Quantum Machine Learning

Leonardo Banchi

University of Florence

1 Oct 2019

About me

- RTD-B @ UniFi (Rita Levi Montalcini)
- INFN Projects: SFT, ML

About me

- RTD-B @ UniFi (Rita Levi Montalcini)
- INFN Projects: SFT, ML

Research Interests

- Quantum Information Theory: Quantum Computation & Quantum Communication
- Quantum Many-Particle Systems
- Machine Learning

Data \ Algorithm	classical	quantum
classical	CC	CQ
quantum	QC	QQ

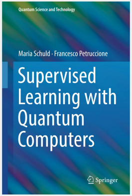

Quantum Machine Learning

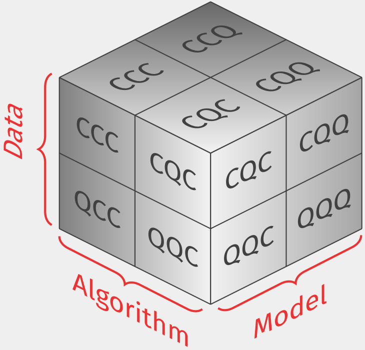

Quantum Machine Learning

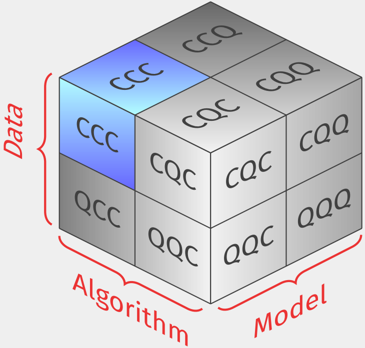

CL. DATA, CL. ALGORITHM, CL. MODEL

ML for data analysis, data interpolation, stochastic optimization

- Deep Neural Networks: TensorFlow, Julia (KNet, Flux)

Cl. Data, Cl. Algorithm, Cl. Model

ML for data analysis, data interpolation, stochastic optimization

- Deep Neural Networks: TensorFlow, Julia (KNet, Flux)
- Random Forests

Cl. Data, Cl. Algorithm, Cl. Model

ML for data analysis, data interpolation, stochastic optimization

- Deep Neural Networks: TensorFlow, Julia (KNet, Flux)
- Random Forests
- XGBoost

CL. DATA, CL. ALGORITHM, CL. MODEL

ML for data analysis, data interpolation, stochastic optimization

- Deep Neural Networks: TensorFlow, Julia (KNet, Flux)
- Random Forests
- XGBoost
- Gaussian Process Regression

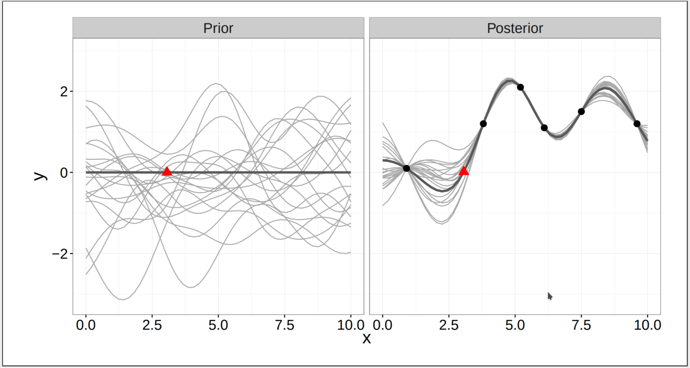

Opening the black box: why ML works?

Information Theory

- Coding theory: best theoretical (de)compression strategies
- Amount of information quantified by the entropy H

Information Theory

- Coding theory: best theoretical (de)compression strategies
- Amount of information quantified by the entropy H

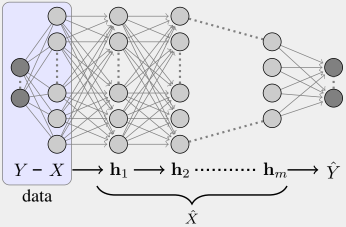

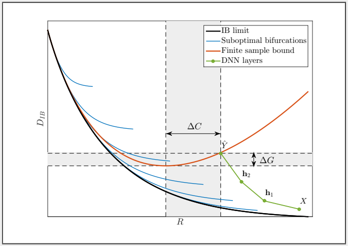

Theoretical limit (supervised learning)

residual information between X and Y not captured by $\hat{X}$

$$D_{\mathrm{IB}} = H(X, Y|\hat{X})$$

Opening the black box: why ML works?

