

Learning Better Representations From Less Data For Propositional Satisfiability

Mohamed Ghanem  

CISPA Helmholtz Center for Information Security, Saarbrücken, Germany

Frederik Schmitt  

CISPA Helmholtz Center for Information Security, Saarbrücken, Germany

Julian Siber  

CISPA Helmholtz Center for Information Security, Saarbrücken, Germany

Bernd Finkbeiner  

CISPA Helmholtz Center for Information Security, Saarbrücken, Germany

Technical University of Munich, Munich, Germany

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.

2012 ACM Subject Classification Computing methodologies \rightarrow Learning latent representations;
Theory of computation \rightarrow Automated reasoning

Keywords and phrases Satisfiability, Neuro-symbolic learning, Resolution proofs, Expert iteration, Graph neural networks, Attention

Digital Object Identifier 10.4230/LIPIcs.CVIT.2016.23

Related Version *Conference Version: NeurIPS 2024 proceedings*

Acknowledgements This work was supported by the European Research Council (ERC) Grant HYPER (No. 101055412).

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 [13]. 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 [12, 47], software verification [17], planning [27], and mathematical proof search [24]. Recently, SAT has also served as a litmus test for assessing the symbolic reasoning capabilities of neural models and a promising domain for



© Mohamed Ghanem, Frederik Schmitt, Julian Siber, and Bernd Finkbeiner;
licensed under Creative Commons License CC-BY 4.0

42nd Conference on Very Important Topics (CVIT 2016).

Editors: John Q. Open and Joan R. Access; Article No. 23; pp. 23:1–23:21

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

44 neuro-symbolic systems [43, 42, 1, 10, 36]. So far, neural models only provide limited, if
 45 any, justification for unsatisfiability predictions. NeuroCore [42], for example, predicts an
 46 unsatisfiable core, the verification of which can be as hard as solving the original problem. No
 47 certificates at all or certificates that are hard to check limit neural methods in a domain where
 48 correctness is critical and prevents close integrations with symbolic methods. Therefore, we
 49 propose a neuro-symbolic model that utilizes *resolution* to solve SAT problems by generating
 50 easy-to-check certificates.

51 A resolution proof is a sequence of case distinctions, each involving two clauses, that
 52 ends in the empty clause (falsum). This technique can also be used to prove satisfiability by
 53 exhaustively applying it until no further new resolution steps are possible and the empty
 54 clause has not been derived. Generating such a proof is an interesting problem from a
 55 neuro-symbolic perspective because unlike other discrete combinatorial problems that have
 56 been considered before [46, 7, 28, 30, 11], it requires selecting compatible *pairs* of clauses
 57 from the dynamically growing pool, as newly derived clauses are naturally considered for
 58 derivation in subsequent steps. In this work, we devise three attention-based mechanisms to
 59 perform this pair-selection needed for generating resolution proofs. In addition, we augment
 60 the architecture to efficiently handle *sat* (satisfiable) formulas with an assignment decoding
 61 mechanism that assigns a truth value to each literal. We hypothesize that, despite their final
 62 goals being in complete opposition, resolution and *sat* assignment finding can form a mutually
 63 beneficial collaboration. On the one hand, clauses derived by resolution incrementally inject
 64 additional information into the network, e.g., deriving a single-literal clause by resolution
 65 directly implies that literal should be true in any possible *sat* assignment. On the other hand,
 66 finding a *sat* assignment absolves the resolution network from having to prove satisfiability
 67 by exhaustion. On that basis, given an input formula, NeuRes proceeds in two parallel tracks:
 68 (1) finding a *sat* assignment, and (2) deriving a resolution proof of unsatisfiability. Both
 69 tracks operate on a shared representation of the problem state. Depending on which track
 70 succeeds, NeuRes produces the corresponding SAT verdict which is guaranteed to be sound
 71 by virtue of its certificate-based design. Since both of our certificate types are efficient to
 72 check, we can afford to perform these symbolic checks at each step. When comparing NeuRes
 73 with NeuroSAT [43], which has been trained to predict satisfiability with millions of samples,
 74 we demonstrate that NeuRes achieves a higher accuracy while providing a proof and requires
 75 only thousands of training samples.

76 As for most problems in theorem proving we are not only interested in finding any proof but
 77 a short proof. Resolution proofs can vary largely in their size depending on the resolution
 78 steps taken. Being able to efficiently check the proof, also allows us to adapt the proof target
 79 while training the model. In particular, we explore an expert iteration mechanism [2, 39]
 80 that pre-rolls the resolution proof of the model and replaces the target proof whenever the
 81 pre-rolled proof is shorter. We demonstrate that this bootstrapping mechanism iteratively
 82 shortens the proofs of our training dataset while further improving the overall performance
 83 of the model.

84 We make the following contributions:

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

92 The implementation of our framework can be found at [https://github.com/Oschart/](https://github.com/Oschart/NeuRes)
93 NeuRes.

94 **2 Related Work**

95 **SAT Solving and Certificates.**

96 We refer to the annual SAT competitions [5] for a comprehensive overview on the ever-evolving
97 landscape of SAT solvers, benchmarks, and proof checkers. SAT solvers are complex systems
98 with a documented history of bugs [9, 26], hence proof certificates have been partially required
99 in this competition since 2013 [3]. Unlike satisfiable formulas, there are several ways to certify
100 unsatisfiable formulas [23]. Resolution proofs [52, 20] are easy to verify [15], but non-trivial
101 to generate from modern solvers based on the paradigm of conflict-driven clause learning [34].
102 Clausal proofs, e.g., in DRAT format [50], are easier to generate and space-efficient, but hard
103 to validate. Verifying the proofs can take longer than their discovery [22] and requires highly
104 optimized algorithms [31].

