



Introduction to Hybrid Quantum-Classical Machine Learning with Applications

Presenter: Samuel Yen-Chi Chen

Quantum Computing and Quantum Machine Learning Algorithms for High Energy Physics Workshop

Date 2022-06-01



- I. Introduction**
- II. Quantum Computing (QC)**
- III. Quantum Machine Learning**
- IV. Applications-Classification**
- V. Applications-Sequential Learning**
- VI. Applications-Reinforcement Learning**
- VII. Conclusion and Outlook**

I. Introduction

II. Quantum Computing (QC)

III. Quantum Machine Learning

IV. Applications-Classification

V. Applications-Sequential Learning

VI. Applications-Reinforcement Learning

VII. Conclusion and Outlook

2016 AlphaGo



ARTICLE

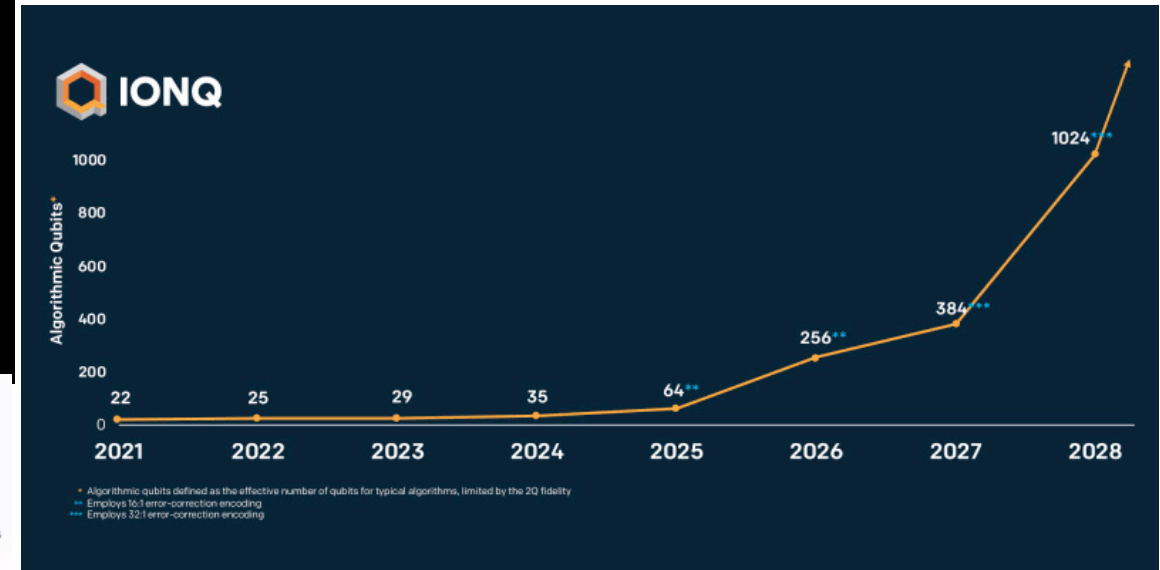
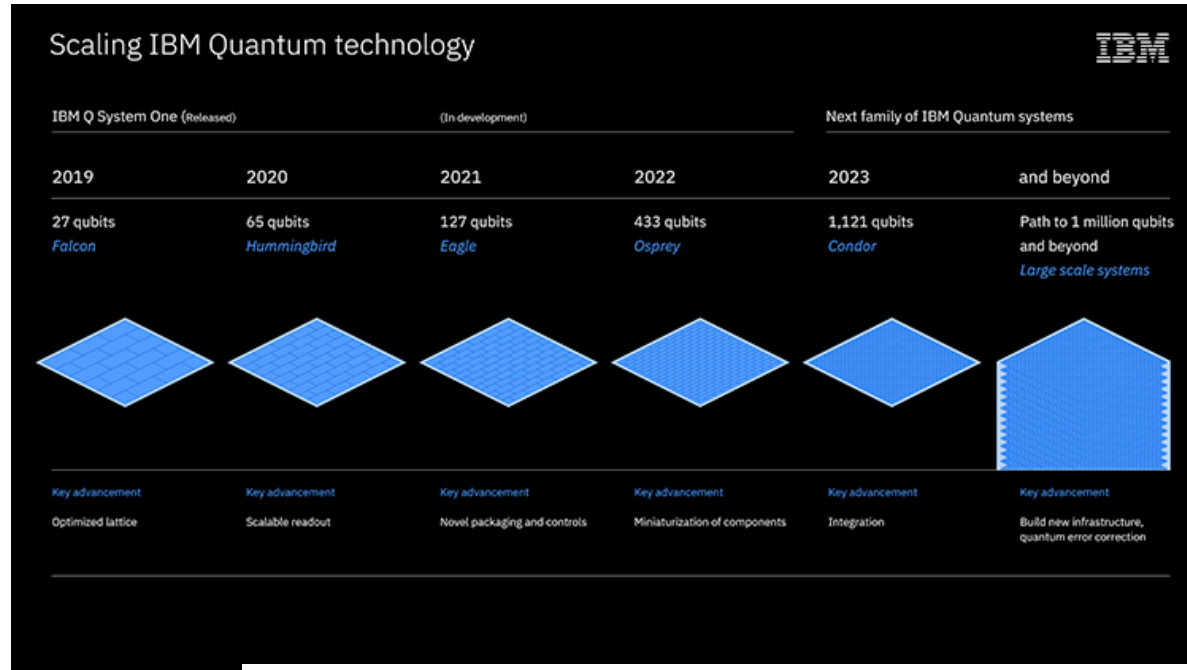
doi:10.1038/nature16961

Mastering the game of Go with deep neural networks and tree search

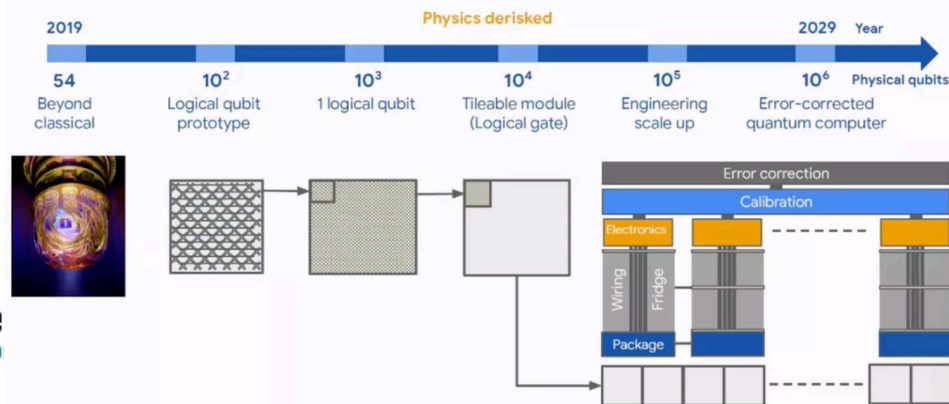
David Silver^{1*}, Aja Huang^{1*}, Chris J. Maddison¹, Arthur Guez¹, Laurent Sifre¹, George van den Driessche¹, Julian Schrittwieser¹, Ioannis Antonoglou¹, Veda Panneershelvam¹, Marc Lanctot¹, Sander Dieleman¹, Dominik Grewe¹, John Nham², Nal Kalchbrenner¹, Ilya Sutskever², Timothy Lillicrap¹, Madeleine Leach¹, Koray Kavukcuoglu¹, Thore Graepel¹ & Demis Hassabis¹

The game of Go has long been viewed as the most challenging of classic games for artificial intelligence owing to its enormous search space and the difficulty of evaluating board positions and moves. Here we introduce a new approach to computer Go that uses ‘value networks’ to evaluate board positions and ‘policy networks’ to select moves. These deep neural networks are trained by a novel combination of supervised learning from human expert games, and reinforcement learning from games of self-play. Without any lookahead search, the neural networks play Go at the level of state-of-the-art Monte Carlo tree search programs that simulate thousands of random games of self-play. We also introduce a new search algorithm that combines Monte Carlo simulation with value and policy networks. Using this search algorithm, our program AlphaGo achieved a 99.8% winning rate against other Go programs, and defeated the human European Go champion by 5 games to 0. This is the first time that a computer program has defeated a human professional player in the full-sized game of Go, a feat previously thought to be at least a decade away.

2017~ QC hardware



Google AI Quantum hardware roadmap



I. Introduction

II. Quantum Computing (QC)

III. Quantum Machine Learning

IV. Applications-Classification

V. Applications-Sequential Learning

VI. Applications-Reinforcement Learning

VII. Conclusion and Outlook

Quantum Computing

- Classical computers: Classical bits 0 vs 1
- Quantum computers: Quantum bits (qubit) $|\Psi\rangle = \alpha|0\rangle + \beta|1\rangle$ where α and β are complex numbers \mathbb{C}
- Quantum **entanglements**: A unique property of quantum physics \rightarrow No analog in the classical computer
- Famous algorithms:
 - Shor's algorithm: Can be used to break the state-of-the-art public key cryptography systems such as RSA
 - Grover's algorithm: Quadratic speedup in unstructured search

- Designing a quantum algorithm is non-trivial task
- Even harder in the noisy quantum machines

Quantum States

Single Qubit State

$$|0\rangle = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$$

$$|1\rangle = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$$

Two Qubit State

$$|0\rangle \otimes |0\rangle = \begin{bmatrix} 1 \\ 0 \end{bmatrix} \otimes \begin{bmatrix} 1 \\ 0 \end{bmatrix} = \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \end{bmatrix}$$

N Qubit State

$$\underbrace{|0\rangle \otimes |0\rangle \otimes \dots \otimes |0\rangle}_N = \underbrace{\begin{bmatrix} 1 \\ 0 \end{bmatrix} \otimes \begin{bmatrix} 1 \\ 0 \end{bmatrix} \otimes \dots \otimes \begin{bmatrix} 1 \\ 0 \end{bmatrix}}_N$$

Quantum States

Density Operators

$$\rho = \sum_j p_j |\psi_j\rangle \langle \psi_j|$$

$|\psi_j\rangle$ **Basis state**

p_j **Probability**

Examples:

$$|0\rangle = \begin{bmatrix} 1 \\ 0 \end{bmatrix} \quad \langle 0| = [1 \quad 0]$$



$$|0\rangle\langle 0| = \begin{bmatrix} 1 \\ 0 \end{bmatrix} [1 \quad 0] = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$$

Quantum Operations

$$\begin{aligned} \boxed{X} & \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \\ \boxed{Y} & \begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix} \\ \boxed{Z} & \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} \\ \boxed{H} & \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \end{aligned}$$

Example:

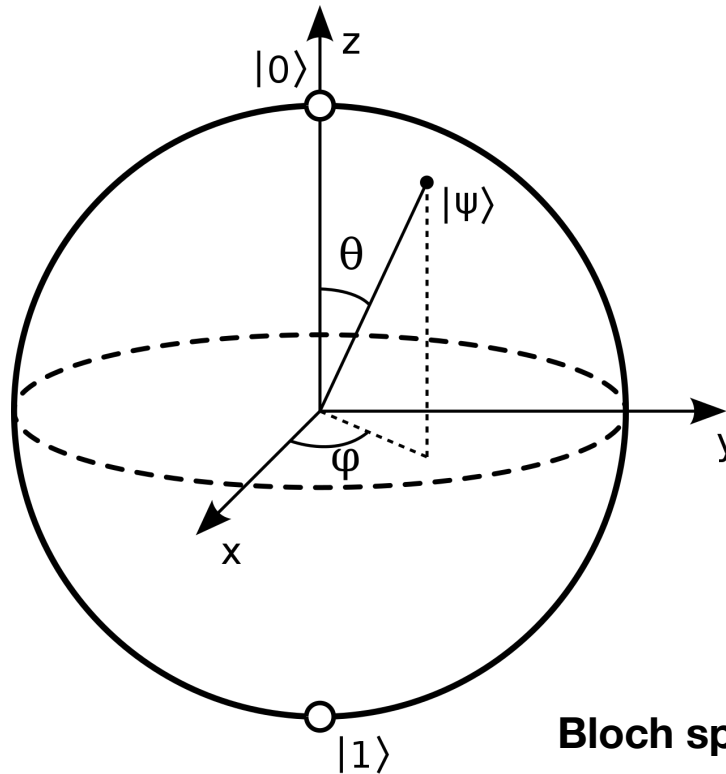
$$|0\rangle \longrightarrow \boxed{X}$$

$$\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} 1 \\ 0 \end{bmatrix} = \begin{bmatrix} 0 \\ 1 \end{bmatrix} = |1\rangle$$

Quantum Operations

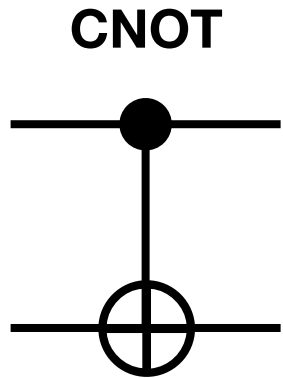
$$R(\phi, \theta, \omega)$$

$$\begin{bmatrix} e^{-i(\phi+\omega)/2} \cos(\theta/2) & e^{-i(\phi-\omega)/2} \sin(\theta/2) \\ e^{-i(\phi-\omega)/2} \sin(\theta/2) & e^{i(\phi+\omega)/2} \cos(\theta/2) \end{bmatrix}$$

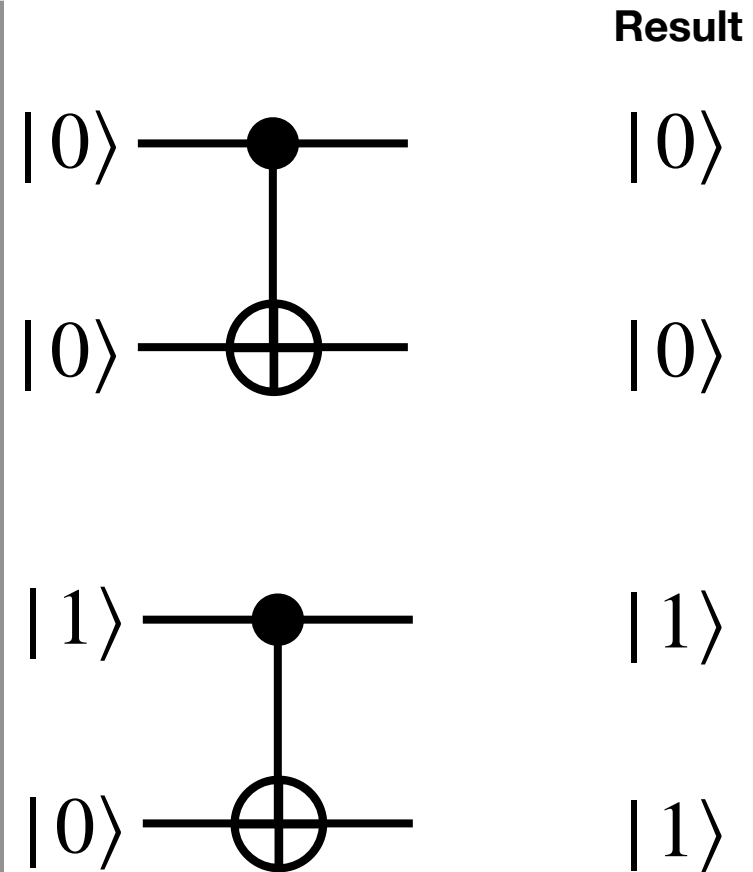


Bloch sphere

Quantum Operations



$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix}$$



I. Introduction

II. Quantum Computing (QC)

III. Quantum Machine Learning

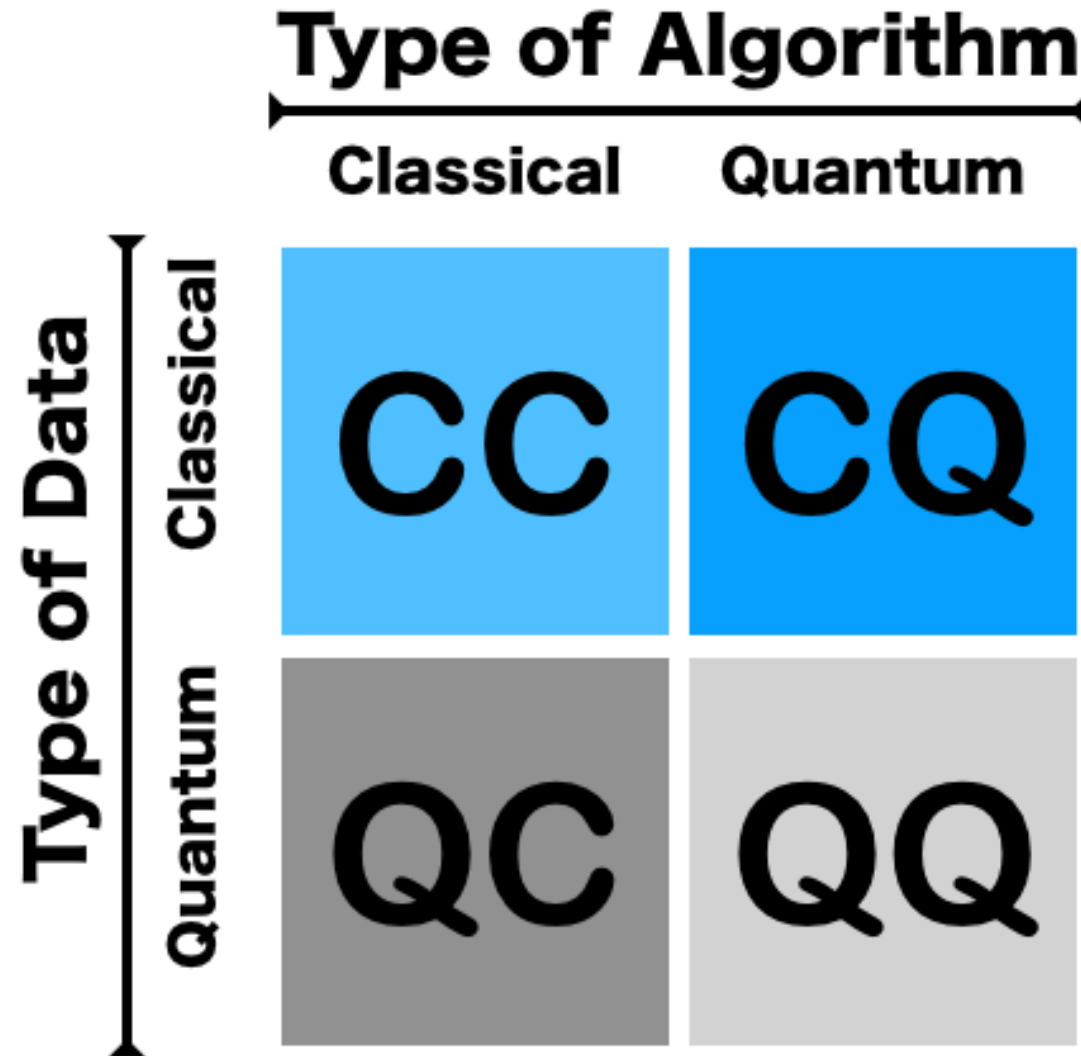
IV. Applications-Classification

V. Applications-Sequential Learning

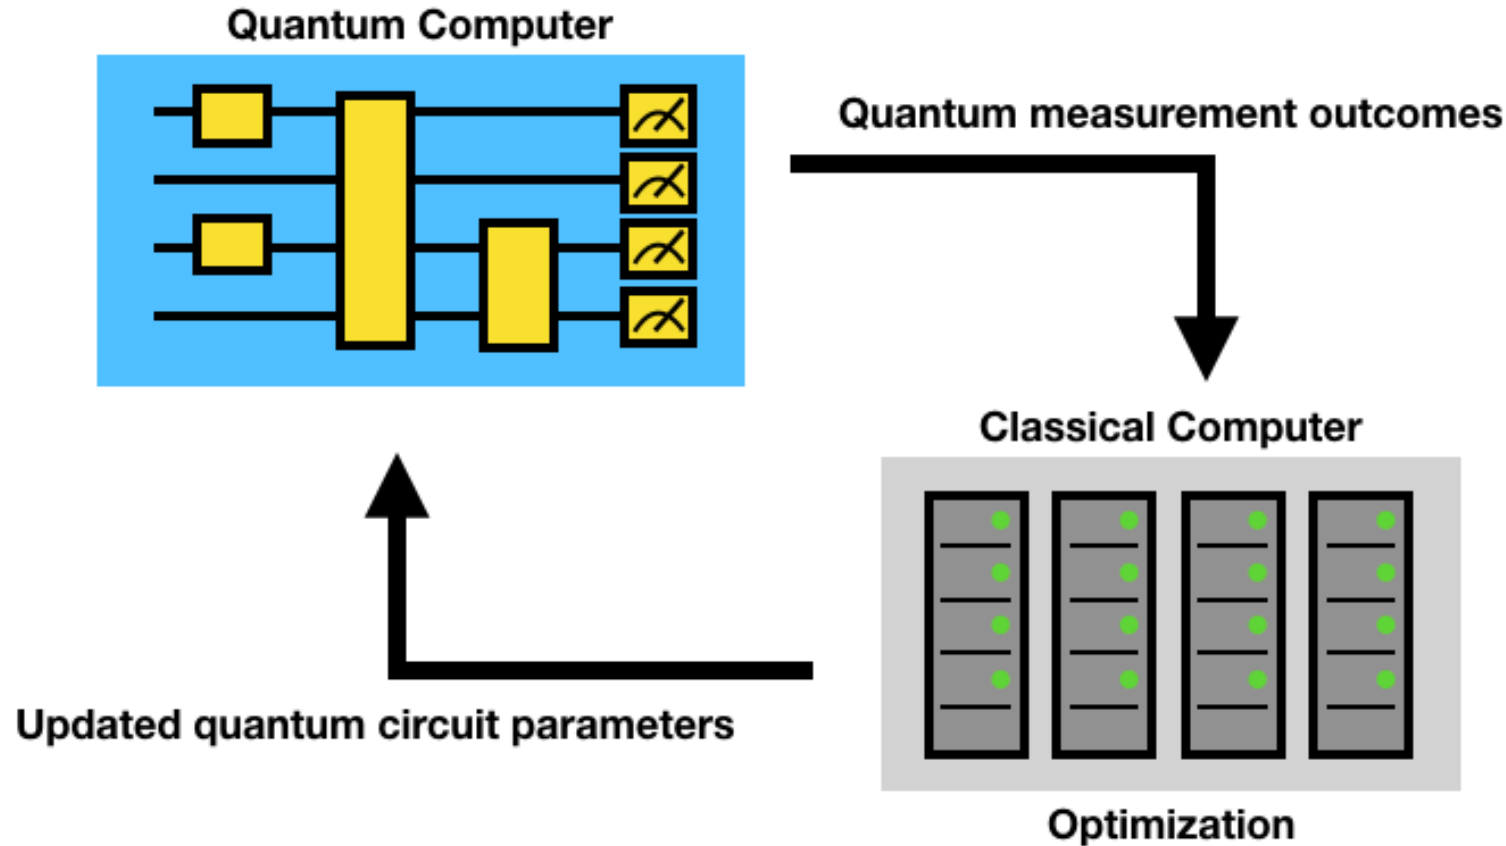
VI. Applications-Reinforcement Learning

VII. Conclusion and Outlook


Quantum Machine Learning

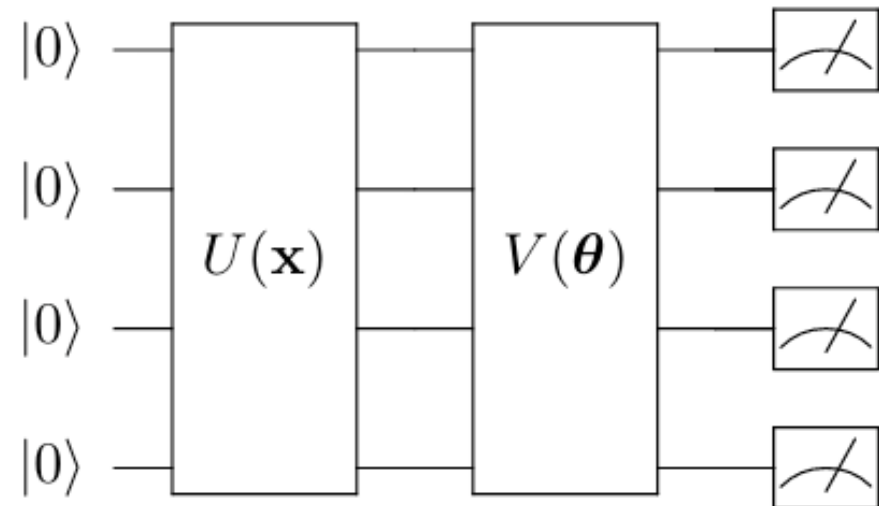


Hybrid Quantum-Classical Paradigm



Variational Quantum Circuits (VQC)

- Quantum circuits with **tunable** parameters.
- Subject to iterative optimization procedures (on classical computers).
- $U(\mathbf{x})$: encoding circuit
- $V(\theta)$: variational circuit
-  : measurement



