ISSN: 2710-0588

MHQ-SAT: An End-to-End SAT Solver with a Multi-Head Hierarchical Query Mechanism

Yingyu Shi*

College of Computer Science and Technology, Henan Polytechnic University, Jiaozuo, China

*Corresponding Author: Yingyu Shi

Abstract

The Boolean Satisfiability Problem (SAT), as a representative NP-complete problem, plays a crucial role in integrated circuit verification, automated reasoning, and combinatorial optimization. In recent years, end-to-end SAT solvers based on graph neural networks (GNNs) have alleviated the reliance on handcrafted heuristics by learning the structural relationships between variables and clauses. However, existing neural SAT solvers with query-based mechanisms typically rely on a single query perspective, which limits their ability to capture diverse reasoning paths within complex constraint structures. As a result, their reasoning stability and generalization performance on large-scale instances remain insufficient. To address these limitations, this paper proposes MHQ-SAT, an end-to-end SAT solving model based on a multi-head hierarchical query mechanism. Built upon the QuerySAT framework, MHQ-SAT introduces a multi-head variable query mechanism and clause-level hierarchical attention modeling. By constructing multiple query perspectives in parallel, the proposed model enhances the expressive capacity of implicit reasoning processes. Meanwhile, a clause-level attention mechanism is employed to dynamically characterize the relative importance of different clauses during reasoning, enabling the model to focus on critical constraint structures. Furthermore, MHQ-SAT integrates query-driven clause updates with gated variable update mechanisms, achieving efficient information interaction between variables and clauses as well as stable iterative reasoning. Extensive experiments on randomly generated 3-SAT and k-SAT benchmarks demonstrate that MHQ-SAT consistently outperforms NeuroSAT, QuerySAT, and other baseline methods across various variable scales and reasoning step settings. In particular, on medium- and large-scale SAT instances, MHQ-SAT exhibits superior robustness and reasoning stability in terms of solution accuracy and the proportion of unsolved instances. Ablation studies further validate the effectiveness of the multi-head query mechanism and clause-level hierarchical modeling in improving overall performance. These results indicate that the proposed multi-head hierarchical query mechanism provides an effective and learnable reasoning enhancement paradigm for end-to-end SAT solving.

Keywords

Boolean Satisfiability Problem, Graph Neural Networks, End-to-End Solving, Multi-Head Query, Hierarchical Attention.

1. Introduction

The Boolean Satisfiability Problem (SAT) is one of the core problems in computer science and was the first problem to be rigorously proven NP-complete[1]. Research on SAT is not only of significant theoretical importance but also plays a critical role in engineering practice, with widespread applications in integrated circuit design and verification[2], software model

ISSN: 2710-0588

checking[3], automated theorem proving[4], planning and scheduling, and combinatorial optimization. Many complex decision problems, such as graph coloring, path planning, and constraint satisfaction problems, can be polynomially reduced to SAT and solved in a unified manner. Consequently, the development of efficient and robust SAT solvers has long been a central topic in artificial intelligence and theoretical computer science. Traditional SAT solvers are primarily based on the DPLL (Davis–Putnam–Logemann–Loveland) algorithm[5] and its extension, Conflict-Driven Clause Learning (CDCL)[6], combined with variable selection heuristics, clause learning mechanisms, and restart strategies. These methods have achieved remarkable success on industrial-scale instances. However, their performance heavily relies on manually designed rules and heuristic strategies, making them highly sensitive to specific problem distributions. When faced with novel SAT instances featuring significant structural changes or substantially increased scales, traditional solvers often require extensive parameter tuning or even the redesign of heuristic strategies[7], which limits their generalization and adaptability. In recent years, with the rapid advancement of deep learning technologies, researchers have begun to explore learning-based SAT solving methods that leverage neural networks to automatically capture structural features and reasoning patterns inherent in SAT problems, thereby reducing reliance on handcrafted heuristics. Such approaches typically model SAT instances as graph-structured data and employ graph neural networks (GNNs) [8] for information propagation and feature updates. Within an end-to-end training framework, these models learn implicit relationships between variable assignments and constraint satisfaction. This line of research provides a new perspective on SAT solving and lays the foundation for investigating learnable reasoning mechanisms.

This paper proposes MHQ-SAT, an end-to-end SAT solving model based on a multi-head hierarchical query mechanism. Built upon the QuerySAT framework, the proposed model systematically extends the query process at the reasoning mechanism level. By jointly leveraging multi-head queries and hierarchical attention, MHQ-SAT effectively models diverse reasoning paths and critical constraint structures in complex SAT instances. The main contributions of this work are summarized as follows:

(1) A multi-head query mechanism is proposed, in which query vectors are constructed in parallel across multiple latent subspaces, enabling the model to represent SAT instances from multiple reasoning perspectives and thereby enhancing the expressiveness and robustness of the implicit search process.

(2) A clause-level hierarchical attention mechanism is designed to explicitly characterize the relative importance of different clauses during the reasoning process, guiding the model to focus on constraint structures that are critical for satisfiability decisions.

