

---

# DeepMath - Deep Sequence Models for Premise Selection

---

**Alexander A. Alemi** \*  
Google Inc.  
alemi@google.com

**François Chollet** \*  
Google Inc.  
fchollet@google.com

**Geoffrey Irving** \*  
Google Inc.  
geoffreyi@google.com

**Christian Szegedy** \*  
Google Inc.  
szegedy@google.com

**Josef Urban** \*  
Czech Technical University in Prague  
josef.urban@gmail.com

## Abstract

We study the effectiveness of neural sequence models for premise selection in automated theorem proving, one of the main bottlenecks in the formalization of mathematics. We propose a two stage approach for this task that yields good results for the premise selection task on the Mizar corpus while avoiding the hand-engineered features of existing state-of-the-art models. To our knowledge, this is the first time deep learning has been applied to theorem proving.

## 1 Introduction

Mathematics underpins all scientific disciplines. Machine learning itself rests on measure and probability theory, calculus, linear algebra, functional analysis, and information theory. Complex mathematics underlies computer chips, transit systems, communication systems, and financial infrastructure – thus the correctness of many of these systems can be reduced to mathematical proofs.

Unfortunately, these correctness proofs are often impractical to produce without automation, and present-day computers have only limited ability to assist humans in developing mathematical proofs and formally verifying human proofs. There are two main bottlenecks: (1) lack of automated methods for *semantic* or *formal* parsing of informal mathematical texts (*autoformalization*), and (2) lack of strong automated reasoning methods to fill in the gaps in already formalized human-written proofs.

The two bottlenecks are related. Strong automated reasoning can act as a semantic filter for autoformalization, and successful autoformalization would provide a large corpus of computer-understandable facts, proofs, and theory developments. Such a corpus would serve as both background knowledge to fill in gaps in human-level proofs and as a training set to guide automated reasoning. Such guidance is crucial: exhaustive deductive reasoning tools such as today’s resolution/superposition automated theorem provers (ATPs) quickly hit combinatorial explosion, and are unusable when reasoning with a very large number of facts without careful selection [5].

In this work, we focus on the latter bottleneck. We develop deep neural networks that learn from a large repository of manually formalized computer-understandable proofs. We learn the task that is essential for making today’s ATPs usable over large formal corpora: the selection of a limited number of most relevant facts for proving a new conjecture. This is known as *premise selection*.

The main contributions of this work are:

- A demonstration for the first time that neural network models are useful for aiding in large scale automated logical reasoning without the need for hand-engineered features.

---

\*Authors listed alphabetically. All contributions are considered equal.

- The comparison of various network architectures (including convolutional, recurrent and hybrid models) and their effect on premise selection performance.
- A method of semantic-aware “definition”-embeddings for function symbols that improves the generalization of formulas with symbols occurring infrequently. This model outperforms previous approaches at relaxed cutoff-thresholds.
- Analysis that shows that the neural network based premise selection models are complementary to those with hand-engineered features and can be ensembled with previous results to produce superior results.

## 2 Formalization and Theorem Proving

In the last two decades, large corpora of complex mathematical knowledge have been *formalized*: encoded in complete detail so that computers can *fully understand the semantics* of complicated mathematical objects. The process of writing such formal and verifiable theorems, definitions, proofs, and theories is called *Interactive Theorem Proving* (ITP).

The ITP field dates back to 1960s [18] and the Automath system by N.G. de Bruijn [10]. ITP systems include HOL (Light) [17], Isabelle [41], Mizar [14], Coq [8], and ACL2 [25]. The development of ITP has been intertwined with the development of its cousin field of *Automated Theorem Proving* (ATP) [33], where proofs of conjectures are attempted fully automatically. Unlike ATP systems, ITP systems allow human-assisted formalization and proving of theorems that are often beyond the capabilities of the fully automated systems.

Large ITP libraries include the Mizar Mathematical Library (MML) with over 50,000 lemmas, and the core Isabelle, HOL, Coq, and ACL2 libraries with thousands of lemmas. These core libraries are a basis for large projects in formalized mathematics and software and hardware verification. Examples in mathematics include the HOL Light proof of the Kepler conjecture (Flyspeck project) [16], the Coq proofs of the Feit-Thompson theorem [13] and Four Color theorem [12], and the verification of most of the Compendium of Continuous Lattices in Mizar [3]. ITP verifications of the seL4 kernel [27] and CompCert compiler [29] show comparable progress in large scale software verification.

