《物理与人工智能》

AI与量子多体物理

授课教师：张亿
frankzhangyi@pku.edu.cn

北京大学物理学院量材中心

2026/11/23 （第11周）

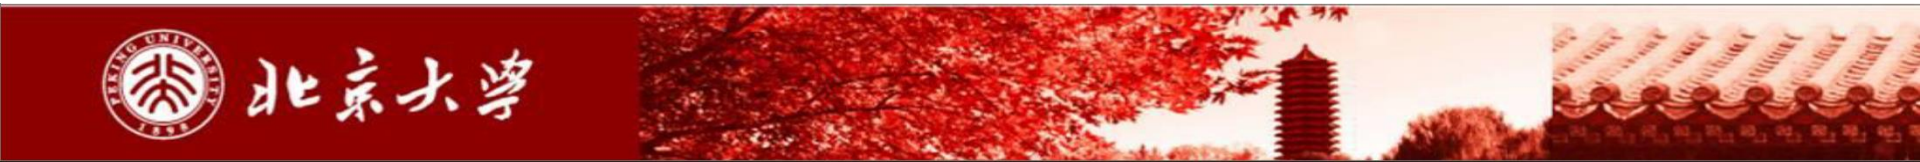

AI on the cutting-edge academic fields

- The Riemann hypothesis:

"all non-trivial zeros of the Riemann zeta function have a real part equal to ½."

—— proven? by *Levent Alpöge & Claude*, Nov. 2026

- The Exponential Time Hypothesis:

"3-SAT cannot be generally solved faster than an exponential time frame."

—— disproven? by *OpenAI & Claude*, Oct. 2026

- So, the fate and prospect

  of our science majors?

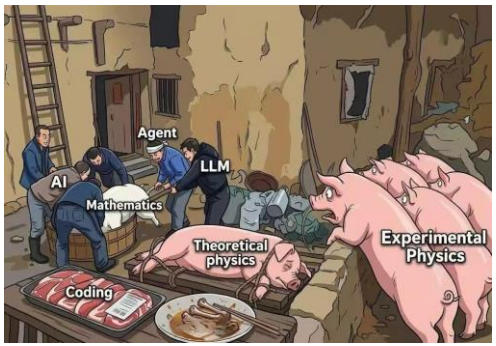

AI vs quantum many-body physics

- Quantum physics and quantum computation is beyond classical simulations.

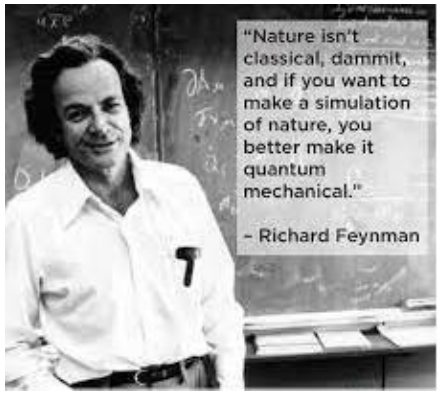



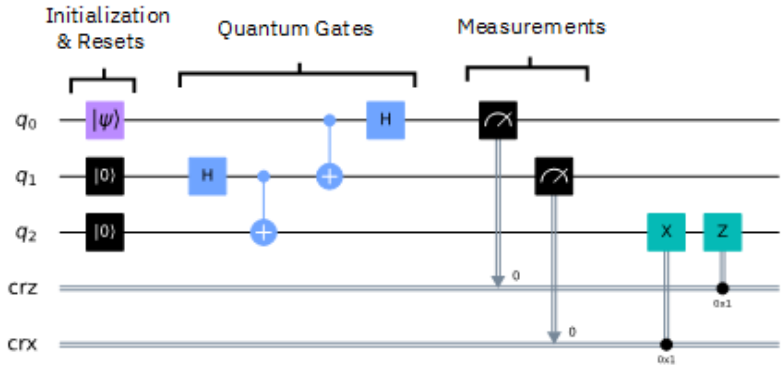

Richard Feynman

A quantum many-body state: $|\Psi\rangle = \sum_{\alpha} \phi_{\alpha}$ 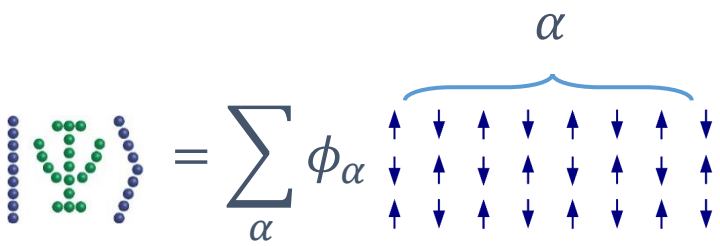

The Fock space is exponentially large $2^{30}\sim 1000000000$, $2^{1000} \sim 100\cdots00$ with >300 zeros!

AI and quantum many-body physics

北京大學 PEKING UNIVERSITY

	Universal macroscopic behaviors	Numerous microscopic degrees of freedom
✓ Image recognition:	flower	a single pixel
✓ Many-body physics:	antiferromagnet	a single spin
✓ And more:	wetness	a water molecule
	ant colony	an ant
	macroeconomy	an individual
	life	a bio molecule
	AI	a node



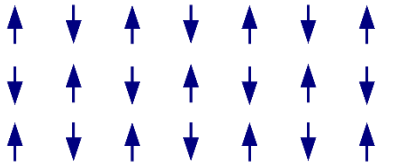





John Hopfield & Geoffrey Hinton
2024 Nobel Prize in Physics

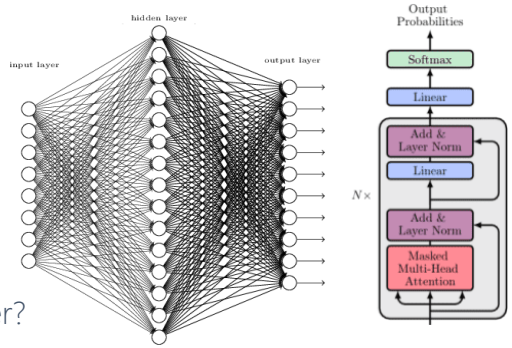

still the era of gun power?

What is emergence?

- Emergence occurs when a complex entity has properties or behaviors that its parts do not have on their own, and emerge only when they interact in a wider whole.

4 August 1972, Volume 177, Number 4047 **SCIENCE**

More Is Different

P. W. Anderson

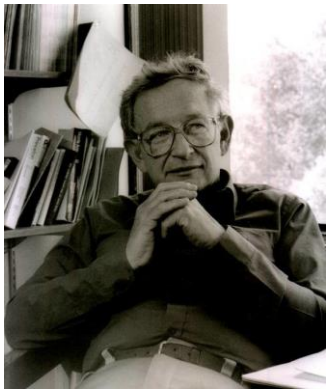

Philip W. Anderson

- Renormalization group: meaningful collective properties



Kenneth G. Wilson

$$H = -J \sum_{\langle ij \rangle} \boldsymbol{\sigma_i} \cdot \boldsymbol{\sigma_j}$$

$J_\infty \to 0$ paramagnetic phase

J_c

$J_\infty \to \infty$ ferromagnetic phase

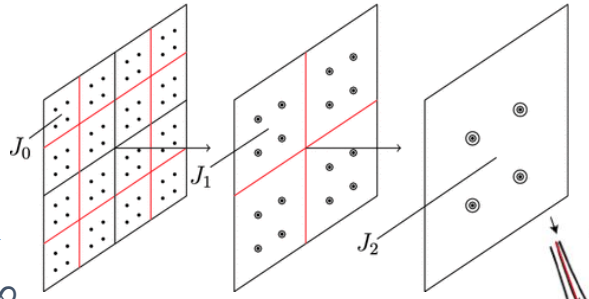

schematic "RG flow" of images:

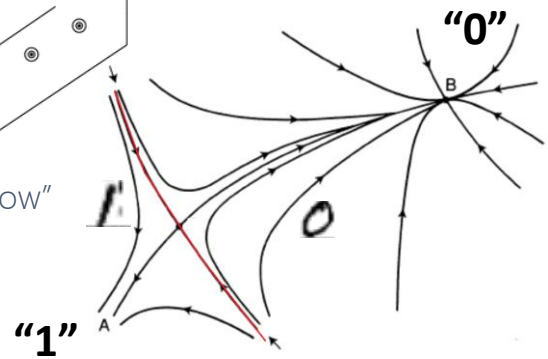

universal properties irrespective of microscopic details

AI for Quantum Phase Recognition

Review: discriminative AI



AI for Quantum Phases Recognition

- Viable – review: discriminative AI and supervised learning
  - Well-defined problems and exact solvable cases
  - With interpretability



Order parameters?
Topological invariants?

