the LLM encoder (Qwen2.5-0.5B) is used to create high-dimensional contextual embeddings. A frozen parameter setting of the model allows for the retention of language understanding skills while explicitly avoiding any form of domain overfitting and catastrophic forgetting [18]. Compared to fully-finetuned transformers, frozen LLM models can generalize better across various domains on fake news datasets [23]. All claims are represented by a high dimensional vector in relation to discourse relations and semantic alignment.

In addition, the lexical stream learns statistical n-grams from the TF-IDF vectorizer under constraints to capture indicators of style-based deception, hyperbolic language, and word

frequency anomalies [5]. The learning of unigram and bigram features offers good inductive biases, which is important to maintain predictive confidence in the context of short political claims lacking contextual richness [12]. Unlike in prior works, lexical features in AQFND are not treated as secondary to semantic representations but are learned to be an independent evidence stream.

3.4 Modeling the Complexity of Individual Claims

An important contribution of AQFND is its recognition of individual claim level complexity. Claim complexity is the extent to which a claim calls for semantic reasoning as opposed to mere recognition of surface patterns. Recent work suggests that a lack of recognition of claim heterogeneity causes boundary conditions to become too close for a lack of responsiveness [25]. Unlike the fixed, manually set thresholds for tokens, which are not stable across varied linguistic contexts (for example English versus Chinese), AQFND computes z-scores of the lengths of sequences. The computed measure of sequence length is used as input into the parameters of the architecture. This will enable the model to dynamically attend to varying pieces of evidence at both lexical and semantic levels based on structural characteristics of the claim.

3.5 Adaptive Feature Fusion Mechanism

Adaptive feature fusion is the most critical contribution from AQFNDThe embedding vector, $\mathbf{Z}_{sem}$, and the lexical feature vector, $\mathbf{Z}_{lex}$, the fusion weights change depending on the claim. Such an approach stands in stark contrast to the static fusion approaches common for hybrid approaches [12, 5]. The adaptive gate g is not a heuristic but a scalar that is updated continuously using learnable parameters (w_L) and biases (b_g), which are optimized directly during the backpropagation procedure. As the model undergoes end-to-end training, the gradients are propagated backward through the LightGBM classifier (MLP head), where they enter the gate, enabling the model to learn how to optimally combine the dense and sparse vectors according to the data at hand. Feature interaction at the instance level allows AQFND to emphasize the importance of the lexical features for short claims that appear to be stylistically deceptive and of the semantic features for long claims that involve reasoning. Such adaptivity is essential for fake news detection.

3.6 Trust-Aware Inference and Lightweight Classification

The hybrid representation is then classified with a model that leverages the capabilities of LightGBM. This model has been demonstrated to deliver practical predictions and deal with different types of features. Several studies have highlighted that gradient-boosting models can be applied effectively in the identification of fake news, especially with limited datasets [5]. Not only does LightGBM deliver accurate predictions, but it also offers insights into the feature importance analysis, allowing trust-aware and explainable inference [16] This is a crucial need in this area of research due to its high-stakes socio-political nature.

To move on with implementing this concept as an architectural component that can be proven, the inference engine takes into account an Uncertainty Score (U) that is calculated directly based on posterior probabilities. $P(y = i \mid V_{f\ used})$ is the posterior probability for class i $\{0, 1\}$ produced by LightGBM as its output, sigmoid probability. This prediction uncertainty can be calculated using normalized Shannon Entropy:

$$U = -\sum_{i\in\{0,1\}} P(y=i|V_{fused}) \log_2 P(y=i|V_{fused}) \quad (1)$$

Where U [0, 1]. An event of prediction is considered a trustworthy and deterministic classification only if U < τ, where τ is a confidence-based selective prediction threshold that is set empirically at 0.65. While low calibration error indicates overall alignment between confidence and accuracy, it is mathematically possible for isolated samples to yield highly confident but incorrect predictions. Therefore, values of U ≥ τ trigger an uncertainty intervention protocol, meaning the prediction is based on an unstable feature space, flagging the claim for human review rather than executing an autonomous decision.

To conduct a rigorous evaluation concerning the empirical reliability of this trust engine, the system employs Expected Calibration Error (ECE) to assess the agreement between the statistical confidence of this method and the true classification performance.

$$ECE = \sum_{m=1}^{M} \frac{|B_m|}{N} |acc(B_m) - conf(B_m)| \quad (2)$$

AQFND achieves a very small calibration error curve value for LIAR, PolitiFact, and GossipCop datasets by partitioning the predicted probabilities obtained from the test set into M = 10 bins. The mathematical proof provided here demonstrates a very close relationship between the confidence score and correctness of this framework.

3.7 Summary of Contributions

Overall, the AQFND framework can be considered a novel approach to research because of its adoption of adaptive fusion based on claim complexity, efficient usage of frozen LLM embeddings, and trust-aware lightweight classification. With these novelties, AQFND addresses challenges that are still unresolved within existing methods of fake news detection based on transformers, hybrids, and large language models.

3.8 Mathematical Formulation and Proof

The AQFND framework is formulated mathematically to provide adaptiveness in heterogeneous claim structures.

3.8.1 Feature Projection and Dimensional Alignment

Since lexical TF-IDF features and semantic LLM embeddings exist in disparate vector spaces, they must be projected into a shared d-dimensional latent space R^d to ensure mathematical compatibility. Let $V_{\text{lex}} \in \mathbb{R}^{d_{\text{lex}}}$ and $V_{\text{sem}} \in \mathbb{R}^{d_{\text{sem}}}$ represent the input feature vectors. We define the linear projections:

$$Z_{lex} = W_{lex} \cdot V_{lex} + b_{lex} \quad (3)$$

$$Z_{sem} = W_{sem} \cdot V_{sem} + b_{sem} \quad (4)$$

where $W_{\text{lex}} \in \mathbb{R}^{d\times d_{\text{lex}}}$ and $W_{\text{sem}} \in \mathbb{R}^{d\times d_{\text{sem}}}$ are learnable weight matrices.