While these large projects mark a coming of age of formalization, ITP remains labor-intensive. For example, Flyspeck took about 20 person-years, with twice as much for Feit-Thompson. Behind this cost are our two bottlenecks: lack of tools for autoformalization and strong proof automation.

Recently the field of *Automated Reasoning in Large Theories* (ARLT) [39] has developed, including AI/ATP/ITP (AITP) systems called *hammers* that assist ITP formalization [5]. Hammers analyze the full set of theorems and proofs in the ITP libraries, estimate the relevance of each theorem, and apply optimized translations from the ITP logic to simpler ATP formalisms. Then they attack new conjectures using the most promising combinations of existing theorems and ATP search strategies. Recent evaluations have proved 40% of all Mizar and Flyspeck theorems fully automatically [22, 23].

This AITP performance speeds up formalization and motivates further research combining statistical learning and deductive tools. However, there is significant room for improvement: with perfect premise selection (a perfect choice of library facts) ATPs can prove at least 56% of Mizar and Flyspeck instead of today’s 40% [5]. In the next section we explain the premise selection task in more detail and explain the experimental setting for measuring such improvements.

## 3 Premise Selection, Experimental Setting and Previous Results

Given a formal corpus of facts and proofs expressed in an ATP-compatible format, our task is

**Definition** (Premise selection problem). *Given a large set of premises  $\mathcal{P}$ , an ATP system  $A$  with given resource limits, and a new conjecture  $C$ , predict those premises from  $\mathcal{P}$  that will most likely lead to an automatically constructed proof of  $C$  by  $A$ .*

We use the Mizar Mathematical Library (MML) version 4.181.1147<sup>2</sup> as the formal corpus and E prover [34] version 1.9 as the underlying ATP system. The following list exemplifies a small

<sup>2</sup>[ftp://mizar.uwb.edu.pl/pub/system/i386-linux/mizar-7.13.01\\_4.181.1147-i386-linux.tar](ftp://mizar.uwb.edu.pl/pub/system/i386-linux/mizar-7.13.01_4.181.1147-i386-linux.tar)

```

:: t99_jordan: Jordan curve theorem in Mizar
for C being Simple_closed_curve holds C is Jordan;

:: Translation to first order logic
fof(t99_jordan, axiom, (! [A] : ( (v1_topreal2(A) & m1_subset_1(A,
k1_zfmisc_1(u1_struct_0(k15_euclid(2)))) => v1_jordan1(A)) ) ).

```

Figure 1: (top) The final statement of the Mizar formalization of the Jordan curve theorem. (bottom) The translation to first-order logic, using name mangling to ensure uniqueness across the entire corpus.

non-representative sample of topics and theorems that are included in the Mizar Mathematical Library:

- Cauchy-Riemann Differential Equations of Complex Functions
- Characterization and Existence of Gröbner Bases
- Maximum Network Flow Algorithm by Ford and Fulkerson
- Gödel’s Completeness Theorem
- Brouwer Fixed Point Theorem
- Arrow’s Impossibility Theorem
- The Borsuk-Ulam Theorem
- Dickson’s Lemma
- The Sylow Theorems
- Hahn Banach Theorem
- Gauss Lemma and Law of Quadratic Reciprocity
- Public-Key Cryptography and Pepin’s Primality Test
- Ramsey’s Theorem

This version of MML was used for the latest AITP evaluation reported in [23]. There are 57,917 proved Mizar theorems and unnamed top-level lemmas in this MML organized into 1,147 articles. This set is chronologically ordered by the order of articles in MML and by the order of theorems in the articles. Proofs of later theorems can only refer to earlier theorems. This ordering also applies to 88,783 other Mizar formulas (encoding the type system and other automations known to Mizar) used in the problems. The formulas have been translated into the TPTP format [37] used by first-order ATPs by the MPTP system [38] (see Figure 1).

Our goal is to automatically prove as many theorems as possible, using at each step all previous theorems and proofs. We can learn from both human proofs and ATP proofs, but previous experiments [28, 22] show that learning only from the ATP proofs is preferable to including human proofs if the set of ATP proofs is sufficiently large. Since for 32,524 (56.2%) of the 57,917 theorems an ATP proof was previously found by a combination of manual and learning-based premise selection [23], we use only these ATP proofs for training.

The 40% success rate from [23] used a portfolio of 14 AITP methods using different learners, ATPs, and numbers of premises. The best single method proved 27.3% of the theorems. Only fast and simple learners such as  $k$ -nearest-neighbors, naive Bayes, and their ensembles were used, based on hand-crafted features such as the set of (normalized) subterms and symbols in each formula.

