Knowledge-infused learning for autonomous driving

K-iLKGC21

May 3, 2021

Cory Henson, Bosch Research

BOSCH

Levels of autonomous driving

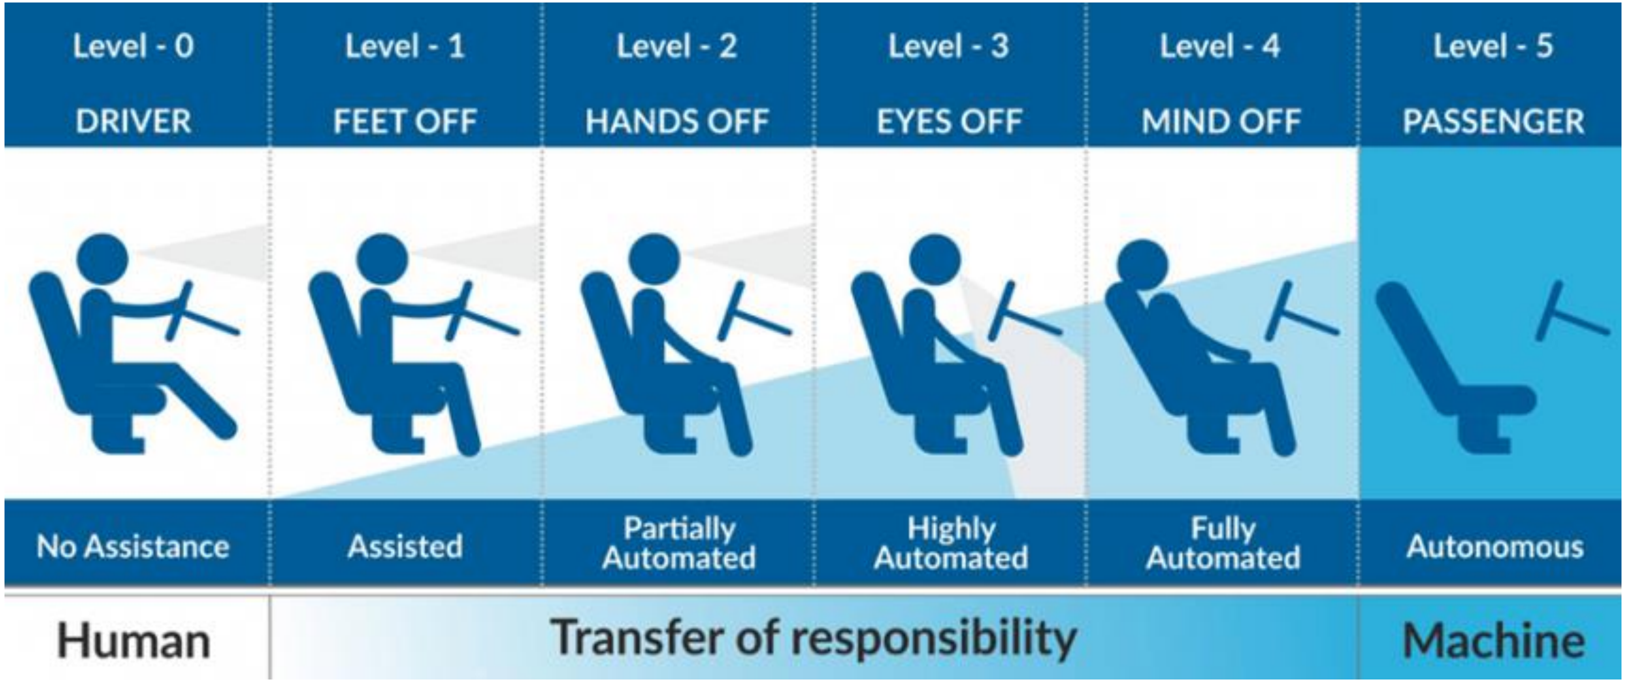

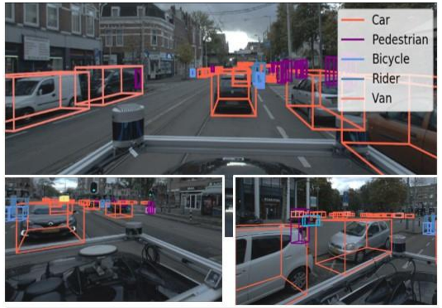

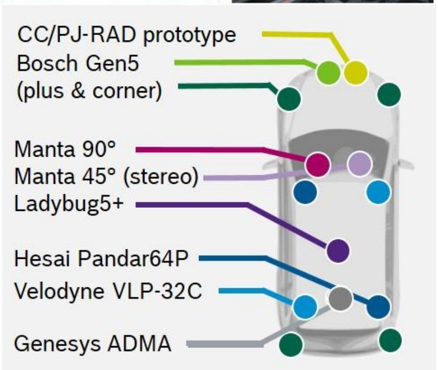

https://blog.netapp.com/how-to-build-a-data-pipeline-for-autonomous-driving/

Sense, think, act

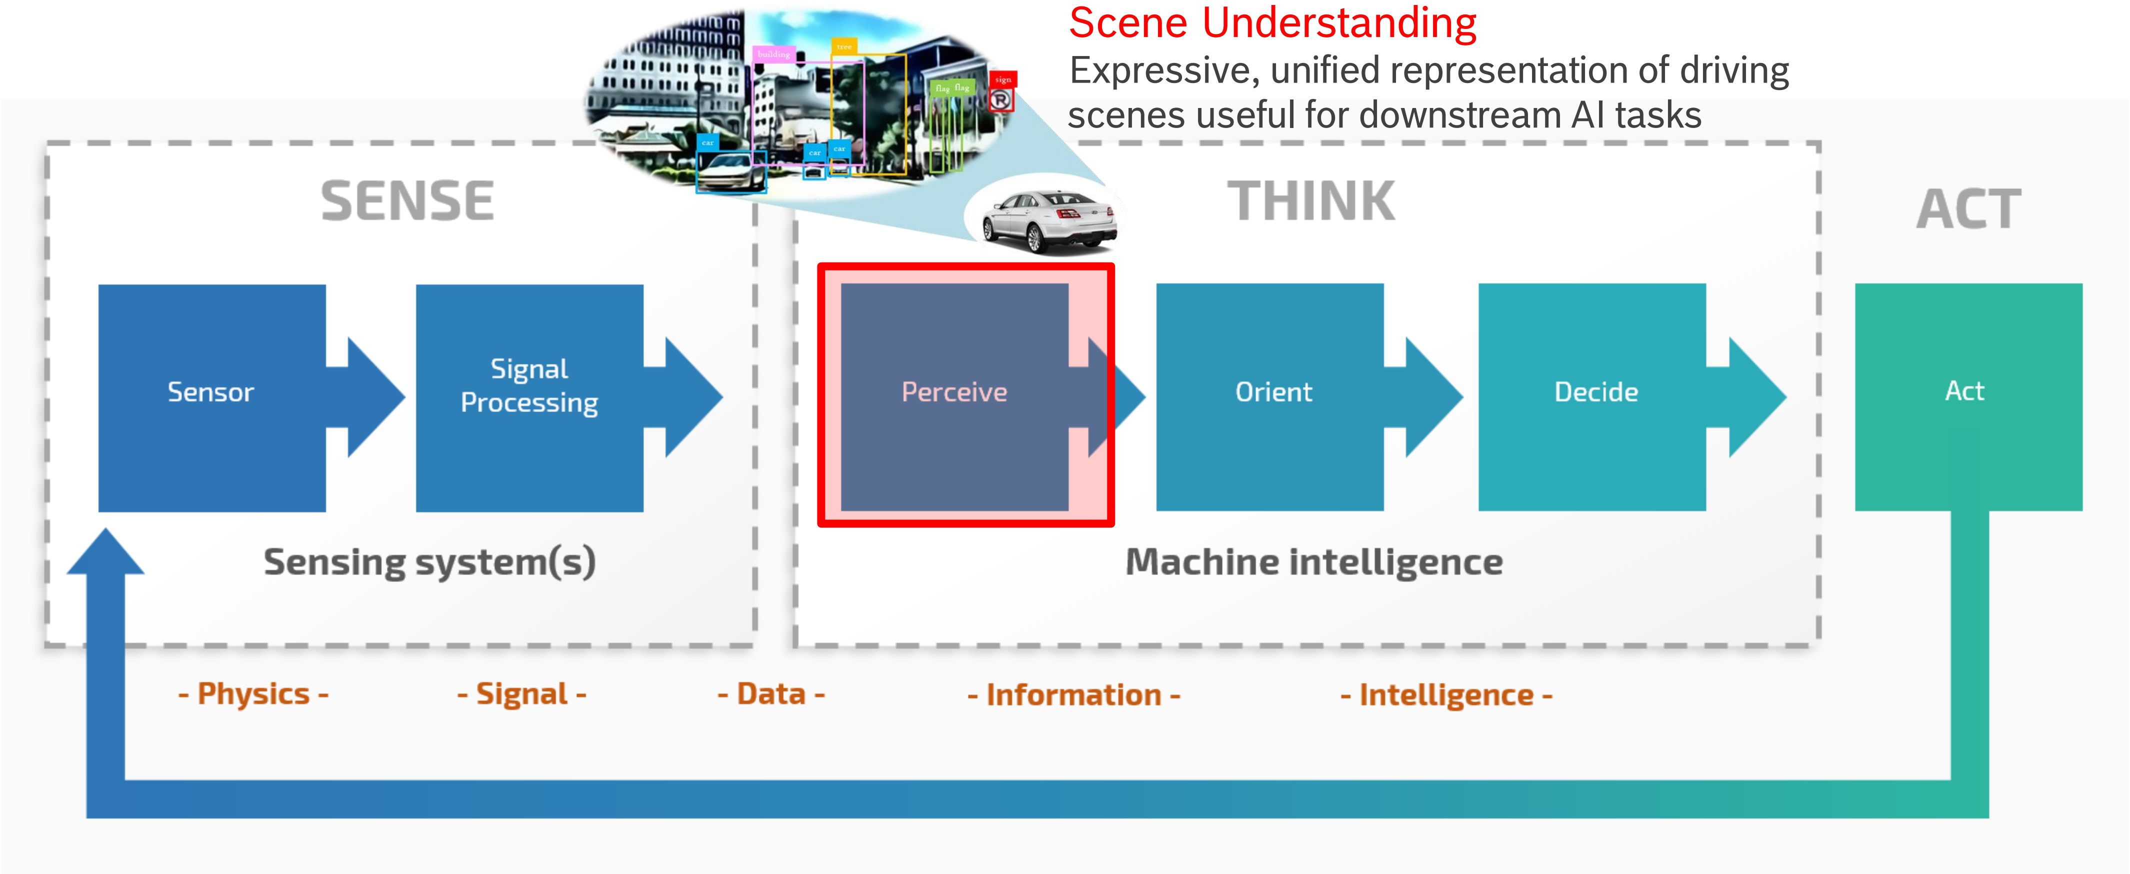

Cory Henson (Bosch Research) | 2021-05-03

https://autonomy.sandia.gov/whatisautonomy/index.html

Relevant projects

Knowledge-infused learning for scene entity prediction Knowledge graph embeddings of driving scenes enable the prediction of additional undetected entities in the scene	AI Institute #AIISC University of South Carolina
Learning visual models using a knowledge graph as a trainer Joint learning of image embeddings with knowledge graph embeddings improves transfer learning on road-sign recognition tasks	Universität Trier
Scene-specific subgraph embeddings Extraction and embedding of subgraphs enables a local, contextualized interpretation of a scene, leading to improved scene understanding	Universität Trier
Explainable embedding-based clustering over knowledge graphs Knowledge graph embeddings of driving scenes create clusters of scenes. Rule-mining techniques are be used to derive descriptions, or explanations, for these clusters	Max Planck Institute for Intelligent Systems

