



RESEARCH ARTICLE

OPEN ACCESS

MULTIMODAL META-LEARNING FOR EARLY RUMOR DETECTION WITH ADAPTIVE DECISION TIMING

El Yazid Gueddoudj^{*1}, Abdelouahab Attia³ and Adel Moussaoui³

¹Institute of Architecture and Urbanism, University of Batna1, Batna 05000, Algeria

¹LRIT laboratory, University of Tlemcen, Tlemcen 13000, Algeria

²Department of Computer Science, University Mohamed El Bachir El Ibrahimi of Bordj Bou Arreridj, 34000 Bordj Bou Arreridj, Algeria

³Faculty of Mathematics and Computer Science, University of M'sila, PO Box 166 Ichebilia, 28000, M'sila, Algeria.

¹<https://orcid.org/0009-0005-1462-8477>, ²<https://orcid.org/0000-0003-1558-7273>, ³<http://orcid.org/0000-0002-6978-1656>

E-mail: *elyazid.gueddoudj@univ-batna.dz, abdelouahab.attia@univ-bba.dz, adel.moussaoui@univ-msila.dz

ARTICLE INFO

ABSTRACT

Article History

Received on: March 28, 2026

Reviewed on: April 29, 2026

Accepted on: July 01, 2026

Published on: July 31, 2026

Keywords:

Rumor Detection,
Early Detection,
Metalearning,
Social Media,
Adaptive Decision Timing.

Early social media rumor detection is essential because it prevents false information from spreading quickly and causing serious harm. Current methods frequently ignore the significance of prompt decision-making in favor of increasing classification accuracy. In this paper, we present a novel framework for early rumor detection that is based on multimodal meta-learning with adaptive decision timing. The method analyzes social media posts as an ongoing stream and learns when enough data has been observed to dynamically identify the best time to make a forecast. By integrating textual, structural, and temporal features, the proposed approach captures complex rumor propagation patterns while jointly optimizing accuracy and detection timeliness. Experimental results on benchmark datasets, including Twitter15, Twitter16, and Weibo, show that the proposed method significantly reduces detection latency, requiring substantially fewer messages for prediction, while maintaining competitive or superior performance compared to state-of-the-art baselines.



Copyright ©2026 by authors and Galileo Institute of Technology and Education of the Amazon (ITEGAM). This work is licensed under the Creative Commons Attribution International License (CC BY 4.0).

I. INTRODUCTION

The danger of false information and rumors spreading has greatly grown due to the quick spread of information on social media platforms. Rapid and widespread rumors can have detrimental effects on society, such as economic disruption, public panic, and disinformation during important events [1]. As a result, creating efficient methods for early rumor detection has emerged as a crucial research task. Existing approaches to rumor detection have achieved notable success by leveraging machine learning and deep learning techniques. Traditional methods such as [2], [3] and [4] rely on handcrafted features and static classifiers, while more recent approaches employ sequential models such as recurrent neural networks and Transformers to capture temporal dependencies in message propagation. Despite these advancements, most existing methods primarily focus on improving classification accuracy and typically assume access to complete or near-complete event sequences. However, prompt detection is just as crucial as precision in practical situations. Rumors might spread far before any action is taken due to delayed detection. Early detection is intrinsically difficult because there is little to no information available at the beginning of an occurrence. Furthermore, the best time to make a trustworthy prediction varies based on the dynamics of each event's propagation. We provide a novel framework based on multimodal meta-learning [5] with adaptive decision timing for early rumor detection in order to overcome these drawbacks. In the suggested method, rumor detection is modeled as a sequential decision-making issue. Incoming messages are processed as a continuous stream by the model, which dynamically determines when enough evidence has been observed to make a prediction. By considering multiple types of information, such as textual content, structural propagation patterns and temporal dynamics, the framework is able to capture comprehensive features of social media events. Furthermore, a meta-learning strategy is applied to facilitate generalization to unseen rumor types, which allows the model to generalize to different propagation patterns. The core of the proposed method is the adaptive decision process, which lets the model automatically find the earliest confident decision time. The proposed model performs early and accurate rumor detection, greatly reducing detection latency while maintaining comparable or higher classification performance compared to existing methods, according to extensive trials carried out on benchmark datasets. The following is a summary of the primary contributions:

1. We propose a multimodal meta-learning framework. It combines textual, structural, and temporal features.
2. We present an adaptive decision timing mechanism that finds the best prediction point.
3. We ran many tests. They show that our method finds rumors earlier. It needs fewer observations and still has strong classification performance.

