AI-Powered Ontology-Based Architecture for Misinformation Detection in Fiction Works

Marcos Arias-González[0009-0007-7955-5808], Juan Manuel Ruiz Muñoz[0009-0008-7529-306X], Pablo Armenteros Cosme[0009-0009-6046-4986], Lucía Isabel Rodríguez González[0009-0002-7334-5141], and Javier Curto Hernández[0009-0002-6279-1189]

BISITE Research Group, Edificio I+D+i - C, C. Espejo, s/n, 37007 Salamanca, Spain
{marcosusal, jmrm30, pablo.arcos, lucia.rodriguez.gonzalez, jcurto}@usal.es

Abstract.
Misinformation is a growing threat in digital environments, especially in fictional media such as films and podcasts, where manipulated content is often subtle yet impactful. This research presents a method based on LLMs and RAG to detect such misinformation. The proposed architecture integrates Large Language Models (LLM) with Retrieval-Augmented Generation (RAG) techniques to construct a dynamic ontology based on verified information, which is subsequently employed to query and authenticate the content of fictional works. This system identifies discrepancies between the verified knowledge base and the information presented in these works. The tool was evaluated using textual data, yielding promising results characterized by high precision and recall, thereby highlighting its potential for detecting misinformation across a variety of fictional media. These findings suggest a robust solution for enhancing content verification in an ever-evolving digital landscape.

Keywords: Misinformation detection · Ontology creation · Retrieval-augmented generation · Fiction works · Large language models

1 Introduction

Online disinformation, especially on social media, has emerged as one of the main threats to the integrity of information. With the rise of digital platforms, fake news spreads more quickly and reaches a wider audience than factual news, being up to 70% more likely to be retweeted and reaching people six times faster [1]. In 2019, organised disinformation campaigns were identified in 70 countries, a 150% increase compared to 2017, many of which used fabricated or manipulated content to influence public opinion [2]. This issue not only affects public perception but also undermines trust in democratic institutions, fosters social polarisation, and compromises decision-making based on objective data [3].

Despite efforts to combat disinformation, it remains a complex, evolving issue with significant societal impact. Even after correction, up to 50% of people continue to be influenced by it [4]. Addressing this challenge requires adaptive methods to keep up with the rapid spread of fake news.

A critical yet overlooked area in disinformation detection is identifying false content in fictional media, such as films, TV shows, and podcasts. These formats, rich in text, imagery, and audio, subtly influence public perceptions and can spread falsehoods. Studies show that individuals are 40% more likely to believe false information when presented visually, such as through memes or videos [5]. During the 2016 U.S. presidential election, the top 20 fake news stories on Facebook gained more engagement than the top stories from major media outlets [6]. These findings emphasize the need for detection tools that analyze not just text, but also visual and auditory content, which can have a stronger and longer-lasting impact.

Artificial intelligence (AI) plays a key role in combating disinformation, particularly in fictional media, which audiences may mistakenly interpret as factual [7]. Convolutional Neural Networks (CNNs) and Long Short-Term Memory (LSTM) networks have proven effective in detecting disinformation in both textual and multimedia content [8] [9]. Data mining and Machine Learning (ML) techniques are essential for identifying patterns in the spread of fake news and must evolve with emerging disinformation forms.

This study presents an automated tool that uses Large Language Model (LLM) agents and Retrieval-Augmented Generation (RAG) techniques to detect and classify disinformation in fictional works. The tool builds a structured knowledge base with verified data, which is used to cross-check texts and multimedia (videos, audio). This method helps identify discrepancies between verified information and content in these media. Addressing disinformation in fictional works is critical, as existing research mainly focuses on fake news, yet fictional content can subtly and powerfully shape perceptions of reality [10].

The paper is organized as follows: Section 2 reviews related work on misinformation detection and the use of ontologies. Section 3 details the methodology, including the system architecture, ontology creation, and verification process. Section 4 presents the results of the system's evaluation, and Section 5 concludes the paper, highlighting the contributions and potential future research directions.

2 Related work

Traditional approaches to fake news detection have relied on NLP and machine learning techniques trained on large annotated corpora. For example, Pérez et al. [11] used SVMs, while other authors such as [12] explored deep learning methods with RNNs and word embeddings to capture contextual nuances.

