

Explaining SAT Solving Using Causal Reasoning

Jiong Yang ✉

National University of Singapore, Singapore

Arijit Shaw ✉

Chennai Mathematical Institute, India

IAI, TCG-CREST, Kolkata, India

Teodora Baluta ✉

National University of Singapore, Singapore

Mate Soos

National University of Singapore, Singapore

Kuldeep S. Meel

National University of Singapore, Singapore

Abstract

The past three decades have witnessed notable success in designing efficient SAT solvers, with modern solvers capable of solving industrial benchmarks containing millions of variables in just a few seconds. The success of modern SAT solvers owes to the widely-used CDCL algorithm, which lacks comprehensive theoretical investigation. Furthermore, it has been observed that CDCL solvers still struggle to deal with specific classes of benchmarks comprising only hundreds of variables, which contrasts with their widespread use in real-world applications. Consequently, there is an urgent need to uncover the inner workings of these seemingly *weak* yet *powerful* black boxes.

In this paper, we present a first step towards this goal by introducing an approach called *CausalSAT*, which employs causal reasoning to gain insights into the functioning of modern SAT solvers. *CausalSAT* initially generates observational data from the execution of SAT solvers and learns a structured graph representing the causal relationships between the components of a SAT solver. Subsequently, given a query such as whether a clause with low literals blocks distance (LBD) has a higher clause utility, *CausalSAT* calculates the causal effect of LBD on clause utility and provides an answer to the question. We use *CausalSAT* to quantitatively verify hypotheses previously regarded as “rules of thumb” or empirical findings, such as the query above or the notion that clauses with high LBD experience a rapid drop in utility over time. Moreover, *CausalSAT* can address previously unexplored questions, like which branching heuristic leads to greater clause utility in order to study the relationship between branching and clause management. Experimental evaluations using practical benchmarks demonstrate that *CausalSAT* effectively fits the data, verifies four “rules of thumb”, and provides answers to three questions closely related to implementing modern solvers.

2012 ACM Subject Classification Theory of computation; Computing methodologies → Artificial intelligence

Keywords and phrases Satisfiability, Causality, SAT solver, Clause management

Digital Object Identifier 10.4230/LIPIcs.SAT.2023.28

Related Version *Full Version*: <https://arxiv.org/abs/2306.06294>

Supplementary Material *Software (Source Code)*: <https://github.com/meelgroup/causalsat>

Funding This work was supported in part by National Research Foundation Singapore under its NRF Fellowship Programme [NRF-NRFFAI1-2019-0004], Ministry of Education Singapore Tier 2 Grant [MOE-T2EP20121-0011], and Ministry of Education Singapore Tier 1 Grant [R-252-000-B59-114]. Teodora Baluta is also supported by the Google PhD Fellowship.



© Jiong Yang, Arijit Shaw, Teodora Baluta, Mate Soos, and Kuldeep S. Meel;
licensed under Creative Commons License CC-BY 4.0

26th International Conference on Theory and Applications of Satisfiability Testing (SAT 2023).

Editors: Meena Mahajan and Friedrich Slivovsky; Article No. 28; pp. 28:1–28:19

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Acknowledgements We are thankful to Tim van Bremen for providing detailed feedback on the early drafts of the paper and grateful to the anonymous reviewers for their constructive comments to improve this paper. We extend sincere gratitude to Jakob Nordström and other participants of the program *Satisfiability: Theory, Practice, and Beyond* at the Simons Institute for the Theory of Computing for insightful discussions. Part of the work was done during Arijit Shaw’s internship at the National University of Singapore. The computational work for this article was performed on resources of the National Supercomputing Centre, Singapore, <https://www.nsc.sg>.

1 Introduction

Boolean Satisfiability (SAT) is a fundamental problem in computer science that involves determining whether there exists an assignment σ that satisfies a given Boolean formula F . The applications of SAT are vast and varied, including but not limited to bioinformatics [29], AI planning [22], and hardware and system verification [7, 11]. The seminal work of Cook [13] demonstrated that SAT is NP-complete. Unsurprisingly, early algorithmic methods such as local search and the DPLL paradigm [14] faced significant scalability challenges in practice. In the early ’90s, the introduction of Conflict Driven Clause Learning (CDCL) [45] ushered in a new era of interest from both theoreticians and practitioners. The outcome of this development is the emergence of effective heuristics that empower SAT solvers to handle challenging instances across various domains. This remarkable advancement is widely recognized as the *SAT revolution* [3, 6, 15, 25, 30–32, 45].

The modern CDCL SAT solvers owe their performance to well-designed and tightly integrated core components: *branching* [25, 30], *phase selection* [38], *clause learning* [4, 28], *restarts* [3, 18, 19, 24], and *learnt clause cleaning* [5, 34]. Consequently, the progress has been driven largely by the continuous improvement of heuristics for these core components. The annual SAT competition [21] has provided evidence of a pattern in which the development of heuristics for one core component necessitates and encourages the design of new heuristics for other components to ensure tight integration. Such progress, however, necessitated the solvers to be complex, whereas a typical modern SAT solver is made up of approximately 15 – 30 thousand lines of code and employs a wide variety of strategies to solve different types of problems. Therefore, to make further progress in the design of SAT solvers, it is crucial to understand the inner workings of these SAT solvers.

Traditional complexity-theoretic studies focus on analyzing the limitations of SAT solvers: such as determining the explanations for why SAT solvers struggle for certain instances [16]. These approaches, however, fail to inform reasons for their success in solving industrial-size instances. Several prior studies provide intuitive hypotheses on why certain SAT-solving heuristics work well on some benchmarks [34, 48, 53]. These hypotheses are mostly derived from the researcher or developer’s intuition and empirically tested in an ad-hoc manner. The tight coupling of heuristics has been a crucial part of the development of modern solvers [47, 49]. For instance, Biere et al. [9] found that certain branching heuristics work particularly well in combination with specific restart heuristics. Similarly, Oh [33] observed that in order to achieve the best solver performance, the decay factor used in activity calculations must be adjusted according to the active restart heuristics.

Motivated by the role played by empirical studies in the development of new heuristics, we take the approach of usage of data-driven methods to develop understandings of the behavior of SAT solvers. Our work relies on the CrystalBall framework [48] that provides white-box access to the execution of a state-of-the-art SAT solver, thereby enabling large-scale data collection. Given the recent development in the field of causal reasoning [35], we investigate

whether it is possible to develop a framework that may use white-box access to the execution of SAT solving to generate causal reasoning to explain the interplay among the features and heuristics of SAT solving?

Several studies have investigated the behavior of modern solvers. Elffers et al. [16] explored the effectiveness of various solver components by employing a large set of theoretical benchmarks, with the running time serving as the parameter to evaluate their utility. Simon [46] focused on analyzing the generated clauses' utility by examining the proofs produced by a SAT solver. Kokkala et al. [23] utilized DRAT-trim-based proofs to investigate the impact of restart and learned clause cleaning heuristics on the length of trimmed proofs. They also investigated the relationship between different parameters and the probability of a clause being included in the trimmed proof. Soos et al. [48] employed CrystalBall to deduce the statistical correlation between the utility of learned clauses and SAT-solving features. Nevertheless, all of these studies were based on correlation analysis and lacked a causal reasoning perspective in understanding solver behavior.

