On the Potential of Logic and Reasoning in Neurosymbolic Systems using OWL-based Knowledge Graphs

A position paper

David Herron [a,*], Ernesto Jiménez-Ruiz [a,b] and Tillman Weyde [a]
[a] *City, University of London, UK*
E-mails: david.herron@city.ac.uk, ernesto.jimenez-ruiz@city.ac.uk, t.e.weyde@city.ac.uk
[b] *SIRIUS, University of Oslo, Norway*

Abstract.
Knowledge graphs feature ever more frequently as symbolic components in neurosymbolic research and systems. But even though a central concern of neurosymbolic AI is to combine neural learning with symbolic reasoning, relatively little neurosymbolic research focuses on leveraging the logical representation and reasoning capabilities of OWL-based knowledge graphs. The objective of this position paper is to inspire more neurosymbolic researchers to embrace the OWL and the Semantic Web by raising awareness of the benefits, capabilities, and applications of OWL-based knowledge graphs, particularly with respect to logical reasoning. We describe the ecosystem of open W3C standards-based resources available that support the adoption and use of OWL-based knowledge graphs; we describe tools that exist for engineering custom OWL ontologies tailored to particular research needs; we discuss how OWL-based knowledge graph reasoning can add value in applications and how knowledge graph embeddings can leverage symbolic knowledge; we discuss and illustrate the reasoning capabilities of OWL-based knowledge graphs; and we describe several promising directions for research that focus on leveraging these capabilities. We also discuss the specialised resources needed to undertake research on OWL-based knowledge graphs in neurosymbolic systems. We use the example of NeSy4VRD, an image dataset with a custom-designed companion OWL ontology. The scarcity of this kind of resource should be addressed to accelerate research in this field.

Keywords: neurosymbolic, AI, deep learning, Semantic Web, OWL, ontologies, knowledge graphs, reasoning

1. Introduction

Following a long gestation spanning decades, neurosymbolic artificial intelligence (NeSy AI) has recently blossomed into a recognised subfield of AI. While neural and symbolic traditions of AI have been tribally rival, recently there is a vibrant diversity of approaches blending the two [1]. Prompted by analysis of the limitations of deep learning (in, *e.g.*, [2–6]), and despite the recent advances resulting from scaling up deep learning, as evidenced in large language models, increasing numbers of researchers are drawn to NeSy AI. The shared motivation is to explore combinations of neural learning and symbolic knowledge representations in order to get the best of both worlds, in a shared belief that this is the best route for advancing AI towards artificial general intelligence.

*Corresponding author. E-mail: david.herron@city.ac.uk.

Knowledge graphs (KGs) are representations of symbolic knowledge that conform to a graph model, where nodes are concepts and entities of interest and edges are relationships between them [7, 8]. As NeSy research has expanded, so has the frequency with which KGs feature as symbolic components in hybrid, NeSy systems [9]. The theme of 'deep deductive reasoning', where neural networks (NNs) are trained to reason over KGs, is progressively developing [10–12]. KGs can be particularly effective when data samples are expensive, difficult or impossible to obtain, so that there is a lack of data to train robust deep learning-based systems, as in few-shot and zero-shot learning scenarios [13–15].

The Web Ontology Language (OWL) [16, 17] is a key component of the Semantic Web technology stack [18, 19]. OWL is used to define ontologies (semantic schemas enriched with logic) that can govern Semantic Web KGs by specifying what assertions of knowledge (types of triples) are admissable and inadmissable. The inference semantics encoded in OWL ontologies allows them to be used in OWL-based KGs to support both the inference of new KG symbolic knowledge (new triples) and the validaton of KG logical consistency.

Given that a central concern of NeSy AI is to explore ways of combining neural learning with symbolic reasoning, it is curious that relatively little NeSy research has focused on exploring how to leverage the logical reasoning capabilities of OWL-based KGs. A mapping study of 476 recent papers that explore approaches to combining Semantic Web technologies with machine learning [20] reports that only 29 (6%) of these papers mention using semantic processing units, and of these only 20 (4%) mention using reasoning capabilities to infer new knowledge. One explanation for this is likely the cross-disciplinary nature of the endeavour: NeSy research with OWL-based KG reasoning requires researchers to be familiar not just with deep learning, KGs and logic but with Semantic Web technologies and specifically OWL, as well.

Our aim with this paper (which is an extended version of [21]) is to promote more NeSy AI research using OWL-based KGs. We hope to both inspire and enable. We hope to inspire more NeSy research using OWL-based KGs by raising awareness of their benefits, capabilities, and applications, especially with respect to logical reasoning and agency. OWL-based KGs are exemplars of the explicit symbolic knowledge representation and symbol manipulation and reasoning machinery that critics of deep learning, such as [2–6], advocate be incorporated in hybrid, NeSy systems. Towards this end, we draw upon illustrative examples from our own research and others to describe promising research directions. We finally discuss how to enable more NeSy research using OWL-based KGs by creating resources such as the recently contributed NeSy4VRD (Neurosymbolic AI for Visual Relationship Detection). NeSy4VRD addresses the scarcity of the specialised resource required for such research: a dataset for neural learning accompanied with a tailor-made OWL ontology describing the domain of the dataset for symbolic reasoning.

2. Benefits and capabilities of OWL-based KGs

In this section, we describe benefits and capabilities of OWL-based KGs. We illustrate capabilities by giving examples showing how and why OWL-based KGs can be utilised in NeSy systems.

2.1. Open standards and reusable resources

The Web Ontology Language (OWL) [16, 17] and OWL-based knowledge graphs (KGs) [7, 8] are key components of the W3C open standards ecosystem of the Semantic Web (SW) [18, 19, 22, 23]. Open standards facilitate interoperability and promote development of reusable, often free, software resources that make it easy to work with OWL-based KGs. Amongst the many such resources are: *(i)* public SW KGs like DBpedia [24], Wikidata [25] and Yago [26]; *(ii)* public repositories of curated OWL ontologies like BioPortal [27] and OBO Foundry [28] in the biomedical domain; *(iii)* RDF stores like GraphDB (not open, but has free version) [29] and RDFox (not open, but has free academic license) [30]; and *(iv)* efficient OWL reasoners like HermiT [31], Pellet [32], RDFox and ELK [33].

