

Automated Explanation Selection for Scientific Discovery

Ashlin Iser*

Karlsruhe Institute of Technology, Germany
ORCID (Ashlin Iser): <https://orcid.org/0000-0003-2904-232X>

Abstract. Automated reasoning is a key technology in the young but rapidly growing field of Explainable Artificial Intelligence (XAI). Explainability helps build trust in artificial intelligence systems beyond their mere predictive accuracy and robustness. In this position paper, we propose a cycle of scientific discovery that combines machine learning with automated reasoning for the generation and the selection of explanations. We present a taxonomy of explanation selection problems that draws on insights from sociology and cognitive science. These selection criteria subsume existing notions and extend them with new properties.

1 Introduction

The trustworthiness of artificial intelligence (AI) systems, including a sound rationale for their deployment, is paramount in safety-critical and high-risk applications such as critical infrastructure operations, healthcare diagnostics and medical devices, law enforcement and criminal justice, biometric identification and categorization of natural persons, or autonomous vehicles such as self-driving cars or drones [46]. There are also legal reasons to want intelligible AI, such as the “European Artificial Intelligence Act” [8]. Discrimination based on gender, race, or origin is a significant challenge in many areas of contemporary society. Discrimination can be exacerbated by the use of AI systems that are trained on data sets that reflect the same biases that exist in the wider world. This becomes a major concern when AI systems are used to make decisions that *affect the life of people*.

Despite their problem-solving capabilities, machine learning (ML) models often suffer from opacity, making their inner workings difficult or impossible for humans to understand. Explainability contributes to building trust in AI systems beyond their pure predictive accuracy and robustness [51]. Efficient automated reasoning enables exact approaches to the explainability of ML models, and is thus a key technology of the young but rapidly growing field of explainable artificial intelligence (XAI) [35]. Formal methods facilitate the early detection of flaws in ML models and the validation of ML models by checking for their coherence with prior knowledge and expectations. This facilitates the identification and prevention of flawed decisions and discriminatory practices through the use of AI systems.

Since the training of an ML model is an inductive process that exploits correlations in a set of observations, i.e., training data, the epistemological status of such learned abstractions is subject to the *problem of induction* [30]. This means that while inductive inferences, such as the rules manifested in ML models, are likely to be true based on the given data, they are not guaranteed to be true. In

contrast, deductive reasoning is characterized by certainty and precision: the conclusions are guaranteed to be true if the premise is true. This implies that if the premise is an accurate representation of an ML model, the conclusions are accurate if the model is accurate. Conversely, if the conclusions are erroneous, the model is erroneous. Therefore, any form of formal explanation extracted from an ML model can merely serve as a hypothesis and can potentially be falsified by new observations.

In this paper, we propose a machine-assisted cycle of scientific discovery which is illustrated in Figure 1. Our model combines machine learning with automated reasoning for the automated selection of explanations which can be used to formulate hypotheses that can be tested by experiments. The blue box represents the inductive phase of the cycle, in which a model is learned from existing training data. The red box represents the deductive phase, where the learned model is encoded into a formal language and automated reasoning is used to deduce explanations for the model. The purple box represents the process of explanation selection which researchers can compare with prior knowledge and even use to formulate hypotheses. These hypotheses can then be tested by further experiments, which in turn produce new data, thereby closing the cycle.

Modern automated reasoning systems are capable of producing certified deductive conclusions that can be checked by formally verified software. This capability enables the reliable automatic analysis of ML models, including verification of ML models and extraction of knowledge from ML models to assist humans in learning from AI systems. The present position paper will focus on this latter aspect.

We begin with a general discussion of the suitability of automated reasoning methods for generating explanations for machine learning models in Section 2. Then, in Section 3, we introduce a taxonomy of explanation selection problems that subsumes existing notions and extends them with new properties. In particular, we discuss desirable properties of explanations drawn from the social sciences and how these can be used to guide automated explanation selection. Finally, we discuss how the evaluation of explanations is simplified by the guarantees provided by automated reasoning methods in Section 4, and then conclude with a summary of our findings.