## 4 Motivation for the use of Deep Learning

Strong premise selection requires models capable of reasoning over mathematical statements, here encoded as variable-length strings of first-order logic. In natural language processing, deep neural networks have proven useful in language modeling [30], text classification [9], sentence pair scoring [4], dependency parsing [2], sentiment analysis [35], conversation modeling [40], and simple question answering [36]. These results have demonstrated the ability of deep networks to extract useful representations from sequential inputs without hand-tuned feature engineering. Neural networks can also mimic some higher-level reasoning on simple algorithmic tasks [15, 42, 20]. Here, we extract learned representations of mathematical statements to assist in premise selection and proof.




Figure 2: Histograms of statement lengths, occurrences of each word, and statement dependencies in the Mizar corpus translated to first order logic. The wide length distribution poses difficulties for RNN models and batching, and the large number of rarely occurring words makes it important to take definitions of words into account.

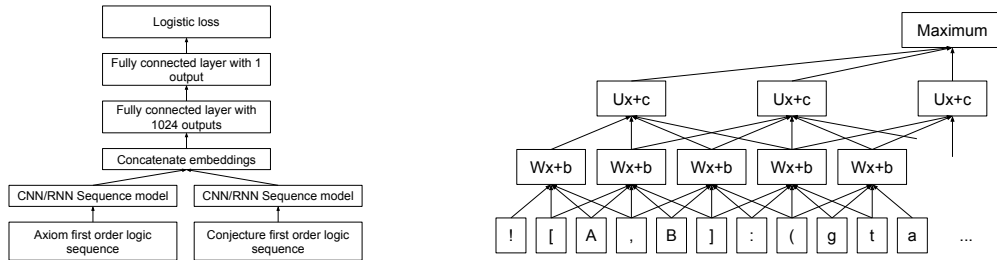


Figure 3: (left) Our network structure. The input sequences are either character-level (section 5.1) or word-level (section 5.2). We use separate models to embed conjecture and axiom, and a logistic layer to predict whether the axiom is useful for proving the conjecture. (right) A convolutional model.

The Mizar data set is also an interesting case study in neural network sequence tasks, as it differs from natural language problems in several ways. It is highly structured with a simple context free grammar – the interesting task occurs only after parsing. The distribution of lengths is wide, ranging from 5 to 84,299 characters with mean 304.5, and from 2 to 21,251 tokens with mean 107.4 (see Figure 2). Fully recurrent models would have to backpropagate through 100s to 1000s of characters or 100s of tokens to embed a whole statement. Finally, there are many rare words – 60.3% of the words occur fewer than 10 times – motivating the definition-aware embeddings in section 5.2.

## 5 Overview of our approach

The full premise selection task takes a conjecture and a set of axioms and chooses a subset of axioms to pass to the ATP. We simplify from subset selection to pairwise relevance by predicting the probability that a given axiom is useful for proving a given conjecture. This approach depends on a relatively sparse dependency graph: we ignore the issue of pruning redundant sets of axioms. Our general architecture is shown in Figure 3(left): the conjecture and axiom sequences are separately embedded into fixed length real vectors, then concatenated and passed to a third network with a few fully connected layers and logistic loss. During training time, the two embedding networks and the joined predictor path is treated as a single large neural network and trained jointly.

As discussed in section 3, we train our models on premise selection data generated by a combination of various methods, including k-nearest-neighbor search on hand-engineered similarity metrics. We start with a first stage of character-level models, and then build second and later stages of word-level models on top of the results of earlier stages.

### 5.1 Stage 1: Character-level models

We begin by avoiding special purpose engineering by treating formulas on the character-level. We use an 80 dimensional one-hot encoding of the character sequence, including 79 unique characters that occur in the translated Mizar corpus and a special character to denote the start and end of each statement. These sequences are passed to a network with weight sharing for variable length input. For the embedding computation, we have explored the following architectures:

1. Pure recurrent LSTM [19] and GRU [7] networks.
2. A pure multi-layer convolutional network with various numbers of convolutional layers (with strides) followed by a global temporal max-pooling reduction (see Figure 3(right)).
3. A recurrent-convolutional network, that uses convolutional layers to produce a shorter sequence which is processed by a LSTM.

The exact architectures used are specified in the experimental section.

Given the cost of sequence embedding, it is computationally prohibitive to compute a large number of (conjecture, axiom)-pairs. Fortunately, our architecture allows caching the embeddings for conjectures and axioms and evaluating only the shared portion of the network for a given pair, making it practical to consider all pairs during evaluation.