Restricted Boltzmann Machines (unsupervised learning)

$$p(\{s_j\}) \propto e^{-\beta\left[\sum_{ij} J_{ij} s_i s_j + \sum_i h_i s_i\right]}$$

Learning as following the renormalization group flow

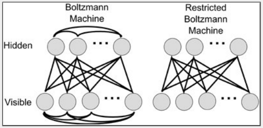

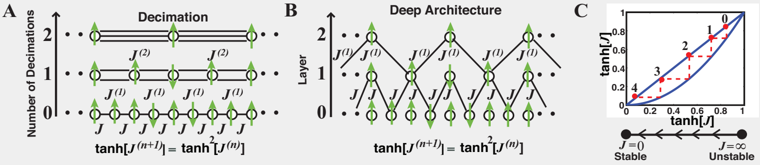

Quantum Machine Learning

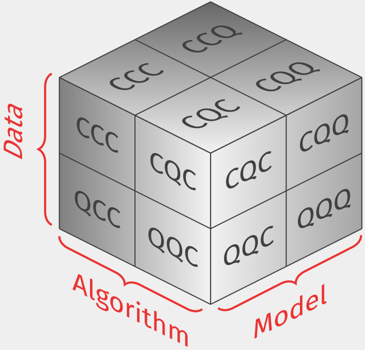

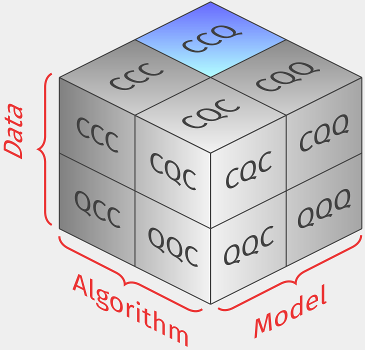
CCQ
CCC
CQQ
CQC
CCC
CQC
CQC
CQQ
QCC
QQC
QQC
QQQ
Data
Algorithm
Model

Problem in ML

effectively model a complex mapping between huge dimensional spaces using few trainable parameters

Problem in ML

effectively model a complex mapping between huge dimensional spaces using few trainable parameters

Analogy: variational methods for quantum many-body systems

$$E_{\mathrm{GS}} \lesssim \min_{\theta} \langle \psi(\theta) | H | \psi(\theta) \rangle$$

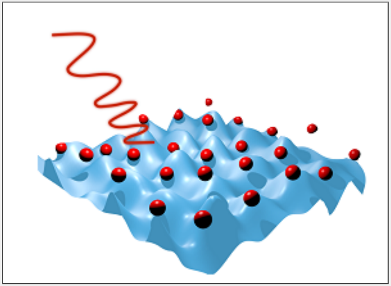

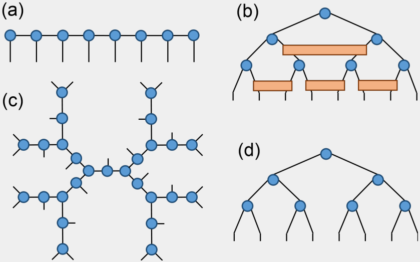

Quantum Machine Learning

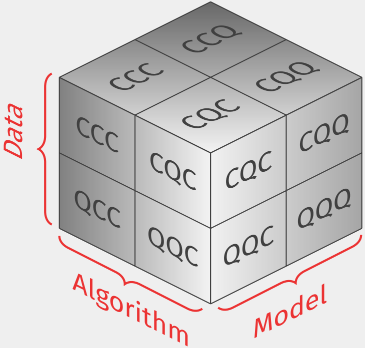

Quantum Machine Learning

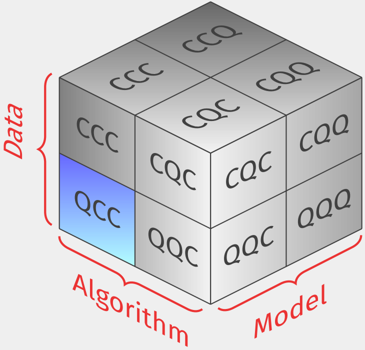

Q. Data, Cl. Algorithm, Cl. Model

ML as a tool to probe and control Quantum Systems

Example applications:

- Extract relevant features of a quantum system from carefully chosen measurements
- Quantum theory is probabilistic: optimising the dynamics of a quantum system requires stochastic optimization (ML)
- Prediction of quantum dynamics via recurrent neural networks
- Optimal tomography via gaussian process regression
- Quantum Control via reinforcement learning

Quantum Machine Learning

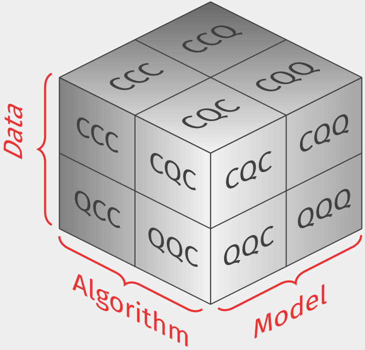

Quantum Computing

"Near-Term" Quantum Devices

- Available nowadays (Google, IBM, etc.)
- Circuit based vs sampling devices (D-Wave, boson samplers)

