

A Benchmark Knowledge Graph of Driving Scenes for Knowledge Completion Tasks

Ruwan Wickramarachchi¹, Cory Henson² and Amit Sheth¹

¹AI Institute, University of South Carolina, Columbia, SC, USA.

²Bosch Center for Artificial Intelligence, Pittsburgh, PA, USA.

Abstract

Knowledge graph completion (KGC) is a problem of significant importance due to the inherent incompleteness in knowledge graphs (KGs). The current approaches for KGC using link prediction (LP) mostly rely on a common set of benchmark datasets that are quite different from real-world industrial KGs. Therefore, the adaptability of current LP methods for real-world KGs and domain-specific applications is questionable. To support the evaluation of current and future LP and KGC methods for industrial KGs, we introduce **DSceneKG**, a suite of real-world driving scene knowledge graphs that are currently being used across various industrial applications. The DSceneKG is publicly available at: <https://github.com/ruwantw/DSceneKG>.

1. Introduction

Knowledge completion is an important problem in knowledge representation and reasoning, particularly due to the inherent incompleteness of KGs. LP methods aim to address this by predicting new links and entities to fill in the missing information. Over the years, various LP methods, categorized into matrix factorization, geometric, or deep learning-based approaches, have been developed and primarily evaluated using common benchmark datasets like Freebase and WordNet.

While such benchmark datasets offer a standardized platform for evaluating LP methods, they do not always accurately capture the complexities of real-world industrial applications. Industries such as automotive, manufacturing, healthcare, and finance are creating large-scale industrial KGs to represent domain-specific data. For instance, in the automotive industry, there is a growing demand for large-scale knowledge graphs to represent multi-modal driving scene data from various sensors and cameras, conforming to domain-specific ontologies developed by subject matter experts (SMEs).

The reliance on benchmark datasets for evaluating LP methods raises concerns about their applicability to industrial KGs. Real-world KGs differ significantly from benchmark datasets in terms of structure, modality, conformance to ontology, in/out degree, cardinality, etc. Industrial KGs often involve multimodal data, including text, images, and sensor data, and exhibit a higher degree of heterogeneity. This discrepancy calls into question the effectiveness of current LP methods when applied to real-world KGs, highlighting the need for new evaluation frameworks and the availability of industrial KGs to the demands of industrial applications.

Posters, Demos, and Industry Tracks at ISWC 2024, November 13–15, 2024, Baltimore, USA

*This work was supported in part by NSF grants #2119654 and #2133842. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the NSF.

✉ ruwan@email.sc.edu (R. Wickramarachchi); cory.henson@us.bosch.com (C. Henson); amit@sc.edu (A. Sheth)



© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

2. DSceneKG: Driving Scenes Knowledge Graph

DSceneKG is a suite of knowledge graphs, developed as a collaborative effort between Bosch and the University of South Carolina, to represent real-world driving data from openly available, real autonomous driving datasets such as Pandaset and NuScenes. These datasets represent diverse driving scenarios, including urban/rural environments, various weather conditions, traffic situations, left/right-hand driving, and scenes from different continents. The Scene Ontology defines high-level object/event classes and their relationships [1]. DSceneKG instantiates all scene elements with metadata such as spatial coordinates, time, and descriptions where available. In the Scene Ontology, scenes are categorized into two types: (1) *Sequence Scene* – A video of 10-20 seconds, with a location region and temporal range; (2) *Frame Scene* – A sampled snapshot from the video, with a location point and timestamp. Table 1 shows some statistics about two KGs developed for the automated driving domain.

Table 1: Statistics of two KG variants

	NuScenes	Pandaset
#triples	6,296,378	3,301,928
#entities	2,108,545	53,248
#relations	14	19
Avg. in-degree	3.0353	62.1387
Avg. out-degree	3.0107	63.3269
Triples/entities	2.9861	62.0104

3. Benchmark Tasks of DSceneKG

The versatility of DSceneKG is demonstrated through various KGC/ knowledge-based tasks both within and outside Bosch. The tasks modeled exclusively as KGC are denoted by †.

1. Knowledge-based entity prediction (KEP)† - enabling a knowledge-based approach for predicting entities in driving scenes [1]
2. Context-based method for labeling unobserved entities (CLUE)† - completing AD datasets with labels for entities[2] that may have gone unobserved or unlabeled.
3. Explainable scene clustering - typing automotive scenes into explainable, high-level semantic clusters[3]
4. Semantic-based scene similarity - identifying automotive scenes that are semantically similar, but may be visually dissimilar[4]
5. Causal discovery† - enabling root-cause analysis/ causal discovery in driving scenes[5]
6. Knowledge-based retrieval - enhancing Bird’s-Eye View (BEV) retrieval by integrating semantic representations with textual descriptions [6]

References

- [1] R. Wickramarachchi, et al., Knowledge-based entity prediction for improved machine perception in autonomous systems, *IEEE Intelligent Systems* (2022).
- [2] R. Wickramarachchi, et al., CLUE-AD: a context-based method for labeling unobserved entities in autonomous driving data, *AAAI* (2023).
- [3] S. Nag Chowdhury, et al., Towards leveraging commonsense knowledge for autonomous driving, 2021.
- [4] R. Wickramarachchi, et al., An evaluation of knowledge graph embeddings for autonomous driving data: Experience and practice, *AAAI-MAKE* (2020).
- [5] U. Jaimini, et al., Causalkg: Causal knowledge graph explainability using interventional and counterfactual reasoning, *IEEE Internet Computing* (2022).
- [6] D. Wei, et al., Bev-clip: Multi-modal bev retrieval methodology for complex scene in autonomous driving, *arXiv* (2024).