

Learning Permutation-Invariant Embeddings for Description Logic Concepts

Caglar Demir and Axel-Cyrille Ngonga Ngomo

Data Science Research Group, Paderborn University

Abstract. Concept learning deals with learning description logic concepts from a background knowledge and input examples. The goal is to learn a concept that covers all positive examples, while not covering any negative examples. This non-trivial task is often formulated as a search problem within an infinite quasi-ordered concept space. Although state-of-the-art models have been successfully applied to tackle this problem, their large-scale applications have been severely hindered due to their excessive exploration incurring impractical runtimes. Here, we propose a remedy for this limitation. We reformulate the learning problem as a multi-label classification problem and propose a neural embedding model (NERO) that learns permutation-invariant embeddings for sets of examples tailored towards predicting F_1 scores of pre-selected description logic concepts. By ranking such concepts in descending order of predicted scores, a possible goal concept can be detected within few retrieval operations, i.e., no excessive exploration. Importantly, top-ranked concepts can be used to start the search procedure of state-of-the-art symbolic models in multiple advantageous regions of a concept space, rather than starting it in the most general concept \top . Our experiments on 5 benchmark datasets with 770 learning problems firmly suggest that NERO significantly (p-value < 1%) outperforms the state-of-the-art models in terms of F_1 score, the number of explored concepts, and the total runtime. We provide an open-source implementation of our approach.¹

Keywords: Description Logics, Concept Learning, Permutation Invariance

1 Introduction

Deep learning based models have been effectively applied to tackle various graph-related problems, including question answering, link prediction [19,30]. Yet, their predictions are not human-interpretable and confined within a fixed set vocabulary terms [9,11]. In contrast, Description Logics (DLs) provide means to derive human-interpretable inference in an infinite setting [1,16,27]. Deriving explanations for DLs concepts has been long understood [5]. For instance, explanations can be derived by using the subsumption hierarchy as a sequence of binary

¹ <https://github.com/dice-group/Nero>

classifiers in a fashion akin to following a path in decision tree [35,4]. Utilizing DLs is considered as a possible backbone for explainable Artificial Intelligence (AI) [33]. Although DLs have become standard techniques to formalize Knowledge Base (KB) [19,20,29], the highly incomplete nature of KBs and impractical runtimes of symbolic models have been a challenge for fulfilling its potential. State-of-the-art Concept Learning (CL) models have been successfully applied to learn DL concepts from a KB and input examples [24,27]. Yet, their practical applications have been severely hindered by their impractical runtimes. This limitation stems from the reliance of myopic heuristic function that often incurs excessive exploration of concepts [16,22,35]. A DL concept is explored by retrieving its individuals and calculating its quality w.r.t. input KB and examples (see Section 2). As the size of an input KB grows, excessive exploration has been a computational bottleneck in practical applications. Here, we propose a remedy for this limitation. We reformulate the learning problem as a multi-label classification problem and propose NERO—a neural permutation-invariant embedding model. Given a set of positive examples E^+ and a set of negative examples E^- , NERO predicts F_1 scores of pre-selected DL concepts as shown in Figure 1.

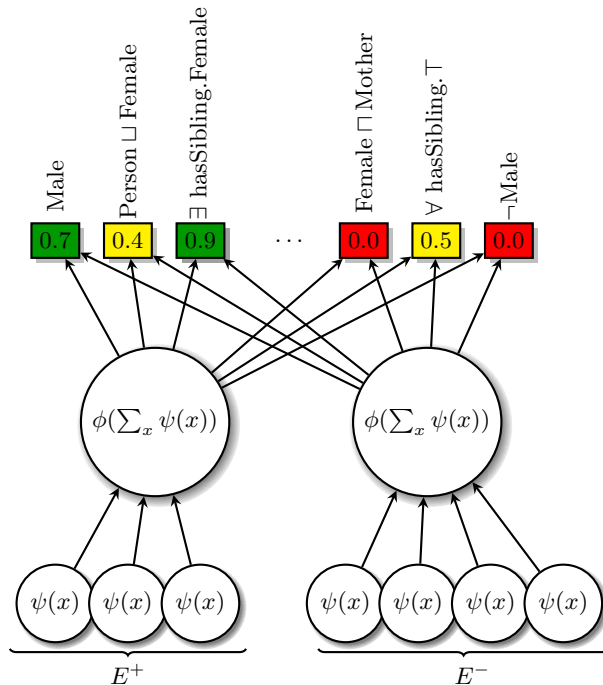


Fig. 1. Visualization of NERO. Boxes and values denote the pre-selected unique DL concepts and their predicted F_1 scores, respectively.

