



ENIGMA Anonymous: Symbol-Independent Inference Guiding Machine (System Description)

Jan Jakubův¹, Karel Chvalovský¹, Miroslav Olšák², Bartosz Piotrowski^{1,3},
Martin Suda¹(✉), and Josef Urban¹

¹ Czech Technical University in Prague, Prague, Czechia
martin.suda@cvut.cz

² University of Innsbruck, Innsbruck, Austria

³ University of Warsaw, Warsaw, Poland

Abstract. We describe an implementation of gradient boosting and neural guidance of saturation-style automated theorem provers that does not depend on consistent symbol names across problems. For the gradient-boosting guidance, we manually create abstracted features by considering arity-based encodings of formulas. For the neural guidance, we use symbol-independent graph neural networks (GNNs) and their embedding of the terms and clauses. The two methods are efficiently implemented in the E prover and its ENIGMA learning-guided framework.

To provide competitive real-time performance of the GNNs, we have developed a new context-based approach to evaluation of generated clauses in E. Clauses are evaluated jointly in larger batches and with respect to a large number of already selected clauses (context) by the GNN that estimates their collectively most useful subset in several rounds of message passing. This means that approximative inference rounds done by the GNN are efficiently interleaved with precise symbolic inference rounds done inside E. The methods are evaluated on the MPTP large-theory benchmark and shown to achieve comparable real-time performance to state-of-the-art symbol-based methods. The methods also show high complementarity, solving a large number of hard Mizar problems.

Keywords: Automated theorem proving · Machine learning · Neural networks · Decision trees · Saturation-style proving

Supported by the ERC Consolidator grant AI4REASON no. 649043 (JJ, BP, MS, and JU), the Czech project AI&Reasoning CZ.02.1.01/0.0/0.0/15_003/0000466 and the European Regional Development Fund (KC and JU), the ERC Project *SMART* Starting Grant no. 714034 (MO), grant 2018/29/N/ST6/02903 of National Science Center, Poland (BP), and the Czech Science Foundation project 20-06390Y (MS).

© Springer Nature Switzerland AG 2020

N. Peltier and V. Sofronie-Stokkermans (Eds.): IJCAR 2020, LNAI 12167, pp. 448–463, 2020.

https://doi.org/10.1007/978-3-030-51054-1_29

1 Introduction: Symbol Independent Inference Guidance

In this work, we develop two *symbol-independent* (anonymous) inference guiding methods for saturation-style automated theorem provers (ATPs) such as E [25] and Vampire [20]. Both methods are based on learning clause classifiers from previous proofs within the ENIGMA framework [5, 13, 14] implemented in E. By *symbol-independence* we mean that no information about the symbol names is used by the learned guidance. In particular, if all symbols in a particular ATP problem are consistently renamed to new symbols, the learned guidance will result in the same proof search and the same proof modulo the renaming.

Symbol-independent guidance is an important challenge for learning-guided ATP, addressed already in Schulz’s early work on learning guidance in E [23]. With ATPs being increasingly used and trained on large ITP libraries [2, 3, 6, 8, 16, 18], it is more and more rewarding to develop methods that learn to reason without relying on the particular terminology adopted in a single project. Initial experiments in this direction using concept alignment [10] methods have already shown performance improvements by transferring knowledge between the HOL libraries [9]. Structural analogies (or even terminology duplications) are however common already in a single large ITP library [17] and their automated detection can lead to new proof ideas and a number of other interesting applications [11].

This system description first briefly introduces saturation-based ATP with learned guidance (Sect. 2). Then we discuss symbol-independent learning and guidance using abstract features and gradient boosting trees (Sect. 3) and graph neural networks (Sect. 4). The implementation details are explained in Sect. 5 and the methods are evaluated on the MPTP benchmark in Sect. 6.

2 Saturation Proving Guided by Machine Learning

Saturation-Based Automated Theorem Provers (ATPs) such as E and Vampire are used to prove goals G using a set of axioms A . They classify the formulas $A \cup \{-G\}$ and try to deduce contradiction using the *given clause loop* [22] as follows. The ATP maintains two sets of processed (P) and unprocessed (U) clauses. At each loop iteration, a given clause g from U is selected, moved to P , and U is extended with new inferences from g and P . This process continues until the contradiction is found, U becomes empty, or a resource limit is reached. The search space grows quickly and selection of the right given clauses is critical.

Learning Clause Selection over a set of related problems is a general method how to guide the proof search. Given a set of FOL problems \mathcal{P} and initial ATP strategy \mathcal{S} , we can evaluate \mathcal{S} over \mathcal{P} obtaining training samples \mathcal{T} . For each successful proof search, training samples \mathcal{T} contain the set of clauses processed during the search. *Positive* clauses are those that were *useful* for the proof search (they appeared in the final proof), while the remaining clauses were *useless*, forming the *negative* examples. Given the samples \mathcal{T} , we can *train* a machine learning *classifier* \mathcal{M} which predicts usefulness of clauses in future proof searches. Some clause classifiers are described in detail in Sects. 3, 4, and 5.