105 **Deep Learning for SAT Solving.**

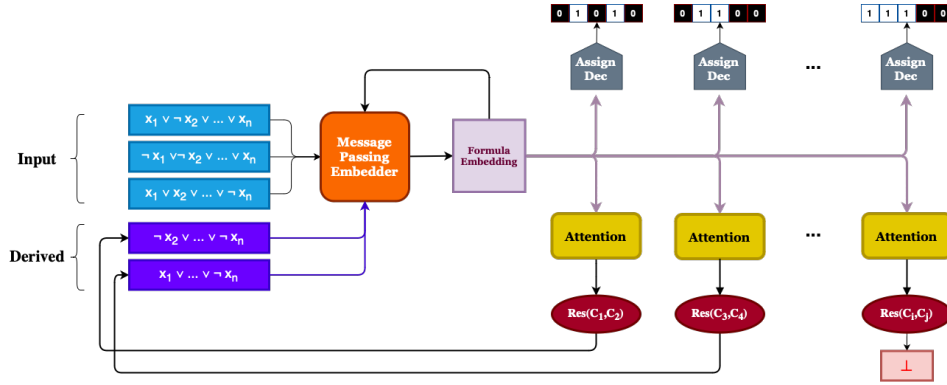
106 NeuroSAT [43] was the first study of the Boolean satisfiability problem as an end-to-end
107 learning task. Building upon the NeuroSAT architecture, a simplified version has been
108 trained to predict unsatisfiable cores and successfully integrated as a branching heuristic in
109 a state-of-the-art SAT solver [42]. Recent work has employed a related architecture as a
110 phase selection heuristic [49]. It has been shown that both the NeuroSAT architecture and a
111 newly introduced deep exchangeable architecture can outperform SAT solvers on instances
112 of 3-SAT problems [10]. The NeuroSAT architecture has also been applied on special classes
113 of crypto-analysis problems [44]. In addition to supervised learning, unsupervised methods
114 have been proposed for solving SAT problems. For Circuit-SAT a deep-gated DAG recursive
115 neural architecture has been presented together with a differentiable training objective to
116 optimize towards solving the Circuit-SAT problem and finding a satisfying assignment [1].
117 For Boolean satisfiability, a differentiable training objective has been proposed together with
118 a query mechanism that allows for recurrent solution trials [36].

119 **Deep Learning for Formal Proof Generation.**

120 In formal mathematics, deep learning has been integrated with theorem proving for clause
121 selection [33, 18], premise selection [25, 48, 6, 35], tactic prediction [51, 37] and whole proof
122 searches [40, 19]. For SMT formulas specifically, deep reinforcement learning has been
123 applied to tactic prediction [4]. In the domain of quantified boolean formulas, heuristics have
124 been learned to guide search algorithms in proving the satisfiability and unsatisfiability of
125 formulas [32]. For temporal logics, deep learning has been applied to prove the satisfiability
126 of linear-time temporal logic formulas and the realizability of specifications [21, 41, 14].

127 **3 Proofs of (Un-)Satisfiability**

128 We start with a brief review of certifying the (un-)satisfiability of propositional formulas in
129 conjunctive normal form. For a set of Boolean variables V , we identify with each variable
130 $x \in V$ the *positive literal* x and the *negative literal* $\neg x$ denoted by \bar{x} . A *clause* corresponds
131 to a disjunction of literals and is abbreviated by a set of literals, e.g., $\{\bar{1}, 3\}$ represents
132 $(\neg x_1 \vee x_3)$. A formula in *conjunctive normal form (CNF)* is a conjunction of clauses and is



■ Figure 1 Overall NeuRes architecture.

133 abbreviated by a set of clauses, e.g., $\{\{\bar{1}, 3\}, \{1, 2, \bar{4}\}\}$ represents $(\neg x_1 \vee x_3) \wedge (x_1 \vee x_2 \vee \neg x_4)$.
 134 Any Boolean formula can be converted to an equisatisfiable CNF formula in polynomial time,
 135 for example with Tseitin transformation [45].

136 A CNF formula is *satisfiable* if there exists an *assignment* $\mathcal{A} : V \rightarrow \{\top, \perp\}$ such that
 137 all clauses are satisfied, i.e., each clause contains a positive literal x such that $\mathcal{A}(x) = \top$ or
 138 a negative literal \bar{x} such that $\mathcal{A}(x) = \perp$. If no such assignment exists we call the formula
 139 *unsatisfiable*. To prove unsatisfiability we rely on resolution, a fundamental inference rule in
 140 satisfiability testing [16]. The resolution rule (*Res*) picks clauses with two opposite literals
 141 and performs the following inference:

$$142 \frac{C_1 \cup \{x\} \quad C_2 \cup \{\bar{x}\}}{C_1 \cup C_2} Res$$

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

147 4 Models

148 4.1 General Architecture

149 NeuRes is a neural network that takes a CNF formula as a set of clauses and outputs either
 150 a satisfying truth assignment or a resolution proof of unsatisfiability. As such, our model
 151 comprises a formula embedder connected to two downstream heads: (1) an attention network
 152 responsible for selecting clause pairs, and (2) a truth assignment decoder. See Figure 1 for an
 153 overview of the NeuRes architecture. After obtaining the initial clause and literal embeddings
 154 (representing the input formula), we continue with the iterative certificate generation phase.
 155 At each step, the model selects a clause pair which gets resolved into a new clause to append
 156 to the current formula graph while decoding a candidate truth assignment in parallel. The
 157 model keeps deriving new clauses until the empty clause is found (marking resolution proof
 158 completion), a satisfying assignment is found (marking a certified *sat* verdict), or the limit
 159 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 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*.