More recently, a study by Panchal et al. [14] adopted a multimodal approach that combines text, images, and audio, improving the accuracy of detection systems. They integrated CNNs for image analysis and LSTM networks for audio processing, demonstrating an increase in accuracy compared to models that analyze only text.

Despite advances, automated detection systems still face challenges, especially with fictional media. Ontologies offer a valuable solution by structuring knowledge to identify inconsistencies and compare with verified information. Alsmadi et al. [15] emphasize the role of ontologies in verification, particularly on social media, while AI technologies, such as synthetic media generation, can enhance this process but must adhere to ethical principles [16].

Tools like OntoStudio [17] and Protégé [18] aid in ontology construction, though automation still requires human validation of relationships. Recent studies, such as Albukhitan et al. [19] with Continuous Bag of Words (CBOW) and Skip-Gram models, and Chen et al. [20] using Deep Belief

Networks (DBNs), explore DL for automated ontology creation, focusing more on relationship extraction than full ontology construction.

LLMs have also been demonstrated to have the capability to generate knowledge graphs, as seen in the works of Pan et al. [21] and the KGfiller system [22]. However, they have yet to achieve fully autonomous and comprehensive management of the information required for disinformation verification. The main challenge lies in developing a dynamic and self-sustaining knowledge base capable of managing and updating information in real time.

This study aims to address these challenges by using LLM agents and RAG techniques for automating ontology creation and disinformation verification, improving accuracy, autonomy, and efficiency in detection while enabling seamless queries and updates to a continuously evolving knowledge base.

Although the automatic construction of knowledge structures has been explored to some extent, most existing research focuses on manually curated ontologies or domain-specific taxonomies. Moreover, the detection of misinformation has traditionally centered on factual news, with fictional narratives remaining largely underexplored. This paper contributes by addressing both gaps: introducing a method for dynamic ontology generation and applying it in a less commonly studied domain.

3 Methodology

The workflow of the tool (Figure 1) can be broadly divided into three main stages: Information Fusion (Section 3.1), Ontology Creation (Section 3.2), and Fictional Works Verification (Section 3.3).

3.1 Architecture for information fusion from different sources and formats

To analyse fictional works in their various forms, it is essential to process different existing formats. This includes .mp4 files, such as films and documentaries; .mp3 files, such as podcasts and audiobooks; as well as written formats, among others.

To address this need, an architecture has been designed to enable the extraction and normalisation of information (Figure 1) from these fictional works. The objective of this architecture is to facilitate processing by converting all formats into a common standard.

Two types of models were used to address this issue. The first model was responsible for converting audio into text, specifically Whisper OpenAI (Base) [23]. The second model, obtained from HuggingFace, generated descriptions from images using the *Salesforce/instructblip-vicuna-7b* model [24], transforming them into text. Additionally, the SceneDetect library [25] was employed to segment videos, breaking them down into audio and images. This allowed for the processing of videos using the two previously mentioned models.

The output format considered, as mentioned earlier, was text-based, as it facilitated the processing of all information using pre-trained large language models.

