A tale of two RLs

For Quantum Systems

Evert van Nieuwenburg

⟨aQa†⟩

https://aqa.liacs.nl/

We have a division for applied quantum algorithms

⟨aQa†⟩

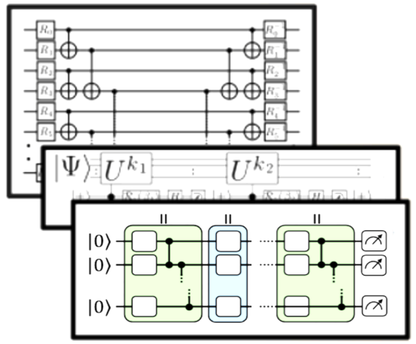

Provable (learning) advantages and separations

Hardware(-aware) design and optimization

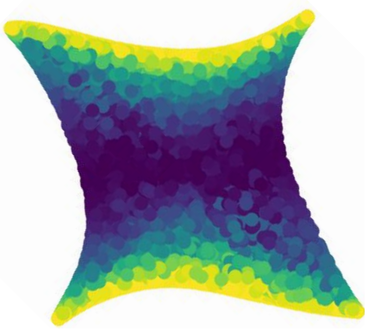

Efficient algorithms and applications

Alfons Laarman

Hao Wang

Vedran Dunjko

Evert van Nieuwenburg

Anna Dawid

Jordi Tura

liacs

Leiden Institute of Physics

Disclaimer

Lots of work-in-progress

I want to get across the main ideas, don't expect to follow details (but ask!)

A tale of two RLs

For Quantum Systems

Evert van Nieuwenburg

⟨aQa†⟩

https://aqa.liacs.nl/

To demonstrate the need for useful representations...

...please perform this division

CCLXI / IX = ...

261 / 9 = 29

Taken from https://www.deeplearningbook.org/

A good representation is...

... one that makes follow-up tasks easier

1. Easier to interpret
2. Less memory needed
3. More symmetric
4. Easier to manipulate
5. ...

For this talk, I'd like to focus on the interpretability aspect

A quick demo: Which (single!) real number would you choose to characterize an arbitrary 2-qubit state?

