



A Study of the Learnability of Relational Properties

Model Counting Meets Machine Learning (MCML)

Muhammad Usman

University of Texas at Austin, USA
muhammadusman@utexas.edu

Kaiyuan Wang*

Google Inc., USA
kaiyuanw@google.com

Wenxi Wang

University of Texas at Austin, USA
wenxiw@utexas.edu

Marko Vasic

University of Texas at Austin, USA
vasic@utexas.edu

Haris Vikalo

University of Texas at Austin, USA
hvikalo@ece.utexas.edu

Sarfraz Khurshid

University of Texas at Austin, USA
khurshid@ece.utexas.edu

Abstract

This paper introduces the *MCML* approach for empirically studying the *learnability* of *relational* properties that can be expressed in the well-known software design language *Alloy*. A key novelty of MCML is quantification of the performance of and semantic differences among trained machine learning (ML) models, specifically decision trees, with respect to *entire* (bounded) input spaces, and not just for given training and test datasets (as is the common practice). MCML reduces the quantification problems to the classic complexity theory problem of *model counting*, and employs state-of-the-art model counters. The results show that relatively simple ML models can achieve surprisingly high performance (accuracy and F1-score) when evaluated in the common setting of using training and test datasets – even when the training dataset is much smaller than the test dataset – indicating the seeming simplicity of learning relational properties. However, MCML metrics based on model counting show that the performance can degrade substantially when tested against the entire (bounded) input space, indicating the high complexity of precisely learning these properties, and the usefulness of model counting in quantifying the true performance.

CCS Concepts: • Mathematics of computing → *Spectra of graphs*; • Computing methodologies → *Cross-validation*.

Keywords: Relational properties, machine learning, model counting, SAT solving, Alloy, ApproxMC, ProjMC

*Research performed while at the University of Texas at Austin.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org. *PLDI '20, June 15–20, 2020, London, UK*

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-7613-6/20/06...\$15.00

<https://doi.org/10.1145/3385412.3386015>

ACM Reference Format:

Muhammad Usman, Wenxi Wang, Marko Vasic, Kaiyuan Wang, Haris Vikalo, and Sarfraz Khurshid. 2020. A Study of the Learnability of Relational Properties: Model Counting Meets Machine Learning (MCML). In *Proceedings of the 41st ACM SIGPLAN International Conference on Programming Language Design and Implementation (PLDI '20), June 15–20, 2020, London, UK*. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3385412.3386015>

1 Introduction

Relational properties which relate abstract entities that can be viewed as vertices in a graph via edges that define the relations, e.g., the connectivity of a social network, or of an object graph on the heap, offer various benefits in the development of software systems [30, 48, 58]. For example, they enable automated analyses of software requirements, designs, specifications, and implementations [37]. However, in most cases, to benefit from the analyses the properties must be written manually. For complex systems [51, 72], writing them correctly is challenging, and faults in the properties’ statements can lead to erroneous confidence in correctness.

Our motivation is to leverage advances in machine learning (ML) to enhance software analyses, thereby improving software quality and reducing the cost of software failures. Our focus is on training binary classifiers that characterize relational properties. Such classifiers, once trained to have high accuracy, can offer much value in automated software analysis. For example, an executable classifier can serve as a *run-time check*, e.g., in an assertion, to validate that the program states at that point conform to the property represented by the classifier. Moreover, executable checks enable automated test input generation [5, 38], static analysis [19, 55], error recovery [14, 35], and automated theorem proving [13].

ML models can also be utilized in tandem with program synthesis techniques such as those based on *sketching* [56], e.g., a decision tree can provide the basis for a sketch that is completed by synthesis, or *holes* in a sketch can be filled in by the decision tree logic [22]. Moreover, learnt properties, even when somewhat imprecise, can be useful for the developer, e.g., a decision tree that approximates the properties of program states at a specific control point can provide insight into likely program behaviors and help seed other analyses.

In this paper, we introduce the *MCML* approach for empirically studying the *learnability* of a key class of relational properties that can be written in the well-known software design language *Alloy* [30]. Our aim is not to formalize learnability [4] of relational properties; rather, we aim to perform controlled experiments and rigorously study various well-known properties over small relations and graphs in order to gain insights into the potential role of ML methods in this important domain. Specifically, we consider training binary classifiers with respect to relational properties such that the trained classifiers accurately represent the properties, e.g., training a decision tree classifier to accurately classify each input as a directed-acyclic graph (DAG) or not a DAG.

A key novelty of MCML is that it allows quantifying the performance of trained decision tree models with respect to the *entire* input space (for a bounded universe of discourse) by utilizing given ground truth formulas, thereby enabling an evaluation of learnability that is not feasible using the commonly used ML approaches based on training and test datasets alone. Likewise, MCML allows quantifying differences among trained decision tree models for the entire (bounded) input space.

To quantify model performance and differences – irrespective of the datasets – MCML reduces the quantification problems to the classic complexity theory problem of *model counting*, which is to compute the number of solutions to a logical formula [21]. The formulas in our case represent semantic differences between the ground truth and the trained model, or between the two models. Given ground truth ϕ and decision tree d , the *false negative* count for d is the model count of “ $\phi \wedge \psi$ ” where ψ is d ’s logic that leads to the output “0” because any solution to “ $\phi \wedge \psi$ ” conforms to the ground truth but leads the decision tree to output “0”; likewise, the *false positive* count for d is the model count of “ $\neg\phi \wedge \tau$ ” where τ is d ’s logic that leads to the output “1”. The *true positive* and *true negative* counts are defined similarly. Using these counts accuracy, precision, recall and F1-score can be derived. Furthermore, given two decision trees d_1 and d_2 , their *semantic difference*, i.e., the number of inputs for which the tree outputs differ, is the sum of the model count of “ $\psi_1 \wedge \tau_2$ ” and the model count of “ $\tau_1 \wedge \psi_2$ ”, where ψ_i is d_i ’s logic that outputs “0”, and τ_i is d_i ’s logic that outputs “1” ($i \in \{1, 2\}$).

Model counting generalizes the propositional satisfiability checking (SAT) problem, and is #P-complete [49]. Despite the theoretical complexity, recent technological and algorithmic advances have led to practical tools that can handle very large state spaces [15, 26]. To embody MCML and quantify the performance and differences by means of model counting we employ two state-of-the-art tools: ApproxMC, which uses approximation techniques to estimate the number of solutions with high precision and provides formal guarantees [57]; and ProjMC [39], which uses an effective disjunctive decomposition scheme to compute the exact number of solutions. As is common with many tools for propositional

logic, ApproxMC and ProjMC take as input propositional formulas in conjunctive normal form (CNF) – the standard input format for SAT solvers. To create model counting problems in CNF, we define a translation from decision trees to CNF in the spirit of previous work [25]. The translation creates succinct CNF formulas that are linear in the size of the input trees.

To train ML models, we use the standard practice of employing training datasets. To evaluate the trained models, we use both the standard practice of employing test datasets and the MCML metrics. Our study has three key distinguishing characteristics: 1) for each property, we use *bounded exhaustive* sets of positive samples which contain *every* positive sample (up to partial symmetry breaking) within a bounded universe of discourse; 2) we leverage ground truth formulas to quantify the performance of trained decision trees with respect to the *entire* input space, not only the given datasets; and 3) we quantify semantic differences among different trained decision trees with respect to the entire input space without the use of ground truth formulas. Moreover, we evaluate different strategies for splitting the datasets into training sets and test sets, including ratios where the amount of training data is much smaller than the amount of test data.

To create the datasets for learning, we rely on logical formulas that describe the relational properties in Alloy – a first-order language with transitive closure [30]. For each property, we use Alloy’s SAT-based back-end to enumerate all the solutions, i.e., valuations that exhibit the property, up to Alloy’s default *symmetry breaking* which heuristically removes many but not all isomorphic solutions. The solutions created by the Alloy analyzer serve as the samples for training and test/evaluation. The solution spaces are very large – even with small bounds on the number of entities in the relations. For non-trivial properties, the number of positive samples is far smaller than the number of negative samples, and exhaustive enumeration of all negative samples is intractable. To avoid erroneously biasing the ML models to simply predict false if the datasets overwhelmingly consist of negative samples, we create *balanced* sets that contain the same number of positive and negative samples [18].

As subjects, we use six ML models, including decision trees, SVMs, and multi-layer perceptrons, and train them using datasets from 16 relational properties over small relations and graphs. We use the adjacency matrix representation for each data item in the training and test datasets; for example, for a relation over 7 entities, i.e., a graph with 7 vertices, we use 49 boolean inputs for the binary classifier which outputs true or false as the predicted value, and the space of all possible inputs has 2^{49} elements.

The results show that relatively simple ML models can achieve surprisingly high performance (accuracy and F1-score) at learning relational properties when evaluated in the common setting of using training and test datasets – even when the training dataset is substantially smaller than the

test dataset – indicating the seeming simplicity of learning these properties. However, the use of MCML metrics based on model counting shows that the performance can degrade substantially when tested against the entire (bounded) input space, indicating the high complexity of precisely learning these properties, and the usefulness of model counting in quantifying the true performance.

The contributions of this paper are as follows. 1) *Learning relational properties*. We present a systematic study of learning 16 relational properties using 6 off-the-shelf machine learning models; 2) *Model counting to quantify performance*. We reduce the problem of evaluating the performance of trained decision tree models over the entire input space with respect to ground truth formulas to the problem of model counting, and employ cutting edge approximate and exact counters to embody the reduction; and 3) *Model counting to quantify semantic differences*. We also introduce the use of model counting to quantify semantic differences between trained decision trees over the entire input space – without the need for ground truth formulas or evaluation datasets.

