



City Research Online

City St George's, University of London

Citation: Lourenço, V. N., Paes, A. & Weyde, T. (2026). Exploring Content and Social Connections of Fake News with Explainable Text and Graph Learning. Paper presented at the Brazilian Conference on Intelligent Systems, 29 Sep - 2 Oct 2025, Fortaleza-CE, Brazil. doi: 10.1007/978-3-032-15990-8_14

This is the accepted version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/37596/>

Link to published version: https://doi.org/10.1007/978-3-032-15990-8_14

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to [City Research Online policy](#).

Exploring Content and Social Connections of Fake News with Explainable Text and Graph Learning

Vítor N. Lourenço¹[0000-0003-0711-1339], Aline Paes¹[0000-0002-9089-7303], and
Tillman Weyde²[0000-0001-8028-9905]

Universidade Federal Fluminense, Niterói, Rio de Janeiro, Brazil
vitorlourenco@id.uff.br, alineapes@ic.uff.br
City St George’s, University of London, United Kingdom
T.E.Weyde@city.ac.uk

Abstract. The global spread of misinformation and concerns about content trustworthiness have driven the development of automated fact-checking systems. Since false information often exploits social media dynamics such as “likes” and user networks to amplify its reach, effective solutions must go beyond content analysis to incorporate these factors. Moreover, simply labelling content as false can be ineffective or even reinforce biases such as automation and confirmation bias. This paper proposes an explainable framework that combines content, social media, and graph-based features to enhance fact-checking. It integrates a misinformation classifier with explainability techniques to deliver complete and interpretable insights supporting classification decisions. Experiments demonstrate that multimodal information improves performance over single modalities, with evaluations conducted on datasets in English, Spanish, and Portuguese. Additionally, the framework’s explanations were assessed for interpretability, trustworthiness, and robustness with a novel protocol, showing that it effectively generates human-understandable justifications for its predictions. The code and experiments are available at <https://github.com/MeLLL-UFF/mu2X/>.

Keywords: Explainability · Interpretability · Fact-checking · Misinformation Detection · Multi-modality · Graph Neural Networks

1 Introduction

The transition from traditional print media to digital platforms has transformed news consumption, with social networks (*e.g.*, Reddit, Twitter, Facebook, WhatsApp, TikTok) becoming primary news-sharing platforms. Recent studies demonstrate this shift’s magnitude: Pew Research Center’s 2021 report¹ reveals 48% of U.S. residents rely on social networks for news, while Statista’s Latin America Dossier² indicates approximately two-thirds of surveyed individuals in Argentina, Brazil, Chile, and Mexico primarily consume news through social me-

¹ <https://www.pewresearch.org/journalism/2021/09/20/news-consumption-across-social-media-in-2021/>

² <https://www.statista.com/topics/8083/news-in-latin-america>

dia. This democratisation of information-sharing, while enabling universal content creation, has led to widespread misinformation. The impact is particularly evident in recent elections and the COVID-19 pandemic, where the lack of distinction between traditional and user-generated content has contributed to significant political and social challenges [1, 7, 40].

A prevalent method for addressing misinformation is the use of fact-checking organisations (such as the British Full Fact, the Brazilian Aos Fatos, and the Argentine Chequeado). These organisations typically rely on various sources and documents, usually encompassing images, videos, and text covering a diverse set of topics and international protocols³. They employ human assessment of the credibility of social media claims and news articles, which is time-consuming and often a bottleneck in providing verified, trustworthy information.

To address scalability challenges associated with manual fact-checking, several studies [8, 10, 12, 21, 24, 25, 27, 30–32, 35, 37] have proposed methods for automatic fact-checking of social media content. For instance, [27] have introduced a multimodal approach for detecting COVID-19-related misinformation on TikTok. Additionally, [20] developed a resource known as MuMiN, which establishes connections between claims and Twitter posts. These automated approaches aim to streamline the fact-checking process and enhance scalability by providing classification approaches and resources (*e.g.*, large-scale datasets) that effectively fill gaps in existing misinformation detection mechanisms.

Whilst aforementioned automatic fact-checking systems improves scalability, they often lack a crucial component: explainability [2, 17, 36], which is essential for user trust and legal compliance. Recent studies propose explainable frameworks for fake news detection [4, 16, 29] and multimodal misinformation detection [28, 39]. However, none of these approaches integrates multiple languages, content types, and social media features, nor do they provide a robust and comprehensive evaluation of their explainability.

This work introduces mu2X, a framework for addressing fact-checking in social media posts from **multimodal**, **multilingual**, and **explainable** perspectives. mu2X tackles explainable misinformation detection from complete and interpretable aspects, as defined by [6]. An *interpretable explanation* describes the system’s inner workings in a manner that humans understand; it must generate sufficiently straightforward explanations for an individual to grasp, employing terminology that resonates with the user. A *complete explanation* describes the system’s inner workings precisely; it must generate explanations that facilitate the anticipation of behaviours across a broader range of scenarios. Striking a balance between interpretability and completeness is challenging, since the most precise explanations are frequently less straightforward for individuals to grasp [6]. Conversely, the most easily understandable explanations often lack the ability to offer a more detailed understanding of the system’s behaviour and so to offer predictive insights that inform decision-making.

