

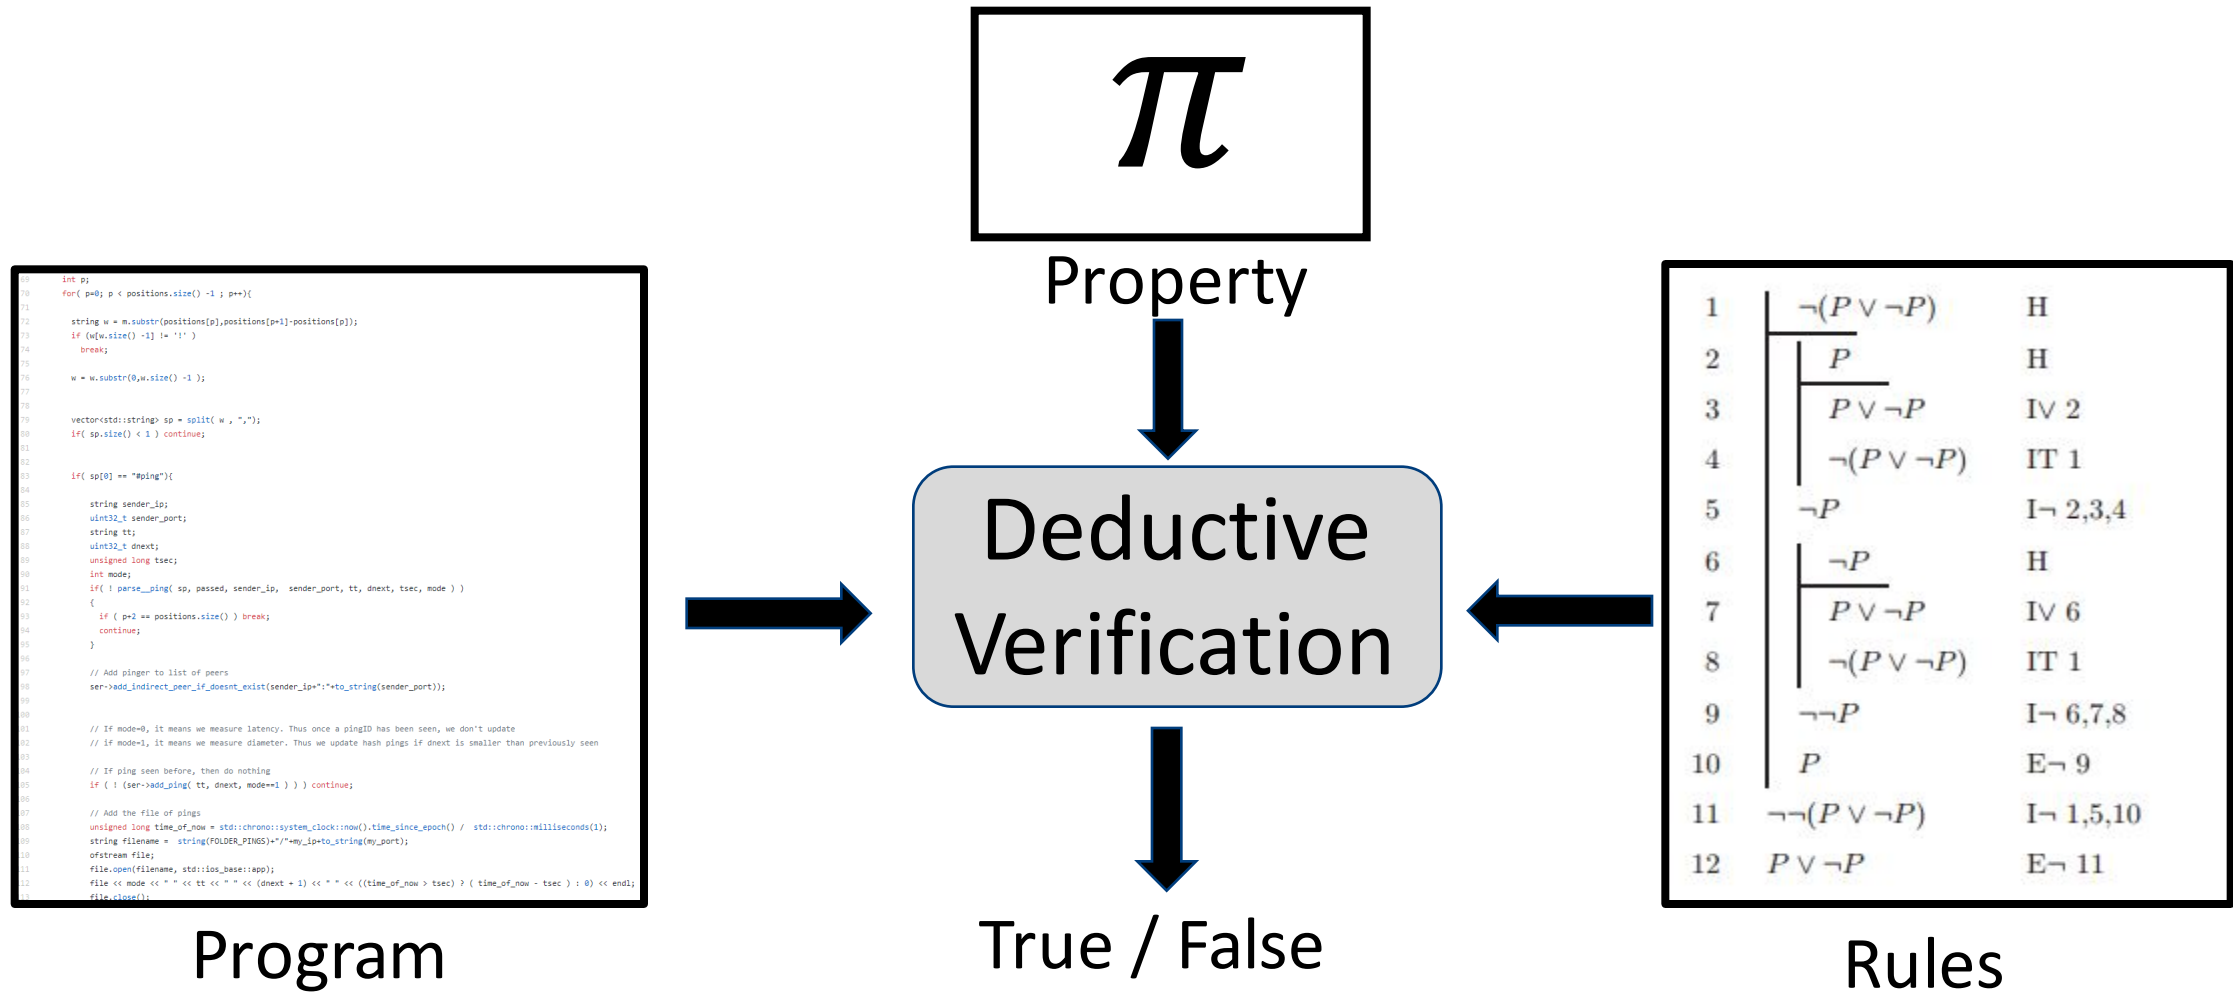
# The Symbiosis of Program Analysis and Machine Learning