- Necessary and helpful:
  - Diverse phases and candidates
  - Hidden, abstract, and complex rules
  - Noises and fluctuations
  - Big data, experimentally or numerically

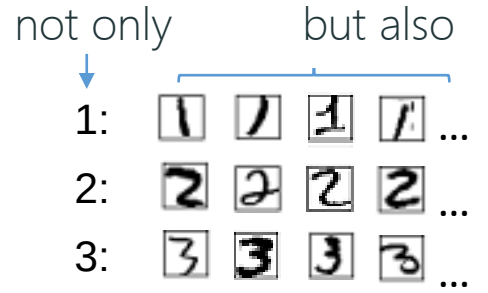

$$|\Psi\rangle = \sum_{\alpha} \phi_{\alpha}$$

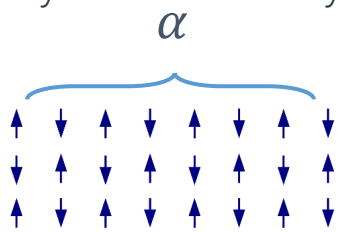

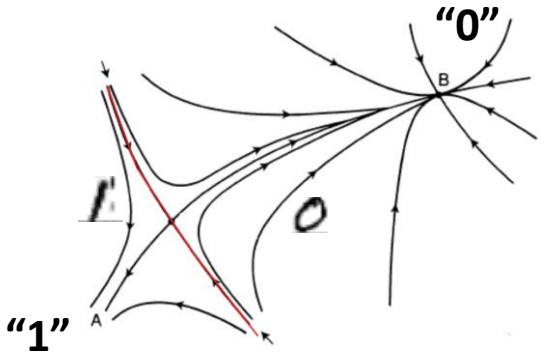

AI for Quantum Phases Recognition

北京大學 PEKING UNIVERSITY

- The Ising model: (Statistical Mechanics 101)

$$H(\sigma) = -J \sum_{\langle i\ j \rangle} \sigma_i \sigma_j$$

two phases: ferromagnetic (ordered) vs paramagnetic (disordered)

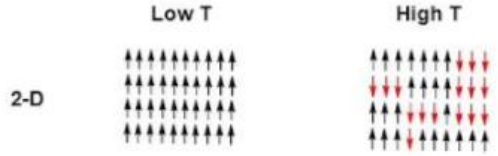

Numerical data: Monte Carlo samples

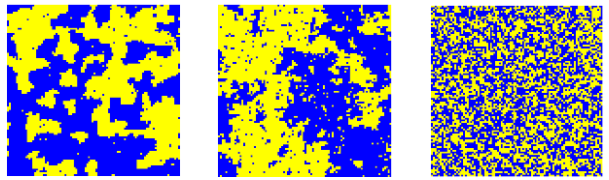

ferromagnetic critical paramagnetic

- The Aubry-Andre model (quasicrystal):

$$H = -J \sum_i \left(c_i^\dagger c_{i+1} + \mathrm{h.c} \right) + 2\lambda \sum_i \cos(2\pi\phi i) c_i^\dagger c_i$$

Numerical data: LDOS

delocalized vs localized

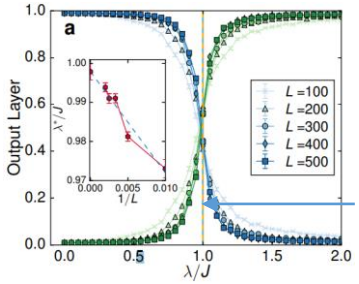

$\lambda_c = J$

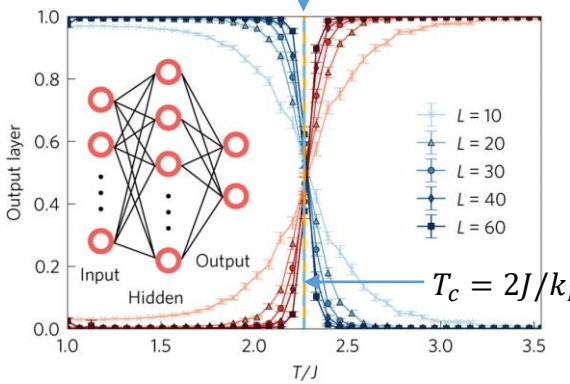

$T_c = 2J/k_B \ln(1+\sqrt{2})$

simple HW assignment: phase diagram (5pt)

Juan Carrasquilla, Roger G. Melko.

AI for Quantum Phases Recognition

北京大學 PEKING UNIVERSITY

- Quantum and topological phases: the compatibility issue

a feature selection layer to bridge between

Quantum many-body states ⟷ 'Informative' operators ⟷ AI

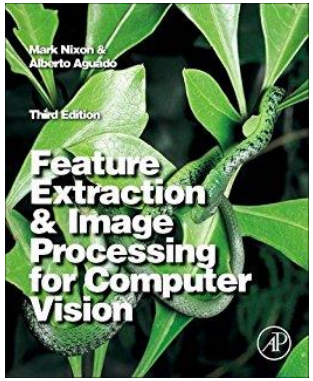

- E.g., topological phases:

relevant operators in the Kubo formula for σ_{xy}

$P_{jk}P_{kl}P_{lj}$ $\quad P_{ij} \equiv \langle c_i^\dagger c_j \rangle$

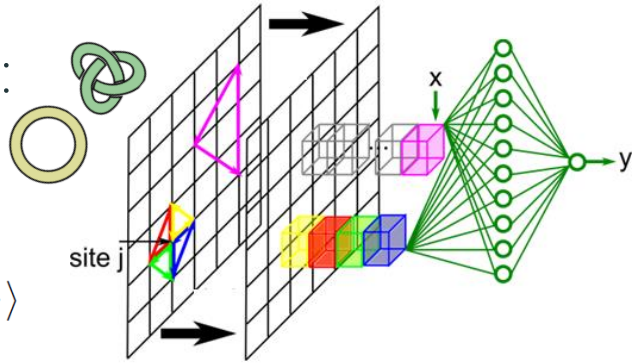

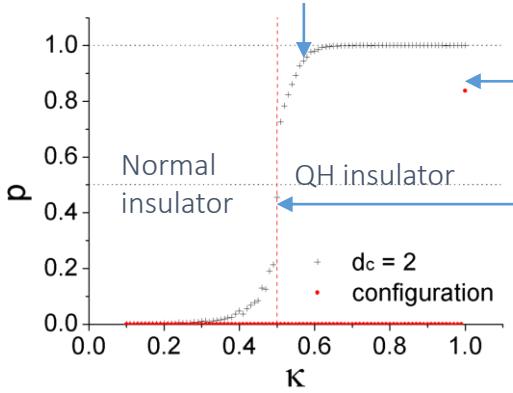

when operators are irrelevant, e.g., $n_r = c_r^\dagger c_r$

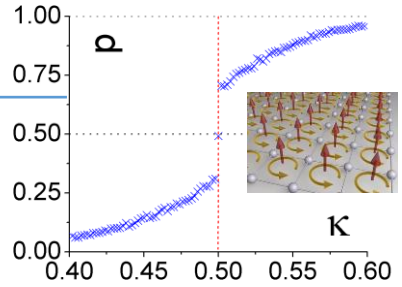

strongly-correlated topological phases:

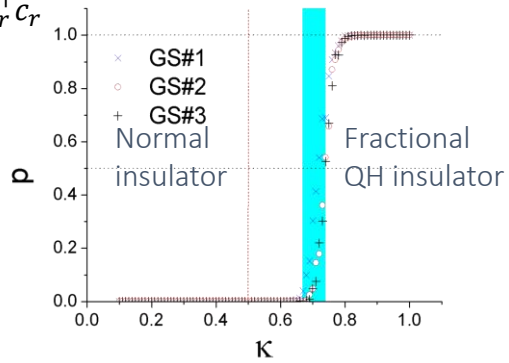

Yi Zhang, E.-A. Kim.

