Structural Resolution of Non-Local SAT Instances

Glass-Box Closure for a Subclass of NP Problems

Simone Mazzoni
RCUBEAI Research

January 2026

Abstract

Boolean Satisfiability (SAT) is a central problem in theoretical computer science and a cornerstone of modern industrial reasoning systems. While state-of-the-art solvers based on Conflict-Driven Clause Learning (CDCL) achieve impressive results on many benchmarks, they remain fundamentally black-box systems and exhibit severe limitations on globally constrained instances. In this paper, we present a Glass-Box resolution engine that treats SAT resolution as a process of structural closure rather than heuristic search. We demonstrate, on a frozen class of non-local SAT instances (global XOR/parity constraints encoded as CNF), that this approach robustly closes all instances under a strict no-order-change constraint, while producing complete, auditable resolution traces. These results do not claim to solve SAT in general, but establish that a meaningful subclass of NP problems is already better handled by a structural, auditable reasoning paradigm, opening the door to practical, governable reasoning systems.

Contents

1 Introduction

Boolean Satisfiability (SAT) occupies a unique position in computer science. As the first problem shown to be NP-complete, it is both a theoretical benchmark for computational hardness and a practical workhorse for applications ranging from hardware verification and planning to cryptography and security analysis.

Over the last decades, the dominant practical approach to SAT has been based on Conflict-Driven Clause Learning (CDCL). Modern CDCL solvers are extraordinarily effective on many classes of instances and form the backbone of contemporary SAT-based systems.

However, two fundamental limitations remain largely unaddressed:

1. **Opacity:** CDCL solvers operate as black boxes. Even when they succeed, they provide little insight into *why* a solution exists or why a contradiction is inevitable.
2. **Non-local structure:** SAT instances with strong global constraints often defeat local reasoning heuristics, leading to exponential blow-up or timeouts.

This paper explores an alternative paradigm: treating SAT resolution as a *structural closure process* with explicit observability, auditability, and replay, rather than as a purely heuristic search.

2 SAT, NP, and Why Resolution Matters

2.1 SAT and NP

The class NP consists of decision problems for which a proposed solution can be verified in polynomial time. SAT is NP-complete, meaning that every problem in NP can be reduced to a SAT instance.

This theoretical status has two important consequences:

- SAT serves as a proxy for the entire NP class.
- Any advance in understanding SAT resolution has implications far beyond SAT itself.

At the same time, NP-completeness does *not* imply that all SAT instances are equally hard. Empirically, SAT exhibits a rich internal structure, with some instances being trivially solvable and others exhibiting extreme difficulty.

2.2 The Practical Gap

From an industrial perspective, SAT solvers are not only judged by whether they return SAT or UNSAT, but by:

- robustness across instance families,
- scalability with problem size,
- and increasingly, **auditability and governance**.

In regulated or safety-critical contexts, a solver that cannot explain its reasoning process is often insufficient, regardless of raw performance.

3 Black-Box SAT Solving and Its Limits

Modern CDCL solvers rely on local decisions, conflict analysis, and learned clauses. While extremely effective in practice, this paradigm has two intrinsic weaknesses:

3.1 Locality Bias

CDCL operates on local clause-level conflicts. When the relevant structure of a problem is global (e.g., parity constraints, global consistency conditions), this structure is not explicitly represented and must be rediscovered indirectly through search.

3.2 Lack of Structural Proof

Although proof formats exist, they are typically low-level, difficult to interpret, and unsuitable as operational explanations. As a result, CDCL solvers provide answers, not reasoning artifacts.

These limitations motivate the search for alternative resolution paradigms.

4 Structural Resolution as Closure

4.1 Resolution as Closure

We adopt the view that resolution is not fundamentally a search for assignments, but a process of achieving *closure* in a representation space.

In this perspective:

- partial representations are progressively refined,
- incompatible structures are eliminated,
- and closure is reached when no residual inconsistency remains.

4.2 Scale Transitions

Within a fixed representational order, two operations are central:

- **Scale In:** isolating local substructures and micro-regimes.
- **Scale Out:** composing stabilized interfaces into higher-level structures.

Importantly, these operations do not change the underlying mathematical order of the problem; they reorganize how structure is made explicit.

5 Problem Class: Global XOR-CNF

To evaluate this paradigm, we focus on a deliberately restricted but meaningful class of SAT instances:

- CNF formulas encoding global XOR (parity) constraints,
- including both satisfiable and unsatisfiable cases,
- with increasing size and density.

