

# Hammering Mizar by Learning Clause Guidance

Jan Jakubův

Czech Technical University in Prague, Czech Republic

Josef Urban

Czech Technical University in Prague, Czech Republic

---

## Abstract

---

We describe a very large improvement of existing hammer-style proof automation over large ITP libraries by combining learning and theorem proving. In particular, we have integrated state-of-the-art machine learners into the E automated theorem prover, and developed methods that allow learning and efficient internal guidance of E over the whole Mizar library. The resulting trained system improves the real-time performance of E on the Mizar library by 70% in a single-strategy setting.

**2012 ACM Subject Classification** Theory of computation → Automated reasoning; Computing methodologies → Theorem proving algorithms; Computing methodologies → Machine learning

**Keywords and phrases** Proof automation, ITP hammers, Automated theorem proving, Machine learning

**Digital Object Identifier** 10.4230/LIPIcs.ITP.2019.34

**Category** Short Paper

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

## 1 Introduction

Proof automation for interactive theorem provers (ITPs) has been a major factor behind the recent progress in formal verification. In particular, *Hammers* linking ITPs with automated theorem provers (ATPs) produce a major speedup of formalization [4]. The main AI component of existing hammers has so far been *premise selection* [1], where only the most relevant facts are chosen from the large ITP libraries as axioms for proving a new conjecture. Machine learning from the large number of proofs in the ITP libraries has resulted in the strongest premise selection methods [1, 17, 19, 8, 3, 2]. Premise selection however does not guide the theorem proving processes once the premises are selected. The success of machine learning in the high-level premise selection task has motivated development of low-level *internal proof search guidance*. This has been recently started both for ATPs [29, 18, 15, 23, 9] and also in the context of tactical ITPs [10, 13].

Recently, we have added [6] two state-of-the-art machine learning methods to the ENIGMA [15, 16] algorithm that efficiently guides saturation-style proof search in ATPs such as E [25, 26]. The first method trains gradient boosted trees on efficiently extracted manually designed (handcrafted) clause features. The second method uses end-to-end training of recursive neural networks, thus removing the need for handcrafted features. While the second method seems very promising and already improves on a simpler linear classifier when used for guidance, its efficient training and use over a large ITP library is still practically challenging. On the other hand, our recent experiments with efficient *feature hashing* have shown that the very good performance of gradient boosted trees is maintained even after significant dimensionality reduction of the feature set [6]. This opens the way to training



© Jan Jakubův and Josef Urban;

licensed under Creative Commons License CC-BY

10th International Conference on Interactive Theorem Proving (ITP 2019).

Editors: John Harrison, John O’Leary, and Andrew Tolmach; Article No. 34; pp. 34:1–34:8

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

learning-based internal guidance of saturation search even on very large ITP libraries, where the hundreds of thousands of handcrafted features would otherwise make the trained guiding systems impractically slow.

In this work we conduct the first practical evaluation of learning-based internal guidance of state-of-the-art saturation provers such as E in a realistic large-library hammer setting, with realistic time limits. The results turn out to be unexpectedly good, improving the real-time performance of E on the whole Mizar Mathematical Library (MML) [12] by 70% in a single-strategy setting. We believe that this is a breakthrough that will quickly lead to ubiquitous deployment of ATPs equipped with learning-based internal guidance in large-theory theorem proving and in hammer-style ITP assistance.

The rest of the paper is organized as follows. Section 2 summarizes the general saturation-style ATP setting and explains how machine learning can be trained and used over a large library of problems to guide the saturation search. Section 3 discusses the practical implementation of ENIGMA, i.e., the features, classifiers, and the feature hashing used to make the ENIGMA guidance both strong and efficient on a large library. Section 4 is our main contribution. We evaluate the latest ENIGMA on the whole Mizar Mathematical Library and show that in several iterations of proving and learning we can develop very strong strategies and solve in low time limits many previously unsolved problems.