AI for Quantum Phases Recognition

北京大學 PEKING UNIVERSITY

- Interpretation of big, complex experimental data

Quasi-particle interference pattern as a Fermi-surface probe of electron liquids

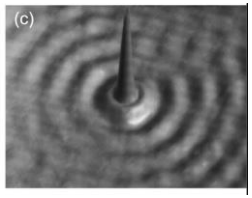

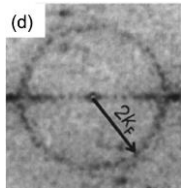

Fourier transform



synergy between experiment and theory

Friedel oscillations as a screening and back-scattering process around a local impurity

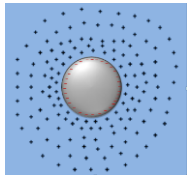

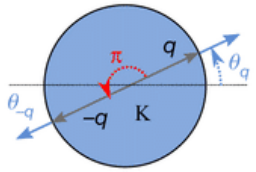

Momentum space

- However, our theoretical capacity largely lags behind our real-world complexity:

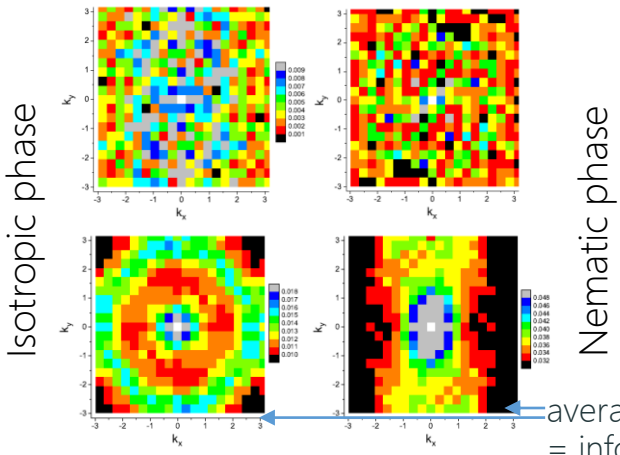

averaged over 1000 FOVs
= information is still present

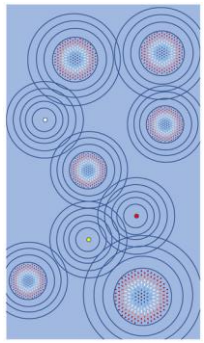

Predict

big data,
noisy data,
many-body,
hidden rules,
etc.

Verify

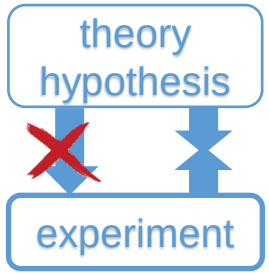

AI interface

Yi Zhang, et al.

AI for Quantum Phases Recognition

北京大學 PEKING UNIVERSITY

- How experiments have evolved: from

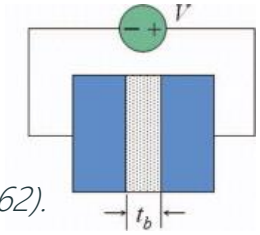

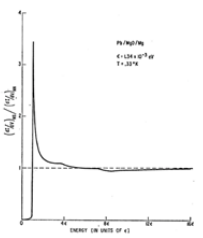

Giaever et al.,
Phys. Rev. ***126****, 941 (1962).*

to

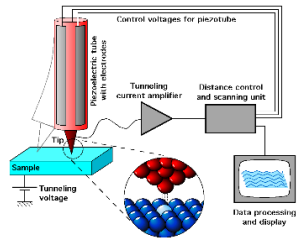

- Idea: Train with the big noisy data, trained for the big noisy data
  - Recognition of nematic phases from STM data

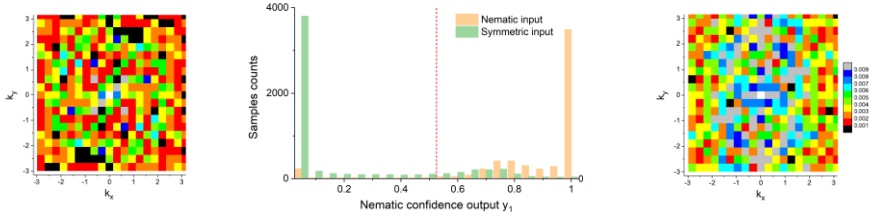

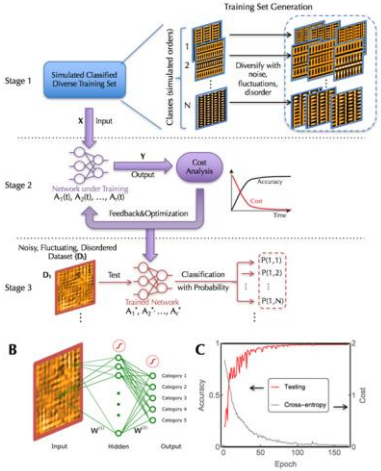

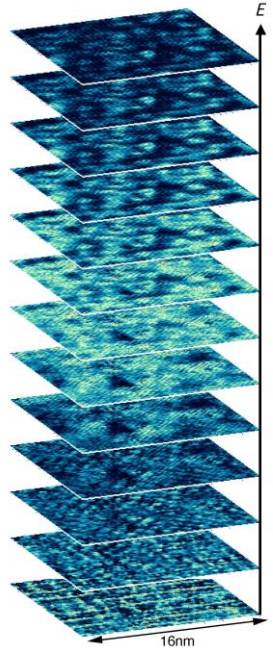

  - Recognition of CDW phases from STM data

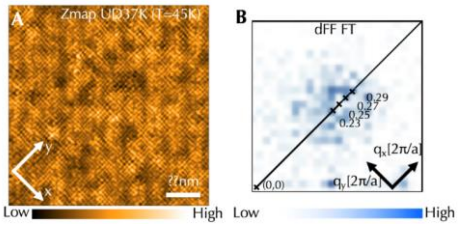

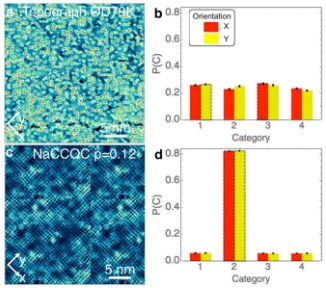

Yi Zhang, et al.

generalizable to other big noisy experimental data, e.g. neutron scattering

Anjana Samarakoon, et al.

AI for Quantum Phases Recognition

北京大學 PEKING UNIVERSITY

A good AI model requires both expressive power and efficient optimization!