(3) By organically integrating the multi-head query mechanism with hierarchical attention modeling, an end-to-end SAT solving framework, MHQ-SAT, is constructed. While preserving the differentiable and control-free nature of QuerySAT, the proposed framework significantly improves solving performance and reasoning stability on complex SAT instances.

Extensive experimental results demonstrate that MHQ-SAT consistently outperforms existing query-based SAT solvers across multiple SAT benchmarks, validating the effectiveness of the proposed multi-head hierarchical query mechanism for end-to-end SAT solving.

2. Related Work

2.1. Research Progress on Neural Network–Based SAT Solving

In neural SAT solving research, NeuroSAT represents one of the earliest works to systematically formulate SAT as a graph-based reasoning task. This approach models literals and clauses as nodes in a bipartite graph and learns structural relationships between variables and constraints through multiple rounds of message passing, demonstrating that neural networks

ISSN: 2710-0588

are capable of implicitly learning logical reasoning patterns to a certain extent. Subsequently, methods such as NNSAT[9] introduced deeper network architectures and improved state update strategies within the NeuroSAT[10] framework to enhance the expressive capacity for complex constraint relationships. However, these approaches still largely rely on a supervised learning paradigm that directly predicts variable assignments or satisfiability labels.To further strengthen structural modeling capabilities, some studies have explored higher-order graph representations. CircuitSAT[11] represents Boolean formulas as logical circuit structures and explicitly models logic gates and their connections within the network, thereby partially preserving the semantic information of logical formulas. In contrast, NeuroCore employs neural networks as scoring functions to provide branching heuristics for traditional CDCL solvers, enabling collaboration between neural models and classical search algorithms. While these hybrid approaches demonstrate improved stability on large-scale instances, they typically depend on external solvers and therefore struggle to form fully end-to-end neural reasoning frameworks. Beyond predictive or hybrid approaches, recent research has also revisited SAT solving from a continuous optimization perspective. DiffSAT[12] relaxes Boolean constraints into continuous and differentiable forms and applies gradient descent to search for satisfiable solutions in continuous spaces, attempting to bridge logical reasoning and numerical optimization. However, such methods are often sensitive to relaxation strategies and hyperparameter settings, and their stability on complex instances remains limited. Against this background, QuerySAT[13] introduces a query-based end-to-end SAT solving framework, providing a novel modeling paradigm for neural SAT solving. Instead of directly predicting discrete variable assignments, QuerySAT constructs loss functions through interactions between differentiable query vectors and clause states, and iteratively updates variable states using gradient information. This enables an implicit reasoning process without explicit search rules. Due to its concise architecture, stable training dynamics, and independence from complex control logic, QuerySAT is widely regarded as one of the representative models in end-to-end neural SAT solving.

Overall, existing neural SAT solving approaches can be broadly categorized into three classes: predictive end-to-end models, neural–symbolic hybrid models, and implicit reasoning models based on continuous optimization. Although these methods alleviate the reliance of traditional SAT solvers on handcrafted heuristics at different levels, they generally suffer from insufficient reasoning capability and degraded generalization performance when confronted with large-scale SAT instances featuring high constraint density or complex structures. This limitation primarily stems from the relatively homogeneous manner in which structural information is exploited during message passing, making it difficult for existing models to adaptively capture the varying importance of different variables and clauses throughout the reasoning process. Consequently, introducing more expressive and selective graph-structured modeling mechanisms to enhance the reasoning ability of neural SAT solvers on complex instances remains a key challenge in current research.

2.2. Limitations of Existing Methods and Research Motivation

A detailed analysis of existing query-based SAT solving methods reveals that their reasoning capability is still significantly constrained by model architecture design, which is mainly reflected in the following aspects.

(1) QuerySAT employs a single query vector to model the entire SAT instance, and its implicit reasoning process relies on gradient feedback from a single perspective. While such a single-query mechanism performs well on small-scale or structurally simple instances, it tends to be restricted by local information when dealing with large-scale SAT instances characterized by complex constraints and diverse clause structures. As a result, the search process struggles to sufficiently explore diverse reasoning paths, leading to degraded solving stability.

ISSN: 2710-0588

(2) SAT problems inherently exhibit a hierarchical structure, where the variable level and clause level are tightly coupled through logical constraints. However, the query mechanism in QuerySAT primarily focuses on updating variable-level states and lacks explicit modeling of the varying importance of constraints at the clause level. During loss construction, all clauses are implicitly treated with equal weights, making it difficult for the model to actively attend to critical clauses that dominate the reasoning process. This limitation restricts the model's expressive power in representing complex constraint structures.

(3) In practical SAT solving, the influence of different clauses on satisfiability decisions often evolves dynamically throughout the iterative reasoning process. Existing query mechanisms lack the capability to adaptively model such dynamic hierarchical relationships, which hampers the formation of stable and efficient search directions-an issue that becomes particularly pronounced in large-scale instances.

Based on the above analysis, introducing a more expressive query mechanism that enhances the modeling of multi-perspective reasoning and hierarchical constraint structures emerges as a key research direction for improving the performance of end-to-end SAT solvers.

