ML FW tests

Naoki Kimura
09/05/2024

My test sample

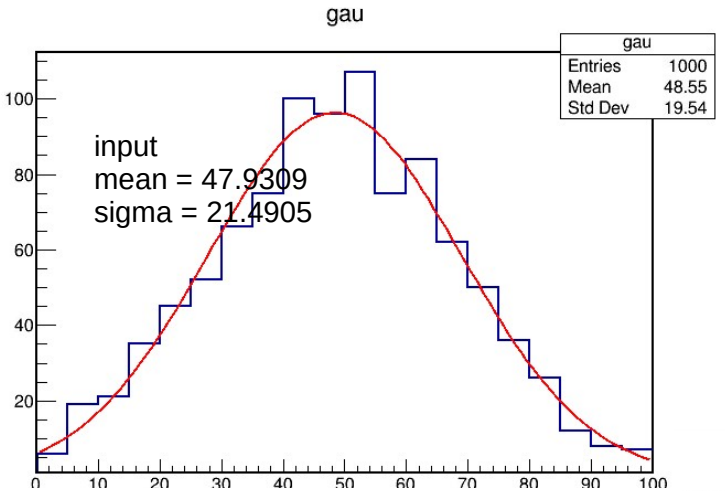

20 bins histogram
1000 entries of Gaussian
Random Mean(10 - 90)
Random Sigma (10 – 30)
Input:
20 bin's data
Output:
Mean and Sigma

Best results should be binned likelihood fit

Uncertainty = ~ 1.0

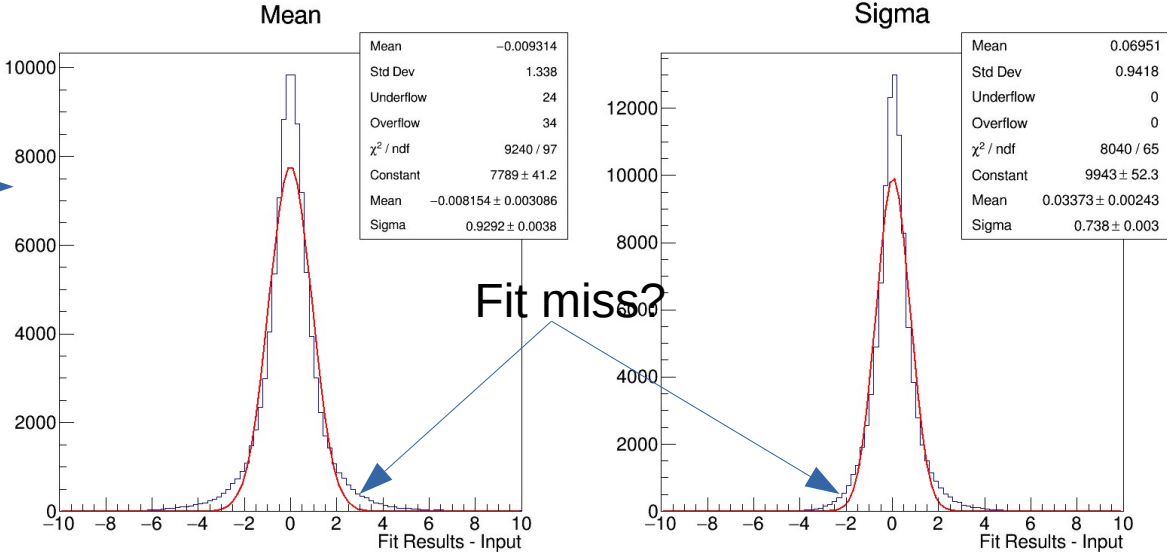

Fit miss?

To get image, make similar performance ML model and check resource and time for FPGA.

```
Model
  Precision:         fixed<20,8>
  ReuseFactor:       1
  Strategy:          Latency
  BramFactor:        1000000000
  TraceOutput:       False
LayerName
  fc1_input
    Trace:           True
    Precision
      result:        ap_fixed<18,6>
  fc1
    Trace:           True
    Precision
      result:        ap_fixed<20,5>
      weight:        ap_fixed<20,2>
      bias:          ap_fixed<12,2>
  fc1_linear
    Trace:           True
    Precision
      result:        ap_fixed<19,5>
  relu1
    Trace:           True
    Precision
      result:        ap_fixed<20,5>
  fc2
    Trace:           True
    Precision
      result:        ap_fixed<20,6>
      weight:        ap_fixed<20,2>
      bias:          ap_fixed<10,2>
  fc2_linear
    Trace:           True
    Precision
      result:        ap_fixed<20,6>
  relu2
    Trace:           True
    Precision
      result:        ap_fixed<20,6>
  fc3
    Trace:           True
    Precision
      result:        ap_fixed<20,7>
      weight:        ap_fixed<20,2>
      bias:          ap_fixed<10,1>
  fc3_linear
    Trace:           True
    Precision
      result:        ap_fixed<20,7>
  relu3
    Trace:           True
    Precision
      result:        ap_fixed<20,7>
  output
    Trace:           True
    Precision
      result:        ap_fixed<12,8>
      weight:        ap_fixed<20,2>
      bias:          ap_fixed<6,2>
  output_linear
    Trace:           True
    Precision
      result:        ap_fixed<12,8>
  linear
    Trace:           True
    Precision
```