## 5.2 Stage 2: Word-level models

The character-level models are limited to word and structure similarity within the axiom or conjecture being embedded. However, many of the symbols occurring in a formula are defined by formulas earlier in the corpus, and we can use these definitions to improve model performance.

Since Mizar is based on first-order set theory, definitions of names can be either explicit or implicit. An explicit definition of  $x$  sets  $x = e$  for some expression  $e$ , while an implicit definition states a property of the defined object, such as defining a function  $f(x)$  by  $\forall x.f(f(x)) = g(x)$ . To avoid manually encoding the structure of implicit definitions, we embed the entire statement defining an identifier  $x$ , and then use these definition embeddings as word-level embeddings.

Ideally, we would train a single network that embeds statements by recursively expanding and embedding the definitions of words. Unfortunately, this recursion would dramatically increase the cost of training since the definition chains can be quite deep. For example, Mizar defines real numbers in terms of nonnegative reals, which are defined as Dedekind cuts of nonnegative rationals, which are defined as ratios of naturals, etc. As an inexpensive alternative, we reuse the axiom embeddings computed by a previously trained character-level model, mapping each defined identifier to the axiom embedding of its defining statement. Other tokens such as brackets and operators are mapped to fixed pseudorandom vectors of the same dimension. An extra binary feature distinguishes between defined tokens and random embeddings so that the network does not have to learn the difference.

Since we embed one token at a time ignoring grammatical structure, our approach does not require a parser, only a trivial lexer implemented in a few lines of Python.

Once we have word-level embeddings, we use the same architectures from stage 1 to reduce down to axiom and conjecture embeddings and then classify whether an (axiom, conjecture) pair is relevant.

We also tried iterating this approach of using definition embeddings as word embeddings multiple times, using the output embeddings of a trained definition-based model as the input word embeddings for another definition-based model. This extension did not result in measurable gains.

# 6 Experiments

## 6.1 Experimental Setup

For training and evaluation we use a subset of 32,524 out of 57,917 theorems that are known to be provable by an ATP given the right set of premises. We split off a random 10% of these (3,124 statements) for holdout testing and validation. This may lead to learning from future proofs: a proof  $P_j$  of theorem  $T_j$  written after theorem  $T_i$  may guide the premise selection for  $T_i$ . However, previous  $k$ -NN experiments show similar performance between a full 10-fold cross-validation and incremental evaluation as long as chronologically preceding formulas participate in proofs of only later theorems.

We additionally held out 400 statements from the 3,124 for monitoring training progress, as well as for model and checkpoint selection. Final evaluation was done on the remaining 2,724 conjectures. Note that we only held out conjectures, but we trained on all statements as axioms. This may lead to learning from future proofs: a proof  $P_j$  of theorem  $T_j$  written after theorem  $T_i$  may guide the premise selection for  $T_i$ . This is comparable to our  $k$ -NN baseline which is also trained on all statements as axioms, and where similar results were obtained from incremental evaluation. This practice is