Relevant projects

Knowledge-infused learning for scene entity prediction

Knowledge graph embeddings of driving scenes enable the prediction of additional undetected entities in the scene

AI INSTITUTE #AIISC UNIVERSITY OF SOUTH CAROLINA

Learning visual models using a knowledge graph as a trainer

Joint learning of image embeddings with knowledge graph embeddings improves transfer learning on road-sign recognition tasks

Universität Trier

Scene-specific subgraph embeddings

Extraction and embedding of subgraphs enables a local, contextualized interpretation of a scene, leading to improved scene understanding

Universität Trier

Explainable embedding-based clustering over knowledge graphs

Knowledge graph embeddings of driving scenes create clusters of scenes. Rule-mining techniques are be used to derive descriptions, or explanations, for these clusters

Max Planck Institute for Intelligent Systems

Cory Henson (Bosch Research) | 2021-05-03

Relevant projects

Knowledge-infused learning for scene entity prediction

Knowledge graph embeddings of driving scenes enable the prediction of additional undetected entities in the scene

AI INSTITUTE #AIISC UNIVERSITY OF SOUTH CAROLINA

Learning visual models using a knowledge graph as a trainer

Joint learning of image embeddings with knowledge graph embeddings improves transfer learning on road-sign recognition tasks

Universität Trier

Scene-specific subgraph embeddings

Extraction and embedding of subgraphs enables a local, contextualized interpretation of a scene, leading to improved scene understanding

Universität Trier

Explainable embedding-based clustering over knowledge graphs

Knowledge graph embeddings of driving scenes create clusters of scenes. Rule-mining techniques are be used to derive descriptions, or explanations, for these clusters

Max Planck Institute for Intelligent Systems

Cory Henson (Bosch Research) | 2021-05-03

Relevant projects

Knowledge-infused learning for scene entity prediction

Knowledge graph embeddings of driving scenes enable the prediction of additional undetected entities in the scene

AI INSTITUTE #AIISC UNIVERSITY OF SOUTH CAROLINA

Learning visual models using a knowledge graph as a trainer

Joint learning of image embeddings with knowledge graph embeddings improves transfer learning on road-sign recognition tasks

Universität Trier

Scene-specific subgraph embeddings

Extraction and embedding of subgraphs enables a local, contextualized interpretation of a scene, leading to improved scene understanding

Universität Trier

Explainable embedding-based clustering over knowledge graphs

Knowledge graph embeddings of driving scenes create clusters of scenes. Rule-mining techniques are be used to derive descriptions, or explanations, for these clusters

Max Planck Institute for Intelligent Systems

Knowledge-infused learning for scene entity prediction

In collaboration with Ruwan Wickramarachchi
AI Institute, University of South Carolina

Entity prediction w/ knowledge-infused learning

Goal
Predict entities in a scene, given current and background knowledge of the scene

Sample question
What is the probability that a child is nearby?

