

NLProlog: Reasoning with Weak Unification for Question Answering in Natural Language

Leon Weber
Humboldt-Universität
zu Berlin
weberple@hu-berlin.de

Pasquale Minervini
University College London
p.minervini@ucl.ac.uk

Jannes Münchmeyer
GFZ German Research Center
for Geoscience Potsdam
munchmej@gfz-potsdam.de

Ulf Leser
Humboldt-Universität
zu Berlin
leser@informatik.hu-berlin.de

Tim Rocktäschel
University College London
t.rocktaschel@cs.ucl.ac.uk

Abstract

Rule-based models are attractive for various tasks because they inherently lead to interpretable and explainable decisions and can easily incorporate prior knowledge. However, such systems are difficult to apply to problems involving natural language, due to its linguistic variability. In contrast, neural models can cope very well with ambiguity by learning distributed representations of words and their composition from data, but lead to models that are difficult to interpret. In this paper, we describe a model combining neural networks with logic programming in a novel manner for solving multi-hop reasoning tasks over natural language. Specifically, we propose to use a Prolog prover which we extend to utilize a similarity function over pretrained sentence encoders. We fine-tune the representations for the similarity function via backpropagation. This leads to a system that can apply rule-based reasoning to natural language, and induce domain-specific rules from training data. We evaluate the proposed system on two different question answering tasks, showing that it outperforms two baselines – BIDAf (Seo et al., 2016a) and FASTQA (Weissenborn et al., 2017b) on a subset of the WIKIHOP corpus and achieves competitive results on the MEDHOP data set (Welbl et al., 2017).

1 Introduction

We consider the problem of multi-hop reasoning on natural language data. For instance, consider the statements “*Socrates was born in Athens*” and “*Athens belongs to Greece*”, and the question “*Where was Socrates born?*”. There are two possible answers following from the given statements, namely “*Athens*” and “*Greece*”. While the answer “*Athens*” follows directly from “*Socrates was born in Athens*”, the answer “*Greece*” requires the reader to combine both statements, using the knowledge

that *a person born in a city X, located in a country Y, is also born in Y*. This step of combining multiple pieces of information is referred to as *multi-hop reasoning* (Welbl et al., 2017). In the literature, such multi-hop reading comprehension tasks are frequently solved via end-to-end differentiable (deep learning) models (Sukhbaatar et al., 2015; Peng et al., 2015; Seo et al., 2016b; Raison et al., 2018; Henaff et al., 2016; Kumar et al., 2016; Graves et al., 2016; Dhingra et al., 2018). Such models are capable of dealing with the linguistic variability and ambiguity of natural language by learning word and sentence-level representations from data. However, in such models, explaining the reasoning steps leading to an answer and interpreting the model parameters to extrapolate new knowledge is a very challenging task (Doshi-Velez and Kim, 2017; Lipton, 2018; Guidotti et al., 2019). Moreover, such models tend to require large amounts of training data to generalise correctly, and incorporating background knowledge is still an open problem (Rocktäschel et al., 2015; Weissenborn et al., 2017a; Rocktäschel and Riedel, 2017; Evans and Grefenstette, 2017).

In contrast, rule-based models are easily interpretable, naturally produce explanations for their decisions, and can generalise from smaller quantities of data. However, these methods are not robust to noise and can hardly be applied to domains where data is ambiguous, such as vision and language (Moldovan et al., 2003; Rocktäschel and Riedel, 2017; Evans and Grefenstette, 2017).

In this paper, we introduce NLPROLOG, a system combining a symbolic reasoner and a rule-learning method with distributed sentence and entity representations to perform rule-based multi-hop reasoning on natural language input.¹ NLPROLOG generates partially interpretable and explain-

¹NLPROLOG and our evaluation code is available at <https://github.com/leonweber/nlprolog>.

able models, and allows for easy incorporation of prior knowledge. It can be applied to natural language without the need of converting it to an intermediate logic form. At the core of NLPROLOG is a backward-chaining theorem prover, analogous to the backward-chaining algorithm used by Prolog reasoners (Russell and Norvig, 2010b), where comparisons between symbols are replaced by differentiable similarity function between their distributed representations (Sessa, 2002). To this end, we use end-to-end differentiable sentence encoders, which are initialized with pretrained sentence embeddings (Pagliardini et al., 2017) and then fine-tuned on a downstream task. The differentiable fine-tuning objective enables us learning domain-specific logic rules – such as transitivity of the relation *is in* – from natural language data. We evaluate our approach on two challenging multi-hop Question Answering data sets, namely MEDHOP and WIKIHOP (Welbl et al., 2017).

Our main contributions are the following: *i*) We show how backward-chaining reasoning can be applied to natural language data by using a combination of pretrained sentence embeddings, a logic prover, and fine-tuning via backpropagation, *ii*) We describe how a Prolog reasoner can be enhanced with a differentiable unification function based on distributed representations (embeddings), *iii*) We evaluate the proposed system on two different Question Answering (QA) datasets, and demonstrate that it achieves competitive results in comparison with strong neural QA models while providing interpretable proofs using learned rules.

2 Related Work

Our work touches in general on weak-unification based fuzzy logic (Sessa, 2002) and focuses on multi-hop reasoning for QA, the combination of logic and distributed representations, and theorem proving for question answering.

