

Learned Provability Likelihood for Tactical Search*

Thibault Gauthier

Czech Technical University in Prague, Prague, Czech Republic

email@thibaultgauthier.fr

We present a method to estimate the provability of a mathematical formula. We adapt the tactical theorem prover TacticToe to factor in these estimations. Experiments over the HOL4 library show an increase in the number of theorems re-proven by TacticToe thanks to this additional guidance. This amelioration in performance together with concurrent updates to the TacticToe framework lead to an improved user experience.

1 Introduction

We take inspiration from the instinct shown by mathematicians when attempting to prove a theorem. An important choice to make in their situation is whether to continue exploring a certain line of work or to switch to another approach entirely. The quality of this decision is influenced greatly by one's experience with other proof attempts. Our aim is to integrate such feedback to improve automation in interactive theorem provers (ITPs). In the majority of ITPs, most proofs are build using tactics in a goal-oriented manner. A tactic is a procedure that takes as input a goal g (a sequent in HOL4) and returns a list of goals whose conjunction implies g . We can classify HOL4 tactics into four categories. Solvers attempts to prove the goal using general or domain-specific knowledge (e.g. *metis_tac* [10] for first-order logic). Rewrite tactics modify subterms of a goal by applying rewrite rules constructed from proven equalities, Induction tactics (e.g. *Induct*) split the goal into multiple cases (typically a base case and an inductive case). Kernel-level tactics (e.g. *exist_tac*) allow for more refined control of the proof state. Given these tactics, it is possible to create an automated prover that searches for the proof of a starting goal by predicting suitable tactics for each intermediate goal. In the following, we refer to such automation as a tactic-based automated theorem prover (ATP). Multiple tactic-based ATPs have been developed in the course of the last five years and are now one of the most effective [8, 4] general proof automation available in an ITP. However, none of the tactic-based ATP so far take into account the provability of the goals produced by each tactic. Therefore, in this project, a tree neural network (TNN) is taught a function estimating the provability of goals in HOL4 [14]. This *value* function produces feedback signals called rewards that further guide the search algorithm of the tactic-based ATP TacticToe [8].

Related Works The most successful related tactic-based ATPs are Tactician [4] for Coq [3], Pamper [12] for Isabelle/HOL [15], and HOList [2] for HOL Light [9]. The Tactician is very user-friendly. In particular, it is the only one that can record tactic calls on the fly. It uses the k -nearest neighbor algorithm for tactic selection as TacticToe does. However, it does not predict argument theorems independently of tactics. In Pamper, the policy predictors are decision trees trained on top of human-engineered features. In HOList, the prediction effort is concentrated on learning the policy for a few selected tactics and their arguments (theorems) using deep reinforcement learning. A related field of research is machine learning

*Supported by the Czech Science Foundation project 20-06390Y

for first-order ATPs. The ENIGMA [5] system for E prover gathers positive and negative clauses from successful proof attempts. We use a similar technique to collect our training examples. The ATP Lean-CoP [11] has been trained via a reinforcement learning loop using boosted random forest predictors. Its proof search relies on the same variant of Monte Carlo Tree Search (MCTS) [13] as TacticToe. Nevertheless, it ignores the goal selection issue and thus could benefit directly from the MCTS adaptations proposed in this paper.

2 Monte Carlo Tree Search with Tactics

We integrate the learned provability estimator into the proof search of TacticToe. Here is a brief summary of how the MCTS algorithm operates in the context of tactic-based theorem proving. The algorithm starts with a list of goals (typically a singleton) to be proven in a root node. In the selection phase, it chooses a goal branch in the current output node and a tactic branch (and possibly an argument branch) leading to the selection of an output node containing the list of goals produced by the tactic. This process is repeated until a leaf is reached. In the extension phase, the tactic t selected in the leaf is applied to its parent goal. In the case of a successful tactic application, an output leaf is created containing the list of goals produced by t . During the backup phase, a feedback signal is propagated from the newly created leaf to the root. The gathered node rewards influence the next selection phase. The three phases are repeated until the algorithm finds a proof for each of the root goals, times out, or saturates. In the following, we present the existing MCTS algorithm for TacticToe and describe ways to improve the quality of the feedback mechanism. Figure 1 illustrates the effect of one iteration of the improved MCTS loop on the search tree for our running example.