Our work leverages the explanations from two post-hoc model-agnostic explainability frameworks to (*i*) enhance awareness of the relevant features that led

³ <https://www.ifcncodeofprinciples.poynter.org/>

to a classification – contributing towards completeness; and *(ii)* enrich and combine the resulting explanations regarding the multiple modalities considered – contributing towards interpretability. In particular, our framework incorporates multiple modalities presented in social media posts by combining the posts’ local network information, textual content, and related metadata to detect misinformation whilst tackling the well-known lack of explainable classifications.

As shown in Fig. 1, the proposed framework comprises three main modules: *(i)* Post Encoding Module, where the posts’ contents are encoded into multimodal vector representations; *(ii)* Misinformation Detection Module, where the encoded representations are subjected to a graph-based classifier model to classify the posts as factual or misinformation; and *(iii)* Classification Explainability Module, where we employ graph-based and text-based explainable frameworks and combine their outcomes to explain the classification, whilst using the multimodal context to maintain the interpretability of generated explanations.

Our experiments employ the MuMiN dataset, to the best of our knowledge, the only multimodal, multilingual, and multi-topical Twitter-based knowledge base. Our goal is to assess mu2X’s misinformation detection and explanations capabilities regarding its interpretability, trustworthiness, and robustness. Our findings show that combining different data modalities improves overall misinformation detection, and the explainable setting enables interpretable, reliable, and robust explanations.

2 Related Work

In recent years, substantial endeavors have been put into automatically detecting misinformation [1, 4, 7, 8, 10, 12, 15, 16, 18, 20, 24, 25, 27–32, 35, 37, 39, 40]. In the present section, we review methods of misinformation classification concerning the covered modality and explainable capabilities. We highlight work that discusses a similar holistic view of misinformation classification in social networks, addressing both the multimodal aspect of information and the explainability of the output classification.

Misinformation Detection. Most contemporary methods for automated misinformation classification are based on single-modality content, neglecting the multimodal nature of news items. For instance, [30, 31] address misinformation detection by a method centered around hierarchical propagation networks, [10] uses graph neural networks supported by external knowledge, [35] rely on word embeddings over linguistic features, [12] use leverage reinforcement learning for fact-checking textual content, and [24] and [32] encode visual information to identify fake and manipulated information. Some recent work also identifies the benefit of encoding multimodal data to tackle misinformation detection [1]. [18] and [20] propose multimodal (image, text, and relational information) datasets for misinformation detection based on Reddit and Twitter posts, respectively. [8] introduce a new multimodal and multilingual dataset of misleading articles, training a model that leverages novelty detection and emotion recognition to

detect fabricated information. [37] propose using multi-reading habits fusion reasoning networks to capture inconsistencies between different types of modalities for news items.

Explainable Misinformation Detection. The decision-making process of identifying a post or news as misinformation or not goes beyond simply classifying it as such. Providing explanations is of foremost importance to prompt human confidence in the classification. To that end, [4] and [29] propose DEFEND, one of the first explainable methods for detecting fake news articles. [15] tailor a domain-specific approach for automated fact-checking of claims regarding public health, whilst providing local model-agnostic explainability using LIME [26]. In the same domain, [27] and [16] propose explainable methods to detect misinformation related to the COVID-19 pandemic on TikTok’s short videos and a knowledge-graph-based crowdsourced dataset. Closest to our work, [28] propose DGExplain, an explainable multimodal framework based on texts and images for COVID-19 misinformation detection. Finally, a multimodal and multi-topical dataset of articles and a model to automate misinformation fact-checking and provide explanations were recently proposed by [39]. Their explainable approach consists of generating a ruling statement using text evidence, the model prediction, and the input claim and conducting a policy learning process.

Our work advances existing research by incorporating relational information and relations’ metadata, leveraging multiple languages, and being domain-agnostic. We also contrast by considering the interpretability of the explanation, an aspect overlooked in prior research, and focus on social media content classification comprehensibility and explainability.

3 Method

This section introduces mu2X, our multimodal and multilingual explainable framework. Our primary goal is to offer an end-to-end solution for this task, highlighting key features that explain the classification, thus enabling better human comprehension. This section covers the problem definition. Following this, we offer an overview of the components within mu2X. Lastly, we present the implementation of the components forming this framework.

3.1 Problem Statement

We define a post-node $\mathcal{P}_i = (\mathcal{T}, \mathcal{G}_k, \mathcal{M}, lang)$, where \mathcal{T} is the textual content; \mathcal{G}_k defines the local k -hop network connections of the post, such as replies, quotes, and re-posts; \mathcal{M} is the post’s shallow metadata set, which includes, for example, the tweet’s owner, the number of likes, the location, etc; and $lang$ is the language in which the tweet was originally written. Our goal is for a model σ to identify and explain the most representative features within \mathcal{P}_i that lead a post to be classified as *misinformation* (i.e., $y_i = \sigma(\mathcal{P}_i) = 0$) or *fact* (i.e., $y_i = \sigma(\mathcal{P}_i) = 1$).