Bit width optimization on Qkeras

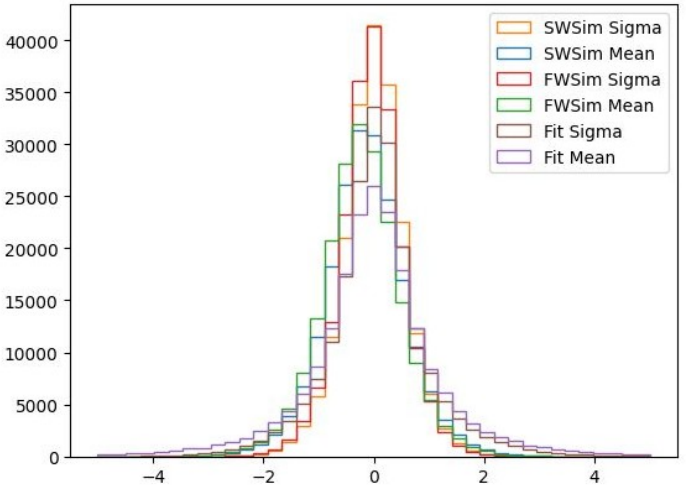

keras 20,8 optimized (reff.)

Summar

Name	BRAM_18K	DSP48E	FF	LUT	URAM
DSP	-	-	-	-	-
Expression	-	-	0	6	-
FIFO	-	-	-	-	-
Instance	-	4038	13096	130595	-
Memory	-	-	-	-	-
Multiplexer	-	-	-	36	-
Register	-	-	4146	-	-
Total	0	4038	17242	130637	0
Available	4320	6840	2364480	1182240	960
Available SLR	1440	2280	788160	394080	320
Utilization (%)	0	59	~0	11	0
Utilization SLR (%)	0	177	2	33	0

Qkeras 20,8 equivalent precision

Utilization Estimates

Summary

Name	BRAM_18K	DSP48E	FF	LUT	URAM
DSP	-	-	-	-	-
Expression	-	-	0	6	-
FIFO	-	-	-	-	-
Instance	-	1133	3728	35596	-
Memory	-	-	-	-	-
Multiplexer	-	-	-	36	-
Register	-	-	3686	-	-
Total	0	1133	7414	35638	0
Available	4320	6840	2364480	1182240	960
Available SLR	1440	2280	788160	394080	320
Utilization (%)	0	16	~0	3	0
Utilization SLR (%)	0	49	~0	9	0

Utilization Post-Synthesis | **Post-Implementation**

Graph | **Table**

Resource	Utilization	Available	Utilization %
LUT	23555	1182240	1.99
LUTRAM	484	591840	0.08
FF	9466	2364480	0.40
BRAM	11	2160	0.51
DSP	1080	6840	15.79
IO	2	832	0.24
BUFG	2	1800	0.11
PLL	1	60	1.67

params

~3 time of # params

~1/3 time of # params

Layer (type)	Output Shape	Param #
fc1 (QDense)	(None, 64)	1344
relu1 (QActivation)	(None, 64)	0
fc2 (QDense)	(None, 32)	2080
relu2 (QActivation)	(None, 32)	0
fc3 (QDense)	(None, 32)	1056
relu3 (QActivation)	(None, 32)	0
output (QDense)	(None, 2)	66
linear (Activation)	(None, 2)	0

Total params: 4,546
Trainable params: 4,546
Non-trainable params: 0

Layer (type)	Output Shape	Param #
fc1 (QDense)	(None, 128)	2688
relu1 (QActivation)	(None, 128)	0
fc2 (QDense)	(None, 64)	8256
relu2 (QActivation)	(None, 64)	0
fc3 (QDense)	(None, 64)	4160
relu3 (QActivation)	(None, 64)	0
output (QDense)	(None, 2)	130
linear (Activation)	(None, 2)	0