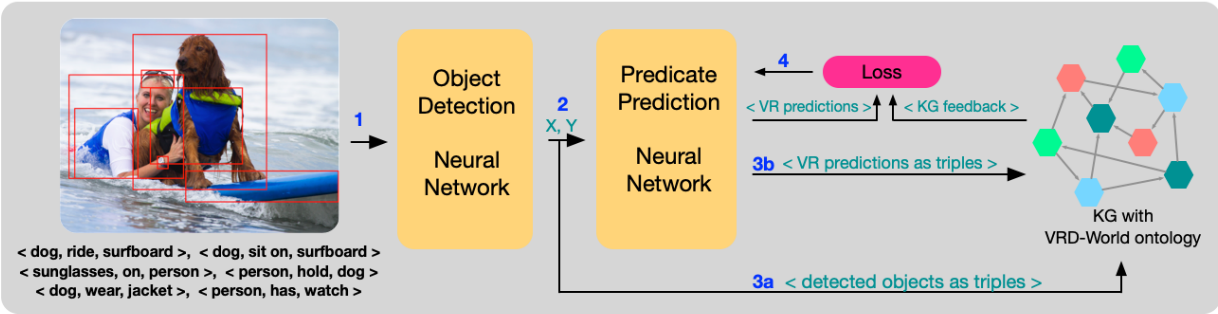

Fig. 1. An example illustration of a hybrid neurosymbolic system architecture for detecting visual relationships in images. Here, reasoning over an OWL-based knowledge graph is used to guide neural learning.

2.2. *Custom ontologies and custom KGs*

Reusing state-of-the-art ontologies and/or public KGs is a good practice option. But researchers can also design their own custom, domain-specific OWL ontologies tailored to their unique needs and use them to govern and enable reasoning within custom OWL-based KGs. Custom ontologies can be aligned with publicly available ontologies to enhance interoperability [34].

This is the approach taken for visual relationship detection in images in the design of a custom OWL ontology, called VRD-World [35]. This ontology describes the domain of the common, everyday images of the VRD dataset [36], as reflected in the object classes and relationships referred to in the `(subject, predicate, object)` visual relationships annotated for the images. As depicted in Figure 1, the VRD-World ontology can govern a custom KG in the hybrid NeSy systems with which we explore using symbolic reasoning to guide neural learning. While designing the VRD-World ontology guidance was taken from the large literature on ontology engineering [*e.g.*, 37–40]. The ontology was specified using the free ontology editor Protégé [41], taking advantage of free Protégé plug-in utilities designed to support ontology development, such as ontology debuggers. Many machine learning tools exist to support various different aspects of ontology development such as, for example, concept learners (see [42]).

We designed two versions of a class hierarchy for our ontology. One version is entirely custom designed and represents the broad range of everyday object classes of the VRD dataset (`person`, `dog`, `jacket`, `surfboard`, *etc.*) feature exclusively as leaf nodes. In the other version, the VRD object classes were first aligned with matching classes in Wikidata [25], and a small number of subsumption paths were selected for each such that this class hierarchy represents a faithful, tractable subset of the Wikidata class hierarchy.

2.3. *KG embeddings, KG completion and knowledge injection*

KGs (of all kinds) have inspired a large amount of NeSy research into encoding KG symbolic background knowledge into vectors as *KG embeddings*. The embeddings preserve semantic similarity and reflect this similarity by proximity within the embedding vector space [42–47]. The primary application area of KG embeddings so far has been tasks relating to *KG completion*: *link prediction* (relating individuals in a KG) or *type prediction* (classifying individuals in a KG). Regardless of the model used to generate the embeddings (of which there are many), these link and type prediction problems are typically cast as neural classification problems, where the embedded KG knowledge is used for training and methods exploiting the proximity principle are applied.

Like all KGs, OWL-based KGs can are readily used in NeSy research that leverages KG embeddings. OWL2Vec* [48] is one embedding model designed for this purpose. Notice, though, that these applications of KG embeddings focus on leveraging KG symbolic background knowledge only. So even if the KG in question is OWL-based, its reasoning capabilities are generally not employed in these applications.

Link inference and *type inference* by logical reasoning are, however, the bread and butter of OWL reasoners. When an OWL reasoner infers the knowledge that is entailed by the inference semantics of a governing OWL ontology in the presence of KG data, it completes the KG by introducing new, explicit (knowledge) triples that were previously implicit (materialisation). The logical soundness of these inferences is guaranteed, whereas embedding-based KG completion is approximate and potentially error prone. The extent to which the KG is extended (completed) is commensurate with the richness of the inference semantics of the governing ontology and the nature of the KG data present at the time of materialisation. Our point is that OWL-based KGs can add important value in any NeSy task associated with KG completion. OWL reasoning can be used to complete a KG automatically, as far as possible, and then NeSy KG completion (NN emulated reasoning) can be used for special cases that the OWL ontology in question does not address or that OWL cannot address in general.

A smaller, but growing, application area of KG embeddings is variously called *knowledge injection* or *knowledge infused learning* [49, 50]. Here, the embedded KG knowledge is injected into the internals of a NN's architecture rather than fed into the front. It is not used as the primary input data; instead, it is used as a knowledge supplement to whatever the primary input data is. OWL-based KGs can readily be used in all such research.

With either application area of KG embeddings (KG completion tasks or knowledge injection tasks), OWL-based KGs and their reasoning capabilities can add value. A fully materialised OWL-based KG (where everything implicit has been made explicit) will contain more knowledge to embed and deliver richer embeddings.

2.4. *OWL-based KG reasoning and agency*

Despite their recent success, large language models are notorious for their lack of reliability in reasoning. In contrast, the reliability of OWL reasoning is guaranteed because it is grounded in formal Description Logics (DLs) that are decidable fragments of first-order logic [51–54], such as the highly expressive DL *SROIQ* that is used in the latest version of OWL, OWL 2 [55]. Sound, logical reasoning based on OWL can be leveraged in different ways. The most common ways are to infer new knowledge (*i.e.* introduce new, explicit triples into a KG) and to check and enforce logical consistency in a KG. Both of these capabilities are commonly used to debug OWL ontologies during development [8, 56]. Crucially, for NeSy systems, they can also be leveraged to enable symbolic reasoning to guide neural learning. Because OWL-based KG technologies support SPARQL interfaces (the Semantic Web query language, akin to SQL for relational databases), OWL-based KGs can be used as active reasoning components, not just data stores.

We depict a generic neurosymbolic system using the VRD dataset in Figure 1, where either ground-truth or predicted visual relationships can be converted into RDF triples and inserted into a KG governed by the VRD-World OWL ontology. The results of symbolic OWL reasoning (whether new inferred knowledge or feedback as to inconsistency) can then be used to guide neural learning by factoring them into the calculation of loss. The availability of an OWL-based KG means that OWL-compatible reasoners can also be leveraged during NN inference. Predictions generated at inference time can be evaluated and semantically invalid ones filtered out.

Link inference in OWL-based KGs is driven by the inference semantics associated with the characteristics and relationships declared for the object and data properties of an OWL ontology. In the VRD-World ontology, the 70 predicates of the VRD dataset (mostly common spatial relations and verbs) are represented by object properties that permit a rich web of characteristics (*e.g.*, symmetry, transitivity) and relationships (*e.g.*, inverses, subproperties, equivalent properties) to be defined that, in turn, enable rich reasoning. Property `beside` is declared to be symmetric; so if `(A, beside, B)`, OWL reasoning can infer `(B, beside, A)`. Property `over` is declared to have as inverse the property `under`; so if `(C, over, D)`, OWL reasoning can infer `(D, under, C)`. Loss can be penalised to encourage a NN to learn these logical implications of its own predictions.

