Berkeley
UNIVERSITY OF CALIFORNIA

BU

Reinforcement Learning in Phases of Quantum Control

A. G.R. Day

D. Sels

P. Mehta

P. Weinberg

A. Polkovnikov

Marin Bukov

arXiv: *1705.00565* (2017)
arXiv: *1711.09109* (2017)

Berkeley
UNIVERSITY OF CALIFORNIA

BU

Reinforcement Learning in Phases of Quantum Control

A. G.R. Day

D. Sels

P. Mehta

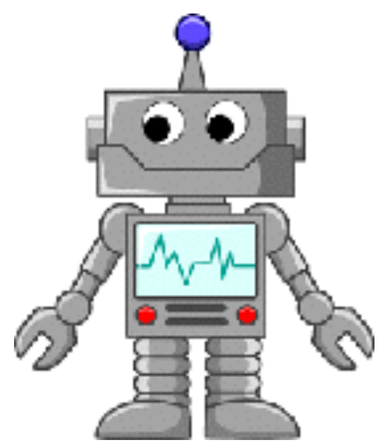

P. Weinberg

A. Polkovnikov

arXiv: *1705.00565* (2017)
arXiv: *1711.09109* (2017)

Marin Bukov

GOAL:

teach reinforcement learning agent to prepare states in non-integrable quantum Ising model

$$H(t) = -\sum_j J S^z_{j+1} S^z_j + h_z S^z_j + h_x(t) S^x_j$$

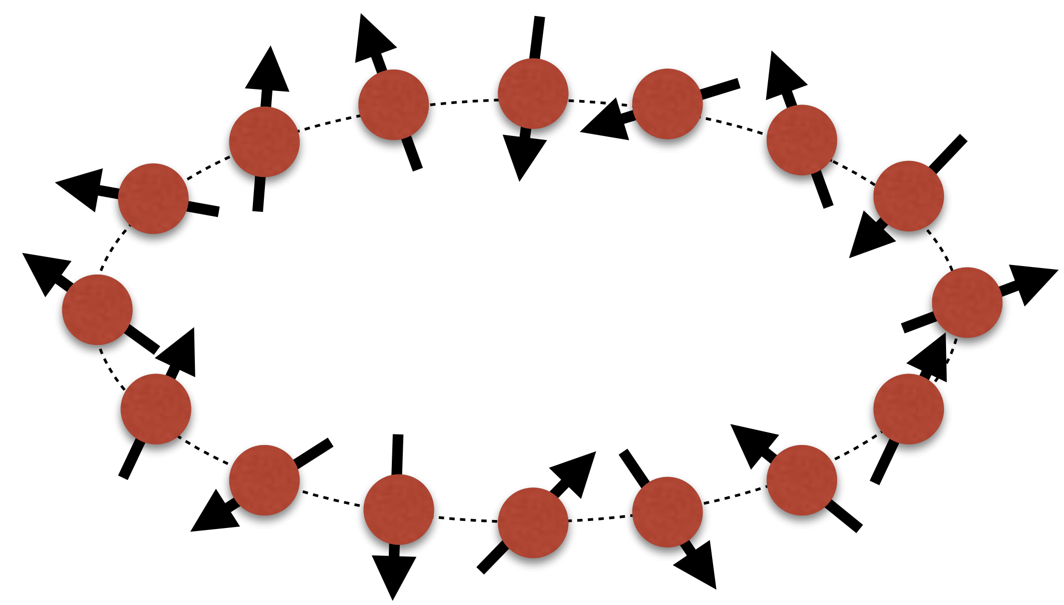

Marin Bukov

GOAL:

teach reinforcement learning agent to prepare states in non-integrable quantum Ising model

$$H(t) = -\sum_j J S^z_{j+1} S^z_j + h_z S^z_j + h_x(t) S^x_j$$

1. Single-qubit System $J = 0$ (THIS TALK)

→ problem setup, RL agent solution

→ control phase transitions

overconstrained phase,
correlated (glassy) phase,
controllable phase

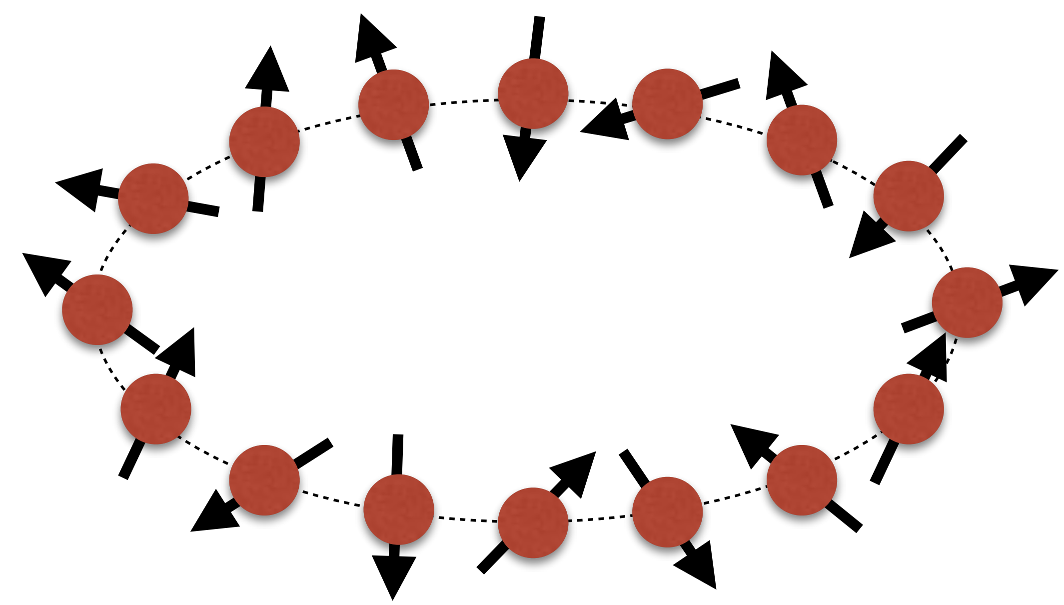

GOAL:

teach reinforcement learning agent to prepare states in non-integrable quantum Ising model

$$H(t) = -\sum_j J S^z_{j+1} S^z_j + h_z S^z_j + h_x(t) S^x_j$$

1. Single-qubit System $J=0$ (THIS TALK)

→ problem setup, RL agent solution

→ control phase transitions
overconstrained phase,
correlated (glassy) phase,
controllable phase

2. Two-qubit System $L=2$ (POSTER)

→ spontaneous symmetry breaking in optimal control landscape

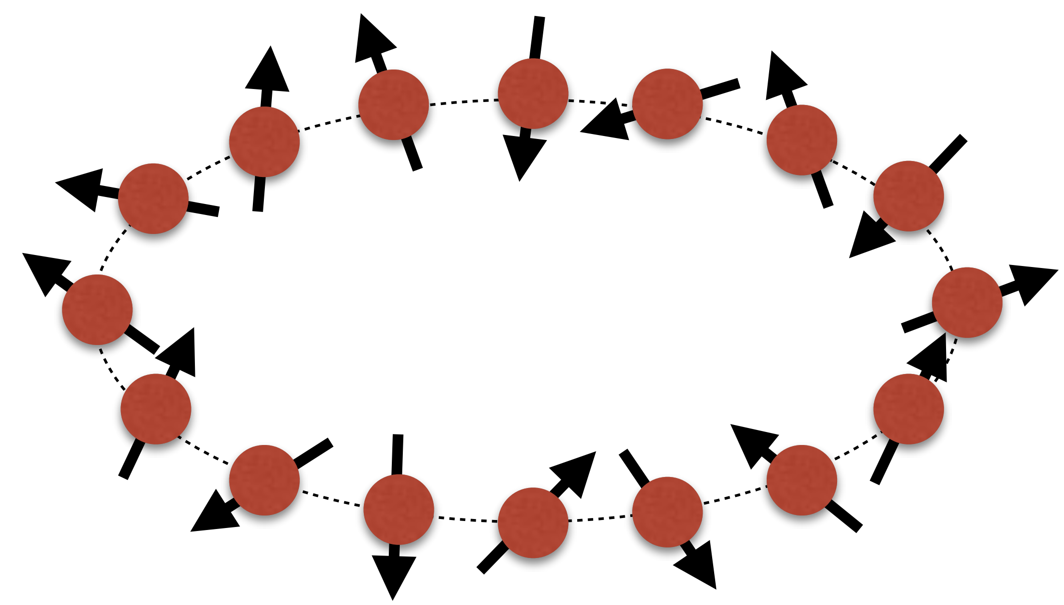

Marin Bukov

GOAL:

teach reinforcement learning agent to prepare states in non-integrable quantum Ising model

$$H(t) = -\sum_j J S^z_{j+1} S^z_j + h_z S^z_j + h_x(t) S^x_j$$

1. Single-qubit System $J = 0$ (THIS TALK)
   - ⟶ problem setup, RL agent solution
   - ⟶ control phase transitions
     overconstrained phase,
     correlated (glassy) phase,
     controllable phase

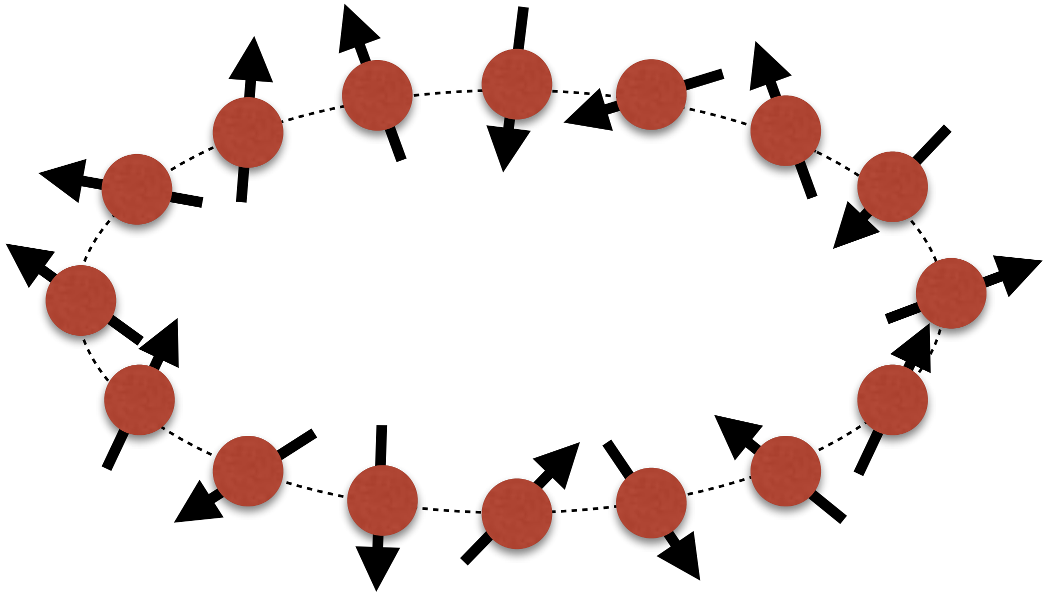

2. Two-qubit System $L = 2$ (POSTER)
   - ⟶ spontaneous symmetry breaking in optimal control landscape

3. Many-Body System (POSTER)
   - ⟶ control phase diagram: *overconstrained phase, glassy/correlated phase*

Optimal Qubit State Preparation

→ Hamiltonian: $H(t) = -S^z - h_x(t)S^x$

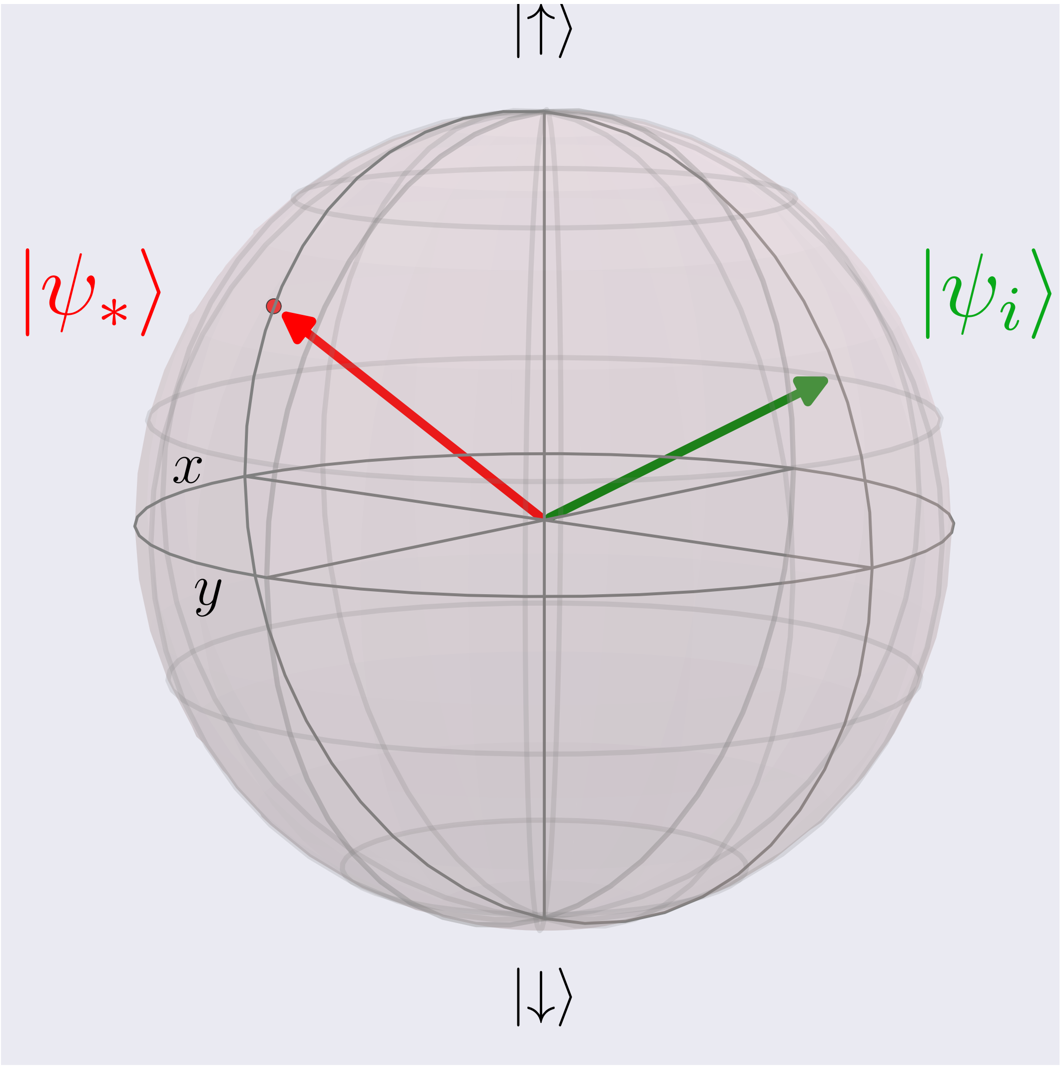

initial state: $|\psi_i\rangle$: $\text{GS of } H_i = -S^z - 2S^x$

target state: $|\psi_*\rangle$: $\text{GS of } H_* = -S^z + 2S^x$

Marin Bukov

Optimal Qubit State Preparation

→ Hamiltonian: $H(t) = -S^z - h_x(t)S^x$

initial state: $|\psi_i\rangle$: GS of $H_i = -S^z - 2S^x$

target state: $|\psi_*\rangle$: GS of $H_* = -S^z + 2S^x$

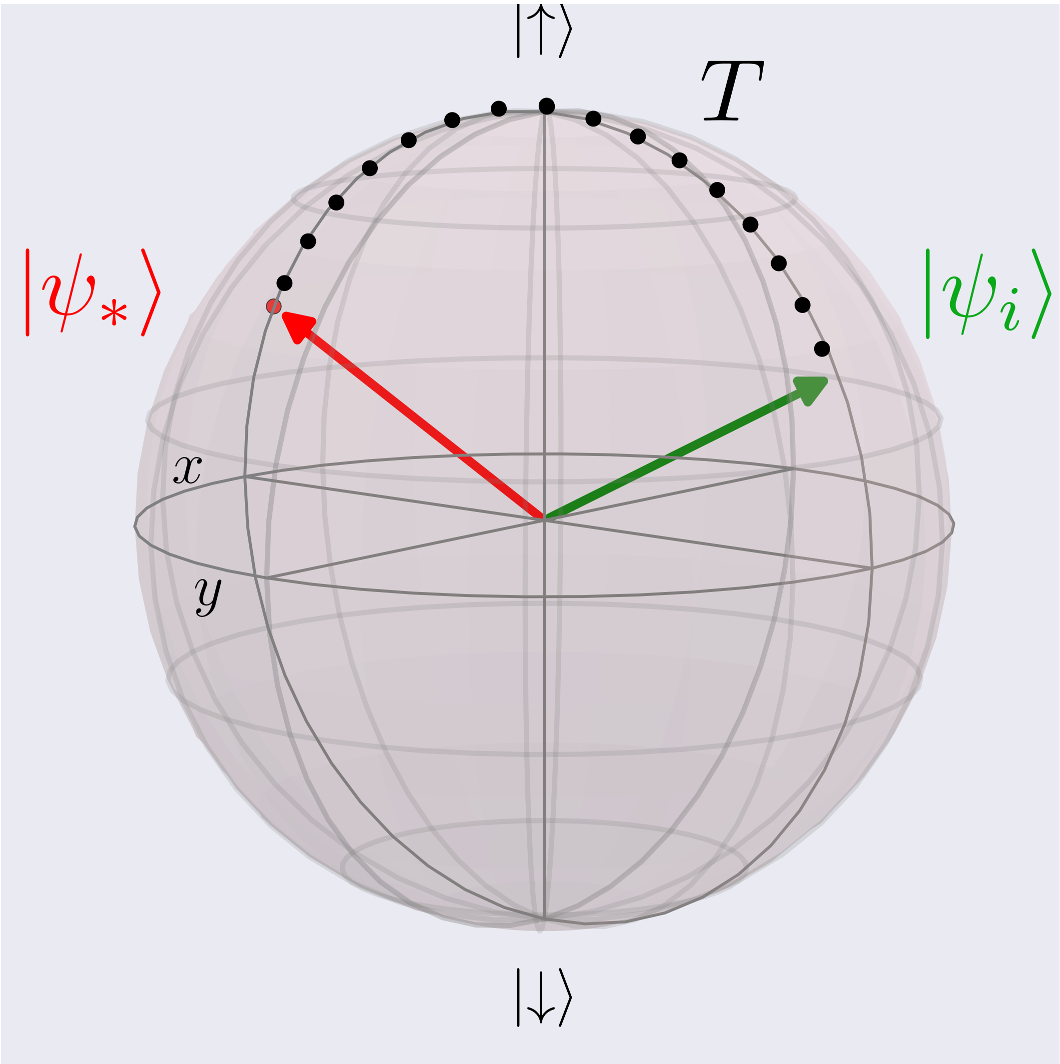

GOAL: find protocol $h(t) \in [-4, 4]$

such that $|\psi(t=0)\rangle = |\psi_i\rangle, \; |\psi(t=T)\rangle = |\psi_*\rangle$

measure: fidelity $F_h(T) = |\langle\psi(T)|\psi_*\rangle|^2$

Marin Bukov

What is Reinforcement Learning?

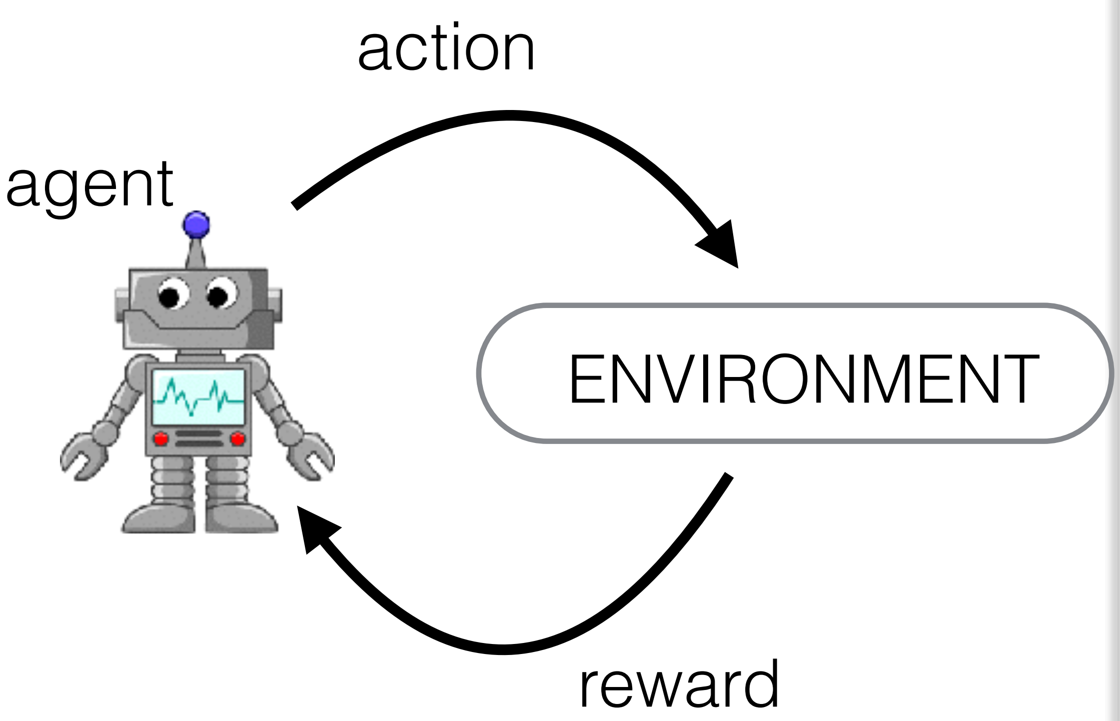

Marin Bukov

Silver et al., Nature 529 (2016) [Google DeepMind]

What is Reinforcement Learning?