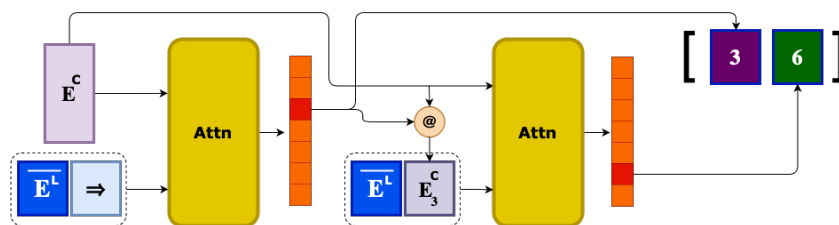
In a *static* approach, we do not change the embeddings of initial clauses upon inferring a new clause. Instead, we exchange local messages between the node corresponding to the new clause and its literal nodes, in both directions. The main advantage of this approach is its low cost. A major drawback 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

23:6 Learning Data-efficient Representations for Propositional Satisfiability

195 of the literal embeddings $\overline{E^L}$ concatenated with a zero vector while performing the second
 196 query using the mean of the literal embeddings concatenated with the embedding vector
 197 $E_{c_1}^C$ of the clause selected in the first query. Formally, Casc-Attn selects a clause index pair
 198 (c_1, c_2) as follows:


$$199 \quad c_i = \underset{j}{\mathbf{argmax}} [u^T \tanh(W_1 q_i + W_2 E_j^C)] \quad \text{with} \quad q_i = \begin{cases} \overline{E^L} \parallel \mathbf{0} & \text{if } i = 1 \\ \overline{E^L} \parallel E_{c_1}^C & \text{if } i = 2 \end{cases} \quad (1)$$

200 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.

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

205 4.3.2 Full Self-Attention (Full-Attn)

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



211 **Figure 3** Full self-attention.

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


$$212 \quad (c_1, c_2) = \underset{(i,j)}{\mathbf{argmax}} S_{i,j} \quad \text{with} \quad Q = E^C W_Q; \quad K = E^C W_K; \quad S = \frac{QK^T}{\sqrt{d}} \quad (2)$$

213 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
 214 cells that correspond to invalid resolution steps (i.e., clause pairs that cannot be resolved),
 215 we mask out the invalid cells from the attention grid to ensure the network selection is valid
 216 at every step.

217 4.3.3 Anchored Self-Attention (Anch-Attn)

218 In Full-Attn, the attention grid grows quadratically with the number of clauses. In this
 219 variant, we relax this cost by exploiting a property of binary resolution where each step
 220 targets a single variable in the two resolvent clauses. This allows us to narrow down candidate
 221 clause pairs by first selecting a variable as an anchor on which our clauses should be resolved.
 222 As such, we do not need to consider the full clause set at once, only the clauses containing
 223 the chosen variable v . We further compress the attention grid by lining clauses containing

224 the literal v on rows while lining clauses containing the literal $\neg v$ on columns. This reduces
 225 the redundancy of the attention grid since clauses containing the variable v with the same
 226 parity cannot be resolved on v , so there is no point in matching them. In this scheme, we
 227 have two attention modules: one attention network to choose an anchor variable followed by
 228 a self-attention network to produce the anchored score grid.



■ **Figure 4** Anchored self-attention.

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

$$237 \quad v = \underset{i}{\operatorname{argmax}} \left[u^T \tanh(W_1 \overline{E^C} + W_2 (E_i^{L^+} + E_i^{L^-})) \right] \quad (3)$$

238 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
 239 pair (c_1, c_2) is then selected according to the same equations of Full-Attn (Eq. 2) using the
 240 v -anchored set of clause embeddings.

241 4.4 Assignment Decoder

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

$$244 \quad \hat{A}(l_i) = \sigma(\psi(E_i^L)) \quad (4)$$

245 Note that since for each variable, we have a positive and a negative literal embeddings,
 246 we can construct two different truth assignments at a time using this method. However,
 247 supervising both assignments did not improve the performance compared to only supervising
 248 the positive assignment (on positive literal embeddings). Thus, to simplify our loss function,
 249 we only derive truth assignments from the positive literal embeddings at train time while
 250 extracting both at test time. Interestingly, at test time, we found using negative literals (in
 251 addition to the positive ones) sometimes produces satisfying assignments before the positive
 252 branch despite receiving no direct supervision during training. Our intuition regarding this
 253 observation attributes it to the fact that the formula graph has no explicit notion of positive
 254 and negative literals, it only represents connections to clauses (positive and negative literals
 255 are connected by an undirected edge that does not distinguish their parity). As such, both
 256 literal nodes have a different local view into the rest of formula, which could result in one of
 257 them leading to a satisfying assignment faster than the other.

258 **5 Training and Hyperparameters**

259 **5.1 Dataset**

260 For our training and testing data, we adopt the same formula generation method as
 261 NeuroSAT [43], namely $\mathbf{SR}(n)$ where n is the number of variables in the formula. This
 262 method was designed to generate a generalized formula distribution that is not limited to a
 263 particular domain of SAT problems. To control our data distributions, we vary the range
 264 on the number of Boolean variables involved in each formula. For our training data, we
 265 use formulas in $\mathbf{SR}(U(10, 40))$ where $U(10, 40)$ denotes the uniform distribution on integers
 266 between 10 and 40 (inclusive). To generate our teacher certificates comprising resolution
 267 proofs and truth assignments, we use the BooleForce solver [8] on the formulas generated on
 268 the \mathbf{SR} distribution.

269 **5.2 Loss Function**

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

$$275 \quad \mathcal{L}_{Res} = -\frac{1}{T} \sum_t \log(p(y_t; \theta)) \cdot \gamma^{(T-t)} \quad (5)$$

276 During *sat* episodes, we minimize \mathcal{L}_{sat} computed as the binary cross-entropy loss between the
 277 sigmoid-activated outputs of assignment decoder $\hat{\mathcal{A}} : V \rightarrow [0, 1]$ and the teacher assignment
 278 $\mathcal{A} : V \rightarrow \{0, 1\}$ as shown in Eq. 6.

$$279 \quad \mathcal{L}_{sat} = \frac{1}{T} \sum_t \left[\frac{\gamma^{(T-t)}}{|V|} \sum_v \text{BCE}(\hat{\mathcal{A}}(v), \mathcal{A}(v)) \right] \quad (6)$$

280 In both types of episodes, step-wise losses are weighted by a time-horizon discounting factor
 281 $\gamma < 1.0$ over the whole episode. The main rationale behind this is that later losses should
 282 have higher weights as the formula tends to get easier to solve with each new clause inferred
 283 by resolution.

