



# Learning Theorem Proving Components

Karel Chvalovský<sup>1</sup> , Jan Jakubův<sup>1,2</sup> , Miroslav Olsák<sup>2</sup> ,  
and Josef Urban<sup>1</sup> 

<sup>1</sup> Czech Technical University in Prague, Prague, Czechia  
karel@chvalovsky.cz

<sup>2</sup> University of Innsbruck, Innsbruck, Austria  
mirek@olsak.net

**Abstract.** Saturation-style automated theorem provers (ATPs) based on the given clause procedure are today the strongest general reasoners for classical first-order logic. The clause selection heuristics in such systems are, however, often evaluating clauses in isolation, ignoring other clauses. This has changed recently by equipping the E/ENIGMA system with a graph neural network (GNN) that chooses the next given clause based on its evaluation in the context of previously selected clauses. In this work, we describe several algorithms and experiments with ENIGMA, advancing the idea of contextual evaluation based on learning important components of the graph of clauses.

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

## 1 Introduction: Clause Selection and Context

Clause selection is a crucial part of saturation-style [29] automated theorem provers (ATPs) such as E [32], Vampire [20], and Prover9 [22]. These systems, implementing the given-clause [21] algorithm, provide the strongest methods for proving lemmas in large interactive theorem prover (ITP) libraries [4], and occasionally prove open conjectures in specialized parts of mathematics [19].

Clause selection heuristics have a long history of research, going back to a number of experiments done with the Otter system [24]. Systems such as Prover9 and E have eventually developed extensive domain-specific languages for clause selection heuristics, allowing application of sophisticated algorithms based on a number of different ideas [13, 23, 27, 31, 35] and their automated improvement [15, 30, 34]. These algorithms are, however, often evaluating clauses in isolation, ignoring other clauses selected in the proof search, and thus largely neglecting the notion of a (*proof*) *state* and its obvious importance for choosing the next *action* (clause).

This has changed recently with equipping the E/ENIGMA [6, 14] system with a logic-aware graph neural network (GNN) [25], where the next given clause is

chosen based on its evaluation in the context of previously selected clauses [12]. In more details, in GNN-ENIGMA, the generated clauses are not ranked immediately and independently on other clauses. Instead, they are judged by the GNN in larger batches and with respect to a large number of already selected clauses (*context*). The GNN estimates collectively the most useful subset of the context and new clauses by several rounds of message passing. The message-passing algorithm takes into account the connections between symbols, terms, subterms, atoms, literals, and clauses. It is trained on many previous proof searches, and it estimates which clauses will collectively benefit the proof search in the best way.

In the rest of the paper, we describe several algorithms and experiments with ENIGMA and GNN-based algorithms, advancing the idea of contextual evaluation. In Sect. 2, we give an overview of the learning-based ENIGMA clause selection in E, focusing on the recently added context-based evaluation by GNNs. Section 3 introduces the first variant of our context-based algorithms called *leapfrogging*. These algorithms interleave saturation-style ATP runs with external context-based evaluations and clause filtering. Section 4 introduces the second variant of our context-based algorithms, based on learning from past interactions between clauses and splitting the proof search into separate components. Section 5 discusses technical details, and Sect. 6 evaluates the methods.

## 2 ENIGMA and Learning Context-Based Guidance

This section summarizes our previous Graph Neural Network (GNN) ENIGMA *anonymization* architecture [12], which was previously successfully used for a given clause guidance within E Prover [16]. In this context, anonymization means guidance independent on specific symbol names.

Saturation-based ATPs, such as E, employ a *given-clause loop*. The input first-order logic problem is translated into a refutationally-equivalent set of clauses, and a search for a contradiction is initiated. Starting with the initial set of clauses, one clause is selected (*given*) for processing, and all possible inferences with all previously processed clauses are derived. This extends the set of clauses available for processing, and the loop repeats until (1) the contradiction (empty clause) is derived, or (2) there are no clauses left for processing (that is, the input problem is not provable), or (3) resources (time, memory, or user patience) are exhausted. As the selection of the right clauses for processing is essential for a success, our approach is to guide the clause selection within an ATP by sophisticated machine learning methods.

For the clause selection with ENIGMA Anonymous, we train a GNN classifier for symbol-independent clause embeddings from a large number of previous successful E proof searches. From every successful proof search, we extract the set of all processed clauses, and we label the clauses that appear in the final proof as *positive* while the remaining (unnecessarily processed) clauses as *negative*. These training data are turned into a tensor representation (one- and two-dimensional variable-length vectors), which encapsulate clause syntax trees by abstracting

from specific symbol names while preserving information about symbol relations. Each tensor represents a set of clauses as a graph with three types of nodes (for terms/subterms, clauses, and symbols), and passes initial embeddings through a fixed number of message-passing (graph convolution) layers. Additionally, the conjecture clauses of the problem to be proved are incorporated into the graph to allow for conjecture-dependent clause classification.

