Title: Quantum-enhanced reinforcement learning

Speakers: Valeria Saggio

Collection: Machine Learning for Quantum Many-Body Systems

Date: June 14, 2023 - 11:30 AM

URL: https://pirsa.org/23060040

Abstract: The field of artificial intelligence (AI) has experienced major developments over the last decade. Within AI, of particular interest is the paradigm of reinforcement learning (RL), where autonomous agents learn to accomplish a given task via feedback exchange with the world they are placed in, called an environment. Thanks to impressive advances in quantum technologies, the idea of using quantum physics to boost the performance of RL agents has been recently drawing the attention of many scientists. In my talk I will focus on the bridge between RL and quantum mechanics, and show how RL has proven amenable to quantum enhancements. I will provide an overview of the most recent results -- for example, the development of agents deciding faster on their next move [1]-- and I will then focus on how the learning time of an agent can be reduced using quantum physics. I will show that such a reduction can be achieved and quantified only if the agent and the environment can also interact quantumly, that is, if they can communicate via a quantum channel [2]. This idea has been implemented on a quantum platform that makes use of single photons as information carriers. The achieved speed-up in the agent's learning time, compared to the fully classical picture, confirms the potential of quantum technologies for future RL applications.

[1] Sriarunothai, T. et al. Quantum Science and Technology 4, 015014 (2018).
[2] Saggio, V. et al. Nature 591, 229-233 (2021).

Machine Learning for Quantum Many-Body Systems
Perimeter Institute, Waterloo, Canada
12-16 June 2023

Quantum-enhanced reinforcement learning

Valeria Saggio

Massachusetts Institute of Technology

A few words about me

Generation of squeezed states of light

Dr. Alessandro Zavatta

Multi-partite entanglement generation and detection

Quantum reinforcement learning

Prof. Philip Walther

Color centers in silicon for scalable quantum computing

Machine learning applications to quantum systems

Prof. Dirk Englund

vsaggio@mit.edu

Let’s pose some questions!

- Why are we interested in machine learning?
- What is the role of quantum mechanics?
- How can we quantize machine learning? (Or what does it mean to quantize machine learning?)
- How can we implement quantum machine learning on quantum platforms?

Where does the hype come from?

Where does the hype come from?

International Journal of Computer Science & Information Technology (IJCSIT), Vol 3, No 1, Feb 2011

MACHINE LEARNING METHODS FOR SPAM E-MAIL CLASSIFICATION

W.A. Awad[1] and S.M. ELseuofi[2]

[1]Math.&Comp.Sci.Dept., Science faculty, Port Said University

wealak@yahoo.com

[2]Inf. System Dept.,Ras El Bar High inst.

sherif.shawki@gmail.com

nature biomedical engineering — REVIEW ARTICLE

https://doi.org/10.1038/s41551-018-0305-z

Artificial intelligence in healthcare

Kun-Hsing Yu[1], Andrew L. Beam[1] and Isaac S. Kohane[1,2*]

DRN: A Deep Reinforcement Learning Framework for News Recommendation

Guanjie Zheng†, Fuzheng Zhang§, Zihan Zheng§, Yang Xiang§

Nicholas Jing Yuan§, Xing Xie§, Zhenhui Li†

RESEARCH ARTICLE

Abusive Language Detection in Online User Content

Authors: Chikashi Nobata, Joel Tetreault, Achint Thomas, Yashar Mehdad, Yi Chang Authors Info & Claims

WWW '16: Proceedings of the 25th International Conference on World Wide Web • April 2016 • Pages 145–153 • https://doi.org/10.1145/2872427.2883062

ARTICLE

doi:10.1038/nature24270

Mastering the game of Go without human knowledge

RESEARCH ARTICLE | MAY 01 2011

Financial Statement Fraud Detection: An Analysis of Statistical and Machine Learning Algorithms

Johan Perols

+ Author & Article Information

AUDITING: A Journal of Practice & Theory (2011) 30 (2): 19–50.

https://doi.org/10.2308/ajpt-50009

TESS Data for Asteroseismology (T'DA) Stellar Variability Classification Pipeline: Setup and Application to the Kepler Q9 Data

J. Audenaert[1], J. S. Kuszlewicz[2,3], R. Handberg[3], A. Tkachenko[1], D. J. Armstrong[4,5], M. Hon[6,7], R. Kgoadi[8], M. N. Lund[3], K. J. Bell[19,9], L. Bugnet[10,11] + Show full author list

Published 2021 October 21 • © 2021. The American Astronomical Society. All rights reserved.

The Astronomical Journal, Volume 162, Number 5

Citation J. Audenaert et al 2021 AJ 162 209

DOI 10.3847/1538-3881/ac166a

Advances in Multi-Variate Analysis Methods for New Physics Searches at the Large Hadron Collider

Anna Stakia, Tommaso Dorigo, Giovanni Banelli, Daniela Bortoletto, Alessandro Casa, Pablo de Castro, Christophe Delaere, Julien Donini, Livio Finos, Michele Gallinaro, Andrea Giammanco, Alexander Held, Fabricio Jiménez Morales, Grzegorz Kotkowski, Seng Pei Liew, Fabio Maltoni, Giovanna Menardi, Ioanna Papavergou, Alessia Saggio, Bruno Scarpa, Giles C. Strong, Cecilia Tosciri, João Varela, Pietro Vischia, Andreas Weiler

… and *many* more applications!

And why considering quantum systems?

Building a (useful) quantum computer in the lab is *not* easy!

It mainly requires:

- Accurate control over very small systems;
- Preservation of quantum coherence through many computational operations (using error correction).

↓

The realization of large fault-taulerant quantum computers is still a challenge.

However, we do have small quantum systems which can test the advantages of quantum computing!

Combining machine learning and quantum computing

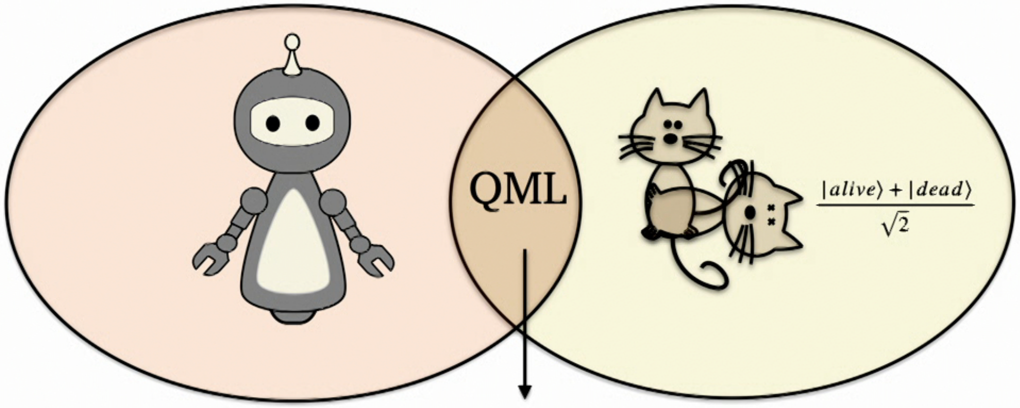

How to interpret this?