Total params: 15,234
Trainable params: 15,234
Non-trainable params: 0

Layer (type)	Output Shape	Param #
fc1 (QDense)	(None, 32)	672
relu1 (QActivation)	(None, 32)	0
fc2 (QDense)	(None, 16)	528
relu2 (QActivation)	(None, 16)	0
fc3 (QDense)	(None, 16)	272
relu3 (QActivation)	(None, 16)	0
output (QDense)	(None, 2)	34
linear (Activation)	(None, 2)	0

Total params: 1,506
Trainable params: 1,506
Non-trainable params: 0

Resour
ce over
the limit

Summary

Name	BRAM_18K	DSP48E	FF	LUT	URAM
DSP	-	-	-	-	-
Expression	-	-	0	6	-
FIFO	-	-	-	-	-
Instance	-	1133	3728	35596	-
Memory	-	-	-	-	-
Multiplexer	-	-	-	36	-
Register	-	-	3686	-	-
Total	0	1133	7414	35638	0
Available	4320	6840	2364480	1182240	960
Available SLR	1440	2280	788160	394080	320
Utilization (%)	0	16	~0	3	0
Utilization SLR (%)	0	49	~0	9	0

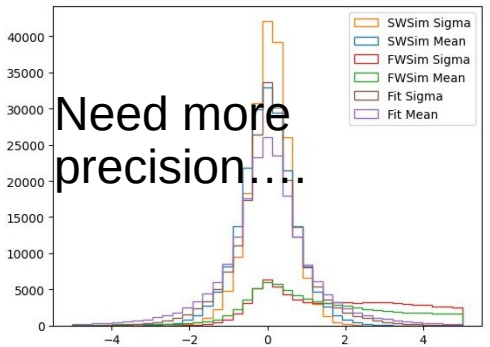

Need more precision….

Summar

Name	BRAM_18K	DSP48E	FF	LUT	URAM
DSP	-	-	-	-	-
Expression	-	-	0	6	-
FIFO	-	-	-	-	-
Instance	-	340	1804	12453	-
Memory	-	-	-	-	-
Multiplexer	-	-	-	36	-
Register	-	-	1998	-	-
Total	0	340	3802	12495	0
Available	4320	6840	2364480	1182240	960
Available SLR	1440	2280	788160	394080	320
Utilization (%)	0	4	~0	1	0
Utilization SLR (%)	0	14	~0	3	0

Also need more precision

Best Results

```
orig Configuration
Model
  Precision:         fixed<20,8>
  ReuseFactor:       1
  Strategy:          Latency
  BramFactor:        1000000000
  TraceOutput:       False
LayerName
  fc1_input
    Trace:           True
    Precision
      result:        ap_fixed<18,6>
  fc1
    Trace:           True
    Precision
      result:        ap_fixed<20,5>
      weight:        ap_fixed<20,2>
      bias:          ap_fixed<12,2>
  fc1_linear
    Trace:           True
    Precision
      result:        ap_fixed<19,5>
  relu1
    Trace:           True
    Precision
      result:        ap_fixed<20,5>
  fc3
    Trace:           True
    Precision
      result:        ap_fixed<20,7>
      weight:        ap_fixed<20,2>
      bias:          ap_fixed<10,1>
  fc3_linear
    Trace:           True
    Precision
      result:        ap_fixed<20,7>
  relu3
    Trace:           True
    Precision
      result:        ap_fixed<20,7>
  output
    Trace:           True
    Precision
      result:        ap_fixed<12,8>
      weight:        ap_fixed<20,2>
      bias:          ap_fixed<6,2>
  output_linear
    Trace:           True
    Precision
      result:        ap_fixed<12,8>
  linear
    Trace:           True
    Precision
      result:        ap_fixed<12,8>
```

Layer (type)	Output Shape	Param #
fc1 (QDense)	(None, 32)	672
relu1 (QActivation)	(None, 32)	0
fc3 (QDense)	(None, 16)	528
relu3 (QActivation)	(None, 16)	0
output (QDense)	(None, 2)	34
linear (Activation)	(None, 2)	0

Total params: 1,234
Trainable params: 1,234
Non-trainable params: 0

- qkeras
- 2 layer (skip layer 2)
- 1200 parameters
- full optimized precision (max 20 bits)
- 3 time more training loop (epoch)

