arXiv:2110.14053v3 [cs.AI] 9 Jun 2022

NeuroComb: Improving SAT Solving with Graph Neural Networks

Wenxi Wang
The University of Texas at Austin
wenxiw@utexas.edu

Yang Hu
The University of Texas at Austin
huyang@utexas.edu

Mohit Tiwari
The University of Texas at Austin
tiwari@austin.utexas.edu

Sarfraz Khurshid
The University of Texas at Austin
khurshid@utexas.edu

Kenneth McMillan
The University of Texas at Austin
kenmcm@cs.utexas.edu

Risto Miikkulainen
The University of Texas at Austin
risto@cs.utexas.edu

Abstract

Propositional satisfiability (SAT) is an NP-complete problem that impacts many research fields, such as planning, verification, and security. Mainstream modern SAT solvers are based on the Conflict-Driven Clause Learning (CDCL) algorithm. Recent work aimed to enhance CDCL SAT solvers by improving their variable branching heuristics through predictions generated by Graph Neural Networks (GNNs). However, so far this approach either has not made solving more effective, or has required online access to substantial GPU resources. Aiming to make GNN improvements practical, this paper proposes an approach called NeuroComb, which builds on two insights: (1) predictions of important variables and clauses can be combined with dynamic branching into a more effective hybrid branching strategy, and (2) it is sufficient to query the neural model *only once* for the predictions before the SAT solving starts. NeuroComb is implemented as an enhancement to a classic CDCL solver called `MiniSat` and a more recent CDCL solver called `Glucose`. As a result, it allowed `MiniSat` to solve 11% and `Glucose` 5% more problems on the recent SATCOMP-2021 competition problem set, with the computational resource requirement of only one GPU. NeuroComb is therefore a both effective and practical approach to improving SAT solving through machine learning.

1 Introduction

Propositional satisfiability (SAT) solvers are designed to check the satisfiability of a given propositional logic formula. SAT solvers have advanced significantly in recent years, which has fueled their applications in a wide range of domains. Solvers need to be both effective, i.e., solve more problems within a given time, and practical, i.e., not require excessive computational resources to do so.

Mainstream modern SAT solvers are based on the CDCL algorithm [31]. They often use a variable branching heuristic called Variable State Independent Decaying Sum (VSIDS) [33] to decide which variables are best to branch on at any given point. This paper aims to make CDCL SAT solvers more effective by using machine learning to improve VSIDS.

Recently, Graph Neural Networks (GNNs) were proposed as a possible way to achieve this goal by training them to predict the best branching variables. This approach has not usually made solvers

Preprint. Under review.

more effective [19, 22, 13]. One exception is NeuroCore [38], which improves three solvers to solve 6% to 11% more problems in the SATCOMP-2018 [14] problem set (without taking into account the network transmission time). NeuroCore performs frequent online model inferences to adjust its predictions with dynamic information extracted from the SAT solving process. However, this online inference is computationally demanding, requiring network access to 20 GPUs distributed over five machines. Such requirement makes it difficult to apply NeuroCore in practice, and especially difficult to scale it to larger problems.

This paper proposes an approach called NeuroComb, which aims both to make solving more effective and to reduce the computational resource requirements, thus making the GNN approach more practical. This goal is achieved through two mechanisms: (1) learning to predict which variables and which clauses are important, and (2) performing that inference only once before the solving process and applying it repeatedly during the process.

First, to obtain useful predictions, instead of focusing on only a single prediction task as NeuroCore does, NeuroComb makes two predictions: (1) which variables are important (i.e., backbone variables and unsatisfiable core variables) and (2) which clauses are important. By converting the input SAT formula into a compact graph representation, the problem of predicting important variables and clauses turns into two binary node classification problems. To get high-accuracy predictions for SAT formulas with diverse scales, NeuroComb employs a densely connected GNN architecture with pooling, both to ensure the model expressiveness and to avoid over-smoothing. To train the model with supervised learning, a dataset called DataComb[1] containing 115,348 labeled SAT formulas with diversity was created from three different sources: Alloy [17], CNFgen [24], and SATLIB [15].

Second, the task of applying the predictions properly to SAT is accomplished in two steps: (1) information about important variables and clauses is statically extracted from the neural predictions before the solving begins, and (2) this information is applied *periodically* using multi-armed bandits during the solving process, guiding the solver to possible good directions. This guidance is achieved by a hybrid branching heuristic that is based not only on the original VSIDS dynamic information (the original dynamic heuristic), but also on the GNN-generated static information (the static heuristic).

NeuroComb is incorporated into two solvers: (1) a classic CDCL SAT solver called `MiniSat` [6] and (2) a more recent CDCL SAT solver called `Glucose` [2, 3], to form two new solvers called `NeuroComb-MiniSat` and `NeuroComb-Glucose`, respectively. The experimental results on all 400 SAT problems from the main track of SATCOMP-2021 [1] show that NeuroComb allows `MiniSat` to solve 11% and `Glucose` to solve 5% more problems. All experiments were run on an ordinary commodity computer with one NVIDIA GeForce RTX 3080 GPU (10GB memory), one AMD Ryzen Threadripper 3970X processor (64 logical cores), and 256GB RAM. The experiments thus demonstrate that NeuroComb is a practical approach to improving SAT solving through machine learning.

2 Background

This section introduces the SAT problem, CDCL algorithm, VSIDS heuristic, and basics of GNNs.

