



# Stateful Premise Selection by Recurrent Neural Networks

Bartosz Piotrowski<sup>1,2\*</sup> and Josef Urban<sup>1†</sup>

<sup>1</sup> Czech Institute of Informatics, Robotics and Cybernetics, Prague, Czech Republic

<sup>2</sup> Faculty of Mathematics, Informatics and Mechanics, University of Warsaw, Poland

## Abstract

In this work we develop a new learning-based method for selecting facts (premises) when proving new goals over large formal libraries. Unlike previous methods that choose sets of facts independently of each other by their rank, the new method uses the notion of *state* that is updated each time a choice of a fact is made. Our stateful architecture is based on recurrent neural networks which have been recently very successful in stateful tasks such as language translation. The new method is combined with data augmentation techniques, evaluated in several ways on a standard large-theory benchmark and compared to state-of-the-art premise approach based on gradient boosted trees. It is shown to perform significantly better and to solve many new problems.

## 1 Introduction: Premise Selection over Large Libraries

*Premise selection* [1] is a critical task in automated theorem proving (ATP) over large theories where typically only a small fraction of the available facts are relevant for proving a new conjecture. One of the main applications is in ITP/ATP *hammers* [2] that assist ITP users to automatically discharge proof obligations in large formalizations. Several heuristic approaches to premise selection such as SiNE [13] and MePo [23], as well as learning-based methods using hand-designed features have been developed so far. The latter include naive Bayes [31], kernel methods [30], k-nearest neighbors [16, 17] and gradient boosted trees [25]. This has been followed by neural architectures that learn the features on their own [14, 20].

The learning-based premise selection methods have been so far based on the same paradigm of *ranking* the available facts (premises) *independently* with respect to the conjecture that is being proved. The highest ranked facts are then used together as axioms and given to the ATP systems together with the conjecture. This approach, although useful and reasonably successful, does not take into account an important aspect of the premise selection problem: premises are *not* independent of each other. There are important logical relations among them.

---

\*Supported by the grant 2018/29/N/ST6/02903 of National Science Center, Poland, and by the *AI4REASON* ERC Consolidator grant nr. 649043.

†Supported by the *AI4REASON* ERC Consolidator grant nr. 649043 and by the Czech project AI&Reasoning CZ.02.1.01/0.0/0.0/15\_003/0000466 and the European Regional Development Fund.

Some premises complement each other better when proving a particular conjecture, while some highly-ranked premises might be just minor variants of one another.

In this work, we first (Section 2) propose the recurrent neural network (RNN) encoder-decoder model [6] as a suitable stateful approach for premise selection and we describe the RNN architecture we have chosen for this task. In Section 3 we develop several data augmentation methods that help training the RNNs for the premise selection task. Section 4 describes the experimental evaluation and Section 5 discusses the results.

## 2 Premise Selection and Neural Machine Translation

In recent years, powerful methods for learning sequences of conditional stateful decisions have been developed in machine translation of the natural languages. In *neural machine translation* (NMT) [6] the source sentence is encoded as a hidden vector representation by the *encoder* and the translated target sentence is produced word-by-word by the *decoder*. Each translated word is conditioned not only on the source sentence but also on the *previously decoded words*. Words and phrases in natural languages are related in many ways and such relations have to be taken into account for the produced sequence of words to be a sensible, grammatically correct sentence. Successful NMT methods using recurrent encoder-decoder architectures [29] are explicitly based on a notion of a learned *hidden state* that is updated with each produced word. This corresponds to our requirement of *stateful premise selection*: We want to have a (learned) hidden state after selecting a particular fact with particular mathematical content that should not be repeated in the next facts but rather suitably complemented by them to justify the conjecture.

Another aspect of neural machine translation that is relevant in premise selection is the *multiplicity of correct outputs*. In translation, there are often multiple correct translations of a given sentence that deliver its *meaning* (perhaps more or less clumsily). Similarly, in mathematics there is typically no single, golden proof of a conjecture. Often there are many different proofs that use various sets of premises and various sequences of inferences. NMT methods accommodate the multiplicity of possible outputs – typically by using *beam search* [9]. Such mechanisms seem directly usable also in premise selection and proof search. NMT systems have already been successfully applied in autoformalization and symbolic settings [33, 26, 21].

**Our Recurrent Neural Architecture:** There are various state-of-the-art neural sequence-to-sequence architectures that can be applied to modelling premise selection. Although ultimately a custom architecture could be designed to capture all aspects of this task, our initial choice was to experiment with an existing established implementation of neural machine translation system. We have chosen the OpenNMT toolkit [19]. It implements the LSTM [12] recurrent cells and several state-of-the-art techniques, including the attention mechanisms [22] and beam search [9]. It has proven to be very good in natural language translation and related tasks [19].

We have decided to use the default parameters for training OpenNMT on our tasks – in this work we mainly investigate the influence of various forms of training data (Section 3) on the predictive performance. The more important values of the OpenNMT hyperparameters chosen by us are as follows: the number of training epochs: 100000, the size of encoder’s and decoder’s LSTM cells: 2 layers of 500 units, word embedding size: 500. Additionally, we have used the attention mechanism by Luong [22].

### 3 Data, Their Representation and Augmentation

A recurrent NMT system is trained on pairs of *source* and *target* sequences. In our case, the source is a statement of a theorem and the target is a list of names of its premises. There are multiple ways how to transform ATP proofs to training examples of such form, and it is not clear which way is the best for training the RNN. In this section, we describe several methods of constructing the training examples<sup>1</sup>. This includes the following topics: (i) representation of the conjecture as a sequence (Section 3.2), (ii) ordering of the premises into a sequence (Section 3.3), (iii) using subproof data for augmenting the training data (Section 3.4), and (iv) oversampling of rare premises (Section 3.5).

