

Differentiable Reasoning on Large Knowledge Bases and Natural Language

Pasquale Minervini,^{*†1} Matko Bošnjak,^{*‡1}
 Tim Rocktäschel,^{1,2} Sebastian Riedel,^{1,2} Edward Grefenstette^{1,2}

¹UCL Centre for Artificial Intelligence, University College London

²Facebook AI Research

{p.minervini, m.bosnjak, t.rocktaschel, s.riedel, e.grefenstette}@cs.ucl.ac.uk

Abstract

Reasoning with knowledge expressed in natural language and Knowledge Bases (KBs) is a major challenge for Artificial Intelligence, with applications in machine reading, dialogue, and question answering. General neural architectures that jointly learn representations and transformations of text are very data-inefficient, and it is hard to analyse their reasoning process. These issues are addressed by end-to-end differentiable reasoning systems such as Neural Theorem Provers (NTPs), although they can only be used with small-scale symbolic KBs. In this paper we first propose Greedy NTPs (GNTPs), an extension to NTPs addressing their complexity and scalability limitations, thus making them applicable to real-world datasets. This result is achieved by dynamically constructing the computation graph of NTPs and including only the most promising proof paths during inference, thus obtaining orders of magnitude more efficient models¹. Then, we propose a novel approach for jointly reasoning over KBs and textual mentions, by embedding logic facts and natural language sentences in a shared embedding space. We show that GNTPs perform on par with NTPs at a fraction of their cost while achieving competitive link prediction results on large datasets, providing explanations for predictions, and inducing interpretable models.

Introduction

The main focus of Artificial Intelligence is building systems that exhibit intelligent behaviour (Levesque 2014). Notably, Natural Language Understanding (NLU) and Machine Reading (MR) aim at building models and systems with the ability to read text, extract meaningful knowledge, and reason with it (Etzioni, Banko, and Cafarella 2006; Hermann et al. 2015; Weston et al. 2015; Das et al. 2017). This ability facilitates both the synthesis of new knowledge and the possibility to verify and update a given assertion. Traditionally, automated reasoning applied to text requires natural language processing tools that compile it into the structured form of a KB (Niklaus et al. 2018). However, the compiled KBs tend to be incomplete, ambiguous, and noisy, impairing the application of

standard deductive reasoners (Huang, van Harmelen, and ten Teije 2005).

A rich and broad literature in MR has approached this problem within a variety of frameworks, including Natural Logic (MacCartney and Manning 2007), Semantic Parsing (Bos 2008), Natural Language Inference and Recognising Textual Entailment (Fyodorov, Winter, and Francez 2000; Bowman et al. 2015), and Question Answering (Hermann et al. 2015). Nonetheless, such methods suffer from several limitations. They rely on significant amounts of annotated data to suitably approximate the implicit distribution from which the data is drawn. In practice, this makes them unable to generalise well in the absence of a sufficient quantity of training data or appropriate priors on model parameters (Evans and Grefenstette 2018). Orthogonally, even when accurate, such methods cannot explain given predictions (Lipton 2018).

A promising strategy for overcoming these issues consists of combining *neural models* and *symbolic reasoning*, given their complementary strengths and weaknesses (d’Avila Garcez et al. 2015; Rocktäschel and Riedel 2017; Yang, Yang, and Cohen 2017; Evans and Grefenstette 2018; Weber et al. 2019). While symbolic models can generalise well from a small number of examples, they are brittle and prone to failure when the observations are noisy or ambiguous, or when the properties of the domain are unknown or hard to formalise, all of which being the case for natural language (Raedt et al. 2008; Garnelo and Shanahan 2019). Contrarily, neural models are robust to noise and ambiguity but not easily interpretable, making them unable to provide explanations or incorporating background knowledge (Guidotti et al. 2018).

Recent work in neuro-symbolic systems has made progress towards end-to-end differentiable reasoning models that can be trained via backpropagation while maintaining interpretability and generalisation, thereby inheriting the best of both worlds. Among such systems, NTPs (Rocktäschel and Riedel 2017; Minervini et al. 2018) are end-to-end differentiable deductive reasoners based on Prolog’s backward chaining algorithm, where discrete unification between atoms is replaced by a differentiable operator computing the similarities between their embedding representations.

NTPs are especially interesting since they allow learning *interpretable rules* from data, by back-propagating the prediction errors to the rule representations. Furthermore, the proving process in NTPs is *explainable* – the proof path asso-

*Equal contribution

†Corresponding author

‡Now at DeepMind

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹The Code and supplementary material are available online at <https://github.com/uclnlp/gntp>.

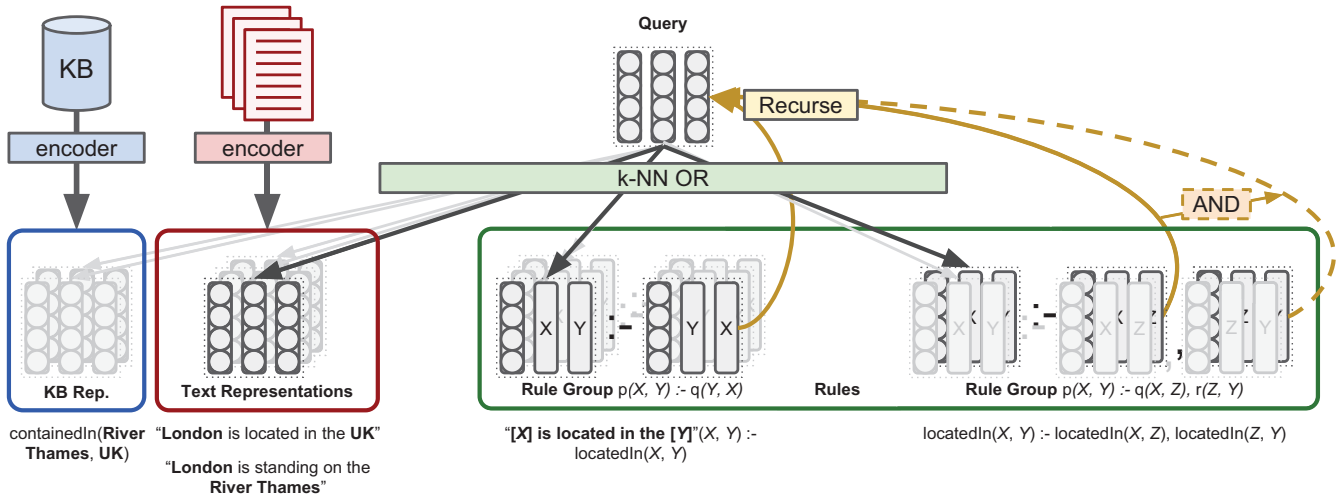


Figure 1: Overall architecture of GNTPs. The two main contributions lie in i) the significantly faster inference mechanism, sped up by the k-NN OR component, and ii) the text encoder.

ciated with the largest proof score denotes which rules and facts are used in the reasoning process. However, NTPs have only been successfully applied to learning tasks involving very small datasets, since their computational complexity makes them unusable on larger, real-world KBs. Furthermore, most human knowledge is not available in KBs, but in natural language texts which are difficult to reason over automatically.

