Quantum noise modeling through Reinforcement Learning

Simone Bordoni
Andrea Papaluca
Piergiorgio Buttarini
Alejandro Sopena
Stefano Carrazza
Stefano Giagu

Preprint:

simone.bordoni@tii.ae

Introduction
Machine learning and quantum computing

- Parametrized quantum circuits
- Quantum anomaly detection
- Quantum generative algorithms
- Speedup of ML training routines

Machine Learning

Quantum computing

- Quantum error correction
- **Noise modelling**
- Transpiling
- Optimal control

Introduction
Table of contents

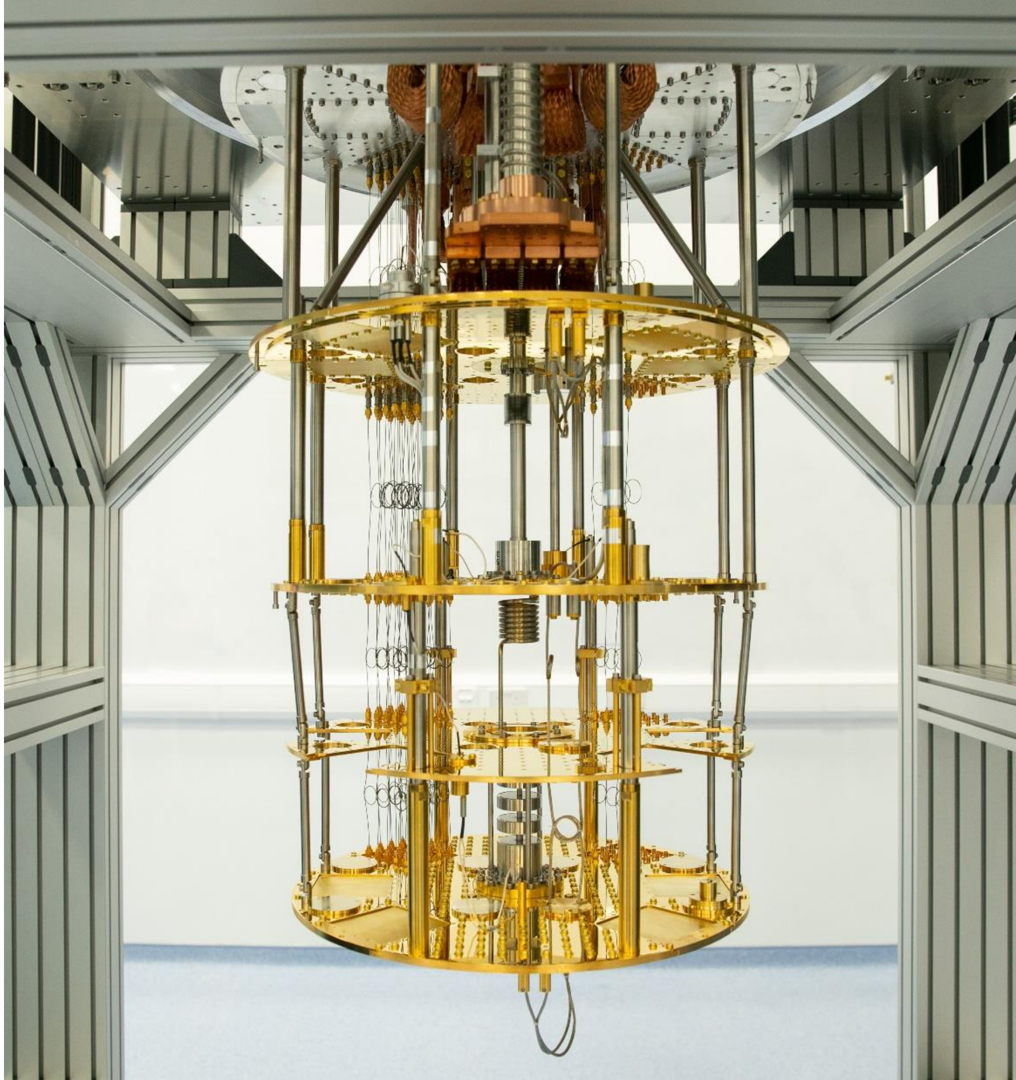

Reinforcement learning
Working principle

Markov decision process: the future state depends solely on the current **state** and **action**, regardless of the history of the system.

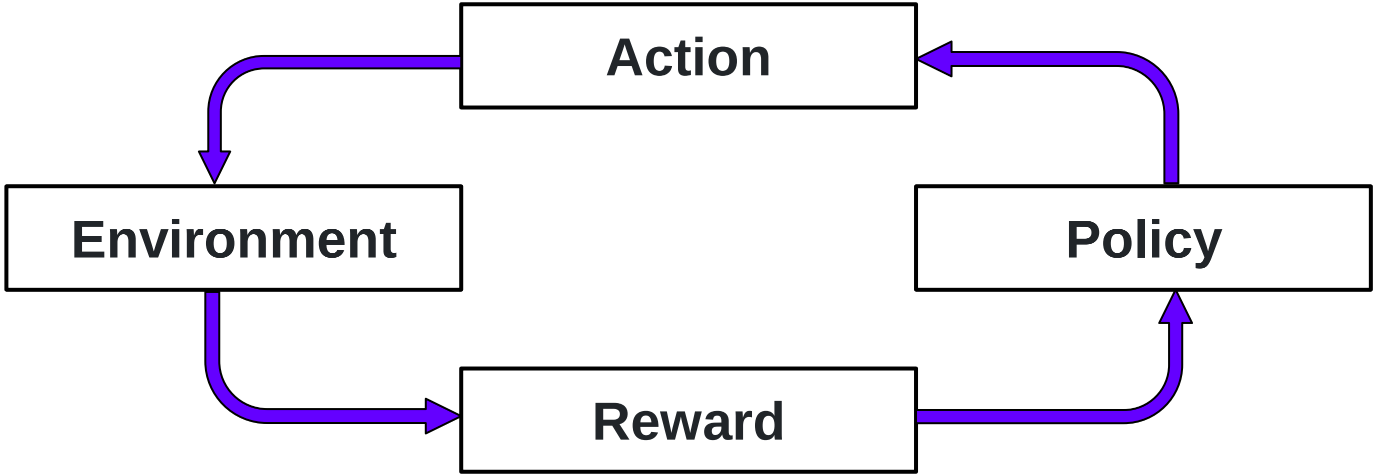

Training a reinforcement learning **agent** involves finding an **optimal policy** that maximizes the expected long-term **cumulative reward**.

