SAT-ACT: LEARNING CDCL BRANCHING FROM OUTCOME-DERIVED ACTION PREFERENCES

Anonymous authors
Paper under double-blind review

ABSTRACT

Learning-based approaches have increasingly been explored for SAT solving, with particular attention to improving branching decisions in conflict-driven clause learning (CDCL) solvers. These approaches derive branching guidance either from static structural properties of SAT formulas or from dynamic signals obtained through solver feedback and execution traces. However, these forms of supervision provide limited evidence about state-dependent branching effects, while different outcome measures may favor different branching choices. To address these limitations, we propose SAT-ACT, an outcome-derived learning framework for CDCL branching. We design three complementary outcome comparison criteria to extract action preferences from independent branching action rollouts initiated at the same solver state. The criteria identify preferences supported by the outcomes and leave comparisons unresolved when the evidence is insufficient or conflicting, instead of forcing an ordering using a predefined scalar metric. A state-conditioned network learns branching priorities from the extracted preferences and uses them to guide CDCL search directly without online action rollouts. We further introduce non-dominated heuristic action supervision, which retains the solver-selected action as a training target only when it remains non-dominated under these preferences. Solver-level experiments show that SAT-ACT consistently reduces CDCL search costs on three SAT families and improves PAR-2 on the SAT Competition benchmarks.

1 INTRODUCTION

The Boolean satisfiability problem (SAT) is a fundamental problem in computer science with broad applications in planning (Büttner & Rintanen, 2005), verification (Vizel et al., 2015), and security (Mironov & Zhang, 2006). As the first problem shown to be NP-complete (Cook, 1971), SAT has motivated decades of research on efficient solving algorithms. Modern SAT solvers are largely built upon conflict-driven clause learning (CDCL), which combines branching, Boolean constraint propagation, conflict analysis, clause learning, and restart mechanisms to solve large and complex instances (Biere et al., 2021). Among these components, branching is particularly consequential because each branching action selects an unassigned variable together with its polarity, and this choice can substantially alter the subsequent search trajectory. Modern CDCL solvers therefore rely on carefully designed branching heuristics to select these actions.

Recently, many studies have explored learning-based approaches to solve SAT instances or enhance SAT solving (Bünz & Lamm, 2017; Selsam et al., 2019; Amizadeh et al., 2019; Li et al., 2023; Cameron et al., 2020; Selsam & Bjørner, 2019; Shi et al., 2023; Wang et al., 2024; Sun et al., 2026). A representative line of work learns structural signals from SAT formulas to guide CDCL heuristics. Such methods predict information related to unsatisfiable cores or backbone variables to assist variable selection and phase initialization (Selsam & Bjørner, 2019; Shi et al., 2023; Wang et al., 2024; Sun et al., 2026). More recent approaches instead derive supervision from the solving process itself. Reinforcement-learning (RL) methods learn branching guidance through solver interaction or solver-level performance feedback (Kurin et al., 2020; Tönshoff & Grohe, 2026), while imitation-based methods learn branching policies from decisions recorded in CDCL execution traces (Zhang et al., 2026). Together, these developments move learning-based branching from static formula-level guidance toward supervision that more directly reflects dynamic CDCL search.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

Although existing learning-based approaches have improved CDCL branching, they construct supervision at different levels of the solving process. Methods that learn structural signals from SAT formulas guide variable selection and phase initialization, but these signals do not describe the downstream effect of a branching decision at a dynamic solver state (Selsam & Bjørner, 2019; Shi et al., 2023; Wang et al., 2024; Sun et al., 2026). RL-based methods learn from solver feedback (Kurin et al., 2020; Tönshoff & Grohe, 2026), while imitation-based methods learn from decisions recorded during CDCL execution (Zhang et al., 2026). These forms of supervision reflect solver behavior, but they do not necessarily provide matched outcomes for multiple branching decisions at the same dynamic state. The effect of a branching decision depends on the current assignments, learned clauses, and search history, so evidence collected from different states or trajectories cannot provide matched comparisons at a particular decision point. Classical SAT lookahead solvers probe candidate branches through propagation-based effects (Heule & van Maaren, 2021; Biere et al., 2021), while strong branching in mixed-integer linear programming (MILP) evaluates candidate variables under a shared branch-and-bound context (Khalil et al., 2016; Gasse et al., 2019). These precedents demonstrate the value of assessing candidate decisions under a shared search context. Existing learning-based CDCL branching methods, however, still provide limited supervision for distinguishing the downstream effects of the branching decisions available at a dynamic solver state.

A further challenge lies in deriving reliable supervision from solver outcomes at a dynamic solver state. Solver outcomes report whether solving terminates and how many conflicts, propagations, decisions, and restarts occur during the evaluation. These measures describe different aspects of the subsequent search and may favor different branching choices. Reducing them to a single scalar criterion would impose a fixed tradeoff among different outcome measures. The key question is therefore when the observed outcomes provide sufficient evidence to prefer one branching choice over another. A suitable supervision scheme should retain only preferences supported by the observed outcomes and leave comparisons with insufficient or conflicting evidence unordered rather than forcing a complete ordering over the available choices.

To address these limitations, we propose SAT-ACT, a framework for learning CDCL branching from outcome-derived action preferences. The framework samples candidate branching actions at selected solver states and runs an independent CDCL rollout for each action from its corresponding state. The rollout outcomes are examined using three complementary outcome comparison criteria to extract supported action preferences. To learn branching decisions from these preferences, we develop a CDCL branching action network that encodes the active CNF and dynamic assignment features and produces state-conditioned scores for branching actions. The extracted preferences supervise the network through a pairwise ranking objective. During solving, the learned scores guide branching directly without requiring action rollouts. To incorporate guidance from the solver's internal branching heuristic, we introduce non-dominated heuristic action supervision. Under this mechanism, the heuristic action is used as a training target only when it remains non-dominated under the extracted preferences.

Our main contributions are summarized below.

- We introduce an outcome-derived learning framework for CDCL branching. The framework uses a state-conditioned model to score branching actions, while outcome-derived action preferences supervise the model through a pairwise ranking objective. During solving, the trained model guides branching directly without online action rollouts.
- We design three complementary outcome comparison criteria for extracting supported preferences between branching actions. These criteria leave action pairs without sufficient evidence unordered and avoid collapsing the observed outcomes into a predefined metric.
- We additionally propose non-dominated heuristic action supervision for the action selected by the solver's internal branching heuristic. The heuristic action is retained as a training target only when it remains non-dominated under the outcome-derived preferences.

2 PRELIMINARIES

Boolean Satisfiability Problem (SAT). We consider a Boolean formula ϕ in conjunctive normal form (CNF) over N Boolean variables and M clauses (Biere et al., 2021). Each variable v_i takes either *true* or *false*, represented by 1 and 0, respectively. A CNF formula is a conjunction of clauses,

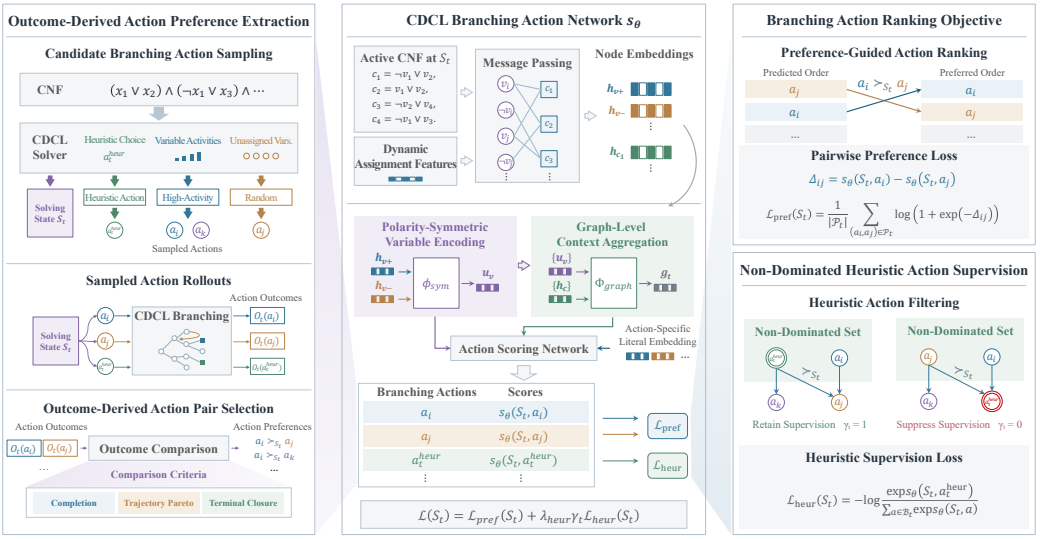


Figure 1: Overview of SAT-ACT, an outcome-derived learning framework for CDCL branching. Candidate branching actions initiate independent CDCL rollouts from the same solver state. The rollout outcomes are compared to extract action preferences, which supervise a state-conditioned branching action network through pairwise ranking.

and each clause is a disjunction of literals. A literal ℓ is either a variable v_i or its negation $\neg v_i$. These two literals are referred to as the positive and negative literals of v_i and have opposite polarities. We use $\mathcal{L} = \{\ell_0, \ldots, \ell_{2N-1}\}$ to denote the set of all literals and $\mathcal{C} = \{c_0, \ldots, c_{M-1}\}$ to denote the set of clauses in ϕ. The SAT problem asks whether there exists an assignment under which ϕ evaluates to *true*. A formula admitting such an assignment is *satisfiable*. Otherwise, it is *unsatisfiable*.

Conflict-Driven Clause Learning (CDCL). CDCL forms the basis of modern SAT solvers (Biere et al., 2021). During solving, CDCL repeatedly performs Boolean constraint propagation and branching decisions, while conflicts trigger conflict analysis, clause learning, and non-chronological backtracking. Let S_t denote the dynamic solver state immediately before the t-th branching decision. The state contains the current partial assignment, dynamic search information, and the active clause set $\mathcal{C}_t$. The set $\mathcal{C}_t$ contains the original and learned clauses retained by the solver at S_t.

CDCL Branching Actions. At each branching decision, CDCL determines an unassigned variable and its polarity. These choices are commonly made by separate variable-selection and phase-selection heuristics. For example, Exponential Variable State-Independent Decaying Sum (EVSIDS) maintains activity scores that emphasize recent conflict participation and selects the eligible unassigned variable with the highest activity (Biere et al., 2021). A phase-selection heuristic then determines its polarity. We represent these choices jointly as a decision literal, referred to as the *branching action*. Applying the action assigns the selected literal to true. Let $\mathcal{B}_t \subseteq \mathcal{L}$ denote the set of legal branching actions at S_t. For each eligible unassigned variable, $\mathcal{B}_t$ contains its positive and negative literals. We denote the action selected by the solver's internal branching heuristic as $a_t^{\text{heur}} \in \mathcal{B}_t$. Different branching actions from the same solver state can produce different propagation effects, search trajectories, and solving outcomes.

3 METHOD

In this section, we present SAT-ACT, an outcome-derived learning framework for CDCL branching. At selected solver states, the framework samples candidate branching actions and runs an independent CDCL rollout for each action from its corresponding state. We compare the rollout outcomes using three complementary criteria to extract supported action preferences, while uncertain comparisons remain unordered. A CDCL branching action network encodes the active CNF and dynamic assignment features to produce state-conditioned scores for branching actions. The extracted preferences supervise the network through a pairwise ranking objective. We also introduce

non-dominated heuristic action supervision for the action selected by the solver's internal branching heuristic. The heuristic action is retained as a training target only when it remains non-dominated under the extracted preferences. Figure 1 presents an overview of the framework.