Preliminaries of SAT In SAT, a propositional logic formula ϕ is usually encoded in Conjunctive Normal Form (CNF), which is a conjunction ($\wedge$) of *clauses*. Each clause is a disjunction ($\vee$) of *literals*. A literal is either a variable v, or its complement $\neg v$. Each variable can be assigned a logical value, 1 (true) or 0 (false). A CNF formula has a satisfying assignment if and only if every clause has at least one true literal. For example, a CNF formula $\phi = (v_1 \vee \neg v_2) \wedge (v_2 \vee v_3)$ consists of two clauses $v_1 \vee \neg v_2$ and $v_2 \vee v_3$, four literals v_1, $\neg v_2$, v_2 and v_3, and three variables v_1, v_2, and v_3. One satisfying assignment of ϕ is $v_1 = 1$, $v_2 = 1$, $v_3 = 0$. The goal of a SAT solver is to check if a formula ϕ is satisfiable (sat) or unsatisfiable (unsat). A *complete* solver either outputs a satisfying assignment for ϕ, or proves that no such assignment exists. In the unsat case, many solvers can produce a subset of the clauses that is unsat, called an *unsat core*.

CDCL Algorithm CDCL makes SAT solvers efficient in practice and is one of the main reasons for the widespread of SAT applications. The general idea of CDCL algorithm is as follows (see [30] for details). First, it picks a variable on which to branch with the branching heuristic and decides a value to assign to it. It then conducts a Boolean propagation based on the decision. In the propagation,

---

[1] The DataComb dataset is publicly available at [anonymous; upon acceptance of the paper], intended as a general resource for researchers aiming to utilize deep learning to enhance SAT solving in the future.

if a conflict occurs (i.e., at least one clause is mapped to 0), it performs a conflict analysis; otherwise, it makes a new decision on another selected variable. In the conflict analysis, CDCL first analyzes the decisions and propagations to investigate the reason for the conflict, then extracts the most relevant wrong decisions, undoes them, and adds the reason to its memory as a learned lesson (encoded in a clause called *learned clause*) in order to avoid making the same mistake in the future. The above steps of clause learning in conflict analysis are called *resolution steps*. Since the memory space is limited, it forgets older learned clauses when necessary and keeps newer ones learned from more recent conflicts. The process continues until all the variables are assigned a value and no conflicts occur (in case the problem is sat), or until it learns the empty clause, or "false" (in case the problem is unsat). Note that modern CDCL solvers frequently restart during the solving, that is, it discards the current search and begins anew but keeping what has been learned.

The VSIDS Branching Heuristic The VSIDS heuristic [33] is the dominant branching heuristic in CDCL SAT solvers. Many high-performance branching heuristics are variants of VSIDS. The essence of VSIDS is an *additive bumping* and *multiplicative decay* behavior. VSIDS maintains an *activity score* for each variable in the Boolean formula. The score is typically initialized to 0. If a variable is involved in a resolution step of the conflict analysis, its activity score is additively bumped by a fixed increment. After every conflict analysis, regardless of the involvements, the activity score of every variable is decayed multiplicatively by a constant factor γ: $0 < \gamma < 1$. VSIDS selects the variable with the highest score on which to branch. The idea is to favor variables that participate in more recent conflict analyses.

Message Passing in GNN GNNs [43, 46] are a family of neural network architectures that operate on graphs [11, 37, 9, 4]. Typical GNNs follow a recursive neighborhood aggregation scheme called message passing [9]. Formally, the input of a GNN is a graph defined as a tuple $G = (V, E, W, H)$, where V denotes the set of nodes; $E \subseteq V \times V$ denotes the set of edges; $W = \{W_{u,v} | (u, v) \in E\}$ contains the feature vector $W_{u,v}$ of each edge (u, v); and $H = \{H_v | v \in V\}$ contains the feature vector H_v of each node v. A GNN maps each node to a vector-space embedding by updating the feature vector of the node iteratively based on its neighbors. For each iteration, a message-passing layer $\mathcal{L}$ takes a graph $G = (V, E, W, H)$ as an input and outputs a graph $G' = (V, E, W, H')$ with updated node feature vectors, i.e., $G' = \mathcal{L}(G)$. Classic GNN models [9, 21, 12] usually stack several message-passing layers to realize iterative updating. Each message-passing layer consists of two operations: message generation and message aggregation. For each edge $(x, y) \in E$, a message $M_{x,y}$ is generated via function $\mathcal{R}$:

$$M_{x,y} = \mathcal{R}(W_{x,y}, H_x, H_y). \tag{1}$$

For each node y, messages from all neighbors of y (denoted as N_y) are aggregated via function $\mathcal{Q}$ as

$$H'_y = \mathcal{Q}(H_y, \{M_{x,y} | x \in N_y\}). \tag{2}$$

Prior work on utilizing GNNs to improve CDCL SAT solving will be reviewed next.

3 Related Work

Recently, several approaches have been developed to utilize GNNs to facilitate CDCL SAT solving. NeuroSAT [39] was the first such framework adapting a neural model into an end-to-end SAT solver, which was not intended as a complete SAT solver. Other applications aim to provide SAT solvers with better heuristics, particularly the branching heuristic. Jaszczur et al. [19] used an architecture similar to NeuroSAT to enhance the branching heuristic in both DPLL [5] and CDCL algorithms. Kurin et al. [22] proposed a branching heuristic called GQSAT trained with value-based reinforcement learning. In NeuroGlue, Han [13] introduced a network to predict glue variables, i.e., those likely to occur in glue clauses, which are conflict clauses identified by the Glucose series of solvers [3]. GVE [45] is built upon NeuroGlue by adding an LSTM module that predicts the value of the glue variables. These approaches either reduce the number of solving iterations or enhance the solving effectiveness on selected small-scale problems with up to a few thousand variables. However, they do not provide obvious improvements in solving effectiveness for the large-scale problems.