## 2 Enhancing ATPs with Machine Learning

**Automated Theorem Proving.** State-of-the-art saturation-based automated theorem provers (ATPs) for first-order logic (FOL), such as E [25] and Vampire [22] are today’s most advanced tools for general reasoning across a variety of mathematical and scientific domains. Many ATPs employ the *given clause algorithm*, translating the input FOL problem  $T \cup \{-C\}$  into a refutationally equivalent set of clauses. The search for a contradiction is performed maintaining sets of *processed* ( $P$ ) and *unprocessed* ( $U$ ) clauses. The algorithm repeatedly selects a *given clause*  $g$  from  $U$ , moves  $g$  to  $P$ , and extends  $U$  with all clauses inferred with  $g$  and  $P$ . This process continues until a contradiction is found,  $U$  becomes empty, or a resource limit is reached. The search space of this loop grows quickly and it is a well-known fact that the selection of the right given clause is crucial for success. Machine learning from a large number of proofs and proof searches may help guide the selection of the given clauses.

E allows the user to select a *proof search strategy*  $\mathcal{S}$  to guide the proof search. An E strategy  $\mathcal{S}$  specifies parameters such as term ordering, literal selection function, clause splitting, paramodulation setting, premise selection, and, most importantly for us, the *given clause selection* mechanism. The given clause selection in E is implemented using a collection of *weight functions*. These weight functions are used in a round robin manner to select the given clause.

**Machine Learning of Given Clause Selection.** To facilitate machine learning research, E implements an option under which each successful proof search gets analyzed and the prover outputs a list of clauses annotated as either *positive* or *negative* training examples. Each processed clause which is present in the final proof is classified as positive. On the other hand, processing of clauses not present in the final proof was redundant, hence they are classified as negative. Our goal is to learn such classification (possibly conditioned on the problem and its features) in a way that generalizes and allows solving related problems.

Given a set of problems  $\mathcal{P}$ , we can run E with a strategy  $\mathcal{S}$  and obtain positive and negative training data  $\mathcal{T}$  from each of the successful proof searches. Various machine learning methods can be used to learn the clause classification given by  $\mathcal{T}$ , each method yielding a

*classifier* or *model*  $\mathcal{M}$ . In order to use the model  $\mathcal{M}$  in E,  $\mathcal{M}$  needs to provide the function to compute the weight of an arbitrary clause. This weight function is then used to guide future E runs.

**Guiding ATPs with Learned Models.** A model  $\mathcal{M}$  can be used in E in different ways. We use two methods to combine  $\mathcal{M}$  with a strategy  $\mathcal{S}$ . Either (1) we use  $\mathcal{M}$  to select *all* the given clauses, or (2) we combine  $\mathcal{M}$  with the given clause guidance from  $\mathcal{S}$  so that roughly half of the clauses are selected by  $\mathcal{M}$ . Proof search settings other than given clause guidance are inherited from  $\mathcal{S}$ . We denote the resulting E strategies as (1)  $\mathcal{S} \odot \mathcal{M}$ , and (2)  $\mathcal{S} \oplus \mathcal{M}$ .

### 3 ENIGMA: Inference Guiding Machine

**Machine Learning in Practice.** ENIGMA [15, 16] is our *efficient* learning-based method for guiding given clause selection in saturation-based ATPs, implementing the framework suggested in the previous Section 2. First-order clauses need to be represented in a format recognized by the selected learning method. While neural networks have been very recently practically used for internal guidance with ENIGMA [6], the strongest setting currently uses manually engineered *clause features* and fast non-neural state-of-the-art gradient boosted trees library [5].

**Clause Features.** Clause features represent a finite set of various syntactic properties of clauses, and are used to encode clauses by a fixed-length numeric vector. Various machine learning methods can handle numeric vectors and their success heavily depends on the selection of correct clause features. Various possible choices of efficient clause features for theorem prover guidance have been experimented with [15, 16, 20, 21]. The original ENIGMA [15] uses term-tree walks of length 3 as features, while the second version [16] reaches better results by employing various additional features.

Since there are only finitely many features in any training data, the features can be serially numbered. This numbering is fixed for each experiment. Let  $n$  be the number of different features appearing in the training data. A clause  $C$  is translated to a feature vector  $\varphi_C$  whose  $i$ -th member counts the number of occurrences of the  $i$ -th feature in  $C$ . Hence every clause is represented by a sparse numeric vector of length  $n$ . Additionally, we embed information about the conjecture currently being proved in the feature vector, yielding vectors of length  $2n$ . See [6, 16] for more details.