ATP Guidance By a Trained Classifier: Once a clause classifier \mathcal{M} is trained, we can use it inside an ATP. An ATP strategy \mathcal{S} is a collection of proof search parameters such as term ordering, literal selection, and also given clause selection mechanism. In E, the given clause selection is defined by a collection of clause *weight functions* which alternate to select the given clauses. Our ENIGMA framework uses two methods of plugging the trained classifier \mathcal{M} into \mathcal{S} . Either (1) we use \mathcal{M} to select all given clauses (*solo mode* denoted $\mathcal{S} \odot \mathcal{M}$), or (2) we combine predictions of \mathcal{M} with clause selection mechanism from \mathcal{S} so that roughly 50% of the clauses is selected by \mathcal{M} (*cooperative mode* denoted $\mathcal{S} \oplus \mathcal{M}$). Proof search settings other than clause selection are inherited from \mathcal{S} in both the cases. See [5] for details. The phases of learning and ATP guidance can be iterated in a *learning/evaluation loop* [29], yielding growing sets of proofs \mathcal{T}_i and stronger classifiers \mathcal{M}_i trained over them. See [15] for such large experiment.

3 Clause Classification by Decision Trees

Clause Features are used by ENIGMA to represent clauses as sparse vectors for machine learners. They are based mainly on vertical/horizontal cuts of the clause syntax tree. We use simple *feature hashing* to handle theories with large number of symbols. A clause C is represented by the vector φ_C whose i -th index stores the value of a feature with hash index i . Values of conflicting features (mapped to the same index) are summed. Additionally, we embed *conjecture features* into the clause representation and we work with vector pairs (φ_C, φ_G) of size $2 * base$, where φ_G is the feature vector of the current goal (conjecture). This allows us to provide goal-specific predictions. See [15] for more details.

Gradient Boosting Decision Trees (GBDTs) implemented by the XGBoost library [4] currently provide the strongest ENIGMA classifiers. Their speed is comparable to the previously used [14] weaker linear logistic classifier, implemented by the LIBLINEAR library [7]. In this work, we newly employ the LightGBM [19] GBDT implementation. A *decision tree* is a binary tree whose nodes contain Boolean conditions on values of different features. Given a feature vector φ_C , the decision tree can be navigated from the root to the unique tree leaf which contains the classification of clause C . GBDTs combine predictions from a collection of follow-up decision trees. While inputs, outputs, and API of XGBoost and LightGBM are compatible, each employ a different method of tree construction. XGBoost constructs trees level-wise, while LightGBM leaf-wise. This implies that XGBoost trees are well-balanced. On the other hand, LightGBM can produce much deeper trees and the tree depth limit is indeed an important learning meta-parameter which must be additionally set.

New Symbol-Independent Features: We develop a feature anonymization method based on symbol arities. Each function symbol name s with arity n is substituted by a special name “ \mathbf{fn} ”, while a predicate symbol name q with arity m is substituted by “ \mathbf{pm} ”. Such features lose the ability to distinguish different symbol names, and many features are merged together. Vector representations

of two clauses with renamed symbols are clearly equal. Hence the underlying machine learning method will provide equal predictions for such clauses. For more detailed discussion and comparison with related work see Appendix B.

New Statistics and Problem Features: To improve the ability to distinguish different anonymized clauses, we add the following features. *Variable statistics* of clause C containing (1) the number of variables in C without repetitions, (2) the number of variables with repetitions, (3) the number of variables with exactly one occurrence, (4) the number of variables with more than one occurrence, (5–10) the number of occurrences of the most/least (and second/third most/least) occurring variable. *Symbol statistics* do the same for symbols instead of variables. Recall that we embed conjecture features in clause vector pair (φ_C, φ_G) . As G embeds information about the conjecture but not about the problem axioms, we propose to additionally embed some statistics of the problem P that C and G come from. We use 22 problem features that E prover already computes for each input problem to choose a suitable strategy. These are (1) number of goals, (2) number of axioms, (3) number of unit goals, etc. See E’s manual for more details. Hence we work with vector triples $(\varphi_C, \varphi_G, \varphi_P)$.

4 Clause Classification by Graph Neural Network

Another clause classifier newly added to ENIGMA is based on graph neural networks (GNNs). We use the symbol-independent network architecture developed in [21] for premise selection. As [21] contains all the details, we only briefly explain the basic ideas behind this architecture here.

Hypergraph. Given a set of clauses \mathcal{C} we create a directed hypergraph with three kinds of nodes that correspond to clauses, function and predicate symbols \mathcal{N} , and unique (sub)terms and literals \mathcal{U} occurring in \mathcal{C} , respectively. There are two kinds of hyperedges that describe the relations between nodes according to \mathcal{C} . The first kind encodes literal occurrences in clauses by connecting the corresponding nodes. The second hyperedge kind encodes the relations between nodes from \mathcal{N} and \mathcal{U} . For example, for $f(t_1, \dots, t_k) \in \mathcal{U}$ we loosely speaking connect the nodes $f \in \mathcal{N}$ and $t_1, \dots, t_k \in \mathcal{U}$ with the node $f(t_1, \dots, t_k)$ and similarly for literals, where their polarity is also taken into account.

