

# EvoLearner: Learning Description Logics with Evolutionary Algorithms

Stefan Heindorf, Lukas Blübaum, Nick Düsterhus, Till Werner, Varun Nandkumar Golani, Caglar Demir, Axel-Cyrille Ngonga Ngomo

DICE Group, Department of Computer Science, Paderborn University, Paderborn, Germany  
{heindorf, caglar.demir, axel.ngonga}@upb.de, {lukasbl, nduester, wtill, vngolani}@mail.upb.de

## ABSTRACT

Classifying nodes in knowledge graphs is an important task, e.g., for predicting missing types of entities, predicting which molecules cause cancer, or predicting which drugs are promising treatment candidates. While black-box models often achieve high predictive performance, they are only post-hoc and locally explainable and do not allow the learned model to be easily enriched with domain knowledge. Towards this end, learning description logic concepts from positive and negative examples has been proposed. However, learning such concepts often takes a long time and state-of-the-art approaches provide limited support for literal data values, although they are crucial for many applications. In this paper, we propose EvoLearner—an evolutionary approach to learn concepts in  $\mathcal{ALCQ}(\mathcal{D})$ , which is the attributive language with complement ( $\mathcal{ALC}$ ) paired with qualified cardinality restrictions ( $\mathcal{Q}$ ) and data properties ( $\mathcal{D}$ ). We contribute a novel initialization method for the initial population: starting from positive examples, we perform biased random walks and translate them to description logic concepts. Moreover, we improve support for data properties by maximizing information gain when deciding where to split the data. We show that our approach significantly outperforms the state of the art on the benchmarking framework SML-Bench for structured machine learning. Our ablation study confirms that this is due to our novel initialization method and support for data properties.

## CCS CONCEPTS

• **Computing methodologies** → *Description logics; Genetic programming; Inductive logic learning; Rule learning; Machine learning.*

## KEYWORDS

Description Logics, Evolutionary Algorithms, Machine Learning

### ACM Reference Format:

Stefan Heindorf, Lukas Blübaum, Nick Düsterhus, Till Werner, Varun Nandkumar Golani, Caglar Demir, Axel-Cyrille Ngonga Ngomo. 2022. EvoLearner: Learning Description Logics with Evolutionary Algorithms. In *Proceedings of the ACM Web Conference 2022 (WWW '22)*, April 25–29, 2022, Virtual Event, Lyon, France. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3485447.3511925>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

WWW '22, April 25–29, 2022, Virtual Event, Lyon, France

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9096-5/22/04...\$15.00

<https://doi.org/10.1145/3485447.3511925>

## 1 INTRODUCTION

While deep learning has become popular over the last decade, its predictions are hardly explainable to humans, and attempts of peeking into the “black box” often focus on *local* explainability rather than *global* explainability, i.e., on single predictions rather than the model as a whole [15]. However, when classifying nodes in knowledge graphs, e.g., predicting the type of entities, predicting which molecules cause cancer, or predicting which drugs are promising treatment candidates, it is often desirable to understand the model as a whole. This enables experts to double-check the model’s plausibility and adapt the learned rules to reflect domain knowledge. Toward this end, interpretable models based on description logics have been proposed [24, 48, 50]. Concepts in description logics can be directly mapped to OWL class expressions allowing to make use of the rich data ecosystem centered around the Web Ontology Language, OWL, the W3C standard at the core of the Semantic Web [19, 27]. Tools for verbalizing OWL class expressions in natural language are readily available to explain the expressions to users [30, 31, 37]. OWL class expressions can be applied to web-based knowledge graphs such as DBpedia [1], Wikidata [49], and YAGO [45], as well as to semantic data on websites, e.g., in the schema.org ontology [18] for type prediction [cf. 38, 52], product classification [53], and discovery and querying of web APIs [29, 51].

State-of-the-art approaches for concept learning are based on inductive logic programming with refinement operators [11, 24, 26]: Starting from the most general concept  $\top$  (also known as Thing), they refine the concept iteratively. For example, a concept is replaced by a subconcept or by multiple subconcepts joined via logical operators such as disjunction, conjunction or negation. In these approaches, the generation of candidate concepts is almost exclusively based on the ontology, and positive and negatives examples are only used to evaluate the generated concepts. In case of large ontologies, these approaches lead to a combinatorial explosion and many generated concepts never appear in the instance data leading to an unnecessarily large search space. For example, in the DBpedia ontology, a person can have over 100 different properties and ILP approaches would try all of them including combinations thereof. However, most of these properties (and combinations thereof) rarely occur in the instance data, e.g., it is seldom stated whether a person is left-handed or right-handed (object property “handedness”) and neither is their hip size (data property “hip size”).

Moreover, many knowledge bases contain large amounts of data properties connecting entities to literals, such as numeric values, which are crucial for good predictions. However, most state-of-the-art approaches for concept learning neglect data properties. Only DL-Learner [26] and SPaCEL [48] have some rudimentary support: DL-Learner divides the value range of a data property

into bins containing approximately the same number of examples; SPaCEL considers all thresholds in a brute-force manner. Neither approach takes the distribution of positive and negative examples into account, resulting in suboptimal thresholds or long runtimes. DL-FOIL [11] does not support data properties at all.

In this work, we speed up and improve concept learning using a bottom-up approach, dubbed EvoLearner, that is based on random walks and evolutionary algorithms. Instead of refining the top concept  $\top$ , we start at the instance data by initializing the initial population of the evolutionary algorithm with biased random walks starting from the positive instances in the knowledge graph. We further refine the initial candidates by means of crossover and mutation operations. For data properties, we determine splits to maximize information gain in a way inspired by decision trees [39]. Our approach significantly outperforms the state-of-the-art approaches CELOE [24] (top-down), Aleph [33] (bottom-up) and SPaCEL [47] (hybrid) on 7 of 9 datasets available in the SML-Bench benchmarking framework [50]. Our ablation study shows that both our random-walk initialization and support for data properties significantly contribute to our strong performance. Investigating predictive performance as a function of available runtime shows that EvoLearner outperforms CELOE and SPaCEL at all time points.

In the following, we briefly discuss related work in Section 2 before introducing our EvoLearner in Section 3. Section 4 evaluates our approach and Section 5 discusses its strength and weaknesses. Section 6 concludes the paper.

## 2 RELATED WORK

Most state-of-the-art approaches [24, 26, 47] for class expression learning employ inductive logic programming (ILP) with *refinement operators*. The first ILP approaches GOLEM [32, 34], PROGOLEM [36], and Aleph [33] employed *upward* refinement operators by means of least generalization. However, as Badea and Nienhuys-Cheng [3] argue, they generate overly specific concepts that tend to overfit. Hence, they propose tackling the problem by *downward* refinement operators with subsequent approaches improving the downward refinement operators by means of heuristics [25, 26] and parallelization [47]. The latest approaches, DL-FOIL [11] and SPaCEL [48], employ a hybrid approach of *upward* and *downward* refinements. DL-FOIL, however, assumes the existence of a “perfect” concept covering exactly all positive and negative instances, which is hardly possible in realistic scenarios and causes the algorithm not to terminate. SPaCEL overcomes this problem by combining many partial descriptions. Our approach, EvoLearner, finds shorter concepts that are more likely to generalize, and in cases without a perfect solution, our approach terminates with a good approximation. In our evaluation, we show that EvoLearner outperforms SPaCEL on most datasets.

While most ILP approaches employ refinement operators, there have been some attempts in the direction of *evolutionary algorithms*. Reiser and Riddle [40] showed that the ability of evolutionary approaches to search in a global space, allowing to escape local minima, can lead to an improvement over ILP algorithms. To our knowledge, Lehmann [23] was the first to apply evolutionary algorithms to the task of concept learning in *description logic*. He combined standard genetic programming (GP) approaches with

**Table 1: Description logic constructs supported by EvoLearner. For their semantics, we refer to Lehmann et al. [26].**

Syntax	Construct	Syntax	Construct
$\mathcal{ALC}$		$\mathcal{Q}$	
$r$	abstract role	$\leq n r.C$	max. cardinality restriction
$b$	Boolean concrete role	$\geq n r.C$	min. cardinality restriction
$d$	numeric concrete role	$(\mathcal{D})$	
$\neg C$	negation	$d \leq v$	max. numeric restriction ( $v \in \mathbb{R}$ )
$C \sqcup C$	union	$d \geq v$	min. numeric restriction ( $v \in \mathbb{R}$ )
$C \sqcap C$	intersection	$b = true$	Boolean value restriction
$\exists r.C$	existential restriction	$b = false$	Boolean value restriction
$\forall r.C$	universal restriction		

genetic refinement operators. For evaluation, he employs a single, small dataset comparing his genetic refinement approach to a standard GP approach. However, we were not able to reproduce his reported results, and in a pilot study, we found that our evolutionary approach works better with standard GP operators than with genetic refinements. Divina [8] gives an overview of evolutionary concept learning in first-order logic. In contrast, we specialize in description logics and contribute a sophisticated initialization strategy based on random walks [14, 44].


## 3 EVOLUTIONARY CONCEPT LEARNING

In this section, we give a brief introduction to description logics and define the task of concept learning before introducing our novel approach EvoLearner for this task. We represent description logic concepts as abstract syntax trees which form the individuals of the evolutionary algorithms. The initial population of individuals is obtained via biased random walks originating from the positive examples. Subsequently, the population evolves from generation to generation by (1) generating offspring via crossover operations, (2) subjecting some of the offspring to mutations, and (3) selecting the fittest individuals. Data properties are handled similarly to classical decision trees by maximizing information gain [39].

