
Learning Better Representations From Less Data For Propositional Satisfiability

Mohamed Ghanem* Frederik Schmitt* Julian Siber* Bernd Finkbeiner*

*CISPA Helmholtz Center for Information Security

{mohamed.ghanem,frederik.schmitt,julian.siber,finkbeiner}@cispa.de

Abstract

Training neural networks on NP-complete problems typically demands very large amounts of training data and often needs to be coupled with computationally expensive symbolic verifiers to ensure output correctness. In this paper, we present NeuRes, a neuro-symbolic approach to address both challenges for propositional satisfiability, being the quintessential NP-complete problem. By combining certificate-driven training and expert iteration, our model learns better representations than models trained for classification only, with a much higher data efficiency – requiring orders of magnitude less training data. NeuRes employs propositional resolution as a proof system to generate proofs of unsatisfiability and to accelerate the process of finding satisfying truth assignments, exploring both possibilities in parallel. To realize this, we propose an attention-based architecture that autoregressively selects pairs of clauses from a dynamic formula embedding to derive new clauses. Furthermore, we employ expert iteration whereby model-generated proofs progressively replace longer teacher proofs as the new ground truth. This enables our model to reduce a dataset of proofs generated by an advanced solver by $\sim 32\%$ after training on it with no extra guidance. This shows that NeuRes is not limited by the optimality of the teacher algorithm owing to its self-improving workflow. We show that our model achieves far better performance than NeuroSAT in terms of both correctly classified and proven instances.

1 Introduction

Boolean satisfiability (SAT) is a fundamental problem in computer science. For theory, this stems from SAT being the first problem proven NP-complete [12]. For practice, this is due to many highly-optimized SAT solvers being used as flexible reasoning engines in a variety of tasks such as model checking [11, 46], software verification [16], planning [26], and mathematical proof search [23]. Recently, SAT has also served as a litmus test for assessing the symbolic reasoning capabilities of neural models and a promising domain for neuro-symbolic systems [42, 41, 1, 9, 35]. So far, neural models only provide limited, if any, justification for unsatisfiability predictions. NeuroCore [41], for example, predicts an unsatisfiable core, the verification of which can be as hard as solving the original problem. No certificates at all or certificates that are hard to check limit neural methods in a domain where correctness is critical and prevents close integrations with symbolic methods. Therefore, we propose a neuro-symbolic model that utilizes *resolution* to solve SAT problems by generating easy-to-check certificates.

A resolution proof is a sequence of case distinctions, each involving two clauses, that ends in the empty clause (falsum). This technique can also be used to prove satisfiability by exhaustively applying it until no further new resolution steps are possible and the empty clause has not been derived. Generating such a proof is an interesting problem from a neuro-symbolic perspective because unlike other discrete combinatorial problems that have been considered before [45, 6, 27, 29, 10], it

requires selecting compatible *pairs* of clauses from the dynamically growing pool, as newly derived clauses are naturally considered for derivation in subsequent steps. In this work, we devise three attention-based mechanisms to perform this pair-selection needed for generating resolution proofs. In addition, we augment the architecture to efficiently handle *sat* (satisfiable) formulas with an assignment decoding mechanism that assigns a truth value to each literal. We hypothesize that, despite their final goals being in complete opposition, resolution and *sat* assignment finding can form a mutually beneficial collaboration. On the one hand, clauses derived by resolution incrementally inject additional information into the network, e.g., deriving a single-literal clause by resolution directly implies that literal should be true in any possible *sat* assignment. On the other hand, finding a *sat* assignment absolves the resolution network from having to prove satisfiability by exhaustion. On that basis, given an input formula, NeuRes proceeds in two parallel tracks: (1) finding a *sat* assignment, and (2) deriving a resolution proof of unsatisfiability. Both tracks operate on a shared representation of the problem state. Depending on which track succeeds, NeuRes produces the corresponding SAT verdict which is guaranteed to be sound by virtue of its certificate-based design. Since both of our certificate types are efficient to check, we can afford to perform these symbolic checks at each step. When comparing NeuRes with NeuroSAT [42], which has been trained to predict satisfiability with millions of samples, we demonstrate that NeuRes achieves a higher accuracy while providing a proof and requires only thousands of training samples.

As for most problems in theorem proving we are not only interested in finding any proof but a short proof. Resolution proofs can vary largely in their size depending on the resolution steps taken. Being able to efficiently check the proof, also allows us to adapt the proof target while training the model. In particular, we explore an expert iteration mechanism [38] that pre-rolls the resolution proof of the model and replaces the target proof whenever the pre-rolled proof is shorter. We demonstrate that the bootstrapping mechanism iteratively shortens the proofs of our training dataset while further improving the overall performance of the model.

We make the following contributions:

1. We introduce novel architectures which combine graph neural networks with attention mechanisms for generating resolution proofs and assignments for CNF formulas (Section 4).
2. We show that for propositional logic, learning to prove rather than predict satisfiability results in better representations and requires far less training samples (Section 6 and 7).
3. We devise a bootstrapped training procedure where our model progressively produces shorter resolution proofs than its teacher (Section 6.2) boosting the model’s overall performance.

2 Related Work

SAT Solving and Certificates. We refer to the annual SAT competitions [4] for a comprehensive overview on the ever-evolving landscape of SAT solvers, benchmarks, and proof checkers. SAT solvers are complex systems with a documented history of bugs [8, 25], hence proof certificates have been partially required in this competition since 2013 [2]. Unlike satisfiable formulas, there are several ways to certify unsatisfiable formulas [22]. Resolution proofs [50, 19] are easy to verify [14], but non-trivial to generate from modern solvers based on the paradigm of conflict-driven clause learning [33]. Clausal proofs, e.g., in DRAT format [48], are easier to generate and space-efficient, but hard to validate. Verifying the proofs can take longer than their discovery [21] and requires highly optimized algorithms [30].