Message-Passing. The hypergraph describes the relation between various kinds of objects occurring in \mathcal{C} . Every node in the hypergraph is initially assigned a constant vector, called the *embedding*, based only on its kind (\mathcal{C} , \mathcal{N} , or \mathcal{U}). These node embeddings are updated in a fixed number of message-passing rounds, based on the embeddings of each node’s neighbors. The underlying idea of such neural message-passing methods¹ is to make the node embeddings encode more and more precisely the information about the connections (and thus various

¹ Graph convolutions are a generalization of the sliding window convolutions used for aggregating neighborhood information in neural networks used for image recognition.

properties) of the nodes. For this to work, we have to learn initial embeddings for our three kinds of nodes and the update function.²

Classification. After the message-passing phase, the final clause embeddings are available in the corresponding clause nodes. The estimated probability of a clause being a good given clause is then computed by a neural network that takes the final embedding of this clause and also aggregated final embeddings of all clauses obtained from the negated conjecture.

5 Learning and Using the Classifiers, Implementation

In order to use either GBDTs (Sect. 3) or GNNs (Sect. 4), a prediction model must be learned. Learning starts with training samples \mathcal{T} , that is, a set of pairs $(\mathcal{C}^+, \mathcal{C}^-)$ of positive and negative clauses. For each training sample $T \in \mathcal{T}$, we additionally know the source problem P and its conjecture G . Hence we can consider one sample $T \in \mathcal{T}$ as a quadruple $(\mathcal{C}^+, \mathcal{C}^-, P, G)$ for convenience.

GBDT. Given a training sample $T = (\mathcal{C}^+, \mathcal{C}^-, P, G) \in \mathcal{T}$, each clause $C \in \mathcal{C}^+ \cup \mathcal{C}^-$ is translated to the feature vector $(\varphi_C, \varphi_G, \varphi_P)$. Vectors where $C \in \mathcal{C}^+$ are labeled as positive, and otherwise as negative. All the labeled vectors are fed together to a GBDT trainer yielding model $\mathcal{D}_{\mathcal{T}}$.

When predicting a generated clause, the feature vector is computed and $\mathcal{D}_{\mathcal{T}}$ is asked for the prediction. GBDT’s binary predictions (positive/negative) are turned into E’s clause weight (positives have weight 1 and negatives 10).

GNN. Given $T = (\mathcal{C}^+, \mathcal{C}^-, P, G) \in \mathcal{T}$ as above we construct a hypergraph for the set of clauses $\mathcal{C}^+ \cup \mathcal{C}^- \cup G$. This hypergraph is translated to a tensor representation (vectors and matrices), marking clause nodes as positive, negative, or goal. These tensors are fed as input to our GNN training, yielding a GNN model $\mathcal{N}_{\mathcal{T}}$. The training works in iterations, and $\mathcal{N}_{\mathcal{T}}$ contains one GNN per iteration epoch. Only one GNN from a selected epoch is used for predictions during the evaluation.

In evaluation, it is more efficient to compute predictions for several clauses at once. This also improves prediction quality as the queried data resembles more the training hypergraphs where multiple clauses are encoded at once as well. During an ATP run on problem P with the conjecture G , we postpone evaluation of newly inferred clauses until we reach a certain amount of clauses \mathcal{Q} to *query*.³ To resemble the training data even more, we add a fixed number of the given clauses processed so far. We call these *context* clauses (\mathcal{X}). To evaluate \mathcal{Q} , we construct the hypergraph for $\mathcal{Q} \cup \mathcal{X} \cup G$, and mark clauses from G as goals. Then model $\mathcal{N}_{\mathcal{T}}$ is asked for predictions on \mathcal{Q} (predictions for \mathcal{X} are dropped). The numeric predictions computed by $\mathcal{N}_{\mathcal{T}}$ are directly used as E’s weights.

Implementation and Performance. We use GBDTs implemented by the XGBoost [4] and LightGBM [19] libraries. For GNN we use Tensorflow [1]. All

² We learn individual components, which correspond to different kinds of hyperedges, from which the update function is efficiently constructed.

³ We may evaluate less than \mathcal{Q} if E runs out of unevaluated unprocessed clauses.

Table 1. Model training and evaluation for anonymous GBDTs (\mathcal{D}_i) and GNN (\mathcal{N}_i).

\mathcal{M}	TPR	TNR	Training			Real time		Abstract time	
	[%]	[%]	Size	Time	Params	$\mathcal{S} \oplus \mathcal{M}$	+%	$\mathcal{S} \oplus \mathcal{M}$	+%
\emptyset	-	-	-	-	-	14 966	0.0	10 679	0.0
\mathcal{D}_0	84.9	68.4	14M	2h29m	X,d12	20 679	38.1	17 917	67.8
\mathcal{D}_1	79.0	79.5	29M	4h33m	X,d12	23 280	58.2	20 760	94.4
\mathcal{D}_2	80.5	79.2	47M	40m	L,d30,l1800	24 347	62.7	22 661	112.2
\mathcal{N}_0	92.1	77.1	14M	17h	e20,q128,c512	20 912	39.7	19 755	84.9
\mathcal{N}_1	90.0	78.6	31M	1d19h	e10,q128,c512	23 156	54.7	21 737	103.5
\mathcal{N}_2	91.3	79.6	50M	1d 8h	e50,q256,c768	23 262	55.4	22 169	107.6

the libraries provide Python interfaces and C/C++ APIs. We use the Python interfaces for training and the C APIs for the evaluation in E. The Python interfaces for XGBoost and LightGBM include the C APIs, while for Tensorflow this must be manually compiled, which is further complicated by poor documentation.