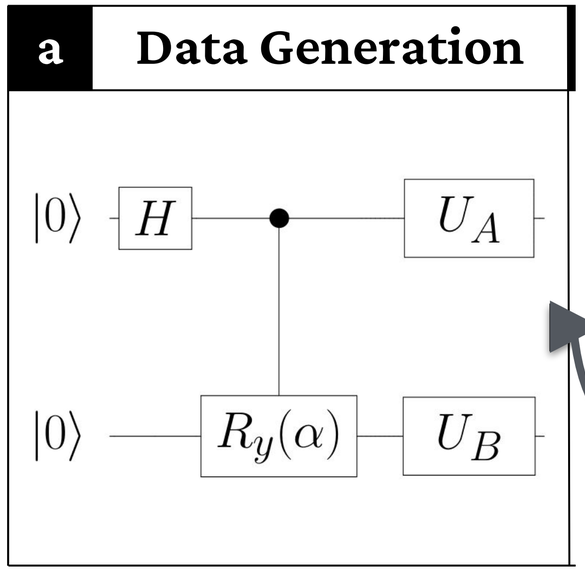

beta-VAE

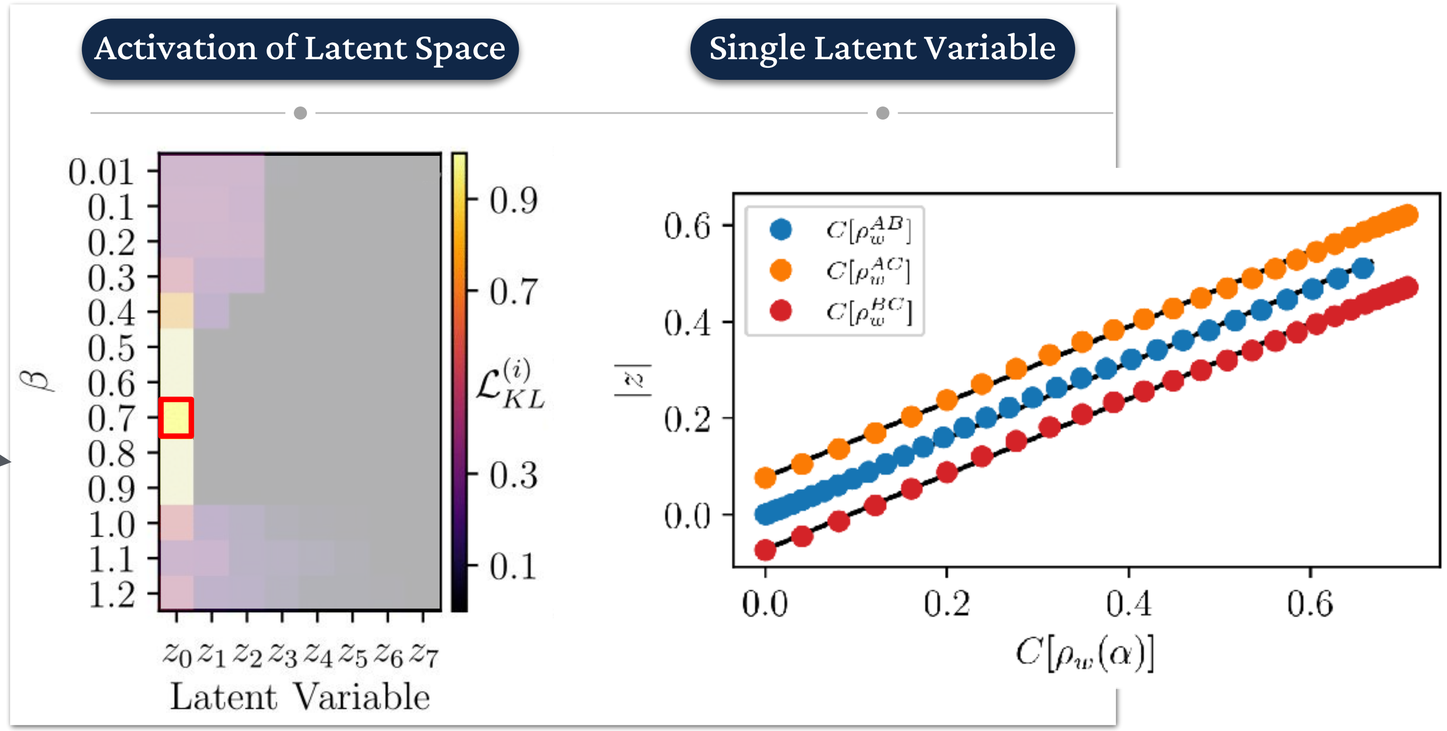

Felix Frohnert

https://arxiv.org/abs/2306.05694

Representations can also be learned in non-bottlenecked models

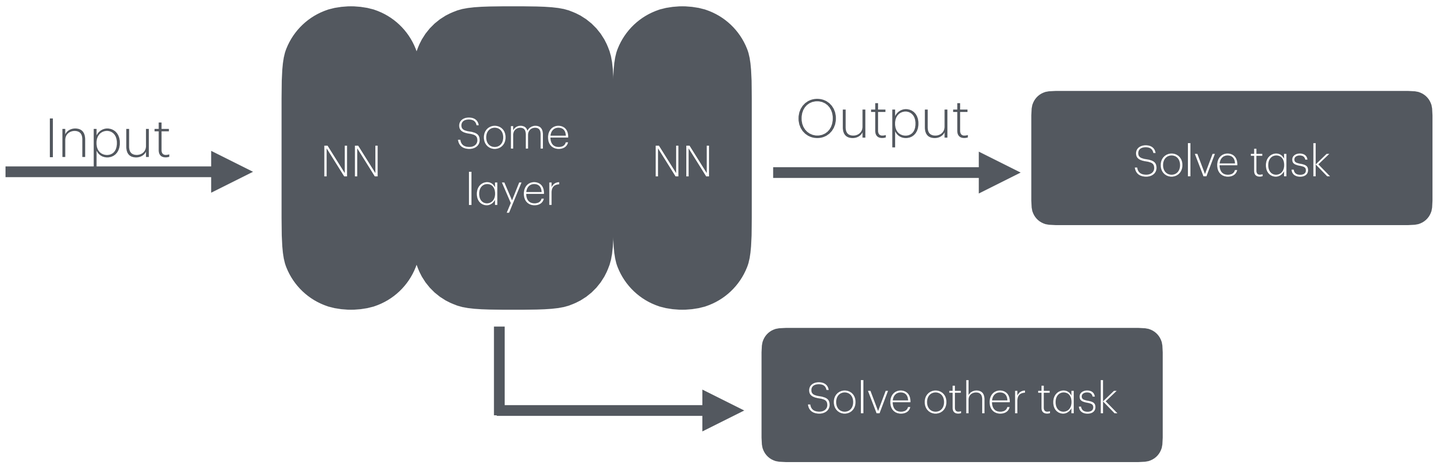

The Forward-Forward algorithm

Training* each layer separately (seems to) give better representations**

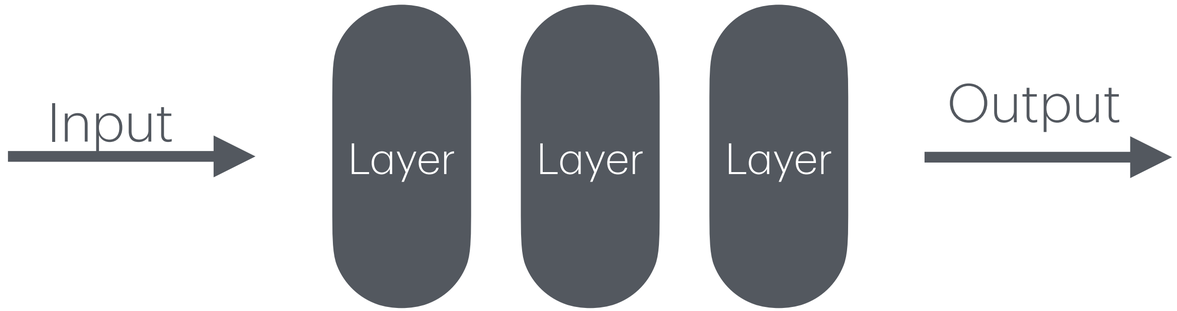

e.g. one-hot encoding *in image*

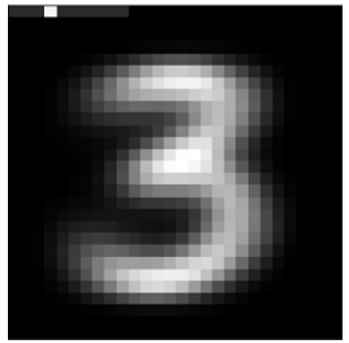

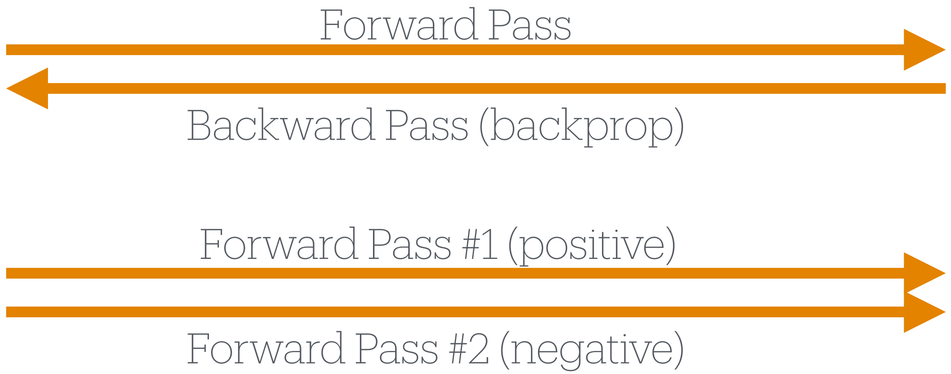

arXiv:2212.13345
arxiv:2305.18353

*Not with back propagation!
**Sparser, and for phase transitions they cluster

The neurons learn sparse representations

MNIST

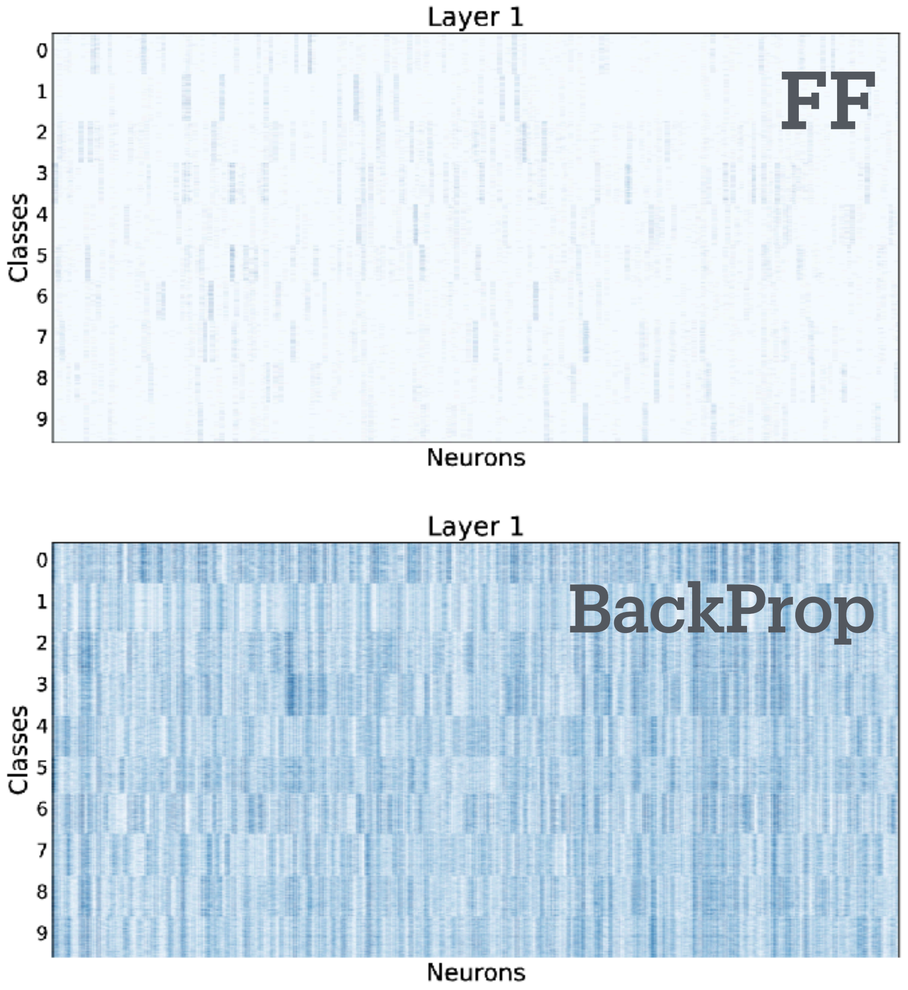

5-state Potts

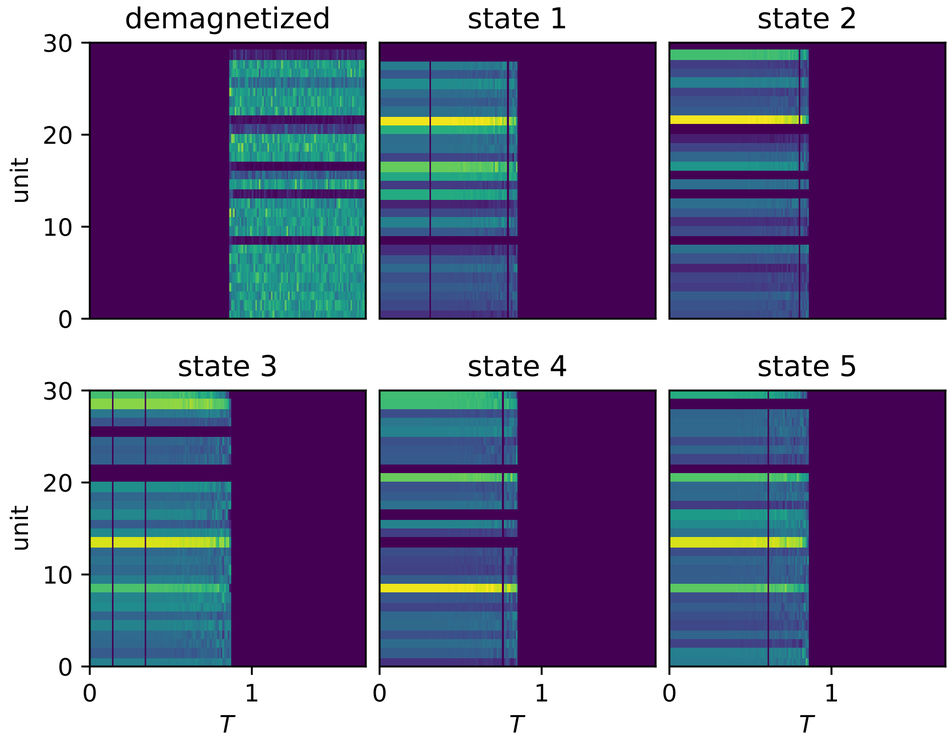

Not shown: the representations found by FF easily cluster using PCA, for various models (i.e. classify phase transitions)

arxiv:2305.18353

The strategy learned by an RL agent…
…is also generally interpretable

Entanglement Distribution in a Quantum Network

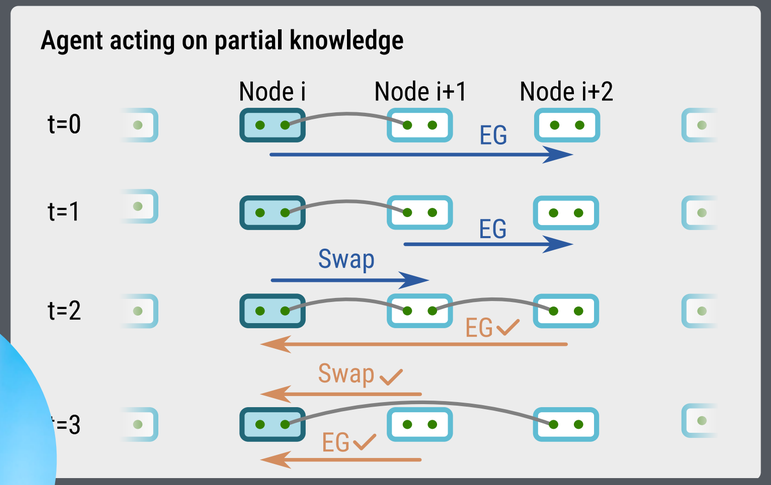

Jan Li

Qubit Control in a Noisy Environment

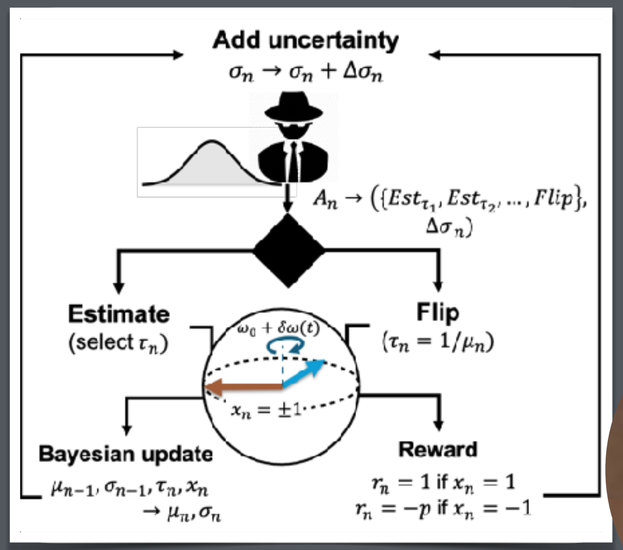

Jan Krzywda (PD)

A basic implementation of an X gate on a qubit

Works by applying a perpendicular field for a specific duration

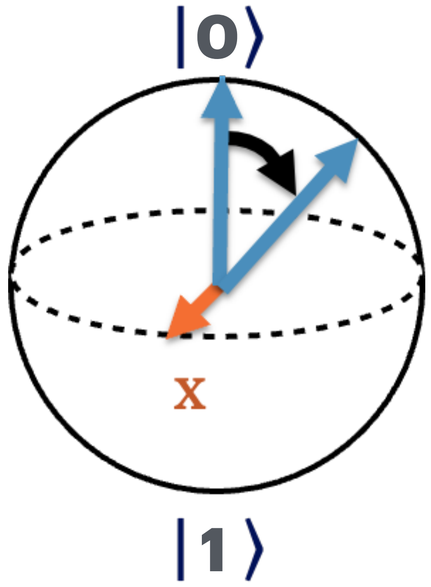

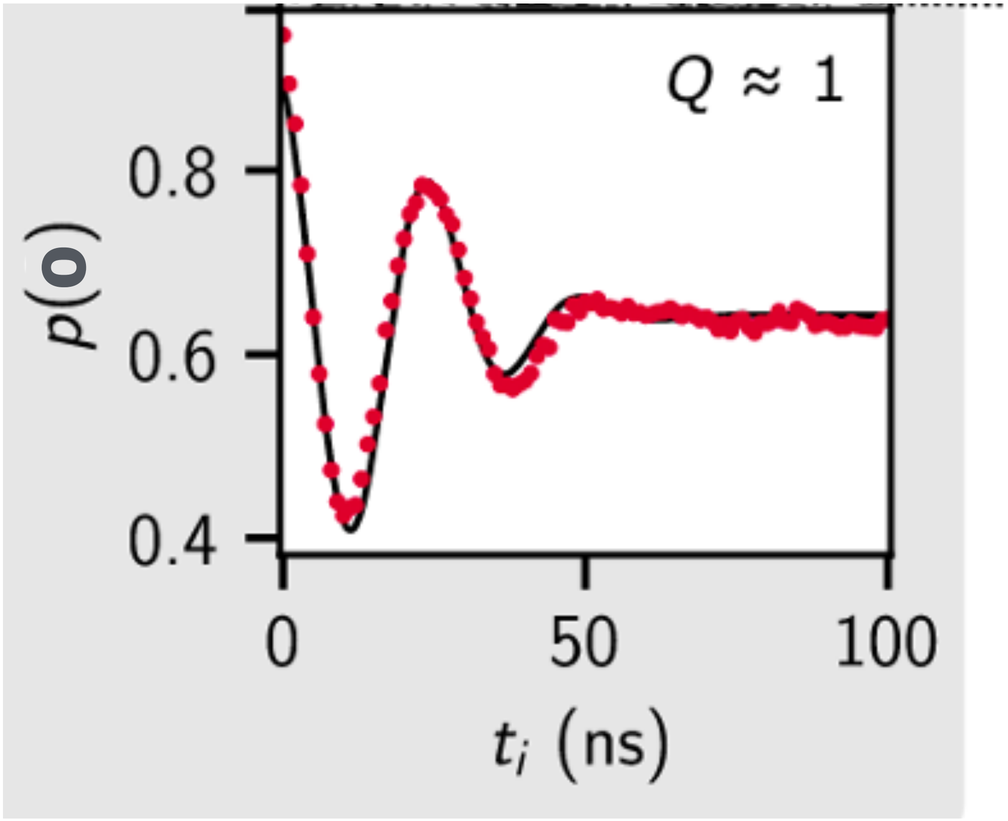

For slow noise, we can use Bayesian estimation

(noise ~ ms, experiments ~ns)

