Quantum Machine Learning at CERN

QUANTUM TECHNOLOGY INITIATIVE

Sofia Vallecorsa

AI and Quantum Research - CERN IT

CERN

Quantum
Hype and Potential…

2019: Google

https://www.nature.com/articles/s41586-019-1666-5

2020: Hefei National Lab

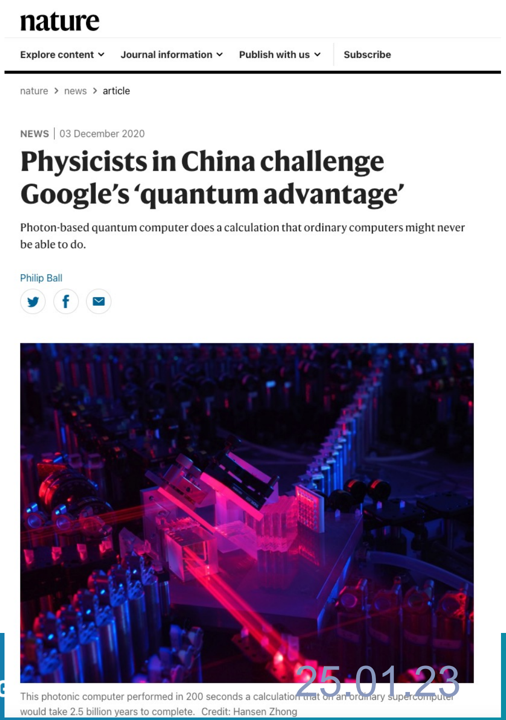

nature

Explore content ⌄ Journal information ⌄ Publish with us ⌄ Subscribe

nature > news > article

NEWS | 03 December 2020

Physicists in China challenge Google's 'quantum advantage'

Photon-based quantum computer does a calculation that ordinary computers might never be able to do.

Philip Ball

This photonic computer performed in 200 seconds a calculation that on an ordinary supercomputer would take 2.5 billion years to complete. Credit: Hansen Zhong

https://www.nature.com/articles/d41586-020-03434-7

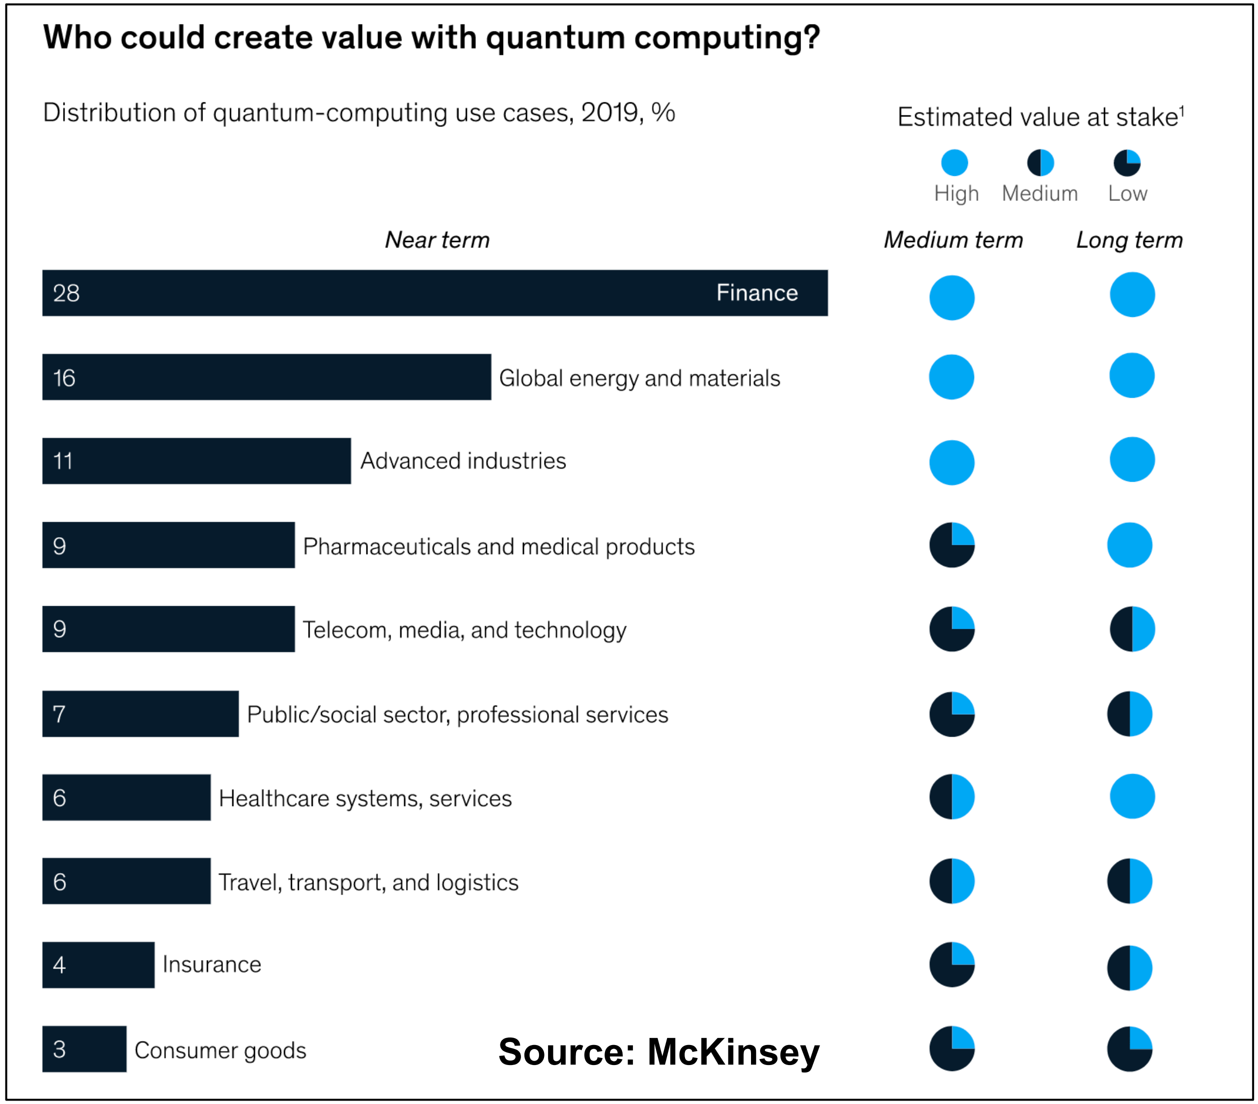

QC use cases in different sectors: the situation in 2019 with the estimated **medium** (2025) **and long** (2035) **term impact.**

Quantum Computing in High Energy Physics

- **Data** is generated by **quantum processes**
- **Quantum correlations** between particles are probed by studying **measured quantities**.

Voir en français

CERN meets quantum technology

The CERN Quantum Technology Initiative will explore the potential of devices harnessing perplexing quantum phenomena such as entanglement to enrich and expand its challenging research programme

30 SEPTEMBER, 2020 | By Matthew Chalmers

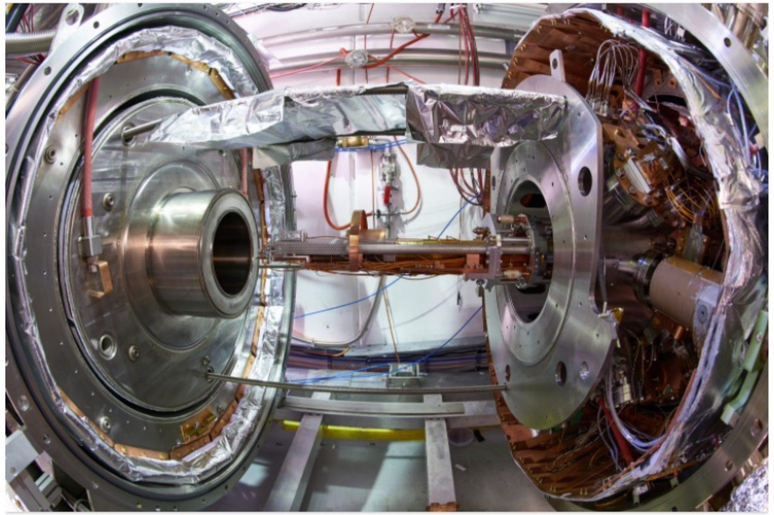

The AEgIS 1T antimatter trap stack. CERN's AEgIS experiment is able to explore the multi-particle entangled nature of photons from positronium annihilation, and is one of several examples of existing CERN research with relevance to quantum technologies. (Image: CERN)

CERN QTI Roadmap:
https://doi.org/10.5281/zenodo.5553774

Can quantum models **naturally deduce** these inherently quantum correlations ?

Quantum Machine Learning

- ML to **study quantum hardware**
- **Quantum computing to improve** ML

Quantum circuits are **differentiable**

- trained minimizing data-dependent cost function

”**Forward pass**” trained circuit

- **Encode data** in the physical state of the device
- **Take measurements**.