**From Logistic Regression to Decision Trees.** So far, the development of ENIGMA was focusing on fast and practically usable methods, allowing E users to directly benefit from our work. Simple but fast linear classifiers such as *linear SVM* and *logistic regression* efficiently implemented by the LIBLINEAR open source library [7] were used in our initial experiments [16]. Our recent experiments [6] report improved performance with *gradient boosted trees*, while maintaining efficiency. Gradient boosted trees are ensembles of decision trees trained by tree boosting. In particular, we use their implementation in the XGBoost library [5].

The model  $\mathcal{M}$  produced by XGBoost consists of a set (*ensemble* [24]) of decision trees. The inner nodes of the decision trees consist of conditions on feature values, while the leafs contain numeric scores. Given a vector  $\varphi_C$  representing a clause  $C$ , each tree in  $\mathcal{M}$  is navigated to the unique leaf using the values from  $\varphi_C$ , and the corresponding leaf scores are aggregated across all trees. The final score is translated to yield the probability that  $\varphi_C$

■ **Table 1** Number of Mizar problems solved in 10 seconds by various ENIGMA strategies.

	$\mathcal{S}$	$\mathcal{S} \odot \mathcal{M}_9^0$	$\mathcal{S} \oplus \mathcal{M}_9^0$	$\mathcal{S} \odot \mathcal{M}_9^1$	$\mathcal{S} \oplus \mathcal{M}_9^1$	$\mathcal{S} \odot \mathcal{M}_9^2$	$\mathcal{S} \oplus \mathcal{M}_9^2$	$\mathcal{S} \odot \mathcal{M}_9^3$	$\mathcal{S} \oplus \mathcal{M}_9^3$
solved	14933	16574	20366	21564	22839	22413	23467	22910	23753
$\mathcal{S}\%$	+0%	+10.5%	+35.8%	+43.8%	+52.3%	+49.4%	+56.5%	+52.8%	+58.4
$\mathcal{S}+$	+0	+4364	+6215	+7774	+8414	+8407	+8964	+8822	+9274
$\mathcal{S}-$	-0	-2723	-782	-1143	-508	-927	-430	-845	-454

  

	$\mathcal{S} \odot \mathcal{M}_{12}^3$	$\mathcal{S} \oplus \mathcal{M}_{12}^3$	$\mathcal{S} \odot \mathcal{M}_{16}^3$	$\mathcal{S} \oplus \mathcal{M}_{16}^3$
solved	24159	24701	25100	25397
$\mathcal{S}\%$	+61.1%	+64.8%	+68.0%	+70.0%
$\mathcal{S}+$	+9761	+10063	+10476	+10647
$\mathcal{S}-$	-535	-295	-309	-183

■ **Table 2** Comparison of several developed strategies in higher time limits.

	$\mathcal{S}$ (30s)	$\mathcal{S} \oplus \mathcal{M}_9^2$ (30s)	$\mathcal{S} \oplus \mathcal{M}_9^2$ (60s)	$\mathcal{S} \oplus \mathcal{M}_9^3$ (60s)	$\mathcal{S} \oplus \mathcal{M}_{12}^3$ (60s)	$\mathcal{S} \oplus \mathcal{M}_{16}^3$ (60s)
solved	15554	24154	24495	24762	25540	26107
hard	75	891	956	1017	1192	1296

represents a positive clause. When using  $\mathcal{M}$  as a weight function in E, the probabilities are turned into binary classification, assigning weight 1.0 for probabilities  $\geq 0.5$  and weight 10.0 otherwise. Our experiments with scaling of the weight by the probability did not yet yield improved functionality.

**Feature Hashing.** The vectors representing clauses have so far had length  $n$  when  $n$  is the total number of features in the training data  $\mathcal{T}$  (or  $2n$  with conjecture features). Experiments revealed that XGBoost is capable of dealing with vectors up to the length of  $10^5$  with a reasonable performance. This might be enough for smaller benchmarks but with the need to train on bigger training data, we might need to handle much larger feature sets. In experiments with the whole translated Mizar Mathematical Library, the feature vector length can easily grow over  $10^6$ . This significantly increases both the training and the clause evaluation times. To handle such larger data sets, we have implemented a simple *hashing* method to decrease the dimension of the vectors.

