Improving ENIGMA-style Clause Selection while Learning From History

Martin Suda

Czech Technical University in Prague, Prague, Czech Republic
`martin.suda@cvut.cz`

Abstract. We re-examine the topic of machine-learned clause selection guidance in saturation-based theorem provers. The central idea, recently popularized by the ENIGMA system, is to learn a classifier for recognizing clauses that appeared in previously discovered proofs. In subsequent runs, clauses classified positively are prioritized for selection. We propose several improvements to this approach and experimentally confirm their viability. For the demonstration, we use a recursive neural network to classify clauses based on their derivation history and the presence or absence of automatically supplied theory axioms therein. The automatic theorem prover Vampire guided by the network achieves a 41 % improvement on a relevant subset of SMT-LIB in a real time evaluation.

Keywords: Saturation-based theorem proving · Clause Selection · Machine Learning · Recursive Neural Networks.

1 Introduction

The idea to improve the performance of saturation-based automatic theorem provers (ATPs) with the help of machine learning (ML), while going back at least to the early work of Schulz [8, 30], has recently been enjoying a renewed interest. Most notable is the ENIGMA system [16, 17] extending the ATP E [31] by machine learned clause selection guidance. The architecture trains a binary classifier for recognizing as positive those clauses that appeared in previously discovered proofs and as negative the remaining selected ones. In subsequent runs, clauses classified positively are prioritized for selection.

A system such as ENIGMA needs to carefully balance the expressive power of the used ML model with the time it takes to evaluate its advice. For example, Loos et al. [22], who were the first to integrate state-of-the-art neural networks with E, discovered their models to be too slow to simply replace the traditional clause selection mechanism. In the meantime, the data-hungry deep learning approaches motivate researchers to augment training data with artificially crafted theorems [1]. Yet another interesting aspect is what features we allow the model to learn from. One could speculate that the recent success of ENIGMA on the Mizar dataset [7, 18] can at least partially be explained by the involved problems sharing a common source and encoding. It is still open whether some new form of general “theorem proving knowledge” could be learned to improve the performance of an ATP across, e.g., the very diverse TPTP library.

In this paper, we propose several improvements to ENIGMA-style clause selection guidance and experimentally test their viability in a novel setting:

- We lay out a set of possibilities for integrating the learned advice into the ATP and single out the recently developed layered clause selection [10, 11, 36] as particularly suitable for the task.
- We speed up evaluation by a new lazy evaluation scheme under which many generated clauses need not be evaluated by the potentially slow classifier.
- We demonstrate the importance of “positive bias”, i.e., of tuning the classifier to rather err on the side of false positives than on the side of false negatives.
- Finally, we propose the use of “negative mining” for improving learning from proofs obtained while relying on previously learned guidance.

To test these ideas, we designed a recursive neural network to classify clauses based solely on their derivation history and the presence or absence of automatically supplied theory axioms therein. This allows us to test here, as a byproduct of the conducted experiments, whether the human-engineered heuristic for controlling the amount of theory reasoning presented in our previous work [11] can be matched or even overcome by the automatically discovered neural guidance.

The rest of the paper is structured as follows. Sect. 2 recalls the necessary ATP theory, explains clause selection and how to improve it using ML. Sect. 3 covers layered clause selection and the new lazy evaluation scheme. In Sect. 4, we describe our neural architecture and in Sect. 5 we bring everything together and evaluate the presented ideas, using the prover Vampire as our workhorse and a relevant subset of SMT-LIB as the testing grounds. Finally, Sect. 6 concludes.

2 ATPs, Clause Selection, and Machine Learning

The technology behind the modern automatic theorem provers (ATPs) for first-order logic (FOL), such as E [31], SPASS [40], or Vampire [21], can be roughly outlined by using the following three adjectives.

Refutational: The task of the prover is to check whether a given conjecture G logically follows from given axioms $A_1, \dots, A_n$, i.e. whether

$$A_1, \dots, A_n \models G, \tag{1}$$

where G and each A_i are FOL formulas. The prover starts by negating the conjecture G and transforming $\neg G, A_1, \dots, A_n$ into an equisatisfiable set of clauses $\mathcal{C}$. It then applies a sound logical calculus to iteratively derive further clauses, logical consequence of $\mathcal{C}$, until the obvious contradiction in the form of the empty clause $\perp$ is derived. This *refutes* the assumption that $\neg G, A_1, \dots, A_n$ could be satisfiable and thus confirms (1).

Superposition-based: The most popular calculus used in this context is superposition [3, 23], an extension of ordered resolution [4] with a built-in support for handling equality. It consists of several inference rules, such as the resolution rule, factoring, subsumption, superposition, or demodulation.

Inference rules in general determine how to derive new clauses from old ones, where by old clauses we mean either the initial clauses $\mathcal{C}$ or clauses derived previously. The clauses that need to be present for a rule to be applicable are called the *premises* and the newly derived clause is called the *conclusion*. By applying the inference rules the prover gradually constructs a *derivation*, a directed acyclic (hyper-)graph (DAG), with the initial clauses forming the leaves and the derived clauses (labeled by the respective applied rules) forming the internal nodes. A *proof* is the smallest sub-DAG of a derivation containing the final empty clause and for every derived clause the corresponding inference and its premises.

Saturation-based: A saturation algorithm is the concrete way of organizing the process of deriving new clauses, such that every applicable inference is eventually considered. Modern saturation-based ATPs employ some variant of the *given-clause algorithm*, in which clauses are selected for inferences one by one [27].

The process employs two sets of clauses, often called the *active* set $\mathcal{A}$ and the *passive* set $\mathcal{P}$. At the beginning all the initial clauses are put to the passive set. Then in every iteration, the prover *selects* and removes a clause C from $\mathcal{P}$, inserts it into $\mathcal{A}$, and performs all the applicable inferences with premises in $\mathcal{A}$ such that at least one of the premises is C . The conclusions of these inferences are then inserted into $\mathcal{P}$. This way the prover maintains (at the end of each iteration) the invariant that inferences among the clauses in the active set have been performed. The selected clause C is sometimes also called the “given clause”.

During a typical prover run, $\mathcal{P}$ grows much faster than $\mathcal{A}$ (the growth is roughly quadratic). Analogously, although for different reasons, when a proof is discovered, its clauses constitute only a fraction of $\mathcal{A}$. Notice that every clause $C \in \mathcal{A}$ that is in the end *not* part of the proof did not need to be selected and represents a wasted effort. This explains why *clause selection*, i.e. the procedure for picking in each iteration the next clause to process, is one of the main heuristic decision points in the prover, which hugely affects its performance [32].