#### 3.1 Initial Data for Training RNNs

The experimental data originate from the Mizar Mathematical Library (MML) [11] translated [31] to the TPTP language [28]. We use the MPTP2078 benchmark [1] – a subset of 2078 Mizar theorems. Using the ATPboost [25] system we have initially proved as many of the MPTP2078 problems as possible, recording each distinct proof. ATPboost in turn relies on the E prover [27] and the XGBoost machine learning system [4] using gradient boosted trees for premise selection. 24087 different proofs of 1469 theorems were found in total. The number of different proofs per theorem ranged between 1 and 265 (on average 16.4). The proofs used in total 2227 different premises. Each proof used between 1 and 50 premises (on average 11.5). Each proof determines a pair  $(t, \{p_1, p_2, \dots, p_n\})$ , where the first element  $t$  is the proved theorem and the second element is the set of premises  $p_i$  used in the proof. These pairs constitute examples for training a machine learning model to propose useful premises for theorems. The 1469 theorems that have an ATP proof were randomly split in proportions 0.75 and 0.25 into training and testing parts. The 1100 training theorems with their proofs resulted in 18361 training pairs. From the set of remaining 369 theorems we filtered out those which contained in all their proofs premises not appearing in the training set. This yields our testing set of 310 theorems.

#### 3.2 Representation of the Statements

The simplest way of constructing the source sequences of the examples is just using tokenized statements in standard TPTP syntax. We label this type of source as **standard**. The tokenized TPTP statements can also be transformed to other formats. A suitable one is the Polish prefix notation (labeled as **prefix**) as shown in Table 1. The motivation is that this format is more compact as the formulas do not contain brackets and commas. In our case, the average length of a **standard** source training sequence is 81, whereas for **prefix** it is only 39. This may be useful for NMT architectures that suffer when using long input (and output) sequences [5]. Related work using NMT in symbolic setting reports improvements when using prefix notation [33].

#### 3.3 Ordering of the Premises

In the abstract premise selection task the order of the premises is not constrained in any way. In practice, ATP systems may be influenced by the order of the premises given in the input. More importantly, existing efficient learning methods that are capable of capturing dependencies

<sup>1</sup>All the data used in this work along with scripts allowing reproduction of the experiments are available at: <https://github.com/BartoszPiotrowski/stateful-premise-selection-with-RNNs>

Table 1: An example of a Mizar statement translated to TPTP, tokenized (`standard`), and additionally transformed to prefix notation (`prefix`).

Mizar:	for A being set st A is empty holds A is finite
TPTP:	![A] : (v1_xboole_0(A) => v1_finset_1(A))
standard:	! [ A ] : ( v1_xboole_0 ( A ) => v1_finset_1 ( A ) )
prefix:	! A => v1_xboole_0 A v1_finset_1 A

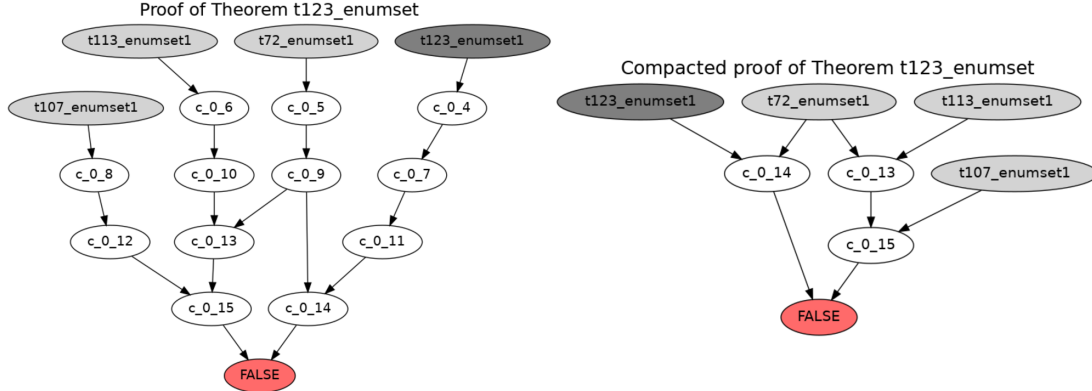


Figure 1: Trees representing a refutational proof of Mizar theorem `t123_enumset` (left) and its compacted version (right). Light-grey nodes are the input premises and a dark-grey node represents the (negated) conjecture. Nodes prefixed by `c_0` are intermediate lemmas.

among the elements of sequences (such as RNNs, used this work) are sensitive to the order of the elements (premises in our case). Preferably, we want to train on examples that illustrate dependencies between the premises. On the other hand, we do not want to rely too much on a particular order of premises in the target sequence. We propose several possible approaches.

**Permutations:** As a baseline approach we permute the target premises randomly, thus not passing to the recurrent neural model any additional information about the order of the premises. We either produce only one permutation (`permuted`) or 100 of them (`permuted_100`).

