

Stacked Structure Learning for Lifted Relational Neural Networks

Gustav Šourek¹ (✉), Martin Svatoš¹, Filip Železný¹, Steven Schockaert², and Ondřej Kuželka³

¹ Czech Technical University, Prague, Czech Republic
{souregus,svatoma1,zelezny}@fel.cvut.cz

² School of CS & Informatics, Cardiff University, Cardiff, UK
SchockaertS1@cardiff.ac.uk

³ Department of Computer Science, KU Leuven, Belgium
ondrej.kuzelka@kuleuven.be

Abstract. Lifted Relational Neural Networks (LRNNs) describe relational domains using weighted first-order rules which act as templates for constructing feed-forward neural networks. While previous work has shown that using LRNNs can lead to state-of-the-art results in various ILP tasks, these results depended on hand-crafted rules. In this paper, we extend the framework of LRNNs with structure learning, thus enabling a fully automated learning process. Similarly to many ILP methods, our structure learning algorithm proceeds in an iterative fashion by top-down searching through the hypothesis space of all possible Horn clauses, considering the predicates that occur in the training examples as well as invented soft concepts entailed by the best weighted rules found so far. In the experiments, we demonstrate the ability to automatically induce useful hierarchical soft concepts leading to deep LRNNs with a competitive predictive power.

1 Introduction

Lifted Relational Neural Networks (LRNNs [15]) are weighted sets of first-order rules, which are used to construct feed-forward neural networks from relational structures. A central characteristic of LRNNs is that a different neural network is constructed for each learning example, but crucially, the weights of these different neural networks are shared. This allows LRNNs to use neural networks for learning in relational domains, despite the fact that training examples may vary considerably in size and structure.

In previous work, LRNNs have been learned from hand-crafted rules. In such cases, only the weights of the first-order rules have to be learned from training data, which can be accomplished using a variant of back-propagation. The use of hand-crafted rules offers a natural way to incorporate domain knowledge in the learning process. In some applications, however, (sufficient) domain knowledge is lacking and both the rules and their weights have to be learned from data. To this end, in this paper we introduce a structure learning method for LRNNs.

Our proposed structure learning method proceeds in an iterative fashion. In each iteration, it may either learn a set of rules that intuitively correspond to a new layer of a neural network template or to learn a set of rules that intuitively correspond to creating new connections among existing layers, a strategy which we refer to as stacked structure learning. The rules that are added in a given iteration either define one of the target predicates, or they define a new predicate that may depend on predicates that were ‘invented’ at earlier layers as well as on predicates from the considered domain. Since the actual meaning of these predicates depends on both the learned rules and their associated weights, structure learning is alternated with weight learning. Intuitively, this means that the definitions of predicates defined in earlier layers can be fine-tuned based on the rules which are added to later layers.

We present experimental result which show that the resulting LRNNs perform comparably to LRNNs that have been learned from hand-crafted rules. We believe that this makes LRNNs a particularly convenient framework for learning in relational domains, without any need for prior knowledge nor for any extensive hypertuning. Somewhat surprisingly, we find that LRNNs with learned rules are often more compact than those with hand-crafted rules. Finally, we also present some initial results which suggest that the use of logical rules enable LRNNs to efficiently learn concepts which neural networks normally struggle with.

The remainder of the paper is structured as follows. In the next section, we first provide the required background on LRNNs. In Section 3, we then present the proposed structure learning method, after which we discuss our experimental results in Section 4.

2 Preliminaries

In this section, we briefly recall the LRNN framework from [15].