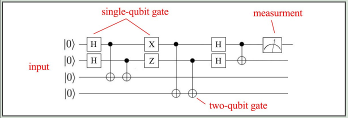

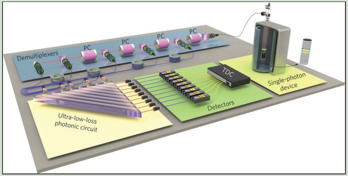

"Full" Quantum Computer

- Can run proper quantum algoritms (Shor, etc.)
- Requires error correction ($10^4 - 10^5$ qubits)
- Up to **exponential speedups** for solving certain problems

Quantum Machine Learning

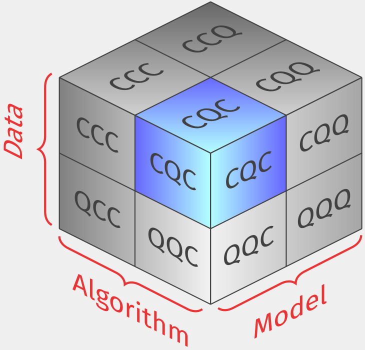

Quantum Neural Networks

classical trainable parameters $\{\theta_j\}$

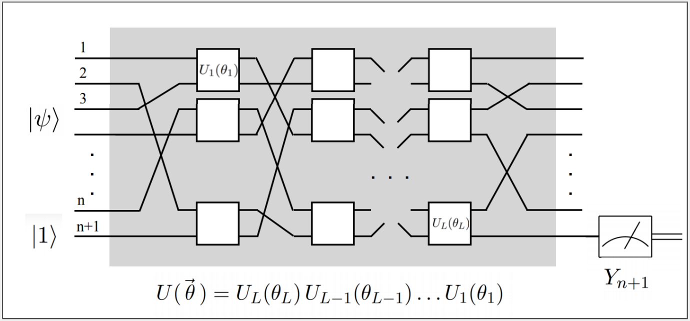

✓ Ok for near-term devices

✗ Possible speedups not known

Unsupervised learning with quantum samplers

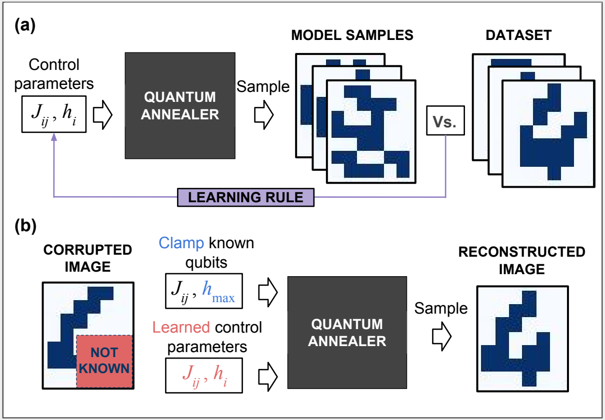

✓ Ok for near-term devices
✗ Speedups might be expected, but not well justified

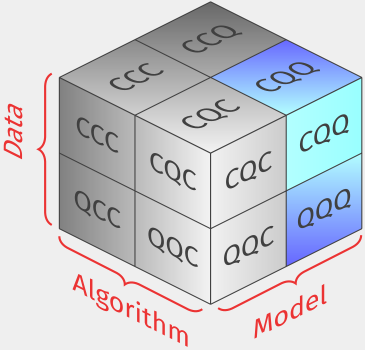
CCQ
CCC
CQQ
CQC
CCC
CQC
CQC
CQQ
QCC
QQC
QQC
QQQ
Data
Algorithm
Model

HHL algorithm

A quantum computer can solve a linear system

$$y = Ax$$

using $\mathcal{O}(\log(N))$ operations. This is an **exponential speedup**

HHL algorithm

A quantum computer can solve a linear system

$$y = Ax$$

using $\mathcal{O}(\log(N))$ operations. This is an **exponential speedup**

Many applications in ML

- Quantum Principal Component Analysis
- Quantum Support Vector Machines
- Quantum Recommendation Systems ("Netflix problem")

HHL algorithm

A quantum computer can solve a linear system

$$y = Ax$$

using $\mathcal{O}(\log(N))$ operations. This is an **exponential speedup**

Many applications in ML

- Quantum Principal Component Analysis
- Quantum Support Vector Machines
- Quantum Recommendation Systems (“Netflix problem”)

Summary:

✓ A quantum computer can speedup ML algorithms

✗ Too difficult for near-term devices

Summary

- Physical theories (information theory, RG-flow, etc.) can help us to understand why ML works
- ML as a tool to probe and optimize quantum systems
- Quantum correlations as “inspiration” to define new ML models
- Quantum sampling devices for classically hard distributions
- Quantum algorithms provide exponential speedups to some ML problems

Summary

- Physical theories (information theory, RG-flow, etc.) can help us to understand why ML works
- ML as a tool to probe and optimize quantum systems
- Quantum correlations as "inspiration" to define new ML models
- Quantum sampling devices for classically hard distributions
- Quantum algorithms provide exponential speedups to some ML problems

Thank you for your attention!