284 **5.3 Hyperparameters**

285 NeuRes has several hyperparameters that influence network size, depth, and loss weighting.
 286 In the experiments we fix the embedding dimension to 128. We train our models with a
 287 batch size of 1 and the Adam optimizer [29] for 50 epochs which took about six days on a
 288 single NVIDIA A100 GPU. We linearly anneal the learning rate from 5×10^{-5} to zero over
 289 the training episodes. This empirically yields better results than using a constant learning
 290 rate. We use a time discounting factor $\gamma = 0.99$ for the episodic loss. We apply global-norm
 291 gradient clipping with a ratio of 0.5 [38].

292 **6 Generating Resolution Proofs**

293 NeuRes uses resolution as the core reasoning technique for certificate generation, both in
 294 the *unsat* and *sat* cases. Hence, we start with an in-depth comparative evaluation of several

295 internal variants for resolution only. In particular, we evaluate the success rate (i.e., problems
 296 solved before timeout) and proof length relative to the teacher, denoted by $\text{p-Len} = \frac{|\mathcal{P}_{\text{NeuRes}}|}{|\mathcal{P}_{\text{teacher}}|}$.
 297 We use a limit of 4 on this ratio as a timeout to avoid simply brute-forcing a resolution
 298 proof. Note that we measure p-Len only for solved formulas to avoid diluting the average
 299 with resolution trails that timed out. For experiments in this section, we train on 8K *unsat*
 300 formulas in $\text{SR}(U(10, 40))$ and test our models on 10K unseen formulas belonging to the
 301 same distribution. We use more formulas than the model was trained on to more reliably
 302 demonstrate its learning capacity.

■ **Table 1** Performance of all attention variants on *unsat* $\text{SR}(U(10, 40))$ 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

303 6.1 Attention Variants

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

315 6.2 Shortening Teacher Proofs with Bootstrapping

316 During our initial experiments, we discovered proofs produced by NeuRes that were shorter
 317 than the corresponding teacher proofs in the training data. Although teacher proofs were
 318 generated by a traditional SAT solver, they are not guaranteed to be size-optimal. The size
 319 of resolution proofs is their only real drawback, hence any method that can reduce this size
 320 would be immensely useful. Upon closer inspection we find that, on average, our previous
 321 best performer trained with regular teacher-forcing manages to shorten $\sim 18\%$ of teacher
 322 proofs by a notable factor (cf. Appendix C). This inspired us to devise a bootstrapped
 323 training procedure to capitalize on this feature: We pre-roll each input problem using model
 324 actions only, and whenever the model proof is shorter than the teacher’s, it replaces the
 325 teacher’s proof in the dataset. In other words, we maximize the likelihood of the shorter
 326 proof. In doing so iteratively, the model progressively becomes its own teacher by exploiting
 327 redundancies in the teacher algorithm.

■ **Table 2** Bootstrapped training data reduction statistics. Reduction statistics are computed on the $\text{SR}(U(10, 40))$ 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

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

339 In Appendix C, we have compiled additional statistics (cf. Table 5) on proof shortening
 340 during the training process, as well as an example proof reduced by the bootstrapped NeuRes
 341 (Figure 7). We only include a small reduction example (from 20 steps to 10 steps) for space
 342 constraints, but we observed many more examples of much larger reductions (e.g., over 400
 343 steps).

344 7 Resolution-Aided SAT Solving

345 In this section, we evaluate the performance of our fully integrated model trained on a
 346 hybrid dataset comprising 8K unsatisfiable formulas (and their resolution proofs) and 8K
 347 satisfiable formulas (and their satisfying assignments). For the unsatisfiable formulas, timeout
 348 ($4 \times |\mathcal{P}_{\text{teacher}}|$) and optimality (p-Len) are measured similarly to previous experiments. For
 349 satisfiable formulas, we set the timeout (maximum #trials) to $2 \times |V|$. Ultimately, this
 350 section aims to investigate the effect of incorporating a certificate-driven downstream head
 351 on the quality of the learnt representations through its impact on the performance of the
 352 complementary task, i.e., proving/predicting satisfiability. We use NeuroSAT as our baseline
 353 as it employs the same formula embedding architecture. Since NeuroSAT proves *sat* but only
 354 *predicts unsat*, we train a classification MLP on top of our trained NeuRes model to further
 355 showcase the benefit of our representations on prediction accuracy.

356 Table 3 confirms this main hypothesis. In essence, this result points to the fact that
 357 learning signals obtained from training on *unsat* certificates largely enhance the ability of
 358 the neural network to extract useful information from the input formula. This is doubly
 359 promising considering NeuroSAT was trained on *millions* of formulas while NeuRes was
 360 trained on only *16K* formulas. Lastly, we find that augmenting *sat* formulas by resolution

■ **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 [43]	70	-	-	73	96	85

361 derivations results in relative improvements ($\sim 2.3\%$) in success rate even though these
 362 derivations are attempting to prove unsatisfiability.

363 8 Utilizing Model Fan-Out

364 In our Full-Attention module, we compute n^2 scores and only perform the top-score resolution
 365 step. This greedy approach arguably underutilizes the attention grid computations as it
 366 ignores other high-scoring steps that might lead to a shorter proof thereby improving the
 367 success rate in addition to reducing the number of queries to the model. The latter leads
 368 to an overall runtime reduction since performing an extra symbolic resolution step is much
 369 faster than a forward model pass. As such, we experiment with performing the top k steps
 370 of the attention grid after each forward pass. It should be noted, however, that this yields
 371 diminishing returns as it leads to a faster growth of the clause base which in turn inflates
 372 the attention grid. For $k > 1$, after deriving the empty clause, only clauses that connect to
 373 it in the resolution graph are kept in the final proof. This post-processing step is linear in
 374 the proof length and eliminates redundant resolution steps resulting from the higher fan-out.
 375 Table 4 shows that taking the top 3 steps already yields a massive reduction in proof lengths
 376 along with a significant boost to the success rate.