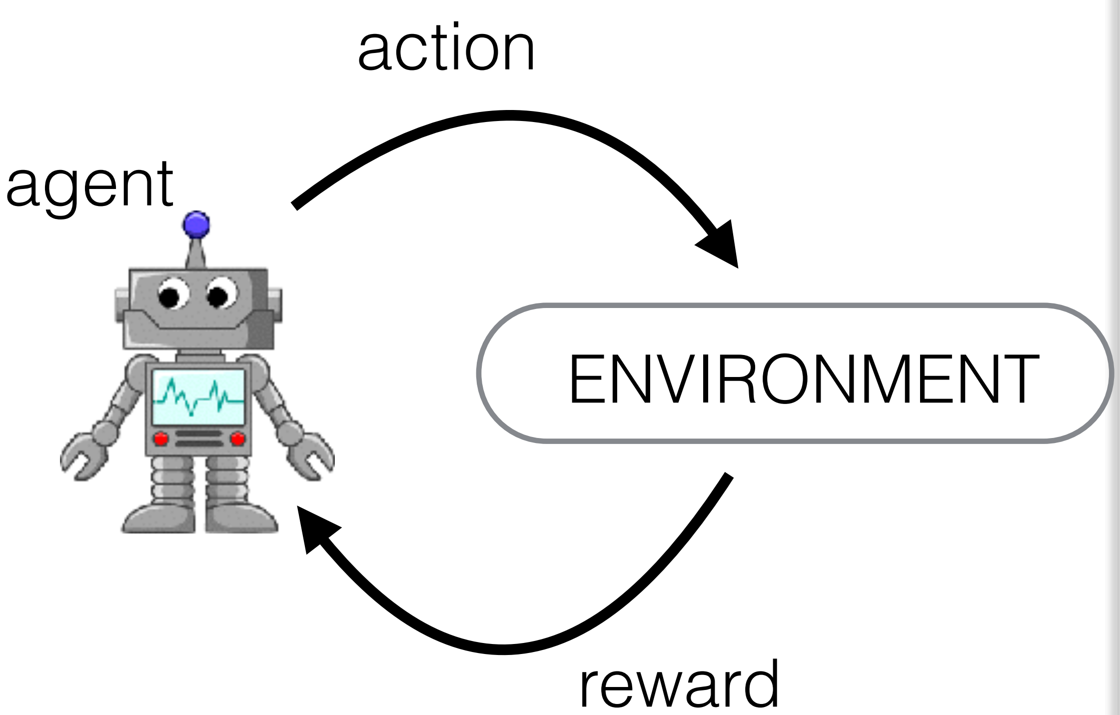

Silver et al., Nature 529 (2016) [Google DeepMind]

Reinforcement Learning in a Nutshell

- **Machine Learning**
  - Supervised Learning
  - ***Reinforcement Learning (RL)***
  - Unsupervised Learning

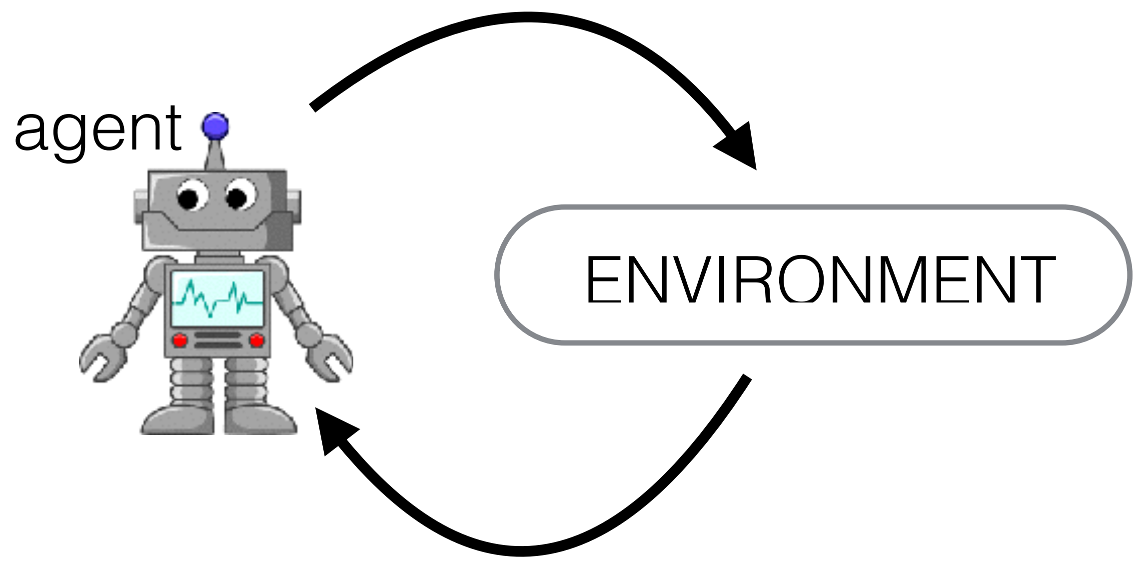

Marin Bukov

Sutton and Barto, *Reinforcement Learning: an Introduction*, MIT press (2015)

Reinforcement Learning in a Nutshell

- **Machine Learning**
  - Supervised Learning
  - ***Reinforcement Learning (RL)***
  - Unsupervised Learning

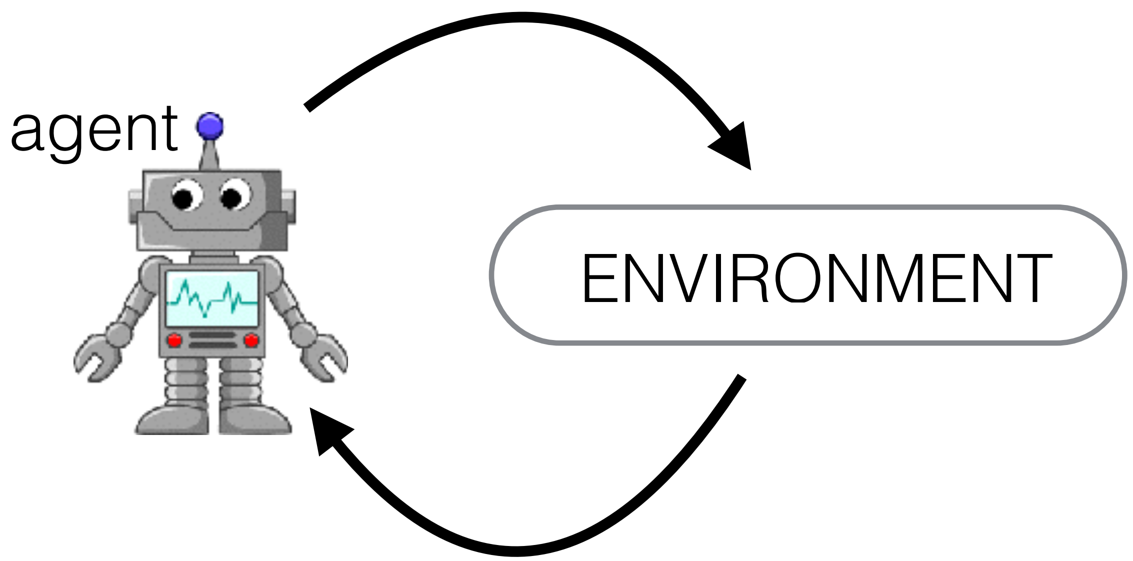

Sutton and Barto, *Reinforcement Learning: an Introduction*, MIT press (2015)

Reinforcement Learning in a Nutshell

- **Machine Learning**
  - Supervised Learning
  - ***Reinforcement Learning (RL)***
  - Unsupervised Learning

RL states $\mathcal{S} = \{s\}$

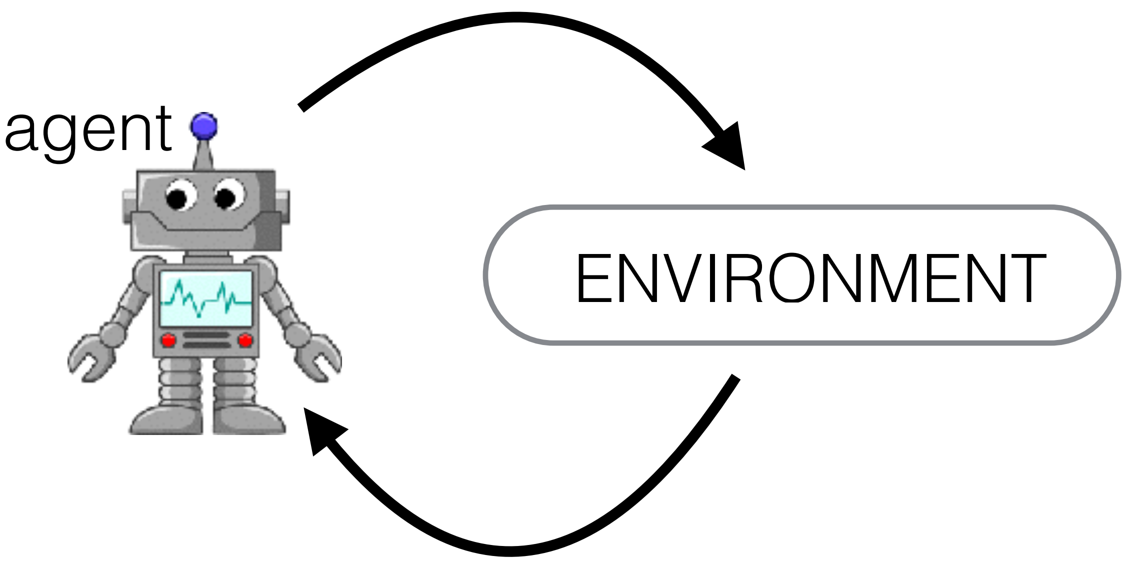

RL states $\mathcal{S} = \{(t, h(t))\}$

Marin Bukov

Sutton and Barto, *Reinforcement Learning: an Introduction*, MIT press (2015)

Berkeley
UNIVERSITY OF CALIFORNIA

Reinforcement Learning in a Nutshell

- **Machine Learning**
  - Supervised Learning
  - ***Reinforcement Learning (RL)***
  - Unsupervised Learning

RL states $\mathcal{S} = \{s\}$

available actions $\mathcal{A} = \{a\}$

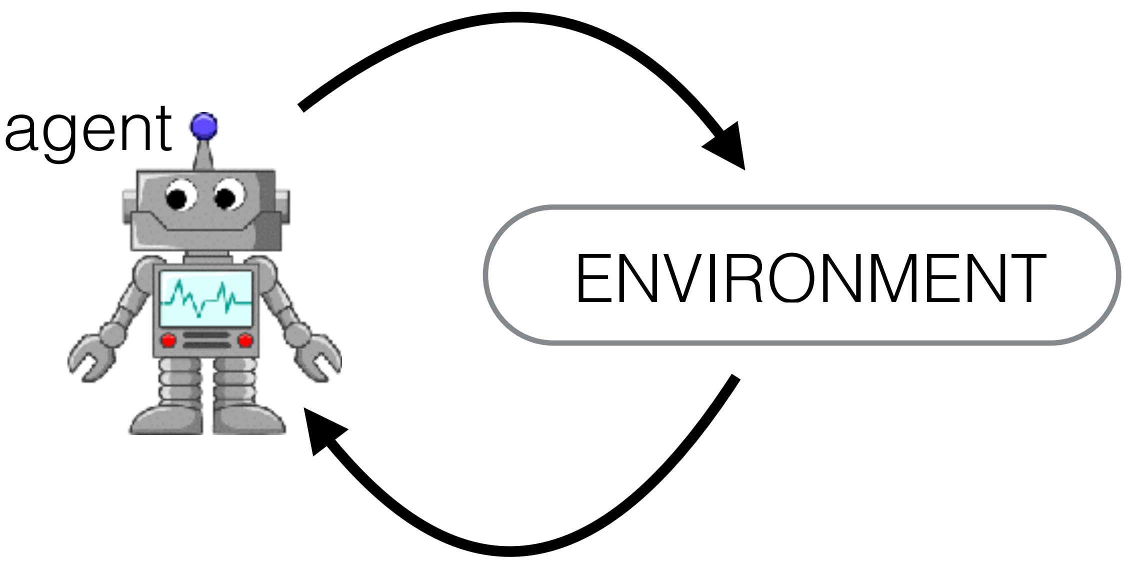

RL states $\mathcal{S} = \{(t, h(t))\}$

available actions $\mathcal{A} = \{a = \delta h\} = \{0, \pm 0.1, \pm 0.2, \pm 0.5, \pm 1.0, \pm 2.0, \pm 4.0, \pm 8.0\}$

Sutton and Barto, *Reinforcement Learning: an Introduction*, MIT press (2015)

Reinforcement Learning in a Nutshell

- **Machine Learning**
  - Supervised Learning
  - ***Reinforcement Learning (RL)***
  - Unsupervised Learning

RL states $\mathcal{S} = \{s\}$

available actions $\mathcal{A} = \{a\}$

rewards $\mathcal{R} = \{r\}$

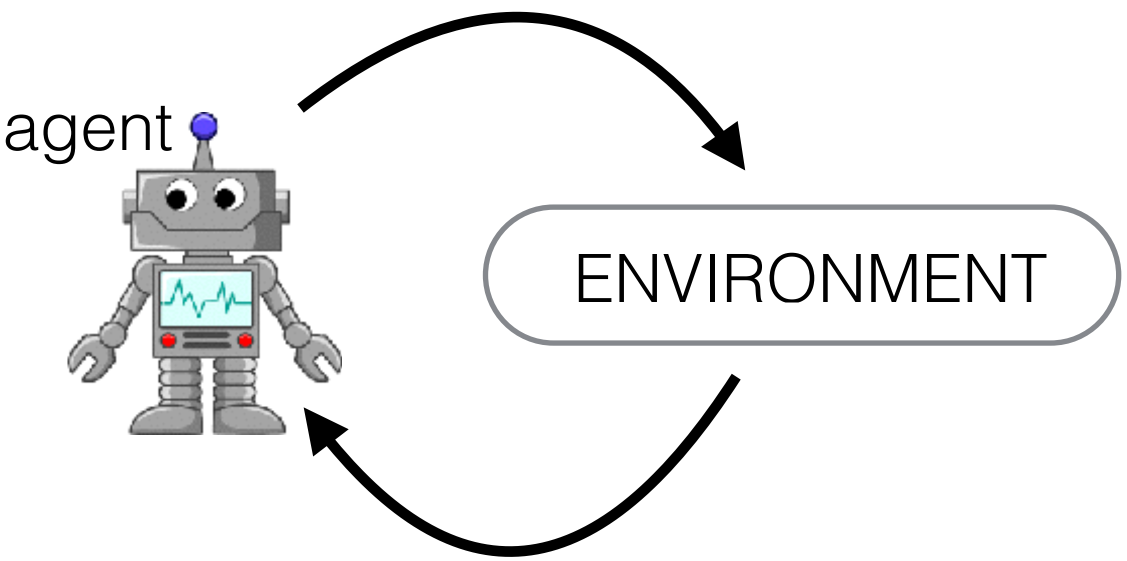

RL states $\mathcal{S} = \{(t, h(t))\}$

available actions $\mathcal{A} = \{a = \delta h\} = \{0, \pm 0.1, \pm 0.2, \pm 0.5, \pm 1.0, \pm 2.0, \pm 4.0, \pm 8.0\}$

rewards $\mathcal{R} = \{r(t) \in [0, 1]\}$

Sutton and Barto, *Reinforcement Learning: an Introduction*, MIT press (2015)

Reinforcement Learning in a Nutshell

- **Machine Learning**
  - Supervised Learning
  - ***Reinforcement Learning (RL)***
  - Unsupervised Learning

RL states $\mathcal{S} = \{s\}$

available actions $\mathcal{A} = \{a\}$

rewards $\mathcal{R} = \{r\}$

cumulative expected reward
or Q-function $Q(s,a)$

encodes experience/knowledge

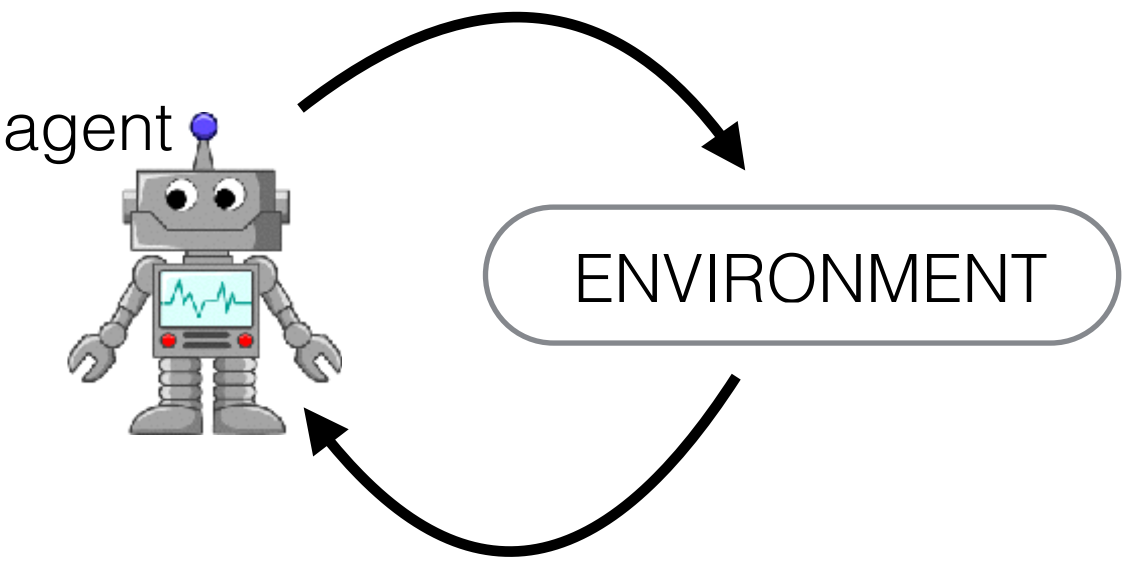

RL states $\mathcal{S} = \{(t, h(t))\}$

available actions $\mathcal{A} = \{a = \delta h\} = \{0, \pm 0.1, \pm 0.2, \pm 0.5, \pm 1.0, \pm 2.0, \pm 4.0, \pm 8.0\}$

rewards $\mathcal{R} = \{r(t) \in [0, 1]\}$

Marin Bukov

Sutton and Barto, *Reinforcement Learning: an Introduction*, MIT press (2015)

Berkeley
UNIVERSITY OF CALIFORNIA

Reinforcement Learning in a Nutshell

- **Machine Learning**
  - Supervised Learning
  - ***Reinforcement Learning (RL)***
  - Unsupervised Learning

RL states $\mathcal{S} = \{s\}$

available actions $\mathcal{A} = \{a\}$

rewards $\mathcal{R} = \{r\}$

cumulative expected reward
or Q-function $Q(s, a)$

encodes experience/knowledge

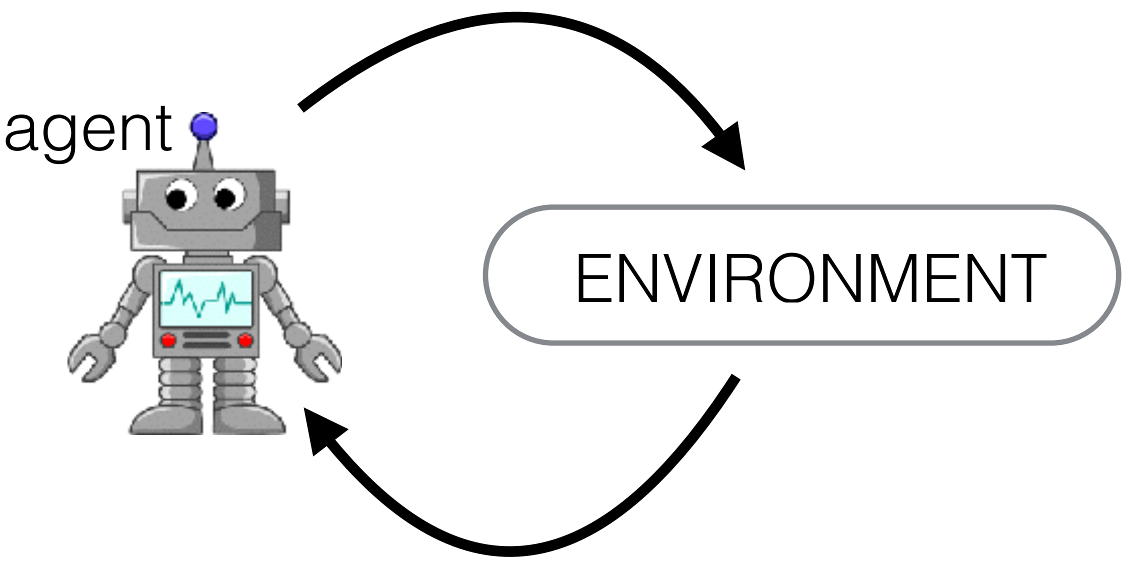

RL states $\mathcal{S} = \{(t, h(t))\}$

available actions $\mathcal{A} = \{a = \delta h\} = \{0, \pm 0.1, \pm 0.2, \pm 0.5, \pm 1.0, \pm 2.0, \pm 4.0, \pm 8.0\}$

rewards $\mathcal{R} = \{r(t) \in [0, 1]\}$

GOAL: maximise cumulative expected reward / Q-function $Q(s, a)$
starting from state s and taking action a