Type inference in OWL-based KGs is driven by the class hierarchy of an OWL ontology, which is declared using the property `rdfs:subClassOf`. If `classA rdfs:subClassOf classB` and it is asserted that individual `X` is a member of `classA`, then OWL reasoning can infer that `X` is also a member of `classB`. One way to leverage type inference is in Datalog rules that extend OWL's reasoning capabilities so that nuanced inference cases beyond OWL's reach can be captured. We intend to explore this opportunity, *e.g.*, with RDFox, which implements a fast engine that seamlessly blends reasoning over the OWL 2 RL profile and Datalog rules. For example, a Datalog rule describing when it is reasonable to infer the visual relationship `(X, wear, Y)` might be represented as

```
wear(X,Y) := WearCapableThing(X), WearableThing(Y), ir(Y,X) > 0.8
```

Suppose individuals `X` and `Y` are asserted to be members of the classes predicted by an object detector (*e.g.*, say, `Dog` and `Jacket`, respectively). The type inference of OWL reasoning will determine the higher-level classes of which individuals `X` and `Y` are also members. In the VRD-World ontology, two such classes are `WearCapableThing` and `WearableThing`. The outcomes of this type inference will in turn determine if the goals `WearCapableThing(X)` and `WearableThing(Y)` in the Datalog rule are satisfied or not. (In this example, the function `ir()` measures an *inclusion ratio* — the extent to which the bounding box for `Y` is included within the bounding box for `X`.)

3. Promising Research Directions

Here we describe areas where the potential for leveraging the reasoning and agency capabilities of OWL-based KGs in hybrid NeSy systems looks particularly promising.

3.1. OWL-based KG reasoning to strengthen weak labelling

The visual relationship annotations of the images of the VRD dataset are sparse and arbitrary rather than exhaustive; the supervision they provide is thus partial and inconsistent, both within and between images. For example, image *A* might have a person riding a horse and an associated annotated visual relationship `(person, ride, horse)` noting that fact, whereas image *B* might have a person riding a horse and there is no associated annotated visual relationship. Additionally, many instances of data conditions suited to few-shot and zero-shot learning exist in the VRD dataset. Taken together, these characteristics represent opportunities for exploring the ability of OWL reasoning (either alone or extended with Datalog rules) to infer knowledge so as to strengthen weak supervision.

Suppose an NN pipeline (per Figure 1) predicts `(person, ride, horse)` for image *B*, and suppose this prediction is entirely valid and correct. The absence of a matching ground-truth annotated visual relationship would typically lead to this prediction being treated as incorrect (a false positive) and to the generation of a large loss contribution. But penalising the loss in this scenario is likely to be counter-productive and to hinder efficient and effective neural learning rather than promote it. If, however, reasoning infers, given the available information (in this case, a person, a horse, and the particular spatial relation between their bounding boxes), that a `(person, ride, horse)` prediction is plausible, then the inferred visual relationship can be used as a stand-in for the missing annotated ground-truth visual relationship, and loss can be calculated as if a matching ground-truth had in fact existed. This way, the loss feedback delivered to the NN to guide its learning will be constructive rather than disruptive.

This notion of using OWL-based KG reasoning to judge whether a prediction is plausible or implausible when ground-truth annotation alone cannot be relied upon to provide such guidance has applications well beyond the VRD dataset. It applies to other non-exhaustively annotated and k-shot supervised learning scenarios, whether within the vision or other domains. It may also be relevant to semi-supervised learning problems (where some examples are labelled, others not), and potentially to unsupervised learning problems as well. Lastly, this line of research might also extend into exploration of probabilistic OWL-based KG reasoning (which has been implemented in proprietary form by remote diagnostics AI startup $HOME \cdot X$[1]), where the probability of a prediction's plausibility might be factored into the calculation of loss when no ground-truth exists.

3.2. Enabling neural networks to emulate OWL-based KG reasoning

One approach to NeSy AI involves equipping NNs with background knowledge and learning biases by introducing structural extensions to their architectures that represent strong priors. An example of this approach is [57]. As part of our research, we have explored this approach to NeSy by considering the feasibility of transferring OWL-based KG reasoning capabilities to NNs in the form of structural extensions representing strong priors. In the process of doing so, we have developed a technique (as a proof-of-concept) for representing the class hierarchy of an

[1] https://homex.com/

OWL ontology as a structural extension to the architecture of an object detection classification NN. This structural extension can equip an object detection NN with the ability to precisely emulate the subsumption reasoning (type inference) capabilities of an OWL-based KG such that it can return not only the predicted class of a detected object (e.g. Dog) but all of its parent classes as well (e.g. Carnivore, Mammal, Animal, LivingThing, etc., as per the class hierarchy of the OWL ontology in question).

Since both NNs and OWL-based KGs (including class hierarchies of OWL ontologies) can be viewed as being directed graphs, we looked to graph theory for inspiration as to how to effect a transfer of subsumption reasoning capability from OWL-based KGs to NNs. One basic result from graph theory is that a graph with n nodes can have its structure encoded in an *adjacency matrix*, where a 1 indicates that two nodes are connected by an edge. An adjacency matrix can readily encode the structure of the directed graph of an OWL class hierarchy, where OWL classes are the nodes and the `rdfs:subClassOf` links between OWL classes are the (directed) edges. It turns out that the adjacency matrix of the graph of the transitive closure of an OWL class hierarchy fully encodes all of the subsumption reasoning capability of an OWL-based KG. By materialising the OWL ontology containing the class hierarchy of interest, OWL reasoning infers the entire transitive closure of the class hierarchy and, in the process, makes all implicit `rdfs:subClassOf` axioms explicit. The explicit `rdfs:subClassOf` axioms can then be extracted (e.g. via SPARQL query) to build the adjacency matrix for the graph of the transitive closure of the OWL class hierarchy. This adjacency matrix can then be used as the (fixed) weight matrix for one extra linear layer (with no activation function) of an object detection NN, enabling normal forward-pass computation to generate (pseudo-infer) the parent classes of the classes of all detected objects.

The idea of transferring OWL-based KG reasoning capabilities to NNs by leveraging results from graph theory, as illustrated by the example just discussed, is ripe for further exploration. One candidate avenue is to explore adding more learnable layers to an object detection NN following the class generalisation layer (as just described) so that learning can proceed driven by generalised class predictions. Another avenue lies in recognising that the solution for transferring subsumption reasoning capabilities to NNs works because the `rdfs:subClassOf` property is *transitive*. However, not only subclass relationships can be transitive in OWL ontologies, so the technique described may be applied to transitive properties in general. Similarly, adjacency matrices for graphs (or subgraphs) of OWL-based KGs can, in theory, be constructed in relation to any given object property of an OWL ontology. Suppose an object property called `beside` that is declared to be *symmetric*, such that if `(:A :beside :B)`, an OWL reasoner can infer that `(:B :beside :A)`. An adjacency matrix encoding these relations will itself be symmetric, and may potentially be capable of being leveraged in the subsymbolic space of a NN to encode the inference semantics associated with a symmetric object property of an OWL ontology, so as to enable the NN to replicate the symmetry-related reasoning of an OWL-based KG. Finally, suppose an OWL object property called `over` that is declared in an OWL ontology to have as inverse the property called `under`, such that if `(:A :over :B)`, an OWL reasoner can infer that `(:B :under :A)`. The adjacency matrices for OWL-based KG data relating to these two properties are the transpose of one another, so only one of them is needed to enable the inference semantics associated with object properties that are inverses of one another to be encoded and replicated in a subsymbolic space.

