

1 Supplementary Materials

2 1 Preliminaries

3 1.1 Background of SAT

4 Let $V = \{x_1, x_2, \dots, x_n\}$ be a set of Boolean variables. A *literal* is either a variable x or its negation $\neg x$. A *clause* is a disjunction
5 of literals. A *conjunctive normal form* (CNF) formula $F = C_1 \wedge C_2 \wedge \dots \wedge C_m$ is a conjunction of clauses. For simplicity we
6 assume all clauses are non-tautological. In other words, no variable x occurs positively ($x \in C$) and negatively ($\neg x \in C$) in the
7 same clause.

8 A (partial) mapping $\alpha : V \rightarrow \{0, 1\}$ is called an *assignment*. If α maps all variables to a boolean value, it is termed a
9 *complete* assignment; otherwise, it is referred to as a *partial* assignment. The value of a variable x under an assignment α is
10 denoted as $\alpha[x]$. An assignment α satisfies a clause if at least one literal evaluates to true under α , and satisfies a CNF formula
11 if it satisfies all its clauses. A CNF formula F is satisfiable if there is at least one satisfying assignment. The empty clause \perp
12 is always unsatisfiable and represents a *conflict*. SAT is the problem of deciding whether a given CNF formula is satisfiable.

13 To solve such the SAT problem, a straightforward approach is to conduct exhaustively search over all possible truth
14 assignments to the variables in V . A complete SAT solver implements this idea by exploring all possible assignments to
15 determine if at least one satisfies F (proving F is satisfiable). If no satisfying assignment exists, the solver conclusively proves
16 F is unsatisfiable. Within this framework, the strategies for branching (for choosing which variable to assign next and what
17 value to assign it) are crucial. Key techniques to enhance the search efficiency include restart strategies, restart strategies, and
18 activity-based heuristics (often implemented by "bumping" variable or clause activity scores).

19 1.2 CDCL solver

20 Conflict-Driven Clause Learning (CDCL)¹ is the most common approach and plays a dominant role in modern high-performance
21 SAT solvers. The core innovation is conflict analysis: when a conflict occurs, the solver analyzes it and derives a new clause
22 that explains the inconsistency. These derived clauses, called learned clauses, enable the solver to prune large portions of the
23 search space, which significantly improve the computational efficiency. However, an excessive number of learnt clauses can
24 degrade unit propagation speed and thus consume excessive memory. Consequently, identifying high-quality learnt clauses
25 and reducing their quantity are essential for maintaining a solver's performance.

26 Specifically, CDCL solvers operate on a propagation and learning mechanism, complemented by decision heuristics. The
27 implementation of propagation and learning is a standard practice and overall similar in different solvers. In contrast, the
28 decision policies (for variable selection) and restart policies (for search restarts) which are crucial for performance vary
29 significantly. Due to the need of manual design and the absence of rigorous mathematical proofs for their optimality, these
30 policies are referred to as important heuristics.

31 There is a long history of research on these heuristics in SAT solvers, among which branching heuristics play a crucial
32 role and continue to impact the performance of CDCL solvers. For example, Variable State Independent Decaying Sum
33 (VSIDS)² is a family of branching heuristics that seek to assign a value to the most promising variable in the Make Decision
34 phase. Another important branching heuristic is Learning-Rate Branching (LRB)³, which frames branching as an optimization
35 problem that picks a variable to maximize a metric called learning rate.

36 Restart heuristics is also essential for enhancing the performance of CDCL solvers. It allows the solver to abandon the
37 current search path and backtrack to a specific decision level. In this process, the learnt clauses are usually maintained for
38 the next search. Fast restart⁴ is a widely used method, and Luby restarts⁵ is also heavily used because they represent a prior
39 optimal strategy. However, most implementations are now switched to Glucose-style restarts⁶, which are widely used in the
40 SAT Competition. For more details, see the extensive overview given by Armin Biere⁷.

41 Rephase is another technique in CDCL solvers. The primary objective is to reset or adjust the current partial assignment,
42 thereby enabling the solver to explore a diverse search space. PrecoSAT and PicoSAT⁸ utilize a Jeroslow-Wang score⁹ to adjust
43 the saved phases either on all or only on irredundant clauses in regular intervals following a luby sequence. StrangeNight¹⁰
44 employs a strategy of flipping values with a certain probability that depends on the depth of the assignment. The motivation
45 is to avoid the heavy-tail phenomenon. These rephasing heuristics have been used and compared in the SAT solver Riss¹¹.

46 2 ModSAT: A Modularized SAT solver

47 2.1 Overview of ModSAT

48 In this section, details of the proposed modularized SAT solver are provided. ModSAT obeys the basic CDCL framework¹
49 in Algorithm 1, which usually initiates with an empty set of partial assignments (line 2). The Unit Propagation (UP), also
50 called Boolean Constraint Propagation, assigns values to variable in clauses which has only one variable. UP can always
51 make a clause satisfied, and this operation will repeat until no more UP is possible (line 4). If no conflicts are detected in \mathcal{X}

Algorithm 1: CDCL Framework

```
1 Input: A CNF  $F$  of SAT instance;
2 Initialization: decision level  $d \leftarrow 0$ , current assignment of variables  $\mathcal{X} \leftarrow \emptyset$ ;
3 while True do
4    $\mathcal{X} \leftarrow \text{Unit Propagation}(F, \mathcal{X})$ ;
5   if Conflicts are detected in  $\mathcal{X}$  then
6     if  $d == 0$  then
7       return UNSAT
8     else
9        $\mathcal{C}_{\text{conflict}}, d_{\text{backtrack}} \leftarrow \text{Analyze Conflict}(F, \mathcal{X})$  ;
10       $C_{\text{learned}} \leftarrow \text{Learn Clause}(\mathcal{C}_{\text{conflict}}, \mathcal{X})$  ;
11       $F \leftarrow F \wedge C_{\text{learned}}$ ;
12       $\mathcal{X} \leftarrow \text{Backtrack}(\mathcal{X}, d_{\text{backtrack}})$ ;
13       $d \leftarrow d_{\text{backtrack}}$ ;
14   else
15     if All variables are assigned with a value then
16       return SAT
17     else
18        $d \leftarrow d + 1$ ;
19        $\mathcal{X} \leftarrow \text{Make Decision}(\mathcal{X})$ 
```

52 during the Decision Detection phase, the algorithm will select a variable to assign a value (line 19). This Make Decision step
53 usually contains a few heuristics. When conflicts are detected (line 5), Analyze Conflict will identify the conflicted clauses
54 (line 9), and a newly learnt clause will be derived based on the clause (i.e., the current partial assignment) of conflicts (line
55 10). Afterwards, the algorithm backtracks to an earlier decision level (lines 12-13). Once all variables are assigned and no
56 conflict is detected, the algorithm obtains a satisfied assignment (line 16); otherwise, detecting conflicts at the decision level 0
57 indicates that the given CNF is unsatisfiable (line 7).

58 Various heuristics have been proposed in ModSAT to enhance its performance. For instance, *reduce* heuristics identify and
59 remove the learnt clauses by controlling the size of the tracking list. In addition, *bump var* heuristics are usually incorporated
60 in the Analyze Conflict step, which affect the choice of variables in the Make Decision function. While the order of choosing
61 variables can determine the search path of branching, *rephase* heuristics can control the polarity in variables to be selected.
62 Also, *restart* heuristics may abandon the current search path, allowing algorithms to explore possibly easier search regions.

63 More precisely, in ModSAT, we have defined *seven* functions which are independently implemented: **restart function**,
64 which manages restart heuristics; **restart condition**, which determines when to execute restart; **reduce condition**, which
65 determines when to reduce; **rephase function**, which manages rephase heuristics; **rephase condition**, which determines
66 when to rephase; **bump var activity**, which governs the order of variables selection in Make Decision; and **bump cla activity**
67 that govern the order of clauses being removed during reduce. These functions collectively improve the solver's ability to
68 handle large and complex SAT instances by balancing exploration, exploitation, and resource management. The combination
69 of conflict-driven learning, backtracking, and heuristic enhancements makes ModSAT an efficient and robust SAT solver.