3.1 Outcome-Derived Action Preference Extraction

CDCL branching is state-dependent, while static supervision or a single observed trajectory provides limited evidence about other legal branching actions available at the same point in the search. We therefore run independent CDCL rollouts for multiple candidate branching actions from the same dynamic solver state and record their solving outcomes. Because all rollouts begin from the same state, their outcomes can be compared under a shared search context. We then apply three complementary outcome comparison criteria to extract supported action preferences while leaving uncertain comparisons unordered.

We first sample solver states along CDCL search trajectories. For each sampled state S_t, we construct a bounded candidate set $\mathcal{A}_t \subseteq \mathcal{B}_t$. The heuristic action a_t^{heur} is retained in $\mathcal{A}_t$. Each action $a \in \mathcal{A}_t$ initiates an independent CDCL rollout from S_t. The action is applied as the next branching decision, after which the solver resumes its internal CDCL procedure under a common configuration and rollout budget. Each rollout produces an outcome

$$O_t(a) = \big(y_t(a), C_t(a), P_t(a), D_t(a), R_t(a)\big), \tag{1}$$

where $y_t(a) \in \{\text{SAT}, \text{UNSAT}, \text{incomplete}\}$ denotes the rollout status, and $C_t(a)$, $P_t(a)$, $D_t(a)$, and $R_t(a)$ record the numbers of conflicts, propagations, decisions, and restarts. The state sampling, candidate construction, and rollout protocol are detailed in Appendix A. These rollouts are used only for training-data collection and are not required during inference.

For two actions $a_i, a_j \in \mathcal{A}_t$, we write $a_i \succ_{S_t} a_j$ when their rollout outcomes provide sufficient evidence for preferring a_i to a_j. The resulting preference set is

$$\mathcal{P}_t = \{(a_i, a_j) \mid a_i \succ_{S_t} a_j\}. \tag{2}$$

Because the observed outcomes may not distinguish every action pair, $\mathcal{P}_t$ defines a partial preference relation. Rather than combining the outcome measures into a predefined scalar score, we use three complementary comparison criteria to determine which preferences are supported.

Criterion 1. Completion (C1). Let $\tau_t(a) = \mathbb{I}[y_t(a) \in \{\text{SAT}, \text{UNSAT}\}]$ indicate whether the rollout of action a reaches a terminal result within the rollout budget. C1 supports $a_i \succ_{S_t} a_j$ when $\tau_t(a_i) > \tau_t(a_j)$. Thus, an action whose rollout reaches SAT or UNSAT is preferred to one whose rollout remains incomplete. If both rollouts have the same completion status, C1 provides no preference evidence.

Criterion 2. Trajectory Pareto (C2). For each completed rollout, we define the search-cost vector $\mathbf{c}_t(a) = (C_t(a), P_t(a), D_t(a), R_t(a))$. C2 supports $a_i \succ_{S_t} a_j$ when $\mathbf{c}_t(a_i) \preceq \mathbf{c}_t(a_j)$ componentwise and $\mathbf{c}_t(a_i) \neq \mathbf{c}_t(a_j)$. This criterion prefers an action only when its rollout requires no more conflicts, propagations, decisions, or restarts and improves at least one of these quantities. If the two rollouts exhibit trade-offs across the search-cost dimensions, C2 provides no preference evidence.

Criterion 3. Terminal Closure (C3). For each completed rollout, we define a terminal-closure level $\kappa_t(a)$ that measures how directly the rollout reaches its final SAT or UNSAT result. A larger value indicates that less additional search is required after applying the branching action. The exact terminal-closure levels are defined in Appendix A.4. C3 supports $a_i \succ_{S_t} a_j$ when $\kappa_t(a_i) > \kappa_t(a_j)$. Equal terminal-closure levels provide no preference evidence.

For each action pair, we combine the evidence provided by the applicable comparison criteria. The pair (a_i, a_j) is added to $\mathcal{P}_t$ only when at least one criterion supports $a_i \succ_{S_t} a_j$ and no criterion supports the opposite direction. If no criterion distinguishes the actions or the available evidence is conflicting, the pair remains unordered.

3.2 Learning CDCL Branching from Action Preferences

After extracting outcome-derived action preferences, we use them to train a state-conditioned CDCL branching action network. At solver state S_t, the network encodes the active CNF and dynamic

assignment features to score each legal branching action $a \in \mathcal{B}_t$ with $s_\theta(S_t, a)$, which determines its relative ranking.

Literal–Clause Graph-Based CDCL State Representation. We represent the active CNF at S_t as a literal–clause bipartite graph $\mathcal{G}_t = (\mathcal{L}_t, \mathcal{C}_t, \mathcal{E}_t)$. Let $\mathcal{V}_t$ denote the variables represented in the active CNF, and let $\mathcal{L}_t = \{v^+, v^- \mid v \in \mathcal{V}_t\}$ contain their positive and negative literals. The set $\mathcal{C}_t$ contains the active clauses, and an edge $(\ell, c) \in \mathcal{E}_t$ indicates that literal ℓ occurs in clause c.

Literal-node features encode the current assignment status and assignment level together with occurrence and global search information. Clause-node features encode the clause size, the current assignment composition of its literals, and its satisfaction and near-unit status. Detailed definitions of these features are provided in Appendix B.

The literal and clause features are projected into a shared embedding space to obtain initial representations $\mathbf{h}_\ell^{(0)}$ and $\mathbf{h}_c^{(0)}$. We then apply a NeuroCore-style message-passing encoder (Selsam & Bjørner, 2019). Literal and clause nodes exchange information along their incidence edges, while each literal also incorporates information from its opposite-polarity counterpart. After several rounds, the encoder produces contextual literal and clause embeddings $\mathbf{h}_\ell$ and $\mathbf{h}_c$.

Polarity-Aware Branching Action Scoring. For each variable v, let $\mathbf{h}_{v^+}$ and $\mathbf{h}_{v^-}$ be its positive and negative literal embeddings. Polarity-symmetric variable encoding combines them as

$$\mathbf{u}_v = \phi_{\text{sym}}\left([\mathbf{h}_{v^+} + \mathbf{h}_{v^-}\,;\, |\mathbf{h}_{v^+} - \mathbf{h}_{v^-}|]\right), \tag{3}$$

where ϕ_{sym} is a shared learnable transformation. The resulting representation is invariant to exchanging the two polarities, while $\mathbf{h}_{v^+}$ and $\mathbf{h}_{v^-}$ retain polarity-specific information.

Graph-level context aggregation summarizes the variable and clause representations as

$$\mathbf{g}_t = \Phi_{\text{graph}}\left(\{\mathbf{u}_v\}_{v\in\mathcal{V}_t}, \{\mathbf{h}_c\}_{c\in\mathcal{C}_t}\right). \tag{4}$$

The graph-level representation $\mathbf{g}_t$ captures the global context of the current solver state.

Graph-conditioned variable encoding then incorporates this context into each variable representation

$$\widetilde{\mathbf{u}}_v = \Phi_{\text{var}}\left(\mathbf{u}_v, \mathbf{g}_t\right), \tag{5}$$

where Φ_{graph} and Φ_{var} are learnable aggregation and conditioning functions. Their implementations are detailed in Appendix C.

For a branching action a associated with variable $v(a)$, the network computes

$$s_\theta(S_t, a) = \phi_{\text{score}}\left(\left[\mathbf{h}_a\,;\, \widetilde{\mathbf{u}}_{v(a)}\,;\, \mathbf{g}_t\right]\right). \tag{6}$$

The action-specific literal embedding $\mathbf{h}_a$ distinguishes the two polarities of the same variable, while $\widetilde{\mathbf{u}}_{v(a)}$ and $\mathbf{g}_t$ provide variable-level and graph-level context.

Preference-Guided Action Ranking. The preference set $\mathcal{P}_t$ supervises the relative ordering of the branching-action scores. For each $(a_i, a_j) \in \mathcal{P}_t$ with $a_i \succ_{S_t} a_j$, we define the score difference

$$\Delta_{ij} = s_\theta(S_t, a_i) - s_\theta(S_t, a_j). \tag{7}$$

We impose the extracted preference relations through a logistic pairwise ranking objective

$$\mathcal{L}_{\text{pref}}(S_t) = \frac{1}{|\mathcal{P}_t|} \sum_{(a_i,a_j)\in\mathcal{P}_t} \log\left(1 + \exp(-\Delta_{ij})\right). \tag{8}$$

This objective encourages the model to assign higher scores to preferred actions without requiring absolute action-quality targets. The loss is averaged within each state so that states with different numbers of supported preference pairs contribute on a comparable scale. We set $\mathcal{L}_{\text{pref}}(S_t) = 0$ when $\mathcal{P}_t$ is empty.

Non-Dominated Heuristic Action Supervision. The action selected by the solver's internal branching heuristic can provide a training target, but it may conflict with the outcome-derived preferences when an evaluated action is preferred to it. We therefore retain this supervision only when the heuristic action remains non-dominated among the evaluated candidates.

Let $\mathcal{N}_t$ denote the non-dominated subset of $\mathcal{A}_t$

$$\mathcal{N}_t = \{a \in \mathcal{A}_t \mid \nexists\, a' \in \mathcal{A}_t \text{ such that } a' \succ_{S_t} a\}\,. \tag{9}$$

Since $\succ_{S_t}$ is a partial preference relation and $\mathcal{A}_t \subseteq \mathcal{B}_t$, membership in $\mathcal{N}_t$ does not imply global optimality. It only indicates that no evaluated action is preferred to that action.

Let a_t^{heur} denote the heuristic action. We define the supervision gate as

$$\gamma_t = \mathbb{I}\left[a_t^{\text{heur}} \in \mathcal{N}_t\right]. \tag{10}$$

When the heuristic action remains non-dominated under the outcome-derived preferences, we use it as a training target through a cross-entropy objective over the legal branching actions

$$\mathcal{L}_{\text{heur}}(S_t) = -\log \frac{\exp s_\theta(S_t, a_t^{\text{heur}})}{\sum_{a \in \mathcal{B}_t} \exp s_\theta(S_t, a)}. \tag{11}$$

Overall Training Objective. The state-level training objective combines preference-guided action ranking with non-dominated heuristic action supervision

$$\mathcal{L}(S_t) = \mathcal{L}_{\text{pref}}(S_t) + \lambda_{\text{heur}} \gamma_t \mathcal{L}_{\text{heur}}(S_t), \tag{12}$$

where λ_{heur} controls the contribution of heuristic action supervision.

Together, the two terms train the network to respect outcome-derived action preferences while retaining compatible guidance from the solver's internal branching heuristic.

4 Experiments

In this section, we present a solver-level experimental analysis of CDCL branching to evaluate the effectiveness of SAT-ACT.

Datasets. We use three SAT families from G4SATBench (Li et al., 2024). These families are SR, Community Attachment (CA), and Popularity-Similarity (PS). For each family, 16,000/2,000/2,000 instances are used for model training, validation, and testing, respectively. Outcome-derived branching supervision is collected with an instrumented CaDiCaL solver (Biere et al., 2024) through independent CDCL rollouts. Dataset and supervision statistics are provided in Appendix D, with solver sampling and evaluation settings in Appendix F.