Once a GNN classifier is trained from a large number of proof searches, it is utilized in a new proof search to evaluate the clauses to be processed and to select the best given clause as follows. Instead of evaluating the clauses one by one, as is the case in the alternative ENIGMA Anonymous decision tree classifiers, we postpone clause evaluation until a specific number of clauses to be evaluated is collected. These clauses form the *query* part and the size of the query is passed to the prover as a parameter. The query clauses are extended with clauses forming a *context*, that is, a specific number of clauses already processed during the current proof search. In particular, we use the first  $n$  clauses processed during the proof search as the context. The context size  $n$  is another parameter passed to the prover. After adding the conjecture and context clauses to the query, their tensor representation is computed and sent to the GNN for evaluation. The GNN applies several graph convolution (message passing) layers getting an embedding of every clause. Each clause is combined through a single fully connected layer with an embedding of the conjecture, and finally transformed into a single score (*logit*), which is sent back to the prover. The prover then processes the clauses with better (higher) scores in advance. For details, see [12, 25].

### 3 Leapfrogging

The first class of algorithms is based on the idea that the graph-based evaluation of a particular clause may significantly change as new clauses are produced and the context changes. It corresponds to the human-based mathematical exploration, in which initial actions can be done with relatively low confidence and following only uncertain hunches. After some amount of initial exploration is done, clearer patterns often appear, allowing re-evaluation of the approach, focusing on the most promising directions, and discarding of less useful ideas.

In tableau-based provers such as leanCoP [26] with a compact notion of state, such methods can be approximated in a reinforcement learning setting by the notion of *big steps* [18] in the Monte-Carlo tree search (MCTS), implementing the standard *explore/exploit* paradigm [10]. In the saturation setting, our proposed algorithm uses short standard saturation runs at the exploration phase, after which the set of processed (selected) clauses is reevaluated and a decision on its most useful subset is made by the GNN. These two phases are iterated in a procedure that we call *leapfrogging*.

In more detail, leapfrogging is implemented as follows (see also Algorithm 1). Given a clausal problem consisting of a set of initial clauses  $S = S_0$ , an initial saturation-style search (in our case E/ENIGMA) is run on  $S$  with an abstract time limit. We may use a fixed limit (e.g., 1000 nontrivial processed clauses)

---

**Algorithm 1:** The Leapfrogging algorithm with a fixed saturation limit

---

**Input:** *AxiomClauses*, *NegConjectureClauses*, *SaturationLimit*,  
*IterationLimit*, *PremiseSelector*;

```

1  $S_0 = \text{AxiomClauses} \cup \text{NegConjectureClauses}$ ;
2 for  $i = 0$  to  $\text{IterationLimit}$  do
3    $(L_{i+1}, \text{Result}) = \text{Saturate}(S_i, \text{SaturationLimit})$ ;
4   if  $\text{Result} = \text{Unsatisfiable}$  then return  $\text{Unsatisfiable}$ ;
5   else if  $\text{Result} = \text{Satisfiable}$  then
6     if  $i=0$  then return  $\text{Satisfiable}$ ;
7     else return  $\text{Unknown}$ ;
8   else // Result = Unknown
9      $S_{i+1} = \text{PremiseSelector}(L_{i+1}, \text{NegConjectureClauses})$  ;
10     $S_{i+1} = S_{i+1} \cup \text{NegConjectureClauses}$ ;
11 return  $\text{Unknown}$ ;
```

---

for all runs, or change (e.g. increase) the limits gradually. If the initial run results in a proof or saturation within the limit, the algorithm is finished. If not, we inspect the set of clauses created in the run. We can inspect the set of all generated clauses, or a smaller set, such as the set of all processed clauses. So far, we used the latter because it is typically much smaller and better suits our training methods. This (*large*) set is denoted as  $L_0$ . Then we apply a trained graph-based predictor to  $L_0$ , which selects a smaller *most promising* subset of  $L_0$ , denoted as  $S_1$ . We may or may not automatically include also the initial negated conjecture clauses or the whole initial set  $S_0$  in  $S_1$ .  $S_1$  is then used as an input to the next limited saturation run of E/ENIGMA. This process is iterated, producing gradually sets  $S_i$  and  $L_i$ .

A particularly simple version of leapfrogging uses GNN-guided ENIGMA for the saturation “jumps”, and omits the external selection, thus setting  $S_{i+1} := L_i$ . This may seem meaningless with deterministic clause selection heuristics that do not use context: the next saturation run may be selecting the same clauses and ending up with  $S_{i+1} = S_i$ . Already in the standard ATP setting this is, however, easy to make less deterministic, as done, for example, in the randoCoP system [28]. The GNN-guided ENIGMA will typically also make different choices with the new input set  $L_0$  than with the input set  $S_0$ .

A more involved version of leapfrogging, however, makes use of a nontrivial trained graph-based predictor that will reduce  $L_i$  to  $S_{i+1}$  such that  $S_{i+1} \subsetneq L_i$ . For this, we use an external evaluation run of a GNN, which has been trained in the same way as the GNN used inside ENIGMA: on sets of *positive* and *negative* processed clauses extracted from many successful proof runs. Here, the positive clauses are those that end up being part of the proof, and the negative ones are the remaining processed clauses. This is also very similar to an external *premise selection* [2] done with the GNNs [25], with the difference that the inputs are now clauses instead of formulas.

## 4 Learning Reasoning Components

The second class of algorithms is based on learning important components in the graph of clauses. This is again motivated by an analogy with solving mathematical problems, which often have well-separated reasoning and computational components. Examples include numerical calculations, computing derivatives and integrals, performing Boolean algebra in various settings, sequences of standard rewriting and normalization operations in various algebraic theories, etc. Such components of the larger problem can be often solved mostly in isolation from the other components, and only their results are then used together to connect them and solve the larger problem.

Human-designed problem solving architectures addressing such decomposition include, e.g., SMT systems, systems such as MetiTarski [1], and a tactic-based learning-guided proof search in systems such as TacticToe [9]. In all these systems, the component procedures or tactics are, however, *human-designed* and (often painstakingly) human-implemented, with a lot of care both for the components and for the algorithms that merge their results. This approach seems hard to scale to the large number of combinations of complex algorithms, decision procedures and reasoning heuristics used in research-level mathematics, and other complex reasoning domains.

Our new approach is to instead start to *learn* such *targeted components*, expressed as sets of clauses that perform targeted reasoning and computation within the saturation framework. We also want to learn the merging of the results of the components automatically. This is quite ambitious, but there seems to be growing evidence that such targeted components are being learned in many iterations of GNN-guided proving followed by retraining of the GNNs in our recent large iterative evaluation over Mizar.<sup>1</sup> In these experiments, we have significantly extended our previously published results [12],<sup>2</sup> eventually automatically proving 73.5% (more than 40k) of the Mizar theorems. In particular, there are many examples shown on the project Github page demonstrating that the GNN is learning to solve more and more involved computations in problems involving differentiation, integration, boolean algebra, algebraic rewriting, etc. Our initial approach is therefore to (i) use the GNN to learn to identify interacting reasoning components, (ii) use graph-based and clustering-based algorithms to split the set of clauses into components based on the GNN predictions, (iii) run saturation on the components independently, (iv) possibly merge the most important parts of the components, and (v) iterate. See the Split and Merge Algorithm 2.

## 5 Clustering Methods

Here we propose two modifications of our previous GNN architecture, described in Sect. 2, for the identification of interacting reasoning components, and we

<sup>1</sup> [https://github.com/ai4reason/ATP\\_Proofs](https://github.com/ai4reason/ATP_Proofs).

<sup>2</sup> The publication of this large evaluation is in preparation.

**Algorithm 2:** The Split and Merge algorithm

---

**Input:** *AxiomClauses*, *NegConjectureClauses*, *SaturationLimit*,  
*IterationLimit*, *PremiseSelector*, *ClusteringAlgo*;

```

1  $S_0 = \text{AxiomClauses} \cup \text{NegConjectureClauses}$ ;
2 for  $i = 0$  to  $\text{IterationLimit}$  do
3    $(L_i, \text{Result}) = \text{Saturate}(S_i, \text{SaturationLimit})$ ;
4   if  $\text{Result} = \text{Unsatisfiable}$  then return  $\text{Unsatisfiable}$ ;
5   else if  $\text{Result} = \text{Satisfiable}$  then
6     if  $i=0$  then return  $\text{Satisfiable}$ ;
7     else return  $\text{Unknown}$ ;
8   else //  $\text{Result} = \text{Unknown}$ 
9      $(C_i^1, \dots, C_i^K) = \text{ClusteringAlgo}(L_i)$ ; // Split to components
10    for  $j = 1$  to  $K$  do
11       $(L_i^j, \text{Result}^j) = \text{Saturate}(C_i^j, \text{SaturationLimit})$ ; // Run each
12      if  $\text{Result}^j = \text{Unsatisfiable}$  then return  $\text{Unsatisfiable}$ ;
13     $S_{i+1} = \text{PremiseSelector}(\bigcup_{j=1}^K L_i^j, \text{NegConjectureClauses})$ ; // Merge
14     $S_{i+1} = S_{i+1} \cup \text{NegConjectureClauses}$ ;
15 return  $\text{Unknown}$ ;
```

---

describe their intended use. The overall methodology to detect and utilize reasoning components is as follows. To produce the training data, we run E with a fixed limit of  $N$  given clause loops. For each solved problem, we output not only the proof, but the full derivation tree of all clauses *generated* during the proof search. These will provide training data to train a GNN classifier. For unproved problems, we output the  $N$  given clauses *processed* during the search. These data from unsuccessful proof searches are then used for the prediction of interacting components. This is the start of the Split and Merge Algorithm 2.