Examples:

- We can use ML to describe the internal state of a quantum system, or to discriminate between quantum states, or to learn phase transitions in many-body quantum systems;

S. Aaronson, In: Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences, vol. 463, pp. 3089–3114. The Royal Society (2007)
A. Bisio et al. *Phys. Rev. A* **81**(3), 032324 (2010)
J. Carrasquilla, et al. *Nature Phys.* **13**, 431–434 (2017)

- We can use quantum computing to speed up a robot's decision-making process or its learning process.

T. Sriarunothai et al. *Quantum Sci. Technol.* **4**, 015014 (2019)
V. Saggio et al. *Nature* **591**, 229–233 (2021)

Machine learning, more specifically…

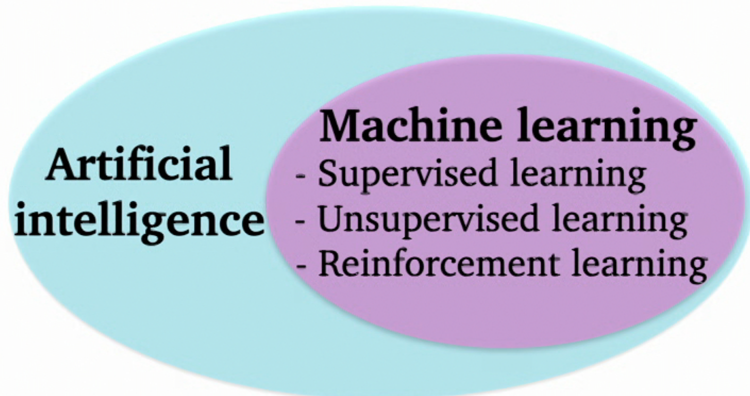

Supervised learning

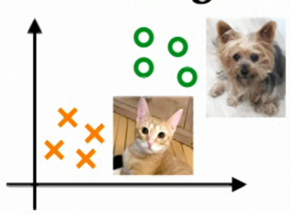

It can classify data (is it a cat or a dog?)

Unsupervised learning

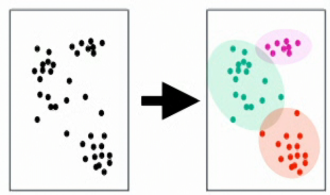

It can find patterns in data (anomaly detection)

Reinforcement learning

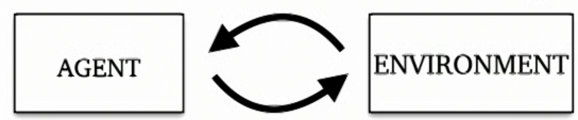

Based on learning via feedback exchange with an environment

Reinforcement learning in one slide (or at least what we need to know)

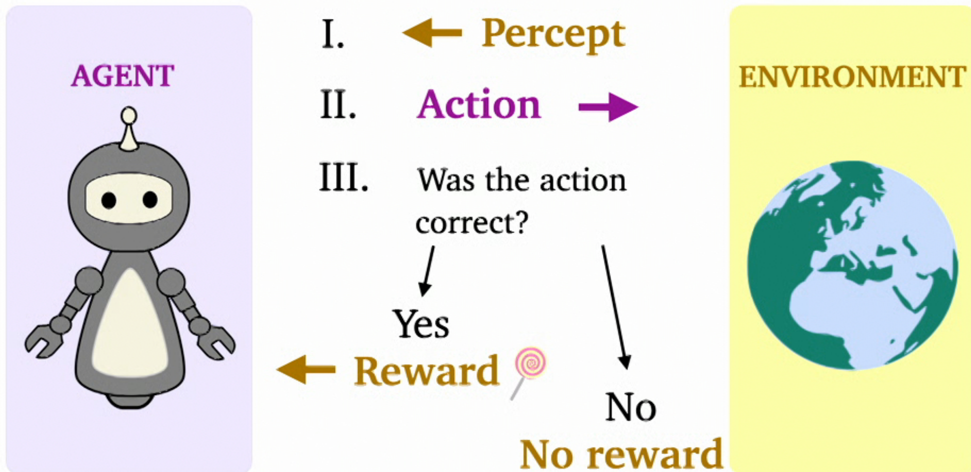

I. The agent receives perceptual input from the environment;

II. The agent processes the input and performs an action;

III. The environment either rewards or punishes the action.

Reinforcement learning in one slide (or at least what we need to know)

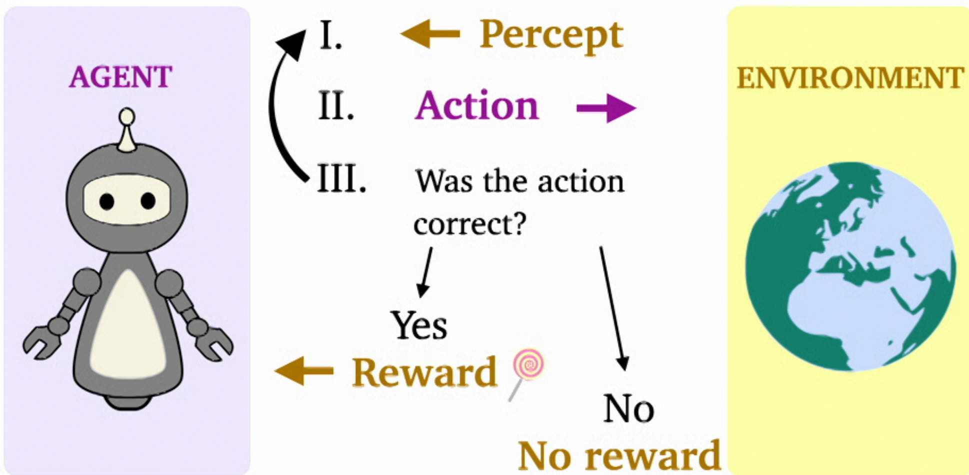

I. The agent receives perceptual input from the environment;

II. The agent processes the input and performs an action;

III. The environment either rewards or punishes the action.

The reward on a certain action increases the likelihood of the agent to perform it again.

⟶ **Learning process, which is *reinforced***

Combining reinforcement learning and quantum mechanics

	
Classical	Quantum
Quantum	Classical
Quantum	Quantum

ML for quantum

Quantum for ML

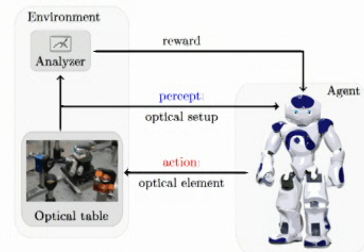

Classical / Quantum → Classical learning techniques facing quantum tasks, for example the design of quantum experiments.

A. Melnikov et al. *PNAS* **115**, 1221 (2018)

Quantum / Classical → Quadratic speed-up in the agent's decision time.

G. D. Paparo et al. *Phys. Rev. X* **4(3)**, 031002 (2014)

T. Sriarunothai et al. *Quantum Sci. Technol.* **4**, 015014 (2019)

Quantum / Quantum → Speed-up in the agent's learning process.

V. Dunjko et al. *Phys. Rev. Lett.* **117**, 130501 (2016)

V. Saggio et al. *Nature* **591,** 229–233 (2021)

What is quantum in the agent and the environment?