Reinforcement learning
Policy training

The **policy** can be implemented with an artificial **neural network** (in this work we use CNN).
Weights of the policy are updated using standard **gradient based optimization** (backpropagation).

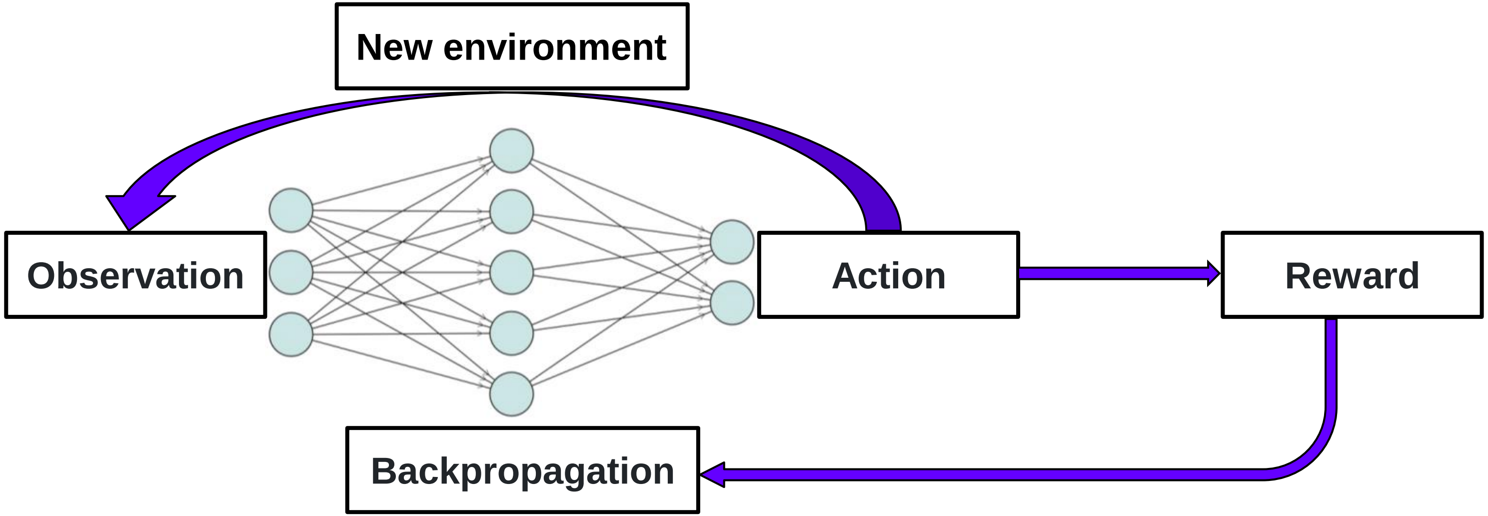

High **instability** during training. To address this issue many training strategies have been developed.
In this work we use **Proximal Policy Optimization** (PPO) [arXiv:1707.06347].

Noise in quantum circuits
Causes

Noise is caused by many possible factors:
- Environment interactions.
- Calibration errors.
- Crosstalk.
- Excited states decay.
- Measurements errors.

Different nature and characteristics of the different noise sources makes the total **noise difficult to characterize and reproduce in simulations**.

Euristic noise models use a set of **noise channels** placed in the circuit to reproduce noisy behavior.
The parameters of these channels are fitted with the **calibration data** (T1, T2, gates fidelity…).

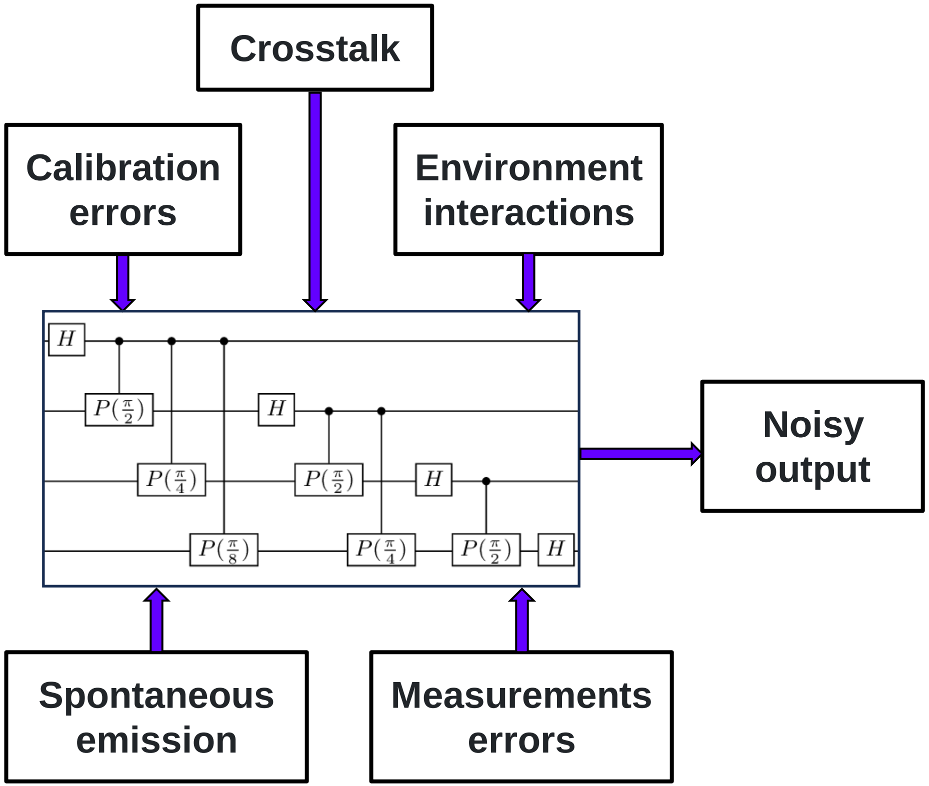

Noise in quantum circuits
Noise channels

Noise channels are **super-operators** acting on state density matrices.
Below the description of two noise channels used in this work:

