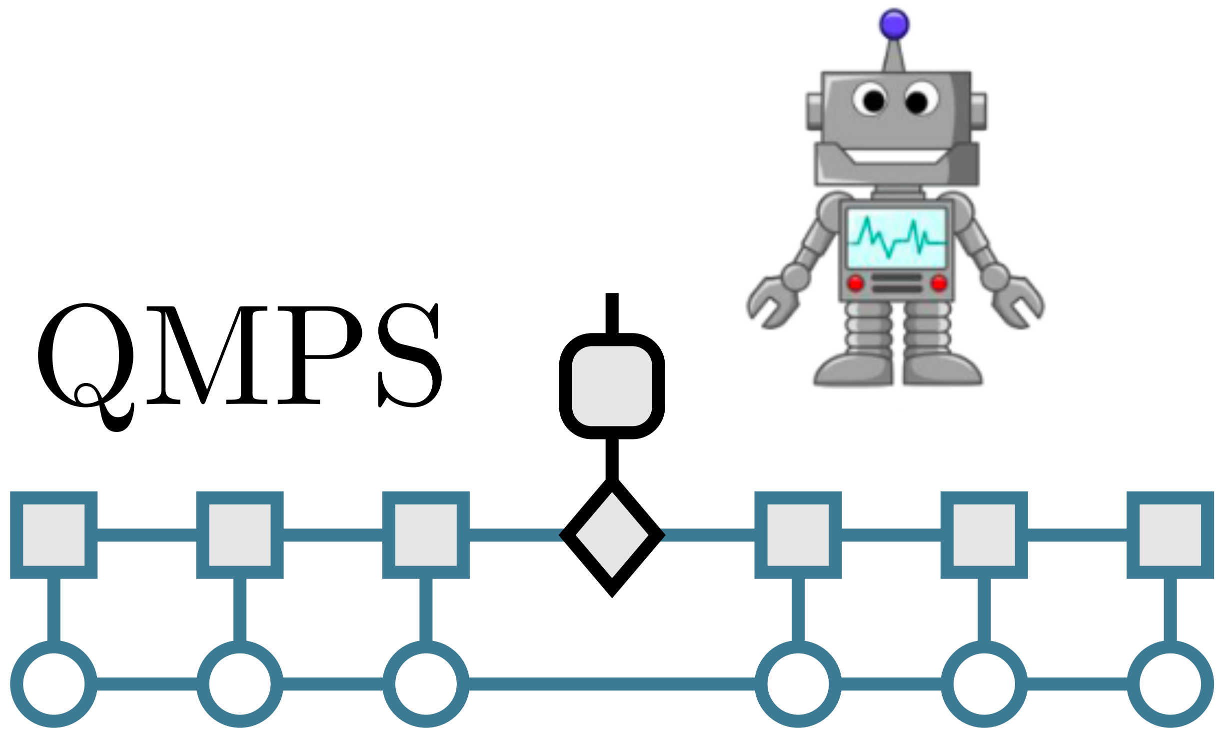

SELF-CORRECTING QUANTUM MANY-BODY CONTROL USING REINFORCEMENT LEARNING WITH TENSOR NETWORKS

Friederike Metz and Marin Bukov

arXiv:2201.11790

OIST

mpipks

SOFIA UNIVERSITY
ST. KLIMENT OHRIDSKI

QUANTUM MANY-BODY CONTROL

Essential for most quantum technologies (computing, simulation, metrology)

e.g. state preparation

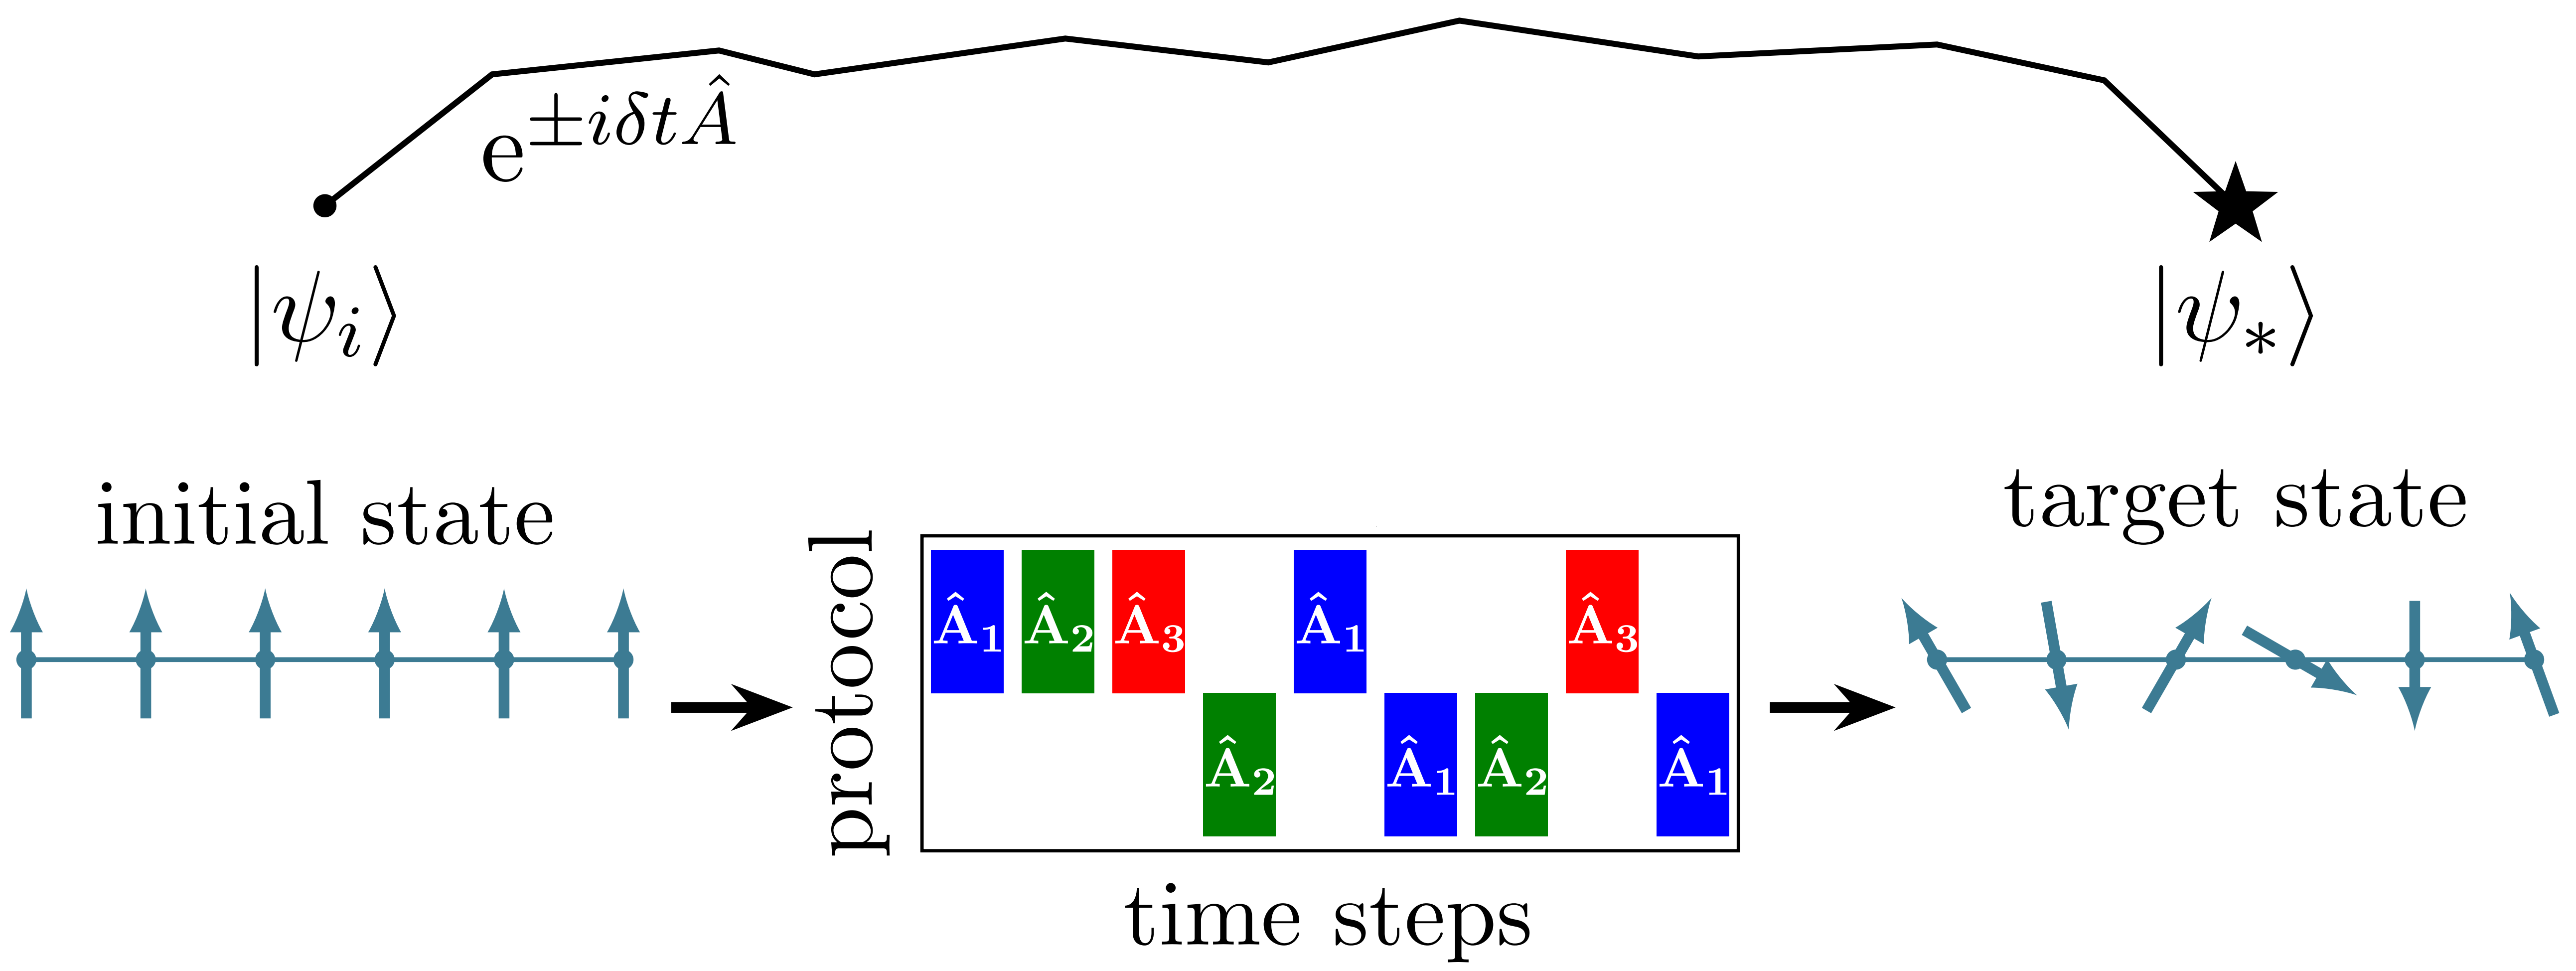

QUANTUM MANY-BODY CONTROL

Essential for most quantum technologies (computing, simulation, metrology)

e.g. state preparation

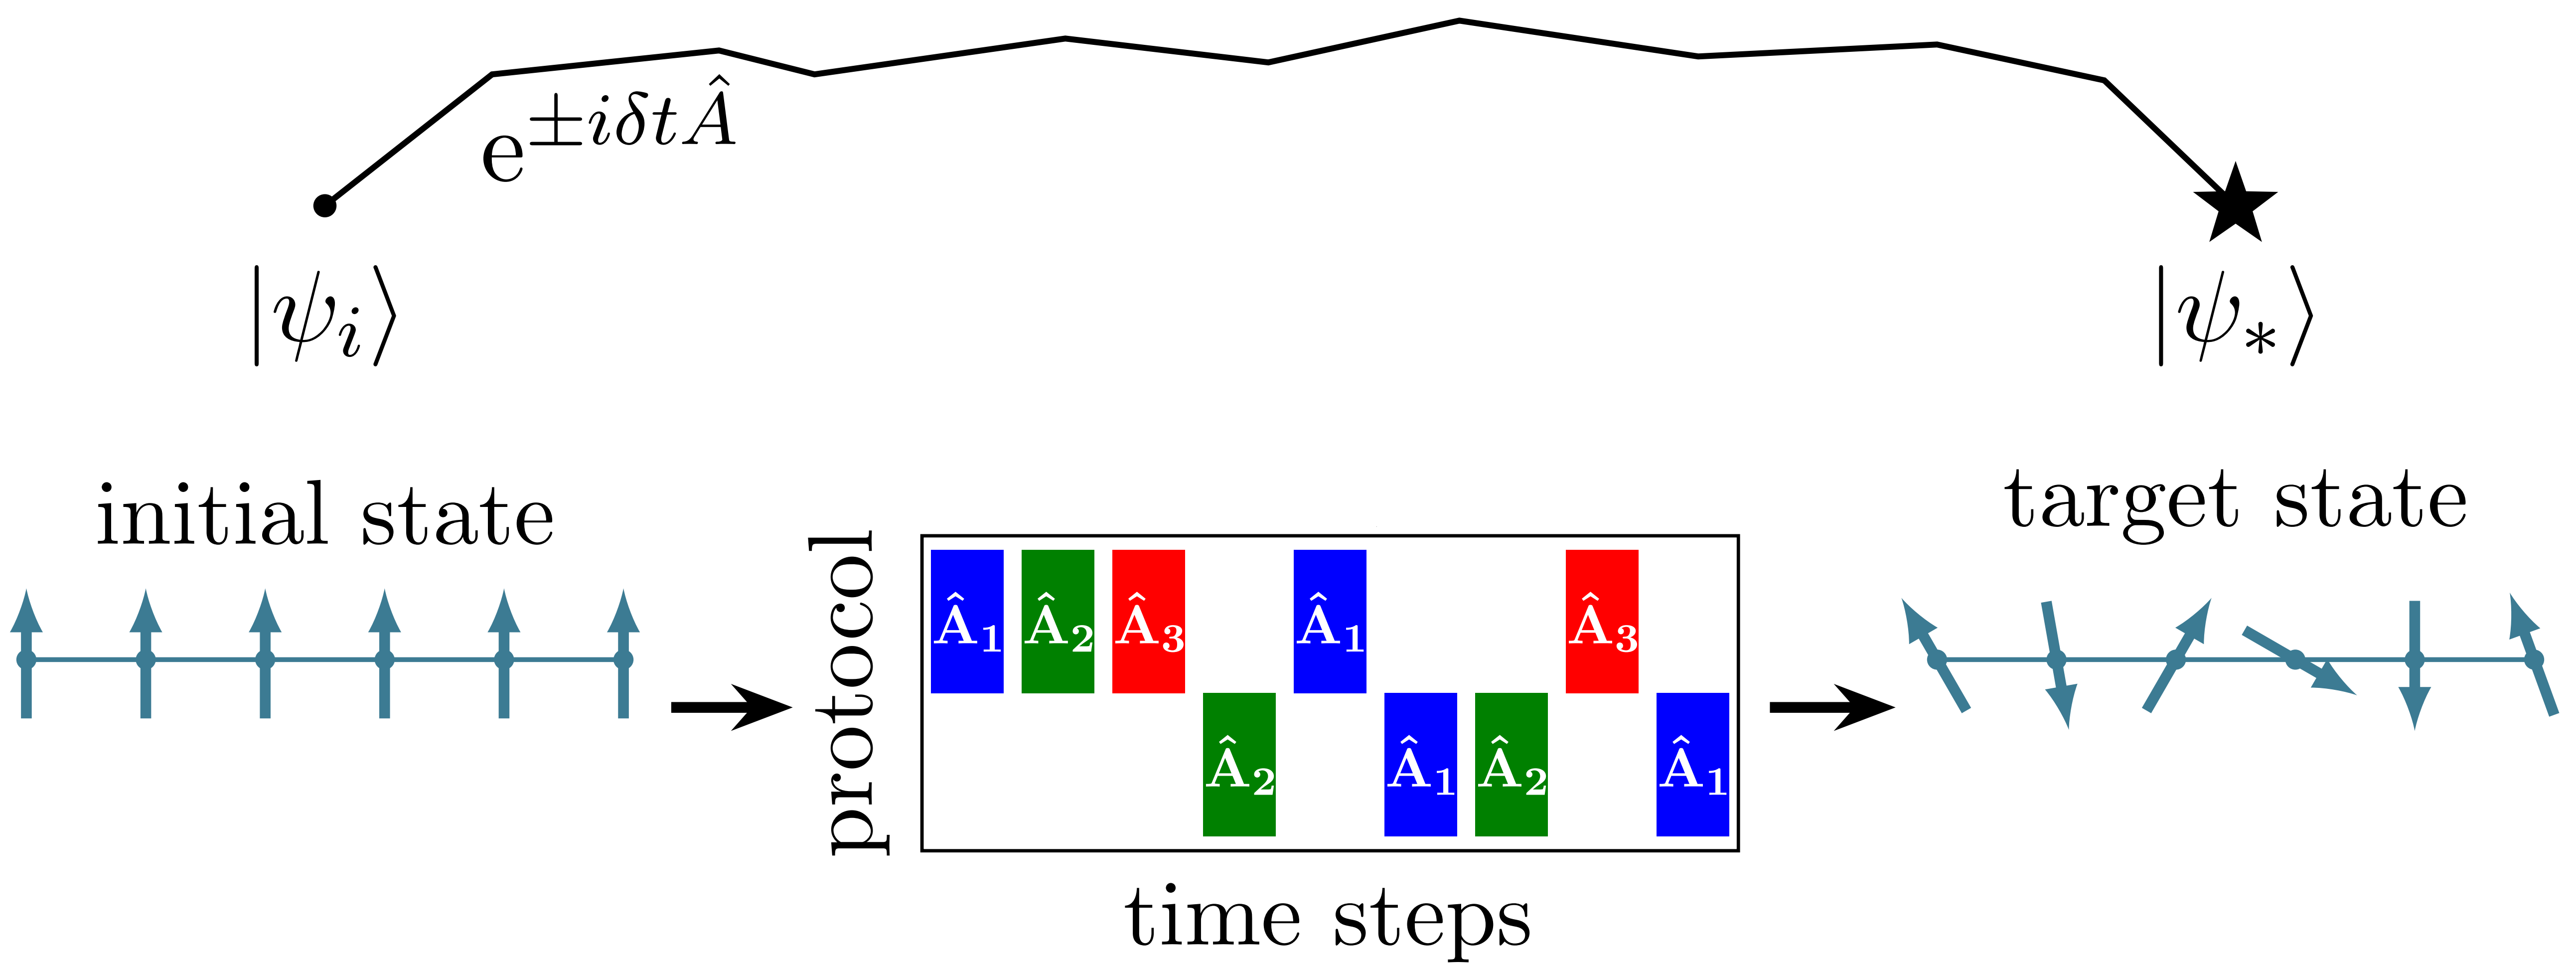

Problem: Hilbert space dimension grows exponentially with system size

QUANTUM MANY-BODY CONTROL

Essential for most quantum technologies (computing, simulation, metrology)

e.g. state preparation

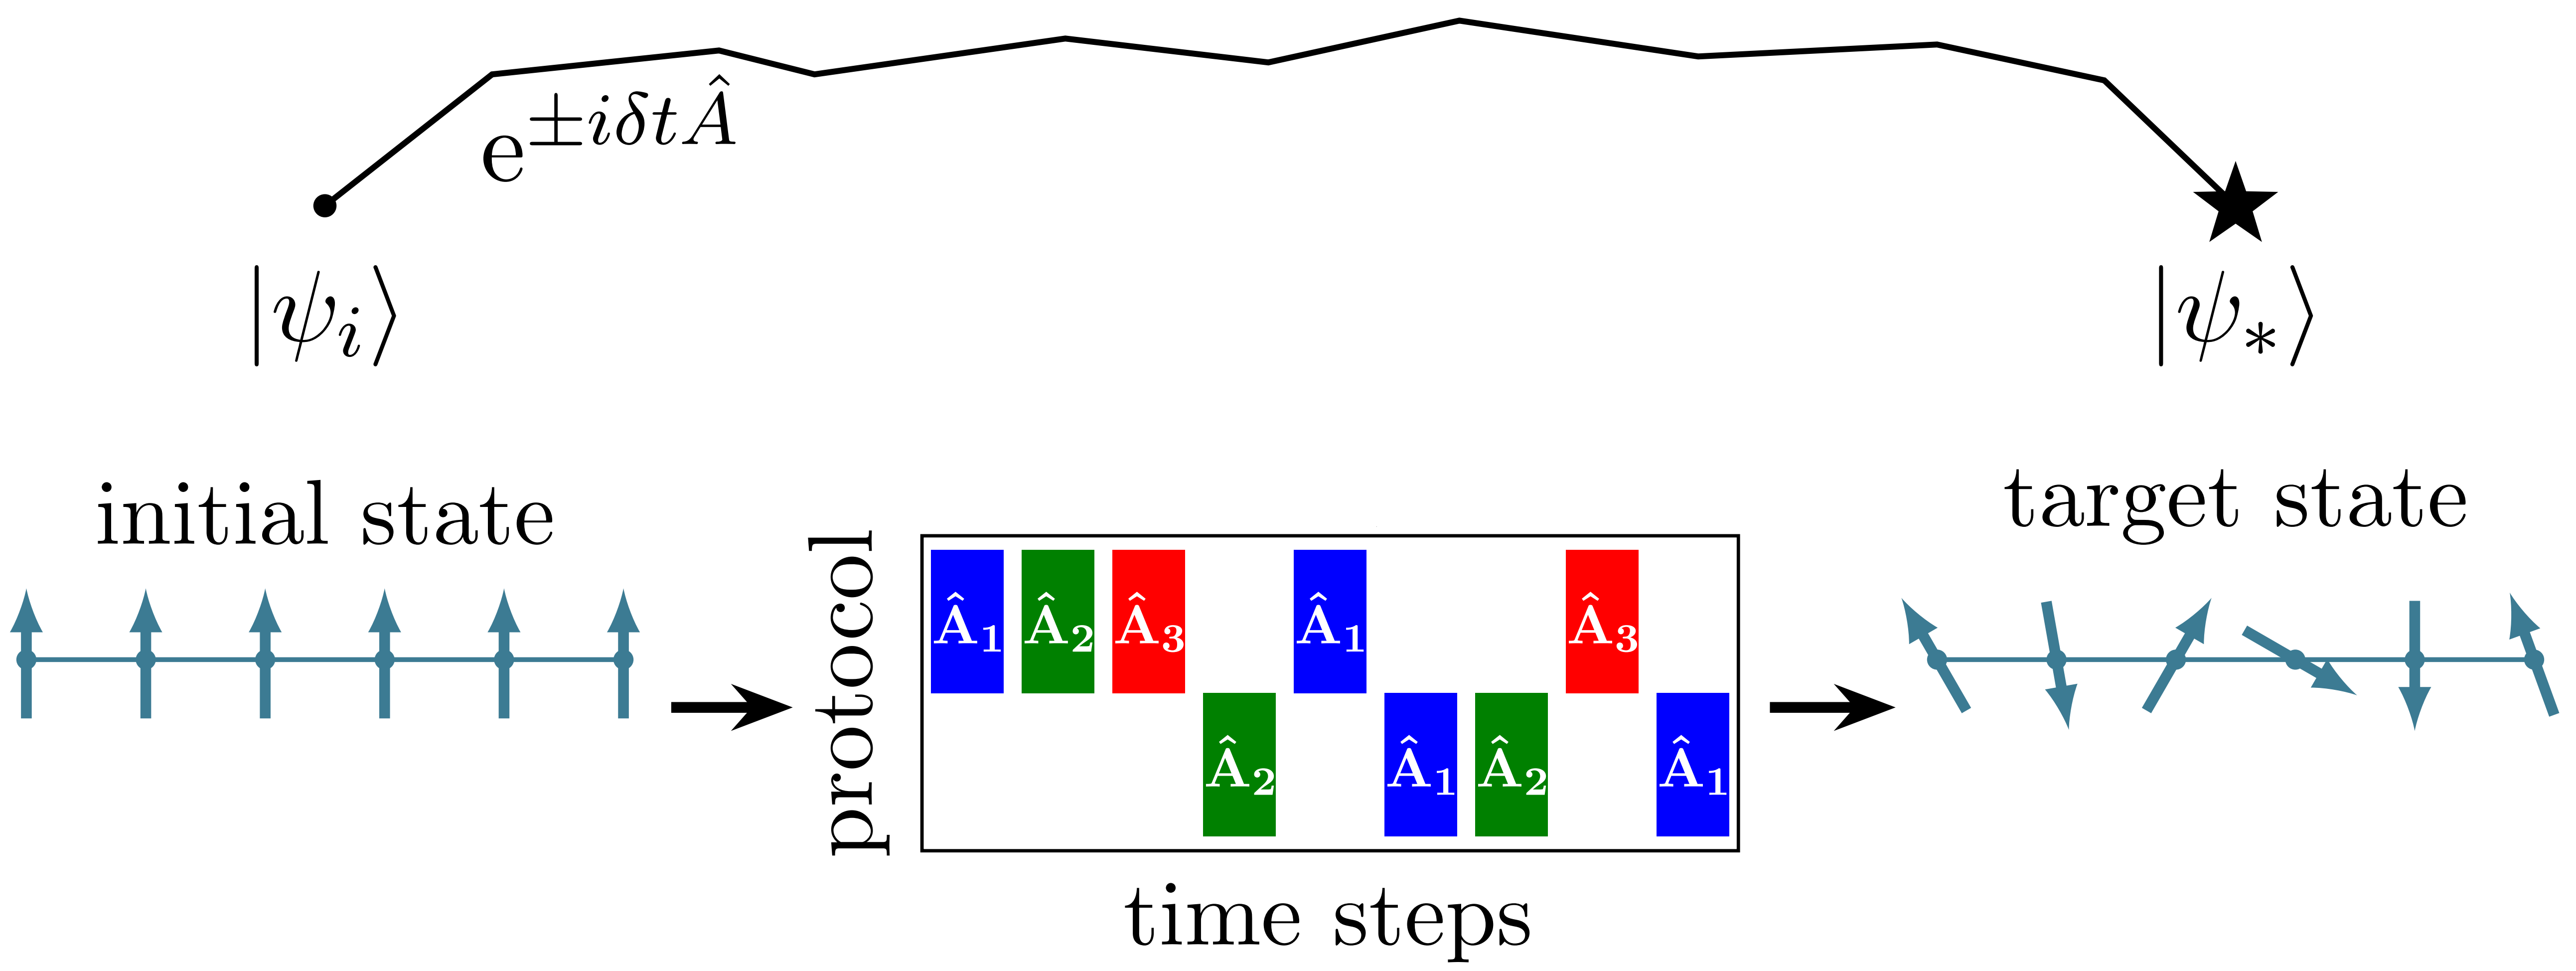

Problem: Hilbert space dimension grows exponentially with system size

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

Compressing the quantum many-body state

Trainable machine learning architecture for the RL agent

QUANTUM MANY-BODY CONTROL

Essential for most quantum technologies (computing, simulation, metrology)