2 Automated Reasoning for Explainability

In recent years, considerable progress has been made in the field of automated reasoning, particularly with regard to algorithmic solutions to the propositional satisfiability problem (SAT) [9]. Automated reasoning and the SAT problem are closely related, with SAT solvers often being a central component in automated reasoning systems. Even though the complexity of SAT is challenging, the ubiquity of NP-hard problems has stimulated research to the extent that modern

* Email: ashlin.iser@kit.edu




Figure 1. Using explainable artificial intelligence to aid scientific discovery.

SAT solvers efficiently solve many SAT instances arising in practice. In addition, modern SAT solvers are suitable for use in safety-critical and high-risk applications because they produce proofs that are given as efficiently checkable sequences of rigorous mathematical reasoning steps [13].

The genericity of the specification of the SAT problem is a key factor for their wide range of applications, including the verification of hardware and software, cryptanalysis, planning, and many more. Advances in SAT solving often lead to advances in related reasoning problems and their optimization variants, such as Maximum Satisfiability (MaxSAT) or Pseudo-Boolean Satisfiability (PBS) and Optimization (PBO). Such advances are also closely related to problems higher up in the polynomial hierarchy, such as Quantified Boolean Formulas (QBF). SAT solvers are also used in the backend of reasoning systems that support richer formal languages, such as SAT Modulo Theory (SMT) solvers [3].

There are three main reasons why SAT-based methods are particularly well suited for this task. First, recent advances in SAT solving show that SAT-based methods have the potential to efficiently solve the instances of the NP-hard explanation selection problems addressed in this paper. Second, SAT-based methods are particularly reliable because they produce proofs that can be efficiently checked for correctness by verified software. Finally, these methods are generic, allowing flexible adaptation to problem variants by modifying the encoding of the problem instance rather than the algorithm.

Efficiency and Trustworthiness

Conflict-driven Clause Learning (CDCL) is the most popular SAT algorithm, because its implementations are extremely efficient in practice. Essential to its success are efficient data structures, implementations, and heuristics for its main pruning techniques, which are unit propagation and clause learning [37].

While solutions to the SAT problem by definition produce certificates of satisfiability where possible, state-of-the-art CDCL solvers also produce certificates of unsatisfiability that can be verified in time polynomial in the length of the proof. Assuming $NP \neq co-NP$, the length of a polynomial time checkable proof is necessarily worst-case exponential in the size of the input formula. But proof checking is feasible in practice, since the length of the proof correlates with the solving time, i.e., if we could generate it, we can check it [13]. Certificate checking is a distinct process from certificate creation and can be accomplished with formally verified software, thereby enhancing the dependability of the solver outcomes [49]. This is particularly crucial when it comes to verifying the safety properties of systems for high-risk applications.

The study of proof complexity provides the theoretical background for understanding some of the practical advances and limitations of logic solvers, in particular with respect to resolution-inspired clause

learning in modern SAT solvers and the emergence of new heuristics for exploiting more powerful proof systems in SAT solving. CDCL solver runtimes are subject to the same domain-specific exponential lower bounds as proof-lengths in the general resolution proof-system they simulate [5]. The study of more powerful proof systems, such as extended resolution, has a long tradition, and recent implementations of new practical heuristics manage to exploit them. This can be achieved by the addition of propagation redundant clauses [44], or symmetry-breaking clauses [7], or variable definitions [11] that help break lower bounds on the length of unsatisfiability proofs is a way to take advantage of more powerful proof systems.

While the SAT solver with the highest average performance may not be optimal in all scenarios, the use of a portfolio of SAT solvers can significantly improve the expected overall performance, indicating that the selection of a solver should be based on the domain to be solved [2]. Recent advances have also built on the capabilities of parallel SAT solvers, with the most efficient being clause-sharing portfolios of CDCL solvers. The massively parallel SAT solver Mallob builds on a portfolio of some of the most successful sequential SAT solvers and employs a unique scalable clause-sharing strategy [47].

