



# ENIGMAWatch: ProofWatch Meets ENIGMA

Zarathustra Goertzel, Jan Jakubův, and Josef Urban<sup>(✉)</sup>

Czech Technical University in Prague, Prague, Czech Republic  
josef.urban@gmail.com

**Abstract.** In this work we describe a new learning-based proof guidance – ENIGMAWatch – for saturation-style first-order theorem provers. ENIGMAWatch combines two guiding approaches for the given-clause selection implemented for the E ATP system: ProofWatch and ENIGMA. ProofWatch is motivated by the watchlist (hints) method and based on symbolic matching of multiple related proofs, while ENIGMA is based on statistical machine learning. The two methods are combined by using the evolving information about symbolic proof matching as additional characterization of the saturation-style proof search for the statistical learning methods. The new system is evaluated on a large set of problems from the Mizar library. We show that the added proof-matching information is considered important by the statistical machine learners, and that it leads to improved performance over ProofWatch and ENIGMA.

## 1 Introduction

This work describes a new learning-based proof guidance – *ENIGMAWatch* – for saturation-style first-order theorem provers. ENIGMAWatch<sup>1</sup> is the combination of two previous guidance methods implemented for the E theorem prover [35]: ProofWatch [11] and ENIGMA [16,17]. Both ProofWatch and ENIGMA learn to guide E’s proof search for a new conjecture based on related proofs.

ProofWatch uses the hints (watchlist) mechanism, which is a form of precise symbolic memory that can allow inference chains done in a former proof to be replayed in the current proof search. It uses standard symbolic subsumption to check which clauses subsume clauses in related proofs. In addition to boosting the priority of these clauses, the completion ratios of the related proofs are computed, and the proof search is biased towards the most completed ones.

ENIGMA uses fast statistical machine learning to learn from related proof-searches to identify good and bad (positive and negative) clauses for the current

<sup>1</sup> The E version used in this paper can be found at <https://github.com/ai4reason/e prover/tree/devel>, and the library for running ENIGMA with E can be found at <https://github.com/ai4reason/enigma>.

---

J. Urban—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.

conjecture. ENIGMA chooses the given clauses based only on features of the problem’s conjecture, which is static throughout the whole proof search. This seems suboptimal: as the proof search evolves, information about the work done so far should influence the selection of the next given clauses.

ENIGMAWatch combines the two approaches by giving the ENIGMA’s learner the ProofWatch completion ratios of the related proofs as an evolving vectorial characterization of the current proof search state. This allows E’s machine learning guidance to have more information about how the proof search is unfolding.

An early version of ENIGMAWatch was tested on the MPTP Challenge<sup>2</sup> [36,39] benchmark. It contains 252 first-order problems extracted from the Mizar Mathematical Library (MML) [14], used in Mizar to prove the Bolzano-Weierstrass theorem. Initially, ENIGMAWatch could not be run on a larger dataset, such as the 57897 Mizar40 [21] benchmark, in a reasonable time. Since then, ENIGMA implemented dimensionality reduction using feature hashing [6], extending its applicability to large corpora. We have additionally improved watchlist mechanism in E through enhanced indexing, first time presented in this work in Sect. 4. This allows also ENIGMAWatch to be applied to larger corpora.

The rest of the paper is organized as follows. Section 2 provides an introduction to saturation-based theorem proving and briefly describes ENIGMA and ProofWatch. Section 3 explains how ENIGMA and ProofWatch are combined into ENIGMAWatch, and how watchlists can be selected. Section 4 describes our improved watchlist indexing in E. Both ENIGMAWatch and the improved watchlist indexing are evaluated in Sect. 5.

## 2 Guiding the Given Clause Selection in ATPs

### 2.1 Automated Theorem Proving and Machine Learning

State-of-the-art saturation-based automated theorem provers (ATPs) for first-order logic (FOL), such as E [33] and Vampire [25] 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 *proof state*  $\Pi$ ). 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 [1–4, 7–10, 15, 16, 19, 20, 22, 26, 29, 31, 32, 38, 40, 41] may help guide the selection of the given clauses.

<sup>2</sup> <http://tptp.cs.miami.edu/~tptp/MPTPChallenge/>.

## 2.2 ENIGMA: Learning from Successful Proof Searches

ENIGMA [6, 16–18] (*Efficient learning-based Internal Guidance MACHine*) is our method for guiding given clause selection in saturation-based ATPs. The method needs to be efficient because it is internally applied to every generated clause. ENIGMA uses E’s capability to analyze successful proof searches, and to output lists of given 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. ENIGMA’s goal is to learn such classification (possibly conditioned on the problem and its features) in a way that generalizes and allows solving new related problems.

**ENIGMA Learning and Models.** 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 a (classification) *model*  $\mathcal{M}$ . In order to use the model  $\mathcal{M}$  in E,  $\mathcal{M}$  is used as a function that computes clause weights. This weight function is then used to guide future E runs.

