

Ontology-guided Semantic Composition for Zero-Shot Learning

Jiaoyan Chen¹, Freddy Lécué^{2,3}, Yuxia Geng⁴, Jeff Z. Pan^{5,6}, Huajun Chen⁴

¹Department of Computer Science, University of Oxford, UK

²INRIA, France

³CortAix@Thales, Canada

⁴College of Computer Science, Zhejiang University, China

⁵School of Informatics, The University of Edinburgh, UK

⁶Department of Computer Science, The University of Aberdeen, UK

Abstract

Zero-shot learning (ZSL) is a popular research problem that aims at predicting for those classes that have never appeared in the training stage by utilizing the inter-class relationship with some side information. In this study, we propose to model the compositional and expressive semantics of class labels by an OWL (Web Ontology Language) ontology, and further develop a new ZSL framework with ontology embedding. The effectiveness has been verified by some primary experiments on animal image classification and visual question answering.

1 Introduction

Machine learning often relies on a large data set, especially for training deep models. However, training and prediction with limited data (a.k.a low-resource learning) is common in applications such as labeling images with new concepts (Nikolaus et al. 2019), learning embeddings of new relations in a knowledge graph (Qin et al. 2020), extracting long-tail information from text (Wu, Hoffmann, and Weld 2008), and Visual Question Answering (VQA) (Teney and van den Hengel 2018) with outside knowledge (Marino et al. 2019). In supervised classification, prediction with new classes that have never appeared in the training data is often known as *zero-shot learning* (ZSL) (Palatucci et al. 2009). Those new classes are called *unseen classes*, while those classes that have labeled training samples are called *seen classes*.

One popular solution for ZSL is building the classifier of an unseen class by combining the classifiers (or learned model parameters) of seen classes with some side information, including (i) text (e.g., (Qin et al. 2020) learns the embedding of zero-shot relations via relation text description, (Lei Ba et al. 2015) predicts images of unseen classes via class text description), (ii) human annotations (e.g., (Lampert, Nickisch, and Harmeling 2009) classifies images of new classes via their annotations of visual characteristics), and (iii) graph data (e.g., (Kampffmeyer et al. 2019) aligns image classes with WordNet nodes and predicts for unseen classes by Graph Neural Network (GNN) while (Geng et al. 2020b) further extends the idea with explanation). However, the semantic expressiveness of such side information is often too weak to represent complex inter-class relationships, thus limiting the performance.

Inspired by compositional learning, where learned components are combined by symbolic operations for a new component (e.g., the language model by (Socher et al. 2012) represents a phrase or a sentence by recursively combining the vectors of its words via its parse tree), we propose to formally represent the semantics of each class label in ZSL by OWL 2 (Web Ontology Language) class definitions. With the ontology we first embed its logic axioms, textual information and paths, and encode classes into vectors (i.e., semantic encodings). We then follow a mapping-based ZSL paradigm which learns a mapping function from the input to the semantic encoding via the training samples, and predicts by comparing the input’s semantic encoding with the semantic encodings of class labels. For example, given an image of a whale, we map it to a semantic encoding and compare it to those semantic encodings of class labels (Blue Whale, Humpback Whale and Killer Whale) – the nearest neighbour is adopted as the image’s class label.

This paper introduces a new ontology-guided neural-symbolic design pattern for the sample shortage problem in machine learning. Specially it develops a framework for using OWL ontology as an expressive side information for improving ZSL. Although this study is in an early stage, some promising results have been achieved in our evaluation with animal image classification and visual question answering.