*Description Logics.* Description logics [21] are widely used to express rules in knowledge bases, e.g., to express which instances belong to a class. These rules are represented by logical expressions such as the union  $A \sqcup B$ , intersection  $A \sqcap B$ , and negation  $\neg A$  of concepts  $A$  and  $B$ . While different description logics offer different trade-offs between expressiveness and reasoning complexity, in this paper, we employ the description logic  $\mathcal{ALCQ}(\mathcal{D})$  because it contains the basic logical operators of all description logics ( $\mathcal{ALC}$ ) as well as cardinality restrictions ( $\mathcal{Q}$ ) and support for data properties ( $\mathcal{D}$ ), which we found to be important for many real-world applications. Table 1 gives an overview of the different constructs.

*Task: Concept Learning in Description Logics.* We define the task of concept learning following Lehmann and Hitzler [26]. Let  $\mathcal{K} = (\mathcal{T}, \mathcal{A})$  be a knowledge base where  $\mathcal{T}$  refers to the terminological box (TBox) expressed in  $\mathcal{ALCQ}(\mathcal{D})$  and  $\mathcal{A}$  refers to the assertion box (ABox). The ABox describes conceptual facts  $C(x)$  where  $x \in N_I$  is an instance,<sup>1</sup> relational facts  $r(x, y)$  with  $x, y \in N_I$ , Boolean

<sup>1</sup>In this paper, we refer to description logic *individuals* as *instances* in order to distinguish them from evolutionary individuals which represent description logic concepts.



**Figure 1: Tree representation of the concept Female**  $\sqcap$   $((\exists \text{hasSibling}.\text{Parent}) \sqcup (\exists \text{married}.\text{Brother}))$  from the Family dataset.

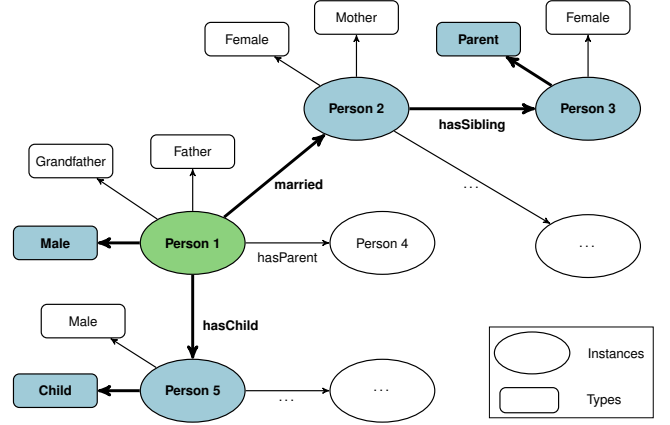
facts  $b(x, y)$  with  $x \in N_I$  and  $y \in \{true, false\}$ , and numeric facts  $d(x, y)$  with  $x \in N_I$  and  $y \in \mathbb{R}$ . Given a knowledge base  $\mathcal{K}$ , positive examples  $E^+ \subseteq N_I$ , and negative examples  $E^- \subseteq N_I$ , the objective is to find a concept  $C$  in  $\mathcal{ALCQ}(\mathcal{D})$  that covers as many of the positive examples as possible while covering as few of the negative examples as possible. In the ideal case, together with  $\mathcal{K}' := \mathcal{K} \cup \{C\}$ , it should follow that  $\mathcal{K}' \models C(e^+)$  for all  $e^+ \in E^+$  and  $\mathcal{K}' \not\models C(e^-)$  for all  $e^- \in E^-$ , i.e., the concept  $C$  provides an explanation for the positive examples but not for the negative examples. Note that it is not always possible to perfectly cover all positives and negatives in practice due to noisy training data with  $E^+ \cap E^- \neq \emptyset$ . Moreover, the perfect solution  $\{e_1^+\} \wedge \dots \wedge \{e_n^+\} \wedge \neg\{e_1^-\} \wedge \dots \wedge \neg\{e_m^-\}$  is often undesirable as it might not generalize to new, unseen data. Hence, the performance is measured with traditional metrics such as accuracy and  $F_1$ -measure [50].

*Reasoning and Graph Representation of Knowledge Base.* Following Lehmann et al. [25, 26], we assume the knowledge base  $\mathcal{K}$  to be static during concept learning and we employ a two-step reasoning process: (1) We employ the OWL reasoner Pellet to derive the instances of named classes and the relationships between them. The resulting graph allows to (approximately) answer all instance checks. (2) Following a closed-world assumption, queries are evaluated by set operations between instances of named classes.

*Tree Representation of Concepts.* As is common in genetic programming, concepts are represented as trees: inner nodes represent operators and leaves represent terminals, which can be atomic concepts in case of object properties and numbers and Booleans in case of data properties. Figure 1 shows an example.

### 3.1 Initialization via Biased Random Walks

We specifically tailored population initialization to concept learning to find good solutions as fast as possible, and we show that our method based on biased random walks significantly outperforms standard initialization methods for genetic programming, like GROW, FULL, and RAMPEDHALFHALF [20]. We seed the initial population of the evolutionary algorithm as follows: starting from a positive example in the knowledge graph, we perform a biased random walk and convert it to a concept which we add to the initial population. In this context, biased means that different outgoing



**Figure 2: Initialization of population: generating a concept via biased random walk (blue) originating from the chosen positive example (green).**

edges are taken non-uniformly with a certain probability. Each generated concept consists of an atomic concept and sequences of role restrictions to describe its type and properties, respectively. Our algorithms are available as pseudocode in the appendix in Section 7.4.

In the following, we describe our steps informally and illustrate them with the example shown in Figure 2 from the Family dataset [41]. Given positive and negative examples, the goal is to learn the concept of an Uncle:

$$\text{Male} \sqcap ((\exists \text{married}.\exists \text{hasSibling}.\text{Parent}) \sqcup (\exists \text{hasSibling}.\text{Parent})) \quad (1)$$

The positive example we use to build an individual is Person 1.

*Input and Precomputation.* Given a knowledge base  $\mathcal{K}$  and positive examples  $E^+$ , we count all their types and super types: for each example  $e \in E^+$ , we increment the count  $ct[t]$  for (super) type  $t$  except for Thing.<sup>2</sup>

- (1) *Select an example and one of its types:* Uniformly randomly pick an example  $e^+ \in E^+$ . Then select one of its types  $t$ , where each of its types is weighted according to its precomputed counts  $ct[t]$ . In a pilot study, we experimented with sampling the types uniformly randomly. However, weighting the types by their relative frequency performed slightly better. In our example, Person 1 with its type Male is picked.

**Example: Male**

- (2) *Randomly select up to  $maxT$  outgoing triples of  $e^+$ , where  $maxT$  is a hyperparameter:* Let  $R = \{r \mid \mathcal{K} \models r(e^+, \cdot)\}$  be the set of outgoing relations of  $e^+$ :<sup>3</sup>

- Uniformly randomly select up to  $maxT$  relations  $r$  from  $R$  without replacement (duplicate relations are not considered to increase variety). For each selected relation  $r$ , uniformly randomly choose an object  $o$ , yielding the triple  $r(e^+, o)$ .

<sup>2</sup>In other words,  $ct[C]$  expresses the number of instances of an atomic concept  $C$ .

<sup>3</sup>These are either object or data properties but not “type relations” since in description logics, the class of instances is expressed in the form of  $C(x)$  rather than by means of relations  $r$ .

- If less than  $maxT$  triples  $r(e^+, o)$  have been selected so far, uniformly randomly choose further triples  $r(e^+, o)$  regardless of  $r$  until either  $maxT$  triples have been selected or all outgoing triples of  $e^+$  have been selected.

Having selected up to  $maxT$  triples  $r(e^+, o)$ , their relations are added to the concept with either a union or an intersection, each with 50% probability, as follows:

- In case  $r$  is an object property, the relation is added as an existential restriction.
- In case  $r$  is a Boolean concrete role, the relation is added as a value restriction  $r = o$ , where  $o \in \{true, false\}$  represents the actual value found in the graph
- In case  $r$  is a numeric concrete role, the split  $\bar{v}_i$  closest to  $o$  is selected from a precomputed set of potential splits (see Section 3.2). Then the value restriction  $r \leq \bar{v}_i$  is added if  $o \leq \bar{v}_i$  and the value restriction  $r \geq \bar{v}_i$  otherwise.

We experimented with different values for  $maxT$  and chose  $maxT := 2$  because smaller values performed significantly worse in terms of  $F_1$ -measure, and larger values did not improve  $F_1$  but tended to increase concept length (Table 9). In our example, we walked along the married relation to Person 2 and along the hasChild yielding the following incomplete concept with placeholders (...):

**Example:** Male  $\sqcap$  ( $\exists$ married...  $\sqcap$   $\exists$ hasChild...)

- (3) *Complete incomplete subconcepts.* For each incomplete subconcept  $C$  with a placeholder, let  $o$  be the entity in the knowledge graph corresponding to the placeholder.
  - With a 50% chance, the subconcept  $C$  is extended by another triple as in Step 2.
  - Finally, subconcept  $C$  is completed by uniformly randomly selecting a type of  $o$ . If  $o$  does not have a specific type, the top concept  $\top$  (a.k.a. Thing) is selected.

