

Knowledge-based Entity Prediction for Improved Machine Perception in Autonomous Systems

Ruwan Wickramarachchi

Artificial Intelligence Institute
University of South Carolina
Columbia, SC, USA
ruwan@email.sc.edu

Cory Henson

Bosch Research and Technology Center
Pittsburgh, PA, USA
cory.henson@us.bosch.com

Amit Sheth

Artificial Intelligence Institute
University of South Carolina
Columbia, SC, USA
amit@sc.edu

Abstract—Knowledge-based entity prediction (KEP) is a novel task that aims to improve machine perception in autonomous systems. KEP leverages relational knowledge from heterogeneous sources in predicting potentially unrecognized entities. In this paper, we provide a formal definition of KEP as a knowledge completion task. Three potential solutions are then introduced, which employ several machine learning and data mining techniques. Finally, the applicability of KEP is demonstrated on two autonomous systems from different domains; namely, autonomous driving and smart manufacturing. We argue that in complex real-world systems, the use of KEP would significantly improve machine perception while pushing the current technology one step closer to achieving full autonomy.

Index Terms—entity prediction, machine perception, autonomous driving, smart manufacturing, event perception, knowledge-infused learning

I. INTRODUCTION

Autonomous systems are becoming an integrated part of everyday life with applications in various domains, including transportation, manufacturing, and healthcare. The primary goal of autonomous systems is to imbue machines with intelligence that enables them to sense and learn to function within changing environments. Achieving this goal, however, entails solving many hard AI problems. For example, autonomous driving (AD), a popular application area in autonomous systems, often functions as a test-bed for tackling such complex AI problems. Scene understanding is a particularly important challenge in machine perception that involves sense-making at different levels; including visual scene detection, recognition, localization, causal reasoning, etc. [1]. In this paper, we define the problem of knowledge-based entity prediction (KEP), a novel formulation that aims to improve current scene understanding technology by leveraging structured relational knowledge.

KEP Definition

KEP is the task of predicting the inclusion of potentially unrecognized entities in a scene, given the current and background knowledge of the scene represented as a knowledge graph.

There are several reasons why a scene may contain unrecognized entities, including (but not limited to): sensor failure, occlusion, low resolution (e.g., from weather), or errors in computer vision models for

object detection and recognition. As an example, consider this real-world scenario for autonomous driving: An autonomous vehicle is driving through a residential neighborhood and a ball is detected bouncing on the road (see Figure 1). KEP may be used to predict the likelihood of encountering a child chasing after the ball, even though no child has yet been detected.

Making such predictions, however, requires a deep understanding of the various aspects of the scene, including the semantic relations among the detected entities. For example, the knowledge that the bouncing ball is a type of toy and children often play with toys and that children often live and play in residential neighborhoods. Note that this type of semantic background knowledge is out-of-scope for current perception systems where computer vision plays the predominant role in assigning semantic types to the objects detected (i.e. semantic segmentation, object detection, and recognition) [2].



Fig. 1. Knowledge-based Entity Prediction Problem in AD

Hypothesis

We hypothesize that potentially unrecognized entities in a scene may be detected through (1) the use of a knowledge graph to provide an expressive, holistic representation of scene knowledge, and (2) the application of knowledge-infused learning for knowledge completion; i.e. predicting missing knowledge in the graph.

Contributions of this paper include:

- 1) Introducing the novel proposition of using KEP for improving machine perception in autonomous systems.
- 2) Defining and formalizing KEP as a knowledge completion problem.

The rest of the paper is structured as follows. Section II discusses the proposition of using KEP to improve the machine perception in autonomous systems, along with its connection to perception in Psychology. Section III formalizes KEP as a knowledge completion problem and introduces three potential solutions. Section IV presents two case studies for using KEP in different application domains, autonomous driving and smart manufacturing, while Section V presents experimental results of a KEP implementation on AD data. Finally, we wrap up in Section VI with conclusions and future work.

