

What can knowledge graph alignment gain with Neuro-Symbolic learning approaches?

Pedro Giesteira Cotovio ^{a,*}, Ernesto Jimenez-Ruiz ^{b,c} and Catia Pesquita ^a

^a *LASIGE, Faculdade de Ciências da Universidade de Lisboa, Lisboa, Portugal*

E-mails: pgcotovio@fc.ul.pt, clpesquita@fc.ul.pt

^b *Department of Computer Science City, University of London, London, UK*

E-mail: ernesto.jimenez-ruiz@city.ac.uk

^c *SIRIUS, University of Oslo, Oslo, Norway*

Abstract. Knowledge Graphs (KG) are the backbone of many data-intensive applications since they can represent data coupled with its meaning and context. Aligning KGs across different domains and providers is necessary to afford a fuller and integrated representation. A severe limitation of current KG alignment (KGA) algorithms is that they fail to articulate logical thinking and reasoning with lexical, structural, and semantic data learning. Deep learning models are increasingly popular for KGA inspired by their good performance in other tasks, but they suffer from limitations in explainability, reasoning, and data efficiency. Hybrid neurosymbolic learning models hold the promise of integrating logical and data perspectives to produce high-quality alignments that are explainable and support validation through human-centric approaches. This paper examines the current state of the art in KGA and explores the potential for neurosymbolic integration, highlighting promising research directions for combining these fields.

Keywords: Knowledge Graph Alignment, Neuro-Symbolic models, Explainable Artificial Intelligence

1. Introduction

With the increasing amount of information available, the need for an effective way to integrate and manage this information has become critical. Knowledge graphs (KGs) have emerged as a solution to this problem, providing graph-based data models that formally represent data objects typically using ontologies. These models can capture knowledge in a graph-based abstraction that presents a significant advantage from a practical point of view compared to other data models [1]. Deductive and inductive reasoning can increase the knowledge base through simple logical rules or combined with machine learning methods like embeddings [2]. KGs are, therefore, an effective way of managing, integrating, and extracting knowledge and value from different data sources.

However, in most practical cases, a single KG is often incomplete, and its usefulness, especially in multi-domain settings, depends on the integration with other KGs. The performance of state-of-the-art KG alignment (KGA) algorithms is often only achieved through hand-crafted configuration since heterogeneous knowledge requires context that lexical, structural, and rule-based methods cannot always grasp [1]. The goal of KGA is to automatically identify and merge semantically equivalent entities, relationships, and attributes across multiple KGs [3]. This process requires comparing the entities and relationships in one graph with those in another and determining their similarity. The resulting correspondences can be used to create a unified representation of the knowledge.

*Corresponding author. E-mail: pgcotovio@fc.ul.pt.

Recent subsymbolic approaches, including linguistic and structural models like Attention Networks and Graph Neural Networks, have been employed to address knowledge graph alignment [4–8]. Although these approaches have shown promise, they lack certain advantages of logical models, such as explainability and robust deduction [9]. In the Artificial Intelligence (AI) community, there is a growing recognition that augmenting subsymbolic AI with symbolic properties can effectively address its current limitations [10–12]. This convergence of symbolic and subsymbolic approaches, known as Neurosymbolic AI, holds significant potential for enhancing the capabilities of AI systems. By combining symbolic reasoning, which excels at knowledge representation and logical inference, with the power of subsymbolic methods in handling large-scale data and extracting patterns, Neurosymbolic AI could address many of the current challenges in KGA. Hybrid networks like Logic Tensor Networks (LTNs) [13], which tightly integrate neural learning with logical reasoning, show particular promise in addressing the limitations of current deep learning algorithms.

In this paper, we provide a more detailed overview of the current state-of-the-art in KGA and neurosymbolic integration and explore the challenges and research opportunities that neurosymbolic models offer to the field of KGA.

2. Knowledge Graph Alignment

2.1. Problem Definition

A Knowledge Graph (KG) is a multi-relational graph of structured knowledge resources, commonly expressed by RDF (Resource Description Framework) triples of the form $\langle \text{subject}, \text{predicate}, \text{object} \rangle$, that can be formally structured according to an *OWL* ontology¹ [5].

An *OWL* ontology, represented as O , is a structured representation of entities within a specific domain. In addition, an ontology O includes a set of valid names for the entities present, known as its signature, $Sig(O)$. Formally, an *OWL* ontology can be seen as a collection of axioms that follow the syntactic rules and constraints defined by their underlying description logics [14].

The signature, $Sig(O)$, is structured as the disjoint union of four finite entity sets: (i) N_C , a collection of unary symbols referred to as named concepts, (ii) N_R , a group of binary symbols known as named object properties, (iii) N_D , an assemblage of binary symbols designated as data properties, and (iv) N_I , a compilation of constant symbols identified as named individuals.

When aligning two KGs, the objective is to identify a mapping function, denoted as f , that maps entities from one KG's signature to entities in another while preserving the structural properties of both KGs. Consider two KGs, KG_1 , KG_2 , along with their respective signatures, $Sig(KG_1)$ and $Sig(KG_2)$. A mapping between entities of KG_1 , KG_2 is typically represented as a tuple $\langle e_1, e_2, r, c \rangle$ such that $e_1 \in Sig(KG_1)$ and $e_2 \in Sig(KG_2)$, $r \in \{\sqsubseteq, \supseteq, \equiv\}$ is a semantic relation, and c is a confidence value (normally $c \in [0, 1]$).

With this understanding, we can conceptualize f as a mapping function denoted as $f : Sig(KG_1) \rightarrow Sig(KG_2)$ that associates the entities in KG_1 with the corresponding entities in KG_2 , and produces a set of mappings M , also referred to as an alignment. Finally, the knowledge graph KG_M , which results from the union of KG_1 , KG_2 , and M , is called the aligned knowledge graph.

2.1.1. Consistency Principle Violations

The aligned knowledge graph, represented as KG_M , has the potential to produce axioms not directly derivable from its sources, KG , KG' , or its alignment, M , in isolation. The deductive difference [15] between KG and KG' , denoted $diff_{\Sigma}(KG, KG')$, is defined as the set of axioms with signature $\Sigma = Sig(KG) \cup Sig(KG')$ that do not hold over KG but are valid over KG' .

Evidence of violations in consistency may signal issues with the mappings or discrepancies between the original KGs. These violations invariably culminate in aligned KGs that are either inconsistent or incoherent, or in other words, KGs comprising unsatisfiable concepts.

¹When referring to OWL we mean OWL 2.

The consistency principle [16] stipulates that every named concept within the aligned knowledge graph KG_M should be satisfiable, under the assumption that the union of the input KGs, represented as $KG^0 = KG \cup KG'$, is devoid of unsatisfiable concepts. Accordingly, if the condition $(A \sqsubseteq \perp) \in \text{diff}_\Sigma(KG^0, KG_M)$ holds for any $A \in \Sigma$, the alignment M is deemed violate the consistency principle. Therefore, a mapping $m \in M$ is said to be involved in the cause of an unsatisfiability $A \sqsubseteq \perp$ if there exists a subset of axioms $\mathcal{K} \subseteq KG_M$ such that $\mathcal{K} \not\models A \sqsubseteq \perp$ and $\mathcal{K} \cup \{m\} \models A \sqsubseteq \perp$.

To validate this principle, we can define a function, $\text{consist} : e' \rightarrow \{0, 1\}$, such that if any mapping $f : e' \rightarrow f(e')$ is involved in the cause of unsatisfiability, then the function $\text{consist}(e')$ will output 0; otherwise, it will yield 1.

This approach can be considered as a stringent or "hard" since a single instance of unsatisfiability results in an output of zero. In contrast, a "soft" alternative approach can be defined as $\text{softconsist} : e' \rightarrow [0, 1]$, wherein $\perp_{e'}$ denotes the number of instances of unsatisfiability involving e' , such that $\lim_{\perp_{e'} \rightarrow \infty} \text{softconsist}(e') = 0$. In this case, the output approaches zero, as the occurrences of unsatisfiability tend towards infinity.