Parity constraints are inherently global: no individual clause reveals the constraint structure, making this class a canonical stress test for local reasoning.

We emphasize that this class is not artificial; it arises naturally in cryptography, error-correcting codes, and consistency verification.

6 Experimental Protocol

The experimental protocol is frozen to ensure interpretability:

- problem sizes up to 1600 variables,
- strict prohibition of inter-order changes,
- fixed seeds and deterministic execution,
- full trace recording for every run.

A modern CDCL solver (Kissat) is used as a baseline under a fixed timeout.

7 Results

Across all tested sizes, the Glass-Box resolution engine:

- closed 100% of instances (SAT and UNSAT),
- exhibited stable contraction behavior,
- and produced complete resolution traces.

By contrast, the CDCL baseline exhibited an increasing rate of timeouts as instance size grew.

The key result is not raw speed, but **robust closure under global constraints**.

8 Auditability and Proof Artifacts

For each resolved instance, the system provides:

- a verifiable certificate,
- a complete event trace,
- a residual trajectory,
- and deterministic replay.

This enables independent verification without exposing internal decision heuristics, demonstrating that auditability and non-trivial NP resolution are compatible.

9 Discussion and Scope

9.1 What This Work Shows

This work establishes that a meaningful subclass of NP problems can already be handled more robustly by a structural, Glass-Box resolution paradigm than by dominant black-box solvers.

Beyond numerical performance, each reported result is accompanied by a complete set of proof artifacts, enabling independent audit and deterministic replay of the resolution process.

Process-Level Superiority and Glass-Box Resolution. A central and distinguishing result of this work lies not only in the ability to close a class of non-local SAT instances, but in the *nature of the resolution process itself.* Each R3-MRM execution produces explicit proof artifacts—namely a verifiable certificate, a complete stepwise resolution trace, and deterministic replay data—thereby rendering the resolution process transparent and independently auditable at the process level.

By contrast, while modern CDCL solvers may emit machine-verifiable certificates (e.g., DRAT or LRAT proofs), these artifacts represent only the *final logical correctness* of the outcome. They do not expose a structured resolution trajectory, a hierarchical organization of reasoning steps, nor any continuous notion of progress toward closure. The internal reasoning dynamics of CDCL remain heuristic, non-deterministic, and largely opaque, even when a formal proof object is available.

This distinction is fundamental. R3-MRM provides an *operational proof* of resolution: the reasoning process itself becomes an observable mathematical object, subject to replay, inspection, and audit. In this sense, the proposed system achieves a level of process-level transparency that is not offered by black-box solving paradigms. The superiority demonstrated here is therefore not merely a matter of performance on a specific problem class, but a qualitative difference in how resolution is represented, validated, and governed.

9.2 What This Work Does Not Claim

We explicitly do not claim:

- that SAT is easy in general,
- that NP problems admit efficient solutions universally,
- or that this approach replaces CDCL on all instances.

9.3 Broader Implications

The results suggest that computational difficulty is not solely a property of problem size, but of representational alignment. Making global structure explicit changes the nature of the resolution process.

10 Critical Analysis and Impact

10.1 Strengths and Limitations

Glass-Box Resolution as a Core Strength. A central contribution of the proposed system is its *Glass-Box* nature. In contrast to classical SAT solvers, which operate as black boxes, each execution produces a structured bundle of proof artifacts, including an explicit event trace, a measurable residual trajectory $\mu(t)$, and a machine-verifiable certificate (for both SAT and UNSAT cases). This bundle enables independent verification, deterministic replay, and post-hoc analysis without requiring access to internal decision heuristics. Such auditability is not presented as a debugging feature, but as a foundational requirement for institutional governance and certified reasoning systems.

Specificity of the Target Class. The reported performance gains are demonstrated on a deliberately chosen class of instances: global XOR constraints encoded as CNF. This class serves as a structural stress test due to its inherently non-local nature, which is known to defeat local reasoning strategies commonly employed by CDCL solvers. The success of the approach must therefore be understood in relation to this structural specificity, rather than as a blanket claim over all SAT instances.