Let's recall the slide

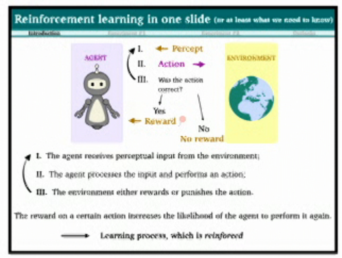

In a quantum-quantum framework, percepts, actions and rewards are promoted to quantum states!

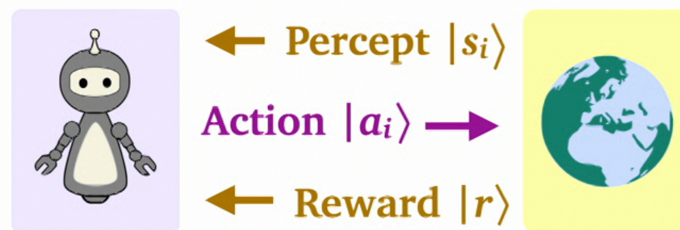

This implies that agent and environment can exchange signals in arbitrary quantum superpositions.

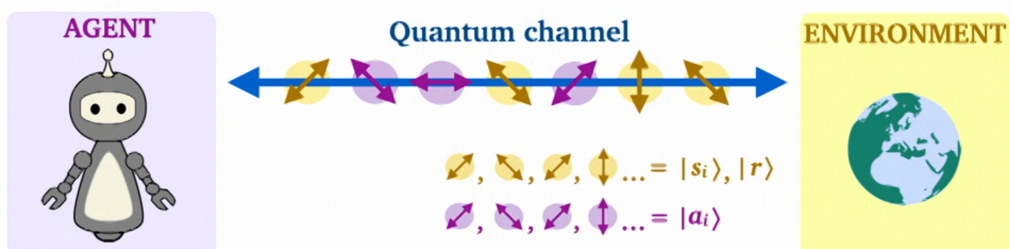

More in detail: the Grover algorithm

- Database of N elements ($x_1, x_2, x_3, \ldots, x_N$)
- Element one wants to find (target element): $|x_2\rangle = |win\rangle$
- Initial state: $|s\rangle = \dfrac{|x_1\rangle + |x_2\rangle + |x_3\rangle + \ldots + |x_N\rangle}{\sqrt{N}}$
- State orthogonal to the target element: $\dfrac{|x_1\rangle + |x_3\rangle + \ldots + |x_{N-1}\rangle}{\sqrt{N-1}} = |lose\rangle$

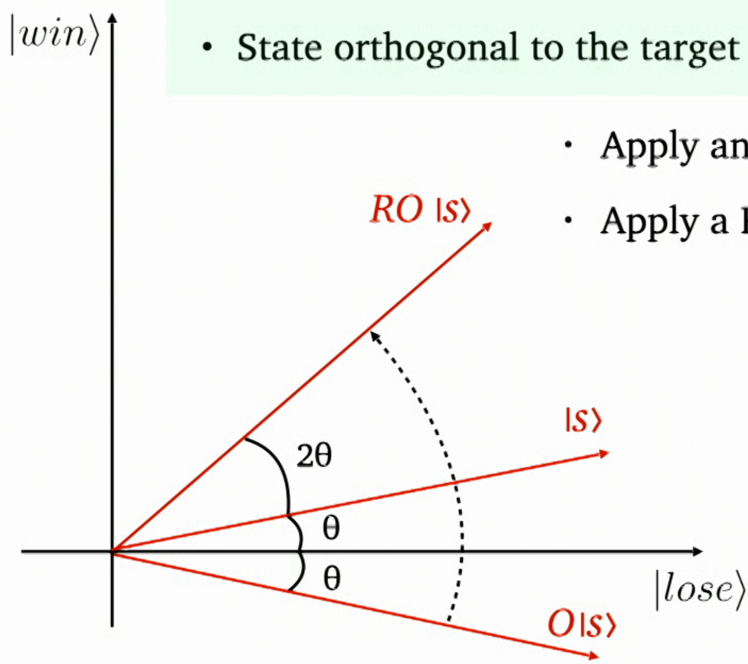

- Apply an Oracle O to $|s\rangle \longrightarrow \dfrac{|x_1\rangle - |x_2\rangle + |x_3\rangle + \ldots + |x_N\rangle}{\sqrt{N}}$
- Apply a Reflection R to $O|s\rangle$

$$\longrightarrow \left(\frac{N-4}{N\sqrt{N}}\right)|x_1\rangle + \left(\frac{3N-4}{N^{3/2}}\right)|x_2\rangle + \ldots + \left(\frac{N-4}{N\sqrt{N}}\right)|x_N\rangle$$