Selection Phase Given a list of tactic $(t_r)_{1 \leq r \leq n}$ with parent node g , we can compute their PUCT [1] (Polynomial Upper Confidence Trees) score as follows:

$$PUCT(t_r) = AverageRewards(t_r) + c \times Policy(t_r) \times \frac{\sqrt{Visits(t_r)}}{Visits(g)}$$

The tactic with the highest PUCT score is selected. The policy *Policy* is given by a nearest neighbor predictor trained from supervised data consisting of goal-tactic pairs [8]. The value of $Policy(t_j)$ is experimentally chosen to be 0.5^{n+1} where n is the number of open tactic branches with parent node g and higher nearest neighbor score. At the start, the search is principally guided towards nodes with higher policy. As the number of iterations increases, the search tends to explore goals with higher reward averages more often. The exploration coefficient c decides how fast this transition happens. We choose c to be 2.0 since this is a suitable value for guiding a goal-oriented first-order ATP [11].

In this version of TacticToe, the theorems predicted by its nearest neighbor algorithm are split. Given a list of predicted arguments x_1, \dots, x_n and a tactic t (except *metis_tac*) that expects a list of theorems as argument, the tactic calls $t[x_1], \dots, t[x_n]$ are constructed instead of $t[x_1, \dots, x_n]$. This adds another branching layer to the tree that functions exactly like the tactic selection layer. To simplify our explanations in the rest of this paper, the branching occurring during argument selection is considered to be part of tactic selection.

In previous developments of TacticToe and in other tactic-based ATPs, the selected goal is always the first unproven goal of the output node. One issue is that the rewards of an output node are exactly the rewards of its first goal until it is proven. To factor the influence of other goals in the output node rewards, each unproven goal branch now receives almost the same number of visits. This is achieved by choosing at each iteration one of the least visited unproven goal branches.