Classical ML/DL are **flexible** algorithms
but rely on **large data sets**

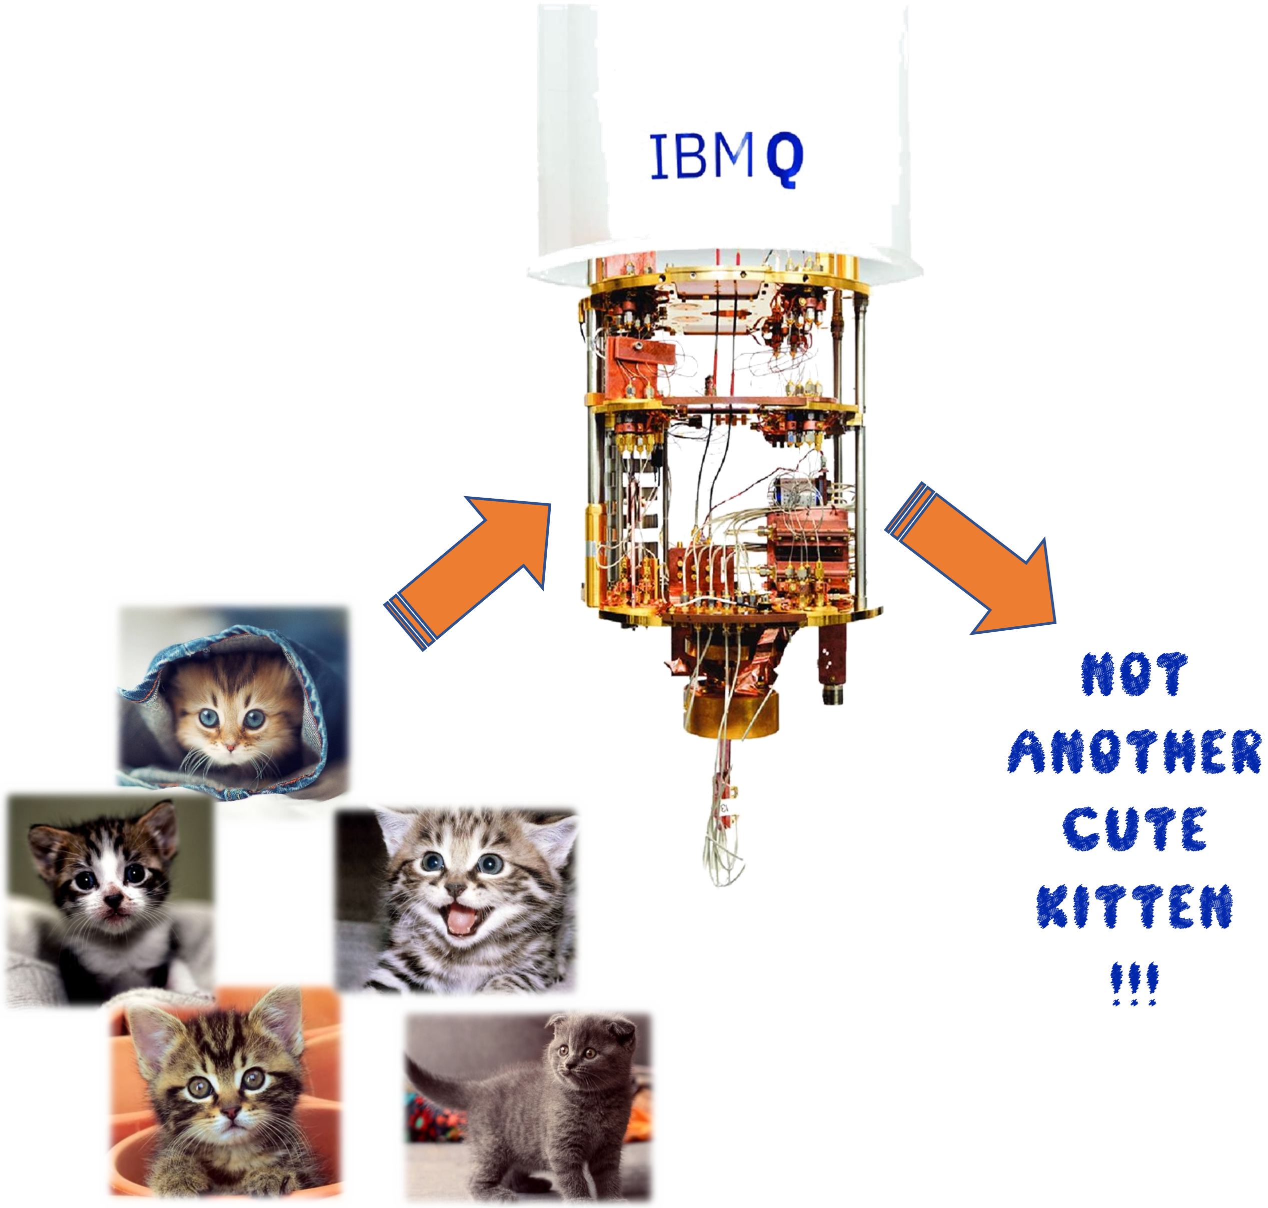

Abbas, Amira, et al. "**The power of quantum neural networks**." *Nature Computational Science* 1.6 (2021): 403-409.

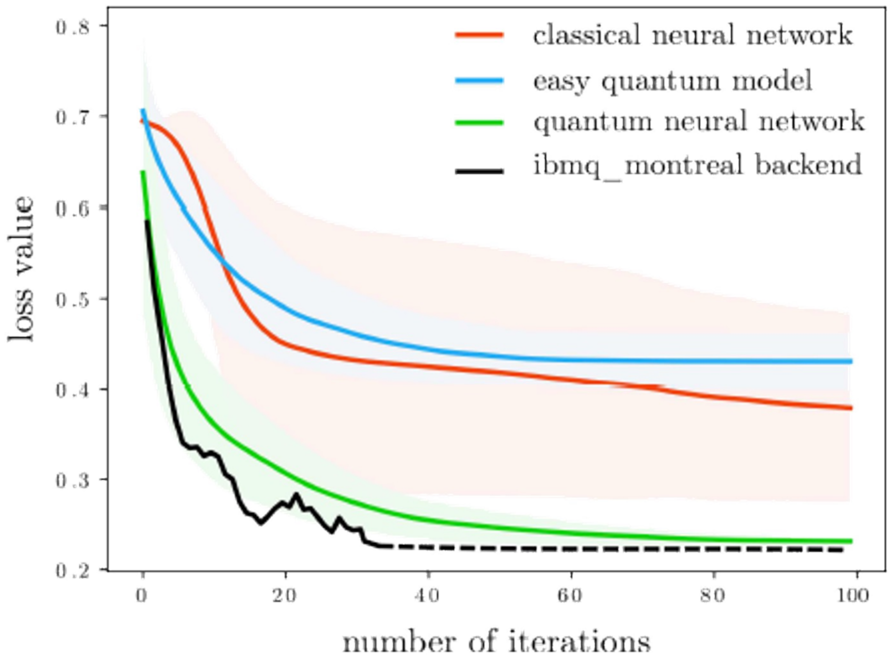

Quantum Advantage for QML

Different definitions

Runtime speedup

Sample complexity

Representational power

Classical Intractability: a quantum algorithm that cannot be efficiently simulated classically

- No established recipe for classical data
- Need to use the whole exponential advantage in Hilbert space, but will it converge ?

(Algorithm expressivity vs convergence and generalization)

Kübler, Jonas, Simon Buchholz, and Bernhard Schölkopf. "**The inductive bias of quantum kernels**." Advances in Neural Information Processing Systems 34 (2021).
Huang, HY., Broughton, M., Mohseni, M. et al. **Power of data in quantum machine learning**. Nat Commun 12, 2631 (2021). https://doi.org/10.1038/s41467-021-22539-9