By ranking pre-selected DL concepts in descending order of predicted scores, a goal concept can be found by only exploring few top-ranked concepts. Importantly, top-ranked concepts can be used to initialize the standard search procedure of state-of-the-art models, if a goal concept is not found. By this, a state-of-the-art CL model is endowed with the capability of starting the search in more advantageous states, instead of starting it in the most general concept \top . Our experiments on 5 benchmark datasets with 770 learning problems indicate that NERO **significantly** (p-value < 1%) outperforms the state-of-the-art models in standard metrics such as F_1 score, the number of explored concepts, and the total runtime. Importantly, equipping NERO with a state-of-the-art model (CELOE) further improves F_1 scores on benchmark datasets with a low runtime cost. The results of Wilcoxon signed rank tests confirm that the superior performance of NERO is significant. We provide an open-source implementation of NERO, including pre-trained models, evaluation scripts as well as a web service.²

2 Background

Knowledge Base: A Knowledge Base (KB) is a pair $\mathcal{K} = (Tbox, Abox)$, where $Tbox$ is a set of terminological axioms describing relations between named concepts N_C [33]. A terminological axiom is in the form of $A \sqsubseteq B$ or $A \equiv B$ s.t. $A, B \in N_C$. $Abox$ is a set of assertions describing relationships among individuals $a, b \in N_I$ via roles $r \in N_R$ as well as concept membership relationships between N_I and N_C . Every assertion in $Abox$ must in the form of $A(x)$ and $r(x, y)$, where $A \in N_C$, $r \in N_R$, and $x, y \in N_I$. An example is visualized in Figure 2.

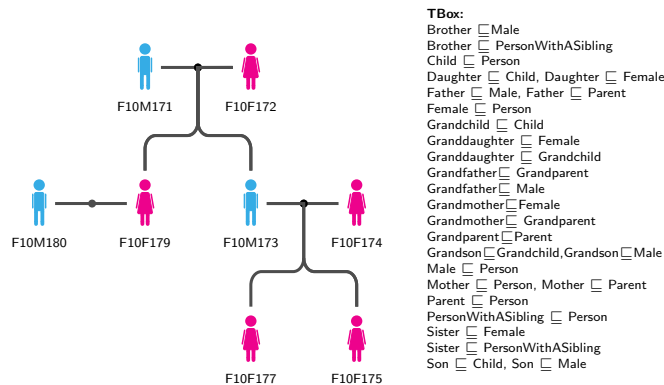


Fig. 2. A visualization of Family KB with $Tbox$ and a subset of $Abox$. Colors denote concept assertions, while (\cdot) and branching from (\cdot) denote role assertions, respectively.

² <https://github.com/dice-group/Nero>

Description Logics: Description Logics (DLs) are fragments of first-order predicate logic using only unary and binary predicates. The unary predicates, the binary predicates and constants are called concepts, roles and individuals, respectively [1]. DL have become standard techniques to formalize background knowledge for many application domains including Semantic Web [20,29]. Leveraging KBs defined over DLs has a potential of being a backbone for explainable AI [33]. Here, we consider KBs in the DL \mathcal{ALC} (Attributive Language with Complements) [1] as in many other works (see Section 3). The model-theoretic semantics of \mathcal{ALC} are given in Table 1.

Table 1. \mathcal{ALC} syntax and semantics. \mathcal{I} stands for an interpretation, $\Delta^{\mathcal{I}}$ for its domain.

Construct	Syntax	Semantics
Atomic concept	A	$A^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$
Role	r	$r^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$
Top concept	\top	$\Delta^{\mathcal{I}}$
Bottom concept	\perp	\emptyset
Conjunction	$C \sqcap D$	$C^{\mathcal{I}} \cap D^{\mathcal{I}}$
Disjunction	$C \sqcup D$	$C^{\mathcal{I}} \cup D^{\mathcal{I}}$
Negation	$\neg C$	$\Delta^{\mathcal{I}} \setminus C^{\mathcal{I}}$
Existential restriction	$\exists r.C$	$\{x \mid \exists y.(x,y) \in r^{\mathcal{I}} \text{ and } y \in C^{\mathcal{I}}\}$
Universal restriction	$\forall r.C$	$\{x \mid \forall y.(x,y) \in r^{\mathcal{I}} \text{ implies } y \in C^{\mathcal{I}}\}$

Concept Learning: Let \mathcal{K} over \mathcal{ALC} , the set $E^+ \subset N_I$ of positive examples, and the set $E^- \subset N_I$ of negative examples be given. The DL concept learning problem is defined as follows

$$\forall p \in E^+, \forall n \in E^- (\mathcal{K} \models \mathbb{H}(p)) \wedge (\mathcal{K} \not\models \mathbb{H}(n)), \quad (1)$$

where $\mathbb{H} \in \mathcal{C}$ denotes an \mathcal{ALC} concept and \mathcal{C} denotes all valid \mathcal{ALC} concepts under the construction rules: $\mathbb{C} ::= \mathbb{A} \mid \neg \mathbb{C} \mid \mathbb{C} \sqcap \mathbb{C} \mid \mathbb{C} \sqcup \mathbb{C} \mid \exists r.C \mid \forall r.C$, where $\mathbb{A} \in N_C$ and $r \in N_R$. $\mathcal{K} \models \mathbb{H}(p)$ implies that an inference of the class membership $\mathbb{H}(p)$ is a logical consequence of \mathcal{K} . Checking whether a \mathbb{H} fulfills Equation (1) is performed by a retrieval function $\mathcal{R} : \mathcal{C} \rightarrow 2^{N_I}$ defined under Open World Assumption (OWA) or Close World Assumption (CWA). This non-trivial learning problem is often transformed into a search problem within a quasi-ordered \mathcal{ALC} concept space (\mathcal{S}, \preceq) [7,13,27,34]. Traversing in \mathcal{S} is commonly conducted via a top-down refinement operator defined as $\rho : \mathcal{S} \rightarrow 2^{\mathcal{S}}$ with