Marin Bukov

Sutton and Barto, *Reinforcement Learning: an Introduction*, MIT press (2015)

RL Applied to Quantum State Preparation

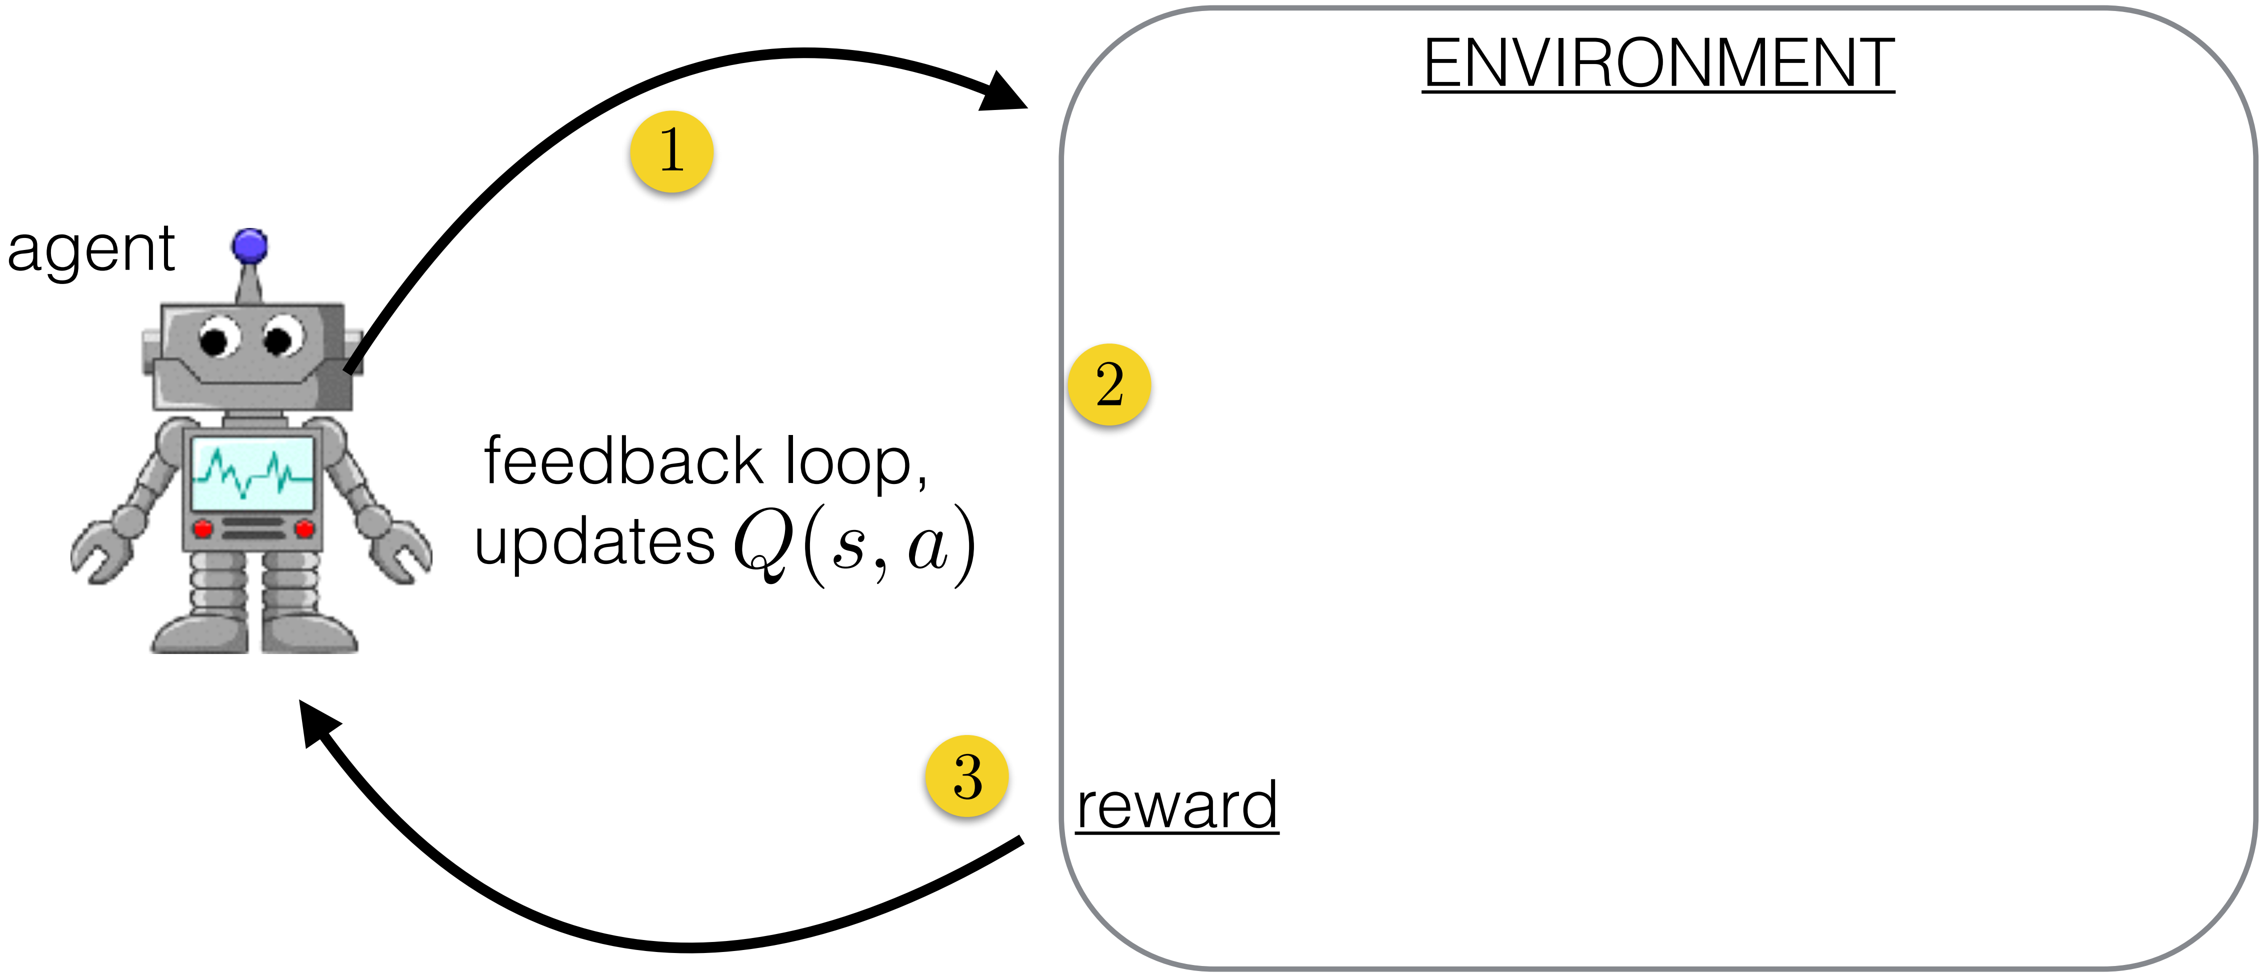

RL Applied to Quantum State Preparation

1

agent

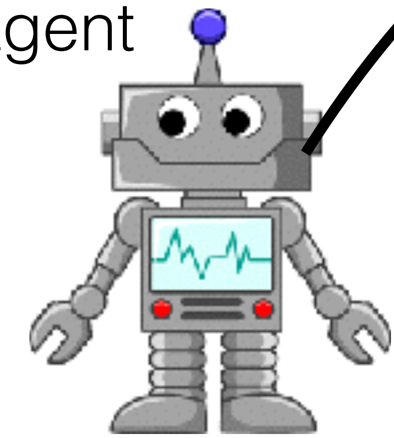

feedback loop,
updates $Q(s,a)$

ENVIRONMENT

2

$$H(t) = S^z + h(t)S^x$$

$$|\psi_i\rangle : \ \text{GS of } H_i = S^z - S^x$$
$$|\psi_*\rangle : \ \text{GS of } H_* = S^z + S^x$$

$$i\partial_t|\psi(t)\rangle = H(t)|\psi(t)\rangle, \ t \in [0,T]$$

3

reward $r = \begin{cases} 0 & , \ 0 \leq t < T \\ |\langle\psi_*|\psi(t=T)\rangle|^2, & t = T \end{cases}$

RL Applied to Quantum State Preparation

agent

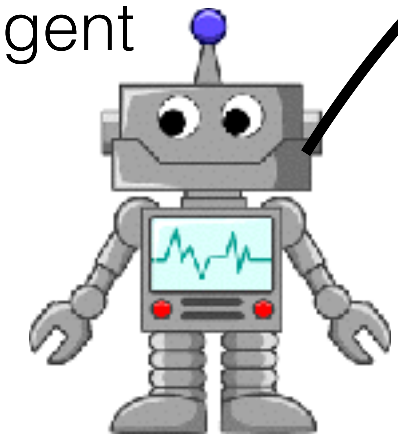

feedback loop, updates $Q(s,a)$

1

2

3

ENVIRONMENT

$$H(t) = S^z + h(t)S^x$$

$$|\psi_i\rangle : \ \text{GS of } H_i = S^z - S^x$$
$$|\psi_*\rangle : \ \text{GS of } H_* = S^z + S^x$$

$$i\partial_t|\psi(t)\rangle = H(t)|\psi(t)\rangle, \ t \in [0,T]$$

reward $r = \begin{cases} 0 & , \ 0 \le t < T \\ |\langle\psi_*|\psi(t=T)\rangle|^2, & t = T \end{cases}$

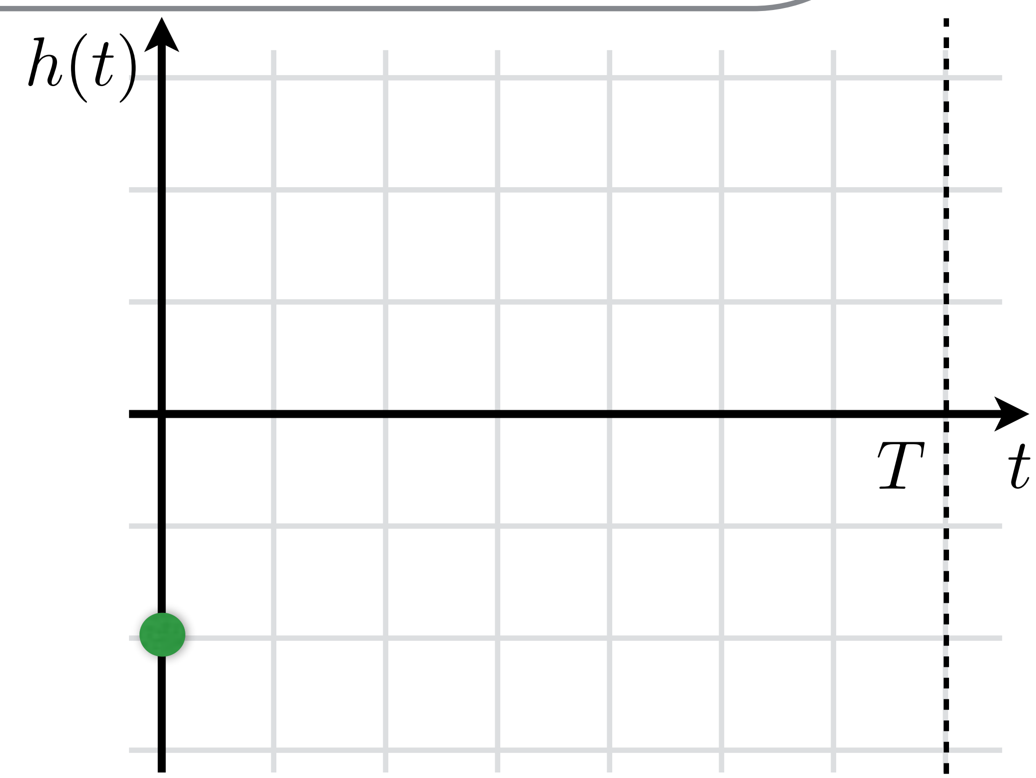

Marin Bukov

Berkeley UNIVERSITY OF CALIFORNIA

RL Applied to Quantum State Preparation

agent

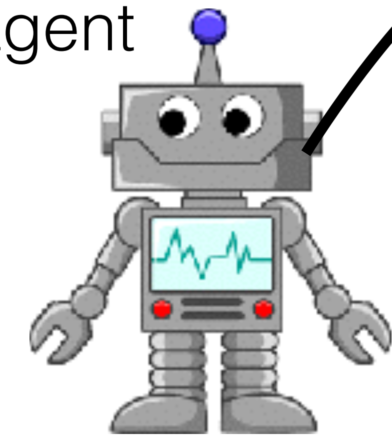

1

feedback loop, updates $Q(s,a)$

3

ENVIRONMENT

2

$$H(t) = S^z + h(t)S^x$$

$$|\psi_i\rangle : \ \text{GS of } H_i = S^z - S^x$$
$$|\psi_*\rangle : \ \text{GS of } H_* = S^z + S^x$$

$$i\partial_t|\psi(t)\rangle = H(t)|\psi(t)\rangle, \ t \in [0,T]$$

reward $r = \begin{cases} 0 & , \ 0 \le t < T \\ |\langle\psi_*|\psi(t=T)\rangle|^2, & t = T \end{cases}$

1 start from state $s_0 = (t = 0, h = -1)$

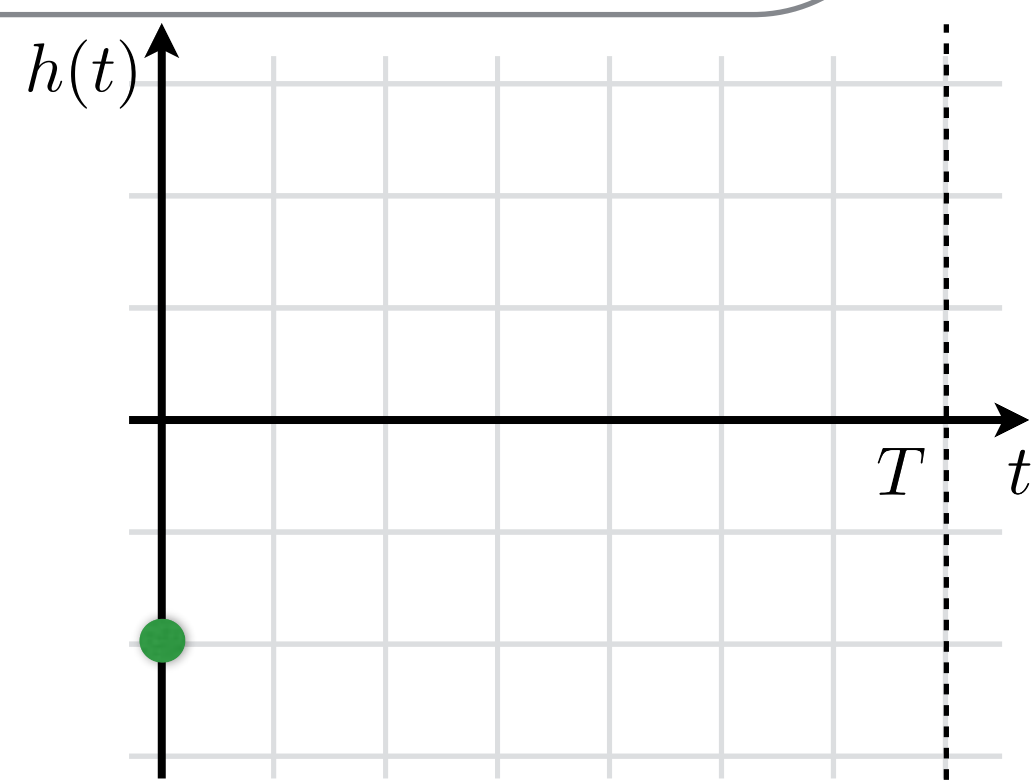

Marin Bukov

RL Applied to Quantum State Preparation

agent

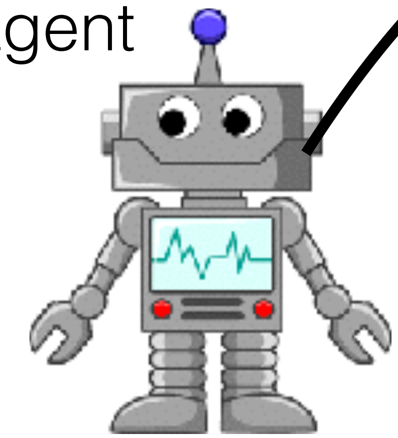

1

feedback loop, updates $Q(s,a)$

3

ENVIRONMENT

2

$$H(t) = S^z + h(t)S^x$$

$$|\psi_i\rangle : \text{ GS of } H_i = S^z - S^x$$
$$|\psi_*\rangle : \text{ GS of } H_* = S^z + S^x$$

$$i\partial_t|\psi(t)\rangle = H(t)|\psi(t)\rangle, \; t \in [0,T]$$

reward $r = \begin{cases} 0 & , \; 0 \le t < T \\ |\langle\psi_*|\psi(t=T)\rangle|^2, & t = T \end{cases}$

1 start from state $s_0 = (t=0, h=-1)$

take action $a_0 : \delta h = 0.8$

go to state $s_1 = (t=0.05, h=-0.2)$

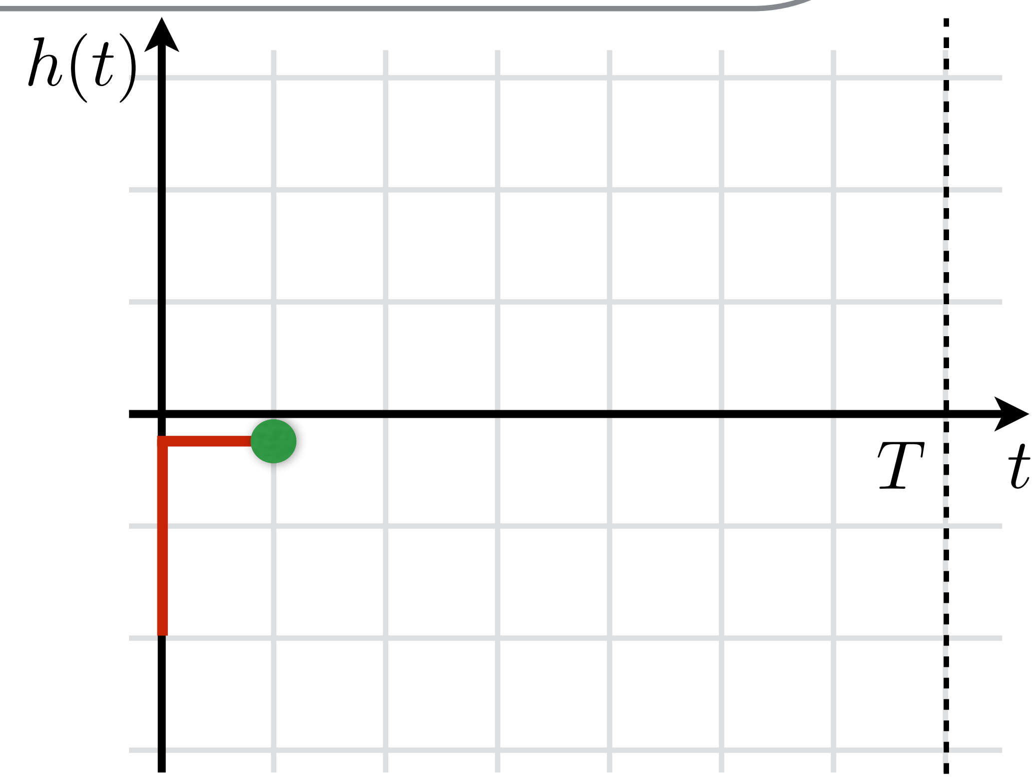

Marin Bukov

RL Applied to Quantum State Preparation

agent

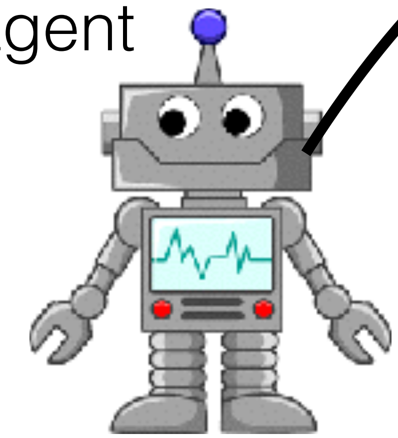

feedback loop, updates $Q(s,a)$

1 2 3

ENVIRONMENT

$$H(t) = S^z + h(t)S^x$$

$$|\psi_i\rangle : \text{ GS of } H_i = S^z - S^x$$

$$|\psi_*\rangle : \text{ GS of } H_* = S^z + S^x$$

$$i\partial_t|\psi(t)\rangle = H(t)|\psi(t)\rangle, \; t \in [0,T]$$

reward $r = \begin{cases} 0 & , \; 0 \le t < T \\ |\langle\psi_*|\psi(t=T)\rangle|^2, & t = T \end{cases}$

1. start from state $s_0 = (t=0, h=-1)$
   take action $a_0 : \delta h = 0.8$
   go to state $s_1 = (t=0.05, h=-0.2)$