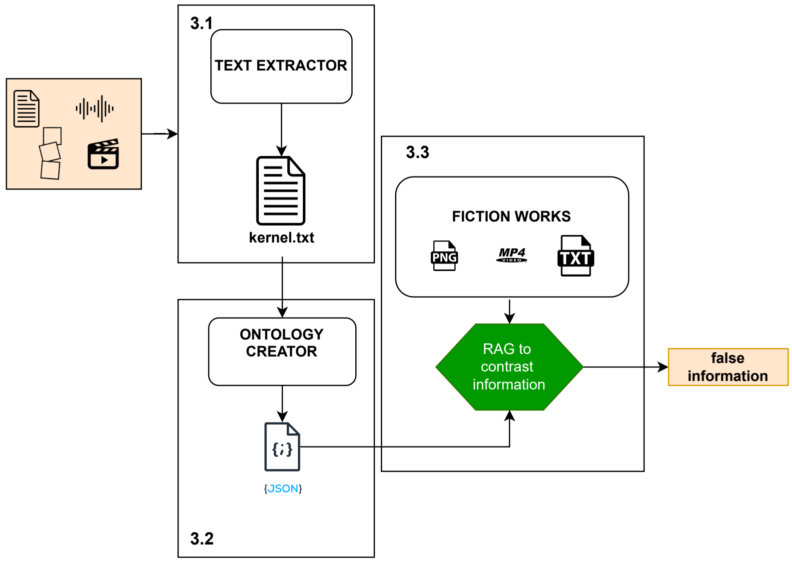

Fig. 1: Overview of the false information detection system for fiction works. The pipeline includes (1) a text extractor, (2) an ontology creator structuring data in JSON, and (3) a RAG-based verification module that analyzes fiction works in various formats to detect false information.

3.2 Ontology creation architecture

This section outlines the iterative process through which the ontology is created and updated, serving as the knowledge base for detecting disinformation in fictional works. Based on the previously obtained and cleaned text, the system sequentially processes each new sentence, dynamically incorporating it into the ontology, which is represented in JSON format. Each node in the ontology corresponds to a class and is organized into three fundamental sections: *Properties*, which store detailed and specific information about the class; *Superclasses*, which group the higher-level classes from which the current class is derived; and *Subclasses*, which contain the specializations or subcategories derived from the class.

The core of the process is based on the calculation and comparison of vector representations (*embeddings*). For each new sentence, an *embedding* is generated using *paraphrase-MiniLM-L6-v2*, a model based on the Sentence-BERT (SBERT) architecture [26], capturing the semantic meaning of the fragment. Simultaneously, the existing ontology is enriched with *embeddings* computed at two levels: at the class level, where a global *embedding* integrates all properties of the class, and at the property level, where a specific *embedding* is calculated for each attribute. These *embeddings* are stored in a vector database enabling efficient queries through cosine similarity calculations.

In each iteration, the *embedding* of the new sentence is compared with the existing *embeddings* in the ontology. A similarity threshold of 0.75 is applied to determine the next action. If the similarity between the *embedding* of the sentence and any of the existing nodes (calculated at either level) exceeds this threshold, the new information is considered to have a strong semantic relationship with that node. In this case, the intelligent agent, leveraging *GPT-4o* [28] with a temperature setting of 0.5 and top_p of 0.9, takes as input the ontology fragment identified as the most similar, along with the new sentence and its corresponding contextual paragraph. Using these elements, the model generates an updated fragment in JSON format, integrating the new information while preserving the defined hierarchical structure.

Conversely, if the similarity falls below the threshold, a new node is created. In this case, in addition to the new sentence and its context, the agent is provided with a list of the names of existing nodes, enabling it to assess whether the information corresponds to an already existing class or if a new class needs to be generated. If the agent generates a node with a name identical to an existing one, a node fusion mechanism is activated, merging the duplicate information, eliminating redundancies, and ensuring the consistency of the ontology.

The process operates in a fully incremental manner, generating and storing new *embeddings* for each fragment—both at the class and property levels—which serve as references for comparing future entries. Additionally, cleaning and validation functions are integrated to rectify potential errors in JSON generation, eliminate redundancies, and ensure the proper maintenance of the hierarchy between nodes (i.e., coherence in the relationships between *Superclasses* and *Subclasses*). This set of techniques ensures that the ontology evolves dynamically and in a structured manner, facilitating subsequent queries and semantic analyses.

3.3 Architecture for integrating new works of fiction and verifying consistency with ontology

The process by which the system integrates new works of fiction and evaluates their consistency with the ontology, considered the verified knowledge base, is described here. The workflow begins with converting the work of fiction into plain text using the information fusion module. Once converted into text, RAG techniques are employed to compare the information with the ontology.

The comparison process is carried out by computing the *embeddings* using the *all-MiniLM-L6-v2* model, also based on SBERT [26], but more efficient due to its training on over a billion sentence pairs. This makes it a high-performance, general-purpose model that balances quality with being five times faster than larger models [27]. Embeddings are calculated at the property level, and similar fragments are retrieved from the vector database, enriched with related classes.