The libraries support training both on CPUs and on GPUs. We train LightGBM on CPUs, and XGBoost and Tensorflow on GPUs. However, we always evaluate on a single CPU as we aim at practical usability on standard hardware. This is non-trivial and it distinguishes this work from evaluations done with large numbers of GPUs or TPUs and/or in prohibitively high real times. The LightGBM training can be parallelized much better – with 60 CPUs it is much faster than XGBoost on 4 GPUs. Neither using GPUs for LightGBM nor many CPUs for XGBoost provided better training times. The GNN training is slower than GBDT training and it is not easy to make Tensorflow evaluate reasonably on a single CPU. It has to be compiled with all CPU optimizations and restricted to a single thread, using Tensorflow’s poorly documented experimental C API.

6 Experimental Evaluation

Setup. We experimentally evaluate⁴ our GBDT and GNN guidance⁵ on a large benchmark of 57880 Mizar40 [18] problems⁶ exported by MPTP [28]. Hence this evaluation is compatible with our previous symbol-dependent work [15]. We evaluate GBDT and GNN separately. We start with a good-performing E strategy \mathcal{S} (see [5, Appendix A]) which solves 14 966 problems with a 10 s limit per problem. This gives us training data $\mathcal{T}_0 = \text{eval}(\mathcal{S})$ (see Sect. 5), and we start three iterations of the learning/evaluation loop (see Sect. 2).

⁴ On a server with 36 hyperthreading Intel(R) Xeon(R) Gold 6140 CPU @ 2.30 GHz cores, 755 GB of memory, and 4 NVIDIA GeForce GTX 1080 Ti GPUs.

⁵ Available at <https://github.com/ai4reason/eprover-data/tree/master/IJCAR-20>.

⁶ <http://grid01.ciirc.cvut.cz/~mptp/1147/MPTP2/problems.small.consist.tar.gz>.

For GBDT, we train several models (with hash base 2^{15}) and conduct a small learning meta-parameters *grid search*. For XGBoost, we try different tree depths ($d \in \{9, 12, 16\}$), and for LightGBM various combinations of tree depths and leaves count ($(d, l) \in \{10, 20, 30, 40\} \times \{1200, 1500, 1800\}$). We evaluate all these models in a cooperative mode with \mathcal{S} on a random (but fixed) 10% of all problems (Appendix A). The best performing model is evaluated on the whole benchmark in both cooperative (\oplus) and solo (\odot) runs. These give us the next samples \mathcal{T}_{i+1} . We perform three iterations and obtain models \mathcal{D}_0 , \mathcal{D}_1 , and \mathcal{D}_2 .

For GNN, we train a model with 100 epochs, obtaining 100 different GNNs. We evaluate GNNs from selected epochs ($e \in \{10, 20, 50, 75, 100\}$) and we try different settings of *query* (q) and *context* (c) sizes (see Sect. 5). In particular, q ranges over $\{64, 128, 192, 256, 512\}$ and c over $\{512, 768, 1024, 1536\}$. All possible combinations of (e, q, c) are again evaluated in a grid search on the small benchmark subset (Appendix A), and the best performing model is selected for the next iteration. We run three iterations and obtain models \mathcal{N}_0 , \mathcal{N}_1 , and \mathcal{N}_2 .

Results are presented in Table 1. For each model \mathcal{D}_i and \mathcal{N}_i we show (1) true positive/negative rates, (2) training data sizes, (3) train times, and (4) the best performing parameters from the grid search. Furthermore, for each model \mathcal{M} we show the performance of $\mathcal{S} \oplus \mathcal{M}$ in (5) real and (6) abstract time. Details follow. (1) Model accuracies are computed on samples extracted from problems newly solved by each model, that is, on testing data not known during the training. Columns TPR/TNR show accuracies on positive/negative testing samples. (2) Train sizes measure the training data in millions of clauses. (4) Letter “X” stands for XGBoost models, while “L” for LightGBM. (5) For real time we use 10 s limit per problem, and (6) in abstract time we limit the number of generated clauses to 5000. We show the number of problems solved and the gain (in %) on \mathcal{S} . The abstract time evaluation is useful to assess the methods modulo the speed of the implementation. The first row shows the performance of \mathcal{S} without learning.

Evaluation. The GNN models start better, but the GBDT models catch up and beat GNN in later iterations. The GBDT models show a significant gain even in the 3rd iteration, while the GNN models start stagnating. The GNN models report better testing accuracy, but their ATP performance is not as good.

For GBDTs, we see that the first two best models (\mathcal{D}_0 and \mathcal{D}_1) were produced by XGBoost, while \mathcal{D}_2 by LightGBM. While both libraries can provide similar results, LightGBM is significantly faster. For comparison, the training time for XGBoost in the third iteration was 7 h, that is, LightGBM is 10 times faster. The higher speed of LightGBM can overcome the problems with more complicated parameter settings, as more models can be trained and evaluated.

For GNNs, we observe higher training times and better models coming from earlier epochs. The training in the 1st and 2nd iterations was done on 1 GPU, while in the 3rd on 4 GPUs. The good abstract time performance indicates that further gain could be obtained by a faster implementation. But note that this is the first time that NNs have been made comparable to GBDTs in real time.