Baselines. We compare SAT-ACT with native CaDiCaL (3.0.0), NeuroCore (Selsam & Bjørner, 2019), GQSAT (Kurin et al., 2020), SATformer (Shi et al., 2023), ImitSAT (Zhang et al., 2026), and RLAF (Tönshoff & Grohe, 2026). For a consistent solver-level comparison, all learning-based methods are integrated with the same CaDiCaL solver while retaining their original guidance mechanisms. NeuroCore, GQSAT, and ImitSAT guide branching decisions during search, whereas SATformer and RLAF provide one-shot guidance before search.

Metrics. We measure solver performance using the relative changes in conflicts (Δ_C), decisions (Δ_D), and propagations (Δ_P) with respect to native CaDiCaL. Following ImitSAT (Zhang et al., 2026), we also report median relative propagation per problem (MRPP) and the one-percent win rate (W(1%)). Detailed metric definitions are provided in Appendix E.

Implementation Details. Our models use 10 message-passing layers with a hidden dimension of 128. They are trained with AdamW (Loshchilov & Hutter, 2019) for 50 epochs using a learning rate of 10^{-4}. We select λ_{heur} from $[0, 0.2]$, where $\lambda_{\text{heur}} = 0$ removes non-dominated heuristic action supervision and corresponds to SAT-ACT (Pref. Only). For the experiments on SR, CA, and PS, the 3-call and 5-call settings allow at most three and five model requests per solver run, respectively, and each run has a 100-second time limit. Complete hyperparameter configurations and solver evaluation settings are provided in Appendix F.

Table 1: Solver-level comparison on SR, CA, and PS. Lower is better for Δ_C, Δ_D, Δ_P, and MRPP, while higher is better for W(1%).

Dataset	Methods	$\Delta_C \downarrow$	$\Delta_D \downarrow$	$\Delta_P \downarrow$	MRPP $\downarrow$	W(1%) $\uparrow$
SR	CaDiCaL (2024)	0.00%	0.00%	0.00%	1.000	0.00%
	NeuroCore (3-call) (2019)	-32.61%	-29.99%	-33.38%	0.663	79.90%
	NeuroCore (5-call) (2019)	-38.44%	-35.29%	-39.34%	0.602	84.30%
	GQSAT (3-call) (2020)	-20.74%	-18.40%	-20.64%	0.800	68.90%
	GQSAT (5-call) (2020)	-24.36%	-21.80%	-24.03%	0.774	70.80%
	SATformer (2023)	-12.22%	-11.52%	-13.02%	1.000	24.40%
	ImitSAT (3-call) (2026)	-0.30%	-0.27%	-0.37%	0.998	46.80%
	ImitSAT (5-call) (2026)	-0.35%	-0.68%	-0.35%	0.995	48.50%
	RLAF (2026)	-26.99%	-26.31%	-27.13%	0.803	73.20%
	SAT-ACT (3-call)	-38.01%	-35.17%	-37.49%	0.582	85.70%
	SAT-ACT (5-call)	**-42.46%**	**-39.63%**	**-41.60%**	**0.534**	**88.30%**
CA	CaDiCaL (2024)	0.00%	0.00%	0.00%	1.000	0.00%
	NeuroCore (3-call) (2019)	-1.44%	2.51%	0.43%	1.029	30.40%
	NeuroCore (5-call) (2019)	-6.12%	0.32%	-5.29%	1.029	34.90%
	GQSAT (3-call) (2020)	-5.94%	-2.08%	-4.93%	1.006	41.50%
	GQSAT (5-call) (2020)	-6.68%	-0.72%	-4.45%	1.005	44.50%
	SATformer (2023)	-10.64%	-4.30%	-9.16%	1.000	35.60%
	ImitSAT (3-call) (2026)	1.04%	2.25%	1.47%	1.011	22.30%
	ImitSAT (5-call) (2026)	1.50%	3.33%	2.83%	1.015	24.30%
	RLAF (2026)	**-23.42%**	-14.04%	**-21.39%**	0.909	63.50%
	SAT-ACT (3-call)	-11.05%	-6.54%	-9.42%	0.961	54.50%
	SAT-ACT (5-call)	-20.85%	**-14.48%**	-18.63%	**0.817**	**70.20%**
PS	CaDiCaL (2024)	0.00%	0.00%	0.00%	1.000	0.00%
	NeuroCore (3-call) (2019)	-17.42%	-17.05%	-16.32%	0.636	76.80%
	NeuroCore (5-call) (2019)	-22.33%	-21.72%	-20.78%	0.585	77.70%
	GQSAT (3-call) (2020)	-2.18%	-2.58%	-2.86%	0.957	52.10%
	GQSAT (5-call) (2020)	-3.90%	-3.91%	-3.42%	0.962	51.30%
	SATformer (2023)	-15.41%	-14.94%	-13.02%	0.999	46.80%
	ImitSAT (3-call) (2026)	-1.26%	-1.34%	-1.72%	1.008	45.30%
	ImitSAT (5-call) (2026)	-1.75%	-1.88%	-2.21%	1.012	45.30%
	RLAF (2026)	-20.12%	-19.93%	-18.55%	0.813	67.10%
	SAT-ACT (3-call)	-33.23%	-32.17%	-31.34%	0.484	86.00%
	SAT-ACT (5-call)	**-43.66%**	**-42.55%**	**-42.40%**	**0.417**	**88.90%**

Exp-1. Comparison with Existing Methods for CDCL Branching. Table 1 compares SAT-ACT with native CaDiCaL and learning-based SAT branching baselines. On SR and PS, SAT-ACT (5-call) improves the three aggregate search-cost metrics over the best competing baseline by 2.26–4.34 and 20.83–21.62 percentage points, respectively. On CA, RLAF obtains lower aggregate conflicts and propagations, whereas SAT-ACT (5-call) achieves better decisions, MRPP, and W(1%). For SAT-ACT, increasing the intervention budget generally improves solver performance.

The performance differences are consistent with the supervision used by each method. Structural and trajectory-level supervision can improve branching, but provides less direct evidence for comparing actions at a dynamic solver state. ImitSAT encodes the CNF and execution trace as sequences, and the difficulty of modeling long sequences may help explain its weaker performance. RLAF benefits from complete-run feedback and persistent global guidance, which appears particularly effective on CA, although its weaker MRPP and W(1%) results suggest that these gains are less broadly distributed across instances. In contrast, SAT-ACT learns from state-matched action outcomes and applies its scores at dynamic decision points, yielding strong performance across the three datasets. Its smaller gains on CA are consistent with the lower coverage of preference pairs reported in Table 5. Additional analyses of SAT/UNSAT instances, neural intervention budgets, and neural inference cost are provided in Appendices G.1, G.4, and G.6, respectively.

Exp-2. Ablation of Outcome Comparison Criteria. We analyze the contribution of the three outcome comparison criteria by progressively incorporating different forms of solving evidence. We compare C1, C1+C2, and C1+C2+C3 under the 5-call setting while keeping the model, training procedure, and solver evaluation protocol unchanged. All variants use preference-guided action ranking

Table 2: Incremental analysis of outcome comparison criteria. Lower is better for all metrics.

Criteria	SR			CA			PS		
	$\Delta_C \downarrow$	$\Delta_D \downarrow$	$\Delta_P \downarrow$	$\Delta_C \downarrow$	$\Delta_D \downarrow$	$\Delta_P \downarrow$	$\Delta_C \downarrow$	$\Delta_D \downarrow$	$\Delta_P \downarrow$
C1	-30.89%	-28.33%	-31.87%	-9.67%	-4.43%	-8.34%	-30.99%	-29.76%	-28.99%
C1 + C2	-39.23%	-36.85%	-38.39%	-13.21%	-8.05%	-11.06%	-34.32%	-33.30%	-31.97%
C1 + C2 + C3	**-41.49%**	**-38.80%**	**-40.30%**	**-20.77%**	**-14.30%**	**-18.12%**	**-36.55%**	**-35.36%**	**-34.24%**

Table 3: Ablation of heuristic action supervision strategies. Lower is better for all metrics.

Variant	SR			CA			PS		
	$\Delta_C \downarrow$	$\Delta_D \downarrow$	$\Delta_P \downarrow$	$\Delta_C \downarrow$	$\Delta_D \downarrow$	$\Delta_P \downarrow$	$\Delta_C \downarrow$	$\Delta_D \downarrow$	$\Delta_P \downarrow$
Pref. Only	-41.49%	-38.80%	-40.30%	-20.77%	-14.30%	-18.12%	-36.55%	-35.36%	-34.24%
Heur. Only	-30.84%	-27.80%	-31.91%	-9.42%	-3.03%	-5.43%	-16.30%	-15.59%	-13.41%
Pref. + Heur.	-40.98%	-38.13%	-39.95%	-17.52%	-11.37%	-13.58%	-33.57%	-32.65%	-30.88%
Pref. + RG	-42.32%	-39.50%	**-41.85%**	-19.97%	-13.83%	-17.36%	-36.97%	-36.09%	-35.36%
SAT-ACT	**-42.46%**	**-39.63%**	-41.60%	**-20.85%**	**-14.48%**	**-18.63%**	**-43.66%**	**-42.55%**	**-42.40%**

without heuristic action supervision. The results are reported in Table 2. These results show that the complete C1+C2+C3 combination achieves the strongest reductions in conflicts, decisions, and propagations on all three datasets. The contribution of each added criterion varies across datasets, but every addition improves all three search-cost metrics. This pattern suggests that the three outcome comparison criteria provide complementary supervision, with terminal-closure information adding useful signals beyond completion and trajectory Pareto evidence. A complete leave-one-out analysis is provided in Appendix G.2.

Exp-3. Ablation of Non-Dominated Heuristic Action Supervision. Under the 5-call setting, we compare SAT-ACT with four supervision variants. *Pref. Only* trains solely from outcome-derived action preferences and removes heuristic action supervision. *Heur. Only* removes preference supervision and trains only on actions selected by the solver's internal branching heuristic. *Pref. + Heur.* combines preference supervision with every heuristic target, regardless of its relation to the evaluated actions. *Pref. + RG* instead retains heuristic targets through random gating. The full SAT-ACT model uses outcome-derived preferences to retain a heuristic target only when its action remains non-dominated. Table 3 reports the results. *Heur. Only* and *Pref. + Heur.* underperform *Pref. Only* across all three datasets, while random gating yields inconsistent improvements. SAT-ACT achieves the best Δ_C and Δ_D on all datasets and the best Δ_P on CA and PS, while *Pref. + RG* is slightly better on SR. These results support filtering heuristic targets with outcome-derived preferences. The complete weight-sensitivity results are provided in Appendix G.3.

Exp-4. Evaluation on SAT Competition Benchmarks. We evaluate SAT-ACT with CaDiCaL on SAT Competition Main Track [1] instances in 2023–2025. Each year's competition contains 400 instances. Both solvers use the same CaDiCaL configuration and per-instance time limit, differing only in neural branching guidance. As shown in Table 4, the guided solver matches the number of solved instances on SC-2023 and solves two and three additional instances on SC-2024 and SC-2025, respectively. It also reduces PAR-2 on all three benchmark sets. The independent training setup, evaluation protocol, and neural-call overhead measurements are provided in Appendix G.7.

