Title: Keynote

Speakers: Friederike Metz

Collection: PSI 15th Anniversary Reunion

Date: June 19, 2024 - 4:00 PM

URL: https://pirsa.org/24060007

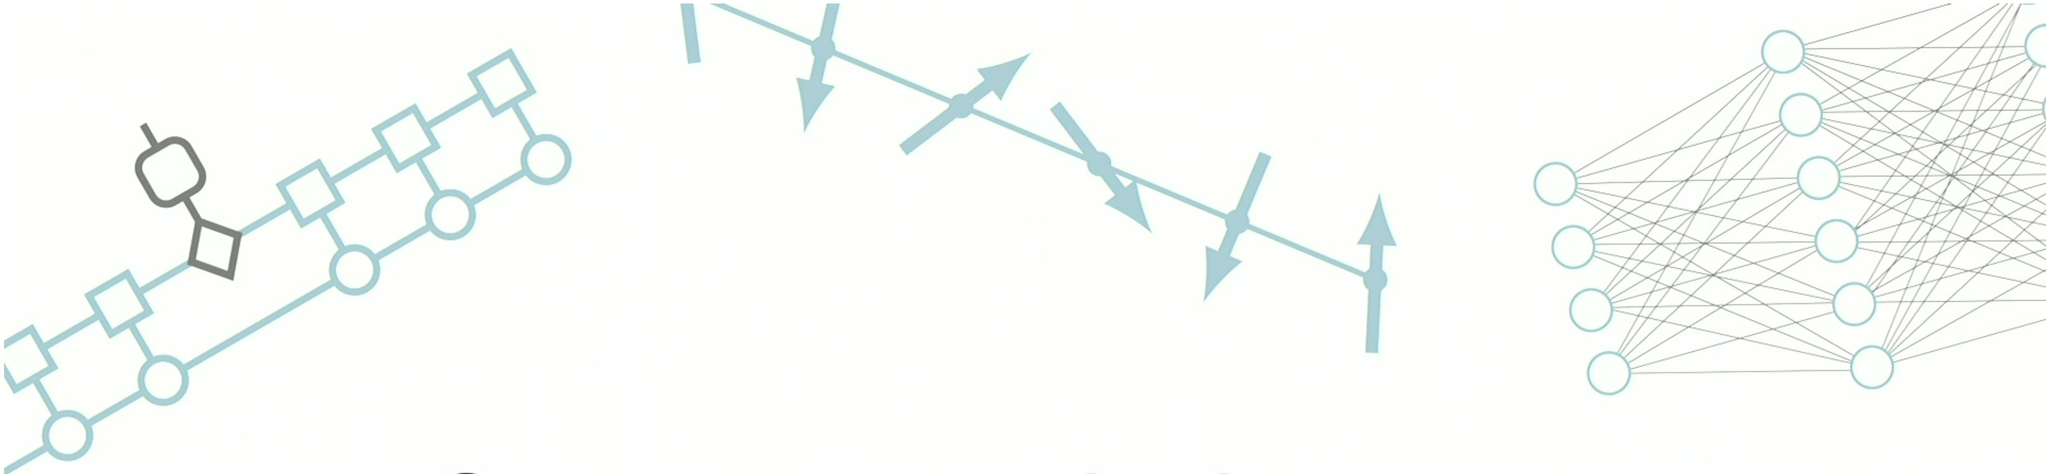

Quantum many-body systems

how to tame and control them

Friederike Metz @ EPFL - Switzerland

19.06.2024 — PSI 15th Anniversary Reunion

OUTLINE

- My academic journey
- My research interests
  - Computational quantum many-body physics
  - Quantum control & reinforcement learning
- Efficient quantum many-body control using reinforcement learning

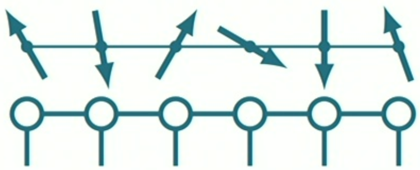

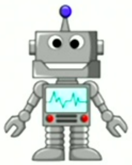

FM, M. Bukov, *Nature Machine Intelligence* **5**, 780-791 (2023)

MY JOURNEY

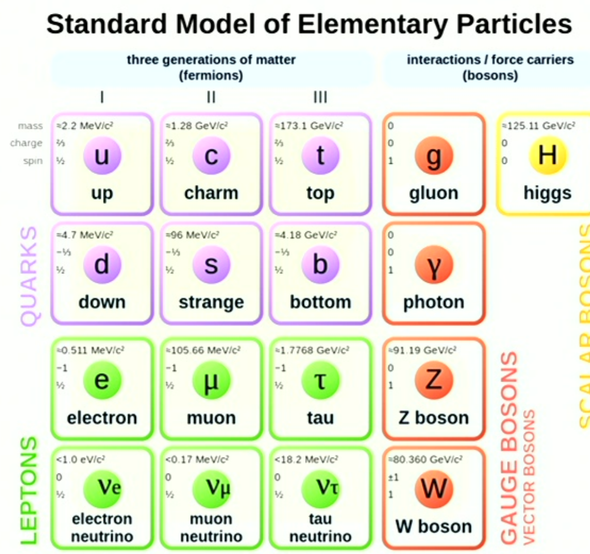

Image from Wikimedia

High school

— 2011

MY JOURNEY

Undergraduate

2011–2015

MY JOURNEY

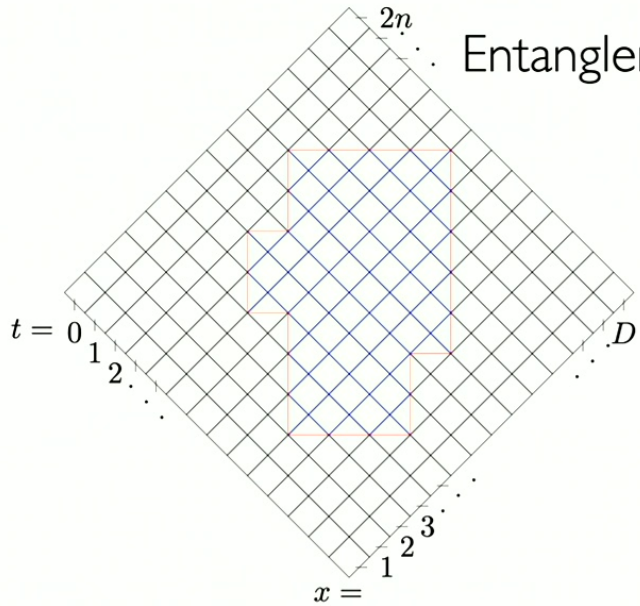

Entanglement entropy of space-time regions

Barbara Terhal

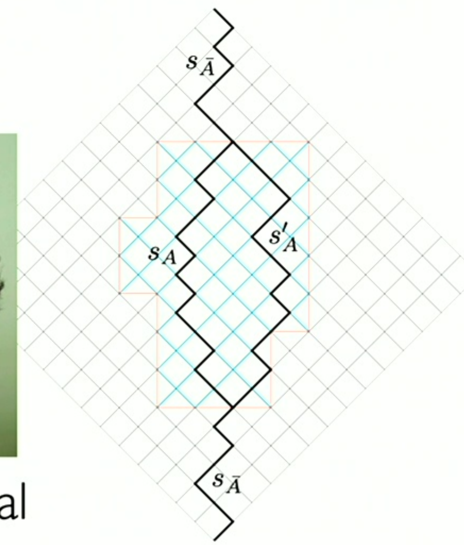

Undergraduate

2011–2015

MY JOURNEY

"PSI is like trying to drink from a firehose."