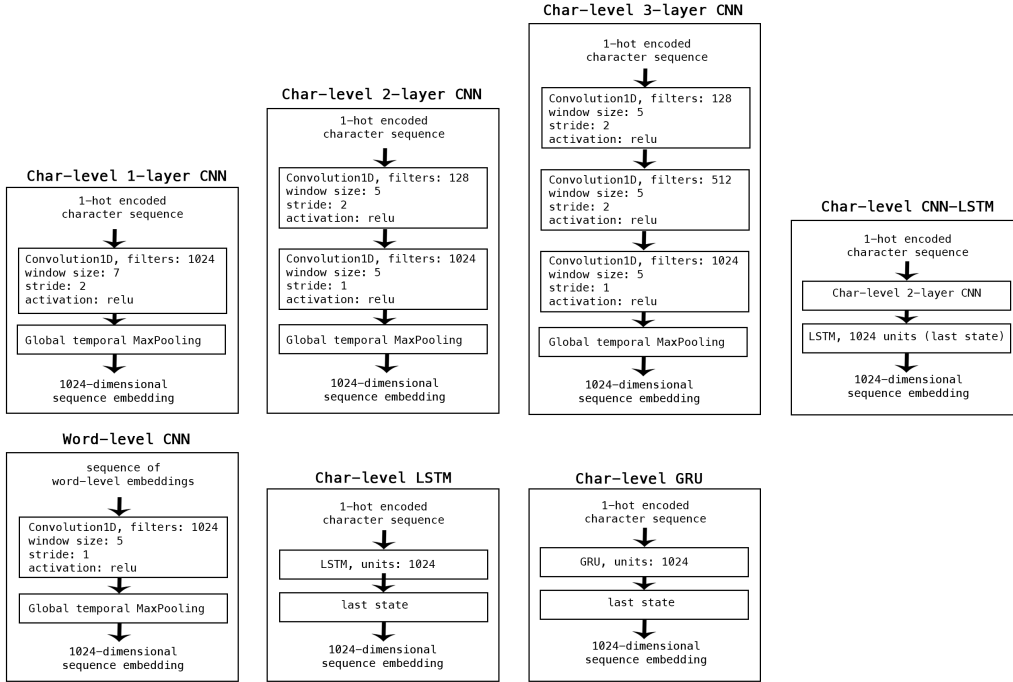


Figure 4: Specification of the different embedder networks.

justified by the fact that the hard problem is to find proofs for new conjectures where we can utilize all our knowledge about the usefulness of statements for other purposes.

## 6.2 Metrics

For each conjecture, our models output a ranking of the possible premises. Our primary metric is the number of conjectures proved from the top- $k$  premises, where  $k = 16, 32, \dots, 1024$ . This metric can accommodate alternative proofs but is computationally expensive. Therefore we additionally measure the ranking quality using the average maximum relative rank of the testing premise set. Formally, average max relative rank is

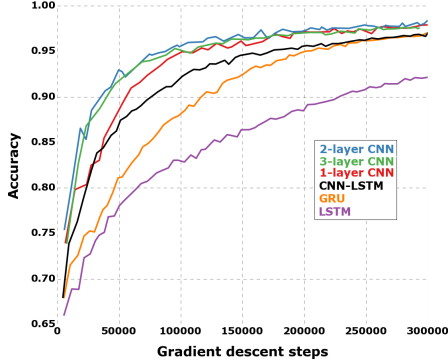
$$\text{aMRR} = \text{mean}_C \max_{P \in \mathcal{P}_{\text{test}}(C)} \frac{\text{rank}(P, \mathcal{P}_{\text{avail}}(C))}{|\mathcal{P}_{\text{avail}}(C)|}$$

where  $C$  ranges over conjectures,  $\mathcal{P}_{\text{avail}}(C)$  is the set of premises available to prove  $C$ ,  $\mathcal{P}_{\text{test}}(C)$  is the set of premises for conjecture  $C$  from the test set, and  $\text{rank}(P, \mathcal{P}_{\text{avail}}(C))$  is the rank of premise  $P$  among the set  $\mathcal{P}_{\text{avail}}(C)$  according to the model. The motivation for aMRR is that conjectures are easier to prove if all their dependencies occur early in the ranking.

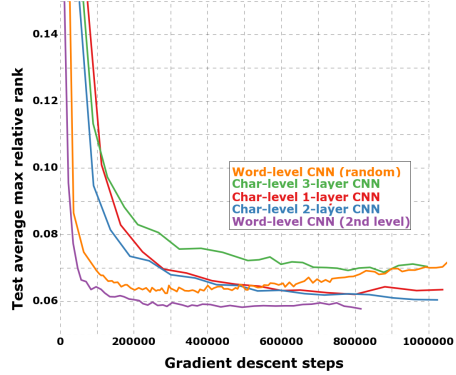
Since it is too expensive to rank all axioms for a conjecture during continuous evaluation, we approximate our objective. For our holdout set of 400 conjectures, we select all true dependencies  $\mathcal{P}_{\text{test}}(C)$  and 128 fixed random false dependencies from  $\mathcal{P}_{\text{avail}}(C) - \mathcal{P}_{\text{test}}(C)$  and compute the average max relative rank in this ordering. Note that aMRR is nonzero even if all true dependencies are ordered before false dependencies; the best possible value is 0.051.

## 6.3 Network Architectures

All our neural network models use the general architecture from Fig 3: a classifier on top of the concatenated embeddings of an axiom and a conjecture. The same classifier architecture was used for all models: a fully-connected neural network with one hidden layer of size 1024. For each model, the axiom and conjecture embedding networks have the same architecture but non-shared weights. The details of the embedding networks are shown in Fig 4.



(a) Training accuracy for different character-level models. Recurrent models seem underperform, while pure convolutional models yield the best results. For each architecture, we trained three models with different random initialization seeds. Only the best runs are shown on this graph; we did not see much variance between runs on the same architecture.