Instead of serially numbering all features, we represent each feature  $f$  by a unique string and apply a general-purpose string hashing function<sup>1</sup> to obtain a number  $n_f$  within a required range (between 0 and an adjustable *hash base*). The value of  $f$  is then stored in the feature vector at the position  $n_f$ . If different features get mapped to the same vector index, the corresponding values are summed up. See [6] for more details.

## 4 Experiments

The experiments are done on a large benchmark of 57880 Mizar40 [19] problems<sup>2</sup> from the MPTP dataset [27]. Since we are here interested in internal guidance rather than in premise selection, we have used the small (*bushy*, re-proving) versions of the problems, however

<sup>1</sup> We use the following hashing function *sdbm*:  $h_i = s_i + (h_{i-1} \ll 6) + (h_{i-1} \ll 16) - h_{i-1}$ .

<sup>2</sup> [http://grid01.ciirc.cvut.cz/~mptp/7.13.01\\_4.181.1147/MPTP2/problems\\_small\\_consist.tar.gz](http://grid01.ciirc.cvut.cz/~mptp/7.13.01_4.181.1147/MPTP2/problems_small_consist.tar.gz)

■ **Table 3** Training statistics and inference speed for different tree depths.

Tree depth	training error	real time	CPU time	model size (MB)	inference speed
9	0.201	2h41m	4d20h	5.0	5665.6
12	0.161	4h12m	8d10h	17.4	4676.9
16	0.123	6h28m	11d18h	54.7	3936.4

■ **Table 4** Effect of looping on 10k randomly selected problems.

	$\mathcal{S}$	$\mathcal{S} \oplus \mathcal{M}^0$	$\mathcal{S} \oplus \mathcal{M}^1$	$\mathcal{S} \oplus \mathcal{M}^2$	$\mathcal{S} \oplus \mathcal{M}^3$	$\mathcal{S} \oplus \mathcal{M}^4$	$\mathcal{S} \oplus \mathcal{M}^5$	$\mathcal{S} \oplus \mathcal{M}^6$
solved	2487	3204	3625	3755	3838	3854	3892	3944
$\mathcal{S}\%$	+0%	+28.8%	+45.7%	+50.9%	+54.3%	+54.9%	+56.4%	+58.5%

without previous ATP minimization. We start with a good evolutionarily optimized [14] E strategy  $\mathcal{S}$  that performed best in previous experiments on the smaller MPTP2078 dataset. We run  $\mathcal{S}$  for 10s on the whole library, producing the first proofs, we learn from them the next guiding strategy, and this is iterated with the growing body of proofs. All problems are run on the same hardware<sup>3</sup> and with the same memory limits employing multiple cores (around 300) for massive parallel evaluation.

Table 1 shows the number of Mizar problems solved in 10 seconds by the baseline strategy  $\mathcal{S}$  and by each iteration of learning and proving with the learned guidance. The model  $\mathcal{M}_9^0$  is trained on the training data coming from the problems solved by  $\mathcal{S}$  with the maximum depth of XGBoost decision trees set to 9. We further *loop* this process and models  $\mathcal{M}_9^n$  are trained on all the problems solved by  $\mathcal{S}$ , and by all the previous  $\mathcal{S} \ominus \mathcal{M}_9^k$  and  $\mathcal{S} \oplus \mathcal{M}_9^k$  for  $k < n$ . Models  $\mathcal{M}_{12}^3$  and  $\mathcal{M}_{16}^3$  are trained on the same data as  $\mathcal{M}_9^3$  but with the tree depth increased to 12 and 16. XGBoost models contain 200 decision trees and the hash base is set to  $2^{15}$ . In the row  $\mathcal{S}\%$  we show the percentage gain over the baseline strategy  $\mathcal{S}$ , while  $\mathcal{S}+$  and  $\mathcal{S}-$  are the additions and missing solutions w.r.t.  $\mathcal{S}$ . We can see that new problems are added with every iteration of looping. Combined versions ( $\oplus$ ) typically perform better and lose less solutions. Increasing the tree depth to 16 leads to a strategy that outperforms the baseline by rather astonishing 70%.