LRNN Structure. A lifted relational neural network (LRNN) \mathcal{N} is a set of weighted definite clauses, i.e. a set of pairs (R_i, w_i) where R_i is a definite clause and $w_i \in \mathbb{R}$. For a LRNN \mathcal{N} , we write \mathcal{N}^* to denote the corresponding set of definite clauses, i.e. $\mathcal{N}^* = \{C \mid (C, w) \in \mathcal{N}\}$. The grounding $\bar{\mathcal{N}}$ of a LRNN \mathcal{N} is defined as $\bar{\mathcal{N}} = \{(C\theta, w) \mid (C, w) \in \mathcal{N}, C\theta \in G(\mathcal{N}^*)\}$, where $G(\mathcal{N}^*)$ is the restriction of the grounding of \mathcal{N}^* to those clauses that correspond to active rules, i.e. rules whose antecedent is satisfied in the least Herbrand model of \mathcal{N}^* . The neural network corresponding to $\bar{\mathcal{N}}$ contains the following types of neurons:

- For each ground atom h occurring in $\bar{\mathcal{N}}$, there is a neuron A_h , called an *atom neuron*.
- For each ground fact $(h, w) \in \bar{\mathcal{N}}$, there is a neuron $F_{(h,w)}$, called a *fact neuron*.
- For every ground rule $(c\theta \leftarrow b_1\theta \wedge \dots \wedge b_k\theta, w) \in \bar{\mathcal{N}}$, there is a neuron $R_{(c\theta \leftarrow b_1\theta \wedge \dots \wedge b_k\theta, w)}$, called a *rule neuron*.
- For every (possibly non-ground) rule $(c \leftarrow b_1 \wedge \dots \wedge b_k, w) \in \mathcal{N}$ and every grounding $h = c\theta$ of c that occurs in \mathcal{H} , there is a neuron $Agg_{(c \leftarrow b_1 \wedge \dots \wedge b_k, w)}^h$, called an *aggregation neuron*.

Forward propagation. Intuitively, the neural network computes for each ground atom h a truth value, which is given by the output of the atom neuron A_h . To obtain these truth values, the network propagates values in a way which closely mimics the immediate consequence operator from logic programming. In particular, when using the immediate consequence operator, there are two ways in which h can become true: if h corresponds to a fact, or if h is the head of a rule whose body is already satisfied. Similarly, the inputs of the atom neuron A_h consist of the fact neurons of the form $F_{(h,w)}$ and aggregation neurons of the form $Agg_{(c \leftarrow b_1 \wedge \dots \wedge b_k, w)}^h$. The output of an atom neuron with inputs i_1, \dots, i_m is given by $g_V(i_1, \dots, i_m)$, where g_V is an activation function that maps the inputs to a real-valued output. In this paper we will use

$$g_V(b_1, \dots, b_k) = \text{sigm} \left(a \cdot \left(\sum_{i=1}^k b_i + b_0 \right) \right)$$

where sigm is the sigmoid function $\text{sigm}(x) = 1/(1+e^{-x})$. We set the parameters $a = 6$ and $b_0 = -0.5$, as g_V then closely approximates the Łukasiewicz fuzzy disjunction [7] (see right panel in Figure 1). This helps with the interpretability of LRNNs, as it means that we can intuitively think of the activation functions as logical connectives, and of LRNNs as (fuzzy) logic programs.

A fact neuron $F_{(h,w)}$ has no input and has the value w as its output. The output of the aggregation neuron $Agg_{(c \leftarrow b_1 \wedge \dots \wedge b_k, w)}^h$ intuitively expresses how strongly h can be derived using the rule $c \leftarrow b_1 \wedge \dots \wedge b_k$. The inputs of the aggregation neuron $Agg_{(c \leftarrow b_1 \wedge \dots \wedge b_k)}^h$ are all rule neurons $R_{(c\theta \leftarrow b_1\theta \wedge \dots \wedge b_k\theta, w)}$ for which $c\theta = h$. The output of this aggregation neuron is given by $w \cdot g_*(i_1, \dots, i_m)$, where i_1, \dots, i_m are its inputs, g_* is an activation function, and w is the weight of the corresponding rule. We will use

$$g_*(b_1, \dots, b_m) = \frac{1}{m} \sum_{i=1}^m b_i.$$

The rule neuron $R_{(c\theta \leftarrow b_1\theta \wedge \dots \wedge b_k\theta, w)}$ intuitively needs to fire if the atoms $b_1\theta, \dots, b_k\theta$ are all true. Accordingly, its inputs i_1, \dots, i_k are given by the atom neurons $A_{b_1\theta}, \dots, A_{b_k\theta}$, and its output is $g_\wedge(i_1, \dots, i_k, w)$, with g_\wedge a third type of activation function. In this paper we will use the activation function

$$g_\wedge(b_1, \dots, b_k) = \text{sigm} \left(a \cdot \left(\sum_{i=1}^k b_i - k + 1 + b_0 \right) \right)$$

where we set $a = 6$ and $b_0 = -0.5$, which approximates Łukasiewicz fuzzy conjunction [7] (see left panel in Figure 1).