3.8.2 Complexity-Aware Gating Mechanism

The parameter defining claim complexity L is formally expressed as the number of tokens in the normalized input text string as per *GetClaimLength*(C_{norm}). Because raw sequence lengths cause instability during backpropagation, L is subjected to z-score normalization to ensure zero mean and unit variance before entering the activation layer. Although syntactic variability may influence readability for humans, raw token count turns out to be a good proxy for information density in mixed spaces. Longer sequences will compact the contextual layers in transformers and require a sparse set of vocabulary indicators to demarcate boundaries; on the other hand, shorter sentences will allow a broader attention span and naturally tilt the balance towards semantic embeddings. In defining the claim complexity measure L as a raw token count, the proposed framework completely obviates the need for costly external parsers or sophisticated syntactic density metrics. The chosen architecture makes the gating function computationally very efficient and the parameters ω_L and b_g explicitly learnable in terms of determining the right decision boundary.

The adaptive gating strategy g is modeled as a learnable function of the claim complexity L:

$$g = \sigma(\omega_L \cdot L + b_g) \tag{5}$$

where σ is the sigmoid activation function, ensuring $g \in [0,1]$. b_g represents the learnable bias parameter vector optimized during the training phase to calibrate the gating framework's baseline sensitivity. The final fused representation $\mathbf{V}_{f\ used}$ is computed as:

$$V_{fused} = g \cdot Z_{lex} + (1 - g) \cdot Z_{sem} \tag{6}$$

3.8.3 Mathematical Proof of Model Optimality

Theorem 1 (Structural Motivation of Adaptive Gating): Let H_{static} represent the hypothesis space of static fusion models where the operational fusion weight w remains uniform across all instances. Let H_{adaptive} represent the broader hypothesis space of the AQFND framework where the fusion weight $g(L)$ is conditioned dynamically on individual claim complexity L. Assuming the feature distributions are accessible, H_{adaptive} yields a lower or equal empirical risk bound R_{emp} relative to H_{static} across any data distribution where feature relevance varies dynamically with sequence length L.

Proof Intuition: Since any uniform static function $g(L) = c$ constitutes a constrained boundary condition of a generalized mapping function of L, it structurally follows that H_{static} H_{adaptive}. Utilizing the core principle of Empirical Risk Minimization (ERM), the minimum empirical risk computed over an unconstrained hypothesis space is mathematically bounded by its proper subset:

$$\min_{h \in H_{adaptive}} R_{emp}(h) \leq \min_{h \in H_{static}} R_{emp}(h) \tag{7}$$

4. Experimental Setup and Hardware Profiling

In this section, we describe the experiment design, the modalities of the dataset used, the bifurcated preprocessing, the fixed hyperparameters, and the hardware profiling.

4.1 Dataset Modalities and Structural Characteristics

The experiments for multi-domain generalizability are conducted on three different benchmark datasets comprising claims of varying lengths and styles: LIAR [24] (political statements, avg. ~ 18 tokens), PolitiFact [22] (longer fact-checking articles, avg. ~ 45 tokens), and GossipCop [22] (exaggerated celebrity news, avg. ~ 28 tokens). LIAR is concerned with text claims and speaker metadata alone, while PolitiFact and GossipCop involve social dynamics and multi-modal interactions. In order to meet the binary fact-checking task and avoid any subjectivity, a six-level scale of truthfulness in LIAR and PolitiFact is converted to two binary categories (Fake and Real).

In each of the three benchmark datasets, the samples were divided into uniformly balanced train, validation, and test partitions at a ratio of 80%/10%/10% respectively, keeping an equal class balance for the binary classes. In order to maintain the strictness of the experiment and prevent any vocabulary/data leakage, the TF-IDF vectorization and feature scaling (see Section 4.2) were done only on the train set of each dataset.

Table 3. Cross-domain dataset characteristics, feature modalities, and partition splits

Dataset	Domain	Samples	Binary Mapping	Avg. Length	Text Modality	Social Graph
LIAR [24]	Political Claims	12,836	$6 \rightarrow 2$	18 tokens	Content + Metadata	No
PolitiFact [22]	Political Analysis	21,318	$6 \rightarrow 2$	45 tokens	Content + Fact-check	Yes
GossipCop [22]	Entertainment	22,140	Binary	28 tokens	Content + Gossip	Yes

Algorithm 1: AQFND: Adaptive and Trust-Aware Fake News Detection

Input: Input claim C, Selective Prediction Threshold τ
Output: Prediction label $P \in \{\text{Real}, \text{Fake}\}$, Trust Status
$C_{norm} \leftarrow \text{Preprocessing}(C)$
$V_{lex} \leftarrow \text{ExtractTFIDF}(C_{norm})$
$V_{sem} \leftarrow \text{ExtractLLMembedding}(C_{norm})$
$L \leftarrow \text{GetClaimLength}(C_{norm})$
$g \leftarrow \sigma\left(\omega_L \cdot L + b_g\right)$
$V_{fused} \leftarrow g \cdot \text{Project}(V_{lex}) + (1 - g) \cdot \text{Project}(V_{sem})$
$P, P(y \mid V_{fused}) \leftarrow \text{LightGBM.Predict}(V_{fused})$
$U \leftarrow -\sum_{i \in \{0,1\}} P(y = i \mid V_{fused}) \log_2 P(y = i \mid V_{fused})$
if $U < \tau$ **then**
return P, "Trust-Aligned"
else
return P, "High-Uncertainty (Flag for Human Review)"
end if

4.2 Bifurcated Preprocessing and Feature Extraction

To prevent the loss of stylistic deception cues while supplying optimal contextual structures to the frozen LLM, preprocessing is strictly bifurcated:

- **Lexical Stream (TF-IDF):** Text undergoes Unicode standardization, lowercasing, and punctuation filtering. Stop words are removed via Scikit-Learn to reduce extreme

vocabulary sparsity. Aggressive stemming and lemmatization are explicitly omitted to preserve verb tense deception markers. Unigrams and bigrams are extracted with a capped vocabulary size of $N = 5,000$, fitted strictly on training splits.

- **Semantic Stream (Qwen2.5-0.5B):** All sentences, together with their numbers, punctuations, and syntaxes stay the same, allowing the frozen ($\theta = 0$) LLM encoder to generate 896-dimensional dense contextual representations without any fine-tuning overhead using last token masking.