One potential heuristic to guide the behaviour of the softconsist function could involve employing the logistic function of the occurrences of unsatisfiability in the following way:

$$\text{softconsist}(e') = \frac{1}{1 + e^{-\perp_{e'}}$$

2.1.2. Unsupervised Setting

The majority of existing KG alignment approaches are fully unsupervised and thus do not require any previously created alignment to learn mapping functions. In practice, within an unsupervised learning setting, the mapping process is typically broken down into two components: a candidate generation function G , which produces a range of possible mappings between entities, and a filtering function F designed to eliminate improbable or irrelevant mappings from the generated candidates.

The candidate generation function G uses different types of information, such as labels, descriptions, types, and relationships with other objects, to pinpoint a group of entities in KG_2 that may correspond to a specific entity in KG_1 . The candidate generation function can be depicted as follows:

$$G(e_1, KG_1, KG_2) = M \quad \forall e_1 \in \text{Sig}(KG_1)$$

Next, the filtering function F evaluates the quality of the generated candidate matches. It achieves this by leveraging the mapping confidence and employing cardinality rules or thresholds to establish a criterion for accepting or rejecting a candidate mapping. Furthermore, filtering functions may also "repair" alignments by mitigating logical inconsistencies between mappings and signatures, thereby enhancing their coherence. In essence, these functions eliminate mappings that exhibit low quality, violate cardinality constraints, or introduce incoherence. We can represent this process symbolically as follows:

$$F(M, KG_1, KG_2) = K^M$$

The complete mapping function f is derived by combining the candidate generation function(s) G with the filtering function(s) F , leading to a final set of high-quality candidate matches. Consequently, the global correspondence function can be expressed as:

$$f(KG_1, KG_2) = F(G(e_1, KG_1, KG_2)) \quad \forall e_1 \in \text{Sig}(KG_1)$$

In this setting, the objective is to maximize the number of mappings while preserving the coherence of KG_M . Given that maximizing the number of mappings and maintaining coherence are two contradicting strategies, they can either be optimized under an assumption of correctness or uncertainty.

An assumption of correctness implies that flawless logic governs the set of axioms constituting KG_1 and KG_2 . Consequently, any incoherent mapping cannot be accurate. Under this assumption, the optimization problem can be formulated as follows:

$$\arg \max_f \sum_{e_1 \in \text{Sig}(KG_1)} \sum_{f(e_1) \subset \text{Sig}(KG_2) \vee |f(e_1)| \leq t} \text{consist}(e_1) * \text{consist}(f(e_1)) * c(e_1, f(e_1))$$

where $f(e_1)$ is the set of entities mapped to e_1 under f , c is a function that returns the confidence value of a given mapping, and t is a cardinality threshold. Therefore, the goal is to maximize the mapping confidence for a given cardinality of coherent mappings.

On the other hand, under an assumption of uncertainty, axioms might be incorrect or incomplete. By viewing logical incoherence as a result of incorrect or incomplete information, we can conceive of a heuristic that aims to minimize logical incoherence while maximizing the confidence value. This heuristic should, in practice, deal with the trade-off between relying on explicit structural information (axioms) and relying on all other sources of information (e.g., lexical/semantic, internal/external). The simplest form of this heuristic, denoted as \mathcal{H} , would involve a threshold θ beyond which the confidence value surpasses logical incoherence. Under this setting, the optimization problem can be formulated as such:

$$\arg \max_f \sum_{e_1 \in \text{Sig}(KG_1)} \sum_{f(e_1) \subset \text{Sig}(KG_2) \vee |f(e_1)| \leq t} \begin{cases} c(e_1, f(e_1)) & c(e_1, f(e_1)) \geq \theta \\ \text{consist}(e_1, f(e_1)) * c(e_1, f(e_1)) & c(e_1, f(e_1)) < \theta \end{cases}$$

An alternative perspective on \mathcal{H} involves adopting a "soft" approach towards addressing unsatisfiability. In this regard, the previously defined *softconsist* function can be employed to formulate the optimization problem as follows:

$$\arg \max_f \sum_{e_1 \in \text{Sig}(KG_1)} \sum_{f(e_1) \subset \text{Sig}(KG_2) \vee |f(e_1)| \leq t} \text{softconsist}(e_1) * \text{softconsist}(f(e_1)) * c(e_1, f(e_1))$$

Within this setup, a mapping's confidence will be subject to a penalty if its constituent entities are involved in mappings that are unsatisfiable, and the penalty can be proportional to the number of unsatisfiability cases it is related to.

2.1.3. Supervised Setting

Supervised learning approaches can be applied to KG alignment when a partial alignment is available, typically manually produced and deemed correct (i.e., a reference alignment). In this setting, the solution may still be divided into a generator and a filter, particularly when combining supervised and unsupervised algorithms. For instance, we could use an unsupervised generator with a supervised filter. However, this division is not strictly necessary. It is entirely possible to attempt to learn the problem holistically.

Let us assume there is a reference alignment symbolized by a function ref , which calculates the truth value for a given relationship between two entities, such that $ref(e_1, e_2, r) \in \{0, 1\}$ and where $r \in \{\sqsubseteq, \supseteq, \equiv\}$ is a semantic relation. In this setting, the correspondence function f can be optimized problem to maximize ref . This optimization problem can be expressed as follows:

$$\arg \max_f \sum_{e_1 \in \text{Sig}(KG_1)} \sum_{f(e_1) \subset \text{Sig}(KG_2) \vee |f(e_1)| \leq t} ref(e_1, f(e_1))$$

The goal is to find the correspondence function f that maximizes the sum of similarities between the corresponding entities in the two KGs.

2.2. Overview/State of the Art

We follow Euzenat & Shvaiko's [17] classification of ontology matching techniques into eight specific categories according to the type of input they take. Content-based matching techniques place their focus on the internal information obtained from the ontologies being matched, while context-based techniques focus on external information.

Context-based approaches comprise two categories defined by their use of either formal or informal resources.

Formal resource-based approaches use formal resources such as previous alignments and other ontologies to either discover new candidate mappings or filter out unlikely candidates. A notable illustration of this approach is put forth by Scharffe et al.[18]. Traditional matching systems, such as AML [19], and several others likewise harness reference alignments and ontological resources within their frameworks. This category also encompasses the majority of supervised or semi-supervised methods, such as VeeAlign [20], as they draw upon reference alignments for training purposes.

Informal resource-based approaches also rely on external resources, but instead of formal resources like ontologies, they utilize informal resources such as annotated relationships between entities or text. A prevalent method involves employing a thesaurus such as thesaurus.com or subject-specific equivalents. For instance, LogMap [21] leverages the UMLS lexicon to expand synonyms. This category also encompasses language models like BERTMap [4], which undergo pre-training on large amounts of external textual content.

Content-based approaches can be categorized into six distinct groups based on their usage of KG content. These methodologies encompass leveraging linguistics across the KG, capitalizing on the structural attributes of the KG, directing attention towards specific instances, and harnessing the semantic richness embedded within the KGs.

String-based approaches involve gauging string similarity within entity names and descriptions to generate plausible candidate mappings. These methods, often referred to as lexical matchers, are prominently featured in classical matching systems like AML [19], and LogMap [21]. A variety of string distance metrics, encompassing Jaccard, Euclidean, Levenshtein, Jaro-Winkler, Jensen-Shannon, and others [22], can be employed by these matchers. String-based techniques are primarily geared towards identifying candidate mappings rather than refining them, as they inherently carry a level of uncertainty. This lack of certainty renders them better suited for assessing the likelihood of a mapping. To discover potential mappings, these techniques can be applied both within the ontology itself and across external resources.

Language-based approaches, unlike string-based approaches, do not treat names or descriptions as mere strings. Instead, they consider them as partial words or sentences in a language by leveraging natural language processing techniques. Prominent examples of these approaches are those that rely on large language models pre-trained on extensive corpora, such as BERTMap [4].