Weight learning. In applications, we usually consider LRNNs of the form $\mathcal{N} \cup \mathcal{E}$, where \mathcal{N} is a weighted set of first-order rules and \mathcal{E} is a weighted set of ground facts. In particular, each \mathcal{E} represents an example, while \mathcal{N} acts as a

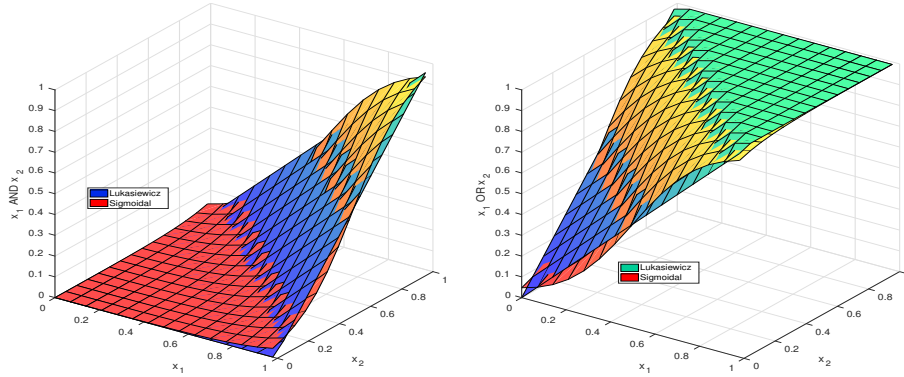


Fig. 1: An approximation of Lukasiewicz conjunction (left) and disjunction (right) by sigmoidal activation functions g_{\wedge} and g_{\vee} for the use in LRNNs.

template for constructing feed-forward neural networks, with $\overline{\mathcal{N} \cup \mathcal{E}}$ being the network corresponding to example \mathcal{E} . While the weights of \mathcal{E} are given, the weights of \mathcal{N} typically need to be learned from training data, as follows.

We are given a list of examples $\mathcal{E} = (\mathcal{E}^1, \dots, \mathcal{E}^m)$ where each \mathcal{E}^j is a LRNN, typically containing only weighted ground facts, and a list of training queries $\mathcal{Q} = (\{(q_1^1, t_1^1), \dots, (q_{k_1}^1, t_{k_1}^1)\}, \dots, \{(q_1^m, t_1^m), \dots, (q_{k_m}^m, t_{k_m}^m)\})$ where each q_i^j is a ground atom, which we call a *training query atom*, and t_i^j is its *target value*. For a query atom q_i^j , let y_i^j denote the output of the atom neuron $A_{q_i^j}$ in the ground neural network of $\overline{\mathcal{N} \cup \mathcal{E}^j}$. The goal of the learning process is to find the weights w_h of the rules (and possibly facts) in \mathcal{N} for which the loss J on the training query atoms $J(\mathcal{Q}) = \sum_{j=1}^m \sum_{i=1}^{k_j} \text{loss}(y_i^j, t_i^j)$ is minimized. This loss function is then optimized using standard stochastic gradient descent algorithm [2]. For details about weight learning of LRNNs, see [15].

3 Structure Learning

In this section we describe a structure learning algorithm for LRNNs. The algorithm receives a list of training examples and a list of training queries, and it produces a LRNN. For simplicity, we will assume that constants are only used as identifiers of objects. In particular, we will assume that attribute values are represented using unary literals, e.g. we would use *red(o)* instead of *color(o, red)*. Besides that we do not put any restrictions on the structure of the training examples.