$$\forall \mathbb{A} \in \mathcal{S} : \rho(\mathbb{A}) \subseteq \{\mathbb{B} \in \mathcal{S} \mid \mathbb{B} \preceq \mathbb{A}\}. \quad (2)$$

State-of-the-art CL models begin their search towards a \mathbb{H} , after a search tree is initialized with the most general DL concept (\top) as a root node. This search tree

is iteratively built by selecting a node containing a quasi-ordered DL concept with the highest heuristic value and adding its qualifying refinements as its children into a search tree [27].

Heuristics: A heuristic function is the key to an efficient search in \mathcal{S} towards a \mathbb{H} [26]. The number of explored concepts and runtimes are used as proxy for the efficiency. Various heuristic functions have been investigated [26,35]. Most heuristic functions of state-of-the-art models can be considered as myopic functions favoring syntactically short and accurate concepts. Hence, they are prone to stuck in a local optimum [35]. For instance, the heuristic function of CELOE is defined as

$$\phi_{\text{CELOE}}(\mathbf{A}, \mathbf{B}) = Q(\mathbf{B}) + \lambda \cdot [Q(\mathbf{B}) - Q(\mathbf{A})] - \beta \cdot |\mathbf{B}|, \quad (3)$$

where $\mathbf{A} \in \mathcal{S}$, $\mathbf{B} \in \rho(\mathbf{A})$. $\beta > \lambda \geq 0$ and $Q(\cdot)$ denotes a quality function (e.g. F_1 score or accuracy). Through $Q(\cdot)$ and $|\cdot|$, the search is steered based on solely \mathbf{A} and \mathbf{B} towards more accurate and syntactically shorter concepts. $F_1(\cdot)$ is defined as

$$F_1(\mathbf{A}) = \frac{|E^+ \cap \mathcal{R}(\mathbf{A})|}{|E^+ \cap \mathcal{R}(\mathbf{A})| + 0.5(|E^- \cap \mathcal{R}(\mathbf{A})| + |E^+ \setminus \mathcal{R}(\mathbf{A})|)}. \quad (4)$$

As the size of KB grows, runtimes of performing retrieval operations $\mathcal{R}(\cdot)$ increase [3,4,25]. Consequently, traversing in \mathcal{S} becomes a computational bottleneck. Therefore, reducing the number of explored concepts plays an important role to tackle CL on KBs. Although state-of-the-art models (e.g. CELOE) apply redundancy elimination and expression simplification rules to reduce the number of explored concepts, impractical long runtimes of state-of-the-art models still prohibit large-scale applications [17]. Moreover, the selected assumption underlying $\mathcal{R}(\cdot)$ also plays a role to tackle CL on large KBs. Due to the incomplete nature of KBs, OWA seems to be a more suitable assumption [31]. Yet, Using OWA often makes membership queries computationally more challenging [12,26]. Consequently, CWA is often adopted in many recent works [16,22,34].

3 Related Work

A plethora of works have investigated learning DLs concepts from a KB and input examples. We refer to [1,18,23] for an introduction. Most symbolic systems differ in the usage of heuristic functions and the design of the refinement operators [2,7,12,14,26,27,21,34]. DL-Learner [24] is regarded as the most mature and recent system for CL [32]. DL-Learner consists of several state-of-the-art models, including ELTL, OCEL, and CELOE. ELTL is based on a refinement operator for the DL \mathcal{EL} and uses a heuristic function that favors syntactically short concepts. CELOE builds on OCEL and ELTL and it applies a more sophisticated heuristic function. CELOE is currently the best CL model available within DL-Learner and often outperforms many state-of-the-art models including OCEL and ELTL

in terms of the quality of learned expression, number explored concepts, and runtimes [35,28]. The aforementioned approaches apply redundancy elimination and expression simplification rules to reduce the number of explored concepts. Although applying redundancy elimination and expression simplification rules often reduce the number of explored concepts, these operations introduces more computation and long runtimes still prohibit large-scale applications [17]. Most recent works have focused on treating the impractical runtimes in CL. CLIP [22] is a neural approach that serves as an addition to refinement-based approaches and supports pruning the search space by predicting the length of a possible goal state. EvoLearner [16] represents a concept as an abstract syntax tree corresponding an individual of an evolutionary algorithm. The initial population of individuals is obtained via biased random walks originating from E^+ . Westphal et al. [35] design a Simulated Annealing based meta-heuristic to balance the exploration-exploitation trade-off during the search process. In this work, we mainly evaluate NERO against CELOE provided in DL-Learner for two reasons: (1) DL-Learner is regarded as the most mature and recent system for CL [32] and (2) most recently developed models are often evaluated w.r.t. the quality of concepts as well as runtimes. Yet, not reporting the number of explored concepts does not permit us to quantify whether a possible improvement through NERO may stem from our novel idea or our efficient implementation. Consequently, in our experiments, we mainly compare NERO against CELOE in terms of number of explored concepts, quality of learned concepts as well as runtimes.