More in detail: the Grover algorithm

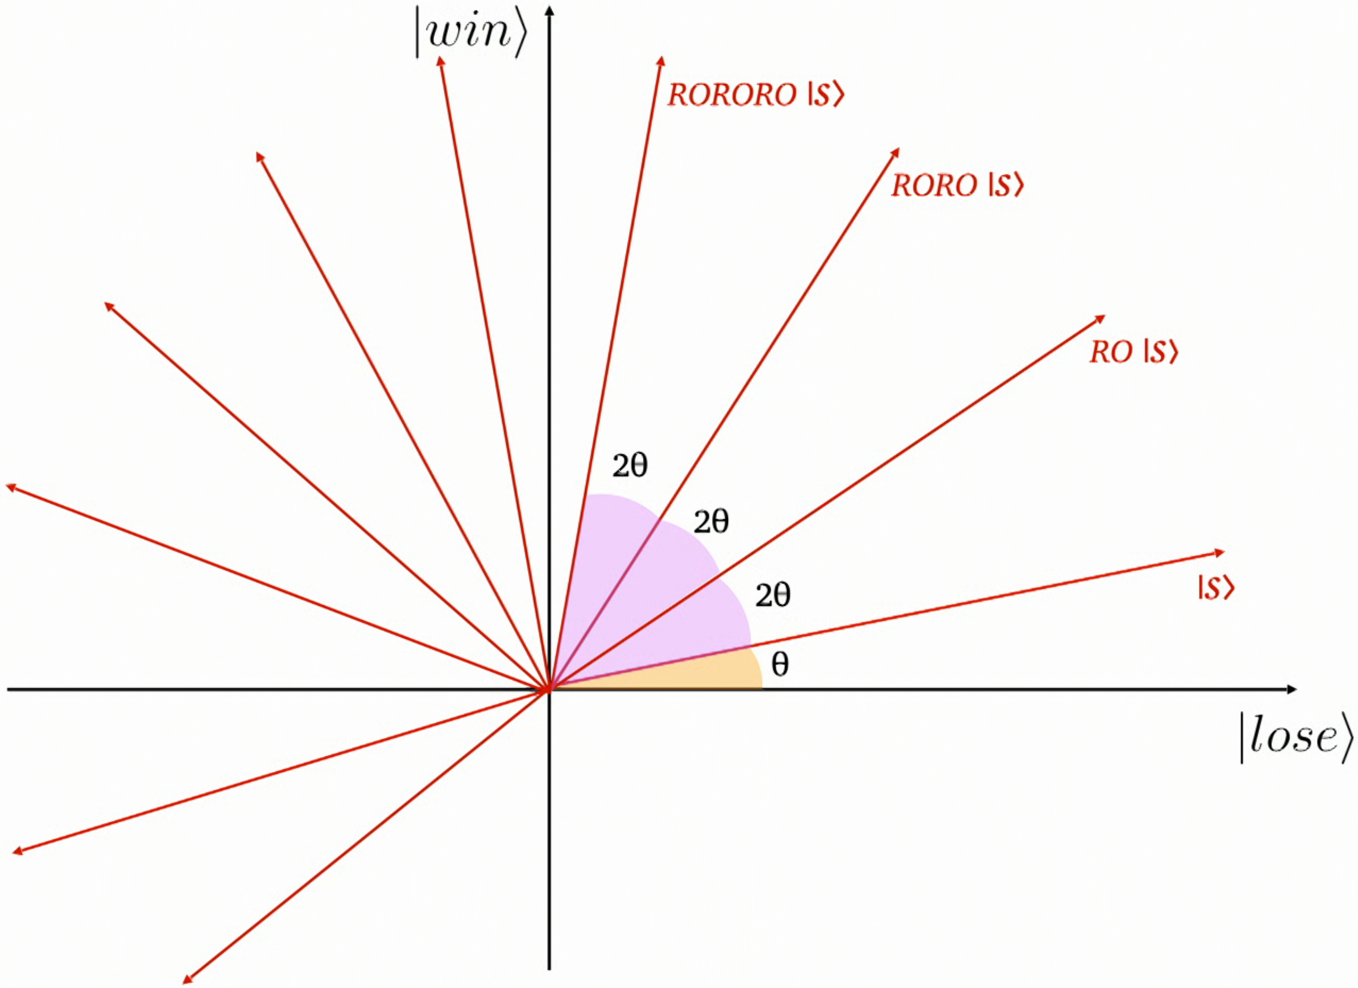

More in detail: the Grover algorithm

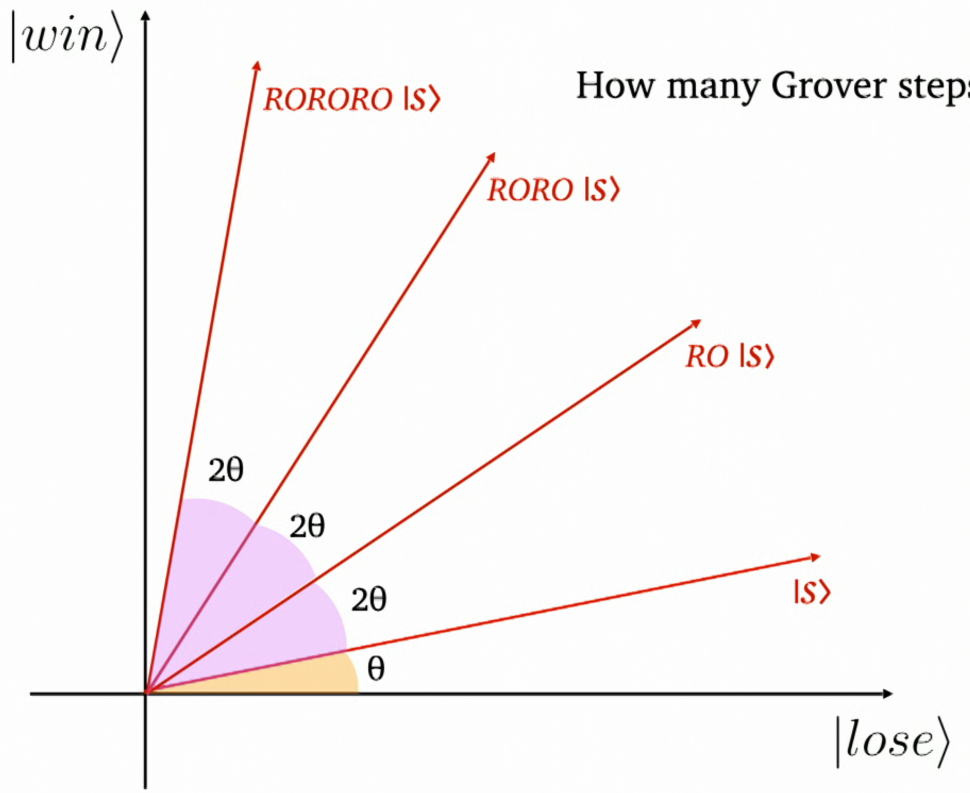

How many Grover steps do we need to find $|win\rangle$?

$O(\sqrt{N})$ trials are needed to maximize the winning probability!

Let's go back to reinforcement learning

Recap

The agent can use the Grover algorithm to find correct actions (winning states) *faster*.

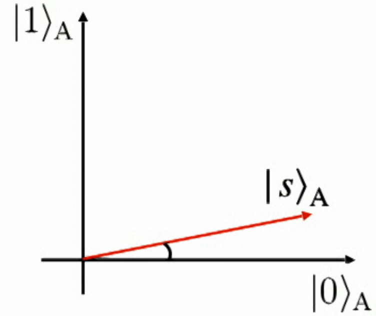

- Target element (correct action) $= |1\rangle_A$
- Other elements (wrong actions) $= |0\rangle_A$

As already seen, the initial state $|s\rangle_A = \sqrt{\varepsilon}\,|1\rangle_A + \sqrt{1-\varepsilon}\,|0\rangle_A$ can be prepared.

But how to experimentally encode the $|1\rangle_A$ and $|0\rangle_A$ states, and how to create the superposition $|s\rangle_A$?

$|s\rangle_A$ is a qubit, which we can implement using photons.

