

# Verifying Quantized Neural Networks using SMT-Based Model Checking

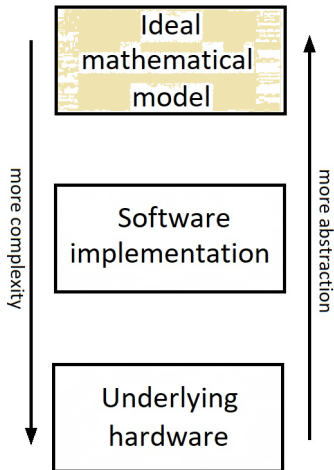
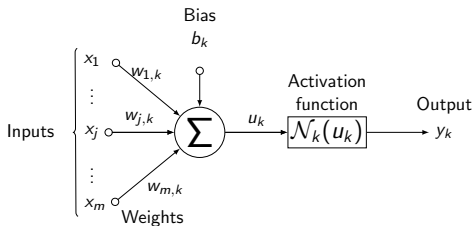
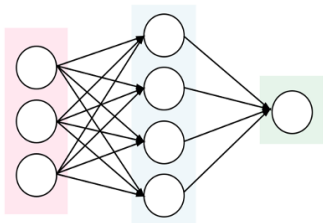
Luiz Sena, Xidan Song, Erickson Alves, Iury Bessa, Edoardo  
Manino, Lucas Cordeiro, Eddie de Lima Filho

EnnCore, University of Manchester, Universidade Federal do Amazonas

February 28, 2022

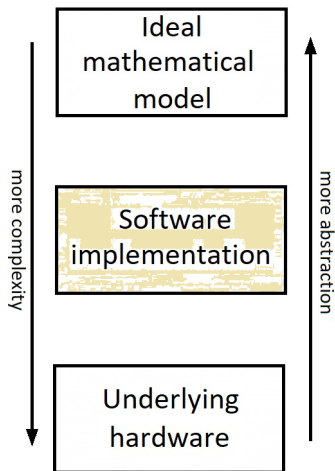


# Neural networks: just mathematical functions?



# Neural networks: looking at the source code

```
1 float potential(float *w,  
2               unsigned int w_len,  
3               float *x,  
4               unsigned int x_len,  
5               float b) {  
6  
7     if (w_len != x_len) {  
8         return 0;  
9     }  
10  
11     float result = 0;  
12  
13     for (unsigned int i = 0; i < w_len;  
14         ++i) {  
15         result += w[i] * x[i];  
16     }  
17     result += b;  
18  
19     return result;  
20 }
```



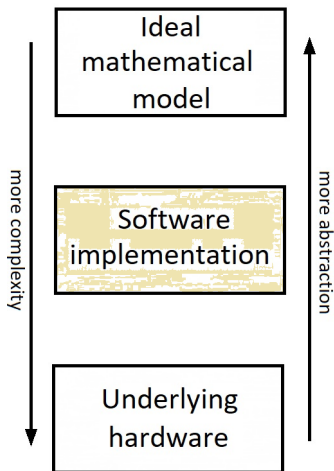
# Research challenges

## Verifying quantized NN

- ▶ Even floating-point is quantized!
- ▶ Fixed-point/integer arithmetics for low-power devices
- ▶ Approximated activation functions
- ▶ Complexity NP  $\rightarrow$  PSPACE-hard

## More software idiosyncrasies

- ▶ NaN, overflow, underflow
- ▶ Memory bugs, buffer overflows
- ▶ Concurrent execution bugs (GPUs)



# Quantization effects

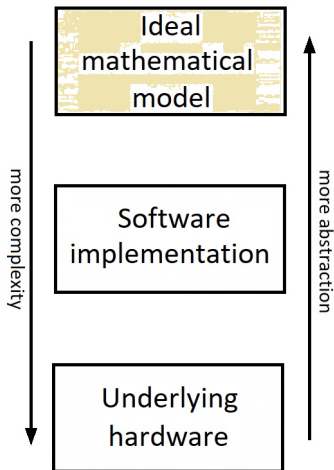
		Number of bits													
Safety Prop.		6	7	8	9	10	11	12	13		28	29	30	31	32
Set.	$R_{40}$	S	S	F	S	S	S	S	S	...	S	S	S	S	S
	$R_{50}$	S	S	F	F	F	F	F	F	...	F	F	F	F	S
Vers.	$R_{20}$	S	F	S	S	S	S	S	S	...	S	S	S	S	S
	$R_{30}$	S	F	S	S	S	S	S	S	...	S	S	S	S	S
	$R_{40}$	S	F	S	F	F	F	S	S	...	S	S	S	S	S
	$R_{50}$	S	F	F	F	F	F	F	F	...	F	F	F	F	F
Virg.	$R_{20}$	S	F	S	S	S	S	S	S	...	S	S	S	S	S
	$R_{30}$	S	F	S	S	S	S	S	S	...	S	S	S	S	S
	$R_{40}$	S	F	S	S	F	S	S	S	...	S	S	S	S	S
	$R_{50}$	S	F	F	F	F	F	F	F	...	F	F	F	F	F