We believe the use of model counting in learning is a promising research area that can help gain deeper insights into the trained models, which can inform practical decisions on how to best utilize the models. For example, if a trained model in a deployed system is to be upgraded to a more sophisticated model, model counting could be a metric that in part informs the decision to upgrade.

2 Related Work

To our knowledge, MCML is the first work to introduce the use of model counting to quantify performance of trained decision trees with respect to ground truth formulas and to quantify semantic differences among different decision trees.

Model Counting Applications in ML. The use of model counting in machine learning has focused largely on probabilistic reasoning [9, 16, 20, 40]. Recent work by Baluta [2] introduced model counting for quantifying differences between binarized neural networks [28] that admit a translation to SAT/CNF [45]. This translation enables our MCML metrics to generalize beyond decision trees and become applicable to quantify the performance of binarized neural networks with respect to the entire input space. The MCML metrics also directly generalize to other techniques that use different solvers for model counting, e.g., techniques based on ODDs and OBDDs [8, 53].

Verification and Testing of ML models. Verification and testing of machine learning models is an active area of research, including work on novel decision procedures such as Reluplex [27, 34], which has been optimized for the analysis of neural networks with ReLU activation functions, testing trained models [46, 60], applying symbolic execution [24, 59], and inferring verifiable policies by mimicking deep reinforcement learning agents [3, 68]. A key difference between

MCML and previous work on verification of properties of trained models is the focus of previous work on either verifying properties of one trained model, or checking equivalence or implication between two models. In contrast, MCML metrics apply in a more general setting, even when two models are neither equivalent nor such that one implies the other.

For translating decision trees to CNF formulas, Hastad [25] introduced the idea to represent a decision tree as a Disjunctive Normal Form (DNF) formula. He also showed that these DNF formulas are convertible to Conjunctive Normal Form (CNF) formulas. MCML leverages this work and uses CNF formulas of trained decision trees to create the quantification problems. While the formulas MCML creates are optimal in terms of the size of the CNF formula with respect to the given input tree in the general case, an alternative approach is to use re-writing, e.g., aka compilation [8, 53], to create smaller CNF formulas in some specific cases. However, re-writing itself has a cost, which can be substantial for non-trivial formulas, e.g., with hundreds of variables (as for our subjects). Moreover, when compilation transforms a decision tree to a simpler decision tree, MCML works *in tandem* with compilation: first apply re-writing to reduce the tree, and then use MCML’s translation to create a reduced CNF formula.

Analyzing Learnability. Efforts to understand the ability of a machine learning model to generalize are at the core of statistical learning theory. The key concepts for establishing such results include the Probably Approximately Correct (PAC) learning framework [65], Vapnik- Chervonenkis (VC) theory [67], and the general learning setting [52, 66]. These techniques enable formal analytical characterization of the number of examples needed to train models for binary classification tasks and may provide useful intuition and offer guidance about the design of learning algorithms. Blumer [4] showed that the finiteness of Vapnik-Chervnenkis (VC) dimension is a basic requirement for distribution-free learning. In particular, by relying on the VC dimension, he analyzed the performance in terms of closure and complexity and provided a detailed set of conditions for learnability.

The PAC learnability concepts, however, provide limited insight in the performance of methods for learning relational properties. This is because the number of positively labeled samples in the domain set (i.e., the space of relational properties) is often orders of magnitude smaller than the number of negatively labeled ones. Therefore, Precision, Recall and F1-score, formally defined in Section 5, are rather more informative performance metrics than the average 0-1 loss; the latter is the focus of the PAC learnability analysis. In contrast, MCML can precisely quantify generalizability both in terms of the accuracy (i.e., the 0-1 loss) as well as precision, recall and F1-score with respect to a given ground truth.

Indeed, our study complements existing theoretical frameworks by providing insights based on empirical evidence from controlled experiments using various relational properties that are common in software systems.

```

sig S { r: set S } // r is a binary relation of type SxS
pred Reflexive() { all s: S | s->s in r }
pred Symmetric() {
  all s, t: S | s->t in r implies t->s in r
}
pred Transitive() { all s, t, u: S |
  s->t in r and t->u in r implies s->u in r
}
pred Equivalence() {
  Reflexive and Symmetric and Transitive
}
E4: run Equivalence for exactly 4 S

```

Figure 1. Alloy specification with one set S , one binary relation r , four predicates (*Reflexive*, *Symmetric*, *Transitive*, and *Equivalence*), and one command (E4).

Learning Program Properties. In the context of learning properties of code, specifically of dynamic data structures in Java programs, two recent projects [43, 64] used a variety of machine learning models to show the effectiveness of off-the-shelf models. However, both these projects used the traditional metrics with training and test datasets, and did not evaluate the performance of the trained models with respect to the entire (bounded) input spaces. More broadly, machine learning enabled automated detection of program errors and repair of faults [6, 41].

Alloy. Alloy has been used in several projects: for design and modeling of software systems [1, 11, 31, 36, 70, 73]; for software analyses, including deep static checking [17, 32], systematic testing [42], data structure repair [50, 71], and automated debugging [23]; for analysis of hardware systems [10, 61, 62]; and testing and studying model counters [44, 69]. MCML introduces the use of Alloy for creating training and test data for machine learning models and leverages Alloy’s backend for creating CNF formulas that represent the ground truth.

3 Example

Figure 1 shows an Alloy specification that declares a set (sig) S of atoms, a binary relation $r : S \times S$, and four predicates (i.e., formulas) that specify reflexive, symmetric, transitive, and equivalence relations. The keyword “all” is universal quantification, “in” is subset, “and” is conjunction, and “implies” is implication. The operator ‘ \rightarrow ’ is Cartesian product; for scalars s and t , $s \rightarrow t$ is the pair (s, t) . The command “E4: run Equivalence for exactly 4 S” instructs the Alloy analyzer to solve the formula(s) that define equivalence relation with respect to a *scope*, i.e., bound, of exactly 4 atoms in set S . The analyzer uses the bound to translate the Alloy specification to a propositional satisfiability formula, and uses off-the-shelf SAT solvers to solve it. The Alloy analyzer supports *incremental* solvers that can enumerate all solutions.

Executing the command E4 and enumerating all solutions creates the 5 solutions illustrated in Figure 2. Note, each solution is non-isomorphic. The Alloy analyzer adds symmetry breaking predicates during the translation, which breaks several (but, in general, not all) symmetries [54]. The

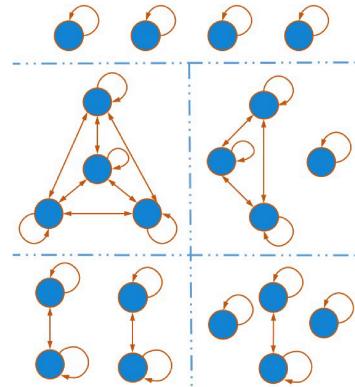


Figure 2. 5 non-isomorphic equivalence relations.

space of all candidate solutions for the command E4 has size $2^{16}=65,536$ since each candidate is a 4×4 matrix of boolean variables. Note, how quickly the state space grows as the number of vertices increases, e.g., for just 7 vertices, the state space has size $2^{(7^2)} = 2^{49}$ which is greater than 10^{14} .

ApproxMC. To illustrate the use of ApproxMC, consider estimating the number of solutions for the Equivalence predicate with scope 20. The Alloy command “E20: run Equivalence for exactly 20 S” defines the constraint solving problem. The Alloy analyzer translates this problem to a CNF formula that has 18,666 variables (of which 400 are primary variables) and 27,202 clauses. ApproxMC solves this CNF formula in 17.8 seconds and reports an approximate model count of 11,264. The exact model count, which we calculate using the Alloy analyzer (which uses an enumerating SAT solver), is 10,946, i.e., the ApproxMC count is within 3% error rate.

ProjMC. To illustrate the use of ProjMC, consider computing the exact model count for the Equivalence predicate for scope 20. Given the CNF formula for the Alloy command E20, ProjMC reports the exact model count of 10,946 in 351.1 seconds. The count reported by ProjMC is, as expected, the exact number of solutions we get using the Alloy analyzer.

4 MCML Approach

Our approach, called *MCML*, introduces the use of model counting for quantifying performance of decision tree models in machine learning (ML). MCML is embodied by two techniques: *Acc_{MC}*, which quantifies the performance of trained ML models with respect to ground truth formulas; and *Diff_{MC}*, which quantifies the semantic differences between two trained models. Both *Acc_{MC}* and *Diff_{MC}* compute the results with respect to the entire input space and do not require any datasets.

We view a decision tree as a set of *paths*; any input follows exactly one path, and each path is a conjunction of *branch conditions* such that each condition contains one input variable [29]. Since our

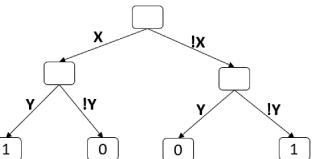
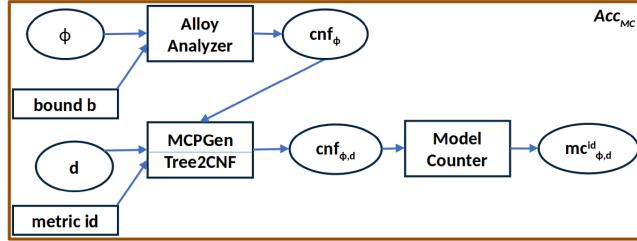
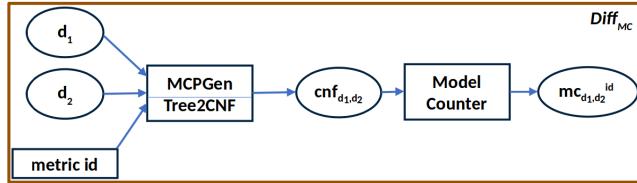


Figure 3. A decision tree.