3.1 Structure of the Learned LRNNs

The structure learning algorithm will create LRNNs having a generic “stacked” structure which we now describe. First, there are rules that define d new predicates, representing *soft clusters* [17] of unary predicates from the dataset. These

can be thought of as the first layer of the LRNN, where the weighted facts from the dataset comprise the zeroth layer. For instance, if the unary predicates in the dataset are A, B, \dots, Z then the LRNN will contain the following rules:

$$\begin{array}{llll} w_{a_1} : \alpha_1^1(X) \leftarrow A(X) & w_{b_1} : \alpha_1^1(X) \leftarrow B(X) & \dots & w_{z_1} : \alpha_1^1(X) \leftarrow Z(X) \\ w_{a_2} : \alpha_2^1(X) \leftarrow A(X) & w_{b_2} : \alpha_2^1(X) \leftarrow B(X) & \dots & w_{z_2} : \alpha_2^1(X) \leftarrow Z(X) \\ \dots & \dots & \dots & \dots \\ w_{a_d} : \alpha_d^1(X) \leftarrow A(X) & w_{b_d} : \alpha_d^1(X) \leftarrow B(X) & \dots & w_{z_d} : \alpha_d^1(X) \leftarrow Z(X) \end{array}$$

Here each α_j^i is a latent predicate representing a soft cluster, the index i denotes the layer in which it appears (in this case, the first layer) and j indexes the individual soft clusters in that level.

In general, the second layer will consist of two types of rules. First, there may be rules introducing new latent predicates. In contrast to the unary predicates that were introduced in the first layer, here the latent predicates could be also of higher arity, although in practice an upper bound will be imposed for efficiency reasons. In the body of these rules, we may find predicates from the dataset itself, or latent predicates that were introduced in the first layer. The new latent predicates introduced in these rules may then be used in the bodies of rules in subsequent layers. Second, there may also be rules that have a predicate from the dataset in their head. These will typically be rules that were learned to predict the target predicates that we want to learn.

Example 1. For instance, in datasets of molecules, unary predicates can be used to represent types of atoms, such as *carbon* or *hydrogen*. An example of a possible second layer rule is:

$$w_{p_1} : p_1(X, Y) \leftarrow \text{bond}(X, Y) \wedge \alpha_1^1(X) \wedge \alpha_2^1(Y)$$

Here p_1 is assumed to be one of the predicates from the dataset. Second layer rules that introduce a new latent predicate could look as follows.

$$\begin{array}{ll} w_{1,1}^2 : \alpha_1^2(V1, V2) \leftarrow \text{bond}(V1, V2) \wedge \alpha_1^1(V1) \wedge \alpha_1^1(V2) \\ w_{1,2}^2 : \alpha_1^2(V1, V3) \leftarrow \text{bond}(V1, V2) \wedge \text{bond}(V2, V3) \wedge \alpha_1^1(V1) \wedge \alpha_1^1(V3) \end{array}$$

The actual intuitive meaning of the predicate α_1^2 will depend on the weights $w_{1,1}^2$, $w_{1,2}^2$. For instance, if both are large enough, the (atom neurons corresponding to the) predicate will have high output whenever its arguments correspond to two atoms which are either one or two steps apart from each other in the molecule, and which have sufficiently high membership in the soft cluster α_1^1 .

Any higher layers have a similar structure to the second layer, where the n^{th} layer contains rules whose bodies only contain predicates from layers 0 to $n - 1$, and whose heads either contain a target predicate or introduce a new latent predicate.

Algorithm 1 General schema of structure learning

```

1:  $\mathcal{E} \leftarrow$  learning examples
2:  $d \leftarrow$  latent concepts' dimension
3:  $\mathcal{W}, \mathcal{V}, \mathcal{R} \leftarrow \emptyset$ 
4:  $\mathcal{R} \leftarrow \text{createLayer1Rules}(\mathcal{E}, d)$ 
5:  $\mathcal{W} \leftarrow \text{initWeights}(\mathcal{R})$ 
6:  $(\mathcal{F}, \mathcal{V}) \leftarrow \text{weightedFacts}(\mathcal{E}, \mathcal{R}, \mathcal{W})$ 
7: while  $\neg \text{StoppingCriterion}$  do
8:    $\text{bestRule} \leftarrow \text{ruleLearning}(\mathcal{F}, \mathcal{V}, \mathcal{R})$ 
9:    $\text{bestRules} \leftarrow \text{predicateInvention}(\text{bestRule})$ 
10:   $\mathcal{R} \leftarrow \mathcal{R} \cup \text{bestRules}$ 
11:   $\mathcal{W} \leftarrow \text{trainWeights}(\mathcal{R}, \mathcal{E}, \mathcal{W})$ 
12:   $(\mathcal{F}, \mathcal{V}) \leftarrow \text{weightedFacts}(\mathcal{E}, \mathcal{R}, \mathcal{W})$ 
13: end while
14: return  $(\mathcal{R}, \mathcal{W})$ 