Prateek Saxena

Associate Professor

*National University of Singapore*

# Program Analysis, Classically




# But, In Practice...



Program



Property



Rules

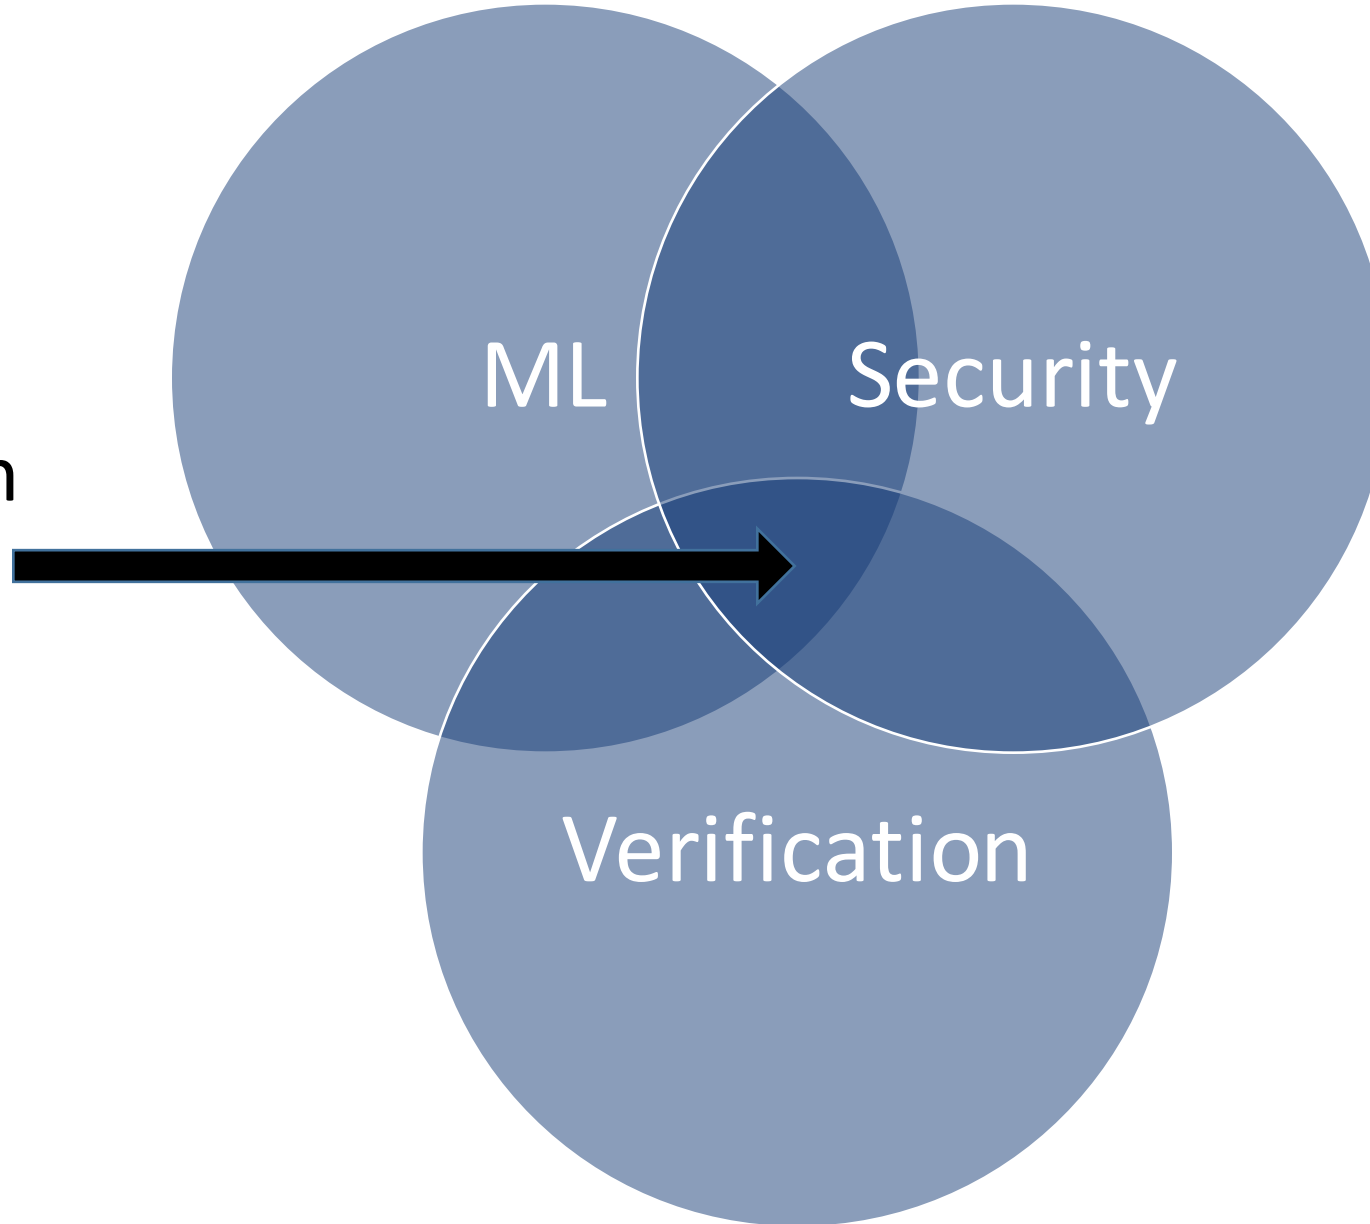
- Too Complex to Analyze / Model
- Probabilistic/ Stochastic Properties
- Not Re-Targetable
- Probabilistic System
- Ambiguous Spec. (Eg. Good patch?)
- Intractable Analysis



Can Machine Learning Help?