The rest of the paper is organized as follows: Section 2 surveys prior work, and Section 3 provides the proposed method and the new framework. Section details the Experimental and the findings. Section 5 provides a discussion of the experimental findings and their analysis. Lastly, Section 6 concludes the paper and future work.

II. RELATED WORK

II.1 RUMOR DETECTION ON SOCIAL MEDIA

The development of manually created characteristics for machine learning algorithms has historically been the primary focus of research on rumor detection. Using both conventional machine learning and deep learning methods, rumor detection on social media has been thoroughly investigated [6]. Early approaches use classifiers like Support Vector Machines and Random Forests in conjunction with manually created features taken from textual content, user profiles, and propagation structures [7]. Although these methods offer preliminary insights, feature engineering and their incapacity to identify intricate patterns restrict their effectiveness. In more recent work, representations are automatically learned from data using deep learning algorithms [8]. The temporal dynamics of message sequences have been extensively modeled using recurrent neural networks, such as LSTM and GRU. By using self-attention methods to capture long-range dependencies, transformer-based models significantly enhance performance. Despite these developments, the majority of current methods, which frequently assume access to full propagation sequences, are primarily concerned with increasing classification accuracy.

II.2 EARLY RUMOR DETECTION

The goal of early rumor detection is to identify rumors when they are just beginning to spread and there is little information available. A number of methods make predictions based on incomplete sequences or preliminary observations in an effort to overcome this difficulty. In order to effectively capture global semantic relations in rumor texts, the authors of [9] built a tweet-word-user heterogeneous graph and proposed a meta-path-based heterogeneous graph attention network framework that uses subgraph attention networks to learn global node relationships and a subgraph-level attention mechanism to fuse node representations. The authors of [10] used transfer learning to extract attitude representations from external stance-classification datasets by utilizing pretrained CNN and BERT models. For early rumor detection, their deep model (CNN-BiLSTM-BERT) combines attitude and content representations. Authors of [11] proposed a deep attention model based on RNNs with soft-attention. Without depending on manually created features, the model accurately identifies long-range dependencies and contextual changes over time. It detects rumors faster and more correctly than state-of-the-art baselines, according to extensive trials on real-world social media datasets. [12] presents a few-shot prototype augmentation framework for low-resource rumor detection. In addition to a metric learning methodology, it uses hierarchical prompt encoding to obtain language/domain-agnostic thread representations and refines prototypes via query-conditioned attention and feature gating. The method is compatible with parallelized streaming settings and achieves cutting-edge cross-lingual/domain performance. To produce predictions at intermediate time steps, sequential models have been modified, as in [13].

Accuracy and timeliness must be balanced, according to recent research, but most methods tackle these goals separately rather than working together to maximize both. A recent approach [14] have integrated timeliness and community structure. It uses a hybrid sensor deployment for effective information gathering, simulates dynamic propagation with a time-varying parameter, and uses maximum likelihood estimation within suspicious communities to identify one or more sources. This is the first framework for community-structured dynamic OSNs, and experiments on both real-world and synthetic networks show correct identification. The authors in [15] have presented an early reinforcement learning-based rumor detection model that extracts characteristics from original tweets and responses using a dual-engine deep learning architecture and double self-attention. While optimizing a reward function that strikes a compromise between timeliness and accuracy, a Deep Recurrent Q-Learning Network (DRQN) learns state sequence features. Nevertheless, a lot of these techniques rely on fixed time intervals or predetermined observation windows, which fail to take into consideration the variation in rumor spreading patterns among various events. As a result, they might either postpone forecasts needlessly or make snap judgments based on incomplete knowledge. [16] presents ERD-DC, a dual-channel early rumor detection model that differentiates between "difficult-to-identify" rumors and "easy-to-identify" rumors. While a slow channel does thorough analysis to ensure overall correctness, a fast channel prioritizes timeliness by performing quick post-evidence matching for simple cases. State-of-the-art performance is demonstrated by experiments on real datasets, which successfully balance detection speed and reliability and achieve an accuracy rate of 84%.

II.3 MULTIMODAL LEARNING FOR RUMOR DETECTION

By combining several information sources, multimodal techniques have been put forth to enhance rumor detection. These include temporal elements that depict how events change over time, as well as textual content, user interactions, and propagation structures. The multimodal rumor detection system presented in [17] introduces trainable input parameters while maintaining a frozen backbone, utilizing quick learning in both the vision and language domains. Early on, different prompts gradually model feature interactions, improving context learning. Superior performance in differentiating between rumors and facts is demonstrated by experiments conducted on real-world multimodal datasets. In order to improve the propagation structure and obtain useful local video segments, [18] uses time-sync comments to detect video rumors. It integrates multimodal characteristics for categorization from both local and global keyframes, subtitles, and comments.