Damping channel

Schematization of spontaneous decay of the excited state. The parameter ɣ controls the decay probability.

$$Damp(\gamma)\,|1\rangle = (1-\gamma)\,|1\rangle + \gamma\,|0\rangle$$

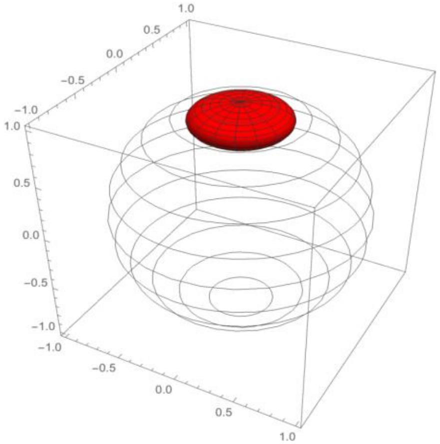

Depolarizing channel

Schematization of the interaction with environment that brings the state closer to the maximally mixed state. The parameter λ is the probability of an error occurring.

$$Dep_{\lambda}(\rho) = (1-\lambda)\rho + \lambda\mathbb{I}/2$$

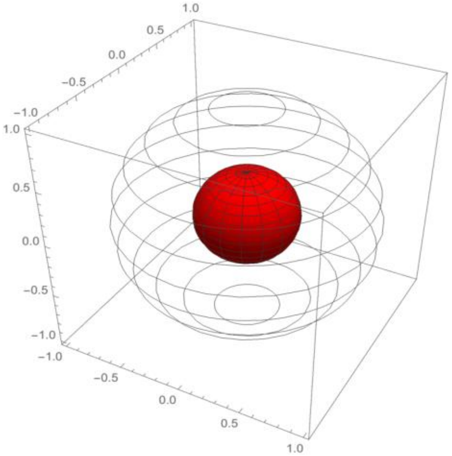

Noise in quantum circuits
Coherent errors

Coherent errors are **unitary**, they don't reduce the purity of quantum states.
These kind of errors don't require noise channels, they can be schematized using **rotation gates** (Rx, Ry, Rz).
These errors do not destroy quantum information, they **can be corrected** once identified.

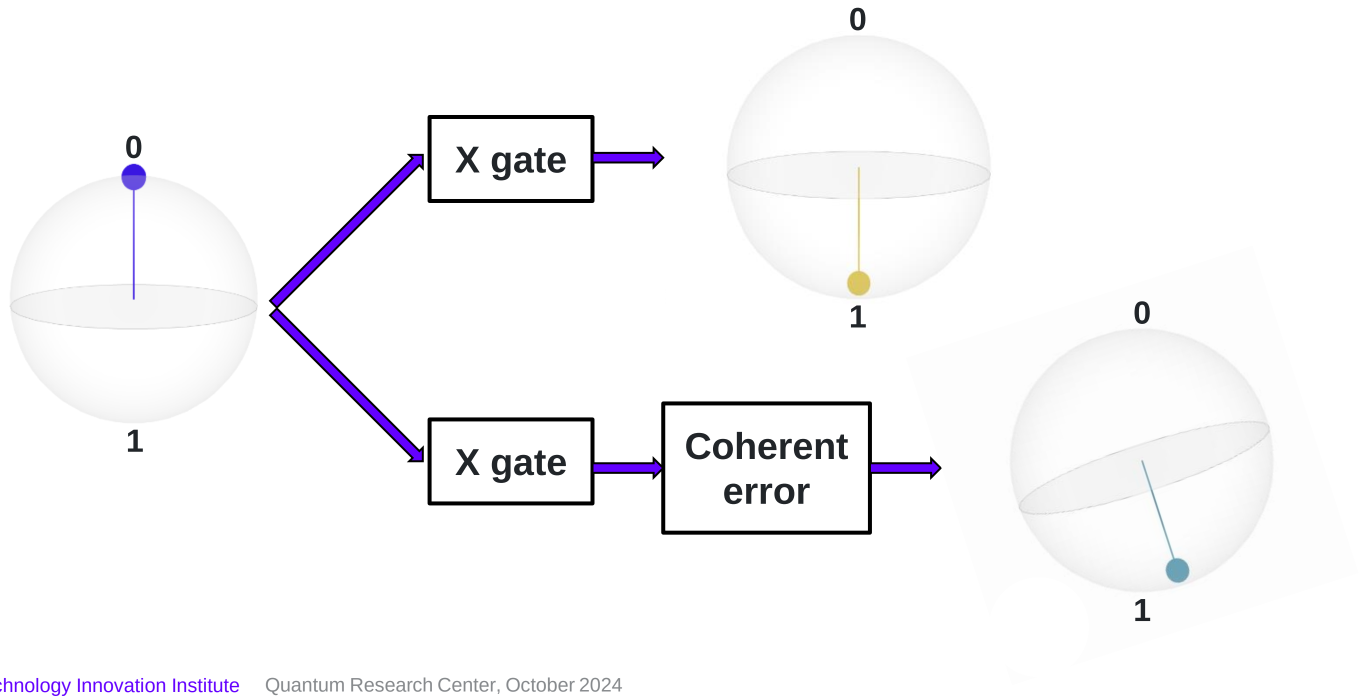

Methodology
Algorithm implementation

- A non noisy quantum circuit representation is given to the agent at each **episode**.
- For every circuit moment the agent can put any number of noise channels with a chosen parameter.
- At the end of an episode the Density Matrix (DM) of the circuit obtained with this process is computed.
- The **trace distance** (TD) between this DM and the ground truth DM is used to compute the reward.
- Weights of the policy NN are update to maximize fidelity.

After many episodes the agent should learn where to **put noise channels in a non-noisy circuit** to reconstruct the DM of the real noisy circuit.