Constraint-based approaches compare the constraints applied to the definition of entities, including constraints on cardinality, types, attributes, or keys. The underlying assumption is that equivalent entities will exhibit similar constraints. Typically, these techniques are used in conjunction with others, as similar constraints on entities are often a common characteristic. Constraint-based approaches are commonly employed as filters to exclude mappings that violate specific constraints, such as cardinality. Examples of such filters can be found in most matching systems [23].

Graph-based approaches leverage the inherent graph structure of KGs and ontologies to generate mappings. State-of-the-art approaches in this category often rely on Graph Neural Networks (GNNs). Examples of such approaches include GCN-align [8], MEDTO [6], and AliNet [24].

Instance-based approaches assess the similarity between classes based on the instances associated with them. The underlying intuition is that similar classes will have similar instances. Several data analysis techniques can be employed to compare instances as proposed by Isaac et al.[25]. However, a major challenge with these approaches is the existence of sparse instances for comparison in some cases.

Model-based approaches leverage the semantic interpretation of paired ontologies by employing formal reasoning techniques, such as description logics reasoning. Due to their precise, albeit computationally intensive, nature, these techniques are primarily used as "repair" algorithms. Their main objective is to filter out logically incoherent mappings from a pre-selected set of candidate mappings. Several works have been published in this area, such as: Solimando et al.[16], Santos et al.[26], Jimenez et al.[27, 28], Shvaiko et al.[29], and Meilicke et al.[30].

1 Until now, we have discussed individual matching techniques. However, state-of-the-art ontology alignment systems typically consist of multiple matching techniques working together. Evaluating the performance of these systems is commonly done through the Ontology Alignment Evaluation Initiative (OAEI) [31, 32]. OAEI is an annual competition that includes various tracks designed to evaluate different tasks and components related to ontology alignment.

2 While many systems participate in the OAEI tracks, two systems have consistently demonstrated excellent performance over the years: AML [19] and its successor, Matcha [33], as well as LogMap [21]. Despite the competition from other systems, AML/Matcha and LogMap have maintained their superiority across most tracks for more than a decade. As a result, these systems have become the go-to choices for ontology alignment, and any neurosymbolic approaches aiming to make an impact in this field will need to surpass their performance.

3 Of particular significance for the neurosymbolic field is the new Bio-ML track [34]. This track is the first to emphasize machine learning-based systems for ontology alignment, making it a crucial arena for neurosymbolic approaches. In the inaugural edition of this track, Matcha-DL [33], a combination of Matcha (a rule-based system) with a simple neural network and BERTMap [4] achieved the best results. These two systems currently serve as the baseline for neurosymbolic approaches to strive for and surpass in performance.

3. Neuro-Symbolic Methods

4 Neurosymbolic AI, also known as Neuro-symbolic computing, is a subfield of artificial intelligence (AI) that envisions integrating symbolic and subsymbolic approaches to this field. While these approaches differ significantly at their extremes (*e.g.*, expert systems vs. multi-layer perceptron), advancements in both domains are slowly bridging the longstanding gap. Although the AI field has, in recent years, predominantly emphasized subsymbolic approaches, such as deep learning techniques, there is a growing consensus that augmenting subsymbolic AI with certain properties of classical symbolic systems could address its current limitations [10–12].

5 Symbolic approaches are those that leverage high-level or explicit knowledge representations, such as those expressed through formal logic and other formal languages, along with symbol manipulation, to accomplish their objectives. During the early stages of AI, classical forms of these approaches gained significant popularity [35]. By employing a formal knowledge source, possess the capability to generate new knowledge and validate assumptions through logical deduction precisely. Since they operate based on manipulations of symbols that are understandable to humans, their deductions can be easily verified even by non-experts. However, symbolic approaches face challenges due to their reliance on human experts to encode knowledge into formal rules, especially when these experts need to represent complex and predominantly continuous domains as discrete rules. This task is often exceedingly difficult, if not impossible. In fact, as early as 1950, Turing expressed his scepticism about the feasibility of representing the world in such a manner in his seminal paper 'Computing Machinery and Intelligence' [36]. Due to these limitations, pure symbolic approaches such as the CYC project [37] have largely fallen out of favour, giving rise to the prominence of subsymbolic approaches.

6 Subsymbolic systems did not garner significant attention from the artificial intelligence community until the emergence of deep learning and big data. This lack of attention primarily stemmed from limitations in computing power and the availability of adequate data to effectively explore these approaches. It was not until the emergence of crucial datasets like WordNet [38] and ImageNet [39] that more robust exploration became feasible. In contrast to symbolic techniques, deep learning focuses on learning through induction from informal data sources, known as datasets. Deep learning models approximate unknown target functions by learning from numerous instances of representations of entities. They can then resort to these approximated functions to make predictions on unseen instances that fall within the learned distribution. When the training dataset is sufficiently representative of the domain being explored, these techniques can achieve exceptional performance [40].

7 The advantage of deep learning models lies in their ability to automatically extract crucial features from raw data, eliminating the need to construct a comprehensive model of the world to solve a given problem. They have achieved remarkable results in various domains such as natural language processing [41, 42], computer vision [43–45], game playing [46, 47], and multimodal learning [48]. This naturally led to a shift towards subsymbolic approaches that has gained significant traction in recent years, acknowledging the need for more flexible and adaptable AI systems

1 capable of handling complex and unstructured information. Similar to the perception surrounding symbolic systems, 1
2 where it was believed that any problem could be solved with enough rules, some initially speculated that logical 2
3 reasoning and consciousness would emerge with sufficient data and computing power. 3

4 However, as deep learning models approach their limits in terms of scale, they face challenges in simple com- 4
5 mon-sense reasoning tasks that can be effortlessly solved by infants. Questions are now being raised about whether 5
6 this approach alone provides the complete answer [10]. Deep learning models also encounter issues with data in- 6
7 efficiency, generalisation beyond the training distribution, fairness concerns, and, perhaps most notably, a lack of 7
8 human comprehensibility [11]. 8

9 The integration of subsymbolic and symbolic approaches initially encountered challenges due to experts in both 9
10 fields who disregarded the potential of each other's techniques. Some individuals continue to advocate exclusively 10
11 for pure subsymbolic approaches, dismissing any symbolic involvement. Fortunately, within the field of AI, an 11
12 increasing number of researchers now recognize the potential of integrating these approaches to address their re- 12
13 spective deficiencies [10–12]. 13

14 One lens through which the differences between symbolic and subsymbolic techniques can be viewed, particularly 14
15 in a neuroscience context, is the framework of System I and System II proposed by Kahneman in his book "Thinking, 15
16 Fast and Slow" [49]. Kahneman suggests that the human brain employs two types of thinking: a fast, unconscious, 16
17 automatic, and effortless thinking attributed to System I, which responds to stimuli, and a slow, logical, and effortful 17
18 thinking attributed to System II, which is utilized to solve complex problems. Notably, this theory aligns with 18
19 the distinctions between subsymbolic and symbolic techniques, indicating that achieving human-like intelligence 19
20 may involve integrating these methodologies. Artificial neural networks can be viewed as abstractions of the inner 20
21 workings of the human brain, while formal logic systems serve as abstractions of the perceptual and reasoning 21
22 processes of our conscious mind [10]. 22

23 By recognizing the value of both subsymbolic and symbolic approaches and understanding the potential benefits 23
24 of integration, the field of Neurosymbolic AI is paving the way for advancing artificial intelligence toward more 24
25 comprehensive and human-like intelligence. 25

26 3.1. Integrating Symbolic and Subsymbolic Approaches 26

27 On the one hand, deep learning systems demonstrate exceptional performance in training on raw data, showcasing 27
28 robustness against errors and outliers. However, they encounter challenges when it comes to incorporating expert 28
29 knowledge, comprehending logical reasoning, and achieving interpretability. On the other hand, symbolic systems 29
30 are easily interpretable and capable of leveraging expert knowledge, but they are susceptible to errors and outliers in 30
31 the knowledge source and pose difficulties in the training process. This naturally raises the question of how we can 31
32 effectively integrate these two abstractions. In [11] Hitzler et al. present intriguing ideas and insights on approaching 32
33 this integration. 33