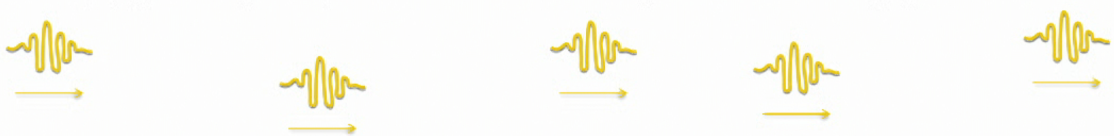

Quantum superposition with photons

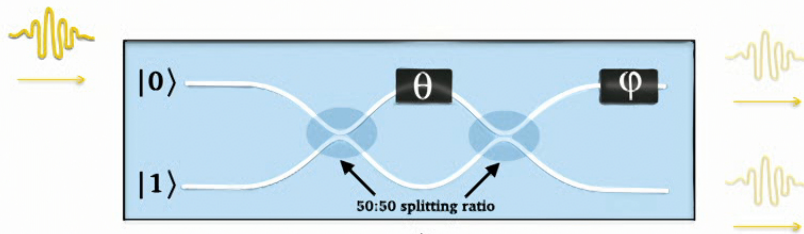

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$$

Tunable beam-splitter

Image of the beam splitter

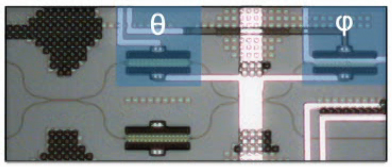

Image of the photonic chip

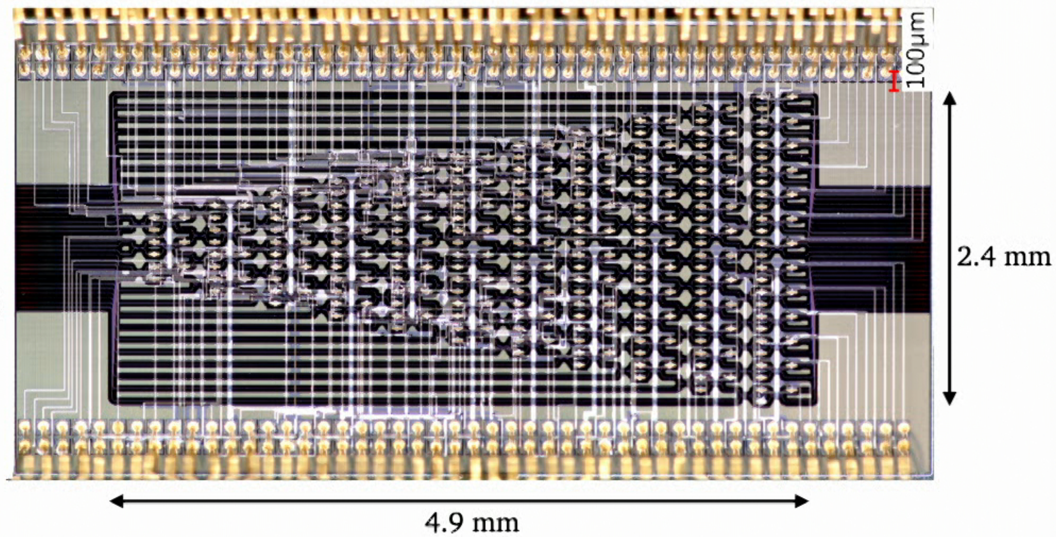

88 tunable beam splitters and 26 spatial modes

Quantum interaction (quantum epoch)

- Prepare the superposition $|s\rangle_A = \sqrt{\varepsilon}\,|1\rangle_A + \sqrt{1-\varepsilon}\,|0\rangle_A$
- Use $|0\rangle_R$ and $|1\rangle_R$ to encode the reward, and put them in superposition
- Apply the oracle $\longrightarrow$ $|s\rangle_A = -\sqrt{\varepsilon}\,|1\rangle_A + \sqrt{1-\varepsilon}\,|0\rangle_A$
- Apply the reflection $\longrightarrow$ $|s\rangle_A = \sqrt{\varepsilon}(3-4\varepsilon)\,|1\rangle_A + \sqrt{1-\varepsilon}(1-4\varepsilon)\,|0\rangle_A$

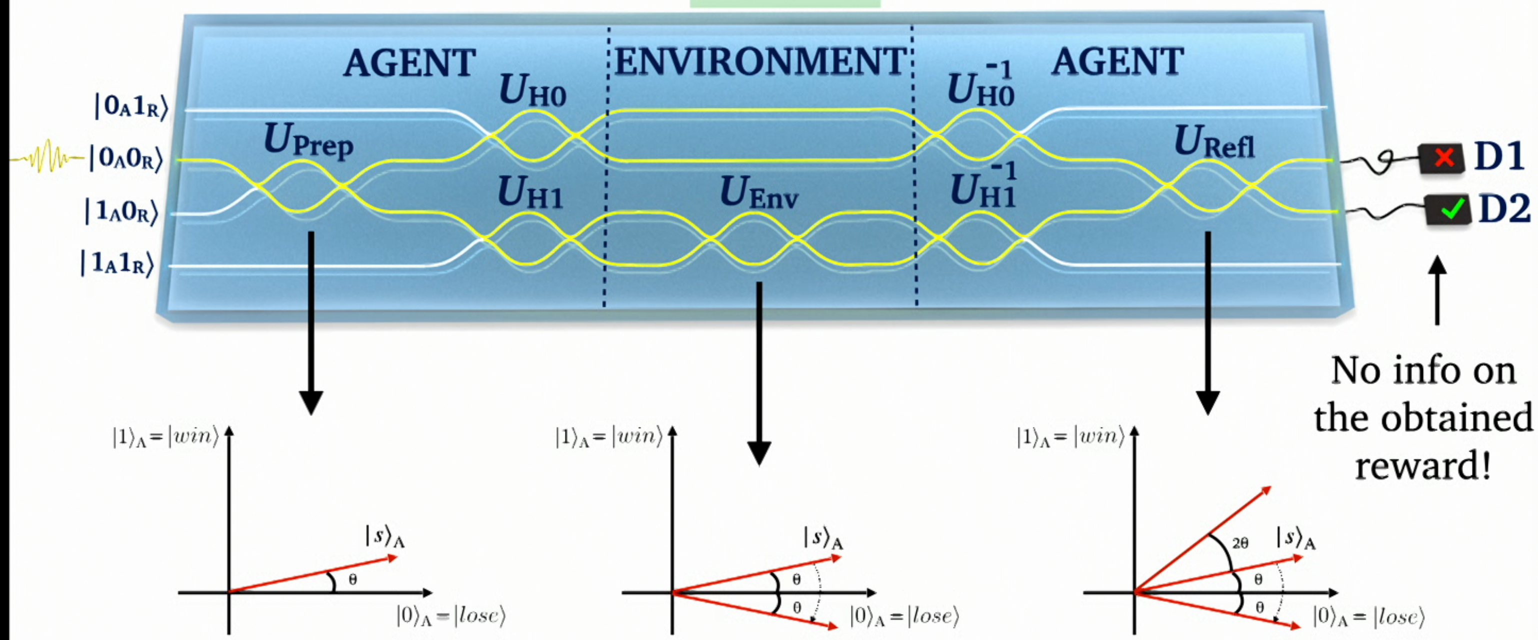

Obtaining the reward classically (classical test epoch)

Use the measured action as input in a classical epoch.