# This Talk...

3 Mainstream  
Security  
Analysis  
Tasks



# Machine Learning $\longrightarrow$ Program Analysis

## New Representations & Inference Tools

$\pi$

Property

Induction  
(ML)

```

10 int p;
11 for( p=0; p < positions.size()-1 ; p++){
12     string w = m.substr(positions[p],positions[p+1]-positions[p]);
13     if (w.size()-1 != '!')
14         break;
15
16     w = w.substr(0,w.size()-1);
17
18     vector<std::string> sp = split(w, ",");
19     if( sp.size() < 1 ) continue;
20
21     if( sp[0] == "ping"){
22
23         string sender_ip;
24         uint32_t sender_port;
25         string tt;
26         uint32_t dnext;
27         unsigned long tsec;
28         int mode;
29         if( ! parse_ping( sp, passed, sender_ip, sender_port, tt, dnext, tsec, mode ) )
30         {
31             if ( p+2 == positions.size() ) break;
32             continue;
33         }
34
35         // Add ping to list of peers
36         ser->add_indirect_peer_if_doesnt_exist(sender_ip+":"+to_string(sender_port));
37
38         // If mode=0, it means we measure latency. Thus once a pingID has been seen, we don't update
39         // If mode=1, it means we measure diameter. Thus we update hash pings if dnext is smaller than previously seen
40
41         // If ping seen before, then do nothing
42         if ( ! (ser->add_ping( tt, dnext, mode+1 ) ) ) continue;
43
44         // Add the file of pings
45         unsigned long time_of_now = std::chrono::system_clock::now().time_since_epoch() / std::chrono::milliseconds(1);
46         string filename = string(FOLDER_PINGS)+"-"+my_ip+to_string(my_port);
47         ofstream file;
48         file.open(filename);
49         file << mode << " " << tt << " " << (dnext + 1) << " " << ((time_of_now > tsec) ? (time_of_now - tsec) : 0) << endl;
50         file.close();
51     }
52 }
    
```

Program Representations

1		$\neg(P \vee \neg P)$	H	
2			$P$	H
3			$P \vee \neg P$	IV 2
4		$\neg(P \vee \neg P)$	IT 1	
5		$\neg P$	$I \neg 2,3,4$	
6			$\neg P$	H
7			$P \vee \neg P$	IV 6
8		$\neg(P \vee \neg P)$	IT 1	
9		$\neg \neg P$	$I \neg 6,7,8$	
10		$P$	$E \neg 9$	
11		$\neg \neg(P \vee \neg P)$	$I \neg 1,5,10$	
12		$P \vee \neg P$	$E \neg 11$	

Rules

# Program Analysis $\longrightarrow$ Machine Learning

## Deductive Reasoning



ML System

$\pi$

Property

Deductive Verification



1	$\neg(P \vee \neg P)$	H
2	$P$	H
3	$P \vee \neg P$	IV 2
4	$\neg(P \vee \neg P)$	IT 1
5	$\neg P$	I $\neg$ 2,3,4
6	$\neg P$	H
7	$P \vee \neg P$	IV 6
8	$\neg(P \vee \neg P)$	IT 1
9	$\neg\neg P$	I $\neg$ 6,7,8
10	$P$	E $\neg$ 9
11	$\neg\neg(P \vee \neg P)$	I $\neg$ 1,5,10
12	$P \vee \neg P$	E $\neg$ 11

Rules

ML  $\longrightarrow$  PA:

# New Representations & Inference Tools For Symbolic Execution


Joint work with:

Shiqi Shen , Shweta Shinde, Soundarya Ramesh, Abhik Roychoudhury (NDSS 2019)

# Symbolic Execution

```
1 def f (x, y):  
2   if (x>y)  
3     x = x+y  
4     if (x-y > 0)  
5       assert false  
6   return (x, y)
```

Dynamic Symbolic Execution (DSE):  
A widely used variation of SE



x and y are symbolic variables  
A and B are symbolic values



# Symbolic Execution for Finding Security Bugs



*Pex*



Angr



Kite

SAGE

jCUTE



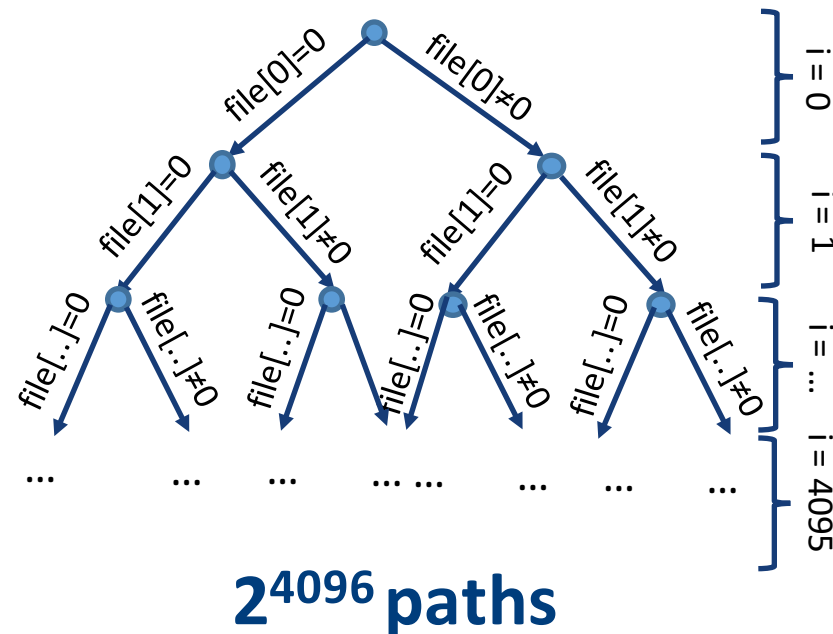
Manticore

TRILON  
Dynamic Binary Analysis

S<sup>2</sup>E

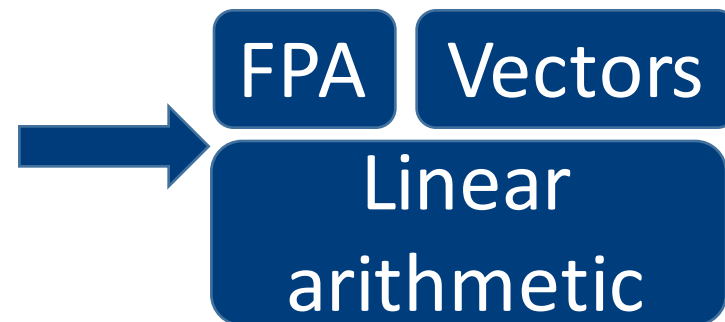
# The Path Explosion Problem

```
1 void copy_data(..., int *file,...) {  
2     static double data[4096], value;  
3     read_double_value(file, ... );  
4     value = fabs (data [0]);  
5     for(i=0; i<4096; i++)  
6         if(file[i] == 0.0) count++;  
7     data[1] /= (value+count-3);  
8     ...  
9 }
```





Prior Approaches: Learn a better representation, symbolically solve!

- Express in SMT theory of floating-point
- Infer that 'count' = # of 0s in input bytes
- Assert:  $value + count - 3 = 0$



# But, why choose this specific constraint representation?

```
1 void copy_data(..., int *file,...) {  
2     static double data[4096], value;  
3     read_double_value(file, ... );  
4     value = fabs (data [0]);  
5     for(i=0; i<4096; i++)  
6         if(file[i] == 0.0) count++;  
7     data[1] /= (value+count-3);  
8     ...  
9 }
```



A Universal Approximate Representation?