2 Background

2.1 OWL Ontology

In this paper, we adopt OWL 2 EL (\mathcal{EL}^{++} Description Logic) ontologies (Baader, Brandt, and Lutz 2005). A signature Σ , noted $(\mathcal{CN}, \mathcal{RN}, \mathcal{IN})$, consists of 3 disjoint sets of (i) atomic concepts \mathcal{CN} , (ii) atomic relations \mathcal{RN} , and (iii) individuals \mathcal{IN} . Given a signature, the top concept \top , the bottom concept \perp , an atomic concept A , an individual a , an atomic relation r , \mathcal{EL}^{++} concept expressions C and D can be composed with the following constructs:

$$\top \mid \perp \mid A \mid C \sqcap D \mid \exists r.C \mid \{a\}$$

An OWL 2 EL ontology consists of a TBox and an ABox, where the TBox is a finite set of General Concept Inclusion axioms (e.g. $C \sqsubseteq D$), relation inclusion and composition axioms (e.g., $r \sqsubseteq s$, $r_1 \circ \dots \circ r_n \sqsubseteq s$), etc. An ABox is a set of concept assertion axioms, relation assertion axioms, individual in/equality axioms, etc.

2.2 The Problem of ZSL

In machine learning classification, a classifier is trained to approximate a target function $f : x \rightarrow y$, where x denotes the input data, y denotes the output i.e., class label. ZSL is to classify data whose possible class labels have been omitted from the training data (Palatucci et al. 2009; Xian et al. 2018). Formally, we denote (i) the samples for training as $\mathcal{D}_{tr} = \{(x, y) | x \in \mathcal{X}_s, y \in \mathcal{Y}_s\}$, where \mathcal{X}_s and \mathcal{Y}_s represent the training data and seen class labels, respectively; (ii) the samples for testing (prediction) as $\mathcal{D}_{te} = \{(x, y) | x \in \mathcal{X}_u, y \in \mathcal{Y}_u\}$ ¹, where \mathcal{X}_u and \mathcal{Y}_u represent the testing data and unseen class labels, respectively, with $\mathcal{Y}_u \cap \mathcal{Y}_s = \emptyset$; and (iii) the side information as a function $h : y \rightarrow z$, where z represents the semantic encoding of class label y . Our ZSL problem is to predict the class labels of samples in \mathcal{X}_u as correctly as possible.

3 Methodology

As shown in Figure 1, our ZSL framework includes (i) an OWL ontology for representing complex concepts, where each class label corresponds to a concept, (ii) ontology embedding which encodes concepts into vectors with semantics of logic axioms, text and paths, and (iii) a mapping-based ZSL method. We mainly describe the core component of the ZSL framework such as the semantic encoding, while the initial architecture for feature learning does need to be changed according to the task.

3.1 Ontology for Semantic Composition

We build an OWL 2 EL ontology (Baader, Brandt, and Lutz 2005) to define the corresponding complex concepts for class labels. Take the animal image classification as an example, the subsumption can model the taxonomic relations, e.g., *Killer_Whale* \sqsubseteq *Toothed_Whale*. The existential quantification can define the visual characteristics of animals, e.g., $\exists hasTexture.Patches$ and $\exists hasHabitat.Ocean$. Concept definitions can then be used to define composed classes, such as *Killer_Whale* \doteq *Toothed_Whale* \sqcap $\exists hasTexture.Patches$ $\sqcap \dots$. We choose OWL 2 EL due to (i) its support of the existential quantification, (ii) polynomial time complexity for entailment reasoning, and (iii) the availability of ontology embedding method for OWL 2 EL, rather than expressive languages like f-SWRL (Pan et al. 2005).

In addition to logical axioms, textual information of concepts in ontologies are useful for ZSL too (Qin et al. 2020; Xian et al. 2018; 2016). It often contains different but complementary semantics as the aforementioned logical axioms. OWL 2 EL ontologies allow the use of rich annotations, including the names, synonyms and descriptions of concepts and relations by e.g., *rdfs:label* and *rdfs:comment*.

Given a specific ZSL task, the supporting ontology can be either edited by domain experts, or (more often) semi-automatically created and curated with expert knowledge, third party data and ontologies. In the above examples:

¹The label set of \mathcal{D}_{te} is often set to $\mathcal{Y}_u \cup \mathcal{Y}_s$, a.k.a generalized ZSL. Our method can deal with both ZSL and generalized ZSL.

the taxonomic relations can come from WordNet (Miller 1995); the textual information can be extracted from DBpedia (Auer et al. 2007); the existential quantification can be created by domain experts or via reusing existing ontologies.