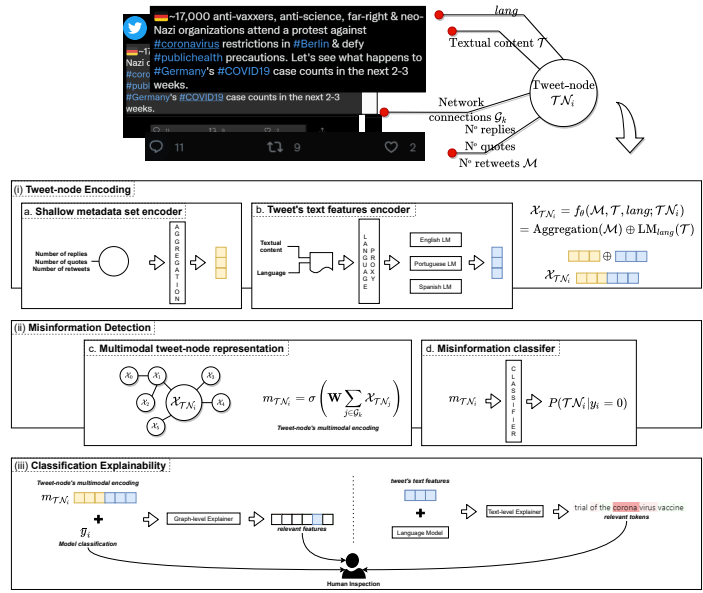


Fig. 1: The proposed multi-modal and multilingual explainable framework instantiated to Twitter posts. Yellow and blue elements illustrate graph-based and text-based feature vectors, respectively.

3.2 mu2X: Multimodal and Multilingual Explainable Framework

The proposed mu2X framework comprises three main modules: (i) post-node encoding; (ii) misinformation detection; and (iii) classification explainability. Fig. 1 illustrates the instantiated framework, and each module is described in the following.

Post Encoding Module. This module comprises the encoding of a given post-node \mathcal{P}_i in its vector representation. The module is divided into two components to represent two modalities: the post’s shallow metadata set \mathcal{M}_i ; and its textual content \mathcal{T}_i . The first part obtains the vector representation of the shallow metadata features with the number of replies, quotes, and re-posts (Fig. 1). a) The second part encompasses a language proxy defined by $lang$ followed by the language-specific text encoders (Fig. 1. b) The language proxy acts as a router, assigning the post’s textual content \mathcal{T} to a pretrained language-specific model (LM). The pre-trained language-specific model encoder generates a vector representation of the textual content. Finally, as described in Equation 1, both feature vectors are concatenated to form a unique multimodal post-node vector representation $\mathcal{X}_{\mathcal{P}_i}$.

$$\begin{aligned} \mathcal{X}_{\mathcal{P}_i} &= f_{\theta}(\mathcal{M}, \mathcal{T}, lang; \mathcal{P}_i) \\ &= \text{Aggregation}(\mathcal{M}) \oplus \text{LM}_{lang}(\mathcal{T}). \end{aligned} \tag{1}$$

Misinformation Detection Module. This module aims to identify misinformation through a graph-based classifier. It leverages the post’s encoded content $\mathcal{X}_{\mathcal{P}_i}$ with its graph connection features \mathcal{G}_k to generate a vector representation called *multimodal post-node representation*. To accomplish this, Equation 2 defines the learning process of the multimodal post-node representation. Given a set of post-node features, $\mathbf{X} = \{\mathcal{X}_{\mathcal{P}_1}, \mathcal{X}_{\mathcal{P}_2}, \dots, \mathcal{X}_{\mathcal{P}_N}\}$, $\mathcal{X}_{\mathcal{P}_j} \in \mathcal{G}_k$, $\mathcal{X}_{\mathcal{P}_j} \in \mathbb{R}^{|\mathcal{X}_{\mathcal{P}_i}|}$,

$$m_{\mathcal{P}_i} = \sigma \left(\mathbf{W} \sum_{j \in \mathcal{G}_k} \mathcal{X}_{\mathcal{P}_j} \right), \quad (2)$$

where σ is a nonlinear function, and \mathbf{W} is a linear transformation.

Next, the multimodal post-node representation $m_{\mathcal{P}_i}$ undergoes a classification process, where the probability distribution computed by a softmax function is taken from the representation:

$$P(\mathcal{P}_i | y_i = 0) = \text{softmax}(m_{\mathcal{P}_i}). \quad (3)$$

Classification Explainability Module. This module explains the predicted label by employing graph-based and text-based post-hoc explainable methods.

The graph-based method \mathcal{E}_g identifies the most important features within the post-node’s multimodal representation $\zeta_{\mathcal{P}_i}$, as displayed in Equation 4:

$$\zeta_{g\mathcal{P}_i} = \text{argmin} \mathcal{E}_g(\text{softmax}(m_{\mathcal{P}_i}), \mathbf{X}_{\mathcal{P}_i}), \quad (4)$$

where $\mathbf{X}_{\mathcal{P}_i}$ is the local information network of the target post-node \mathcal{P}_i .

The text-based method \mathcal{E}_t generates a scoring vector $\zeta_{t\mathcal{P}_i}$, where each word in the text is associated with a positive or negative importance score:

$$\zeta_{t\mathcal{P}_i} = \mathcal{E}_t(\text{LM}_{lang}, \mathcal{T}), \quad (5)$$

where $\zeta_{t\mathcal{P}_i} \in [-1, 1]^{|\mathcal{T}|}$, \mathcal{T} is the textual content of the post-node \mathcal{P}_i , and LM_{lang} is the text encoding method. Finally, both explanatory factors and the classification are provided to the user for inspection.

3.3 Framework Instantiation

This section details the mechanisms used to instantiate the previously introduced components illustrated in Fig. 1. Here, we focus on Twitter posts, calling them *tweet-nodes*.