34 One suggested proposition is the use of pure subsymbolic models to solve classical symbolic problems, such as 34
35 algebraic problems, logical deduction and abduction, rule learning, term rewriting, or planning. These challenging 35
36 problems, even for complex symbolic systems, provide an opportunity to test the limits of current deep learning 36
37 approaches. 37

38 Another avenue proposed is the incorporation of symbolic knowledge bases, such as KGs and other forms of 38
39 explicit metadata, to enhance subsymbolic systems. Neurosymbolic models could leverage KGs in combination 39
40 with deep learning models to improve zero-shot learning, as shown in the work by Lee et al.[50]. Other approaches, 40
41 such as K-Bert [51] and the model proposed by Ma et al.[52], fuse KGs with large language models to provide 41
42 domain context, thereby enhancing coherence and consistency within a given domain. 42
43 43

44 A further suggested path involves the use of background knowledge, such as KGs, to explain the input-output 44
45 behaviour of black-box neural networks. Numerous works have delved into this area, exploiting KGs or other forms 45
46 of expert knowledge. Examples of such research include the works by Sarker et al.[53], Geng et al.[54, 55], and 46
47 Dalal et al.[56]. 47
48 48

49 Lastly, a key research avenue in the field of Neurosymbolic AI involves the coupling of deep learning models with 49
50 symbolic components. This approach aligns closely with the objectives of Neurosymbolic AI, aiming to integrate 50
51 subsymbolic and symbolic models. To provide a framework for understanding the different levels of integration 51

1 between symbolic and subsymbolic methodologies, Kautz [12] proposed a taxonomy for neurosymbolic systems. 1
 2 This taxonomy is further explored by Garcez and Lamb [10]. 2

3 Within this taxonomy, Type I integration, referred to as **Symbolic Neurosymbolic** integration, represents the 3
 4 use of standard deep learning techniques. The argument for considering this as neurosymbolic lies in the fact that 4
 5 deep learning techniques often begin with symbols (*e.g.*, natural language processing), encode them into a contin- 5
 6 uous space, process this vector space using a neural engine, and then decode the result back into symbols, thereby 6
 7 demonstrating Symbolic Neurosymbolic integration. 7

8 A Type II integration, known as **Symbolic[Neuro]** integration, involves loosely coupled hybrid systems that use 8
 9 a core neural model within a symbolic problem solver. This type of integration is exemplified by systems like 9
 10 DeepMind’s AlphaGo [46]. 10

11 Type III neurosymbolic integration, referred to as **NeuroSymbolic** integration, entails a neural model and a 11
 12 symbolic reasoning system working together to achieve a complex goal by handling complementary tasks. In these 12
 13 hybrid systems, a neural model receives non-symbolic input and produces a symbolic data structure, which is then 13
 14 processed by a symbolic reasoner to verify complex statements. The Neuro Symbolic Concept Learner [57] serves 14
 15 as a prime example of this type of integration. 15

16 A Type IV neural-symbolic system (**Neuro:Symbolic** → **Neuro** integration) involves a standard neural network 16
 17 with symbolic constraints on training, achieved either by training on symbolic datasets or by implementing con- 17
 18 straints on the learning process itself. Arabshahi et al.[58] propose an approach to learning and reasoning over 18
 19 mathematical programs, which is then extended in [59] to an architecture capable of extrapolating to more challeng- 19
 20 ing symbolic reasoning problems. Garcez and Lamb [10] also argue that systems with learning constraints, such as 20
 21 those explored in Logical Neural Networks (LNN) [60], should fall under this type of integration. LNNs establish a 21
 22 one-to-one correspondence between neurons and logical formulas, where each neuron’s activation function can be 22
 23 constrained to approximate a logical operation. The main challenges associated with these networks lie in mapping 23
 24 one-to-one networks for large-scale problems and integrating them with standard deep learning architectures. 24

25 Type V systems (**Neuro_{Symbolic}** integration) are the most tightly coupled without achieving full integration. 25
 26 In these systems, symbolic rules are mapped onto embeddings that serve as regularizers for the network’s loss 26
 27 functions. Current examples of this integration include tensorization methods like LTNs [13] and Tensor Product 27
 28 Representations [61]. LTNs combine deep neural networks with first-order fuzzy logic by embedding elements of 28
 29 a formal language in a vector space. Unlike LNNs, LTNs enable full integration with other neural architectures 29
 30 for solving complex problems. LTNs are trained on facts and rules and can infer novel logical statements from 30
 31 data. Recent versions of LTNs have demonstrated excellent performance across various AI tasks, such as multi- 31
 32 label classification, relational learning, data clustering, semi-supervised learning, regression, embedding learning, 32
 33 and query answering [13]. Wagner and Garcez [62] also show that LTNs can significantly improve fairness metrics 33
 34 while maintaining state-of-the-art performance by incorporating fairness constraints into the neural engine. 34

35 Lastly, Type VI integration, known as **Neuro[Symbolic]** integration, refers to fully integrated systems where a 35
 36 symbolic engine is embedded into a neural network, representing a comprehensive integration of System 1 and Sys- 36
 37 tem 2. Kautz suggests that such architectures should allow for combinatorial reasoning, potentially by leveraging 37
 38 attention-based approaches. While fully integrated systems do not currently exist, Lamb et al.[63] argue that inte- 38
 39 grating graph neural networks with Neurosymbolic AI and incorporating attention mechanisms may be a promising 39
 40 direction towards achieving this type of system. 40

41 4. Challenges and Opportunities 41

42 43 44 45 Table 1 outlines the correlation between existing challenges within KGA and emerging Neuro-Symbolic AI ap- 45
 46 proaches. These approaches are classified based on their primary functions: mapping or repair. 46

47 4.1. Tackling Large Size 47

48 49 50 One of the current challenges facing the KGA field is scalability. As KGs continue to grow in size, symbolic 50
 51 systems that often rely on complex sets of formal rules and require meticulous fine-tuning became increasingly 51

Table 1
Table matching KGA challenges to Neuro-Symbolic approaches that offer promising potential solutions.

Challenges	Mapping	Repair
Tackling Large Size	Large-Language Models (BERTMap)	Integration of mapping and repair (Adversarial Training)
Handling Heterogeneity	Formal knowledge injection (K-BERT, Instruction Training)	Subsymbolic approximators (MLP, GCN)
Achieving Correctness	Formal knowledge injection (K-BERT, Instruction Training)	Soft Repair (LNN, LTN) Integration of mapping and repair (Adversarial Training)
Minimizing Human Effort	Data efficient approaches (K-Bert, Instruction Training, LTN)	Human-in-the-loop approaches (RL)
Explaining Decisions	Post-hoc explanations (concept induction, path-based explanations) Tightly coupled NSai (LNN, LTN)	

harder to scale to larger problems. Furthermore, when applied to different KGs, these systems essentially start from scratch, resulting in impractically long run times for very large KGs and rendering them unsuitable for real-time applications.

Prominent examples like AML [19] and LogMap [21] highlight the challenges faced by pure symbolic systems. While they continue to achieve state-of-the-art performance in specific domains and have shown effectiveness over the years, they lack some of the advantages offered by more data-driven approaches.

KG alignment systems that explore **Large Language Models** (LLM), like BERTMap [4], undergo initial training on extensive corpus and subsequent fine-tuning using reference alignments for various KGs. Although the initial training phase may require considerable time, the subsequent inference process benefits from the efficiency of deep learning approaches, allowing them to outperform symbolic approaches in terms of speed when matching new, unseen KGs. This natural scalability can lead to improved performance on larger and more complex KGs while maintaining the ability to rapidly infer mappings. Therefore, LLMs offer advantages for scaling KGA system’s mapping process, as the inference time experiences only negligible changes as the KGs grow in size and complexity.