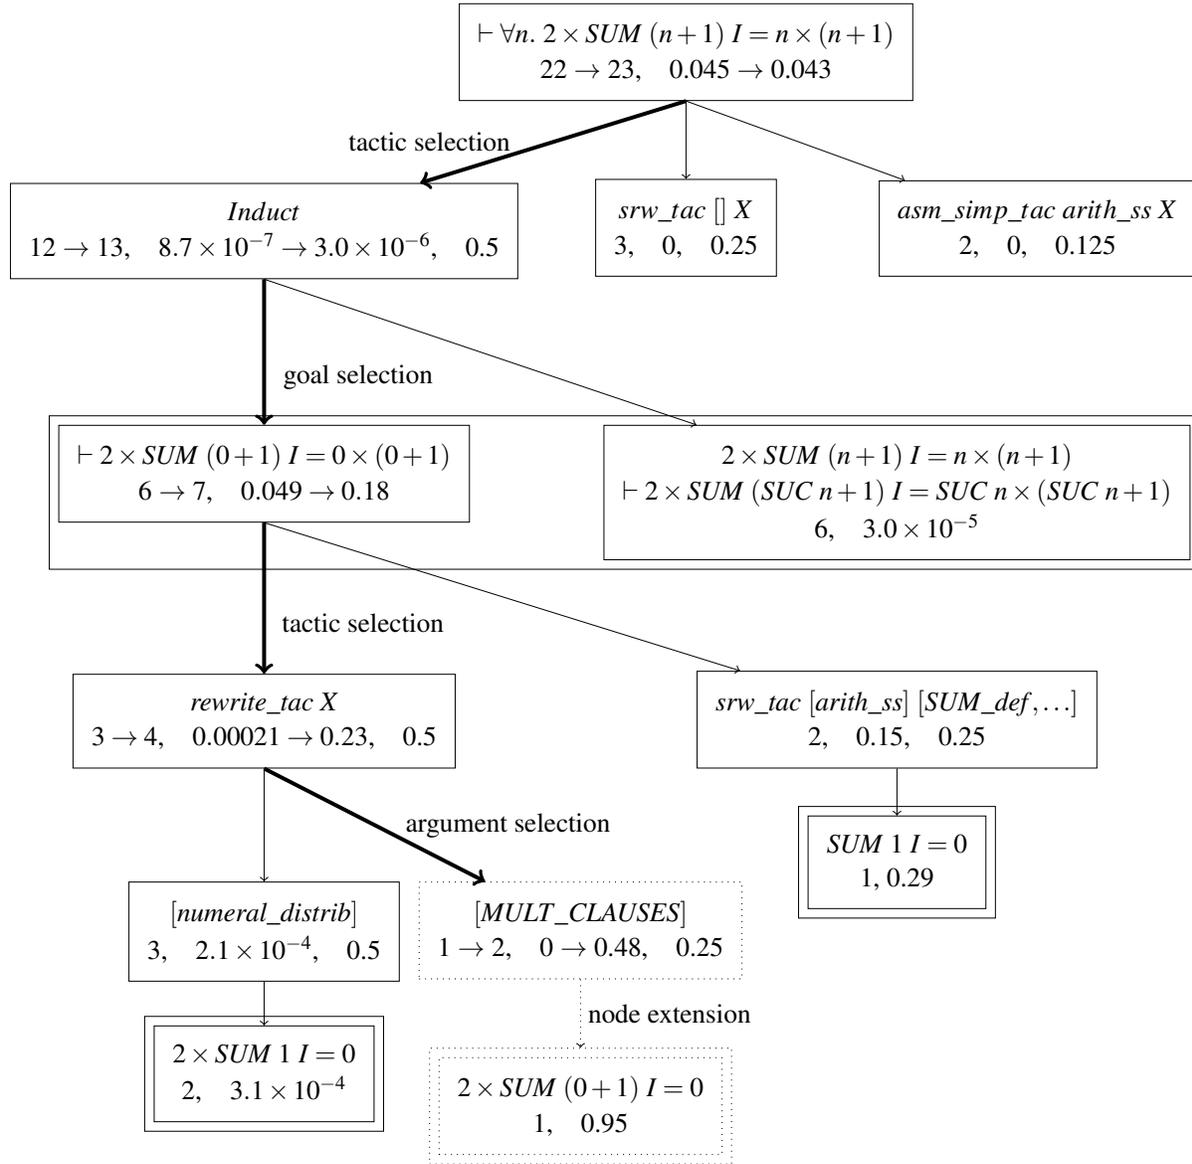


Figure 1: Iteration 22 of the value-guided MCTS loop on the goal $\vdash \forall n. 2 \times \text{SUM } (n+1) I = n \times (n+1)$. The term $\text{SUM } n f$ stands for $\sum_{x=0}^{n-1} f(x)$ and I is the identity function. Each node contains in the following order: the number of visits with a possible update, the average of the rewards and the prior policy for tactics (and arguments). The presence of an arrow after these numbers indicates a backup update. The selection path is made bold and created nodes are dotted. Saturated tactic (and argument) nodes and the subtree of the inductive case are omitted. Tactics may contain the placeholder X [8] to indicate that an argument has to be provided. To avoid dividing by 0 in the PUCT formula, all tactic (and argument) nodes are initialized with one visit and a reward of 0. A proof is found after 427 iterations.

Backup Phase There are three possible outcomes of the extension phase. If the tactic t applied proves the parent goal g , t receives a reward of 1. If t fails or induces a loop then t receives a reward of 0. Otherwise, a new output node is created and the value network is called on each of the goals. The reward of the newly created leaf is computed by multiplying the inferred values and backed up unchanged to t . Each tactic selection layer propagates the reward of the selected tactic unchanged to its parent goal. To capture the influence of multiple explored goal branches in the goal selection layer, we back up the rewards through this layer in the following manner. Given a list of goals $(g_r)_{1 \leq r \leq n}$ composing an output node p and a selected goal g_i , the reward for p is given by:

$$Reward_k(p) = Reward_k(g_i) \times \prod_{1 \leq r \leq n \wedge r \neq i} AverageRewards(g_r)$$

If a goal g_r is proven, the value of $AverageRewards(g_r)$ in the formula is overridden and set to 1. This formula multiplies the reward of the goal g_i at iteration k of the search loop with the existing reward average $AverageRewards$ for the other goals in p . Thus the feedback of different branches is merged although only one branch is explored at each iteration of the loop. Overall, p receives a reward that is a lower bound estimation for its provability as it assumes independence of its goals.