In contrast, NeuroCore [38], the most closely related approach to this paper, aims to make the solving more effective especially for large-scale problems as in SAT competitions. It enhances VSIDS branching heuristic for CDCL using supervised learning to map unsat problems to unsat core variables (i.e., the variables involved in the unsat core). The authors are aware that perfect predictions of the unsat core do not always yield a useful branching heuristic. In some problems, the smallest core

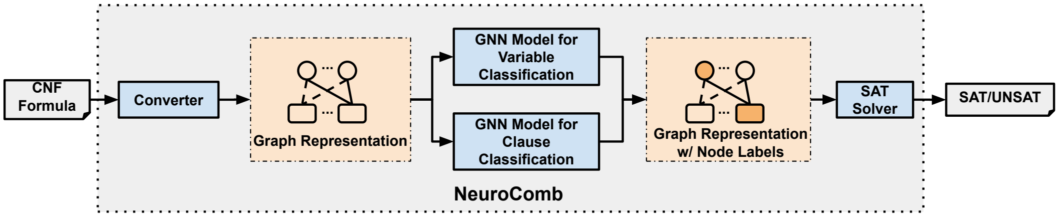
CNF Formula
Converter
Graph Representation
GNN Model for Variable Classification
GNN Model for Clause Classification
Graph Representation w/ Node Labels
SAT Solver
SAT/UNSAT
NeuroComb

Figure 1: Overview of NeuroComb. First, the input CNF formula is converted into a compact graph representation. Two trained GNN models are then applied once on the graph before SAT solving begins: one model classifies the variables in the graph and the other classifies the clauses. The SAT solver utilizes the resulting labeled graph periodically to inform its solving process. Thus, with the offline process of creating informative labels and the online process of utilizing them properly, NeuroComb makes the solving more effective and practical.

may include every variable, and in sat problems, there are no cores at all. Therefore, NeuroCore relies on *imperfect* prediction, with the hope that the variables assigned with higher probability correlate well with the variables that form a good basis for branching. Based on the dynamically learned clauses during the solving process, NeuroCore performs frequent online model inferences to tune the predictions. However, this online inference is computationally demanding, requiring 20 GPUs evenly distributed over five machines in their experiments. To mitigate this limitation, NeuroComb is proposed to utilize more informative offline predictions and apply them efficiently enough to make the approach practical.

4 NeuroComb

In order to reduce the computational cost of the online model inference and to make SAT solving more effective, NeuroComb utilizes static high-quality model predictions. Instead of focusing on only one term (e.g., unsat core variables), NeuroComb integrates information on the predictions of three terms (i.e., two terms for important variables and one term for important clauses), thus improving reliability. The model inference is performed only once before the SAT solving process, and the resulting offline predictions is applied periodically with multi-armed bandits during the solving. Fig. 1 shows the overview of NeuroComb. Its components are described in detail in the subsections below.

4.1 Graph Representation of CNF Formulas

As in recent work [23, 44], a SAT formula is represented using a more compact undirected bipartite graph than the one adopted in NeuroCore. Two node types represent the variables and clauses, respectively. Each edge connects a variable node to a clause node, representing that the clause contains the variable. Two edge types represent two polarities of a variable appearing in the clause, i.e., the variable itself and its complement. Formally, for a graph representation $G = (V, E, W, H)$ of a SAT formula, the edge feature $W_{u,v}$ of each edge (u, v) is initialized by its edge type with the value 1 representing positive polarity and -1 negative polarity; the node feature H_v of each node v is initialized by its node type with 0 representing a variable and 1 a clause.

4.2 Classifying Variables and Clauses

Important Variables We identify two types of important variables in SAT: *unsat core variables* for unsat problems as NeuroCore, and *backbone variables* for sat problems. The unsat core variables are important because they occur in the proof of unsatisfiability. Backbone variables [34] refer to variables that have the same value in every satisfying assignment. Note that backbone variables are seldom considered to improve the efficiency of modern SAT solvers, because computing the backbone variables is hard in general [36, 18, 20]. However, if backbone variables can be efficiently estimated, they can be helpful in reducing the search space and thus facilitating SAT solving.

Important Clauses An important insight in NeuroComb is that clauses involved in more resolution steps during conflict analysis are more important. Consequently, an importance score s_c of the clause c can be defined as $s_c = k_c/K$, where K denotes the total number of resolution steps in the entire SAT solving and k_c denotes the number of times that clause c participates in resolution steps. In contrast to the activity score with decay in CDCL SAT solvers which indicates how actively a variable

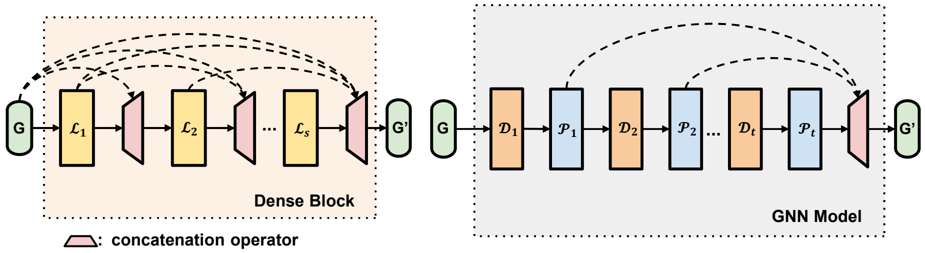

Figure 2: The architecture of the GNN model in NeuroComb. Labels $\mathcal{L}_1, \ldots, \mathcal{L}_s$ refer to message-passing layers, $\mathcal{D}_1, \ldots, \mathcal{D}_t$ to dense blocks, and $\mathcal{P}_1, \ldots, \mathcal{P}_t$ to pooling layers.

participates *locally* in recent conflict analyses, this importance score indicates how often the clause participates *globally* in conflict analyses during the entire SAT solving process. Clauses with an importance score higher than a threshold θ are deemed important. In the current implementation, θ is set to 0, indicating that any clauses participating at least once in the resolution steps are taken as important.