3. The MHQ-SAT Method

3.1. Overall Model Architecture

To enhance the reasoning capability of neural SAT solvers on large-scale and complex formulas, this paper proposes an end-to-end SAT solving model, MHQ-SAT, based on a Multi-Head Query (MHQ) mechanism. The proposed model incorporates multi-head variable queries, clause-level hierarchical attention, and query-driven message passing to iteratively approximate satisfying assignments through multiple rounds of reasoning. MHQ-SAT adopts a variable–clause bipartite graph as the input representation and performs alternating updates between variable nodes and clause nodes. As illustrated in Fig. 1, the overall architecture mainly consists of variable and clause state initialization, clause-level hierarchical attention modeling, multi-head variable querying, query-driven clause updating and message passing, as well as gated variable updating and assignment prediction modules.

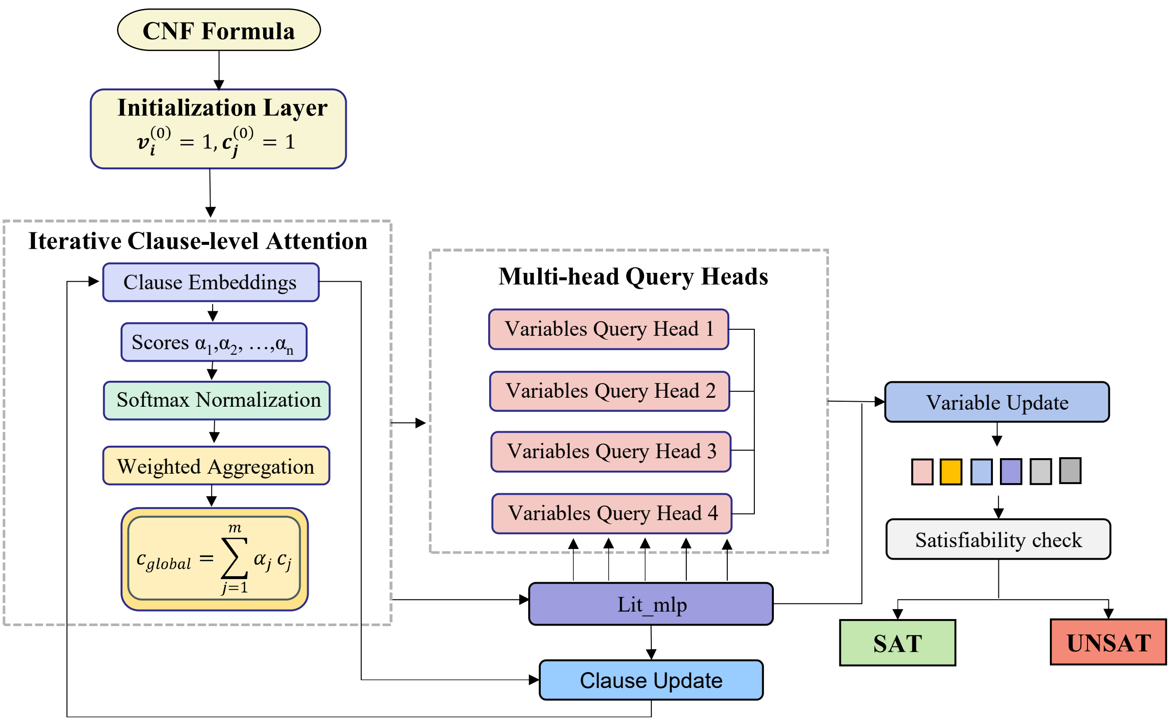

Figure 1. Overall Architecture of the MHQ-SAT Model

ISSN: 2710-0588

3.2. Model Implementation Process

3.2.1. Macros.

Given a CNF formula, the model first represents it as a variable–clause bipartite graph. Suppose the formula contains n variables and m clauses. A sparse adjacency matrix $A \in \{0,1\}^{2n\times m}$ is constructed, where each variable is decomposed into two literal nodes corresponding to its positive and negative literals, in order to distinguish positive and negative constraint relationships. At the initialization stage, the hidden states of all variable nodes and clause nodes are initialized as constant vectors, as shown in Eq. (1):

$$v_i^{(0)} = 1, c_j^{(0)} = 1 \tag{1}$$

where d_v and d_c represent the dimensions of the variable states and clause states, respectively.

3.2.2. Clause-Level Hierarchical Attention Mechanism.

The clause-level hierarchical attention mechanism enables the model to dynamically perceive the global constraint structure formed by all clauses at each reasoning iteration, thereby alleviating local information bias and improving the global consistency of variable reasoning. In each iteration, the model first performs global modeling over the current states of all clauses. Specifically, for each clause feature vector c_j, a mean pooling operation is applied to compress its high-dimensional representation into a scalar value. This scalar reflects the overall activation strength of the clause at the current iteration and serves as an initial measure of clause importance. A larger mean pooling value indicates that the constraint information contained in the clause is more prominent during the current reasoning step. To obtain the relative importance of each clause within the global context, the mean pooling results of all clauses are normalized using a Softmax function, as shown in Eq. (2):

$$\alpha_j = \frac{exp(mean(c_j))}{\sum_k exp(mean(c_k))} \tag{2}$$

Here, $\exp(\mathrm{mean}(c_j))$ amplifies the importance differences among clauses, making the contributions of critical clauses more pronounced, while the denominator normalizes the importance scores of all clauses to ensure that the attention weights sum to one. The resulting weight α_j represents the relative importance of the j-th clause in the current global constraint structure, with higher values indicating a stronger guiding role in subsequent reasoning.