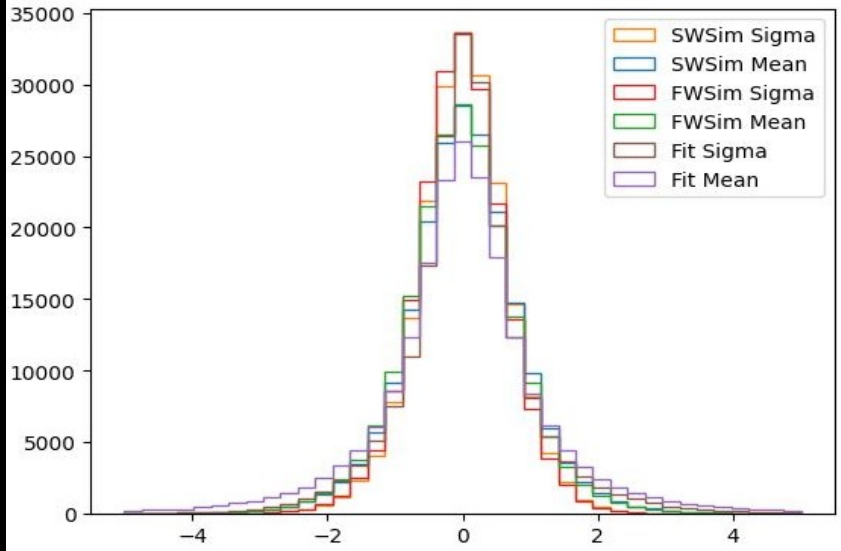

Obs.-Truth

Summar

Name	BRAM_18K	DSP48E	FF	LUT	URAM
DSP	-	-	-	-	-
Expression	-	-	0	6	-
FIFO	-	-	-	-	-
Instance	-	287	642	9702	-
Memory	-	-	-	-	-
Multiplexer	-	-	-	36	-
Register	-	-	1658	-	-
Total	0	287	2300	9744	0
Available	4320	6840	2364480	1182240	960
Available SLR	1440	2280	788160	394080	320
Utilization (%)	0	4	~0	~0	0
Utilization SLR (%)	0	12	~0	2	0

More tight opt. (before plateau)

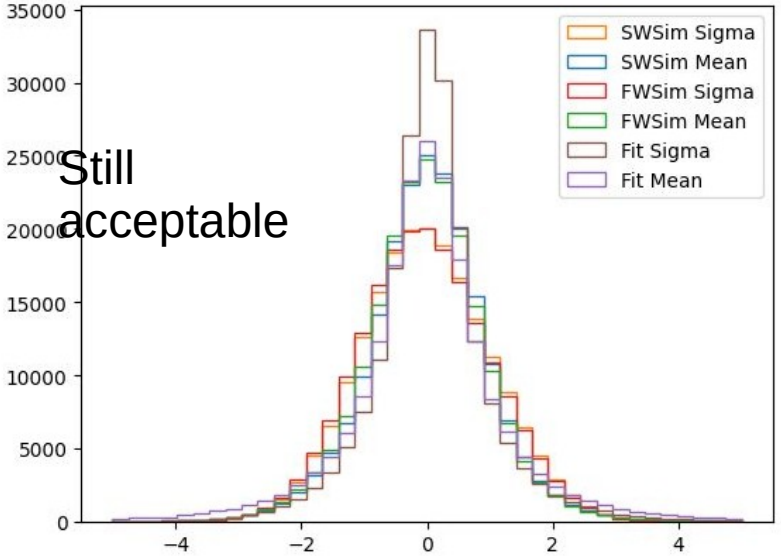

Still acceptable

Obs.-Truth

Summary

Name	BRAM_18K	DSP48E	FF	LUT	URAM
DSP	-	-	-	-	-
Expression	-	-	0	6	-
FIFO	-	-	-	-	-
Instance	-	111	672	3996	-
Memory	-	-	-	-	-
Multiplexer	-	-	-	36	-
Register	-	-	706	-	-
Total	0	111	1378	4038	0
Available	4320	6840	2364480	1182240	960
Available SLR	1440	2280	788160	394080	320
Utilization (%)	0	1	~0	~0	0
Utilization SLR (%)	0	4	~0	1	0

```
Layer (type)                 Output Shape              Param #
=================================================================
fc1 (QDense)                 (None, 16)                336

relu1 (QActivation)          (None, 16)                0

fc3 (QDense)                 (None, 8)                 136

relu3 (QActivation)          (None, 8)                 0

output (QDense)              (None, 2)                 18

linear (Activation)          (None, 2)                 0

=================================================================
Total params: 490
Trainable params: 490
Non-trainable params: 0
```

Still precision has improvement point

very small resource usage, enough performance

BDT

- Original hls4ml’s conifer is old and did not work well.
- BDT para name was changed, e.g. n_feature → max_feature
- GradientBoostingRegressor did not support multi output
- MultiOutputRegressor is not supported in conifer (?)
- XGBRegressor with Conifer env setup then conifer.converters.convert_from_xgboost work
- Can not check FW 2nd output….

