From Reinforcement Learning to Spin Glasses
The many surprises in Quantum State preparation

Phys. Rev. X 8, 031086 (2018) , Physical Review A, 97, 052114 (2018), arXiv:1803.10856

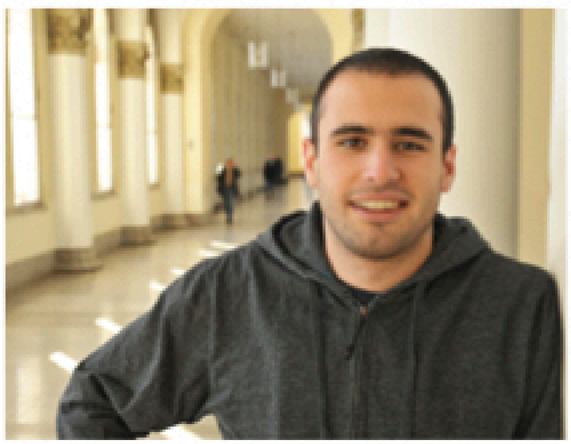

Marin Bukov

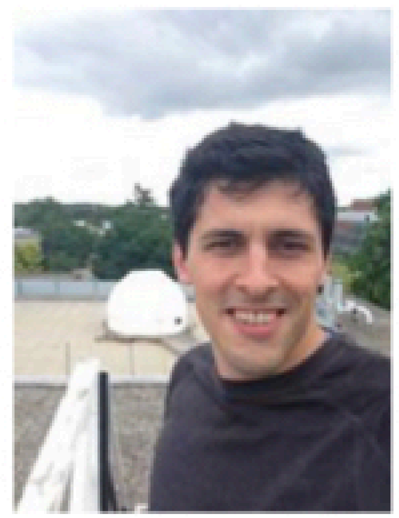

Alex Day

Dries Sels

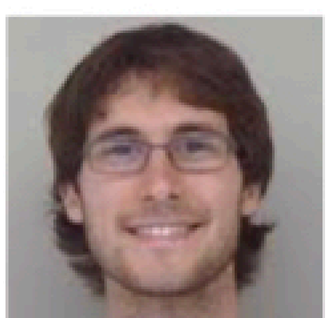

Phil Weinberg

A. Polkovnikov

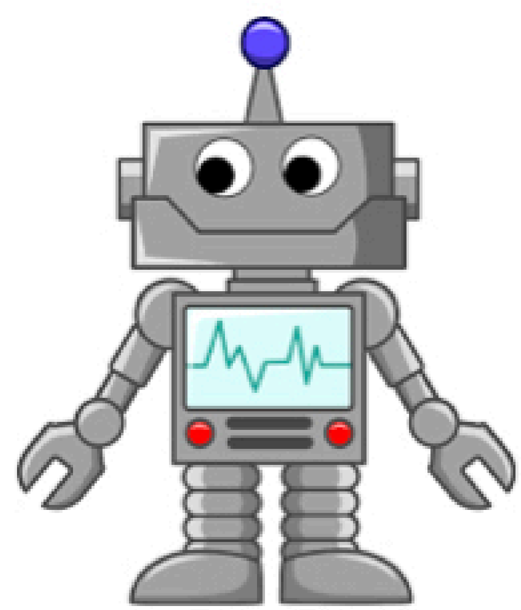

A high-bias, low-variance introduction to Machine Learning for physicists

Pankaj Mehta, Ching-Hao Wang, Alexandre G. R. Day, and Clint Richardson

Department of Physics, Boston University, Boston, MA 02215, USA[*]

Marin Bukov

Department of Physics, University of California, Berkeley, CA 94720, USA

Charles K. Fisher

Unlearn.AI, San Francisco, CA 94108

David J. Schwab

Initiative for the Theoretical Sciences, The Graduate Center, City University of New York, 365 Fifth Ave., New York, NY 10016

(Dated: March 26, 2018)

Notebooks

A zip file containing all notebooks can be downloaded here. Individual notebooks can be downloaded below. We also include links to html versions of the notebook. **RBM notebooks have been updated on 6/6/18.**

1. **Section II: Machine Learning is difficult** python html
2. **Section IV: Gradient Descent** python html
3. **Section VI: Linear Regression (Diabetes)** python html
4. **Section VI: Linear Regression (Ising)** python html
5. **Section VII: Logistic Regression (SUSY)** python html
6. **Section VII: Logistic Regression (Ising)** python html
7. **Section VII: Logistic Regression (MNIST)** python html
8. **Section VII: Bagging** python html
9. **Section VIII: Random Forests (Ising)**python html
10. **Section VIII: XGBoost (SUSY)** python html
11. **Section IX: Keras DNN (MNIST)** python html
12. **Section IX: TensorFlow DNN (Ising)** python html
13. **Section IX: Pytorch DNN (SUSY)** python html
14. **Section X: Pytorch CNN (Ising)** python html
15. **Section XII: Clustering** python html
16. **Section XIV: Expectation Maximization** python html
17. **Section XVI: Restricted Boltzmann Machines (MNIST)** python html
18. **Section XVI: Restricted Boltzmann Machines (Ising)** python html
19. **Section XVII: Keras Variational Autoencoders (MNIST)** python html
20. **Section XVII: Keras Variational Autoencoders (Ising)** python html

118 double column pages
20+ Python Notebooks
87 figures

ArXiv: 1803:08823

physics.bu.edu/~pankajm/MLnotebooks.html

Amazing what progress we have made in 4 years! Maybe I should have more hope

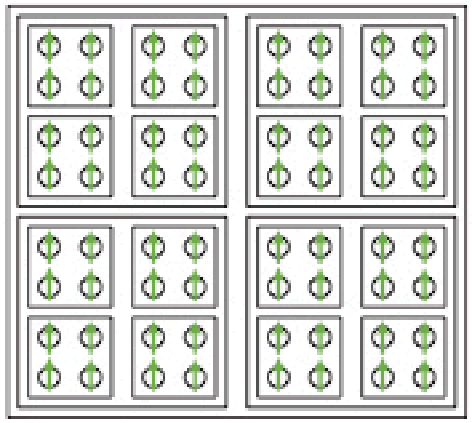

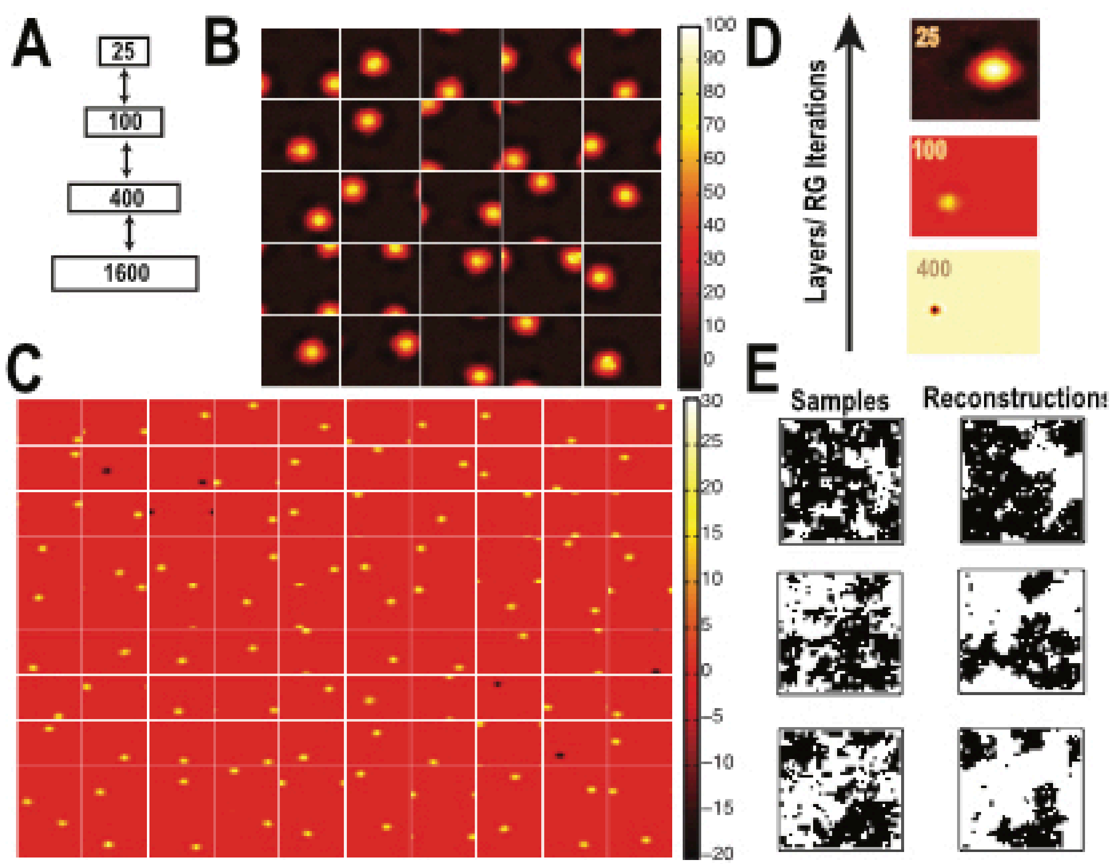

ArXiv: 1410.3831

Recently, it was suggested that modern RG techniques developed in the context of quantum systems such as matrix product states and tensor networks have a natural interpretation in terms of variational RG [17]. These new techniques exploit ideas such as entanglement entropy and disentanglers which create a features with a minimum amount of redundancy. It is an open question to see whether these ideas can be imported into deep learning algorithms. Our mapping also suggests a route for applying real space renormalization techniques to complicated physical systems.