- From quantum many-body models, for quantum many-body models:
  - Observation 1: various model parameters

$$H = \sum_{rr'} t_{rr'} c_r^\dagger c_{r'} + \sum_r \mu_r c_r^\dagger c_r + \sum_r U_r c_r^\dagger c_r n_r^f + \cdots$$

  - Observation 2: nonlinear physical properties

LDOS: $y_k = -\frac{1}{\pi}\Im\left[\hat{\mathbf{G}}_{L,L}^{(L)}\right]_{k,k}$ Conductance $y = \sum_m \left|\left[\hat{\mathbf{G}}_{0,L}^{(L)}\right]_{m,1}\right|^2$

  - Observation 3: efficient recursive methods

$$\begin{cases} \hat{\mathbf{G}}_{N,N}^{(N)} = \left[z\hat{I} - \hat{H}_N - \hat{T}_{N-1}^\dagger \hat{\mathbf{G}}_{N-1,N-1}^{(N-1)} \hat{T}_{N-1}\right]^{-1} \\ \hat{\mathbf{G}}_{i,N}^{(N)} = \hat{\mathbf{G}}_{i,N-1}^{(N-1)} \hat{T}_{N-1} \hat{\mathbf{G}}_{N,N}^{(N)} \end{cases}$$

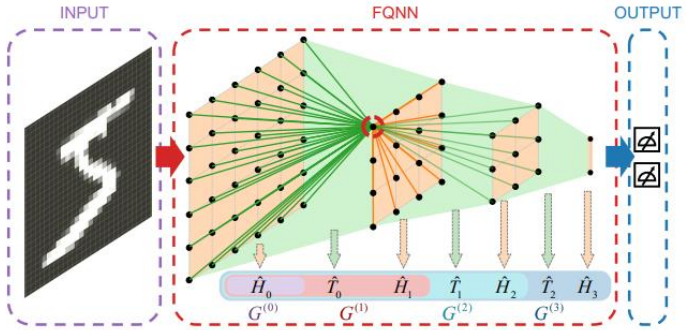

- Direct quantum and topological phase recognitions without presumptions:

map and recognize topological phases via collective properties of FQNN models:

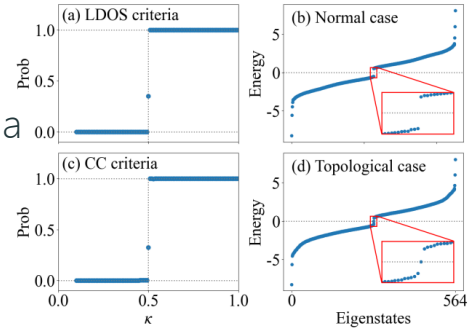

recognize emergent charge density wave phases in strongly correlated models

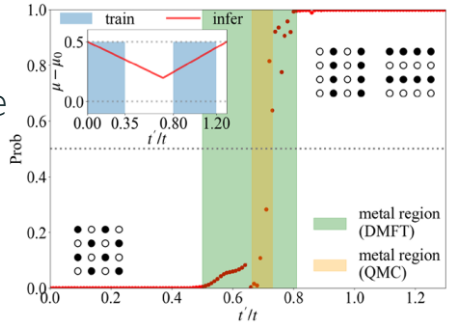

Pei-Lin Zheng, et al.

- Solutions and insights into the vanishing gradient problem:

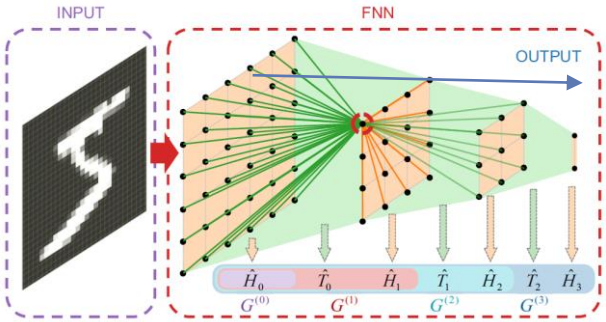

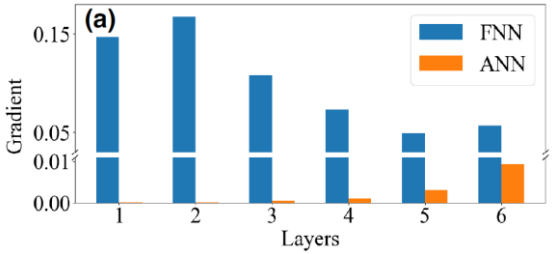

e.g., with global properties (instead of local responses) as output

performance and mechanism comparable to residual blocks ➡ - formal end of the second AI winter

He, K., Zhang, X., Ren, S. and Sun, J., Proceedings of the IEEE conference on computer vision and pattern recognition, 770 (2016). —— *350,000 citations*

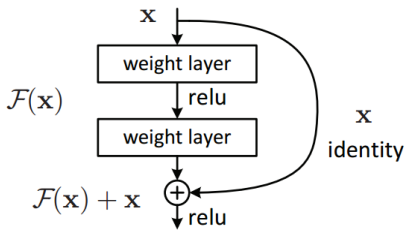

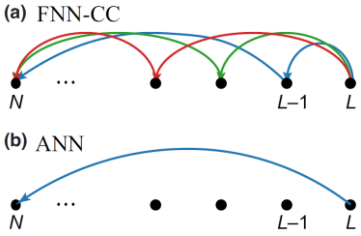

- Quantum toolbox for analysis:

—— e.g., quantum entanglement

- More efficient AI architecture?

—— sample correlations? criticality?

$$S_A = -\mathrm{tr}(\hat{\rho}_A \log \hat{\rho}_A)$$

$$\hat{\rho}_A = \mathrm{tr}_{\bar{A}}(\hat{\rho})$$

$$I_{AB} = S_A + S_B - S_{AB}$$

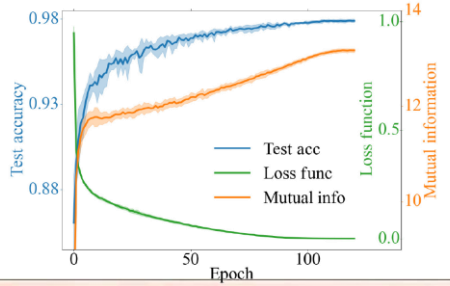

Quantum Many-body Physics for AI

- Strong randomness - Anderson localization:
  - good expressive power
  - hard to optimize

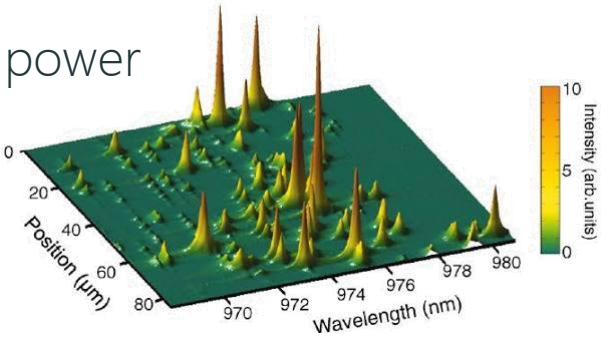



- Weak randomness – metallic states with global response
  - simpler to optimize
  - fewer degrees of freedom – poor expression

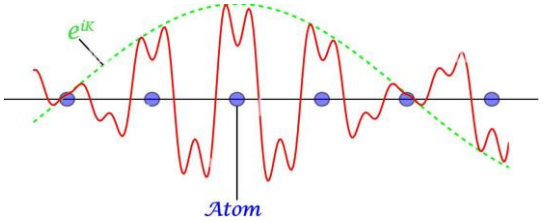

- Critical is good:
  - weak anti-localization
  - topological transitions

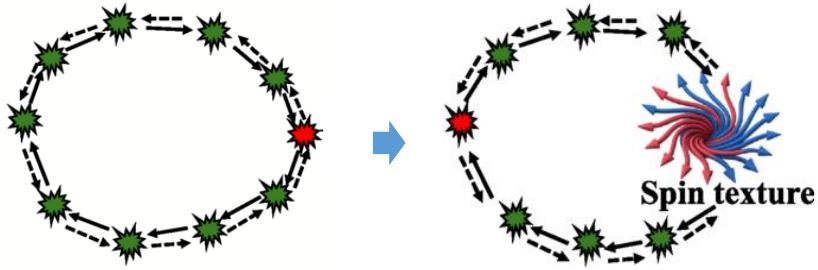

AI for Quantum Monte Carlo Methods

Review: generative AI



- Statistical mechanics: $W(\vec{x}) = \frac{1}{Z}\exp\left[-\frac{E(\vec{x})}{k_B T}\right]$, $Z = \sum_{\vec{x}} \exp\left[-\frac{E(\vec{x})}{k_B T}\right]$
  - Example: Ising model $E(\vec{x}) = -\sum_{\langle ij \rangle} J x_i x_j$, $x_i = \pm 1$

