Using a Knowledge Graph of Scenes to Enable Search of Autonomous Driving Data

Cory Henson[1], Stefan Schmid[1], Tuan Tran[2], Antonios Karatzoglou[2]

[1] Corporate Research, Robert Bosch GmbH
[2] Chassis Systems Control, Robert Bosch GmbH
cory.henson@us.bosch.com,
{stefan.schmid,anhtuan.tran2,antonios.karatzoglou}@de.bosch.com

Abstract. With the recent push to develop autonomous driving technologies, automotive companies are investing heavily into machine learning and AI. Training machine learning models for this task requires access to lots of data. In this talk, we discuss our experience in using semantic technologies to organize and manage this data within a large enterprise company. More specifically, we have developed a knowledge graph of driving scenes and will demonstrate its utility for representing, integrating, and querying large amounts of autonomous driving data.

Keywords: Knowledge Graph, Ontology, Semantic Search, Big Data.

1 Background

Every major automotive company is racing to develop and deploy autonomous driving technology in the coming years. These technologies may be categorized into five distinct levels of automation: (1) driver assistance, (2) partial automation, (3) conditional automation, (4) high automation, and (5) full automation; see [1] for further detail. Bosch is currently developing autonomous driving technologies at all 5 levels. The research and development activities related to these technologies are spread across many different projects, groups, and departments within the company. This distribution presents an obstacle to data-scientists and engineers who require access to the data in order to train machine learning models. To mitigate this issue, Bosch is building an enterprise data lake that centralizes the storage and access of automotive data [2]. This data lake, however, does not solve the difficult and persistent challenge of finding data that's relevant for a particular project or use-case; resulting in a significant lack of data re-use.

2 Approach

Our approach to solving this challenge focuses on representing, integrating, and querying for information about *scenes*. For the purposes of autonomous driving, a scene may be understood as a situation in which an event occurs (e.g. emergency braking) within

some context (e.g. on the highway, with snow on the road). Three primary technologies are developed for the task of representing, integrating, and querying scenes:

1. **Scene Ontology** provides a common definition for the concept of *scene*. This ontology is used to semantically annotate data generated by various use-cases, projects, and departments focused on autonomous driving.
2. **Scene Knowledge Graph** provides a unified semantic representation of scenes. The knowledge graph integrates heterogeneous data and meta-data about a scene from various sources, including sensor data (e.g. video, LIDAR, RADAR), inferences from sensor data (e.g. object recognition), and relevant information available on the web (e.g. map data from Open Street Map [3]); see Figure 1.
3. **Scene-based Data Access** provides an API for ontology-based search. The API utilizes links between the knowledge graph and the data lake to enable query of the autonomous driving data based on semantic descriptions of scenes (e.g. find sensor data related to emergency braking maneuvers, on a highway, with snow on the road).

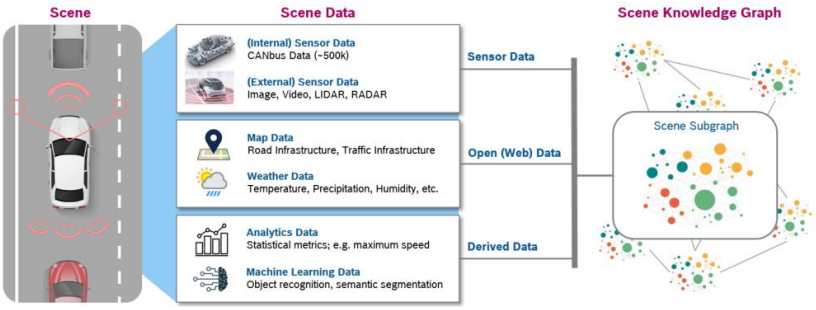

Fig. 1. Illustration of the types of heterogeneous information that are collected, annotated, and integrated into the scene knowledge graph.

3 Conclusion

At Bosch, we are using semantic technologies to improve our ability to represent, integrate, and query autonomous driving data. Our approach focuses on generating a knowledge graph of scenes and using it to support semantic query of data stored in a large-scale enterprise data lake. This knowledge graph is currently deployed and used by various projects and departments, enabling our data-scientists and engineers to find and re-use the data that's relevant for their particular application.

References

1. The 5 Autonomous Driving Levels Explained, https://www.iotforall.com/5-autonomous-driving-levels-explained/ (accessed: 06/25/2019).
2. Schmid, S., Henson, C., Tran, T.: Using Knowledge Graphs to Search an Enterprise Data Lake. In: Proceedings of the 16th Extended Semantic Web Conference (2019).
3. Haklay, M., Weber, P.: OpenStreetMap: User-Generated Street Maps. In: IEEE Pervasive Computing, vol. 7, no. 4, pp. 12-18 (2008).