Limits of the v1 Study. The present work is intentionally conservative in scope. It explicitly excludes any claim of general superiority over all SAT problem classes, as well as any implication toward resolving the $P = NP$ question. The experimental dataset is restricted to a minimum validation set (five instances per size), designed to establish the existence of a stable closure regime rather than to provide exhaustive statistical coverage. As such, the results should be interpreted as a proof of concept and a foundational benchmark, not as a comprehensive performance study.

10.2 Scope and Broader Impact

A Proof of Concept for Auditable AI. Beyond SAT, the primary impact of these results lies in demonstrating that it is possible to resolve non-trivial NP-class problems while simultaneously emitting artifacts suitable for independent audit. This directly challenges the widespread assumption that high-performance reasoning systems must necessarily be opaque. The work provides a concrete counterexample: a system that resolves complex global constraints while remaining fully auditable.

A Methodological Bridge Toward Language Representation Models. The SAT phases presented here are explicitly positioned as a methodological prototype. The same discipline of proof—combining explicit traces, residual contraction, and verifiable certificates—is intended to transfer to Language Representation Models (LRM). In that context, the goal is not merely syntactic or probabilistic continuation, but semantic or grammatical closure under global constraints. The SAT results therefore serve as an existence proof for a broader class of Glass-Box representational systems.

From Search to Structural Closure. Perhaps the most significant theoretical implication is a shift in perspective. Traditional approaches treat resolution as a combinatorial search problem. The present results suggest an alternative view: resolution as a process of representational closure. From this standpoint, certain forms of exponential difficulty may arise not from intrinsic intractability, but from a misalignment between the problem's global structure and the solver's local reasoning framework. Structural linearization, as demonstrated here, provides a different axis along which such difficulties can be addressed.

Summary. In summary, these results validate the QDE/R3-MRM architecture as a non-trivial, implemented system capable of deterministically and audibly linearizing complex global constraints. While limited in scope by design, the work establishes a foundation for applying the same principles to broader domains, including formal mathematics and language, where closure, auditability, and governance are critical requirements.

11 Conclusion

We have presented a Glass-Box structural resolution engine that robustly closes a class of non-local SAT instances while producing auditable proofs. This demonstrates that explainability, governance, and robustness are achievable even in NP-complete settings, and opens the path toward certified reasoning systems in broader domains.

Appendix A — Detailed Experimental Results and Baseline Comparison (Global XOR-SAT)

A.1 Problem Class

The experiments reported in this appendix concern the following fixed problem class:

- Class: **SAT_global / XOR-CNF**
- Description:
  - Parity (XOR) constraints encoded in conjunctive normal form (CNF)
  - Global, non-local constraint structure
- Instance types:
  - SAT
  - UNSAT
- Sizes tested:
  - $n = 400$
  - $n = 800$
  - $n = 1600$
- Number of instances per size:
  - 5 (3 SAT, 2 UNSAT)

A.2 Experimental Environment

All experiments were conducted under the following fixed conditions:

- Machine:
  - CPU: recorded in `env_manifest.txt`
  - RAM: recorded in `env_manifest.txt`
- Operating system:
  - Linux (Ubuntu), version recorded in `env_manifest.txt`
- Execution mode:
  - Single-threaded
- Seeds:
  - Fixed and recorded per instance
- Timeouts:
  - Kissat: 20 seconds per instance
- Software versions:
  - Python: recorded in `env_manifest.txt`
  - Kissat: recorded in `env_manifest.txt`

A.3 R3-MRM Results — Per-Instance Details

Tables 1, 2, and 3 report detailed results for all R3-MRM runs, with one row per instance.

Instance ID	Size	SAT/UNSAT	Steps	Scales	Runtime (s)	$\hat{\rho}$
sat_01.cnf	400	SAT	36	35	94.93	0.9578
sat_02.cnf	400	SAT	38	37	94.27	0.9573
sat_03.cnf	400	SAT	36	35	95.98	0.9571
unsat_01.cnf	400	UNSAT	35	34	93.51	0.9544
unsat_02.cnf	400	UNSAT	32	31	85.61	0.9836

Table 1: R3-MRM detailed results for XOR-CNF instances at size $n = 400$.

Size $n = 400$

Instance ID	Size	SAT/UNSAT	Steps	Scales	Runtime (s)	$\hat{\rho}$
sat_01.cnf	800	SAT	74	73	404.79	0.9895
sat_02.cnf	800	SAT	70	69	330.52	0.9914
sat_03.cnf	800	SAT	66	65	368.66	0.9894
unsat_01.cnf	800	UNSAT	75	74	417.33	0.9859
unsat_02.cnf	800	UNSAT	68	67	373.93	0.9893

Table 2: R3-MRM detailed results for XOR-CNF instances at size $n = 800$.