Quantum Encoding and State Preparation

A general N qubit quantum state can be represented as:

$$|\psi\rangle = \sum_{(q_1, q_2, \dots, q_N) \in \{0,1\}} c_{q_1, q_2, \dots, q_N} |q_1\rangle \otimes |q_2\rangle \otimes \dots \otimes |q_N\rangle$$

where $c_{q_1, \dots, q_N} \in \mathbb{C}$ is the complex amplitude for each basis state and each $q_i \in \{0,1\}$

The total probability is equal to 1: $\sum_{(q_1, \dots, q_N) \in \{0,1\}} \|c_{q_1, \dots, q_N}\|^2 = 1$

Quantum Encoding and State Preparation

Amplitude Encoding

Encode a vector $(\alpha_0, \dots, \alpha_{2^n-1})$ into a n -qubit quantum state:

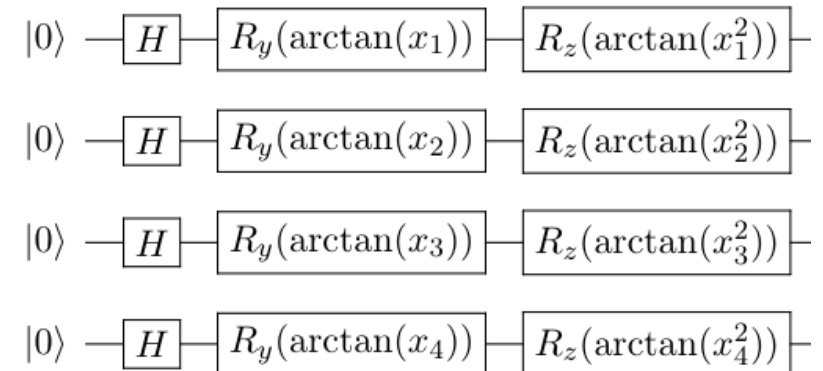
$$|\Psi\rangle = \alpha_0 |00\dots 0\rangle + \dots + \alpha_{2^n-1} |11\dots 1\rangle$$

where α_i are real numbers and $(\alpha_0, \dots, \alpha_{2^n-1})$ is normalized

N -dimensional vector will require only $\log_2(N)$ qubits to encode

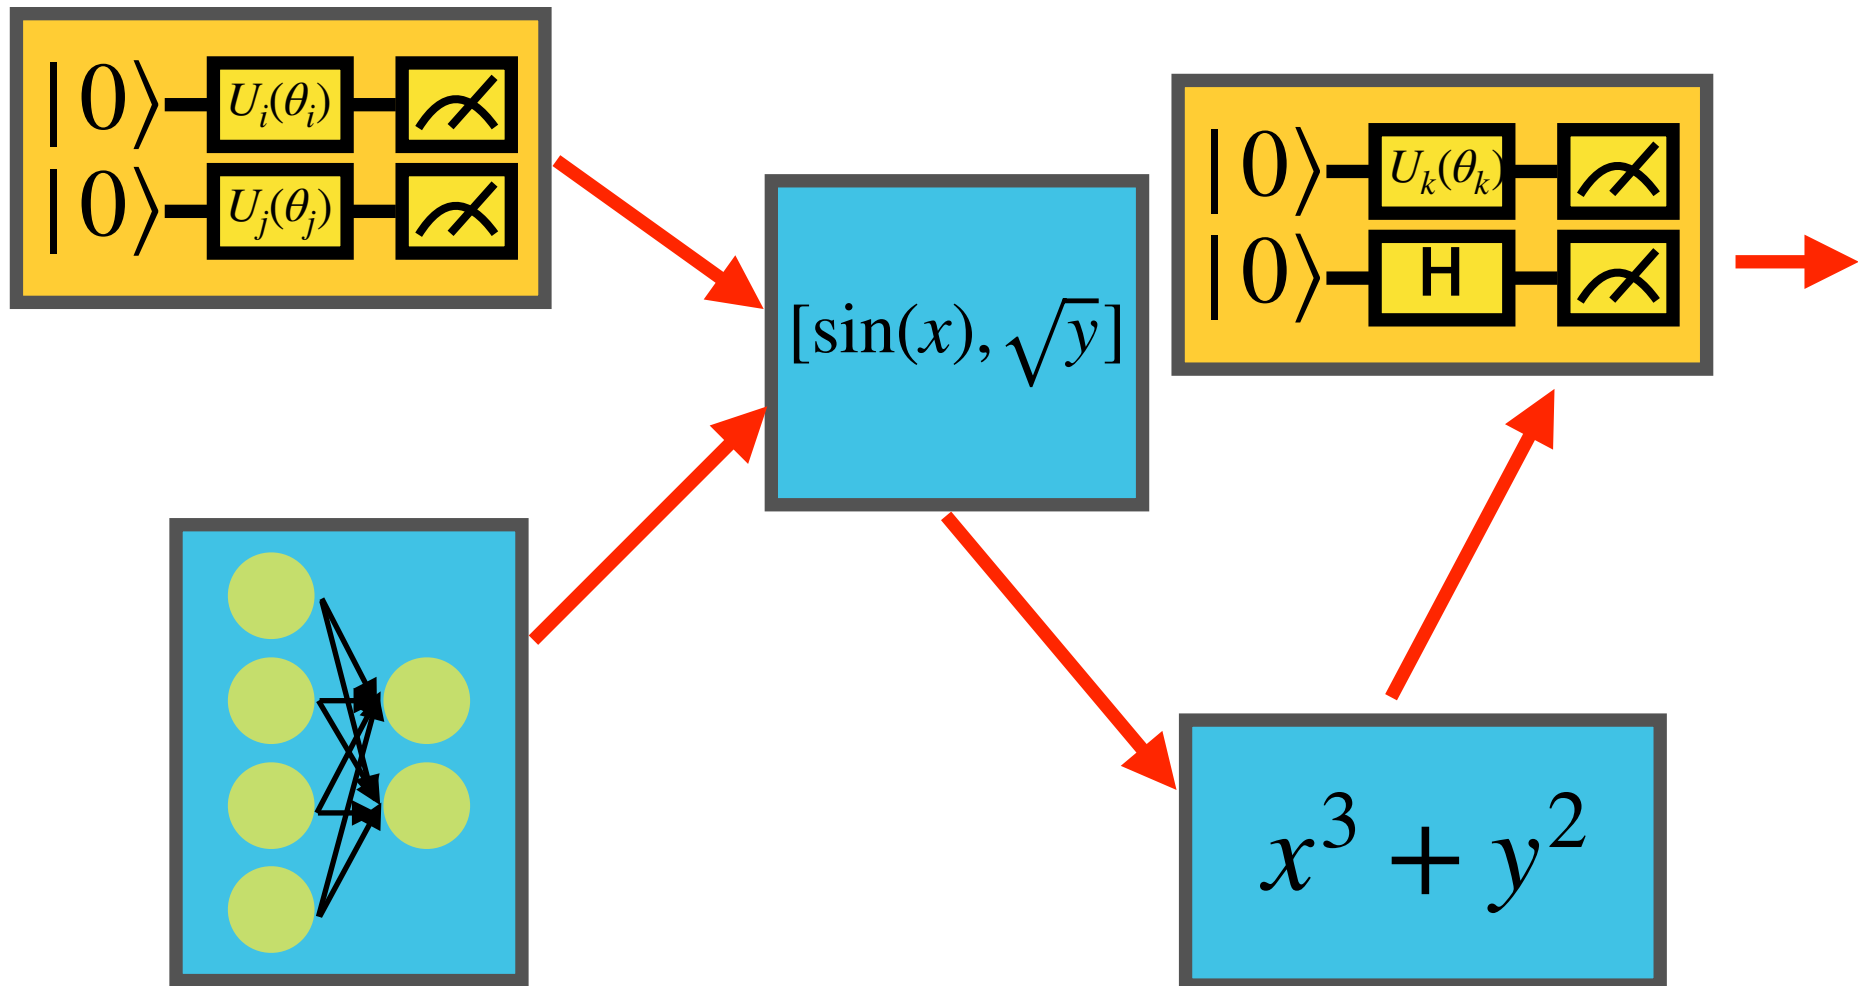
Variational Encoding

Input numbers $x_1 \dots x_n$ are used as quantum rotation angles



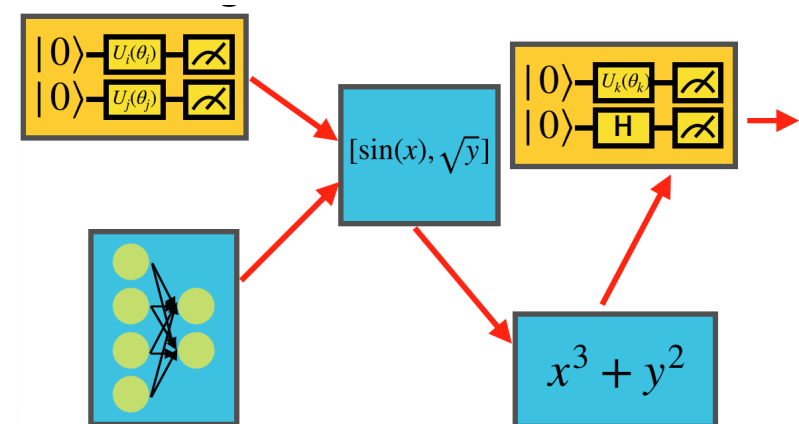
Simpler implementation than amplitude encoding

Interfacing with Classical ML

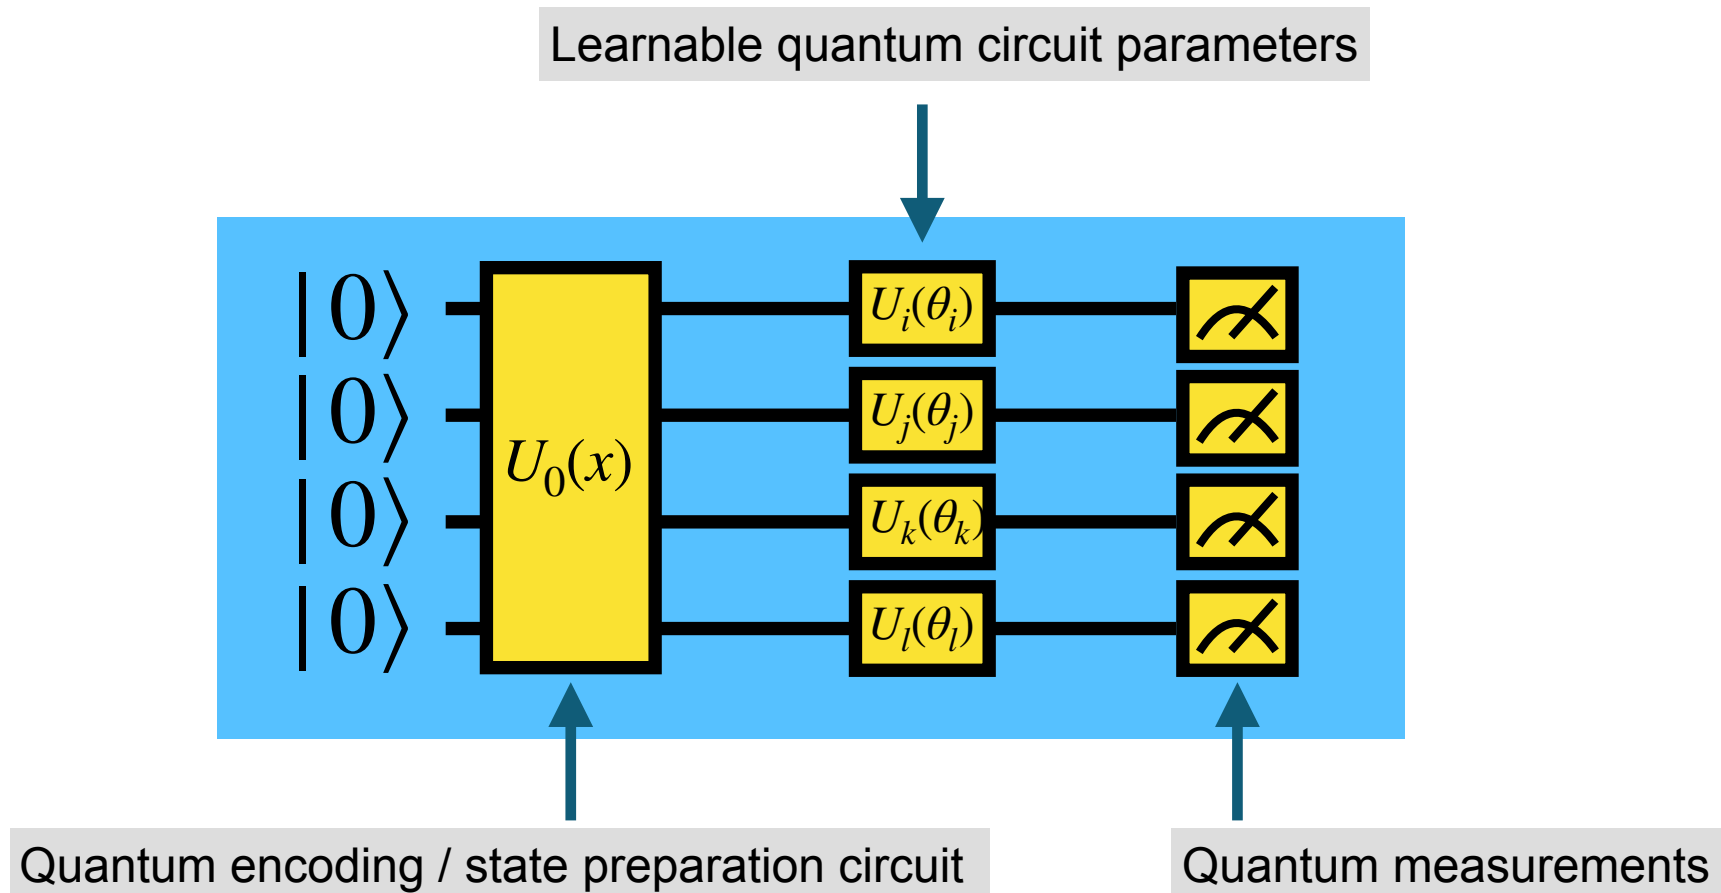


Interfacing with Classical ML

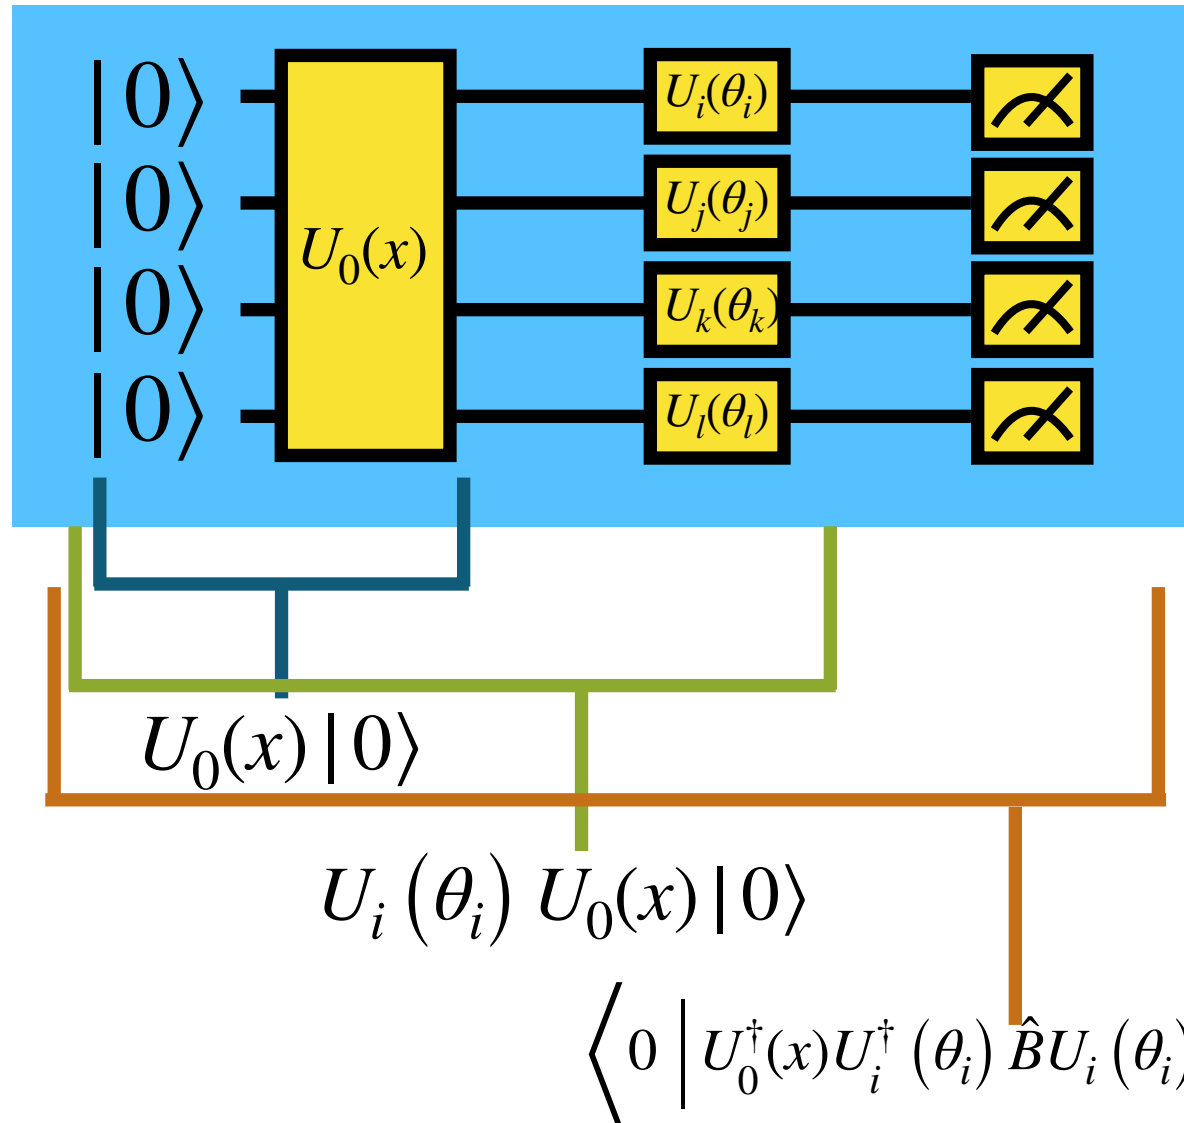
- Mixing *classical* and *quantum* computing components
- These classical and quantum nodes are arranged in a *directed acyclic graph* (DAG)
- The hybrid architecture is similar to the one in *deep learning* models
- The whole model can be trained with *backpropagation* method or other gradient-free methods such as evolutionary optimization
- The next question is “*How to calculate the gradient of a quantum node?*”



Quantum Gradients



Quantum Gradients



Quantum Gradients

$$f(x; \theta_i) = \left\langle 0 \left| U_0^\dagger(x) U_i^\dagger(\theta_i) \hat{B} U_i(\theta_i) U_0(x) \right| 0 \right\rangle = \left\langle x \left| U_i^\dagger(\theta_i) \hat{B} U_i(\theta_i) \right| x \right\rangle$$

- x : input value
- $U_0(x)$: encoding circuit
- i : circuit parameter index
- $U_i(x_i)$: single-qubit rotation generated by the Pauli operators

Quantum Gradients

The gradient of f with respect to the parameter θ_i is:

$$\nabla_{\theta_i} f(x; \theta_i) = \frac{1}{2} \left[f\left(x; \theta_i + \frac{\pi}{2}\right) - f\left(x; \theta_i - \frac{\pi}{2}\right) \right]$$

This value can be calculated via running two quantum circuits with shifted parameters, the so-called *parameter-shift* rule.

PRA 98, 032309 (2018)

PRA 99, 032331 (2019)