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 libraries such as XGBoost [5]. The model  $\mathcal{M}$  produced by XGBoost consists of a set (*ensemble* [30]) of decision trees. Given a clause  $C$ , the model  $\mathcal{M}$  yields the probability that  $C$  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.

**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 [16, 17, 22, 23]. The original ENIGMA [16] uses term-tree walks of length 3 as features, while the second version [17] 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, 17] for more details.

**Feature Hashing.** Experiments revealed that XGBoost is capable of dealing with vectors up to the length of  $10^5$  with a reasonable performance. 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, a simple *hashing* method has previously been implemented 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 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.

### 2.3 ProofWatch: Proof Guidance by Clause Subsumption

In this section we explain the ProofWatch guiding mechanisms. Unlike the statistical approach in ENIGMA, ProofWatch implements a form of *symbolic* memory and guidance. It produces a notion of *proof-state vector* that is dynamically created and updated.

**Standard Watchlist Guidance.** The watchlist (hint list) mechanism itself does not perform any statistical machine learning. It steers given clause selection via symbolic matching between generated clauses and a set of clauses called a *watchlist*. This technique has been originally developed by Veroff [42] and implemented in Otter [27] and Prover9 [28]. Since then, it has been extensively used in the AIM project [24] for obtaining long and advanced proofs of open algebraic conjectures. The watchlist mechanism is nowadays implemented also in E. All the above implementations use only a single watchlist, as opposed to ProofWatch discussed below.

Recall that a clause  $C$  *subsumes* a clause  $D$ , written  $C \sqsubseteq D$ , when there exists a substitution  $\sigma$  such that  $C\sigma \subseteq D$  (where clauses are considered to be sets of literals). The watchlist guidance then works as follows. Every generated clause  $C$  is checked for subsumption with every watchlist clause  $D \in W$ . When  $C$  subsumes at least one of the watchlist clauses, then  $C$  is considered important for the proof search and is processed with high priority. The idea behind this is that the watchlist  $W$  contains clauses which were processed during a previous successful proof search of a related conjecture. Hence processing of similar clauses may lead to success again.

In E, the watchlist mechanism is implemented using a priority function<sup>3</sup> which takes precedence over the weight function used to select the next given clause. Priority functions assign the priority to each clause, and clauses with higher priority are selected as given before clauses with lower priority<sup>4</sup>.

<sup>3</sup> See the priority function `PreferWatchlist` in the E manual.

<sup>4</sup> Numerically the lower the priority, the better. Hence 0 is the best priority.

When clauses from previous proofs are put on a watchlist, E thus prefers to follow steps from the previous proofs whenever it can.

**ProofWatch.** Our approach [11, Sect. 5] extends standard watchlist guidance by allowing for multiple watchlists  $W_1, \dots, W_n$ , for example, one corresponding to each related proof found before. We say that a generated clause  $C$  *matches* the watchlist  $W_i$ , written  $C \sqsubseteq W_i$ , iff  $C$  subsumes some clause  $D \in W_i$  ( $C \sqsubseteq D$ ). Similarly, the above watchlist clause  $D$  is said to be *matched* by  $C$ .

The reason to include multiple watchlists is that during a proof search, clauses from some watchlists might get matched more often than clauses from others. The more clauses are matched from some watchlist  $W_i$ , the more the current proof search resembles  $W_i$ , and hence  $W_i$  might be more relevant for this proof search. Thus the idea of ProofWatch is to prioritize clauses that match more relevant watchlists (proofs).

Watchlist *relevance* is dynamically computed as follows. We define  $progress(W_i)$  to be the count of clauses from  $W_i$  that have been matched in the proof search thus far. The *completion ratio*,  $c_i = \frac{progress(W_i)}{|W_i|}$ , measures how much of the watchlist  $W_i$  has been matched. The *dynamic relevance* of each generated clause  $C$  is defined as the maximum completion ratio over all the watchlists  $W_i$  that  $C$  matches:

$$relevance(C) = \max_{W \in \{W_i: C \sqsubseteq W_i\}} \left( \frac{progress(W)}{|W|} \right)$$

The higher the dynamic relevance  $relevance(C)$ , the higher the priority of  $C$ . The dynamic watchlist mechanism is implemented using the E priority function.<sup>5</sup> The results of experiments in [11, Sect. 6.3] on the same dataset as this work (Mizar40 [21]) indicate that dynamic relevance improves performance over an ensemble of strategies, whereas the single watchlist approach is stronger on each individual strategy.

When using a large problem library such as Mizar40, it is practically useful to choose only some proofs for watchlists. First, E's speed decreases with each additional proof on the watchlist, so if working on a large dataset, loading all available proofs as watchlists will lead to a large slowdown (cf. Sect. 4). Second, it's not guaranteed that all proofs will help E with proving the problem at hand.