Table 2 compares several of our new strategies with higher time limits and also shows the number of solved *hard* problems, i.e., the problems unsolved by any method developed previously in [19]. Our best strategy  $\mathcal{S} \oplus \mathcal{M}_{16}^3$  solves 26107 problems in 60s. Note that the 60s portfolio of our six best previous evolutionarily developed strategies for Mizar (i.e., each run for 10s) solves only 22068 problems, i.e., the single new strategy is 18.3% better. Vampire in the CASC (best portfolio) mode run in 300s has solved 27842 of these problems in 300s in [19].

Table 3 shows the training times, model sizes and inference speeds of XGBoost in the 4th iteration of proving and learning, using different tree depths. The training data is a sparse matrix with 65536 ( $= 2 * 2^{15}$ ) columns (features) consisting of 63M examples. The total number of non-empty entries in the matrix is 5B (40GB). The inference speed is the average of the generated clauses per second measured on problems that timed out in all three runs. Note that despite the decrease of the inference speed with the more complicated XGBoost models, their accuracy and real-time performance grows (cf. Table 2). Training of better models on the millions of proof search examples however already requires significant resources – almost 12 CPU days for the best model with tree depth 16.

<sup>3</sup> Intel(R) Xeon(R) CPU E5-2698 v3 @ 2.30GHz with 256G RAM.

Table 4 presents additional shorter experiments with more looping performed on a randomly selected 10k problems. The tree depth is set to 9. Again, the model  $\mathcal{M}^0$  is trained only on the problems solved by  $\mathcal{S}$  and the next models are obtained by looping. The highest improvement is achieved after the first learning ( $\mathcal{M}^0$ ), however, the next iterations continue to add improvements.

## 5 Conclusion and Future Work

We have taken a good previously tuned E strategy and turned it into a learning-guided strategy that is 70% stronger in real time. We have done that by several iterations of MaLAREa-style [28] feedback loop between proving and learning over a large mathematical library. The iterations here are however not done for learning premise selection as in MaLAREa, but for learning efficient internal guidance. While developing this kind of efficient internal guidance for state-of-the-art saturation ATPs has been challenging and took time, the very large gains obtained here show that this has been very well invested effort. Future work will certainly focus on even stronger learning methods and also on more dynamic proof state characterization such as ENIGMAWatch [11]. It is however clear that this is the point when machine learning guidance has very strongly overtaken the human development of ATP strategies over large problem corpora.

---

### References

- 1 Jesse Alama, Tom Heskes, Daniel Kühlwein, Evgeni Tsivtsivadze, and Josef Urban. Premise Selection for Mathematics by Corpus Analysis and Kernel Methods. *J. Autom. Reasoning*, 52(2):191–213, 2014. doi:10.1007/s10817-013-9286-5.
- 2 Alexander A. Alemi, François Chollet, Niklas Eén, Geoffrey Irving, Christian Szegedy, and Josef Urban. DeepMath - deep sequence models for premise selection. In Daniel D. Lee, Masashi Sugiyama, Ulrike V. Luxburg, Isabelle Guyon, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 2235–2243, 2016. URL: <http://papers.nips.cc/paper/6280-deepmath-deep-sequence-models-for-premise-selection>.
- 3 Jasmin Christian Blanchette, David Greenaway, Cezary Kaliszyk, Daniel Kühlwein, and Josef Urban. A Learning-Based Fact Selector for Isabelle/HOL. *J. Autom. Reasoning*, 57(3):219–244, 2016. doi:10.1007/s10817-016-9362-8.
- 4 Jasmin Christian Blanchette, Cezary Kaliszyk, Lawrence C. Paulson, and Josef Urban. Hammering towards QED. *J. Formalized Reasoning*, 9(1):101–148, 2016. doi:10.6092/issn.1972-5787/4593.
- 5 Tianqi Chen and Carlos Guestrin. XGBoost: A Scalable Tree Boosting System. In *KDD*, pages 785–794. ACM, 2016.
- 6 Karel Chvalovský, Jan Jakubuv, Martin Suda, and Josef Urban. ENIGMA-NG: efficient neural and gradient-boosted inference guidance for E. *CoRR*, abs/1903.03182, 2019. arXiv:1903.03182.
- 7 Rong-En Fan, Kai-Wei Chang, Cho-Jui Hsieh, Xiang-Rui Wang, and Chih-Jen Lin. LIBLINEAR: A library for large linear classification. *Journal of Machine Learning Research*, 9:1871–1874, 2008. doi:10.1145/1390681.1442794.
- 8 Michael Färber and Cezary Kaliszyk. Random Forests for Premise Selection. In Carsten Lutz and Silvio Ranise, editors, *Frontiers of Combining Systems - 10th International Symposium, FroCoS 2015, Wrocław, Poland, September 21-24, 2015. Proceedings*, volume 9322 of *Lecture Notes in Computer Science*, pages 325–340. Springer, 2015. doi:10.1007/978-3-319-24246-0\_20.