XGBoostRegressor

n_estimators=120, max_depth=3

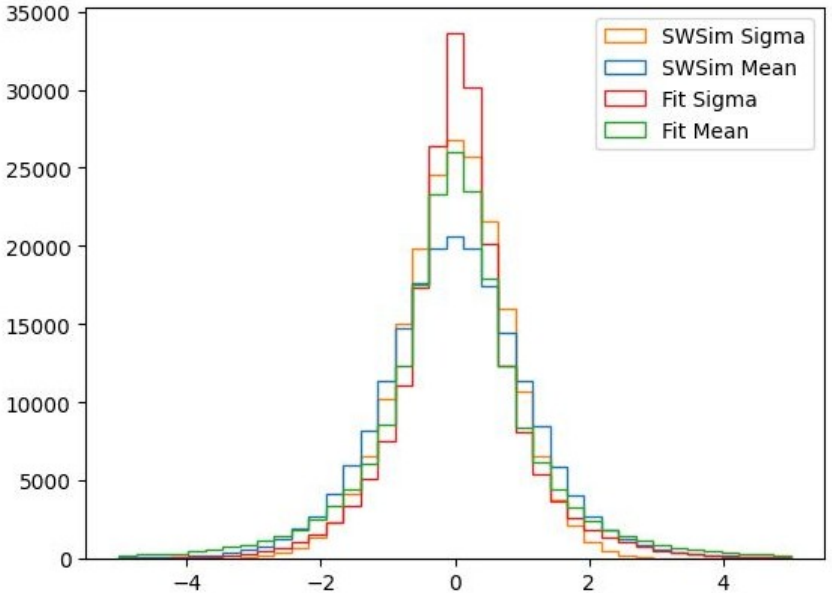

n_estimators=40, max_depth=3

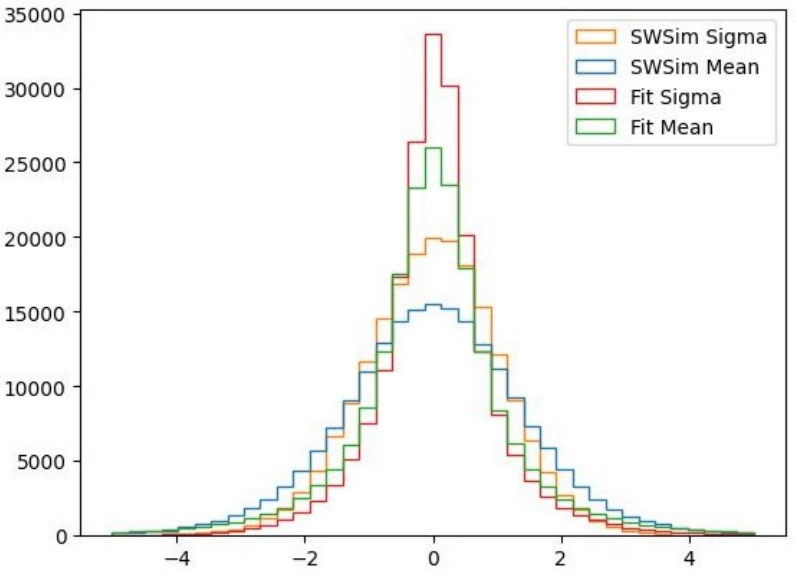

XGB Conifer transfer

```
test
[[11.0591 10.1829]
 [32.4552 29.2091]
 [32.1445 26.2538]
 [78.6988 13.4342]
 [66.3327 11.4321]
 [37.8392 28.8438]
 [62.6105 22.6655]
 [37.8469 26.6227]
 [60.1769 12.5171]
 [27.6076 29.584 ]
 [89.8138 22.2686]
 [39.9187 11.25  ]
 [45.8845 11.9377]
 [77.7424 29.2906]
 [80.5033 27.3695]]
skl
[[11.403256 11.425323]
 [32.249268 26.790737]
 [29.819868 24.552683]
 [78.529106 13.474673]
 [67.420906 11.437288]
 [41.162796 28.331532]
 [65.00193  20.876379]
 [34.598995 25.571938]
 [62.122936 11.396581]
 [27.362764 29.004648]
 [88.53254  22.799871]
 [39.49116  10.862266]
 [45.787045 11.793477]
 [76.49551  27.92482 ]
 [81.46407  27.805416]]
cnf
[ 21.82857631  58.04000799  53.37255412  91.00378441  77.85819469
  68.49433274  84.87831232  59.17093968  72.51951524  55.36741129
 110.33241539  49.35343489  56.58053129 103.420331   108.26949423]
```

true

SW

FW (mean)

Precision: ap_fixed<64,32>

?