### 3 ENIGMAWatch: ProofWatch Meets ENIGMA

#### 3.1 Completion Ratios as Semantic Embeddings of the Proof Search

The watchlist completion ratios ( $c_0, \dots, c_N$ ) ( $N$  ranges over the watchlist proofs) at each step in E's proof search can be taken as a vectorial representation of the current proof state  $\Pi$ . The general motivation for this approach is to come up with an *evolving* characterization of the saturation-style proof state  $\Pi$ , preferably in a vectorial form  $\varphi_\Pi$  suitable for machine learning tools, such as ENIGMA.

<sup>5</sup> See `PreferWatchlistRelevant` in [11].

Recall that the proof state  $\Pi$  is a set of processed clauses  $P$  and unprocessed clauses  $U$ . The vector of watchlist completion ratios thus maintains a running tally of where clauses in  $P \cup U$  match the different related proofs. In general, this could be replaced, e.g., by a vector of more abstract similarities of the current proof state to other proofs measured in various (possibly approximate) ways. In ENIGMAWatch we use the ProofWatch based *proof-state vector* for a proof state  $\Pi$  defined by the completion ratios, i.e.,  $\varphi_{\Pi} = (c_0, \dots, c_N)$ . This is the first practical implementation of the general idea: using *semantic embeddings* (i.e., representations in  $R^n$ ) of the proof state  $\Pi$  for guiding statistical learning methods. ENIGMAWatch uses the proof-state vectors  $\varphi_{\Pi}$  as follows. The positive  $\mathcal{C}^+$  and negative  $\mathcal{C}^-$  given clauses are output along with  $\varphi_{\Pi}$ , the proof-state vector at the time of their selection, and used as added features of the proof state when training ENIGMA-style classifiers.

**Table 1.** Example of the proof-state vector for 8 (of 32) (serially numbered) proofs loaded to guide the proof of **YELLOW\_5:36**. The three columns are the watchlist  $i$ , the completion ratio of  $i$ , and  $progress(W_i)/|W_i|$ .

0	0.438	42/96	1	0.727	56/77	2	0.865	45/52	3	0.360	9/25
4	0.750	51/68	5	0.259	7/27	6	0.805	62/77	7	0.302	73/242

Table 1 shows a sample proof-state vector based on 32 related proofs<sup>6</sup> for the Mizar theorem **YELLOW 5:36**<sup>7</sup> (De Morgan’s law<sup>8</sup>) at the end of the proof search. Note that some related proofs, such as #2, were almost fully matched, while others, such as #7 were mostly not matched in the proof search.

### 3.2 Proof Vector Construction

**Data Construction.** In the ProofWatch [11] experiments, the best method for selecting related proofs (watchlists) was to use k-nearest neighbor (k-NN) to recommend 32 proofs per problem. The watchlists there are thus problem specific. In ENIGMAWatch, we want the watchlists to be globally fixed across the whole library, so that the proof completion ratios have the same meaning in all proofs. To construct the proof vectors, we first use a strong E strategy to produce a set of initial proofs (14882 over the 57897 Mizar40 problems). Then we run E with ProofWatch and the same strategy over the full 57897 problems with the 14882 proofs loaded into the watchlist. The time limit for both runs was *T60-G10000*, which means that E stops after 60s or 10000 generated clauses. This data provides information on how often each watchlist was encountered in each successful proof search. The training data then

<sup>6</sup> The proofs were chosen via k-NN. See [11, Sec. 6.1] for details.

<sup>7</sup> [http://grid01.ciirc.cvut.cz/~mptp/7.13.01\\_4.181.1147/html/yellow\\_5#T36](http://grid01.ciirc.cvut.cz/~mptp/7.13.01_4.181.1147/html/yellow_5#T36).

<sup>8</sup>  $\neg(P \vee Q) \iff (\neg P) \wedge (\neg Q)$ .

consists of a proof vector for each given clause (for each conjecture/problem): (*conjecture, given-clause, proof-state vector*).

**Dimensionality Reduction.** Next, we experiment with various pre-processing methods to reduce the *proof-state vector* dimension and thus decrease the number of watchlists loaded in E. For each problem we compute the mean of proof-state vectors over all given clauses  $g$ :  $\frac{1}{\#g} \sum_g \varphi_{\Pi_g}$ . This vector consists of the averaged completion ratios for each watchlist, which will be higher if the watchlist was matched earlier in the proof. This results in the mean proof-state matrix  $M$  consisting of row vectors (*mean-proof-vector*) (one for each conjecture/problem).

The following are methods experimented with in this paper for constructing the globally fixed vector of 512 watchlists from matrix  $M$ :

- *Mean*: compute the mean of  $M$  across the rows to obtain a mean proof-state vector that contains for each watchlist its average use across all problems. Then we take the top 512 watchlists.
- *Corr*: compute the Pearson correlation matrix<sup>9</sup> based on (the transpose of)  $M$ , and find a relatively uncorrelated set of 512 watchlists.
- *Var*: compute the variance (across the rows) of each column in  $M$ , and take the 512 watchlists with the highest variance. The intuition is that watchlists whose completion ratio vary more over the problem corpus may be more useful for learning.
- *Rand*: randomly select 512 watchlists.