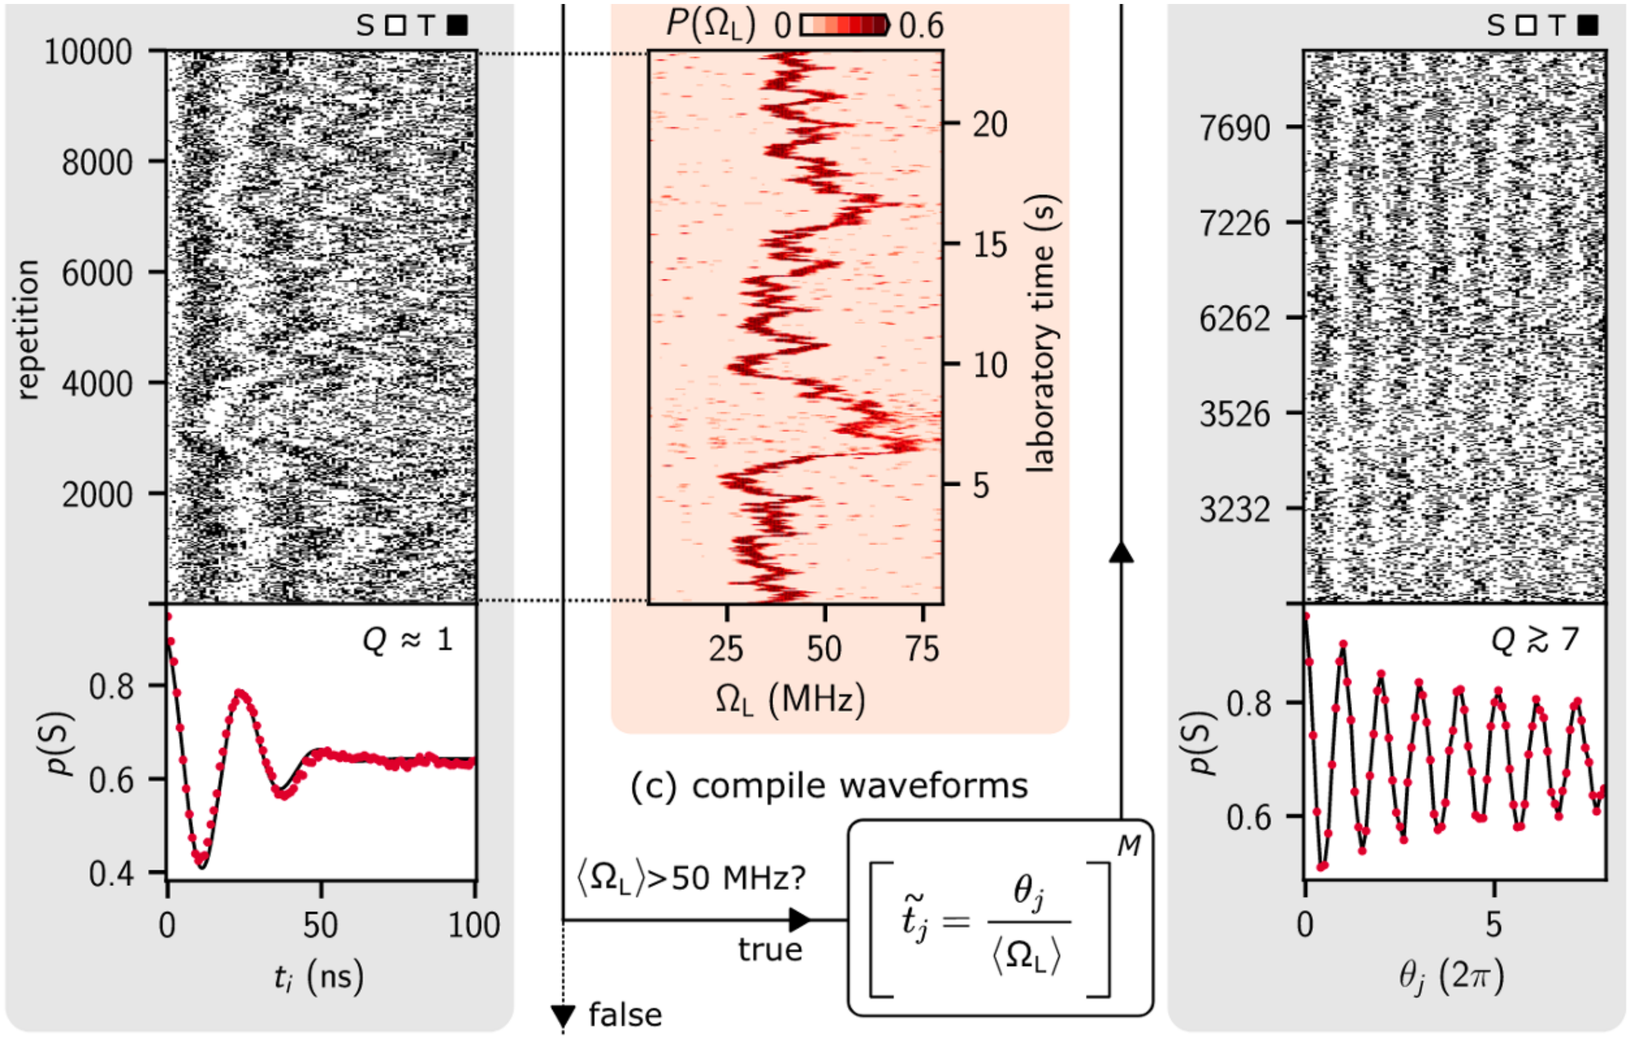

There is an inherent tradeoff between using shots to estimate and choosing to run an experiment (here: X gate)