In order to address multimodal misinformation detection, [19] presents CPLS (Contrastive Prompt Learning in Structured Graph Networks), a model that builds structured graph networks with semantic-aware pruning and fusing. Accuracy gains of 2.4% on Weibo and 4.5% on Twitter are demonstrated in experiments conducted on public datasets. In recent work [20] presents LoRAD (Logic-Reasoning Augmented DeBERTa). It uses DeBERTa's disentangled attention method to combine logical evidence produced by an LLM with the original text. Evaluated on three publicly available datasets as well as a recently created Twitter26-Mini dataset, LoRAD outperforms LLMs in terms of robustness and inference speed, making it appropriate for real-time rumor detection. These methods produce stronger representations and better classification results by merging complimentary modalities. Nevertheless, the problem of early detection in streaming contexts is not specifically addressed by current multimodal models, which are usually built for static or fully observed data. A successful model for enhancing generalization across tasks is meta-learning [21]. Meta-learning allows models to swiftly adjust to new patterns with few observations in the setting of sequential data. Although meta-learning has been used in many different fields, it is still not widely used for rumor detection. Its potential for simulating various rumor dissemination dynamics or for enhancing early decision-making is hardly explored in existing studies. For instance, a meta-learning framework was developed by [22] that jointly learns social bot detection as an auxiliary job to improve rumor detection. The model ensures that information from the auxiliary task efficiently serves the main goal by using an outer-loop optimization on rumor detection and an inner-loop pseudo-update on bot detection. The findings of the experiment show excellent performance, reaching 90% accuracy in just two hours.

II.4 META-LEARNING IN SEQUENTIAL MODELING

Additionally, there is still much to learn about the potential of meta-learning in the field of rumor detection. In order to fill these deficiencies, this paper presents a unified framework that combines adaptive decision timing and multimodal meta-learning to enable early and accurate rumor identification in dynamic information contexts. Using the ISOT dataset [23], [24] proposed the application of meta-learning in conjunction with BERT for the detection of fake news. While highlighting constraints and the significance of responsible AI design, an ensemble transformer-based model achieves a high accuracy of 98.29%, showcasing the promise of meta-learning in enhancing detection performance. Meanwhile, [25] proposes a multimodal meta-learning method based on few-shot learning (FSL) for early rumor detection, addressing the challenges of limited spread and small sample sizes. It uses a meta-learning method with quick, multi-step updates to improve flexibility while designing a multimodal feature extraction layer, a hidden information extraction layer, and a fusion output layer. In conclusion, even though previous studies have shown significant improvements in multimodal learning, sequential modeling, early detection, and rumor detection, a number of important constraints still exist. Current approaches lack support for streaming decision-making pipelines, rely on fixed observation windows that limit adaptation to dynamic information dispersion, and primarily prioritize classification accuracy at the expense of decision speed.

III. METHODOLOGY

In this section, we present the proposed framework. The model handles social media events as a series of data streams. It then finds the best moment to make a prediction. We start by stating the problem. Next, we outline the framework, model architecture, and the proposed algorithm.

III.1 PROBLEM STATEMENT

Over a stream of social media posts, we formulate early rumor detection as a sequential decision-making issue. Let an event be represented as a chronologically ordered series of communications, such as a discussion thread or a collection of posts pertaining to a particular claim:

$$E = \{m_1, m_2, \dots, m_T\} \quad (1)$$

Where m_t denotes the $t - th$ message arriving over time. Only the partial observation is available to the model at each time step t . $E_t = \{m_1, m_2, \dots, m_t\}$ and must decide whether to: Continue observing further messages, or Stop and assign a binary classification (0: non-rumor, 1: rumor) to $y \in \{0, 1\}$. The goal is to reduce both the detection delay and the classification error simultaneously. This is expressed formally as:

$$\mathcal{L} = \alpha \cdot \mathcal{L}_{cls} + \beta \cdot \mathcal{L}_{time} \quad (2)$$

where:

- $\mathcal{L}_{cls}$: classification loss (cross-entropy)
- $\mathcal{L}_{time} = \frac{t}{T}$: normalized detection delay
- α, β : trade-off parameters

The system aims to achieve the earliest possible reliable detection by learning a policy that maps partial observations to halting choices and predictions. This allows for prompt action while maintaining high classification fidelity.