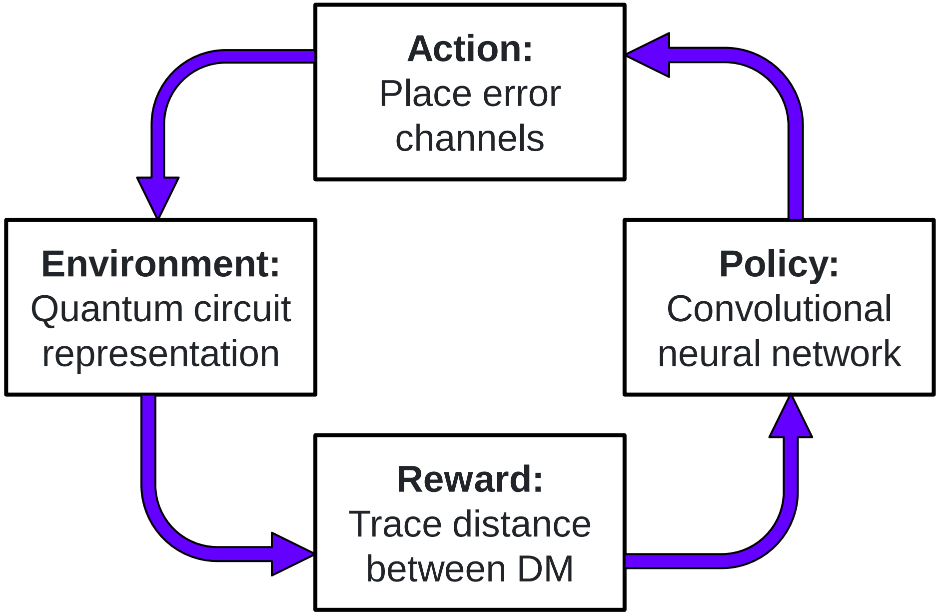

Methodology
Quantum circuit representation

To train the RL agent it is necessary to **represent the quantum circuit as an array**, that can be used as input of the policy neural network. In the general case this vector has three dimensions: the first entrance identifies the qubit, the second the circuit moment and the third the features of gates and noise channels.

Features are organized as follows:

- Presence of single qubit gates.
- Presence of two qubit gates.
- Rotation angles of single qubit gates.
- Depolarizing channel parameter.
- Amplitude damping channel parameter.
- Coherent error parameters.

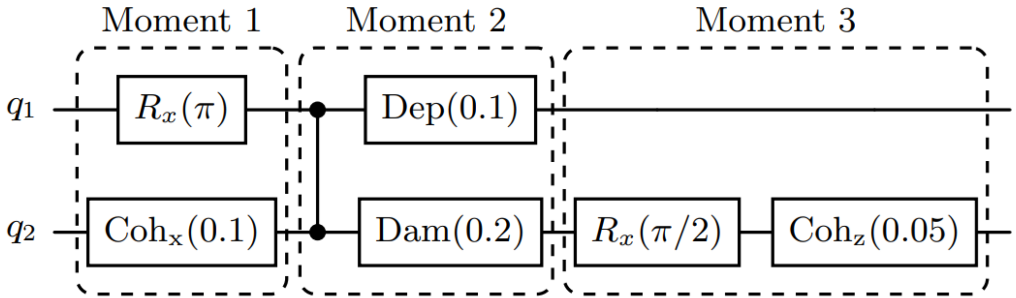

		R_z	R_x	CZ	θ	Dep	Dam	$\mathrm{Coh_z}$	$\mathrm{Coh_x}$
Moment 1	q_1	0	1	0	0.5	0	0	0	0
	q_2	0	0	0	0	0	0	0	0.1
Moment 2	q_1	0	0	1	0	0.1	0	0	0
	q_2	0	0	1	0	0	0.2	0	0
Moment 3	q_1	0	0	0	0	0	0	0	0
	q_2	1	0	0	0.25	0	0	0.05	0

Methodology
Reward shaping

The **reward** is the key element for good **convergence** of the RL algorithm during training.
Best results obtained using: $1/\left(TD(x,x')^2+\delta\right)$
where *x* and *x'* are the result and ground truth density matrices, δ avoids infinite values.
On **quantum hardware** the ground truth DM can be computed using **state tomography**.

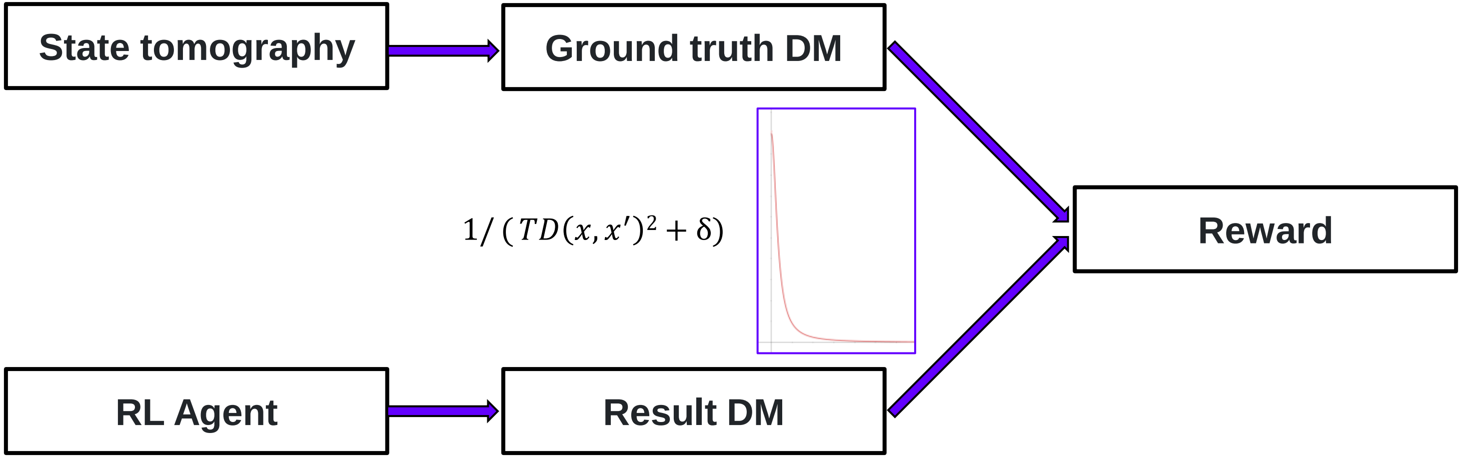

Results
Training

Top: 1 Qubit.
Bottom: 3 Qubits.