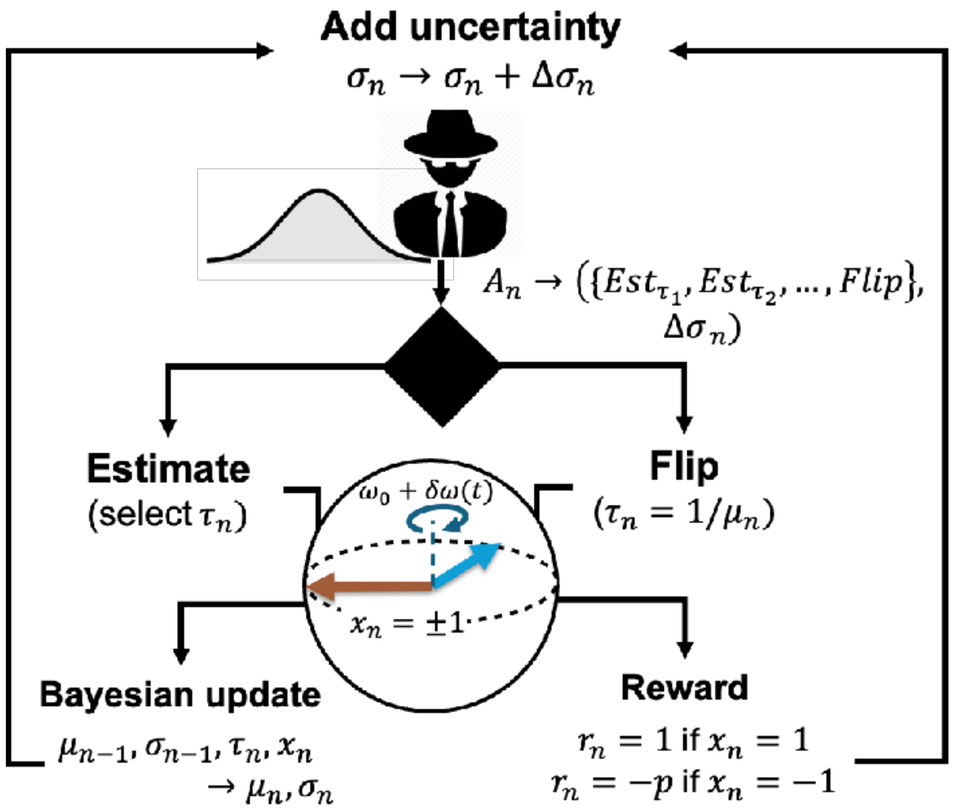

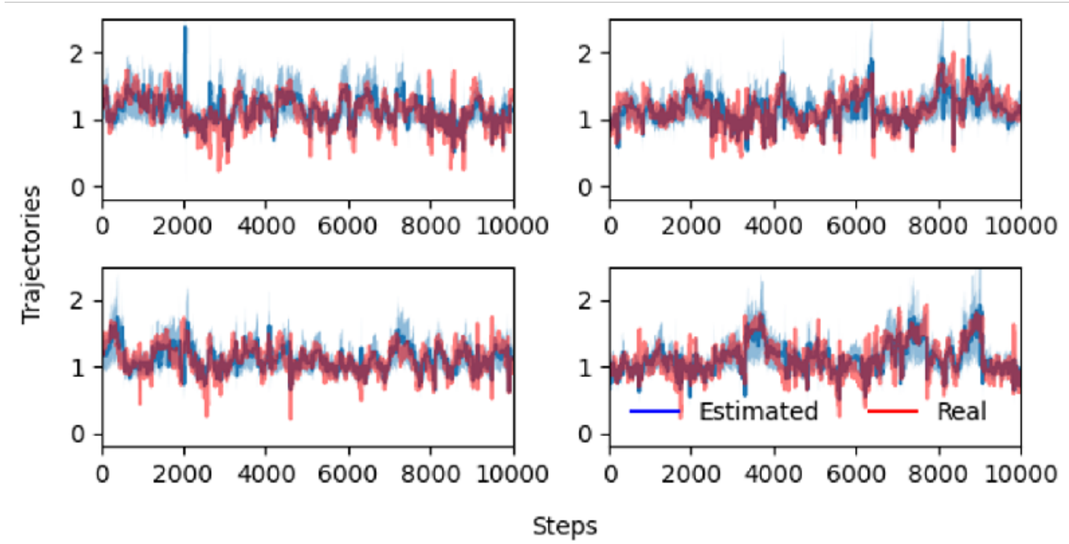

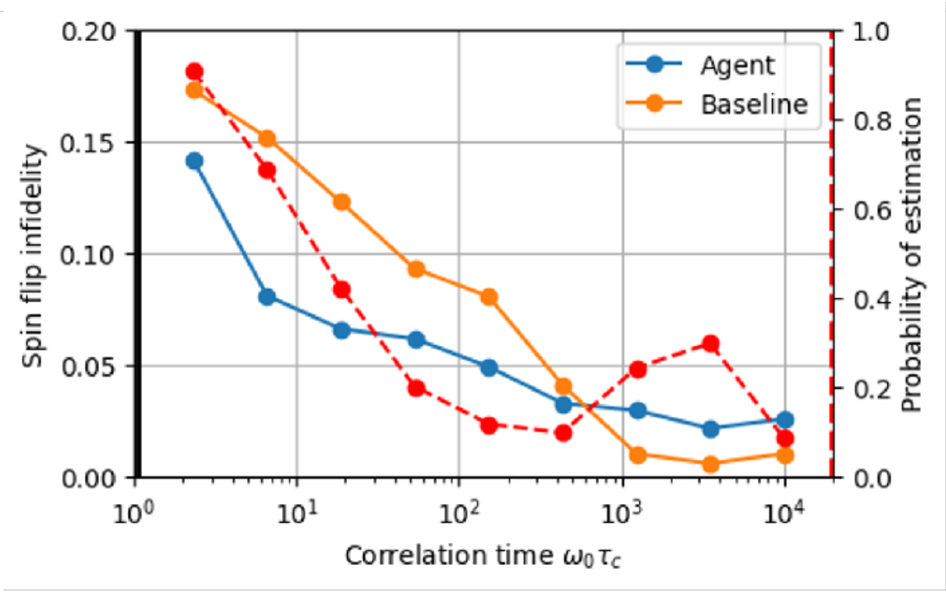

The agent learns to estimate when it is uncertain

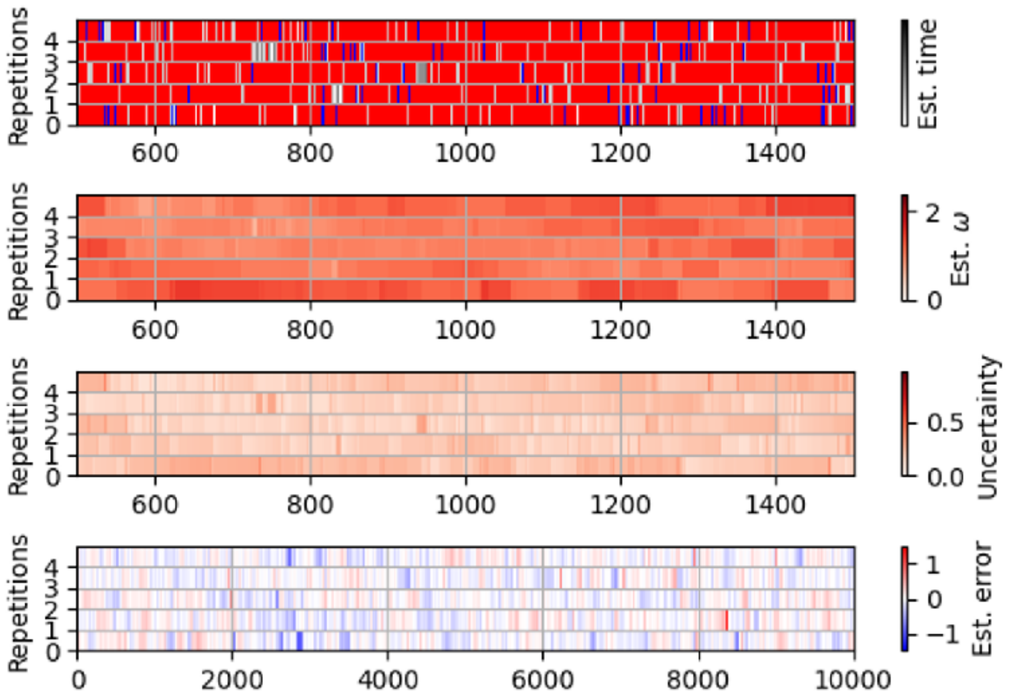

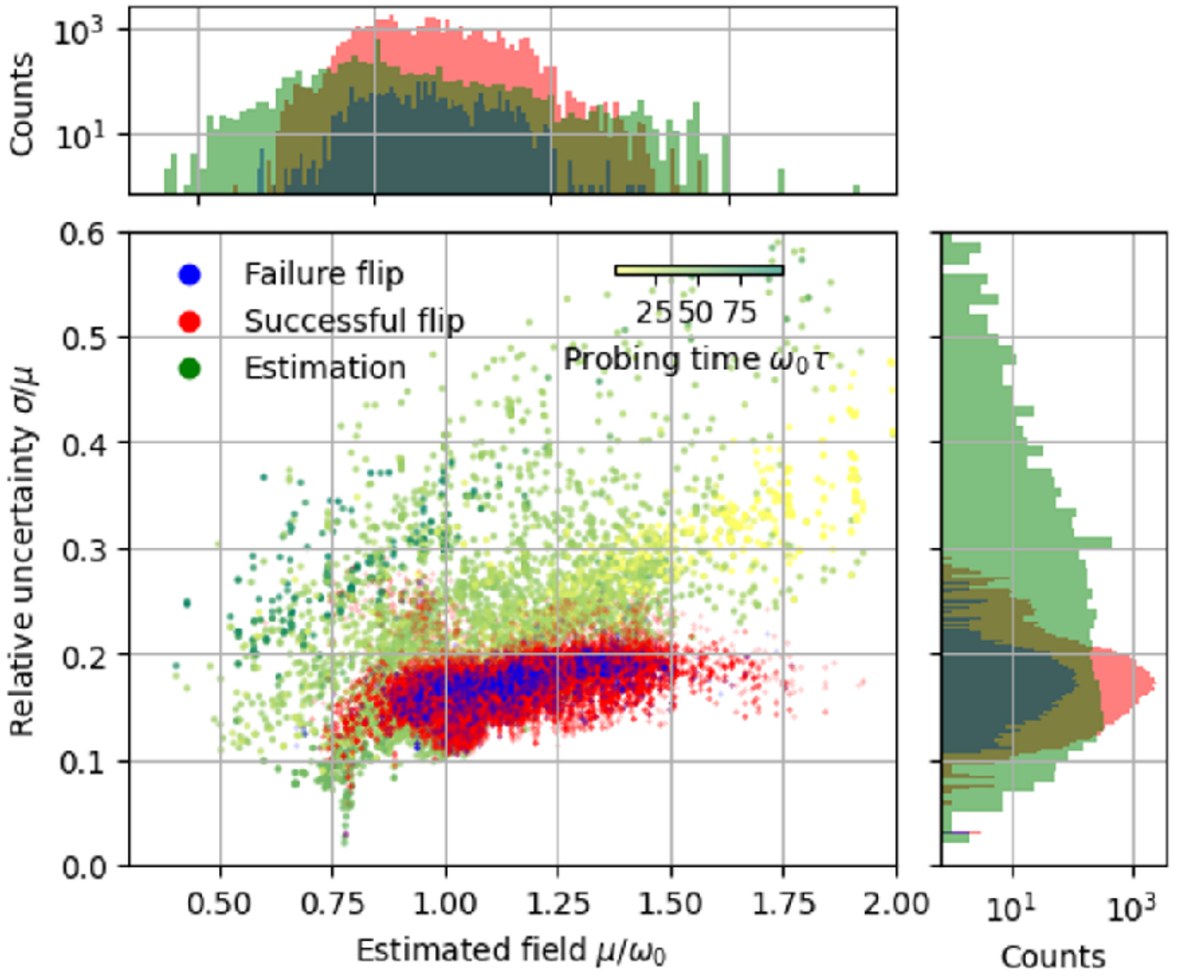

The strategy learned by an RL agent…
…is also generally interpretable

Entanglement Distribution in a Quantum Network

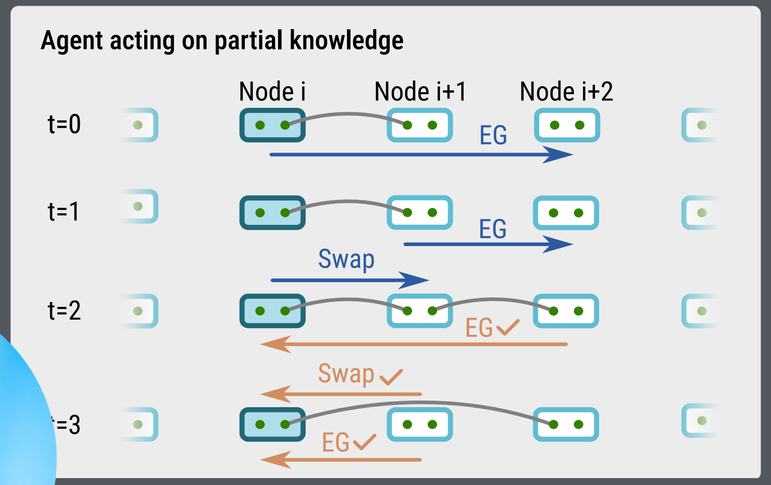

Jan Li

Qubit Control in a Noisy Environment

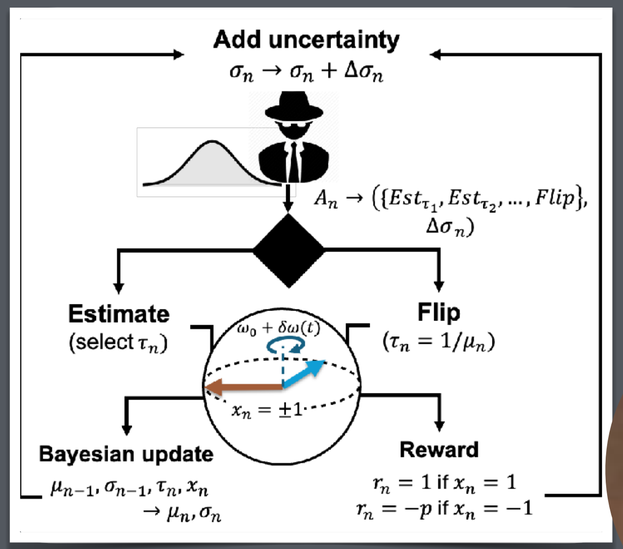

Jan Krzywda (PD)