e.g. state preparation

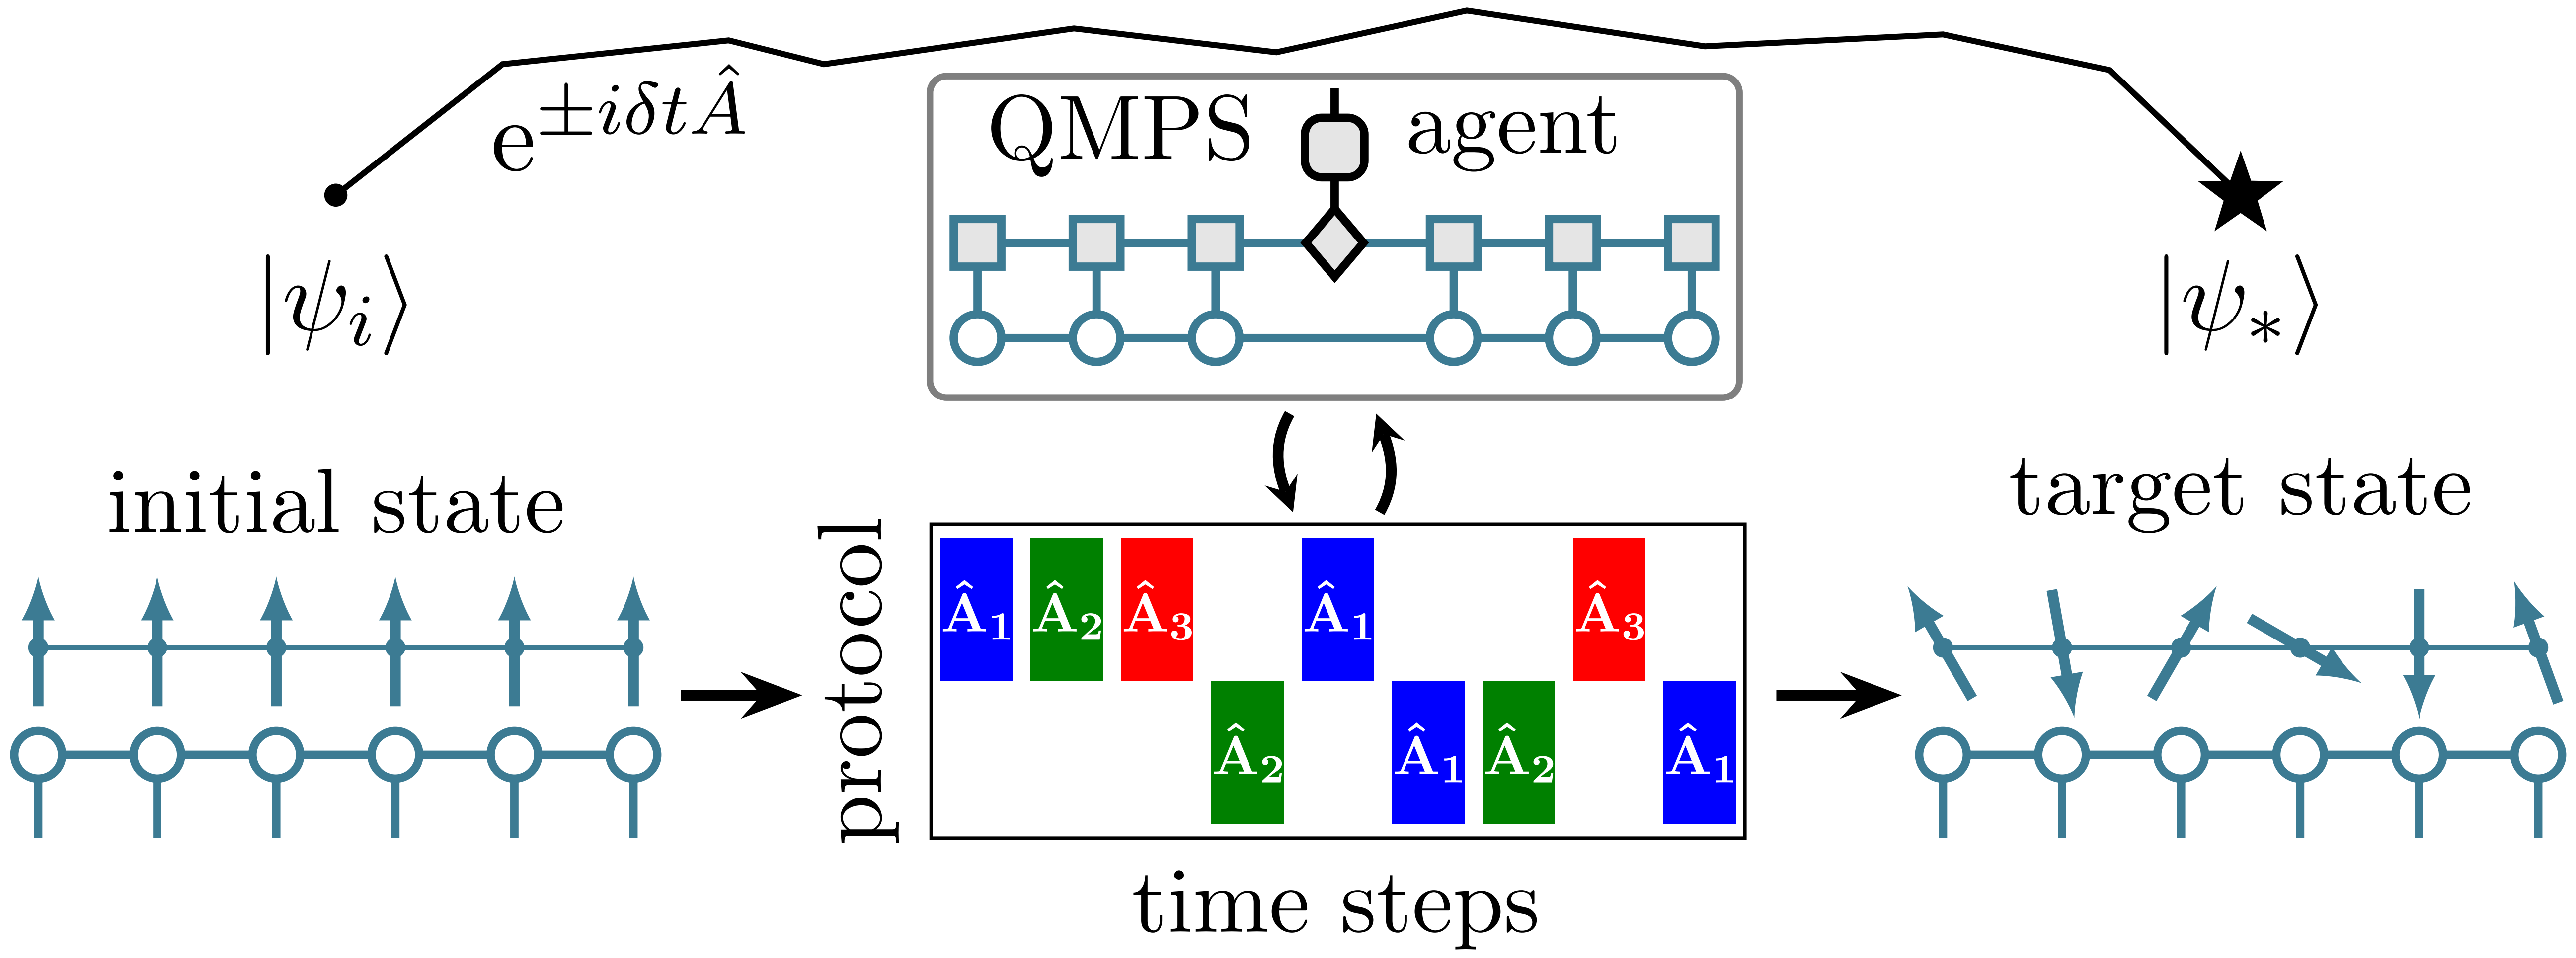

Problem: Hilbert space dimension grows exponentially with system size

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

Compressing the quantum many-body state

Trainable machine learning architecture for the RL agent

REINFORCEMENT LEARNING

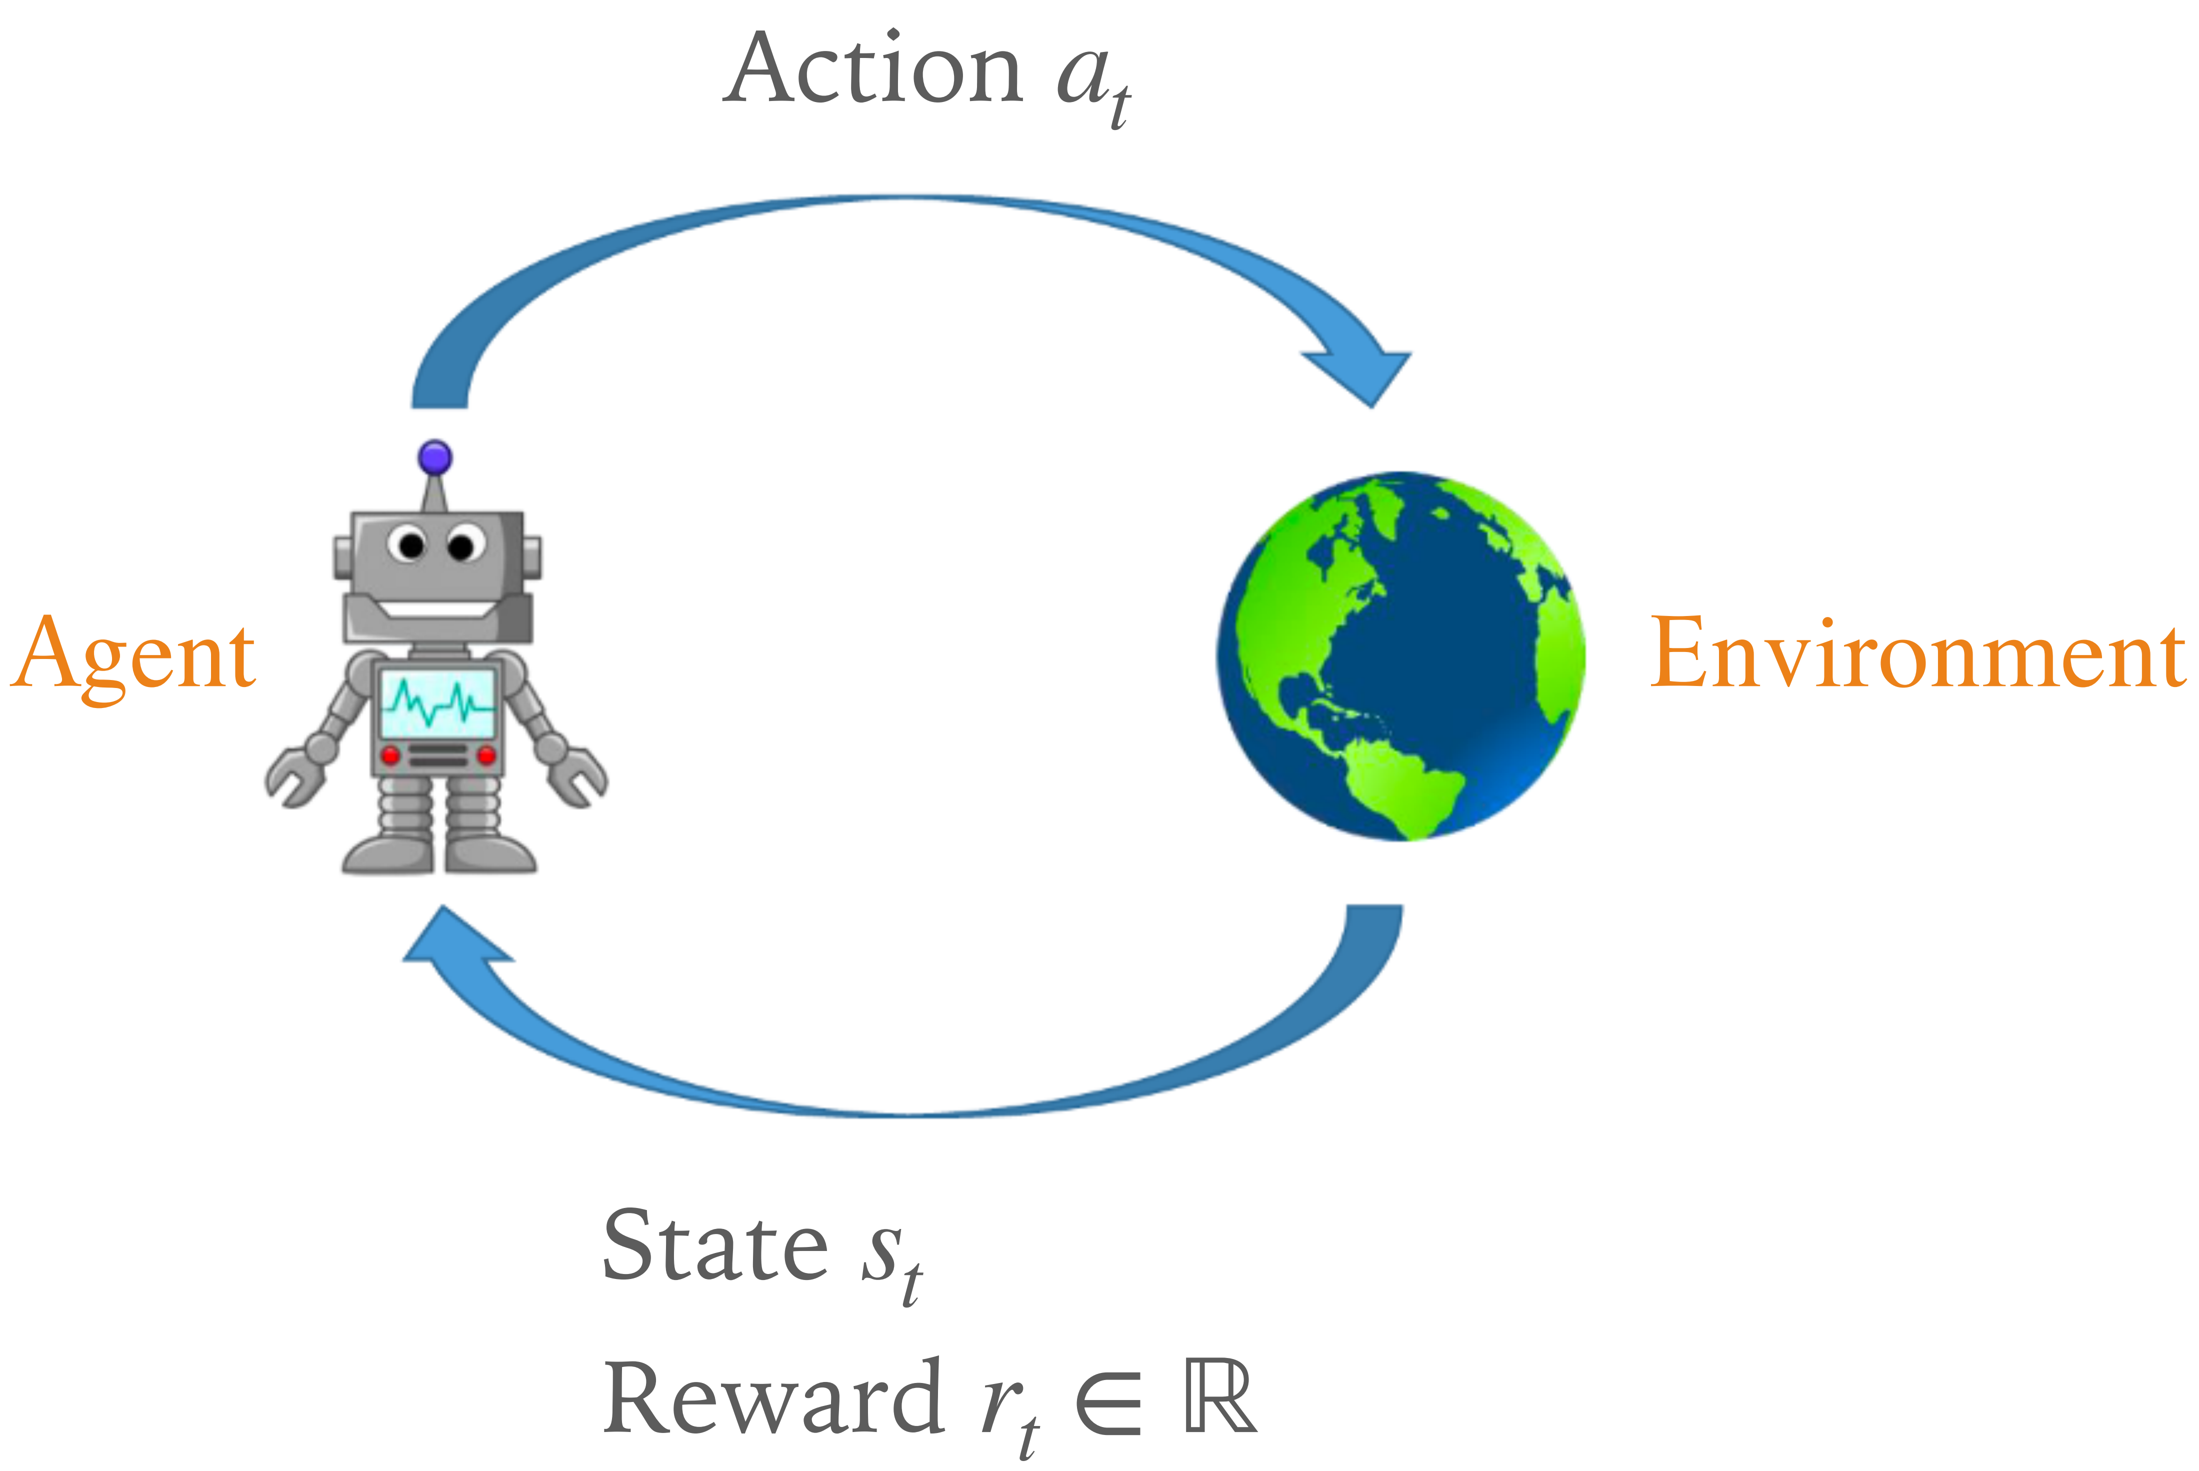

Q-learning: Learn optimal Q values $Q^*(s,a)$

Maximum expected sum of future rewards if you start in state s and take action a

Mnih et al., Nature 518, 529–533 (2015)
R. S. Sutton, A. G. Barto, Reinforcement learning: An introduction. MIT press

REINFORCEMENT LEARNING

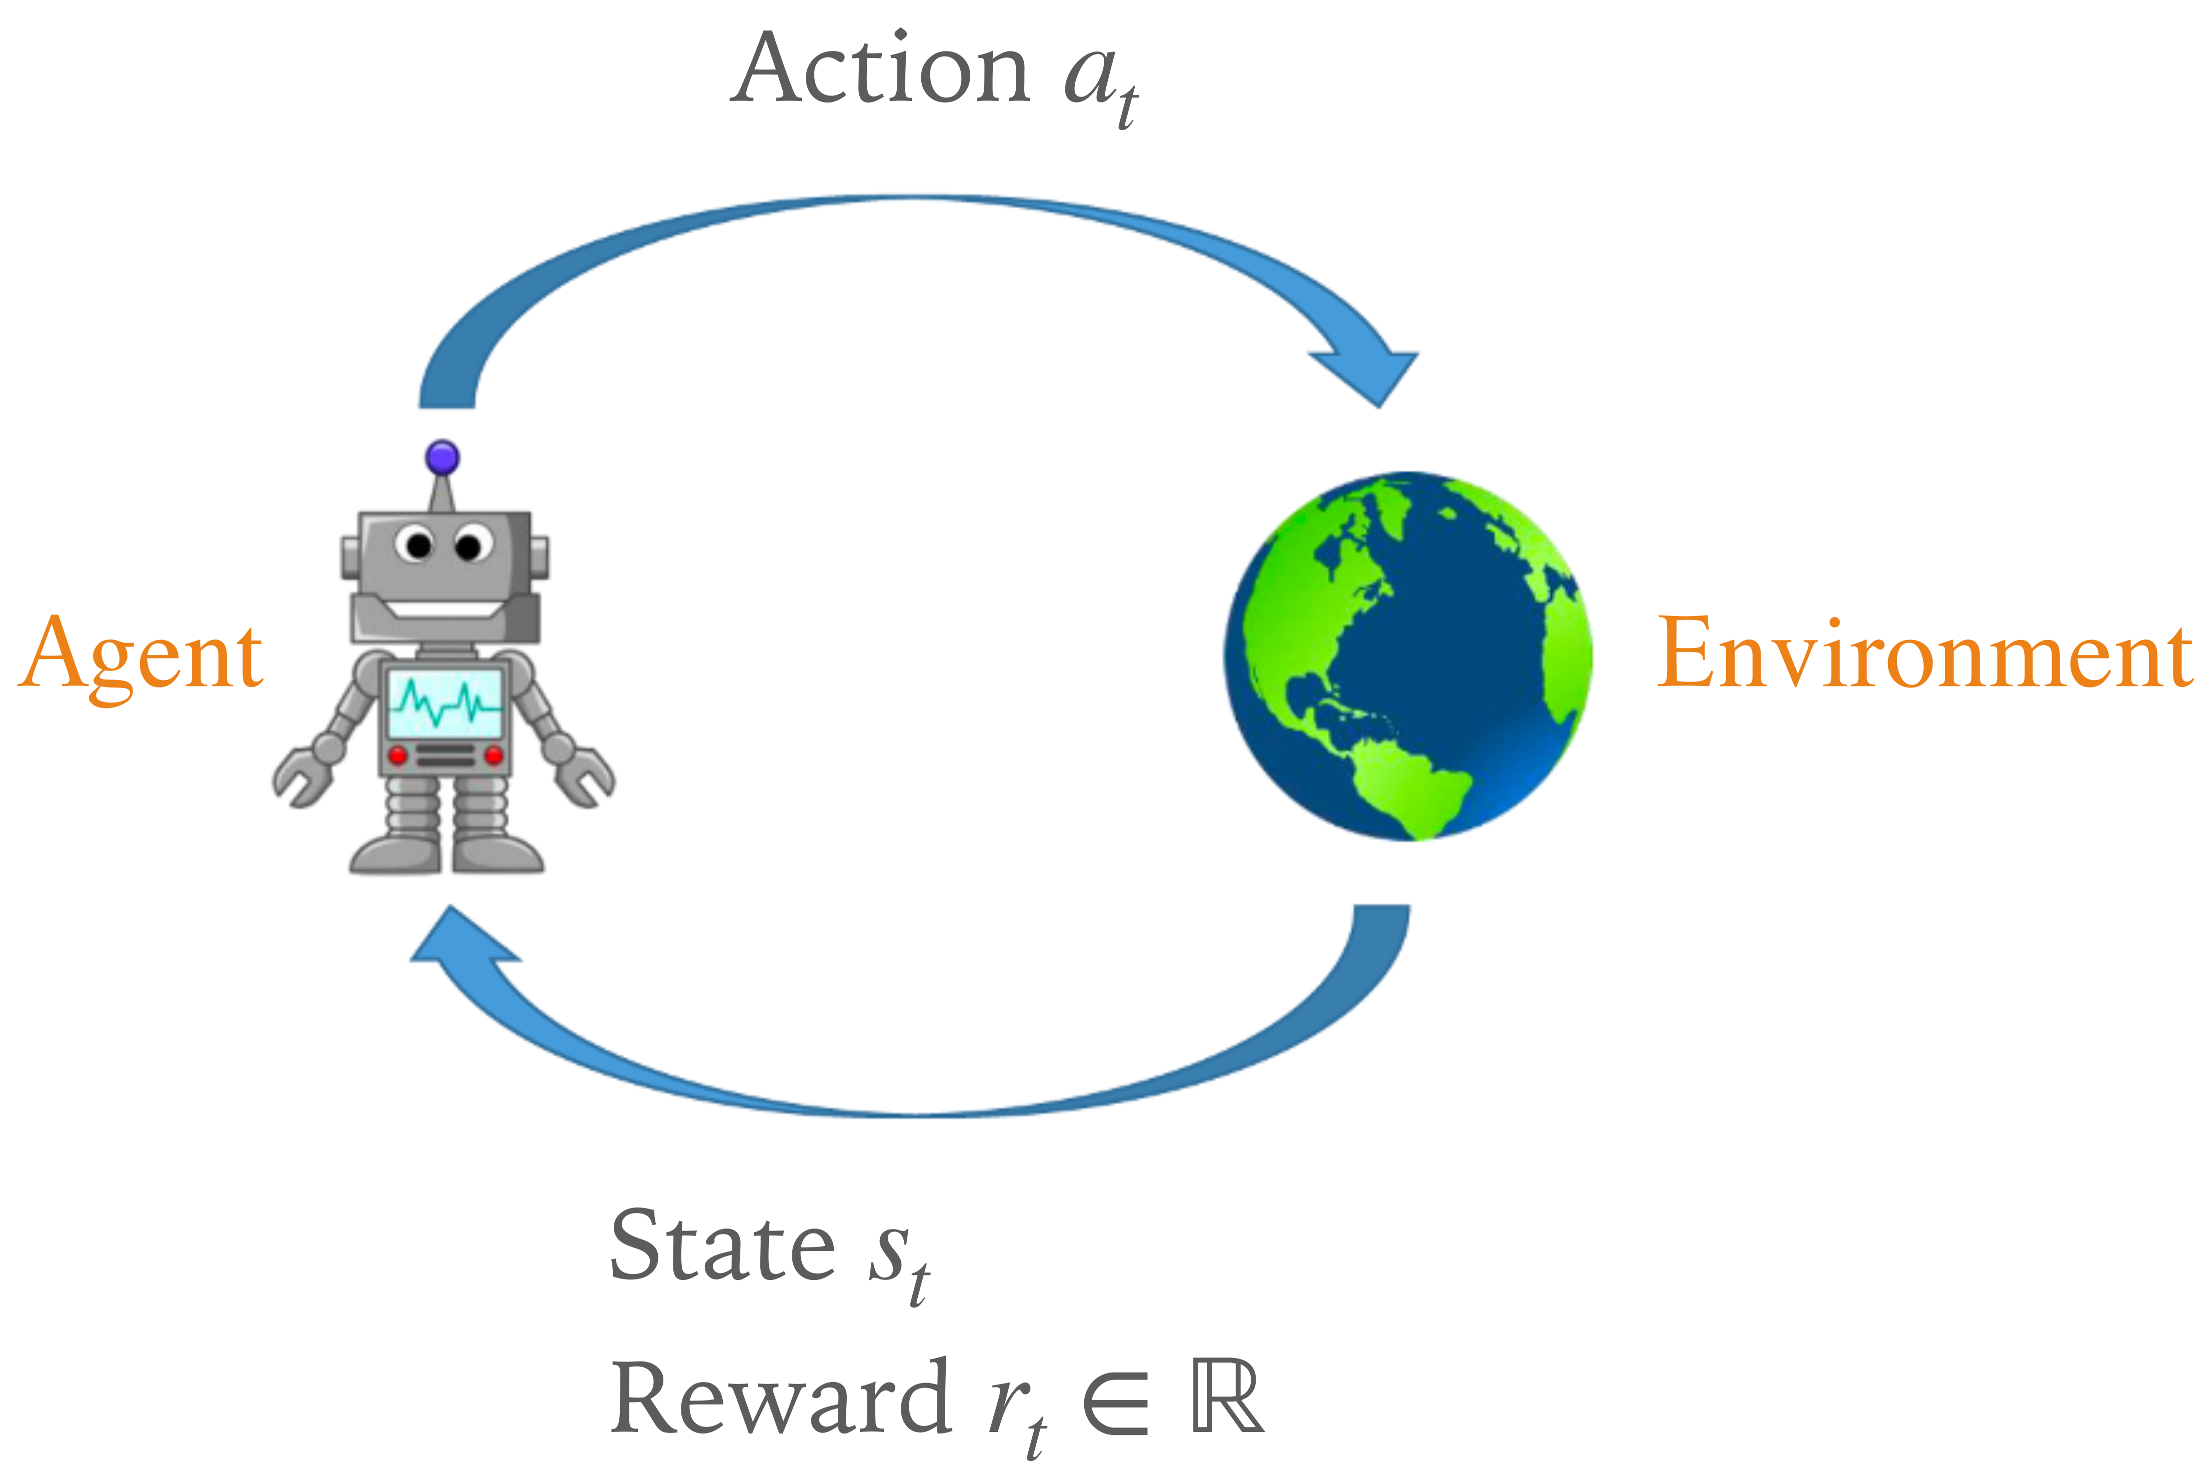

Q-learning: Learn optimal Q values $Q^*(s, a)$

Maximum expected sum of future rewards if you start in state s and take action a

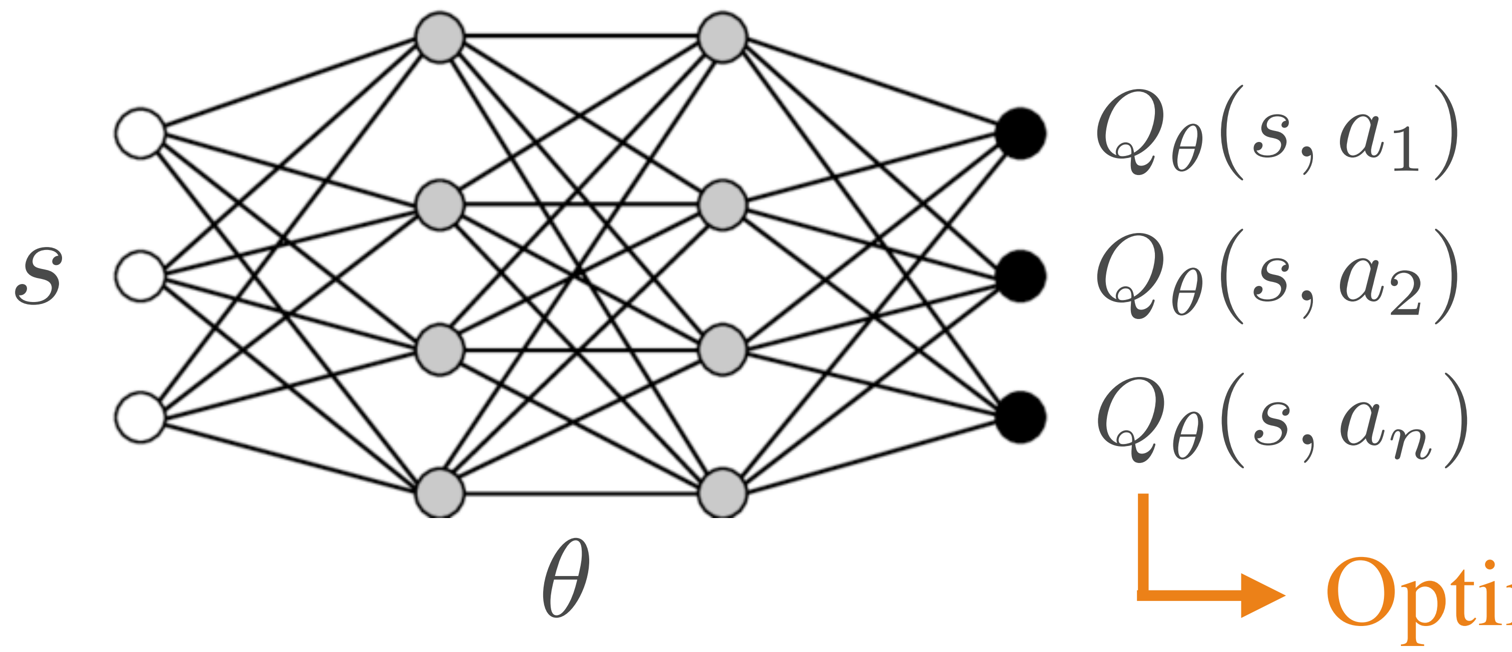

Optimal policy/action

$$\pi^*(s) = a$$

Mnih et al., Nature 518, 529–533 (2015)
R. S. Sutton, A. G. Barto, Reinforcement learning: An introduction. MIT press

MATRIX PRODUCT STATE ANSATZ FOR Q-LEARNING (QMPS)

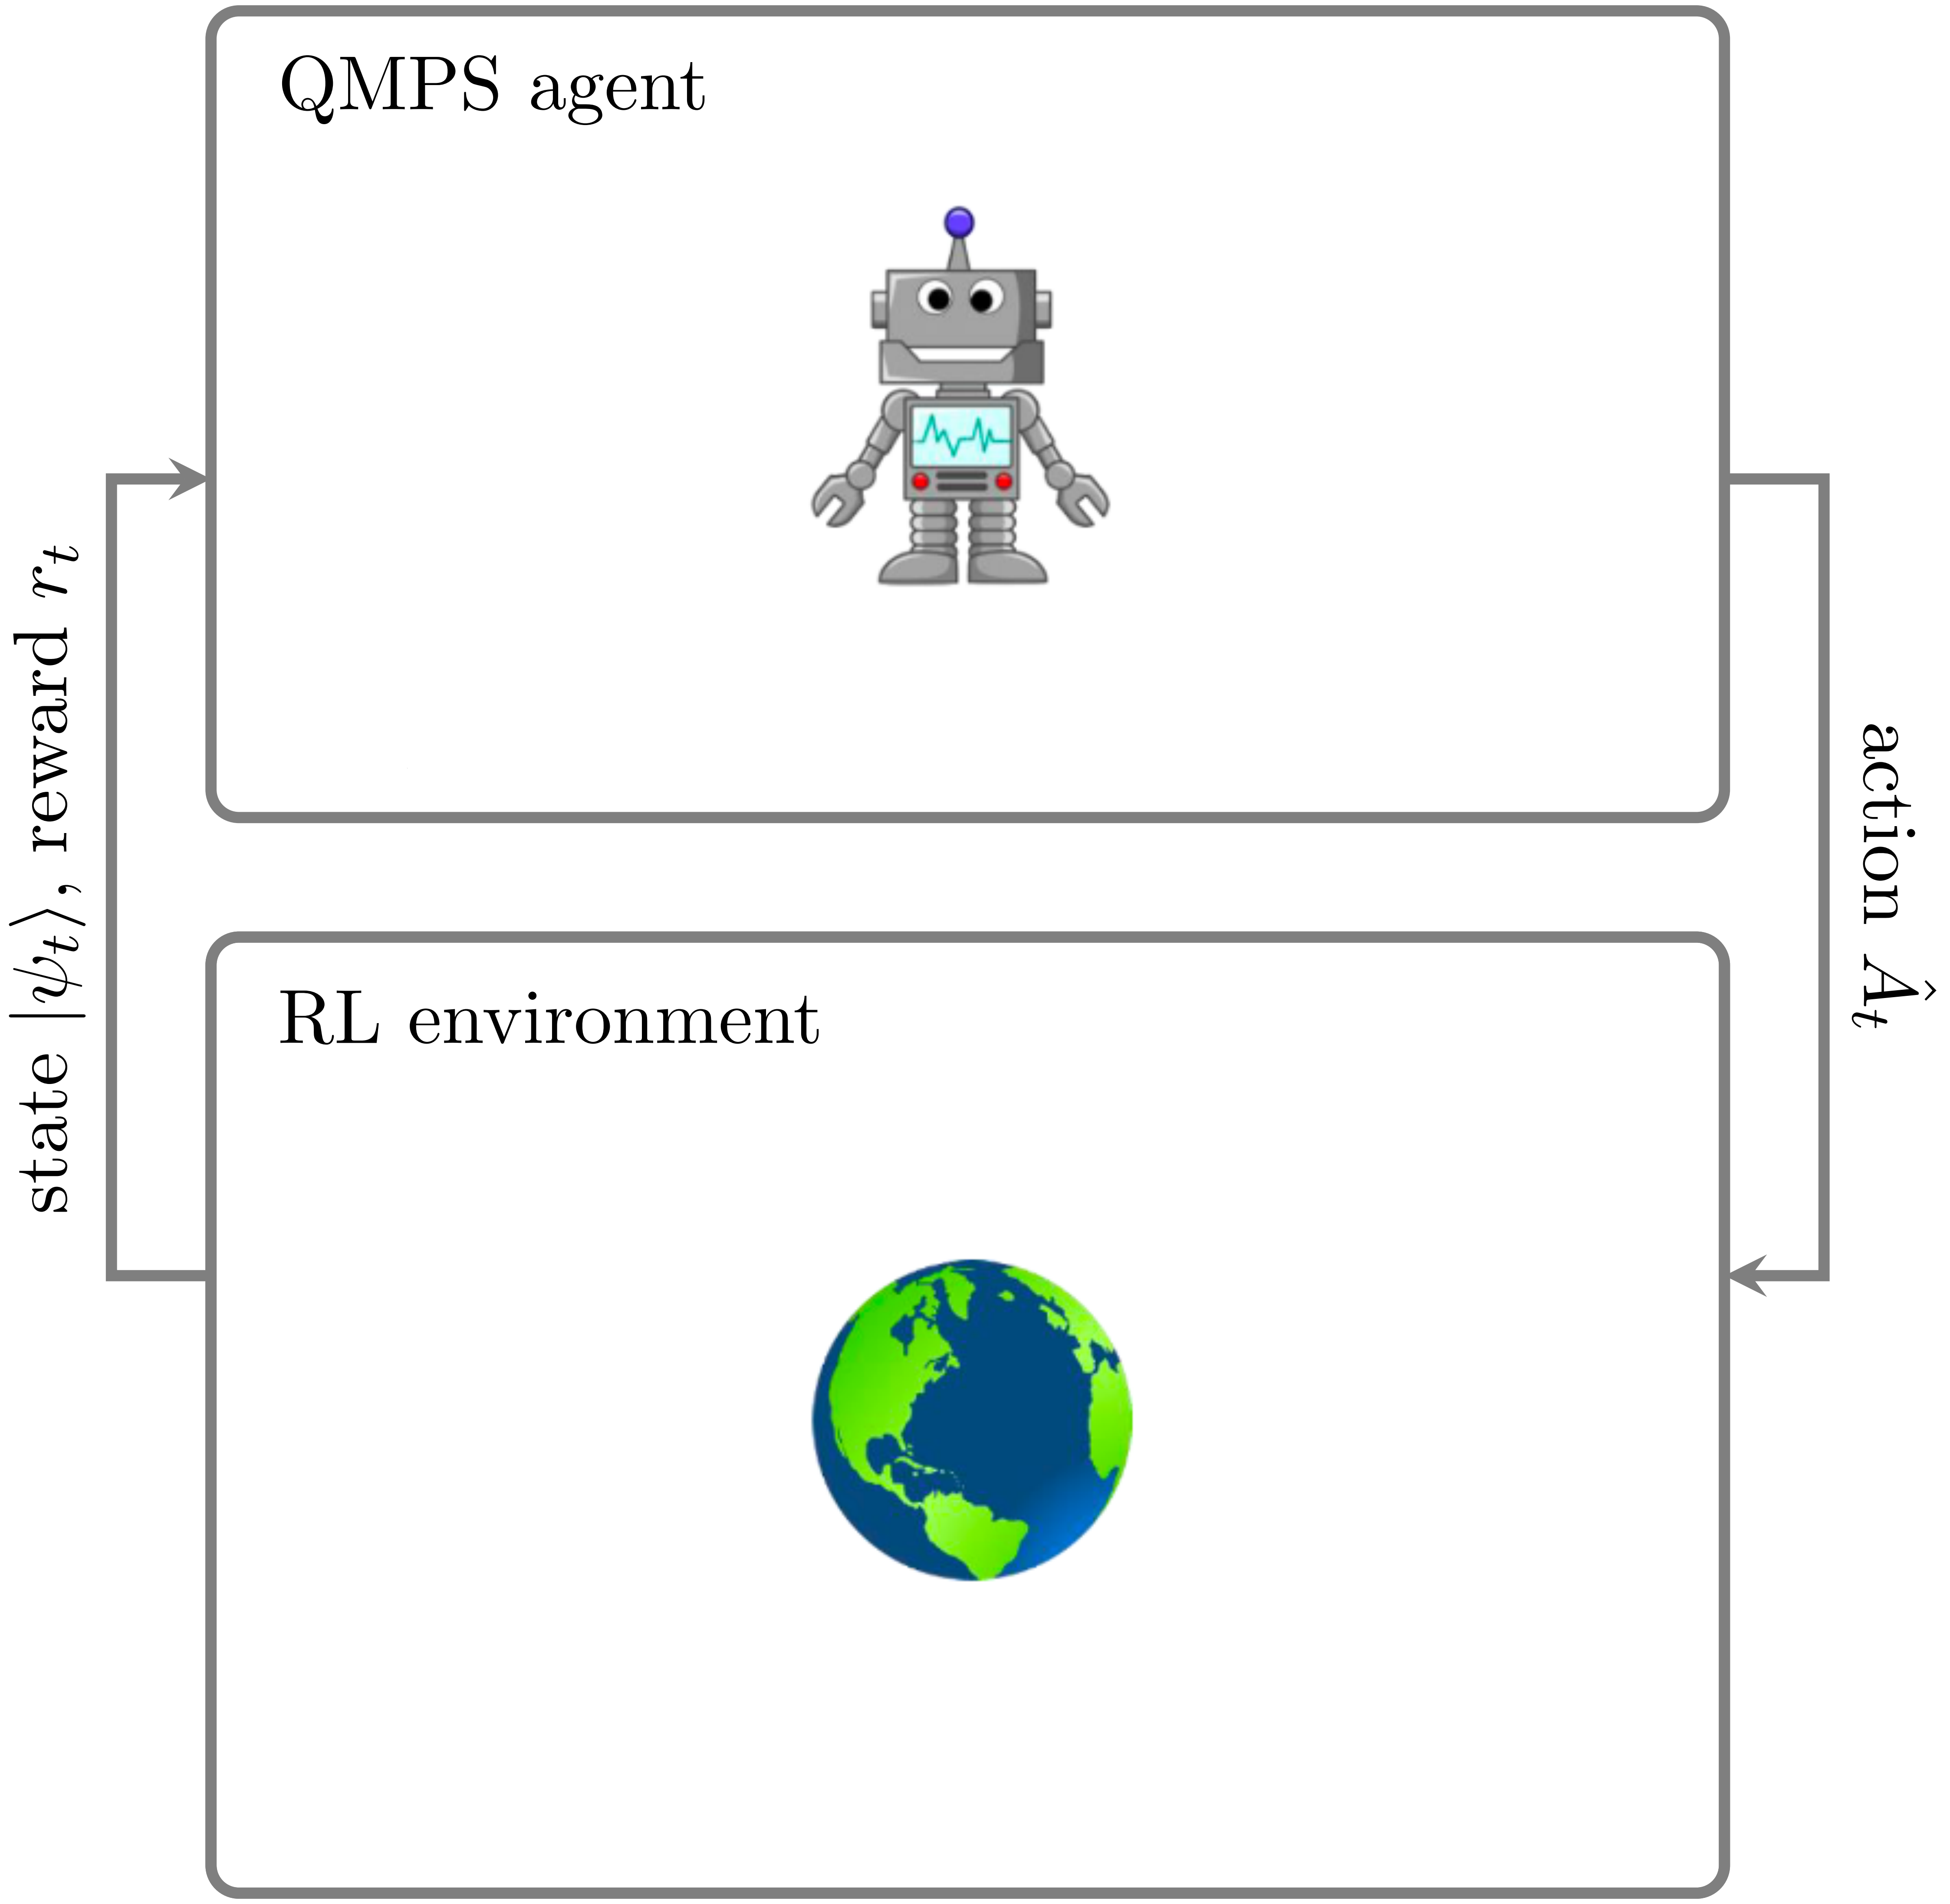

MATRIX PRODUCT STATE ANSATZ FOR Q-LEARNING (QMPS)

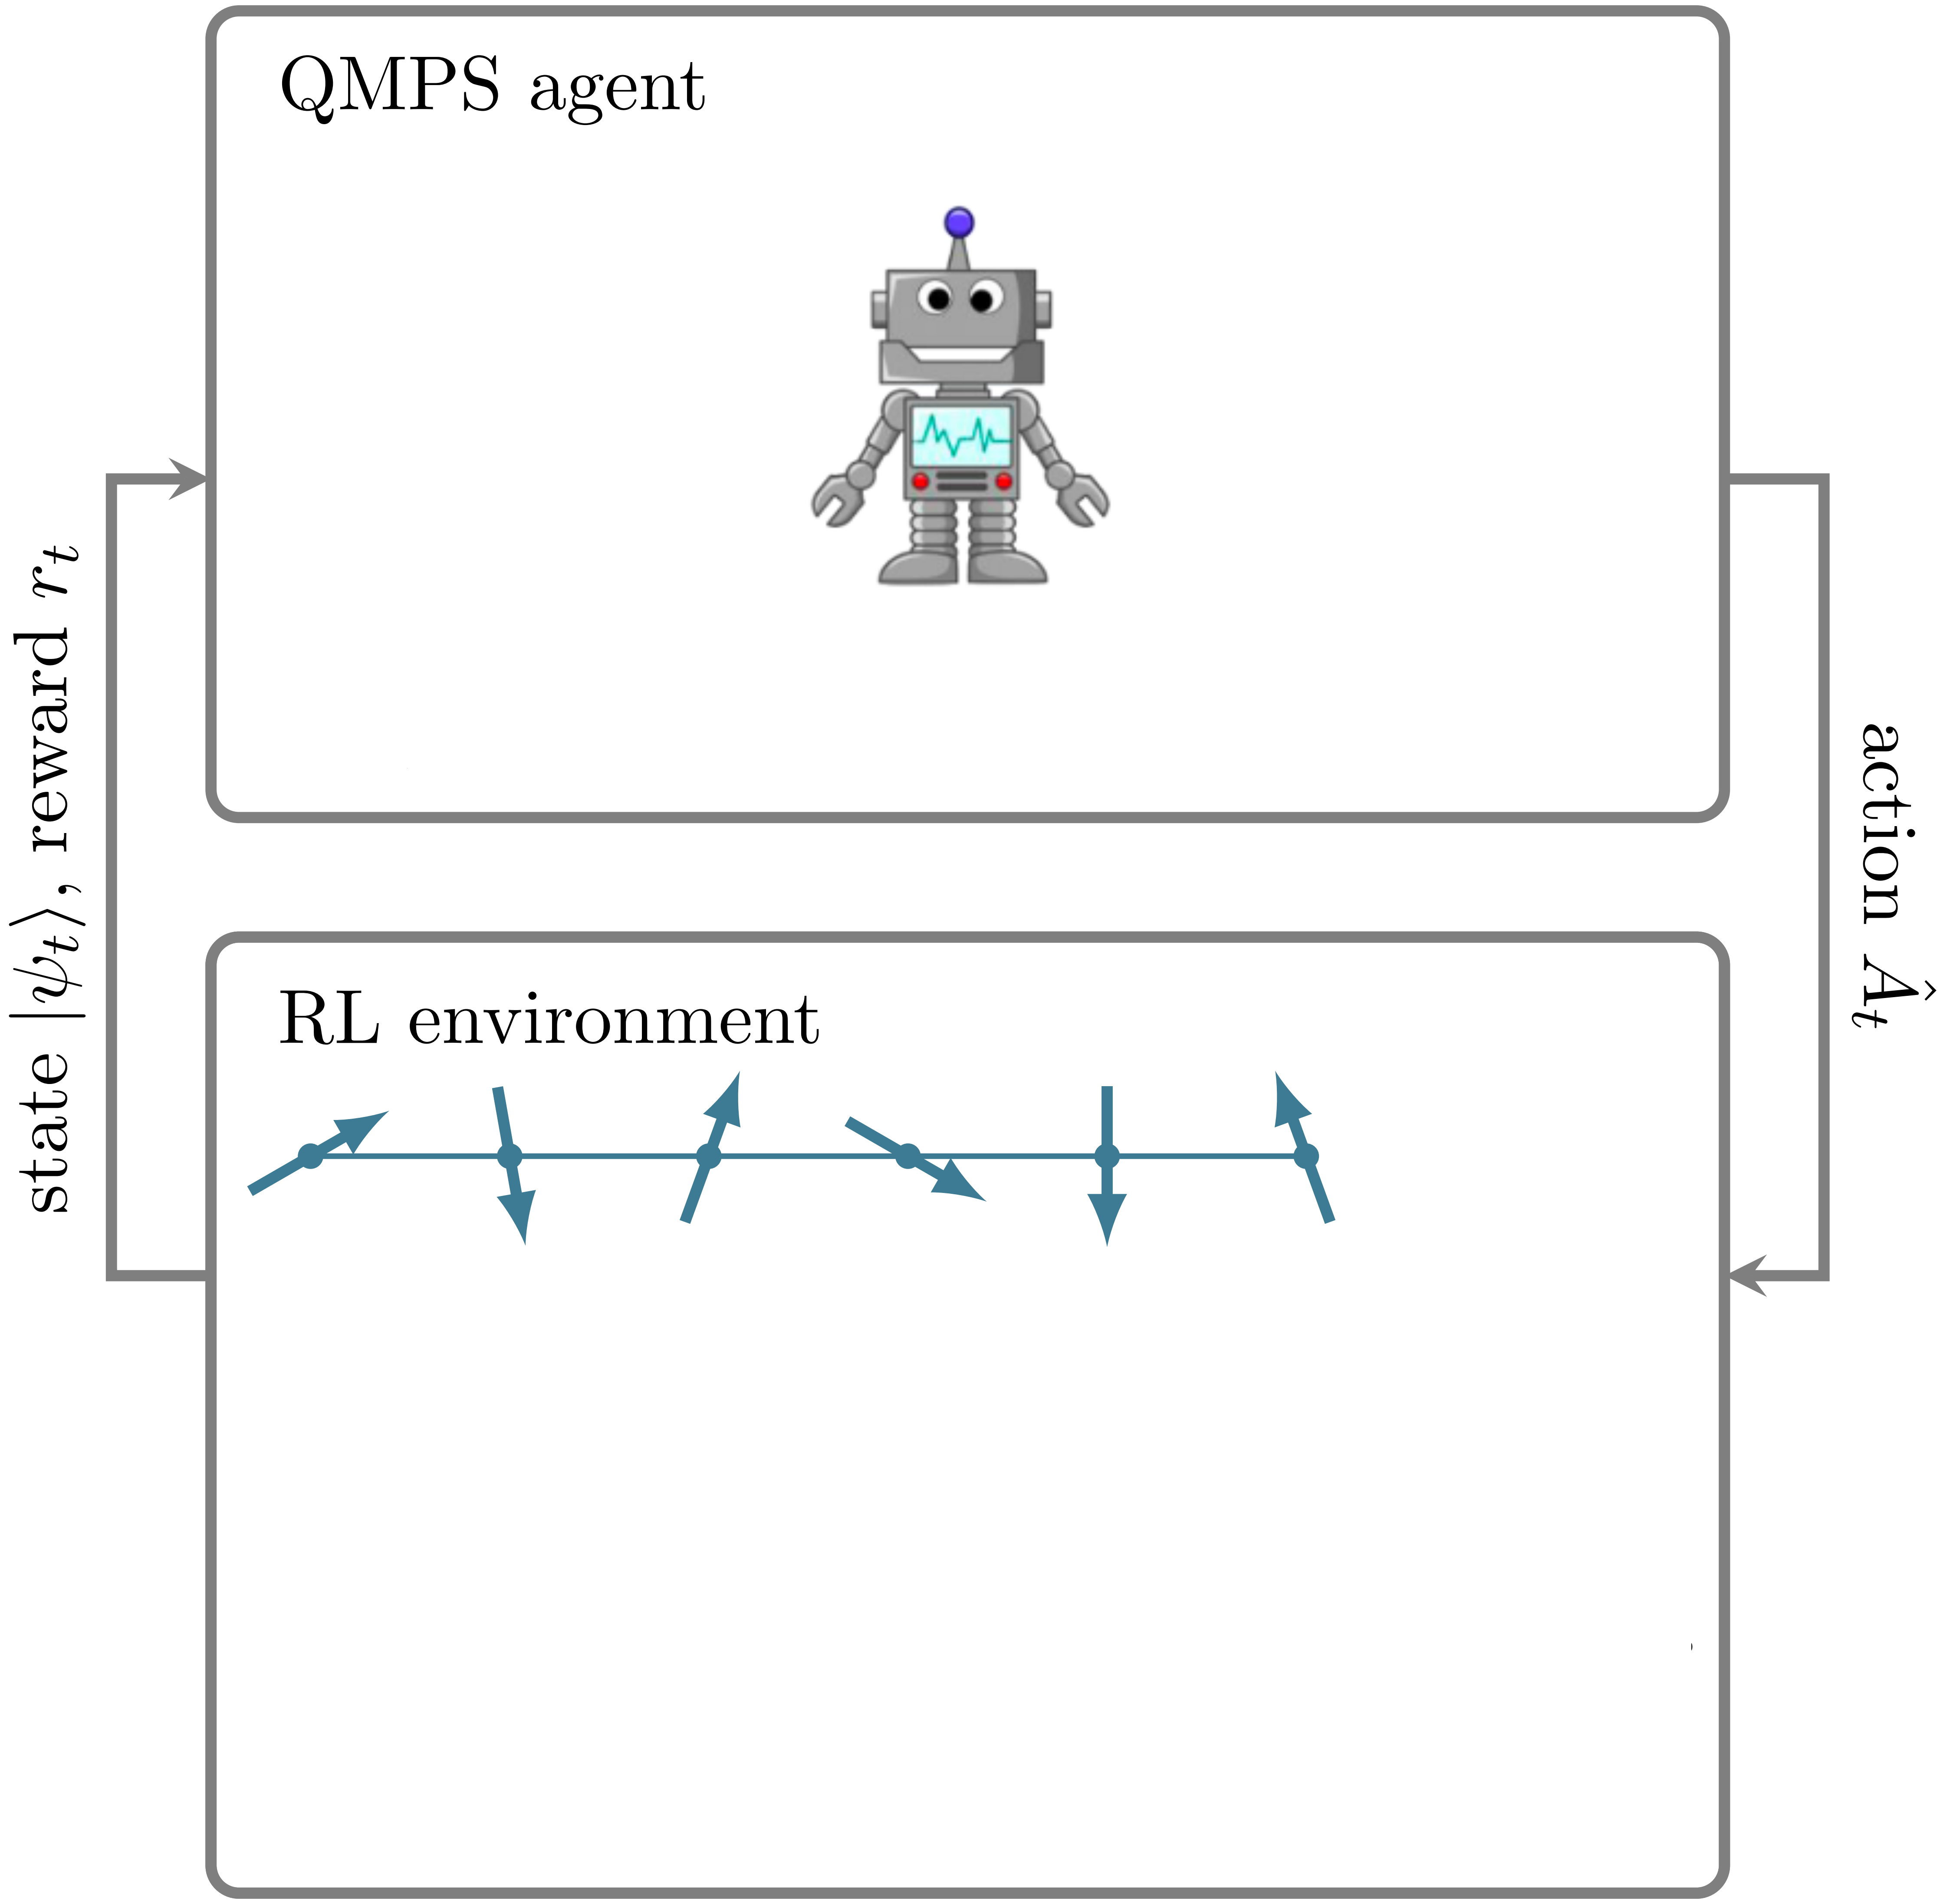

RL environment

MATRIX PRODUCT STATE ANSATZ FOR Q-LEARNING (QMPS)