QML Lifecycle

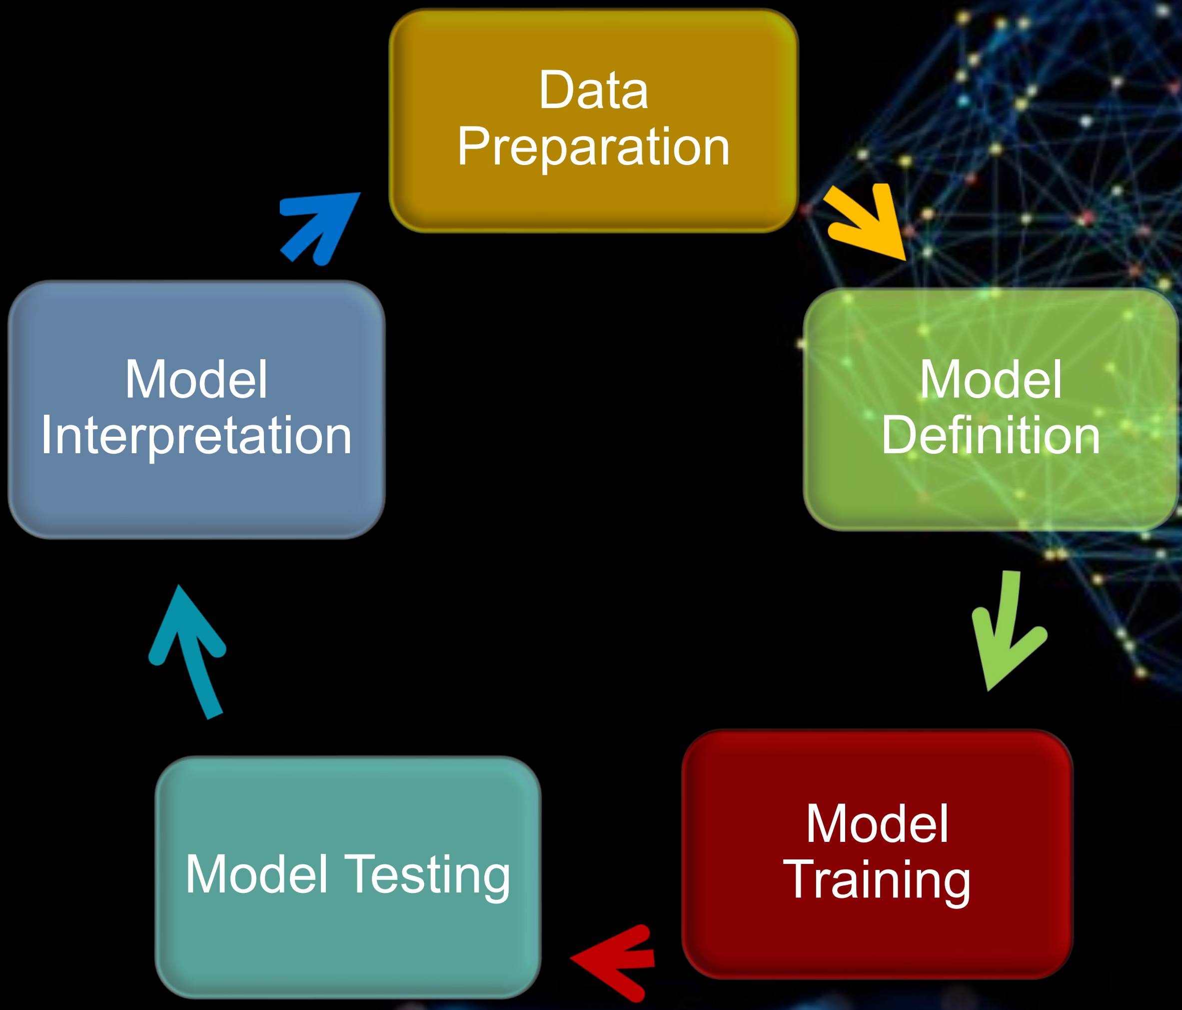

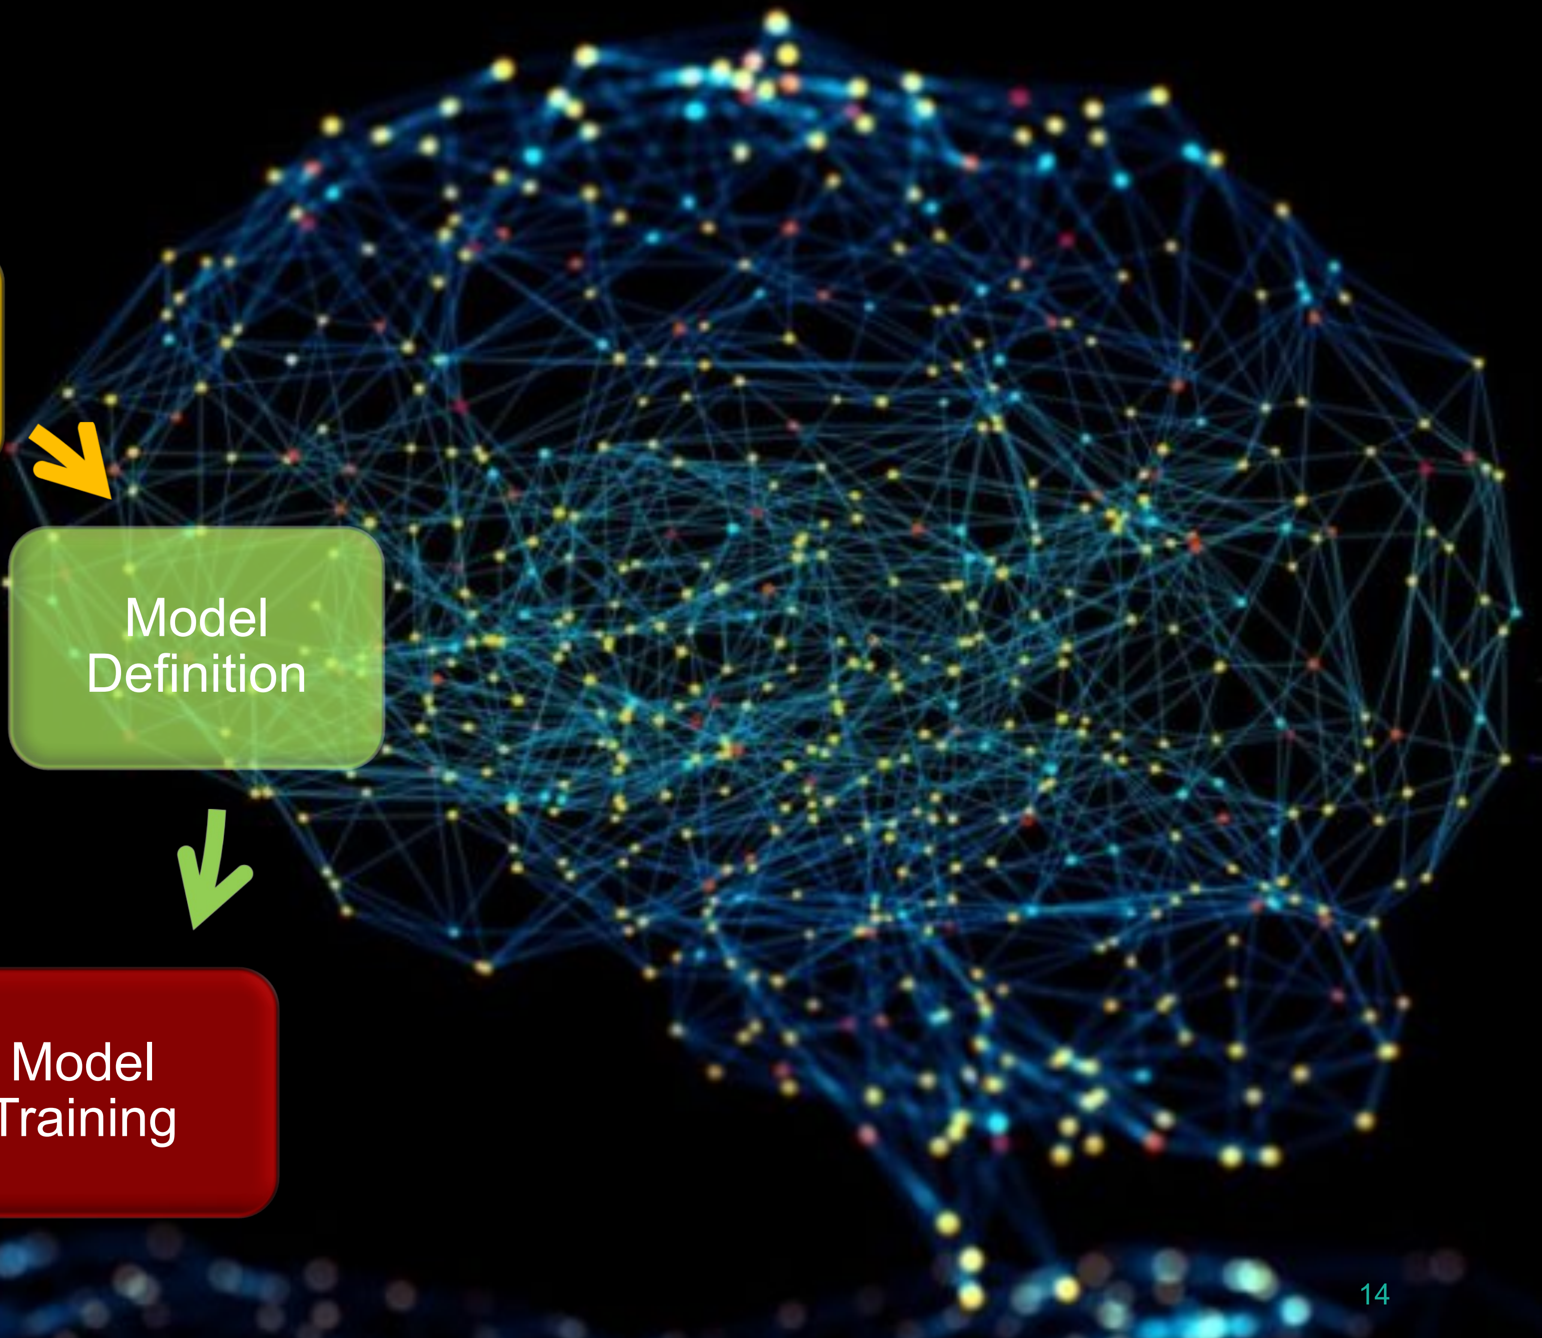

Quantum Machine Learning models

General algorithms applicable to different problems, implemented as quantum-classical hybrids, noise robust