Custom noise model**:**
- Depolarizing channels.
- Damping channels.
- Coherent RX and RZ.

Dataset:
- Clifford circuits.
- Train set 400 circuits.
- Test set 100 circuits.

Training and CNN:
- 3*10^6 timesteps.
- 2 convolutional layers + 2 dense layers.
- ReLU activation.
- PPO algorithm.

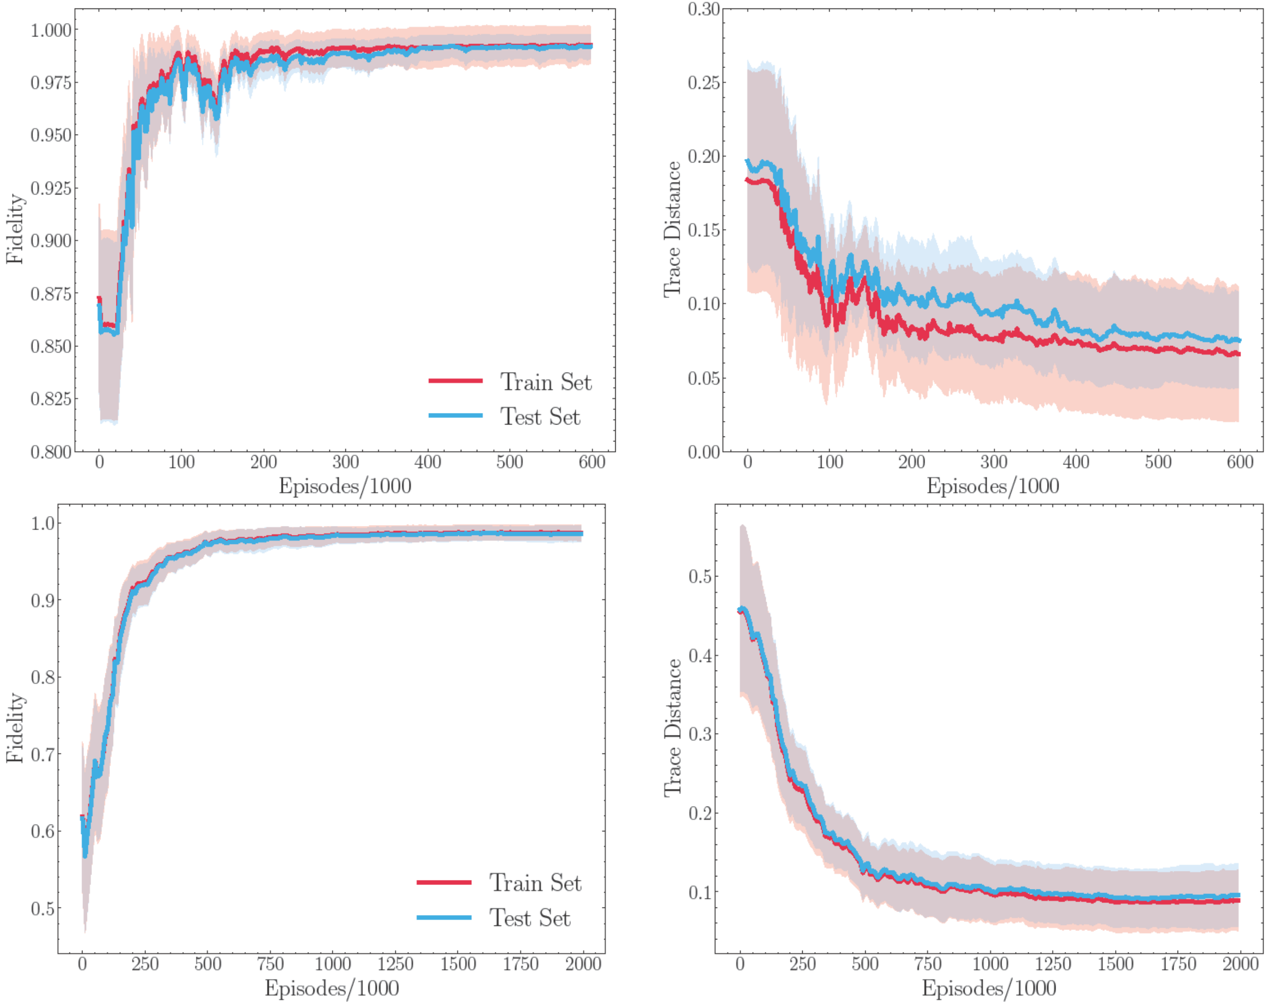

Results
Benchmarking

Performance comparison of RL agent and **Randomized Benchmarking** (RB) [arXiv:0707.0963] for circuits with different depths (from 3 to 30), also maximally mixed state (MMS) and circuits with no noise added.
1 Qubit (top), 3 Qubits (bottom).

RB considers all the noise sources as depolarizing. The RL agent shows better performance because it is able to **identify specific features of the noise**.