**Permutations Preserving the Proof Tree:** Each proof produced by a refutational prover (such as E) is a tree (more precisely, a DAG), with `FALSE` in its root and the premises and the conjecture in its leaves. See Figure 1 for an example. We produce a compacted version of the proof by removing all intermediate nodes that have only one parent<sup>2</sup> (right side of Figure 1). The premises `t72_enumset1` and `t113_enumset1` there interact directly resulting in an intermediate lemma `c_0_13`. This lemma subsequently interacts with `t107_enumset1`. Putting `t72_enumset1` and `t113_enumset1` closer together than `t107_enumset1` and `t72_enumset1` tells the learner how much the premises interact with each other. This idea is implemented in the following way. Each tree induces its *nested list representation*, which we define recursively in the following way: (i) list representation of a tree rooted in a non-leaf node  $n$  is a list of list representations of parents of  $n$ , (ii) list representation of a leaf node is its label. For instance,  $[[t123, t72], [[t72, t113], t107]]$  is a list representation of the tree from Figure 1. Each tree has many list representations depending on how the elements of the lists are ordered.

<sup>2</sup>Parent/child terminology is used in such a way that children are derived from their parents.

We say that a sequence  $s$  of labels of the leaves *respects the tree* if  $s$  resulted from the flattening of a list representation of the tree. (In the example shown in Figure 1 the sequence (t123, t72, t72, t113, t107) respects the tree but the sequence (t123, t107, t72, t113, t72) does not.) Such sequences have the property of keeping closer the premises that interacted closer. Note that the sequences may contain repetitions, as in the example above, and each tree has many sequences respecting it. For each proof tree we remove from its sequences the conjecture. We take either only one such sequence for each proof or (up to) 100 different sequences, which yields the `permuted_tree` and `permuted_tree_100` sets of training examples.

**ATP Induced Order:** We also experiment with using the proofs as linearized by E prover. For a given E proof  $\mathcal{P}$  we first order its internal lemmas: `lemma_1`  $<_L$  `lemma_2` iff `lemma_1` appears in  $\mathcal{P}$  (linearized by E) before `lemma_2`. Then we define a non-strict ordering of the premises of  $\mathcal{P}$ :  $p_1 \leq_P p_2$  iff the  $<_L$ -minimal child of  $p_1$  is smaller than the  $<_L$ -minimal child of  $p_2$ , where both children are taken from the *compacted* tree of  $\mathcal{P}$ . For our example proof of `t123_enumset` we have: `c_0_13`  $<_L$  `c_0_14`  $<_L$  `c_0_15`. Hence `t72`  $\leq_P$  `t107` (since  $\min_L(\{\text{c}_0_14, \text{c}_0_13\}) <_L \text{c}_0_15$ ), and `t113`  $=_P$  `t107` (since  $\min_L(\{\text{c}_0_14, \text{c}_0_13\}) = \text{c}_0_13$ ). We break the ties randomly. Different E proofs of one theorem may result in different premise orderings. This way of ordering premises in the target of the examples is labeled as `order_from_proof`.

### 3.4 Augmentation with Subproof Data

We can also augment the training data by extracting many subproofs from the original training proofs. For this, we use the intermediate lemmas from the compacted representations of proofs that are not derived from the negated conjecture. The pairs  $(l, \{p_1, p_2, \dots, p_n\})$  where  $p_i$  are all premise ancestors of lemma  $l$  constitute additional examples that can be used for augmenting the main training data. From the subproofs of the training theorems we extracted 46094 such different training pairs. The data set that includes these examples together with the main ones is marked as `augmented`. The experiments with sublemmas only are described in Appendix A.

### 3.5 Oversampling Rare Examples

Some premises appear frequently in the training examples while some are rare. *Oversampling* is a general method that often improves the performance of neural architectures on imbalanced data [3]. We experiment with a simple oversampling scheme: training examples that contain rare premises are used multiple times. More precisely, for an example  $e = (t, P) \in \mathcal{T}$ , where  $\mathcal{T}$  is the training set, we define the *occurrence rate* (OR) of  $e$  as:

$$\text{OR}(e) = \frac{1}{|P|} \sum_{p \in P} \text{OR}_{\text{premise}}(p), \quad \text{where } \text{OR}_{\text{premise}}(p) = \frac{|\{(t, P) \in \mathcal{T} : p \in P\}|}{\sum_{(t, P) \in \mathcal{T}} |P|}.$$

The idea is simple: OR of an example is the average  $\text{OR}_{\text{premise}}$  of its target premises ( $P$ ).  $\text{OR}_{\text{premise}}$  measures how often a premise appears in all the targets of all the training examples.

The set of training examples  $\mathcal{T}$  is split into 10 evenly sized chunks  $\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_{10}$  according to their occurrence rate so that a higher index of  $\mathcal{T}_i$  implies lower OR:

$$x \in \mathcal{T}_i \wedge y \in \mathcal{T}_j \wedge i < j \Rightarrow \text{OR}(x) > \text{OR}(y).$$

Each example  $e \in \mathcal{T}_i$  is oversampled  $i$  times: the more rare premises an example  $e$  contains the more often  $e$  appears in the oversampled training set. This scheme was applied both to the main training set and to the augmented one (described in Section 3.4), resulting in data sets `oversampled` and `augmented_oversampled`.

## 4 Experimental Evaluation

We train<sup>3</sup> and evaluate RNNs using the OpenNMT toolkit with its default hyper-parameters (Section 2) on the various premise-selection data described in Section 3. When evaluating on the testing sets, we use OpenNMT’s beam search with width 10 to get for each conjecture its 10 most probable sequences of premises. We want to compare the results also with state-of-the-art premise selection based on gradient boosted trees using the XGBoost toolkit. For that, we use the features and settings developed in our ATPboost system [25]. As explained in Section 1, the training data used for XGBoost are unordered. XGBoost produces a ranking of the premises and we use several segments of the top-ranked premises for the XGBoost evaluation. While OpenNMT needs only positive examples, XGBoost also needs negative examples. We produce them by sampling negatives randomly, which performed well in ATPboost.