Subsequently, the model computes a global clause context vector by performing a weighted summation over all clause state vectors $\{c_j\}$ using the attention weights $\{\alpha_j\}$, as shown in Eq. (3):

$$c_{global} = \sum_{j=1}^{m} \alpha_j c_j \tag{3}$$

The resulting global clause context vector integrates constraint information from all clauses while emphasizing the contributions of key clauses through attention weighting, thereby providing a comprehensive representation of the global constraint structure. After obtaining c_{global}, the model broadcasts this vector and concatenates it with the feature representations of

ISSN: 2710-0588

each variable node. This design allows the subsequent variable query generation process to directly perceive the overall constraint structure, preventing variable reasoning from being limited to local information. By incorporating global clause constraints, the model can dynamically adjust its reasoning direction, enhancing the coherence and accuracy of the overall reasoning process.

3.2.3. Multi-Head Variable Query Mechanism

The multi-head variable query mechanism enhances the model's ability to capture diverse constraint patterns by generating and integrating queries from multiple perspectives, thereby improving reasoning robustness on complex SAT instances. For each variable node, the model concatenates three types of information-the current variable state, the global clause context, and a random noise vector-to form the query input vector x_i:

$$x_i = \left[v_i^{(t)}, c_{global}, \epsilon\right] \tag{4}$$

Here, $v_i^{(t)}$ denotes the state vector of the i-th variable at the t-th reasoning iteration, which encapsulates historical reasoning information such as the current assignment tendency and interaction strength with other variables, and serves as the core representation of the variable's reasoning process. The global clause context c_{global}, generated by the clause-level hierarchical attention mechanism, conveys globally aggregated constraint information to the variable, enabling reasoning beyond local neighborhoods and allowing the model to perceive the overall constraint structure of the problem. The random perturbation vector $\epsilon$is a randomly sampled noise vector. Since the SAT solution space is discrete and highly complex, fixed inputs may cause the model to converge to locally optimal assignment configurations. Introducing random perturbations encourages exploration of diverse potential assignment directions and reduces the risk of being trapped in local optima. The model then employs Kindependent query-head MLPs to generate multiple query vectors:

$$\mathrm{q}_i^{(k)} = \mathrm{MLP}_k(\mathrm{x}_i), k = 1, \dots, K \tag{5}$$

Different query heads are expected to learn to capture constraint patterns from distinct dimensions. For example, some heads may focus on the number of variables contained in clauses, others may attend to the frequency of variable occurrences across clauses, while certain heads may emphasize the difficulty of clause satisfaction. These multiple query vectors provide the variable with K different "perceptual perspectives," enabling a more comprehensive understanding of clause constraints.

Each query head independently computes a Softplus-based unsupervised clause loss to measure the consistency between the generated query and clause constraints. If the variable assignment implied by the query is more likely to satisfy a clause, the corresponding loss value is smaller; otherwise, the loss increases. As SAT problems lack explicit supervision labels, this unsupervised loss serves as the core optimization signal for model training. By integrating multiple query perspectives, the multi-head query mechanism effectively alleviates the expressive limitations of single-query approaches on complex SAT instances.

3.2.4. Query-Driven Clause Update and Message Passing.

In the MHQ-SAT model, query-driven clause updating and message passing serve as a critical component that connects variable reasoning with the evolution of clause states. This module

ISSN: 2710-0588

receives the outputs of multi-head variable queries and enables bidirectional information interaction between variables and clauses through gradient feedback and structured message passing. After obtaining the loss values from the multi-head queries, the model first transforms these losses into gradient signals for variable updates. For each query head k, the gradient of the query loss with respect to the corresponding query vector is computed. Since different query heads capture constraint patterns from different perspectives, the gradients from all K query heads are averaged to obtain a unified variable-level gradient signal. The resulting gradient is then directly applied to update the states of variable nodes. Meanwhile, the model performs variable-to-clause information propagation through a literal-based message passing mechanism, enabling clauses to perceive the updated variable states. Specifically, each variable state vector is first mapped to its corresponding literal representations. Using the predefined sparse adjacency matrix, all literal representations are then aggregated and propagated to their associated clause nodes. The aggregated results constitute structured messages received by each clause node. The model concatenates the current clause state vector (which encodes historical constraint features), the aggregated literal messages, and the query loss features to form the input to the clause update MLP. After a nonlinear transformation by the clause update MLP, a new clause state vector is generated.

The graph structures corresponding to SAT problems often exhibit severe node degree imbalance, which can lead to unstable gradients and hinder model convergence during training. To address this issue, the MHQ-SAT model incorporates PairNorm normalization during the clause update stage and employs residual connections to preserve historical information.

