

Adversarial Learning to Reason in an Arbitrary Logic

Stanisław J. Purgał and Cezary Kaliszyk

University of Innsbruck
Austria

Abstract

Existing approaches to learning to prove theorems focus on particular logics and datasets. In this work, we propose Monte-Carlo simulations guided by reinforcement learning that can work in an arbitrarily specified logic, without any human knowledge or set of problems. Since the algorithm does not need any training dataset, it is able to learn to work with any logical foundation, even when there is no body of proofs or even conjectures available. We practically demonstrate the feasibility of the approach in multiple logical systems. The approach is stronger than training on randomly generated data but weaker than the approaches trained on tailored axiom and conjecture sets. It however allows us to apply machine learning to automated theorem proving for many logics, where no such attempts have been tried to date, such as intuitionistic logic or linear logic.

Introduction

In the last decade for many logical systems machine learning approaches have managed to improve on the best human heuristics. This worked well for example in classical first-order logic guiding the superposition calculus (Jakubuv et al. 2020), tableaux calculus (Olsák, Kaliszyk, and Urban 2020), or even in higher-order logic (Färber and Brown 2016). In all these works strategies based on machine learning can significantly improve on the best human-designed ones.

To train such machine-learned strategies, datasets of problems and baselines on these problems are required. In particular successful proofs (and in some cases also unsuccessful proofs (Kaliszyk et al. 2018)) are gathered and used to train a machine-learned version of the prover.

In this work, we consider the same problem but without a fixed logic and without a dataset of problems given. We apply a policy-guidance algorithm (known for example from AlphaZero (Silver et al. 2017)) to proving in an arbitrary logic without a given problem set. In particular we:

- propose a theorem-construction game that allows for learning theorem proving with AlphaZero, without relying on training data;
- propose the first dataset for learning for multiple logics, together with learning baselines for this dataset.

- propose an adjusted Monte-Carlo Tree Search that is able to take into account certain (sure) information, when a player makes multiple moves (explained in “Certain Value Propagation” section);
- evaluate the trained prover on the dataset showing that it improves proof capability in various considered logics; and
- as many other problems and games can be directly encoded as logical problems we show that the proposed universal learning for logic also works on some encoded games, such as Sokoban.

Related Work

(Silver et al. 2017) have shown that Monte-Carlo tree search combined with reinforcement learning applied to policy and value functions can generalize to multiple logical games (Go, Chess, Shogi).

The work of (Firoiu et al. 2021) focuses on classical first-order logic without equality and the resolution calculus (so only one of the many logics we consider) and (like us) does not use the complete TPTP problems, but only the axiom sets to learn. There is however a significant overlap between the axioms and conjectures in other problems. The resulting prover does learn to prove theorems but is significantly weaker than E-prover on the TPTP problems. The idea to use only the axioms has already been considered by (Bansal et al. 2019). Even if no original conjectures are exposed to the prover, the dataset used for training is quite large, in comparison with ours, where no formulas are given at all.

As already discussed in the introduction, there are many approaches to applying MCTS with policy and guidance learning in various fixed calculi and on fixed datasets (Kaliszyk et al. 2018; Rawson and Reger 2019). The results are better than those we are able to get here, but no new logics or problems are tried and generalization and transfer have been very limited so far. The AlphaZero algorithm has also been applied in theorem proving to the synthesis of formulas (Brown and Gauthier 2019) and functions (Gauthier 2020).

Kaiser et al. demonstrated that theorem proving can be used to solve many games, such as the ones in the General Game Playing competition (Kaiser and Stafiniak 2011). With the current work, we show that learning for logic can be also applied to these games.

Preliminaries

AlphaZero

The core of the AlphaZero algorithm (Silver et al. 2017) is learning from self-play. It trains a neural network to evaluate the states of a game to estimate the final outcome of the game as well as a policy maximizing the expected outcome. Using the neural network in the current stage of learning a lot of playouts are generated, then this data is used as training data for further improvement.

For training value estimation, the algorithm uses the actual outcomes of the games. To train policy estimation, a Monte-Carlo Tree Search (MCTS) is used to compute a better policy, then the network is trained to return this better policy.