In our example, the married-subconcept is extended by another relation and a type, whereas the hasChild-subconcept is extended only by a type:

**Example:** Male  $\sqcap$  (( $\exists$ married. $\exists$ hasSibling.Parent)  $\sqcap$  ( $\exists$ hasChild.Child))

Steps 1–3 generate one concept, i.e., one evolutionary individual, and are repeated until the specified population size of the evolutionary algorithm is reached. Since the initialization is random, one positive example can serve as the starting point of multiple concepts. Our objective is not to find a perfect solution immediately, but to start with a diverse population providing good crossover material. As can be seen in the generated example solution, the part  $\exists$ hasChild.Child is not part of the desired concept of an Uncle, see Solution (1), whereas the second part is:  $\exists$ married. $\exists$ hasSibling.Parent. Apart from starting at positive examples, we also experimented with starting at negative examples and negating the concept; however, this did not improve  $F_1$ -measure. The runtime of our method scales linearly with the population size, and in all our experiments with population sizes up to 1,000 individuals, the initialization took less than two seconds. Atomic concepts and relations not selected to be part of the initial population can be introduced later by our mutation operator, which deletes and inserts new concepts (see Section 3.3).

## 3.2 Data Properties

To overcome the problems of suboptimal data property thresholds and long runtimes outlined in Section 1, we adopt the idea of information gain [39], originally proposed for decision tree induction, and transfer it to concept learning: when generating value restrictions of the form  $d \leq v$  and  $d \geq v$ , where  $d$  is a numeric concrete role and  $v$  is a threshold (see Table 1), we determine  $v$  as to maximize information gain. Our evaluation results show that this approach significantly improves the performance.

Given a data property  $d$  from the knowledge base  $\mathcal{K}$  and a set  $E$  of positive and negative examples, the split maximizing information gain is computed following Quinlan [39]: Let  $V_d^E = \{v | \mathcal{K} \models d(e, v) \wedge e \in E\}$  be the set of all values for the data property  $d$  and examples  $E$ . We sort all values  $v_i \in V_d^E$ , and for each pair of adjacent values  $(v_i, v_{i+1})$ , we compute the pair's mean  $\bar{v}_i = (v_i + v_{i+1})/2$  for consideration as a potential split with information gain

$$IG(E, \bar{v}_i) = H(E) - H(E|\bar{v}_i) = H(E) - \left( \frac{|E_L|}{|E|} H(E_L) + \frac{|E_R|}{|E|} H(E_R) \right), \quad (2)$$

where  $E_i^L = \{e | (\mathcal{K} \models d(e, u) \wedge e \in E \wedge u \leq \bar{v}_i)\}$  and  $E_i^R = \{e | (\mathcal{K} \models d(e, u) \wedge e \in E \wedge u > \bar{v}_i)\}$  are the two sets produced by splitting  $E$  on  $\bar{v}_i$ . In the formula,  $H$  refers to the entropy

$$H(E) = - \sum_{c \in \{+, -\}} \Pr(E = c) \log \Pr(E = c), \quad (3)$$

with  $\Pr(E = c)$  denoting the fraction of positive/negative examples in  $E$ ; for example, with 4 positive and 3 negative examples,  $H(E) = -4/7 \log(4/7) - 3/7 \log(3/7)$ . As proven by Fayyad and Irani [12], if all examples with values  $v_i$  and  $v_{i+1}$  belong to the same class, a split between the two values cannot lead to a partition maximizing information gain. Hence, at most  $\min(|E^+|, |E^-|) < |E|$  splits need to be tried, yielding a runtime of  $O(|E| \cdot \log |E|)$  due to the sorting.

When computing the graph representation of the knowledge base (see beginning of Section 3), we precompute up to  $k$  splits  $\bar{v}_i$  per data property. We do so in a greedy fashion and take interactions between data properties into account as follows: We start with the set  $E = E^- \cup E^+$  of positive and negative examples, and for each data property  $d$  we compute the split  $\bar{v}_i$  maximizing information gain along with its sets  $E_i^L$  and  $E_i^R$  which we add to a list of example sets. Once we have computed the best split for each data property, we sort the list of example sets according to entropy in descending order. Then for each data property  $d$  and for each example set  $E$  in the list, we compute the split that maximizes information gain. The process terminates once we have found  $k$  splits per data property or there are less than  $k$  splits per property. Intuitively, our computation is equivalent to computing one decision tree for each data property and traversing these decision trees in a breadth-first search from tree level to tree level, where each level is traversed in the order of decreasing entropy until we have found  $k$  splits per data property. In our experiments, we chose  $k = 10$ , and the whole computation of splits took less than a second for all learning problems. In future work, further non-precomputed splits might be obtained on the fly.

Within our approach, we add the precomputed splits to our graph representation of the knowledge base and employ them at two points: (1) during the initialization phase described in Section 3.1, when a numeric property is selected, we select the *closest* split to the actual value of the data property; (2) during the mutation

described in Section 3.3, when a new subtree is randomly generated, a *random* split from the precomputed splits is selected. For Boolean value restrictions, we pick the actual data value in the initialization and the mutation operator randomly picks a value *true* or *false*; for cardinality restrictions, we pick a value  $\{1, 2, \dots, N\}$  where  $N$  is a hyperparameter, which was set to 5 in our experiments. In future work, this hyperparameter might be determined automatically based on the underlying data and it might vary from property to property.

### 3.3 Selection, Crossover, and Mutation

Starting from an initial population of evolutionary individuals represented as concept trees, selection, crossover, and mutation operators randomly generate new offspring based on their parents (algorithm `eaSimple[2]`).

To select the fittest individuals, a tournament selection is performed in a first pass over the population (`selectTournament`): Let  $n$  be the population size. Then  $n$  tournaments are performed, where in each tournament a small number of individuals  $k$  is chosen uniformly randomly with replacement and the fittest individual among them is selected. In a second pass over the population, a crossover operator combines the information of two parents generating two descendants. In a third pass, a mutation operator randomly modifies single individuals. In our experiments, we employ one-point crossover (`cxOnePoint`) and the uniform mutation operation (`mutUniform`).<sup>4</sup> The one-point crossover operator randomly chooses a cut point (i.e., a node in the parse tree) in both individuals and swaps the subtrees below the cut points. The mutation operation uniformly randomly selects a node in the tree and replaces it with a randomly generated tree: the tree is generated with the `RAMPEDHALFHALF` method with height 1 to 3 [20]. Following Koza [20], we employ a tournament size of  $k = 7$  and apply the crossover operation to two consecutive individuals in a population with 90% probability and the mutation operation with 10% probability. In a pilot study, we experimented with further crossover operators (e.g., `cxOnePointLeafBiased`), mutation operators (e.g., `mutShrink`, `mutNodeReplacement`, `mutInsert`, Table 8), and probabilities. Many combinations yielded good results and our chosen combination only slightly outperformed other combinations we tried. The whole process is called a generation and is repeated until the maximum number of generations is reached or a timeout occurs. The population size stays constant.

### 3.4 Fitness Function and Bloat Control

We measure the fitness of each evolutionary individual, i.e., each concept  $C$ , in terms of *accuracy* on the positive and negative examples in the training set. *Accuracy* is defined as the ratio of correct predictions among all predictions

$$\text{accuracy}(C) = \frac{|E^+ \cap R(C)| + |E^- \setminus R(C)|}{|E^+| + |E^-|}, \quad (4)$$

with  $R(C)$  denoting the individuals retrieved for concept  $C$  via our graph representation. To control bloat, i.e., overly long concepts that do not generalize to new, unseen data, we apply linear parametric

<sup>4</sup><https://deap.readthedocs.io/en/master/api/tools.html>

**Table 2: Overview of the datasets in terms of number of instances ( $N_I$ ), axioms, atomic concepts, properties, expressiveness and positive and negative examples ( $E^+$ ,  $E^-$ ). The statistics were obtained with Protégé.**

Dataset	Instances	Axioms	Atomic Concepts	Object Prop.	Data Prop.	Expressiveness	$ E^+ $	$ E^- $
Carcinogenesis	22,372	74,566	142	4	15	$\mathcal{ALC}(\mathcal{D})$	162	136
Family	202	1,829	18	4	0	$\mathcal{ALC}$	38	38
Hepatitis	6,812	79,935	14	5	12	$\mathcal{ALE}(\mathcal{D})$	206	294
Lymphography	148	2,193	49	0	0	$\mathcal{AL}$	81	67
Mammographic	975	6,808	19	3	2	$\mathcal{AL}(\mathcal{D})$	445	516
Mutagenesis	14,145	62,066	86	5	6	$\mathcal{AL}(\mathcal{D})$	13	29
NCTREC	10,209	103,070	37	9	50	$\mathcal{ALCI}(\mathcal{D})$	131	93
Premier League	11,319	2,153,818	9	13	202	$\mathcal{ALEH}(\mathcal{D})$	40	41
Pyrimidine	74	2,080	1	0	27	$\mathcal{AL}(\mathcal{D})$	20	20

parsimony pressure [28] yielding the fitness function

$$\text{fitness}(C) := \text{accuracy}(C) \cdot x - \text{length}(C), \quad (5)$$

where the  $\text{length}(C)$  of an individual denotes the number of nodes in the tree (e.g., length 7 in Figure 1) and where  $x$  is a weighting factor that can be set as a hyperparameter. In a pilot study, as per Luke and Panait [28], we experimented with different values for  $x$  between 1,024 and 8,192 and decided to use  $x = 2,048$ . We obtained similar results for a wide range of values. However, smaller values slightly decreased performance in terms of  $F_1$ -measure. Larger values were only beneficial for few datasets and learning problems.