# Key Insights

## Desired Representation:

$$\begin{aligned} & \text{count} == \\ & \sum_{i \in [0, 4095]} \text{sign}(\text{file}[i] == 0) \end{aligned}$$

```
1 void copy_data(..., int *file ..) {
2     static double data[4096], value;
3     read_double_value(file, ... );
4     value = fabs (data [0]);
5     for(i=0; i<4096; i++)
6         if(file[i] == 0.0) count++;
7     data[1] /= (value+count-3);
8     ...
9 }
```

A neural network is an approximate representation of the desired...

## Remarks:

- Neural Networks are universal approximators
- Increasing practical success

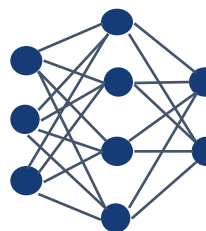
# Key Insights

Values of Symbolic  
Variables



Values of Variables in  
CVP


Learn an approximation  
with small number of  
I/O examples

```
1 int main (...) {  
2   if (strlen filename >1 && filename[0]=='-')  
3     exit(1)  
4   copy_data(...);  
5   ...  
6 }  
7 void copy_data(..., int *file ...) {  
8   Approximate Constraint (as a neural net):  
9  
10  file →  → count & value  
11  
12  
13  data[1] /= (value+count-3); CVP: Divide-by-zero  
14  ...  
15 }
```

# A New Approach: Neuro-symbolic Execution

```
int p;  
for( p=0; p < positions.size() - 1; p++){  
    string w = m.substr(positions[p],positions[p+1]-positions[p]);  
    if (w.size() != 1) break;  
    w = w.substr(0,w.size() - 1);  
    vector<string> sp = split( w, "," );  
    if( sp.size() < 1 ) continue;  
    if( sp[0] == "#ping"){  
        string sender_ip;  
        uint32_t sender_port;  
        string tt;  
        uint32_t dnext;  
        unsigned long tsec;  
        int mode;  
        if( ! parse_ping( sp, passed, sender_ip, sender_port, tt, dnext, tsec, mode ) )  
        {  
            if ( p+2 == positions.size() ) break;  
            continue;  
        }  
        ser->add_indirect_peer_if_doesnt_exist(sender_ip+":"+to_string(sender_port));  
        // If mode=0, it means we measure ping time. If we have a pingID we have seen, we don't update  
        // If mode=1, it means we measure ping time. If we have a pingID we have seen, we don't update  
        // If ping seen before, then do nothing  
        if ( ! (ser->add_ping( tt, dnext, mode=1 ) ) ) continue;  
        // Add the ping time  
        unsigned long time_of_ping = system_clock::time_point::to_time_t() + milliseconds(1);  
        string filename = string("ping_time_") + to_string( time_of_ping );  
        ofstream f( filename, ios::app );  
        f << mode << " " << tt << " " << dnext << " " << ((time_of_ping > tsec) ? (time_of_ping - tsec) : 0) << endl;  
        f.close();  
    }  
}
```

Program



Property

Symbolic (SMT) Constraints

Neural Network

Symbolic Exploit Condition

SATISFIABLE?

# Constraint Solving: Satisfiability Checking

## 1. Reachability constraints:

$$\begin{aligned} & \text{strlen}(\text{filename}) \leq 1 \\ & \forall \text{filename} \neq \text{'-'} \end{aligned}$$

$\wedge$

## 2. Vulnerability condition:

$$\text{value} + \text{count} - 3 == 0$$

**Purely symbolic constraints:**

➔ SMT solver

No variable shared with neural constraints

**Mixed constraints:**

Including both neural constraints and symbolic constraints with shared variables




# Solving SMT + Neural Constraints: Encode SMT constraints as the loss function

$N: infile \rightarrow (value, count) \wedge \boxed{value + count - 3 == 0}$   
Symbolic constraint

**Criterion for crafting the loss function:**  
The minimum point of the loss function satisfies the symbolic constraints.