To allow a meaningful comparison of the two approaches, we shorten the rankings produced by XGBoost according to the lengths of the sequences produced by OpenNMT for a given conjecture. In more detail: if  $\hat{R}$  is a ranking produced by XGBoost and  $\hat{P}_1, \hat{P}_2, \dots, \hat{P}_{10}$  are sequences of premises produced by OpenNMT, we take  $\hat{R}_1, \hat{R}_2, \dots, \hat{R}_{10}$  to be the top slices of the ranking  $\hat{R}$  of lengths  $|\hat{P}_1|, |\hat{P}_2|, \dots, |\hat{P}_{10}|$ , respectively. These 10 top slices are treated as predictions from the XGBoost system and compared with the OpenNMT predictions. We always do both the standard *ML evaluation* and the *ATP evaluation*.

**ML Evaluation:** The Jaccard index and Coverage metrics are used, defined as below:

$$\text{Jaccard}(A, B) = \frac{|A \cap B|}{|A \cup B|}, \quad \text{Coverage}(A, B) = \frac{|A \cap B|}{|B|}.$$

Each theorem  $T$  from the test set is associated with  $n_T$  sets of premises  $P_1^T, P_2^T, \dots, P_{n_T}^T$  which were used as axioms in its  $n_T$  known proofs and with 10 sets of premises  $\hat{P}_1^T, \hat{P}_2^T, \dots, \hat{P}_{10}^T$  predicted by a given machine learning model. We estimate the quality of the predictions with  $\text{Jaccard}(\bigcup_i \hat{P}_i^T, \bigcup_i P_i^T)$  and  $\text{Coverage}(\bigcup_i \hat{P}_i^T, \bigcup_i P_i^T)$ . The Jaccard metric emphasises the *precision* of the prediction: it measures how much the predicted premises intersect with the premises used in the known proofs. At the same time the score decreases when the predicted set is large. The Coverage does not penalize large predicted sets. It takes into account the fact that *true positives* are typically more important than *true negatives*. The prover may deal with some redundant axioms while the lack of relevant premises may make conjectures unprovable.

**ATP Evaluation:** The simple ML metrics may not directly translate to ATP performance. They compare unions of premises whereas an ATP is run for each premise selection separately. Additionally, during the ATP evaluation new proofs are often discovered. Such new proofs are not taken into account by the similarity metrics.

We perform an ATP evaluation using the E prover [27] run with a time limit of 10 s and a memory limit of 2 GB, keeping the rest of the settings in its default values. This limits the power of the prover, preventing e.g. its own axiom pruning methods such as SInE [13]. To establish a simple ATP baseline, E prover was run for all the testing theorems with all available premises as axioms, proving 9% of the theorems. For a given machine learning method and for each testing theorem  $T$  we run 11 proof attempts. One for each of the 10 sets of predicted premises  $\hat{P}_1^T, \hat{P}_2^T, \dots, \hat{P}_{10}^T$  and one with  $\bigcup_i \hat{P}_i^T$ .

<sup>3</sup>We train on a single GeForce RTX 2080 Ti GPU. Each training took between 2 and 4 hours.

Table 2: Performance of the neural model on the test set, trained on examples with different formats of source sequences (statements) and differently ordered target sequences (premises), expressed with the similarity metrics (Jaccard index and Coverage) and with ATP success rate.

Target ordering	Source format					
	standard			prefix		
	Jaccard	Coverage	Proved	Jaccard	Coverage	Proved
<code>permuted</code>	0.18	0.39	0.25	0.20	0.36	0.23
<code>permuted_100</code>	0.09	0.14	0.10	0.18	0.26	0.19
<code>permuted_tree</code>	0.16	0.23	0.16	0.17	0.25	0.18
<code>permuted_tree_100</code>	0.04	0.05	0.03	0.11	0.15	0.09
<code>order_from_proof</code>	0.22	0.46	0.29	0.22	0.43	0.27

## 5 Results and Discussion

### 5.1 Source and Target Combinations

First, we evaluate combinations of the statement format (`standard` or `prefix` – Section 3.2) and orderings of the target premises (Section 3.3). The results are shown in Table 2. The first, simplest way of ordering premises – `permuted` – performs well: in combination with the two formats of source statements it resulted in predictions with ATP success rates 25% and 23%. Many permutations of the same target sequence are bad for the NMT learner: `permuted_100` performed much worse than `permuted`. Using multiple permutations of the target sequences was motivated by the fact that the order of premises should not matter in the abstract premise selection task. The recurrent neural network however likely sees them as *contradictory data* detrimental to its training.

The `permuted_tree` ordering was meant to reflect the distance of interaction between the premises. It however performed much worse than `permuted`. This may be caused by many repetitions in the target sequences, which also increase their lengths, making the task for the neural decoder more difficult. These repetitions appear because of the repeated premises in the leaves of the proof trees. Similarly to `permuted_100`, adding multiple permutations (`permuted_tree_100`) is detrimental also in the case of `permuted_tree`.

The best performing way of ordering the premises for the NMT learner is to use the `order_from_proof`. This is true both for the similarity metrics and ATP performance and in combination with both formats of the conjecture. This likely means that extracting the premise ordering from the proofs brings useful and consistent information which the neural model is able to take advantage of during the training.

There is only a small difference between the ATP performance of the `standard` and `prefix` encoding of the conjecture. This is somewhat surprising in the context of related work [33], where prefix notation is useful for NMT architectures. Shorter sequences should be easier to process by the recurrent encoder. In our case, however, the different structure of the `prefix` statements seems to be reducing the benefit of conciseness.