Table: Effects of quantization on the safety of a NN for the Iris dataset.

# Existing works

## NN as an ideal mathematical function

- ▶ See last year's VNN-COMP'21
- ▶ Winner:  $\alpha\beta$ -CROWN
- ▶ Runners-up: VeriNet, Oval, ERAN...



# Existing works

## NN as an ideal mathematical function

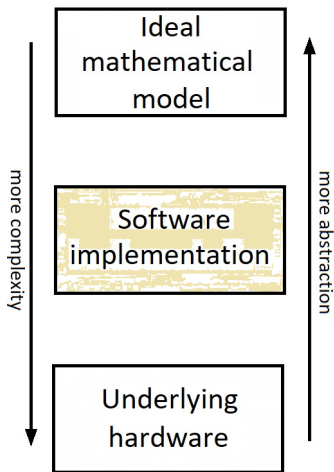
- ▶ See last year's VNN-COMP'21
- ▶ Winner:  $\alpha\beta$ -CROWN
- ▶ Runners-up: VeriNet, Oval, ERAN...

## Quantization effects

- ▶ Giacobbe *et al.*, 2019 (ReLU-N)
- ▶ Henzinger *et al.*, 2021 (ReLU-N)
- ▶ Baranowski *et al.*, 2020 (fixed-point)

## Other implementation effects

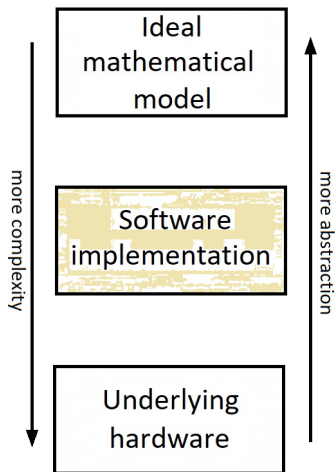
- ▶ Odena *et al.*, 2019 (fuzz testing)
- ▶ Sena *et al.*, 2019 (CUDA)



# Our verification framework (high-level view)

## Goal

- ▶ Support floating-point, fixed-point, integer and binary arithmetic
- ▶ Support all activation functions
- ▶ Let the user specify a wide range of safety properties





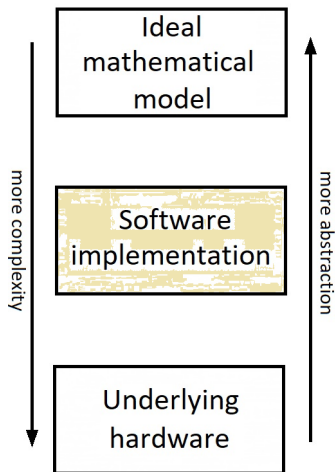
# Our verification framework (high-level view)

## Goal

- ▶ Support floating-point, fixed-point, integer and binary arithmetic
- ▶ Support all activation functions
- ▶ Let the user specify a wide range of safety properties

## Main ideas

- ▶ Apply model checking techniques
- ▶ C code as a low level abstraction
- ▶ Safety property with `assume()` and `assert()` instructions
- ▶ Convert code + property into SMT
- ▶ Check satisfiability of SMT formula



# Our verification framework (SMT encoding)

```
1 int main() {
2   _Bool x, y;
3   int a, b, f;
4   x = nondet_bool();
5   y = nondet_bool();
6   a = ((2*x) - (3*y));
7   a = a < 0 ? 0 : a;
8   b = (x + (4*y));
9   b = b < 0 ? 0 : b;
10  f = ((3*x) + y);
11  f = f < 0 ? 0 : f;
12  assert(a <= 2 && b <= 5 && f <= 4);
13  return 0;
14 }
```