Monte-Carlo Tree Search

For the purpose of training the policy evaluating network we need to provide it with a somewhat better policy. This is done by exploring a tree of possible moves. It is a guided exploration, biased towards the moves pointed to by the policy and to where the value estimations are higher.

A tree is constructed with every node representing a state of the game. Each of those states is evaluated using the neural network. When deciding where to add a new node (thus exploring a branch of the game further) the MCTS algorithm takes into account both the value and the policy estimations from the neural network (biased toward following the policy and higher values), as well as how well a branch was already explored (biased to explore yet unexplored branches more).

After adding a set amount of nodes to the tree, the new better policy is defined to be proportional to the number of nodes explored below each of the immediate children of the root node (representing the state for which we are computing a better policy).

Approach

The theorem-construction game

We propose a two-player game such that a system trained to play the game well could be used to effectively prove the theorems of a given logic system. The first player (referred to as *adversary*) constructs a provable statement, while the other player (referred to as *prover*) tries to prove it. The goal of the adversary is to construct such a theorem, that the prover will fail to prove it. However, because of the available game moves (construction steps), the statements are always provable.

We represent the game objects using Prolog-like *terms*, where a term can be either a *variable* or a pair of an atom and a list of subterms. In the examples, we use the convention of marking variables with capital letters, and denoting compound terms as an atom name followed by a list of subterms in brackets (skipped when the list is empty). For example $\text{tee}(A, \text{implies}(B, \text{false}))$, which can also be expressed with operators like $(A \vdash (B \rightarrow \perp))$.

The construction game is defined for a given set of *inference rules*. An inference rule is a pair of a term and a list of terms, that can share variables. For example

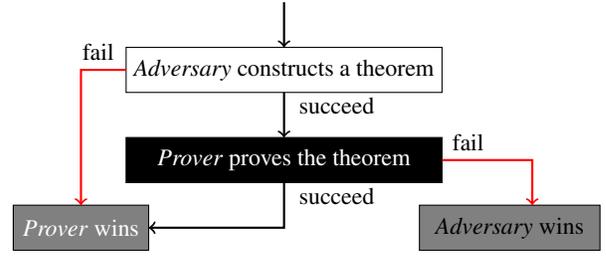


Figure 1: High-level overview of the theorem-construction game.

$\text{tee}(A, \text{and}(B, C)) \leftarrow \text{tee}(A, B), \text{tee}(A, C)$, equivalent to $(A \vdash (B \wedge C)) \leftarrow (A \vdash B), (A \vdash C)$.

For the prover, a game state consists of a list of terms that need to be proven (together with the information that the prover is making the move). During their move, a player can choose one of the given inference rules (the *action space* is the set of *inference rules*), and apply it to the first term of the list. The left side of the rule is then unified with that term. If the unification fails, the player making the move loses. If it succeeds, the term is removed from the list, and the right side of the rule (after unification) is added.

The *adversary* (the player constructing a theorem) makes moves in much the same way, except instead of starting with a theorem to be proven, it starts with a single variable. Applying inference rules to prove this variable will unify it with some term. If the proof is successfully completed, this variable we started with will be unified with a provable theorem. We keep track of this variable in the game state. When the adversary finishes its proof, we pass the constructed theorem to the prover, after replacing all remaining variables with fresh constants.

The second player tries to prove the theorem, winning when the list is empty. To better illustrate the working of our theorem-construction game we present the rules of a concrete game in Figure 2 and an example playout in Figure 3.

$$A, B \vdash A \leftarrow \quad (1)$$

$$A \vdash (B \rightarrow C) \leftarrow (B, A \vdash C) \quad (2)$$

$$A \vdash (B \wedge C) \leftarrow (A \vdash B), (A \vdash C) \quad (3)$$

$$A \vdash B \leftarrow (A \vdash (B \wedge C)) \quad (4)$$

$$A \vdash B \leftarrow (A \vdash (C \wedge B)) \quad (5)$$

$$A \vdash B \leftarrow (A \vdash \perp) \quad (6)$$

Figure 2: A subset of propositional logic inference rules used in the example in figure 3