5 Related Work

In this section, we review related studies on learning-based approaches for SAT solving.

End-to-End Learning for SAT Solving. End-to-end approaches learn to reason directly over neural representations of SAT instances. Early methods focus on CNF representations. Bünz & Lamm (2017) apply GNNs to satisfiability prediction using a signed-literal graph with clause co-occurrence and complementarity edges. NeuroSAT (Selsam et al., 2019) uses a literal–clause graph with connections between complementary literals and performs iterative message passing between literal and clause nodes. Permutation-invariant architectures have also been developed to learn directly from CNF formulas (Cameron et al., 2020). Another group of methods targets circuit representations.

[1] https://satcompetition.github.io/

Table 4: Solver integration results on SAT Competition 2023–2025 instances. Solved is the sum of SAT and UNSAT results, and lower PAR-2 is better.

Dataset	Solver	#Solved	#SAT	#UNSAT	#Timeout	PAR-2 (s)
SC-2023	CaDiCaL (Native)	**251**	**116**	135	**149**	4171.07
	CaDiCaL (SAT-ACT)	**251**	115	**136**	**149**	**4150.55**
SC-2024	CaDiCaL (Native)	264	**137**	127	136	3989.14
	CaDiCaL (SAT-ACT)	**266**	135	**131**	**134**	**3952.30**
SC-2025	CaDiCaL (Native)	253	**144**	109	147	4214.04
	CaDiCaL (SAT-ACT)	**256**	143	**113**	**144**	**4094.72**

DG-DAGRNN (Amizadeh et al., 2019) represents circuit SAT instances as directed acyclic graphs and learns satisfying assignments with gated recursive networks. DeepSAT (Li et al., 2023) instead operates on AND-inverter graphs and incorporates polarity-aware representations for Boolean reasoning. Neural reasoning has also been extended to related satisfiability tasks. HyperSAT (Chen et al., 2025) develops an unsupervised hypergraph neural framework for Weighted MaxSAT, while NSNet (Li & Si, 2022) formulates SAT and #SAT as probabilistic inference with neural belief propagation. These approaches perform prediction, assignment construction, or optimization directly from the input instance. Our work instead retains the CDCL solving process and learns to guide branching decisions from dynamic solver states.

Learning-Guided SAT Solving. Learning-guided methods integrate neural predictions into established SAT solvers. Beyond branching, learned models have been used to guide restart policies (Liang et al., 2018) and clause deletion (Liu et al., 2024). Learning-guided techniques have also been applied to stochastic local search for variable selection (Yolcu & Póczos, 2019) and assignment initialization (Zhang et al., 2020). GraSS (Zhang et al., 2024) learns solver selection from graph-based representations. For CDCL branching, several methods learn structural properties of SAT formulas. NeuroCore (Selsam & Bjørner, 2019) predicts unsat-core variables to guide variable selection, while SATformer (Shi et al., 2023) improves this prediction through transformer-based representations. Sun et al. (2026) develop polarity-aware representations over hypergraphs for this task, and NeuroBack (Wang et al., 2024) predicts backbone phases for phase initialization. Other methods obtain branching guidance from the solving process. Graph-Q-SAT (Kurin et al., 2020) learns through reinforcement learning and solver interaction. RLAF (Tönshoff & Grohe, 2026) uses solver-level algorithm feedback to predict variable weights and polarities, while ImitSAT (Zhang et al., 2026) learns branching decisions from CDCL execution traces.

Candidate evaluation under a shared search context also has precedents in combinatorial search. Classical SAT lookahead solvers probe candidate branches through propagation-based effects (Heule & van Maaren, 2021), while strong branching in MILP evaluates candidate variables under a shared branch-and-bound context (Khalil et al., 2016; Gasse et al., 2019). These precedents demonstrate the value of comparing candidate decisions under a shared search context. Inspired by these precedents, we construct partial action preferences from state-matched rollout outcomes at dynamic CDCL states and use them to train a branching policy that operates without online candidate probing.

6 Conclusion

In this paper, we presented SAT-ACT, an outcome-derived learning framework for CDCL branching. The framework compares independent rollouts initiated from the same solver state and applies three complementary outcome comparison criteria to derive partial preferences without a predefined scalar action-quality score. A state-conditioned branching action network learns branching priorities from these preferences, while non-dominated heuristic action supervision retains solver-heuristic guidance only when the selected choice remains non-dominated.

Experiments on three SAT families show that SAT-ACT consistently reduces CDCL search costs relative to native CaDiCaL and compares favorably with existing learning-based methods. It also improves PAR-2 across three SAT Competition benchmarks and increases overall solving coverage. Ablation studies further demonstrate the contributions of the outcome comparison criteria and non-dominated heuristic action supervision. Future work includes reducing offline rollout costs, improving candidate and intervention selection, and extending the framework to broader SAT distributions and solver configurations.

AI USE STATEMENT

Generative AI tools were used to support research ideation and execution by assisting with brainstorming, refining methodological choices, and providing feedback on experimental plans and result interpretation. They were also used to draft portions of the manuscript and revise the authors' text for clarity and readability.

We did not use generative AI tools to generate synthetic datasets or formulate or prove mathematical claims. The remaining required disclosure tasks were either not performed with generative AI or were not applicable to this work.

The authors reviewed all AI-assisted suggestions and manuscript text. Methodological decisions, implementations, experiments, results, and claims were checked by the authors against the source code and experimental records. The authors determined the final content and take full responsibility for the work.

REFERENCES

Saeed Amizadeh, Sergiy Matusevych, and Markus Weimer. Learning to solve circuit-sat: An unsupervised differentiable approach. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019.

Armin Biere, Marijn Heule, Hans van Maaren, and Toby Walsh (eds.). *Handbook of Satisfiability - Second Edition*, volume 336 of *Frontiers in Artificial Intelligence and Applications*. IOS Press, 2021. ISBN 978-1-64368-160-3.

Armin Biere, Tobias Faller, Katalin Fazekas, Mathias Fleury, Nils Froleyks, and Florian Pollitt. CaDiCaL 2.0. In Arie Gurfinkel and Vijay Ganesh (eds.), *Computer Aided Verification - 36th International Conference, CAV 2024, Montreal, QC, Canada, July 24-27, 2024, Proceedings, Part I*, volume 14681 of *Lecture Notes in Computer Science*, pp. 133–152. Springer, 2024. doi: 10.1007/978-3-031-65627-9_7.

Benedikt Bünz and Matthew Lamm. Graph neural networks and boolean satisfiability. *CoRR*, abs/1702.03592, 2017.

Markus Büttner and Jussi Rintanen. Satisfiability planning with constraints on the number of actions. In Susanne Biundo, Karen L. Myers, and Kanna Rajan (eds.), *Proceedings of the Fifteenth International Conference on Automated Planning and Scheduling (ICAPS 2005), June 5-10 2005, Monterey, California, USA*, pp. 292–299. AAAI, 2005.

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, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pp. 3324–3331. AAAI Press, 2020.

Qiyue Chen, Shaolin Tan, Suixiang Gao, and Jinhu Lü. Hypersat: Unsupervised hypergraph neural networks for weighted maxsat problems. *CoRR*, abs/2504.11885, 2025. doi: 10.48550/ARXIV.2504.11885.

Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using RNN encoder–decoder for statistical machine translation. In Alessandro Moschitti, Bo Pang, and Walter Daelemans (eds.), *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 1724–1734, Doha, Qatar, October 2014. Association for Computational Linguistics. doi: 10.3115/v1/D14-1179. URL https://aclanthology.org/D14-1179/.

Stephen A. Cook. The complexity of theorem-proving procedures. In Michael A. Harrison, Ranan B. Banerji, and Jeffrey D. Ullman (eds.), *Proceedings of the 3rd Annual ACM Symposium on Theory of Computing, May 3-5, 1971, Shaker Heights, Ohio, USA*, pp. 151–158. ACM, 1971. doi: 10.1145/800157.805047.

Maxime Gasse, Didier Chételat, Nicola Ferroni, Laurent Charlin, and Andrea Lodi. Exact combinatorial optimization with graph convolutional neural networks. *Advances in neural information processing systems*, 32, 2019.

Jesús Giráldez-Cru and Jordi Levy. A modularity-based random SAT instances generator. In Qiang Yang and Michael J. Wooldridge (eds.), *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015*, pp. 1952–1958. AAAI Press, 2015.

Jesús Giráldez-Cru and Jordi Levy. Locality in random sat instances. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, pp. 638–644, 2017. doi: 10.24963/ijcai.2017/89.

Marijn J. H. Heule and Hans van Maaren. Look-ahead based SAT solvers. In Armin Biere, Marijn Heule, Hans van Maaren, and Toby Walsh (eds.), *Handbook of Satisfiability - Second Edition*, volume 336 of *Frontiers in Artificial Intelligence and Applications*, pp. 183–212. IOS Press, 2021. doi: 10.3233/FAIA200988.

Robert G. Jeroslow and Jinchang Wang. Solving propositional satisfiability problems. *Ann. Math. Artif. Intell.*, 1:167–187, 1990. doi: 10.1007/BF01531077.

Elias Khalil, Pierre Le Bodic, Le Song, George Nemhauser, and Bistra Dilkina. Learning to branch in mixed integer programming. In *Proceedings of the AAAI conference on artificial intelligence*, volume 30, 2016.

Vitaly Kurin, Saad Godil, Shimon Whiteson, and Bryan Catanzaro. Can q-learning with graph networks learn a generalizable branching heuristic for a SAT solver? In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.

Min Li, Zhengyuan Shi, Qiuxia Lai, Sadaf Khan, Shaowei Cai, and Qiang Xu. On eda-driven learning for SAT solving. In *60th ACM/IEEE Design Automation Conference, DAC 2023, San Francisco, CA, USA, July 9-13, 2023*, pp. 1–6. IEEE, 2023. doi: 10.1109/DAC56929.2023.10248001.

Zhaoyu Li and Xujie Si. Nsnet: A general neural probabilistic framework for satisfiability problems. In Sanmi Koyejo, S. Mohamed, A. Agarwal, Danielle Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*, 2022.

Zhaoyu Li, Jinpei Guo, and Xujie Si. G4satbench: Benchmarking and advancing SAT solving with graph neural networks. *Trans. Mach. Learn. Res.*, 2024, 2024.

Jia Hui Liang, Chanseok Oh, Minu Mathew, Ciza Thomas, Chunxiao Li, and Vijay Ganesh. Machine learning-based restart policy for cdcl sat solvers. In *International Conference on Theory and Applications of Satisfiability Testing*, pp. 94–110. Springer, 2018.

Hongduo Liu, Peng Xu, Yuan Pu, Lihao Yin, Hui-Ling Zhen, Mingxuan Yuan, Tsung-Yi Ho, and Bei Yu. Neuroselect: Learning to select clauses in SAT solvers. In Vivek De (ed.), *Proceedings of the 61st ACM/IEEE Design Automation Conference, DAC 2024, San Francisco, CA, USA, June 23-27, 2024*, pp. 131:1–131:6. ACM, 2024. doi: 10.1145/3649329.3656250.

Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019.

Ilya Mironov and Lintao Zhang. Applications of sat solvers to cryptanalysis of hash functions. In *International Conference on Theory and Applications of Satisfiability Testing*, pp. 102–115. Springer, 2006.