3.2 Ontology Embedding

With a given ontology, we implement the label encoding function $h : y \rightarrow z$ by (i) embedding the logical axioms according to their geometric construction (Kulmanov et al. 2019), and (ii) fine-tuning the word vector model by textual information and graph paths.

Embedding of Logical Axioms Ontology embedding is to approximate the interpretation of a given ontology by mapping logical axioms into a geometric space, based on which numeric losses are calculated and vector representations of concepts and relations are learned by optimization.

Specially, a concept C is modeled as a ball, namely a vector denoted as $\nu(C) \in \mathbb{R}^n$, where n is the embedding dimension that can be configured, and a radius denoted as $\gamma(C) \in \mathbb{R}$. A relation r is modeled as a vector $\nu(r) \in \mathbb{R}^n$.

To extract samples for training, OWL 2 EL axioms are first normalized into four normal forms: $C \sqsubseteq D$, $C \sqsubseteq \exists r.D$, $\exists r.D \sqsubseteq C$ and $C \sqcap D \sqsubseteq E$. This can be implemented by OWL ontology reasoners such as jCel (Mendez 2012), TrOWL (Thomas, Pan, and Ren 2010) and HermiT (Glimm et al. 2014). A loss function which separately deals with the axioms of each normal form is defined, denoted as \mathcal{L} . For simplicity, we present the loss calculation of $C \sqsubseteq D$ and $C \sqsubseteq \exists r.D$ below as they are the most commonly used in defining the complex concept for a ZSL problem:

$$\begin{aligned} \mathcal{L}(C \sqsubseteq D) = & \max(0, \|\nu(C) - \nu(D)\| + \gamma(C) - \gamma(D) - \epsilon) \\ & + \left| \|\nu(C)\| - 1 \right| + \left| \|\nu(D)\| - 1 \right|, \end{aligned} \quad (1)$$

$$\begin{aligned} \mathcal{L}(C \sqsubseteq \exists r.D) = & \max(0, \\ & \|\nu(C) + \nu(r) - \nu(D)\| + \gamma(C) - \gamma(D) - \epsilon) \\ & + \left| \|\nu(C)\| - 1 \right| + \left| \|\nu(D)\| - 1 \right|, \end{aligned} \quad (2)$$

where $\|\cdot\|$ denotes L2-norm, $|\cdot|$ denotes the absolute value, $\max(\cdot, \cdot)$ denotes selecting the maximum value, ϵ is a hyper parameter for the max margin loss. The subsumption \sqsubseteq is modeled by geometric inclusion of two balls (violation will be punished by increasing the loss), while the existential logic $\exists r.C$ is modeled by translation (i.e., a directed movement of the concept vector). The concept vectors are also normalized. Readers are referred to (Kulmanov et al. 2019) for the loss functions of the remaining normal forms.

In training, axioms of the normal forms are extracted from the ontology as samples. A stochastic gradient descent algorithm is used to learn the embeddings of concepts and relations by minimizing the overall loss. The embedding vector of the concept that corresponds to a ZSL class label y is denoted as $\nu(y)$.

Embedding of Text and Paths We can directly use the word vector of a class label (or the average of vectors of its compositional words), using a pre-trained language model by a general corpus like Wikipedia articles; however, this