Classification Classifying important variables and clauses is taken as two binary node classification tasks. Two GNN models with the same architecture are trained with the same loss function type called Binary Cross Entropy (BCE) but targeting the corresponding node types. Given a graph G representing an input CNF formula, each GNN model outputs a graph G' with updated node feature vectors, which are then fed to a multi-layer perceptron model for node classification.

4.3 The GNN Model Architecture of NeuroComb

In this subsection, the architecture of the GNN model is introduced with three key components: message passing, dense blocks, and pooling.

Message Passing As introduced in Equation 1 and Equation 2, each message-passing layer $\mathcal{L}$ consists of two operations: message generation via function $\mathcal{R}$ and message aggregation via function $\mathcal{Q}$. Inspired by GraphSAGE [12], in NeuroComb, the message generation function $\mathcal{R}$ is realized with a multi-layer perceptron model $\mathcal{M}$:

$$M_{x,y} = \mathcal{R}(W_{x,y}, H_x, H_y) = \mathcal{M}(W_{x,y} \cdot H_x), \tag{3}$$

the aggregation function Q for each node y is realized with the mean of messages from all neighbors N_y:

$$\mathcal{Q}(H_y, \{M_{x,y} | x \in N_y\}) = \mathcal{A}(\mathcal{I}(\mathcal{B}(H_y || \frac{1}{|N_y|} \sum_{x \in N_y} M_{x,y}))), \tag{4}$$

where $\mathcal{A}$ denotes an activation function, $\mathcal{I}$ denotes a normalization function, $\mathcal{B}$ denotes a linear model, and $||$ denotes vector concatenation.

Dense Blocks Inspired by the design of Graph Convolutional Networks [21], multiple message passing layers are stacked in NeuroComb to make the model more expressive. One key hyperparameter is *model depth*, i.e., the number of message-passing layers. In order to let the messages pass sufficiently through the graph and thus enable each node to learn the whole graph structure, model depth is usually set to at least the diameter of the input graph [29]. However, due to the diverse CNF formula scales, corresponding graph representations have various diameters, which makes it difficult to determine what an ideal model depth is for all potential graph inputs. Deeper models usually lead to *over-smoothing* [26] when applied to smaller graphs, while shallower models may not be sufficiently expressive when applied to larger graphs.

In NeuroComb, a sufficiently deep model is used to make sure it is expressive enough, and *dense connectivity* is used to overcome over-smoothing. Inspired by DenseNet [16] and DeepGCN [25], a *dense block* is constructed by stacking and connecting the message-passing layers in a dense manner, as shown in Fig. 2 (left). Each message-passing layer obtains its inputs by concatenating the outputs from all its preceding message-passing layers. Formally, a dense block is recursively defined as

$$\mathcal{D}^{(s)}(G) = \begin{cases} \mathcal{L}_s(\mathcal{D}^{(s-1)}(G)) \parallel \mathcal{D}^{(s-1)}(G), & s > 0 \\ G, & s = 0, \end{cases} \tag{5}$$

where $\mathcal{D}^{(s)}$ denotes the dense block containing s message-passing layers, $\mathcal{L}_s$ denotes the s-th message passing layer, $\|$ is the graph concatenation operator satisfying $(V, E, W, H) \parallel (V, E, W, H') = (V, E, W, \{H_y \parallel_h H'_y | y \in V\})$, and $\|_h$ denotes the vector concatenation operator.

Pooling With dense blocks, a deep and expressive GNN model can be built. One straightforward way is to construct it as one deep dense block. However, this approach tends to produce *high-dimensional* node feature vectors. Such vectors would reduce the performance of the node classifier because of the curse of dimensionality [42]. To address this problem, *shallow* dense blocks are stacked and one pooling layer is added to the end of each block to reduce the dimension of the block output, as shown in the right side of Fig. 2. Instead of concatenating the outputs directly from all dense blocks, NeuroComb concatenates the outputs of all pooling layers as the final node feature vectors. Formally, a GNN model with pooling is defined as

$$\mathcal{J}^{(t)}(G) = \big\Vert_{i=1}^{t} \mathcal{P}_i(\mathcal{D}_i^{(s)}(\ldots \mathcal{P}_1(\mathcal{D}_1^{(s)}(G)))), \tag{6}$$

where t denotes the number of dense blocks, $\|$ the graph concatenation operator, $\mathcal{P}_i$ the i-th pooling layer, and $\mathcal{D}_i^{(s)}$ the i-th dense block, which consists of s message-passing layers. Since using the common local pooling layers could be less helpful in improving the GNN performance (as shown in the recent work by Mesquita et al. [32]), the message-passing layer, which outputs lower dimensional node feature vectors, is utilized as the pooling layer.

Implementation The current implementation of the GNN model consists of two dense blocks with three message-passing layers each ($s = 3$, $t = 2$). The activation function is Rectified Linear Unit (ReLU) [10]. The normalization function is instance normalization [41]. The model is implemented using PyTorch [35] and PyTorch Geometric [7].

Training The two GNN models for the classification tasks were trained on our entire DataComb dataset, using the AdamW optimizer [28] with an initial learning rate of 10^{-3} and a weight decay coefficient of 10^{-2}. The batch size was set to six and the number of epochs to 20 for each training run. The training for both tasks took 12 hours in total on the commodity computer.

4.4 Applying GNN Predictions in SAT

The goal is to utilize the important variable and clause predictions from the GNN models to improve the VSIDS variable branching heuristic of CDCL solvers. The predictions are used as static information characterizing the SAT problem. The insight is to utilize such static information effectively to guide the dynamic VSIDS variable branching heuristic.