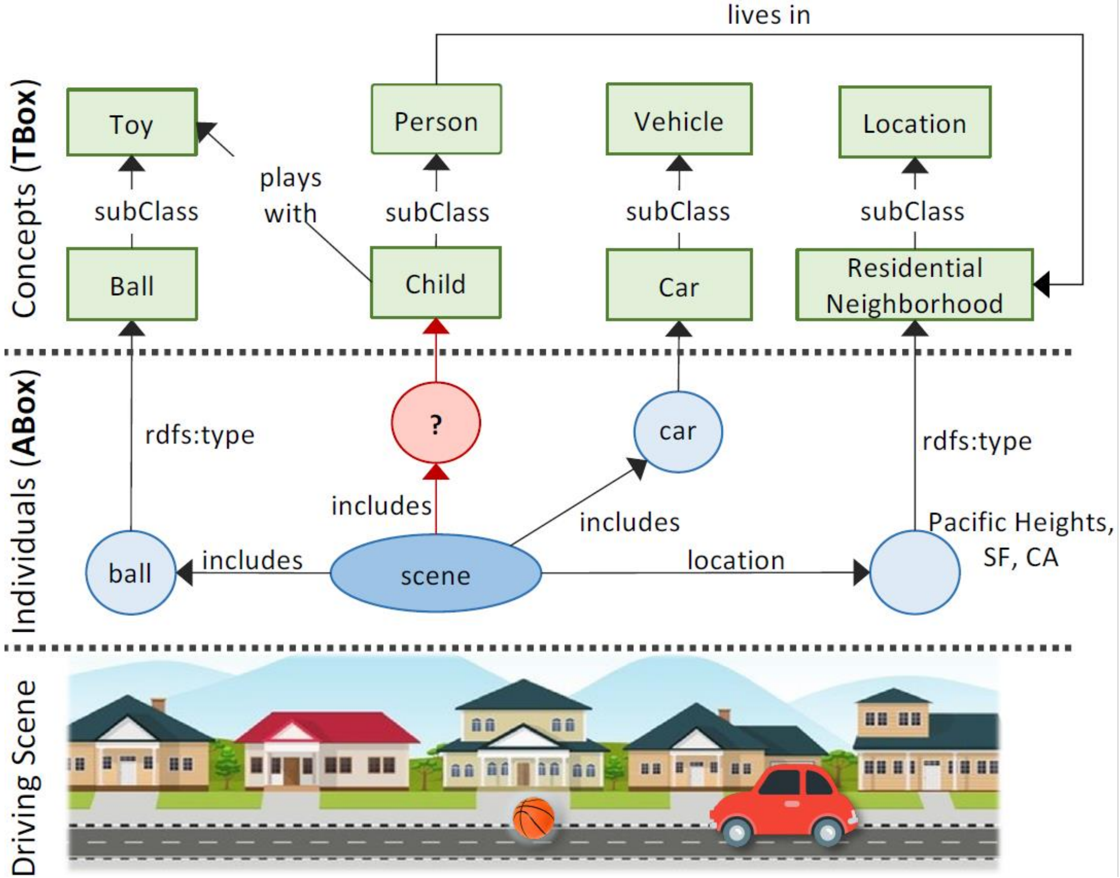

Cory Henson (Bosch Research) | 2021-05-03

Entity prediction w/ knowledge-infused learning

Goal
Predict entities in a scene, given current and background knowledge of the scene

Sample question
What is the probability that a child is nearby?

Solution
Train an entity prediction model with a KG of scenes, applying knowledge-infused learning techniques

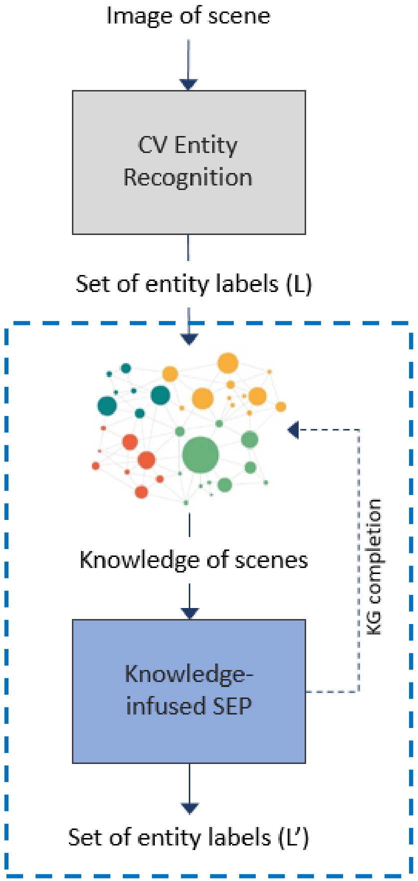

Cory Henson (Bosch Research) | 2021-05-03

Architecture for scene entity prediction (SEP)