Language Models. To encode our text features, we leverage both pre-trained domain and language specific models. For English tweets we use HuggingFace’s pre-trained BERTweet⁴ model [19] as encoder; we used BERTweet.BR⁵ for Portuguese tweets and RoBERTuito [23] for Spanish tweets.

⁴ https://huggingface.co/docs/transformers/model_doc/bertweet

⁵ <https://huggingface.co/mell1-uff/bertweetbr>

Graph-based Classifier. As misinformation classifier, we use the widely used Graph Attention Networks (GAT) [34], followed by a softmax classifier as described in Equation 3. As such, we compute Equation 2 employing a self-attention $\alpha_{kj}, \alpha : \mathbb{R}^{|\mathcal{X}_{P_i}|} \times \mathbb{R}^{|\mathcal{X}_{P_i}|} \rightarrow \mathbb{R}$, mechanism over the tweet-node features:

$$p_i = \sigma \left(\mathbf{w} \sum_{j \in \mathcal{G}_k} \alpha_{kj} \mathcal{X}_{P_j} \right).$$

Classification Explanation. To identify the key features that contribute to the classification, we used GraphLIME [11], a nonlinear model-agnostic explanation framework that extends LIME [26] for Graph Neural Networks. This approach uses Hilbert-Schmidt Independence Criterion (HSIC) Lasso [38] over k -hop neighbour nodes of a given target node. Furthermore, we employ Integrated Gradients [33] to assign an importance score to each word of a given tweet’s textual content. Integrated Gradients is an axiomatic model that assigns importance scores to individual input features. This approach consists of approximating the integral of the model’s output gradients concerning the input.

4 Experiments

We conduct experiments to show the capabilities of mu2X, our proposed multimodal and multilingual explainable framework. In our experiments, we measure and assess the framework’s capabilities of detecting misinformation, providing robust and trustworthy explanatory factors for these detections, and relating multiple explanatory factors to enhance interpretability.

We carry out our evaluations on a dataset including multilingual and multi-topical texts and network data that contextualise the texts in the social media environment. As far as we know, the MuMiN [20] dataset is the only dataset that fulfils all requirements of our evaluation. In addition, our focus is on characterising the representation of the textual content and relational features. The MuMiN dataset is a publicly available dataset of multimodal data from Twitter. MuMiN links tweets covering multiple topics and languages with fact-checked claims, including text, metadata, and visual content from the tweets. For our study, we relied on the *MuMiN-small* version of this dataset, which consists of 2,183 claims and 7,202,506 tweets. To streamline our analysis, we filtered the dataset to consider tweets written in the three most prevalent languages (English, Portuguese, and Spanish), four types of entities (*Claim*, *Tweet*, *Reply*, and *User*), and six relations (*Posted*, *Mentions*, *Retweeted*, *Quote_Of*, *Reply_To*, and *Discusses*).

We developed our framework using PyTorch [22]. The full code of the framework, the described instantiation and example notebooks are available for review. The dataset split used was the same as proposed in the MuMiN dataset. As described in Section 3.3, our misinformation classifier is a GAT to account for the

Table 1: Misinformation classification F1 for each language and modality.

GAT’s input Feature representation	Multilingual	English	Spanish	Portuguese
Graph-based	0.9488 ± 0.0032	0.9419 ± 0.0045	0.9477 ± 0.0029	0.9920 ± 0.0038
Text-based	0.9487 ± 0.0032	0.9409 ± 0.0044	0.9486 ± 0.0028	0.9923 ± 0.0042
Multimodal	0.9738 ± 0.0057	0.9738 ± 0.0053	0.9500 ± 0.0004	0.9951 ± 0.0053

relational information. The classifier was trained using a single Nvidia RTX4070 Mobile GPU, with a learning rate of 0.005, 16-dimensional hidden layer, and Adam optimiser [13] for a total of 400 epochs – the optimal setting we found in our experiments. The textual features were encoded using pre-trained language-specific models. The textual encoding is then submitted to a linear mapping of its original 768-dimensional space vector to a 812-dimensional space vector to match the dimensions of the shallow metadata set representation. The experiments in Sections 4.2-4.4 regard only the multilingual configuration, whilst comparing the various modalities.

4.1 Explainable Tweet Misinformation Detection

The empirical results of tweet classification are reported in Table 1, in which the model assigns each input information to the *misinformation* or *fact* label. We assess three different inputs: graph-based features, text-based features, and multimodal features, which are the concatenation of both graph-based and text-based features, and measure the obtained F1-score. The reported results are generated by 1000 bootstrap resampling with 95% confidence interval.

The results in each language demonstrate that the multimodal feature vector achieves better overall performance in misinformation detection than text encoding alone, confirming the results obtained by Nielsen *et al.* [20]. A secondary point is that text-based features outperform graph-based ones in low-resource languages and sparse graphs (*e.g.*, Spanish and Portuguese).

4.2 Interpretability of Explanations

We further examined the explanatory features through a combination of qualitative and quantitative analyses.

Our qualitative examination focused on the most relevant features and their ability to establish a connection between the tweet-node and its associated label. The results of our investigations are presented in Table 2, where we showcase the tweet-node text and a limited set of metadata, along with the classification produced by our GAT model, the most relevant features identified by GraphLime [11], and our interpretation of these explanatory factors.