Further efficiency can be achieved by tackling a common time-consuming feature of contemporary KGA systems: the separation of mapping and repair processes. Given the fundamental requirement of most repair algorithms, outlined in section 2.1, for a complete alignment, state-of-the-art KGA systems divide the task into distinct mapping and repair phases. This approach introduces the possibility of redundant efforts between mapping and repair algorithms.

To tackle this issue, we propose an approach for the **integration of mapping and repair algorithms within an adversarial training framework**, such as a simple generator-discriminator architecture. This architecture would allow the learning and execution of mapping and repair functions in tandem, streamlining the KGA process to a single inference task across KGs. While optimal performance might still require two inference runs – one for source-to-target and another for target-to-source – this approach holds promise for more efficient scaling on larger KGs compared to existing state-of-the-art systems.

A similar architecture was already introduced by [7], showing promising results. Notably, this approach employs relatively straightforward generator and discriminator networks in contrast to the complex mapping and repair algorithms used by competing systems. Exploring whether this approach’s efficiency endures when applied to more complex KGA tasks would be an interesting avenue for further research.

4.2. Handling Heterogeneity

Addressing heterogeneity across diverse domains is of crucial importance in the context of KGA challenges. Real-world KGs exhibit considerable disparity stemming from a multitude of origins, domains, and even languages. Notably, the most intriguing KGA tasks involve the alignment of exceptionally dissimilar KGs. In order to effectively manage such diversity, a substantial degree of generalisation ability becomes imperative.

Subsymbolic algorithms can effectively leverage patterns inherent in raw data to produce high-quality mappings, while symbolic algorithms heavily rely on expert-defined rules for generating mappings. This has led to the emergence of LLM-based methods, such as BERTMap [4], which are capable of effectively generalising across domains.

1 Their scale enables them to store and leverage a significant amount of information within their networks. This, in
2 theory, should mean that deep learning models have the potential to continuously enhance their generalisation power
3 as more data becomes available.

4 However, the very characteristics that contribute to the success of LLMs can also be limitations. These models
5 heavily rely on large corpora and numerous reference mappings to perform well, creating requirements that are not
6 always feasible to meet. Additionally, training these models effectively demands substantial computational power.
7 Much of the knowledge embedded in the structure of KGs, which could potentially be inferred, must instead be
8 induced from thousands of examples, resulting in an overall data-inefficient model.

9 Neurosymbolic integration offers a promising pathway for achieving efficient data usage. **Formal knowledge**
10 **injection** models like K-Bert [51] demonstrate comparable data induction capabilities to pure symbolic models
11 while incorporating symbolic knowledge bases, resulting in improved performance when faced with the out-of-
12 distribution challenges associated with smaller datasets. This integration reduces the reliance on raw data, making
13 it a valuable approach for KGA problems.

14 Another emerging avenue for leveraging symbolic knowledge within LLM frameworks involves zero-shot learn-
15 ing via instruction-based training. This innovative approach has gained prominence more recently, facilitated by
16 the subsequent wave of LLMs that followed in the footsteps of ChatGPT. It entails the refinement of general lan-
17 guage models like GPT-3 [42] for specific tasks, achieved through instructions and illustrative examples provided
18 as textual context. This technique circumvents direct alterations to the model’s parameters through supervised or
19 semi-supervised learning, thereby enabling the democratization of LLM-based research.

20 In the context of KGA tasks, this novel fine-tuning methodology could be effectively applied by incorporating
21 the KG’s structural information into the mapping process, such as an entity’s subclasses. OntoLAMA [64] explores
22 a similar approach for subsumption tasks, yielding promising results. While the adaptation of such a technique to
23 complex matching and equivalence tasks might not follow a straightforward trajectory, we hold the belief that it
24 presents a promising avenue for research, warranting further exploration.

25 In the domain of repair, acquiring the skill to correct these mappings could be facilitated by the use of **sub-**
26 **symbolic approximators**, trained on incomplete reference mappings. Notably, this approach has had encouraging
27 results, as showcased by the MatchaDL system [19]. In this system, a neural network functions as a heuristic, learn-
28 ing to adjust the weighting of scores generated by symbolic algorithms. Further advancements in this domain could
29 explore the potential of GNNs for learning these weights, considering the structural positions of entities within their
30 respective graphs.

31 4.3. Achieving Correctness 32

33
34 Another major challenge of KGA tasks is the requirement for models that can effectively reason over and interpret
35 KGs to achieve correct alignments. Meaning that the mappings between entities not only need to be semantically
36 accurate but also logically satisfiable when reasoning over KGs.

37 KG alignment symbolic systems commonly employ separate algorithms for scoring mappings, followed by a
38 heuristic to consolidate these scores and generate candidates. While these techniques perform well individually,
39 they lack integration, as information sharing among algorithms is limited. The final heuristic used for consolidation
40 becomes crucial while often relying on domain experts’ rules, which may not be optimal and can vary in perfor-
41 mance across different domains.

42 On the other hand, subsymbolic systems employ complex networks that can combine multiple architectures to
43 explore several properties of the KG when generating candidate mappings. GNNs can capture structural information
44 similarly to symbolic algorithms, while LLMs, as discussed earlier, can learn from large corpora and make better
45 semantic interpretations. Recent subsymbolic works that explore this approach to mapping, such as [4, 5, 7], have
46 achieved state-of-the-art performance in certain tasks. However, subsymbolic models have limitations, such as their
47 difficulty in leveraging logical reasoning and semantics in commonsense reasoning tasks. As such, without incor-
48 porating formal forms of external knowledge, their potential to fully replace symbolic approaches in mapping tasks
49 is limited.

50 Neurosymbolic integration can play a crucial role in harnessing the advantages of both symbolic and subsymbolic
51 systems. As we have shown previously, approaches like instruction training and context incorporation into the

attention mechanism of language models (as exemplified by K-Bert [51]) can be effectively employed to assimilate structured knowledge from aligned KGs or external references into the mapping procedure. Moreover, we have also demonstrated that subsymbolic approximators exhibit the capacity to rectify discrepancies in a range of symbolic and subsymbolic matching algorithms.

Nevertheless, this type of loosely coupled neurosymbolic approach might not be enough to achieve correctness. Real-world KGs often contain errors, including incomplete information or incompatibilities among KGs.

Let's consider the following example [65], from the FMA and NCI ontologies, henceforward denoted as O_{FMA} and O_{NCI} . In O_{FMA} , "Lymphokine" is expressed as $Lymphokine \sqsubseteq Protein$, while in O_{NCI} , an homologous "Therapeutic_Lymphokine" is expressed as $Therapeutic_Lymphokine \sqsubseteq Pharmacologic_Substance$. Additionally, a disjointness axiom $\alpha : Protein \sqcap Pharmacologic_Substance \sqsubseteq \perp$ is defined in O_{NCI} . Current algorithms would typically generate two candidate mappings: $\langle Protein \in O_{FMA} \equiv Protein \in O_{NCI} \rangle$ and $\langle Lymphokine \in O_{FMA} \equiv Therapeutic_Lymphokine \in O_{NCI} \rangle$. However, state-of-the-art logical reasoners (repair algorithms) would discard one of these mappings due to the logical inconsistency derived from α . In reality, both mappings are correct, and the problem lies with either α , or the compatibility between the ontologies.

Tightly coupled neurosymbolic systems, such as LNNs [60] or LTNs [13], could offer a promising approach to handle logical satisfiability more flexibly. These systems integrate adjustable logical constraints into their learning procedure, allowing the learning of an acceptable level of logical unsatisfiability. Under the idea of soft and hard constraints in the repair process explored in section 2.1, these approaches could present a promising heuristic to learn soft constraints in the repair process, effectively producing a **soft repair**

Currently, the practical applications of LNNs [60] are limited due to constraints in current implementations. However, they present intriguing research avenues for KGA by mapping KG information to the network architecture and incorporating adjustable constraints on logical behaviour. While LNNs heavily rely on receiving formal knowledge as input, which may pose challenges in other applications, it aligns well with KGA, where the inputs are predominantly symbolic data structures. Nevertheless, scalability issues related to network mapping remain a concern for the field. Overcoming implementation challenges should be prioritized in future research into these models, as it can pave the way for more effective and efficient neuro-symbolic models.

