
HOList: An Environment for Machine Learning of Higher-Order Theorem Proving

Kshitij Bansal^{*1} Sarah Loos^{*1} Markus Rabe^{*1} Christian Szegedy^{*1} Stewart Wilcox^{*1}

Abstract

We present an environment, benchmark, and deep learning driven automated theorem prover for higher-order logic. Higher-order interactive theorem provers enable the formalization of arbitrary mathematical theories and thereby present an interesting, open-ended challenge for deep learning. We provide an open-source framework based on the HOL Light theorem prover that can be used as a reinforcement learning environment. HOL Light comes with a broad coverage of basic mathematical theorems on calculus and the formal proof of the Kepler conjecture, from which we derive a challenging benchmark for automated reasoning. We also present a deep reinforcement learning driven automated theorem prover, DeepHOL, with strong initial results on this benchmark.

1. Introduction

Formalization of mathematics and the automated creation of new mathematical content is at the frontier of current AI techniques. Given the fundamental nature of mathematics and its importance for most scientific disciplines, the capability for high level formal mathematical reasoning is both an important practical task as well as one of the most challenging case studies in AI. However, traditional formal computer mathematics has been a fragmented domain, exploring various approaches for different logical foundations. This has led to a large number of incompatible theorem proving systems, which added extra challenges for AI researchers trying to push the limits of formal reasoning using machine learning.

Well-defined, large-scale benchmarks were instrumental for unifying disparate efforts in machine learning research: Lib-

^{*}Equal contribution ¹Google Research, Mountain View, California, USA. Correspondence to: Sarah Loos <smloos@google.com>, Christian Szegedy <szegedy@google.com>.

riSpeech (Panayotov et al., 2015) for speech recognition, the Netflix prize (Bennett et al., 2007) for recommendation, ImageNet (Deng et al., 2009) for object recognition, MSCOCO (Lin et al., 2014) for object detection and segmentation, WMT (Bojar et al., 2014) for machine translation, and SQuAD (Rajpurkar et al., 2016) for question answering - just to name a couple of examples. Benchmarks have fostered collaboration and competition and provide a means to measure progress, contributing significantly to accelerated progress and reproducible science.

This paper provides a benchmark and reinforcement learning environment for theorem proving. The long-term goal is to enable the automatic formalization of large theories, and hence we want to start with a theorem proving system that has a track-record of large-scale formalization efforts and includes a large corpus of foundational mathematics for benchmarking and learning. Our choice fell on HOL Light, the interactive theorem prover (ITP) in which the proof of the Kepler conjecture (Hales et al., 2017) has been formalized. The formalization of the proof of the Kepler conjecture has been a huge effort, taking over 20 person-years to complete, and required formalizing a significant part of arithmetic, linear algebra, and multivariate analysis. The resulting benchmark consists of 2199 definitions and 29462 theorems and lemmata, which capture a variety of interesting mathematics and should be a practical seed for new (auto-)formalization efforts.

To demonstrate the feasibility of the proposed learning task, we present an automated theorem prover powered by deep learning, called DeepHOL. Based on a simple solver architecture, DeepHOL learns to prove theorems based on imitating human proofs and improves itself using reinforcement learning. Given a proof goal (represented as a string) DeepHOL learns to predict the tactic (and its arguments) that leads to a successful proof. Thereby, DeepHOL achieves theorem proving capabilities that are comparable to much more complicated state-of-the-art automated theorem proving systems. In our open-source release, available at <http://deephhol.org>, we expose the APIs of our modular theorem prover. This simplifies the development of new provers significantly and allows researchers to focus on the machine learning aspects.

The contributions of our work are the following:

- An instrumented, pre-packaged version of HOL Light that can be used as a reinforcement learning environment for theorem proving using our well-defined, stable Python API. Our solution comes with optimized startup capabilities for proof search, while allowing replay and strict verification of the produced proofs.
- Proof export and import capabilities that allow for managing large theories programmatically from the Python interface.
- A full-fledged, competitive automated neural theorem proving system that can automatize theorem proving in higher-order logic at tactic level directly.
- A large scale reinforcement learning system that was used for training our prover.
- Comparison of neural model architectures for theorem proving purposes.
- Well-defined benchmarks on our HOL Light based environment to enable research and measuring progress of AI driven theorem proving in large theories.

This paper is organized as follows. We discuss related work in Section 2 before we describe our theorem proving environment in Section 3. In Section 4 we present the organization of the benchmark. The DeepHOL automated theorem prover is described in Section 5 and we discuss first experimental results for it in Section 6. Then we conclude in Section 7.

2. Related Work

The earliest work of applying machine learning on reasoning in large theories is (Urban et al., 2008). The most similar works to ours are TacticToe (Gauthier et al., 2017) and GamePad (Huang et al., 2018). TacticToe is the first published result on machine learning tackling higher-order theorem proving at a relatively large scale at tactic level (Gauthier et al., 2017). Although TacticToe is a great success that came with significant improvements over previous automated theorem proving systems, they do not propose an easy to use benchmark or environment for machine learning researchers. TacticToe does not employ deep learning nor reinforcement learning. They rely on the HOL4 (Slind & Norrish, 2008) system that has a significantly less theorems with more complex human proof scripts with a larger number of more elementary tactics.