Multi-hop Reasoning for QA. One prominent approach for enabling multi-hop reasoning in neural QA models is to iteratively update a query embedding by integrating information from embeddings of context sentences, usually using an attention mechanism and some form of recurrence (Sukhbaatar et al., 2015; Peng et al., 2015; Seo et al., 2016b; Raison et al., 2018). These models have achieved state-of-the-art results in a number of reasoning-focused QA tasks. Henaff et al. (2016) employ a differentiable memory structure

that is updated each time a new piece of information is processed. The memory slots can be used to track the state of various entities, which can be considered as a form of temporal reasoning. Similarly, the Neural Turing Machine (Graves et al., 2016) and the Dynamic Memory Network (Kumar et al., 2016), which are built on differentiable memory structures, have been used to solve synthetic QA problems requiring multi-hop reasoning. Dhingra et al. (2018) modify an existing neural QA model to additionally incorporate coreference information provided by a coreference resolution model. De Cao et al. (2018) build a graph connecting entities and apply Graph Convolutional Networks (Kipf and Welling, 2016) to perform multi-hop reasoning, which leads to strong results on WIKIHOP. Zhong et al. (2019) propose a new neural QA architecture that combines a combination of coarse-grained and fine-grained reasoning to achieve very strong results on WIKIHOP.

All of the methods above perform reasoning implicitly as a sequence of opaque differentiable operations, making the interpretation of the intermediate reasoning steps very challenging. Furthermore, it is not obvious how to leverage user-defined inference rules during the reasoning procedure.

Combining Rule-based and Neural Models.

In Artificial Intelligence literature, integrating symbolic and sub-symbolic representations is a longstanding problem (Besold et al., 2017). Our work is very related to the integration of Markov Logic Networks (Richardson and Domingos, 2006) and Probabilistic Soft Logic (Bach et al., 2017) with word embeddings, which was applied to Recognizing Textual Entailment (RTE) and Semantic Textual Similarity (STS) tasks (Garrette et al., 2011, 2014; Beltagy et al., 2013, 2014), improving over purely rule-based and neural baselines.

An area in which neural multi-hop reasoning models have been investigated is Knowledge Base Completion (KBC) (Das et al., 2016; Cohen, 2016; Neelakantan et al., 2015; Rocktäschel and Riedel, 2017; Das et al., 2017; Evans and Grefenstette, 2018). While QA could be in principle modeled as a KBC task, the construction of a Knowledge Base (KB) from text is a brittle and error prone process, due to the inherent ambiguity of natural language.

Very related to our approach are Neural Theorem Provers (NTPs) (Rocktäschel and Riedel, 2017): given a goal, its truth score is computed via a continuous relaxation of the backward-chaining rea-

soning algorithm, using a differentiable unification operator. Since the number of candidate proofs grows exponentially with the length of proofs, NTPs cannot scale even to moderately sized knowledge bases, and are thus not applicable to natural language problems in its current form. We solve this issue by using an external prover and pretrained sentence representations to efficiently discard all proof trees producing proof scores lower than a given threshold, significantly reducing the number of candidate proofs.

Theorem Proving for Question Answering.

Our work is not the first to apply theorem proving to QA problems. Angeli et al. (2016) employ a system based on Natural Logic to search a large KB for a single statement that entails the candidate answer. This is different from our approach, as we aim to learn a set of rules that combine multiple statements to answer a question.

Systems like Watson (Ferrucci et al., 2010) and COGEX (Moldovan et al., 2003) utilize an integrated theorem prover, but require a transformation of the natural language sentences to logical atoms. In the case of COGEX, this improves the accuracy of the underlying system by 30%, and increases its interpretability. While this work is similar in spirit, we greatly simplify the preprocessing step by replacing the transformation of natural language to logic with the simpler approach of transforming text to triples by using co-occurrences of named entities.

Fader et al. (2014) propose OPENQA, a system that utilizes a mixture of handwritten and automatically obtained operators that are able to parse, paraphrase and rewrite queries, which allows them to perform large-scale QA on KBs that include Open IE triples. While this work shares the same goal – answering questions using facts represented by natural language triples – we choose to address the problem of linguistic variability by integrating neural components, and focus on the combination of multiple facts by learning logical rules.

3 Background

In the following, we briefly introduce the backward chaining algorithm and unification procedure (Russell and Norvig, 2016) used by Prolog reasoners, which lies at the core of NLPROLOG. We consider Prolog programs that consists of a set of rules in

the form of Horn clauses:

$$h(f_1^h, \dots, f_n^h) \Leftarrow p_1(f_1^1, \dots, f_m^1) \wedge \dots \wedge p_B(f_1^B, \dots, f_l^B),$$

where h, p_i are predicate symbols, and f_j^i are either function (denoted in lower case) or variable (upper case) symbols. The domain of function symbols is denoted by \mathcal{F} , and the domain of predicate symbols by \mathcal{P} . $h(f_1^h, \dots, f_n^h)$ is called the *head* and $p_1(f_1^1, \dots, f_m^1) \wedge \dots \wedge p_B(f_1^B, \dots, f_l^B)$ the *body* of the rule. We call B the *body size* of the rule and rules with a body size of zero are named *atoms* (short for *atomic formula*). If an atom does not contain any variable symbols it is termed *fact*.

For simplicity, we only consider function-free Prolog in our experiments, *i.e.* Datalog (Gallaire and Minker, 1978) programs where all function symbols have arity zero and are called *entities* and, similarly to related work (Sessa, 2002; Julián-Iranzo et al., 2009), we disregard negation and disjunction. However, in principle NLPROLOG also supports functions with higher arity.

A central component in a Prolog reasoner is the *unification* operator: given two atoms, it tries to find variable substitutions that make both atoms syntactically equal. For example, the atoms $country(Greece, Socrates)$ and $country(X, Y)$ result in the following variable substitutions after unification: $\{X/Greece, Y/Socrates\}$.

Prolog uses *backward chaining* for proving assertions. Given a goal atom g , this procedure first checks whether g is explicitly stated in the KB – in this case, it can be proven. If it is not, the algorithm attempts to prove it by applying suitable rules, thereby generating subgoals that are proved next. To find applicable rules, it attempts to unify g with the heads of all available rules. If this unification succeeds, the resulting variable substitutions are applied to the atoms in the rule body: each of those atoms becomes a subgoal, and each subgoal is recursively proven using the same strategy.