3.2.5. Gated Variable Update and Assignment Prediction.

In the inference process of the MHQ-SAT model, gated variable updating and assignment prediction constitute the final critical step that bridges constraint perception and feasible solution output. Through an adaptive gating mechanism, this module integrates multi-source information to update variable states and enables early termination of the reasoning process via dynamic satisfiability detection. During the variable update stage, the model does not simply perform iterative updates on variable states. Instead, it employs a gated update network to generate more robust variable representations by jointly considering multiple signals:

$$\mathrm{v}_i^{(t+1)} = \mathrm{Gate}\left(\left[\nabla \mathrm{q}_i, \mathrm{v}_i^{(t)}, \mathrm{m}_i^{+}, \mathrm{m}_i^{-}\right]\right) \tag{6}$$

where $\nabla \mathrm{q}_i$ denotes the gradient signal derived from the multi-head query losses, reflecting the optimization direction indicating how the variable state should be adjusted to better satisfy clause constraints; $\mathrm{v}_i^{(t)}$ represents the state vector of variable i at the t-th reasoning iteration, encoding its historical inference information and current assignment tendency; and m_i^{+} and m_i^{-} denote the aggregated messages from the positive and negative literals of the variable, respectively.

The key advantage of the gating network lies in its adaptive fusion capability. It learns to assign dynamic importance weights to different input signals at each iteration, selectively emphasizing the most informative sources for variable updates while avoiding information redundancy or gradient conflicts that may arise from naive feature concatenation. To further enhance inference stability, the updated variable states are also processed with normalization operations and residual connections, ensuring smoother state transitions and more stable convergence.

Finally, the model employs a variable output head to predict multi-channel logits for each variable and dynamically checks whether the current assignment satisfies all clause constraints

ISSN: 2710-0588

at each reasoning iteration. Once a satisfying assignment is detected, the inference process is terminated early.

4. Experimental Design and Results Analysis

4.1. Experimental Setup

To verify the effectiveness and generalization capability of the proposed MHQ-SAT method for solving SAT problems, a series of comparative and ablation experiments are conducted under a unified experimental environment. This section provides a detailed description of the experimental configuration from the perspectives of the experimental environment, dataset construction, and training and testing settings.

4.1.1. Experimental Environment.

All experiments are performed on the same computing platform. The proposed model is implemented using the TensorFlow framework. The experimental system is equipped with an Intel(R) Xeon(R) Gold 6330 CPU @ 2.00GHz and an NVIDIA GeForce RTX 4090 GPU. To ensure the reproducibility and fairness of the experimental results, all models adopt identical initialization strategies and training hyperparameter settings.

4.1.2. Dataset Construction.

To systematically evaluate the reasoning capability and mechanism effectiveness of the proposed MHQ-SAT model, randomly generated SAT instances are used as experimental datasets. All experiments are conducted on SAT instances generated in Conjunctive Normal Form (CNF), including both satisfiable (SAT) and unsatisfiable (UNSAT) samples, ensuring a balanced distribution for training and evaluation. In the 3-SAT experiments, each clause consists of exactly three literals. By controlling the ratio between the number of variables and clauses, SAT instance sets with different scales and constraint densities are constructed. This dataset is mainly used to evaluate the reasoning performance and convergence behavior of MHQ-SAT under standard SAT problem settings. To further verify the generalization ability of the model under more complex constraint conditions, a k-SAT dataset is additionally constructed, where the clause length $k > 3$. Compared with 3-SAT problems, k-SAT instances exhibit a larger search space and higher reasoning difficulty. Therefore, these experiments more comprehensively demonstrate the advantages of the proposed multi-head query mechanism and graph attention structure when handling complex SAT problems.

4.2. Experimental Results and Analysis

This section provides a systematic evaluation of the overall performance, generalization capability, and effectiveness of the key structural designs of the proposed MHQ-SAT model on randomly generated 3-SAT and k-SAT datasets. The experimental results are analyzed from three aspects: overall model comparison, ablation studies, and parameter sensitivity analysis.

4.2.1. Solving Accuracy Comparison.

Table 1 reports the solving accuracy of NeuroSAT, QuerySAT, and the proposed MHQ-SAT model on k-SAT instances with different variable scales. It can be observed that, as the number of variables increases, the solving accuracy of all models exhibits varying degrees of degradation, indicating that large-scale SAT instances impose higher demands on the reasoning capability of neural solvers. From an overall performance perspective, MHQ-SAT achieves the best or near-best solving accuracy across almost all experimental settings. On small-scale instances with 50 variables, the accuracy of all models approaches saturation. However, on medium- and large-scale instances with 100–400 variables, the advantage of MHQ-SAT becomes increasingly evident. For example, under the setting of 300 variables with $T = 4096$reasoning steps, MHQ-SAT attains an accuracy of 0.8809, achieving an improvement of

ISSN: 2710-0588

nearly 15 percentage points compared to QuerySAT. Further analysis on the influence of reasoning steps shows that increasing the number of iterations from 512 to 4096 consistently improves the solving accuracy of all models. Notably, MHQ-SAT exhibits a more significant performance gain on medium- and large-scale instances. This phenomenon suggests that the proposed multi-head query mechanism is able to continuously provide effective gradient signals during long-horizon reasoning, thereby guiding the model to progressively approach feasible solutions within the satisfiable solution space. In summary, the above results demonstrate that MHQ-SAT effectively enhances reasoning stability and solving capability on complex k-SAT instances by incorporating multi-head query mechanisms and clause-level information aggregation. In particular, the proposed model exhibits stronger robustness and generalization performance under large-scale problem settings.

