Title

Efficient representation and learning of quantum many-body states

Author

Hu, Hongye

Publication Date

2022

UNIVERSITY OF CALIFORNIA SAN DIEGO

Efficient representation and learning of quantum many-body states

A dissertation submitted in partial satisfaction of the
requirements for the degree
Doctor of Philosophy

in

Physics

by

Hongye Hu

Committee in charge:

Professor Yi-Zhuang You, Chair
Professor Chunhui Du
Professor Tarun Grover
Professor John McGreevy
Professor Kesong Yang

2022

The dissertation of Hongye Hu is approved, and it is acceptable in quality and form for publication on microfilm and electronically.

University of California San Diego

2022

DEDICATION

This thesis is dedicated to my family, my fiancée, and all my great mentors.

EPIGRAPH

To learn, read.

To know, write.

To master, teach.

—Hindu proverb

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGEMENTS

Time flies and the past six years are among the most memorable periods in my life. I am grateful that I can spend the time with those lovely, intellectual, inspiring people here at the University of California, San Diego, which is located at one of the best places on earth, the sunny La Jolla.

First and foremost, I am deeply indebted to my advisor Dr. Yi-Zhuang You, for his guidance and support. I feel extremely lucky to study under his wing over the past few years. When I started my Ph.D. study at UCSD, I didn't know where is my destination. It feels like standing inside a forest with my bare hands, and I am fortunate to have Yi-Zhuang as my mentor, dear friend, and inspiring role model along this journey. He always encourages me to be myself and follow my heart. Inside the forest of knowledge, there are many paths. There are times when I am afraid I am choosing the wrong one and wasting my time. With Yi-Zhuang's encouragement echoing in my head, I gradually realized maybe there is no right or wrong in the choice. The important aspect is to start my first step, follow my heart, and then make sure I put faith in my choice, work hard and keep going. When looking back, I had tremendous joy in doing research in computational neuroscience, renormalization group and neural network, quantum many-body physics with disorder, and quantum information. All of them form a unique path which leads me to this thesis, *the efficient representation and learning of quantum many-body states*. As a theoretical physicist, Yi-Zhuang has his own way of doing research and learning new knowledge, which is very hands-on. This technique inspires me all the time. He always starts with simple models or problems that one can do some math or solve numerically and play with them. I find this method is very efficient in learning new concepts and gaining intuitions. In addition, Yi-Zhuang is an excellent educator. He can explain everything very clearly to non-expert people. Even hard math becomes very attractive with his words, and I find his clever usage of analogy is very useful in understanding and expressing hard concepts. Yi-Zhuang loves science, and his passion influences me a lot. Whenever I feel blue in real life, talking science with Yi-Zhuang is

one of the best remedies for me. Yi-Zhuang sets a very high standard for me in many aspects. I hope one day, I can be a good scientist, dear friend, and an excellent educator as Yi-Zhuang.

During my Ph.D. study, I feel fortunate to have many good professors, without whom I cannot imagine completing this incredible expedition. Every journey has a starting point. For this, I would like to thank Prof. Congjun Wu, who gave me the chance to study at UCSD. It was a night at the airport, Prof. Wu and I had an excellent discussion, during which time he explained his insights on how to understand electrons as individuals living as a society. I am really grateful I had many chances to discuss with Prof. Wu about physics and science. No matter it is here, or I am at home in China, he always invites me to discuss some science with him. I enjoy those discussions a lot, and I am very impressed by his profound insights in many aspects, such as condensed matter physics, mathematics, and history. In addition, I would thank Prof. John McGreevy, Prof. Tarun Grover, and Prof. Daniel Arovas. All of them are extremely talented scientists. I am grateful that they put the time in developing many wonderful classes, preparing lecture notes, and are willing to pass their knowledge to the younger generation. I really appreciate the fact that they answer all my questions with patience and clearly, even if I know my questions are sometimes naive to them. And, of course, I really thank Prof. Arovas for bringing much joy into the department. I think I will really miss that. I would also like to thank my committee members: Prof. Chunhui Du and Prof. Kesong Yang. I really thank them for supporting me during my study here and meeting me regularly during qualification and other annual meetings. Especially, Prof. Du provided me with much guidance on experimental-related questions and certainly developed my interest in experimental-related research. She is very kind to include me in her group meeting, from where I certainly learned a lot about NV-related experiments and I am very impressed by the techniques developed in her group.

During this journey, I am indebted to many good friends and collaborators, from whom I learned a lot. First, I would like to thank Prof. Soonwon Choi from MIT. I am fortunate to have Soonwon as my collaborator. He is an excellent physicist and has a very high standard in research which is that we should always have a clear physical picture. I think all his collaborators will agree with me on this, and his research style is very apparent from many of his papers. I really learned a lot from our discussion and collaboration. And I hope I can try to reach and keep his standard in all my future research. I would also like to thank Soonwon as a friend. Especially during my postdoc application, he put many personal hours with me in preparing my presentation and going over slides. He would challenge me with any questions and help me understand what is the big picture of the study, which not many people discuss with me. I am really grateful, and I hope I will continue learning with him. Also, I would like to thank Prof. Lei Wang from the Institute of Physics. I really learned most of the machine learning from him, and I am amazed by his profound knowledge in both physics and machine learning. Whenever I have questions regarding numerical studies, he is my go-to person. And Lei is very generous. At the beginning of my research on machine learning, I don't have any computational tools. He let me use many of his resources for free to give me a boosted start during my early research time. In addition, I would like to thank Prof. Koji Hashimoto from Kyoto University. I have been deeply infected by his passion for machine learning and physics since I met him at Microsoft. And I am really thankful that I can collaborate with him on machine learning and holographic duality, during which time I enjoyed learning a lot about the holographic universe from him. I feel really lucky that Koji invited me to give a talk about ”machine learning holography” at the second international conference on machine learning and physics at the Yukawa Institute. During my visit to Japan, I had a lot of fun. For example, I thank Koji for teaching me how to fish and cook Japanese food. I am really lucky to write the Chinese book of *Machine Learning and Physics* with Koji, and hopefully, this book will be useful to other college students in China. Also, my friend Dr. Xun Gao influences me a lot during my Ph.D. study. He has a unique taste of what good research is, and I certainly like his style. He always tells me that good physics researchers need good physical pictures and shouldn't rely on the heavy machinery of mathematics. Especially in the study of quantum machine learning, I think he achieves his own standard, and I am very grateful that I have a chance to collaborate with him. And I would like to thank my mentor, Dr. Zhihui Wang, during my internship with the NASA quantum artificial intelligence lab (QuAIL). I met Zhihui six years ago during the APS March Meeting when I was still an undergraduate. I am surprised that she still remembers me years later. I want to thank her for her support during my internship and my postdoc application. I really learned a lot from her. In addition, I would like to express my gratitude to Prof. Liang Jiang at the University of Chicago. Liang met me seven years ago when I was an undergraduate traveling in a foreign land. He kindly invited me to lunch together at Yale and discussed interesting science with me, including quantum computing and the behavior economics class at Caltech when he was an undergraduate. For many years, he has been a role model for me, which he probably does not know. I am really grateful for that lunch and the friendship we have built. Moreover, I would like to thank Prof. Charles Stevens (Chuck) and Prof. Tatyana Sharpee at Salk institute for biological studies. I really appreciate my first two-year experience at Salk Institute to study neuroscience under their wings. It was fascinating and exciting. Especially, I am very impressed by Chuck's personality. He continued to learn new things such as general relativity and computer programming in his 90s and wrote very clear notes. I am very lucky to have many personal conversations with him and his personal notes.

In addition, I would like to express my sincere gratitude to my parents, Kebiao Hu, Ying Zhang, my fiancée, Ziyan Xu, and my younger brother Hongkai Hu. Even though my parents may not fully understand my research now, they value education over most things and have supported my study since a very young age. So I was fortunate to have many books when I was a boy. During those time, I came across the book, *Surely you are joking, Mr. Feynman!*, that opened a new world for me. I remember I disassembled my 8th birthday gift from my mom, which was a controllable robot, to figure out why the robot could move, and made an electric fan using the engine of the robot. Looking back, I am so lucky that they were not angry with me but encourage me being curious. Also, I appreciate my dad buying those English books for me whenever he has a chance during his busy oversea errands. They really place the seed for my future scientific journey intentionally or unintentionally. In addition, my parents support me doing research and I cannot achieve what I have now without their support. I feel so lucky to such a family! And, of course, I love my little brother Hongkai. He is very lovely and smart. It is fortunate to have a such a brother. During my adulthood, I am really lucky to have Ziyan. She is a very caring and supportive partner and a bright young immunologist. I am indebted to her in many ways. Sometimes, I work very late or have discussions with my collaborators during weekends or late at night. She usually sacrifices her time and compromises with me. From time to time, she jokes about it: ”I have a finding that physicists do not need sleep.” Sometimes life can be rough, and I am really grateful for her companionship. She really has her own magical way to make me feel better. I grow into a better person with her. And we both are very lucky to have a lovely corgi, Tinkerbell.

Last but not the least, I deeply appreciate all my friends and collaborators, with whom I shared lifelong stories: Prof. Biao Wu, Prof. Frank Wilczek, Prof. Susanne Yelin, Prof. Romain Vasseur, Prof. Vedika Khemani, Prof. Bruno Olshausen, Prof. Xiao-Liang Qi, Prof. Chao-Ming Jian, Prof. Charles Stevens, Prof. Tatyana Sharpee, Dr. Hongchen Jiang, Dr. Eleanor Rieffel, Shang Liu, Robert Hsin-Yuan Huang, Xue-Yang Song, Eric Anschuetz, Wei-Ting Guo, Jordan Cotler, Gurtej Kanwar, Dachuan Lu, Ahmed Akhtar, Ruben Verresen, Jinlong Huang, Handuo Shi, Shenxiu Liu, Yi Liu, Yijun Jiang, Hengyun Zhou, Shengtao Wang, Wenbo Fu, Dian Wu, Yijian Zou, Wenzhuo Huang, Xiong Lin, Meng Zeng, Yilun Zhang, Di Luo, Junyi Shan, Honglie Ning, Xiao Zhang, Fangli Liu, Xiang Ji, Li Jing, Xizhi Han, Wenfan Chen, Jinguo Liu, Tianci Zhou, Xunchen You, Xudong Lv, Fangming Xie, Charles Chunjun Cao, Yubei Chen, Lun-Hui Hu, Shuo-Hui Li, Zengyi Li, Pan Li, Ziming Liu, Tailin Wu, Wang Yang, Yifei Zhang, Donghao Wang, Xiaoyan Xu, Qingdong Jiang, Wenjie Yao, Tsung-Cheng Lu, Ryan LaRose, Ryan Levy, the Swarma Club, Senrui Chen, Zeng Pei, Junyu Liu, Shangnan Zhou, Xiaoyang Huang, Dian Wu, Dan Mao, Jinghui Liu, Xiu-Zhe Luo, Boqun Shi, Dr. Amanda Han, Yimu Bao, Gushu Li, Zhaoran Wang, Junyi Cao, Sharmila Poddar, Bingtian Ye, Jack Berkowitz, Wei-Mien Mendy Hsu, Shon Grabbe, James Sud, Raimel A. Medina, Sona Najafi, Jahan Claes, Ze-Pei Cian, Sisi Zhou, Haoxiang Wang, Wanda Hou, Kaifeng Bu, Wenping Cui, Yao Wang, Chenguang Guan.

Chapter 5, in total, is a reprint of the material as it appears in Hong-Ye Hu, Yi-Zhuang You, Physical Review Research 4, 013043. The dissertation author was the primary investigator and author of this paper.

Chapter 6, in total, is a reprint of the material as it appears in Hong-Ye Hu, Soonwon Choi, Yi-Zhuang You, arXiv:2107.04817. The dissertation author was the primary investigator and author of this paper.

Chapter 7, in total, is a reprint of the material as it appears in Hong-Ye Hu*, Carlos M. Duque*, Yi-Zhuang You, Vedika Khemani, Ruben Verresen, Romain Vasseur, Physical review B 103, L100207. The dissertation author was the primary investigator and author of this paper.

Chapter 8, in total, is a reprint of the material as it appears in Hong-Ye Hu, Shuo-Hui Li, Lei Wang, Yi-Zhuang You, Physical Review Research 2, 023369. The dissertation author was the primary investigator and author of this paper.

VITA

2016	B. S. in Physics, *cum laude*, Peking University, China
2022	Ph. D. in Physics, University of California San Diego

PUBLICATIONS

Hong-Ye Hu, Yi-Zhuang You, *Hamiltonian-Driven Shadow Tomography of Quantum States*, Physical Review Research 4, 013043 (2022).

Chenhua Geng, **Hong-Ye Hu**, Yijian Zou, *Differentiable Programming of Isometric Tensor Networks*, Machine Learning Science and Technology 3, 015020 (2022).

Hong-Ye Hu, Soonwon Choi, Yi-Zhuang You, *Classical Shadow Tomography with Locally Scrambled Quantum Dynamics*, arXiv:2107.04817 (2021).

Hong-Ye Hu*, Carlos M. Duque*, Yi-Zhuang You, Vedika Khemani, Ruben Verresen, Romain Vasseur, *Topological and Symmetry-enriched Random Quantum Critical Points*, Physical Review B 103, L100207 (Editors' Suggestion 2021).

Koji Hashimoto, **Hong-Ye Hu**, Yi-Zhuang You, *Neural Ordinary Differential Equation and Holographic Quantum Chromodynamics*, Machine Learning: Science and Technology 2, 035011 (2021).

Meng Zeng, Lun-Hui Hu, **Hong-Ye Hu**, Yi-Zhuang You, Congjun Wu, *Phase-fluctuation Induced Time-Reversal Symmetry Breaking Normal State*, arXiv:2102.06158 (2021).

Hong-Ye Hu, Dian Wu, Yi-Zhuang You, Bruno Olshausen, Yubei Chen, *RG-Flow: A hierarchical and explainable flow model based on renormalization group and sparse prior*, arXiv:2010.00029 (2020). Under review by Journal of Machine Learning Research.

Hong-Ye Hu, Shuo-Hui Li, Lei Wang, Yi-Zhuang You, *Machine Learning Holographic Mapping by Neural Network Renormalization Group*, Physical Review Research 2, 023369 (2020).

Frank Wilczek, **Hong-Ye Hu**, Biao Wu, *Resonant Quantum Search with Monitor Qubits*, Chinese Physics Letters 37 (5), 050304. (Editors' Suggestion 2020).

Xiao-Yang Huang, Taige Wang, Shang Liu, **Hong-Ye Hu**, Yi-Zhuang You, *Quantum Magnetism in Wannier-Obstructed Mott Insulators*, arXiv:2005.01439 (2020).

Hong-Ye (Hongye) Hu, Biao Wu, *Optimizing Quantum Adiabatic Algorithm*, Physical Review A 93 (1), 012345 (2016).

ABSTRACT OF THE DISSERTATION

Efficient representation and learning of quantum many-body states

by

Hongye Hu

Doctor of Philosophy in Physics

University of California San Diego, 2022

Professor Yi-Zhuang You, Chair

Quantum mechanics is one of the most successful and striking theories in physics. It predicts atomic particles can have exotic properties, such as quantum entanglement, that any classical local theory cannot describe. This phenomenon dramatically increases the complexity of nature, and it indicates there is no classical algorithm that can universally simulate all quantum many-body states. On the other hand, as opposed to classical systems, we never observe quantum properties directly since the measurement for the quantum systems is destructive. People can only determine the quantum black box by the statistics of classical readouts. The complexity of quantum objects implies exponentially many measurements and classical data to figure out the quantum states fully. It is underlying those challenges to find an efficient classical representation

of quantum many-body states. An efficient (classical) representation will require fewer classical data of quantum states and learn many of its properties. And an efficient representation can also be served as a classical simulation algorithm for the quantum states. The is no universal, efficient representation for all the quantum states, and it usually depends on the learning properties or underlying quantum states. This thesis will give two efficient representations: the classical shadow representation of quantum states and the hierarchical representation of quantum states. We will see that those efficient representations will help us learn and simulate quantum many-body states and lead to many critical applications in quantum information technology, condensed matter physics, and quantum field theory.

Part I

Introduction

Chapter 1

Background and introduction

Quantum mechanics is one of the most successful, striking, and controversial theories in physics. It predicts atomic particles can have exotic properties, such as quantum entanglement, a quantum correlation that does not have any classical analog [EPR35, Sch35, Sch36]. Einstein suspected that quantum mechanics was an incomplete description of reality and the troublesome fundamental probabilities of quantum mechanics may be originated from the ignorance of some hidden degrees of freedom. Therefore, he hoped to construct some classical stochastic theory, like statistical mechanics, within the framework of classical mechanics, which he dubbed ”local hidden variable theory”. However, this dream was shattered by John Bell in 1964 [Bel64a]. Incredibly, he shows that any local hidden variable theory is incompatible with quantum mechanics, and quantum mechanical probabilities cannot arise from the ignorance of pre-existing local variables. Otherwise, superluminal communication is possible. Nowadays, people know quantum entanglement is an essential property that gives rise to the Bell’s inequality. This phenomenon dramatically increases the complexity of nature. And the complexity poses two difficulties in understanding quantum many-body states: 1) it is really hard to fully characterize what is the quantum state in the black box by experimental measurements, and 2) it is really hard to find algorithms to simulate quantum states on a classical computer.

In classical mechanics, a system containing N particles only needs $6N$ or, in general, $\text{Poly}(N)$ parameters to fully characterize the system's state. In contrast, a quantum system composed of N qubits (quantum two-level systems) requires approximately 2^N complex numbers to represent its state fully. On top of the exponential scaling of parameters, quantum measurements are destructive. The quantum state will collapse to one of the eigenstates of the observable each time after measurement. Combining those factors, people find it hard to unveil the secret of quantum states. More rigorously, people prove that the number of measurements needed to characterize quantum states fully grows exponentially with system size. And this process of figuring out quantum states by experiments is called *quantum state tomography*. Quantum state tomography is essential in many quantum technology applications, such as quantum device benchmarking and discrimination of quantum states. While reconstructing the full classical representation, the density matrix, becomes unfeasible with increasing system size, predicting a collection of (possibly exponentially many) properties of quantum states with only a polynomial number of state copies and measurements is still possible. Therefore, we would like to find a representation of quantum states such that it only requires a polynomial number of classical data for us to learn (possibly exponentially) many properties of quantum states. Those representations are efficient. There is no universal, efficient representation for all quantum states. Usually, that depends on the assumption of underlying states, such as low entanglement states, stabilizer states, or depends on the collection of properties that we want to learn, such as low-rank observables, local observables. Some efficient representations based on quantum state assumptions include matrix product state tomography [CPF$^+$10b, LMH$^+$17a, WHW$^+$17], reduced density matrix tomography [LPW02, LW02, CJZZ12, XLC$^+$19] and machine learning tomography [CTMA19, TM18, TTv$^+$19]. And more recently, a new efficient representation called shadow tomography has attracted much attention [HKP20]. Based on statistical learning theory, this new scheme does not rely on assumptions of underlying quantum states, and its efficiency depends on the collection of observables.

When studying physical states beyond the exactly solvable limit, physicists use computers to simulate their behavior to get a better understanding. However, this approach breaks down when the system of interest is quantum mechanical. As we have seen, the general quantum states require exponentially many complex numbers to be characterized. Only saving those data on a classical computer can be troublesome. An efficient representation of quantum states can also be served as a good ansatz for numerical simulation, which does not require exponential many computational resources. For example, matrix product state is efficient representation for low-entanglement states [Vid07a, Vid08a], and stabilizer formalism is efficient representation for stabilizer states [AG04, Got98a], which can be created with Clifford circuits on product all-zero states.

As we can see, finding efficient representations of quantum many-body states is an essential task in modern physics. It can help physicists not only learn properties of underlying states from very few experimental measurements but also serve as an efficient ansatz for numerical simulations. Motivated by this overarching question, I will introduce two important efficient representations: the classical shadow representation of quantum states and the hierarchical representation of quantum states. In Chapter 2 to Chapter 4, I will review the background knowledge in quantum mechanics, random unitary designs, and machine learning. Chapter 5 and 6 form part I of the thesis, where I will introduce classical shadow tomography with chaotic dynamics [HY22, HCY21]. Especially, I solve the reconstruction map of classical shadow representation under arbitrary locally scrambled unitary dynamics. The new solution explicitly shows how the entanglement created in quantum dynamics can help with quantum measurements. And this new proposal of classical shadow tomography is the first one that does not rely on the group structure of random measurements. Chapter 7 and 8 form the second part of the thesis, where hierarchical representation of quantum states is discussed. The hierarchical representation is very closely related to holographic duality and renormalization group. Especially in chapter 7, I will introduce the spectrum bifurcation renormalization group and its application to simulation of strongly disordered quantum systems [DHY$^+$21]. In chapter 8, I will introduce the machine learning renormalization group and show the relation between this invertible renormalization group and exact holographic mapping [HLWY20]. And I will also illustrate some applications of this machine learning holographic mapping, such as efficient simulation of critical field theory and finding holographic dual theory.

Chapter 2

Quantum mechanics

2.1 Closed quantum systems

2.1.1 Axioms of quantum mechanics

Quantum mechanics is a mathematical model for the physical world. It successfully describes many aspects of fundamental particles that no classical theory can achieve. The theory of a closed quantum system is based on *five axioms* for 1) quantum states, 2) physical observables, 3) measurement, 4) quantum evolution, and 5) composition of systems. In this section, we review the basics of axioms for closed quantum systems. More discussions can be found in [Pre18a].

Axiom 1: Quantum states. A state is the complete description of the physical system, and it is described by a *ray* in the *Hilbert space*.

First, a Hilbert space is a linear vector space on complex numbers $\mathbb{C}$ with an inner product. The vectors in the Hilbert space can be written as $|\psi\rangle$ using the Dirac notation. The inner product $\langle\psi|\phi\rangle$ is the map that takes two vectors $|\psi\rangle$ and $|\phi\rangle$ to a complex number $\mathbb{C}$, which has the following properties:

1. $\langle\psi|\psi\rangle > 0$ for any non-zero vectors.

2. $\langle\psi|(a|\phi\rangle + b|\varphi\rangle) = a\langle\psi|\phi\rangle + b\langle\psi|\varphi\rangle$.

3. $\langle\psi|\phi\rangle = \langle\phi|\psi\rangle^{*}$[1]

Not all different vectors in the Hilbert space describe different physical states. If two vectors differ by a non-zero complex number, i.e., $|\psi\rangle = \alpha|\phi\rangle$ where $\alpha \in \mathbb{C}$, then they are in the same ray. So a *ray* is the equivalent class for vectors that only differ by a non-zero complex number. Usually, we choose the representation of a *ray* to be normalized, such that $\langle\psi|\psi\rangle = 1$. The representation vector $|\psi\rangle$ is the description of a physical system. And one should also notice that the normalization condition doesn't fix the phase of vector, so $|\psi\rangle$ and $e^{i\theta}|\psi\rangle$ describe the same physical system. This is also called gauge redundancy.

Axiom 2: Physical observables. Physical observables are quantities that can be measured, and they are described by *self-adjoint* operators.

Operators are linear maps in the Hilbert space $\mathbf{O} : |\psi\rangle \to \mathbf{O}|\psi\rangle$, and self-adjoint operators satisfy $\mathbf{O}^\dagger = \mathbf{o}$. Self-adjoint operators can be diagonalized or have spectral representation, such that

$$\mathbf{o} = \sum_n a_n \mathbf{E}_n, \tag{2.1}$$

where $\mathbf{E}_n$ are orthogonal projectors, and a_n are eigenvalues. And eigenvalues a_n are possible readouts for measuring physical observable $\mathbf{O}$.

Axiom 3: Measurements. A measurement is the process of getting information of the physical system through a observable $\mathbf{O}$. After the measurement, the observer get a classical readout that is one of the eigenvalues a_n of observable $\mathbf{O}$, and the state *collapse* to the corresponding eigenstate. The probability of getting readout a_n is $p(a_n) = \langle\psi|\mathbf{E}_n|\psi\rangle$. So the expectation value of the observable $\mathbf{O}$ is

$$\langle\mathbf{O}\rangle = \sum_n a_n p(a_n) = \sum_n a_n \langle\psi|\mathbf{E}_n|\psi\rangle = \langle\psi|\mathbf{O}|\psi\rangle \tag{2.2}$$

---

[1] Here, asterisk means the complex conjugate.

Axiom 4: Dynamics. The evolution of a *closed* quantum system is governed by unitary evolution: $|\psi(t')\rangle = U(t',t)|\psi(t)\rangle$. The infinitesimal evolution is generated by a *self-adjoint* operator $H(t)$, which is called the Hamiltonian of the system. In particular, $U(t+dt,t) = \mathbb{1} - iH(t)dt$. And we can also write down the partial differential equation for the dynamics of quantum states, which is the Shrodinger equation:

$$\mathbb{i}\frac{\partial|\psi(t)\rangle}{\partial t} = H(t)|\psi(t)\rangle. \tag{2.3}$$

Axiom 5: Composition of systems. Suppose there are two systems: A and B with corresponding Hilbert spaces $\mathcal{H}_A$ and $\mathcal{H}_B$. Then the Hilbert space of the composite system $A+B$ will be the tensor product $\mathcal{H}_A \otimes \mathcal{H}_B$. Those five axioms conclude the mathematical foundation of quantum theory for a closed quantum system. And such a model already comprises a lot of interesting phenomena that cannot be captured by classical models, such as entanglement.

2.1.2 Schmidt decomposition

Before I introduce open quantum states, let me review a beneficial standard form for pure bipartite quantum states, called Schmidt form or Schmidt decomposition. Suppose we have a bipartite system composed of A and B. According to the fifth axiom, the total Hilbert space is $\mathcal{H}_A \otimes \mathcal{H}_B$, and any quantum state can be written as

$$|\psi\rangle = \sum_{i,a} \psi_{a,i}|a\rangle_A \otimes |i\rangle_B, \tag{2.4}$$

where $\{|a\rangle_A\}$ and $\{|i\rangle_B\}$ are orthonormal basis in Hilbert space $\mathcal{H}_A$ and $\mathcal{H}_B$ respectively. We can view the coefficients $\psi_{a,i}$ as a matrix, which is not necessarily square. Any matrix will have a singular value decomposition (SVD), so that

$$\psi_{n\times p} = U_{n\times n}\Lambda_{n\times p}V^{\dagger}_{p\times p}, \tag{2.5}$$

where U and V are unitary matrices, i.e. $UU^{\dagger} = \mathbb{I}$, $VV^{\dagger} = \mathbb{I}$, and $V_{n\times p}$ is a diagonal matrix. If we plug the SVD decomposition into Eq.(2.4), then we have

$$\begin{aligned} |\psi\rangle &= \sum_{i,a,\alpha} U_{a,\alpha}\lambda_{\alpha,\alpha}V^{*}_{i,\alpha}|a\rangle_A \otimes |i\rangle_B \\ &= \sum_{\alpha} \lambda_{\alpha,\alpha}|\alpha\rangle_A \otimes |\alpha\rangle_B, \end{aligned} \tag{2.6}$$

where we use U and V to define new basis as $|\alpha\rangle_A = \sum_a |a\rangle_A U_{a,\alpha}$ and $|\alpha\rangle_B = \sum_i |i\rangle_B V^{*}_{i,\alpha}$. Or in matrix form, the state in new basis is

$$|\psi\rangle = \begin{pmatrix} |0\rangle_A & \cdots & |\alpha\rangle_A & \cdots \end{pmatrix} \begin{pmatrix} \lambda_0 & 0 & 0 & 0 \\ 0 & \ddots & 0 & 0 \\ 0 & 0 & \lambda_\alpha & 0 \\ 0 & 0 & 0 & \ddots \end{pmatrix} \begin{pmatrix} |0\rangle_B \\ \vdots \\ |\alpha\rangle_B \\ \vdots \end{pmatrix}. \tag{2.7}$$

The above equation is the Schmidt form of a pure quantum state, and it exists for any pure bipartite quantum state. However, in general, one cannot write two states $|\psi\rangle_{AB}$ and $|\phi\rangle_{AB}$ in Schmidt form simultaneously. The basis can be different. One should also notice the ambiguity in the Schmidt form. If there is degeneracy in the singular values λ_i, then one can apply any unitary transformation O in the degenerate subspace to rotate the basis, such that $U \to UO$ and $V \to VO$. This simultaneous transformation changes the basis $\{|i\rangle_A\}$, and $\{|i\rangle_B\}$ for A and B but leaves $|\psi\rangle$ unchanged.

2.2 Tensor diagram: a Lego for understanding quantum mechanics

Before we further discuss open quantum systems, I would like to introduce a handy tool which is called *tensor diagram*. In the literature, people also call it *tensor networks* [BB17, Oru14], or *graphical calculus* [WBC15, MPM17]. It will help us understand concepts like the Feynman diagram for quantum field theory. And we will find it very helpful in the following discussions. First, let me clarify some notations. For Hilbert spaces, instead of using $\mathcal{H}_X$,$\mathcal{H}_Y$, I will use $\mathcal{X}$, $\mathcal{Y}$ for simplification. And $L(\mathcal{X},\mathcal{Y})$ is the space of bounded linear operators: $A : \mathcal{X} \to \mathcal{Y}$. Tensors are the basic objects we will be dealing with. They are multi-dimensional arrays of complex numbers with a fixed standard basis. For example, scalars are rank-0 tensors, and vectors are rank-1 tensors with one index, $v_i = \langle i|v\rangle$, and matrices are rank-2 tensors with two indexes, $A = \sum_{i,j} A_{i,j}|i\rangle\langle j|$, where $A_{i,j} = \langle i|A|j\rangle$. If the distinct index can be thought of as legs, then we can represent elementary tensors as graphs as shown in Fig. 2.1.

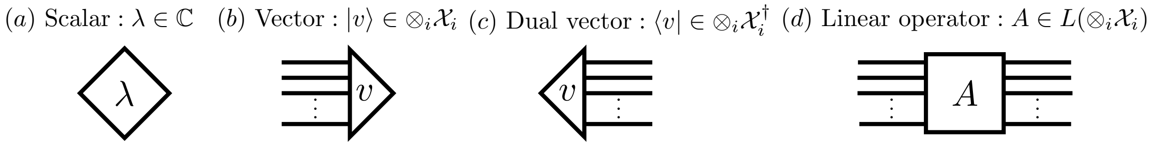
(a) Scalar : $\lambda \in \mathbb{C}$ (b) Vector : $|v\rangle \in \otimes_i \mathcal{X}_i$ (c) Dual vector : $\langle v| \in \otimes_i \mathcal{X}_i^\dagger$ (d) Linear operator : $A \in L(\otimes_i \mathcal{X}_i)$

Figure 2.1: Tensor diagrams for elementary tensors.

As we see in Fig. 2.1, to avoid confusion, we should read the whole diagram from left to right. And opening legs pointing to the left indict vectors in basis $|i\rangle$, while the opening legs pointing to the right indict vectors in the dual-basis $\langle i|$. In the next chapter, I will also introduce another convention with arrows on the legs. Now let me introduce the color summation convection. If elementary objects share the same color, they will be summed over the shared basis. For example, the following diagram is the same as $I = \sum_i |i\rangle\langle i|$:

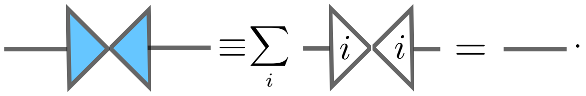

(2.8)

And the (unnormalized) Bell state $|\Phi^+\rangle = \sum_i |i\rangle \otimes |i\rangle$ and its dual-state $\langle\Phi^+| = \sum_i \langle i| \otimes \langle i|$ are very useful:

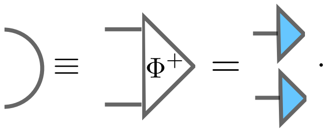

(2.9)

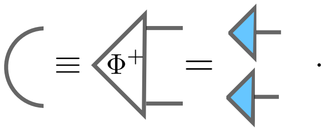

(2.10)

With these basic definition, we can already get some interesting results:

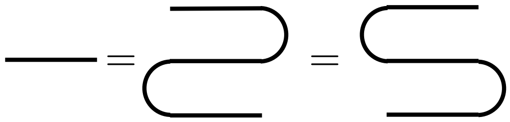

(2.11)

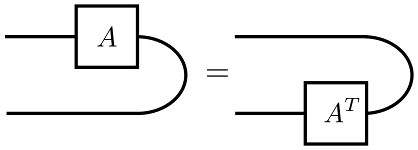

(2.12)

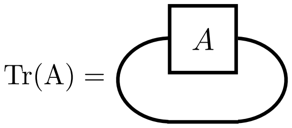

(2.13)

Another important concept is vectorization, which reshapes a $m \times n$ matrix into a mn vector. There is a freedom to choose the order of reshaping in reshaping elements. Typically, there are row-first and column-first orders. For row(column)-first ordering, we stack rows (columns) of the matrix on top of each other. In literature [WBC11], they are sometimes called row-vector and column vector respectively. I will choose the following convention in defining them:

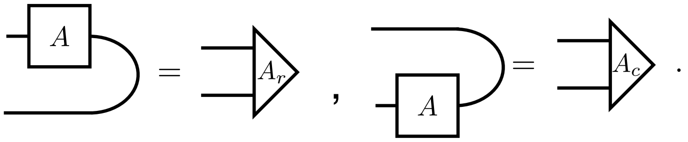

(2.14)

Or in mathematics, it means

$$|A\rangle\!\rangle_r = (A \otimes \mathbb{1})|\Phi^+\rangle, \quad |A\rangle\!\rangle_c = (\mathbb{1} \otimes A)|\Phi^+\rangle \tag{2.15}$$

Following the definition, we also have the following identity,

$$|\Phi^+\rangle = |\mathbb{1}\rangle\!\rangle_r = |\mathbb{1}\rangle\!\rangle_c \tag{2.16}$$

In general, we can the vectorization procedure with respect to an arbitrary basis for $L(\mathcal{X}, \mathcal{Y})$. Suppose $\mathcal{X} \cong \mathbb{C}^{d_x}$, $\mathcal{Y} \cong \mathbb{C}^{d_y}$, and $\mathcal{Z} \cong \mathbb{C}^{D}$, where $D = d_x d_y$. We choose an arbitrary orthonormal operator basis $\{\sigma_\alpha : \alpha = 0, \cdots, D-1\}$ and $\mathrm{Tr}(\sigma_\alpha^\dagger \sigma_\beta) = \delta_{\alpha,\beta}$. Then vectorization with respect to this basis for $L(\mathcal{X}, \mathcal{Y})$ is given by

$$\sigma - \mathrm{vec} : L(\mathcal{X}, \mathcal{Y}) \to \mathcal{Z} : A \to |A\rangle\!\rangle_\sigma \tag{2.17}$$

and

$$|A\rangle\!\rangle_\sigma = \sum_{\alpha=0}^{D-1} \mathrm{Tr}(\sigma_\alpha^\dagger A)|\sigma\rangle \tag{2.18}$$

So for row(col)-vector, the orthonormal basis are elementary matrix basis $\{E_{i,j} = |i\rangle|j\rangle\}$ and making assignment to $\alpha = di + j$ for row-vector, $\alpha = i + dj$ for col-vector. Given two choices of orthonormal basis $\{\sigma_\alpha\}$ and $\{\omega_\alpha\}$ for $L(\mathcal{X}, \mathcal{Y})$, we can define a basis transformation operator that transforms vectorized operators from σ-vec to $\omega-$vec:

$$T_{\sigma\to\omega}:\mathcal{Z}\to\mathcal{Z}:|A\rangle\rangle_{\sigma}\to|A\rangle\rangle_{\omega}. \tag{2.19}$$

The transformation is given by

$$T_{\sigma\to\omega}=\sum_{\alpha}|\sigma_{\alpha}\rangle\rangle_{\omega}\langle\alpha|=\sum_{\alpha}|\alpha\rangle\langle\langle\omega_{\alpha}|_{\sigma}. \tag{2.20}$$

We can check the two definition in Eq.(2.20) are equal and

$$\begin{aligned}
T_{\sigma\to\omega}|A\rangle\rangle_{\sigma}&=\sum_{\alpha}|\sigma_{\alpha}\rangle\rangle_{\omega}\langle\alpha|\sum_{\beta}\mathrm{Tr}(\sigma_{\beta}^{\dagger}A)|\beta\rangle\\
&=\sum_{\beta}\left(\sum_{\alpha}\mathrm{Tr}(\omega_{\beta}^{\dagger}\sigma_{\alpha})\,\mathrm{Tr}(\sigma_{\alpha}^{\dagger}A)\right)|\beta\rangle\\
&=\sum_{\beta}\mathrm{Tr}(\omega_{\beta}^{\dagger}A)|\beta\rangle\\
&=|A\rangle\rangle_{\omega}.
\end{aligned} \tag{2.21}$$

In the derivation, we have used $\sum_{\alpha}\mathrm{Tr}(\omega_{\beta}^{\dagger}\sigma_{\alpha})\,\mathrm{Tr}(\sigma_{\alpha}^{\dagger}A)=\mathrm{Tr}(\omega_{\beta}^{\dagger}A)$. Especially, the transformation between row-vector and col-vector is the SWAP transformation:

$$T_{c\to r}=T_{r\to c}=\mathrm{SWAP}. \tag{2.22}$$

We can prove this by the following diagrams:

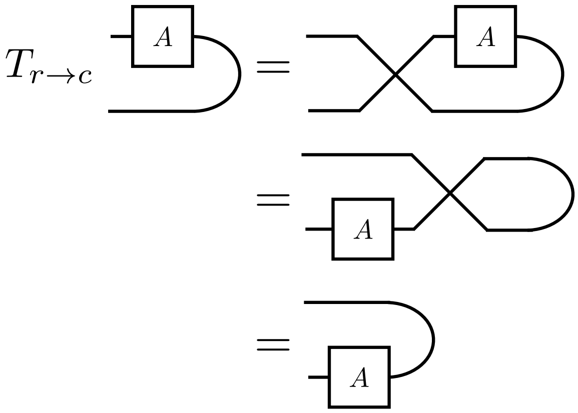

(2.23)

2.3 Open quantum systems

2.3.1 Density matrix: a complete representation of quantum states

In the previous section, I have introduced the five axioms of quantum theory for a *closed* quantum system. The states are a ray of vectors in the Hilbert space. And the quantum dynamics are unitary, meaning there is no loss of information. However, it is very rare to have a closed quantum system that does not interact with the environment. If that is the case, people usually call the system is open. For an open quantum system, the axioms do not hold in general. The states are not rays anymore, and the dynamics are not unitary. Let me first review how to characterize a quantum state in the open quantum system.

Suppose $\{|a\rangle_A\}$ and $\{|b\rangle_B\}$ are orthonomal basis for Hilbert space $\mathcal{H}_A$ and $\mathcal{H}_B$ respectively. Then a quantum state in the total system is

$$|\psi\rangle = \sum_{a,b} \psi_{a,b} |a\rangle_A \otimes |b\rangle_B. \tag{2.24}$$

An observable only acting on subsystem A will be $M_A \otimes I_B$. Suppose we want to get the expectation value of M_A, what we have is

$$\begin{aligned}\langle M_A\rangle &= \langle\psi|M_A\otimes I_B|\psi\rangle \\ &= \sum_{a',a,b}\langle a'|_A M_A|a\rangle_A \psi^*_{a',b}\psi_{a,b} \\ &= \mathrm{Tr}(\rho_A M_A),\end{aligned} \tag{2.25}$$

where

$$\rho_A = \mathrm{Tr}_B(|\psi\rangle\langle\psi|) = \sum_{a',a,b}\psi^*_{a',b}\psi_{a,b}|a\rangle_A\langle a'|_A. \tag{2.26}$$

We see if an observer only can prove the information in the subsystem A, then the subsystem is fully described by $\rho_A = \mathrm{Tr}_B(|\psi\rangle\langle\psi|)$, which is called (reduced) density matrix. The density matrix is the general description of a quantum system, and any properties of subsystem A can be predicted with ρ_A. From the definition of the density matrix in Eq.(2.26), we can infer the density matrix ρ has the following properties:

1. Self-adjoint: $\rho^\dagger = \rho$
2. Positivity: for any $|\varphi\rangle$, $\langle\varphi|\rho|\varphi\rangle = \sum_b|\sum_a\psi_{a,b}\langle\varphi|a\rangle|^2 \geq 0$,
3. $\mathrm{Tr}(\rho) = 1$.

If the state is pure, $\rho = |\psi\rangle\langle\psi|$, then the density matrix is a projector to the one-dimensional subspace expanded by $|\psi\rangle$. And $\mathrm{Tr}(\rho^2) = \mathrm{Tr}(\rho) = 1$. And if this is not the case, we say the state is *mixed*. Since ρ is self-adjoint, it can be diagonalized or has a spectral decomposition:

$$\rho = \sum_a p_a|a\rangle\langle a|, \tag{2.27}$$

where $0 < p_a \leq 1$ and $\sum_a p_a = 1$. We can interpret p_a as the *classical* probability of finding the system in a pure state $|a\rangle\langle a|$. This is also called the ensemble interpretation of the mixed state.

This interpretation also gives us a protocol for preparing ρ. Suppose Alice has a random number generator that can generate samples a according to probability p_a, and a machine to prepare any pure state in $\{|a\rangle\}$. Suppose Bob wants to measure the expectation value of M. Every time, Alice prepares a pure state $|a\rangle$ with probability p_a and hands it to Bob for the measurement. The expectation for M will be

$$\langle M\rangle = \sum_a p_a \langle a|M|a\rangle = \mathrm{Tr}(\rho M) \tag{2.28}$$

So Bob cannot distinguish this probabilistic preparation procedure with another direct preparation of ρ if it exists.

2.3.2 Ensemble interpretation is not unique

A fascinating consequence of the ensemble interpretation is that the ensemble is not unique. According to Eq.(2.27), one may think the ensemble interpretation is only the probabilistic combination of $|a\rangle$ determined by spectral decomposition. However, this is not true. The tricky part is that any convex combination of the density matrix is a valid density matrix. Suppose ρ_1 and ρ_2 are density matrices, then $\rho(\lambda) = \lambda\rho_1 + (1-\lambda)\rho_2$ with $0 \le \lambda \le 1$ is also a valid density matrix. It is straightforward to check it is self-adjoint, non-negative, and normalized. Therefore, all the possible density matrices form a convex set as denoted in Fig.2.2.

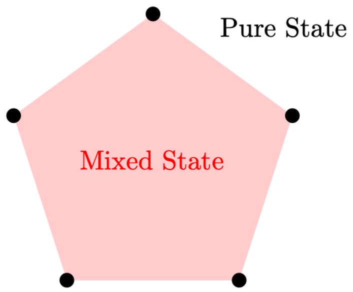

Figure 2.2: The convex set of density matrix. The extremal points of the set are pure states. And all the other points are mixed states.