I. Introduction

II. Quantum Computing (QC)

III. Quantum Machine Learning

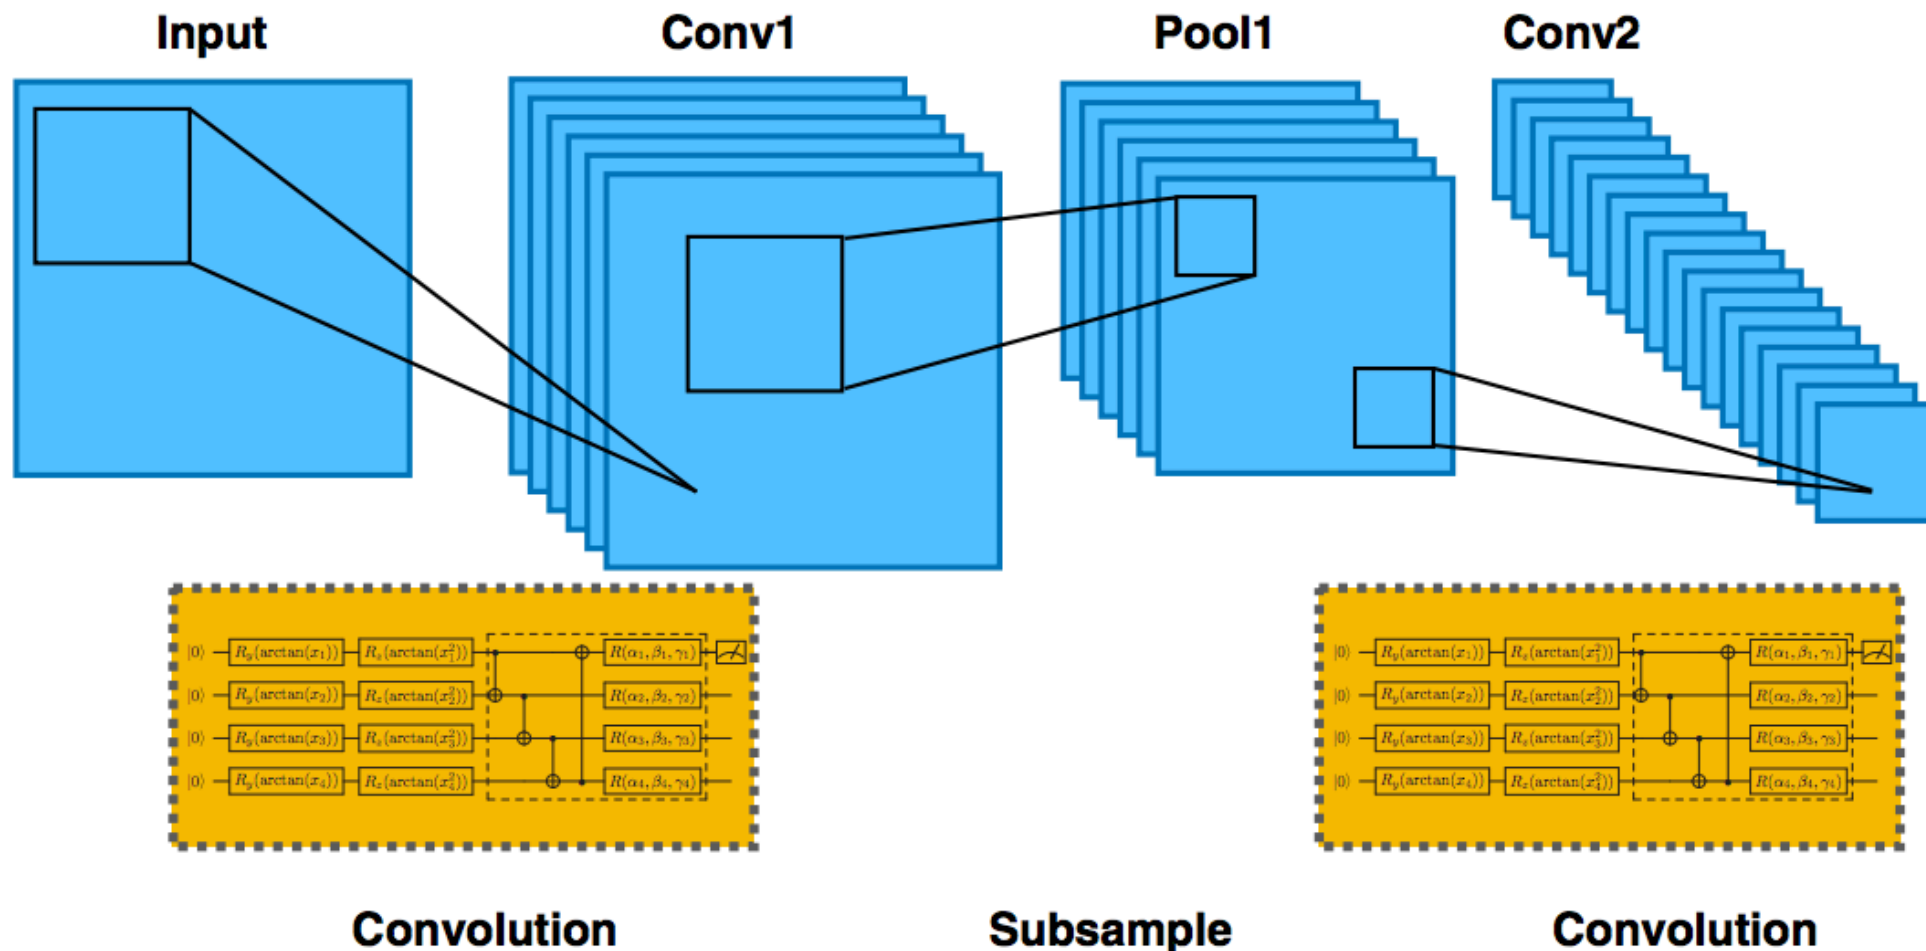
IV. Applications-Classification

V. Applications-Sequential Learning

VI. Applications-Reinforcement Learning

VII. Conclusion and Outlook

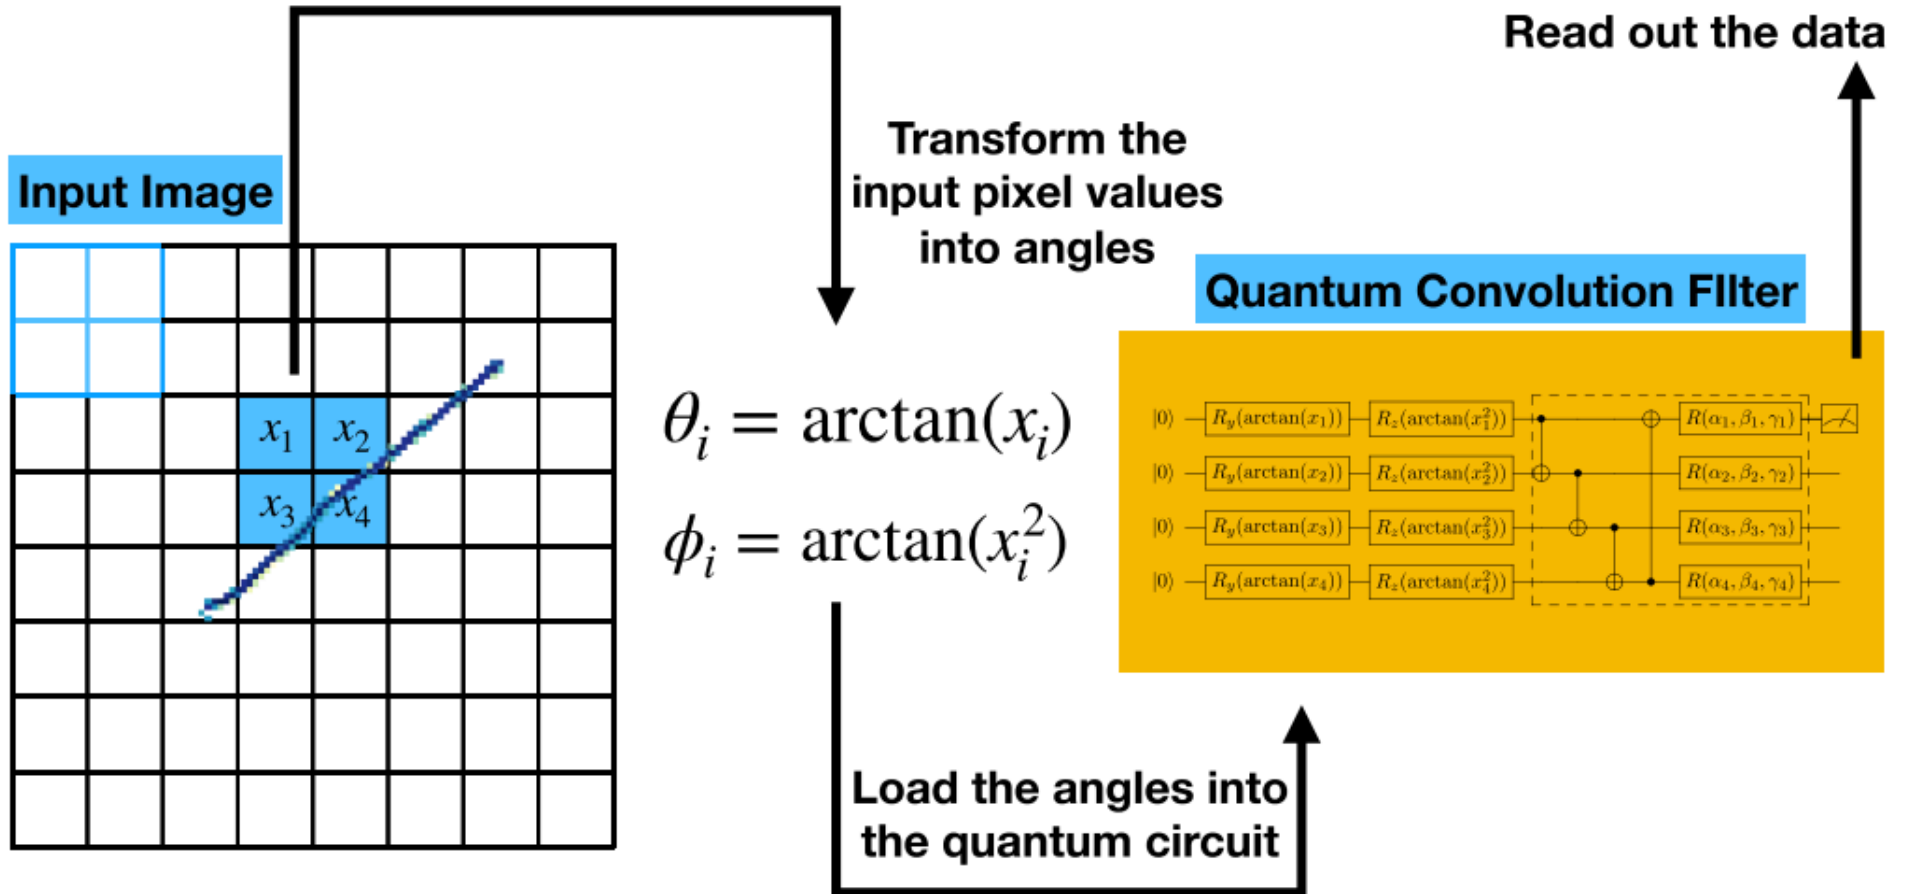
Quantum Convolutional Neural Network

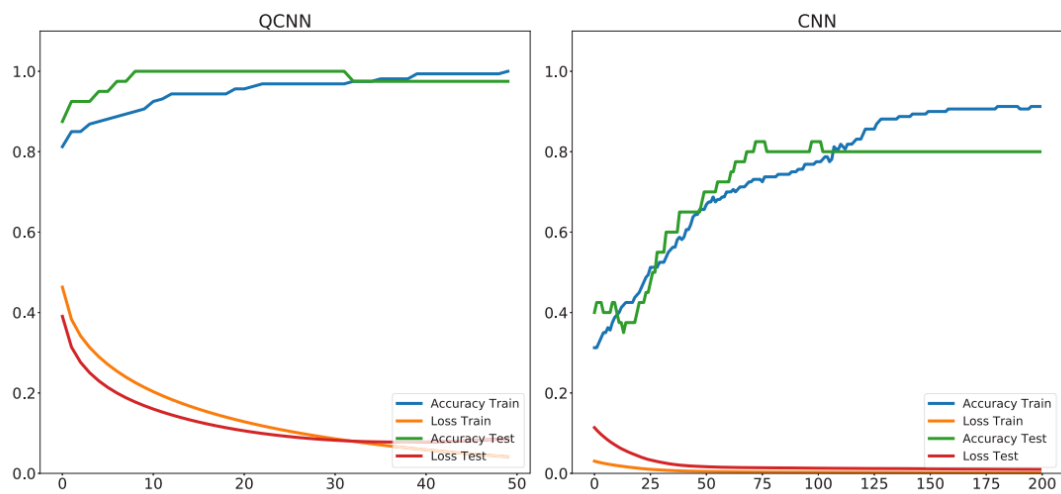
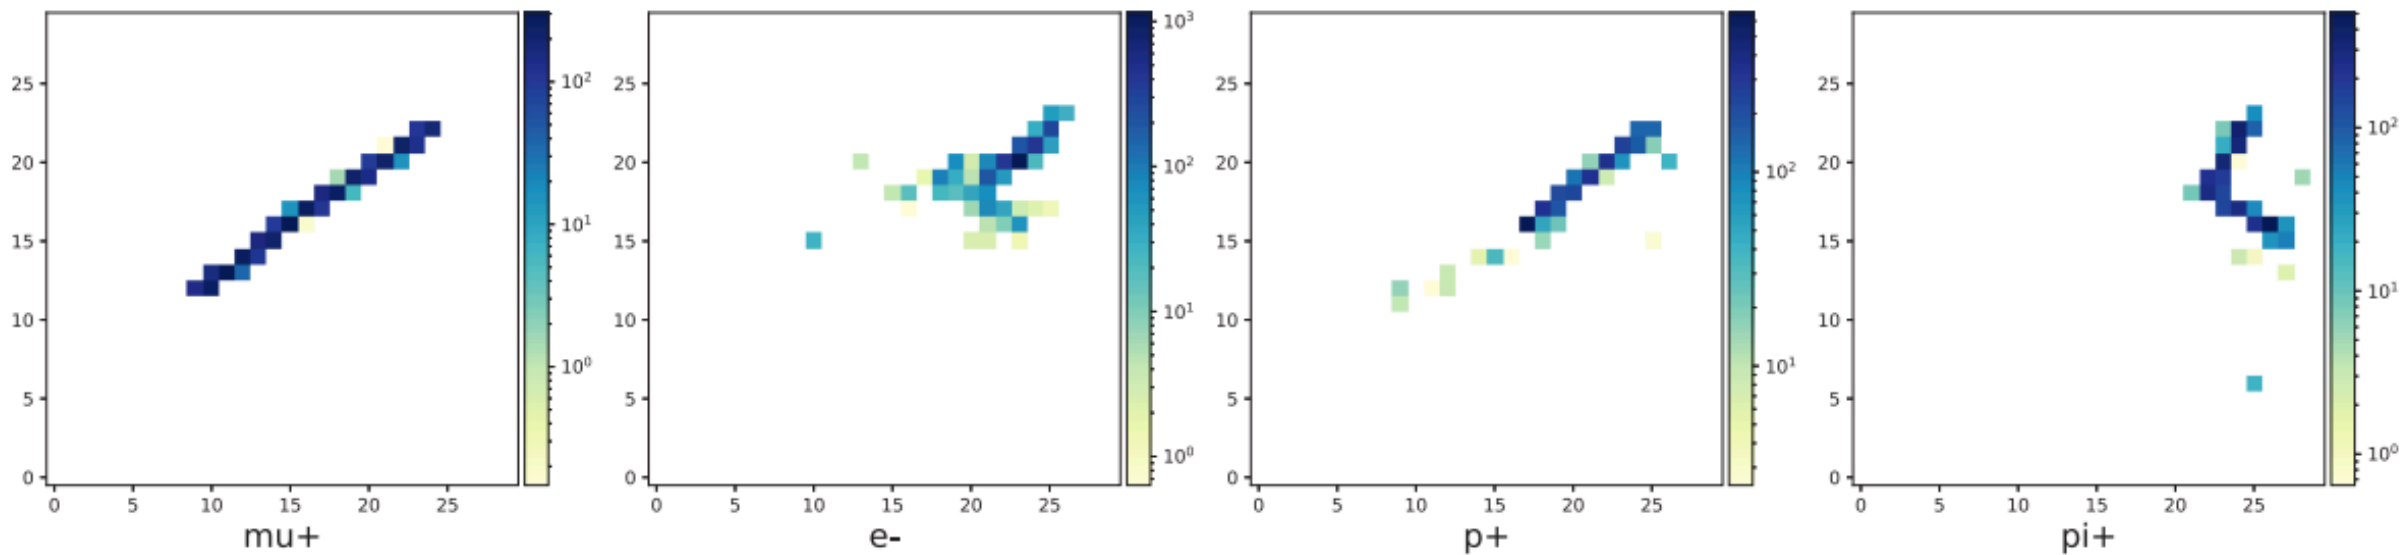


Quantum Convolutional Neural Network

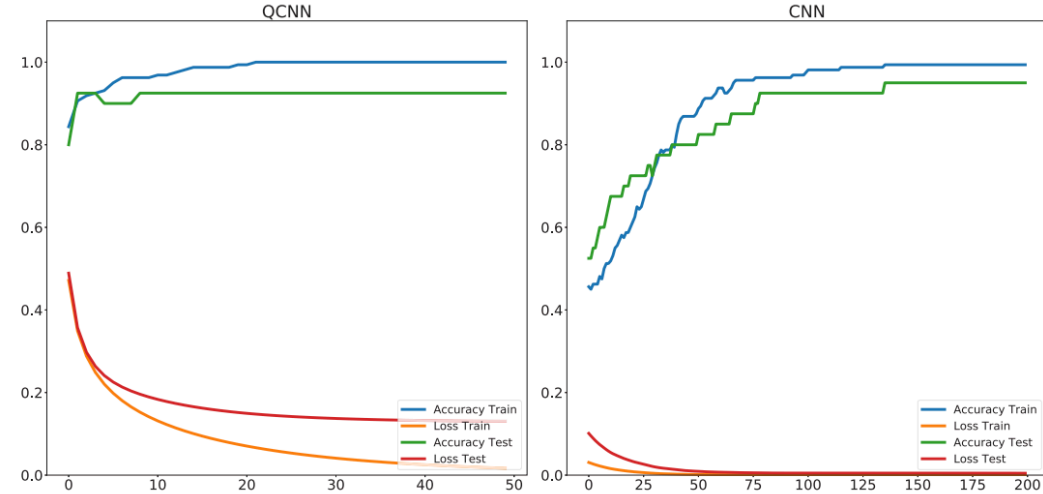
Scan over the input image

Pixel values (x_1, x_2, x_3, x_4)





mu+ vs p

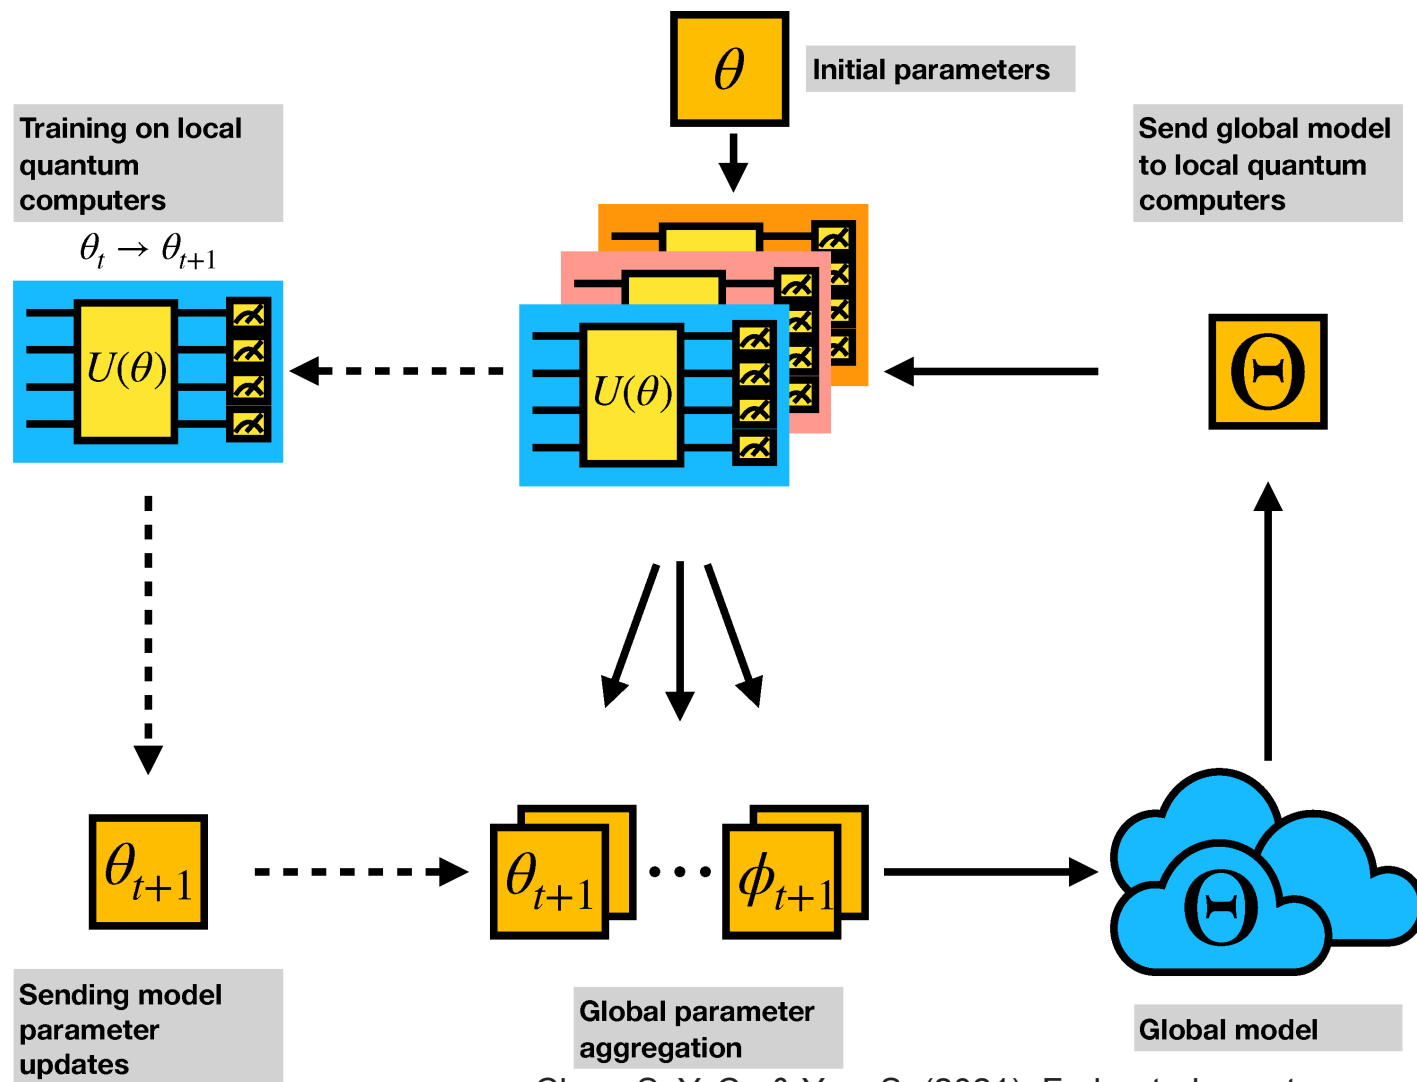


mu+ vs e-

TABLE I. The comparison of model architectures between QCNN and CNN used in this paper.

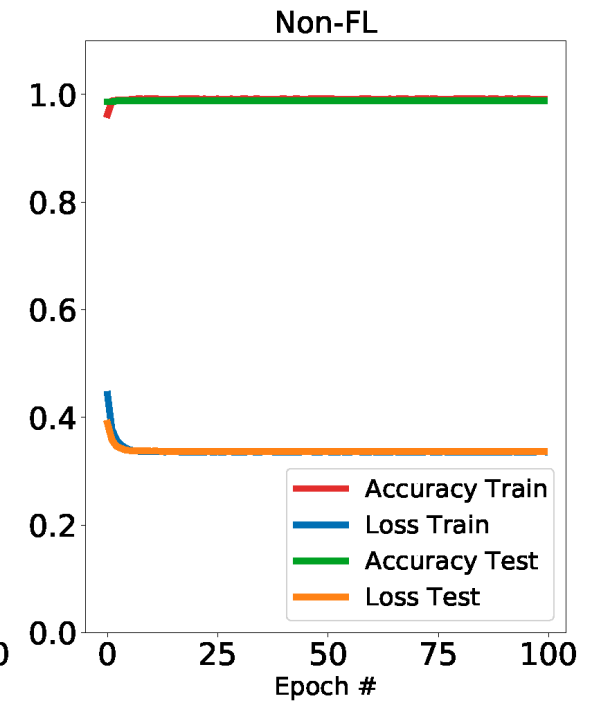
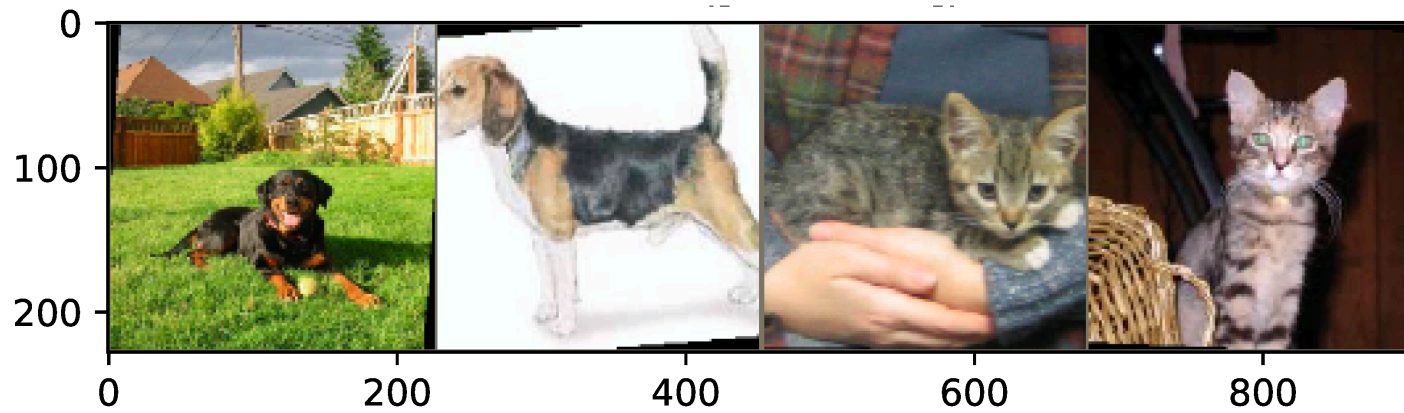
	First conv layer (# of channels)	Filter size	Second conv layer (# of channels)	Filter size	Classical part (# of params)	Total number of params
QCNN	1	3 × 3	1	2 × 2	394	472
CNN	4	5 × 5	2	5 × 5	198	498

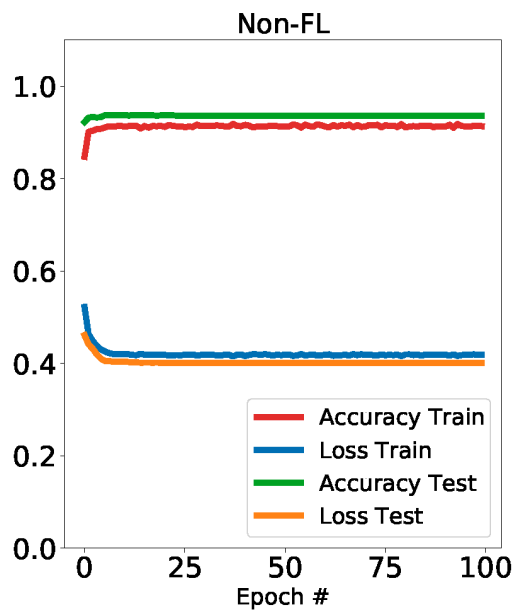
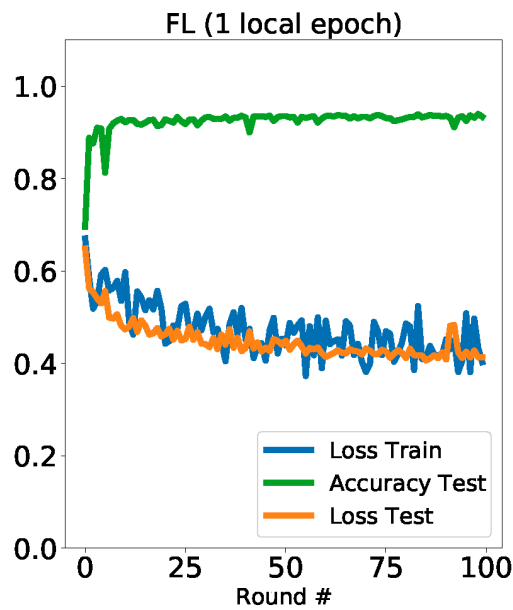
Federated Quantum Machine Learning