70 2.2 Three principles behind ModSAT

71 As already mentioned in the paper, we follow three principles when developing ModSAT to ensure it is LLM-friendly:

- 72 • **Maintain functions simple and focus.** The function optimized by LLMs should be simple and explicit, unlike common
73 implementation in complex solvers.
- 74 • **Utilize class variables for shared information.** Local variables should be declared as class member variables to give
75 LLMs access to them.
- 76 • **Proactively prevent bugs during heuristics discovery.** The bugs written by LLMs should be fixed proactively, so that
77 the solver could compile correctly with the same heuristics, which helps LLMs to generate more diverse correct codes.

78 The following example demonstrates the application of the principle *Maintain functions simple and focus*, with the original
79 function given in Figure 1. Instead of modifying the whole search function by LLMs, we modularize it into three distinct

Algorithm 2: ModSAT

```
1 Input: A CNF  $F$  of SAT instance;
2 Initialization: decision level  $d \leftarrow 0$ , current assignment of variables  $\mathcal{X} \leftarrow \emptyset$ ;
3 while True do
4    $\mathcal{X} \leftarrow \text{Unit Propagation}(F, \mathcal{X})$ ;
5   if Conflicts are detected in  $\mathcal{X}$  then
6     if  $d == 0$  then
7       | return UNSAT
8     else
9       |  $\mathcal{C}_{\text{conflict}}, d_{\text{backtrack}} \leftarrow \text{Analyze Conflict}(F, \mathcal{X})$  ;
10      |  $C_{\text{learned}} \leftarrow \text{Learn Clause}(\mathcal{C}_{\text{conflict}}, \mathcal{X})$  ;
11      |  $F \leftarrow F \wedge C_{\text{learned}}$ ;
12      |  $\mathcal{X} \leftarrow \text{Backtrack}(\mathcal{X}, d_{\text{backtrack}})$ ;
13      |  $d \leftarrow d_{\text{backtrack}}$ ;
14    if Restart condition then
15      |  $d \leftarrow \text{Restart}(d)$ 
16    else
17      | continue
18    if Rephase condition then
19      |  $\mathcal{X} \leftarrow \text{Rephase}(\mathcal{X})$ 
20    else
21      | continue
22    if Reduce condition then
23      |  $C_{\text{learned}} \leftarrow \text{Reduce}(C_{\text{learned}})$ 
24    else
25      | continue
26  else
27    if All variables are assigned with a value then
28      | return SAT
29    else
30      |  $d \leftarrow d + 1$ ;
31      |  $\mathcal{X} \leftarrow \text{Make Decision}(\mathcal{X})$ 
```

functions by isolating the components that significantly impact the performance and are suitable for LLM to modify, see Figure 2. The decomposition also enables LLMs to focus on refining one kind of heuristic at a time, thereby enhancing the code generation capability.

For the principle *use Class Variables for Shared information*, we move variables that may reside in local scope into class members to ensure that LLMs can access this shared information, see Figure 3.

Figures 4 and 5 provide two example on how to *prevent bugs during heuristics discovery proactively*. The first one is to include some extra packages to prevent LLMs from implementing extra functions, while the second one is to overload common functions to prevent LLMs from incorrectly implementing simple functions because of a misunderstanding of data structures.

3 Automatic Prompt Optimization

To tackle the problem of prohibitive execution time (5000s timeout per instance) and with large performance variance across different datasets, we propose an unsupervised automatic prompt optimization method using Shannon entropy as the evaluation metric. This approach eliminates the dependency on explicit labels while maintaining the adaptability to the evolution of LLMs, particularly suitable for compute-intensive optimization tasks with ambiguous success criteria.

The basic prompt template follows the instructions in OpenAI framework docs¹², which obeys the following format:

- Define the **Role** of an agent as a solver expert who needs to assess and improve the heuristics function in the SAT solver.

Maintain functions simple: Original function to modify

```
1 lbool Solver::search(int nof_conflicts){
2     // if there is a conflict
3     .....
4     // if there is no conflict
5     if ((lbd_queue_size == 50 && 0.8 * fast_lbd_sum / lbd_queue_size > slow_lbd_sum / conflicts) || !withinBudget())
6         restart_function();
7
8     // Simplify the set of problem clauses:
9     if (decisionLevel() == 0 && !simplify())
10        return l_False;
11
12    // Reduce the set of learnt clauses:
13    if (learnts.size()-nAssigns() >= max_learnts)
14        reduceDB();
15    if (rephase_condition())
16        rephase_function();
17
18    Lit next = lit_Undef;
19    while (decisionLevel() < assumptions.size()){
20        // Perform user provided assumption:
21        Lit p = assumptions[decisionLevel()];
22        if (value(p) == l_True){
23            // Dummy decision level:
24            newDecisionLevel();
25        }else if (value(p) == l_False){
26            analyzeFinal(~p, conflict);
27            return l_False;
28        }else{
29            next = p;
30            break;
31        }
32    }
33
34    if (next == lit_Undef){
35        // New variable decision:
36        decisions++;
37        next = pickBranchLit();
38
39        if (next == lit_Undef)
40            // Model found:
41            return l_True;
42    }
43    newDecisionLevel();
44    uncheckedEnqueue(next);
45}
```

Figure 1. Illustration of principle *Maintain functions simple and focus*.

- Clearly state the **Goal**, such as providing optimization suggestions, writing code, or feedback.
- Enhance the agents' capabilities by inserting optional **Tips** that guide them to avoid common mistakes during code generation. Additionally, through this flexible interface, agents can effectively utilize external codes and results and can be instructed to specify the types of modification directions such as changing parameters, modifying heuristics, or adding new heuristics.

Maintain functions simple: modularized function to modify

```
1 // original function which has been modularized
2 lbool Solver::search(int nof_conflicts){
3     .....
4
5     if (restart_condition())
6         restart_function();
7     if (reduce_condition())
8         reduceDB();
9
10    if (rephase_condition())
11        repause_function();
12
13    .....
14}
15
16 // functions to modify
17 bool Solver::rephase_condition() {
18     if (rephases >= repause_limit)
19         return true;
20     else
21         return false;
22}
23
24 bool Solver::reduce_condition() {
25     if (rephases >= repause_limit)
26         return true;
27     else
28         return false;
29}
30
31 bool Solver::restart_condition(){
32     if ((lbd_queue_size == 50 && 0.8 * fast_lbd_sum / lbd_queue_size > slow_lbd_sum / conflicts) || !withinBudget())
33         return true;
34     else
35         return false;
36}
```

Figure 2. Illustration of principle *Maintain functions simple and focus*.

- Key code of SAT solver is appended at the end of each prompt to ensure all agents are in the same context. Note that the key code includes member parameters in cpp class of the solver (LLMs may utilize), along with the main loop function, and all corresponding functions LLMs may need to understand (e.g. such as the functions call the target optimized function).

The previous three components, role, goal and tips can be automatically optimized by LLMs. In each iteration, we randomly select one component to optimize by LLMs, compute the correctness and diversity (entropy) of the generated codes, and update the component in prompt with better performance in both correctness and diversity. This loop will iterate until the stopping criteria is met, see Algorithm 3 for details. An empirical evaluation of the original and optimized prompts is presented in Figure 6. It is clear that optimized prompt achieves better performance than the original one.

4 Presearch strategy

To overcome this combinatorial explosion when searching over all candidate functions, we propose a novel two-phase optimization strategy:

Add class member variables

```
1 // LBD heuristics
2 int lbd_queue[500],      // circled queue saved the recent 500 LBDs.
3     lbd_queue_size,      // The number of LBDs in this queue
4     lbd_queue_pos;
5 double fast_lbd_sum, slow_lbd_sum;
6
7 // rephase heuristics
8 int rephases, rephase_limit, rephase_count, threshold;
9 double last_rephase_progress;
10
11 // restart heuristics
12 int curr_restarts;
13 double last_restart_progress;
```

Figure 3. Illustration of principle *Utilize class variables for shared information*

Prevent bugs: add more packages

```
1 #include <math.h>
2 #include <unordered_set>
3 #include <algorithm>
4 using namespace std;
```

Figure 4. Illustration of principle *Proactively prevent bugs during heuristics discovery*

- Presearch function candidate: Conduct small-scale preliminary testing to identify and eliminate functions that consistently degrade the performance.
- Refined evolutionary search: Execute a focused $(1 + \lambda)$ Evolutionary Algorithm (EA) search only on the high-impact candidate functions identified in the last phase.