Build error…..
Vivado_hls
ERROR: [HLS 200-101] 'open_solution': Unknown option '-flow_target'.
Vitis_hls
ERROR: [SIM 211-100] 'csim_design' failed: compilation error(s).

Noah’s egamma BDT

	trial	pooling	reduce	n_estimators	max_depth	...	test_loss	test_rejection	test_quant_rejection	FF	LUT
91	91	max	True	63	5	...	0.183702	10.398537	9.284124	7902	36990
93	93	max	True	87	4	...	0.188241	10.572739	9.186414	1933	26645
39	39	max	False	40	6	...	0.184578	11.017203	9.142433	7309	34081
27	27	max	False	91	4	...	0.192623	10.108413	9.044053	2027	27252
92	92	max	False	86	4	...	0.189861	10.355819	9.006847	1931	25263
89	89	max	False	66	5	...	0.186505	11.172620	8.952708	7103	33568
23	23	max	False	98	4	...	0.186270	10.406617	8.729719	2226	29546
59	59	max	False	95	4	...	0.186796	10.342754	8.548206	2103	28734
11	11	max	False	80	4	...	0.195152	10.223503	8.472520	1883	24557
19	19	max	False	80	4	...	0.195174	10.223655	8.472180	1887	24615
54	54	max	False	78	4	...	0.198135	10.103973	8.344769	1837	23445
97	97	max	False	49	5	...	0.193440	10.525800	8.320928	6263	28067
43	43	max	True	87	4	...	0.195688	9.242605	8.241387	1649	21118
37	37	max	True	83	4	...	0.194680	9.875176	8.051194	1761	23041
33	33	max	False	59	4	...	0.214661	8.828334	7.802401	1505	17996
5	5	max	True	51	4	...	0.215044	7.871490	7.677613	1212	14937
58	58	max	False	58	4	...	0.211757	8.916071	7.624726	1502	17713
44	44	max	False	48	4	...	0.220270	8.469055	7.518219	896	14510

FF 1933 LUT 26645

Check Trial #93

SW and FW outputs

Trial 93
2nd higher rejection

	trial	pooling	reduce	n_estimators	max_depth	gamma	n_nodes	train_loss	train_rejection	test_loss	test_rejection	test_quant_rejection	FF	LUT
91	91	max	True	63	5	1.638448	1622	0.130788	26.159525	0.183702	10.398537	9.284124	7902	36990
93	93	max	True	87	4	0.570012	1216	0.148949	17.987314	0.188241	10.572739	9.186414	1933	26645