## 4 Multi-indices Subsumption Indexing

In order to determine whether a generated clause matches a watchlist, the generated clause must be checked for subsumption with every watchlist clause. A major limitation of previous work [11, 12] was the slowdown of E as the watchlist size increased beyond 4000 clauses. Including more than 128 proofs was impractical. This section describes a method we have developed to speed up watchlist matching.

E already implements feature vector indexing [34] used also for the purpose of watchlist matching. The watchlist clauses are inserted into an indexing data structure and various properties of clauses are used to prune possible subsumption candidates. In this way, the number of possibly expensive subsumption calls is reduced. We build upon this, and further limit the number of required subsumption checks by using multiple indices instead of a single index.<sup>10</sup>

We take advantage of the fact that a clause  $C$  cannot subsume a clause  $D$  if the top-level predicate symbols do not match. In particular,  $C \sqsubseteq D$  can only hold if all the predicate symbols from  $C$  also appear in  $D$ , because substitution can neither introduce nor remove predicate symbols from a clause.

<sup>9</sup> <https://docs.scipy.org/doc/numpy/reference/generated/numpy.corrcoef.html>.

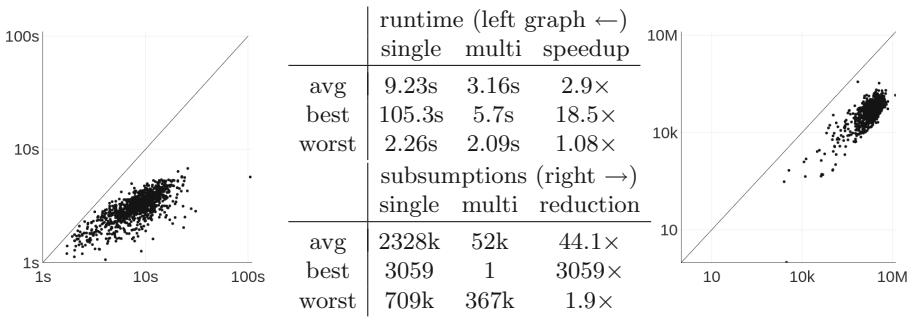
<sup>10</sup> Even with multiple watchlists, all the watchlist clauses are inserted into a single index, and only the name of the original watchlist is additionally stored.

We define the *code* of a clause  $C$ , denoted  $\text{code}(C)$ , as the set of predicate symbols with their logical signs (either  $+$  for positive predicates, or  $-$  for negated ones). For example, the code of the clause “ $P(a) \vee \neg P(b) \vee P(f(x))$ ” is the set  $\{+P, -P\}$ . The following holds because codes are preserved under substitution.

**Lemma 1.** *Given clauses  $C$  and  $D$ ,  $C \sqsubseteq D$  implies  $\text{code}(C) \subseteq \text{code}(D)$ .*

We create a separate index for every different clause code. Each watchlist clause  $D$  is inserted only to the index corresponding to  $\text{code}(D)$ . In order to check whether some clause  $C$  matches a watchlist, we only need to search in the indices whose codes are supersets of (or equal to)  $\text{code}(C)$ . Each index is implemented using E’s native feature vector indexing structure. Evaluation of this simple indexing method is provided in Sect. 5.1.

**Table 2.** Evaluation of multi-indices subsumption indexing.



## 5 Experiments

This section describes the experimental evaluation<sup>11</sup> of

1. the improved watchlist mechanism from Sect. 4
2. the watchlist selection for ENIGMAWatch from Sect. 3

### 5.1 Multi-indices Subsumption Indexing Evaluation

We propose a simple experiment to evaluate our implementation of multi-indices subsumption indexing from Sect. 4. We take a random sample of 1000 problems from the Mizar40 [21] data set and create a watchlist with around 60 k clauses coming from proofs of problems similar to the sample problems. We then run E

<sup>11</sup> Experiments code and data are available at <https://github.com/ai4reason/eprover-data/tree/master/TABLEAUX-19>

All experiments are run on the same hardware: Intel(R) Xeon(R) Gold 6140 CPU @ 2.30 GHz with 188 GB RAM.



on the sample problems with a fixed limit of 1000 generated clauses. This gives us a measure of how fast the single-index and multi-indices versions are, that is, how fast they can generate the first 1000 clauses. As the watchlist indexing does not influence the proof search, both versions process the same clauses and output the same result. Each generated clause has to be checked for watchlist subsumption and hence the limit on generated clauses is also the limit on different watchlist checks. We expect the number of clause-to-clause subsumption checks to decrease with multi-indices, as the method prunes possible subsumption candidates.