GamePad has very similar objectives to ours (Huang et al., 2018). They also provide an easy-to-use Python API for an interactive theorem prover, and they present test and training sets. They chose to base their system on Coq (Bertot & Castéran, 2013), an interactive theorem prover based on the calculus of inductive constructions. While enabling automatic code extraction, it comes with a much smaller

coverage of fundamental mathematics. Even including the formalization of the Feit-Thompson theorem, their benchmark comprises only 1602 theorems and lemmas, while ours features 29462 theorems and lemmas. Besides presenting a much larger data set, we also demonstrate the feasibility of achieving state-of-the-art prover performance based on our data and environment by presenting a deep learning based theorem prover. We also report the results as theorem proving performance instead of proxy metrics.

Other interactive theorem provers we could have based a learning environment on include Mizar (Mizar), Isabelle (Wenzel et al., 2008), HOL4 (Slind & Norrish, 2008), and Lean (de Moura et al., 2015). The Mizar mathematical library is probably the most comprehensive formalization effort, but its declarative style makes it hard to employ proof search, and its source code is not freely available. Like Coq and HOL Light, also Isabelle (Wenzel et al., 2008) was used for major formalization efforts, such as the formalization of the seL4 microkernel (Klein et al., 2009). We are not aware of a comprehensive coverage of fundamental mathematics in Isabelle, HOL4, or Lean.

In closely related work, Kaliszyk & Urban (2014) translate from HOL Light and Flyspeck to automated theorem provers and SMT solvers, for which they learn a premise selector. In contrast to our work, they use neither deep learning nor reinforcement learning. Similar methods for premise selection on the HOL Light corpora were proposed in (Kaliszyk & Urban, 2012).

The first use of deep neural networks for large scale theorem proving was proposed in (Alemi et al., 2016). They have used convolutional networks for premise selection in large theories, particularly on Mizar mathematical library (Mizar). Those methods were used as a pre-selection for applying the first order logic automated theorem prover E (Schulz, 2002). We have reused several ideas from that paper, including some aspects of our neural network architecture and the hard negative mining methodology.

Whalen (2016) proposed a purely deep reinforcement learning based solution for theorem proving for the Metamath prover (Megill, 1997). This work was moderately successful, finding mostly proofs for very simple theorems, especially in propositional logic. On the other hand, Metamath is not considered to be a serious contender for large scale mathematical formalization work.

Loos et al. (2017) proposed deep neural networks to augment theorem prover E (Schulz, 2002) to rank given clauses during proof search. Here, we propose a neural prover written from scratch, relying solely on a small set of preexisting tactics and neural networks for all high level decisions.

Kaliszyk et al. (2017) proposed a machine learning benchmark for higher-order logic reasoning based on the HOL

Light corpus. It features a few static datasets and it remains unclear how performance of machine learning models on this dataset relates to real world prover performance. (Kaliszyk et al., 2018) demonstrated the viability of reinforcement learning with XGBoost and LIBLINEAR (Fan et al., 2008) on hand engineered features in first order logic context using leanCoP (Otten & Bibel, 2003) on Mizar mathematical library (Mizar).

Earlier works on employing (non-deep) machine learning for theorem proving in general and for reasoning in large theories include (Schulz, 2000; Duncan et al., 2004; Urban et al., 2011; Kühlwein et al., 2012; Kaliszyk & Urban, 2013; Kühlwein et al., 2013; Alama et al., 2014; Bridge et al., 2014; Kaliszyk et al., 2014b;a; Färber & Kaliszyk, 2015; Kaliszyk & Urban, 2015c; Kaliszyk et al., 2015; Kaliszyk & Urban, 2015a;b; Gauthier & Kaliszyk, 2015; Blanchette et al., 2016). Recently, Wang et al. (2017) proposed a premise selection method utilizing deep graph embeddings.

3. Architecture of the Environment

Here we describe the architecture of the evaluation and training environment. The goal of the environment is to enable artificial agents to interact with the HOL Light interactive theorem prover (ITP) in a replicable manner.

3.1. ITP Terminology

In order to describe our changes to HOL Light, it is helpful to establish some common terminology. To prove a theorem in an ITP, the human user starts with entering the theorem’s statement as the *goal* of a new proof. The ITP provides a small number of *tactics* to manipulate the goal. Tactics may have *tactic arguments*, which can be a previously proven theorem or a list of previously proven theorems. (There are also tactics that take terms as arguments, but we do not support them currently.) Applying a tactic to a goal can lead to a failure, when not all conditions are met, or is successful and produces a list of subgoals. The goal is only proven successfully, if all its subgoals are proven. In particular, if the goal is proven if the tactic application produces an empty list of subgoals. We refer to tactic applications sometimes also as *proof steps*.

We can think of proofs as trees, where goals are nodes and tactic applications are (hyper-)edges to other goals. In a successful proof, all leaves are goals with a tactic application that produced an empty list of subgoals.