Fig. 1. Left: the number of problems solved in time; Right: the number of processed clauses (the x -axis for S , and the y -axis for $S \oplus D_0$ and $S \oplus \mathcal{N}_0$, respectively).

Figure 1 summarizes the results. On the left, we observe a slower start for GNNs caused by the initial model loading. On the right, we see a decrease in the number of processed clauses, which suggests that the guidance is effective.

Complementarity. The twelve (solo and cooperative) versions of the methods compared in Fig. 1 solve together 28271 problems, with the six GBDTs solving 25255 and the six GNNs solving 26571. All twenty methods tested by us solve 29118 problems, with the top-6 greedy cover solving (in 60 s) 28067 and the top-15 greedy cover solving (in 150 s) 29039. The GNNs show higher complementarity – varying the epoch as well as the size of the query and context produces many new solutions. For example, the most complementary GNN method adds to the best GNN method 1976 solutions. The GNNs are also quite complementary to the GBDTs. The second (GNN) strategy in the greedy cover adds 2045 solutions to the best (GBDT) strategy. Altogether, the twenty strategies solve (in 200 s) 2109 of the Mizar40 *hard* problems, i.e., the problems unsolved by any method developed previously in [18].

7 Conclusion

We have developed and evaluated symbol-independent GBDT and GNN ATP guidance. This is the first time symbol-independent features and GNNs are tightly integrated with E and provide good real-time results on a large corpus. Both the GBDT and GNN predictors display high ability to learn from previous proof searches even in the symbol-independent setting.

To provide competitive real-time performance of the GNNs, we have developed context-based evaluation of generated clauses in E. This introduces a new paradigm for clause ranking and selection in saturation-style proving. The generated clauses are not ranked immediately and independently of other clauses. Instead, they are judged in larger batches and with respect to a large number of already selected clauses (context) by a neural network that estimates their collectively most useful subset by several rounds of message passing. This also allows

new ways of parameterizing the search that result in complementary methods with many new solutions.

The new GBDTs show even better performance than their symbol-dependent versions from our previous work [15]. This is most likely because of the parameter grid search and new features not used before. The union of the problems solved by the twelve ENIGMA strategies (both \odot and \oplus) in real time adds up to 28 247. When we add \mathcal{S} to this portfolio we solve 28 271 problems. This shows that the ENIGMA strategies learned quite well from \mathcal{S} , not losing many solutions. When we add eight more strategies developed here we solve 29 130 problems, of which 2109 are among the hard Mizar40. This is done in general in 200s and without any additional help from premise selection methods. Vampire in 300s solves 27 842 problems. Future work includes joint evaluation of the system on problems translated from different ITP libraries, similar to [9].

Acknowledgments. We thank Stephan Schulz and Thibault Gauthier for discussing with us their methods for symbol-independent term and formula matching.

A Additional Data From the Experiments

This appendix presents additional data from the experiments in Section 6. Figure 3 shows the results of the grid search for GNN models on one tenth of all benchmark problems done in order to find the best-performing parameters for *query* and *context* sizes. The x -axis plots the query size, the y -axis plots the context size, while the z -axis plots the ATP performance, that is, the number of solved problems. Recall that the grid search was performed on a randomly selected but fixed tenth of all benchmark problems with a 10s real-time limit per problem. For \mathcal{N}_0 and \mathcal{N}_1 , there is a separate graph for each iteration, showing only the best epochs. For \mathcal{N}_2 , there are two graphs for models from epoch 20 and 50. Note how the later epoch 50 becomes more independent on the context size. The ranges of the grid search parameters were extended in later iterations when the best-performing value was at the graph edge.

Figure 4 shows the grid search results for the best LightGBM’s GBDT models from iterations 1, 2, and 3 (denoted here \mathcal{D}_0 , \mathcal{D}_1 , and \mathcal{D}_2). The x -axis plots the number of tree leaves, the y -axis plots the tree depth, while the z -axis plots the number of solved problems. There are two models from the second iteration (\mathcal{D}_1), showing the effect of different learning rate (η). Again, the ranges of meta-parameters were updated in between the iterations by a human engineer.

Figure 5 shows the training accuracies and training loss for the LightGBM model \mathcal{D}_2 . Accuracies (TPR and TNR) of the training data are computed from the first iteration (\mathcal{T}_0). The values for loss (z) are inverted ($1 - z$) so that higher values correspond to better models which makes a visual comparison easier. We can see a clear correlation between the accuracies and the loss, but not so clear correlation with the ATP performance. The ATP performance of \mathcal{D}_2 is the same as in Figure 4, repeated here for convenience.




Fig. 2. Scatter plots for the lengths of the discovered proofs (the x -axis for \mathcal{S} , and the y -axis for $\mathcal{S} \oplus \mathcal{D}_2$ and $\mathcal{S} \oplus \mathcal{N}_2$, respectively).

Figure 2 compares the lengths of the discovered proofs. We can see that there is no systematic difference in this metric between the base strategy and the ENIGMA ones.

Finally, we have compared the feature vectors of the symbol-dependent and symbol-independent versions of the GBDTs. On the same data, we observe roughly 2x more collisions. The symbol-independent version has around 1% of colliding feature vectors, while the symbol-dependent version has 0.42%.