A detailed description of this strategy can be found in Algorithm 4. For the phase of presearch function candidate, we construct a compact and representative evaluation subset comprising the 50% problem instances from the original dataset. This subset can capture the essential solver behavior while minimizing the evaluation cost. Each candidate function is evaluated separately on the sub-dataset for its impact on the solver's PAR-2 score (where lower values indicate better performance), which measures the standalone effect of adding the function to a baseline solver. Then functions that degrade the PAR-2 score are identified and pruned. For each dataset, the pruned functions will not be further considered for evolutionary search. This phase typically retains a small set of high-impact functions (e.g., 4 functions) that consistently show a positive or neutral effect on PAR-2 in the preliminary tests. For the evolutionary optimization phase, we use the significantly refined function set (e.g., about 4 functions) to execute a standard $(1 + \lambda)$ Evolutionary Algorithm on the full target dataset to find optimal combinations.

It is also helpful to investigate the contribution of each function to the final results. To this end, we have calculated the number of times that a function contributes to the final performance improvement in each experiment, see Figure 8, where each subfigure shows the results for one dataset. It can be observed that almost all functions we select could contribute substantially to the final performance,

5 Heuristics discovery

The details of heuristics discovery in AutoModSAT are presented in Figure 9, which contains three agents, LLMs coder, LLMs evaluator, LLMs repairer. Since the LLM coder agent is already optimized in automatic prompt optimization, here we give a brief discussion of the LLMs evaluator and LLMs repairer.

In our experiments, we find that sometimes LLMs coder would generate synonymous code compared with original one, even though we have optimized the prompt. This will lead to redundant iterations in the evolutionary search. Thus, we

Prevent Bugs: overloading functions

```
1 #include <type_traits>
2
3 template <typename T1, typename T2>
4 typename std::common_type<T1, T2>::type max(T1 a, T2 b) {
5     static_assert(std::is_integral<T1>::value && std::is_integral<T2>::value,
6                 "max: Both types must be integers (int or long int)");
7     return (a < b) ? b : a;
8 }
9
10 template <typename T1, typename T2>
11 typename std::common_type<T1, T2>::type min(T1 a, T2 b) {
12     static_assert(std::is_integral<T1>::value && std::is_integral<T2>::value,
13                 "min: Both types must be integers (int or long int)");
14     return (a < b) ? a : b;
15 }
16
17 template <typename T1, typename T2>
18 typename std::common_type<T1, T2>::type max(T1 a, T2 b) {
19     static_assert(std::is_floating_point<T1>::value && std::is_floating_point<T2>::value,
20                 "max: Both types must be floating-point (float or double)");
21     return (a < b) ? b : a;
22 }
23
24 template <typename T1, typename T2>
25 typename std::common_type<T1, T2>::type min(T1 a, T2 b) {
26     static_assert(std::is_floating_point<T1>::value && std::is_floating_point<T2>::value,
27                 "min: Both types must be floating-point (float or double)");
28     return (a < b) ? a : b;
29 }
```

Figure 5. Illustration of principle *Proactively prevent bugs during heuristics discovery*.

136 develop an LLMs evaluator agent, which identifies whether the generated code is synonymous. If the code is synonymous,
137 LLMs coder needs to regenerate a code; otherwise, the code will be sent to compile.

138 In addition, LLMs sometimes make mistakes, such as missing parentheses or utilizing incorrect data types. To address this
139 issue, an LLM repairer agent is introduced. LLM repairer analyzes the LLM-generated code and the corresponding errors, and
140 then seeks to fix the bug. This agent can successfully resolves some common errors. While there are some errors remaining
141 uncorrectable, LLM Coder will re-generate the code.

142 6 Experimental details

143 6.1 Dataset description

144 Here we give more details about the datasets, including 7 datasets selected from SAT Competition 2023 and 2024, 3 generated
145 ones by Picat, and another one from an industrial EDA scenario.

- 146 • **Argumentation problem** involves finding acceptable sets of arguments in a directed graph where attacks between
147 arguments are represented by edges.
- 148 • **Social Golfer problem** is a combinatorial problem that aims to assign golfers to groups over several weeks, ensuring
149 no two golfers play in the same group more than once.
- 150 • **Hashtable Safety problem** focuses on verifying the correctness of operations in a hash table to avoid collisions and
151 ensure the integrity of the structure.

Algorithm 3: Automatic Prompt Optimization

Input: initial prompt template P , solver codebase S , max_iterations i , prompt optimized part $R = \{Role, Goal, Tips\}$

Output: optimized prompt template P^*

```

1  $i \leftarrow 10$                                      // prompt optimization iterations
2  $j \leftarrow 20$                                      // number of code generation in each iteration
3 while  $i \geq 0$  do
4   select prompt part  $r$  from  $R$  uniformly at random
5    $P' \leftarrow \text{refine\_prompt}(r, P)$                 // LLMs refine the prompt template
6   while  $j \geq 0$  do
7     generated code  $c \leftarrow \text{call\_llm}(current\_prompt)$ 
8     compilation error  $e \leftarrow \text{compile}(c, S)$ 
9     if  $\neg e$  then
10      |  $C \leftarrow C \cup c$ 
11    else
12      |  $\text{execute\_code}(corrected\_code)$ 
13    end
14     $j \leftarrow j - 1$ 
15  end
16  diversity  $d_i \leftarrow \text{compute\_code\_diversity}(C)$ 
17  success rate  $s_i \leftarrow \text{compute\_code\_success}(C)$ 
18  if  $d_i > d$  and  $s_i > \text{threshold}$  then
19    |  $d \leftarrow d_i$ 
20    |  $P \leftarrow \text{update\_prompt}(S, P')$ 
21  else
22  end
23   $i \leftarrow i - 1$ 
24 end

```

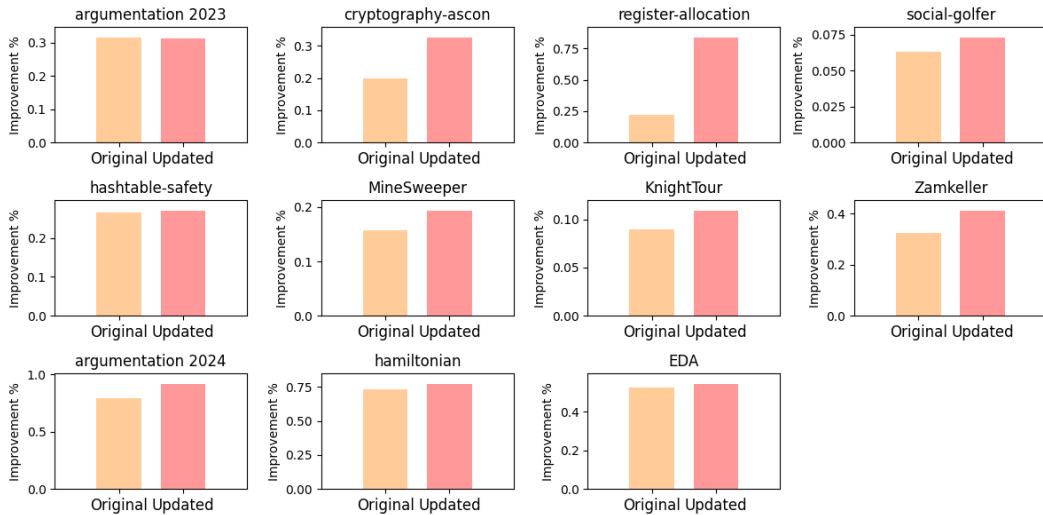


Figure 6. Comparison over Different Prompts. This figure shows the performance of AutoModSAT using different prompts. The vertical axis shows the PAR-2 improvement ratio compared to the original solver, while the two coordinates in the horizontal axis correspond to the original prompt (Original) and optimized prompt (Updated), respectively.