The current version of LTNs [13] appears more feasible for immediate application in KGA problems. Due to their tighter integration with subsymbolic models, they can be more easily combined with state-of-the-art subsymbolic systems. A fundamental research direction is to evaluate their performance in large-scale settings and assess their ability to reduce the amount of raw data and computational power required to achieve state-of-the-art results.

An interesting research direction for tightly coupled neurosymbolic systems would be to explore combining the advantages of clearly defining adjustable logical constraints, such as in LNNs, with the practical integration found in tensorization-based methods like LTNs.

An alternative avenue to address the challenge of ensuring correctness involves revisiting the integration between mapping and repair functions. The current dilemma of achieving correctness stems from the segmentation of mapping generation and logical filtering tasks. This division, while commonly practised, carries inherent limitations as it assumes the correctness of KGs. This assumption can undermine the quality of alignment due to its overly restrictive nature. Therefore the adoption of an adversarial training framework is, yet again, a potentially promising approach. This idea becomes particularly compelling when considering the use of a soft repair heuristic, such as an LNN [60] or an LTN [13], as the discriminator. The interplay between mapping and repair networks under this framework could potentially yield significant advancements in addressing the problem of correctness in alignment tasks.

4.4. Minimizing Human Effort

The challenge of minimizing human effort within the domain of KGA is intricately linked to the notion of streamlining systems to require as little human intervention as possible. Present-day state-of-the-art systems still rely on human involvement during both the mapping and repair phases.

In the context of the mapping process, human intervention primarily centres on dataset curation and manual calibration. This rings especially true for symbolic systems, where formal rules require meticulous curation by domain experts and heuristics demand manual tuning. As a result, these systems demand significant human effort to operate and maintain, rendering them challenging for non-experts to use. Despite this, contemporary machine

learning systems have also failed to provide significant improvement in terms of human effort. This paper stresses that these systems often suffer from severe sample inefficiency, demanding copious examples even for elementary tasks. While pre-training has mitigated the need for extensive curated samples in large models such as LLMs, the inefficiency of fine-tuning remains notable, especially considering the large amount of effort required to generate reference alignments for model training.

As discussed earlier, **data efficient approaches** emphasizing sample efficiency and tapping into formal knowledge within subsymbolic models, like zero-shot learning through instruction training or approaches like that proposed in K-Bert [51], present potential avenues for alleviating human effort in the mapping phase. These approaches can potentially reduce the size of the reference alignments and external curated data needed to achieve state-of-the-art results.

Shifting the focus to the repair phase, human intervention traditionally occurs toward the end of the alignment process. Here, certain systems flag mappings with low confidence for validation by human experts. While sometimes indispensable, this verification process introduces hurdles when striving to seamlessly integrate these algorithms into plug-and-play applications. In response, we advocate for the exploration of alternative interactive methodologies. One intriguing avenue is **human-in-the-loop reinforcement learning (HLRL)** [66–68]. HLRL leverages human feedback to refine reward functions, demonstrating success, as seen in the case of ChatGPT’s chat functionalities, which were trained using a similar approach based on the GPT-3 base generator model [42]. Within the domain of KGA, this feedback-driven function could be integrated into an inclusive KGA system, either as part of a discriminator in an adversarial training setup or as a reward in a reinforcement learning framework. Regardless of the specific integration strategy, learning a human preference function empowers the automation of expert feedback, eventually leading to the realization of a fully automated system.

4.5. Explaining Decisions

Explainability is another significant challenge in KGA tasks, as automated algorithms often produce non-transparent results, making it difficult for end-users to understand the alignment decisions or identify any inconsistencies or biases in the results. This limits end-users’ trust in the results of KGA.

From the outset of deep learning, it became clear that explaining the results of these black box models would be a challenge. Practical applications of these algorithms for decision-making require some comprehension of the reasoning behind the models, but in most cases, understanding the decision-making of deep learning models is a challenging task even for experts. This task becomes almost impossible when we consider large language models with billions of activation parameters. In the last few years, some approaches have risen to prominence to analyze the activation patterns of neural networks [69, 70], but this task remains very complex, and human interpretable explanations are only attainable in some select cases.

Although symbolic systems inherently offer more interpretability compared to subsymbolic approaches, they can still be challenging for non-experts to grasp if decision explanations are not readily generated.

This challenge has not gone unnoticed by the semantic web community, which has proposed strategies to tackle this challenge [71]. A particularly intriguing approach, especially for explaining deep learning models, involves **Post-Hoc Explanations**. This methodology leverages external formal knowledge to shed light on the decisions made by black box models. Within post-hoc explanations, several approaches were already explored. Particularly pertinent to KGA are Concept Induction [56, 72] and Path-based explanations [73, 74].

Concept Induction employs a symbolic reasoner to attach formal knowledge to neural activations. This approach facilitates comprehension of subsymbolic constituents that constitute a deep learning model. This aligns closely with the KGA domain, where formal knowledge is already abundant and can be seamlessly augmented.

Path-based explanations focus instead on generating logical pathways or "reasoning lines" across ontologies to formulate plausible rationales for a given decision. This proves especially valuable for KGA since ontologies already exist for reasoning purposes, and this knowledge foundation can be augmented, if necessary, with supplementary formal resources. These approaches offer the added advantage of being model agnostic, meaning they explain the decision outcome rather than interpreting the model itself. This flexibility enables their use for both symbolic and subsymbolic systems.

Tightly coupled neurosymbolic algorithms offer another promising solution to the limitations of purely neural or symbolic models. By incorporating the reasoning power of symbolic models, they become inherently capable of generating explanations for their decisions. The most notable examples of tightly coupled neurosymbolic networks with this inherent explainability are LNNs [60] and LTNs [13]. These networks can offer explanations irrespective of their size or complexity, paving a promising path for the development of explainable systems.

5. Outlook

The wake of deep learning has rekindled interest in neurosymbolic integration. This paper presents a comprehensive overview of how this promising approach can enhance the performance of state-of-the-art KGA algorithms. By bridging the gap between symbolic reasoning and data-driven learning, neurosymbolic integration offers a promising pathway to overcome many of the current challenges in KGA.

We have outlined several critical challenges in the KGA and promising research paths for unlocking them by leveraging neurosymbolic methods. These paths comprise a spectrum of innovative strategies, such as: using symbolic knowledge within subsymbolic systems; incorporating soft logical constraints into the repair process; exploring the potential of tightly coupled hybrid systems; and integrating the mapping and repair process into a single framework. Each of these directions represents an opportunity to improve the performance of KGA algorithms in one or more of the key challenges identified.

While our paper showcases several approaches, each focusing on specific aspects of KGA tasks, we envision future research trajectories in this domain to be even more ambitious. Beyond whether individual approaches can successfully tackle specific challenges, we advocate for research endeavours that transcend singular methodologies. The aspiration is to witness the emergence of KGA systems that combine several neurosymbolic approaches, culminating in a system that collectively surpasses prevailing state-of-the-art algorithms. This consolidation of methodologies holds the promise of reshaping the landscape of KGA research and significantly propelling the field forward.