3 Learning Provability

We explain here how a TNN can be trained to estimate the provability of a goal. These estimates are to be used as rewards in our improved MCTS algorithm. We choose a TNN as our machine learning model because it performs well on arithmetic and propositional formulas [7] as well as on Diophantine equations and combinators [6]. In our TNN, each HOL4 operator of arity a has a neural network associated with it modeling a function from $\mathbb{R}^{a \times d}$ to \mathbb{R}^d , where d is a globally fixed embedding size. When $a = 0$, the associated neural network is a trainable embedding (vector in \mathbb{R}^d). Networks are composed in a manner that reflects the tree structure of a given goal g . They gradually construct embeddings for sub-trees of g . A head network models a function from \mathbb{R}^d to \mathbb{R} targeting a provability estimation for g from an embedding for g . Figure 2 shows part of the computation graph produced by a TNN on the running example.

Until now, the value function of TacticToe was uniform. It gave a reward of 1 independently of the provability of the goal produced. Thus, the search was guided solely by a policy trained from human-written proof scripts. The lack of negative examples in these proofs makes this dataset unsuitable for training the value function. That is why we extract training examples from TacticToe searches instead. To do so, we collect search trees from the successful proof attempts. The goal nodes extracted from these trees create our training examples with proven goals labeled positively and other goals negatively. In order to speed up the training process, we keep at most one example for each goal with a preference for the positive ones. We also remove examples with large goals (≥ 80 operators) and select the 600 most visited negative examples per proof attempt. From this dataset, we train a TNN with embedding size ($d = 16$) and one layer per operator. The training is scheduled to take 100 epochs at a learning rate of 0.08 and the batch size is increased regularly according to the sequence [16, 24, 32, 48, 64]. These parameters were optimized by experimenting with a fifth of the HOL4 library. There, the parameters used in [7] were gradually mutated to yield higher accuracy on HOL4 goals.

4 Results

We experiment on the HOL4 library included in the HOL4 repository ¹. This library contains 168 theories about mathematical and computer science concepts such as lists, trees, probabilities, measures and

¹<https://github.com/HOL-Theorem-Prover/HOL>

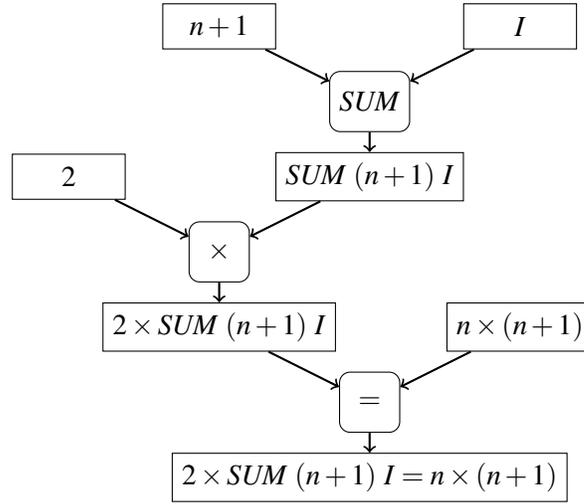


Figure 2: Computation of the embedding of $2 \times \text{SUM}(n+1) I = n \times (n+1)$ by a TNN. To simplify the diagram, the computation of the embeddings of 1, 2, $n+1$, $n \times (n+1)$ are not shown. Rectangles represent embeddings and squares represent neural networks.

integration. During the development of HOL4 library, the proof of a theorem is usually found as an argument of one of the following functions: `store_thm`, `maybe_thm`, `Store_thm`, `asm_store_thm`, `prove` or `TAC_PROOF`. We introduce a hook in these functions in order to evaluate `TacticToe` on their argument goal as the library is built. In this way, only tactics, theorems and simplification sets created before a goal g is proven are available to `TacticToe` when attempting to prove g . The code for this experiment is available at this commit². Replication instructions are given in the file `src/tacticToe/EVALUATION`.