II. IMPROVED MACHINE PERCEPTION THROUGH KEP

The stated hypothesis allows us to view scene understanding in autonomous systems from a new perspective. For example, consider the following case from the autonomous driving domain. The perception module, one of the three main functional components in an AD system, uses the raw data from sensors to perform two critical tasks: (i) environmental perception – obtaining percepts (i.e. detecting and recognizing high-level entities) in the environment (P_{CV}), and (ii) localization – determining the relative position of the ego vehicle and other entities in the environment. The outputs of the perception module are then used by subsequent processes (see Figure 2(a)) such as planning and control. The planning module performs tasks such as behavioral and motion planning of the ego vehicle while the control module governs the actuators that change the vehicle state and trajectory. This current process, however, has a serious vulnerability. If there is a mistake in the perception process, an unrecognized entity for example, this mistake will propagate through the subsequent processes and change the desired outcome. In the case of autonomous driving, this could lead to a fatal accident.

Proposition. With this background, consider the following proposition for improved machine perception. Environmental percepts of scenes (recognized by computer vision algorithms) (P_{CV}) are now represented within a KG. The KEP process then predicts a new set of percepts (P_{KEP}) for the potentially unrecognized entities – i.e. unrecognized objects and events. These percepts are used to enrich the scene KG by adding new nodes and relations. Note that, as shown in Figure 2(b), this is a *cyclical process* in which the scene KG is updated with new facts continuously over time. From a scene understanding perspective, the perception module’s understanding of a scene improves as the KG becomes more complete. As a result of this process, (i) the perception module is enriched with potentially missing entities and (ii) the subsequent processes (i.e. planning and control) will have the benefit of using a richer set of percepts about the environment $P_{CV} + P_{KEP}$ along with other contextual and background cues from the relational knowledge in the scene

KG. For example, consider the case where the perception module detects a pedestrian (P_{CV}) on the road. It does not, however, recognize that the pedestrian is *jaywalking*. Even if no *jaywalking* events have been seen while training the CV perception module, representing knowledge of this event – i.e. (*Pedestrian, participatesIn, Jaywalking*) – in the scene KG could provide a new insight or cue for handling this *edge-case* with KEP (i.e. P_{KEP}). We believe KEP has the potential to improve the machine perception in autonomous systems by significantly reducing the omission and misclassification errors due to sensor failures, occlusions, false negatives in CV models, etc. in the current perception modules. It should be noted that this proposition does not demand radical change to any of the current CV-based perception systems. Rather, KEP would complement the CV models by having the ability to leverage a machine-readable, unified knowledge representation in scene KG that is central to the KEP process.

Connection to Perception in Psychology

Event perception is a particularly interesting area of research in Psychology where it is considered that “the perception of events depends on both sensory cues and knowledge structures that represent previously learned information about event parts and inferences about actors’ goals and plans” [3]. The event segmentation theory (EST), a theory on the mechanism that segments ongoing activities into meaningful events, starts with a *perceptual processing stream* that takes a representation of the current state of the world (e.g., perceptual information, language, external knowledge, etc.) and **produces a set of predictions** about “what will happen a short time in the future” [4]. This can be seen as an *ongoing* comprehension process where the predictive processing is modulated by an (event) model that considers perceptual streams together with a representation of the world.

Corollary to KEP: At a higher level, KEP functions on the same key fundamentals on which EST is built. First, they both start with a given set of percepts (or perceptual stream) as the input. Then, they rely on a rich representation of the world that is comprised of higher-level concepts and relations, similar to scene KG used by the KEP process. Given the perceptual stream and the representation of the world, EST predicts *events* that will happen in a short time while KEP predicts both *objects and events* that could be missing from the scene or may be seen within a short time. Viewed in this manner – i.e. as an ongoing comprehension process of event perception – KEP can be seen as an implementation of EST for improving machine perception in autonomous systems.