(a) Quantifying the performance of decision tree d w.r.t. ground truth ϕ .



(b) Quantifying differences between trees d_1 and d_2 .

Figure 4. MCML Framework.

focus is on relational properties over graphs that are represented using adjacency matrices, the input variables for the decision tree are all binary and so is the output. Moreover, each branch condition on any path is simply of the form either “(v)”, i.e., input v is “0” (*false*), or “(v)”, i.e., input v is “1” (*true*). Therefore, each branch condition is simply a *literal*, i.e., a variable or its negation. Figure 3 illustrates a decision tree with 2 inputs (x and y) and 4 paths ($[x, y]$, $[x, !y]$, $[!x, y]$, and $[!x, !y]$).

Acc_{MC}: Quantifying model performance. Assume property ϕ (ground truth) is learned by a decision tree d that is a binary classifier with n inputs and so the input space has size 2^n . Assume d has t unique paths p_1, \dots, p_t ($t \geq 0$) that predict label *true*, and f unique paths q_1, \dots, q_f ($f \geq 0$) that predict label *false*. For path p , let $\psi(p)$ be the conjunction of branch conditions along p ; we refer to $\psi(p)$ as the *path condition* for p [47]. We define the following four metrics based on model counting to generalize (for the entire 2^n input space) the traditional metrics of *true positives* (*tp*), *false positives* (*fp*), *true negatives* (*tn*), and *false negatives* (*fn*), where $mc(a, b)$ represents the model count for the formula a with respect to bound b :

$$tp(\phi, d) = mc(\phi \wedge \bigvee_{i=1}^t \psi(p_i), n) \quad (1)$$

$$fp(\phi, d) = mc(\neg\phi \wedge \bigvee_{i=1}^t \psi(p_i), n) \quad (2)$$

$$tn(\phi, d) = mc(\neg\phi \wedge \bigvee_{i=1}^f \psi(q_i), n) \quad (3)$$

$$fn(\phi, d) = mc(\phi \wedge \bigvee_{i=1}^f \psi(q_i), n) \quad (4)$$

The formula $\bigvee_{i=1}^t \psi(p_i)$ characterizes the inputs that the decision tree t classifies as *true*; the formula is a disjunction

of the path conditions for the paths that predict *true* and any input that is classified as *true* must be a solution to exactly one of these paths. For example, for the decision tree depicted in Figure 3, the path condition formula for label *true* should be $(x \wedge y) \vee (\neg x \wedge \neg y)$. Likewise, the formula $\bigvee_{i=1}^f \psi(q_i)$ characterizes the inputs that the tree t classifies as *false*.

While a direct application of our metrics requires the ground truth formula ϕ that characterizes the property of interest, the metrics are also applicable when ϕ is not known. Specifically, the metrics naturally generalize to allow quantifying *differences* in two trained models, again without requiring any datasets for quantification.

Diff_{MC}: Quantifying model differences. Let d_1 and d_2 be decision trees that are trained as binary classifiers (using the same or different datasets) with n inputs. Let d_1 have paths p_1, \dots, p_{t_1} ($t_1 \geq 0$) that predict *true* and paths q_1, \dots, q_{f_1} ($f_1 \geq 0$) that predict *false*. Let d_2 have paths r_1, \dots, r_{t_2} ($t_2 \geq 0$) that predict *true* and paths s_1, \dots, s_{f_2} ($f_2 \geq 0$) that predict *false*. The following four metrics quantify the number of inputs (in the entire 2^n input space) for which d_1 and d_2 make the same decision – *tt*, i.e., both predict *true*, or *ff*, i.e., both predict *false* – and numbers of inputs for which the decisions differ – *tf*, i.e., d_1 predicts *true* but d_2 predicts *false*, and *ft*, i.e., d_1 predicts *false* but d_2 predicts *true*:

$$tt(d_1, d_2) = mc(\bigvee_{i=1}^{t_1} \psi(p_i) \wedge \bigvee_{j=1}^{t_2} \psi(r_j), n) \quad (5)$$

$$tf(d_1, d_2) = mc(\bigvee_{i=1}^{t_1} \psi(p_i) \wedge \bigvee_{j=1}^{f_2} \psi(s_j), n) \quad (6)$$

$$ff(d_1, d_2) = mc(\bigvee_{i=1}^{f_1} \psi(q_i) \wedge \bigvee_{j=1}^{f_2} \psi(s_j), n) \quad (7)$$

$$ft(d_1, d_2) = mc(\bigvee_{i=1}^{f_1} \psi(q_i) \wedge \bigvee_{j=1}^{t_2} \psi(r_j), n) \quad (8)$$

The semantic *difference* (*diff*) in d_1 and d_2 is quantified as the ratio of the number of inputs x such that decisions for x differ between d_1 and d_2 , to the total number of inputs; and the *similarity* (*sim*) in d_1 and d_2 is quantified as the ratio of the number of inputs x such that predicted labels for x match between d_1 and d_2 , to the total number of inputs:

$$diff(d_1, d_2) = \frac{tf(d_1, d_2) + ft(d_1, d_2)}{2^n} \quad (9)$$

$$sim(d_1, d_2) = \frac{tt(d_1, d_2) + ff(d_1, d_2)}{2^n} \quad (10)$$

$$\text{i.e., } sim(d_1, d_2) = 1 - diff(d_1, d_2) \quad (11)$$

Framework. We embody the metrics into a prototype framework that translates decision trees to CNF and leverages off-the-shelf model counters. Figure 4a illustrates the key steps

to quantify the performance of a trained decision tree model d with respect to the ground truth ϕ written in Alloy. ϕ is translated by the Alloy analyzer with respect to bound b (e.g., 20 vertices) into a CNF formula cnf_{ϕ} . The module *MCPGen* takes as input d and cnf_{ϕ} , translates the relevant parts of d with respect to the desired metric (tp , fp , tn , and fn) into CNF using the sub-module *Tree2CNF*, and outputs the CNF formula $cnf_{\phi,d}$ which defines the model counting problem. The CNF is an input to the model counter (e.g., *ApproxMC* and *ProjMC*) that outputs the approximate model count and exact model count respectively. Figure 4b illustrates the key steps to quantifying the semantic differences between two trained decision tree models d_1 and d_2 .

Translating decision tree logic to CNF. The goal of the sub-module *Tree2CNF* is to translate a formula ρ of the form $\vee_{i=1}^k \psi(p_i)$, which represents either all the paths that predict label *true* or all the paths that predict label *false*, to CNF. ρ is originally in *disjunctive normal form* and can be translated to CNF using various techniques. One standard technique is to apply propositional equivalences and De Morgan's laws [33]; however, this technique can lead to a blow-up in the size of the formula that can negatively impact the back-end solver's performance. Another standard technique is to apply the Tseitin transformation [63] which creates formulas linear in the size of the input; however, this technique uses auxiliary variables and creates an equisatisfiable (but not necessarily equivalent) formula which can have a different model count than the original formula.

Our translation uses the following observation [25] that allows a translation which does not cause a blow up in the formula size and preserves the model counts: any input that does not get classified as *true* gets classified as *false* and vice versa. Therefore, if each $\psi(p_i)$ is a path condition for a path that leads to label *true* and $\rho = \vee_{i=1}^k \psi(p_i)$ represents the decision tree logic that predicts *true*, the negation of ρ characterizes the decision tree logic that predicts *false*. The formula $\neg\rho$ immediately simplifies to CNF because $\neg\rho = \neg \vee_{i=1}^k \psi(p_i) = \wedge_{i=1}^k \neg\psi(p_i)$ and each $\psi(p_i)$ is a conjunction of literals; therefore, each $\neg\psi(p_i)$ is a disjunction of literals. For example, for the decision tree in Figure 3, the path condition formula for label *false* is the negation of the path condition formula for label *true*, that is $!((x \wedge y) \vee (!x \wedge !y))$ which can be directly translated into CNF formula $(!x \vee !y) \wedge (x \vee y)$.

Analysis. Our translation creates a compact formula *directly* in CNF without introducing any auxiliary variables – both for output label *true* and for label *false*. For a decision tree with n leaf nodes, the number of path conditions (PCs) is n , i.e., *linear* in the tree size. If there are k features, the size of the CNF formula is $O(nk)$: each PC has $\leq k$ conditions, each condition is a literal, and the formula is a conjunction of the negation of each PC. Thus, in terms of the size of the formula for the back-end solver, the translation is an optimal choice for CNF with no auxiliary variables used in the general case.

Our tool embodiment of MCML supports two state-of-the-art back-ends: 1) the exact model counter *ProjMC* [39]; and 2) the approximate model counter *ApproxMC* [57].

5 Study

This section describes our study methodology and summarizes the evaluation results for the 16 relational properties (Table 1) and 6 machine learning models (decision trees (*DT*), random forest decision trees (*RFT*), Adaboost decision trees (*ABT*), gradient boosting decision trees (*GBDT*), support vector machines (*SVM*) and multi-layer perceptrons (*MLP*) studied. The properties include: anti-symmetric, bijective, connex, equivalence, function, functional, injective, irreflexive, non-strict order, partial order, pre-order, reflexive, strict order, surjective, total order and transitive.