In addition, we set a static depth limit  $d$  for the tree depth of individuals since bloat control often performs better if combined with such a limit [28]. In our experiments, we set  $d = 17$  as previously done in the literature [20, 28]. We also experimented with other bloat control methods such as proportional and double tournament [28] without significant improvements.

## 4 EVALUATION

After introducing our evaluation setup in Section 4.1, we compare our approach with state-of-the-art approaches in Section 4.2. In Sections 4.3 and 4.4, we perform an ablation analysis to identify the components that contribute most towards our high performance. Section 4.5 studies the performance based on dataset characteristics. How the learned concepts evolve from generation to generation is exemplified in Section 4.6, and Section 4.7 investigates performance as a function of runtime.

### 4.1 Evaluation Setup

Our evaluation was performed with SML-Bench [50], a benchmarking framework for structured machine learning.

*Datasets and Learning Problems.* Table 2 gives an overview of the nine datasets we employ for our evaluation. In contrast to SML-Bench’s original configuration [50], we added the Family dataset [24], which is provided as part of DL-Learner and often used to evaluate structured machine learning approaches [48]. Per dataset, the SML-Bench framework provides one learning problem, i.e., one set of positive and negative examples. For the Family dataset, the concept of an Uncle is to be learned.

*Baselines.* We compare our novel approach EvoLearner with four baselines: DL-Learner (CELOE), DL-Learner (OCEL), Aleph, and SPaCEL. Except for SPaCEL, all approaches are part of the SML-Bench framework [50]. CELOE and OCEL employ inductive logic programming with downward refinement operators with different heuristics. Aleph serves as an example of a bottom-up approach [33]. SPaCEL [48] employs a hybrid approach of upward and downward refinements. We also experimented with DL-FOIL [10, 11, 42]; however, it assumes the existence of a perfect solution, which covers all positive examples and none of the negative examples. Since almost none of the SML-Bench learning problems have such a perfect solution with a reasonable length, DL-FOIL did not terminate and we excluded it from our experiments. Additionally, we experimented with Metagol [35]. However, Metagol did not terminate within our timeout for any of the datasets, corroborating previous findings that it has problems with larger datasets and that finding a suitable set of metarules for each dataset is a challenging, unsolved task [4, 9]. Hence, we excluded it from our experiments.

*Hyperparameters.* As outlined in Section 3, we employ a population size of 800 with 200 generations, a tournament size of 7, a crossover probability of 90%, a mutation probability of 10%, and a static depth limit of 17.

*Cross-validation and Evaluation Metrics.* As SML-Bench [50], we employ 10-fold cross-validation, i.e., all positive and negative examples for each learning problem are randomly shuffled and split into 10 folds. In each round, the examples of one fold are used as the validation set and the other folds are used as the training set. Finally, the performance on all validation folds is averaged. As commonly done [50], we evaluate the predictive performance of the approaches in terms of accuracy and  $F_1$ -measure, and the explainability in terms of concept length [48]. Following SML-Bench, we employ a timeout of 5 minutes per fold.

*Reproducibility.* To ensure the reproducibility of our experiments, the code underlying our research is publicly available.<sup>5</sup> It enables those who wish to follow up on our work to replicate the plots and performance values reported. We implemented EvoLearner in Python 3.6.9 with the evolutionary framework DEAP [13] 1.3.1 and the ontology framework Owlready2 [22] 0.29. For evaluation, we employ SML-Bench [50] 0.3.0, DL-Learner [24] 1.4.0, and the latest version of SPaCEL [48]. Our experiments were run on a machine with 32 GB memory and an Intel Core i7-9750H with 2.6 GHz.

## 4.2 Evaluation of EvoLearner

Table 3 evaluates our EvoLearner approach on the benchmarking datasets and compares it with state-of-the-art baselines. Our approach outperforms each of the baselines on at least 7 out of 9 datasets in terms of accuracy and  $F_1$ -measure. In each of the 8 pairwise comparisons of EvoLearner with a baseline according to  $F_1$ -measure or accuracy, we can reject the null hypothesis that both approaches are equally good with  $p < 0.05$  (according to the non-parametric Wilcoxon signed-rank test). Looking at the lengths of learned concepts, we find that our approach generates much shorter concepts than SPaCEL and slightly longer concepts than CELOE

<sup>5</sup><https://github.com/dice-group/EvoLearner>

**Table 3: Evaluation results of EvoLearner and four state-of-the-art baselines in terms of accuracy,  $F_1$ -measure, and concept length obtained via 10-fold cross-validation on 9 learning problems with a maximum execution time of 5 minutes per fold. In addition to the mean value, we report the folds’ standard deviation. EvoLearner is initialized via random walks and supports data properties. The results of EvoLearner and SPaCEL are averaged over 3 runs since they contain randomness.**

Accuracy					
Learn. Problem	EvoLearner (ours)	DL-Learner (CELOE)	DL-Learner (OCEL)	Aleph	SPaCEL
Carcinogenesis	<b>0.64 ± 0.16</b>	0.55 ± 0.02	<i>no results</i>	0.48 ± 0.10	0.51 ± 0.10
Family	1.00 ± 0.01	0.97 ± 0.05	<b>1.00 ± 0.00</b>	—	0.97 ± 0.08
Hepatitis	<b>0.82 ± 0.06</b>	0.49 ± 0.06	<i>no results</i>	0.67 ± 0.05	<i>no results</i>
Lymphography	0.81 ± 0.12	0.70 ± 0.15	<b>0.85 ± 0.09</b>	0.83 ± 0.10	0.71 ± 0.14
Mammographic	<b>0.83 ± 0.05</b>	0.49 ± 0.02	0.82 ± 0.04	0.65 ± 0.04	0.70 ± 0.05
Mutagenesis	<b>1.00 ± 0.00</b>	0.94 ± 0.13	<i>timeout</i>	0.72 ± 0.25	<b>1.00 ± 0.00</b>
NCTRER	<b>1.00 ± 0.00</b>	0.59 ± 0.03	0.94 ± 0.06	0.72 ± 0.14	<b>1.00 ± 0.00</b>
Premier League	<b>1.00 ± 0.00</b>	0.99 ± 0.04	0.85 ± 0.01	0.95 ± 0.09	0.98 ± 0.04
Pyrimidine	0.90 ± 0.15	0.82 ± 0.17	0.85 ± 0.24	<b>0.95 ± 0.16</b>	0.87 ± 0.24
$F_1$ -measure					
Learn. Problem	EvoLearner (ours)	DL-Learner (CELOE)	DL-Learner (OCEL)	Aleph	SPaCEL
Carcinogenesis	0.70 ± 0.12	<b>0.71 ± 0.01</b>	<i>no results</i>	0.46 ± 0.12	0.60 ± 0.08
Family	1.00 ± 0.01	0.98 ± 0.05	<b>1.00 ± 0.00</b>	—	0.97 ± 0.11
Hepatitis	<b>0.79 ± 0.08</b>	0.61 ± 0.03	<i>no results</i>	0.38 ± 0.12	<i>no results</i>
Lymphography	0.84 ± 0.09	0.78 ± 0.10	<b>0.85 ± 0.10</b>	0.84 ± 0.09	0.75 ± 0.13
Mammographic	<b>0.81 ± 0.06</b>	0.64 ± 0.01	0.78 ± 0.08	0.48 ± 0.08	0.64 ± 0.06
Mutagenesis	<b>1.00 ± 0.00</b>	0.93 ± 0.14	<i>timeout</i>	0.43 ± 0.47	<b>1.00 ± 0.00</b>
NCTRER	<b>1.00 ± 0.00</b>	0.74 ± 0.01	0.94 ± 0.06	0.71 ± 0.18	<b>1.00 ± 0.00</b>
Premier League	<b>1.00 ± 0.00</b>	0.99 ± 0.04	0.81 ± 0.13	0.94 ± 0.11	0.98 ± 0.04
Pyrimidine	<b>0.91 ± 0.14</b>	0.84 ± 0.15	0.84 ± 0.22	0.90 ± 0.32	0.86 ± 0.29
Concept Length					
Learn. Problem	EvoLearner (ours)	DL-Learner (CELOE)	DL-Learner (OCEL)	SPaCEL	
Carcinogenesis	23.47 ± 4.10	3.90 ± 0.32	<i>no results</i>	1093.30 ± 82.39	
Family	10.87 ± 1.90	9.00 ± 0.00	13.20 ± 0.63	15.57 ± 1.68	
Hepatitis	19.77 ± 7.16	4.30 ± 0.95	<i>no results</i>	<i>no results</i>	
Lymphography	14.37 ± 7.31	9.40 ± 0.70	11.20 ± 2.10	172.30 ± 48.32	
Mammographic	20.43 ± 4.03	7.00 ± 0.00	8.00 ± 0.00	1547.13 ± 76.04	
Mutagenesis	3.00 ± 0.00	3.00 ± 0.00	<i>timeout</i>	3.00 ± 0.00	
NCTRER	3.00 ± 0.00	3.70 ± 1.16	7.00 ± 0.00	3.00 ± 0.00	
Premier League	7.00 ± 0.00	9.00 ± 0.00	5.00 ± 0.00	20.03 ± 3.75	
Pyrimidine	11.40 ± 1.61	10.60 ± 1.26	6.60 ± 1.26	19.67 ± 4.49	

and OCEL on most datasets. A manual inspection revealed that SPaCEL generates overly verbose expressions with many redundant and unnecessary parts, whereas CELOE’s and OCEL’s concepts are often incomplete. In the following, we show example solutions found by EvoLearner for the *Premier League* and *Carcinogenesis* learning problems. In case of *Premier League*, the objective is to find descriptions of soccer goalkeepers based on their statistics; for example, the concept of a goalkeeper can be characterized by the number of shots on the target conceded. In case of *Carcinogenesis*, molecules causing cancer can be characterized by certain chemical structures. Another example is the solution (1) found by EvoLearner for the *Uncle* problem.

