# Exploring Knowledge Graphs for Automatic Fake News Detection in Portuguese

**Lucas dos Santos**
Federal University of Bahia (UFBA)
Camaçari, Bahia, Brazil
lukas.santtos53@gmail.com

**Manoel Rodrigues Euclides Santos**
Federal University of Bahia (UFBA)
Camaçari, Bahia, Brazil
santos.manoel@ufba.br

**Yuri Silva Souza**
Federal University of Bahia (UFBA)
Camaçari, Bahia, Brazil
souza.yuri@ufba.br

**João Pedro Holanda Sousa**
Universidade Norte do Paraná (UNOPAR)
Londrina, Paraná, Brazil
joaocmdt505@gmail.com

**Roney Lira de Sales Santos**
Federal University of Bahia
Camaçari, Bahia, Brazil
roneysantos@ufba.br

## Abstract

The proliferation of fake news in digital environments poses serious challenges to democratic processes, particularly in morphologically rich languages such as Portuguese. While most existing approaches focus on stylistic cues or propagation patterns in social networks, this paper proposes an automated fake news verification methodology grounded in Knowledge Graphs (KGs). Instead of treating news as raw text, we represent each article as a set of factual events encoded as semantic triples of subject, predicate, and object. A proprietary knowledge graph is built from Brazilian data sources, and a verification algorithm is introduced to estimate the veracity of news articles based on graph connectivity evidence. Experimental results confirm the feasibility of the proposed approach and highlight its inherent explainability as a key advantage over deep learning black-box models. Error analysis further indicates that the main limitation stems from the syntactic complexity of Open Information Extraction in Portuguese, suggesting that improvements at this extraction stage are essential to increase system robustness.

## 1 Introduction

Fake news dissemination represents a significant challenge for today's information society, as social networks amplify its reach and enable rapid circulation of digital content (Vosoughi et al., 2018; Lazer et al., 2018). In academic contexts, fake news is typically characterized as deliberately deceptive content produced to induce distorted interpretations of reality (Rubin et al., 2015). This phenomenon includes both fully fabricated stories and the distortion or decontextualization of real facts. It also goes beyond the term's popular political use to discredit journalistic outlets.

Automated detection of this type of content raises substantial technical challenges, largely because of how fake news is crafted. Fake news is rarely entirely invented; it is often built from fragments of real facts presented out of context, which blurs the boundary between truth and falsehood. This characteristic limits the effectiveness of automated systems based solely on lexical or stylometric patterns. Moreover, the architecture of social networks increases the scale of the problem, demanding verification methods that are both accurate and scalable.

For Portuguese, these challenges are compounded by the scarcity of adequate linguistic and computational resources. While the English-language literature offers large corpora and mature tools for Open Information Extraction (OIE), Portuguese still lacks robust annotated datasets for complex semantic tasks (Monteiro et al., 2018). As a result, many existing approaches focus on style-based classification or propagation signals, with factual verification playing a secondary role.

To address this gap, this work proposes an approach to automatic fake news detection in Portuguese based on Knowledge Graphs. The methodology represents each news article as a set of factual events expressed as semantic triples composed of subject, predicate, and object, and estimates veracity from the degree of compatibility between these events and a reference graph built



*Proceedings of the 17th International Conference on Computational Processing of Portuguese (PROPOR 2026)*
*Vol. 1*, pages 446–455

from trusted sources. Accordingly, the paper focuses not only on classification outcomes but also on the challenges of Open Information Extraction in Portuguese, which is critical for building event-based graphs.

This paper is organized as follows: Section 2 reviews related work, focusing on fake news detection in Portuguese, Knowledge Graphs for fact checking, and Open Information Extraction. Section 3 presents the proposed methodology, detailing graph construction and the veracity scoring procedure. Experimental results and error analysis are discussed in Section 4. Finally, Section 5 concludes the paper and outlines directions for future work.

## 2 Related Work

Automatic fake news detection is a multidisciplinary research area that has evolved substantially over the past decade, drawing contributions from Computational Linguistics, Information Science, and Network Science. In this section, we review related work from three complementary perspectives: (i) fake news detection approaches for Portuguese; (ii) the use of Knowledge Graphs for fact checking; and (iii) the challenges of Open Information Extraction in Portuguese, which is fundamental for building event-based graphs.