```

3.2 Structure Learning Algorithm

The structure learning algorithm (Algorithm 1) iteratively constructs LRNNs that have the structure described in the previous section. It alternates weight learning steps with rule learning steps⁴. In the weight learning steps, the algorithm uses stochastic gradient descent to minimise the squared loss of the LRNN by optimising the weights of the rules, as described in Section 2. In the rule learning steps, the algorithm fixes the weights of all rules which define latent predicates and it searches for some *good* rule R . This rule R should be such that the squared loss of the LRNN decreases after we add R to it and after we retrain the weights of all rules with non-latent head predicates. Next we describe this algorithm in detail.

The first step of the structure learning algorithm (lines 4–5) is the construction of the first level of the LRNN, which defines the unary predicates representing soft clusters of object properties, as described in Section 3.1.

After the first step, the algorithm repeats the following procedure for a given number of iterations or until no suitable rules can be found anymore. It fixes the weights of all rules defining latent predicates (line 6). Then it runs a beam search algorithm searching through the space of possible rules⁵ (line 8). The scoring function which is used by the beam search algorithm is computed as follows. Given a rule R , the algorithm creates a copy of the current LRNN to which the given candidate rule R is added. It then optimises the log-loss of this new LRNN (which corresponds to maximum-likelihood estimation for logistic regression), training just the non-fixed weights, i.e. the weights of the rules with non-latent predicates in their heads. The score of the rule R is then defined to

⁴ Variants of this strategy are employed by many structure learning algorithms in the context of statistical relational learning, e.g. [4, 8, 5].

⁵ The space of rules is defined by two user-specified constraints: maximum rule length and maximum number of variables in a rule.

be the log-loss after training the non-fixed weights. The reason why we do not retrain all weights of the LRNN when checking score of a rule R are efficiency considerations because training the weights of the whole LRNN corresponds to training a deep neural network. After the beam search algorithm finishes, the rule R^* that it returned is added to the original LRNN.

Note that R^* contains one of the target predicates in its head. However, in addition to adding R^* , we also add a set of related rules that have latent predicates in their head (line 9), as follows. Here, we will assume for simplicity that all latent predicates have the same arity k , but the same method can still be used when the latent predicates are allowed to have different arities. Let i be the highest index such that R^* contains a latent predicate of the form α_j^i (i.e. a latent predicate from layer i) in its body, where we assume $i = 1$ if R^* does not contain any latent predicates. Then for each latent predicate α_j^{i+1} from the $(i + 1)$ -th layer, the algorithm adds to the LRNN all rules which have $\alpha_j^{i+1}(V_1, \dots, V_k)$ in the head and which can be obtained by unifying V_1, \dots, V_k with the variables in R^* . This process is illustrated in the following example.

Example 2. Revisiting the example of molecular datasets, let $R^* = p(A, B) \leftarrow bond(A, B) \wedge \alpha_2^1(A) \wedge \alpha_5^2(B)$ and let $k = 1$. Then the algorithm will add the following latent-predicate rules:

$$\begin{array}{lll} w_{1,1}^3 : & \alpha_1^3(V_1) & \leftarrow bond(V_1, B) \wedge \alpha_2^1(V_1) \wedge \alpha_5^2(B) \\ w_{1,2}^3 : & \alpha_1^3(V_1) & \leftarrow bond(A, V_1) \wedge \alpha_2^1(A) \wedge \alpha_5^2(V_1) \\ w_{2,1}^3 : & \alpha_2^3(V_1) & \leftarrow bond(V_1, B) \wedge \alpha_2^1(V_1) \wedge \alpha_5^2(B) \\ w_{2,2}^3 : & \alpha_2^3(V_1) & \leftarrow bond(A, V_1) \wedge \alpha_2^1(A) \wedge \alpha_5^2(V_1) \\ \dots & \dots & \dots \\ w_{d,1}^3 : & \alpha_d^3(V_1) & \leftarrow bond(V_1, B) \wedge \alpha_2^1(V_1) \wedge \alpha_5^2(B) \\ w_{d,2}^3 : & \alpha_d^3(V_1) & \leftarrow bond(A, V_1) \wedge \alpha_2^1(A) \wedge \alpha_5^2(V_1) \end{array}$$

Note that the algorithm has to add the new rules to the layer 3 because R^* already contained predicates from the layer 2.

After the LRNN has been extended by all these rules obtained from R^* , the weights of all the rules, including those corresponding to latent predicates, are retrained using stochastic gradient descent (line 11). Note that typically there will be some latent predicates which are not used in any rules; their weights are not considered during training. Subsequently, the algorithm again fixes the weights of the rules corresponding to the latent predicates, and repeats the same process to find an additional rule. This is repeated until a given stopping condition is met.

4 Experiments

In this section we describe the results of experiments performed with the structure learning algorithm on a real-life molecular dataset and on a difficult artificial learning problem.




Fig. 2: Comparison of crossvalidated test errors of LRNNs produced by structure learning with nFoil and kFoil learners as baselines.

4.1 Molecular Datasets

We performed experiments on 72 NCI datasets [13], each of which contains several thousands of molecules, labeled by their ability to inhibit the growth of different types of tumors. We compare the performance of the proposed LRNN structure learning method with the best previously published LRNNs, which contain large generic, yet manually constructed weighted rule sets [15]. For further comparison we include the relational learners kFOIL [10] and nFOIL [9], which respectively combine relational rule learning with support vector machines and with naive Bayes learning.

The results are shown in Figure 2 and Figure 3. The automatically learned LRNNs outperform both kFOIL and nFOIL in terms of predictive accuracy (measured using cross-validation). The learned LRNNs are also competitive with the manually constructed LRNNs from [16, 15], although they do not outperform them. They are slightly worse than the largest of the manually constructed LRNNs, based on graph patterns with 3 vertices, enumerating all possible combinations of soft cluster types of the three atoms and soft cluster types of the two bonds connecting them. Figure 4 displays statistics of the learned LRNN rule sets. These statistics show that the structure learner turned out to produce quite complex LRNNs having multiple layers of invented latent predicates.

The weights of the rules defining the latent predicates in the first layer of the LRNN can be interpreted as coordinates of a vector-space embedding of the properties (atom types in our case). In Figure 5, we plot the evolution of these embeddings as new rules are being added by the structure learning algorithm. The left panel of Figure 5 displays the evolution of the embeddings of atom types after these have been pre-trained using an unsupervised method which was originally used for statistical predicate invention in [17]. The right panel of the same figure displays the evolution of the embeddings when starting from random initialization without any unsupervised pre-training. What can be seen from these figures is how, as the model becomes more complex, the atom




Fig. 3: Comparison of test errors of LRNNs produced automatically by structure learning with 3 handcrafted LRNNs with varying lengths of chain patterns from [15].




Fig. 4: Statistics of the learned LRNN rule sets from experiments with the 72 NCI datasets. We display (i) the number of rules (including zeroth layer soft clusters), (ii) the number of conjunctive rules (patterns) learned, (iii) the average length of these rules (patterns), and (iv) the overall number of layers (depth of template).

types start to make more visible clusters. Interestingly and perhaps somewhat against intuition, the use of the unsupervised pre-training seemed to consistently decrease predictive performance (we omit details due to limited space).