After an evaluation of `TacticToe` with a 30 seconds timeout over 15948 theorems of the library resulting in 8812 successful searches, 188409 examples (34021 positives and 154388 negatives) are collected. Correcting for the imbalance between the positive and negative examples by oversampling positive examples did not seem to improve the overall performance when experimenting with a small part of the library. Therefore, no balancing method is applied in this full-scale experiment. The efficiency of the learning phase is then assessed by splitting the examples into a training set (90%) and a test set (10%). Following the training of the TNN, we measure an accuracy of 97.5% on the training set and 84.7% on the test set. For the final evaluation, the TNN is retrained on all examples.

The performance of value-guided `TacticToe` is compared with its default version in Table 1. All changes to the proof search algorithm proposed in this paper are included in both versions. The only difference between the two versions is whether a uniform or trained value function gives the reward signal. Overall, the results show a small increase in the number of theorems re-proven. Calls to the TNN increase the node creation time from 3.7% to 13.3% of the total search time. Such a small footprint could not have been achieved without a Standard ML implementation of TNNs [7]. Thus, the positive impact of learning is not severely dampened by the neural network overhead.

4.1 Proofs

Among the 239 theorems solely proven by `TacticToeTNN`, 30 of them belong to the theory `real_topology`. Two such theorems with their tactical proof are analyzed here. Compared to the previous version of

²13fbff7b94bb1a5b51881a5aeee63ffc44d3be1

	TacticToe	TacticToe _{TNN}	Combined
Standard library (15948 theorems)	8812 (138)	8913 (239)	9051

Table 1: Number of theorems in the HOL4 standard library re-proven within 30 seconds. The number of theorems re-proven by one strategy but not by the other is shown in parentheses.

TacticToe, the proofs are now automatically printed with the preferred style of many HOL4 users. Tactics are written in lowercase when possible and the tacticals `>>` (one goal), `>|` (multiple goals) compose tactics in the final proof script.

The theorem $\vdash \forall a. \text{diameter } \{a\} = 0$ is re-proven in 28.2 seconds. It states that the diameter of any singleton is equal to 0. The diameter of a set of reals s is by definition 0 if s is empty and otherwise is the supremum $\{\text{abs}(x - y) \mid x \in s \wedge y \in s\}$. The proof starts by rewriting the goal with the definition of *diameter*, continues by distinguishing whether s is empty or not and proves each case by calling a first-order solver and a simplification tactic with an appropriate lemma.

```
rewrite_tac [diameter] >> REPEAT strip_tac >>
COND_CASES_TAC >| [metis_tac [], srw_tac [] [REAL_SUP_UNIQUE]]
```

The theorem $\vdash \forall f s. (\text{linear } f \wedge \text{subspace } s) \Rightarrow \text{subspace } (\text{IMAGE } f s)$ is re-proven in 16.7 seconds. It states that the image of a subspace by a linear function is a subspace. A subspace s is a subset of \mathbb{R} that satisfies three conditions $0 \in s$, $(x \in s \wedge y \in s) \Rightarrow x + y \in s$, and $\forall c \in \mathbb{R}. x \in s \Rightarrow c \times x \in s$. The proof consists of rewriting the goal with this definition and then solving each of the three cases by calling *metis_tac* with suitable premises.

```
srw_tac [] [subspace] >|
[metis_tac [REAL_MUL_LZERO, linear],
 first_assum (X_CHOOSE_TAC ``B`` o MATCH_MP (LINEAR_BOUNDED)) >>
  pop_assum (mp_tac o Q.SPEC `x`) >> metis_tac [LINEAR_ADD],
 metis_tac [LINEAR_CMUL]]
```

5 Usage

To make TacticToe more attractive to new users, tactic-goal pairs from the HOL4 library are prerecorded. Therefore, the users can now use TacticToe right after building HOL4 without spending multiple hours recording the data themselves (see `HOL/src/tactictoe/README`). Multiple new functions relying on the TNN trained in this paper are implemented and available at the commit mentioned in Section 4. To run these functions, one first need to download the file `tnn_for_tactictoe` from <http://grid01.ciirc.cvut.cz/~thibault/> to a desired location `path_to_foo`. The function `ttt_tnn` runs TacticToe_{TNN} on a chosen goal with the advice of a TNN imported by `mlTreeNeuralNetwork.read_tnn`. The following commands demonstrate how to execute `ttt_tnn` in an interactive session:

```
load "tacticToe"; open tacticToe;
val tnn = mlTreeNeuralNetwork.read_tnn "path_to_foo";
load "sum_numTheory"; open sum_numTheory; open arithmeticTheory;
set_timeout 60.0;
ttt_tnn tnn ([, ``!n. 2 * SUM (n+1) I = n * (n+1) ``);
```