(b) Test average max relative rank for different models. The best is a word-level CNN using definition embeddings from a character-level 2-layer CNN. An identical word-embedding model with random starting embedding overfits after only 250,000 iterations and underperforms the best character-level model.

## 6.4 Network Training

The neural networks were trained using asynchronous distributed stochastic gradient descent using the Adam optimizer [26] with up to 20 parallel NVIDIA K-80 GPU workers per model. We used the TensorFlow framework [1] and the Keras library [6]. We used [11] to initialize convolutional and fully connected layers and Polyak averaging with 0.9999 decay to produce the final weights [32]. We experimented with gradient clipping at various thresholds, which helped stabilize training but yielded inferior results vs. non-clipped models. The character level models were trained with maximum sequence length 2048 characters, where the word-level (and definition embedding) based models were trained with a maximum sequence length of 500 words.

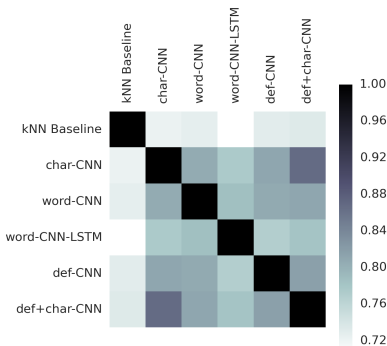
## 6.5 Experimental Results

Our best selection pipeline uses a stage-1 character-level convolutional neural network model to produce word-level embeddings for the second stage. The baseline uses distance-weighted  $k$ -NN [21, 23] with handcrafted semantic features [24]. For all conjectures in our holdout set, we consider each preceding statement (lemma, definition axiom) in the chronological ordering as a premise candidate. In the DeepMath case, premises were ordered by their logistic scores. E prover was applied to the top- $k$  of the premise-candidates for each of the cutoffs  $k \in (16, 32, \dots, 1024)$  until a proof is found or  $k = 1024$  fails. Table 1 reports the number of theorems proved with a cutoff value *at most* the  $k$  in the leftmost column. For E prover, we use a soft time limit of 90 seconds, a hard time limit of 120 seconds, a memory limit of 4 GB, and a processed clauses limit of 500,000. These settings are generous: even dramatically increasing the limits proves at most 10 extra theorems.

Our most successful models employ simple convolutional networks followed by max pooling (as opposed to recurrent networks like LSTM/GRU), and the two stage definition-based **def-CNN** outperforms the naïve **word-CNN** word embedding significantly. In the latter the word embeddings were learned in a single pass; in the former they are fixed from the stage-1 character-level model. For each architecture (cf. Figure 4) two convolutional layers perform best. Although our models differ significantly from each other, they differ even more from the  $k$ -NN baseline based on hand-crafted features. The right column of Table 1 shows the result if we average the prediction score of the stage-1 model with that of the definition based stage-2 model. We also experimented with shorter character-based RNN models using shorter sequences: these lagged behind our long-sequence (at max 2048) character models but performed significantly better than those trained on longer sequences. This suggest that these RNNs could be improved by more sophisticated optimization techniques such as curriculum learning. Finally, our definition-based CNN model is capable of proving 415 theorems improved in E by previous methods, including those using dependencies from human Mizar proofs.

Cutoff	$k$ -NN Baseline	char-CNN	word-CNN	word-CNN-LSTM	def-CNN	def+char-CNN
16	<b>503</b>	390	364	256	375	473
32	<b>810</b>	660	638	502	669	778
64	<b>1153</b>	1013	932	780	1029	1120
128	1403	1336	1254	1080	1339	<b>1454</b>
256	1541	1547	1454	1297	1549	<b>1617</b>
512	1626	1627	1558	1447	1644	<b>1696</b>
1024	1627	1670	1608	1526	1682	<b>1726</b>
% proved	59.34	60.90	58.64	55.65	61.34	<b>62.95</b>

Table 1: Results for ATP premise selection experiments. Each entry is the number of theorems automatically proved with E prover using that particular model to sort the premises, out of a total of 2,742 statements. Taking a union of all generated proofs we prove 2,108 statements (77%); taking the union of just the neural network models proves 2,038 statements (74%).



(c) Jaccard similarities between proved sets of conjectures across models. Each of the neural network models are more like each other than they are like the  $k$ -NN baseline.