2. solve Schrödinger Eq. and obtain the QM state $|\psi(\delta t)\rangle$

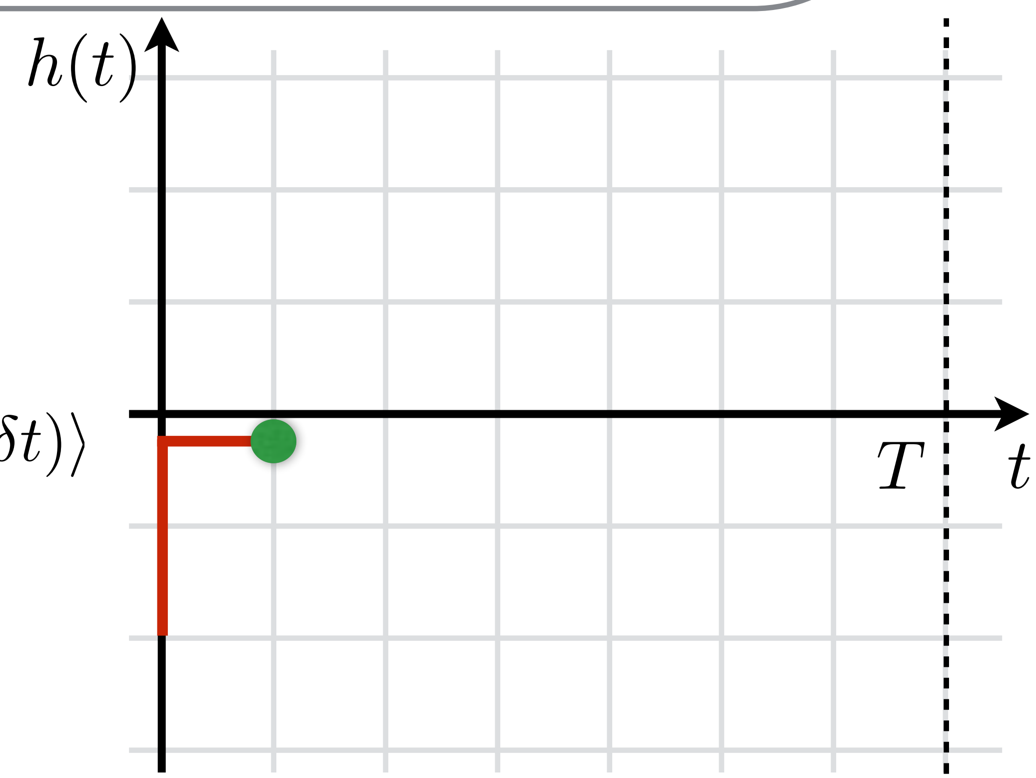

Marin Bukov

Berkeley
UNIVERSITY OF CALIFORNIA

RL Applied to Quantum State Preparation

agent

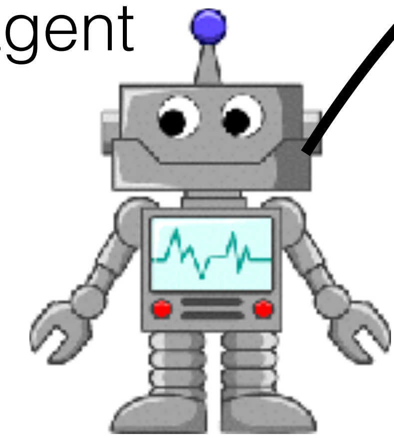

1

feedback loop,
updates $Q(s,a)$

ENVIRONMENT

$$H(t) = S^z + h(t)S^x$$

2

$$|\psi_i\rangle : \ \text{GS of } H_i = S^z - S^x$$
$$|\psi_*\rangle : \ \text{GS of } H_* = S^z + S^x$$

$$i\partial_t|\psi(t)\rangle = H(t)|\psi(t)\rangle, \ t \in [0,T]$$

3

reward $r = \begin{cases} 0 & , \ 0 \le t < T \\ |\langle\psi_*|\psi(t=T)\rangle|^2, & t = T \end{cases}$

1 start from state $s_0 = (t=0, h=-1)$
take action $a_0 : \delta h = 0.8$
go to state $s_1 = (t=0.05, h=-0.2)$

2 solve Schrödinger Eq. and obtain the QM state $|\psi(\delta t)\rangle$

3 calculate reward r
and use it to update $Q(s,a)$
which in turn is used to choose subsequent actions

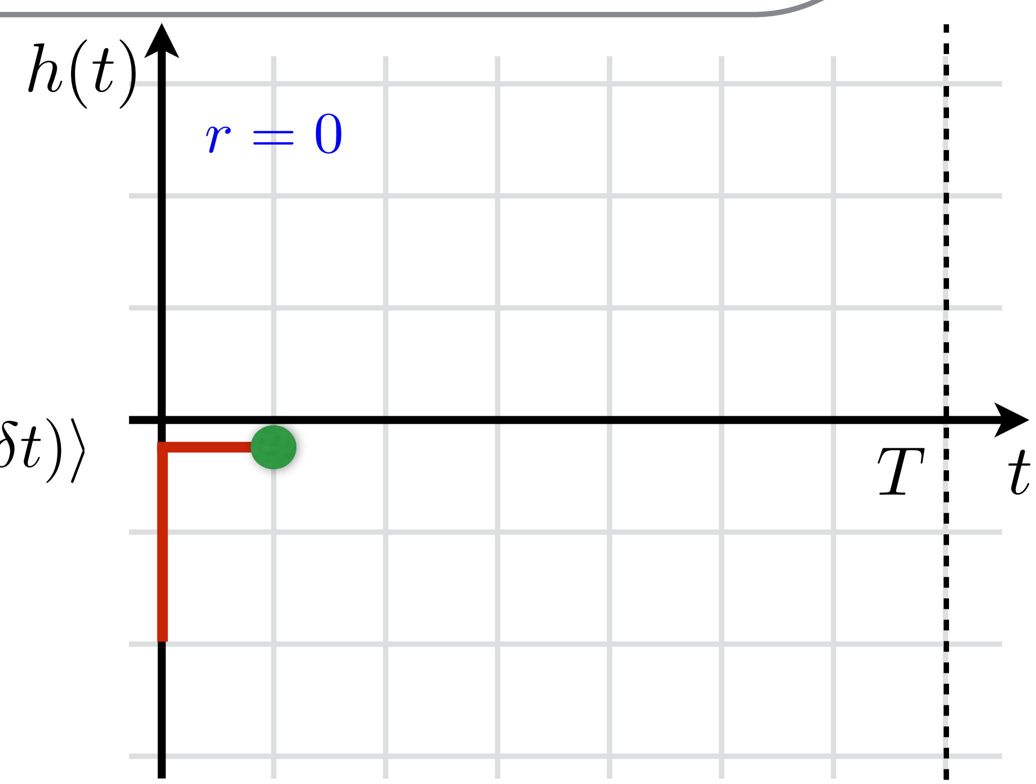

Marin Bukov

RL Applied to Quantum State Preparation

agent

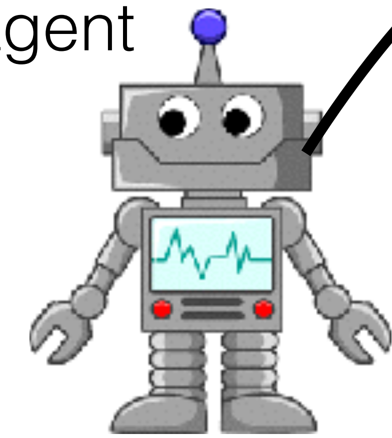

feedback loop,
updates $Q(s,a)$

1

ENVIRONMENT

2

$$H(t) = S^z + h(t)S^x$$

$$|\psi_i\rangle : \text{ GS of } H_i = S^z - S^x$$

$$|\psi_*\rangle : \text{ GS of } H_* = S^z + S^x$$

$$i\partial_t|\psi(t)\rangle = H(t)|\psi(t)\rangle, \; t\in[0,T]$$

3

reward $r = \begin{cases} 0 & , \; 0 \le t < T \\ |\langle\psi_*|\psi(t=T)\rangle|^2, & t = T \end{cases}$

1. start from state $s_0 = (t=0, h=-1)$
   take action $a_0 : \delta h = 0.8$
   go to state $s_1 = (t=0.05, h=-0.2)$

2. solve Schrödinger Eq. and obtain the QM state $|\psi(\delta t)\rangle$

3. calculate reward r
   and use it to update $Q(s,a)$
   which in turn is used to choose subsequent actions

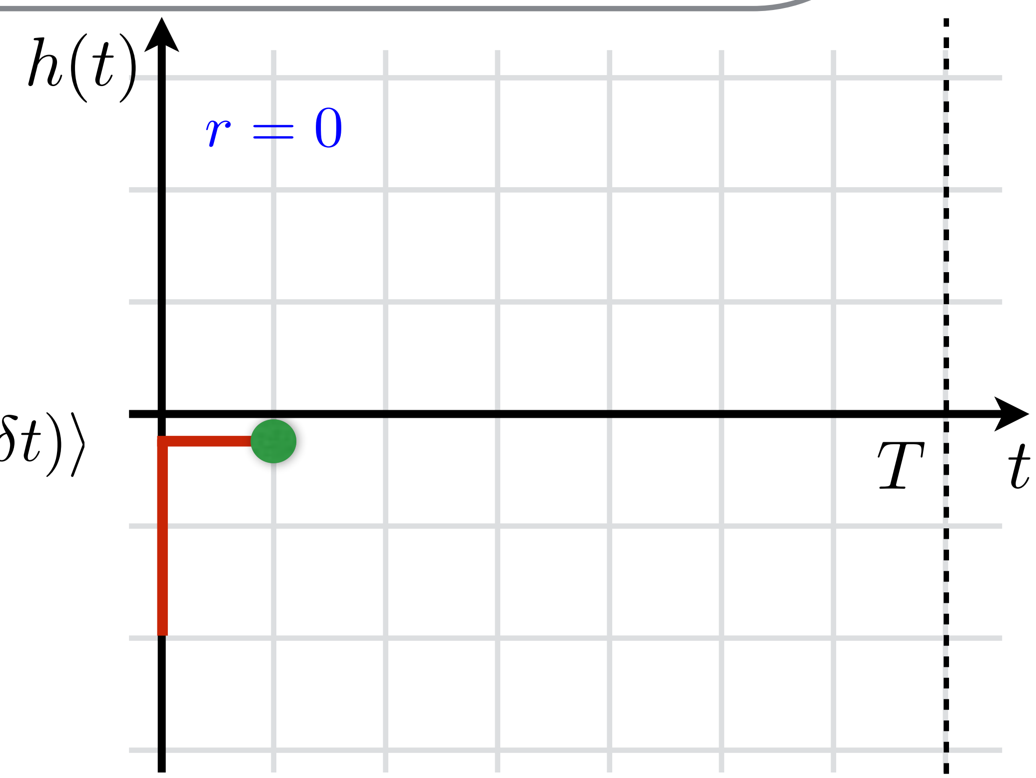

Marin Bukov

Berkeley
UNIVERSITY OF CALIFORNIA

RL Applied to Quantum State Preparation

agent

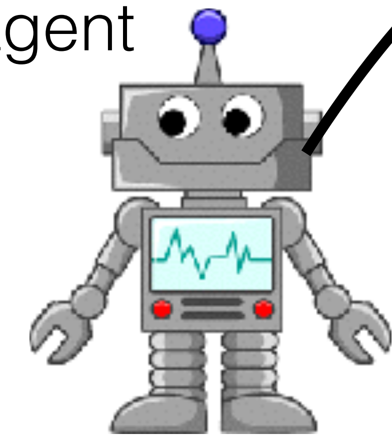

feedback loop, updates $Q(s,a)$

1

ENVIRONMENT

2

$$H(t) = S^z + h(t)S^x$$

$$|\psi_i\rangle : \ \text{GS of } H_i = S^z - S^x$$
$$|\psi_*\rangle : \ \text{GS of } H_* = S^z + S^x$$

$$i\partial_t|\psi(t)\rangle = H(t)|\psi(t)\rangle, \ t \in [0,T]$$

3

reward $r = \begin{cases} 0 & , \ 0 \leq t < T \\ |\langle\psi_*|\psi(t=T)\rangle|^2, & t = T \end{cases}$

1. start from state $s_0 = (t=0, h=-1)$
   take action $a_0 : \delta h = 0.8$
   go to state $s_1 = (t=0.05, h=-0.2)$

2. solve Schrödinger Eq. and obtain the QM state $|\psi(\delta t)\rangle$

3. calculate reward r
   and use it to update $Q(s,a)$
   which in turn is used to choose subsequent actions

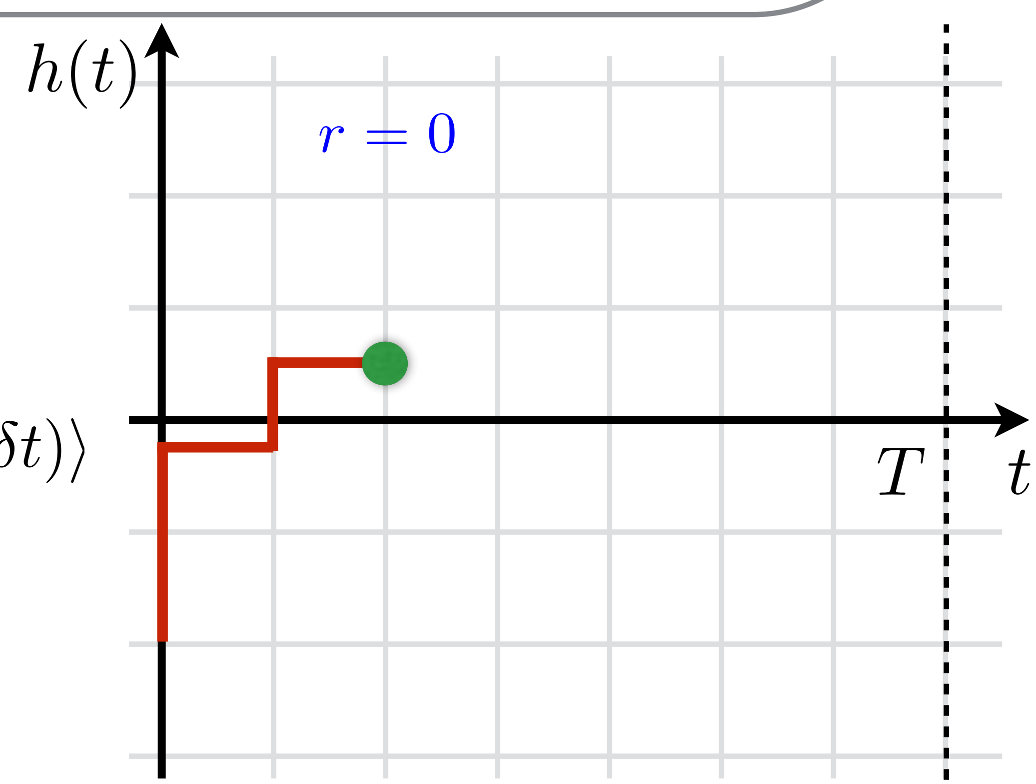

Marin Bukov

Berkeley UNIVERSITY OF CALIFORNIA

RL Applied to Quantum State Preparation

agent

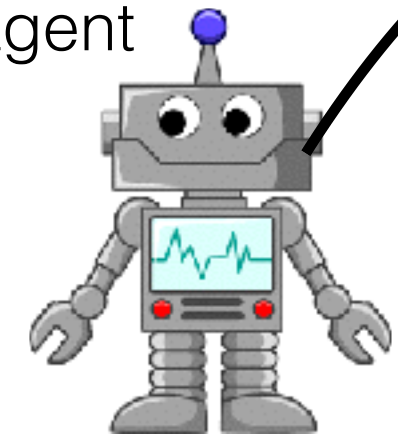

feedback loop, updates $Q(s,a)$

1

2

3

ENVIRONMENT

$$H(t) = S^z + h(t)S^x$$

$$|\psi_i\rangle : \text{ GS of } H_i = S^z - S^x$$

$$|\psi_*\rangle : \text{ GS of } H_* = S^z + S^x$$

$$i\partial_t|\psi(t)\rangle = H(t)|\psi(t)\rangle,\ t \in [0,T]$$

reward $r = \begin{cases} 0 & , \ 0 \le t < T \\ |\langle\psi_*|\psi(t=T)\rangle|^2, & t = T \end{cases}$

1 start from state $s_0 = (t=0, h=-1)$

take action $a_0 : \delta h = 0.8$

go to state $s_1 = (t=0.05, h=-0.2)$

2 solve Schrödinger Eq. and obtain the QM state $|\psi(\delta t)\rangle$

3 calculate reward r

and use it to update $Q(s,a)$

which in turn is used to choose subsequent actions

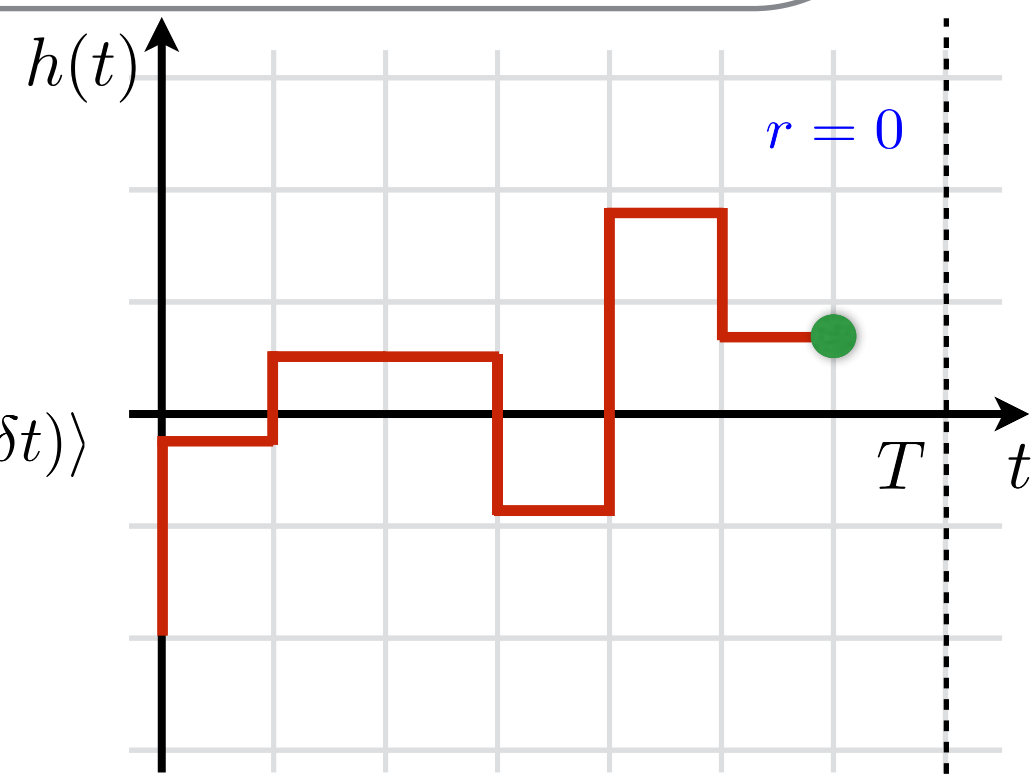

Marin Bukov

RL Applied to Quantum State Preparation

agent

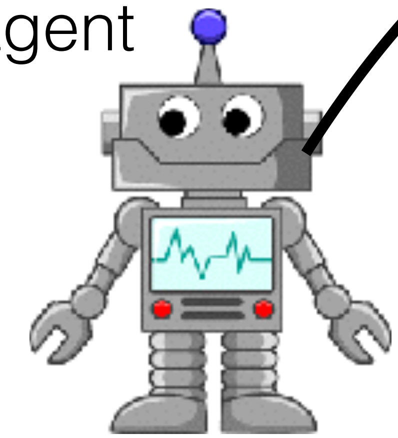

feedback loop, updates $Q(s,a)$

ENVIRONMENT

$$H(t) = S^z + h(t)S^x$$

$$|\psi_i\rangle : \ \text{GS of } H_i = S^z - S^x$$
$$|\psi_*\rangle : \ \text{GS of } H_* = S^z + S^x$$

$$i\partial_t|\psi(t)\rangle = H(t)|\psi(t)\rangle, \ t \in [0,T]$$

reward $r = \begin{cases} 0 & , \ 0 \le t < T \\ |\langle\psi_*|\psi(t=T)\rangle|^2, & t = T \end{cases}$

1 start from state $s_0 = (t=0, h=-1)$
take action $a_0 : \delta h = 0.8$
go to state $s_1 = (t=0.05, h=-0.2)$

2 solve Schrödinger Eq. and obtain the QM state $|\psi(\delta t)\rangle$

3 calculate reward r
and use it to update $Q(s,a)$
which in turn is used to choose subsequent actions