2.1 Traditional Approaches to Clause Selection

There are two basic criteria that have been identified as generally correlating with the likelihood of a clause contributing to the yet-to-be discovered proof.

One is clause’s *age* or, more precisely, its “date of birth”, typically implemented as an ever increasing timestamp. Preferring for selection old clauses to more recently derived ones corresponds to a breadth-first strategy and ensures fairness. The other criterion is clause’s size, referred to as *weight* in the ATP lingo, and is realized by some form of symbol counting. Preferring for selection small clauses to large ones is a greedy strategy, based on the observation that small conclusions typically belong to inferences with small premises and that the ultimate conclusion—the empty clause—is the smallest of all. The best results are achieved when these two criteria (or their variations) are combined [32].

To implement efficient clause selection by numerical criteria such as age and weight, an ATP represents the passive set $\mathcal{P}$ as a set of priority queues. A queue contains (pointers to) the clauses in $\mathcal{P}$ ordered by its respective criterion. Selection typically alternates between the available queues under a certain ratio.

A successful strategy is, for instance, to select 10 clauses by weight for every clause selected by age, i.e., with an *age-to-weight* ratio of 1:10.

2.2 ENIGMA-style Machine-Learned Clause Selection Guidance

The idea to improve clause selection by learning from previous prover experience goes, to the best of our knowledge, back to Schulz [8, 30] and has more recently been successfully employed by the ENIGMA system and others [7, 15–17, 22].

The experience is collected from successful prover runs, where each selected clause constitutes a training example and the example is marked as *positive*, if the clause ended-up in the discovered proof, and *negative* otherwise. A machine learning (ML) algorithm is then used to *fit* this data and produce a *model* $\mathcal{M}$ for *classifying* clauses into positive and negative, accordingly. A good learning algorithm produces a model $\mathcal{M}$ which not only accurately classifies the training data but also *generalizes* well to unseen examples. The computational costs of both training and evaluation are also important.

While clauses are logical formulas, i.e., discrete objects forming a countable set, ML algorithms, rooted in mathematical statistics, are primarily equipped to dealing with fixed-sized real-valued vectors. Thus the question of how to *represent* clauses for the learning is the first obstacle that needs to be overcome, before the whole idea can be made to work. In the beginning, the authors of ENIGMA experimented with various forms of hand-crafted numerical clause *features* [16, 17]. An attractive alternative explored in later work [7, 15, 22] is the use of artificial *neural networks*, which can be understood as extracting the most relevant features automatically.

An important distinction can in both cases be made between approaches which have access to the concrete identity of predicate and function symbols (i.e., the signature) that make up the clauses, and those that do not. For example: Is the ML algorithm allowed to assume that the symbol `grp_mult` is used to represent the multiplication operation in a group or does it only recognize a general binary function? The first option can be much more powerful, but we need to ensure that the signature symbols are *aligned* and used consistently across the problems in our benchmark. Otherwise the learned advice cannot meaningfully carry over to previously unsolved problems. While the assumption of aligned signature has been employed by the early systems [16, 22], the most recent version of ENIGMA [15, 24] can work in a “signature agnostic” mode.

In this work we represent clauses solely by their derivation history, deliberately ignoring their logical content. Thus we do not require the assumption of an aligned signature, per se. However, we rely on a fixed set of distinguished axioms to supply features in the derivation leaves.

2.3 Integrating the Learned Advice

Once we have a trained model $\mathcal{M}$, an immediate possibility for integrating it into the clause selection procedure is to introduce a new queue that will order the clauses using $\mathcal{M}$. Two basic versions of this idea have been described:

“Priority”: The ordering puts all the clauses classified by $\mathcal{M}$ as positive before those classified negatively. Within the two classes, older clauses are preferred.

Let us for the purposes of future reference denote this scheme $\mathcal{M}^{1,0}$. It has been successfully used by the early ENIGMAs [7, 16, 17].

“Logits”: Even models officially described as binary classifiers typically internally compute a real-valued estimate L of how much “positive” or “negative” an example appears to be and only turn this estimate into a binary decision in the last step, by comparing it against a fixed threshold t , most often 0. A machine learning term for this estimate L is the *logit*.¹

The second version orders the clauses on the new queue by the “raw” logits produced by a model. We denote it $\mathcal{M}^{-\mathbb{R}}$ to stress that clauses with high L are treated as small from the perspective of the selection and therefore preferred. This scheme has been used by Loos et al. [22] and in the latest ENGIMA [15, 37].

Combining with a traditional strategy. While it is possible to rely exclusively on selection governed by the model, it turns out to be better [7] to combine it with the traditional heuristics. The most natural choice is to take $\mathcal{S}$, the original strategy that was used to generate the training data, and extend it by adding the new queue, be it $\mathcal{M}^{1,0}$ or $\mathcal{M}^{-\mathbb{R}}$, next to the already present queues. We then again supply a ratio under which the original selection from $\mathcal{S}$ and the new selection based on $\mathcal{M}$ get alternated. We will denote this kind of combination with the original strategy as $\mathcal{S} \oplus \mathcal{M}^{1,0}$ and $\mathcal{S} \oplus \mathcal{M}^{-\mathbb{R}}$, respectively.

3 Layered Clause Selection and Lazy Model Evaluation

Layered clause selection (LCS) is a recently developed method [10, 11, 36] for smoothly incorporating a categorical preference for certain clauses into a base clause selection strategy $\mathcal{S}$. In this paper, we will readily use it in combination with the binary classifier advice from a trained model $\mathcal{M}$.

When we instantiate LCS to our particular case,² its function can be summarized by the expression

$$\mathcal{S} \oplus \mathcal{S}[\mathcal{M}^1].$$

In words, the base selection strategy $\mathcal{S}$ is alternated with $\mathcal{S}[\mathcal{M}^1]$, the same selection scheme $\mathcal{S}$ but applied only to clauses classified positively by $\mathcal{M}$. Implicit here is a convention that whenever there is no positively classified passive clause, a fallback to plain $\mathcal{S}$ occurs. Additionally, we again specify a “second-level” ratio to govern the alternation between pure $\mathcal{S}$ and $\mathcal{S}[\mathcal{M}^1]$.

The main advantage of LCS, compared to the options outlined in the previous section, is that the original, typically well-tuned, base selection mechanism $\mathcal{S}$ is also applied to $\mathcal{M}^1$, the clauses classified positively by $\mathcal{M}$.

¹ A logit can be turned into a (formal) probability, i.e. a value between 0 and 1, by passing it, as is typically done, through the *sigmoid* function $\sigma(x) = 1/(1 + e^{-x})$.