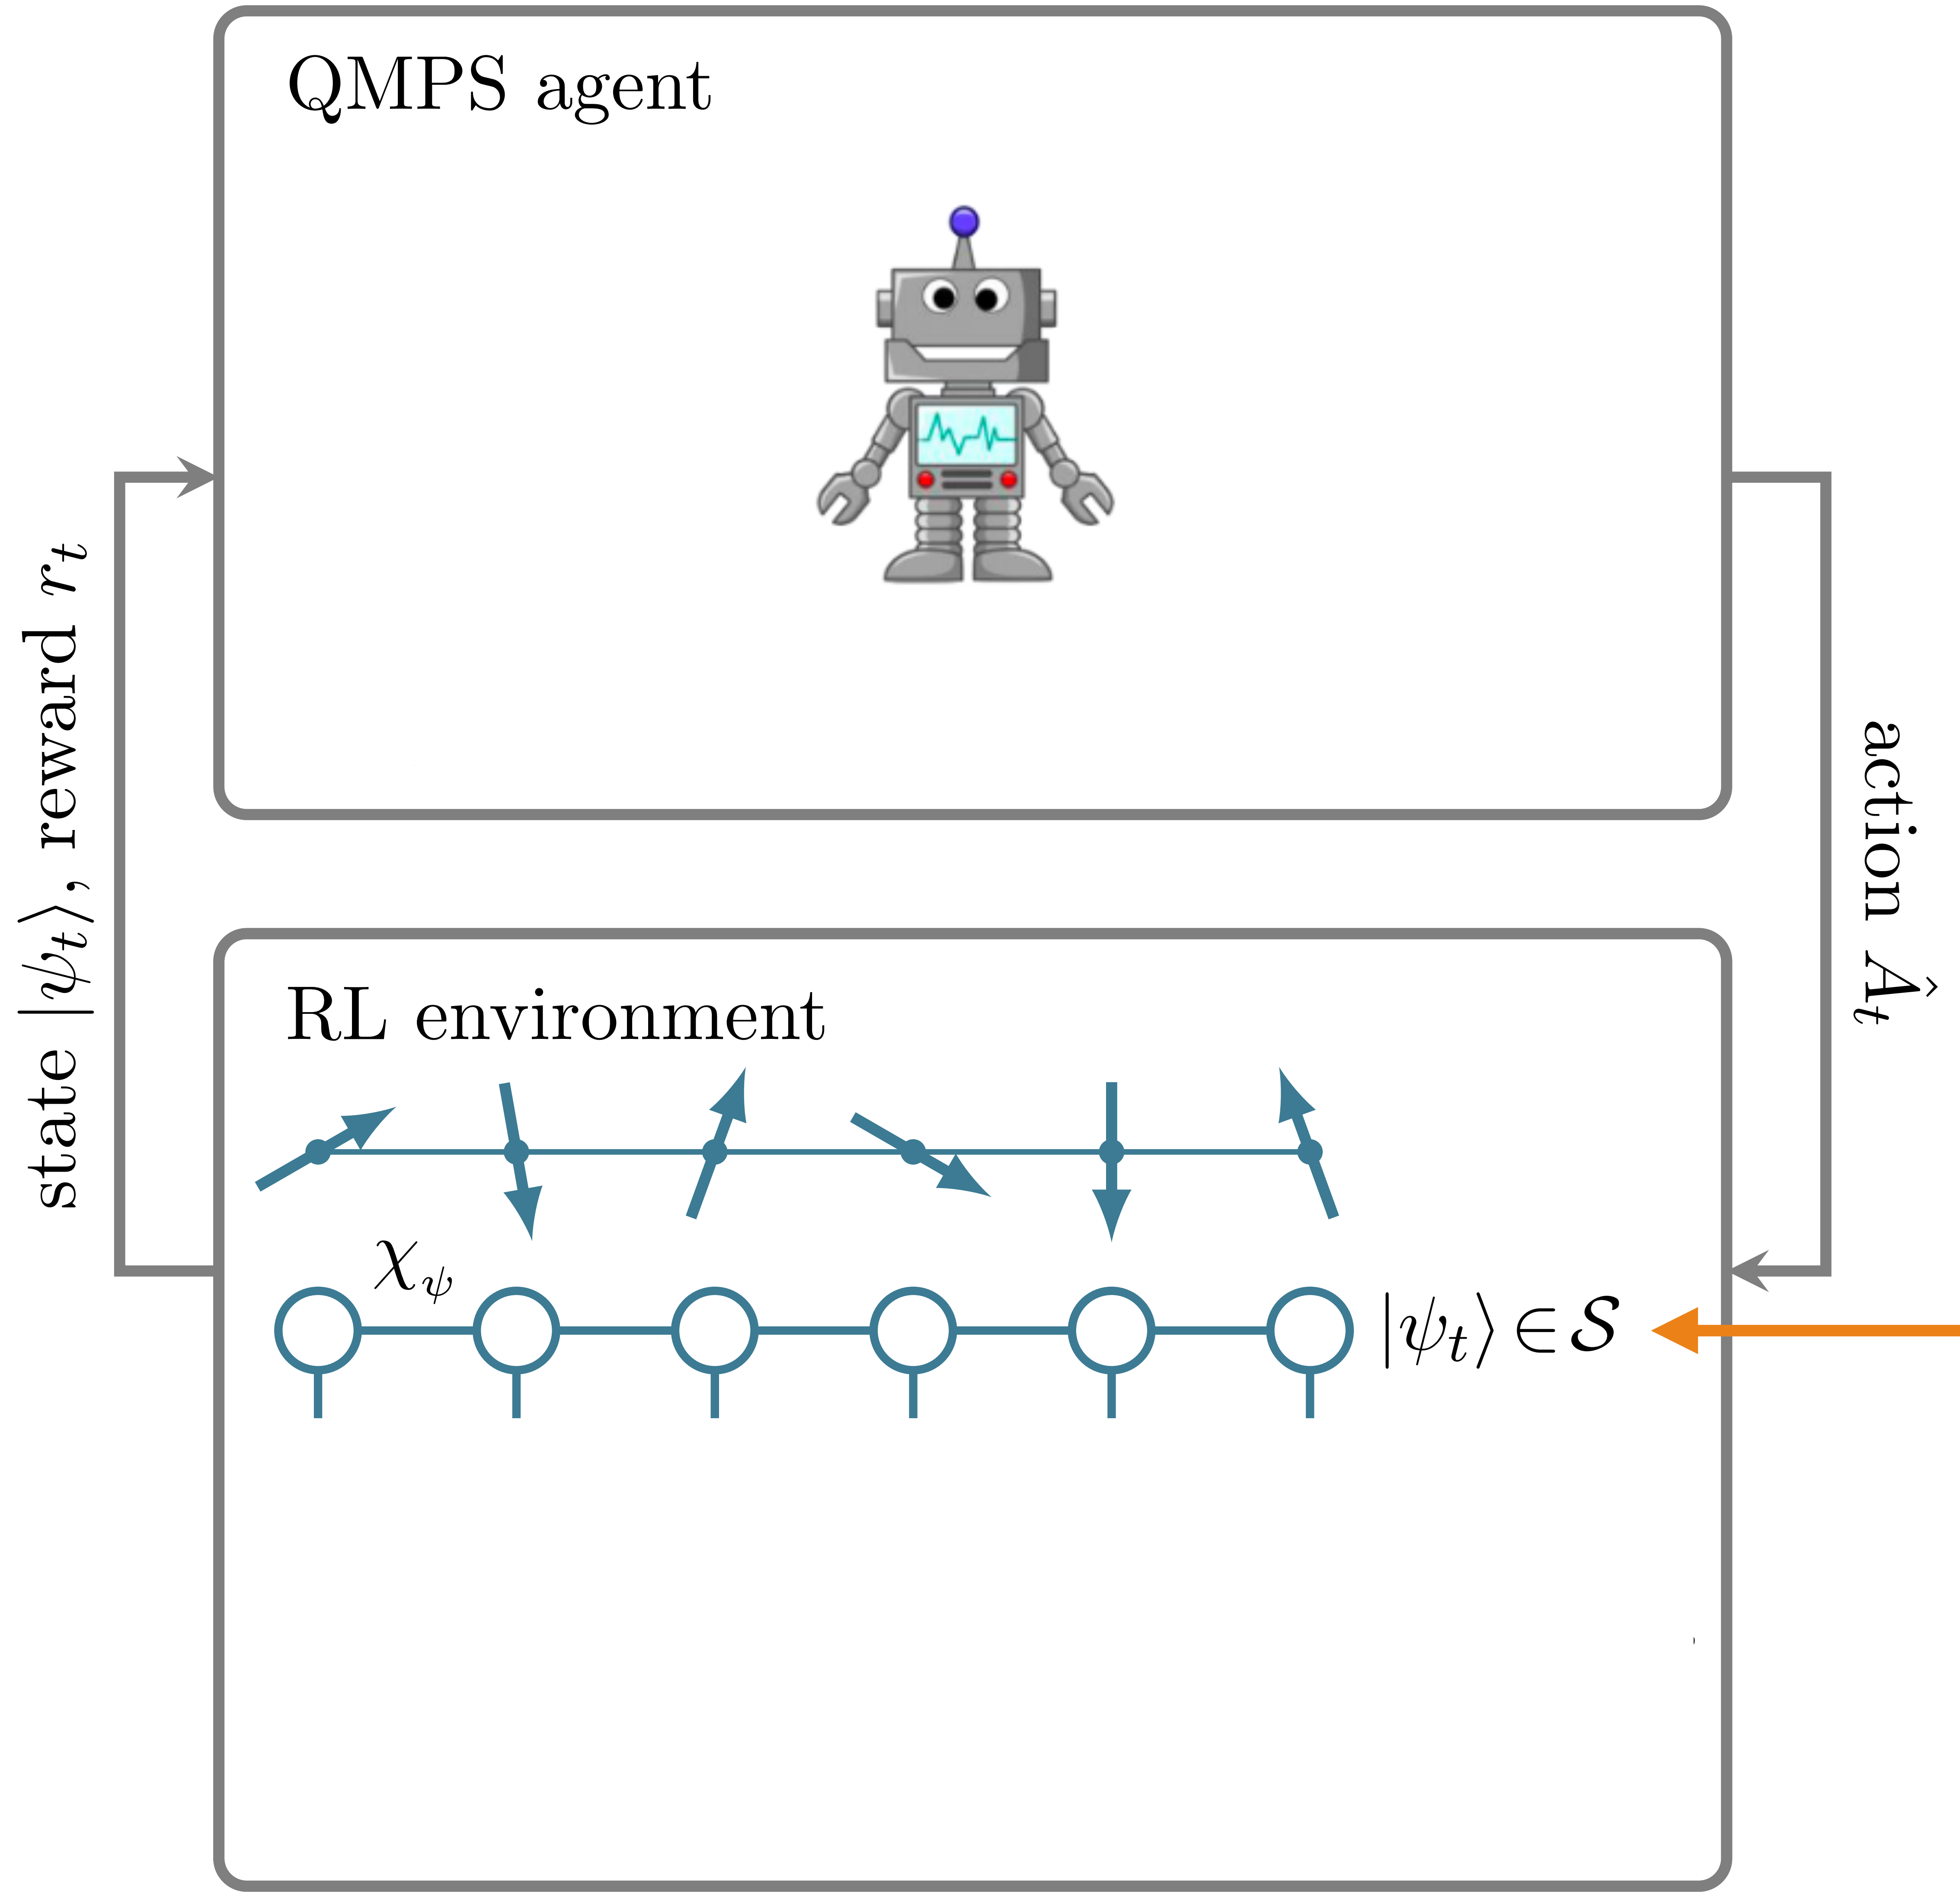

RL environment

State

MATRIX PRODUCT STATE ANSATZ FOR Q-LEARNING (QMPS)

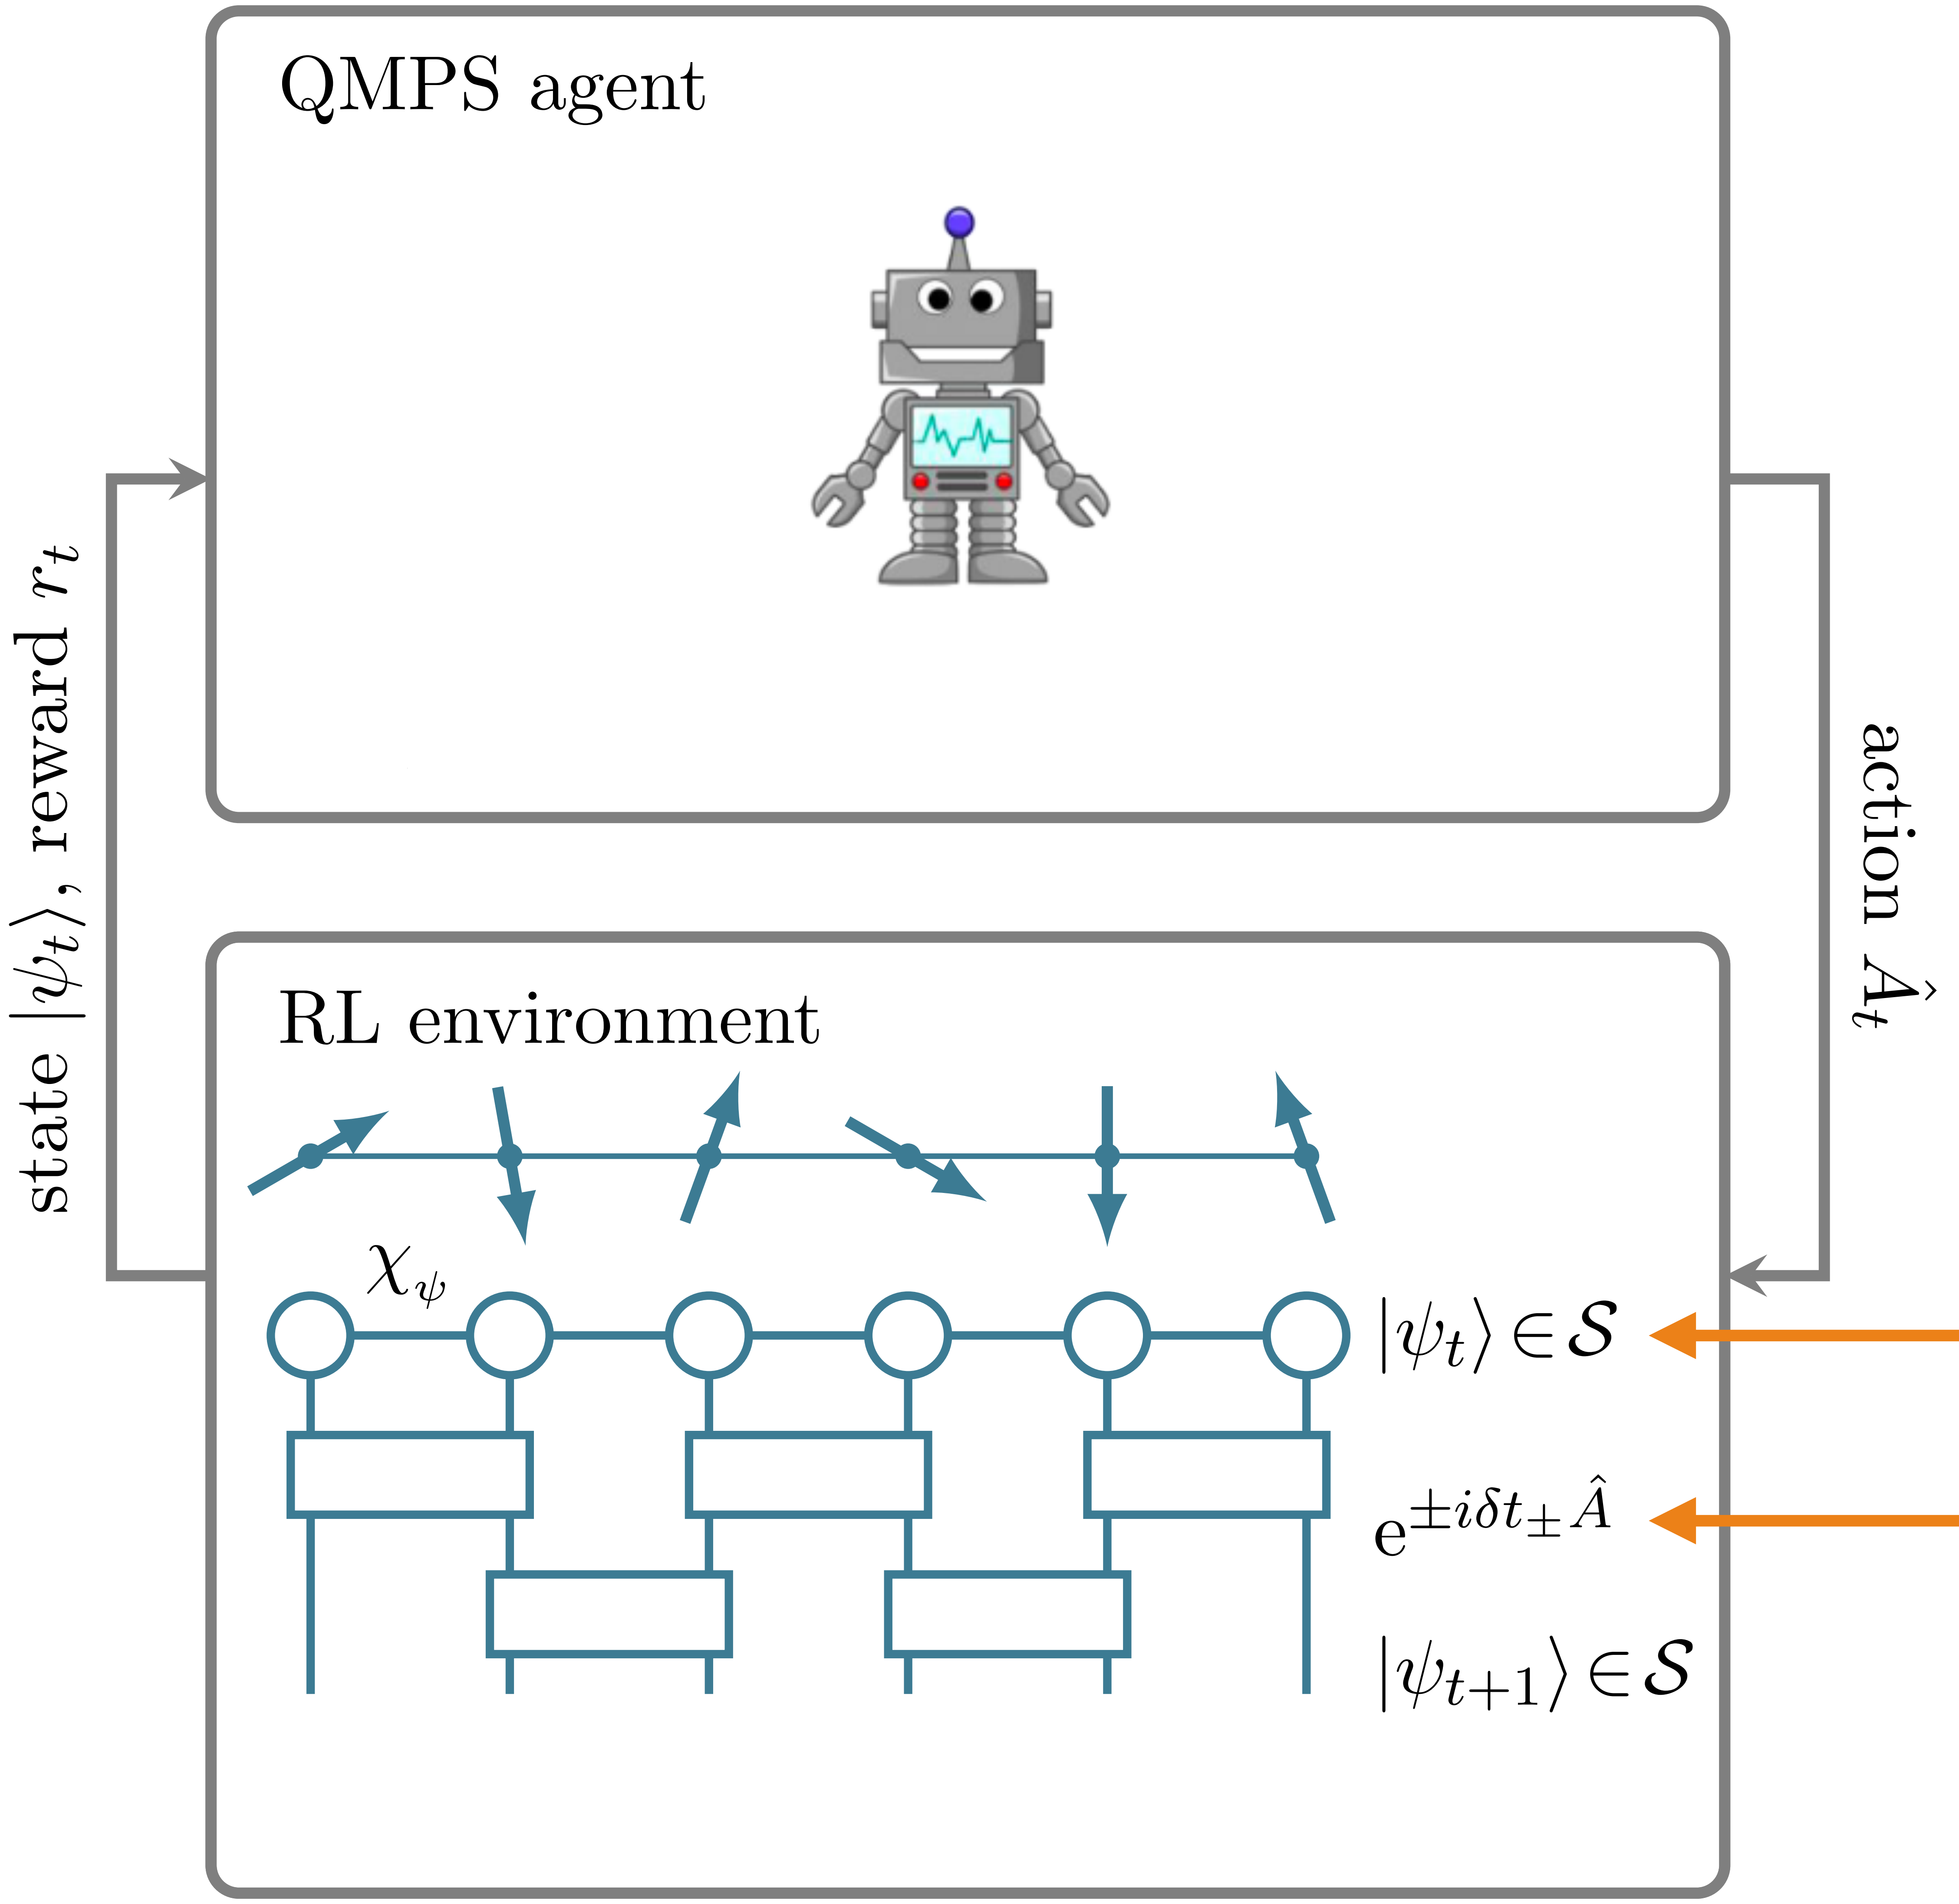

RL environment

State

Action

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

MATRIX PRODUCT STATE ANSATZ FOR Q-LEARNING (QMPS)

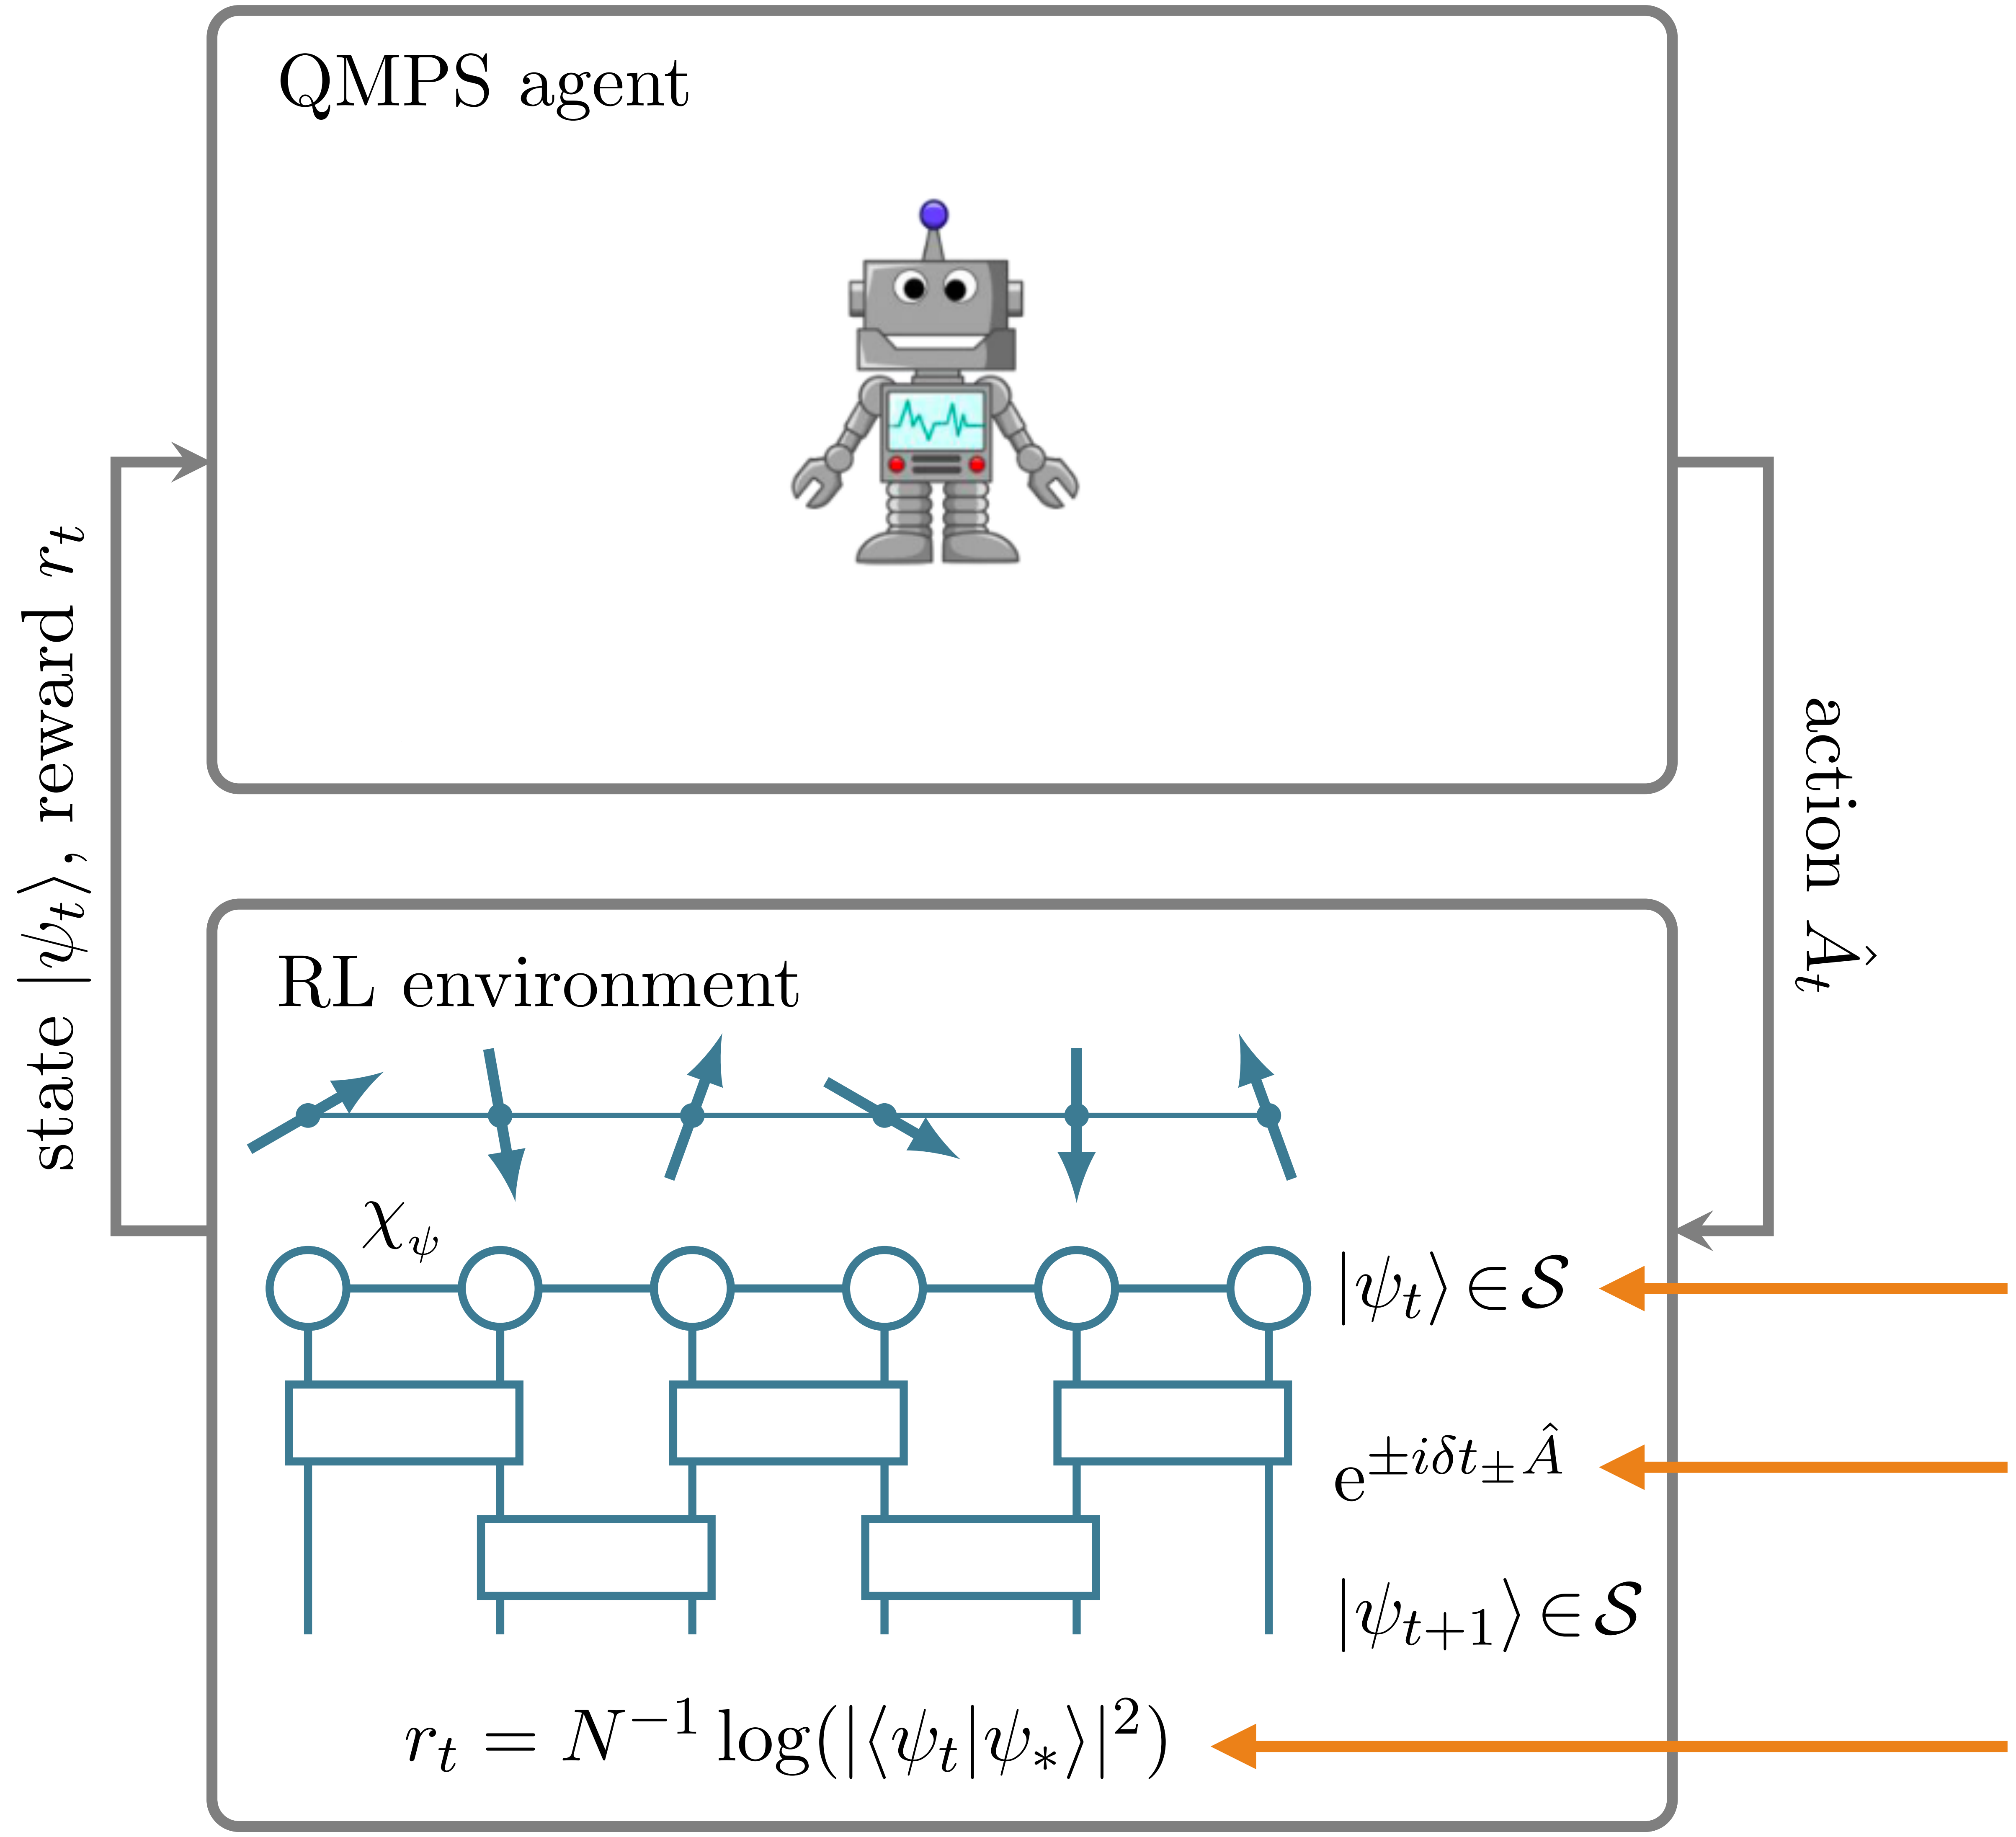

RL environment

State

Action

Reward (Goal: Prepare target state $|\psi_*\rangle$)