Variational algorithms

”equivalent” of a neural network

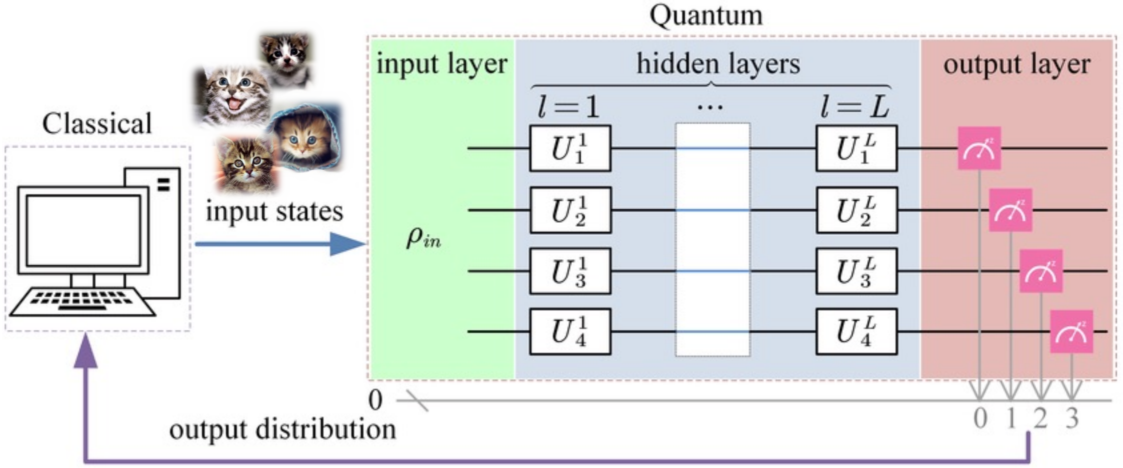

Kernel methods

kernels defined by a quantum feature map:

$$k(x_i, x) := \phi(x_i)^\dagger \cdot \phi(x)$$

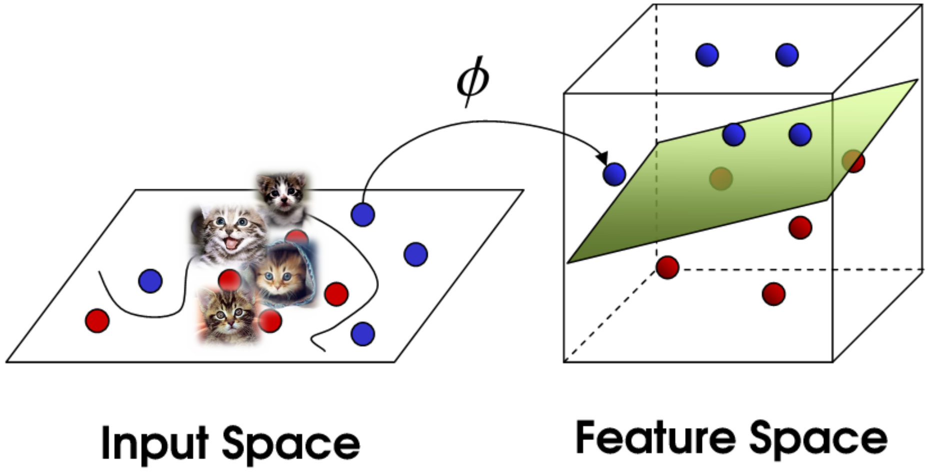

Model convergence in the quantum space

Gradient-based optimization suffers from ”barren plateaus”

Quantum NN are strongly affected by barren plateaus

Need compromise between “power” and convergence

J. McClean *et al*., arXiv:1803.11173

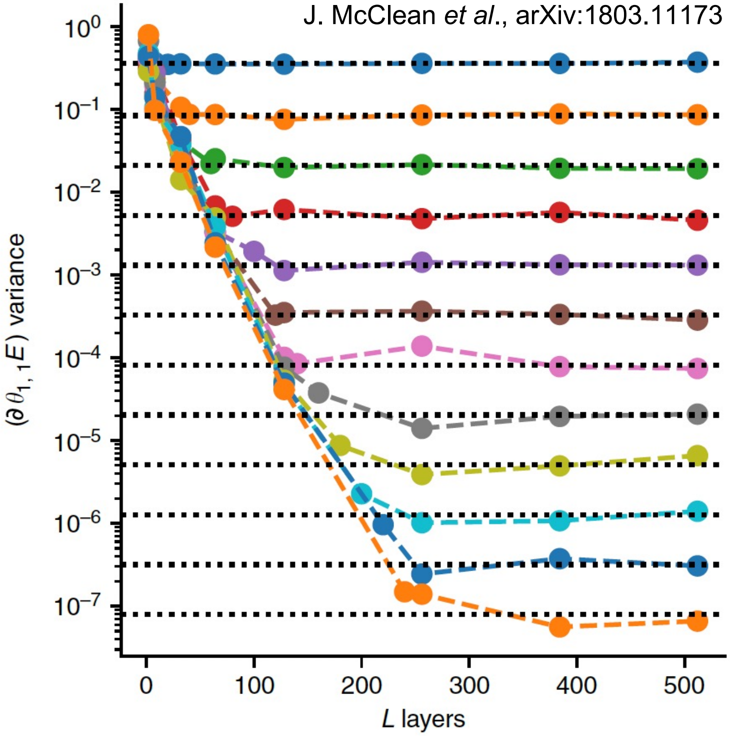

Our results so far..

- ⊕ Multiple QML prototypes for different applications
- ⊕ Increasing level of precision
- ⊕ Robustness against noise
- ⊕ Same initial hints at advantages

- ⊖ Scale is still a problem on current quantum hardware
- ⊖ Complex data pre-processing

Michael Schenk et al., **Hybrid actor-critic algorithm for quantum reinforcement learning at CERN beam lines.** arXiv:2209.11044

Quantum Reinforcement Learning

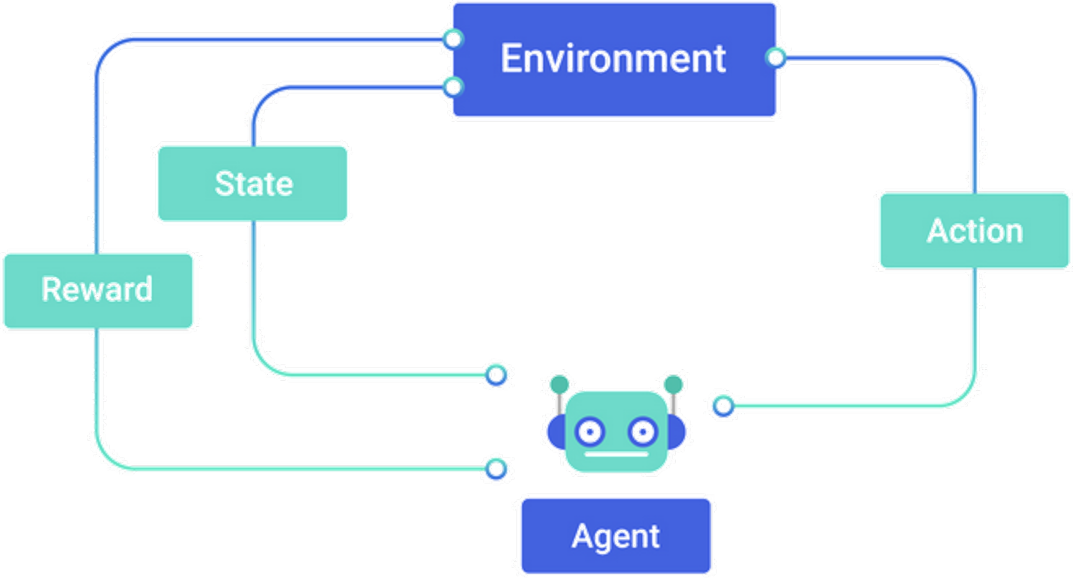

Agent interacts with environment
- Follow **policy**
- Find policy that **maximizes**

Expected reward is estimated by ***value function*** $Q(s, a)$
- **DQN:** Deep Q-learning *(NN-based)*
- **FERL:** Free energy-based RL *(clamped Quantum Boltzmann Machine)*

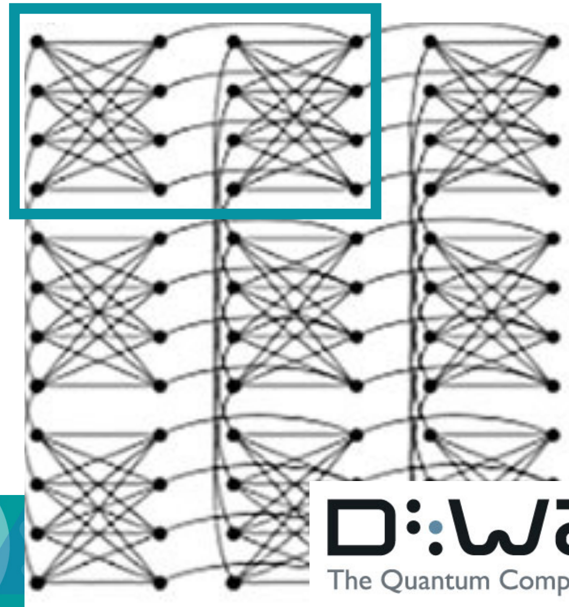

Implement the **quantum NN on a set of qubits**

Quantum computer calculates the **reward as the energy** of the qubit system

In this framework the **agent is classical**

D:wave The Quantum Computing Company™

Michael Schenk et al., **Hybrid actor-critic algorithm for quantum reinforcement learning at CERN beam lines**, e-Print: 2209.11044 [quant-ph]