In this work, we propose to use *causal reasoning* to answer what factors influence the utility of clauses in memory management in SAT solvers. Our approach combines prior knowledge and experimental data to derive a causal model, a graph whose nodes represent different components and heuristics of the solver, and its edges represent causal relations between them. The derived causal model explicitly encodes assumptions about the underlying phenomenon in clause memory management. The causal model allows us to query it by applying a set of principled rules known as do-calculus to compute the effect of a particular factor on the utility of clauses. These rules ensure that each query computes the right statistical quantity given the model without introducing bias in the estimates. We find that our derived models fit the data reasonably well even when less data is available. Using our framework, we were able to confirm several well-known hypotheses about solvers, which highlights its potential and assures its correctness. Additionally, we have used the framework to inquire about some unresolved questions related to the solving process, and it has provided us with answers to those inquiries. This highlights that our framework provides a fresh perspective on how to approach and analyze solvers.

We demonstrate the application of our approach to study four hypotheses that stem from observations or assumptions in prior work and encode these as causal queries. We find that clauses with lower Literal Block Distance (LBD) have greater utility; thus, lower LBD clauses should be prioritized in clause memory management. We find that small clauses have greater utility, regardless of their LBD. We also justify the hierarchical clause memory management that high LBD clauses should be cached for a short time since their utility experiences a rapid drop over time. Moreover, we verify that LBD has a larger causal impact than the size on clause utility, and hence LBD is a preferable choice for prioritizing clauses in memory management. We additionally study three novel questions that have not been formulated before, such as which branching heuristic results in a greater clause utility, to study the relationship between branching and clause management. We find that as a branching heuristic, Maple [24] leads to a greater clause utility than VSIDS [31]. To summarize, we demonstrate that causal reasoning allows one to evaluate the crucial choices in the design of SAT solvers.

2 Background and Motivation

2.1 CDCL Solvers

CDCL-based SAT solvers begin with an initially empty set of assignments, maintaining a partial assignment at each step. The solver incrementally assigns a subset of variables until the current partial assignment is unable to satisfy the current formula. At this point,

the solver employs a backtracking mechanism to trace the cause of unsatisfiability, which is expressed as a conflict clause. Modern solvers frequently perform restarts, resetting the partial assignment to empty.

Modern SAT solvers may run for millions of conflicts during a single execution, with each conflict resulting in the learning of one or more clauses that are subsequently added to a database. As this database grows over time, it becomes necessary to periodically remove some of the less useful clauses to maintain solver efficiency. To accomplish this, various heuristics have been proposed that predict the usefulness of learned clauses. These heuristics, collectively known as *learned clause cleaning* heuristics, enable the solver to identify and remove unnecessary clauses, improving performance and scalability. A *branching heuristic* is employed to determine which variable to select for each decision point in the search process. A *restart heuristic* is used to determine when the solver should restart the search process from the beginning.

2.2 Learnt Clause Cleaning

One of the most crucial parts of a CDCL-based SAT solver is the learning of clauses, which occurs rapidly. In fact, a typical SAT solver can learn millions of clauses within a minute of operation. However, the sheer volume of these clauses can hinder the solver’s efficiency. To address this issue, *learnt clause cleaning* heuristics were developed. These heuristics evaluate the “utility” of a clause based on some parameters and determine whether to retain or discard it. A heuristic may prioritize keeping clauses with higher utility while discarding those with lower utility to optimize the solver’s performance. Previous literature proposed various parameters to determine the usefulness of a clause. Een et al. [15] employed the notion of *activity (or VSIDS)* as a surrogate for utility. Activity is a parameter that indicates a clause’s involvement in recent conflicts. On the other hand, Audemard et al. [3] utilized *literal block distance (LBD)* as a proxy for utility, where LBD represents the number of unique decision levels in the clause. Despite its success, it is difficult to tell *why* LBD is a good indicator of utility and disentangle its impact from other solving heuristics. To illustrate this, we highlight some of the existing hypotheses along with novel questions about the utility of clauses below.

- Clauses with small LBD have greater utility.
- Small clauses have greater utility.
- High-LBD clause experiences a rapid drop in clause utility over time.
- LBD has a greater impact on clause utility than clause size.
- What factor, other than LBD, size, and activity, has the greatest impact on clause utility?
- Which branching heuristic results in the greatest clause utility?
- Which restart heuristic results in the greatest clause utility?

It is thus natural to ask if these hypotheses and questions explain the underlying SAT-solving behavior well. Do they provide new insights into other aspects of solving? We show that in order to correctly answer such questions, we require a principled framework to even compute the right statistical quantities from data.

2.3 Causal Model

A causal model represents the set of features that cause an outcome and the function that measures the effect of the features on the outcome. Prior work proposes prior domain knowledge such as causal invariant transformations [51], score-based learning algorithms [42], interventional causal representation learning [1, 10, 39] or learning invariant relationships from training datasets from different distributions [2, 37] to construct causal models.

As with probabilistic graphical models, the underlying representation of the causal model is (1) a directed acyclic graph (DAG) $G = (V, E)$ where the vertices V are the set of random variables $V = \{X_1, \dots, X_n\}$ and the edges $(X_i, X_j) \in E$ represent causal relationships X_i causes X_j and (2) set of parameters Θ quantifying the relationships between the variables in V . The set of variables V can take either discrete or continuous values. We say that a subset of variables $S \subset V$ cause X_i , if $X_i = g(\{X_j\}_{j \in S}, N_Y)$ for some function g and N_Y a random variable independent of all other variables. The set of functions is parameterized by $\theta \in \Theta$. Henceforth, we will use causal models and causal graphs interchangeably. The absence of an edge from X to Y represents (conditional) independence. Due to this, causal graphs are also known as causal Bayesian networks. Causal graphs still have the probabilistic interpretation of their associative (Bayesian) counterpart: they represent a joint probability distribution of the variables in V , determined by their conditional independence relationships, i.e., $\mathcal{P}_\theta(V) = \prod_i \Pr[X_i | pa_{X_i}]$, $X_i \in V$, where pa_{X_i} represents the parents of the X_i vertex.

Causal Inference

The causal graph has a quantitative interpretation: we use it to answer *do-queries* that estimate the causal effect between a *treatment* variable and an *outcome* variable.

do-query. We say there is a causal effect between a treatment variable X_i and an outcome variable Y if under an intervention on the treatment variable X_i , the outcome variable changes. Given a set of variables $V = \{X_1, \dots, X_n, Y\}$, an intervention on a variable X_i is an experiment where the experimenter controls the variable $X_i \in S$ to take a value u of another independent (from other variables $\in V$) variable, i.e., $X_i = u$. More generally, we can intervene on a subset of variables $S \subset V$. This operation has been formalized as the *do* operator by Pearl [35]. The *do* operator thus changes the value of X_i while keeping every other variable in V the same, except for those directly or indirectly affected by X_i . This is akin to removing the edges of the nodes in S to their parents in the graph G , resulting in a manipulated graph. The intervention effect on the outcome variable, e.g., $\Pr[Y | do(X_i = u)]$, is then measured by data generated under the distribution of the manipulated graph. *Do-queries* are thus fundamentally different from observational or conditional queries in that it requires being able to set the value of the variable X_i to a potentially unobserved value and keeping every other unaffected variable in V the same. The framework to reason about *do-queries* has been formalized in a set of rules known as the *do-calculus* which has been shown to be sound and complete [20].

Average Treatment Effect (ATE). Given a causal graph, a treatment variable X , and an outcome variable Y , the average treatment effect measures the change of the expected value of the outcome variable when we intervene on the treatment variable and change it from a constant value b to a .