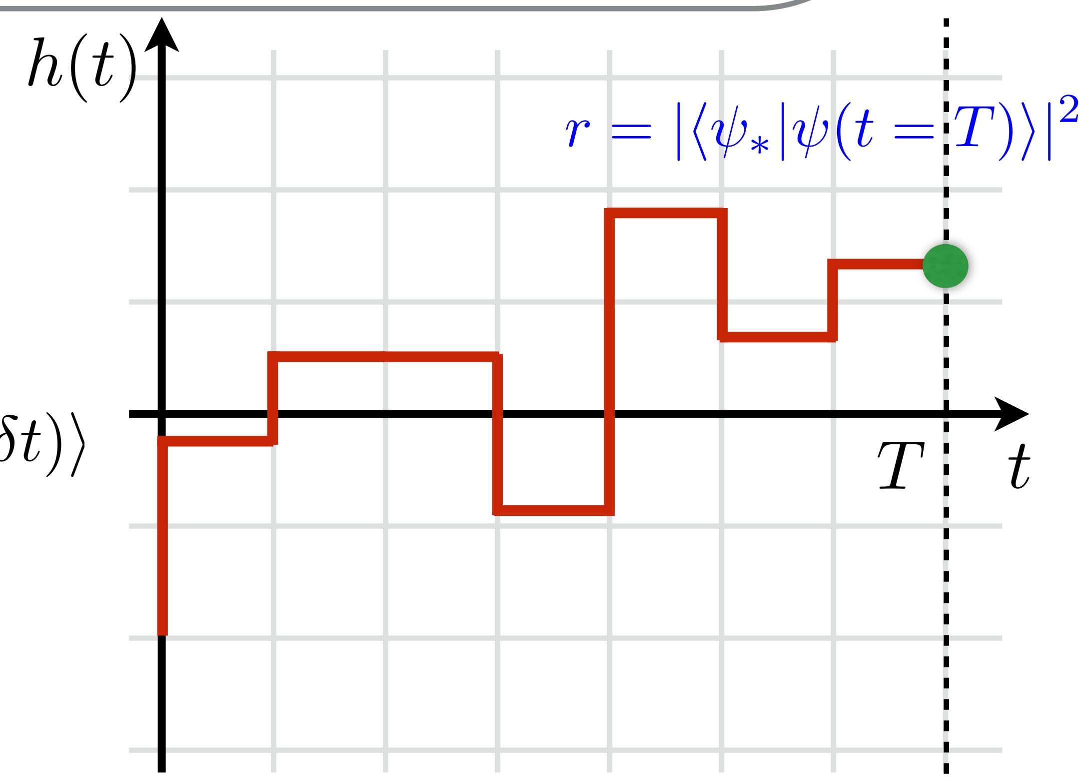

episode completed

Marin Bukov

RL Applied to Quantum State Preparation

agent

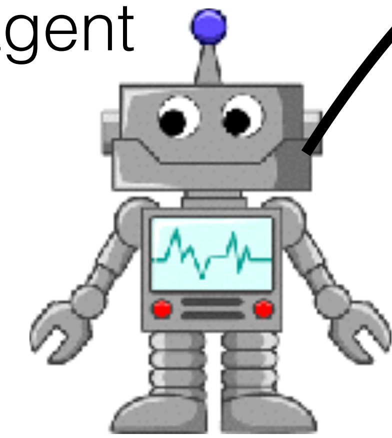

1

feedback loop,
updates $Q(s,a)$

3

<u>ENVIRONMENT</u>

2

$$H(t) = S^z + h(t)S^x$$

$$|\psi_i\rangle : \text{ GS of } H_i = S^z - S^x$$
$$|\psi_*\rangle : \text{ GS of } H_* = S^z + S^x$$

$$i\partial_t|\psi(t)\rangle = H(t)|\psi(t)\rangle,\ t\in[0,T]$$

<u>reward</u> $r = \begin{cases} 0 & ,\ 0 \le t < T \\ |\langle\psi_*|\psi(t=T)\rangle|^2, & t = T \end{cases}$

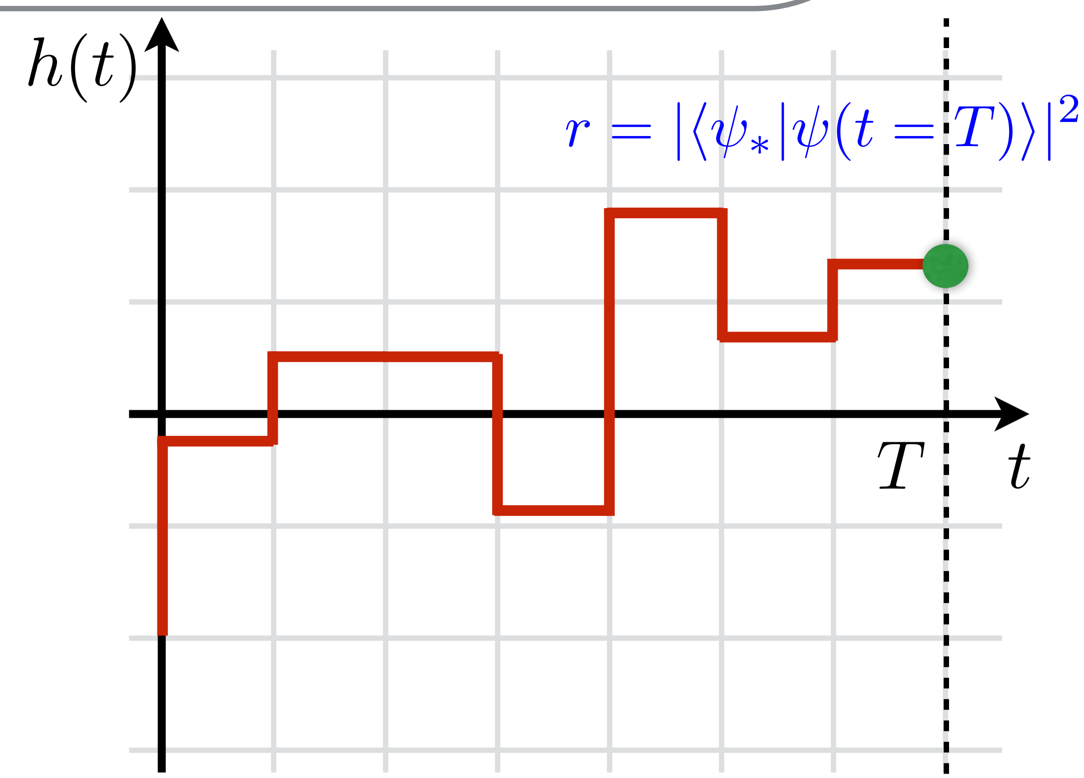

problems:

→ state space *exponentially* big

→ how do we choose actions?

Marin Bukov

RL Applied to Quantum State Preparation

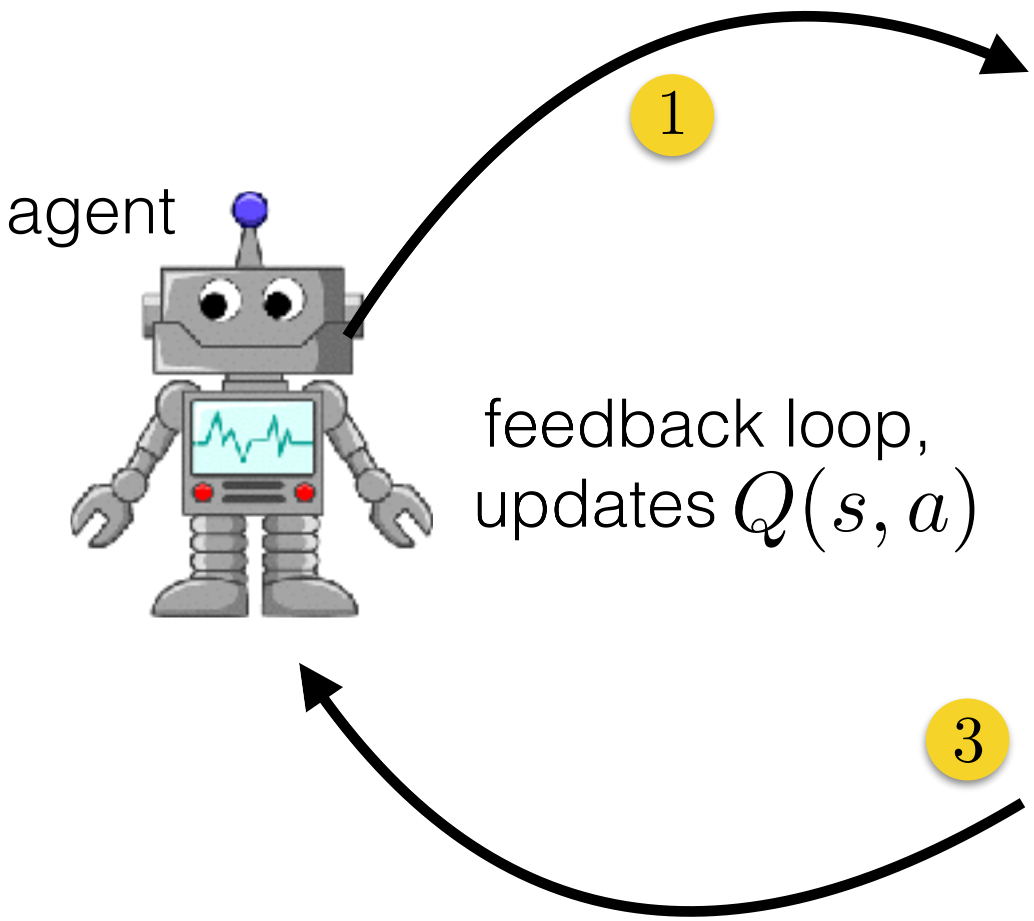

ENVIRONMENT

2

$$H(t) = S^z + h(t)S^x$$

$$|\psi_i\rangle : \text{ GS of } H_i = S^z - S^x$$
$$|\psi_*\rangle : \text{ GS of } H_* = S^z + S^x$$

$$i\partial_t|\psi(t)\rangle = H(t)|\psi(t)\rangle, \; t\in[0,T]$$

reward $r = \begin{cases} 0 & , \; 0\le t < T \\ |\langle\psi_*|\psi(t=T)\rangle|^2, & t = T \end{cases}$

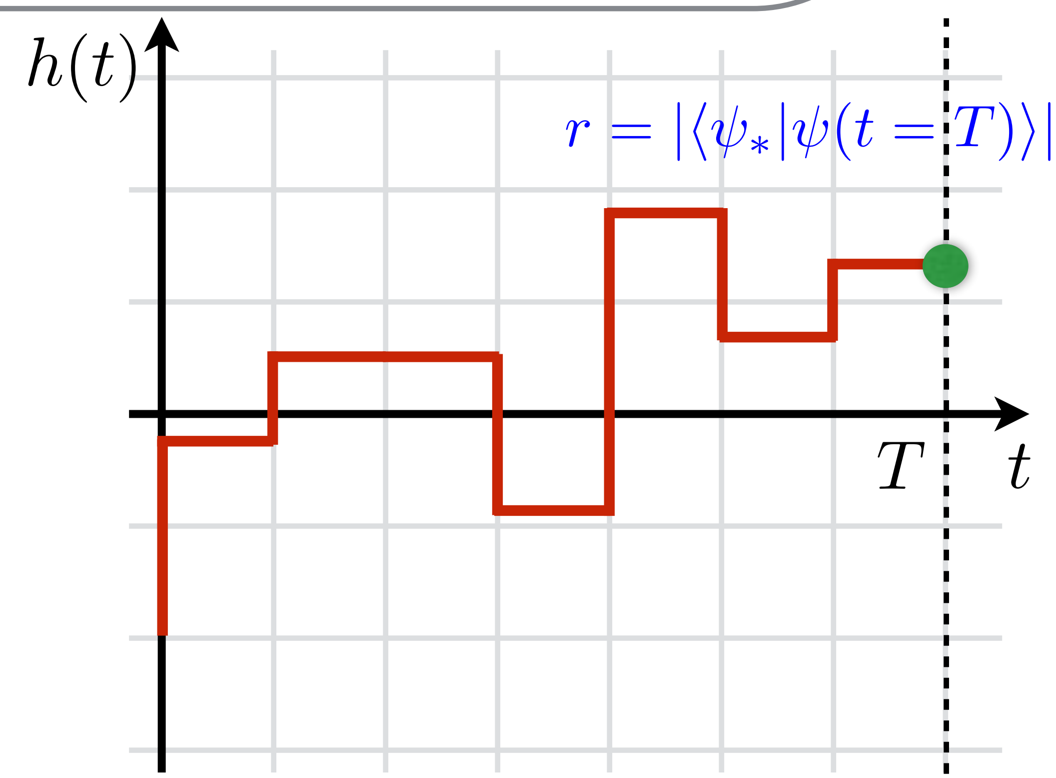

problems:

→ state space *exponentially* big

biased MC sampling

→ how do we choose actions?

Berkeley
UNIVERSITY OF CALIFORNIA

RL Applied to Quantum State Preparation

agent

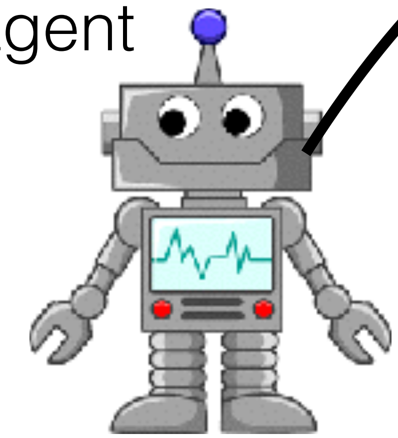

1

feedback loop,
updates $Q(s,a)$

ENVIRONMENT

2

$$H(t) = S^z + h(t)S^x$$

$$|\psi_i\rangle : \ \text{GS of } H_i = S^z - S^x$$

$$|\psi_*\rangle : \ \text{GS of } H_* = S^z + S^x$$

$$i\partial_t|\psi(t)\rangle = H(t)|\psi(t)\rangle, \ t\in[0,T]$$

3

reward $r = \begin{cases} 0 & , \ 0 \le t < T \\ |\langle\psi_*|\psi(t=T)\rangle|^2, & t = T \end{cases}$

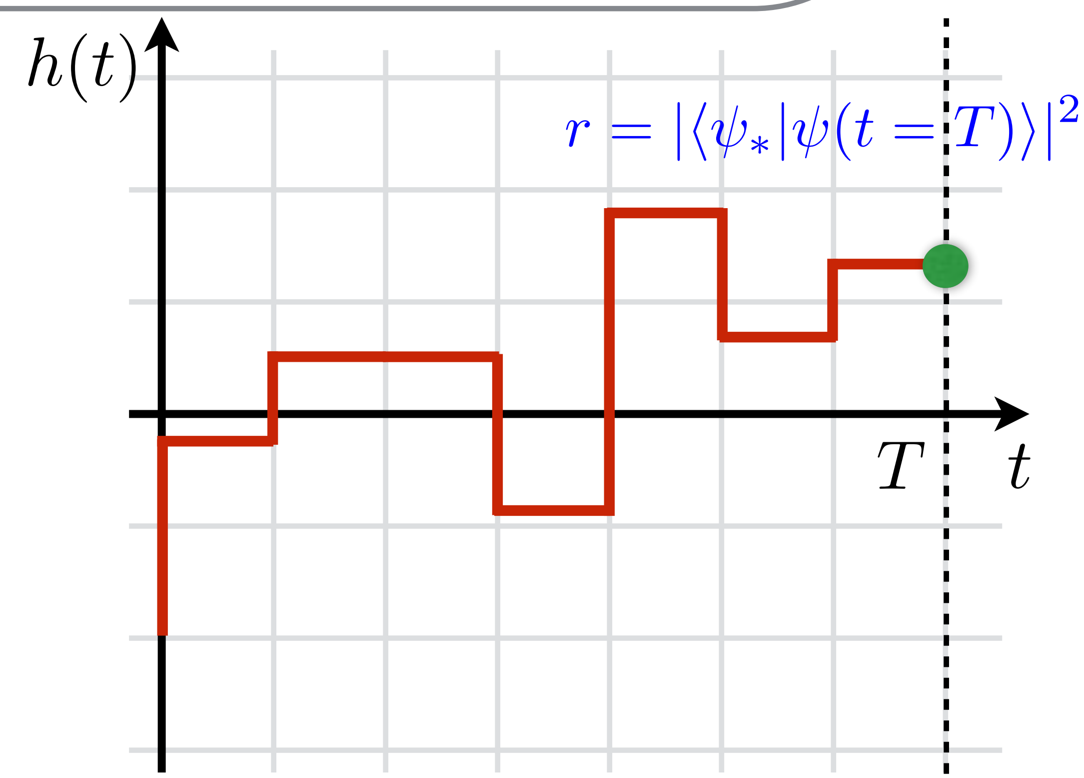

problems:

→ state space *exponentially* big

biased MC sampling

→ how do we choose actions?

exploration exploitation dilemma

Marin Bukov

Reinforcement Learning
Quantum State Preparation

$$H(t) = -S^z - h_x(t)S^x$$

the learning process

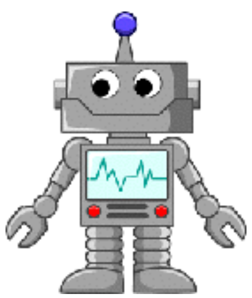

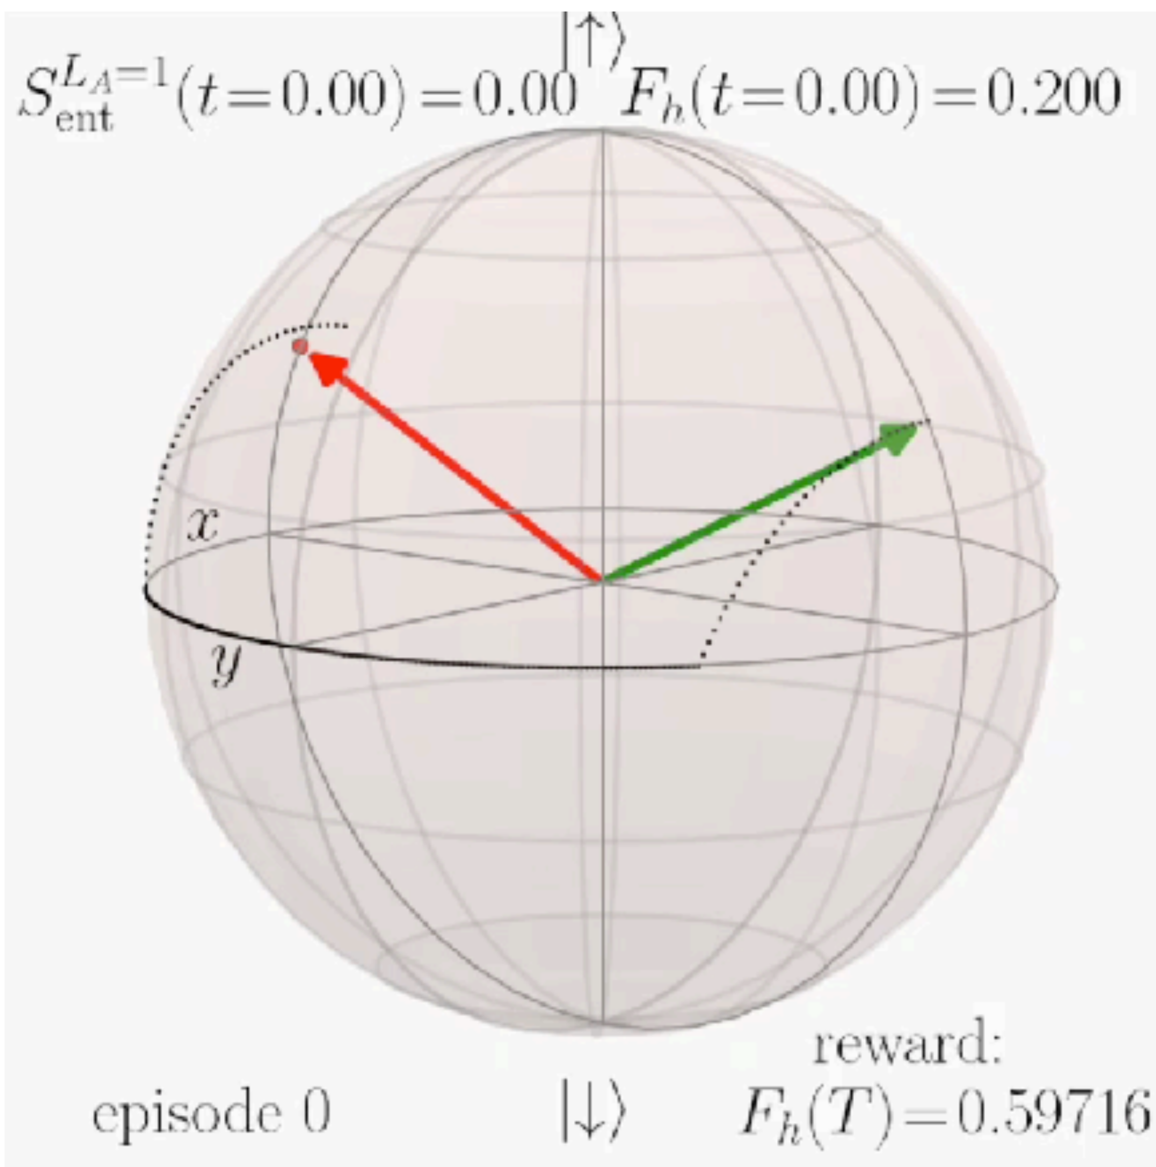

$h \in \{\pm 4\}$ bang-bang protocols

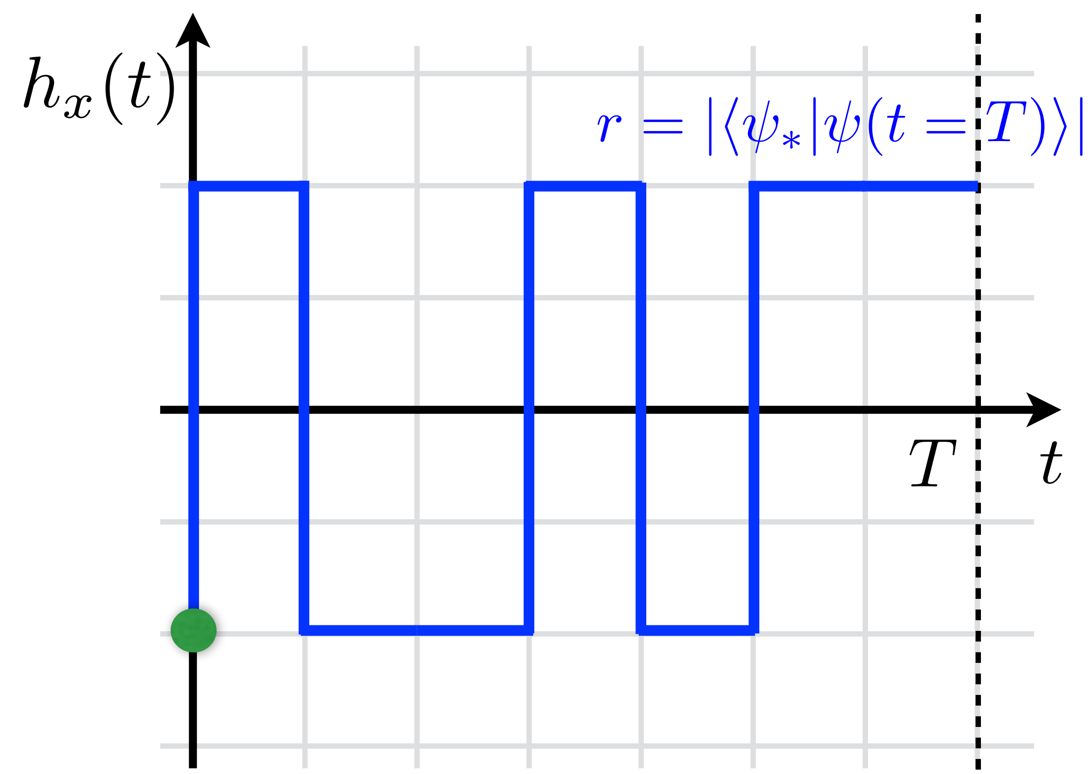

episode completed

Chen et al, IEEE 25, 90 *9920* (2014)

arXiv: *1705.00565* (2017)