152

- **Register Allocation problem** is a problem that arises in compiler optimization, where the goal is to assign a limited number of CPU registers to variables in a program.

153

- **Cryptography-Ascon problem** is a lightweight cryptographic algorithm challenge focused on implementing and ver-

Original Prompt

(Role) You are a SAT solver researcher trying to rewrite the {{ func_name }} function(s).

(Goal) Your goal is to improve the SAT solver by rewriting the {{ func_name }} function(s), after reading and understanding the <key code> of SAT solver below.

(Tips) Tips:

- 1) Your rewritten function code must start with "/* start function name" and end with "/* end function name"
- 2) Your rewritten function(s) code must be different from original code, not just rewrite code synonymous!
- 3) You are not allowed to create your own new function(s) in the rewritten function(s). You are not allowed to create your own new global variables, but you can use the global variables existing in the <key code>.
- 4) Make sure the rewritten function(s) code can be executed correctly.

<key code> of SAT solver is:

```
 {{ replace_key_code }}
```

Updated Prompt

(Role) You are a SAT solver researcher trying to improve the {{ func_name }} function.

(Goal) Objective:

Your goal is to improve the SAT solver by rewriting the {{ func_name }} function.

Instructions:

1. Carefully read and comprehend the <key code> of the SAT solver provided below.
2. Analyze potential improvements and devise a strategy for optimizing the heuristics of function.
3. Deliver your improved function(s) with the following format:
 - Begin with: '/* start function name'
 - End with: '/* end function name'

(Tips) Tips:

1. Ensure that your rewritten function(s) are substantially different from the original, beyond mere synonym replacements.
2. You may utilize existing global variables from the <key code>, but refrain from introducing new global variables.
3. Verify that the rewritten function(s) execute correctly.

Take a deep breath and think it step by step.

<key code> of SAT solver is:

```
 """ {{ replace_key_code }} """ ...
```

Figure 7. Comparison of original prompt and optimized prompt

155 ifying the Ascon cipher, a NIST-standardized authenticated encryption scheme, for resource-constrained IoT devices,
156 balancing security against differential attacks with minimal computational overhead.

157

- 158 • **Hamiltonian problem** is a graph theory problem that involves determining whether a given graph contains a Hamiltonian cycle, a closed loop visiting each vertex exactly once, which is NP-complete and often applied to route optimization or circuit design verification.
- 160 • **MineSweeper problem** is derived from the classic MineSweeper game, where the objective is to determine the placement of hidden mines on a grid based on numerical clues.
- 162 • **Langford problem** is a combinatorial mathematics problem that involves finding a specific permutation of the sequence

Algorithm 4: PreSearch Strategy in AutoModSAT

```

1 Input: Datasets  $P$ , modularized SAT solver with seven functions  $\{h_1, \dots, h_7\}$ , prompt template, baseline functions
    $\{b_1, \dots, b_7\}$ 
2 Phase 1: PreSearch Function Candidate  $P_{\text{compact}} \leftarrow 50\%$  representative instances from  $P$ 
3  $R \leftarrow \emptyset$ ; // Retained function set
4 for each function  $h_i \in \{h_1, \dots, h_7\}$  do
5    $A_{\text{test}} \leftarrow$  build solver replacing  $h_i$  with baseline  $b_i$ 
6    $f_i, s_i \leftarrow \text{evaluate}(A_{\text{test}}, P_{\text{compact}})$ ; // Get PAR-2 metric
7    $F \leftarrow \text{sort}(f_1, \dots, f_7)$ ; // Sort PAR-2 metric in different functions
8   Get top 4 function index  $R$  from  $F$ 
9 Phase 2: Evolutionary Algorithm Optimization  $A \leftarrow$  solver with functions:  $(\forall i \in R : h_i) \cup (\forall i \notin R : b_i)$ 
10  $f^* \leftarrow \text{evaluate}(A, P)$ ; // Full dataset evaluation
11  $\text{evalBudget} \leftarrow 50$ ; // Maximum evaluations
12 while  $\text{evalBudget} > 0$  do
13    $M \leftarrow \emptyset$ 
14    $\ell \sim \text{Bin}(|R|, \frac{1}{|R|})$ ; // Sample modification count
15   chosen  $\ell$  distinct values  $M \leftarrow \{m_0, \dots, m_\ell\}$  from  $R$  uniformly at random;
16   Generate new functions  $\{h'_m\}_{m \in M}$  via LLM using  $A$  and  $M$ 
17    $A' \leftarrow \text{update } A \text{ with } \{h'_m\}_{m \in M}$ 
18    $f(A') \leftarrow \text{evaluate}(A', P)$ 
19   if  $f(A') \leq f^*$  then
20      $A \leftarrow A'$ 
21      $f^* \leftarrow f(A')$ 
22    $\text{evalBudget} \leftarrow \text{evalBudget} - 1$ 

```

163 1, 1, 2, 2, ..., n, n where the two copies of each number k are exactly k units apart.

164 • **KnightTour problem** aims to find a path for a knight on a chessboard that visits every square exactly once, with
165 possible extensions to different board sizes and types.

166 • **Zamkeller problem** involves finding a permutation of integers from 1 to n that maximizes the number of differential
167 alternations in subsequences divisible by integers from 1 to k , where $(1 < k < n)$.

168 • **EDA problem** involves formally proving whether two design specifications are functionally equivalent, which is one of
169 the most essential techniques in Electronic Design Automation and digital IC design. It has a wide range of applications,
170 such as functional equivalent logic removal, sequential equivalence checking, circuit-based method for symmetries
171 detection, engineering change orders, among others.

172 For the generated instances using Picat¹³, we adopt the settings in Chapters 2 and 3 of the book by NengFa Zhou¹³ and
173 conduct grid sampling within a parameter space. Specifically, the parameter space $\Theta = \{\theta_1, \theta_2, \dots\}$, $\theta_i^L \leq \theta_i \leq \theta_i^U$ is defined to
174 ensure that three baseline solvers including EasySAT, MiniSat and Kissat, can obtain a solution within proper cputime range,
175 i.e., [1s, 5000s]. The we apply a space Θ' to generate the dataset by enlarging the upper bound of Θ , e.g., $\theta_i'^U = \theta_i^U * 1.2$,
176 such that we can test whether AutoModSAT can solve the instances where the baseline solvers cannot.

177 • **MineSweeper problem**
178 Parameter Θ : $\{m, n, k, p\}$
179 Parameter Space: $[500, 1, 600] \times [400, 3200] \times [72, 689, 1, 572, 118] \times [0.32, 0.38]$
180 Notes: m, n represent the grid size of the Minesweeper game. k is the total number of mines. p is the probability that a
181 given cell contains a mine (range: 0.32 to 0.38).

182 • **KnightTour problem**
183 Parameter Θ : $\{k\}$
184 Parameter Space: $[12, 75]$
185 Notes: A $k \times k$ chessboard where a knight's tour is attempted, covering all squares and returning to the start point. (A
186 solution is not possible for odd-sized boards, i.e., they are unsatisfiable.)