4 Methodology

Motivation: The goal in the CL problem is to find a DL concept $\mathbb{H} \in \mathcal{C}$ maximizing Equation (4). Here, we are interested in achieving this goal by learning permutation-invariant embeddings tailored towards predicting F_1 scores of pre-selected concepts. Through exploring top-ranked concepts at first, we aim to find a goal concept can only with few retrieval operations. If a goal state is not found within top-ranked concepts, the search tree of a state-of-the-art CL model can be initialized with top-ranked concepts and \top concept along with corresponding heuristic values. By this, the standard search procedure can be started in more advantageous states, than the most general concept \top .

Approach: Equation (4) indicates that $F_1(\cdot)$ is invariant to the order of individuals in E^+ , E^- , and $\mathcal{R}(\cdot)$. Previously, Zaheer et al. [36] have proven that all functions being invariant to the order in inputs can be decomposed into

$$f(\mathbf{x}) = \phi\left(\sum_{x \in \mathbf{x}} \psi(x)\right), \quad (5)$$

where $\mathbf{x} = \{x_1, \dots, x_m\} \in 2^{\mathcal{X}}$ and $\phi(\cdot)$ and $\psi(\cdot)$ denote a set of input and two parameterized continuous functions, respectively. A permutation-invariant neural network defined via Equation (5) still abides by the universal approximation

theorem [36]. We conjecture that such neural network can learn permutation-invariant embeddings for sets of individuals (e.g. E^+ and E^-) tailored towards predicting F_1 scores of pre-selected concepts. Through accurately predicting F_1 scores of pre-selected DL concepts, possible goal concepts from pre-selected concepts can be detected without using $F_1(\cdot)$ and $\mathcal{R}(\cdot)$. With these considerations, we define NERO as follows

$$\text{NERO}(E^+, E^-) = \sigma \left(\phi \left(\sum_{x \in E^+} \psi(x) \right) - \phi \left(\sum_{x \in E^-} \psi(x) \right) \right), \quad (6)$$

where $\psi(\cdot) : N_I \rightarrow \mathbb{R}^m$ and $\phi : \mathbb{R}^m \rightarrow [0, 1]^{|T|}$ denote an embedding look-up operation and an affine transformation, respectively. \mathcal{T} represents the pre-selected DL concepts. The result of the translation operation denoted with $\mathbf{z} \in \mathbb{R}^m$ is normalized via the logistic sigmoid function $\sigma(\mathbf{z}) = \frac{1}{1 + \exp(-\mathbf{z})}$. Hence, $\text{NERO} : 2^{N_I} \times 2^{N_I} \mapsto [0, 1]^{|T|}$ can be seen as a mapping from two sets of individuals to $|\mathcal{T}|$ unit intervals. NERO can be seen as a multi-task learning approach that leverages the similarity between multi-tasks, where a task in our case corresponds to accurately predicting the F_1 score of a pre-selected DL concept [8].

The importance of learning representations tailored towards *related* tasks has been well investigated [15,8]. Motivated by this, we elucidate the process of selecting DL concepts in Algorithm 1. We select such concepts that their canonical interpretations do not fully overlap (see the 4.th line). As shown therein, NERO can be trained on knowledge base defined over any DLs provided that $\mathcal{R}(\cdot)$ and $\rho(\cdot)$ are given.

Algorithm 1 Constructing target DL concepts

Input: $\mathcal{R}(\cdot)$, $\rho(\cdot)$, d , maxlength **Output:** \mathcal{T}

- 1: $\mathcal{T} := \{C \mid C \in \rho(\mathcal{T}) \wedge |C| \leq \text{maxlength} \wedge 0 < |\mathcal{R}(C)|\}$
- 2: **for each** $A \in \mathcal{T}$ **do**
- 3: **for each** $B \in \mathcal{T}$ **do**
- 4: **if** $\mathcal{R}(A) \neq \mathcal{R}(B)$ **then**
- 5: **for each** $X \in \{A \sqcap B, A \sqcup B\}$ **do**
- 6: **if** $|\mathcal{R}(X)| > 0 \wedge \mathcal{R}(X) \notin \{\mathcal{R}(E) \mid E \in \mathcal{T}\}$ **then**
- 7: Add X to \mathcal{T} .
- 8: **end if**
- 9: **if** $|\mathcal{T}| = d$ **then**
- 10: **return** \mathcal{T}
- 11: **end if**
- 12: **end for**
- 13: **end if**
- 14: **end for**
- 15: **end for**
- 16: **if** $|\mathcal{T}| < d$ **then**
- 17: Go to the step (2).
- 18: **end if**