Extracting Static Information The first step is to extract the static information from the GNN predictions. The idea is to make the final static scores robust to prediction errors by combining the important variable prediction with the important clause prediction. Higher scores are assigned to variables that are important (i.e., unsat core variables or backbone variables) and also appear more frequently in important clauses. Formally, the static score s_v of the variable v is defined as

$$s_v = \alpha p_v + (1 - \alpha)\mathcal{N}(\sum_{c \in C_v} p_c), \tag{7}$$

where $p_v \in [0, 1]$ denotes the important variable prediction score of variable v; $p_c \in [0, 1]$ denotes the important clause prediction score of clause c; C_v denotes the set of all clauses containing variable v; and $\alpha \in (0, 1)$ is a hyperparameter for the weighted mean between two prediction scores. Further, $\mathcal{N}$ denotes the function that normalizes the accumulated score $q_v = \sum_{c \in C_v} p_c$ to the range $[0, 1]$:

$$\mathcal{N}(q_v) = \frac{q_v - \min_{u \in U}(q_u)}{\max_{u \in U}(q_u) - \min_{u \in U}(q_u)}, \tag{8}$$

where U denotes the set of all variables in a SAT formula.

Applying Static Information with Hybrid Branching Next, the static information needs to be utilized to enhance the VSIDS branching heuristic. A hybrid branching heuristic is proposed to branch on variables not only based on their VSIDS activity scores (referred to as *dynamic branching*) but also based on their static scores (referred to as *static branching*, i.e., branching on the undecided variable with the highest static score). Dynamic branching remains the primary branching strategy but

Table 1: Details of the Source and Combined Datasets

Dataset	**Alloy**			**CNFgen**			**SATLIB**			**DataComb**		
# cnf	58,201			34,483			41,526			115,348		
Stats	**min**	**max**	**avg**	**min**	**max**	**avg**	**min**	**max**	**avg**	**min**	**max**	**avg**
# var	829	41,674	11,482	5	32,766	1,886	48	6,325	114	5	41,647	4,470
# cla	1,170	68,207	19,209	15	105,116	5,363	50	134,621	517	15	134,621	8,290

is interrupted periodically (with a period of β restarts) by static branching that lasts only a short time at each turn (with a duration denoted as γ). This hybrid strategy aims to guide the solver periodically to possible good directions by drawing attention to the important variables indicated by the static scores. The insight is that small and periodic perturbations produced by the static branching should not only change a number of branching decisions, but also slowly and profoundly influence the solver states because the new decisions that lead to conflicts affect the future VSIDS scores.

Deterministic Effect of Static Branching The time-limited effect of static branching is inspired by NeuroCore. However, this mechanism could make the solving process non-deterministic: SAT solvers conduct thousands of searches per second, the solver status could change even within a microsecond, and there is no accurate timing mechanism providing the guarantee that it does not. To make sure that the process remains deterministic, the effect duration γ is quantified with the number of branching decisions (denoted as γ_d), by estimating the frequency of decision generation (denoted as m_d) measured in decisions per second: $\gamma_d = \gamma m_d$.

This estimation is a regression problem and can be solved e.g., by the Random Forest machine learning method [27] trained to predict the decision frequency m_d based on features that describe the current solver status. These features include the number of undecided variables, the number of not-yet-satisfied original clauses, the number of literals in these clauses, the number of learned clauses, the average number of literals in each learned clause, the number of literals in conflict clauses, the number of decisions, and the number of propagations. To generate data, the target SAT solver was run on 1,200 SAT formulas randomly selected from DataComb with a timeout of 100 seconds, collecting data every 10 seconds, until 11,000 samples are collected. The mean squared error is applied as the loss function. The average R^2 score of the ten-fold cross validation on all generated samples for each targeted SAT solver is 0.991. The final decision frequency predictor is obtained by training the model on all samples.

Multi-Armed Bandit based Static Branching In static information extraction, there is a hyperparameter α for the weighted mean between two offline prediction scores. Choosing the value of α from candidate values $\alpha_1, ..., \alpha_k$ to apply in static branching can be seen as a k-armed bandits problem. Each arm $i \in \{1, ..., k-1\}$ represents applying the static information with the value α_i; a special arm k with the value $\alpha_k = -1$ represents applying the original VISIDS branching, considering the solving status to which the original branching can be beneficial. The reward $r_{i,t}$ is the SAT solving progress (estimated in the target SAT solver) made by pulling the arm i at time step t. The goal is to maximize the expected total reward, i.e., the expected total progress. Thompson Sampling [40] with Gaussian priors is used to solve the k-armed bandits problem. The distribution of reward $r_{i,t}$ is assumed to follow the Gaussian distribution $N(\mu_i, \sigma_i)$, where μ_i and σ_i are estimated as the mean and standard deviation, respectively, of all the rewards that arm i received up until time step t. In particular, for the first few steps, $r_{i,t}$ is assumed to follow the standard Gaussian distribution (i.e., μ_i is estimated as 0 and σ_i as 1).

Implementation There are two target SAT solvers: the classic `MiniSat` solver and the more recent `Glucose` solver, adapted to support the hybrid branching strategy. The implementations are called `NeuroComb-MiniSat` and `NeuroComb-Glucose`, respectively. For hyperparameter tuning, β was tried with 10, 20, and 30 restarts; γ with 1000^{-1}, 900^{-1}, and 800^{-1} seconds. Based on the average solving time on 100 selected hard problems from DataComb, β was set to 10 restarts for `NeuroComb-MiniSat` and 20 restarts for `NeuroComb-Glucose`; and γ was set to 800^{-1} seconds (1.25 milliseconds) for both solvers. In addition, 0.3, 0.4, 0.5, 0.6, 0.7, and -1 were used as candidate values for α, thus establishing a six-armed bandit problem.