Quantum networks consist of nodes with qubits

and entanglement distribution is one of the core concepts

Entanglement generation

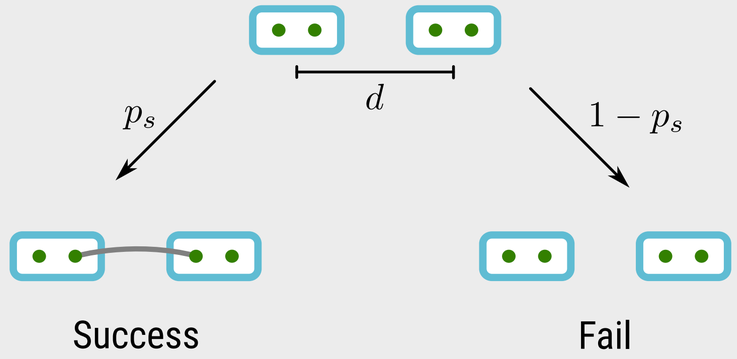

Entanglement Swap

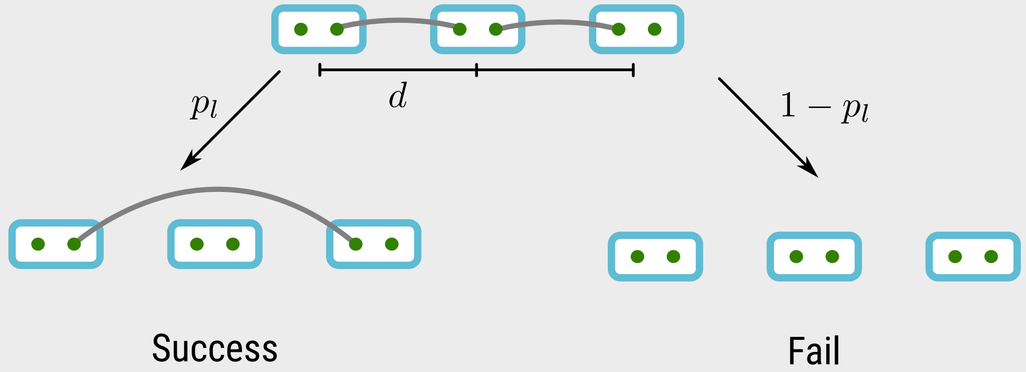


Bell state

Quantum Repater

The SWAP-asap protocol is the benchmark

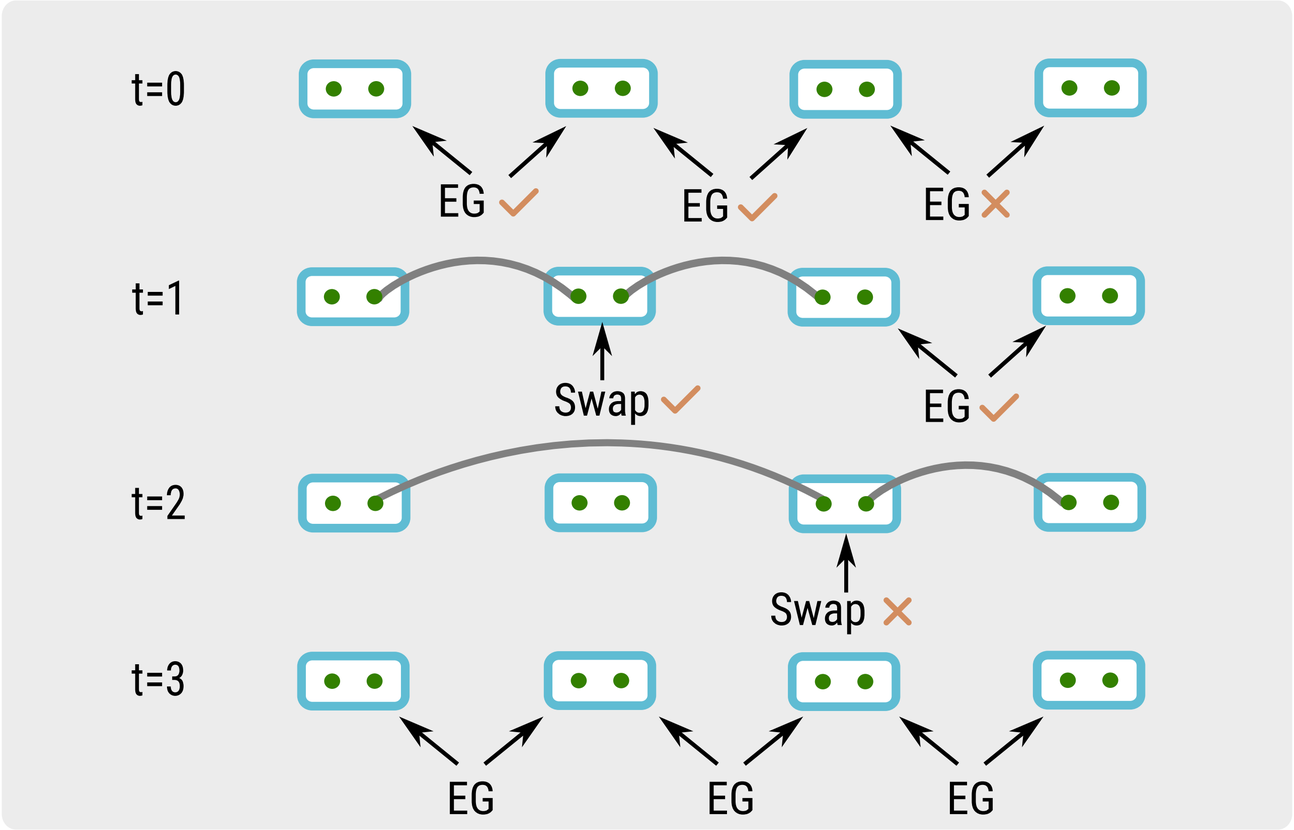

But in reality, we need to take classical communication into account

(We are considering a single agent for now, not one at each node)