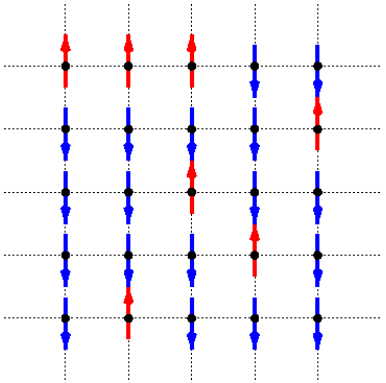

- The Metropolis Algorithm: (also used in simulated annealing)

1. Generate a random initial state $\vec{x}_{t=0}$ with energy $E(\vec{x}_{t=0})$;
2. Flip a random spin $x_i \to -x_i$ and calculate the energy $E(\vec{x}_?)$ of this trial state $\vec{x}_?$;
3. Calculate the difference in energy generated by the spin flip, $\Delta E = E(\vec{x}_?) - E(\vec{x}_t)$;
   - If $\Delta E \leq 0$ (the trial spin state is energetically favorable), accept the spin flip;
   - If $\Delta E > 0$, accept the spin flip with probability $p = \exp(-\Delta E / k_B T)$;
4. Measure the target physical quantities, e.g., energy, magnetization, etc.
5. Repeat steps (2) to (4) until sufficient number N of uncorrelated samples are obtained.

The target probabilities are guaranteed by detailed balance:

$$\frac{W(A)}{W(B)} = \frac{P(B \to A)}{P(A \to B)} = \exp\left(-\frac{E_A - E_B}{k_B T}\right)$$

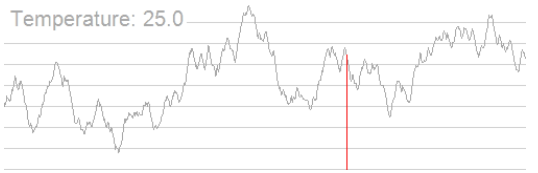

optimization via simulated annealing

MC and Quantum MC Methods

- Also applicable to certain quantum many-body models
  - Auxiliary-Field Quantum Monte Carlo
  - Path Integral Monte Carlo
  - Determinantal Monte Carlo
  - Stochastic Series Expansion Quantum Monte Carlo, etc.

  commonly sampling configurations in (d+1)-dim space-(imaginary)-time

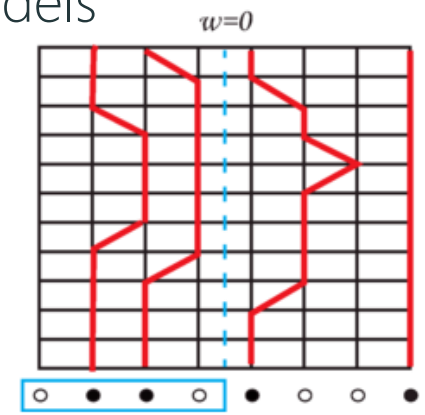

- However, local minima cause critical slowing down – cluster update:
  1. Choose a random site x_i.
  2. Add neighbor site $x_j = x_i$ into the cluster with probability $p = 1 - e^{-2\beta J}$.
  3. Grow the cluster until all neighbors are considered. Flip cluster.

Ulli Wolff, 1989.

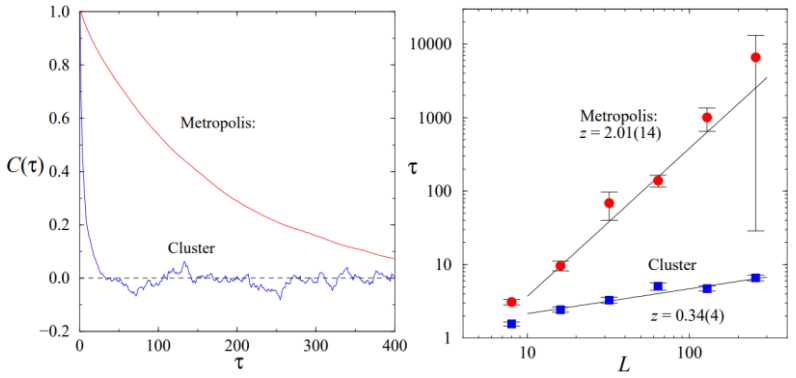

AI for Quantum MC Methods

- Pros and cons of cluster update:
  - probability $W(\vec{x})$ ensured via detailed balance
  - global updates with high efficiency (100% acceptance rate)
  - yet, heavily reliant on the model

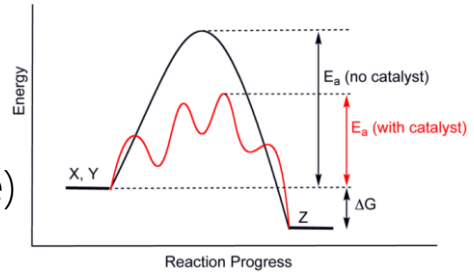

We cannot flip a random cluster with detailed-balance probability, which equals exponentially small acceptance rate!➡globally distinctive states with similar weights ~catalyst

- Idea: fitting $W(\vec{x})$ with an AI model:

  then accept cluster with acceptance rate:

  $$\alpha(A \to B) = \min\{1, e^{-\beta[(E_B - E_B^{\text{eff}}) - (E_A - E_A^{\text{eff}})]}\}$$

  after which the MC is **exact**.

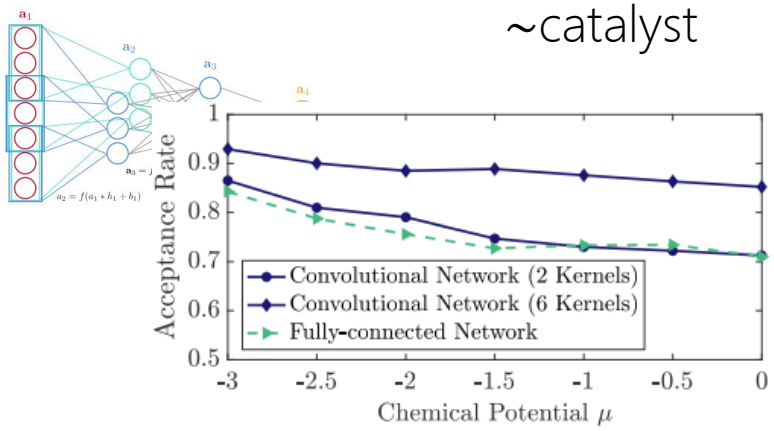

Huitao Shen, Junwei Liu, and Liang Fu.

Generative models

- Graph models: probability distribution with statistical mechanics

$$\mu(v) = \frac{1}{Z} \exp\left\{ \sum_i \theta_i v_i + \sum_{(i,j)\in E} \theta_{ij} v_i v_j \right\}$$

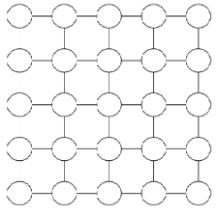

- Restricted Boltzmann Machine

"Restricted"

Similar to and trainable as ANN:

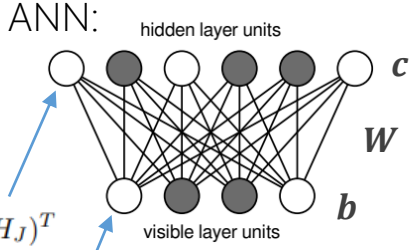

  - A *binary* graphic model with no intra-layer connections
  - The configuration probability follows Boltzmann distribution

$$P(\boldsymbol{X}, \boldsymbol{H}) = \frac{1}{Z} \exp\left(-E(\boldsymbol{X}, \boldsymbol{H})\right)$$

$$E(\boldsymbol{X}, \boldsymbol{H}) = -\boldsymbol{X}^T \boldsymbol{b} - \boldsymbol{c}^T \boldsymbol{H} - \boldsymbol{X}^T \boldsymbol{W} \boldsymbol{H}$$

$$\boldsymbol{H} = (H_1, ..., H_J)^T$$

$$\boldsymbol{X} = (X_1, ..., X_I)^T$$

  - $\boldsymbol{W}$, $\boldsymbol{b}$, and $\boldsymbol{c}$ as model parameters, after training:

maximize the likelihood of given data or fit to a given distribution

Generating handwritten digits:





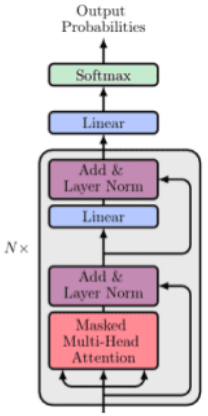

- Generative Pre-trained Transformer (GPT)

AI for Quantum MC Methods

- Example: the Falicov-Kimball model on 2D square lattice

$$\hat{H}_{\mathrm{FK}} = \sum_{i,j} \hat{c}_i^{\dagger} \mathcal{K}_{ij} \hat{c}_j + U \sum_{i=1}^{N} \left(\hat{n}_i - \frac{1}{2}\right)\left(x_i - \frac{1}{2}\right) \qquad \begin{cases} x_i \in \{0,1\} \\ \mathcal{K}_{ij} = -t \quad U/t = 4 \\ \beta = 1/T \end{cases}$$

$$p_{\mathrm{FK}}(\mathbf{x}) = e^{-F_{\mathrm{FK}}(\mathbf{x})} / Z_{\mathrm{FK}}$$

- The trained RBM successfully captures the probability distribution:

compensate with: $A(\mathbf{x} \to \mathbf{x}') = \min\left[1, \frac{p(\mathbf{x})}{p(\mathbf{x}')} \cdot \frac{p_{\mathrm{FK}}(\mathbf{x}')}{p_{\mathrm{FK}}(\mathbf{x})}\right]$

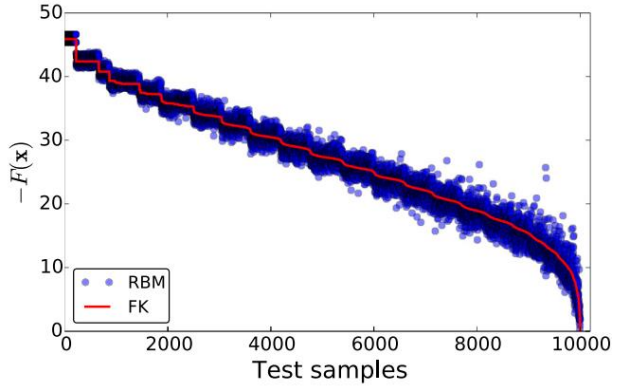

- Nonlocal updates from hidden variables:

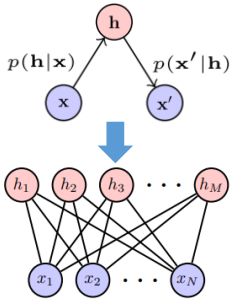

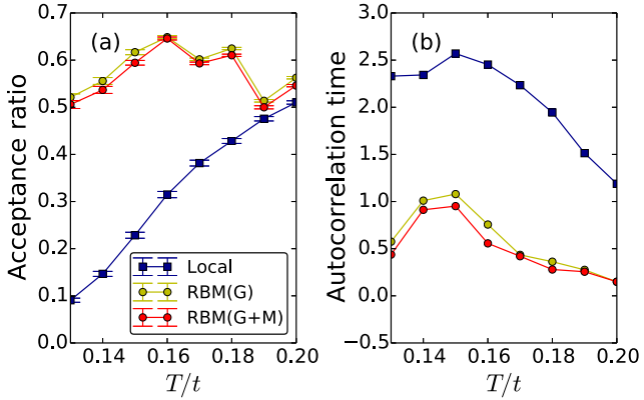

drastically improved acceptance rate and autocorrelation time

Li Huang and Lei Wang.

AI for Quantum Control and Optimization

Review: reinforcement learning



- Classical computer:



Logical gates:

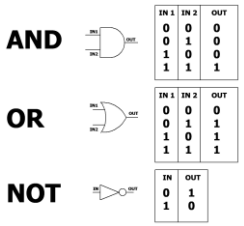

Software ⟺ Logical layer ⟺ Physical Layer

- Quantum computer:

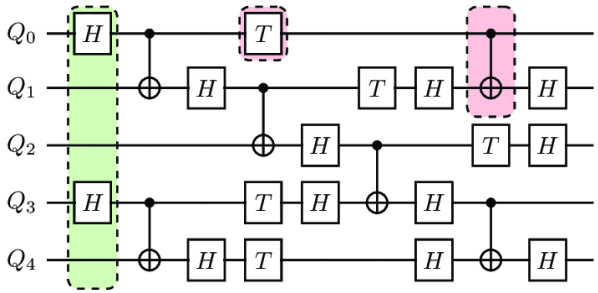

Fundamental quantum gates:

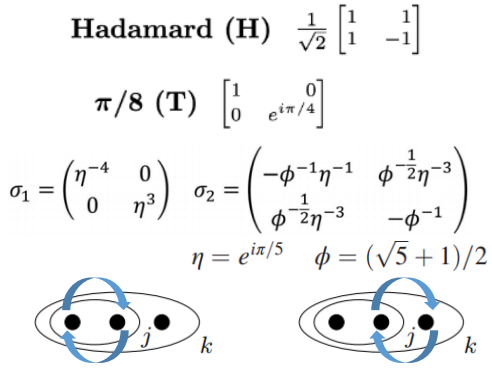

braiding of Fibonacci anyons

Chetan Nayak, et al.

- Goal: find *fast* a *short* sequence $U \approx U_1^{n_1} U_2^{n_2} U_1^{n_3} U_2^{n_4} \cdots$ *close* to U_{tar}
  - brute-force: good length complexity but bad time complexity
  - Solovay-Kitaev (recursive): good time complexity but bad length complexity

Reinforcement learning

- Reinforcement learning: An agent that interacts with an environment and maximizes reward (minimizes penalty)

AlphaGo





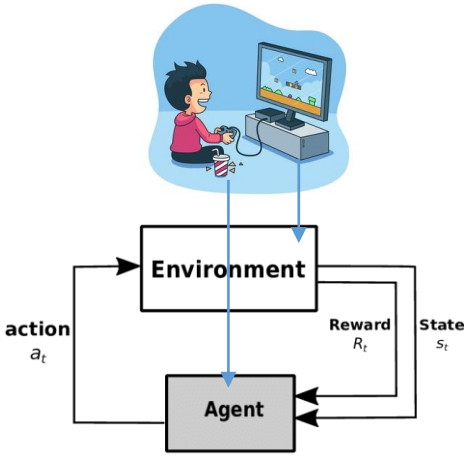

S: current state; A: action upon state; R: reward

$$S \xrightarrow{A} S' \xrightarrow{A'} S'' \xrightarrow{A''} \cdots \rightarrow \text{target}$$

AI ← accumulated reward R

- Video games: S: screen; A: joystick input; R: score, life, cleared levels …
- Chess, Go: S: current board configuration; A: next move; R: win …
- Rubik's cube: S: current colorings; A: next twist; R: (minus) steps taken …

Training the AI model self-consistently with the Bellman equation:

$$Q^{new}(s_t, a_t) \leftarrow \underbrace{Q(s_t, a_t)}_{\text{old value}} + \underbrace{\alpha}_{\text{learning rate}} \cdot \overbrace{\Big(\underbrace{r_t}_{\text{reward}} + \underbrace{\gamma}_{\text{discount factor}} \cdot \underbrace{\max_a Q(s_{t+1}, a)}_{\text{estimate of optimal future value}} - \underbrace{Q(s_t, a_t)}_{\text{old value}}\Big)}^{\text{temporal difference}}$$

$$\underbrace{\qquad\qquad\qquad\qquad\qquad\qquad}_{\text{new value (temporal difference target)}}$$