Image from DALL·E

PSI Program

2015–2016

MY JOURNEY

Effective Theory of Monopole
- Fermion Scattering

Cliff Burgess

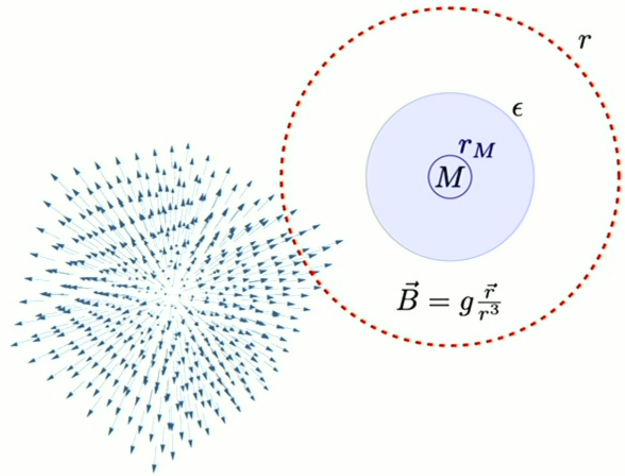

PSI Program

2015–2016

MY JOURNEY

Studying ≠ Doing research

PSI Program

2015–2016

MY JOURNEY

OIST

Okinawa Institute of
Science and Technology

PhD

2017–2023

MY JOURNEY

Reinforcement learning

D. Silver et al. Nature 529, 484 (2016)

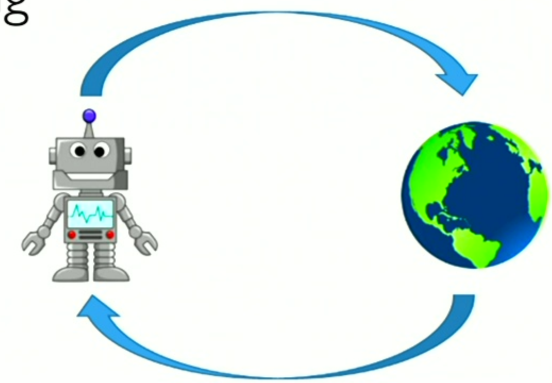

OIST

Okinawa Institute of
Science and Technology

PhD

2017–2023

MY JOURNEY

Machine learning applications to study and control quantum systems

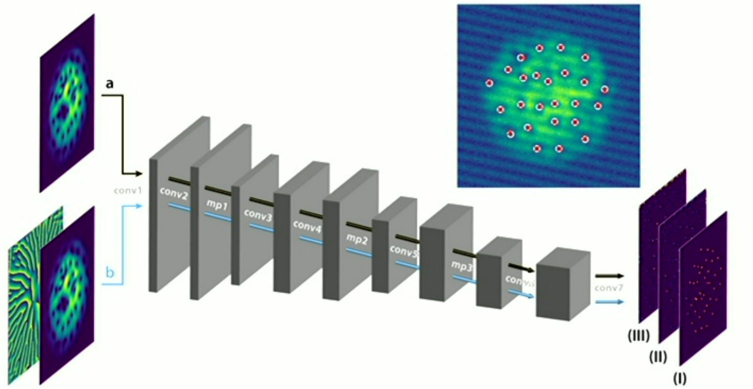

FM, J. Polo, N. Weber, T. Busch, *Mach. Learn.: Sci. Technol.* 2, 035019 (2021)

Deep-learning based vortex detection in rotating BECs

PhD

2017–2023

MY JOURNEY

Machine learning applications to study and control quantum systems

Deep-learning bas
detection in rotat

Variational quantum ar
Unsupervised mapping

Quantum many-body control using reinforcement learning with tensor networks

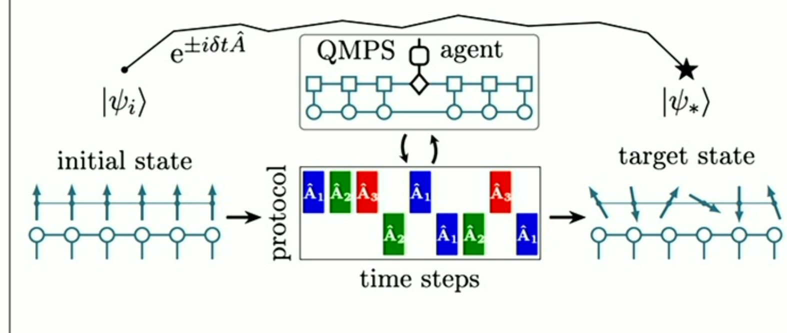

FM, M. Bukov, *Nature Machine Intelligence* 5, 780–791 (2023)

PhD

2017–2023

MY JOURNEY

Machine learning applications to study and control quantum systems

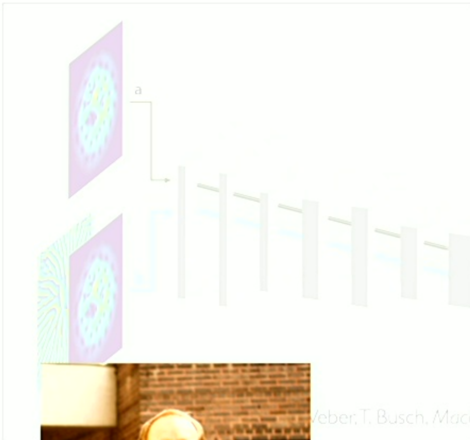

Variational quantum ar
Unsupervised mapping

K. Kottmann, FM, J. Fraxanet, N. Baldelli, *Phys. Rev. Rese*

rning bas
n in rotat

Thomas Busch

Quantum many-body control using reinforcement learning with tensor networks

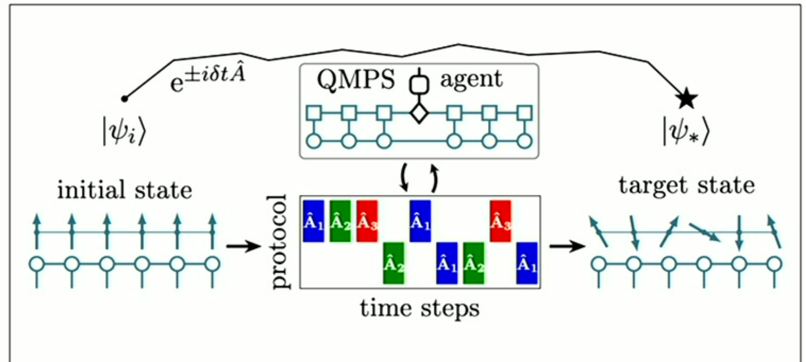

FM, M. Bukov, *Nature Machine Intelligence* 5, 780–791 (2023)