Quantum State Preparations

- Cold atoms: preparing Bose-Einstein condensates
  - Paul Wigley et al., Nature, 2016

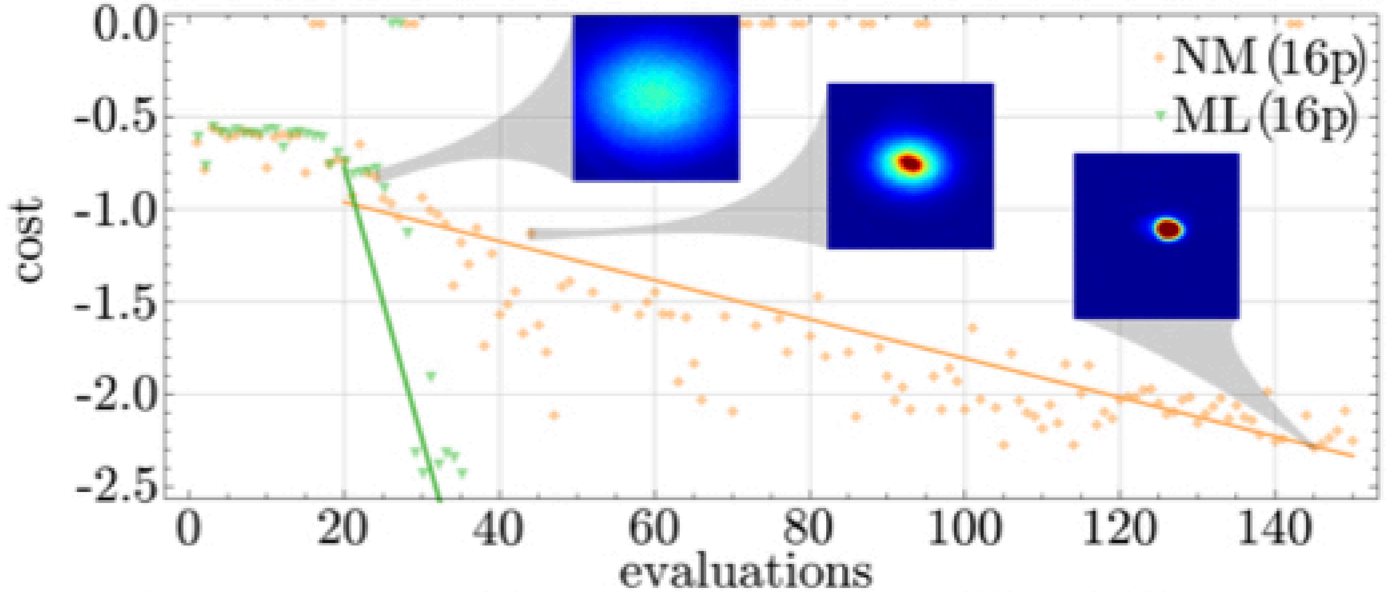

- Quantum computing with SQUIDS, NV centers, etc.
  - Brian B. Zhou et al. Nature Physics, 2016
  - R. Barends et al. Nature, 2016
- (perhaps soon) Quantum machine learning (HHL)
  - P. Rebenfrost et al., PRL 2014

Requires preparing complicated quantum states with a robust and fast procedure

The quantum state preparation problem

GOAL: **Given some constraints (i.e. *T*) determine the Hamiltonian (*H*(*t*)) such that the overlap to a target state is maximized:**