Deep Learning for SAT Solving. NeuroSAT [42] was the first study of the Boolean satisfiability problem as an end-to-end learning task. Building upon the NeuroSAT architecture, a simplified version has been trained to predict unsatisfiable cores and successfully integrated in a state-of-the-art SAT solver [41]. It has been showed that both the NeuroSAT architecture and a newly introduced deep exchangeable architecture can outperform SAT solvers on instances of 3-SAT problems [9]. The NeuroSAT architecture has also been applied on special classes of crypto-analysis problems [43]. In addition to supervised learning, unsupervised methods have been proposed for solving SAT problems. For Circuit-SAT a deep-gated DAG recursive neural architecture has been presented together with a differentiable training objective to optimize towards solving the Circuit-SAT problem and finding a satisfying assignment [1]. For Boolean satisfiability, a differentiable training objective has been proposed together with a query mechanism that allows for recurrent solution trials [35].

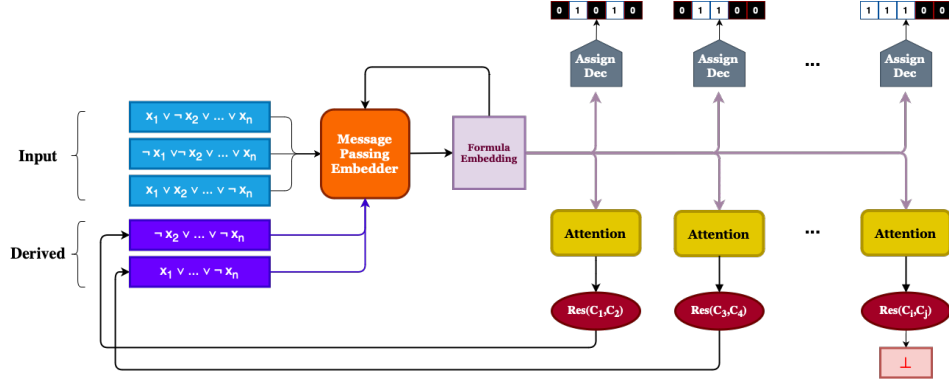


Figure 1: Overall NeuRes architecture

Deep Learning for Formal Proof Generation. In formal mathematics, deep learning has been integrated with theorem proving for clause selection [32, 17], premise selection [24, 47, 5, 34], tactic prediction [49, 36] and whole proof searches [39, 18]. For SMT formulas specifically, deep reinforcement learning has been applied to tactic prediction [3]. In the domain of quantified boolean formulas, heuristics have been learned to guide search algorithms in proving the satisfiability and unsatisfiability of formulas [31]. For temporal logics, deep learning has been applied to prove the satisfiability of linear-time temporal logic formulas and the realizability of specifications [20, 40, 13].

3 Proofs of (Un-)Satisfiability

We start with a brief review of certifying the (un-)satisfiability of propositional formulas in conjunctive normal form. For a set of Boolean variables V , we identify with each variable $x \in V$ the *positive literal* x and the *negative literal* $\neg x$ denoted by \bar{x} . A *clause* corresponds to a disjunction of literals and is abbreviated by a set of literals, e.g., $\{\bar{1}, 3\}$ represents $(\neg x_1 \vee x_3)$. A formula in *conjunctive normal form (CNF)* is a conjunction of clauses and is abbreviated by a set of clauses, e.g., $\{\{\bar{1}, 3\}, \{1, 2, \bar{4}\}\}$ represents $(\neg x_1 \vee x_3) \wedge (x_1 \vee x_2 \vee \neg x_4)$. Any Boolean formula can be converted to an equisatisfiable CNF formula in polynomial time, for example with Tseitin transformation [44].

A CNF formula is *satisfiable* if there exists an *assignment* $\mathcal{A} : V \rightarrow \{\top, \perp\}$ such that all clauses are satisfied, i.e., each clause contains a positive literal x such that $\mathcal{A}(x) = \top$ or a negative literal \bar{x} such that $\mathcal{A}(x) = \perp$. If no such assignment exists we call the formula *unsatisfiable*. To prove unsatisfiability we rely on resolution, a fundamental inference rule in satisfiability testing [15]. The resolution rule (*Res*) picks clauses with two opposite literals and performs the following inference:

$$\frac{C_1 \cup \{x\} \quad C_2 \cup \{\bar{x}\}}{C_1 \cup C_2} \text{ Res}$$

Resolution effectively performs a case distinction on the value of variable x : Either it is assigned to *false*, then C_1 has to evaluate to *true*, or it is assigned to *true*, then C_2 has to evaluate to *true*. Hence, we may infer the clause $C_1 \cup C_2$. A *resolution proof* for a CNF formula is a sequence of applications of the *Res* rule ending in the empty clause.

4 Models

4.1 General Architecture

NeuRes is a neural network that takes a CNF formula as a set of clauses and outputs either a satisfying truth assignment or a resolution proof of unsatisfiability. As such, our model comprises a formula embedder connected to two downstream heads: (1) an attention network responsible for selecting clause pairs, and (2) a truth assignment decoder. See Figure 1 for an overview of the NeuRes architecture. After obtaining the initial clause and literal embeddings (representing the input formula), we continue with the iterative certificate generation phase. At each step, the model selects a clause

pair which gets resolved into a new clause to append to the current formula graph while decoding a candidate truth assignment in parallel. The model keeps deriving new clauses until the empty clause is found (marking resolution proof completion), a satisfying assignment is found (marking a certified *sat* verdict), or the limit on episode length is reached (marking timeout).