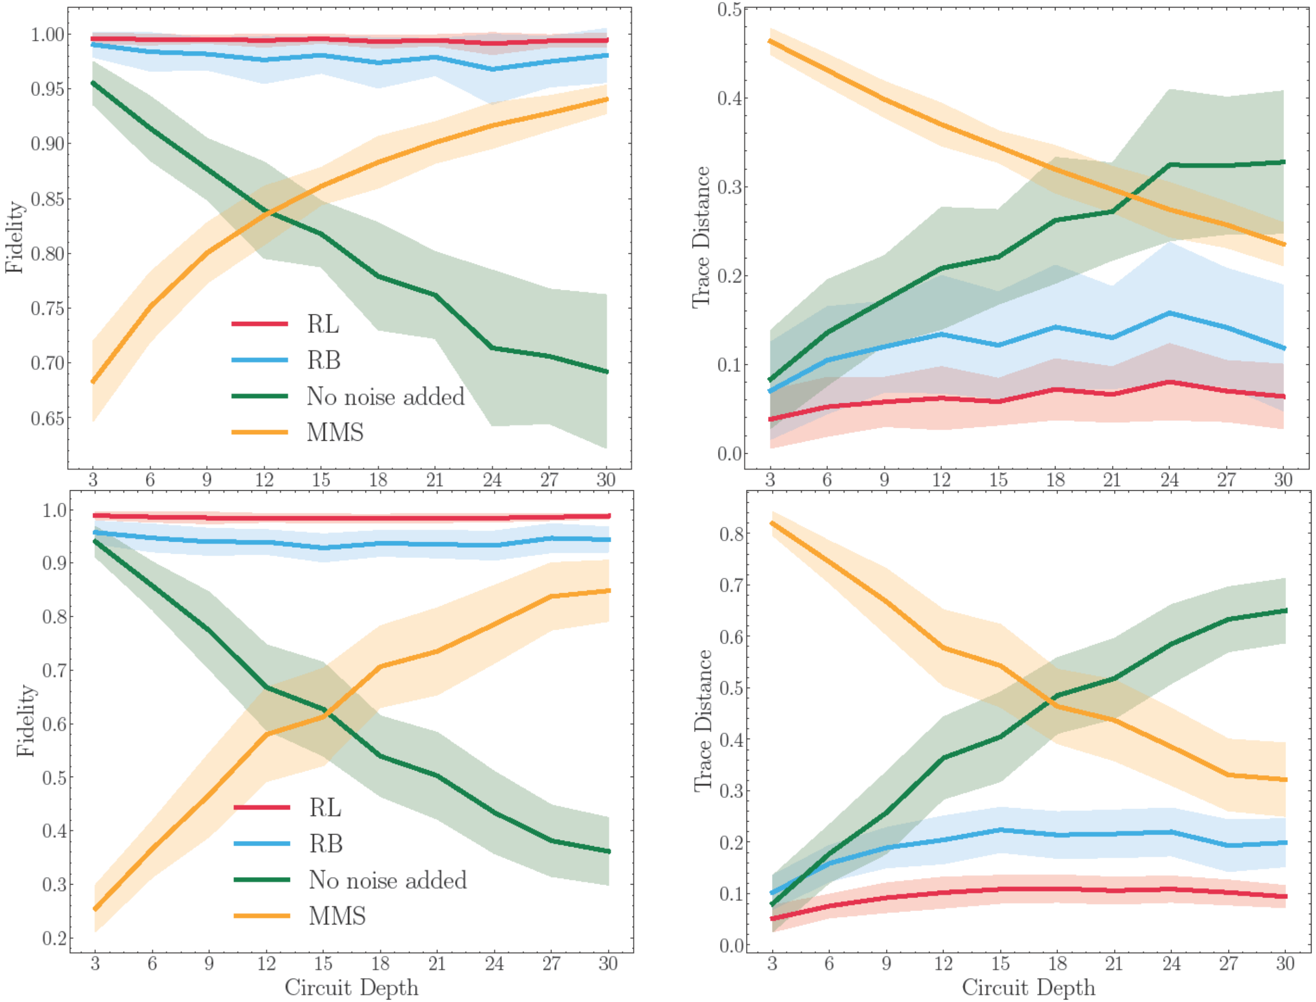

Results
Quantum algorithms

Top: Grover's algorithm (low noise).

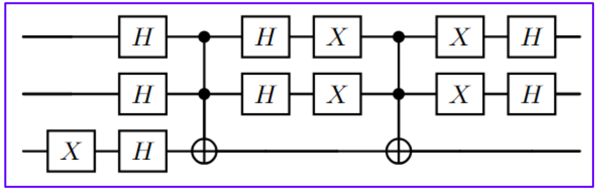

Bottom: QFT (high noise).

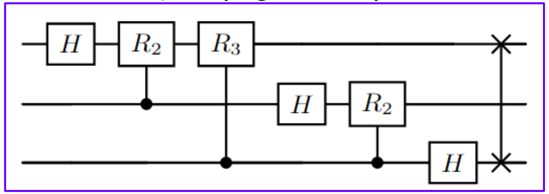

Circuit	Total gates	CZ gates	Moments
QFT	23	6	15
Grover	40	7	25

Circuit	Noise	RL	RB	No noise added	MMS
QFT	High	0.99	0.97	0.59	0.70
	Low	0.99	0.99	0.78	0.52
Grover	High	0.98	0.95	0.40	0.83
	Low	0.98	0.96	0.65	0.64