Federated Quantum Machine Learning







I. Introduction

II. Quantum Computing (QC)

III. Quantum Machine Learning

IV. Applications-Classification

V. Applications-Sequential Learning

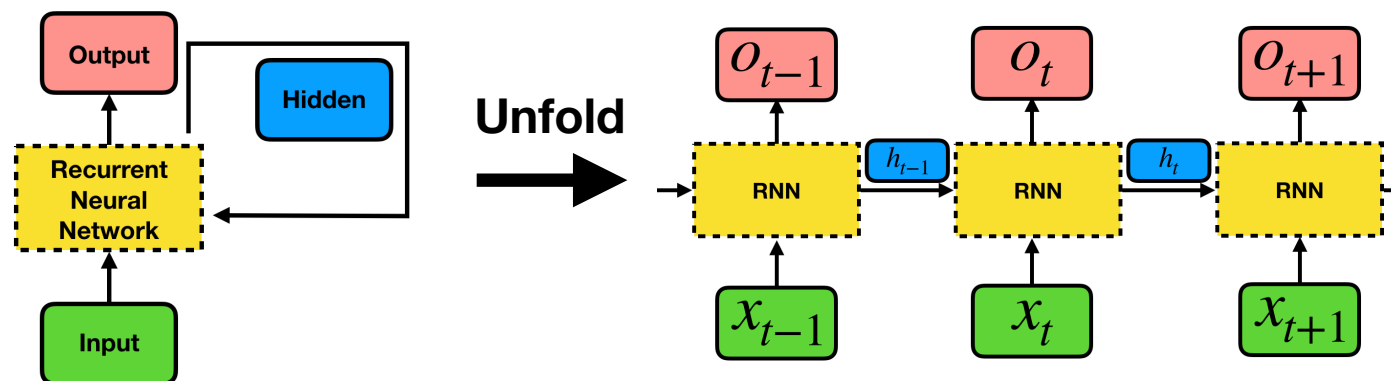
VI. Applications-Reinforcement Learning

VII. Conclusion and Outlook

Sequential Learning: Motivation

- Natural language processing: question answering, machine translation, ...
- Physical dynamics: quantum dynamics, quantum optimal control, ...
- Reinforcement learning (RL) agents also need a **memory**
- Goal #1: to predict time series accurately
- Goal #2: to explore if quantum computers could deliver better outcome

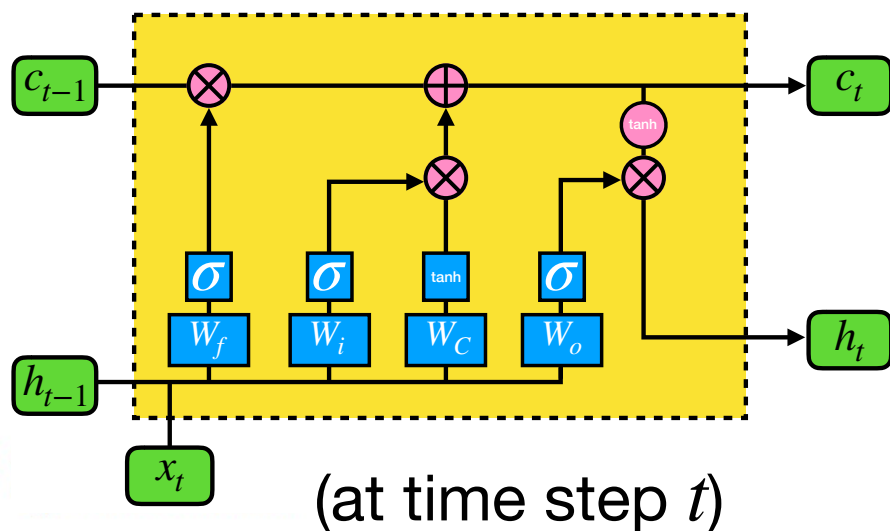
Recurrent Neural Network (RNN) and LSTM



Cons of RNN:

1. not suitable for long memory time
2. not easy to train (vanishing gradient)

(Classical) Long short-term memory (LSTM)



c_t : Cell state

h_t : Hidden state

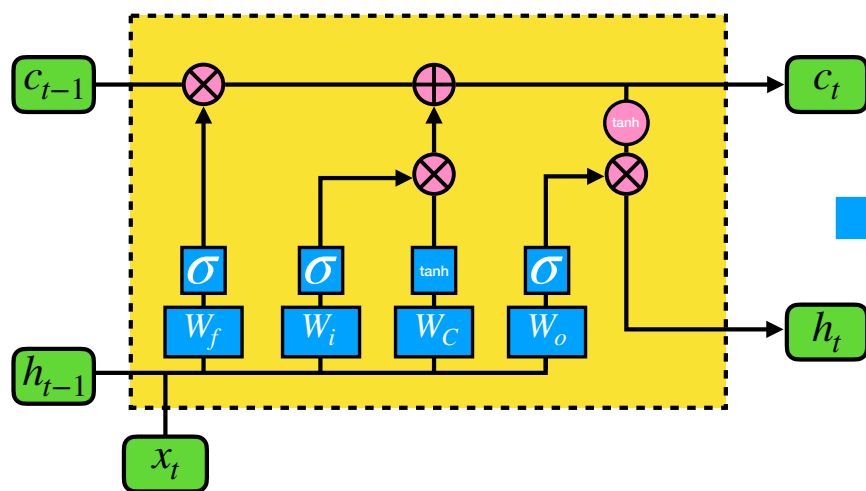
x_t : Input

$\{W\}$: Neural networks

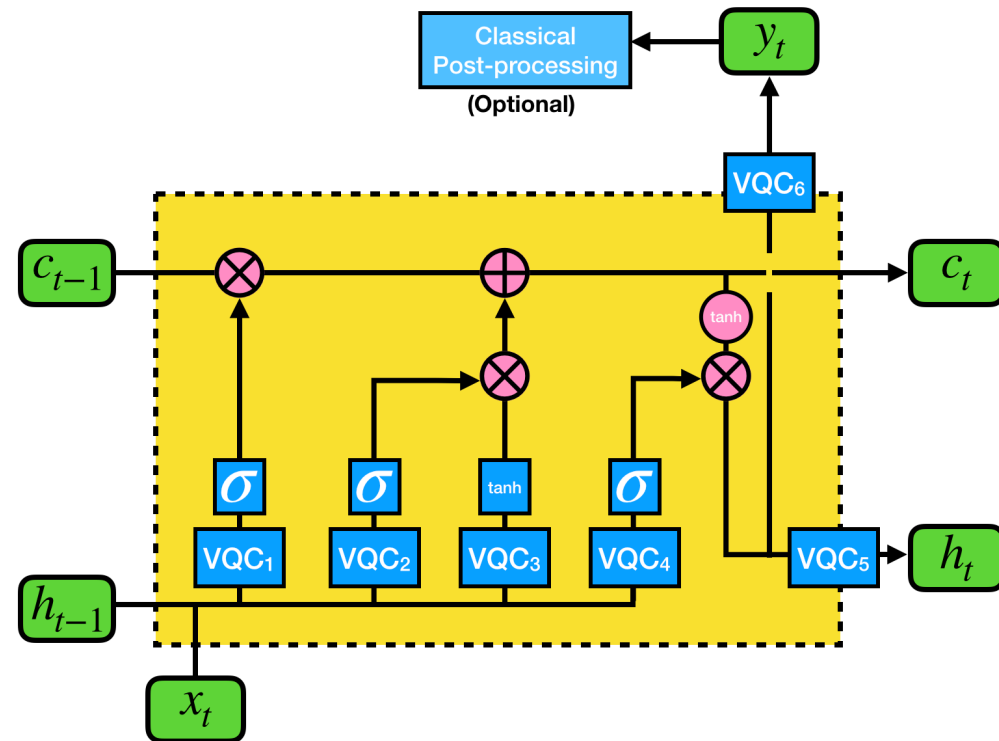
Use nonlinear activation functions (σ & \tanh)

Quantum LSTM (QLSTM)

- Replace the classical neural networks with variational quantum circuits (VQC)
- Other components (input, output, activation functions, internal states) are still kept “classical” to avoid no-clone

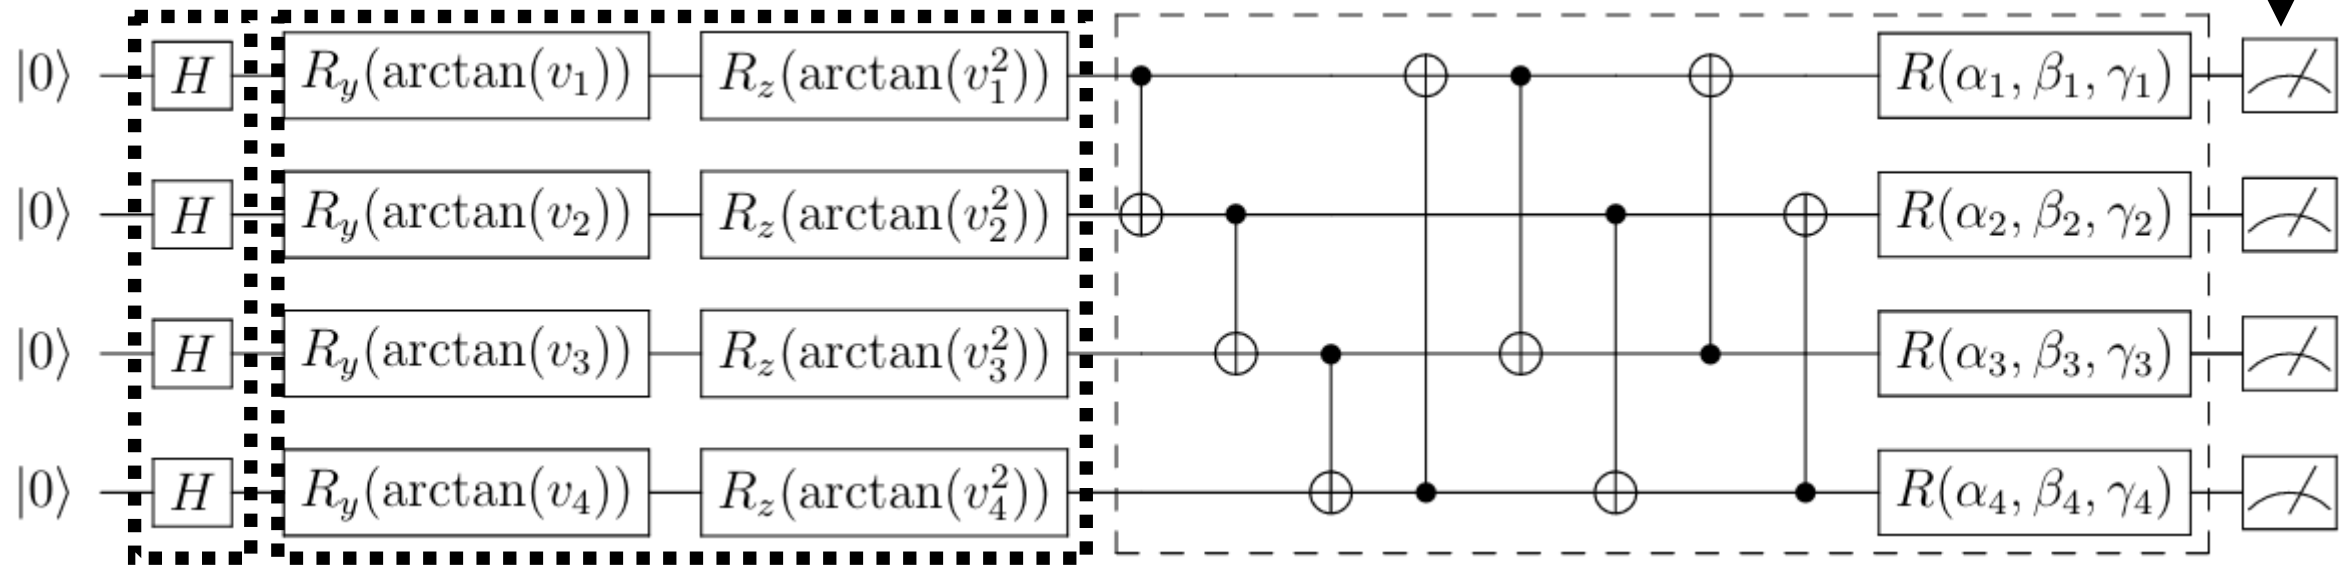


(at time step t)



VQC in QLSTM

Circuit component for QLSTM



Measurements

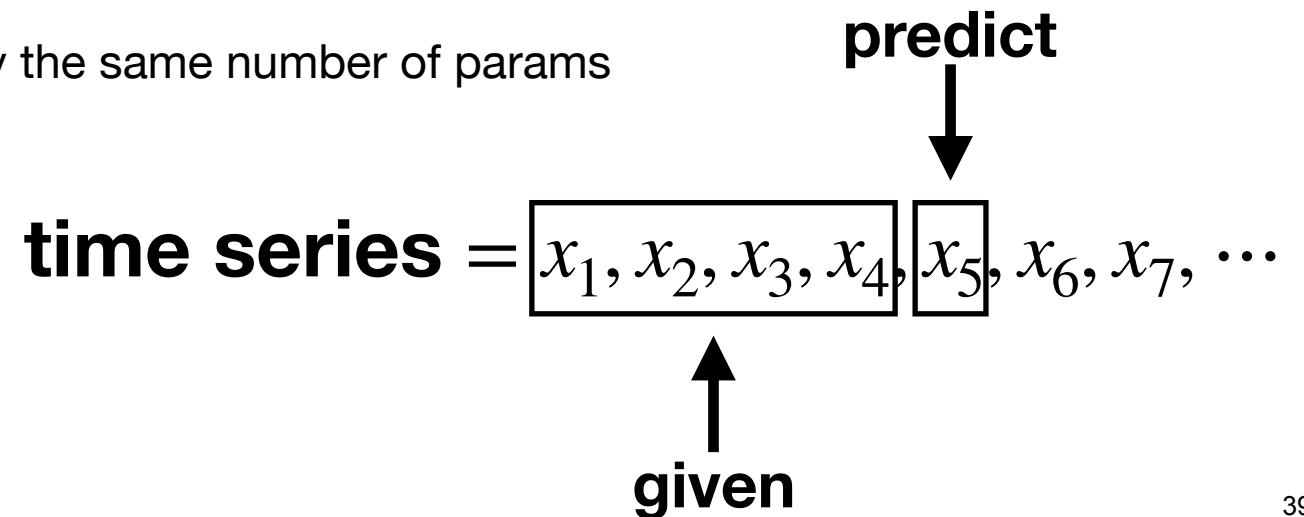
Initialize the unbiased state

Variational Encoding

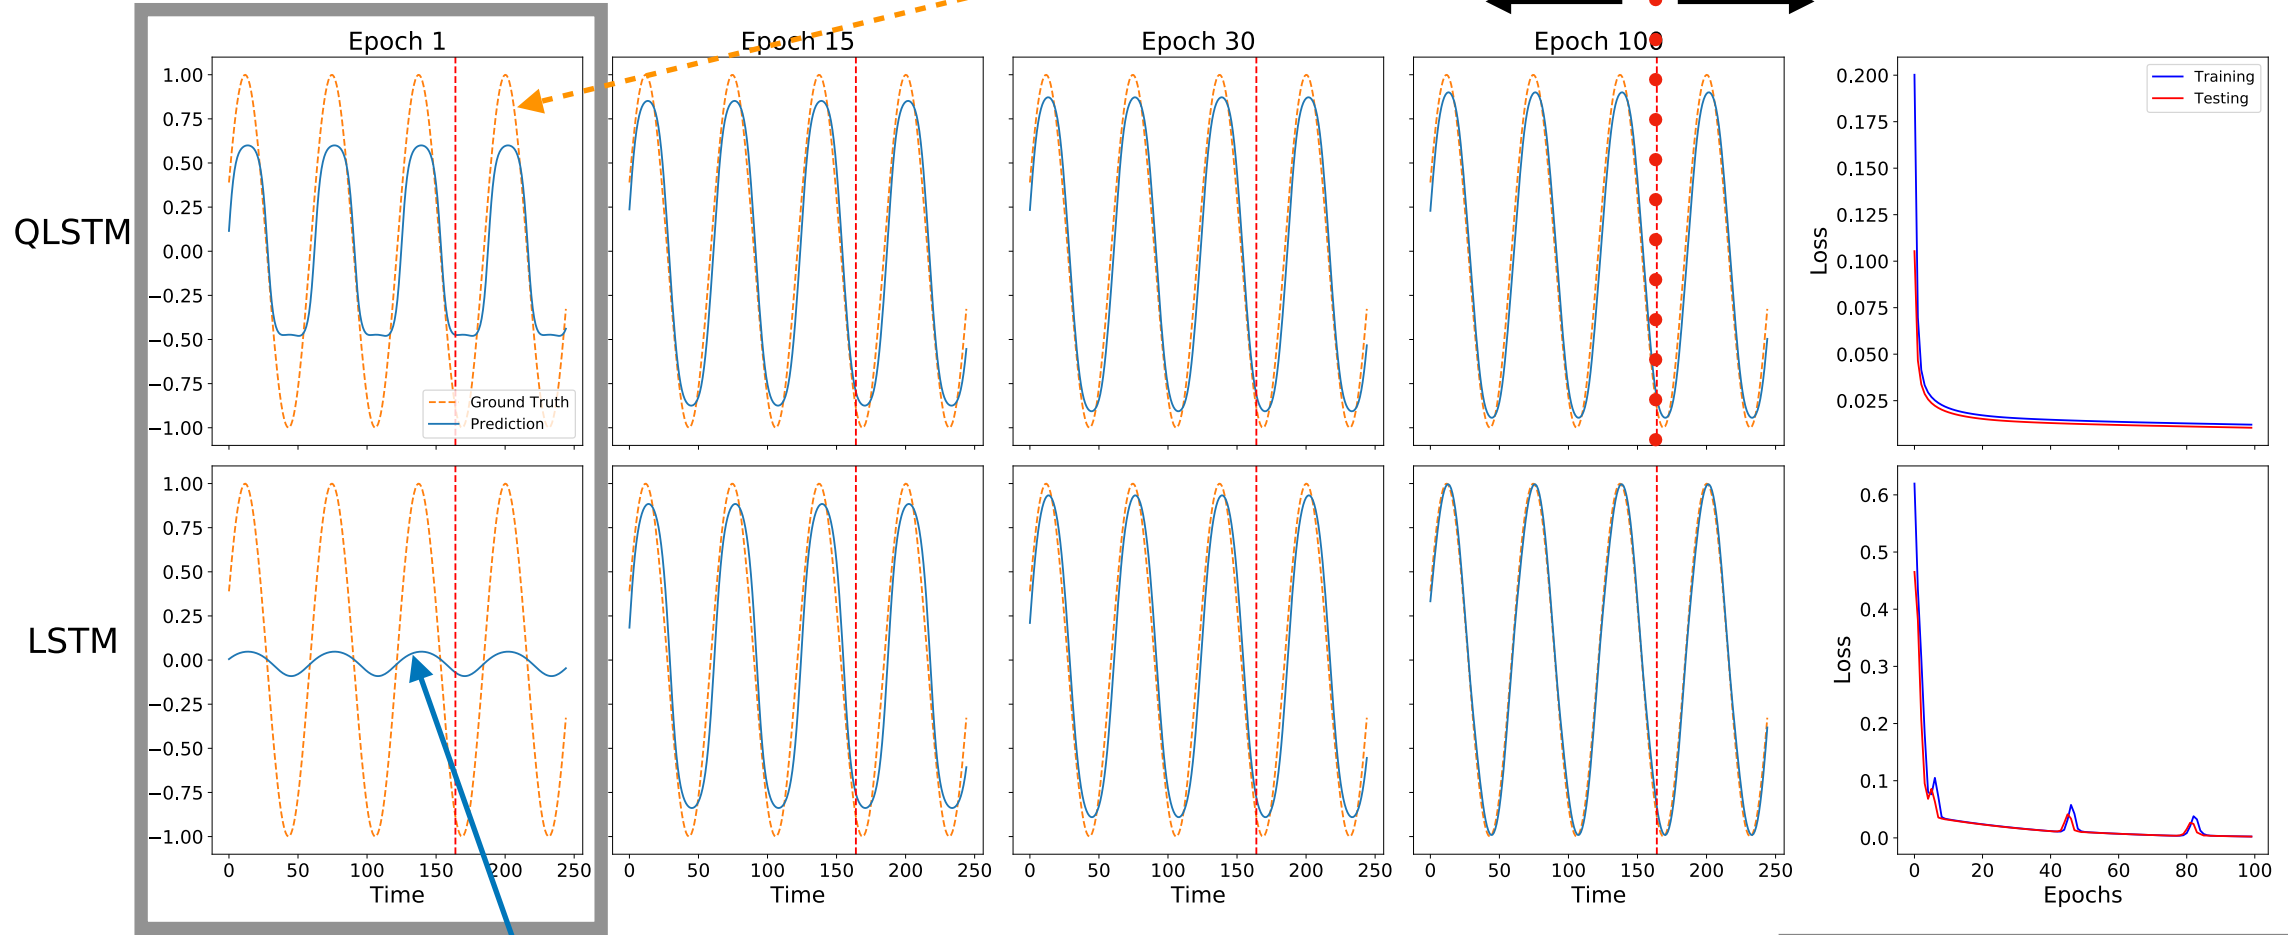
Learning parts

Experiments

- Goal is to model/predict time series in a sliding window fashion
- Given 4 points from previous time steps $x_{t-4}, x_{t-3}, x_{t-2}, x_{t-1}$, predict x_t
- The time series data is loaded into the (Q)LSTM sequentially
- In each time step t , the concatenation of x_t and hidden state from previous time step h_{t-1} , $[x_t, h_{t-1}]$ is loaded into the quantum circuit
- Compare QLSTM with LSTM that has roughly the same number of params
- Loss function = Mean Square Error (MSE)
- Optimizer = RMSProp
- Noiseless quantum simulator



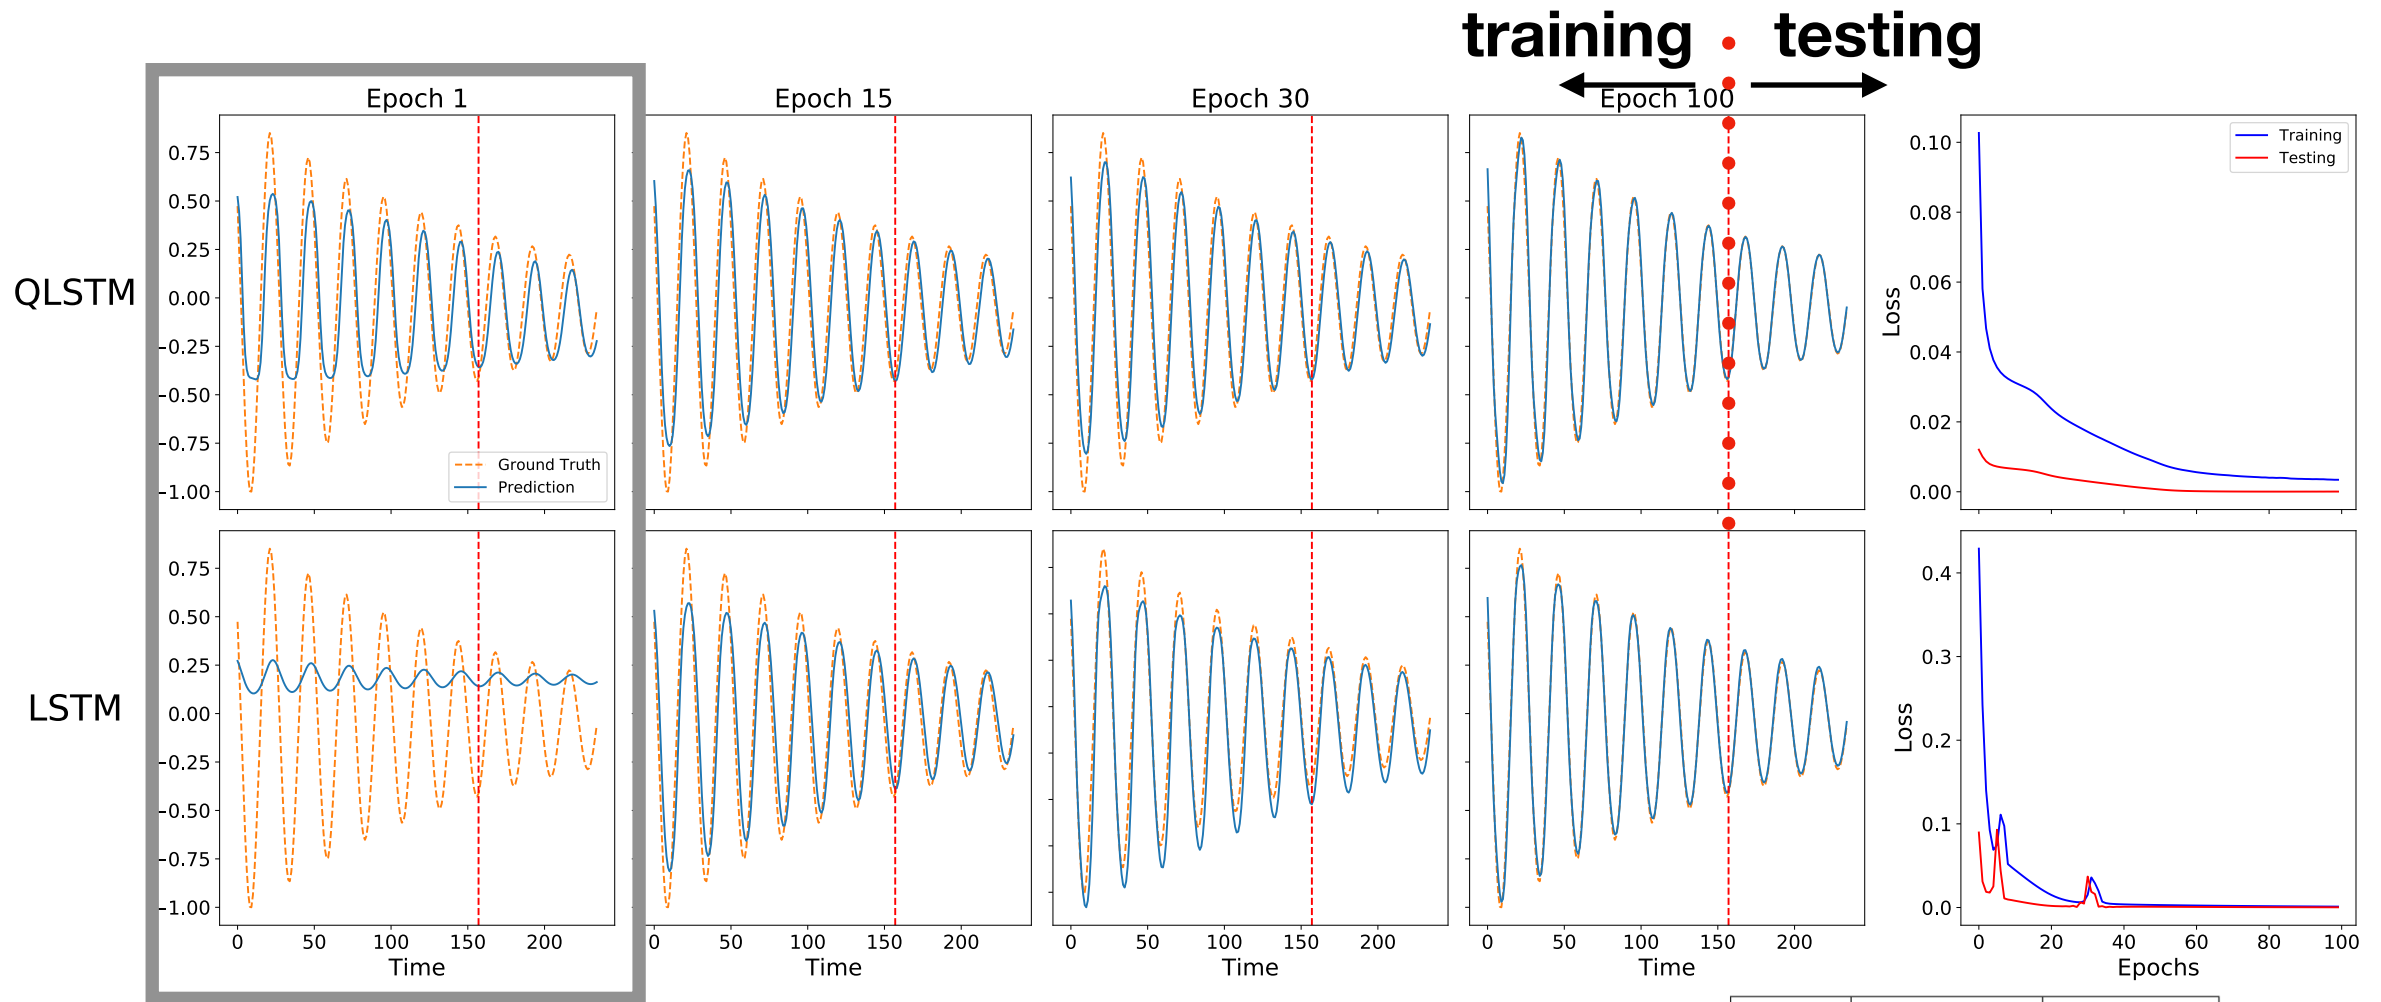
Test case: $\sin(t)$ ground truth training • testing



prediction

	Training Loss	Testing Loss
QLSTM	1.89×10^{-2}	1.69×10^{-2}
LSTM	2.86×10^{-2}	2.81×10^{-2}