The number of words per query is limited to provide the LLM model with sufficient contextual information while preventing overload. The LLM agent, implemented using *GPT-4o Mini* with a temperature of 0 [29], is configured for consistent and deterministic responses. A temperature of 0 minimizes output variability, ensuring that the model always generates similar responses for the same input. This reduces fluctuations, providing stable and comparable metrics. If a discrepancy is detected, the model generates a report detailing whether the cause is a lack of information or misinformation in the work of fiction.

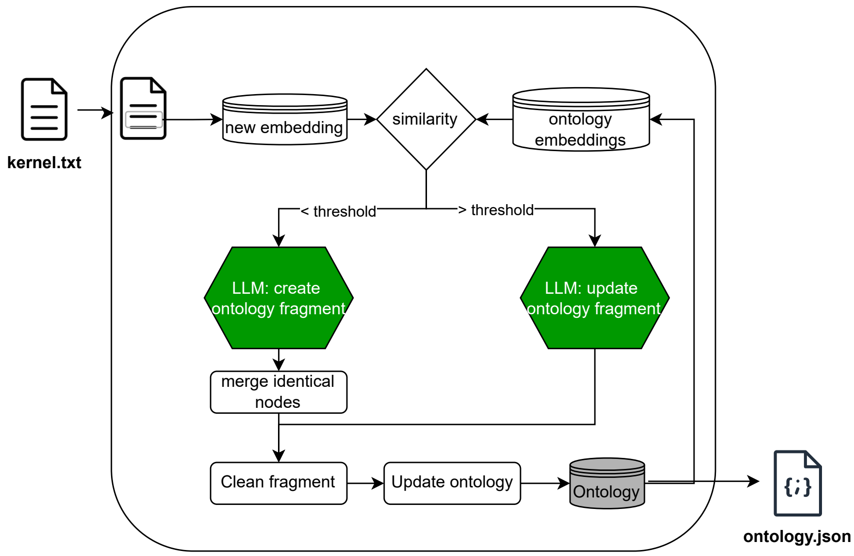

Fig. 2: General workflow of the false information detection tool. The system generates embeddings from input data and compares them to existing ontology embeddings. Based on a similarity threshold, a LLM either creates a new ontology fragment or updates an existing one. The ontology is then refined, updated, and stored in `ontology.json`.

As illustrated in Figure 2, this process enables the efficient integration of new information into the ontology and facilitates the accurate verification of concordance between the content of the work of fiction and the knowledge base.

4 Results

To evaluate the performance of the ontology generation and concordance verification modules, the extracted text from the Wikipedia entry on the Solar System [30] was used as a case study. This text comprises approximately 7,300 words. Using the creator module (Section 3.2), an ontology was generated, structured as shown in Listing 1, consisting of 246 classes. It is important to highlight that the content length within each class varies significantly: some classes contain a large amount of information in their *Properties* section, while others, especially those referring to more specific aspects, present more concise information. In this regard, the average number of words per class is 18.66, with a standard deviation of 31.4.

Listing 1: Reduced JSON fragment of the *SolarSystem* class from the generated ontology

```
{
  "SolarSystem": {
    "Properties": {
      "Description": "The gravitationally bound system of the Sun and the
          objects that orbit it.",
      "..."
    },
    "Superclasses": [
      "MilkyWay",
      "..."
    ],
    "Subclasses": [
      "Planets",
      "..."
    ]
  }
}
```

To assess the system's ability to verify the concordance between new information and the knowledge base, the same text used for ontology generation was divided into 80 paragraphs. From each paragraph, two sentences containing accurate information and two sentences with altered information (by modifying a concept or numerical value) were extracted, resulting in a total of 320 evaluation sentences. The verification was conducted using the concordance module, producing the following confusion matrix and evaluation metrics:

	Pred. 0 (No)	**Pred. 1 (Yes)**
True 0 (No)	VN (158)	FP (2)
True 1 (Yes)	FN (38)	VP (122)

Table 1: Confusion matrix resulting from the verification of concordance between evaluation sentences and the knowledge base.

	Prec.	**Recall**	**F1-Score**	**Accuracy**
Value	0.98	0.76	0.86	0.88

Table 2: Evaluation metrics derived from the confusion matrix, assessing the system's ability to verify concordance.

These results indicate that, although the system does not detect all available accurate information, it is highly precise in avoiding the acceptance of misinformation. The low incidence of false positives (only 2 cases) confirms its reliability in this regard.