### 2.1 Fake News Detection in Portuguese

In Portuguese, the development of benchmark corpora has been crucial to the field's progress. The Fake.Br Corpus (Monteiro et al., 2018; Silva et al., 2020) became a reference dataset by aligning true and false news on the same topics, enabling controlled classification experiments and comparative analyses of linguistic characteristics. Later, additional datasets were introduced, such as FakeRecogna (Garcia et al., 2022) and, more recently, FakeTrueBR (Chavarro et al., 2023), which aimed to mitigate temporal and topical misalignment by pairing news items using Sentence-BERT-based models (Reimers and Gurevych, 2019).

Most Portuguese-language approaches focus on stylometric and linguistic features. Prior studies show that classical classifiers such as Support Vector Machines and Random Forests, when combined with features like bag-of-words, part-of-speech patterns, and readability metrics, can achieve consistent performance in fake news detection (Monteiro et al., 2018; Silva et al., 2020). These results highlight the strength of style-based methods under controlled conditions.

More recently, researchers have explored deep learning models and Large Language Models (LLMs) for fake news detection in Portuguese. Recent work indicates that Transformer-based architectures, such as BERTimbau and instruction-tuned variants, reach high performance in controlled settings, particularly when paired with supervised fine-tuning strategies (Gôlo et al., 2024; Garcia et al., 2024b). However, these studies also report limitations related to interpretability and reliance on textual correlations, reinforcing interest in complementary approaches based on explicit factual verification.

Nevertheless, a fundamental limitation of these approaches is their dependence on linguistic style. Sophisticated fake news that closely mimics journalistic writing can reduce the effectiveness of style-driven methods. This challenge has grown with the popularization of large language models capable of generating fluent, well-structured text that is nevertheless factually incorrect, strengthening the case for strategies that explicitly verify factual content.

### 2.2 Knowledge Graphs for Fact Checking

In contrast to purely text-based approaches, Knowledge Graph-based methods aim to model veracity through semantic consistency between entities and relations. A pioneering work in this direction is presented by Ciampaglia et al. (2015). In that study, the authors propose computing truth values from the semantic proximity between entities in graphs derived from DBpedia (Auer et al., 2007).

Later studies extended this idea by incorporating explicit verification of triples extracted from news articles. For example, Pan et al. (2018) use knowledge graphs to validate claims represented as subject, relation, and object triples. Despite these advances, a recurring limitation is the reliance on encyclopedic and static graphs such as Wikidata or DBpedia, which do not always capture domain-specific news events or timely information.

In the Brazilian context, Santos and Pardo (2020) explore the construction of Portuguese-specific graphs and propose methods based on shortest-path searches between entities in graphs derived from trusted sources. While these efforts are relevant advances, they still emphasize entities more than the explicit modeling of factual events.

### 2.3 Challenges of Open Information Extraction in Portuguese

Open Information Extraction (OIE) (Banko et al., 2007) refers to the task of automatically identifying semantic relations from unstructured text without assuming a predefined relation schema. Unlike supervised extraction, OIE systems aim to extract factual statements in an open-ended way, typically represented as subject, predicate, and object triples, enabling broader coverage across domains (Etzioni et al., 2011; Del Corro and Gemulla, 2013; Angeli et al., 2015). This flexibility makes OIE especially attractive for fact-checking applications, but it also introduces challenges related to semantic precision and linguistic normalization.

The viability of event-graph-based approaches depends directly on the ability to extract structured triples from unstructured text via OIE. While English has well-established tools such as OpenIE (Etzioni et al., 2011) and ClausIE (Del Corro and Gemulla, 2013), Portuguese poses specific syntactic challenges, including more flexible constituent order, frequent subject omission, and high verbal inflection complexity.

Generic tools based on machine translation or multilingual models often fail to capture semantic relations accurately in more complex Portuguese sentences. To address these limitations, work such as DptOIE (Oliveira and Claro, 2019; Glauber et al., 2019) proposes dependency-based approaches, improving the extraction of compound verbal relations and handling subordinate clauses. However, applying these tools to build graphs for news verification, where relational precision is critical, remains relatively underexplored.

### 2.4 Research Gap

The literature reveals a clear gap. On the one hand, there are robust Portuguese fake news classifiers that rely primarily on textual features or propagation patterns (Chavarro et al., 2023). On the other hand, knowledge-graph-based verification approaches are often constrained to encyclopedic graphs or predominantly English-language scenarios.