Test case: damped oscillation

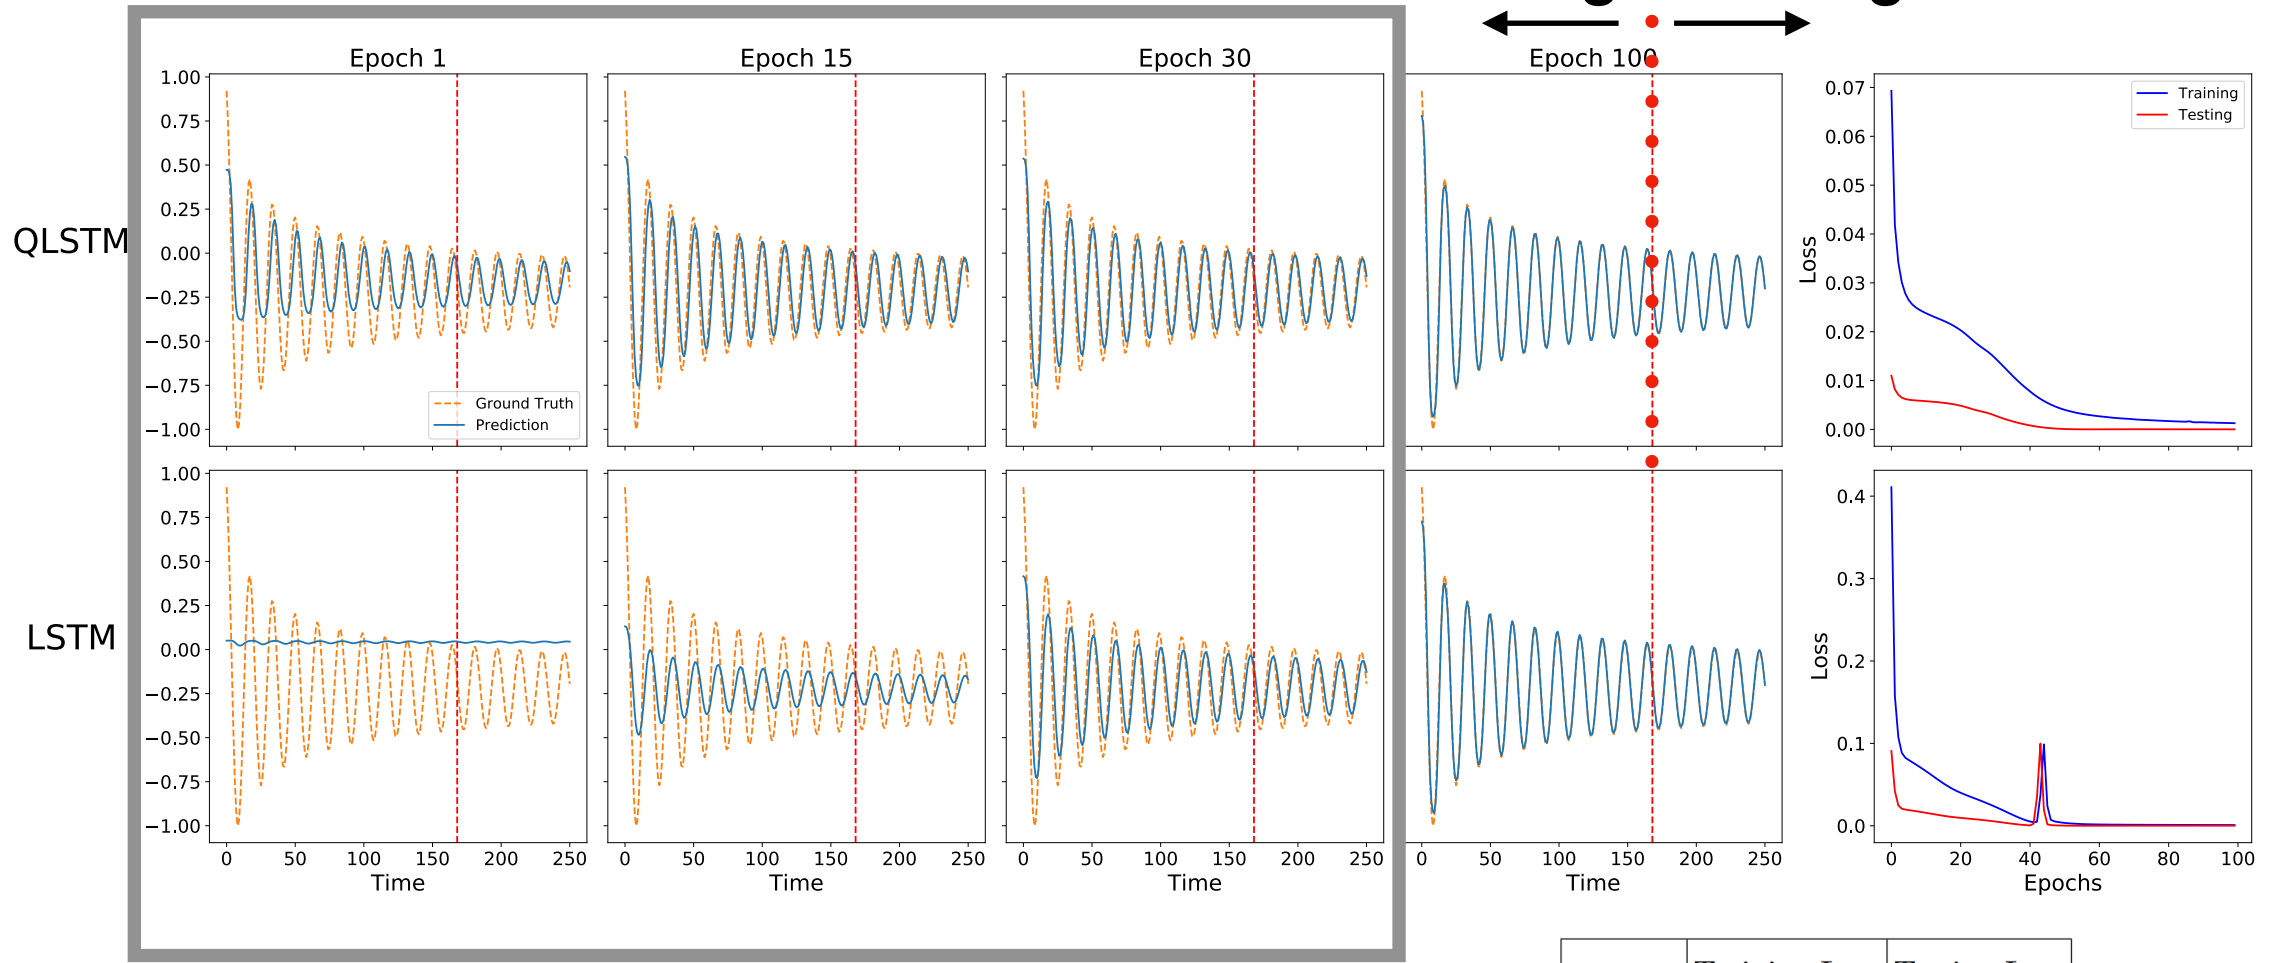


$$\frac{d^2\theta}{dt^2} + \frac{b}{m} \frac{d\theta}{dt} + \frac{g}{L} \sin \theta = 0$$

	Training Loss	Testing Loss
QLSTM	2.92×10^{-2}	6×10^{-3}
LSTM	3.15×10^{-2}	5×10^{-3}

Test case: Bessel function $J_2(t)$

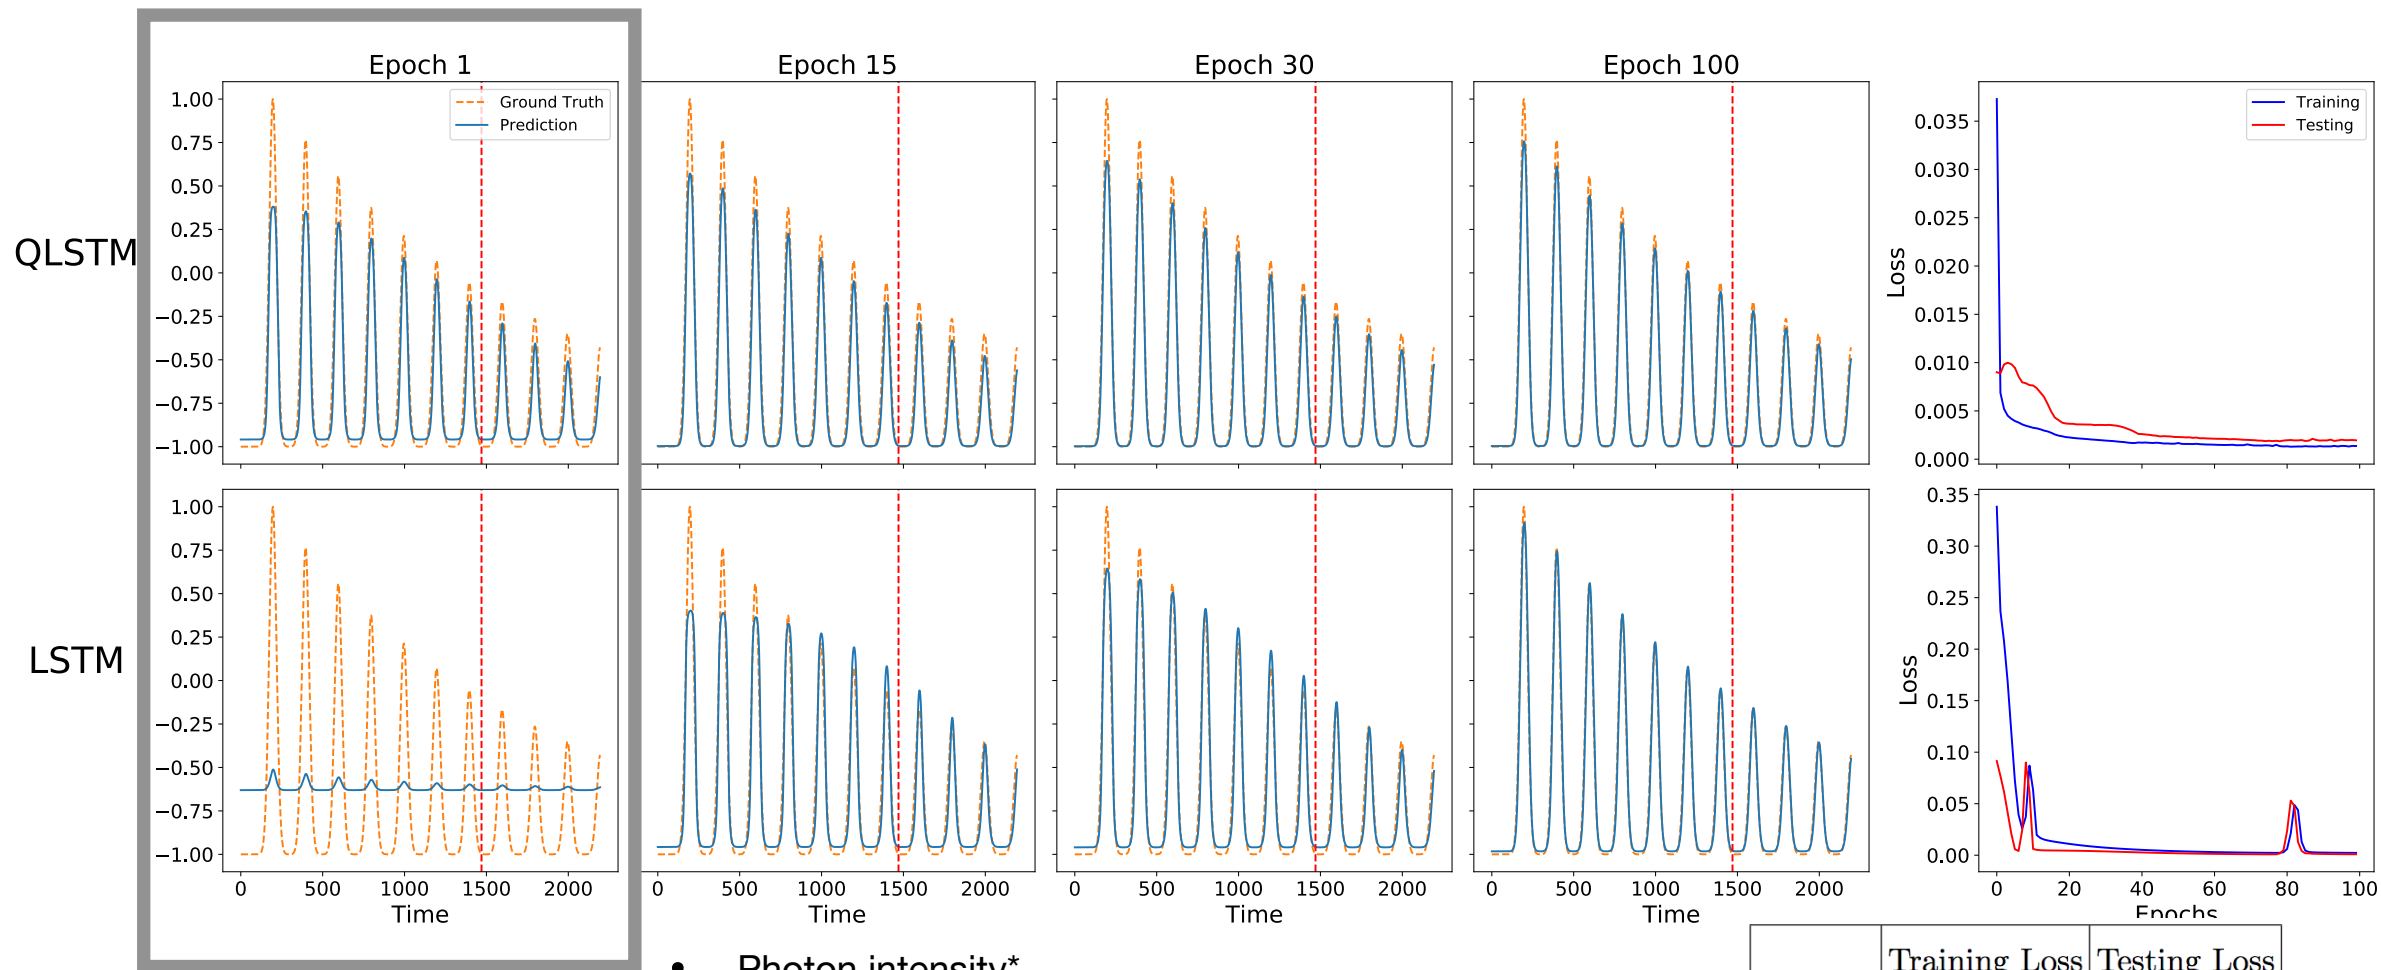
training • testing



$$J_\alpha(x) = \sum_{m=0}^{\infty} \frac{(-1)^m}{m! \Gamma(m + \alpha + 1)} \left(\frac{x}{2}\right)^{2m + \alpha}$$

	Training Loss	Testing Loss
QLSTM	2.26×10^{-2}	5.5×10^{-3}
LSTM	5.43×10^{-2}	1.28×10^{-2}

Test case: photon trapping



- Photon intensity*
- Example of non-Markovian quantum dynamics (due to delayed quantum feedback)

*See, e.g., Tufarelli et al, PRA 87, 013820 (2013)

	Training Loss	Testing Loss
QLSTM	2.88×10^{-3}	5.7×10^{-3}
LSTM	1.44×10^{-2}	4.7×10^{-3}

I. Introduction

II. Quantum Computing (QC)

III. Quantum Machine Learning

IV. Applications-Classification

V. Applications-Sequential Learning

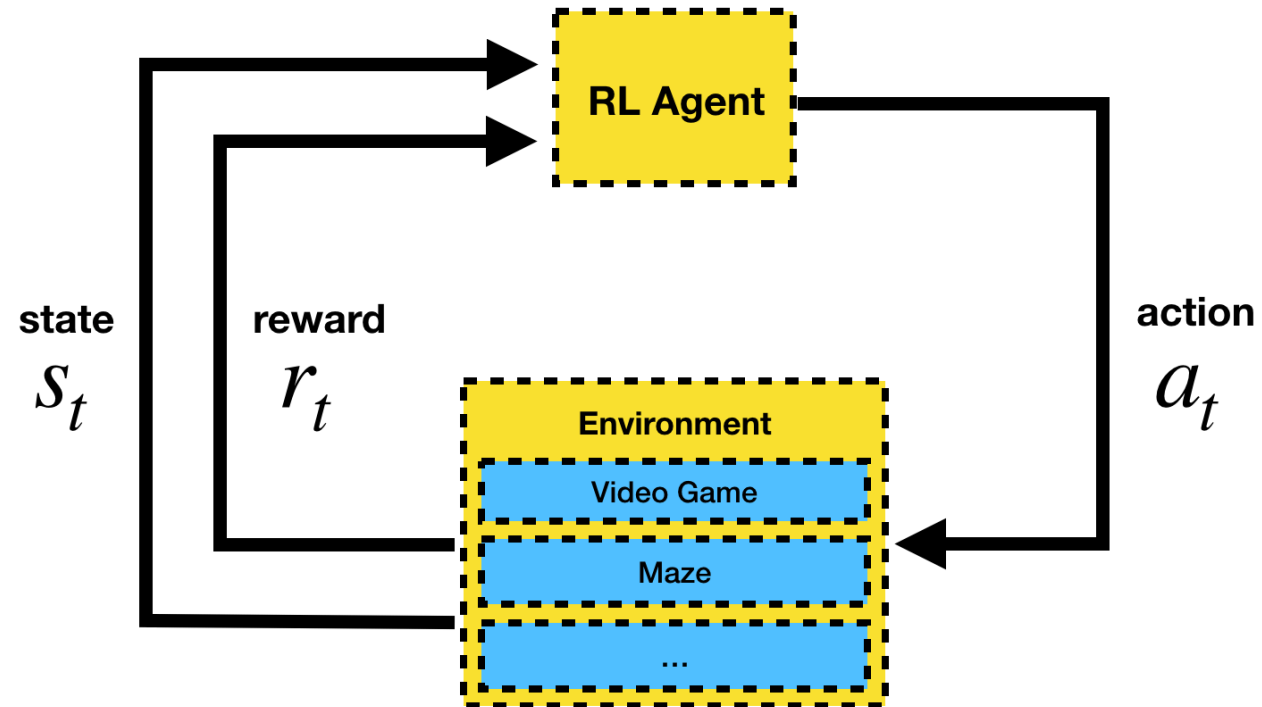
VI. Applications-Reinforcement Learning

VII. Conclusion and Outlook

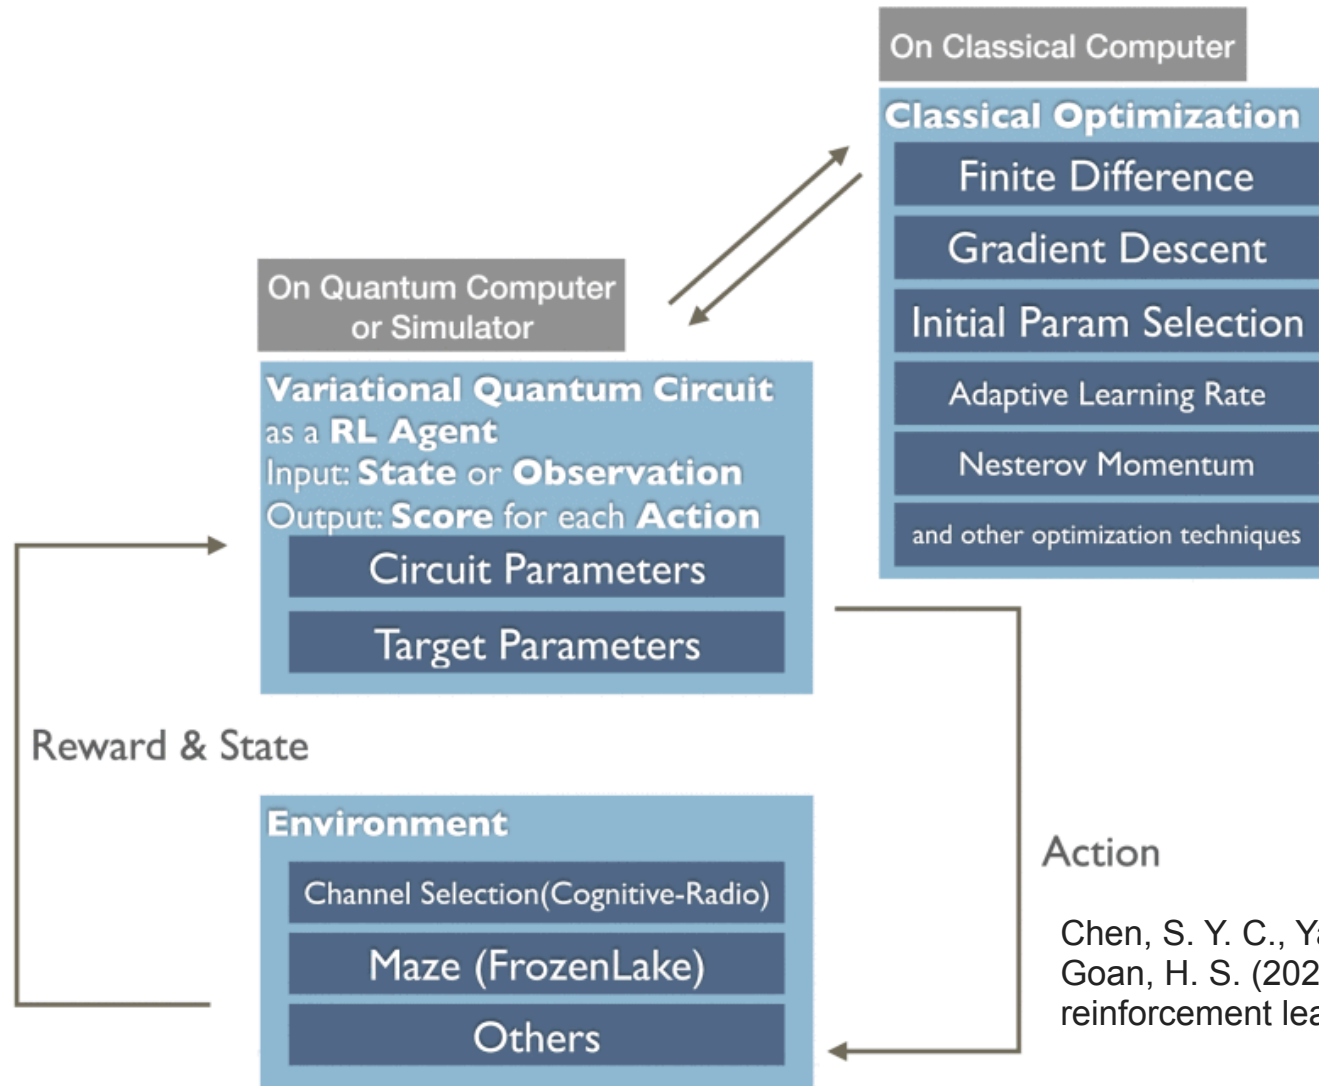
Reinforcement Learning (RL)

- RL: An agent interacts with an **environment** \mathcal{E} over a number of discrete time steps.
- The agent receives **state** or **observation** s_t and then chooses an **action** a_t from a set of actions \mathcal{A} according to its **policy** π .
- Goal: Maximize the **total discounted**

$$\text{return } R_t = \sum_{t'=t}^T \gamma^{t'-t} r_{t'}$$



Quantum Reinforcement Learning



Chen, S. Y. C., Yang, C. H. H., Qi, J., Chen, P. Y., Ma, X., & Goan, H. S. (2020). Variational quantum circuits for deep reinforcement learning. *IEEE Access*, 8, 141007-141024.