In Table 2, the first entry is classified by the GAT model as misinformation. GraphLime highlighted the number of retweets and the number of replies as the most relevant features. Upon analysing these retweets and replies, we observed

Table 2: Report of two inferred cases with their classification, features, and explanatory factors identified by our framework.

Text	Shallow metadata set	Classification	GraphLime most representative features	Human interpretation
Great news! Corona virus vaccine ready. Able to cure patient within 3 hours after injection. Hats off to US Scientists. Right now Trump announced that Roche Medical Company will launch the vaccine next Sunday, and millions of doses are ready from it !!! VIA: @wajjh79273180 https://t.co/BZJCLtwuXq	Number of retweets 26 Number of replies 42 Number of quotes 7	Misinformation	Number of retweets Number of replies	Users who retweet tend to spread misinformation
17,000 anti-vaxxers, anti-science, far-right & neo-Nazi organizations attend a protest against #coronavirus restrictions in #Berlin & defy #publichealth precautions. Let's see what happens to #Germany's #COVID19 case counts in the next 2-3 weeks. https://t.co/TD5xIoT5sV	Number of retweets 11 Number of replies 9 Number of quotes 2	Fact	Text	Word importance illustrated in Fig. 2.



Fig. 2: Word importance for text explainability of the second tweet from the table above.

that users who engage in replies or retweets tend to have other tweets classified as misinformation, indicating a trend in the spread of misinformation. The second entry in Table 2 was classified as factual, with GraphLime’s most relevant feature being the text embedding. In addition, we explored textual factors using Captum [14] to assist in explaining the label (Fig. 2). We noted that phrases like “protest against #coronavirus” and “Let’s see what happens to #Germany” contributed to the accurate classification.

For our quantitative analysis, we fixed the multimodal setting to study how language and dataset size influence explanatory factors. Fig. 3 shows, for each language, the distribution of features selected as explanations by modality (metadata, graph-based, or text-based) and by the number of features selected (one, two, or three) and the overall (features together). For languages with fewer examples, like Spanish (Fig. 3.c) and Portuguese (Fig. 3.d), metadata and graph-based features are used more frequently, followed by text-based features. This is likely due to fewer nodes and less dense relational connections, making these features more discriminative. In contrast, the multilingual setting (Fig. 3.a) relies more on text-based features, likely due to its denser and more complex tweet-node network. English (Fig. 3.b) shows a balanced use across modalities, reflecting its intermediate node count and density. Regardless of language, metadata features are the most common explanatory factor when only one feature is selected, aligning with Table 2’s first entry and reinforcing that reply and retweet counts signal misinformation bubbles.

4.3 Trustworthiness of Explanations

In real-world applications, it is fundamental that the classifier’s predictions are credible, *i.e.*, the classifier should be consistently accurate in associating labels so that the end-user can relate to the prediction as trustworthy. Moreover, explanations play an essential role in assessing predictions. Therefore, we develop

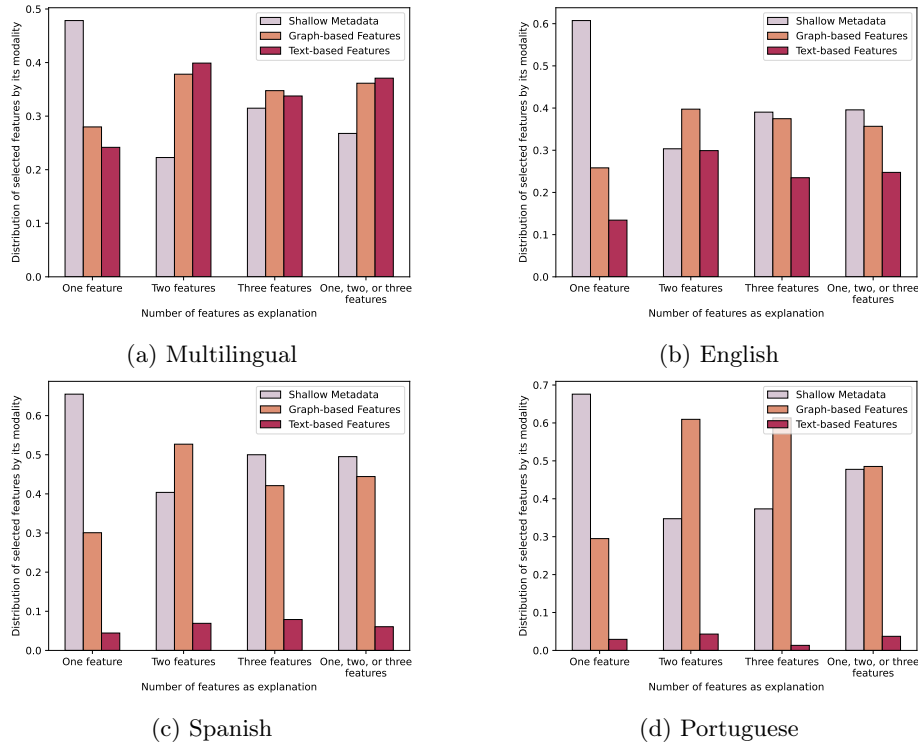


Fig. 3: Distribution of features selected as explanations by their modality. Results are separated by language configuration, and colour indicates each set of features.


a protocol to measure the extent to which the explanations have the capacity to assist users in determining the reliability of the prediction.