Function Contribution

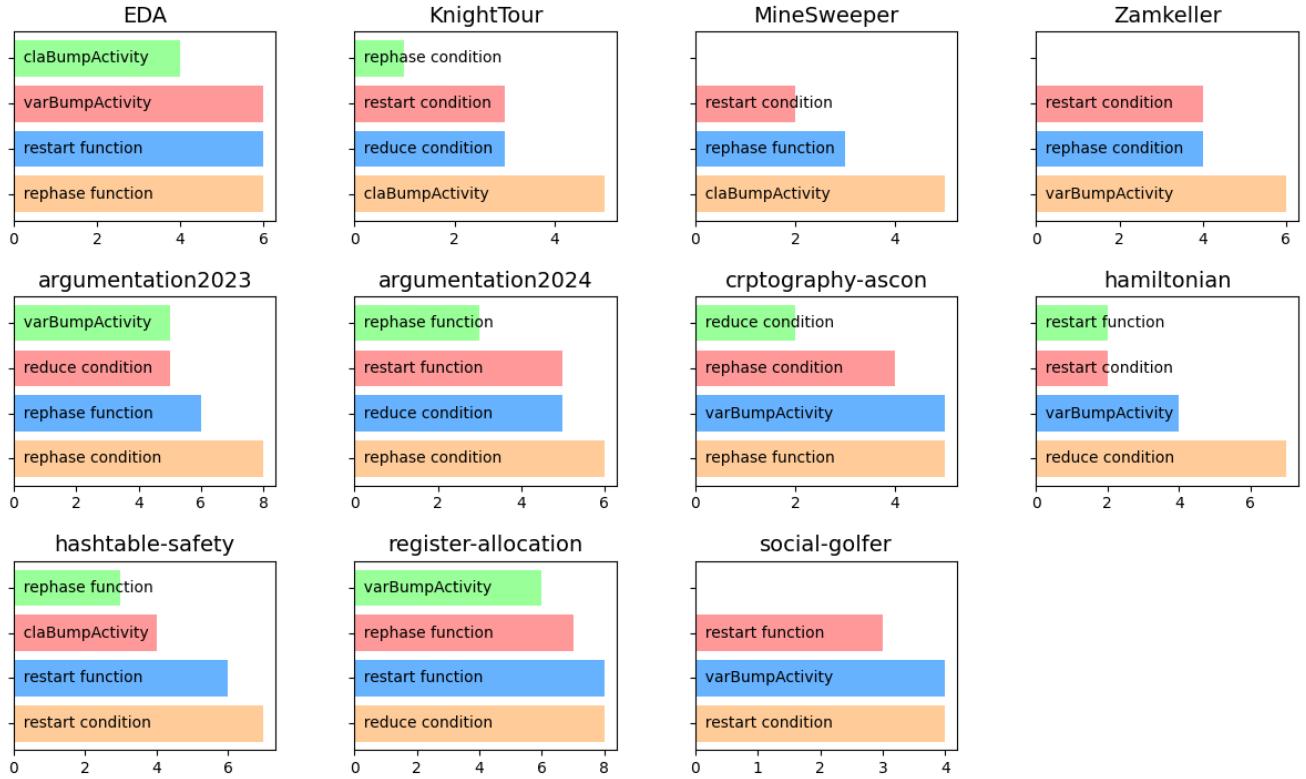


Figure 8. Function Contribution: where the vertical axis lists different function candidates and the horizontal axis denotes the frequency of contributions.

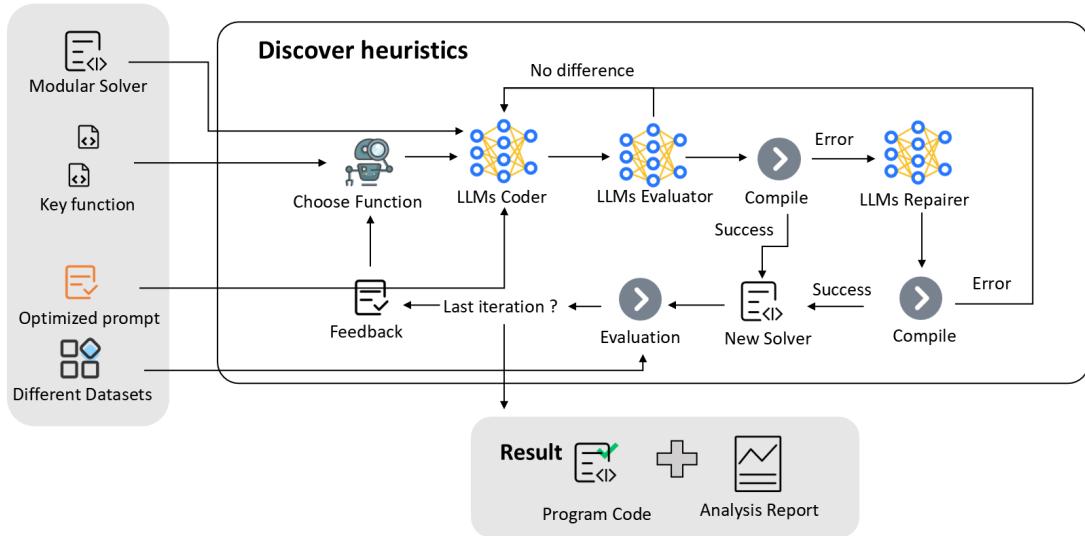


Figure 9. Details about heuristics discovery

187 • **Zamkeller problem:**

188 Parameter Θ : $\{k, n\}$

189 Parameter Space: $[3, 34] \times [25, 100]$

190 Notes: k represents the total sequence length, and n represents the subsequence length. For all subsequences of length
191 k , the goal is to change them into the minimum number of distinct sequences.

192 The detailed configuration of the training dataset and the function candidates in heuristics discovery are presented in
193 Table 1.

194 **6.2 Evaluation Metric**

195 We consider two specific metrics for evaluation of a SAT solver: (1) the number of SAT instances solved within the given
196 time-out bound, and (2) the Penalized Average Runtime with a factor of 2 score (PAR-2). Both metrics are commonly used in
197 the SAT Competitions.

198 Consider a dataset of n instances. Let t_i be the runtime of the SAT solver on the instance i . The PAR-2 score is formally
199 defined as:

$$\text{PAR-2} = \frac{1}{n} \sum_{i=1}^n \tau_i \quad \text{where: } \tau_i = \begin{cases} t_i, & \text{if } t_i \leq \mathcal{T} \\ 2\mathcal{T}, & \text{if } t_i > \mathcal{T} \text{ or the solver fails to return a result,} \end{cases}$$

200 where \mathcal{T} is the predefined time-out bound. For example, consider a benchmark dataset with three instances and a time-out
201 bound $\mathcal{T} = 100$ seconds. The runtimes (in seconds) for the three instances are: $t_1 = 80$ for instance 1, $t_2 = 120$ for instance
202 2, the solver fails to return a result for instance 3. Then

- 203 • since $t_1 \leq \mathcal{T}$, we have $\tau_1 = 80$;
- 204 • since $t_2 > \mathcal{T}$, we have $\tau_2 = 200$ (penalized);
- 205 • since the solver fails for instance 3, we have $\tau_3 = 200$ (penalized).

206 Therefore, The PAR-2 score is given by

$$\text{PAR-2} = \frac{1}{3} (80 + 200 + 200) = \frac{480}{3} = 160.$$

Table 1. Configuration of training set, where the indices 1 to 7 represent the following function candidates in order:
rephase_condition, rephase_function, reduce_condition, restart_condition, restart_function, varBumpActivity,
claBumpActivity.

Dataset	Training Timeout	Function candidate
cryptography-ascon	800	1, 2, 3, 6
register-allocation	5000	2, 3, 5, 6
social-golfer	2000	1, 4, 5, 6
hashtable-safety	500	2, 4, 5, 7
argumentation 2023	2000	1, 2, 3, 6
argumentation 2024	2000	1, 2, 3, 5
hamiltonian	800	3, 4, 5, 6
MineSweeper	500	2, 4, 3, 7
KnightTour	2000	1, 3, 4, 7
Zamkeller	2000	1, 3, 4, 6
EDA	800	2, 5, 6, 7

207 **6.3 Search Space for Parameter Tuning**

208 In this paper, we adopt SMAC3 to optimize the baseline SAT solver parameters across different datasets. The search space for
209 each solver, including the parameter name, type, description and range, are presented in Tables 2, 3, and 4.