The problem of enumerating all models of a satisfiable formula (AllSAT) has been addressed by a variety of approaches that can be broadly classified into two categories: blocking approaches, which are based on incremental SAT solving using minimized blocking clauses [54], and non-blocking approaches, which ensure exclusive enumeration through chronological backtracking [10]. A recent approach to AllSAT outperforms existing approaches by avoiding no-good learning through the combination of chronological backtracking and implicant shrinking [48].

The propositional optimization problem (MaxSAT), is defined as the problem of finding an assignment that maximizes the number of satisfied clauses. This problem has two notable variants: a partial variant, in which a subset of its clauses must be satisfied, and a weighted variant, in which the objective is to maximize a sum of clause weights. The MaxSAT algorithms that remain relevant today include model-improving approaches, also known as Linear SAT UNSAT (LSU), core-guided approaches (CG), and implicit hitting set approaches (IHS). These algorithms essentially solve a sequence of SAT instances with an incremental CDCL SAT solver [1]. LSU algorithms solve a sequence of satisfiable instances by adding bound-tightening constraints until the instance is UNSAT. Conversely, CG algorithms solve a sequence of unsatisfiable instances by relaxing the instance in each iteration until the instance is SAT. Similar to core-guided approaches, IHS algorithms solve a sequence of unsatisfiable instances, but tighten the lower bound by computing a minimum-cost hitting set of the accumulated unsatisfiable cores. Recent advances in MaxSAT solvers have been made in multiobjective optimization [28] and in incremental MaxSAT solving [41].

Formal Explanation Approaches

In naturally interpretable models such as decision trees, evaluation traces can serve as direct explanations of prediction results [46]. In practice, such models often have limited predictive power because they can only generate step-like decision boundaries, are mostly restricted to categorical data, and are prone to overfitting. Moreover, according to recent research in formal explainability, explanations by interpretable models are generally not concise and exhibit redundancy [27].

Heuristic explanation approaches sample the input space around the instance to be explained and observe the effect on the output space. As a result, these approaches can be applied to any class of prediction models. Due to their probabilistic nature, such model-agnostic approaches do not provide strict guarantees of correctness or minimality of explanations. A prominent example is LIME, a sampling-based approach that locally approximates a given model by an intrinsically interpretable model, such as a decision tree [45]. However, heuristic explainers are not guaranteed to be correct, and their explanations can be misleading. A formal analysis of the quality of the well-known Shapley values has shown that they can make misleading predictions about the importance of features [19]. Anchors is an approach that computes sufficient conditions along with an estimate of their accuracy for a given prediction. Also for Anchors, a formal analysis has shown that it gives unrealistically high accuracy estimates of explanations [39].

In contrast, explanations by formal methods are derived from an exact formal language encoding of the learned model and a formalization of the desired properties of the sought explanations. The guarantees of formal approaches are mathematically exact, as they are based on sound deduction systems [21]. In recent years, the emerging formal XAI community has focused on the notions of abductive explanation (AXp) and contrastive explanation (CXp). Here, an AXp, also known as prime implicant explanation (PI-explanation) or minimally sufficient reason, denotes a minimal set of features sufficient to guarantee a given prediction. In contrast, a CXp, also known as minimal required change, denotes a minimal set of features sufficient to change a prediction [23]. There are weak and strong variants in the literature to distinguish *subset minimality* from *cardinality minimality* for both AXps and CXps. The computation of both AXps and CXps for tree ensembles is an NP-hard problem [26].

Nevertheless, efficient formal explainers for tree ensembles have been developed, one based on recent SMT and one based on recent MaxSAT approaches. In terms of runtime, both exact approaches outperform even the heuristic explainer Anchor [25]. An important property of AXps and CXps is their duality, i.e., contrastive explanations are minimal hitting sets of abductive explanations and vice versa [24]. Thus, AXp/CXp enumeration is closely related to the well-studied enumeration problems of Minimum Unsatisfiable Subsets (MUS) and Minimum Correction Sets (MCS) in formal logic which is Σ_2^P -hard [36]. The state-of-the-art formal explanation enumeration algorithm for decision lists is one that exploits this duality, analogous to the generic MARCO algorithm for MUS/MCS enumeration [22, 32].