Marin Bukov

Reinforcement Learning
Quantum State Preparation

$$H(t) = -S^z - h_x(t) S^x$$

the learning process

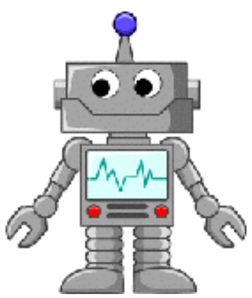

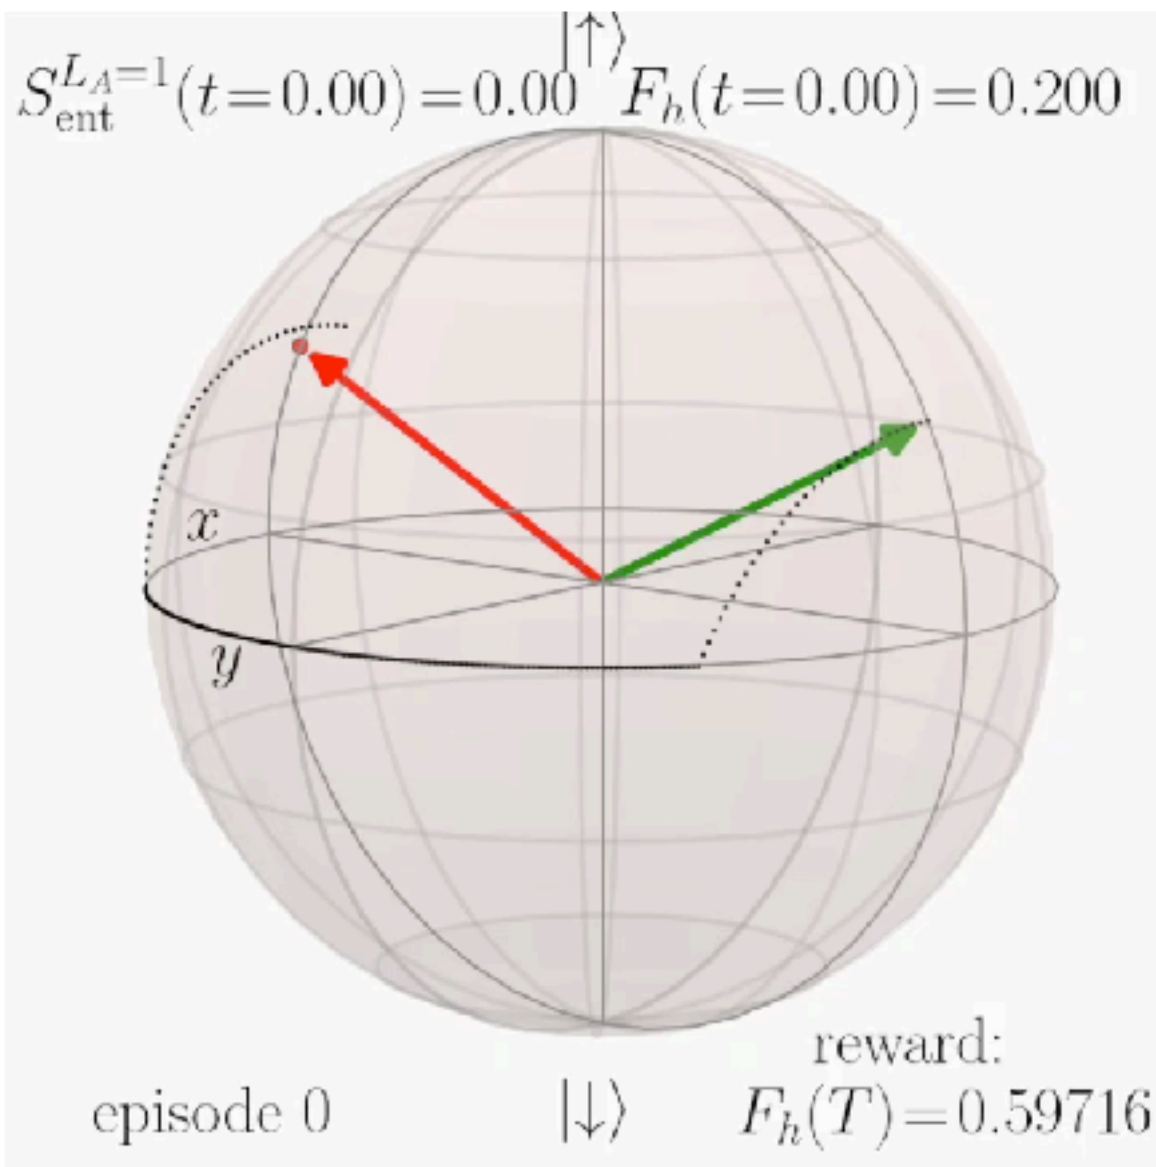

$h \in \{\pm 4\}$ bang-bang protocols

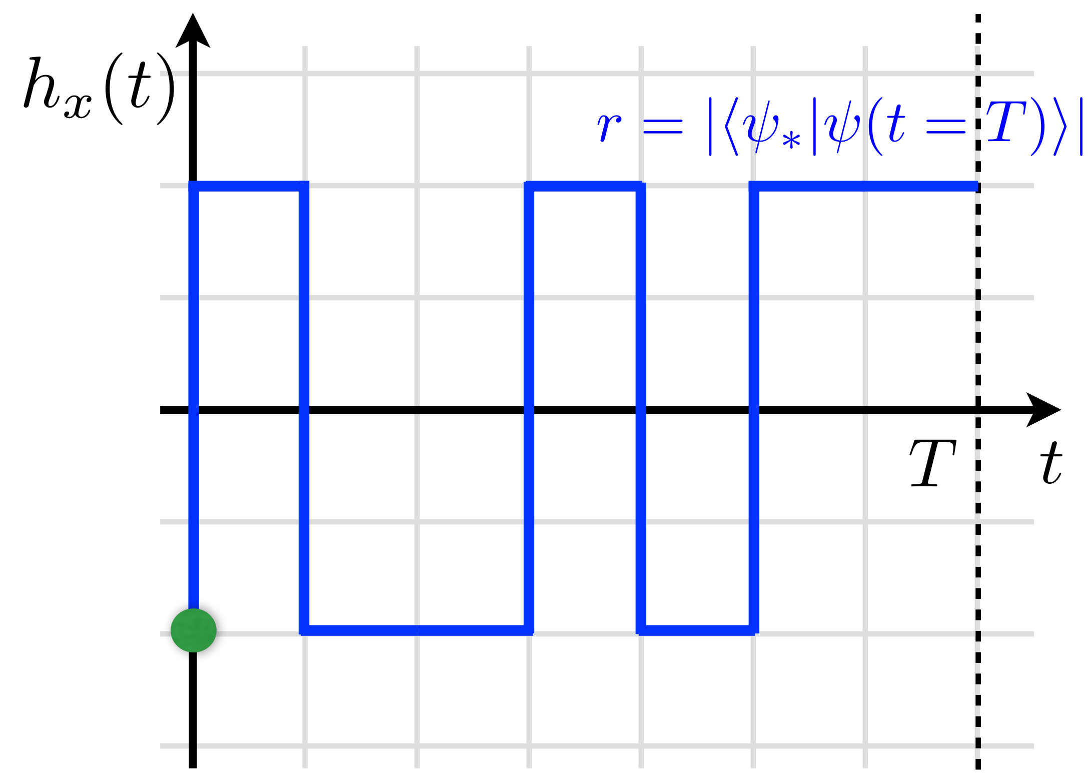

episode completed

Chen et al, IEEE 25, 90 *9920* (2014)

arXiv: *1705.00565* (2017)

Marin Bukov

Berkeley UNIVERSITY OF CALIFORNIA

Reinforcement Learning Quantum State Preparation

$$H(t) = -S^z - h_x(t)S^x$$

the learning process

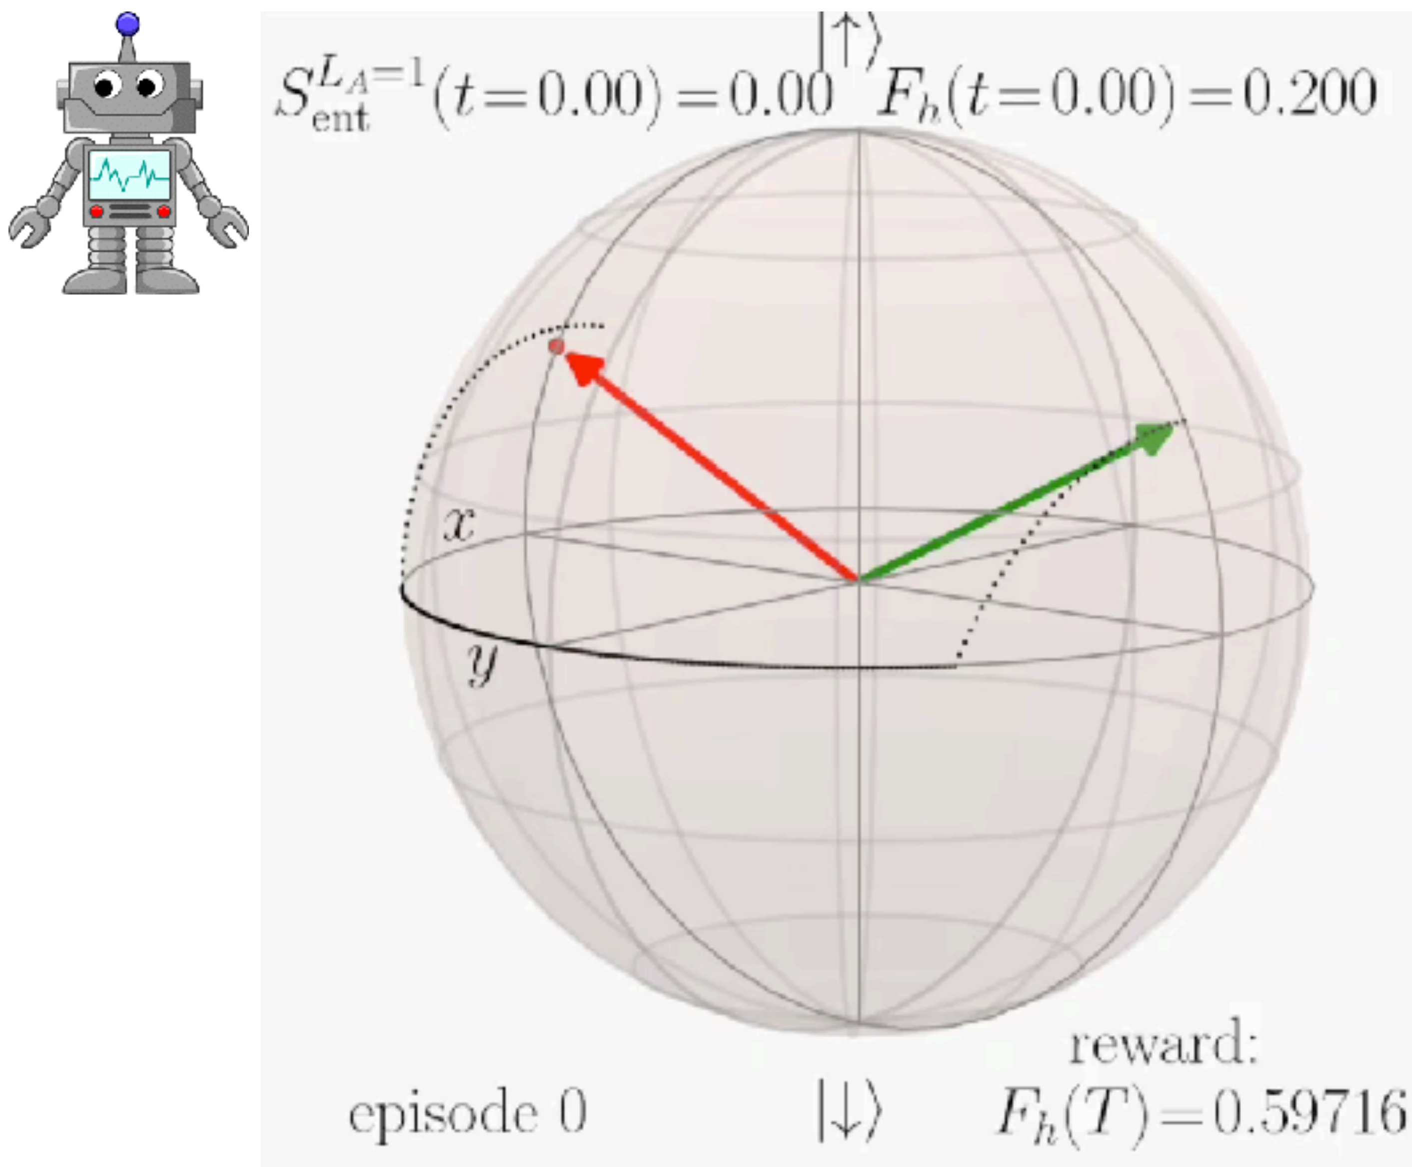

$h \in \{\pm 4\}$ bang-bang protocols

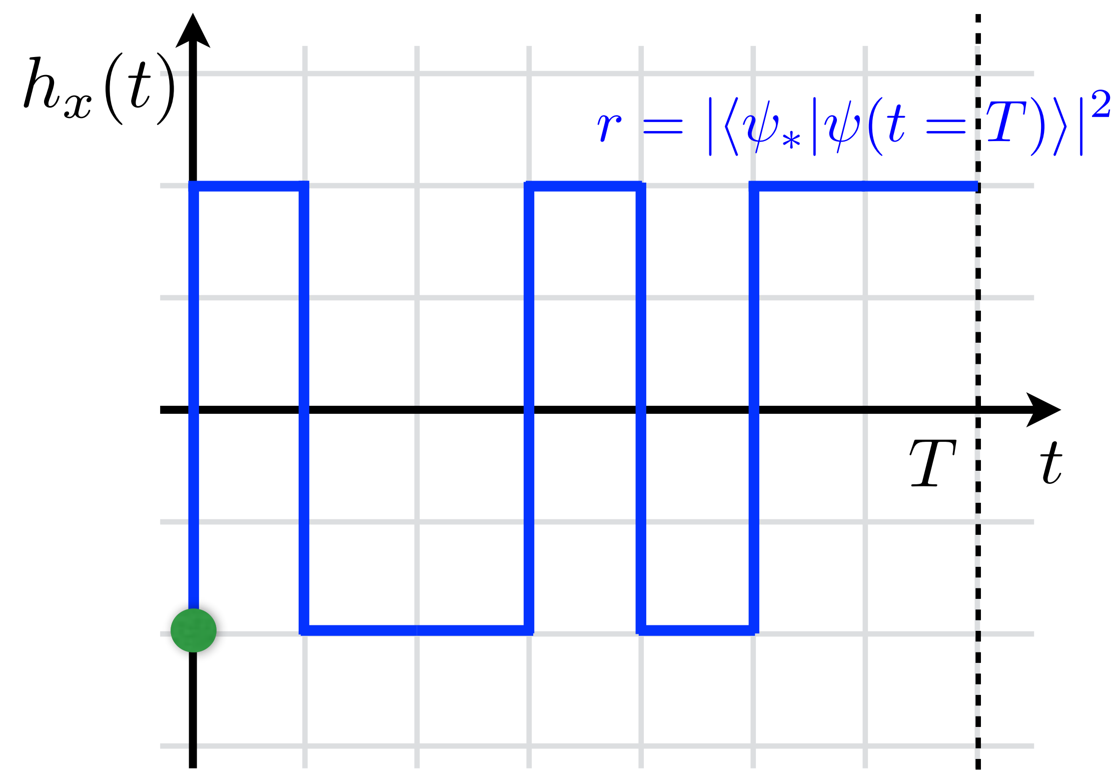

episode completed

Chen et al, IEEE 25, 90 *9920* (2014)
arXiv: *1705.00565* (2017)

How hard is this optimisation problem?

Marin Bukov

Berkeley
UNIVERSITY OF CALIFORNIA

RL-inspired Discovery: Phase Diagram of Quantum Control

$$H(t) = -S^z - h_x(t) S^x$$

bang-bang protocols

$$h \mapsto 1 - F_h(T)$$

RL-inspired Discovery: Phase Diagram of Quantum Control

infidelity landscape (schematic)

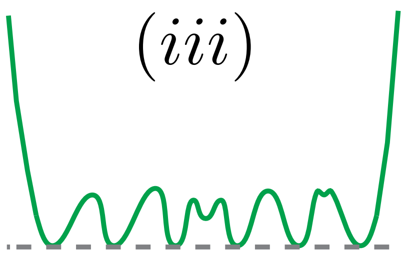

$$H(t) = -S^z - h_x(t)S^x$$

bang-bang protocols

$$h \mapsto 1 - F_h(T)$$

infidelity landscape minima: $\{h^\alpha\}$

RL-inspired Discovery: Phase Diagram of Quantum Control

infidelity landscape (schematic)

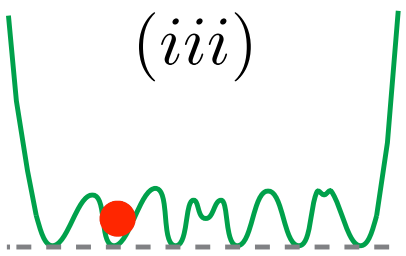

$$H(t) = -S^z - h_x(t)S^x$$

bang-bang protocols

$$h \mapsto 1 - F_h(T)$$

infidelity landscape
minima: $\{h^\alpha\}$

RL-inspired Discovery: Phase Diagram of Quantum Control

infidelity landscape (schematic)

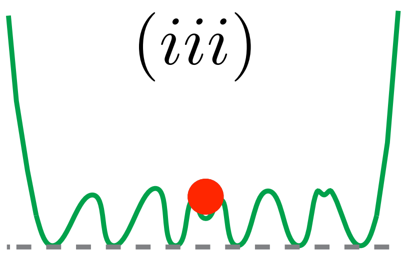

$$H(t) = -S^z - h_x(t)S^x$$

bang-bang protocols

$$h \mapsto 1 - F_h(T)$$

infidelity landscape minima: $\{h^\alpha\}$

RL-inspired Discovery: Phase Diagram of Quantum Control

infidelity landscape (schematic)

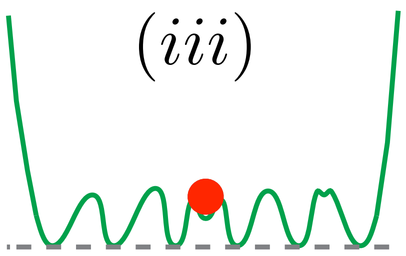

$$H(t) = -S^z - h_x(t)S^x$$

bang-bang protocols

$$h \mapsto 1 - F_h(T)$$

infidelity landscape minima: $\{h^\alpha\}$

$$\bar{h}(t) = \frac{1}{\#_{\mathrm{real}}} \sum_\alpha h^\alpha(t)$$

Edwards-Anderson-like order parameter:

$$q(T) \sim \overline{\sum_t \{h(t) - \bar{h}(t)\}^2}$$

RL-inspired Discovery: Phase Diagram of Quantum Control

infidelity landscape (schematic)

