

Learning Branching Heuristics for Propositional Model Counting

Pashootan Vaezipoor,^{1*} Gil Lederman,^{2*} Yuhuai Wu,^{1,3} Chris Maddison,^{1,3}
Roger B. Grosse,^{1,3} Sanjit A. Seshia,² Fahiem Bacchus¹

¹ University of Toronto

² UC Berkeley

³ Vector Institute

{pashootan, ywu, cmaddis, rgrosse, fbacchus}@cs.toronto.edu

{gilled, ssesia}@eecs.berkeley.edu

Abstract

Propositional model counting, or #SAT, is the problem of computing the number of satisfying assignments of a Boolean formula. Many problems from different application areas, including many discrete probabilistic inference problems, can be translated into model counting problems to be solved by #SAT solvers. Exact #SAT solvers, however, are often not scalable to industrial size instances. In this paper, we present `Neuro#`, an approach for learning branching heuristics to improve the performance of exact #SAT solvers on instances from a given family of problems. We experimentally show that our method reduces the step count on similarly distributed held-out instances and generalizes to much larger instances from the same problem family. It is able to achieve these results on a number of different problem families having very different structures. In addition to step count improvements, `Neuro#` can also achieve orders of magnitude wall-clock speedups over the vanilla solver on larger instances in some problem families, despite the runtime overhead of querying the model.

1 Introduction

Propositional model counting is the problem of counting the number of satisfying solutions to a Boolean formula (Gomes, Sabharwal, and Selman 2009). When the Boolean formula is expressed in conjunctive normal form (CNF), this problem is known as the #SAT problem. #SAT is a #P-complete problem, and by Toda’s theorem (Toda 1991) any problem in the polynomial-time hierarchy (PH) can be solved by a polynomial number of calls to a #SAT oracle. This means that effective #SAT solvers, if they could be developed, have the potential to help solve problems whose complexity lies beyond NP, from a range of applications. The tremendous practical successes achieved by encoding problems to SAT and using modern SAT solvers (Marques-Silva 2018) demonstrate the potential of such an approach.

Modern exact #SAT solvers are based on the DPLL algorithm (Davis, Logemann, and Loveland 1962) and have been successfully applied to solve certain problems, e.g., inference in Bayes Nets (Bacchus, Dalmao, and Pitassi 2003b; Li, Poupart, and van Beek 2011; Sang, Beame, and Kautz 2005b;

Domshlak and Hoffmann 2007) and bounded-length probabilistic planning (Domshlak and Hoffmann 2006); however, many applications remain out of reach of current solvers. For example, in problems such as inference in Markov Chains, which have a temporal structure, exact model counters are still generally inferior to earlier methods such as Binary Decision Diagrams (BDDs). In this paper we show that machine learning methods can be used to greatly enhance the performance of exact #SAT solvers.

In particular, we learn problem family specific branching heuristics for the 2012 version of the DPLL-based #SAT solver `SharpSAT` (Thurley 2006) which uses a state-of-the-art search procedure. We cast the problem as a Markov Decision Process (MDP) in which the task is to select the best literal to branch on next. We use a Graph Neural Network (GNN) (Scarselli et al. 2009) to represent the particular component of the residual formula the solver is currently working on. The model is trained end-to-end via Evolution Strategies (ES), with the objective of minimizing the mean number of branching decisions required to solve instances from a *given distribution of problems*. In other words, given a training set of instances drawn from a problem distribution, the aim is to automatically tailor the solver’s branching decisions for better performance on unseen problems of that distribution.

We found that our technique, which we call `Neuro#`, can generalize not only to unseen problem instances of similar size but also to much larger instances than those seen at training time. Furthermore, despite `Neuro#`’s considerable runtime overhead from querying the learnt model, on some problem domains `Neuro#` can achieve *orders-of-magnitude improvements* in the solver’s *wall-clock* runtime. This is quite remarkable in the context of prior related work (Yolcu and Póczos 2019; Selsam and Bjørner 2019; Balcan et al. 2018; Gasse et al. 2019; Khalil et al. 2016; Hansknecht, Joormann, and Stiller 2018; Lederman et al. 2020), where using ML to improve combinatorial solvers has at best yielded modest wall-clock time improvements (less than a factor of two), and positions this line of research as a viable path towards improving the practicality of exact model counters.¹

*Equal Contribution

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹Our code and the extended version of the paper (with the appendix) are available at: github.com/NeuroSharp.

1.1 Related Work

The first successful application of machine learning to propositional satisfiability solvers was the *portfolio-based* SAT solver `SATzilla` (Xu et al. 2008). Equipped with a set of standard SAT solvers, a classifier was trained offline to map a given SAT instance to the solver best suited to solve it.

Recent work has been directed along two paths: *heuristic improvement* (Selsam and Bjørner 2019; Kurin et al. 2019; Lederman et al. 2020; Yolcu and Póczos 2019), and purely ML-based solvers (Selsam et al. 2019; Amizadeh, Matussevych, and Weimer 2019). In the former, a model is trained to replace a heuristic in a standard solver, thus the model is embedded as a module within the solver’s framework and guides the search process. In the latter, the aim is to train a model that acts as a *stand-alone* “neural” solver. These neural solvers are inherently stochastic and often *incomplete*, meaning that they can only provide an estimate of the satisfiability of a given instance. This is often undesirable in applications of SAT solvers where an exact answer is required, e.g., in formal verification. In terms of functionality, our work is analogous to the first group, in that we aim at improving the branching heuristics of a standard solver. More concretely, our work is similar to (Yolcu and Póczos 2019), who used Reinforcement Learning (RL) and GNNs to learn variable selection heuristics for the *local search-based* SAT solver `WalkSAT` (Selman, Kautz, and Cohen 1993). Local search cannot be used to solve exact #SAT, and empirically (Yolcu and Póczos 2019) obtained only modest improvements on much smaller problems. Our method is also related to (Lederman et al. 2020) and (Gasse et al. 2019), where similar techniques were used in solving quantified Boolean formulas and mixed integer programs, respectively.

Recently, Abboud, Ceylan, and Lukasiewicz (2020) trained GNNs as a stand-alone approximate solver for Weighted DNF Model Counting (#DNF). However, approximating #DNF is a much easier problem: it has a fully polynomial randomized approximation scheme (Karp, Luby, and Madras 1989). So the generalization to larger problem instances demonstrated in that paper is not comparable to the scaling on exact #SAT our approach achieves.