Size $n = 800$

Instance ID	Size	SAT/UNSAT	Steps	Scales	Runtime (s)	$\hat{\rho}$
sat_01.cnf	1600	SAT	146	145	1774.79	0.9889
sat_02.cnf	1600	SAT	141	140	1670.49	0.9951
sat_03.cnf	1600	SAT	148	147	1952.69	0.9894
unsat_01.cnf	1600	UNSAT	146	145	2296.40	0.9888
unsat_02.cnf	1600	UNSAT	143	142	2373.07	0.9892

Table 3: R3-MRM detailed results for XOR-CNF instances at size $n = 1600$.

Size $n = 1600$

A.3bis Proof Artifacts Produced per Run

For each R3-MRM execution reported in Tables 1–3, the solver produced the following artifacts:

- a machine-verifiable certificate (assignment for SAT instances, contradiction object for UNSAT instances),
- a complete, ordered event trace recording every resolution step,
- a residual trajectory $\mu(t)$ enabling independent verification of contraction,
- deterministic replay hashes allowing full re-execution consistency checks.

These artifacts were generated systematically for all runs, without exception, and are available for independent audit under the same experimental protocol.

A.4 CDCL Baseline Results — Kissat

Tables 4, 5, and 6 report Kissat results on the same instances, under a fixed timeout of 20 seconds per instance.

Instance ID	Size	Result	Runtime (s)
sat_01.cnf	400	TIMEOUT	20.0
sat_02.cnf	400	TIMEOUT	20.0
sat_03.cnf	400	TIMEOUT	20.0
unsat_01.cnf	400	UNSAT	0.0236
unsat_02.cnf	400	UNSAT	0.0130

Table 4: Kissat results for XOR-CNF instances at size $n = 400$.

Size $n = 400$

Instance ID	Size	Result	Runtime (s)
sat_01.cnf	800	TIMEOUT	20.0
sat_02.cnf	800	TIMEOUT	20.0
sat_03.cnf	800	TIMEOUT	20.0
unsat_01.cnf	800	UNSAT	0.0202
unsat_02.cnf	800	UNSAT	0.0201

Table 5: Kissat results for XOR-CNF instances at size $n = 800$.

Size $n = 800$

Instance ID	Size	Result	Runtime (s)
sat_01.cnf	1600	TIMEOUT	20.0
sat_02.cnf	1600	TIMEOUT	20.0
sat_03.cnf	1600	TIMEOUT	20.0
unsat_01.cnf	1600	UNSAT	0.0381
unsat_02.cnf	1600	UNSAT	0.0371

Table 6: Kissat results for XOR-CNF instances at size $n = 1600$.

Size $n = 1600$

A.5 Synthetic Comparison

A.6 Factual Observations

- All R3-MRM runs terminated with closure.
- No R3-MRM run exceeded the step or scale budget.

Size	R3-MRM Closure Rate	Kissat Closure Rate	Kissat Timeout Rate
400	100%	40%	60%
800	100%	40%	60%
1600	100%	40%	60%

Table 7: Aggregate comparison between R3-MRM and Kissat on XOR-CNF instances.

- Kissat failed to close all SAT instances at all tested sizes, resulting in systematic timeouts on satisfiable cases.

A.7 Data Availability

All instances, logs, and public verification traces are available upon request or via a controlled API under the same experimental protocol.

Metadata

- **Author:** Simone Mazzoni
- **Affiliation:** RCUBEAI Research Lab, Paris, France
- **ORCID:** 0009-0004-0400-4823

- **Title:** Structural Resolution of Non-Local SAT Instances

- **Version:** v1.0
- **Date:** January 2026
- **Type of Paper:** Foundational Theoretical Paper
- **Validation Status:** Issued (Author-Issued)
- **DOI:**

- **Keywords:** Representation Models, explicit representation, controlled reasoning, non-closure, hierarchical computation, AI governance

How to Cite

Mazzoni, S. (2026). *Structural Resolution of Non-Local SAT Instances.* Foundational architectural paper. Version v1.0. RCUBEAI Research Lab. License: CC BY 4.0.