4.2 A Hard Artificial Problem

We consider the following variant of graph colorability, which can be seen as a relational generalization of the problem of learning the XOR function. For a graph, where each node may take on different “shades” $\{sh_1 \dots sh_n\}$ of colors $\{col_1 \dots col_m\}$ that are not observed (i.e. it is not given to which color each shade corresponds), the task is to learn to classify graphs that are correctly colored, i.e. where each edge in the graph connects two nodes of shades of different colors. In this problem, learning a correct representation of the colors (as sets of shades) is completely decorrelated from the target, but to correctly learn to classify

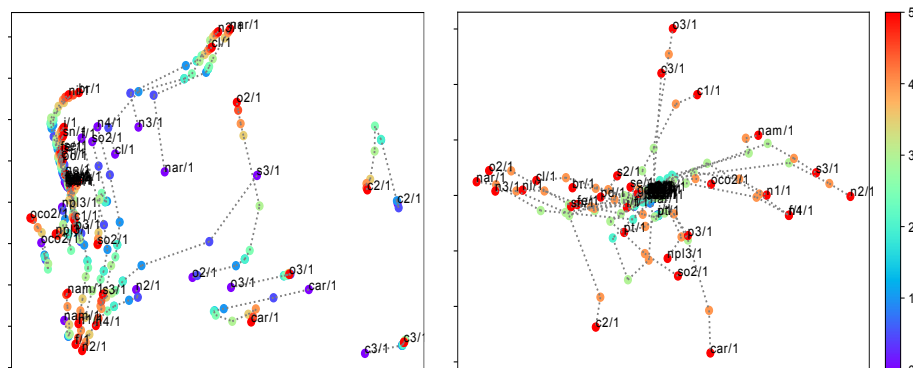


Fig. 5: PCA projection of evolution of atom embeddings during first 6 iterations (denoted by colors) of structure learning of a LRNN, with initialization based on unsupervised pre-training (left) and with completely random initialization (right).

correctly colored graphs, we need to learn some such color concepts. An ideal learned LRNN correctly solving the problem could be very compact, for instance for 3 colors it might look like

$$\begin{aligned}
 w_1 : \text{notCorrectlyColored} &\leftarrow cl_0(X), \text{edge}(X, Y), cl_0(Y) \\
 w_2 : \text{notCorrectlyColored} &\leftarrow cl_1(X), \text{edge}(X, Y), cl_1(Y) \\
 w_3 : \text{notCorrectlyColored} &\leftarrow cl_2(X), \text{edge}(X, Y), cl_2(Y)
 \end{aligned}$$

together with rules defining the color concepts cl_0 , cl_1 and cl_2 .

Initial experiments with these artificial problems showed that the structure learning algorithm is able to learn appropriate LRNNs. The rule sets were typically different from the rule set shown above but they also encoded correct solutions that were comparably short. The performance results, shown in Table 1, suggest that the LRNN structure learning method is able to efficiently produce accurate and compact solutions without extensive hyper-tuning.

5 Related Work

LRNNs are related to many older works on using neural networks for relational learning such as [1] and more recent approaches such as [14, 3]. The structure learning strategy that we employ in the methods presented in this paper is in many respects similar to structure learning methods from statistical relational learning such as [4, 8, 5]. However, what clearly distinguishes it from all these previous SRL approaches is its ability to automatically induce hierarchies of latent concepts. In this respect, it is also related to meta-interpretive learning [11]. However, meta-interpretive learning is only applicable to the learning of crisp logic programs. The structure learning approach is also related to works on refining architectures of neural networks [6, 12]. However, from these it differs in its ability to handle relational data.

#colors-#shades	majority error	training error	% of perfect solutions
2c-1s	0.5	0.025	0.9
2c-2s	0.5	0.0	1
3c-1s	0.33	0.0	1
3c-2s	0.33	0.014	0.6
3c-3s	0.33	0.111	0.4
4c-1s	0.25	0.1375	0.0
4c-2s	0.25	0.160	0.0
4c-3s	0.25	0.129	0.1
4c-4s	0.25	0.044	0.1

Table 1: Results of the structure learning algorithm (SL) on the graph coloring classification problem. Reported statistics are majority train error, training error, and ratio of cases with zero learning error. All problems were run 10 times with different random initialization seeds.