2 Background

#SAT A propositional Boolean formula consists of a set of propositional (true/false) variables composed by applying the standard operators “and” (\wedge), “or” (\vee) and “not” (\neg). A *literal* is any variable v or its negation $\neg v$. A *clause* is a disjunction of literals $\bigvee_{i=1}^n l_i$. A clause is a *unit clause* if it contains only one literal. Finally, a Boolean formula is in *Conjunctive Normal Form* (CNF) if it is a conjunction of clauses. We denote the set of literals and clauses of a CNF formula ϕ by $\mathcal{L}(\phi)$ and $\mathcal{C}(\phi)$, respectively. We assume that all formulas are in CNF.

A *truth assignment* for any formula ϕ is a mapping of its variables to $\{0, 1\}$ (**false/true**). Thus there are 2^n different truth assignments when ϕ has n variables. A truth assignment π satisfies a literal ℓ when ℓ is the variable v and $\pi(v) = 1$ or when $\ell = \neg v$ and $\pi(v) = 0$. It satisfies a clause when at least one of its literals is satisfied. A CNF formula ϕ is satisfied

when all of its clauses are satisfied under π in which case we call π a *satisfying assignment* for ϕ .

The #SAT problem for ϕ is to compute the number of satisfying assignments. If ℓ is a unit clause of ϕ then all of ϕ ’s satisfying assignments must make ℓ true. If another clause $c' = \neg \ell \vee \ell'$ is in ϕ , then every satisfying assignment must also make ℓ' true since $\neg \ell \in c'$ must be false. This process of finding all literals whose truth value is forced by unit clauses is called *Unit Propagation* (UP) and is used in all SAT and #SAT solvers. Such solvers traverse the search tree by employing a branching heuristic. This heuristic selects an unforced variable and branches on it by setting it to **true** or **false**. When a literal ℓ is set to **true** the formula ϕ can be reduced by finding all forced literals using UP (this includes ℓ), removing all clauses containing a true literal, and finally removing all false literals from all clauses. The resulting formula is denoted by $\text{UP}(\phi, \ell)$.

Two sets of clauses are called *disjoint* if they share no variables. A component $C \subset \mathcal{C}(\phi)$ is a subset of ϕ ’s clauses that is disjoint from its complement $\mathcal{C}(\phi) - C$. A formula ϕ can be efficiently broken up into a maximal number of disjoint components C_1, \dots, C_k . Although most formulas initially consist of only one component, as variables are set by branching decisions and clauses are removed, the reduced formulas will often break up into multiple components. Components are important for improving the efficiency of #SAT solving as each component can be solved separately and their counts multiplied: $\text{COUNT}(\phi) = \prod_{i=1}^k \text{COUNT}(C_i)$. In contrast, solving the formula as a monolith takes $2^{\Theta(n)}$ where n is the number of variables in the input formula, and so is not efficient for large n .

A formula ϕ or component C_i can be represented by a *literal-clause incidence graph* (LIG). This graph contains a node for every clause and every literal of ϕ (i.e., v and $\neg v$ for every variable v of ϕ). An edge connects a clause node n_c and a literal node n_ℓ if and only if $\ell \in c$. Note that if C_i is a component of ϕ , then C_i ’s LIG will be a disconnected sub-graph of ϕ ’s LIG (Figure 1).

Both exact (Thurley 2006; Sang et al. 2004; Oztok and Darwiche 2015), approximate (Chakraborty et al. 2014; Meel and Akshay 2020), and probabilistically correct (Sharma et al. 2019) model counters have been developed. In this paper, we focus on exact model counting using the `SharpSAT` solver (Thurley 2006). `SharpSAT` and other modern exact #SAT solvers are based on DPLL (Davis, Logemann, and Loveland 1962) augmented with *clause learning* and *component caching* (Bacchus, Dalmao, and Pitassi 2003a, 2009). A simplified version of the algorithm with the clause learning parts omitted is given in Algorithm 1.

The #DPLL`CACHE` algorithm works on one component at a time. If that component’s model count has already been cached it returns the cached value. Otherwise it selects a literal to branch on (line 4) and computes the model count under each value of this literal by calling `COUNTSIDE()`. The sum of these two counts is the model count of the passed component ϕ , and so is stored in the cache (line 7). The `COUNTSIDE` function first unit propagates the input literal. If an empty clause is found, then the current formula ϕ_ℓ is




Figure 1: An example Literal-Clause Incidence Graph (LIG) for formula: $(x_1 \vee \neg x_4) \wedge (\neg x_1 \vee \neg x_2) \wedge (x_3 \vee x_5)$.

Algorithm 1 Component Caching DPLL

```

1: function #DPLLCACHE( $\phi$ )
2:   if INCACHE( $\phi$ ) then
3:     return CACHELOOKUP( $\phi$ )
4:   Pick a literal  $\ell \in \mathcal{L}(\phi)$ 
5:    $\#l = \text{COUNTSIDE}(\phi, \ell)$ 
6:    $\#\neg l = \text{COUNTSIDE}(\phi, \neg \ell)$ 
7:   ADDTOCACHE( $\phi, \#l + \#\neg l$ )
8:   return  $\#l + \#\neg l$ 

9: function COUNTSIDE( $\phi, \ell$ )
10:   $\phi_\ell = \text{UP}(\phi, \ell)$ 
11:  if  $\phi_\ell$  contains an empty clause then
12:    return 0
13:  if  $\phi_\ell$  contains no clauses then
14:     $k = \#$  of unset variables
15:    return  $2^k$ 
16:   $K = \text{FINDCOMPONENTS}(\phi_\ell)$ 
17:  return  $\prod_{\kappa \in K} \#DPLLCACHE(\kappa)$ 

```

unsatisfiable and has zero models. Otherwise, ϕ_ℓ is divided into its sub-components which are independently solved and the product of their model counts is returned. Critical to the performance of the algorithm is the choice of which literal from the current formula ϕ to branch on. This choice affects the efficiency of clause learning and the effectiveness of component generation and caching lookup success. SharpSAT uses the VSADS heuristic (Sang, Beame, and Kautz 2005a) which is a linear combination of a heuristic aimed at making clause learning effective (VSIDS) and a count of the number of times a variable appears in the current formula.