- 9 Michael Färber, Cezary Kaliszyk, and Josef Urban. Monte Carlo tableau proof search. In Leonardo de Moura, editor, *Automated Deduction - CADE 26 - 26th International Conference on Automated Deduction, Gothenburg, Sweden, August 6-11, 2017, Proceedings*, volume 10395 of *Lecture Notes in Computer Science*, pages 563–579. Springer, 2017. doi:10.1007/978-3-319-63046-5\_34.
- 10 Thibault Gauthier, Cezary Kaliszyk, and Josef Urban. TacticToe: Learning to reason with HOL4 tactics. In Thomas Eiter and David Sands, editors, *LPAR-21, 21st International Conference on Logic for Programming, Artificial Intelligence and Reasoning, Maun, Botswana, May 7-12, 2017*, volume 46 of *EPiC Series in Computing*, pages 125–143. EasyChair, 2017. URL: <http://www.easychair.org/publications/paper/340355>.
- 11 Zarathustra Goertzel, Jan Jakubův, and Josef Urban. ProofWatch meets ENIGMA: First experiments. In Gilles Barthe, Konstantin Korovin, Stephan Schulz, Martin Suda, Geoff Sutcliffe, and Margus Veanes, editors, *LPAR-22 Workshop and Short Paper Proceedings*, volume 9 of *Kalpa Publications in Computing*, pages 15–22. EasyChair, 2018. doi:10.29007/z7qx.
- 12 Adam Grabowski, Artur Korniłowicz, and Adam Naumowicz. Mizar in a nutshell. *J. Formalized Reasoning*, 3(2):153–245, 2010.
- 13 Thomas Gransden, Neil Walkinshaw, and Rajeev Raman. SEPIA: search for proofs using inferred automata. In *Automated Deduction - CADE-25 - 25th International Conference on Automated Deduction, Berlin, Germany, August 1-7, 2015, Proceedings*, pages 246–255, 2015. doi:10.1007/978-3-319-21401-6\_16.
- 14 Jan Jakubův and Josef Urban. BliStrTune: hierarchical invention of theorem proving strategies. In Yves Bertot and Viktor Vafeiadis, editors, *Proceedings of the 6th ACM SIGPLAN Conference on Certified Programs and Proofs, CPP 2017, Paris, France, January 16-17, 2017*, pages 43–52. ACM, 2017. doi:10.1145/3018610.3018619.
- 15 Jan Jakubův and Josef Urban. ENIGMA: efficient learning-based inference guiding machine. In Herman Geuvers, Matthew England, Osman Hasan, Florian Rabe, and Olaf Teschke, editors, *Intelligent Computer Mathematics - 10th International Conference, CICM 2017, Edinburgh, UK, July 17-21, 2017, Proceedings*, volume 10383 of *Lecture Notes in Computer Science*, pages 292–302. Springer, 2017. doi:10.1007/978-3-319-62075-6\_20.
- 16 Jan Jakubův and Josef Urban. Enhancing ENIGMA Given Clause Guidance. In Florian Rabe, William M. Farmer, Grant O. Passmore, and Abdou Youssef, editors, *Intelligent Computer Mathematics - 11th International Conference, CICM 2018, Hagenberg, Austria, August 13-17, 2018, Proceedings*, volume 11006 of *Lecture Notes in Computer Science*, pages 118–124. Springer, 2018. doi:10.1007/978-3-319-96812-4\_11.
- 17 Cezary Kaliszyk and Josef Urban. Learning-Assisted Automated Reasoning with Flyspeck. *J. Autom. Reasoning*, 53(2):173–213, 2014. doi:10.1007/s10817-014-9303-3.
- 18 Cezary Kaliszyk and Josef Urban. FEMaLeCoP: Fairly efficient machine learning connection prover. In Martin Davis, Ansgar Fehnker, Annabelle McIver, and Andrei Voronkov, editors, *Logic for Programming, Artificial Intelligence, and Reasoning - 20th International Conference, LPAR-20 2015, Suva, Fiji, November 24-28, 2015, Proceedings*, volume 9450 of *Lecture Notes in Computer Science*, pages 88–96. Springer, 2015. doi:10.1007/978-3-662-48899-7\_7.
- 19 Cezary Kaliszyk and Josef Urban. MizAR 40 for Mizar 40. *J. Autom. Reasoning*, 55(3):245–256, 2015. doi:10.1007/s10817-015-9330-8.
- 20 Cezary Kaliszyk, Josef Urban, Henryk Michalewski, and Miroslav Olsák. Reinforcement Learning of Theorem Proving. In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, 3-8 December 2018, Montréal, Canada.*, pages 8836–8847, 2018. URL: <http://papers.nips.cc/paper/8098-reinforcement-learning-of-theorem-proving>.