The results of the experiments are presented in Table 2. For each problem, we measure the runtime (left graph) and the number of different clause subsumption calls (right graph). The suffix “s” stands for seconds, “k” stands for thousands, and “M” stands for millions. Although subsumption is also used for purposes other than watchlist matching, we should be able to observe a decrease in the number of calls. Each point in the graphs corresponds to one sample problem, and is drawn at the position  $(x, y)$  corresponding to the results of single-index  $(x)$  and multi-indices  $(y)$  versions. Hence points below the diagonal signify an improvement. Also note logarithmic axes. The table shows the average improvement, and also the best and the worst cases. From the results, we can see that an average speed-up is almost 3 times. Furthermore, the average reduction of subsumption calls is more than 44 times and the number is reduced even in the worst case.

**Table 3.** ProofWatch evaluation: Problems solved by different versions.

Baseline	Mean	Var	Corr	Rand	Baseline $\cup$ Mean	Total
1140	1357	1345	1337	1352	1416	1483

**Table 4.** ENIGMAWatch evaluation: Problems solved and the effect of looping.

loop	ENIGMA	Mean	Var	Corr	Rand	ENIGMA $\cup$ Mean	Total
0	1557	1694	1674	1665	1690	1830	1974
1	1776	1815	1812	1812	1847	1983	2131
2	1871	1902	1912	1882	1915	2058	2200
3	1931	1954	1946	1920	1926	2110	2227

The number of watchlist clauses in the experiments was 61501, and the multi-indices version used 11442 different indices. This means that there were less than 6 clauses per index in average, although the count of clauses in different indices varied from 1 to 3837. The most crowded index was for the code  $\{+ =\}$ , that is, for positive equality clauses. Finally, 6955 indices contained only a single clause.

## 5.2 Experimental Evaluation of ENIGMAWatch


The experiments are done on a random subset of 5000 Mizar40 [21] problems. The time limit of 60s and 30000 generated clauses is used to allow a comparison to be done without regard for the differences in clause processing speed. The 30000 is approximately the average number of clauses that the baseline strategy generates in 10s. Table 3 provides the evaluation of different watchlist selection mechanisms using ProofWatch (without ENIGMA) and making use of the improved watchlist indexing. The last two columns show the number of problems solved by (1) the Baseline together with Mean, and by (2) all the five methods. This shows the relative complementarity of the methods. We can see that the Mean method yields the best results, reaching more than 15% improvement over the baseline strategy. The Rand method is however quite competitive.

Table 4 provides the evaluation of ENIGMAWatch and its comparison to ENIGMA. The experiments are done in multiple loops, where in each loop all the proof-runs in prior loops can be used as training data. This way ENIGMA can learn increasingly effective models.

We can see that ENIGMAWatch can attain superior performance to ENIGMA. The relation of looping and results is interesting. The largest absolute improvement over ENIGMA is in loop 0 – 8.8% by the Mean method. This however drops to 1.2% in loop 4. In loops 1 and 2, Rand is the strongest, but Mean ends up being the best in loop 3. In total, all the ENIGMA and ENIGMAWatch methods solve together nearly twice as many problems as the baseline strategy. Figure 1 shows the results of running ENIGMA and Mean for 13 loops. The rate of improvement slows down, both methods eventually converge to a similar level of performance, and the union of the two is ca. 150 problems better.

**Table 5.** ENIGMA and ENIGMAWatch: Model and training statistics.

Model	Pos. acc	Neg. acc	Features	Watchlist F	Train size	Train time
ENIGMA0	99.12%	92.16%	5061	0	0.4 GB	14 min
ENIGMA1	97.39%	86.82%	7071	0	0.8 GB	31 min
ENIGMA2	96.13%	83.92%	8089	0	1.4 GB	55 min
ENIGMA3	95.39%	82.5%	8662	0	2.0 GB	85 min
Mean0	99.05%	92.59%	5424	308	2.9 GB	19 min
Mean1	96.92%	88.16%	6950	316	6.2 GB	29 min
Mean2	95.75%	86.46%	7809	331	9.6 GB	38 min
Mean3	95.04%	85.24%	8313	330	13.0 GB	39 min



**Fig. 1.** Convergence: The improvement of ENIGMA and Mean decreases over 13 loops, and their performance converges. The Union is consistently ca. 150 problems better.

### 5.3 Training, Model Statistics and Analysis

The XGBoost models used in our experiments are trained with a maximum tree depth of 9 and 200 rounds (which means 200 trees are learned). There are 300000 features in the 5000 problem dataset hashed into  $2^{15}$  buckets. Combining clause and conjecture features with the watchlist completion ratios, XGBoost makes its predictions based on 66048 features ( $2 \cdot 2^{15}$  plus the count of completion ratios).

Table 5 provides various training and model statistics of the ENIGMA and ENIGMAWatch models and their loops. The columns “Pos. Acc.” and “Neg. Acc.” describe the training accuracy of the models on positive and negative training examples. The column “Features” presents the number of features referenced in the decision trees. We see that the models use a small fraction of all the 66048 available features. The column “Watchlist F.” provides the number of watchlist features out of all the used features. Finally, “Train Size” and “Train Time” specify the size of the input training file (in GB) and training times (in minutes). The XGBoost models after the training are smaller than 4 MB.