² We rely here on the *monotone* mode of split; there is also a *disjoint* mode [10].

3.1 Lazy Model Evaluation

It is often the case that evaluating a clause by the model $\mathcal{M}$ is a relatively expensive operation [22]. As we explain here, however, this operation can be avoided in many cases, especially when using LCS to integrate the advice.

We propose the following *lazy evaluation approach* to be used with $\mathcal{S} \oplus \mathcal{S}[\mathcal{M}^1]$. Every clause entering the passive set $\mathcal{P}$ is initially inserted to both $\mathcal{S}$ and $\mathcal{S}[\mathcal{M}^1]$ *without* being evaluated by $\mathcal{M}$. Then, whenever (as governed by the second-level ratio) it is the moment to select a clause from $\mathcal{S}[\mathcal{M}^1]$, the algorithm

1. picks (as usual, according to $\mathcal{S}$) the best clause C in $\mathcal{S}[\mathcal{M}^1]$,
2. only then evaluates C by $\mathcal{M}$, and
3. if C gets classified as negative, it forgets C , a goes back to 1.

This repeats until the first positively classified clause is found, which is then returned. Note that this way the “observable behaviour” of $\mathcal{S}[\mathcal{M}^1]$ is preserved.

The power of lazy evaluation lies in the fact that not every clause needs to be evaluated before a proof is found. Indeed, recall the remark that the passive set $\mathcal{P}$ is typically much larger than the active set $\mathcal{A}$, which also holds on a typical successful termination. Every clause left in passive at that moment is a clause that did not need to be evaluated by $\mathcal{M}$ thanks to lazy evaluation.

We remark that lazy evaluation can similarly be used with the integration mode $\mathcal{M}^{1,0}$ based on priorities.

We experimentally demonstrate the effect of the technique in Sect. 5.4.

4 A Neural Classification of Clause Derivations

In this work we choose to represent a clause, for the purpose of learning, solely by its derivation history. Thus a clause can only be distinguished by the axioms from which it was derived and by the precise way in which these axioms interacted with each other through inferences in the derivation. This means we deliberately ignore the clause’s logical content.

We decided to focus on this representation, because it promises to be fast. Although an individual clause’s derivation history may be large, it is a simple function of its parents’ histories (just one application of an inference rule). Moreover, before a clause with a complicated history can be selected, most of its ancestors will have been selected already.³ This guarantees the amortised cost of evaluating a single clause to be constant.

A second motivation comes from our recent work [11], where we have shown that theory reasoning facilitated by automatically adding theory axioms for axiomatising theories, while in itself a powerful technique, often leads the prover to unpromising parts of the search space. We developed a heuristic for controlling the amount of theory reasoning in the derivation of a clause [11]. Our goal here is to test whether a similar or even stronger heuristic can be automatically discovered by a neural network.

³ Exceptions are caused by simplifying inferences applied eagerly outside of the governance of the main clause selection mechanism.

Examples of axioms that Vampire uses to axiomatise theories include the commutativity or associativity axioms for the arithmetic operations, an axiomatization of the theory of arrays [6] or of the theory of term algebras [20]. For us it is mainly important that the axioms are introduced internally by the prover and can therefore be consistently identified across individual problems.

4.1 Recursive Neural Networks

A recursive neural network (RvNN) is a network created by *recursively* composing a finite set of neural building blocks over a structured input [12]. A general neural block is a function $N_\theta : \mathbb{R}^k \rightarrow \mathbb{R}^l$ depending on a vector of parameters θ that can be optimized during training (see below in Section 4.3).

In our case, the structured input is a clause derivation, i.e. a DAG with nodes identified with the derived clauses. To enable a recursion, an RvNN represents each node C by a real vector v_C (of a fixed dimension n) called a (learnable) *embedding*. During training a network learns to embed the space of derivable clauses into $\mathbb{R}^n$ in some a priori unknown, but still useful way.

We assume that each initial clause C , a leaf of the derivation DAG, is labeled as belonging to one of the automatically added theory axioms or coming from the user input. Let these labels form a finite set of *axiom origin labels* $\mathcal{L}_A$. Furthermore, let the applicable inference rules that label the internal nodes of the DAG form a finite set of *inference rule labels* $\mathcal{L}_R$. The specific building blocks of our neural architecture are the following three (indexed families of) functions:

- for every axiom label $l \in \mathcal{L}_A$, a nullary *init* function $I_l \in \mathbb{R}^n$ which to an initial clause C labeled by l assigns its embedding $v_C := I_l$,
- for every inference rule $r \in \mathcal{L}_R$, a *deriv* function, $D_r : \mathbb{R}^n \times \dots \times \mathbb{R}^n \rightarrow \mathbb{R}^n$ which to a conclusion clause C_c derived by r from premises $(C_1, \dots, C_k)$ with embeddings $v_{C_1}, \dots, v_{C_k}$ assigns the embedding $v_{C_c} := D_r(v_{C_1}, \dots, v_{C_k})$,
- and, finally, a single *eval* function $E : \mathbb{R}^n \rightarrow \mathbb{R}$ which evaluates an embedding v_C such that the corresponding clause C is classified as *positive* whenever $E(v_C) \geq t$, with the threshold t set, by default, to 0.

By recursively composing the init and deriv functions, any derived clause C can be assigned an embedding v_C and also evaluated by E to see whether the network recommends it as positive, that should be preferred in proof search.

4.2 Architecture Details

Here we outline the details of our architecture for the benefit of neural network practitioners. All the used terminology is standard (see, e.g., [13]).

We realized each init function I_l as an independent learnable vector. Similarly, each deriv function D_r was independently defined. For a rule of arity two, such as resolution, we used:

$$D_r(v_1, v_2) = \text{LayerNorm}(y), y = W_2^T \cdot x + b_2^T, x = \text{ReLU}(W_1^T \cdot [v_1, v_2] + b_1^T),$$

where $[\cdot, \cdot]$ denotes vector concatenation, ReLU is the rectified linear unit non-linearity ($f(x) = \max\{0, x\}$) applied component-wise, and the learnable matrices

W_1^r, W_2^r and vectors b_1^r, b_2^r are such that $x \in \mathbb{R}^{2n}$ and $y \in \mathbb{R}^n$. (We took inspiration from Sandler et al. [29] for doubling the embedding size before applying the non-linearity.) Finally, LayerNorm is a *layer normalization* [2] module, without which training often became numerically unstable for deeper derivation DAGs.⁴

For unary inference rules, such as factoring, we used an equation analogous to the above, except for the concatenation operation. We did not need to model an inference rule with a variable number of premises, but one option would be to arbitrarily “bracket” its arguments into a tree of binary applications.