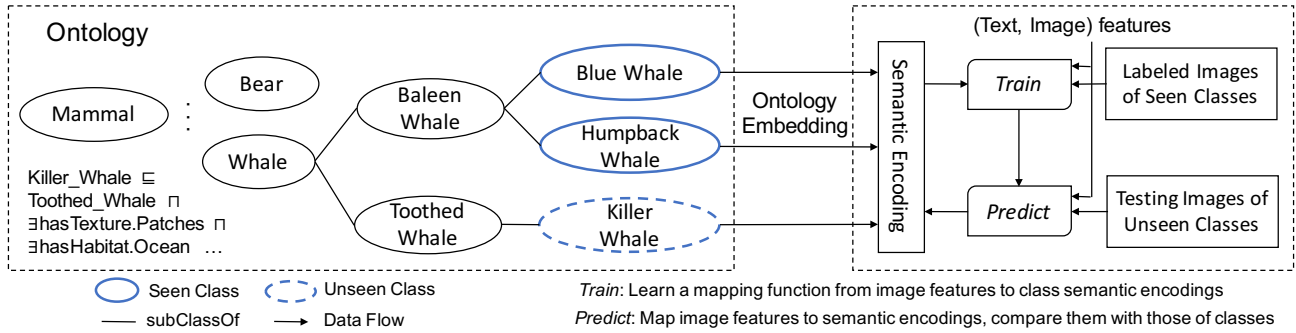


Figure 1. The ZSL Framework with Ontology and Ontology Embedding. The framework is illustrated on an image classification task but could be adapted to any machine learning architectures in *Train* and *Predict*. For instance we could envision applying ZSL on a VQA task where joint features of both question text and image are used as the input.

would ignore the statistical correlation between words and concepts in the ontology. Thus we propose to fine-tune the pre-trained model by a corpus extracted from the ontology which includes the semantics of text and paths, following and extending the idea of OWL2Vec (Holter et al. 2019).

To extract the local corpus we first project the original OWL ontology into an RDF (Resource Description Framework) graph using the method in (Agibetov et al. 2018). For example, the concept inclusion axiom $C \sqsubseteq D$ is transformed into $\langle C, rdfs:subClassOf, D \rangle$, while the existential logic axioms $C \sqsubseteq \exists r.D$ and $\exists r.D \sqsubseteq C$ are both transformed into $\langle C, r, D \rangle$. As we only aim at keeping the relative position (statistical correlation) of concepts, relations and words in the ontology, such approximations are reasonable. We then apply a random walk algorithm over the graph to extract paths, and transform them into word sequences as the new corpus by concatenating the concept and relation labels.

The corpus from the ontology is finally used to fine-tune a pre-trained language model, and its resulting vector of a class label y is denoted as $\omega(y)$. We merge the embedding of logic axioms and the word vector by concatenating them; i.e., $h(y) = [\nu(y), \omega(y)]$.

3.3 Mapping-based Training and Prediction

We adopt a mapping-based ZSL paradigm (see Figure 1) which includes (i) the training of a mapping function from the input features (like image features) to the semantic encoding of seen class labels, and (ii) the prediction of unseen class labels according to nearest neighbour.

Instead of directly learning the target function $f : x \rightarrow y$ as in normal supervised learning, a mapping-based algorithm learns a function from the input to the class label’s semantic encoding space, denoted as $g : x \rightarrow z$. Given the original training samples \mathcal{D}_{tr} , it first calculates the semantic encoding of the class labels and gets $\mathcal{D}'_{tr} = \{(x, z) | z = h(y), (x, y) \in \mathcal{D}_{tr}\}$, and then uses \mathcal{D}'_{tr} to train the function g . Some state-of-the-art mapping-based ZSL algorithms such as DeViSe (Frome et al. 2013) and SAE (Semantic Autoencoder) (Kodirov, Xiang, and Gong 2017) can be adopted. Take SAE as an example, briefly it learns a linear encoder from x to z and a decoder from z back to x , by minimizing the loss of both the encoder (on z) and the

decoder (on x), and adopts the encoder as the function g .

In prediction, for each testing sample x_u in \mathcal{X}_u , the learned function g maps it to a semantic encoding $g(x_u)$, and $g(x_u)$ then is compared with the semantic encoding of each unseen class label in \mathcal{Y}_u . The unseen class label whose semantic encoding leads to the lowest distance to $g(x_u)$ is determined as the label of x_u :

$$f(x_u) = \arg \min_{y_u \in \mathcal{Y}_u} dist(h(y_u), g(x_u)), \quad (3)$$

where $dist(\cdot, \cdot)$ denotes a distance function such as L2-norm and cosine similarity.

4 Experimental Results

Preliminary but promising results are reported on two different ZSL tasks to demonstrate the broad potential impact of ontology embeddings.