Beam optimisation in linear accelerators

- **Action: (discrete) deflection** angle
- **State: (continuous) BPM** position
- **Reward: integrated beam intensity on target**
- **Optimality**: fraction of states in which the agent takes the right decision

- **Quantum RL massively** outperforms classical Q-learning *(8±2 vs. 320±40 steps with e. r.)*

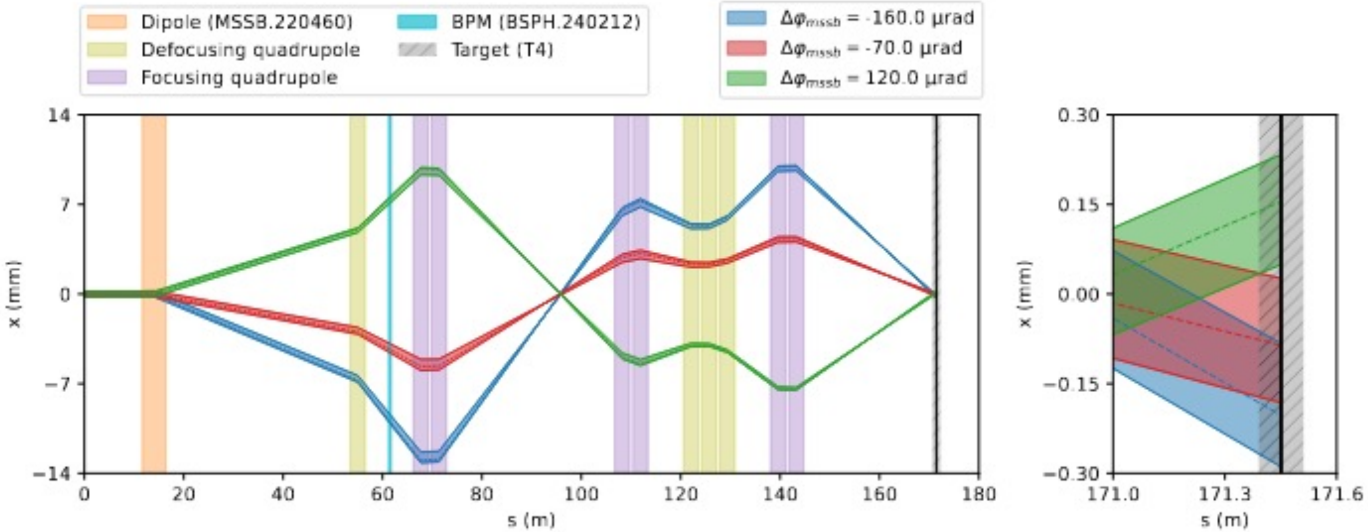

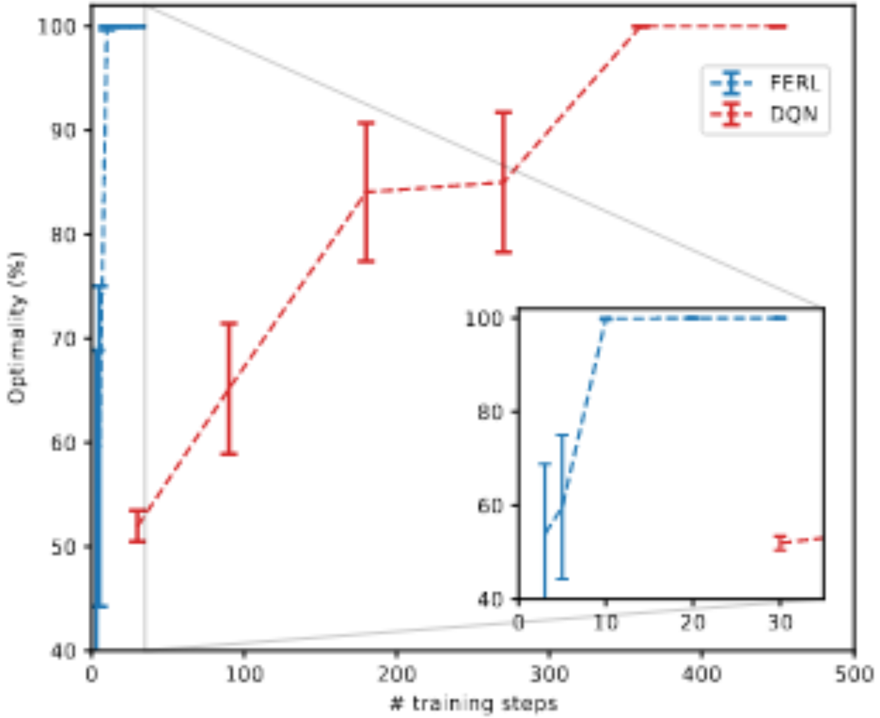

Sample complexity and representational power

Early work pointed toward possible advantage in terms of **sample complexity** and/or fast convergence

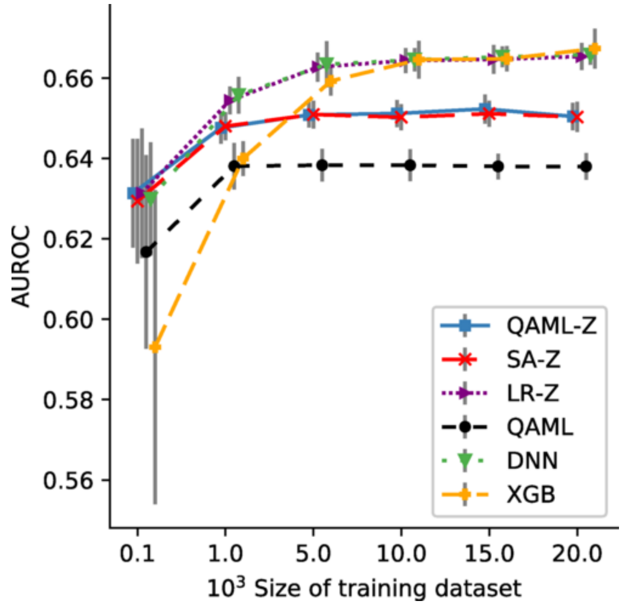

- Alexander Zlokapa, Alex Mott, Joshua Job, Jean-Roch Vlimant, Daniel Lidar, and Maria Spiropulu. **Quantum adiabatic machine learning by zooming into a region of the energy surface.** Physical Review A, 102:062405, 2020. DOI:10.1103/PhysRevA.102.062405.

QRL use cases confirms advantage in terms of **model size** and **training steps**

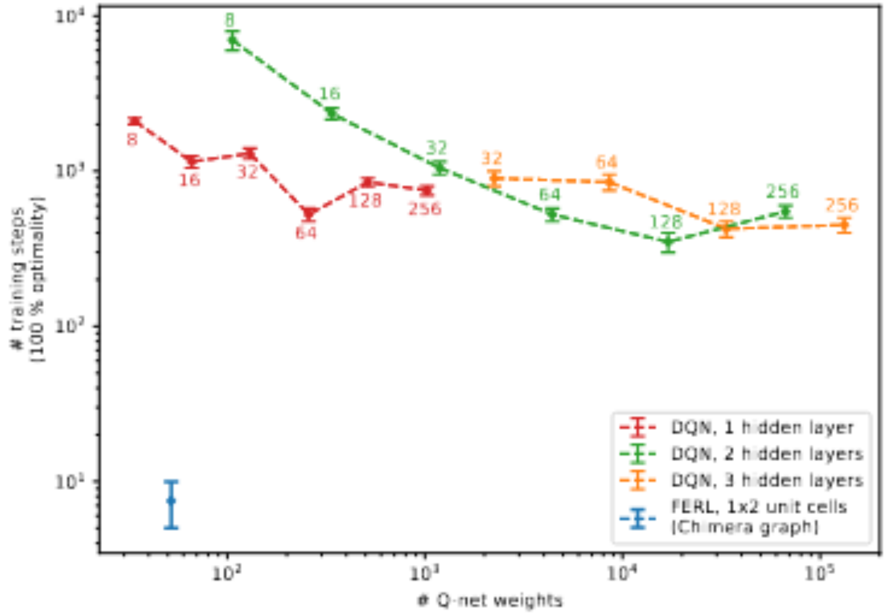

Michael Schenk, Elías F. Combarro, Michele Grossi, Verena Kain, Kevin Shing Bruce Li, Mircea-Marian Popa, Sofia Vallecorsa, **Hybrid actor-critic algorithm for quantum reinforcement learning at CERN beam lines.** arXiv:2209.11044