Daniel Selsam and Nikolaj Bjørner. Guiding high-performance sat solvers with unsat-core predictions. In *International conference on theory and applications of satisfiability testing*, pp. 336–353. Springer, 2019.

Daniel Selsam, Matthew Lamm, Benedikt Bünz, Percy Liang, Leonardo de Moura, and David L. Dill. Learning a SAT solver from single-bit supervision. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019.

Zhengyuan Shi, Min Li, Yi Liu, Sadaf Khan, Junhua Huang, Hui-Ling Zhen, Mingxuan Yuan, and Qiang Xu. Satformer: Transformer-based UNSAT core learning. In *IEEE/ACM International Conference on Computer Aided Design, ICCAD 2023, San Francisco, CA, USA, October 28 - Nov. 2, 2023*, pp. 1–4. IEEE, 2023. doi: 10.1109/ICCAD57390.2023.10323731.

Zhenchao Sun, Shuai Ma, Ping Lu, and Chongyang Tao. Unsat core prediction through polarity-aware representation learning over clause-literal hypergraphs. In *Forty-third International Conference on Machine Learning*, 2026. URL https://openreview.net/forum?id=YdQq7DvUfI.

Jan Tönshoff and Martin Grohe. Learning from algorithm feedback: One-shot sat solver guidance with gnns. In *International Conference on Learning Representations*, volume 2026, pp. 98621–98642, 2026.

Yakir Vizel, Georg Weissenbacher, and Sharad Malik. Boolean satisfiability solvers and their applications in model checking. *Proc. IEEE*, 103(11):2021–2035, 2015. doi: 10.1109/JPROC.2015.2455034.

Wenxi Wang, Yang Hu, Mohit Tiwari, Sarfraz Khurshid, Kenneth L. McMillan, and Risto Miikkulainen. Neuroback: Improving CDCL SAT solving using graph neural networks. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024.

Emre Yolcu and Barnabás Póczos. Learning local search heuristics for boolean satisfiability. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d'Alché-Buc, Emily B. Fox, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pp. 7990–8001, 2019.

Wenjie Zhang, Zeyu Sun, Qihao Zhu, Ge Li, Shaowei Cai, Yingfei Xiong, and Lu Zhang. Nlocalsat: Boosting local search with solution prediction. In Christian Bessiere (ed.), *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020*, pp. 1177–1183. ijcai.org, 2020. doi: 10.24963/IJCAI.2020/164.

Zewei Zhang, Huan Liu, Yuanhao Yu, Jun Chen, and Xiangyu Xu. Boolean satisfiability via imitation learning. In *International Conference on Learning Representations*, volume 2026, pp. 43867–43889, 2026.

Zhanguang Zhang, Didier Chételat, Joseph Cotnareanu, Amur Ghose, Wenyi Xiao, Hui-Ling Zhen, Yingxue Zhang, Jianye Hao, Mark Coates, and Mingxuan Yuan. Grass: Combining graph neural networks with expert knowledge for SAT solver selection. In Ricardo Baeza-Yates and Francesco Bonchi (eds.), *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD 2024, Barcelona, Spain, August 25-29, 2024*, pp. 6301–6311. ACM, 2024. doi: 10.1145/3637528.3671627.

A Additional Details of Outcome Collection and Action Preference Extraction

This section provides additional details of outcome collection and action preference extraction in Section 3.1. We describe how decision states and candidate branching actions are selected, how independent rollouts are run from the same solver state, how the rollout budget is defined, and how terminal-closure levels are determined.

A.1 Decision-State Sampling

For each solver run, we retain a small number of decision states distributed across different stages of the search. Early decision states are always represented, while additional states are sampled from later parts of the search trajectory. We also retain states reached after restarts so that the collected states cover both the initial search phase and later search contexts.

Each retained state corresponds to the solver state S_t immediately before a branching decision. The state contains the active clauses, current partial assignment, decision trail, variable activities, saved phases, and other dynamic search information accumulated during CDCL solving.

A.2 Candidate Branching Action Sampling and Rollouts

Let $\mathcal{B}_t$ denote the legal branching actions at S_t. We construct a bounded candidate set

$$\mathcal{A}_t = \{a_t^1, \ldots, a_t^{K_t}\}, \qquad \mathcal{A}_t \subseteq \mathcal{B}_t. \tag{13}$$

Candidate construction. We combine the branching variable selected by the solver's internal heuristic with activity-based, residual Jeroslow–Wang (JW) (Jeroslow & Wang, 1990), and pseudo-random proposals. The heuristic variable preserves the solver's reference action. Activity reflects accumulated search history, residual JW captures the current unsatisfied-clause structure, and random proposals broaden coverage beyond the heuristic rankings. This mixture is a practical design under a limited rollout budget; we do not claim that these sources are optimal or exhaustive.

We first retain the heuristic variable. We then construct three proposal pools with up to eight legal branching variables in each pool. The pools contain the highest-activity variables, the highest-scoring variables under residual JW, and a reproducible pseudorandom selection. For residual JW, each unsatisfied active clause contributes 2^{-r} to each of its unassigned active literals, where r is the number of such literals in the clause. A variable is ranked by the larger of its two literal scores.

Starting with the heuristic variable, we add previously unselected variables in round-robin order across the activity, JW, and random proposals until eight distinct variables have been retained or the pools are exhausted. This interleaving prevents an earlier proposal source from filling the entire candidate budget before the others are considered. We run a separate rollout for each polarity of every retained variable, yielding $K_t \leq 16$ distinct branching actions. The heuristic action is therefore retained in $\mathcal{A}_t$. An additional unforced solver continuation is recorded separately as a baseline and does not define another action in $\mathcal{A}_t$.

State reconstruction. For every action rollout, we reconstruct the target pre-action state by deterministic re-execution rather than by loading a solver snapshot. We launch a fresh solver process, load the original CNF, and re-execute the original search using the same solver binary, configuration, and random seed as in state collection. No candidate intervention is applied before reaching the target. Re-executing the search prefix regenerates the active clause database, partial assignment, decision trail, variable activities, and saved phases through the solver's normal updates; it is not merely a reassignment of recorded decision literals. At the target decision point, we check consistency with the recorded state locator before applying the candidate action. Variable activities and saved phases are reconstructed through deterministic execution rather than separately loaded or directly checked by the locator. For offline data collection, we use CaDiCaL 3.0.0 in `plain` mode with seed 1 and `stabilizeonly=1`.

Action intervention and rollout. At the reconstructed state S_t, we override exactly one branching decision with the candidate literal $a \in \mathcal{A}_t$. This is an ordinary, backtrackable decision, not

an assumption or an added unit clause. Subsequent decisions follow the solver's internal CDCL procedure, and normal CDCL backtracking may undo the intervention and earlier non-root assignments. Thus, a completed rollout reports the global SAT/UNSAT result of the original instance, not the satisfiability of a formula conditioned on the candidate literal. Completed rollouts for the same instance therefore share its global SAT/UNSAT result; actions differ in whether they finish within the budget and in the search effort required.

Each rollout records its solving status and the numbers of conflicts, propagations, decisions, and restarts, corresponding to $O_t(a)$ in Section 3.1. These counters are measured relative to the pre-action state, excluding the replay prefix, and the decision count includes the intervened decision itself. A rollout is labeled SAT or UNSAT if it reaches a terminal result within the rollout budget, and incomplete if a budget limit terminates it before a terminal result is reached. All candidate actions rolled out from the same state use the same solver configuration and rollout budget, as specified in Appendix A.3. Preference pairs are extracted only between actions rolled out from the same pre-action state; no cross-state or cross-instance cost comparisons are used to assign preference labels.

A.3 Rollout Budget

To bound the cost of action rollouts, each rollout is executed under a fixed rollout budget. We represent the budget as

$$\mathbf{b}_{\mathrm{roll}} = (C_{\mathrm{lim}}, D_{\mathrm{lim}}, P_{\mathrm{lim}}, T_{\mathrm{lim}}), \tag{14}$$

where C_{lim}, D_{lim}, and P_{lim} denote the maximum numbers of conflicts, decisions, and propagations allowed during the rollout, respectively, and T_{lim} denotes the maximum rollout time. For offline supervision collection, we use propagation horizons of 1,000, 15,000, and 5,000 for CA, PS, and SR, respectively, with a 300-second rollout timeout and no conflict or decision limit.

These quantities are measured from the point at which the candidate branching action is applied. The rollout terminates when the solver reaches SAT or UNSAT or when one of the prescribed limits is reached. Accordingly, the solving status is defined as

$$y_t(a) = \begin{cases} \mathrm{SAT}, & \text{if the rollout reaches a satisfying assignment,} \\ \mathrm{UNSAT}, & \text{if the rollout proves unsatisfiability,} \\ \mathrm{incomplete}, & \text{if a rollout limit is reached before termination.} \end{cases} \tag{15}$$

The same rollout budget is applied to the candidate actions rolled out from the same solver state. For an incomplete rollout, the solver statistics observed before reaching the rollout limit are retained as part of $O_t(a)$.

A.4 Terminal Closure

Terminal closure provides additional evidence about how directly a completed rollout reaches its final SAT or UNSAT result. The decision count $D_t(a)$ includes the forced branching decision itself. Therefore,

$$D_t(a) = 1 \tag{16}$$

indicates that no additional branching decision is required after applying a.

For a SAT rollout, we define the terminal-closure level as

$$\kappa_t(a) = \begin{cases} 2, & C_t(a) = 0,\ R_t(a) = 0,\ D_t(a) = 1, \\ 1, & C_t(a) = 0,\ R_t(a) = 0,\ D_t(a) > 1, \\ 0, & C_t(a) > 0 \text{ or } R_t(a) > 0. \end{cases} \tag{17}$$

The highest level corresponds to a satisfying assignment reached without any additional branching decision, conflict, or restart. The intermediate level corresponds to a conflict-free and restart-free rollout that still requires additional branching. A rollout containing a conflict or restart receives the lowest SAT closure level.

For an UNSAT rollout, we define

$$\kappa_t(a) = \begin{cases} 1, & D_t(a) = 1, \\ 0, & D_t(a) > 1. \end{cases} \tag{18}$$

The higher UNSAT closure level is assigned when the forced action leads to the terminal contradiction without another branching decision. Otherwise, the action receives the lower level. Conflicts are not used to distinguish UNSAT closure levels because reaching an UNSAT result in CDCL necessarily involves conflict-based reasoning.

For two completed actions a_i and a_j, a higher terminal-closure level provides preference evidence in favor of that action. Formally,

$$\kappa_t(a_i) > \kappa_t(a_j) \quad \Longrightarrow \quad a_i \succ_{S_t} a_j. \tag{19}$$

Equal closure levels provide no additional preference evidence.

B Node Features for Dynamic CDCL States

We provide the detailed node features used for the dynamic CDCL state representation in Section 3.2. The features describe the current assignment, local formula structure, and global search context. Normalized quantities are denoted by $\widetilde{\cdot}$.