■ **Table 4** Performance of different model fan-outs on **SR**(40) test data. Proof length (p-Len) and #Model Calls are both normalized by the length of the teacher proof.

Full-Attn Fan-out	p-Len	#Model Calls	Total Proven (%)
Top-1	1.15	1.15	98.2
Top-3	0.57	0.49	99.9
Top-5	0.52	0.43	100.0

377 One way to offset the attention grid inflation with higher fan-out would be to keep a
 378 saliency map for all clauses then discarding k clauses with the least saliency scores after
 379 each forward pass. One simple way to compute this saliency score for a clause would be the
 380 sum/mean/max of its respective row in the attention scores grid. Another proxy for saliency
 381 could be the recency score reflecting how many steps have elapsed since the last time a given
 382 clause was used.

9 Generalizing to Larger Problems

In order to test our model’s out-of-distribution performance, we evaluate our NeuRes model 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 [43], 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 **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.

10 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 representations.

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 rather large, 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 into 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

413 results of traditional solvers. This proof reduction is facilitated by a bootstrapped training
 414 procedure that uses teacher proofs as a guide as opposed to a golden standard.

 415 ——— **References** ———

- 416 1 Saeed Amizadeh, Sergiy Matuskevych, and Markus Weimer. Learning to solve circuit-sat: An
 417 unsupervised differentiable approach. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. URL:
 418 <https://openreview.net/forum?id=BJxgz2R9t7>.
 419
- 420 2 Thomas Anthony, Zheng Tian, and David Barber. Thinking fast and slow with
 421 deep learning and tree search. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman 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 5360–5370, 2017. URL: <https://proceedings.neurips.cc/paper/2017/hash/d8e1344e27a5b08cdfd5d027d9b8d6de-Abstract.html>.
 422
 423
 424
 425
 426
- 427 3 Adrian Balint, Anton Belov, Marijn J.H. Heule, and Matti 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.
 428
 429
- 430 4 Mislav Balunovic, Pavol Bielik, and Martin T. Vechev. Learning to solve SMT formulas. In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman 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. URL: <https://proceedings.neurips.cc/paper/2018/hash/68331ff0427b551b68e911eebe35233b-Abstract.html>.
 431
 432
 433
 434
 435
- 436 5 Tomas Balyo, Marijn Heule, Markus Iser, Matti Järvisalo, and Martin 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.
 437
 438
 439
- 440 6 Kshitij Bansal, Sarah M. Loos, Markus N. Rabe, Christian Szegedy, and Stewart Wilcox. Holist: An environment for machine learning of higher order logic theorem proving. In Kamalika Chaudhuri and Ruslan 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. URL: <http://proceedings.mlr.press/v97/bansal19a.html>.
 441
 442
 443
 444
 445
- 446 7 Irwan Bello, Hieu Pham, Quoc V. Le, Mohammad Norouzi, and Samy 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. URL: <https://openreview.net/forum?id=Bk9mx1SfX>.
 447
 448
 449
- 450 8 Armin Biere. Booleforce sat solver. <https://fmv.jku.at/booleforce/>, 2010.
 451
- 452 9 Robert Brummayer, Florian Lonsing, and Armin Biere. Automated testing and debugging of SAT and QBF solvers. In Ofer Strichman and Stefan 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. doi:10.1007/978-3-642-14186-7_6.
 453
 454
 455
- 456 10 Chris Cameron, Rex Chen, Jason S. Hartford, and Kevin 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. doi:10.1609/aaai.v34i04.5733.
 457
 458
 459
- 460 11 Quentin Cappart, Didier Chételat, Elias B. Khalil, Andrea Lodi, Christopher Morris, and Petar Velickovic. Combinatorial optimization and reasoning with graph neural networks. *J. Mach. Learn. Res.*, 24:130:1–130:61, 2023. URL: <http://jmlr.org/papers/v24/21-0449.html>.
 461
 462