Quantum Deep Q-Learning

Algorithm 1 Variational Quantum Deep Q Learning

Initialize replay memory \mathcal{D} to capacity N

Initialize action-value function quantum circuit Q with random parameters

for episode = 1, 2, ..., M **do**

 Initialise state s_1 and encode into the quantum state

for $t = 1, 2, \dots, T$ **do**

 With probability ϵ select a random action a_t

 otherwise select $a_t = \max_a Q^*(s_t, a; \theta)$ from the output of the quantum circuit

 Execute action a_t in emulator and observe reward r_t and next state s_{t+1}

 Store transition (s_t, a_t, r_t, s_{t+1}) in \mathcal{D}

 Sample random minibatch of transitions (s_j, a_j, r_j, s_{j+1}) from \mathcal{D}

 Set $y_j = \begin{cases} r_j & \text{for terminal } s_{j+1} \\ r_j + \gamma \max_{a'} Q(s_{j+1}, a'; \theta) & \text{for non-terminal } s_{j+1} \end{cases}$

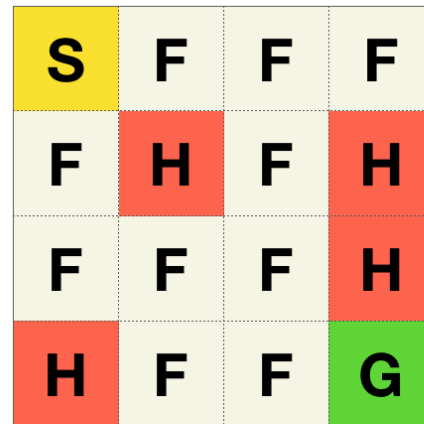
 Perform a gradient descent step on $(y_j - Q(s_j, a_j; \theta))^2$

end for

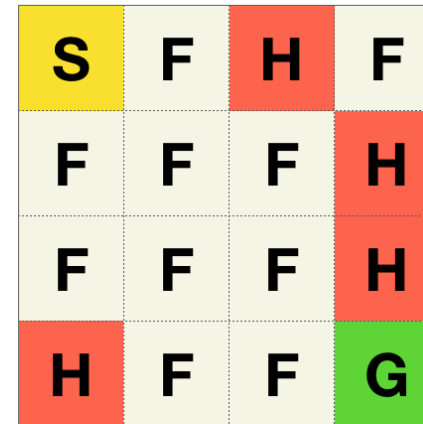
end for

Quantum Deep Q-Learning

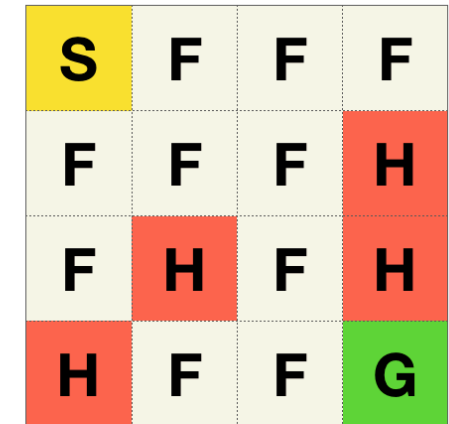
- Env: FrozenLake
- 16 discrete states
- 4 actions



(a)



(b)



(c)

Location	Reward
HOLE	-0.2
GOAL	1.0
OTHER	-0.01

Environment with 16 states.

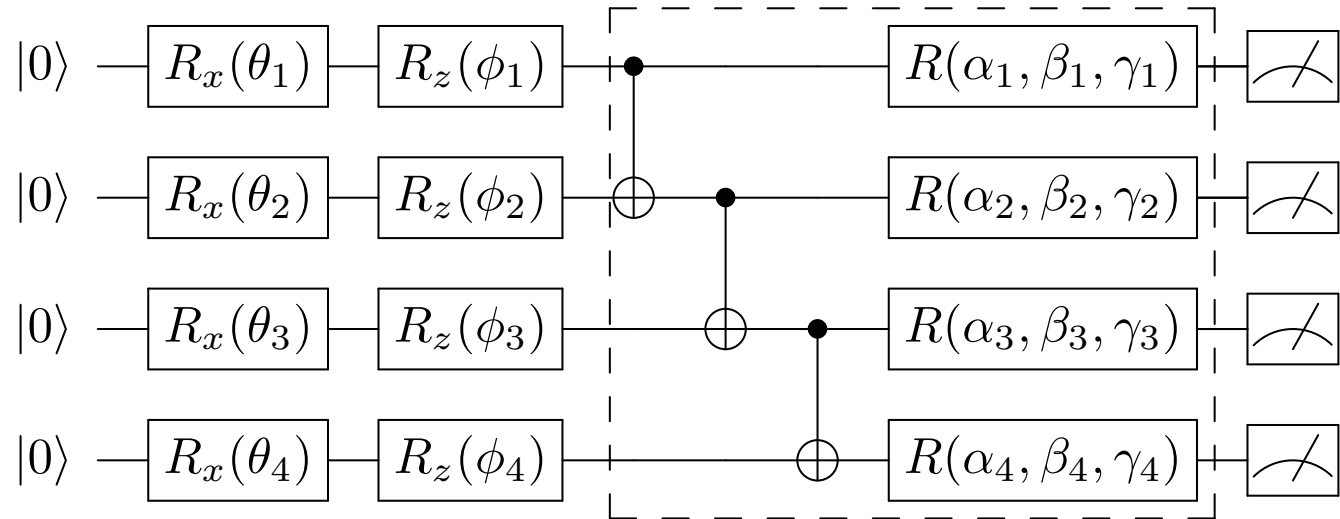
States numbered as 0~15

Example: State 12 : 1100 ->1,1,0,0

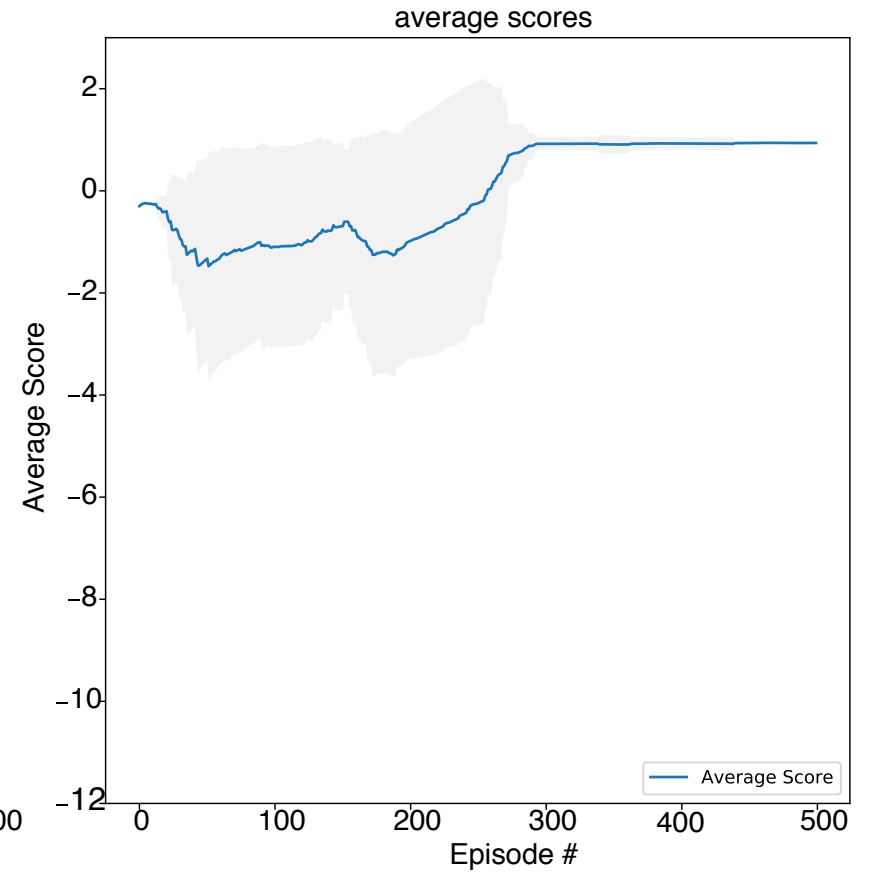
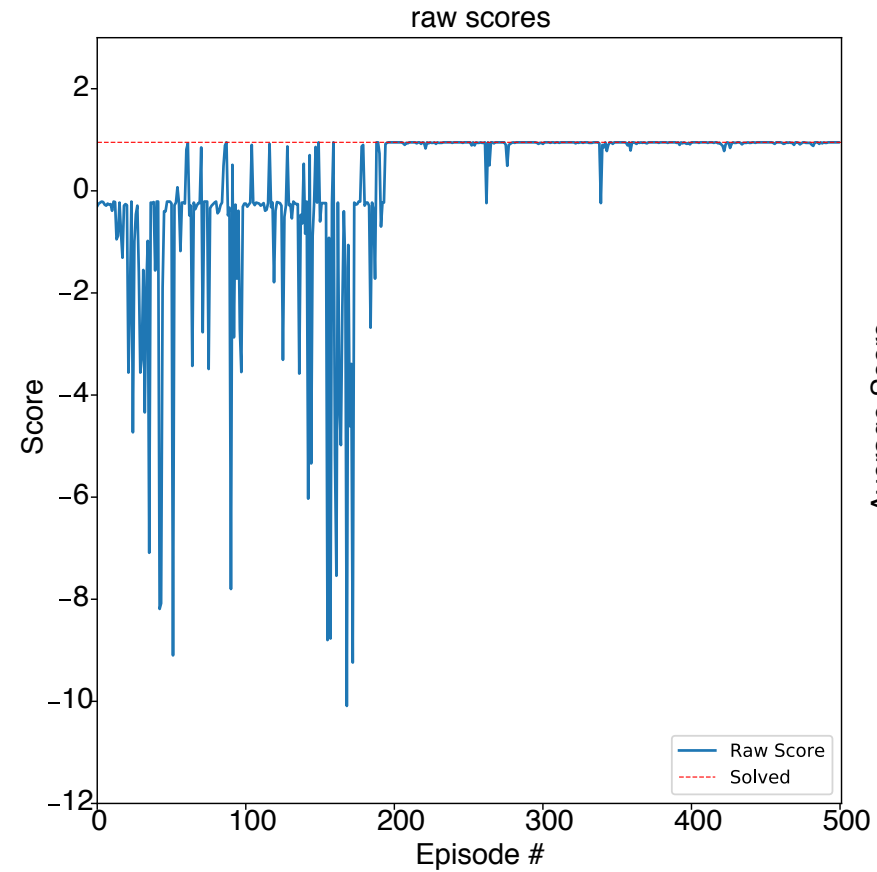
Rotation : $\theta_i = \pi \times b_i$

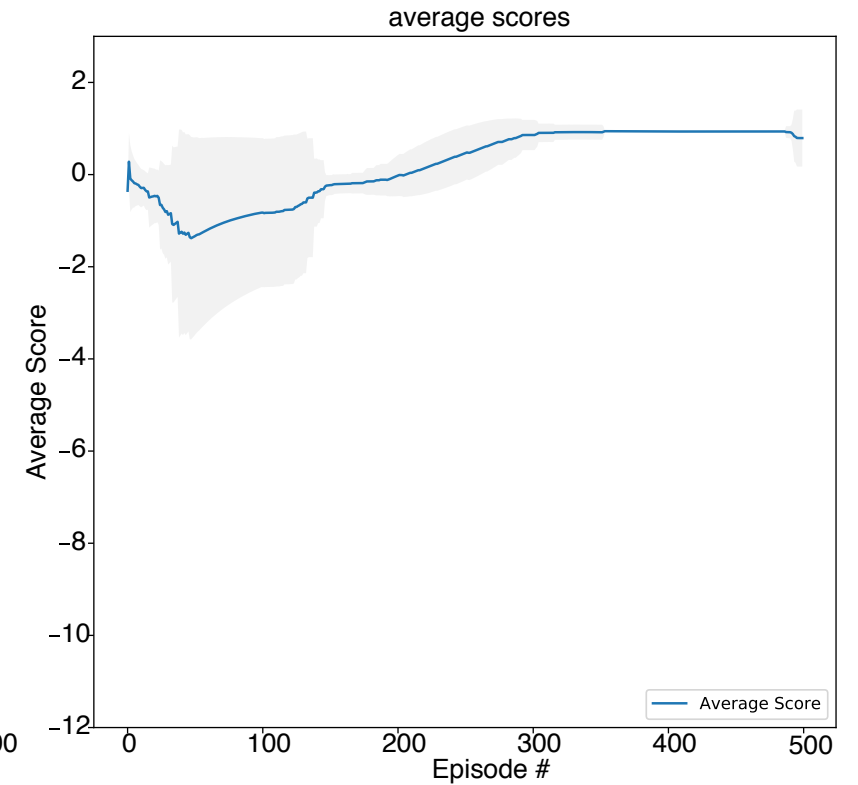
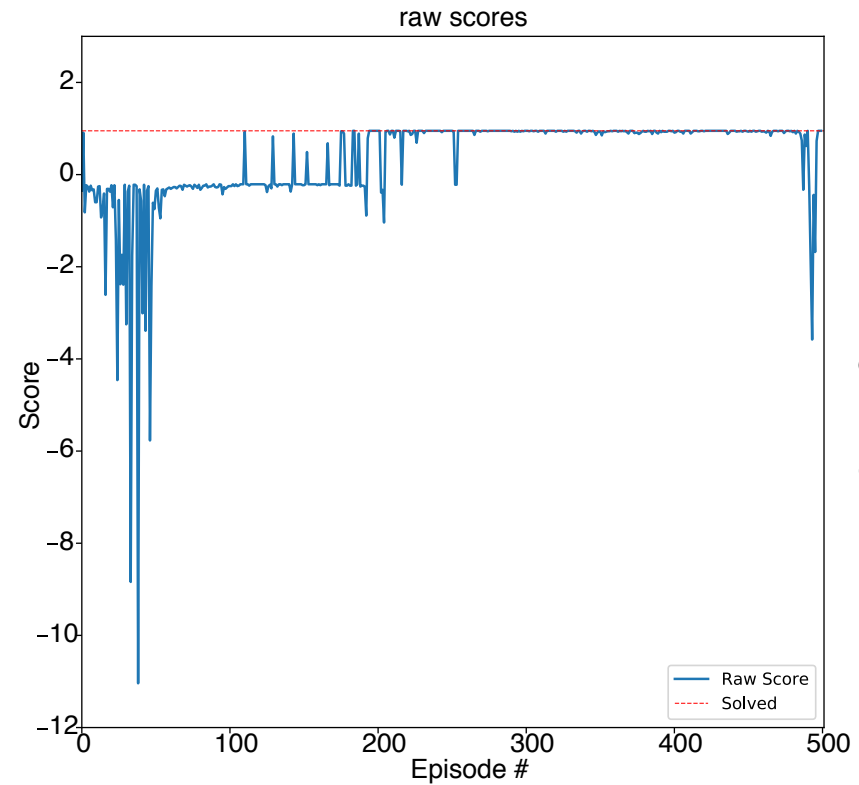
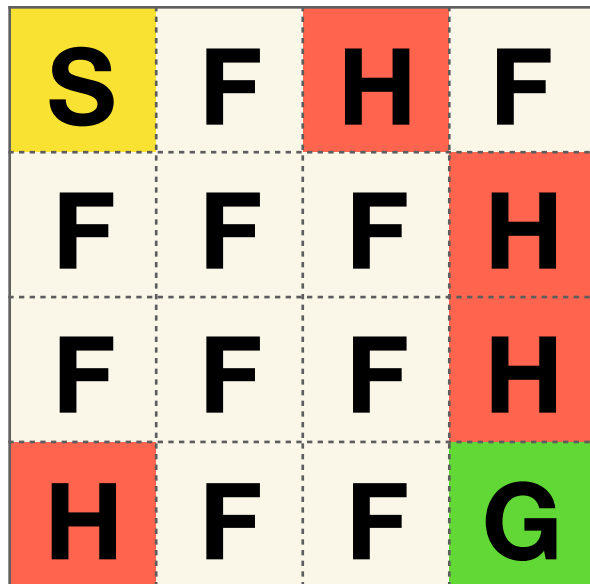
$\phi_i = \pi \times b_i$

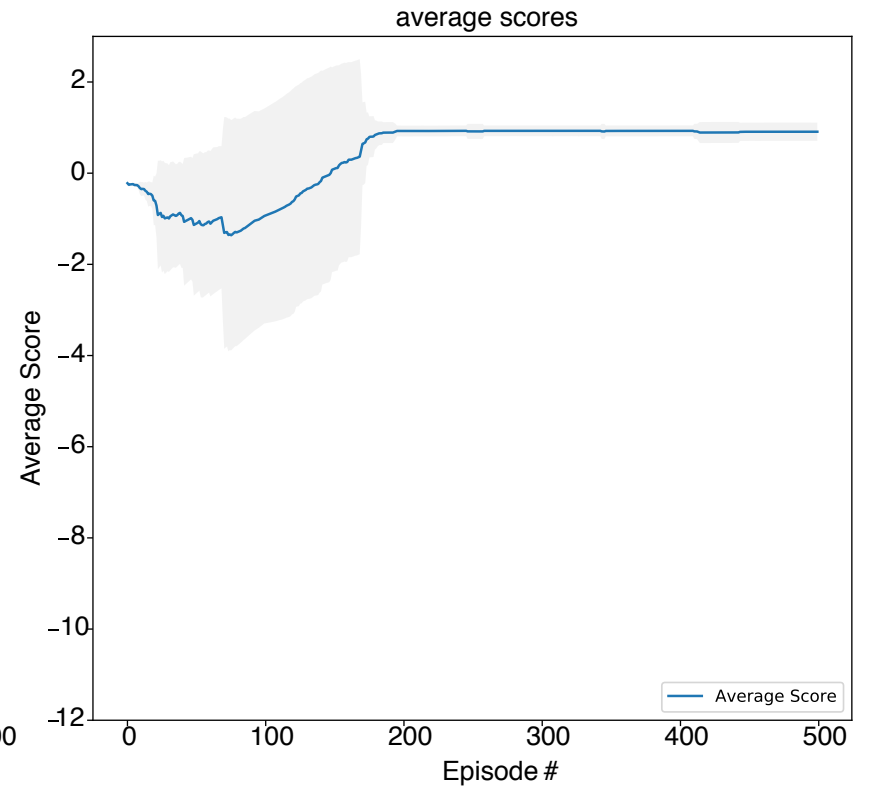
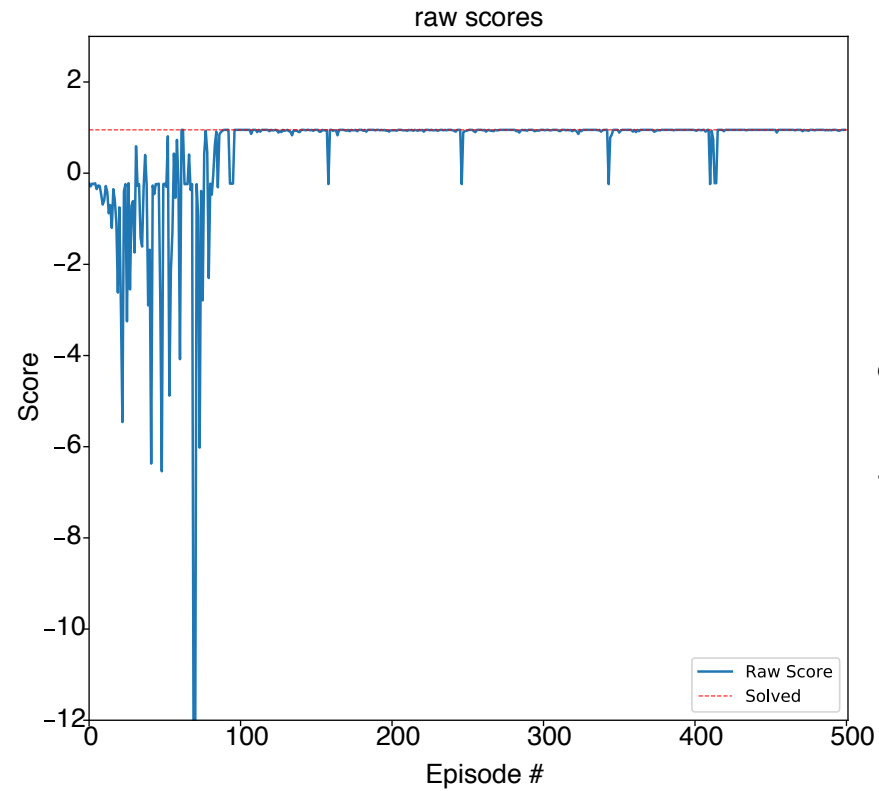
Result : $|1\rangle \otimes |1\rangle \otimes |0\rangle \otimes |0\rangle$