Table 1. Accuracy Comparison of Models on the K-SAT Task

var	NeuroSAT		QuerySAT		MHQ-SAT	
	T=512	T=4096	T=512	T=4096	T=512	T=4096
50	0.8923	0.9346	0.9956	**0.9990**	0.9984	0.9908
100	0.8548	0.8772	0.9688	0.9918	0.9825	**0.9953**
150	0.7315	0.8065	0.8888	0.9625	0.9402	**0.9836**
200	0.6827	0.7131	0.7697	0.9000	0.8850	**0.9728**
250	0.5932	0.6513	0.6984	0.8079	0.8100	**0.9302**
300	0.5566	0.6068	0.5830	0.7494	0.7421	**0.8809**
350	0.4421	0.5565	0.5257	0.6848	0.6747	**0.8296**
400	0.3470	0.4378	0.4572	0.5516	0.6139	**0.7720**

Figure 2 illustrates the comparison of solving accuracy among different SAT solvers on the Random 3-SAT dataset as the number of variables increases from 50 to 300. From the overall trend, the solving accuracy of all three models decreases to varying degrees as the variable scale grows, indicating that the expansion of problem size significantly increases the reasoning difficulty of SAT instances. However, MHQ-SAT consistently achieves the highest solving accuracy across all variable scales, with particularly pronounced advantages on large-scale instances. On small-scale problems (50–100 variables), both QuerySAT and MHQ-SAT achieve nearly 100% solving accuracy, significantly outperforming NeuroSAT. This observation suggests that query-based models possess stronger expressive capability in modeling clause constraints. As the number of variables further increases beyond 150, the performance of NeuroSAT degrades rapidly, and QuerySAT also begins to exhibit noticeable performance deterioration. In contrast, MHQ-SAT maintains relatively high solving accuracy on medium- and large-scale instances (200–300 variables). For example, at the scale of 300 variables, MHQ-SAT still achieves an accuracy close to 0.8, whereas QuerySAT and NeuroSAT drop to approximately 0.58 and 0.52, respectively. These results demonstrate that the proposed multi-head query mechanism and clause-level information modeling effectively alleviate the performance degradation caused by increasing problem scale. Overall, MHQ-SAT benefits from parallel modeling of multiple constraint subspaces through multi-head queries, combined with hierarchical clause-level information aggregation, enabling the model to exhibit stronger reasoning robustness and generalization capability under complex constraint conditions, especially for large-scale SAT instances.