To the best of our knowledge, this is the first work to propose building and using an event-based Knowledge Graph rather than an entity-centered graph, specifically for the Portuguese news domain. The approach integrates Portuguese-adapted Open Information Extraction with a graph-topology-based veracity scoring algorithm. Unlike methods that classify the entire text as an indivisible unit, our method seeks to validate individual factual claims that compose the news narrative.

## 3 Methodology

The proposed methodology follows a two-stage pipeline, as illustrated in Figure 1. The first stage builds a base knowledge graph, referred to as KG-base, from true news articles. The second stage verifies an input news article. In both stages, texts undergo preprocessing and event extraction into subject, predicate, and object triples. In the first stage, these triples are used to build the reference graph. In the second, triples extracted from the input article are queried against the graph to compute a truth value that supports the final decision.

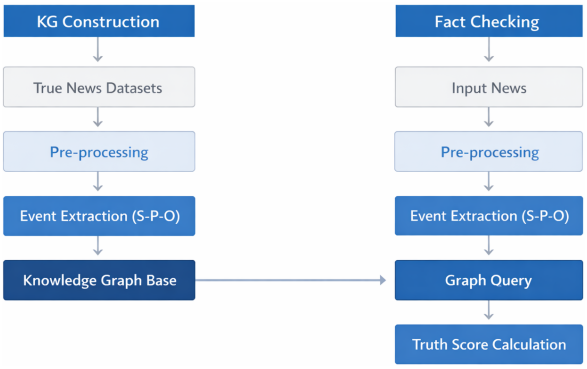


Figure 1: Pipeline of the proposed method for KG-base construction and automatic news verification.

### 3.1 Dataset Preprocessing

For the experiments, we consolidated a Portuguese-language collection composed of three public benchmark datasets widely used in the literature: Fake.Br (Monteiro et al., 2018), FakeRecogna (Garcia et al., 2022), and FakeTrueBR (Chavarro et al., 2023). Table 1 summarizes their main characteristics.

To construct the KG-base, we used only the articles labeled as true, totaling 11,342 news items. These data form the reference set of facts used during verification. Articles labeled as false were used only for evaluating the method and were not incorporated into the graph construction.

During preprocessing, texts underwent tokenization, basic normalization, and stopword removal. We intentionally did not apply named entity recognition or coreference resolution in order to isolate the impact of dependency-based event extraction,

| Dataset | Total News | True News | Usage |
|---|---|---|---|
| Fake.Br | 7200 | 3600 | KG-base + Evaluation |
| FakeTrueBR | 3582 | 1791 | KG-base + Evaluation |
| FakeRecogna | 11902 | 5951 | KG-base + Evaluation |

Table 1: Summary of the datasets used in this study.

which is the central focus of this work.

### 3.2 Event Extraction (Subject, Relation, Object)

Events are represented as semantic triples defined as $T = (s, p, o)$, where $s$ and $o$ correspond to the subject and object of the event, respectively, and $p$ denotes the predicate that links them. This representation reduces each news article to a discrete set of verifiable factual statements.

We extracted triples using a deterministic, rule-based morphosyntactic algorithm implemented with the spaCy `pt_core_news_lg` model (Honnibal et al., 2020). For each sentence, the procedure consists of: (i) segmentation and morphosyntactic tagging; (ii) identifying the subject via subject dependencies (e.g., `nsubj`); (iii) identifying the predicate as the main verb, with lemmatization; and (iv) identifying the direct object via object dependencies (e.g., `obj`). A triple is kept only when all three elements are successfully identified.

This strategy prioritizes precision over coverage and is especially effective for simple factual sentences, which are common in headlines and leads. In contrast, sentences containing coordination, ellipsis, negation, or subordinate structures may yield incomplete or semantically imprecise extractions.

### 3.3 Knowledge Graph Construction

We model the graph using NetworkX (Hagberg et al., 2008) and define it as $G = (V, E)$, where $V$ is the set of extracted entities and $E$ is the set of semantic relations between them. Nodes represent entities or noun phrases, and edges represent predicates.

Edges are treated as undirected, enabling bidirectional traversal during verification, following the formulation of Ciampaglia et al. (2015). Unlike encyclopedic graphs, the resulting knowledge graph is event oriented and reflects interactions described in true news articles.