Generation of positive and negative samples. For each property and scope, we use Alloy to create two sets of solutions – positive solutions, which satisfy the property, and negative solutions, which negate it. For positive solutions, we use the set of *all* solutions enumerated by Alloy's backend SAT engine (with respect to the scope). Due to the complex nature of the properties, the number of negative solutions is much larger than the number of positive solutions, and it is infeasible to exhaustively enumerate all negative solutions. To create the set of negative solutions, we sample them at random from the entire state space (with respect to the scope). Specifically, to sample a negative solution, we first create a candidate solution at random and then confirm that it does not satisfy the property using the Alloy Evaluator, which simply evaluates the corresponding Alloy formula by replacing all the variables in the formula with the values in the given candidate solution, and using constant propagation (without any constraint solving).

Each data sample has a feature vector and a binary label that is 1 for the positive class and 0 for the negative class. The features linearly represent the adjacency matrix, e.g., if the scope of a property is 7, a 49-bit feature vector is used.

Selection of scope and symmetry breaking. We study each property under two settings for symmetry breaking. The first setting is when Alloy's default symmetry breaking constraints are used; in this setting we choose the smallest scope such that there are $\geq 10,000$ positive solutions. The second setting is when no symmetry breaking constraints are added by the Alloy analyzer; in this setting, we choose the smallest scope such that there are $\geq 90,000$ positive solutions.

Table 1 tabulates for each property, the scope (i.e., number of atoms in set S), the size of the input space, the number of positive samples created by Alloy, the count of positive samples estimated by *ApproxMC*, the count of positive samples estimated by *ApproxMC* when symmetry breaking is turned off, the count of positive samples given by *ProjMC* and the count of positive samples given by *ProjMC* when symmetry breaking is turned off.

Table 1. Subject properties and model counts. For each property, the scope used, the size of the state space, the number of positive samples created by Alloy using its default enumeration with symmetry breaking, the count estimated by ApproxMC with symmetry breaking, the count estimated by ApproxMC without symmetry breaking, the count computed by projMC with symmetry breaking, and the count computed by ProjMC without symmetry breaking are shown. “-” indicates time-out.

Property	Scope	State Space	Valid-SymBr (Alloy)	Est-Valid-SymBr (ApproxMC)	Est-Valid-NoSymBr (ApproxMC)	Valid-SymBr (ProjMC)	Valid-NoSymBr (ProjMC)
<i>Antisymmetric</i>	5	2^{25}	56723	55296	1998848	56723	1889568
<i>Bijective</i>	14	2^{196}	25043	25088	-	25043	-
<i>Connex</i>	6	2^{36}	148884	147456	14680064	148884	14348907
<i>Equivalence</i>	20	2^{400}	10946	11264	-	10946	-
<i>Function</i>	8	2^{64}	16531	16640	17563648	16531	16777216
<i>Functional</i>	8	2^{64}	35017	35328	48234496	35017	43046721
<i>Injective</i>	8	2^{64}	16531	16640	17563648	16531	16777216
<i>Irreflexive</i>	5	2^{25}	35886	36352	2686976	35886	1048576
<i>NonStrictOrder</i>	7	2^{49}	26387	26112	6422528	26387	6129859
<i>PartialOrder</i>	6	2^{36}	82359	88064	8126464	82359	8321472
<i>PreOrder</i>	7	2^{49}	43651	43008	9175040	43651	9535241
<i>Reflexive</i>	5	2^{25}	35886	35840	1048576	35886	1048576
<i>StrictOrder</i>	7	2^{49}	26387	29184	6815744	26387	6129859
<i>Surjective</i>	14	2^{196}	25043	25088	-	25043	-
<i>TotalOrder</i>	13	2^{169}	15511	14848	5502926848	15511	-
<i>Transitive</i>	6	2^{36}	95564	102400	9306112	95564	9415189

Training models. A key decision in evaluating ML models is to select the training and test ratio. We use 5 different ratios for *training:test*, specifically 75:25, 50:50, 25:75 10:90 and 1:99 to evaluate a wide range of settings. There is no overlap in the training and test datasets. We performed the experiments using basic out-of-the-box models. We explored tuning the hyper-parameters but the results were only marginally improved and the small increase in accuracy was offset by a much larger increase in time spent on tuning. We report the results obtained using the basic models without tuning of hyper-parameters. We used Python programming language and Scikit-Learn Library to implement machine learning models. We kept time-out of 5000 seconds which is common in field of model counting.

Performance metrics. We used four standard metrics to determine the quality of classification results: *accuracy*, i.e., $\frac{TP+TN}{TP+FP+TN+FN}$; *precision*, i.e., $\frac{TP}{TP+FP}$; *recall*, i.e., $\frac{TP}{TP+FN}$; and *F1-score*, i.e., $\frac{2*Precision*Recall}{Precision+Recall}$.

Due to space limitations we show results for select properties and focus on ApproxMC. Complete results of all properties for ApproxMC and ProjMC with respect to all ratios and list of metadata items such as number of primary variables, total variables and total clauses for each property are available at : <https://github.com/muhammadusman93/MCML-PLDI2020>. Alloy specifications of the properties and the detailed results are also available at the GitHub repository. All experiments were performed on Ubuntu 16.04 with an Intel Core-i7 8750H CPU (2.20 GHz) and 16GB RAM.

Research questions. We answer the following research questions in our empirical study.

- **RQ1:** How effective are ML models in learning relational properties?

Table 2. Classification results on the test set for *PartialOrder*. “*Training:Test*” ratio is shown in the left column, and default Alloy symmetry breaking is used to create datasets.

Ratio	Model	Accuracy	Precision	Recall	F1-score
75:25	DT	0.9996	0.9992	1.0000	0.9996
	RFT	0.9999	0.9998	1.0000	0.9999
	GBDT	0.9951	0.9905	0.9997	0.9951
	ABT	0.9412	0.9396	0.9425	0.9411
	SVM	0.9993	0.9985	1.0000	0.9993
	MLP	0.9996	0.9993	1.0000	0.9996
25:75	DT	0.9983	0.9968	0.9999	0.9983
	RFT	0.9996	0.9993	1.0000	0.9996
	GBDT	0.9952	0.9906	0.9998	0.9952
	ABT	0.9419	0.9401	0.9441	0.9421
	SVM	0.9987	0.9973	1.0000	0.9987
	MLP	0.9991	0.9982	1.0000	0.9991
1:99	DT	0.9798	0.9656	0.9949	0.9801
	RFT	0.9931	0.9937	0.9926	0.9931
	GBDT	0.9904	0.9846	0.9964	0.9904
	ABT	0.9387	0.9351	0.9428	0.9389
	SVM	0.9769	0.9767	0.9772	0.9770
	MLP	0.9909	0.9867	0.9953	0.9910

- **RQ2:** How well do ML (decision tree) models generalize (outside of the test set)?
- **RQ3:** How do symmetries in the dataset affect performance of the ML models?
- **RQ4:** How does mismatch in symmetries present in the training and evaluation sets affect the performance?
- **RQ5:** What is the quantitative difference between two decision trees models?

5.1 Answers to the Research Questions

- 5.1.1 **RQ1: How Effective are ML Models in Learning Relational Properties?** Table 2 summarizes performance of ML classifiers for one selected property (*PartialOrder*). The

Table 3. Decision tree performance on the test set (symmetries broken) and in respect to the entire state space encoded by ground truth formula ϕ (constrained by symmetry breaking).

Property	Accuracy Test	Precision Test	Recall Test	F1-score Test	Accuracy ϕ	Precision ϕ	Recall ϕ	F1-score ϕ	Time[s]
<i>Antisymmetric</i>	0.9913	0.9871	0.9957	0.9914	0.9690	0.4576	0.9960	0.6271	0.9
<i>Bijective</i>	0.9971	0.9942	1.0000	0.9971	0.9729	0.0000	1.0000	0.0000	219.3
<i>Connex</i>	0.9996	0.9993	1.0000	0.9996	0.9992	0.0698	1.0000	0.1304	1.3
<i>Equivalence</i>	0.9958	0.9920	0.9997	0.9958	-	-	0.9997	-	-
<i>Function</i>	0.9954	0.9928	0.9980	0.9954	0.9924	0.0000	0.9982	0.0000	6.0
<i>Functional</i>	0.9991	0.9983	0.9998	0.9991	0.9943	0.0000	0.9998	0.0000	6.3
<i>Injective</i>	0.9979	0.9961	0.9998	0.9980	0.8889	0.0000	0.9998	0.0000	6.1
<i>Irreflexive</i>	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	0.4
<i>NonStrictOrder</i>	0.9994	0.9989	1.0000	0.9994	0.9944	0.0000	1.0000	0.0000	3.1
<i>PartialOrder</i>	0.9963	0.9936	0.9990	0.9963	0.9675	0.0059	0.9991	0.0116	2.3
<i>PreOrder</i>	0.9992	0.9985	0.9999	0.9992	0.9909	0.0000	0.9999	0.0000	3.5
<i>Reflexive</i>	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	0.4
<i>Strictorder</i>	0.9991	0.9982	1.0000	0.9991	0.9915	0.0000	1.0000	0.0000	3.0
<i>Surjective</i>	0.9980	0.9961	1.0000	0.9980	0.9993	0.0000	1.0000	0.0000	208.1
<i>TotalOrder</i>	0.9994	0.9988	1.0000	0.9994	0.9983	0.0000	1.0000	0.0000	95.7
<i>Transitive</i>	0.9949	0.9910	0.9989	0.9949	0.9866	0.0030	0.9990	0.0059	2.1

dataset is generated using default Alloy symmetry breaking. The dataset is further split into training and test datasets using different ratios (75:25, 25:75 and 1:99) to measure the performance change when using different splits of the datasets. All models exhibit high accuracy and F1-score, where the accuracy is in range [0.94, 1.00], the precision is in range [0.94, 1.00], the recall is in range [0.94, 1.00], and the F1-score is in range [0.94, 1.00]. Overall, all models achieve good performance, and surprisingly even for a small training:test ratio (1 : 99) models achieve good performance. We have also performed similar experiments for all other relational properties (available at GitHub repository), and results are similar as for the *PartialOrder* – accuracy is in range [0.92, 1.00], and the F1-score is in range [0.92, 1.00], where the lowest accuracy is for *Antisymmetric* property. It is surprising to note that on most of the properties, all models report accuracy ≥ 0.92 even when trained on only 1% of dataset. The results achieved demonstrate that relational properties are learnable, and that even simple ML models are effective in learning them.