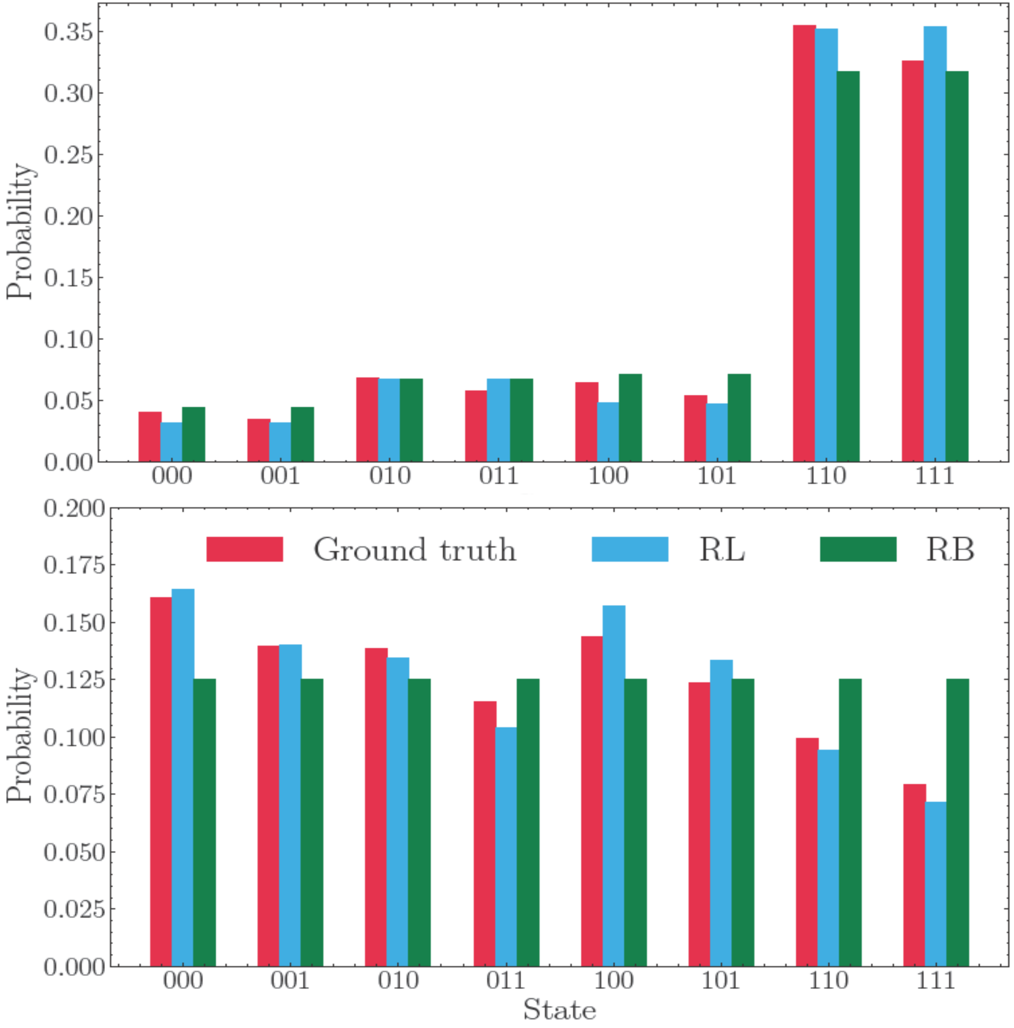

Results
Quantum hardware

Training on single qubit **superconducting chip** at **TII**.
Training set composed of 60 circuits of depth 10. Ground truth density matrices have been obtained with quantum **state tomography**, 2048 shots for circuit.

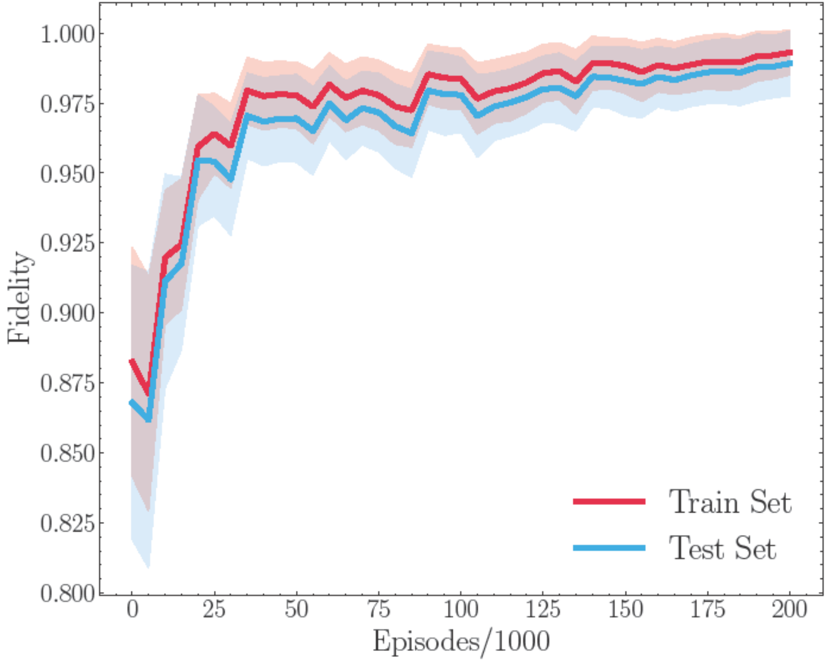

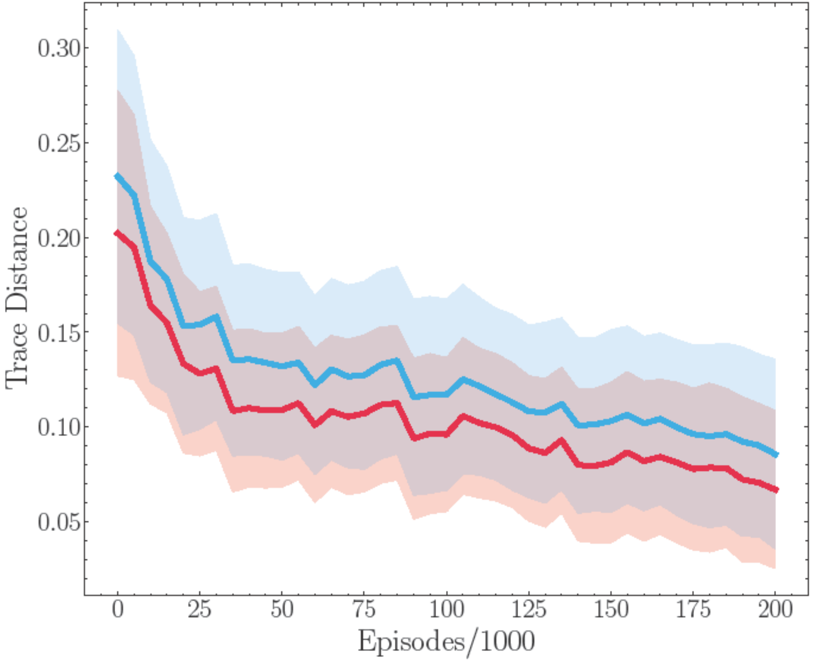

Results
Quantum hardware benchmarking

Benchmarking with circuits with depth spanning from 5 to 50.