- 463 12 Edmund M. Clarke, Armin Biere, Richard Raimi, and Yunshan Zhu. Bounded model checking
464 using satisfiability solving. *Formal Methods Syst. Des.*, 19(1):7–34, 2001. doi:10.1023/A:
465 1011276507260.
- 466 13 Stephen A. Cook. The complexity of theorem-proving procedures. In Michael A. Harrison,
467 Ranan B. Banerji, and Jeffrey D. Ullman, editors, *Proceedings of the 3rd Annual ACM*
468 *Symposium on Theory of Computing, May 3-5, 1971, Shaker Heights, Ohio, USA*, pages
469 151–158. ACM, 1971. doi:10.1145/800157.805047.
- 470 14 Matthias Cosler, Frederik Schmitt, Christopher Hahn, and Bernd Finkbeiner. Iterative circuit
471 repair against formal specifications. In *The Eleventh International Conference on Learning*
472 *Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023. URL:
473 <https://openreview.net/pdf?id=SEcSah10Q1>.
- 474 15 Ashish Darbari, Bernd Fischer, and João Marques-Silva. Industrial-strength certified SAT
475 solving through verified SAT proof checking. In Ana Cavalcanti, David Déharbe, Marie-
476 Claude Gaudel, and Jim Woodcock, editors, *Theoretical Aspects of Computing - ICTAC*
477 *2010, 7th International Colloquium, Natal, Rio Grande do Norte, Brazil, September 1-3, 2010*.
478 *Proceedings*, volume 6255 of *Lecture Notes in Computer Science*, pages 260–274. Springer,
479 2010. doi:10.1007/978-3-642-14808-8_18.
- 480 16 Martin Davis and Hilary Putnam. A computing procedure for quantification theory. *J. ACM*,
481 7(3):201–215, 1960. doi:10.1145/321033.321034.
- 482 17 Leonardo Mendonça de Moura and Nikolaj S. Bjørner. Satisfiability modulo theories: introduc-
483 tion and applications. *Commun. ACM*, 54(9):69–77, 2011. doi:10.1145/1995376.1995394.
- 484 18 Vlad Firoiu, Eser Aygün, Ankit Anand, Zafarali Ahmed, Xavier Glorot, Laurent Orseau,
485 Lei M. Zhang, Doina Precup, and Shibl Mourad. Training a first-order theorem prover from
486 synthetic data. *CoRR*, abs/2103.03798, 2021. URL: <https://arxiv.org/abs/2103.03798>,
487 arXiv:2103.03798.
- 488 19 Emily First, Markus N. Rabe, Talia Ringer, and Yuriy Brun. Baldur: Whole-proof generation
489 and repair with large language models. *CoRR*, abs/2303.04910, 2023. arXiv:2303.04910,
490 doi:10.48550/arXiv.2303.04910.
- 491 20 Evgenii I. Goldberg and Yakov Novikov. Verification of proofs of unsatisfiability for CNF
492 formulas. In *2003 Design, Automation and Test in Europe Conference and Exposition*
493 *(DATE 2003), 3-7 March 2003, Munich, Germany*, pages 10886–10891. IEEE Computer
494 Society, 2003. URL: <https://doi.ieeecomputersociety.org/10.1109/DATE.2003.10008>,
495 doi:10.1109/DATE.2003.10008.
- 496 21 Christopher Hahn, Frederik Schmitt, Jens U. Kreber, Markus Norman Rabe, and Bernd
497 Finkbeiner. Teaching temporal logics to neural networks. In *9th International Conference on*
498 *Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net,
499 2021. URL: <https://openreview.net/forum?id=d0cQK-f4byz>.
- 500 22 Marijn Heule, Warren A. Hunt Jr., and Nathan Wetzler. Bridging the gap between easy
501 generation and efficient verification of unsatisfiability proofs. *Softw. Test. Verification Reliab.*,
502 24(8):593–607, 2014. doi:10.1002/stvr.1549.
- 503 23 Marijn J. H. Heule. Proofs of unsatisfiability. In Armin Biere, Marijn Heule, Hans van
504 Maaren, and Toby Walsh, editors, *Handbook of Satisfiability - Second Edition*, volume 336
505 of *Frontiers in Artificial Intelligence and Applications*, pages 635–668. IOS Press, 2021.
506 doi:10.3233/FAIA200998.
- 507 24 Marijn J. H. Heule and Oliver Kullmann. The science of brute force. *Commun. ACM*,
508 60(8):70–79, 2017. doi:10.1145/3107239.
- 509 25 Geoffrey Irving, Christian Szegedy, Alexander A. Alemi, Niklas Eén, François Chol-
510 let, and Josef Urban. Deepmath - deep sequence models for premise selection. In
511 Daniel D. Lee, Masashi Sugiyama, Ulrike von Luxburg, Isabelle Guyon, and Roman Gar-
512 nett, editors, *Advances in Neural Information Processing Systems 29: Annual Confer-*
513 *ence on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona*,

- 514 *Spain*, pages 2235–2243, 2016. URL: <https://proceedings.neurips.cc/paper/2016/hash/f197002b9a0853eca5e046d9ca4663d5-Abstract.html>.
- 515
- 516 **26** Matti Järvisalo, Marijn Heule, and Armin Biere. Inprocessing rules. In Bernhard Gramlich,
517 Dale Miller, and Uli Sattler, editors, *Automated Reasoning - 6th International Joint Conference,*
518 *IJCAR 2012, Manchester, UK, June 26-29, 2012. Proceedings*, volume 7364 of *Lecture Notes*
519 *in Computer Science*, pages 355–370. Springer, 2012. doi:10.1007/978-3-642-31365-3_28.
- 520 **27** Henry A. Kautz and Bart Selman. Planning as satisfiability. In Bernd Neumann, editor, *10th*
521 *European Conference on Artificial Intelligence, ECAI 92, Vienna, Austria, August 3-7, 1992.*
522 *Proceedings*, pages 359–363. John Wiley and Sons, 1992.
- 523 **28** Elias B. Khalil, Hanjun Dai, Yuyu Zhang, Bistra Dilkina, and Le Song. Learning com-
524 binatorial optimization algorithms over graphs. In Isabelle Guyon, Ulrike von Luxburg,
525 Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Gar-
526 nett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference*
527 *on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA,*
528 *USA*, pages 6348–6358, 2017. URL: <https://proceedings.neurips.cc/paper/2017/hash/d9896106ca98d3d05b8cbdf4fd8b13a1-Abstract.html>.
- 529
- 530 **29** Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv*
531 *preprint arXiv:1412.6980*, 2014.
- 532 **30** Wouter Kool, Herke van Hoof, and Max Welling. Attention, learn to solve routing problems!
533 In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA,*
534 *USA, May 6-9, 2019*. OpenReview.net, 2019. URL: [https://openreview.net/forum?id=](https://openreview.net/forum?id=ByxBFsRqYm)
535 [ByxBFsRqYm](https://openreview.net/forum?id=ByxBFsRqYm).
- 536 **31** Peter Lammich. Efficient verified (UN)SAT certificate checking. *J. Autom. Reason.*, 64(3):513–
537 532, 2020. doi:10.1007/s10817-019-09525-z.
- 538 **32** Gil Lederman, Markus N. Rabe, Sanjit Seshia, and Edward A. Lee. Learning heuristics for
539 quantified boolean formulas through reinforcement learning. In *8th International Confer-*
540 *ence on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020.*
541 OpenReview.net, 2020. URL: [https://openreview.net/forum?id=](https://openreview.net/forum?id=BJluxREKDB)
542 [BJluxREKDB](https://openreview.net/forum?id=BJluxREKDB).
- 543 **33** Sarah M. Loos, Geoffrey Irving, Christian Szegedy, and Cezary Kaliszyk. Deep network
544 guided proof search. In Thomas Eiter and David Sands, editors, *LPAR-21, 21st International*
545 *Conference on Logic for Programming, Artificial Intelligence and Reasoning, Maun, Botswana,*
546 *May 7-12, 2017*, volume 46 of *EPiC Series in Computing*, pages 85–105. EasyChair, 2017.
547 URL: <https://doi.org/10.29007/8mwc>, doi:10.29007/8MWC.
- 548 **34** João Marques-Silva, Inês Lynce, and Sharad Malik. Conflict-driven clause learning SAT solvers.
549 In Armin Biere, Marijn Heule, Hans van Maaren, and Toby Walsh, editors, *Handbook of*
550 *Satisfiability - Second Edition*, volume 336 of *Frontiers in Artificial Intelligence and Applications*,
551 pages 133–182. IOS Press, 2021. doi:10.3233/FAIA200987.
- 552 **35** Maciej Mikula, Szymon Antoniak, Szymon Tworkowski, Albert Qiaochu Jiang, Jin Peng
553 Zhou, Christian Szegedy, Lukasz Kucinski, Piotr Milos, and Yuhuai Wu. Magnusham-
554 mer: A transformer-based approach to premise selection. *CoRR*, abs/2303.04488, 2023.
555 URL: <https://doi.org/10.48550/arXiv.2303.04488>, arXiv:2303.04488, doi:10.48550/
ARXIV.2303.04488.
- 556 **36** Emils Ozolins, Karlis Freivalds, Andis Draguns, Eliza Gaile, Ronalds Zakovskis, and Sergejs
557 Kozlovics. Goal-aware neural SAT solver. In *International Joint Conference on Neural*
558 *Networks, IJCNN 2022, Padua, Italy, July 18-23, 2022*, pages 1–8. IEEE, 2022. doi:10.1109/
559 IJCNN55064.2022.9892733.
- 560 **37** Aditya Paliwal, Sarah M. Loos, Markus N. Rabe, Kshitij Bansal, and Christian Szegedy.
561 Graph representations for higher-order logic and theorem proving. In *The Thirty-Fourth AAAI*
562 *Conference on Artificial Intelligence, AAAI 2020, New York, NY, USA, February 7-12, 2020*,
563 pages 2967–2974. AAAI Press, 2020. URL: <https://doi.org/10.1609/aaai.v34i03.5689>,
564 doi:10.1609/AAAI.V34I03.5689.