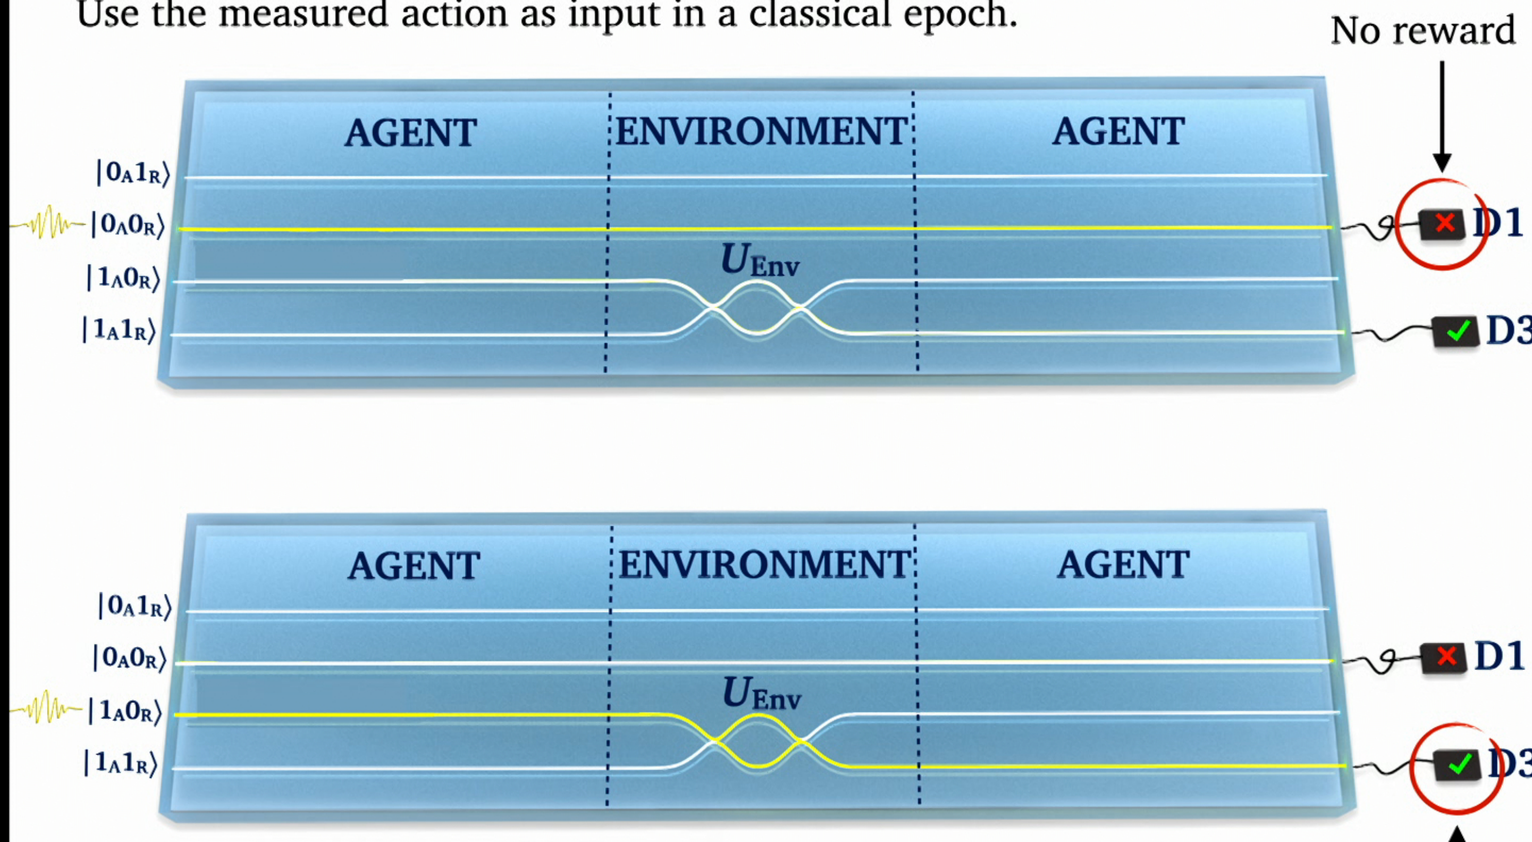

The agent can now update ε (to ε') and start a new round preparing $|s\rangle_A = \sqrt{\varepsilon'}|1\rangle_A + \sqrt{1-\varepsilon'}|0\rangle_A$.

The hybrid agent

The agent alternates between quantum and classical test epochs.

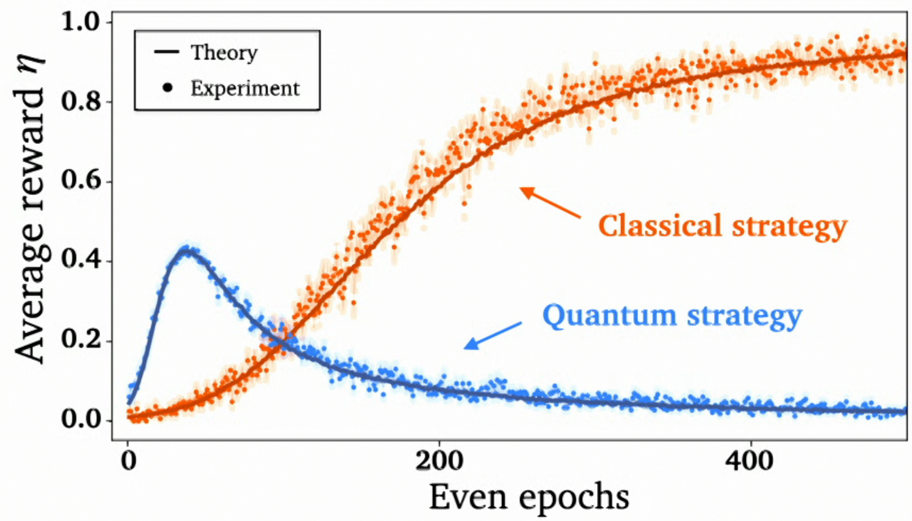

The hybrid agent

The agent alternates between quantum and classical test epochs, only as long as $\varepsilon_{\text{switch}} < 0.396$, and plays classically from that point on.