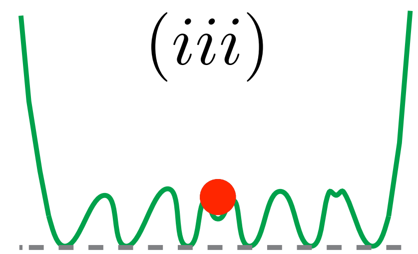

$$H(t) = -S^z - h_x(t)S^x$$

bang-bang protocols

$$h \mapsto 1 - F_h(T)$$

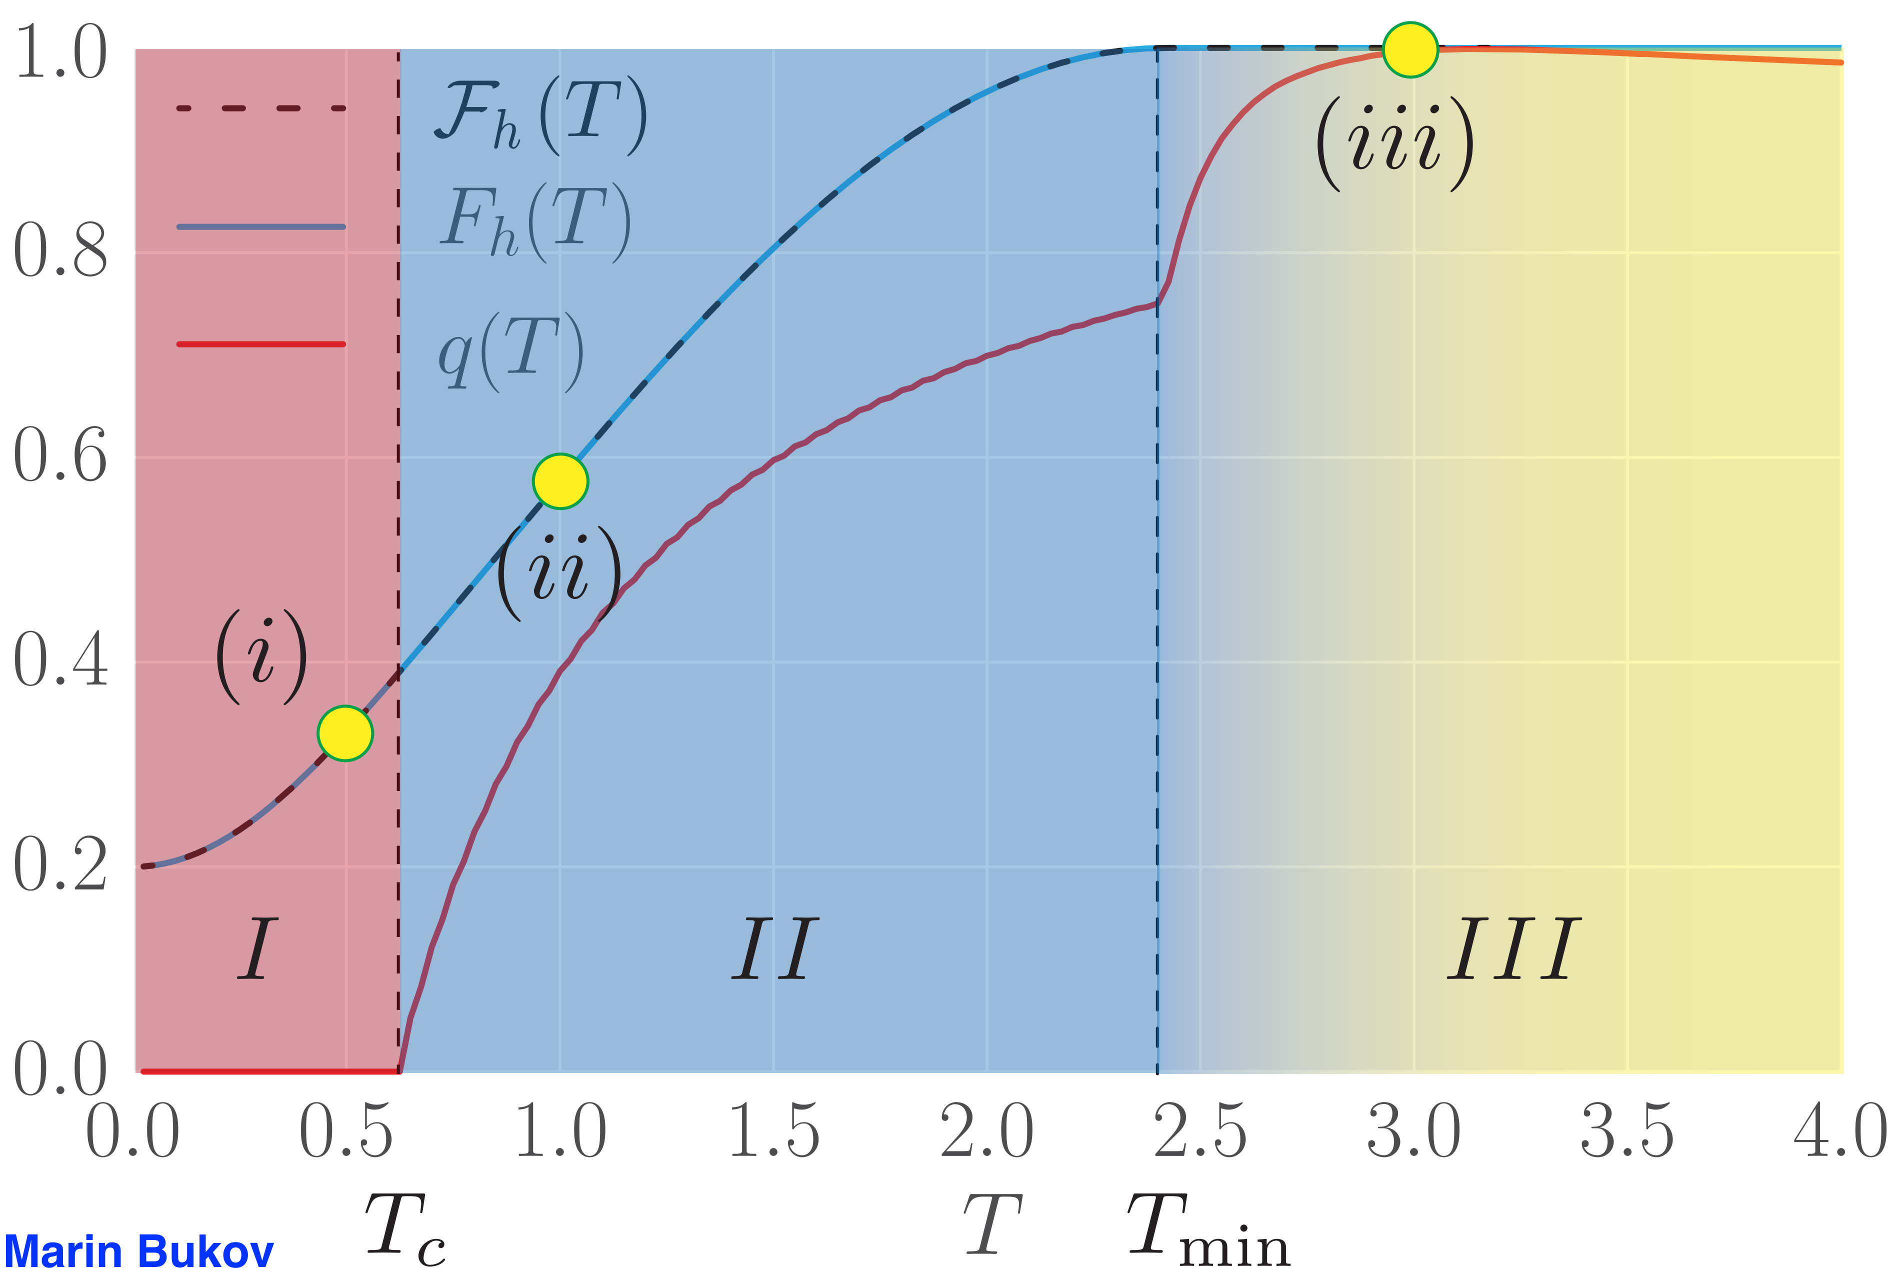

infidelity landscape minima: $\{h^\alpha\}$

$$\bar{h}(t) = \frac{1}{\#_{\mathrm{real}}} \sum_\alpha h^\alpha(t)$$

Edwards-Anderson-like order parameter:

$$q(T) \sim \overline{\sum_t \{h(t) - \bar{h}(t)\}^2}$$

Marin Bukov

arXiv: *1705.00565* (2017)

RL-inspired Discovery: Phase Diagram of Quantum Control

infidelity landscape (schematic)

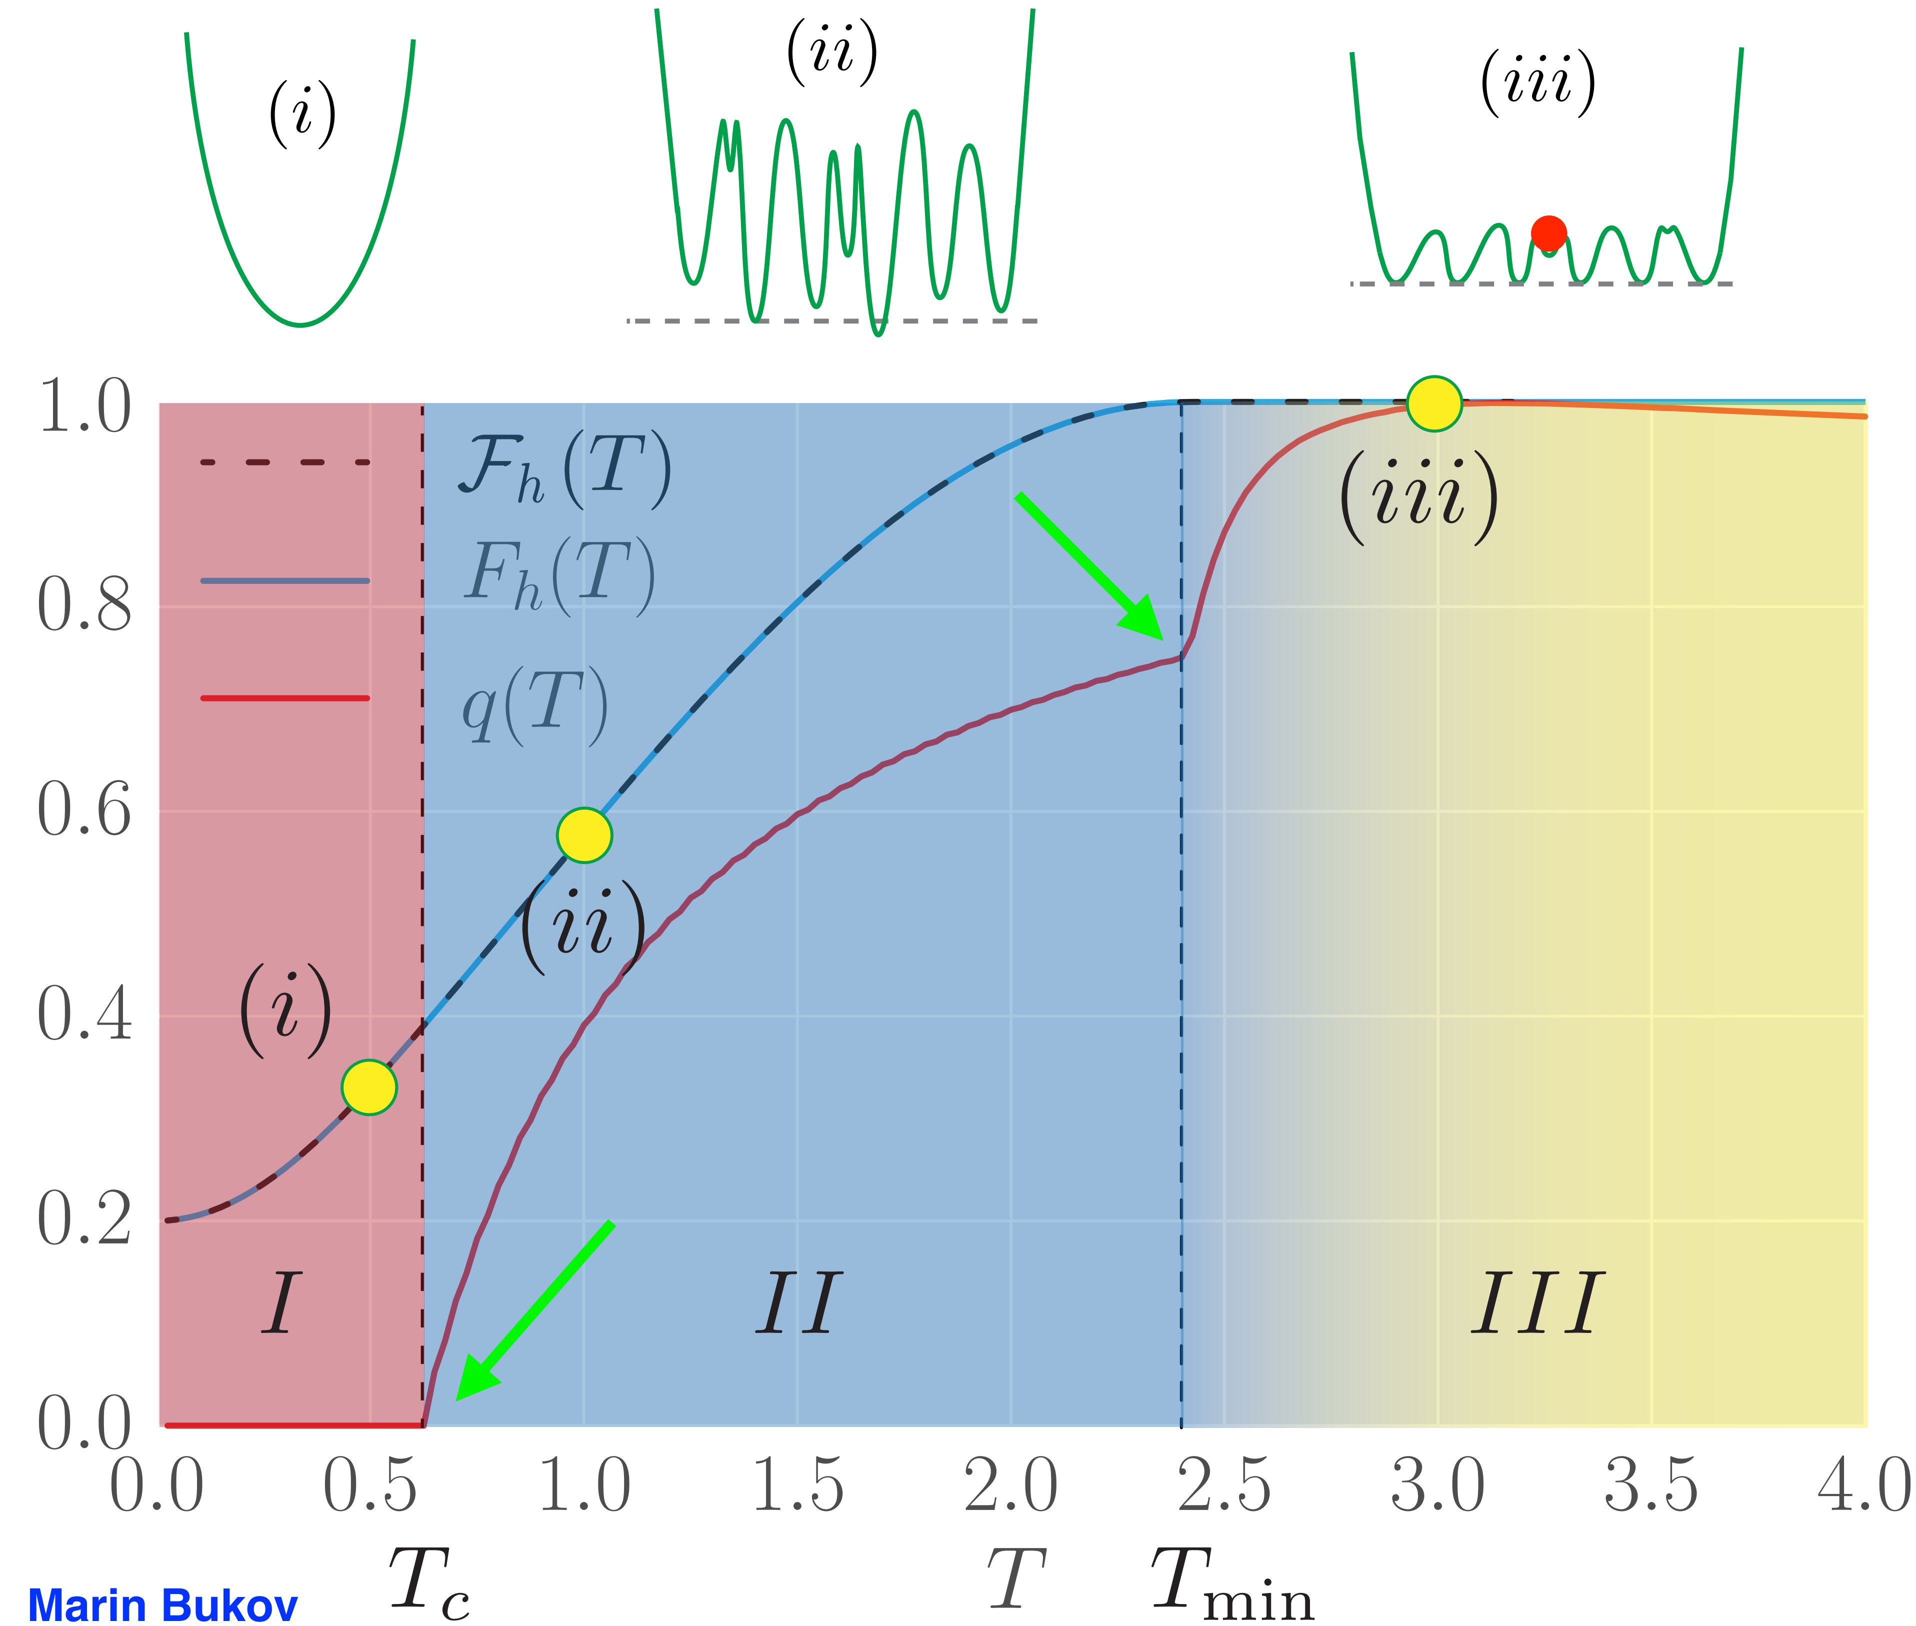

$$H(t) = -S^z - h_x(t)S^x$$

bang-bang protocols

$$h \mapsto 1 - F_h(T)$$

infidelity landscape minima: $\{h^\alpha\}$

$$\bar{h}(t) = \frac{1}{\#_{\text{real}}}\sum_\alpha h^\alpha(t)$$

Edwards-Anderson-like order parameter:

$$q(T) \sim \overline{\sum_t \{h(t) - \bar{h}(t)\}^2}$$

Marin Bukov

arXiv: *1705.00565* (2017)