Training Process: Let $\mathcal{D} = \{(E_i^+, E_i^-, \mathbf{y}_i)\}_{i=1}^N$ represent a training dataset, where a data point (E^+, E^-, \mathbf{y}) is obtained in four consecutive steps: (i) Sample \mathcal{C} from \mathcal{T} uniformly at random, (ii) Sample k individuals $E^+ \subset \mathcal{R}(\mathcal{C})$ uniformly at random, (iii) Sample k individuals $E^- \subset N_I \setminus E^+$ uniformly at random, and (iv) Compute F_1 scores \mathbf{y} via Equation (4) w.r.t. E^+, E^- , for \mathcal{T} . For a given (E^+, E^-, \mathbf{y}) and predictions $\hat{\mathbf{y}} := \text{NERO}(E^+, E^-)$, an incurred binary cross entropy loss. Important to note that after training process, permutation-invariant embeddings of any \mathcal{ALC} DL concepts can be readily obtained omitting the translation operation in NERO, e.g. embeddings of a DL concept (e.g. $\text{Male} \sqcap \exists \text{hasSibling.Female}$) can be obtained via $\phi(\sum_{x \in \mathcal{R}(\text{Male} \sqcap \exists \text{hasSibling.Female})} \psi(x))$. In our project page, we provided a 2D visualization of learned embeddings for the Family KB.

5 Experiments

We based our experimental setup on [6,26,7] and used learning problems provided therein. An overview of the datasets is provided in Table 2. To perform extensive comparisons between models, additional learning problems are generated by randomly sampling E^+ and E^- . We ensured that none of the learning problems used in our evaluation has been used in the unsupervised training phase. In our experiments, we evaluated all models in \mathcal{ALC} for Class Expression Learning (CEL) on the same hardware.

Table 2. An overview of class expression learning benchmark datasets.

Dataset	$ N_I $	$ N_C $	$ N_R $
Family	202	18	4
Carcinogenesis	22372	142	21
Mutagenesis	14145	86	11
Biopax	323	28	49
Lymphography	148	49	1

We evaluated models via the F_1 score, the runtime and number of explored concepts. The F_1 score is used to measure the quality of the concepts found w.r.t. positive and negative examples, while the runtime and the number of explored concepts are to measure the efficiency. We measured the full computation time including the time spent preprocessing time of the input data and tackling the learning problem. Moreover, we used two standard stopping criteria for state-of-the-art models. (i) We set the maximum runtime to 10 seconds although models often reach good solutions within 1.5 seconds [27]. (ii) The models are configured to terminate as soon as they found a goal concept. In our experiments, we evaluate all models in \mathcal{ALC} for CL on the same hardware. During training, we set $|\mathcal{T}| = 1000$, $N = 50$ and used Adam optimizer for NERO. We only considered top-100 ranked concepts to evaluate NERO.

6 Results

Results with Benchmark Learning Problems: Table 3 reports the concept learning results with benchmark learning problems. Table 3 suggests that equipping NERO with the standard search procedure improves the state-of-the-art performance in terms of F_1 scores even further with a small cost of runtimes. CELOE and ELTL require at least $14.7\times$ more time than NERO to find accurate concepts on Family. This stems from the fact that NERO explores on average only 21 concepts, whereas CELOE explored 1429. On Mutagenesis and Carcinogenesis, NERO finds more accurate concepts, while exploring less, hence, achieving better runtime performance. Runtime gains stem from the fact that NERO explores at least $2.3\times$ fewer concepts.

Table 3. Results on benchmark learning problems. F_1 , T, and Exp. denote F_1 score, total runtime in seconds, and the number of explored concepts, respectively. NERO[†] denotes equipping NERO with CELOE. ELTL does not report the Exp.

Dataset	NeRo [†]			NeRo			CELOE			ELTL	
	F_1	T	Exp.	F_1	T	Exp.	F_1	T	Exp.	F_1	T
Family	.987	.83	26	.984	.68	21	.980	4.65	1429	.964	4.12
Mutagenesis	.714	17.30	200	.704	13.18	100	.704	23.05	516	.704	21.04
Carcinogenesis	.725	32.23	200	.720	26.26	100	.714	37.18	230	.719	36.29

Important to note we did not use parallelism in NERO and we reload parameters of NERO for each single learning problem. To conduct more extensive evaluation, we generated total 750 random learning problems on five benchmark datasets. Since Lymphography and Biopax datasets do not contain any learning problems, they are not included in Table 3.

Results with Random Learning Problems: Table 4 reports the concept learning results with random learning problems. Table 4 suggests that CELOE explores at least $3.19\times$ more concepts than NERO. Importantly, NERO finds on-par or more accurate concepts, while exploring less. Here, we load the parameters of NERO only once per dataset and are used to tackle learning problems sequentially. This resulted in reducing the total computation time of NERO by $3 - 6\times$ on Family, Mutagenesis and Carcinogenesis benchmark datasets. Although NERO can tackle learning problems in parallel (e.g. through multiprocessing), we did not use any parallelism, since CELOE and ELTL do not abide by parallelism [7]. Loading the learning problems in a standard mini-batch fashion and using multi-GPUs may further improve the runtimes of NERO. These results suggest that NERO can be more suitable than CELOE and ELTL on applications requiring low latency.