**Table 4:  $F_1$ -measure of 10-fold cross-validation of different configurations of EvoLearner. Note that the datasets *Family* and *Lymphography* do not contain data properties.**

Learning Problem	EvoLearner (ours)	Without Rand. Walk Init.	Without Data Properties	Without Both
Carcinogenesis	<b>0.70 ± 0.12</b>	0.60 ± 0.21	0.63 ± 0.13	0.62 ± 0.13
Family	<b>1.00 ± 0.01</b>	0.87 ± 0.13	—	0.86 ± 0.14
Hepatitis	<b>0.79 ± 0.08</b>	0.67 ± 0.15	0.46 ± 0.14	0.47 ± 0.13
Lymphography	<b>0.84 ± 0.09</b>	0.83 ± 0.11	—	0.83 ± 0.09
Mammographic	<b>0.81 ± 0.06</b>	0.78 ± 0.08	0.77 ± 0.07	0.75 ± 0.06
Mutagenesis	<b>1.00 ± 0.00</b>	<b>1.00 ± 0.00</b>	0.44 ± 0.48	0.50 ± 0.51
NCTREER	<b>1.00 ± 0.00</b>	<b>1.00 ± 0.00</b>	0.74 ± 0.05	0.75 ± 0.05
Premier League	<b>1.00 ± 0.00</b>	0.98 ± 0.04	0.50 ± 0.23	0.50 ± 0.22
Pyrimidine	<b>0.91 ± 0.14</b>	0.83 ± 0.22	0.67 ± 0.00	0.67 ± 0.00

*Premier League:*

Player  $\sqcap \exists$ has\_action.(shots\_on\_target\_conceded  $\geq 2$ )

*Carcinogenesis:*

(drosophila\_srl1 = true)  $\sqcup$  (amesTestPositive = true)  $\sqcup$  ( $\geq 4$  hasStructure.Halide)  $\sqcup$  (chromaberr = false)

### 4.3 Initialization and Data Properties

To investigate the cause of EvoLearner’s high performance, we perform an ablation study and investigate the following variants of EvoLearner: (1) without random walk initialization, (2) without data properties, and (3) without both. Instead of our random walk initialization, we employ the RAMPEDHALFHALF initialization method with a maximal depth of 6, which combines the FULL and GROW methods and has been found to be one of the best default initialization methods [20].


Table 4 shows that both our novel initialization method and our support for data properties lead to an increase in performance on almost all datasets. For tackling the *Hepatitis*, *Mammographic*, *Mutagenesis*, *NCTREER* and *Pyrimidine* learning problems, sufficient support for data properties and cardinality restrictions is crucial. Unsurprisingly, on the *Lymphography* dataset, which does not contain any properties, neither our initialization nor support for data properties could boost performance. On *Mutagenesis* and *NCTREER*, we achieve an  $F_1$ -measure of 1.00 with and without our initialization technique since the optimal solution only requires finding the suitable value for one data property.

### 4.4 Variants of Biased Random Walks

We examined why the biased random walk method for initialization works so well by experimenting with different variants thereof: leaving out type or path information (see Section 3.1). In the first variant, we omitted Step (1) from the random walk method, so we did not select a type during the initialization and only selected the paths afterward. For the second variant, we omitted Step (2) and Step (3) so the resulting concepts in the initial population only consisted of the selected type in Step (1). Table 5 shows that both type and path information are important to get the highest performance. However, path information was more important than type information, especially for *Hepatitis* and *Pyrimidine*. Note that the variant without types still outperforms the variant without random walks and the variant without paths is about as good.

**Table 5:  $F_1$ -measure of 10-fold cross-validation of different variants of the random walk initialization. Note that *Lymphography* only contains types but no properties.**

Learn. Probl.	Rand. Walk Init. (EvoLearner)	Without Paths	Without Types	Without Rand. Walk Init.
Carcinogenesis	<b>0.70 ± 0.12</b>	0.61 ± 0.20	0.65 ± 0.17	0.60 ± 0.21
Family	<b>1.00 ± 0.01</b>	0.86 ± 0.14	0.90 ± 0.16	0.87 ± 0.13
Hepatitis	<b>0.79 ± 0.08</b>	0.66 ± 0.16	0.76 ± 0.13	0.67 ± 0.15
Lymphography	<b>0.84 ± 0.09</b>	<b>0.84 ± 0.09</b>	—	0.83 ± 0.11
Mammographic	<b>0.81 ± 0.06</b>	0.78 ± 0.08	0.81 ± 0.07	0.78 ± 0.08
Mutagenesis	<b>1.00 ± 0.00</b>	<b>1.00 ± 0.00</b>	<b>1.00 ± 0.00</b>	<b>1.00 ± 0.00</b>
NCTREER	<b>1.00 ± 0.00</b>	<b>1.00 ± 0.00</b>	<b>1.00 ± 0.00</b>	<b>1.00 ± 0.00</b>
Premier League	<b>1.00 ± 0.00</b>	0.98 ± 0.05	<b>1.00 ± 0.00</b>	0.98 ± 0.04
Pyrimidine	<b>0.91 ± 0.14</b>	0.82 ± 0.21	0.87 ± 0.21	0.83 ± 0.22

**Figure 3: Average  $F_1$ -measure depending on runtime over all datasets except *Hepatitis* and *Family* since not all approaches yielded results for them.**

### 4.5 Dataset Characteristics

We observed that our random walk/bottom-up initialization starting at the ABox tends to be particularly good for datasets whose solutions require long class expressions and for datasets that contain many object properties in the TBox (Table 2 shows the number of object properties, Table 3 [bottom] the concept lengths, and Table 3 [top] the performance). This can be explained as follows: Top-down approaches randomly refining the top concept  $\top$  have a hard time finding sensible concepts by chance if the class expressions are long and the TBox contains many object properties due to combinatorial explosion. On the other hand, our bottom-up initialization based on random walks only yields combinations of object properties that do appear in the ABox. As shown in Table 4, this effect is particularly pronounced on the datasets *Carcinogenesis*, *Family* and *Hepatitis*.

### 4.6 Example of Emerging Concepts

Table 6 provides an example of emerging concepts from generation to generation on the *Uncle* learning problem. From each generation’s population, the best concept according to the fitness function is shown. It can be seen that up to generation 11, concepts become more accurate on the training set while also becoming longer. Afterwards, concepts maintain their good accuracy while becoming shorter. The best concept is reached after 29 generations. This pattern of growing and shrinking concepts is typical and we observed it for many datasets and learning problems.

**Table 6: Emerging concepts for the Uncle learning problem.**

Gen. Concept	Len.	Train		Test	
		Acc.	F <sub>1</sub>	Acc.	F <sub>1</sub>
1 Brother $\sqcup$ $\exists$ married.Sister	5	0.941	0.944	0.875	0.889
9 Male $\sqcap$ (( $\exists$ married. $\exists$ hasSibling.Daughter) $\sqcup$ ( $\exists$ hasSibling. $\exists$ hasChild.( $\leq 1$ hasParent.Male)))	16	0.970	0.969	0.750	0.667
11 Male $\sqcap$ (( $\exists$ hasSibling. $\exists$ hasChild.( $\leq 1$ hasParent.Daughter)) $\sqcup$ ( $\exists$ married. $\exists$ hasSibling. $\exists$ hasChild( $\leq 2$ hasParent.Male)))	21	1.000	1.000	1.000	1.000
18 Male $\sqcap$ (( $\exists$ hasSibling. $\exists$ hasChild.Child) $\sqcup$ ( $\exists$ married. $\exists$ hasSibling. $\exists$ hasChild.Child))	15	1.000	1.000	1.000	1.000
23 Male $\sqcap$ (( $\exists$ hasSibling. $\exists$ hasChild.Child) $\sqcup$ ( $\exists$ married. $\exists$ hasSibling.Parent))	13	1.000	1.000	1.000	1.000
29 Male $\sqcap$ (( $\exists$ hasSibling.Parent) $\sqcup$ ( $\exists$ married. $\exists$ hasSibling.Parent))	11	1.000	1.000	1.000	1.000

#### 4.7 Effectiveness Depending on Runtime

In Figure 3, we examine the relationship between performance and runtime. The time for loading the datasets is not taken into account. For EvoLearner, the time starts right before the calculation of the splits and the random-walk initialization; for SPaCEL and CELOE, time starts after the initialization of the search and before the start of the refinements; for Aleph, time starts immediately as per SML-Bench. After the random-walk initialization, EvoLearner achieves an  $F_1$ -measure of 0.83 averaged over all datasets. This number increases to 0.88 after the evolutionary algorithm has run for 5 seconds. Except for *Carcinogenesis*, EvoLearner reaches 200 generations within the first 60 seconds, with further generations hardly increasing performance. A manual analysis revealed that for *NCTRER* and *Mutagenesis*, our initialization almost always generated an individual with the correct data property and split within the initial population. SPaCEL and CELOE initially predict Thing, yielding an  $F_1$ -measure of 0.65. SPaCEL finishes in under a second for most datasets, except for *Carcinogenesis*, *Lymphography* and *Mammographic*. Therefore, performance hardly improves afterwards. CELOE struggles to improve on *Carcinogenesis*, *Lymphography*, *Mammographic*, *Mutagenesis* and *NCTRER*. On *Pyrimidine*, the performance of CELOE even decreases over time due to overfitting. The jump between 2 and 3 minutes is due to an increase in  $F_1$ -measure on *Premier League*. Aleph does not provide partial solutions, and the jumps of average performance correspond to Aleph’s solution to *NCTRER* after about 10 seconds and Aleph’s solution to *Carcinogenesis* and *Premier League* after about 40 seconds.