Graph Neural Networks *Graph Neural Networks* (GNNs) are a class of neural networks used for representation learning over graphs (Gori, Monfardini, and Scarselli 2005; Scarselli et al. 2009). Utilizing a neighbourhood aggregation (or message passing) scheme, GNNs map the nodes of the input graph to a vector space. Let $G = (V, E)$ be an undirected graph with node feature vectors $h_v^{(0)}$ for each node $v \in V$. GNNs use the graph structure and the node features to learn an embedding vector h_v for every node. This is done through iterative applications of a neighbourhood aggregation function. In each iteration k , the embedding of a node $h_v^{(k)}$ is updated by aggregating the embeddings of its neighbours from iteration $k - 1$ and passing the result through a nonlin-

ear aggregation function A parameterized by $W^{(k)}$:

$$h_v^{(k)} = A\left(h_v^{(k-1)}, \sum_{u \in \mathcal{N}(v)} h_u^{(k-1)}; W^{(k)}\right), \quad (1)$$

where $\mathcal{N}(v) = \{u | u \in V \wedge (v, u) \in E\}$. After K iterations, $h_v^{(K)}$ is extracted as the final node embedding h_v for node v . Through this scheme, v 's node embedding at step k incorporates the structural information of all its k -hop neighbours.

Evolution Strategies *Evolution Strategies* (ES) are a class of zeroth order black-box optimization algorithms (Beyer and Schwefel 2002; Wierstra et al. 2014). Inspired by natural evolution, a population of parameter vectors (genomes) is perturbed (mutated) at every iteration, giving birth to a new generation. The resulting offspring are then evaluated by a predefined fitness function. Those offspring with higher fitness score will be selected for producing the next generation.

We adopt a version of ES that has shown to achieve great success in the standard RL benchmarks (Salimans et al. 2017): Let $f : \Theta \rightarrow \mathbb{R}$ denote the fitness function for a parameter space Θ , e.g., in an RL environment, f computes the stochastic episodic reward of a policy π_θ . To produce the new generation of parameters of size n , (Salimans et al. 2017) uses an additive Gaussian noise with standard deviation σ to perturb the current generation: $\theta_{t+1}^{(i)} = \theta_t + \sigma \epsilon^{(i)}$, where $\epsilon^{(i)} \sim \mathcal{N}(0, I)$. We then evaluate every new generation with fitness function $f(\theta_{t+1}^{(i)})$ for all $i \in [1, \dots, n]$. The update rule of the parameter is as follows,

$$\begin{aligned} \theta_{t+1} &= \theta_t + \eta \nabla_{\theta} \mathbb{E}_{\theta \sim \mathcal{N}(\theta_t, \sigma^2 I)} [f(\theta)] \\ &\approx \theta_t + \eta \frac{1}{n\sigma} \sum_i^n f(\theta_{t+1}^{(i)}) \epsilon^{(i)}, \end{aligned}$$

where η is the learning rate. The update rule is intuitive: each perturbation $\epsilon^{(i)}$ is weighted by the fitness of the corresponding offspring $\theta_{t+1}^{(i)}$. We follow the rank-normalization and mirror sampling techniques of (Salimans et al. 2017) to scale the reward function and reduce the variance of the gradient, respectively.

3 Method

We formalize the problem of learning the branching heuristic for #DPLLCACHE as a *Markov Decision Process* (MDP). In our setting, the environment is SharpSAT, which is deterministic except for the initial state, where an instance (CNF formula) is chosen randomly from a given distribution. A time step t is equivalent to an invocation of the branching heuristic by the solver (Algorithm 1: line 4). At time step t the agent observes state s_t , consisting of the component ϕ_t that the solver is operating on, and performs an action from the action space $\mathcal{A}_t = \{l | l \in \mathcal{L}(\phi_t)\}$. The objective function is to reduce the number of decisions the solver makes, while solving the counting problem. In detail, the reward function is defined by:

$$R(s) = \begin{cases} 1 & \text{if } s \text{ is a terminal state with "instance solved" status} \\ -r_{penalty} & \text{otherwise} \end{cases}$$

If not finished, episodes are aborted after a predefined max number of steps, without receiving the termination reward.

Training with Evolution Strategies With the objective defined, we observe that for our task, the potential action space as well as the horizon of the episode can be quite large (up to 20,000 and 1,000, respectively). As Vemula, Sun, and Bagnell (2019) show, the exploration complexity of an action-space exploration RL algorithm (e.g. Q-Learning, Policy Gradient) increases with the size of the action space and the problem horizon. On the other hand, a parameter-space exploration algorithm like ES is independent of these two factors. Therefore, we choose to use a version of ES proposed by Salimans et al. (2017) for optimizing our agent.

SharpSAT Components as GNNs As the task for the neural network agent is to pick a literal l from the component ϕ , we opt for a LIG representation of the component (see Section 2) and we use GNNs to compute a literal selection heuristic based on that representation. In detail, given the LIG $G = (V, E)$ of a component ϕ , we denote the set of clause nodes as $C \subset V$, and the set of literal nodes as $L \subset V$, $V = C \cup L$. The initial vector representation is denoted by $h_c^{(0)}$ for each clause $c \in C$ and $h_l^{(0)}$ for each literal $l \in L$, both of which are learnable model parameters. We run the following message passing steps iteratively:

- Literal to Clause (L2C):

$$h_c^{(k+1)} = \mathcal{A}\left(h_c^{(k)}, \sum_{l \in c} [h_l^{(k)}, h_{\bar{l}}^{(k)}]; W_C^{(k)}\right), \quad \forall c \in C,$$

- Clause to Literal (C2L):

$$h_l^{(k+1)} = \mathcal{A}\left(h_l^{(k)}, \sum_{c, l \in c} h_c^{(k)}; W_L^{(k)}\right), \quad \forall l \in L,$$

where \mathcal{A} is a nonlinear aggregation function, parameterized by $W_C^{(k)}$ for clause aggregation and $W_L^{(k)}$ for literal aggregation at the k^{th} iteration. Following Selsam et al. (2019); Lederman et al. (2020), to ensure negation invariance (i.e. that the graph representation is invariant under literal negation), we concatenate the literal representations corresponding to the same variable $h_l^{(k)}, h_{\bar{l}}^{(k)}$ when running L2C message passing. After K iterations, we obtain a d -dimensional vector representation for every literal in the graph. We pass each representation through a policy network, a Multi-Layer Perceptron (MLP), to obtain a score, and we choose the literal with the highest score. Recently, Xu et al. (2019) developed a GNN architecture named *Graph Isomorphism Network* (GIN), and proved that it achieves maximum expressiveness among the class of GNNs. We hence choose GIN for the parameterization of \mathcal{A} . Specifically, $\mathcal{A}(x, y; W) = \text{MLP}((1+\epsilon)x+y; W)$, where ϵ is a hyperparameter.