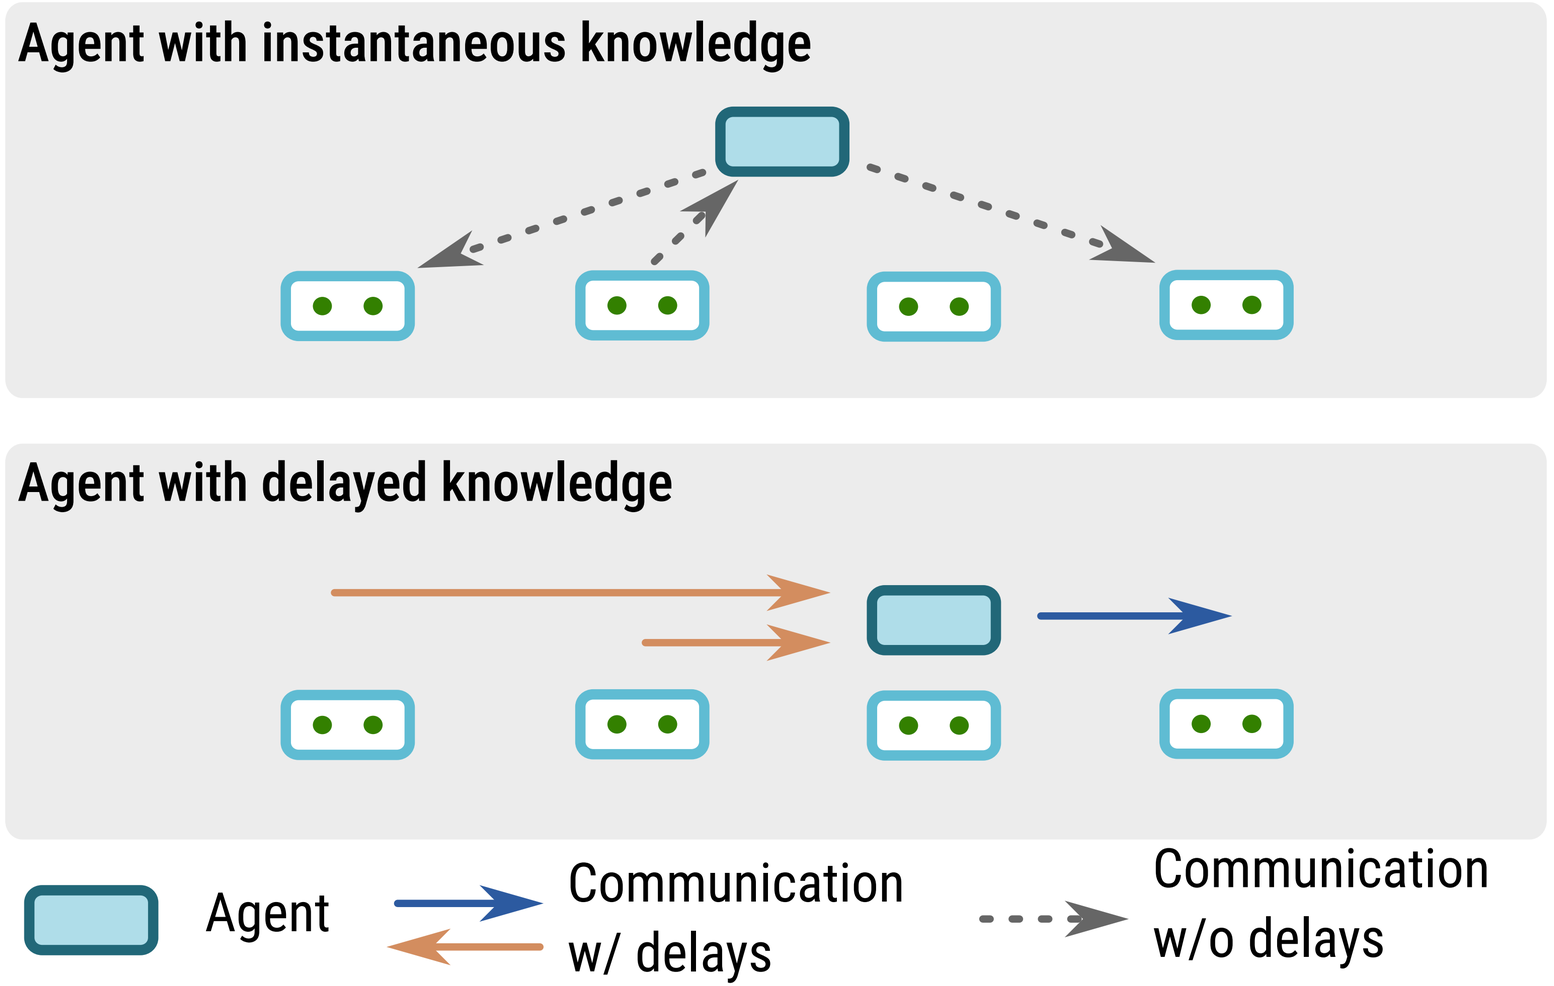

An RL agent *could* help us find a better strategy with CC

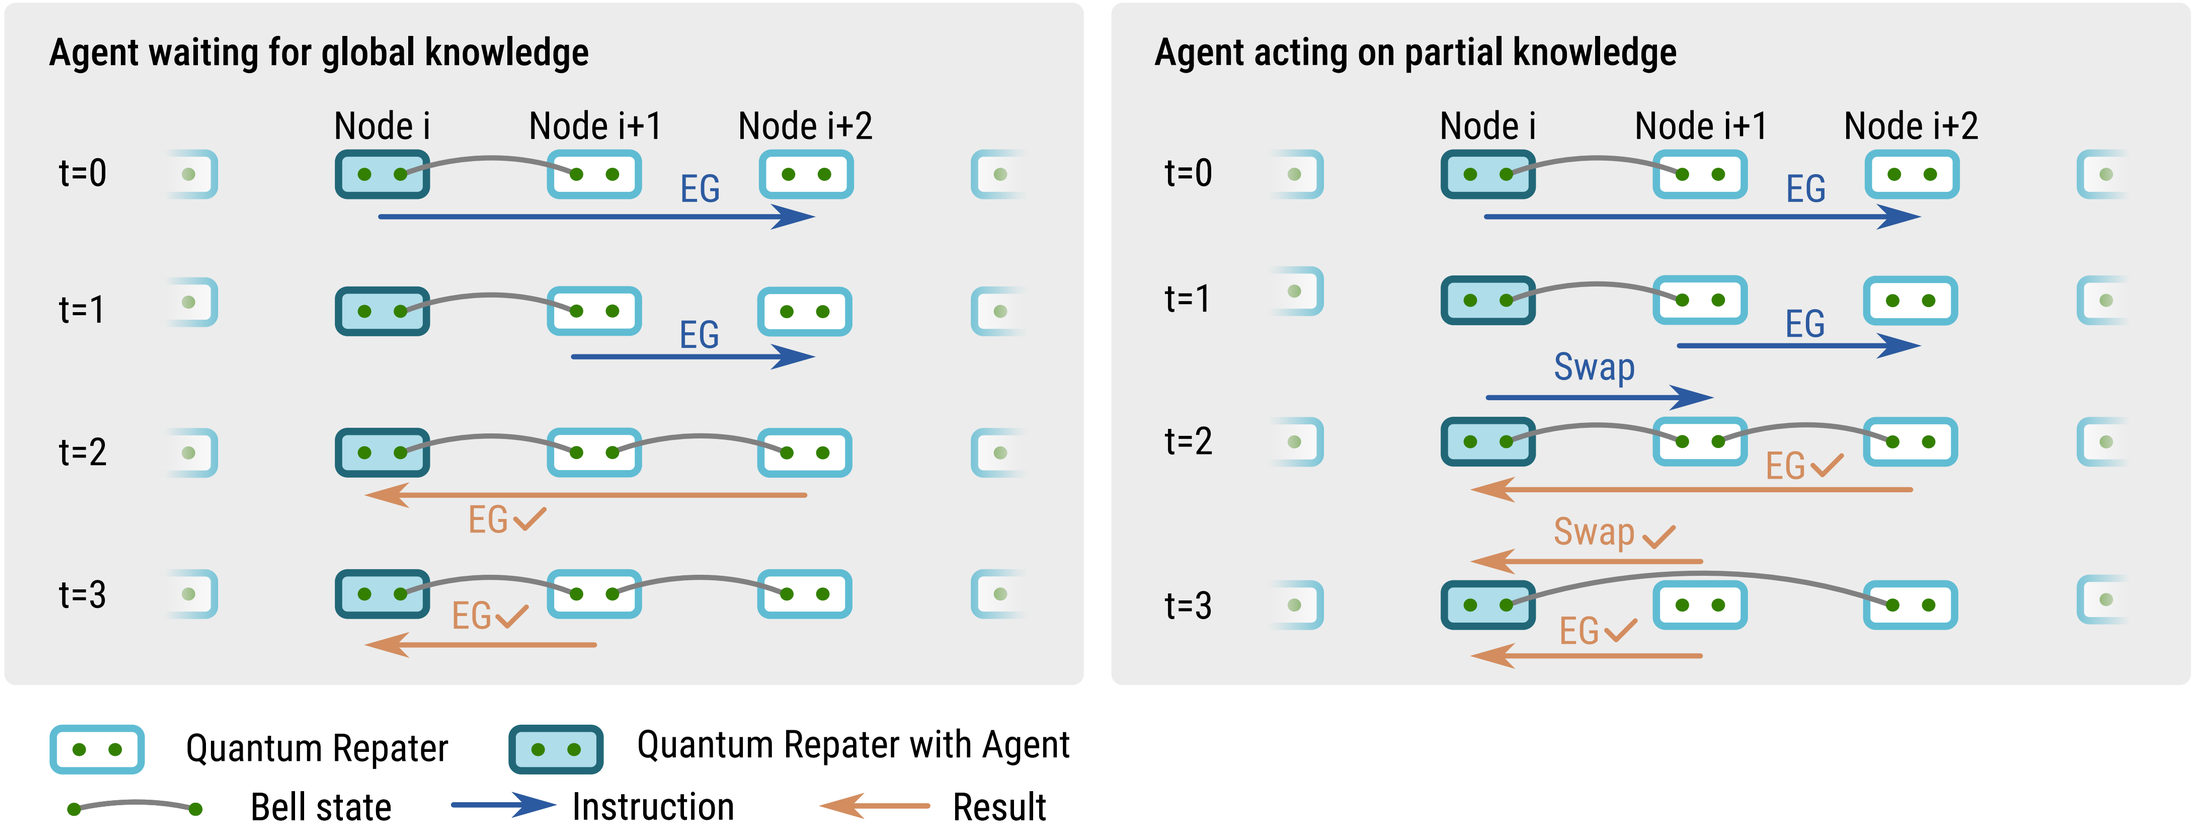

Compared to heuristic benchmarks, the RL agent performs well

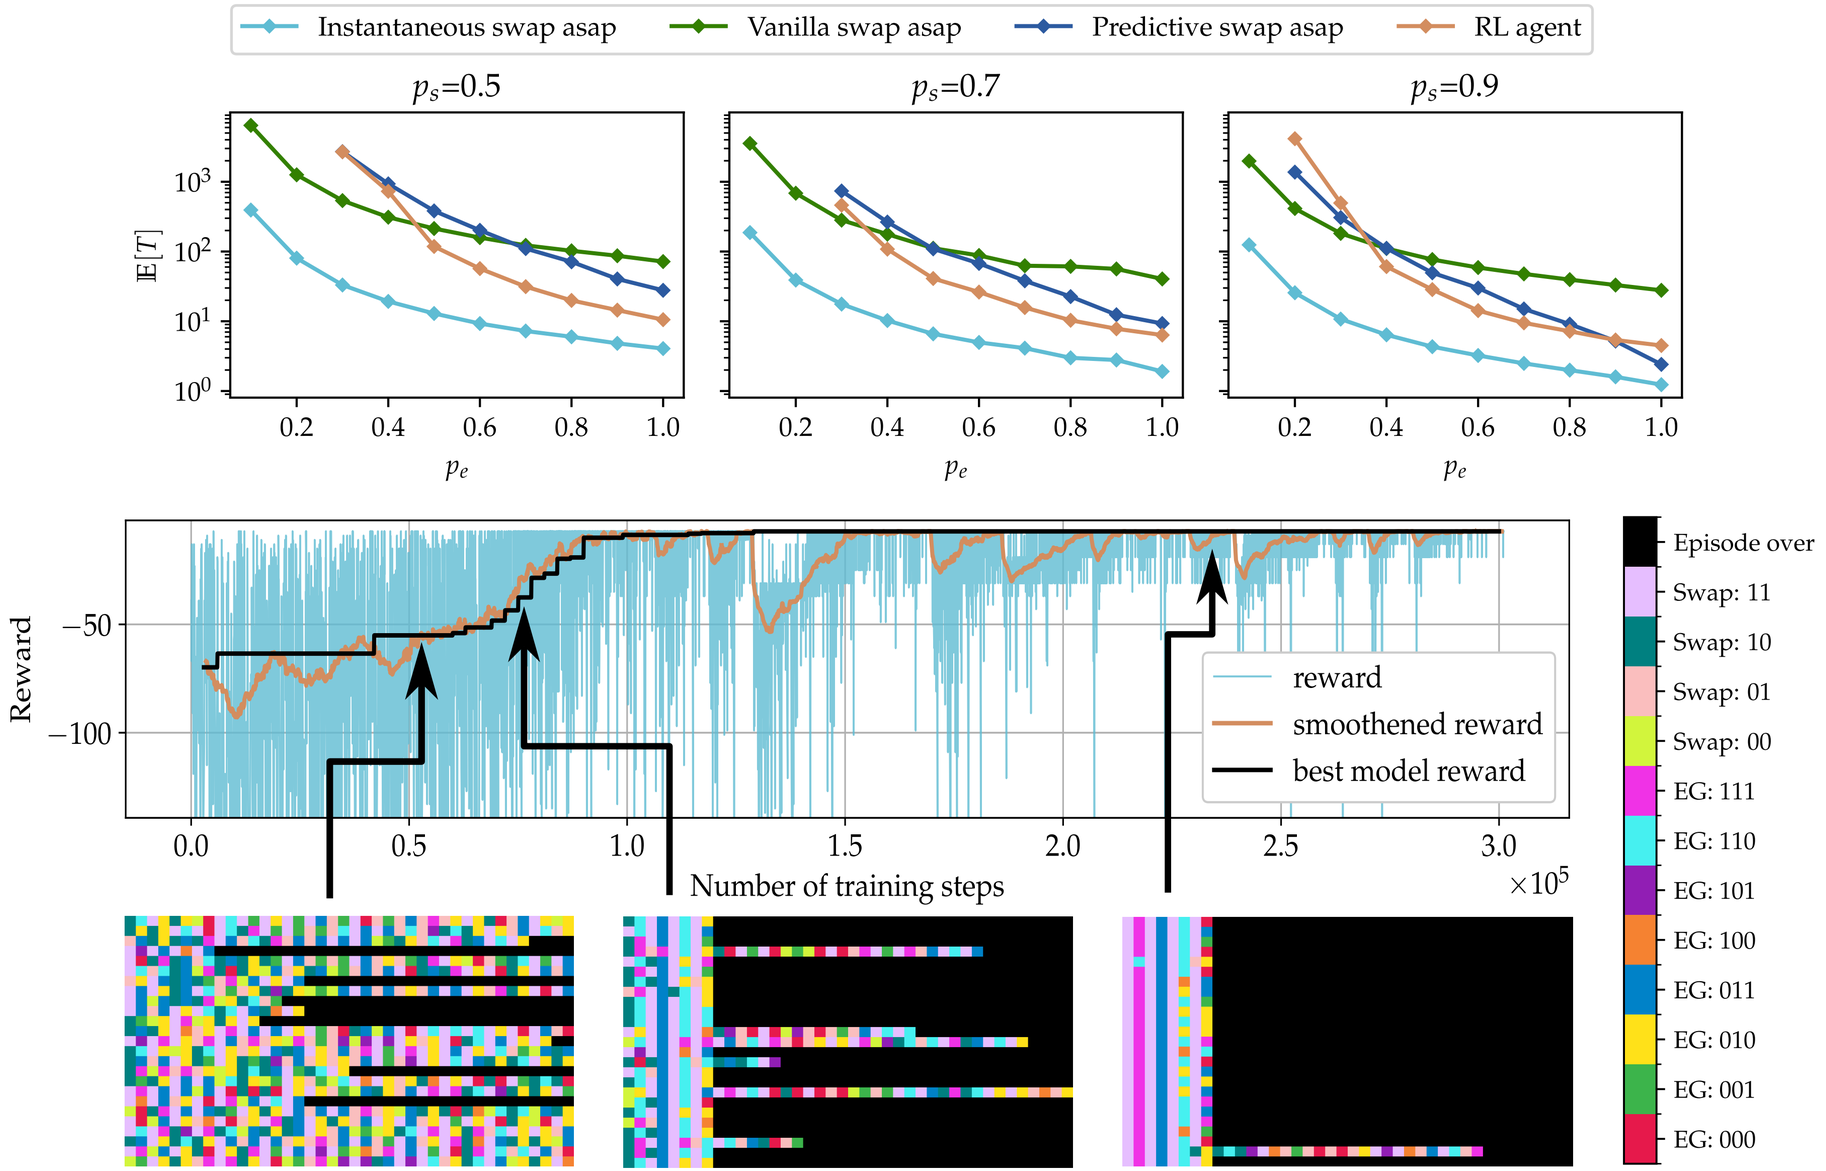

Quantum Games for/as Research

Gamify Quantum Problem
Solve game with AI

Quantum AI for
(Quantum) Games

Quantum Game Theory

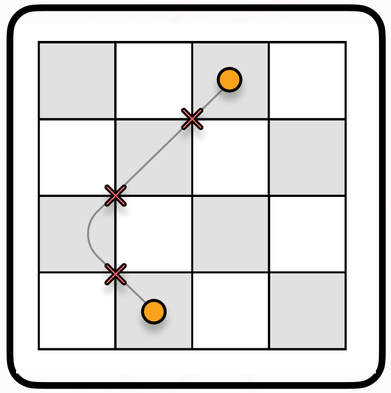

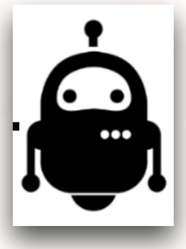

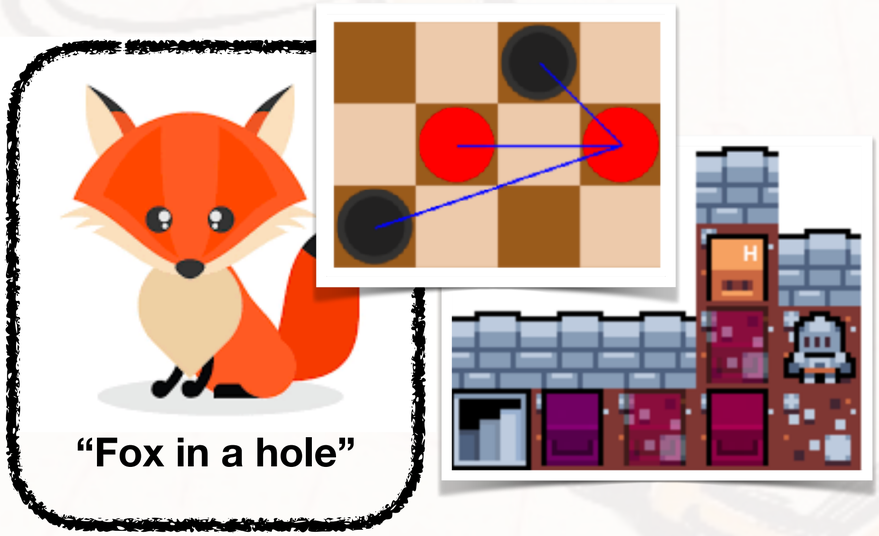

"Fox in a hole"

$$\rho_{i\ldots N}$$

$$\downarrow$$

$$u_i(\rho) = \mathrm{Tr}(\rho R_i)$$

$$\downarrow$$

$$u_i(\rho) \geq u_i(\rho_i' \otimes \rho_{\neg i}) \; \forall i$$

Quantum Nash Equilibrium