Certain Value Propagation

The AlphaZero (Silver et al. 2017) algorithm utilizes a neural network to estimate state values (a number in range $(-1, 1)$, we will call it v_θ) and policies (a vector with as many dimensions as the size of the action space). Then a

rule	terms to be proven	constructed theorem
	X	X
2	$A, B \vdash C$	$B \vdash A \rightarrow C$
3	$(A, B \vdash D), (A, B \vdash E)$	$B \vdash A \rightarrow (D \wedge E)$
4	$(A, B \vdash (D \wedge F)), (A, B \vdash E)$	$B \vdash A \rightarrow (D \wedge E)$
1	$(D \wedge F), B \vdash E$	$B \vdash (D \wedge F) \rightarrow (D \wedge E)$
6	$(D \wedge F), B \vdash \perp$	$B \vdash (D \wedge F) \rightarrow (D \wedge E)$
5	$(D \wedge F), B \vdash (G \wedge \perp)$	$B \vdash (D \wedge F) \rightarrow (D \wedge E)$
1		$B \vdash (D \wedge \perp) \rightarrow (D \wedge E)$
	$b \vdash (d \wedge \perp) \rightarrow (d \wedge e)$	
2	$(d \wedge \perp), b \vdash (d \wedge e)$	
6	$(d \wedge \perp), b \vdash \perp$	
5	$(d \wedge \perp), b \vdash (A \wedge \perp)$	
1		Prover won

Figure 3: An example of a playout of the theorem-construction game with the inference rules shown in figure 2

Monte Carlo Tree Search (MCTS) (Kocsis and Szepesvári 2006) is used to compute better estimates of value and policy. This is done by exploring the tree of possible playouts, with a bias toward where value and policy lead to.

During this exploration, MCTS maintains a better value estimation of every state (we will refer to it as v), which is defined to be the average of v_θ of all explored descendants. We will use an equivalent definition, as the weighted average of immediate children, with weights being the number of visits of a given node (the difference will become important).

$$\begin{aligned}
 v(n) &= \frac{v_\theta(n) + \sum_{d < n} v_\theta(d)}{|d : d < n| + 1} = \\
 &= \frac{v_\theta(n) + \sum_{c \in \text{children}(n)} v(c) * (|d : d < c| + 1)}{|d : d < n| + 1}
 \end{aligned}$$

In our version of MCTS, for every node we keep track of a lower and upper bound for possible node values. For non-final nodes, these are simply $(-1, 1)$ (as this is the range of possible outcomes), but for the final nodes, the bounds are both equal to the final reward. These bounds are propagated up the tree in a natural way (taking into account state ownership). Then, for every node we compute a new v_c value, which is simply v adjusted to fall within lower-upper bounds – so eg. if the lower bound is higher than v , then v_c will be equal to the lower bound. Then, when computing v for the nodes above we use this new v_c value rather than the old v .

$$\begin{aligned}
 v_c(n) &= \max(\text{lower}(n), \min(\text{upper}(n), v(n))) \\
 v(n) &= \frac{v_\theta(n) + \sum_{c \in \text{children}(n)} v_c(c) * (|d : d < c| + 1)}{|d : d < n| + 1}
 \end{aligned}$$

Additionally, whenever the value estimation of a state is determined to be -1 (the lowest possible outcome), this state will be avoided. This avoidance is applied both to MCTS exploration and the choice of an action to take during playouts.

The impact of certain value propagation on the final performance of the prover is shown in Fig. 5.

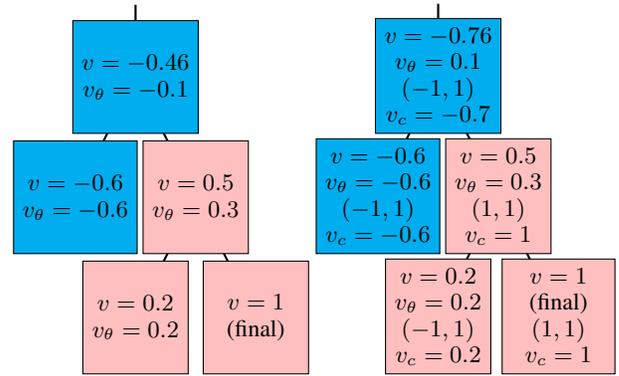


Figure 4: Examples of value estimation in MCTS without Certain Value Propagation (left) and with (right). State ownership (which player is making a move) is marked with color.