However, the recall rate of 76% highlights some limitations in the system's ability to identify true information, as evidenced by the 38 false negatives. This relatively low recall suggests that, while the system demonstrates high precision in rejecting misinformation (as reflected by the minimal false positives), it occasionally fails to capture all accurate content. While this work presents a single case study, the proposed methodology is designed to be domain-independent. Future work will focus on evaluating its performance across diverse contexts to assess its generalizability. This may result from missing concepts during ontology generation, creating gaps in the knowledge base. The verification module also struggles when new information subtly differs from existing content. These factors contribute to missed true information, pointing to the need for further optimization of both the ontology creation and verification processes in order to enhance recall and ensure more comprehensive detection of accurate content.

To mitigate these issues and improve recall, we propose two main strategies: (1) enriching the ontology with external verified sources beyond the initial dataset, to broaden its coverage and reduce omissions; and (2) implementing feedback-based refinement cycles, where undetected true positives (false negatives) are reintegrated iteratively into the knowledge base, allowing the system to learn from its mistakes and fill semantic gaps over time.

One key issue is the LLM agent's tendency to omit details despite explicit prompt instructions. The ontology creator includes mechanisms to detect and correct such errors, but these can also discard newly generated classes that do not meet predefined criteria, further contributing to information loss. Likewise, the verification module, also LLM-based, may struggle to extract the correct semantic relationships for accurate validation. These limitations can result in incomplete ontologies, which directly impact the system's performance by reducing recall and increasing semantic ambiguity. For instance, the omission of intermediate concepts can hinder the system's ability to link related entities during verification, leading to missed detections even when partial evidence exists in the knowledge base. Improving recall may require refining the ontology creation process and enhancing the verification module's ability to capture deeper semantic connections.

5 Conclusions and future work

The system has shown highly promising results, with strong performance in both information integration and verification. Its efficient ontology update mechanism, regardless of knowledge base size, marks a significant step forward in managing reliable information. Scalability is supported through sharding and parallel processing, with embeddings precomputed and indexed to maintain near-constant update times. For video, content is segmented into scenes and a small subset is analyzed, focusing on the most informative parts and minimizing data load. The system's low rate of false positives further confirms its robustness.

Despite these strengths, several areas require optimization. The unification module could better integrate data from different modalities, especially in multimedia processing. In video analysis, integrating audio and visual information simultaneously would allow more precise correlation and improve detection accuracy. To this end, future evaluations will involve real multimedia inputs, using Whisper [23] for audio segmentation and PySceneDetect [25] for selecting representative video frames. These tools will enable the creation of aligned multimodal datasets to assess misinformation detection beyond textual data. Multimodal analysis also allows the extraction of visual patterns and audio cues, contributing to a deeper understanding of the message and enhancing detection in complex scenarios.

While LLMs facilitate semantic structuring, they also present challenges such as hallucinations and performance variability [31]. We consider integrating more advanced embedding models and fine-tuning techniques—e.g., with DeepSeek—to improve the fidelity of generated ontologies and the quality of verification reports. These improvements aim to capture complex elements like irony, ambiguity, and implicit relationships. Although the system currently focuses on factual mismatches, interpreting abstract or metaphorical content will require deeper contextual modeling and, potentially, emotion recognition or figurative language analysis.

Acknowledgements

This publication is part of the project "TRUstworthy artificial intElligence over NPL to fight againST disinfORmation InstrumEnts in fiction (TRUESTORIES)" (CPP2021-008358), funded by MICIU/AEI/10.13039/501100011033 and by the European Union "NextGenerationEU"/PRTR".