References

- [1] C. Gutiérrez and J.F. Sequeda, Knowledge graphs, *Communications of the ACM* **64**(3) (2021), 96–104.
- [2] P. Kolyvakis, A. Kalousis and D. Kiritsis, Hyperbolic knowledge graph embeddings for knowledge base completion, in: *The Semantic Web: 17th International Conference, ESWC 2020, Heraklion, Crete, Greece, May 31–June 4, 2020, Proceedings 17*, Springer, 2020, pp. 199–214.
- [3] J. Euzenat, P. Shvaiko et al., *Ontology matching*, Vol. 18, Springer, 2007.
- [4] Y. He, J. Chen, D. Antonyrajah and I. Horrocks, BERTMap: a BERT-based ontology alignment system, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 36, 2022, pp. 5684–5691.
- [5] J. Chen, E. Jiménez-Ruiz, I. Horrocks, D. Antonyrajah, A. Hadian and J. Lee, Augmenting ontology alignment by semantic embedding and distant supervision, in: *The Semantic Web: 18th International Conference, ESWC 2021, Virtual Event, June 6–10, 2021, Proceedings 18*, Springer, 2021, pp. 392–408.
- [6] J. Hao, C. Lei, V. Efthymiou, A. Quamar, F. Özcan, Y. Sun and W. Wang, MEDTO: Medical Data to Ontology Matching Using Hybrid Graph Neural Networks, in: *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining, KDD '21*, Association for Computing Machinery, New York, NY, USA, 2021, pp. 29462954–. ISBN 9781450383325. doi:10.1145/3447548.3467138.
- [7] M. Qu, J. Tang and Y. Bengio, Weakly-supervised knowledge graph alignment with adversarial learning, *arXiv preprint arXiv:1907.03179* (2019).
- [8] Z. Wang, Q. Lv, X. Lan and Y. Zhang, Cross-lingual Knowledge Graph Alignment via Graph Convolutional Networks, in: *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, Association for Computational Linguistics, Brussels, Belgium, 2018, pp. 349–357. doi:10.18653/v1/D18-1032. <https://aclanthology.org/D18-1032>.
- [9] T.R. Besold, A.d. Garcez, S. Bader, H. Bowman, P. Domingos, P. Hitzler, K.-U. Kühnberger, L.C. Lamb, D. Lowd, P.M.V. Lima et al., Neural-symbolic learning and reasoning: A survey and interpretation, *arXiv preprint arXiv:1711.03902* (2017).
- [10] A.d. Garcez and L.C. Lamb, Neurosymbolic AI: The 3rd wave, *Artificial Intelligence Review* (2023), 1–20.
- [11] P. Hitzler, A. Eberhart, M. Ebrahimi, M.K. Sarker and L. Zhou, Neuro-symbolic approaches in artificial intelligence, *National Science Review* **9**(6) (2022), nwac035.
- [12] H. Kautz, The third AI summer: AAAI Robert S. Engelmore memorial lecture, *AI Magazine* **43**(1) (2022), 105–125.
- [13] S. Badreddine, A.d. Garcez, L. Serafini and M. Spranger, Logic tensor networks, *Artificial Intelligence* **303** (2022), 103649.
- [14] I. Horrocks, O. Kutz and U. Sattler, The Even More Irresistible SROIQ, in: *Proceedings of the Tenth International Conference on Principles of Knowledge Representation and Reasoning, KR'06*, AAAI Press, 2006, pp. 5767–. ISBN 9781577352716.

- [15] B. Konev, D. Walther and F. Wolter, The Logical Difference Problem for Description Logic Terminologies, in: *Proceedings of the 4th International Joint Conference on Automated Reasoning, IJCAR '08*, Springer-Verlag, Berlin, Heidelberg, 2008, pp. 259274–. ISBN 9783540710691. doi:10.1007/978-3-540-71070-7_21. https://doi.org/10.1007/978-3-540-71070-7_21.
- [16] A. Solimando, E. Jimenez-Ruiz and G. Guerrini, Minimizing conservativity violations in ontology alignments: Algorithms and evaluation, *Knowledge and Information Systems* **51** (2017), 775–819.
- [17] J. Euzenat and P. Shvaiko, *Ontology matching*, 2nd edn, Springer-Verlag, Heidelberg (DE), 2013.
- [18] F. Scharffe, O. Zamazal and D. Fensel, Ontology alignment design patterns, *Knowledge and information systems* **40** (2014), 1–28.
- [19] D. Faria, C. Pesquita, E. Santos, M. Palmonari, I.F. Cruz and F.M. Couto, The agreementmakerlight ontology matching system, in: *On the Move to Meaningful Internet Systems: OTM 2013 Conferences: Confederated International Conferences: CoopIS, DOA-Trusted Cloud, and ODBASE 2013, Graz, Austria, September 9-13, 2013. Proceedings*, Springer, 2013, pp. 527–541.
- [20] V. Iyer, A. Agarwal and H. Kumar, VeeAlign: multifaceted context representation using dual attention for ontology alignment, in: *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 2021, pp. 10780–10792.
- [21] E. Jiménez-Ruiz and B. Cuenca Grau, Logmap: Logic-based and scalable ontology matching, in: *The Semantic Web–ISWC 2011: 10th International Semantic Web Conference, Bonn, Germany, October 23-27, 2011, Proceedings, Part I 10*, Springer, 2011, pp. 273–288.
- [22] W.W. Cohen, P. Ravikumar and S.E. Fienberg, A Comparison of String Distance Metrics for Name-Matching Tasks, in: *Proceedings of the 2003 International Conference on Information Integration on the Web, IIWEB'03*, AAAI Press, 2003, pp. 7378–.
- [23] I. Megdiche, O. Teste and C. Trojahn, An extensible linear approach for holistic ontology matching, in: *The Semantic Web–ISWC 2016: 15th International Semantic Web Conference, Kobe, Japan, October 17–21, 2016, Proceedings, Part I 15*, 2016, pp. 393–410.
- [24] Z. Sun, C. Wang, W. Hu, M. Chen, J. Dai, W. Zhang and Y. Qu, Knowledge Graph Alignment Network with Gated Multi-Hop Neighborhood Aggregation, *Proceedings of the AAAI Conference on Artificial Intelligence* **34**(01) (2020), 222–229. doi:10.1609/aaai.v34i01.5354. <https://ojs.aaai.org/index.php/AAAI/article/view/5354>.
- [25] A. Isaac, L. Van Der Meij, S. Schlobach and S. Wang, An empirical study of instance-based ontology matching, in: *International Semantic Web Conference*, Springer, 2007, pp. 253–266.
- [26] E. Santos, D. Faria, C. Pesquita and F.M. Couto, Ontology alignment repair through modularization and confidence-based heuristics, *PloS one* **10**(12) (2015), e0144807.
- [27] E. Jiménez-Ruiz, B.C. Grau and I. Horrocks, On the Feasibility of Using OWL 2 DL Reasoners for Ontology Matching Problems., in: *ORE*, 2012.
- [28] E. Jiménez-Ruiz, C. Meilicke, B.C. Grau and I. Horrocks, Evaluating Mapping Repair Systems with Large Biomedical Ontologies., *Description Logics* **13** (2013), 246–257.
- [29] P. Shvaiko, F. Giunchiglia and M. Yatskevich, Semantic matching with s-match, in: *Semantic Web Information Management: a Model-Based Perspective*, Springer, 2009, pp. 183–202.
- [30] C. Meilicke and H. Stuckenschmidt, Applying logical constraints to ontology matching, in: *Annual Conference on Artificial Intelligence*, Springer, 2007, pp. 99–113.
- [31] Ontology Alignment Evaluation Initiative, Accessed: August 17, 2023.
- [32] M.A.N. Pour, A. Algergawy, P. Buche, L.J. Castro, J. Chen, H. Dong, O. Fallatah, D. Faria, I. Fundulaki, S. Hertling et al., Results of the ontology alignment evaluation initiative 2022, *CEUR Workshop Proceedings*, 2023.
- [33] D. Faria, M.C. Silva, P. Cotovio, P. Eugénio and C. Pesquita, Matcha and Matcha-DL results for OAEI 2022 (2022).
- [34] Y. He, J. Chen, H. Dong, E. Jiménez-Ruiz, A. Hadian and I. Horrocks, Machine Learning-Friendly Biomedical Datasets for Equivalence and Subsumption Ontology Matching (2022).
- [35] E.A. Feigenbaum, J. Feldman et al., *Computers and thought*, Vol. 7, New York McGraw-Hill, 1963.
- [36] A.M. Turing, *Computing machinery and intelligence*, Springer, 2009.
- [37] D.B. Lenat, CYC: A Large-Scale Investment in Knowledge Infrastructure, *Commun. ACM* **38**(11) (1995), 3338–. doi:10.1145/219717.219745.
- [38] G.A. Miller, WordNet: a lexical database for English, *Communications of the ACM* **38**(11) (1995), 39–41.
- [39] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li and L. Fei-Fei, Imagenet: A large-scale hierarchical image database, in: *2009 IEEE conference on computer vision and pattern recognition*, Ieee, 2009, pp. 248–255.
- [40] Y. LeCun, Y. Bengio and G. Hinton, Deep learning, *nature* **521**(7553) (2015), 436–444.
- [41] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, Ł. Kaiser and I. Polosukhin, Attention is all you need, *Advances in neural information processing systems* **30** (2017).
- [42] T. Brown, B. Mann, N. Ryder, M. Subbiah, J.D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever and D. Amodei, Language Models are Few-Shot Learners, in: *Advances in Neural Information Processing Systems*, Vol. 33, H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan and H. Lin, eds, Curran Associates, Inc., 2020, pp. 1877–1901. <https://proceedings.neurips.cc/paper/2020/file/1457c9d66fcb4967418bfb8ac142f64a-Paper.pdf>.
- [43] S. Ren, K. He, R. Girshick and J. Sun, Faster r-cnn: Towards real-time object detection with region proposal networks, *Advances in neural information processing systems* **28** (2015).
- [44] T.N. Kipf and M. Welling, Semi-supervised classification with graph convolutional networks, *arXiv preprint arXiv:1609.02907* (2016).
- [45] K. He, X. Zhang, S. Ren and J. Sun, Deep Residual Learning for Image Recognition, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [46] D. Silver, A. Huang, C.J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot et al., Mastering the game of Go with deep neural networks and tree search, *nature* **529**(7587) (2016), 484–489.