The protocol consists of, first, simulating the trust in individual predictions. For that, we randomly select 30% of the features and assume these as “untrustworthy”, *i.e.*, the user would likely recognise these features and would not want them as explanations. The second part of the protocol embraces the constructions of an *oracle* “trustworthiness”. We labelled the classifier’s prediction as *untrustworthy* if the prediction changes after removing the features selected as *untrustworthy* from an instance, and *trustworthy* otherwise. The oracle assumes that if removing *untrustworthy* features changes a prediction, those features are decisive for the classifier. In this simulated setup, the oracle serves as the true label for evaluating prediction credibility. The third step simulates user judgment by considering a prediction *untrustworthy* if it changes after removing all *untrustworthy* features identified in the explanation using a linear approximation model. Finally, we compare the oracle’s label with the simulated users’ decisions.

Given the trustworthiness evaluation protocol, we report and compare the F1-score on the trustworthy predictions for the classifier trained with graph-

Table 3: F1-Score of mu2X’s trustworthiness based on different modalities.

Modality	F1-Score given top- k most relevant features				
	Top-1	Top-2	Top-3	Top-5	Top-10
Graph-based	0.7576 ± 0.2199	0.8059 ± 0.0158	0.8035 ± 0.0188	0.8076 ± 0.0137	0.8031 ± 0.0187
Text-based	0.8188 ± 0.0164	0.8096 ± 0.0166	0.8111 ± 0.0154	0.8060 ± 0.0156	0.8029 ± 0.0118
Multimodal	0.8207 ± 0.0185	0.8061 ± 0.0153	0.8080 ± 0.0203	0.8062 ± 0.0167	0.8115 ± 0.0175


 Fig. 4: F1-Score of trustworthiness box-plot visualisation. Lines are the average F1-Score in each Top- K .

based features, text-based features, and multimodal features over 25 rounds, considering the top- K most relevant explainable features ($K \in \{1, 2, 3, 5, 10\}$). The results are compiled in Fig. 4 and Table 3. The multimodal classifier achieved the most stable results, having the best F1-score in top-1 and top-10, whilst being the second-best in the other cases. The classifier that was solely trained with graph-based features displayed the worst results, indicating that it usually mistrusts trustworthy predictions, whilst trusting untrustworthy predictions too often. Despite the used modality, the results show F1-Score above 80% (except the graph-based top-1), confirming the reliability of the mu2X explanations.

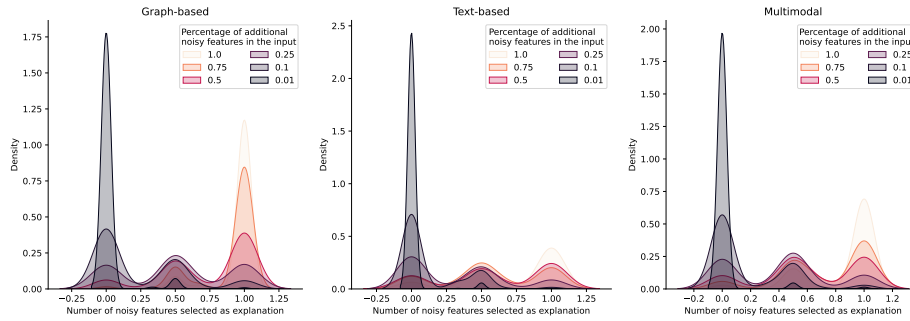
4.4 Robustness of Explanations

Besides credible predictions, robust explanations are key for the success of misinformation classifiers. We propose a protocol to evaluate the capacity of the explainable method to select relevant features, *i.e.*, non-noisy features. Firstly, we artificially added $N = |\mathcal{X}| * p$ amount of randomly generated features, *i.e.*, *noisy features*, where $N \in \mathcal{N}$, $|\mathcal{X}|$ is the number of features in a sample \mathcal{X} , and $p \in \{0.01, 0.1, 0.25, 0.5, 0.75, 1.0\}$ is the the overall proportion of noisy features to be introduced in each sample. Then, we train the classifier with the added noisy features. Finally, we observe the amount of noisy features selected as explanations.

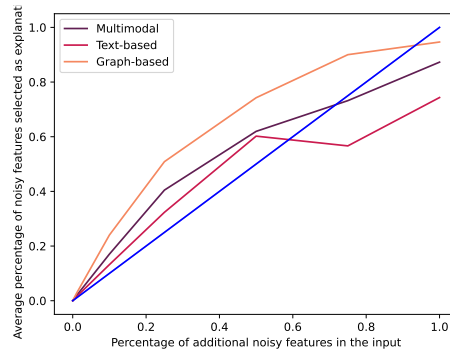
Given the robustness evaluation protocol, we report the average results of 25 rounds of the noisy features selected as explanation patterns for the three classifiers. In Fig. 5, we show the kernel density estimation of the frequency of

selected noisy features given the classifier and the proportion of noisy features added. The explanations based on the text-based features classifier were more robust, having a high density at zero noisy features selected in the lower percentages and even curves at zero or one noisy feature selected in higher noisy environments. The explanations grounded on the multimodal features classifier showed average results, where it keeps the higher density at zero for lower noisy environments and is even out on higher noisy environments.