► **Definition 1** (Average Treatment Effect). *The average treatment effect of a variable X (called the treatment) on the target variable Y (called the outcome) is:*

$$ATE(X, Y, a, b) = \mathbb{E}[Y | do(X = a)] - \mathbb{E}[Y | do(X = b)],$$

where a, b are constants for which X is defined. We omit the constants when the query is over the domain of X . Following Definition 1, the conditional average treatment effect (CATE) measures the average outcome conditioned on a variable W of interest.

► **Definition 2** (Conditional Average Treatment Effect). *The conditional average treatment effect of X on Y conditioned on a variable of interest W is:*

$$CATE(X, Y, W, a, b) = \mathbb{E}[Y|do(X = a), W] - \mathbb{E}[Y|do(X = b), W].$$

Why Causal Analysis?

Deriving the right quantities to compute is challenging without a proper framework. As an example, we first consider `LBD`, `Propagation`, i.e., the number of times a clause is involved in propagation, and `LastTouch`, i.e., the number of conflicts since a clause is involved in conflict analysis. We collect data corresponding to these variables during the SAT-solving process of 80 instances by varying several of the existing heuristics. Suppose we want to compute the impact of `Propagation` on `LastTouch`, that is, the expected change of `LastTouch` if `Propagation` increases by one. This quantity can be written as the expected value of `LastTouch` when `Propagation` increased from one to two: $\mathbb{E}[\text{LastTouch} \mid \text{Propagation} = 2] - \mathbb{E}[\text{LastTouch} \mid \text{Propagation} = 1] = -431.29$.

However, there is a common cause, `LBD`, that affects both `Propagation` and `LastTouch`. Having this knowledge at hand, we need to “control” for `LBD`. “Controlling” for a common cause in statistical measurements means binning over the observed values of the common cause. Thus, the right quantity to compute is: $\sum_z (\mathbb{E}[\text{LastTouch} \mid \text{Propagation} = 2, \text{LBD} = z] - \mathbb{E}[\text{LastTouch} \mid \text{Propagation} = 1, \text{LBD} = z]) \Pr[\text{LBD} = z]$. Notice that we can represent such relationships between variables using a *causal model*, i.e., a graph where the nodes are variables and the edges represent causal relations. We show the graph corresponding to our example in Figure 1a. Estimating this statistical quantity using a linear regression model that fits the data yields -2.92 . Thus, our first attempt at estimating the effect did not model the underlying process precisely.

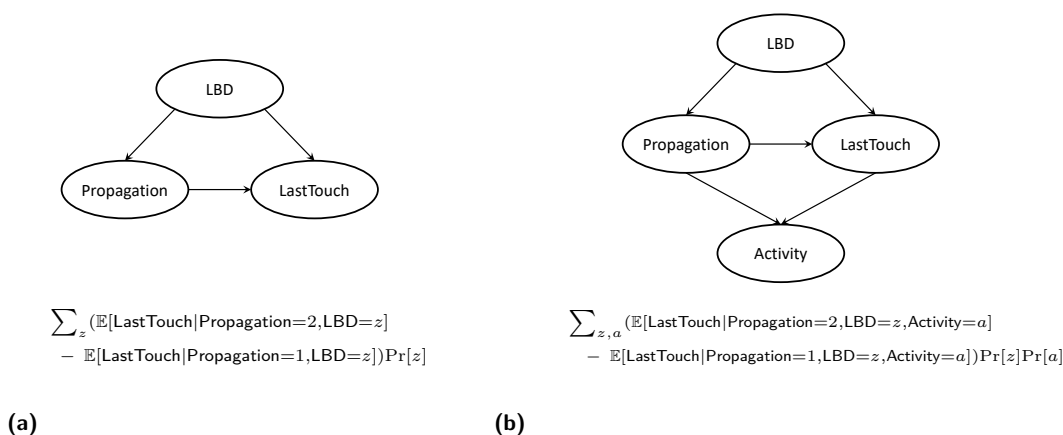
Let us now consider a fourth variable, `Activity` that is also correlated with `Propagation` and `LastTouch` (see Figure 1b). Should we control for this variable as well? Suppose we controlled for it and estimated again the effect of `Propagation` on `LastTouch` as the following:

$$\sum_{z,a} (\mathbb{E}[\text{LastTouch} \mid \text{Propagation} = 2, \text{LBD} = z, \text{Activity} = a] - \mathbb{E}[\text{LastTouch} \mid \text{Propagation} = 1, \text{LBD} = z, \text{Activity} = a]) \Pr[\text{LBD} = z] \Pr[\text{Activity} = a]$$

On the collected data, this results in an estimated effect of -1.96 . However, `Activity` should not be controlled for since it is not a common cause, but rather it is affected by both `Propagation` and `LastTouch`. Intuitively, during SAT solving, the `Activity` is computed *after* these two variables. The correct estimate is $\sum_z (\mathbb{E}[\text{LastTouch} \mid \text{Propagation} = 2, \text{LBD} = z] - \mathbb{E}[\text{LastTouch} \mid \text{Propagation} = 1, \text{LBD} = z]) \Pr[\text{LBD} = z] = -2.92$. Over-controlling biased the estimated effect. It is thus crucial to identify the correct set of variables to control.

3 Causal Reasoning for SAT Solving

`LBD`, `size`, and `activity` have long been employed to predict the utility of learned clauses in SAT solvers. For instance, Kissat [8] consistently retains clauses with $\text{LBD} \leq 2$, while CryptoMiniSat [49] always preserves clauses with $\text{LBD} \leq 3$. These heuristics constitute a critical component and substantially enhance the performance of modern SAT solvers. Despite being considered a “rule of thumb” for solver design, whether a low-`LBD` clause yields higher utility than a high-`LBD` clause still remains uncertain.



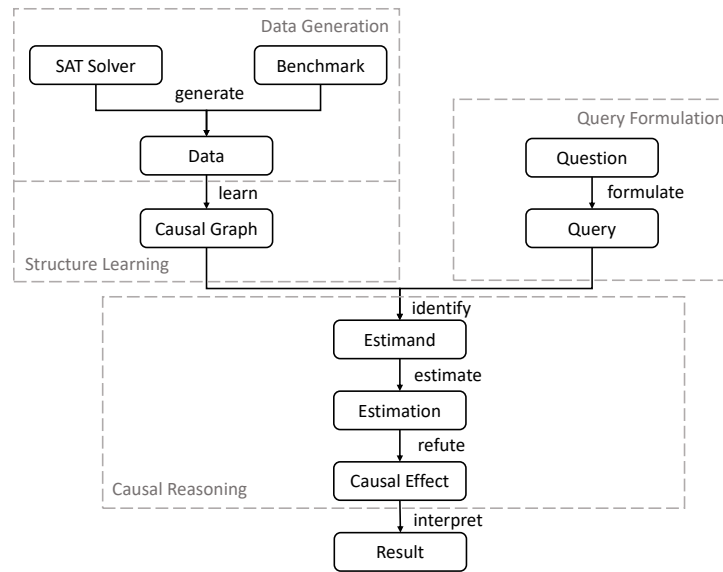
■ **Figure 1** Causal models help correctly identify the set of variables to control for. (a) Estimated effect of **Propagation** on **LastTouch** when **LBD** is a cause of both. (b) Over-controlling while estimating the effect of **Propagation** on **LastTouch** when **Activity** is introduced. We use the notation $\Pr[z]$ and $\Pr[a]$ as a shorthand for $\Pr[\text{LBD} = z]$ and $\Pr[\text{Activity} = a]$, respectively.

To better comprehend the functionality of these factors and promote the discovery of new factors, we propose a causality-based approach to examine the influence of these factors on the utility of learned clauses. Our objective is to use this method to address queries closely tied to the implementation of modern SAT solvers, such as whether a high or low LBD clause offers greater utility and what other factors significantly impact clause utility. The findings can then be leveraged to inform the design of SAT solvers.