The training data are extracted from successful proof searches as follows. From each derivation tree, we extract all clause pairs  $C_i$  and  $C_j$  which interacted during the proof search, that is, the pairs which were used to infer another clause. All pairs  $(C_i, C_j)$  which were used to infer a proof clause are marked as *positive* while the remaining clause pairs as *negative*. Such clauses with the information about their positive/negative pairing are used to train a GNN predictor.

The trained GNN predictor will guide the construction of clusters, where clauses resembling positively linked clauses should end up within the same cluster. We obtain the data for predictions from the above unsuccessful proof searches (with the fixed limit of  $N$  processed clauses), and they contain  $N$  processed clauses for every problem. We want to assign each pair of clauses  $(C_i, C_j)$  a score  $l_{i,j}$  which describes the likelihood of inferring a useful clause from  $C_i$  and  $C_j$ . These scores are the basis for the clustering algorithms.

We experiment with two slightly different GNN architectures for the identification of reasoning components. Let  $d$  be the dimension of the final clause

embedding, and let  $c_i, c_j$  be the embeddings of clauses  $C_i, C_j$  respectively. Then the two architectures—differently computing the score  $l_{i,j}$ —are as follows:

1. We pass both  $c_i, c_j$  through a linear layer (with biases, without an activation function) with the output dimension  $n$ , resulting in  $d_i, d_j$ . Then we calculate  $l_{i,j} = d_i \cdot d_j / \sqrt{n}$ .
2. We pass both  $c_i, c_j$  through a linear layer with the output dimension  $2n$ , resulting in  $d'_i, d'_j$ . Then we calculate  $l_{i,j} = d_i \cdot \text{rev}(d_j) / \sqrt{n}$  where  $\text{rev}$  represents reversing the vector.

Mathematically, this corresponds to  $l_{i,j} = c_i^T A c_j + v^T (c_i + c_j) + b$  where  $v^T$  are  $n$ -dimensional vectors for clause evaluation,  $b$  is a scalar bias, and  $A$  is an  $n \times n$  matrix which is symmetric and positive definite in architecture 1, and just symmetric in architecture 2. For training, we pass the value  $l_{i,j}$  through sigmoid and binary cross entropy loss.

## 5.1 Clustering

To split the clauses into separate components, we use standard clustering algorithms. However, in our case, it is likely that some clauses should be shared among various components, and hence we are also interested in methods capable of such overlapping assignments.

Of course, the crucial precondition for splitting clauses into components is defining the similarity between clauses, or even better, a distance between them. We have at least two straightforward options here—to define the distance between two clauses as the distance between their embeddings (vectors) or use the matrix  $L = (l_{ij})$  as a similarity measure, which approximates the likelihood that clauses  $i$  and  $j$  interact in the proof. A simple way to produce distances from  $L$  is to treat each row of  $L$  as a vector and define the distance between two clauses as the (Euclidean) distance between the corresponding rows of  $L$ .

Another approach is to use directly the intended meaning of matrix  $L$ , the likelihood that two given clauses appear in a proof, and to produce a weighted graph from  $L$ , where vertices are clauses and edges are assigned weights according to  $L$ . Moreover, we can remove edges that have weights below some threshold, expressing that such clauses do not interact. In this way, we obtain a weighted graph that can be clustered into components. The following paragraphs briefly describe the clustering algorithms used in the experiments.

**$k$ -Means.** A widely used clustering method is  $k$ -means. The goal is to separate vectors into  $k$  clusters in such a way that their within-cluster variance is minimal. Although  $k$ -means is a popular clustering method, it suffers from numerous well-known problems. For example, it assumes that we know the correct number of clusters in advance, the clusters are of similar sizes, and they are nonoverlapping. Although these assumptions are not satisfied in our case, we used  $k$ -means from SciPY [36] as a well-known baseline.

**Table 1.** Four leapfrogging runs with different GNN-ENIGMAs

GNN-strategy	Original-60s-run	Leapfrogging (300-500-60s)	Union	Added-by-lfrg
$G_1$	2711	2218	3370	659
$G_2$	2516	2426	3393	877
$G_3$	2655	2463	3512	857
$G_4$	2477	2268	3276	799

**Soft  $k$ -Means.** It is possible to modify  $k$ -means in such a way that overlapping clusters (also called soft clusters) are allowed.<sup>3</sup> An example is the Fuzzy C-Means (FCM) algorithm [3] that generalizes  $k$ -means by adding the membership function for each point. This function scores how much each point belongs to a cluster, and it is possible to adjust the degree of overlap between the clusters. We used the `fuzzy-c-means` package [7] for our experiments.