Auxiliary replays

To facilitate the prover learning to prove theorems constructed by the adversary we add additional *auxiliary replays*. These come from the games won by the adversary when the prover fails to prove a constructed theorem. Because of the way the theorem was constructed we know how to prove it – we just need to apply the same moves that the adversary used to construct it. Using this fact, we create a replay that shows how the theorem could be proven. In this replay, the policy is not computed using MCTS, but rather is just a one-hot vector pointing to the move that the adversary made when constructing the theorem.

With these auxiliary replays, our algorithm can be considered to train on artificially constructed theorems, that at first come from simply randomly applying inference rules, but later on uses neural guidance to find theorems that the prover cannot yet prove. However, as mentioned earlier, we only do this for theorems that the prover failed to prove.

The impact of including auxiliary replays on the final performance of the prover is shown in figure 5.

Balancing training data

Since the theorem proving game (explain in section “The theorem-construction game”) is asymmetrical, simply using all replay data for training would result in an unbalanced dataset. On top of this, we use auxiliary replays (explained in section “Auxiliary replays”), further disturbing the training data.

To deal with this imbalance we apply training data balancing. Replays are split into parts according to which player won, and a third set of auxiliary replays. All training batches contain the same number of examples from each part.

However, this means disturbing the way the Mean Square Error loss works for value estimation. Consider a value estimation of the starting state. Normally, optimal loss for it would be achieved if the estimate was the average outcome of the game, but with balancing the optimal loss will be achieved by estimating the value to be 0 (mean between losing and winning). This problem affects every state that occurs multiple times in the training dataset.

To counteract this problem, the value loss is weighted in proportion to the size of the part of the data, from which the point originates. So if a player A won in proportion $4 : 1$, the training batches would include games won by this player in proportion $1 : 1$, but $\frac{4}{5}$ of the loss (and therefore gradients) would be determined by data from games won by the player A . For auxiliary replays, this weight is set to 0 and only policy is learned from them.

The impact of balancing training data and weighing the value loss on the final performance of the prover is shown in figure 5.

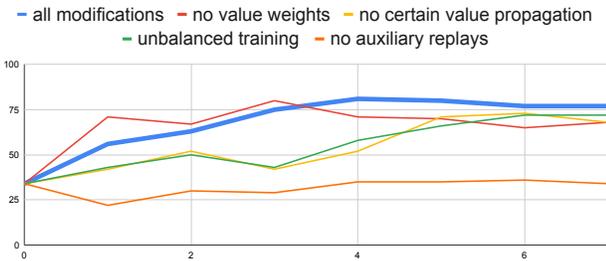


Figure 5: The impact of our modifications on our algorithm. Solved first-order classical tableaux test problems over time (episodes) with 5k games per episode.

Applicability

As mentioned in section “The theorem-construction game”, our system works with a logic system defined by a set of inference rules. This set of rules can be thought of as a Prolog program, and since Prolog is Turing-complete our method can (at least theoretically) be used to learn to reason in any formally defined (and decidable) context. As an example of wide applicability, we train our system to solve Sokoban puzzles. This can be done by defining rules of the game as inference rules of a pseudo-“logic system”.

This of course does not mean that the system will always work well. For example, a saturation prover requires

all terms used in the proof to be fully determined from the start. This negates the advantage of the adversary player, who normally can still modify the constructed theorem late into its proof. Because of this, the probability of constructing a theorem that an untrained prover cannot prove becomes really low – so low that potentially no such theorems will be generated for the initial training set. In such a situation, nothing can be learned from such data and the system is stuck. This problem could potentially be overcome by simply generating enough playouts, but in our experiments with using a saturation prover, the system got stuck after the first step, with the prover winning all games. This was the case in our experiments using a saturation proving method, with 10^4 games generated per episode (possibly more games could help).