Sequential Semantics Many problems, such as dynamical systems and bounded model checking, are iterative in nature, with a distinct temporal dimension to them. In the original problem domain, there is often a state that is evolved

through time via repeated applications of a state transition function. A structured CNF encoding of such problems usually maps every state s_t to a set of variables, and adds sets of clauses to represent the dynamical constraints between every transition (s_t, s_{t+1}) . Normally, all temporal information is lost in reduction to CNF. However, with a learning-based approach, the time-step feature from the original problem can be readily incorporated as an additional input to the network, effectively annotating each variable with its time-step. In our experiments, we represented time by appending to each literal embedding a scalar value (representing the normalized time-step t) before passing it through the output MLP. We perform an ablation study to investigate the impact of this additional feature in Section 5.

Engineering Trade-offs and Constraints Directly training on challenging #SAT instances of enormous size is computationally infeasible. We tackle this issue by training `Neuro#` on small instances of a problem (fast rollouts) and relying on generalization to solve the more challenging instances from the same problem domain. Thus the main requirement is the availability of the generative process that lets us sample problems with desired level of difficulty. Access to such generative process is not an unreasonable assumption in industry and research.

Although reducing the number of branching steps is itself interesting, to beat `SharpSAT` in wall-clock time, `Neuro#`'s lead needs to be wide enough to justify the imposed overhead of querying the GNN. Since the time-per-step ratio is relatively constant, for the method to be effective it is desirable that the step count reduction be *superlinear*, meaning it becomes more effective compared to the vanilla heuristic the larger the problem becomes.

4 Data Generation

Our goal was to evaluate our method on more structured and much larger instances than the small random instances typically used in other related works (Yolcu and Póczos 2019; Selsam et al. 2019; Kurin et al. 2019). To that end, we searched SAT and planning benchmarks for problems whose generative processes were publicly available or feasible to implement. To test the versatility of our method, we made sure that these problems cover a diverse set of domains: sudoku, blocked n-queens, cell (*combinatorial*); sha-1 preimage attack (*cryptography*); island, grid_wrlld (*planning*), bv_expr, it_expr (*circuits*). For brevity, we explain two of the problems that we will use in later sections to discuss the behaviour of our trained model and provide a more detailed description of the other datasets and their generation process in Appendix B:

- `cell(R, n, r)`: Elementary (i.e., one-dimensional, binary) Cellular Automata are simple systems of computation where the cells of an n -bit binary state vector are progressed through time by repeated applications of a rule R (seen as a function on the state space). Figure 2a shows the evolution grid of rules 9, 35 and 49 for 20 iterations.

Reversing Elementary Cellular Automata: Given a randomly sampled state T , compute the number of initial




Figure 2: Contrary to SharpSAT, Neuro# branches earlier on variables of the bottom rows. (a) Evolution of a bit-vector through repeated applications of Cellular Automata rules. The result of applying the rule at each iteration is placed under the previous bit-vector, creating a two-dimensional, top-down representation of the system’s evolution; (b) The initial formula simplification on a *single* formula. Yellow indicates the regions of the formula that this process prunes; (c & d) Variable selection ordering by SharpSAT and Neuro# averaged over the entire dataset. Lighter colours show that the corresponding variable is selected earlier on average.

states I that would lead to that terminal state T in r applications of R , i.e., $|\{I : R^r(I) = T\}|$. The entire r -step evolution grid is encoded by mapping each cell to a Boolean variable ($n \times r$ in total). The clauses impose the constraints between cells of consecutive rows as given by R . The variables corresponding to T (last row of the evolution grid) are assigned as unit clauses. This problem was taken from SATCOMP 2018 (Heule, Järvisalo, and Suda 2018).

- `grid_wrlld(s, t)`: This bounded horizon planning problem from (Vazquez-Chanlatte et al. 2018; Vazquez-Chanlatte, Rabe, and Seshia 2019) is based on encoding a grid world with different types of squares (e.g., lava, water, recharge), and a formal specification such as “Do not recharge while wet” or “avoid lava”. We randomly sample a grid world of size s and a starting position I for an agent. We encode to CNF the problem of counting the number of trajectories of length t beginning from I that always avoid lava.

5 Experiments

To evaluate our method, we designed experiments to answer the following questions: **1) I.I.D. Generalization:** Can a model trained on instances from a given distribution generalize to unseen instances of the same distribution? **2) Upward Generalization:** Can a model trained on small instances generalize to larger ones? **3) Wall-Clock Improvement:** Can the model improve the runtime substantially? **4) Interpretation:** Does the sequence of actions taken by the model exhibit any discernible pattern at the problem level? Our baseline

Dataset	# vars	# clauses	Random	SharpSAT	Neuro#
sudoku(9, 25)	182	3k	338	220	195(1.1x)
n-queens(10, 20)	100	1.5k	981	466	261(1.7x)
sha-1(28)	3k	13.5k	2,911	52	24(2.1x)
island(2, 5)	1k	34k	155	86	30(1.8x)
cell(9, 20, 20)	210	1k	957	370	184(2.0x)
cell(35, 128, 110)	6k	25k	867	353	198(1.8x)
cell(49, 128, 110)	6k	25k	843	338	206(1.6x)
grid_wrlld(10, 5)	329	967	220	195	66(3.0x)
bv_expr(5, 4, 8)	90	220	1,316	328	205(1.6x)
it_expr(2, 2)	82	264	772	412	266(1.5x)

Table 1: Neuro# generalizes to unseen i.i.d. test problems often with a large margin compared to SharpSAT.