B.1 Literal Features

For a literal ℓ associated with variable v, we use

$$\mathbf{x}_\ell^L = \left[u_\ell, t_\ell, f_\ell, \widetilde{d}_v^{\,\mathrm{asg}}, \widetilde{o}_\ell, \log(1+o_\ell), \log(1+|\mathcal{V}_t|), \log(1+|\mathcal{C}_t|), \rho_t, \widetilde{d}_t \right]. \tag{20}$$

Here, u_ℓ, t_ℓ, and f_ℓ indicate whether the literal is unassigned, true, or false under the current partial assignment. $\widetilde{d}_v^{\,\mathrm{asg}}$ denotes the normalized assignment level of variable v. o_ℓ and $\widetilde{o}_\ell$ denote the occurrence count of ℓ in the active clauses and its normalized counterpart. ρ_t is the fraction of currently assigned variables, and $\widetilde{d}_t$ is the normalized current decision level. The graph-size terms provide the current numbers of variables and active clauses on a logarithmic scale.

B.2 Clause Features

For a clause c, let n_c^T, n_c^F, and n_c^U denote the numbers of true, false, and unassigned literals. Its feature vector is

$$\mathbf{x}_c^C = \left[\log(1+|c|), \frac{n_c^T}{|c|}, \frac{n_c^F}{|c|}, \frac{n_c^U}{|c|}, q_c^{\mathrm{sat}}, q_c^{(1)}, q_c^{(2)} \right]. \tag{21}$$

Here, q_c^{sat} indicates whether the clause is currently satisfied, while $q_c^{(1)}$ and $q_c^{(2)}$ indicate unsatisfied clauses with one and two unassigned literals, respectively.

C Details of Polarity-Aware Branching Action Scoring

This section provides additional details of the representation modules used for branching-action scoring in Section 3.2.

For each variable v, polarity-symmetric variable encoding constructs

$$\mathbf{u}_v = \phi_{\mathrm{sym}}\left([\mathbf{h}_{v^+} + \mathbf{h}_{v^-} \,;\, |\mathbf{h}_{v^+} - \mathbf{h}_{v^-}|] \right). \tag{22}$$

Graph-level context aggregation first pools the variable and clause representations

$$\mathbf{m}_t = \frac{\sum_{v \in \mathcal{V}_t} \mathbf{u}_v + \sum_{c \in \mathcal{C}_t} \mathbf{h}_c}{|\mathcal{V}_t| + |\mathcal{C}_t|}. \tag{23}$$

The pooled representation is combined with a learnable global token to obtain the graph-level context

$$\mathbf{g}_t = \Phi_{\text{graph}}\left(\mathbf{z}_g, \phi_{\text{ctx}}(\mathbf{m}_t)\right), \tag{24}$$

where $\mathbf{z}_g$ is a learnable graph token, ϕ_{ctx} projects the pooled representation, and Φ_{graph} is a learnable context aggregation module.

Graph-conditioned variable encoding then incorporates this context into each variable representation

$$\widetilde{\mathbf{u}}_v = \Phi_{\text{var}}\left(\mathbf{u}_v, \mathbf{g}_t\right), \tag{25}$$

where Φ_{var} is a shared learnable conditioning module. In our implementation, both context aggregation modules use GRUs (Cho et al., 2014).

Finally, for a branching action a associated with variable $v(a)$, the action score is

$$s_\theta(S_t, a) = \phi_{\text{score}}\left(\left[\mathbf{h}_a \,;\, \widetilde{\mathbf{u}}_{v(a)} \,;\, \mathbf{g}_t\right]\right), \tag{26}$$

where ϕ_{score} is shared across both polarities. The action-specific literal representation $\mathbf{h}_a$ distinguishes the two polarities of the same variable, while $\widetilde{\mathbf{u}}_{v(a)}$ and $\mathbf{g}_t$ provide variable-level and graph-level context.

D Datasets and Supervision Statistics

We use the SR, Community Attachment (CA), and Popularity-Similarity (PS) families from G4SATBench (Li et al., 2024). Each family contains both SAT and UNSAT instances. The instances are generated using the SR generator from NeuroSAT (Selsam et al., 2019), the CA model (Giráldez-Cru & Levy, 2015), and the PS model (Giráldez-Cru & Levy, 2017). For each family, we use 16,000/2,000/2,000 instances for training, validation, and offline testing, respectively, with balanced SAT and UNSAT instances. The supervision statistics below aggregate the training, validation, and offline test splits.

For each instance, an instrumented CaDiCaL solver (Biere et al., 2024) retains up to six decision states and runs rollouts for up to sixteen candidate branching actions at each retained state. The resulting upper bounds are 360,000 retained states and 5.76 million independent CDCL rollouts. Recording one additional unforced solver continuation at each state gives an upper bound of 6.12 million solver runs. These are design upper bounds. An instance may yield fewer than six retained states, and a state may contain fewer than sixteen candidate actions. The unforced solver continuation is recorded separately and does not define an additional branching action.

Each action rollout reconstructs the target state by re-executing the original search prefix, as described in Appendix A.2. The three outcome comparison criteria in Section 3.1 extract partial action preferences from the resulting outcomes. Raw rollout generation was executed sequentially by family with 20 workers, taking 2 h 22 min for CA, 2 h 07 min for SR, and 2 h 27 min for PS, or approximately 6 h 55 min in total.

Table 5 summarizes the CNF statistics and the resulting outcome-derived supervision.

E Evaluation Metrics

We evaluate solver search cost using conflicts (C), decisions (D), and propagations (P). For these quantities, we report the percentage changes Δ_C, Δ_D, and Δ_P relative to native CaDiCaL, computed from the aggregated counts over the evaluated instances. Negative values indicate reductions in solver cost. Following ImitSAT (Zhang et al., 2026), we additionally report the median relative propagation per problem (MRPP) and the one-percent win rate W(1%). Let $\mathcal{I}$ denote the set of evaluated instances, and let P_i and P_i^{nat} denote the propagation counts of the evaluated method and native CaDiCaL on instance i, respectively. MRPP is defined as

$$\text{MRPP} = \text{median}_{i \in \mathcal{I}} \left(\frac{P_i}{P_i^{\text{nat}}} \right). \tag{27}$$

Values below one indicate fewer propagations than native CaDiCaL on the median instance.

Table 5: Dataset and outcome-derived supervision statistics. CNF statistics are reported per instance. Decision states and rolled-out actions are reported after supervision filtering. A state is included only if it produces at least one preference pair. Rolled-out actions count candidate branching actions and exclude the additional unforced solver baseline. Criterion coverage is measured over all action pairs within these retained states, and the three criteria are not mutually exclusive.

Statistic	**SR**	**CA**	**PS**
Formula Statistics (Mean / Min / Max)			
Variables	119.78 / 40 / 200	120.81 / 40 / 200	118.70 / 40 / 200
Clauses	652.90 / 159 / 1,233	1,671.15 / 175 / 2,999	823.15 / 224 / 1,599
Clause/Variable Ratio	5.48 / 3.54 / 8.72	13.64 / 4.38 / 15.00	6.93 / 5.13 / 8.00
Outcome-Derived Supervision			
Instances	20,000	20,000	20,000
Decision States	107,448	109,609	108,751
Rolled-Out Actions	1,712,508	1,753,744	1,738,400
Candidate Pairs	12,817,702	13,153,080	13,031,332
Preference Pairs	9,970,806	9,336,502	10,274,426
Criterion Coverage			
C1 Completion	521,391 (4.07%)	560,030 (4.26%)	311,065 (2.39%)
C2 Trajectory Pareto	9,411,303 (73.42%)	8,756,335 (66.57%)	9,912,676 (76.07%)
C3 Terminal Closure	270,985 (2.11%)	105,060 (0.80%)	268,018 (2.06%)
Combined Coverage	77.79%	70.98%	78.84%

To capture instance-wise improvements, we further report the one-percent win rate. Following the definition in ImitSAT, with margin $\delta = 0.01$,

$$\mathrm{W}(1\%) = \frac{100}{|\mathcal{I}|} \sum_{\substack{i \in \mathcal{I} \\ P_i^{\mathrm{nat}} > 0}} \mathbb{I}\left[P_i \le (1-\delta) P_i^{\mathrm{nat}}\right]. \tag{28}$$

W(1%) therefore measures the percentage of evaluated instances for which the method reduces propagations by at least 1% relative to native CaDiCaL. Higher values are better.

F Experimental Setup

F.1 Implementation Details

Our models are implemented in PyTorch with 10 message-passing rounds and two-layer MLPs. The hidden dimension is 128. Models are trained using AdamW for 50 epochs with a learning rate of 10^{-4}. For SR, CA, and PS, we use a batch size of 128, select the weight decay from $\{0, 10^{-5}, 5 \times 10^{-5}, 10^{-4}\}$, and retain the checkpoint with the lowest validation loss. The weight λ_{heur} for non-dominated heuristic action supervision is selected from $[0, 0.2]$. Setting $\lambda_{\mathrm{heur}} = 0$ removes this supervision and corresponds to SAT-ACT (Pref. Only). All hyperparameters are selected exclusively on the validation split and fixed before solver-level test evaluation. All models are trained on a machine equipped with an Intel Xeon Gold 6148 CPU at 2.40 GHz and an NVIDIA Tesla V100 PCIe GPU with 32 GB memory. For the SR, CA, and PS experiments, CaDiCaL runs on the CPU, while neural inference is performed on the GPU. Code is available at `https://anonymous.4open.science/r/SAT-ACT-Review-1DA2`.

F.2 Solver Evaluation

All solver-level evaluations use CaDiCaL 3.0.0 with `stabilizeonly=1` and `preprocesslight=0`, while the remaining solver options retain their default values. For neural intervention, both the node and cell size checks require $2|\mathcal{V}_t| + |\mathcal{C}_t| + |\mathcal{E}_t| \le 2{,}000{,}000$ for the literal–clause representation of each active CNF.

For each SAT family, we randomly sample 200 instances from the test split, with equal numbers of SAT and UNSAT instances. We repeat the sampling and solver evaluation five times, with all

Table 6: SAT and UNSAT subset analysis of SAT-ACT. Lower is better for all metrics.

Dataset	Budget	SAT			UNSAT		
		$\Delta_C \downarrow$	$\Delta_D \downarrow$	$\Delta_P \downarrow$	$\Delta_C \downarrow$	$\Delta_D \downarrow$	$\Delta_P \downarrow$
SR	3-call	-45.82%	-36.35%	-43.17%	-33.70%	-34.11%	-34.08%
	5-call	-51.43%	-41.42%	-47.33%	-37.34%	-37.95%	-37.88%
CA	3-call	-14.11%	-7.23%	-12.35%	-6.77%	-5.02%	-4.65%
	5-call	-18.47%	-9.46%	-15.82%	-23.87%	-25.94%	-22.95%
PS	3-call	-40.53%	-37.10%	-38.53%	-26.68%	-27.14%	-23.72%
	5-call	-59.47%	-54.92%	-58.22%	-29.31%	-29.98%	-26.15%

Table 7: Leave-one-out analysis of outcome comparison criteria. Lower is better for all metrics.