Table 2. ModSAT configuration parameters

Parameter	Type	Description	Search Space
var-decay	double	Variable activity decay factor	(0, 1)
cla-decay	double	Clause activity decay factor	(0, 1)
rnd-freq	double	Frequency for random variable selection	[0, 1]
rnd-init	bool	Randomize initial activities	{true, false}
rfirst	int	Base restart interval	[1, 1e4]
rinc	double	Restart interval increase factor	(1.5, 4)
gc-frac	double	Wasted memory fraction triggering garbage collection	(0, 1)
min-learnts	int	Minimum learnt clause limit	[0, 1e6]

Table 3. Kissat configuration parameters

Parameter	Type	Description	Search Space
chrono	bool	Enable chronological backtracking	{0, 1}
eliminate	bool	Enable variable elimination	{0, 1}
forcephase	bool	Force initial phase assignment	{0, 1}
minimize	bool	Enable clause minimization	{0, 1}
phase	bool	Set initial decision phase	{0, 1}
phasesaving	bool	Enable phase saving during restarts	{0, 1}
probe	bool	Enable failed literal probing	{0, 1}
reduceint	int	Conflict interval for clause DB reduction	{10 ¹ , 10 ² , 10 ³ , 10 ⁴ , 10 ⁵ }
rephaseint	int	Conflict interval for phase resetting	{10 ¹ , 10 ² , 10 ³ , 10 ⁴ , 10 ⁵ }
restartint	int	Base restart interval (conflicts)	{1, 10 ² , 10 ³ , 10 ⁴ }
restartmargin	int	Rapid restart margin threshold	{0, 5, 10, 15, 20, 25}
simplify	bool	Enable periodic simplification	{0, 1}
stable	int	Search stability mode (0=focused, 1=stable, 2=switching)	{0, 1, 2}
target	int	Target phase selection strategy (0=negative, 1=positive, 2=best)	{0, 1, 2}
tier1	int	Tier 1 glue limit for learned clauses	{2, 3, 4, 5}
tier2	int	Tier 2 glue limit for learned clauses	{6, 7, 8, 9, 10, 20, 50}

210 7 More Examples of Discovered Heuristics

211 In this section, to demonstrate LLMs’ ability of generating effective heuristics, we provide more contrastive examples (one
212 for each function candidate), along with the explanations for the changes.

213 Figure 10 provides an example for the updated claBumpActivity function, which introduces two key enhancements in con-
214 trast to the original implementation. First, during activity rescaling, it enforces a minimum activity threshold ($min_activity =$
215 $1e - 20$) to prevent clauses from becoming numerically insignificant after scaling. This preserves the relevance of histori-
216 cally important clauses and avoids premature elimination from the learning process. Second, it incorporates dynamic decay
217 adjustment based on recent conflict rates: when conflicts exceed 1000 and the LBD queue surpasses 50 entries, cla_inc is
218 scaled down proportionally to the conflict density (with a floor of 0.8). This adaptively moderates activity growth during high-
219 conflict phases, prioritizing recent impactful clauses while maintaining stability. Together, these refinements yield a more
220 balanced clause management strategy preventing underutilization of valuable learned clauses while dynamically optimizing
221 activity decay for solver efficiency.

222 Figure 11 provides an example of the updated varBumpActivity function, which introduces three key improvements over
223 the original. First, it scales the increment by $(1.0 + 0.1 * decisionLevel())$, prioritizing variables involved in recent decisions
224 to accelerate conflict-driven learning. Second, the rescaling mechanism uses a larger threshold (1e100) and finer scale factor
225 (1e-100), while preserving variable relevance by enforcing a minimum activity floor (1e-100) to maintain relative ordering and
226 prevent premature underflow. Third, it optimizes heap management through conditional updates: dynamically adjusting the
227 variable’s heap position only when its activity exceeds the current maximum, or inserting undefined decision variables lazily.

Table 4. CaDiCaL configuration parameters

Parameter	Type	Description	Search Space
chrono	int	Chronological backtracking mode (0: none, 1: limited, 2: always)	$\{0, 1, 2\}$
elim	bool	Enables variable elimination during simplification	$\{0, 1\}$
forcephase	bool	Forces phase saving for decision variables	$\{0, 1\}$
minimize	bool	Enables clause minimization during conflict analysis	$\{0, 1\}$
phase	bool	Initial decision phase assignment (0: negative, 1: positive)	$\{0, 1\}$
probe	bool	Enables probing (failed literal detection)	$\{0, 1\}$
reduceint	int	Conflict interval for clause database reduction	$\{10^2, 10^3, 10^4, 10^5\}$
rephaseint	int	Conflict interval for resetting variable phases	$\{10^1, 10^2, 10^3, 10^4, 10^5\}$
restartint	int	Base restart interval (conflicts between restarts)	$\{2, 10^2, 10^3, 10^4\}$
restartmargin	int	Restart margin percentage (Luby sequence scaling)	$\{0, 5, 10, 15, 20, 25\}$
stabilize	bool	Stabilizes search by limiting activity updates	$\{0, 1\}$
target	int	Search target (0: SAT, 1: UNSAT, 2: balanced)	$\{0, 1, 2\}$

228 These enhancements collectively improve search guidance, reduce floating-point stability issues, and minimize unnecessary
229 data structure operations.

230 Figure 12 provides an example of the updated restart_condition function significantly improves upon the original by
231 replacing its static threshold approach with a dynamic, performance-driven restart strategy that adapts to real-time solver
232 behavior. Instead of relying on fixed queue sizes and hardcoded multipliers, the new version intelligently calculates restart
233 thresholds using multiple runtime metrics: it combines average LBD (measuring clause quality) with conflict rates (tracking
234 solver progress) to dynamically adjust restart timing based on problem difficulty. Crucially, it introduces a progress-sensitive
235 mechanism that aggressively lowers thresholds when stagnation is detected (progressEstimate changes < 0.01), enabling
236 proactive recovery from plateaus a capability absent in the original. This multi-factor approach yields more precise restart
237 decisions, reduces wasteful recomputations, and enhances solver adaptability across diverse SAT instances while maintaining
238 robustness during initialization through default thresholds.

239 Figure 13 provide an example of the updated restart_function, which introduces significant improvements over the original
240 implementation by incorporating adaptive restart strategies based on real-time solver performance metrics. Unlike the original
241 version, which always resets to decision level 0 (a full restart), the enhanced function dynamically calculates two exponential
242 moving averages of conflict difficulty (fast_avg and slow_avg) using Literal Block Distance (LBD) scores. By analyzing the
243 ratio between these averages, it intelligently selects one of three restart depths: full restart (level 0) for deteriorating conflict
244 quality, partial restart (mid-level) for moderately harder conflicts, or minimal restart (current level -1) for stable conditions.
245 This adaptability preserves useful learned clauses during partial/minimal restarts, reducing redundant recomputation. Ad-
246 ditionally, periodic clause database reduction (every 16 restarts) curbs memory growth, while rebuilding the variable order
247 heap ensures branching decisions reflect updated activity scores. Collectively, these optimizations balance exploration and
248 exploitation, enhancing solver efficiency through context-aware restarts and resource management.

249 Figure 14 provides and examples of the updated rephase_condition function, which introduces an adaptive rephasing mech-
250 anism that significantly enhances the original static threshold approach. Unlike the prior version which solely relied on a fixed
251 rephase limit, the new implementation dynamically adjusts rephasing intervals based on real-time search progress and conflict
252 density. By calculating normalized progress through trail size changes and setting variable-driven thresholds (e.g., 2% of total
253 variables), it detects stagnation when progress falls below expectations and responds by reducing subsequent rephase inter-
254 vals exponentially. Conversely, substantial progress triggers gradual interval expansion. This self-tuning capability optimizes
255 computational efficiency: it minimizes unnecessary rephasing during productive search phases while aggressively countering
256 stagnation, thereby improving solution convergence without compromising robustness.