B Discussion of Anonymization

Our use of symbol-independent arity-based features for GBDTs differs from Schulz’s anonymous *clause patterns* [23, 24] (CPs) used in E for proof guidance and from Gauthier and Kaliszyk’s (GK) anonymous abstractions used for their concept alignments between ITP libraries [10] in two ways:

1. In both CP and GK, serial (de Bruijn-style) numbering of abstracted symbols of the same arity is used. I.e., the term $h(g(a))$ will get abstracted to $F11(F12(F01))$. Our encoding is just $F1(F1(F0))$. It is even more lossy, because it is the same for $h(h(a))$.
2. ENIGMA with gradient boosting decision trees (GBDTs) can be (approximately) thought of as implementing weighted feature-based clause classification where the feature weights are learned. Whereas both in CP and GK, exact matching is used after the abstraction is done.⁷ In CP, this is used for hint-style guidance of E. There, for clauses, such serial numbering however isn’t stable under literal reordering and subsumption. Partial heuristics can be used, such as normalization based on a fixed global ordering done in both CP and GK.

Addressing the latter issue (stability under reordering of literals and subsumption) leads to the NP hardness of (hint) matching/subsumption. I.e.,

⁷ We thank Stephan Schulz for pointing out that although CPs used exact matching by default, matching up to a certain depth was also implemented.

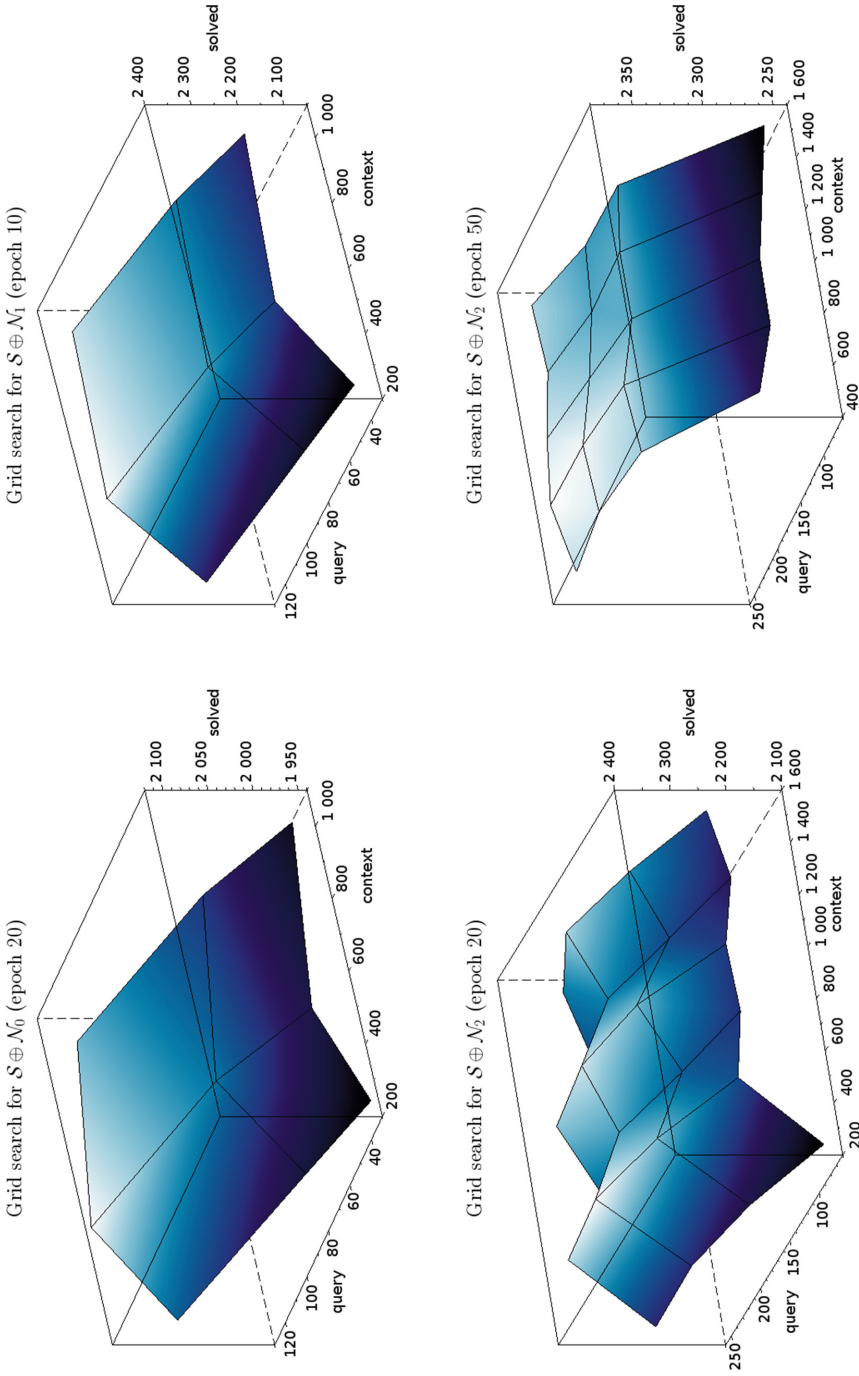


Fig. 3. Grid search results for GNN models (M_i).