### 3.4 Truth Value Scoring

Given an input news article $N$, we extract a set of triples $E_N = \{e_1, e_2, \ldots, e_m\}$ using the procedure described in Section 3.2. For each triple $e = (s, p, o)$, we compute a truth value $\tau(e) \in [0, 1]$ based on the semantic proximity measure proposed by Ciampaglia et al. (2015). This measure relates a claim's veracity to the connectivity between $s$ and $o$ in the graph, penalizing paths that traverse highly generic nodes.

The overall score of an article is defined as the arithmetic mean of the truth values assigned to its extracted triples:

$$Score(N) = \frac{1}{|E_N|} \sum_{e \in E_N} \tau(e).$$

Intuitively, articles whose triples are more compatible with the KG-base tend to receive higher scores, whereas articles containing events not supported by the graph tend to receive lower scores.

We use a two-threshold decision scheme that explicitly defines an uncertainty region. An article $N$ is classified as *true* if $Score(N) \geq 0.60$ and as *fake* if $Score(N) < 0.40$. Intermediate values in the interval $[0.40, 0.60)$ are labeled as *uncertain*, reflecting cases where the method finds partial evidence in the graph but not enough to support a reliable binary decision. Threshold selection and the empirical score distribution are discussed in Section 4.

## 4 Results and Discussion

This section presents the experimental evaluation of the proposed knowledge-graph-based method for fake news detection. We report quantitative results across different scenarios and discuss, through error analysis, the main factors that limit performance, with emphasis on Portuguese event extraction and KG-base coverage.

### 4.1 Experimental Setup

The evaluation was conducted using three Portuguese datasets widely adopted in the fake news detection literature, namely Fake.Br, FakeTrueBR, and FakeRecogna, as well as a consolidated scenario that merges these corpora into a single collection. In each scenario, the KG-base was built

from articles labeled as true, while the verification procedure was applied to the articles in the corresponding test set.

Performance was measured using standard binary metrics: accuracy, precision, recall, $F_1$, and AUC-ROC. Among these, $F_1$ is particularly relevant because it summarizes the trade-off between precision and recall, an important aspect in fake news detection, where the costs of false positives and false negatives can be asymmetric depending on the intended use case.

We adopt a decision policy based on the article's mean truth value. Each article receives a final value in $[0, 1]$, defined as the average truth value of its extracted triples. Because the method explicitly defines an *uncertain* class in the interval $[0.40, 0.60)$, binary metrics are reported under a conservative policy in which intermediate cases are mapped to the fake class. This choice reduces false positives, meaning it avoids labeling as true an article with weak factual support, at the cost of increasing false negatives when true articles exhibit only partial evidence in the KG-base.

### 4.2 Quantitative Results

Table 2 reports the performance of the proposed method across the evaluated scenarios. The best results are obtained on FakeTrueBR and FakeRecogna, while performance drops in the consolidated setting. This behavior is consistent with increased lexical, topical, and temporal heterogeneity from merging distinct corpora, which expands the variety of entities, relations, and syntactic structures, amplifying the impact of extraction noise and KG-base coverage gaps.

Overall, these results suggest that knowledge-graph-based verification can achieve competitive performance in Portuguese when focusing on high-confidence decisions. Unlike purely textual classifiers, the proposed method grounds its predictions on explicitly extracted events and graph connectivity evidence, enabling inspection of which claims are supported by the KG-base and which remain unsupported.

It is important to note that the KG-base used in this work is built exclusively from articles labeled as true. This methodological choice reflects the core goal of the approach, which is to verify whether the events described in a news article are supported by a set of previously observed facts considered reliable. In this sense, the KG-base serves as a factual reference against which the semantic compatibility of extracted triples is evaluated.

This choice has direct implications for the observed results. On the one hand, it strengthens the interpretability and verifiability of the method, since high scores correspond to explicit factual support in the graph. On the other hand, it makes the system naturally sensitive to coverage limitations, especially for recent, rare, or sparsely documented events. Therefore, performance drops in more heterogeneous scenarios should not be interpreted as failures of the verification criterion itself, but rather as consequences of factual incompleteness in the graph and the inherent difficulty of event extraction in Portuguese.