$$F[H(t)] = |\langle \psi(T)|\psi^\star\rangle\,|^2$$

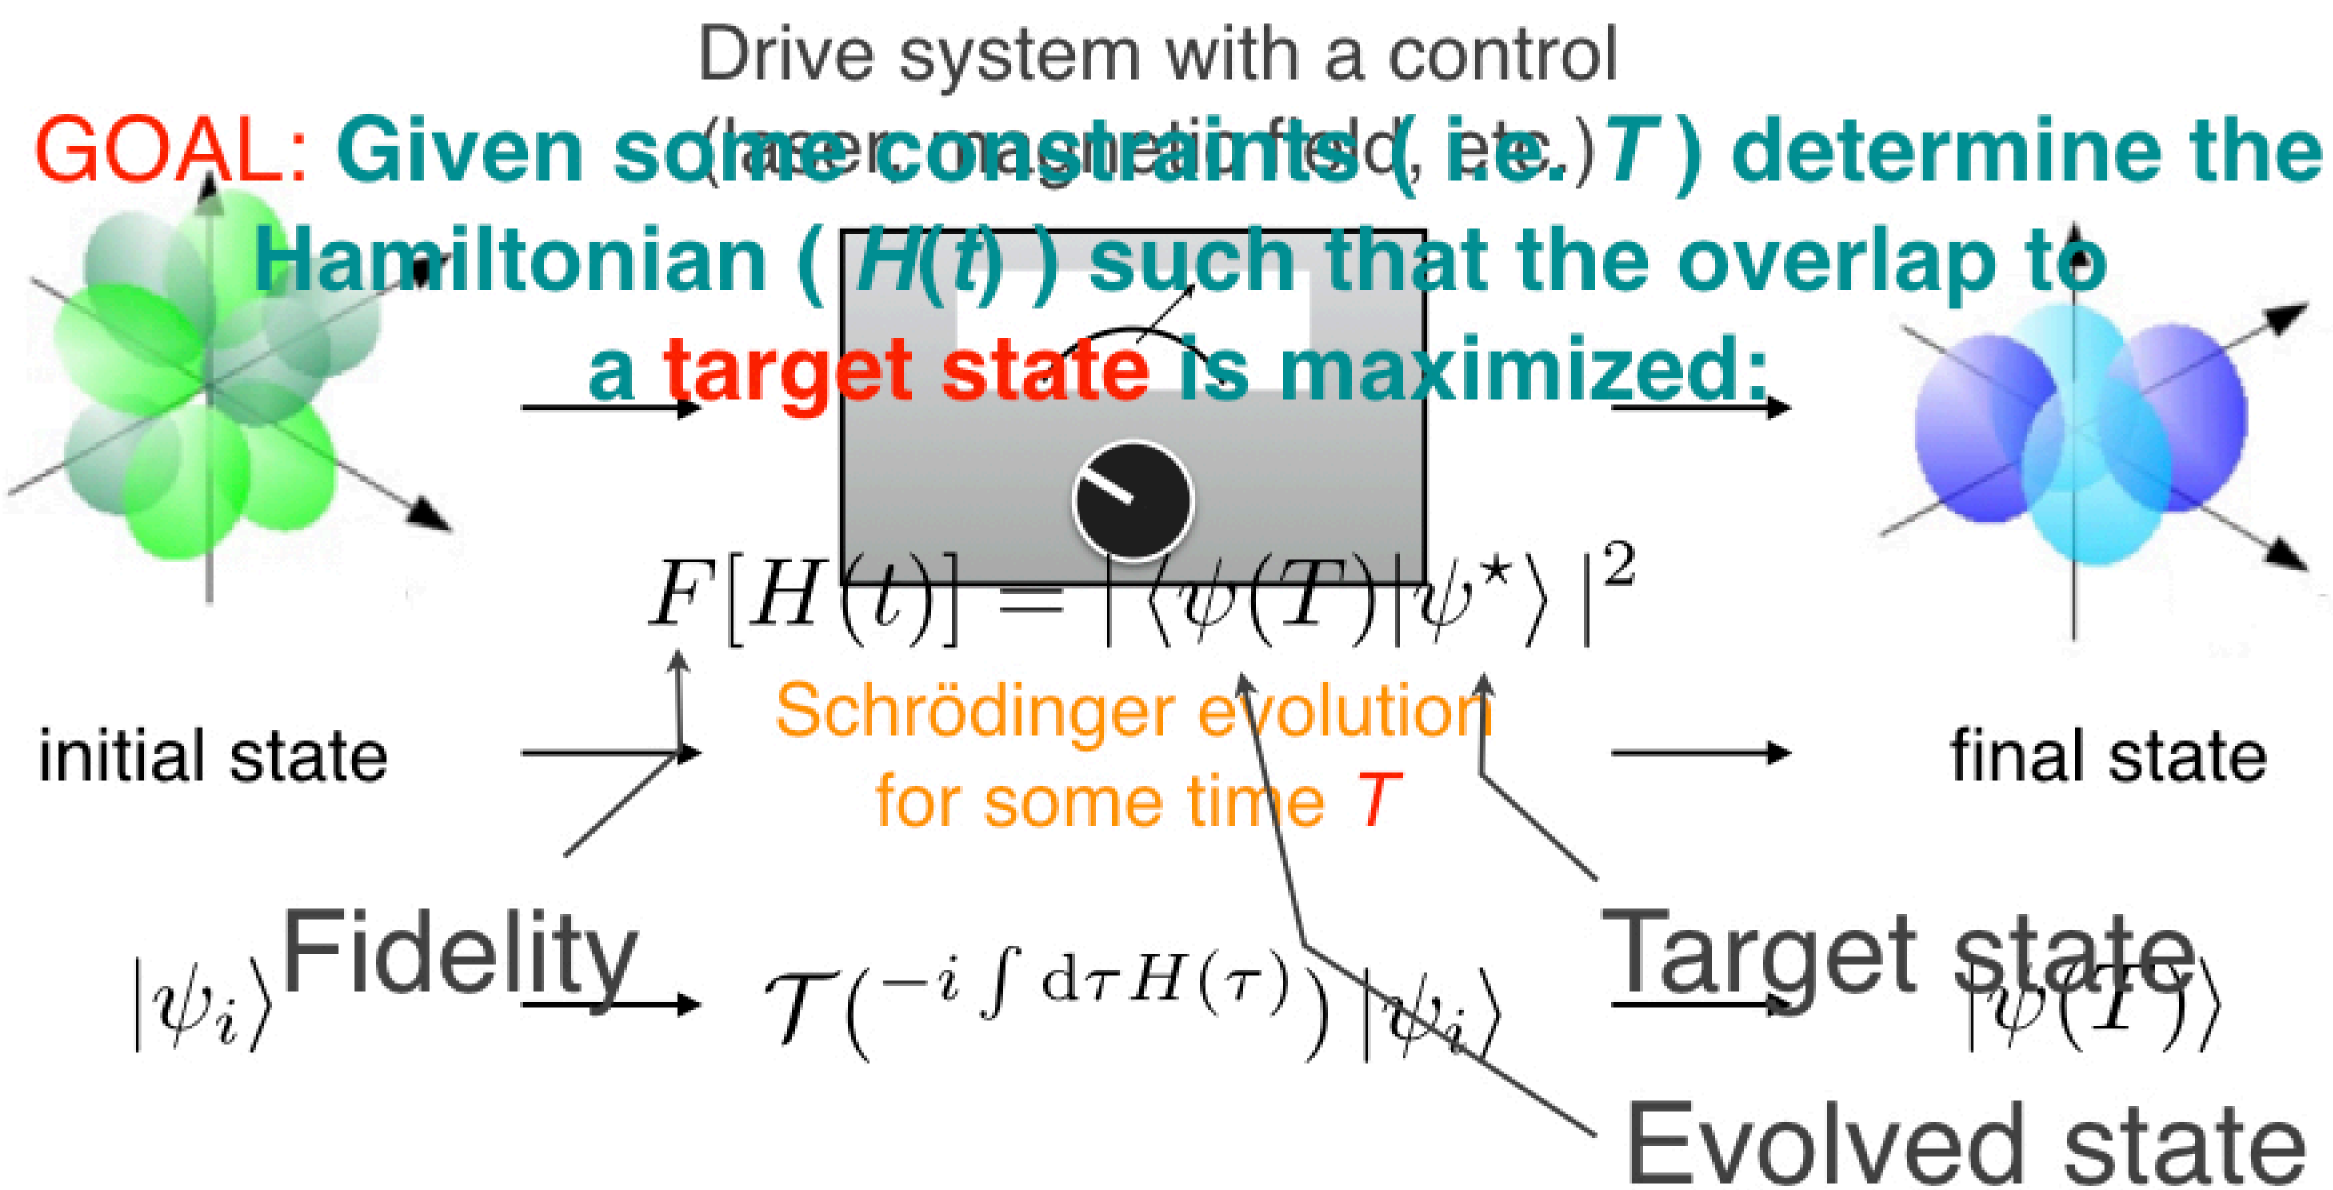

GOAL
prepare states in non-integrable Ising model

1. Two-Level System $H(t) = S^z + h(t)S^x$

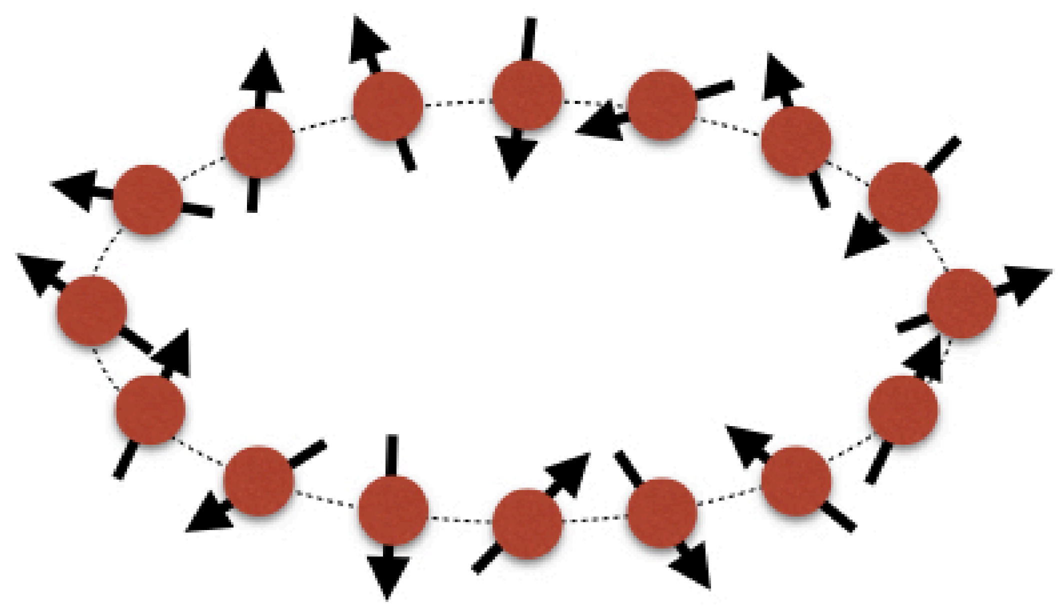

2. Many-Body System $H(t) = \sum_j J S^z_{j+1} S^z_j + h_z S^z_j + h_x(t) S^x_j$

Reinforcement Learning in a Nutshell

- Machine Learning
  - ...
  - ***Reinforcement Learning (RL)***
  - ...

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

agent

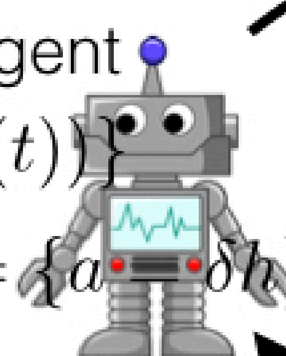

ENVIRONMENT

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

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

RL Applied to the State Preparation Problem

BU

agent

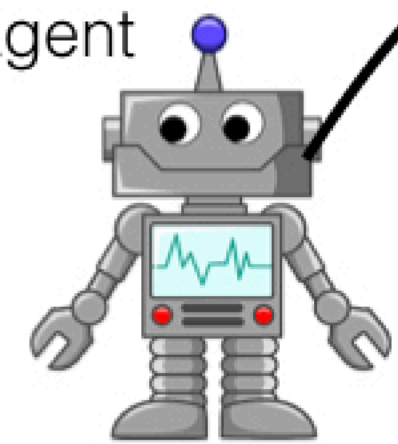

1

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 $a_0 : \delta h = 0.8$

problems:

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

→ state space *exponentially* big

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

biased MC sampling

→ how do we chose actions?

3 calculate reward r

exploration exploitation dilemma

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

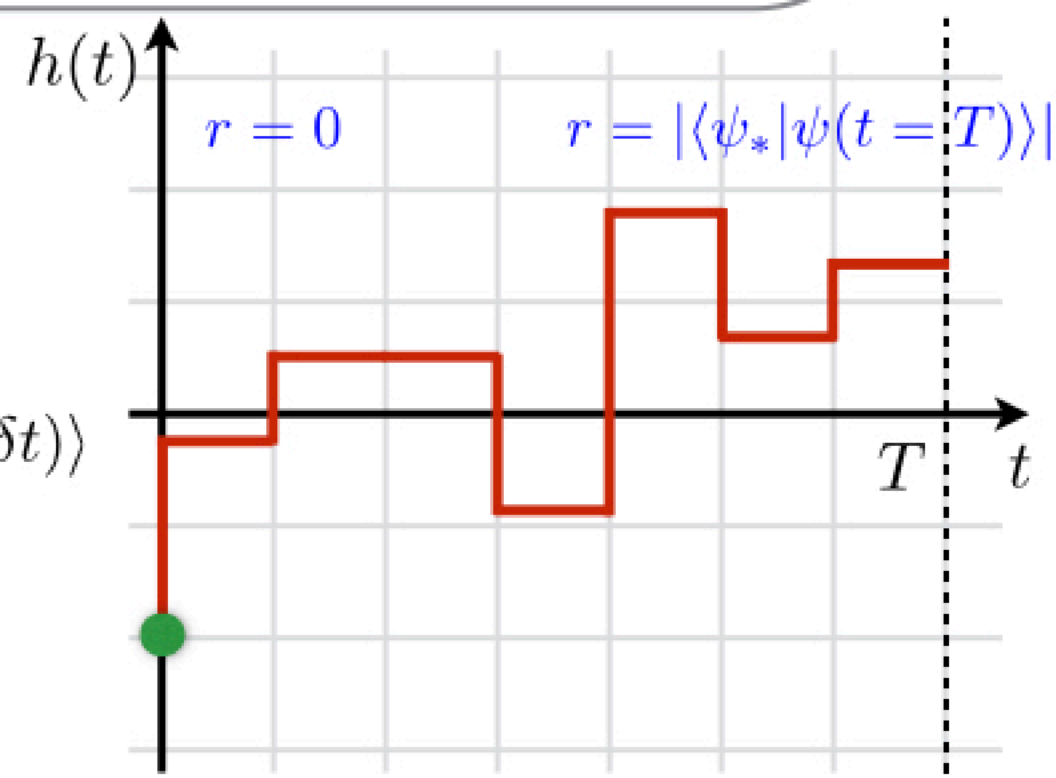

episode completed

Fidelity optimization

- We discretize the protocols (bang-bang protocols):

$$h(t) \in \{-h^\star, h^\star\}$$

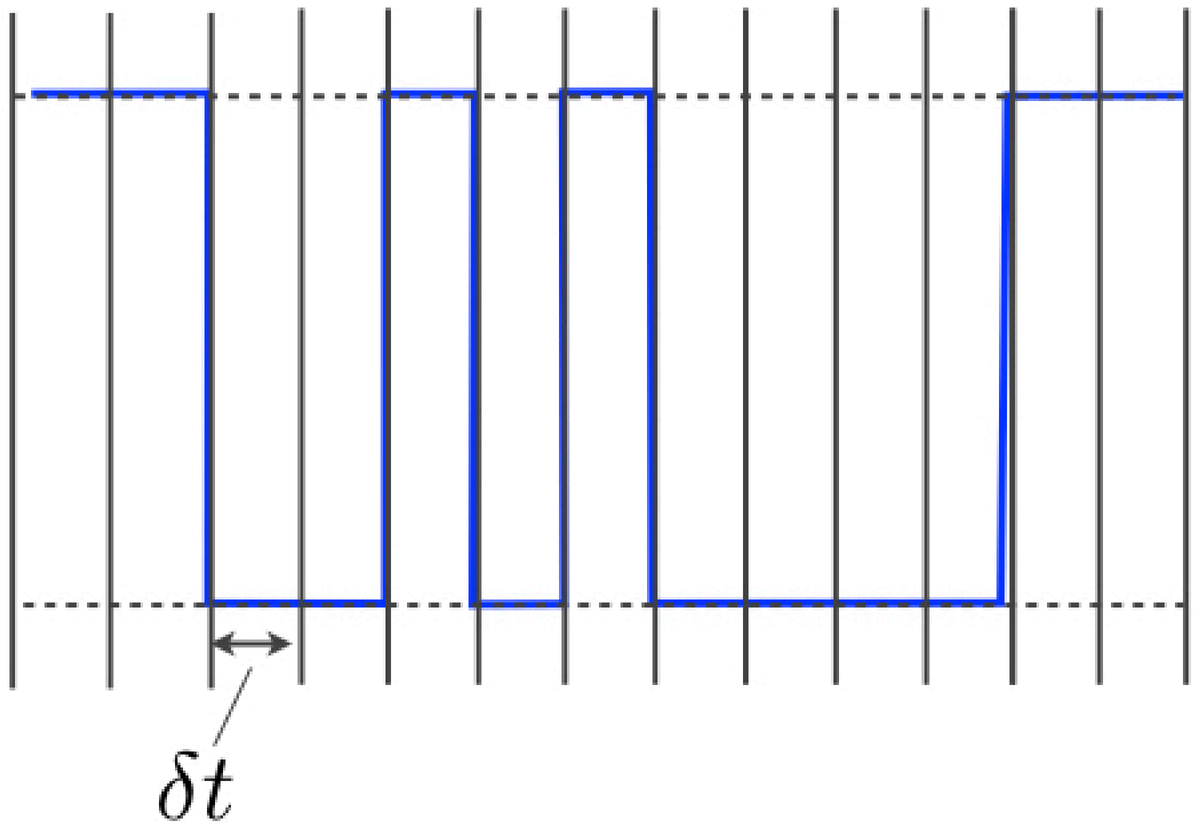

h = (1, 1, 0, 0, 1, 0, 1, 0, 0 ,0 ,0, 1 ,1)

We will take:

$$h^\star = 4$$

Initially, the agent has no knowledge about the problem and tries out random strategies...

Every episode (see bottom left corner), the computer learns new bits of information about the problem.

How hard is this optimization problem?

Many-body problem

- Non-integrable chain of qubits:

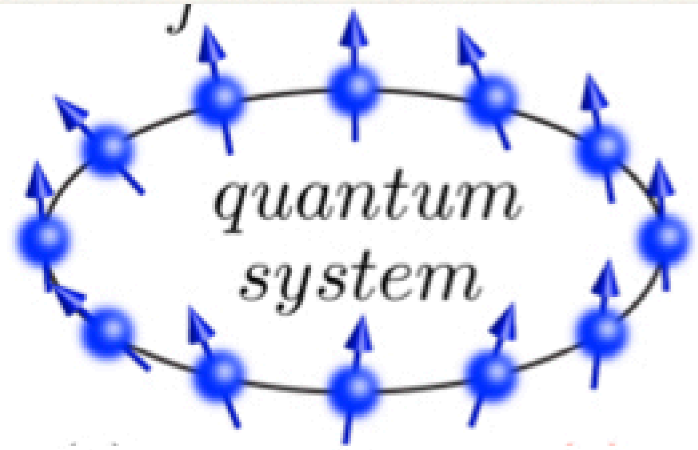

$$H(t) = \sum_{i}^{L} (-J S_i^z S_{i+1}^z + h(t) S_i^x + h_z S_i^z)$$

$|\psi_i\rangle$ Ground-state of *H*, with *h* =-2

$|\psi^\star\rangle$ Ground-state of *H*, with *h* =+2

we take L>=6

Stochastic descent and local minima

- Start from an initial random protocol
    - c.g. : $h = (1, 1, 0, 0, 1, 0, 1, 0, 0 ,0 ,0, 1 ,1)$
- Compute fidelity between final and target state $F_0 = 0.521$
- Propose updating a single bit $h = (1, 1, 1, 0, 1, 0, 1, 0, 0 ,0 ,0, 1 ,1)$
- If $F_1 > F_0$, accept update
- Continue until a local maximum is reached