### 5.2 Augmentation with Subproof Data and Oversampling

Next we use the best performing combination (from now on called `basic`) from Section 5.1 for evaluating the augmentation and oversampling methods (Sections 3.4 and 3.5). The results are

Table 3: Performance of the NMT and XGBoost systems trained on examples augmented with subproof data and with oversampling applied. The non-modified data set is denoted as **basic**. The examples used **standard** source format and **order\_from\_proof** ordering of the target sequences.

Training data	Machine learning system					
	NMT			XGBoost		
	Jaccard	Coverage	Proved	Jaccard	Coverage	Proved
<b>basic</b>	0.22	0.46	0.29	0.26	0.56	0.25
<b>oversampled</b>	0.21	0.48	0.31	0.24	0.61	0.30
<b>augmented</b>	0.27	0.51	0.40	0.26	0.51	0.27
<b>augmented_oversampled</b>	0.26	0.47	0.39	0.25	0.51	0.31

shown in Table 3 which also contains the XGBoost results.

Oversampling trained the learner using data with changed distribution – less frequent premises were appearing more often. This made the predictions more diverse and less precise compared to **basic**. This is reflected in the change of the similarity metrics: Jaccard index decreased and Coverage increased. Importantly, oversampling translates to better ATP performance of the NMT predictions.

Augmentation with subproof data improved the ATP performance by a large margin of 11% points over **basic**. It means that the RNN is helped a lot by the additional subproof training data despite their slightly different origin and shape (internal clauses instead of input formulas). The last row of Table 3 shows the result of applying oversampling on top of the augmented training set. This does not improve the NMT performance compared to **augmented**.

ATP performance of XGBoost was worse than that of NMT for all 4 data sets, and the best XGBoost ATP result (0.31) is significantly (29%) worse than the best NMT ATP result (0.40). However, XGBoost tends to show better values than NMT in similarity metrics. This can be explained by the following effects:

1. The initial data come from ATPboost – a meta-system using XGBoost. The XGBoost predictions in our experiments may be correlated with the initial testing set.
2. Even though XGBoost achieves a higher similarity between unions of the predicted premises and the premises used in the known proofs, the recurrent neural network wins with its diverse (but stateful and therefore complementary) predictions for a given conjecture. I.e., when making several ATP attempts, it seems better to use several plausible premise sets (proof ideas) that are orthogonal to one another, rather than making incremental additions to the initial set of premises. This is the effect that we wanted to achieve with RNN-based premise selection, and indeed, it makes a significant difference. XGBoost instead just extends its single most plausible set of premises more and more according to the single ranking of premises. It seems nontrivial to instruct XGBoost to produce multiple alternative rankings with good complementary properties in the same way as RNN does.

The NMT predictions are quite orthogonal to those from XGBoost. In all the experiments related to the results shown in Table 3 there were 167 theorems proved with predictions from NMT and 142 theorems proved with the use of XGBoost. The size of the intersection of these sets is 121, and there were 46 theorems that were proved with NMT but not with XGBoost.



## 6 Conclusions and Future Work

We have shown that state-of-the-art recurrent neural architectures – designed originally for natural language tasks such as machine translation – are very useful for premise selection. In particular, they can be used to implement (i) stateful / conditional premise selection and (ii) beam search with multiple output sequences that may differ a lot while being meaningful as a whole. Our experiments show significant improvement over state of the art obtained by such methods. We have also developed several data representation and augmentation methods that result in additional improved performance of both the old and new premise selection methods.

NMT architectures are also more natural in some aspects. There is no need for hand-designed features of formulas and no need to construct negative training examples. This is important, because in theorem proving it is often difficult (or impossible) to say that a particular selection of premises *cannot* lead to a proof. Once the recurrent neural network is trained, it directly outputs the most probable sequences of candidate premises – not just their rankings. We have used 10 most probable sequences for the experiments described here, but larger numbers can be used and given to ATPs, depending on available resources.

Future work includes integration into AI/TP metasystems interleaving premise selection with learning such as ATPboost and MaLAREa [32]. Another direction is tighter integration between the ATPs and the neural network so that the prover can take advantage of the order in which the premises are presented. Conditional selection might also be implemented in learning-guided ATP systems such as rlCoP [18], plCoP [34], TacticToe [10] and ENIGMA [15, 7]. Neural network research is advancing quickly and experiments with other stateful neural architectures may bring further improvement. Finally, we could provide the neural networks with more information about the premises. Currently, the premises are just names (words) and NMT can only learn their *latent semantics* [8]. Adding more information about their logical representation and meaning [24] may be useful.

## References

- [1] J. Alama, T. Heskes, D. Kühlwein, E. Tsvitvadze, and J. Urban. Premise selection for mathematics by corpus analysis and kernel methods. *J. Autom. Reasoning*, 52(2):191–213, 2014.
- [2] J. C. Blanchette, C. Kaliszyk, L. C. Paulson, and J. Urban. Hammering towards QED. *J. Formalized Reasoning*, 9(1):101–148, 2016.
- [3] M. Buda, A. Maki, and M. A. Mazurowski. A systematic study of the class imbalance problem in convolutional neural networks. *Neural Networks*, 106:249–259, 2018.
- [4] T. Chen and C. Guestrin. XGBoost: A scalable tree boosting system. In *Proceedings of the 22Nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD 2016, pages 785–794, New York, NY, USA, 2016. ACM.
- [5] K. Cho, B. van Merriënboer, D. Bahdanau, and Y. Bengio. On the properties of neural machine translation: Encoder-decoder approaches. In D. Wu, M. Carpuat, X. Carreras, and E. M. Vecchi, editors, *Proceedings of SSST@EMNLP 2014, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation, Doha, Qatar, 25 October 2014*, pages 103–111. Association for Computational Linguistics, 2014.
- [6] K. Cho, B. van Merriënboer, Ç. Gülçehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio. Learning phrase representations using RNN encoder-decoder for statistical machine translation. In A. Moschitti, B. Pang, and W. Daelemans, editors, *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25–29, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL*, pages 1724–1734. ACL, 2014.