Can a classical AI optimize strategies for quantum systems?
Does a quantum circuit do better than a classical neural network?

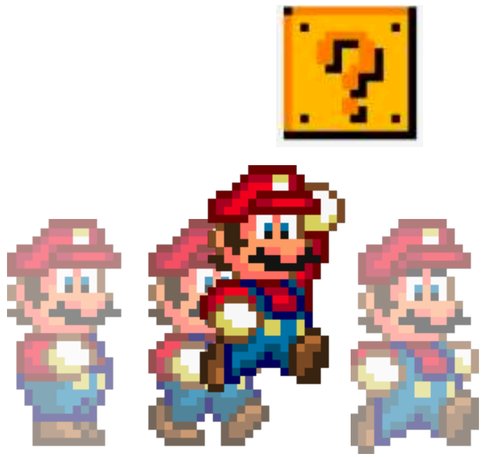

SciGym RL library Discord!

Please reach out!

Evert van Nieuwenburg

evert.vn@lorentz.leidenuniv.nl

e.p.l.van.nieuwenburg@liacs.leidenuniv.nl

⟨aQa⟩

https://aqa.liacs.nl/

AI for Quantum Control

How do you control something you cannot look at?

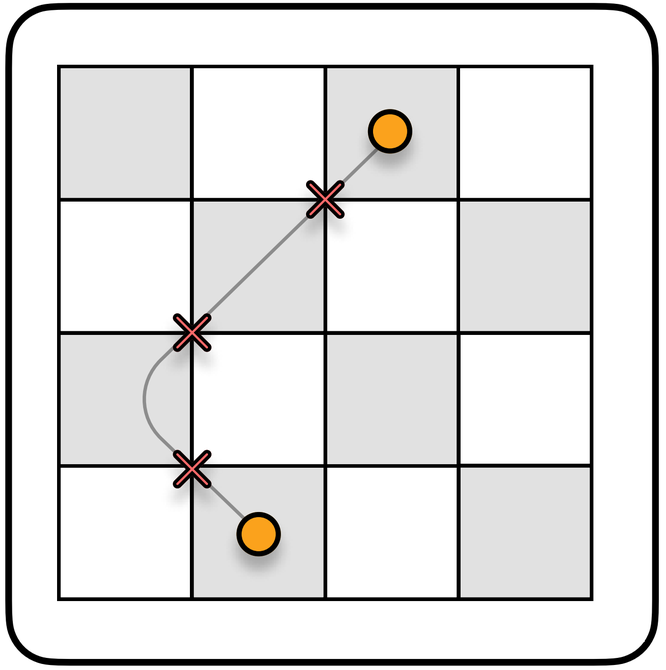

Quantum Error Correction

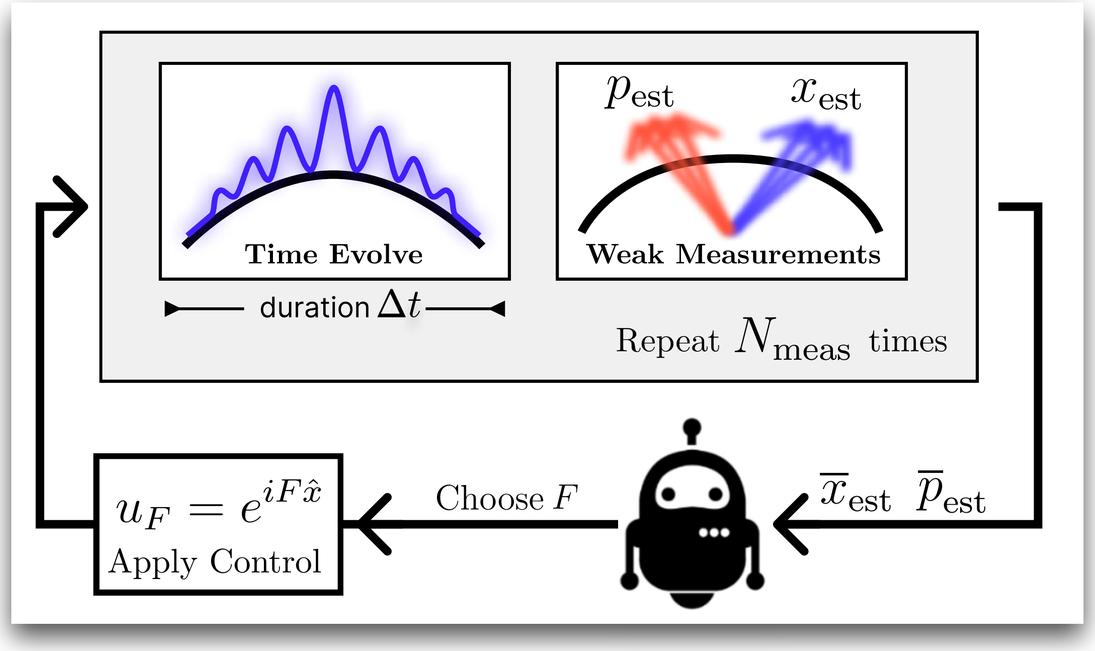

The Quantum Cartpole

Can a classical AI optimize strategies for quantum systems?
Does a quantum circuit do better than a classical neural network?

Dimensionality reduction...