#### 4.2.1 Abstention Based Decisions and Metric Evaluation

Beyond the aggregated metrics reported above, we analyzed the classifier behavior by explicitly considering the uncertainty class defined by the intermediate truth value interval. In this work, this class is interpreted as an abstention mechanism rather than as a classification error, signaling cases where the KG-base provides partial factual support but not enough to justify a reliable binary decision.

Under this setting, the system operates internally with three labels: fake (0), true (1), and uncertain (2). For quantitative evaluation and direct comparison with binary approaches in the literature, we adopt a conservative policy in which uncertain predictions are mapped to the fake class. This choice prioritizes reducing false positives and aligns with automated fact-checking scenarios where accepting fake news as true is considered the most severe error.

Table 3 reports performance metrics computed from the full confusion matrices of each scenario. Metrics were calculated from the original three-label matrices after mapping uncertain predictions (2) to the fake class (0). Even under this restrictive policy, the method maintains high $F_1$ and accuracy, particularly on FakeTrueBR and FakeRecogna, where $F_1$ exceeds $0.96$.

The consistently high recall across scenarios indicates that the method is particularly effective at identifying true news when the graph provides strong factual evidence. However, mapping uncertain cases to the fake class directly affects the precision and recall balance. It reduces the risk of labeling as true an article with weak factual support, but it can increase false negatives when true articles have only partial support in the KG-base.

| Scenario | Acc | Prec | Rec | $F_1$ | AUC |
|---|---|---|---|---|---|
| Corpora (Overall) | 0.8365 | 0.9420 | 0.8365 | 0.8725 | 0.7375 |
| Fake.Br | 0.8193 | 0.9458 | 0.8193 | 0.8612 | 0.7000 |
| FakeTrueBR | 0.9176 | 0.9663 | 0.9176 | 0.9378 | 0.8713 |
| FakeRecogna | 0.9234 | 0.9690 | 0.9234 | 0.9432 | 0.8789 |

Table 2: Results of the proposed method across evaluation scenarios.

| Scenario | Acc | Prec | Rec | $F_1$ |
|---|---|---|---|---|
| Corpora (Overall) | 0.9359 | 0.8887 | 0.9966 | 0.9396 |
| Fake.Br | 0.9390 | 0.8917 | 0.9994 | 0.9425 |
| FakeTrueBR | 0.9640 | 0.9333 | 0.9994 | 0.9652 |
| FakeRecogna | 0.9681 | 0.9449 | 0.9941 | 0.9689 |

Table 3: Performance with abstention mapped to the fake class, with metrics derived from confusion matrices.

This analysis shows that the model behaves cautiously and selectively, preferring abstention over potentially incorrect decisions. Explicitly modeling uncertainty also enables future extensions such as selective classification, human in the loop workflows, or hybrid decision strategies that treat the uncertain class as an informative output rather than forcing binary predictions.

#### 4.2.2 Comparison with the State of the Art

Within stylometric and linguistic approaches, studies on Fake.Br report $F_1$ values around 0.88 (Monteiro et al., 2018), while more refined models achieve values close to 0.97 in controlled settings (Silva et al., 2020). These results show the strength of style signals as discriminators, but also indicate reliance on linguistic regularities that can be mimicked, especially with the spread of generative language models.

On the factual verification side, Santos and Pardo (2020) validate claims by combining knowledge graphs with external evidence retrieved via search engines. In contrast, our method restricts inference to the topological consistency of a graph built exclusively from true news, prioritizing reproducibility and experimental control at the cost of greater sensitivity to factual coverage and temporal updates of the KG-base.

To the best of our knowledge, no previous work has employed event-based knowledge graphs for fake news detection in Brazilian Portuguese in a way that allows direct performance comparison. Nevertheless, it is possible to situate our results in relation to recent approaches based on Large Language Models (LLMs). In a recent study, Gôlo et al. (2025) evaluated the use of models such as GPT-4 and Llama for detecting misinformation in Brazilian political content, reporting $F_1$ scores ranging from 0.85 to 0.92. Our results, which achieved an $F_1 = 0.87$ using a knowledge-graph approach with triple extraction, are therefore competitive and, in some scenarios, comparable to these black-box models. An important advantage of the proposed approach lies in explainability: while LLMs operate through probabilistic and opaque decision processes, the graph-based method enables explicit tracing of the factual evidence (subject–predicate–object triples) that supports the classification of each news article, reducing the risk of hallucinations inherent to generative models.