References

1. Vosoughi, S., Roy, D., Aral, S.: The spread of true and false news online. Science **359**(6380), 1146–1151 (2018). https://doi.org/http://doi.org/10.1126/science.aap9559
2. Bradshaw, S., Howard, P. N.: The Global Disinformation Disorder: 2019 Global Inventory of Organised Social Media Manipulation. Oxford Internet Institute, Oxford, UK (2019). https://www.oii.ox.ac.uk/news-events/reports/the-global-disinformation-order-2019-global-inventory-of-organised-social-media-manipulation/
3. Lazer, D.M.J., Baum, M.A., Benkler, Y., Berinsky, A.J., Greenhill, K.M., Menczer, F., Metzger, M.J., Nyhan, B., Pennycook, G., Rothschild, D., Schudson, M., Sloman, S.A., Sunstein, C.R., Thorson, E.A., Watts, D.J., Zittrain, J.L.: The science of fake news. Science **359**(6380), 1094–1096 (2018). https://doi.org/http://doi.org/10.1126/science.aao2998
4. Lewandowsky, S., Ecker, U.K.H., Seifert, C.M., Schwarz, N., Cook, J.: Misinformation and its correction: Continued influence and successful debiasing. Psychol. Sci. Public Interest **13**(3), 106–131 (2012). https://doi.org/http://doi.org/10.1177/1529100612451018
5. Gordon Pennycook, David G. Rand: The Psychology of Fake News. Trends in Cognitive Sciences **25**(5), 388–402 (2021). https://doi.org/http://doi.org/10.1016/j.tics.2021.02.007
6. Graham, D.A.: Some real news about fake news. The Atlantic, Section: Ideas (2019). https://www.theatlantic.com/ideas/archive/2019/06/fake-news-republicans-democrats/591211/, last accessed 2025/03/13
7. Noémi Bontridder, Yves Poullet: The Role of Artificial Intelligence in Disinformation. ResearchGate (2021). https://doi.org/http://doi.org/10.13140/RG.2.2.28805.27365.
8. Güler, G., Gündüz, S.: Deep learning based fake news detection on social media. Int. J. Inf. Secur. Sci. **12**(2), 1–21 (2023). https://doi.org/http://doi.org/10.55859/ijiss.1231423
9. Alim Al Ayub Ahmed, Ayman Aljarbouh, Praveen Kumar Donepudi, Myung Suh Choi: Detecting Fake News using Machine Learning: A Systematic Literature Review. Psychology and Education Journal **58**(1), 1932–1939 (2021). https://doi.org/http://doi.org/10.17762/pae.v58i1.1046
10. Biarritz Alain Evangelista, Erix Laud Reyes, Bernard Jezua Tandang: Reality Check: A Comprehensive Review of AI-Generated Media and Detection Methods. SSRN Scholarly Paper (2024). https://doi.org/http://doi.org/10.2139/ssrn.5059062.
11. Pérez-Rosas, V., Kleinberg, B., Lefevre, A., Mihalcea, R.: Automatic detection of fake news. arXiv preprint arXiv:1708.07104 (2017). https://doi.org/http://doi.org/10.48550/arXiv.1708.07104
12. Kolluru, V., Mungara, S., Chintakunta, A.N.: Combating misinformation with machine learning: Tools for trustworthy news consumption. Mach. Learn. Appl. Int. J. **7**(4), 28–39 (2020). https://doi.org/http://doi.org/10.5121/mlaij.2020.7403
13. Eslam Amer, Kyung-Sup Kwak, Shaker El-Sappagh: Context-Based Fake News Detection Model Relying on Deep Learning Models. Electronics **11**(8), 1255 (2022). https://doi.org/http://doi.org/10.3390/electronics11081255
14. Panchal, S.: Research paper on AI-driven detection of fake news in non-textual content. Int. J. Res. Appl. Sci. Eng. Technol. **12**(12), 150–161 (2024). https://doi.org/http://doi.org/10.22214/ijraset.2024.65492