Pure states are special cases. It *cannot* be expressed as a convex combination of two other states. In order to prove this statement, suppose we have a pure state $\rho = |\psi\rangle\langle\psi|$ and a orthogonal state $|\psi_\perp\rangle$ such that $\langle\psi|\psi_\perp\rangle = 0$. If the pure state can be expressed as a convex combination of two other states, then

$$\rho = \lambda\rho_1 + (1-\lambda)\rho_2,\ \lambda \neq 0. \tag{2.29}$$

$$0 = \langle\psi_\perp|\rho|\psi_\perp\rangle = \lambda\langle\psi_\perp|\rho_1|\psi_\perp\rangle + (1-\lambda)\langle\psi_\perp|\rho_2|\psi_\perp\rangle \tag{2.30}$$

Therefore, we conclude $\langle\psi_\perp|\rho_1|\psi_\perp\rangle = 0$ and $\langle\psi_\perp|\rho_2|\psi_\perp\rangle = 0$. Since $|\psi_\perp\rangle$ can be any orthogonal states, $\rho = \rho_1 = \rho_2$. The pure states are unique. They are the extremal points in the convex set, which are denoted as black points in Fig.2.2.

Now we can ask whether the ensemble interpretation of a density matrix is unique. Mathematically, the question is how many ways can we write ρ as a convex combination of pure states. In fact, the answer is there are *infinitely* many ways to write a mixed state as a convex combination of pure states. Without proving this statement, let's state spin-$\frac{1}{2}$ system as an example. Any quantum state can be written as

$$\rho(\vec{P}) = \frac{1}{2}(I + \vec{P}\vec{\sigma}). \tag{2.31}$$

If $|\vec{P}| = 1$, the state is pure; and if $|\vec{P}| < 1$, then the state is mixed. This is also called the Bloch sphere representation of a qubit. Suppose $\vec{n}_1$ and $\vec{n}_2$ are two vectors on the Bloch shpere, then $\rho(\vec{n}_1)$ and $\rho(\vec{n}_2)$ are pure states. If $\vec{P} = \lambda\vec{n}_1 + (1-\lambda)\vec{n}_2$, then we have $\rho(\vec{P}) = \lambda\rho(\vec{n}_1) + (1-\lambda)\rho(\vec{n}_2)$. As a consequence, if point $\vec{P}$ lies on the line segment connecting $\vec{n}_1$ and $\vec{n}_2$, then it is a valid ensemble interpretation. We immediately see there are infinitely many plausible interpretations.

Therefore, one should always be careful to interpret p_a in $\rho = \sum_a p_a|a\rangle\langle a|$ as a classical probability. And we should remember there are infinitely many such ensemble interpretations for the same mixed state. The ambiguity of the ensemble interpretation is quite fascinating since we will never see this fact in the classical probabilistic system. In the classical world, if an object can be interpreted as 20% A and 80% B, there is no other ensemble interpretation. One may think the classical world only has one probabilistic interpretation, while the quantum world can be viewed from many aspects. Each aspect has its ensemble interpretation, and they are equally right.

2.3.3 Generalized measurements

In the previous section, we have seen that the complete representation of a quantum state is the density matrix. And in this section, we will see that the general measurements are not orthogonal projections. Here, I will follow a simple argument or model of generalized measurement envisioned by von Neumann. And more detailed discussion can be found in Ref.[NC10].

To see how generalized measurements are not orthogonal, let's first take a look at a qubit system A, that is described by $\alpha|0\rangle_A + \beta|1\rangle_A$. The generalized measurements can be achieved with an ancilla system B, which can be viewed as a model of the instrument. As a simple example, we also assume the instrument is modeled by a two-level quantum system, with orthonormal basis $\{|0\rangle_B, |1\rangle_B\}$. Without loss of generality, we assume the instrument is initially in the state $|0\rangle_B$. In order to probe the system A, the instrument couples ancilla system B with A, and evolve the total system with a simple unitary dynamics:

$$U : (\alpha|0\rangle_A + \beta|1\rangle_A) \otimes |0\rangle_B \to \alpha|0\rangle_A \otimes |0\rangle_B + \beta|1\rangle_A \otimes |1\rangle_B. \tag{2.32}$$

And the instrument measure the ancilla qubit to give the readout. If the instrument measures spin-z of B, then it equivalently measures the spin-z information of A. And it is still an orthogonal measurement. However, this is not always the case. Suppose the instrument measures $|+\rangle_B$ and $|-\rangle_B$. If the instrument has readout ±1, then the state of A changes to $\alpha|0\rangle_A \pm \beta|1\rangle_A$ after the measurement. And the post-measurement states are not orthogonal in general.

This idea can be generalized to ancilla system has N orthonormal states:

$$\{|0\rangle_B, \cdots, |N-1\rangle_B\}.$$

And the instrument couples the system A with ancilla B by a unitary dynamics:

$$U : |\psi\rangle_A \otimes |0\rangle_B \to \sum_a M_a |\psi\rangle_A \otimes |a\rangle_B, \tag{2.33}$$

where M_a is usually called *Kraus operators*. Since unitary dynamics preserves the norm of vector, we have a constraints on the Kraus operators:

$$1 = \sum_{a,b} \langle\psi|_A M_a^\dagger M_b |\psi\rangle_A \langle a|b\rangle_B = \langle\psi|_A \sum_a M_a^\dagger M_a |\psi\rangle_A. \tag{2.34}$$

Since $|\psi\rangle_A$ is arbitrary, we require $\sum_a M_a^\dagger M_a = I$. An orthogonal measurements on ancilla system is described by $I \otimes |a\rangle_B \langle a|_B$. Therefore, the probability of getting readout a is

$$p(a) = \langle\psi|_A M_a^\dagger M_a |\psi\rangle_A = ||M_a|\psi\rangle_A||^2. \tag{2.35}$$

And after the measurement, the state changes to

$$|\psi\rangle_A \to \frac{M_a|\psi\rangle_A}{\sqrt{\langle\psi|_A M_a^\dagger M_a |\psi\rangle_A}}. \tag{2.36}$$

One should notice that if $M_a M_b \neq \delta_{a,b} M_a$, then two consecutive measurements can give different readout, which is in contrast to orthogonal measurements. In the literature, people usually define $F_a \equiv M_a^\dagger M_a$ as the generalized measurement operator, and dubbed it the positive operator-valued measure (POVM). By definition, it satisfies the following properties:

1. Self-adjoint: $F_a^\dagger = F_a$,
2. Positivity: for any state $|\psi\rangle$, $\langle\psi|F_a|\psi\rangle \geq 0$,
3. Completeness: $\sum_a F_a = I$.

So the POVM operators F_a give us the probability of getting the readout a but do not provide information on the post-measurement state. Since any $M_a = U_a\sqrt{F_a}$ with unitary operator U_a can be the corresponding Kraus operator. And we can design any Kraus operator we want to evolve the total system. By doing projective measurements on the ancilla system, the corresponding POVM general measurement can be achieved.

2.3.4 Quantum Channels

For open quantum system, the dynamics is not unitary in general. In Eq.(2.33), we have introduced Kraus operators $\{M_a\}$ and seen the formal general evolution of the total system:

$$U : |\psi\rangle_A \otimes |0\rangle_B \to \sum_a M_a|\psi\rangle_A \otimes |a\rangle_B, \tag{2.37}$$

If we don't have any knowledge about environment B, then for open quantum system A, the general evolution is

$$\mathcal{E} : \rho_A \to \sum_a M_a \rho_A M_a^\dagger, \tag{2.38}$$

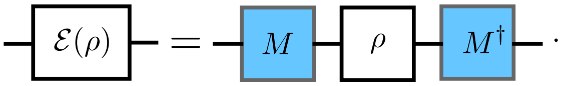

(2.39)

Such a linear map $\mathcal{E}(\rho)$ is also call *quantum channel*. And the Kraus operators must satisfy the completeness condition $\sum_a M_a^\dagger M_a = I$. Similar to the ambiguity in ensemble interpretation for density matrix, the Kraus operator representation for a quantum channel $\mathcal{E}$ is also not unique.

The root of this phenomenon is also from the fact that there are more than one way to look at the quantum system. Suppose we perform a basis transformation on ancilla system B with unitary V: $|a\rangle_B = \sum_i |i\rangle_B V_{i,a}$. Therefore Eq.(2.33) can be rewritten as

$$U : |\psi\rangle_A \otimes |0\rangle_B \to \sum_{a,i} M_a |\psi\rangle_A \otimes V_{i,a}|i\rangle_B = \sum_i (\sum_a V_{i,a} M_a)|\psi\rangle_A \otimes |i\rangle_B. \tag{2.40}$$

If we do a partial trace on ancilla system with the new basis $\{|i\rangle_B\}$, the same quantum channel is represented as

$$\mathcal{E} : \rho_A \to \sum_a \tilde{M}_a \rho_A \tilde{M}_a^\dagger, \tag{2.41}$$

where the new Kraus operators are $\tilde{M}_a = \sum_i V_{a,i} M_i$. In mathematics, people also call the map of the quantum channel the *Complete Positive Trace Preserving* map or *CPTP* map for short. Using the Kraus representation of quantum channels, we can easily verify a quantum channel is: 1) positive and 2) trace-preserving. Physically, it guarantees the quantum channel maps quantum states to states. Complete positivity is a stronger requirement. If a map $\mathcal{E}$ acting on $\mathcal{H}_A$ is positive, and for any composite system $\mathcal{H}_A \otimes \mathcal{H}_B$, the composite map $\mathcal{E} \otimes \mathbb{1}$ is also positive, then we say $\mathcal{E}$ is completely positive. With Kraus representation, we can easily verify $M \otimes \mathbb{1}$ is completely positive. The physical intuition behind these strong constraints is that even the system is part of the big world, the quantum channel should still be a valid dynamic.

2.4 Other representations of quantum channel

In the previous section, we have seen that the quantum channel is a CPTP map and the Kraus representation of a quantum channel. And there are other representations of quantum channels, which can be useful. In this section, without proof, I will briefly review some other representations: 1) *Louiville representation*, 2) *Choi matrix representation*, and 3) *process matrix* *representation*. A detailed discussion on different representations can be found in [WBC15].

2.4.1 Louiville representation

The Louiville representation is based on the vectorization of density matrix $\rho \to |\rho\rangle\rangle_\sigma$ with respect to some orthonormal basis $\{\sigma_\alpha, \alpha = 0, \cdots, d^2-1\}$. Then the quantum channel can be defined as a superoperator S, which maps operators to operators:

$$S : \mathcal{X} \otimes \mathcal{X} \to \mathcal{Y} \otimes \mathcal{Y} : |\rho\rangle\rangle_\sigma \to |\mathcal{E}(\rho)\rangle\rangle_\sigma. \tag{2.42}$$

In the following, I will choose col-vector ordering and drop the subscript. Otherwise, I will mention the ordering of vectorization. The pictorial description of the quantum channel can be viewed as:

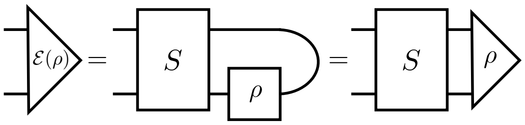

(2.43)

and, for col-vector, the mathematical description is

$$\mathcal{E}(\rho)_{mn} = \sum_{\mu,\nu} S_{mn,\nu\mu}\rho_{\mu\nu}. \tag{2.44}$$

In the previous section, we have introduced the basis transformation operator:

$$T_{\sigma\to\omega} : |A\rangle\rangle_\sigma \to |A\rangle\rangle_\omega, \tag{2.45}$$

with

$$T_{\sigma\to\omega} = \sum_\alpha |\alpha\rangle\langle\langle\omega_\alpha|_\sigma = \sum_\alpha |\sigma_\alpha\rangle\rangle_\omega\langle\alpha|. \tag{2.46}$$

The graphical representation of the equation is below:

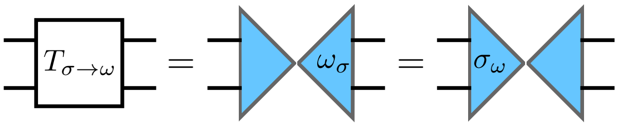

(2.47)

Since most of the time we use col-vector convention, I also give the tensor diagram of transformation $T_{c\to\sigma}$ in Fig. 2.3.

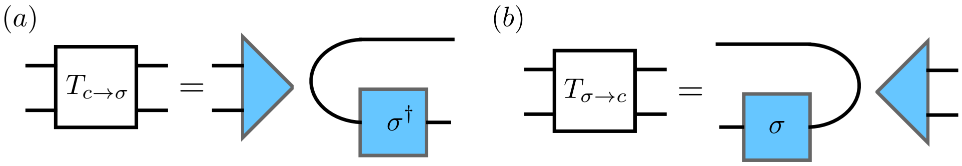

Figure 2.3: The tensor diagram of basis transformation between col-order and arbitrary σ-order.

Then we can transform superoperator S to other basis by

$$S_\sigma = T_{c\to\sigma} S T^\dagger_{c\to\sigma} \tag{2.48}$$

2.4.2 Choi matrix representation

Another representation of quantum channel is called *Choi matrix*[Cho75]. It is an application of the Choi-Jamiolkowski isomorphism, which bijectively map linear maps to linear operators[Jam72]. For $\mathcal{X} \cong \mathbb{C}^d$, there are two conventions of the definition of Choi matrix:

$$\begin{aligned}
\Lambda_c &= \sum_{i,j=0}^{d-1} |i\rangle\langle j| \otimes \mathcal{E}(|i\rangle\langle j|) = (I\otimes\mathcal{E})|\Phi^+\rangle\langle\Phi^+|,\\
\Lambda_r &= \sum_{i,j=0}^{d-1} \mathcal{E}(|i\rangle\langle j|) \otimes |i\rangle\langle j| = (\mathcal{E}\otimes I)|\Phi^+\rangle\langle\Phi^+|,
\end{aligned} \tag{2.49}$$

where $\{|i\rangle, i=0,\cdots,d\}$ is the orthonormal basis for $\mathcal{X}$. In the following, I will use the column-vector convention and drop the subscript. The diagrammatic representation of Λ_c with Kraus operators is

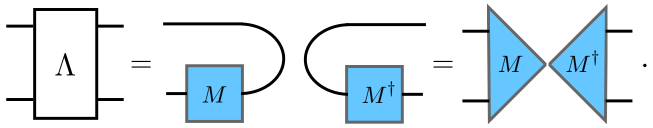

(2.50)

The evolution of the density matrix in terms of Choi-matrix is given by

$$\mathcal{E}(\rho) = \mathrm{Tr}_{\mathcal{X}}[(\rho^T \otimes \mathbb{1}_{\mathcal{Y}})\Lambda], \tag{2.51}$$

which can be represented as

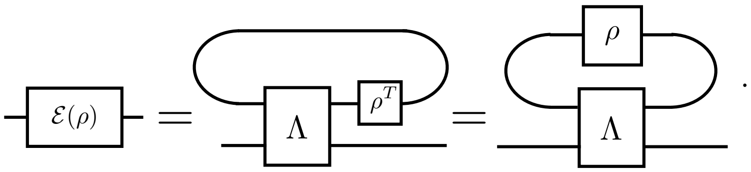

(2.52)

A diagrammatic proof of the evolution equation Eq.(2.51) is:

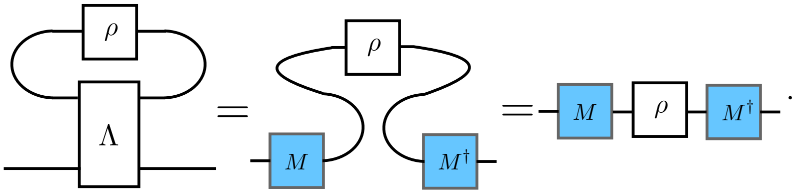

(2.53)

Unlike the Kraus representation, once we fixed the orthonormal basis $\{|i\rangle_{\mathcal{X}}, i = 0, \cdots d_x - 1\}$ for the system, the Choi representation of a quantum channel is fixed. In Kraus representation, we have seen the ambiguity even for the system with a fixed basis. This ambiguity is rooted in the ambiguous choice for the environmental basis.

2.4.3 Process matrix representation

Another representation for the quantum channel is called *process matrix representation*. It is related to the Choi representation through the change of basis for $L(\mathcal{X},\mathcal{X})$. Suppose the Hilbert space $\mathcal{X} \cong \mathbb{C}^d$, and let $D = d^2$. If one choose a basis $\{\sigma_\alpha : \alpha = 0,\ldots,D-1\}$, then the quantum channel can be represented as

$$\mathcal{E}(\rho) = \sum_{\alpha,\beta=0}^{D-1} \chi_{\alpha,\beta}\sigma_\alpha\rho\sigma_\beta^\dagger, \tag{2.54}$$

where $\chi_{\alpha,\beta}$ is the process matrix with respect to orthonormal basis $\{\sigma_\alpha : \alpha = 0,\ldots,D-1\}$. And the Choi matrix is related to general process matrix by a basis transformation. In other word, the Choi matrix can be viewed as the process matrix in the column order basis where $\{\sigma_\alpha = |i\rangle\langle j| : \alpha = i + jd\}$. Mathematically, it is

$$\begin{aligned} \chi &= T_{c\to\sigma} \cdot \Lambda \cdot T_{c\to\sigma}^\dagger \\ \Lambda &= \sum_{\alpha,\beta} \chi_{\alpha,\beta} |\sigma_\alpha\rangle\!\rangle_c \langle\!\langle\sigma_\beta|_c. \end{aligned} \tag{2.55}$$

This relation can be simply proved by the diagrammatic approach:

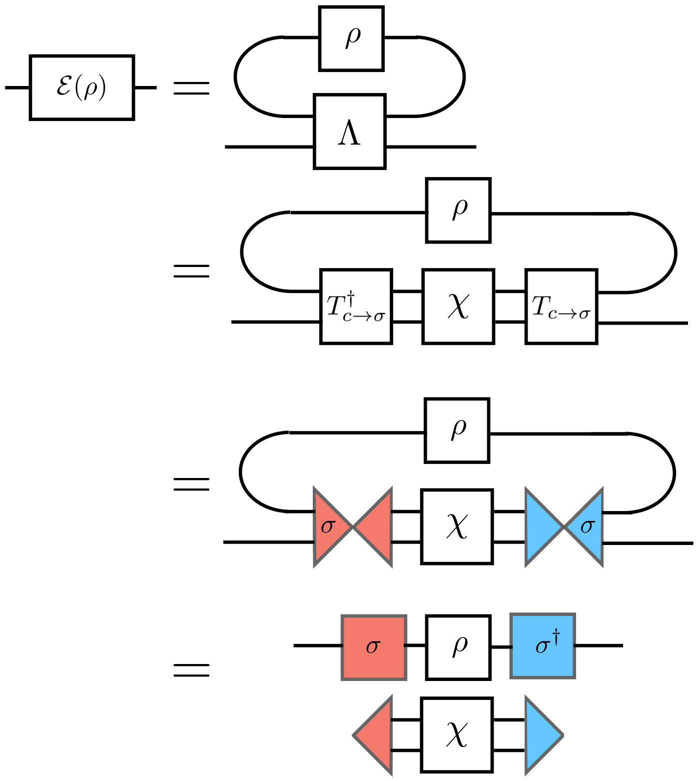

(2.56)

2.4.4 Relation among representations

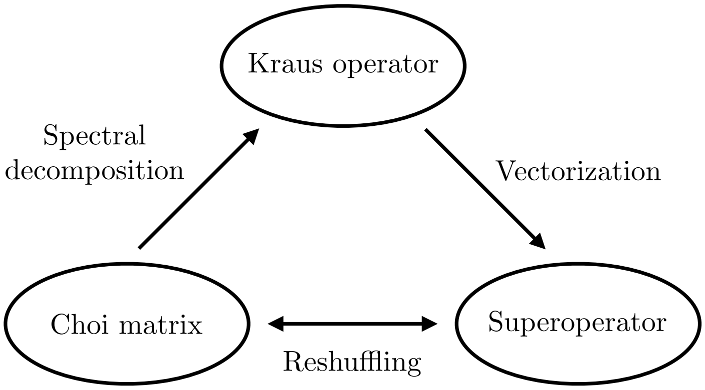

Figure 2.4: Relations between different representations of a quantum channel.

In the previous sections, I have introduced the Kraus representation, Choi matrix ($\chi-$matrix) representation, and superoperator representation of a quantum channel. The relation among them is summarized in Fig. 2.4. Here, I present the transformation between different representations and the proof using tensor diagrams. This section is largely inspired by [WBC15].

Relations between Choi matrix and superoperator: reshuffling

The Choi matrix Λ and superoperator S can be transformed back and forth with a reshuffling of the indexes:

$$\Lambda_{mn,\mu\nu} = S_{\nu n,\mu m},\ S_{mn,\mu\nu} = \Lambda_{\nu n,\mu m}. \tag{2.57}$$

or diagrammatically it can be represented as:

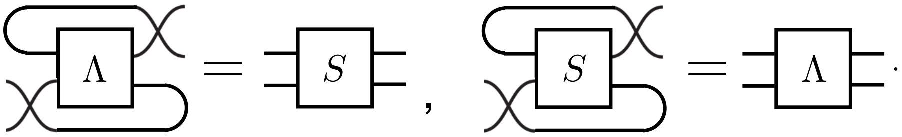

(2.58)

I will prove the first equation with tensor diagrams, and the second equality can be proved in a similar way:

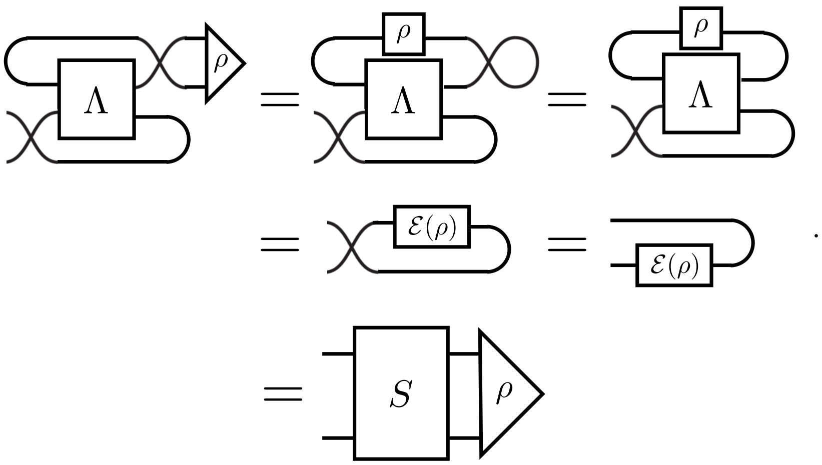

(2.59)

Kraus operator to superoperator: vectorization

Once the Kraus representation of a given quantum channel is given by $\{M_\alpha\}$, the superoperator representation is constructed via vectorization:

$$S = \sum_{\alpha} M_{\alpha}^{*} \otimes M_{\alpha}, \tag{2.60}$$

where M_{α}^{*} is the hermitian conjugate.

Choi matrix to Kraus operators: spectral decomposition

Since the Choi matrix describes a CPTP-map, it is hermitian and positive. So it always has a spectral decomposition form:

$$\Lambda = \sum_{\alpha} \lambda_{\alpha} |V_{\alpha}\rangle\langle V_{\alpha}|, \tag{2.61}$$

with $\lambda_{\alpha} > 0$. $|V_{\alpha}\rangle$ is the eigenvector and there is a unique operator K_{α} satisfies $|K_{\alpha}\rangle\rangle_{c} = |V_{\alpha}\rangle$. K_{α} can be find by reordering elements of $|V_{\alpha}\rangle$. Then the Kraus operator can be defined as

$$M_{\alpha} = \sqrt{\lambda_{\alpha}} K_{\alpha} \tag{2.62}$$

2.5 Hidden variable models

As discussed, quantum mechanics is a mathematical model that describes the physical world based on axioms. And certainly, the axiom about measurement dynamics and quantum state dynamics seems not very satisfying. Einstein suspected that quantum mechanics is an incomplete description of the physical world [EPR35]. And some of the uncertainty in the measurements are probably rooted in some unknown hidden degree of freedom. Since people never "observe" a quantum state directly, we never witness magic properties, like entanglement, predicted by quantum mechanics. People can only infer what is in the black box through the classical measurement data. All the gems are hidden in the statistics of the readouts. Having introduced quantum mechanics, now let me introduce what is not quantum mechanics. Especially, I will focus on the (local) hidden variable model.

The hidden variable model is a classical stochastic model. First, it assumes that every quantum state or quantum black box labelled by $|\psi\rangle$ is described by a probability distribution of hidden variable λ: $P_{|\psi\rangle}(\lambda)$. For every physical observable O, the probability of getting readout y_i is given by a conditional probability: $P(y_i|\lambda, O)$. So the total probability of having readout y_i from observing O is

$$P(y_i|O) = \sum_{\lambda} P(y_i|\lambda, O) P_{|\psi\rangle}(\lambda). \tag{2.63}$$

Suppose we have a qubit system $|\psi\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$. From quantum mechanics, we know if one measures spin-z operator σ_z, then half of the chance we get $+1$ and the other half we get -1. And if one measures spin-x operator σ_x, then we get $+1$ all the time. This phenomenon can also be described by a hidden variable model. Suppose there is a hidden variable that we don't know, and it can take four values $0_z0_x, 0_z1_x, 1_z0_x, 1_z1_x$. The measurement probability is defined as

$$\begin{aligned}
&P(+1|0_z0_x, \sigma_z) = 1,\ P(+1|0_z0_x, \sigma_x) = 1 \\
&P(+1|0_z1_x, \sigma_z) = 1,\ P(-1|0_z1_x, \sigma_x) = 1 \\
&P(-1|1_z0_x, \sigma_z) = 1,\ P(+1|1_z0_x, \sigma_x) = 1 \\
&P(-1|1_z1_x, \sigma_z) = 1,\ P(-1|1_z1_x, \sigma_x) = 1
\end{aligned} \tag{2.64}$$

Then, the above measurement scheme can be reproduced with a hidden variable system that $P_{|\psi\rangle}(0_z0_x) = 0.5$, and $P_{|\psi\rangle}(1_z0_x) = 0.5$. There is nothing that goes beyond classical hidden variable theory about one qubit. So what properties are purely quantum mechanical? In the following, I will introduce two aspects of quantum mechanical systems that cannot be described by any local hidden variable theory: quantum non-locality and quantum contextuality.

2.5.1 Bell's inequality

When we have a composite system, magical things can happen. One of the most exciting properties in quantum mechanics is called *quantum entanglement*. Quantum entanglement is a quantum version of correlation that goes beyond classical correlation. Suppose we have a pure state $|\psi\rangle$ for a bipartite system $A \otimes B$. With Schmidt decomposition, any pure bipartite state has the Schmidt form:

$$|\psi\rangle = \sum_{\alpha} \lambda_{\alpha} |\alpha\rangle_A \otimes |\alpha\rangle_B, \tag{2.65}$$

where $\{|\alpha\rangle_A\}$ and $\{|\alpha\rangle_B\}$ are orthonormal basis for $\mathcal{H}_A$ and $\mathcal{H}_B$ respectively. Then the reduced density matrices are

$$\begin{aligned} \rho_A &= \mathrm{Tr}_B(\rho) = \sum_{\alpha} \lambda_{\alpha}^2 |\alpha\rangle_A \langle\alpha|_A, \\ \rho_B &= \mathrm{Tr}_A(\rho) = \sum_{\alpha} \lambda_{\alpha}^2 |\alpha\rangle_B \langle\alpha|_B. \end{aligned} \tag{2.66}$$

There are many definitions of entanglement entropy, such as von Neumann entropy and Renyi entropy. Here, I will introduce the von Neumann entropy: $S = -\mathrm{Tr}(\rho \log(\rho))$. As we can see the entanglement between A and B depends on the spectral of Schmidt decomposition, and $S = -\sum_{\alpha} \lambda_{\alpha}^2 \log(\lambda_{\alpha}^2)$. Another important property of entanglement entropy is that it does not rely on the choice of basis for subsystems.

Now let me generalize the previous example to a two-qubit state: $|\psi\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$, which is also called Einstein–Podolsky–Rosen (EPR) state. One can check that the entanglement entropy for part A and part B are one bit. Therefore, part A and part B are entangled and not separable. Moreover, if one calculates the mutual information, we see the EPR state can share two bits of mutual information between A and B. This phenomenon is quite interesting because the perfectly correlated two classical bits cannot share mutual information that exceeds one bit. This indicates the EPR state has some correlation beyond classical correlations. In the previous discussion, we have seen that the classical hidden variable model can perfectly fit the data produced by one qubit. Now let me prove this will not be true for two qubits EPR state.

Suppose we have a EPR state $|\mathrm{EPR}\rangle$ that described by a hidden variable theory $P_{\mathrm{EPR}}(\lambda)$. After state preparation, we send one qubit to Alice, and the other qubit to Bob. We assume there is no superluminal communication, so Alice and Bob have the same hidden variable distributions, $P_{\mathrm{EPR}}(\lambda)$. This is also called local hidden variable theory. And we assume Alice and Bob have a instrument that can measure at least three quantities: $\vec{n}_1 \cdot \vec{\sigma}$, $\vec{n}_2 \cdot \vec{\sigma}$, and $\vec{n}_3 \cdot \vec{\sigma}$. Suppose there is a hidden variable theory with hidden variable λ that can take value $\lambda = \pm 1_{n_1} \pm 1_{n_2} \pm 1_{n_3}$. And if one measures $\vec{n}_i \cdot \vec{\sigma}$, the probability of readout is $p(\pm 1| \cdots \pm 1_{n_i} \cdots, \vec{n}_i \cdot \vec{\sigma}) = 1$. For example $p(+1|+1_{n_1} - 1_{n_2} + 1_{n_3}, \vec{n}_1 \cdot \vec{\sigma}) = 1$, $p(-1|+1_{n_1} - 1_{n_2} + 1_{n_3}, \vec{n}_2 \cdot \vec{\sigma}) = 1$, and $p(+1|+1_{n_1} - 1_{n_2} + 1_{n_3}, \vec{n}_3 \cdot \vec{\sigma}) = 1$. If such a hidden variable theory exists, there must be a preassigned probability distribution $P_{\mathrm{EPR}}(\lambda)$, which is illustrated in Table.2.1.

Table 2.1: The hidden variable theory for EPR state. If such a theory exists, then there is a preassigned probability distribution.

Alice	Bob	$P_{\mathrm{EPR}}(\lambda)$
$+1_{n_1} + 1_{n_2} + 1_{n_3}$	$+1_{n_1} + 1_{n_2} + 1_{n_3}$	p_1
$+1_{n_1} + 1_{n_2} - 1_{n_3}$	$+1_{n_1} + 1_{n_2} - 1_{n_3}$	p_2
$+1_{n_1} - 1_{n_2} + 1_{n_3}$	$+1_{n_1} - 1_{n_2} + 1_{n_3}$	p_3
$+1_{n_1} - 1_{n_2} - 1_{n_3}$	$+1_{n_1} - 1_{n_2} - 1_{n_3}$	p_4
$-1_{n_1} + 1_{n_2} + 1_{n_3}$	$-1_{n_1} + 1_{n_2} + 1_{n_3}$	p_5
$-1_{n_1} + 1_{n_2} - 1_{n_3}$	$-1_{n_1} + 1_{n_2} - 1_{n_3}$	p_6
$-1_{n_1} - 1_{n_2} + 1_{n_3}$	$-1_{n_1} - 1_{n_2} + 1_{n_3}$	p_7
$-1_{n_1} - 1_{n_2} - 1_{n_3}$	$-1_{n_1} - 1_{n_2} - 1_{n_3}$	p_8

Suppose Alice measures $\vec{n}_1$ and Bob measures $\vec{n}_2$, the probability they have the same readout is $p(\vec{n}_1 \cdot \vec{\sigma}_A = \vec{n}_2 \cdot \vec{\sigma}_B) = p_1 + p_2 + p_7 + p_8$. With the same logic, we can write down

$$p(\vec{n}_1 \cdot \vec{\sigma}_A = \vec{n}_2 \cdot \vec{\sigma}_B) + p(\vec{n}_2 \cdot \vec{\sigma}_A = \vec{n}_3 \cdot \vec{\sigma}_B) + p(\vec{n}_3 \cdot \vec{\sigma}_A = \vec{n}_1 \cdot \vec{\sigma}_B) = 1 + 2p_1 + 2p_8 \geq 1. \quad (2.67)$$

If we use the quantum mechanical model to calculate the same probability, we will get:

$$\begin{aligned}
&p(\vec{n}_1\cdot\vec{\sigma}_A=\vec{n}_2\cdot\vec{\sigma}_B)+p(\vec{n}_2\cdot\vec{\sigma}_A=\vec{n}_3\cdot\vec{\sigma}_B)+p(\vec{n}_3\cdot\vec{\sigma}_A=\vec{n}_1\cdot\vec{\sigma}_B)\\
&=\frac{1+\vec{n}_1\cdot\vec{n}_2}{2}+\frac{1+\vec{n}_2\cdot\vec{n}_3}{2}+\frac{1+\vec{n}_3\cdot\vec{n}_1}{2}.
\end{aligned} \tag{2.68}$$

If the classical hidden variable model is compatible with quantum mechanical model, and be able to explain everything, then we will have

$$\frac{1+\vec{n}_1\cdot\vec{n}_2}{2}+\frac{1+\vec{n}_2\cdot\vec{n}_3}{2}+\frac{1+\vec{n}_3\cdot\vec{n}_1}{2}\geq 1. \tag{2.69}$$

This is a simple version of the Bell's inequality[Bel64b]. Especially, if we choose $\vec{n}_1$,$\vec{n}_2$ and $\vec{n}_3$ to be vectors that lie in the same plane and has angle $2\pi/3$ with each other, then $\vec{n}_1\cdot\vec{n}_2=\vec{n}_2\cdot\vec{n}_3=\vec{n}_3\cdot\vec{n}_1=-1/2$. And we see the Bell's inequality is violated:

$$\frac{1+\vec{n}_1\cdot\vec{n}_2}{2}+\frac{1+\vec{n}_2\cdot\vec{n}_3}{2}+\frac{1+\vec{n}_3\cdot\vec{n}_1}{2}=\frac{3}{4}\leq 1. \tag{2.70}$$

If we can test the violation of Bell's inequality in experiments, then it will confirm quantum mechanics cannot be described by any classical hidden variable model, and quantum entanglement is beyond classical correlations. Currently, there are several experiments that confirm the violation of Bell's inequality[ADR82, GMS18, CLZ$^+$18].

2.5.2 Quantum contextuality

As we have seen, the Bell's inequality rules out the possibility of classical local hidden variable theory. It shows us quantum entanglement and non-locality is a pure quantum effect. We should also notice that the Bell's inequality is a statement that depends on the quantum state. Besides quantum non-locality, there is another interesting quantum property called quantum contextuality that is independent of the underlying quantum state. Quantum contextuality focuses on the non-commuting nature of the measurements in quantum mechanics, and shows this non-commuting property or quantum uncertainty principle is intrinsic in quantum mechanics. Formally, the Kochen–Specker theorem[SK68] is a ‘no go’ result for noncontextual hidden variable theories. The theorem states that the observables of a quantum system can’t have definite, non-contextual, pre-existing values before they are measured, not even for certain finite sets of observables. Here, I will use a game, called Mermin magic square[Mer90, Mer94], to prove a simpler version of Kochen-Specker theorem.

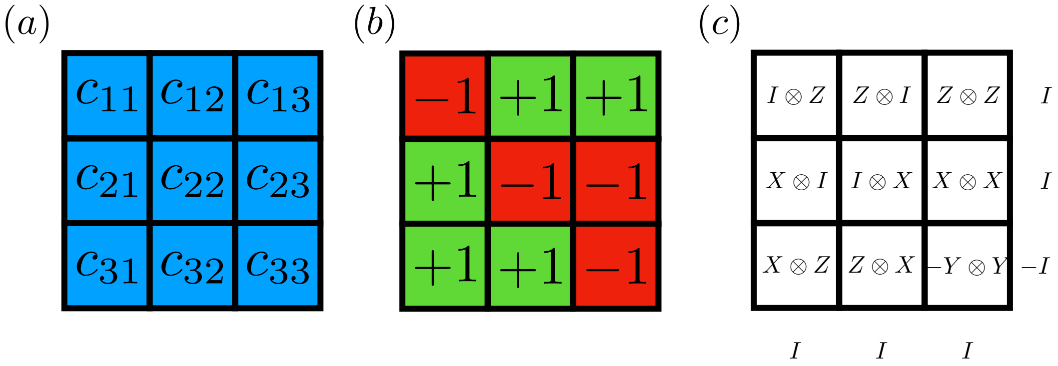

Figure 2.5: (a) A square with 9 entries. Each entry can be assigned with number ±1 or a probability $p(\pm 1)$ in general. (b) An example of the assignment. (c) A quantum magic square. Each entry is assigned with an operator. The value of entry will be obtained after measurements on some quantum state.

Suppose we have a 3×3 square as shown in Fig. 2.5 (a). Each of the entry c_{ij} can be assigned with ±1. Fig. 2.5 (b) is an example of one assignment. Let us consider the assignment of numbers are generated from a quantum mechanical process. Each element of the square is assigned with a two qubit Pauli operator, whose eigenvalues are ±1. For example, we assign $\sigma_x \otimes \sigma_z$ or $X\otimes Z$ to c_{31}. We can check the operators in each row and column commute with each other. Remember the commuting operators can be measured simultaneously. Therefore, we call each row and column a *context*, a concept from natural languages. Regardless of the underlying quantum state, we ask whether can we pre-assign probability $p(\pm 1)$ to each operator. If such a goal can be achieved, then we say there exist a non-contextual hidden variable theory.

Otherwise, there is no pre-existing probability that can be assigned to the readout of those quantum observables. And the value of each operator depends on the context that we are measuring.

Like the Bell's inequality, we define a statistical quantity,

$$\begin{aligned} C = \langle c_{11}c_{12}c_{13}\rangle + \langle c_{21}c_{22}c_{23}\rangle - \langle c_{31}c_{32}c_{33}\rangle \\ + \langle c_{11}c_{21}c_{31}\rangle + \langle c_{12}c_{22}c_{32}\rangle + \langle c_{13}c_{23}c_{33}\rangle, \end{aligned} \tag{2.71}$$

where $\langle\cdot\rangle$ means the statistical average. It is easy to check that for quantum model, $C = 1$. For non-contextual hidden variable model, it is also easy to check that by assigning ± 1 to c_{ij}, the largest C value one can get is $C = 4$. Therefore, any non-contextual hidden variable model will have $C \leq 4$. The experimental verification of C value indicates the non-communing nature of physical observables indicates the contextual property of quantum mechanics[GMS18, DRLB20]. In recent years, people have shown that quantum contextuality can be a potential resources for quantum advantage in various computational tasks[GAW^{+}21]. It is interesting to see art and science are often intertwined. Last but not the least, I would like to use a famous art piece by Escher to illustrate the contextuality. Just as shown in Fig. 2.6, local agreement will not guarantee the global agreement.

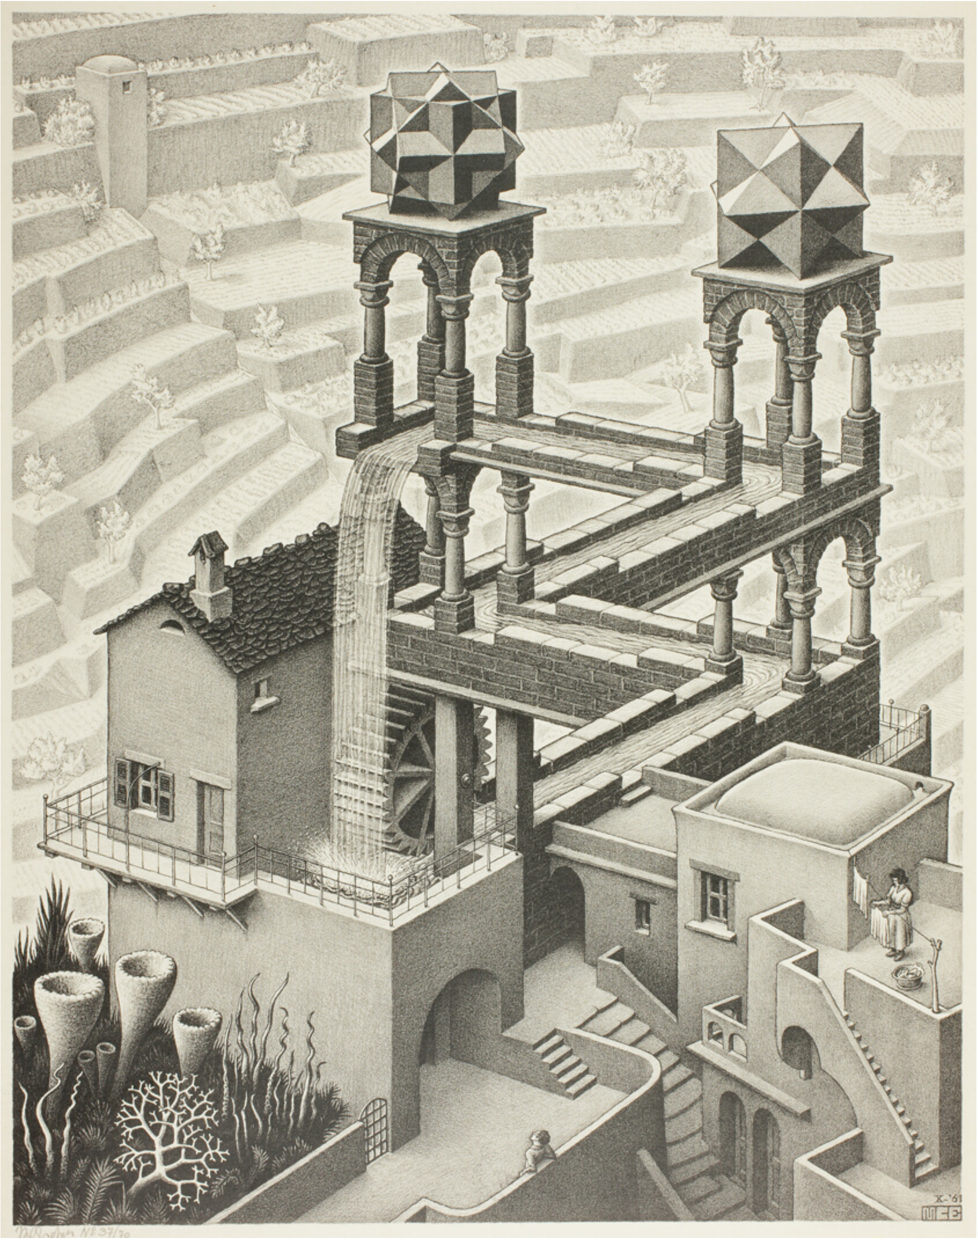

Figure 2.6: Escher's *Waterfall*. It shows local agreement will not guarantee the global agreement. It can be viewed as an art illustration of quantum contextuality.