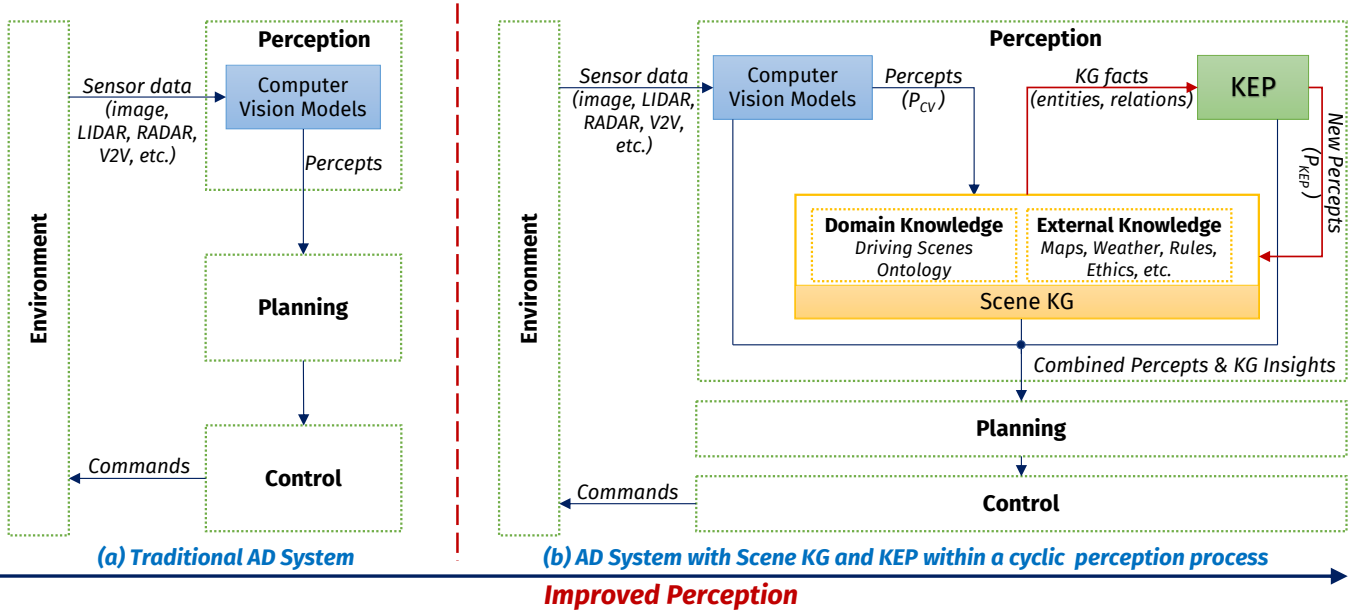


Fig. 2. Improved machine perception through KEP in Autonomous Driving

A. Other Benefits of KEP

In addition to improving machine perception, the proposed KEP approach and the use of scene KG will have several other benefits. We briefly discuss them here.

1) **External Knowledge Integration:** Currently, all processes from perception to control primarily rely on large, diverse, and high-quality training data to function robustly. Despite the availability of a variety of useful knowledge sources, it remains a challenge to leverage them within current deep learning and computer vision methods. Interestingly, for KEP, the scene knowledge is represented in a KG and conformant to an ontology. This allows for easy integration of additional information, including geo-spatial data (e.g., Open Street Maps (OSM)), commonsense knowledge (e.g., ConceptNet, CSKG [5]), driving rules and law (e.g., from driving manuals), etc. With this information represented as a KG, many existing knowledge-infused learning [6]–[8] algorithms – e.g., Knowledge Graph Embeddings (KGEs) [9], Graph Neural Networks (GNNs) [10] – can learn latent representations of this knowledge to be used with several downstream applications. Therefore, a scene KG continually evolving with KEP would create a rich internal representation of the external world (i.e. “world model”) for autonomous systems. Recent work in the autonomous driving domain has shed some light into the development of KGs from real AD data – e.g. the Driving Scene Knowledge Graph (DSKG) and enhancing it with geo-spatial data [11] and commonsense knowledge [12].

2) **Interpretability and Explainability:** Another benefit of KEP is the ability to use relational and contextual knowl-

edge to generate explanations for the predictions. Consider the following example: Similar to the way a flight recorder (i.e. a black box) is used to trace the steps in-flight crash investigation, AD accident investigations look at replays of sensor readings, detected entities, and decisions made for each time point. Now, due to having semantic concepts, relations defined between entities, and raw sensor data mapped to their corresponding semantic types in the scene KG, it will allow investigators to interpret the raw sensor reading. The sensor data at the time of the crash can be easily interpreted as (i) *high-speed* by comparing the recorded speed with the legal speed limit of the location, (ii) the applicable traffic rules can be obtained by first decoding the latitude/longitude readings as the higher-level concept “*School Zone*” and then querying the rules applicable to school zones. Through such rich and contextual interpretations, we believe KEP could significantly improve the domain explainability of autonomous systems.