**Graph Clustering.** We have experimented with the cluster application from the popular Graphviz visualisation software [8], which can split a graph into clusters using the methods described in [5]. The graphs are clustered based on the *modularity* measure which considers the density of links inside a cluster compared to links between clusters. It is possible to either directly specify the intended number of clusters (soft constraint), or base the number of clusters on their modularity. We also experimented with clustering using the *modularity quality*.<sup>4</sup> Moreover, by removing some highly connected vertices (clauses) before clustering and adding them into all clusters, we can produce overlapping clusters.

## 6 Evaluation

### 6.1 Leapfrogging

The first leapfrogging experiment is done as follows:

1. We stop GNN-ENIGMA after 300 processed clauses and print them.
2. We restart with the 300 clauses used as input, stop at 500 clauses and print the 500 clauses.
3. We restart with the 500 clauses, and do a final run for 60s.

This is done on a set of 28k *hard* Mizar problems that we have been trying to prove with many different methods in a large ongoing evaluation over the full Mizar corpus.<sup>5</sup> We try with four differently trained and parameterized GNNs, denoted as  $G_1, \dots, G_4$ . The summary of the runs is given in Table 1.

<sup>3</sup> Another popular way how to generalize  $k$ -means (and assign a point to more than one cluster) is to use Gaussian mixture models.

<sup>4</sup> <https://gitlab.com/graphviz/graphviz/-/blob/main/lib/sparse/mq.h>.

<sup>5</sup> Details are at [https://github.com/ai4reason/ATP\\_Proofs](https://github.com/ai4reason/ATP_Proofs).

**Table 2.** Clustering 3000 problems for evaluation

Method	#clusters	Newly solved problems
<i>k</i> -means	2	67
<i>k</i> -means	3	78
Soft <i>k</i> -means	2	63
Soft <i>k</i> -means	3	93
Graphviz	≤4	111

We see that the methods indeed achieve high complementarity to the original GNN strategies. This is most likely thanks to the different context in which the GNN sees the initial clauses in the subsequent runs.

### 6.2 Splitting and Merging

The initial experimental evaluation<sup>6</sup> is done on a large benchmark of 57880 Mizar40 [17] problems<sup>7</sup> exported to first-order logic by MPTP [33]. We use a subset of 52k Mizar40 [17] problems for training. To produce the training data, we run E with a well-performing GNN guidance, and with the limit of 1000 given clause loops. Within this limit, around 20k of the training problems are solved. For the 32k unproved training problems, we output the 1000 given clauses *processed* during the search. As described in Sect. 5, we train a

```

::: 16 f.x=sin.x+x^1/2
theorem FDIFF_5:16
for Z being open Subset of REAL st Z c= dom (sin + (#R (1 / 2))) holds
( sin + (#R (1 / 2)) is differentiable_on Z & ( for x being Real st x in Z holds
((sin + (#R (1 / 2))) ` Z) . x = (cos . x) + ((1 / 2) * (x #R (- (1 / 2)))) ) )
proof
let Z be open Subset of REAL; :: thisis:
assume A1: Z c= dom (sin + (#R (1 / 2))) ; :: thisis:
then Z c= (dom (#R (1 / 2))) /\ (dom sin) by VALUED_1:def 1;
then A2: Z c= dom (#R (1 / 2)) by XBOOLE_1:18;
then A3: #R (1 / 2) is differentiable_on Z by Lm3;
A4: sin is differentiable_on Z by FDIFF_1:26, SIN_COS:68;
now :: thisis:
let x be Real; :: thisis:
assume A5: x in Z ; :: thisis:
then ((sin + (#R (1 / 2))) ` Z) . x = (diff (sin,x) + (diff ((#R (1 / 2)),x))) by A1, A3, A4, FDIFF_1:18
.= (cos . x) + ((diff ((#R (1 / 2)),x))) by SIN_COS:64
.= (cos . x) + ((#R (1 / 2)) ` Z) . x by A3, A5, FDIFF_1:def 7
.= (cos . x) + ((1 / 2) * (x #R (- (1 / 2)))) by A2, A5, Lm3 ;
hence ((sin + (#R (1 / 2))) ` Z) . x = (cos . x) + ((1 / 2) * (x #R (- (1 / 2)))) ; :: thisis:
end;
hence ( sin + (#R (1 / 2)) is differentiable_on Z & ( for x being Real st x in Z holds
((sin + (#R (1 / 2))) ` Z) . x = (cos . x) + ((1 / 2) * (x #R (- (1 / 2)))) ) ) by A1, A3, A4, FDIFF_1:18;
::: thisis:
end;

```

**Fig. 1.** Differentiation – T16.FDIFF\_5

<sup>6</sup> 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.