Finally, the eval function was $E(v) = W_2 \cdot \text{ReLU}(W_1 \cdot v + b) + c$ with trainable $W_1 \in \mathbb{R}^{n \times n}, b \in \mathbb{R}^n, W_2 \in \mathbb{R}^{1 \times n}$, and $c \in \mathbb{R}$.

4.3 Training the Network

To train a network means to find values for the trainable parameters such that it accurately classifies the training data and ideally also generalises to unseen future cases. We follow a standard methodology for training our RvNN.

In particular, we use the gradient descent (GD) optimization algorithm (with the Adam optimiser [19]) minimising the typical binary cross-entropy *loss*, composed as a sum of contributions, for every selected clause C , of the form

$$-y_C \cdot \log(\sigma(E(v_C))) - (1 - y_C) \cdot \log(1 - \sigma(E(v_C))),$$

with $y_C = 1$ for the positive and $y_C = 0$ for the negative examples.

These contributions are weighted such that each derivation DAG (corresponding to a prover run on a single problem) receives equal weight. Moreover, within each DAG we re-scale the influence of positive versus the negative examples such that these two categories contribute evenly. The scaling is important as our training data is highly unbalanced (cf. Sect. 5.1).

We split the available successful derivations into a *training* set and a *validation* set, and only train on the first set using the second to observe generalisation to unseen examples. As the GD algorithm progresses, iterating over the training data in rounds called *epochs*, we evaluate the loss on the validation set and stop the process early if this loss does not decrease for a specified period. This *early stopping* criterion was important to produce a model that generalizes well.

As another form of regularisation, i.e. a technique for preventing overfitting to the training data, we employ *dropout* [35] (independently for each “read” of a clause embedding by one of the deriv or eval functions). Dropout means that at training time each component v_i of the embedding v has a certain probability of being zero-ed out. This “voluntary brain damage” makes the network more robust as it prevents neurons from forming too complex co-adaptations [35].

Finally, we experimented with using non-constant learning rates as suggested by Smith et al. [33, 34]. In the end, we used a schedule with a linear warmup for the first 50 epochs followed by a hyperbolic cooldown [38] (cf. Fig. 1 in Sect. 5.2).

⁴ We also tried to skip LayerNorm and replace ReLU by the hyperbolic tangent function. This restores stability, but does not train or classify so well.

4.4 An Abstraction for Compression and Caching

Since our representation of clauses deliberately discards information, we end up encountering distinct clauses indistinguishable from the perspective of the network. For example, every initial clause C originating from the input problem (as opposed to being added as a theory axiom) receives the same embedding $v_C = I_{input}$. Indistinguishable clauses also arise as conclusions of an inference that can be applied in more than one way to certain premises.

Mathematically, we deal with an equivalence relation $\sim$ on clauses based on “having the same derivation tree”: $C_1 \sim C_2 \leftrightarrow \text{derivation}(C_1) = \text{derivation}(C_2)$. The “fingerprint” $\text{derivation}(C)$ of a clause could be defined as a formal expression recording the derivation history of C using the labels from $\mathcal{L}_A$ as nullary operators and those from $\mathcal{L}_R$ as operators with arities of the corresponding inference rules. For example: $\text{Resolution}(\text{thax_inverse_assoc}, \text{Factoring}(\text{input}))$.

We made use of this equivalence in our implementation in two places:

1. When preparing the training data. We “compressed” each derivation DAG as a factorisation by $\sim$, keeping only one representative of each class. A class containing a positive example was marked as a positive example.
2. When interfacing the trained model from the ATP. We cached the embeddings (and evaluated logits) for the already encountered clauses under their class identifier. Sect. 5.4 evaluates the effect of this technique.

5 Experiments

We implemented the infrastructure for training an RvNN clause derivation classifier (as described in Sect. 4) in Python, relying on the PyTorch (version 1.7) library [25] and its TorchScript extension for interfacing the trained model from C++. We modified the automatic theorem prover Vampire (version 4.5.1) to (1) optionally record to a log file the constructed derivation, including information on selected clauses and clauses found in the discovered proof (the *logging-mode*), (2) to be able to load a trained TorchScript model and use it for clause selection guidance under various modes of integration (detailed in Sects. 2.3 and 3).⁵

We took the same subset of 20 795 problems from the SMT-LIB library [5] as in previous work [11]: formed as the largest set of problems in a fragment supported by Vampire, excluding problems known to be satisfiable and those provable by Vampire’s default strategy in 10s either without adding theory axioms or while performing clause selection by age only.

As the baseline strategy $\mathcal{S}$ we took Vampire’s implementation of the DISCOUNT saturation loop under the age-to-weight ratio 1:10 (which typically performs well with DISCOUNT), keeping all other settings default, including the enabled AVATAR architecture. We later enhanced this $\mathcal{S}$ with various forms of guidance. All the benchmarking was done using a 10s time limit.⁶

⁵ Supplementary materials can be found at <https://git.io/JtHNL>.

⁶ Running on an Intel(R) Xeon(R) Gold 6140 CPUs @ 2.3GHz server with 500GB RAM, using no more than 30 of the available 72 cores to reduce mutual influence.

5.1 Data Preparation

During an initial run, the baseline strategy $\mathcal{S}$ was able to solve 734 problems under the 10s time limit. We collected the corresponding successful derivations using the logging-mode (and lifting the time limit, since the logging causes a non-negligible overhead) and processed them into a form suitable for training a neural model. The derivations contained approximately 5.0 million clauses in total (the overall context), out of which 3.9 million were selected⁷ (the training examples) and 30 thousand of these appeared in a proof (the positive examples). In these derivations, Vampire used 31 distinct theory axioms to facilitate theory reasoning. Including the “user input” label for clauses coming from the actual problem files, there were in total 32 distinct labels for the derivation leaves. In addition, we recorded 15 inference rules, such as resolution, superposition, backward and forward demodulation or subsumption resolution and including one rule for the derivation of a component clause in AVATAR [26, 39]. Thus we obtained 15 distinct labels for the internal nodes.

We compressed these derivations identifying clauses with the same “abstract derivation history” dictated by the labels, as described in Sect. 4.4. This reduced the derivation set to 0.7 million nodes (i.e. abstracted clauses) in total. Out of the 734 derivations 242 were still larger than 1000 nodes (the largest had 6426 nodes) and each of these gave rise to a separate “mini-batch”. We grouped the remaining 492 derivations to obtain an approximate size of 1000 nodes per mini-batch (the maximum was 12 original derivations grouped in one mini-batch). In total, we obtained 412 mini-batches and randomly singled out 330 (i.e., 80 %) of these for training, keeping 82 aside for validation.