Failure states

Another consequence of the game being asymmetrical is the possibility of the training getting stuck when one player starts winning every time. The mechanism of auxiliary replays mentioned above counteracts this to some extent, allowing the prover to still learn even if the adversary is always successfully constructing a hard enough theorem. If the prover was winning every game, however, we would need to rely on exploration for the adversary to find something hard to prove. This situation is however virtually impossible, because of the exploration noise used during playouts. This should lead to adversary towards theorems where the prover is uncertain and sometimes loses due to exploration noise, and then to theorems where the prover fails.

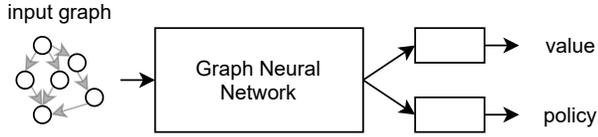
There is however another failure state which if reached would be entirely stable. It is possible because the construction of theorems is inherently easier than proving. Consider a theorem $\exists x \text{HASH}(x) = y$. It is easy to prove such a theorem if one can choose what y is going to be. If y is already decided, proving such a theorem becomes extremely difficult. So difficult in fact, that we cannot hope that a neural network would be able to learn to do this.

If the adversary found such a space of *Uninteresting Hard Theorems*, it would never learn to do anything else. After all, it is a winning strategy for this game. The prover, even using the auxiliary replays would never learn to do anything useful in this situation, and would gradually forget all the useful knowledge learned previously.

This does not seem to happen in any of our experiments. In some of our considered logic systems, it is not even clear that such an Uninteresting Hard Theorem space exists.

Neural architecture

For evaluating state value and policy we use a Graph Neural Network similar to the one described in (Purgal 2020). It is essentially a Graph Attention Network (Veličković et al. 2017) using dot-product attention from the Transformer model (Vaswani et al. 2017) with different attention masks for different attention heads. One Graph Neural Network is used to create a single vector representation of the graph, which is then fed to the final layers to estimate policy and value.



Game states are represented as syntactic graphs. One graph contains all terms that need to be proven, together with information about which player the state belongs to, and (for the adversary player) the state of the constructed theorem. An example of such a graph is shown in figure 6.

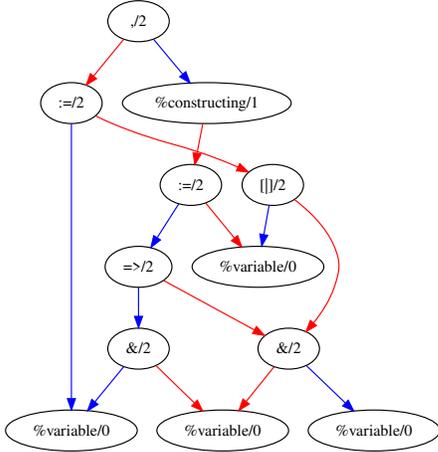


Figure 6: An example graph representing the game state from the game in figure 3 after initial 4 moves.

A single Graph Neural Network is used to evaluate the states for both players, the prover and the adversary.

Evaluation

To test the impact of our method of adversarial training we compare an algorithm trained using our theorem-construction game with a prover trained using uniformly generated random data (an approach somewhat similar to (Firoiu et al. 2021)).

The methods are tested on a dataset not seen by either approach. This test dataset is human-generated (see section “Considered logics” for details on each test dataset) and is often very far outside the training distribution.

Baseline

We generate baseline training data by applying inference rules randomly. This is essentially the adversary from our game doing random moves. Because this does not require evaluating states with a neural network, generating such data is much cheaper, so we generate more playouts – 10^6 (we note that not all playouts result in a constructed theorem).