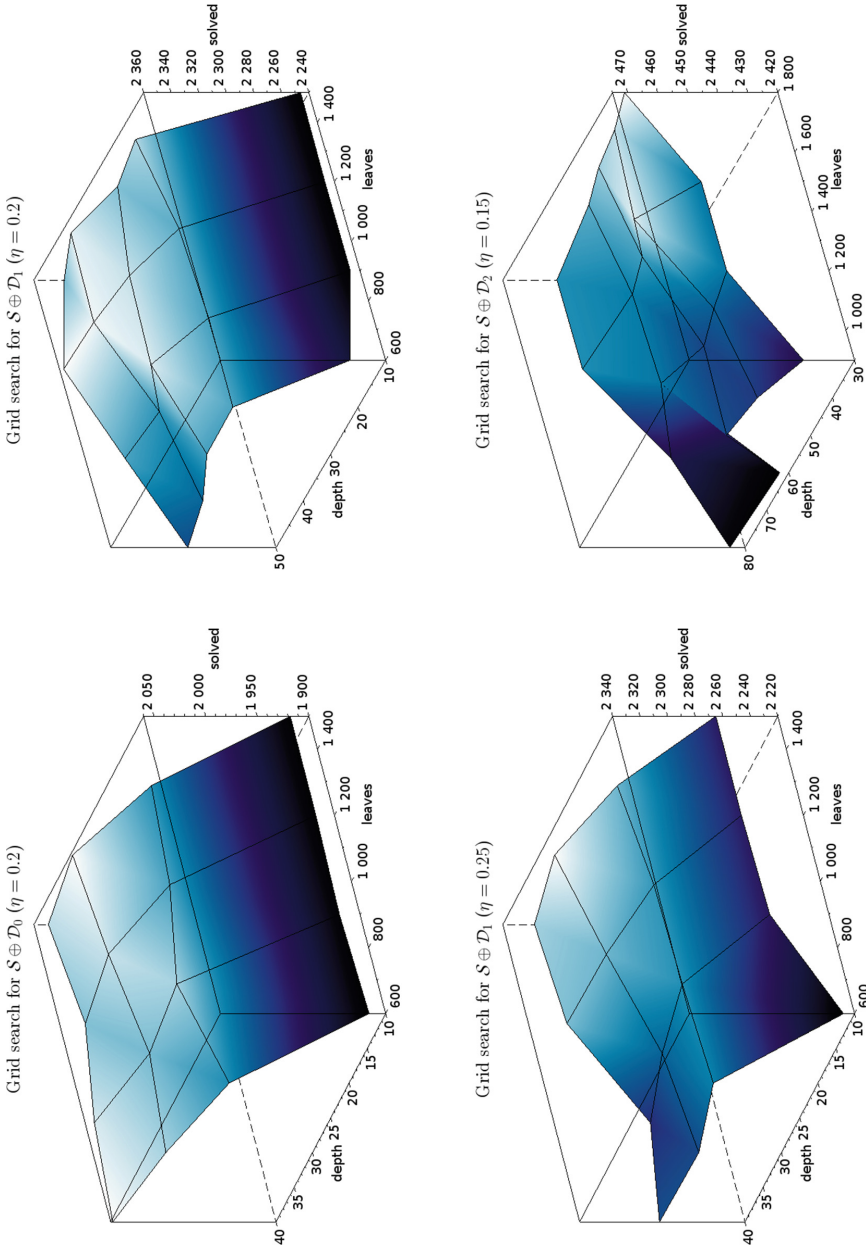


Fig. 4. Grid search results for LightGBM GBDT models (\mathcal{D}_i).




Fig. 5. Accuracies for LightGBM model D_2 .

the abstracted subsumption task can be encoded as standard first-order subsumption for clauses where terms like $F11(F12(F01))$ are encoded as $apply1(X1, apply1(X2, apply0(X3)))$. The NP hardness of subsumption is however here more serious in practice than in standard ATP because only applications behave as non-variable symbols during the matching.

Thus, the difference between our anonymous approach and CP is practically the same as between the standard symbol-based ENIGMA guidance and standard hint-based [30] guidance. In the former the matching (actually, clause classification) is approximate, weighted and learned, while with hints the clause matching/classification is crisp, logic-rooted and preprogrammed, sometimes running into the NP hardness issues. Our latest comparison [12] done over the Mizar/MPTP corpus in the symbol-based setting showed better performance of ENIGMA over using hints, most likely due to better generalization behavior of ENIGMA based on the statistical (GBDT) learning.

Note also that the variable and symbol statistics features to some extent alleviate the conflicts obtained with our encoding. E.g., $h(g(a))$ and $h(h(a))$ will have different *symbol statistics* (Section 3) features. To some extent, such features are similar to Schulz’s feature vector and fingerprint indexing [26, 27].