4.2 Message-Passing Embedder

Similar to NeuroSAT, we use a message-passing GNN to obtain clause and literal embeddings by performing a predetermined number of rounds. Our formula graph is also constructed in a similar fashion to NeuroSAT graphs where clause nodes are connected to their constituent literal nodes and literals are connect to their complements (cf. Appendix A). For a formula F in m variables and n clauses, the outputs of this GNN are two matrices: $E^L \in \mathbb{R}^{m \times d}$ for literal embeddings and $E^C \in \mathbb{R}^{n \times d}$ for clause embedding, where $d \in \mathbb{N}^+$ is the embedding dimension. Here we have two key differences from NeuroSAT. Firstly, NeuroSAT uses these embeddings as voters to predict satisfiability through a classification MLP. In our case, we use these embeddings as clause tokens for clause pair selection and literal tokens for truth value assignment. Secondly, since our model derives new clauses with every resolution step, we need to embed these new clauses, as well as update existing embeddings to reflect their relation to the newly inferred clauses. Consequently, we need to introduce a new phase to the message-passing protocol, for which we explore two approaches: *static embeddings* and *dynamic embeddings*.

In a *static* approach, we do not change the embeddings of initial clauses upon inferring a new clause. Instead, we exchange local messages between the node corresponding to the new clause and its literal nodes, in both directions. The main advantage of this approach is its low cost. A major drawback is that initial clauses never learn information about their relation to newly inferred clauses.

In a *dynamic* approach, we do not only generate a new clause and its embedding, we also update the embeddings of all other clauses. This accounts for the fact that the utility of an existing clause may change with the introduction of a new clause. We perform one message-passing round on the mature graph for every newly derived clause, which produces the new clause embedding and updates other clause embeddings. Since message-passing rounds are parallel across clauses, a single update to the whole embedding matrix is reasonably efficient.

4.3 Selector Networks

After producing clause and literal embeddings, NeuRes enters the derivation stage. At each step, our model needs to select two clauses to resolve, produce the resultant clause, and add it to the current formula. To realize our clause-pair selection mechanism, we employ three attention-based designs.

4.3.1 Cascaded Attention (Casc-Attn)




Figure 2: Cascaded attention

In this design, pairs are selected by making two consecutive attention queries on the clause pool. We condition the second attention query on the outcome (i.e., the clause) of the first query. Figure 2 shows this scheme where we perform the first query using the mean of the literal embeddings $\overline{E^L}$ concatenated with a zero vector while performing the second query using the mean of the literal embeddings concatenated with the embedding vector $E^c_{c_1}$ of the clause selected in the first query. Formally, Casc-Attn selects a clause index pair (c_1, c_2) as follows:

$$c_i = \underset{j}{\mathbf{argmax}} \left[u^T \tanh(W_1 q_i + W_2 E_j^C) \right] \quad \text{with} \quad q_i = \begin{cases} \overline{E^L} \parallel \mathbf{0} & \text{if } i = 1 \\ \overline{E^L} \parallel E_{c_1}^C & \text{if } i = 2 \end{cases} \quad (1)$$

where $W_1 \in \mathbb{R}^{2d \times d}$, $W_2 \in \mathbb{R}^{d \times d}$, $u \in \mathbb{R}^d$ are trainable network parameters.

The advantage of this design is that it is not limited to pair selection and can be used to select a tuple of arbitrary length. The main downside, however, is that this design chooses c_1 independently from c_2 , which is undesirable because the utility of a resolution step is determined by both clauses simultaneously (not sequentially).

4.3.2 Full Self-Attention (Full-Attn)

To address the downside of independent clause selection, this variant performs self-attention between all clauses to obtain a matrix $S \in \mathbb{R}^{n \times n}$ where $S_{i,j}$ represents the attention score of the clause pair (c_i, c_j) as shown in Figure 3. The model selects clause pairs by choosing the cell with the maximal score. In this attention scheme, the clause embeddings are used as both queries and keys.




Figure 3: Full self-attention

Formally, Full-Attn selects a clause index pair (c_1, c_2) as follows:

$$(c_1, c_2) = \underset{(i,j)}{\mathbf{argmax}} S_{i,j} \quad \text{with} \quad Q = E^C W_Q; \quad K = E^C W_K; \quad S = \frac{QK^T}{\sqrt{d}} \quad (2)$$

where $W_Q \in \mathbb{R}^{d \times d}$, $W_K \in \mathbb{R}^{d \times d}$ are trainable network parameters. Since S contains many cells that correspond to invalid resolution steps (i.e., clause pairs that cannot be resolved), we mask out the invalid cells from the attention grid in ensure the network selection is valid at every step.

4.3.3 Anchored Self-Attention (Anch-Attn)

In Full-Attn, the attention grid grows quadratically with the number of clauses. In this variant, we relax this cost by exploiting a property of binary resolution where each step targets a single variable in the two resolvent clauses. This allows us to narrow down candidate clause pairs by first selecting a variable as an anchor on which our clauses should be resolved. As such, we do not need to consider the full clause set at once, only the clauses containing the chosen variable v . We further compress the attention grid by lining clauses containing the literal v on rows while lining clauses containing the literal $\neg v$ on columns. This reduces the redundancy of the attention grid since clauses containing the variable v with the same parity cannot be resolved on v , so there is no point in matching them. In the worst case, this relaxed grid is of size $\frac{n}{2} \times \frac{n}{2} = \frac{n^2}{4}$ instead of n^2 . In this scheme, we have two attention modules: one attention network to choose an anchor variable followed by a self-attention network to produce the anchored score grid.