3.2. Instrumentation to HOL Light

In order to create a stable, well-defined environment, we fix a particular version of HOL Light with a pre-selected subset of tactics and a fixed library of basic theorems, which are proved in one well-defined order. This is the ITP part

of the environment which is written in OCaml with a few additional C++ functions. Since it is non-trivial to find and build the exact correct set of libraries for this environment, we provide a prepackaged docker image. It can be used as a reliable black box for proof search and as reinforcement learning environment, communicated with using a simple API. We have also open sourced all the changes to the HOL Light system so that new modifications and forks are possible by third parties.

The prepackaged version we provide has the following additional instrumentation, which we describe below in detail:

- Logging of human-written proofs shipped with HOL Light.
- A new API to interact with HOL Light for proof search.
- Fast startup for distributed proof search.
- A proof checker to remove the need to trust search algorithms.

3.3. Proof Logging for Human Proofs

We want to utilize the existing human proofs for both training and evaluation. To that effect, we have instrumented the `prove` method in HOL Light with extra logging code. If HOL Light is executed in proof-dump mode, each invocation of the `prove` function dumps the proven theorems and their proofs into files. These *proof logs* can then be converted to training examples (see Section 4.1).

3.4. Proof Assistant API

The API provides two functions: (1) to apply tactics to goals and (2) to register theorems for future use in tactic applications. Tactic applications are completely stateless and contain the goal, the tactic to be applied, and the tactic arguments. The proof assistant (i.e. HOL Light in our implementation) returns the outcome of the tactic application, including the list of subgoals for successful applications. The stateless tactic application interface frees us from the strict order on subgoals that HOL Light enforces in the human interface, and allows us to easily implement more advanced proof search strategies.

The tactic arguments can consist of a list of theorems. Implemented naively, this list could make the tactic application request very large and could slow down the prover. In the argument list of tactics we therefore allow theorems to be referenced by a fingerprint number. The second API function allows us to register theorems such that HOL Light can resolve the fingerprints to theorems. The registration of theorems is hence stateful, in contrast to tactic applications.

3.5. Fast Startup

Starting HOL Light and loading all the potentially needed libraries can take a long time - we measured it at up to 20 minutes. This would be inhibitive long for proof search, especially in a distributed setting with thousands of workers and the startup time has to be paid for every worker. The Proof Assistant API allows us to load only a minimal core of HOL Light and register the remaining theorems from the libraries using the API. This brings the startup time of our HOL Light to mere seconds.

3.6. Proof Checking

Any bug in the implementation of a theorem prover could make its reasoning unsound, rendering the whole formalization effort futile. For that reason, HOL Light is designed around a small trusted core of about 400 lines of OCaml code that builds proofs from few very basic rules. OCaml’s type system guarantees that a theorem object can only be constructed by this trusted core, and the rest of the HOL Light system can be seen as mere convenience features.

Our API allows researchers to implement proof search algorithms outside of OCaml. The correctness of any proof found through the API thus relies on the correctness of our API implementation and the proof search itself. We thus implemented a proof checker that avoids the need for trusting the proof search and even the API. The proof checker compiles proofs into OCaml code that can be loaded in HOL Light, where they have to pass through the trusted core.

4. Benchmark

We present three different corpora: “core”, “complex”, and “fleyspeck.” The core corpus contains the basic theorems that needed to define the tactics and the complex corpus consists of theorems of complex calculus. While proofs of core theorems are useful for training, we omit them in validation, since some tactics assume those theorems. Flayspeck contains most of the lemmas and theorems of the Kepler conjecture. Together these three corpora encompass almost 30k theorems and proofs (see Table 1).

We propose two tasks that can be measured on these benchmarks:

- Predict the tactic and tactic arguments that were employed in the human proof.
- Prove each of the theorems in the corpora while utilizing only those theorems as tactic arguments that also humans had available. For that purpose, we provide all theorems in the three corpora in one unified list, in the order they were proven by humans.

	Definitions	Theorems	Proof states
core	239	2320	23512
complex	398	16623	509621
fleyspeck	1563	10519	538540
all	2200	29462	1071673

Table 1. The three corpora of the benchmark.

4.1. Training Examples

Our training examples consist of a *goal*, a *tactic*, an *arglist*, and a *negarglist*. The *goal* is a provable statement, i.e. it is either a theorem from one of the corpora or a subgoal of a successful proof. The *tactic* is the ID of one of a preselected small set of tactics (currently consisting of 41 tactics) that led to a successful proof. The *arglist* is the list of theorems that were passed to a tactic application as arguments. Additionally, there is a special argument signifying that the argument list was empty.

The *negarglist* is an optional list of non-arguments that is not actually necessary for any proof. *negarglist* consists of high-scoring theorems that were not actually needed as arguments. They are collected during proof search in our reinforcement learning pipeline, and the list is empty for all the examples generated from the human proof logs.

4.2. Splits

Before training and evaluation, we have split the top level theorems into three subsets: training, validation and test set in a 60:20:20 ratio. Since the goals occurring in the proof of a theorem are likely correlated with the theorem itself, we assign them the same split as the theorem. The validation set can be used for continuous monitoring for proxy metrics of the model during training. The validation set is also occasionally used to measure the end-to-end prover performance of the models during training. The test set, on the other hand, must only be used extremely rarely for final assessment of a few models before publishing a paper alongside their validation set performance.