We can see that the accuracy decreases with the increase of the training data size, but the number of theorems proved increases. About 62% of the watchlists are judged as useful by XGBoost and used in the decision trees. Figure 2 shows the root of the first decision tree of the Mean model in loop 3. Green means “yes” (the condition holds), red means “no”, and blue means that the feature is not present. The multi line box is a (shortened) bucket of features, and single line boxes correspond to watchlists (#194, etc.). We can see that ENIGMAWatch uses a watchlist feature for the very first decision when judging newly generated clauses. This shows that the features that characterize the evolving proof state are indeed considered very significant by the methods that automatically learn given clause guidance.

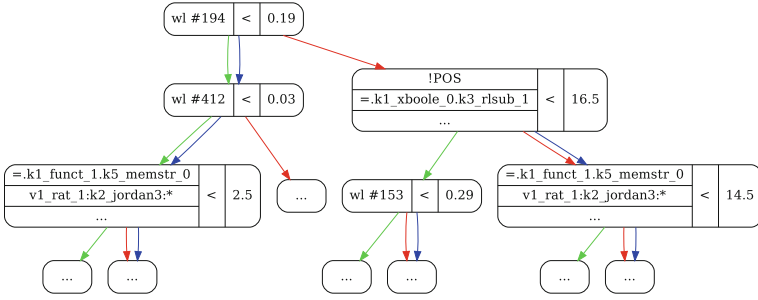


Fig. 2. Example of an XGBoost decision tree.

## 6 Conclusion and Future Work

We have produced and evaluated the first practically usable version of the ENIGMAWatch system which can now be efficiently used over large mathematical datasets. The previous experiments with the first prototype on the small MPTP Challenge [12] demonstrated that ENIGMAWatch can find proofs faster (in terms of how many processed clauses are needed). The work presented here shows that with improved subsumption indexing, feature hashing, and suitable global watchlist selection, ENIGMAWatch outperforms ENIGMA on the large Mizar40 dataset. In particular, ENIGMAWatch significantly outperforms both ProofWatch and ENIGMA when used without looping. With several MaLAREa-style [37, 40] iterations of proving and learning, the difference to ENIGMA gets smaller, however the two methods are still quite complementary, providing solutions to a large number of different problems. In total, all the ENIGMA and ENIGMAWatch methods (Table 4) together solve almost twice as many problems as the baseline strategy after four iterations of learning and proving.

The system is ready to be used on hard problems and to expand the set of Mizar problems for which an ATP proof has been found. Future work includes refining the watchlist selection, defining more sophisticated methods of computing the proof completion ratios, analyzing the learned decision tree models to see which watchlists are the most useful, and also defining further and more abstract meaningful representations and embeddings of saturation-style proof search.

## References

1. Alama, J., Heskes, T., Kühlwein, D., Tsvitshivadze, E., Urban, J.: Premise selection for mathematics by corpus analysis and kernel methods. *J. Autom. Reasoning* **52**(2), 191–213 (2014)
2. Alemi, A.A., Chollet, F., Eén, N., Irving, G., Szegedy, C., Urban, J.: DeepMath - deep sequence models for premise selection. In: Lee, D.D., Sugiyama, M., Luxburg, U.V., Guyon, I., Garnett, R. (eds.) *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016*, 5–10 December 2016, Barcelona, Spain, pp. 2235–2243 (2016)