In light of Figure 4, this approach combines structural elements from Casc-Attn and Full-Attn; however, both elements are used differently in Anch-Attn. Firstly, the attention mechanism in Casc-Attn is used to select clauses whereas Anch-Attn uses it to select variables. Secondly, self-attention in Full-Attn matches any pair of clauses (c_i, c_j) in both directions as the row and column dimensions in the attention score grid reflect the same clauses (all clauses). By contrast, Anch-Attn computes self-attention scores for clause pairs in only one order (positive instance to negative instance). Formally, Anch-Attn selects an anchor variable v as follows:

$$v = \underset{i}{\mathbf{argmax}} \left[u^T \tanh(W_1 \overline{E^C} + W_2 (E_i^{L^+} + E_i^{L^-})) \right] \quad (3)$$




Figure 4: Anchored self-attention

where $W_1 \in \mathbb{R}^{d \times d}$, $W_2 \in \mathbb{R}^{d \times d}$, $u \in \mathbb{R}^d$ are trainable network parameters. The clause index pair (c_1, c_2) is then selected according to the same equations of Full-Attn (Eq. 2) using the v -anchored set of clause embeddings.

4.4 Assignment Decoder

To extract satisfying assignments, we use a sigmoid-activated MLP ψ on top of the literal embeddings E^L to assign a truth value $\hat{\mathcal{A}}(l_i)$ to a literal l_i as shown in Eq. 4.

$$\hat{\mathcal{A}}(l_i) = \sigma(\psi(E_i^L)) \quad (4)$$

Note that since for each variable, we have a positive and a negative literal embeddings, we can construct two different truth assignments at a time using this method. However, to simplify our loss function, we only derive truth assignments from the positive literal embeddings at train time while extracting both at test time.

5 Training and Hyperparameters

5.1 Dataset

For our training and testing data, we adopt the same formula generation method as NeuroSAT [42], namely $\mathbf{SR}(n)$ where n is the number of variables in the formula. This method was designed to generate a generalized formula distribution that is not limited to a particular domain of SAT problems. To control our data distributions, we vary the range on the number of Boolean variables involved in each formula. For our training data, we use formulas in $\mathbf{SR}(U(10, 40))$ where $U(10, 40)$ denotes the uniform distribution on integers between 10 and 40 (inclusive). To generate our teacher certificates comprising resolution proofs and truth assignments, we use the BooleForce solver [7] on the formulas generated on the \mathbf{SR} distribution.

5.2 Loss Function

We train our model in a supervised fashion using teacher-forcing on solver certificates. During *unsat* episodes, teacher actions (clause pairs) are imposed over the whole run. The length of the teacher proof dictates the length of the respective episode, denoted as T . Model parameters θ maximize the likelihood of teacher choices y_t thereby minimizing the resolution loss \mathcal{L}_{Res} shown in Eq 5.

$$\mathcal{L}_{Res} = -\frac{1}{T} \sum_t \log(p(y_t; \theta)) \cdot \gamma^{(T-t)} \quad (5)$$

During *sat* episodes, we minimize \mathcal{L}_{sat} computed as the binary cross-entropy loss between the sigmoid-activated outputs of assignment decoder $\hat{\mathcal{A}} : V \rightarrow [0, 1]$ and the teacher assignment $\mathcal{A} : V \rightarrow \{0, 1\}$ as shown in Eq. 6.

$$\mathcal{L}_{sat} = \frac{1}{T} \sum_t \left[\frac{\gamma^{(T-t)}}{|V|} \sum_v \text{BCE}(\hat{\mathcal{A}}(v), \mathcal{A}(v)) \right] \quad (6)$$

In both types of episodes, step-wise losses are weighted by a time-horizon discounting factor $\gamma < 1.0$ over the whole episode. The main rationale behind this is that later losses should have higher weights as the formula tends to get easier to solve with each new clause inferred by resolution.

Table 1: Performance of all attention variants on *unsat* $\mathbf{SR}(U(10, 40))$ test problems.

VARIANT	STATIC-EMBED		DYNAMIC-EMBED	
	PROVEN (%)	P-LEN	PROVEN (%)	P-LEN
CASC-ATTN	14.72	1.87	37.33	1.79
FULL-ATTN	25.38	1.61	95.2	1.67
ANCH-ATTN	28.72	2.12	60.5	2.28

Table 2: Bootstrapped training data reduction statistics. Reduction statistics are computed on the $\mathbf{SR}(U(10, 40))$ training set while p-Len and success rate are computed on a test set of the same distribution.

REDUCTION DEPTH	MAX: 23, AVG: 6.6
PROOF REDUCTION (%)	MAX: 86.11, AVG: 33.51
PROOFS REDUCED (%)	90.08
TOTAL REDUCTION (%)	31.85
P-LEN	1.15
SUCCESS RATE (%)	100.0

5.3 Hyperparameters

NeuRes has several hyperparameters that influence network size, depth, and loss weighting. In the experiments we fix the embedding dimension to 128. We train our models with a batch size of 1 and the Adam optimizer [28] for 50 epochs which took about six days on a single NVIDIA A100 GPU. We linearly anneal the learning rate from 5×10^{-5} to zero over the training episodes. This empirically yields better results than using a constant learning rate. We use a time discounting factor $\lambda = 0.99$ for the episodic loss. We apply global-norm gradient clipping with a ratio of 0.5 [37].

6 Generating Resolution Proofs

NeuRes uses resolution as the core reasoning technique for certificate generation, both in the *unsat* and *sat* cases. Hence, we start with an in-depth comparative evaluation of several internal variants for resolution only. In particular, we evaluate the success rate (i.e., problems solved before timeout) and proof length relative to teacher, denoted by p-Len = $\frac{|\mathcal{P}_{\text{NeuRes}}|}{|\mathcal{P}_{\text{teacher}}|}$. We use a limit of 4 on this ratio as a timeout to avoid simply brute-forcing a resolution proof. Note that we measure p-Len only for solved formulas to avoid diluting the average with resolution trails that timed out. For experiments in this section, we train on 8K *unsat* formulas in $\mathbf{SR}(U(10, 40))$ and test our models on 10K unseen formulas belonging to the same distribution. We use more formulas than the model was trained on to more reliably demonstrate its learning capacity.