4.3. Representation of Expressions

All expressions are presented as S-expressions that have only few types of non-leaf nodes: function applications, abstractions (i.e. lambda functions), variables, constants, and function types. All other information, such as variable names, constant names, and type names, is given as leaf-nodes. For example, the expression $f(x)$ for a function $f : \mathbb{R} \rightarrow \mathbb{R}$ looks as follows: `(a (v (fun (real) real) f) (v real x))`. These S-expressions have a unique correspondence to terms in HOL Light and are easy to parse into a tree. However, our current models only observe the string version of these expressions. Expressions

are quite long in this representation: The average number of tokens in the goals is around 500, and the median is around 300.

For many operations, HOL Light automatically invents new generic types (e.g. `?345882`) and generic variables (e.g. `GEN%PVAR%9675`) on the fly. This leads to thousands of types and variables that often occur in only one (or a few) expressions, and hence would hardly get meaningful embeddings in typical deep learning approaches. Further, tokens that are shared only between few expressions bear the risk of unintentionally giving away information about the relations between these statements. We therefore decided to *normalize* the data sets by mapping generic types and generic variables to a much smaller set of names while maintaining the semantics of all expressions. After normalization, the number of distinct tokens is 1254.

5. DeepHOL Prover architecture

In this section, we describe the high-level architecture of our reference neural prover. The intelligence is fully learned without any hand-crafted features, and with very simple data preprocessing. In particular, we have not implemented any tweaks for the particular logic or interactive theorem prover (ITP). All the engineering went into the neural network architecture, which is very generic, and into maintaining the proof search graph without any special regard for the particular ITP system. In other words, DeepHOL currently uses HOL Light and its logic (HOL), but is not specialized to it. We believe that our solution would also work with other goal-tactic based prover like Coq (Bertot & Castéran, 2013), HOL4 (Slind & Norrish, 2008), or Lean (de Moura et al., 2015). Here we describe the details of our reference prover solution in detail.

5.1. Action Generator

The most crucial part of our prover is the action generator that produces a list of tactic applications for a given goal. We have split this into two subtasks:

- To rank the tactics, and
- to create an argument list for each of the tactics (comprised of a list of theorems).

As noted earlier, DeepHOL is currently not using tactics that take arbitrary terms (formulas) as parameters.

For both subtasks, the action generator employs a neural network, which we describe in Section 5.4. The ranking of tactic applications it produces is used in the proof search (Section 5.3) to expand the proof search graph (Section 5.2).

5.2. Proof Search Graph

The proof search graph is our data structure that captures the state of the proof search, and allows us to detect when a proof for the original goal is available. The nodes of the proof search graph are the goals that we have seen in the proof search, including the original goal statement that we want to prove. Each goal can have multiple alternative tactic applications, each of which might result in multiple subgoals. That is, tactic applications are labelled hyperedges in the proof search graph.

The proof search graph provides some features that allow us to prune some subtrees of the search early: First, whenever a tactic application closes a subgoal, this information is traced back to the parent subgoals and each alternate tactic application (and its whole sub-branch) is marked as closed and is discarded from the queue to be processed. If during this recursive process the root node is reached, then the proof is closed and the proof process stops. Second, when all tactic applications for a goal fail we mark that goal as unsuccessful. Similar to tracing closed goals, the proof search graph automatically traces the siblings of unsuccessful subgoals that become superfluous, and mark them unsuccessful as well. Third, when tactic applications produce identical subgoals, we let them point to the same node in the proof search graph. We refer to this as *subgoal sharing*, and once a subgoal is newly shared, previously stored information about subgoals being closed or ignored is propagated through the search graph.

5.3. Proof Search

Our proof search is a simple breadth first search. In each iteration, it expands all leaf nodes (i.e. goals that have not been expanded yet). To expand a goal, it calls the action generator to generate a list of tactic applications, and applies them in order. It stops applying tactics to a goal, when it reaches a maximum number of unsuccessful tactic applications or a minimum number of successful tactic applications. Whenever a complete proof is found for the top level goal, the proof search is stopped and the whole proof search graph is serialized and stored as the result. Also, the proof search finishes if the search graph reaches a prescribed limit on the number of subgoals or the proof search times out.

Note that subgoal sharing, as explained in Section 5.2, is crucial for our proof search: Without subgoal sharing the search process could end up oscillating between two formulas by rewriting the same subterm back and forth using the same equation.

5.4. Neural Architectures

For the generation and ranking of actions in the action generator, we use a deep, two-tower neural network depicted in

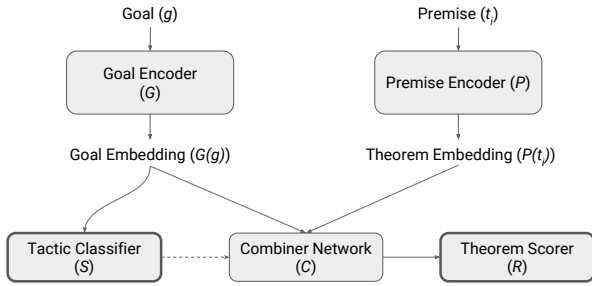


Figure 1. Two-tower neural architecture for ranking actions.