3. Blanchette, J.C., Greenaway, D., Kaliszzyk, C., Kühlwein, D., Urban, J.: A learning-based fact selector for Isabelle/HOL. *J. Autom. Reasoning* **57**(3), 219–244 (2016)
4. Bridge, J.P., Holden, S.B., Paulson, L.C.: Machine learning for first-order theorem proving - learning to select a good heuristic. *J. Autom. Reasoning* **53**(2), 141–172 (2014)
5. Chen, T., Guestrin, C.: Xgboost: a scalable tree boosting system. In: *KDD*, pp. 785–794. ACM (2016)
6. Chvalovský, K., Jakubuv, J., Suda, M., Urban, J.: ENIGMA-NG: efficient neural and gradient-boosted inference guidance for E. CoRR, abs/1903.03182 (2019)
7. Denzinger, J., Fuchs, M., Goller, C., Schulz, S.: Learning from Previous Proof Experience. Technical Report AR99-4, Institut für Informatik, Technische Universität München (1999)
8. Ertel, W., Schumann, J., Suttner, C.B.: Learning heuristics for a theorem prover using back propagation. In: Retti, J., Leidlmaier, K. (eds.) *Österreichische Artificial Intelligence-Tagung, Igls, Tirol*, vol. 208, pp. 87–95. Springer, Heidelberg (1989). [https://doi.org/10.1007/978-3-642-74688-8\\_10](https://doi.org/10.1007/978-3-642-74688-8_10)
9. Färber, M., Brown, C.: Internal Guidance for Satallax. In: Olivetti, N., Tiwari, A. (eds.) *IJCAR 2016. LNCS (LNAI)*, vol. 9706, pp. 349–361. Springer, Cham (2016). [https://doi.org/10.1007/978-3-319-40229-1\\_24](https://doi.org/10.1007/978-3-319-40229-1_24)
10. Gauthier, T., Kaliszzyk, C.: Premise selection and external provers for HOL4. In: *Certified Programs and Proofs (CPP 2015)*. LNCS. Springer, Berlin (2015). <https://doi.org/10.1145/2676724.2693173>
11. Goertzel, Z., Jakubuv, J., Schulz, S., Urban, J.: ProofWatch: watchlist guidance for large theories in E. In: Avigad, J., Mahboubi, A. (eds.) *ITP 2018. LNCS*, vol. 10895, pp. 270–288. Springer, Cham (2018). [https://doi.org/10.1007/978-3-319-94821-8\\_16](https://doi.org/10.1007/978-3-319-94821-8_16)
12. Goertzel, Z., Jakubuv, J., Urban, J.: ProofWatch meets ENIGMA: first experiments. In: Barthe, G., Korovin, K., Schulz, S., Suda, M., Sutcliffe, G., Veanes, M. (eds.) *LPAR-22 Workshop and Short Paper Proceedings*, vol. 9, pp. 15–22. Kalpa Publications in Computing. EasyChair (2018)
13. Gottlob, G., Sutcliffe, G., Voronkov, A. (eds.): *Global Conference on Artificial Intelligence, GCAI 2015, Tbilisi, Georgia, October 16–19, 2015*, vol. 36, EPiC Series in Computing. EasyChair (2015)
14. Grabowski, A., Kornilowicz, A., Naumowicz, A.: Mizar in a nutshell. *J. Formalized Reasoning* **3**(2), 153–245 (2010)
15. Jakubuv, J., Urban, J.: Hierarchical invention of theorem proving strategies. *AI Commun.* **31**(3), 237–250 (2018)
16. Jakubuv, J., Urban, J.: ENIGMA: efficient learning-based inference guiding machine. In: Geuvers, H., England, M., Hasan, O., Rabe, F., Teschke, O. (eds.) *CICM 2017. LNCS (LNAI)*, vol. 10383, pp. 292–302. Springer, Cham (2017). [https://doi.org/10.1007/978-3-319-62075-6\\_20](https://doi.org/10.1007/978-3-319-62075-6_20)
17. Jakubuv, J., Urban, J.: Enhancing ENIGMA given clause guidance. In: Rabe, F., Farmer, W.M., Passmore, G.O., Youssef, A. (eds.) *CICM 2018. LNCS (LNAI)*, vol. 11006, pp. 118–124. Springer, Cham (2018). [https://doi.org/10.1007/978-3-319-96812-4\\_11](https://doi.org/10.1007/978-3-319-96812-4_11)
18. Jakubuv, J., Urban, J.: Hammering Mizar by learning clause guidance. *CoRR*, abs/1904.01677 (2019)
19. Kaliszzyk, C., Urban, J.: Learning-assisted automated reasoning with Flyspeck. *J. Autom. Reasoning* **53**(2), 173–213 (2014)