5.2 Training

Since the size of the training set is relatively small, we instantiated the architecture described in Sect. 4.2 with embedding size $n = 64$ and dropout probability $p = 0.3$. We trained for 100 epochs, with a non-constant learning rate peaking at $\alpha = 2.5 \times 10^{-4}$ in epoch 50. Every epoch we computed the loss on the validation set and selected the model which minimizes this quantity. This was the model from epoch 45 in our case, which we will denote $\mathcal{M}$ here.

The development of the training and validation loss throughout training, as well as that of the learning rate, is plotted in Fig. 1. Additionally, the right side of the figure allows us to compare the validation loss—an ML estimate of the model’s ability to generalize—with the ultimate metric of practical generalization, namely the number of in-training-unseen problems solved by Vampire equipped with the corresponding model for guidance.⁸ We can see that the “proxy” (i.e. the minimisation of the validation loss) and the “target” (i.e. the maximisation of ATP performance) correspond quite well, at least to the degree that we measured the highest ATP gain with the validation-loss-minimizing $\mathcal{M}$.

⁷ Ancestors of selected clauses are sometimes not selected clauses themselves if they arise through immediate simplifications or through reductions.

⁸ Integrated using the layered scheme with a second level ratio 2:1 (cf. Sect. 5.3).

Fig. 1. Training the neural model. Red: the training (left) and validation (right) loss as a function training time; shaded: per problem weighted standard deviations. Blue (left): the supplied non-constant learning rate (cf. Sect. 4.3). Green (right): in training unseen problems solved by Vampire equipped with the corresponding model.

We remark that this assurance was not cheap to obtain. While the whole 100 epoch training took 45 minutes to complete (using 20 workers and 1 master process in a parallel training setup), each of the 20 ATP evaluation data points corresponds to approximately 2 hours of 30 core computation.

5.3 Advice Integration

In this part of the experiment we tested the various ways of integrating the learnt advice as described in Sects. 2.3 and 3. Let us recall that these are the single queue schemes $\mathcal{M}^{-\mathbb{R}}$ and $\mathcal{M}^{1,0}$ based on the raw logits and the binary decision, respectively, their combinations $\mathcal{S} \oplus \mathcal{M}^{-\mathbb{R}}$ and $\mathcal{S} \oplus \mathcal{M}^{1,0}$ with the base strategy $\mathcal{S}$ under some second level ratio, and, finally, $\mathcal{S} \oplus \mathcal{S}[\mathcal{M}^1]$, the integration of the guidance by the layered clause selection scheme.

Our results are shown in Table 1. It starts by reporting on the performance of the baseline strategy $\mathcal{S}$ and then compares it to the other strategies (the gained and lost columns are w.r.t. the original run of $\mathcal{S}$).⁹ We can see that the two single queue approaches are quite weak, with the better $\mathcal{M}^{1,0}$ solving only 25 % of the baseline. Nor can the combination $\mathcal{S} \oplus \mathcal{M}^{-\mathbb{R}}$ be considered a success, as it only solves more problems when less and less advice is taken, seemingly approaching the performance of $\mathcal{S}$ from below. This trend repeats with $\mathcal{S} \oplus \mathcal{M}^{1,0}$, although here an interesting number of problems not solved by the baseline is gained by strategies which rely on the advice more than half of the time.

With our model $\mathcal{M}$, only the layered clause selection integration $\mathcal{S} \oplus \mathcal{S}[\mathcal{M}^1]$ is able to improve on the performance of the baseline strategy $\mathcal{S}$. In fact, it

⁹ We had to switch to a different machine after producing the training data. There, a rerun of $\mathcal{S}$ gave a slightly better performance than the 734 solved problems used for training. We still used the original run’s results to compute the gained and lost values here; the percentage solved is with respect to the new run of $\mathcal{S}$.

Table 1. Performance results of various forms of integrating the model advice.

strategy	ratio	$\mathcal{M}$ eval. time%	#solved (percent $\mathcal{S}$)		gained	lost
$\mathcal{S}$	—	0 %	756	100 %	26	4
$\mathcal{M}^{-\mathbb{R}}$	—	25 %	55	7 %	25	704
$\mathcal{M}^{1,0}$	—	13 %	190	25 %	30	574
$\mathcal{S} \oplus \mathcal{M}^{-\mathbb{R}}$	5:1	57 %	543	71 %	86	277
	2:1	48 %	445	58 %	78	367
	1:1	41 %	335	44 %	54	453
	1:2	32 %	248	32 %	39	525
	1:5	32 %	140	18 %	28	622
$\mathcal{S} \oplus \mathcal{M}^{1,0}$	10:1	11 %	686	90 %	80	128
	2:1	14 %	602	79 %	112	244
	1:1	14 %	555	73 %	111	290
	1:2	14 %	519	68 %	132	347
	1:10	14 %	520	68 %	132	346
$\mathcal{S} \oplus \mathcal{S}[\mathcal{M}^1]$	2:1	27 %	855	113 %	210	89
	1:1	32 %	1032	136 %	411	113
	1:2	33 %	1036	137 %	430	128
	1:3	30 %	1026	135 %	428	136
	1:5	25 %	989	130 %	405	150

Table 2. Performance decrease caused by turning off abstraction caching and lazy evaluation, and both; demonstrated on $\mathcal{S} \oplus \mathcal{S}[\mathcal{M}^1]$ under the second level ratio 1:2.

	$\mathcal{M}$ eval. time%	#solved (percent $\mathcal{S}$)	
both techniques enabled	33 %	1036	137 %
without abstraction caching	45 %	1007	133 %
without lazy evaluation	58 %	905	119 %
both techniques disabled	73 %	782	103 %

improves on it very significantly: with the second level ratio of 1:2 we achieve 137 % performance of the baseline and gain 430 problems unsolved by $\mathcal{S}$.

5.4 Evaluation Speed, Lazy Evaluation, and Abstraction Caching

Table 1 also shows the percentage of computation time the individual strategies spent evaluating the advice, i.e. interfacing $\mathcal{M}$.

A word of warning first. These number are hard to interpret across different strategies. It is because different guidance steers the prover to different parts of the search space. For example, notice the seemingly paradoxical situation most pronounced with $\mathcal{S} \oplus \mathcal{M}^{-\mathbb{R}}$, where the more often is the advice from $\mathcal{M}$ nominally requested, the less time the prover spends interfacing $\mathcal{M}$. Looking closely at a few problems, we discovered that in strategies relying a lot on $\mathcal{M}^{-\mathbb{R}}$, such as $\mathcal{S} \oplus \mathcal{M}^{-\mathbb{R}}$ under the ratio 1:5, most of the time is spent performing forward subsumption. An explanation is that the guidance becomes increasingly bad and the prover slows down, processing larger and larger clauses for which the subsumption checks are expensive and dominate the runtime.¹⁰

¹⁰ A similar experience with bad guidance has been made by the authors of ENIGMA.

