Cyprus University of Technology
Department of Electrical Engineering and Computer Engineering and Informatics

Bachelor Thesis

An Evaluation of Automated Fact-Checking using Grounding Large Language Models with Knowledge Graphs

Sotiris Gypsiotis

Limassol, May 2025

CYPRUS UNIVERSITY OF TECHNOLOGY

Faculty of Engineering and Technology

Department of Electrical Engineering, Computer Engineering, and Informatics

Bachelor Thesis

An Evaluation of Automated Fact-Checking using Grounding Large Language Models with Knowledge Graphs

Sotiris Gypsiotis

Advisor:

Dr. Michael Sirivianos,

Dean of the School of Engineering and Technology

Limassol, May 2025

Copyrights

Acknowledgements

I would like to especially thank Dr. Michael Sirivianos, Dr. Nikos Salamanos, and PhD student Pantelitsa Leonidou of the Department of Electrical Engineering, Computer Engineering and Informatics, for their help and support throughout this thesis. Thank you for giving me the opportunity to undertake a topic of such interest and importance.

ABSTRACT

Automated fact-checking has become an essential tool in combating the widespread spread of misinformation in the digital age. This thesis explores the integration of Large Language Models with Knowledge Graphs to enhance fact-checking systems, for both structured and unstructured data, in order to improve accuracy and scalability. LLMs, particularly Transformer-based architectures such as GPT and DeepSeek, have demonstrated remarkable capabilities in understanding and generating natural language. However, their sensitivity to hallucinations and reliance on probabilistic reasoning present challenges in ensuring factual accuracy. In contrast, KGs provide structured, relational data that encodes factual knowledge in a deterministic manner. By combining the contextual language understanding of LLMs with the structured reasoning of KGs, automated fact-checking systems can achieve greater robustness and reliability.

This thesis investigates the mechanisms of both LLMs and KGs, detailing their strengths and limitations in the fact-checking pipeline. The research first provides an overview of AI, ML, and DL, focusing on their applications in misinformation detection. It then examines the construction and representation of KGs, including methods for entity recognition, relationship extraction, and integration with language models. Furthermore, we explore the role of LLMs in claim parsing, evidence retrieval, and truthfulness assessment. The integration of these technologies introduces challenges such as knowledge alignment, consistency, computational efficiency, and understandability.

To evaluate the effectiveness of this integration, this work implements an automated fact-checking pipeline using the Ollama framework and three DeepSeek R1 models: 1.5B, 7B, and 8B. A dataset of 200 factual statements was sourced from the Wikidata KG and used to test the models' ability to extract, validate, and compare facts. Evaluation was conducted using classification metrics (Precision, Recall, F1-score) and lexical overlap metrics (ROUGE score). The 7B model achieved the highest F1-score (0.58), outperforming both the smaller 1.5B model and the larger 8B model. Notably, the 8B model showed signs of overfitting and increased hallucination, despite its larger parameter count. These results suggest that model size does not linearly correlate with performance, and that mid-sized models, like the 7B, may provide better generalization and factual alignment in structured verification tasks.

The findings of this thesis contribute to the field of AI-driven misinformation detection, offering insights into how the cooperation between LLMs and KGs can lead to more accurate, transparent, and scalable fact-checking solutions. This thesis concludes by discussing future directions, including the potential for real-time claim verification and ethical considerations in AI-driven fact-checking. By advancing the integration of structured and unstructured data sources, this work lays the groundwork for more reliable automated verification systems, ultimately for a more trustworthy information system.

Keywords: Automated Fact-Checking, Large Language Models (LLM), Misinformation, Knowledge Graph (KG)