Dataset	# vars	# clauses	Random	SharpSAT	Neuro#
sudoku(16, 105)	1k	31k	7,654	2,373	2,300 (1.03x)
n-queens(12, 20)	144	2.6k	31,728	12,372	6,272 (1.9x)
sha-1(40)	5k	25k	15k	387	83 (4.6x)
island(2, 8)	1.5k	73.5k	1,335	193	46 (4.1x)
cell(9, 40, 40)	820	4k	39,000	53,349	42,325 (1.2x)
cell(35, 192, 128)	12k	49k	36,186	21,166	1,668 (12.5x)
cell(35, 256, 200)	25k	102k	41,589	26,460	2,625 (10x)
cell(35, 348, 280)	48k	195k	54,113	33,820	2,938 (11.5x)
cell(49, 192, 128)	12k	49k	35,957	24,992	1,829 (13.6x)
cell(49, 256, 200)	25k	102k	47,341	30,817	2,276 (13.5x)
cell(49, 348, 280)	48k	195k	53,779	37,345	2,671 (13.9x)
grid_wrlld(10, 10)	740	2k	22,054	13,661	367 (37x)
grid_wrlld(10, 12)	2k	6k	100k \leq	93,093	1,320 (71x)
grid_wrlld(10, 14)	2k	7k	100k \leq	100k \leq	2,234 (-)
grid_wrlld(12, 14)	2k	8k	100k \leq	100k \leq	2,782 (-)
bv_expr(7, 4, 12)	187	474	35,229	5,865	2,139 (2.7x)
it_expr(2, 4)	162	510	51,375	7,894	2,635 (3x)

Table 2: Neuro# generalizes to much larger problems than what it was trained on, sometimes achieving orders of magnitude improvements over SharpSAT.

in all comparisons is SharpSAT’s heuristic. Also, to make sure that our model’s improvements are not trivially attainable without training we tested a Random policy that simply chooses a literal uniformly at random. We also studied the impact of the trained model on a variety of solver-specific quality metrics (e.g., cache-hit rate, ...), the results of which are in Appendix D.

The `grid_wrlld`, being a problem of an iterative nature (i.e., steps in the planning problem), was a natural candidate for testing our hypothesis regarding the effect of adding the “time” feature of Section 3, so we report the results for `grid_wrlld` with that feature included and later in this section we perform an ablation study on that feature.

Experimental Setup For each dataset, we sampled 1,800 instances for training and 200 for testing. We trained for 1000 ES iterations. At each iteration, we sampled 8 formulas and 48 perturbations ($\sigma = 0.02$). With mirror sampling, we obtained in total $96 = 48 \times 2$ perturbations. For each




Figure 3: `Neuro#` generalizes well to larger problems. Compare the robustness of `Neuro#` vs. `SharpSAT` as the problem sizes increase. Solid and dashed lines correspond to `SharpSAT` and `Neuro#`, respectively. All episodes are capped at 100k steps.

perturbation, we ran the agent on the 8 formulas (in parallel), to a total of $768 = 96 \times 8$ episodes per parameter update. All episodes, unless otherwise mentioned, were capped at 1k steps during training and 100k during testing. The agent received a negative reward of $r_{penalty} = 10^{-4}$ at each step. We used the Adam optimizer (Kingma and Ba 2015) with default hyperparameters, a learning rate of $\eta = 0.01$ and a weight decay of 0.005.

GNN messages were implemented by an MLP with ReLU non-linearity. The size of literal and clause embeddings were 32 and the dimensionality of C2L (*resp.* L2C) messages was $32 \times 32 \times 32$ (*resp.* $64 \times 32 \times 32$). We used $T = 2$ message passing iterations and final literal embeddings were passed through the MLP policy network of dimensions $32 \times 256 \times 64 \times 1$ to get the final score. When using the extra “time” feature, the first dimension of the decision layer was 33 instead of 32. The initial ($T = 0$) embeddings of both literals and clauses were trainable model parameters.

5.1 Results

I.I.D. Generalization Table 1 summarizes the results of the i.i.d. generalization over the problem domains of Section 4. We report the average number of branching steps on the test set. `Neuro#` outperformed the baseline across all datasets. Most notably, on `grid_wrlld`, it reduced the number of branching steps by a factor of 3.0 and on `cell`, by an average factor of 1.8 over the three cellular rules.

Upward Generalization We created instances of larger sizes (up to an order of magnitude more clauses and variables) for each of the datasets in Section 4. We took the models trained from the previous i.i.d. setting and directly




Figure 4: Cactus plots comparing `Neuro#` to `SharpSAT` on `cell` and `grid_wrlld`. Lower and to the right is better: for any point t on the y axis, the plot shows the number of benchmark problems that are individually solvable by the solver, within t steps (top) and seconds (bottom).

evaluated on these larger instances without further training. The evaluation results in Table 2 show that `Neuro#` generalized to the larger instances across all datasets and in almost all of them achieved substantial gains compared to the baseline as we increased the instance sizes. Figure 3 shows this effect for multiple sizes of `cell(49)` and `grid_wrlld` by plotting the percentage of the problems solved within a number of steps. The superlinear gaps get more pronounced once we remove the cap of 10^5 steps, i.e., let the episodes run to completion. In that case, on `grid_wrlld(10, 12)`, `Neuro#` took an average of 1,320 branching decisions, whereas `SharpSAT` took 809,408 (613x improvement).

Wall-Clock Improvement Given the scale of improvements on the upward generalization benchmark, in particular `cell(49)` and `grid_wrlld`, we measured the runtime of `Neuro#` vs. `SharpSAT` on those datasets (Figure 4). On both problems we observe that as `Neuro#` widens the gap in the number of steps, it manages to beat `SharpSAT` in wall-clock. Note that the query overhead could still be greatly reduced in our implementation through GPU utilization, loading the model in the solver’s code in C++ instead of making out-of-process calls to Python, etc.

5.2 Model Interpretation

Formal analysis of the learned policy and its performance improvements is difficult, however we can form some high-level conjectures regarding the behaviour of `Neuro#` by how well it decomposes the problem. Here we look at `cell`. The reason is that this problem has a straightforward encoding that directly relates the CNF representation to an easy-to-visualize evolution grid that coincides with the standard representation of Elementary Cellular Automata. The natural way to decompose this problem is to start from the known state T (bottom row) and to “guess” the preimage, row by row from the bottom up through variable assignment. Different preimages can




Figure 5: Full-sized variable selection heatmap on dataset `cell(35, 348, 280)`. We show the 99th percentile for each row of the heatmap in the last column.

be computed *independently* upwards, and indeed, this is how a human would approach the problem. Heat maps in Figure 2 (c & d) depict the behaviour under `SharpSAT` and `Neuro#` respectively. The heat map aligns with the evolution grid, with the terminal state T at the bottom. The hotter coloured cells indicate that, on average, the corresponding variable is branched on earlier by the policy. The cooler colours show that the variable is often selected later or not at all, meaning that its value is often inferred through UP either initially or after some variable assignments. That is why the bottom row T and adjacent rows are completely dark, because they are simplified by the solver before any branching happens. We show the effect of this early simplification on a single formula per dataset in Figure 2 (b). Notice that in `cell(35&49)` the simplification shatters the problem space into few small components (dark triangles), while in `cell(9)` which is a more challenging problem, it only chips away a small region of the problem space, leaving it as a single component. Regardless of this, as conjectured, we can see a clear trend with `Neuro#` focusing more on branching early on variables of the bottom rows in `cell(9)` and in a less pronounced way in `cell(35&49)`. Moreover, as more clearly seen in the heatmap for the larger problem in Figure 5, `Neuro#` actually branches early according to the pattern of the rule.