Figure 1. The predictions of the neural network are based on a single goal, represented as an S-expression of the HOL Light term (i.e. a string). (In HOL Light, each goal consists of a list of hypotheses and a conclusion, and we currently drop the hypotheses before we feed a goal to the neural network.)

The neural network has two separate prediction heads S and R . The goal tower G computes an embedding $G(g)$ of the current goal g and infers a scoring vector $S(G(g))$ for the fixed set of tactics where the tactic classifier S is a linear layer producing logits of a softmax classifier. The premise tower P computes a fixed size embedding $P(t_i)$ of all possible tactic arguments t_i in the scope of the goal to be proved. The ranking of the premises is performed by a combiner network C that takes the concatenation of the goal embedding, the premise embedding and possibly that of the tactic T_j to be applied: $r(t_i) = C(G(g), P(t_i), T_j)$, where $r(t_i)$ is the score of theorem t_i for its being a useful tactic argument for transforming the current goal g towards a closed proof. We have also tried the unconditioned setup, in which the ranking of the tactic arguments is independent of that of the tactic to be applied, that is $r(t_i) = C(G(g), P(t_i))$. In essence, we propose a hybrid architecture that both predicts the correct tactic to be applied, as well as rank the premise parameters required for meaningful application of the tactics.

5.5. Supervised Learning

We started training DeepHOL in a supervised learning setup, for which we use the human proof logs. We have split our data into test, train, and validation set on the theorem level, as described in Section 4. We always report both validation and test set performance for the final result to verify that we did not over-fit on the validation set. Continuous measurements and ablation analyses are reported only on the validation or training set.

5.6. Reinforcement Learning Loop

In the reinforcement learning loop, we have both a trainer and multiple provers running continuously. The training

is (optionally) seeded with training examples from existing (human/generated) proof logs. Then, we run the neural prover in *rounds*, each round trying to prove a random sample of theorems in the training set. Training examples extracted from successful proof logs of each round of our neural prover are mixed in continuously. Training examples of more recent rounds (*fresh* examples) can be weighed differently from older rounds (*historical* examples) during the training process.

To summarize, our loop works with the following four kinds of training example pools:

1. (optional) Human training examples as seed.
2. (optional) Inherited computer generated examples as seed: in addition to using human training examples as seed, examples generated during any previous experiments with our prover can also be used as seed. In our current experiments, we used examples that were generated by a prover that was run on the whole training set utilizing a model that was trained in purely supervised manner.
3. Fresh generated loop examples (examples that were produced in the last k rounds, where k is a user-settable parameter).
4. Historical training loop examples (examples that were produced in all but the last k rounds).

During training, batches are filled with examples from each pool according to a prescribed split ratio. This means that the ratio of different kinds of examples the model is trained on does not shift as more examples are generated by the loop. Most importantly, it also ensures that examples from freshly constructed new proofs show up quickly and deterministically during the training process. Note that although we can make use of human and inherited proof traces, the system can learn without any supervision or initial seed data. However, preliminary experiments have shown that, in its current form, it learns inferior models compared to those that were seeded with human proofs.

5.6.1. PROOF PRUNING

The argument lists of tactic applications in the reinforcement learning loop are quite long, and they contain superfluous elements. In order to obtain high quality training data for tactic argument prediction, we *prune* the parameter list before using them for training. For all tactics that take a list of theorems as an argument, our current implementation generates a list of fixed length. For successful tactic applications, we then iterate over the arguments in reverse score order and greedily omit those arguments that do not change the outcome of the tactic application. While a non-greedy approach might yield even shorter argument lists, it would also take longer to compute. In practice, our approach produces short argument lists with minimal effort. Removed parameters

Description	Proof success
ASM_MESON_TAC	6.1%
ASM_MESON_TAC + argument selection	9.2%
WaveNet	31.72%
Deeper WaveNet	32.65%
Wider WaveNet	27.60%
Loop	36.3%
Trained on loop output	36.8%
Loop tactic dependent	38.9%

Table 2. Percentage of theorems closed using various models on the validation set of the *complex* corpus comprising of 3225 theorems. First two lines are trivial baselines that call HOL Light’s built-in first order theorem prover with and without utilizing our argument selection model. The middle section shows results of models trained in a supervised scenario on human proofs. The last four lines report results using our reinforcement learning loops.

are stored as “hard negatives” and utilized during training.

6. Results and Comparisons

In this section, we first present several baseline results based on imitation (i.e. fully supervised) learning. Then we come to our reinforcement learning results using a WaveNet (Van Den Oord et al., 2016) based encoder architecture, but with three different training methodologies.

6.1. Model Training Hyperparameters

All models were trained with the Adam optimizer (Kingma & Ba, 2014) and exponentially decreasing learning rate starting at 0.0001 with decay rate 0.98 at every 100000 steps. For evaluation, we use moving exponential parameter averaging at a rate of 0.9999 per step (Polyak, 1990; Polyak & Juditsky, 1992). First, we established trivial baselines by running the built-in first-order theorem prover ASM_MESON_TAC on each theorem on the dataset with empty argument list and with an argument list predicted with our baseline WaveNet model. Next, we compare the performance of various WaveNet style architectures. Finally, we report our reinforcement learning experiments on the complex analysis corpus. Our final prover performance numbers are summarized in Table 2.