4.1 Settings

Animal Image Classification (AIC): We are using the AwA2² (Xian et al. 2018) benchmark, including 50 mammals that are partitioned into 40 seen classes and 10 unseen classes. Each class on average has 746 training / testing images whose 2048-dimension features (used as the input x in our method) have been extracted by a deep residual neural network. The corresponding ontology, constructed with the taxonomy from WordNet and expertise of animal visual characteristics (cf. concrete examples in Methodology), includes 1488 axioms, 180 classes and 12 relations.

Visual Question Answering (VQA): We are considering a VQA task which is to predict an answer for a given pair of image and question. The answer is an object (concept) in the image or a Boolean for closed-ended questions. Features of the image-question pair are learned from both question text and image by a BERT-based Concept-Vision-Language model (cf. (Lu et al. 2019) for the details of the model). The ZSL framework here is to predict answers that have never appeared in the training data. We adopt the Outside Knowledge-VQA (OK-VQA) dataset (Marino et al. 2019)

²<http://cvml.ist.ac.at/AwA2/>

which contains 14,031 images and 14,055 questions (see example in Figure 2). There are 768 seen classes and 339 unseen classes. The ontology is built with ConceptNet knowledge graph (Speer, Chin, and Havasi 2016) and a core \mathcal{EL}^{++} schema which has been operated on ConceptNet properties that connect answer concepts in OK-VQA questions. Specifically 2.3% of ConceptNet properties are involved. Figure 2 presents an example of VQA, where the potential answers of Truck and Firefighter are two seen classes in our ZSL setting while Firetruck is an unseen class.



Figure 2. OK-VQA Example. Question: *What vehicle uses this item?* Answer: Firetruck. In our ontology Firetruck is related to FireHydrant by axioms e.g., $Firetruck \sqsubseteq Vehicle \sqcap \exists hasWaterSupply.(ConnectionPoint \sqcap FireHydrant)$.

4.2 Evaluation

AIC: We calculate the micro accuracy of each unseen class as the ratio of its correctly annotated images, and then average the micro accuracy of all the unseen classes as the metric *accuracy*. Table 1 presents the results of using different semantic encodings: Label_W2V denotes the average word vector of the class label; Attribute denotes the continuous attribute vector, each of whose slots denotes a real value degree of an attribute annotation on visual characteristics (Lampert, Nickisch, and Harmeling 2009); and EL_Embed denotes the embedding of ontology logic axioms. Note that the Label_W2V now adopts *Word2Vec* (Mikolov et al. 2013) pre-trained by the Wikipedia dump in June 2018, while the language model fine-tuned by the ontology corpus will be implemented and evaluated in our future work.

VQA: We evaluate the performance of our ZSL framework by comparing the predicted answer with the ground truth answer, and calculating the metric *accuracy* as the ratio of the testing image-question pairs whose answers are correct. Image-question pairs of both seen and unseen answers are tested. As in AIC, we compare different semantic encodings of the answer label. The state-of-the-art result comes from (Lu et al. 2019) where the word vector of the answer label is adopted (cf. results on VQA in row Label_W2V in Table 1).

4.3 Results and Discussion

In both cases i.e., AIC and VQA (cf. Table 1), the embedding of logic axioms outperforms both label word vector and attribute vector, both of which are widely applied and studied semantic encoding solutions in ZSL (Xian et al. 2018). Meanwhile, we also find concatenating the embedding to

both label word vector and attribute vector significantly improves the accuracy. For example adding EL_Embed improves the accuracy of Label_W2V by respectively 47.3% and 45.5% in AIC and VQA. These observations indicate that the embedding of ontology logic axioms which model the compositional semantics of class labels is an effective and complementary side information for ZSL.

On AIC-specific Results: The best ZSL accuracy on AIC case is now achieved by GNNs (Kampffmeyer et al. 2019; Wang, Ye, and Gupta 2018), but they rely on a big graph with around 21k nodes (concepts) extracted from WordNet. Such a big graph is often unavailable or hard to be constructed in many real word applications. In contrast, our ontology only contains 96 mammal concepts and uses expressive logic axioms to specify the compositional semantics. With this small graph, (Kampffmeyer et al. 2019) only achieves an accuracy of 6.2% due to sample shortage.