5.1.2 RQ2: How Well do ML (Decision Tree) Models Generalize (Outside of the Test Set)? In *RQ1* we show that ML models can exhibit high accuracy when learning relational properties. Now, we move to the question of how well do those models generalize outside of the test set. Our framework allows us to answer that question for decision tree models, which is focus of this section. Specifically, we answer the question of how well the previously trained decision tree models perform in respect to the entire input space. Table 3 compares performance of the (same) decision trees models when evaluated: (1) on the test set, and (2) in respect to the entire input space. Decision trees are trained using 10% of the constructed dataset (used in the previous section), while the test set represents the remaining part of the dataset.

Table 4. Classification results on the test set for *PartialOrder* property. “*Training:Test*” ratio is shown in the left column, and symmetry breaking is turned off when creating datasets.

Ratio	Model	Accuracy	Precision	Recall	F1-score
75:25	DT	0.9985	0.9970	1.0000	0.9985
	RFT	0.9981	0.9963	1.0000	0.9981
	GBDT	0.9788	0.9593	1.0000	0.9792
	ABT	0.8414	0.8636	0.8098	0.8359
	SVM	0.9940	0.9881	1.0000	0.9940
	MLP	0.9989	0.9978	1.0000	0.9989
25:75	DT	0.9966	0.9935	0.9997	0.9966
	RFT	0.9963	0.9927	1.0000	0.9963
	GBDT	0.9780	0.9578	1.0000	0.9784
	ABT	0.8394	0.8613	0.8085	0.8341
	SVM	0.9901	0.9806	1.0000	0.9902
	MLP	0.9977	0.9954	1.0000	0.9977
1:99	DT	0.9692	0.9511	0.9893	0.9698
	RFT	0.9821	0.9699	0.9950	0.9823
	GBDT	0.9734	0.9495	1.0000	0.9741
	ABT	0.8352	0.8409	0.8268	0.8338
	SVM	0.9635	0.9320	1.0000	0.9648
	MLP	0.9842	0.9704	0.9988	0.9844

Time shown is the total time taken by MCML for computing all 4 performance metrics, where MCML leverages the model counting techniques to predict the number of true positives and negatives, and false positives and negatives. To evaluate in respect to the entire input space we use the ground truth formula constrained with symmetry breaking conditions ¹.

When decision trees are evaluated on the test set, accuracy, recall and F1-score is ≥ 0.99 while precision is ≥ 0.98 . However, when decision trees are evaluated on ground truth, minimum accuracy decreases to 0.89, minimum precision and F1-score decreases to 0.00 and minimum recall remains at 0.99. Precision and F1-score is spread all over the range

¹Symmetry breaking conditions are added so as to make distributions of examples similar to the ones present in the training set. We later show evaluation where we remove this constraint.

Table 5. Decision tree performance on the test set and in respect to the entire state space encoded by ground truth formula ϕ

Property	Accuracy Test	Precision Test	Recall Test	F1-score Test	Accuracy ϕ	Precision ϕ	Recall ϕ	F1-score ϕ	Time[s]
<i>Antisymmetric</i>	0.9997	0.9996	0.9998	0.9997	0.9996	0.9935	0.9998	0.9967	2.1
<i>Bijective</i>	0.9991	0.9982	1.0000	0.9991	0.9981	0.0000	1.0000	0.0000	225.8
<i>Connex</i>	0.9957	0.9933	0.9982	0.9957	0.9935	0.2258	0.9985	0.3683	0.8
<i>Equivalence</i>	0.9997	0.9994	1.0000	0.9997	0.9995	0.0000	1.0000	0.0000	34.1
<i>Function</i>	0.9946	0.9900	0.9993	0.9946	0.9899	0.0001	0.9993	0.0001	2.4
<i>Functional</i>	0.9968	0.9940	0.9997	0.9969	0.9945	0.0003	0.9997	0.0006	2.6
<i>Injective</i>	0.9968	0.9940	0.9997	0.9969	0.9877	0.0001	0.9989	0.0001	2.3
<i>Irreflexive</i>	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	0.5
<i>NonStrictOrder</i>	0.9990	0.9985	0.9994	0.9990	0.9983	0.0011	0.9995	0.0022	1.9
<i>PartialOrder</i>	0.9934	0.9879	0.9991	0.9935	0.9864	0.2407	0.9992	0.3879	1.2
<i>PreOrder</i>	0.9985	0.9974	0.9996	0.9985	0.9972	0.0012	0.9997	0.0024	2.0
<i>Reflexive</i>	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	0.5
<i>StrictOrder</i>	0.9988	0.9979	0.9997	0.9988	0.9979	0.0009	0.9998	0.0019	1.9
<i>Surjective</i>	0.9988	0.9979	0.9997	0.9988	0.9984	0.0000	1.0000	0.0000	283.3
<i>TotalOrder</i>	0.9999	0.9997	1.0000	0.9999	0.9997	0.0000	1.0000	0.0000	8.4
<i>Transitive</i>	0.9999	0.9997	1.0000	0.9999	0.9760	0.1588	0.9902	0.2737	2.0

0 to 1, and is generally low. In-fact, out of 15 properties that did not time-out, precision is around 0 on 12 properties. This is because of high false positive rate showing that the decision trees are classifying many negative instances as positives. This indicates that the models are biased towards classifying an example as positive, likely learning patterns present in training dataset but *not generalizable*—applying the models outside of the dataset will likely incur many false positives. For the properties Reflexive and Irreflexive, the models continue to have perfect performance since establishing these properties requires only checking the diagonal. To summarize the findings, while the results in previous section showed encouraging results of learnability of relational properties, evaluation in respect to the entire input space shows concerning issues (with false positives) if models are to be used in the wild.

In summary, MCML’s ability to quantify w.r.t. the entire state space is of unique value as it quantifies model’s generalizability, and avoids a false sense of confidence in traditional ML metrics. In addition, MCML is also time efficient as it reported results for 12 properties within 10 seconds and 3 properties within 220 seconds. It timed-out on only one property (*Equivalence*), where the state space for possible solutions is around 2^{400} .

5.1.3 RQ3: How do Symmetries in the Dataset Affect Performance of the ML Models? We now move on to study the effect of symmetries in the dataset on ML models performance. We perform experiments similar as in the *RQ1* and *RQ2* now shown in Table 4 and Table 5, but without performing symmetry breaking on training and evaluation sets.

Table 4 summarizes performance of ML classifiers for one selected property (*PartialOrder*). The dataset is generated without performing symmetry breaking. Across all ratios, the accuracy is ≥ 0.83 , the precision is ≥ 0.84 , the recall is ≥ 0.80 , and the F1-score is ≥ 0.83 . We have performed

experiments for other properties (available at GitHub repository), and noticed similar trends. Across all 16 properties, all 5 ratios, and all 6 ML models, the accuracy is ≥ 0.82 , the precision is ≥ 0.81 , the recall is ≥ 0.81 , and the F1-score is ≥ 0.82 . Similar to before, all models achieve good performance, and good performance is preserved even for small training:test ratio. However, in comparison to the previous results (Table 2) there is a noticeable decrease in terms of accuracy and F1-score, introduced by symmetries. This shows that if symmetries are broken (in both train and test set) than model can better learn the properties, as it only uses the representatives of distinct groups for learning. This is analogous to training a digit classifier where all digits are upright. Introducing symmetries, e.g., digits at different orientations makes it harder for classifier to learn well.

Table 5 compares results of the decision tree models on the test set and in respect to the entire input space. Unlike in *RQ2*, now the training set is generated without symmetry breaking, and the entire input space encoded by the formula ϕ is now not constrained with symmetry breaking. On the test set accuracy, recall and F1-score is ≥ 0.99 and precision is ≥ 0.98 . However, in respect to the entire input space, minimum accuracy is 0.98, minimum precision and F1-score decreases to 0.00, but minimum recall remains at 0.99. The outliers are properties Reflexive and Irreflexive, where the models continue to have perfect performance. In summary, the results shows a similar trend as before, where even with enhancing the training set with symmetric examples, decision tree models still generalize poorly.

5.1.4 RQ4: How does Mismatch in Symmetries Present in the Training and Evaluation Sets Affect the Performance? We now move to the next question of how the mismatch in presence of symmetries in training and evaluation sets affect the performance. Specifically, we look at the two scenarios: (1) symmetries are not present in the training set

Table 6. Decision tree performance on the test set (symmetries broken) and in respect to the entire state space encoded by ground truth formula ϕ .