6.2. Comparison of Model Architectures

We trained and evaluated a large number of networks and present a sample of our findings. During our experiments, we looked at the following proxy metrics:

1. Accuracy of tactic prediction out of the 41 possible tactics. (Ranging between 38% and 42% for most




Figure 2. This figure presents the cumulative number of proofs closed by the tactic dependent loop. The total number of theorems in the training set is 10199.

models.)

2. Success rate of selecting a positive tactic argument over a randomly selected negative argument. (Around 1% error rate).

For the decoders, we have tried WaveNet (Van Den Oord et al., 2016) style networks with different hyper-parameters. The various results on the complex analysis corpus are based on imitation learning and the combination of imitation learning and reinforcement learning. In the base model we used two WaveNet blocks of four layers each. The number of filters in each block was either 128 or 256. As one can see in Table 2, the network with less filters did better. Then we tried a deeper variant with four blocks of five layers each, in this case with depth 128. Here the deeper network with more blocks turned out to be superior. Both architectures incorporate fully connected combiner layers with additional dropout layers before each of them. The ratio of dropped out neurons during training was 0.3. Note that the reinforcement learning experiments was performed earlier and was ran with the narrow architecture (with 128 filters in each layer) and with two wavenet blocks.

6.3. Reinforcement Learning

In our reinforcement learning set up, the model training runs on a single GPU, while theorem proving is performed in a distributed manner: we attempt to prove 2000 randomly selected theorems from the training set of the union of the complex and core corpora in every round. At the start of each round, we fetch the latest trained model checkpoint and precompute the theorem argument embedding for each theorem in the complex and core libraries. This precomputation greatly accelerates the ranking of the tactic arguments. The proof search is distributed over 1000 cores and we set a computation limit of 100 explored proof states and a total timeout of 300 seconds. Each individual tactic application

Maximum number of top tactics explored	[6, 16]
Maximum <i>successful</i> tactic applications	[3, 6]
Number of selected tactic arguments	[1, 32]

Table 3. Randomized proof search parameters and their ranges.

Name	Theorems proved (% of training set)
Loop	5679 (55.7%)
Loop tactic dependent	5518 (54.1%)
Loop on subgoals	1988 (19.5%)
Union	5919 (58.0%)

Table 4. Total count of proofs found by each loop.

has a timeout of 5 seconds. Additionally, for each example, we pick prover options uniformly in the ranges described by Table 3, to increase the diversity of the generated proofs. This also increases the chance of finding a proof at all for harder statements.

Given the high computational cost of running the reinforcement learning loop, we have only tried a couple of variants. Each of our these experiments use the same version of WaveNet (Van Den Oord et al., 2016) architecture (with 128 filters in each layer). In our first loop experiment, “Loop”, we trained a loop with tactic independent argument selection. That is, the tactic argument ranking was independent of the tactic chosen, and we pick only top level theorems to be proved by proof search in the reinforcement learning scenario. Alongside our first loop, we trained a separate model “Trained on loop output” that was not used in the loop for proof search guidance, but did benefit from a curriculum-style learning, since it trained in parallel to the loop. In our second loop experiment, “Loop tactic dependent”, we have trained a model in which the arguments ranking depends on the selected tactic. In our third loop experiment, “Loop on subgoals”, the proof search can pick from any of the internal proof states from the training set of the joined core+complex corpus. This was motivated partially by the success of (Zombori et al., 2019), we tried to run a reinforcement learning loop in which we train for solving each subgoal separately, hoping that it will help for learning longer proofs. This means, that we expected a bigger variety of theorems to be generated during proof search. However, our naive implementation did not seem to end up with improved results. Performance of each loop’s final checkpoint on the validation set is presented in Table 2. We also ran the final checkpoint of the “Loop” on a sample of 2000 proofs from the *flyspeck* dataset; we closed 752 (37.0%) of these proofs automatically.

While it was too computationally expensive to track the validation performance on every round of the loop, we did




Figure 3. Percentage of theorems proved in each round of the loop. Each round samples 2000 theorems from the training set.

record the performance on the training set. In Fig. 2, we show the cumulative number of proofs closed by the tactic dependent loop at each round. Recall that in each round we sample theorems from the training set and use the most recent checkpoint to guide the proof search. In Fig. 3, we show the percentage of the sampled theorems that are proved in each round.

7. Conclusion

We presented a machine learning oriented open source environment for higher-order theorem proving as well as a neural network based automated prover, trained on a large-scale reinforcement learning system. We also suggest a benchmark for machine reasoning in higher-order logic on a relatively large and practically relevant corpus of theorems with varying complexity. Our benchmark includes purely neural network based baselines, which demonstrate strong automated reasoning capabilities, including premise selection from a large number of theorems. We hope that our initial effort fosters collaboration and paves the way for strong and practical AI systems that can learn to reason efficiently in large formal theories.