PhD

2017–2023

MY JOURNEY

MAX-PLANCK-GESELLSCHAFT

Marin Bukov

Max Planck Institute for the
Physics of Complex Systems

Visiting researcher

2022

MY JOURNEY

MAX-PLANCK-GESELLSCHAFT

Marin Bukov

Quantum many-body control & reinforcement learning

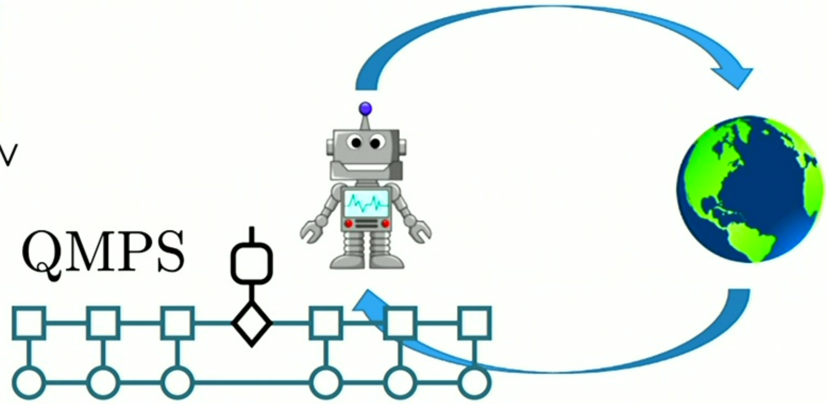

Max Planck Institute for the Physics of Complex Systems

Visiting researcher

2022

MY JOURNEY

Quantum software development

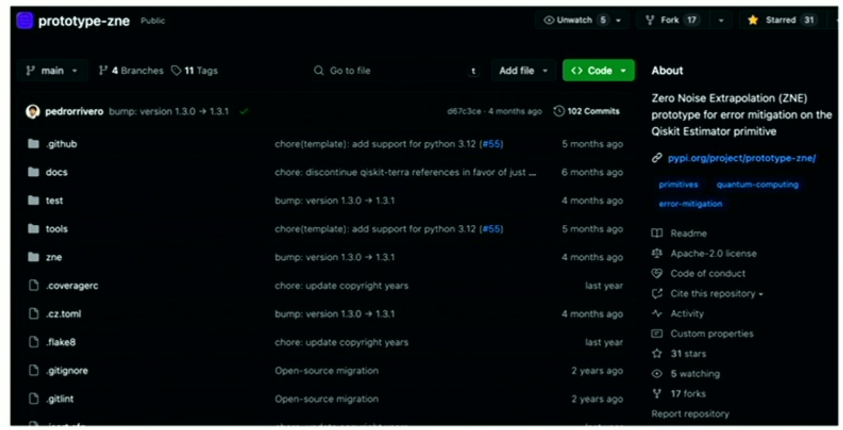

Internship

2022

BIG PICTURE

Quantum chemistry

- Better catalysts
- New drugs
- More efficient batteries
- Carbon capture & storage

Image from DALL·E

Materials science

Condensed matter physics

- High temperature superconductivity
- Topologically ordered states
- Quantum spin liquids
- Exotic states of matter

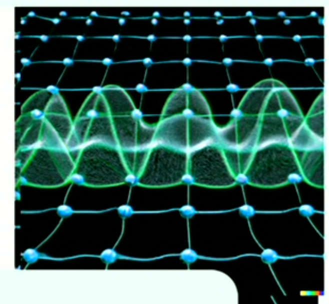

DALL·E

Quantum physics

More is different. - *Philip W. Anderson*

Quantum information

Atomic, molecular, and optical physics

BIG PICTURE

Quantum chemistry
- Better catalysts
- New drugs
- More efficient batteries
- Carbon capture & storage

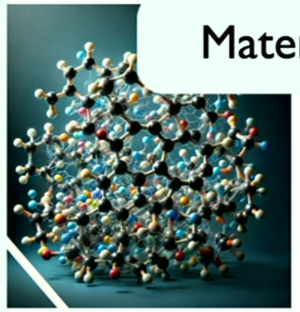

Image from DALL·E

Materials science

Condensed matter physics
- High temperature superconductivity
- Topologically ordered states
- Quantum spin liquids
- Exotic states of matter

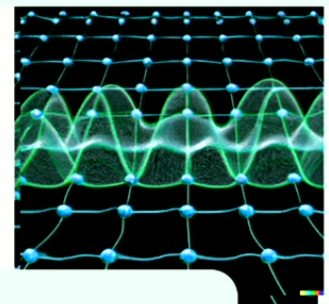

DALL·E

Quantum many-body physics

More is different. - *Philip W. Anderson*

Quantum information

Atomic, molecular, and optical physics

CURSE OF DIMENSIONALITY

Hilbert space dimension grows exponentially in system size!

N Spin-$1/2$ particles:

$$\uparrow \; \downarrow \; \downarrow \; \uparrow \; \downarrow \; \uparrow \; \uparrow \; \downarrow$$

$$|\boldsymbol{\sigma}\rangle = |\; 0 \;\; 1 \;\; 1 \;\; 0 \;\; 1 \;\; 0 \;\; 0 \;\; 1 \;\rangle$$

2^N dimensional

$$|\Psi\rangle = \Psi_1|00\ldots0\rangle + \Psi_2|10\ldots0\rangle + \Psi_3|01\ldots0\rangle + \cdots + \Psi_{2^N}|11\ldots1\rangle = \sum_{\boldsymbol{\sigma}} \Psi_{\boldsymbol{\sigma}}|\boldsymbol{\sigma}\rangle$$

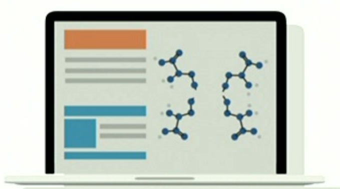

1 Petabyte: $10^{15} \sim 2^{50}$

TENSOR NETWORKS

Compressed representations/approximations of exponentially large quantum states

Wave function vector describing N spin-$1/2$: $\Psi_{\boldsymbol{\sigma}} = \Psi_{\sigma_1, \sigma_2, \ldots, \sigma_N}$ $\qquad \sigma_i = 0, 1$ $\quad \uparrow \downarrow$

TENSOR NETWORKS

Compressed representations/approximations of exponentially large quantum states

Tensor diagrams 101 *Roger Penrose 1971*

Tensor with N indices = Shape with N legs

$$\Psi_{\sigma_1,\sigma_2,\ldots,\sigma_N} =$$

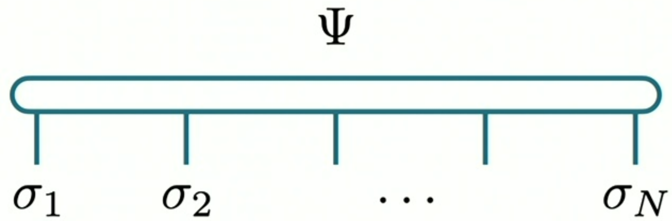

Important examples:

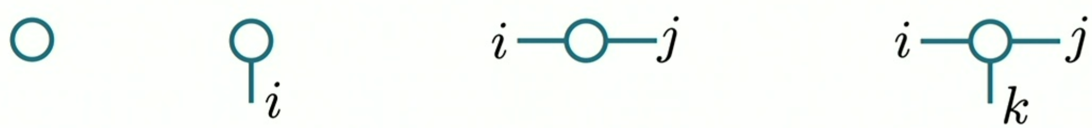

S $\quad$ V_i $\quad$ M_{ij} $\quad$ T_{ijk}

J. C. Bridgeman, C. T. Chubb, J. Phys. A: Math. Theor. 50 223001 (2017)

TENSOR NETWORKS

Compressed representations/approximations of exponentially large quantum states

Tensor diagrams 101 *Roger Penrose 1971*

Linear algebra operations:

$$\sum_i V_i W_i$$

$$\sum_j A_{ij} B_{jk}$$

Important examples:

S

V_i

M_{ij}

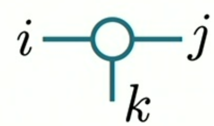

T_{ijk}

J. C. Bridgeman, C. T. Chubb, J. Phys. A: Math. Theor. 50 223001 (2017)

TENSOR NETWORKS

Compressed representations/approximations of exponentially large quantum states

Tensor diagrams 101 *Roger Penrose 1971*

Linear algebra operations:

$$\sum_i V_i W_i$$

$$\sum_j A_{ij} B_{jk}$$

Singular value decomposition & truncation

$M \quad = \quad U \quad \tilde{S} \quad V^\dagger \quad = \quad A_1 \quad A_2$ → lower dimensional

J. C. Bridgeman, C. T. Chubb, J. Phys. A: Math. Theor. 50 223001 (2017)

TENSOR NETWORKS

Compressed representations/approximations of exponentially large quantum states

Matrix product state (MPS)

Tensor train (TT)

$$\Psi_{\sigma_1,\sigma_2,\ldots,\sigma_N} \approx$$

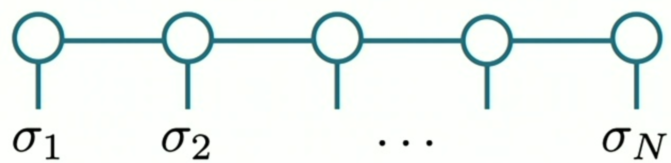

MPS reduces number of parameters: $2^N \rightarrow 2\chi^2 N$

J. C. Bridgeman, C. T. Chubb, J. Phys. A: Math. Theor. 50 223001 (2017)

TENSOR NETWORKS

Compressed representations/approximations of exponentially large quantum states

$$\Psi_{\sigma_1,\sigma_2,\ldots,\sigma_N} \approx$$

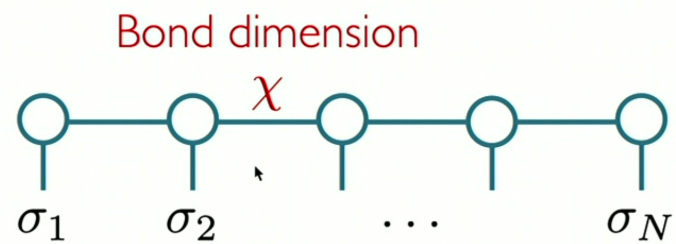

Matrix product state (MPS)

Tensor train (TT)

MPS reduces number of parameters: $2^N \rightarrow 2\chi^2 N$

Linear scaling!