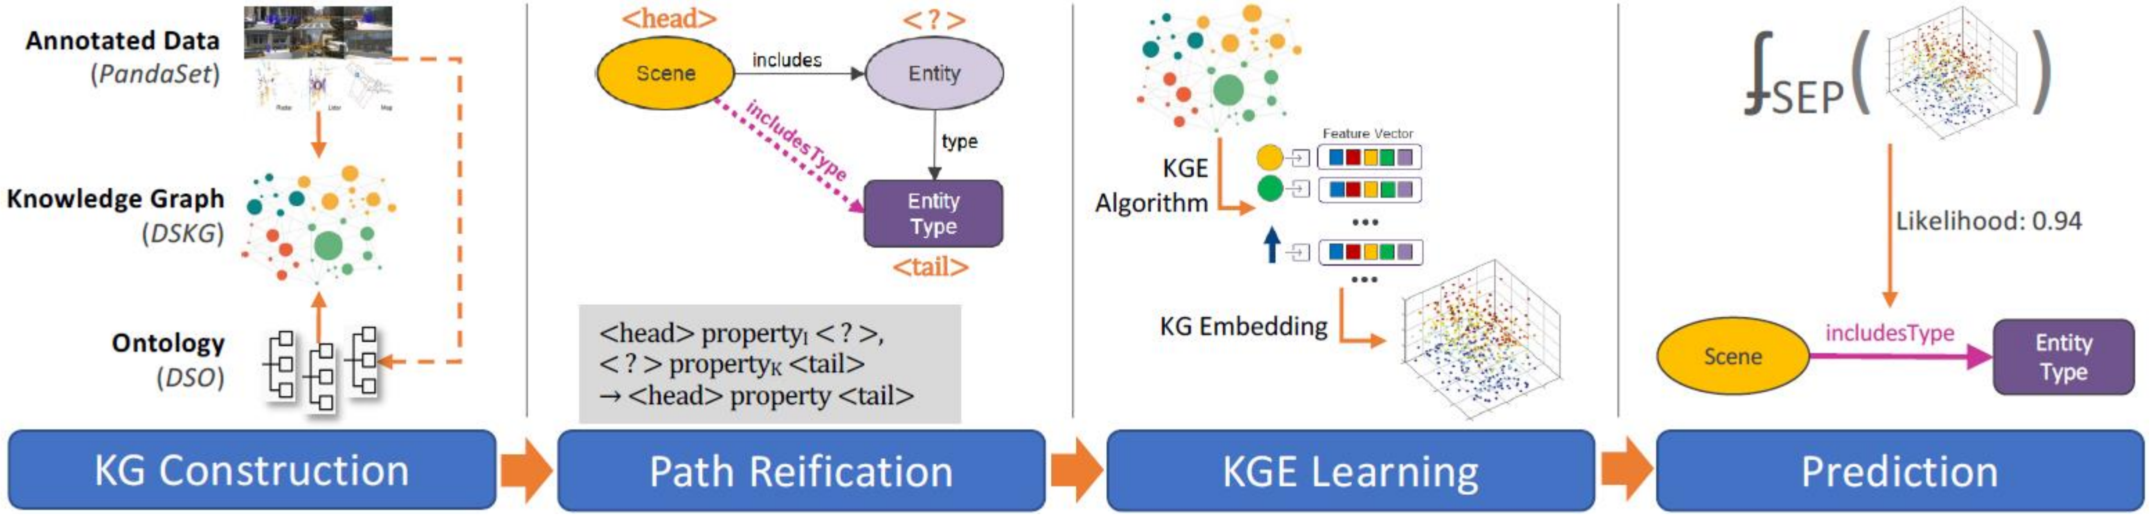

Construct a knowledge graph of driving scenes (DSKG)

Transform KG to a form that is more usable for SEP ($DSKG_R$)

Learn a knowledge graph embedding for the $DSKG_R$

Predict entities in a scene, given a KGE for $DSKG_R$

Knowledge graph transformation

Goal: Map entity prediction to a KG link prediction problem

- Link prediction is a well known problem in KG completion literature
- Link prediction is the primary objective of knowledge graph embeddings (KGEs)

Challenge: SEP is a *path prediction* problem

- For a given scene, predict the ***entity types*** linked through an entity instance
- KG link prediction *DOES NOT* handle path prediction