Table 4. Random learning problems with different sizes per benchmark dataset. Each row reports the mean and standard deviations attained in 50 learning problems. $|E|$ denotes $|E^+| + |E^-|$.

Dataset	$ E $	NeRo			CELOE			ELTL	
		F_1	T	Exp.	F_1	T	Exp.	F_1	T
Family	10	.913 ± .06	.16 ± .51	74 ± 43	.903 ± .06	11.61 ± 3.58	5581 ± 2375	.718 ± .01	4.45 ± 2.84
	20	.807 ± .04	.16 ± .49	100 ± 00	.795 ± .05	13.28 ± 1.47	7586 ± 645	.678 ± .02	3.59 ± 1.27
	30	.775 ± .03	.15 ± .41	100 ± 00	.760 ± .03	13.24 ± 1.42	7671 ± 575	.672 ± .01	3.46 ± 1.59
Lymphography	10	.968 ± .07	.12 ± .43	75 ± 41	.968 ± .07	6.63 ± 4.29	5546 ± 5169	.733 ± .09	3.07 ± .30
	20	.828 ± .04	.13 ± .40	100 ± 00	.826 ± .05	13.01 ± 1.23	11910 ± 1813	.678 ± .02	3.08 ± .50
	30	.780 ± .04	.13 ± .01	100 ± 00	.780 ± .04	13.02 ± 1.69	13138 ± 2601	.672 ± .01	3.09 ± .72
Biopax	10	.859 ± .08	.19 ± .71	86 ± 34	.806 ± .07	13.26 ± 1.94	4752 ± 2153	.685 ± .06	3.71 ± .10
	20	.793 ± .05	.19 ± .52	100 ± 00	.746 ± .04	13.63 ± .10	4151 ± 748	.668 ± .06	3.72 ± .10
	30	.749 ± .03	.18 ± .52	100 ± 00	.718 ± .02	13.91 ± .44	3843 ± 963	.668 ± .06	3.90 ± .22
Mutagenesis	10	.777 ± .05	3.47 ± 1.61	100 ± 00	.753 ± .06	20.27 ± 1.39	546 ± 613	.670 ± .02	10.29 ± .40
	20	.746 ± .05	3.09 ± 1.75	100 ± 00	.712 ± .02	20.38 ± 1.30	430 ± 28	.667 ± .00	10.73 ± 1.10
	30	.721 ± .03	2.89 ± 1.60	100 ± 00	.700 ± .02	20.39 ± 1.06	429 ± 38	.667 ± .00	11.74 ± .97
Carcinogenesis	10	.768 ± .06	5.39 ± 2.98	98 ± 14	.764 ± .06	29.90 ± 1.02	401 ± 125	.673 ± .05	19.99 ± .67
	20	.722 ± .03	5.40 ± 1.87	100 ± 00	.713 ± .02	30.30 ± .19	318 ± 152	.667 ± .00	20.00 ± 1.11
	30	.704 ± .05	4.70 ± 2.78	100 ± 00	.697 ± .02	29.99 ± .58	319 ± 43	.667 ± .00	20.38 ± .85

Table 5. Performance comparison with different number of explored concepts. Each row reports the mean and standard deviations attained in 50 learning problems.

Dataset	$ E $	NeRo-1		NeRo-10		NeRo-1000	
		F_1	T	F_1	T	F_1	T
Family	10	.906 ± .07	.08 ± .06	.910 ± .05	.09 ± .06	.916 ± .06	.81 ± .50
	20	.793 ± .05	.08 ± .05	.806 ± .04	.09 ± .05	.807 ± .04	1.17 ± .50
	30	.742 ± .05	.08 ± .05	.773 ± .03	.09 ± .05	.775 ± .03	1.15 ± .50
Lymphography	10	.882 ± .07	.08 ± .06	.905 ± .05	.08 ± .06	.916 ± .06	.77 ± .50
	20	.793 ± .05	.07 ± .05	.827 ± .04	.08 ± .05	.828 ± .04	1.03 ± .50
	30	.738 ± .05	.07 ± .06	.777 ± .04	.08 ± .05	.780 ± .03	1.00 ± .60
Biopax	10	.853 ± .08	.09 ± .06	.856 ± .05	.97 ± .59	.868 ± .08	1.31 ± .80
	20	.779 ± .05	.09 ± .06	.791 ± .04	.10 ± .62	.793 ± .04	1.35 ± .60
	30	.708 ± .07	.09 ± .06	.742 ± .04	.10 ± .63	.749 ± .03	1.39 ± .60
Mutagenesis	10	.733 ± .07	.32 ± 2.03	.785 ± .06	.57 ± 1.81	.803 ± .06	36.98 ± 5.81
	20	.689 ± .08	.31 ± 1.99	.734 ± .05	.52 ± 1.82	.751 ± .04	34.29 ± 5.91
	30	.673 ± .08	.31 ± 2.03	.712 ± .04	.49 ± 1.77	.728 ± .03	32.88 ± 6.13
Carcinogenesis	10	.717 ± .09	.41 ± 2.53	.740 ± .09	.89 ± 2.89	.783 ± .02	56.941 ± 9.55
	20	.680 ± .06	.40 ± 2.49	.707 ± .05	.82 ± 2.45	.731 ± .02	57.205 ± 5.08
	30	.610 ± .11	.41 ± 2.83	.671 ± .06	.77 ± 2.66	.716 ± .02	52.872 ± 7.58