Dataset	Comparison Criteria	3-call			5-call		
		$\Delta_C \downarrow$	$\Delta_D \downarrow$	$\Delta_P \downarrow$	$\Delta_C \downarrow$	$\Delta_D \downarrow$	$\Delta_P \downarrow$
SR	C1 + C2 + C3	-36.84%	-33.57%	-36.45%	-41.49%	-38.80%	-40.30%
	w/o C1	-26.66%	-24.56%	-26.00%	-40.01%	-37.55%	-39.03%
	w/o C2	-29.71%	-27.03%	-30.00%	-32.76%	-30.81%	-32.35%
	w/o C3	-28.85%	-27.23%	-28.76%	-39.23%	-36.85%	-38.39%
	C1 only	-30.30%	-27.60%	-30.91%	-30.89%	-28.33%	-31.87%
CA	C1 + C2 + C3	-10.99%	-6.51%	-8.97%	-20.77%	-14.30%	-18.12%
	w/o C1	-7.78%	-4.24%	-6.48%	-14.74%	-9.33%	-12.51%
	w/o C2	-7.94%	-3.64%	-6.95%	-8.97%	-3.17%	-5.93%
	w/o C3	-9.83%	-5.39%	-7.55%	-13.21%	-8.05%	-11.06%
	C1 only	-5.83%	-2.35%	-5.32%	-9.67%	-4.43%	-8.34%
PS	C1 + C2 + C3	-30.75%	-29.68%	-28.38%	-36.55%	-35.36%	-34.24%
	w/o C1	-28.64%	-27.60%	-25.95%	-28.08%	-27.35%	-24.40%
	w/o C2	-19.51%	-18.73%	-16.39%	-31.74%	-30.80%	-30.14%
	w/o C3	-26.63%	-25.87%	-23.62%	-34.32%	-33.30%	-31.97%
	C1 only	-22.54%	-21.52%	-20.14%	-30.99%	-29.76%	-28.99%

methods evaluated on the same sampled instances and under the same solver settings. The reported results are averaged across the five runs.

Each solver run is limited to 100 seconds. Native CaDiCaL without neural guidance serves as the reference solver for computing Δ_C, Δ_D, Δ_P, MRPP, and W(1%).

For online branching methods, the 3-call and 5-call settings allow at most three and five model requests per solver run, respectively. Each request counts toward this intervention budget regardless of whether the proposed branching action is ultimately applied by the solver.

G Additional Experimental Results

G.1 Analysis of SAT and UNSAT Results

We separately evaluate SAT-ACT on SAT and UNSAT instances under both 3-call and 5-call settings. We report the same solver-level cost metrics used in the main comparison.

Table 6 shows that SAT-ACT consistently reduces conflicts, decisions, and propagations on both SAT and UNSAT instances across all three datasets. The magnitude of the improvement varies across datasets and termination regimes. This may reflect the different search characteristics of SAT and UNSAT solving, since SAT solving can terminate after finding one satisfying assignment, whereas UNSAT solving requires completing a proof. Overall, the results show that SAT-ACT remains effective under both termination regimes.

G.2 Additional Ablations of Outcome Comparison Criteria

We further analyze the three outcome comparison criteria using leave-one-out ablations. C1, C2, and C3 denote completion, trajectory Pareto, and terminal closure, respectively. In addition to the full C1+C2+C3 combination, we remove each criterion individually and include the C1-only variant used in the incremental analysis. All variants share the same model architecture and solver evaluation protocol, with preference-guided action ranking as the sole training objective. Table 7 provides a consistent picture across datasets and intervention budgets. The full C1+C2+C3 combination achieves the largest reductions in conflicts, decisions, and propagations in every evaluated setting, while removing any individual criterion degrades all three search-cost metrics. This indicates that C1, C2, and C3 contribute complementary supervision rather than being interchangeable.

The relative importance of individual criteria nevertheless varies across datasets and budgets. C2 is particularly important for CA under the 5-call setting and for PS under the 3-call setting, where removing it causes the largest degradation. By contrast, C1 has a stronger effect on SR with three calls and on PS with five calls. Removing C3 also consistently weakens performance, although the magnitude of the degradation varies across settings. Together, these results show that no single outcome comparison criterion explains the overall gain. Combining completion, trajectory-level, and terminal evidence yields the most reliable solver-level improvements across different SAT distributions and intervention budgets.

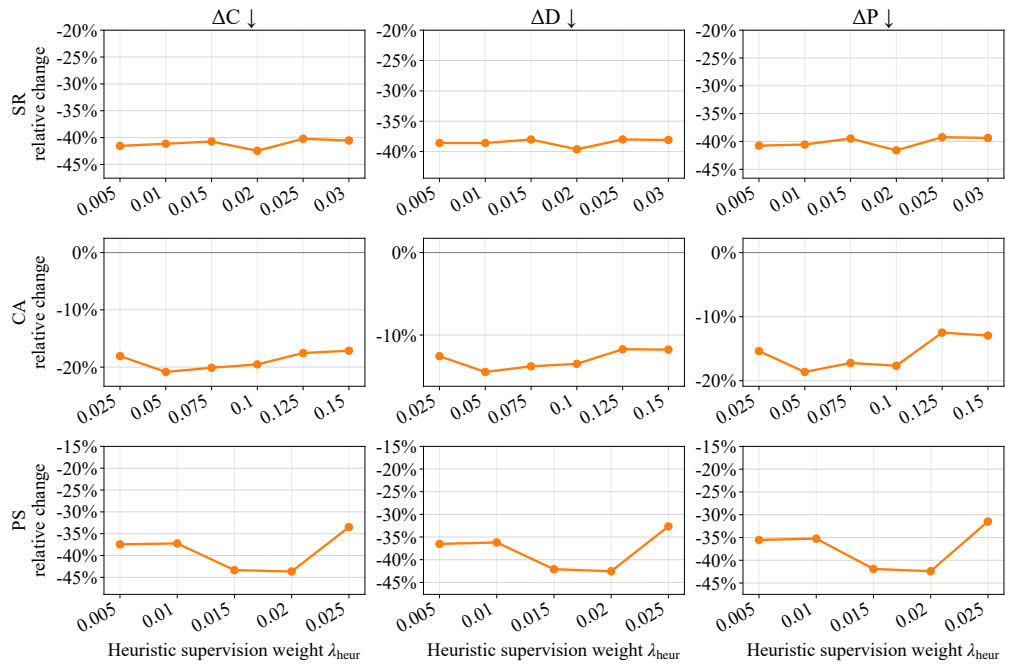


Figure 2: Complete sensitivity analysis of the heuristic action supervision weight on SR, CA, and PS under the 5-call setting. The plots report relative changes in conflicts (Δ_C), decisions (Δ_D), and propagations (Δ_P) with respect to native CaDiCaL; lower is better.

G.3 Sensitivity to Heuristic Action Supervision Weight

We study the sensitivity of non-dominated heuristic action supervision to its weight λ_{heur} on SR, CA, and PS under the 5-call setting. The preference-guided ranking objective, model architecture, and evaluation protocol remain fixed while λ_{heur} varies. Figure 2 reports the relative changes in conflicts, decisions, and propagations. The three solver-cost metrics exhibit broadly consistent trends across the tested weights. SR remains relatively stable over the evaluated range. CA shows comparable performance across moderate weights, with its strongest reductions generally occurring over $\lambda_{\text{heur}} \in [0.05, 0.1]$, while PS performs best over $\lambda_{\text{heur}} \in [0.015, 0.02]$. Increasing the heuristic action supervision weight beyond these ranges does not consistently improve performance.

These results indicate that non-dominated heuristic action supervision performs best with a moderate weight.

G.4 Effect of Neural Intervention Budget

We evaluate SAT-ACT (Pref. Only) and SAT-ACT under different neural intervention budgets. The budget controls the maximum number of model requests during each solver run. Figure 3 reports relative changes in conflicts, decisions, and propagations over the complete budget range. Overall, the three solver-cost metrics show broadly consistent behavior. Larger neural intervention budgets usually lead to larger reductions in search cost, but the gains are not strictly monotonic and local fluctuations appear across datasets and budgets.

The clearest budget sensitivity is observed on SR and CA, where all three metrics improve substantially as the neural intervention budget increases. PS shows a less smooth trend, with noticeable fluctuations across neighboring budgets, especially under the smaller and intermediate budget settings. Nevertheless, compared with very small budgets, larger budgets still provide markedly stronger reductions overall.

Across the three SAT families, SAT-ACT generally performs at least as well as SAT-ACT (Pref. Only) and often better, particularly on SR and CA. Its advantage is not uniform at every individual budget, especially on PS, but the overall pattern still supports the benefit of non-dominated heuristic action supervision across a broad range of intervention budgets. These results indicate that additional neural guidance remains useful beyond the first few decisions, while also showing that the marginal benefit becomes increasingly instance- and dataset-dependent at higher budgets.

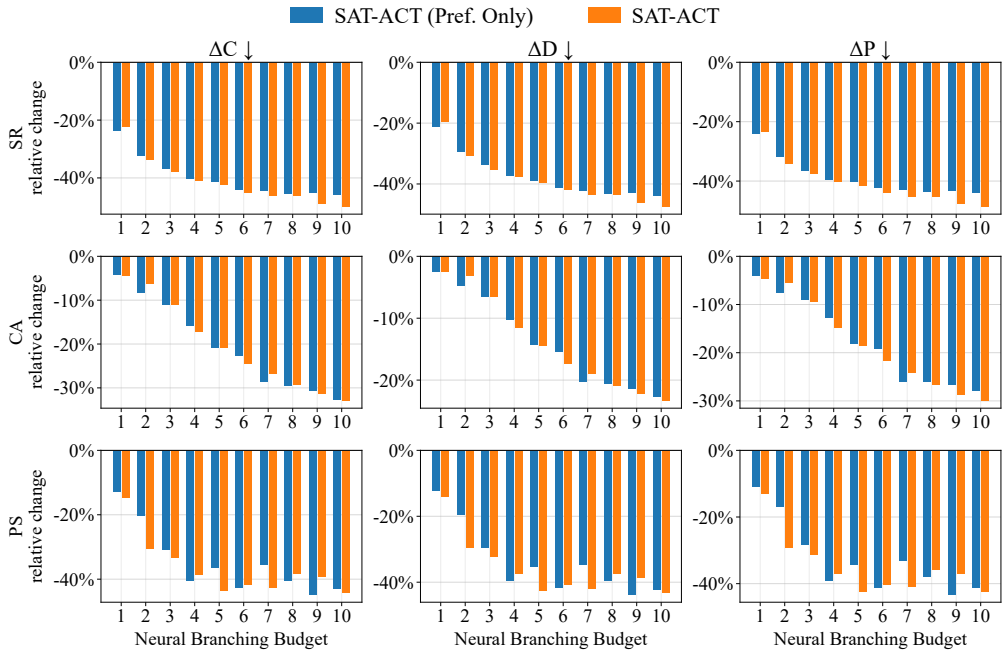


Figure 3: Complete neural intervention budget results on SR, CA, and PS. Bars show the relative changes in conflicts (Δ_C), decisions (Δ_D), and propagations (Δ_P) with respect to native CaDiCaL; lower is better.

G.5 Generalization to Hard Instances

We further evaluate whether the models used in the main experiments generalize to substantially harder SAT instances. Specifically, we directly apply the same trained models used for the main-experiment split to the corresponding Hard split, without retraining, fine-tuning, or adaptation on any Hard instances. Therefore, this experiment evaluates cross-difficulty generalization of the learned branching policies rather than performance after additional training on the target split.

Table 8: Direct generalization of the same trained models used in the main experiments to the corresponding Hard splits. Δ_C, Δ_D, and Δ_P denote relative changes in conflicts, decisions, and propagations, respectively. Lower is better for Δ_C, Δ_D, Δ_P, and MRPP; higher is better for W(1%).