<sup>7</sup> [http://grid01.ciirc.cvut.cz/~mptp/1147/MPTP2/problems\\_small\\_consist.tar.gz](http://grid01.ciirc.cvut.cz/~mptp/1147/MPTP2/problems_small_consist.tar.gz).

```

theorem : NEWTON:48
for m, n, k being Nat holds m gcd (n gcd k) = (m gcd n) gcd k
proof
let m, n, k be Nat; :: thesis:
set M = n gcd k;
set K = m gcd (n gcd k);
set N = m gcd n;
set L = (m gcd n) gcd k;
A1: m gcd (n gcd k) divides n gcd k by NAT_D:5;
A2: m gcd (n gcd k) divides m by NAT_D:5;
n gcd k divides n by NAT_D:5;
then m gcd (n gcd k) divides n by A1, NAT_D:4;
then A3: m gcd (n gcd k) divides m gcd n by A2, NAT_D:5;
A4: (m gcd n) gcd k divides m gcd n by NAT_D:5;
A5: (m gcd n) gcd k divides k by NAT_D:5;
m gcd n divides n by NAT_D:5;
then (m gcd n) gcd k divides n by A4, NAT_D:4;
then A6: (m gcd n) gcd k divides n gcd k by A5, NAT_D:5;
m gcd n divides m by NAT_D:5;
then (m gcd n) gcd k divides m by A4, NAT_D:4;
then A7: (m gcd n) gcd k divides m gcd (n gcd k) by A6, NAT_D:5;
n gcd k divides k by NAT_D:5;
then m gcd (n gcd k) divides k by A1, NAT_D:4;
then m gcd (n gcd k) divides (m gcd n) gcd k by A3, NAT_D:5;
hence m gcd (n gcd k) = (m gcd n) gcd k by A7, NAT_D:5; :: thesis:
end;

```

Fig. 2. Associativity of gcd by many rewrites – T48\_NEWTON

GNN predictor on the 20k successful runs, and use it to predict the interactions between the processed clauses of the unsuccessful runs. Since the evaluation on the full set of the 32k unsolved problems would be too resource-intensive, we limit this to its randomly chosen 3000-big subset. Table 2 shows the performance of the clustering methods in solving the problems in the first Split phase. The strongest method is the Graphviz-based graph clustering. In more detail, the `cluster` tool gives us on the GNN graph predictions up to four graph components. We run again with a 1000-given clause limit on them newly solving altogether 111 problems inside the components of the 3000. Then we choose this clustering for an experiment with the Merge phase. We merge the components of the remaining unsolved 2889 problems and use our GNN for a premise-selection-style final choice of the jointly best subset of the clauses produced by all the components (line 13 of Algorithm 2). We use four thresholds for the premise selection, and run again with a 1000-given clause limit on each of such premise selections (line 3 of Algorithm 2). This run on the merged components yields another 66 new proofs. Many of the newly found proofs indeed show frequent computational patterns. Examples include the proofs of Mizar problems T16\_FDIFF\_5 (Fig. 1),<sup>8</sup> T48\_NEWTON (Fig. 2),<sup>9</sup> T10\_MATRIX\_4,<sup>10</sup> T11\_VECTSP\_2,<sup>11</sup> T125\_RVSUM\_1,<sup>12</sup> T13\_BICIALG\_3,<sup>13</sup> and T14\_FUZZY\_2.<sup>14</sup>

<sup>8</sup> [http://grid01.ciirc.cvut.cz/~mptp/7.13.01\\_4.181.1147/html/fdiff\\_5.html#T16](http://grid01.ciirc.cvut.cz/~mptp/7.13.01_4.181.1147/html/fdiff_5.html#T16).

<sup>9</sup> [http://grid01.ciirc.cvut.cz/~mptp/7.13.01\\_4.181.1147/html/newton.html#T48](http://grid01.ciirc.cvut.cz/~mptp/7.13.01_4.181.1147/html/newton.html#T48).

<sup>10</sup> [http://grid01.ciirc.cvut.cz/~mptp/7.13.01\\_4.181.1147/html/matrix\\_4.html#T10](http://grid01.ciirc.cvut.cz/~mptp/7.13.01_4.181.1147/html/matrix_4.html#T10).

<sup>11</sup> [http://grid01.ciirc.cvut.cz/~mptp/7.13.01\\_4.181.1147/html/vectsp\\_2.html#T11](http://grid01.ciirc.cvut.cz/~mptp/7.13.01_4.181.1147/html/vectsp_2.html#T11).

<sup>12</sup> [http://grid01.ciirc.cvut.cz/~mptp/7.13.01\\_4.181.1147/html/rvsum\\_1.html#T125](http://grid01.ciirc.cvut.cz/~mptp/7.13.01_4.181.1147/html/rvsum_1.html#T125).

<sup>13</sup> [http://grid01.ciirc.cvut.cz/~mptp/7.13.01\\_4.181.1147/html/bcialg\\_3.html#T13](http://grid01.ciirc.cvut.cz/~mptp/7.13.01_4.181.1147/html/bcialg_3.html#T13).

<sup>14</sup> [http://grid01.ciirc.cvut.cz/~mptp/7.13.01\\_4.181.1147/html/fuzzy\\_2.html#T14](http://grid01.ciirc.cvut.cz/~mptp/7.13.01_4.181.1147/html/fuzzy_2.html#T14).

