Fast deep learning on FPGAs

The team:
Vladimir, Sioni, Jennifer
Dejan, Petr, Adrian, Hasib

Conv2D in hls4ml

- C/HLS implementation in place in the hls4ml <u>hack7</u> branch, but

  - still some troubles in matching the HLS output with expected values from Keras —> work in progress

- Adding support for different Conv2D cases

  - different data format (channels_last vs channels_first)

  - Binary/Ternary Conv2D

- Working to optimize the HLS code for latency

- Need to prepare few Conv1D models to ensure support for that too
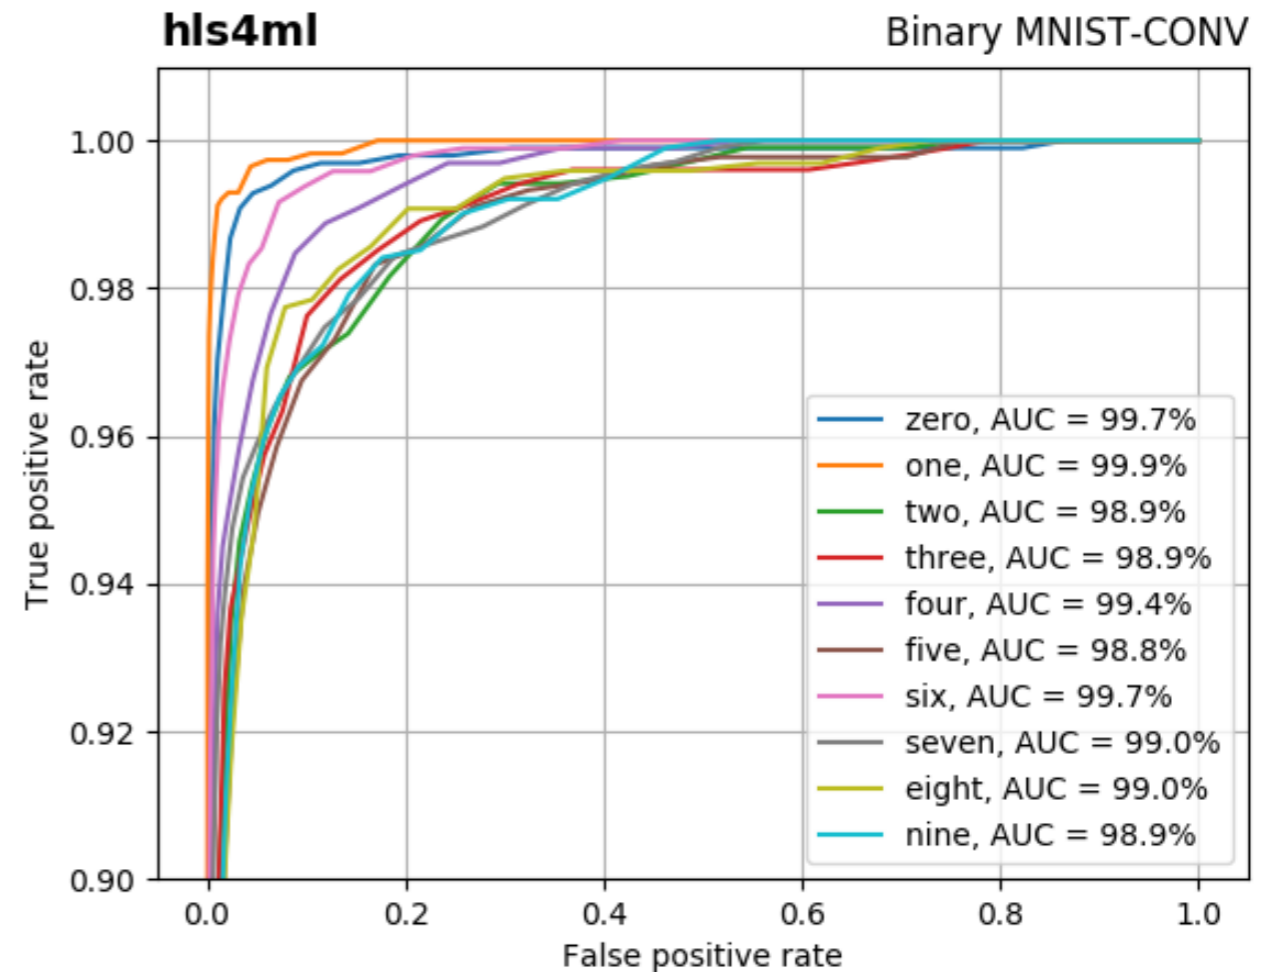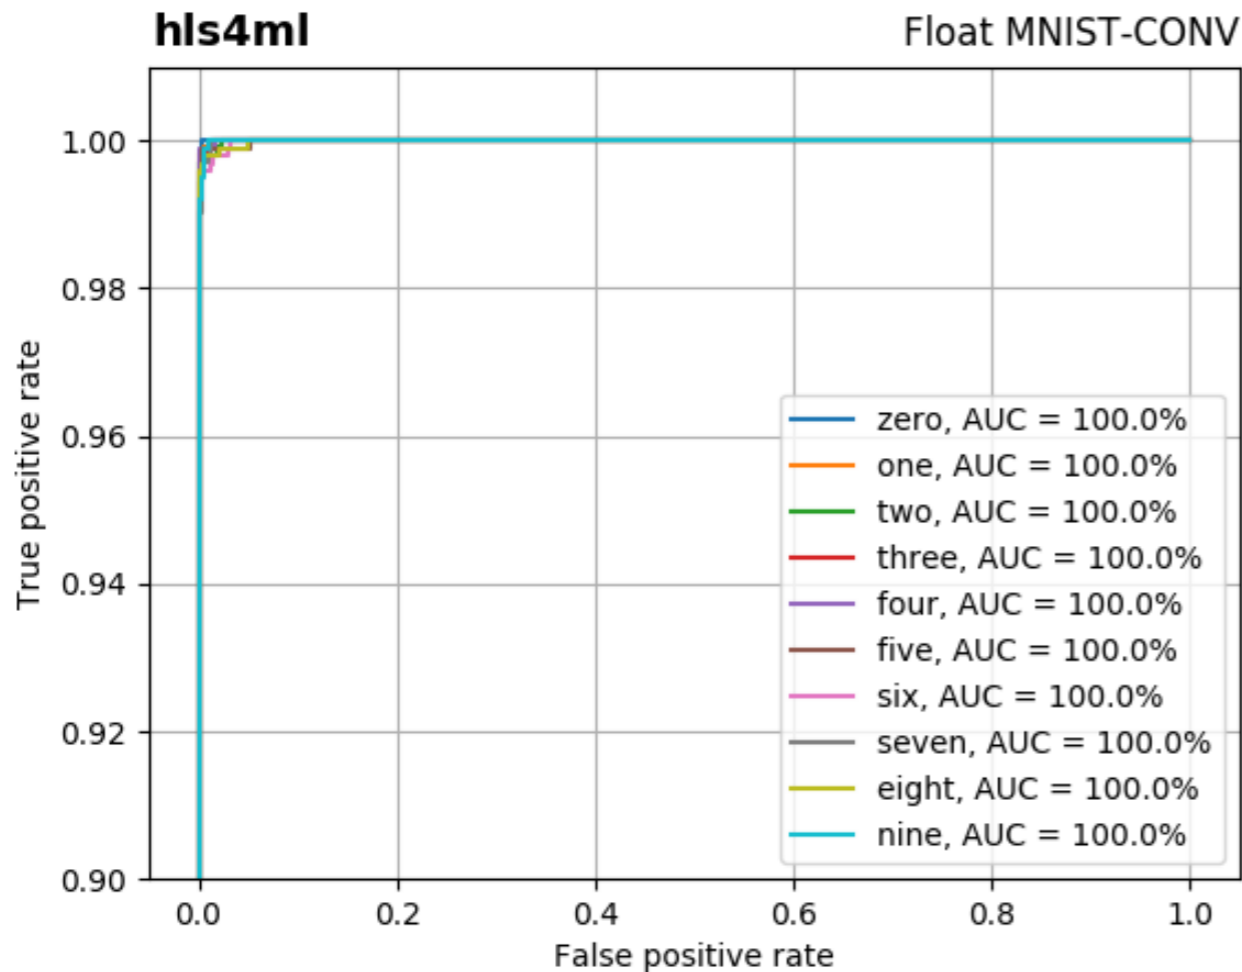
BinaryConv2D

Conv2D

Test accuracy: 0.99

BinaryConv2D

Test accuracy: 0.93

For same number of convolutions, reduced accuracy as expected
FPGA synthesis to be checked next for the two cases

01.10.2019