Fig. 2. The receiver operating characteristic curve (left) and a related plot with explicit threshold (right) for the selected model $\mathcal{M}$; both based on validation data.

When the guidance is the same, however, we can use the eval. time percentage to estimate the efficiency of the integration. The results shown in Table 1 were obtained using both lazy evaluation¹¹ and abstraction caching (as described in sections 3.1 and 4.4). Taking the best performing $\mathcal{S} \oplus \mathcal{S}[\mathcal{M}^1]$ under the second level ratio 1:2, we selectively disabled: first abstraction caching, then lazy evaluation and finally both techniques, obtaining the values shown in Table 2.

We can see that the techniques considerably contribute to the overall performance. Indeed, without them Vampire would spend the whole 73 % of computation time evaluating the network (compared to only 33 %) and the strategy would barely match (with 103 %) the performance of the baseline $\mathcal{S}$.

5.5 Positive Bias

Two important characteristics, from a machine learning perspective, of an obtained model are the *true positive rate* (TPR) (also called sensitivity) and the *true negative rate* (TNR) (also specificity). TPR is defined as the fraction of positively labeled examples which the model also classifies as such. TNR is, analogously, the fraction of negatively labeled examples. Our model $\mathcal{M}$ achieves (on the validation set) 86 % TPR and 81 % TNR.

The final judgement of a neural classifier follows from a comparison to a threshold value t , set by default to $t = 0$ (recall Sect. 4.1). Changing this threshold allows us to trade TPR for TNR and vice versa in straightforward way. The interdependence of these two values on the varied threshold is traditionally captured by the so called *receiver operating characteristic* (ROC) curve, shown for our model in Fig. 2 (left). The tradition dictates that the x axis be labeled by the false positive rate (FPR) (also called fall-out) which is simply $1 - \text{TNR}$. Under such presentation, one generally strives to pick a threshold value at which the

¹¹ With the exception of the $\mathcal{M}^{-\mathbb{R}}$ guidance, with which it is incompatible.

Table 3. The performance of $\mathcal{S} \oplus \mathcal{S}[\mathcal{M}^1]$ under the second level ratio 1:2 while changing the logit threshold. A smaller threshold means more clauses classified as positive.

threshold	#solved (percent $\mathcal{S}$)		gained lost	
-0.50	1063	140 %	427	98
-0.25	1066	141 %	439	107
0.00	1036	137 %	430	128
0.25	945	125 %	375	164
0.50	825	109 %	278	187

curve is the closest to the upper left corner of the plot.¹² However, this is not necessarily the best configuration for every application.

In the Fig. 2 (right), we “decompose” the ROC curve by using the threshold t for the independent axis x . We also highlight, for every problem (again, in the validation set), what is the minimal logit value across all positively labeled examples belonging to that problem. In other words, what is the logit of the “least positively classified” clause from the problem’s proof. We can see that for the majority of the problems these minima are below the threshold $t = 0$. This means that for those problems at least one clause from the original proof is getting classified as negative by $\mathcal{M}$ under $t = 0$.

These observations motivated us to experiment with non-zero values of the threshold in an ATP evaluation. Particularly promising seemed the use of a threshold t smaller than zero with the intention of classifying more clauses as positive. The results of the experiment are in shown Table 3. Indeed, we could further improve the best performing strategy from Table 1 with both $t = -0.25$ and $t = -0.5$. It can be seen that smaller values lead to fewer problems lost, but even the ATP gain is better with $t = -0.25$ than with the default $t = 0$, leading to the overall best improvement of 141 % with respect to the baseline $\mathcal{S}$.

5.6 Learning from Guided Proofs and Negative Mining

As previously unsolved problems get proven with the help of the trained guidance, the new proofs can be used to enrich the training set and potentially help obtaining even better models. This idea of alternating the training and the ATP evaluation steps in a reinforcing *loop* has been proposed and successfully realized by the authors of ENIGMA on the Mizar dataset [18]. Here we propose an enhancement of the idea and repeat an analogous experiment in our setting.

By collecting proofs discovered by a selection of 8 different configurations tested in the previous sections, we grew our set of solved problems from 734 to 1528. We decided to keep one proof per problem, strictly extending the original training set. We then repeated the same training procedure as described in Sect. 5.2 on this new set and on an extension of this set obtained as follows.

Negative mining: We suspected that the successful derivations obtained with the help of $\mathcal{M}$ might not contain enough “typical wrong decisions” from the

¹² Minimizing the standard cross entropy loss should actually automatically “bring the curve” close to that corner for the threshold $t = 0$.

Table 4. The performance of new models learned from guided proofs. $\mathcal{U}$ is the set of 1528 problems used for the training. The gained and lost counts are here w.r.t. $\mathcal{U}$.

	#solved (percent $\mathcal{S}$) (percent $ \mathcal{U} $)			gained lost	
plain	1268	167 %	82 %	90	350
with negative mining	1394	184 %	91 %	140	274

perspective of $\mathcal{S}$ to provide for good enough training. We therefore logged the *failing* runs of $\mathcal{S}$ on the (1528 – 734) problems only solved by one of the guided strategies and augmented the corresponding derivations with these.¹³

Table 4 confirms¹⁴ that negative mining indeed helps to produce a better model. Mainly, however, it shows that training from additional derivations further dramatically improves the performance of the obtained strategy.

6 Conclusion

We revisited the topic of ENIGMA-style clause selection guidance by a machine learned binary classifier and proposed four improvements to previous work: (1) the use of layered clause selection for integrating the advice, (2) the lazy evaluation trick to reduce the overhead of interfacing a potentially expensive model, (3) the “positive bias” idea suggesting to be really careful not to discard potentially useful clauses, and (4) the “negative mining” technique to provide enough negative examples when learning from proofs obtained with previous guidance.

We have also shown that a strong advice can be obtained by looking just at the derivation history to discriminate a clause. The automatically discovered neural guidance significantly improves upon the human-engineered heuristic [11] under identical conditions. Rerunning $\mathcal{S}$ with the theory heuristic enabled in its default form [10] resulted here in 816 (107 %) solved problems.

By deliberately focusing of the representation of clauses by their derivations, we obtained some nice properties, such as relative speed of evaluation. However, in situations where theory reasoning by automatically added theory axioms is not prevalent, such as on most of the TPTP library, we expect guidance based on derivations with just a single axiom origin label, the *input*, to be quite weak.