We use all data generated this way to train a network to estimate policy and value. The policy is a one-hot vector pointing to what the adversary did to construct a theorem, and the value is 0.99^n with n being the number of moves left to do.

logic	baseline	best during game training	total solved during training
int. prop. sequent	12	12	13
classical FO sequent	42	39	40
classical FO tableaux	73	79	83
classical FO Hilbert	38	37	38
modal K prop. sequent	2	5	7
modal T prop. sequent	6	12	13
modal S4 prop. sequent	2	6	8
modal S5 prop. sequent	4	24	24
linear prop. sequent	37	34	39
sokoban solving	4	10	12

Table 1: Results of training in all tested logics – *int.* and *prop.* stand for intuitionistic and propositional

Setup

We implemented the proposed theorem-construction game engine SWI-Prolog (Wielemaker et al. 2010) using PyTorch (Paszke et al. 2019) for the proposed adversarial neural architecture.

We train our system in episodes, first generating 10^4 playouts, then training the neural network using these playouts as training data. This step is repeated multiple times, and after every one, we evaluate the system using the test dataset.

Experiments

To test how much the prover has learned, we play the game in a similar way to when training, except skipping the construction phase, and instead using a theorem from the test set. During such testing we forgo forcing exploration – we do not add exploration noise in Monte-Carlo Tree Search and use the most probable action instead of choosing randomly. Also, when a final state is found during MCTS exploration, we just follow a path to it.

For termination during testing we use a limit on explored states – nodes added to the MCTS tree. Because a part of the tree can be reused for the next state (the part below the node that was chosen) this does not imply any strict turn limit.

Considered logics

Intuitionistic We train our prover on sequent calculus in propositional intuitionistic logic (Huth and Ryan 2000). For test theorems, we use a part of the ILTP library (Raths, Otten, and Kreitz 2005).

Classical We run three experiments with classical first-order logic, trying out three different proof systems. One is sequential calculus, the same as used with intuitionistic logic, another is the Tableaux connection prover (Hähnle 2001), and lastly the rather unwieldy Hilbert system. For the test set, we use a small subset of the Mizar40 dataset (Kaliszyk and Urban 2015) of formulas that do not use equality.

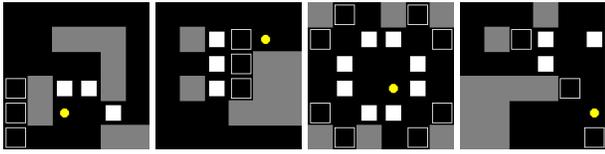


Figure 7: A few examples of Sokoban problems

Linear We also train in linear logic (Girard 1987), only in the propositional setting. For the evaluation we use the LLTP (OlarTE et al. 2020) library, most of which is taken from ILLTP (OlarTE et al. 2019). We also use a few hand-written examples.

Modal In another experiment we train the prover to work with modal propositional logic (Blackburn, van Benthem, and Wolter 2007), in four variants: K, T, S4, S5. Each of those extends the definition of the logic by an additional rule.

For evaluation we use the propositional part of the QMTLP library (Raths and Otten 2011). The set of test theorems is expanded for each consecutive added rule.

Sokoban Sokoban is a classic puzzle game, where the goal is to push boxes into their target positions. The puzzle is PSPACE-complete (Culberson 1997). We only generate puzzles of size 6 by 6. For testing we use a dataset available online¹ (only the problems that can fit into a 6 by 6 grid). A few examples of such problems are shown in figure 7.

Results and Discussion

The results of the evaluation are presented in table 1. We compare the baseline against the last model trained in the adversarial setup, and additionally list the number of unique problems solved in all training epochs. For all modal logics as well as for classical first-order Tableaux the proposed adversarial theorem-construction game leads to many more solved problems. For a number of other logical calculi, the adversarial version is slower but leads to finding solutions different than those trained in the supervised setting, therefore leading to a large number of total solutions found. This is the case for intuitionistic sequent calculus and linear logic. Among the tried calculi, only for the classical sequent-calculus and Hilbert-calculus there is no advantage - this is likely due to the fact that the calculus is much closer to the syntax and the learned baseline can generalize enough. Finally, for the encoded Sokoban games the results are particularly good, with many games solved only in the adversarial-logical setting. We believe, that the adversary learns to construct more and more complex Sokoban-encoded proof games, while the player learns to solve them, in a way similar to curriculum learning.