- 565 38 Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. Understanding the exploding gradient
566 problem. *CoRR*, abs/1211.5063, 2(417):1, 2012.
- 567 39 Stanislas Polu, Jesse Michael Han, Kunhao Zheng, Mantas Baksys, Igor Babuschkin, and
568 Ilya Sutskever. Formal mathematics statement curriculum learning. In *The Eleventh International
569 Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*.
570 OpenReview.net, 2023. URL: <https://openreview.net/pdf?id=-P7G-8dmSh4>.
- 571 40 Stanislas Polu and Ilya Sutskever. Generative language modeling for automated theorem
572 proving. *CoRR*, abs/2009.03393, 2020. URL: <https://arxiv.org/abs/2009.03393>, arXiv:
573 2009.03393.
- 574 41 Frederik Schmitt, Christopher Hahn, Markus N. Rabe, and Bernd Finkbeiner. Neural
575 circuit synthesis from specification patterns. In Marc’Aurelio Ranzato, Alina Beygelz-
576 imer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan, editors, *Ad-
577 vances in Neural Information Processing Systems 34: Annual Conference on Neural
578 Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*,
579 pages 15408–15420, 2021. URL: [https://proceedings.neurips.cc/paper/2021/hash/
580 8230bea7d54bcdf99cdf99cb07313d5-Abstract.html](https://proceedings.neurips.cc/paper/2021/hash/8230bea7d54bcdf99cdf99cb07313d5-Abstract.html).
- 581 42 Daniel Selsam and Nikolaj S. Bjørner. Guiding high-performance SAT solvers with unsat-
582 core predictions. In Mikolás Janota and Inês Lynce, editors, *Theory and Applications of
583 Satisfiability Testing - SAT 2019 - 22nd International Conference, SAT 2019, Lisbon, Portugal,
584 July 9-12, 2019, Proceedings*, volume 11628 of *Lecture Notes in Computer Science*, pages
585 336–353. Springer, 2019. doi:10.1007/978-3-030-24258-9_24.
- 586 43 Daniel Selsam, Matthew Lamm, Benedikt Bünz, Percy Liang, Leonardo de Moura, and
587 David L. Dill. Learning a SAT solver from single-bit supervision. In *7th International
588 Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*.
589 OpenReview.net, 2019. URL: https://openreview.net/forum?id=HJMC_iA5tm.
- 590 44 Ling Sun, David Gérard, Adrien Benamira, and Thomas Peyrin. Neurogift: Using a machine
591 learning based sat solver for cryptanalysis. In Shlomi Dolev, Vladimir Kolesnikov, Sachin
592 Lodha, and Gera Weiss, editors, *Cyber Security Cryptography and Machine Learning - Fourth
593 International Symposium, CSCML 2020, Be’er Sheva, Israel, July 2-3, 2020, Proceedings*,
594 volume 12161 of *Lecture Notes in Computer Science*, pages 62–84. Springer, 2020. doi:
595 10.1007/978-3-030-49785-9_5.
- 596 45 Grigori S Tseitin. On the complexity of derivation in propositional calculus. *Automation of
597 reasoning: 2: Classical papers on computational logic 1967–1970*, pages 466–483, 1983.
- 598 46 Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. Pointer networks. *Advances in neural
599 information processing systems*, 28, 2015.
- 600 47 Yakir Vizel, Georg Weissenbacher, and Sharad Malik. Boolean satisfiability solvers and their
601 applications in model checking. *Proc. IEEE*, 103(11):2021–2035, 2015. doi:10.1109/JPROC.
602 2015.2455034.
- 603 48 Mingzhe Wang, Yihe Tang, Jian Wang, and Jia Deng. Premise selection for theorem
604 proving by deep graph embedding. In Isabelle Guyon, Ulrike von Luxburg, Samy Ben-
605 gio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, ed-
606 itors, *Advances in Neural Information Processing Systems 30: Annual Conference on
607 Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA,
608 USA*, pages 2786–2796, 2017. URL: [https://proceedings.neurips.cc/paper/2017/hash/
609 18d10dc6e666eab6de9215ae5b3d54df-Abstract.html](https://proceedings.neurips.cc/paper/2017/hash/18d10dc6e666eab6de9215ae5b3d54df-Abstract.html).
- 610 49 Wenxi Wang, Yang Hu, Mohit Tiwari, Sarfraz Khurshid, Kenneth L. McMillan, and Risto
611 Miikkulainen. Neuroback: Improving CDCL SAT solving using graph neural networks. In
612 *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna,
613 Austria, May 7-11, 2024*. OpenReview.net, 2024. URL: [https://openreview.net/forum?id=
614 samyfu6G93](https://openreview.net/forum?id=samyfu6G93).
- 615 50 Nathan Wetzler, Marijn Heule, and Warren A. Hunt Jr. Drat-trim: Efficient checking and
616 trimming using expressive clausal proofs. In Carsten Sinz and Uwe Egly, editors, *Theory*

- 617 *and Applications of Satisfiability Testing - SAT 2014 - 17th International Conference, Held*
618 *as Part of the Vienna Summer of Logic, VSL 2014, Vienna, Austria, July 14-17, 2014.*
619 *Proceedings*, volume 8561 of *Lecture Notes in Computer Science*, pages 422–429. Springer,
620 2014. doi:10.1007/978-3-319-09284-3_31.
- 621 **51** Kaiyu Yang and Jia Deng. Learning to prove theorems via interacting with proof assistants. In
622 Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International*
623 *Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*,
624 volume 97 of *Proceedings of Machine Learning Research*, pages 6984–6994. PMLR, 2019. URL:
625 <http://proceedings.mlr.press/v97/yang19a.html>.
- 626 **52** Lintao Zhang and Sharad Malik. Validating SAT solvers using an independent resolution-based
627 checker: Practical implementations and other applications. In *2003 Design, Automation and*
628 *Test in Europe Conference and Exposition (DATE 2003), 3-7 March 2003, Munich, Germany*,
629 pages 10880–10885. IEEE Computer Society, 2003. URL: [https://doi.ieeecomputersociety.](https://doi.ieeecomputersociety.org/10.1109/DATE.2003.10014)
630 [org/10.1109/DATE.2003.10014](https://doi.ieeecomputersociety.org/10.1109/DATE.2003.10014), doi:10.1109/DATE.2003.10014.

631 Appendix

632 **A** NeuroSAT Formula Graph Construction

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

633 NeuroSAT-style formula graphs have two designated node types: clause nodes connected
 634 to the literal nodes corresponding to their constituent literals [43]. For example, in Figure 6,
 635 the clause contents are as follows: $c_1 = (a \vee b)$, $c_2 = (a \vee \bar{b})$, $c_3 = (\bar{a} \vee \bar{b})$. Each message-
 636 passing round involves two exchange phases: (1) Literal-to-Clause, and (2) Clause-to-Literal
 637 (and implicitly Literal-to-Complement). This construction is particularly efficient as it allows
 638 the message-passing protocol to cover the entire graph connectivity in at most $|V| + 1$ rounds
 639 where V is the set of variables in the formula.

640 **B** Clause Connectivity Under Static Embeddings

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

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

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

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

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

¹ We 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

One rather interesting observation on Table 5 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. Definitely confirming this would require a more in-depth investigation.

Table 5 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

D Runtime Comparison with Traditional Solver

In the following table, we compare the average runtimes of our top-1 Full-Attn, top-3 Full-Attn, and the traditional solver we used as a teacher (BooleForce) on our main $\mathbf{SR}(40)$ test dataset. For both Full-Attn models, we use our Python prototype implementation; for BooleForce, we use an official C implementation.

Table 6 Average time (ms) to solve an instance by neural model vs. teacher solver.

Solver	SAT (ms)	UNSAT (ms)	Total (ms)
FULL-ATTN TOP-1	2.3	88	45.15
FULL-ATTN TOP-3	3	54.4	28.7
BOOLEFORCE	4	5	4.5

E Model Size Comparison with NeuroSAT

In terms of the model architecture, both NeuRes and NeuroSAT models can be broken down to:

- **Embedding/Representation Network:** for both models, this network is an LSTM-based GNN that embeds the formula graph by message-passing. We use the exact same architecture and model size to ensure that our improved representations are a result of our fully certificate-based learning objective as opposed to a tweak in the model architecture. This GNN has 429,824 parameters in total.
- **Downstream Networks:** NeuroSAT: uses a 3-layer MLP applied on the literal embeddings (width = 128) to extract the literal votes to predict if the formula is satisfiable or


23:20 Learning Data-efficient Representations for Propositional Satisfiability

685 not. This MLP has $128 \times 128 \times 3 = 49,152$ parameters. NeuRes (Full-Attn): uses an
686 attention module to select clause pairs. This attention network is composed of two 1-layer
687 MLPs for the query and key transformations on the clauses embeddings (width = 128).
688 The whole attention module has $128 \times 128 \times 2 = 32,768$ parameters. To decode the
689 variable assignments, NeuRes uses a 2-layer scalar MLP with $128 \times 128 + 128 = 16,512$
690 parameters


691 Total NeuroSAT size = $429,824 + 49,152 = 478,976$ parameters

692 Total NeuRes size = $429,824 + 49,280 = 479,104$ parameters

693 All in all, NeuRes only learns 128 more parameters.



(a) Teacher Proof



(b) NeuRes Proof

Figure 7 Teacher Proof Reduction Example