J. C. Bridgeman, C. T. Chubb, J. Phys. A: Math. Theor. 50 223001 (2017)

TENSOR NETWORKS

Compressed representations/approximations of exponentially large quantum states

$\Psi_{\sigma_1,\sigma_2,\ldots,\sigma_N} \approx$

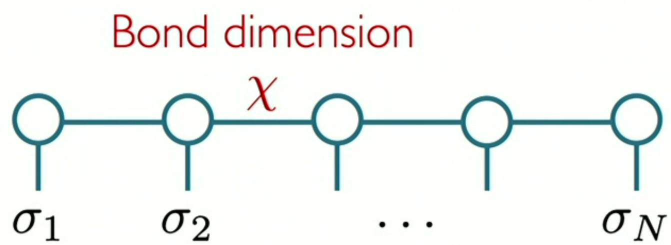

Matrix product state (MPS)

Tensor train (TT)

MPS reduces number of parameters: $2^N \rightarrow 2\chi(N)^2 N$

$$\chi \sim \begin{cases} \text{const} & \text{low-entangled} \\ 2^N & \text{highly-entangled} \end{cases}$$

J. C. Bridgeman, C. T. Chubb, J. Phys. A: Math. Theor. 50 223001 (2017)

MATRIX PRODUCT STATES (MPS)

MPS allow for efficient computations:

- E.g. wave function overlaps

$$\langle\Psi|\Psi'\rangle = \sum_{\boldsymbol{\sigma}} \Psi^*_{\boldsymbol{\sigma}} \Psi'_{\boldsymbol{\sigma}}$$

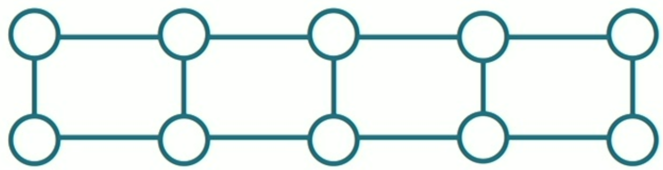

- Computing ground states: Density matrix renormalization group (DMRG) *Steven R. White 1992*

J. C. Bridgeman, C. T. Chubb, J. Phys. A: Math. Theor. 50 223001 (2017)

MATRIX PRODUCT STATES (MPS)

MPS allow for efficient computations:

- E.g. wave function overlaps

$$\langle\Psi|\Psi'\rangle = \sum_{\boldsymbol{\sigma}} \Psi^*_{\boldsymbol{\sigma}} \Psi'_{\boldsymbol{\sigma}}$$

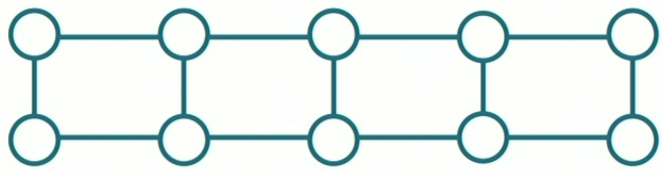

- Computing ground states: Density matrix renormalization group (DMRG) *Steven R. White 1992*

- Time evolution: $e^{-i\hat{H}t}|\Psi\rangle$

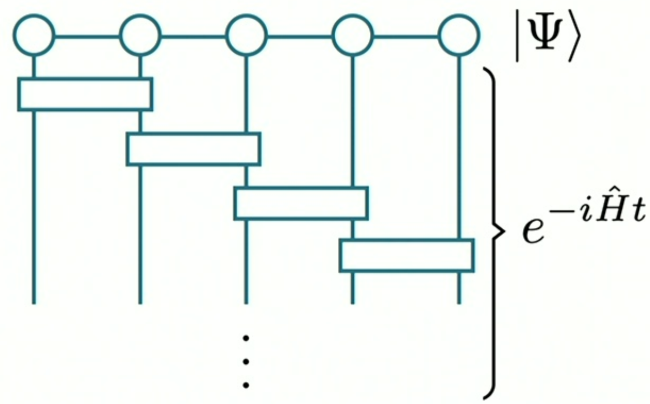

WAVE FUNCTION ANSATZE

Compressed representations/approximations of exponentially large quantum states

Tensor networks

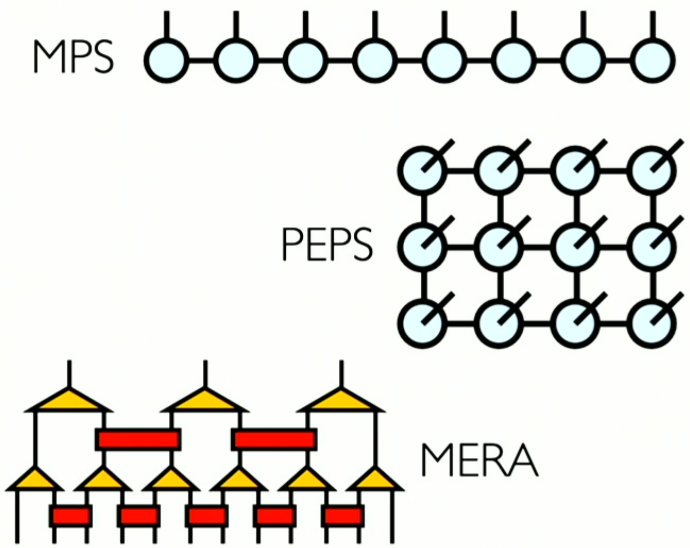

Images from *J. Reyes, M. Stoudenmire, arXiv:2001.08286 (2020)*

Neural quantum states

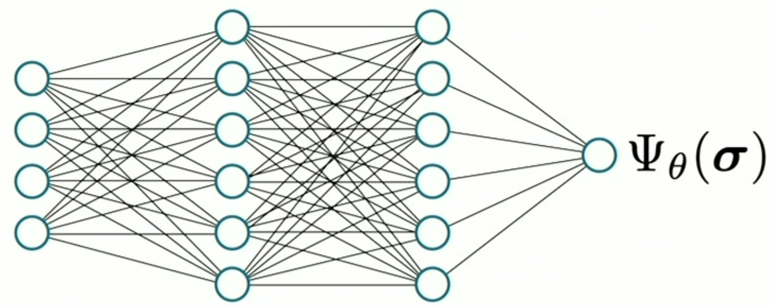

G. Carleo, M. Troyer, Science 355, 602 (2017)

Giuseppe Carleo

QUANTUM SIMULATORS

Nature isn't classical, dammit, and if you want to make a simulation of nature, you'd better make it quantum mechanical. - *Richard Feynman 1981*

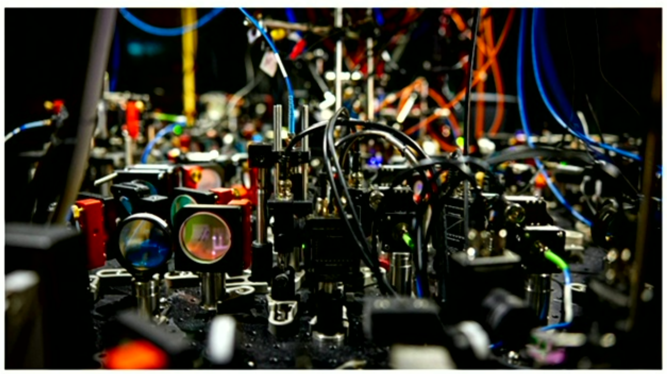

Optical table from Mikhail Lukin's cold atom experiment at Harvard
Image from Quanta Magazine