Ensemble learning is a popular approach to improve the accuracy of ML models by combining multiple weak models into a single strong model [12]. Tree ensembles, such as random forests and gradient boosting machines, are well-known examples of ensemble learning methods in which the weak models are decision trees. Tree ensembles achieve high model accuracy with a minimal number of hyperparameters, and their generalization performance is robust even

in high-dimensional feature spaces due to randomization in the training process and the combination of many decision trees with limited depth. Although tree ensembles are formed from a large number of intrinsically interpretable decision trees, they themselves are *not* intrinsically interpretable.

Studying the formal explainability of tree ensembles is a relevant direction of research because tree ensembles are a common class of ML models used in many practical applications, such as diagnosis [53], or marketing [43]. Moreover, they can be used to explain the predictions of any black-box or otherwise hard-to-codify prediction model, such as a large neural network, as is done in Asteryx, which combines sampling-based explanation with formal explanation. Asteryx uses sampling to locally approximate a given black-box model by a random forest trained on the samples drawn. This random forest then serves as a surrogate model for the subsequent computation of exact formal explanations [4].

3 Desirable Properties of Explanations

We find that the problem of explanation selection has been understudied, resulting in a literature that contains redundant and overlapping notions, some of which even have misleading names. To alleviate these problems, we have reviewed the literature on relevant notions of explanation in the social sciences and have identified a set of desirable properties that can be used to guide the selection of explanations.

There is an extensive body of research from philosophy, psychology, sociology and cognitive science on the nature of explanations and what constitutes a good, useful or acceptable explanation [38]. On the nature of explanations, we first note that an explanation is an answer to a “why?” question. While the concrete question “Why P?” can be extended to “Why P instead of not P?”, in informal communication a contrast case Q is given by the context, so that implicitly it is almost always a question of the form “Why P instead of Q?” [42]. Contrast cases, also known as foils, are key to providing the context for an explanation, and different contrast cases lead to different explanations [6]. An effect can have multiple necessary and sufficient causes but full causal chains are often too large to comprehend [15]. The contrast between *the fact and the foil* is the primary criterion for selecting an explanation which highlights the differences between the contrasted cases [33, 14].

Moreover, *necessity and sufficiency* are strong criteria for selecting explanations [33]. *Minimality and generality* are other desirable properties of explanations, as people tend to prefer shorter explanations and those that apply to many situations rather than just one particular situation [50, 52]. Explanation selection can also be based on *anomaly*, i.e., people tend to select unusual events as explanations [16]. In contrast, the most probable explanations are often assumed to be known [14]. Explanations based on statistical generalization require an additional explanation of the *causes for the statistical generalization* [29].

Selection of useful explanations effectively reduces the exponential number of explanations to a digestible set. This involves a systematic formalization of desirable properties, thereby extending and subsuming existing notions of explanation selection. The diverse (and sometimes multiple) objectives can be then formulated as (multi-objective) MaxSAT problems. This also allows us to formulate objective thresholds that can be used to reduce the number of explanations for enumeration approaches.

Necessity, sufficiency, and minimality have been studied in the context of formal explanation selection for machine learning mod-

els [23]. For minimality, we can formulate both weak and strong versions of some of those desirable properties to disambiguate subset-minimality and cardinality-minimality. In addition, contrast cases can be specified in each of the resulting explanation selection problems. Some desirable properties require reasoning about an underlying data set, which poses additional challenges for formal encoding of models. This particularly affects *generality*, i.e., conditions explaining the largest number of data points, and *anomaly*, i.e., conditions explaining the smallest number of data points.

Relationship with Theories of Causality

Even though explanation and causality denote two different things, explanation can be associated with the dependence theories of causality, which include Hume’s theory of regularity [20], later refined by Lewis, who defined causality as a matter of counterfactuals [31]. This concept is closely related to explanation selection, but in this context we use counterfactuals to isolate desired explanations from the set of possible explanations.

Relationship with Theories of Abduction

Explanation selection is also closely related to the notion of abduction, which was introduced by Peirce as a third mode of reasoning in addition to deduction and induction. Deduction is the process of deriving a result from a known fact and a rule. Induction, on the other hand, is the process of generalizing a rule from a set of observed facts and results. Abduction, on the other hand, is the process of proposing a fact to explain a known result using a known rule [34].