Still, we see a great opportunity in using statistical methods for analyzing ATP behaviour; not only for improving prover performance with a black box guidance, but also as a tool for discovering regularities that could be exploited to improve our understanding of the technology on a deeper level.

Acknowledgement

This work was supported by the Czech Science Foundation project 20-06390Y and the project RICAIP no. 857306 under the EU-H2020 programme. We also thank the anonymous reviewers for suggesting numerous improvements.

¹³ Negative mining has, for instance, been previously used when training deep models for the premise selection task [14].

¹⁴ The ATP eval was again integrating via $\mathcal{S} \oplus \mathcal{S}[\mathcal{M}^1]$ under the second level ratio 1:2.

References

1. Aygün, E., Ahmed, Z., Anand, A., Firoiu, V., Glorot, X., Orseau, L., et al.: Learning to prove from synthetic theorems. CoRR **abs/2006.11259** (2020)
2. Ba, L.J., Kiros, J.R., Hinton, G.E.: Layer normalization. CoRR **abs/1607.06450** (2016)
3. Bachmair, L., Ganzinger, H.: Rewrite-based equational theorem proving with selection and simplification. J. Log. Comput. **4**(3), 217–247 (1994). <https://doi.org/10.1093/logcom/4.3.217>
4. Bachmair, L., Ganzinger, H.: Resolution theorem proving. In: Robinson and Voronkov [28], pp. 19–99. <https://doi.org/10.1016/b978-044450813-3/50004-7>
5. Barrett, C., Fontaine, P., Tinelli, C.: The Satisfiability Modulo Theories Library (SMT-LIB) (2016), www.SMT-LIB.org
6. Bradley, A.R., Manna, Z., Sipma, H.B.: What’s decidable about arrays? In: Emerson, E.A., Namjoshi, K.S. (eds.) Verification, Model Checking, and Abstract Interpretation, 7th International Conference, VMCAI 2006, Charleston, SC, USA, January 8–10, 2006, Proceedings. LNCS, vol. 3855, pp. 427–442. Springer (2006). https://doi.org/10.1007/11609773_28
7. Chvalovský, K., Jakubuv, J., Suda, M., Urban, J.: ENIGMA-NG: efficient neural and gradient-boosted inference guidance for E. In: Fontaine [9], pp. 197–215. https://doi.org/10.1007/978-3-030-29436-6_12
8. Denzinger, J., Schulz, S.: Learning Domain Knowledge to Improve Theorem Proving. In: McRobbie, M., Slaney, J. (eds.) Proc. of the 13th CADE, New Brunswick. pp. 62–76. No. 1104 in LNAI, Springer (1996)
9. Fontaine, P. (ed.): Automated Deduction - CADE 27 - 27th International Conference on Automated Deduction, Natal, Brazil, August 27–30, 2019, Proceedings, LNCS, vol. 11716. Springer (2019). <https://doi.org/10.1007/978-3-030-29436-6>
10. Gleiss, B., Suda, M.: Layered clause selection for saturation-based theorem proving. In: Fontaine, P., Korovin, K., Kotsireas, I.S., Rümmer, P., Tournet, S. (eds.) Joint Proceedings of the 7th Workshop on Practical Aspects of Automated Reasoning (PAAR) and the 5th Satisfiability Checking and Symbolic Computation Workshop (SC-Square), co-located with the 10th International Joint Conference on Automated Reasoning (IJCAR 2020), Paris, France, June–July, 2020 (Virtual). CEUR Workshop Proceedings, vol. 2752, pp. 34–52. CEUR-WS.org (2020), <http://ceur-ws.org/Vol-2752/paper3.pdf>
11. Gleiss, B., Suda, M.: Layered clause selection for theory reasoning - (short paper). In: Peltier, N., Sofronie-Stokkermans, V. (eds.) Automated Reasoning - 10th International Joint Conference, IJCAR 2020, Paris, France, July 1–4, 2020, Proceedings, Part I. LNCS, vol. 12166, pp. 402–409. Springer (2020). https://doi.org/10.1007/978-3-030-51074-9_23
12. Goller, C., Küchler, A.: Learning task-dependent distributed representations by backpropagation through structure. In: Proceedings of International Conference on Neural Networks (ICNN’96), Washington, DC, USA, June 3–6, 1996. pp. 347–352. IEEE (1996). <https://doi.org/10.1109/ICNN.1996.548916>
13. Goodfellow, I.J., Bengio, Y., Courville, A.C.: Deep Learning. Adaptive computation and machine learning, MIT Press (2016), <http://www.deeplearningbook.org/>
14. Irving, G., Szegedy, C., Alemi, A.A., Eén, N., Chollet, F., Urban, J.: Deepmath - deep sequence models for premise selection. In: Lee, D.D., Sugiyama, M., von Luxburg, U., Guyon, I., Garnett, R. (eds.) Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5–10, 2016, Barcelona,