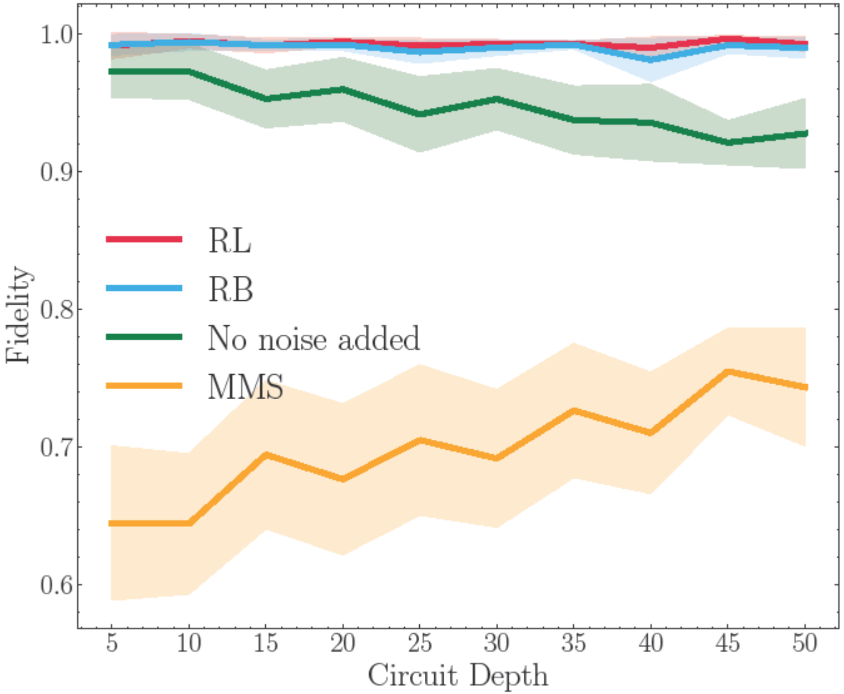

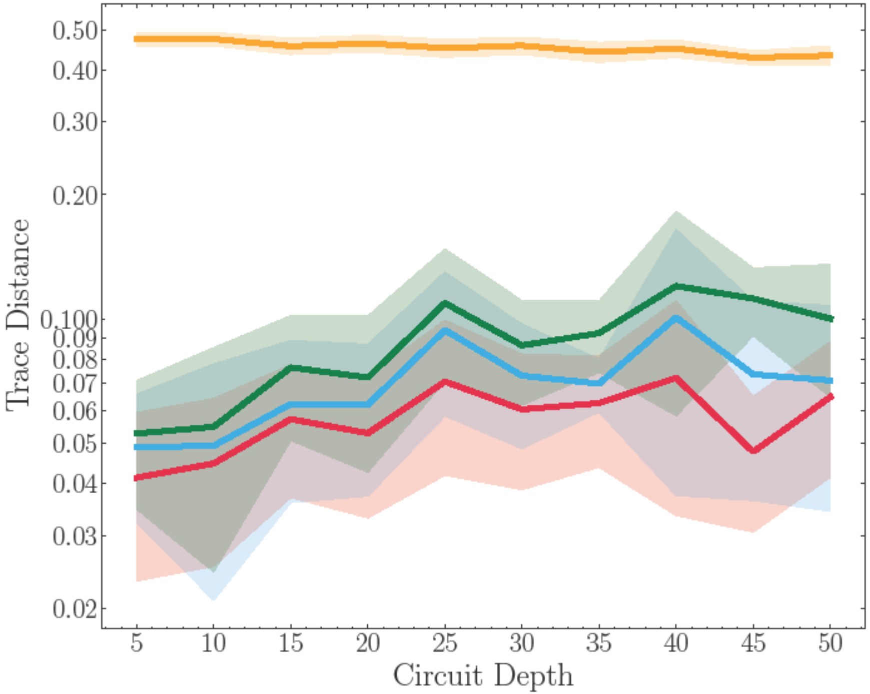

Conclusions

- RL agents can learn complex noise patterns.
- Better adaptability to different noise models.
- Better generalization to new circuits.
- Possibility to learn long-term correlations.
- Identification of coherent errors and other calibration problems.

Future work

- Scale to simulations with more than three qubits.
- Scale to multi qubit hardware.
- Solve bottleneck of quantum state tomography.
- Benchmark with other state of the art noise simulation techniques.
- Error mitigation with similar working principle.

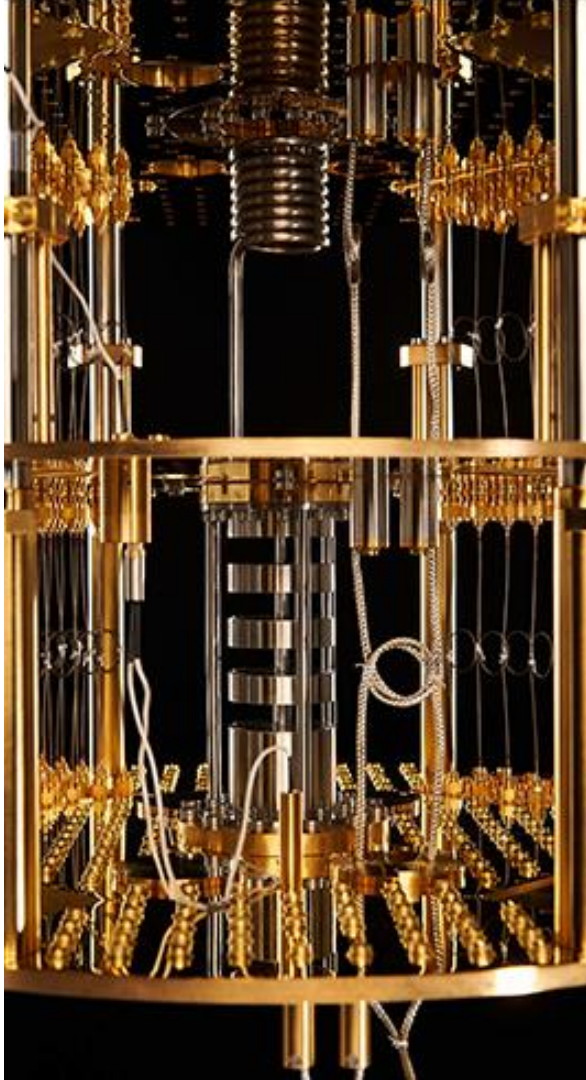

Technology Innovation Institute

SAPIENZA UNIVERSITÀ DI ROMA

Thank you for your attention!

Preprint: