

SMT Solvers and Formulas

SMT Theories

Integers
Reals
Arrays
BitVectors
Strings
...

Φ
First-order
logic formula



SAT
UNSAT
UNKNOWN

Applications


SW & HW Verification
Neural Network Verification
Program Synthesis
Symbolic Execution
Program Analysis
...

SMT Formula Example

$(\sin(x)^3 = \cos(\log(y) \cdot x) \vee b \vee -x^2 \geq 2.3y) \wedge (-b \vee y < -34.4 \vee \exp(y) > \frac{y}{x})$ where $b \in \mathbb{B}, x, y \in \mathbb{R}$

How are Formulas Solved?

Formulas are solved by applying sequence of transformations until a base form (true/false) is reached



To be fast, SMT solvers rely on a set of handcrafted strategies which define how to transform given formula

Our Work

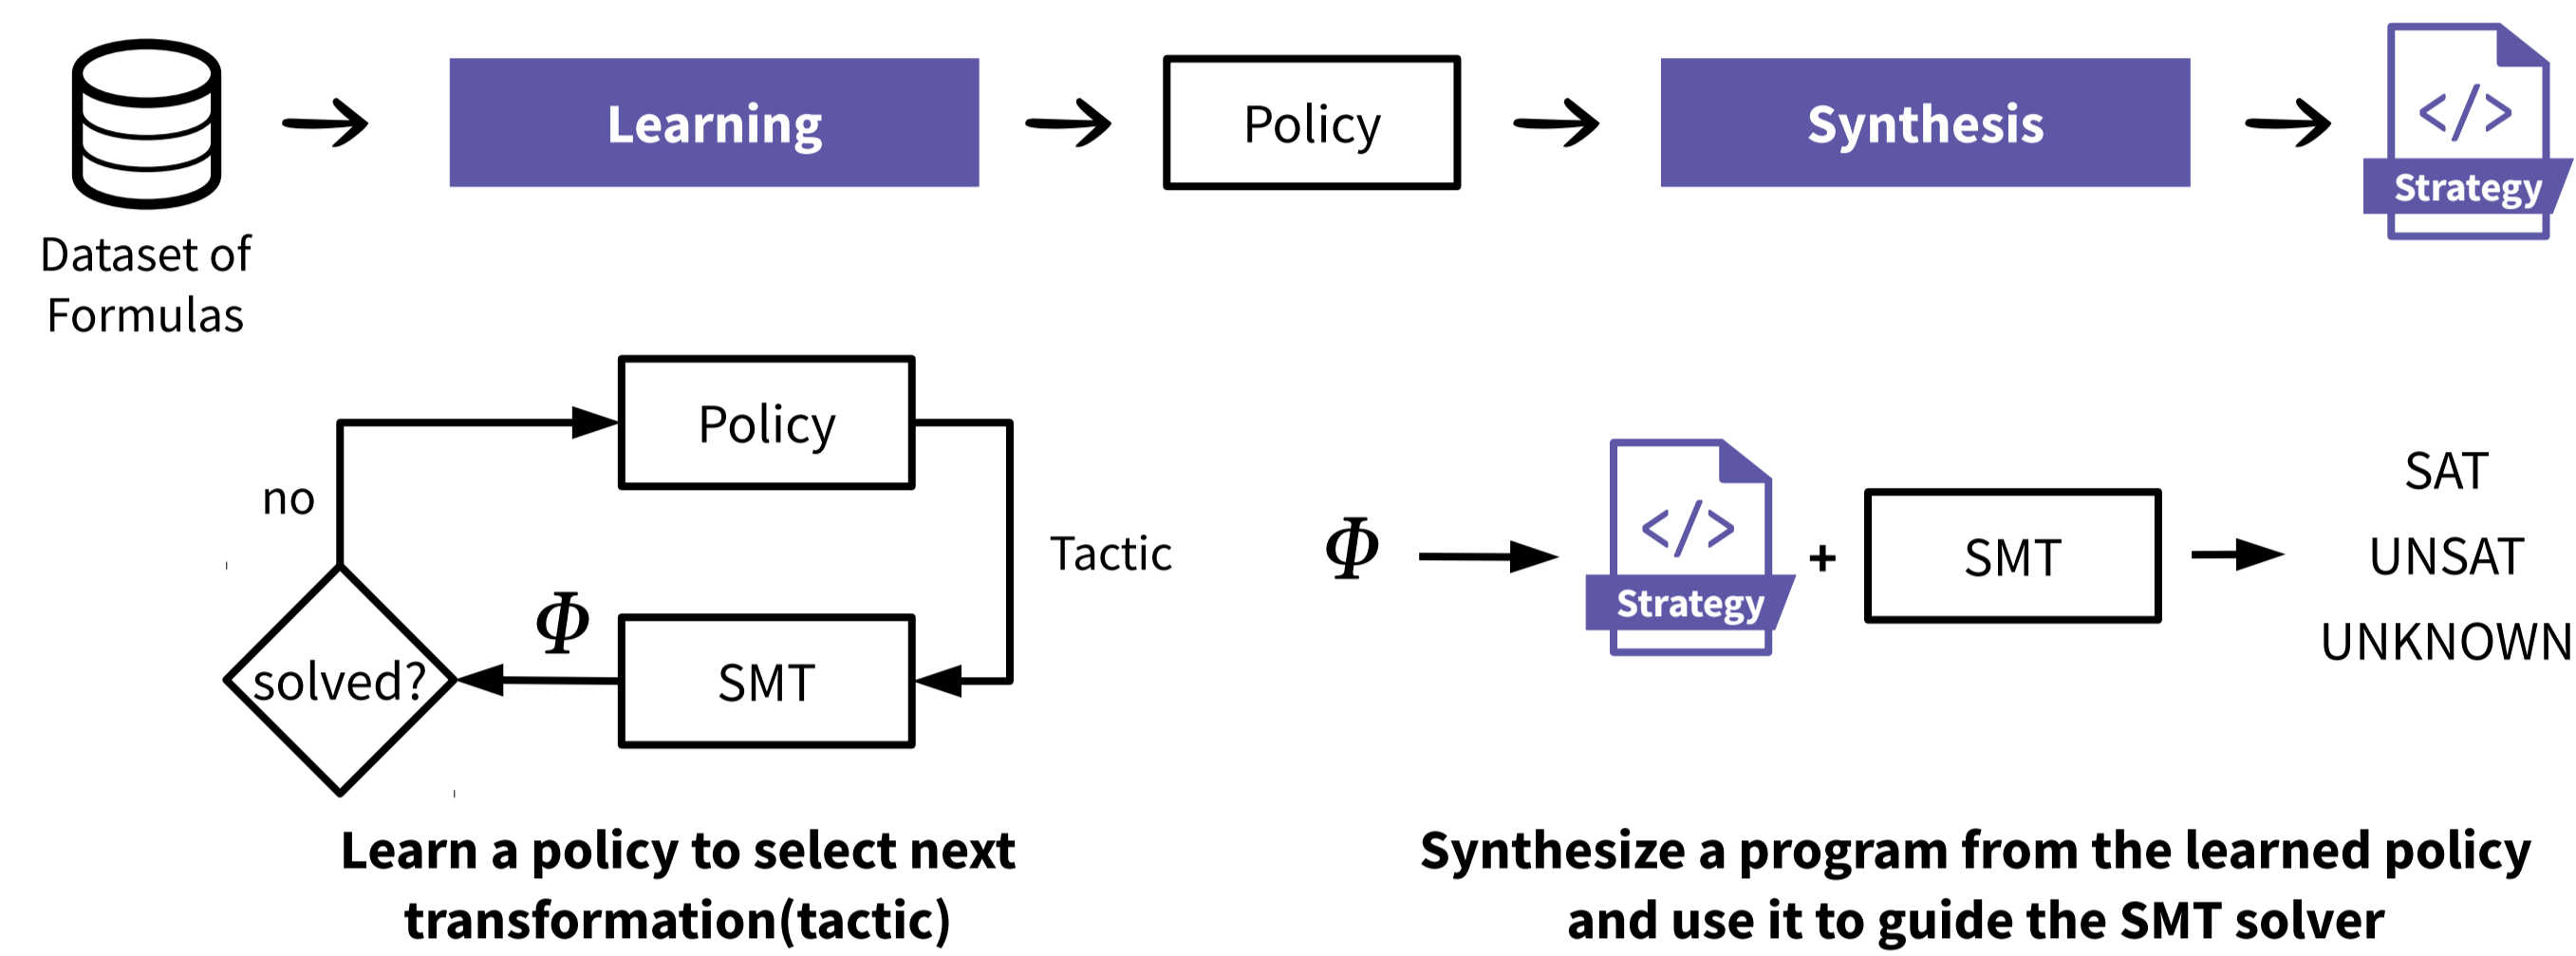
Replace the handcrafted strategies by learning to solve the SMT formulas



such that:

- 1) Learning does not assume any prior knowledge (initially no formula is solvable)
- 2) The learned strategies can be integrated into existing state-of-the-art SMT solvers

Our Approach

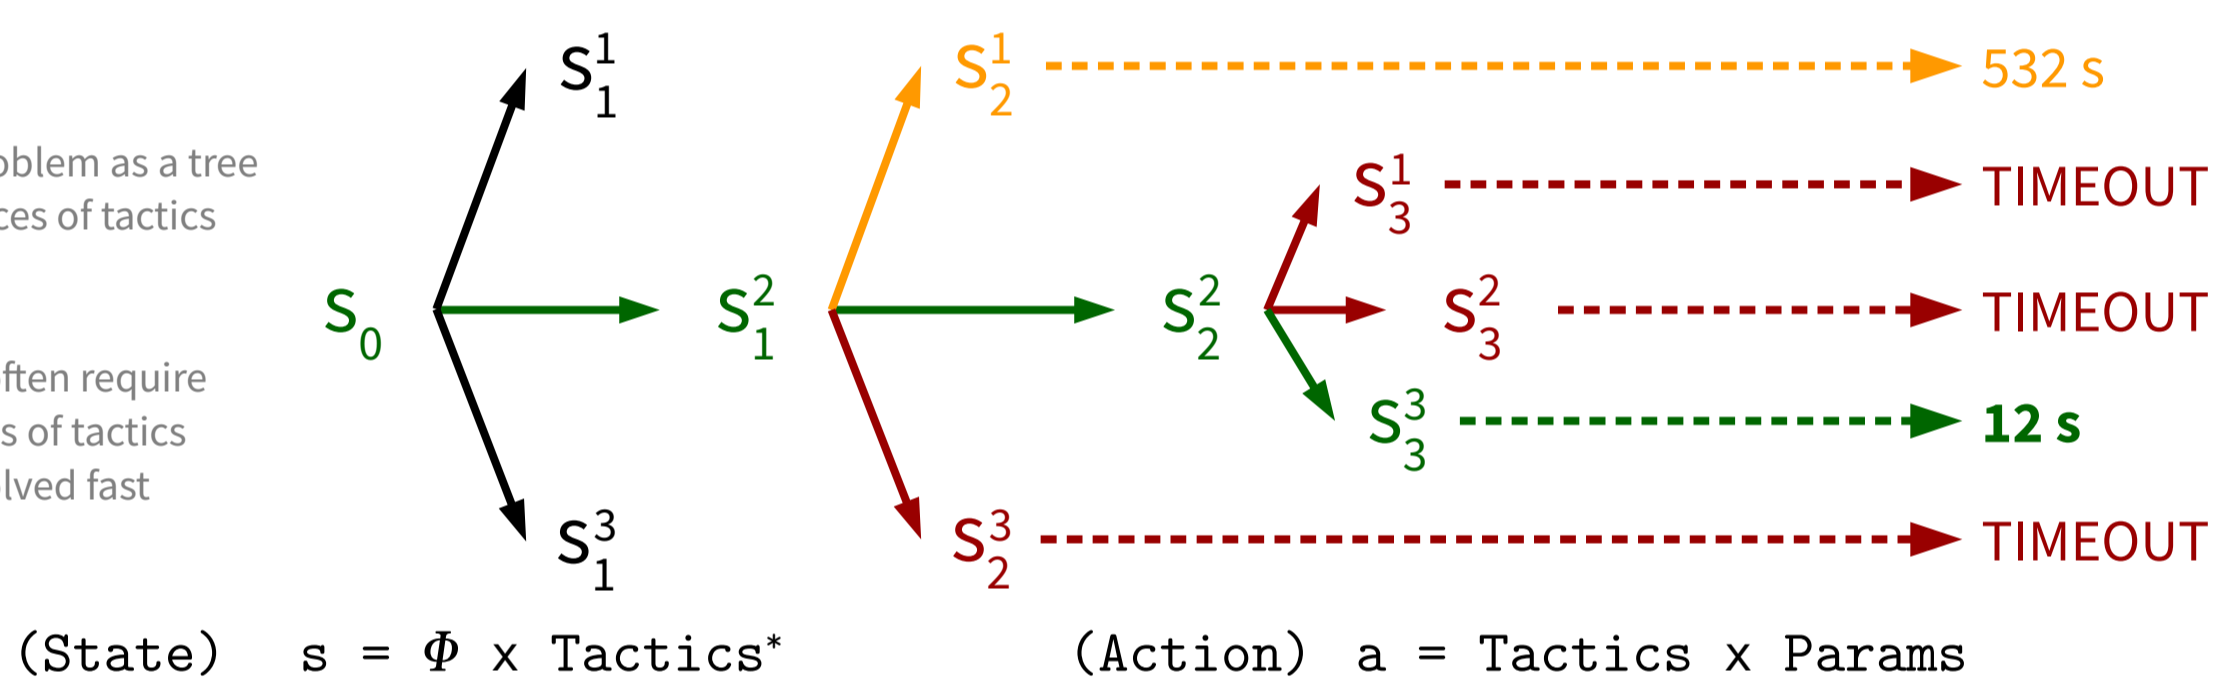


Learning

Given a formula Φ find sequential strategies that are fast at solving Φ

Phrase the learning problem as a tree search over sequences of tactics

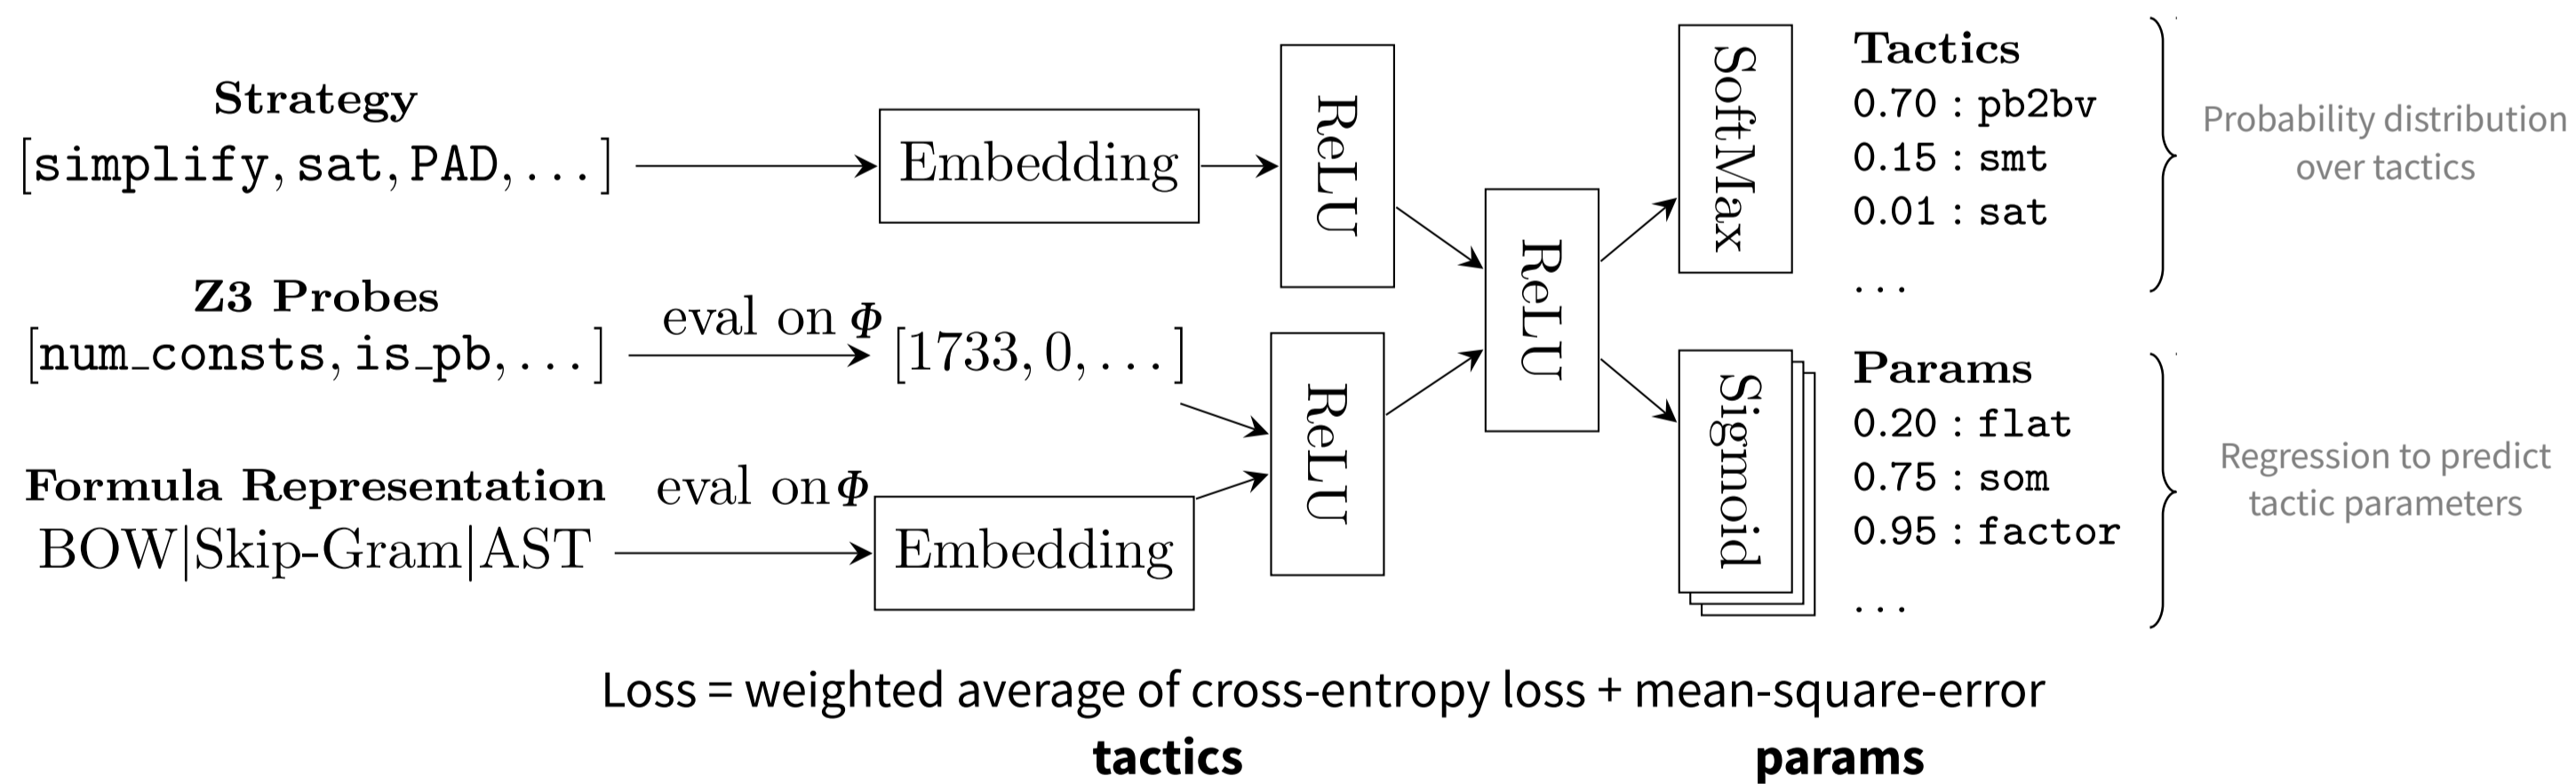
Different formulas often require different sequences of tactics in order to be solved fast



Examples of Strategies

simplify with {arith_lhs:true, som:true}; norm_bounds; lia2pb; pb2bv; bit_blast; sat
simplify with {local_ctx: true}; sat; bit_blast; sat

Learn a policy to select next tactic



Training algorithm based on DAGger [1]

Algorithm 1: Iterative algorithm used to train policy π

Data: Formulas \mathcal{F} , Number of iterations N , Number of formulas to sample K , Exploration rates β , Exploration policy π_{explore} (e.g., random policy)

Result: Trained policy π , Explored strategies

- 1 $\mathcal{D} \leftarrow \emptyset$; $\pi \leftarrow \emptyset$; $\pi \leftarrow$ policy initialization
- 2 for $i = 1$ to N do
- 3 $\hat{\pi} \leftarrow \beta_i \pi + (1 - \beta_i) \pi_{\text{explore}}$ \triangleright policy $\hat{\pi}$ explores with probability $(1 - \beta_i)$
- 4 $Q \leftarrow Q \cup$ Top K most likely strategies for each formula in \mathcal{F} according to $\hat{\pi}$
- 5 $\mathcal{D} \leftarrow \mathcal{D} \cup$ Extract training dataset from strategies
- 6 $\pi \leftarrow$ Retrain model π on \mathcal{D}

[1] S. Ross, G. Gordon, and D. Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. AISTATS'11

Synthesis


Synthesize an interpretable program with branches that selects a strategy that performs best on a dataset of formulas

✓ avoids overhead of running the policy

✓ enables integration with state-of-the-art SMT solvers

Sequential Strategies

Obtained by running the learned policy on a dataset of formulas



Strategy with Branches

Single strategy with synthesized branch for each state with multiple outgoing edges



Decision Tree Learning

Synthesize predicates for each node in the tree in a top-down fashion

Multi-label entropy of a dataset of formulas

$$H(\mathcal{F}) = - \sum_{q \in \mathcal{Q}} p(q) \log(p(q)) + (1 - p(q)) \log(1 - p(q))$$

$p(q)$ denotes ratio of formulas solved by strategy q

Cost associated with branch b

$$\text{cost}(b, \mathcal{F}_{\text{true}}, \mathcal{F}_{\text{false}}) = \frac{|\mathcal{F}_{\text{true}}|}{|\mathcal{F}|} H(\mathcal{F}_{\text{true}}) + \frac{|\mathcal{F}_{\text{false}}|}{|\mathcal{F}|} H(\mathcal{F}_{\text{false}})$$

$H(\mathcal{F}_{\text{true}})$ is entropy of formulas for which branch predicate evaluates to false

Example

	Φ_1	Φ_2	Φ_3	Φ_1	Φ_2	Φ_3	Candidate Branch	Cost
$a_1; a_2; a_3$	10	100	100	200	250	30	if true then a_2 else -	0.276
$a_1; a_4; a_5$	30	TIMEOUT	20	num_expr	if num_expr > 100 then a_2 else a_4			0.2

Syntax of Strategy Language used by Z3 solver

(Strategy) $q ::= t | q; q | \text{if } p \text{ then } q \text{ else } q | q \text{ else } q | \text{repeat } q; c | \text{try } q \text{ for } c | \text{using } t \text{ with params}$

(Predicates) $p ::= p \wedge p | p \vee p | \text{expr} \bowtie \text{expr}$

(Expressions) $\text{expr} ::= c | \text{probe} | \text{expr} \oplus \text{expr}$

(Constants) $c \in \text{Consts} = \mathbb{Q}$

(Probes) $\text{probe} ::= \text{Probe} \rightarrow \mathbb{Q}$, $\text{Probe} = \{ \text{num.consts, is.pb, ...} \}$

(AOperators) $\oplus ::= + | - | * | /$

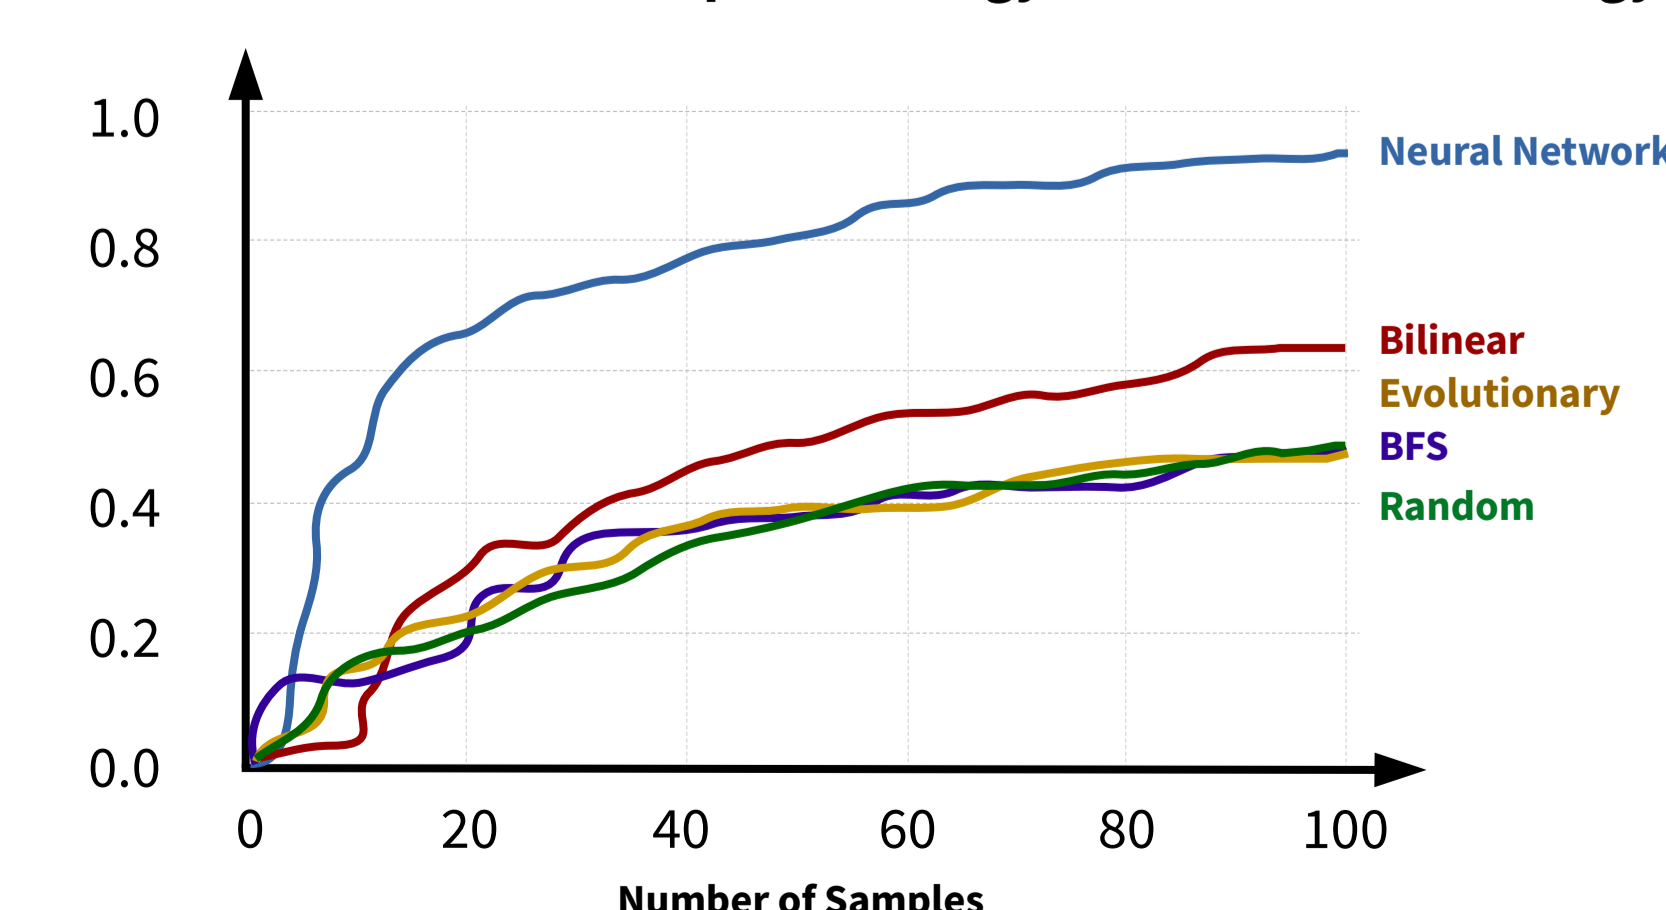
(BOperators) $\bowtie ::= > | < | \geq | \leq | = | \neq$

(Parameter) $\text{param} ::= (\text{Param}, \mathbb{Q})$, $\text{Param} = \{ \text{hoist.mul, flat, som, ...} \}$

(Parameters) $\text{params} ::= \epsilon | \text{param}; \text{params}$

Evaluation: Learning

Runtime Ratio of Best Sampled Strategy vs Best Known Strategy



Find fastest strategy for each formula


2.6x
90th percentile

31x
50th percentile

310x
10th percentile

90th percentile denotes that 90% of the formulas have speedup at-least 2.6x

Speedup over State-of-the-art Z3 Solver



Synthesized Strategy vs Handcrafted Z3 Strategy

	Formulas Solved				Speedup Percentile		
	Both	Only Z3	Only Learned	None	90 th	50 th	10 th
leipzig	55	5	1	7	1.4x	9.9x	21.7x
Sage2	2488	200	705	3051	0.8x	1.2x	3.1x
core	270	0	0	0	1.2x	1.3x	1.8x
hycomp	1547	93	112	230	0.4x	1.1x	2.3x
APROVE	1365	76	112	159	3.2x	65.1x	860.8x
Total	53.6%	3.5%	8.7%	34.1%	1.4x	15.7x	178.0x

Strategy Generalization: 10 seconds Training, 10 seconds Testing → 31x Average Speedup; 10 seconds Training, 10 minutes Testing → 32x Average Speedup