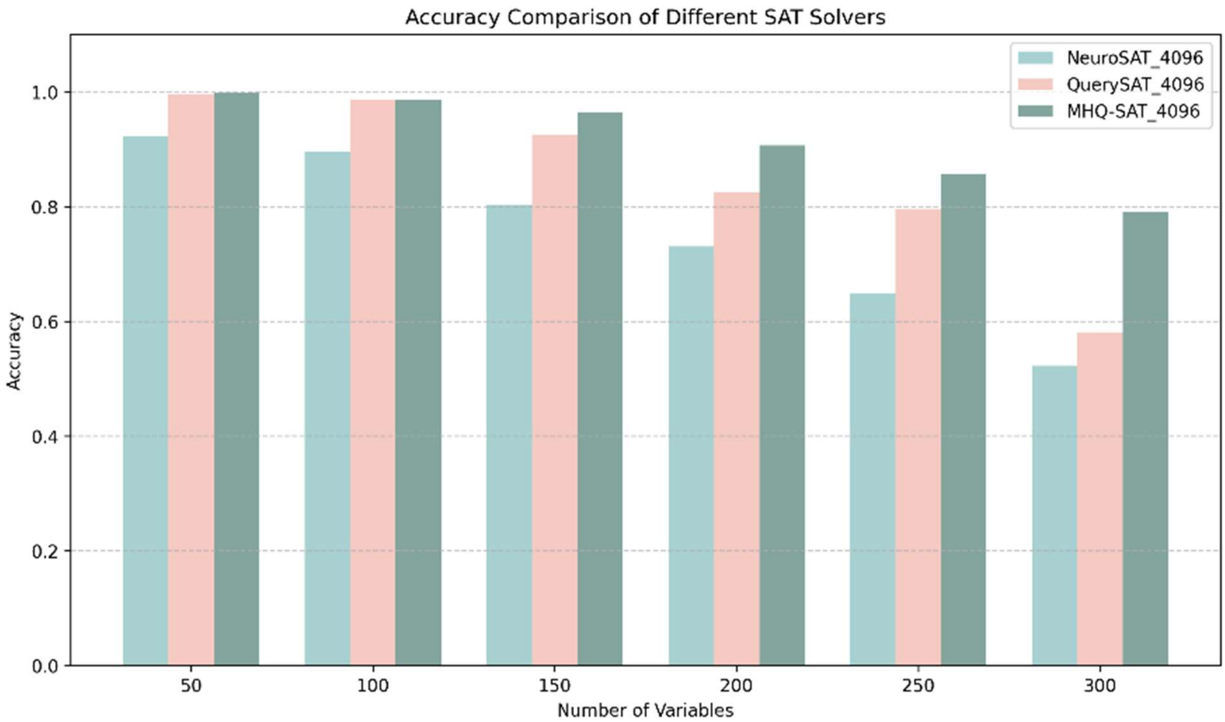

Figure 2. Accuracy of MHQ-SAT on the 3-SAT test dataset

4.2.2. Analysis of Unsolved Instance Ratio.

The ratio of unsolved instances provides an alternative perspective on a model's ability to find feasible solutions within a limited number of reasoning steps. A lower unsolved ratio indicates higher reasoning efficiency and stability. Table 2 reports the unsolved instance ratios of different models on SAT instances with varying variable scales under a fixed maximum reasoning step constraint. From the overall trend, as the variable scale increases from 100 to 300, the unsolved instance ratio of all models rises significantly, indicating that larger-scale SAT problems pose substantially greater challenges for neural reasoning. Under the same reasoning step settings, NeuroSAT consistently exhibits the highest unsolved ratios, while QuerySAT achieves moderate improvements across all scales by introducing the query mechanism. In contrast, MHQ-SAT demonstrates notably stronger stability on medium- and large-scale instances, with unsolved instance ratios significantly lower than those of the other two models. Notably, MHQ-SAT achieves the lowest unsolved instance ratio across all experimental settings, and this advantage becomes more pronounced for larger problem sizes and smaller reasoning budgets. For example, under the setting of 300 variables with $T = 512$reasoning steps, MHQ-SAT achieves an unsolved ratio of 16.0%, representing a reduction of approximately 12.5 percentage points compared to QuerySAT. When the reasoning steps are increased to $T = 4096$, the unsolved ratio further decreases to 7.7%, indicating that MHQ-SAT possesses a stronger capability to explore the solution space under sufficient reasoning iterations. Further analysis of the impact of reasoning steps shows that increasing the number of steps from 512 to 4096 leads to a reduction in unsolved instance ratios for all models; however, the improvement is substantially more significant for MHQ-SAT. This observation suggests that the multi-head query mechanism can continuously provide diverse and effective guidance signals throughout the reasoning process, enabling the model to cover critical clauses more efficiently within limited iterations and reducing the likelihood of reasoning stagnation or convergence to local optima. In summary, by incorporating multi-head query mechanisms and clause-level

ISSN: 2710-0588

information aggregation, MHQ-SAT effectively enhances reasoning stability and solving capability on complex K-SAT instances, exhibiting superior robustness and generalization performance, particularly in large-scale problem settings.

Table 2. Unsolved Instance Ratio under the Maximum Reasoning Steps

var	reasoning steps	NeuroSAT(%)	QuerySAT(%)	MHQ-SAT（%）
100	512	18.6	9.4	4.5
	4096	10.2	4.1	2.0
200	512	31.4	18.7	9.2
	4096	19.8	9.3	3.0
300	512	44.7	29.5	16.0
	4096	32.6	18.4	7.7

4.2.3. Impact of the Number of Multi-Head Queries.

Table 3 reports the solving accuracy of the MHQ-SAT model on SAT instances with 300 variables under different configurations of the number of query heads. It can be observed that as the number of query heads increases from 1 to 4, the solving accuracy consistently improves, indicating that the multi-head query mechanism is able to capture effective information from multiple potential reasoning directions in parallel, thereby enhancing the model's ability to represent complex constraint structures. When the number of query heads is set to 4, MHQ-SAT achieves its highest solving accuracy, reaching 0.8809. However, when the number of query heads is further increased to 8, the model performance no longer improves and instead exhibits a noticeable decline. This phenomenon suggests that an excessive number of query heads may introduce redundant or even mutually interfering query information, increasing noise during the information aggregation process and consequently impairing the stability and effectiveness of the reasoning procedure. Overall, an appropriate number of query heads enables a favorable balance between model expressiveness and training stability. Under the experimental settings considered in this work, four query heads are demonstrated to be a well-suited configuration.

Table 3. Solving Accuracy of the Model under Different Numbers of Query Heads

Heads	Accuracy (300 vars)
1	0.7494
2	0.8267
4	**0.8809**
8	0.8749

4.2.4. Ablation Study.

To verify the effectiveness of the proposed multi-head hierarchical query mechanism and its constituent modules, an ablation study is conducted on SAT instances with 300 variables under a maximum of 4096 reasoning iterations. Specifically, QuerySAT, QuerySAT with Clause

ISSN: 2710-0588

Attention (+Clause Attention), and QuerySAT with Clause Attention and Clause Update (+Clause Update) are evaluated for comparison. As the baseline model, QuerySAT achieves a solving accuracy of 0.7494 on the test set. After incorporating the clause-level attention mechanism, the accuracy increases to 0.8012, indicating that explicitly weighting critical clauses helps enhance the model's global perception of constraint structures. Furthermore, by introducing the clause update module, the solving accuracy is further improved to 0.8247, demonstrating that iteratively updating clause states during reasoning effectively strengthens the propagation of clause-level information in variable inference. When clause-level attention, clause update, and the multi-head query mechanism are jointly incorporated, the complete MHQ-SAT model achieves a solving accuracy of 0.8809 on the test set, representing a substantial improvement over the baseline. These results show that the multi-head query mechanism is able to guide the variable assignment search from multiple potential reasoning directions in parallel, thereby significantly enhancing the model's reasoning capability and stability on complex SAT instances. Overall, the ablation results indicate that clause-level modeling components provide a stable structural improvement for MHQ-SAT, while the multi-head query mechanism further amplifies the overall reasoning performance on this basis. Rather than contributing independently, these modules interact synergistically during the reasoning process, jointly promoting the model's effectiveness on large-scale SAT problems.