SSA

```
1 x1 == nondet_symbol(nondet0)
2 y1 == nondet_symbol(nondet1)
3 a1 == 2 * (int)x1 - 3 * (int)y1
4 a2 == (a1 < 0 ? 0 : a1)
5 b1 == (int)x1 + 4 * (int)y1
6 b2 == (b1 < 0 ? 0 : b1)
7 f1 == 3 * (int)x1 + (int)y1
8 f2 == (f1 < 0 ? 0 : f1)
9 (assert) a2 <= 2
0 (assert) b2 <= 5
1 (assert) f2 <= 4
```

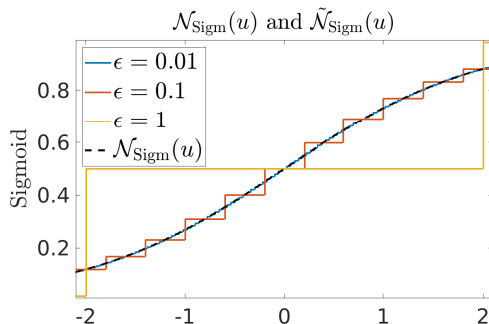
aggressive  
simplifications

SMT

```
6 ...
7 (assert(= a1 (- (* 2 x1) (* 3 y1))))
8 (assert (and (< a1 0) (= a2 0)))
9 (assert (and (> a1 0) (= a2 a1)))
10 ...
```

```
1 x1 == nondet_symbol(nondet0)
2 y1 == nondet_symbol(nondet1)
3 a1 == 2 * (int)x1 - 3 * (int)y1
4 a2 == (a1 < 0 ? 0 : a1)
```

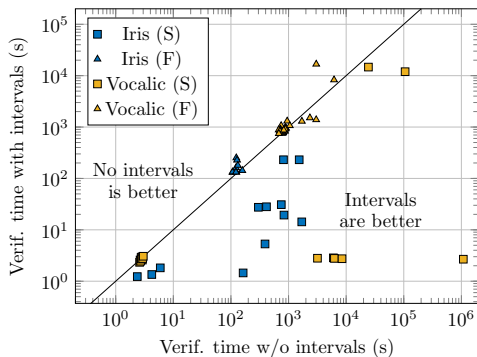
# Our verification framework (activation functions)



## Encoding non-linear functions

- ▶ Piecewise linear (e.g. ReLU) → if-then-else
- ▶ Others (e.g. sigmoid, tanh) → lookup table (DSP-style)
- ▶ Speeds up both inference and verification!

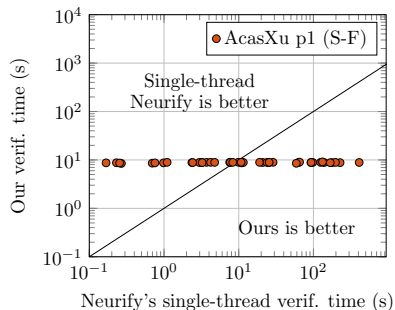
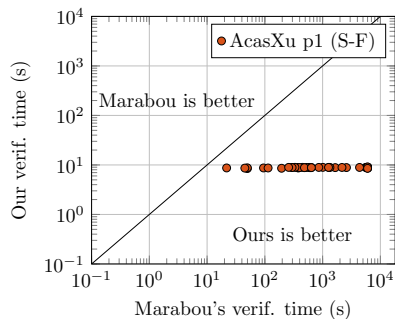
# Our verification framework (interval analysis)



## Input set propagation

- ▶ Transferable from verification of ideal NNs
- ▶ Generates an overapproximation of the neuron values
- ▶ Reduces the search space for safe (S) instances

# Our verification framework (comparison with SOTA)



Warning: this is not an equal contest!

- ▶ Comparison between infinite precision and fixed-point
- ▶ Useful as a qualitative result

# Conclusions

## Summary

- ▶ Implementations of NNs are software!
- ▶ Quantization effects, finite arithmetic, other potential bugs
- ▶ Higher theoretical complexity than verifying ideal NNs
- ▶ Positive note: similar verification time in practice

## Further resources

- ▶ Try our QNNVerifier tool:
- ▶ <https://arxiv.org/abs/2111.13110>
- ▶ Read our pre-print journal paper:
- ▶ <https://arxiv.org/abs/2106.05997JournalPaper>

Thank you!