Attention-Augmented Knowledge Graph Alignment for Evaluating LLM Factual Consistency

Gunapala S.A.C.H[1], Mamta N[1], Subavarshana A[1], Kithuni Deva Wickramasinghe[1], Pragatheeswaran V[2,*], Uthayasanker Tahayasivam[1], Kamal Premaratne[2] [1]University of Moratuwa, Colombo, Sri Lanka [2]University of Miami, Florida, USA

Large language models (LLMs) frequently generate responses that appear semantically coherent yet subtly diverge from ground-truth facts. To systematically diagnose such deviations, we propose a unified framework for knowledge graph (KG) comparison and triplet-level alignment that balances structural rigor with semantic flexibility. First, we revisit graph-kernel–based comparison methods such as KEA, which combine semantic embedding, clustering, and Weisfeiler–Lehman (WL) structural kernels. While effective, semantic clustering introduces threshold sensitivity and may collapse contextually distinct labels. To address this limitation, we introduce an attention-augmented alignment mechanism that replaces rigid clustering with context-sensitive neighborhood attention, enabling implicit semantic alignment prior to structural kernel comparison. Second, at the triplet level, we complement structural similarity with sentence-level semantic alignment to improve recall under relation paraphrasing. By separating broad semantic alignment from interpretable error typing (entity substitution versus relation substitution), the framework provides both quantitative similarity scores and fine-grained diagnostic insights. Empirical evaluation demonstrates that attention-based alignment reduces preprocessing complexity while maintaining structural fidelity, and sentence-level matching significantly improves recall without sacrificing interpretability. Together, the proposed system offers a robust and scalable approach for evaluating LLM factual consistency through knowledge graph similarity and triplet alignment analysis.

Keywords: Hallucination; Graph Kernels; Semantic Similarity; Factual Consistency.