Chapter 3

Complexity and unitary designs

3.1 Measure of Haar

In mathematics, a measure can be thought of as a ruler that tells you how things are distributed or concentrated in a mathematical set or space. For example, if we want to calculate the volume of a sphere, we can choose to use Cartesian coordinate (x,y,z) and do integral over $dx \cdot dy \cdot dz$. Or we can choose the spherical coordinate (r,θ,ϕ) and do integral over $r^2\sin(\theta)dr \cdot d\theta \cdot d\phi$. The extra $r^2\sin(\theta)$ can be viewed as the measure. Here, I am not going to define Haar measure rigorously. The intuitive picture for Haar measure of the unitary group can be viewed as a uniform distribution over the unitary group. It is the unique probability measure on the unitary group $U(N)$ that is both left-invariant and right-invariant:

$$\begin{aligned} &\int_{\text{Haar}} dU = 1 \\ &\int_{\text{Haar}} f(VU)dU = \int_{\text{Haar}} f(UV)dU = \int_{\text{Haar}} f(U)dU, \end{aligned} \tag{3.1}$$

for all $V \in U(N)$, where f is an arbitrary function. So Haar random unitaries are extremely random, and they are distributed evenly for the unitary group. In this section, I will review some mathematics that we will use for Haar integral, and some parts are largely inspired by [RY17].

3.2 Schur-Weyl duality

Theorem 1 *Let $L(\mathcal{H}^{\otimes k})$ be the space for all linear maps $A : \mathcal{H}^{\otimes k} \to \mathcal{H}^{\otimes k}$. Let $U(\mathcal{H})$ be the unitary group on $\mathcal{H}$. An operator $A \in L(\mathcal{H}^{\otimes k})$ commute with all $V^{\otimes k}$ with $V \in U(\mathcal{H})$ if and only if A is a linear combination of permutation operators W_π*

$$[A, V^{\otimes k}] = 0, \forall V \in U(\mathcal{H}) \Longleftrightarrow A = \sum_{\pi \in S_k} c_\pi \cdot W_\pi, \tag{3.2}$$

where S_k is the permutation group for k element. A permutation operator W_π with permutation $\pi = \pi(1)\pi(2)\cdots\pi(k)$ is defined as

$$W_\pi |a_1, a_2, \cdots, a_k\rangle = |a_{\pi(1)}, a_{\pi(1)}, \cdots, a_{\pi(k)}\rangle. \tag{3.3}$$

The Pauli string is the tensor product of Pauli operators without any global phase. For example, a Pauli string for four qubits can be $\sigma^{1203} = X \otimes Y \otimes I \otimes Z$. The Pauli operators provide a basis for $L(\mathcal{H})$, the space of linear operators acting on Hilbert space $\mathcal{H}$. Since $\mathrm{Tr}(P_i^\dagger P_j) = 2\delta_{ij}$, we can expand any operator A as

$$A = \sum_j a_j P_j, \;\; a_j = \frac{1}{2}\mathrm{Tr}(P_j^\dagger A). \tag{3.4}$$

One interesting example is the cyclic permutation operator W_{cyc} for k copies of qubit system. We can easily find the Pauli decomposition as

$$W_{\text{cyc}} = \frac{1}{2^{k-1}} \sum_{P_1, \ldots, P_{k-1} \in \mathcal{P}} P_1 \otimes P_2 \otimes \cdots P_{k-1} \otimes (P_1 P_2 \cdots P_{k-1})^\dagger, \tag{3.5}$$

where the sum is over $k-1$ copies of $\mathcal{P}$. In particular, we see the SWAP operator has the decomposition:

$$\mathrm{SWAP} = \frac{1}{2}\sum_{P\in\mathcal{P}} P\otimes P^{\dagger} = \frac{1}{2}(\mathbb{1}\otimes\mathbb{1} + X\otimes X + Y\otimes Y + Z\otimes Z). \tag{3.6}$$

Using tensor diagrams, we can express the above equation as

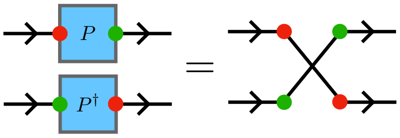

(3.7)

Here, I introduce the arrows on the line as a ”new” ingredient for the tensor diagram. The inward and outward arrows are a good indication of the left and right basis for matrices. With this notation, we can freely move tensors around without causing confusion. In contrast, the tensor diagram introduced in Sec.2.2 is strictly read from left to right. As an example, we can convert the above diagram to the following without confusion:

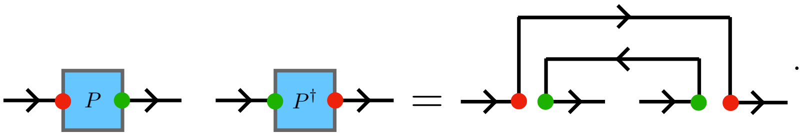

(3.8)

A natural application of (3.8) tells us the average Pauli channel is the identity channel:

$$\frac{1}{2}\sum_{P\in\mathcal{P}} P^{\dagger}AP = \mathbb{1}\,\mathrm{Tr}(A). \tag{3.9}$$

3.3 Haar integral

When we are dealing with random unitary channels, we usually need to integrate the unitary matrices uniformly from unitary group,

$$\int_{\text{Haar}} (U^{\otimes k})^{\dagger} \otimes U^{\otimes k} dU. \tag{3.10}$$

Here I will introduce a diagrammatic approach to calculate the Haar integral above. First, let me define the k-fold channel of an operator $A \in L(\mathcal{H}^{\otimes k})$ as

$$\Phi^{(k)}_{\text{Haar}}(A) := \int_{\text{Haar}} (U^{\otimes k})^{\dagger} A U^{\otimes k} dU. \tag{3.11}$$

By using the right-invariance of the Haar measure, we can show

$$(V^{\otimes k})^{\dagger} \Phi^{(k)}_{\text{Haar}}(A) V^{\otimes k} = \int_{\text{Haar}} f(UV) dU = \Phi^{(k)}_{\text{Haar}}(A), \tag{3.12}$$

and by using the left-invariance of the Haar measure, we can show

$$\Phi^{(k)}_{\text{Haar}}\left((V^{\otimes k})^{\dagger} A V^{\otimes k}\right) = \int_{\text{Haar}} f(VU) dU = \Phi^{(k)}_{\text{Haar}}(A), \tag{3.13}$$

where V is any unitary matrix defined in $\mathcal{H}$. From the (3.12) and (3.12), we can see $\Phi^{(k)}_{\text{Haar}}(A)$ commute with all operators $V^{\otimes k}$. Utilizing Schur-Weyl theorem, we can write $\Phi^{(k)}_{\text{Haar}}(A)$ as a linear combination of permutation operators,

$$\Phi^{(k)}_{\text{Haar}}(A) = \sum_{\pi \in S_k} W_{\pi} \cdot u_{\pi}(A). \tag{3.14}$$

$u_\pi(A)$ is a linear function of A and depending on the permutation π. So we can write it as

$$u_\pi(A) = \mathrm{Tr}(C_\pi A). \tag{3.15}$$

Using (3.13) again, we can find

$$\mathrm{Tr}(C_\pi (V^{\otimes k})^\dagger A V^{\otimes k}) = \mathrm{Tr}(V^{\otimes k} C_\pi (V^{\otimes k})^\dagger A) = \mathrm{Tr}(C_\pi A). \tag{3.16}$$

Therefore, $[C_\pi, V^{\otimes k}] = 0$ for $\forall V$. By using Schur-Weyl theorem again, we can write

$$C_\pi = \sum_{\sigma \in S_k} c_{\pi,\sigma} W_\sigma. \tag{3.17}$$

And we have

$$\Phi^{(k)}_{\mathrm{Haar}}(A) = \sum_{\pi,\sigma \in S_k} c_{\pi,\sigma} W_\pi \,\mathrm{Tr}(W_\sigma A). \tag{3.18}$$

The coefficient $c_{\pi,\sigma}$ depends on two permutation indexes, and it is called the Weingarten matrix. To find the Weingarten matrix, we can use the fact that $\Phi^{(k)}_{\mathrm{Haar}}(W_\lambda) = W_\lambda$,

$$\begin{aligned} W_\lambda &= \sum_{\pi,\sigma} c_{\pi,\sigma} W_\pi \,\mathrm{Tr}(W_\sigma W_\lambda) \\ &= \sum_{\pi,\sigma} c_{\pi,\sigma} W_\pi d^{\#\mathrm{cycles}(\sigma\lambda)} \end{aligned} \tag{3.19}$$

Therefore, we have a linear equation for the Weingarten matrix,

$$\delta_{\pi,\lambda} = \sum_{\sigma \in S_k} c_{\pi,\sigma} Q_{\sigma,\lambda}, \tag{3.20}$$

with $Q_{\sigma,\lambda} = d^{\#\text{cycles}(\sigma\lambda)}$. As we can see, the Weingarten matrix is the inverse matrix of the cycle counting matrix Q, which is symmetric. So the $k-$fold channel solution is

$$\begin{aligned}\Phi^{(k)}_{\text{Haar}}(A) &= \sum_{\pi,\sigma\in S_k} (Q^{-1})_{\pi,\sigma} W_\pi \operatorname{Tr}(W_\sigma A) = \sum_{\pi,\sigma\in S_k} \mathsf{Wg}[\pi,\sigma] W_\pi \operatorname{Tr}(W_\sigma A) \\ &= \sum_{\pi,\sigma\in S_k} \mathsf{Wg}[\pi\sigma^{-1}] W_\pi \operatorname{Tr}(W_\sigma A),\end{aligned} \tag{3.21}$$

where I have introduced the Weingarten function $\mathsf{Wg}[\pi,\sigma] = (Q^{-1})_{\pi,\sigma}$. And in some literature, it also has been written as $\mathsf{Wg}[\pi,\sigma] = \mathsf{Wg}[\pi\sigma^{-1}]$. It will be more useful and intuitive to write down Haar integral as tensor diagrams and give a diagrammatic rule for Haar integral. Like the Feynman diagram, this will help people calculate Haar integral easier. First, I can express Haar integral as a tensor diagram:

$$\int_{\text{Haar}} U^{\otimes k} \otimes (U^\dagger)^{\otimes k} dU$$

$$= \mathbb{E}_U \tag{3.22}$$

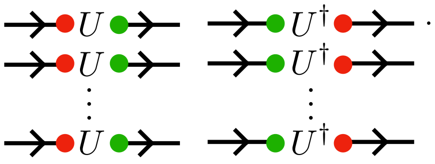

The inward(outward) arrows in the tensor diagram indicate the left(right) indexes of a matrix. The inward nodes of U are labeled by one color, red, as an example. And the outward node of U is labeled by another color, such as green. The coloring rules for $U^\dagger$ are the opposite. And the

Haar integral result can be represented diagrammatically in the following:

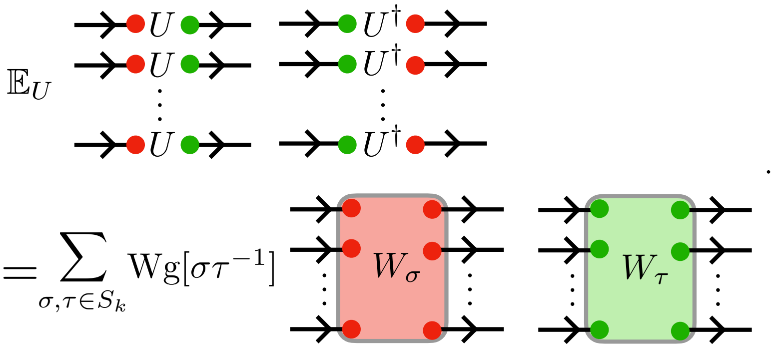

(3.23)

The red dots are connected with red dots with some permutation operator W_σ where $\sigma \in S_k$, and the green dots are connected with green dots with some other permutation operator W_τ where $\tau \in S_k$. We need to sum σ and τ over all the permutations in permutation group S_k and each term is weighted by Weingarten function $\mathrm{Wg}[\sigma\tau^{-1}]$.

3.4 Haar random states

Similarly, we can define the $k-$fold average of Haar random states. If $|\psi\rangle = U|0\rangle$, and U is a Haar random unitary, then $|\psi\rangle$ is the Haar random state or Page state. Now let's use the diagrammatic approach to calculate the $k-$fold average of Haar random states,

$$\int_{\text{Haar}} (|\psi\rangle\langle\psi|)^{\otimes k} d\psi := \Phi^{(k)}_{\text{Haar}}(|\psi\rangle\langle\psi|^{\otimes k}) \tag{3.24}$$

$$\int_{\text{Haar}} (|\psi\rangle\langle\psi|)^{\otimes k} d\psi$$

$$= \mathbb{E}_U$$

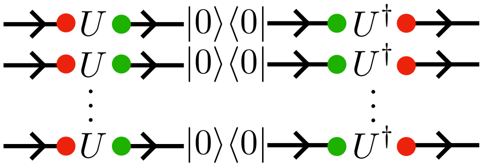

$$= \sum_{\sigma,\tau\in S_k} \mathrm{Wg}[\sigma\tau^{-1}] \tag{3.25}$$

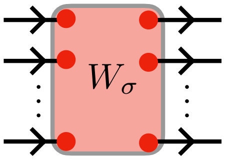

$$= \sum_{\sigma\in S_k} \left(\sum_{\lambda\in S_k} \mathrm{Wg}[\lambda] \right)$$

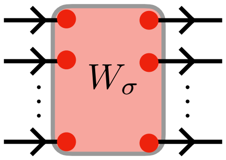

$$= c \sum_{\sigma} W_\sigma$$

In the second equal sign, I have used the fact that no matter what is the permuation connection for the green dots, it will always result in one. In addition, for a fixed permutation σ, $\sum_{\tau\in S_k} \mathrm{Wg}[\sigma\tau^{-1}] = \sum_{\lambda\in S_k} \mathrm{Wg}[\lambda]$. In the last equal sign, I define $c = \sum_{\lambda\in S_k} \mathrm{Wg}[\lambda]$. If we take a trace of (3.25). We see

$$c = \frac{1}{\sum_{\sigma\in S_k} d^{\#\text{cycles}(\sigma)}} = \frac{1}{k!C(k+d-1,k)}, \tag{3.26}$$

where $C(k+d-1,k) = \dfrac{(k+d-1)!}{k!(d-1)!}$ is the binomial coefficient. As a byproduct, we also derived the following equality,

$$\sum_{\sigma\in S_k} \mathrm{Wg}[\sigma] = \frac{1}{\sum_{\sigma\in S_k} d^{\#\text{cycles}(\sigma)}} = \frac{1}{k!C(k+d-1,k)}. \tag{3.27}$$

3.5 Unitary design and frame potential

As we have discussed, a Haar random matrix ensemble can be viewed as a uniform distribution of random unitary matrices. Considering an arbitrary ensemble of unitary $\mathcal{E} = \{U_i | p(U_i) = p_i\}$, *unitary design* is a quantitative measure of how random the ensemble is. We call an ensemble unitary t-design if and only if the ensemble cannot be distinguished with Haar random ensemble up to tth moment. That is the t-fold channel with respect to $\mathcal{E}$,

$$\Phi_{\mathcal{E}}^{(t)}(A) := \sum_j p_j (U_j^{\otimes t})^\dagger A (U_j)^{\otimes t}, \tag{3.28}$$

is the same as Haar random t-fold channel up to t, i.e. $\Phi_{\mathcal{E}}^{(t)}(A) = \Phi_{\text{Haar}}^{(t)}(A)$. Since Pauli strings are basis for operators A, an ensemble is t-design if and only if $\Phi_{\mathcal{E}}^{(t)}(P) = \Phi_{\text{Haar}}^{(t)}(P), \forall P \in (\mathcal{P})^{\otimes t}$. Therefore, we have another useful criteria,

$$\mathcal{E} \text{ is } t\text{-design} \Longleftrightarrow \Phi_{\mathcal{E}}^{(t)}(P) \text{ is a linear combination of } W_\pi \text{ for } \forall P \in (\mathcal{P})^{\otimes t}. \tag{3.29}$$

The definition of unitary t-design does not tell people how to check whether an ensemble is t-design efficiently, since the Pauli group $\mathcal{P}^{\otimes t}$ has exponentially many elements. Now let me introduce *frame potential* as a single quantity for characterizing unitary t-design.

Considering an ensemble of unitaries $\mathcal{E}$, the tth frame potential is defined as

$$F_{\mathcal{E}}^{(t)} := \frac{1}{|\mathcal{E}|^2} \sum_{U,V \in \mathcal{E}} |\operatorname{Tr}(U^\dagger V)|^{2k}, \tag{3.30}$$

where $|\mathcal{E}|$ is the cardinality of $\mathcal{E}$.

Theorem 2 *For any ensemble $\mathcal{E}$ of unitaries,*

$$F_{\mathcal{E}}^{(t)} \geq F_{Haar}^{(t)}. \tag{3.31}$$

The equal sign is satisfied if and only if $\mathcal{E}$ is t-design.

A simple and insightful proof is given by [RY17, Sco08]. Considering $S = \int_{\mathcal{E}} (U^\dagger)^{\otimes t} \otimes U^{\otimes t} dU - \int_{\text{Haar}} (U^\dagger)^{\otimes t} \otimes U^{\otimes t} dU$, we have

$$\begin{aligned} 0 \le \operatorname{Tr}(S^\dagger S) = &\int_{U\in\mathcal{E}} \int_{V\in\mathcal{E}} dU dV |\operatorname{Tr}(U^\dagger V)|^{2t} - 2\int_{U\in\mathcal{E}} \int_{V\in\text{Haar}} dU dV |\operatorname{Tr}(U^\dagger V)|^{2t} \\ &+ \int_{U\in\text{Haar}} \int_{V\in\text{Haar}} dU dV |\operatorname{Tr}(U^\dagger V)|^{2t}. \end{aligned} \tag{3.32}$$

By utilizing the invariance of Haar measure,

$$\begin{aligned} &2\int_{U\in\mathcal{E}} \int_{V\in\text{Haar}} dU dV |\operatorname{Tr}(U^\dagger V)|^{2t} \\ &= 2\int_{U\in\mathcal{E}} \int_{V\in\text{Haar}} \int_{W,W'\in\text{Haar}} dW dW' dU dV |\operatorname{Tr}(U^\dagger W'^\dagger V W^\dagger)|^{2t} \\ &= 2\int_{U\in\mathcal{E}} \int_{V\in\text{Haar}} \int_{W,W'\in\text{Haar}} dW dW' dU dV |\operatorname{Tr}(W^\dagger U^\dagger W'^\dagger V)|^{2t} \\ &= 2\int_{V\in\text{Haar}} \int_{W\in\text{Haar}} dW dV |\operatorname{Tr}(W^\dagger V)|^{2t} \\ &= 2F_{\text{Haar}}^{(t)} \end{aligned} \tag{3.33}$$

Therefore, we have

$$\begin{aligned} &0 \le \operatorname{Tr}(S^\dagger S) = F_{\mathcal{E}}^{(t)} - 2F_{\text{Haar}}^{(t)} + F_{\text{Haar}}^{(t)} = F_{\mathcal{E}}^{(t)} - F_{\text{Haar}}^{(t)} \\ &\iff F_{\mathcal{E}}^{(t)} \ge F_{\text{Haar}}^{(t)}. \end{aligned} \tag{3.34}$$

In particular, using tensor diagram, we can easily calculate $F_{\text{Haar}}^{(t)} = t!$. So the frame potential can be a character for the complexity of an unitary ensemble $\mathcal{E}$.

Chapter 4

Machine learning

4.1 Basic concepts

Artificial intelligence is not a new concept. After the birth of computers and computer programs, people are interested in whether a machine can achieve human performance in some tasks, such as chess playing and recognition of objects. The traditional methods mainly focus on designing algorithms or computer programs that can be carried out by machines in order to achieve some tasks. However, this path is not very successful, even for solving tasks that seem to be easy for a human, such as recognizing objects in pictures. In the meantime, there is another school that advocates artificial intelligence can be built with units emulating the function of neurons in the brain [CHhH02, KSH17b, LBBH98, LBD$^+$89, LBD$^+$90]. With millions of artificial neurons connected together, complicated behaviors of the (artificial) neural network can emerge. This philosophy is really similar to that in modern condensed matter physics, where emergence is an important topic [And72]. A simple neural network is shown in Fig. 4.1, where each colored nodes are the artificial neuron, and the arrows indicate the connectivity between neurons. Now let me explain what I mean by artificial neurons and their connectivity.

In neuroscience, it is widely believed that information is carried by the electrical signal

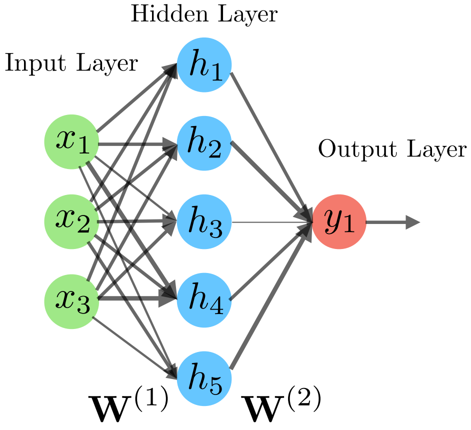

Figure 4.1: A cartoon of feedforward neural network. Each colored node represents a neuron, and its state can be labeled by a real number $x_i \in \mathbb{R}$, $h_i \in \mathbb{R}$, and $y_i \in \mathbb{R}$. The arrows indicate connectivity between neurons, and they can be representated using real matrices $\mathbf{W}^{(i)}$.

or action potential of neurons. Each neuron receives the signal from other neurons through its dendrites and processes the signal. Then it will pass information to other neurons via synapses. A simple mathematical model for the neural is the *linear-nonlinear model*. The state of each neuron is specified by a real number $y \in \mathbb{R}$. First, it will receive a signal from its parent neurons. For example, the neuron h_1 in Fig. 4.1 has three parents: x_1, x_2 and x_3. They pass their state information to h_1 through a linear channel:

$$h_1 = W_{11}^{(1)} x_1 + W_{12}^{(1)} x_2 + W_{13}^{(1)} x_3 + b_1^{(1)}. \tag{4.1}$$

Each coefficient can be viewed as connectivity between neurons. If the coefficient is positive and large, it means they have a strong bond. Then the neuron will process the received signal via a nonlinear activation function $f(h_1)$, such as sigmoid function $f_{a,b}(h_1) = \dfrac{1}{1+e^{-a(h_1-b)}}$. If we list the state of neurons as a column vector, the hidden layer neurons will have states

$$\vec{h} = f_\theta(W^{(1)}\vec{x} + \vec{b}^{(1)}), \tag{4.2}$$

and the output neurons will have states,

$$\vec{y} = F_\Theta(\vec{x}) = f_\theta\left(W^{(2)}(f_\theta(W^{(1)}\vec{x} + \vec{b}^{(1)})) + \vec{b}^{(2)}\right), \tag{4.3}$$

where θ denotes some parameters in the activation function, and Θ is all parameters, including W and $\vec{b}$ in general. As we can imagine, if we have more and more units in the system, the behavior of $F_\Theta(\vec{x})$ is more and more complicated. And it is amazing that people prove the *universal approximation theorem* for a single hidden layer feed-forward neural network, whose architecture is shown in Fig. 4.1. The theorem states that with enough hidden neurons in the middle, $F_\Theta(\vec{x})$ can approximate any continuous function on compact space of the input space $\mathbb{R}^n$. This universal approximation theorem lies the foundation of modern machine learning. The neural network $F_\Theta(\vec{x})$ can approximate any continuous functions with nice properties don't tell people how to find the parameters Θ. Before we discuss this problem in the next section, let me first define some common classes of machine learning models.

The classes of machine learning models are defined based on the tasks. And the three common machine learning models are: *supervised machine learning* models, *unsupervised machine learning* models, and *reinforcement learning* models.

In supervised machine learning, the dataset constitutes pairs of data $\{\vec{x}_i, \vec{y}_i\}$, and the learning task is to predict new $\vec{y}$ given input $\vec{x}$ that is not in the dataset. People usually assume there is a pattern between $\vec{x}$ and $\vec{y}$, and loosely speaking there is a map $\vec{y} = f(\vec{x})$. So mathematically, machine learning of this pattern from the data is equivalent to finding the parameters Θ of the neural network such that $F_\Theta(\cdot)$ approximate $f(\cdot)$. It is called supervised machine learning because the machine has access to correct $\{\vec{y}_i\}$ in the dataset, which serves as a supervisor. For example, pattern recognition in computer vision is supervised machine learning. $\{\vec{x}_i\}$ can be the input images, and $\{\vec{y}_i\}$ are the category of object in the images, such as car, cat or dog. In condensed matter physics, the classification of phases of matter is an important task. Here we could assign system configurations as input $\{\vec{x}_i\}$ and the output $\{y_i\}$ is the phase label of the system [CM17b, KDK$^+$21, HG20, CCMK17, DSDD19].

As the name suggested, in unsupervised machine learning, the dataset is pure $\{\vec{x}_i\}$. We assume all the data are independently sampled from identical hidden distribution $P(\vec{x})$. So the learning task for the machine is to find the best parameters Θ for neural network ansatz $P_\Theta(\vec{x})$, such that $P_\Theta(\cdot)$ approximates $P(\cdot)$. For example, in computer vision, people have billions of human face images, and those images form an empirical distribution of human faces. Computer scientists want the neural network ansatz $P_\Theta(\cdot)$ to approximate the empirical distribution, such that they can use $P_\Theta(\cdot)$ to generate new human faces that look really real. And it is also called the generative model. In physics, usually the distribution of field configurations is determined by the Hamiltonian of the system, $P(\vec{x}) = e^{-\beta H(\vec{x})}/\mathcal{Z}$. However, sampling field configurations from this distribution can be hard. Especially when the system is critical and possesses (quasi-) long-range correlations, the Monte Carlo simulation can suffer from the critical slowing down. So physicists have use neural network models $P_\Theta(\vec{x})$ to learn the true distribution and propose efficient sampling scheme with unsupervised machine learning [LW18, HLWY20, AKS19, KAB$^+$20, AKR$^+$21, HHA$^+$21].

In reinforcement learning, usually, there is no dataset beforehand. What we have is an environment, and we want the agent to achieve some tasks in the environment. For example, we want the agent to drive the car safely. The neural network usually constitutes a policy $\pi_\Theta(a|s)$, which tells the agent what the action a to take if the current environment is in state s. Usually there is reward function $r(s,a)$ for the agent based on the current environment s and the action chosen by the agent a. The agent learns optimal policy by setting parameters Θ in the neural network through the interactions with the environment. For example, game playing is a perfect task for reinforcement learning. And it is amazing to see a trained machine with no human knowledge can beat the world champion in the game of Go. In physics, there are various control problems, such as how to find the best quantum control to mitigate quantum errors and decoherence. Those type of questions are suitable for reinforcement learning models

[AN17, WMD21, KPR20, SQC+21, KHMZ+20].

4.2 Differential programming

In the previous section, I have introduced basic concepts of neural networks and different machine learning models. In all of them, we can view the neural network as a parameterized function $F_\Theta(\cdot)$. The learning tasks usually define an objective function $\mathcal{L}(\Theta)$ that we want to minimize. And the learning of neural networks is to find the best parameters Θ^* such that $\mathcal{L}(\Theta^*)$ is minimized. There are three engines in modern machine learning: 1) big data, 2) efficient hardware, and 3) differential programming platform. The first two are easy to understand. Big data make the empirical loss function more accurate. And in order to handle big data faster, we need efficient hardware, such as GPU and TPU. But what is *differential programming*? In order to answer this, we need to understand how to find optimal parameters Θ in the neural network.

All the fancy methods used in neural network optimization today have the root from *gradient descent*. In order to find the (sub-)optimal Θ, we can update each parameter according to its gradient,

$$\Theta_i \to \Theta_i - \alpha \frac{\partial \mathcal{L}}{\partial \Theta_i}. \tag{4.4}$$

However, it would be very slow if one chose to use finite difference method $\frac{\Delta \mathcal{L}}{\Delta \Theta_i} \approx \frac{\partial \mathcal{L}}{\partial \Theta_i}$, since a neural network usually contains millions of parameters. In addition, the finite difference method also introduces errors in the estimation of gradients. If we don't want to run the program a million times to get the gradient of all the parameters, what should we do? The answer is differential programming.

Differential programming is a new paradigm of computing. By utilizing the differentiable nature of elementary operations with chain rule in calculus, the computer program can

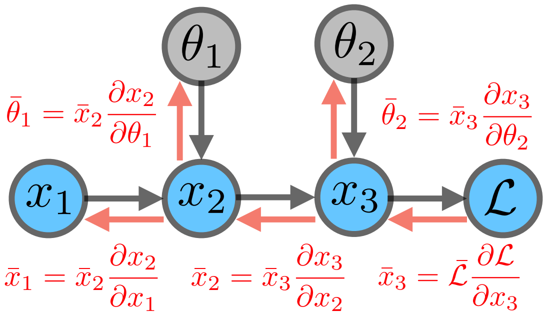

Figure 4.2: A cartoon of computation graph. Every computer program can be represented as a direct acyclic graph, which is called computation graph. The nodes are variables or parameters in the computation. The black arrows define the operations in the forward pass. And the red arrows define the gradient flow in the backward pass.

automatically construct an adjoint program that calculates the exact gradient with respect to any parameters in the program. The meaning of this definition will be more clear with an example. Suppose we want to write a computer program that calculates

$$\mathcal{L} = (x_1^2\theta_1 + \theta_2)^2, \tag{4.5}$$

where θ_1 and θ_2 are parameters to be optimized, and $\mathcal{L}$ is the minimization objective function. In order to do gradient descent, we need $\partial\mathcal{L}/\partial\theta_1$ and $\partial\mathcal{L}/\partial\theta_1$. As we said before, the finite difference estimation is not satisfied because we need to do finite-difference $O(N)$ times for N parameters, and it introduces finite difference errors. If we have learned calculus, then we can calculate the gradient analytically,

$$\begin{aligned} \frac{\partial\mathcal{L}}{\partial\theta_1} &= 2(x_1^2\theta_1 + \theta_2)\cdot x_1^2, \\ \frac{\partial\mathcal{L}}{\partial\theta_2} &= 2(x_1^2\theta_1 + \theta_2). \end{aligned} \tag{4.6}$$

In deriving (4.6), we have implicitly used the chain rules. When the number of parameters θ_i is small and the program is easy, we can perform this analysis manually. However, in modern neural networks, there are millions or billions of parameters, and the structure of the program is complicated. Therefore, differential programming utilizes the calculus chain rule and helps us calculate those gradients in a systematic way. In differential programming, the computer will represent the program as a direct acyclic graph (DAG), where each node stores the value of the parameters or variables, and the black arrows are elementary differentiable operations, such as $\sin(\cdot)$, and $(\cdot)^2$. For example, the computer graph of (4.5) is shown in Fig. 4.2. The black arrows are elementary differentiable operators. For example, the computer transform (4.5) into the following steps: (1) $x_2 = x_1^2\theta_1$, (2) $x_3 = x_2 + \theta_2$, and (3) $\mathcal{L} = x_3^2$.

After the computer compiles the computation graph, it will construct the adjoint computation graph indicated by the red arrows. The adjoint variable is defined as the gradient $\bar{x} = \partial\mathcal{L}/\partial x$, and the output adjoint variable is defined as one, i.e. $\bar{\mathcal{L}} = \partial\mathcal{L}/\partial\mathcal{L} = 1$. The computer will construct the adjoint computation graph backward, and

$$\bar{x} = \sum_{y_i \in \{\text{parent(x)}\}} \bar{y}_i \frac{\partial y_i}{\partial x}, \tag{4.7}$$

For any node, the its parent nodes {parent(x)} can be found via computation graph, and the differential programming platform already encodes the gradient of every elementary computations, for example $d(x^2)/dx = 2x$. So $\partial y_i/\partial x$ is carried out automatically by the differential programming platform. Therefore, after we write the computer program, it will be compiled into a forward computation graph and a backward adjoint computation graph. Each time when we run the forward graph with input x_1, θ_1, θ_2, intermediate results like x_2, and x_3 will be stored. And it will only run the backward graph *once* to get all the gradient $\partial\mathcal{L}/\partial\theta_i$. This new paradigm of computing solves two important issues: 1) gradients with respect to many parameters will be calculated with just one run, and 2) the gradients are accurate without finite-difference errors. In modern days, with the help of differential programming platforms, like PyTorch [PGM+19], TensorFlow [AAB+15b], Jax [BFH+18], people only need to program the forward operations like what they used to do, and the backward computation graph is automatically constructed. And this new paradigm of computing will influence many aspects of physics, such as optimization, variational inference, and optimal control.

4.3 Backpropagating through continuous and discrete samples

As we have seen in the previous section, the main engine of modern deep learning is the *gradient chain rule*. And if the neural network is simply a parameterized continuous function $F_\Theta(\vec{x})$, then it is straightforward to calculate the gradients $\frac{\partial F_\Theta(\vec{x})}{\partial \Theta}$ with auto-differentiation. This is usually the case for supervised machine learning, where the dataset usually contains pairs $\{(\vec{x}_i, y_i)\}$. The y_i are supervised signals, and it is usually assumed there is a underlying relation between inputs and signals, i.e. $F(\vec{x}_i, y_i)$. The goal of supervised machine learning is to use neural network as a variational ansatz $F_\Theta(\vec{x})$ and choose the parameters Θ such that $F_\Theta \approx F$. However, not all neural networks are deterministic. The idea of machine learning is more interesting and powerful when it involves stochastic process and sampling, which is often the case in *unsupervised machine learning*, *generative modeling*, and *reinforcement learning*. In those cases, we may need to backpropagate the gradient through samples from a parameterized distribution. More specifically, we want to minimize an expected loss

$$\mathcal{L}(\theta, \phi) = \mathbb{E}_{x \sim p_\phi(x)}[f_\theta(x)]. \tag{4.8}$$

In order to optimize this loss function with gradient descent, we will need an efficient method to calculate

1. $\nabla_\theta \mathcal{L}(\theta, \phi) = \mathbb{E}_{x \sim p_\phi(x)}[\nabla_\theta f_\theta(x)]$,
2. $\nabla_\phi \mathcal{L}(\theta, \phi) = \int_x \nabla_\phi p_\phi(x) f_\theta(x) dx$.

We see the first type of gradients is easier because we can use Monte Carlo simulation to estimate $\nabla_\theta \mathcal{L}(\theta, \phi)$. So the only two requirements are: 1. $f_\theta(x)$ is a differentiable function with respect to x and 2. we can get samples from $p_\phi(x)$ efficiently. The second type of gradients is not straightforward to calculate, since it does not form an expectation value. Therefore, we cannot use Monte Carlo sampling to get it. In the following, I am going to briefly review three techniques to tackle the gradient of the second type: 1. Score function estimator (REINFORCE), 2. Reparameterization trick and 3. Gumbel-softmax trick.

REINFORCE [Wil92]: The score function estimator is also called *REINFORCE* because it is a trick that usually appears in the reinforcement learning. By utilizing the property, $\nabla_\phi p_\phi(x) = p_\phi(x)\nabla_\phi \log p_\phi(x)$, we can rewrite the gradient as

$$\begin{aligned}\nabla_\phi \mathcal{L}(\theta,\phi) &= \int_x f_\theta(x)\nabla_\phi p_\phi(x)dx \\ &= \int_x f_\theta(x)\nabla_\phi \log(p_\phi(x))p_\phi(x)dx \\ &= \mathbb{E}_{x\sim p_\phi}\left[f_\theta(x)\nabla_\phi \log p_\phi(x)\right].\end{aligned} \tag{4.9}$$

With the above trick, we can again use Monte Carlo to estimate the gradient. The above equation can be incorporated with the auto-differentiation (auto-diff) platform. In the forward calculation, we want to sample x from $p_\phi(x)$ and estimate $\mathbb{E}_{x\sim p_\phi}[f_\theta(x)]$. And in the backward calculation, we want the auto-diff platform automatically calculate $\mathbb{E}_{x\sim p_\phi}\left[f_\theta(x)\nabla_\phi \log p_\phi(x)\right]$. To achieve this, we define

$$r_\phi(x) = \log p_\phi(x) + (1-\log p_\phi(x)).\text{detach}, \tag{4.10}$$

where $(A).detach$ creates a flag such that A is detached from the backpropgation graph. REINFORCE method sometimes can have large variance. To mitigate this effect, one can substract a baseline from it,

$$\nabla_\phi \mathcal{L}(\theta,\phi) = \mathbb{E}_{x\sim p_\phi}\left[(f_\theta(x)-b(x))\nabla_\phi \log p_\phi(x)\right] + \mathbb{E}_{x\sim p_\phi}[b(x)]. \tag{4.11}$$

Reparameterization trick [KW14]: Sometime, the random variable x can be reparameterized as a *deterministic* function g that takes ϕ and a random variable ε that does not depend on ϕ as input, i.e. $x = g(\phi, \varepsilon)$. One famous example is the Gaussian random variable $x \sim \mathcal{N}(\mu(\phi), \sigma^2(\phi))$ can be rewritten as $x = \mu(\phi) + \sigma^2(\phi) * \varepsilon$, and $\varepsilon \sim \mathcal{N}(0,1)$. With reparameterization trick, we can rewrite the gradient as

$$\begin{aligned}\nabla_\phi \mathcal{L}(\theta, \phi) &= \nabla_\phi \mathbb{E}_{\varepsilon \sim p(\varepsilon)}\left[f_\theta(g(\phi, \varepsilon))\right] \\ &= \mathbb{E}_{\varepsilon \sim p(\varepsilon)}\left[f_\theta'(g(\phi, \varepsilon)) \nabla_\phi g(\phi, \varepsilon)\right].\end{aligned} \tag{4.12}$$

One should notice, the requirement for reparameterization trick is $f_\theta(x)$ must be differentiable with respect to x, and $g(\phi, \varepsilon)$ must exist and be differentiable with respect to ϕ.

Gumbel-softmax [JGP17]: This method can be viewed as a reparameterization trick for the discrete samples. Let discrete random variable $x \sim \text{Cat}(\pi_1, \ldots, \pi_k)$ follow a categorical distribution, where π_i is the probability of x in category i. One way of sampling this discrete variable is utilizing the accumulated distribution. Let $\eta \sim \text{Uniform}[0,1]$, and $x = \{\max i | \pi_1 + \pi_2 + \cdots + \pi_{i-1} \leq \eta\}$. With the reparameterization trick, a categorical random variable with distribution $\text{Cat}(\pi_1, \ldots, \pi_k)$ can also be reparameterized as a deterministic function acting on probability values π_i and a random variable $\varepsilon_i \sim \text{Gumbel}[0,1]$ following the Gumbell distribution. The Gumbell random variable can be obtained via $\varepsilon_i \sim -\log(-\log(\eta_i))$, where $\eta_i \sim \text{Uniform}[0,1]$. Then it can be shown that the random variable

$$x := g(\vec{\pi}, \vec{\varepsilon}) = \text{argmax}_k(\log \pi_k + \varepsilon_k), \tag{4.13}$$

with $\varepsilon_k \sim \text{Gumbel}[0,1]$ will follow the correct categorical distribution $\text{Cat}(\vec{\pi})$. The reparameterization trick has successfully separate the effect of randomness and parameter of distributions. And it is usually referred as *Gumbell-max* trick. However, the argmax function is not differentiable.

To make it differentiable, people use softmax function to approximate it. More precisely, one can define a Gumbell-softmax random variable,

$$x_k = \text{Softmax}\left((\log \pi_k + \varepsilon_k)/\tau\right) = \frac{\exp\left((\log \pi_k + \varepsilon_k)/\tau\right)}{\sum_k \exp\left((\log \pi_k + \varepsilon_k)/\tau\right)}. \tag{4.14}$$

In above, we introduce a hyper-parameter τ called temperature, since it resembles the temperature of Fermi-Dirac distribution. When $\tau \to 0$, one can see random variable x_k will become a *one-hot* vector that follows the categorical distribution $\text{Cat}(\vec{\pi})$. And when $\tau \to \infty$, it will become a uniform distribution. In practice, people usually choose $\tau \sim 1/N$, where N is the number of categories. With Gumbell-max and softmax, now the gradient can be backpropagate through the discrete samples. As a summary, I list the comparison of different tricks in the Table.4.1.

Table 4.1: Summary of different methods to backpropagate gradient through samples.

Method	**Continuous** or **Discrete**	**Exact** $p(x)$	$\nabla_x f_\theta(x)$ **must exist**
REINFORCE	Both	Yes	No
Reparameterization	Continuous	Yes	Yes
Gumbell-softmax	Discrete	No	Yes

Part II

Classical shadow representation of quantum states

Chapter 5

Hamiltonian-driven shadow tomography of quantum states

5.1 Introduction

Quantum state tomography [VR89, JKMW01, CFS02, PR04, RLR$^+$04, CW20] is an essential quantum technology underlying the characterization of quantum devices and the discrimination of quantum states. It aims to reconstruct the density matrix from repeated measurements of identically prepared copies of a quantum system. While the complexity of exact tomography of the full density matrix scales exponentially with the system size due to the curse of dimensionality[HHJ$^+$17a], approximate tomography with polynomial complexity has been developed with assumptions of the underlying quantum state, including matrix product state tomography[CPF$^+$10a, LMH$^+$17b, WHW$^+$17], reduced density matrix tomography[LPW02, LW02, CJR$^+$12, CJZZ12, XLC$^+$19, Di4], and machine learning tomography[TMC$^+$18b, TM18, CTMA18, XX18, QFK18, TTv$^+$19, CLP$^+$19, CGW$^+$20, NFJ$^+$20].

Among various tomography schemes, *shadow tomography*[BKL$^+$17, Aar17, AR19, HKP20, ZRM20, EG20a, EKH$^+$20, CYZF20a, ZZL20, SZK$^+$21, GZJ21] has recently attracted much research attention. Given a copy of a N-qubit quantum system described by the density matrix ρ, the shadow tomography protocol first performs a random unitary transformation U on the state $\rho \to \rho' = U\rho U^\dagger$, then measures the transformed state ρ' in the computational basis (i.e. simultaneously measuring the Pauli-Z operator on every qubit), as illustrated in Fig. 6.1. After the measurement, the system collapses to a pure state $|b\rangle$ labeled by the bit-string $b \in \{0,1\}^N$ of measurement outcomes. Classical snapshots $\hat{\sigma} = U^\dagger|b\rangle\langle b|U$ of the quantum system can be collected through repeated measurements. Given the knowledge about the random ensemble of the unitary transformation U, the density matrix ρ can be reconstructed as a particular linear combination of the ensemble average of classical snapshots $\hat{\sigma}$, where the linear channel only depends on properties of the unitary ensemble.

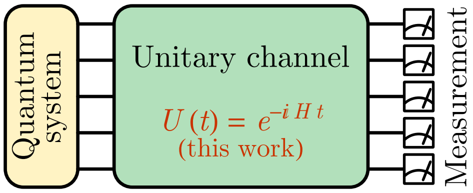

Figure 5.1: Shadow tomography protocol. Specially in this work, the unitary channel is considered to be a time evolution generated by a random Hamiltonian H for time t.