- 21 Cezary Kaliszyk, Josef Urban, and Jirí Vyskočil. Efficient Semantic Features for Automated Reasoning over Large Theories. In *IJCAI*, pages 3084–3090. AAAI Press, 2015.
- 22 Laura Kovács and Andrei Voronkov. First-Order Theorem Proving and Vampire. In Natasha Sharygina and Helmut Veith, editors, *CAV*, volume 8044 of *LNCS*, pages 1–35. Springer, 2013. doi:10.1007/978-3-642-39799-8\_1.
- 23 Sarah M. Loos, Geoffrey Irving, Christian Szegedy, and Cezary Kaliszyk. Deep Network Guided Proof Search. In Thomas Eiter and David Sands, editors, *LPAR-21, 21st International Conference on Logic for Programming, Artificial Intelligence and Reasoning, Maun, Botswana, May 7-12, 2017*, volume 46 of *EPiC Series in Computing*, pages 85–105. EasyChair, 2017. URL: <http://www.easychair.org/publications/paper/340345>.
- 24 Robi Polikar. Ensemble based systems in decision making. *Circuits and Systems Magazine, IEEE*, 6(3):21–45, 2006.
- 25 Stephan Schulz. E - A Brainiac Theorem Prover. *AI Commun.*, 15(2-3):111–126, 2002. URL: <http://iospress.metapress.com/content/n908n94nmvk59v3c/>.
- 26 Stephan Schulz. System Description: E 1.8. In Kenneth L. McMillan, Aart Middeldorp, and Andrei Voronkov, editors, *LPAR*, volume 8312 of *LNCS*, pages 735–743. Springer, 2013. doi:10.1007/978-3-642-45221-5\_49.
- 27 Josef Urban. MPTP 0.2: Design, implementation, and initial experiments. *J. Autom. Reasoning*, 37(1-2):21–43, 2006. doi:10.1007/s10817-006-9032-3.
- 28 Josef Urban, Geoff Sutcliffe, Petr Pudlák, and Jirí Vyskočil. MaLAREa SG1 - Machine Learner for Automated Reasoning with Semantic Guidance. In *IJCAR*, pages 441–456, 2008. doi:10.1007/978-3-540-71070-7\_37.
- 29 Josef Urban, Jirí Vyskočil, and Petr Štěpánek. MaLeCoP: Machine learning connection prover. In Kai Brunnler and George Metcalfe, editors, *TABLEAUX*, volume 6793 of *LNCS*, pages 263–277. Springer, 2011. doi:10.1007/978-3-642-22119-4\_21.