Property	Accuracy Test	Precision Test	Recall Test	F1-score Test	Accuracy ϕ	Precision ϕ	Recall ϕ	F1-score ϕ	Time[s]
<i>Antisymmetric</i>	0.9913	0.9871	0.9957	0.9914	0.9442	0.5098	0.2241	0.3114	1.3
<i>Bijective</i>	0.9971	0.9942	1.0000	0.9971	-	-	-	-	-
<i>Connex</i>	0.9996	0.9993	1.0000	0.9996	0.9992	0.0520	0.1731	0.0800	1.9
<i>Equivalence</i>	0.9958	0.9920	0.9997	0.9958	-	-	-	-	-
<i>Function</i>	0.9954	0.9928	0.9980	0.9954	0.9922	0.0000	0.0667	0.0000	86.6
<i>Functional</i>	0.9991	0.9983	0.9998	0.9991	0.9980	0.0000	0.2907	0.0000	277.5
<i>Injective</i>	0.9979	0.9961	0.9998	0.9980	0.9961	0.0000	0.2565	0.0000	79.6
<i>Irreflexive</i>	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	0.5
<i>NonStrictOrder</i>	0.9994	0.9989	1.0000	0.9994	0.9990	0.0000	0.6263	0.0000	8.9
<i>PartialOrder</i>	0.9963	0.9936	0.9990	0.9963	0.9937	0.0068	0.3435	0.0134	2.9
<i>PreOrder</i>	0.9992	0.9985	0.9999	0.9992	0.9987	0.0000	0.5180	0.0000	16.6
<i>Reflexive</i>	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	0.6
<i>StrictOrder</i>	0.9991	0.9982	1.0000	0.9991	0.9983	0.0000	0.4660	0.0000	8.3
<i>Surjective</i>	0.9980	0.9961	1.0000	0.9980	-	-	-	-	-
<i>TotalOrder</i>	0.9994	0.9988	1.0000	0.9994	0.9990	0.0000	0.4737	0.0000	3059.5
<i>Transitive</i>	0.9949	0.9910	0.9989	0.9949	0.9914	0.0038	0.2394	0.0074	3.6

Table 7. Decision tree performance on the test set and in respect to the entire state space encoded by ground truth formula ϕ (constrained by symmetry breaking).

Property	Accuracy Test	Precision Test	Recall Test	F1-score Test	Accuracy ϕ	Precision ϕ	Recall ϕ	F1-score ϕ	Time[s]
<i>Antisymmetric</i>	0.9997	0.9996	0.9998	0.9997	1.0000	1.0000	1.0000	1.0000	0.7
<i>Bijective</i>	0.9991	0.9982	1.0000	0.9991	0.9992	0.0000	1.0000	0.0000	10.1
<i>Connex</i>	0.9957	0.9933	0.9982	0.9957	0.9959	0.1510	1.0000	0.2624	0.6
<i>Equivalence</i>	0.9997	0.9994	1.0000	0.9997	0.9991	0.0000	1.0000	0.0000	15.7
<i>Function</i>	0.9946	0.9900	0.9993	0.9946	0.9922	0.0000	0.9912	0.0001	1.6
<i>Functional</i>	0.9968	0.9940	0.9997	0.9969	0.9964	0.0003	1.0000	0.0005	1.5
<i>Injective</i>	0.9968	0.9940	0.9997	0.9969	0.9914	0.0002	1.0000	0.0003	1.5
<i>Irreflexive</i>	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	0.4
<i>NonStrictOrder</i>	0.9990	0.9985	0.9994	0.9990	0.9972	0.0016	1.0000	0.0032	1.5
<i>PartialPrder</i>	0.9934	0.9879	0.9991	0.9935	0.9862	0.2784	1.0000	0.4356	1.0
<i>PreOrder</i>	0.9985	0.9974	0.9996	0.9985	0.9971	0.0028	0.9997	0.0056	1.6
<i>Reflexive</i>	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	0.5
<i>StrictOrder</i>	0.9988	0.9979	0.9997	0.9988	0.9977	0.0021	1.0000	0.0041	1.5
<i>Surjective</i>	0.9988	0.9979	0.9997	0.9988	0.9968	0.0000	1.0000	0.0000	10.9
<i>TotalOrder</i>	0.9999	0.9997	1.0000	0.9999	0.9996	0.0000	1.0000	0.0000	9.4
<i>Transitive</i>	0.9999	0.9997	1.0000	0.9999	0.9906	0.1673	0.9912	0.2863	0.8

Table 8. Evaluating differences between decision tree models.

Subject	TT	TF	FT	FF	Diff	Time [s]
<i>Antisymmetric</i>	7.86E+05	3.28E+04	3.17E+04	3.36E+07	0.19	1.1
<i>Bijective</i>	9.19E+54	5.88E+56	3.86E+56	1.00E+59	0.96	117.8
<i>Connex</i>	3.62E+07	1.15E+07	1.21E+07	6.87E+10	0.03	1.7
<i>Equivalence</i>	7.88E+116	1.95E+118	3.91E+118	2.54E+120	2.25	852.8
<i>Function</i>	6.31E+16	8.11E+16	8.11E+16	1.84E+19	0.87	8.1
<i>Functional</i>	2.03E+16	1.58E+16	1.58E+16	1.84E+19	0.17	7.4
<i>Injective</i>	7.21E+16	0	0	1.84E+19	0.00	4.2
<i>Irreflexive</i>	1.05E+06	0	0	3.25E+07	0.00	0.5
<i>NonStrictOrder</i>	2.92E+11	2.92E+11	2.92E+11	5.63E+14	0.10	4
<i>PartialOrder</i>	3.02E+08	1.07E+08	1.17E+08	6.87E+10	0.32	2.4
<i>PreOrder</i>	3.78E+11	3.78E+11	3.78E+11	5.63E+14	0.13	3.8
<i>Reflexive</i>	1.05E+06	0	0	3.25E+07	0.00	0.5
<i>StrictOrder</i>	6.18E+11	3.44E+11	3.44E+11	5.63E+14	0.12	3.8
<i>Surjective</i>	2.45E+55	3.68E+56	3.68E+56	1.00E+59	0.73	119.2
<i>TotalOrder</i>	3.65E+47	3.65E+47	3.65E+47	7.48E+50	0.10	79.5
<i>Transitive</i>	4.36E+08	2.01E+08	1.93E+08	6.87E+10	0.57	2.4

but are in the evaluation set, and (2) symmetries are present in the training set but not in the evaluation.

Table 6 shows results of decision tree models trained and tested on datasets with symmetry breaking, but evaluated on the entire input space (without symmetry breaking). When decision trees are evaluated on the test set accuracy, recall and F1-score is ≥ 0.99 , while precision is ≥ 0.98 . However, when decision trees are evaluated on ground truth (without symmetry breaking constraints), minimum accuracy decrease to 0.94, minimum precision and F1-score decreases to 0.00, while minimum recall decrease to 0.06. Since symmetries are present in the entire input space, the performance decreases compared to the case where symmetries are also present in the training set (Table 5) These experiments show that decision trees perform worst when they are trained on the dataset without symmetries and evaluated w.r.t entire state space (with no symmetry breaking). This is expected result since the trained model did not see any symmetrical instances in the training dataset and is likely to perform incorrectly when tested on permutations of same instances. This is similar to an example of neural networks failing to generalize on digit recognition task, when seeing digits at different orientation than observed during the training [7].

Table 7 shows results of decision tree models trained and tested on datasets without symmetry breaking, and evaluated on the entire input space with added symmetry breaking constraints. When decision trees are evaluated on the test set, accuracy, recall and F1-score is ≥ 0.99 while precision is ≥ 0.98 . However, when decision trees are evaluated on ground truth (with added symmetry breaking), minimum accuracy remains at 0.99, minimum precision and F1-score decreases to 0.00 while minimum recall remains at 0.99. These results show that even when training set is richer (contains symmetries) than evaluation set, ML models still fail to generalize in respect to the entire input space.

5.1.5 RQ5: What is the Quantitative Difference Between Two Decision Tree Models? We next employ MCML to quantify the difference between the two decision tree models. Our framework allows us to get rigorous measure of the differences between the two models, i.e., the measure in respect to the entire input space (not just the train/test datasets). We trained two decision tree models for each property, using different values of hyperparameters, and measured difference between the two.

Table 8 shows the results for quantifying differences between 2 trained models, where we show number of examples in which both models predict true (TT), in which the first predicts true and second false (TF), and the other combinations denoted by FT and FF. The *Diff* column shows the percentage of cases in which the two models make a different prediction. In all cases, the difference is close to 0. MCML is able to quantify the differences between the two models in respect to the entire input space, which makes it a powerful

Table 9. Comparison of performance between traditional metrics and MCML for *antisymmetric* property for different class ratios (ratio of valid samples to invalid samples in training dataset); true ratio for the entire state space is 1:99.

Ratio of Valid:Invalid in training dataset	Traditional Precision	MCML Precision
99:1	0.98	0.19
90:10	0.97	0.21
75:25	0.98	0.44
50:50	0.99	0.46
25:75	0.99	0.56
10:90	0.99	0.75
1:99	1.00	0.97

technique for evaluation of models. For example, one can take a smaller (compressed) model and rigorously quantify it in respect to the larger model to see if it can be used as a replacement. MCML is able to detect this for 12 properties within 10 seconds and is able to detect all properties within 1000 seconds.

5.2 Discussion

5.2.1 Traditional Metrics and MCML. Tables 2 and 4 evaluate off-the-shelf models using traditional ML metrics based on training and test datasets. These metrics give a false sense of confidence, which our proposed MCML framework addresses. Tables 3, 5, 6 and 7 show that the majority of precision scores and F1-scores are low w.r.t the entire state space. An important characteristic of the properties we consider is that the number of positive cases is far too small compared to the number of negative cases. We believe new classifiers are needed to handle complex relational properties.