Forwards vs. backwards conjecturing

During our experiments we briefly considered implementing a different method of constructing theorems, namely

¹<https://sourceforge.net/projects/sokoban-solver-statistics/>

forward constructing: starting from assumptions, working towards the theorem. This method is used in (Firoiu et al. 2021) to generate synthetic data (though without any training for the generator).

The problem with forward constructing (and the reason we decided not to use it) can be illustrated using the following inference rule (disjunction elimination):

$$(T \vdash C) \leftarrow (T \vdash A \vee B), (T \vdash A \rightarrow C), (T \vdash B \rightarrow C)$$

For the adversary to ever successfully apply such a rule, it would first have to construct three statements specifically fitting the rule. In the initial phase of random exploration that would be extremely unlikely. Moreover, the policy predictor will quickly learn that applying this rule always ends in failure, thus making its use even more unlikely after some training. In backwards construction the rule can easily be applied, and simply results in three new statements that consequently need to be proven.

This problem essentially does not exist in the case of a saturation prover, which uses a single inference rule that can be applied anywhere.

Comparison with existing methods

Saturation provers are the state-of-the-art for first-order theorem proving. These are designed to narrow down the search space compared to possibly applying any inference rule at any point. Moreover, their solutions often involve millions of inference steps, while in our case the limit of moves is in the order of 10^2 . As such many problems from the test dataset may not even technically be solvable by our system.

For these reasons, our system performs a lot worse than these in their respective domains. It can however be applied to any formally defined domain and is (as far as the authors know) the only proposed theorem-proving system that may continuously learn and improve without any dataset.

Conclusions

We presented an algorithm for learning to reason in an arbitrary logic. The system, given only a formal definition of a logic, learns to construct increasingly harder problems in the logic and learns to prove them. We show that the system does learn to perform better than a baseline system trained using uniformly generated logical problems. The performance is of course weaker than that of domain-specific Automated Theorem Provers and provers trained on tailored datasets. We are, however, able to construct automatically the first efficient learned automated theorem provers for some logics where none existed before, including various modal logics.

Future work includes encoding more intricate theorem proving calculi, in order to compare them with the more tailored machine-learned systems. Furthermore, for most of the considered logics, the performance on the test sets has stagnated after a few episodes. It remains an open question if trying a compute power comparable with AlphaZero (Silver et al. 2017) would produce significantly better results.

Acknowledgements This work has been supported by the ERC starting grant no. 714034 SMART.