Overview

We introduce an approach designed to address SAT-related inquiries. Figure 2 displays our approach’s overview, while Algorithm 1 depicts the pseudo-code for our prototype. Initially, we generate observational data from a SAT solver at Line 1. This data records factor values and clause utility (formally defined in Section 3.1). Subsequently, we construct a causal graph representing the causal relationship between all variables, including factors and clause utility, Line 3, based on our prior knowledge and observational data. Meanwhile, we formulate each SAT-related question into a causal query at Line 4. For instance, the LBD-related question can be formulated as a computation of the causal effect from LBD to clause utility. From Lines 5 to 9, we calculate a causal effect estimate. First, we identify a mathematical expression for the estimation known as an *estimand*. Next, the estimand is evaluated on observational data to calculate the causal effect estimate. If the resulting estimate passes the refutation test, it is returned as the final causal effect; otherwise, a failure is reported. Lastly, we return the causal effect corresponding to the original question.

In the subsequent subsections, we will delve into each component depicted in Figure 2. Section 3.1 outlines the process of generating observational data from a SAT solver. Section 3.2 elaborates on the construction of a causal graph using our prior knowledge of SAT solvers and the generated data. Section 3.3 introduces the SAT-related questions of interest and describes their formulation into a causal query. Lastly, Section 3.4 demonstrates the computation of the causal effect for a given query.



■ **Figure 2** Our approach overview, from data generation to the causal estimate.

■ **Algorithm 1** Our approach (CausalSAT).

```

1: data ← generate observational data from a SAT solver;
2: whitelist, blacklist ← decode predefined and blocked edges from prior knowledge;
3: causalGraph ← HillClimbing(whitelist, blacklist, data);
4: query ← formulate a question;
5: estimand ← Identify(query, causalGraph)
6: estimates ← Estimate(estimand, causalGraph, data);
7: passRefutation ← Refute(estimates, query, causalGraph, data);
8: if passRefutation then return estimates;
9: else return Failure;

```

3.1 Data Generation

To accomplish our goal of exploring causality and assess the utility of a clause, we collected data that included traces of the SAT solver’s execution. These traces captured various characteristics of a clause at different stages of the instance-solving process, as well as the future utility of the clause. We work with a set of unsatisfiable instances. For each instance, we run the SAT solver multiple times with different combinations of heuristics. To complete the data collection procedure, we employ a two-pass system. During the *forward pass*, we run the SAT solver until it reaches the conclusion of unsatisfiability. To gather data on all the learned clauses, we deactivated the learned clause cleaning heuristic during the execution of the SAT solver. Consequently, all the learned clauses were retained throughout the solver’s execution. We store data about all the learned clauses at different snapshots of solving. However, this data lacks labeling about whether each clause was useful in proof generation. To address this issue, we execute a *backward pass* using the proof generator DRAT-trim. In this pass, we label all the data collected in the forward pass by how many times the clause will be used in the next 10k conflicts. By combining the data from both passes, we obtain a complete dataset for our study. As a result, each data point in our dataset includes information about the instance, the heuristics used, the learned clauses, some of its features, and the usefulness of each clause in proof generation.

Features and Heuristics. In general, the accuracy of a causal model is often improved with a larger number of features and a larger training dataset. However, many algorithms for detecting causal structures are limited in their ability to handle a large number of features. Therefore, we have selected a subset of important features that are commonly used in modern SAT solver literature. In Table 1, we list all the features we have in our data. In our study, we adopted the definition of clause utility as proposed in [48], which relies on the frequency of clause usage observed in the proofs reconstructed by DRAT-trim [52]. It should be noted that the proof generated by DRAT-trim may differ significantly from the proof generated by the SAT solver, necessitating caution in interpretation.

■ **Table 1** Data collected about clauses during solving.

Branching	Heuristics to determine the order in which variables are assigned values during the search. We consider VSIDS [31] and Maple [24,26].
Restart	Heuristics to determine when the solver should restart. We consider Geometric [50], LBD-based [4], and Luby [19,27] heuristics.
Size	number of literals in a clause.
LBD	number of distinct decision levels of literals in a clause.
Activity	measure of clause’s importance in the search process, based on the clause’s involvement in recent conflicts.
UIP	number of times that the clause took part in a 1st-UIP conflict generation since its creation.
Propagation	number of times the clause was used in propagations.
LastTouch	number of conflicts since the clause was used during a 1st-UIP conflict clause generation.
Time	number of conflicts since the generation of this clause.
Utility	within the next 10,000 conflicts, the number of times this clause has been used in DRAT proof generation. The number is weighted based on at which points the clauses are used.

3.2 Structure Learning

A causal graph can be built following our prior knowledge. For example, the clause management heuristic is a known cause for the average LBD of learned clauses. If we preserve the clause of small LBD but remove that of large LBD during clause reduction, the average LBD will go down. Therefore, we can add to the causal graph an edge from clause management heuristic to average LBD to represent the known causality between them. For another instance, both branching and restart heuristics are user-defined parameters, and therefore no other variables inside the SAT solver can cause them. We can then block incoming edges for both heuristics. Overall, we add the following constraints based on our prior knowledge:

- No incoming edges are allowed for **Branching** and **Restart** since they are given by users.
- No incoming edges are allowed for **Time** because time ticks are fixed in our experiments.
- No incoming edges allowed for **LBD** and **Size** except for edges from **Branching** and **Restart** because **LBD** and **Size** are determined before other variables.
- No outgoing edges are allowed for **Utility** because **Utility** follows other variables.

However, prior knowledge is insufficient to build a complete causal graph in many scenarios because the knowledge we have about the system is always incomplete, like a SAT solver. Hence, a complete causal graph has to be learned from the observational data. We use the hill-climbing algorithm [44] in our approach to learn a causal graph for its practical performance. The algorithm starts from an initial directed graph. In each step, we randomly add, remove, or reverse an edge of the graph to form candidate graphs. A score function called Bayesian Information Criterion [40], which favors a graph with a simple structure and lower log-likelihood over the data, is then applied to these candidate graphs to select a locally optimal graph. The process continues until a fixed point where we do not find a better graph.

Furthermore, we apply k -fold cross-validation to the learning process to lower the variance of the returned graph. The dataset is initially partitioned into k subsets. Subsequently, for each iteration, we learn a graph on $k - 1$ subsets and evaluate its fitness over the remaining subset. Finally, we obtain an averaged graph by taking a majority vote from the k candidates on the existence and the direction of each edge.

There are other structure-learning algorithms available such as the PC Stable algorithm [12]. However, the PC Stable algorithm always returns a graph with a few undirected edges in our experiments, which means the algorithm failed to infer the causal direction between some variables. Due to its algorithmic nature, the hill-climbing algorithm always returns a directed graph, meeting our requirements. A detailed comparison between structure-learning algorithms is out of the scope of this paper, and we refer the interested readers to [41]. Additionally, we present a pseudo-code of the hill-climbing algorithm in Appendix A [54].

3.3 Queries on SAT Solving

Are you curious about whether a clause with a high or low LBD possesses a greater utility? Alternatively, you might be interested in determining which factor – size or LBD – exerts a more significant influence on clause utility. Previously, one would have to rely on personal experience or experiment with heuristics to reach a conclusion. However, with our proposed method, you simply need to transform these inquiries into causal queries, and our causal model will provide the answers.