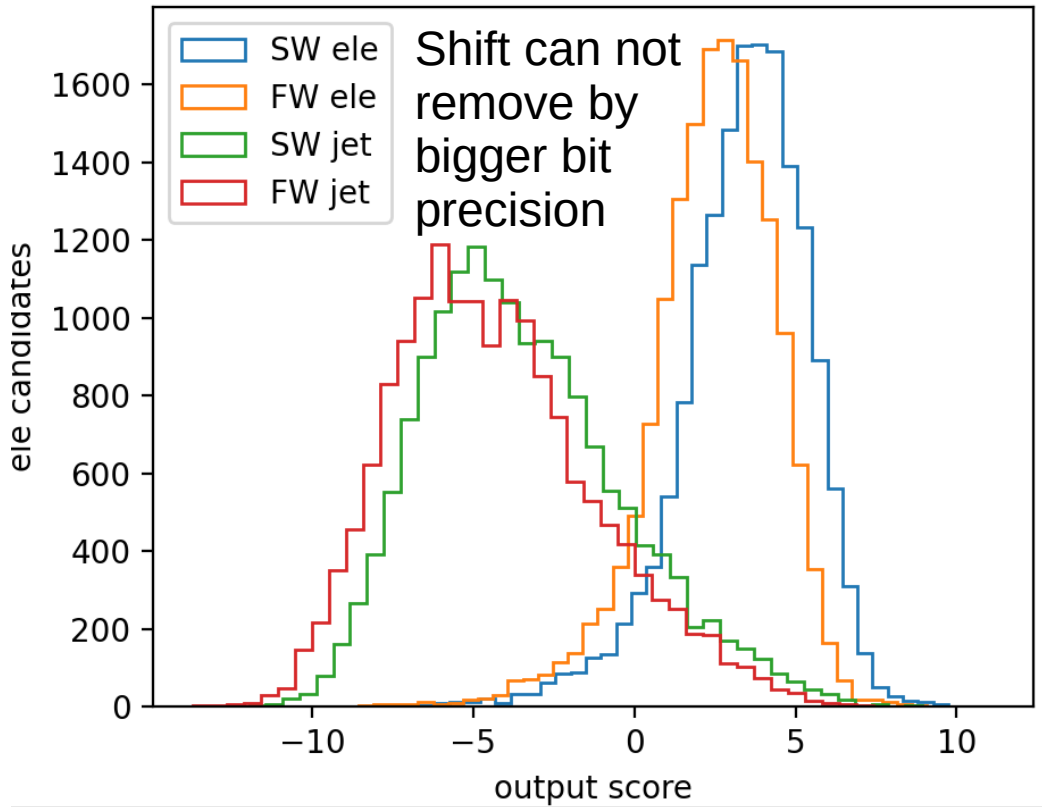

Simpler bit precision :
All 10 bits
Input int<10>
Threshold and output fixed<10,5>

QCD rejection @ 95% ele eff. :
SW 10.57
FW 9.19

Eratio 4.0

Others?

Trial 4 (40 th rejection)

[32602 rows x 66 columns]

	trial	pooling	reduce	n_estimators	max_depth	gamma	n_nodes	train_loss	train_rejection	test_loss	test_rejection	test_quant_rejection	FF	LUT
4	4	max	True	63	5	2.347221	1581	0.130723	25.600381	0.181747	10.147131	5.899807	7746	36289