III. THE FORMALIZATION AND SOLUTIONS FOR KEP

Next, we formalize KEP task as a generic knowledge completion problem. Then we devise possible solutions for KEP considering different machine learning and data mining techniques.

A. Formalization of KEP

First, we define the notation for a formalization of KEP that would be generic enough to allow for a wide range of possible solutions. Let the KG $G : N \times R \times N$ where N, R denote nodes and relations, respectively. Considering the different types of nodes in G , let $s^{(i)} \in S \subseteq N$ (where S represents the set of scene nodes), $e_j \in E \subseteq N$ (where E represents set of entity type nodes), and $\{E_{obs}^{(i)}, E_{inf}^{(i)}\} \subseteq E \subseteq N$ (where $E_{obs}^{(i)}$ represents the subset of E that are observed and included in



Fig. 3. Three possible solutions devised for KEP based on different machine learning/data mining techniques

$s^{(i)}$, and $E_{inf}^{(i)}$ as the subset of E to be inferred). Scenes and entities are connected in G through $(s^{(i)}, e_j) \in I \subseteq R$ where I denotes the set of `includes` relations. Entities are connected with their types in G through $T \subseteq R$ where T denotes the set of `type` relations – e.g., `rdf:type`. Now, for a given scene $s^{(i)}$, the objective of KEP is to find $E_{inf}^{(i)}$ such that $E_{inf}^{(i)} \leftarrow \phi(G, s^{(i)}, E_{obs}^{(i)})$ where ϕ is the learned inference model generated and used by the solution.

B. Possible Solutions for KEP

1) **Link Prediction:** In the KG literature, Link Prediction (LP) is commonly used to address the KG incompleteness issue. While there are many techniques for LP, Knowledge Graph Embedding (KGE) methods, which represent KG nodes and relations in low-dimensional latent space, achieve the state-of-the-art performance for this task. Specifically, given an incomplete triple $(h, r, ?)$, KGE-based LP methods attempt to predict the missing component (?) by querying the KGE model $\phi(h, r, ?)$ to obtain the tail (t) of this triple. Given a KG with scenes and related entities, it begs the question of whether KEP can be formulated as a LP problem. In a scene KG, however, if $h = s^{(i)}$, predicting t would mean predicting an entity instance (i.e. a specific entity) rather than an entity class (i.e. an entity type). One potential solution would be to reify the path from $s^{(i)}$ to the associated class of an entity included in the scene (e.g., `Car` included in a scene) [11]. This path could be reified using a direct `includesType` relation as shown in Eq. 1. Note that this would allow the reuse of any existing KGE technique for predicting the tail of an `includesType` relation, thus predicting the entity type included in a scene.

$$\langle s^{(i)}, \text{includes}, e_j \rangle \wedge \langle e_j, \text{rdfs:type}, ? \rangle \Rightarrow \langle s^{(i)}, \text{includesType}, ? \rangle$$

2) **Association Rule Mining:** Association rule mining (ARM) is a popular data mining technique that is widely used to uncover interesting correlations, frequent patterns, and associations among sets of items. For example, in retail market basket analysis, ARM could be used to find associations between items customers frequently buy together. The objective of ARM is to generate *association rules* in the

form $r_i : \{A, B\} \Rightarrow \{C\} | c$ where $\{A, B\}$ (i.e. antecedents) imply the co-occurrence of $\{C\}$ (i.e. consequent) (here, $\{A, B\}, \{C\} \subset I$, referred to as *itemsets*) with a confidence factor of $c; 0 \leq c \leq 1$. This would mean, at the minimum, $c\%$ transactions in the set of transactions T satisfies the r_i rule. Considering the objective of KEP, we may postulate whether the co-occurrences of entities could help predict the missing set of entities in a scene. If so, then KEP could be formulated as a market basket analysis problem, where the sets of items represent the set of observed entities in scenes. As shown in [11], the ARM approach for KEP consists of 3 main steps. First, an ARM algorithm (e.g., Apriori algorithm [13]) is run on the training dataset (i.e. a scene KG) to generate a set of association rules. Second, a mask is created to filter the relevant set of rules for a given scene. A rule is considered relevant if its antecedent is a subset of the observed entities of the scene (i.e. $E_{obs}^{(i)}$). Finally, the missing set of entities of this scene (i.e. $E_{inf}^{(i)}$) is obtained by considering the aggregation of consequents that satisfy the mask created above.

3) **Collective Classification:** Collective classification is a combinatorial optimization problem that works on the premise that knowledge of the correct label for one node in a graph (i.e. semantic type in scene KG) improves the ability to assign correct labels to the other unlabeled nodes it connects to. In its most generic form, the goal of collective classification is to jointly determine the correct label assignments of all the nodes in a graph/network. Let $V = v_1, v_2, \dots, v_n$ be the set of nodes in the graph, and let N be the neighborhood function ($N \subseteq V$) that describes the underlying network structure [14]. Each node in V is a random variable that can take a value from an appropriate domain, $L, L = l_1, l_2, \dots, l_m$. Considering the node label space, V is divided into two sets: X , the nodes for which the correct label assignments are known (observed variables), and Y , the nodes whose values need to be inferred (unobserved variables). The inference task is to then label the nodes $y_j \in Y$ with a subset of predefined labels in L . When considering KEP, nodes $X^{(i)}$ would denote the set of observed entities (i.e. $E_{obs}^{(i)}$) for scene $s^{(i)}$, while $y_j \in Y$ would denote the additional set of entities to be predicted/inferred (i.e. $E_{inf}^{(i)}$) where $E = L$ in the collective classification formalization.

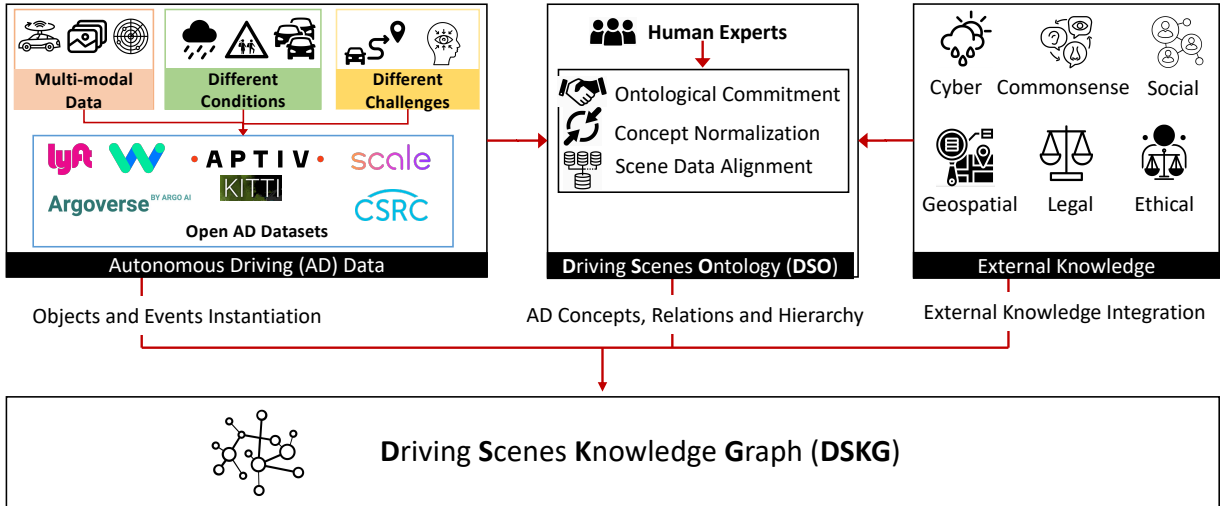


Fig. 4. The process of creating a unified knowledge representation for knowledge-based downstream applications in AD domain

IV. CASE STUDIES

In this section, we demonstrate the applicability of KEP in two application domains of autonomous systems: autonomous driving and smart manufacturing. We will also highlight the different types of knowledge KEP exploits in each use case.

A. Autonomous Driving

Autonomous driving is a key technical domain of interest in AI with the objective of achieving full (level-5) autonomy where the car drives itself without any human involvement. Key obstacles for achieving full autonomy center around the incompleteness in learning, dealing with edge-cases, and challenges associated with open-world driving environments. Currently, the computer vision based models solely rely on training from observational data, and there is no clear way to *infuse* the innate understanding or knowledge that humans have about the world, including knowledge of different aspects of a scene, such as time, space, causality, semantics, etc. We believe one possible way of tackling this problem could be to change the way the scenes are represented in AD. We argue that AD requires a more holistic representation of driving scenes composed of higher-level semantic concepts and relations, that evolves over time, and contains high-quality multi-modal data from heterogeneous sources. Knowledge Graphs (KGs) naturally fit these requirements as they are capable of representing expressive and meaningful relations among entities (i.e. objects, events, and abstract concepts) in the real world.

Figure 4 shows a schematic overview of how the scene KG is developed as a unified representation for AD knowledge. KEP, through a *cyclic process*, updates this knowledge over time by predicting the unseen entities leveraging the underlying relational structure and diverse knowledge in the scene KG. We believe this would significantly improve the machine perception in AD (i.e. by reducing misclassification errors, omissions, object occlusions, etc.). Additionally, the latent

transformations of the scene KG, such as knowledge graph embeddings [15], would allow the subsequent processes in AD to benefit from the relational knowledge in the scene KG. This could push the current technology one step closer to achieving full autonomy.

B. Smart Manufacturing

Smart manufacturing (or Industry 4.0/I4.0) [16] employs smart robots that function fully autonomously for a defined task. Precise understanding of manufacturing events plays a critical role in such systems. It requires analyzing an array of sensor feeds (e.g., video streams from various cameras, pressure, thermal, and humidity sensors, etc.) as well as understanding the “manufacturing context” through the use of other sensor data, external knowledge and information relevant to the production line. Similar to the way driving scenes are represented above for autonomous driving, we argue the manufacturing process can be seen as a stream of “manufacturing scenes”. In this case, a KG can be developed to represent manufacturing scenes that are composed of entities (i.e. objects and events) along with their spatial and temporal attributes and relevant meta-data. In addition to scene data, the KG allows for the seamless integration of existing knowledge sources such as ontologies for I4.0 (e.g., Standards Ontology, Resource Reconfiguration Ontology, Manufacturing Resource Capability Ontology (MaRCO), I4.0 Components Ontology, etc.) [17] as well as relevant domain knowledge (e.g., data from robot product specifications, safe working conditions, etc.). In this context, the KEP task could be used for improving event understanding by predicting the possible events of failures and/or outliers in components (e.g., prognostic component) of smart manufacturing systems.

V. EXPERIMENTAL RESULTS

To provide a synopsis of how KEP performs on real-world autonomous systems, we present experimental results of an

TABLE I
 KEP RESULTS OF PANDASET (DSKG- P_R) [11] ON 3 ALGORITHMS, EACH EXPERIMENT AVERAGED WITH STANDARD DEVIATION ACROSS 5 RUNS.
 EVALUATION METRICS: MRR=MEAN RECIPROCAL RANK, H@K= HITS@K, ACCU. = KEP ACCURACY, MICRO/MACRO F1 =
 MICRO/MACRO-AVERAGED-F1-SCORE

	Ranking Metrics				KEP Performance Metrics		
	MRR	H@1	H@3	H@10	Accu. (%)	Micro F1	Macro F1
TransE [18]	0.32 ±0.03	0.16 ±0.05	0.35±0.04	0.71±0.03	22.98 ±4.33	0.26 ±0.04	0.20 ±0.02
HolE [19]	0.93 ±0.00	0.87±0.01	0.98 ±0.00	1.00±0.00	88.91±0.64	0.90 ±0.01	0.87 ±0.00
ConvKB [20]	0.29±0.01	0.11 ±0.02	0.31 ±0.02	0.86±0.02	17.83±1.99	0.22±0.02	0.17 ±0.02

LP-based implementation of KEP considering autonomous driving data. For this purpose, we refer to the Driving Scenes Knowledge Graph (DSKG) introduced by Wickramarachchi et al. [11] which represents the scene data from the Pandaset [21] dataset and is conformant to the Driving Scenes Ontology (DSO). Table I summarizes the KEP results with three KGE algorithms and two sets of evaluation metrics – i.e. ranking metrics (hits@K, Mean Reciprocal Rank (MRR)) and KEP performance metrics (Accuracy, Micro/Macro F1). This evaluation supports the hypothesis for KEP, predicting missing entities with a 0.87 precision (hits@1) and 88.91% accuracy. A more detailed and thorough evaluation with other AD datasets, different KG structures, and other implementations can be found in [11]. When considering the space and time complexities of this solution, it is most efficient when TransE is used (Time: $\mathcal{O}(n_t d)$, Space: $\mathcal{O}(md + nd)$) and least efficient when ConvKB is used (Space: $\mathcal{O}(md + nd + (\tau + 3)d)$, Time: $\mathcal{O}(n_t \tau d)$). It should be noted, however, the space/time complexities of KEP are dependent on the complexity of the underlying solution and its sub-components.

VI. CONCLUSIONS AND FUTURE WORK

This paper introduces the novel task of knowledge-based entity prediction (KEP) that leverages knowledge from heterogeneous sources for inferring and/or predicting unrecognized entities. A new proposition has been presented for using KEP within a *cyclic perception process* for improving machine perception in autonomous systems. We define and formalize KEP as a knowledge completion problem and present three possible solutions considering several data mining and machine learning techniques. Finally, we demonstrate the applicability of KEP for two use-cases of autonomous systems from different domains. In the near future, we plan to implement and evaluate the solutions discussed in this paper with real autonomous systems data. More generally, however, we hope the proposed reformulation of KEP may inspire other researchers to devise, apply, and share new and novel solutions to the KEP problem.

ACKNOWLEDGEMENTS

We thank Dr. Valerie Shalin for the helpful suggestions and discussion on perception in Psychology. This work was supported in part by Bosch Research, NSF grants #2133842, and #2119654. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding organizations. We thank the anonymous reviewers for their

constructive comments.