257 Figure 16 provides an example of the updated rephase_function, which introduces several key improvements over the
258 original implementation, enhancing adaptability and search efficiency. First, it implements dynamic rephase limit adjustment
259 by scaling rephase_limit based on progress measured through conflict resolution (conflictR). If progress occurs, the limit in-
260 creases by 50% to exploit productive phases more aggressively; otherwise, it decays by 10% (with a lower bound of 512) to
261 conserve resources during stagnation. This replaces the originals static increment ($+ = 8192$) and fixed decay (threshold $* = 0.9$),

enabling context-sensitive resource allocation. Second, the refined phase selection strategy uses weighted probabilities with four distinct policies: local-best phases (40%), global phase inversion (30%), randomized phases for low-activity variables (20%), and user-specified phases (10%). This replaces the originals rigid three-policy cascade, adding targeted randomization for less-active variables which helps escape local optima and reintroducing user phases for domain-specific guidance. Finally, adaptive threshold reset (threshold = $\text{trail.size()} * 0.8$) dynamically scales with the solvers state, replacing the fixed decay, while verbosity-controlled logging aids debugging. These changes collectively improve the solvers ability to balance exploration versus exploitation, mitigate stagnation, and leverage problem-specific knowledge.

Figure 17 provides an example of the updated reduce_condition function, which significantly enhances the original version through four key improvements that collectively the optimize memory management and solver adaptability. First, it retains the core check for absolute learnt clause limits ($\text{learnts.size()} \geq \text{max_learnts}$), ensuring baseline constraint adherence. Second, it introduces memory pressure awareness by triggering reduction when wasted clause memory exceeds 80% of the garbage collection threshold ($\text{ca.wasted()} > \text{ca.size()} * \text{garbage_frac} * 0.8$). This proactively mitigates memory bloat and improves cache efficiency. Third, a learnt-to-original clause ratio check ($\text{learnts.size()} > 2 * \text{nClauses}()$) prevents learnt clauses from disproportionately dominating the formula, maintaining balanced reasoning. Finally, a conflict-driven heuristic ($\text{conflictR} > 1000 \ \&\& \text{learnts.size()} > \text{max_learnts} * 0.8$) adapts to high-conflict phases by initiating earlier reduction, thus accelerating recovery from solver stagnation. These layered criteria synergistically boost robustness: they minimize redundant computation through memory-sensitive garbage collection, preserve clause quality via ratio controls, and dynamically respond to runtime behavior ultimately yielding faster, more memory-efficient SAT solving.

References

1. Marques-Silva, J. P. & Sakallah, K. A. Grasp: A search algorithm for propositional satisfiability. *IEEE Transactions on Comput.* **48**, 506–521 (1999).
2. Biere, A., Heule, M., van Maaren, H. & Walsh, T. Conflict-driven clause learning sat solvers. *Handb. Satisf. Front. Artif. Intell. Appl.* **4**, 131–153 (2009).
3. Liang, J. H., Ganesh, V., Poupart, P. & Czarnecki, K. Learning rate based branching heuristic for SAT solvers. In *Theory and Applications of Satisfiability Testing–SAT 2016: 19th International Conference*, 123–140 (2016).
4. Ramos, A., Van Der Tak, P. & Heule, M. J. Between restarts and backjumps. In *Theory and Applications of Satisfiability Testing–SAT 2011: 14th International Conference, SAT*, 216–229 (2011).
5. Luby, M., Sinclair, A. & Zuckerman, D. Optimal speedup of las vegas algorithms. *Inf. Process. Lett.* **47**, 173–180 (1993).
6. Audemard, G. & Simon, L. Glucose 2.1: Aggressive, but reactive, clause database management, dynamic restarts. In *Pragmatics of SAT 2012* (2012).
7. Biere, A. & Fröhlich, A. Evaluating cdcl restart schemes. In *Proceedings of Pragmatics of SAT*, 1–17 (2015).
8. Biere, A. Lingeling, plingeling, picosat and precosat at sat race 2010. *FMVReport Ser. Tech. Rep.* **10** (2010).
9. Jeroslow, R. G. & Wang, J. Solving propositional satisfiability problems. *Annals mathematics Artif. Intell.* **1**, 167–187 (1990).
10. Soos, M. Strangelight. In *Proceedings of SAT Competition 2013: Solver, Benchmark and Proof Checker Descriptions*, 89–90 (Department of Computer Science, University of Helsinki, 2013).
11. Balint, A., Belov, A., Järvisalo, M. & Sinz, C. Overview and analysis of the sat challenge 2012 solver competition. *Artif. Intell.* **223**, 120–155 (2015).
12. OpenAI. OpenAI API Documentation. <https://platform.openai.com/docs> (2023).
13. Zhou, N.-F., Kjellerstrand, H. & Fruhman, J. *Constraint solving and planning with Picat*, vol. 11 (Springer, 2015).

original claBumpActivity

```
1 inline void Solver::claBumpActivity (Clause& c) {
2     if ( (c.activity() += cla_inc) > 1e20 ) {
3         // Rescale:
4         for (int i = 0; i < learnts.size(); i++)
5             ca[learnts[i]].activity() *= 1e-20;
6         cla_inc *= 1e-20; }
7 }
```

discovered claBumpActivity

```
1 inline void Solver::claBumpActivity (Clause& c) {
2     // Adaptive scaling based on clause length and conflict participation
3     if ((c.activity() += cla_inc) > 1e20) {
4         // Maintain minimum activity while scaling to preserve relevance
5         double scale_factor = 1e-20;
6         double min_activity = 1e-20;
7         for (int i = 0; i < learnts.size(); i++) {
8             ca[learnts[i]].activity() *= scale_factor;
9             if (ca[learnts[i]].activity() < min_activity)
10                 ca[learnts[i]].activity() = min_activity;
11         }
12         cla_inc *= scale_factor;
13
14         // Stabilize scaling by maintaining lower bound on clause activity
15         if (cla_inc < min_activity)
16             cla_inc = min_activity;
17     }
18     // Dynamic decay adjustment based on recent conflict rate
19     if (conflicts > 1000 && lbd_queue_size > 50) {
20         double conflict_scale = 1.0 - 0.01 * (lbd_queue_size / 50.0);
21         cla_inc *= conflict_scale > 0.8 ? conflict_scale : 0.8;
22     }
23 }
```

Figure 10. Example of LLM discovered heuristics of claBumpActivity

original varBumpActivity

```
1 inline void Solver::varBumpActivity(Var v, double inc) {
2     activity[v] += inc;
3     // Rescale condition to prevent floating-point overflow earlier
4     if (activity[v] > 1e50) {
5         for (int i = 0; i < nVars(); i++)
6             activity[i] *= 1e-50;
7         var_inc *= 1e-50;
8     }
9 }
```

discoverd varBumpActivity

```
1 inline void Solver::varBumpActivity(Var v, double inc) {
2     // Scale the increment based on decision level to prioritize recent variables
3     double scaled_inc = inc * (1.0 + 0.1 * decisionLevel());
4     activity[v] += scaled_inc;
5     // More gradual rescaling that preserves relative ordering better
6     if (activity[v] > 1e100) {
7         double scale_factor = 1e-100;
8         for (int i = 0; i < nVars(); i++) {
9             activity[i] *= scale_factor;
10            // Prevent variables from becoming completely inactive
11            if (activity[i] < 1e-100) activity[i] = 1e-100;
12        }
13        var_inc *= scale_factor;
14    }
15    // Optimized heap update with early exit for common case
16    if (order_heap.inHeap(v)) {
17        if (activity[v] > activity[order_heap[0]]) {
18            order_heap.decrease(v);
19        }
20    } else if (decision[v] && value(v) == l_Undef) {
21        order_heap.insert(v);
22    }
23 }
```