Dataset	Method	$\Delta_C \downarrow$	$\Delta_D \downarrow$	$\Delta_P \downarrow$	MRPP $\downarrow$	W(1%) $\uparrow$
SR Hard	CaDiCaL	0.00%	0.00%	0.00%	1.000	0.00%
	SAT-ACT (PREF. ONLY)	-25.95%	-25.28%	-25.69%	0.704	78.10%
	SAT-ACT	**-28.24%**	**-27.44%**	**-27.87%**	**0.680**	**81.40%**
CA Hard	CaDiCaL	0.00%	0.00%	0.00%	**1.000**	0.00%
	SAT-ACT (PREF. ONLY)	-1.99%	-1.54%	-0.68%	1.001	38.70%
	SAT-ACT	**-3.72%**	**-3.78%**	**-3.22%**	1.001	**41.00%**
PS Hard	CaDiCaL	0.00%	0.00%	0.00%	1.000	0.00%
	SAT-ACT (PREF. ONLY)	-3.29%	-3.37%	-3.00%	**0.461**	80.80%
	SAT-ACT	**-3.92%**	**-4.00%**	**-3.60%**	0.482	**81.00%**

The Hard instances are substantially larger than those in the main-experiment split. SR Hard and CA Hard contain 200–400 variables, with average clause counts of 1616.47 and 4197.51, respectively, while PS Hard contains 200–300 variables with an average of 1761.98 clauses. We use the 3-call intervention setting for this generalization evaluation. The results are reported in Table 8.

The same models used in the main experiments retain clear solver-level improvements when directly transferred to the Hard splits. On SR Hard, SAT-ACT reduces conflicts, decisions, and propagations by 28.24%, 27.44%, and 27.87%, respectively, and also improves MRPP and W(1%) over SAT-ACT (PREF. ONLY). On CA Hard, SAT-ACT similarly achieves larger reductions in all three aggregate search-cost metrics, although the overall gains are smaller than on SR Hard.

On PS Hard, SAT-ACT improves the reductions in conflicts, decisions, and propagations from 3.29%, 3.37%, and 3.00% with SAT-ACT (PREF. ONLY) to 3.92%, 4.00%, and 3.60%, respectively, while slightly improving W(1%). Its MRPP is slightly worse than that of SAT-ACT (PREF. ONLY), showing that the instance-level metrics do not improve uniformly.

Overall, the results show that the branching policies learned in the main experiments remain effective when directly applied to substantially larger Hard instances. Since exactly the same trained models are used without any specific training or adaptation, these results provide evidence that the trained branching policies generalize beyond the instance scale observed in the main-experiment split.

G.6 NEURAL INFERENCE COST

Measurement Protocol. We compare the neural inference cost of SAT-ACT with NeuroCore, GQSAT, SATformer, ImitSAT, and RLAF. For methods that provide online branching guidance, we report results under matched 3-call and 5-call intervention budgets. SATformer and RLAF instead perform a single initialization before search. Inference time is measured by the model service and reported as average accumulated milliseconds per instance.

Inference-Cost Comparison. Table 9 reports the neural inference cost on SR, CA, and PS. These datasets contain relatively small instances that native CaDiCaL often solves within milliseconds. The neural inference cost therefore outweighs the runtime benefit obtained from the reduced search cost on these instances. The main results on SR, CA, and PS should consequently be interpreted as improvements in search efficiency rather than end-to-end runtime acceleration.

Under matched intervention budgets, the inference cost of SAT-ACT remains comparable to NeuroCore, GQSAT, and ImitSAT and is often lower than that of the latter two methods. SATformer and RLAF use one-shot initialization and do not provide repeated guidance at dynamic solver states. The online inference overhead of SAT-ACT is therefore comparable to existing neural methods with similar deployment mechanisms.

The runtime tradeoff changes on the substantially harder SAT Competition benchmarks analyzed in Appendix G.7. As shown in Table 4, neural guidance reduces PAR-2 on all three benchmark sets

Table 9: Neural inference cost of learning-based SAT branching methods. All values are average accumulated milliseconds per instance. Lower is better.

Method	Budget	Inference (ms)		
		SR	CA	PS
NeuroCore	3-call	29.6	65.9	37.2
NeuroCore	5-call	47.8	109.2	58.3
GQSAT	3-call	66.1	153.8	84.3
GQSAT	5-call	103.1	251.2	134.5
SATformer	one-shot	48.5	97.1	59.9
ImitSAT	3-call	41.3	80.1	53.7
ImitSAT	5-call	73.3	135.5	94.1
RLAF	one-shot	15.4	23.7	17.3
SAT-ACT	3-call	42.1	60.4	46.7
SAT-ACT	5-call	69.1	100.8	72.6

and solves additional SC-2024 and SC-2025 instances. On the 752 instances solved by both configurations, it reduces total runtime by 10.15% after accounting for neural-call overhead. These results show that the search-time reduction can outweigh neural inference cost on sufficiently challenging SAT instances.

G.7 SAT Competition Evaluation

Training Instances and Data Split. This experiment uses an independently constructed training set derived from the SAT Competition Main Track instances released between 2017 and 2022. We partition the instances into training, validation, and test splits using an 80/10/10 ratio and a split seed of 42. Before data collection, we retain formulas with at most 200,000 graph nodes and 1,000,000 literal occurrences. The graph contains two literal nodes for each variable and one node for each clause. No additional limits are imposed on the CNF file size or maximum clause length. The filtering process retains 956 formulas that satisfy these constraints and pass DIMACS parsing.

State and Action Collection. The SAT Competition data collection uses a larger candidate set and a different rollout budget from the G4SATBench setup. Each retained state is restricted to 200,000 graph nodes and 1,000,000 literal occurrences. At most 10 candidate variables are retained at each state, and both polarities initiate separate rollouts.

State collection and action rollouts use CaDiCaL 3.0.0 with its default solver configuration. Each forced-action rollout has a timeout of 1,800 seconds. No conflict, decision, or propagation limit is imposed. The resulting collection contains 114,120 reliable forced-action rollouts.

After removing states without a supported action preference, the resulting dataset contains 2,831 training states, 313 validation states, and 329 test states. The training, validation, and test splits contain 418,283, 45,312, and 48,272 directed preference pairs, respectively.

Solver Evaluation. Each of SC-2023, SC-2024, and SC-2025 contains 400 Main Track instances. We evaluate native CaDiCaL and CaDiCaL guided by SAT-ACT on all instances from the three benchmark sets. Both configurations use CaDiCaL 3.0.0 with `stabilizeonly=1` and `preprocessslight=0`. All remaining solver options retain their default values. Both CaDiCaL and neural inference run on the CPU in this evaluation. Each solver run has a time limit of 5,000 seconds and uses the 3-call setting, allowing at most three neural requests. The two configurations otherwise use identical solver settings.

A neural request is permitted only when both the node and cell size checks satisfy $2|\mathcal{V}_t| + |\mathcal{C}_t| + |\mathcal{E}_t| \leq 2{,}000{,}000$. Here, $\mathcal{E}_t$ denotes the literal–clause incidence cells, which correspond to literal occurrences.

Table 10: Neural-call overhead for the SAT Competition evaluation.

Dataset	Solver	NN Wait Ratio	NN Call Latency (s)
SC-2023	CaDiCaL (Native)	0.00%	–
	CaDiCaL (SAT-ACT)	11.99%	9.61
SC-2024	CaDiCaL (Native)	0.00%	–
	CaDiCaL (SAT-ACT)	5.79%	8.10
SC-2025	CaDiCaL (Native)	0.00%	–
	CaDiCaL (SAT-ACT)	7.92%	8.43

We evaluate aggregate solving time using PAR-2,

$$\mathrm{PAR2} = \frac{1}{N}\sum_{i=1}^{N}\begin{cases} T_i, & \text{if instance } i \text{ is solved,} \\ 2T_{\max}, & \text{otherwise,} \end{cases} \tag{29}$$

where T_i is the runtime of a solved instance and $T_{\max} = 5{,}000$ seconds. Each timed-out instance therefore contributes 10,000 seconds.

Neural-Call Overhead. We report the macro-averaged NN wait ratio,

$$\text{NN Wait Ratio} = \frac{1}{N}\sum_{i=1}^{N}\frac{T_{\text{NN-wait},i}}{T_{\text{solve},i}}, \tag{30}$$

which assigns equal weight to each evaluated instance. We also report the NN call latency,

$$\text{NN Call Latency} = \frac{\sum_i T_{\text{NN-wait},i}}{\sum_i N_{\text{RPC-attempts},i}}. \tag{31}$$

The latency measures the end-to-end solver waiting time for each neural request and covers communication, queueing, input processing, and model inference. Table 10 reports these measurements over all 400 instances in each benchmark set.

The NN wait ratios are 11.99%, 5.79%, and 7.92% on SC-2023, SC-2024, and SC-2025, respectively. The corresponding NN call latencies are 9.61, 8.10, and 8.43 seconds. These measurements cover both solved instances and runs that reach the 5,000-second time limit.

Analysis of Commonly Solved Instances. We define the commonly solved set as the instances for which both native and SAT-ACT-guided CaDiCaL return SAT or UNSAT. The two solvers produce identical SAT/UNSAT results on every instance in this set. Across the three benchmark sets, 752 instances are solved by both solvers, 16 are solved only by native CaDiCaL, 21 are solved only by the guided solver, and 411 remain unsolved by both. On SC-2023, both solvers solve 251 instances, but each uniquely solves four instances. The guided solver obtains net gains of two solved instances on SC-2024 and three on SC-2025, resulting in five additional solved instances overall.

Among the 752 commonly solved instances, the guided solver is faster on 448 instances, corresponding to 59.57%, while native CaDiCaL is faster on 304 instances. The guided solver reduces the total runtime by 61,101.75 seconds, or 10.15%, corresponding to an aggregate speedup of $1.113\times$. We define the instance-level runtime ratio as $R_i = T_{\text{native},i}/T_{\text{guided},i}$, where $R_i > 1$ favors SAT-ACT. The median ratios are $1.051\times$, $1.128\times$, and $1.055\times$ on SC-2023, SC-2024, and SC-2025, respectively. The corresponding aggregate speedups are $1.080\times$, $1.073\times$, and $1.193\times$. SC-2024 shows the largest median improvement, whereas SC-2025 shows the largest aggregate speedup. This difference suggests that the SC-2025 gains are concentrated more heavily on instances with longer solving times.

We also recompute the neural-call overhead on the commonly solved instances. The NN wait ratios are 19.26%, 8.76%, and 12.62% on SC-2023, SC-2024, and SC-2025, while the corresponding NN call latencies are 12.33, 7.61, and 9.09 seconds. Across all 752 instances, the NN wait ratio is 13.48% and the NN call latency is 9.67 seconds. These wait ratios are higher than the full-set values because the full-set averages contain many runs that reach the 5,000-second time limit. A small

number of neural calls accounts for only a small fraction of the total runtime on these timed-out runs. The commonly solved set therefore provides a more direct view of neural overhead during successful solving.

Together, these results show that neural guidance maintains outcome consistency, modestly expands solving coverage, and reduces end-to-end runtime on commonly solved instances after accounting for neural-call overhead, although the gains are not uniform across all instances.