These patterns can also be observed in Fig. 5d. The figure shows the distribution between the average percentage of noisy features selected as explanations and the percentage of noisy features added. Here, it is observed that the features selected as explanations for the text-based classifier increase linearly with the percentage of noisy features added. Additionally, it can be observed that the features selected as explanations using the multimodal classifier soften the non-linear curve presented by the explanations over the graph-based classifier.



(a) Graph-based features (b) Text-based features (c) Multimodal features



(d) Relation between the percentage of additional noisy features and the average percentage of noisy features selected as an explainable factor.

Fig. 5: Distribution of noisy features selected as explanation given the percentage of noisy features added.

5 Conclusion

This paper introduced mu2X, a multimodal and multilingual explainable framework for misinformation detection in social media. It leverages multiple data modalities and two explainability methods to detect misinformation while clarifying the features that support each classification. Our experiments show that multimodal features enhance detection performance. The proposed explainable methodology adds completeness to the framework by providing trustworthy, robust, and relevant features, along with interpretability, particularly when combined with multimodal data.

There are multiple directions for future work. First, we aim to enhance and further explore mu2X by incorporating additional modalities and modality-specific explainability methods. We hypothesize that leveraging more modalities will improve both classifier performance and the completeness and interpretability of explanations. Second, we plan to develop a protocol for quantitatively evaluating interpretability based on human assessment, since existing ones do not handle multimodality [3, 9]. Lastly, we envision integrated explanations that combine content snippets with explanatory factors or incorporate prior knowledge to improve clarity and user understanding.

Limitations. Our work has two main limitations: framework generalization and user interpretability. First, although it supports various multimodal aggregation methods, we only tested vector concatenation. Exploring other strategies could improve performance. Second, the feature-based explainability may be hard for some users to interpret. Still, our approach balances accuracy and interpretability by offering combined explanations across modalities, aligning with the trade-off discussed in [5].

Declaration on Use of Generative AI. The author(s) used Claude and Amazon Nova family models for grammar and spelling checks during the preparation of this work. After reviewing and editing the content as needed, the author(s) take full responsibility for the publication’s content.

Funding The second author thanks the grants from CNPq (National Council for Scientific and Technological Development), FAPERJ – *Fundação Carlos Chagas Filho de Amparo à Pesquisa do Estado do Rio de Janeiro*, processes SEI-260003/002930/2024, SEI-260003/000614/2023, and CAPES.

References

1. Alam, F., Cresci, S., *et al.*: A survey on multimodal disinformation detection. In: Proc. 29th Int. Conf. Comput. Linguist., COLING (2022)
2. Augenstein, I., Baldwin, T., Cha, M., *et al.*: Factuality challenges in the era of large language models and opportunities for fact-checking. Nat. Mach. Intell. (2024)
3. Colin, J., FEL, T., Cadene, R., Serre, T.: What i cannot predict, i do not understand: A human-centered evaluation framework for explainability methods. In: Adv. Neural Inf. Process. Syst., NeurIPS (2022)

4. Cui, L., Shu, K., Wang, S., Lee, D., Liu, H.: Defend: A system for explainable fake news detection. In: Proc. 28th ACM Int. Conf. Inf. Knowl. Manag., CIKM (2019)
5. Garouani, M., Mothe, J., Barhrhouj, A., Aligon, J.: Investigating the duality of interpretability and explainability in machine learning. In: IEEE Int. Conf. on Tools with Artif. Intell., ICTAI (2024)
6. Gilpin, L.H., Bau, D., Yuan, B.Z., Bajwa, A., Specter, M., Kagal, L.: Explaining explanations: An overview of interpretability of machine learning. Proc. IEEE 5th Int. Conf. Data Sci. Adv. Anal., DSAA pp. 80–89 (2018)
7. Guo, Z., Schlichtkrull, M., Vlachos, A.: A survey on automated fact-checking. Transactions of the Association for Computational Linguistics **10**, 178–206 (2021)
8. Gupta, V., Kumari, R., *et al.*: MMM: An emotion and novelty-aware approach for multilingual multimodal misinformation detection. In: Find. Assoc. Comput. Linguist., ACL (2022)
9. Hase, P., Bansal, M.: Evaluating explainable AI: Which algorithmic explanations help users predict model behavior? In: Proc. 58th Annu. Meet. Assoc. Comput. Linguist., ACL (2020)
10. Hu, L., Yang, T., *et al.*: Compare to the knowledge: Graph neural fake news detection with external knowledge. In: Proc. 59th Annu. Meet. Assoc. Comput. Linguist. 11th Int. Joint Conf. Nat. Lang. Process., ACL-IJCNLP (2021)
11. Huang, Q., Yamada, M., Tian, Y., Singh, D., Chang, Y.: Graphlime: Local interpretable model explanations for graph neural networks. IEEE Trans. Knowl. Data Eng., TKDE **35**(7), 6968–6972 (2023)
12. Kawintiranon, K., Singh, L.: Demis: Data-efficient misinformation detection using reinforcement learning. In: Mach. Learn. Knowl. Discov. Databases, ECML PKDD (2023)
13. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: Int. Conf. Learn. Represent., ICLR (2015)
14. Kokhlikyan, N., Miglani, V., Martin, M., *et al.*: Captum: A unified and generic model interpretability library for pytorch. CoRR **abs/2009.07896** (2020)
15. Kotonya, N., Toni, F.: Explainable automated fact-checking for public health claims. In: Proc. 2020 Conf. Empir. Methods Nat. Lang. Process., EMNLP (2020)
16. Kou, Z., Shang, L., Zhang, Y., Wang, D.: Hc-covid: A hierarchical crowdsourced knowledge graph approach to explainable covid-19 misinformation detection. Proc. ACM Hum.-Comput. Interact., HCI (2022)
17. Lourenço, V., Paes, A.: A modality-level explainable framework for misinformation checking in social networks. In: LatinX AI Workshop at 36th Conf. Neural Inf. Process. Syst., LXAI-NeurIPS (2022)
18. Nakamura, K., Levy, S., Wang, W.Y.: Fakeddit: A new multimodal benchmark dataset for fine-grained fake news detection. In: Proc. 12th Lang. Resour. Eval. Conf., LREC (2020)
19. Nguyen, D.Q., Vu, T., Tuan Nguyen, A.: BERTweet: A pre-trained language model for English tweets. In: Proc. 2020 Conf. Empir. Methods Nat. Lang. Process.: Syst. Demonstr., EMNLP. pp. 9–14 (2020)
20. Nielsen, D.S., McConville, R.: Mumin: A large-scale multilingual multimodal fact-checked misinformation social network dataset. Proc. 45th Int. ACM SIGIR Conf. Res. Dev. Inf. Retr., SIGIR pp. 3141–3153 (2022)
21. Panchendraran, R., Zubiaga, A.: Claim detection for automated fact-checking: A survey on monolingual, multilingual and cross-lingual research. Natural Language Processing Journal **7**, 100066 (2024)