5.3 Ablation Study

We tested the degree to which the “time” feature contributed to the upward generalization performance of `grid_wrlld`. We compared three architectures with `SharpSAT` as the baseline: **1. *GNN***: The standard architecture proposed in Section 3, **2. *GNN+Time***: Same as *GNN* but with the variable embeddings augmented with the “time” feature and **3. *Time***: Where no variable embedding is computed and only the “time” feature is fed into the policy network. As depicted in Figure 6, we discovered that “time” is responsible for most of the improvement over `SharpSAT`. This fact is encouraging, because it demonstrates the potential gains that could be achieved by simply utilizing problem-level data, such as “time”, that otherwise would have been lost during the CNF encoding.

5.4 Discussions

We observed varying degrees of success on different problem families. This raises the question of what traits make a problem more amenable for improvement via `Neuro#`. One of the main contributing factors is the model’s ability to observe similar components many times during training. In other




Figure 6: Ablation study on the impact of the “time” feature on upward generalization on `grid_wrlld(10, 12)`.

words, if a problem gets shattered into smaller components either by the initial simplification (e.g., UP) or after a few variable assignments, there is a high chance that the model fits to such distribution of components. If larger problems of the same domain also break down into similar component structures, then `Neuro#` can generalize well on them. This explains why sequential problems like `grid_wrlld` benefit from our method, as they are naturally decomposable into similar iterations and addition of the “time” feature demarcates the boundaries between iterations even more clearer.

Recently (Bliem and Jarvisalo 2019) showed branching according to the centrality scores of variable nodes of the CNF graph leads to significant performance improvements. We obtained their centrality enhanced version of `SharpSAT` and compared it with `Neuro#` trained on specific problem family. We found that, although centrality enhanced `SharpSAT`, `Neuro#` retained its orders of magnitude superiority over it. This indicates that whatever structure `Neuro#` is exploiting from the graph, it is not exclusively centrality. Also we compared the performance of `Neuro#` against the state-of-the-art *probabilistic* model counter GANAK (Sharma et al. 2019). Note that this solver performs the easier task of providing a model count that is only *probably correct within a given error tolerance*. Thus to make the comparison more fair we set the error tolerance of GANAK to a small value of 10^{-3} and observed that its performance was again inferior to `Neuro#`. An interesting future direction would be to investigate if our method could also be used to customize GANAK’s heuristics. Appendix C includes cactus plots for every problem in our dataset along with comparison to both centrality and GANAK.

6 Conclusion

We studied the feasibility of enhancing the branching heuristic in propositional model counting via learning. We used solver’s branching steps as a measure of its performance and trained our model to minimize that measure. We demonstrated experimentally that the resulting model not only is capable of generalizing to the unseen instances from the same problem distribution, but also maintains its lead relative to `SharpSAT` on larger problems. For certain problems, this lead widens to a degree that the trained model achieves wall-clock time improvement over the standard heuristic, in spite of the imposed runtime overhead of querying the model. This is an exciting first step and it positions this line of research as a potential path towards building better model counters and thus broadening their application horizon.

Acknowledgements

Dr. Seshia's contribution to this work was funded by DARPA contract FA8750-20-C-0156 (LOGiCS), NSF grant CNS-1545126 (VeHiCaL), Microsoft and Intel. Gil Lederman's work was partially funded by NSF award CNS-1836601.