There is a large body of work on abductive reasoning in the context of the philosophy of logic and scientific discovery [34]. Peirce developed a theory of science in which abduction is the step in which a hypothesis is generated, i.e., a fact that explains an observation is suggested by scientific intuition, then deduction is used to evaluate the plausibility of the hypothesis, and then induction is used to provide further evidence for the hypothesis through experimentation.

In AI research, interpretations that capture abduction as a form of “Inference to the Best Explanation” (IBE) are common [29]. In contrast, Hintikka’s notion of abduction is embedded in an interrogative model of inquiry [18] and understood as the introduction of new explanatory hypotheses, as opposed to IBE’s reasoning with a fixed hypothesis space [17]. This view is closer to Peirce’s notion of abduction, emphasizing hypothesis formation over selection among given alternatives; whereas IBE reduces to deduction over a fixed set of hypotheses with additional preference criteria.

4 Evaluation of Explanations

The authors of a recent comprehensive survey of evaluation methods in Explainable Artificial Intelligence identified a total of *twelve evaluation criteria* for the explanations generated by XAI tools [40]. In the following, we argue that many of these criteria can be fully satisfied a priori by formal reasoning methods.

The criteria *correctness*, *completeness*, and *consistency*, are fully satisfied by SAT-based formal reasoning methods, since they produce certified results that can be validated with verified software. Exact formal approaches can also explicitly maximize the *continuity* and *compactness* of the explanations by incorporating them as objective functions in the encoding of the explanation selection problem, while at the same time preserving the correctness and completeness of the encoding.

However, *coherence* with background knowledge, beliefs, etc. is an inappropriate criterion for evaluating explanations in the context of formal reasoning. This is because, due to the exactness of the formal method, any incoherence of explanations necessarily follows from an incoherence of the model. In the context of heuristic explainers, which are not guaranteed to be exact, it may make sense to check whether the resulting explanations are consistent with what is expected from the background knowledge. In the formal setting, however, if the exact explanation is not consistent with the background knowledge, then the model is not consistent with the background knowledge. Thus, coherence is not a criterion for evaluating the explanation, but rather a criterion for evaluating the model. Formal verification of ML models is a natural complementary application to explanation selection. Verification of ML models can be achieved by encoding expectations as additional constraints in the model encoding to allow proof by contradiction.

Since SAT-based formal reasoning produces exact and verified results, the *confidence* in the generated explanations is generally high. However, the whole process still depends on the correctness and completeness of the SAT encoding of the model, as illustrated by the edge connecting the red and blue boxes in Figure 1. This involves parsing the data structures of the ML model M and generating a formal language representation F that preserves the semantics of M . This means that for any input I and output O , the formula F has a model if and only if $M(I) = O$, provided that the input and output vectors I and O are represented by variables of F . This encoding phase must be verified by additional means, such as formal proofs of the correctness and completeness of the encoding and, for testing practical implementations, automated consistency tests.

The *contrastivity* criterion evaluates the possibility of specifying contrast cases. Formal encodings of ML models should therefore explicitly allow the specification of contrast cases in the encoding of the explanation selection problem.

Other evaluation criteria such as *controllability*, *covariate complexity*, *composition*, and *user satisfaction* are not directly addressed by the formal reasoning methods themselves, but can be achieved through the design of an appropriate user interface for specification of user constraints, and for presentation and visualization of explanations and explanation statistics.

5 Conclusion

We proposed a cycle of scientific discovery that combines machine learning with automated reasoning for generating and selecting explanations. Our findings indicate that automated reasoning methods are well-suited for generating explanations for machine learning models and can guide their selection. We also introduced a taxonomy of explanation selection problems, drawing on insights from sociology and cognitive science, which subsumes existing notions and extends them with new properties. We showed that evaluation is greatly simplified through the guarantees provided by automated reasoning methods. Overall, applying automated reasoning to explanation selection for machine learning models represents a promising direction for future research.