Results with Limited Exploration: Table 5 reports concept learning results with limited exploration on five benchmark datasets. Table 5 suggests that NERO-10 often outperforms CELOE and ELTL (see Table 4) in all metrics even when exploring solely 10 top-ranked concepts.

Significance Testing: To validate the significance of our results, we performed Wilcoxon signed-rank tests (one and two-sided) on F_1 scores, runtimes and the number of explored concepts. Our null hypothesis was that the performances of NERO and CELOE come from the same distribution. We were able to reject the null hypothesis with a p-value $< 1\%$ across all the datasets, hence, the superior performance of NERO is statistically significant.

6.1 Discussion

Our results uphold our hypothesis: F_1 scores of DL concepts can be accurately predicted by means of learning permutation-invariant embeddings for sets of individuals. Through considering top-ranked DL concepts at first, the need of excessive number of retrieval operations to find a goal concept can be mitigated. Throughout our experiments, NERO consistently outperforms state-of-the-art models w.r.t. the F_1 score, the number of explored concepts and the total computational time. Importantly, starting the standard search procedure on these top-ranked concepts further improves the results. Hence, NERO can be applied within state-of-the-art models to decrease their runtimes. However, it is important to note that Lehmann et al. [26] have previously proved the completeness of CELOE in the CL problem, i.e., for a given learning problem, CELOE finds a goal expression if it exists provided that there are no upper-bounds on the time and memory requirements. Although these requirements are simply not practical, equipping NERO with the search procedure of CELOE is necessary to achieve the completeness in CL.

7 Conclusion

We introduced a permutation-invariant neural embedding model (NERO) to efficiently tackle the description logic concept learning problem. For given learning problem, NERO accurately predicts F_1 scores of pre-selected description logic concepts in a multi-label classification fashion. Through ranking concepts in descending order of predicted F_1 scores, a goal concept can be learned within few retrieval operations. Our experiments showed that NERO outperforms state-of-the-art models in 770 concept learning problems on 5 benchmark datasets w.r.t. the quality of predictions, number of explored concepts and the total computational time. Equipping NERO with the standard search procedure further improves the F_1 scores across learning problems and benchmark datasets.

We believe that incorporating neural models in concept learning problems is worth pursuing further. In future, we will work on using NERO on more expressive description logics and integrating embeddings for concepts in non-myopic heuristics [10].

Acknowledgments

This work has been supported by the European Union’s Horizon Europe research and innovation programme (GA No 101070305), by the Ministry of Culture and Science of North Rhine-Westphalia within the project SAIL (GA No NW21-059D), and the Deutsche Forschungsgemeinschaft (GA No TRR 318/1 2021 – 438445824).

References

1. Baader, F., Calvanese, D., McGuinness, D., Patel-Schneider, P., Nardi, D., et al.: The description logic handbook: Theory, implementation and applications. Cambridge university press (2003)
2. Badea, L., Nienhuys-Cheng, S.H.: A refinement operator for description logics. In: *ILP* (2000)
3. Bin, S., Bühmann, L., Lehmann, J., Ngomo, A.C.N.: Towards sparql-based induction for large-scale rdf data sets. In: *ECAI* (2016)
4. Bin, S., Westphal, P., Lehmann, J., Ngonga, A.: Implementing scalable structured machine learning for big data in the sake project. In: *2017 IEEE International Conference on Big Data (Big Data)*. pp. 1400–1407. IEEE (2017)
5. Borgida, A., Franconi, E., Horrocks, I., McGuinness, D.L., Patel-Schneider, P.F.: Explaining alc subsumption. In: *ECAI*. pp. 209–213 (2000)
6. Bühmann, L., Lehmann, J., Westphal, P.: Dl-learner—a framework for inductive learning on the semantic web. *Journal of Web Semantics* **39**, 15–24 (2016)
7. Bühmann, L., Lehmann, J., Westphal, P., Bin, S.: Dl-learner structured machine learning on semantic web data. In: *Companion Proceedings of the The Web Conference 2018*. pp. 467–471. WWW ’18, International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland (2018)
8. Caruana, R.: *Multitask learning*. Springer (1998)
9. Demir, C., Moussallem, D., Heindorf, S., Ngomo, A.C.N.: Convolutional hypercomplex embeddings for link prediction. In: *Asian Conference on Machine Learning*. pp. 656–671. PMLR (2021)
10. Demir, C., Ngomo, A.: Drill-deep reinforcement learning for refinement operators in alc. *CoRR* abs/2106.15373 **223**, 224 (2021)
11. Dettmers, T., Minervini, P., Stenetorp, P., Riedel, S.: Convolutional 2d knowledge graph embeddings. In: *AAAI* (2018)
12. Fanizzi, N., d’Amato, C., Esposito, F.: Dl-foil concept learning in description logics. In: *International Conference on Inductive Logic Programming*. pp. 107–121. Springer (2008)
13. Fanizzi, N., Rizzo, G., d’Amato, C.: Boosting dl concept learners. In: *European Semantic Web Conference*. pp. 68–83. Springer (2019)
14. Fanizzi, N., Rizzo, G., d’Amato, C., Esposito, F.: Dlfoil: Class expression learning revisited. In: *European Knowledge Acquisition Workshop*. pp. 98–113. Springer (2018)
15. Goller, C., Kuchler, A.: Learning task-dependent distributed representations by backpropagation through structure. In: *Proceedings of International Conference on Neural Networks (ICNN’96)*. vol. 1, pp. 347–352. IEEE (1996)