MATRIX PRODUCT STATE ANSATZ FOR Q-LEARNING (QMPS)

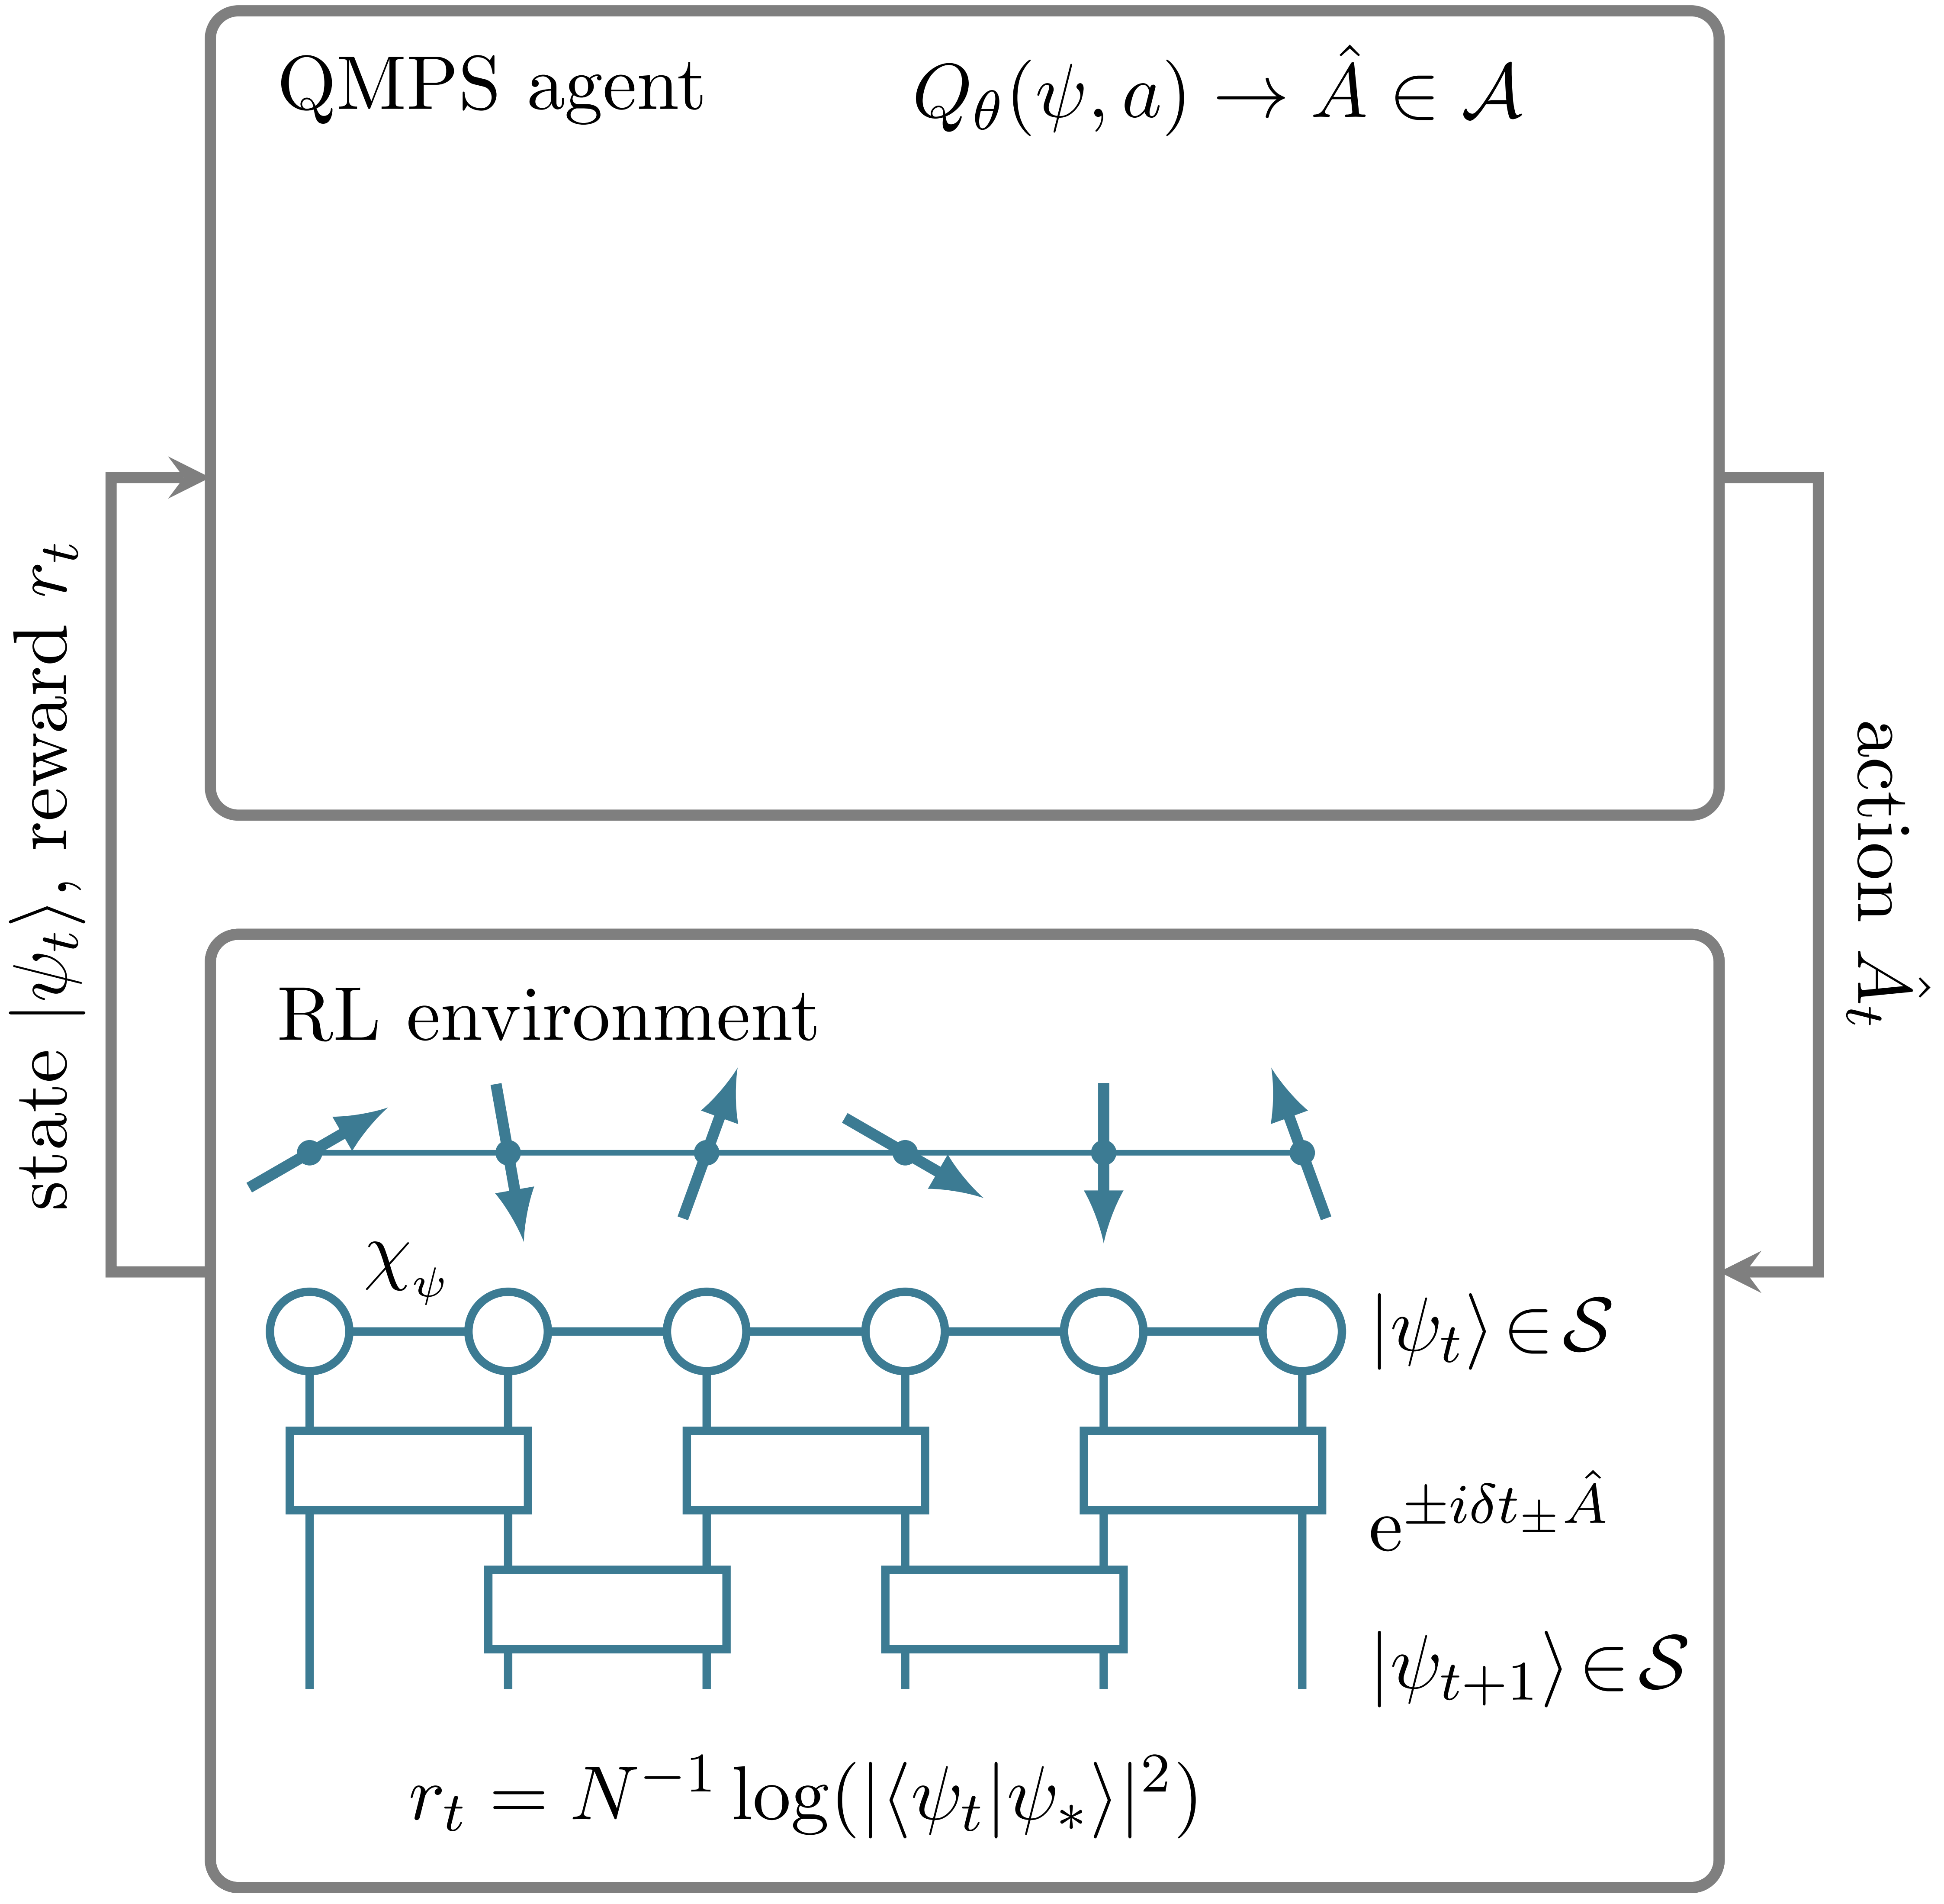

QMPS agent

MATRIX PRODUCT STATE ANSATZ FOR Q-LEARNING (QMPS)

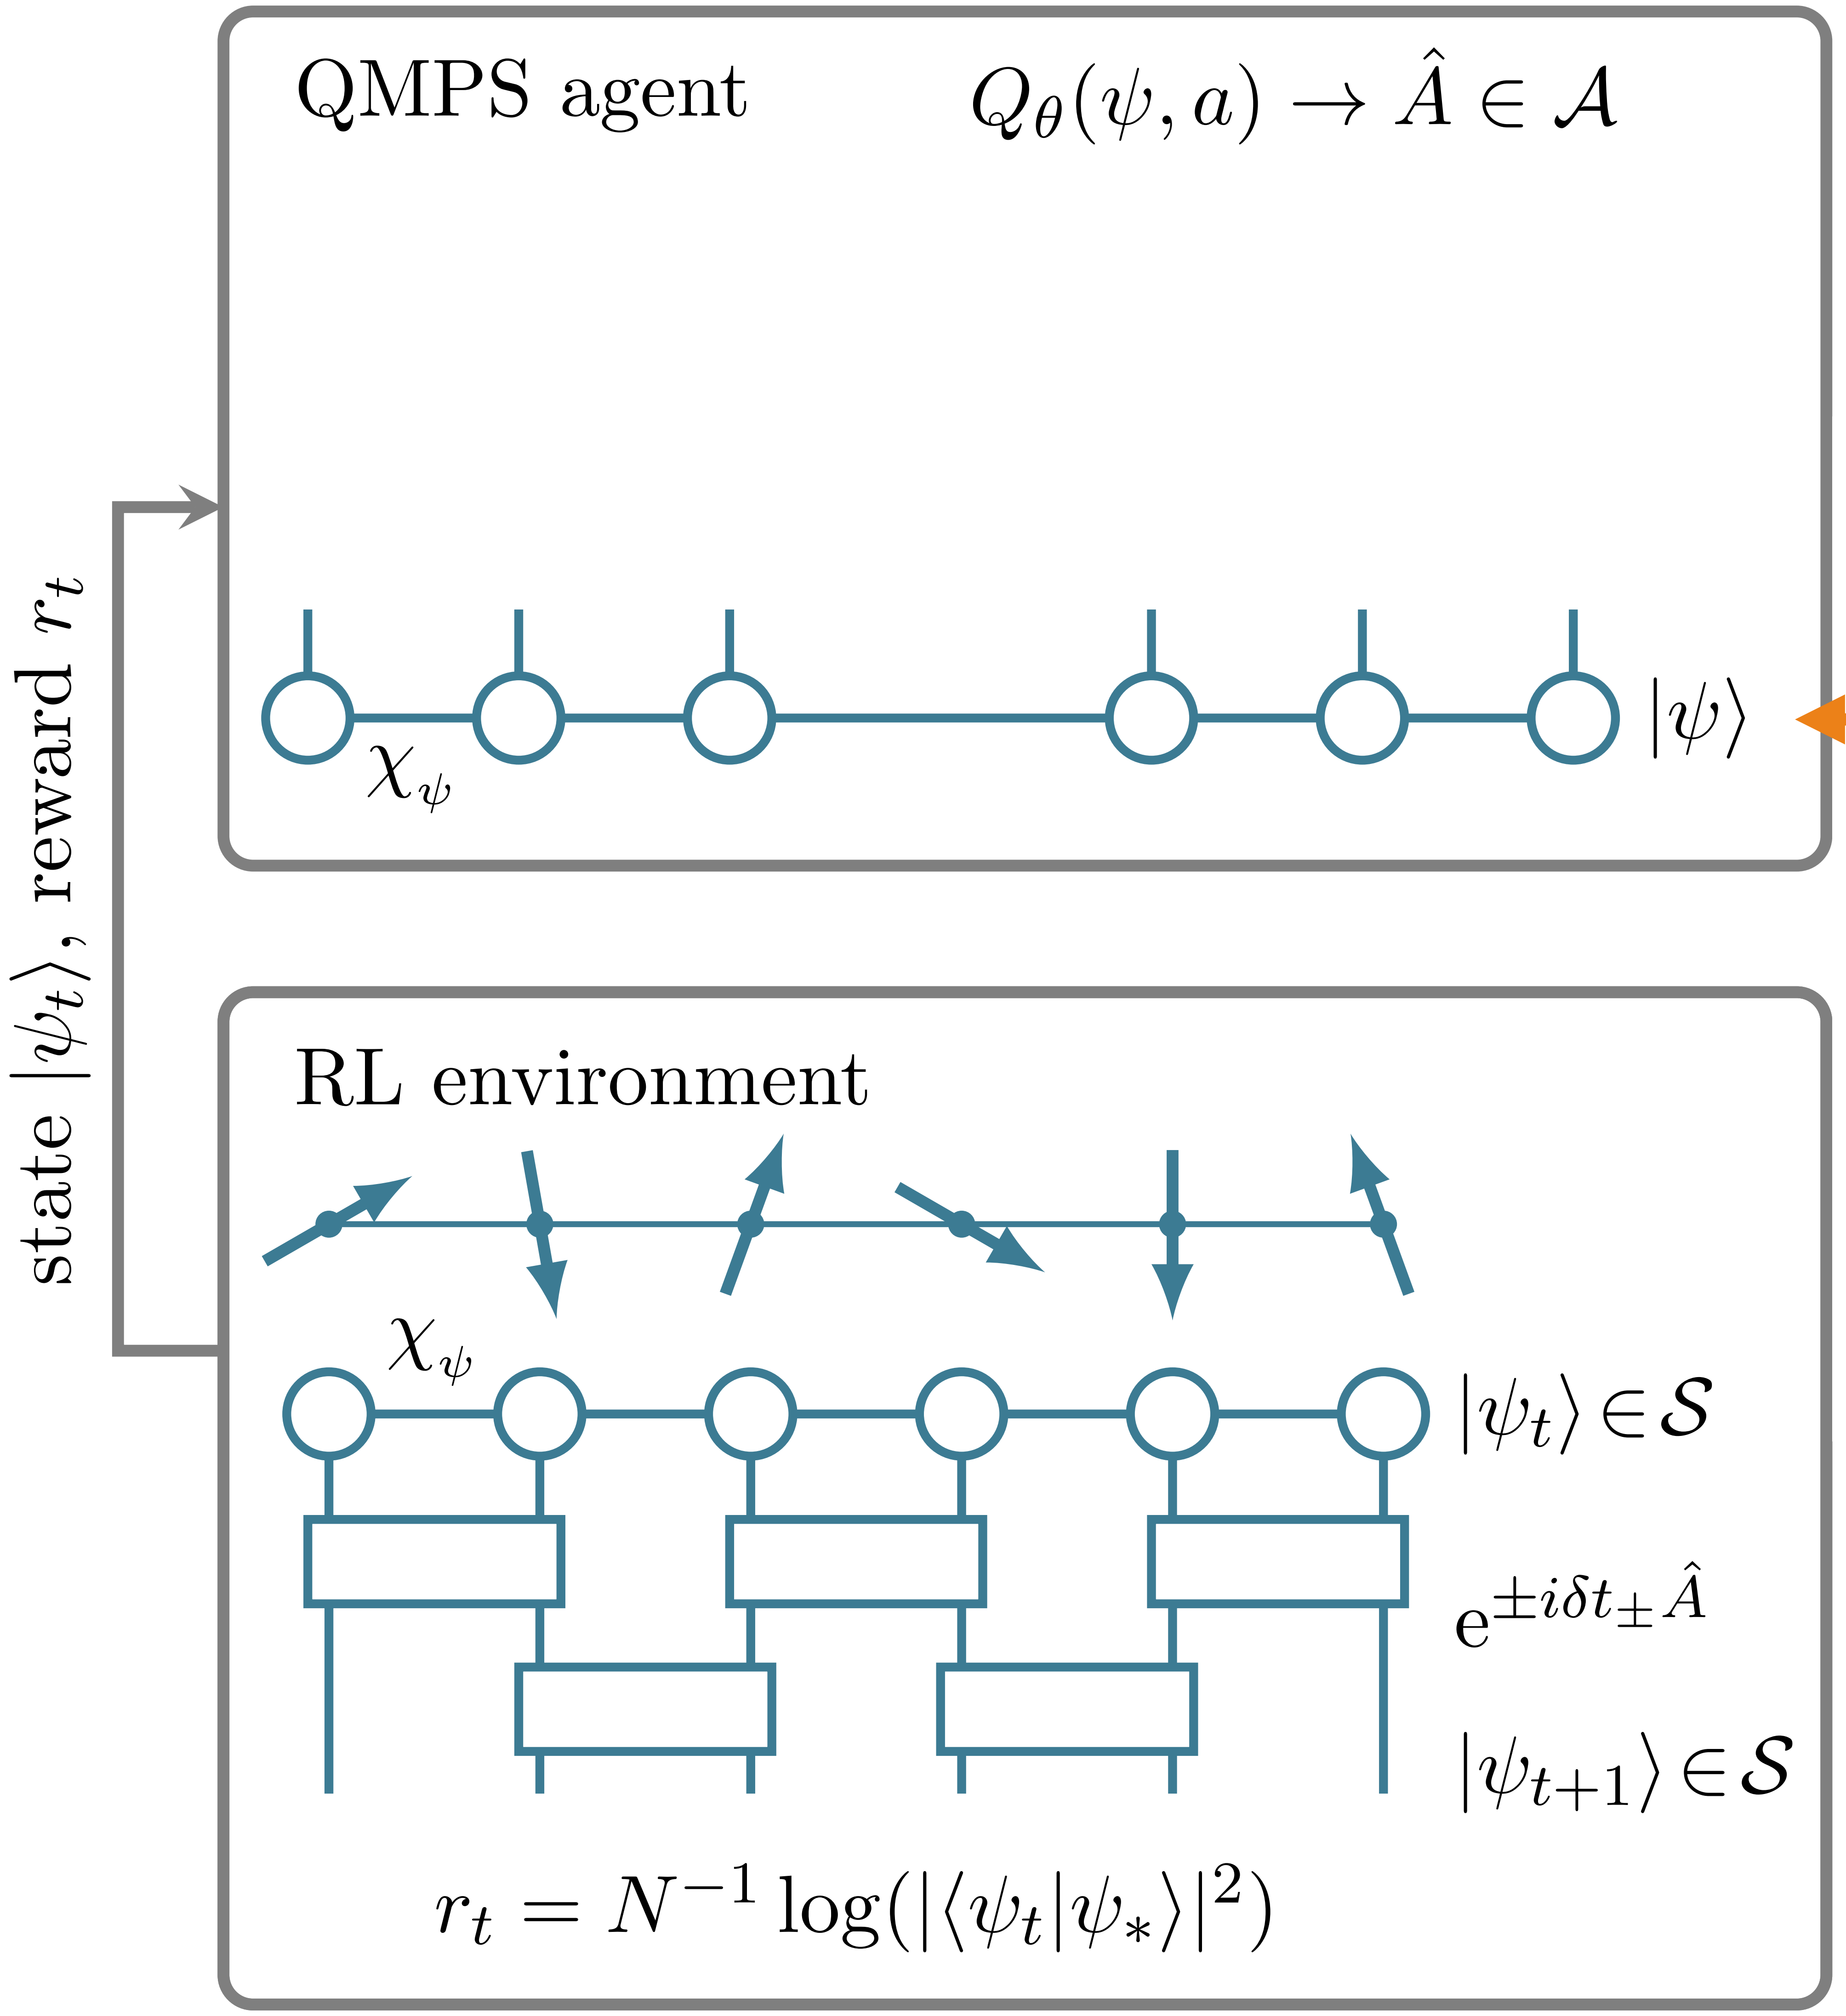

QMPS agent

Quantum state input (MPS)

MATRIX PRODUCT STATE ANSATZ FOR Q-LEARNING (QMPS)

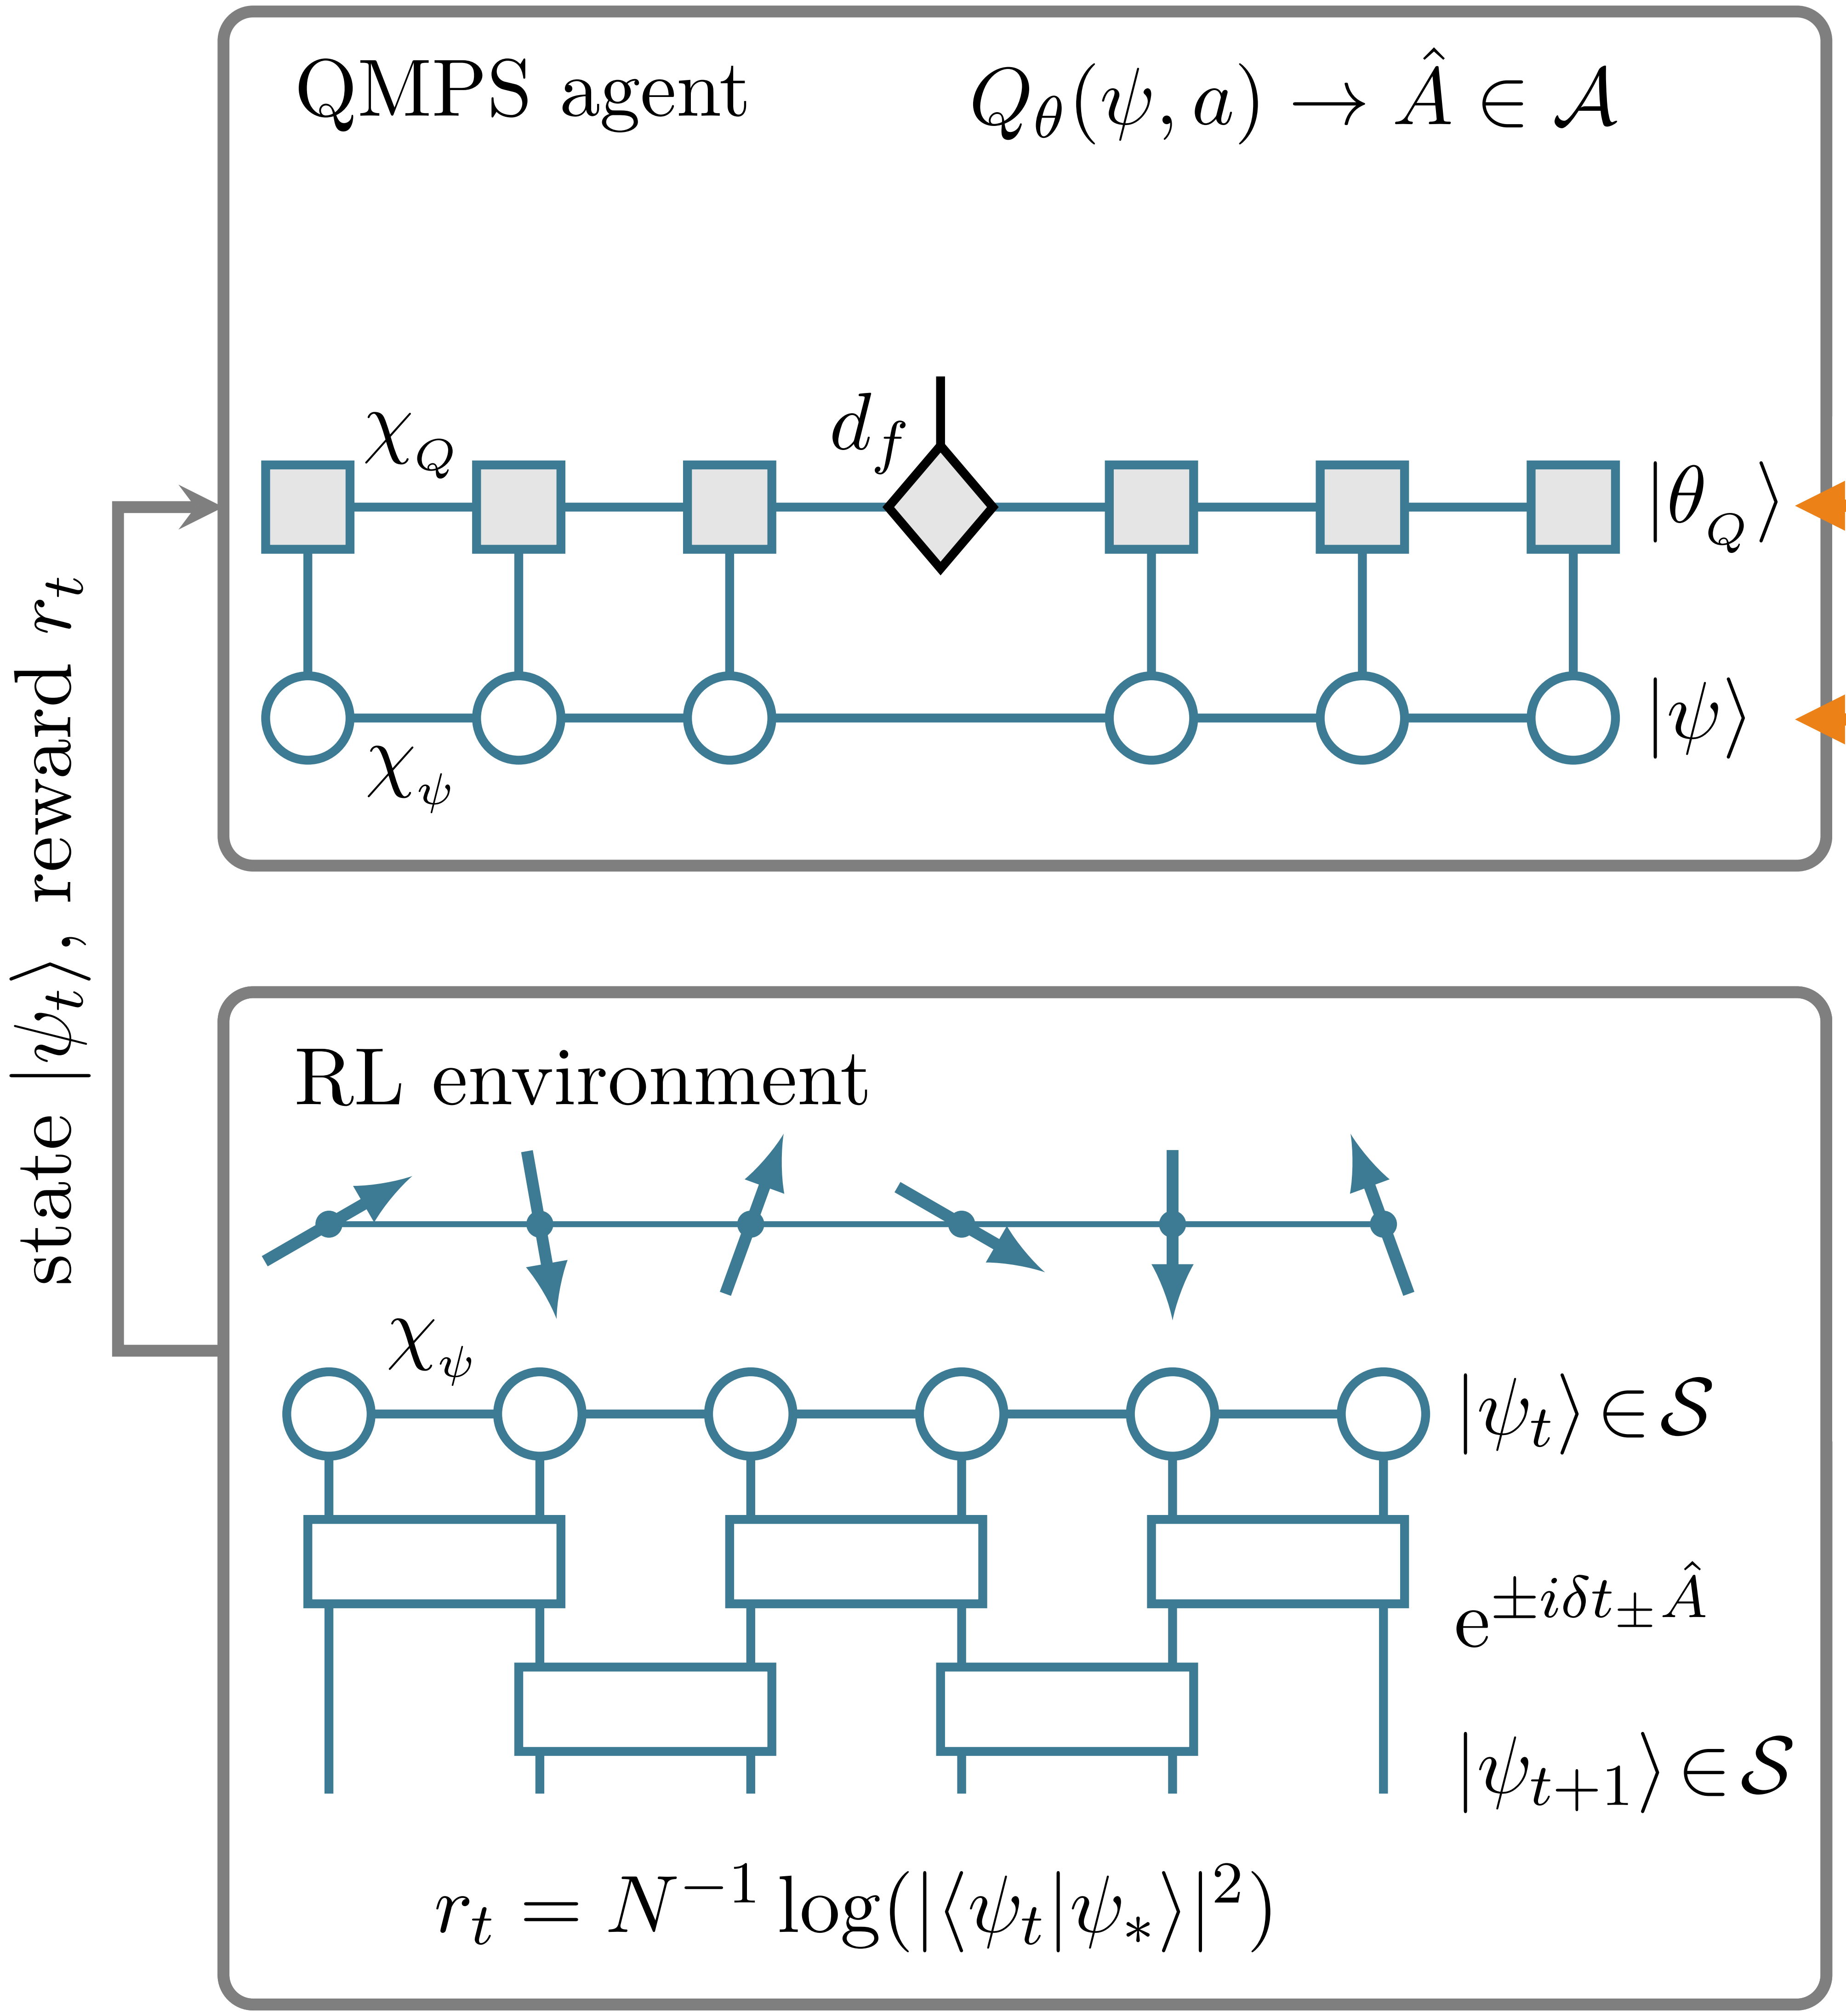

QMPS agent

Trainable parameters (MPS)

Quantum state input (MPS)

MATRIX PRODUCT STATE ANSATZ FOR Q-LEARNING (QMPS)