AI for Quantum Control and Optimization

- Comparison between Rubik's cube and quantum compiling:
  - S: current configuration / unitary U
  - A: applied rotation / elementary gate U_i
  - R: expected distance towards solution

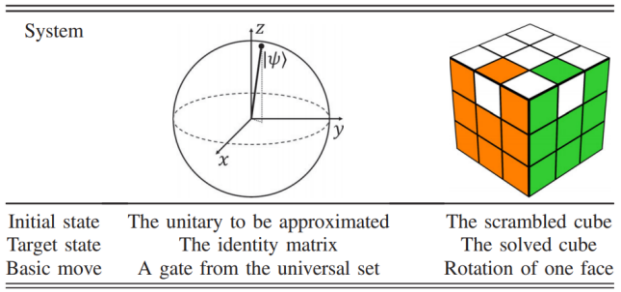

- Combine the cost(-to-go) function:

Q-learning: $J\prime(s) = \min_a(1 + J(A(s,a)))$

with the weighted A* search: $f(x) = \lambda g(x) + J(x)$

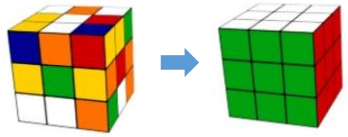

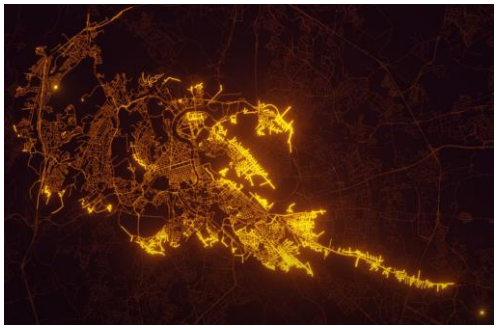

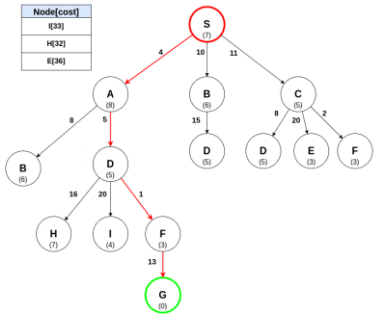

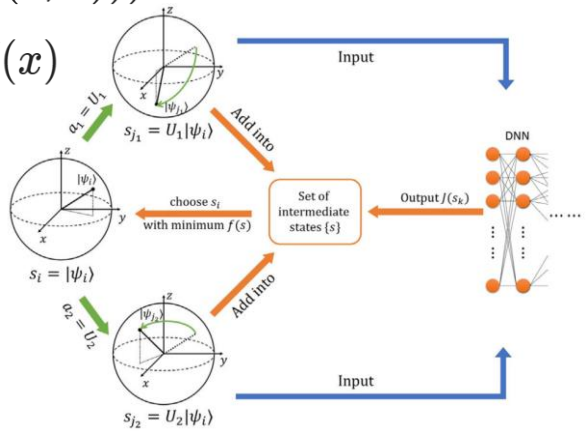

AI for Quantum Control and Optimization

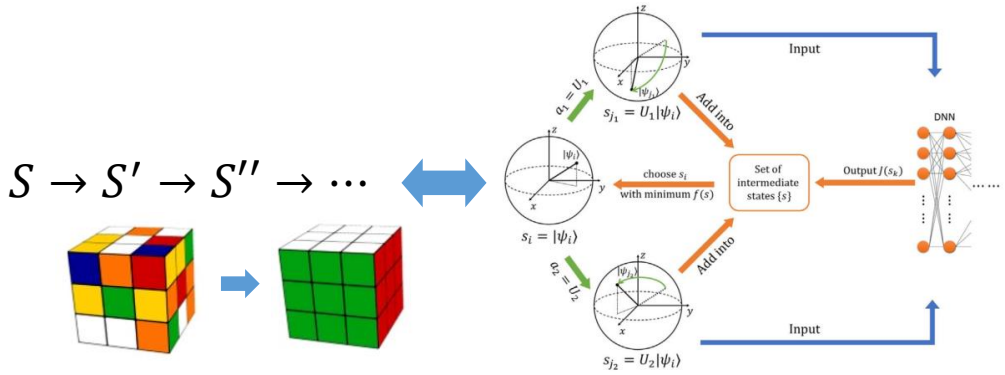

- Time complexity: comparable to the SK recursion, very efficient
- Length complexity: comparable to brute force

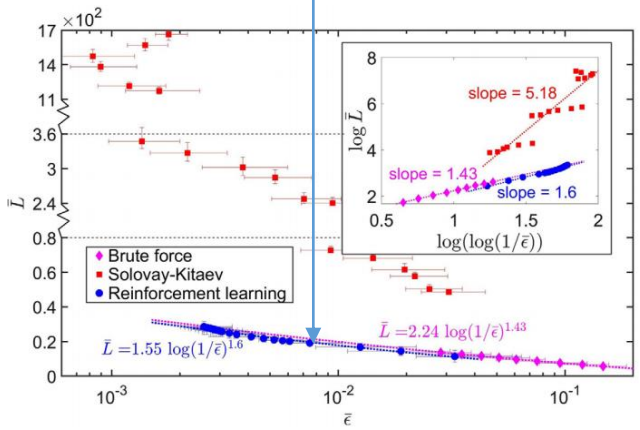

Typical target-unitary examples :

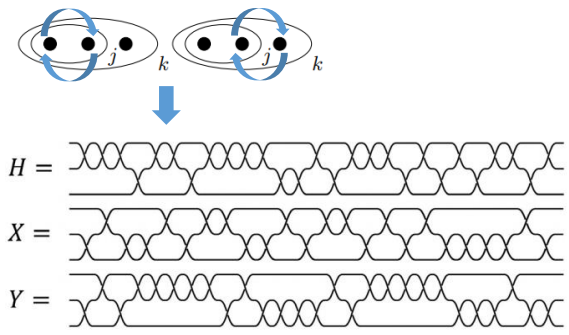

better than $O(10^{-3})$ precision

A novel *good-enough* solver

Yuan-Hang Zhang, Pei-Lin Zheng, Yi Zhang, and Dong-Ling Deng.

AI for Quantum Control and Optimization

- Quantum process – quantum state preparation, e.g., the Dicke state

$$H = \frac{c_2}{2N}\mathbf{L}^2 - q(t)N_0 \qquad c_2 < 0 \qquad |\psi^{(0)}_{\text{Dicke}}\rangle \equiv |N, L_z = 0\rangle$$

  - Conventionally, adiabatic evolution, slowly turn off $q(t)$ to keep at ground state

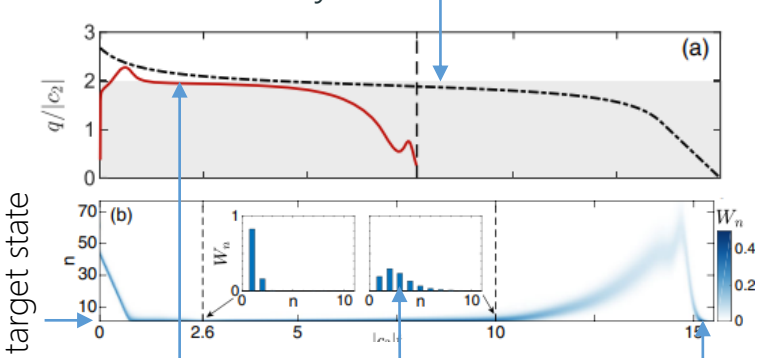

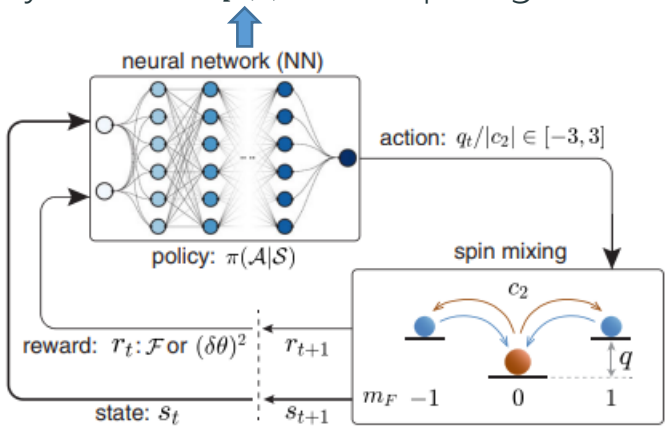

  - AI allows to think outside the box:

A faster process is obtained via reinforcement learning

s_t: $\rho_0 = \langle N_0\rangle/N$, $\langle \delta N_0^2\rangle/N^2$, $|\langle a_{+1}^\dagger a_{-1}^\dagger a_0^2\rangle|/N^2$
$\theta_s = \arg\langle a_{+1}^\dagger a_{-1}^\dagger a_0^2\rangle$

Excited states are generated in the meantime – no adiabaticity

Nevertheless, final state large overlap with target $\mathcal{F} = |\langle\psi(t)|\psi^{(0)}_{\text{Dicke}}\rangle|^2$

Shuai-Feng Guo, et al.

Summary

- AI for quantum phases: numerical and experimental data and models
- AI for quantum methods: synergy and catalyst for algorithmic efficiency
- AI for quantum control: quantum compiling and state preparation
- Discussions:

'A good chef cannot make a decent meal with no ingredients.'

  - No black magic: performance bounded from above by the quality of the samples.
  - Even for the best case scenario, AI methods are approximate.
  - Use the knowledge and intuition to improve, every bit helps!
  - Reverse thinking and consider AI for reverse thinking
  - Sometimes, trying an idea out is the best way to verify its practicality.

Know your target and limitations!

- We are still at an early stage of AI ⊗ Physics.



谢谢

北京大学
PEKING UNIVERSITY

Email: frankzhangyi@pku.edu.cn

AI for Quantum Many-body States

- Quantum state tomography: reconstructs the quantum state by measurements
  - essential for quantum computation, quantum memory, etc.
  - Difficulty: exponential scaling both in the representation and the analysis
  - Setup and requirements:
    - ✓ A quantum system $\hat{\rho}$ that can be prepared repeatedly
    - ✓ A set of quantum measurements
    - ✓ A model that analyze the outcome of the measurements to produce an estimate ρ^*
  - Example: single spinor state $\langle\vec{\sigma}\rangle = tr(\rho\vec{\sigma})$ ← assuming infinitely precise

also, quantum fluctuations and noises

$$\rho = aI + \mathbf{b}\cdot\boldsymbol{\sigma} \Rightarrow \begin{cases} a = \frac{1}{2}\operatorname{tr}\rho = \frac{1}{2} \\ b_i = \frac{1}{2}\operatorname{tr}(\sigma_i\rho) = \frac{1}{2}\langle\sigma_i\rangle \end{cases} \Rightarrow \rho = \frac{1}{2}\left(I + \langle\boldsymbol{\sigma}\rangle\cdot\boldsymbol{\sigma}\right)$$

For $N = 8$ qubits, a brute-force QST requires almost 10^6 measurements.

Article | Published: 29 November 2017

Probing many-body dynamics on a 51-atom quantum simulator

Hannes Bernien, Sylvain Schwartz, Alexander Keesling, Harry Levine, Ahmed Omran, Hannes Pichler, Soonwon Choi, Alexander S. Zibrov, Manuel Endres, Markus Greiner, Vladan Vuletić & Mikhail D. Lukin

Nature **551**, 579–584 (30 November 2017) | Download Citation

Letter | Published: 29 November 2017

Observation of a many-body dynamical phase transition with a 53-qubit quantum simulator

J. Zhang, G. Pagano, P. W. Hess, A. Kyprianidis, P. Becker, H. Kaplan, A. V. Gorshkov, Z.-X. Gong & C. Monroe

Nature **551**, 601–604 (30 November 2017) | Download Citation

AI for Quantum Many-body States

北京大學 PEKING UNIVERSITY

- Neural network quantum state tomography

1. Parameterize (RBMs) the quantum many-body state: $\phi_\mu = \log p_\mu(\mathbf{x})$

$$\psi_{\lambda,\mu}(\mathbf{x}) = \sqrt{\frac{p_\lambda(\mathbf{x})}{Z_\lambda}}\, e^{i\phi_\mu(\mathbf{x})/2} \qquad p_\kappa(\boldsymbol{\sigma},\mathbf{h}) = e^{\sum_{ij} W_{ij}^K h_i \sigma_j + \sum_j b_j^K \sigma_j + \sum_i c_i^K h_i}$$

2. Train to maximize the likelihood of data-set measurements: $|\psi_{\lambda,\mu}(\mathbf{x}^{[b]})|^2 \simeq |\Psi(\mathbf{x}^{[b]})|^2$

Example: 1. a W-state: $|\Psi_W\rangle = \frac{1}{\sqrt{N}}(|100\ldots\rangle + \ldots + |\ldots001\rangle)$

convergence of amplitudes

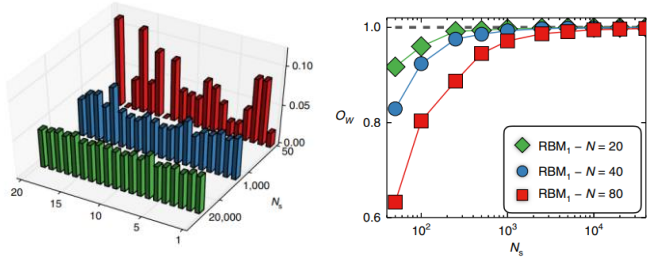

as well as variable phases

$\exp(i\theta(\boldsymbol{\sigma}_k^z)/2)$

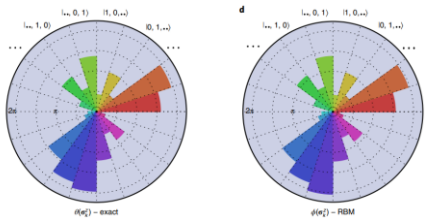

2. ground states of interacting many-body problems

$$\mathcal{H} = \sum_{ij} J_{ij}\sigma_i^z\sigma_j^z - h\sum_i \sigma_i^x \qquad \mathcal{H} = \sum_{ij}\left[\Delta(\sigma_i^x\sigma_j^x + \sigma_i^y\sigma_j^y) + \sigma_i^z\sigma_j^z\right]$$

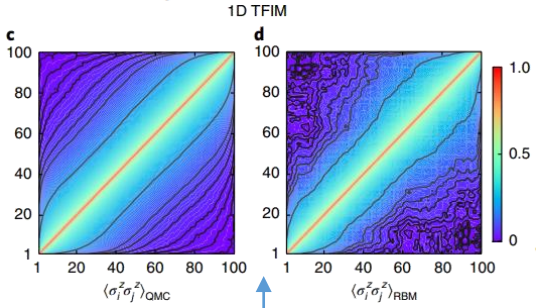

consistent correlations and entanglement entropy

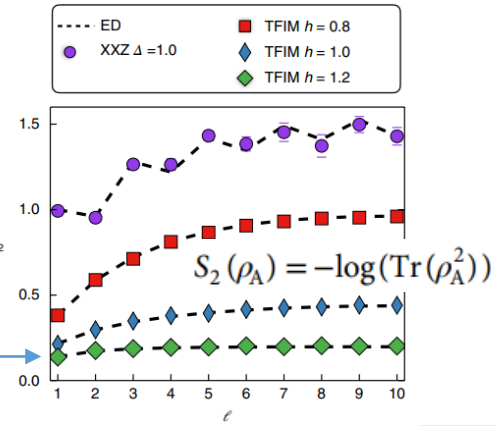

G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. Melko, and G. Carleo.