6.1 Attention Variants

To assess the basic resolution performance of NeuRes, we evaluate each attention variant using both static and dynamic embeddings. For this experiment, we perform 32 rounds of message-passing for each input formula. As shown in Table 1, dynamic embedding is decisively better for all three attention variants, thereby confirming its conceptual merit. While anchored attention leads over other variants under static embeddings, full attention performs significantly better for dynamic embeddings, albeit at the cost of longer proofs on average. We believe that Anch-Attn’s better performance in the static setting can be explained through the full connectivity of its attention grid (proven in Appendix B). Since dynamic-embedding Full-Attn is the best-performing configuration over in-distribution test settings, we will demonstrate the remaining evaluation experiments exclusively on this variant.

6.2 Shortening Teacher Proofs with Bootstrapping

During our initial experiments, we discovered proofs produced by NeuRes that were shorter than the corresponding teacher proofs in the training data. Although teacher proofs were generated by a traditional SAT solver, they are not guaranteed to be size-optimal. The size of resolution proofs is their only real drawback, hence any method that can reduce this size would be immensely useful. Upon closer inspection we find that, on average, our previous best performer trained with regular teacher-forcing manages to shorten $\sim 18\%$ of teacher proofs by a notable factor (cf. Appendix C).

This inspired us to devise a bootstrapped training procedure to capitalize on this feature: We pre-roll each input problem using model actions only, and whenever the model proof is shorter than the teacher’s, it replaces the teacher’s in the dataset. In other words, we maximize the likelihood of the shorter proof. In doing so iteratively, the model progressively becomes its own teacher by exploiting redundancies in the teacher algorithm.

The outcome of this bootstrapped training process is summarized in Table 2. We find that bootstrapping results in notable gains in terms of both success rate and optimality. The sharp decline in proof length (relatively quantified by p-Len) at test time shows that the models transfers the bootstrapped knowledge to unseen test formulas, as opposed to merely overfitting on training formulas. In addition to success rate and p-Len, we inspect the reduction statistics of our bootstrapped variant (first three rows of Table 2). Since the bootstrapped model performs multiple reduction scans over the training dataset, we add a metric for reduction depth computed as the number of progressive reductions made to a proof. To further quantify this effect, we report the maximum and average reduction ratios of reduced proofs relative to teacher proofs. Finally, we report the total reduction made to the dataset size in terms of total number of proof steps.

In Appendix C, we have compiled additional statistics (cf. Table 4) on proof shortening during the training process, as well as an example proof reduced by the bootstrapped NeuRes (Figure 7). We only include a small reduction example (from 20 steps to 10 steps) for space constraints, but we have many more examples of much larger reductions (e.g., from ~ 800 to ~ 400 steps).

7 Resolution-Aided SAT Solving

In this section, we evaluate the performance of our fully integrated model trained on a hybrid dataset comprising 8K unsatisfiable formulas (and their resolution proofs) and 8K satisfiable formulas (and their satisfying assignments). For the unsatisfiable formulas, timeout ($4 \cdot |\mathcal{P}_{\text{teacher}}|$) and optimality (p-Len) are measured similarly to previous experiments. For satisfiable formulas, we set the timeout (maximum #trials) to $2 \cdot |V|$. Ultimately, this section aims to investigate the effect of incorporating a certificate-driven downstream head on the quality of the learnt representations through its impact on the performance of the complementary task, i.e., proving/predicting satisfiability. We use NeuroSAT as our baseline as it employs the same formula embedding architecture. Since NeuroSAT proves *sat* but only *predicts unsat*, we train a classification MLP on top of our trained NeuRes model to further showcase the benefit of our representations on prediction accuracy.

Table 3: Performance of full solver mode tested on **SR(40)** problems and trained on **SR($U(10, 40)$)** problems where PREDICTED refers to the satisfiability prediction without certificate.

MODEL	PROVEN (%)			PREDICTED (%)		
	SAT	UNSAT	TOTAL	SAT	UNSAT	TOTAL
NEURES	96.8	99.6	98.2	84.28	99.2	91.65
NEUROSAT [42]	70	-	-	73	96	85

Table 3 confirms this main hypothesis. In essence, this result points to the fact that learning signals obtained from training on *unsat* certificates largely enhance the ability of the neural network to extract useful information from the input formula. This is doubly promising considering NeuroSAT was trained on *millions* of formulas while NeuRes was trained on only *16K* formulas. Lastly, we find that augmenting *sat* formulas by resolution derivations results in relative improvements ($\sim 2.3\%$) in success rate even though these derivations are attempting to prove unsatisfiability.

8 Generalizing to Larger Problems

We also evaluate our Res-Aided model (trained on $\mathbf{SR}(U(10, 40))$) on five datasets comprising formulas with up to 5 times more variables than encountered during training. We use the same distributions reported by NeuroSAT and we run our model for the same maximum number of iterations (1000).




Figure 5: SAT success rate over iterations.

Figure 5 shows the scalability of NeuRes to larger problems by letting it run for more iterations. Compared to NeuroSAT [42], NeuRes scores a much higher first-try success rate on all 5 problem distributions, and a higher final success rate on all of them except for $\mathbf{SR}(40)$ on which both models nearly score 100%. Particularly, NeuRes shows higher first-try success on the 3 largest problem sizes where NeuroSAT solves zero or near-zero problems on the first try.