Stochastic descent

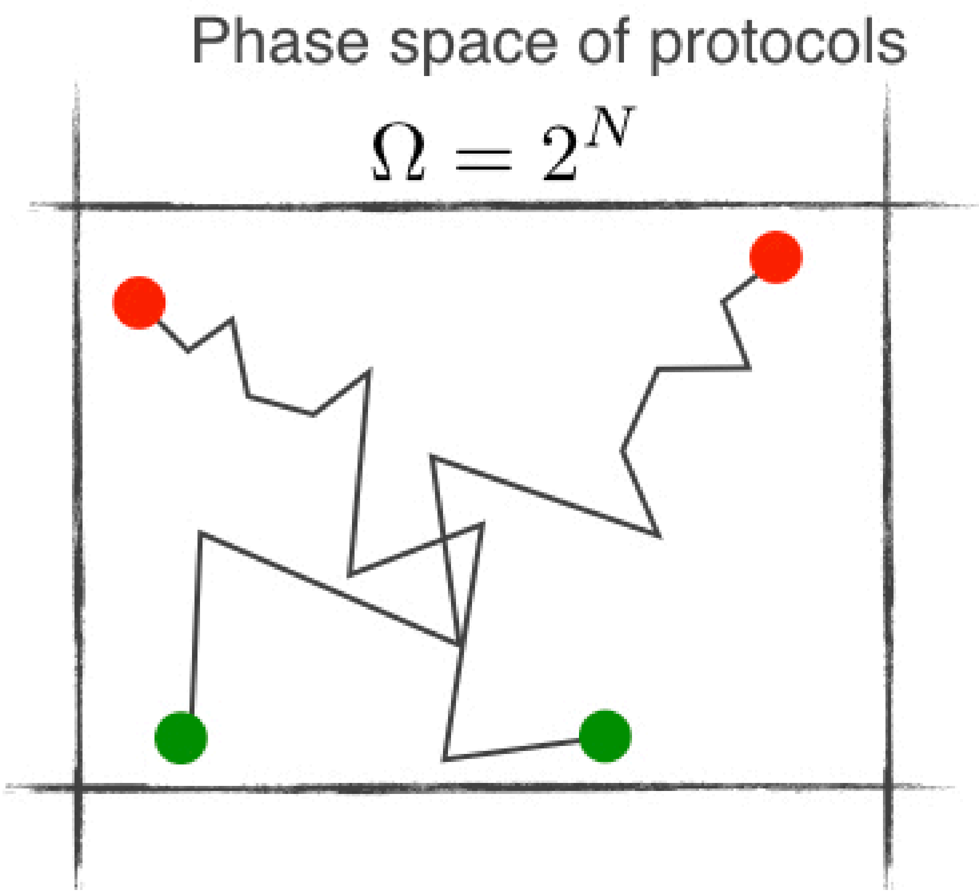

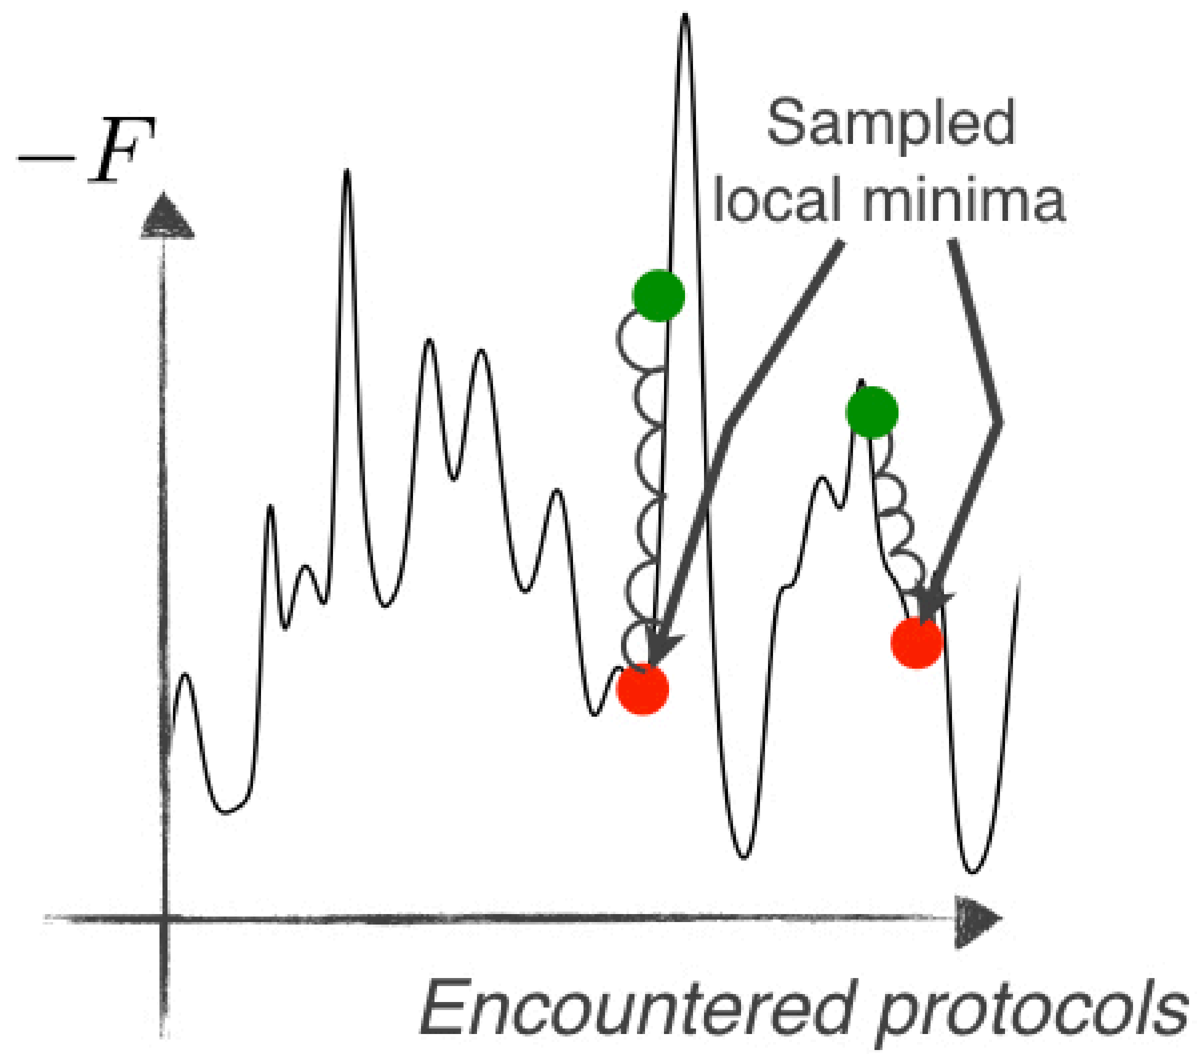

- For a non-convex landscape, multiple minima are found

Measuring order parameters

- We can define a variance order parameter:

$$q(T) = \frac{1}{N}\sum_{t}\sum_{i}^{\#sample}\left(h_i(t) - \langle h(t)\rangle\right)^2$$