The function `confidence_tnn` returns the provability estimate of a goal as computed by the TNN in 1 to 10 milliseconds. This might encourage the user to call TacticToe_{TNN} if this value becomes high on an open goal during a manual proof attempt.

The function `suggest` creates partial proofs from failed proof attempts. By limiting the numbers of iterations of the MCTS loop to 22 (`tttSearch.looplimit := SOME 22;`), a call to `ttt_tnn` now fails on the running example but the search tree is stored in the reference `searchtree_glob`. Executing `suggest ()`; at this point returns the most promising partial proof tree in `searchtree_glob`. To construct the tree, this function traverses down the tree and selects the first proved tactic branch or the most visited open tactic branch if no branch has been proved. A tactic built from this proof tree can then be applied to the original goal and the user might continue working on the produced open goals. The resulting partial proof suggested from our failed attempt is:

```
Induct >| [rewrite_tac [numeralTheory.numeral_distrib] >> all_tac,
          srw_tac [ARITH_ss] [MULT_CLAUSES, LEFT_ADD_DISTRIB] >> all_tac]
```

The depth of this advice may be controlled by updating the reference `tttSearch.suggest_depth`. The function `suggest` can also be used in combination with the original `ttt` function but the advice will be of lesser value as the most promising proof branches are fixed by the policy and therefore the most promising proof tree is gradually extended as the search tree grows. In contrast, the decision on which path is most promising might be revised when running `ttt_tnn` with higher timeouts.

6 Conclusion

Training a TNN on intermediate goals generated by `TacticToe` leads to the creation of a goal provability estimator with good accuracy. Since a goal can be split into multiple goals during tactic-based searches, we propose improvements to the selection and backup phases of the underlying MCTS algorithm in order to get more accurate feedback from the learned estimator. Experiments on the full HOL4 library measure an increase in the success rate of `TacticToe` and demonstrate the scalability of our approach. This resulted in the creation of new tools for HOL4 users that leverage the advice given by the estimator. In particular, `TacticToe` is now able to suggest how to start a proof even when it fails to find one.

References

- [1] David Auger, Adrien Couëtoux & Olivier Teytaud (2013): *Continuous Upper Confidence Trees with Polynomial Exploration - Consistency*. In: *Machine Learning and Knowledge Discovery in Databases - European Conference, ECML PKDD 2013, Prague, Czech Republic, September 23-27, 2013, Proceedings, Part I*, pp. 194–209, doi:10.1007/978-3-642-40988-2_13.
- [2] Kshitij Bansal, Sarah M. Loos, Markus N. Rabe, Christian Szegedy & Stewart Wilcox (2019): *HOList: An Environment for Machine Learning of Higher Order Logic Theorem Proving*. In Kamalika Chaudhuri & Ruslan Salakhutdinov, editors: *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA, Proceedings of Machine Learning Research 97*, PMLR, pp. 454–463. Available at <http://proceedings.mlr.press/v97/bansal19a.html>.
- [3] Yves Bertot (2008): *A Short Presentation of Coq*. In Otmane Ait Mohamed, César Muñoz & Sofiène Tahar, editors: *Conference on Theorem Proving in Higher Order Logics (TPHOLs)*, LNCS 5170, Springer, pp. 12–16, doi:10.1007/978-3-540-71067-7_3.
- [4] Lasse Blaauwbroek, Josef Urban & Herman Geuvers (2020): *The Tactician - A Seamless, Interactive Tactic Learner and Prover for Coq*. In Christoph Benzmüller & Bruce R. Miller, editors: *Intelligent Computer Mathematics - 13th International Conference, CICM 2020, Bertinoro, Italy, July 26-31, 2020, Proceedings*, LNCS 12236, Springer, pp. 271–277, doi:10.1007/978-3-030-53518-6_17.
- [5] Karel Chvalovský, Jan Jakubuv, Martin Suda & Josef Urban (2019): *ENIGMA-NG: Efficient Neural and Gradient-Boosted Inference Guidance for E*. In: *Automated Deduction - CADE 27 - 27th International Conference on Automated Deduction, Natal, Brazil, August 27-30, 2019, Proceedings*, pp. 197–215, doi:10.1007/978-3-030-29436-6_12.