↓


$$L = abs(value + count - 3)$$





# Optimize using Gradient Descent

Gradient:  $\nabla_{file} L$


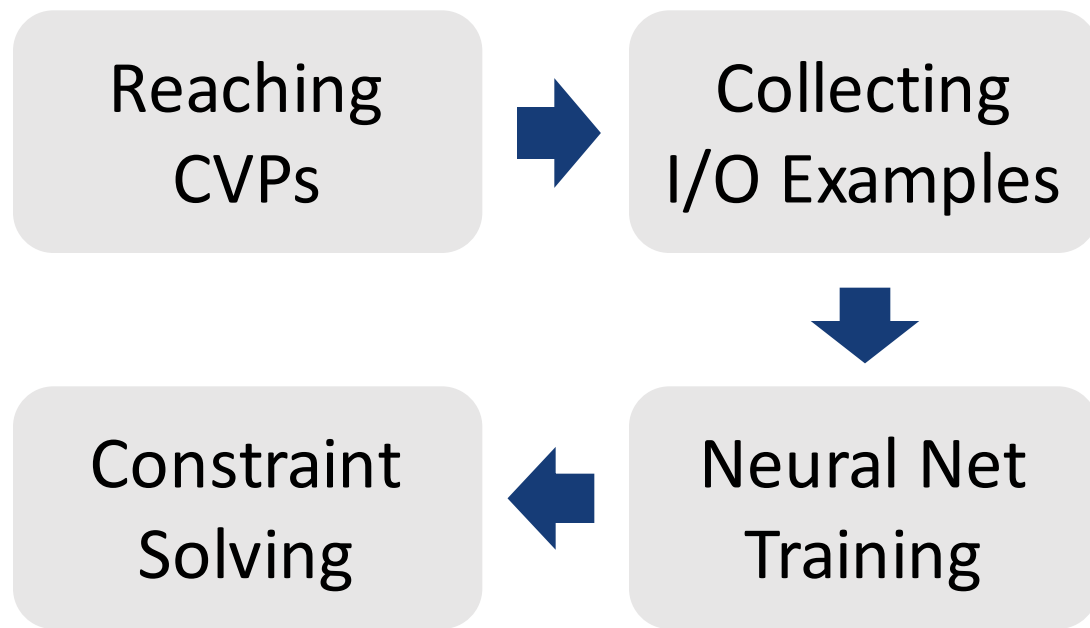


count	value	loss
257	20	274
1	20	18
0	20	17
...	...	...
0	3	0

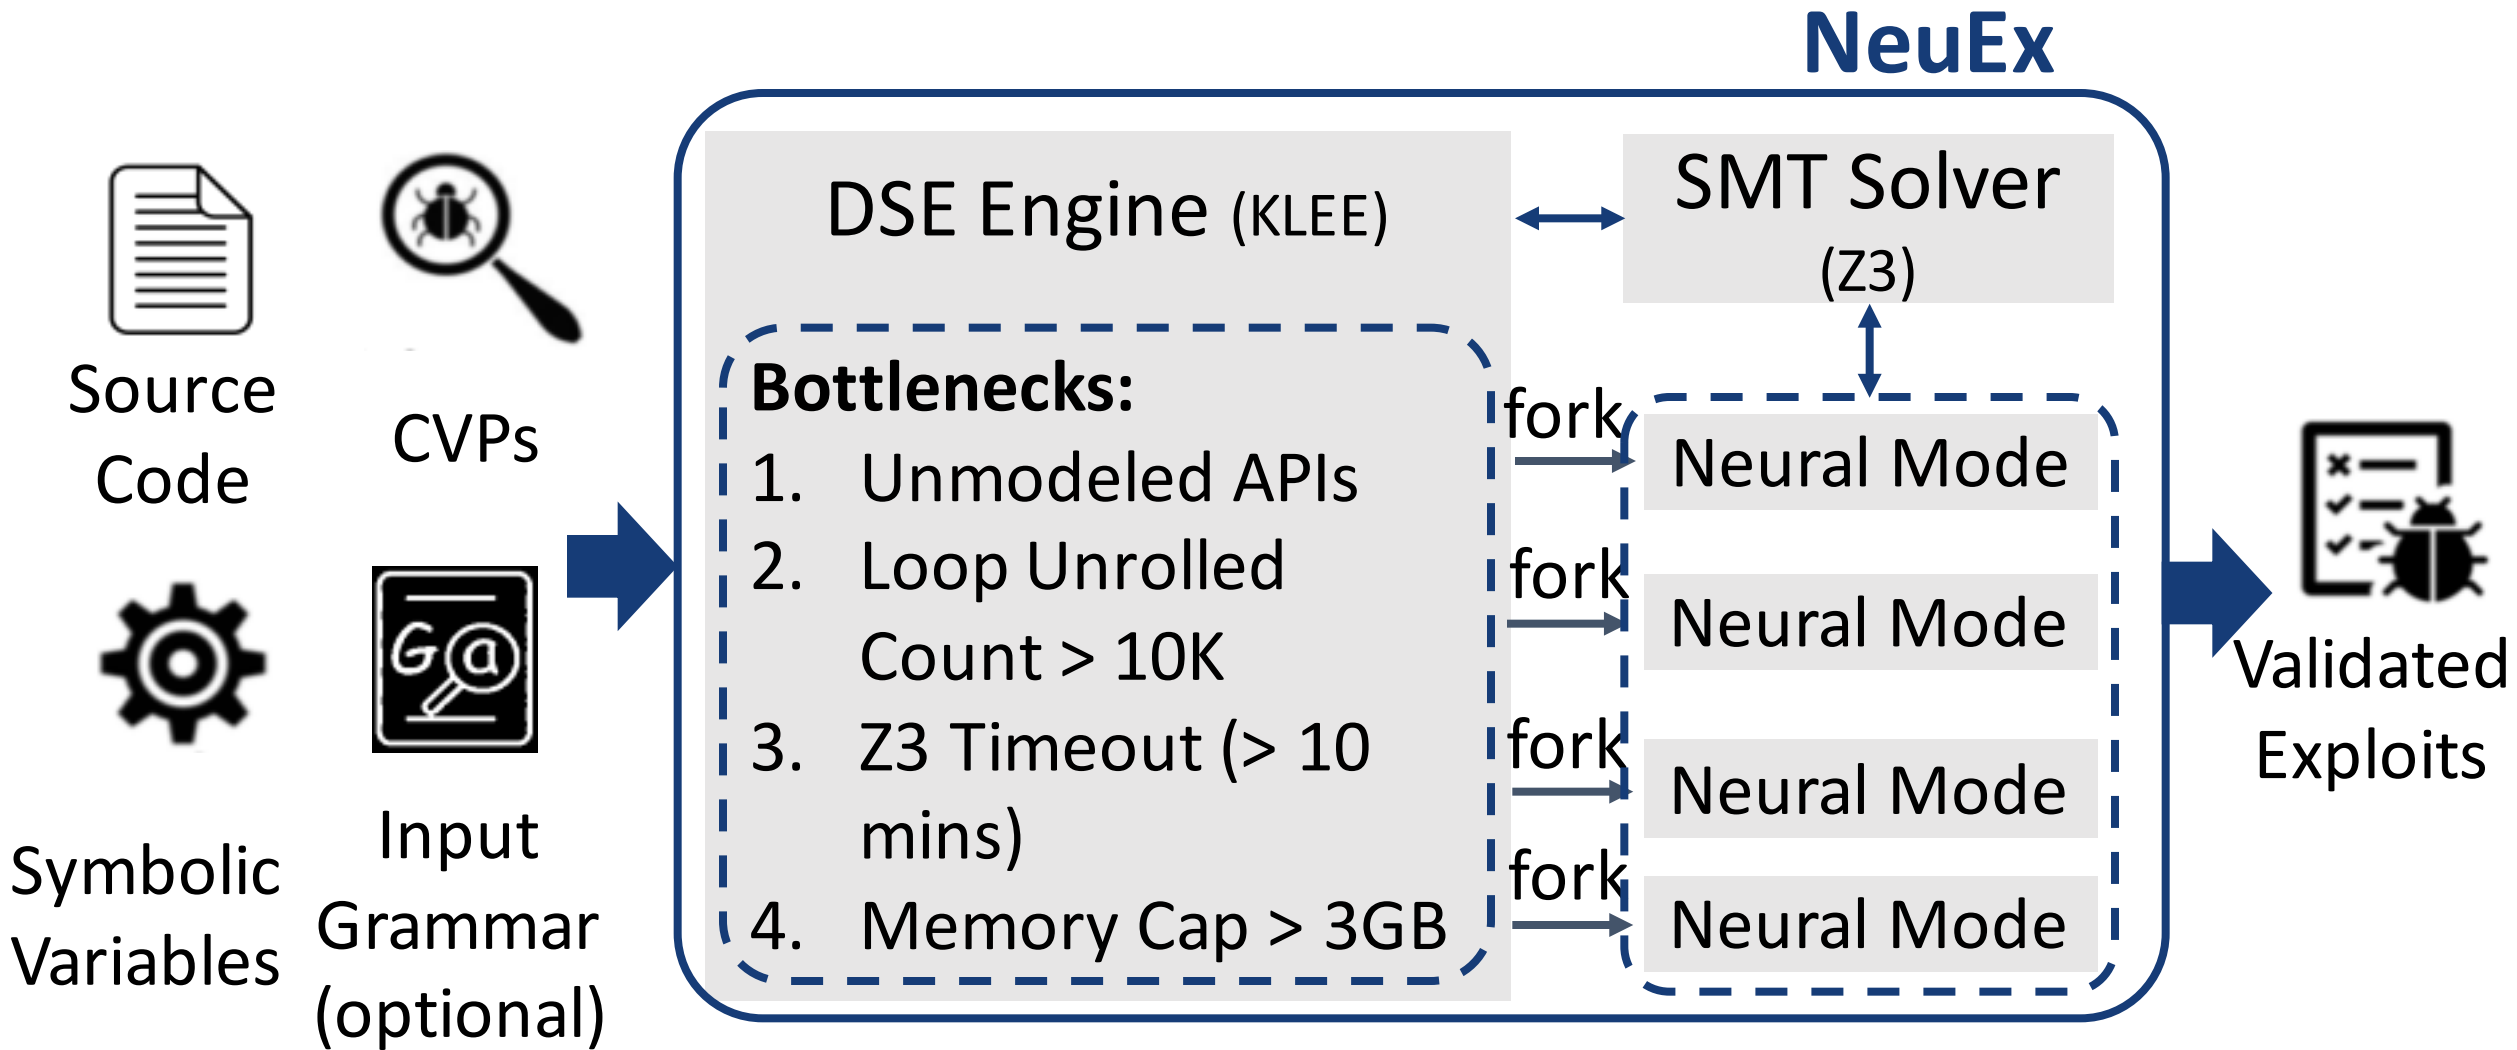
Concretely validate the exploit

← file: 000...

# NeuEX = Neuro-symbolic Execution + KLEE





# NeuEx Tool Overview



# Evaluation

- **Recall:** Neural mode is only triggered when DSE encounters bottlenecks
  - **Benchmarks:** 7 Programs known to be difficult for classic DSE
    - 4 Real programs
      - cURL: Data transferring
      - SQLite: Database
      - libTIFF: Image processing
      - libsndfile: Audio processing
    - LESE benchmarks
      - BIND, Sendmail, and WuFTP
- Include:
1. Complex loops
  2. Floating-point variables
  3. Unmodeled APIs

# NeuEx vs KLEE



#CVPs reached or covered by NeuEx is 25% higher than vanilla KLEE.

KLEE gets stuck (e.g. complex loops)


NeuEx finds 94% and 89% more bugs than vanilla KLEE in BFS and RAND mode in 12 hours.

ML  $\longrightarrow$  PA:  
New Representations & Inference  
For Taint Analysis

Joint work with:

Shiqi Shen , Shweta Shinde, Soundarya Ramesh, Abhik Roychoudhury (NDSS 2019)

# Taint Analysis



- Taint analysis tracks the information flow within a program:
  - E.g. T[v] is the taint bit for operand “v”
- Taint analysis is the basis for many security applications
  - Information leakage detection
  - Enforcing program integrity
  - Vulnerability detection
  - ...

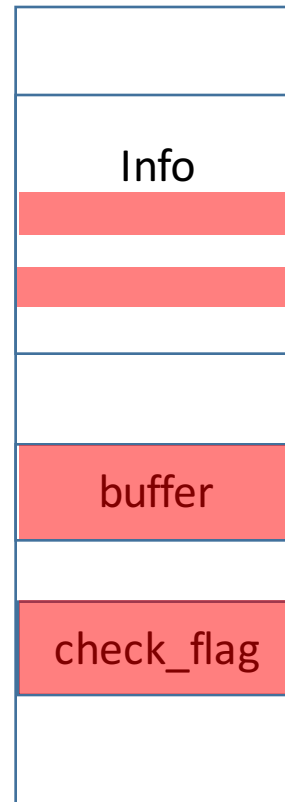
```