9 Conclusion

In this paper, we introduced a deep learning approach for proving and predicting propositional satisfiability. We proposed an architecture that combines graph neural networks with attention mechanisms to generate resolution proofs of unsatisfiability. Unlike methods that merely predict **unsatisfiability**, our models provide easily verifiable certificates for their verdicts. We demonstrated that our certificate-based training and resolution-aided mode of operation surpass previous approaches in terms of performance and data efficiency, which we attribute to learning better representation.

Despite its promising benchmark performance, our model cannot solely outperform highly-engineered industrial solvers, as is currently the case for all neural methods as standalone tools. The gap between neural networks and symbolic algorithms is still several breakthroughs wide, and our hope is to bring deep learning methods one concrete step closer to filling this gap. For NeuRes, this step is recognizing the immense value of carefully integrating certificates into the model design and training as opposed to using shallow supervision labels. Last but not least, it is worth noting that even at their present state, neural networks stand great potential to advance traditional solvers by combining them in hybrid solvers that utilize the deep long-range dependencies captured by neural networks along with the exploration speed of symbolic algorithms. Moreover, we demonstrated a unique potential to advance SAT solving through proof reduction, as proof size is a major challenge in certifying the results of traditional solvers. This proof reduction is facilitated by a bootstrapped training procedure that uses teacher proofs as a guide as opposed to a golden standard.

References

[1] S. Amizadeh, S. Matushevych, and M. Weimer. Learning to solve circuit-sat: An unsupervised differentiable approach. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019.

- [2] A. Balint, A. Belov, M. Heule, and M. Järvisalo, editors. *Proceedings of SAT Competition 2013: Solver and Benchmark Descriptions*, volume B-2013-1 of *Department of Computer Science Series of Publications B*. University of Helsinki, Finland, 2013.
- [3] M. Balunovic, P. Bielik, and M. T. Vechev. Learning to solve SMT formulas. In S. Bengio, H. M. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 10338–10349, 2018.
- [4] T. Balyo, M. Heule, M. Iser, M. Järvisalo, and M. Suda, editors. *Proceedings of SAT Competition 2023: Solver, Benchmark and Proof Checker Descriptions*. Department of Computer Science Series of Publications B. Department of Computer Science, University of Helsinki, Finland, 2023.
- [5] K. Bansal, S. M. Loos, M. N. Rabe, C. Szegedy, and S. Wilcox. Holist: An environment for machine learning of higher order logic theorem proving. In K. Chaudhuri and R. Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pages 454–463. PMLR, 2019.
- [6] I. Bello, H. Pham, Q. V. Le, M. Norouzi, and S. Bengio. Neural combinatorial optimization with reinforcement learning. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Workshop Track Proceedings*. OpenReview.net, 2017.
- [7] A. Biere. Booleforce sat solver. <https://fmv.jku.at/booleforce/>, 2010.
- [8] R. Brummayer, F. Lonsing, and A. Biere. Automated testing and debugging of SAT and QBF solvers. In O. Strichman and S. Szeider, editors, *Theory and Applications of Satisfiability Testing - SAT 2010, 13th International Conference, SAT 2010, Edinburgh, UK, July 11-14, 2010. Proceedings*, volume 6175 of *Lecture Notes in Computer Science*, pages 44–57. Springer, 2010.
- [9] C. Cameron, R. Chen, J. S. Hartford, and K. Leyton-Brown. Predicting propositional satisfiability via end-to-end learning. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 3324–3331. AAAI Press, 2020.
- [10] Q. Cappart, D. Chételat, E. B. Khalil, A. Lodi, C. Morris, and P. Velickovic. Combinatorial optimization and reasoning with graph neural networks. *J. Mach. Learn. Res.*, 24:130:1–130:61, 2023.
- [11] E. M. Clarke, A. Biere, R. Raimi, and Y. Zhu. Bounded model checking using satisfiability solving. *Formal Methods Syst. Des.*, 19(1):7–34, 2001.
- [12] S. A. Cook. The complexity of theorem-proving procedures. In M. A. Harrison, R. B. Banerji, and J. D. Ullman, editors, *Proceedings of the 3rd Annual ACM Symposium on Theory of Computing, May 3-5, 1971, Shaker Heights, Ohio, USA*, pages 151–158. ACM, 1971.
- [13] M. Cosler, F. Schmitt, C. Hahn, and B. Finkbeiner. Iterative circuit repair against formal specifications. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023.
- [14] A. Darbari, B. Fischer, and J. Marques-Silva. Industrial-strength certified SAT solving through verified SAT proof checking. In A. Cavalcanti, D. Déharbe, M. Gaudel, and J. Woodcock, editors, *Theoretical Aspects of Computing - ICTAC 2010, 7th International Colloquium, Natal, Rio Grande do Norte, Brazil, September 1-3, 2010. Proceedings*, volume 6255 of *Lecture Notes in Computer Science*, pages 260–274. Springer, 2010.
- [15] M. Davis and H. Putnam. A computing procedure for quantification theory. *J. ACM*, 7(3):201–215, 1960.
- [16] L. M. de Moura and N. S. Bjørner. Satisfiability modulo theories: introduction and applications. *Commun. ACM*, 54(9):69–77, 2011.