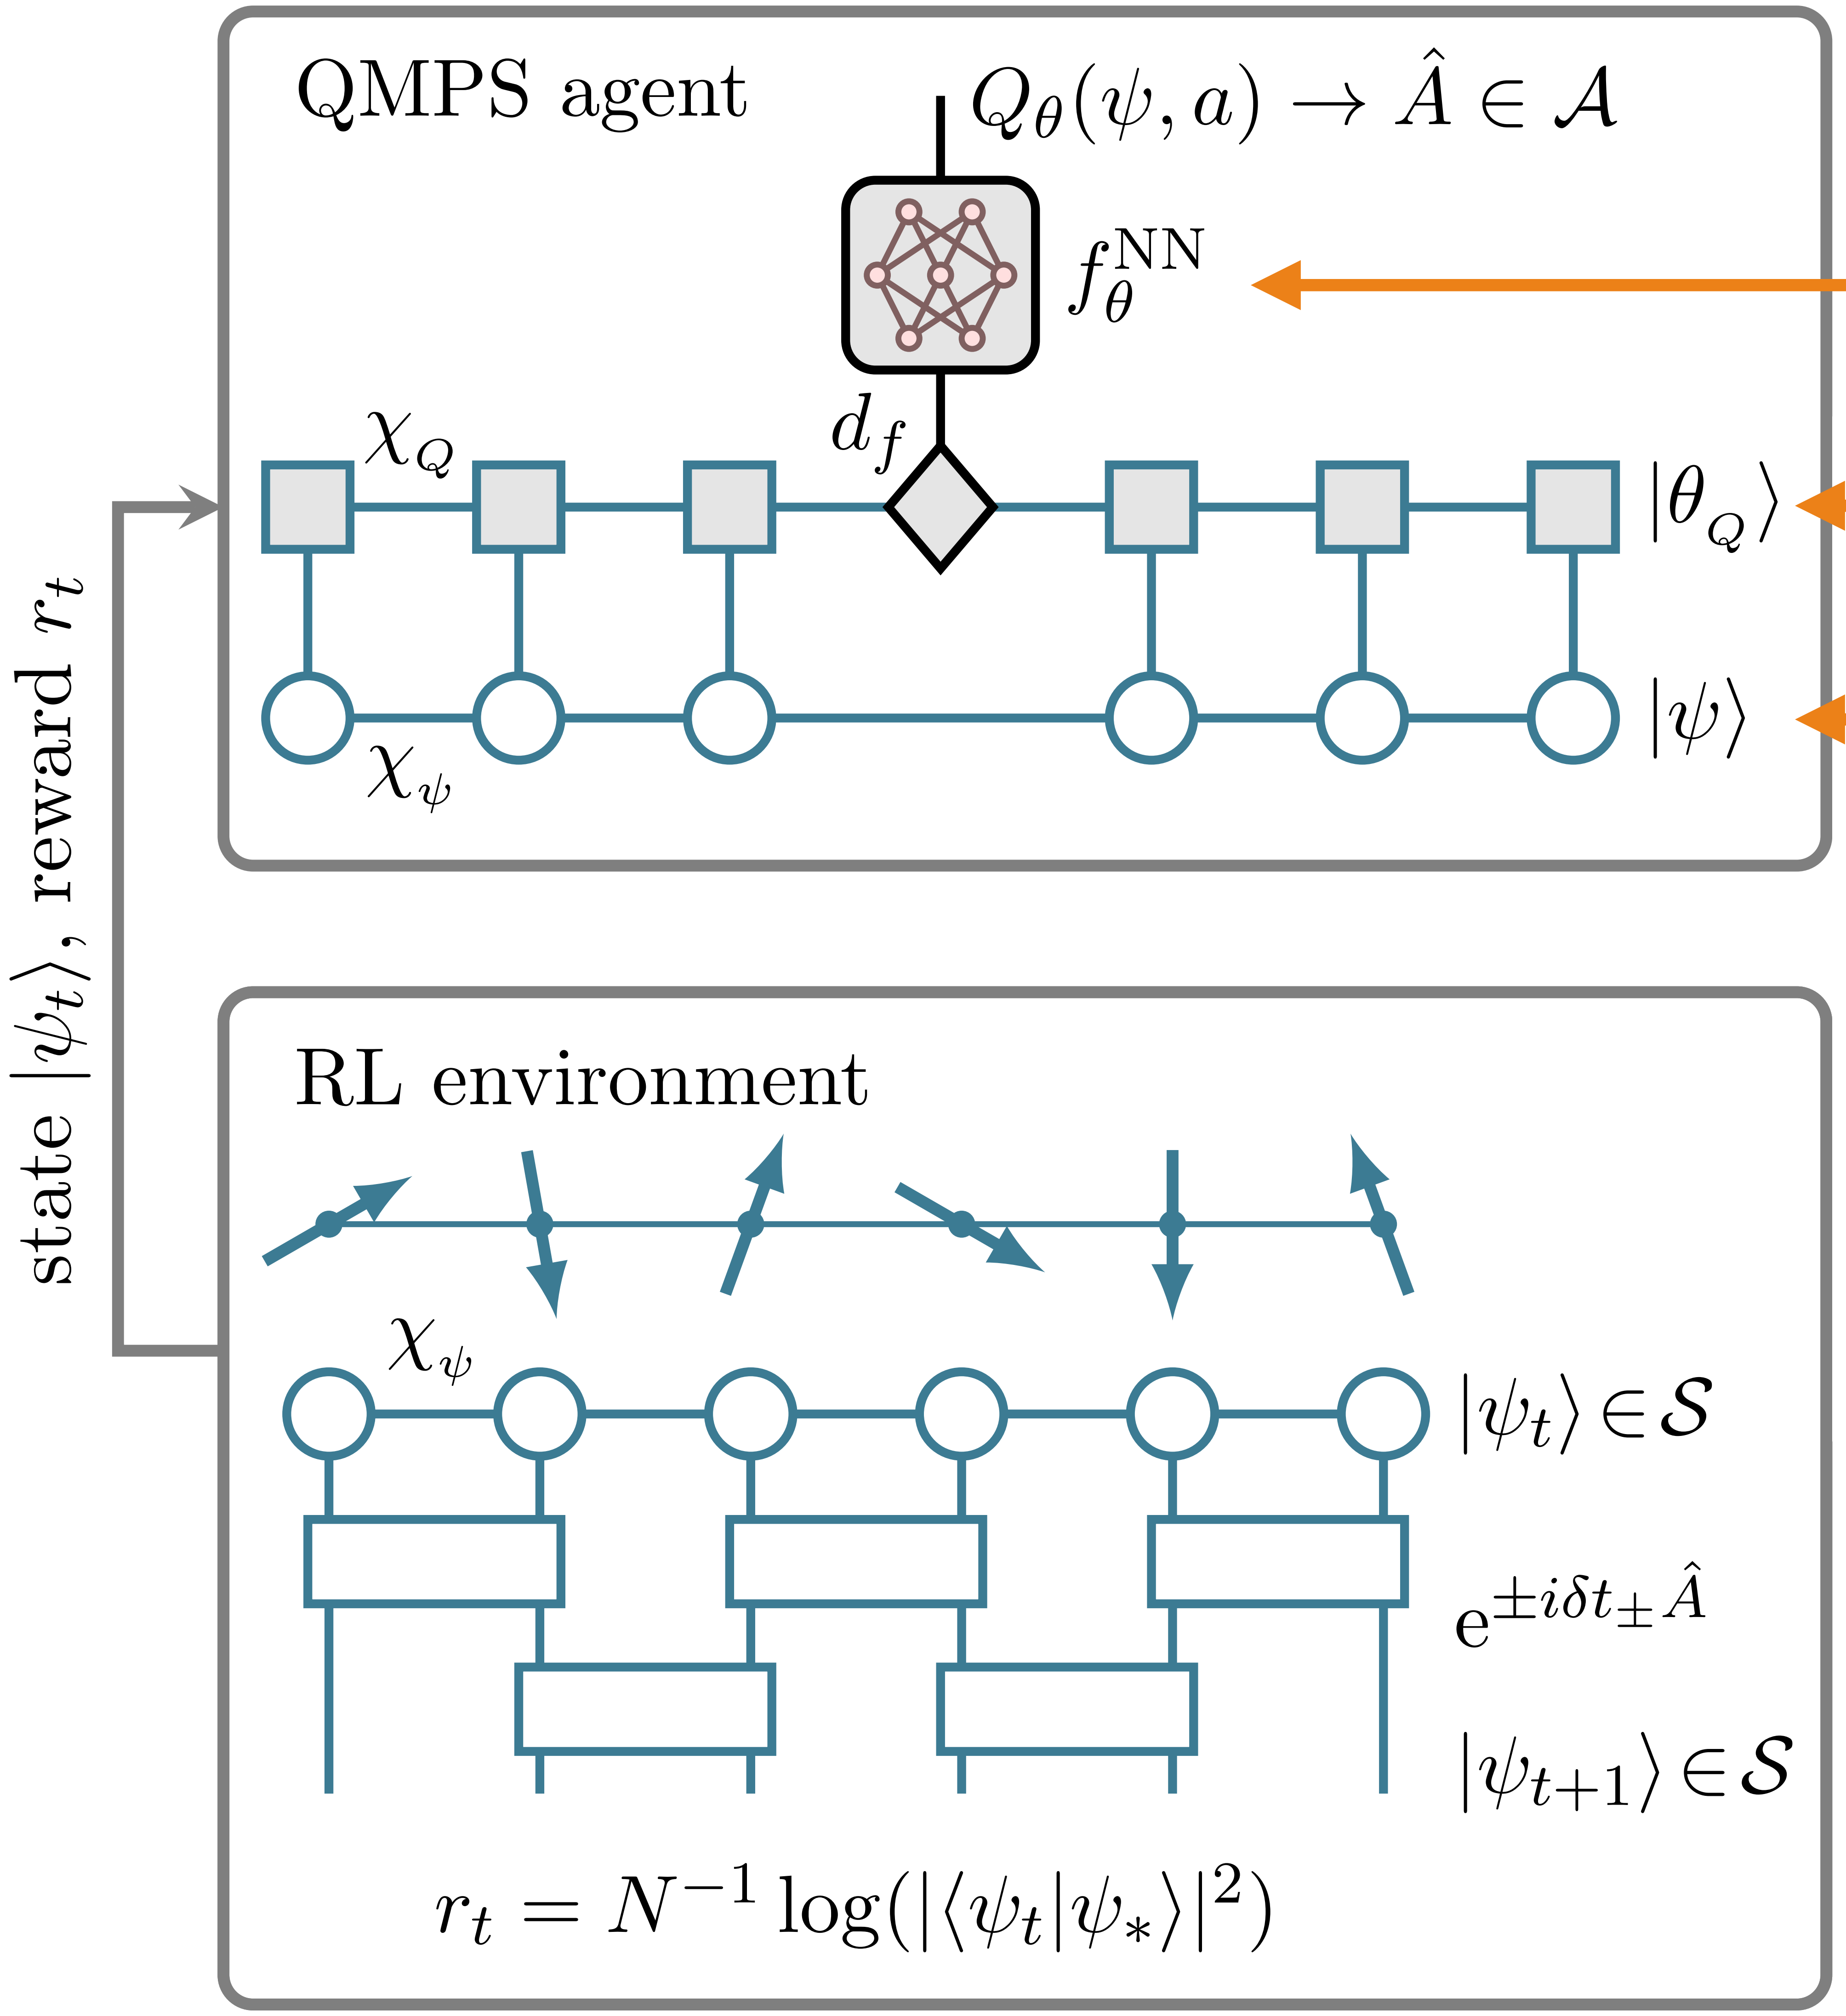

QMPS agent

Neural network (NN)

Trainable parameters (MPS)

Quantum state input (MPS)

MATRIX PRODUCT STATE ANSATZ FOR Q-LEARNING (QMPS)

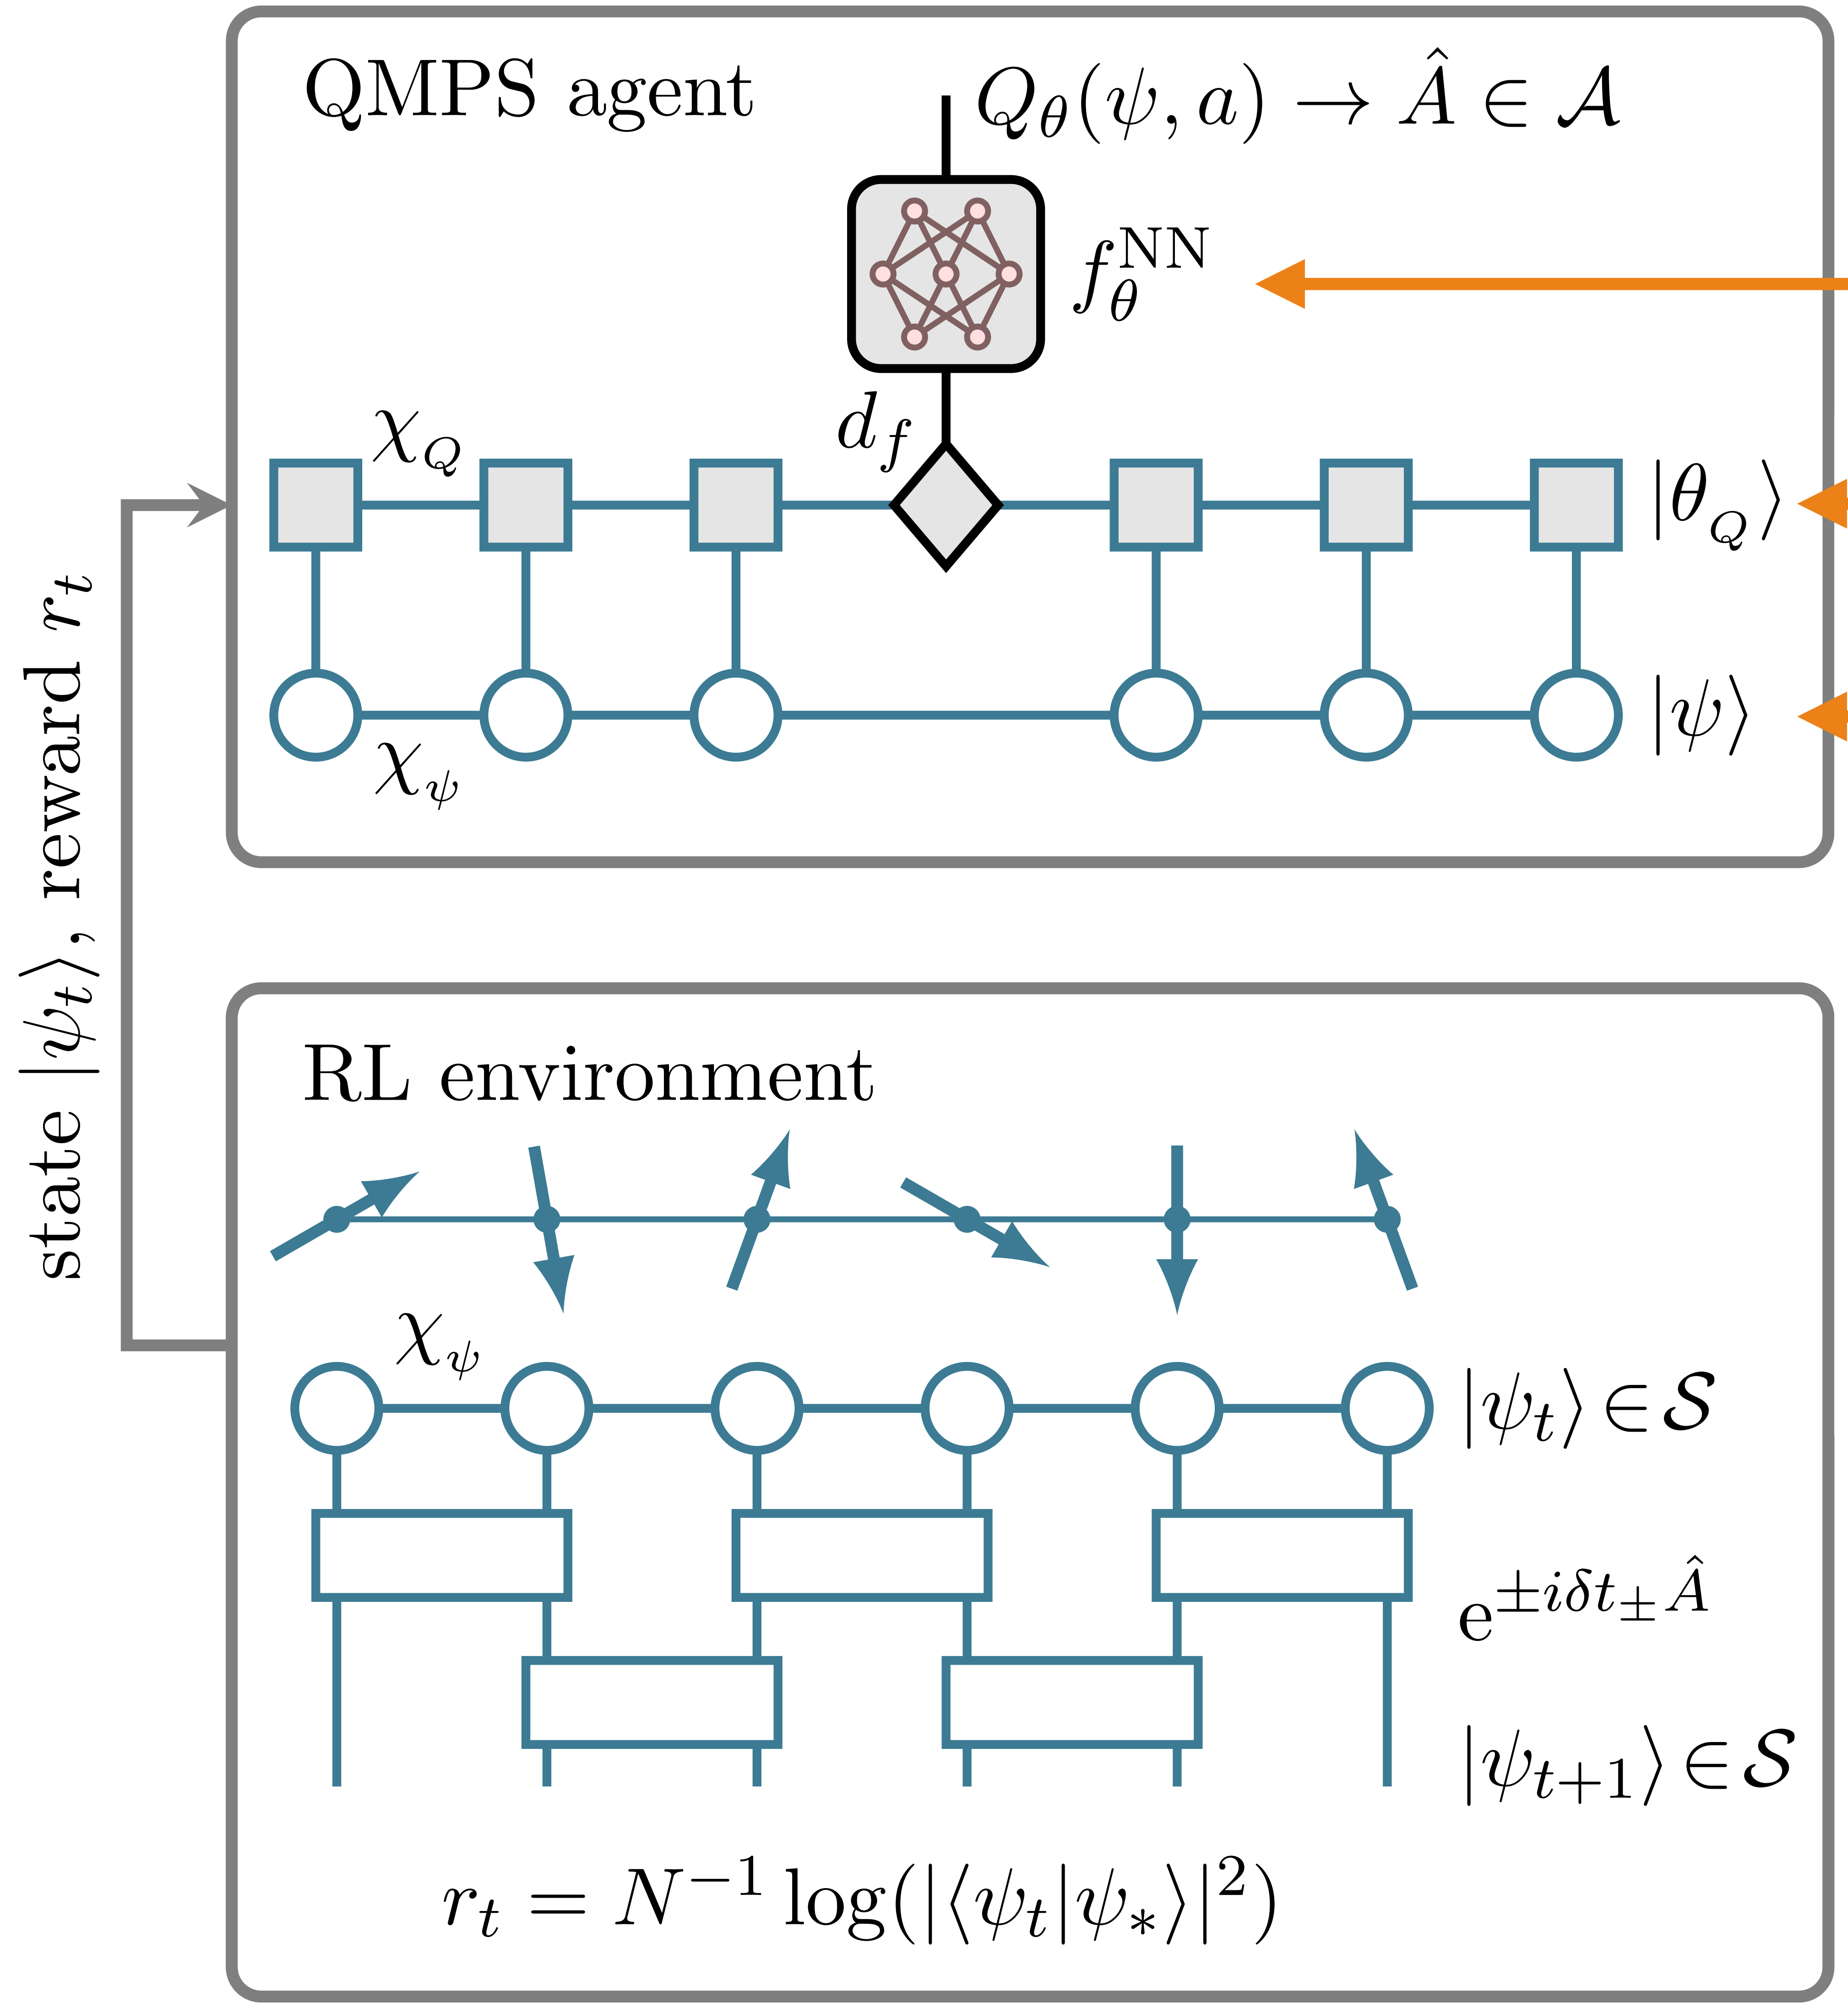

QMPS agent

Neural network (NN)

Trainable parameters (MPS)

Quantum state input (MPS)

→ QMPS: Resources scale only linearly in system size N

ROBUST CRITICAL STATE PREPARATION

ROBUST CRITICAL STATE PREPARATION

Mixed-field Ising: $\hat{H}_{\text{Ising}} = J \sum_{i=1}^{N-1} \hat{Z}_i \hat{Z}_{i+1} - g_x \sum_{i=1}^{N} \hat{X}_i - g_z \sum_{i=1}^{N} \hat{Z}_i$

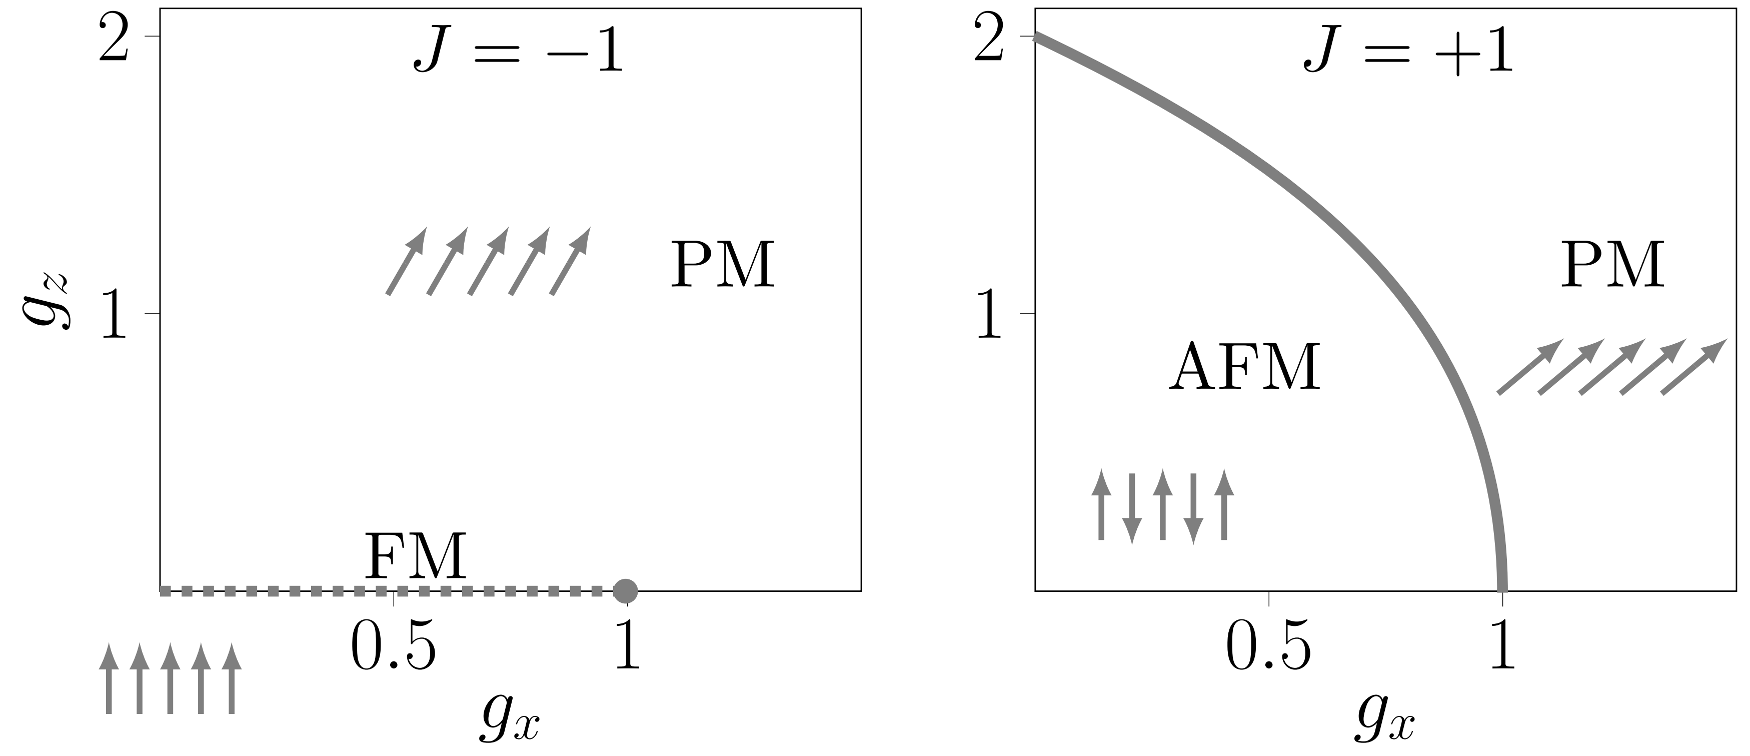

ROBUST CRITICAL STATE PREPARATION

Mixed-field Ising: $\hat{H}_{\mathrm{Ising}} = J \sum_{i=1}^{N-1} \hat{Z}_i \hat{Z}_{i+1} - g_x \sum_{i=1}^{N} \hat{X}_i - g_z \sum_{i=1}^{N} \hat{Z}_i$

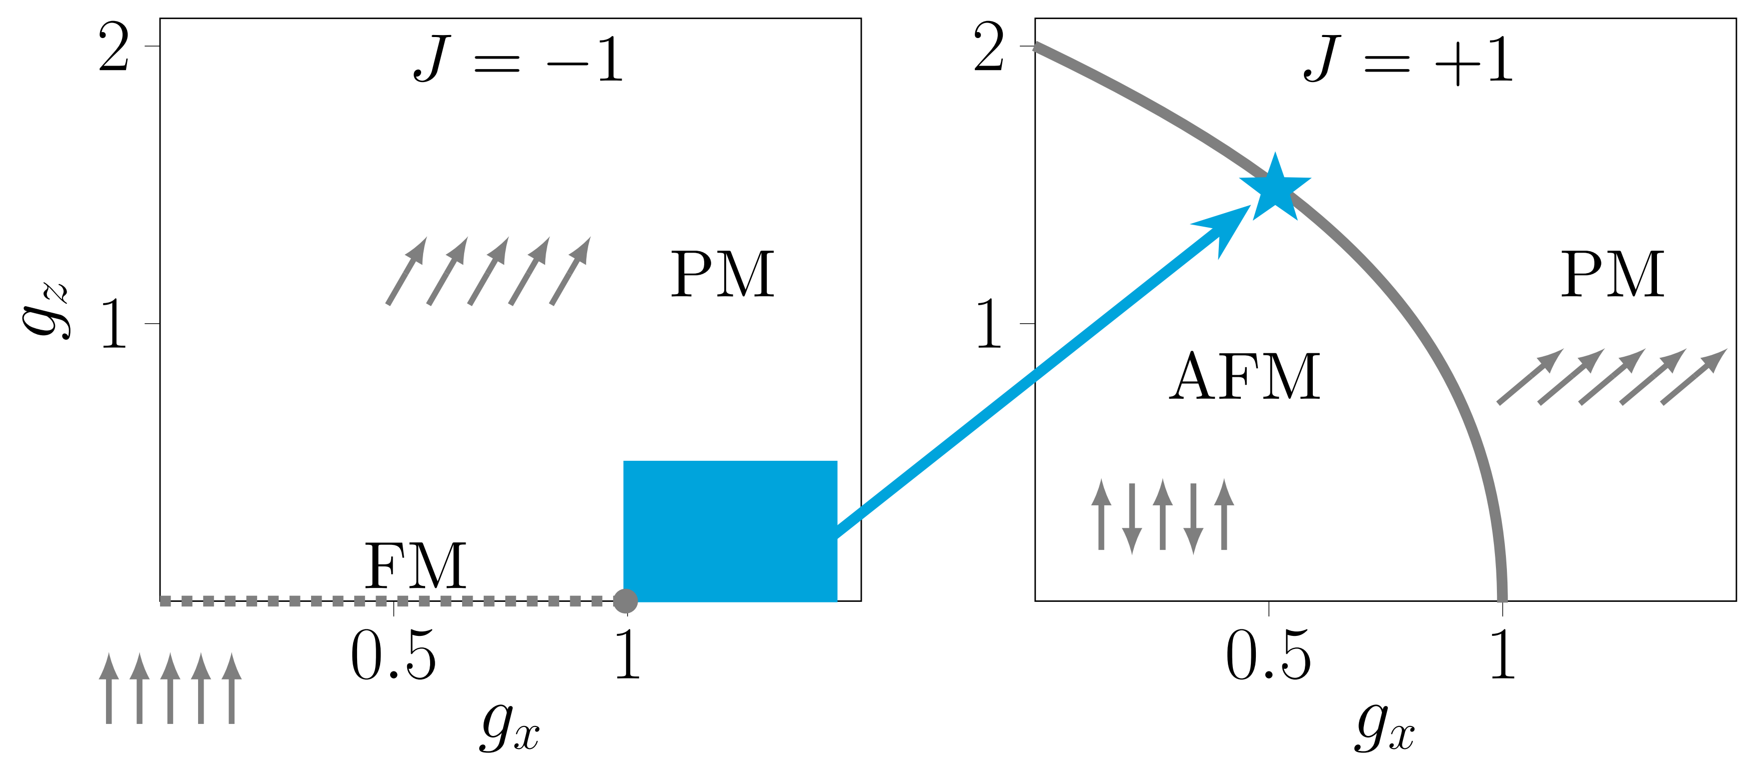

ROBUST CRITICAL STATE PREPARATION

Mixed-field Ising: $\hat{H}_{\text{Ising}} = J\sum_{i=1}^{N-1}\hat{Z}_i\hat{Z}_{i+1} - g_x\sum_{i=1}^{N}\hat{X}_i - g_z\sum_{i=1}^{N}\hat{Z}_i$

$\rightarrow N = 16$ spins

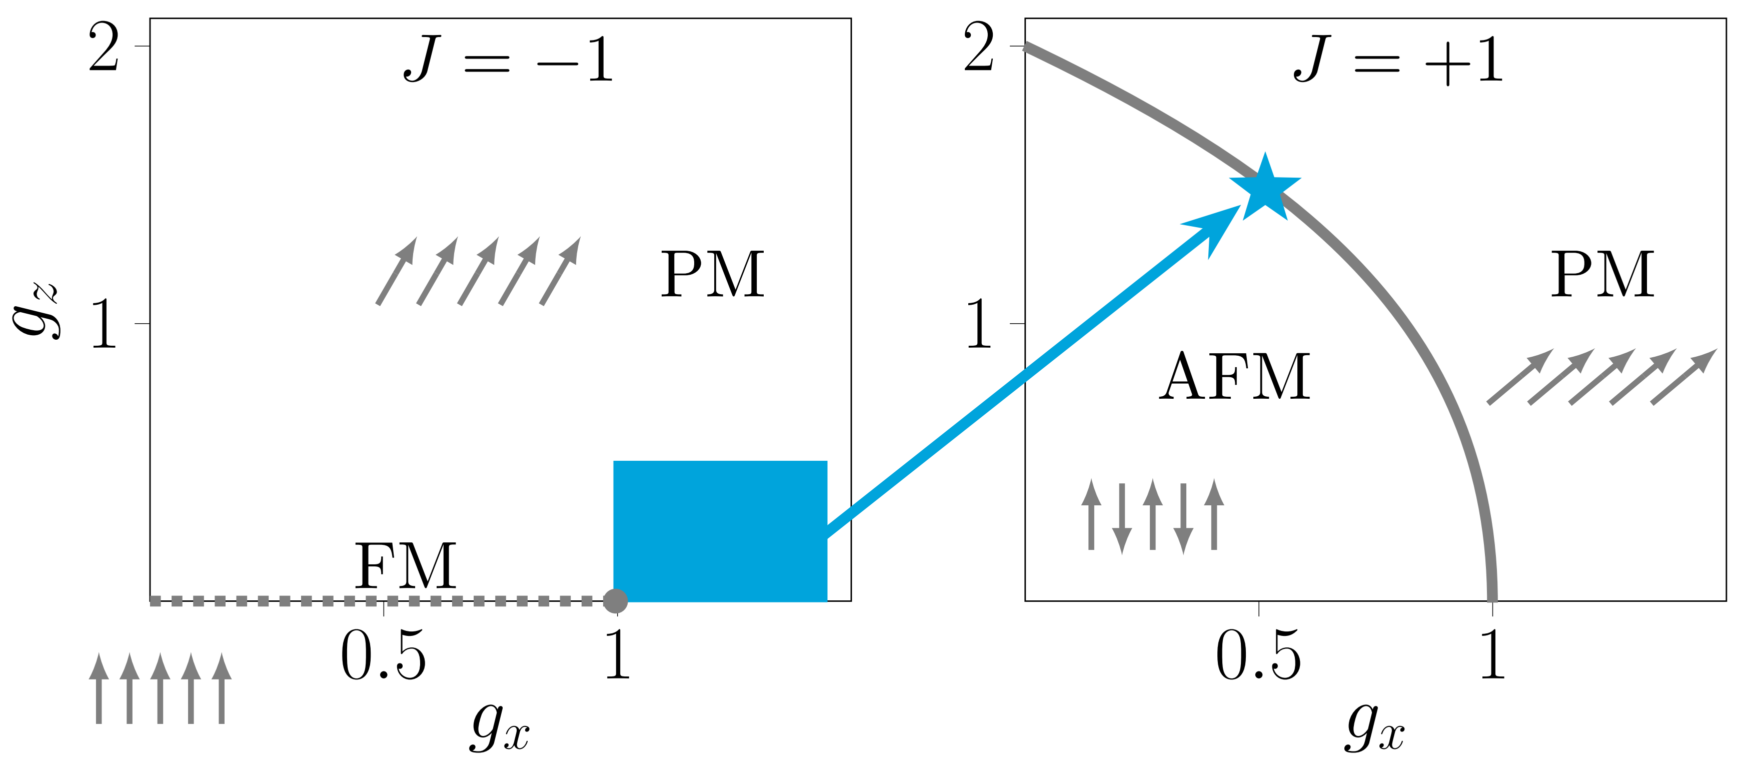

ROBUST CRITICAL STATE PREPARATION

Mixed-field Ising: $\hat{H}_{\text{Ising}} = J \sum_{i=1}^{N-1} \hat{Z}_i \hat{Z}_{i+1} - g_x \sum_{i=1}^{N} \hat{X}_i - g_z \sum_{i=1}^{N} \hat{Z}_i$

$\rightarrow N = 16$ spins

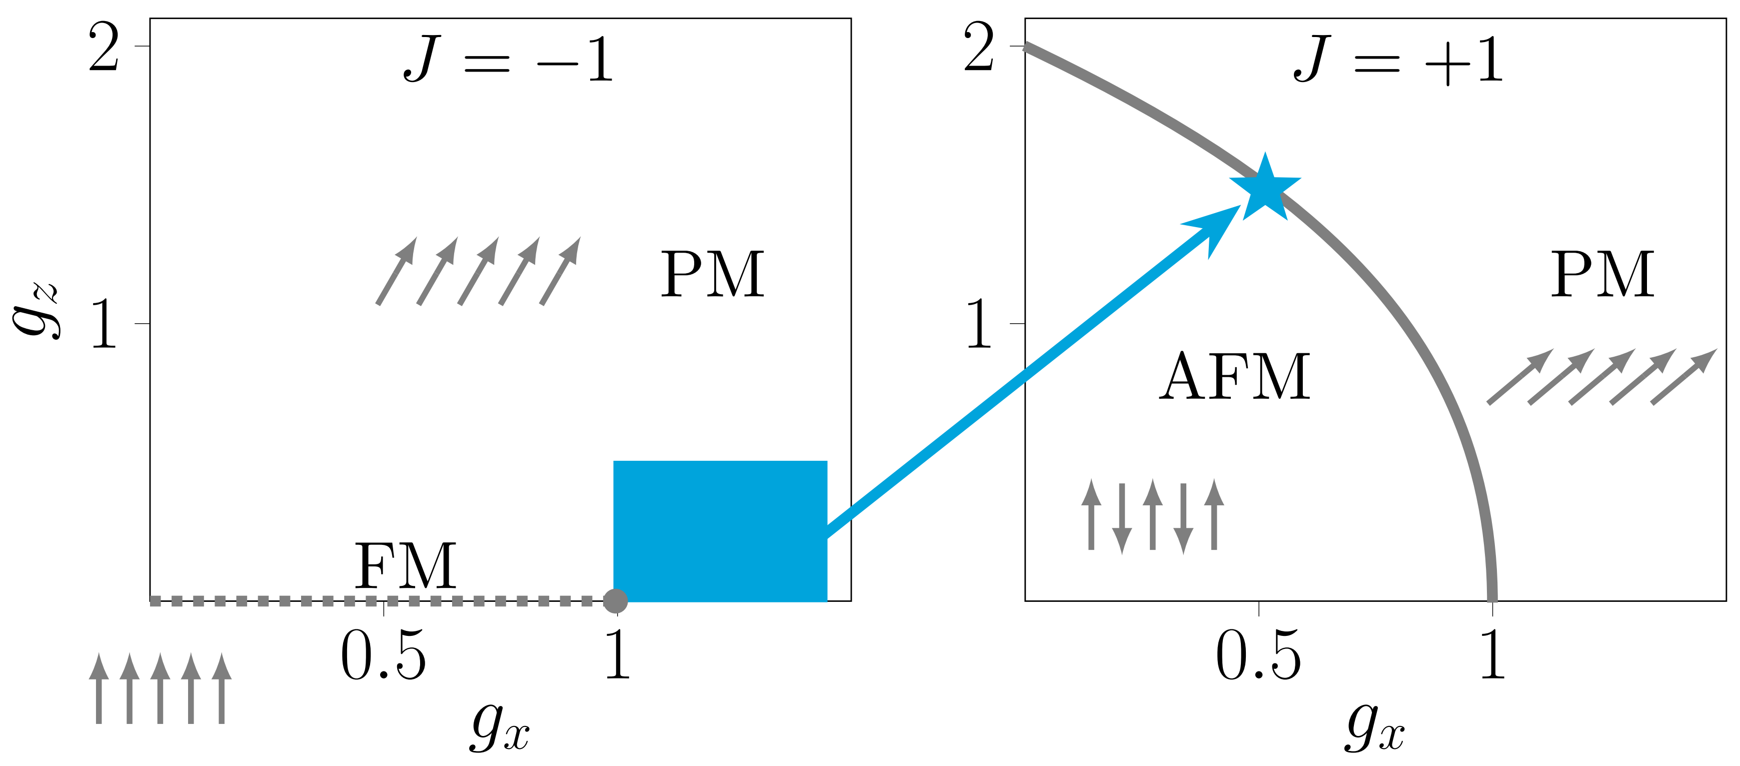

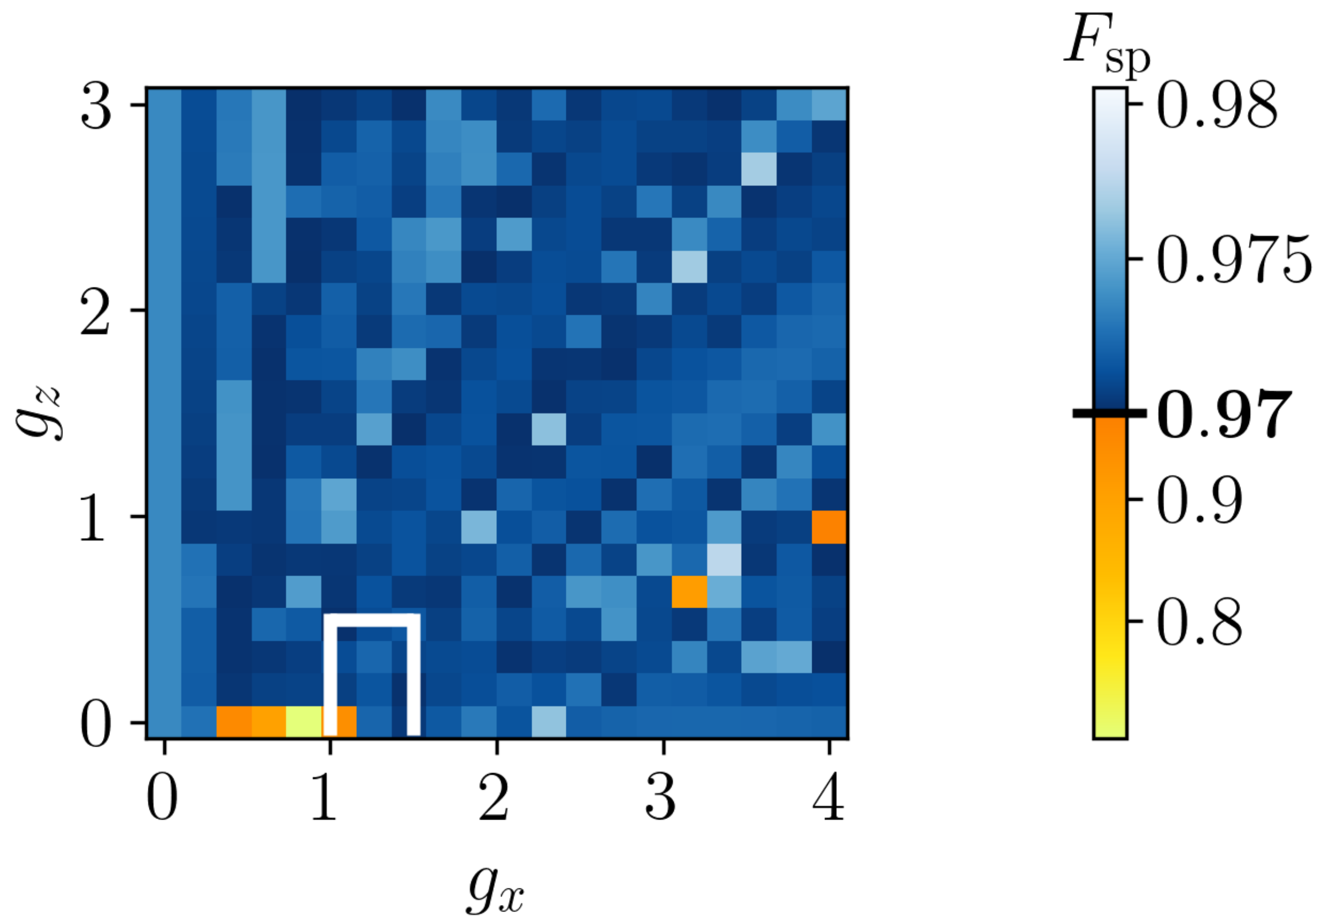

ROBUST CRITICAL STATE PREPARATION

Mixed-field Ising: $\hat{H}_{\text{Ising}} = J \sum_{i=1}^{N-1} \hat{Z}_i \hat{Z}_{i+1} - g_x \sum_{i=1}^{N} \hat{X}_i - g_z \sum_{i=1}^{N} \hat{Z}_i$