For instance, the application of the rule $country(X, Y) \Leftarrow born_in(Y, X)$ to the goal $country(Greece, Socrates)$ would yield the subgoal $born_in(Socrates, Greece)$. Then the process is repeated for all subgoals until no subgoal is left to be proven. The result of this procedure is a set of rule applications and variable substitutions referred to as *proof*. Note that the number of possible proofs grows exponentially with its depth, as every rule might be used in the proof of each subgoal.

Pseudo code for weak unification can be found in Appendix A – we refer the reader to (Russell and Norvig, 2010a) for an in-depth treatment of the unification procedure.

4 NLProlog

Applying a logic reasoner to QA requires transforming the natural language paragraphs to logical representations, which is a brittle and error-prone process.

Our aim is reasoning with natural language representations in the form of triples, where entities and relations may appear under different surface forms. For instance, the textual mentions *is located in* and *lies in* express the same concept. We propose replacing the exact matching between symbols in the Prolog unification operator with a *weak unification* operator (Sessa, 2002), which allows to unify two different symbols s_1, s_2 , by comparing their representations using a differentiable similarity function $s_1 \sim_{\theta} s_2 \in [0, 1]$ with parameters θ .

With the weak unification operator, the comparison between two logical atoms results in an *unification score* resulting from the aggregation of each similarity score. Inspired by fuzzy logic t-norms (Gupta and Qi, 1991), aggregation operators are *e.g.* the minimum or the product of all scores. The result of backward-chaining with weak unification is a set of proofs, each associated with a proof score measuring the truth degree of the goal with respect to a given proof. Similarly to backward chaining, where only successful proofs are considered, in NLPROLOG the final proof success score is obtained by taking the maximum over the success scores of all found proofs. NLPROLOG combines inference based on the weak unification operator and distributed representations – such as embeddings – obtained from natural language statements.

Each natural language statement is first translated into a *triple*, where the first and third element denote the entities involved in the sentence, and the second element denotes the *textual surface pattern* connecting the entities. All elements in each triple – both the entities and the textual surface pattern – are then embedded into a vector space. These vector representations are used by the similarity function \sim_{θ} for computing similarities between two entities or two textual surface patterns and, in turn, by the backward chaining algorithm with

the weak unification operator for deriving a proof score for a given assertion. Note that the resulting proof score is fully end-to-end differentiable with respect to the model parameters θ : we can train NLPROLOG using gradient-based optimisation by back-propagating the prediction error to θ . Fig. 1 shows an outline of the model, its components and their interactions.

4.1 Triple Extraction

To transform the support documents to natural language triples, we first detect entities by performing entity recognition with SPACY (Honnibal and Montani, 2017). From these, we generate triples by extracting all entity pairs that co-occur in the same sentence and use the sentence as the predicate blinding the entities. For instance, the sentence “Socrates was born in Athens and his father was Sophronicus” is converted in the following triples: *i) (Socrates, ENT1 was born in ENT2 and his father was Sophronicus, Athens), ii) (Socrates, ENT1 was born in Athens and his father was ENT2, Sophronicus), and iii) (Athens, Socrates was born in ENT1 and his father was ENT2, Sophronicus)*. We also experimented with various Open Information Extraction frameworks (Niklaus et al., 2018): in our experiments, such methods had very low recall, which led to significantly lower accuracy values.

4.2 Similarity Computation

Embedding representations of the symbols in a triple are computed using an encoder $e_{\theta} : \mathcal{F} \cup \mathcal{P} \mapsto \mathbb{R}^d$ parameterized by θ – where \mathcal{F}, \mathcal{P} denote the sets of entity and predicate symbols, and d denotes the embedding size. The resulting embeddings are used to induce the similarity function $\sim_{\theta} : (\mathcal{F} \cup \mathcal{P})^2 \mapsto [0, 1]$, given by their cosine similarity scaled to $[0, 1]$:

$$s_1 \sim_{\theta} s_2 = \frac{1}{2} \left(1 + \frac{\mathbf{e}_{\theta}(s_1)^{\top} \mathbf{e}_{\theta}(s_2)}{\|\mathbf{e}_{\theta}(s_1)\| \cdot \|\mathbf{e}_{\theta}(s_2)\|} \right) \quad (1)$$

In our experiments, for using textual surface patterns, we use a sentence encoder composed of a static pre-trained component – namely, SENT2VEC (Pagliardini et al., 2017) – and a Multi-Layer Perceptron (MLP) with one hidden layer and Rectified Linear Unit (ReLU) activations (Jarrett et al., 2009). For encoding predicate symbols and entities, we use a randomly initialised embedding matrix. During training, both the MLP and the embedding matrix are learned via backpropagation, while the sentence encoder is kept fixed.

Additionally, we introduce a third lookup table and MLP for the predicate symbols of rules and goals. The main reason of this choice is that semantics of goal and rule predicates may differ from the semantics of fact predicates, even if they share the same surface form. For instance, the query (X, parent, Y) can be interpreted either as $(X, \text{is the parent of}, Y)$ or as $(X, \text{has parent}, Y)$, which are semantically dissimilar.

4.3 Training the Encoders

We train the encoder parameters θ on a downstream task via gradient-based optimization. Specifically, we train NLPROLOG with backpropagation using a learning from entailment setting (Muggleton and Raedt, 1994), in which the model is trained to decide whether a Prolog program \mathcal{R} entails the truth of a candidate triple $c \in C$, where C is the set of candidate triples. The objective is a model that assigns high probabilities $p(c|\mathcal{R}; \theta)$ to true candidate triples, and low probabilities to false triples. During training, we minimize the following loss:

$$L(\theta) = -\log p(a|\mathcal{R}; \theta) - \log \left(1 - \max_{c \in C \setminus \{a\}} p(c|\mathcal{R}; \theta) \right), \quad (2)$$

where $a \in C$ is the correct answer. For simplicity, we assume that there is only one correct answer per example, but an adaptation to multiple correct answers would be straight-forward, e.g. by taking the minimum of all answer scores.

To estimate $p(c|\mathcal{R}; \theta)$, we enumerate all proofs for the triple c up to a given depth D , where D is a user-defined hyperparameter. This search yields a number of proofs, each with a success score S_i . We set $p(c|\mathcal{R}; \theta)$ to be the maximum of such proof scores:

$$p(c|\mathcal{R}; \theta) = S_{\max} = \max_i S_i \in [0, 1].$$

Note that the final proof score $p(c|\mathcal{R}; \theta)$ only depends on the proof with maximum success score S_{\max} . Thus, we propose to first conduct the proof search by using a prover utilizing the similarity function induced by the current parameters \sim_{θ_i} , which allows us to compute the maximum proof score S_{\max} . The score for each proof is given by the aggregation – either using the *minimum* or the *product* functions – of the weak unification scores, which in turn are computed via the differentiable similarity function \sim_{θ} . It follows that $p(c|\mathcal{R}; \theta)$

is end-to-end differentiable, and can be used for updating the model parameters θ via Stochastic Gradient Descent.

4.4 Runtime Complexity of Proof Search

The worst case complexity vanilla logic programming is exponential in the depth of the proof (Russell and Norvig, 2010a). However, in our case, this is a particular problem because weak unification requires the prover to attempt unification between all entity and predicate symbols.

To keep things tractable, NLPROLOG only attempts to unify symbols with a similarity greater than some user-defined threshold λ . Furthermore, in the search step for one statement q , for the rest of the search, λ is set to $\max(\lambda, S)$ whenever a proof for q with success score S is found. Due to the monotonicity of the employed aggregation functions, this allows to prune the search tree without losing the guarantee to find the proof yielding the maximum success score S_{\max} , provided that $S_{\max} \geq \lambda$. We found this optimization to be crucial to make the proof search scale on the considered data sets.

4.5 Rule Learning

In NLPROLOG, the reasoning process depends on rules that describe the relations between predicates. While it is possible to write down rules involving natural language patterns, this approach does not scale. Thus, we follow Rocktäschel and Riedel (2017) and use rule templates to perform Inductive Logic Programming (ILP) (Muggleton, 1991), which allows NLPROLOG to learn rules from training data. In this setting, a user has to define a set of rules with a given structure as input. Then, NLPROLOG can learn the rule predicate embeddings from data by minimizing the loss function in Eq. (2) using gradient-based optimization methods.

For instance, to induce a rule that can model transitivity, we can use a rule template of the form $p_1(X, Z) \Leftarrow p_2(X, Y) \wedge p_3(Y, Z)$, and NLPROLOG will instantiate multiple rules with randomly initialized embeddings for p_1 , p_2 , and p_3 , and fine-tune them on a downstream task. The exact number and structure of the rule templates is treated as a hyperparameter.

Unless explicitly stated otherwise, all experiments were performed with the same set of rule templates containing two rules for each of the forms $q(X, Y) \Leftarrow p_2(X, Y)$, $p_1(X, Y) \Leftarrow p_2(Y, X)$ and $p_1(X, Z) \Leftarrow p_2(X, Y) \wedge p_3(Y, Z)$,

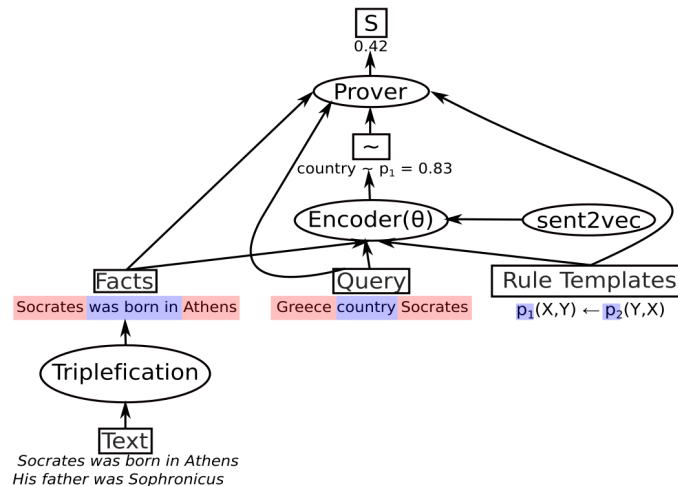


Figure 1: Overview of NLPROLOG – all components are depicted as ellipses, while inputs and outputs are drawn as squares. Phrases with red background are entities and blue ones are predicates.

where q is the query predicate. The number and structure of these rule templates can be easily modified, allowing the user to incorporate additional domain-specific background knowledge, such as $born_in(X, Z) \Leftarrow born_in(X, Y) \wedge located_in(Y, Z)$

5 Evaluation

We evaluate our method on two QA datasets, namely MEDHOP, and several subsets of WIKIHOP (Welbl et al., 2017). These data sets are constructed in such a way that it is often necessary to combine information from multiple documents to derive the correct answer.

In both data sets, each data point consists of a query $p(e, X)$, where e is an entity, X is a variable – representing the entity that needs to be predicted, C is a list of candidate entities, $a \in C$ is an answer entity and p is the query predicate. Furthermore, every query is accompanied by a set of support documents which can be used to decide which of the candidate entities is the correct answer.

5.1 MedHop

MEDHOP is a challenging multi-hop QA data set, and contains only a single query predicate. The goal in MEDHOP is to predict whether two drugs interact with each other, by considering the interactions between proteins that are mentioned in the support documents. Entities in the support documents are mapped to data base identifiers. To compute better entity representations, we reverse this mapping and replace all mentions with the

drug and proteins names gathered from DRUG-BANK (Wishart et al., 2006) and UNIPROT (Apweiler et al., 2004).

5.2 Subsets of WikiHop

To further validate the effectiveness of our method, we evaluate on different subsets of WIKIHOP (Welbl et al., 2017), each containing a single query predicate. We consider the predicates *publisher*, *developer*, *country*, and *record_label*, because their semantics ensure that the annotated answer is unique and they contain a relatively large amount of questions that are annotated as requiring multi-hop reasoning. For the predicate *publisher*, this yields 509 training and 54 validation questions, for *developer* 267 and 29, for *country* 742 and 194, and for *record_label* 2305 and 283. As the test set of WIKIHOP is not publicly available, we report scores for the validation set.

5.3 Baselines

Following Welbl et al. (2017), we use two neural QA models, namely BiDAF (Seo et al., 2016a) and FASTQA (Weissenborn et al., 2017b), as baselines for the considered WIKIHOP predicates. We use the implementation provided by the JACK² QA framework (Weissenborn et al., 2018) with the same hyperparameters as used by Welbl et al. (2017), and train a separate model for each predicate.³ To ensure that the performance of the

²<https://github.com/uclmr/jack>

³We also experimented with the AllenNLP implementation of BiDAF, available at <https://github.com/allenai/allennlp/blob/master/allennlp/>

baseline is not adversely affected by the relatively small number of training examples, we also evaluate the BiDAF model trained on the whole WIKIHOP corpus. In order to compensate for the fact that both models are extractive QA models which cannot make use of the candidate entities, we additionally evaluate modified versions which transform both the predicted answer and all candidates to vectors using the *wiki-unigrams* model of SENT2VEC (Pagliardini et al., 2017). Consequently, we return the candidate entity which has the highest cosine similarity to the predicted entity. We use the normalized version of MEDHOP for training and evaluating the baselines, since we observed that denormalizing it (as for NLPROLOG) severely harmed performance. Furthermore on MEDHOP, we equip the models with word embeddings that were pretrained on a large biomedical corpus (Pyysalo et al., 2013).

5.4 Hyperparameter Configuration

On MEDHOP we optimize the embeddings of predicate symbols of rules and query triples, as well as of entities. WIKIHOP has a large number of unique entity symbols and thus, learning their embeddings is prohibitive. Thus, we only train the predicate symbols of rules and query triples on this data set. For MEDHOP we use bigram SENT2VEC embeddings trained on a large biomedical corpus⁴, and for WIKIHOP the *wiki-unigrams* model⁵ of SENT2VEC. All experiments were performed with the same set of rule templates containing two rules for each of the forms $p(X, Y) \Leftarrow q(X, Y)$, $p(X, Y) \Leftarrow q(Y, X)$ and $p(X, Z) \Leftarrow q(X, Y) \wedge r(Y, Z)$ and set the similarity threshold λ to 0.5 and maximum proof depth to 3. We use Adam (Kingma and Ba, 2014) with default parameters.

5.5 Results

The results for the development portions of WIKIHOP and MEDHOP are shown in Table 1. For all predicates but *developer*, NLPROLOG strongly outperforms all tested neural QA models, while achieving the same accuracy as the best performing QA model on *developer*. We evaluated NLPROLOG on the hidden test set of MedHop and obtained

models/reading_comprehension/bidaf.py,
obtaining comparable results.

⁴<https://github.com/ncbi-nlp/BioSentVec>

⁵<https://drive.google.com/open?id=0B6VhzidiLvjSa19uYw1LUEkzX3c>

an accuracy of 29.3%, which is 6.1 pp better than FastQA and 18.5 pp worse than BiDAF.⁶ As the test set is hidden, we cannot diagnose the exact reason for the inconsistency with the results on the development set, but observe that FastQA suffers from a similar drop in performance.

5.6 Importance of Rules

Exemplary proofs generated by NLPROLOG for the predicates *record_label* and *country* can be found in Fig. 2.

To study the impact of the rule-based reasoning on the predictive performance, we perform an ablation experiment in which we train NLPROLOG without any rule templates. The results can be found in the bottom half of Table 1. On three of the five evaluated data sets, performance decreases markedly when no rules can be used and does not change on the remaining two data sets. This indicates that reasoning with logic rules is beneficial in some cases and does not hurt performance in the remaining ones.

5.7 Impact of Entity Embeddings

In a qualitative analysis, we observed that in many cases multi-hop reasoning was performed via aligning entities and not by applying a multi-hop rule. For instance, the proof of the statement *country(Oktabrskiy Big Concert Hall, Russia)* visualized in Figure 2, is performed by making the embeddings of the entities *Oktabrskiy Big Concert Hall* and *Saint Petersburg* sufficiently similar. To gauge the extent of this effect, we evaluate an ablation in which we remove the MLP on top of the entity embeddings. The results, which can be found in Table 1, show that fine-tuning entity embeddings plays an integral role, as the performance degrades drastically. Interestingly, the observed performance degradation is much worse than when training without rules, suggesting that much of the reasoning is actually performed by finding a suitable transformation of the entity embeddings.

5.8 Error Analysis

We performed an error analysis for each of the WIKIHOP predicates. To this end, we examined all instances in which one of the neural QA models (with SENT2VEC) produced a correct prediction

⁶Note, that these numbers are taken from Welbl et al. (2017) and were obtained with different implementations of BiDAF and FASTQA

Model	MedHop	<i>publisher</i>	<i>developer</i>	<i>country</i>	<i>recordlabel</i>
BiDAF	42.98	66.67	65.52	53.09	68.90
+ Sent2Vec	—	75.93	68.97	61.86	75.62
+ Sent2Vec + wikihop	—	74.07	62.07	66.49	78.09
FastQA	52.63	62.96	62.07	57.21	70.32
+ Sent2Vec	—	75.93	58.62	64.95	78.09
NLProlog	65.78	83.33	68.97	77.84	79.51
- rules	64.33	83.33	68.97	74.23	74.91
- entity MLP	37.13	68.52	41.38	72.16	64.66

Table 1: Accuracy scores in percent for different predicates on the development set of the respective predicates. +/- denote independent modifications to the base algorithm.

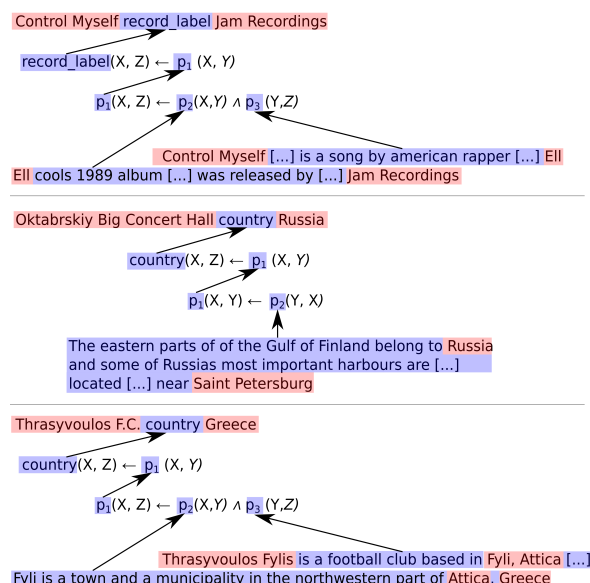


Figure 2: Example proof trees generated by NLPROLOG, showing a combination of multiple rules. Entities are shown in red and predicates in blue. Note, that entities do not need to match exactly. The first and third proofs were obtained without the entity MLP (as described in Section 5.7), while the second one was obtained in the full configuration of NLPROLOG.

and NLPROLOG did not, and labeled them with predefined error categories. Of the 55 instances, 49% of the errors were due to NLPROLOG unifying the wrong entities, mainly because of an over-reliance on heuristics, such as predicting a record label if it is from the same country as the artist. In 25% of the cases, NLPROLOG produced a correct prediction, but another candidate was defined as the answer. In 22% the prediction was due to an error in predicate unification, *i.e.* NLPROLOG identified the correct entities, but the sentence did not express the target relation. Furthermore, we performed an evaluation on all problems of the studied WIKI-

HOP predicates that were unanimously labeled as containing the correct answer in the support texts by Welbl et al. (2017). On this subset, the micro-averaged accuracy of NLPROLOG shows an absolute increase of 3.08 pp, while the accuracy of BiDAF (FASTQA) augmented with SENT2VEC decreases by 3.26 (3.63) pp. We conjecture that this might be due to NLPROLOG’s reliance on explicit reasoning, which could make it less susceptible to spurious correlations between the query and supporting text.

6 Discussion and Future Work

We proposed NLPROLOG, a system that is able to perform rule-based reasoning on natural language, and can learn domain-specific rules from data. To this end, we proposed to combine a symbolic prover with pretrained sentence embeddings, and to train the resulting system using backpropagation. We evaluated NLPROLOG on two different QA tasks, showing that it can learn domain-specific rules and produce predictions which outperform those of the two strong baselines BiDAF and FASTQA in most cases.

While we focused on a subset of First Order Logic in this work, the expressiveness of NLPROLOG could be extended by incorporating a different symbolic prover. For instance, a prover for temporal logic (Orgun and Ma, 1994) would allow to model temporal dynamics in natural language. We are also interested in incorporating future improvements of symbolic provers, triple extraction systems and pretrained sentence representations to further enhance the performance of NLPROLOG. Additionally, it would be interesting to study the behavior of NLPROLOG in the presence of multiple WIKIHOP query predicates.

Acknowledgments

Leon Weber and Jannes Münchmeyer acknowledge the support of the Helmholtz Einstein International Berlin Research School in Data Science (HEIB-RiDS). We would like to thank the anonymous reviewers for the constructive feedback. We gratefully acknowledge the support of NVIDIA Corporation with the donation of a Titan X Pascal GPU used for this research.

References

- Gabor Angeli, Neha Nayak, and Christopher D Manning. 2016. Combining natural logic and shallow reasoning for question answering. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 442–452.
- Rolf Apweiler, Amos Bairoch, Cathy H Wu, Winona C Barker, Brigitte Boeckmann, Serenella Ferro, Elisabeth Gasteiger, Hongzhan Huang, Rodrigo Lopez, Michele Magrane, et al. 2004. Uniprot: the universal protein knowledgebase. *Nucleic acids research*, 32(suppl_1):D115–D119.
- Stephen H. Bach, Matthias Broecheler, Bert Huang, and Lise Getoor. 2017. Hinge-loss markov random fields and probabilistic soft logic. *Journal of Machine Learning Research*, 18:109:1–109:67.
- Islam Beltagy, Cuong Chau, Gemma Boleda, Dan Garrette, Katrin Erk, and Raymond Mooney. 2013. Montague meets markov: Deep semantics with probabilistic logical form. In *Second Joint Conference on Lexical and Computational Semantics (* SEM), Volume 1: Proceedings of the Main Conference and the Shared Task: Semantic Textual Similarity*, volume 1, pages 11–21.
- Islam Beltagy, Katrin Erk, and Raymond Mooney. 2014. Probabilistic soft logic for semantic textual similarity. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 1210–1219.
- Tarek R Besold, Artur d’Avila Garcez, Sebastian Bader, Howard Bowman, Pedro Domingos, Pascal Hitzler, Kai-Uwe Kühnberger, Luis C Lamb, Daniel Lowd, Priscila Machado Vieira Lima, et al. 2017. Neural-symbolic learning and reasoning: A survey and interpretation. *arXiv preprint arXiv:1711.03902*.
- William W Cohen. 2016. Tensorlog: A differentiable deductive database. *arXiv preprint arXiv:1605.06523*.
- Rajarshi Das, Shehzaad Dhuliawala, Manzil Zaheer, Luke Vilnis, Ishan Durugkar, Akshay Krishnamurthy, Alex Smola, and Andrew McCallum. 2017. Go for a walk and arrive at the answer: Reasoning over paths in knowledge bases using reinforcement learning. *arXiv preprint arXiv:1711.05851*.
- Rajarshi Das, Arvind Neelakantan, David Belanger, and Andrew McCallum. 2016. Chains of reasoning over entities, relations, and text using recurrent neural networks. *arXiv preprint arXiv:1607.01426*.
- Nicola De Cao, Wilker Aziz, and Ivan Titov. 2018. Question answering by reasoning across documents with graph convolutional networks. *arXiv preprint arXiv:1808.09920*.
- Bhuwan Dhingra, Qiao Jin, Zhilin Yang, William W Cohen, and Ruslan Salakhutdinov. 2018. Neural models for reasoning over multiple mentions using coreference. *arXiv preprint arXiv:1804.05922*.
- Finale Doshi-Velez and Been Kim. 2017. *Towards a rigorous science of interpretable machine learning*. *arXiv*.
- Richard Evans and Edward Grefenstette. 2017. *Learning explanatory rules from noisy data*. *CoRR*, abs/1711.04574.
- Richard Evans and Edward Grefenstette. 2018. Learning explanatory rules from noisy data. *J. Artif. Intell. Res.*, 61:1–64.
- Anthony Fader, Luke Zettlemoyer, and Oren Etzioni. 2014. Open question answering over curated and extracted knowledge bases. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD ’14*, pages 1156–1165, New York, NY, USA. ACM.
- David Ferrucci, Eric Brown, Jennifer Chu-Carroll, James Fan, David Gondek, Aditya A Kalyanpur, Adam Lally, J William Murdock, Eric Nyberg, John Prager, and Others. 2010. Building watson: An overview of the DeepQA project. *AI magazine*, 31(3):59–79.
- Hervé Gallaire and Jack Minker, editors. 1978. *Logic and Data Bases, Symposium on Logic and Data Bases, Centre d’études et de recherches de Toulouse, 1977*, Advances in Data Base Theory. Plenum Press, New York.
- Dan Garrette, Katrin Erk, and Raymond Mooney. 2011. Integrating logical representations with probabilistic information using markov logic. In *Proceedings of the Ninth International Conference on Computational Semantics*, pages 105–114. Association for Computational Linguistics.
- Dan Garrette, Katrin Erk, and Raymond Mooney. 2014. A formal approach to linking logical form and vector-space lexical semantics. In *Computing meaning*, pages 27–48. Springer.
- Alex Graves, Greg Wayne, Malcolm Reynolds, Tim Harley, Ivo Danihelka, Agnieszka Grabska-Barwińska, Sergio Gómez Colmenarejo, Edward Grefenstette, Tiago Ramalho, John Agapiou, et al.

2016. Hybrid computing using a neural network with dynamic external memory. *Nature*, 538(7626):471.
- 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):93:1–93:42.
- M. M. Gupta and J. Qi. 1991. Theory of T-norms and Fuzzy Inference Methods. *Fuzzy Sets and Systems*, 40(3):431–450.
- Mikael Henaff, Jason Weston, Arthur Szlam, Antoine Bordes, and Yann LeCun. 2016. Tracking the world state with recurrent entity networks. *arXiv preprint arXiv:1612.03969*.
- Matthew Honnibal and Ines Montani. 2017. spacy 2: Natural language understanding with bloom embeddings, convolutional neural networks and incremental parsing. *To appear*.
- Kevin Jarrett, Koray Kavukcuoglu, Marc’ Aurelio Ranzato, and Yann LeCun. 2009. What is the best multi-stage architecture for object recognition? In *ICCV*, pages 2146–2153. IEEE Computer Society.
- Pascual Julián-Iranzo, Clemente Rubio-Manzano, and Juan Gallardo-Casero. 2009. Bousi prolog: a prolog extension language for flexible query answering. *Electron. Notes Theor. Comput. Sci.*, 248(Supplement C):131–147.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Thomas N Kipf and Max Welling. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*.
- Ankit Kumar, Ozan Irsoy, Peter Ondruska, Mohit Iyyer, James Bradbury, Ishaan Gulrajani, Victor Zhong, Romain Paulus, and Richard Socher. 2016. Ask me anything: Dynamic memory networks for natural language processing. In *International Conference on Machine Learning*, pages 1378–1387.
- Zachary C. Lipton. 2018. The mythos of model interpretability. *Commun. ACM*, 61(10):36–43.
- Dan Moldovan, Christine Clark, Sanda Harabagiu, and Steve Maiorano. 2003. COGEX: A logic prover for question answering. In *Proceedings of the 2003 Conference of the North American Chapter of the Association for Computational Linguistics on Human Language Technology - Volume 1, NAACL ’03*, pages 87–93, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Stephen Muggleton. 1991. Inductive logic programming. *New generation computing*, 8(4):295–318.
- Stephen Muggleton and Luc De Raedt. 1994. **Inductive logic programming: Theory and methods**. *J. Log. Program.*, 19/20:629–679.
- Arvind Neelakantan, Benjamin Roth, and Andrew McCallum. 2015. Compositional vector space models for knowledge base completion. *arXiv preprint arXiv:1504.06662*.
- Christina Niklaus, Matthias Cetto, André Freitas, and Siegfried Handschuh. 2018. A survey on open information extraction. In *COLING*, pages 3866–3878. Association for Computational Linguistics.
- Mehmet A Orgun and Wanli Ma. 1994. An overview of temporal and modal logic programming. In *Temporal logic*, pages 445–479. Springer.
- Matteo Pagliardini, Prakhar Gupta, and Martin Jaggi. 2017. Unsupervised learning of sentence embeddings using compositional n-gram features. *arXiv preprint arXiv:1703.02507*.
- Baolin Peng, Zhengdong Lu, Hang Li, and Kam-Fai Wong. 2015. Towards neural network-based reasoning. *arXiv preprint arXiv:1508.05508*.
- Sampo Pyysalo, Filip Ginter, Hans Moen, Tapio Salakoski, and Sophia Ananiadou. 2013. Distributional semantics resources for biomedical text processing.
- Martin Raison, Pierre-Emmanuel Mazaré, Rajarshi Das, and Antoine Bordes. 2018. Weaver: Deep co-encoding of questions and documents for machine reading. *arXiv preprint arXiv:1804.10490*.
- Matthew Richardson and Pedro M. Domingos. 2006. Markov logic networks. *Machine Learning*, 62(1-2):107–136.
- Tim Rocktäschel and Sebastian Riedel. 2017. End-to-end differentiable proving. In *Advances in Neural Information Processing Systems*, pages 3788–3800.
- Tim Rocktäschel, Sameer Singh, and Sebastian Riedel. 2015. Injecting logical background knowledge into embeddings for relation extraction. In *NAACL HLT 2015, The 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Denver, Colorado, USA, May 31 - June 5, 2015*, pages 1119–1129.
- Stuart J. Russell and Peter Norvig. 2010a. *Artificial Intelligence - A Modern Approach (3. internat. ed.)*. Pearson Education.
- Stuart J Russell and Peter Norvig. 2010b. *Artificial Intelligence: A Modern Approach*.
- Stuart J Russell and Peter Norvig. 2016. *Artificial intelligence: a modern approach*. Malaysia; Pearson Education Limited,.
- Minjoon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. 2016a. Bidirectional attention flow for machine comprehension. *arXiv preprint arXiv:1611.01603*.

Minjoon Seo, Sewon Min, Ali Farhadi, and Hananeh Hajishirzi. 2016b. Query-reduction networks for question answering. *arXiv preprint arXiv:1606.04582*.

Maria I Sessa. 2002. Approximate reasoning by similarity-based sld resolution. *Theoretical computer science*, 275(1-2):389–426.

Sainbayar Sukhbaatar, Jason Weston, Rob Fergus, et al. 2015. End-to-end memory networks. In *Advances in neural information processing systems*, pages 2440–2448.

Dirk Weissenborn, Tomas Kocisky, and Chris Dyer. 2017a. Dynamic integration of background knowledge in neural nlu systems. *CoRR*, abs/1706.02596.

Dirk Weissenborn, Pasquale Minervini, Isabelle Augenstein, Johannes Welbl, Tim Rocktäschel, Matko Bosnjak, Jeff Mitchell, Thomas Demeester, Tim Dettmers, Pontus Stenetorp, and Sebastian Riedel. 2018. Jack the reader - A machine reading framework. In *Proceedings of ACL 2018, Melbourne, Australia, July 15-20, 2018, System Demonstrations*, pages 25–30.

Dirk Weissenborn, Georg Wiese, and Laura Seiffe. 2017b. Fastqa: A simple and efficient neural architecture for question answering. *arXiv preprint arXiv:1703.04816*.

Johannes Welbl, Pontus Stenetorp, and Sebastian Riedel. 2017. Constructing datasets for multi-hop reading comprehension across documents. *arXiv preprint arXiv:1710.06481*.

David S Wishart, Craig Knox, An Chi Guo, Savita Shrivastava, Murtaza Hassanali, Paul Stothard, Zhan Chang, and Jennifer Woolsey. 2006. Drugbank: a comprehensive resource for in silico drug discovery and exploration. *Nucleic acids research*, 34(suppl_1):D668–D672.

Victor Zhong, Caiming Xiong, Nitish Shirish Keskar, and Richard Socher. 2019. Coarse-grain fine-grain coattention network for multi-evidence question answering. *arXiv preprint arXiv:1901.00603*.

Appendices

A Algorithms

```

fun unify(x, y,  $\theta$ , S)
  Input:
  x: function  $f(\dots)$  | atom  $p(\dots)$  |
  variable | list  $x_1 :: x_2 :: \dots :: x_n$ 
  y: function  $f'(\dots)$  | atom  $p'(\dots)$  |
  variable | list  $y_1 :: y_2 :: \dots :: y_m$ 
   $\theta$ : current substitutions, default = {}
  S: current success score, default = 1.0
  Output: (Unifying substitution  $\theta'$  or
  failure, Updated success score
  S')
  if  $\theta = \text{failure}$  then return (failure, 0)
  else if  $S < \lambda$  then return (failure, 0)
  else if  $x = y$  then return ( $\theta$ , S)
  else if x is Var then return
    unify_var(x, y,  $\theta$ , S)
  else if y is Var then return
    unify_var(y, x,  $\theta$ , S)
  else if x is  $f(x_1, \dots, x_n)$ , y is
   $f'(y_1, \dots, y_n)$ , and  $f \sim f' \geq \lambda$  then
     $S' := S \wedge f \sim f'$ 
    return unify( $x_1 :: \dots :: x_n, y_1 ::$ 
     $\dots :: y_n, \theta, S'$ )
  end
  else if x is  $p(x_1, \dots, x_n)$ , y is
   $p'(y_1, \dots, y_n)$ , and  $p \sim p' \geq \lambda$  then
     $S' := S \wedge p \sim p'$ 
    return unify( $x_1 :: \dots :: x_n, y_1 ::$ 
     $\dots :: y_n, \theta, S'$ )
  end
  else if x is  $x_1 :: \dots :: x_n$  and y is
   $y_1 :: \dots :: y_n$  then
    ( $\theta', S'$ ) := unify( $x_1, y_1, \theta, S$ )
    return unify( $x_2 :: \dots :: x_n, y_2 ::$ 
     $\dots :: y_n, \theta', S'$ )
  end
  else if x is empty list and y is empty list
  then return ( $\theta$ , S)
  else return (failure, 0)
fun unify_var(v, o,  $\theta$ , S)
  if {v/val}  $\in \theta$  then return
    unify(val, o,  $\theta$ , S)
  else if {o/val}  $\in \theta$  then return
    unify(var, val,  $\theta$ , S)
  else return ({v/o} +  $\theta$ , S)

```

Algorithm 1: The weak unification algorithm in NLPROLOG without occurs check