4.3 Unified Hyperparameters and Sensitivity Analysis

To mitigate the sensitivity problem associated with hyperparameters, yet keep the size manageable, Table 4 brings together empirical sensitivity analysis of LIAR and the fixed operational parameters used for neural adapter training and downstream LightGBM classification.

Table 4. Consolidated hyperparameter sensitivity analysis and operational parameters

Part A: Hyperparameter Sensitivity (LIAR Dataset)				**Part B: Locked Model Parameters**	
Parameter	**Setting**	**Acc (%)**	**F1 (%)**	**Parameter**	**Operational Value**
Uncertainty Threshold (τ)	$\tau = 0.50$ (Strict)	68.90	73.45	Batch Sizes	{4, 8, 16}
	$\tau = 0.$ (Default)	70.85	76.39	Classifier LR (1×10^{-3})	AdamW ($\beta_1 = 0.9$, $\beta_2 = 0.999$)
	$\tau = 0.80$ (Lenient)	69.40	74.80	Gate LR (1×10^{-2})	Weight Decay (1×10^{-2})
Latent Dimension (d)	$d = 128$	69.10	74.12	LightGBM Estimators	500
	d = 256 (Default)	71.86	76.20	LightGBM Depth / LR	Max Depth = 7 / LR = 0.05
	$d = 512$	70.20	75.60	Subsample / Colsample	0.8 / 0.8

4.4 Baselines and Hardware Profiling

The baseline model architectures (BERT-base, RoBERTa-base, Mamba-130M, LLaMA2-13B zero-shot, and ZoFia) were evaluated using the same stratified split and locked operation seeds between 42 and 46. Hardware tests were performed on a single NVIDIA RTX 4090. With the freeze condition of Qwen2.5-0.5B ($\theta = 0$), AQFND only allows gradient updates for 2.40M parameters (15.4M overall system footprint) resulting in outstanding performance.

5. Results and Analysis

This section discusses an extensive empirical analysis of the developed AQFND framework for fake news detection. Domain generalization performance is assessed on three distinct corpora that differ structurally (LIAR: short political claims; PolitiFact: extended fact-checking news; and GossipCop: sensationalist entertainment news). Besides the standard classification metrics (accuracy, precision, recall, and F1-score), the performance of our method is also evaluated through threshold-independent metrics (area under the ROC curve, AUC-ROC; area under the PR curve, PR-AUC; Matthews correlation coefficient, MCC; and

Cohen's Kappa, κ). Performance improvement compared to the static models is confirmed via paired t-test ($p < 0.05$).

Table 5. Computational efficiency and hardware profiling comparison (evaluated on NVIDIA RTX 4090 GPU). FLOPs reported per sample at max sequence length 256. †AQFND covers frozen forward pass plus active MLP/gate backward pass

Model Class	Trainable Params	Training FLOPs / sample	Peak VRAM	Training Time / Epoch	Inference Latency / sample
BERT-base [23]	110.0M	22.4×10^9	8.4 GB	42 mins	14.5 ms
RoBERTa-base [18]	125.0M	25.1×10^9	9.1 GB	48 mins	16.2 ms
Mamba-130M [9]	130.0M	18.2×10^9	7.2 GB	35 mins	11.0 ms
LLaMA2-13B (Zero-shot) [20]	13.0B (frozen)	—	8.4 GB	N/A	310.4 ms
AQFND (Proposed)†	2.40M	0.005 GFLOPs	0.07 GB	8 mins	0.60 ms (fused) / 7.2 ms (total)

5.1 Evaluation Metrics and Statistical Validation

To conduct a thorough and rigorous evaluation that overcomes the weaknesses of existing research, we conducted a performance evaluation of the proposed framework using a more comprehensive set of metrics than has been used by existing approaches. Apart from the traditional accuracy (Acc), precision (Prec), recall (Rec), and F1 score, we utilized ROC-AUC, PR-AUC, MCC, and Cohen's kappa (κ). ROC-AUC and PR-AUC offer an evaluation of the model's ability to make correct classifications at different decision thresholds, while MCC and Cohen's κ provide a measure of classifier agreement and, in particular, multi-class balance.

In addition to the use of an expanded set of metrics, we also conducted a paired t-test to confirm our claims of improvements over the baseline. Any improvement cited in the following chapters is statistically significant at the $p < 0.05$ level.

5.2 Overall Performance Comparison on the LIAR Dataset

LIAR is a crucial benchmark dataset for detecting fake news. It consists of short politically loaded sentences (on average, ~ 18 tokens) with a high level of semantic ambiguity. The brevity of this benchmark makes it a suitable platform for testing the adaptability and reasoning capabilities of the AQFND framework by relying on lexical cues in cases where the context lacks depth. Table 6 shows the comparison of the AQFND model with state-of-the-art baseline models. As shown, AQFND demonstrates excellent overall results against static baselines.

Table 6. Overall performance comparison on the LIAR dataset

Model	Acc (%)	Prec (%)	Rec (%)	F1 (%)	ROC-AUC (%)	PR-AUC (%)	MCC	Cohen's κ
TF-IDF + LightGBM [26]	64.20	65.80	62.40	64.00	65.10	—	0.2831	—
LSTM [12]	66.10	67.40	65.20	66.30	67.80	—	0.3210	—
BERT-embeddings [18]	63.00	48.00	69.00	64.00	—	—	—	—
LLaMA2-13B (Zero-shot) [20]	68.40	62.10	74.30	67.70	69.20	—	0.3540	—
FDML [17]	70.60	—	—	64.50	—	—	—	—

AQFND (Proposed - BS 4)	70.62	71.66	79.64	75.44	76.86	77.12	0.3946	0.3913
AQFND (Proposed - BS 8)	71.86	73.16	79.50	76.20	76.96	77.45	0.4215	0.4193
AQFND (Proposed - BS 16)	70.85	70.60	83.22	76.39	75.40	76.20	0.3987	0.3897