Alloy formulas are intuitively simple but semantically quite complex, and getting high precision with respect to the entire state space is hard. In fact, in only 2 cases (reflexive and irreflexive) the precision is 1, and that is because the properties can be checked simply by inspecting the diagonal entries – the trained decision trees indeed do so. In all other cases, the precision is very low (<0.1). MCML provides a new tool to rigorously study properties that seem “easy” to learn in the traditional setting of training and test data, but are actually “hard” to learn when viewed in the context of the entire state space. For example, for partial order, precision was 0.9936 in traditional setting whereas MCML reported that the precision is 0.0059 for the entire state space. This indicates that the model is biased towards classifying an example as a partial order, likely learning patterns present in training dataset but not generalizable – when using data outside of dataset one can anticipate many false positives. MCML’s ability to quantify w.r.t. the entire state space is of unique value as it quantifies model’s generalizability, and avoids a false sense of confidence in traditional ML metrics.

Varying class ratios (ratio of valid samples to invalid samples in training dataset). A key utility of MCML is that it allows quantifying model generalizability even if the training distribution is different from the true distribution. Table 9 illustrates this for the antisymmetric property. The traditional metric reports precision to be ≥ 0.97 for all ratios whereas MCML metrics report that the precision is as low as 0.19 (for 99:1 ratio) and is ≥ 0.97 only when the models are trained on datasets having a class ratio of 1:99, which is very close to the true distribution ratio. Thus, for this property, traditional metrics fail to capture the performance for almost all class ratios; in contrast, MCML metrics allow precise quantification of true performance of the trained model for each class ratio.

The primary limitation of MCML is its requirement for ϕ , which characterizes the ground truth and is not required by traditional metrics (which are not able to utilize it). We expect property synthesis methods can help MCML apply more generally; moreover, even when the exact ϕ is not available, an approximation of ϕ that is human understandable can be of much use.

5.2.2 Use of Alloy/SAT. We used the Alloy toolset for writing the properties and generating the datasets for this study. The use of Alloy introduces a potential for bias in the results due to the specific techniques employed by Alloy. Our study mitigates this potential threat to validity as follows. There are two sources of potential bias: 1) translation from Alloy to propositional logic; and 2) solving the propositional formula. When Alloy’s symmetry breaking is turned off, the propositional formula Alloy creates faithfully represents the original property (with respect to the scope). Specifically, the set of solutions at the Alloy level is the same (modulo data representation) as the set of solutions at the propositional level with respect to the *primary* variables. For positive data samples, recall, we use every solution enumerated by the SAT backend. Any SAT solver that enumerates every solution will create the same set of solutions. Indeed, different solvers may create that set in different orders. However, our experiments do not depend on the order in which the solutions are created. For example, when we use 10% of valid solutions for training, we do not select the first 10% of the solutions created by the solver, rather we select a random subset with the desired size from all the solutions. For negative data samples, recall, we do not use any constraint solving at all, and instead select a set of negative solutions at random (as described earlier in this section). Hence, when symmetry breaking is turned off, our results do not suffer from bias from the use of Alloy.

When symmetry breaking is turned on, Alloy adds symmetry breaking predicates to the propositional formula before it is solved by the backend SAT engine. Alloy’s motivation of adding these predicates, which preserve the satisfiability of the original formula, is to enable faster solving (by helping SAT prune more effectively, especially for formulas that are

unsatisfiable). A consequence of adding them for our study is that they remove *valid* solutions and can substantially reduce the number of solutions, although the actual reduction depends on the specific formula and scope. To illustrate, consider the common example of enumerating binary trees [12]. For trees with 3 nodes, Alloy removes all symmetries and gives a 6X reduction in the number of solutions; but for 7 nodes, it gives a 1160X reduction whereas full symmetry breaking gives 5040X reduction [69]. Thus, under symmetry breaking turned on, our results are specific to the symmetry breaking predicates added by Alloy’s default setting for each property and scope. A different setting in Alloy or a different tool may lead to different results. We plan to more deeply study the impact of symmetry breaking on learnability in future work.

Note also that MCML uses a very efficient encoding of decision trees directly into CNF that does not use any *auxiliary* variables and is linear in the size of the tree. Thus, the translation is not only faithful but also allows the resulting formulas to be solved readily by off-the-shelf model counters.

6 Conclusion

This paper introduced the MCML approach for empirically studying the learnability of relational properties that can be expressed in the software design language Alloy. A key novelty of MCML is quantification of the performance of and semantic differences among trained machine learning (ML) models, specifically decision trees, with respect to entire input spaces (up to a bound on the input size), and not just for given training and test datasets (as is the common practice). MCML reduces the quantification problems to the classic complexity theory problem of model counting. The results show that relatively simple ML models can achieve surprisingly high performance (accuracy and F1-score) when evaluated in the common setting of using training and test datasets, indicating the seeming simplicity of learning these properties. However, the use of MCML metrics shows that the performance can degrade substantially when tested against the entire (bounded) input space, indicating the high complexity of precisely learning these properties, and the usefulness of model counting in quantifying the true performance.

MCML offers exciting new directions for leveraging model counting in machine learning, e.g., 1) it provides quantitative answers to key questions like "did we train enough?", "how much did we overfit?", and "is this model basically the same as this other model (I have in mind)?"; and 2) it allows informed decision making for ML-based systems, e.g., "should a deployed model be replaced with another (newer) model?".

Acknowledgments

We thank Darko Marinov and the reviewers for very helpful comments and feedback. This work was partially supported by the National Science Foundation grant CCF-1718903.