3.3. Using OWL-based KG reasoning as logical constraints

Much NeSy research explores using background knowledge expressed in first-order or propositional logic axioms as constraints to guide neural learning, often by manipulating loss to encourage constraint satisfaction. Examples are the NN training framework Logic Tensor Networks (LTN) [58, 59] which allows fuzzy, first-order *Real Logic* knowledge axioms (constraints) to be defined over training data, the set of propositional logic constraints specified for the ROAD-R dataset [60], and [61]. OWL reasoners can do more than infer new knowledge entailed by a KG's OWL ontology and data — they can check and enforce the logical consistency of a KG. The Pellet reasoner, for example, can detect when a KG is inconsistent and can explain why by listing the logical justifications. Some KG technologies enforce logical consistency as data triples are inserted and reject insertion if permitting it would make the KG inconsistent. A triple insertion rejection is thus equivalent to (one or more) logical constraints being violated and thus can be used to penalise loss. These abilities to check and enforce logical consistency mean that many of the ontological rules expressed in a KG's OWL ontology can be used as direct counterparts of logical constraints. As a consequence, OWL-based KGs are well-positioned to participate in research associated with the *logical constraints* approach to NeSy.

One category of ontological rules where this assertion applies relates to the domain and range restrictions that can be defined for the object properties of OWL ontologies. In [62], Donadello uses LTN and negative domain/range LTN *Real Logic* axioms (constraints) to train predicate predictors on the VRD dataset. For example, given that the VRD dataset has 100 object classes, to train a binary classification NN to predict predicate `wear` using this logical constraint training strategy, approximately 96 negative domain LTN axioms would have been required, such as

$$\forall xy\ wear(x,y) \rightarrow \neg Laptop(x) \quad \forall xy\ wear(x,y) \rightarrow \neg Sofa(x) \quad \forall xy\ wear(x,y) \rightarrow \neg Tree(x) \quad \ldots,$$

expressing the background knowledge that the data classes on the right of these logical implications are not capable of wearing things. Similarly, approximately 80 negative range LTN axioms would have been required, such as

$$\forall xy\ wear(x,y) \rightarrow \neg Table(y) \quad \forall xy\ wear(x,y) \rightarrow \neg Car(y) \quad \forall xy\ wear(x,y) \rightarrow \neg Oven(y) \quad \ldots,$$

expressing the background knowledge that the data classes on the right of these implications are not wearable things.

The VRD-World ontology expresses the logical constraint equivalent of these 96 negative domain LTN *Real Logic* axioms by defining the class `WearCapableThing` in its class hierarchy and by declaring that the domain of object property `wear` is restricted to members of this class, using the single OWL axiom

```
vrd:wear rdfs:domain vrd:WearCapableThing .
```

Similarly, the logical constraint equivalent of the 80 negative range LTN *Real Logic* axioms is established by defining the class `WearableThing` in the class hierarchy of VRD-World, and by declaring that the range of object property `wear` is restricted to members of this class, again using a single OWL axiom:

```
vrd:wear rdfs:range vrd:WearableThing .
```

Figure 1 shows how an OWL-based KG with an appropriate ontology (such as VRD-World) can be used, in the guise of an active reasoning agent, to leverage ontological rules as logical constraints to guide neural learning. Suppose the Object Detection neural network predicts that `objectX` is a dog and `objectY` is a surfboard. If the multi-class, multi-label Predicate Prediction neural network shows a tendency to predict a visual relationship such as `(dog, wear, surfboard)`, the RDF triples representing this prediction

```
vrd:objectX  rdf:type  vrd:Dog
vrd:objectY  rdf:type  vrd:Surfboard
vrd:objectX  vrd:wear  vrd:objectY
```

can be inserted into the KG for evaluation. Type inference (subsumption reasoning) will infer that while `:objectX` is a `WearCapableThing` (*i.e.*, in VRD-World, dogs can wear things), `:objectY` (a surfboard) is not a `WearableThing`. Given that VRD-World also declares classes `WearCapableThing` and `WearableThing` to be suitably *disjoint* from other classes, the OWL-based KG will detect that the range restriction of property `wear` has been violated. This feedback, that an ontological rule (logical constraint) has been violated, can be used to penalise loss to help a NN learn which visual relationships are semantically invalid.

In addition to illustrating that OWL-based KGs can emulate the logical constraints approach to NeSy AI, this example also illustrates an important advantage possessed by OWL-based KGs over the logical constraints approach. The research in [62] shows that the logical constraints approach to NeSy AI is exposed to the risk of combinatorial explosion, where the number of constraints requiring expression grows too rapidly with the number of classes in the dataset. Almost 200 LTN `Real Logic` axioms were needed in relation to just one VRD predicate, `wear`. And about 30 of the 70 VRD predicates admit domain and/or range restrictions of some kind. Indeed, [62] reports implementing a "tractable sample" only of the LTN *Real Logic* axioms implied by the negative domain/range constraints training strategy he selected. In contrast, once an appropriate class hierarchy is defined, expressing powerful domain and range restrictions in OWL is easy.

Another category of ontological rules that can be used as counterparts of logical constraints relates to the use of OWL constructs for declaring OWL classes and OWL properties to be disjoint. The autonomous vehicle driving videos and annotated bounding boxes of the ROAD-R dataset [60] are accompanied by 243 manually specified propositional logic constraints that define the permissible combinations of labels for 10 agent classes, 19 agent action classes, and 12 agent location classes. The 243 logical constraints contain 45 with a format such as (¬Car ∨ ¬Bus), meaning "a car cannot be a bus", that express mutual exclusiveness between the 10 agent classes. Precise counterparts of these 45 propositional constraints can be represented in OWL with two axioms that declare the set of agent classes (from the class hierarchy of an appropriate OWL ontology) to be mutually disjoint, such as the axioms

```
:DisjointAgents rdf:type owl:AllDisjointClasses .
:DisjointAgents owl:members (:Car :Bus :Motorbike :Pedestrian ... ) .
```

Similarly, 66 of the ROAD-R propositional constraints express pairwise mutual exclusiveness amongst the 12 agent location classes. Counterparts of these can be represented in OWL using two more such axioms.