In this section, we initially present the different query types supported by our approach. Following that, we explore questions closely related to SAT solver implementation and demonstrate the process of formulating them into causal queries. Our primary focus lies on queries concerning clause utility, with the intention of examining the variables that exert a substantial influence on clause utility.

Query Type

Our causal framework accommodates three distinct query categories, as detailed below:

- Average Treatment Effect (ATE) quantifies the variation in the outcome variable when the treatment variable shifts from one value to another.
- Conditional Average Treatment Effect (CATE) encapsulates the ATE, considering a condition for an additional variable.
- We propose the Averaged Conditional Average Treatment Effect (ACATE) concept to account for the need to average the CATE across all values of the conditional variable.

► **Definition 3** (Averaged Conditional Average Treatment Effect).

$$ACATE(X, Y, W, a, b) = \sum_w (\mathbb{E}[Y|do(X = a), W = w] - \mathbb{E}[Y|do(X = b), W = w]) \Pr[W = w]$$

LBD. In modern SAT solvers, LBD serves as a key factor for clause memory management heuristics after being introduced in Glucose [4]. It is assumed that the clause of small LBD has greater utility and is therefore favored during clause reduction. The assumption drives us to ask the question: which clause, with low or high LBD, has greater utility? The question can be translated to calculate the causal effect from LBD to clause utility, corresponding to an ATE query (Q1) in Table 2. The ATE measures the average change of the outcome (clause utility) if the treatment (LBD) changes from one to two. Following the assumption that a clause of small LBD has greater utility, increasing LBD will decrease the utility, and hence the ATE is expected to be negative. If the ATE is indeed negative, we prove the assumption. Otherwise, if ATE is positive, we reach the opposite conclusion that the clause of large LBD has greater utility. Finally, if ATE equals zero, the LBD has no effect on clause utility, and clauses with large or small LBD have the same utility.

■ **Table 2** Causal queries for questions on SAT solving.

Question	Query
Q1 Which clause, with low or high LBD, has greater utility?	$ATE(LBD, Utility, 2, 1) < 0$
Q2 Which type of clause, large or small, has greater utility? What if the LBD is fixed?	$\begin{cases} ATE(Size, Utility, 2, 1) < 0 \\ ACATE(Size, Utility, LBD, 2, 1) > 0 \end{cases}$
Q3 Which clause, with low or high LBD, experiences a rapid drop in utility over time?	$\begin{cases} CATE(Time, Utility, LBD \leq 6, 10000, 0) \geq 0 \\ CATE(Time, Utility, LBD > 6, 10000, 0) < 0 \end{cases}$
Q4 Which factor, size or LBD, has a greater impact on clause utility?	$ ATE(Size, Utility, 2, 1) > ATE(LBD, Utility, 2, 1) $
Q5 Which factor, besides size, LBD, and activity, has the greatest impact on clause utility?	$\arg \max_{Treatment \neq Utility} \{ ATE(Treatment, Utility, 2, 1) \}$
Q6 Which branching heuristic, VSIDS or Maple, results in a greater clause utility?	$ATE(Branching, Utility, Maple, VSIDS)$
Q7 Which restart heuristic, Geometric, LBD-based, or Luby, results in the greatest clause utility?	$\begin{cases} ATE(Restart, Utility, Luby, Geometric) \\ ATE(Restart, Utility, Luby, LBD-based) \\ ATE(Restart, Utility, Geometric, LBD-based) \end{cases}$

Size. Modern solvers assume a small clause has higher utility than a large clause, which brings us to the question: which type of clause, large or small, has a greater utility? The question is formulated into the first ATE query (Q2) in Table 2, which implies a decreasing utility when the clause size increases from one to two. If the resulting ATE is negative, a small clause is proved to have greater utility. Otherwise, a large clause has a greater utility if the ATE is positive. Lastly, the size does not affect utility if the ATE is zero.

A historical implementation of Glucose [17] prioritizing large clauses of the same LBD outperformed the one favoring small clauses, which implies that large clauses have greater utility given a fixed LBD. The observation is counter-intuitive because small clauses are usually assumed to be more useful, and people think it still holds when LBD is fixed.

Consequently, we ask the question: which type of clause, large or small, has greater utility given a fixed LBD? The question is formulated into an ACATE query (Q2) in Table 2, which implies a decreasing utility when clause size increases from one to two conditioned on a fixed LBD. The hypothesis will be verified if the ACATE is indeed positive or disproved otherwise.

LBD over Time. Modern SAT solvers adopt multi-tier clause memory management heuristic: more useful clauses are preserved for a long time, while less useful clauses are cached for a short time. Clauses are classified based on LBD. For example, Kissat [8] stores learned clauses with $LBD \leq 6$ in the long-time tier with the rest in the short-time tier. This observation implies that low-LBD clauses are expected to maintain their utility for a long time, while conversely, the utility of large-LBD clauses is expected to drop quickly over time.

The implication drives us to ask the question: Which clause, with low or high LBD, experiences a rapid drop in utility over time? Following that, we formulate the question into (Q3) in Table 2. The first query formulates that small-LBD clauses do not experience a decrease in utility when time ticks from 0 to 10,000, while the second one encodes that large-LBD clauses experience a utility drop over time. We model the time period as 10,000 conflicts because the data is collected per 10,000 conflicts. If the answers to two queries are both yes, the hypothesis is verified from a causal perspective.

Size vs. LBD. In modern SAT solvers, variants of clause management heuristics are based on size and LBD, but a question remains open: which has a greater impact on clause utility? We formulate the question into (Q4) in Table 2 to compare the causal effects on clause utility between size and LBD. If the size (resp. LBD) has a greater effect on clause utility, we suggest clause size (resp. LBD) be involved more in clause management. Note that we compare the effects over normalized data to eliminate the unfairness due to different orders of magnitude for size and LBD.

Beyond Size, LBD, and Activity. Modern SAT solvers consider size, LBD, and activity as the key factors in predicting clause utility, while other factors received less attention. To motivate new clause management heuristics, we utilize CausalSAT to identify new factors for predicting clause utility. As shown in (Q5) of Table 2, we formulate an optimization query to find the treatment that has the largest effect on clause utility. The discovery may inspire the design of new clause management heuristics.

Branching Heuristics. Modern SAT solvers can only deploy one branching heuristic at a time, which defers the hard choice among heuristics to users. In this work, we provide suggestions from a causal perspective. We ask the question: which branching heuristic results in a greater clause utility? We consider two prevalent heuristics, VSIDS [31] and Maple [24,26]. Following that, we formulated the question into an ATE query (Q6) in Table 2. The query evaluates the effect on clause utility when the branching heuristic changes from VSIDS to Maple. A positive ATE indicates Maple yields a higher utility, while a negative value implies VSIDS demonstrates a superior utility. Otherwise, the ATE equals zero, which implies the branching heuristic has no impact on clause utility.

Restart Heuristics. Considering three dominant restart heuristics Geometric [50], LBD-based [4], and Luby [19,27], we provide suggestions on the choice of restart heuristics by asking the question: Which restart heuristic, Geometric, LBD-based, or Luby, results in the greatest clause utility? Following that, we formulate the question into three ATE queries (Q7)

in Table 2, which compares the causal effects on clause utility among three heuristics. For example, if the first two ATEs are positive, Luby leads to higher utility than Geometric and LBD-based. Hence Luby is a preferred heuristic for high clause utility. Similar arguments apply to Geometric and LBD-based heuristics.