Solution: Path reification

- Create a direct link between a scene and an entity class whenever the following path is available: Scene-Instance → Entity-Instance → Entity-Class
- The KG with these additional reified links is called DSKG_R
- Use LP to make predictions about this new relation: includesType
- *Path reification rule:*

$$\langle s_i, \mathtt{includes}, e_j\rangle \wedge \langle e_j, \mathtt{rdfs{:}type}, ?\rangle \Rightarrow \langle s_i, \mathbf{\mathtt{includesType}}, ?\rangle$$

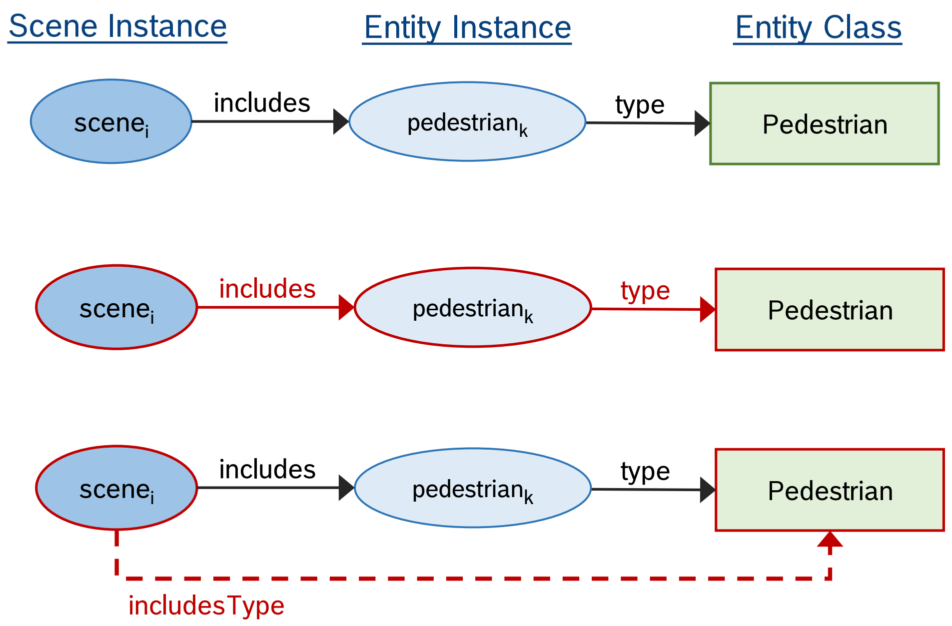

Cory Henson (Bosch Research) | 2021-05-03

BOSCH

Performance of scene entity prediction

Learn KGEs and evaluate performance of SEP

Choose one prototype KGE algorithm from each class

- Geometric: TransE
- Matrix Factorization: HolE
- Deep Learning: ConvKB

KGE algorithm classes

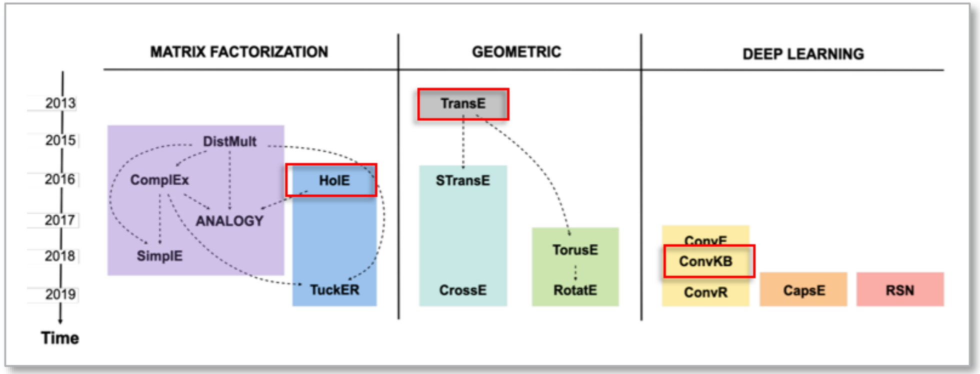

Performance results

Algo.	Entity LP Results Hits@1	Hits@3	Hits@10
TransE	0.1723	0.3620	0.7445
HolE	**0.8988**	**0.9896**	**0.9999**
ConvKB	0.1142	0.3331	0.9004