## 7 Conclusion

We have described several algorithms advancing the idea of contextual evaluation based on learning important components of the graph of clauses. The first leapfrogging experiments already show very encouraging results on the Mizar dataset, providing many complementary solutions. The component-based algorithm also produces new solutions and there are clearly many further methods and experiments that can be tried in this setting. We believe that this approach may eventually lead to using large mathematical libraries for automated learning of nontrivial components, algorithms, and decision procedures involved in mathematical reasoning.

**Acknowledgments.** Supported by the ERC Consolidator grant *AI4REASON* no. 649043 (JJ, JU), by the Czech project AI & Reasoning CZ.02.1.01/0.0/0.0/15\_003/0000466 and the European Regional Development Fund (KC, JU), by the ERC Starting grant *SMART* no. 714034 (JJ, MO), and by the Czech MEYS under the ERC CZ project *POSTMAN* no. LL1902 (JJ).

## References

1. Akbarpour, B., Paulson, L.C.: MetiTarski: an automatic theorem prover for real-valued special functions. *J. Autom. Reasoning* **44**(3), 175–205 (2010)
2. Alama, J., Heskens, T., Kühlwein, D., Tsvitsovadze, E., Urban, J.: Premise selection for mathematics by corpus analysis and kernel methods. *J. Autom. Reason.* **52**(2), 191–213 (2013). <https://doi.org/10.1007/s10817-013-9286-5>
3. Bezdek, J.C., Ehrlich, R., Full, W.: FCM: the fuzzy *c*-means clustering algorithm. *Comput. Geosci.* **10**(2), 191–203 (1984). [https://doi.org/10.1016/0098-3004\(84\)90020-7](https://doi.org/10.1016/0098-3004(84)90020-7). <https://www.sciencedirect.com/science/article/pii/0098300484900207>
4. Blanchette, J.C., Kaliszky, C., Paulson, L.C., Urban, J.: Hammering towards QED. *J. Formalized Reason.* **9**(1), 101–148 (2016). <https://doi.org/10.6092/issn.1972-5787/4593>. <http://dx.doi.org/10.6092/issn.1972-5787/4593>
5. Blondel, V.D., Guillaume, J.L., Lambiotte, R., Lefebvre, E.: Fast unfolding of communities in large networks. *J. Stat. Mech: Theory Exp.* **2008**(10), P10008 (2008)
6. 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](https://doi.org/10.1007/978-3-030-29436-6_12)
7. Dias, M.L.D.: fuzzy-*c*-means: an implementation of fuzzy *c*-means clustering algorithm. (2019). <https://doi.org/10.5281/zenodo.3066222>. <https://git.io/fuzzy-c-means>
8. Ellson, J., Gansner, E.R., Koutsofios, E., North, S.C., Woodhull, G.: Graphviz and dynagraph - static and dynamic graph drawing tools. In: Jünger, M., Mutzel, P. (eds.) *Graph Drawing Software*, pp. 127–148. Springer, Heidelberg (2004). [https://doi.org/10.1007/978-3-642-18638-7\\_6](https://doi.org/10.1007/978-3-642-18638-7_6)

9. Gauthier, T., Kaliszyk, C., Urban, J.: TacticToe: learning to reason with HOL4 tactics. 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. 125–143. EasyChair (2017). <http://www.easychair.org/publications/paper/340355>
10. Gittins, J.C.: Bandit processes and dynamic allocation indices. *J. Roy. Stat. Soc. Ser. B (Methodol.)* 148–177 (1979)
11. Gottlob, G., Sutcliffe, G., Voronkov, A. (eds.): Global Conference on Artificial Intelligence, GCAI 2015, Tbilisi, Georgia, 16–19 October 2015. EPiC Series in Computing, vol. 36. EasyChair (2015). <http://www.easychair.org/publications/volume/GCAI2015>
12. Jakubův, J., Chvalovský, K., Olšák, M., Piotrowski, B., Suda, M., Urban, J.: ENIGMA anonymous: symbol-independent inference guiding machine (system description). In: Peltier, N., Sofronie-Stokkermans, V. (eds.) IJCAR 2020. LNCS (LNAI), vol. 12167, pp. 448–463. Springer, Cham (2020). [https://doi.org/10.1007/978-3-030-51054-1\\_29](https://doi.org/10.1007/978-3-030-51054-1_29)
13. Jakubův, J., Urban, J.: Extending E prover with similarity based clause selection strategies. In: Kohlhase, M., Johansson, M., Miller, B., de Moura, L., Tompa, F. (eds.) CICM 2016. LNCS (LNAI), vol. 9791, pp. 151–156. Springer, Cham (2016). [https://doi.org/10.1007/978-3-319-42547-4\\_11](https://doi.org/10.1007/978-3-319-42547-4_11)
14. 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](https://doi.org/10.1007/978-3-319-62075-6_20)
15. Jakubův, J., Urban, J.: Hierarchical invention of theorem proving strategies. *AI Commun.* **31**(3), 237–250 (2018). <https://doi.org/10.3233/AIC-180761>
16. 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, 9–12 September 2019, Portland, OR, USA. LIPIcs, vol. 141, pp. 34:1–34:8. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2019). <https://doi.org/10.4230/LIPIcs.ITP.2019.34>
17. Kaliszyk, C., Urban, J.: MizAR 40 for Mizar 40. *J. Autom. Reason.* **55**(3), 245–256 (2015). <https://doi.org/10.1007/s10817-015-9330-8>
18. Kaliszyk, C., Urban, J., Michalewski, H., Olšák, M.O.: Reinforcement learning of theorem proving. In: Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, 3–8 December 2018, Montréal, Canada, pp. 8836–8847 (2018). <http://papers.nips.cc/paper/8098-reinforcement-learning-of-theorem-proving>
19. 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)
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](https://doi.org/10.1007/978-3-642-39799-8_1)
21. McCune, W.: Otter 2.0. In: Stickel, M.E. (ed.) CADE 1990. LNCS, vol. 449, pp. 663–664. Springer, Heidelberg (1990). [https://doi.org/10.1007/3-540-52885-7\\_131](https://doi.org/10.1007/3-540-52885-7_131)
22. McCune, W.: Prover9 and Mace4 (2005–2010). <http://www.cs.unm.edu/~mccune/prover9/>