6 Conclusions and Future Work

In this paper we have introduced a method for learning the structure of LRNNs, capable of learning deep weighted rule sets with invented latent predicates. The predictive accuracies obtained by the learned LRNNs were competitive with results that we obtained in our previous work using manually constructed LRNNs. The method presented in this paper therefore has the potential to make LRNNs useful in domains where it would otherwise be difficult to come up with a rule set manually. It also makes the adoption of LRNNs by non-expert users more straightforward, as the proposed method can learn competitive LRNNs without requiring any user input (besides the dataset).

Acknowledgements GŠ, MS and FŽ acknowledge support by project no. 17-26999S granted by the Czech Science Foundation. This work was done while OK was with Cardiff University and supported by a grant from the Leverhulme Trust (RPG-2014-164). SS is supported by ERC Starting Grant 637277. Computational resources were provided by the CESNET LM2015042 and the CERIT Scientific Cloud LM2015085, provided under the programme “Projects of Large Research, Development, and Innovations Infrastructures”.

References

1. Blockeel, H., Uwents, W.: Using neural networks for relational learning. In: ICML-2004 Workshop on Statistical Relational Learning and its Connection to Other Fields, pp. 23–28 (2004)
2. Bottou, L.: Stochastic gradient descent tricks. In: Neural networks: Tricks of the trade, pp. 421–436. Springer (2012)
3. Cohen, W.W.: Tensorlog: A differentiable deductive database. arXiv preprint arXiv:1605.06523 (2016)

4. Davis, J., Burnside, E.S., de Castro Dutra, I., Page, D., Costa, V.S.: An integrated approach to learning bayesian networks of rules. In: Proceedings of the 16th European Conference on Machine Learning, pp. 84–95 (2005)
5. Dinh, Q.T., Exbrayat, M., Vrain, C.: Generative structure learning for markov logic networks based on graph of predicates. In: IJCAI Proceedings-International Joint Conference on Artificial Intelligence, vol. 22, p. 1249 (2011)
6. Fahlman, S.E., Lebiere, C.: The cascade-correlation learning architecture (1989)
7. Hájek, P.: Metamathematics of fuzzy logic, vol. 4. Springer Science & Business Media (1998)
8. Kok, S., Domingos, P.: Learning the structure of markov logic networks. In: Proceedings of the 22nd International Conference on Machine Learning, pp. 441–448 (2005)
9. Landwehr, N., Kersting, K., Raedt, L.D.: Integrating naive bayes and foil. The Journal of Machine Learning Research **8**, 481–507 (2007)
10. Landwehr, N., Passerini, A., De Raedt, L., Frasconi, P.: kFOIL: learning simple relational kernels. In: AAAI’06: Proceedings of the 21st national conference on Artificial intelligence, pp. 389–394. AAAI Press (2006)
11. Muggleton, S.H., Lin, D., Tamaddoni-Nezhad, A.: Meta-interpretive learning of higher-order dyadic datalog: predicate invention revisited. Machine Learning **100**(1), 49–73 (2015)
12. Opitz, D.W., Shavlik, J.W.: Heuristically expanding knowledge-based neural networks. In: IJCAI, pp. 1360–1365 (1993)
13. Ralaivola, L., Swamidass, S.J., Saigo, H., Baldi, P.: Graph kernels for chemical informatics. Neural Netw. **18**(8), 1093–1110 (2005)
14. Rocktäschel, T., Riedel, S.: Learning knowledge base inference with neural theorem provers. In: NAACL Workshop on Automated Knowledge Base Construction (AKBC) (2016)
15. Šourek, G., Aschenbrenner, V., Železný, F., Kuželka, O.: Lifted relational neural networks. In: Proceedings of the NIPS Workshop on Cognitive Computation: Integrating Neural and Symbolic Approaches (2015)
16. Šourek, G., Aschenbrenner, V., Železný, F., Kuželka, O.: Lifted Relational Neural Networks. arXiv preprint (2015). URL <http://arxiv.org/abs/1508.05128>
17. Šourek, G., Manandhar, S., Železný, F., Schockaert, S., Kuželka, O.: Learning predictive categories using lifted relational neural networks. In: ILP’16, 26th International Conference on Inductive Logic Programming (2016)