Acknowledgements

We would like to thank Alex Alemi, Geoffrey Irving, Cezary Kaliszyk, Ramana Kumar, Viktor Toman, and Josef Urban for their insightful comments and contributions to early versions of this work.

References

Alama, J., Heskes, T., Kühlwein, D., Tsvitvadze, E., and Urban, J. Premise selection for mathematics by corpus

- analysis and kernel methods. *Journal of Automated Reasoning*, 52(2):191–213, 2014.
- Alemi, A. A., Chollet, F., Irving, G., Eén, N., Szegedy, C., and Urban, J. Deepmath-deep sequence models for premise selection. In *Advances in Neural Information Processing Systems*, pp. 2235–2243, 2016.
- Bennett, J., Lanning, S., et al. The netflix prize. In *Proceedings of KDD cup and workshop*, volume 2007, pp. 35. New York, NY, USA, 2007.
- Bertot, Y. and Castéran, P. *Interactive theorem proving and program development: Coq’Art: the calculus of inductive constructions*. Springer Science & Business Media, 2013.
- Blanchette, J. C., Greenaway, D., Kaliszyk, C., Kühlwein, D., and Urban, J. A learning-based fact selector for isabelle/hol. *Journal of Automated Reasoning*, 57(3): 219–244, 2016.
- Bojar, O., Buck, C., Federmann, C., Haddow, B., Koehn, P., Leveling, J., Monz, C., Pecina, P., Post, M., Saint-Amand, H., et al. Findings of the 2014 workshop on statistical machine translation. In *Proceedings of the ninth workshop on statistical machine translation*, pp. 12–58, 2014.
- Bridge, J. P., Holden, S. B., and Paulson, L. C. Machine learning for first-order theorem proving. *J. Autom. Reasoning*, pp. 1–32, 2014. ISSN 0168-7433. doi: 10.1007/s10817-014-9301-5. URL <http://dx.doi.org/10.1007/s10817-014-9301-5>.
- de Moura, L., Kong, S., Avigad, J., Van Doorn, F., and von Raumer, J. The lean theorem prover (system description). In *International Conference on Automated Deduction*, pp. 378–388. Springer, 2015.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. Imagenet: A large-scale hierarchical image database. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pp. 248–255. Ieee, 2009.
- Duncan, H., Bundy, A., Levine, J., Storkey, A., and Pollet, M. The use of data-mining for the automatic formation of tactics. 2004.
- Fan, R.-E., Chang, K.-W., Hsieh, C.-J., Wang, X.-R., and Lin, C.-J. Liblinear: A library for large linear classification. *Journal of machine learning research*, 9(Aug): 1871–1874, 2008.
- Färber, M. and Kaliszyk, C. Random forests for premise selection. In *International Symposium on Frontiers of Combining Systems*, pp. 325–340. Springer, 2015.
- Gauthier, T. and Kaliszyk, C. Premise selection and external provers for hol4. In *Proceedings of the 2015 Conference on Certified Programs and Proofs*, pp. 49–57. ACM, 2015.
- Gauthier, T., Kaliszyk, C., and Urban, J. Tactictoe: Learning to reason with hol4 tactics. In *LPAR-21. 21st International Conference on Logic for Programming, Artificial Intelligence and Reasoning*, volume 46, pp. 125–143, 2017.
- Hales, T., Adams, M., Bauer, G., Dang, T. D., Harrison, J., Le Truong, H., Kaliszyk, C., Magron, V., McLaughlin, S., Nguyen, T. T., et al. A formal proof of the kepler conjecture. In *Forum of Mathematics, Pi*, volume 5. Cambridge University Press, 2017.
- Huang, D., Dhariwal, P., Song, D., and Sutskever, I. Gamepad: A learning environment for theorem proving. *arXiv preprint arXiv:1806.00608*, 2018.
- Kaliszyk, C. and Urban, J. Initial experiments with external provers and premise selection on hol light corpora. 2012.
- Kaliszyk, C. and Urban, J. Stronger automation for flyspeck by feature weighting and strategy evolution. 2013.
- Kaliszyk, C. and Urban, J. Learning-assisted automated reasoning with flyspeck. *Journal of Automated Reasoning*, 53(2):173–213, 2014.
- Kaliszyk, C. and Urban, J. Femalecop: Fairly efficient machine learning connection prover. In *Logic for Programming, Artificial Intelligence, and Reasoning*, pp. 88–96. Springer, 2015a.
- Kaliszyk, C. and Urban, J. Learning-assisted theorem proving with millions of lemmas. *Journal of symbolic computation*, 69:109–128, 2015b.
- Kaliszyk, C. and Urban, J. Mizar 40 for mizar 40. *Journal of Automated Reasoning*, 55(3):245–256, 2015c.
- Kaliszyk, C., Mamane, L., and Urban, J. Machine learning of coq proof guidance: First experiments. *arXiv preprint arXiv:1410.5467*, 2014a.
- Kaliszyk, C., Urban, J., and Vyskočil, J. Machine learner for automated reasoning 0.4 and 0.5. *arXiv preprint arXiv:1402.2359*, 2014b.
- Kaliszyk, C., Urban, J., and Vyskočil, J. Efficient semantic features for automated reasoning over large theories. In *IJCAI*, 2015.
- Kaliszyk, C., Chollet, F., and Szegedy, C. Holstep: A machine learning dataset for higher-order logic theorem proving. *arXiv preprint arXiv:1703.00426*, 2017.