Using suitable declarations of disjoint classes, disjoint properties, functional properties and domain/range restrictions, it is entirely conceivable that an OWL ontology can be designed which expresses (the set or a superset of) the 243 propositional logical constraints specified for the ROAD-R dataset. During NN training, the labels predicted for each bounding box of the ROAD-R dataset could be converted into appropriate RDF triples and inserted into a KG hosting this ontology for evaluation. If OWL reasoning determines that the KG has become inconsistent, thus indicating that one or more ontological rules (aka logical constraints) has been violated, the loss could be penalised appropriately, as per the ROAD-R experiments described in [60].

An alternate strategy for using OWL to emulate the propositional logical constraints of the ROAD-R dataset is to extend OWL with the concept of *integrity constraints* described in [63]. Here, instead of checking constraint violation using OWL reasoning to evaluate KG logical consistency, integrity constraints employ Datalog rules (constructed in relation to an OWL ontology) to represent logical constraints. Constraint violation is signalled through the inference of new knowledge, *e.g.*, by inferring a new instance of class `Violation`, say, as in

```
Violation(X) :- Car(X), Bus(X)
```

meaning an agent cannot be both a car and a bus.

One advantage of the propositional logical constraints associated with the ROAD-R dataset is that they lend themselves to statistical analysis as to which ones have been violated and with what relative frequency. Feedback from OWL reasoning saying the KG has become inconsistent does not permit this same granularity of analysis quite as readily. But OWL reasoners also return logical justifications that explain why a KG has become inconsistent, so analysis of these should permit the logical constraint violation metrics described in [60] to be closely approximated, if not precisely replicated. In summary, repeating the ROAD-R experiments described in [60] by replacing their propositional logical constraints with a custom-designed OWL ontology hosted in a KG operating as a reasoning agent represents a feasible and interesting NeSy research project.

3.4. Integrating OWL-based KG reasoning with existing NeSy frameworks

OWL-based KG symbolic knowledge and deductive reasoning can be integrated with and leveraged by existing logic-based NeSy frameworks such as LTN. So long as *(i)* there is sufficient contextual information contained in the tensors of NN input data (or otherwise) to permit meaningful SPARQL queries to be constructed, and *(ii)* the KG's responses to those SPARQL queries can be mapped to fuzzy truth values in $[0, 1]$, then functions encapsulating interactions with OWL-based KGs that leverage their reasoning capabilities can participate in the LTN *Real Logic* knowledge axioms used to train NNs.

One application of this concept involves using OWL-based KG reasoning to solve the combinatorial explosion problem (described in Section 3.3) to which the *logical constraints* approach to NeSy is exposed. A prime cause of exposure to this risk derives from the fact that logical constraints (as used by the LTN training framework, and the ROAD-R dataset, for example) are restricted to being expressed in terms of the low-level, granular object classes present in the data and their annotations. The inability to express constraints in terms of higher-level, more general classes that have wide catchment areas that cover multiple, lower-level data class cases means that the number of logical constraints requiring expression cannot be managed (kept small). In contrast, as we have seen in examples described above, OWL ontologies routinely possess rich class hierarchies that permit ontological rules to be defined in terms of high-level, general classes, which affords simplicity and parsimony.

To illustrate, let us revisit the research undertaken in [62], where Donadello uses LTN with a training strategy based upon defining an intractable number of negative domain/range LTN *Real Logic* knowledge axioms (constraints). This time, however, suppose that instead of accepting that we must express our logical constraints in terms of the low-level data classes present in the data, we opt to integrate an OWL-based KG into our LTN *Real Logic* knowledge axioms in order to transform the data by mapping the low-level data classes to higher-level, more general classes defined in the class hierarchy of an OWL ontology (in this case, our VRD-World ontology). Using this strategy, we can imagine replacing the original (close to) 200 negative domain/range constraints associated with VRD predicate `wear` (as described in Section 3.3) with one positive LTN *Real Logic* domain constraint and one positive LTN range constraint that are (effectively) precise counterparts of the domain and range restrictions defined for OWL object property `wear` defined in the VRD-World ontology:

$$\forall xy\ wear(x,y) \rightarrow WearCapableThing(x) \qquad \forall xy\ wear(x,y) \rightarrow WearableThing(y)$$

One implementation of this strategy would involve encapsulating within the functions `WearCapableThing` and `WearableThing` SPARQL interactions with a KG (acting as an active reasoning agent) hosting our VRD-World ontology. For each training example in the $n \times m$ input data tensors x and y, the low-level object class (represented as a binary one-hot vector) would be converted to an RDF triple, such as (`:objectP rdf:type :ClassQ`), and inserted into the KG. A corresponding SPARQL query would then be executed to enquire whether `:objectP` had been deemed to be a member of class `WearCapableThing` (or `WearableThing`, respectively). The functions would each return a 1D binary tensor reflecting the results of the data mapping (class generalisation), where 0 indicates that the object class for that training example is not a member of the high-level class in question, and a 1 indicates that it is a member.

An alternate and more compute-efficient implementation that gives identical data transformation results would be to encapsulate within the functions `WearCapableThing` and `WearableThing` use of the *adjacency matrix* of the transitive closure of the class hierarchy of the VRD-World OWL ontology, as described in Section 3.2. This matrix encodes the full class subsumption reasoning capabilities of the VRD-World ontology. From input tensor x (or y, respectively), the function in question would fashion an $n \times c$ tensor, D, containing binary one-hot representations (of size c) of the low-level data classes present in x (or y, respectively). Matrix D would then be multiplied against the $c \times c$ adjacency matrix A. The result, DA, would be an $n \times c$ matrix containing the required subsumption reasoning outcomes. The column of matrix DA that corresponds to class `WearCapableThing` (or `WearableThing`, respectively) will contain the binary outcomes reflecting the results of the data mapping (class generalisation) exercise and this column can be returned by the function, ready for use in the fuzzy logic execution of the `Real Logic` axiom in which it participates.

The two implementations of the data transformation (class generalisation) strategy described here will give identical results. Hence, either approach can be used to overcome the combinatorial explosion problem encountered in [62]. In this example, since we require to leverage only the subsumption reasoning (class generalisation) capabilities of an OWL-based KG, we can entertain the use of the adjacency matrix and enjoy the speed benefits without sacrificing any reasoning flexibility. In other scenarios, however, where OWL-based KG reasoning cannot be adequately replicated (such as being encoded in a matrix), real-time KG access may still present a viable and useful option.

4. Enabling NeSy Research using OWL-based KGs with NeSy4VRD

Sections 2 and 3 focus on inspiring more NeSy research using OWL-based KGs by highlighting their benefits, capabilities, and applications, especially with respect to deductive reasoning and agency. But inspiration alone may not be sufficient because to undertake such NeSy research in a practical way, researchers need a specialised type of resource, which is scarce. A resource is needed that combines a dataset for neural learning with a strongly-aligned companion OWL ontology that describes the domain of the dataset and supports symbolic reasoning that is directly pertinent to the application task for which the dataset is designed. We suspect that the scarcity of such resources represents a barrier to entry that is likely inhibiting NeSy research using OWL-based KGs that might otherwise be undertaken. A resource of this kind, *NeSy4VRD* (Neurosymbolic AI for Visual Relationship Detection), was co-developed and published by the authors of this paper [35].