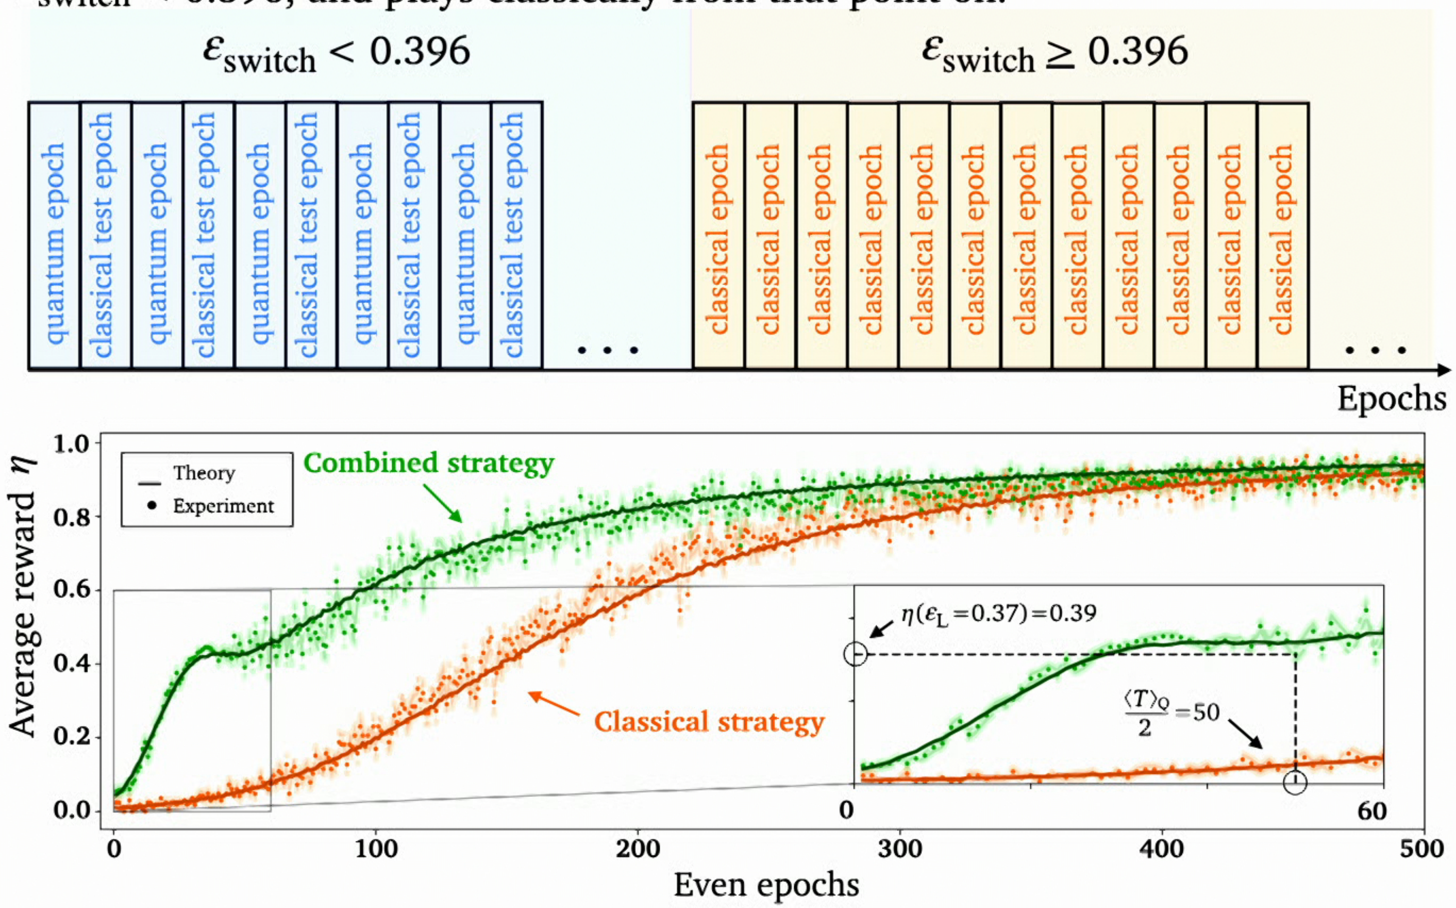

Reduction in the learning time

Learning time: number of epochs necessary to achieve, on average, a certain probability ε_{L} (smaller than 0.396).

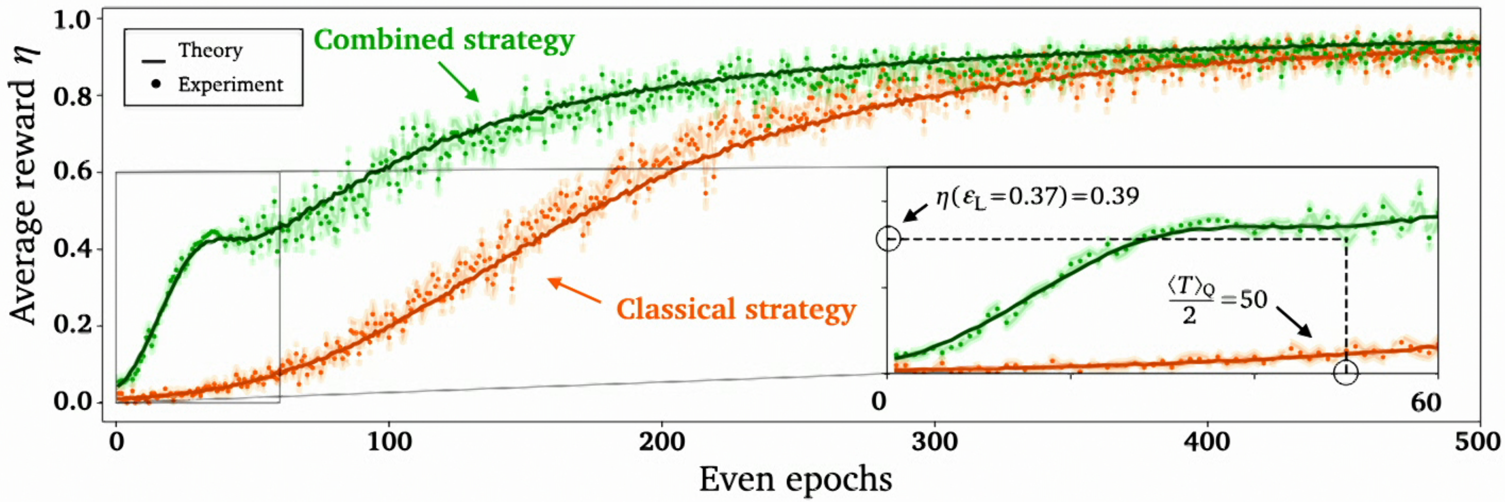

Combining reinforcement learning and quantum mechanics

		
Classical	Quantum	→ Classical learning techniques facing quantum tasks, for example the design of quantum experiments. A. Melnikov et al. *PNAS* **115**, 1221 (2018)
Quantum	Classical	→ Quadratic speed-up in the agent's decision time. G. D. Paparo et al. *Phys. Rev. X* **4(3)**, 031002 (2014); T. Sriarunothai et al. *Quantum Sci. Technol.* **4**, 015014 (2019)
Quantum	Quantum	→ Speed-up in the agent's learning process. V. Dunjko et al. *Phys. Rev. Lett.* **117**, 130501 (2016); V. Saggio et al. *Nature* **591,** 229–233 (2021)

ML for quantum

Quantum for ML

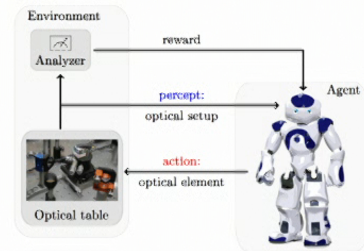

Designing quantum experiments with ML

Suppose we want to create a specific type of complex entangled state in the lab.

(e.g. high-dimensional multi-particle states)

We'd need to figure out the specific optical components (not always easy!)

What if we assign this task to a machine?

RESEARCH ARTICLE | PHYSICAL SCIENCES

Active learning machine learns to create new quantum experiments

Alexey A. Melnikov, Hendrik Poulsen Nautrup, Mario Krenn, Vedran Dunjko, Markus Tiersch, Anton Zeilinger, and Hans J. Briegel +3 Authors Info & Affiliations

Contributed by Anton Zeilinger, November 14, 2017 (sent for review August 24, 2017; reviewed by Jacob D. Biamonte and Jonathan P. Dowling)