$\rightarrow N = 16$ spins

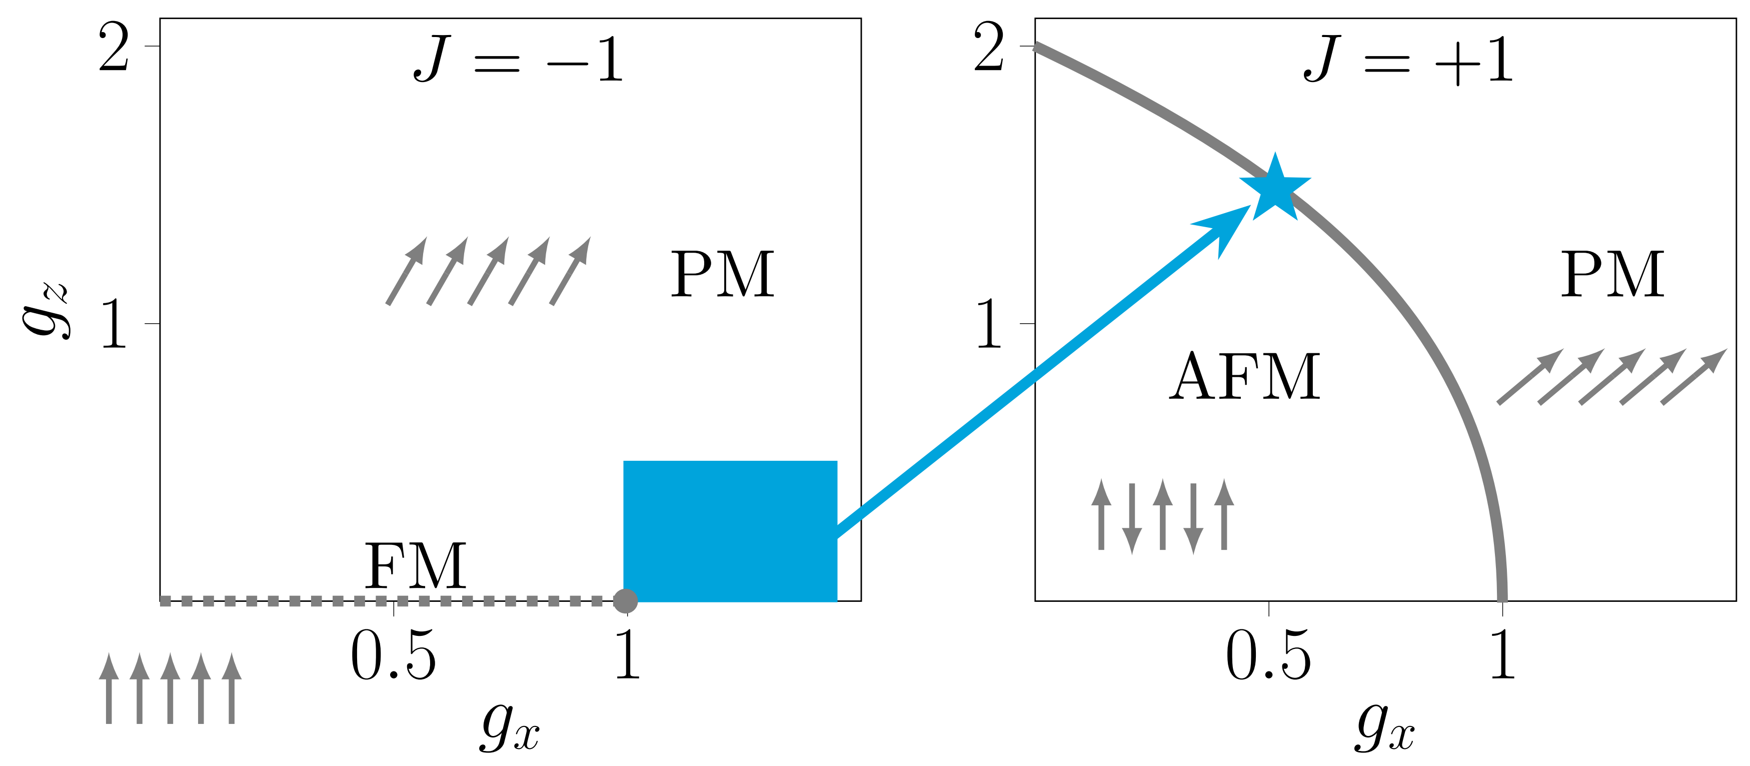

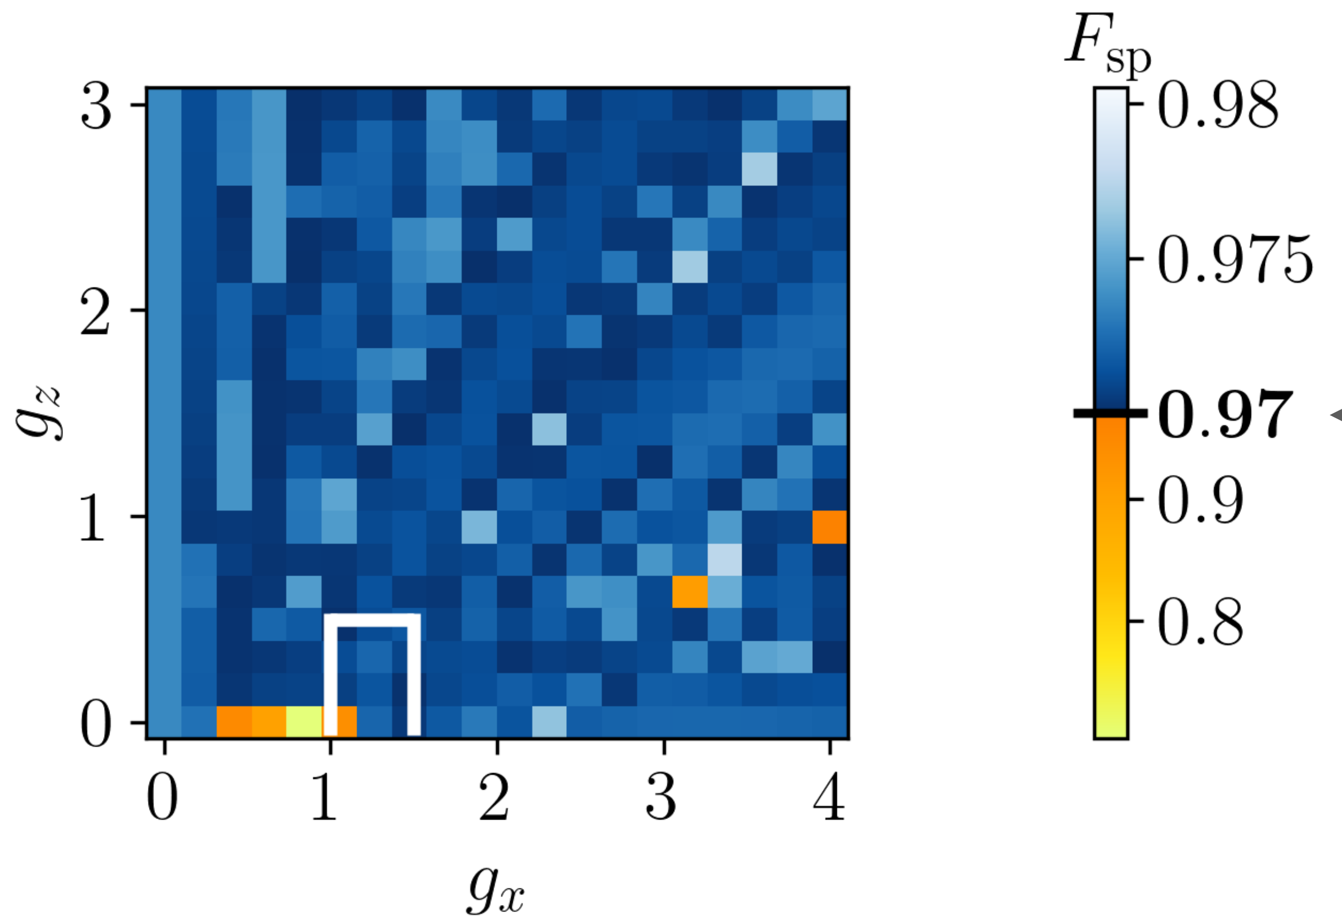

Fidelity threshold

ROBUST CRITICAL STATE PREPARATION

Mixed-field Ising: $\hat{H}_{\text{Ising}} = J \sum_{i=1}^{N-1} \hat{Z}_i \hat{Z}_{i+1} - g_x \sum_{i=1}^{N} \hat{X}_i - g_z \sum_{i=1}^{N} \hat{Z}_i$

$\rightarrow N = 16$ spins

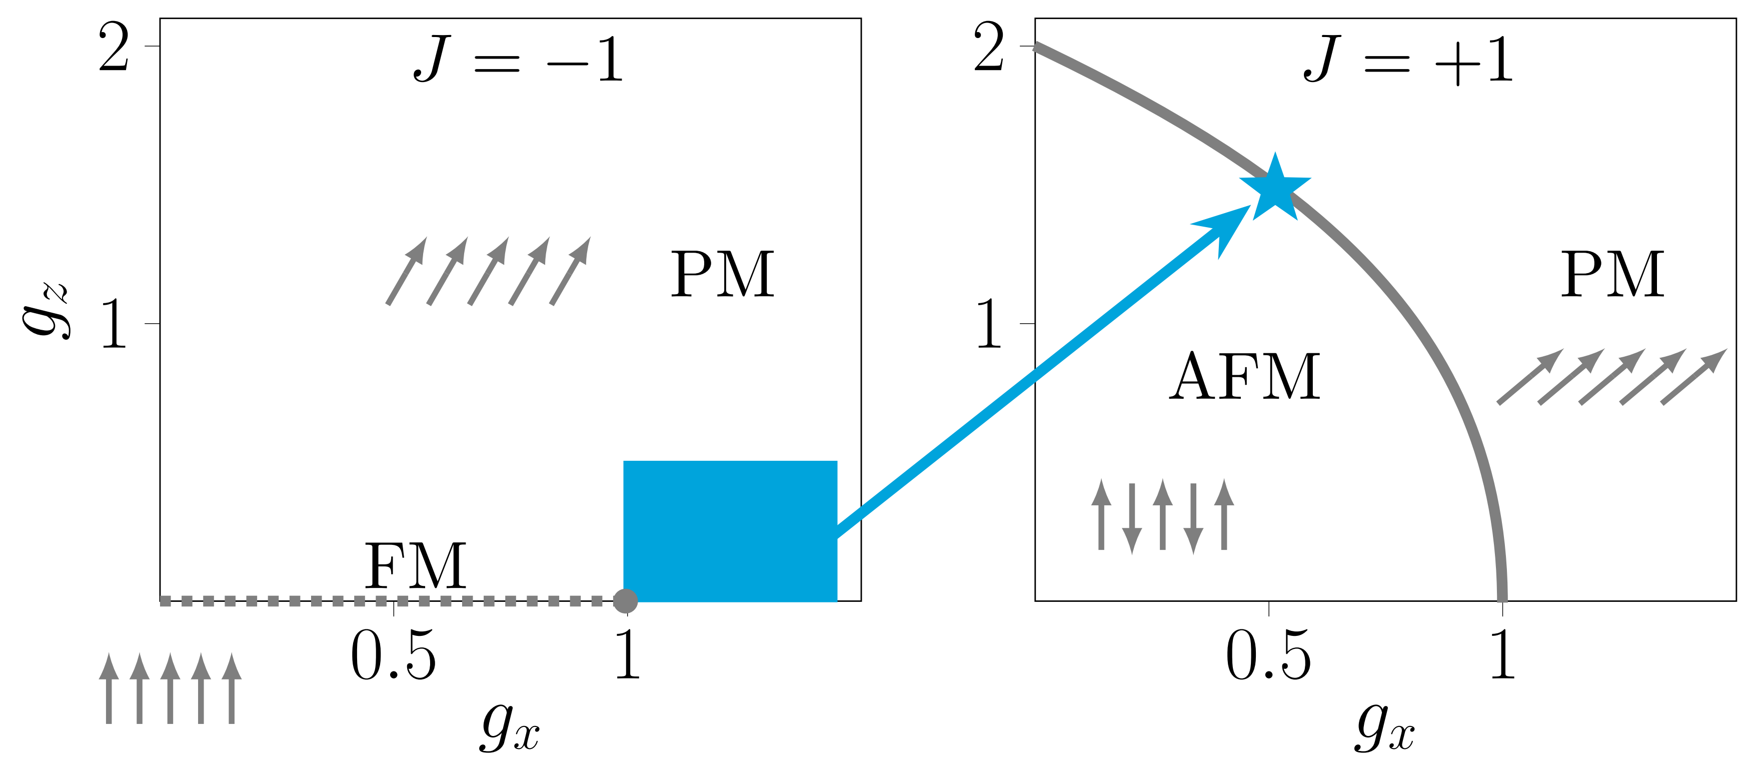

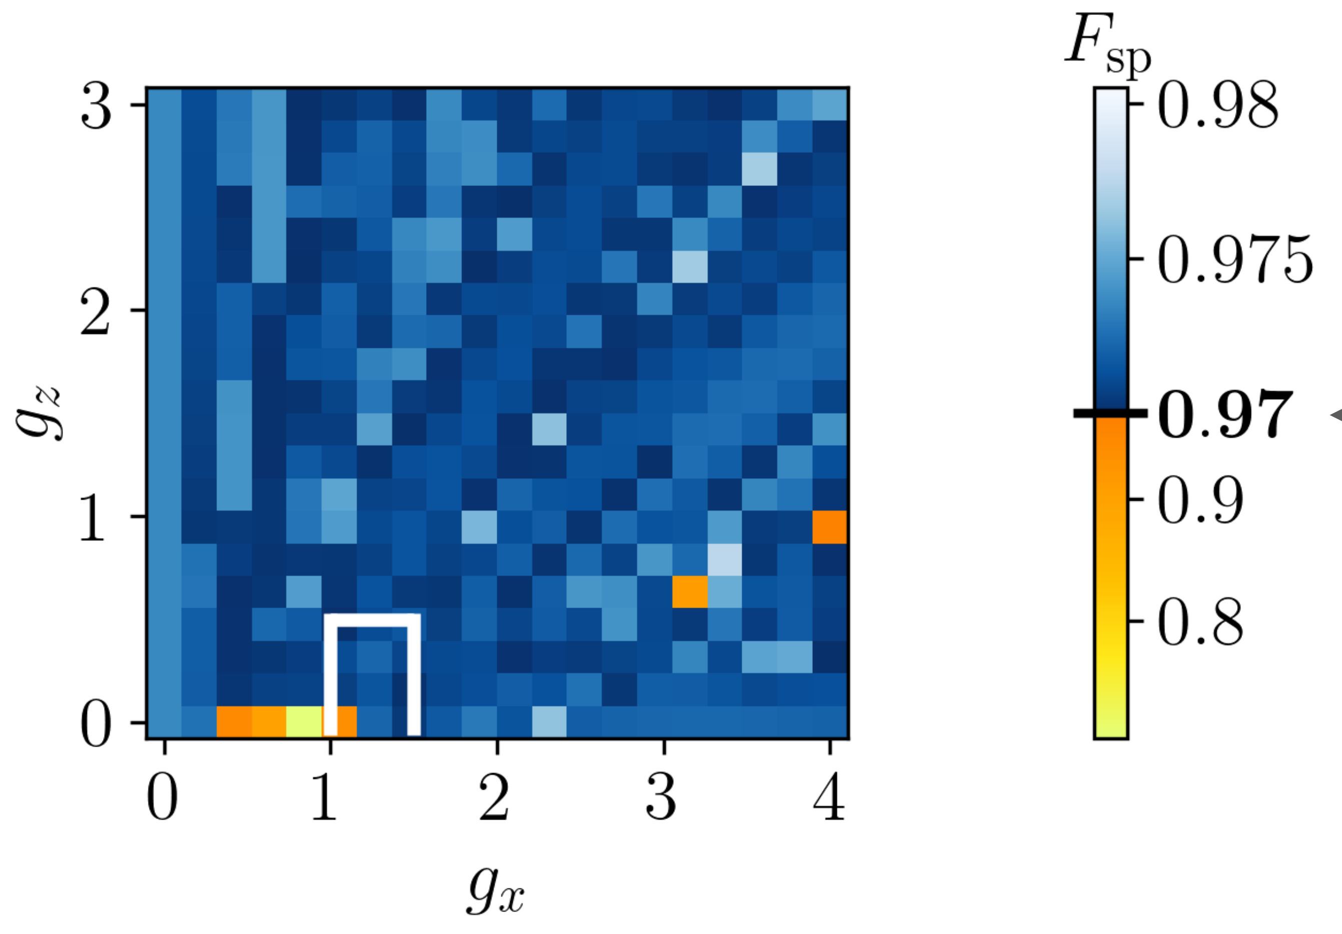

Fidelity threshold

QMPS agent can devise optimal protocols from various initial states

ROBUST CRITICAL STATE PREPARATION

Mixed-field Ising: $\hat{H}_{\text{Ising}} = J \sum_{i=1}^{N-1} \hat{Z}_i \hat{Z}_{i+1} - g_x \sum_{i=1}^{N} \hat{X}_i - g_z \sum_{i=1}^{N} \hat{Z}_i$

$\rightarrow N = 16$ spins

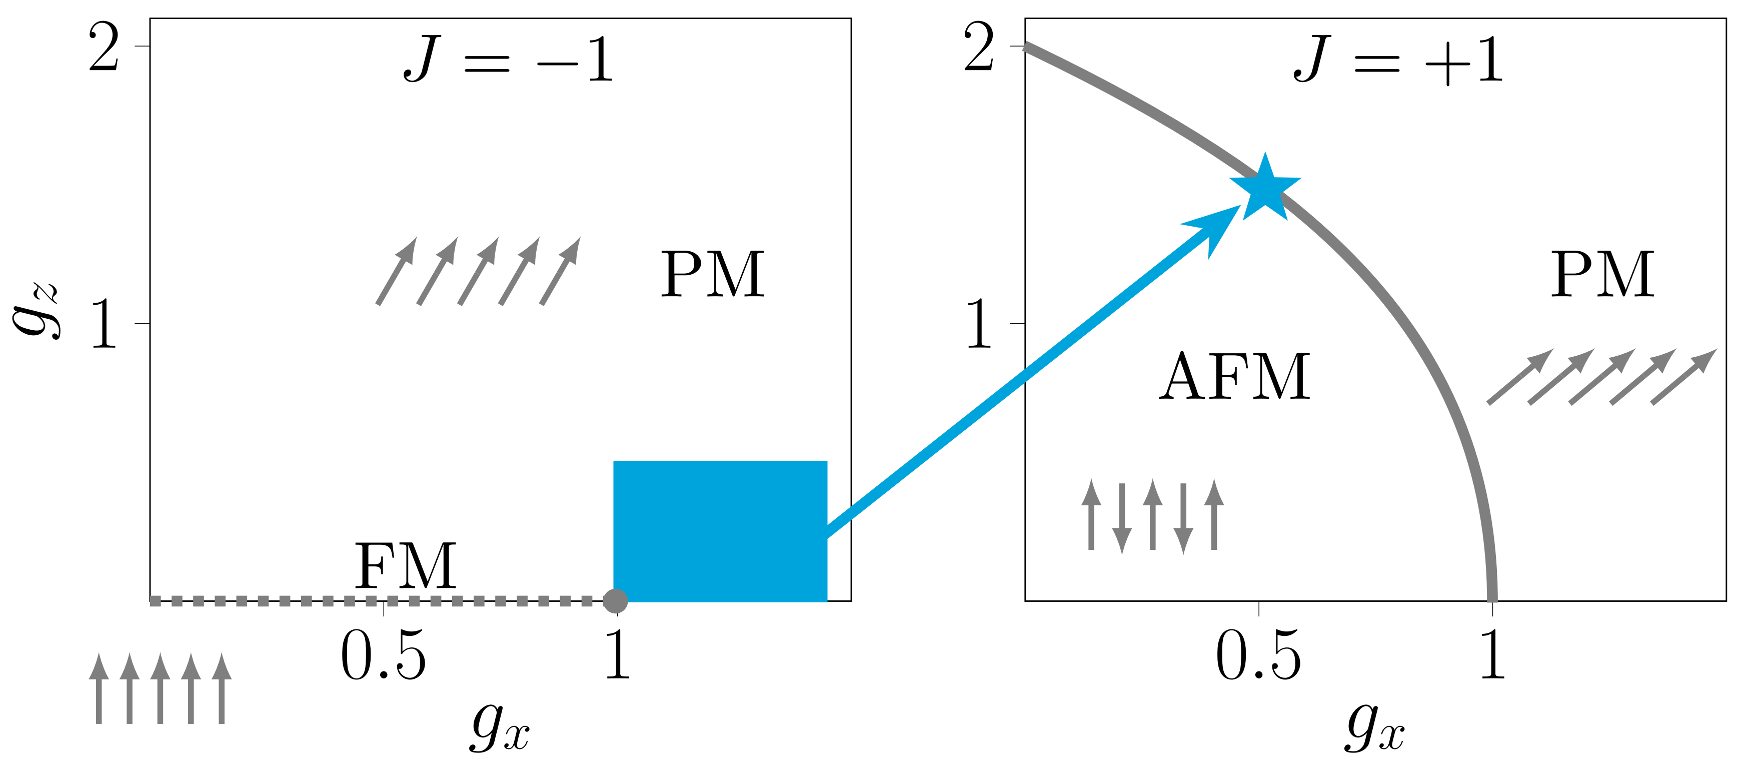

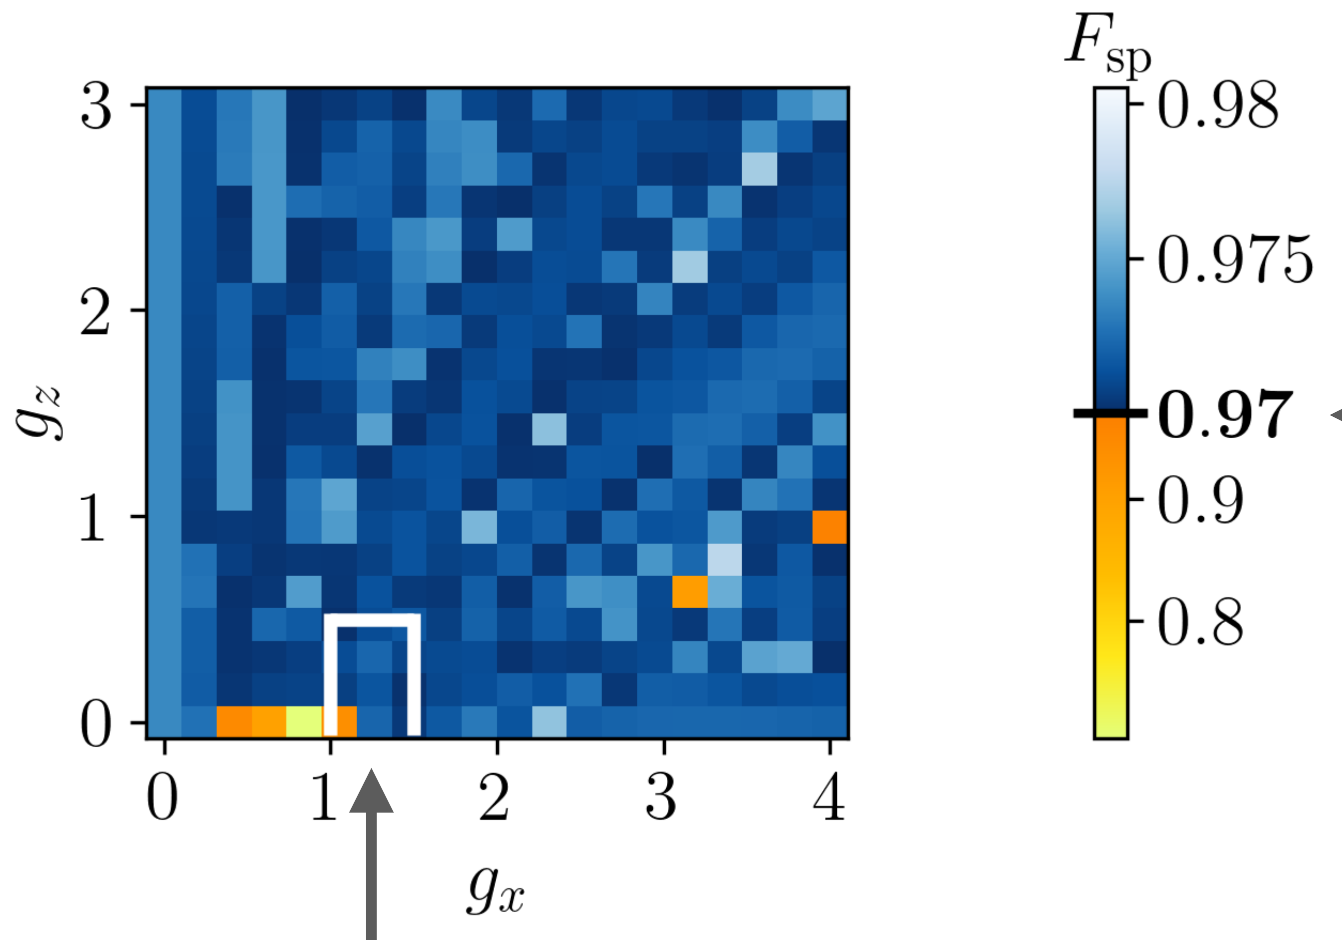

Fidelity threshold

Training region

ROBUST CRITICAL STATE PREPARATION

Mixed-field Ising: $\hat{H}_{\text{Ising}} = J\sum_{i=1}^{N-1}\hat{Z}_i\hat{Z}_{i+1} - g_x\sum_{i=1}^{N}\hat{X}_i - g_z\sum_{i=1}^{N}\hat{Z}_i$

$\rightarrow N = 16$ spins

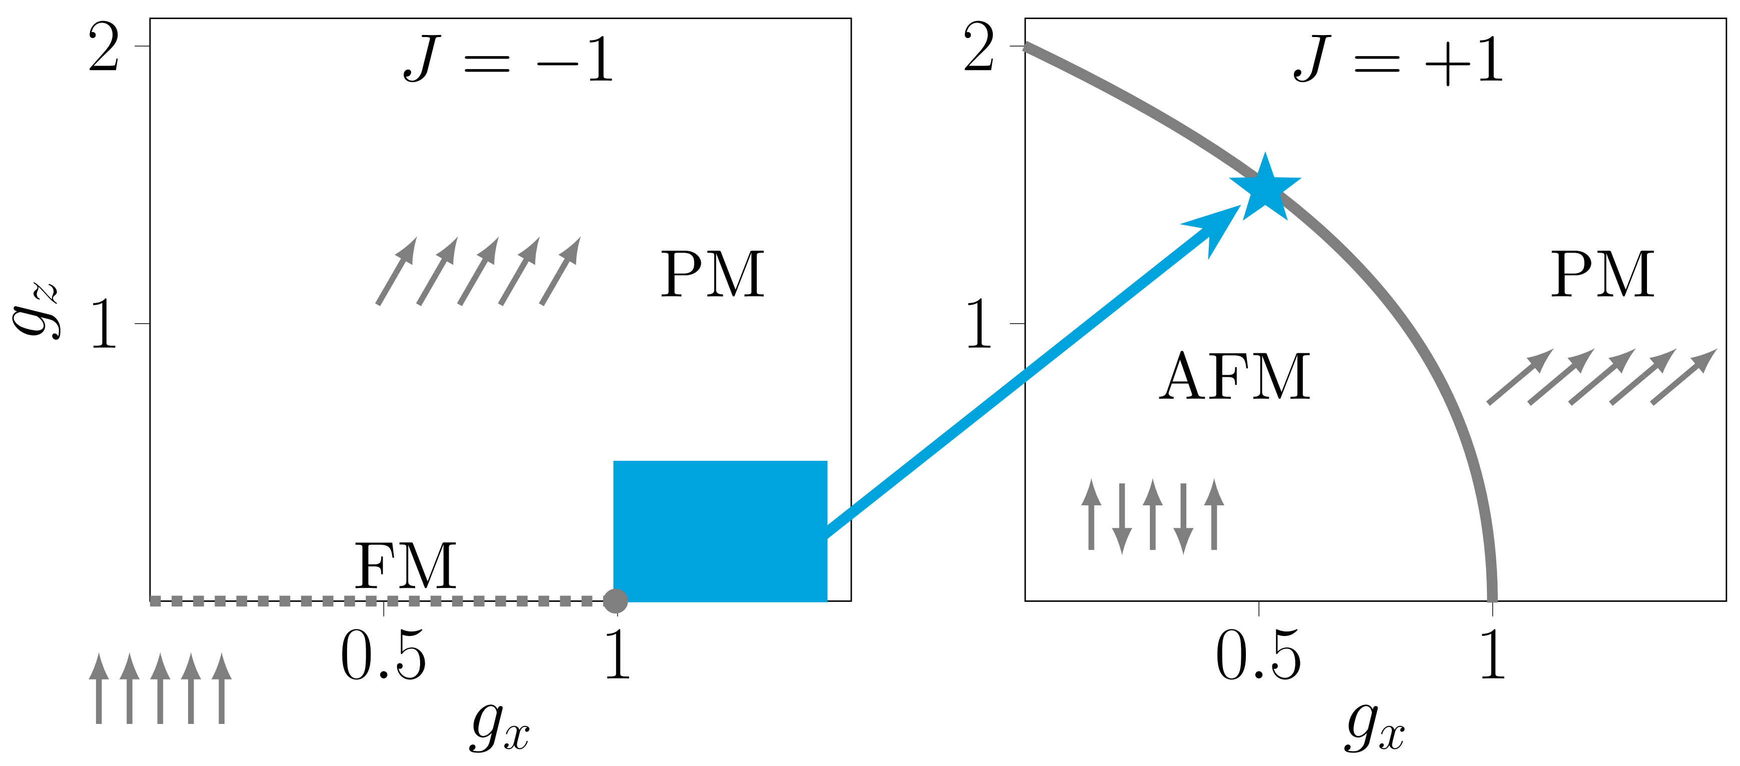

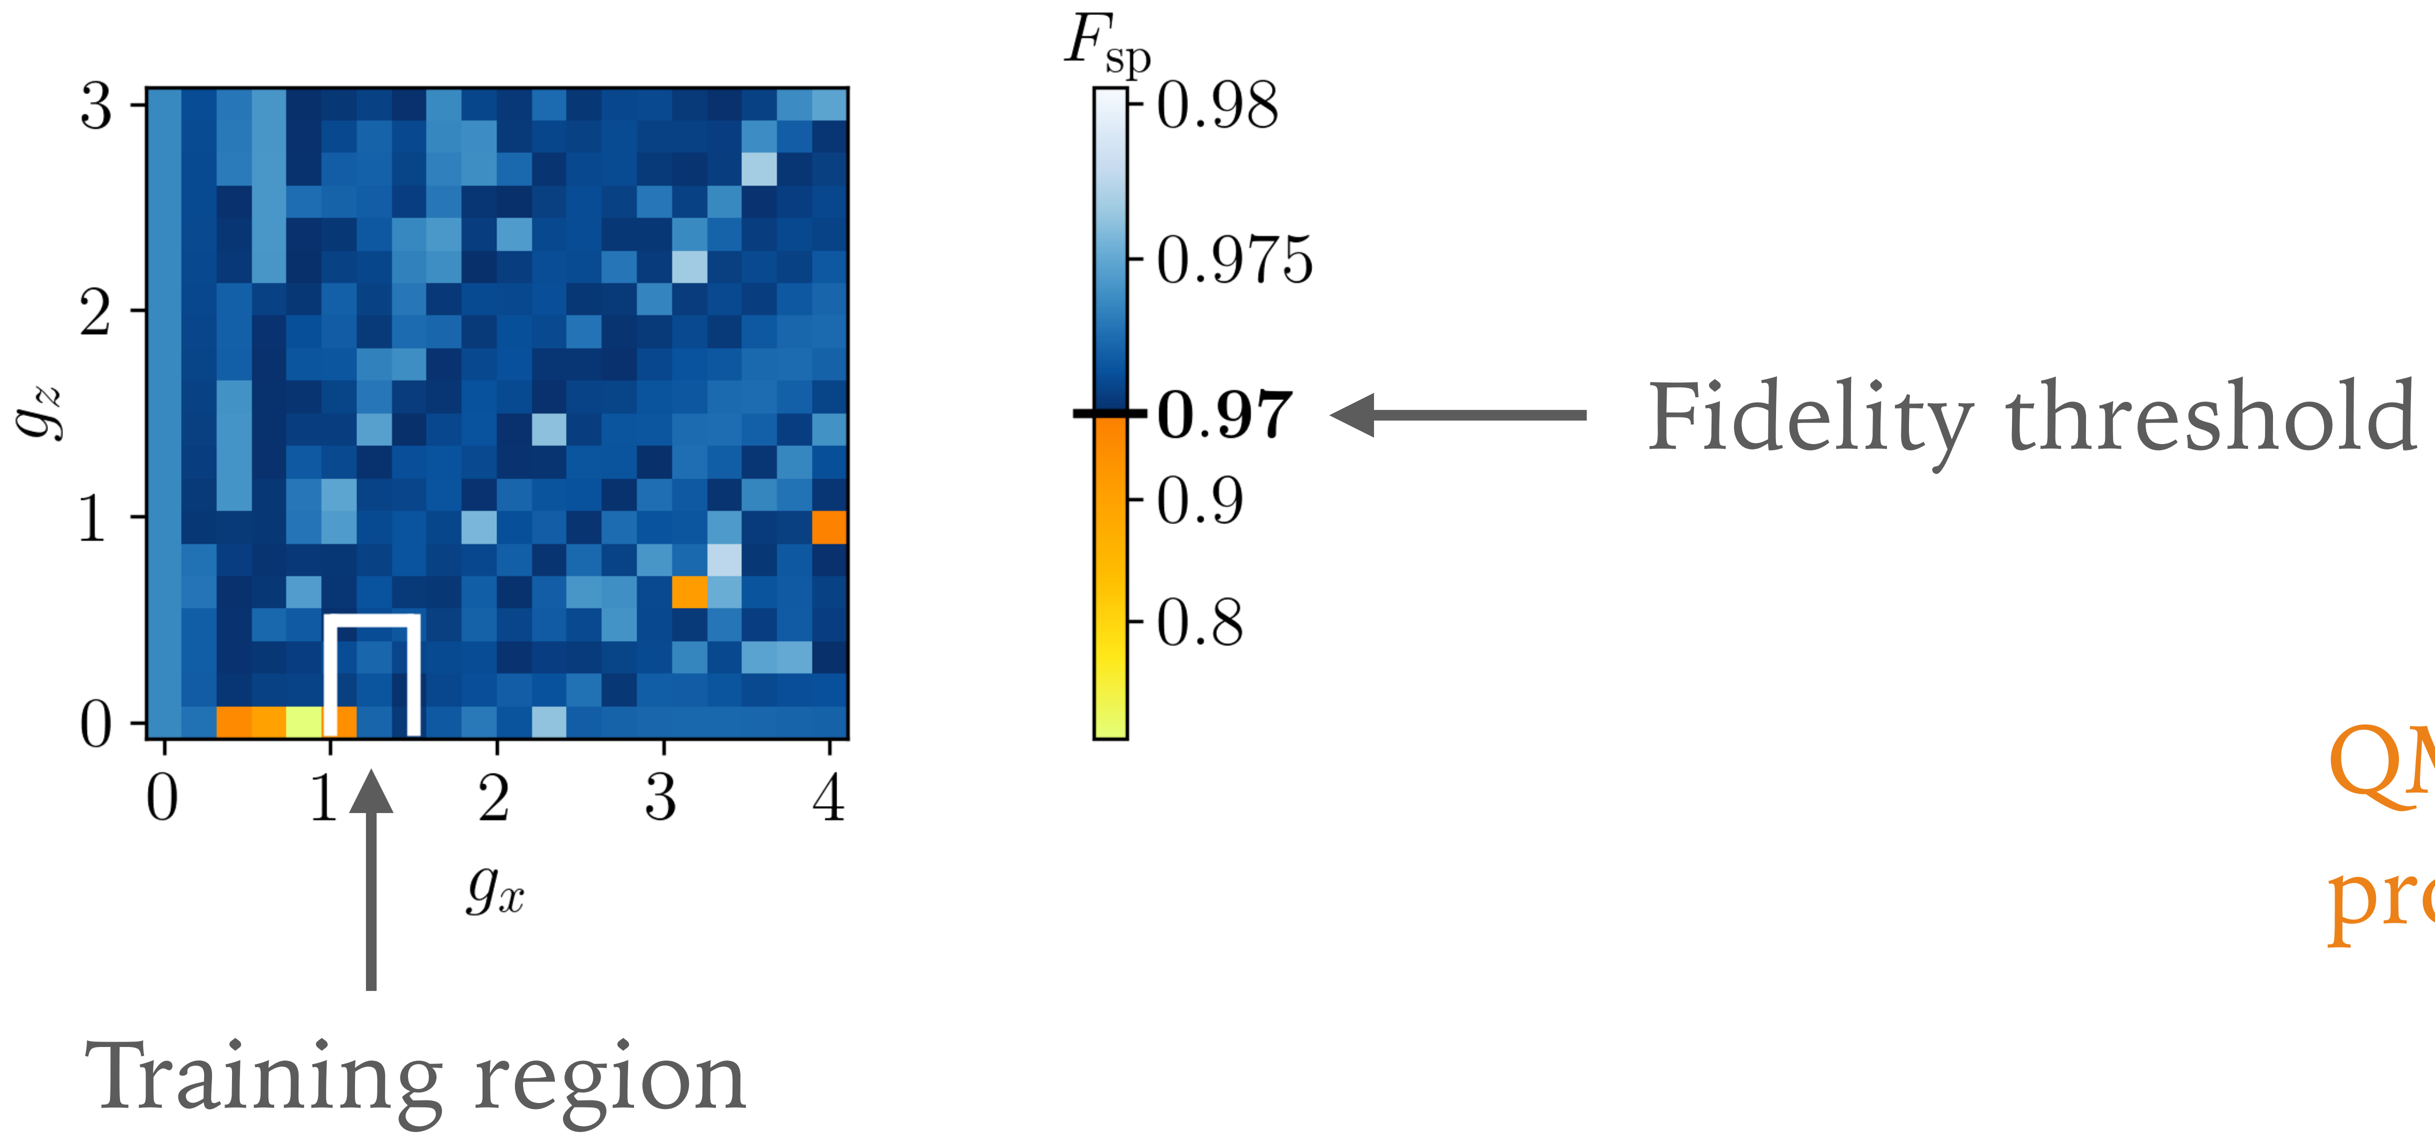

QMPS agent can extrapolate optimal protocols well beyond training region

ROBUST CRITICAL STATE PREPARATION

Mixed-field Ising: $\hat{H}_{\text{Ising}} = J\sum_{i=1}^{N-1}\hat{Z}_i\hat{Z}_{i+1} - g_x\sum_{i=1}^{N}\hat{X}_i - g_z\sum_{i=1}^{N}\hat{Z}_i$