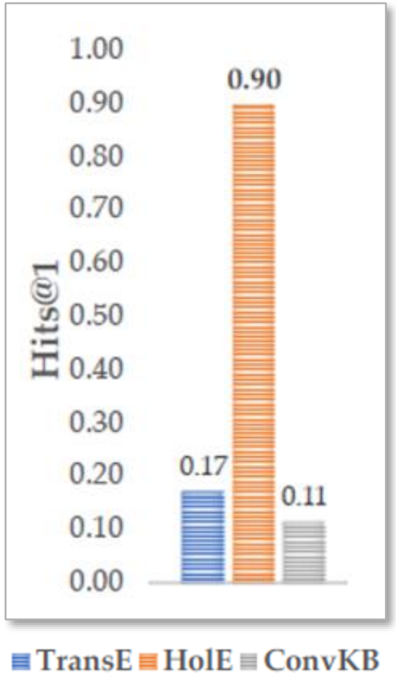

Cory Henson (Bosch Research) | 2021-05-03

BOSCH

Opportunities for knowledge-infused learning in AD

Unification and integration
AD data is multimodal, requiring significant integration for a unified and complete scene understanding; at both the symbolic and sub-symbolic level.

Coherency, consistency and correctness
AD decision making has meaningful consequences and must be of high quality; thus the quality of determinant information is paramount.

Explainability
Trust is essential for adoption of AD technology. The ability to explain ‘*why*’ a decision is made is necessary for enabling such trust.

Ethics, values, accountability and law
Rules of the road are encoded as laws and regulations. AD systems should understand and abide by all rules.

Efficiency
Within an 'open world' environment like driving, many encountered situations have only sparse relevant data. The integration of background knowledge, i.e. common-sense, could fill this gap.

Completion
Most knowledge about a scene is derived through ML and thus probabilistic in nature. KG completion should take this into consideration.

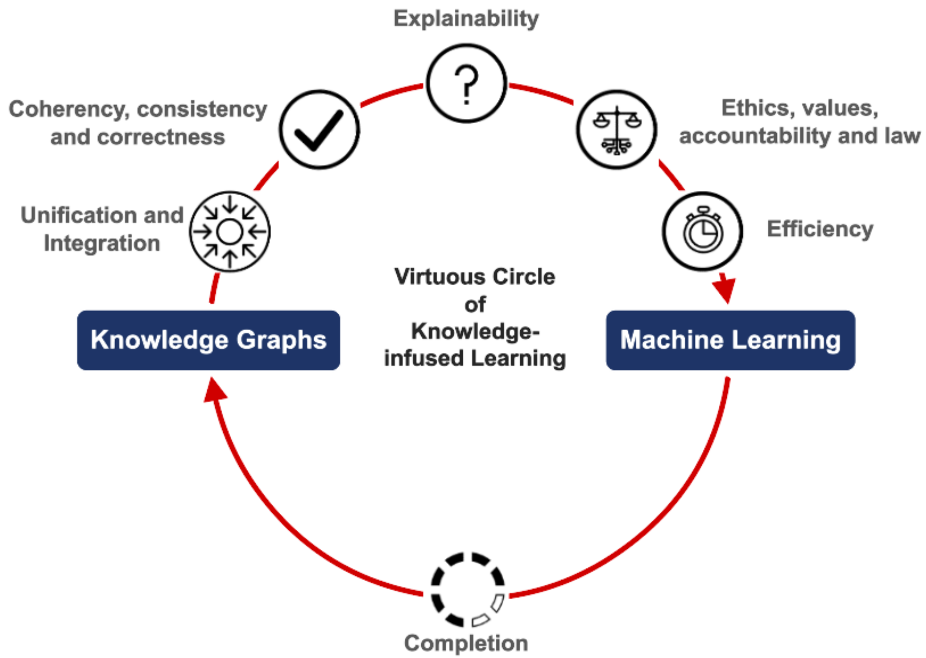


Thank
You

EXTRAS