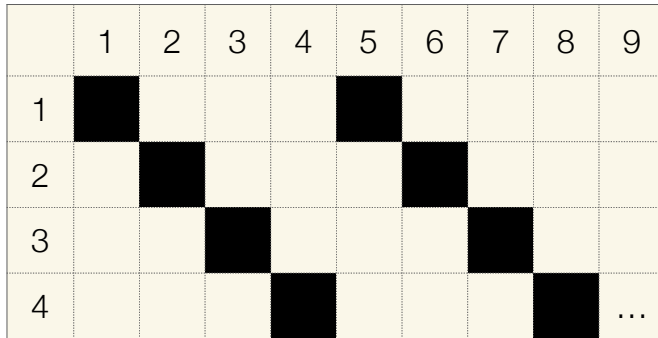
S	F	F	F
F	H	F	H
F	F	F	H
H	F	F	G



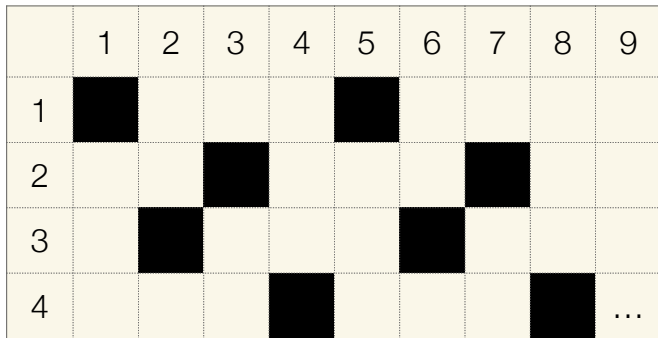




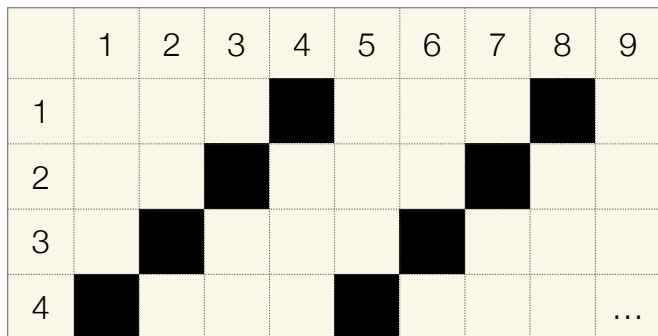
- Env: CognitiveRadio



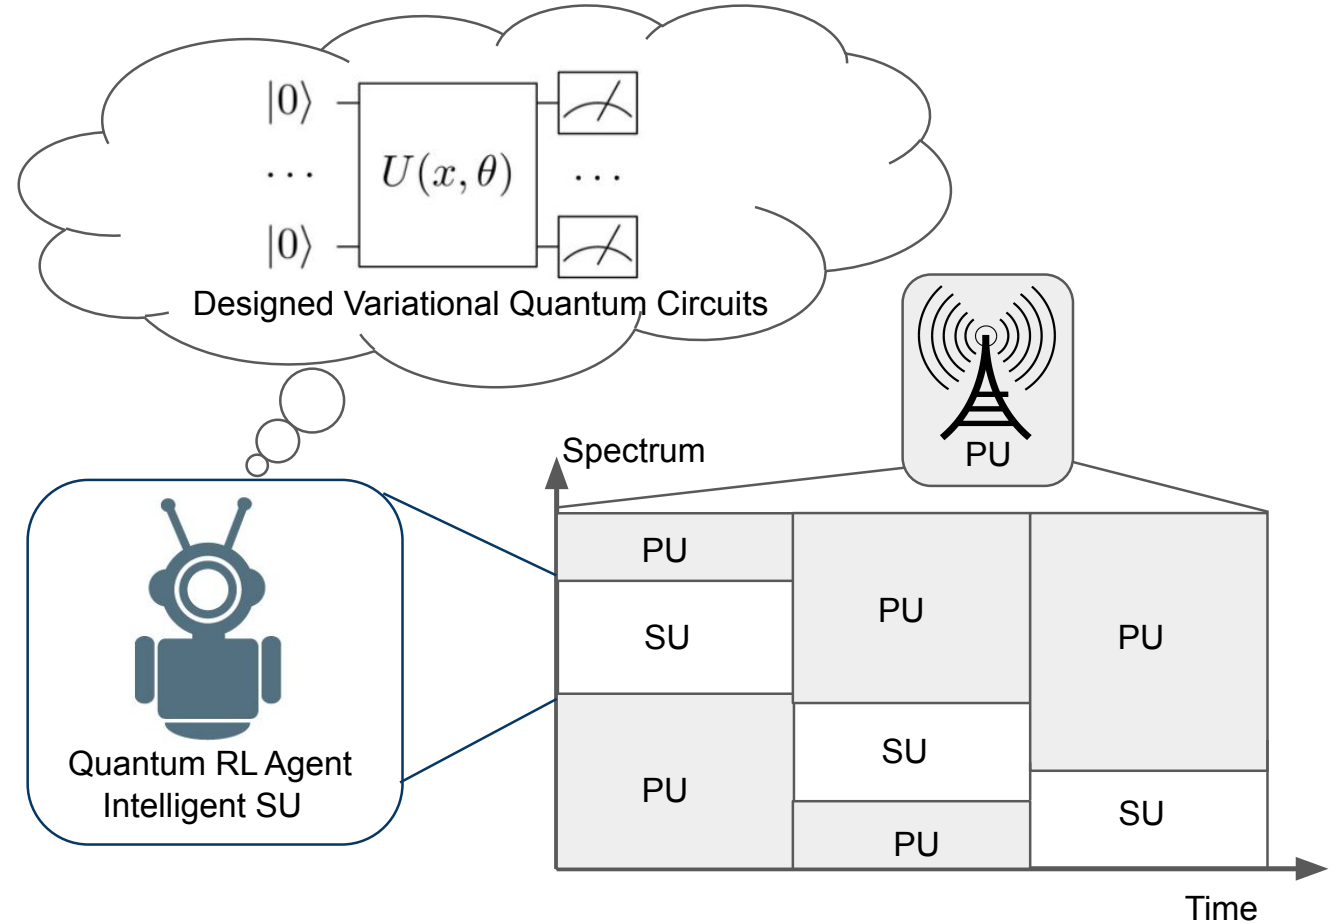
(a)



(b)

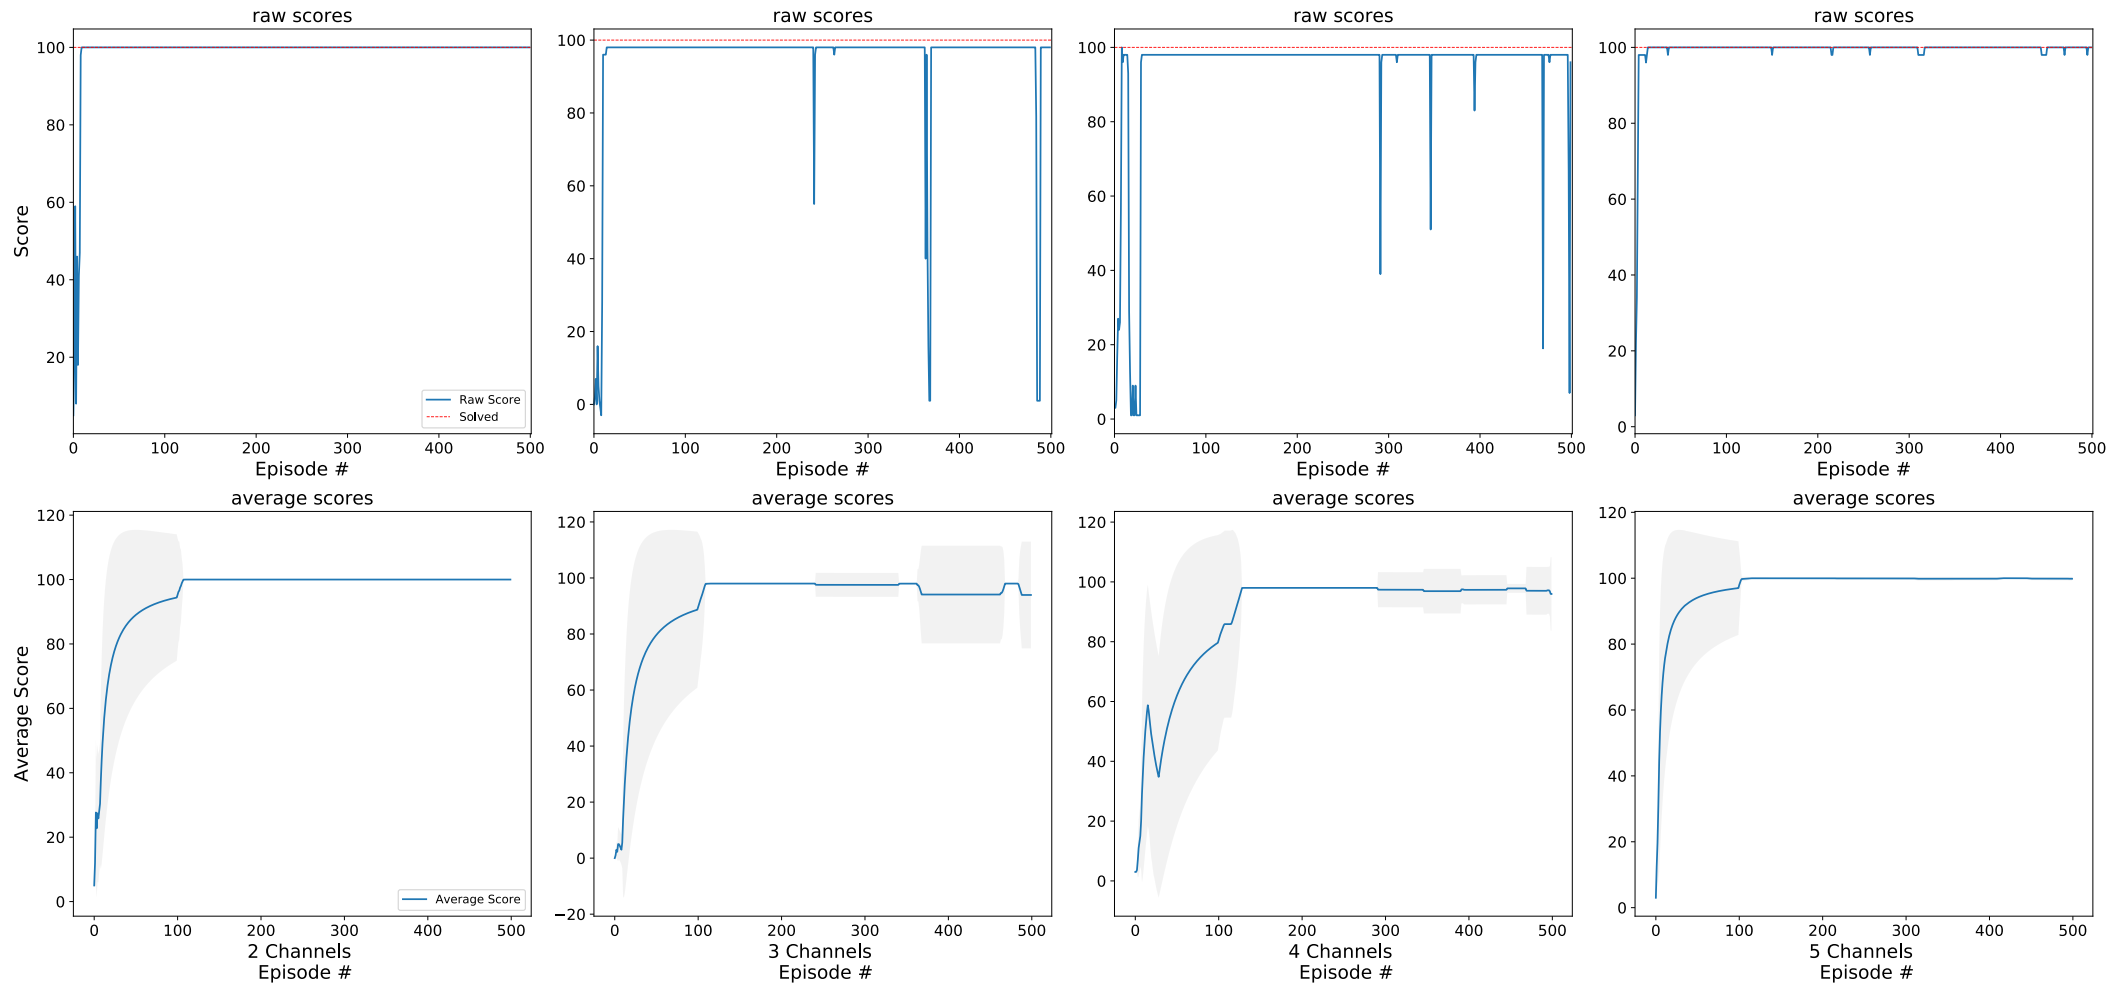


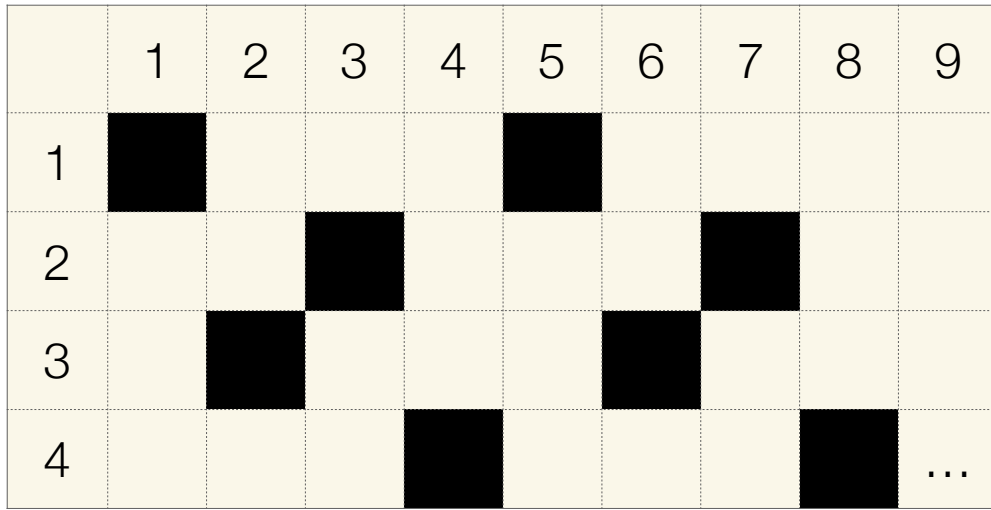
(c)



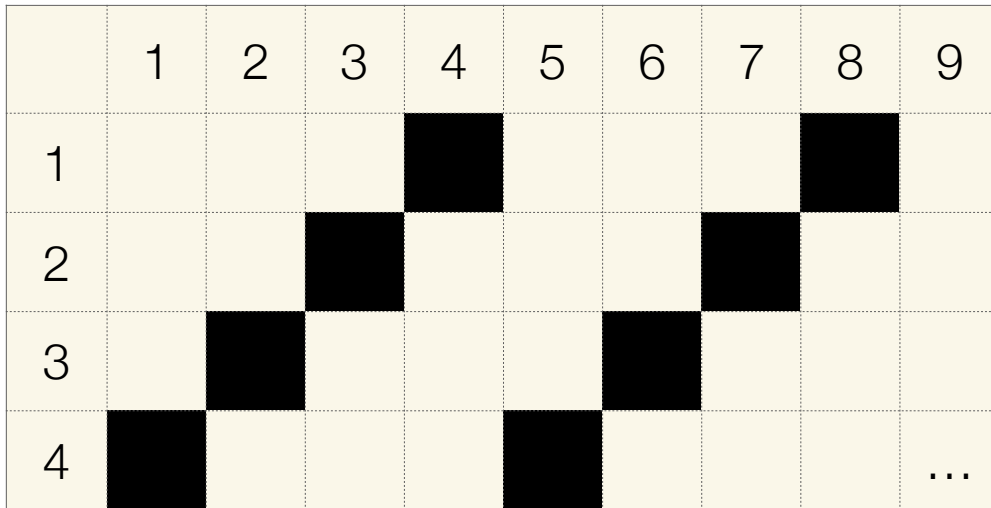
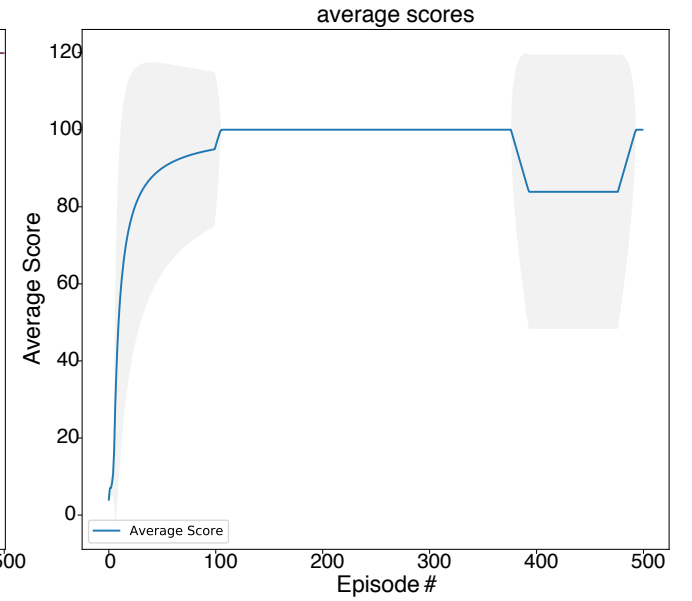
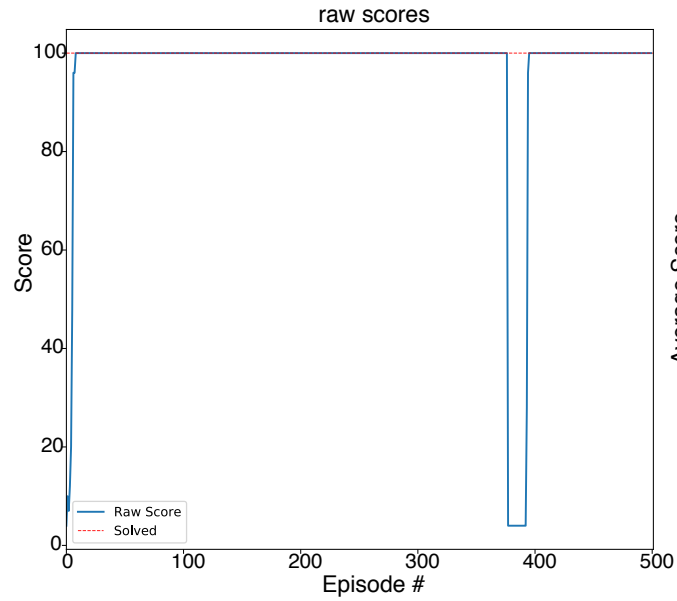
	1	2	3	4	5	6	7	8	9
1	■				■				
2		■				■			
3			■				■		
4				■				■	...

(a)

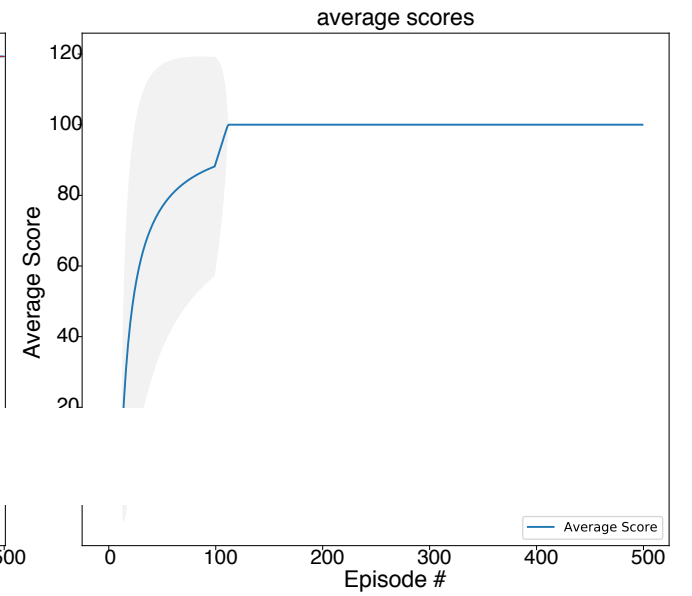
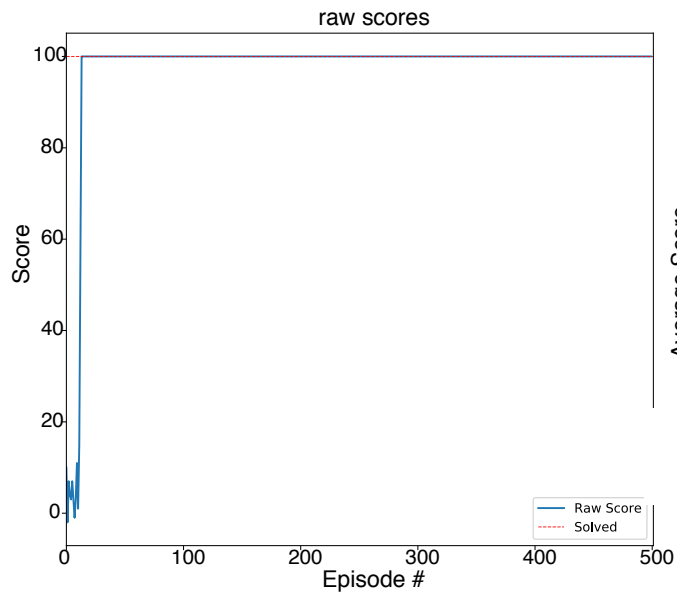


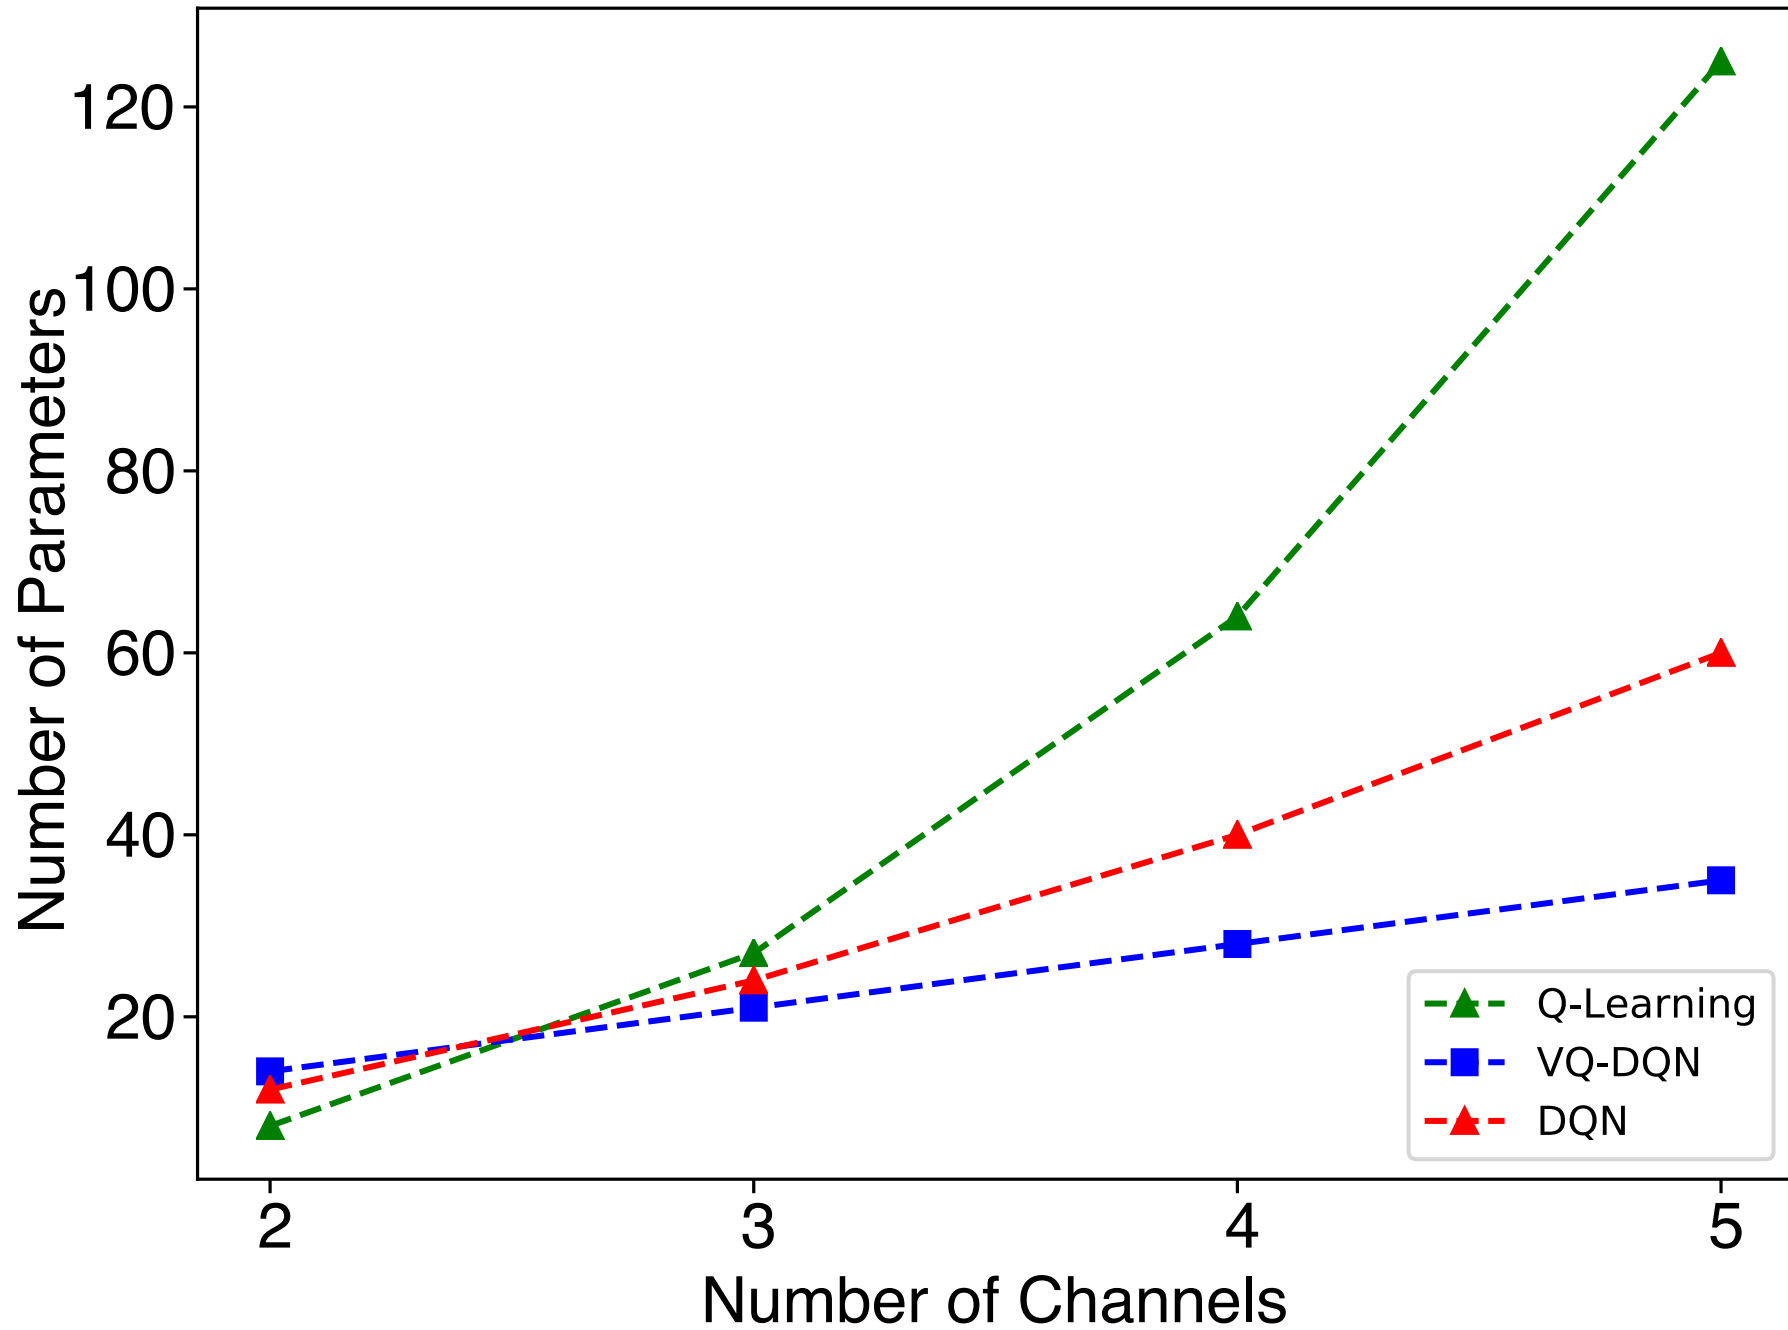


(b)



(c)





Evolutionary Quantum RL

- Why?
 - Gradient-based methods may suffer from local optima.
 - Certain QRL models are difficult to train via gradient-based methods.
 - In classical RL, evolutionary optimization can beat gradient-based methods in some hard tasks.