Acknowledgements

I would like to express my gratitude to Prof. Arnaud Durand (Univ. Paris Cité) and Prof. Christine Roussat (Univ. Clermont Auvergne) for their invaluable feedback, which they provided through discussions or comments on earlier drafts of this paper.

References

- [1] F. Bacchus, M. Järvisalo, and R. Martins. Maximum satisfiability. In *Handb. of Satisf. (2nd Ed.)*, pages 929–991. 2021. doi: 10.3233/FAIA201008.
- [2] J. Bach, M. Iser, and K. Böhm. A Comprehensive Study of k-Portfolios of Recent SAT Solvers. In *Conf. on Theo. and Appl. of Satisf. Test., SAT*, 2022. doi: 10.4230/LIPICS.SAT.2022.2.
- [3] A. Biere, M. Heule, H. van Maaren, and T. Walsh, editors. *Handb. of Satisf. (Second Edition)*, Frontiers in AI and Applications, 2021.
- [4] R. Boumazouza, F. Cheikh-Alili, B. Mazure, and K. Tabia. Asteryx: A model-agnostic SAT-based approach for symbolic and score-based expl. In *Int. Conf. on Inf. & Knowl. Man., CIKM*, pages 120–129, 2021. doi: 10.1145/3459637.3482321.
- [5] S. Buss and J. Nordström. Proof complexity and SAT solving. In *Handb. of Satisf. (2nd Ed.)*, pages 233–350. 2021. doi: 10.3233/FAIA200990.
- [6] S. Chin-Parker and J. Cantelon. Contrastive constraints guide explanation-based category learning. *Cogn. Sci.*, 2017. doi: 10.1111/cogs.12405.
- [7] J. Devriendt, B. Bogaerts, M. Bruynooghe, and M. Denecker. Improved static symmetry breaking for SAT. In *Conf. on Theo. and Appl. of Satisf. Test., SAT*, pages 104–122, 2016. doi: 10.1007/978-3-319-40970-2_8.
- [8] EU. AI Act, 2021. URL <https://artificialintelligenceact.eu/the-act/>.
- [9] N. Froleys, M. Heule, M. Iser, M. Järvisalo, and M. Suda. SAT competition 2020. *AI J.*, 2021. doi: 10.1016/J.ARTINT.2021.103572.
- [10] O. Grumberg, A. Schuster, and A. Yadgar. Memory efficient all-solutions SAT solver and its application for reachability analysis. In *Conf. on Form. Meth. in Comp.-Aided Des., FMCAD*, pages 275–289, 2004. doi: 10.1007/978-3-540-30494-4_20.
- [11] A. Haberlandt, H. Green, and M. J. H. Heule. Effective auxiliary variables via structured reencoding. In *Conf. on Theo. and Appl. of Satisf. Test., SAT*, pages 1–19, 2023. doi: 10.4230/LIPICS.SAT.2023.11.
- [12] T. Hastie, R. Tibshirani, and J. H. Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction, 2nd Edition*. 2009. doi: 10.1007/978-0-387-84858-7.
- [13] M. Heule. Proofs of unsatisfiability. In *Handb. of Satisf. (2nd Ed.)*, pages 635–668. 2021. doi: 10.3233/FAIA200998.
- [14] D. Hilton. Mental models and causal explanation: Judgements of probable cause and explanatory relevance. *Thinking and Reasoning*, pages 273–308, 1996. doi: 10.1080/135467896394447.
- [15] D. Hilton. Social attribution and explanation. In *The Oxford Handb. of Causal Reas.*, pages 645–674. 2016. doi: 10.1093/oxfordhb/9780199399550.001.0001.
- [16] D. Hilton and B. Slugoski. Knowledge-based causal attribution: The abnormal conditions focus model. *Psychol. Rev.*, pages 75–88, 1986. doi: 10.1037/0033-295x.93.1.75.
- [17] J. Hintikka. What is abduction? the fundamental problem of contemporary epistemology. *Transactions of the Charles S. Peirce Society*, 1998. URL <http://www.jstor.org/stable/40320712>.