$\rightarrow N = 16$ spins

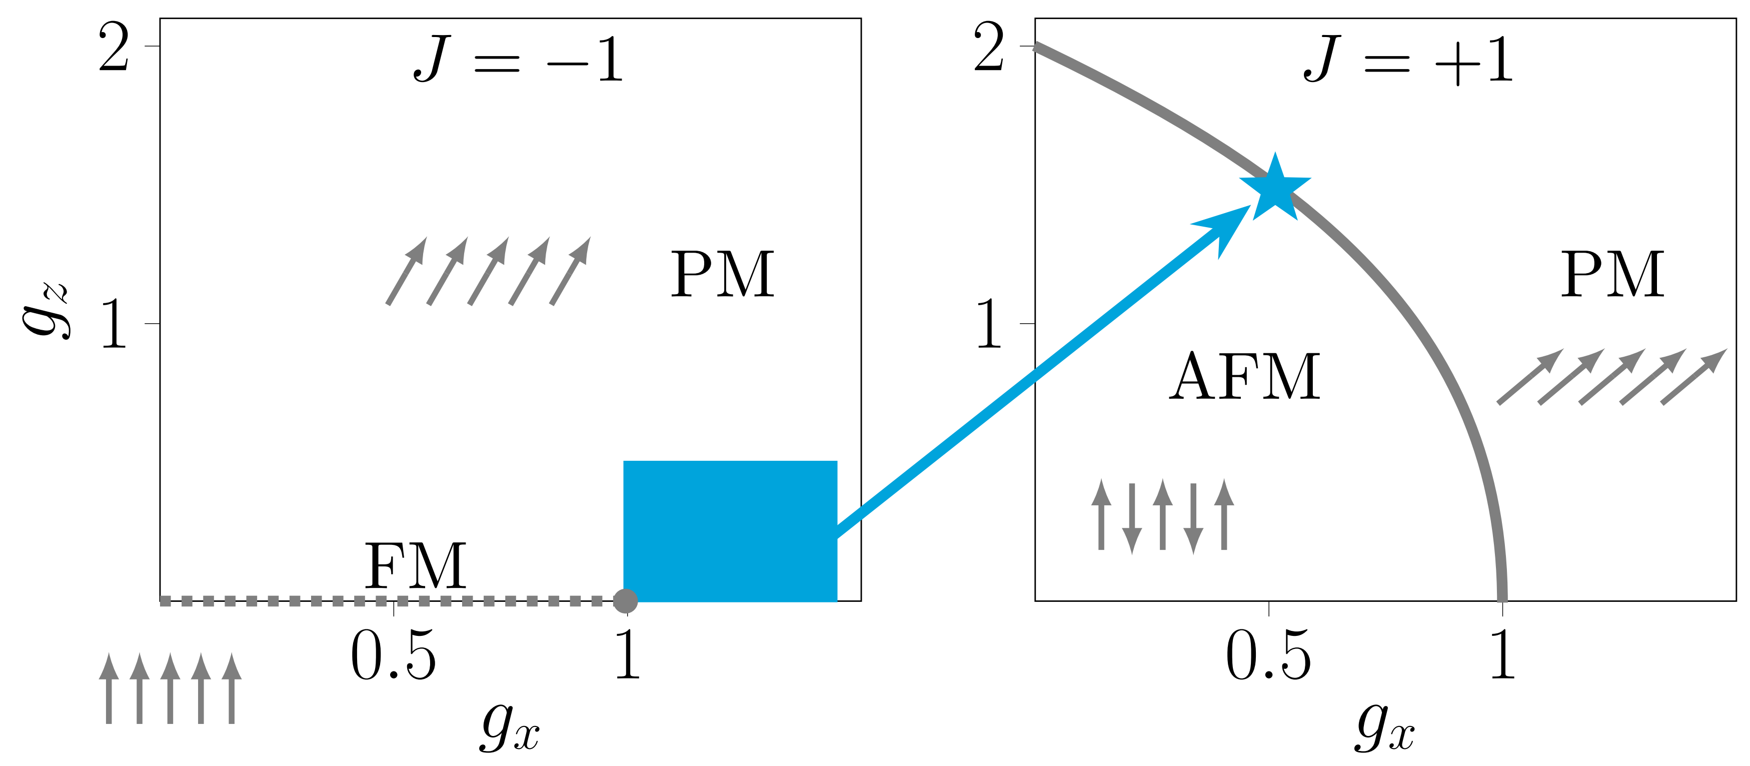

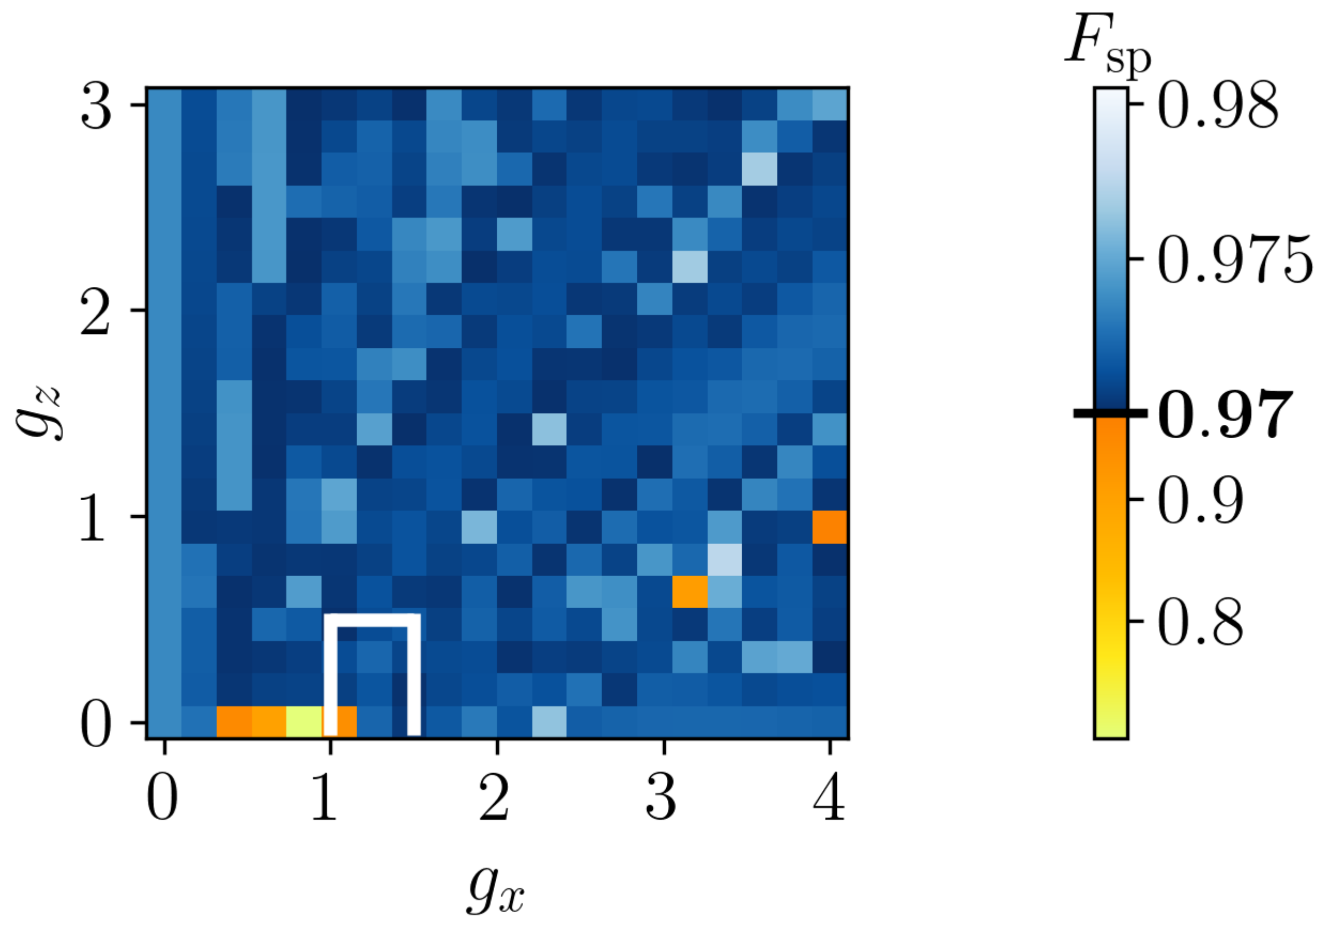

Noise: At each step, white Gaussian noise with std σ is added to step duration δt $\quad (e^{\pm i\delta t_{\pm}\hat{A}})$

ROBUST CRITICAL STATE PREPARATION

Mixed-field Ising: $\hat{H}_{\text{Ising}} = J \sum_{i=1}^{N-1} \hat{Z}_i \hat{Z}_{i+1} - g_x \sum_{i=1}^{N} \hat{X}_i - g_z \sum_{i=1}^{N} \hat{Z}_i$

$\rightarrow N = 16$ spins

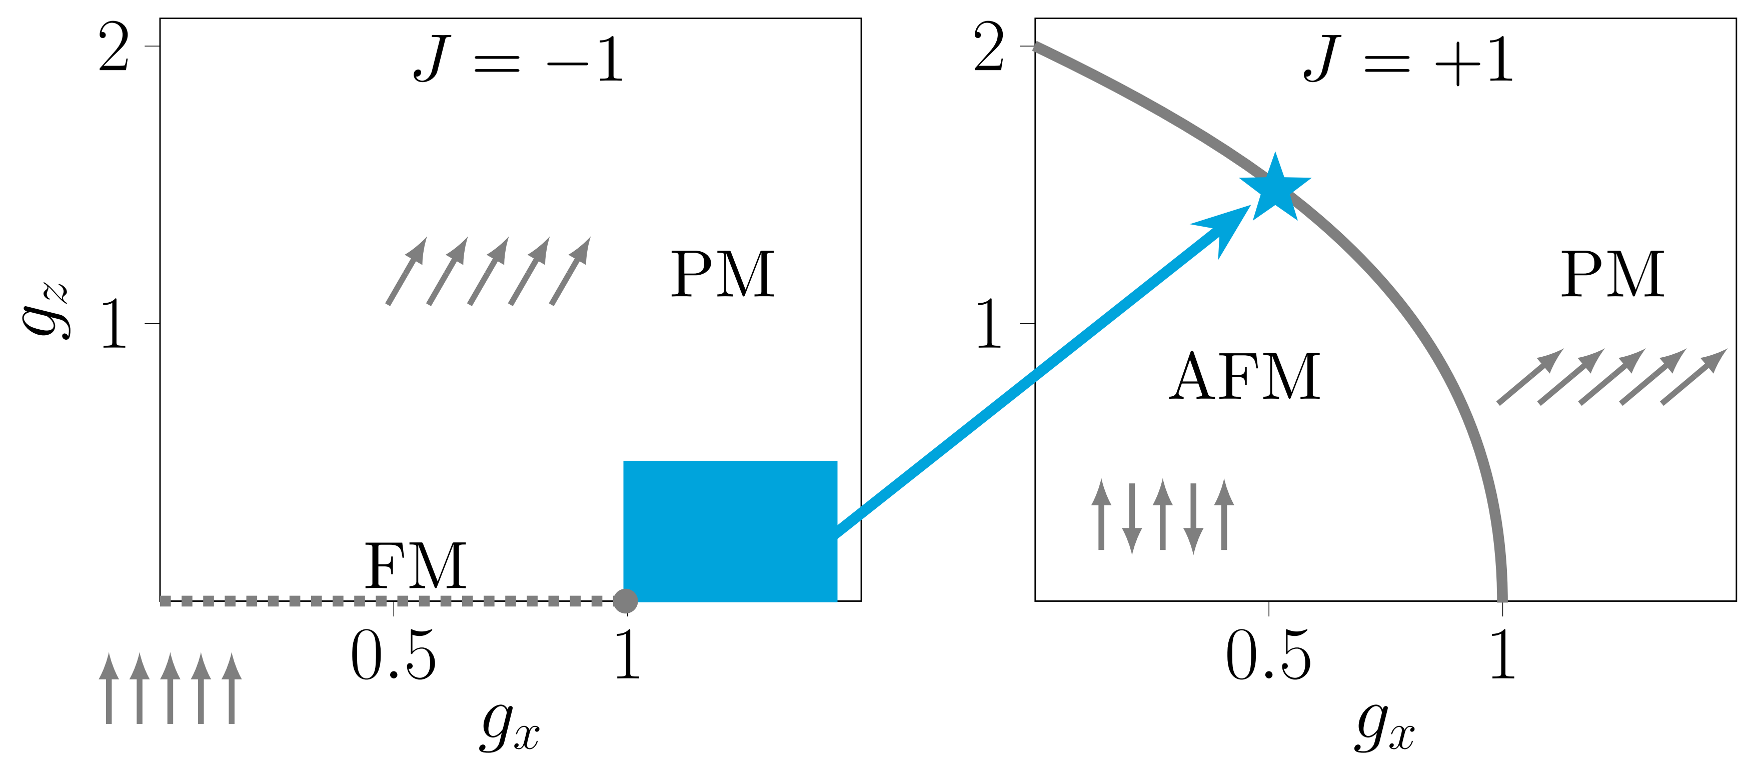

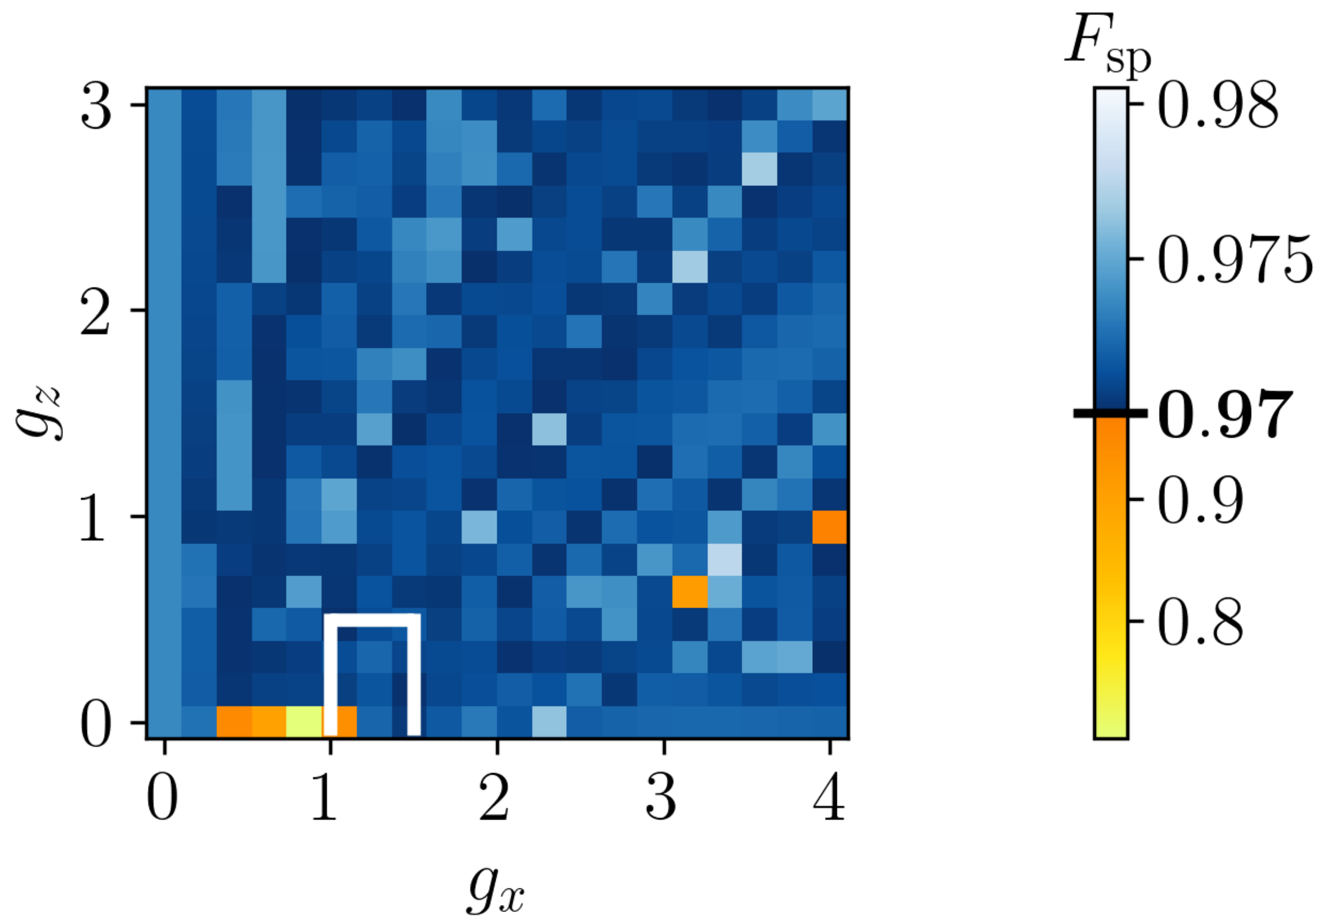

Noise: At each step, white Gaussian noise with std σ is added to step duration δt $(e^{\pm i \delta t_{\pm} \hat{A}})$

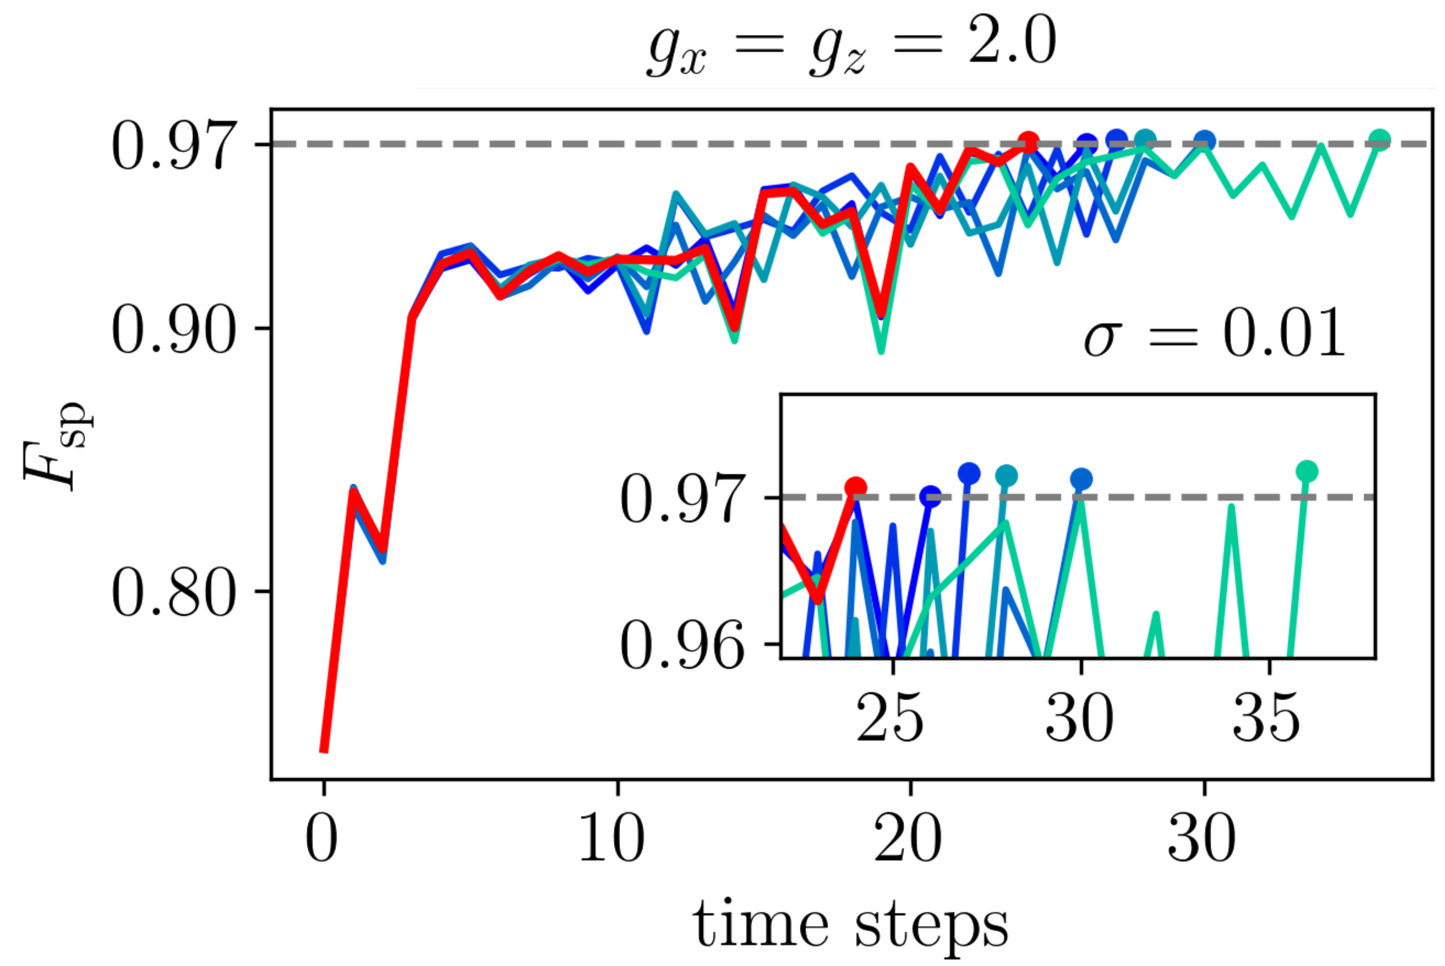

ROBUST CRITICAL STATE PREPARATION

Mixed-field Ising: $\hat{H}_{\text{Ising}} = J\sum_{i=1}^{N-1}\hat{Z}_i\hat{Z}_{i+1} - g_x\sum_{i=1}^{N}\hat{X}_i - g_z\sum_{i=1}^{N}\hat{Z}_i$

$\rightarrow N = 16$ spins

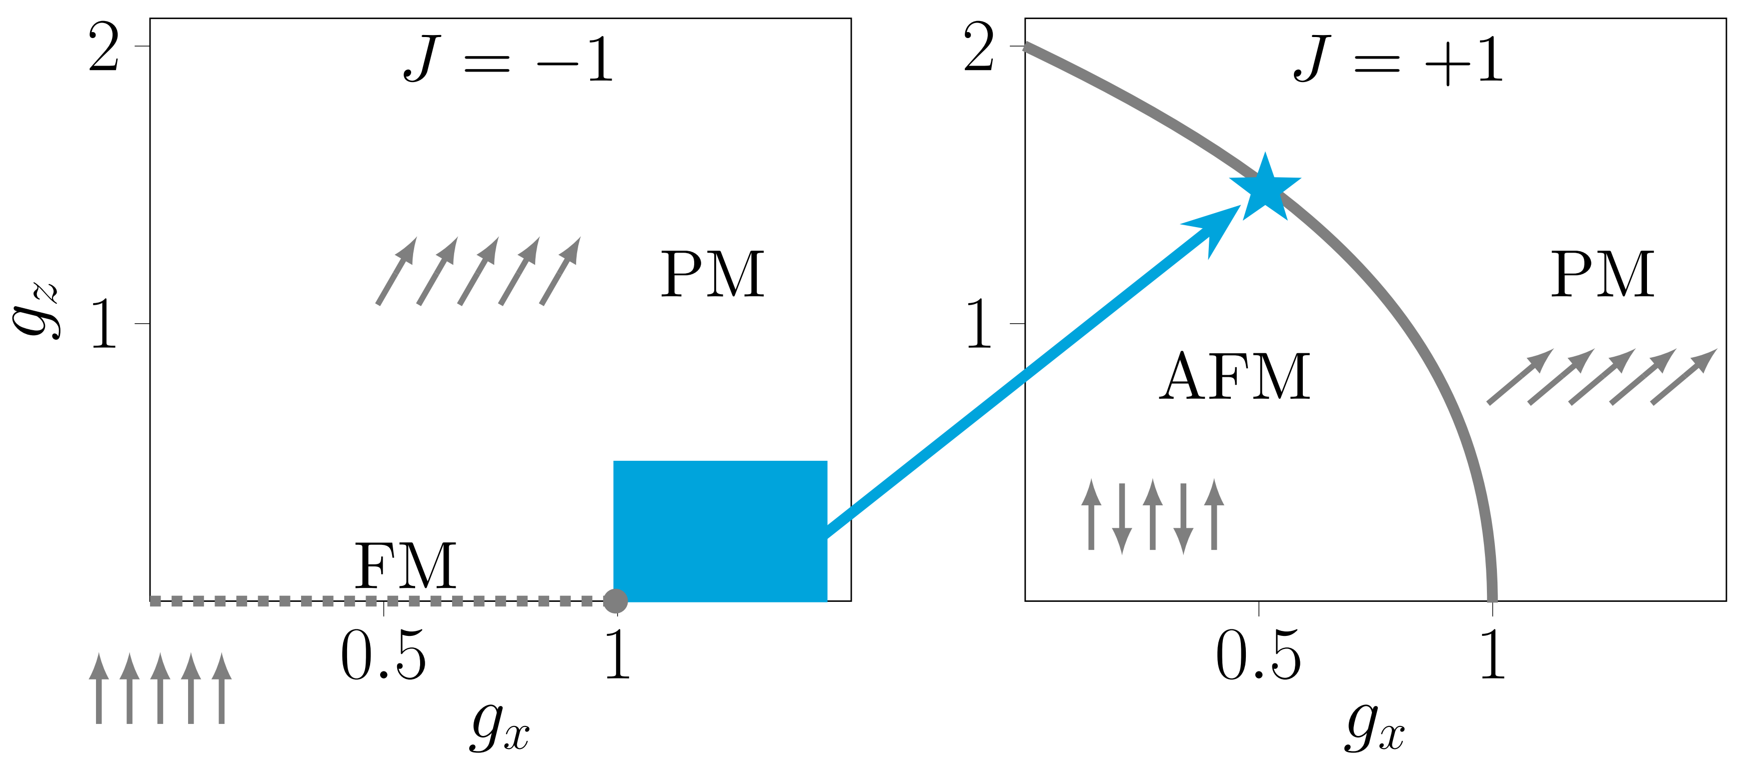

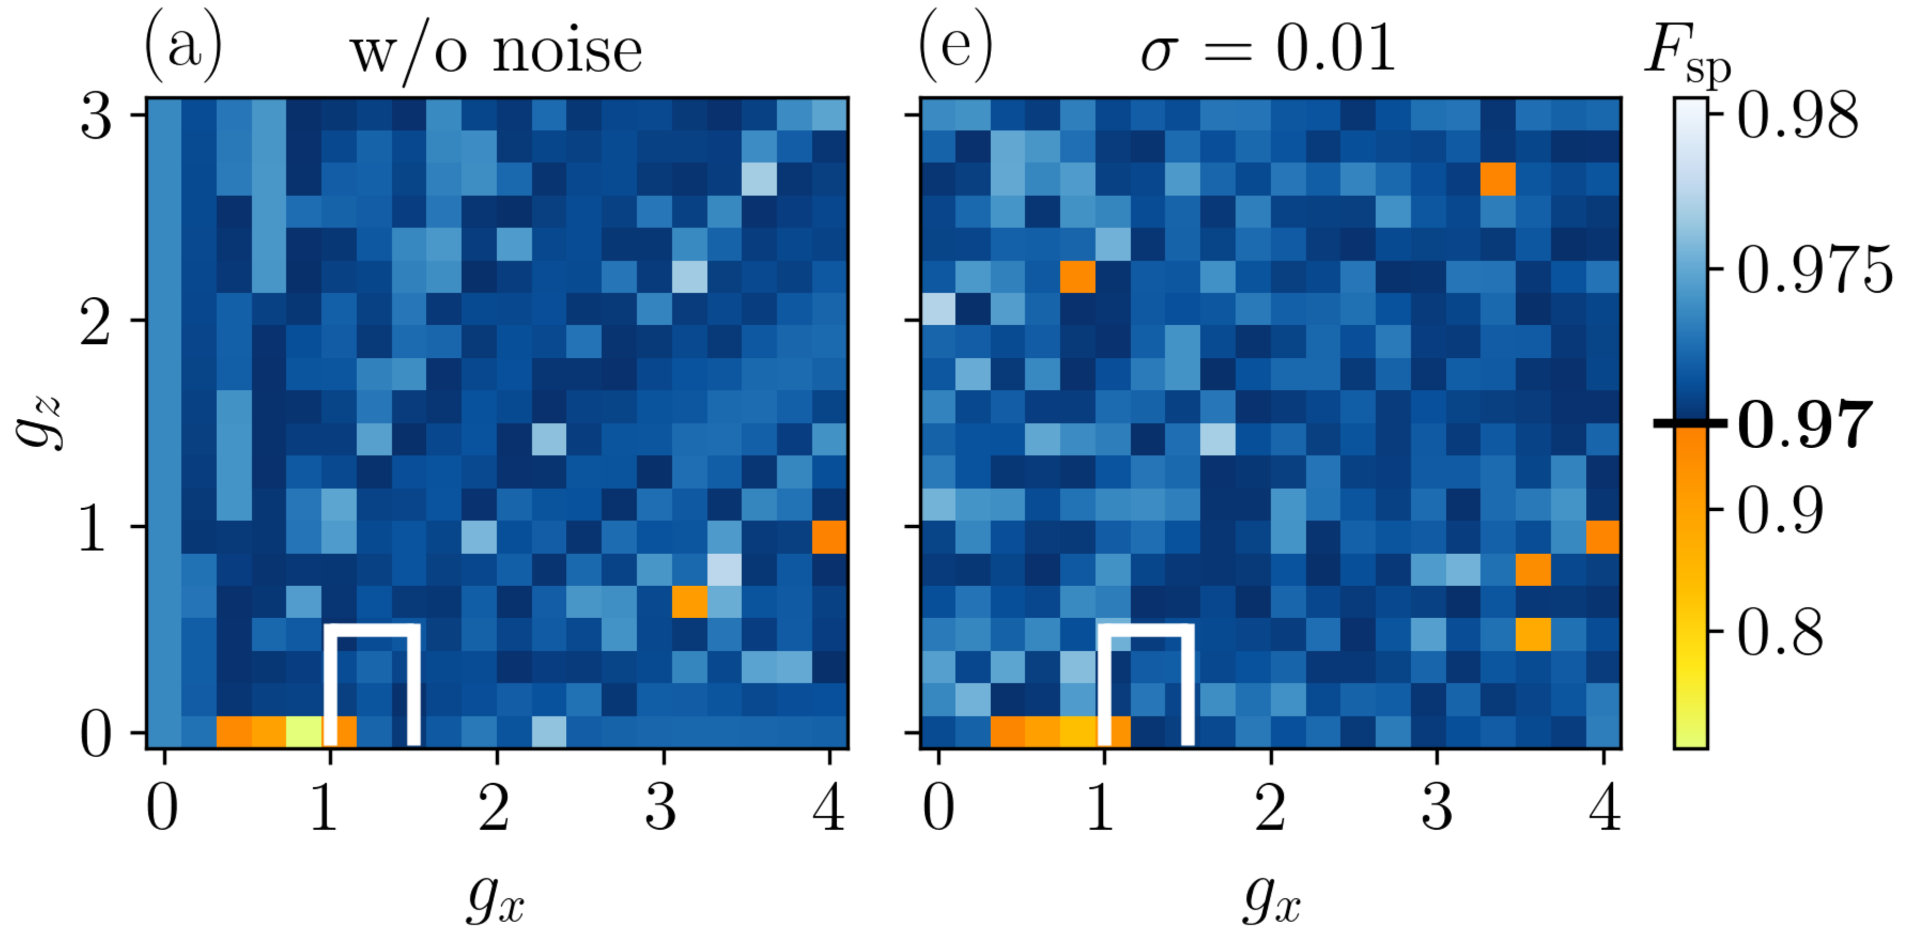

Noise: At each step, white Gaussian noise with std σ is added to step duration δt $(e^{\pm i\delta t_\pm \hat{A}})$

ROBUST CRITICAL STATE PREPARATION

Mixed-field Ising: $\hat{H}_{\text{Ising}} = J\sum_{i=1}^{N-1}\hat{Z}_i\hat{Z}_{i+1} - g_x\sum_{i=1}^{N}\hat{X}_i - g_z\sum_{i=1}^{N}\hat{Z}_i$