1 int parse_buffer(char buffer[100], struct
pkt_info *info) {
2     char check_flag;
3
4     check_flag = buffer[5] & 0x16;
5
6     err = init_pkt_info(info);
7     if (!err)
8         return err;
9     info->flag = check_flag;
10    /* ... */
11    strncpy(info->data, buffer + 6);
12    info->seq = get_current_seq();
13    return OK;
14 }
```

Is Return Address Tainted?

# Taint Analysis on Binaries

```
/* tainted input from network socket */
1 int parse_buffer(char buffer[100], struct
pkt_info *info) {
2     char check_flag;
3
4     check_flag = buffer[5] & 0x16;
5
6     err = init_pkt_info(info);
7     if (!err)
8         return err;
9     info->flag = check_flag;
10    /* ... */
11    strncpy(info->data, buffer + 6, 50);
12    info->seq = get_current_seq();
13    return OK;
14 }
```

Taint Map  
T[]



Write binary taint rules based on instruction semantics

```
movsx    eax, byte ptr [rsi + 5]
and      eax, 16
mov      cl, al
mov      byte ptr [rbp - 25], cl
```

$T[\text{check\_flag}] = T[\text{buffer}+5]$



# Taint Rule Representations in Existing Systems

- What is the taint rule for `and eax, 16` on the x86 architecture?

## Taint Engine 1

$T[\text{eax}] = T[\text{eax}]$




## Taint Engine 2

$T[\text{eax}] = T[\text{eax}]$   
 $T[\text{pf}] = T[\text{sf}] = T[\text{zf}] = T[\text{eax}]$   
 $T[\text{of}] = T[\text{cf}] = 0$



## Taint Engine 3

$T[\text{eax}] = T[\text{eax}]$   
 $T[\text{pf}] = T[\text{sf}] = T[\text{zf}] = T[\text{eax}]$   
 $T[\text{of}] = T[\text{cf}] = 0$   
if `imm == 0` {  $T[\text{eax}] = 0$  }

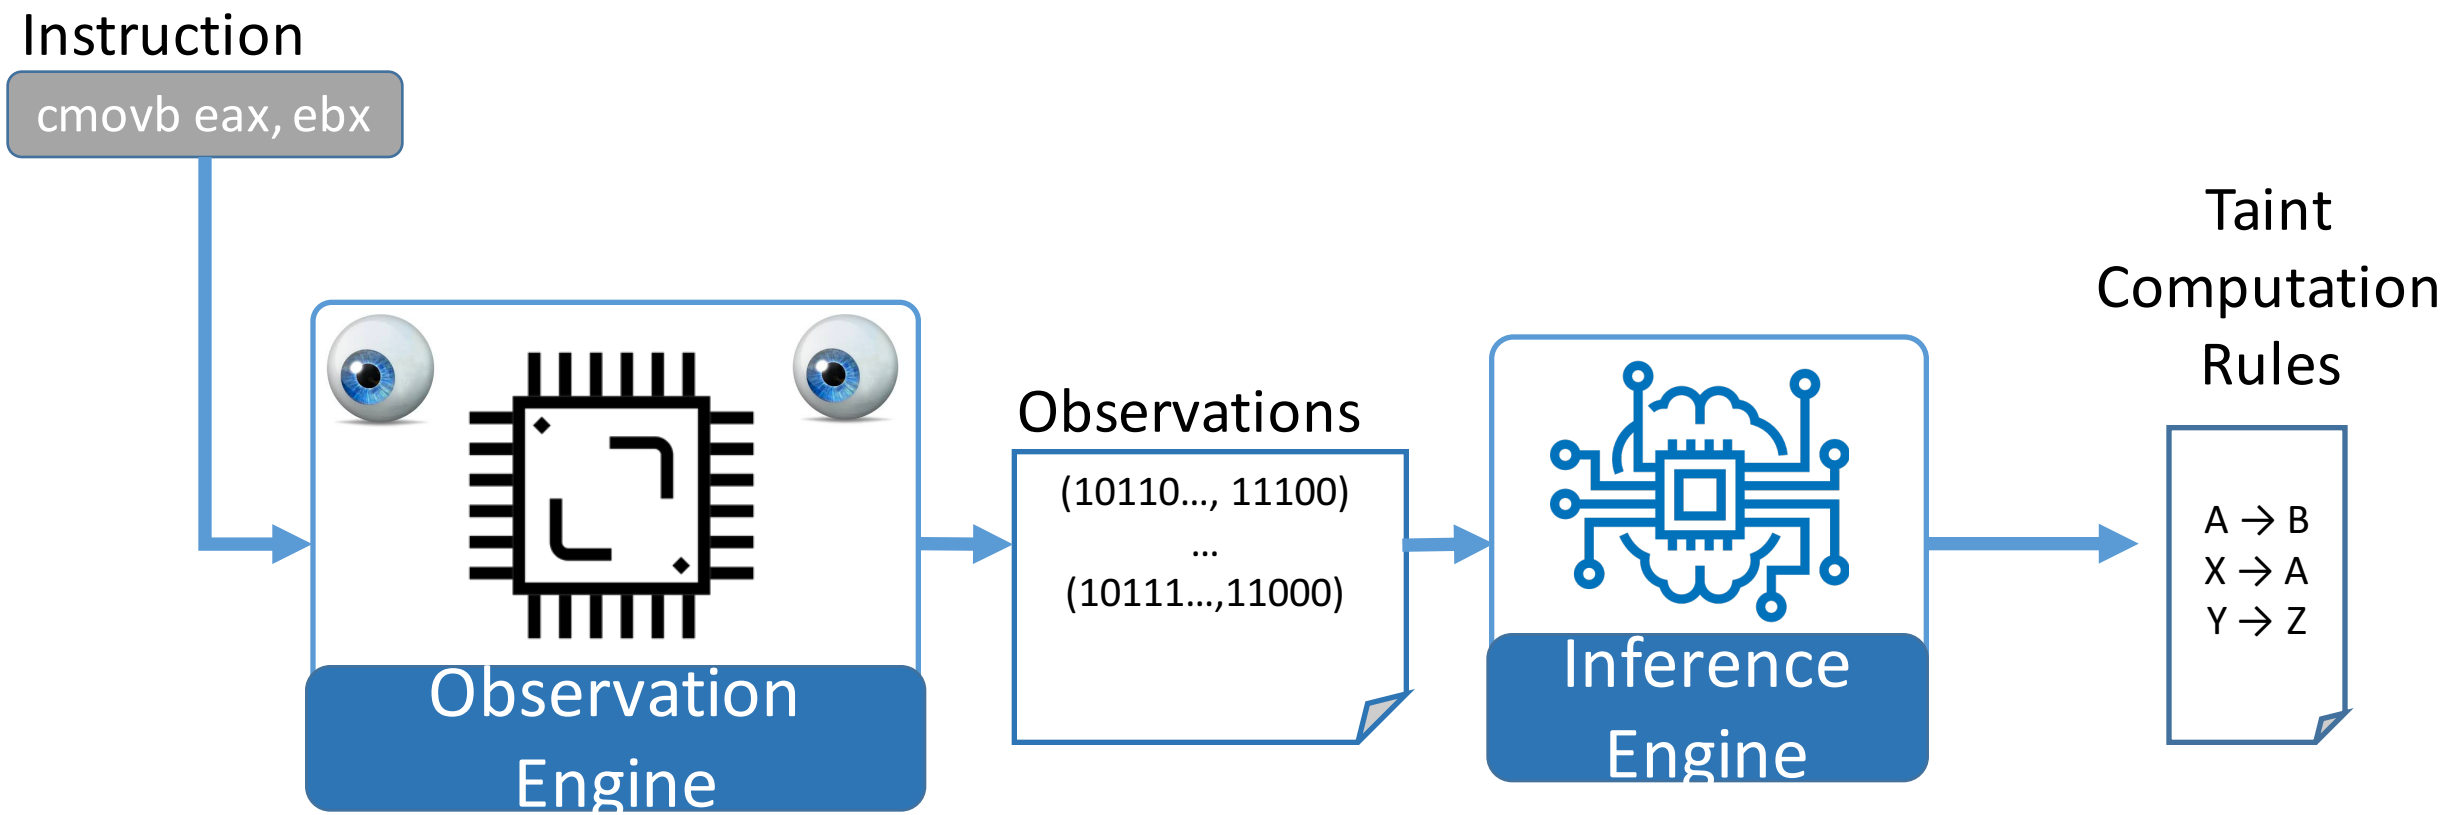


# Complexity of “Real” Taint Rules

- Input dependent propagation
- Size dependent propagation
- Architectural quirks for backwards compatibility
- There can be 1000s of opcodes per instruction set!

```
if (size == 64 || size == 32 || size == 16) {
  for (x = 0; x < size / 8; x++) {
    if (t1[x] & t2[x]) t1[x] = 1;
    else if (t1[x] and !t2[x])
      t1[x] = t1[x] & op2[x];
    else if (!t1[x] & t2[x])
      t1[x] = t2[x] & op1[x];
    else t1[x] = 0;
  } else if (size == 8) {
// 0 if it's lower 8 bits, 1 if it's upper 8 bits
    pos1 = isUpper(op1); pos2 = isUpper(op2);
    if (t1[pos1] & t2[pos2]) t1[pos1] = 1;
    else if (t1[pos1] & !t2[pos2])
      t1[pos1] = t1[pos1] & op2[pos2];
    else if (!t1[pos1] & t2[pos2])
      t1[pos1] = t2[pos2] & op1[pos1];
    else t1[pos1] = 0;}}
if (mode64bit == 1 and size == 64)
  for (x = 32; x < size; x++) t1[x] = 0;
```


# Learning Taint Rules Automatically



# TaintInduce: Sample and Learn


- Flip a bit and observe the output for changes.
  - $\Delta EBX_0 \rightarrow \Delta EAX_0$
  - $\Delta EBX_0 \rightarrow \Delta EBX_0$
- Influence (Inf) only valid if :
  - $EAX = 11100011, EBX = 00101000$
- Form a truth table with all of the collected observations.
  - True if there is a change, False otherwise
- Unseen values are conservatively set to “Don’t-Cares”

mov eax, ebx



EAX <sub>0</sub>	EAX <sub>1</sub>	...	EBX <sub>0</sub>	EBX <sub>1</sub>	...	Inf
1	1	...	0	0	...	1
1	1	...	1	0	...	1
0	0	...	1	1	...	1
0	0	...	0	0	...	1
...	...	...	...	...	...	0

# Captures Conditional & Indirect Dependencies



ebx → eax	eax → eax
CF=1, EAX=542, EBX=19, ECX=7, ...	CF=0, EAX=12, EBX=4, ECX=1023...
CF=1, EAX=32, EBX=3, ECX=0, ...	CF=0, EAX=42, EBX=11, ECX=13, ...
CF=1, EAX=873, EBX=32, ECX=1, ...	CF=0, EAX=2, EBX=3, ECX=33, ...
...	...