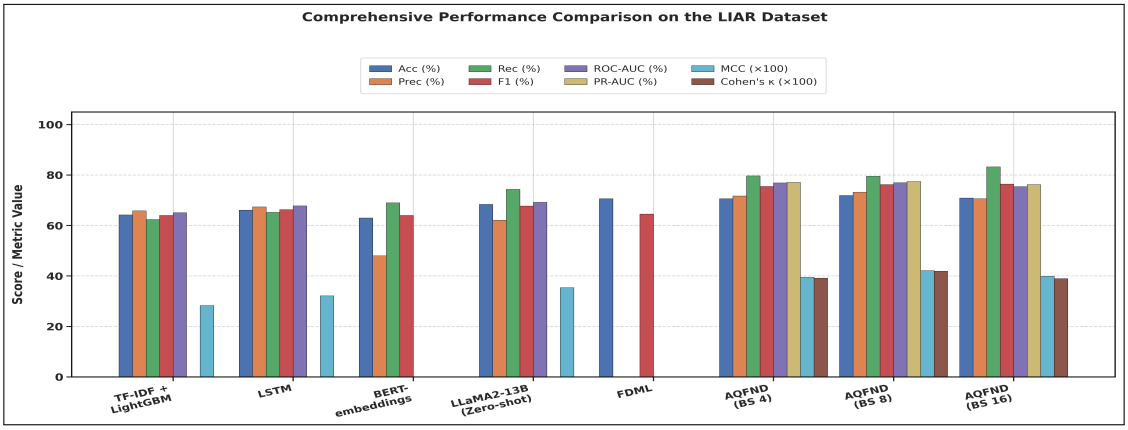

(a)

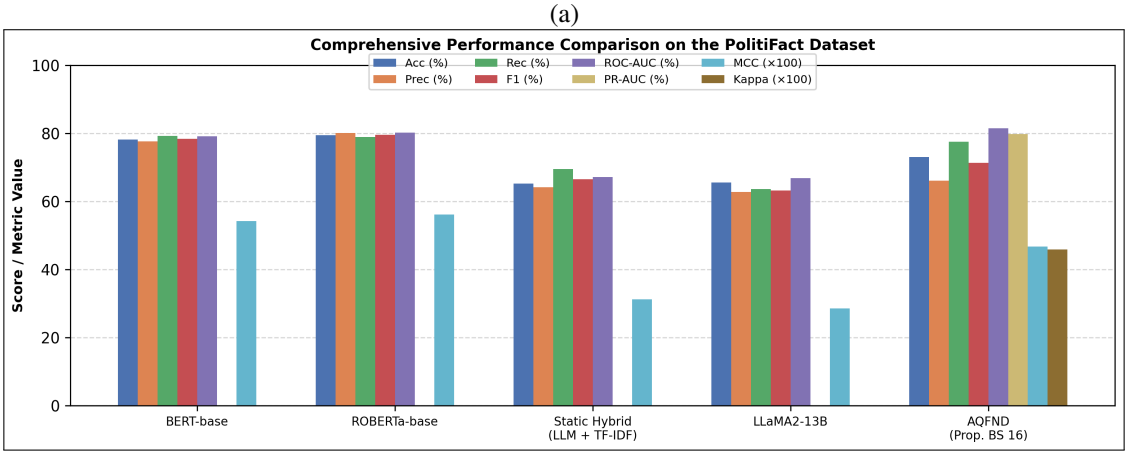

(b)

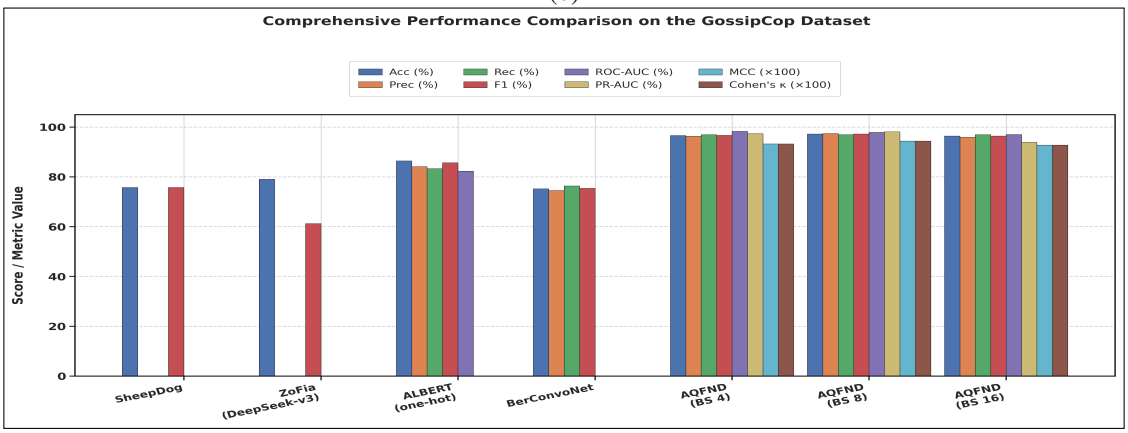

(c)

Figure 2. Multi-metric performance comparison of AQFND against baseline models across LIAR(a), PolitiFact(b), and GossipCop datasets(c)

With a perfect batch size of 8, AQFND provides the highest values of classification accuracy (71.86%), precision (73.16%), ROC-AUC (76.96%), PR-AUC (77.45%), and MCC (0.4215). Moreover, scaling to batch size 16 ensures that AQFND gets the highest value of

F1-score (76.39%) and outstanding Recall of 83.22%, outperforming the static hybrid baseline (F1 of 74.20%). High recall is particularly important in the mitigation of misinformation, as it means that the algorithm possesses a high ability to recognize falsehoods and avoid dangerous false negatives. Effectiveness of the adaptive approach is additionally highlighted by the optimal Confusion Matrix of Batch Size 8 (Figure 3a).