Experimental platforms

- Rydberg atoms trapped in optical tweezers
- Ultracold atoms in optical lattices
- Trapped ions
- Superconducting circuits
- Photonic systems
- …

TOY QUANTUM SIMULATOR

$$\hat{H}_{\mathrm{Ising}} = -J \sum_i \hat{Z}_i \hat{Z}_{i+1} - h_x \sum_i \hat{X}_i - h_z \sum_i \hat{Z}_i$$

Interactions

External magnetic fields

TOY QUANTUM SIMULATOR

$$\hat{H}_{\text{Ising}} = -J\sum_i \hat{Z}_i\hat{Z}_{i+1} - h_x\sum_i \hat{X}_i - h_z\sum_i \hat{Z}_i$$

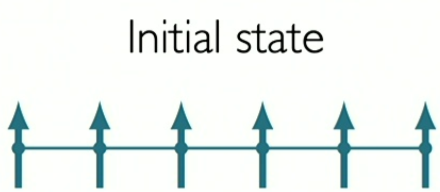

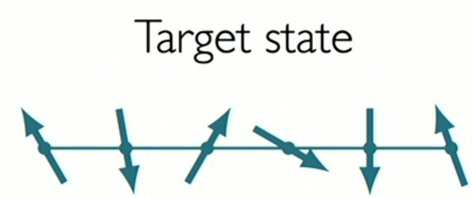

TOY QUANTUM SIMULATOR

$$\hat{H}_{\text{Ising}} = -J\sum_i \hat{Z}_i \hat{Z}_{i+1} - h_x \sum_i \hat{X}_i - h_z \sum_i \hat{Z}_i$$

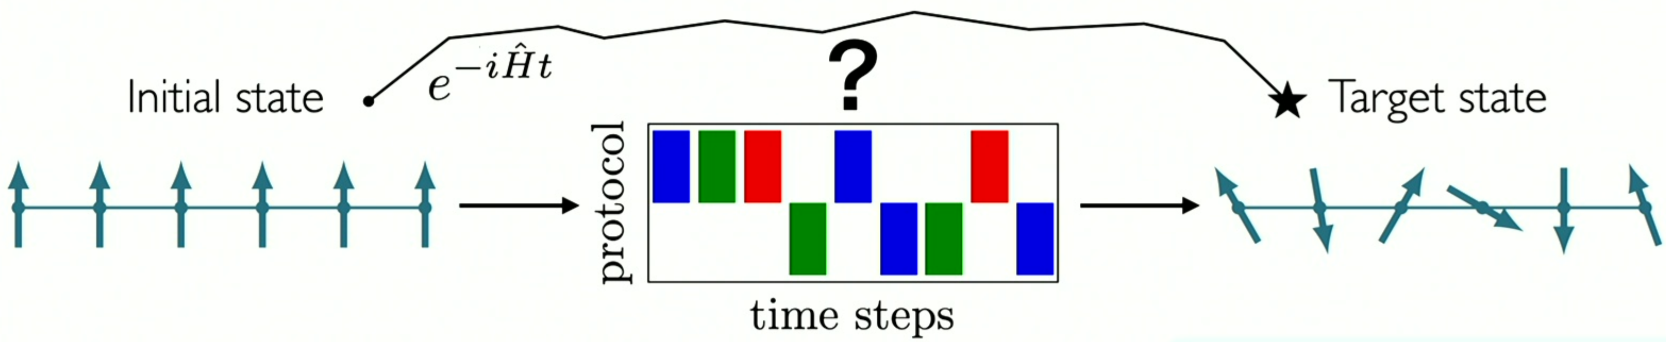

Optimal control problem!

Optimal control

- Greedy search
- Beam search
- …
- **Reinforcement learning**

REINFORCEMENT LEARNING (RL)

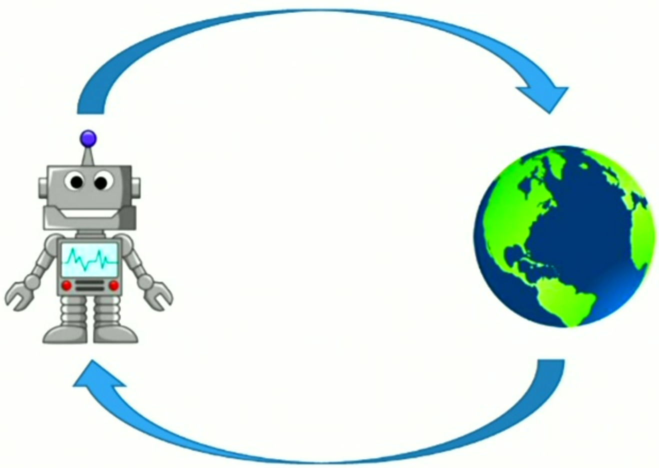

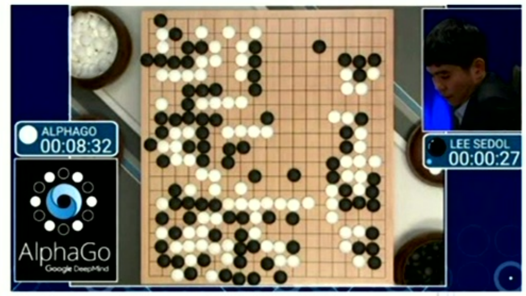

D. Silver et al. Nature 529, 484 (2016)

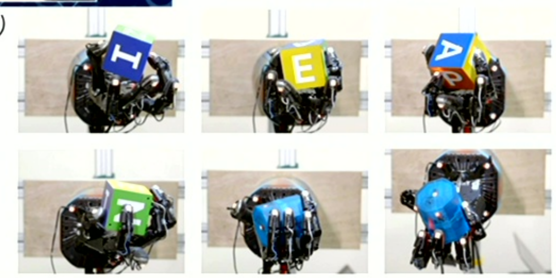

OpenAI et al. arXiv:1808.00177 (2019)

Explain reinforcement learning in three short bullet points

- **Interaction-based Learning**: Agents learn by interacting with an environment, receiving rewards or penalties for actions.
- **Trial and Error**: Agents improve strategies over time by trying actions and learning from outcomes.
- **Exploration vs. Exploitation**: Balances discovering new actions with using known, rewarding actions.

ChatGPT 4o by OpenAI

R. S. Sutton, A. G. Barto, Reinforcement learning: An introduction. MIT press

REINFORCEMENT LEARNING (RL)

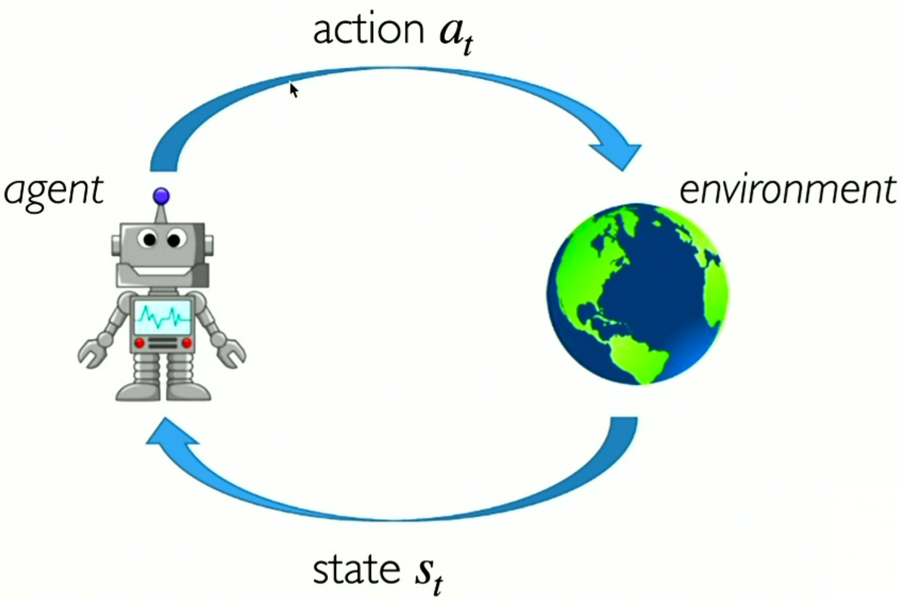

R. S. Sutton, A. G. Barto, Reinforcement learning: An introduction. MIT press

D. Silver et al. Nature 529, 484 (2016)

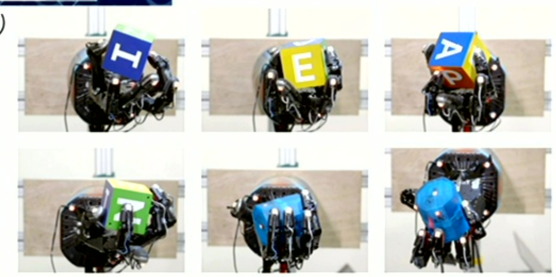

OpenAI et al. arXiv:1808.00177 (2019)