# TaintInduce: Outputs a succinct rule

- Use inference technique to learn a succinct rule for the observed function

CF=0, EAX=12, ...Z	False
CF=1, EAX=333, ...	True
CF=0, EAX=42, ...	False
CF=0, EAX=44, ...	False
CF=1, EAX=873, ...	True
CF=0, EAX=1023, ...	False
CF=0, EAX=33, ...	False
CF=1, EAX=32, ...	True
CF=0, EAX=2, ...	False
...	DC

Inference



IF

CF=1

True

THEN (EBX<sub>0</sub> → EAX<sub>0</sub>)

ELSE (EAX<sub>0</sub> → EAX<sub>0</sub>)

# Results: Comparison with State-of-the-art

Matches: 93.27% - 99.5% with existing hand-written tools.  
Only 0.28% discrepancies are errors in TaintInduce.

X86 Instructionsxw	Arith	Comp	Jump	Move	Cond	FPU	SIMD	Misc	Total
TaintInduce	43	9	33	33	60	85	259	28	550
libdft	15	5	1	30	32	X	X	8	91
Triton	38	9	19	33	32	X	144	13	288
TEMU	7	1	2	3	X	X	X	X	13

# Results: Coverage and Correctness

Auto-generated taint rules for 4 architectures: x86, x64, AArch 64, MIPS-I  
with no mistakes for ~71% of the instructions

Methodology: train for 100 seeds, test on 1000 random inputs for each instruction

	Arith	Comp	Jump	Move	Cond	FPU	SIMD	Misc
x86	✓	✓	✓	✓	✓	✓	✓	✓
x64	✓	✓	✓	✓	✓	✓	✓	✓
AArch64	✓	✓	✓	✓	✓	✓	✓	✓
MIPS-I	✓	✓	✓	✓	-	-	-	-

Room for Future Work: Learn precise rules for all the instructions....




# PA $\longrightarrow$ ML: Deductive Reasoning

Joint work with:

Teodora Baluta, Shiqi Shen, Shweta Shinde, Kuldeep S. Meel (CCS 2019)\_

# Concerns with ML Systems

## Robustness




[Adversarial Examples – Goodfellow et al.]

## Fairness




## Memorization



[Trojaning Attacks - Liu et. al]

# Qualitative Verification

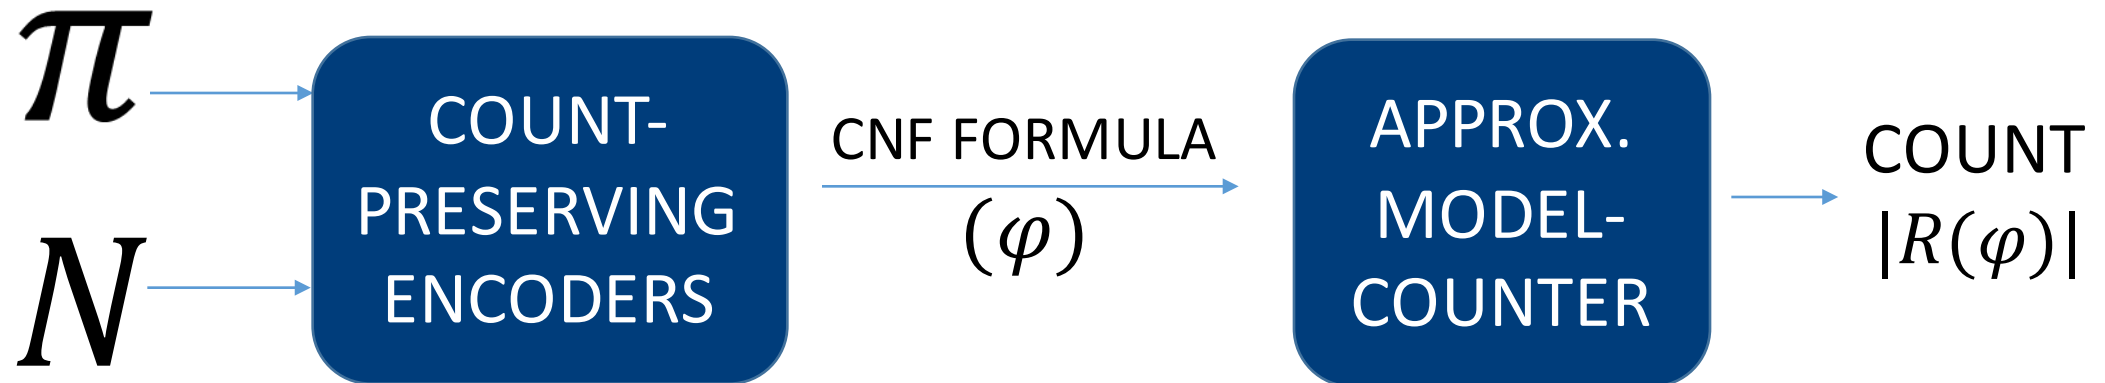


But, the network or property is often stochastic...

# Quantitative Verification



# NPAQ: A Quantitative Verifier For Neural Nets



PAC-style Soundness Guarantees:

$$\Pr[(1 + \epsilon)^{-1} |R(\varphi)| \leq r \leq (1 + \epsilon)|R(\varphi)|] \geq 1 - \delta$$

Error  
Tolerance


True  
Count

Approx.  
Count

Confidence

# NPAQ Results

- 84 models of up to 51,410 parameters
- 1,056 encoded CNF formulae
- 3 applications on MNIST and UCI Adult datasets:
  - **Robustness:** How many adversarial samples within some perturbation distance?
  - **Fairness:** How often will a prediction change (favorably) if the gender of the applicant is changed, keeping all else constant?
  - **Trojan attack efficiency:** How often will an image with a trigger result in desired misclassification?



97.1% of the encoded formula solved within 24 hours timeout each

# Key Takeaways

- **Machine Learning Helps Program Analysis**

- When program, property or analysis rules are uncertain
- Provides powerful approximate representations and solving tools
- Specific Applications:
  - Neuro-Symbolic Execution
  - Automatically Learning Taint Rules

- **Program Analysis Helps Machine Learning**

- By verifying properties
- SAT/SMT-based quantitative reasoning is a powerful tool
- Specific Applications: Fairness, Robustness, Memorization

**Thank you!**