CERN AWAKE facility

2GeV electron beam line

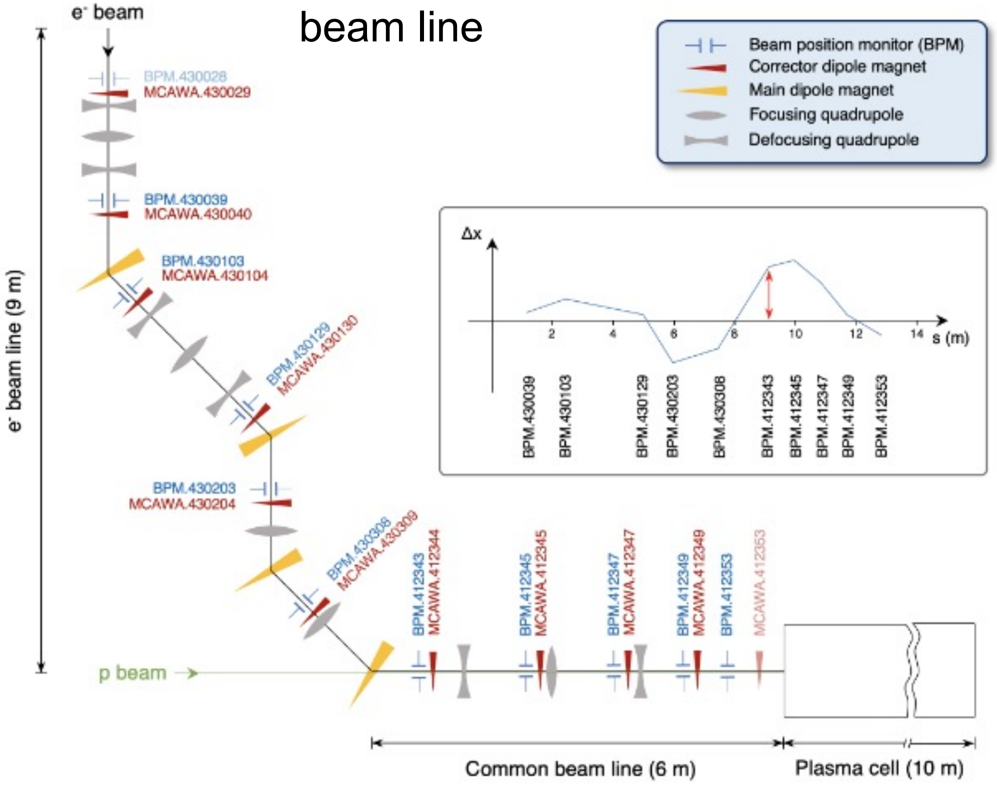

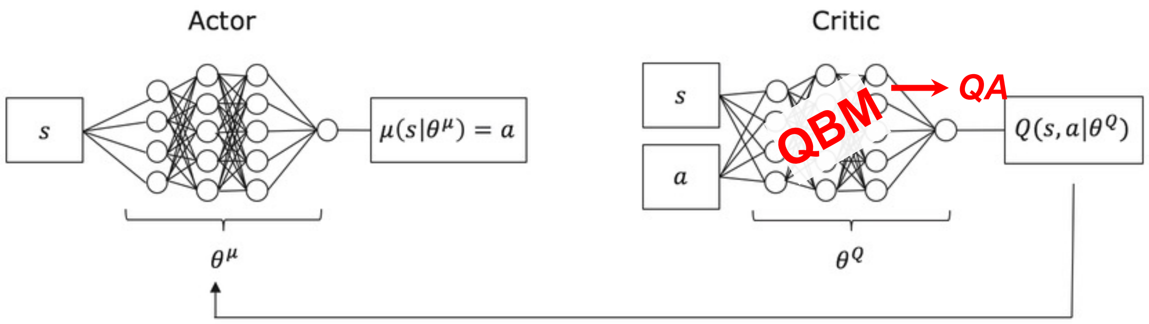

Policy Gradient: $\nabla_{\theta^\mu}\mu = \mathbb{E}_\mu[\nabla_{\theta^\mu} Q(s,\mu(s|\theta^\mu)|\theta^Q)] = \mathbb{E}_\mu[\nabla_a Q(s,a|\theta^Q) \cdot \nabla_{\theta^\mu}\mu(s|\theta^\mu)]$

Actor-Critic Q-learning training D-Wave Advantage

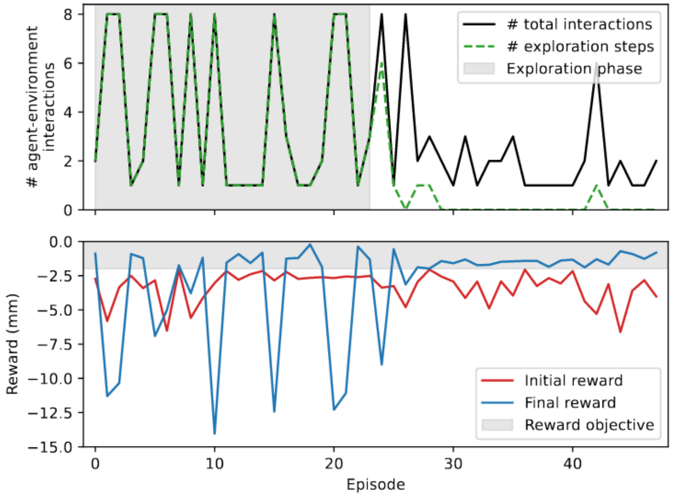

Figure 11: Single RL agent training evolution on D-Wave Advantage Systems using the simulated AWAKE environment with a reward objective of -2 mm.

Successful evaluation on the real beam-line

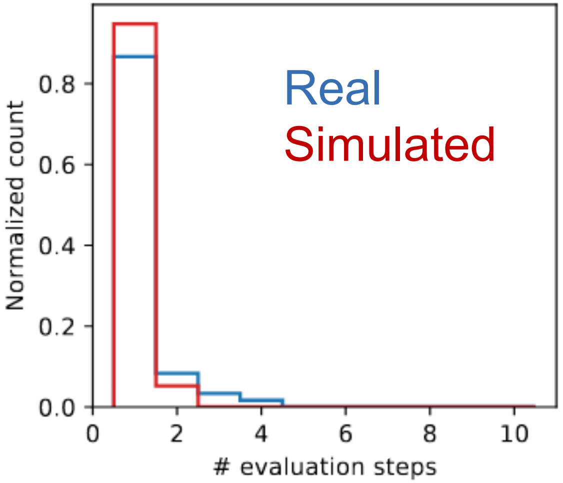

Michael Schenk et al., **Hybrid actor-critic algorithm for quantum reinforcement learning at CERN beam lines**, e-Print: 2209.11044 [quant-ph]

Saverio Monaco et al., **Quantum phase detection generalisation from marginal quantum neural network models**, arXiv:2208.08748v1.

Classifying quantum data

Generate **quantum states** directly on the device

Train QCNN to **classify quantum states**

Use marginal datasets → **OOD generalization !**

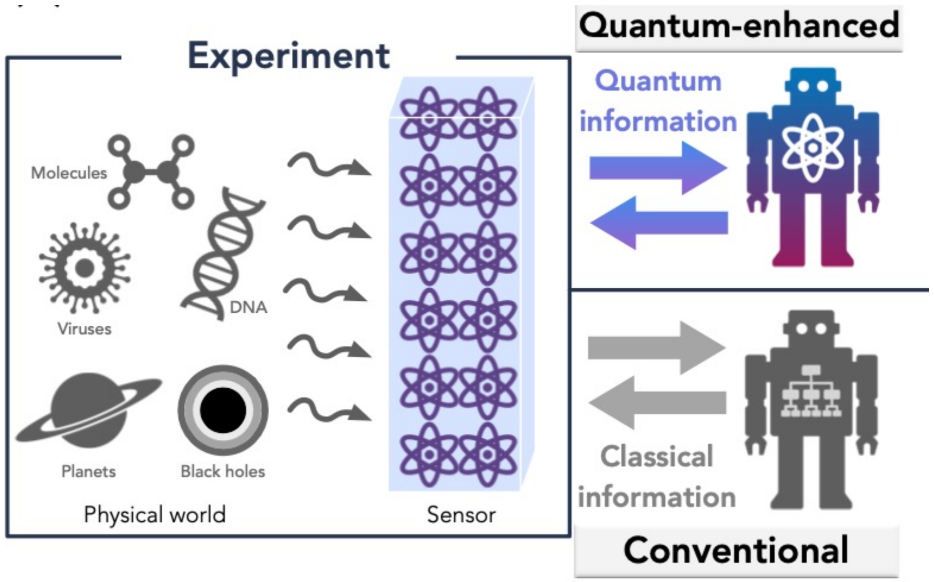

Out of Distribution Generalization

M..Caro et al., Out-of-distribution generalization for learning quantum dynamics, arxiv:2204.10268