3.4 Causal Reasoning

In this section, we outline the process for calculating a query with the aid of a causal graph and observational data. Initially, we demonstrate the identification of controlled variables that bias the effect estimation. Following that, we exemplify the calculation of a causal effect using linear regression and the identified variables. Lastly, we discuss the utilization of refutation tests for validating the employed estimation technique.

3.4.1 Identification

A causal query asks to calculate the effect of treatment on the outcome (with some conditions). The naive way to do the calculation is to intervene in the treatment, i.e., change the treatment from one value to the other and observe the variation in outcome. However, intervening in treatment is infeasible in many applications like SAT solving, where it's impossible to increase the size of a clause and then observe the changes in utility. Alternatively, we can estimate the effect from the observational data. As motivated in Section 2.2, we have to identify a set of *controlled* variables to eliminate bias.

There are three methods to identify the controlled variables: backdoor, frontdoor, and instrumental variable identification. We use backdoor identification in our approach because frontdoor and instrumental variable identification always failed to find these variables in our experiments. Interested readers may refer to [36] for the details of other methods. The backdoor algorithm returns a set of variables called *backdoor set*, which is the smallest set separating the treatment and effect after removing the paths from the treatment to the outcome. These paths carry the pure effect from treatment to the outcome. A backdoor set represents a set of variables that exert an implicit effect on the treatment and outcome, such as the common cause LBD in Figure 1a. Controlling these variables ensures an unbiased estimation of the causal effect. Using a backdoor set, we can convert an ATE query into an expression called *estimand*, which can be evaluated over observational data to estimate the causal effect. Details of the backdoor algorithm are deferred to Appendix B [54].

3.4.2 Effect Estimation

We utilize linear regression as the estimator, given its effectiveness and interpretability. As a widely used estimator for causal effects, linear regression delivers remarkably strong performance. Its estimate is simple to interpret, signifying the average shift in the outcome when the treatment rises by one unit. Consequently, we demonstrate that the coefficient β corresponding to the treatment variable X signifies the ACATE value.

► **Lemma 4.** *Suppose $Z \cup W$ is a backdoor set and $\mathbb{E}[Y \mid X = x, Z = z, W = w]$ is linear, i.e., $\mathbb{E}[Y \mid X = x, Z = z, W = w] = \beta_0 + \beta_1 x + \beta_2^T z + \beta_3^T w$. We obtain*

$$\text{ACATE}(X, Y, W, a, b) = (a - b)\beta_1$$

We defer the proof to Appendix C [54] due to the page limit.

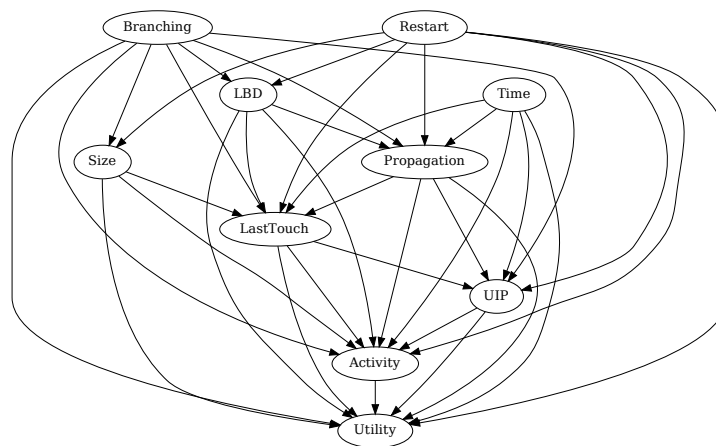
To identify a backdoor set that incorporates the conditional variable W , we can designate W as an element of the set and seek a valid backdoor set. In practice, we presume $\mathbb{E}[Y \mid X = x, Z = z, W = w]$ to be linear and employ $(a - b)\beta_1$ as an estimate of ACATE.

This conclusion is also applicable to ATE and CATE, with the proof detailed in Appendix [54]. To enhance confidence in statistical estimations, we conduct refutation tests to verify our estimation method. For example, we can assess whether the estimate changes substantially upon adding an independent random variable to the data as a common cause of the treatment and outcome. Interested readers may refer to Appendix D [54] for more details.

4 Experimental Evaluation

We implement a prototype called CausalSAT to evaluate our approach over practical instances. CausalSAT is built on top of a structure-learning package, bnlearn [41], and a causal-reasoning package, DoWhy [43]. We first present CausalSAT fits the dataset and then use it to answer the queries in Table 2.

For our experiments, we used a high-performance computing cluster. We collected traces of the run using instances from the SAT competition benchmarks. Because of the construction of our framework, we could only use unsatisfiable instances. Additionally, the process of collecting traces added significant overhead to the SAT-solving process, preventing us from gathering data from instances that took more than an hour to solve in a modern SAT solver or instances that could be solved within a few seconds. To ensure a variety of data, we chose to use 80 UNSAT instances from the SAT competition benchmarks from 2014-'18. We allowed a 2-hour timeout for all instances to run and record the traces. Afterward, we sampled data from those traces to ensure an equal representation of all types of instances. In total, we gathered 4 million data points. Lastly, we ran the hill-climbing algorithm using 10-fold cross-validation and obtained the causal graph in Figure 3.



■ **Figure 3** The causal graph CausalSAT derives using 10-fold cross-validation and 4 million points.

Summary. Our model fits the dataset, attaining a comparable loss to both linear models and decision trees while exhibiting a higher Pearson correlation with clause utility than any individual variable. With the model in place, we answer the questions in Table 2. For LBD, a small value leads to a larger clause utility, but large value results in a smaller utility that drops rapidly over time. For size, a large clause has greater utility, and the statement also holds when the LBD is fixed. Overall, the LBD has a greater impact than the size on clause utility. Besides LBD, size, and activity, the number of propagations has the greatest impact. Last but not least, Maple has a greater utility between branching heuristics, and Luby achieves the highest utility among restart heuristics.

4.1 Fitness on Data

We evaluate the fitness of CausalSAT on the dataset by showing a comparable prediction error and a stronger correlation with clause utility than any individual variable. We present the competence of CausalSAT with a linear model and decision tree by evaluating their performance based on the mean squared error (MSE) metric (the lower, the better). We further demonstrate that the Pearson correlation of CausalSAT with clause utility is higher than that of any individual variable.

CausalSAT delivers a comparable prediction error with a linear model and decision tree. We evaluate the mean squared error (MSE) of predicting clause utility for the three models. We adopt 10-fold cross-validation where the dataset is split into ten subsets, and every time the model is trained on nine subsets and evaluated on the remaining subset. The overall MSE is averaged over the 10 evaluations. Table 3 presents the MSE on different data sizes. When 0.5 million data points are available, CausalSAT obtains an MSE of 10050.41, which is smaller than both the linear model and decision tree. The result shows that CausalSAT still fits the data when we don't have much data available. On the other hand, the MSE of the decision tree decreases significantly as the data size increases, and the decision tree finally becomes the best predictor. Note that the prediction is not the focus of a causal model so it is expected to observe a smaller MSE from the decision tree when large data is available.

■ **Table 3** Mean squared error in 10-fold cross-validation.

Data Size	Linear Model	Decision Tree	CausalSAT
0.5 million	10087.79	13400.43	10050.41
1 million	9848.14	10731.84	9810.78
2 million	9986.05	7555.20	9951.33
4 million	9977.85	4335.32	9943.03

Additionally, we evaluate the correlation with Utility for any individual variable and CausalSAT. We compute the Pearson correlation coefficient between any individual variable and Utility except for Branching and Restart. Pearson correlation coefficient is a measure of linear correlation between two variables and has a value between -1 and 1, with 1 (resp. -1) indicating a strong positive (resp. negative) linear relationship and 0 suggesting no linear relationship between the variables. The Pearson correlation for CausalSAT measures the linear relationship between the Utility and the prediction generated by CausalSAT. As shown in Table 4, CausalSAT's prediction has a higher correlation with Utility than any individual variable. Consequently, any individual variable is not a good indicator for Utility while CausalSAT, considering all variables and their causal relationship, delivers a better correlation with clause utility.