## 5 DISCUSSION

*Open-world vs. Closed-world Reasoning.* Closed-world reasoners [21] assume every fact not explicitly stated to be true to be actually false. Open-world reasoners keep unspecified information open and require a logical consequence to be true in all conceivable states of the world. In this work, we employ a closed-world reasoner for the following reasons: (1) Our baseline systems employ closed-world reasoning. (2) Closed-world reasoners may be less surprising to new users<sup>6</sup> [43]. (3) Closed-world reasoners are a lot faster, allowing the evaluation of many more concepts in the same amount of time. Our code can be adapted to employ an open-world reasoner, which is available as part of the Owlready2 library [22].

<sup>6</sup><https://spinrdf.org/shacl-and-owl.html>

*Datatypes.* Our current implementation supports Booleans as well as numbers including double values and integers. Extending this to dates is straightforward as dates can be sorted, and thresholds can be computed in the same way as for numbers. An extension to strings can be realized by treating them as atomic concepts.

*Initialization of Baselines.* Our initialization strategy is specifically tailored towards evolutionary algorithms and our crossover and mutation operators. Applying our initialization strategy to our baselines would lead to a significantly smaller search space since concepts that are more general than our initial concepts would not be found by downward refinement operators. We leave it to future work to develop novel initialization strategies specifically tailored to inductive logic programming with refinement operators.

## 6 CONCLUSION

In this paper, we introduce EvoLearner, an approach to learn description logic concepts with evolutionary algorithms. In particular, our contributions are twofold: (1) a novel initialization method based on biased random walks, and (2) support for data properties, which maximizes information gain. Our results show that EvoLearner significantly outperforms the state of the art, and our ablation study confirms that this can be attributed to our initialization and support for data properties. An analysis of performance over runtime indicates that we find equally good concepts faster than the state of the art, and our approach often keeps improving when other approaches are already stagnating.

In future work, we will explore ways to combine our approach with knowledge graph embeddings [5, 6] and reinforcement learning [7], e.g., guiding the crossover and mutation in a promising direction. Moreover, we will further increase the expressiveness of our approach towards OWL 2 Full, i.e., *SROIQ(D)*. Applications may include fact checking [46] and vandalism detection [16, 17].

## ACKNOWLEDGMENTS

This work has been supported by the German Federal Ministry for Economic Affairs and Energy (BMWi) within the project RAKI under the grant no 01MD19012B and by the German Federal Ministry of Education and Research (BMBF) within the project DAIKIRI under the grant no 01IS19085B. We are grateful to Pamela Heidi Douglas for editing the manuscript.



## REFERENCES

- [1] Sören Auer, Christian Bizer, Georgi Kobilarov, Jens Lehmann, Richard Cyganiak, and Zachary G. Ives. 2007. DBpedia: A Nucleus for a Web of Open Data. In *ISWC/ASWC*. Springer, 722–735.
- [2] Thomas Bäck, David B Fogel, and Zbigniew Michalewicz. 2018. *Evolutionary computation 1: Basic algorithms and operators*. CRC press.
- [3] Liviu Badea and Shan-Hwei Nienhuys-Cheng. 2000. A Refinement Operator for Description Logics. In *ILP*. Springer, 40–59.
- [4] Andrew Cropper and Sebastian Dumancic. 2020. Inductive logic programming at 30: a new introduction. *CoRR abs/2008.07912* (2020).
- [5] Caglar Demir, Diego Moussallem, Stefan Heindorf, and Axel-Cyrille Ngonga Ngomo. 2021. Convolutional Hypercomplex Embeddings for Link Prediction. In *ACML*. PMLR, 656–671.
- [6] Caglar Demir and Axel-Cyrille Ngonga Ngomo. 2021. Convolutional Complex Knowledge Graph Embeddings. In *ESWC*. Springer, 409–424.
- [7] Caglar Demir and Axel-Cyrille Ngonga Ngomo. 2021. DRILL- Deep Reinforcement Learning for Refinement Operators in ALC. *CoRR abs/2106.15373* (2021).
- [8] Federico Divina. 2006. Evolutionary concept learning in First Order Logic: An overview. *AI Commun.* 19, 1 (2006), 13–33.
- [9] Sebastijan Dumancic, Alberto García-Durán, and Mathias Niepert. 2019. A Comparative Study of Distributional and Symbolic Paradigms for Relational Learning. In *IJCAI*. 6088–6094.
- [10] Nicola Fanizzi, Claudia d’Amato, and Floriana Esposito. 2008. DL-FOIL Concept Learning in Description Logics. In *ILP*. Springer, 107–121.
- [11] Nicola Fanizzi, Giuseppe Rizzo, Claudia d’Amato, and Floriana Esposito. 2018. DLFOIL: Class expression learning revisited. In *European Knowledge Acquisition Workshop*. Springer, 98–113.
- [12] Usama M Fayyad and Keki B Irani. 1992. On the handling of continuous-valued attributes in decision tree generation. *Machine learning* 8, 1 (1992), 87–102.
- [13] Félix-Antoine Fortin, François-Michel De Rainville, Marc-André Gardner, Marc Parizeau, and Christian Gagné. 2012. DEAP: evolutionary algorithms made easy. *J. Mach. Learn. Res.* 13 (2012), 2171–2175.
- [14] Agata Fronczak and Piotr Fronczak. 2009. Biased random walks in complex networks: The role of local navigation rules. *Physical Review E* 80, 1 (2009), 016107.
- [15] Riccardo Guidotti, Anna Monreale, Salvatore Ruggieri, Franco Turini, Fosca Giannotti, and Dino Pedreschi. 2019. A Survey of Methods for Explaining Black Box Models. *ACM Comput. Surv.* 51, 5 (2019), 93:1–93:42.
- [16] Stefan Heindorf, Martin Potthast, Benno Stein, and Gregor Engels. 2016. Vandalism Detection in Wikidata. In *CIKM*. ACM, 327–336.
- [17] Stefan Heindorf, Yan Scholten, Gregor Engels, and Martin Potthast. 2019. Debiasing Vandalism Detection Models at Wikidata. In *WWW*. ACM, 670–680.
- [18] André Hernich, Carsten Lutz, Ana Ozaki, and Frank Wolter. 2015. Schema.org as a Description Logic. In *IJCAI*. AAAI Press, 3048–3054.
- [19] Aidan Hogan, Eva Blomqvist, Michael Cochez, Claudia d’Amato, Gerard de Melo, Claudio Gutiérrez, Sabrina Kirrane, José Emilio Labra Gayo, Roberto Navigli, Sebastian Neumaier, Axel-Cyrille Ngonga Ngomo, Axel Polleres, Sabbir M. Rashid, Anisa Rula, Lukas Schmelzeisen, Juan F. Sequeda, Steffen Staab, and Antoine Zimmermann. 2021. Knowledge Graphs. *ACM Comput. Surv.* 54, 4 (2021), 71:1–71:37.
- [20] John R. Koza. 1992. *Genetic programming: On the programming of computers by means of natural selection*. MIT Press.
- [21] Markus Krötzsch, Frantisek Simancik, and Ian Horrocks. 2012. A Description Logic Primer. *CoRR abs/1201.4089* (2012).
- [22] Jean-Baptiste Lamy. 2017. Owlready: Ontology-oriented programming in Python with automatic classification and high level constructs for biomedical ontologies. *Artificial Intelligence in Medicine* 80 (2017), 11–28.
- [23] Jens Lehmann. 2007. Hybrid Learning of Ontology Classes. In *MLDM*. Springer, 883–898.
- [24] Jens Lehmann. 2009. DL-Learner: Learning Concepts in Description Logics. *J. Mach. Learn. Res.* 10 (2009), 2639–2642.
- [25] Jens Lehmann, Sören Auer, Lorenz Bühmann, and Sebastian Tramp. 2011. Class expression learning for ontology engineering. *J. Web Semant.* 9, 1 (2011), 71–81.
- [26] Jens Lehmann and Pascal Hitzler. 2010. Concept learning in description logics using refinement operators. *Mach. Learn.* 78, 1-2 (2010), 203–250.
- [27] Jiaqi Li, Xuan Wu, Chang Lu, Wenxing Deng, and Yizheng Zhao. 2021. Computing Views of OWL Ontologies for the Semantic Web. In *WWW*. ACM / IW3C2, 2624–2635.
- [28] Sean Luke and Liviu Panait. 2006. A Comparison of Bloat Control Methods for Genetic Programming. *Evol. Comput.* 14, 3 (2006), 309–344.
- [29] Franck Michel, Catherine Faron-Zucker, Olivier Corby, and Fabien Gandon. 2019. Enabling Automatic Discovery and Querying of Web APIs at Web Scale using Linked Data Standards. In *WWW (Companion Volume)*. ACM, 883–892.
- [30] Diego Moussallem, Dwaraknath Gnaneshwar, Thiago Castro Ferreira, and Axel-Cyrille Ngonga Ngomo. 2020. NABU - Multilingual Graph-Based Neural RDF Verbalizer. In *ISWC*. Springer, 420–437.
- [31] Diego Moussallem, René Speck, and Axel-Cyrille Ngonga Ngomo. 2020. Generating Explanations in Natural Language from Knowledge Graphs. In *Knowledge Graphs for eXplainable Artificial Intelligence*. Studies on the Semantic Web, Vol. 47. 213–241.
- [32] Stephen Muggleton. 1991. Inductive Logic Programming. *New Gener. Comput.* 8, 4 (1991), 295–318.
- [33] Stephen Muggleton. 1995. Inverse Entailment and Progol. *New Gener. Comput.* 13, 3&4 (1995), 245–286.
- [34] Stephen Muggleton and Cao Feng. 1990. Efficient Induction of Logic Programs. In *ALT*. Springer/Ohmsa, 368–381.
- [35] Stephen Muggleton, Dianhuan Lin, Niels Pahlavi, and Alireza Tamaddoni-Nezhad. 2013. Meta-interpretive learning: application to grammatical inference. *Machine Learning* 94 (2013), 25–49.
- [36] Stephen Muggleton, José Carlos Almeida Santos, and Alireza Tamaddoni-Nezhad. 2009. ProGolem: A System Based on Relative Minimal Generalisation. In *ILP*. Springer, 131–148.
- [37] Axel-Cyrille Ngonga Ngomo, Diego Moussallem, and Lorenz Bühmann. 2019. A Holistic Natural Language Generation Framework for the Semantic Web. In *RANLP*. 819–828.
- [38] Heiko Paulheim and Christian Bizer. 2013. Type Inference on Noisy RDF Data. In *ISWC*, Vol. 8218. Springer, 510–525.
- [39] J. Ross Quinlan. 1986. Induction of Decision Trees. *Mach. Learn.* 1, 1 (1986), 81–106.
- [40] Philip GK Reiser and Patricia J Riddle. 1999. Evolution of logic programs: part-of-speech tagging. In *Cat. No. 99TH8406*, Vol. 2. IEEE, 1338–1345.
- [41] Bradley L. Richards and Raymond J. Mooney. 1995. Automated Refinement of First-Order Horn-Clause Domain Theories. *Mach. Learn.* 19, 2 (1995), 95–131.
- [42] Giuseppe Rizzo, Nicola Fanizzi, and Claudia d’Amato. 2020. Class expression induction as concept space exploration: From DL-Foil to DL-Focl. *Future Gener. Comput. Syst.* 108 (2020), 256–272.
- [43] Alan Ruttenberg, Jonathan Rees, and Joanne Luciano. 2005. Experience Using OWL DL for the Exchange of Biological Pathway Information. In *OWLED (CEUR Workshop Proceedings, Vol. 188)*.
- [44] Roberta Sinatra, Jesús Gómez-Gardenes, Renaud Lambiotte, Vincenzo Nicosia, and Vito Latora. 2011. Maximal-entropy random walks in complex networks with limited information. *Physical Review E* 83, 3 (2011), 030103.
- [45] Fabian M. Suchanek, Gjergji Kasneci, and Gerhard Weikum. 2007. Yago: a core of semantic knowledge. In *WWW*. ACM, 697–706.
- [46] Zafar Habeeb Syed, Michael Röder, and Axel-Cyrille Ngonga Ngomo. 2019. Unsupervised Discovery of Corroborative Paths for Fact Validation. In *ISWC*. Springer, 630–646.
- [47] An C. Tran, Jens Dietrich, Hans W. Guesgen, and Stephen Marsland. 2012. An Approach to Parallel Class Expression Learning. In *RuleML*. Springer, 302–316.
- [48] An C. Tran, Jens Dietrich, Hans W. Guesgen, and Stephen Marsland. 2017. Parallel Symmetric Class Expression Learning. *J. Mach. Learn. Res.* 18 (2017), 64:1–64:34.
- [49] Denny Vrandečić and Markus Krötzsch. 2014. Wikidata: a free collaborative knowledgebase. *Commun. ACM* 57, 10 (2014), 78–85.
- [50] Patrick Westphal, Lorenz Bühmann, Simon Bin, Hajira Jabeen, and Jens Lehmann. 2019. SML-Bench - A benchmarking framework for structured machine learning. *Semantic Web* 10, 2 (2019), 231–245.
- [51] Dennis Wolters, Stefan Heindorf, Jonas Kirchhoff, and Gregor Engels. 2017. Linking Services to Websites by Leveraging Semantic Data. In *ICWS*. IEEE, 668–675.
- [52] Hamada M. Zahera, Stefan Heindorf, and Axel-Cyrille Ngonga Ngomo. 2021. ASSET: A Semi-supervised Approach for Entity Typing in Knowledge Graphs. In *K-CAP*. ACM, 261–264.
- [53] Ziqi Zhang and Monica Lestari Paramita. 2019. Product Classification Using Microdata Annotations. In *ISWC*, Vol. 11778. Springer, 716–732.