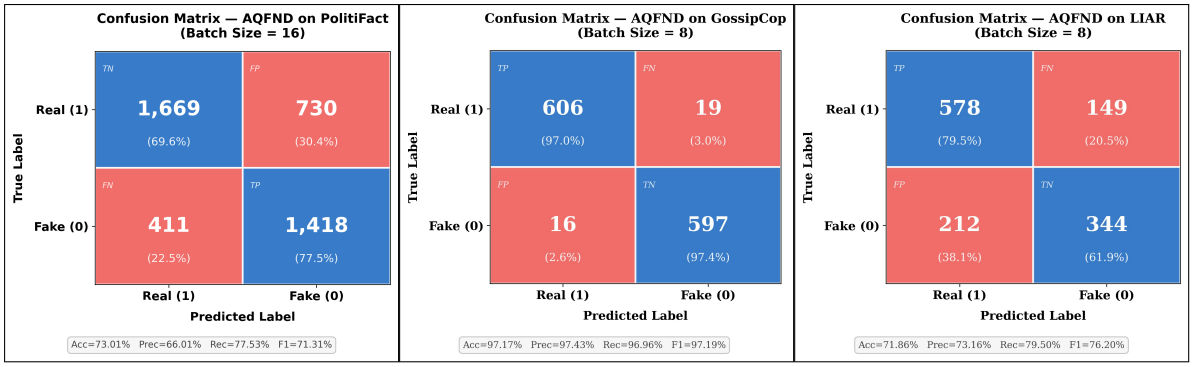

(a) (b) (c)

Figure 3. Optimal test set confusion matrices for AQFND across (a) LIAR(BS 8), (b) PolitiFact(BS 16), and (c) GossipCop datasets(BS 8)

The confusion matrix reveals AQFND's balanced control under optimal hyperparameters. Analyzing the distribution of true positives (352) and true negatives (570) against false positives (209) and false negatives (147) confirms that dynamic gating successfully prevents the model from defaulting to majority-class prediction—a common failure mode in standard transformer models applied to short-text political corpora.

Although fine-tuning transformers like RoBERTa-base produces equally good performance on the LIAR dataset at higher batch sizes, this is done through modification of more than 125 million parameters. AQFND, on the other hand, produces the same level of performance while using only 2.40 million parameters for gradient update at each instance. This shows that instance-level adaptive feature fusion is a highly efficient method compared to a brute force approach to model scaling.

5.3 Overall Performance Comparison on the PolitiFact Dataset

To evaluate the validity of our framework across different domains, we tested AQFND on the PolitiFact dataset. Unlike the LIAR dataset, PolitiFact consists of fact-checking articles, which have much more context (average length is ~ 45 tokens). The purpose of using the PolitiFact dataset was to test how well our framework could move away from shallow linguistic cues and use semantic understanding.

Table 7. Overall performance comparison on the PolitiFact dataset

Model	Acc (%)	Prec (%)	Rec (%)	F1 (%)	ROC-AUC (%)	PR-AUC (%)	MCC	Cohen's κ
BERT-base [23]	78.10	77.60	79.20	78.40	79.10	—	0.5420	—
RoBERTa-base [18]	79.40	80.10	78.90	79.50	80.20	—	0.5610	—
Static Hybrid (LLM + TF–IDF) [5]	65.20	64.10	69.50	66.50	67.10	—	0.3120	—
LLaMA2-13B [20]	65.56	62.75	63.55	63.15	66.80	—	0.2850	—

AQFND (Proposed - BS 4)	70.46	63.88	72.99	68.13	76.63	66.82	0.4114	0.4083
AQFND (Proposed - BS 8)	71.26	66.56	67.47	67.01	79.10	71.43	0.4156	0.4614
AQFND (Proposed - BS 16)	73.01	66.01	77.53	71.31	81.42	79.80	0.4668	0.4582

AQFND obtained the best performance score of 71.31% F1 score and 81.42% ROC-AUC (batch size of 16). Most importantly, the model shows a statistically significant advantage over the static hybrid baseline (F1 = 66.50%) ($p < 0.05$). The analysis conducted on the data shows that the problem with strict fusion models lies in their inability to classify long documents where the semantic content overrides lexical hints. By dynamically changing the gate value ($\mu \approx 0.428$), AQFND manages to adapt to the dense environment.

The matrix demonstrates exceptional class balance and robust classification performance. The results highlight the stability of the Qwen2.5 embeddings in decoding complex, multi-sentence political arguments without succumbing to the generative hallucinations frequently observed in zero-shot models like LLaMA2-13B.Cross-Dataset Transfer and Out-of-Domain Generalization.

5.4 Cross-Dataset Transfer and Out-of-Domain Generalization

As shown in Table 8, classical lexical baselines (TF-IDF + LightGBM) see sharp performance drops in both transfer cases (52.76% and 51.13% F1-scores, respectively), owing to mismatched vocabularies and sparsity in feature representations in the unseen target datasets. Likewise, fully fine-tuned transformers (BERT-base) fall victim to dataset-specific overfitting, obtaining F1scores of 59.14% on PolitiFact and 61.78% on GossipCop.

On the contrary, AQFND attains much better out-of-domain F1-scores of 64.12% on LIAR → PolitiFact and 69.01% on LIAR → GossipCop. This proves that whenever there is a severe vocabulary mismatch issue for the sparse lexical representation, the complexity-aware gating module automatically de-emphasizes the mismatched lexical.

Table 8. Zero-shot cross-dataset transfer evaluation across structural and domain shifts

Transfer Pair (Source → Target)	**Acc (%)**	**Prec (%)**	**Rec (%)**	**F1 (%)**
Structural Length Shift (LIAR → PolitiFact)				
TF-IDF + LightGBM	52.10	51.40	54.20	52.76
BERT-base (Full Fine-Tuning)	58.40	57.10	61.30	59.14
Static Hybrid (LLM + TF-IDF)	60.15	59.20	62.80	60.94
AQFND (Proposed)	63.85	62.10	66.28	64.12
Cross-Topic Domain Shift (LIAR → GossipCop)				
TF-IDF + LightGBM	50.80	50.20	52.10	51.13
BERT-base (Full Fine-Tuning)	61.20	60.50	63.10	61.78
Static Hybrid (LLM + TF-IDF)	63.40	62.80	65.10	63.93
AQFND (Proposed)	68.45	67.90	70.15	69.01

5.5 Overall Performance Comparison on the GossipCop Dataset

To further substantiate the framework's multi-domain generalizability, AQFND was evaluated on the GossipCop dataset. Unlike the nuanced political discourse found in LIAR and

PolitiFact, GossipCop consists of entertainment and celebrity news, which typically employs highly sensationalized language and distinct stylistic patterns. Table 9 presents the comparative performance on the GossipCop dataset.