Finally, transformer-based deep learning approaches tend to report the highest absolute performance. Garcia et al. (2024a) report $F_1$ scores above 0.96, and the BERTimbau model (Souza et al., 2020) reaches approximately $F_1 = 0.984$. However, quantitative gains should be interpreted in light of the type of evidence that supports the decision. Deep models act primarily as correlational classifiers and rarely provide explicit event-level justifications. In contrast, the knowledge-graph-based approach emphasizes explicit factual verification, with triple-level traceability and connectivity evidence, which is a key property in decision support tools and scenarios that require auditability. Thus, observed performance differences reflect distinct methodological goals and underscore the potential of graph-based methods as complementary components in fact-checking systems.

### 4.3 Error Analysis

The qualitative analysis indicates that the main errors arise from two structural limitations of the

method: the quality of extracted triples and the factual coverage of the KG-base. While deep language models can absorb local noise through dense representations, our approach depends directly on the completeness and precision of triples, shifting the primary bottleneck to Open Information Extraction in Portuguese.

Recurring syntactic phenomena in Portuguese, such as subject omission, flexible constituent order, and frequent subordinate clauses, make it difficult to identify semantic relations precisely. As a result, incomplete or semantically imprecise triples become more common. When events present in a true article are not properly represented, the system fails to find sufficient support in the graph, leading to a lower truth value. In many such cases, the article is correctly flagged as uncertain rather than being labeled as true without solid factual support.

The empirical distribution of scores reinforces this diagnosis. We observe a substantial concentration of articles with intermediate values, approximately in the range $[0.40, 0.59]$, which characterizes cases of partial factual evidence in the KG-base. This uncertainty region reflects both graph coverage gaps and event extraction limitations. The conservative policy adopted here, which treats such cases as abstentions during inference and as the negative class in binary evaluation, reduces the risk of labeling as true an article with fragile factual support, at the cost of penalizing traditional metrics such as precision and accuracy.

False negatives occur predominantly in contextual fake news, where true facts are combined with fabricated inferences or conclusions. In these scenarios, validating isolated triples can raise the article's mean score and mask the central misleading claim. This behavior highlights a structural limitation of approaches based solely on local factual validation and suggests the need for aggregation strategies that are more sensitive to global narrative consistency or for explicit contradiction detection mechanisms.

Taken together, this analysis indicates that the observed errors do not stem from isolated flaws in the truth value computation, but rather reflect current limitations of event extraction and explicit knowledge modeling in Portuguese. Strengthening this layer through more robust extractors, hybrid strategies, or integration with neural models is the primary path toward closing the gap to the state of the art while preserving the method's advantages in interpretability and explicit factual verification.

## 5 Conclusion

This paper presented a methodology for automatic fake news detection in Portuguese based on structured factual verification through Knowledge Graphs. In contrast to approaches centered on stylistic patterns or propagation signals, the proposed method represents each article as a set of factual events expressed as semantic subject, predicate, and object triples. From this representation, we build a KG-base exclusively from articles labeled as true, which serves as a factual reference to verify whether events described in a given article are supported by previously observed facts. This design choice makes the method's core principle explicit: veracity is estimated by the compatibility between extracted events and the structured knowledge encoded in the graph.

Experimental results indicate that the approach is viable for fake news detection in Portuguese and can achieve competitive performance across different scenarios, reaching $F_1 = 0.8725$ in the consolidated setting. Beyond quantitative performance, the main contribution lies in explainability. Decisions are grounded in explicit triples and graph connectivity evidence, making it possible to identify which claims are factually supported and which remain unsupported by the KG-base. This distinguishes the approach from purely correlational models and supports its use as a decision support tool for fact checking. In addition, explicitly treating uncertain cases as abstentions reinforces the method's conservative stance by avoiding affirmative decisions when factual support in the KG-base is insufficient.

Error analysis shows that system performance is primarily limited by the quality of Open Information Extraction in Portuguese and by the factual and temporal coverage of the graph. Furthermore, because the KG-base is built only from true news available in the datasets used, recent or sparsely documented events tend to have lower connectivity, negatively affecting veracity scores.