...is a good way of representation learning

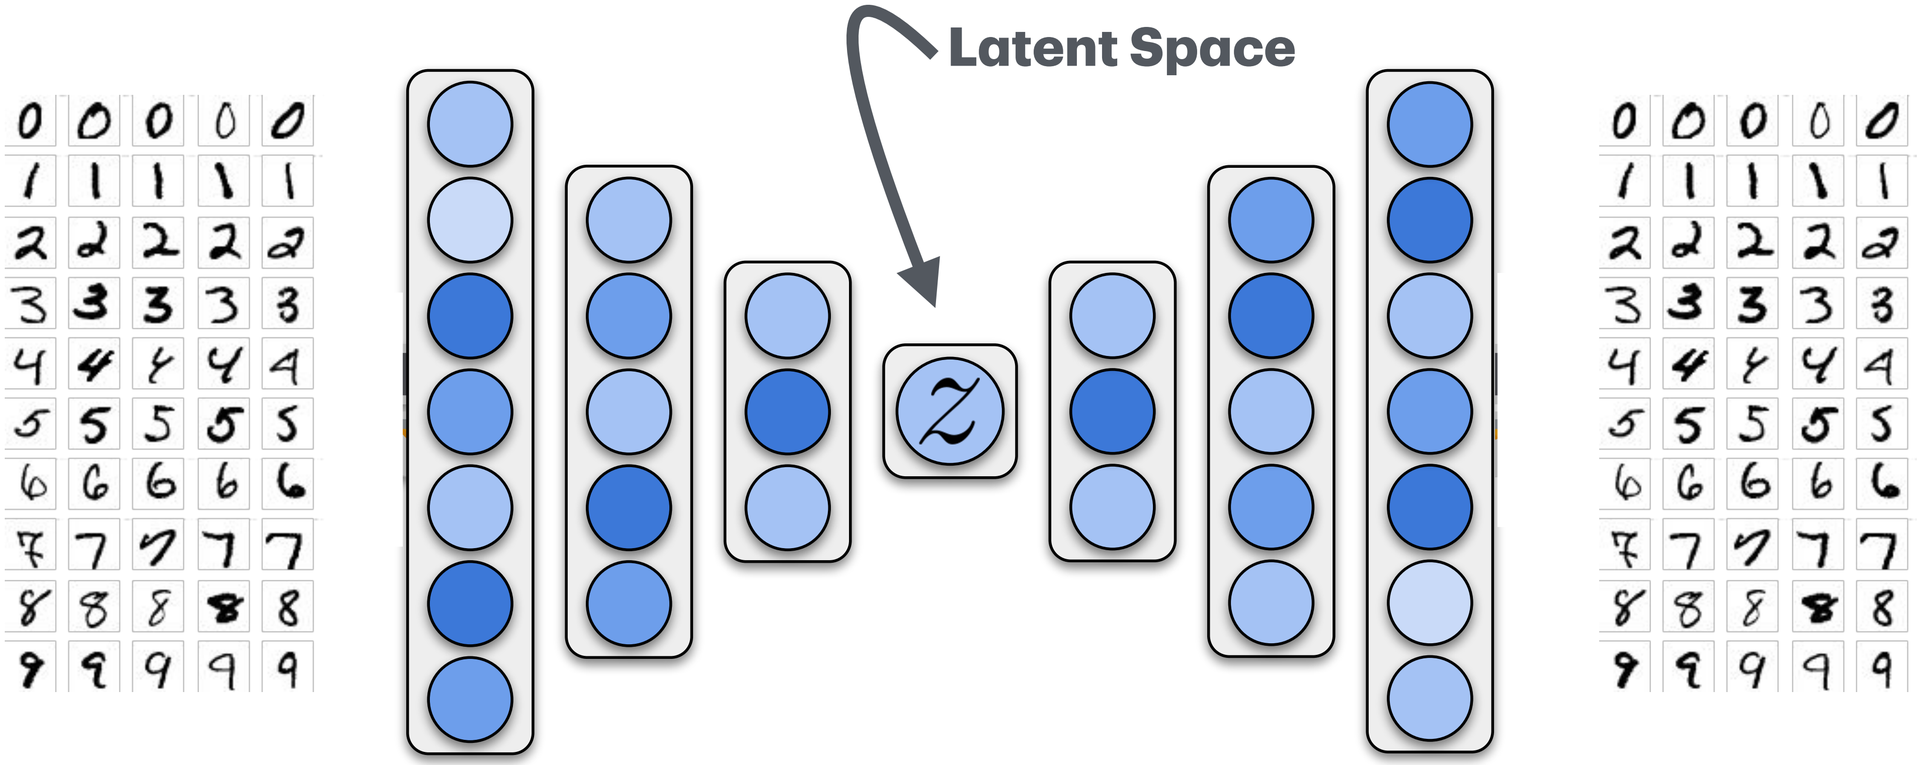

Information Bottleneck

(AutoEncoder)

There is an inherent tradeoff...

...in representation learning

Information Bottleneck

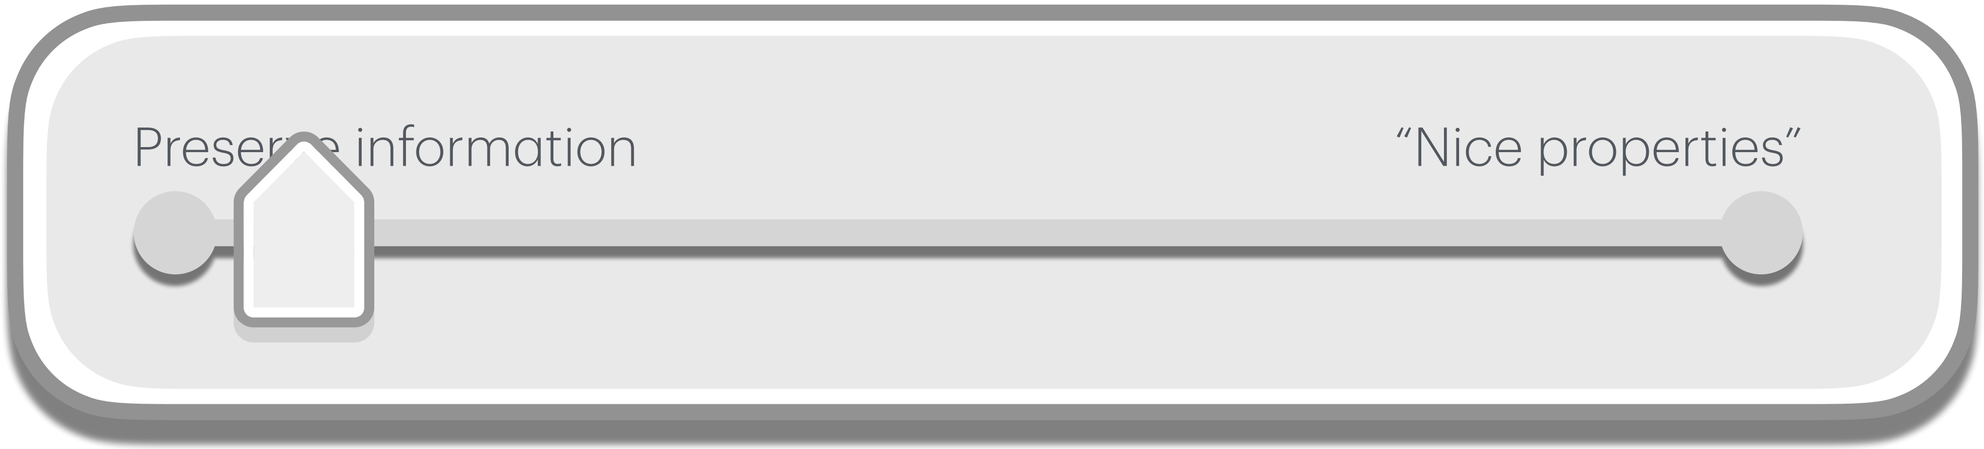

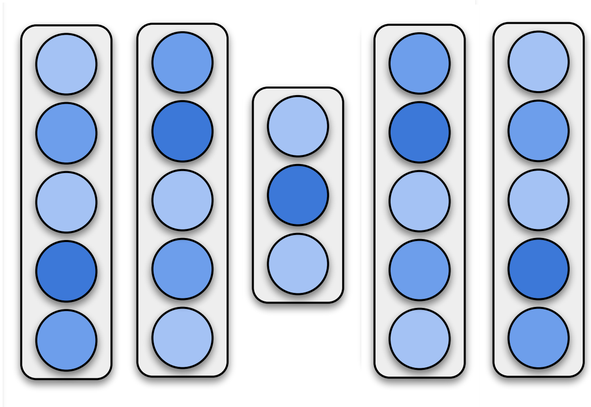

There is an inherent tradeoff...

...in representation learning

Information Bottleneck

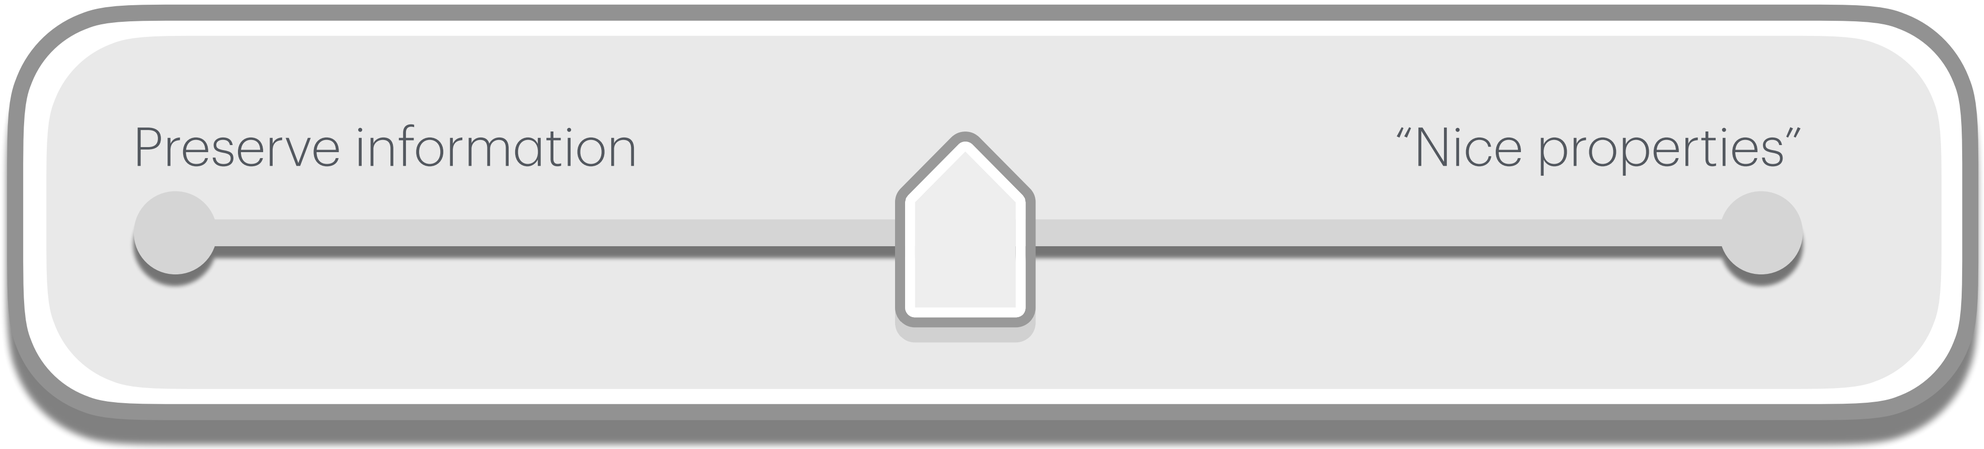

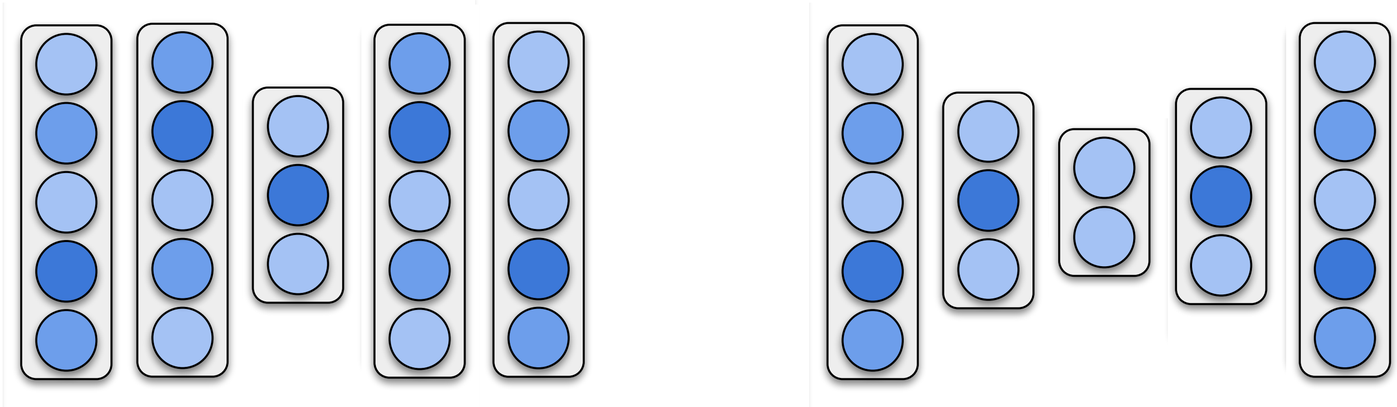

There is an inherent tradeoff...

...in representation learning