In this paper we address these issues by proposing: *i)* two efficiency improvements for significantly reducing the time and space complexity of NTPs by reducing the number of candidate proof paths and introducing an attention mechanism for rule induction, and *ii)* an extension of NTPs towards natural language, jointly embedding predicates and textual surface patterns in a shared space by using an end-to-end differentiable reading component.

End-to-end Differentiable Proving

NTPs (Rocktäschel and Riedel 2017) recursively build a neural network enumerating all the possible proof paths for proving a query (or *goal*) on a given KB, and aggregate all their proof scores via max pooling. They do so by relying on three modules—a *unification module*, which compares sub-symbolic representations of logic atoms, and mutually recursive *or* and *and modules*, which jointly enumerate all possible proof paths, before the final aggregation selects the highest-scoring one.

In the following, we briefly overview these modules, and the training process used for learning the model parameters from data. We assume the existence of a function-free Datalog KB \mathfrak{K} containing *ground facts* in the form $[p, A, B]$ ², representing the logical atom $p(A, B)$ where p is a relation type, and A, B are its arguments.³ It also contains *rules* in

²For consistency, we use the same notation as Rocktäschel and Riedel (2017).

³We consider binary predicates, without loss of generality.

the form $H :- B$ such as $[p, X, Y] :- [[q, X, Z], [r, Z, Y]]$, denoting the rule $p(X, Y) :- q(X, Z), r(Z, Y)$, meaning that $q(X, Z), r(Z, Y)$ implies $p(X, Y)$, where X, Y, Z are universally quantified variables.

Unification Module. In the backward chaining reasoning algorithm, *unification* is the operator that matches two logic atoms, such as $\text{locatedIn}(\text{LONDON}, \text{UK})$ and $\text{situatedIn}(X, Y)$. Discrete unification checks for equality between the elements composing the two atoms (*e.g.* $\text{locatedIn} \neq \text{situatedIn}$), and binds variables to symbols via substitutions (*e.g.* $\{X/\text{LONDON}, Y/\text{UK}\}$). In NTPs, unification matches two atoms by comparing their *embedding representations* via a differentiable similarity function – a Gaussian kernel – which enables matching different symbols with similar semantics.

More formally, $\text{unify}_{\theta}(H, G, S) = S'$ creates a neural network module that matches two atoms H and G by comparing their embedding vectors. For instance, given a goal $G = [\text{locatedIn}, \text{LONDON}, \text{UK}]$, a fact $H = [\text{situatedIn}, X, Y]$, and a proof state $S = (S_{\psi}, S_{\rho})$ consisting of a set of substitutions S_{ψ} and a proof score S_{ρ} , the *unify* module compares the embedding representations of locatedIn and situatedIn with a Gaussian kernel k , updates the variable binding substitution set $S'_{\psi} = S_{\psi} \cup \{X/\text{LONDON}, Y/\text{UK}\}$, and calculates the new proof score $S'_{\rho} = \min(S_{\rho}, k(\theta_{\text{locatedIn}}, \theta_{\text{situatedIn}}))$ and proof state $S' = (S'_{\psi}, S'_{\rho})$.

OR Module. The *or* module computes the unification between a goal and all facts and rule heads in a KB, and then recursively invokes the *and* module on the corresponding rule bodies. Formally, for each rule $H :- B$ ⁴ in a KB \mathfrak{K} , $\text{or}_{\theta}^{\mathfrak{K}}(G, d, S)$ unifies the goal G with the rule head H , and invokes the *and* module to prove atoms in the body B , keeping

⁴Facts are seen as rules with no body and variables, *i.e.* $F :- []$.

track of the maximum proof depth d :

$$\text{or}_{\theta}^{\mathfrak{K}}(G, d, S) = [S' \mid H :- B \in \mathfrak{K}, \\ S' \in \text{and}_{\theta}^{\mathfrak{K}}(B, d, \text{unify}_{\theta}(H, G, S))] \quad (1)$$

For example, given a goal $G = [\text{situatedIn}, Q, \text{UK}]$ and a rule $H :- B$ with $H = [\text{locatedIn}, X, Y]$ and $B = [[\text{locatedIn}, X, Z], [\text{locatedIn}, Z, Y]]$, the model would unify the goal G with the rule head H , and invoke the `and` modules to prove the sub-goals in the rule body B .

AND Module. The `and` module recursively proves a list of sub-goals in a rule body. Given the first sub-goal B and the following sub-goals \mathbb{B} , the $\text{and}_{\theta}^{\mathfrak{K}}(B : \mathbb{B}, d, S)$ module will substitute variables in B with constants according to the substitutions in S , and invoke the `or` module on B . The resulting state is used to prove the atoms in \mathbb{B} , by recursively invoking the `and` module:

$$\text{and}_{\theta}^{\mathfrak{K}}(B : \mathbb{B}, d, S) = [S'' \mid d > 0, \\ S'' \in \text{and}_{\theta}^{\mathfrak{K}}(\mathbb{B}, d, S'), \\ S' \in \text{or}_{\theta}^{\mathfrak{K}}(\text{sub}(B, S_{\psi}), d - 1, S)] \quad (2)$$

For example, when invoked on the rule body B of the example mentioned above, the `and` module will substitute variables with constants for the sub-goal $[\text{locatedIn}, X, Z]$ and invoke the `or` module, whose resulting state will be the basis of the next invocation of `and` module on $[\text{locatedIn}, Z, Y]$.

Proof Aggregation. After building a neural network that evaluates all the possible proof paths of a goal G on a KB \mathfrak{K} , NTPs select the proof path with the largest proof score:

$$\text{ntp}_{\theta}^{\mathfrak{K}}(G, d) = \max_S S_{\rho} \\ \text{with } S \in \text{or}_{\theta}^{\mathfrak{K}}(G, d, (\emptyset, 1)) \quad (3)$$

where $d \in \mathbb{N}$ is a predefined maximum proof depth. The initial proof state is set to $(\emptyset, 1)$ corresponding to an empty substitution set and to a proof score of 1.

Training. In NTPs, embedding representations are learned by minimising a cross-entropy loss $\mathcal{L}^{\mathfrak{K}}(\theta)$ on the final proof score, by iteratively masking facts in the KB and trying to prove them using other available facts and rules.

Negative examples are obtained via a corruption process, denoted by $\text{corrupt}(\cdot)$, by modifying the subject and object of triples in the KB (Nickel et al. 2016):

$$\mathcal{L}^{\mathfrak{K}}(\theta) = - \sum_{F := \square \in \mathfrak{K}} \log \text{ntp}_{\theta}^{\mathfrak{K}}(F, d) \\ - \sum_{\tilde{F} \sim \text{corrupt}(F)} \log[1 - \text{ntp}_{\theta}^{\mathfrak{K}}(\tilde{F}, d)] \quad (4)$$

NTPs can also learn *interpretable rules*. Rocktäschel and Riedel (2017) show that it is possible to learn rules from data by specifying *rule templates*, such as $H :- B$ with $H = [\theta_p, X, Y]$ and $B = [[\theta_q, X, Z], [\theta_r, Z, Y]]$.

Parameters $\theta_p, \theta_q, \theta_r \in \mathbb{R}^k$, denoting rule-predicate embeddings, can be learned from data by minimising the loss in Eq. 4, and decoded by searching the closest representation of known predicates.

Efficient Differentiable Reasoning on Large-Scale KBs

NTPs are capable of deductive reasoning, and the proof paths with the highest score can provide human-readable explanations for a given prediction. However, enumerating and scoring all bounded-depth proof paths for a given goal, as given in Eq. 3, is computationally intractable. For each goal and sub-goal G , this process requires to unify G with the representations of *all* rule heads and facts in the KB, which quickly becomes computationally prohibitive even for moderately sized KBs. Furthermore, the expansion of a rule like $p(X, Y) :- q(X, Z), r(Z, Y)$ via backward chaining causes an increase of the sub-goals to prove, both because all atoms in the body need to be proven, and because Z is a free variable with many possible bindings (Rocktäschel and Riedel 2017). We consider two problems – given a sub-goal G such as $[p, A, B]$, we need to efficiently select *i*) the k_f facts that are most likely to prove a sub-goal G , and *ii*) the k_r rules to expand to reach a high-scoring proof state.

Fact Selection. Unifying a sub-goal G with all facts in the KB \mathfrak{K} may not be feasible in practice. The number of facts in a real-world KB can be in the order of millions or billions. For instance, Freebase contains over 637×10^6 facts, while the Google Knowledge Graph contains more than 18×10^9 facts (Nickel et al. 2016). Identifying the facts $F \in \mathfrak{K}$ that yield the maximum proof score for a sub-goal G reduces to solving the following optimisation problem:

$$\text{ntp}_{\theta}^{\mathfrak{K}}(G, 1) = \max_{F := \square \in \mathfrak{K}} S_{\rho}^F = S_{\rho}^* \\ \text{with } S^F = \text{unify}_{\theta}(F, G, (\emptyset, 1)) \quad (5)$$

Hence, the fact $F \in \mathfrak{K}$ that yields the maximum proof score for a sub-goal G is the fact F that yields the maximum unification score with G . Recall that the unification score between a fact F and a goal G is given by the similarity of their embedding representations θ_F and θ_G , computed via a Gaussian kernel $k(\theta_F, \theta_G)$. Given a goal G , NTPs will compute the unification score between G and every fact $F \in \mathfrak{K}$ in the KB. This is problematic, since computing the similarity between the representations of the goal G and every fact $F \in \mathfrak{K}$ is computationally prohibitive – the number of comparisons is $\mathcal{O}(|\mathfrak{K}|n)$, where n is the number of (sub-)goals in the proving process. However, $\text{ntp}_{\theta}^{\mathfrak{K}}(G, d)$ only returns the single largest proof score. This means that, at inference time, we only need the largest proof score for returning the correct output. Similarly, during training, the gradient of the proof score with respect to the parameters θ can also be calculated exactly by using the single largest proof score:

$$\frac{\partial \text{ntp}_{\theta}^{\mathfrak{K}}(G, 1)_{\rho}}{\partial \theta} = \frac{\partial \max_{F \in \mathfrak{K}} S_{\rho}^F}{\partial \theta} = \frac{\partial S_{\rho}^*}{\partial \theta} \\ \text{with } S_{\rho}^* = \max_{F \in \mathfrak{K}} S_{\rho}^F$$

In this paper, we propose to efficiently compute S^* , the highest unification score between a given sub-goal G and a fact $F \in \mathfrak{K}$, by casting it as a Nearest Neighbour Search (NNS) problem. This is feasible since the Gaussian kernel used by NTPs is a monotonic transformation of the negative Euclidean distance.

Identifying S^* permits to reduce the number of neural network sub-structures needed for the comparisons between each sub-goal and facts from $\mathcal{O}(|\mathfrak{K}|)$ to $\mathcal{O}(1)$. We use the exact and approximate NNS framework proposed by Johnson, Douze, and Jégou (2017) for efficiently searching \mathfrak{K} for the best supporting facts for a given sub-goal. Specifically we use the exact L2-nearest neighbour search and, for the sake of efficiency, we update the search index every 10 batches, assuming that the small updates made by stochastic gradient descent do not necessarily invalidate previous search indexes.

Rule Selection. We use a similar idea for selecting which rules to activate for proving a given goal G . We empirically notice that unifying G with the closest rule heads, such as $G = [\text{locatedIn}, \text{LONDON}, \text{UK}]$ and $H = [\text{situatedIn}, X, Y]$, is more likely to generate high-scoring proof states. This is a trade-off between symbolic reasoning, where proof paths are expanded only when the heads exactly match with the goals, and differentiable reasoning, where all proof paths are explored.

This prompted us to implement a heuristic that dynamically selects rules among rules sharing the same template during both inference and learning. In our experiments, this heuristic for selecting proof paths was able to recover valid proofs for a goal when they exist, while drastically reducing the computational complexity of the differentiable proving process.

More formally, we generate a partitioning $\mathfrak{P} \in 2^{\mathfrak{K}}$ of the KB \mathfrak{K} , where each element in \mathfrak{P} groups all facts and rules in \mathfrak{K} sharing the same template, or high-level structure – e.g. an element of \mathfrak{P} contains all rules with structure $\theta_p(X, Y) :- \theta_q(X, Z), \theta_r(Z, Y)$, with $\theta_p, \theta_q, \theta_r \in \mathbb{R}^k$.⁵ We then redefine the or_{θ} operator as follows:

$$\text{or}_{\theta}^{\mathfrak{K}}(G, d, S) = [S' \mid H :- B \in \mathcal{N}_{\mathcal{P}}(G), \mathcal{P} \in \mathfrak{P}, \\ S' \in \text{and}_{\theta}^{\mathfrak{K}}(B, d, \text{unify}_{\theta}(H, G, S))]$$

where, instead of unifying a sub-goal G with all rule heads, we constrain the unification to only the rules where heads are in the neighbourhood $\mathcal{N}_{\mathcal{P}}(G)$ of G .

Learning to Attend Over Predicates. Although NTPs can be used for *learning interpretable rules* from data, the solution proposed by Rocktäschel and Riedel (2017) can be quite inefficient, as the number of parameters associated to rules can be quite large. For instance, the rule $H :- B$, with $H = [\theta_p, X, Y]$ and $B = [[\theta_q, X, Z], [\theta_r, Z, Y]]$, where $\theta_p, \theta_q, \theta_r \in \mathbb{R}^k$, introduces $3k$ parameters in the model,

⁵Grouping rules with the same structure together makes allows parallel inference to be implemented very efficiently on GPU. This optimisation is also present in Rocktäschel and Riedel (2017).

where k denotes the embedding size, and it may be computationally inefficient to learn each of the embedding vectors if k is large.

We propose using an *attention mechanism* (Bahdanau, Cho, and Bengio 2015) for attending over known predicates for defining the rule-predicate embeddings $\theta_p, \theta_q, \theta_r$. Let \mathcal{R} be the set of known predicates, and let $R \in \mathbb{R}^{|\mathcal{R}| \times k}$ be a matrix representing the embeddings for the predicates in \mathcal{R} . We define θ_p as $\theta_p = \text{softmax}(\mathbf{a}_p)^\top R$, where $\mathbf{a}_p \in \mathbb{R}^{|\mathcal{R}|}$ is a set of trainable *attention weights* associated with the predicate p . This sensibly improves the parameter efficiency of the model in cases where the number of known predicates is low, i.e. $|\mathcal{R}| \ll k$, by introducing $c|\mathcal{R}|$ parameters for each rule rather than ck , where c is the number of trainable predicate embeddings in the rule.

Jointly Reasoning on Knowledge Bases and Natural Language

In this section, we show how GNTPs can jointly reason over KBs and natural language corpora. In the following, we assume that our KB \mathfrak{K} is composed of facts, rules, and *textual mentions*. A fact is composed of a predicate symbol and a sequence of arguments, e.g. $[\text{locationOf}, \text{LONDON}, \text{UK}]$. On the other hand, a *mention* is a textual pattern between two co-occurring entities in the KB (Toutanova et al. 2015), such as “LONDON is located in the UK”.

We represent mentions jointly with facts and rules in \mathfrak{K} by considering each textual surface pattern linking two entities as a new predicate, and embedding it in a d -dimensional space by means of an end-to-end differentiable reading component. For instance, the sentence “United Kingdom borders with Ireland” can be translated into the following mention: $[[[\text{arg1}], \text{borders}, \text{with}, [\text{arg2}]], \text{UK}, \text{IRELAND}]$, by first identifying sentences or paragraphs containing KB entities, and then considering the textual surface pattern connecting such entities as an extra relation type. While predicates in \mathcal{R} are encoded by a look-up operation to a predicate embedding matrix $R \in \mathbb{R}^{|\mathcal{R}| \times k}$, textual surface patterns are encoded by an $\text{encode}_{\theta} : \mathcal{V}^* \rightarrow \mathbb{R}^k$ module, where \mathcal{V} is the vocabulary of words and symbols occurring in textual surface patterns.

More formally, given a textual surface pattern $t \in \mathcal{V}^*$ – such as $t = [[[\text{arg1}], \text{borders}, \text{with}, [\text{arg2}]]]$ – the encode_{θ} module first encodes each token w in t by means of a token embedding matrix $V \in \mathbb{R}^{|\mathcal{V}| \times k'}$, resulting in a pattern matrix $W_t \in \mathbb{R}^{|t| \times k'}$. Then, the module produces a textual surface pattern embedding vector $\theta_t \in \mathbb{R}^k$ from W_t by means of an end-to-end differentiable encoder. For assessing whether a simple encoder architecture can already provide benefits to the model, we use an encode_{θ} module that aggregates the embeddings of the tokens composing a textual surface pattern via mean pooling: $\text{encode}_{\theta}(t) = \frac{1}{|t|} \sum_{w \in t} V_w \in \mathbb{R}^k$. Albeit the encoder can be implemented by using other differentiable architectures, for this work we opted for a simple but still very effective Bag of Embeddings model (White et al. 2015; Arora, Liang, and Ma 2017) showing that, even in this case,

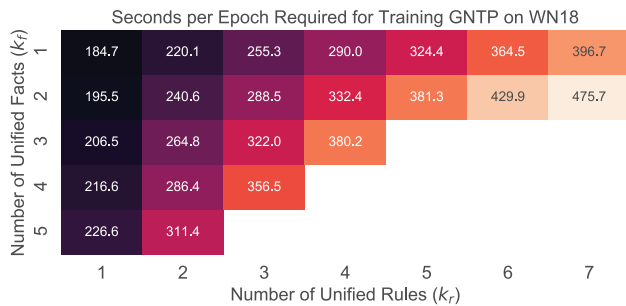


Figure 2: Number of seconds per epoch required for training on the WN18 dataset using batches of 1000 examples on a GPU. Missing entries denote out-of-memory errors.

the model achieves very accurate results.

Related Work

A notable corpus of literature aims at addressing the limitations of neural architectures in terms of generalisation and reasoning abilities. A line of research consists of enriching neural network architectures with a differentiable *external memory* (Sukhbaatar et al. 2015; Graves, Wayne, and Danihelka 2014; Joulin and Mikolov 2015; Grefenstette et al. 2015; Kaiser and Sutskever 2016). The underlying idea is that a neural network can learn to represent and manipulate complex data structures, thus disentangling the algorithmic part of the process from the representation of the inputs. By doing so, it becomes possible to train such models from enriched supervision signals, such as from *program traces* rather than simple input-output pairs.

A related field is *differentiable interpreters*—program interpreters where declarative or procedural knowledge is compiled into a neural network architecture (Bošnjak et al. 2017; Rocktäschel and Riedel 2017; Evans and Grefenstette 2018). This family of models allows imposing strong inductive biases on the models by partially defining the program structure used for constructing the network, *e.g.*, in terms of instruction sets or rules. A major drawback of differentiable interpreters, however, is their computational complexity, so far deeming them unusable except for smaller learning problems. Rae et al. (2016) use an approximate nearest neighbour data structures for sparsifying read operations in memory networks.

Riedel et al. (2013) pioneered the idea of jointly embedding KB facts and textual mentions in shared embedding space, by considering mentions as additional relations in a KB factorisation setting, and more elaborate mention encoders were investigated by McCallum, Neelakantan, and Verga (2017).

Our work is also related to path encoding models (Das et al. 2017) and random walk approaches (Lao, Mitchell, and Cohen 2011; Gardner et al. 2014), both of which lack a rule induction mechanisms, and to approaches combining observable and latent features of the graph (Nickel, Jiang, and Tresp 2014; Minervini et al. 2016). Lastly, our work is related to Yang, Yang, and Cohen (2017), a scalable rule induction approach for KB completion, but has not been applied to

textual surface patterns.

Experiments

Datasets and Evaluation Protocols. We report the results of experiments on benchmark datasets — Countries (Bouchard, Singh, and Trouillon 2015), Nations, UMLS, and Kinship (Kemp et al. 2006) — following the same evaluation protocols as Rocktäschel and Riedel (2017). Furthermore, since GNTPs allows to experiment on significantly larger datasets, we also report results on the WN18 (Bordes et al. 2013), WN18RR (Dettmers et al. 2018) and FB122 (Guo et al. 2016) datasets. Results are reported in terms of the Area Under the Precision-Recall Curve (AUC-PR) (Davis and Goadrich 2006), Mean Reciprocal Rank (MRR), and HITS@ m (Bordes et al. 2013). Datasets and hyperparameters are described in the Appendix.⁶

Baselines. On benchmark datasets, we compare GNTPs with NTPs and two other neuro-symbolic reasoning systems, MINERVA (Das et al. 2018), which employs a reinforcement learning algorithm to reach answers by traversing the KB graph, and NeuralLP (Yang, Yang, and Cohen 2017), which compiles inference tasks in a sequence of differentiable operations. In addition, we consider DistMult (Yang et al. 2015) and ComplEx (Trouillon et al. 2016), two state-of-the-art black-box neural link predictors suited for large datasets.

Run-Time Evaluation. To assess the benefits of GNTPs in terms of computational complexity and range of applications, we consider the best hyperparameters we found for the WN18 dataset, and measured the time needed for each training epoch varying the number of unified facts and rules during inference. Results, outlined in Fig. 2, show that learning on WN18 quickly becomes infeasible by increasing the number of unified facts and rules. NTPs are a special case of GNTPs where, during the forward pass, there is no pruning of the proof paths.

From Fig. 2 we can see that even for KBs a fraction the size of WordNet and Freebase, NTPs rapidly run out of memory, deeming them inapplicable to reasonably sized KBs. Instead, sensible pruning of proof paths in GNTPs drastically increases the efficiency of both the learning and the inference process, allowing to train on large KBs like WordNet. We refer to the Appendix⁶ for additional experiments showing run-time improvements by several orders of magnitude.

Link Prediction Results. We compare GNTPs and NTPs on a set of link prediction benchmarks, also used in Rocktäschel and Riedel (2017). Results, presented in Table 1, show that GNTPs achieves better or on-par results in comparison with NTPs and baselines MINERVA (Das et al. 2018) and NeuralLP (Das et al. 2018), consistently through all benchmark datasets. We can also see that models learned by GNTPs are *interpretable*: in Table 1 we show the decoded rules learned by the model, and learn about the domain at

⁶The Appendix can be found at <https://github.com/uclnlp/gntp>

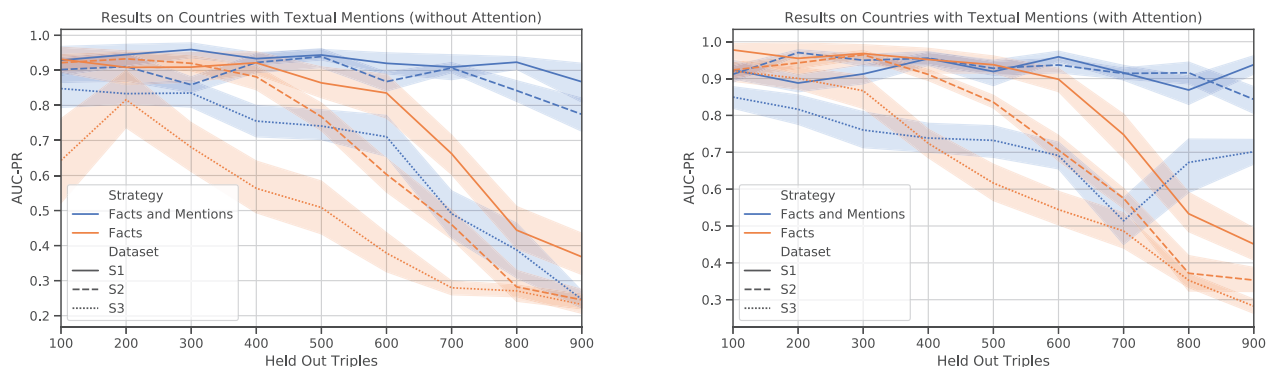


Figure 3: GNTPs on Countries with generated mentions. We replaced a varying number of relations with textual mentions and integrated them by encoding the mentions using a text encoder (*Facts and Mentions*) and by simply adding them to the KB (*Facts*). Two figures contrast the effects of rule learning without attention (left) and with it (right).

Table 1: Comparison of GNTPs, NTPs, NeuralLP (Yang, Yang, and Cohen 2017), and MINERVA (Das et al. 2018) (from Das et al. (2018)) on benchmark datasets, with and without attention.

Datasets	Metrics	Models					Rules Learned by GNTP
		NTP ⁷	GNTP		NeuralLP	MINERVA	
			Standard	Attention			
Countries	S1	90.83 ± 15.4	99.98 ± 0.05	100.0 ± 0.0	100.0 ± 0.0	locatedIn(X,Y) :- locatedIn(X,Z), locatedIn(Z,Y)	
	S2	87.40 ± 11.7	90.82 ± 0.88	93.48 ± 3.29	75.1 ± 0.3	92.36 ± 2.41	neighborOf(X,Y) :- neighborOf(X,Z), locatedIn(Z,Y)
	S3	56.68 ± 17.6	87.70 ± 4.79	91.27 ± 4.02	92.20 ± 0.2	95.10 ± 1.20	neighborOf(X,Y) :- neighborOf(Y,X)
Kinship	MRR	0.35	0.719	0.759	0.619	0.720	term0(X, Y) :- term0(Y, X)
	HITS@1	0.24	0.586	0.642	0.475	0.605	term4(X, Y) :- term4(Y, X)
	HITS@3	0.37	0.815	0.850	0.707	0.812	term13(X,Y) :- term13(X, Z), term10(Z, Y)
	HITS@10	0.57	0.958	0.959	0.912	0.924	term2(X,Y) :- term4(X, Z), term7(Z, Y)
Nations	MRR	0.61	0.658	0.645	—	—	commonbloc1(X, Y) :- relngo(Y, X)
	HITS@1	0.45	0.493	0.490	—	—	timesincewar(X,Y) :- independence(X,Y)
	HITS@3	0.73	0.781	0.736	—	—	unweightedunvote(X,Y) :- relngo(X,Y)
	HITS@10	0.87	0.985	0.975	—	—	ngo(X, Y) :- independence(Y, X)
UMLS	MRR	0.80	0.841	0.857	0.778	0.825	isa(X,Y) :- isa(X,Z), isa(Z,Y)
	HITS@1	0.70	0.732	0.761	0.643	0.728	complicates(X,Y) :- affects(X,Y)
	HITS@3	0.88	0.941	0.947	0.869	0.900	affects(X, Y) :- affects(X, Z), affects(Z, Y)
	HITS@10	0.95	0.986	0.983	0.962	0.968	process_of(X,Y) :- affects(X,Y)

hand. For instance, we can see that on UMLS, a biomedical KB, the *isa* and *affects* relation are transitive.

Experiments with Generated Mentions. For evaluating different strategies of integrating textual surface patterns, in the form of mentions, in NTPs, we proceeded as follows. We replaced a varying number of training set triples from each of the Countries S1-S3 datasets with human-generated textual mentions (for more details, see Appendix).⁶ For instance, the fact *neighbourOf*(UK, IRELAND) may be replaced by the textual mention “UK is neighbouring with IRELAND”. The entities UK and IRELAND become the arguments, while the text between them is treated as a new logic predicate, forming a new fact “X is neighbouring with Y”(UK, IRELAND).

Then, we evaluate two ways of integrating textual mentions in GNTPs: *i*) adding them as facts to the KB, and *ii*) parsing the mention by means of an encoder. The results, presented in Fig. 3, show that the proposed encoding module yields consistent improvements of the ranking accuracy in comparison to simply adding the mentions as facts to the KB. This

is especially evident in cases where the number of held-out facts is higher, as it is often the case in real-world use cases, where there is an abundance of text but the KBs are sparse and incomplete (Nickel et al. 2016). GNTPs are extremely efficient at learning rules involving both *logic atoms* and *textual mentions*.

For instance, by analysing the learned models and their explanations, we can see that GNTPs learn rules such as

```
neighborOf(X, Y) :- “Y is a neighboring state to X”(X, Y)
locatedIn(X, Y) :- “X is a neighboring state to Z”(X, Z),
“Z is located in Y”(Z, Y)
```

and leverage them during their reasoning process, providing human-readable explanations for a given prediction.

⁷Results reported in Rocktäschel and Riedel (2017) were calculated with an incorrect evaluation function, causing artificially better results. We corrected the issues, and recalculated the results.

Table 2: Link prediction results on the Test-I, Test-II and Test-ALL on FB122. Note that KALE, ASR methods, and KBLR have access to a set of rules provided by Guo et al. (2016), while neural link predictors and GNTPs do not. Test-II (6,186 triples) denotes a subset of FB122 that can be inferred via logic rules, while Test-I (5,057 triples) denotes all other test triples. We can see that, even without providing any rule to the model, GNTPs yields better ranking results in comparison with neural link prediction models—since it is able to learn such rules from data—and it is comparable with models that can leverage the provided rules.

		Test-I				Test-II				Test-ALL			
		Hits@N (%)			MRR	Hits@N (%)			MRR	Hits@N (%)			MRR
		3	5	10		3	5	10		3	5	10	
With Rules	KALE-Pre (Guo et al. 2016)	35.8	41.9	49.8	0.291	82.9	86.1	89.9	0.713	61.7	66.2	71.8	0.523
	KALE-Joint (Guo et al. 2016)	38.4	44.7	52.2	0.325	79.7	84.1	89.6	0.684	61.2	66.4	72.8	0.523
	ASR-DistMult (Minervini et al. 2017)	36.3	40.3	44.9	0.330	98.0	99.0	99.2	0.948	70.7	73.1	75.2	0.675
	ASR-ComplEx (Minervini et al. 2017)	37.3	41.0	45.9	0.338	99.2	99.3	99.4	0.984	71.7	73.6	75.7	0.698
	KBLR (García-Durán and Niepert 2018)	–	–	–	–	–	–	–	–	74.0	77.0	79.7	0.702
Without Rules	TransE (Bordes et al. 2013)	36.0	41.5	48.1	0.296	77.5	82.8	88.4	0.630	58.9	64.2	70.2	0.480
	DistMult (Yang et al. 2015)	36.0	40.3	45.3	0.313	92.3	93.8	94.7	0.874	67.4	70.2	72.9	0.628
	ComplEx (Trouillon et al. 2016)	37.0	41.3	46.2	0.329	91.4	91.9	92.4	0.887	67.3	69.5	71.9	0.641
	GNTPs	33.7	36.9	41.2	0.313	98.2	99.0	99.3	0.977	69.2	71.1	73.2	0.678

Table 3: Explanations, in terms of rules and supporting facts, for the queries in the validation set of WN18 provided by GNTPs by looking at the proof paths yielding the largest proof scores.

	Query	Score S_p	Proofs / Explanations
WN18	part_of(CONGO.N.03, AFRICA.N.01)	0.995	part_of(X, Y) :- has_part(Y, X) has_part(AFRICA.N.01, CONGO.N.03)
		0.787	part_of(X, Y) :- instance_hyponym(Y, X) instance_hyponym(AFRICAN_COUNTRY.N.01, CONGO.N.03)
	hyponym(EXTINGUISH.V.04, DECOUPLE.V.03)	0.987	hyponym(X, Y) :- hypernym(Y, X) hypernym(DECOUPLE.V.03, EXTINGUISH.V.04)
	has_part(TEXAS.N.01, ODESSA.N.02)	0.961	has_part(X, Y) :- part_of(Y, X) part_of(ODESSA.N.02, TEXAS.N.01)

Results on Freebase and WordNet

Link prediction results for FB122 are summarised in Table 2. The FB122 dataset proposed by Guo et al. (2016) is fairly large scale: it comprises 91,638 triples, 9,738 entities, and 122 relations, as well as 47 rules that can be leveraged by models for link prediction tasks. For such a reason, we consider a series of models that can leverage the presence of such rules, namely KALE (Guo et al. 2016), DistMult and ComplEx using Adversarial Sets (ASR) (Minervini et al. 2017)—a method for incorporating rules in neural link predictors via adversarial training—and the recently proposed KBLR (García-Durán and Niepert 2018). Note that, unlike these methods, GNTPs do not have access to such rules and need to learn them from data.

Table 2 shows that GNTP, whilst not having access to rules, performs significantly better than neural link predictors, and on-par with methods that have access to all rules. In particular, we can see that on Test-II, a subset of FB122 directly related to logic rules, GNTP yields competitive results. GNTP is able to induce rules relevant for accurate predictions, such as:

```
timeZone(X, Y) :- containedBy(X, Z), timeZone(Z, Y).
nearbyAirports(X, Y) :- containedBy(X, Z), contains(Z, Y).
children(X, Y) :- parents(Y, X).
spouse(X, Y) :- spouse(Y, X).
```

We also evaluate GNTP on WN18 (Bordes et al. 2013) and WN18RR (Dettmers et al. 2018). In terms of ranking accuracy, GNTPs is comparable to state-of-the-art models, such as ComplEx and KBLR. In García-Durán and Niepert

(2018) authors report a 94.2 MRR for ComplEx and 93.6 MRR for KBLR, while NeuralLP (Yang, Yang, and Cohen 2017) achieves 94.0, with hits@10 equal to 94.5. GNTP achieves 94.2 MRR and 94.31, 94.41, 94.51 hits@3, 5, 10, which is on par with state-of-the-art neural link prediction models, while being interpretable via proof paths. Table 3 shows an excerpt of validation triples together with their GNTP proof scores and associated proof paths for WN18. On WN18RR, GNTP with MRR of 43.4 performs close to ComplEx (Dettmers et al. 2018) (44.0 MRR) but lags behind NeuralLP (46.3 MRR).

We can see that GNTPs is capable of learning and utilising rules, such as `has_part(X, Y) :- part_of(Y, X)`, and `hyponym(X, Y) :- hypernym(Y, X)`. Interestingly, GNTP is able to find non-trivial explanations for a given fact, based on the similarity between entity representations. For instance, it can explain that CONGO is part of AFRICA by leveraging the semantic similarity with AFRICAN_COUNTRY.

Conclusions

NTPs combine the strengths of rule-based and neural models but, so far, they were unable to reason over large KBs and natural language. In this paper, we overcome such limitations by considering only the subset of proof paths associated with the largest proof scores during the construction of a dynamic computation graph.

The proposed model, GNTP, is more computationally efficient by several orders of magnitude, while achieving similar or better predictive performance than NTPs. GNTPs enable

end-to-end differentiable reasoning on large KBs and natural language texts, by embedding logic atoms and textual mentions in the same embedding space. Furthermore, GNTPs are interpretable and can provide explanations in terms of logic proofs at scale.

References

- Arora, S.; Liang, Y.; and Ma, T. 2017. A simple but tough-to-beat baseline for sentence embeddings. In *ICLR*.
- Bahdanau, D.; Cho, K.; and Bengio, Y. 2015. Neural Machine Translation by Jointly Learning to Align and Translate. In *ICLR*.
- Bordes, A.; Usunier, N.; García-Durán, A.; Weston, J.; and Yakhnenko, O. 2013. Translating Embeddings for Modeling Multi-relational Data. In *NIPS*, 2787–2795.
- Bos, J. 2008. Wide-coverage semantic analysis with boxer. In *STEP*. Association for Computational Linguistics.
- Bouchard, G.; Singh, S.; and Trouillon, T. 2015. On approximate reasoning capabilities of low-rank vector spaces. In *AAAI Spring Symposia*. AAAI Press.
- Bošnjak, M.; Rocktäschel, T.; Naradowsky, J.; and Riedel, S. 2017. Programming with a Differentiable Forth Interpreter. In *ICML*, volume 70, 547–556.
- Bowman, S. R.; Angeli, G.; Potts, C.; and Manning, C. D. 2015. A large annotated corpus for learning natural language inference. In *EMNLP*.
- Das, R.; Neelakantan, A.; Belanger, D.; and McCallum, A. 2017. Chains of Reasoning over Entities, Relations, and Text using Recurrent Neural Networks. In *EACL*, 132–141.
- Das, R.; Dhuliawala, S.; Zaheer, M.; Vilnis, L.; Durugkar, I.; Krishnamurthy, A.; Smola, A. J.; and McCallum, A. 2018. Go for a Walk and Arrive at the Answer: Reasoning Over Paths in Knowledge Bases using Reinforcement Learning. In *ICLR*.
- d’Avila Garcez, A. S.; Besold, T. R.; Raedt, L. D.; Földiák, P.; Hitzler, P.; Icard, T.; Kühnberger, K.; Lamb, L. C.; Miiikkulainen, R.; and Silver, D. L. 2015. Neural-symbolic learning and reasoning: Contributions and challenges. In *AAAI Spring Symposia*.
- Davis, J., and Goadrich, M. 2006. The relationship between Precision-Recall and ROC curves. In *ICML*, volume 148.
- Dettmers, T.; Minervini, P.; Stenetorp, P.; and Riedel, S. 2018. Convolutional 2D Knowledge Graph Embeddings. In *AAAI*.
- Etzioni, O.; Banko, M.; and Cafarella, M. J. 2006. Machine reading. In *AAAI*, 1517–1519. AAAI Press.
- Evans, R., and Grefenstette, E. 2018. Learning Explanatory Rules from Noisy Data. *JAIR* 61:1–64.
- Fyodorov, Y.; Winter, Y.; and Francez, N. 2000. A Natural Logic Inference System. In *Proceedings of the of the 2nd Workshop on Inference in Computational Semantics*.
- García-Durán, A., and Niepert, M. 2018. KBlrn: End-to-End Learning of Knowledge Base Representations with Latent, Relational, and Numerical Features. In *UAI*, 372–381.
- Gardner, M.; Talukdar, P. P.; Krishnamurthy, J.; and Mitchell, T. M. 2014. Incorporating Vector Space Similarity in Random Walk Inference over Knowledge Bases. In *EMNLP*, 397–406.
- Garnelo, M., and Shanahan, M. 2019. Reconciling deep learning with symbolic artificial intelligence: representing objects and relations. *Current Opinion in Behavioral Sciences* 29:17 – 23.
- Graves, A.; Wayne, G.; and Danihelka, I. 2014. Neural Turing Machines. *CoRR* abs/1410.5401.
- Grefenstette, E.; Hermann, K. M.; Suleyman, M.; and Blunsom, P. 2015. Learning to Transduce with Unbounded Memory. In *NIPS*, 1828–1836.
- Guidotti, R.; Monreale, A.; Ruggieri, S.; Turini, F.; Giannotti, F.; and Pedreschi, D. 2018. A Survey of Methods for Explaining Black Box Models. *ACM CSUR* 51(5):93:1–93:42.
- Guo, S.; Wang, Q.; Wang, L.; Wang, B.; and Guo, L. 2016. Jointly Embedding Knowledge Graphs and Logical Rules. In *EMNLP*, 192–202.
- Hermann, K. M.; Kocisky, T.; Grefenstette, E.; Espeholt, L.; Kay, W.; Suleyman, M.; and Blunsom, P. 2015. Teaching Machines to Read and Comprehend. In *NIPS*, 1693–1701.
- Huang, Z.; van Harmelen, F.; and ten Teije, A. 2005. Reasoning with Inconsistent Ontologies. In *IJCAI*, 454–459.
- Johnson, J.; Douze, M.; and Jégou, H. 2017. Billion-scale similarity search with gpus. *arXiv preprint arXiv:1702.08734*.
- Joulin, A., and Mikolov, T. 2015. Inferring Algorithmic Patterns with Stack-Augmented Recurrent Nets. In *NIPS*.
- Kaiser, L., and Sutskever, I. 2016. Neural GPUs Learn Algorithms. In *ICLR*.
- Kemp, C.; Tenenbaum, J. B.; Griffiths, T. L.; Yamada, T.; and Ueda, N. 2006. Learning Systems of Concepts with an Infinite Relational Model. In *AAAI*, 381–388.
- Lao, N.; Mitchell, T. M.; and Cohen, W. W. 2011. Random Walk Inference and Learning in A Large Scale Knowledge Base. In *EMNLP*, 529–539.
- Levesque, H. J. 2014. On our best behaviour. *Artificial Intelligence* 212:27–35.
- Lipton, Z. C. 2018. The mythos of model interpretability. *Commun. ACM* 61(10):36–43.
- MacCartney, B., and Manning, C. D. 2007. Natural logic for textual inference. In *ACL-PASCAL@ACL*, 193–200. *ACL*.
- McCallum, A.; Neelakantan, A.; and Verga, P. 2017. Generalizing to Unseen Entities and Entity Pairs with Row-less Universal Schema. In *EACL*, 613–622.
- Minervini, P.; d’Amato, C.; Fanizzi, N.; and Esposito, F. 2016. Leveraging the schema in latent factor models for knowledge graph completion. In *SAC*, 327–332. *ACM*.
- Minervini, P.; Demeester, T.; Rocktäschel, T.; and Riedel, S. 2017. Adversarial Sets for Regularising Neural Link Predictors. In *UAI*.
- Minervini, P.; Bosnjak, M.; Rocktäschel, T.; and Riedel, S. 2018. Towards neural theorem proving at scale. *CoRR* abs/1807.08204.

- Nickel, M.; Murphy, K.; Tresp, V.; and Gabrilovich, E. 2016. A Review of Relational Machine Learning for Knowledge Graphs. *Proceedings of the IEEE* 104(1):11–33.
- Nickel, M.; Jiang, X.; and Tresp, V. 2014. Reducing the rank in relational factorization models by including observable patterns. In *NIPS*, 1179–1187.
- Niklaus, C.; Cetto, M.; Freitas, A.; and Handschuh, S. 2018. A Survey on Open Information Extraction. In *CICLing*.
- Rae, J. W.; Hunt, J. J.; Danihelka, I.; Harley, T.; Senior, A. W.; Wayne, G.; Graves, A.; and Lillicrap, T. 2016. Scaling memory-augmented neural networks with sparse reads and writes. In *NIPS*, 3621–3629.
- Raedt, L. D.; Frasconi, P.; Kersting, K.; and Muggleton, S., eds. 2008. *Probabilistic Inductive Logic Programming - Theory and Applications*, volume 4911 of *LNCS*. Springer.
- Riedel, S.; Yao, L.; McCallum, A.; and Marlin, B. M. 2013. Relation extraction with matrix factorization and universal schemas. In *HLT-NAACL*, 74–84. ACL.
- Rocktäschel, T., and Riedel, S. 2017. End-to-end Differentiable Proving. In *NIPS*, 3791–3803.
- Sukhbaatar, S.; Szlam, A.; Weston, J.; and Fergus, R. 2015. End-To-End Memory Networks. In *NIPS*, 2440–2448.
- Toutanova, K.; Chen, D.; Pantel, P.; Poon, H.; Choudhury, P.; and Gamon, M. 2015. Representing Text for Joint Embedding of Text and Knowledge Bases. In *EMNLP*, 1499–1509.
- Trouillon, T.; Welbl, J.; Riedel, S.; Gaussier, É.; and Bouchard, G. 2016. Complex Embeddings for Simple Link Prediction. In *ICML*, volume 48, 2071–2080.
- Weber, L.; Minervini, P.; Münchmeyer, J.; Leser, U.; and Rocktäschel, T. 2019. Nlprolog: Reasoning with weak unification for question answering in natural language. In *ACL (1)*, 6151–6161. Association for Computational Linguistics.
- Weston, J.; Bordes, A.; Chopra, S.; and Mikolov, T. 2015. Towards AI-Complete Question Answering: A Set of Prerequisite Toy Tasks. *CoRR* abs/1502.05698.
- White, L.; Togneri, R.; Liu, W.; and Bannamoun, M. 2015. How well sentence embeddings capture meaning. In *ADCS*.
- Yang, B.; Yih, W.; He, X.; Gao, J.; and Deng, L. 2015. Embedding Entities and Relations for Learning and Inference in Knowledge Bases. In *ICLR*.
- Yang, F.; Yang, Z.; and Cohen, W. W. 2017. Differentiable Learning of Logical Rules for Knowledge Base Reasoning. In *NIPS*, 2316–2325.