Explain reinforcement learning in three short bullet points

- **Interaction-based Learning**: Agents learn by interacting with an environment, receiving rewards or penalties for actions.
- **Trial and Error**: Agents improve strategies over time by trying actions and learning from outcomes.
- **Exploration vs. Exploitation**: Balances discovering new actions with using known, rewarding actions.

ChatGPT 4o by OpenAI

REINFORCEMENT LEARNING (RL)

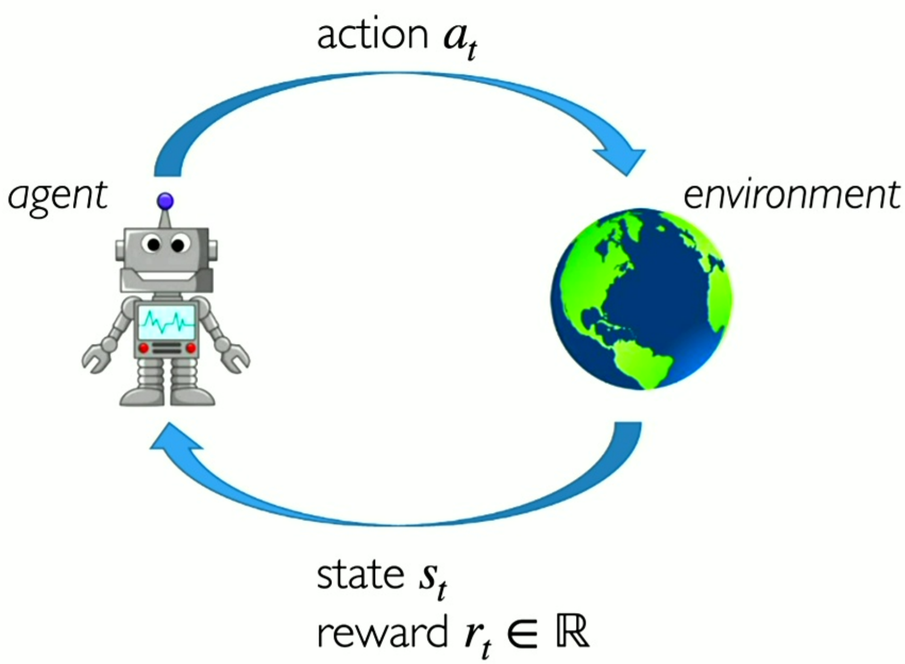

Agent chooses actions according to a policy (i.e. strategy):

$$\pi(s) = a$$

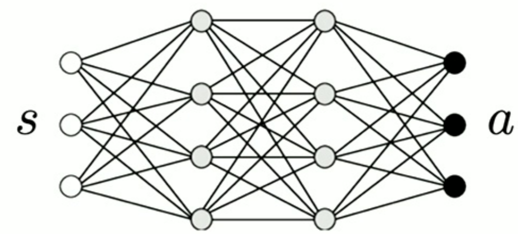

Goal: Maximize expected cumulative reward by finding an optimal policy

R. S. Sutton, A. G. Barto, Reinforcement learning: An introduction. MIT press

REINFORCEMENT LEARNING (RL)

RL agents can

- learn w/o knowing the physical laws of the environment
- produce protocols from different initial states w/o further optimization
- adapt their protocols in dynamically changing or stochastic environments

Agent chooses actions according to a policy (i.e. strategy):

$$\pi(s) = a$$

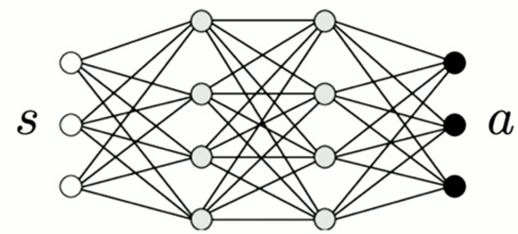

Goal: Maximize expected cumulative reward by finding an optimal policy

R. S. Sutton, A. G. Barto, Reinforcement learning: An introduction. MIT press

RL FOR QUANTUM MANY-BODY CONTROL

Challenge: **Optimization is expensive**

→ Perform numerical optimization in advance on classical computer

→ But: curse of dimensionality of quantum many-body systems

Our control framework: Deep reinforcement learning (RL) with matrix product states (MPS)