5 Dataset

Training Dataset A new dataset containing CNF formulas with labeled important variables and clauses, called DataComb, was created for training our GNN model. The CNF formulas are obtained using three sources: a software modeling tool called Alloy [17], a recent CNF generator called

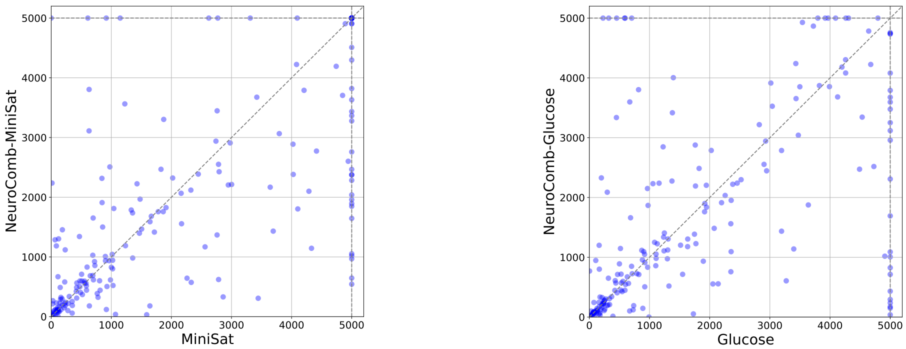

Figure 3: Time taken by `NeuroComb-MiniSat` vs. `MiniSat` (left) and `NeuroComb-Glucose` vs. `Glucose` (right) to solve each test problem in seconds (for problems that are solved by either solver). Each problem is represented by a dot whose location indicates the solution time of each method. The dots on the dashed lines at 5,000 seconds indicate failures. Baseline solvers fail on many instances that NeuroComb solvers solve in under 4,000 seconds, demonstrating the power of the approach.

CNFgen [24], and an online benchmark library called SATLIB [15]. Important variables are labeled in each CNF formula using two publicly available tools: MiniBones [36] to provide the backbone variables, and Kissat [8] to provide the unsat core variables. Important clauses are labeled based on the clause importance score provided by `MiniSat`. The timeout for generating the CNF formulas is 60 seconds and for getting the labels is 200 seconds.

Table 1 shows the details of the resulting dataset. Alloy is a mature toolset that models real-world problems from a wide range of applications and translates them into CNF formulas. As a base, 119 problems were first selected, ranging from security in protocols and type checking in programming languages to media asset management and network topology. Alloy models typically have scale parameters, such as the number of processes in a protocol, allowing us to generate problems of varying size. The sizes of 119 problems were then varied systematically, generating a total of 58,201 CNF formulas (with labels). CNFgen complements Alloy well in that it instead generates CNF formulas that appear in proof complexity literature. CNFgen generated 34,483 CNF formulas (with labels), grouped by six categories: pigeonhole principle, subset cardinality, counting principle, pebbling game, random k-CNF, and clique-coloring. SATLIB, on the other hand, is a library of known benchmark instances and therefore provides a baseline challenge. From SATLIB, 41,526 CNF formulas (with labels) were included, encoding four kinds of combinatorial problems: random-3-SAT, graph coloring, planning, and all-interval series. Altogether from these three sources, 134,210 labeled CNF formulas were obtained, consisting of 57,674 sat and 76,536 unsat formulas. To make the dataset balanced, 57,674 unsat formulas were randomly selected from the initial set, for a final DataComb dataset of 115,348 labeled CNF formulas. In Table 1, the number of variables and the number of clauses are used to characterize the size of CNF formulas. The formulas generated by Alloy are generally larger than those generated by CNFgen, which are in turn larger than those originating from SATLIB. DataComb thus contains a diverse set of formulas.

Testing Dataset To evaluate the effectiveness of `NeuroComb-MiniSat` and `NeuroComb-Glucose`, we choose all 400 CNF formulas from the main track of SATCOMP-2021 [1] as the testing data. Note that the testing dataset is not only an unseen set, but also completely independent from the training set, DataComb, whose generation does not consider any of the problems from previous SAT competitions (in contrast with e.g., NeuroCore experiments, where part of the training set was generated based on problems from SATCOMP-2017, and the testing set was from the main track of SATCOMP-2018). In particular, the average number of variables per formula is 42 times larger, and the average number of clauses 311 times larger, in the testing set than in the training set.

6 Experiments

The experiments aim to answer two research questions:
RQ1: *How accurately does the GNN model identify important variables and clauses?*
RQ2: *How effective is the NeuroComb approach?*

RQ1: GNN Model Performance The GNN model was evaluated with the ten-fold cross validation on DataComb. Table 2 shows the average (± stdev) prediction performance for both tasks in terms of precision, recall, F_1 score and accuracy. The model is able to classify 90.2% of the variables and

Table 2: Performance of the GNN model in Classifying Important Variables and Clauses

Stats Task	**Precision**	**Recall**	$\mathbf{F_1}$	**Accuracy**
Variable	0.941 ±0.015	0.778 ±0.014	0.852 ±0.005	0.902 ±0.003
Clause	0.929 ±0.005	0.820 ±0.008	0.871 ±0.003	0.937 ±0.001

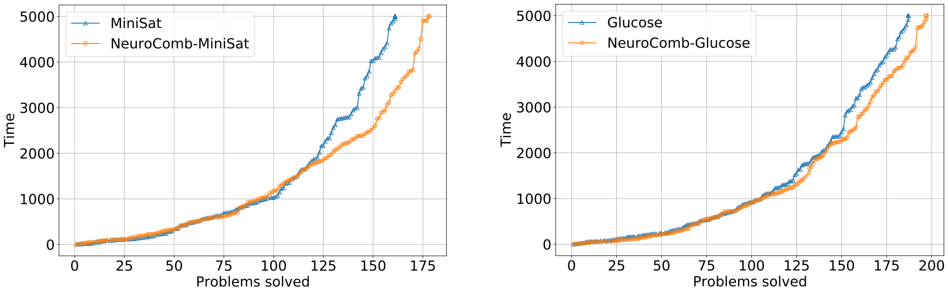

Figure 4: Progress of `NeuroComb-MiniSat` vs. `MiniSat` (left), and `NeuroComb-Glucose` vs. `Glucose` (right) over time in solving the 400 test problems (time in seconds). Both `NeuroComb-MiniSat` and `NeuroComb-Glucose` take the lead at around 2,000 seconds and the differences become more pronounced as the instances take longer to solve.