16. Heindorf, S., Blübaum, L., Düsterhus, N., Werner, T., Golani Nandkumar, V., Demir, C., Ngonga Ngomo, A.C.: Evolearner: Learning description logics with evolutionary algorithms. In: WWW. ACM (2022)
17. Hitzler, P., Bianchi, F., Ebrahimi, M., Sarker, M.K.: Neural-symbolic integration and the semantic web. *Semantic Web* **11**(1), 3–11 (2020)
18. Hitzler, P., Krotzsch, M., Rudolph, S.: Foundations of semantic web technologies (2009)
19. Hogan, A., Blomqvist, E., Cochez, M., d’Amato, C., Melo, G.d., Gutierrez, C., Kirrane, S., Gayo, J.E.L., Navigli, R., Neumaier, S., et al.: Knowledge graphs. *ACM Computing Surveys (CSUR)* **54**(4), 1–37 (2021)
20. Horrocks, I., Patel-Schneider, P.F., Van Harmelen, F.: From shiq and rdf to owl: The making of a web ontology language. *Journal of web semantics* **1**(1), 7–26 (2003)
21. Iannone, L., Palmisano, I., Fanizzi, N.: An algorithm based on counterfactuals for concept learning in the semantic web. *Applied Intelligence* **26**(2), 139–159 (2007)
22. Kouagou, N., Heindorf, S., Demir, C., Ngomo, A.C.N.: Learning concept lengths accelerates concept learning in alc. In: Nineteenth Extended Semantic Web Conference - Research Track. Springer (2022)
23. Krötzsch, M., Simancik, F., Horrocks, I.: A description logic primer. arXiv preprint arXiv:1201.4089 (2012)
24. Lehmann, J.: DL-learner: learning concepts in description logics. *The Journal of Machine Learning Research* **10**, 2639–2642 (2009)
25. Lehmann, J.: Learning OWL class expressions, vol. 22. IOS Press (2010)
26. Lehmann, J., Auer, S., Bühmann, L., Tramp, S.: Class expression learning for ontology engineering. *Journal of Web Semantics* **9**(1), 71–81 (2011)
27. Lehmann, J., Hitzler, P.: Concept learning in description logics using refinement operators. *Machine Learning* **78**(1-2), 203 (2010)
28. Lehmann, J., et al.: DL-learner manual (2016)
29. Michel, F., Turhan, A.Y., Zarriß, B.: Efficient tbox reasoning with value restrictions—introducing the flower reasoner. In: Proceedings of the 3rd International Joint Conference on Rules and Reasoning (RuleML+ RR 2019). LNCS, Bolzano, Italy. Springer (2019)
30. Nickel, M., Murphy, K., Tresp, V., Gabrilovich, E.: A review of relational machine learning for knowledge graphs. *Proceedings of the IEEE* **104**(1), 11–33 (2015)
31. Rudolph, S.: Foundations of description logics. In: Reasoning Web International Summer School. pp. 76–136. Springer (2011)
32. Sarker, M.K., Hitzler, P.: Efficient concept induction for description logics. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 33, pp. 3036–3043 (2019)
33. Schockaert, S., Ibanez-Garcia, Y., Gutierrez-Basulto, V.: A description logic for analogical reasoning. In: Proceedings of IJCAI-21. pp. 2040–2046 (8 2021)
34. Tran, A.C., Dietrich, J., Guesgen, H.W., Marsland, S.: Parallel symmetric class expression learning. *J. Mach. Learn. Res.* **18**, 64:1–64:34 (2017)
35. Westphal, P., Vahdati, S., Lehmann, J.: A simulated annealing meta-heuristic for concept learning in description logics. In: International Conference on Inductive Logic Programming. pp. 266–281. Springer (2021)
36. Zaheer, M., Kottur, S., Ravanbakhsh, S., Poczos, B., Salakhutdinov, R.R., Smola, A.J.: Deep sets. In: Advances in Neural Information Processing Systems. vol. 30. Curran Associates, Inc. (2017)