$$q \approx 0$$

sample #1 ↑↑↑↑↓↓↓↓

sample #2 ↑↑↑↑↓↓↓↓

sample #3 ↑↑↑↑↑↓↓↓

sample #4 ↑↑↑↓↓↓↓↓

sample #5 ↑↑↑↑↓↓↓↓

$$q \approx 1$$

sample #1 ↑↓↑↓↓↑↓↑

sample #2 ↓↑↓↑↑↓↑↓

sample #3 ↑↓↓↓↑↓↓↓

sample #4 ↓↑↓↓↓↑↑↓

sample #5 ↑↑↑↓↑↑↓↓

Measuring order parameters

- Order parameter: $$f = \frac{\text{Number of unique minima}}{M}$$

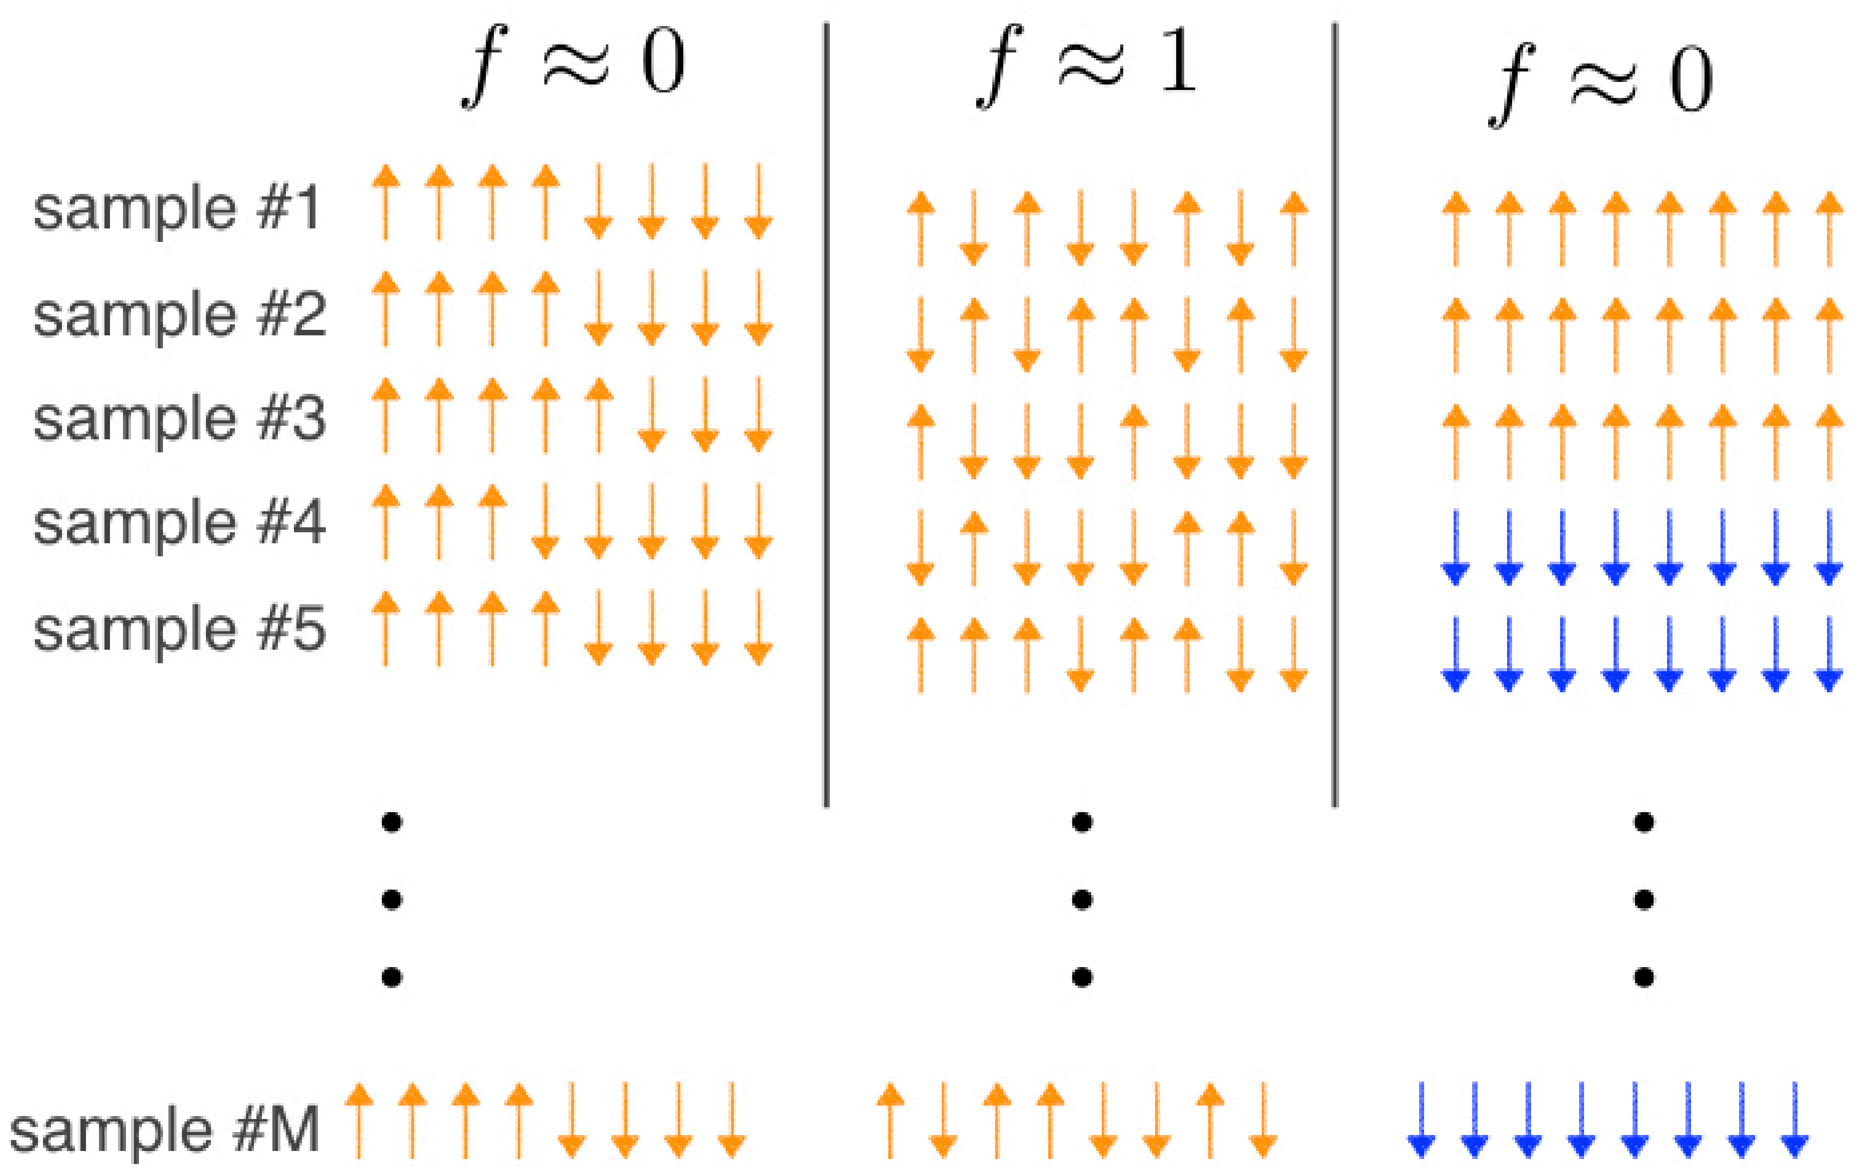

The single qubit phase diagram