20. Kaliszyk, C., Urban, J.: FEMaLeCoP: fairly efficient machine learning connection prover. In: Davis, M., Fehnker, A., McIver, A., Voronkov, A. (eds.) LPAR 2015. LNCS, vol. 9450, pp. 88–96. Springer, Heidelberg (2015). [https://doi.org/10.1007/978-3-662-48899-7\\_7](https://doi.org/10.1007/978-3-662-48899-7_7)
21. Kaliszyk, C., Urban, J.: MizAR 40 for Mizar 40. *J. Autom. Reasoning* **55**(3), 245–256 (2015)
22. Kaliszyk, C., Urban, J., Michalewski, H., Olsák, M.: Reinforcement learning of theorem proving. In: *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, Montréal, Canada, 3–8 December 2018*, pp. 8836–8847 (2018)
23. Kaliszyk, C., Urban, J., Vyskocil, J.: Efficient semantic features for automated reasoning over large theories. In: IJCAI, pp. 3084–3090. AAAI Press (2015)
24. Kinyon, M., Veroff, R., Vojtěchovský, P.: Loops with abelian inner mapping groups: an application of automated deduction. In: Bonacina, M.P., Stickel, M.E. (eds.) *Automated Reasoning and Mathematics*. LNCS (LNAI), vol. 7788, pp. 151–164. Springer, Heidelberg (2013). [https://doi.org/10.1007/978-3-642-36675-8\\_8](https://doi.org/10.1007/978-3-642-36675-8_8)
25. Kovács, L., Voronkov, A.: First-order theorem proving and VAMPIRE. In: Sharygina, N., Veith, H. (eds.) CAV 2013. LNCS, vol. 8044, pp. 1–35. Springer, Heidelberg (2013). [https://doi.org/10.1007/978-3-642-39799-8\\_1](https://doi.org/10.1007/978-3-642-39799-8_1)
26. Loos, S.M., Irving, G., Szegedy, C., Kaliszyk, C.: Deep network guided proof search. In: Eiter, T., Sands, D. (eds.) LPAR-21, 21st International Conference on Logic for Programming, Artificial Intelligence and Reasoning, Maun, Botswana, 7–12 May 2017. EPiC Series in Computing, vol. 46, pp. 85–105. EasyChair (2017)
27. McCune, W., Wos, L.: Otter: the CADE-13 competition incarnations. *J. Autom. Reasoning* **18**(2), 211–220 (1997). Special Issue on the CADE 13 ATP System Competition
28. McCune, W.W.: Prover9 and Mace4. <http://www.cs.unm.edu/~mccune/prover9/>, 2005–2010. Accessed 29 Mar 2016
29. Piotrowski, B., Urban, J.: ATPBOOST: learning premise selection in binary setting with ATP feedback. In: Galmiche, D., Schulz, S., Sebastiani, R. (eds.) IJCAR 2018. LNCS (LNAI), vol. 10900, pp. 566–574. Springer, Cham (2018). [https://doi.org/10.1007/978-3-319-94205-6\\_37](https://doi.org/10.1007/978-3-319-94205-6_37)
30. Polikar, R.: Ensemble based systems in decision making. *IEEE Circuits Syst. Mag.* **6**(3), 21–45 (2006)
31. Schäfer, S., Schulz, S.: Breeding theorem proving heuristics with genetic algorithms. In: Gottlob et al. [13], pp. 263–274 (2015)
32. Schulz, S.: Learning search control knowledge for equational deduction. In: DISKI, vol. 230, Infix Akademische Verlagsgesellschaft (2000)
33. Schulz, S.: E - a brainiac theorem prover. *AI Commun.* **15**(2–3), 111–126 (2002)
34. Schulz, S.: Simple and efficient clause subsumption with feature vector indexing. In: Bonacina, M.P., Stickel, M.E. (eds.) *Automated Reasoning and Mathematics*. LNCS (LNAI), vol. 7788, pp. 45–67. Springer, Heidelberg (2013). [https://doi.org/10.1007/978-3-642-36675-8\\_3](https://doi.org/10.1007/978-3-642-36675-8_3)
35. Schulz, S.: System description: E 1.8. In: McMillan, K., Middeldorp, A., Voronkov, A. (eds.) LPAR 2013. LNCS, vol. 8312, pp. 735–743. Springer, Heidelberg (2013). [https://doi.org/10.1007/978-3-642-45221-5\\_49](https://doi.org/10.1007/978-3-642-45221-5_49)
36. Urban, J.: MPTP 0.2: Design, implementation, and initial experiments. *J. Autom. Reasoning* **37**(1–2), 21–43 (2006)
37. Urban, J.: MaLARea: a metasystem for automated reasoning in large theories. In: Sutcliffe, G., Urban, J., Schulz, S. (eds.) ESARLT, vol. 257, CEUR Workshop Proceedings. CEUR-WS.org (2007)

38. Urban, J.: BliStr: the blind strategymaker. In: Gottlob et al. [13], pp. 312–319 (2013)
39. Urban, J., Sutcliffe, G.: ATP cross-verification of the mizar MPTP challenge problems. In: Dershowitz, N., Voronkov, A. (eds.) LPAR 2007. LNCS (LNAI), vol. 4790, pp. 546–560. Springer, Heidelberg (2007). [https://doi.org/10.1007/978-3-540-75560-9\\_39](https://doi.org/10.1007/978-3-540-75560-9_39)
40. Urban, J., Sutcliffe, G., Pudlák, P., Vyskočil, J.: MaLAREa SG1 - machine learner for automated reasoning with semantic guidance. In: Armando, A., Baumgartner, P., Dowek, G. (eds.) IJCAR 2008. LNCS (LNAI), vol. 5195, pp. 441–456. Springer, Heidelberg (2008). [https://doi.org/10.1007/978-3-540-71070-7\\_37](https://doi.org/10.1007/978-3-540-71070-7_37)
41. Urban, J., Vyskočil, J., Štěpánek, P.: MaLeCoP machine learning connection prover. In: Brünnler, K., Metcalfe, G. (eds.) TABLEAUX 2011. LNCS (LNAI), vol. 6793, pp. 263–277. Springer, Heidelberg (2011). [https://doi.org/10.1007/978-3-642-22119-4\\_21](https://doi.org/10.1007/978-3-642-22119-4_21)
42. Veroff, R.: Using hints to increase the effectiveness of an automated reasoning program: case studies. *J. Autom. Reasoning* **16**(3), 223–239 (1996)