- [7] K. Chvalovský, J. Jakubuv, M. Suda, and J. Urban. ENIGMA-NG: efficient neural and gradient-boosted inference guidance for E. In P. Fontaine, editor, *Automated Deduction - CADE 27 - 27th International Conference on Automated Deduction, Natal, Brazil, August 27-30, 2019, Proceedings*, volume 11716 of *Lecture Notes in Computer Science*, pages 197–215. Springer, 2019.
- [8] S. C. Deerwester, S. T. Dumais, T. K. Landauer, G. W. Furnas, and R. A. Harshman. Indexing by Latent Semantic Analysis. *JASIS*, 41(6):391–407, 1990.
- [9] M. Freitag and Y. Al-Onaizan. Beam search strategies for neural machine translation. In T. Luong, A. Birch, G. Neubig, and A. M. Finch, editors, *Proceedings of the First Workshop on Neural Machine Translation, NMT@ACL 2017, Vancouver, Canada, August 4, 2017*, pages 56–60. Association for Computational Linguistics, 2017.
- [10] T. Gauthier, C. Kaliszyk, J. Urban, R. Kumar, and M. Norrish. Learning to prove with tactics. *CoRR*, abs/1804.00596, 2018.
- [11] A. Grabowski, A. Kornilowicz, and A. Naumowicz. Mizar in a nutshell. *J. Formalized Reasoning*, 3(2):153–245, 2010.
- [12] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997.
- [13] K. Hoder and A. Voronkov. Sine qua non for large theory reasoning. In N. Bjørner and V. Sofronie-Stokkermans, editors, *CADE*, volume 6803 of *LNCS*, pages 299–314. Springer, 2011.
- [14] G. Irving, C. Szegedy, A. A. Alemi, N. Eén, F. Chollet, and J. Urban. DeepMath – Deep sequence models for premise selection. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 2235–2243, 2016.
- [15] J. Jakubuv and J. Urban. ENIGMA: Efficient learning-based inference guiding machine. In H. Geuvers, M. England, O. Hasan, F. Rabe, and O. Teschke, editors, *Intelligent Computer Mathematics - 10th International Conference, CICM 2017*, volume 10383 of *Lecture Notes in Computer Science*, pages 292–302. Springer, 2017.
- [16] C. Kaliszyk and J. Urban. Learning-assisted automated reasoning with Flyspeck. *J. Autom. Reasoning*, 53(2):173–213, 2014.
- [17] C. Kaliszyk and J. Urban. MizAR 40 for Mizar 40. *J. Autom. Reasoning*, 55(3):245–256, 2015.
- [18] C. Kaliszyk, J. Urban, H. Michalewski, and M. Olsák. Reinforcement learning of theorem proving. In S. Bengio, H. M. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, 3-8 December 2018, Montréal, Canada*, pages 8836–8847, 2018.
- [19] G. Klein, Y. Kim, Y. Deng, J. Senellart, and A. Rush. OpenNMT: Open-source toolkit for neural machine translation. In *Proceedings of ACL 2017, System Demonstrations*, pages 67–72, Vancouver, Canada, July 2017. Association for Computational Linguistics.
- [20] A. S. Kucik and K. Korovin. Premise selection with neural networks and distributed representation of features. *CoRR*, abs/1807.10268, 2018.
- [21] G. Lample and F. Charton. Deep learning for symbolic mathematics. *CoRR*, abs/1912.01412, 2019.
- [22] T. Luong, H. Pham, and C. D. Manning. Effective approaches to attention-based neural machine translation. In L. Màrquez, C. Callison-Burch, J. Su, D. Pighin, and Y. Marton, editors, *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing, EMNLP 2015, Lisbon, Portugal, September 17-21, 2015*, pages 1412–1421. The Association for Computational Linguistics, 2015.
- [23] J. Meng and L. C. Paulson. Lightweight relevance filtering for machine-generated resolution problems. *J. Applied Logic*, 7(1):41–57, 2009.
- [24] M. Olsák, C. Kaliszyk, and J. Urban. Property invariant embedding for automated reasoning.