- Spain. pp. 2235–2243 (2016), <https://proceedings.neurips.cc/paper/2016/hash/f197002b9a0853eca5e046d9ca4663d5-Abstract.html>
15. Jakubuv, J., Chvalovský, K., Olsák, M., Piotrowski, B., Suda, M., Urban, J.: ENIGMA anonymous: Symbol-independent inference guiding machine (system description). In: Peltier, N., Sofronie-Stokkermans, V. (eds.) *Automated Reasoning - 10th International Joint Conference, IJCAR 2020, Paris, France, July 1-4, 2020, Proceedings, Part II*. LNCS, vol. 12167, pp. 448–463. Springer (2020). https://doi.org/10.1007/978-3-030-51054-1_29
 16. Jakubuv, J., Urban, J.: ENIGMA: efficient learning-based inference guiding machine. In: Geuvers, H., England, M., Hasan, O., Rabe, F., Teschke, O. (eds.) *Intelligent Computer Mathematics - 10th International Conference, CICM 2017, Edinburgh, UK, July 17-21, 2017, Proceedings*. LNCS, vol. 10383, pp. 292–302. Springer (2017). 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.) *Intelligent Computer Mathematics - 11th International Conference, CICM 2018, Hagenberg, Austria, August 13-17, 2018, Proceedings*. LNCS, vol. 11006, pp. 118–124. Springer (2018). https://doi.org/10.1007/978-3-319-96812-4_11
 18. Jakubuv, J., Urban, J.: Hammering Mizar by learning clause guidance (short paper). In: Harrison, J., O’Leary, J., Tolmach, A. (eds.) *10th International Conference on Interactive Theorem Proving, ITP 2019, September 9-12, 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>
 19. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: Bengio, Y., LeCun, Y. (eds.) *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings* (2015), <http://arxiv.org/abs/1412.6980>
 20. Kovács, L., Robillard, S., Voronkov, A.: Coming to terms with quantified reasoning. In: Castagna, G., Gordon, A.D. (eds.) *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2017, Paris, France, January 18-20, 2017*. pp. 260–270. ACM (2017). <https://doi.org/10.1145/3009837.3009887>
 21. Kovács, L., Voronkov, A.: First-order theorem proving and Vampire. In: Sharygina, N., Veith, H. (eds.) *Computer Aided Verification - 25th International Conference, CAV 2013, Saint Petersburg, Russia, July 13-19, 2013. Proceedings*. LNCS, vol. 8044, pp. 1–35. Springer (2013). https://doi.org/10.1007/978-3-642-39799-8_1
 22. 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, May 7-12, 2017. EPIc Series in Computing*, vol. 46, pp. 85–105. EasyChair (2017). <https://doi.org/10.29007/8mwc>
 23. Nieuwenhuis, R., Rubio, A.: Paramodulation-based theorem proving. In: Robinson and Voronkov [28], pp. 371–443. <https://doi.org/10.1016/b978-044450813-3/50009-6>
 24. Olsák, M., Kaliszyk, C., Urban, J.: Property invariant embedding for automated reasoning. In: Giacomo, G.D., Catalá, A., Dilkina, B., Milano, M., Barro, S., Bugarín, A., Lang, J. (eds.) *ECAI 2020 - 24th European Conference on Artificial Intelligence, 29 August-8 September 2020, Santiago de Compostela, Spain, August 29 - September 8, 2020 - Including 10th Conference on Prestigious Applications of Artificial Intelligence (PAIS 2020)*. *Frontiers in Artificial Intelligence and Applications*, vol. 328. IOS Press, 2020. https://doi.org/10.3233/978-1-60750-131-1_10

- cial Intelligence and Applications, vol. 325, pp. 1395–1402. IOS Press (2020). <https://doi.org/10.3233/FAIA200244>
25. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., et al.: Pytorch: An imperative style, high-performance deep learning library. In: Wallach, H., Larochelle, H., Beygelzimer, A., dAlché-Buc, F., Fox, E., Garnett, R. (eds.) *Advances in Neural Information Processing Systems 32*, pp. 8024–8035. Curran Associates, Inc. (2019), <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>
 26. Reger, G., Suda, M., Voronkov, A.: Playing with AVATAR. In: Felty, A.P., Middeldorp, A. (eds.) *Automated Deduction - CADE-25 - 25th International Conference on Automated Deduction*, Berlin, Germany, August 1-7, 2015, *Proceedings. LNCS*, vol. 9195, pp. 399–415. Springer (2015). https://doi.org/10.1007/978-3-319-21401-6_28
 27. Riazanov, A., Voronkov, A.: Limited resource strategy in resolution theorem proving. *J. Symb. Comput.* **36**(1-2), 101–115 (2003). [https://doi.org/10.1016/S0747-7171\(03\)00040-3](https://doi.org/10.1016/S0747-7171(03)00040-3)
 28. Robinson, J.A., Voronkov, A. (eds.): *Handbook of Automated Reasoning* (in 2 volumes). Elsevier and MIT Press (2001)
 29. Sandler, M., Howard, A.G., Zhu, M., Zhmoginov, A., Chen, L.: Mobilenetv2: Inverted residuals and linear bottlenecks. In: *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*. pp. 4510–4520. IEEE Computer Society (2018). <https://doi.org/10.1109/CVPR.2018.00474>
 30. Schulz, S.: *Learning Search Control Knowledge for Equational Deduction*. No. 230 in *DISKI*, Akademische Verlagsgesellschaft Aka GmbH Berlin (2000)
 31. Schulz, S., Cruanes, S., Vukmirovic, P.: Faster, higher, stronger: E 2.3. In: Fontaine [9], pp. 495–507. https://doi.org/10.1007/978-3-030-29436-6_29
 32. Schulz, S., Möhrmann, M.: Performance of clause selection heuristics for saturation-based theorem proving. In: Olivetti, N., Tiwari, A. (eds.) *Automated Reasoning - 8th International Joint Conference, IJCAR 2016, Coimbra, Portugal, June 27 - July 2, 2016, Proceedings. LNCS*, vol. 9706, pp. 330–345. Springer (2016). https://doi.org/10.1007/978-3-319-40229-1_23
 33. Smith, L.N.: Cyclical learning rates for training neural networks. In: *2017 IEEE Winter Conference on Applications of Computer Vision, WACV 2017, Santa Rosa, CA, USA, March 24-31, 2017*. pp. 464–472. IEEE Computer Society (2017). <https://doi.org/10.1109/WACV.2017.58>
 34. Smith, L.N., Topin, N.: Super-convergence: Very fast training of residual networks using large learning rates. *CoRR* **abs/1708.07120** (2017)
 35. Srivastava, N., Hinton, G.E., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.: Dropout: a simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.* **15**(1), 1929–1958 (2014), <http://dl.acm.org/citation.cfm?id=2670313>
 36. Tammiet, T.: GKC: A reasoning system for large knowledge bases. In: Fontaine [9], pp. 538–549. https://doi.org/10.1007/978-3-030-29436-6_32
 37. Urban, J.: personal communication
 38. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., et al.: Attention is all you need. In: Guyon, I., von Luxburg, U., Bengio, S., Wallach, H.M., Fergus, R., Vishwanathan, S.V.N., Garnett, R. (eds.) *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*. pp. 5998–6008 (2017), <https://proceedings.neurips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html>

39. Voronkov, A.: AVATAR: the architecture for first-order theorem provers. In: Biere, A., Bloem, R. (eds.) *Computer Aided Verification - 26th International Conference, CAV 2014, Held as Part of the Vienna Summer of Logic, VSL 2014, Vienna, Austria, July 18–22, 2014. Proceedings. LNCS*, vol. 8559, pp. 696–710. Springer (2014). https://doi.org/10.1007/978-3-319-08867-9_46
40. Weidenbach, C., Dimova, D., Fietzke, A., Kumar, R., Suda, M., Wischniewski, P.: SPASS version 3.5. In: Schmidt, R.A. (ed.) *Automated Deduction - CADE-22, 22nd International Conference on Automated Deduction, Montreal, Canada, August 2–7, 2009. Proceedings. LNCS*, vol. 5663, pp. 140–145. Springer (2009). https://doi.org/10.1007/978-3-642-02959-2_10

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.