Evolutionary Optimization

- **Initialization:**

Initialize the population \mathcal{P} of N agents with each of them given randomly generated initial parameters θ , which are sampled from $\mathcal{N}(0, I)$

- **Running and evaluating the agents:**

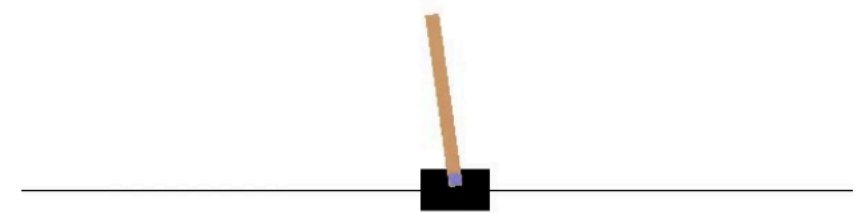
- Each agent plays the game R_1 times and get the average score $S_i^{avg} = \frac{1}{R_1} \sum_{r=1}^{R_1} S_{i,r}$
- Top T agents are selected to be the *parents* to generate the next generation

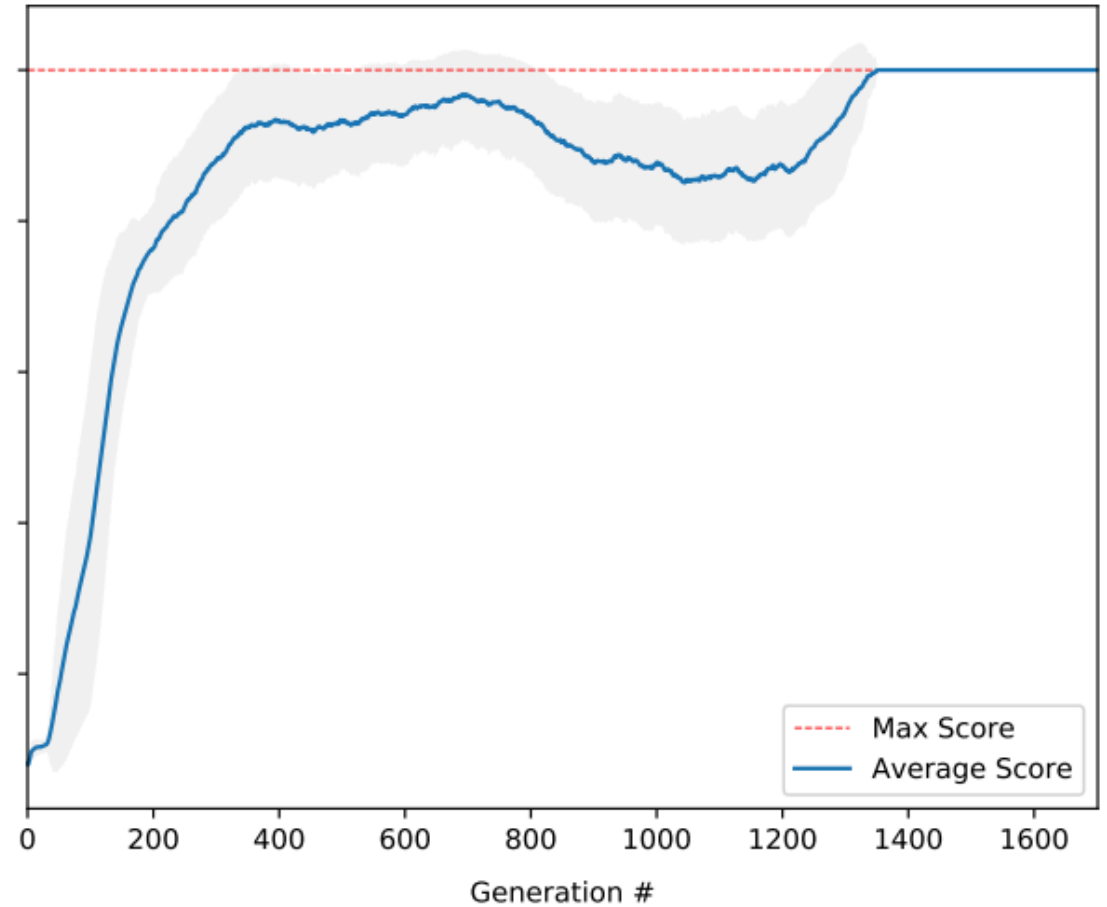
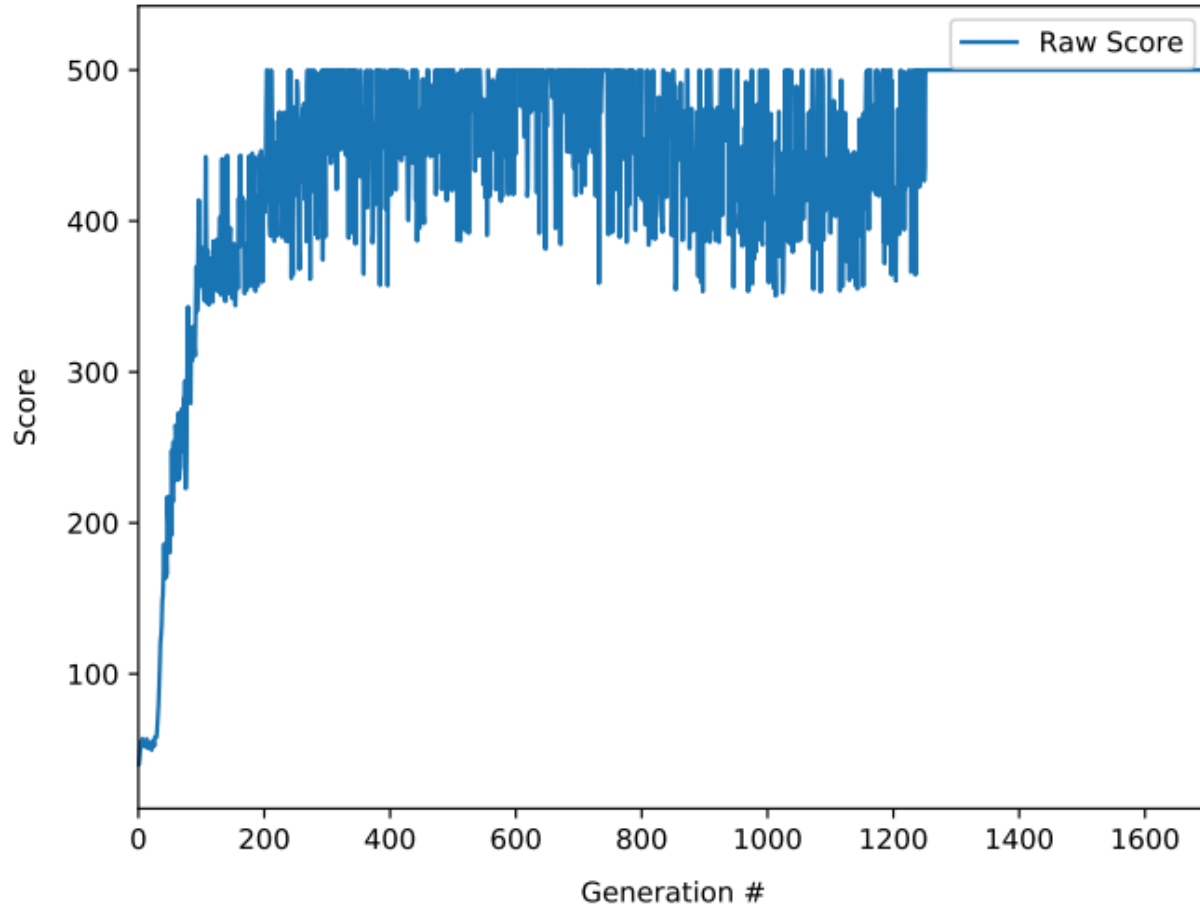
- **Mutation and the next generation:**

- $N - 1$ children: Each child is generated via a randomly selected agent from the parent group and slightly mutated according to $\theta \leftarrow \theta + \sigma \epsilon$ where σ is the mutation power and ϵ is the Gaussian noise
- The *elite* or N^{th} - child is the best performing from the parent group

Environments-CartPole

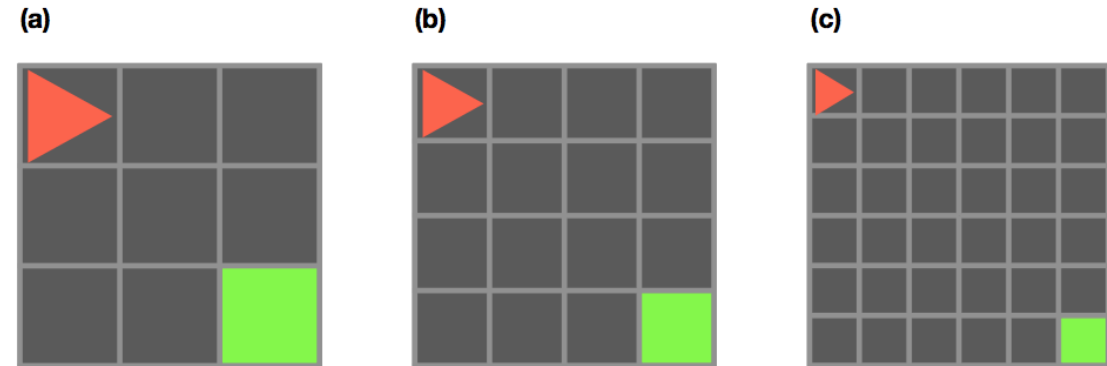
- **Observation:** A four dimensional vector s_t comprising values of the cart position, cart velocity, pole angle and pole velocity at the top.
- **Action:** There are two actions: pushing to the *right* or *left*.
- **Reward:** A reward +1 is given for every time step where the pole close to being upright.



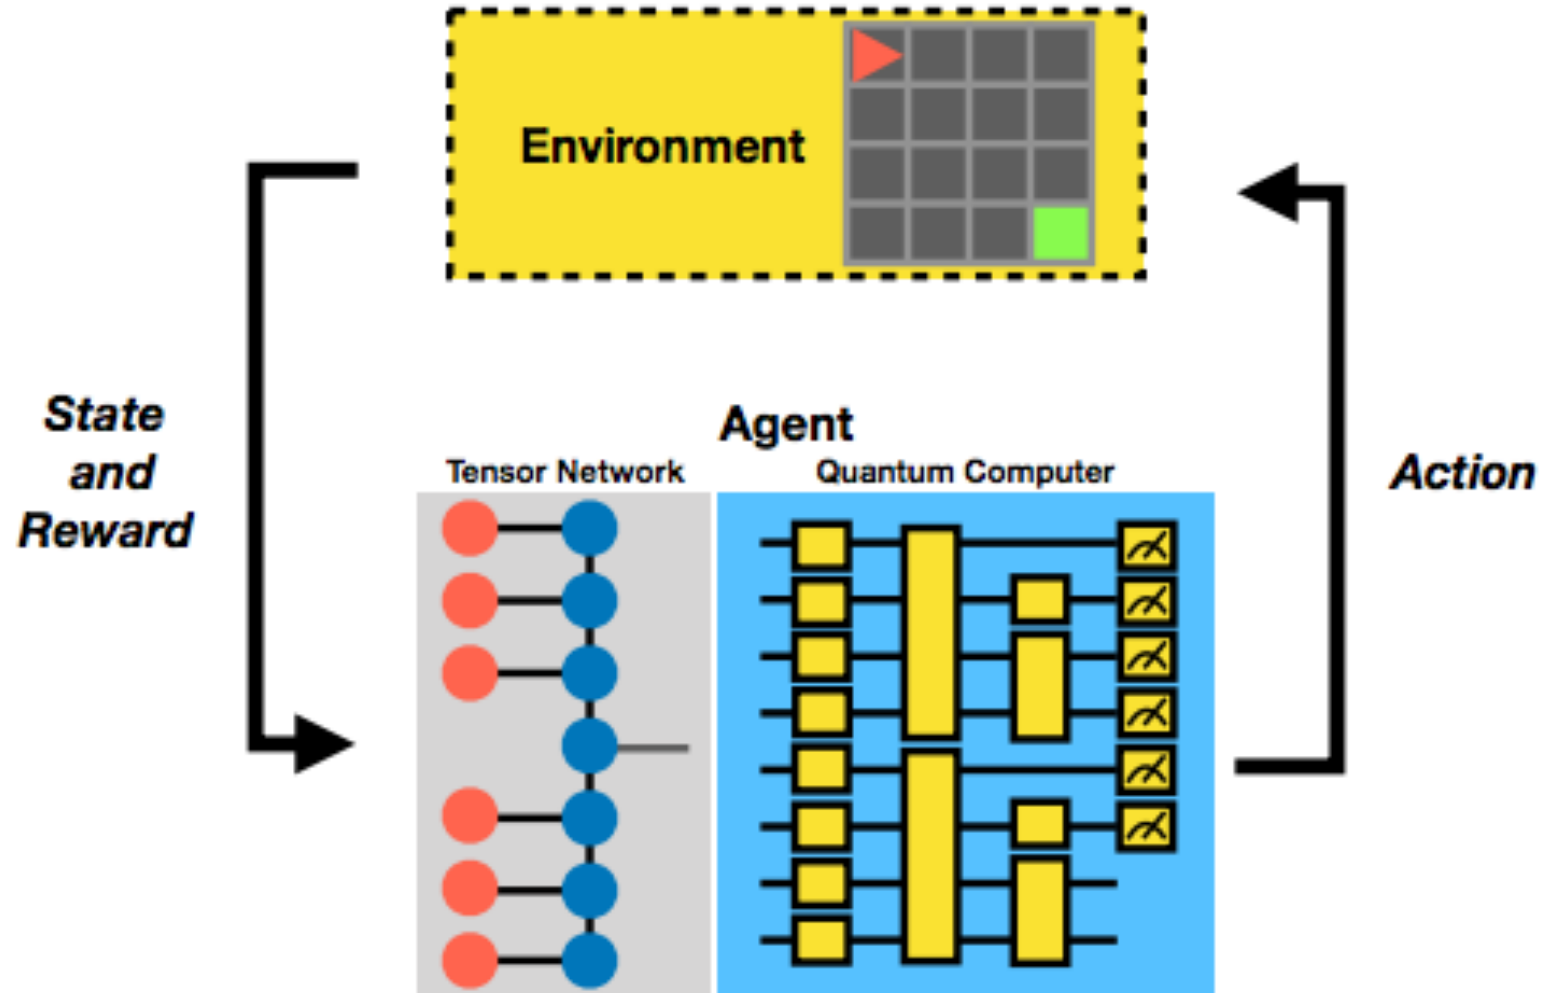


Environments-MiniGrid

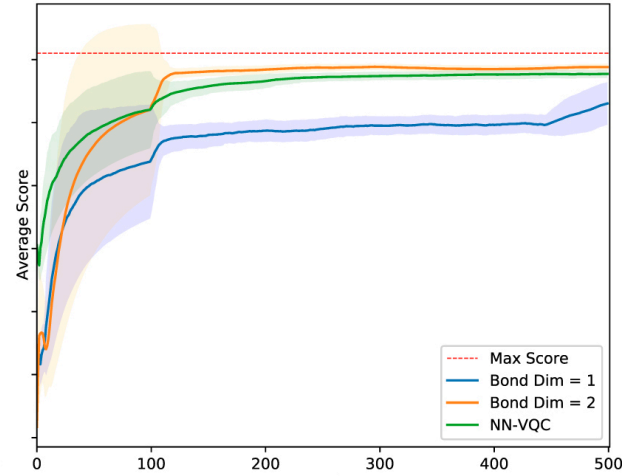
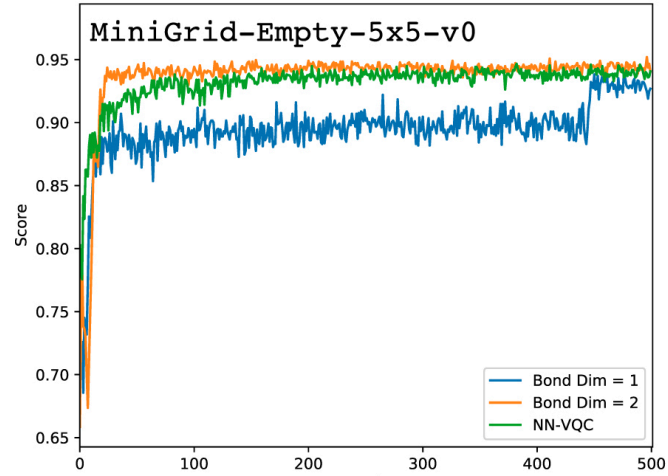
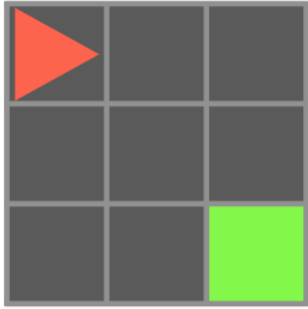
- **Observation:** A 147 dimensional vector s_t
- **Action:** There are 6 actions:
 - Turn left
 - Turn right
 - Move forward
 - Pick up an object
 - Drop the object
 - Toggle
- **Reward:** A reward of 1 is given when the agent reaches the goal. A penalty is subtracted from the reward according to:
 $1 - 0.9 \times (\text{number of steps}/\text{max steps allowed})$



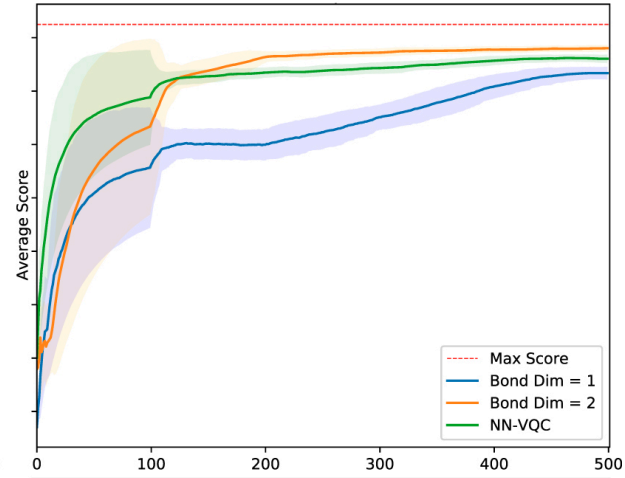
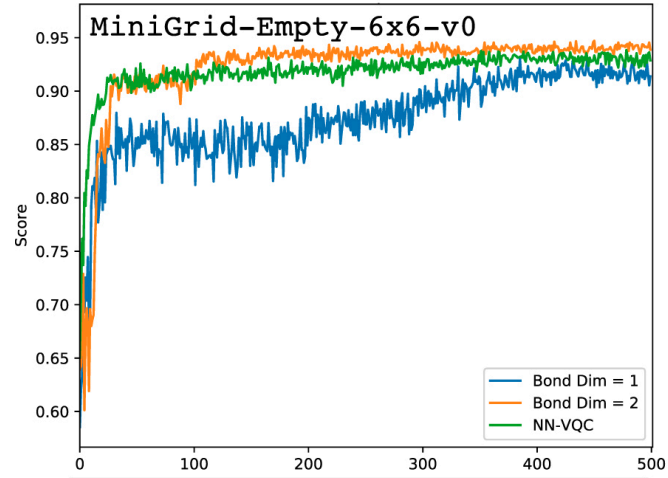
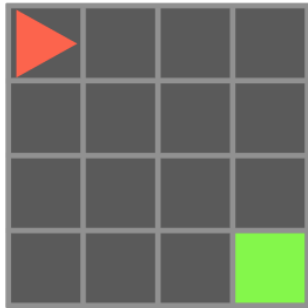
Hybrid TN-VQC model



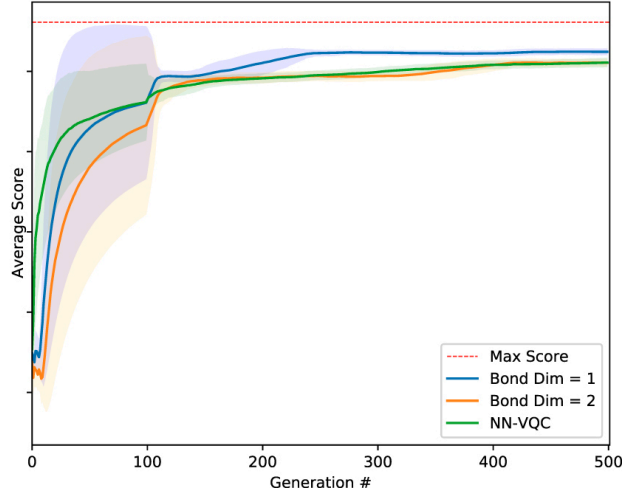
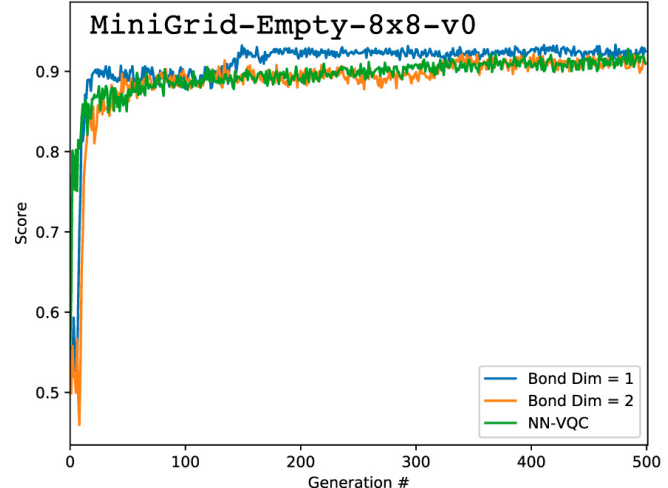
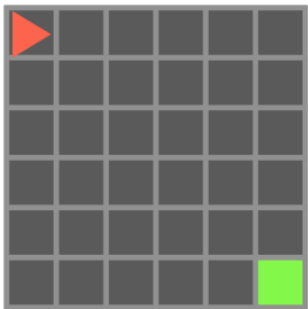
(a)



(b)



(c)



- I. Introduction
- II. Quantum Computing (QC)
- III. Quantum Machine Learning
- IV. Applications-Classification
- V. Applications-Sequential Learning
- VI. Applications-Reinforcement Learning
- VII. Conclusion and Outlook**

- Quantum encoding / embedding methods are critical.
- Using VQC to replace NN in CNN can improve the performance when the number of parameters is similar.
- With careful design, quantum reinforcement learning can learn a similar task with fewer model parameter.
- QML models can be trained in federation to preserve data privacy.

Acknowledgements

- Brookhaven National Laboratory LDRD #20-024
- U.S. Department of Energy, Office of Science, DE-SC-0012704
- U.S. Air Force Office of Scientific Research, FA2386-20-1-4033
- Ministry of Science and Technology (MOST) of Taiwan:
 - 107-2112-M-002-016-MY3
 - 108-2112-M-002-020-MY3
 - 109-2112-M-002-023-MY3
 - 109-2627-M-002-003
 - 107-2627-E-002-001-MY3
 - 109- 2622-8-002-003
- National Taiwan University grant No. NTUCC-110L890102