- CoRR*, abs/1911.12073, 2019.
- [25] B. Piotrowski and J. Urban. ATPboost: Learning premise selection in binary setting with ATP feedback. In D. Galmiche, S. Schulz, and R. Sebastiani, editors, *Automated Reasoning - 9th International Joint Conference, IJCAR 2018, Held as Part of the Federated Logic Conference, FloC 2018, Oxford, UK, July 14-17, 2018, Proceedings*, volume 10900 of *Lecture Notes in Computer Science*, pages 566–574. Springer, 2018.
- [26] B. Piotrowski, J. Urban, C. E. Brown, and C. Kaliszyk. Can neural networks learn symbolic rewriting? *CoRR*, abs/1911.04873, 2019.
- [27] S. Schulz. System description: E 1.8. In K. L. McMillan, A. Middeldorp, and A. Voronkov, editors, *LPAR*, volume 8312 of *LNCS*, pages 735–743. Springer, 2013.
- [28] G. Sutcliffe. The TPTP world - infrastructure for automated reasoning. In E. M. Clarke and A. Voronkov, editors, *LPAR (Dakar)*, volume 6355 of *LNCS*, pages 1–12. Springer, 2010.
- [29] I. Sutskever, O. Vinyals, and Q. V. Le. Sequence to sequence learning with neural networks. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 3104–3112. Curran Associates, Inc., 2014.
- [30] E. Tsivtsivadze, J. Urban, H. Geuvers, and T. Heskes. Semantic graph kernels for automated reasoning. In *SDM*, pages 795–803. SIAM / Omnipress, 2011.
- [31] J. Urban. MPTP 0.2: Design, implementation, and initial experiments. *J. Autom. Reasoning*, 37(1-2):21–43, 2006.
- [32] J. Urban, G. Sutcliffe, P. Pudlák, and J. Vyskočil. MaLAREa SG1 - Machine Learner for Automated Reasoning with Semantic Guidance. In A. Armando, P. Baumgartner, and G. Dowek, editors, *IJCAR*, volume 5195 of *LNCS*, pages 441–456. Springer, 2008.
- [33] Q. Wang, C. Kaliszyk, and J. Urban. First experiments with neural translation of informal to formal mathematics. In F. Rabe, W. M. Farmer, G. O. Passmore, and A. Youssef, editors, *Intelligent Computer Mathematics - 11th International Conference, CICM 2018, Hagenberg, Austria, August 13-17, 2018, Proceedings*, volume 11006 of *Lecture Notes in Computer Science*, pages 255–270. Springer, 2018.
- [34] Z. Zombori, J. Urban, and C. E. Brown. Prolog technology reinforcement learning prover. *CoRR*, abs/2004.06997, 2020.

## A Subproofs as Standalone Data Set

From all the proofs in the initial data set (not only the training part) we extracted 60299 different pairs of the form

$$(\text{lemma}, \{\text{premise}_1, \text{premise}_2, \dots, \text{premise}_n\}),$$

same as those used for augmenting the training set (Section 3.4). This means that:

- `lemma` is an intermediate sublemma appearing in the compacted representation of the proof (see Section 3.3 and Figure 1), which has no negated conjecture of the original proof among its ancestors,
- `{premise_1, premise_2, ..., premise_n}` is a set of all the premises being among ancestors of the `lemma`.

These pairs have 29616 different lemmas (each lemma may have several different proofs). We split these lemmas into training and testing parts in proportions approximately 0.75 and 0.25, respectively (independently from the main training/testing split of the Mizar theorems). We also recorded information about the heights of the subproof trees from which the lemmas were

Table 4: Performance of the NMT and XGBoost models evaluated on the testing part of the subproofs data set. We use our similarity metrics (Jaccard index and Coverage) and ATP evaluation (columns named Proved). The first column contains information about the average height of the proof subtrees the sublemmas originated from. The second column is the number of lemmas in a given subset. The first numeric row refers to all the testing examples, independently of the height. The largest numbers in the columns are marked in bold.

Height	#Lemmas	Machine learning model					
		NMT			XGBoost		
		Jaccard	Coverage	Proved	Jaccard	Coverage	Proved
$[1, \infty)$	7300	0.30	0.80	0.83	0.27	0.74	0.61
1	1610	0.19	0.84	0.83	0.18	<b>0.82</b>	0.66
(1, 2]	<b>1803</b>	0.29	<b>0.87</b>	0.84	0.25	0.79	0.60
(2, 3]	1431	0.34	0.82	0.81	0.30	0.75	0.57
(3, 4]	936	0.36	0.75	0.82	0.31	0.69	0.56
(4, 5]	580	0.34	0.70	0.83	0.31	0.64	0.60
(5, 6]	319	0.37	0.70	0.86	0.33	0.64	0.62
(6, 7]	223	0.34	0.64	0.83	0.32	0.60	0.64
(7, 8]	124	0.37	0.70	0.89	0.34	0.66	<b>0.79</b>
(8, 9]	90	<b>0.40</b>	0.69	<b>0.92</b>	<b>0.36</b>	0.63	0.71

extracted – to investigate how the height of the tree correlates with the difficulty of learning premise selection. The heights of the subtrees vary between 1 and 35, where lower subtrees are much more frequent than higher ones.

The NMT system was trained on the training examples, with the same settings as in the main experiments. Additionally, we trained the XGBoost system for comparison.

The results of the evaluation on the testing examples, in terms of the similarity metrics (Jaccard index and Coverage) as well as an ATP evaluation, are presented in Table 4. The table presents the performance of the machine learning methods with respect to all the testing examples as well as on subsets of them selected according to the heights of subtrees of proof trees a given sublemma originated from.

Overall, in comparison to the main data, premise selection for subproofs turned out to be a significantly easier task for both machine learning methods. On the whole testing set the ATP performance was 83% for NMT and 61% for XGBoost. When running the automated prover without any premise selection advice, with all available premises as axioms, the ATP success rate was 13%.

As for the results depending on the heights of the subtrees, in the table we present them up to the height 9 – for larger values the number of lemmas becomes very small. There are two trends visible for both NMT and XGBoost: with increased height the Jaccard metric goes up and Coverage goes down. The likely explanation is that the lower trees contain fewer premises in their leaves and precise selection of them by the predictor is less likely, hence the low Jaccard metric. On the other hand, the higher trees have more premises in their leaves and covering them by the predictor is more difficult, which is reflected by the low Coverage. The dependence of ATP performance on the heights is unclear. Surprisingly, it is *not* the case that the smallest subtrees contained the easiest premise selection examples.