As a practical outcome, we developed a functional web interface that implements the proposed method and allows users to submit texts for automatic fake news verification (Souza et al., 2025; Santos et al., 2025). The application was designed with usability and accessibility in mind and is ready for deployment, pending hosting on appropriate computational infrastructure to enable public access.

Future work falls into three main directions. First, improving event extraction with more robust neural models, including Transformer-based methods and instruction tuned large language models designed specifically for triple generation, with proper handling of negation, coreference resolution, and semantic normalization. Second, building a dynamic knowledge graph with continuous updates and incremental ingestion of new facts, reducing temporal coverage limitations. Third, developing hybrid decision strategies that combine the factual veracity score with complementary textual signals, preserving interpretability while improving robustness against increasingly sophisticated fake news.

As part of our commitment to reproducibility and transparency, the code used in the experiments is publicly available online[1]. This study is also part of a broader research initiative called FAKTUS-BR[2], which aims to investigate computational approaches for detecting fake news in Portuguese using knowledge graphs, natural language processing, and large-scale language models.

In summary, this work shows that explicit modeling of events and relations is a promising alternative for fake news detection in Portuguese, particularly in applications that require transparency, auditability, and factual justifications. Within the broader context of the FAKTUS-BR project, the proposed methodology represents an initial step toward building interpretable and verifiable computational tools to support automated misinformation detection.

## Acknowledgments

The authors gratefully acknowledge the financial support provided by the Universidade Federal da Bahia (UFBA) and the Brazilian National Council for Scientific and Technological Development (CNPq).

[1]Available at https://github.com/FAKTUS-BR/propor26

[2]Available at https://github.com/FAKTUS-BR

## References


Gabor Angeli, Melvin Jose Johnson Premkumar, and Christopher D. Manning. 2015. Leveraging linguistic structure for open domain information extraction. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics (ACL)*.

Sören Auer, Christian Bizer, Georgi Kobilarov, Jens Lehmann, Richard Cyganiak, and Zachary Ives. 2007. Dbpedia: A nucleus for a web of open data. In *Proceedings of the 6th International Semantic Web Conference (ISWC)*.

Michele Banko, Michael J. Cafarella, Stephen Soderland, Matt Broadhead, and Oren Etzioni. 2007. Open information extraction from the web. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 2670–2676.

Juan Pablo Chavarro, Jonata Tyska Carvalho, Tarlis Tortelli Portela, and Jonathan Cardoso Silva. 2023. Faketruebr: um corpus brasileiro de notícias falsas. In *Anais da Escola Regional de Banco de Dados (ERBD)*, Palmas, PR.

Giovanni Luca Ciampaglia, Prashant Shiralkar, Luis M. Rocha, Johan Bollen, Filippo Menczer, and Alessandro Flammini. 2015. Computational fact checking from knowledge networks. *PLOS ONE*, 10(6):e0128193.

Luciano Del Corro and Rainer Gemulla. 2013. Clausie: Clause-based open information extraction. In *Proceedings of the 22nd International Conference on World Wide Web (WWW)*, pages 355–366.

Oren Etzioni, Anthony Fader, Janara Christensen, Stephen Soderland, and Mausam. 2011. Open information extraction: The second generation. *AI Magazine*, 32(3):3–13.

Gabriel Lino Garcia, Pedro Henrique Paiola, Danilo Samuel Jodas, Luis Afonso Sugi, and João Paulo Papa. 2024a. Text summarization and temporal learning models applied to portuguese fake news detection in a novel brazilian corpus dataset. In *Proceedings of the 16th International Conference on Computational Processing of Portuguese (PROPOR)*, Santiago de Compostela, Spain. Association for Computational Linguistics.

Gustavo L. Garcia, Lucas C. Afonso, and João P. Papa. 2022. Fakerecogna: a new brazilian corpus for fake news detection. In *Computational Processing of the Portuguese Language*, Lecture Notes in Artificial Intelligence, pages 57–67. Springer.

Gustavo L. Garcia, Lucas C. Afonso, and João P. Papa. 2024b. Fakerecogna 2.0: Advances in brazilian fake news detection. In *PROPOR*.

Rafael Glauber, Diego B. Claro, and Leandro S. Oliveira. 2019. Dependency parser on open information extraction for portuguese texts: Dptoie and dependentie on iberlef. In *Proceedings of IberLEF*, pages 442–448.