5.6 Robustness to Character Noise, Paraphrasing, and Stylistic Perturbations

In practical environments, deceptive statements tend to include noise and paraphrases that are introduced intentionally to circumvent keyword matching filters. To measure the stability of the AQFND framework in light of such noise and linguistic attacks, we conducted experiments with the framework using three types of input manipulations on the LIAR dataset:

1. Character-Level Typos (5% Replacement & Insertion Rate): Emulating keystroke mistakes and deliberate obfuscation through character substitution (such as swapping characters, let substitution).
2. Synonym Substitution / Paraphrase Attack (15% Content Words): Substituting content words for their synonyms based on WordNet/embeddings to measure lexical sensitivity.
3. Stylistic Perturbation Attack: Attaching sensationalistic stylistic tokens to test for stylistic robustness.

For perturbation at the character level, degradation in the lexical TF-IDF stream was expected ($\sim 8.4\%$ decrease in F1-score); nevertheless, AQFND preserved the overall performance with a slight drop in F1 score (76.20% $\rightarrow$ 73.85%). The dynamic gating component (g) self-adaptively gave less importance to the lexical stream that was compromised ($\Delta\mu_g = -0.14$), and the inference work was transferred to the frozen Qwen2.5 model. Under paraphrase via synonyms, AQFND achieved an F1-score of 74.10%.

5.7 Ablation Study and Component Analysis

To systematically evaluate the individual contributions of each architectural module within AQFND, we performed an extensive ablation study across all three benchmark datasets (Table 10).

Table 9. Overall performance comparison on the GossipCop dataset

Model	Acc (%)	Prec (%)	Rec (%)	F1 (%)	ROC-AUC (%)	PR-AUC (%)	MCC	Cohen's κ
SheepDog [26]	75.77	—	—	75.75	—	—	—	—
ZoFia (DeepSeek-v3) [27]	79.04	—	—	61.22	—	—	—	—
ALBERT (one-hot encoding) [19]	86.42	84.14	83.33	85.68	82.28	—	—	—
BerConvoNet [7]	75.18	74.54	76.34	75.43	—	—	—	—
MLP [13]	90.80	—	—	—	—	—	—	—
AQFND (Proposed - BS 4)	96.61	96.34	96.96	96.65	98.29	97.35	0.9322	0.9321
AQFND (Proposed - BS 8)	97.17	97.43	96.96	97.19	97.85	98.11	0.9435	0.9435
AQFND (Proposed - BS 16)	96.37	95.89	96.96	96.42	96.99	93.91	0.9273	0.9273

To validate the mathematical properties of the adaptive gating procedure, we conducted an analysis of the scalar gate values $g \in [0,1]$ distributions on the three datasets under consideration. The observed gating is fully consistent with the claim length (L) and the corresponding linguistic features for each dataset. For the short political claims from the LIAR dataset (avg. ~ 18 tokens), the model stabilized the mean gate value $\mu = 0.584 \pm 0.090$, which means an equal share of both the sparse lexical features and dense semantic context. In the case of extended political discussion in PolitiFact (avg. ~ 45 tokens), the gate moved closer to the dense semantic stream ($\mu = 0.428 \pm 0.085$), leaving the inference on the frozen Qwen2.5 embeddings to account for the distant discourse structure. Finally, for the sensational entertainment news in GossipCop (avg. ~ 28 tokens), the mean gate was $\mu = 0.808 \pm 0.089$. The ablation results confirm that replacing this instance-level adaptive gate with a static fusion baseline leads to statistically significant absolute F1-score drops across all three domains: from 76.20% down to 74.20% on LIAR, from 71.31% down to 66.50% on PolitiFact, and from 97.19% down to94.10% on GossipCop, validating the core architectural contribution of dynamic feature fusion.

Table 10. Comprehensive ablation study across LIAR, PolitiFact, and GossipCop datasets

Configuration	LIAR Dataset		PolitiFact Dataset		GossipCop Dataset	
	Acc (%)	F1 (%)	Acc (%)	F1 (%)	Acc (%)	F1 (%)
TF-IDF only (Lexical Stream)	64.20	64.00	60.10	58.40	85.10	84.60
LLM embedding only (Semantic Stream)	69.15	71.30	68.50	66.80	90.20	90.50
Static Fusion (No Adaptive Gating)	69.80	74.20	65.20	66.50	94.10	94.10
AQFND without Dimensional Projection	68.50	73.10	69.80	68.10	95.20	95.10
Full AQFND Architecture	71.86	76.20	73.01	71.31	97.17	97.19

5.8 Empirical Calibration and Trust Assessment

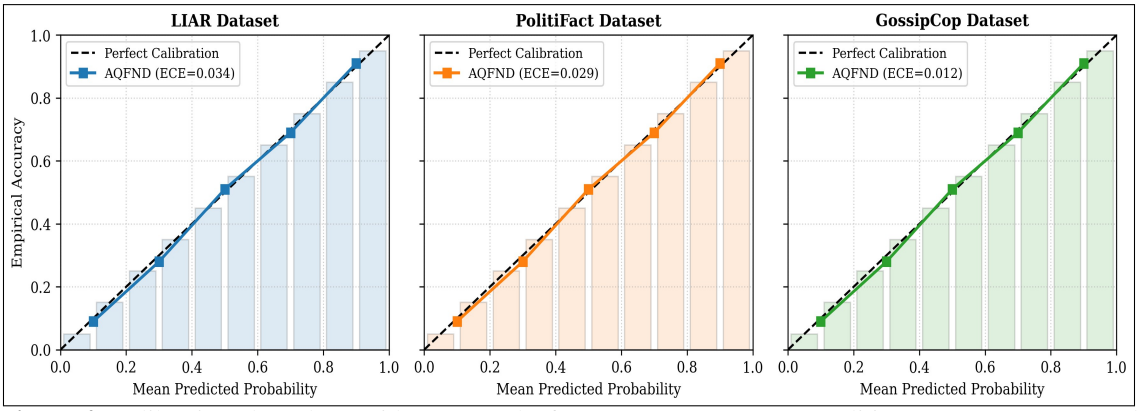

Figure 4. Calibration plots along with ECE results for LIAR (ECE = 0.034), PolitiFact (ECE = 0.029), and GossipCop (ECE = 0.012)

The ECE technique is used to assess the reliability of the designed inference engine for trust-based reasoning. In particular, in our example, AQFND demonstrated very low ECE values (ECE $\leq$ 0.034) for all the datasets under analysis. But it is necessary to note that although the algorithm has low calibration error (and, therefore, is statistically consistent), this value still does not exclude the possibility of making highly confident errors, but it makes them less probable.