## 7 APPENDIX

In the following, we show the results of additional experiments we performed regarding initialization methods, mutation operators, and hyperparameters. Moreover, we provide pseudocode for our random walk initialization and our calculation of splitting values for data properties.

### 7.1 Comparison of Initialization Methods

In Table 7, we compare our random walk initialization to three widely used initialization methods [20]: GROW, FULL, and RAMPED-HALFHAF. GROW generates trees where leaves can have different depths, FULL generates trees where all leaves have the same depth, and RAMPEDHALFHAF generates half of a population’s individuals with the GROW or FULL method. We observe that the performance of the three methods is similar, with FULL having a slight edge, while our random walk initialization performs significantly better. We hypothesize that this can be attributed to the fact that our generated class expressions, i.e., combinations of atomic concepts and operators, do actually appear in the instance data while the other three methods generate random/artificial combinations unrelated to the instance data. Manual spot checks confirmed our hypothesis.

### 7.2 Comparison of Mutation Operators

Table 8 compares the performance of different mutation operators which are part of the DEAP framework [13]: mutUniform, mutShrink, mutNodeReplacement and mutInsert. mutUniform uniformly randomly selects a node in the tree and replaces it by a randomly generated tree as described in Section 3.3; mutShrink randomly selects two nodes in the tree which are on a path from the root to a leaf and replaces the subtree induced by the higher node by the subtree induced by the lower node. mutNodeReplacement randomly replaces a single node in the tree by a single random node; mutInsert randomly selects an inner node, replaces it with a random inner node, and adds the original subtree induced by the node as a child.

While the differences are small on most datasets, mutUniform performs much better than other approaches on *Hepatitis*. We attribute this to the fact that mutUniform introduces new constructs, like universal restrictions, negations, and cardinality restrictions, which mutShrink does not. Similarly, mutNodeReplacement and mutInsert can only introduce them to a limited extent by replacing or adding single nodes. These constructs seem to be particularly important for the *Hepatitis* dataset where almost all the good solutions contain them. In case of *Lymphography*, mutUniform learns long class expressions with training scores exceeding 0.9 (not shown in tables) which cannot be matched on test data (0.84 in Table 8), hinting at slight overfitting.

### 7.3 Comparison of maxT Settings

The hyperparameter *maxT* considers the number of outgoing triples of positive examples (c.f., Section 3.1 and Section 7.4). Table 9 evaluates different values of the *maxT* parameter. It shows that a value of two is needed since setting *maxT* := 1 yielded worse  $F_1$ -measures. However, increasing the parameter beyond *maxT* := 2 did not significantly increase the  $F_1$ -measure and resulted in longer solutions.

**Table 7: Different initialization methods in terms of  $F_1$ -measure for 10-fold cross-validation.**

Learning Problem	Rand. Walk	GROW	FULL	RAMPED
Carcinogenesis	<b>0.70 ± 0.12</b>	0.61 ± 0.20	0.63 ± 0.19	0.60 ± 0.21
Family	<b>1.00 ± 0.01</b>	0.86 ± 0.14	0.87 ± 0.14	0.87 ± 0.13
Hepatitis	<b>0.79 ± 0.08</b>	0.63 ± 0.20	0.76 ± 0.14	0.67 ± 0.15
Lymphography	<b>0.84 ± 0.09</b>	0.83 ± 0.08	<b>0.84 ± 0.09</b>	0.83 ± 0.11
Mammographic	<b>0.81 ± 0.06</b>	0.78 ± 0.08	0.78 ± 0.08	0.78 ± 0.08
Mutagenesis	<b>1.00 ± 0.00</b>	<b>1.00 ± 0.00</b>	<b>1.00 ± 0.00</b>	<b>1.00 ± 0.00</b>
NCTRRER	<b>1.00 ± 0.00</b>	<b>1.00 ± 0.00</b>	<b>1.00 ± 0.00</b>	<b>1.00 ± 0.00</b>
Premier League	<b>1.00 ± 0.00</b>	0.99 ± 0.04	0.98 ± 0.05	0.98 ± 0.04
Pyrimidine	<b>0.91 ± 0.14</b>	0.84 ± 0.22	0.84 ± 0.22	0.83 ± 0.22