- [17] V. Firoiu, E. Aygün, A. Anand, Z. Ahmed, X. Glorot, L. Orseau, L. M. Zhang, D. Precup, and S. Mourad. Training a first-order theorem prover from synthetic data. *CoRR*, abs/2103.03798, 2021.
- [18] E. First, M. N. Rabe, T. Ringer, and Y. Brun. Baldur: Whole-proof generation and repair with large language models. *CoRR*, abs/2303.04910, 2023.
- [19] E. I. Goldberg and Y. Novikov. Verification of proofs of unsatisfiability for CNF formulas. In *2003 Design, Automation and Test in Europe Conference and Exposition (DATE 2003), 3-7 March 2003, Munich, Germany*, pages 10886–10891. IEEE Computer Society, 2003.
- [20] C. Hahn, F. Schmitt, J. U. Kreber, M. N. Rabe, and B. Finkbeiner. Teaching temporal logics to neural networks. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021.
- [21] M. Heule, W. A. H. Jr., and N. Wetzler. Bridging the gap between easy generation and efficient verification of unsatisfiability proofs. *Softw. Test. Verification Reliab.*, 24(8):593–607, 2014.
- [22] M. J. H. Heule. Proofs of unsatisfiability. In A. Biere, M. Heule, H. van Maaren, and T. Walsh, editors, *Handbook of Satisfiability - Second Edition*, volume 336 of *Frontiers in Artificial Intelligence and Applications*, pages 635–668. IOS Press, 2021.
- [23] M. J. H. Heule and O. Kullmann. The science of brute force. *Commun. ACM*, 60(8):70–79, 2017.
- [24] G. Irving, C. Szegedy, A. A. Alemi, N. Eén, F. Chollet, and J. Urban. Deepmath - deep sequence models for premise selection. In D. D. Lee, M. Sugiyama, U. von Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 2235–2243, 2016.
- [25] M. Järvisalo, M. Heule, and A. Biere. Inprocessing rules. In B. Gramlich, D. Miller, and U. Sattler, editors, *Automated Reasoning - 6th International Joint Conference, IJCAR 2012, Manchester, UK, June 26-29, 2012. Proceedings*, volume 7364 of *Lecture Notes in Computer Science*, pages 355–370. Springer, 2012.
- [26] H. A. Kautz and B. Selman. Planning as satisfiability. In B. Neumann, editor, *10th European Conference on Artificial Intelligence, ECAI 92, Vienna, Austria, August 3-7, 1992. Proceedings*, pages 359–363. John Wiley and Sons, 1992.
- [27] E. B. Khalil, H. Dai, Y. Zhang, B. Dilkina, and L. Song. Learning combinatorial optimization algorithms over graphs. In I. Guyon, U. von Luxburg, S. Bengio, H. M. Wallach, R. Fergus, S. V. N. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 6348–6358, 2017.
- [28] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [29] W. Kool, H. van Hoof, and M. Welling. Attention, learn to solve routing problems! In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019.
- [30] P. Lammich. Efficient verified (UN)SAT certificate checking. *J. Autom. Reason.*, 64(3):513–532, 2020.
- [31] G. Lederman, M. N. Rabe, S. Seshia, and E. A. Lee. Learning heuristics for quantified boolean formulas through reinforcement learning. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020.
- [32] S. M. Loos, G. Irving, C. Szegedy, and C. Kaliszyk. Deep network guided proof search. In T. Eiter and D. Sands, editors, *LPAR-21, 21st International Conference on Logic for Programming, Artificial Intelligence and Reasoning, Maun, Botswana, May 7-12, 2017*, volume 46 of *EPiC Series in Computing*, pages 85–105. EasyChair, 2017.

- [33] J. Marques-Silva, I. Lynce, and S. Malik. Conflict-driven clause learning SAT solvers. In A. Biere, M. Heule, H. van Maaren, and T. Walsh, editors, *Handbook of Satisfiability - Second Edition*, volume 336 of *Frontiers in Artificial Intelligence and Applications*, pages 133–182. IOS Press, 2021.
- [34] M. Mikula, S. Antoniak, S. Tworowski, A. Q. Jiang, J. P. Zhou, C. Szegedy, L. Kucinski, P. Milos, and Y. Wu. Magnushammer: A transformer-based approach to premise selection. *CoRR*, abs/2303.04488, 2023.
- [35] E. Ozolins, K. Freivalds, A. Draguns, E. Gaile, R. Zakovskis, and S. Kozlovics. Goal-aware neural SAT solver. In *International Joint Conference on Neural Networks, IJCNN 2022, Padua, Italy, July 18-23, 2022*, pages 1–8. IEEE, 2022.
- [36] A. Paliwal, S. M. Loos, M. N. Rabe, K. Bansal, and C. Szegedy. Graph representations for higher-order logic and theorem proving. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 2967–2974. AAAI Press, 2020.
- [37] R. Pascanu, T. Mikolov, and Y. Bengio. Understanding the exploding gradient problem. *CoRR*, abs/1211.5063, 2(417):1, 2012.
- [38] S. Polu, J. M. Han, K. Zheng, M. Baksys, I. Babuschkin, and I. Sutskever. Formal mathematics statement curriculum learning. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023.
- [39] S. Polu and I. Sutskever. Generative language modeling for automated theorem proving. *CoRR*, abs/2009.03393, 2020.
- [40] F. Schmitt, C. Hahn, M. N. Rabe, and B. Finkbeiner. Neural circuit synthesis from specification patterns. In M. Ranzato, A. Beygelzimer, Y. N. Dauphin, P. Liang, and J. W. Vaughan, editors, *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 15408–15420, 2021.
- [41] D. Selsam and N. S. Bjørner. Guiding high-performance SAT solvers with unsat-core predictions. In M. Janota and I. Lynce, editors, *Theory and Applications of Satisfiability Testing - SAT 2019 - 22nd International Conference, SAT 2019, Lisbon, Portugal, July 9-12, 2019, Proceedings*, volume 11628 of *Lecture Notes in Computer Science*, pages 336–353. Springer, 2019.
- [42] D. Selsam, M. Lamm, B. Bünz, P. Liang, L. de Moura, and D. L. Dill. Learning a SAT solver from single-bit supervision. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019.
- [43] L. Sun, D. Gérard, A. Benamira, and T. Peyrin. Neurogift: Using a machine learning based sat solver for cryptanalysis. In S. Dolev, V. Kolesnikov, S. Lodha, and G. Weiss, editors, *Cyber Security Cryptography and Machine Learning - Fourth International Symposium, CSCML 2020, Be'er Sheva, Israel, July 2-3, 2020, Proceedings*, volume 12161 of *Lecture Notes in Computer Science*, pages 62–84. Springer, 2020.
- [44] G. S. Tseitin. On the complexity of derivation in propositional calculus. *Automation of reasoning: 2: Classical papers on computational logic 1967–1970*, pages 466–483, 1983.
- [45] O. Vinyals, M. Fortunato, and N. Jaitly. Pointer networks. *Advances in neural information processing systems*, 28, 2015.
- [46] Y. Vazel, G. Weissenbacher, and S. Malik. Boolean satisfiability solvers and their applications in model checking. *Proc. IEEE*, 103(11):2021–2035, 2015.
- [47] M. Wang, Y. Tang, J. Wang, and J. Deng. Premise selection for theorem proving by deep graph embedding. In I. Guyon, U. von Luxburg, S. Bengio, H. M. Wallach, R. Fergus, S. V. N. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 2786–2796, 2017.