NeSy4VRD consists of the following components and services:

1. the images of the original VRD dataset [36] (distributed with permission from one of the principals associated with its creation) in order to make them publicly available once again;
2. quality-improved versions of the original VRD visual relationship annotations that have been comprehensively customised and extended to enable the engineering of a robust ontology;
3. a strongly-aligned, custom-designed companion OWL ontology, called VRD-World, that precisely describes the domain of the images and visual relationships;
4. sample Python code for loading the annotated visual relationships into a knowledge graph hosting the VRD-World ontology, and for extracting them from a knowledge graph and restoring them to their native format;
5. support for extensibility of the annotations (and, thereby, the ontology) in the form of *(a)* comprehensive Python code enabling deep but easy analysis of the images and their annotations, *(b)* a custom, text-based protocol for specifying annotation customisation instructions declaratively, and *(c)* a configurable, managed Python workflow for customising annotations in an automated, repeatable process;
6. comprehensive documentation describing *(a)* how to use the extensibility support infrastructure, *(b)* how to share annotation/ontology extensibility projects undertaken by researchers in pursuit of their private research interests, *(c)* how to reuse shared extensibility projects and use the NeSy4VRD workflow to compose them in novel combinations, and *(d)* how the ability to undertake, share, reuse and compose NeSy4VRD extensibility projects represents a new model of collaborative data annotation that we call Distributed Annotation Enhancement.

The NeSy4VRD dataset package (VRD images, quality-improved visual relationship annotations, and companion VRD-World OWL ontology) is distributed on Zenodo[2]. The NeSy4VRD extensibility support infrastructure and comprehensive documentation are available on GitHub[3].

5. Conclusion

A central concern of NeSy AI research is to explore ways of combining neural learning with symbolic background knowledge and reasoning. OWL-based KGs are exemplars of symbolic knowledge representation and reasoning technology and machinery. They can do everything that general KGs can do in terms of representing symbolic knowledge and generating embeddings, plus they can perform sound deductive reasoning to both infer new knowledge and enforce logical consistency, and they can do so in the guise of active, reasoning agents. Given these attractive features, OWL-based KGs warrant more research attention from the NeSy community than they have received to date. Their potential for contributing to NeSy AI is not being fully explored. By describing and illustrating their benefits, capabilities and flexible applications, we hope to inspire more such research. By contributing NeSy4VRD to the NeSy community, a specialised and scarce dataset resource, we hope to lower barriers to entry and thereby enable more such research. A recent overview of NeSy systems [64] reports success using an OWL-based KG to boost expert user satisfaction with large language model performance. Like us, the authors strongly advocate the use of KGs (general and OWL-based) as symbolic components in hybrid NeSy systems.

[2]https://doi.org/10.5281/zenodo.7916355
[3]https://github.com/djherron/NeSy4VRD

References

[1] M.K. Sarker, L. Zhou, A. Eberhart and P. Hitzler, Neuro-Symbolic Artificial Intelligence: Current Trends, *CoRR* (2021). https://arxiv.org/abs/2105.05330.

[2] F. Chollet, Deep Learning: Current Limits and What Lies Beyond Them, 2018. https://raais.co/speakers-2018.

[3] G. Marcus, Deep Learning: A Critical Appraisal, *CoRR* (2018). http://arxiv.org/abs/1801.00631.

[4] G. Marcus, Innateness, AlphaZero, and Artificial Intelligence, *CoRR* (2018). http://arxiv.org/abs/1801.05667.

[5] G. Marcus, The Next Decade in AI: Four Steps Towards Robust Artificial Intelligence, *CoRR* (2020). https://arxiv.org/abs/2002.06177.

[6] H. Kautz, The third ai summer: Aaai robert s. engelmore memorial lecture, Vol. 43, 2022, pp. 105–125.

[7] A. Hogan, E. Blomqvist, M. Cochez, C. d'Amato, G. de Melo, C. Gutiérrez, S. Kirrane, J.E. Labra Gayo, R. Navigli, S. Neumaier, A.-C. Ngonga Ngomo, A. Polleres, S.M. Rashid, A. Rula, L. Schmelzeisen, J.F. Sequeda, S. Staab and A. Zimmermann, *Knowledge Graphs*, Synthesis Lectures on Data, Semantics, and Knowledge, Springer, 2021. https://kgbook.org/.

[8] M. Kejriwal, C.A. Knoblock and P. Szekely, *Knowledge Graphs: Fundamentals, Techniques, and Applications*, The MIT Press, 2021.

[9] P. Hitzler, Knowledge Graphs in Neuro-Symbolic AI, 2021, (Presentation slides). https://people.cs.ksu.edu/~hitzler/pub2/2021-10-AKBC.pdf.

[10] F. Bianchi and P. Hitzler, On the Capabilities of Logic Tensor Networks for Deductive Reasoning, in: *Proceedings of the AAAI-MAKE*, Vol. 2350, 2019.

[11] M. Ebrahimi, A. Eberhart, F. Bianchi and P. Hitzler, Towards Bridging the Neuro-Symbolic Gap: Deep Deductive Reasoners, *Appl. Intell.* **51**(9) (2021), 6326–6348.

[12] M. Ebrahimi, M.K. Sarker, F. Bianchi et al., Neuro-Symbolic Deductive Reasoning for Cross-Knowledge Graph Entailment, in: *Proceedings of the AAAI-MAKE*, Vol. 2846, 2021. http://ceur-ws.org/Vol-2846/paper8.pdf.

[13] Z. Chen, J. Chen et al., Zero-Shot Visual Question Answering using Knowledge Graph, *CoRR* (2021). https://arxiv.org/abs/2107.05348.

[14] Y. Geng, J. Chen, Z. Ye, Z. Yuan, W. Zhang and H. Chen, Explainable Zero-shot Learning via Attentive Graph Convolutional Network and Knowledge Graphs, *Semantic Web* **12**(5) (2021), 741–765. https://doi.org/10.3233/SW-210435.

[15] J. Chen, Y. Geng et al., Low-Resource Learning with Knowledge Graphs: A Comprehensive Survey, *CoRR* (2021). https://arxiv.org/abs/2112.10006.

[16] P. Hitzler, M. Krötzsch, B. Parsia, P.F. Patel-Schneider and S. Rudolph, OWL 2 Web Ontology Language Primer, W3C Recommendation, World Wide Web Consortium, 2012. http://www.w3.org/TR/owl2-primer/.