■ **Table 4** Correlation with clause utility.

	Activity	LBD	LastTouch	Size	Propagation	UIP	Time	CausalSAT
Correlation	0.2819	-0.0285	-0.0326	-0.0283	0.2381	0.2557	0.0933	0.3425

4.2 Query Answering

We use CausalSAT to answer the questions in Table 2 by calculating the causal effects for the queries. Table 5 lists the calculations and answers, where the causal estimates all pass the three refutation tests. We use the original data to calculate ATE for a single variable due to its interpretability while using normalized data in Q4 and Q5 to compare ATEs between variables to eliminate the unfairness of their different orders of magnitude. We apply standard score normalization to the data. As follows, we explain the answers in detail:

- (A1) LBD has a negative effect on clause utility, indicating that clauses with low LBD have greater utility. Hence, a low-LBD clause should be prioritized over a high-LBD clause in clause memory management.
- (A2) Both ATE and ACATE values are negative, indicating that small clauses have greater utility, which also holds when LBD is fixed, though with a smaller effect. The answer to the fixed-LBD case contradicts the previous observation that a large clause is assumed to have greater utility when LBD is fixed.
- (A3) Time has a positive effect on Utility for a low-LBD clause and a negative effect for a high-LBD clause. The results show that a high-LBD clause experiences a rapid drop in utility while the utility of a low-LBD clause even increases over time. Therefore, low-LBD clauses should be preserved in memory for a longer time than high-LBD clauses.
- (A4) LBD has a larger absolute ATE value than size, indicating a greater impact on clause utility. The result suggests the usage of LBD in a clause memory management heuristic against the size from a causal perspective.
- (A5) Appendix E [54] lists the ATE of all variables on clause utility over the normalized data. Among them, Propagation has the greatest effect. The result suggests the introduction of Propagation into the future design of clause memory management heuristics.
- (A6) Branching has a positive effect on Utility when changing the strategy from VSIDS to Maple. Hence, Maple leads to greater utility and serves as a recommended branching heuristic for high-clause-utility scenarios.
- (A7) Luby restart leads to a greater utility than Geometric and LBD-based, achieving the greatest utility. Hence, Luby is a recommended restart heuristic for high-clause-utility scenarios.

■ **Table 5** Query estimates and their interpretations.

Query	Answer
Q1 $\text{ATE}(\text{LBD}, \text{Utility}, 2, 1) = -0.2610 < 0$	Low-LBD clause has greater utility.
Q2 $\begin{cases} \text{ATE}(\text{Size}, \text{Utility}, 2, 1) = -0.0314 < 0 \\ \text{ACATE}(\text{Size}, \text{Utility}, \text{LBD}, 2, 1) = -0.0202 < 0 \end{cases}$	Small clause has greater utility, which also holds when LBD is fixed.
Q3 $\begin{cases} \text{CATE}(\text{Time}, \text{Utility}, \text{LBD} \leq 6, 10000, 0) = 0.3772 > 0 \\ \text{CATE}(\text{Time}, \text{Utility}, \text{LBD} > 6, 10000, 0) = -0.0884 < 0 \end{cases}$	High-LBD clause experiences a rapid drop in utility over time.
Q4 $ \text{ATE}(\text{Size}, \text{Utility}, 2, 1) = 2.6660 < \text{ATE}(\text{LBD}, \text{Utility}, 2, 1) = 3.3754$	LBD has a greater impact than size.
Q5 $\arg \max_{\text{Treatment} \neq \text{Utility}} \{ \text{ATE}(\text{Treatment}, \text{Utility}, 2, 1) \}$	Propagation has the greatest impact on utility.
Q6 $\text{ATE}(\text{Branching}, \text{Utility}, \text{Maple}, \text{VSIDS}) = 18.5745 > 0$	Maple leads to greater utility.
Q7 $\begin{cases} \text{ATE}(\text{Restart}, \text{Utility}, \text{Luby}, \text{Geometric}) = 6.7347 > 0 \\ \text{ATE}(\text{Restart}, \text{Utility}, \text{Luby}, \text{LBD-based}) = 17.7385 > 0 \\ \text{ATE}(\text{Restart}, \text{Utility}, \text{Geometric}, \text{LBD-based}) = 11.0037 > 0 \end{cases}$	Luby leads to the greatest utility.

5 Conclusion

This work serves as the first work to utilize causality to uncover the inner workings of modern SAT solvers. We proposed a framework to calculate the causal effects on clause utility from key factors such as LBD, size, activity, and the choice of heuristics. The causal results verified the “rules of thumb” in solver implementation, for example, that a low-LBD clause is assumed to have greater utility than a high-LBD clause. We also answered questions closely related to SAT solving, for instance, which factor, size or LBD, has a greater impact on clause utility? The causal framework provides a systematical way to quantitatively investigate the relationship among components of a SAT solver, which paves the way to further understanding how modern SAT solvers work.

This work opens several promising directions for future research. Our immediate focus will be enhancing our framework by introducing a new definition for clause utility. This revised definition will consider the frequency of a learned clause’s usage by a SAT solver rather than its usage in DRAT proof generation. Additionally, We aim to include a definition for clause utility that encompasses both satisfiable and unsatisfiable instances, expanding beyond our current analysis of unsatisfiable cases. Looking ahead, we intend to incorporate a broader range of features into our study and extend a further investigation into the obtained causal graph’s structure and causal estimates’ value to enrich our analysis.

References

- 1 Kartik Ahuja, Yixin Wang, Divyat Mahajan, and Yoshua Bengio. Interventional causal representation learning. In *Proc. of NeurIPS Workshop on Causality for Real-world Impact*, 2022.
- 2 Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization. *arXiv preprint*, 2019. [arXiv:1907.02893](https://arxiv.org/abs/1907.02893).
- 3 Gilles Audemard and Laurent Simon. Glucose: a solver that predicts learnt clauses quality. *SAT Competition*, 2009.
- 4 Gilles Audemard and Laurent Simon. Predicting learnt clauses quality in modern SAT solvers. In *Proc. of IJCAI*, 2009.
- 5 Gilles Audemard and Laurent Simon. On the glucose SAT solver. *International Journal on Artificial Intelligence Tools*, 2018.
- 6 Armin Biere. Lingeling, Plingeling and Treengeling entering the SAT competition 2013. *Proc. of SAT competition*, 2013.
- 7 Armin Biere, Alessandro Cimatti, Edmund Clarke, and Yunshan Zhu. Symbolic model checking without BDDs. In *Proc. of TACAS*, 1999.
- 8 Armin Biere, Katalin Fazekas, Mathias Fleury, and Maximillian Heisinger. CaDiCaL, Kissat, Paracooba, Plingeling and Treengeling entering the SAT Competition 2020. In *Proc. of SAT Competition – Solver and Benchmark Descriptions*, 2020.
- 9 Armin Biere and Andreas Fröhlich. Evaluating CDCL variable scoring schemes. In *Proc. of DAC*, 2015.
- 10 Johann Brehmer, Pim De Haan, Phillip Lippe, and Taco Cohen. Weakly supervised causal representation learning. In *Proc. of UAI Workshop on Causal Representation Learning*, 2022.
- 11 Edmund Clarke, Daniel Kroening, and Flavio Lerda. A tool for checking ANSI-C programs. In *Proc. of TACAS*, 2004.
- 12 Diego Colombo and Marloes H. Maathuis. Order-Independent Constraint-Based Causal Structure Learning. *Journal of Machine Learning Research*, 2014.
- 13 Stephen A Cook. The complexity of theorem-proving procedures. In *Proc. of STOC*, 1971.
- 14 Martin Davis, George Logemann, and Donald Loveland. A machine program for theorem-proving. *Communications of the ACM*, 1962.