Model	Test max average relative rank
char-CNN	0.060
word-CNN	0.063
word-CNN-LSTM	0.068
def-CNN	<b>0.059</b>

(d) Test results obtained by our different character-level and word-level models. Lower values are better. The evaluation was performed continuously during training with 400 theorems, using all positive premises from the training set and 128 randomly selected negatives. With this setup, the optimal max average relative rank with perfect predictions is 0.051.

## 7 Conclusions

In this work we provide evidence that even simple neural models can compete with hand-engineered features for premise selection, helping to find many new proofs. This translates to real gains in automatic theorem proving. Despite these encouraging results, our models are relatively shallow networks with inherent limitations to representational power and are incapable of capturing high level properties of mathematical statements. We believe theorem proving is a challenging and important domain for deep learning methods, and that more sophisticated optimization techniques and training methodologies will prove more useful than in less structured domains.

## References

- [1] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.
- [2] D. Andor, C. Alberti, D. Weiss, A. Severyn, A. Presta, K. Ganchev, S. Petrov, and M. Collins. Globally normalized transition-based neural networks. *arXiv preprint arXiv:1603.06042*, 2016.
- [3] G. Bancerek and P. Rudnicki. A Compendium of Continuous Lattices in MIZAR. *J. Autom. Reasoning*, 29(3-4):189–224, 2002.
- [4] P. Baudiš, J. Pichl, T. Vyskočil, and J. Šedivý. Sentence pair scoring: Towards unified framework for text comprehension. *arXiv preprint arXiv:1603.06127*, 2016.



- [5] J. C. Blanchette, C. Kaliszyk, L. C. Paulson, and J. Urban. Hammering towards QED. *J. Formalized Reasoning*, 9(1):101–148, 2016.
- [6] F. Chollet. Keras. <https://github.com/fchollet/keras>, 2015.
- [7] J. Chung, C. Gulcehre, K. Cho, and Y. Bengio. Gated feedback recurrent neural networks. *arXiv preprint arXiv:1502.02367*, 2015.
- [8] The Coq Proof Assistant. <http://coq.inria.fr>.
- [9] A. M. Dai and Q. V. Le. Semi-supervised sequence learning. In *Advances in Neural Information Processing Systems*, pages 3061–3069, 2015.
- [10] N. de Bruijn. The mathematical language AUTOMATH, its usage, and some of its extensions. In M. Laudet, editor, *Proceedings of the Symposium on Automatic Demonstration*, pages 29–61, Versailles, France, Dec. 1968. Springer-Verlag LNM 125.
- [11] X. Glorot and Y. Bengio. Understanding the difficulty of training deep feedforward neural networks. In *International conference on artificial intelligence and statistics*, pages 249–256, 2010.
- [12] G. Gonthier. The four colour theorem: Engineering of a formal proof. In D. Kapur, editor, *Computer Mathematics, 8th Asian Symposium, ASCM 2007, Singapore, December 15-17, 2007. Revised and Invited Papers*, volume 5081 of *Lecture Notes in Computer Science*, page 333. Springer, 2007.
- [13] G. Gonthier, A. Asperti, J. Avigad, Y. Bertot, C. Cohen, F. Garillot, S. L. Roux, A. Mahboubi, R. O’Connor, S. O. Biha, I. Pasca, L. Rideau, A. Solovyev, E. Tassi, and L. Théry. A machine-checked proof of the Odd Order Theorem. In S. Blazy, C. Paulin-Mohring, and D. Pichardie, editors, *ITP*, volume 7998 of *LNCS*, pages 163–179. Springer, 2013.
- [14] A. Grabowski, A. Kornilowicz, and A. Naumowicz. Mizar in a nutshell. *J. Formalized Reasoning*, 3(2):153–245, 2010.
- [15] A. Graves, G. Wayne, and I. Danihelka. Neural Turing machines. *arXiv preprint arXiv:1410.5401*, 2014.
- [16] T. C. Hales, M. Adams, G. Bauer, D. T. Dang, J. Harrison, T. L. Hoang, C. Kaliszyk, V. Magron, S. McLaughlin, T. T. Nguyen, T. Q. Nguyen, T. Nipkow, S. Obua, J. Pleso, J. Rute, A. Solovyev, A. H. T. Ta, T. N. Tran, D. T. Trieu, J. Urban, K. K. Vu, and R. Zumkeller. A formal proof of the Kepler conjecture. *CoRR*, abs/1501.02155, 2015.
- [17] J. Harrison. HOL Light: A tutorial introduction. In M. K. Srivas and A. J. Camilleri, editors, *FMCAD*, volume 1166 of *LNCS*, pages 265–269. Springer, 1996.
- [18] J. Harrison, J. Urban, and F. Wiedijk. History of interactive theorem proving. In J. H. Siekmann, editor, *Computational Logic*, volume 9 of *Handbook of the History of Logic*, pages 135 – 214. North-Holland, 2014.
- [19] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [20] Ł. Kaiser and I. Sutskever. Neural gpu learn algorithms. *arXiv preprint arXiv:1511.08228*, 2015.
- [21] C. Kaliszyk and J. Urban. Stronger automation for Flyspeck by feature weighting and strategy evolution. In J. C. Blanchette and J. Urban, editors, *PxTP 2013*, volume 14 of *EPiC Series*, pages 87–95. EasyChair, 2013.
- [22] C. Kaliszyk and J. Urban. Learning-assisted automated reasoning with Flyspeck. *J. Autom. Reasoning*, 53(2):173–213, 2014.
- [23] C. Kaliszyk and J. Urban. MizAR 40 for Mizar 40. *J. Autom. Reasoning*, 55(3):245–256, 2015.
- [24] C. Kaliszyk, J. Urban, and J. Vyskocil. Efficient semantic features for automated reasoning over large theories. In Q. Yang and M. Wooldridge, editors, *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015*, pages 3084–3090. AAAI Press, 2015.
- [25] M. Kaufmann and J. S. Moore. An ACL2 tutorial. In Mohamed et al. [31], pages 17–21.
- [26] D. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [27] G. Klein, J. Andronick, K. Elphinstone, G. Heiser, D. Cock, P. Derrin, D. Elkaduwe, K. Engelhardt, R. Kolanski, M. Norrish, T. Sewell, H. Tuch, and S. Winwood. seL4: formal verification of an operating-system kernel. *Commun. ACM*, 53(6):107–115, 2010.
- [28] D. Kuehlwein and J. Urban. Learning from multiple proofs: First experiments. In P. Fontaine, R. A. Schmidt, and S. Schulz, editors, *PAAR-2012*, volume 21 of *EPiC Series*, pages 82–94. EasyChair, 2013.
- [29] X. Leroy. Formal verification of a realistic compiler. *Commun. ACM*, 52(7):107–115, 2009.
- [30] T. Mikolov, M. Karafiát, L. Burget, J. Cernocký, and S. Khudanpur. Recurrent neural network based language model. In *INTERSPEECH*, volume 2, page 3, 2010.