[17] B. Motik, P.F. Patel-Schneider, B. Parsia, C. Bock, A. Fokoue, P. Haase, R. Hoekstra, I. Horrocks, A. Ruttenberg, U. Sattler and M. Smith, OWL 2 Web Ontology Language: Structural Specification and Functional-Style Syntax, W3C Recommendation, World Wide Web Consortium, 2012. http://www.w3.org/2007/OWL/draft/owl2-syntax/.

[18] The Semantic Web Wiki, 2001. https://www.w3.org/2001/sw/wiki/Main_Page.

[19] The Semantic Web Stack, 2022. https://en.wikipedia.org/wiki/Semantic_Web_Stack.

[20] A. Breit, L. Waltersdorfer, F.J. Ekaputra, M. Sabou, A. Ekelhart, A. Iana, H. Paulheim, J. Portisch, A. Revenko, A.t. Teije and F. van Harmelen, Combining Machine Learning and Semantic Web: A Systematic Mapping Study, *ACM Computing Surveys* (2023). https://doi.org/10.1145/3586163.

[21] D. Herron, E. Jiménez-Ruiz and T. Weyde, On the Benefits of OWL-based Knowledge Graphs for Neural-Symbolic Systems, in: *NeSy 2023: 17th International Workshop on Neural-Symbolic Learning and Reasoning, July 3–5, La Certosa di Pontignano, Siena, Italy*, CEUR Workshop Proceedings, Vol. 3432, CEUR-WS.org, 2023. https://ceur-ws.org/Vol-3432/paper28.pdf.

[22] T. Berners-Lee, J. Hendler and O. Lassila, The Semantic Web, *Scientific American* **284**(5) (2001), 34–43. https://www.jstor.org/stable/pdf/26059207.pdf.

[23] P. Hitzler, M. Krötzsch and S. Rudolph, *Foundations of Semantic Web Technologies*, 3rd edn, CRC Press, 2010.

[24] J. Lehmann, R. Isele, M. Jakob, A. Jentzsch, D. Kontokostas, P.N. Mendes, S. Hellmann, M. Morsey, P. van Kleef, S. Auer and C. Bizer, DBpedia - A Large-scale, Multilingual Knowledge Base Extracted from Wikipedia., *Semantic Web* **6**(2) (2015), 167–195. http://dx.doi.org/10.3233/SW-140134.

[25] D. Vrandečić and M. Krötzsch, Wikidata: A Free Collaborative Knowledgebase, *Communications of the ACM* **57**(10) (2014), 78–85.

[26] T.P. Tanon, G. Weikum and F.M. Suchanek, YAGO 4: A Reason-able Knowledge Base, in: *The Semantic Web - 17th International Conference, ESWC 2020, Proceedings*, Lecture Notes in Computer Science, Vol. 12123, Springer, 2020, pp. 583–596. https://doi.org/10.1007/978-3-030-49461-2_34.

[27] P.L. Whetzel, N.F. Noy, N.H. Shah, P.R. Alexander, C. Nyulas, T. Tudorache and M.A. Musen, BioPortal: Enhanced Functionality via new Web Services from the National Center for Biomedical Ontology to Access and use Ontologies in Software Applications, *Nucleic Acids Research* **39**(suppl 2) (2011), W541–W545, https://bioportal.bioontology.org/.

[28] R. Jackson, N. Matentzoglu, J.A. Overton, R. Vita, J.P. Balhoff, P.L. Buttigieg, S. Carbon, M. Courtot, A.D. Diehl, D.M. Dooley, W.D. Duncan, N.L. Harris, M.A. Haendel, S.E. Lewis, D.A. Natale, D. Osumi-Sutherland, A. Ruttenberg, L.M. Schriml, B. Smith, C.J. Stoeckert Jr., N.A. Vasilevsky, R.L. Walls, J. Zheng, C.J. Mungall and B. Peters, OBO Foundry in 2021: Operationalizing Open Data Principles to Evaluate Ontologies, *Database* **2021** (2021), http://obofoundry.org/. https://doi.org/10.1093/database/baab069.

[29] Ontotext GraphDB, 2023, A popular RDF store; https://www.ontotext.com/products/graphdb/. https://en.wikipedia.org/wiki/Ontotext_GraphDB.

[30] Y. Nenov, R. Piro, B. Motik, I. Horrocks, Z. Wu and J. Banerjee, RDFox: A Highly-Scalable RDF Store, in: *The Semantic Web - 14th International Conference, ISWC 2015, Proceedings*, Lecture Notes in Computer Science, Vol. 9367, Springer Verlag, 2015, pp. 3–20, https://www.oxfordsemantic.tech/product. https://ora.ox.ac.uk/objects/uuid:2a08b023-77be-431a-a08c-89b47381586a.

[31] B. Glimm, I. Horrocks, B. Motik, G. Stoilos and Z. Wang, HermiT: An OWL 2 Reasoner, *J. Autom. Reason.* **53**(3) (2014), 245–269. https://doi.org/10.1007/s10817-014-9305-1.

[32] E. Sirin, B. Parsia, B.C. Grau, A. Kalyanpur and Y. Katz, Pellet: A Practical OWL-DL Reasoner, *J. Web Semant.* **5**(2) (2007), 51–53. https://doi.org/10.1016/j.websem.2007.03.004.

[33] Y. Kazakov, M. Krötzsch and F. Simančík, Unchain My *EL* Reasoner, in: *Proceedings of the 24th International Workshop on Description Logics (DL'11)*, CEUR Workshop Proceedings, Vol. 745, CEUR-WS.org, 2011.

[34] J. Euzenat and P. Shvaiko, *Ontology Matching, Second Edition*, Springer, 2013. ISBN 978-3-642-38720-3.

[35] D. Herron, E. Jiménez-Ruiz, G. Tarroni and T. Weyde, NeSy4VRD: A Multifaceted Resource for Neurosymbolic AI Research using Knowledge Graphs in Visual Relationship Detection, *CoRR* (2023). http://arxiv.org/abs/2305.13258.

[36] C. Lu, R. Krishna, M. Bernstein and L. Fei-Fei, Visual Relationship Detection with Language Priors, in: *European Conference on Computer Vision*, Springer, 2016, pp. 852–869. https://cs.stanford.edu/people/ranjaykrishna/vrd/.

[37] N. Noy and D. McGuinness, Ontology Development 101: A Guide to Creating Your First Ontology, Technical Report, Knowledge Systems Laboratory, Stanford University, 2001. http://www.ksl.stanford.edu/people/dlm/papers/ontology-tutorial-noy-mcguinness.pdf.

[38] E.F. Kendall and D.L. McGuinness, *Ontology Engineering*, Synthesis Lectures on The Semantic Web: Theory and Technology, Morgan & Claypool Publishers, 2019.

[39] D. Allemang, J. Hendler and F. Gandon, *Semantic Web for the Working Ontologist*, 3rd edn, ACM Books, 2020.

[40] C.M. Keet, *An Introduction to Ontology Engineering*, 'v1.5' edn, 2020. https://people.cs.uct.ac.za/~mkeet/OEbook/.