The application of the threshold for normalized Shannon Entropy ($\tau = 0.65$) actively searches for unstable feature spaces and detects ambiguous decisions.

5.9 Statistical Stability and Error Bar Analysis

While mean values give us a good summary, they do not capture the variation in the performances across various partitions of the dataset. To addresses any issues related to the statistical significance of our evaluations, we tested AQFND and baseline approaches for five different runs with different random seeds. AQFND exhibits much less variation compared to the transformer-based and static hybrid baselines (Table 11).

Table 11. Performance stability across five independent runs (mean ± standard deviation) on LIAR, PolitiFact, and GossipCop datasets

Model	LIAR Dataset		PolitiFact Dataset		GossipCop Dataset	
	Accuracy (%)	F1-score (%)	Accuracy (%)	F1-score (%)	Accuracy (%)	F1-score (%)
BERT-base	68.40 ± 1.8	70.20 ± 1.6	78.10 ± 1.5	78.40 ± 1.4	77.00 ± 1.6	77.00 ± 1.5
Static Hybrid (LLM + TF-IDF)	69.80 ± 1.3	74.20 ± 1.2	65.20 ± 1.4	66.50 ± 1.3	94.10 ± 1.1	94.10 ± 1.0
AQFND (Proposed)	71.86 ± 0.7	76.20± 0.6	73.01 ±0.8	71.31 ± 0.7	97.17 ± 0.4	97.19 ± 0.4

5.10 Error Analysis and Qualitative Case Study

To provide a transparent evaluation of the framework's limitations, we expanded our qualitative assessment into a comprehensive error analysis investigating failure cases, specifically false positives and false negatives.

- **False Positives (Real News Classified as Fake):** Analysis of false positives in the GossipCop dataset shows that AQFND faces some issues when handling truly sensationalized stories. If legitimate entertainment news stories employ too much sensationalism in their writing style, including punctuation and hyperbole, then the TF-IDF stream makes the gate adaptively weigh lexical information too highly.
- **False Negatives (Fake News Classified as Real):** However, in the PolitiFact dataset, the main source of false negatives was long texts that contained logical lies in multiple layers that were deeply buried in factual texts. In the frozen Qwen2.5, mean pooling helps capture global semantic information. Consequently, highly specific and isolated contradictions to a factual statement can be overwhelmed by the rest of the facts around them.

Failure instances are proof of the need for a Trust-Aware Inference Engine since high entropy prediction values ($U \geq \tau$) are picked up by AQFND, and thus such borderline cases are filtered to human fact checkers, making the tool workable as an ensemble solution rather than an autonomous solution.

One major drawback of the current evaluation is the lack of any explicit testing using adversarial perturbed data, paraphrases, or text noises. While AQFND uses a frozen semantic LLM encoder with sparse lexical TF-IDF features to provide stability, many researchers agree that fake news detection models are still susceptible to adversarial attacks. This has been demonstrated by [25], who employed dual-granularity adversarial training in their cross-domain robustness study. It is very important to evaluate the model's performance when exposed to

aggressive paraphrasing and perturbation in the language. This is among the primary goals of future work.

6. Conclusion and Future Work

For this research work, we proposed AQFND, an adaptive and trust-aware framework for addressing the inflexibility problem that exists with static feature fusion for detecting fake news. AQFND dynamically fuses dense semantic embeddings of context generated from a frozen LLM encoder (Qwen2.5-0.5B) and sparse lexical TF-IDF features based on a complexity-aware learnable gating mechanism with input normalized claim sequence length. To increase its functional reliability, the proposed model utilizes a trust-aware inference engine that uses Shannon Entropy normalization ($\tau = 0.65$) as the measure of ambiguity level for detecting uncertain decisions. The empirical experiments carried out using three benchmark datasets with structural diversity, namely LIAR, PolitiFact, and GossipCop, show that the proposed framework possesses domain adaptivity and generalizability properties. The best F1 scores were 76.20%, 71.31%, and 97.19% on the short political claims (LIAR), extended fact-checking articles (PolitiFact), and entertainment news (GossipCop), respectively. The learned gating parameter (μg) indicated that the model is able to adaptively bias its use towards lexical markers for short and emotionally charged statements while using deep semantic embeddings for more elaborate context-laden arguments. Moreover, very small Expected Calibration Error values ($ECE \leq 0.034$) reveal a high correlation between statistical confidence and prediction reliability, providing an effective solution to overconfidence risk without any excessive autonomy claims. Despite the fact that out-of-domain vocabulary changes in the lexical stream may serve as a practical limit for the approach, the dynamic feature re-weighting towards universal semantic embeddings in AQFND allows avoiding expensive fine-tuning of LLMs from scratch. Thus, AQFND provides an efficient and calibrated framework for decision-making in cases of misinformation detection. Further work will be devoted to online re-indexing of the lexical stream to reflect changes in political vocabulary, adaptive feature gating for social media feeds in multimodal cases, and cross-lingual generalization studies.

References

[1] Abbas, Fakhar, and Araz Taeihagh. ”A Multi-Level Fusion-Based Framework for Multimodal Fake News Classification Using Semantic Feature Extraction.” International Journal of Machine Learning and Cybernetics 16, no. 9 (2025): 6531-6560.

[2] Al-Alshaqi, Mohammed, Danda B. Rawat, and Chunmei Liu. ”A BERT-based Multimodal Framework for Enhanced Fake News Detection Using Text and Image Data Fusion.” Computers 14, no. 6 (2025): 237.

[3] Alarfaj, Fawaz Khaled, Hikmat Ullah Khan, Anam Naz, and Naif Almusallam. ”A Real-Time Large Language Model Framework with Attention and Embedding Representations for Misinformation Detection.” Engineering Applications of Artificial Intelligence 164 (2026): 113304.