References

- Abboud, R.; Ceylan, İ. İ.; and Lukasiewicz, T. 2020. Learning to Reason: Leveraging Neural Networks for Approximate DNF Counting. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference*, 3097–3104. AAAI Press. URL <https://aaai.org/ojs/index.php/AAAI/article/view/5705>.
- Amizadeh, S.; Matushevych, S.; and Weimer, M. 2019. Learning To Solve Circuit-SAT: An Unsupervised Differentiable Approach. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net. <https://openreview.net/forum?id=BJxgz2R9t7>.
- Bacchus, F.; Dalmao, S.; and Pitassi, T. 2003a. Algorithms and Complexity Results for #SAT and Bayesian Inference. In *44th Symposium on Foundations of Computer Science (FOCS 2003), 11-14 October 2003, Cambridge, MA, USA, Proceedings*, 340–351. IEEE Computer Society. doi:10.1109/SFCS.2003.1238208. <https://doi.org/10.1109/SFCS.2003.1238208>.
- Bacchus, F.; Dalmao, S.; and Pitassi, T. 2003b. Value Elimination: Bayesian Inference via Backtracking Search. In *UAI '03, Proceedings of the 19th Conference in Uncertainty in Artificial Intelligence, Acapulco, Mexico, August 7-10 2003*. Morgan Kaufmann. https://dslpitt.org/uai/displayArticleDetails.jsp?mmnu=1&smnu=2&article_id=909&proceeding_id=19.
- Bacchus, F.; Dalmao, S.; and Pitassi, T. 2009. Solving #SAT and Bayesian Inference with Backtracking Search. *J. Artif. Intell. Res.* 34. doi:10.1613/jair.2648. <https://doi.org/10.1613/jair.2648>.
- Balcan, M.; Dick, T.; Sandholm, T.; and Vitercik, E. 2018. Learning to Branch. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*. PMLR. <http://proceedings.mlr.press/v80/balcan18a.html>.
- Bayardo, Jr., R. J.; and Pehoushek, J. D. 2000. Counting Models Using Connected Components. In *Proceedings of the Seventeenth National Conference on Artificial Intelligence and Twelfth Conference on Innovative Applications of Artificial Intelligence, July 30 - August 3, 2000, Austin, Texas, USA*, 157–162. AAAI Press / The MIT Press. <http://www.aaai.org/Library/AAAI/2000/aaai00-024.php>.
- Beyer, H.; and Schwefel, H. 2002. Evolution strategies - A Comprehensive Introduction. *Nat. Comput.* 1(1): 3–52. <https://doi.org/10.1023/A:1015059928466>.
- Birnbaum, E.; and Lozinskii, E. L. 1999. The Good Old Davis-Putnam Procedure Helps Counting Models. *J. Artif. Intell. Res.* 10: 457–477. <https://doi.org/10.1613/jair.601>.
- Bliem, B.; and Järvisalo, M. 2019. Centrality Heuristics for Exact Model Counting. In *31st IEEE International Conference on Tools with Artificial Intelligence, ICTAI 2019, Portland, OR, USA, November 4-6, 2019*, 59–63. IEEE. URL <https://doi.org/10.1109/ICTAI.2019.00017>.
- Chakraborty, S.; Fremont, D. J.; Meel, K. S.; Seshia, S. A.; and Vardi, M. Y. 2014. Distribution-Aware Sampling and Weighted Model Counting for SAT. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence, July 27 -31, 2014, Québec City, Québec, Canada*, 1722–1730. AAAI Press.
- Davis, M.; Logemann, G.; and Loveland, D. W. 1962. A machine program for theorem-proving. *Commun. ACM* 5(7): 394–397. doi:10.1145/368273.368557. URL <https://doi.org/10.1145/368273.368557>.
- Domshlak, C.; and Hoffmann, J. 2006. Fast Probabilistic Planning through Weighted Model Counting. In *Proceedings of the Sixteenth International Conference on Automated Planning and Scheduling, ICAPS 2006, Cumbria, UK, June 6-10, 2006*, 243–252. AAAI. URL <http://www.aaai.org/Library/ICAPS/2006/icaps06-025.php>.
- Domshlak, C.; and Hoffmann, J. 2007. Probabilistic Planning via Heuristic Forward Search and Weighted Model Counting. *J. Artif. Intell. Res.* 30. doi:10.1613/jair.2289. <https://doi.org/10.1613/jair.2289>.
- Gasse, M.; Chételat, D.; Ferroni, N.; Charlin, L.; and Lodi, A. 2019. Exact Combinatorial Optimization with Graph Convolutional Neural Networks. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, 8-14 December 2019, Vancouver, BC, Canada*. <http://papers.nips.cc/paper/9690-exact-combinatorial-optimization-with-graph-convolutional-neural-networks>.
- Geffner, T.; and Geffner, H. 2018. Compact Policies for Fully Observable Non-Deterministic Planning as SAT. In *Proceedings of the Twenty-Eighth International Conference on Automated Planning and Scheduling, ICAPS 2018, Delft, The Netherlands, June 24-29, 2018*, 88–96. AAAI Press.
- Gomes, C. P.; Sabharwal, A.; and Selman, B. 2009. Model Counting. In *Handbook of Satisfiability*, volume 185 of *Frontiers in Artificial Intelligence and Applications*, 633–654. IOS Press. <https://doi.org/10.3233/978-1-58603-929-5-633>.
- Gori, M.; Monfardini, G.; and Scarselli, F. 2005. A New Model for Learning in Graph Domains. In *Proceedings. 2005 IEEE International Joint Conference on Neural Networks, 2005.*, volume 2.
- Hansknecht, C.; Joormann, I.; and Stiller, S. 2018. Cuts, Primal Heuristics, and Learning to Branch for the Time-Dependent Traveling Salesman Problem. Technical report, arXiv. <https://arxiv.org/abs/1805.01415>.
- Heule, M. J.; Järvisalo, M. J.; and Suda, M. 2018. Proceedings of SAT Competition 2018: Solver and Benchmark Descriptions <http://hdl.handle.net/10138/237063>.
- Karp, R. M.; Luby, M.; and Madras, N. 1989. Monte-Carlo Approximation Algorithms for Enumeration Problems. *J. Algorithms* 10(3): 429–448. doi:10.1016/0196-6774(89)90038-2. URL [https://doi.org/10.1016/0196-6774\(89\)90038-2](https://doi.org/10.1016/0196-6774(89)90038-2).
- Khalil, E. B.; Bodic, P. L.; Song, L.; Nemhauser, G. L.; and Dilkina, B. 2016. Learning to Branch in Mixed Integer Programming. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA*, 724–731. AAAI Press. <http://www.aaai.org/ocs/index.php/AAAI/AAAI16/paper/view/12514>.
- Kingma, D. P.; and Ba, J. 2015. Adam: A Method for Stochastic Optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*. <http://arxiv.org/abs/1412.6980>.
- Kurin, V.; Godil, S.; Whiteson, S.; and Catanzaro, B. 2019. Improving SAT Solver Heuristics with Graph Networks and Reinforcement Learning. *CoRR* abs/1909.11830. <http://arxiv.org/abs/1909.11830>.