$\rightarrow N = 16$ spins

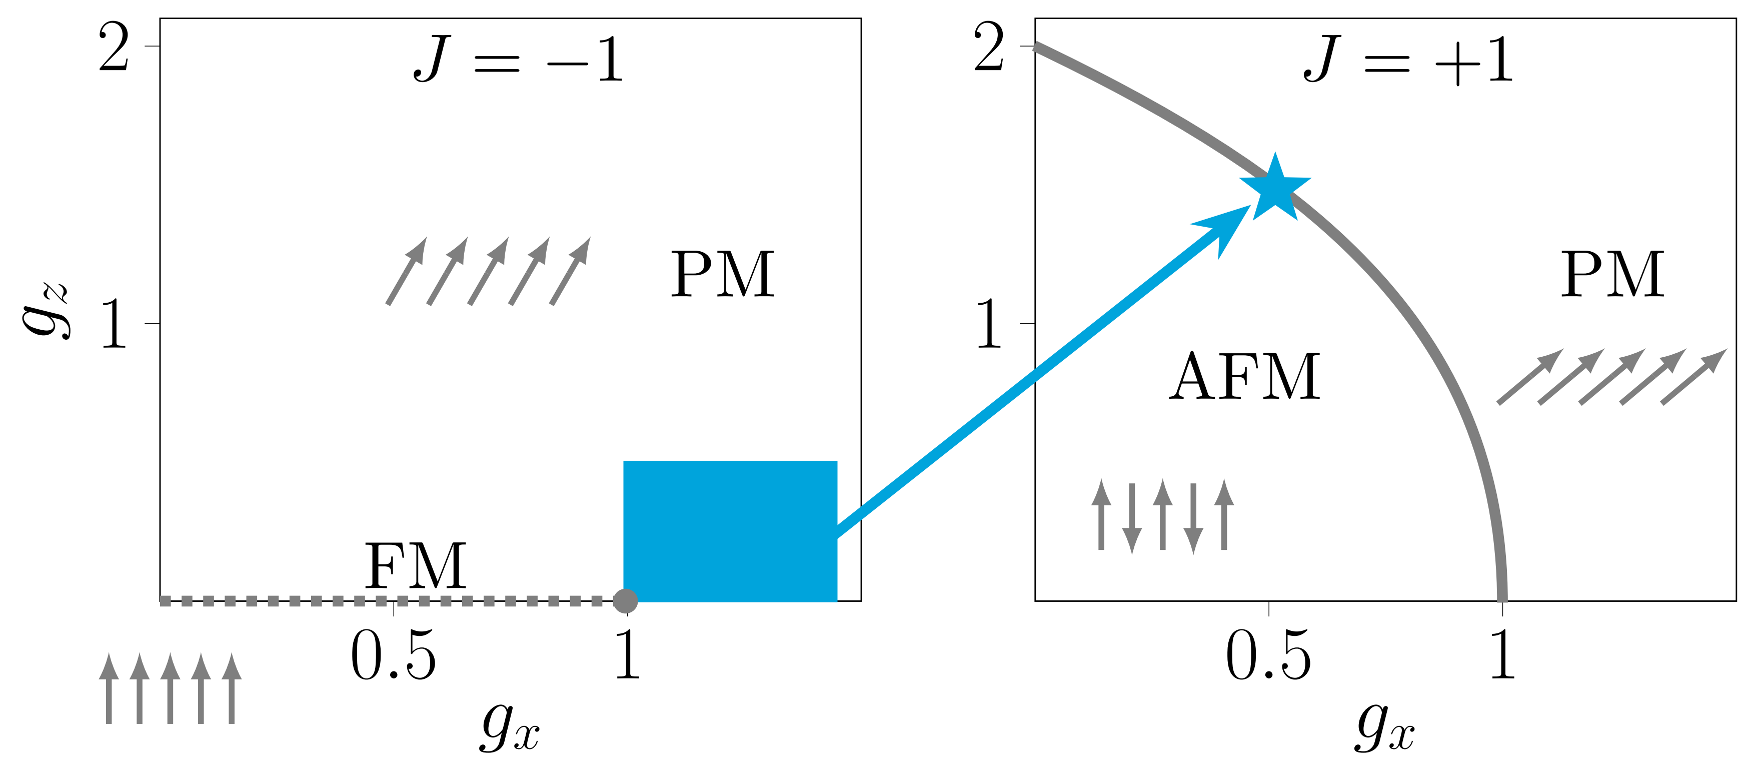

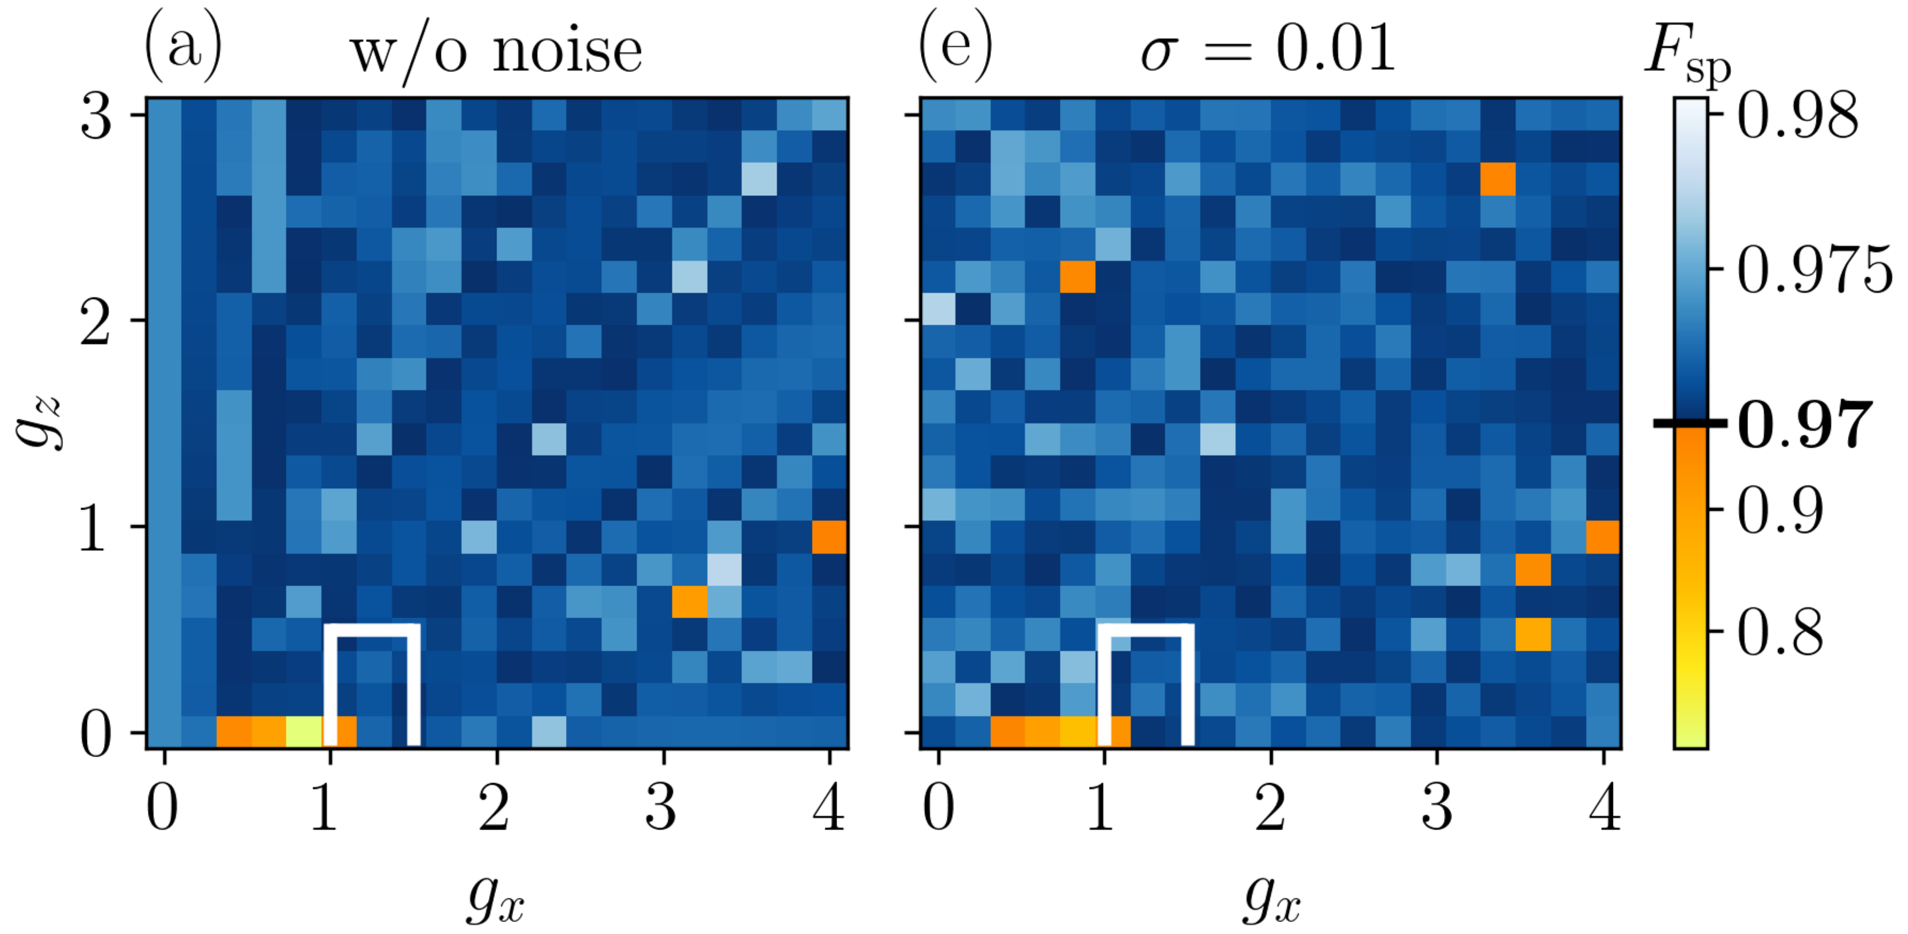

Noise: At each step, white Gaussian noise with std σ is added to step duration δt $\quad(e^{\pm i\delta t_{\pm}\hat{A}})$

QMPS agent can self-correct protocols on-the-fly

MORE EXAMPLES

Universal state preparation

→ Learn to prepare target from arbitrary initial quantum state (N=4)

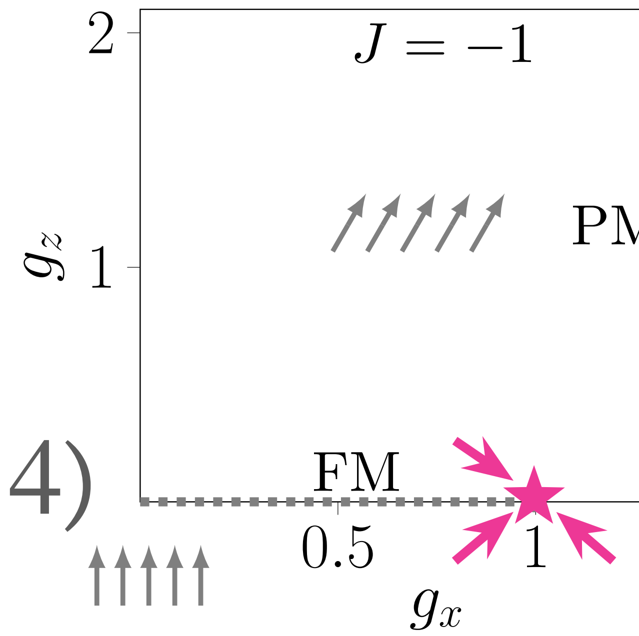

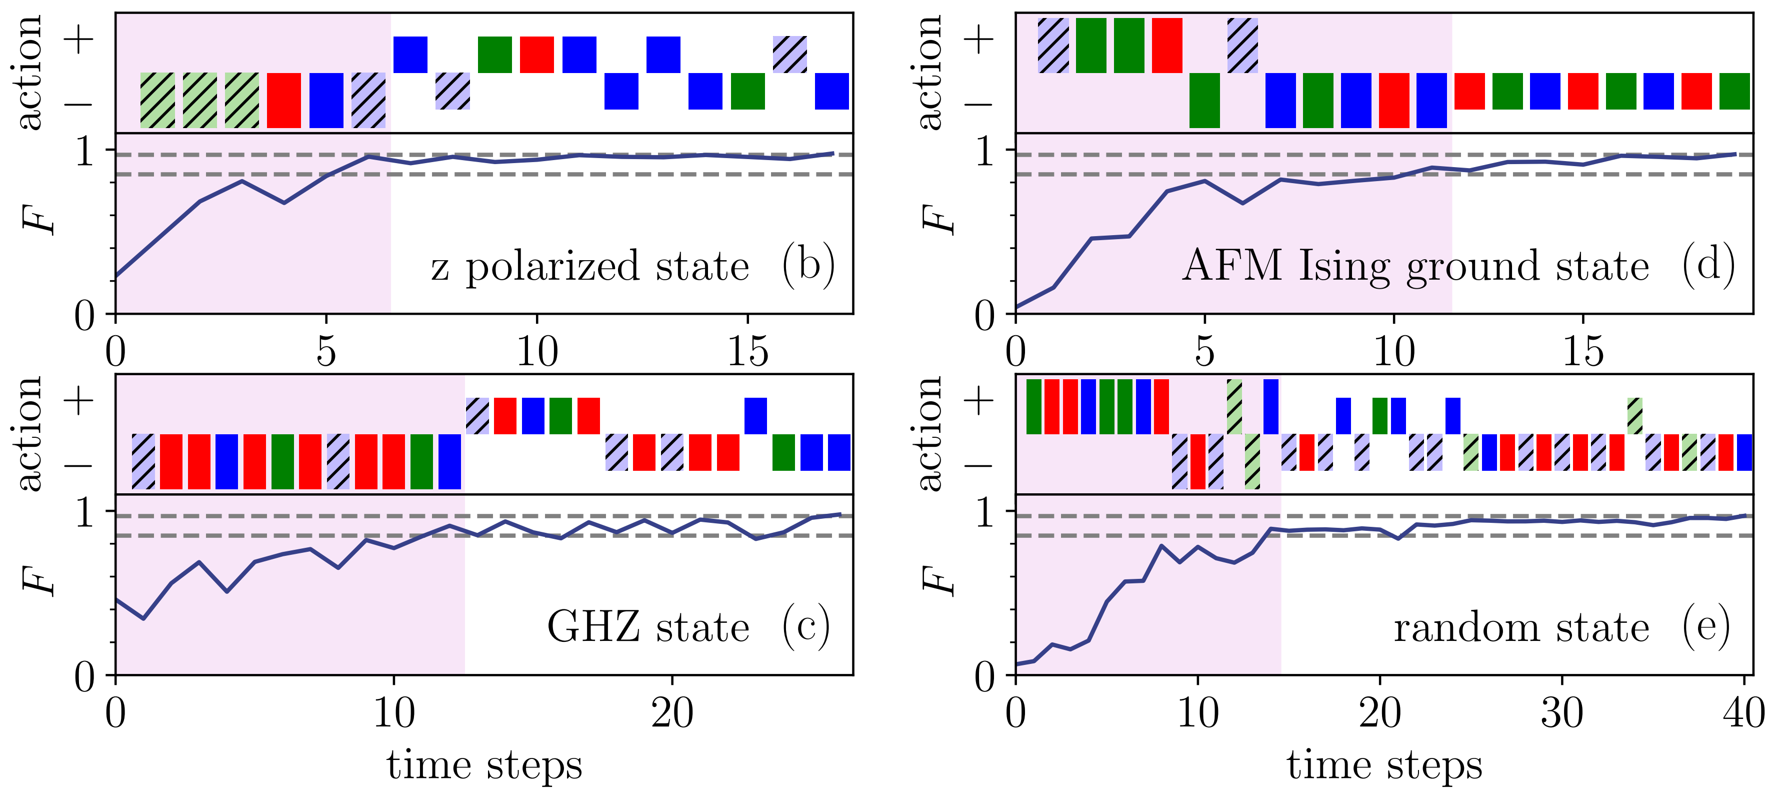

Many-body state preparation

→ $N = 32$ spins (beyond ED)

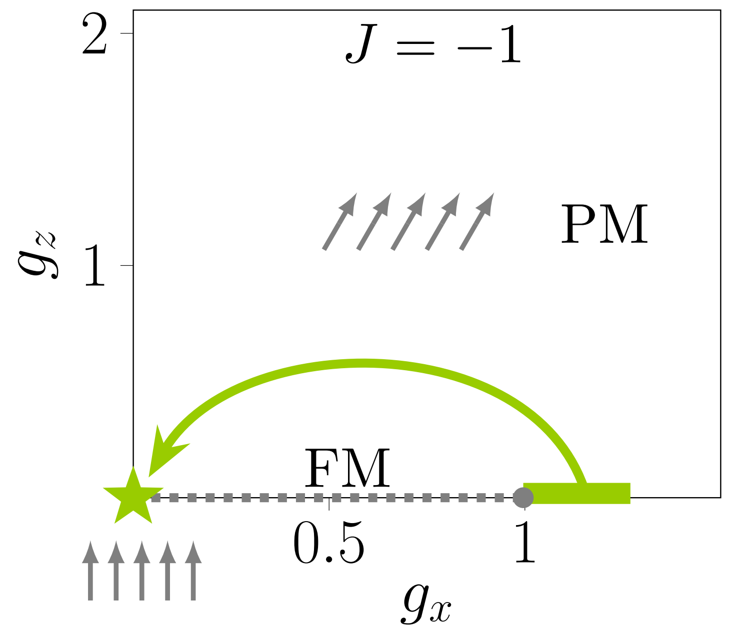

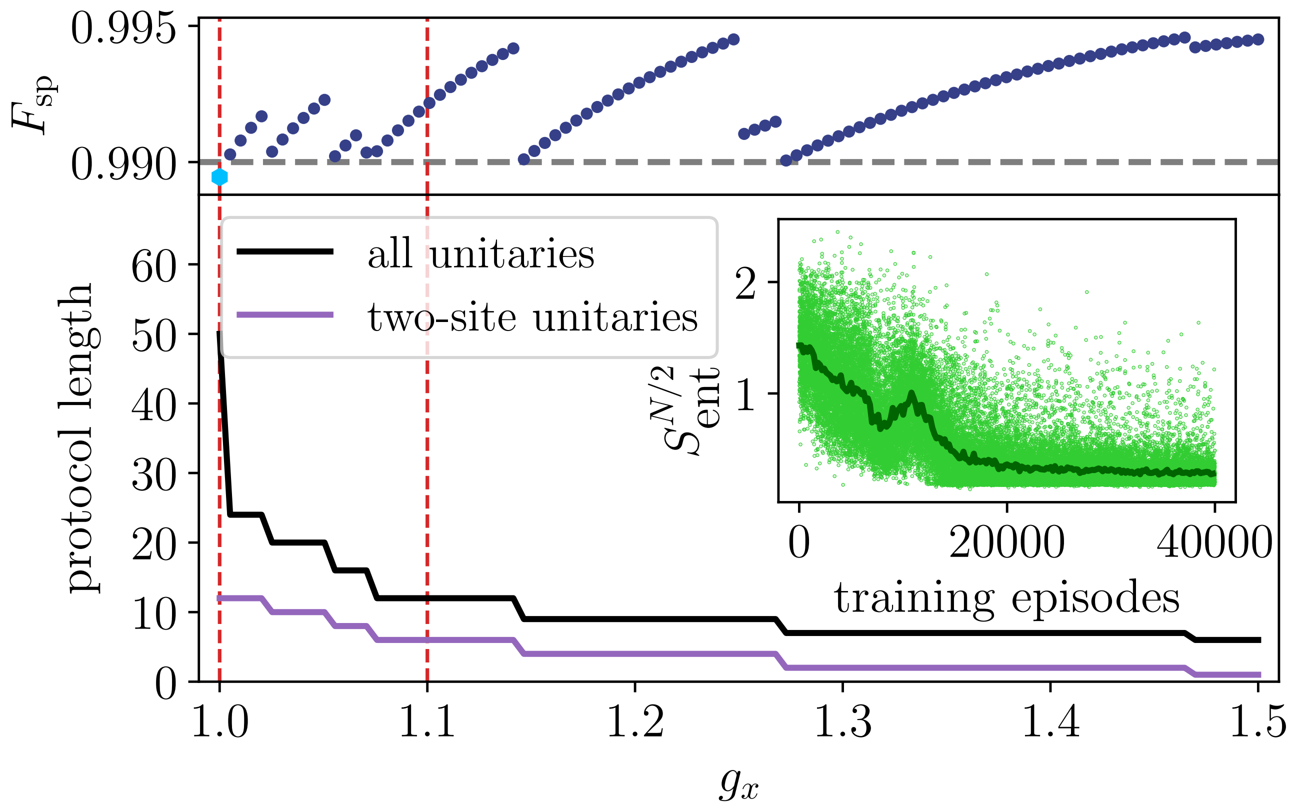

MORE EXAMPLES

Universal state preparation

→ Learn to prepare target from arbitrary initial quantum state (N=4)

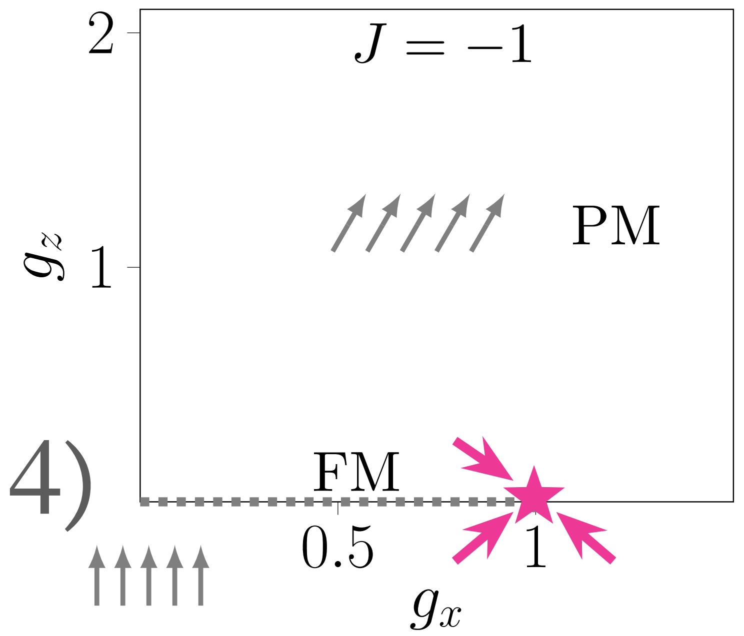

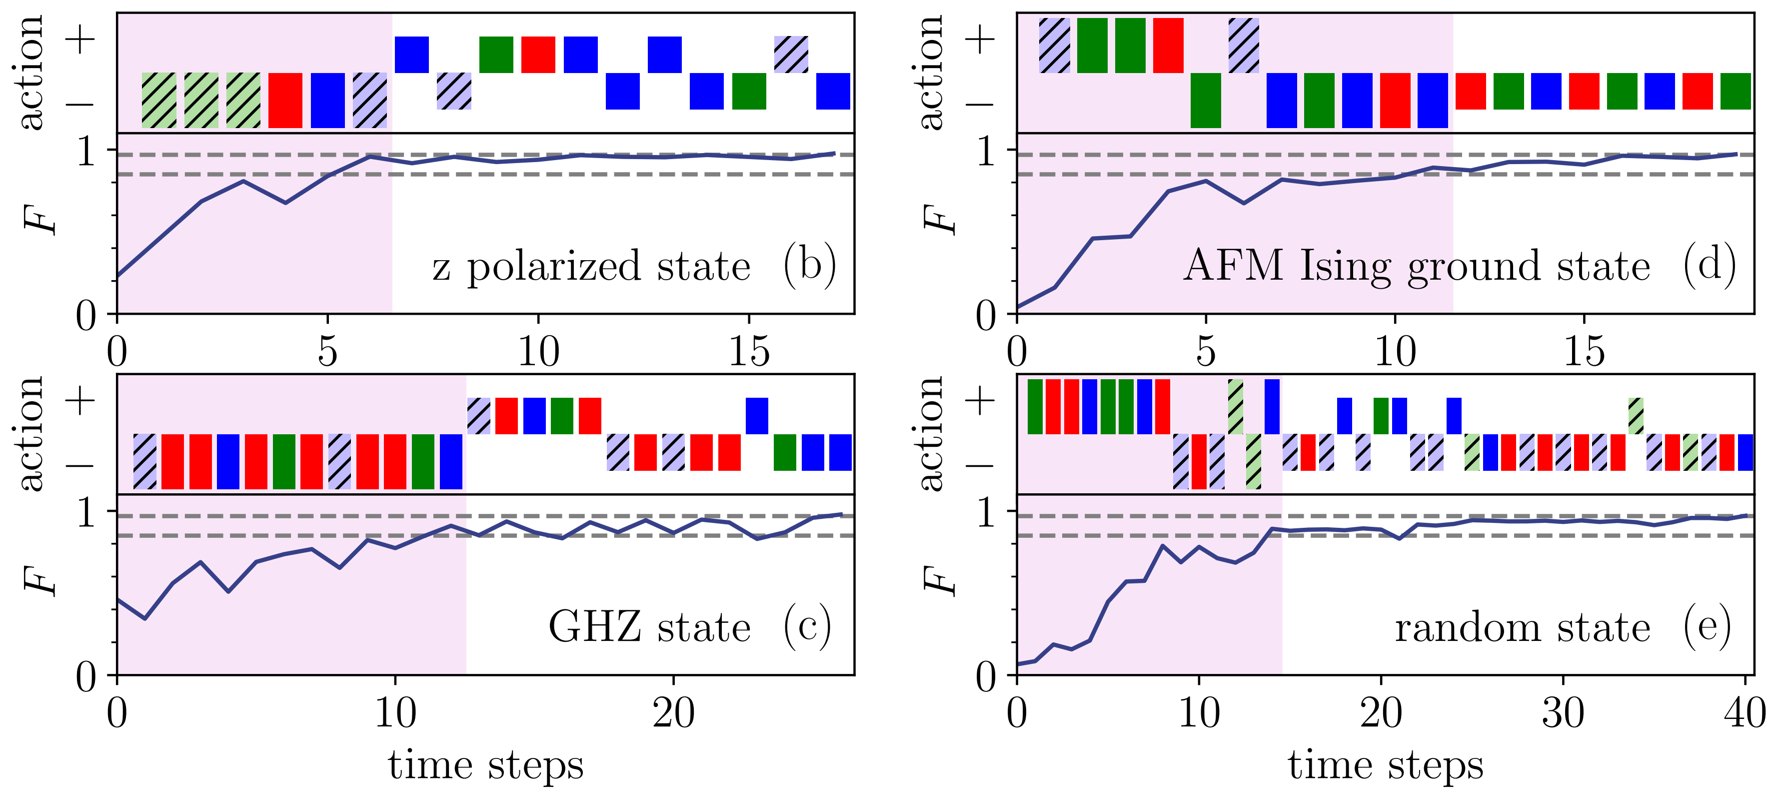

Many-body state preparation

→ $N = 32$ spins (beyond ED)

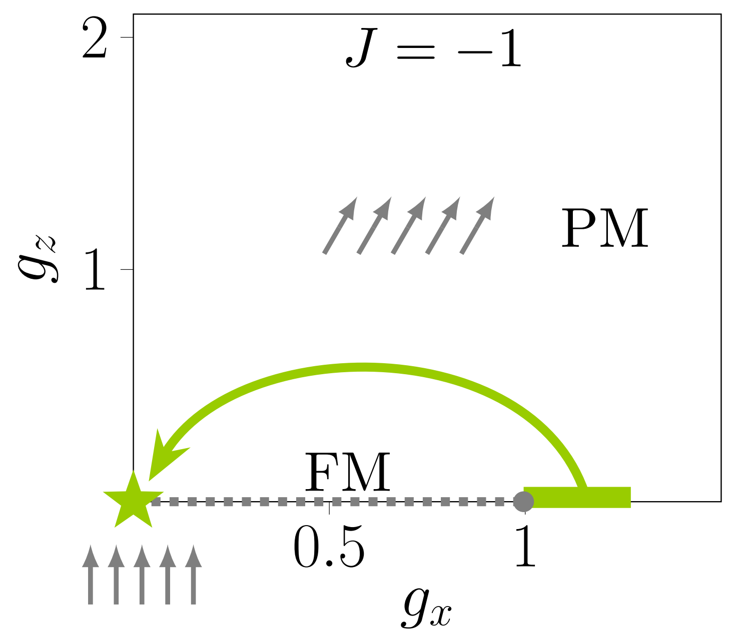

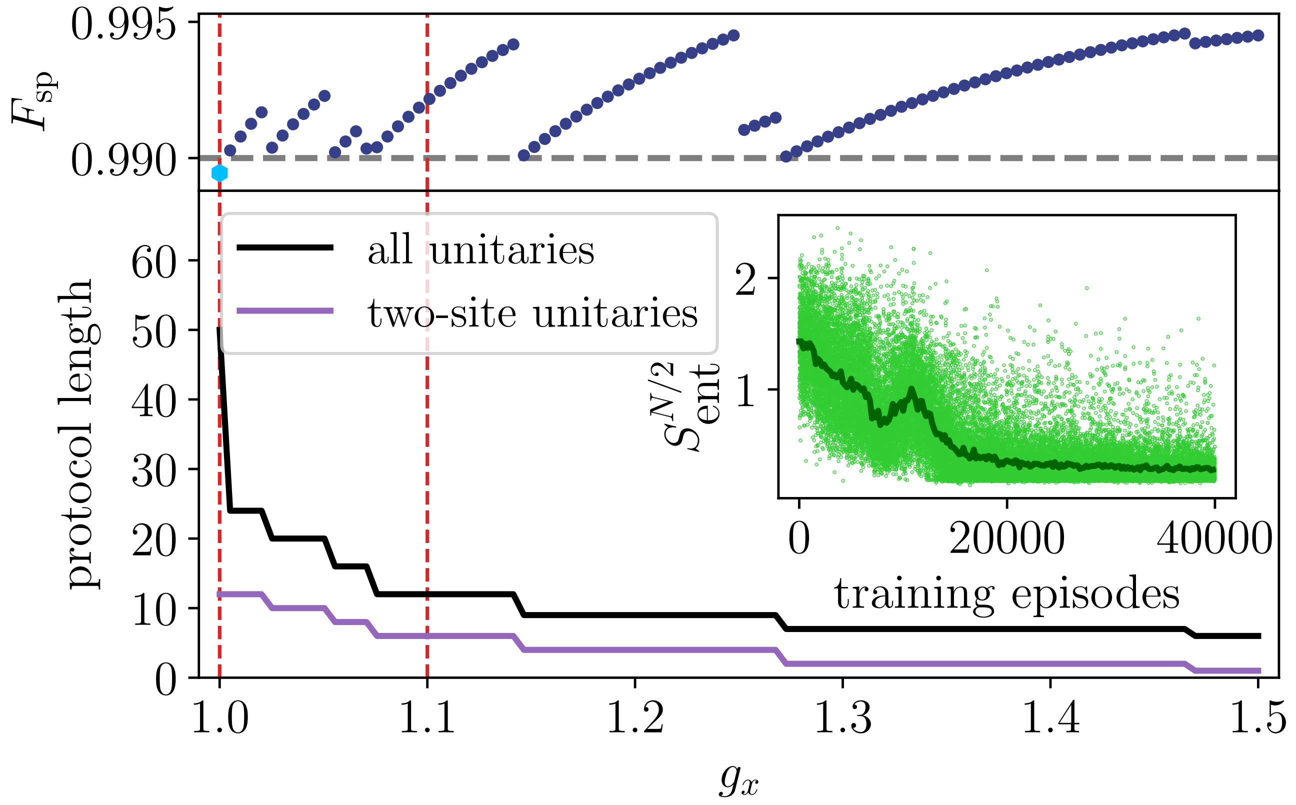

Outlook

- ➤ Map MPS to quantum circuit → integrate with NISQ device simulations
- ➤ Study ansatz/data/training using MPS toolbox → interpretable machine learning

MORE EXAMPLES

Universal state preparation

→ Learn to prepare target from arbitrary initial quantum state (N=4)

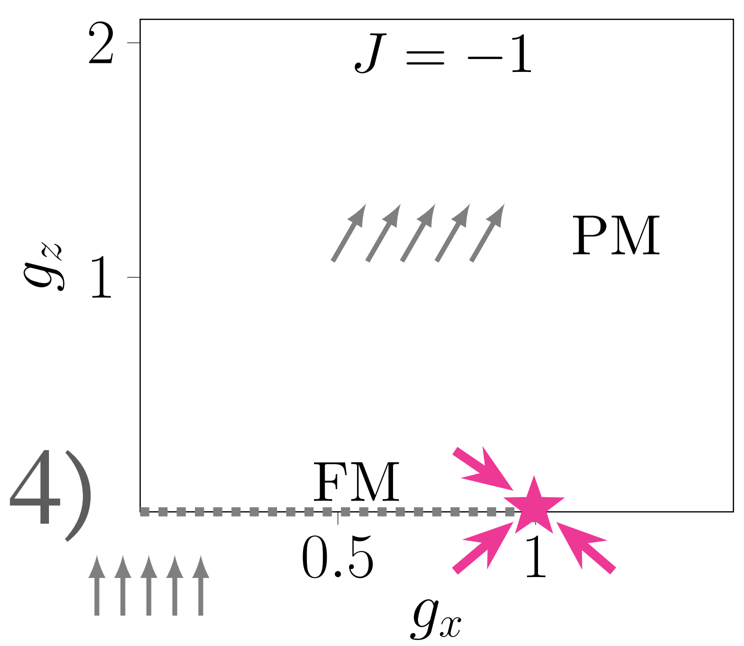

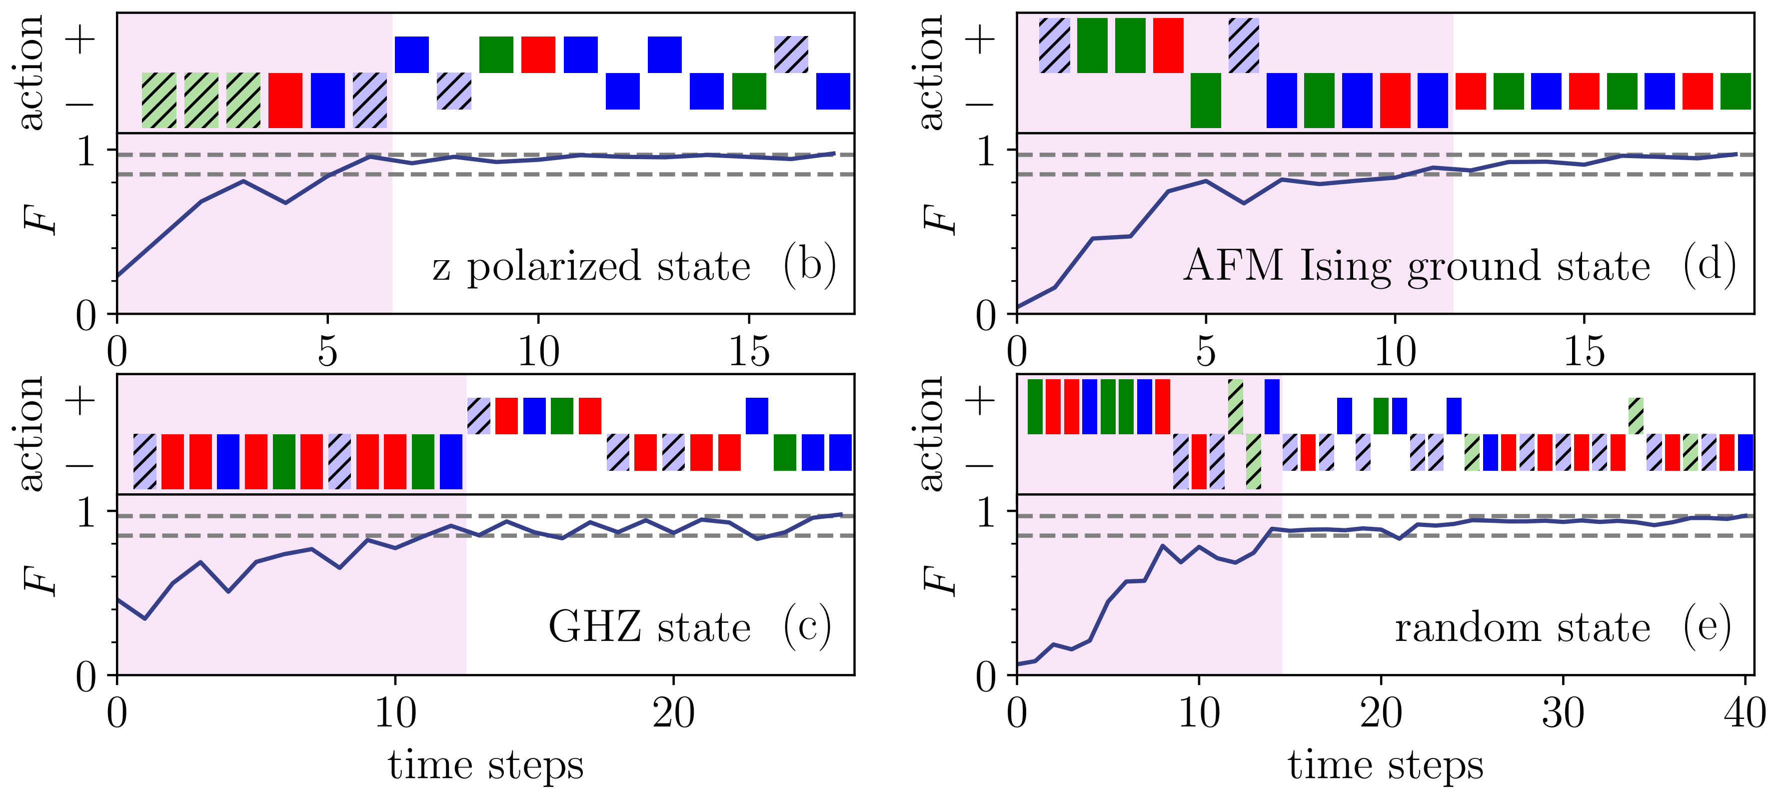

Many-body state preparation

→ $N = 32$ spins (beyond ED)

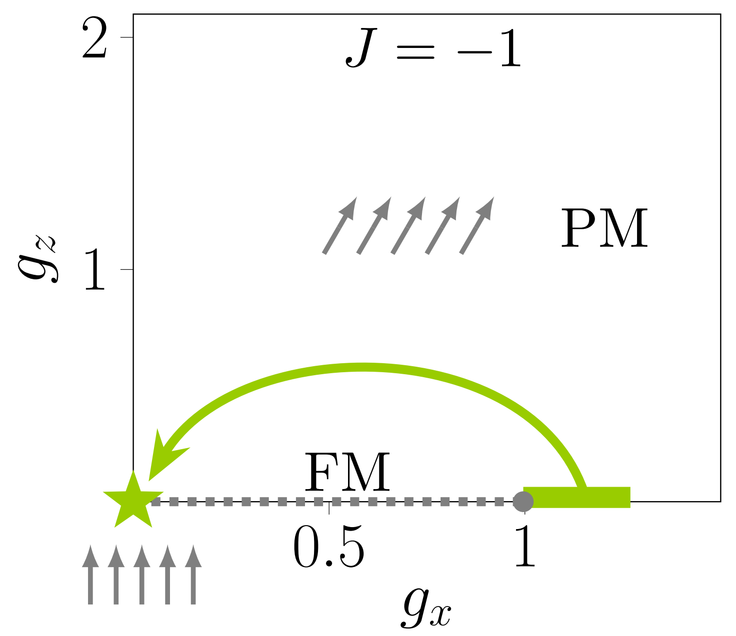

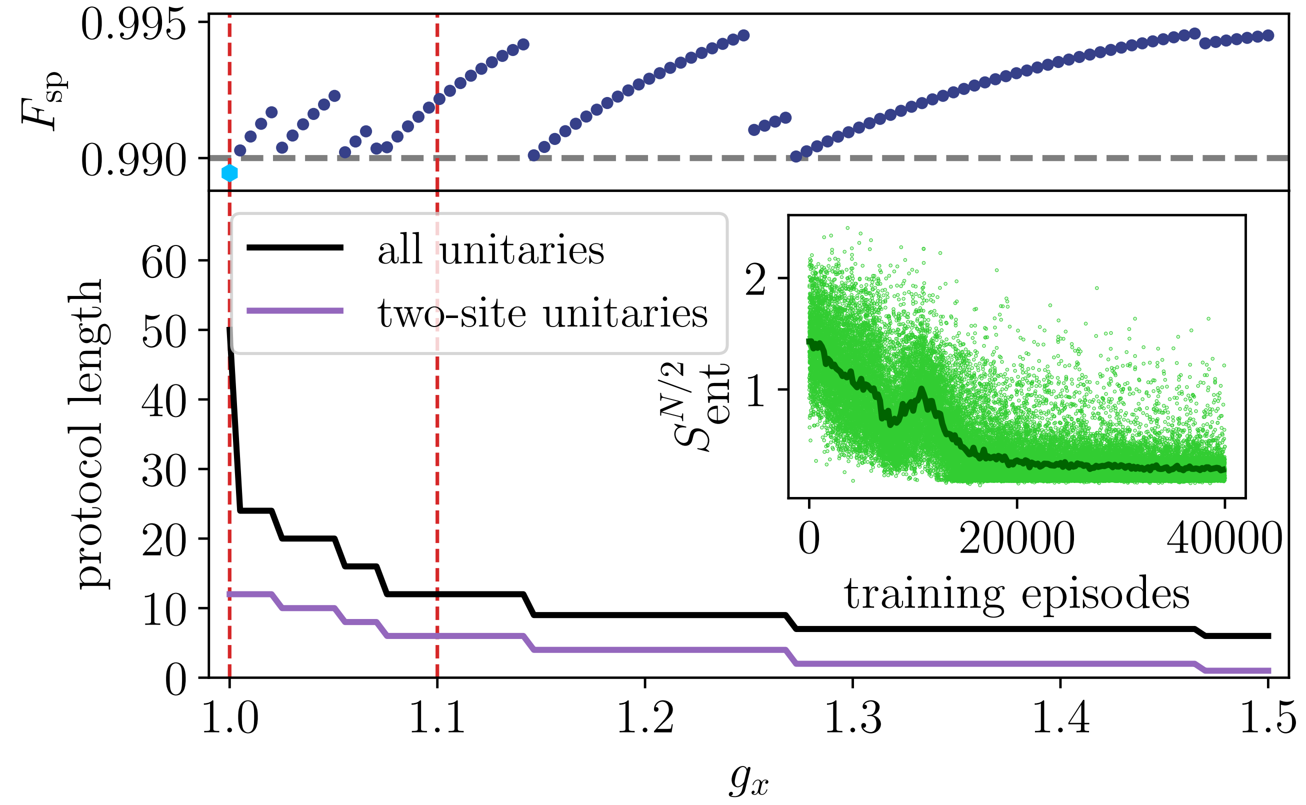

Marin Bukov

Outlook

- ➤ Map MPS to quantum circuit → integrate with NISQ device simulations
- ➤ Study ansatz/data/training using MPS toolbox → interpretable machine learning

THANK YOU!