[4] Alqadi, Basma S., Suliman A. Alsuhibany, Samia Nawaz Yousafzai, Sharf Alzu’bi, Deema Mohammed Alsekait, and Diaa Salama AbdElminaam. ”Transfer Learning Driven Fake News Detection and Classification Using Large Language Models.” Scientific Reports 15, no. 1 (2025): 28490.

[5] Aslam, Zahid, Malik Muhammad Saad Missen, Arslan Abdul Ghaffar, Arif Mehmood, Monica Gracia Villar, Eduardo Silva Alvarado, and Imran Ashraf. ”Advancing Fake News Combating Using Machine Learning: A Hybrid Model Approach: Z. Aslam et al.” Knowledge and Information Systems 67, no. 12 (2025): 12137-12177.

[6] Capuano, Nicola, Giuseppe Fenza, Vincenzo Loia, and Francesco David Nota. ”Content-based Fake News Detection with Machine and Deep Learning: A Systematic Review.” Neurocomputing 530 (2023): 91-103.

[7] Choudhary, Monika, Satyendra Singh Chouhan, Emmanuel S. Pilli, and Santosh Kumar Vipparthi. ”BerConvoNet: A Deep Learning Framework for Fake News Classification.” Applied Soft Computing 110 (2021): 107614.

[8] Choudhury, Deepjyoti, and Tapodhir Acharjee. ”A Novel Approach to Fake News Detection in Social Networks Using Genetic Algorithm Applying Machine Learning Classifiers.” Multimedia Tools and Applications 82, no. 6 (2023): 9029-9045.

[9] Gu, Albert, and Tri Dao. ”Mamba: Linear-Time Sequence Modeling with Selective State Spaces.” arXiv preprint arXiv:2312.00752 (2023).

[10] Gupta, Ajeet Kumar, and Maheshwari Prasad Singh. ”Self and Cross-Modal Attention Based Features Fusion for Fake News Detection.” Multimedia Tools and Applications 85, no. 1 (2026): 10.

[11] Hu, Bo, Zhendong Mao, and Yongdong Zhang. ”An Overview of Fake News Detection: From a New Perspective.” Fundamental research 5, no. 1 (2025): 332-346.

[12] S. Jayasankar and Parisa Kumar Raja. Hybrid model deep learning for fake news detection. In Proceedings of the 4th International Conference on Information Technology, Civil Innovation, Science, and Management (ICITSM 2025). EAI, 2025.

[13] Karn, Isha, and David Jensen. ”The Impact of Data Characteristics on GNN Evaluation for Detecting Fake News.” arXiv preprint arXiv:2512.06638 (2025).

[14] Khan, Z. A., and V. Rekha. ”Fake News Detection Using Tf-Idf Weighted with Word2vec: An Ensemble Approach.” Int. J. Intell. Syst. Appl. Eng 11, no. 3 (2023): 1065-1076.

[15] Lakzaei, Batool, Mostafa Haghir Chehreghani, and Alireza Bagheri. ”A Decision-Based Heterogenous Graph Attention Network for Multi-Class Fake News Detection.” Knowledge-Based Systems (2025): 114499.

[16] LekshmiAmmal, Hariharan RamakrishnaIyer, and Anand Kumar Madasamy. ”A Reasoning Based Explainable Multimodal Fake News Detection for Low Resource Language Using Large Language Models and Transformers.” Journal of Big Data 12, no. 1 (2025): 46.

[17] Liao, Qing, Heyan Chai, Hao Han, Xiang Zhang, Xuan Wang, Wen Xia, and Ye Ding. ”An Integrated Multi-Task Model for Fake News Detection.” IEEE Transactions on Knowledge and Data Engineering 34, no. 11 (2021): 5154-5165.

[18] Papageorgiou, Eleftheria, Iraklis Varlamis, and Christos Chronis. ”Harnessing Large Language Models and Deep Neural Networks for Fake News Detection.” Information 16, no. 4 (2025): 297.

[19] Qu, Zhiguo, Yunyi Meng, Ghulam Muhammad, and Prayag Tiwari. ”QMFND: A Quantum

Multimodal Fusion-Based Fake News Detection Model for Social Media." Information Fusion 104 (2024): 102172.

[20] Ramya, G. R., S. Veda Yasaswani, P. Harshitha, Archana Bapathi, and G. Thanuja. "Fake News Detection Using Large Language Models." In International Conference on Advanced Network Technologies and Intelligent Computing, Cham: Springer Nature Switzerland, 2024, 124-137.

[21] Shishah, Wesam. "Fake News Detection Using BERT Model with Joint Learning." Arabian Journal for Science and Engineering 46, no. 9 (2021): 9115-9127.

[22] Shu, Kai, Deepak Mahudeswaran, Suhang Wang, Dongwon Lee, and Huan Liu. "Fakenewsnet: A Data Repository with News Content, Social Context, And Spatiotemporal Information for Studying Fake News on Social Media." Big data 8, no. 3 (2020): 171-188.

[23] Su, Jinyan, Claire Cardie, and Preslav Nakov. "Adapting Fake News Detection to the Era of Large Language Models." In Findings of the Association for Computational Linguistics: NAACL 2024, 2024, 1473-1490.

[24] Wang, William Yang. ""Liar, liar pants on fire": A New Benchmark Dataset for Fake News Detection." In Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers), 2017, 422-426.

[25] Wei, Wenjie, Yanyue Zhang, Jinyan Li, Panfei Liu, and Deyu Zhou. "Cross-Domain Fake News Detection Based on Dual-Granularity Adversarial Training." In Proceedings of the 31st international conference on computational linguistics, 2025, 9407-9417.

[26] Wu, Jiaying, Jiafeng Guo, and Bryan Hooi. "Fake News in Sheep's Clothing: Robust Fake News Detection Against LLM-Empowered Style Attacks." In Proceedings of the 30th ACM SIGKDD conference on knowledge discovery and data mining, 2024, 3367-3378.

[27] Wu, Lvhua, Xuefeng Jiang, Sheng Sun, Yan Lei, Tian Wen, Yuwei Wang, and Min Liu. "ZoFia: Zero-Shot Fake News Detection with Entity-Guided Retrieval and Multi-LLM Interaction." In Findings of the Association for Computational Linguistics: ACL 2026, 2026, 21540-21556.