15. Alsmadi, I., Alazzam, I., AlRamahi, M.A.: An ontological analysis of misinformation in online social networks. arXiv preprint arXiv:2102.11362 (2021). https://doi.org/http://doi.org/10.48550/arXiv.2102.11362
16. Maia, C.H., Ariel, P., Nunes, S.: Adding human values on the deepfake: co-designing fact-checking solutions to combat misinformation. AI Ethics (2024). https://doi.org/http://doi.org/10.1007/s43681-024-00619-y
17. Weiten, M.: OntoSTUDIO® as a ontology engineering environment. In: Davies, J., Grobelnik, M., Mladenić, D. (eds.) Semantic Knowledge Management: Integrating Ontology Management, Knowledge Discovery, and Human Language Technologies, pp. 51–60. Springer, Berlin, Heidelberg (2009). https://doi.org/http://doi.org/10.1007/978-3-540-88845-1_5
18. Hernandez, J., Vegega, C., Pollo-Cattaneo, M.F.: Revisión sistemática de literatura sobre herramientas ontológicas para representar conocimientos. In: Actas del Congreso CACIC 2024, pp. 135–145. https://cacic2024.info.unlp.edu.ar/wp-content/uploads/2024/10/Libro-de-Actas-CACIC-2024-Ebook_.pdf (2024).
19. Saeed Albukhitan, Tarek Helmy, Ahmed Alnazer: Arabic ontology learning using deep learning. In: Proceedings of the International Conference on Web Intelligence, pp. 1138–1142. Association for Computing Machinery, New York, NY, USA (2017). https://doi.org/http://doi.org/10.1145/3106426.3109052
20. Yu Chen, Wenjie Li, Yan Liu, Dequan Zheng, Tiejun Zhao: Exploring Deep Belief Network for Chinese Relation Extraction. In: CIPS-SIGHAN Joint Conference on Chinese Language Processing (2010). https://aclanthology.org/W10-4115/
21. Pan, J.Z., Razniewski, S., Kalo, J.-C., Singhania, S., Chen, J., Dietze, S., Jabeen, H., Omeliyanenko, J., Zhang, W., Lissandrini, M., Biswas, R., de Melo, G., Bonifati, A., Vakaj, E., Dragoni, M., Graux, D.: Large language models and knowledge graphs: Opportunities and challenges. Trans. Graph Data Knowl. **1**(1), 2:1–2:38 (2023). https://doi.org/http://doi.org/10.4230/TGDK.1.1.2
22. Giovanni Ciatto, Andrea Agiollo, Matteo Magnini, Andrea Omicini: Large language models as oracles for instantiating ontologies with domain-specific knowledge. Knowledge-Based Systems **310**, 112940 (2025). https://doi.org/http://doi.org/10.1016/j.knosys.2024.112940
23. Radford, A., Kim, J. W., Xu, T., Brockman, G., McLeavey, C., Sutskever, I.: Robust speech recognition via large-scale weak supervision. arXiv preprint arXiv:2212.04356 (2022). https://doi.org/http://doi.org/10.48550/arXiv.2212.04356
24. Dai, W., Li, J., Li, D., Tiong, A. M. H., Zhao, J., Wang, W., Li, B., Fung, P., Hoi, S.: InstructBLIP: towards general-purpose vision-language models with instruction tuning. arXiv preprint arXiv:2305.06500 (2023). https://doi.org/http://doi.org/10.48550/arXiv.2305.06500
25. Castellano, B.: PySceneDetect: tool for scene detection in videos. (2025). https://www.scenedetect.com/
26. Reimers, N., Gurevych, I.: Sentence-BERT: sentence embeddings using Siamese BERT-networks. In: Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing. Association for Computational Linguistics (2019). https://doi.org/http://doi.org/10.48550/arXiv.1908.10084
27. Reimers, N., Gurevych, I.: Sentence-BERT: Sentence embeddings using Siamese BERT-networks. arXiv preprint arXiv:1908.10084 (2019). https://doi.org/http://doi.org/10.48550/arXiv.1908.10084
28. OpenAI: Hello GPT-4o. (2024). https://openai.com/index/hello-gpt-4o/
29. OpenAI: GPT-4o Mini. (2024). https://openai.com/index/gpt-4o-mini-advancing-cost-efficient-intelligence/
30. Wikipedia contributors: Solar system. Wikipedia, The Free Encyclopedia. (2024). https://en.wikipedia.org/wiki/Solar_System
31. Bang, Y., Cahyawijaya, S., Lee, N., Dai, W., Su, D., Wilie, B., Lovenia, H., Ji, Z., Yu, T., Chung, W., Do, Q.V., Xu, Y., Fung, P.: A Multitask, Multilingual, Multimodal Evaluation of ChatGPT on Reasoning, Hallucination, and Interactivity. arXiv preprint arXiv:2302.04023 (2023). https://doi.org/http://doi.org/10.48550/arXiv.2302.04023