The unitary transformation U plays an important role in the protocol to scramble the quantum information, such that the computational basis measurement on the scrambled state ρ' is effectively a simultaneous measurement of a random set of N commuting operators $\{U^\dagger Z_i U\}_{i=1}^N$ on the original state ρ. In this way, each measurement returns measurement outcomes for 2^N observables (as products of arbitrary subsets of the commuting operators, e.g. $U^\dagger(Z_i Z_j \cdots)U$), which provides an efficient way to extract information from the quantum system. The shadow tomography saturates the fundamental lower bound on the minimal number of independent samples required for tomography,[HHJ$^+$17a, HKP20] achieving the maximal efficiency. In particular, when the unitary ensemble is Haar random (or any unitary 2-design such as random Clifford circuits), the scrambling is strongest to enable estimating rank-1 observables (such as quantum fidelity) with a constant number of samples that does not scale with the system size N.

However, insufficient scrambling power of the unitary channel will affect the efficiency of shadow tomography. To investigate the effect of scrambling power on the tomography efficiency, we consider the random unitary channel $U(t) = e^{-\mathrm{i}Ht}$ to be the time evolution generated by a chaotic Hamiltonian H, which enables us to tune the scrambling power of the unitary channel by the evolution time t. Such a tomography scheme will be called the *Hamiltonian-driven shadow tomography*. When $t = 0$, the unitary channel is an identity that has no scrambling power, then the shadow tomography is only able to reconstruct the diagonal part of the density matrix ρ (in the computational basis). As the quantum system evolves for some time, the computational basis measurement will be able to probe the off-diagonal part of ρ more efficiently, but it also becomes less efficient to infer the diagonal part of ρ as the diagonal information starts to get scrambled with the off-diagonal information. In the long-time limit, the unitary ensemble approaches the Haar limit, and the efficiency for both diagonal and off-diagonal parts converges to the same limit. In this work, we derive the reconstruction channel for the Hamiltonian-driven shadow tomography and analyze its efficiency as a function of the evolution time t and the total Hilbert space dimension $D = 2^N$. We find that, given observables are Pauli operators, there exists an intermediate time range $1 \lesssim t \lesssim D^{1/6}$, in which the Hamiltonian-driven shadow tomography only need $\sim \mathrm{Tr}(O_d^2)/D$ samples to estimate a diagonal observable O_d, and $\sim \mathrm{Tr}(O_o^2)$ samples to estimate an off-diagonal observable O_o, which is superior to the shadow tomography based on 2-design random unitaries. Our result may find applications in shadow tomography with shallow quantum circuits, which are feasible in the noisy intermediate-scale quantum (NISQ)[Pre18b] era.

5.2 Brief Review of Shadow Tomography

Consider a N-qubit system described by an unknown density matrix ρ which can be prepared repeatedly. The goal of the quantum state tomography is to infer ρ from repeated measurements on independent copies of the state. In each experiment, the state ρ is first evolved by a unitary operator U drawn from the distribution $P(U)$ and then measured in the computational basis. The measurement will prepare a classical snapshot of the quantum system $\hat{\sigma} \equiv U^\dagger |b\rangle\langle b|U$ with $b \in \{0,1\}^N$ labeling the measurement outcome. After M repeated experiments, we will collect M classical snapshots $\{\hat{\sigma}_1, \cdots, \hat{\sigma}_M\}$. We can view the average mapping from the quantum state ρ to its classical snapshots as a measure-and-prepare quantum channel $\mathcal{M}$,

$$\mathcal{M}(\rho) = \mathbb{E}\hat{\sigma} = \mathbb{E}\left[U^\dagger |b\rangle\langle b|U\right], \tag{5.1}$$

where the averaging is taken over both the unitary ensemble $P(U)$ and the possible measurement outcomes following the probability distribution $P(b|U) = \langle b|U\rho U^\dagger|b\rangle$.

The quantum state ρ can be reconstructed by applying the inverse channel $\mathcal{M}^{-1}$ (i.e. the reconstruction channel)

$$\rho = \mathcal{M}^{-1}(\mathbb{E}\hat{\sigma}) = \mathbb{E}\mathcal{M}^{-1}(\hat{\sigma}) = \mathbb{E}\hat{\rho}, \tag{5.2}$$

where $\hat{\rho} \equiv \mathcal{M}^{-1}(\hat{\sigma}) = \mathcal{M}^{-1}(U^\dagger|b\rangle\langle b|U)$ is called the *classical shadow* of the original quantum state ρ. The reconstruction channel $\mathcal{M}^{-1}$ does not admit physical implementation, as it is not completely positive in general. Nevertheless, given the distribution $P(U)$ of the unitary ensemble, the reconstruction channel $\mathcal{M}^{-1}$ can be calculated and applied to the classical snapshots $\{\hat{\sigma}_1, \cdots, \hat{\sigma}_M\}$ by classical post-processing[HKP20] to obtain the set of classical shadows $\{\hat{\rho}_1, \cdots, \hat{\rho}_M\}$, which can then be used to estimate both linear and nonlinear functions of the underlying quantum state ρ. For example, the expectation value $o = \mathrm{Tr}(O\rho)$ of a physical observable O is a linear function of ρ, which can be estimated as

$$o = \mathrm{Tr}(O\rho) = \mathbb{E}[\mathrm{Tr}(O\hat{\rho})] \simeq \frac{1}{M}\sum_i \mathrm{Tr}(O\hat{\rho}_i). \tag{5.3}$$

Nonlinear functions, such as $\mathrm{Tr}(O\rho\otimes\rho)$, can also be estimated similarly,

$$\begin{aligned}\mathrm{Tr}(O\rho\otimes\rho) &= \mathrm{Tr}(O\mathbb{E}[\hat{\rho}]\otimes\mathbb{E}[\hat{\rho}]) \\ &\simeq \frac{1}{M(M-1)}\sum_{i\neq j}\mathrm{Tr}(O\hat{\rho}_i\otimes\hat{\rho}_j),\end{aligned} \tag{5.4}$$

given the fact that $\hat{\rho}_i$ and $\hat{\rho_j}$ are statistically independent. When number of experiments M is large enough, the statistical averages over classical shadows in Eq. (5.3) and Eq. (5.4) will converge to their corresponding expectation values without bias.

5.3 Hamiltonian-Driven Shadow Tomography.

The shadow tomography approach crucially relies on the realization of the unitary channel with sufficient scrambling power. Tomography schemes using global Haar/Clifford unitary ensemble have been proposed in Ref. [HKP20]. In practice, it remains challenging to realize these unitary ensembles on NISQ devices. We propose to generate the scrambling unitary channel by some quantum chaotic Hamiltonian H through time-evolution $U(t) = e^{-\mathrm{i}Ht}$. With this setup, the quantum dynamics also enter the discussion, as it becomes meaningful to discuss how the reconstruction channel $\mathcal{M}^{-1}$ and the tomography efficiency depends on the evolution time t.

To analyze the problem, we model the chaotic Hamiltonian generally by a random Hermitian operator H drawn from the Gaussian unitary ensemble (GUE) in $D = 2^N$ dimensional Hilbert space, following the probability distribution $P(H) \propto \exp(-\frac{D}{2}\mathrm{Tr}H^2)$. The energy scale is such normalized that the spectral density of H approaches the Wigner semicircle law $\rho(E) = \frac{1}{2\pi}\sqrt{4-E^2}$ of the spectral radius 2 as $D\to\infty$. The unitary time-evolution generated by H admits the following eigen decomposition

$$U(t) = e^{-iHt} = V\Lambda(t)V^{\dagger}. \tag{5.5}$$

where V is the matrix diagonalizing H and $\Lambda(t)$ is the diagonal matrix $\Lambda_{nn'}(t) = e^{-iE_n t}\delta_{nn'}$ with E_n being the eigen energies of H. For the GUE random matrix H, the unitary V is Haar random.

Substitute the eigen decomposition of U in Eq. (5.5) to Eq. (5.1), the quantum channel $\mathcal{M}$ can be expanded as

$$\begin{aligned}\mathcal{M}(\rho) = \mathop{\mathbb{E}}_{V,\Lambda} \sum_{b\in\{0,1\}^N} & V\Lambda(t)^{\dagger}V^{\dagger}|b\rangle\langle b|V\Lambda(t)V^{\dagger} \\ & \langle b|V\Lambda(t)V^{\dagger}\rho V\Lambda(t)^{\dagger}V^{\dagger}|b\rangle.\end{aligned} \tag{5.6}$$

Using the results of Haar measure integral[Wei78, CŚ06] and the spectral form factor of GUE matrices[CHLY17], to the leading order of D, the quantum channel simplifies to

$$\begin{aligned}\mathcal{M}(\rho) &= \mathcal{M}(\mathbb{1}/D + \rho_o + \rho_d) \\ &= \frac{\mathbb{1}}{D} + \frac{\rho_o}{\alpha_D(t)} + \frac{\rho_d}{\beta_D(t)},\end{aligned} \tag{5.7}$$

where $\mathbb{1}$ stands for the identity matrix, ρ_o is the off-diagonal part of ρ, and ρ_d is the traceless diagonal part of ρ. The coefficients $\alpha_D(t)$, and $\beta_D(t)$ are defined as

$$\begin{aligned}\alpha_D(t) &= \Big(\frac{1}{D+1} - \lambda_D(t)\Big)^{-1}, \\ \beta_D(t) &= \Big(\frac{1}{D+1} + D\lambda_D(t)\Big)^{-1}, \\ \lambda_D(t) &= \frac{(Dr^2(t) + r(2t))^2 - 4r^2(t)}{(D+3)(D^2-1)},\end{aligned} \tag{5.8}$$

and $r(t) = J_1(2t)/t$ with J_1 being the Bessel function of the first kind (which captures the leading-D behavior of spectral form factors). See Appendix A.1 for detail derivations.

In the short time limit $t \to 0$, the quantum channel becomes

$$\mathcal{M}(\rho) \to \frac{\mathbb{1}}{D} + \rho_d. \tag{5.9}$$

As expected, only the diagonal part of the density matrix will be transmitted through this quantum channel. Because in the absence of time-evolution, the computational basis measurement can only extract the diagonal information of the density matrix, and the off-diagonal information is completely lost. As the measurement is not tomographically complete, the channel $\mathcal{M}$ is not invertible at $t = 0$. However, after a finite time-evolution, a finite fraction of the off-diagonal information will be scrambled to the diagonal part of the density matrix and become accessible to the measurement, then the channel $\mathcal{M}$ will be invertible (as long as $t \neq 0$). In the long-time limit $t \to +\infty$, the quantum channel converges to the known result[GKKT18] for shadow tomography with unitary 2-design

$$\mathcal{M}(\rho) \to \frac{\mathbb{1}}{D} + \frac{\rho_d}{D+1} + \frac{\rho_o}{D+1}. \tag{5.10}$$

The diagonal and off-diagonal parts converge to the same channel transmission rate in this limit, indicating that the quantum information has been fully scrambled.

Since $\mathbb{1}/D$, ρ_o, and ρ_d are all orthogonal to each other, the inverse channel of Eq. (5.7) is simply obtained by inverting the coefficient of each term,

$$\hat{\rho} = \mathcal{M}^{-1}(\hat{\sigma}) = \frac{\mathbb{1}}{D} + \alpha_D(t)\hat{\sigma}_o + \beta_D(t)\hat{\sigma}_d, \tag{5.11}$$

where $\hat{\sigma}_o$ and $\hat{\sigma}_d$ are respectively the off-diagonal and the traceless diagonal part of the classical snapshot $\hat{\sigma} = U(t)^\dagger |b\rangle\langle b| U(t)$. Coefficients $\alpha_D(t)$ and $\beta_D(t)$ were defined in Eq. (5.8). Given the reconstruction channel $\mathcal{M}^{-1}$ for the Hamiltonian-driven shadow tomography, we can use classical shadows $\hat{\rho} = \mathcal{M}^{-1}(\hat{\sigma})$ to reconstruct the density matrix $\rho = \mathbb{E}\hat{\rho}$ and to estimate physical properties of the quantum system following Eq. (5.3) and Eq. (5.4).

5.4 Tomography Efficiency Analysis.

We now analyze the efficiency of the Hamiltonian-driven shadow tomography, i.e. how many independent copies of ρ are typically needed to predict functions of ρ to a suitable precision. We will mainly focus on the efficiency of predicting linear functions, but our result can be generalized to nonlinear function predictions systematically.

According to Eq. (5.3), the linear function $o = \mathrm{Tr}(O\rho)$ can be estimated from the classical shadow $\hat{\rho}$ via $o = \mathbb{E}[\mathrm{Tr}(O\hat{\rho})] = \mathbb{E}\hat{o}$, where $\hat{o} \equiv \mathrm{Tr}(O\hat{\rho})$ can be viewed as a random variable derived from the classical shadow. In practice, we conduct M experiments to collect classical shadows $\hat{\rho}_i$, and estimate o using

$$o_{\mathrm{avg}} = \frac{1}{M}\sum_{i=1}^{M}\hat{o}_i = \frac{1}{M}\sum_{i=1}^{M}\mathrm{Tr}(O\hat{\rho}_i). \tag{5.12}$$

Based on Chebyshev's inequality, the probability of the estimation o_{avg} to deviate from its expectation value o is bounded by its variance $\mathrm{Var}(o_{\mathrm{avg}})$ as $\Pr(|o_{\mathrm{avg}} - o| \geq \delta) \leq \mathrm{Var}(o_{\mathrm{avg}})/\delta^2$. To control the deviation probability within a desired statistical accuracy ε, we require $\mathrm{Var}(o_{\mathrm{avg}})/\delta^2 = \mathrm{Var}(\hat{o})/(M\delta^2) \leq \varepsilon$, where δ bounds the additive error in o_{avg} and M is the number of identity copies of ρ used. In other words, the number of experiments needed to achieved the desired tomography accuracy is given by

$$M \geq \mathrm{Var}(\hat{o})/\varepsilon\delta^2. \tag{5.13}$$

Therefore the problem boils down to analyzing the variance $\mathrm{Var}(\hat{o})$ of the single-shot random variable $\hat{o} = \mathrm{Tr}(O\hat{\rho})$. We will assume the physical observable O to be traceless ($\mathrm{Tr}\, O = 0$), since adding O by $c\mathbb{1}$ only shifts $\hat{o}$ by a constant c (given $\mathrm{Tr}(\hat{\rho}) \equiv 1$), which does not affect its variance.

We can further bound the variance by

$$\mathrm{Var}(\hat{o}) = \mathbb{E}[\hat{o}^2] - \mathbb{E}[\hat{o}]^2 \leq \mathbb{E}[\hat{o}^2], \tag{5.14}$$

By decomposing the observable $O = O_o + O_d$ to its off-diagonal part O_o and traceless diagonal part O_d, we can evaluate $\mathbb{E}[\hat{o}^2]$, which takes the form of

$$\mathbb{E}[\hat{o}^2] = V[O^{\otimes 2}] = V[O_o^{\otimes 2}] + V[O_d^{\otimes 2}] + V[O_o \otimes O_d], \tag{5.15}$$

where $V[O^{\otimes 2}]$ is a linear function of the double-operator $O^{\otimes 2}$, whose explicit form is given in Appendix A.2. More explicitly, we can express $V[O_o^{\otimes 2}] = \mathrm{Tr}(O_o^2)F_o(t)$ and $V[O_d^{\otimes 2}] = \mathrm{Tr}(O_d^2)F_d(t)$ with the dynamic form factors given by

$$\begin{aligned} F_o(t) &= f_1(t) + f_2(t)\frac{\mathrm{Tr}(O_o^2\rho)}{\mathrm{Tr}(O_o^2)} + f_3(t)\frac{\mathrm{Tr}(O_o^2\rho_d)}{\mathrm{Tr}(O_o^2)}, \\ F_d(t) &= f_4(t) + f_5(t)\frac{\mathrm{Tr}(O_d^2\rho)}{\mathrm{Tr}(O_d^2)}, \end{aligned} \tag{5.16}$$

where the time-dependent functions $f_{1,2,3,4,5}$ are given in (A.13) of Appendix A.2. Combining the results in Eq. (5.13), Eq. (5.14) and Eq. (5.15), we arrive at the following theorems regarding the efficiency of Hamiltonian-driven shadow tomography, which are central results of this work.

Theorem 3 *Given an off-diagonal operator O_o, the Hamiltonian-driven shadow tomography with an evolution time t uses $O(\frac{1}{\varepsilon\delta^2}\mathrm{Tr}(O_o^2)F_o(t))$ independent copies of ρ to estimate the expectation value o_o of the observable O_o to the precision that $Pr(|o_o - \mathrm{Tr}(O_o\rho)| \geq \delta) \leq \varepsilon$.*

Theorem 4 *Given a traceless diagonal operator O_d, the Hamiltonian-driven shadow tomography with an evolution time t uses $O(\frac{1}{\varepsilon\delta^2}\mathrm{Tr}(O_d^2)F_d(t))$ independent copies of ρ to estimate the expectation value o_d of the observable O_d to the precision that $Pr(|o_d - \mathrm{Tr}(O_d\rho)| \geq \delta) \leq \varepsilon$.*

In the long-time limit ($t \to \infty$),

$$F_o(\infty) = F_d(\infty) = 1 + 2\frac{\mathrm{Tr}(O^2\rho)}{\mathrm{Tr}(O^2)} \leq 3. \tag{5.17}$$

There is no difference between diagonal and off-diagonal observables in terms of tomography efficiency. The required number of samples (i.e. copies of ρ) scales as $M \sim \mathrm{Tr}(O^2)/(\varepsilon\delta^2)$ which agrees with the result for shadow tomography using Haar/Clifford random unitaries.[HKP20] For rank-1 observables, such as quantum fidelity, $\mathrm{Tr}(O^2)$ is independent of the system size N, then the shadow tomography only needs a constant number of experiments to achieve the desired accuracy.

We dubbed $F_o(t)$ and $F_d(t)$ as *sample form factors* since they can be interpreted as the ratio of the required number of samples in the Hamiltonian-driven shadow tomography to that in the random-unitary-based shadow tomography (with 2-design unitary channels). A sample form factor less (or greater) than its long-time limit indicates the Hamiltonian-driven shadow tomography is more (or less) efficient than the random-unitary-based shadow tomography. With these understandings, we investigate the early time behavior of $F_o(t)$ and $F_d(t)$.

At early time, the behavior of sample form factors can be rather complicated. However, we found that for Pauli observables (i.e. O is a Pauli operator), they take particularly simple forms (to the leading order in D)

$$F_o(t) = \frac{1}{1-r^4(t)},\ F_d(t) = \frac{1}{1+Dr^4(t)}. \tag{5.18}$$

Recall that $r(t) = J_1(2t)/t$, the time-dependence of $F_o(t)$ and $F_d(t)$ for Pauli observables are plotted in Fig. 5.2. In the following, we will mainly focus on the Pauli observables. The general cases are discussed in Appendix A.2 and A.4.

For off-diagonal observables, the sample form factor $F_o(t)$ diverges as $t \to 0$, as it is impossible to infer the off-diagonal information from computational basis measurement in the absence of information scrambling. As time evolves, the off-diagonal information gets scrambled to the diagonal part, then $F_o(t)$ decays with t as t^{-2}, as shown in Fig. 5.2(a). $F_o(t)$ quickly approaches 1 after a characteristic time $t_o \sim 1$ set by the inverse energy scale of the Hamiltonian H, which was identified as the scrambling time of the system in Ref. [YG18]. Thus in the Hamiltonian-driven shadow tomography, one just needs to wait for the scrambling time to achieve

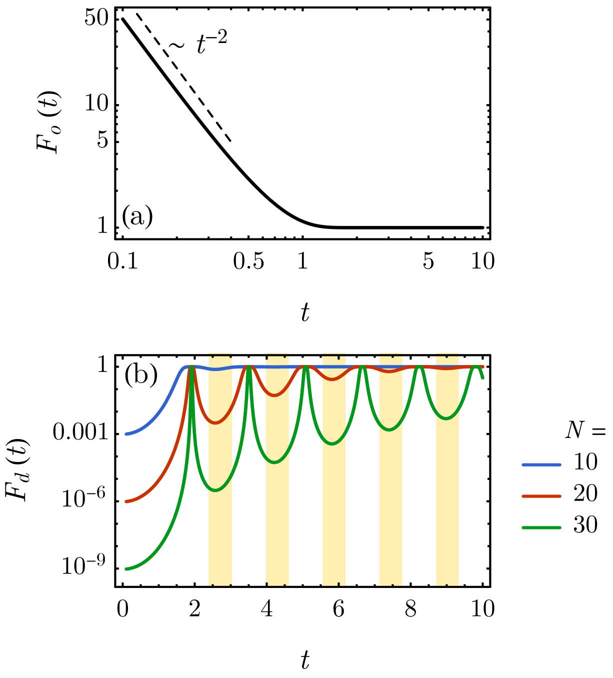

Figure 5.2: Time-dependence of the sample form factor: (a) $F_o(t)$ for off-diagonal observables and (b) $F_d(t)$ for diagonal observables. The time scales in unit of the inverse energy scale of the chaotic Hamiltonian H.

effectively the same efficiency as random-unitary-based shadow tomography in terms of off-diagonal observable. Moreover, the scrambling time t_o is independent of D (or the system size N) in our model, given the non-local nature of the GUE random Hamiltonian.

For diagonal observables, the sample form factor $F_d(t)$ is of the order $F_d(t) \sim D^{-1} = 2^{-N}$ at $t = 0$, which is exponentially small in system size N. Without any unitary scrambling, the computational basis measurement is directly measuring the diagonal information of the density matrix, therefore it requires much fewer samples to infer diagonal observables as compared to that of the general-purpose random-unitary-based shadow tomography. As time evolves, the diagonal information is scrambled away, hence more samples are required to achieve the accuracy goal. So the tomography efficiency decreases with time for diagonal observables, in contrast to the increasing efficiency for off-diagonal observables.

Interestingly, $F_d(t)$ peaks to its maximal value periodically before it saturates to its long-time limit, as shown in Fig. 5.2(b). At these peaks, the diagonal information is maximally scrambled, therefore we name this phenomenon as *scrambling beats*, which was first reported in Ref. [YG18]. The peaks occur at times $t_k = x_k/2$, where x_k is the kth zero of the Bessel function $J_1(x)$. Under coherent Hamiltonian evolution, the scrambled information can partially bounce back in a finite-size system, leading to the beat behavior of $F_d(t)$. But how long scrambling beats will last depends on the system size. The characteristic time for scrambling beats to die off is of the order $t_d \sim D^{1/6} = 2^{N/6}$, when $Dr^4(t_d) \sim 1$. Before this time scale, there exist time windows between peaks, as yellow-shaded regions in Fig. 5.2(b), when the sample form factor maintains at a low level of $F_d(t) \sim D^{-1}$.

For a large enough system, the time scales $t_d \sim D^{1/6}$ and $t_o \sim 1$ are well-separated, which admits an intermediate time range $t_o \lesssim t \lesssim t_d$ where the Hamiltonian-driven shadow tomography can simultaneously achieve exponentially higher efficiency for diagonal observables and the same efficiency for off-diagonal observables, as compared to the random-unitary-base shadow tomography. Such behavior could potentially be advantageous when the diagonal observables are of more interest in certain tomography tasks.

5.5 Efficiency for Nonlinear Functions.

Our result can be generalized to analyze the tomography efficiency of predicting nonlinear functions of the density matrix ρ. For nonlinear function involving k copies of ρ, which generally takes the form of $\mathrm{Tr}(O\rho^{\otimes k})$, the variance of the shadow estimation can be bounded by

$$\begin{aligned}\mathrm{Var}(\mathrm{Tr}(O\hat{\rho}^{\otimes k})) &\leq \mathbb{E}[\mathrm{Tr}(O\hat{\rho}^{\otimes k})^2] \\ &\leq \sum_{\boldsymbol{\alpha}} \mathrm{Tr}(O_{\alpha_1^L\alpha_2^L\cdots\alpha_k^L} O_{\alpha_1^R\alpha_2^R\cdots\alpha_k^R}) \prod_{\alpha_i} F_{\alpha_i}(t),\end{aligned} \tag{5.19}$$

where $\boldsymbol{\alpha} = (\alpha_1, \cdots, \alpha_k)$ has k-component, and each

$$\alpha_i = (\alpha_i^L, \alpha_i^R) \in \{(o,o),(d,d),(I,o),(o,I),(I,d),(d,I),(o,d),(d,o)\}$$

is a pair of labels where $\alpha_i^{L/R}$ labels the i-th tensor leg of left/right O operator. The summation is over all combinations of $\boldsymbol{\alpha}$. The details are discussed in Appendix A.5. Plugging the variance into Eq. (5.13), the number of required experiments can be bounded as well for nonlinear functions.

5.6 Summary and Discussions.

We propose to use Hamiltonian generated unitary evolution to scramble the quantum information for shadow tomography. We provide an unbiased estimator of the density matrix for all ranges of evolution time. We investigated the efficiency of the Hamiltonian-driven shadow tomography. In particular, for Pauli observables, we showed that it can be superior to the shadow tomography based on 2-design random unitaries within an intermediate time window. Although our analysis is based on the GUE random Hamiltonian, we expect that the result could be generalized to other types of quantum chaotic Hamiltonians[CSM$^+$21, CMH$^+$21]. In addition, we expect quantum dynamics generated by generic chaotic Hamiltonians are more feasible to existing quantum devices, such as nuclear spin systems and cold atom quantum simulators, than deep quantum circuits. In experiments, the Sachdev-Ye-Kiteav model which exhibits the same level statistics as GUE Hamiltonians has been successfully simulated on those quantum platforms recently [LYL$^+$19, WS21]. In a follow-up work [HCY21], we demonstrate that approximate shadow tomography can be performed using a simple spin chain Hamiltonian modeled after programmable trapped ions or Rydberg atom array systems, implying that our method is hardware efficient for existing quantum devices [Saf16a, MCD$^+$21]. Another interesting possibility is to consider Hamiltonians consist of random Pauli strings with random coefficients. In the strong disorder regime, such Hamiltonians can be asymptotically diagonalized by Clifford unitaries using the spectrum bifurcation renormalization group approach[YQX16], which enables efficient classical post-processing of the classical shadow data and makes the Hamiltonian-driven shadow tomography computationally feasible. Machine learning techniques may also find useful application in the classical post-processing phase to construct unbiased reconstruction channels based on data statistics, which helps to mitigate the influence of noises and imperfections of NISQ devices. Finally, we would like to mention that our current analysis is limited to non-local Hamiltonians. How to include locality into the discussion will be a challenging problem for future research.

5.7 Acknowledgements

This work is supported by a startup fund from UC San Diego. Chapter 3, in full, is a reprint of the material as it appears in Hong-Ye Hu, Yi-Zhuang You, Physical Review Research 4, 013054 (2022). The dissertation author was the primary investigator and author of this paper.

Chapter 6

Classical shadow tomography with locally scrambled quantum dynamics

6.1 Introduction

Quantum state tomography[VR89, JKMW01, CFS02] is an essential task in many quantum technology applications. It seeks to reconstruct a quantum state from experimental data of repeated measurements. While reconstructing the full density matrix of a many-body system quickly becomes unfeasible with increasing system size due to the curse of dimensionality[OW16, HHJ$^+$17b], predicting a collection of (possibly exponentially many) properties of the quantum system can still be efficiently achieved with an only polynomial number of state copies, which was the idea of *shadow tomography* proposed by Aaronson[Aar17, AR19]. The idea is further improved by the recent work[HKP20] to propose the *classical shadow tomography*, which significantly reduces the demand on the quantum hardware and enables efficient classical post-processing.

Given a copy of an unknown quantum state ρ of N qubits, the classical shadow tomography protocol (see Fig. 6.1) first transforms the state $\rho \to \rho' = U \rho U^\dagger$ by a unitary U, which is randomly sampled (independently each time) from some probability distribution $P(U)$, then measures the transformed state ρ' in the computational basis, $\rho' \to |b\rangle\langle b|$, which collapses the system to a product state $|b\rangle$ labeled by a bit-string $b \in \{0,1\}^{\times N}$ of measurement outcomes $b = (b_1, \cdots, b_N)$ with the probability $P(b|\rho') = \langle b|\rho'|b\rangle$. Based on the observed bit-string b and the classical description of the unitary U, a classical snapshot $\hat{\sigma}_{U,b} = U^\dagger |b\rangle\langle b|U$ can be constructed in principle, which essentially encodes the measurement outcomes together with their basis choice (pulled back through the unitary evolution). Repeating such measurements on independent and identical copies of ρ for a few times, a collection of classical snapshots $\mathcal{E}_{\sigma|\rho} = \{\hat{\sigma}_{U,b}\}$ can be obtained (which correlates with ρ). Ref. [ONE13] showed that as long as the unitary ensemble is expressive enough (i.e. tomographically complete), there exist a linear reconstruction map $\mathcal{M}^{-1}$ such that the density matrix ρ can be formally recovered as $\rho = \mathbb{E}_{\hat{\sigma} \in \mathcal{E}_{\sigma|\rho}} \mathcal{M}^{-1}[\hat{\sigma}]$. This also enables the prediction of many properties of ρ, like the expectation value of any physical observable O as: $\langle O\rangle = \mathrm{Tr}(O\rho) = \mathbb{E}_{\hat{\sigma} \in \mathcal{E}_{\sigma|\rho}} \mathrm{Tr}(O\mathcal{M}^{-1}[\hat{\sigma}])$. The construction of classical snapshots $\hat{\sigma}_{U,b}$ and the computation of their associated properties are performed on a classical computer.

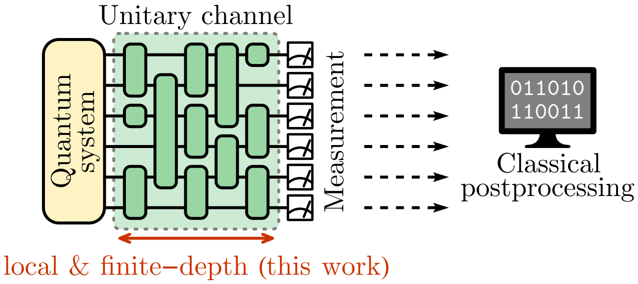

Figure 6.1: Illustration of classical shadow tomography protocol. This work focuses on the case when the unitary channel is of finite depth and respects locality.

However, the existing methods[Aar17, AR19, HKP20, HY21] have limitations in applying to near-term quantum devices. First, depending on the type of observables O that we are interested in, one needs to employ different strategies to design the unitary circuit U. Two limiting cases have been analyzed in Ref. [HKP20]: (i) if the observable is low-rank (such as many-body overlap fidelity), it is most efficient to adopt deep circuits, such that U effectively forms a global Haar random ensemble; (ii) if the observable is high-rank and quasi-local, it would be more efficient to adopt shallow circuits (e.g. the on-site Haar random). Otherwise the sample complexity will be high. However, the flexibility to interpolate between these two limits has not been available yet, such that the tomography protocol can not adjust to the target observables in a more adaptive manner. Second, more importantly, in existing quantum simulation platforms, applying random unitary circuits is very challenging, because it requires high degrees of sophisticated quantum controls. In particular, for programmable quantum simulators of large systems based on trapped ions or Rydberg atom systems,[EWL$^+$20, SSW$^+$20, ZPH$^+$17] a certain set of entangling unitary evolution is much more favorable to implement than typical random unitaries that require fine-tuned control. Therefore, it is desirable to develop a method applicable for systems with limited controls.

In this work, we address these challenges by generalizing the classical shadow tomography methods to a broad class of unitary ensembles. In our approach, the specific details of the unitary ensemble is not important as long as the ensemble generates *locally scrambled* quantum dynamics[KAAY20]. Rigorously speaking, the probability distribution $P(U)$ of evolution unitaries is invariant under local basis transformations, i.e. $\forall V \in \mathrm{U}(2)^N : P(UV) = P(U)$ where $V = \prod_i V_i$ is a product of local unitary operator V_i on each qubit. This basically means that the unitary evolution U is efficient in scrambling local quantum information, such that the initial local basis choice is quickly "forgotten" under the quantum dynamics. Examples of locally scrambled quantum dynamics includes random unitary circuits (including random Clifford circuit at the 3-design level)[NRVH17, ZN18, NVH18, CBQA19, BCA19, FVVY20] and quantum Brownian dynamics[LSH$^+$13, XS18, GHST18, ZC19, CZ19]. As the unitary ensemble does not care about local basis choice, the only information that matters will be the quantum entanglement that the unitary dynamics can create in the quantum system. Therefore, for locally scrambled quantum dynamics, the reconstruction map only depends on the entanglement property of the classical snapshots. The density matrix ρ can be reconstructed as a linear superposition of the classical snapshot $\hat{\sigma}$ reduced in different subsystems. The combination coefficient can be calculated from the entanglement feature[YYQ18b, YG18] of the classical snapshots, which is simply the collection of average purities of classical snapshots in all possible subregions.

Since our method is applicable to a broad class of quantum dynamics, it is natural to consider an ensemble of realistic Hamiltonian evolutions that are readily available in near-term quantum devices. To this end, we introduce an *approximate* classical shadow tomography (with a non-vanishing but small bias) applicable to an ensemble of time-dependent Hamiltonian evolution that generates approximately locally scrambled dynamics. We numerically demonstrate this idea by using a simple spin chain Hamiltonian modeled after programmable trapped ions or Rydberg atom array systems. We introduce the local frame potential to characterize the bias and we show the bias decreases rapidly for the initial short period of time, and reaches a vanishingly small plateau value for the proposed Hamiltonian. Surprisingly, we find even a single instance from an ensemble of Hamiltonian evolution suffices to perform an approximate tomography, implying that our method is hardware efficient for existing quantum devices[Saf16b, MCD^{+}21].

In the following, we will first establish the general theoretical framework to calculate the reconstruction map in Sec. 6.2.1 and to bound the sample complexity in Sec. 6.2.2. We also provide a two-qudit toy model to analytically demonstrate our construction in Sec. 6.2.3. We comment on how to carry out the computation efficiently in Sec. 6.2.4. Then we apply our construction for local unitary circuits and numerically demonstrates its accuracy in quantum fidelity and Pauli observable estimation tasks in Sec. 6.3.1, as well as their scaling of sample complexity in Sec. 6.3.2. Finally, we show in Sec. 6.3.3 that our approach can be extended to broader classes of unitary ensembles that are approximately locally scrambled. We propose a frame potential to characterize the level of approximation, which serves as a powerful indicator to design nearly-locally-scrambled unitary ensembles that are available for existing analog quantum simulators[Saf16b, MCD^{+}21]. We summarize our classical post-processing protocol and outline a few interesting future applications in Sec. 6.4

6.2 Theoretical Framework

6.2.1 Reconstruction Map from Entanglement Features

To be general, we consider a quantum system consists of N qudits, where each qudit has the Hilbert space dimension d (where $d = 2$ corresponds to the qubit system). The protocol of classical shadow tomography describes a process that first measures the unknown quantum state ρ in a random basis specified by the unitary transformation U and then prepare the classical snapshot $\hat{\sigma}_{U,b} \equiv U^{\dagger}|b\rangle\langle b|U$ based on the measurement outcome b. The randomness involved in the process includes (i) sampling U from the distribution $P(U)$ and (ii) obtaining the measurement outcome b conditioned on the evolved state $\rho' = U\rho U^{\dagger}$ with the probability $P(b|\rho') = \langle b|\rho'|b\rangle = \mathrm{Tr}(\hat{\sigma}_{U,b}\rho)$. As pointed out by Ref. [ASS21], the ensemble of classical snapshots forms a positive operator valued measure (POVM), defined as

$$\mathcal{E}_{\sigma|\rho} = \{\hat{\sigma}_{U,b} \mid P(\hat{\sigma}_{U,b}|\rho) = \mathrm{Tr}(\hat{\sigma}_{U,b}\rho)P(U)\}. \tag{6.1}$$

We will call $\mathcal{E}_{\sigma|\rho}$ the *posterior POVM*, as it is conditioned on the observation of ρ. The posterior POVM reduces to the *prior POVM*

$$\mathcal{E}_{\sigma} = \{\hat{\sigma}_{U,b} \mid P(\hat{\sigma}_{U,b}) = d^{-N}P(U)\}, \tag{6.2}$$

when there is no knowledge contained in ρ, i.e. $\rho = d^{-N}\mathbb{1}$. For the prior distribution $P(\hat{\sigma}_{U,b})$, the outcome b is uniformly drawn from all possible outcomes in $\{0, 1, \cdots, d-1\}^{\times N}$ (independent of U, ρ). The prior POVM $\mathcal{E}_{\sigma}$ only depends on the unitary ensemble $\mathcal{E}_U = \{U|P(U)\}$.

With the notation introduced above, the expected classical snapshot σ can be expressed as

$$\sigma \equiv \mathop{\mathbb{E}}_{\hat{\sigma}\in\mathcal{E}_{\sigma|\rho}} \hat{\sigma} = \mathop{\mathbb{E}}_{\hat{\sigma}\in\mathcal{E}_{\sigma}} \hat{\sigma}\operatorname{Tr}(\hat{\sigma}\rho)d^N = \mathcal{M}[\rho], \tag{6.3}$$

which is related to the original state ρ by a quantum channel $\mathcal{M}$, called the *measurement channel*. It is easy to check that the measurement channel $\mathcal{M}$ is trace-preserving, completely positive and self-adjoint. It is generally difficult to obtain an explicit expression of $\mathcal{M}$ for generic unitary ensemble $\mathcal{E}_U$ (or for generic prior POVM $\mathcal{E}_\sigma$). Results of $\mathcal{M}$ are known for global and on-site 2-design unitaries[ONE13, GKKT18, EVRZ19] (possibly with noise[EG20b, CYZF20b]), fermionic Gaussian unitaries[ZRM20], and many-body Gaussian unitaries[HY21].

We can make progress in computing the measurement channel $\mathcal{M}$ (and its inverse) for yet another class of unitary ensemble, namely the *locally scrambled unitaries*[KAAY20], for which $P(U)$ obeys the (right-)local-basis invariance condition[1]

$$\forall V \in \mathrm{U}(d)^N : P(U) = P(UV), \tag{6.4}$$

where $\mathrm{U}(d)^N$ is the group of all on-site unitary operators. This condition is sufficient to ensure the prior ensemble $\mathcal{E}_\sigma$ of snapshot states $\hat{\sigma}$ to be invariant under $\hat{\sigma} \to V^\dagger \hat{\sigma} V$,

$$\forall V \in \mathrm{U}(d)^N : P(\hat{\sigma}) = P(V^\dagger \hat{\sigma} V). \tag{6.5}$$

In this case, we say that $\mathcal{E}_\sigma$ is a *locally scrambled POVM*. In fact, our following derivation only requires the weaker condition Eq. (6.5) at the state level, instead of Eq. (6.4) at the channel level, though it will be practically more straight forward to design unitary circuits that satisfies Eq. (6.4) by assembling locally scrambled unitary gates.

---

[1] Strictly speaking, the locally scrambled condition originally introduced in Ref. [KAAY20] requires both the left- and the right-invariance, i.e. $P(VU) = P(UV) = P(U)$. However, for the purpose of this work, we only require a weaker condition that the distribution only need to be invariant under the right-multiplication of the local basis transformation V.

Nevertheless, as long as the states $\hat{\sigma}$ are locally scrambled (even if the unitaries U may or may not be locally scrambled), we will be able to insert local basis transformations V in Eq. (6.3), and average V over any ensemble of our choice,

$$\sigma = \mathop{\mathbb{E}}_{V\in \mathrm{U}(d)^N}\ \mathop{\mathbb{E}}_{\hat{\sigma}\in\mathcal{E}_\sigma} V^\dagger \hat{\sigma} V \,\mathrm{Tr}(V^\dagger \hat{\sigma} V \rho) d^N. \tag{6.6}$$

We can choose the ensemble of $V = \prod_i V_i$ to be such that every V_i is independently a local 2-design unitary. With this choice, the ensemble average of V can be evaluated by averaging every V_i over the Haar unitary measure following Ref. [Wei78, CŚ06], and the result can be written as (see Appendix B.1 for derivation)

$$\sigma = \sum_{B,C\in 2^{\Omega_N}} d^{2N-|B|} \rho_B \mathrm{Wg}_{B,C} W^{(2)}_{\mathcal{E}_\sigma,C}, \tag{6.7}$$

with B, C summing over all possible subregions of the N qudit system, where each subregion is labeled by a subset of $\Omega_N = \{1,\cdots,N\}$ (as an element in the power set 2^{Ω_N}). $|B|$ denotes the size (cardinality) of the region B. $\rho_B = (\mathrm{Tr}_{\bar{B}}\,\rho)\otimes(\mathbb{1}_{\bar{B}}/d^{|\bar{B}|})$ is the reduced density matrix of ρ in region B embedded back into the total Hilbert space. $\bar{B}$ denotes the complement of region B. Note that B and $\bar{B}$ do not need to be consecutive regions in the space, and they can intertwine with each other in general. $\mathrm{Wg}_{B,C} = (d^2-1)^{-N}(-1/d)^{|B\ominus C|}$ is the Weingarten function of regions B and C, where $B\ominus C$ denotes the subregions that belong to either B or C but not both.

$$W^{(2)}_{\mathcal{E}_\sigma,C} \equiv \mathop{\mathbb{E}}_{\hat{\sigma}\in\mathcal{E}_\sigma} \mathrm{Tr}_C(\mathrm{Tr}_{\bar{C}}\,\hat{\sigma})^2 = \mathop{\mathbb{E}}_{\hat{\sigma}\in\mathcal{E}_\sigma} e^{-S^{(2)}_C(\hat{\sigma})} \tag{6.8}$$

is the 2nd *entanglement feature*[YYQ18b, YG18] of the prior POVM $\mathcal{E}_\sigma$, where $S^{(2)}_C(\hat{\sigma})$ denotes the 2nd Rényi entanglement entropy of the state $\hat{\sigma}$ in region C. The entanglement feature $W^{(2)}_{\mathcal{E}_\sigma,C}$ is merely a property of the unitary ensemble $\mathcal{E}_U$ (which determines $\mathcal{E}_\sigma$). It describes how the unitary channel entangles a product state in general. It depends on neither the underlying state ρ

to be reconstructed nor any particular snapshot state $\hat{\sigma}$ collected in the tomography process.