23. McCune, W.: Semantic guidance for saturation provers. In: Calmet, J., Ida, T., Wang, D. (eds.) AISC 2006. LNCS (LNAI), vol. 4120, pp. 18–24. Springer, Heidelberg (2006). [https://doi.org/10.1007/11856290\\_4](https://doi.org/10.1007/11856290_4)
24. McCune, W.: Otter 3.3 reference manual. Technical report ANL/MS-C-TM-263, Argonne National Laboratory, Argonne, USA (2003)
25. Olšák, M., Kaliszyk, C., Urban, J.: Property invariant embedding for automated reasoning. In: Giacomo, G.D., et al. (eds.) ECAI 2020–24th European Conference on Artificial Intelligence, 29 August–8 September 2020, Santiago de Compostela, Spain, 29 August–8 September 2020 - Including 10th Conference on Prestigious Applications of Artificial Intelligence (PAIS 2020). *Frontiers in Artificial Intelligence and Applications*, vol. 325, pp. 1395–1402. IOS Press (2020). <https://doi.org/10.3233/FAIA200244>
26. Otten, J., Bibel, W.: leanCoP: lean connection-based theorem proving. *J. Symb. Comput.* **36**(1–2), 139–161 (2003)
27. Quaipe, A.: *Automated Development of Fundamental Mathematical Theories*. Kluwer Academic Publishers (1992)
28. Raths, T., Otten, J.: randocop: randomizing the proof search order in the connection calculus. In: Konev, B., Schmidt, R.A., Schulz, S. (eds.) *Proceedings of the First International Workshop on Practical Aspects of Automated Reasoning*, Sydney, Australia, 10–11 August 2008. *CEUR Workshop Proceedings*, vol. 373. CEUR-WS.org (2008). <http://ceur-ws.org/Vol-373/paper-08.pdf>
29. Robinson, J.A., Voronkov, A. (eds.): *Handbook of Automated Reasoning* (in 2 volumes). Elsevier and MIT Press (2001)
30. Schäfer, S., Schulz, S.: Breeding theorem proving heuristics with genetic algorithms. In: Gottlob et al. [11], pp. 263–274. [http://www.easychair.org/publications/paper/Breeding\\_Theorem\\_Proving\\_Heuristics\\_with\\_Genetic\\_Algorithms](http://www.easychair.org/publications/paper/Breeding_Theorem_Proving_Heuristics_with_Genetic_Algorithms)
31. Schulz, S.: E - a brainiac theorem prover. *AI Commun.* **15**(2–3), 111–126 (2002)
32. 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)
33. Urban, J.: MPTP 0.2: design, implementation, and initial experiments. *J. Autom. Reasoning* **37**(1–2), 21–43 (2006)
34. Urban, J.: BliStr: the blind strategymaker. In: Gottlob et al. [11], pp. 312–319. [http://www.easychair.org/publications/paper/BliStr\\_The\\_Blind\\_Strategymaker](http://www.easychair.org/publications/paper/BliStr_The_Blind_Strategymaker)
35. Veroff, R.: Using hints to increase the effectiveness of an automated reasoning program: case studies. *J. Autom. Reason.* **16**(3), 223–239 (1996)
36. Virtanen, P., et al.: SciPy 1.0 contributors: SciPy 1.0: fundamental algorithms for scientific computing in python. *Nat. Methods* **17**, 261–272 (2020). <https://doi.org/10.1038/s41592-019-0686-2>