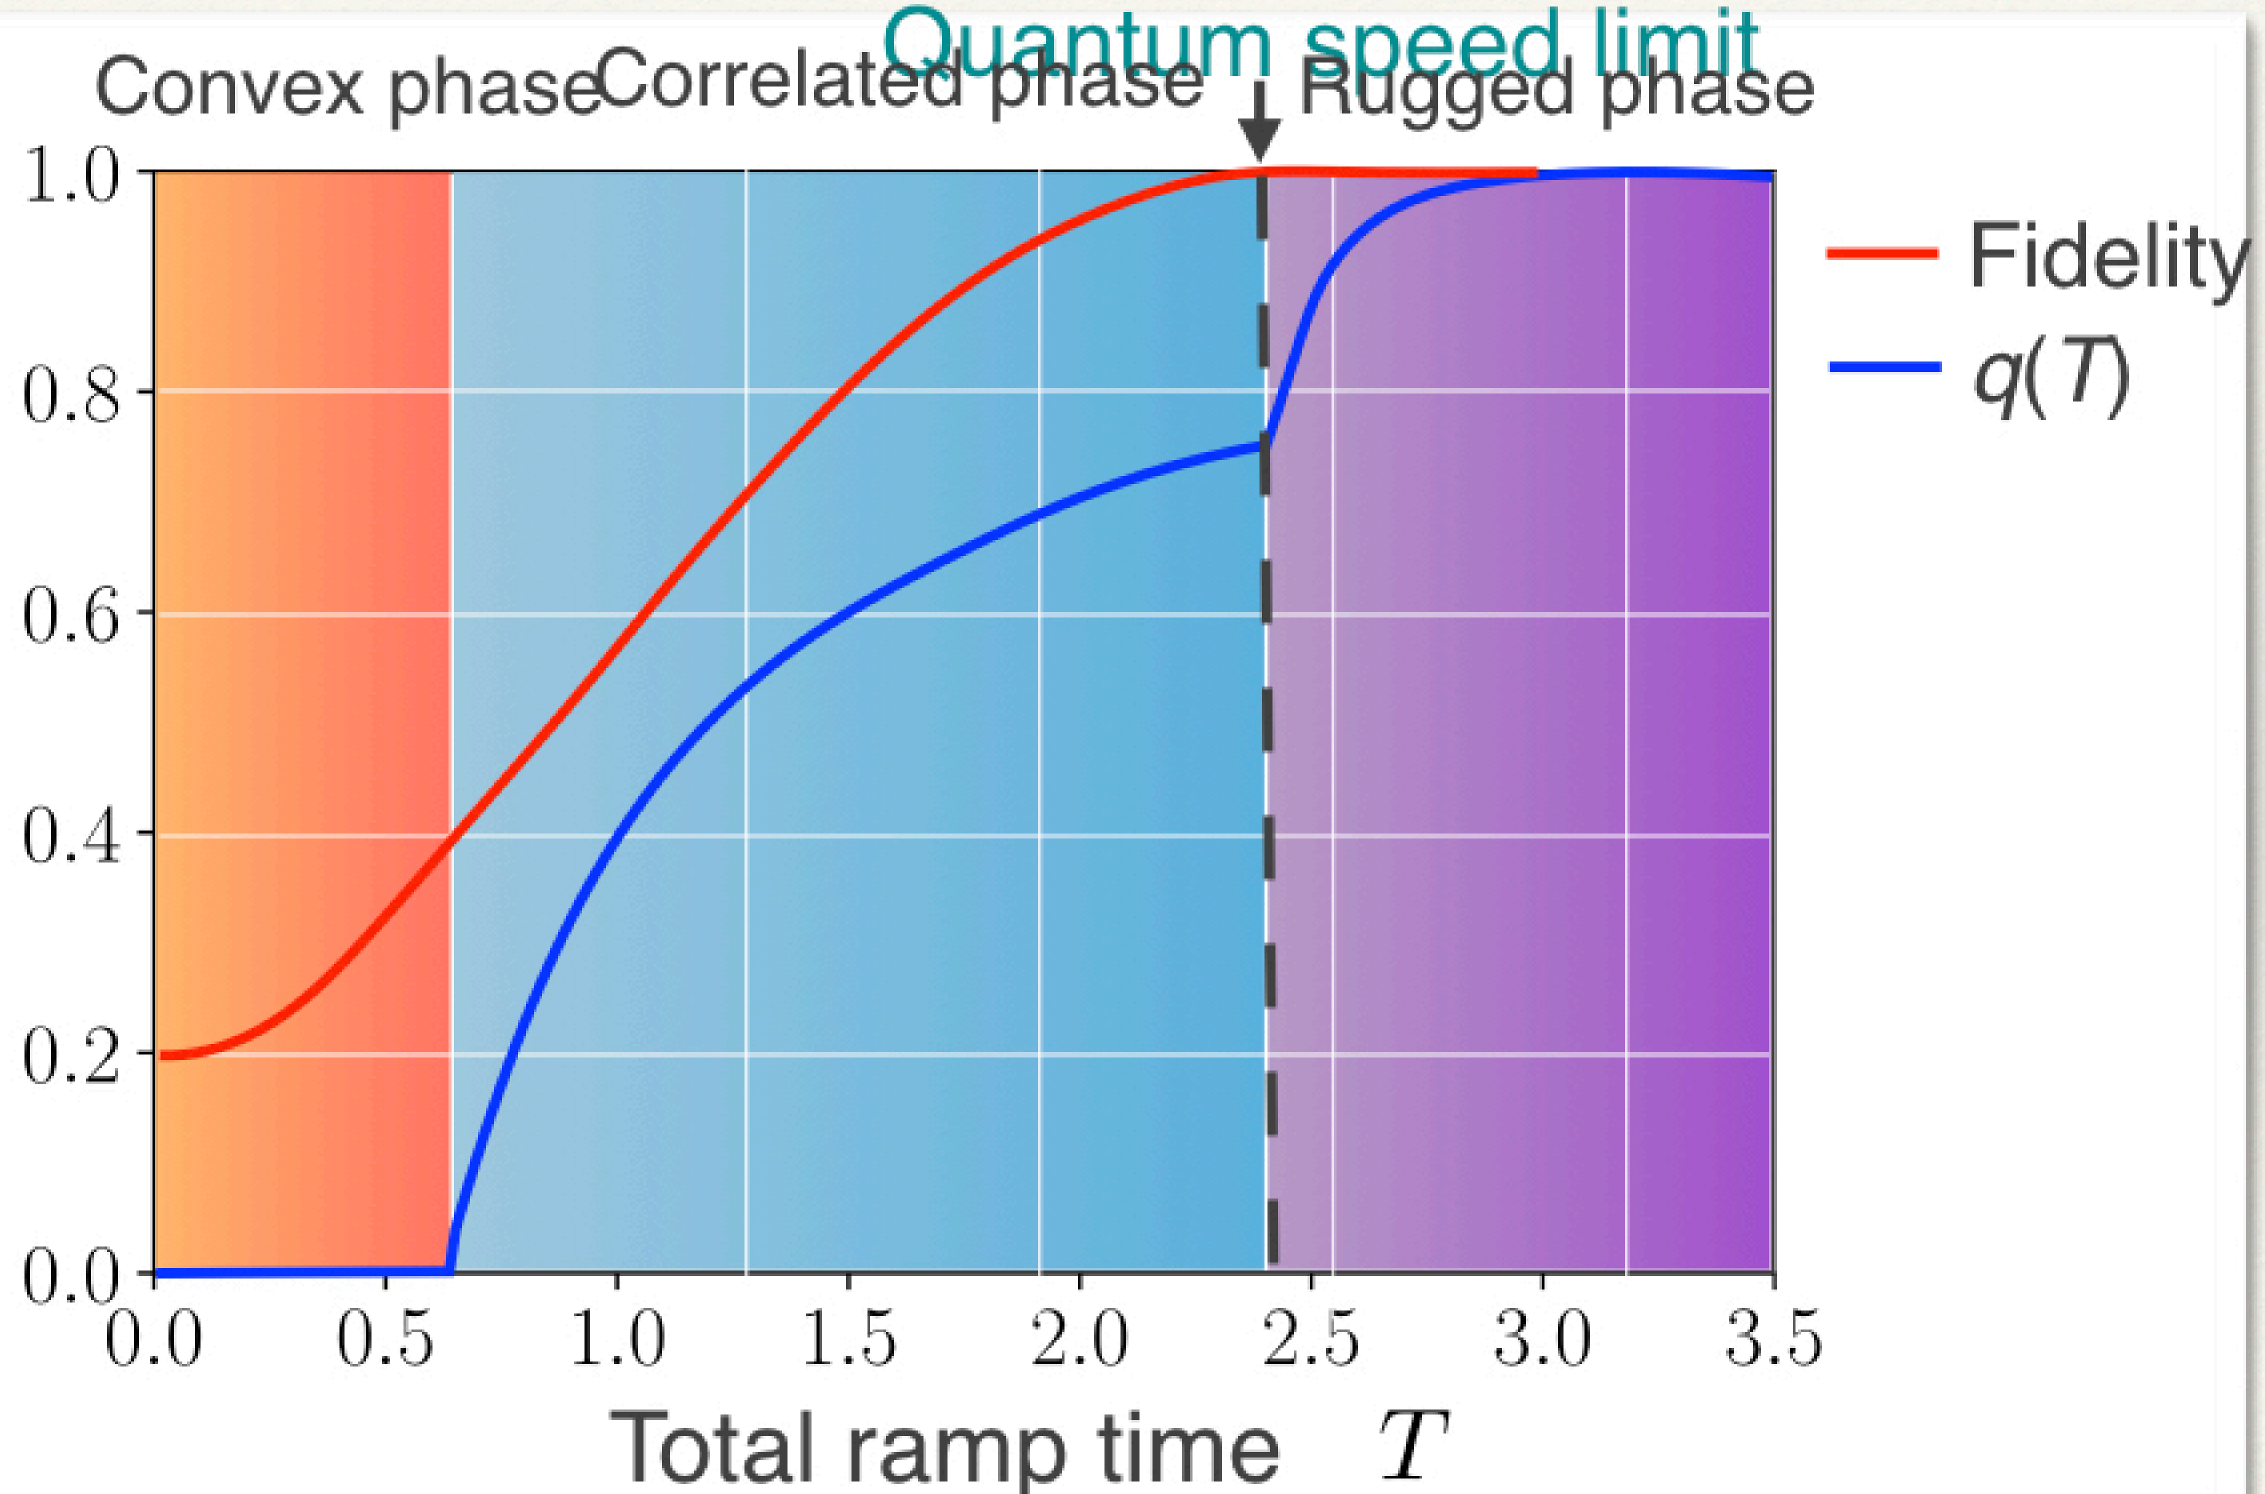

Many-body system phase diagram

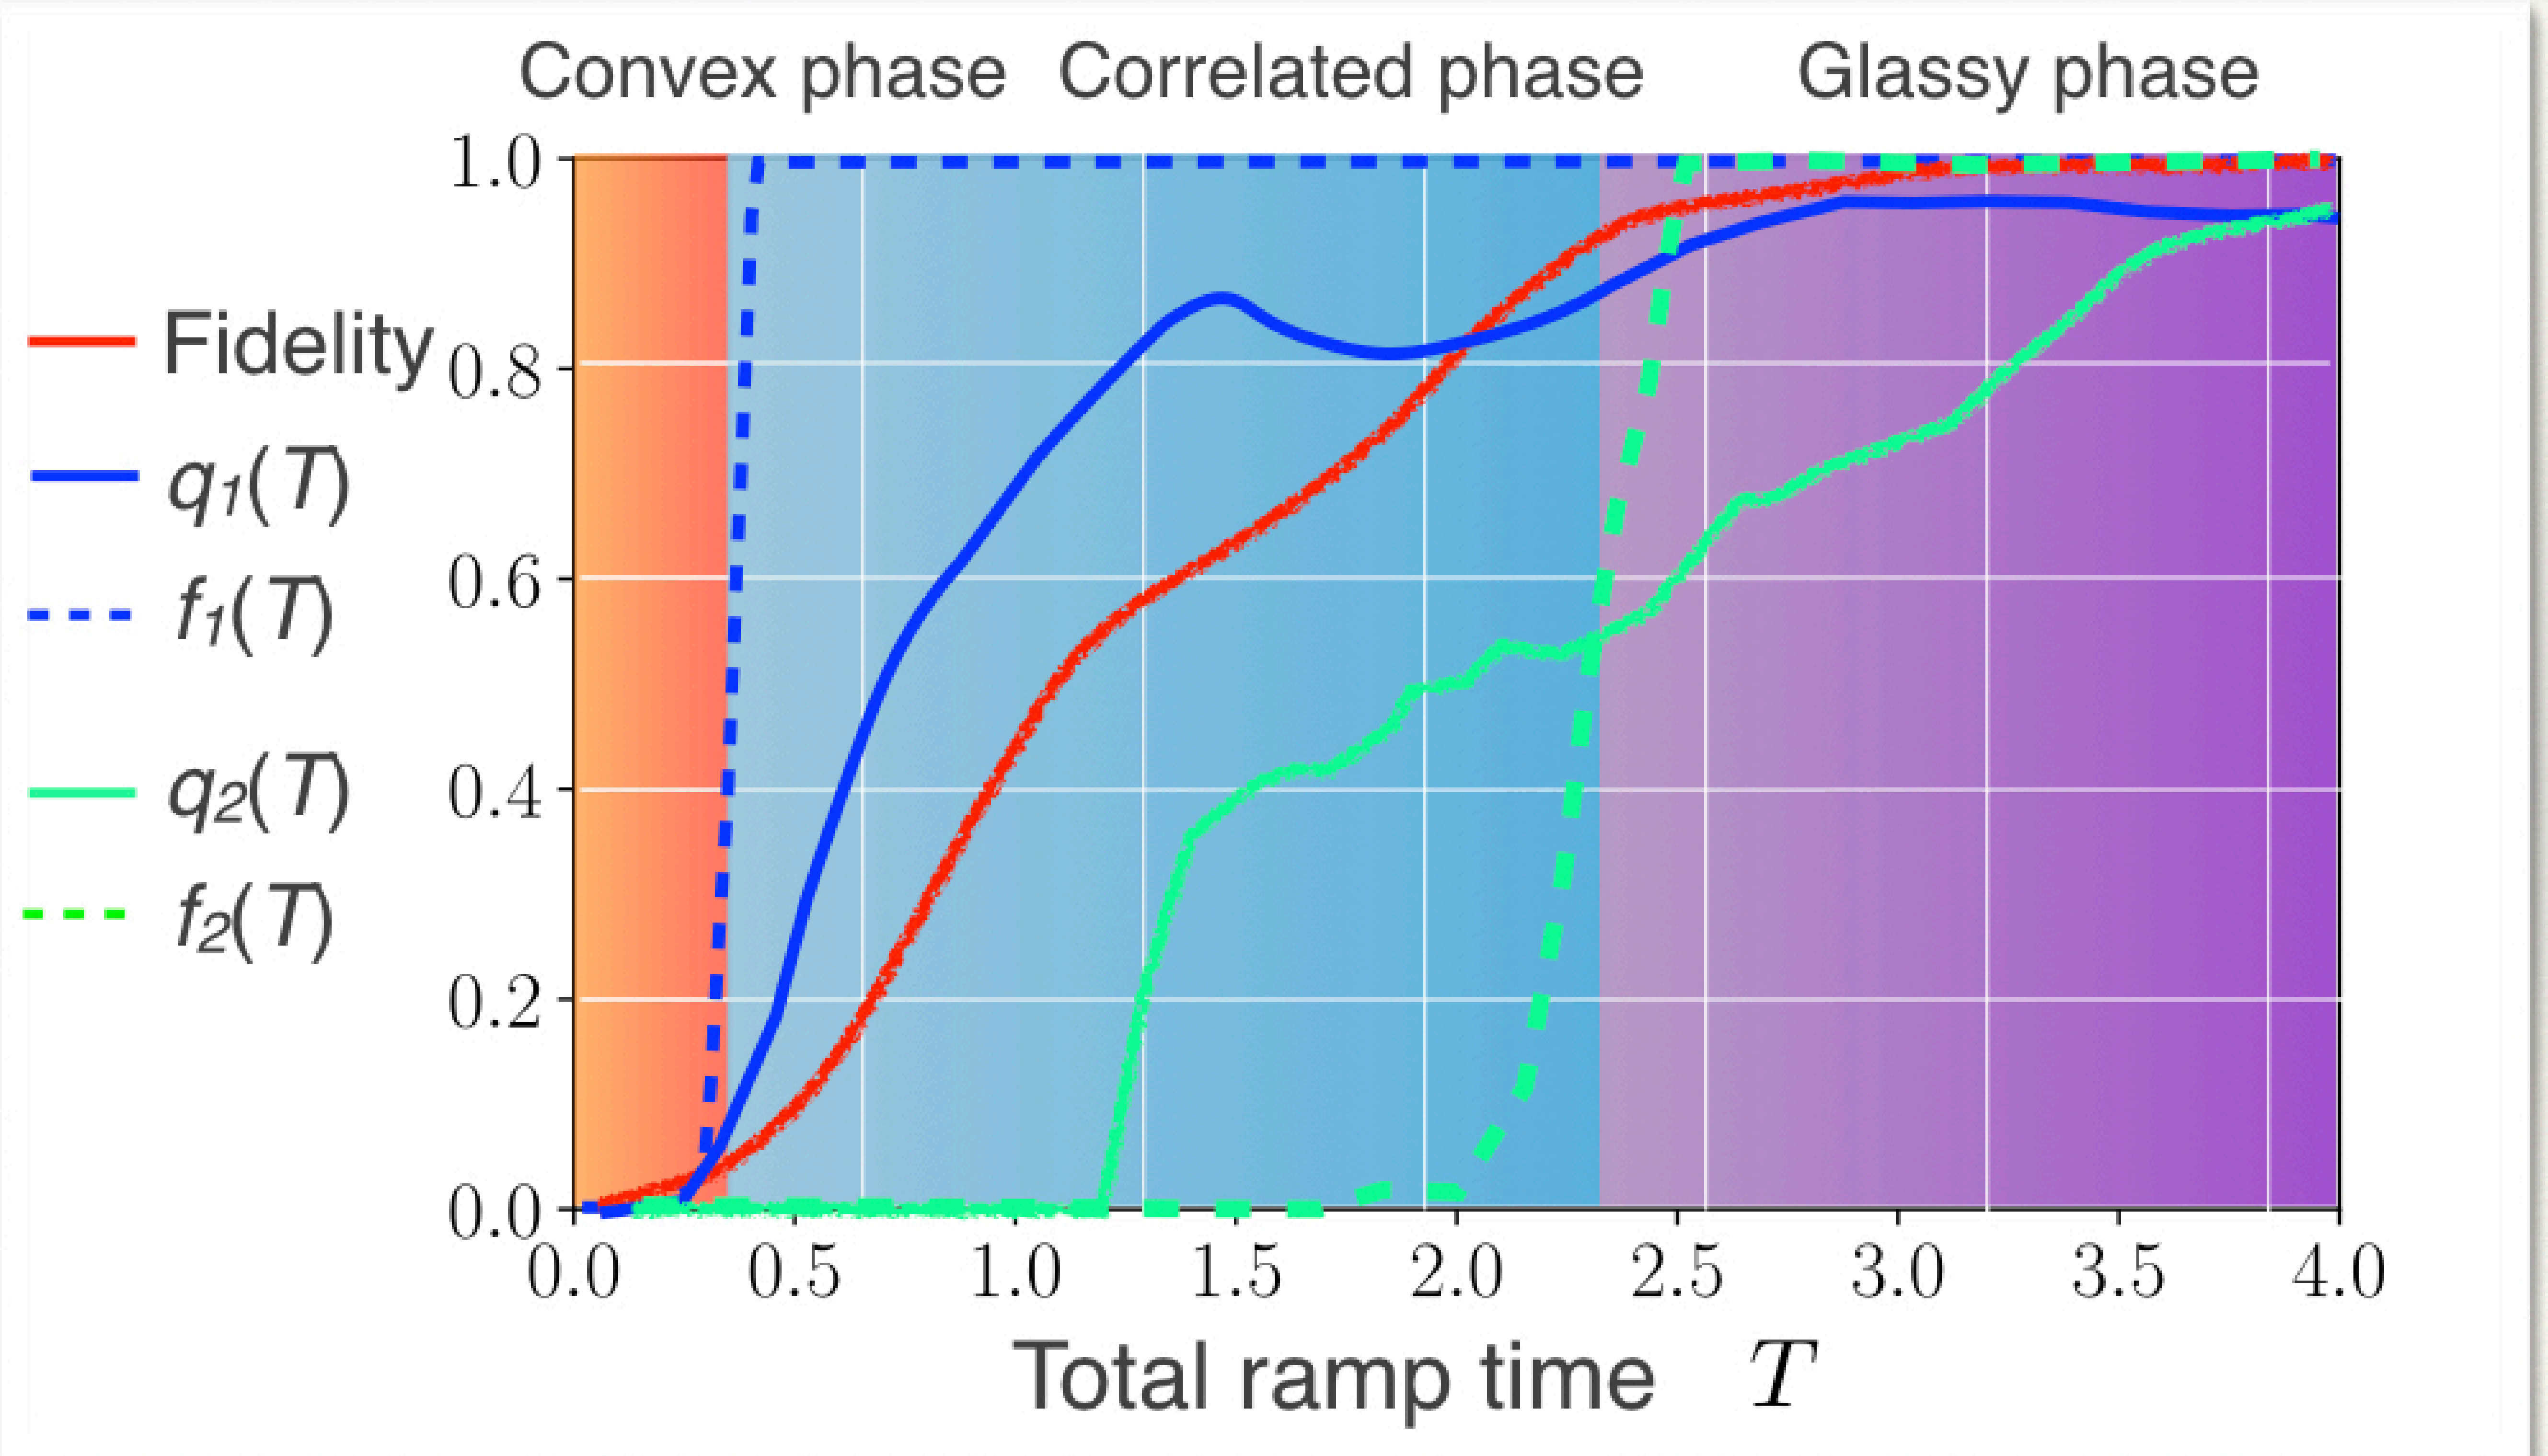

Manifold learning and clustering

- t-SNE: t-distributed stochastic neighbor embedding

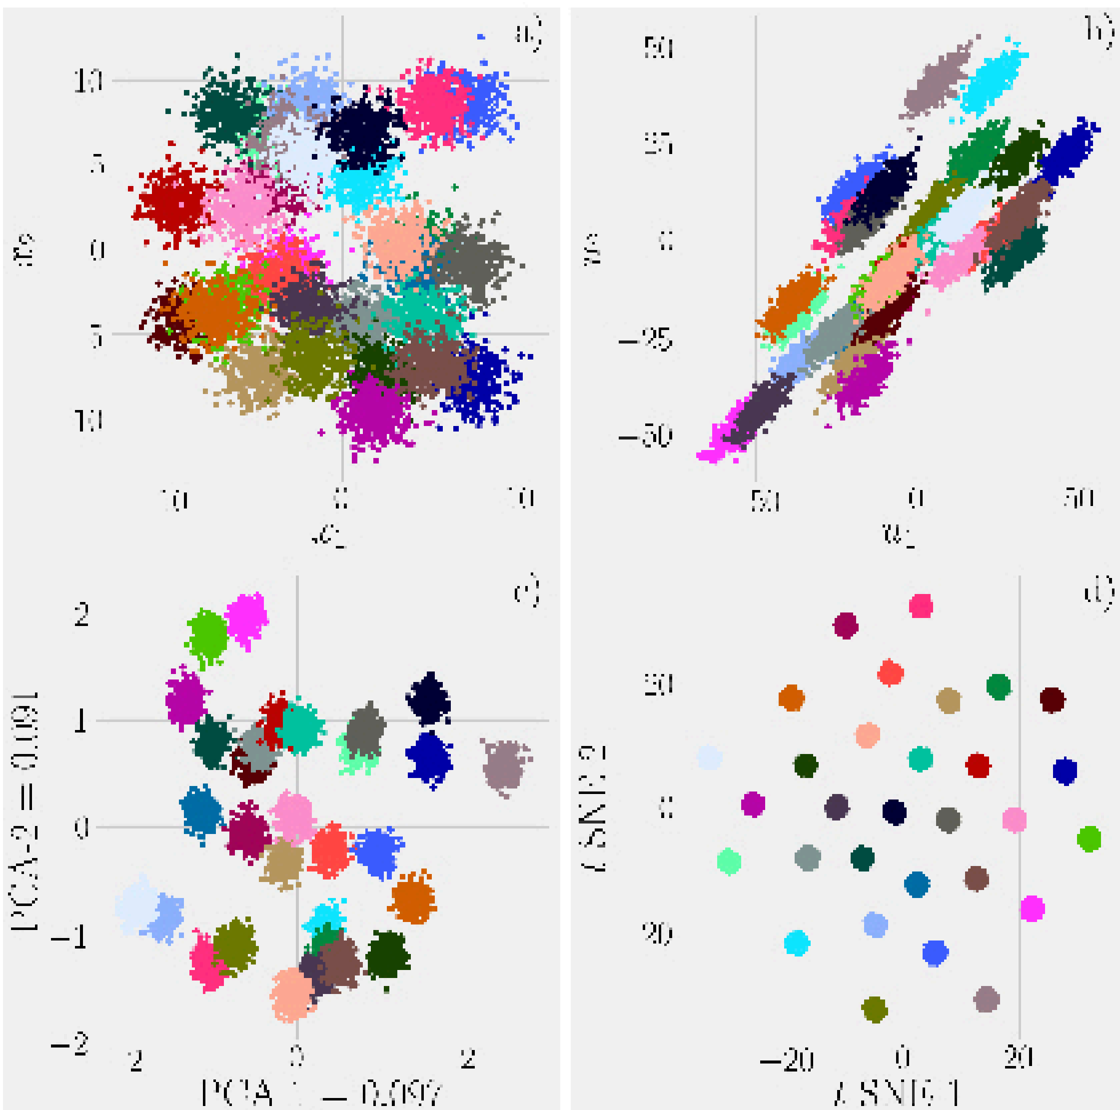

Gaussian Mixture in 40 dim

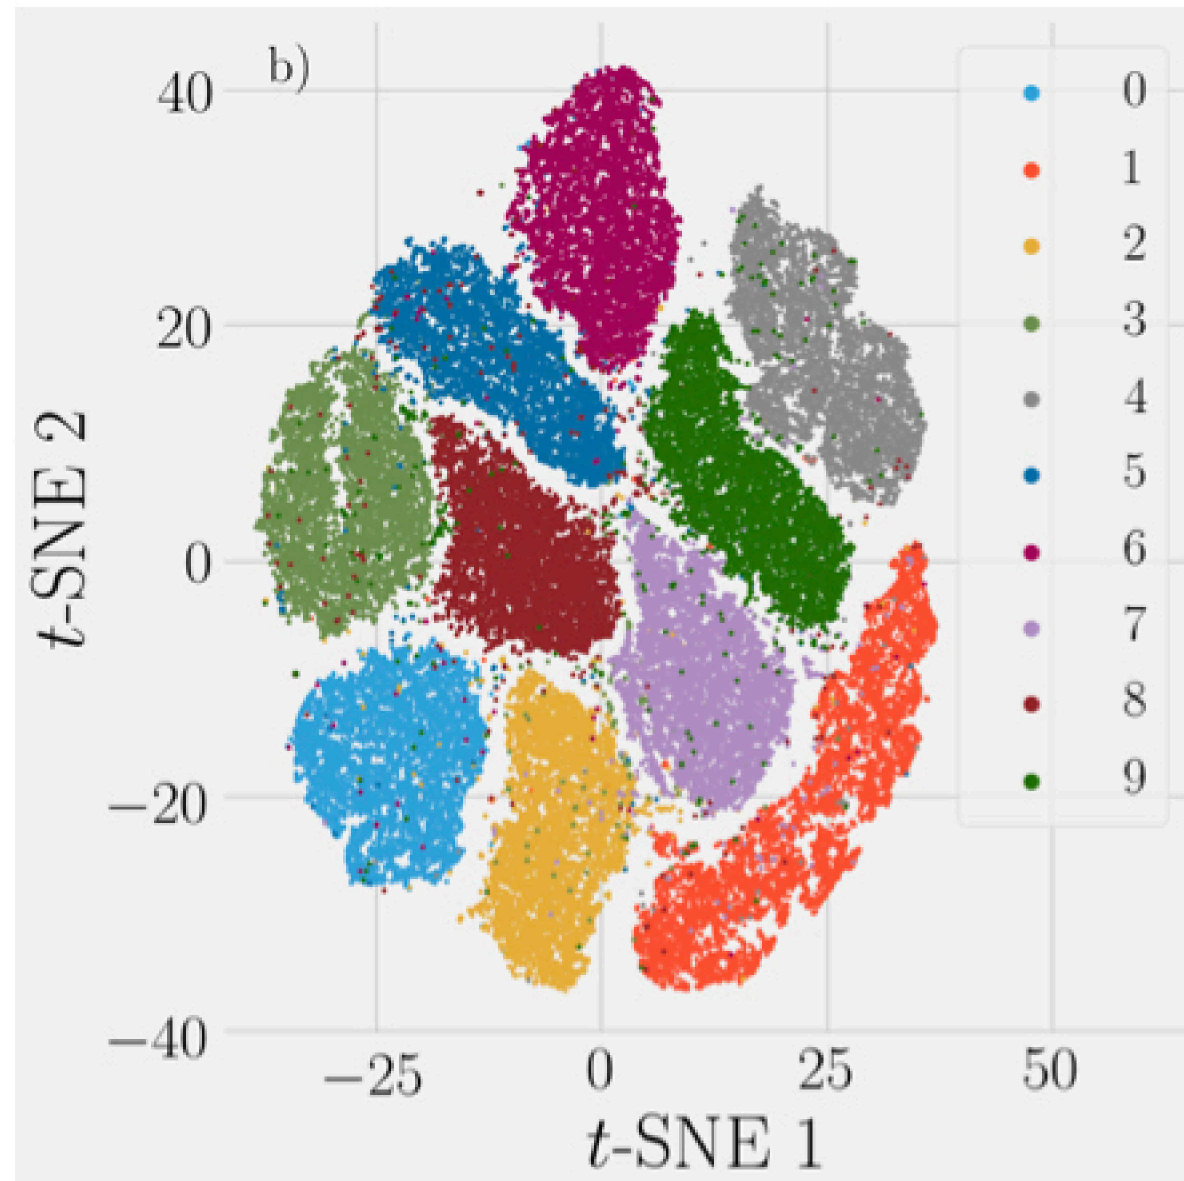

MNIST

Clustering in Glassy Phase

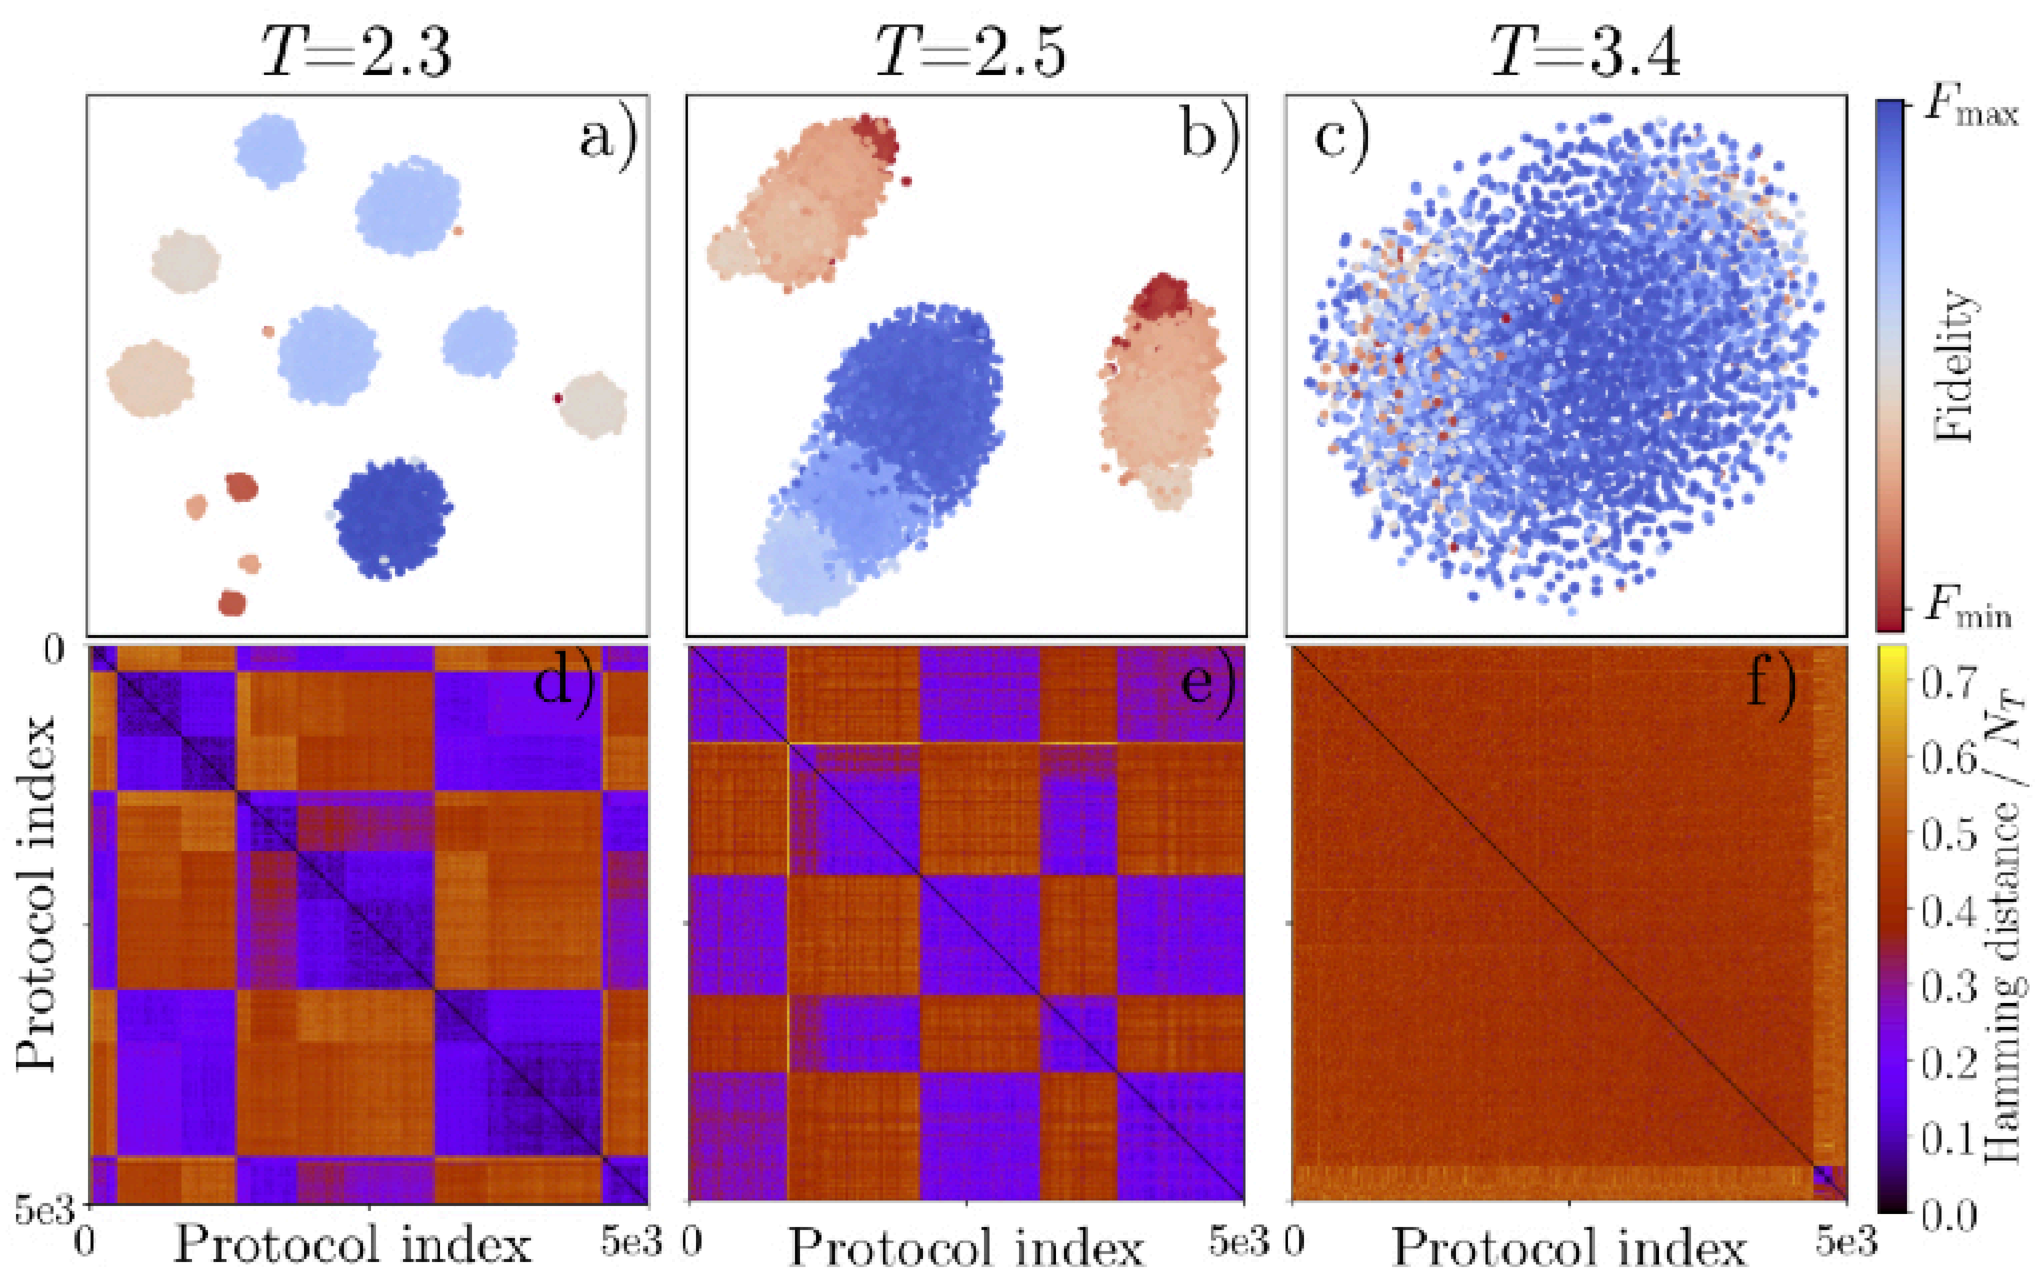

Conclusion