- Lederman, G.; Rabe, M. N.; Seshia, S.; and Lee, E. A. 2020. Learning Heuristics for Quantified Boolean Formulas through Reinforcement Learning. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net. <https://openreview.net/forum?id=BJluxREKDB>.
- Li, W.; Poupart, P.; and van Beek, P. 2011. Exploiting Structure in Weighted Model Counting Approaches to Probabilistic Inference. *J. Artif. Intell. Res.* 40. <http://jair.org/papers/paper3232.html>.
- Marques-Silva, J. 2018. Computing with SAT Oracles: Past, Present and Future. In *Sailing Routes in the World of Computation - 14th Conference on Computability in Europe, CiE 2018, Kiel, Germany, July 30 - August 3, 2018, Proceedings*, volume 10936 of *Lecture Notes in Computer Science*. Springer. https://doi.org/10.1007/978-3-319-94418-0_27.
- Meel, K. S.; and Akshay, S. 2020. Sparse Hashing for Scalable Approximate Model Counting: Theory and Practice. *CoRR* abs/2004.14692. <https://arxiv.org/abs/2004.14692>.
- Oztok, U.; and Darwiche, A. 2015. A Top-Down Compiler for Sentential Decision Diagrams. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015*, 3141–3148. AAAI Press. <http://ijcai.org/Abstract/15/443>.
- Robertson, N.; and Seymour, P. D. 1991. Graph Minors. X. Obstructions to Tree-Decomposition. *J. Comb. Theory, Ser. B* 52(2): 153–190. [https://doi.org/10.1016/0095-8956\(91\)90061-N](https://doi.org/10.1016/0095-8956(91)90061-N).
- Robertson, N.; and Seymour, P. D. 2010. Graph Minors XXIII. Nash-Williams’ Immersion Conjecture. *J. Comb. Theory, Ser. B* 100(2): 181–205. <https://doi.org/10.1016/j.jctb.2009.07.003>.
- Salimans, T.; Ho, J.; Chen, X.; and Sutskever, I. 2017. Evolution Strategies as a Scalable Alternative to Reinforcement Learning. *CoRR* abs/1703.03864. URL <http://arxiv.org/abs/1703.03864>.
- Sang, T.; Bacchus, F.; Beame, P.; Kautz, H. A.; and Pitassi, T. 2004. Combining Component Caching and Clause Learning for Effective Model Counting. In *SAT 2004 - The Seventh International Conference on Theory and Applications of Satisfiability Testing, 10-13 May 2004, Vancouver, BC, Canada, Online Proceedings*. <http://www.satisfiability.org/SAT04/programme/21.pdf>.
- Sang, T.; Beame, P.; and Kautz, H. A. 2005a. Heuristics for Fast Exact Model Counting. In *Theory and Applications of Satisfiability Testing, 8th International Conference, SAT 2005, St. Andrews, UK, June 19-23, 2005, Proceedings*, volume 3569 of *Lecture Notes in Computer Science*, 226–240. Springer. https://doi.org/10.1007/11499107_17.
- Sang, T.; Beame, P.; and Kautz, H. A. 2005b. Performing Bayesian Inference by Weighted Model Counting. In *Proceedings, The Twentieth National Conference on Artificial Intelligence and the Seventeenth Innovative Applications of Artificial Intelligence Conference, July 9-13, 2005, Pittsburgh, Pennsylvania, USA*. AAAI Press / The MIT Press. <http://www.aaai.org/Library/AAAI/2005/aaai05-075.php>.
- Scarselli, F.; Gori, M.; Tsoi, A. C.; Hagenbuchner, M.; and Monfardini, G. 2009. The Graph Neural Network Model. *IEEE Trans. Neural Networks* 20(1): 61–80. <https://doi.org/10.1109/TNN.2008.2005605>.
- Selman, B.; Kautz, H. A.; and Cohen, B. 1993. Local Search Strategies for Satisfiability Testing. In *Cliques, Coloring, and Satisfiability, Proceedings of a DIMACS Workshop, New Brunswick, New Jersey, USA, October 11-13, 1993*, volume 26 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, 521–531. DIMACS/AMS. <https://doi.org/10.1090/dimacs/026/25>.
- Selsam, D.; and Bjørner, N. 2019. NeuroCore: Guiding High-Performance SAT Solvers with Unsat-Core Predictions. *CoRR* abs/1903.04671. <http://arxiv.org/abs/1903.04671>.
- Selsam, D.; Lamm, M.; Bünz, B.; Liang, P.; de Moura, L.; and Dill, D. L. 2019. Learning a SAT Solver from Single-Bit Supervision. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net. https://openreview.net/forum?id=HJMC_iA5tm.
- Sharma, S.; Roy, S.; Soos, M.; and Meel, K. S. 2019. GANAK: A Scalable Probabilistic Exact Model Counter. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, 1169–1176. ijcai.org. URL <https://doi.org/10.24963/ijcai.2019/163>.
- Thurley, M. 2006. SharpSAT - Counting Models with Advanced Component Caching and Implicit BCP. In Biere, A.; and Gomes, C. P., eds., *Theory and Applications of Satisfiability Testing - SAT 2006, 9th International Conference, Seattle, WA, USA, August 12-15, 2006, Proceedings*, volume 4121 of *Lecture Notes in Computer Science*, 424–429. Springer. https://doi.org/10.1007/11814948_38.
- Toda, S. 1991. PP is as Hard as the Polynomial-Time Hierarchy. *SIAM J. Comput.* 20(5): 865–877. <https://doi.org/10.1137/0220053>.
- Vazquez-Chanlatte, M.; Jha, S.; Tiwari, A.; Ho, M. K.; and Seshia, S. A. 2018. Learning Task Specifications from Demonstrations. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, 3-8 December 2018, Montréal, Canada*, 5372–5382. <http://papers.nips.cc/paper/7782-learning-task-specifications-from-demonstrations>.
- Vazquez-Chanlatte, M.; Rabe, M. N.; and Seshia, S. A. 2019. A Model Counter’s Guide to Probabilistic Systems. *CoRR* abs/1903.09354. <http://arxiv.org/abs/1903.09354>.
- Vemula, A.; Sun, W.; and Bagnell, J. A. 2019. Contrasting Exploration in Parameter and Action Space: A Zeroth-Order Optimization Perspective. In *The 22nd International Conference on Artificial Intelligence and Statistics, AISTATS 2019, 16-18 April 2019, Naha, Okinawa, Japan*, volume 89 of *Proceedings of Machine Learning Research*, 2926–2935. PMLR. <http://proceedings.mlr.press/v89/vemula19a.html>.
- Wierstra, D.; Schaul, T.; Glasmachers, T.; Sun, Y.; Peters, J.; and Schmidhuber, J. 2014. Natural Evolution Strategies. *J. Mach. Learn. Res.* 15(1): 949–980. <http://dl.acm.org/citation.cfm?id=2638566>.
- Xu, K.; Hu, W.; Leskovec, J.; and Jegelka, S. 2019. How Powerful are Graph Neural Networks? In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net. <https://openreview.net/forum?id=ryGs6iA5Km>.
- Xu, L.; Hutter, F.; Hoos, H. H.; and Leyton-Brown, K. 2008. SATzilla: Portfolio-based Algorithm Selection for SAT. *J. Artif. Intell. Res.* 32: 565–606. <https://doi.org/10.1613/jair.2490>.
- Yolcu, E.; and Póczos, B. 2019. Learning Local Search Heuristics for Boolean Satisfiability. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, 8-14 December 2019, Vancouver, BC, Canada*, 7990–8001. <http://papers.nips.cc/paper/9012-learning-local-search-heuristics-for-boolean-satisfiability>.