- [47] D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton et al., Mastering the game of go without human knowledge, *nature* **550**(7676) (2017), 354–359.
- [48] T. Baltruaitis, C. Ahuja and L.-P. Morency, Multimodal Machine Learning: A Survey and Taxonomy, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **41**(2) (2019), 423–443. doi:10.1109/TPAMI.2018.2798607.
- [49] D. Kahneman, *Thinking fast and slow*, Farrar, 2000.
- [50] C.-W. Lee, W. Fang, C.-K. Yeh and Y.-C.F. Wang, Multi-label zero-shot learning with structured knowledge graphs, in: *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 1576–1585.
- [51] W. Liu, P. Zhou, Z. Zhao, Z. Wang, Q. Ju, H. Deng and P. Wang, K-bert: Enabling language representation with knowledge graph, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34, 2020, pp. 2901–2908.
- [52] K. Ma, J. Francis, Q. Lu, E. Nyberg and A. Oltramari, Towards Generalizable Neuro-Symbolic Systems for Commonsense Question Answering, in: *Proceedings of the First Workshop on Commonsense Inference in Natural Language Processing*, 2019, pp. 22–32.
- [53] M.K. Sarker, N. Xie, D. Doran, M. Raymer and P. Hitzler, Explaining trained neural networks with semantic web technologies: First steps, *arXiv preprint arXiv:1710.04324* (2017).
- [54] Y. Geng, J. Chen, E. Jimenez-Ruiz and H. Chen, Human-centric transfer learning explanation via knowledge graph, *arXiv preprint arXiv:1901.08547* (2019).
- [55] 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.
- [56] A. Dalal, M.K. Sarker, A. Barua and P. Hitzler, Explaining Deep Learning Hidden Neuron Activations using Concept Induction, *arXiv preprint arXiv:2301.09611* (2023).
- [57] J. Mao, C. Gan, P. Kohli, J.B. Tenenbaum and J. Wu, The neuro-symbolic concept learner: Interpreting scenes, words, and sentences from natural supervision, *arXiv preprint arXiv:1904.12584* (2019).
- [58] F. Arabshahi, S. Singh and A. Anandkumar, Combining symbolic expressions and black-box function evaluations in neural programs, *arXiv preprint arXiv:1801.04342* (2018).
- [59] F. Arabshahi, Z. Lu, S. Singh and A. Anandkumar, Memory augmented recursive neural networks, *arXiv* (2019).
- [60] R. Riegel, A. Gray, F. Luus, N. Khan, N. Makondo, I.Y. Akhalwaya, H. Qian, R. Fagin, F. Barahona, U. Sharma et al., Logical neural networks, *arXiv preprint arXiv:2006.13155* (2020).
- [61] P. Smolensky, M. Lee, X. He, W.-t. Yih, J. Gao and L. Deng, Basic reasoning with tensor product representations, *arXiv preprint arXiv:1601.02745* (2016).
- [62] B. Wagner and A. d’Avila Garcez, Neural-symbolic integration for fairness in AI, in: *CEUR Workshop Proceedings*, Vol. 2846, 2021.
- [63] L.C. Lamb, A. Garcez, M. Gori, M. Prates, P. Avelar and M. Vardi, Graph neural networks meet neural-symbolic computing: A survey and perspective, *arXiv preprint arXiv:2003.00330* (2020).
- [64] Y. He, J. Chen, E. Jiménez-Ruiz, H. Dong and I. Horrocks, Language Model Analysis for Ontology Subsumption Inference (2023).
- [65] E. Jiménez-Ruiz, B.C. Grau, I. Horrocks and R. Berlanga, Logic-based assessment of the compatibility of UMLS ontology sources, *Journal of biomedical semantics* **2**(1) (2011), 1–16.
- [66] D.J. Hejna III and D. Sadigh, Few-shot preference learning for human-in-the-loop rl, in: *Conference on Robot Learning*, PMLR, 2023, pp. 2014–2025.
- [67] P.F. Christiano, J. Leike, T. Brown, M. Martic, S. Legg and D. Amodei, Deep reinforcement learning from human preferences, *Advances in neural information processing systems* **30** (2017).
- [68] K. Lee, L. Smith and P. Abbeel, Pebble: Feedback-efficient interactive reinforcement learning via relabeling experience and unsupervised pre-training, *arXiv preprint arXiv:2106.05091* (2021).
- [69] M.T. Ribeiro, S. Singh and C. Guestrin, "Why should i trust you?" Explaining the predictions of any classifier, in: *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, 2016, pp. 1135–1144.
- [70] S.M. Lundberg and S.-I. Lee, A Unified Approach to Interpreting Model Predictions, in: *Advances in Neural Information Processing Systems*, Vol. 30, I. Guyon, U.V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan and R. Garnett, eds, Curran Associates, Inc., 2017. <https://proceedings.neurips.cc/paper/2017/file/8a20a8621978632d76c43dfd28b67767-Paper.pdf>.
- [71] C. Pesquita, Towards Semantic Integration for Explainable Artificial Intelligence in the Biomedical Domain., 2021.
- [72] C. Widmer, M.K. Sarker, S. Nadella, J. Fiechter, I. Juvina, B. Minnery, P. Hitzler, J. Schwartz and M. Raymer, Towards Human-Compatible XAI: Explaining Data Differentials with Concept Induction over Background Knowledge, *arXiv preprint arXiv:2209.13710* (2022).
- [73] W. Xiong, T. Hoang and W.Y. Wang, Deeppath: A reinforcement learning method for knowledge graph reasoning, *arXiv preprint arXiv:1707.06690* (2017).
- [74] A. Zhu, D. Ouyang, S. Liang and J. Shao, Step by Step: A Hierarchical Framework for Multi-Hop Knowledge Graph Reasoning with Reinforcement Learning, *Know.-Based Syst.* **248**(C) (2022). doi:10.1016/j.knosys.2022.108843.