- [31] O. A. Mohamed, C. A. Muñoz, and S. Tahar, editors. *Theorem Proving in Higher Order Logics, 21st International Conference, TPHOLs 2008, Montreal, Canada, August 18-21, 2008. Proceedings*, volume 5170 of *LNCS*. Springer, 2008.
- [32] B. T. Polyak and A. B. Juditsky. Acceleration of stochastic approximation by averaging. *SIAM Journal on Control and Optimization*, 30(4):838–855, 1992.
- [33] J. A. Robinson and A. Voronkov, editors. *Handbook of Automated Reasoning (in 2 volumes)*. Elsevier and MIT Press, 2001.
- [34] S. Schulz. E - A Brainiac Theorem Prover. *AI Commun.*, 15(2-3):111–126, 2002.
- [35] R. Socher, J. Pennington, E. H. Huang, A. Y. Ng, and C. D. Manning. Semi-supervised recursive autoencoders for predicting sentiment distributions. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 151–161. Association for Computational Linguistics, 2011.
- [36] S. Sukhbaatar, J. Weston, R. Fergus, et al. End-to-end memory networks. In *Advances in Neural Information Processing Systems*, pages 2431–2439, 2015.
- [37] 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.
- [38] J. Urban. MPTP 0.2: Design, implementation, and initial experiments. *J. Autom. Reasoning*, 37(1-2):21–43, 2006.
- [39] J. Urban and J. Vyskočil. Theorem proving in large formal mathematics as an emerging AI field. In M. P. Bonacina and M. E. Stickel, editors, *Automated Reasoning and Mathematics: Essays in Memory of William McCune*, volume 7788 of *LNAI*, pages 240–257. Springer, 2013.
- [40] O. Vinyals and Q. Le. A neural conversational model. *arXiv preprint arXiv:1506.05869*, 2015.
- [41] M. Wenzel, L. C. Paulson, and T. Nipkow. The Isabelle framework. In Mohamed et al. [31], pages 33–38.
- [42] W. Zaremba and I. Sutskever. Learning to execute. *arXiv preprint arXiv:1410.4615*, 2014.