Table 4. Ablation Study Results

	300(T=4096)
QuerySAT	0.7494
+Clause Attention	0.8012
+ Clause Update	0.8247
MHQ-SAT	**0.8809**

5. Conclusion

In this paper, we proposed MHQ-SAT, an end-to-end neural SAT solving framework based on a multi-head hierarchical query mechanism, aiming to improve the reasoning capability and stability of neural SAT solvers on large-scale and structurally complex instances. Motivated by the limitations of existing query-based approaches such as QuerySAT, which rely on a single query perspective and lack explicit modeling of clause-level importance, MHQ-SAT systematically extends the query-driven reasoning process from both multi-view and hierarchical perspectives. Specifically, MHQ-SAT introduces a multi-head variable query mechanism to enable parallel exploration of diverse implicit reasoning directions, thereby enhancing the expressive power and robustness of gradient-based implicit search. In addition, a clause-level hierarchical attention mechanism is designed to dynamically model the relative importance of clauses during reasoning, allowing the model to focus on key constraints that dominate satisfiability decisions. By integrating query-driven clause updates and gated variable updates, MHQ-SAT achieves coordinated evolution of variable and clause states while preserving full differentiability and avoiding explicit control logic. Extensive experiments on randomly generated 3-SAT and k-SAT datasets demonstrate that MHQ-SAT consistently outperforms NeuroSAT and QuerySAT across different variable scales and reasoning step budgets. In particular, MHQ-SAT shows significant advantages on medium- and large-scale instances in terms of both solving accuracy and the ratio of unsolved instances. Ablation studies

ISSN: 2710-0588

further confirm that the multi-head query mechanism, clause-level attention, and clause update modules contribute complementary improvements to the overall reasoning performance.

Overall, MHQ-SAT provides an effective and scalable solution for end-to-end neural SAT solving by enhancing multi-view reasoning and hierarchical constraint modeling. This work highlights the importance of expressive query mechanisms in neural reasoning and opens up new directions for designing more robust and generalizable neural SAT solvers.

References

[1] Alyahya, Tasniem Nasser, Mohamed El Bachir Menai, and Hassan Mathkour. "On the structure of the boolean satisfiability problem: a survey." ACM Computing Surveys (CSUR) 55.3 (2022): 1-34.

[2] Qian, Yuhang, et al. "X-SAT: An Efficient Circuit-Based SAT Solver." 2025 62nd ACM/IEEE Design Automation Conference (DAC). IEEE, 2025.

[3] Ivančić, Franjo, et al. "Efficient SAT-based bounded model checking for software verification." Theoretical Computer Science 404.3 (2008): 256-274.

[4] Holden, Sean B. "Machine learning for automated theorem proving: Learning to solve SAT and QSAT." Foundations and Trends® in Machine Learning 14.6 (2021): 807-989.

[5] Nieuwenhuis, Robert, Albert Oliveras, and Cesare Tinelli. "Solving SAT and SAT modulo theories: From an abstract Davis--Putnam--Logemann--Loveland procedure to DPLL (T)." Journal of the ACM (JACM) 53.6 (2006): 937-977.

[6] Cai, Shaowei, and Xindi Zhang. "Deep cooperation of CDCL and local search for SAT." International Conference on Theory and Applications of Satisfiability Testing. Cham: Springer International Publishing, 2021.

[7] Markidis, Stefano. "The old and the new: Can physics-informed deep-learning replace traditional linear solvers?." Frontiers in big Data 4 (2021): 669097.

[8] Liang, Fan, et al. "Survey of graph neural networks and applications." Wireless Communications and Mobile Computing 2022.1 (2022): 9261537.

[9] Mejia-Lavalle, Manuel, et al. "Modified Neural Net for the Boolean Satisfiability Problem." 2015 International Conference on Mechatronics, Electronics and Automotive Engineering (ICMEAE). IEEE, 2015.

[10] Selsam, Daniel, et al. "Learning a SAT solver from single-bit supervision." arXiv preprint arXiv:1802.03685 (2018).

[11] Chukharev, Konstantin, et al. "Effective Partitioning Method with Predictable Hardness for CircuitSAT." IEEE Access (2025).

[12] Zhang, Yu, et al. "DiffSAT: Differential MaxSAT Layer for SAT Solving." Proceedings of the 43rd IEEE/ACM International Conference on Computer-Aided Design. 2024.

[13] Ozolins, Emils, et al. "Goal-aware neural SAT solver." 2022 International joint conference on neural networks (IJCNN). IEEE,