93.7% of the clauses correctly. Both tasks are predicted with over 90% precision and over 0.85 F_1 score. Thus, *the GNN model learns to classify both important variables and important clauses well.*

RQ2: NeuroComb Performance To evaluate the solving effectiveness, NeuroComb solvers `NeuroComb-MiniSat` and `NeuroComb-Glucose`, and their corresponding baseline solvers `MiniSat` and `Glucose`, were applied to all 400 SAT problems in the testing dataset, with the standard time out of 5,000 seconds. Each solver utilized up to 64 different processes in parallel on the dedicated 64-core machine. The solving time of NeuroComb solvers includes both the model inference time and the SAT solving time (the static information extraction time is negligible). The inference was done on the GPU for those 297 problems that fit into GPU memory and on the CPU for the rest 103 problems. The total time of the GPU inference ranged from 0.2 to 27.3 seconds, with an average of 4.6 seconds, and the CPU inference from 61.5 to 2,456.6 seconds, with an average of 367.2 seconds.

As a result, `NeuroComb-MiniSat` solved 178 problems, whereas `MiniSat` solved 161, which amounts to an improvement of 11%. Similarly, `NeuroComb-Glucose` solved 197 problems, whereas `Glucose` solved 187, which amounts to an improvement of 5%. The scatter plot is commonly used in SAT community for comparing the solving effectiveness of two solvers on each problem, as shown in Fig. 3. There are many problems that can be solved by both NeuroComb solvers under 4,000 seconds on which baseline solvers fail; the opposite rarely happens. In addition, more dots appear in the lower triangular area in both plots, indicating that there are more problems on which NeuroComb solvers outperform the baseline solvers. The cactus plot is another commonly used plot in SAT community for demonstrating the solving progress over time, as shown in Fig. 4. Both NeuroComb solvers improve upon the baseline solvers at around 2,000 seconds and the differences become more pronounced as the instances take longer to solve. This phenomenon can be explained by the way the neural predictions are applied using multi-armed bandits: The approach does more exploration at the early stages and more exploitation at the later stages, with useful knowledge learned along the way. Overall, the results show that *NeuroComb approach is able to make the baseline solvers more effective* on the SATCOMP-2021 problems.

7 Discussion and Future Work

The current implementation of NeuroComb has two main limitations. First, training the GNN model requires a large amount of data. Once trained, however, the model is able to generalize to new problem categories and to much larger problems, so the investment in putting together a large dataset is warranted. The DataComb dataset will also be publicly available to benefit future research. Second, only limited hyperparameter tuning of NeuroComb solvers has been done so far, and it is possible that better settings exist. A compelling direction for future work is to use search methods such as reinforcement learning or evolutionary optimization both to expand the dataset and to optimize the hyperparameters.

8 Conclusion

This paper proposes a machine learning approach, NeuroComb, to make CDCL SAT solving both more effective and more practical. The main idea is to train a GNN model to identify both important variables and important clauses, and use the neural predictions as static information to inform the VSIDS branching heuristic. The model inference is performed only once before the SAT solving starts and the resulting offline predictions are applied periodically using multi-arm bandits during the SAT solving. Incorporated in both the classic `MiniSat` solver and the more recent `Glucose` solver, this approach makes it possible to solve more instances. NeuroComb is thus a promising and practical approach to improving the SAT solving effectiveness through machine learning.

References

[1] Sat competition 2021. `https://satcompetition.github.io/2021/index.html`, 2021. Accessed: 2022-05-15.

[2] Gilles Audemard, Jean-Marie Lagniez, and Laurent Simon. Improving glucose for incremental sat solving with assumptions: Application to mus extraction. In *International conference on theory and applications of satisfiability testing*, pages 309–317. Springer, 2013.

[3] Gilles Audemard and Laurent Simon. On the glucose sat solver. *International Journal on Artificial Intelligence Tools*, 27(01):1840001, 2018.

[4] Peter W Battaglia, Jessica B Hamrick, Victor Bapst, Alvaro Sanchez-Gonzalez, Vinicius Zambaldi, Mateusz Malinowski, Andrea Tacchetti, David Raposo, Adam Santoro, Ryan Faulkner, et al. Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261*, 2018.

[5] Martin Davis, George Logemann, and Donald Loveland. A machine program for theorem-proving. *Communications of the ACM*, 5(7):394–397, 1962.

[6] Niklas Eén and Niklas Sörensson. An extensible sat-solver. In *International conference on theory and applications of satisfiability testing*, pages 502–518. Springer, 2003.

[7] Matthias Fey and Jan Eric Lenssen. Fast graph representation learning with pytorch geometric. *arXiv preprint arXiv:1903.02428*, 2019.

[8] Armin Biere Katalin Fazekas Mathias Fleury and Maximilian Heisinger. Cadical, kissat, paracooba, plingeling and treengeling entering the SAT competition 2020. *SAT COMPETITION 2020*, page 50, 2020.

[9] Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. In *International conference on machine learning*, pages 1263–1272. PMLR, 2017.

[10] Xavier Glorot, Antoine Bordes, and Yoshua Bengio. Deep sparse rectifier neural networks. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 315–323. JMLR Workshop and Conference Proceedings, 2011.

[11] Marco Gori, Gabriele Monfardini, and Franco Scarselli. A new model for learning in graph domains. In *Proceedings. 2005 IEEE International Joint Conference on Neural Networks, 2005.*, volume 2, pages 729–734. IEEE, 2005.