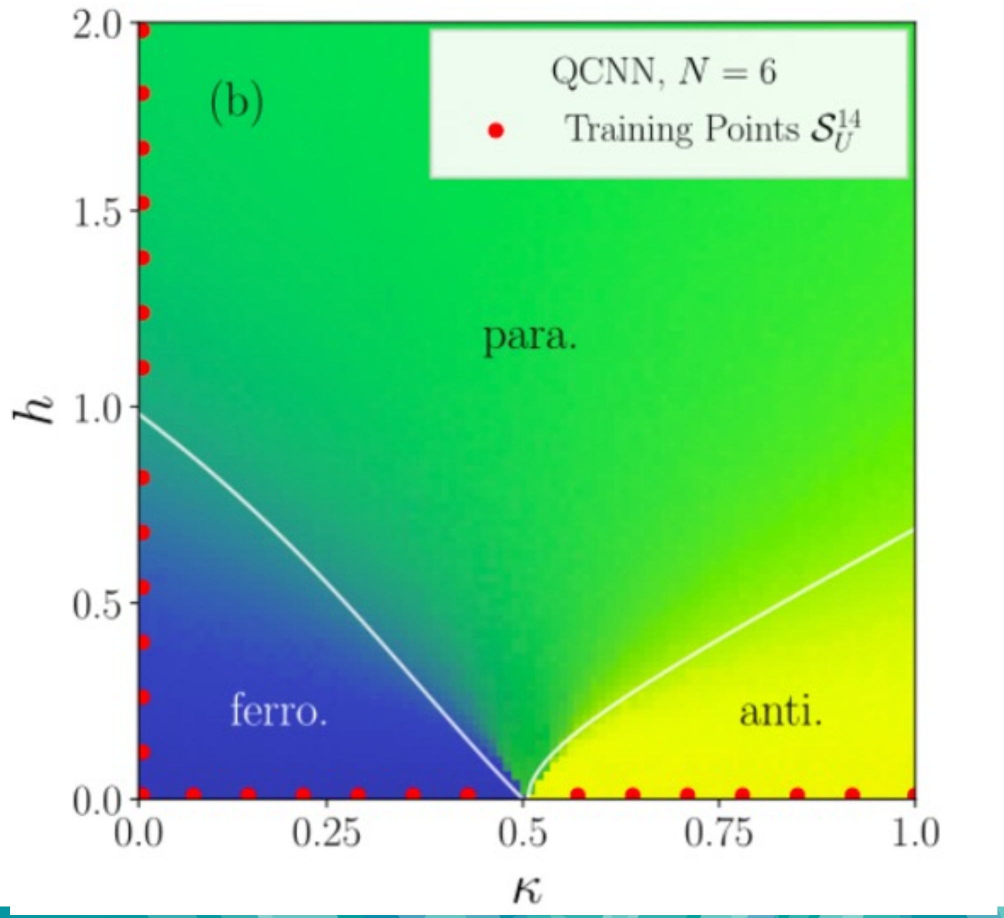

Quantum Generative Models

Delgado and Hamilton, arXiv:2203.03578 (2022)
Zoufal, et al., *npj Quantum Inf* **5**, 103 (2019)
Leadbeater et al., *Entropy* **2021**, *23*, 1281.
Amin, et al. *Physical Review X* 8.2 (2018): 021050.

QCBM

Sample variational pure state $|\psi(\theta)\rangle$ by projective measurement through **Born rule**: $\mathbf{p_\theta}(\mathbf{x}) = |\langle \mathbf{x}|\boldsymbol{\psi}(\boldsymbol{\theta})\rangle|^2$.

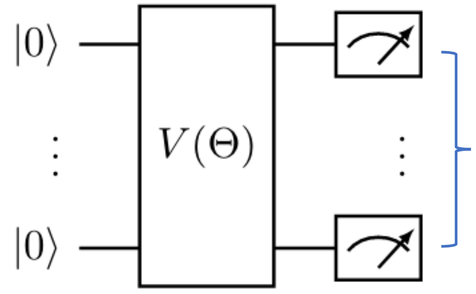

QGAN

Multiple implementations, mostly classical-quantum hybrid

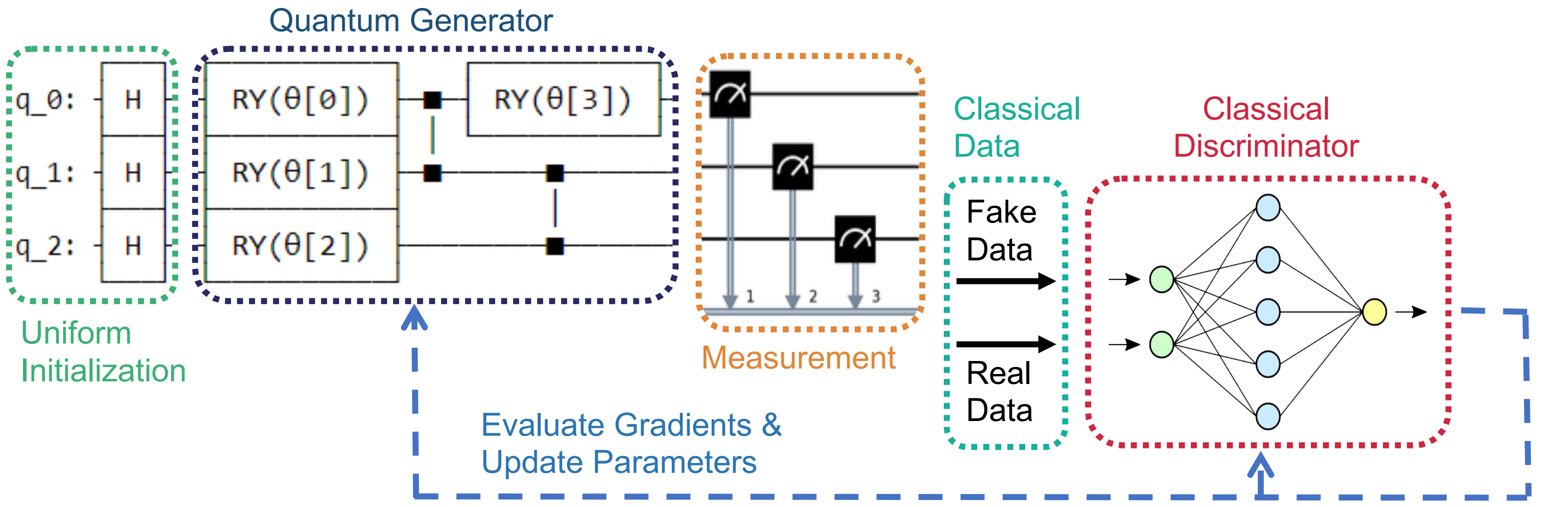

QBM

Network of stochastic binary units with a quadratic energy function that follows the Boltzman distribution (Ising Hamiltonian)

$$H = -\sum_a b_a \sigma_a^z - \sum_{a,b} w_{ab}\sigma_a^z\sigma_b^z$$

Typical metrics:

$$D_{\mathrm{KL}}(P\|Q) = \sum_i P(i)\,\log\left(\frac{P(i)}{Q(i)}\right)$$

$$\mathrm{MMD}(\mathbb{P}_r,\mathbb{P}_g) = \left(\mathbb{E}_{\substack{\mathbf{x}_r,\mathbf{x}'_r\sim\mathbb{P}_r,\\ \mathbf{x}_g,\mathbf{x}'_g\sim\mathbb{P}_g}}\left[k(\mathbf{x}_r,\mathbf{x}'_r) - 2k(\mathbf{x}_r,\mathbf{x}_g) + k(\mathbf{x}_g,\mathbf{x}'_g)\right]\right)^{\frac{1}{2}}$$

QCBM for event generation

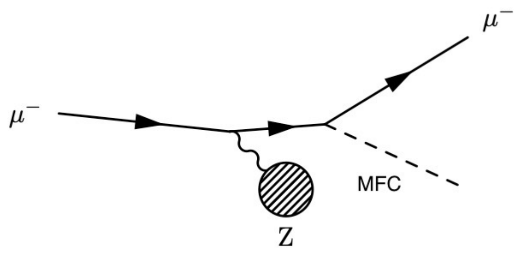

Muon Force Carriers, in muon fixed-target experiments or muon interactions in calorimeters[1]

Generate multivariate distribution (E, p_t, η)