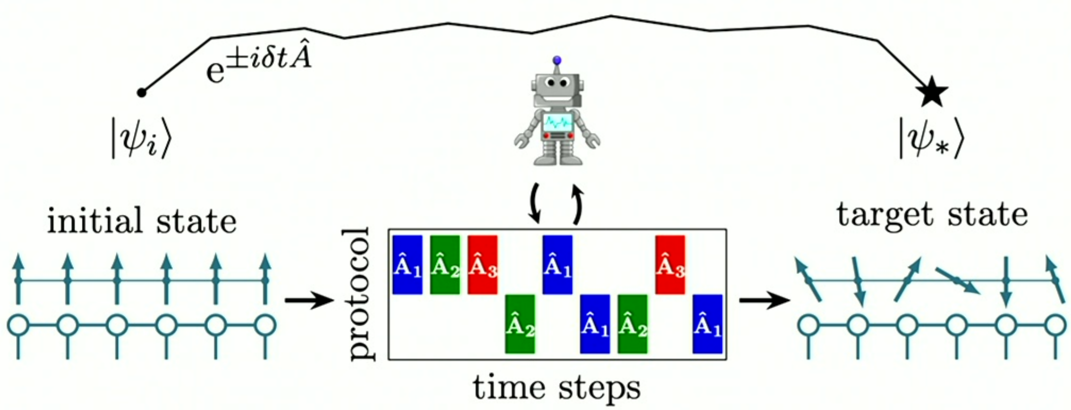

RL FRAMEWORK

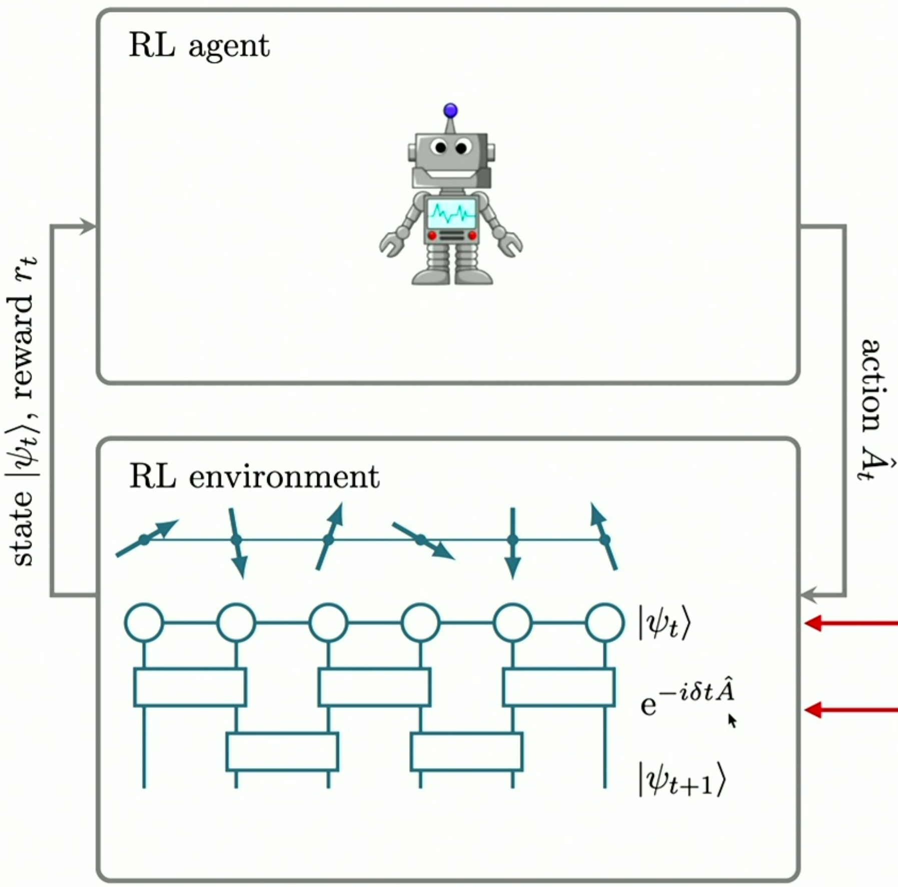

RL environment

state

action

$$\hat{A} \in \left\{ \sum_i \hat{Z}_i \hat{Z}_{i+1} , \sum_i \hat{X}_i , \sum_i \hat{Z}_i \right\}$$