[12] William L Hamilton, Rex Ying, and Jure Leskovec. Inductive representation learning on large graphs. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 1025–1035, 2017.

[13] Jesse Michael Han. Enhancing sat solvers with glue variable predictions. *arXiv preprint arXiv:2007.02559*, 2020.

[14] Marijn JH Heule, Matti Juhani Järvisalo, Martin Suda, et al. Proceedings of SAT competition 2018: Solver and benchmark descriptions, 2018.

[15] Holger H Hoos and Thomas Stützle. Satlib: An online resource for research on sat. *Sat*, 2000:283–292, 2000.

[16] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017.

[17] Daniel Jackson. Alloy: a lightweight object modelling notation. *ACM Transactions on Software Engineering and Methodology (TOSEM)*, 11(2):256–290, 2002.

[18] Mikoláš Janota, Inês Lynce, and Joao Marques-Silva. Algorithms for computing backbones of propositional formulae. *Ai Communications*, 28(2):161–177, 2015.

[19] Sebastian Jaszczur, Michał Łuszczyk, and Henryk Michalewski. Neural heuristics for sat solving. *arXiv preprint arXiv:2005.13406*, 2020.

[20] Philip Kilby, John Slaney, Sylvie Thiébaux, Toby Walsh, et al. Backbones and backdoors in satisfiability. In *AAAI*, volume 5, pages 1368–1373, 2005.

[21] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.

[22] Vitaly Kurin, Saad Godil, Shimon Whiteson, and Bryan Catanzaro. Improving SAT solver heuristics with graph networks and reinforcement learning, 2019.

[23] Vitaly Kurin, Saad Godil, Shimon Whiteson, and Bryan Catanzaro. Improving SAT solver heuristics with graph networks and reinforcement learning. In *Advances in Neural Information Processing Systems*, 2020.

[24] Massimo Lauria, Jan Elffers, Jakob Nordström, and Marc Vinyals. Cnfgen: A generator of crafted benchmarks. In *International Conference on Theory and Applications of Satisfiability Testing*, pages 464–473. Springer, 2017.

[25] Guohao Li, Matthias Muller, Ali Thabet, and Bernard Ghanem. Deepgcns: Can gcns go as deep as cnns? In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9267–9276, 2019.

[26] Qimai Li, Zhichao Han, and Xiao-Ming Wu. Deeper insights into graph convolutional networks for semi-supervised learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.

[27] Andy Liaw, Matthew Wiener, et al. Classification and regression by randomforest. *R news*, 2(3):18–22, 2002.

[28] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.

[29] Andreas Loukas. What graph neural networks cannot learn: depth vs width. *arXiv preprint arXiv:1907.03199*, 2019.

[30] Joao Marques-Silva, Inês Lynce, and Sharad Malik. Conflict-driven clause learning sat solvers. In *Handbook of Satisfiability: Second Edition. Part 1/Part 2*, pages 133–182. IOS Press BV, 2021.

[31] Joao P Marques-Silva and Karem A Sakallah. Grasp—a new search algorithm for satisfiability. In *The Best of ICCAD*, pages 73–89. Springer, 2003.

[32] Diego Mesquita, Amauri Souza, and Samuel Kaski. Rethinking pooling in graph neural networks. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 2220–2231. Curran Associates, Inc., 2020.

[33] Matthew W. Moskewicz, Conor F. Madigan, Ying Zhao, Lintao Zhang, and Sharad Malik. Chaff: Engineering an efficient sat solver. In *Proceedings of the 38th Annual Design Automation Conference*, DAC '01, page 530–535, New York, NY, USA, 2001. Association for Computing Machinery.

[34] Andrew J Parkes. Clustering at the phase transition. In *AAAI/IAAI*, pages 340–345. Citeseer, 1997.

[35] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *arXiv preprint arXiv:1912.01703*, 2019.

[36] Alessandro Previti and Matti Järvisalo. A preference-based approach to backbone computation with application to argumentation. In *Proceedings of the 33rd Annual ACM Symposium on Applied Computing*, pages 896–902, 2018.

[37] Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. The graph neural network model. *IEEE Transactions on Neural Networks*, 20(1):61–80, 2008.

[38] Daniel Selsam and Nikolaj Bjørner. Guiding high-performance SAT solvers with unsat-core predictions. In *International Conference on Theory and Applications of Satisfiability Testing*, pages 336–353. Springer, 2019.

[39] Daniel Selsam, Matthew Lamm, Benedikt Bünz, Percy Liang, Leonardo de Moura, and David L Dill. Learning a SAT solver from single-bit supervision. *arXiv preprint arXiv:1802.03685*, 2018.

[40] William R Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3-4):285–294, 1933.

[41] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Instance normalization: The missing ingredient for fast stylization. *arXiv preprint arXiv:1607.08022*, 2016.

[42] Stephan Wojtowytsch and Weinan E. Can shallow neural networks beat the curse of dimensionality? a mean field training perspective. *IEEE Transactions on Artificial Intelligence*, 1(2):121–129, 2020.

[43] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and S Yu Philip. A comprehensive survey on graph neural networks. *IEEE transactions on neural networks and learning systems*, 32(1):4–24, 2020.

[44] Emre Yolcu and Barnabás Póczos. Learning local search heuristics for boolean satisfiability. In *Advances in Neural Information Processing Systems*, pages 7992–8003, 2019.

[45] Ziwei Zhang and Yang Zhang. Elimination mechanism of glue variables for solving sat problems in linguistics. In *The Asian Conference on Language*, pages 147–167, 2021.

[46] Jie Zhou, Ganqu Cui, Shengding Hu, Zhengyan Zhang, Cheng Yang, Zhiyuan Liu, Lifeng Wang, Changcheng Li, and Maosong Sun. Graph neural networks: A review of methods and applications. *AI Open*, 1:57–81, 2020.