III.2 OVERALL FRAMEWORK

The suggested framework is intended for real-time rumor detection in social media streams, where posts are constantly coming in and judgments need to be taken as soon as possible to stop the propagation of false information. Our proposed framework uses a streaming paradigm that gradually consumes each incoming post and dynamically determines the best time to make a forecast, in contrast to traditional batch-based methods that process predefined windows of data. By striking a balance between two conflicting goals maximizing detection accuracy and decreasing decision latency the main purpose is to put an end to rumors. The design consists of three closely related modules:

Adaptive Decision, Meta Learning, and Multimodal Representation Learning, as shown in Figure 1. A new social media post is encoded at each time step t , the internal model state is updated, and a confidence-based decision policy decides if there is enough evidence to categorize the material as rumor or not. Each component is described in detail in the ensuing subsections.




Figure 1: The proposed framework architecture.

Source: Authors, (2026).

1. Multimodal Representation Module allows for incremental state updates by encoding each incoming post into a single embedding that combines textual, structural, and temporal data. Each message m_t is represented using multiple modalities:

Textual Features: Extracted using: BERT / LSTM encoder, $h_t^{text} = f_{text}(m_t)$

Structural Features: Propagation patterns (retweets, replies) with user credibility, $h_t^{struct} = f_{struct}(m_t)$

Temporal Features: Time intervals between messages with posting dynamics, $h_t^{temp} = f_{temp}(t)$

Multimodal Fusion: The final representation is obtained via fusion: $h_t = Fusion(h_t^{text}, h_t^{struct}, h_t^{temp})$, with fusion strategies.

2. Meta-Learning Module recognizes when enough evidence has been observed by quickly adjusting the internal state at each timestep after learning a previous over rumor detection tasks. The model is trained using a meta-learning strategy where each event is treated as a task: $\theta^* = \arg \min_{\theta} \sum_{E \sim D} \mathcal{L}_E(\theta)$. This enables the model to generalize across different rumor propagation patterns.

3. Adaptive Decision Module employs an ideal stopping strategy that dynamically chooses the predicted time to strike a compromise between timeliness and accuracy. At each timestep t , the model computes: The classification probability: $p_t = P(y | E_t)$ and the confidence score: $c_t = \max(p_t)$.

III.3 ADAPTIVE DECISION ALGORITHM FOR EARLY RUMOR DETECTION

The proposed algorithm (Algorithm 1) adds a number of innovative features that together improve early rumor detection. It learns adaptive decision timing to determine the best time to make a prediction per event, in contrast to conventional approaches that depend on full event observation or fixed-length inputs. This is accomplished by using multimodal meta-learning, which permits quick adaptation to changing dynamics and generalization across various propagation patterns. In order to represent real-world social media circumstances, the problem is stated as a sequential streaming process where judgments are made progressively as posts arrive. While event-specific adaptation enables the system to act quickly on clear cases and collect additional evidence when ambiguity persists, a joint optimization of accuracy and timeliness explicitly incorporates detection delay into the goal, striking a principled balance between early intervention and reliable classification.

Algorithm 1 Early Rumor Detection with Adaptive Decision Timing

Input: Event stream $E = \{m_1, \dots, m_T\}$

Output: Predicted label y , detection time t^*

Initialize model parameters θ

$t \leftarrow 1$

while $t \leq T$ do

 Receive message m_t

 Encode multimodal features:

$$h_t = Fusion(f_{text}(m_t), f_{struct}(m_t), f_{temp}(t))$$

 Update representation using previous context

 Compute prediction:

$$p_t = P(y | E_t)$$

$$c_t = \max(p_t)$$

if $c_t \geq \tau$ then

 return prediction y and detection time t

end if

$$t \leftarrow t + 1$$

end while

Return final prediction using full sequence

IV. EXPERIMENTS

The efficacy of the proposed framework for early rumor identification is assessed in this section. The datasets, baselines, and assessment measures are described first, and then the implementation specifics. We then present the experimental results in terms of both classification performance and detection timeliness, along with ablation studies to analyze the contribution of each component.

IV.1 DATASETS

We assess the proposed framework using three popular benchmark datasets [26] for rumor detection: Weibo [27], Twitter15, and Twitter16 [28]. These datasets characteristics (Table 1) include social media events that have been tagged as rumors or not, along with the propagation threads that go along with them. A series of replies and reposts follow the first post in each event, creating a temporal structure appropriate for early detection analysis. Twitter15 and Twitter16 have a variety of dissemination patterns. Weibo is a sizable dataset that was gathered from a Chinese social media site and has a notably greater quantity of posts and events. It makes it possible to assess cross-domain robustness by offering more diversity in language, user behavior, and propagation dynamics.

Table 1: Dataset Characteristics.

Dataset	Events	Posts	Classes
Twitter15	1,490	331,612	2
Twitter16	818	204,820	2
Weibo	4,664	3,805,656	2

Source: Authors, (2026).

We use conventional preparation procedures for all datasets: tokenization, text normalization, and the removal of URLs, mentions, and non-informative tokens. In order to accurately replicate the streaming environment necessary for early rumor detection, messages inside each event are arranged chronologically, treating each event as a time-ordered sequence.

IV.2 BASELINES

We evaluate the proposed framework against a wide range of baselines that include early detection, propagation-aware, sequential deep learning, and conventional machine learning. Conventional approaches include a propagation tree kernel-based SVM PTK [26] that captures structural similarities between rumor propagation trees and a Random Forest classifier RFC [29] that uses manually created textual, user, and structural variables. LSTM [30], GRU [31], and Transformer are sequential deep learning models that use whole sequence observation to capture temporal dependencies. PPC [32], a hybrid RNN-CNN architecture that learns from propagation paths and user features, represents propagation-aware deep models. We use early LSTM, a time-aware variation that generates predictions at predetermined intermediate steps, for early detection.

IV.3 EVALUATION METRICS

We evaluate the proposed method utilizing timeliness-aware metrics as well as classification performance metrics to evaluate both prediction accuracy and early rumor detection. **Metrics for Classification:** We employ common metrics to assess prediction performance:

- Accuracy: the total percentage of events that are correctly classified
- Precision: the percentage of accurately anticipated rumors among positives
- Recall: the percentage of true rumors that are accurately identified
- F1-score: precision and recall harmonic mean
- Area under the ROC curve, or AUC

Timeliness Metrics: We utilize the following metrics to evaluate early detection capability:

- Average Detection Time (ADT): $ADT = \frac{1}{T} \sum_{i=t}^N t_i^*$, Measures the average number of messages required before making a decision.
 - Normalized Detection Time (NDT): $NDT = t^*/T$, Captures the relative detection delay with respect to the full event length.
- Early Detection Rate (EDR): Proportion of rumors correctly detected before a predefined observation threshold. Messages Used (%), $t^*/T \times 100$, this indicates the fraction of messages required for prediction.

IV.4 IMPLEMENTATION DETAILS

The proposed framework is trained on a workstation with an Intel Core i9 processor, 32 GB of RAM, and an NVIDIA GPU. It is implemented in Python 3.9 using PyTorch 1.12.1. To guarantee consistent and repeatable training, we employ the Adam optimizer [13] with standard settings, $\beta_1 = 0.9$ and $\beta_2 = 0.999$. The learning rate is set to 5×10^{-4} . To mitigate overfitting, dropout is applied with a probability of 0.5, and weight decay is set to 5×10^{-4} .

The attention parameters for Twitter15 and Twitter16 are $K_1 = 4$ and $K_2 = 2$. These parameters are raised to $K_1 = 8$ and $K_2 = 4$ for the Weibo dataset, which has longer propagation sequences and denser user interactions. To enhance generalization across various events, the meta-learning component is trained end-to-end in tandem with the primary model. We use a five-fold cross-validation approach, dividing the data into training, validation, and test sets at 80%, 10%, and 10% ratios in each step.

IV.5 OVERALL PERFORMANCE RESULTS


In terms of standard classification performance across all datasets, we compare the proposed model with baseline methods. The performance of the suggested framework and the baselines on Twitter 15, Twitter 16, and Weibo are shown in Tables 2.

Table 2: The overall performance comparison on datasets.


Method	Twitter15 (F1)	Twitter16 (F1)	Weibo (F1)	AUC
PTK	0.76	0.74	0.78	0.80
RFC	0.74	0.72	0.75	0.78
PPC	0.85	0.83	0.87	0.89
LSTM	0.81	0.79	0.83	0.85
GRU	0.83	0.81	0.85	0.87
Transformer	0.86	0.84	0.88	0.90
Early LSTM	0.84	0.82	0.86	0.88
Proposed	0.88	0.86	0.90	0.93

Source: Authors, (2026).


As shown in Figure 2, the findings show that the suggested framework maintains good predictive performance while successfully capturing intricate rumor dissemination patterns. Specifically, the suggested model demonstrates the efficacy of integrating multimodal learning with adaptive decision time, outperforming propagation-based models like PPC and PTK.




(a) Accuracy Twitter 15.




(b) F1 Score Twitter 15.




(c) Accuracy Twitter 16.



(d) F1 Score Twitter 16.



(e) Accuracy Weibo.



(f) F1 Score Weibo.

Figure 2: Results of the proposed structure.

Source: Authors, (2026).

IV.6 ACCURACY VS TIME

The proposed model achieves excellent classification performance at early observation stages, as shown by the accuracy progression study. In comparison to baseline approaches, which require far more observations to achieve equivalent performance, the framework achieves competitive accuracy across all investigated datasets after processing only a small fraction of the message sequence. Plotting accuracy as a function of the percentage of seen messages. Figure 3 shows this trend and consistently demonstrates the advantage of the proposed model in early detection phases.




Figure 3: Accuracy vs Time across Multiple Datasets.
Source: Authors, (2026).

IV.7 DETECTION TIMELINESS

In this section, we assess how well the suggested model predicts the early phases of the message propagation process. To provide a more thorough comparison, we incorporate propagation-aware models like PPC in addition to traditional baselines. As described in Table 3 the proposed framework demonstrates its efficacy for early rumor detection by achieving the maximum early detection rate (EDR), the lowest detection delay ADT (Average Detection Time) NDT (Normalized Detection Time), and the fewest messages.

Table 3: The comparison of detection timeliness across methods.

Model	ADT	NDT	Messages Used	EDR
LSTM	18.5	0.92	100%	0.65
GRU	17.2	0.89	100%	0.67
Transformer	16.8	0.87	100%	0.69
PPC	13.6	0.66	78%	0.75
Early LSTM	12.4	0.61	72%	0.78
Proposed Model	8.1	0.39	48%	0.87

Source: Authors, (2026).




Figure 4: Detection timeliness comparison across methods.
Source: Authors, (2026).

As shown in Figure 4, across all assessed criteria, the proposed model significantly improves detection timeliness. Less than half of the available sequence length is used, and ADT and NDT are greatly decreased, indicating that judgments are taken earlier in the event lifecycle with noticeably fewer input messages on average. A greater EDR further reflects this efficiency, indicating the model's ability to correctly detect rumors in their early phases. However, conventional methods, on the other hand, rely on almost full observation sequences, resulting in significant delays that restrict their usefulness for real-time applications. A boxplot comparison of detection times for several models is shown in Figure 5. In comparison to baseline techniques, the suggested model has a much lower median and a more compact interquartile range, suggesting both faster and more stable detection behavior.




Figure 5: Boxplot comparison of detection times across different models.
Source: Authors, (2026).

IV.8 ABLATION STUDY

We perform an ablation research by eliminating important modules from the suggested framework in order to assess each component's contribution the results are shown in Table 4. Figure 6 shows how each component affects F1-score and detection time.

Table 4: Ablation results.

Variant	F1	ADT	Messages Used (%)
Full Model	0.88	8.1	48
w/o Meta-Learning	0.84	12.9	71
w/o Multimodal	0.82	13.5	74
w/o Adaptive Timing	0.86	18.0	100

Source: Authors, (2026).




Figure 6: Performance vs Detection time.
Source: Authors, (2026).

The ablation study validates each module's contribution by confirming that eliminating any one component causes a discernible decline in performance. In particular, removing meta-learning slows detection and forces the model to rely on more observations, which reflects less generalization across heterogeneous rumor events. The classification performance significantly declines when multimodal characteristics are removed, highlighting the complimentary benefit of combining textual, structural, and temporal information. Lastly, the model becomes a standard sequential classifier when adaptive decision timing is disabled, requiring entire sequence observation and resulting in the largest detection delay. These results verify that each component tackles a unique and non-redundant part of the detection problem, and that only when all modules work together can optimal performance be attained.

V. DISCUSSION

The experimental results show that the proposed framework successfully strikes a balance between classification performance and detection timeliness, two goals that are often addressed separately in previous research. By dynamically modifying the decision point based on prediction confidence, the adaptive decision timing method significantly reduces detection delay without sacrificing accuracy, producing more realistic behavior in streaming contexts. By combining textual, structural, and temporal information, multimodal feature integration further improves rumor representation, and the meta-learning approach improves generalization across diverse events by recording common dissemination patterns. When taken as a whole, these elements promote early rumor detection as a joint optimization problem where timeliness and accuracy are viewed as complementing goals. However, some limitations should be acknowledged, such as the computational complexity of multimodal processing and meta-learning in large-scale real-time scenarios, and the sensitivity of the preset confidence threshold across datasets with different distributional features.

VI. CONCLUSION

In this paper, we proposed a framework for early rumor detection that integrates multimodal representation learning with meta-learning and adaptive decision timing, modeling the problem as a sequential decision-making process over streaming input. The value of adaptive stopping in real-world intervention scenarios is highlighted by experimental results on Twitter15, Twitter16, and Weibo, which show that our approach significantly reduces detection latency while maintaining competitive or superior classification performance compared to state-of-the-art baselines. In the future, the framework will be expanded to large-scale streaming settings, richer modalities including images and user interaction graphs will be included, and reinforcement learning for completely adaptive halting rules will be investigated. To further advance the integration of meta-learning, multimodal analysis, and adaptive decision-making for next-generation rumor detection systems, we also want to look into uncertainty estimation for more accurate early predictions and cross-platform domain adaptation for wider generalization.

VII. AUTHOR'S CONTRIBUTION

Conceptualization: El Yazid Gueddoudj.

Methodology: El Yazid Gueddoudj and Abdelouahab Attia.

Investigation: El Yazid Gueddoudj, Abdelouahab Attia and Adel Moussaoui.

Discussion of results: El Yazid Gueddoudj, Abdelouahab Attia and Adel Moussaoui.

Writing – Original Draft: El Yazid Gueddoudj and Abdelouahab Attia.

Writing – Review and Editing: El Yazid Gueddoudj and Abdelouahab Attia.

Resources: El Yazid Gueddoudj.

Supervision: El Yazid Gueddoudj, Abdelouahab Attia and Adel Moussaoui.

Approval of the final text: El Yazid Gueddoudj, Abdelouahab Attia and Adel Moussaoui.

VIII. REFERENCES

- [1] M. Shahbazi and D. Bunker, "Social media trust: Fighting misinformation in the time of crisis," *International Journal of Information Management*, vol. 77, p. 102780, 2024.
- [2] H. Ran, C. Jia, P. Zhang, and X. Li, "MGAT-ESM: Multi-channel graph attention neural network with event-sharing module for rumor detection," *Information Sciences*, vol. 592, pp. 402–416, 2022.
- [3] C.-Y. Li, "Persuasive messages on information system acceptance: A theoretical extension of elaboration likelihood model and social influence theory," *Computers in human behavior*, vol. 29, no. 1, pp. 264–275, 2013.
- [4] B. G. Robbins, "From the general to the specific: How social trust motivates relational trust," *Social Science Research*, vol. 55, pp. 16–30, 2016.
- [5] A. Vettoruzzo, M.-R. Bouguelia, and T. Rögndalsson, "Multimodal meta-learning through meta-learned task representations," *Neural Computing and Applications*, vol. 36, no. 15, pp. 8519–8529, 2024.
- [6] B. Janssens, L. Schetgen, M. Bogaert, M. Meire, and D. den Poel, "360 Degrees rumor detection: When explanations got some explaining to do," *European Journal of Operational Research*, vol. 317, no. 2, pp. 366–381, 2024.
- [7] Z. Zhao, P. Resnick, and Q. Mei, "Enquiring minds: Early detection of rumors in social media from enquiry posts," in *Proceedings of the 24th international conference on world wide web*, 2015, pp. 1395–1405.
- [8] W. Zhang et al., "Robust rumor detection against noise," *Neurocomputing*, p. 132741, 2026.
- [9] Q. Huang, J. Yu, J. Wu, and B. Wang, "Heterogeneous graph attention networks for early detection of rumors on twitter," in *2020 international joint conference on neural networks (IJCNN)*, 2020, pp. 1–8.
- [10] L. Tian, X. Zhang, Y. Wang, and H. Liu, "Early detection of rumours on twitter via stance transfer learning," in *European Conference on Information Retrieval*, 2020, pp. 575–588.
- [11] T. Chen, X. Li, H. Yin, and J. Zhang, "Call attention to rumors: Deep attention based recurrent neural networks for early rumor detection," in *Pacific-Asia conference on knowledge discovery and data mining*, 2018, pp. 40–52.
- [12] Y. Han, X. Li, H. Ran, and Y. Feng, "Few-shot with prototype augmentation for low-resource domains rumor detection," *The Journal of Supercomputing*, vol. 82, no. 2, p. 52, 2026.
- [13] L. Zhang, Y. Chen, X. Luo, H. Zhu, and G. He, "Early detection of rumors based on propagation prediction in social media," *Complex & Intelligent Systems*, vol. 12, no. 1, p. 23, 2026.
- [14] D.-W. Huang, W. Wu, J. Bi, J. Li, C. Gan, and W. Zhou, "Timeliness-aware rumor sources identification in community-structured dynamic online social networks," *Information Sciences*, vol. 689, p. 121508, 2025.
- [15] W. Wang, Y. Qiu, S. Xuan, and W. Yang, "Early Rumor Detection Based on Deep Recurrent Q-Learning," *Security and Communication Networks*, vol. 2021, no. 1, p. 5569064, 2021.
- [16] Y. Wu, J. Sun, X. Yuan, Z. Huang, and J. Dai, "@inproceedings{chen2024multimodal, title={Multimodal rumor detection via multi modal prompt learning}, author={Chen, Fan and Li, Xiuhong and Li, Zhe and Zhou, Chenyu and Sheng, Jiabao}, booktitle={2024 International Joint Conference on Neural Networks (IJCNN)}, Expert Systems with Applications, vol. 238, p. 121928, 2024.
- [17] X. Shen, M. Huang, Z. Hu, S. Cai, and T. Zhou, "Multimodal Fake News Detection with Contrastive Learning and Optimal Transport," *Frontiers in Computer Science*, vol. 6, p. 1473457, 2024.