- [18] J. Hintikka. *Inquiry as Inquiry: A Logic of Scientific Discovery*. Springer, 1999. URL <https://doi.org/10.1007/978-94-015-9313-7>.
- [19] X. Huang and J. Marques-Silva. On the failings of shapley values for explainability. *Int. J. of Approx. Reas.*, page 109112, 2024. doi: <https://doi.org/10.1016/j.ijar.2023.109112>.
- [20] D. Hume. *Of the Idea of Necessary Connexion*, chapter 7. 2000. URL https://fitelson.org/confirmation/hume_enquiry.pdf.
- [21] A. Ignatiev. Towards trustworthy explainable AI. In *Int. Joint Conf. on AI, IJCAI*, pages 5154–5158, 2020. doi: 10.24963/ijcai.2020/726.
- [22] A. Ignatiev and J. Marques-Silva. Sat-based rigorous explanations for decision lists. In *Conf. on Theo. and Appl. of Satisf. Test., SAT*, pages 251–269, 2021. doi: 10.1007/978-3-030-80223-3_18.
- [23] A. Ignatiev, N. Narodytska, and J. Marques-Silva. Abduction-based explanations for machine learning models. In *AAAI Conf. on AI*, pages 1511–1519, 2019. doi: 10.1609/AAAI.V33I01.33011511.
- [24] A. Ignatiev, N. Narodytska, N. Asher, and J. Marques-Silva. From contrastive to abductive explanations and back again. In *Adv. in AI: Int. Conf. of the It. Assoc. for AI*, pages 335–355, 2020. doi: 10.1007/978-3-030-77091-4_21.
- [25] A. Ignatiev, Y. Izza, P. Stuckey, and J. Marques-Silva. Using maxsat for efficient explanations of tree ensembles. In *AAAI Conf. on AI*, pages 3776–3785, 2022. doi: 10.1609/AAAI.V36I4.20292.
- [26] Y. Izza and J. Marques-Silva. On explaining random forests with SAT. In *Int. Joint Conf. on AI, IJCAI*, pages 2584–2591, 2021. doi: 10.24963/IJCAI.2021/356.
- [27] Y. Izza, A. Ignatiev, and J. Marques-Silva. On tackling explanation redundancy in decision trees. *J. of AI Res.*, pages 261–321, 2022. doi: 10.1613/JAIR.1.13575.
- [28] C. Jabs, J. Berg, H. Ihalainen, and M. Järvisalo. Preprocessing in sat-based multi-objective combinatorial optimization. In *Conf. on Theo. and Appl. of Satisf. Test., SAT*, pages 18:1–18:20, 2023. doi: 10.4230/LIPICS.CP.2023.18.
- [29] J. R. Josephson and S. G. Josephson, editors. *Abductive Inference: Computation, Philosophy, Technology*. 1994. URL <https://books.google.de/books?id=uu6zXrogwWAC&lpg=PA1>.
- [30] M. Lange. Hume and the problem of induction. In *Inductive Logic, Handb. of the Hist. of Logic*, pages 43–91. 2011. doi: 10.1016/B978-0-444-52936-7.50002-1.
- [31] D. Lewis. Causation. *J. of Philos.*, pages 556–567, 1973. doi: 10.2307/2025310.
- [32] M. Liffiton, A. Previti, A. Malik, and J. Marques-Silva. Fast, flexible MUS enumeration. *Constraints*, pages 223–250, 2016. doi: 10.1007/S10601-015-9183-0.
- [33] P. Lipton. Contrastive explanation. *Royal Institute of Philosophy Supplement*, pages 247–266, 1990. doi: 10.1017/s1358246100005130.
- [34] L. Magnani, editor. *Handbook of Abductive Cognition*. Springer, 2023. URL <https://doi.org/10.1007/978-3-031-10135-9>.
- [35] J. Marques-Silva. *Logic-Based Explainability in Machine Learning*, pages 24–104. 2023. doi: 10.1007/978-3-031-31414-8_2.
- [36] J. Marques-Silva and C. Mencía. Reasoning about inconsistent formulas. In *Int. Joint Conf. on AI, IJCAI*, pages 4899–4906, 2020. doi: 10.24963/IJCAI.2020/682.