Given the entanglement feature $W^{(2)}_{\mathcal{E}_\sigma,C}$, Eq. (6.7) spells out how the expected classical snapshot σ is written as a linear combination of reduced density matrices ρ_B in all regions, which explicitly specifies the measurement channel $\mathcal{M}$ as a linear map $\sigma = \mathcal{M}[\rho]$ from ρ to σ. Therefore, any reduced classical snapshot σ_A must also be a linear combination of reduced density matrices ρ_B, which implies that the measurement channel can be represented as a matrix $\mathcal{M}_{AB}$ such that $\sigma_A = \sum_B \mathcal{M}_{AB}\rho_B$. Suppose the map $\mathcal{M}$ is invertible (i.e. the unitary ensemble is tomographically complete), the inverse map $\mathcal{M}^{-1}$ (the reconstruction map) must also be a linear map that combines all reduced classical snapshots σ_A to reconstruct $\rho_B = \sum_A (\mathcal{M}^{-1})_{BA}\sigma_A$. In particular, we are most interested to reconstruct the full density matrix ρ (because all reduced density matrices follows from its partial trace), which must also be a linear combination of σ_A with some coefficients $r_A \in \mathbb{R}$,

$$\rho = \mathcal{M}^{-1}[\sigma] = d^N \sum_{A\in 2^{\Omega_N}} r_A \sigma_A, \tag{6.9}$$

where $\sigma_A = (\operatorname{Tr}_{\bar{A}}\sigma)\otimes(\mathbb{1}_{\bar{A}}/d^{|\bar{A}|})$ follows the same definition as the reduced density matrix. The reconstruction map $\mathcal{M}^{-1}$ is not a physical channel, because the reconstruction coefficients r_A may not be positive definite in general. Nevertheless, $\mathcal{M}^{-1}$ is still trace-preserving and self-adjoint. Since $\mathcal{M}^{-1}$ is linear, we have $\rho = \mathcal{M}^{-1}[\mathbb{E}_{\hat{\sigma}\in\mathcal{E}_{\sigma|\rho}}\hat{\sigma}] = \mathbb{E}_{\hat{\sigma}\in\mathcal{E}_{\sigma|\rho}}\mathcal{M}^{-1}[\hat{\sigma}]$, which enables us to reconstruct the underlying state ρ from the ensemble of classical snapshots. The collection of $\hat{\rho} = \mathcal{M}^{-1}[\hat{\sigma}]$ is also called the *classical shadow*[HKP20] of ρ, which can then be used to predict many properties of ρ efficiently.

Now the key problem is to compute r_A from $W^{(2)}_{\mathcal{E}_\sigma,C}$. For a system of N qudits, there will be 2^N many reconstruction coefficients r_A. To determine them, we substitute Eq. (6.7) to Eq. (6.9) and find

$$\rho = \sum_{A,B,C\in 2^{\Omega_N}} f_{A,B,C}\, r_A \rho_B W^{(2)}_{\mathcal{E}_\sigma,C}, \tag{6.10}$$

with the fusion coefficient $f_{A,B,C}$ given by

$$\begin{aligned} f_{A,B,C} &= \sum_{D\in 2^{\Omega_N}} \delta_{B,A\cap D} d^{2N+|A|-|B|+|\bar{A}\cap\bar{D}|} \mathsf{Wg}_{D,C} \\ &= \Big(\frac{d^3}{d^2-1}\Big)^N \sum_{D\in 2^{\Omega_N}} \delta_{B,A\cap D} d^{-|D|} \Big(-\frac{1}{d}\Big)^{|C\ominus D|}, \end{aligned} \tag{6.11}$$

which is universally determined by the qudit dimension d. Here $\delta_{A,B}$ denotes the Kronecker delta of two regions A and B, s.t. $\delta_{A,B} = 1$ (or 0) if $A = B$ (or $A \neq B$). Eq. (6.10) will hold for any choice of ρ if and only if

$$\sum_{A,C\in 2^{\Omega_N}} r_A f_{A,B,C} W^{(2)}_{\mathcal{E}_\sigma,C} = \delta_{B,\Omega_N}, \tag{6.12}$$

where $\Omega_N = \{1,\cdots,N\}$ is the full set that labels the full system of N qudits. By solving this linear equation, we can determine the reconstruction coefficients r_A in terms of of the entanglement feature $W^{(2)}_{\mathcal{E}_\sigma,C}$, such that the reconstruction map $\mathcal{M}^{-1}$ can be constructed according to Eq. (6.9).

In conclusion, we provide a general framework to compute the reconstruction map for the classical shadow tomography with locally scrambled quantum dynamics. The protocol is summarized as:

1. Given the prior POVM $\mathcal{E}_\sigma$, first calculate its entanglement feature by

$$W^{(2)}_{\mathcal{E}_\sigma,C} = \mathop{\mathbb{E}}_{\hat{\sigma}\in\mathcal{E}_\sigma} \mathrm{Tr}_C(\mathrm{Tr}_{\bar{C}}\,\hat{\sigma})^2.$$

2. Solve for the reconstruction coefficient r_A by

$$\sum_{A,C\in 2^{\Omega_N}} r_A f_{A,B,C} W^{(2)}_{\mathcal{E}_\sigma,C} = \delta_{B,\Omega_N}.$$

3. Then the reconstruction map is given by

$$\rho = \mathcal{M}^{-1}[\sigma] = d^N \sum_{A\in 2^{\Omega_N}} r_A \sigma_A.$$

All computations are supposed to be carried out on a classical computer in the post-processing procedure. Although solving for r_A may be computationally demanding for large systems, it only needed to be done once and its result can be applied to process all classical snapshots collected from all possible states ρ to be learned.

6.2.2 Variance Estimation and Sample Complexity

Given the ensemble $\mathcal{E}_{\sigma|\rho}$ of classical snapshots collected from measuring the unknown state ρ, we can use the reconstruction map $\mathcal{M}^{-1}$ to predict properties of ρ. For example, let O be a traceless Hermitian operator representing a physical observable. Its expectation value $\langle O\rangle \equiv \mathrm{Tr}(O\rho)$ can be predicted via

$$\langle O\rangle = \mathop{\mathbb{E}}_{\hat{\sigma}\in\mathcal{E}_{\sigma|\rho}} \mathrm{Tr}(O\mathcal{M}^{-1}[\hat{\sigma}]) = \mathop{\mathbb{E}}_{\hat{\sigma}\in\mathcal{E}_{\sigma|\rho}} \mathrm{Tr}(\mathcal{M}^{-1}[O]\hat{\sigma}), \tag{6.13}$$

where we have used the self-adjoint property of $\mathcal{M}^{-1}$ to transpose its action from $\hat{\sigma}$ to O. We can interpret $\hat{o}(\hat{\sigma}) \equiv \operatorname{Tr}(\mathcal{M}^{-1}[O]\hat{\sigma})$ as the single-shot estimation of the observable (based on a particular classical snapshot $\hat{\sigma}$), such that $\langle O \rangle = \mathbb{E}_{\hat{\sigma}\in\mathcal{E}_{\sigma|\rho}}\, \hat{o}(\hat{\sigma})$.

The variance of the single-shot estimation is defined as

$$\operatorname{Var}\hat{o} \equiv \mathop{\mathbb{E}}_{\hat{\sigma}\in\mathcal{E}_{\sigma|\rho}} \hat{o}(\hat{\sigma})^2 - (\mathop{\mathbb{E}}_{\hat{\sigma}\in\mathcal{E}_{\sigma|\rho}} \hat{o}(\hat{\sigma}))^2,$$

which can be bounded by (the first term in $\operatorname{Var}\hat{o}$)

$$\begin{aligned} \operatorname{Var}\hat{o} \leq \|O\|^2_{\mathcal{E}_{\sigma|\rho}} &\equiv \mathop{\mathbb{E}}_{\hat{\sigma}\in\mathcal{E}_{\sigma|\rho}} \hat{o}(\hat{\sigma})^2 \\ &= \mathop{\mathbb{E}}_{\hat{\sigma}\in\mathcal{E}_{\sigma}} (\operatorname{Tr}\mathcal{M}^{-1}[O]\hat{\sigma})^2 \operatorname{Tr}(\hat{\sigma}\rho) d^N. \end{aligned} \tag{6.14}$$

The bound $\|O\|_{\mathcal{E}_{\sigma|\rho}}$ can be considered as a generalized ρ-dependent notion of the (squared) *shadow norm*[HKP20] of an operator O (whereas the shadow norm originally defined in Ref. [HKP20] further maximizes over all possible underlying states ρ to remove the dependence on ρ). Assuming $\mathcal{E}_\sigma$ is locally scrambled, following the same approach of inserting and averaging local-basis transformations as in Eq. (6.6), the bound in Eq. (6.14) becomes

$$\|O\|^2_{\mathcal{E}_{\sigma|\rho}} = \sum_{g,h\in S_3^N} \|O\|^2_{\rho,g} \mathsf{Wg}_{g,h} W^{(3)}_{\mathcal{E}_\sigma,h}, \tag{6.15}$$

where g, h are group elements in the S_3^N (product of 3-fold permutation groups over N qudits), $\mathsf{Wg}_{g,h}$ is the Weingarten function of permutations g and h. $\|O\|^2_{\rho,g}$ is a generalized operator norm for O, which is defined as

$$\|O\|^2_{\rho,g} \equiv d^N \operatorname{Tr}((\mathcal{M}^{-1}[O]^{\otimes 2} \otimes \rho)\chi_g), \tag{6.16}$$

where χ_g is the representation of the S_3^N permutation g in the 3-fold Hilbert space. $W^{(3)}_{\mathcal{E}_\sigma,h}$ is the

3rd *entanglement feature* of the ensemble $\mathcal{E}_\sigma$, defined as

$$W^{(3)}_{\mathcal{E}_\sigma,h} \equiv \mathop{\mathbb{E}}_{\hat{\sigma}\in\mathcal{E}_\sigma} \operatorname{Tr}(\hat{\sigma}^{\otimes 3}\chi_h). \tag{6.17}$$

Note that the 2nd entanglement feature previously defined in Eq. (6.8) can be consistently cast into the form of Eq. (6.17) in terms of permutation operators (see Ref. [YG18]).

In practice, the expectation value $\langle O\rangle$ is always estimated based on a finite collection of the snapshot states. Let M be the number of samples of ρ measured in the data acquisition stage (each sample results in a snapshot state $\hat{\sigma}_k$). The finite average estimation $\bar{o} = \frac{1}{M}\sum_{k=1}^{M}\hat{o}(\hat{\sigma}_k)$ will fluctuate around the true expectation value $\langle O\rangle$ with a variance that scales as $(\operatorname{Var}\hat{o})/M$. By the Chebyshev inequality, the probability for $\bar{o}$ to deviate from $\langle O\rangle$ by more than ϵ amount is bounded by

$$\Pr(|\bar{o}-\langle O\rangle| \geq \epsilon) \leq \frac{\operatorname{Var}\hat{o}}{\epsilon^2 M} \leq \frac{\|O\|^2_{\mathcal{E}_{\sigma|\rho}}}{\epsilon^2 M}. \tag{6.18}$$

Therefore, to control the failure probability within a threshold δ, i.e. $\Pr(|\bar{o}-\langle O\rangle| \geq \epsilon) \leq \delta$, sufficient number of samples is required

$$M \geq \frac{\|O\|^2_{\mathcal{E}_{\sigma|\rho}}}{\epsilon^2\delta}. \tag{6.19}$$

A larger (smaller) shadow norm $\|O\|^2_{\mathcal{E}_{\sigma|\rho}}$ indicates that more (less) samples are needed.

However, the ρ-dependent shadow norm $\|O\|^2_{\mathcal{E}_{\sigma|\rho}}$ is generally complicated to evaluate. If we are not interested in the shadow norm for a specific state ρ, but rather the expectation of the shadow norm over an ensemble of states $\{V\rho V^\dagger\}$ that are similar to ρ by local basis transformations $V \in \mathrm{U}(d)^N$, we can actually define a ρ-independent shadow norm by averaging over V,

$$\|O\|^2_{\mathcal{E}_\sigma} \equiv \mathop{\mathbb{E}}_{V\in\mathrm{U}(d)^N} \|O\|^2_{\mathcal{E}_{\sigma|V\rho V^\dagger}}. \tag{6.20}$$

The expected shadow norm can be expressed purely in terms of the entanglement features of $\mathcal{E}_\sigma$ and $\mathcal{E}_O$ (see Appendix B.2 for derivation),

$$\|O\|^2_{\mathcal{E}_\sigma} = \sum_{A,B,C,D\in 2^{\Omega_N}} v_{A,B,C,D} W^{(2)}_{\mathcal{E}_\sigma, A\cap B\cap C} W^{(2)}_{\mathcal{E}_O,D}, \tag{6.21}$$

where the coefficient $v_{A,B,C,D}$ is given by

$$v_{A,B,C,D} = r_A r_B \Big(\frac{d^2}{d^2-1}\Big)^N d^{|A\cap B\cap C|-|C|}\Big(-\frac{1}{d}\Big)^{|C\ominus D|}, \tag{6.22}$$

and $\mathcal{E}_O = \{V^\dagger O V | V \in \mathrm{U}(d)^N\}$ denotes the locally scrambled ensemble (or known as $\mathrm{U}(d)^N$-twirling) associated with the observable O in question.

In conclusion, given a traceless Hermitian operator O, its expected shadow norm $\|O\|^2_{\mathcal{E}_\sigma}$ provides a typical lower bound for the number of samples needed

$$M \gtrsim \frac{\|O\|^2_{\mathcal{E}_\sigma}}{\epsilon^2 \delta}, \tag{6.23}$$

in order to control the error of the prediction $\bar{o}$ given by the classical shadow tomography within the probability bound $\Pr(|\bar{o} - \langle O\rangle| \geq \epsilon) \leq \delta$. Here we have only analyzed the sample complexity for a single linear observable. For the analysis of multiple and/or non-linear observables, we refer to the original paper of Ref. [HKP20]. Their result applies to our case simply by replacing the shadow norm with our version.

6.2.3 A Toy Example of Two-Qudit System

To demonstrate our framework and to gain some analytical intuition, we present a toy example to compute the reconstruction map in a two-qudit ($N = 2$) system. We assume that the two-qudit system always evolves under a locally scrambled quantum dynamics, which can be modeled (for example) by a finite-time Brownian evolution[2] driven by random Hamiltonians. Every classical snapshot $\hat{\sigma}_{U,b} = U^\dagger|b\rangle\langle b|U$ is generated by the reversed evolution from the product state $|b\rangle\langle b|$. In the long-time limit (Fig. 6.2(a)), the entanglement feature $W^{(2)}_{\mathcal{E}_\sigma} = (1, \frac{2d}{d^2+1}, \frac{2d}{d^2+1}, 1)$ follows from that of Page states, where the subregion basis are arranged in the order of $\{\}, \{1\}, \{2\}, \{1,2\}$. This is because the evolution of entanglement feature under any locally scrambled quantum dynamics always converges to the Page state, regardless of the initial state, as proven in Ref. [KAAY20]. In the short-time limit (Fig. 6.2(b)), $\hat{\sigma}$ remains as a product state, therefore the entanglement entropy vanishes for all regions, which translates to $W^{(2)}_{\mathcal{E}_\sigma} = (1,1,1,1)$. In general, for any intermediate time, the entanglement feature should take the form of

$$W^{(2)}_{\mathcal{E}_\sigma} = (1, w, w, 1), \tag{6.24}$$

with w varies between $\frac{2d}{d^2+1}$ (the long-time limit) and 1 (the short-time limit). The physical meaning of w is the average single-qudit purity in the snapshot state $\hat{\sigma}$.

Given $W^{(2)}_{\mathcal{E}_\sigma}$ in Eq. (6.24), Eq. (6.12) reads

$$d^2 \begin{bmatrix} 1 & \frac{d(d-w)}{d^2-1} & \frac{d(d-w)}{d^2-1} & \frac{d^2(d^2-2dw+1)}{(d^2-1)^2} \\ 0 & \frac{dw-1}{d^2-1} & 0 & \frac{d(d^2w-2d+w)}{(d^2-1)^2} \\ 0 & 0 & \frac{dw-1}{d^2-1} & \frac{d(d^2w-2d+w)}{(d^2-1)^2} \\ 0 & 0 & 0 & \frac{d^2-2dw+1}{(d^2-1)^2} \end{bmatrix} \begin{bmatrix} r_{\{\}} \\ r_{\{1\}} \\ r_{\{2\}} \\ r_{\{1,2\}} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \end{bmatrix}. \tag{6.25}$$

---

[2] The Brownian unitary evolution is a product of a sequence of infinitesimal time-evolution $U = \prod_t e^{-\mathrm{i}H_t\,\delta t}$, but the Hamiltonian H_t at each time step is independent drawn from a random Hamiltonian ensemble (unlike the coherent quantum dynamics, where the same Hamiltonian drives the dynamics though all time.)

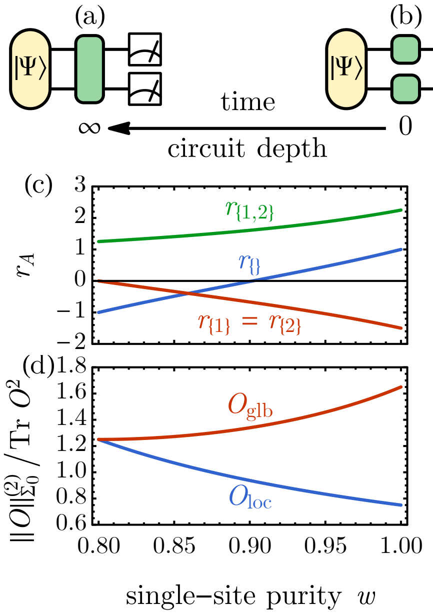

Figure 6.2: Two-qudit unitary channel in (a) the long-time (Page state) limit and (b) the short-time (product state) limit. (c) Reconstruction coefficients r_A and (d) the shadow norm $\|O\|^2_{\mathcal{E}_\sigma}$ v.s. the single-qudit purity w, for $d = 2$. w varying from 1 to $4/5$ effectively models the circuit depth (or evolution time) growing from 0 to ∞.

By solving this linear equation, the reconstruction coefficient r_A can be obtained

$$r=\begin{bmatrix} r_{\{\}} \\ r_{\{1\}} \\ r_{\{2\}} \\ r_{\{1,2\}} \end{bmatrix}=\begin{bmatrix} \frac{d^3w-3d^2+3dw-2w^2+1}{(dw-1)(d^2-2dw+1)} \\ \frac{-d^4w+2d^3-2d+w}{d(dw-1)(d^2-2dw+1)} \\ \frac{-d^4w+2d^3-2d+w}{d(dw-1)(d^2-2dw+1)} \\ \frac{(d^2-1)^2}{d^2(d^2-2dw+1)} \end{bmatrix}. \tag{6.26}$$

The behavior of r_A as a function of w is shown in Fig. 6.2(c), which continuously interpolates the two limits.

In the short-time limit, $w=1$ and Eq. (6.26) reduces to $r_A=(1,-(d+1)/d,-(d+1)/d,(d+1)^2/d^2)$, corresponding to the reconstruction map

$$\mathcal{M}^{-1}[\sigma]=\bigotimes_{i=1,2}((d+1)\sigma_i-\mathbb{1}_i), \tag{6.27}$$

matching the result of on-site 2-design circuits[GKKT18, EVRZ19]. In the long-time limit, $w=\frac{2d}{d^2+1}$ and Eq. (6.26) reduces to $r_A=(-1,0,0,(d^2+1)/d^2)$, corresponding to the reconstruction map

$$\mathcal{M}^{-1}[\sigma]=(d^2+1)\sigma-\mathbb{1}, \tag{6.28}$$

matching the result of global 2-design circuits[GKKT18, EVRZ19]. The general result in Eq. (6.26) provides the reconstruction map that interpolates these two limits, which allows us to perform classical shadow tomography for intermediate unitary channels that are neither on-site nor global 2-design.

To investigate the sample complexity of the tomography scheme in the two-qudit system, we consider a traceless Hermitian operator O (i.e. $\operatorname{Tr} O = 0$) and assume:

$$\begin{aligned} \operatorname{Tr}_{\{1\}}(\operatorname{Tr}_{\{2\}} O)^2 &= k_1 \operatorname{Tr} O^2 \\ \operatorname{Tr}_{\{2\}}(\operatorname{Tr}_{\{1\}} O)^2 &= k_2 \operatorname{Tr} O^2, \end{aligned} \tag{6.29}$$

then the entanglement feature of the observable O can be arranged as the following vector

$$W^{(2)}_{\mathcal{E}_O} = (0, k_1, k_2, 1) \operatorname{Tr} O^2, \tag{6.30}$$

with the same choice of region basis as in Eq. (6.24). Given r, $W^{(2)}_{\mathcal{E}_O}$ and $W^{(2)}_{\mathcal{E}_\sigma}$, we have all the information needed to calculate the shadow norm, according to Eq. (6.21)

$$\|O\|^2_{\mathcal{E}_\sigma} = \tfrac{d^2-1}{d^3}\left(\tfrac{k_{\text{tot}}}{dw-1} + \tfrac{(d^2-1)(d-k_{\text{tot}})}{d^2-2dw+1}\right) \operatorname{Tr} O^2, \tag{6.31}$$

where $k_{\text{tot}} = k_1 + k_2$.

The operator locality crucially affects k_{tot}. Consider modeling a local operator O_{loc} by a random operator drawn from the Gaussian unitary ensemble (GUE) and acting on the first qudit only, we have

$$W^{(2)}_{\mathcal{E}_{O_{\text{loc}}}} = (0, d, 0, 1) \operatorname{Tr} O^2_{\text{loc}}, \tag{6.32}$$

hence $k_{\text{tot}} = d$. On the other hand, for a global operator O_{glb} modeled by a global GUE random operator acting on both qudits simultaneously, we have

$$W^{(2)}_{\mathcal{E}_{O_{\text{glb}}}} = (0, \tfrac{d}{d^2+1}, \tfrac{d}{d^2+1}, 1) \operatorname{Tr} O^2_{\text{glb}}, \tag{6.33}$$

hence $k_{\text{tot}} = \frac{2d}{d^2+1}$. In these two cases, the shadow norm in Eq. (6.31) becomes

$$\begin{aligned}\|O_{\text{loc}}\|^2_{\mathcal{E}_\sigma} &= \tfrac{d^2-1}{d^2(dw-1)} \operatorname{Tr} O^2_{\text{loc}},\\ \|O_{\text{glb}}\|^2_{\mathcal{E}_\sigma} &= \tfrac{d^2-1}{d^2(d^2+1)} \big(\tfrac{(d^2-1)^2}{d^2-2dw-1} + \tfrac{2}{dw-1}\big) \operatorname{Tr} O^2_{\text{glb}}.\end{aligned} \tag{6.34}$$

Their dependence in w is plotted in Fig. 6.2(d). In the short-time limit ($w = 1$), $\|O_{\text{loc}}\|^2_{\mathcal{E}_\sigma} < \|O_{\text{glb}}\|^2_{\mathcal{E}_\sigma}$, meaning that the shallow circuit is more efficient in predicting local observables. In the long-time limit ($w = \frac{2d}{d^2+1}$), $\|O_{\text{loc}}\|^2_{\mathcal{E}_\sigma} = \|O_{\text{glb}}\|^2_{\mathcal{E}_\sigma} = (1 + d^{-2}) \operatorname{Tr} O^2$, such that there is no difference in predicting both local and global observables in terms of the sample efficiency, because all operators are equally scrambled in this limit.

6.2.4 Additional Remarks on Computational Methods and Future Directions

Efficient numerical methods have been developed[FVVY20, AY20] to calculate the evolution of entanglement feature $W^{(k)}_{\mathcal{E}_\sigma}$ under locally scrambled quantum dynamics by solving the corresponding entanglement dynamics equation (without simulating the quantum dynamics using brute force). However, we will leave this approach for future exploration. In this work, we will compute the entanglement feature beforehand based on the definition Eq. (6.8), by direct sampling from the prior POVM $\mathcal{E}_\sigma$. For experimentally generated random unitaries whose distribution is a priori unknown, it is also possible to estimate the entanglement feature efficiently from Rényi entropy measurements[EVD$^+$18, VED$^+$18, BEJ$^+$19] following the definition Eq. (6.8).

As shown in Ref. [AY20], for one-dimensional quantum systems, the entanglement feature vector $W^{(k)}_{\mathcal{E}_\sigma}$ admits efficient matrix product state (MPS) representation, even if snapshot states in $\mathcal{E}_\sigma$ are volume-law entangled. Combining the MPS representation of $W^{(2)}_{\mathcal{E}_\sigma}$ with the fact that $f_{A,B,C}$ is factorizable to every qudit, one can develop efficient MPS-based numerical approach to find the solution of r_A (also as a MPS). However, we will defer the development of this approach to future work. In the following numerical demonstrations, we will directly solve Eq. (6.12) for small systems as a proof of concept.

The MPS representations for r, $W^{(2)}_{\mathcal{E}_\sigma}$ and $W^{(2)}_{\mathcal{E}_O}$ also enables us to calculate the shadow norm $\|O\|^2_{\mathcal{E}_\sigma}$ efficiently by a four-way MPS contraction. The ability to compute the shadow norm efficiently will be particularly useful if we want to design optimal unitary channels to minimize the sample complexity for a given set of designated observables. It possible to apply machine learning approaches (such as deep reinforcement learning) to perform the circuit structure optimization. Therefore, our construction provides the flexibility to allow the classical shadow tomography to adapt to designated observables, which has not been possible before. We will also leave this promising direction to future research.

6.3 Numerical Demonstrations

To demonstrate the effectiveness of our approach, we consider two types of unitary ensembles for the unitary channel in the data acquisition protocol, as illustrated in Fig. 6.3.

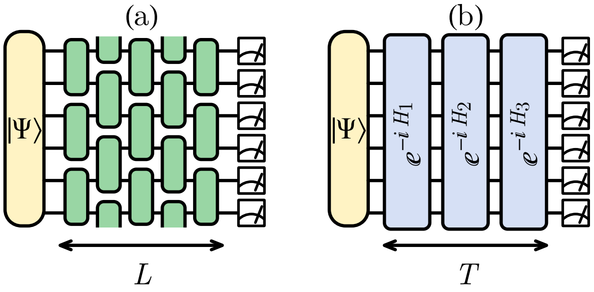

Figure 6.3: Classical shadow tomography with (a) finite-depth random unitary circuits (of L layers), (b) discrete-time Hamiltonian dynamics (of T steps).

6.3.1 Classical Shadow Tomography with Shallow Random Unitary Circuits

We first consider using random unitary circuits (RUCs) [NRVH17] for the unitary channel. As illustrated in Fig. 6.3(a), the unitary circuit consists of two-qubit local unitary gates arranged in the brick-wall pattern with a periodic boundary condition. Each gate in the circuit is independently drawn from the Haar random unitary ensemble. The depth L of the circuit can be adjusted. Obviously, RUCs are locally scrambled, as any local-basis transformation (from both left and right) can be absorbed by the Haar random unitary gates in the circuit. Therefore, we expect our reconstruction map to work perfectly in this case for any choice of the circuit depth L.

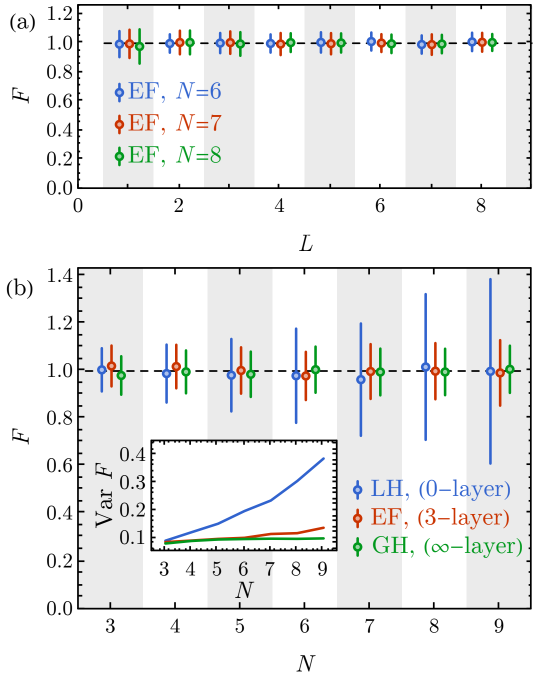

Figure 6.4: (a) Fidelity estimation of GHZ state with RUC of different circuit depth L using entanglement-feature-based reconstruction $\mathcal{M}_{\text{EF}}^{-1}$ (denoted by EF) over different number N of qubits. (b) Fidelity estimation of GHZ state using shallow RUC (3-layer, with $\mathcal{M}_{\text{EF}}^{-1}$, denoted by EF), random on-site (local Haar) gates (0-layer, with $\mathcal{M}_{\text{LH}}^{-1}$, denoted by LH) and global Haar unitary (∞-layer, with $\mathcal{M}_{\text{GH}}^{-1}$, denoted by GH). The inset shows the variance $\text{Var}\,F$ of the predicted fidelity as a function of system size N. In both subfigures, the sample size is 5000. Error bar indicates 3-standard-deviation estimated by the bootstrap method. Points are split horizontally to avoid the overlap of markers.

For illustration purpose, we start with a Greenberger-Horne-Zeilinger (GHZ) state $\rho = |\Psi\rangle\langle\Psi|$, where $|\Psi\rangle = \frac{1}{\sqrt{2}}(|00\cdots0\rangle + |11\cdots1\rangle)$. For every given circuit depth L, we first calculate the entanglement feature $W^{(2)}_{\mathcal{E}_\sigma,C}$ to determine the reconstruction map $\mathcal{M}^{-1}$. This calculation is done for once and stored in the classical memory for future reference. In our numerical simulation of the data acquisition process, we sample the RUC, apply it to the GHZ state $|\Psi\rangle$, and perform the computational basis measurement. We generate a collection of classical snapshots $\mathcal{E}_{\sigma|\rho} = \{\hat{\sigma}\}$ of size M by repeated measurements. We then estimate the fidelity F of the reconstructed state by

$$F = \sqrt{\frac{1}{M} \sum_{\hat{\sigma} \in \mathcal{E}_{\sigma|\rho}} \langle\Psi|\mathcal{M}^{-1}[\hat{\sigma}]|\Psi\rangle}. \tag{6.35}$$

Following the philosophy of classical shadow tomography, one should view Eq. (6.35) as a prediction task. If the shadow tomography is successful, then the estimated fidelity should converge to $F = 1$. This estimation can be achieved accurately by a few measurements, even though the full density matrix estimation $\operatorname{avg}_{\hat{\sigma}\in\mathcal{E}_{\sigma|\rho}} \mathcal{M}^{-1}[\hat{\sigma}]$ may still have large fluctuations. In addition, when the reconstruction is biased, for example the experimental channel $\mathcal{M}$ doesn't match the theoretical assumption of the unitary ensemble, then the fidelity estimation will deviate from one. Fig. 6.4 (a) shows that the entanglement-feature-based reconstruction map $\mathcal{M}^{-1}_{\text{EF}}$ indeed gives unbiased estimation of fidelity F for different circuit depths L and for different system sizes N. Furthermore, when the GHZ state is prepared with Z errors, our method can give the correct fidelity estimation that decreases linearly with the probability of Z error, which is challenging for the current state-of-art machine-learning quantum state tomography method[CTMA19, HKP20] (see Appendix B.3 for more discussions).

To compare with the existing classical shadow tomography method[HKP20], we consider the reconstruction maps $\mathcal{M}^{-1}_{\text{GH}}[\sigma] = (d^N + 1)\sigma - \mathbb{1}$ and $\mathcal{M}^{-1}_{\text{LH}}[\sigma] = \bigotimes_i((d+1)\sigma_i - \mathbb{1}_i)$, where $\mathcal{M}^{-1}_{\text{GH}}$ (or $\mathcal{M}^{-1}_{\text{LH}}$) assumes the unitary ensemble is global (or on-site local) Haar random. They can be viewed as special limits where circuit depth L tends to infinity and zero respectively. Although

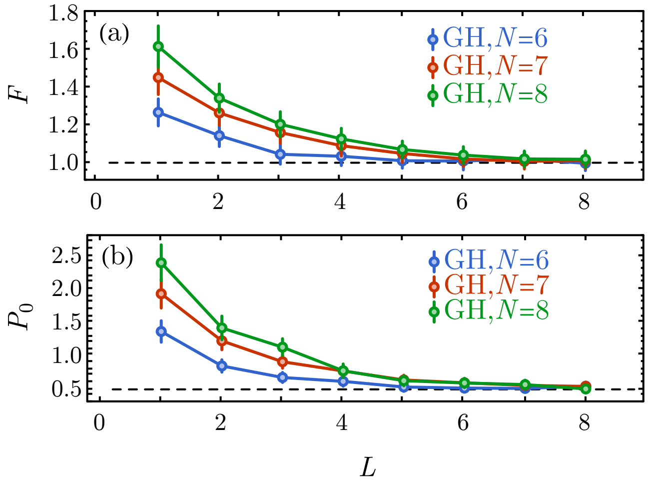

Figure 6.5: (a) Fidelity estimation of the reconstructed GHZ state with RUC of finite depth L. (b) Estimation of observable $P_0 = |\langle\Psi|00\cdots0\rangle|^2$ (the projection operator to the $|00\cdots0\rangle$ state) on the reconstructed GHZ state with RUC of finite depth L. In both cases, the reconstruction uses the global Haar reconstruction map. The sample size is 5000. Error bar indicates 3-standard-deviation estimated by the bootstrap method.

they can also achieve an unbiased estimation of quantum fidelity, the tomography efficiency differs. In Fig. 6.4(b), the error bar shows how the (3-times) standard deviation of the estimated fidelity scales with the number of qubits N at 5000 sample size. As we can see (both from the error bar and from the inset of Fig. 6.4(b)), the variance of (on-site) local Haar estimation increases drastically as N increases, which implies an increasingly high sample complexity for large systems.

In the other limit, the variance of global Haar estimation is independent of system size, achieving the optimal sample complexity as advocated in Ref. [HKP20]. However, to realize the global Haar ensemble, the circuit depth needs to be at least of order $O(N)$, which is quite demanding for quantum devices. If we approximate the global Haar ensemble with finite-depth circuits and use the reconstruction map $\mathcal{M}_{\text{GH}}^{-1}$ on data collected from finite-depth circuit measurements,[ONE13] this will yield systematically biased predictions for physical quantities when the circuit is not deep enough. In Fig. 6.5(a), we show that the biased prediction tends to over-estimate the fidelity, leading to the unphysical result of $F > 1$ (the correct behavior is $F = 1$).

This occurs because, when the measurement channel $\mathcal{M}$ in data acquisition protocol disagrees with the reconstruction channel $\mathcal{M}^{-1}$ in classical post-processing protocol, the reconstructed density matrix $\frac{1}{M}\sum_{\hat{\sigma}\in\mathcal{E}_{\sigma|\rho}}\mathcal{M}^{-1}[\hat{\sigma}]$ may not be positive-definite (see Appendix B.4 for detailed discussions), resulting in the unphysical fidelity estimation. This bias gets worse for larger system size. In Fig. 6.5(b), we also show the estimation of $P_0 = |\langle\Psi|00\cdots 0\rangle|^2$. For GHZ state, the correct behavior is $P_0 = 0.5$, and we still see significant bias when applying $\mathcal{M}_{\text{GH}}^{-1}$ for shallow circuits.

However, with the entanglement-feature-based reconstruction map $\mathcal{M}_{\text{EF}}^{-1}$, as demonstrated in Fig. 6.4 (b), we are able to achieve an *unbiased* fidelity estimation with a 3-layer shallow circuit, approaching *similar variance level* (i.e. similar sample efficiency) as global Haar ensemble while keeping a *low circuit complexity*. This clearly demonstrates the advantage of our approach.

6.3.2 Scaling of Variance and Tomography Complexity

The above discussion motivate us to define the *tomography complexity* as $\mathcal{C} = (L+1)M$, where L is the *circuit complexity* (the number of layers in the quantum circuit), and M is the *sample complexity* (the number of sample needed). M will be proportional to the single-shot variance $\operatorname{Var}\hat{o}$. Suppose applying each layer of quantum gates and performing measurements both take a unit of time on the quantum device, then $\mathcal{C}$ is roughly the total amount of time needed to collect the classical shadow from M copies of the quantum state, which characterizes the complexity of the data acquisition protocol. This notion of complexity is consistent with the Quantum Algorithmic Measurement (QUALM) complexity introduced in Ref. [ACQ21] (with LM and M being their gate and query complexities respectively). In the following, we will investigate the scaling of single-shot variance $\operatorname{Var}\hat{o}$ as a function of circuit depth L and system size N for both low-rank operators (such as fidelity) and full-rank operators (such as Pauli operators), and show how the tomography complexity $\mathcal{C}$ can guide us to find the optimal circuit depth L_*.

For low-rank operators, we will focus on quantum fidelity, which is important in many quantum information applications, such as (variational) state preparation. We will define the

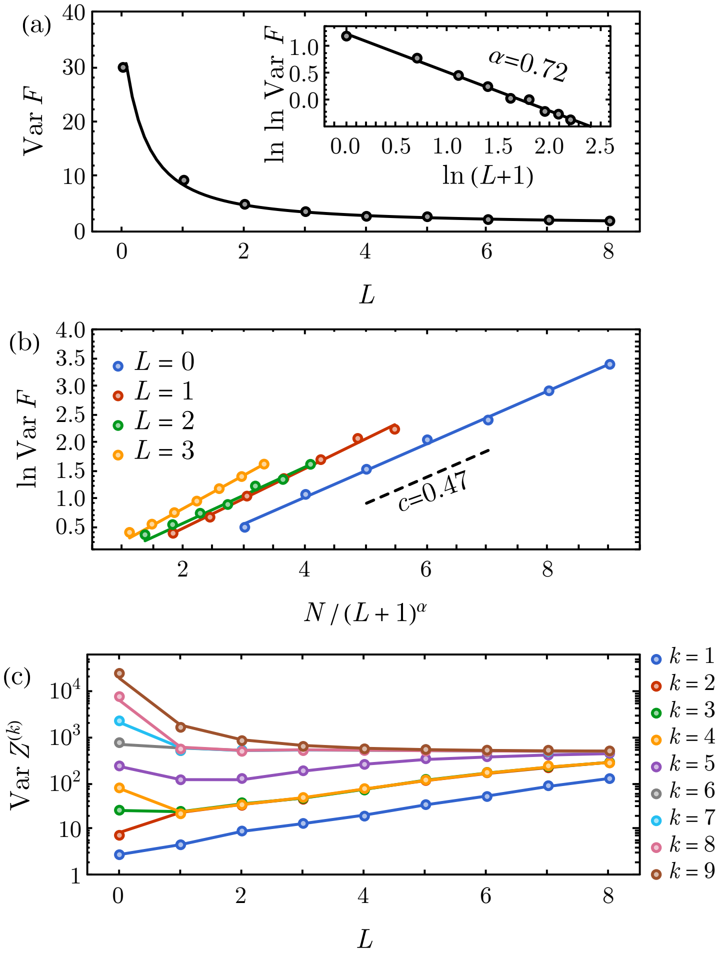

Figure 6.6: (a) Single-shot variance of estimated fidelity v.s. circuit depth L for a 9-qubits GHZ state. (b) Single-shot variance of estimated fidelity as a function of the effective system size N_{eff}. The best fit for $\mathrm{Var}\,F \propto \exp\left(c\frac{N}{(L+1)^{\alpha}}\right)$ gives $c = 0.47 \pm 0.08$ and $\alpha = 0.72 \pm 0.1$. (c) Variance of full rank operator estimation on 9-qubit GHZ state. The full rank operators are Pauli-Z operators of the form: $Z^{(k)} = Z^{\otimes k} I^{\otimes (N-k)}$ with different support k. The dots are experimental results from simulation, and the lines are theoretical prediction using operator shallow norm by Eq. (6.21). They match perfectly.

zero-depth limit ($L \to 0$) of the RUC to be a single layer of on-site Haar-random gates because even if there is no two-qubit gate in the "zero-depth" circuit, we still assume that the unitary ensemble is locally scrambled such that on-site scrambling unitaries continue to persist. In this limit, the single-shot variance $\mathrm{Var}F$ of fidelity estimation scales exponentially with the number of qubits N. On the other hand, in the deep circuit limit ($L \to \infty$), RUCs will approach the global Haar unitary ensemble, and the variance $\mathrm{Var}F$ will be independent of system size. We are interested to investigate how $\mathrm{Var}F$ behaves in the shallow circuit regime. In Fig. 6.6(a), we calculated $\mathrm{Var}F$ numerically using the bootstrap method for the 9-qubit GHZ state. It shows the variance $\mathrm{Var}F$ will decrease quickly in the shallow circuit regime. Interestingly, we found that an empirical formula $\mathrm{Var}F \propto \exp\left(c\frac{N}{(L+1)^\alpha}\right)$ fits the data well in the shallow circuit regime, with $\alpha = 0.7 \pm 0.1$. In Fig. 6.6(b), we plot $\mathrm{Var}F$ as a function of $\frac{N}{(L+1)^\alpha}$ for different fixed circuit depth L. We find curves with different choices of circuit depth L all collapse together with the same coefficient $c = 0.47 \pm 0.08$.

The physical intuition behind the empirical formula has to do with the operator growth in RUCs. If the quantum circuit is very shallow, then the computational basis measurement will only probe local information in the original basis. If the circuit becomes deeper, computational basis measurement can probe information in larger regions in the original basis, because the measurement operator has grown under the (backward) circuit evolution. Suppose the size of the measurement operator grows in a power-law manner $\sim (L+1)^{\alpha}$[3] with respect to the depth L of the RUC, the relative size of the system will effectively shrink to $N_{\text{eff}} = \frac{N}{(L+1)^\alpha}$, such that $\mathrm{Var}F$ should scale universally with N_{eff}, as proposed in the empirical formula. We might expect $\alpha = 1/2$ (or $\alpha = 1$), if the operator grew diffusively (or ballisticaly). However, the best fit of our numerical result seems to indicate an effective operator growth between the diffusive and ballistic limits. Due to the limited system size in this study, we are unable to determine whether our observation persists to the thermodynamic limit. We will leave this intriguing scaling behavior

---

[3] In the $L \to 0$ limit, the measurement operator is still of at least size 1, which motivates the "$+1$" regularization in $(L+1)^\alpha$.

for further investigation in the future. Nevertheless, for any α, the variance decreases faster than exponential with L in the shallow circuit regime, which already speaks for the advantage of applying shallow circuits in classical shadow tomography.