Figure 11. Example of LLM discovered heuristics of varBumpActivity

original restart_condition

```
1 bool Solver::restart_condition(){
2     if ((lbd_queue_size == 50 && 0.8 * fast_lbd_sum / lbd_queue_size > slow_lbd_sum /
3         conflicts) || !withinBudget())
4         return true;
5     else
6         return false;
7 }
```

discoverd restart_condition

```
1 bool Solver::restart_condition() {
2     // Base restart condition on conflict count
3     if (conflicts <= 0) return false;
4
5     // Dynamic restart threshold based on recent solver performance
6     double restart_threshold;
7
8     if (lbd_queue_size > 0) {
9         // Use both average LBD and conflict rate for adaptive restart
10        double avg_lbd = fast_lbd_sum / lbd_queue_size;
11        double conflict_rate = (double)conflictR / (double)conflicts;
12
13        // Adjust threshold based on problem difficulty (higher LBD = harder problem)
14        restart_threshold = restart_first * (0.8 + 0.4 * avg_lbd) * (1.0 + 0.5 *
15        conflict_rate);
16
17        // More aggressive restart if we're making little progress
18        if (progressEstimate() - last_rephase_progress < 0.01) {
19            restart_threshold *= 0.7;
20        }
21    } else {
22        // Default threshold for initial phase
23        restart_threshold = restart_first;
24    }
25
26    // Check if we've reached the threshold
27    if (conflictR >= restart_threshold) {
28        conflictR = 0;
29        return true;
30    }
31    return false;
32 }
```

Figure 12. Example of LLM discovered heuristics of restart condition

original restart_function

```
1 void Solver::restart_function(){
2     fast_lbd_sum = lbd_queue_size = lbd_queue_pos = 0;
3     int level = 0;
4     cancelUntil(level);
5 }
```

discovered restart_function

```
1 void Solver::restart_function() {
2
3     // Update restart statistics
4     static double fast_avg = 0, slow_avg = 0;
5     static int restart_count = 0;
6
7     // Calculate moving averages of LBD scores
8     if (lbd_queue_size > 0) {
9         double curr_fast = fast_lbd_sum / lbd_queue_size;
10        fast_avg = 0.9 * fast_avg + 0.1 * curr_fast;
11        slow_avg = 0.99 * slow_avg + 0.01 * curr_fast;
12    }
13
14    // Determine restart level based on recent performance
15    int restart_level = 0;
16    if (fast_avg > 0 && slow_avg > 0) {
17        double ratio = fast_avg / slow_avg;
18        if (ratio > 1.2) {
19            // Recent conflicts are harder - do a full restart
20            restart_level = 0;
21        } else if (ratio > 1.0) {
22            // Partial restart - keep some learned clauses
23            restart_level = std::max(0, decisionLevel() / 2);
24        } else {
25            // Recent conflicts are easier - minimal restart
26            restart_level = std::max(0, decisionLevel() - 1);
27        }
28    }
29
30    // Reset LBD statistics
31    fast_lbd_sum = lbd_queue_size = lbd_queue_pos = 0;
32
33    // Perform the restart
34    cancelUntil(restart_level);
35
36    // Occasionally reduce the clause database after restarts
37    if (restart_count++ \% 16 == 15) {
38        reduceDB();
39    }
40
41    // Rebuild the order heap to reflect any activity changes
42    rebuildOrderHeap();
43 }
```

Figure 13. Example of LLM discovered heuristics of restart function

original rephase_condition

```
1 bool Solver::rephase_condition() {
2     if (rephases >= rephase_limit) return true;
3     else return false;
4 }
```

discovered rephase_condition

```
1 bool Solver::rephase_condition() {
2     // Adaptive rephasing based on progress and conflict density
3     const int base_rephase_limit = 1024;
4     const double progress_factor = 0.02; // 2% of total variables
5
6     if (conflictR < rephase_limit)
7         return false;
8
9     // Calculate normalized progress since last rephase
10    int progress = trail.size() - last_rephase_progress;
11    int progress_threshold = std::max(50, (int)(nVars() * progress_factor));
12
13    // Dynamic limit adjustment based on search progress
14    if (progress < progress_threshold) {
15        // Stagnation detected - reduce next rephase interval
16        rephase_limit = std::max(base_rephase_limit, rephase_limit * 2/3);
17    } else {
18        // Good progress - expand next rephase interval
19        rephase_limit = std::min(base_rephase_limit * 16, rephase_limit * 3/2);
20    }
21
22    // Reset conflict counter and update progress tracker in rephase_function
23    return true;
24 }
```

Figure 14. Example of LLM discovered heuristics of rephase condition

original rephase_function

```
1 void Solver::rephase_function() {
2     int var_nums = nVars();
3     conflictR = 0, rephases = 0, threshold *= 0.9, rephase_limit += 8192;Add commentMore
4     actions
5     int phase_rand = rand() \% 100;
6     if ((phase_rand -= 40) < 0){
7         for (int i = 0; i < var_nums; i++){
8             polarity[i] = local_best[i];
9         }
10    }
11    else if ((phase_rand -= 25) < 0){
12        for (int i = 0; i < var_nums; i++){
13            polarity[i] = 1-local_best[i];
14        }
15    }
16    else if ((phase_rand -= 15) < 0){
17        for (int i = 0; i < var_nums; i++){
18            polarity[i] = !polarity[i];
19        }
20    }
21    else {
22        for (int i = 0; i < var_nums; i++){
23            polarity[i] = saved[i];
24        }
25    }
}
```

Figure 15. Example of LLM discovered heuristics of rephase function

discovered rephase_function

```
1 void Solver::rephase_function() {
2     // Adaptive rephase limit adjustment
3     if (rephases > 0 && conflictR > last_rephase_progress) {
4         // Progress was made - be more aggressive with rephasing
5         rephase_limit = rephase_limit * 1.5;
6     } else {
7         // No progress - be more conservative
8         rephase_limit = rephase_limit * 0.9;
9         if (rephase_limit < 512) rephase_limit = 512;
10    }
11    last_rephase_progress = conflictR;
12    rephase_count++;
13
14    // Dynamic phase selection strategy
15    double rand_val = drand(random_seed);
16    if (rand_val < 0.4) {
17        // 40% chance: Use local best phases
18        for (int v = 0; v < nVars(); v++) {
19            polarity[v] = local_best[v];
20        }
21        if (verbosity >= 1) printf("| Rephase: Local best phases |\n");
22    } else if (rand_val < 0.7) {
23        // 30% chance: Invert all phases
24        for (int v = 0; v < nVars(); v++) {
25            polarity[v] = !polarity[v];
26        }
27        if (verbosity >= 1) printf("| Rephase: Inverted all phases |\n");
28    } else if (rand_val < 0.9) {
29        // 20% chance: Random phases for variables with low activity
30        double activity_threshold = 0.2 * var_inc;
31        for (int v = 0; v < nVars(); v++) {
32            if (activity[v] < activity_threshold) {
33                polarity[v] = drand(random_seed) < 0.5;
34            }
35        }
36        if (verbosity >= 1) printf("| Rephase: Randomized low-activity phases |\n");
37    } else {
38        // 10% chance: Original user phases
39        for (int v = 0; v < nVars(); v++) {
40            if (user_pol[v] != l_Undef) {
41                polarity[v] = (user_pol[v] == l_True);
42            }
43        }
44        // Add commentMore actions
45        if (verbosity >= 1) printf("| Rephase: Reset to user phases |\n");
46    }
47
48    // Reset the threshold for next local best phase tracking
49    threshold = trail.size() * 0.8;
50    cancelUntil(0);
51}
```

Figure 16. Example of LLM discovered heuristics of rephase function

original reduce_condition

```
1 bool Solver::reduce_condition()
2 {
3     if (learnts.size() >= max_learnts)  return true;
4     else return false;
5 }
```

discovered reduce_condition

```
1 bool Solver::reduce_condition()
2 {
3     // Check if we've reached the absolute limit of learnt clausesAdd commentMore actions
4     if (learnts.size() >= max_learnts)  return true;
5
6     // Consider memory pressure and garbage collection needs
7     if (ca.wasted() > ca.size() * garbage_frac * 0.8) return true;
8
9     // Consider the ratio between learnt and original clauses
10    if (learnts.size() > 0 && learnts.size() > 2 * nClauses()) return true;
11
12    // Consider recent solver performance (conflict rate)
13    if (conflictR > 1000 && learnts.size() > max_learnts * 0.8) return true;
14
15    return false;
16 }
```

Figure 17. Example of LLM discovered heuristics of reduce condition