[41] M.A. Musen, The Protégé Project: A Look Back and a Look Forward, *AI Matters* **1**(4) (2015), 4–12. doi:10.1145/2757001.2757003. https://protege.stanford.edu/.

[42] C. d'Amato, Machine Learning for the Semantic Web: Lessons learnt and next research directions, *Semantic Web* **11**(1) (2020), 195–203. https://doi.org/10.3233/SW-200388.

[43] M. Nickel, K. Murphy, V. Tresp and E. Gabrilovich, A Review of Relational Machine Learning for Knowledge Graphs, *CoRR* (2015). http://arxiv.org/abs/1503.00759.

[44] Y. Dai, S. Wang, N.N. Xiong and W. Guo, A Survey on Knowledge Graph Embedding: Approaches, Applications and Benchmarks, *Electronics* **9**(5) (2020). doi:10.3390/electronics9050750. https://www.mdpi.com/2079-9292/9/5/750.

[45] M. Wang, L. Qiu and X. Wang, A Survey on Knowledge Graph Embeddings for Link Prediction, *Symmetry* **13**(3) (2021). https://doi.org/10.3390/sym13030485.

[46] Z. Chen, Y. Wang, B. Zhao, J. Cheng, X. Zhao and Z. Duan, Knowledge Graph Completion: A Review, *IEEE Access* **8** (2020), 192435–192456. https://doi.org/10.1109/ACCESS.2020.3030076.

[47] A. Rossi, D. Barbosa, D. Firmani, A. Matinata and P. Merialdo, Knowledge Graph Embedding for Link Prediction: A Comparative Analysis, *ACM Trans. Knowl. Discov. Data* **15**(2) (2021), 14:1–14:49. https://doi.org/10.1145/3424672.

[48] J. Chen, P. Hu, E. Jiménez-Ruiz, O.M. Holter, D. Antonyrajah and I. Horrocks, OWL2Vec*: Embedding of OWL Ontologies, *Mach. Learn.* **110**(7) (2021), 1813–1845. https://doi.org/10.1007/s10994-021-05997-6.

[49] U. Kursuncu, M. Gaur and A.P. Sheth, Knowledge Infused Learning (K-IL): Towards Deep Incorporation of Knowledge in Deep Learning, in: *Proceedings of the AAAI 2020 Spring Symposium on Combining Machine Learning and Knowledge Engineering in Practice, AAAI-MAKE 2020*, A. Martin, K. Hinkelmann, H. Fill, A. Gerber, D. Lenat, R. Stolle and F. van Harmelen, eds, CEUR Workshop Proceedings, Vol. 2600, CEUR-WS.org, 2020. https://ceur-ws.org/Vol-2600/paper14.pdf.

[50] R. Wickramarachchi, C.A. Henson and A.P. Sheth, Knowledge-infused Learning for Entity Prediction in Driving Scenes, *Frontiers in Big Data* **4** (2021), 759110. https://doi.org/10.3389/fdata.2021.759110.

[51] R. Brachman and H. Levesque, *Knowledge Representation and Reasoning*, Morgan Kaufman, 2004.

[52] F. Baader, D. Calvanese, D. McGuinness, D. Nardi and P.F. Patel-Schneider (eds), *The Description Logic Handbook*, 2nd edn, Cambridge University Press, 2007.

[53] F. Baader, I. Horrocks, C. Lutz and U. Sattler, *An Introduction to Description Logic*, Cambridge University Press, 2017.

[54] D. Nardi and R.J. Brachman, An Introduction to Description Logics, in: *The Description Logic Handbook: Theory, Implementation, and Applications*, F. Baader, D. Calvanese, D.L. McGuinness, D. Nardi and P.F. Patel-Schneider, eds, Cambridge University Press, 2003, pp. 1–40.

[55] I. Horrocks, O. Kutz and U. Sattler, The Even More Irresistible SROIQ, in: *Proceedings, Tenth International Conference on Principles of Knowledge Representation and Reasoning*, P. Doherty, J. Mylopoulos and C.A. Welty, eds, AAAI Press, 2006, pp. 57–67. http://www.aaai.org/Library/KR/2006/kr06-009.php.

[56] E. Jiménez-Ruiz, Logic-based Support for Ontology Development in Open Environments, PhD thesis, Jaume I University, Spain, 2010. http://hdl.handle.net/10803/10493.

[57] R. Kopparti and T. Weyde, Weight Priors for Learning Identity Relations, in: *Workshop Knowledge Representation & Reasoning Meets Machine Learning at NeurIPS*, 2019, pp. 8–14.

[58] L. Serafini and A.S. d'Avila Garcez, Logic Tensor Networks: Deep Learning and Logical Reasoning from Data and Knowledge, in: *Proceedings of the 11th International Workshop on Neural-Symbolic Learning and Reasoning (NeSy'16)*, CEUR Workshop Proceedings, Vol. 1768, CEUR-WS.org, 2016. http://ceur-ws.org/Vol-1768/NESY16_paper3.pdf.

[59] S. Badreddine, A. d'Avila Garcez, L. Serafini and M. Spranger, Logic Tensor Networks, *Artificial Intelligence* **303** (2022). https://doi.org/10.1016/j.artint.2021.103649.

[60] E. Giunchiglia, M.C. Stoian, S. Khan, F. Cuzzolin and T. Lukasiewicz, ROAD-R: The Autonomous Driving Dataset with Logical Requirements, *Machine Learning* (2023). https://doi.org/10.1007/s10994-023-06322-z.

[61] C. Cornelio, J. Stuehmer, S.X. Hu and T. Hospedales, Learning Where and When to Reason in Neuro-Symbolic Inference, in: *International Conference on Learning Representations*, 2023. https://openreview.net/forum?id=en9V5F8PR-.

[62] I. Donadello and L. Serafini, Compensating Supervision Incompleteness with Prior Knowledge in Semantic Image Interpretation, in: *IJCNN Hungary, July 14-19*, IEEE, 2019, pp. 1–8. https://doi.org/10.1109/IJCNN.2019.8852413.

[63] E. Kharlamov, B.C. Grau, E. Jiménez-Ruiz, S. Lamparter, G. Mehdi, M. Ringsquandl, Y. Nenov, S. Grimm, M. Roshchin and I. Horrocks, Capturing Industrial Information Models with Ontologies and Constraints, in: *The Semantic Web - ISWC 2016 - 15th International Semantic Web Conference, Kobe, Japan, October 17-21, 2016, Proceedings, Part II*, P. Groth, E. Simperl, A.J.G. Gray, M. Sabou, M. Krötzsch, F. Lécué, F. Flöck and Y. Gil, eds, Lecture Notes in Computer Science, Vol. 9982, 2016, pp. 325–343. doi:10.1007/978-3-319-46547-0_30.

[64] A.P. Sheth, K. Roy and M. Gaur, Neurosymbolic AI - Why, What, and How, *CoRR* (2023). https://doi.org/10.48550/arXiv.2305.00813.