For full-rank operators, we mainly focus on consecutive strings of Pauli operators of the form

$$Z^{(k)} = Z^{\otimes k} I^{\otimes (N-k)} = \underbrace{ZZ\cdots Z}_{k}\underbrace{II\cdots I}_{N-k}, \tag{6.36}$$

where Z is the Pauli-Z operator, and I is the identity operator. We define the locality of the Pauli string operator by its length k. In the shallow circuit limit ($L \to 0$), the variance of estimation for $Z^{(k)}$ scales $\operatorname{Var} Z^{(k)} \propto 4^k$. So shallow circuit is only efficient for predicting the local observables, and becomes inefficient for non-local observables. In the deep circuit limit ($L \to \infty$), as the unitary ensemble becomes globally Haar, there is no difference between local and non-local operators in this limit, and $\operatorname{Var} Z^{(k)} \propto 2^N$. A simple comparison indicates: when $k \gtrsim N/2$, $\operatorname{Var} Z^{(k)}$ will decrease with L, thus deep circuits will have lower sample complexity; when $k \lesssim N/2$, $\operatorname{Var} Z^{(k)}$ will increase with L, thus shallow circuits will have lower sample complexity. In Fig. 6.6(c), the dots shows the variance $\operatorname{Var} Z^{(k)}$ as a function of circuit depth L for different support k. The trend agrees with our simple argument. For non-local operators, their variance will quickly decrease with the circuit depth L, while the variance for local operators will mildly increase with L. The behavior is theoretically described by how the operator shadow norm $\|O\|^2_{\mathcal{E}_\sigma}$ depends on both the circuit depth L and the operator locality k, which are separately encoded in the entanglement features of $\mathcal{E}_\sigma$ and $\mathcal{E}_O$. We calculate the shadow norm $\|Z^{(k)}\|^2_{\mathcal{E}_\sigma}$ based on the entanglement feature formalism using Eq. (6.21), and plot the result as lines in Fig. 6.6(c). The theoretical calculation agrees perfectly with the numerical results, which also indicates that the shadow norm bounds the single-shot variance (and hence the sample complexity) quite tightly.

Based on the discussion in Sec. 6.2.2, the sample complexity M is proportional to the single-shot variance. Given the scaling of variance with the the circuit depth L, we can study the

(a) Low–rank observables

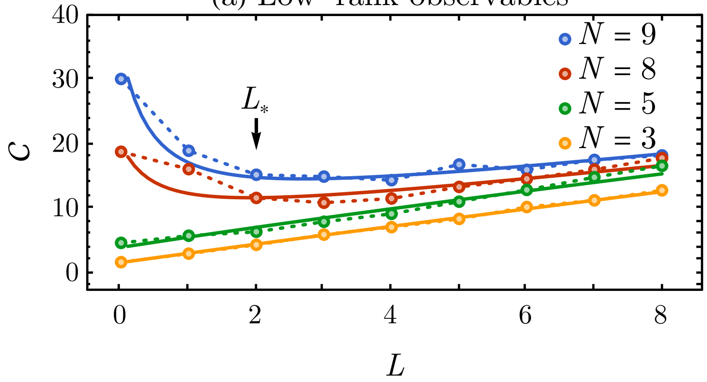

(b) Full–rank observables

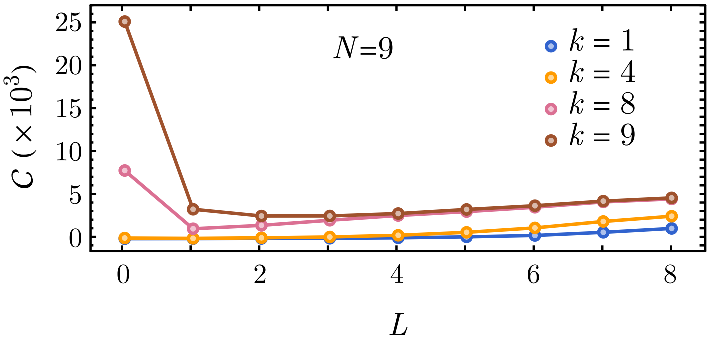

Figure 6.7: (a) Tomography complexity $\mathcal{C} \propto (L+1)\,\mathrm{Var}\,F$ as a function of circuit depth L for the fidelity (low-rank observable) estimation task. Dots are tomography complexities for GHZ states of qubit number N by our numerical simulation. Solid curves are best fits based on the empirical formula Eq. (6.37). (b) Tomography complexity $\mathcal{C} \propto (L+1)\,\mathrm{Var}\,Z^{(k)}$ for the Pauli string (full-rank observable) estimation task. Dots are numerical simulation results. Solid curves are analytic calculations using the operator shadow norm formula Eq. (6.21).

scaling of the sample complexity, as well as that of the tomography complexity $\mathcal{C}$. For fidelity estimation task, $\mathcal{C}$ scales as

$$\begin{aligned} \mathcal{C} &= (L+1)M \propto (L+1)\,\mathrm{Var}\,F \\ &\propto (L+1)\exp\left(\frac{cN}{(L+1)^{\alpha}}\right). \end{aligned} \tag{6.37}$$

For sufficiently large systems, the complexity $\mathcal{C}$ can have a non-trivial minimum at a finite circuit depth $L_{*} \simeq (\alpha c N)^{1/\alpha} - 1$. Our simulation result in Fig. 6.7(a) verifies such behavior. For small systems ($N \lesssim 5$), random single-qubit measurements can efficiently benchmark the quantum state, so we do not need to use a finite-depth circuit for data acquisition. However, as the system size N gets larger, to maintain the prediction accuracy, single-qubit measurements will require more and more samples that have to grow exponentially with N. As shown in Fig. 6.6(a), applying a few layers of quantum circuits before the measurement can quickly bring down the single-shot variance (and hence reduce the sample complexity). However, we also do not want to go too far in the circuit depth, because that would increase the circuit complexity. Therefore, we expect an optimal circuit depth L_{*} where the sample complexity and the circuit complexity reach a balance, and the total tomography complexity is minimized. This explains the advantage of shallow circuits in classical shadow tomography, as compared to the existing method that requires either on-site Haar random ($L \to 0$) or global Haar random ($L \to \infty$) unitaries.

We also study the tomography complexity $\mathcal{C} \propto (L+1)\,\mathrm{Var}\,Z^{(k)}$ for the full-rank observables, such as Pauli strings $Z^{(k)}$, as shown in Fig. 6.6(b). In this case, what matters is the locality k of the full-rank operator (the length k of the Pauli string). For local operators (small k), on-site measurement will be most efficient. However, for non-local operators (large k), we observe that the tomography complexity is minimized at some finite circuit depth, again demonstrating the advantage of employing shallow circuits in classical shadow tomography. For different classes of physical observables, we can use the tomography complexity $\mathcal{C}$ as an objective function to guide the design of the optimal circuit structure. We will leave this promising direction for future investigation.

6.3.3 Approximate Classical Shadow Tomography with Local Hamiltonian Dynamics

Requiring an unitary ensemble to be strictly locally scrambled could be restrictive. To this end, we would like to explore a broader class of unitary ensembles that are only approximately locally scrambled. In particular, we study unitary evolutions $U = e^{-\mathrm{i}HT}$ generated by a local Hamiltonian H for finite amount of time T, as depicted in Fig. 6.3(b). Two classes of Hamiltonians are of particular interest. In the first class, we consider a model of random local Hamiltonians

$$H = \sum_i H_{i,i+1}, \tag{6.38}$$

where each term $H_{i,i+1}$ is independently sampled as 2-local GUE random matrices. We dub this class the *GUE2 ensemble* to remind ourselves that the Hamiltonian is 2-local. The local Hamiltonian describes a disordered one-dimensional quantum system in general. Once every $H_{i,i+1}$ term is sampled, we will use the Hamiltonian H to drive the time evolution without changing H during the evolution. The unitary GUE2 ensemble is only invariant under $U \to V^\dagger UV$ (not $U \to UV$) for $V \in U(d)^N$, such that that its corresponding prior POVM $\mathcal{E}_\sigma$ will transform as $\hat{\sigma}_{U,b} = U^\dagger |b\rangle\langle b| U \to V^\dagger U^\dagger V |b\rangle\langle b| V^\dagger UV \neq V^\dagger \hat{\sigma}_{U,b} V$, which does not satisfy the locally scrambling condition at the state level (i.e. the invariance under $\hat{\sigma} \to V^\dagger \hat{\sigma} V$). However, we anticipate that under a sufficient amount of time evolution, the original local basis choice (of $|b\rangle$) will be quickly randomized given the chaotic nature of the local Hamiltonian, such that the initial choice of $V|b\rangle\langle b|V^\dagger$ or $|b\rangle\langle b|$ will not make a substantial difference statistically, so the GUE2 ensemble will become approximately locally scrambled after some local thermalization (scrambling) time T_{Th}.

Another more realistic class of random Hamiltonians to be considered is based on the quantum Ising model with both disordered coupling in space and random fields in time

$$H_t = \sum_{\langle ij\rangle} J_{ij} X_i X_j + h \sum_i (\cos\theta_t X_i + \sin\theta_t Y_i), \tag{6.39}$$

where the local coupling $J_{ij} \sim \mathrm{Uni}[J - \frac{J}{2}, J + \frac{J}{2}]$ is drawn from a uniform distribution, and the angle of magnetic field $\theta_t \sim \mathrm{Uni}[0, 2\pi]$ is also random. We use this Hamiltonian to drive the quantum dynamics in discrete time steps. In each period of time, the magnetic field h will be applied along a different random direction θ_t in the x-y plane for all spins uniformly. However, J_{ij} will remain the same throughout the time evolution. The ensemble of unitary consists of

$$U = \prod_{t=1}^{T} e^{-iH_t}. \tag{6.40}$$

which we name as the *Disordered Quantum Ising Model* or DQIM for short. The DQIM ensemble is friendly for quantum technology such as Rydberg-atom-based[Saf16b] or trapped-ion-based[MCD$^+$21] quantum simulators. Similar construction of approximate unitary designs by Hamiltonian evolution with random quenches in time was also proposed in Ref. [EVD$^+$18, VED$^+$18]. We would like to investigate how well our framework applies to these two cases.

Each approximately locally scrambled unitary ensembles $\mathcal{E}_U$ leads to a prior POVM $\mathcal{E}_\sigma = \{\hat{\sigma}_{U,b} | b \in \{0,1\}^{\times N}, U \in \mathcal{E}_U\}$ that is also approximately locally scrambled. We propose to characterize how close the prior POVM $\mathcal{E}_\sigma$ is towards its local-basis invariant limit by the following frame potential

$$\mathcal{F}^{(k)}_{\mathcal{E}_\sigma} = \mathop{\mathbb{E}}_{\hat{\sigma},\hat{\sigma}' \in \mathcal{E}_\sigma} (\mathrm{Tr}\,\hat{\sigma}\hat{\sigma}')^k. \tag{6.41}$$

Recall that in deriving Eq. (6.6) from Eq. (6.3), we only require the 2nd moment to match, i.e.

$$\mathop{\mathbb{E}}_{\hat{\sigma} \in \mathcal{E}_\sigma} \hat{\sigma}^{\otimes 2} = \mathop{\mathbb{E}}_{V \in U(d)^N} \mathop{\mathbb{E}}_{\hat{\sigma} \in \mathcal{E}_\sigma} (V^\dagger \hat{\sigma} V)^{\otimes 2}, \tag{6.42}$$

therefore we will be most interested in the 2nd frame potential $\mathcal{F}^{(2)}_{\mathcal{E}_\sigma}$. The frame potential $\mathcal{F}^{(2)}_{\mathcal{E}_\sigma}$ for any ensemble $\mathcal{E}_\sigma$ is lower bounded by its locally-scrambled ($\mathrm{U}(d)^N$-twirled) limit $\mathcal{F}^{(2)}_{\mathcal{E}^{\mathrm{LS}}_\sigma}$ as

$$\mathcal{F}^{(2)}_{\mathcal{E}_\sigma} \geq \mathcal{F}^{(2)}_{\mathcal{E}^{\mathrm{LS}}_\sigma} = \sum_{A,B} W^{(2)}_{\mathcal{E}_\sigma,A} \mathsf{Wg}_{A,B} W^{(2)}_{\mathcal{E}_\sigma,B}. \tag{6.43}$$

The fact that $\mathcal{F}^{(2)}_{\mathcal{E}^{\mathrm{LS}}_\sigma}$ is expressed purely in terms of the entanglement feature of $\mathcal{E}_\sigma$ indicates that it is indeed free of any local-basis-dependent information. We can define the gap between the frame potential and its locally-scrambled limit as

$$\begin{aligned}\Delta^{(2)}_{\mathcal{E}_\sigma} &= \mathcal{F}^{(2)}_{\mathcal{E}_\sigma} - \mathcal{F}^{(2)}_{\mathcal{E}^{\mathrm{LS}}_\sigma} \\ &= \mathrm{Tr}\Big(\mathop{\mathbb{E}}_{\hat{\sigma}\in\mathcal{E}_\sigma}\big(\hat{\sigma}^{\otimes 2} - \mathop{\mathbb{E}}_{V\in\mathrm{U}(d)^N}(V^\dagger\hat{\sigma}V)^{\otimes 2}\big)\Big)^2,\end{aligned} \tag{6.44}$$

which turns out to match the trace-square-difference between the 2nd moment $\mathbb{E}_{\hat{\sigma}}\hat{\sigma}^{\otimes 2}$ and its local twirling $\mathbb{E}_{V,\hat{\sigma}}(V^\dagger\hat{\sigma}V)^{\otimes 2}$. The frame potential gap $\Delta^{(2)}_{\mathcal{E}_\sigma}$ serves as an indicator of the validity of our approach, as it vanishes if $\mathcal{E}_\sigma$ is locally scrambled such that our construction becomes exact.

Different unitary ensembles can lead to different frame potential gaps of $\mathcal{E}_\sigma$, which can be used to evaluate the quality of the unitary ensemble in obeying the local scrambling condition. In Fig. 6.8(a), we first focus on the frame potential gap $\Delta^{(2)}$ for the GUE2 ensemble. We find the gap will first decay exponentially and then saturate to a plateau at a very low level. The quickly vanishing gap implies that the GUE2 ensemble quickly becomes approximately locally scrambled as time evolves. We define the characteristic time associated with the exponential decay as T_{Th}, i.e. $\Delta^{(2)}(T) \propto \exp(-T/T_{\mathrm{Th}})$, which can be considered as the local scrambling (thermalization) time. In addition, the inset plot in Fig. 6.8(a) shows that T_{Th} is independent of the system size N, as the slope remains mostly the same for different N. Unlike the global scrambling (global thermalization) which requires a long time ($\sim N$) to achieve, achieving local scrambling only

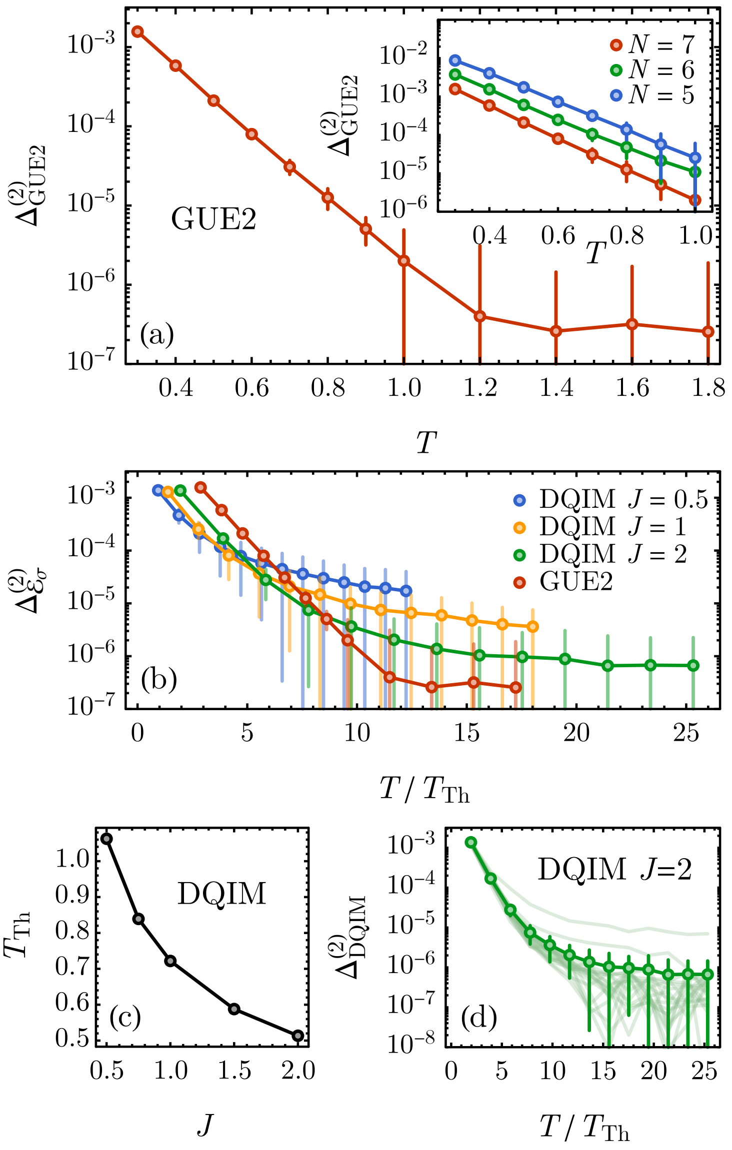

Figure 6.8: (a) Frame potential gap $\Delta^{(2)}_{\mathcal{E}_\sigma}$ of the GUE2 ensemble as a function of evolution time T. The inset shows the decay behavior for different system sizes N. (b) Frame potential gap of DQIM ensemble at different coupling strength J, in comparison with that of the GUE2 ensemble. (c) The dependence of the local scrambling time T_{Th} on the coupling strength J. (d) Frame potential gap for single instances in the DQIM ensemble. Each instance corresponds to a light-green curve in the background.

requires a fixed amount of time set by the ultra-violet energy scale that is independent of the system size N. This is another advantage of using locally scrambled quantum dynamics for classical shadow tomography in practice.

As for the DQIM ensemble, we fix the strength of the magnetic field at $h = \pi/4$, since this value produces the fastest on-site scrambling of a single qubit. According to the definition Eq.(6.39), the only tuning parameter will be the mean value J of Ising couplings (which also sets their disorder strength). We calculate the frame potential gap for DQIM ensemble with different J, and use the exponential decay regime to define the local scrambling time T_{Th}. The result is shown in Fig. 6.8(b). The DQIM ensemble also approaches local scrambling as time evolves, although the final saturation plateau is not as low as the GUE2 ensemble. Larger Ising coupling J will lower the saturation plateau and shorter the local scrambling time T_{Th}, as shown in Fig. 6.8(c). In addition, as shown in Fig. 6.8(d), we find the frame potential gap for a single realization quenched-disorder Hamiltonian does not deviate significantly from the ensemble mean value. This indicates that a single fixed disordered Ising chain under a randomly rotating uniform magnetic field is already good to generate an approximately locally scrambled ensemble that can be used for classical shadow tomography.

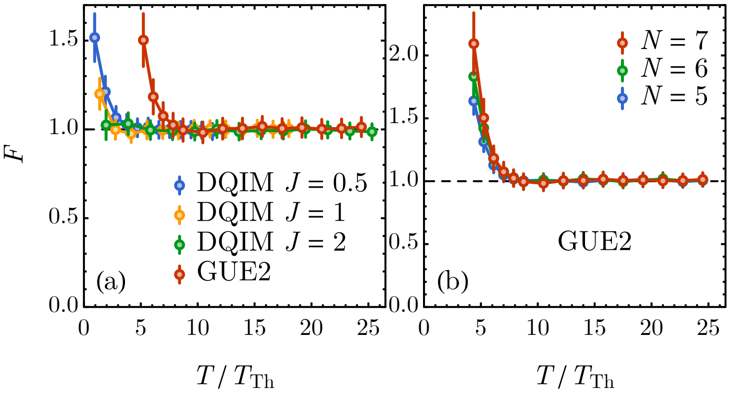

Figure 6.9: Fidelity prediction by (a) different approximated locally scrambled ensembles, and (b) the GUE ensemble at different system sizes N. Sample number is 10000 and error bar indicates 3-standard deviation.

In practice, we use the two proposed approximated ensembles, (i) the GUE2 ensemble and (ii) a single instance of the DQIM ensemble, to perform the tomography task and predict the fidelity of a 7-qubit GHZ state. In Fig. 6.9(a), we see the predicted fidelity will be biased in the beginning (the biased fidelity can be greater than one, see Appendix B.4 for more discussions), due to the fact that the quantum dynamics is still on its way to establish local scrambling. After around $T \sim 10T_{\mathrm{Th}}$, the local scrambling condition is approximately established, then the entanglement-feature-based reconstruction map $\mathcal{M}_{\mathrm{EF}}^{-1}$ can provide a good reconstruction of the quantum state, as indicated by the convergence of the quantum fidelity to identity. In Appendix B.4, we further investigate the quantum fidelity of $\tilde{\rho}$ projected to the physical space (to tame the unphysical $F > 1$ behavior) and show that the reconstruction is perfect after around $T \sim 10T_{\mathrm{Th}}$. In addition, Fig. 6.9(b) also shows the local scrambling time for GUE2 is independent of system size, which is consistent with the same behavior in Fig. 6.8(a). The results in Fig. 6.9 suggest that the entanglement-feature-based approach could be applicable for approximately locally scrambled unitary ensembles. The reconstruction bias vanishes as the frame potential gap decays. As long as the frame potential gap is low enough, the bias is also expected to be vanishingly small for all predictions. This significantly broadens the application of classical shadow tomography to a large class of quantum dynamics that can be achieved on NISQ devices.

6.4 Summary and Discussions

Our result can be further extended to more general measurement channels, which can involve ancilla qubits and partial measurements. The unitary channel can be noisy and the measurements can be weak. Under generalized measurements, the state ρ collapses to $\rho \to K_a \rho K_a^\dagger / (\operatorname{Tr} K_a \rho K_a^\dagger)$, where K_a is the Kraus operators[Kra71] associated with the measurement outcome a. We can define the measurement operator $\hat{\sigma}_a = K_a^\dagger K_a$ (with the standard normalization $\sum_a \hat{\sigma}_a = \mathbb{1}$), which forms the prior POVM $\mathcal{E}_\sigma = \{\hat{\sigma}_a | P(\hat{\sigma}_a) = d^{-N}\}$, and the posterior POVM will be $\mathcal{E}_{\sigma|\rho} = \{\hat{\sigma}_a | P(\hat{\sigma}_a|\rho) = \operatorname{Tr} \hat{\sigma}_a \rho\}$ correspondingly. As long as the generalized prior POVM $\mathcal{E}_\sigma$ is locally scrambled, i.e. $\forall V \in \mathrm{U}(d)^N : P(\hat{\sigma}) = P(V^\dagger \hat{\sigma} V)$, our theoretical framework automatically applies, and all formulations in this work remain valid in the same form. This enables us to consider classical shadow tomography with very general data acquisition protocols.

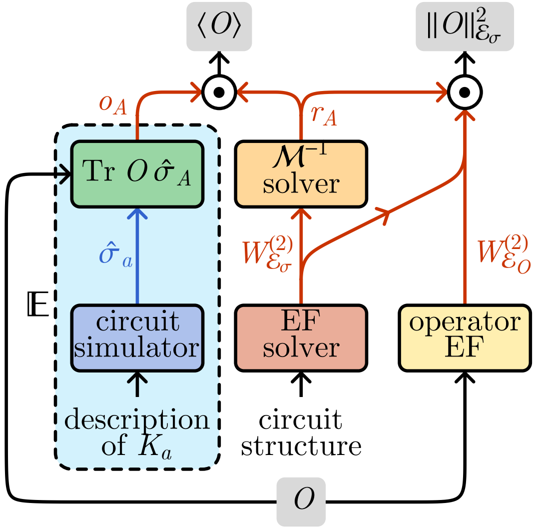

Figure 6.10: Classical post-processing protocol to estimate the operator expectation value and shadow norm.

The entanglement feature formalism plays a central role in our approach. Fig. 6.10 summarizes the proposed classical post-processing protocol to predict the expectation value $\langle O \rangle$ of a physical observable O, together with its estimated variance (given by the shadow norm $\|O\|^2_{\mathcal{E}_\sigma}$ divided by the sample size M). Given the circuit structure, the entanglement feature (EF) solver calculates the entanglement feature $W^{(2)}_{\mathcal{E}_\sigma}$ of the prior POVM as defined in Eq. (6.8) (the algorithm is developed in previous works[KAAY20, FVVY20, AY20]). The result is passed to the inverse channel solver to calculate the reconstruction coefficients r_A by solving Eq. (6.12). With r_A, we can predict any physical observable O by $\langle O \rangle = d^N \sum_A r_A o_A$ where $o_A = \mathbb{E}_{\hat{\sigma} \in \mathcal{E}_{\sigma|\rho}} \operatorname{Tr} O \hat{\sigma}_A$ (the median-of-means trick[HKP20] can be used here if multiple observables are to be predicted). For every sample of classical description of the Kraus operator K, a quantum circuit simulator (running on a classical computer) is needed to construct the (efficient representation of) measurement operator $\hat{\sigma} = K^{\dagger}K$. The classical simulation could be efficient if the circuit is Clifford[Got98b] (our formalism applies to random Clifford circuits with no problem). The part of computation in the dashed box of Fig. 6.10 should be repeated for every sample to evaluate the ensemble average. Finally, given the reconstruction coefficient r and the entanglement features $W^{(2)}_{\mathcal{E}_\sigma}$ and $W^{(2)}_{\mathcal{E}_O}$, the shadow norm $\|O\|_{\mathcal{E}_\sigma}$ can be calculated, which provides an estimation for variance of the predicted observable. Although it takes some effort to process the entanglement feature data and to calculate the reconstruction coefficients, such computation (everything outside the dashed box in Fig. 6.10) only occurs once for a given circuit structure, therefore this computational effort is usually affordable (especially when efficient tensor-network approaches are developed and employed).

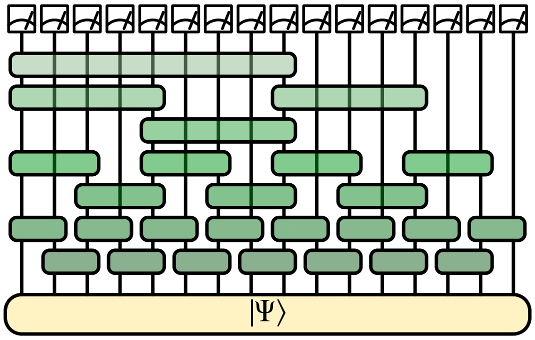

Figure 6.11: Illustration of holographic classical shadow tomography scheme, where the quantum circuit is arranged in a hierarchical structure (forming the hyperbolic bulk space).

The theoretical framework established in this work extends the classical shadow tomography to general quantum circuits, which opens up many possible applications. As one interesting example, we consider performing the classical shadow tomography in the "holographic bulk" by transforming the original state by a random Clifford circuit arranged in a hierarchical structure (see Fig. 6.11), similar to the multi-scale entanglement renormalization ansatz (MERA) network[Vid07b, Vid08b] or the holographic quantum error-correcting code[PYHP15a]. Following the idea of holographic duality, local measurements in the holographic bulk translate to measurements at all different scales on the holographic boundary. Therefore it is conceivable that the holographic classical shadow tomography could achieve high sample efficiency for operators of all scales, potentially evading the dichotomy between sample complexity and circuit complexity.

Another interesting application is to consider random circuits hybrid with random measurements inserted into the circuit at a fixed rate[LCF18, LCF19, SRS19, CNPS19, SRN19]. Conditioned on the intermediate measurement outcomes, the hybrid quantum circuit forms a quantum channel that transmits quantum information from end to end. Driven by the measurement rate, the final state can undergo an entanglement transition[VPYL18, JYVL19, BCA19] (or the quantum channel can undergo a purification transition[GH19] equivalently). When the measurement rate is high, the quantum information in the initial state can be efficiently extracted by intermediate measurements (eavesdroppers), such that the channel has zero transmission capacity. When the measurement rate is lower than a critical threshold, the channel will have a finite capacity and can transmit quantum information in an error-correcting manner.[CBQA19, FVVY20, GKH$^+$20] However, it is unclear how to take advantage of the self-organized quantum error correction in these hybrid quantum circuits. We anticipate that the classical shadow tomography with a flexible measurement scheme can help decoding the measurement-induced quantum error-correcting code. We will leave these interesting applications for future explorations.

Finally, the classical shadow tomography provides an efficient interface that converts quantum states to classical shadow data, which enables us to exploit the power of classical computation, especially data-driven and machine learning approaches, to advance our understanding of complex quantum systems and to solve challenging quantum many-body problems. As shown in Ref. [HKT$^+$21], classical algorithms that learns from the classical shadow data has provable performance advantages over conventional numerical approaches that do not learn of data. Our work further adds to this promising direction by providing a more flexible classical shadow tomography scheme that works with very general measurement protocols (beyond on-site Pauli measurements), which could lead to potentially more efficient classical-shadow-based learning algorithms.

6.5 Acknowledgements

This work is supported supported by a UC Hellman fellowship. Chapter 4, in full, is a reprint of the material as it appears in Hong-Ye Hu, Soonwon Choi, Yi-Zhuang You, arXiv:2107.04817. The dissertation author was the primary investigator and author of this paper.

Part III

Hierarchical representation of quantum states

Chapter 7

Topological and symmetry-enriched random quantum critical points

7.1 Introduction

Topological phases form a cornerstone of modern condensed matter physics, extending beyond the Landau-Ginzburg paradigm of symmetry-breaking order. An especially important class of topological states are Symmetry-Protected Topological (SPT) phases [GW09, CGW11, TPB11, FK11, CLW11, PBTO12, LV12, LG12, CGLW12, CGLW13, Sen15], which are gapped systems characterized by non-local order parameters and symmetry-protected topological edge modes. Prominent examples of SPT phases include fermionic topological insulators [VBA85, SV88, MNZ04, KM05, FK07, Kit09, SRFL09, HK10]—protected by time-reversal and charge conservation symmetry—or the Haldane phase in quantum spin chains [Hal83, AKLT88, DR93, PTBO10]—protected by spin-rotation symmetry.

Recently, the concept of SPT order was extended to gapless systems [KWSDS11, FLNF11, SHFDS11, Tsv11, CT11, RDTHA12, GV12, KDB$^+$13, ODC$^+$14, RBA15, IMR$^+$15, LB15, KB15, KC15, ZW17, OC16, MDIR17, WXPH17, KSH17a, RA17, KSGC17, SPV17, GLB17,

CYT$^+$18, VJP18, ZL18, JLSL18, PSV18, KBA18, JV19, VTJP19, JSW19, Ver20]: surprisingly, many of the key features of SPT physics carry over to the gapless case, despite the non-trivial coupling between topological edge modes and bulk critical fluctuations. It is also helpful to think of gapless SPT (gSPT) states [SPV17] as symmetry-enriched quantum critical points (SEQCP) [VTJP19], where global symmetries can enrich the critical behavior of critical systems. This led to the discovery of new critical points and phases with unusual nonlocal scaling operators which imply an anomalous surface critical behavior, and symmetry-protected topological edge modes. In certain cases, such SEQCPs are naturally realized as phase transitions separating SPT and symmetry-broken phases: while the bulk universality class is locally dictated by the Landau-Ginzburg theory of spontaneous symmetry-breaking, the nonlocal operators and the surface critical behavior are affected by the neighboring SPT phase.

In this work, we show that the mechanism protecting gapless SPT phases persists upon adding disorder. We focus on one-dimensional systems, where the bulk criticality flows to infinite-randomness fixed points [MDH79, DM80, Fis92, Fis94, Fis95, MMHF00]. We first discuss the paradigmatic infinite-randomness Ising criticality, where we find that—similar to the clean case [VTJP19]—there are topologically distinct versions in the presence of an additional $\mathbb{Z}_2^T$ symmetry. We find that one of these classes has topologically-protected edge states. Whilst this is a finite-tuned critical point, our second example is a stable "random singlet" phase of matter. Moreover, in the latter case, there are additional gapped degrees of freedom which are able to make the edge mode exponentially-localized. We also illustrate how this topological random quantum criticality can emerge naturally in periodically driven (Floquet) systems.

7.2 Ising* transition

We consider the spin-1/2 chain

$$H = -\sum_i J_i Z_i Z_{i+1} - \sum_i h_i X_i - \sum_i g_i Z_{i-1} X_i Z_{i+1}, \tag{7.1}$$

where X, Y, Z denote the Pauli matrices. The model has a $\mathbb{Z}_2$ spin-flip symmetry (generated by $P = \prod_i X_i$) and a time-reversal symmetry $\mathbb{Z}_2^T$ (acting as the complex conjugation $T = K$). Let us first consider the clean case, where the coefficients $J_i \equiv J$, $h_i \equiv h$ and $g_i \equiv g$ are site-independent. In this case, the $J, h, g \geq 0$ terms respectively drive the system towards ferromagnetic (FM), trivial paramagnetic (PM) and $\mathbb{Z}_2 \times \mathbb{Z}_2^T$ symmetry protected topological (SPT) [Suz71, RB01, KM04, SAF$^+$11b, VMP17] phases, the latter sometimes being called the cluster or Haldane SPT phase. The phase diagram is shown in Fig. 7.1(a), with the gray solid lines indicating Ising criticalities.

Although the FM-PM and FM-SPT transition are both described by the Ising conformal field theory (CFT), the time-reversal symmetry acts differently on the disorder operator, leading to different symmetry enriched CFTs (or gapless SPTs) [SPV17, VJP18, PSV18, VTJP19, Ver20]. To briefly review this, note that an Ising CFT has a unique local and a unique nonlocal scaling operator with scaling dimension $\Delta = 1/8$, commonly denoted by σ and μ, respectively. These are the order parameters of the nearby phases, i.e., $\sigma(n) \sim Z_n$ is the Ising order parameter, whereas the disorder operator $\mu(n)$ is the Kramers-Wannier-dual string order parameter of the symmetry-preserving phase. In the trivial PM, $\mu(n) \sim \prod_{j=-\infty}^n X_j$, whereas in the SPT phase, $\mu(n) \sim \prod_{j=-\infty}^n Z_{j-1} X_j Z_{j+1} = \cdots X_{n-2} X_{n-1} Y_n Z_{n+1}$ [SAF$^+$11a, BV14, JV19, VTJP19]. We see that the two Ising critical lines are distinguished by the discrete invariant $T\mu T = \pm\mu$ [VTJP19]. This means they must be separated by a phase transition. Indeed, in Fig. 7.1(a) they meet at a multicritical point where the central charge is $c = 1$.

We refer to the non-trivial case, where the nonlocal bulk operator is charged $T\mu T = -\mu$, as Ising*. This supports a localized zero-energy edge state [VTJP19]. Intuitively, the edge of the

Ising* criticality spontaneously breaks the Ising $\mathbb{Z}_2$ symmetry. This unusual degenerate boundary fixed point is stable since μ is charged and hence cannot be used to disorder the boundary. The finite-size splitting of this edge mode is parametrically faster than the finite-size bulk gap $\sim 1/L$. In particular, if the model is dual to free-fermions (such as Eq. (7.1)) then the edge mode is exponentially-localized [VJP18] whereas with interactions, the splitting becomes $\sim 1/L^{14}$ [VTJP19].

7.3 Random Ising* transition

We now study the fate of Ising* upon disordering the system. The coefficients J_i, h_i and g_i in Eq. (7.1) are now independently drawn from power-law distributions $P(J) = (J/J_0)^{1/\Gamma}/(\Gamma J)$ for $J \in [0, J_0]$ (similarly for $P(h)$ and $P(g)$), where Γ controls the width of the distribution in logarithmic scale. The limit $\Gamma \to 0$ would recover the clean case. We will take $\Gamma = 1$, i.e., the uniform distribution.

In the presence of randomness, the Ising CFT flows towards the infinite-randomness fixed point ($\Gamma \to \infty$) [Fis92, Fis95]. We will explore the symmetry enriched infinite-randomness fixed point as the many-body localized counterpart of gapless SPT states. The disordered phase diagram is shown in Fig. 7.1(a), which is qualitatively unchanged from the clean case. This was obtained by mapping Eq. (7.1) to free fermions (using a Jordan-Wigner transformation) and using the transfer matrix method to determine the topological winding number ω [MDH01]; in this case the PM, FM and SPT phases map to the trivial ($\omega = 0$), Kitaev chain ($\omega = 1$) and two Kitaev chains ($\omega = 2$). In the original spin chain language, one can interpret ω as encoding the ground state degeneracy 2^ω with open boundary conditions, which is 0, 2 and 4, respectively.

Similar to the Ising CFT, the infinite-randomness Ising fixed point also has a local σ and nonlocal μ scaling operator. While their scaling dimensions have changed ($\Delta^{\text{bulk}} = 1 - \varphi/2 \approx 0.191$ where $\varphi = \frac{1}{2}(1+\sqrt{5})$ is the golden ratio [Fis95]), their lattice expressions are as before—

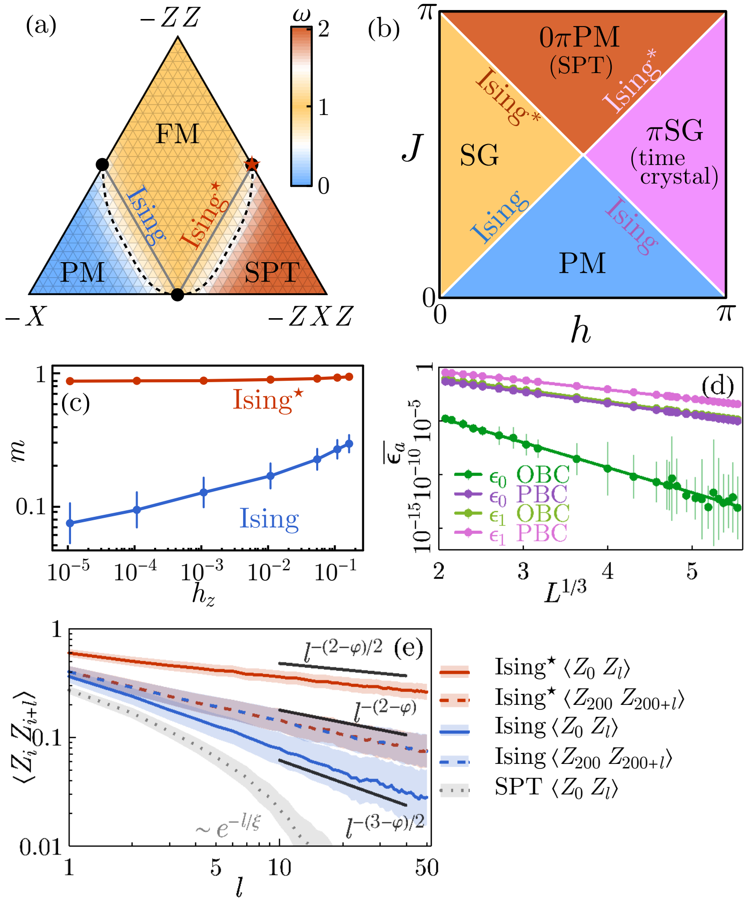

Figure 7.1: **Random Ising* transition** (a) Phase diagram of the random Ising Hamiltonian (7.1) for clean (solid lines) and disordered (dotted lines), showing the topological winding number ω for the dual fermionic description (see main text). (b) Floquet phase diagram of eq. (7.4) which shows *two* topologically non-trivial Ising* transitions. (c) Boundary magnetization under small Zeeman field, showing spontaneous magnetization at the Ising* transition (red star in (a)). (d) Finite-size energy splitting of boundary spins at the Ising* transition. (e) Spin-spin correlations involving bulk and boundary spins (averaged over 1.5×10^5 realizations), compared to theory predictions (solid black lines), where φ is the golden ratio. Calculations are performed using the SBRG method on a 512-site spin chain.

indeed, the nearby gapped phases are still characterized by the same order parameters. We thus still have the bulk topological invariant $T\mu T = \pm\mu$, distinguishing two distinct symmetry-enriched infinite-randomness Ising fixed points, which we refer to as the Ising and Ising*. For the same reasons as before, we expect that the disordered Ising* criticality has spontaneously-fixed boundary conditions. This would come with at least three physical fingerprints: (i) a nonzero spontaneous magnetization at the boundary, (ii) a degenerate edge mode whose finite-size splitting is parametrically smaller than the bulk gap, and (iii) spin-spin correlations near the boundary should have a boundary scaling dimension [McC69, IR98] $\Delta_\sigma^{\text{bdy}} = 1/2$ (or 0) for free (or spontaneously-fixed) boundary condition, characterizing the Ising (Ising*) case.

We now test these predictions numerically. Because we will be interested in including interactions, we use the spectrum bifurcation renormalization group (SBRG) method [YQX16, SYX16, SBYX17, sup], which is a numerical real-space renormalization group approach that progressively transforms the original Hamiltonian H to its diagonal form $H_{\text{MBL}} = \sum_a \varepsilon_a \tau_a + \sum_{ab} \varepsilon_{ab} \tau_a \tau_b + \cdots$ as a many-body localization (MBL) effective Hamiltonian [SPA13, HNO14, Swi13], and constructs the (approximate) local integrals of motion τ_a of the MBL system in the form of Pauli strings. The approximation is asymptotically exact in the strong-disorder limit. The rescaled parameters $(\tilde{J}, \tilde{h}, \tilde{g}) \equiv (J_0, h_0, g_0)^{1/\Gamma}$ are invariant under the renormalization group (RG) flow, and should be considered as tuning parameters. SBRG can be thought of as an implementation of the strong disorder real space renormalization group (RSRG) [MDH79, DM80, Fis92, Fis94, Fis95, MMHF00] and its generalization to excited states (RSRG-X) [PRA$^+$14, VA13, VA14, VPP15] in operator space. While SBRG can be used to study MBL physics and excited states, in the following we focus on $T = 0$ groundstate properties.

7.4 SBRG results

We focus on the Ising* transition at $(\tilde{J}, \tilde{h}, \tilde{g}) = (1, 0, 1)$ (red star in Fig. 7.1(a)). We have verified [sup] that in the bulk, the Ising* transition flows to an infinite-randomness fixed point with dynamical scaling $l \sim (\log t)^2 \sim (-\log \varepsilon)^2$ that relates the length scale l and the energy scale ε [Fis95], and logarithmic scaling of the entanglement entropy [RM04, RM09]. This is not surprising since with periodic boundary conditions, Ising and Ising* are unitarily equivalent.

We now probe the boundary properties. To include the effect of interactions, we follow Ref. [VTJP19] and add a generic $\mathbb{Z}_2 \times \mathbb{Z}_2^T$-symmetric boundary perturbation $H_V = -V(X_0 Z_1 Z_2 + Z_{L-2} Z_{L-1} X_L)$, with V a random variable ten times smaller than the bulk couplings. Microscopically, this perturbation can flip the boundary Ising spin. Nevertheless, if we study the boundary magnetization $m = \langle Z_0 \rangle$ in response to a small Zeeman field h_z applied along the z-axis, we find that it tends to a nonzero limit as $h_z \to 0$ (with h_z smaller than the finite-size bulk gap, but larger than the groundstates splitting, see below), shown in Fig. 7.1(d). This is in contrast to the trivial Ising fixed point, where the boundary magnetization is known to vanish as $m(h_z) \sim 1/|\log h_z|$ [McC69].