- [18] M. Yin, W. Chen, D. Zhu, and J. Jiang, "Enhancing video rumor detection through multimodal deep feature fusion with time-sync comments," *Information Processing & Management*, vol. 62, no. 1, p. 103935, 2025.
- [19] Y. Guo, K. Zhen, and J. Liu, "Contrastive Prompt Learning in Structured Graph Networks for Multimodal Fake News Detection," *IEEE Transactions on Big Data*, 2026.
- [20] Y. Zhang, F. Muchtar, M. K. Mohd Sidik, and J. Mohamad Sharif, "LoRAD: Logic-Reasoning Augmented DeBERTa with Adversarial Training for Robust Rumor Detection," *Electronics*, vol. 15, no. 5, p. 1021, 2026.
- [21] Y. Li et al., "Bridging NIP and MLM: A Unified Meta-Learning Framework for Sequential Recommendation," *ACM Transactions on Knowledge Discovery from Data*, vol. 20, no. 2, pp. 1–19, 2026.
- [22] Z. Lv et al., "Meta Learning Based Rumor Detection with Awareness of Social Bot," in *Knowledge Science, Engineering and Management*, C. Cao, H. Chen, L. Zhao, J. Arshad, T. Asyhari, and Y. Wang, Eds., Springer Nature Singapore, 2024.
- [23] H. Ahmed, I. Traore, and S. Saad, "Detecting Opinion Spams and Fake News Using Text Classification," 2018.
- [24] U. Y. Ibrahim, R. Prasad, B. H. Ibrahim, and O. F. Nonyelum, "Ensemble Meta-Learning Framework with BERT for Fake News Detection in Social Media.," *International Journal of Performability Engineering*, vol. 22, no. 1, 2026.
- [25] Y. Ye, H. Chen, Q. Cai, and H. Su, "Multimodal Meta-Learning for Early Rumor Detection Based on Few-Shot Learning," *IEEE Transactions on Computational Social Systems*, 2025.
- [26] J. Ma, W. Gao, and K.-F. Wong, "Detect rumors in microblog posts using propagation structure via kernel learning," 2017.
- [27] J. Ma, W. Gao, and K.-F. Wong, "Weibo Rumor Detection Dataset," 2016.
- [28] SyntheticProgrammer, "Rumor Detection Dataset (Twitter15 and Twitter16)," 2017.
- [29] S. Kwon, M. Cha, K. Jung, W. Chen, and Y. Wang, "Prominent features of rumor propagation in online social media," in *2013 IEEE 13th international conference on data mining*, 2013, pp. 1103–1108.
- [30] H. Han, Z. Ke, X. Nie, L. Dai, and W. Slamu, "Multimodal fusion with dual-attention based on textual double-embedding networks for rumor detection," *Applied Sciences*, vol. 13, no. 8, p. 4886, 2023.
- [31] J. MA et al., "Detecting rumors from microblogs with recurrent neural networks (2016)," in *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI 2016)*, pp. 3818–3824.
- [32] Y. Liu and Y.-F. Wu, "Early detection of fake news on social media through propagation path classification with recurrent and convolutional networks," in *Proceedings of the AAAI conference on artificial intelligence*, 2018.