- 15 Niklas Eén and Niklas Sörensson. An extensible SAT-solver. In *Proc. of SAT*, 2003.
- 16 Jan Elffers, Jesús Giráldez-Cru, Stephan Gocht, Jakob Nordström, and Laurent Simon. Seeking practical cdcl insights from theoretical sat benchmarks. In *IJCAI*, pages 1300–1308, 2018.
- 17 Laurent Simon Gilles Audemard. Glucose 1.0, 2009. sources/glucose/core/Solver.C:Lines 631–635. URL: https://www.labri.fr/perso/lsimon/downloads/software/glucose_1.0.zip.
- 18 Carla P Gomes, Bart Selman, Henry Kautz, et al. Boosting combinatorial search through randomization. *Proc. of AAAI/IAAI*, 1998.
- 19 Jinbo Huang et al. The Effect of Restarts on the Efficiency of Clause Learning. In *Proc. of IJCAI*, 2007.
- 20 Yimin Huang and Marco Valtorta. Pearl’s calculus of intervention is complete. In *Proc. of UAI*, 2006.
- 21 Matti Järvisalo, Daniel Le Berre, Olivier Roussel, and Laurent Simon. The international SAT solver competitions. *AI Magazine*, 2012.
- 22 Henry A Kautz, Bart Selman, et al. Planning as Satisfiability. In *Proc. of ECAI*, 1992.
- 23 Janne I Kokkala and Jakob Nordström. Using resolution proofs to analyse CDCL solvers. In *Proc. of CP*, 2020.
- 24 Jia Hui Liang, Vijay Ganesh, Pascal Poupart, and Krzysztof Czarnecki. Learning rate based branching heuristic for SAT solvers. In *Proc. of SAT*, 2016.
- 25 Jia Hui Liang, Chanseok Oh, Minu Mathew, Ciza Thomas, Chunxiao Li, and Vijay Ganesh. Machine learning-based restart policy for CDCL SAT solvers. In *Proc. of SAT*, 2018.
- 26 Jia Hui Liang, Pascal Poupart, Krzysztof Czarnecki, and Vijay Ganesh. An empirical study of branching heuristics through the lens of global learning rate. In *Proc. of SAT*, 2017.
- 27 Michael Luby, Alistair Sinclair, and David Zuckerman. Optimal speedup of Las Vegas algorithms. *Information Processing Letters*, 1993.
- 28 Mao Luo, Chu-Min Li, Fan Xiao, Felip Manyà, and Zhipeng Lü. An effective learnt clause minimization approach for CDCL SAT solvers. In *Proc. of IJCAI*, 2017.
- 29 Inês Lynce and Joao Marques-Silva. SAT in bioinformatics: Making the case with haplotype inference. In *Proc. of SAT*, 2006.
- 30 João P Marques-Silva and Kareem A Sakallah. GRASP: A search algorithm for propositional satisfiability. *IEEE Transactions on Computers*, 1999.
- 31 Matthew W Moskewicz, Conor F Madigan, Ying Zhao, Lintao Zhang, and Sharad Malik. Chaff: Engineering an efficient SAT solver. In *Proc. of DAC*, 2001.
- 32 Alexander Nadel and Vadim Ryvchin. Chronological backtracking. In *Proc. of SAT*, 2018.
- 33 Chanseok Oh. Between SAT and UNSAT: the fundamental difference in CDCL SAT. In *Proc. of SAT*, 2015.
- 34 Chanseok Oh. *Improving SAT solvers by exploiting empirical characteristics of CDCL*. PhD thesis, New York University, 2016.
- 35 Judea Pearl. Causal inference in statistics: An overview. *Statistics Surveys*, 2009.
- 36 Judea Pearl. *Causality*. Cambridge University Press, 2009.
- 37 Jonas Peters, Peter Bühlmann, and Nicolai Meinshausen. Causal inference by using invariant prediction: identification and confidence intervals. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, pages 947–1012, 2016.
- 38 Knot Pipatsrisawat and Adnan Darwiche. A lightweight component caching scheme for satisfiability solvers. In *Proc. of SAT*, 2007.
- 39 Bernhard Schölkopf, Francesco Locatello, Stefan Bauer, Nan Rosemary Ke, Nal Kalchbrenner, Anirudh Goyal, and Yoshua Bengio. Toward causal representation learning. *Proc. of IEEE*, 109(5):612–634, 2021.
- 40 Gideon Schwarz. Estimating the dimension of a model. *The Annals of Statistics*, 1978.
- 41 Marco Scutari. Learning Bayesian Networks with the bnlearn R Package. *Journal of Statistical Software*, 2010.
- 42 Marco Scutari and Radhakrishnan Nagarajan. Identifying significant edges in graphical models of molecular networks. *Artificial Intelligence in Medicine*, 2013.

- 43 Amit Sharma, Emre Kiciman, et al. DoWhy: A Python package for causal inference, 2019. URL: <https://github.com/microsoft/dowhy>.
- 44 Qiang Shen and Rónán Daly. Methods to accelerate the learning of bayesian network structures. In *Proc. of UK Workshop on Computational Intelligence*, 2007.
- 45 JP Marques Silva and Karem A Sakallah. Conflict analysis in search algorithms for satisfiability. In *Proc. of ICTAI*, 1996.
- 46 Laurent Simon. Post mortem analysis of sat solver proofs. In *POS@ SAT*, pages 26–40, 2014.
- 47 Mate Soos, Jo Devriendt, Stephan Gocht, Arijit Shaw, and Kuldeep S Meel. Cryptominisat with CCAnr at the SAT competition 2020. *SAT COMPETITION*, 2020.
- 48 Mate Soos, Raghav Kulkarni, and Kuldeep S Meel. CrystalBall: Gazing in the Black Box of SAT Solving. In *Proc. of SAT*, 2019.
- 49 Mate Soos, Karsten Nohl, and Claude Castelluccia. Extending SAT solvers to cryptographic problems. In *Proc. of SAT*, 2009.
- 50 Toby Walsh. Search in a small world. In *Proc. of IJCAI*, pages 1172–1177, 1999.
- 51 Ruoyu Wang, Mingyang Yi, Zhitang Chen, and Shengyu Zhu. Out-of-distribution generalization with causal invariant transformations. In *Proc. of CVPR*, pages 375–385, 2022.
- 52 Nathan Wetzler, Marijn JH Heule, and Warren A Hunt Jr. DRAT-trim: Efficient checking and trimming using expressive clausal proofs. In *Proc. of SAT*, 2014.
- 53 Lin Xu, Frank Hutter, Holger H Hoos, and Kevin Leyton-Brown. SATzilla: portfolio-based algorithm selection for SAT. *Journal of artificial intelligence research*, 2008.
- 54 Jiong Yang, Arijit Shaw, Teodora Baluta, Mate Soos, and Kuldeep S. Meel. Explaining sat solving using causal reasoning. *arXiv*, 2023. [arXiv:2306.06294](https://arxiv.org/abs/2306.06294).