January 18, 2018 | 115 (6) 1221-1226 | https://doi.org/10.1073/pnas.1714936115

The theoretical scheme

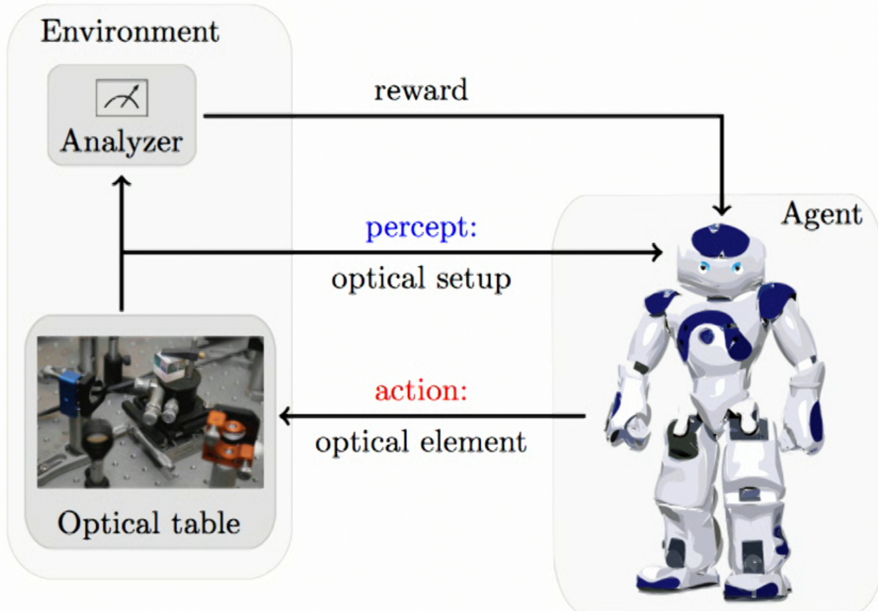

1) The agent places a chosen element on the optical table;

2) The quantum state generated by the setup is analyzed;

3) If the experiment is successful, a reward is given.

In a little more detail

- A maximum number of optical elements is considered (due to accumulation of imperfections).

- The produced quantum states are analyzed considering the Schmidt-Rank vector.

- Tripartite entangled states are considered - examples of Schmidt-Rank vectors are e.g. (3,3,2) or (3,3,3).

- Other than successfully generating the desired state, the agent should also use the lowest possible number of optical elements.

Designing short experiments and learning to create new ones

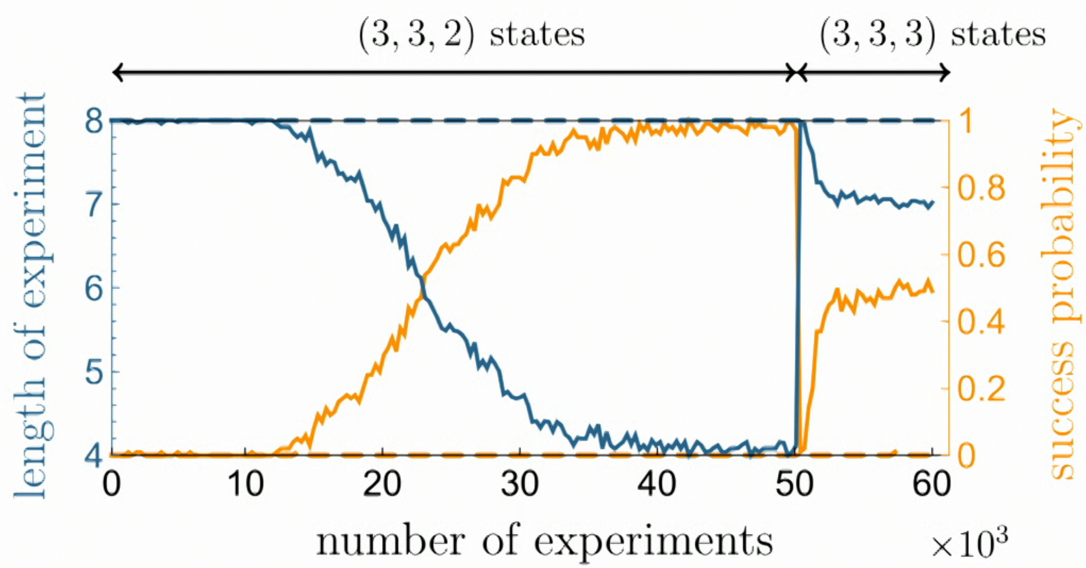

Interestingly, the agent without previous training on (3,3,2) states does not succeed in creating (3,3,3) states.

↓

The training is beneficial also for new experiments!

Not only reinforcement learning

Also supervised and unsupervised learning for quantum (and vice versa):

- Supervised learning with quantum computers.
  M. Schuld et al. Vol. 17. Berlin: Springer, 2018.
- Quantum (exponential) speed-up for supervised and unsupervised machine learning algorithms for cluster assignment and cluster finding.
  S. Lloyd et al. *arXiv:1307.0411* (2013)
- Accelerating unsupervised learning algorithms by quantizing some of their subroutines.
  E. Aïmeur et al. *Machine Learning* **90**, 261-287 (2013)
- Experimental learning of quantum states.
  A. Rocchetto et al. *Science advances* **5**, 3 (2019)
- Training Gaussian boson sampling by quantum machine learning.
  C. Conti, *Quantum Machine Intelligence* **3**, 26 (2021)
- Detecting entanglement with unsupervised learning (potential scalability advantage).
  Y. Chen et al., *Quantum Science and Technology*, **7**(1), 015005 (2021)
- Reconstructing unknown quantum processes (quantum process tomography).
  G. Torlai et al., *arXiv:2006.02424* (2020)

Conclusions

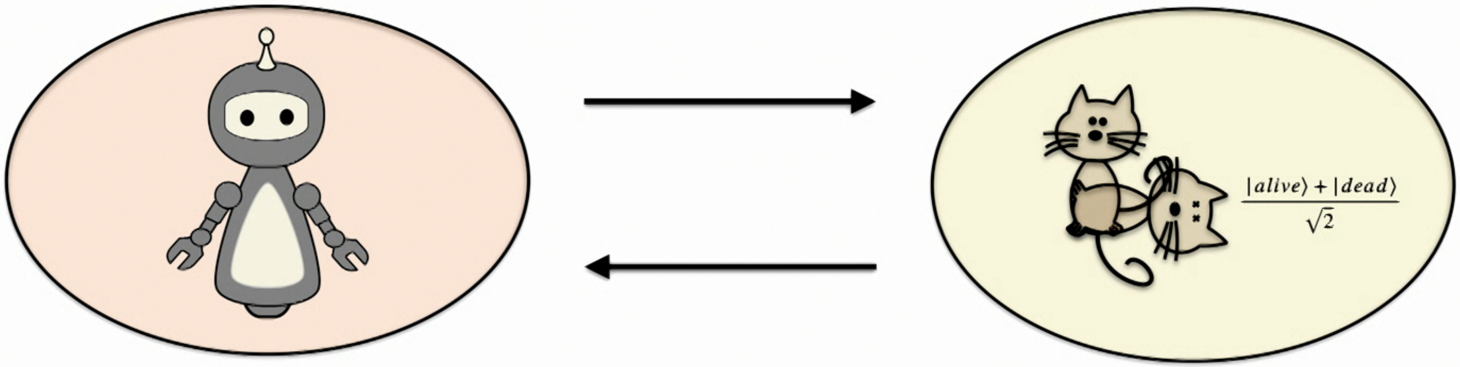

- Embed quantum algorithms in a ML framework to prove quantum advantage;
- Use classical ML in quantum experiments.

However, there are still tons of possibilities to explore (not only with photonics)!