On VQA-specific Results: The result with ontology embedding outperforms the state-of-the-art approach, specifically for unseen answers. The semantic encoding has been crucial for encoding semantic relationships among properties and concepts, and derive out-of-knowledge answers i.e., answers which are not given in training data, and only available in the ontology.

Semantic Encoding	AIC	VQA
Label_W2V	34.2%	30.5%
Attribute	47.6%	32.1%
EL_Embed	48.3%	41.6%
Label_W2V + EL_Embed	50.4%	44.4%
Attributes + EL_Embed	56.7%	48.1%
Label_W2V + Attribute	54.6%	34.9%
Label_W2V + Attribute + EL_Embed	58.9%	50.1%

Table 1. The Accuracy of SAE with Different Semantic Encodings (*h*). + denotes vector concatenation.

5 Conclusion and Future Work

In this paper we present a new ZSL framework which utilises an OWL ontology to express the compositional semantics of class labels. It first learns ontology embeddings that encode the semantics of logical axioms, graph paths and text, and then incorporates the embeddings into a mapping-based ZSL paradigm. Such a framework provides a new neural symbolic design pattern for dealing with low-resource learning, via synergistic integration of well expressed formal semantics and neural models. Some preliminary evaluation results have shown the effectiveness of introducing OWL ontology in ZSL. In the future, we will (*i*) develop more robust ontology embeddings, (*ii*) evaluate the framework with more ZSL algorithms such as the generation based (Geng et al. 2020a), and (*iii*) explain the prediction and feature transfer with ontologies (Chen et al. 2018).

Acknowledgments

We want to thank Ian Horrocks from University of Oxford for helpful discussions. The work is supported by the AIDA project (Alan Turing Institute), the SIRIUS Centre for Scalable Data Access (Research Council of Norway), Samsung Research UK, Siemens AG, and the EPSRC projects AnaLOG (EP/P025943/1), OASIS (EP/S032347/1) and UK FIRES (EP/S019111/1).

References

- Agibetov, A.; Jiménez-Ruiz, E.; Ondrésik, M.; Solimando, A.; Banerjee, I.; Guerrini, G.; Catalano, C. E.; Oliveira, J. M.; Patanè, G.; Reis, R. L.; et al. 2018. Supporting shared hypothesis testing in the biomedical domain. *Journal of biomedical semantics* 9(1):9.
- Auer, S.; Bizer, C.; Kobilarov, G.; Lehmann, J.; Cyganiak, R.; and Ives, Z. 2007. DBpedia: A nucleus for a web of open data. In *The semantic web*. Springer. 722–735.
- Baader, F.; Brandt, S.; and Lutz, C. 2005. Pushing the EL envelope. In *IJCAI*, 364–369.
- Chen, J.; Lécué, F.; Pan, J. Z.; Horrocks, I.; and Chen, H. 2018. Knowledge-based transfer learning explanation. In *Sixteenth International Conference on Principles of Knowledge Representation and Reasoning*.
- Frome, A.; Corrado, G. S.; Shlens, J.; Bengio, S.; Dean, J.; Ranzato, M.; and Mikolov, T. 2013. DeViSE: A deep visual-semantic embedding model. In *Advances in neural information processing systems*, 2121–2129.
- Geng, Y.; Chen, J.; Chen, Z.; Ye, Z.; Yuan, Z.; Jia, Y.; and Chen, H. 2020a. Generative adversarial zero-shot learning via knowledge graphs. *arXiv preprint arXiv:2004.03109*.
- Geng, Y.; Chen, J.; Ye, Z.; Zhang, W.; and Chen, H. 2020b. Explainable zero-shot learning via attentive graph convolutional network and knowledge graphs.
- Glimm, B.; Horrocks, I.; Motik, B.; Stoilos, G.; and Wang, Z. 2014. HermiT: an OWL 2 reasoner. *Journal of Automated Reasoning* 53(3):245–269.
- Holter, O. M.; Myklebust, E. B.; Chen, J.; and Jimenez-Ruiz, E. 2019. Embedding OWL ontologies with OWL2Vec. In *CEUR Workshop Proceedings*, volume 2456, 33–36.
- Kampffmeyer, M.; Chen, Y.; Liang, X.; Wang, H.; Zhang, Y.; and Xing, E. P. 2019. Rethinking knowledge graph propagation for zero-shot learning. In *CVPR*, 11487–11496.
- Kodirov, E.; Xiang, T.; and Gong, S. 2017. Semantic autoencoder for zero-shot learning. In *CVPR*, 3174–3183.
- Kulmanov, M.; Liu-Wei, W.; Yan, Y.; and Hoehndorf, R. 2019. El embeddings: geometric construction of models for the description logic EL++. In *IJCAI*, 6103–6109.
- Lampert, C. H.; Nickisch, H.; and Harmeling, S. 2009. Learning to detect unseen object classes by between-class attribute transfer. In *CVPR*, 951–958. IEEE.
- Lei Ba, J.; Swersky, K.; Fidler, S.; et al. 2015. Predicting deep zero-shot convolutional neural networks using textual descriptions. In *Proceedings of the IEEE International Conference on Computer Vision*, 4247–4255.
- Lu, J.; Batra, D.; Parikh, D.; and Lee, S. 2019. ViLBERT: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks.
- Marino, K.; Rastegari, M.; Farhadi, A.; and Mottaghi, R. 2019. OK-VQA: A visual question answering benchmark requiring external knowledge. In *CVPR*.
- Mendez, J. 2012. jcel: A modular rule-based reasoner. In *ORE*.
- Mikolov, T.; Sutskever, I.; Chen, K.; Corrado, G. S.; and Dean, J. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, 3111–3119.
- Miller, G. A. 1995. WordNet: a lexical database for english. *Communications of the ACM* 38(11):39–41.
- Nikolaus, M.; Abdou, M.; Lamm, M.; Aralikkatte, R.; and Elliott, D. 2019. Compositional generalization in image captioning. In *CoNLL*, 87–98.
- Palatucci, M.; Pomerleau, D.; Hinton, G. E.; and Mitchell, T. M. 2009. Zero-shot learning with semantic output codes. In *Advances in neural information processing systems*, 1410–1418.
- Pan, J. Z.; Stamou, G.; Tzouvaras, V.; and Horrocks, I. 2005. f-SWRL: A fuzzy extension of SWRL. In *ICANN*.
- Qin, P.; Wang, X.; Chen, W.; Zhang, C.; Xu, W.; and Wang, W. Y. 2020. Generative adversarial zero-shot relational learning for knowledge graphs.
- Socher, R.; Huval, B.; Manning, C. D.; and Ng, A. Y. 2012. Semantic compositionality through recursive matrix-vector spaces. In *EMNLP-CoNLL*, 1201–1211. Association for Computational Linguistics.
- Speer, R.; Chin, J.; and Havasi, C. 2016. Conceptnet 5.5: An open multilingual graph of general knowledge.
- Teney, D., and van den Hengel, A. 2018. Visual question answering as a meta learning task. In Ferrari, V.; Hebert, M.; Sminchisescu, C.; and Weiss, Y., eds., *Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part XV*, volume 11219 of *Lecture Notes in Computer Science*, 229–245. Springer.
- Thomas, E.; Pan, J. Z.; and Ren, Y. 2010. TrOWL: Tractable OWL 2 reasoning infrastructure. In *ESWC*.
- Wang, X.; Ye, Y.; and Gupta, A. 2018. Zero-shot recognition via semantic embeddings and knowledge graphs. In *CVPR*, 6857–6866.
- Wu, F.; Hoffmann, R.; and Weld, D. S. 2008. Information extraction from Wikipedia: Moving down the long tail. In *SIGKDD*, 731–739.
- Xian, Y.; Akata, Z.; Sharma, G.; Nguyen, Q.; Hein, M.; and Schiele, B. 2016. Latent embeddings for zero-shot classification. In *CVPR*, 69–77.
- Xian, Y.; Lampert, C. H.; Schiele, B.; and Akata, Z. 2018. Zero-shot learning—a comprehensive evaluation of the good, the bad and the ugly. *IEEE transactions on pattern analysis and machine intelligence* 41(9):2251–2265.