Single Photon Image Classification

T. Fischbacher, L. Sbaiz, Google Research, ICLR 2021-05-(03-07)

Single Photon Image Classification - Context

- Our work is about building bridges between Quantum and ML communities *via mutually accessible simple examples*.

- Our own key insight: even "small" quantum state (vector) spaces can capture high level notions such as "image shows a dress" remarkably well. Quantum ML *might* possibly become a reality before there are large-scale quantum computers.

- Interpretation of QM sometimes calls for thought experiments where we make an observer part of a quantum system ("Wigner's Friend").
  - Not technically feasible
  - Alternative "dream": Use Artificial General Intelligence implemented on Quantum Computer instead. *Still way out of reach!*
  - *What we have*: Simplistic ML simulacrum for intelligent observer, still may allow studying some relevant aspects, computationally very manageable ("Complex 1000x1000 matrices").

- Alternatively, our work also is an early example for *using ML for quantum circuit discovery/planning*.

Quantum Mechanics

Quantum Mechanics is not just "a small correction term in some equations".

Quantum interference has profound implications.

Consider this problem [Deutsch, Josza, 1992]: Given a function f that maps a n-bit word to a single bit which is zero on exactly 0, 2^n, or 2^{n-1} inputs, *then a single quantum evaluation of f can tell us if f is a constant function*.

Quantum ML

Our paper illustrates ML implications of quantum interference in the simplest possible setting: "Textbook QM problem meets textbook ML problem".

Starting question: "What is the fastest image classification system one can theoretically build in our Universe, and what accuracy can that have?"

"Fastest possible" = "need to make a call after sensing the first photon that came from the image."

Quantum ML

Simplifying our paper's set-up even further:

- Send photons one-after-another towards a mask.
- Mask either has one wide or two narrow slits in it.
- QM: Distribution of photons landing on a screen placed behind the mask depends on shape of the mask.
- How accurately can we infer from observing a single photon if we used the 1-slit or 2-slit mask?
- How much better can we do if we can optically transform the photon quantum state after it passed through the mask?

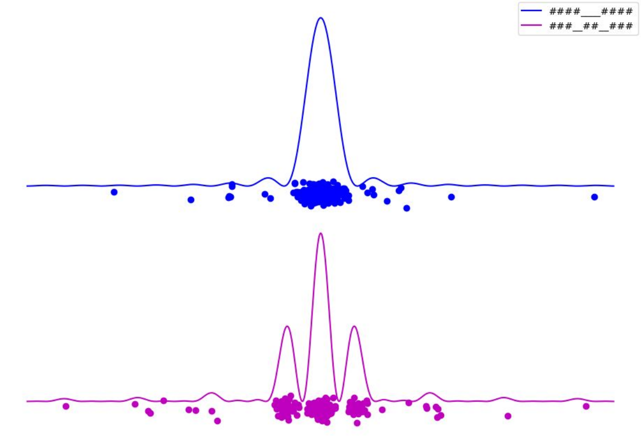

Quantum ML

"Single-Photon MNIST", classical case:

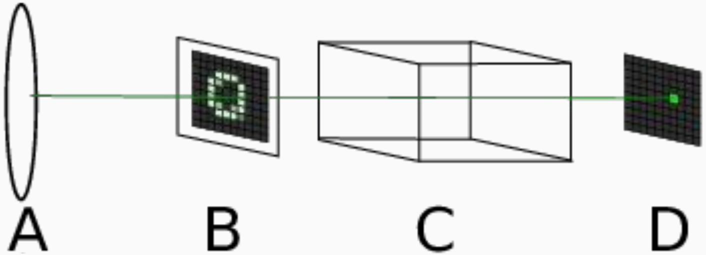

Laser illuminates screen, photon travels to single-photon detector array (SPAD), provably-optimal accuracy is MLE, *hard upper bound of 22.96%* for MNIST, "with cheating" i.e. achieved when training on the test set.
Theorem: No classical single-photon classifier can outperform lookup on this table (right).
"Fastest possible": Photon flight time + SPAD detection + 1 binary-or: ~5 ns to detection.

```
          2 2 2 2 6 6 6 6 6 6 6 6 6 6 6 6
          2 2 2 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6
        5 3 2 2 2 2 2 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6
      4 3 3 2 2 2 2 2 2 2 2 2 2 2 2 6 6 6 5 5 5 5 5 5
  3 3 4 3 3 3 3 3 3 3 3 3 3 1 1 1 1 1 1 5 5 5 5 5 5 5
    7 3 3 3 3 3 3 3 3 3 3 3 1 1 1 1 1 5 5 5 5 5 5 5 5 8
  7 7 7 3 3 7 7 3 3 7 7 9 9 1 1 1 1 1 0 5 5 5 5 5 5 5 8
  7 7 7 7 7 7 7 7 7 7 7 7 7 1 1 1 7 7 7 0 0 5 5 5 5 5 5
7 7 7 7 7 7 7 7 7 7 7 7 7 7 1 1 1 7 7 7 0 0 0 5 5 5 5 8
7 7 7 7 7 7 7 7 7 9 9 9 7 1 1 1 1 7 7 7 0 0 0 0 0 5 5 5
7 7 7 7 7 7 7 7 9 9 9 5 5 1 1 1 1 7 7 7 0 0 0 0 0 5 5 8
  4 7 7 7 7 4 4 9 9 5 5 3 1 1 1 1 7 7 7 0 0 0 0 0 0 5
  7 7 7 4 0 0 4 4 4 5 5 3 1 1 1 9 9 7 7 0 0 0 0 0 6 6
    7 4 4 0 0 4 4 4 4 4 8 1 1 1 9 9 4 6 0 0 0 0 0 6 7
    5 4 0 0 0 4 4 4 4 4 4 1 1 1 4 4 4 6 0 0 0 0 0 2 7 4
  0 3 2 0 0 0 0 4 4 4 4 1 1 1 1 7 4 4 6 6 0 0 0 0 2 2 2
    3 2 0 0 0 0 0 6 6 2 1 1 1 7 7 4 3 6 6 0 0 0 2 2 2 2
  2 3 3 2 0 0 0 0 6 6 1 1 1 1 7 7 6 6 6 6 0 2 2 2 2 2 2
  2 3 3 2 0 0 0 0 2 6 1 1 1 1 7 6 6 6 6 2 2 2 2 2 2 2 2
  2 3 3 2 2 2 0 2 2 2 1 1 1 1 6 6 3 3 2 2 2 2 2 2 2 2
    3 3 3 2 2 2 2 2 2 1 1 1 1 1 3 3 3 2 2 2 2 2 2 2 2
    3 3 3 3 3 3 3 3 3 1 1 1 1 3 3 3 3 2 2 2 2 2 2 2 2
  3 3 3 3 3 3 3 3 3 3 3 3 1 1 1 8 9 9 4 4 2 2 2 2 2 2
      3 3 3 3 7 7 7 7 7 7 7 7 9 9 9 9 9 9 9 9 9 9 2
      3 7 7 7 7 7 7 7 7 7 7 7 7 9 9 9 9 9 9 9 9 4 9
        7 7 7 7 7 7 7 7 7 7 7 7 7 7 9 9 9 9 9 9 9
            7 7 7 7 7 7 7 7 7 7 7 7 7 7 7 9
```

Quantum ML

"Single-Photon MNIST", quantum case:

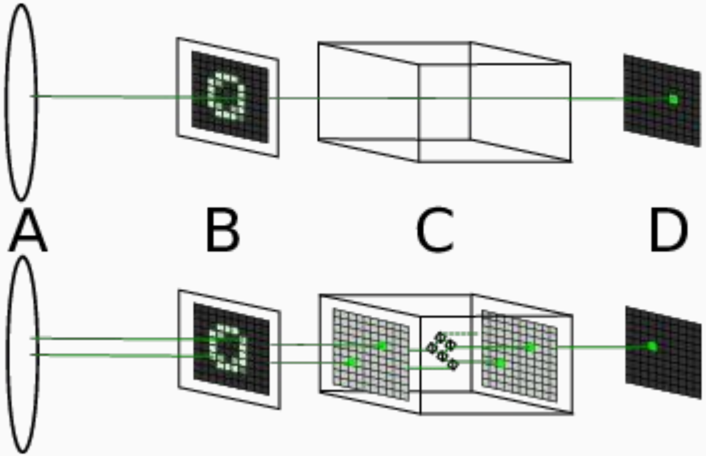

We show: If the image is illuminated coherently, and one is permitted to put an optical device in front of the sensor, one can theoretically achieve an accuracy of more than 41% on MNIST.

- Still only one photon!
- Still fast - optical transform contributes < 1 ns.
- Training on training set only, evaluating on test set.
- Uses 1-to-1 equivalence between linear optical devices and elements of the Lie group U(N): Optimize over all unitary matrices, then translate to optical circuit network list.

Conclusions

Quantum Mechanics is un-intuitive, and coming up with quantum algorithms is hard.

Our work can be seen as a pedagogical example for how to use classical ML to design quantum circuits that solve complex problems.