References

1. Abadi, M., et al.: TensorFlow: large-scale machine learning on heterogeneous systems (2015). <http://tensorflow.org>
2. Blanchette, J.C., Greenaway, D., Kaliszyk, C., Kühlwein, D., Urban, J.: A learning-based fact selector for Isabelle/HOL. *J. Autom. Reasoning* **57**(3), 219–244 (2016). <https://doi.org/10.1007/s10817-016-9362-8>
3. Blanchette, J.C., Kaliszyk, C., Paulson, L.C., Urban, J.: Hammering towards QED. *J. Formalized Reasoning* **9**(1), 101–148 (2016)
4. Chen, T., Guestrin, C.: XGBoost: a scalable tree boosting system. In: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD 2016), pp. 785–794. ACM, New York (2016)
5. Chvalovský, K., Jakubův, J., Suda, M., Urban, J.: ENIGMA-NG: efficient neural and gradient-boosted inference guidance for E. In: Fontaine, P. (ed.) CADE 2019. LNCS (LNAI), vol. 11716, pp. 197–215. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-29436-6_12
6. Czajka, L., Kaliszyk, C.: Hammer for Coq: automation for dependent type theory. *J. Autom. Reasoning* **61**(1–4), 423–453 (2018). <https://doi.org/10.1007/s10817-018-9458-4>
7. Fan, R.-E., Chang, K.-W., Hsieh, C.-J., Wang, X.-R., Lin, C.-J.: Liblinear: a library for large linear classification. *J. Mach. Learn. Res.* **9**, 1871–1874 (2008)
8. Gauthier, T., Kaliszyk, C.: Premise selection and external provers for HOL4. In: Leroy, X., Tiu, A. (eds.) Proceedings of the 2015 Conference on Certified Programs and Proofs (CPP 2015), Mumbai, India, 15–17 January (2015), pp. 49–57. ACM (2015)
9. Gauthier, T., Kaliszyk, C.: Sharing HOL4 and HOL light proof knowledge. In: Davis, M., Fehner, A., McIver, A., Voronkov, A. (eds.) LPAR 2015. LNCS, vol. 9450, pp. 372–386. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-48899-7_26

10. Gauthier, T., Kaliszzyk, C.: Aligning concepts across proof assistant libraries. *J. Symb. Comput.* **90**, 89–123 (2019)
11. Gauthier, T., Kaliszzyk, C., Urban, J.: Initial experiments with statistical conjecturing over large formal corpora. In: Kohlhase, A. (eds.) *Joint Proceedings of the FM4M, MathUI, and ThEdu Workshops, Doctoral Program, and Work in Progress at the Conference on Intelligent Computer Mathematics 2016 co-located with the 9th Conference on Intelligent Computer Mathematics (CICM 2016) of CEUR Workshop Proceedings, Bialystok, Poland, 25–29 July 2016*, vol. 1785, pp. 219–228. CEUR-WS.org (2016)
12. Goertzel, Z., Jakubův, J., Urban, J.: ENIGMAWatch: proofWatch meets ENIGMA. In: Cerrito, S., Popescu, A. (eds.) *TABLEAUX 2019. LNCS (LNAI)*, vol. 11714, pp. 374–388. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-29026-9_21
13. Jakubův, 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
14. Jakubův, 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
15. Jakubův, J., Urban, J.: Hammering Mizar by learning clause guidance. In: Harrison, J., O’Leary, J., Tolmach, A. (eds.) *10th International Conference on Interactive Theorem Proving (ITP 2019) of LIPIcs*, 9–12 September 2019, Portland, OR, USA, vol. 141, pp. 34:1–34:8. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2019)
16. Kaliszzyk, C., Urban, J.: Learning-assisted automated reasoning with FLYSPECK. *J. Autom. Reasoning* **53**(2), 173–213 (2014). <https://doi.org/10.1007/s10817-014-9303-3>
17. Kaliszzyk, C., Urban, J.: HOL(y)Hammer: online ATP service for HOL light. *Math. Comput. Sci.* **9**(1), 5–22 (2015). <https://doi.org/10.1007/s11786-014-0182-0>
18. Kaliszzyk, C., Urban, J.: MizAR 40 for Mizar 40. *J. Autom. Reasoning* **55**(3), 245–256 (2015). <https://doi.org/10.1007/s10817-015-9330-8>
19. Ke, G., et al.: Lightgbm: a highly efficient gradient boosting decision tree. In: *NIPS*, pp. 3146–3154 (2017)
20. 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
21. Olsák, M., Kaliszzyk, C., Urban, J.: Property invariant embedding for automated reasoning. *CoRR*, abs/1911.12073 (2019)
22. Overbeek, R.A.: A new class of automated theorem-proving algorithms. *J. ACM* **21**(2), 191–200 (1974)
23. Schulz, S.: *Learning Search Control Knowledge For Equational Deduction of DISKI*, vol. 230. Infix Akademische Verlagsgesellschaft, Frankfurt (2000)
24. Schulz, S.: Learning search control knowledge for equational theorem proving. In: Baader, F., Brewka, G., Eiter, T. (eds.) *KI 2001. LNCS (LNAI)*, vol. 2174, pp. 320–334. Springer, Heidelberg (2001). https://doi.org/10.1007/3-540-45422-5_23
25. Schulz, S.: E - a brainiac theorem prover. *AI Commun.* **15**(2–3), 111–126 (2002)
26. Schulz, S.: Fingerprint indexing for paramodulation and rewriting. In: Gramlich, B., Miller, D., Sattler, U. (eds.) *IJCAR 2012. LNCS (LNAI)*, vol. 7364, pp. 477–483. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-31365-3_37

27. 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
28. Urban, J.: MPTP 0.2: design, implementation, and initial experiments. *J. Autom. Reasoning* **37**(1–2), 21–43 (2006)
29. 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
30. Veroff, R.: Using hints to increase the effectiveness of an automated reasoning program: case studies. *J. Autom. Reasoning* **16**(3), 223–239 (1996)