Thus the boundary is spontaneously magnetized in the Ising* case despite the Hamiltonian (7.1) being symmetric. Schematically, on a finite system we have two spontaneously-fixed ferromagnetic (FM) ground states $|\uparrow_L \uparrow_R\rangle$ and $|\downarrow_L \downarrow_R\rangle$, where L and R denote the configurations of the left and right edge modes (note that these are split from $|\uparrow_L \downarrow_R\rangle$ and $|\downarrow_L \uparrow_R\rangle$ by the critical bulk penalizing antiferromagnetic states) [SPV17, VTJP19]. The above perturbation H_V can couple these FM states at second order in V, which should lead to a finite-size splitting. The claim that we have a ground state degeneracy is only meaningful if this splitting is smaller than the bulk finite-size gap. To confirm this, we arrange the energy coefficients ε_a obtained from SBRG in the ascending order $\varepsilon_0 < \varepsilon_1 < \cdots$, and focus on the lowest two. For the Ising* transition with open boundary condition (OBC), ε_0 characterizes the smallest energy splitting between $|\uparrow_L \uparrow_R\rangle \pm |\downarrow_L \downarrow_R\rangle$ whereas ε_1 characterizes the bulk excitation gap. As shown in Fig. 7.1(e), both splittings $\overline{\varepsilon_0}$ and $\overline{\varepsilon_1}$ follow $\overline{\varepsilon_a} \sim \exp(-\alpha_a L^{1/3})$ but with different exponents $\alpha_0 = 5.4 \pm 0.6$ and $\alpha_1 = 2.51 \pm 0.02$, i.e., $\overline{\varepsilon_0} \approx \overline{\varepsilon_1}^2$. The finite-size splitting $\overline{\varepsilon_0}$ of the symmetry-protected edge modes decays significantly faster with the system size L compared to $\overline{\varepsilon_1}$. This provides a quantitative distinction between the topological edge modes and the bulk excitations. To further verify this interpretation, we switch to the periodic boundary condition (PBC), the fast-decaying topological splitting disappears and the smallest splitting decays with the "bulk" exponent as

$\alpha_0 = 2.45 \pm 0.02$.

The Ising and Ising* states can be further distinguished by their average boundary-bulk spin-spin correlation functions $\overline{\langle Z_0 Z_l \rangle}$, which decay as $\sim 1/l^{\Delta_\sigma^{\text{bdy}} + \Delta^{\text{bulk}}}$, where $\Delta_\sigma^{\text{bdy}}$ (Δ^{bulk}) is the boundary (bulk) scaling dimension of the Ising order parameter mentioned before. We thus predict

$$\overline{\langle Z_0 Z_l \rangle} \sim \begin{cases} l^{-(3-\varphi)/2} \approx l^{-0.69} & \text{Ising}, \\ l^{-(2-\varphi)/2} \approx l^{-0.19} & \text{Ising}^\star. \end{cases} \tag{7.2}$$

In Fig. 7.1(f), we find that the boundary-bulk correlation follows $\overline{\langle Z_0 Z_l \rangle} \sim l^{-(0.67 \pm 0.08)}$ for Ising and $l^{-(0.20 \pm 0.02)}$ for Ising*, which matches Eq. (7.2) within error bars. We also checked that the bulk-bulk correlation $\overline{\langle Z_i Z_{i+l} \rangle} \sim l^{-(0.42 \pm 0.05)}$ decays with the expected exponent $2\Delta^{\text{bulk}} = 2 - \varphi \approx 0.38$ for both Ising and Ising* transitions.

7.5 Symmetry-enriched random singlet phase

The Ising* transition provides a clear example of symmetry-enriched random quantum critical point, with stretched-exponentially localized edge modes. It is natural to ask whether this notion can be extended to random *critical phases*, and whether the topological edge modes can be made *exponentially* localized despite the absence of a bulk gap. Here, we answer both questions in the positive, by introducing a symmetry-enriched random singlet phase.

In order to obtain a critical phase in one dimension, we consider a system with charge conservation and particle-hole symmetry. For concreteness, we will focus on the random antiferromagnetic spin-1/2 XXZ spin chain $H_{\text{A}} = \sum_i J_i (X_i^A X_{i+1}^A + Y_i^A Y_{i+1}^A + \Delta_i Z_i^A Z_{i+1}^A)$, with $J_i > 0$ and $0 < \Delta_i < 1$ random couplings specified later. It has a symmetry group $G_A = U(1) \rtimes \mathbb{Z}_2^A$ with the $\mathbb{Z}_2^A$ spin flip generated by $\prod_i X_i^A$, while the $U(1)$ part corresponds to $\sum_i Z_i^A$ conservation. For uniform couplings, this spin chain is in a Luttinger liquid phase; while for random couplings, its low energy properties can be captured by a real-space renormalization group (RSRG) procedure very similar to the SBRG approach above (but restricted to the groundstate). The random XXZ spin chain forms a *random singlet phase* [Fis94], where the groundstate is asymptotically made of non-crossing pairs of singlets of all ranges, with quantum critical properties similar to the random Ising transition (which itself can be thought of as a random singlet state of Majorana fermions). In particular, the entanglement entropy grows logarithmically with effective central charge $c_{\text{eff}} = \log 2$ [RM07, RM09], and the gap closes stretched-exponentially with system size (dynamical exponent $z = \infty$).

To obtain a topological random singlet phase, we use the *decorated domain walls* construction [CLV14] to "twist" the random XXZ chain. To that effect, we introduce another spin species B, with Ising symmetry $G_B = \mathbb{Z}_2^B$, with Hamiltonian $H_B = -\sum_i X_i^B + g_B Z_i^B Z_{i+1}^B$. We take $g_B \ll 1$ so that the B spins are disordered, deep into a quantum paramagnetic phase. We then couple the two models by attaching charges of the $G_B = \mathbb{Z}_2^B$ symmetry to the domain walls of the A spins. This is achieved by the unitary transformation $U = \prod_{\text{DW}(A)} (-1)^{(1-Z_i^B)/2}$, where the product runs over all the domain walls of the A spins in the Z basis, with $U^2 = 1$. After unitary rotation ("twist") of $H_A + H_B + V$, we find

$$\begin{aligned} H = &\sum_i J_i \left[Z_{i-1}^B (X_i^A X_{i+1}^A + Y_i^A Y_{i+1}^A) Z_{i+1}^B + \Delta_i Z_i^A Z_{i+1}^A \right] \\ &- \sum_i Z_i^A X_i^B Z_{i+1}^A + g_B Z_i^B Z_{i+1}^B + V', \end{aligned} \tag{7.3}$$

where $V' = UVU$ represents arbitrary small perturbations that preserves the $G_A \times G_B$ symmetry. Following the terminology of Ref. [SPV17, PSV18], we refer to Eq. (7.3) and $H_A + H_B + V = UHU$ as the gSPT and gTrivial (gapless, topologically trivial) Hamiltonians, respectively.

For periodic boundary conditions, H is unitarily related to $H_A + H_B$ plus perturbations, and thus corresponds to random singlet A spins coupled to the gapped paramagnetic B spins. Nevertheless, the two models are topologically distinct. Like Ising and Ising* above, they can be distinguished by the charges of nonlocal scaling operators. In fact, since now there are

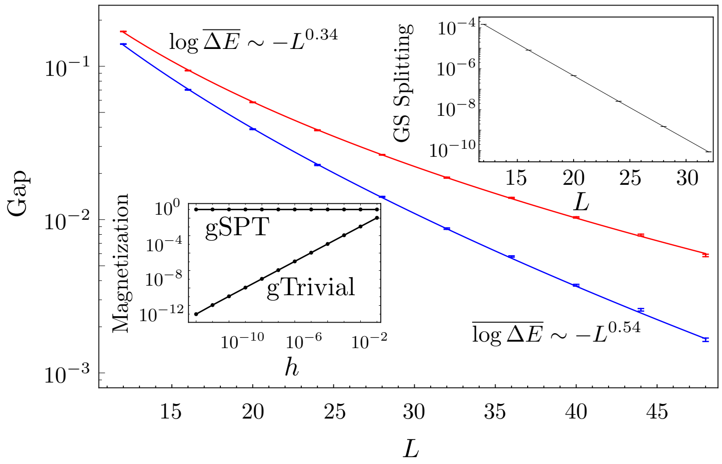

Figure 7.2: **Symmetry-enriched random singlet phase.** DMRG results on eq. (7.3) including various perturbations [sup]. Fits of the typical and average finite size gaps, showing a scaling compatible with the random-singlet $z = \infty$ scalings $\Delta E_{\text{typical}} \equiv \mathrm{e}^{\overline{\log \Delta E}} \sim \mathrm{e}^{-\sqrt{L}}$ and $\overline{\Delta E} \sim \mathrm{e}^{-L^{1/3}}$. *Top-right Inset:* the splitting between the two ground states vanishes exponentially with system size, indicating exponentially-localized edge modes. *Bottom-left Inset:* spontaneous boundary magnetization in the presence of a small symmetry-breaking magnetic field h.

additional gapped degrees of freedom, one can consider a string order parameter with long-range order: in the trivial case $H_A + H_B$ this is $\cdots X^B_{j-2} X^B_{j-1} X^B_j$ whereas in the topological case H it is $\cdots X^B_{j-2} X^B_{j-1} X^B_j Z^A_{j+1}$. In the latter case, this string order parameter for the gapped B variables is charged under G_A. This discrete invariant shows that we have two distinct symmetry-enriched versions of the same underlying infinite-randomness fixed point. Relatedly, for open boundary conditions, we have $H = J_0 \Delta_0 Z^A_0 Z^A_1 + Z^A_0 X^B_0 Z^B_1 + Z^B_0 Z^B_1 + \ldots$, and in the absence of additional perturbations ($V = 0$), we see that $[Z^A_0, H] = 0$, providing an exact edge mode.

Going away from this special limit, we expect exponentially-localized topological edge modes to be protected by the finite gap of the B spins, as in the clean case [SPV17, PSV18]. We confirmed numerically the presence of exponentially localized edge modes coexisting with bulk random singlet criticality using density-matrix renormalization group (DMRG) [Whi93, ITe] techniques (Fig. C.3), including generic symmetry-preserving perturbations [sup].

7.6 Floquet Ising* criticality

To close this letter, we illustrate how such novel universality classes emerge naturally in the context of periodically driven (Floquet) systems. We focus on the driven quantum Ising chain characterized by the single-period evolution (Floquet) operator [KLMS16]

$$F = \mathrm{e}^{-\frac{i}{2}\sum_i J_i Z_i Z_{i+1} + \cdots} \mathrm{e}^{-\frac{i}{2}\sum_i h_i X_i + \cdots} \tag{7.4}$$

where the dots represent small but arbitrary interactions preserving the $\mathbb{Z}_2$ symmetry $G = \prod_i X_i$. For strong enough disorder, this system admits four dynamical phases protected by MBL [KLMS16]. In addition to the familiar paramagnetic (PM) and spin glass (SG) Ising phases, there are two more phases called π-SG (a.k.a. time crystal [KLMS16, EBN16, vKKS16, ZHK$^+$17, KMS19]) and 0π PM (a non-trivial SPT phase); see Fig. 7.1(b). This phase structure is due to an emergent $\mathbb{Z}_2$ symmetry inherited from time translation symmetry. The transitions between those phases have been argued to be in the random Ising universality class [YPPV17, BKPV18] (ignoring potential instabilities towards thermalization in the presence of interactions [MHK20, SMY$^+$20, WAV20]). Here we note that the transitions out of the 0π PM are actually in the random Ising* universality class described above, protected by $\mathbb{Z}_2 \times \mathbb{Z}_2$ symmetry (one of the $\mathbb{Z}_2$'s being emergent). This is because the 0π PM is closely related to the $\mathbb{Z}_2 \times \mathbb{Z}_2$ equilibrium SPT [KLMS16, vS16, EN16, PMV16, HR17]. We find exponentially localized edge modes at the transitions separating the 0π PM to either the SG or π-SG, which are protected due to the disorder operator μ for the critical $\mathbb{Z}_2$ symmetry again being charged with respect to the second $\mathbb{Z}_2$ symmetry, as detailed in the supplemental material [sup]. (The edge mode localization is exponential as in the random singlet phase above, as the protecting symmetry is $\mathbb{Z}_2 \times \mathbb{Z}_2$ instead of $\mathbb{Z}_2 \times \mathbb{Z}_2^T$.)

7.7 Discussion

We have demonstrated the existence of symmetry-enriched infinite-randomness fixed points with robust topological edge modes coexisting with all the characteristics of strong disorder quantum criticality. In particular, we have shown that the paradigmatic random Ising critical point and XXZ random singlet phase come in topologically distinct versions in the presence of an additional $\mathbb{Z}_2^T$ or $\mathbb{Z}_2$ symmetry. The topological edge modes couple non-trivially to gapless bulk fluctuations, leading to anomalous boundary critical behavior. We expect our findings to extend to essentially all known strong- and infinite-randomness critical points: finding examples of symmetry-enriched random critical points in 2+1d [MMHF00, KPP$^+$] and 3+1d represents an interesting direction for future works. It would also be interesting to investigate the consequences of our results for dynamical properties [PVS19, BKPV18, KYL19, YAM20].

7.8 Acknowledgements

This work is supported by a startup fund from UC San Diego. Chapter 5, in full, is a reprint of the material as it appears in Hong-Ye Hu*, Carlos M.Duque*, Yi-Zhuang You, Vedika Khemani, Ruben Verresen, Romain Vasseur, Physical review B 103, L100207. The dissertation author was the primary investigator and author of this paper.

Chapter 8

Machine learning holographic mapping by neural network renormalization group

8.1 Introduction

The holographic duality, also known as the anti-de-Sitter space and conformal field theory correspondence (AdS/CFT) [Wit98b, Wit98a, GKP98, Mal99], is a duality between a CFT on a flat boundary and a gravitational theory in the AdS bulk with one higher dimension. It is intrinsically related to the renormalization group (RG) flow [dVV00, Ske02, HP11, Swi12a, Swi12b, NRT12, BGL13] of the boundary quantum field theory, since the dilation transformation, as a part of the conformal group, naturally corresponds to the coarse-graining procedure in the RG flow. The extra dimension emergent in the holographic bulk can be interpreted as the RG scale. In the traditional real-space RG [Kad66], the coarse-graining procedure decimates irrelevant degrees of freedom along the RG flow, therefore the RG transformation is irreversible due to the information loss. However, if the decimated degrees of freedom are collected and hosted in the bulk, the RG transformation becomes a *bijective* map between the degrees of freedom on the CFT boundary and the degrees of freedom in the AdS bulk. Such mappings, generated by information-preserving RG transforms, are called exact holographic mappings (EHM) [Qi13, LQ15, GLW$^+$16], which were first formulated for free fermion CFT. Under the EHM, the boundary features of a quantum field theory of different scales are mapped to different depths in the bulk, and vice versa. The field variable deep in the bulk represents the overall or infrared (IR) feature, while the variable close to the boundary controls the detailed or ultraviolet (UV) feature. Such a hierarchical arrangement of information is often observed in deep neural networks, particularly in convolutional neural networks (CNN) [LBH15]. The similarity between renormalization group and deep learning has been discussed in several works [B13, MS14, BO15, OT17, LTR17, GS17]. With the developement of machine learning in quantum many-body physics [TM16, Wan16, vLH17, CT17, CM17a, TMC$^+$18a], deep learning techniques have also been employed to construct the optimal RG transformations [LW18, KR18] and to uncover the holographic geometry [YYQ18a, HSTT18b, HSTT18a, Has19]. In this work, we further explore the possibility of designing the EHM for interacting quantum field theories using deep learning approaches. By training a flow-based hierarchical generative model [DSB16, KD18] to generate field configurations following the probability distribution specified by the field theory action, the model converges to the optimal EHM with an emergent bulk gravitational description, where neural network parameters and latent variables correspond respectively to the gravity (geometry) and matter degrees of freedom in the holographic bulk. The learned holographic mapping is useful for both sampling and inference tasks. In the sampling task, the generative model is used to propose efficient global-updates for boundary field configurations, which helps to boost the Monte Carlo simulation efficiency of the CFT. In the inference task, the boundary field theory is pulled back to an effective theory of the bulk field, which enables us to probe the emergent dual geometry (on the classical level) by measuring the mutual information in bulk field.

8.2 Renormalization Group and Generative Model

Renormalization group (RG) plays a central role in the study of quantum field theory (QFT) and many body physics. The RG transformation progressively coarse-grains the field configuration to extract relevant features. The coarse-graining rules (or the RG schemes) are generally model-dependent and requires human design. Take the real-space RG[Kad66] for example: for a ferromagnetic Ising model, the RG rule should progressively extract the uniform spin components as the most relevant feature; however for an antiferromagnetic Ising model, the staggered spin components should be extracted instead; if the spin couplings are randomly distributed on the lattice, the RG rule can be more complicated. When it comes to the momentum-space RG[Wil83], the rule becomes to renormalize the low-energy degrees of freedom by integrating out the high-energy degrees of freedom. What is the general designing principle behind all these seemly different RG schemes? Can a machine learns to design the optimal RG scheme based on the model action?

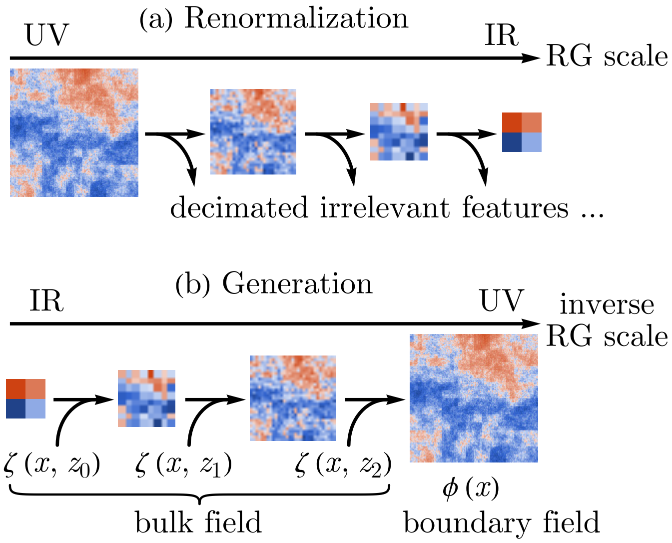

Figure 8.1: Relation between (a) RG and (b) generative model. The inverse RG can be viewed as a generative model that generates the ensemble of field configurations from random sources. The random sources can are supplied at different RG scales (coordinated by z), which can be viewed as a field $\zeta(x,z)$ living in the holographic bulk with one more dimension. The original field $\phi(x)$ will be generated on the holographic boundary.

With these questions in mind, we take a closer look at the RG procedure in a lattice field theory setting. In the traditional RG approach, the RG transformation is invertible due to the information loss at each RG step when the irrelevant features are decimated, as illustrated in Fig. 8.1(a). However, if the decimated features are kept at each RG scale, the RG transformation can be inverted. Under the inverse RG flow, the decimated degrees of freedom $\zeta(x,z)$ are supplied to each layer (step) of the inverse RG transformation, such that the field configuration $\phi(x)$ can be regenerated, as shown in Fig. 8.1(b). Here we assume that the $\phi(x)$ field is defined in a flat Euclidean spacetime coordinated by $x=(x_1,x_2,\cdots)\in\mathbb{R}^d$, then $\zeta(x,z)$ will live on a manifold with one higher dimension, and the extra dimension z corresponds to the RG scale. Given its close analogy to the holographic duality, we may view $\zeta(x,z)$ as the field in the holographic bulk and $\phi(x)$ as the field on the holographic boundary. The inverse RG can be considered as a deep generative model G, which organizes the bulk field $\zeta(x,z)$ to generate the boundary field $\phi(x)$,

$$\phi(x) = G[\zeta(x,z)]. \tag{8.1}$$

The renormalization G^{-1} and generation G procedures are thus unified as the forward and backward maps of a bijective (invertible) map between the boundary and the bulk, known as the EHM.[Qi13, LQ15]

At the first glance, such an information-preserving RG does not seem to have much practical use, because $\zeta(x,z)$ is simply a rewriting of $\phi(x)$, which does not seem to reduce the degrees of freedom. However, since the bulk field $\zeta(x,z)$ represents the irrelevant feature to be decimated under RG, it should look like independent random noise, which contains minimal amount of information. So instead of memorizing the bulk field configuration $\zeta(x,z)$ at each RG scale for reconstruction purpose, we can simply sample $\zeta(x,z)$ from uncorrelated (or weakly correlated) random source and serve them to the inverse RG transformation. Suppose the bulk field $\zeta(x,z)$ is drawn from a prior distribution $P_{\text{prior}}[\zeta]$, the transformation $\phi=G[\zeta]$ will deform the prior distribution to a posterior distribution $P_{\text{post}}[\phi]$ for the boundary field $\phi(x)$,

$$P_{\text{post}}[\phi] = P_{\text{prior}}[\zeta] \left| \det \left(\frac{\delta G[\zeta]}{\delta \zeta} \right) \right|^{-1}, \tag{8.2}$$

where $|\det(\delta_\zeta G)|^{-1}$ is the Jacobian determinant of transformation. In such manner, the objective of the inverse RG is not to reconstruct a particular original field configuration, but to generate an ensemble of field configurations $\phi(x)$, whose probability distribution $P_{\text{post}}[\phi]$ should better match the Boltzmann distribution

$$P_{\text{target}}[\phi] = e^{-S_{\text{QFT}}[\phi]}/Z_{\text{QFT}} \tag{8.3}$$

specified by the action functional $S_{\text{QFT}}[\phi(x)]$ of the boundary field theory, where $Z_{\text{QFT}} = \sum_{[\phi]} e^{-S_{\text{QFT}}[\phi]}$ denotes the partition function.

This setup provides us a theoretical framework to discuss the designing principles of a good RG scheme. We propose two objectives for a good RG scheme (or EHM): the RG transformation should aim at decimating irrelevant features and preserving relevant features, and the inverse RG must aim at generating field configurations matching the target field theory distribution $P_{\text{target}}[\phi]$ in Eq. (8.3). An information theoretic criterion for "irrelevant" features is that they should have minimal mutual information, so the prior distribution $P_{\text{prior}}[\zeta]$ should be chosen to minimize the mutual information between bulk fields at different points, i.e. $\min I(\zeta(x,z) : \zeta(x',z'))$. We will refer to this designing principle as the minimal bulk mutual information (minBMI) principle, which is a general information theoretic principle behind different RG schemes and is independent of the notion of field pattern or energy scale. The close relation between RG and deep learning has been discussed in several early works[MS14, OT17, LTR17]. However, as pointed out in Ref. [KR18, LRHK18], the hierarchical architecture itself cannot guarantee the emergence of RG transformation in a deep neural network. Additional information theoretic principles must be imposed to guide the learning. In light of this observation, Ref. [KR18, LRHK18] proposed the maximal real-space mutual information (maxRSMI) principle, which aims at maximizing the mutual information between the coarse-grained field and the fine-grained field in the surrounding environment. Our minBMI principle is consistent with and is more general than the maxRSMI principle (see Appendix D.1 for detailed discussion about the relation between these two principles).

In the simplest setting, we can hard code the minBMI principle by assigning the prior distribution to the uncorrelated Gaussian distribution,

$$P_{\text{prior}}[\zeta] = \mathcal{N}[\zeta; 0, 1] \propto e^{-\|\zeta\|^2}, \tag{8.4}$$

where $\|\zeta\|^2 = \sum_{x,z} |\zeta(x,z)|^2$. Hence the mutual information vanishes for every pair of points in the holographic bulk. Given the prior distribution, the problem of finding the optimal EHM boils down to training the optimal generative model G to minimize the Kullback-Leibler (KL) divergence between the posterior distribution $P_{\text{post}}[\phi]$ in Eq. (8.2) and the target distribution $P_{\text{target}}[\phi]$ in Eq. (8.3), i.e. $\min \mathcal{L}$ with

$$\begin{aligned} \mathcal{L} &= \mathsf{KL}(P_{\text{post}}[\phi] \parallel P_{\text{target}}[\phi]) \\ &= \mathop{\mathbb{E}}_{\zeta \sim P_{\text{prior}}} S_{\text{QFT}}[G[\zeta]] + \ln P_{\text{prior}}[\zeta] - \ln \det \left(\frac{\delta G[\zeta]}{\delta \zeta} \right), \end{aligned} \tag{8.5}$$

where $\mathbb{E}_{\zeta \sim P_{\text{prior}}}$ denotes the average over the ensemble of ζ drawn from the prior distribution. This fits perfectly to the framework of flow-based generative models[DSB16, KD18] in machine learning, which can be trained efficiently thanks to its tractable and differentiable posterior likelihood. We model the bijective map G by a neural network (to be detailed later) with trainable network parameters. We initiate the sampling from the bulk $\zeta \sim P_{\text{prior}}$ and push the bulk field to the boundary by $\phi = G[\zeta]$, collecting the logarithm of the Jacobian determinant along the way. Given the action $S_{\text{QFT}}[\phi]$, we can evaluate the loss function $\mathcal{L}$ in Eq. (8.5) and back propagate its gradient with respect to the network parameters. We then update the network parameters by stochastic gradient descent. We iterate the above steps to train the neural network. In this way, simply by presenting the QFT action S_{QFT} to the machine, the machine learns to design the optimal RG transformation G by keep probing S_{QFT} with various machine generated field configurations. Thus our algorithm may be called the neural network renormalization group (neural RG)[LW18], which can be implemented using the deep learning platforms such as TensorFlow[AAB$^+$15a].

8.3 Holographic Duality and Classical Approximation

We would like to provide an alternative interpretation of the loss function $\mathcal{L}$ in Eq. (8.5) in the context of holographic duality, which will deepen our understanding of the capabilities and limitations of our approach. Suppose we can sample the boundary field configuration $\phi(x)$ from the target distribution $P_{\text{target}}[\phi]$ and map $\phi(x)$ to the bulk by applying the EHM along the RG direction $\zeta = G^{-1}[\phi]$, the obtained bulk field $\zeta(x,z)$ will follow the distribution

$$\begin{aligned} P_{\text{bulk}}[\zeta] &= P_{\text{target}}[\phi] \det(\delta_\phi G^{-1}[\phi])^{-1} \\ &= Z_{\text{QFT}}^{-1} e^{-S_{\text{QFT}}[G[\zeta]]} \det(\delta_\zeta G). \end{aligned} \tag{8.6}$$

where we have used Eq. (8.3) to express $P_{\text{target}}[\phi]$ in terms of ζ. The normalization of the bulk field probability distribution $\sum_{[\zeta]} P_{\text{bulk}}[\zeta] = 1$ further implies that the QFT partition function Z_{QFT}, which was originally defined on the holographic boundary, can now be written in terms of the bulk field ζ as well

$$Z_{\text{QFT}} = \sum_{[\zeta]} e^{-S_{\text{QFT}}[G[\zeta]] + \ln \det(\delta_\zeta G)}. \tag{8.7}$$

Note that Z_{QFT} is by definition independent of G, we are allowed to sum over all possible G on both sides of Eq. (8.7), which establishes a duality between the following two partition functions

$$Z_{\text{QFT}} = \sum_{[\phi]} e^{-S_{\text{QFT}}[\phi]} \leftrightarrow Z_{\text{grav}} = \sum_{[\zeta, G]} e^{-S_{\text{grav}}[\zeta, G]}, \tag{8.8}$$

with the bulk theory S_{grav} given by

$$S_{\text{grav}}[\zeta, G] = S_{\text{QFT}}[G[\zeta]] - \ln\det(\delta_{\zeta} G). \tag{8.9}$$

By "duality" we mean that Z_{QFT} and Z_{grav} only differ by a proportionality constant (as $Z_{\text{grav}} = \sum_{[G]} Z_{\text{QFT}}$), so they are equivalent descriptions of the same physics theory. $S_{\text{grav}}[\zeta, G]$ describes how the bulk variables ζ (matter field) and the neural network G (geometry) would fluctuate and interact with each other, which resembles a "quantum gravity" theory in the holographic bulk. The bulk has more degrees of freedom than the boundary, as there can be many different choices of ζ and G that leads to the same boundary field configuration $\phi = G[\zeta]$. This is a gauge redundancy in the bulk theory, which covers the diffeomorphism invariance as well as the interchangeable role between matter and spacetime geometry in a gravity theory. At this level, the bulk theory looks intrinsically nonlocal and the geometry can fluctuate strongly.

However, it is usually more desired to work with quantum gravity theories with a classical limit, which describe weak fluctuations (matter fields and gravitons) around a classical geometry. Although not every CFT admits a classical gravity dual, we still attempt to find the *classical approximation* of the dual quantum gravity theory, neglecting the fluctuation of G. Aiming at a classical geometry, we look for the optimal G that maximizes its marginal probability $P_{\text{EHM}}[G] = Z_{\text{grav}}^{-1} \sum_{[\zeta]} e^{-S_{\text{grav}}[\zeta, G]}$ with the bulk matter field ζ traced out. This optimization problem seems trivial, because according to Eq. (8.7), $P_{\text{EHM}}[G] = Z_{\text{QFT}}/Z_{\text{grav}}$ is independent of G. It is understandable that any choice of G is equally likely if we have no preference on the prior distribution $P_{\text{prior}}[\zeta]$ of the bulk matter field, because there is a trade-off between G and P_{prior} that one can always adjust P_{prior} to compensate the change in G. Such a trade-off behavior is fundamentally required by the gauge redundancy in the bulk gravity theory. To fix the gauge, we evoke the minBMI principle to bias the bulk matter field towards independent random noise, such that the classical solution of G will look like a RG transformation, in line with our expectation for a holographic mapping. Choosing a minBMI prior distribution such as Eq. (8.4) and replace $\det(\delta_\zeta G)$ in Eq. (8.9) by $P_{\text{prior}}/P_{\text{post}}$, $P_{\text{EHM}}[G]$ can be cast into

$$P_{\text{EHM}}[G] = \frac{Z_{\text{QFT}}}{Z_{\text{grav}}} \underset{\zeta\sim P_{\text{prior}}}{\mathbb{E}} \frac{P_{\text{target}}[G[\zeta]]}{P_{\text{post}}[G[\zeta]]} \geq \frac{Z_{\text{QFT}}}{Z_{\text{grav}}} e^{-\mathcal{L}}, \tag{8.10}$$

which is bounded by $e^{-\mathcal{L}}$ from below, with $\mathcal{L}$ being the KL divergence between P_{post} and P_{target} as defined in Eq. (8.5). Therefore the objective of maximizing $P_{\text{EHM}}[G]$ can be approximately replaced by minimizing the loss function $\mathcal{L}$, which is no longer a trivial optimization problem. From this perspective, the loss function $\mathcal{L}$ can be approximately interpreted as the action (negative log-likelihood) for the holographic bulk geometry associated to the EHM G. Minimizing the loss function corresponds to finding the classical saddle point solution of the bulk geometry. We will build a flow-based generative model to parameterize G and train the neural network using deep learning approaches. The fluctuation of neural network parameters in the learning dynamics reflects (at least partially) the gravitational fluctuation in the holographic bulk.

At the classical saddle point $G_* = \operatorname{argmin}_G \mathcal{L}$, we may extract an effective theory for the bulk matter field

$$\begin{aligned} S_{\text{eff}}[\zeta] &\equiv S_{\text{grav}}[\zeta, G_*] \\ &= \|\zeta\|^2 + \ln P_{\text{post}}[G_*[\zeta]] - \ln P_{\text{target}}[G_*[\zeta]]. \end{aligned} \tag{8.11}$$

As the KL divergence $\mathcal{L} = \mathsf{KL}(P_{\text{post}} \parallel P_{\text{target}})$ is minimized after training, we expect P_{post} and P_{target} to be similar, such that their log-likelihood difference $\ln P_{\text{post}} - \ln P_{\text{target}}$ will be small, so the effective theory $S_{\text{eff}}[\zeta]$ will be dominated by the first term $\|\zeta\|^2$ in Eq. (8.11), implying that the bulk field ζ will be massive. The small log-likelihood difference further provides kinetic terms (and interactions) for the bulk field ζ, allowing it to propagate on a classical background that is implicitly specified by G_*. In this way, the bulk field will be correlated in general. Even though one of our objectives is to minimize the bulk mutual information as much as possible, the learned

EHM typically cannot resolve all correlations in the original QFT, so the residual correlations will be left in the bulk field ζ as described by the log-likelihood difference in Eq. (8.11). The mismatch between P_{post} and P_{target} may arise from several reasons: first, limited by the design of the neural network, the generative model G may not be expressive enough to precisely deform the prior distribution to the target distribution; second, even if G has the sufficient representation power, the training may not be able to converge to the global minimum; finally and perhaps the most fundamental reason is not every QFT has a classical gravitational dual, the bulk theory should be quantum gravity in general. Taking the classical approximation and ignoring the gravitational fluctuation leads the unresolvable correlation and interaction for the matter field ζ that has to be kept in the bulk.

Nevertheless, our framework could in principle include fluctuations of G by falling back to Z_{grav} in Eq. (8.8). We can either model the marginal distribution $P_{\text{EHM}}[G]$ by techniques like graph generative models, or directly analyze the gravitational fluctuations by observing the fluctuations of neural network parameters in the learning dynamics as mentioned below Eq. (8.10). We will leave these ideas for future exploration. In the following, we will use a concrete example, a 2D compact boson CFT on a lattice, to illustrate our approach of learning the EHM as a generative model and to demonstrate its applications in both the sampling and the inference tasks.

8.4 Application to Complex ϕ^4 Model

We consider a lattice field theory defined on a 2D square lattice, described by the Euclidean action

$$S_{\text{QFT}}[\phi] = -t\sum_{\langle ij\rangle}\phi_i^*\phi_j + \sum_i(\mu|\phi_i|^2 + \lambda|\phi_i|^4), \tag{8.12}$$

where $\phi_i \in \mathbb{C}$ is a complex scalar field defined on each site i of a square lattice and $\langle ij\rangle$ denotes the summation over all nearest neighbor sites. The model has a global $\mathrm{U}(1)$ symmetry, under which the field rotates by $\phi_i \to e^{\mathrm{i}\varphi}\phi_i$ on every site. We choose $\mu = -200 + 2t$ and $\lambda = 25$ to create a deep Mexican hat potential that basically pins the complex field on a circle $\phi_i = \sqrt{\rho}e^{\mathrm{i}\theta_i}$ of radius $\sqrt{\rho} = 2$. In this way, the field theory falls back to the XY-model $S_{\text{QFT}} = -\frac{1}{T}\sum_{\langle ij\rangle}\cos(\theta_i - \theta_j)$ with an effective temperature $T = (\rho t)^{-1}$. By tuning the temperature T, the model exhibits two phases: the low-T algebraic liquid phase with a power-law correlation $\langle\phi_i^*\phi_j\rangle \sim |x_i - x_j|^{-\alpha}$ and the high-T disordered phase with a short-range correlation. The two phases are separated by the Kosterlitz-Thouless (KT) transition. Several recent works[BGM18, ZLW18, RS18, ZEPS18] have focused on applying machine learning method to identify phase transitions or topological defects (vortices). Our purpose is different here: we stay in the algebraic liquid phase, described by a Luttinger liquid CFT, and seek to develop the optimal holographic mapping for the CFT.

We design the generative model G as a bijective deep neural network following the architecture of the neural network renormalization group (neural-RG) proposed by Ref. [LW18]. Its structure resembles the MERA network [Vid07c] as depicted in Fig. 8.2(a). Each RG step contains a layer of disentangler blocks (like CNN convolutional layer) to resolve local correlations, and a layer of decimator blocks (like CNN pooling layer) to separate the renormalized and decimated variables. Given that the spacetime dimension is two on the boundary, we can overlay decimators on top of disentanglers in an interweaving manner as shown in Fig. 8.2(b). Both the disentangler and the decimator are made of three bijective layers: a linear scaling layer, an orthogonal transformation layer and an invertible non-linear activation layer, as arranged in Fig. 8.3 (see Appendix D.2 for more details). They are designed to be invertible, non-linear and U(1)-symmetric transformations, which are used to model generic RG transformations for the complex ϕ^4 model. The bijector parameters are subject to training (training procesure and number of parameters are specified in Appendix D.3). The Jacobian matrix of these transformations are calculable. After each decimator, only one renormalized variable flows to the next RG layer, and the other three decimated variables are positioned into the bulk as little crosses as shown in Fig. 8.2(a) and (b). The entire network constitutes an EHM between the original boundary field $\phi(x)$ and the dual field $\zeta(x,z)$ in the holographic bulk.

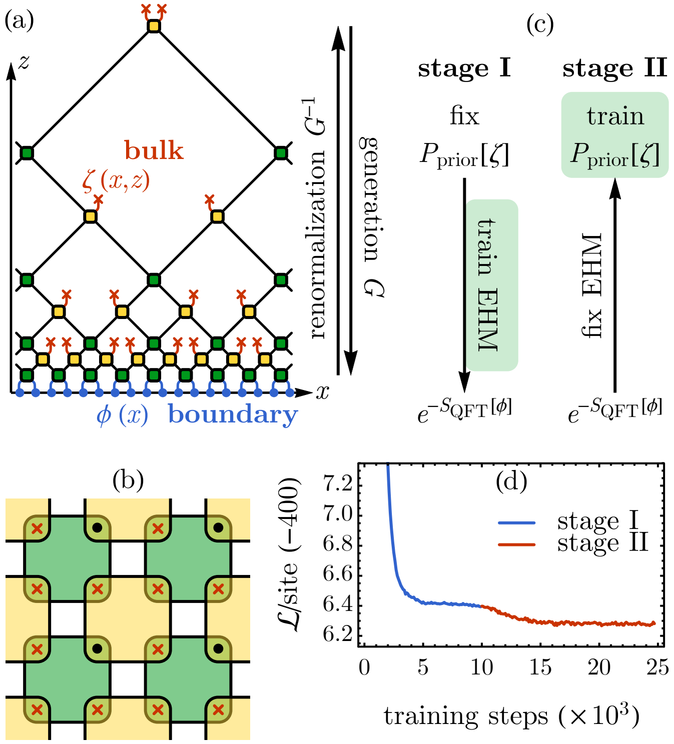

Figure 8.2: (a) Side view of the neural-RG network. x is the spatial dimension(s) and z corresponds to the RG scale. There are two types of blocks: disentanglers (dark green) and decimators (light yellow). The network forms an EHM between the boundary variables (blue dots) and the bulk variables (red crosses). (b) Top view of one RG layer in the network. Disentanglers and decimators interweave in the spacetime (taking two-dimensional spacetime for example). Each decimator pushes the coarse-grained variable (black dot) to the higher layer and leaves the decimated variables (red crosses) in the holographic bulk. (c) The training contains two stages. In the first stage, we fix the prior distribution $P[\zeta]$ to be uncorrelated Gaussian and train the EHM G to bring it to the Boltzmann distribution of the CFT. In the second stage, we learn the prior distribution with the trained EHM held fixed. (d) The behavior of the loss function $\mathcal{L}$ in the two training stages.

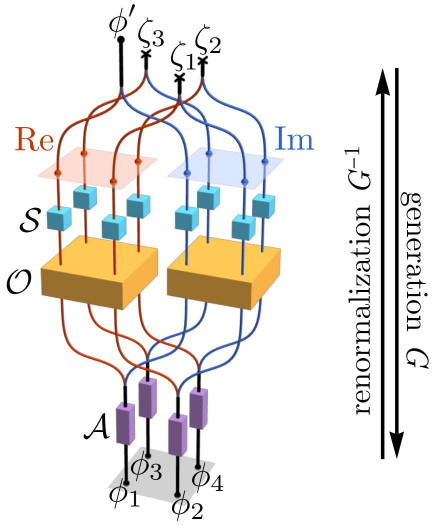

Figure 8.3: Neural network architecture within a decimator block (the disentangler block shares the same architecture). Starting from the renormalized variable ϕ' and the bulk noise $\zeta_{1,2,3}$ as complex variables, the Re and Im channels are first separated, then $\mathcal{S}$ applies the scaling separately to the four variables within each channel and $\mathcal{O}$ implements the $\mathrm{O}(4)$ transformation that mixes the four variables together. $\mathcal{S}$ and $\mathcal{O}$ are identical for Re and Im channels to preserve the $\mathrm{U}(1)$ symmetry. Then the channels merge into complex variables followed by element wise non-linear activation describe by an invertible $\mathrm{U}(1)$-symmetric map $\phi_i \mapsto (\phi_i/|\phi_i|)\sinh|\phi_i|$.

We start with a 32×32 square lattice as the holographic boundary and build up the neural-RG network. The network will have five layers in total. Since the boundary field theory has a gobal $\mathrm{U}(1)$ symmetry, the bijectors in the neural network are designed to respect the $\mathrm{U}(1)$ symmetry (see Fig. 8.3), such that the bulk filed also preserves the $\mathrm{U}(1)$ symmetry. The training will be divided into two stages, as pictured in Fig. 8.2(c). In the training stage I, we fix the prior distribution in Eq. (8.4) and train the network parameters in the generative model G to minimize the loss function $\mathcal{L}$. The training method is outlined below Eq. (8.5). The loss function $\mathcal{L}$ decays with training steps, whose typical behavior is shown in Fig. 8.2(d). We will discuss the stage II training later.