- Kaliszyk, C., Urban, J., Michalewski, H., and Olšák, M. Reinforcement learning of theorem proving. *arXiv preprint arXiv:1805.07563*, 2018.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Klein, G., Elphinstone, K., Heiser, G., Andronick, J., Cock, D., Derrin, P., Elkaduwe, D., Engelhardt, K., Kolanski, R., Norrish, M., Sewell, T., Tuch, H., and Winwood, S. sel4: Formal verification of an os kernel. In *Proceedings of the ACM SIGOPS 22Nd Symposium on Operating Systems Principles, SOSP '09*, pp. 207–220, New York, NY, USA, 2009. ACM. ISBN 978-1-60558-752-3. doi: 10.1145/1629575.1629596. URL <http://doi.acm.org/10.1145/1629575.1629596>.
- Kühlwein, D., van Laarhoven, T., Tsvitsov, E., Urban, J., and Heskes, T. Overview and evaluation of premise selection techniques for large theory mathematics. In *International Joint Conference on Automated Reasoning*, pp. 378–392. Springer, 2012.
- Kühlwein, D., Blanchette, J. C., Kaliszyk, C., and Urban, J. Mash: machine learning for sledgehammer. In *International Conference on Interactive Theorem Proving*, pp. 35–50. Springer, 2013.
- Lin, T.-Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., and Zitnick, C. L. Microsoft coco: Common objects in context. In *European conference on computer vision*, pp. 740–755. Springer, 2014.
- Loos, S., Irving, G., Szegedy, C., and Kaliszyk, C. Deep network guided proof search. *arXiv preprint arXiv:1701.06972*, 2017.
- Megill, N. Metamath: A computer language for pure mathematics. 1997.
- Mizar. The Mizar Mathematical Library. URL <http://mizar.org>. Accessed: 2018/01/18.
- Otten, J. and Bibel, W. leanCoP: lean connection-based theorem proving. *J. Symb. Comput.*, 36(1-2):139–161, 2003.
- Panayotov, V., Chen, G., Povey, D., and Khudanpur, S. Librispeech: an asr corpus based on public domain audio books. In *Acoustics, Speech and Signal Processing (ICASSP), 2015 IEEE International Conference on*, pp. 5206–5210. IEEE, 2015.
- Polyak, B. T. A new method of stochastic approximation type. *Avtomatika i telemekhanika*, 7:98–107, 1990.
- Polyak, B. T. and Juditsky, A. B. Acceleration of stochastic approximation by averaging. *SIAM Journal on Control and Optimization*, 30(4):838–855, 1992.
- Rajpurkar, P., Zhang, J., Lopyrev, K., and Liang, P. Squad: 100,000+ questions for machine comprehension of text. *arXiv preprint arXiv:1606.05250*, 2016.
- Schulz, S. *Learning search control knowledge for equational deduction*, volume 230 of *DISKI*. Infix Akademische Verlagsgesellschaft, 2000. ISBN 978-3-89838-230-4.
- Schulz, S. E - A Brainiac Theorem Prover. *AI Commun.*, 15(2-3):111–126, 2002.
- Slind, K. and Norrish, M. A brief overview of hol4. In *International Conference on Theorem Proving in Higher Order Logics*, pp. 28–32. Springer, 2008.
- Urban, J., Sutcliffe, G., Pudlák, P., and Vyskočil, J. Malarea sg1-machine learner for automated reasoning with semantic guidance. In *International Joint Conference on Automated Reasoning*, pp. 441–456. Springer, 2008.
- Urban, J., Vyskočil, J., and Štěpánek, P. Malecop machine learning connection prover. In *International Conference on Automated Reasoning with Analytic Tableaux and Related Methods*, pp. 263–277. Springer, 2011.
- Van Den Oord, A., Dieleman, S., Zen, H., Simonyan, K., Vinyals, O., Graves, A., Kalchbrenner, N., Senior, A., and Kavukcuoglu, K. Wavenet: A generative model for raw audio. *CoRR abs/1609.03499*, 2016.
- Wang, M., Tang, Y., Wang, J., and Deng, J. Premise selection for theorem proving by deep graph embedding. In *Advances in Neural Information Processing Systems*, pp. 2786–2796, 2017.
- Wenzel, M., Paulson, L. C., and Nipkow, T. The isabelle framework. In Mohamed, O. A., Muñoz, C. A., and Tahar, S. (eds.), *Theorem Proving in Higher Order Logics, 21st International Conference, TPHOLs 2008, Montreal, Canada, August 18-21, 2008. Proceedings*, volume 5170 of *Lecture Notes in Computer Science*, pp. 33–38. Springer, 2008.
- Whalen, D. Holophrasm: a neural automated theorem prover for higher-order logic. *arXiv preprint arXiv:1608.02644*, 2016.
- Zombori, Z., Csiszár, A., Michalewski, H., Kaliszyk, C., and Urban, J. Curriculum learning and theorem proving. In *AITP*, 2019.