- [37] J. Marques-Silva, I. Lynce, and S. Malik. Conflict-driven clause learning SAT solvers. In *Handb. of Satisf. (2nd Ed.)*, pages 133–182. 2021. doi: 10.3233/FAIA200987.
- [38] T. Miller. Explanation in AI: Insights from the social sciences. *AI J.*, pages 1–38, 2019. doi: 10.1016/J.ARTINT.2018.07.007.
- [39] N. Narodytska, A. Shrotri, K. Meel, A. Ignatiev, and J. Marques-Silva. Assessing heuristic machine learning expl. with model counting. In *Conf. on Theo. and Appl. of Satisf. Test., SAT*, pages 267–278, 2019. doi: 10.1007/978-3-030-24258-9_19.
- [40] M. Nauta, J. Trienes, S. Pathak, E. Nguyen, M. Peters, Y. Schmitt, J. Schlotterer, M. van Keulen, and C. Seifert. From anecdotal evidence to quantitative evaluation methods: A systematic review on evaluating explainable AI. *ACM Comput. Surv.*, pages 295:1–295:42, 2023. doi: 10.1145/3583558.
- [41] A. Niskanen, J. Berg, and M. Järvisalo. Enabling incrementality in the implicit hitting set approach to maxsat under changing weights. In *Conf. on Princ. and Prac. of Constr. Prog., CP*, volume 210, pages 44:1–44:19, 2021. doi: 10.4230/LIPICS.CP.2021.44.
- [42] J. Pearl and D. Mackenzie. *The Book of Why: The New Science of Cause and Effect*. 2018. URL <https://dl.acm.org/doi/10.5555/3238230>.
- [43] B. Prabadevi, R. Shalini, and B. Kavitha. Customer churning analysis using machine learning algorithms. *International Journal of Intelligent Networks*, pages 145–154, 2023. doi: 10.1016/j.ijin.2023.05.005.
- [44] J. E. Reeves, M. J. H. Heule, and R. E. Bryant. Preprocessing of propagation redundant clauses. *J. Autom. Reason.*, 2023. doi: 10.1007/S10817-023-09681-3.
- [45] M. T. Ribeiro, S. Singh, and C. Guestrin. "why should i trust you?": Explaining the predictions of any classifier. In *Conf. on Knowl. Disc. and Data Min., KDD*, pages 1135–1144, 2016. doi: 10.1145/2939672.2939778.
- [46] C. Rudin. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, pages 206–215, 2019. doi: 10.1038/S42256-019-0048-X.
- [47] D. Schreiber. *Scalable SAT Solving and its Application*. PhD thesis, Karlsruhe Institut für Technologie (KIT), 2023.
- [48] G. Spallitta, R. Sebastiani, and A. Biere. Enumerating disjoint partial models without blocking clauses. In *AAAI Conf. on AI*, pages 1–10, 2023. doi: 10.48550/ARXIV.2306.00461.
- [49] Y. K. Tan, M. Heule, and M. O. Myreen. Verified propagation redundancy and compositional UNSAT checking in cakeml. *Softw. Tools Technol. Transf.*, pages 167–184, 2023. doi: 10.1007/S10009-022-00690-Y.
- [50] P. Thagard. Explanatory coherence. *Behavioral and Brain Sciences*, pages 435–467, 1989. doi: 10.1017/S0140525X00057046.
- [51] D. Weld and G. Bansal. The challenge of crafting intelligible intelligence. pages 70–79, 2019. doi: 10.1145/3282486.
- [52] J. Woodward. Sensitive and insensitive causation. *The Philos. Rev.*, pages 1–50, 2006. URL <http://www.jstor.org/stable/20446880>.
- [53] Z. Yang, X. Cui, and Z. Song. Predicting sepsis onset in icu using machine learning models: a systematic review and meta-analysis. *BMC Infectious Diseases*, 2023. doi: 10.1186/s12879-023-08614-0.
- [54] Y. Yu, P. Subramanyan, N. Tsiskaridze, and S. Malik. All-sat using minimal blocking clauses. In *Conf. on VLSI Des., VLSID*, pages 86–91, 2014. doi: 10.1109/VLSID.2014.22.