- RL is an interesting paradigm for studying nonequilibrium physics
- Quantum control can be understood in the language of phase diagrams and phase transitions
- For a large class of control problems, it is equivalent to a spin-glass:
  - Extremely rugged energy-like landscape
  - Exponential complexity in finding optimal control protocol
  - Effective Ising glass - non-local frustrated couplings
  - Exhibit a clustering transition we think is analogous to constraint satisfaction problems

Acknowledgments

Marin Bukov

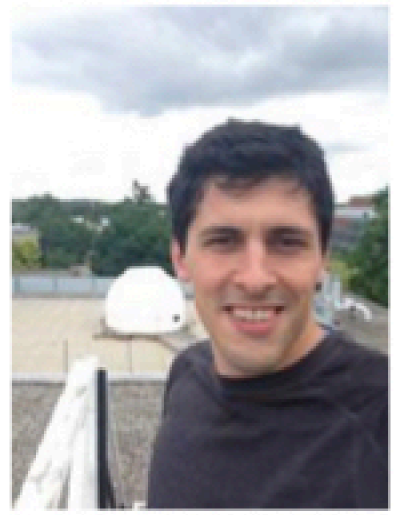

Alex D

Dries Sels

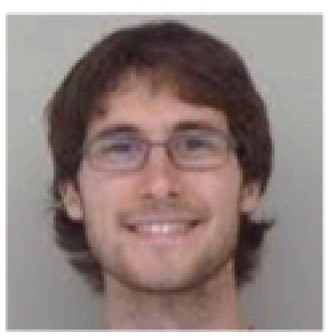

Phil Weinberg

A. Polkovnikov

SIMONS FOUNDATION

NATIONAL INSTITUTES OF HEALTH

RESEARCH CORPORATION for SCIENCE ADVANCEMENT

GORDON AND BETTY MOORE FOUNDATION

ALFRED P. SLOAN FOUNDATION 1934