RL FRAMEWORK

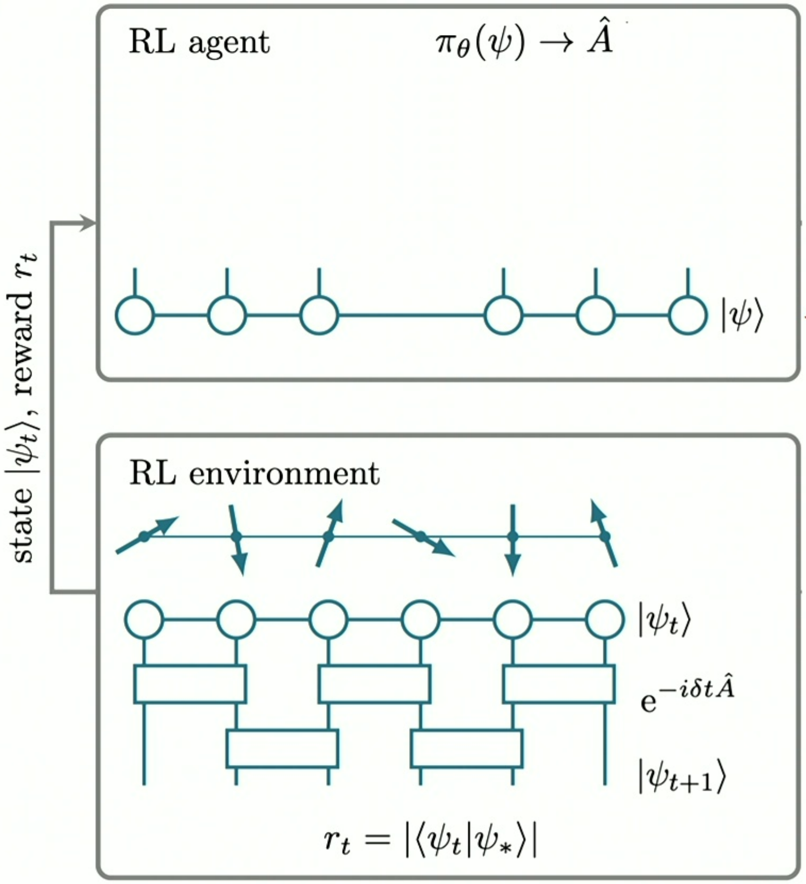

RL agent

Quantum state input

RL FRAMEWORK

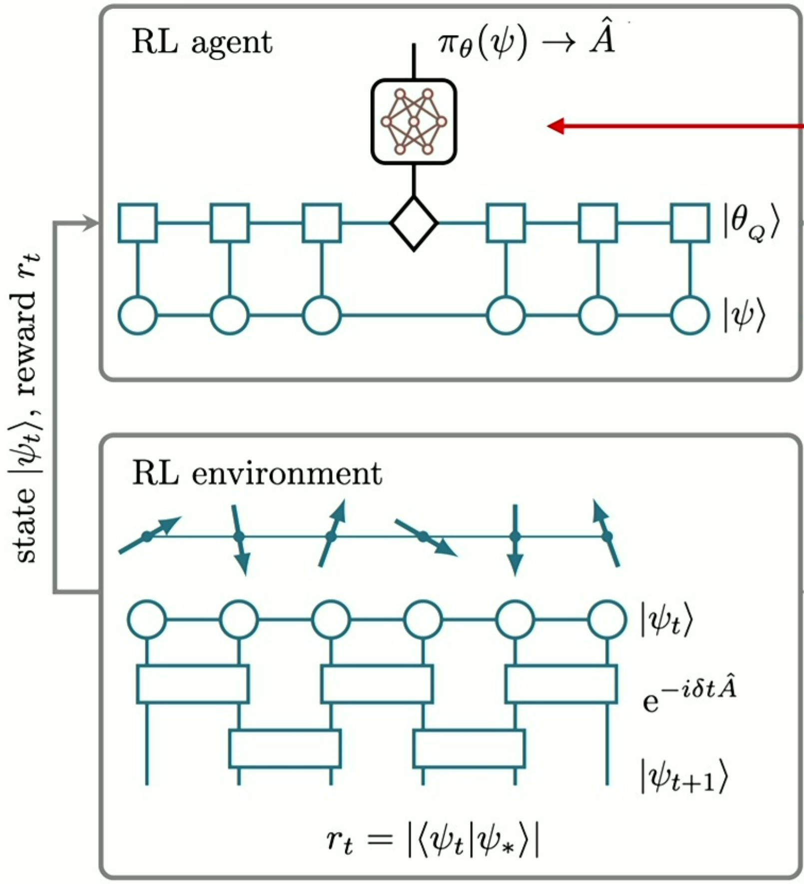

RL agent

Neural network

Trainable parameter MPS

Quantum state input

RL FRAMEWORK

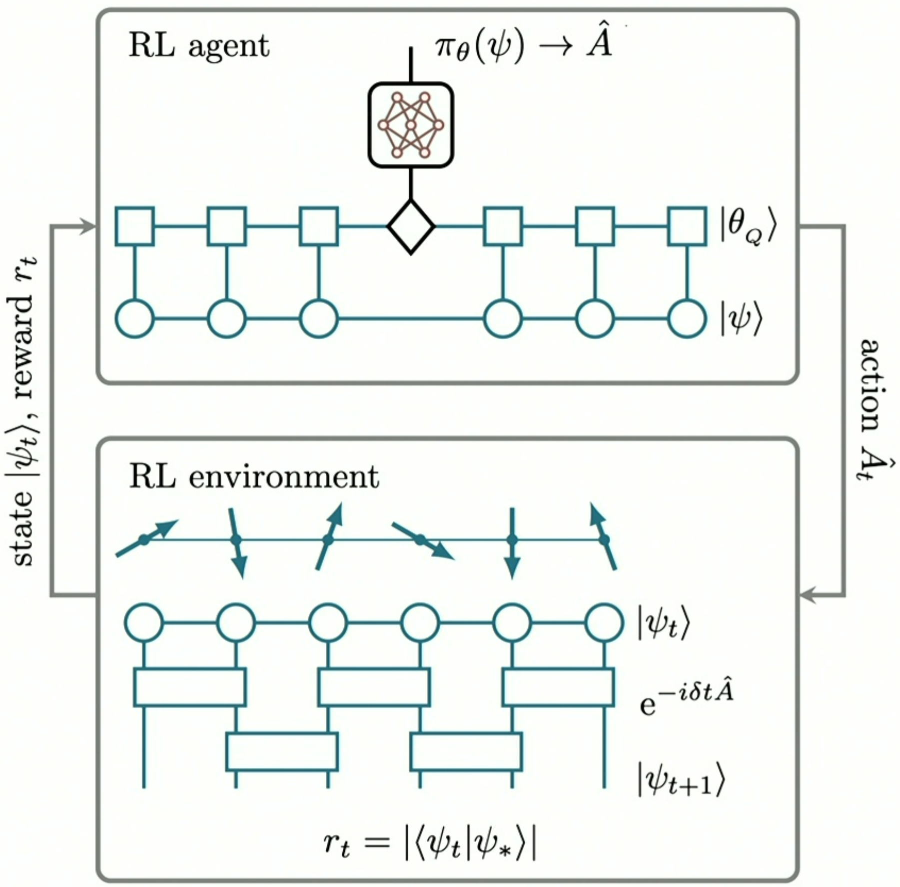

Training

RL FRAMEWORK

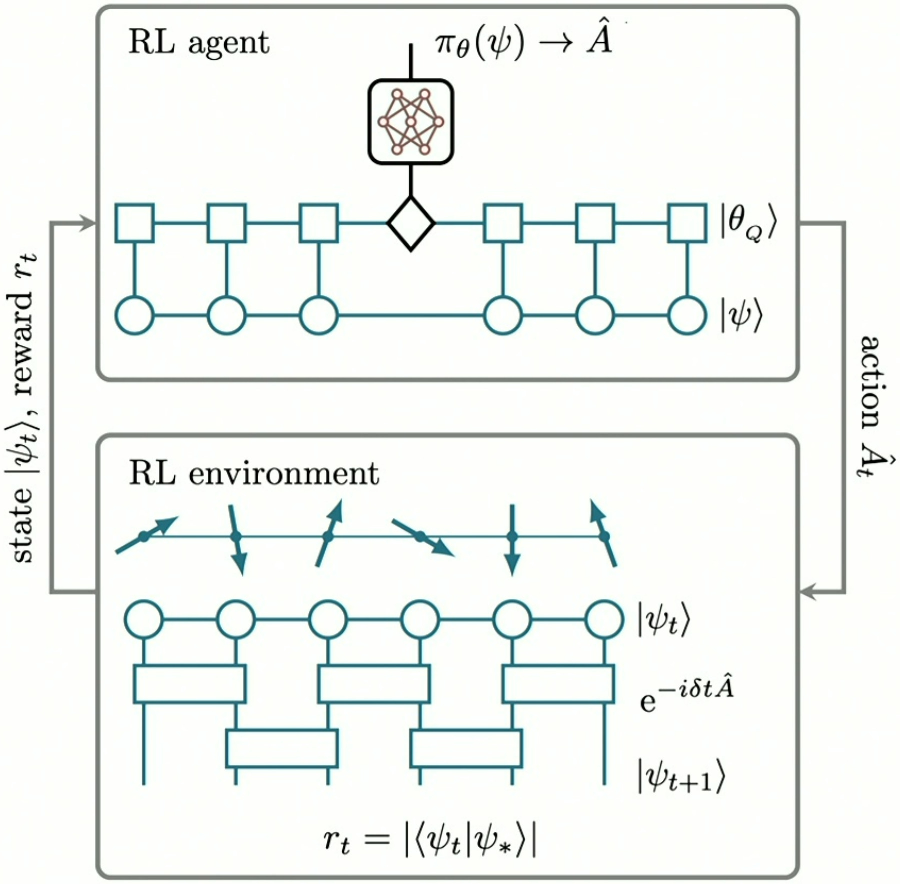

Training

Start each episode from a (random) initial state

Terminate episode when fidelity threshold is reached, e.g. $F = 0.99$

GROUND STATE PREPARATION

$$\hat{H}_{\text{Ising}} = J \sum_{i=1}^{N-1} \hat{Z}_i \hat{Z}_{i+1} - h_x \sum_{i=1}^{N} \hat{X}_i - h_z \sum_{i=1}^{N} \hat{Z}_i$$

Initial ground state $|\psi_{\text{initial}}\rangle$

$J = -1, h_x = 1.2, h_z = 0.2$

Target ground state $|\psi^*\rangle$

$J = +1, h_x = 0.5, h_z = 1.5$

The agent can prepare the target state from (nearly) any Ising ground state

GROUND STATE PREPARATION

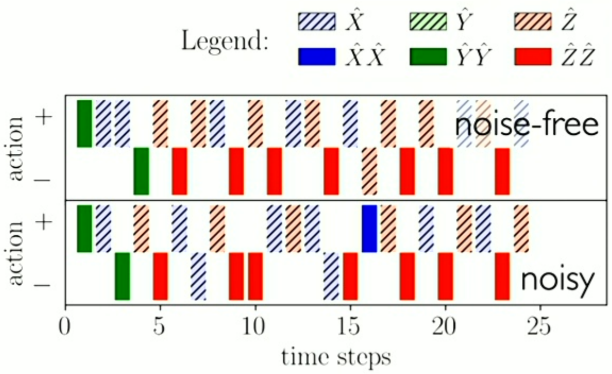

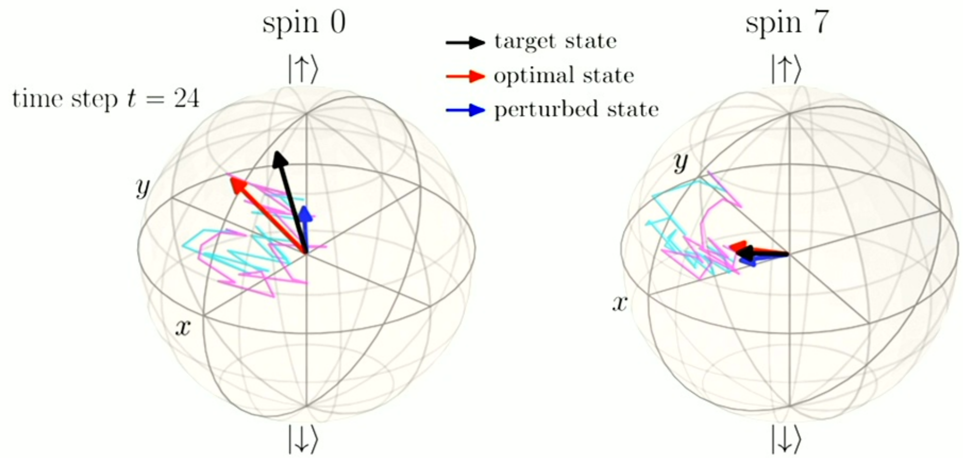

Make environment noisy

$$A \to A + \epsilon$$

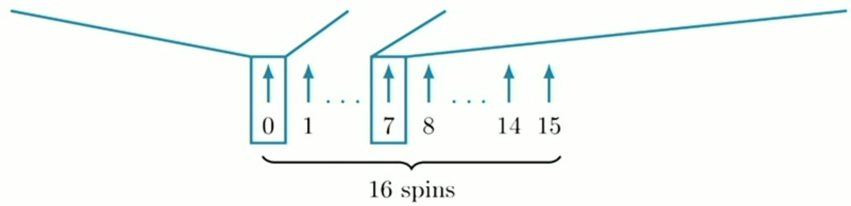

Fidelity threshold $F = 0.97$

SUMMARY

- **Matrix product states** (MPS): Compressed representations of exponentially large quantum many-body states

- **Reinforcement learning** (RL): An agent dynamically learns optimal controls in a trial-and-error fashion to maximize long term rewards

- RL + MPS → **Efficient quantum many-body control** *Nature Machine Intelligence* **5**, 780-791 (2023)

Marin Bukov

THANK YOU

EPFL OIST mpipks SOFIA UNIVERSITY ST. KLIMENT OHRIDSKI