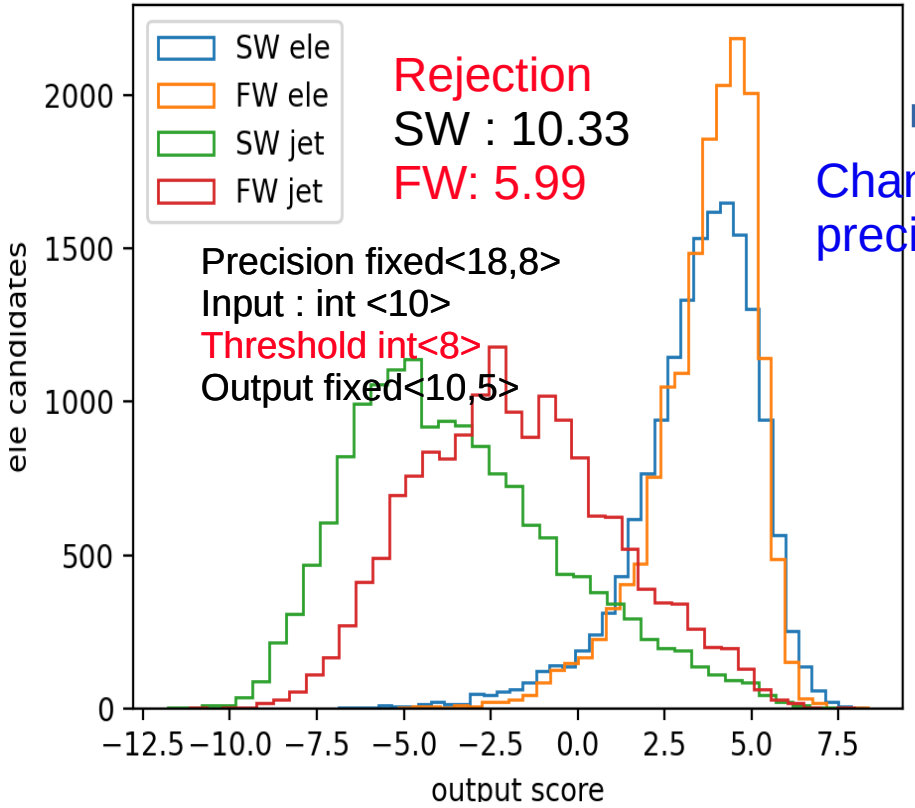

Change bit precision

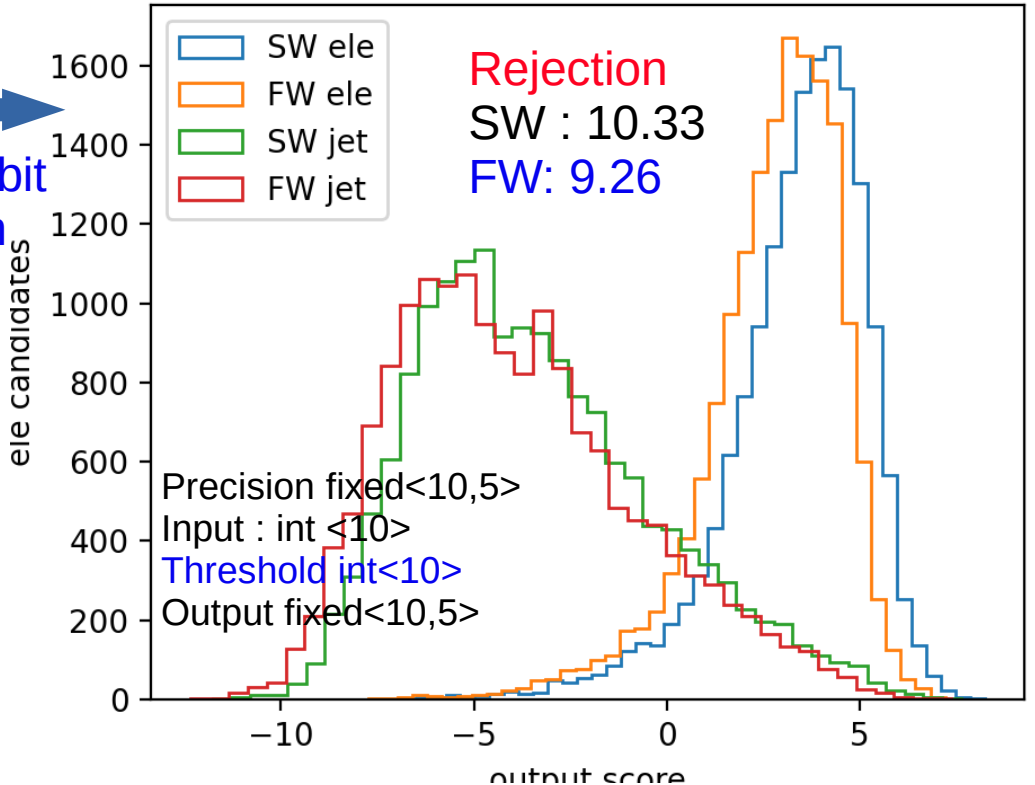

All2023.2.

Egamma BDT

Vitis synthesis on Conifer... FF 1933 LUT 26645

Vitis synthesis

Pipelined	BRAM	DSP	FF	LUT	URAM
yes	0	0	1246	37584	0
yes	0	0	1241	37582	0

Vitis implementation

	Verilog
SLICE	0
LUT	3253
FF	814
DSP	0
BRAM	0
URAM	0
LATCH	0
SRL	0
CLB	570

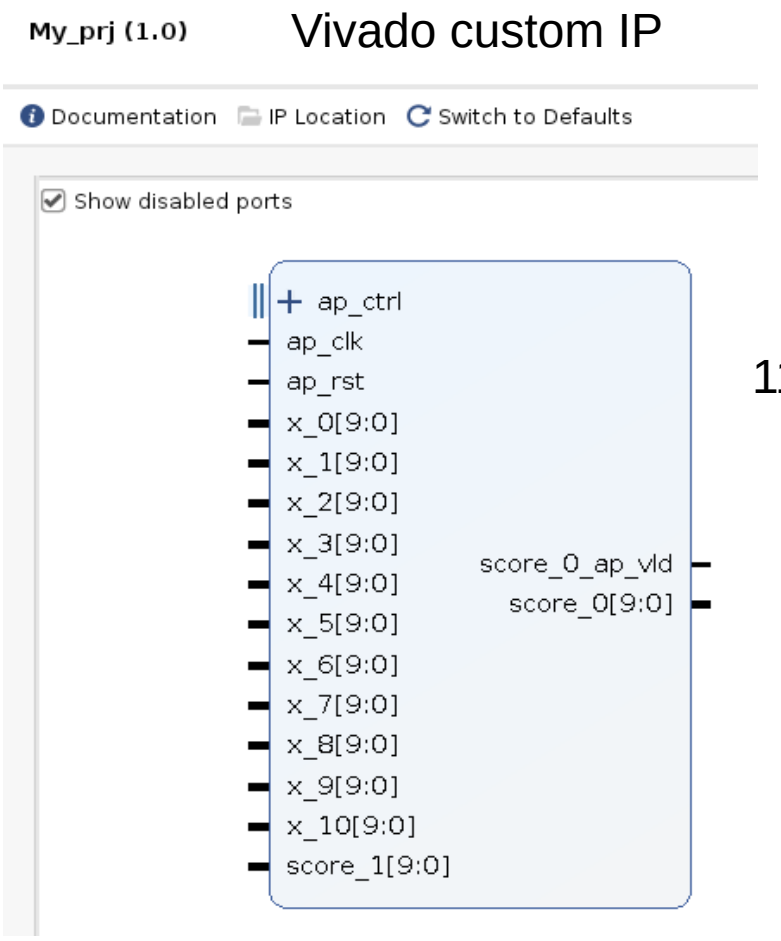

11 inputs

Next plan
-Real FW test
- cpp out check?

backup