Berkeley UNIVERSITY OF CALIFORNIA

Nature of Control Phase Transitions

$$H(t) = -S^z - h_x(t)S^x$$

Nature of Control Phase Transitions

$$H(t) = -S^z - h_x(t)S^x$$

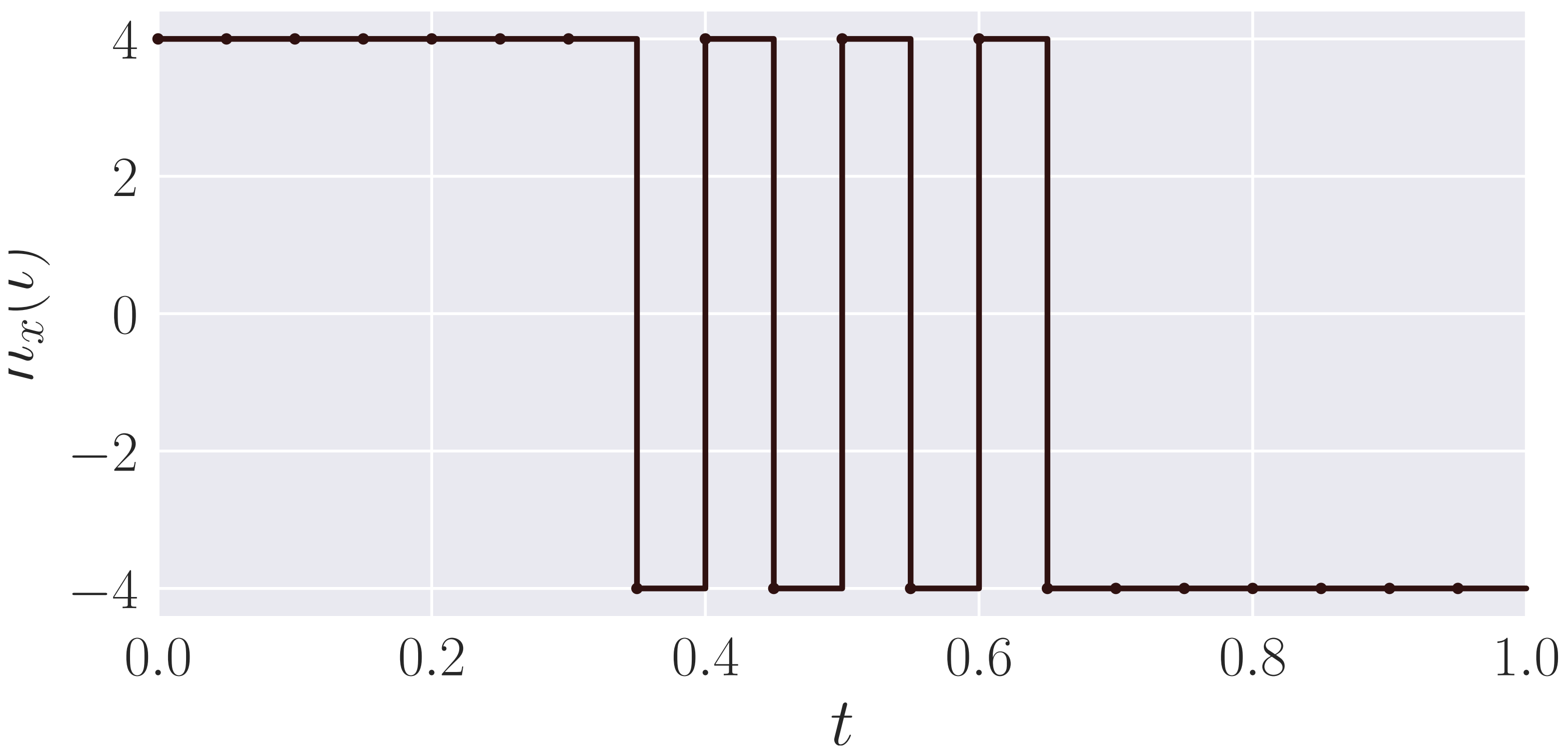

Marin Bukov

Berkeley
UNIVERSITY OF CALIFORNIA

Nature of Control Phase Transitions

$$H(t) = -S^z - h_x(t) S^x$$

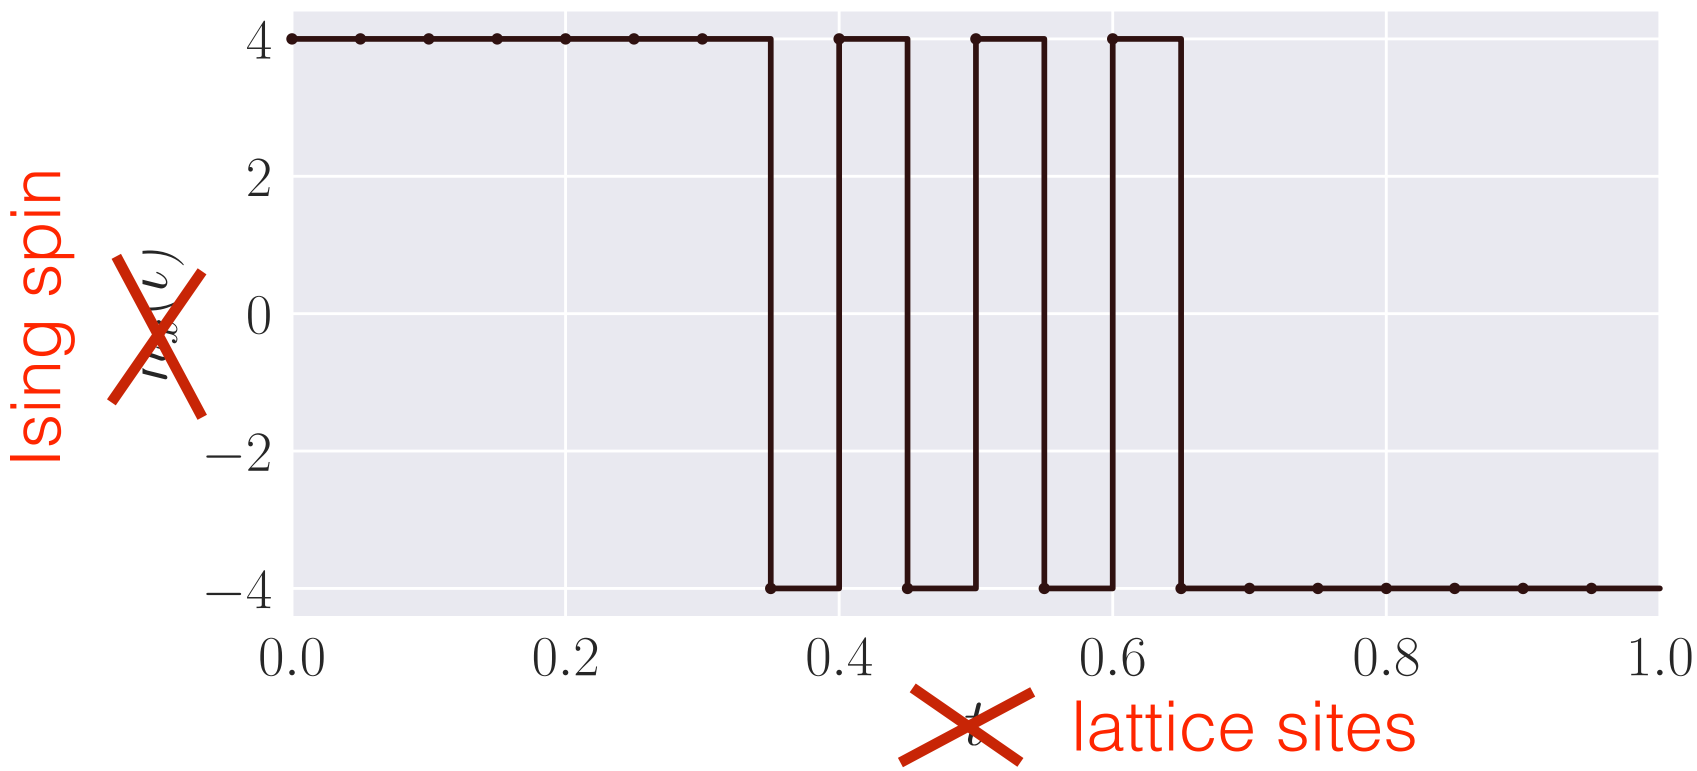

Marin Bukov

Nature of Control Phase Transitions

$$H(t) = -S^z - h_x(t)S^x$$

→ one-to-one correspondence:

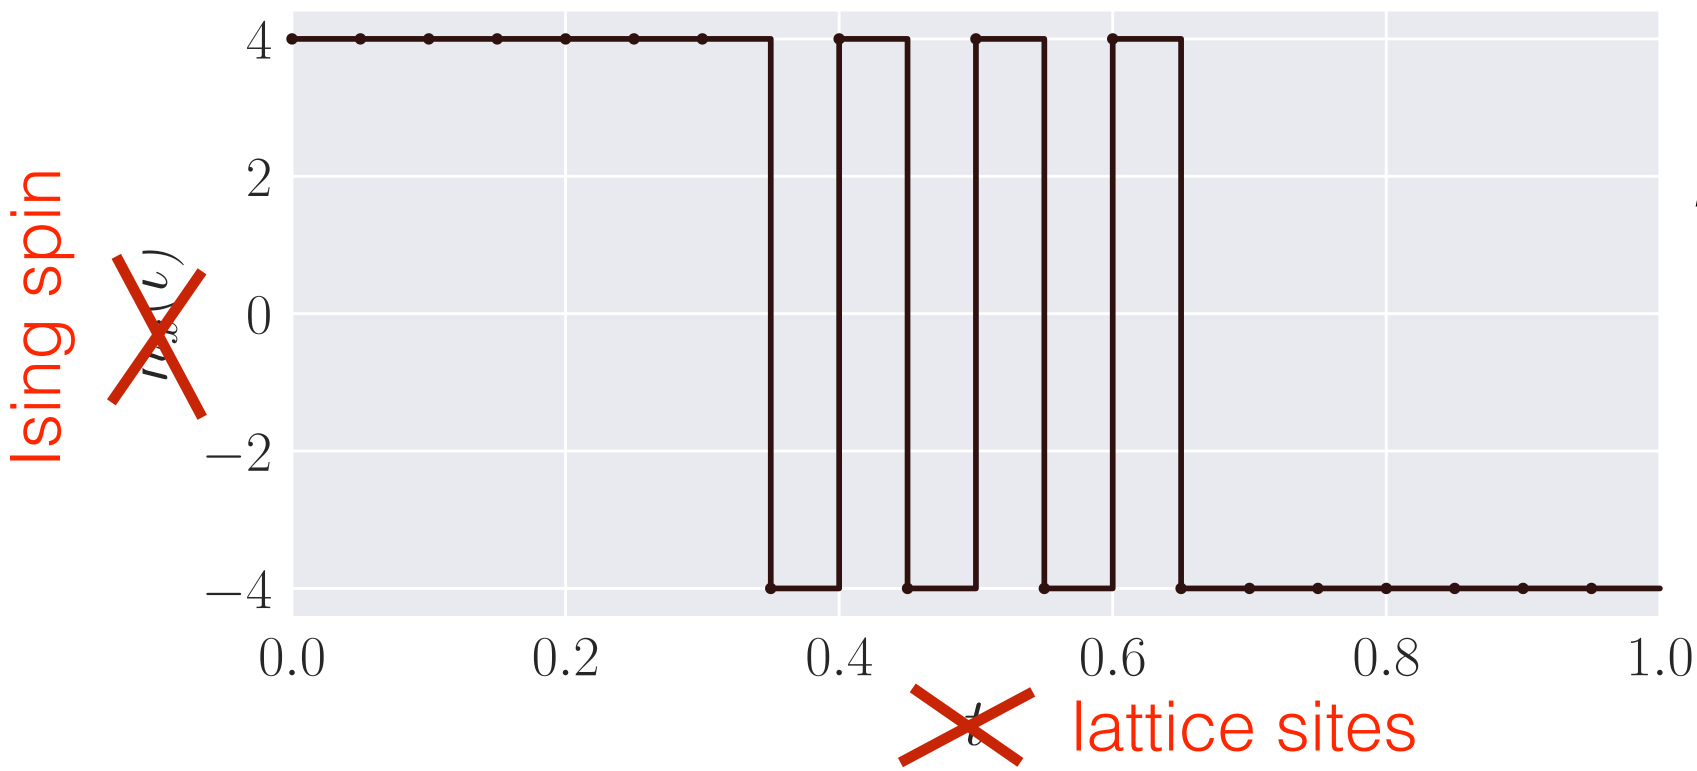

bang-bang protocol ↔ ***classical*** *spin state*

infidelity ↔ *energy*

Marin Bukov

Berkeley
UNIVERSITY OF CALIFORNIA

Nature of Control Phase Transitions

$$H(t) = -S^z - h_x(t) S^x$$

→ one-to-one correspondence:

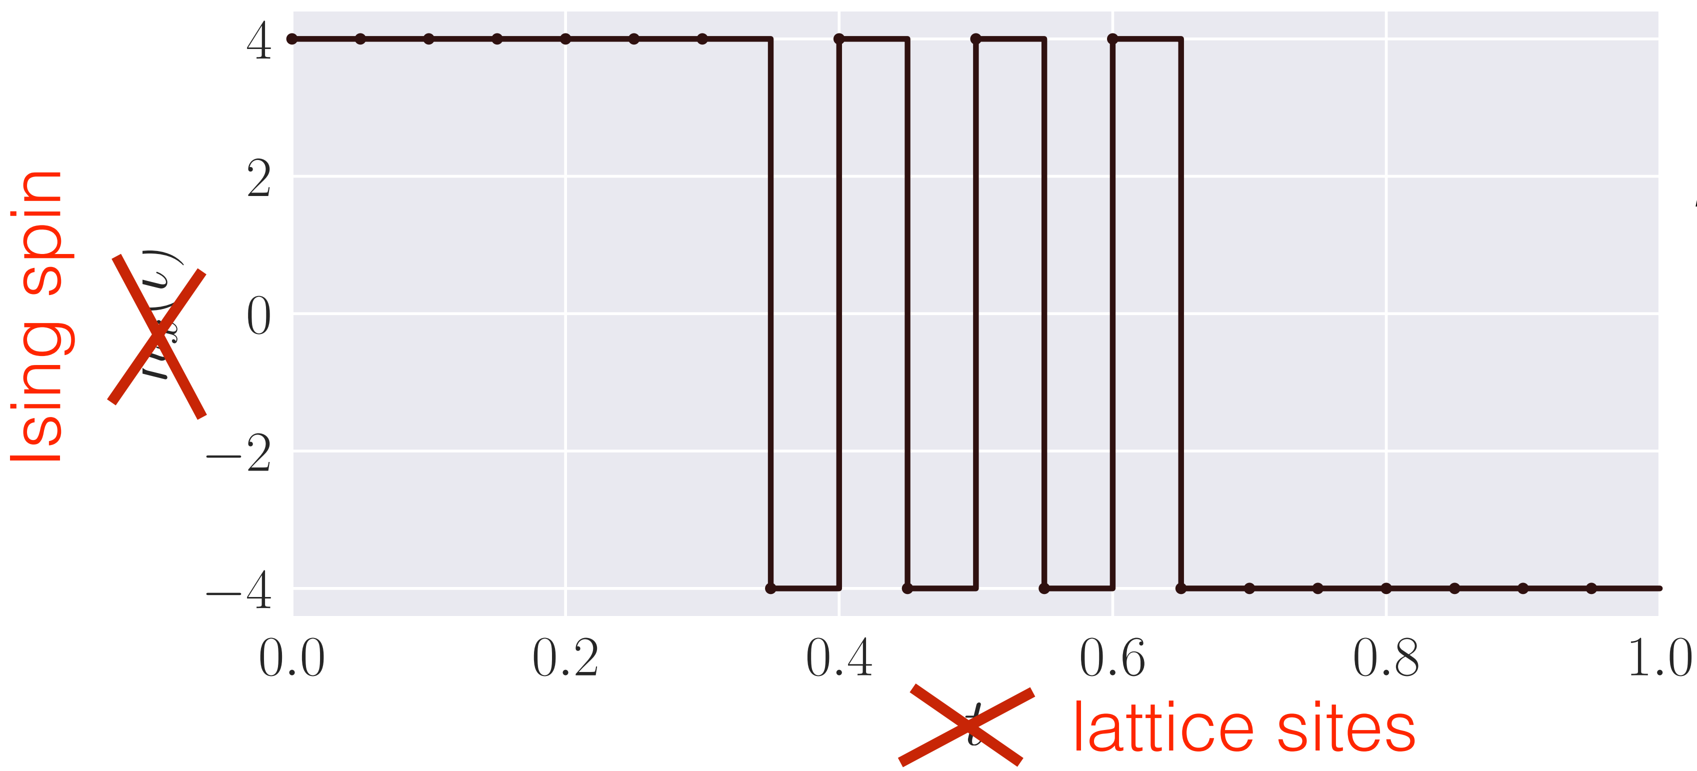

bang-bang protocol ⟷ ***classical*** *spin state*

infidelity ⟷ *energy*

→ effective *classical* energy function governs control phase transitions

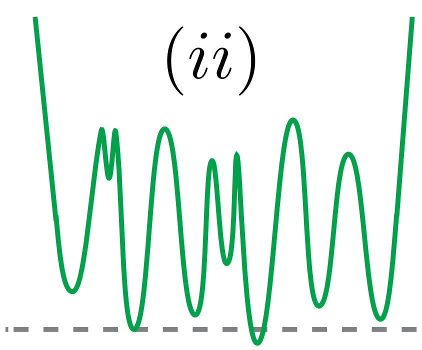

$$\mathcal{H}_{\text{eff}}(T) = I(T) + \sum_j G_j(T) h_j + \sum_{ij} J_{ij}(T) h_i h_j + \sum_{ijk} K_{ijk}(T) h_i h_j h_k + \ldots$$

j : sites on time lattice

Berkeley UNIVERSITY OF CALIFORNIA

Nature of Control Phase Transitions

$$H(t) = -S^z - h_x(t)S^x$$

→ one-to-one correspondence:

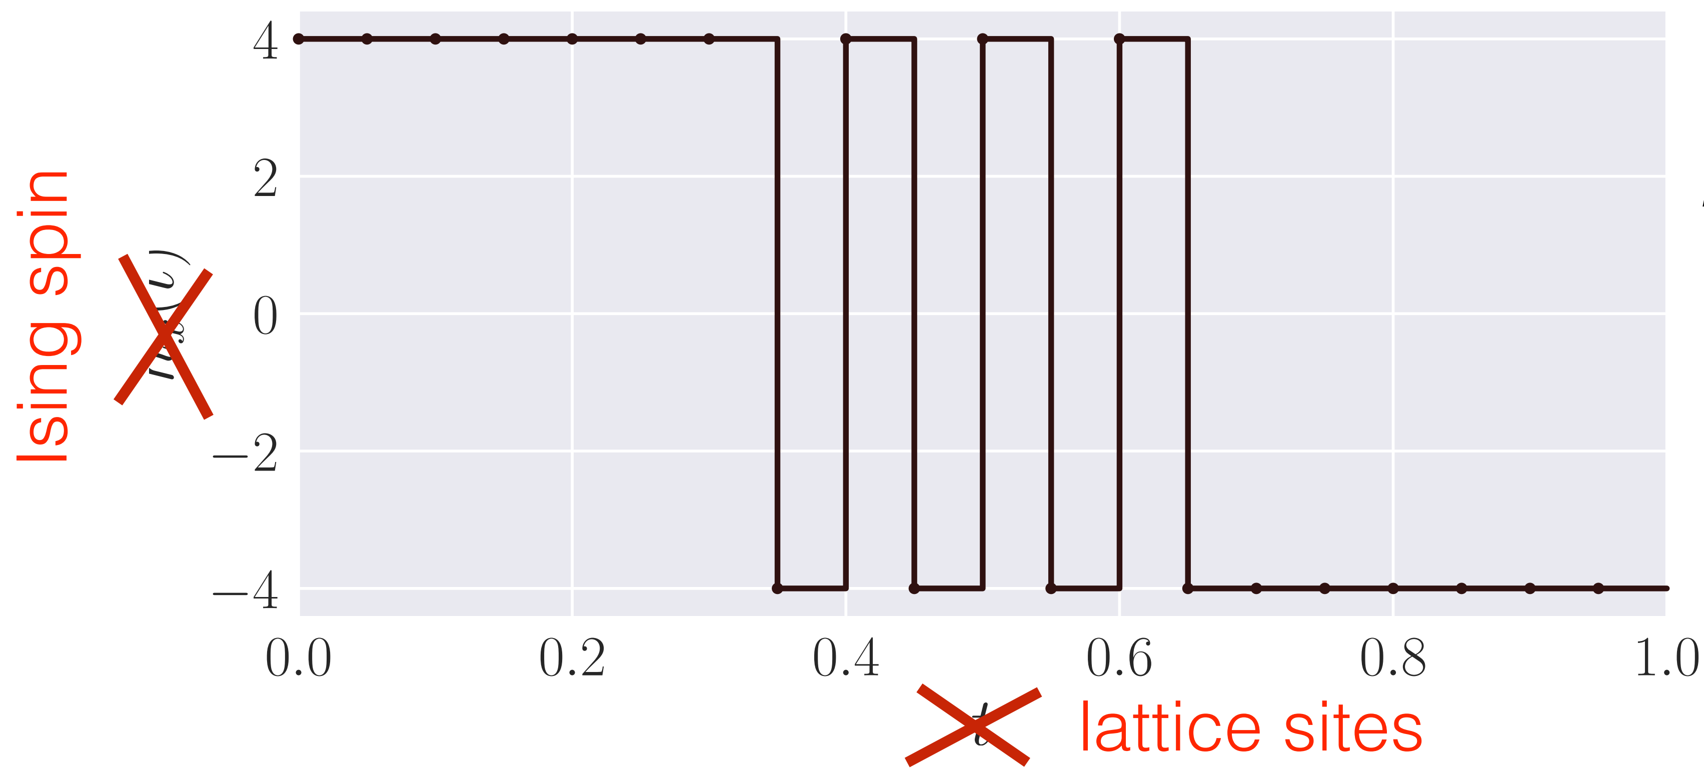

bang-bang protocol ↔ ***classical** spin state*

infidelity ↔ *energy*

→ effective *classical* energy function governs control phase transitions

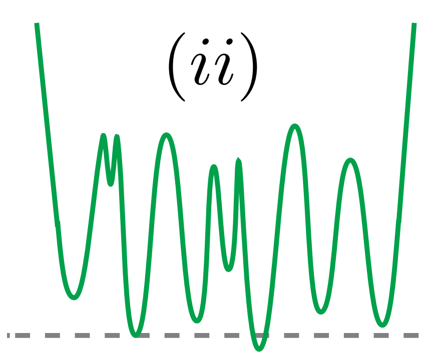

$$\mathcal{H}_{\text{eff}}(T) = I(T) + \sum_j G_j(T)h_j + \sum_{ij} J_{ij}(T)h_ih_j + \sum_{ijk} K_{ijk}(T)h_ih_jh_k + \ldots$$

$j:$ sites on time lattice

→ control phase transitions: *classical (?)*, ***non-equilibrium***

Marin Bukov

Berkeley
UNIVERSITY OF CALIFORNIA

Outlook

→ one can teach a reinforcement learning agent to prepare quantum states at short times with high fidelity

→ finding optimal driving protocol as hard as searching for absolute GS of a spin glass (even if system is disorder-free)

→ quantum control problems have extremely rich phase diagrams with overconstrained, controllable,
correlated and glassy phases —> POSTER (Alex Day)
exhibit symmetry breaking —> POSTER (MB)

→ control phase transitions: classical & nonequilibrium, generic?

Berkeley
UNIVERSITY OF CALIFORNIA

BOSTON
UNIVERSITY

GORDON AND BETTY
MOORE
FOUNDATION

arXiv: *1705.00565* (2017)
arXiv: *1711.09109* (2017)

web: *mgbukov.github.io*

Outlook

QuSpin: weinbe58.github.io/QuSpin/

open-source Python package *for ED and quantum dynamics of arbitrary boson, fermion and spin many-body systems, supporting various (user-defined) symmetries and time evolution.*

SciPost Phys. 2, 003 (2017)

web: *mgbukov.github.io*

arXiv: *1705.00565* (2017)
arXiv: *1711.09109* (2017)