**Table 8: Different mutation operators in terms of  $F_1$ -measure for 10-fold cross-validation.**

Learning Problem	mutUniform	mutShrink	mutN.Repl.	mutInsert
Carcinogenesis	<b>0.70 ± 0.12</b>	0.65 ± 0.17	0.68 ± 0.15	0.66 ± 0.14
Family	<b>1.00 ± 0.01</b>	0.96 ± 0.06	0.99 ± 0.02	0.97 ± 0.06
Hepatitis	<b>0.79 ± 0.08</b>	0.40 ± 0.21	0.53 ± 0.16	0.49 ± 0.23
Lymphography	<b>0.84 ± 0.09</b>	0.81 ± 0.09	0.81 ± 0.09	0.83 ± 0.09
Mammographic	<b>0.81 ± 0.06</b>	0.80 ± 0.07	0.81 ± 0.07	0.80 ± 0.06
Mutagenesis	<b>1.00 ± 0.00</b>	0.97 ± 0.11	<b>1.00 ± 0.00</b>	0.97 ± 0.11
NCTRRER	<b>1.00 ± 0.00</b>	0.99 ± 0.02	<b>1.00 ± 0.00</b>	0.99 ± 0.02
Premier League	<b>1.00 ± 0.00</b>	0.98 ± 0.05	0.99 ± 0.03	0.98 ± 0.06
Pyrimidine	0.91 ± 0.14	0.91 ± 0.15	0.92 ± 0.13	<b>0.93 ± 0.14</b>

**Table 9: Different settings for *maxT* in terms of  $F_1$ -measure for 10-fold cross-validation.**

Learning Problem	maxT = 1	maxT = 2	maxT = 4	maxT = 6
Carcinogenesis	0.64 ± 0.18	<b>0.70 ± 0.12</b>	0.67 ± 0.14	0.68 ± 0.14
Family	0.93 ± 0.08	1.00 ± 0.01	0.98 ± 0.05	<b>1.00 ± 0.00</b>
Hepatitis	0.71 ± 0.15	0.79 ± 0.08	0.79 ± 0.12	<b>0.82 ± 0.10</b>
Lymphography	<b>0.84 ± 0.09</b>	<b>0.84 ± 0.09</b>	<b>0.84 ± 0.09</b>	<b>0.84 ± 0.09</b>
Mammographic	0.79 ± 0.08	0.81 ± 0.06	<b>0.82 ± 0.05</b>	0.81 ± 0.07
Mutagenesis	<b>1.00 ± 0.00</b>	<b>1.00 ± 0.00</b>	<b>1.00 ± 0.00</b>	<b>1.00 ± 0.00</b>
NCTRRER	<b>1.00 ± 0.00</b>	<b>1.00 ± 0.00</b>	<b>1.00 ± 0.00</b>	<b>1.00 ± 0.00</b>
Premier League	<b>1.00 ± 0.00</b>	<b>1.00 ± 0.00</b>	1.00 ± 0.02	<b>1.00 ± 0.00</b>
Pyrimidine	0.85 ± 0.21	<b>0.91 ± 0.14</b>	0.91 ± 0.15	0.87 ± 0.17

### 7.4 Pseudocode

Algorithm 1 depicts the main loop of our random walk initialization. The PopulationFromExamples method takes a set of positive examples  $E^+$  to generate the initial population. Therefore, it randomly selects positive examples  $e^+ \in E^+$  and calls ConceptFromExample, which performs a biased random walk on  $e^+$  to generate a concept. Algorithm 2 turns a role/object pair  $(r, o)$  into an existential restriction in case of an object property. In case of a data property, it generates a Boolean value restriction or a min/max numeric value restriction. Algorithm 3 shows the pseudocode for the calculation of splitting values for data properties. If not otherwise specified, sampling in the algorithms is performed uniformly at random.

```

1 Input: Knowledge base  $\mathcal{K}$ , Positive examples  $E^+$ , Max.
  adjacent triples  $maxT$ , Set of splits for each data property
2 Output: Set of concepts population
3 Function PopulationFromExamples( $\mathcal{K}, E^+, maxT,$ 
  splits):
4   Compute the frequencies  $ct[C]$  of the types (=atomic
  concepts)  $C$  of the positive examples  $E^+$ 
5   population = []
6   while  $len(population) < POPULATIONSIZE$  do
7     Uniformly randomly pick  $e^+ \in E^+$ 
8     population.append(
  ConceptFromExample( $\mathcal{K}, e^+, ct, maxT, splits$ ))
9   end
10 return population

11 Input: Knowledge base  $\mathcal{K}$ , Positive example  $e^+$ , Type
  frequencies  $ct$ , Max. adjacent triples  $maxT$ , Set of splits for
  each data property
12 Output: Concept conc
13 Function ConceptFromExample( $\mathcal{K}, e^+, ct, maxT, splits$ ):
14   Sample an atomic concept  $C$  from  $\{C \mid \mathcal{K} \models C(e^+)\}$ 
  weighted by the frequencies  $ct$  // Step (1)
15   conc =  $C$ 
16    $R = \{r \mid \mathcal{K} \models r(e^+, \cdot)\}$  // Step (2)
17   Sample up to  $maxT$  roles from  $R$  without replacement
18   role_obj =  $\emptyset$ 
19   for each  $r \in roles$  do
20     Sample  $o$  from  $\{o \mid \mathcal{K} \models r(e^+, o)\}$ 
21     role_obj = role_obj  $\cup \{(r, o)\}$ 
22   end
23   if  $|role\_obj| < maxT$  then
24      $num = maxT - |role\_obj|$ 
25     Sample up to  $num$  new role_obj from
   $\{(r, o) \mid \mathcal{K} \models r(e^+, o) \wedge r \in R\}$  without replacement
26     role_obj = role_obj  $\cup new\_role\_obj$ 
27   end
28   for each  $(r, o) \in role\_obj$  do // Step (3)
29     Sample an operator  $op$  from  $\{\sqcup, \sqcap\}$ 
30     if  $type(r) = DataProperty$  or  $\{s \mid \mathcal{K} \models s(o, \cdot)\} = \emptyset$ 
  or  $random() < 0.5$  then
31       conc = conc  $op$  RoleObjToConc( $\mathcal{K}, (r, o), splits$ )
32     else
33       Sample a role  $s$  from  $\{s \mid \mathcal{K} \models s(o, \cdot)\}$ 
34       Sample  $v$  from  $\{(v \mid \mathcal{K} \models s(o, v) \wedge v \neq e^+)\}$ 
35       conc = conc  $op$   $\exists r.RoleObjToConc(\mathcal{K}, (s, v),$ 
  splits)
36   end
37 end
38 return conc

```

**Algorithm 1:** Generates the initial population by creating concepts from positives examples via biased random walks.

```

1 Input: Knowledge base  $\mathcal{K}$ , Role/Object Pair  $(r, o)$ , Set of
  splits for each data property
2 Output: Concept conc
3 Function RoleObjToConc( $\mathcal{K}, (r, o), splits$ ):
4   if  $type(o) = Individual$  then
5     Sample an atomic concept  $D$  from  $\{D \mid \mathcal{K} \models D(o)\}$ ,
  if the set is empty take Thing
6     conc =  $(\exists r.D)$ 
7   else if  $type(o) = Boolean$  then
8     conc =  $(r = o)$ 
9   else if  $type(o) = numeric$  then
10    Find the split  $v \in splits[r]$  that is closest to  $o$ 
11    if  $o \geq v$  then
12      conc =  $(r \geq v)$ 
13    else
14      conc =  $(r \leq v)$ 
15    end
16 return conc

```

**Algorithm 2:** Turns the role/object pair into a concept.

```

1 Input: Knowledge base  $\mathcal{K}$ , Data properties  $P$ , Number of
  splits  $k$ , Positive examples  $E^+$ , Negative examples  $E^-$ 
2 Output: Set of splits for each data property  $d \in P$ 
3 Function CalculateSplits( $\mathcal{K}, P, k, E^+, E^-$ ):
4   For each  $d \in P$  set splits[ $d$ ] =  $\emptyset$ 
5   current_sets =  $\{E^+ \cup E^-\}$ 
6   while  $P \neq \emptyset$  and current_sets  $\neq \emptyset$  do
7     next_level =  $\emptyset$ 
8     for each  $d \in P$  do
9       for each  $E \in current\_sets$  do
10         $V_d^E = \{v \mid \mathcal{K} \models d(e, v) \wedge e \in E\}$ 
11        Sort  $V_d^E$  in ascending order
12         $\bar{V}_d^E = \{(v_i + v_{i+1})/2 \mid i \in \{0, 1, \dots, |V_d^E| - 2\}\}$ 
13        Find the best split  $\bar{v} \in \bar{V}_d^E$  on  $E$  depending on
  the information gain
14        splits[ $d$ ] = splits[ $d$ ]  $\cup \{\bar{v}\}$ 
15        Let  $E^L$  and  $E^R$  be the sets produced by
  splitting  $E$  on  $\bar{v}$ 
16        next_level = next_level  $\cup \{E^L, E^R\}$ 
17        if  $|splits[d]| \geq k$  then
18           $P.remove(d)$ 
19          break
20        end
21      end
22    end
23    current_sets = next_level
24    Sort current_sets in descending order on the
  remaining entropy of the sets and remove sets with
  0 entropy
25  end
26 return splits

```

**Algorithm 3:** Calculation of splitting values for data properties.