REFERENCES

- [1] V. Ramanishka, Y.-T. Chen, T. Misu, and K. Saenko, "Toward driving scene understanding: A dataset for learning driver behavior and causal reasoning," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 7699–7707.
- [2] S. Grigorescu, B. Trasnea, T. Cocias, and G. Macesanu, "A survey of deep learning techniques for autonomous driving," *Journal of Field Robotics*, vol. 37, no. 3, pp. 362–386, 2020.
- [3] J. M. Zacks, N. K. Speer, K. M. Swallow, T. S. Braver, and J. R. Reynolds, "Event perception: a mind-brain perspective," *Psychological bulletin*, vol. 133, no. 2, p. 273, 2007.
- [4] J. M. Zacks, "Event perception and memory," *Annual Review of Psychology*, vol. 71, pp. 165–191, 2020.
- [5] F. Ilievski, P. Szekely, and B. Zhang, "Cskg: The commonsense knowledge graph," in *The Semantic Web*, R. Verborgh, K. Hose, H. Paulheim, P.-A. Champin, M. Maleshkova, O. Corcho, P. Ristoski, and M. Alam, Eds. Cham: Springer International Publishing, 2021, pp. 680–696.
- [6] A. Sheth, M. Gaur, U. Kursuncu, and R. Wickramarachchi, "Shades of knowledge-infused learning for enhancing deep learning," *IEEE Internet Computing*, vol. 23, no. 6, pp. 54–63, Nov 2019.
- [7] L. G. Valiant, "Knowledge infusion," 2006.
- [8] A. d. Garcez and L. C. Lamb, "Neurosymbolic ai: The 3rd wave," *arXiv preprint arXiv:2012.05876*, 2020.
- [9] A. Rossi, D. Barbosa, D. Firmani et al., "Knowledge graph embedding for link prediction: A comparative analysis," *ACM Trans. Knowl. Discov. Data*, vol. 15, no. 2, Jan. 2021.
- [10] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and S. Y. Philip, "A comprehensive survey on graph neural networks," *IEEE transactions on neural networks and learning systems*, vol. 32, no. 1, pp. 4–24, 2020.
- [11] R. Wickramarachchi, C. Henson, and A. Sheth, "Knowledge-infused Learning for Entity Prediction in Driving Scenes," *Frontiers in Big Data*, vol. 4, p. 759110, 2021.
- [12] S. N. Chowdhury, R. Wickramarachchi, M. H. Gad-Elrab, D. Stepanova, and C. Henson, "Towards leveraging commonsense knowledge for autonomous driving," in *The 20th International Semantic Web Conference (ISWC)*, 2021.
- [13] R. Agrawal, H. Mannila, R. Srikant, H. Toivonen, and A. I. Verkamo, "Fast discovery of association rules," in *Advances in Knowledge Discovery and Data Mining*, U. M. Fayyad, G. Piatetsky-Shapiro, P. Smyth, and R. Uthurusamy, Eds. AAAI/MIT Press, 1996, pp. 307–328.
- [14] P. Sen, G. Namata, M. Bilgic, and L. Getoor, *Collective Classification*. Boston, MA: Springer US, 2010, pp. 189–193.
- [15] R. Wickramarachchi, C. Henson, and A. Sheth, "An evaluation of knowledge graph embeddings for autonomous driving data: Experience and practice," *AAAI 2020 Spring Symposium on Combining Machine Learning and Knowledge Engineering in Practice (AAAI-MAKE 2020)*, 2020.
- [16] B. Wang, F. Tao, X. Fang, C. Liu, Y. Liu, and T. Freiheit, "Smart manufacturing and intelligent manufacturing: A comparative review," *Engineering*, vol. 7, no. 6, pp. 738–757, 2021.
- [17] M. Yahya, J. G. Breslin, and M. I. Ali, "Semantic web and knowledge graphs for industry 4.0," *Applied Sciences (Switzerland)*, vol. 11, no. 11, pp. 1–23, 2021.
- [18] A. Bordes, N. Usunier, A. Garcia et al., "Translating embeddings for modeling multi-relational data," in *Adv. in Neural Information Processing Systems*, 2013, pp. 2787–2795.

- [19] M. Nickel, L. Rosasco, and T. Poggio, "Holographic embeddings of knowledge graphs," in *Thirtieth AAAI Conference on Artificial Intelligence*, 2016.
- [20] D. Q. Nguyen, T. D. Nguyen, D. Q. Nguyen *et al.*, "A novel embedding model for knowledge base completion based on convolutional neural network," in *Proceedings of NAACL-HLT*, 2018, pp. 327–333.
- [21] Scale AI, "Pandaset open datasets - scale," <https://scale.com/open-datasets/pandaset>, 2020, (Accessed on 01/25/2021).

AUTHORS

Ruwan Wickramarachchi: Ruwan is a Ph.D. student at the AI Institute, University of South Carolina. His dissertation research focuses on introducing expressive knowledge representation and knowledge-infused learning techniques to improve machine perception and context understanding in autonomous systems (e.g., autonomous driving and smart manufacturing). Contact him at: ruwan@email.sc.edu

Cory Henson: Cory is a lead research scientist at the Bosch Research and Technology Center in Pittsburgh, PA. His research focuses on the application of knowledge representation and knowledge-infused learning to enable autonomous driving. He also holds an Adjunct Faculty position at Wright State University. Prior to joining Bosch, he earned a Ph.D. in Computer Science from WSU, where he worked at the Kno.e.sis Center applying semantic technologies to represent and manage sensor data on the Web. Contact him at: cory.henson@us.bosch.com

Amit Sheth: Amit is the founding director of the AI Institute, University of South Carolina. His current core research is in knowledge-infused learning and explanation, and transnational research includes personalized and public health, social good, education, and future manufacturing. He is a fellow of IEEE, AAAI, AAAS, and ACM. Contact him at: amit@sc.edu (<http://aiisc.ai/amit>)