22. Paszke, A., Gross, S., Massa, F., *et al.*: Pytorch: An imperative style, high-performance deep learning library. In: Adv. Neural Inf. Process. Syst., NeurIPS (2019)
23. Pérez, J.M., Furman, D.A., *et al.*: Robertuito: a pre-trained language model for social media text in spanish. In: Proc. Lang. Resour. Eval. Conf., LREC (2022)
24. Qi, P., Cao, J., Yang, T., Guo, J., Li, J.: Exploiting multi-domain visual information for fake news detection. In: IEEE Int. Conf. Data Min., ICDM (2019)
25. Qian, S., Hu, J., Fang, Q., Xu, C.: Knowledge-aware multi-modal adaptive graph convolutional networks for fake news detection. ACM Trans. Multimedia Comput. Commun. Appl. (2021)
26. Ribeiro, M.T., Singh, S., Guestrin, C.: "why should i trust you?": Explaining the predictions of any classifier. In: Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Discov. Data Min., SIGKDD (2016)
27. Shang, L., Kou, Z., Zhang, Y., Wang, D.: A multimodal misinformation detector for covid-19 short videos on tiktok. In: 2021 IEEE Int. Conf. Big Data, Big Data. pp. 899–908 (2021)
28. Shang, L., Kou, Z., Zhang, Y., Wang, D.: A duo-generative approach to explainable multimodal covid-19 misinformation detection. Proc. ACM Web Conf., WWW (2022)
29. Shu, K., Cui, L., Wang, S., Lee, D., Liu, H.: Defend: Explainable fake news detection. In: Proc. 25th ACM SIGKDD Int. Conf. Knowl. Discov. Data Min., SIGKDD (2019)
30. Shu, K., Mahudeswaran, D., Wang, S., Liu, H.: Hierarchical propagation networks for fake news detection: Investigation and exploitation. In: Int. Conf. Web Soc. Media, ICWSM (2020)
31. Shu, K., Wang, S., Liu, H.: Beyond news contents: The role of social context for fake news detection. In: Proc. 12th ACM Int. Conf. Web Search Data Min., WSDM (2019)
32. Singh, B., Sharma, D.K.: Siteforge: Detecting and localizing forged images on microblogging platforms using deep convolutional neural network. Comput. Ind. Eng. (2021)
33. Sundararajan, M., Taly, A., Yan, Q.: Axiomatic attribution for deep networks. In: Proc. 34th Int. Conf. Mach. Learn., ICML (2017)
34. Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., Bengio, Y.: Graph Attention Networks. Int. Conf. Learn. Represent., ICLR (2018)
35. Verma, P.K., Agrawal, P., Amorim, I., Prodan, R.: Welfake: Word embedding over linguistic features for fake news detection. IEEE Trans. Comput. Soc. Syst. (2021)
36. Warren, G., Shklovski, I., Augenstein, I.: Show me the work: Fact-checkers' requirements for explainable automated fact-checking. In: Conf. on Human Factors in Comput. Syst., CHI (2025)
37. Wu, L., Liu, P., Zhang, Y.: See how you read? multi-reading habits fusion reasoning for multi-modal fake news detection. Proc. AAAI Conf. Artif. Intell., AAAI **37**(11) (2023)
38. Yamada, M., Jitkrittum, W., Sigal, L., Xing, E.P., Sugiyama, M.: High-Dimensional Feature Selection by Feature-Wise Kernelized Lasso. Neural Comput. **26**(1), 185–207 (2014)
39. Yao, B.M., Shah, A., *et al.*: End-to-end multimodal fact-checking and explanation generation: A challenging dataset and models. In: Proc. 46th Int. ACM SIGIR Conf. Res. Dev. Inf. Retr., SIGIR (2023)
40. Zhou, X., Zafarani, R.: A survey of fake news: Fundamental theories, detection methods, and opportunities. ACM Comput. Surv. **53**(5) (2020)