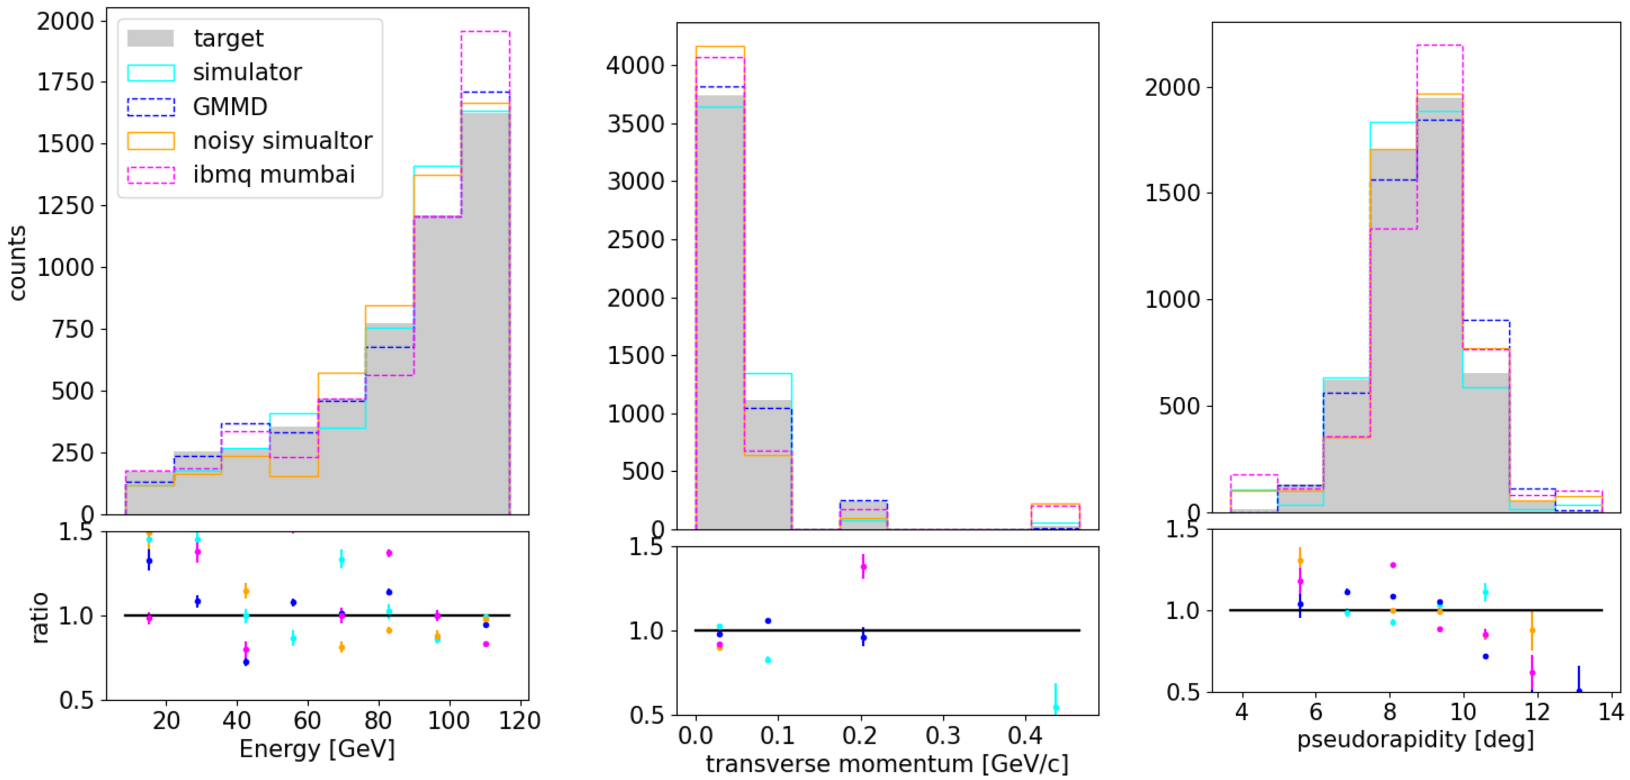

\+ Implement **conditional** $p(y|x)$ wrt incoming particle energy E_{in}.

1 Galon, I, Kajamovitz, E et al. "Searching for muonic forces with the ATLAS detector". In: Phys. Rev. D 101, 011701 (2020)

Realistic performance

Borras, Kerstin, et al. "**Impact of quantum noise on the training of quantum Generative Adversarial Networks**." submitted to QTML2022

QML can simulate the energy deposited by particles in a detector

Ex. Hybrid quantum GAN model

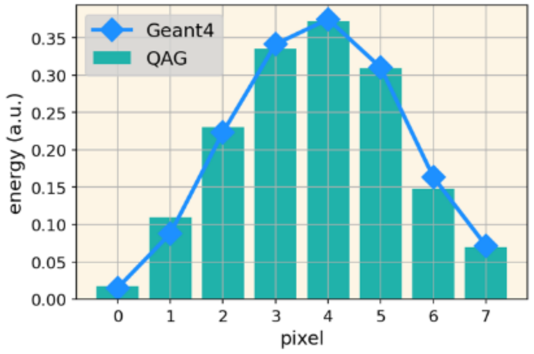

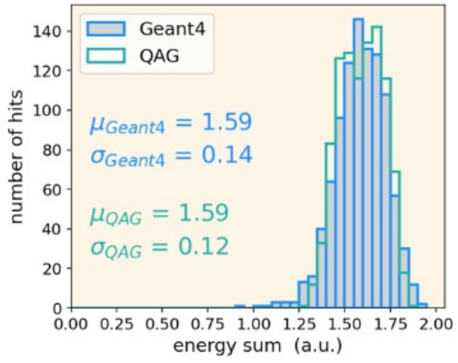

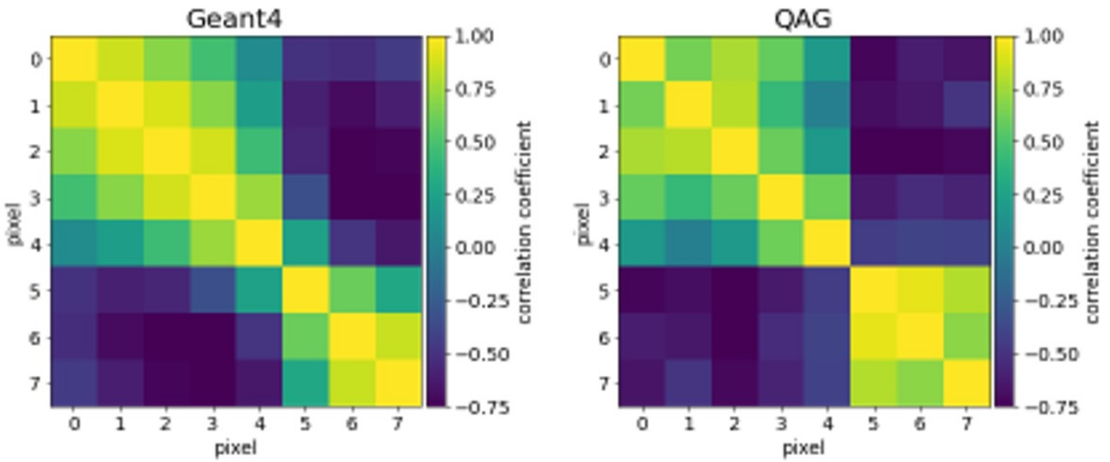

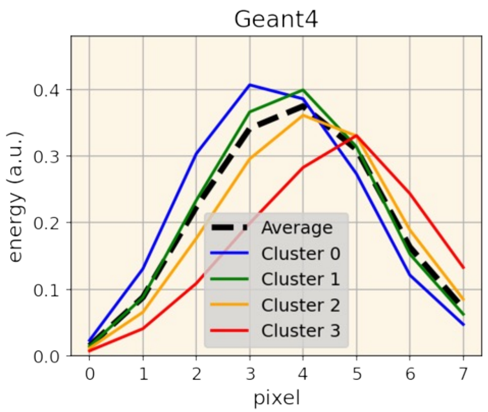

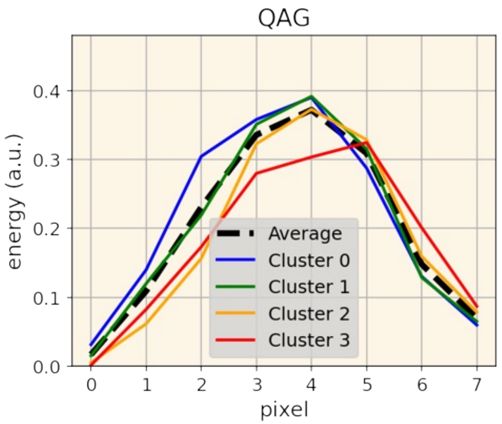

Scale is still a problem for some use cases (ex. Detector simulation)

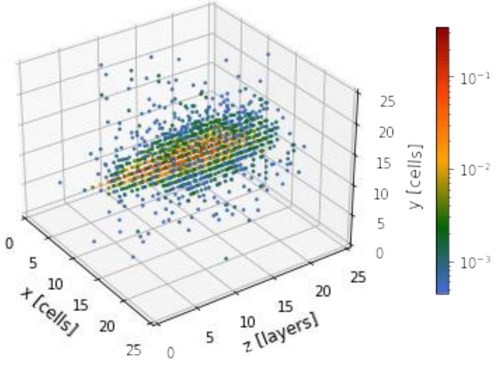

Perspective

Quantum Machine Learning is a broad-lively research field

- Some **preliminary hints** of advantage
- Need more **robust theoretical studies** to interpret experimental results
- Need to establish **«fair comparison»** to classical tools

CERN is formulating a **longer term research plan dedicated to understanding impact for physics**

- Identify cases where quantum approach could be **more effective** than classical algorithms**...**
- Study performance **beyond near-term hardware**
- …

- CERN has joined proposition of an **Open Quantum Institute.**
- Quantum technologies for key societal challenges
- Proposal made through **GESDA**, the Geneva Science and Diplomacy Anticipator Foundation
- **UniGe, ETH, EPFL, Microsoft and IBM** are among supporters

QUANTUM MACHINE LEARNING

2023

Thanks!

Sofia.Vallecorsa@cern.ch

CERN

QUANTUM
TECHNOLOGY
INITIATIVE