References

- [1] Hamid Bagheri, Eunsuk Kang, Sam Malek, and Daniel Jackson. 2018. A formal approach for detection of security flaws in the Android permission system. *Formal Asp. Comput.* (2018).
- [2] Teodora Baluta, Shiqi Shen, Shweta Shinde, Kuldeep S Meel, and Prateek Saxena. 2019. Quantitative Verification of Neural Networks And its Security Applications. *arXiv preprint arXiv:1906.10395* (2019).
- [3] Osbert Bastani, Yewen Pu, and Armando Solar-Lezama. 2018. Verifiable reinforcement learning via policy extraction. In *ASIACCS*.
- [4] Anselm Blumer, A. Ehrenfeucht, David Haussler, and Manfred K. Warmuth. 1989. Learnability and the Vapnik-Chervonenkis Dimension. *JACM* 36, 4 (1989).
- [5] Chandrasekhar Boyapati, Sarfraz Khurshid, and Darko Marinov. 2002. Korat: Automated Testing Based on Java Predicates. In *ISSTA*.
- [6] Y. Brun and M. D. Ernst. 2004. Finding latent code errors via machine learning over program executions. In *ICSE*.
- [7] Can Deep Networks Learn Invariants [n. d.]. Can Deep Networks Learn Invariants. ([n. d.]). <https://blog.singularitynet.io/can-deep-networks-learn-invariants-1e06a5052555>.
- [8] Hei Chan and Adnan Darwiche. 2003. Reasoning About Bayesian Network Classifiers. In *UAI*.
- [9] Mark Chavira and Adnan Darwiche. 2008. On probabilistic inference by weighted model counting. *JAI* 172, 6-7 (2008).
- [10] CheckMate GitHub. 2019. <https://github.com/ctrippel/checkmate>. (2019).
- [11] Nathan Chong, Tyler Sorensen, and John Wickerson. 2018. The Semantics of Transactions and Weak Memory in x86, Power, ARM, and C++. In *PLDI*.
- [12] Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. 2009. *Introduction to Algorithms, Third Edition* (3rd ed.). The MIT Press.
- [13] Leonardo Mendonça de Moura, Soonho Kong, Jeremy Avigad, Floris van Doorn, and Jakob von Raumer. 2015. The Lean Theorem Prover (System Description). In *CADE*.
- [14] Brian Demsky and Martin C. Rinard. 2003. Automatic detection and repair of errors in data structures. In *OOPSLA*.
- [15] Niklas Een and Niklas Sörensson. 2004. An Extensible SAT-solver. In *Theory and Applications of Satisfiability Testing*, Enrico Giunchiglia and Armando Tacchella (Eds.) 502–518.
- [16] Daan Fierens, Guy Van den Broeck, Ingo Thon, Bernd Gutmann, and Luc De Raedt. 2012. Inference in Probabilistic Logic Programs using Weighted CNF's. *CoRR* abs/1202.3719 (2012).
- [17] J. P. Galeotti, N. Rosner, C. G. López Pombo, and M. F. Frias. 2013. TACO: Efficient SAT-Based Bounded Verification Using Symmetry Breaking and Tight Bounds. *TSE* (2013).
- [18] Salvador GarcíAa and Francisco Herrera. 2009. Evolutionary Undersampling for Classification with Imbalanced Datasets: Proposals and Taxonomy. *Evolutionary Computation* 17, 3 (2009).
- [19] Pranav Garg, Daniel Neider, P. Madhusudan, and Dan Roth. 2016. Learning Invariants Using Decision Trees and Implication Counterexamples. In *POPL*.
- [20] Robert Gens and Pedro Domingos. 2013. Learning the Structure of Sum-product Networks. In *ICML*.
- [21] Carla P. Gomes, Ashish Sabharwal, and Bart Selman. 2008. Model Counting. (2008).
- [22] Divya Gopinath, Sarfraz Khurshid, Diptikalyan Saha, and Satish Chandra. 2014. Data-guided repair of selection statements. In *36th International Conference on Software Engineering (ICSE)*. 243–253.
- [23] Divya Gopinath, Muhammad Zubair Malik, and Sarfraz Khurshid. 2011. Specification-Based Program Repair Using SAT. In *TACAS*.
- [24] Divya Gopinath, Kaiyuan Wang, Mengshi Zhang, Corina S. Pasareanu, and Sarfraz Khurshid. 2018. Symbolic Execution for Deep Neural Networks. *CoRR* abs/1807.10439 (2018).
- [25] Johan Håstad. 1987. *Computational Limitations of Small-depth Circuits*. MIT Press, Cambridge, MA, USA.
- [26] Marijn J. H. Heule, Matti Juhani Järvisalo, and Martin Suda. 2018. Proceedings of SAT Competition 2018: Solver and Benchmark Descriptions (*Department of Computer Science Series of Publications B*). <http://hdl.handle.net/10138/237063>
- [27] Xiaowei Huang, Marta Kwiatkowska, Sen Wang, and Min Wu. 2017. Safety Verification of Deep Neural Networks. In *CAV*.
- [28] Itay Hubara, Matthieu Courbariaux, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. 2016. Binarized neural networks. In *NIPS*.
- [29] Ronald L. Iman, Jon C. Helton, and James E. Campbell. 1981. An Approach to Sensitivity Analysis of Computer Models: Part I: Introduction, Input Variable Selection and Preliminary Variable Assessment. *JQT* 13, 3 (1981).
- [30] Daniel Jackson. 2002. Alloy: A Lightweight Object Modeling Notation. *TOSEM* 11, 2 (April 2002).
- [31] Daniel Jackson and Kevin J. Sullivan. 2000. COM revisited: Tool-assisted modelling of an architectural framework. In *SIGSOFT FSE*.
- [32] Daniel Jackson and Mandana Vaziri. 2000. Finding Bugs with a Constraint Solver. In *ISSTA*. Portland.
- [33] P. T. Johnstone. 1979. *Conditions related to de Morgan's law*.
- [34] G. Katz, C. Barrett, D. Dill, K. Julian, and M. Kochenderfer. 2017. Reluplex: An Efficient SMT Solver for Verifying Deep Neural Networks. In *CAV*.
- [35] Yalin Ke, Kathryn T. Stolee, Claire Le Goues, and Yuriy Brun. 2015. Repairing Programs with Semantic Code Search (T). In *ASE*.
- [36] Sarfraz Khurshid and Daniel Jackson. 2000. Exploring the Design of an Intentional Naming Scheme with an Automatic Constraint Analyzer. In *ASE00*. Grenoble, France.
- [37] Moonzoo Kim and Yunho Kim. 2011. Automated Analysis of Industrial Embedded Software. In *Automated Technology for Verification and Analysis*, Tevfik Bultan and Pao-Ann Hsiung (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 51–59.
- [38] B. Korel. 1990. Automated software test data generation. *TSE* 16, 8 (1990).
- [39] Jean-Marie Lagniez and Pierre Marquis. 2019. A Recursive Algorithm for Projected Model Counting. *AAAI* 33 (2019), 1536–1543.
- [40] Yitao Liang, Jessa Bekker, and Guy Van den Broeck. 2017. Learning the Structure of Probabilistic Sentential Decision Diagrams. In *UAI*.
- [41] Fan Long and Martin Rinard. 2016. Automatic Patch Generation by Learning Correct Code. In *43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*. 298–312.
- [42] Darko Marinov and Sarfraz Khurshid. 2001. TestEra: A novel framework for automated testing of Java programs. In *ASE*. 22–31.
- [43] Facundo Molina, Renzo Degiovanni, Pablo Ponzio, German Regis, Nazareno Aguirre, and Marcelo Frias. 2019. Training Binary Classifiers as Data Structure Invariants. In *ICSE*.
- [44] Wenxi Wang Muhammad Usman and Sarfraz Khurshid. 2020. TestMC: A Framework for Testing Model Counters using Differential and Metamorphic Testing. (2020). Under submission.
- [45] Nina Narodytska, Shiva Prasad Kasiviswanathan, Leonid Ryzhyk, Mooly Sagiv, and Toby Walsh. 2018. Verifying Properties of Binarized Deep Neural Networks. In *AAAI*.
- [46] Kexin Pei, Yinzhi Cao, Junfeng Yang, and Suman Jana. 2017. DeepXplore: Automated Whitebox Testing of Deep Learning Systems. In *SOSP*.
- [47] J. R. Quinlan. 1987. Generating Production Rules from Decision Trees. In *IJCAI*.
- [48] J. Rumbaugh, I. Jacobson, and G. Booch. 1998. *The Unified Modeling Language Reference Manual*. Addison-Wesley Object Technology Series.
- [49] Sartaj Sahni and Teofilo Gonzalez. 1976. P-Complete Approximation Problems. *JACM* 23, 3 (July 1976). <https://doi.org/10.1145/321958.321975>
- [50] Hesam Samimi, Ei Darli Aung, and Todd D. Millstein. 2010. Falling Back on Executable Specifications. In *ECOOP*.

[51] Neeraj Sangal, Ev Jordan, Vineet Sinha, and Daniel Jackson. 2005. Using Dependency Models to Manage Complex Software Architecture. *SIGPLAN Not.* 40, 10 (Oct. 2005), 167–176. <https://doi.org/10.1145/1103845.1094824>

[52] Shai Shalev-Shwartz, Ohad Shamir, Nathan Srebro, and Karthik Sridharan. 2010. Learnability, Stability and Uniform Convergence. *JMLR* 11 (Dec. 2010). <http://dl.acm.org/citation.cfm?id=1756006.1953019>

[53] Andy Shih, Arthur Choi, and Adnan Darwiche. 2018. Formal Verification of Bayesian Network Classifiers. In *PGM*.

[54] Ilya Shlyakhter. 2001. Generating Effective Symmetry-Breaking Predicates for Search Problems. In *SAT*.

[55] Xujie Si, Hanjun Dai, Mukund Raghothaman, Mayur Naik, and Le Song. 2018. Learning Loop Invariants for Program Verification. In *NIPS*.

[56] Armando Solar-Lezama, Rodric Rabbah, Rastislav Bodík, and Kemal Ebcioğlu. 2005. Programming by sketching for bit-streaming programs. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*. 281–294.

[57] Mate Soos and Kuldeep S. Meel. 2019. BIRD: Engineering an Efficient CNF-XOR SAT Solver and its Applications to Approximate Model Counting. In *AAAI*.

[58] J. M. Spivey. 1992. *The Z Notation: A Reference Manual* (second ed.). Prentice Hall.

[59] Youcheng Sun, Min Wu, Wenjie Ruan, Xiaowei Huang, Marta Kwiatkowska, and Daniel Kroening. 2018. Concolic Testing for Deep Neural Networks. In *ASE*.

[60] Yuchi Tian, Kexin Pei, Suman Jana, and Baishakhi Ray. 2018. DeepTest: Automated testing of deep-neural-network-driven autonomous cars. In *ICSE*.

[61] Caroline Trippel, Daniel Lustig, and Margaret Martonosi. 2018. Check-Mate: Automated Synthesis of Hardware Exploits and Security Litmus Tests. In *MICRO*.

[62] Caroline Trippel, Daniel Lustig, and Margaret Martonosi. 2019. Security Verification via Automatic Hardware-Aware Exploit Synthesis: The CheckMate Approach. *IEEE Micro* 39, 3 (2019).

[63] G. S. Tseytin. 1966. On the complexity of derivation in propositional calculus. (1966). Presented at the Leningrad Seminar on Mathematical Logic.

[64] Muhammad Usman, Wenxi Wang, Kaiyuan Wang, Cagdas Yelen, Nima Dini, and Sarfraz Khurshid. 2019. A Study of Learning Data Structure Invariants Using Off-the-shelf Tools. In *SPIN*.

[65] L. G. Valiant. 1984. A Theory of the Learnable. *CACM* 27, 11 (Nov. 1984). <https://doi.org/10.1145/1968.1972>

[66] Vladimir N. Vapnik. 1995. *The Nature of Statistical Learning Theory*.

[67] V. N. Vapnik and A Ya. Chervonenkis. 1971. On the Uniform Convergence of Relative Frequencies of Events to Their Probabilities. *Theory of Probability and its Applications* (1971).

[68] Marko Vasic, Andrija Petrovic, Kaiyuan Wang, Mladen Nikolic, Rishabh Singh, and Sarfraz Khurshid. 2019. MoET: Interpretable and Verifiable Reinforcement Learning via Mixture of Expert Trees. *CoRR* (2019).

[69] Wenxi Wang, Muhammad Usman, Alyas Almaawi, Kaiyuan Wang, Kuldeep S. Meel, and Sarfraz Khurshid. 2020. A Study of Symmetry Breaking Predicates and Model Counting. In *TACAS*. To appear.

[70] John Wickerson, Mark Batty, Tyler Sorensen, and George A. Constantinides. 2017. Automatically Comparing Memory Consistency Models. In *POPL*.

[71] Razieh Nokhbeh Zaeem and Sarfraz Khurshid. 2010. Contract-Based Data Structure Repair Using Alloy. In *ECOOP*.

[72] Pamela Zave. 2012. Using Lightweight Modeling to Understand Chord. *SIGCOMM CCR* 42, 2 (2012).

[73] P. Zave. 2017. Reasoning About Identifier Spaces: How to Make Chord Correct. *IEEE Transactions on Software Engineering* (2017).