- [48] N. Wetzler, M. Heule, and W. A. H. Jr. Drat-trim: Efficient checking and trimming using expressive clausal proofs. In C. Sinz and U. Egly, editors, *Theory and Applications of Satisfiability Testing - SAT 2014 - 17th International Conference, Held as Part of the Vienna Summer of Logic, VSL 2014, Vienna, Austria, July 14-17, 2014. Proceedings*, volume 8561 of *Lecture Notes in Computer Science*, pages 422–429. Springer, 2014.
- [49] K. Yang and J. Deng. Learning to prove theorems via interacting with proof assistants. In K. Chaudhuri and R. Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pages 6984–6994. PMLR, 2019.
- [50] L. Zhang and S. Malik. Validating SAT solvers using an independent resolution-based checker: Practical implementations and other applications. In *2003 Design, Automation and Test in Europe Conference and Exposition (DATE 2003), 3-7 March 2003, Munich, Germany*, pages 10880–10885. IEEE Computer Society, 2003.

Appendix

A NeuroSAT Formula Graph Construction




Figure 6: Two-phase message-passing round on NeuroSAT formula graph.

NeuroSAT-style formula graphs have two designated node types: clause nodes connected to the literal nodes corresponding to their constituent literals [42]. For example, in Figure 6, the clause contents are as follows: $c_1 = (a \vee b)$, $c_2 = (a \vee \bar{b})$, $c_3 = (\bar{a} \vee b)$. Each message-passing round involves two exchange phases: (1) Literal-to-Clause, and (2) Clause-to-Literal (and implicitly Literal-to-Complement). This construction is particularly efficient as it allows the message-passing protocol to cover the entire graph connectivity in at most $|V| + 1$ rounds where V is the set of variables in the formula.

B Clause Connectivity Under Static Embeddings

In Section 4.2, we stated that under static embeddings for a derived clause, as the embedder creates its embedding, it only updates the representations of the variables involved in it – leaving other clause embeddings intact. This might present a problem for Full-Attn where the attention grid contains all clauses including disconnected¹ pairs. An example of such a pair would be two derived clauses that do not share a variable. This could potentially lower the efficacy of the Full-Attn mechanism as it tries to match clauses that are unaware of each other. Interestingly, despite being a relaxation on Full-Attn, Anch-Attn has a distinct edge over Full-Attn under static embeddings in form of the following property:

Lemma B.1. *Clauses in the variable-anchored attention grid of Anch-Attn are guaranteed to be connected under both static and dynamic embeddings.*

Proof. Let v be a variable in the input formula, and the set of clauses of a v -anchored attention grid be A . We show that we always have at least one clause $A_i \in A$ that reaches all other clauses in A on the formula graph. We make two case distinctions:

Case 1: All clauses in A are input clauses (in the original formula). Here, the lemma follows trivially since all these clause were connected during the input-phase message-passing protocol as they share at least one variable v .

Case 2: A contains derived clauses. Let A_i be the most recently derived clause in A . Since A_i shares variable v with all other clauses in A , then A_i would be connected to them all during the derivation-phase message-passing protocol immediately after A_i was derived. This is because A_i receives a message from V (under both static and dynamic embeddings) containing information about all other clauses containing v , which is precisely $A \setminus \{A_i\}$. Therefore, the lemma holds. \square


¹We use the terms *connected* and *disconnected* here to refer to the fact of whether two nodes have exchanged messages (in either direction) or not, respectively.

C Teacher Proof Reduction


Table 4: Teacher proof reduction statistics of non-bootstrapped model trained on unreduced $\mathbf{SR}(U(10, 40))$ dataset. Note that all rows, except for Total Reduction, are computed over the *reduced portion* of the dataset, i.e., the proofs that were successfully shortened by NeuRes.

(%)	TRAIN	TEST
PROOFS REDUCED	17.82	18.29
MAX. REDUCTION	86.11	76.4
AVG. REDUCTION	23.55	23.65
TOTAL REDUCTION	3.07	3.15

On rather interesting observation on Table 4 is that the model appears to be marginally better at producing shorter proofs for unseen (test) formulas than for training formulas. While we would normally expect the opposite, a fair speculation would be that the trained model was teacher-forced to match teacher proofs during training over multiple epochs while the same does not hold for unseen formulas where the bias towards teacher behavior is significantly lower. To definitively confirm this would require a more in-depth investigation.



(a) Teacher Proof



(b) NeuRes Proof

Figure 7: Teacher Proof Reduction Example