- [6] Thibault Gauthier (2020): *Deep Reinforcement Learning for Synthesizing Functions in Higher-Order Logic*. In Elvira Albert & Laura Kovács, editors: *LPAR 2020: 23rd International Conference on Logic for Programming, Artificial Intelligence and Reasoning, Alicante, Spain, May 22-27, 2020, EPIc Series in Computing 73*, EasyChair, pp. 230–248, doi:10.29007/7jmg.
- [7] Thibault Gauthier (2020): *Tree Neural Networks in HOL4*. In Christoph Benzmüller & Bruce R. Miller, editors: *Intelligent Computer Mathematics - 13th International Conference, CICM 2020, Bertinoro, Italy, July 26-31, 2020, Proceedings, LNCS 12236*, Springer, pp. 278–283, doi:10.1007/978-3-030-53518-6_18.
- [8] Thibault Gauthier, Cezary Kaliszzyk, Josef Urban, Ramana Kumar & Michael Norrish (2021): *TacticToe: Learning to Prove with Tactics*. *J. Autom. Reason.* 65(2), pp. 257–286, doi:10.1007/s10817-020-09580-x.
- [9] John Harrison (2009): *HOL Light: An Overview*. In Stefan Berghofer, Tobias Nipkow, Christian Urban & Makarius Wenzel, editors: *Conference on Theorem Proving in Higher Order Logics (TPHOLs), LNCS 5674*, Springer, pp. 60–66, doi:10.1007/978-3-642-03359-9_4.
- [10] Joe Hurd (2005): *System Description: The Metis Proof Tactic*. In Carsten Schuermann Christoph Benzmueller, John Harrison, editor: *Workshop on Empirically Successful Automated Reasoning in Higher-Order Logic (ESHOL)*, pp. 103–104. Available at <https://arxiv.org/pdf/cs/0601042>.
- [11] Cezary Kaliszzyk, Josef Urban, Henryk Michalewski & Miroslav Olsák (2018): *Reinforcement Learning of Theorem Proving*. In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi & Roman Garnett, editors: *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pp. 8836–8847. Available at <https://proceedings.neurips.cc/paper/2018/hash/55acf8539596d25624059980986aaa78-Abstract.html>.
- [12] Yutaka Nagashima & Yilun He (2018): *PaMpeR: proof method recommendation system for Isabelle/HOL*. In Marianne Huchard, Christian Kästner & Gordon Fraser, editors: *Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering, ASE 2018, Montpellier, France, September 3-7, 2018*, ACM, pp. 362–372, doi:10.1145/3238147.3238210.
- [13] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, Yutian Chen, Timothy Lillicrap, Fan Hui, Laurent Sifre, George van den Driessche, Thore Graepel & Demis Hassabis (2017): *Mastering the game of Go without human knowledge*. *Nature* 550, pp. 354–359, doi:10.1038/nature24270.
- [14] Konrad Slind & Michael Norrish (2008): *A Brief Overview of HOL4*. In Otmane Aït Mohamed, César A. Muñoz & Sofïène Tahar, editors: *Conference on Theorem Proving in Higher Order Logics (TPHOLs), LNCS 5170*, Springer, pp. 28–32, doi:10.1007/978-3-540-71067-7_6.
- [15] Makarius Wenzel, Lawrence C. Paulson & Tobias Nipkow (2008): *The Isabelle Framework*. In Otmane Aït Mohamed, César A. Muñoz & Sofïène Tahar, editors: *Conference on Theorem Proving in Higher Order Logics (TPHOLs), LNCS 5170*, Springer, pp. 33–38, doi:10.1007/978-3-540-71067-7_7.