Marcos P. S. Gôlo, Adriel L. V. Mori, William G. Oliveira, Jacson R. Barbosa, Valdemar V. Graciano-Neto, Eliomar A. de Lima, and Ricardo M. Marcacini. 2024. On the use of large language models to detect brazilian politics fake news. In *Proceedings of the Encontro Nacional de Inteligência Artificial e Computacional (ENIAC)*, pages 1–12.

Marcos P. S. Gôlo, Adriel L. V. Mori, William G. Oliveira, Jacson R. Barbosa, Valdemar V. Graciano-Neto, Eliomar A. Lima, and Ricardo M. Marcacini. 2025. Evaluating large language models for brazilian political fake news detection. In *Proceedings of the Encontro Nacional de Inteligência Artificial e Computacional (ENIAC)*.

Aric Hagberg, Pieter J. Swart, and Daniel A. Schult. 2008. Exploring network structure, dynamics, and function using networkx. Technical report, Los Alamos National Laboratory.

Matthew Honnibal, Ines Montani, Sofie Van Landeghem, and Adriane Boyd. 2020. spacy: Industrial-strength natural language processing in python.

David M. J. Lazer, Matthew A. Baum, Yochai Benkler, Adam J. Berinsky, Kelly M. Greenhill, Filippo Menczer, Miriam J. Metzger, Brendan Nyhan, Gordon Pennycook, David Rothschild, Michael Schudson, Steven A. Sloman, Cass R. Sunstein, Emily A. Thorson, Duncan J. Watts, and Jonathan L. Zittrain. 2018. The science of fake news. *Science*, 359(6380):1094–1096.

Rafael A. Monteiro, Roney L. S. Santos, Thiago A. S. Pardo, Tiago A. Almeida, Evandro E. S. Ruiz, and Oto A. Vale. 2018. Contributions to the study of fake news in portuguese: new corpus and automatic detection results. In *Computational Processing of the Portuguese Language*, Lecture Notes in Artificial Intelligence, pages 324–334. Springer.

Leandro de Oliveira and Diego B. Claro. 2019. Dptoie: a portuguese open information extraction system based on dependency analysis. Disponível no Repositório da UFBA.

Jeff Z. Pan, Siyana Pavlova, Chenxi Li, Ning Li, Yang Li, and Jie Liu. 2018. Content-based fake news detection using knowledge graphs. In *The Semantic Web – ISWC 2018*, pages 669–683. Springer.

Nils Reimers and Iryna Gurevych. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.

Victoria L. Rubin, Yimin Chen, and Niall J. Conroy. 2015. Deception detection for news: Three types of fakes. In *Proceedings of the Association for Information Science and Technology*, pages 1–4. Wiley.

Lucas Santos, Yuri Silva Souza, Manoel Santos, and Roney Santos. 2025. Análise e uso de grafos de conhecimento para detecção de eventos em notícias

falsas. In *Congresso UFBA 2025: Imaginar futuros – desafios da emergência climática*, pages 214–214, Salvador, Brazil. Universidade Federal da Bahia. Caderno de Resumos.

Roney L. S. Santos and Thiago A. S. Pardo. 2020. Fact-checking for portuguese: knowledge graph and google search-based methods. In *Computational Processing of the Portuguese Language*, Lecture Notes in Artificial Intelligence, pages 195–205. Springer.

Renato M. Silva, Roney L. Santos, Tiago A. Almeida, and Thiago A. Pardo. 2020. Towards automatically filtering fake news in portuguese. *Expert Systems with Applications*, 146:113199.

Fábio Souza, Rodrigo Nogueira, and Roberto Lotufo. 2020. Bertimbau: Pretrained bert models for brazilian portuguese. In *Proceedings of the Brazilian Conference on Intelligent Systems (BRACIS)*. Springer.

Yuri Silva Souza, Lucas Santos, Manoel Santos, and Roney Santos. 2025. Fakecheck 2.0: Aplicativo para detecção de notícias falsas em português. In *Congresso UFBA 2025: Imaginar futuros – desafios da emergência climática*, pages 558–558, Salvador, Brazil. Universidade Federal da Bahia. Caderno de Resumos.

Soroush Vosoughi, Deb Roy, and Sinan Aral. 2018. The spread of true and false news online. *Science*, 359(6380):1146–1151.