## B Examples of predictions from RNN

In this section, we show two examples of predictions from the recurrent neural NMT system and compare them with the respective XGBoost predictions. All the presented predictions come from the systems trained on the `basic` data set. Note that in both cases below, the NMT predictions are more diverse, expressing different proof approaches and allowing quite different proof attempts. On the other hand, as soon as XGBoost ranks high a bad set of lemmas that mislead E prover, adding more premises does not help in these cases.

### B.1 Theorem `t128_zfmisc_1`

**The Mizar statement of the theorem:**

for  $x, y, z, Y$  being set holds (  $[x,y]$  in  $[:\{z\},Y:]$  iff (  $x = z$  &  $y$  in  $Y$  ) )

**The NMT predictions:**

```
1: d1_enumset1 t71_enumset1 t69_enumset1 t70_enumset1 l54_zfmisc_1
2: d2_tarski t77_enumset1 t79_enumset1 t76_enumset1 t84_enumset1 l54_zfmisc_1
3: l38_zfmisc_1 t69_enumset1 t70_enumset1 t71_enumset1 t20_zfmisc_1 l54_zfmisc_1
4: d1_tarski t69_enumset1 l54_zfmisc_1 d2_tarski
5: d1_tarski t69_enumset1 l54_zfmisc_1
6: l33_zfmisc_1 t69_enumset1 t70_enumset1 t71_enumset1 t20_zfmisc_1 l54_zfmisc_1
7: t76_enumset1 d1_enumset1 l54_zfmisc_1
8: t20_zfmisc_1 t69_enumset1 t70_enumset1 t71_enumset1 t65_zfmisc_1 l54_zfmisc_1
9: d2_tarski t70_enumset1 t71_enumset1 t69_enumset1 l54_zfmisc_1
10: l24_zfmisc_1 t69_enumset1 t70_enumset1 t71_enumset1 d1_enumset1 l54_zfmisc_1
```

**The XGBoost predictions:**

Ranking: `d1_tarski t69_enumset1 t70_enumset1 t71_enumset1 l54_zfmisc_1 t106_zfmisc_1 ...`

**Comparison:** E prover without auto mode was able to prove `t128_zfmisc_1` with the 5th prediction proposed by NMT:

```
5: d1_tarski t69_enumset1 l54_zfmisc_1
```

but no proof could be found with the top slices of the ranking proposed by XGBoost. The reason for this is that premises appearing in the top part of the ranking:

```
t69_enumset1 t70_enumset1 t71_enumset1
```

are very similar, and E prover is stuck with them, even with higher CPU time limits. Below we show the Mizar statements of the discussed premises:

```
d1_tarski: for x being set holds ( x in it iff x = y );
l54_zfmisc_1: [x,y] in [:X,Y:] iff x in X & y in Y
t69_enumset1: for x1 being set holds {x1,x1} = {x1}
t70_enumset1: for x1, x2 being set holds {x1,x1,x2} = {x1,x2}
t71_enumset1: for x1, x2, x3 being set holds {x1,x1,x2,x3} = {x1,x2,x3}
```

## B.2 Theorem t30\_tops\_1

**The Mizar statement of the theorem:**

for GX being TopStruct for R being Subset of GX holds  
 ( R is open iff R ' is closed )

**The NMT predictions:**

```
1: dt_k3_subset_1 d1_tops_1 t52_pre_topc t29_tops_1 d8_tops_1 d7_tops_1
2: t100_xboole_1 t12_setfam_1 t28_xboole_1 t48_xboole_1 commutativity_k2_tarski d5_subset_1 ...
3: involutiveness_k3_subset_1 t29_tops_1 t101_tops_1 t52_pre_topc dt_k3_subset_1 d8_tops_1 ...
4: t100_xboole_1 t12_setfam_1 t36_xboole_1 t48_xboole_1 t7_ordinal1 t2_xboole_1 d5_xboole_0 ...
5: t100_xboole_1 t12_setfam_1 d5_subset_1 t2_boole t91_tops_1 involutiveness_k3_subset_1 ...
6: t28_xboole_1 t12_setfam_1 commutativity_k3_xboole_0 t100_xboole_1 t22_xboole_1 t36_xboole_1...
7: t12_setfam_1 t70_enumset1 t100_xboole_1 d5_subset_1 t71_enumset1 t72_enumset1 t73_enumset1 ...
8: d10_xboole_0 t2_xboole_1 t43_xboole_1 t12_xboole_1 commutativity_k2_xboole_0 t41_xboole_1 ...
9: d10_xboole_0 t2_xboole_1 t43_xboole_1 t12_xboole_1 commutativity_k2_xboole_0 t41_xboole_1 ...
10: d10_xboole_0 t2_xboole_1 t43_xboole_1 t12_xboole_1 commutativity_k2_xboole_0 t41_xboole_1 ...
```

**The XGBoost predictions:**

```
Ranking: d10_xboole_0 t3_subset d4_subset_1 d3_struct_0 d5_subset_1 involutiveness_k3_subset_1
          dt_k3_subset_1 dt_k2_pre_topc dt_l1_pre_topc d3_tarski dt_k2_subset_1 ...
```

**Comparison:** E prover without auto mode was able to prove t30\_tops\_1 with the 3rd prediction proposed by NMT:

```
3: involutiveness_k3_subset_1 t29_tops_1 t101_tops_1 t52_pre_topc dt_k3_subset_1
   d8_tops_1 d7_tops_1
```

actually using these 3 premises:

```
t29_tops_1 involutiveness_k3_subset_1 dt_k3_subset_1
```

E prover was not able to prove the theorem with any top slice of the XGBoost ranking.