References

- Bansal, K.; Loos, S. M.; Rabe, M. N.; and Szegedy, C. 2019. Learning to reason in large theories without imitation. *ArXiv abs/1905.10501*.
- Blackburn, P.; van Benthem, J. F. A. K.; and Wolter, F., eds. 2007. *Handbook of Modal Logic*, volume 3 of *Studies in logic and practical reasoning*. North-Holland.
- Brown, C. E., and Gauthier, T. 2019. Self-learned formula synthesis in set theory.
- Culberson, J. 1997. Sokoban is pspace-complete.
- Färber, M., and Brown, C. E. 2016. Internal guidance for sattallax. In Olivetti, N., and Tiwari, A., eds., *Automated Reasoning - 8th International Joint Conference, IJCAR 2016, Coimbra, Portugal, June 27 - July 2, 2016, Proceedings*, volume 9706 of *Lecture Notes in Computer Science*, 349–361. Springer.
- Firoiu, V.; Aygun, E.; Anand, A.; Ahmed, Z.; Glorot, X.; Orseau, L.; Zhang, L.; Precup, D.; and Mourad, S. 2021. Training a first-order theorem prover from synthetic data.
- Gauthier, T. 2020. Deep reinforcement learning for synthesizing functions in higher-order logic. In Albert, E., and Kovács, L., eds., *LPAR 2020: 23rd International Conference on Logic for Programming, Artificial Intelligence and Reasoning, Alicante, Spain, May 22-27, 2020*, volume 73 of *EPiC Series in Computing*, 230–248. EasyChair.
- Girard, J.-Y. 1987. Linear logic. *Theoretical computer science* 50(1):1–101.
- Hähnle, R. 2001. Tableaux and related methods. In Robinson, J. A., and Voronkov, A., eds., *Handbook of Automated Reasoning (in 2 volumes)*. Elsevier and MIT Press. 100–178.
- Huth, M., and Ryan, M. D. 2000. *Logic in computer science - modelling and reasoning about systems*. Cambridge University Press.
- Jakubov, J.; Chvalovský, K.; Olsák, M.; Piotrowski, B.; Suda, M.; and Urban, J. 2020. ENIGMA anonymous: Symbol-independent inference guiding machine (system description). In Peltier, N., and Sofronie-Stokkermans, V., eds., *Automated Reasoning - 10th International Joint Conference, IJCAR 2020, Paris, France, July 1-4, 2020, Proceedings, Part II*, volume 12167 of *Lecture Notes in Computer Science*, 448–463. Springer.
- Kaiser, L., and Stafiniak, L. 2011. First-order logic with counting for general game playing. In Burgard, W., and Roth, D., eds., *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2011, San Francisco, California, USA, August 7-11, 2011*. AAAI Press.
- Kaliszyk, C., and Urban, J. 2015. MizAR 40 for Mizar 40. *J. Autom. Reasoning* 55(3):245–256.
- Kaliszyk, C.; Urban, J.; Michalewski, H.; and Olsák, M. 2018. Reinforcement learning of theorem proving. In Bengio, S.; Wallach, H.; Larochelle, H.; Grauman, K.; Cesa-Bianchi, N.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 31*, 8836–8847. Curran Associates, Inc.
- Kocsis, L., and Szepesvári, C. 2006. Bandit based monte-carlo planning. In *European conference on machine learning*, 282–293. Springer.
- Olarte, C.; de Paiva, V.; Pimentel, E.; and Reis, G. 2019. The illtp library for intuitionistic linear logic. *arXiv preprint arXiv:1904.06850*.
- Olarte, C.; de Paiva, V.; Pimentel, E.; and Reis, G. 2020. Linear logic theorem proving. <https://github.com/meta-logic/lltp>.
- Olsák, M.; Kaliszyk, C.; and Urban, J. 2020. Property invariant embedding for automated reasoning. In Giacomo, G. D.; Catalá, A.; Dilkina, B.; Milano, M.; Barro, S.; Bugarín, A.; and Lang, J., eds., *ECAI 2020 - 24th European Conference on Artificial Intelligence and Applications*, 1395–1402. IOS Press.
- Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; et al. 2019. Pytorch: An imperative style, high-performance deep learning library. In *Advances in neural information processing systems*, 8026–8037.
- Purgal, S. J. 2020. Improving expressivity of graph neural networks. *2020 International Joint Conference on Neural Networks (IJCNN)* 1–7.
- Raths, T., and Otten, J. 2011. The qmltp library: Benchmarking theorem provers for modal logics. Technical report, Technical Report, University of Potsdam.
- Raths, T.; Otten, J.; and Kreitz, C. 2005. The iltp library: Benchmarking automated theorem provers for intuitionistic logic. In *International Conference on Automated Reasoning with Analytic Tableaux and Related Methods*, 333–337. Springer.
- Rawson, M., and Reger, G. 2019. A neurally-guided, parallel theorem prover. In *FroCos*.
- Silver, D.; Hubert, T.; Schrittwieser, J.; Antonoglou, I.; Lai, M.; Guez, A.; Lanctot, M.; Sifre, L.; Kumaran, D.; Graepel, T.; Lillicrap, T. P.; Simonyan, K.; and Hassabis, D. 2017. Mastering chess and shogi by self-play with a general reinforcement learning algorithm. *ArXiv abs/1712.01815*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention is all you need.
- Veličković, P.; Cucurull, G.; Casanova, A.; Romero, A.; Lio, P.; and Bengio, Y. 2017. Graph attention networks. *arXiv preprint arXiv:1710.10903*.
- Wielemaker, J.; Schrijvers, T.; Triska, M.; and Lager, T. 2010. Swi-prolog.