We perform the stage I training for several neural networks at different temperatures T separately, i.e. we use $S_{\mathrm{QFT}}[\phi]$ of different parameters to train different neural networks. After training, each neural network can generate configurations of the boundary field ϕ from the bulk

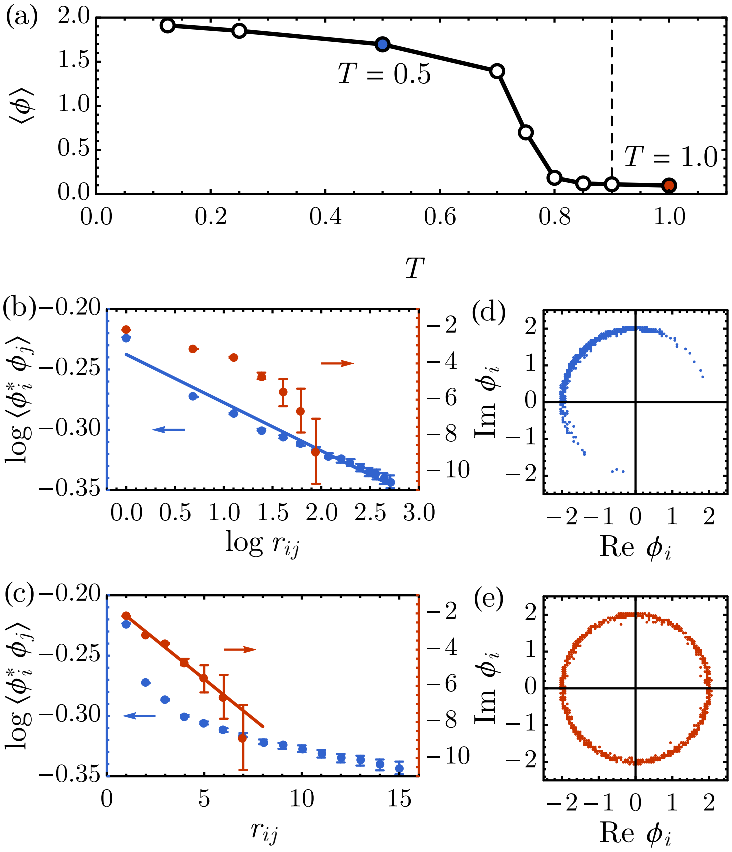

Figure 8.4: Performance of the trained EHM for the complex ϕ^4 theory. (a) Order parameter $\langle \phi \rangle$ v.s. temperature T. Different models are trained separately at different temperature. For finite-sized system, $\langle \phi \rangle$ crosses over to zero around the KT transition. Correlation function $\langle \phi_i^* \phi_j \rangle$ scaling in log-log plot (b) and log-linear plot (c). Distribution of ϕ_i in a single sample generated by the neural network trained in (d) the algebraic liquid phase and (e) the disordered phase.

uncorrelated Gaussian field ζ efficiently. To test how well these generative models work, we measured the order parameter $\langle\phi\rangle$ and the correlation function $\langle\phi_i^*\phi_j\rangle$ using the field configurations generated by the neural network. Although the order parameter $\langle\phi\rangle$ is expected to vanish in the thermodynamic limit, for our finite-size system, it is not vanishing and can exhibit a crossover around the KT transition, as shown in Fig. 8.4(a). The cross over temperature $T \simeq 0.9$ agrees with the previous Monte Carlo study [Ols95, HP97, Has05] of the KT transition temperature $T_{\text{KT}} = 0.8929$ in the two-dimensional XY model. We measure the correlation function $\langle\phi_i^*\phi_j\rangle$ at two different temperatures: one at $T = 0.5$ in the algebraic liquid phase, one at $T = 1.0$ in the disordered phase. We plot the two-point function $\langle\phi_i^*\phi_j\rangle$ as a function of the Euclidean distance $r_{ij} \equiv |x_i - x_j|$ (on the square lattice) in both the log-log scale as Fig. 8.4(b) and the log-linear scale as Fig. 8.4(c). The comparison shows that the correlation function in the algebraic liquid (or the disordered) phase fits better to the power-law (or the exponential) decay. Fig. 8.4(d) shows the statistics of ϕ_i in one sample generated by the machine trained in the algebraic liquid phase. It exhibits the “spontaneous symmetry breaking” behavior due to the finite-size effect, although accumulating over multiple samples will restore the $\mathrm{U}(1)$ symmetry. However, similar plot Fig. 8.4(e) in the disordered phase respects the $\mathrm{U}(1)$ symmetry in every single sample. Based on these tests, we can conclude that the neural network has learned to generate field configurations $\phi(x)$ that reproduce the correct physics of the complex ϕ^4 model. The trained generative model G maps an almost uncorrelated bulk field ζ to a correlated boundary field ϕ, and vice versa, therefore G provides a good EHM for the ϕ^4 theory.

The learned EHM can be useful in both the backward and forward directions. The backward mapping from bulk to boundary provides efficient sampling of the CFT configurations, which can be used to boost the Monte Carlo simulation of the CFT. The forward mapping from boundary to bulk enables direct inference of bulk field configurations, allowing us to study the bulk effective theory and to probe the bulk geometry. Let us first discuss the sampling task. The EHM establishes a mapping between the massive bulk field ζ and the massless boundary field ϕ.

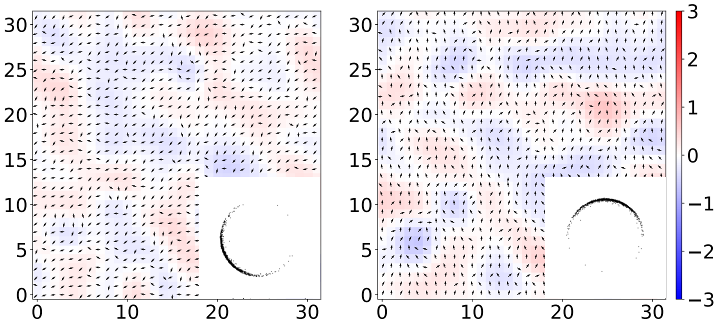

Figure 8.5: The boundary field configuration ϕ before (left) and after (right) a local update in the most IR layer of the bulk field ζ. The complex field ϕ_i is represented by the small arrow on each site. The background color represents the vorticity. The inset shows the distribution of ϕ_i in the complex plane.

The bulk field admits efficient sampling in terms of local update, because the it is uncorrelated (or short-range correlated). Local updates in the bulk gets maps to global updates on the boundary, which allows us to sample the critical boundary field efficiently, minimizing the effect of critical slowdown. To demonstrate this, we first generate one field configuration as shown in the left panel of Fig. 8.5. Then we push this field configuration back into the bulk field ζ. We tweak the bulk field in the most IR degree of freedom by adding a small random Gaussian number. The variance of the added Gaussian variable is five times smaller than the variance of the bulk field degree of freedom. Under the EHM, we observe a global change of the boundary field configuration as shown in Fig. 8.5. It is interesting to note that the change of the IR bulk field basically induces a global $\mathrm{U}(1)$ rotation of ϕ_i (see the insets of Fig. 8.5), which corresponds the "Goldstone mode" associated to the "spontaneous symmetry breaking" in the fine-sized system, showing that the machine can identify the order parameter as the relevant IR degrees of freedom without prior knowledge about low-energy modes of the system. We also check that the Hamiltonian Monte Carlo sampling in the bulk converges much faster compared to applying the same algorithm on the boundary (see Appendix D.4 for more evidences). In connection to several recent works, our neural-RG architecture can be integrated to self-learning Monte Carlo approaches[LSQ$^+$16, AK16, HW17, LQMF17, NSQ$^+$17, TT17, NOT18] to boost the numerical efficiency in simulating CFTs. The inverse RG transformation can also be used to generate super-resolution samples[EBM18] for finite-size extrapolation of thermodynamic observables.

Now let us turn to the inference task. We can use the optimal EHM to push the boundary field back into the bulk and investigate the effective bulk theory $S_{\text{eff}}[\zeta]$ induced by the boundary CFT. As analyzed below Eq. (8.11), the mismatch between P_{post} and P_{target} will give rise to the residual correlation (mutual information) of the bulk matter field, which can be used to probe the holographic bulk geometry. Assuming an emergent locality in the holographic bulk, the expectation is that the bulk effective theory $S_{\text{eff}}[\zeta]$ will take the following form in the continuum limit,

$$S_{\text{eff}}[\zeta] = \int_{\mathcal{M}} g^{\mu\nu}\partial_\mu\zeta^*\partial_\nu\zeta + m^2|\zeta|^2 + u|\zeta|^4 + \cdots , \tag{8.13}$$

which describes the bulk field ζ on a curved spacetime background $\mathcal{M}$ equipped with the metric tensor $g^{\mu\nu}$. Strictly speaking, ζ is not a single field but contains a tower of fields corresponding to different primary operators in the CFT. We choose to focus on the lightest component and model it by a scalar field, as it will dominate the bulk mutual information at large scale. Because the bulk field excitation is massive and can not propagate far, we expect the mutual information between the bulk variables at two different points to decay exponentially with their geodesic distance in the bulk. Following this idea, suppose $\zeta_i = \zeta(x_i, z_i)$ and $\zeta_j = \zeta(x_j, z_j)$ are two bulk field variables, then their distance $d(\zeta_i : \zeta_j)$ can be inferred from their mutual information $I(\zeta_i : \zeta_j)$ as follows

$$d(\zeta_i : \zeta_j) = -\xi \ln \frac{I(\zeta_i : \zeta_j)}{I_0}, \tag{8.14}$$

where the correlation length ξ and the information unit I_0 are global fitting parameters.

To estimate the mutual information among bulk field variables, we take a quadratic approximation of the bulk effective action $S_{\text{eff}}[\zeta] \simeq \sum_{ij} \zeta_i^* K_{ij} \zeta_j = \zeta^\dagger K \zeta$, ignoring the higher order interactions of ζ for now. This amounts to relaxing the prior distribution of the bulk field ζ to a correlated Gaussian distribution

$$P'_{\text{prior}}[\zeta] = \frac{1}{\sqrt{\det(2\pi K^{-1})}} e^{-\zeta^\dagger K \zeta}. \tag{8.15}$$

The kernel matrix K is carefully designed to ensure positivity and bulk locality (see Appendix D.5 for more details). To determine the best fit of K, we initiate the stage II training to learn the prior distribution with the EHM fixed at its optimal solution obtained in the stage I training, as illustrated in Fig. 8.2(c). We use the reparametrization trick [KW13] to sample the bulk field ζ from the correlated Gaussian in Eq. (8.15), then ζ is pushed to the boundary by the fixed EHM to evaluate the loss function $\mathcal{L}$ in Eq. (8.5), and the gradient signal can back-propagate to train the kernel K. As we relax the Gaussian kernel K for training, we can see that the loss function will continue to drop in the stage II, as shown in Fig. 8.2(d). This indicates that the Gaussian model is learning to capture the residual bulk field correlation (at least partially), such that the overall performance of generation gets improved. One may wonder why we do not train the generative model G and bulk field distribution $P_{\text{prior}}[\zeta]$ jointly. This is because there is a trade-off between these two objectives. For example, one can weaken the disentanglers in G and push more correlation to the bulk field distribution $P_{\text{prior}}[\zeta]$. Such trade-off will undermine our objective of minimizing bulk mutual information in training a good EHM, therefore the two training stages should be separated, or at least assigned very different learning rates. Intuitively, the machine learns the background geometry in the stage I training and the bulk field theory (to the quadratic order) in the stage II training. The trade-off between the two training stages resembles the interchangeable roles between matter and spacetime geometry in a gravity theory.

After the stage II training, we obtain the fitted kernel matrix K. The mutual information $I(\zeta_i : \zeta_j)$ can be evaluated from

$$I(\zeta_i : \zeta_j) = -\frac{1}{2} \ln \Big(1 - \frac{\langle \zeta_i^* \zeta_j \rangle}{\langle \zeta_i^* \zeta_i \rangle \langle \zeta_j^* \zeta_j \rangle} \Big), \tag{8.16}$$

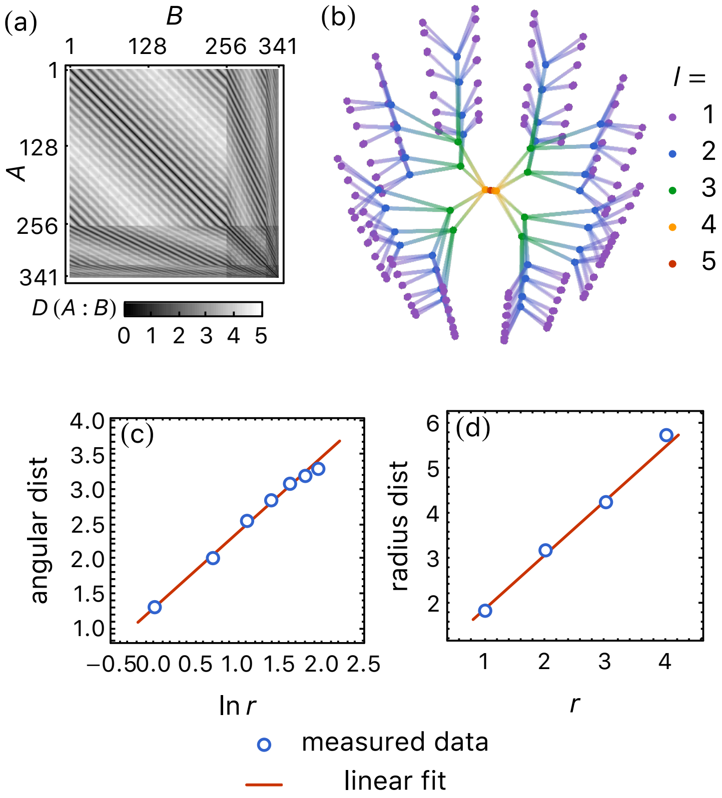

Figure 8.6: (a) Distance matrix $D(A : B)$, indexed by the decimator indices A, B, obtained based on Eq. (8.17). (b) Visualization of the bulk geometry by multidimensional scaling projected to the leading three principle dimensions. Each point represent a decimator in the neural network, colored according to layers from UV to IR. The neighboring UV-IR links are add to guide the eye. Subplot (c) & (d) shows the distance scaling along (c) the radius and (b) the angular direction.

where the bulk correlation $\langle\zeta_i^*\zeta_j\rangle = (K^{-1})_{ij}$ is simply given by the inverse of the kernel matrix K. Then we can measure the holographic distance $d(\zeta_i : \zeta_j)$ between any pair of bulk variables ζ_i and ζ_j following Eq. (8.14). To probe the bulk geometry, we further define the distance between two decimators A and B to be the average distance between all pairs of bulk variables separately associated to them,

$$D(A : B) = \underset{\zeta_i\in A,\zeta_j\in B}{\text{avg}} d(\zeta_i : \zeta_j). \tag{8.17}$$

The result is presented in Fig. 8.6(a). To visualize the bulk geometry qualitatively, we perform a multidimensional scaling to obtain a three-dimensional embedding of the decimators in Fig. 8.6(b). One can see a hyperbolic geometry emerges in the bulk. To be more quantitative, we label each decimator by three coordinates (x^1,x^2,z), where $x = (x^1,x^2)$ denotes its center position projected to the holographic boundary and $z = 2^l$ is related to its layer depth l (ascending from UV to IR). We found that the measured distance function follows the scaling behavior

$$\begin{aligned} D(x^1,x^2,z : x^1+r,x^2,z) &\propto \ln r,\\ D(x^1,x^2,z : x^1,x^2,z+r) &\propto r, \end{aligned} \tag{8.18}$$

as demonstrated in Fig. 8.6(c)&(d). These scaling behaviors agree with the geometry of a three-dimensional hyperbolic space H^3, which corresponds to the AdS_3 spacetime under the Wick rotation of the time dimension. This indicates that the emergent bulk geometry is indeed hyperbolic at the classical level.

Our result demonstrates that the Luttinger liquid CFT can be approximately dual to a massive scalar fields on AdS_3 background geometry. The duality is only approximate because we have assumed a classical geometry in the bulk, ignoring all the gravitational fluctuations. In $\text{AdS}_3/\text{CFT}_2$ correspondence, the bulk gravitational coupling $G_N = 3\ell/2c$ is inversely proportional to the central charge c of the CFT.[BH86] The Luttinger liquid CFT has a relatively small central charge $c = 1$ and hence a large gravitational coupling in the bulk, so we should not expect a classical dual description. It would be more appropriate to consider holographic CFTs which admit classical duals. However, our current method only applies to lattice field theories of bosons with explicit action functionals, which prevent us to study interesting holographic CFTs. Generalizing the neural RG approach to involve fermions and gauge fields and to work with continuous spacetime will be important directions for future development.

8.5 Summary and Discussions

In conclusion, we introduced the neural RG algorithm to allow automated construction of EHM by machine learning instead of human design. Previously, the EHM was only designed for free fermion CFT. Using machine learning approaches, we are able to develop more general EHMs that also apply to interacting field theories. Given the QFT action as input, the machine effective digests the information contained in the action and encode it into the structure of the EHM network, which represents the emergent holographic geometry. Our result provides a concrete example that the holographic spacetime geometry can emerge as the optimal generative network of a quantum field theory.[DZ18] The obtained EHM simultaneously provides an information-preserving RG scheme and a generative model to reproduce the QFT, which could be useful for both inference and sampling tasks.

However, as a version of EHM, our approach also bares the limitations of EHM. By construction, the bulk geometry is discrete and classical, such that the model can not resolve the sub-AdS geometry and can not capture gravitational fluctuations. Recent development of neural ordinary differential equation approaches[CRBD18, ZEW18, GCB$^+$18] are natural ways to extend our flow-based generative model to the continuum limit. Continuous formulation of real-space RG has been discussed in the context of gradient flows[Fuj16, AF18, CHN18] and trivializing maps[L10], where the RG flow equations are human-designed. Our research may pave way for machine-learned RG flow equations for continuous holographic mappings. Our formalism also allows the inclusion of gravitational fluctuations in principle, by relaxing optimization to allow superposition of different EHMs. We conjecture that the bulk gravitational fluctuation could be partially captured by the fluctuation of neural network parameters. The learned EHM provides us a starting point to investigate the corrections on top of the classical geometry approximation, which may enable us to go beyond holographic CFTs and study the quantum gravity dual of generic QFTs. Another feature of EHM is that it is a one-to-one mapping of field configurations (operators) between bulk and boundary, while in holographic duality, a local bulk operator can be mapped to multiple boundary operators in different regions. A resolution[ADH15, PYHP15b] of the paradox is that the non-unique bulk-boundary correspondence only applies to the low-energy freedoms in the bulk, which can be encoded on the boundary in a redundant and error-correcting manner. The bidirectional holographic code (BHC)[YHQ16, QYY17] was proposed as an extension of the EHM to capture the error-correction property of the holographic mapping. Extending our current network design to realize machine-learned BHC will be another open question for future research.

8.6 Acknowledgement

This work is supported by a startup fund from UC San Diego. Chapter 6, in full, is a reprint of the material as it appears in Hong-Ye Hu, Shuo-Hui Li, Lei Wang, Yi-Zhuang You, Physical Review Research 2, 023369. The dissertation author was the primary investigator and author of this paper.

Chapter 9

Summary and outlook

Quantum mechanics is one of the most successful and striking theories in physics. It predicts atomic particles can have exotic properties, such as quantum entanglement, that cannot be described by any classical local theory. This phenomenon dramatically increases the complexity of nature in two ways: first, it is really hard to fully characterize the quantum many-body states in a black box and it usually takes exponentially many measurements; second, it is really hard to simulate quantum many-body states on a classical computer, and it usually takes exponentially classical resources, such as time and memory. These two tasks are essential in both theoretical and experimental physics.

To tackle these problems, I investigate two efficient (classical) representations of quantum many-body states in this thesis: 1) the classical shadow representation and 2) hierarchical representation. In chapter 5 and 6, I introduced classical shadow tomography with chaotic dynamics. My new proposal explicitly shows how the entanglement created in quantum dynamics can help with quantum measurements, and it is the first classical shadow tomography protocol does not rely on the group structure of random measurements. In chapter 7, I introduced spectrum bifurcation renormalization group, and showed strongly disordered quantum system can be approximated solved by a hierarchical Clifford circuit, therefore those systems can be efficiently simulated classically. Especially I applied this technique to study the interplay between strong disorder and topological critical point. In chapter 8, I introduced machine learning holographic mapping. I showed the invertible renormalization group scheme, exact holographic mapping and flow-based generative model in unsupervised learning are related to each other. Especially, I introduced an information theoretical goal for machines to learn optimal (real-space) renormalization group solely based on the action or Hamiltonian description of physical systems. And I showed this new method can not only find interesting holographic dual theory of critical field theory on the boundary but also boost the efficiency of classical simulation of critical systems.

Statistical machine learning theory provides a new scope in understanding quantum many-body states and still have many interesting open questions. Despite many practical applications of machine learning to quantum many-body states, a theoretical understanding of what properties of quantum systems are hard to learn is still opening [Aar07, Roc18, RAS^{+}19]. In recent literature [HKT^{+}21], the authors showed the provable advantage of machine learning models in solving quantum many-body problems. It is an interesting open question about how the results generalized to self-supervised or unsupervised machine learning settings, where less or no human prior knowledge is given to the machine. In the future, I am interested in investigating the unsupervised machine learning theory and its implication on quantum many-body physics.

In addition, there lacks a standard test and database for machine learning quantum many-body problems. Over the past a few years, there are many machine learning models proposed for quantum phase classifications, quantum controls, and etc. But most of them are specifically designed for one task in quantum many-body physics. The lack of a standard test makes it hard to compare the performance of different models and their generalization performance in real-world applications. Therefore, I propose to design a standard test and dataset for quantum many-body problems. Like the ImageNet challenge [RDS^{+}15] in computer vision, I envision a standard dataset for quantum many-body states should contain many quantum wave functions that gathered from different quantum phase of matter. Each quantum wave function can be saved as a set of classical shadows, which contain unbiased information of underlying quantum states. A standard dataset challenge would benefit the whole field of machine learning quantum many-body physics in many ways. It will provide a fair comparison among machine learning models, and their generalization performance. In analog to the natural language applications and BERT [DCLT19], a standard dataset can be used by machines to learn the representations of quantum states and perform self-supervised learning and few-shot learning.

A successfully pretrained model for quantum many-body states will soon find applications in the near-term quantum machines. Over the past a few years, we have witnessed the great process made in programmable quantum simulators (PQS). We will soon enter a region where interesting quantum phase of matter can be realized in PQS but doing a full quantum state tomography is impossible [SHP$^+$21, SLK$^+$21, MSE$^+$21]. Therefore, it is an interesting and challenging questions to predict properties of those quantum states in PQS and understanding their phases by classical measurement readouts. The pretrained machine learning model can help with those important tasks with fewer experimental data.

Last but not the least, quantum systems can be naturally viewed as (quantum) generative models, as under each measurement, a random outcome is generated based on the Born's rule. It is interesting to investigate whether quantum properties, such as quantum non-locality and contextuality, can enhance neural network generative models. In a recent work [GAW$^+$21] and my work in preparation, we study the relative expressive power between classical recurrent neural networks and quantum recurrent models based on Gaussian operations. We show there is an unconditional quantum advantage in quantum sequence generative models. More importantly, we can interpret the quantum advantage is inherited from quantum contextuality. In addition, we demonstrate this quantum advantage through the popular natural language translation tasks, implying the potential impact of quantum recurrent neural networks for real-world applications. In the future, I want to investigate whether other intrinsic quantum many-body effect, such as topological quantum effect, can also enhance the performance of quantum neural networks. I am also interested in the relation between barren plateau problem and information scrambling in the training of quantum neural networks [PNGY21, SMM$^+$22, SZYZ20].

Appendix A

Hamiltonian-driven shadow tomography of quantum states

A.1 Diagrammatic approach towards quantum channel $\mathcal{M}$

In the shadow tomography of quantum states, for each experiment, the state ρ is first evolved by a random unitary operator $U = e^{-iHt} = V\Lambda(t)V^\dagger$ generated from GUE random Hamiltonian H, then measured in the computational basis. The measurement will prepare a classical snapshot of quantum system $\hat{\sigma} = U^\dagger|b\rangle\langle b|U$ with $b \in \{0,1\}^N$ labeling the measurement outcome. We can view the average mapping from the quantum state ρ to its classical snapshots as a quantum channel $\mathcal{M}$,

$$\begin{aligned}
\mathcal{M}(\rho) &= \mathbb{E}[U^\dagger|b\rangle\langle b|U] \\
&= \mathop{\mathbb{E}}_{V,\Lambda} \sum_{b\in\{0,1\}^N} V\Lambda(t)^\dagger V^\dagger|b\rangle\langle b|V\Lambda(t)V^\dagger\langle b|V\Lambda(t)V^\dagger\rho V\Lambda(t)^\dagger V^\dagger|b\rangle \\
&= \mathop{\mathbb{E}}_{V,\Lambda} \begin{tikzcd}
{} & V^\dagger \arrow[l] & \Lambda(t) \arrow[l] & V \arrow[l] & {} \arrow[l] & {} \arrow[d] & V^\dagger \arrow[l, dash] & \Lambda^\dagger(t) \arrow[l] & V \arrow[l] & {} \arrow[l] \\
& & & & {|b\rangle\langle b|} \arrow[u, dash] \arrow[r, dotted, dash] & {|b\rangle\langle b|} \arrow[d, dash] & & & & \rho \arrow[u, dash] \\
{} \arrow[r] & V \arrow[r] & \Lambda^\dagger(t) \arrow[r] & V^\dagger \arrow[r, dash] & {} \arrow[u] & {} \arrow[r] & V \arrow[r] & \Lambda(t) \arrow[r] & V^\dagger \arrow[r, dash] & {} \arrow[u]
\end{tikzcd} ,
\end{aligned} \tag{A.1}$$

in the last line we use the tensor network to represent the equation, the dashed line means summing over the bit-string b, and the arrow indicates the order that operators are multiplied together. As we can see this is a 4-fold twirl of Haar random matrix V, and $\Lambda(t)$ is a diagonal matrix with $\Lambda_{nn}(t) = e^{-iE_n t}$ with E_n being eigen energies of H. It can be further simplified as

$$\mathcal{M}(\rho) = \mathop{\mathbb{E}}_{\Lambda} \sum_{\sigma,\tau\in S_4} \mathrm{Wg}[\sigma\tau^{-1}g_0] A[\sigma] B[\tau], \tag{A.2}$$

where σ, τ are permutations from permutation group S_4, $\mathrm{Wg}[g]$ is the Weingarten function[Wei78] of the permutation group element g, $g_0 = (1,3)(2,4)$ is a fixed permutation to match the tensor network connection, and $A[\sigma]$, $B[\tau]$ are defined as:

$$A[\sigma] =$$

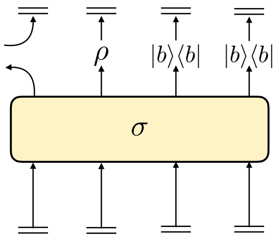

$$, \; B[\tau] =$$

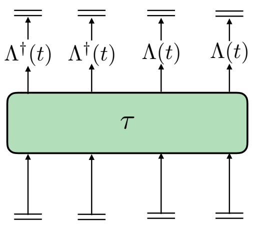

$$. \tag{A.3}$$

In the above tensor diagram, the double line means the periodic boundary condition for top and bottom legs. After specifying the choice of permutation σ and τ, one can make a connection according to σ and τ in the yellow and green block, and those tensor diagram can be evaluated.

The average over $\Lambda_{nn'}(t) = e^{-iE_n t}\delta_{n,n'}$ can be calculated using the joint probability distribution of eigen energies, and the spectral form factor of GUE matrices[CHLY17, YG18].

$$P_{\mathrm{GUE}}[E] \propto \prod_{m>m'} \left(E_m - E_{m'}\right)^2 e^{-\frac{D}{2}\sum_m E_m^2}. \tag{A.4}$$

The summation of $\sigma, \tau \in S_4$ permutation group and averaging over eigen energies can be carries out and gives

$$\begin{aligned}\mathcal{M}(\rho) &= \mathcal{M}(\mathbb{1}/D+\rho_o+\rho_d) \\ &= \frac{\mathbb{1}}{D}+\frac{\rho_o}{\alpha_D(t)}+\frac{\rho_d}{\beta_D(t)},\end{aligned} \tag{A.5}$$

where $\mathbb{1}$ stands for the identity matrix, ρ_o is the off-diagonal part of ρ, and ρ_d is the traceless diagonal part of ρ. The coefficients $\alpha_D(t)$, and $\beta_D(t)$ are defined as

$$\begin{aligned}\alpha_D(t) &= \Big(\frac{1}{D+1}-\lambda_D(t)\Big)^{-1}, \\ \beta_D(t) &= \Big(\frac{1}{D+1}+D\lambda_D(t)\Big)^{-1}, \\ \lambda_D(t) &= \frac{(Dr^2(t)+r(2t))^2-4r^2(t)}{(D+3)(D^2-1)},\end{aligned} \tag{A.6}$$

and $r(t) = J_1(2t)/t$ with J_1 being the Bessel function of the first kind.

A.2 Diagrammatic approach towards variance calculation

In the main text, we have shown that the efficiency of using shadow tomography to predict physical observables $o = F(O,\rho)$ is closely related to the variance of $\hat{o} = F(O,\hat{\rho})$, where $\hat{\rho}$ is the classical shadow, and F is a function that depends on density matrix ρ, and observable O. In general, we have

$$\mathrm{Var}(\hat{o}) \leq \mathbb{E}[\hat{o}^2], \tag{A.7}$$

and we define a $V[\cdots]$ linear function of the double-operator $O^{\otimes 2}$ as

$$\mathbb{E}[\hat{o}^2] = \mathbb{E}[\mathrm{Tr}(O\hat{\rho})^2] = V[O^{\otimes 2}]. \tag{A.8}$$

Diagrammatically, the above equation can be expressed as

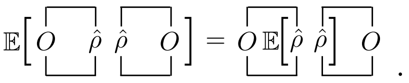

(A.9)

The $V[\cdots]$ function that take $O^{\otimes 2}$ as input can be expressed diagrammatically as

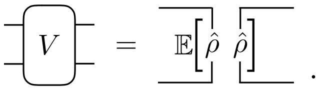

(A.10)

The calculation of Eq.A.10 involves 6-design random unitaries and we showed that is involves 8 diagrams:

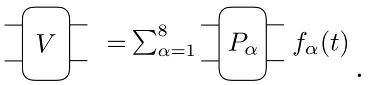

(A.11)

We name them P-diagrams, which are defined as:

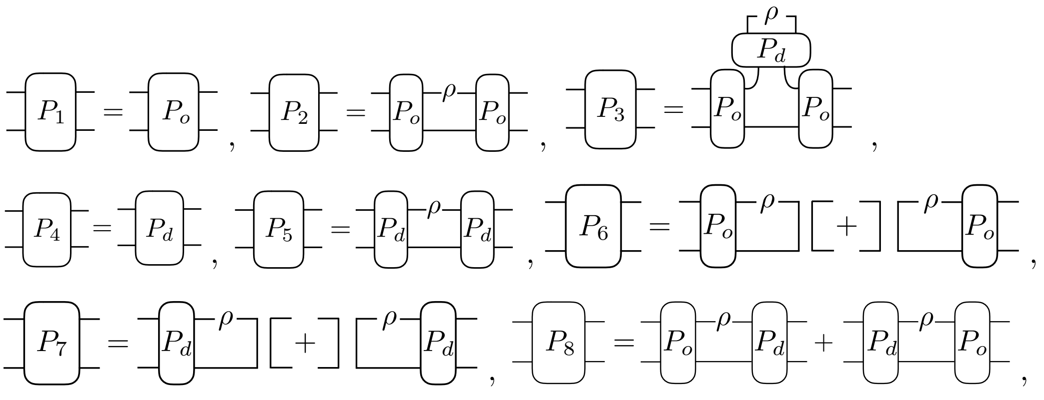

(A.12)

where P_o is a projection operator that takes the off-diagonal part of an operator, P_d is a projection operator that takes the traceless diagonal part of an operator. The *form factors* $f_\alpha(t)$ are given by (with $r(t) = J_1(2t)/t$)

$$\begin{aligned}
f_1(t) &= \frac{1-3r^4(t)+2r^6(t)}{(1-r^4(t))^2} \\
f_2(t) &= 2f_1(t) \\
f_3(t) &= \frac{2D(r^4(t)-r^6(t))}{(1-r^4(t))^2} \\
f_4(t) &= \frac{1}{1+Dr^4(t)} \\
f_5(t) &= \frac{2+D^2r^6(t)+6Dr^4(t)}{(1+Dr^4(t))^2} \\
f_6(t) &= \frac{r^2(t)-r^6(t)}{D(1-r^4(t))} \\
f_7(t) &= \frac{r^6(t)}{1+Dr^4(t)} \\
f_8(t) &= \frac{2D(r^4(t)-r^6(t))+2}{(1+Dr^4(t))(1-r^4(t))}.
\end{aligned} \tag{A.13}$$

The shapes of these form factors are shown in Fig. A.1.

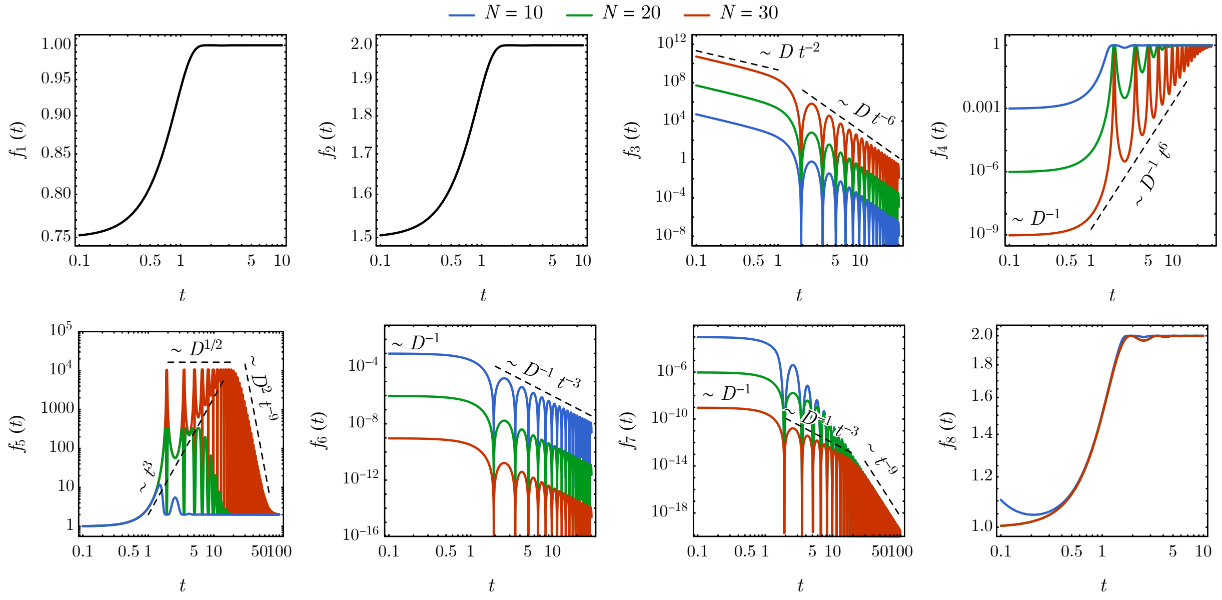

Figure A.1: Form factors $f_\alpha(t)$ and their asymptotic behaviors. Colored curves corresponds to different $D = 2^N$ with $N = 10, 20, 30$ (blue, green, red). For form factors independent of D, the curve is plotted in black.

Collecting the P-diagrams, $V[O^{\otimes 2}]$ can be decomposed into five parts:

$$V[O^{\otimes 2}] =$$

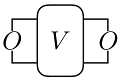

$$= V[O_o^{\otimes 2}] + V[O_d^{\otimes 2}] + V[O_o \otimes O_d] + 2V[\mathbb{1} \otimes O_o] + 2V[\mathbb{1} \otimes O_d] \tag{A.14}$$

where $V[O_o^{\otimes 2}]$ involves diagram P_1, P_2, and P_3; $V[O_d^{\otimes 2}]$ involves diagram P_4, and P_5; $V[\mathbb{1} \otimes O_o]$ involves P_6; $V[\mathbb{1} \otimes O_d]$ involves P_7; $V[O_o \otimes O_d]$ involves P_8. Here O_d, O_o stand for the traceless diagonal part and the off-diagonal part of the observable O. More explicitly, we can write down:

$$\begin{aligned}
&V[O_o^{\otimes 2}] = f_1(t)\,\mathrm{Tr}(O_o^2) + f_2(t)\,\mathrm{Tr}(O_o^2\rho) + f_3(t)\,\mathrm{Tr}(O_o^2\rho_d) \\
&V[O_d^{\otimes 2}] = f_4(t)\,\mathrm{Tr}(O_d^2) + f_5(t)\,\mathrm{Tr}(O_d^2\rho) \\
&V[\mathbb{1} \otimes O_o] = f_6(t)\,\mathrm{Tr}(O_o\rho)\,\mathrm{Tr}(O) \\
&V[\mathbb{1} \otimes O_d] = f_7(t)\,\mathrm{Tr}(O_d\rho)\,\mathrm{Tr}(O) \\
&V[O_o \otimes O_d] = f_8(t)(\mathrm{Tr}(O_dO_o\rho) + \mathrm{Tr}(O_oO_d\rho)).
\end{aligned} \tag{A.15}$$

We analyze the asymptotic behavior of form factor $f_\alpha(t)$ in the following, and provide numerical evidence in the following appendix. Given $r(t) = J_1(2t)/t$ and the envelop behavior of the Bessel function $J_\nu(t) \sim \sqrt{2/(\pi t)}$, we can obtain the following asymptotic behavior

$$r(t) \overset{t\to 0}{=} 1 - \frac{1}{2}t^2, \quad r(t) \overset{t\to\infty}{\simeq} t^{-3/2}. \tag{A.16}$$

Off-diagonal terms: The off-diagonal terms involve P_1, P_2, and P_3, corresponding to the form factors f_1, f_2 and f_3. As shown in Fig. A.1, the form factors $f_1(t)$ and $f_2(t)$ are bounded between their $t = 0$ and $t \to \infty$ values,

$$\begin{aligned}
&\lim_{t\to 0} f_1(t) = 0.75, \ \lim_{t\to\infty} f_1(t) = 1, \\
&\lim_{t\to 0} f_2(t) = 1.5, \ \lim_{t\to\infty} f_2(t) = 2.
\end{aligned} \tag{A.17}$$

The early-time divergence behavior of off-diagonal terms comes from f_3. As we can see,

$$f_3(t) \overset{t\to 0}{=} \frac{1}{2}Dt^{-2}. \tag{A.18}$$

So for short time, $f_3(t)$ diverges as t^{-2} as shown in Fig. A.1. The long time behavior of $f_3(t)$ is

$$f_3(t) \overset{t\to\infty}{\simeq} Dt^{-6}. \tag{A.19}$$

Diagonal traceless terms: The diagonal traceless terms involve two P-diagrams, P_4 and P_5, corresponding to f_4 and f_5 form factors.

For $f_4(t)$, the form factor reaches local maximal $f_4(t_k) = 1$ at $t_k = x_k/2$, where x_k is the kth zero point of the Bessel function $J_1(x)$, see Fig. A.1. And in the window between those points, $f_4(t) \sim D^{-1}t^6$. We name this phenomenon as *scrambling beats*. And for the long time,

$$f_4(t) \overset{t\to\infty}{\simeq} 1 - Dt^{-6}. \tag{A.20}$$

Therefore the beats behavior of $f_4(t)$ will last for a characteristic time $T_4 \sim D^{1/6}$.

For $f_5(t)$, in the long time, it becomes

$$f_5(t) \overset{t\to\infty}{\simeq} 2 + 2Dt^{-6} + D^2t^{-9}. \tag{A.21}$$

We can define two characteristic times $T_5^{(1)} \sim D^{1/6}$ and $T_5^{(2)} \sim D^{2/9}$. In Fig. A.1, we show the asymptotic behavior of $f_5(t)$. We can see the oscillation behavior separates to two stages. In the first stage, it will peak at a constant value that scales as $D^{1/2}$. This stage ends around $T_5^{(1)} \sim D^{1/6}$. Then, in the second stage, the peak of $f_5(t)$ will decay and eventually reaches its long-time value. Time $T_5^{(2)} \sim D^{2/9}$ characters the total time for $f_5(t)$ to reach its long-time value.

We now analyze how the peak value scales with D in the first stage. We observe the peak happens in vicinity of $t_k = x_k/2$. First, we Talyor expand $r(t)$ around t_k as

$$\begin{aligned} r(t_k+\delta t) &= a_k \delta t \\ f_5(t_k+\delta t) &= \frac{2+D^2 a_k^6 \delta t^6 + 6 D a_k^4 \delta t^4}{(1+D a_k^4 \delta t^4)^2} \end{aligned} \tag{A.22}$$

By analyzing the behavior of $f_5(t_k+\delta t)$, we can find the maximal peak happens at $\delta t = \pm \frac{3^{1/4}}{a_k D^{1/4}}$ with the peak value $\max(f_5) = \frac{5}{4} + \frac{3\sqrt{3}\sqrt{D}}{16}$.

A.3 Variance of linear functions

With the diagrammatic tool developed in Appendix A.2, we can discuss the efficiency of shadow tomography in predicting linear and nonlinear functions. In this section, we are going to focus on the linear function prediction, which has the form $o = \mathrm{Tr}(O\rho)$, and in Appendix A.4, we will provide numerical evidence of variance behavior under various conditions. In addition, in Appendix A.5, we will discuss the prediction of nonlinear function $o = \mathrm{Tr}(O\rho^{\otimes k})$, which involves k copies of ρ.

The linear function $o = \mathrm{Tr}(O\rho)$ can be estimated from classical shadows $\hat{\rho}$ via $o = \mathbb{E}[\mathrm{Tr}(O\hat{\rho})] = \mathbb{E}[\hat{o}]$, where $\hat{o} \equiv \mathrm{Tr}(O\hat{\rho})$ can be viewed as a random variable derived from the classical shadow. In the main text, we have shown the efficiency of shadow tomography is closely related the variance of estimation $\mathrm{Var}(\hat{o})$, which is upper bounded by

$$\mathrm{Var}(\hat{o}) \leq \mathbb{E}[\hat{o}^2]. \tag{A.23}$$

Diagrammatically, it can be expressed as

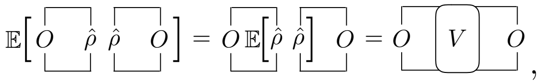

(A.24)

where V operator is defined in Eq.A.11. Furthermore, since $\mathrm{Tr}(\hat{\rho}) = 1$, adding O by $c\mathbb{1}$ only shifts $\hat{o}$ by a constant c, and it will not affect variance. So for linear function, we can assume the observable O to be traceless, i.e. $\mathrm{Tr}(O) = 0$. Then

$$\mathrm{Var}(\hat{o}) \leq \mathbb{E}[\hat{o}^2] = V[O^{\otimes 2}] = V[O_o^{\otimes 2}] + V[O_d^{\otimes 2}] + V[O_o \otimes O_d]. \tag{A.25}$$

A.3.1 Off-diagonal Pauli observables

We define the off-diagonal dynamical form factor $F_o(t)$ as

$$F_o(t) = \frac{V[O_o^{\otimes 2}]}{\mathrm{Tr}(O_o^2)}. \tag{A.26}$$

If the observables are off-diagonal Pauli operator O_o, then $F_o(t)$ can be simplified as

$$F_o(t) = \frac{f_1(t)\,\mathrm{Tr}(O_o^2) + f_2(t)\,\mathrm{Tr}(O_o^2\rho) + f_3(t)\,\mathrm{Tr}(O_o^2\rho)}{\mathrm{Tr}(O_o^2)} = \frac{f_1(t)D + f_2(t) + f_3(t)}{D} \simeq \frac{1}{1 - r^4(t)}. \tag{A.27}$$

In the last step, we only keep the leading D terms.

A.3.2 Diagonal Pauli observables

We define the diagonal dynamical form factor $F_d(t)$ as

$$F_d(t) = \frac{V[O_d^{\otimes 2}]}{\mathrm{Tr}(O_d^2)}. \tag{A.28}$$

If the observables are diagonal Pauli operator O_d, then $F_d(t)$ can be simplified as

$$F_d(t) = \frac{f_4(t)\,\mathrm{Tr}(O_d^2) + f_5(t)\,\mathrm{Tr}(O_d^2\rho)}{\mathrm{Tr}(O_d^2)} = \frac{f_4(t)D + f_5(t)}{D} \simeq f_4(t) = \frac{1}{1 + Dr^4(t)}. \tag{A.29}$$

In the second last step, we use the fact that $\max(f_5(t)/D) \sim D^{-1/2}$. For large D, this contribution can be ignored.

A.4 Numerical results and case studies

A.4.1 Numerical studies of reconstruction channel

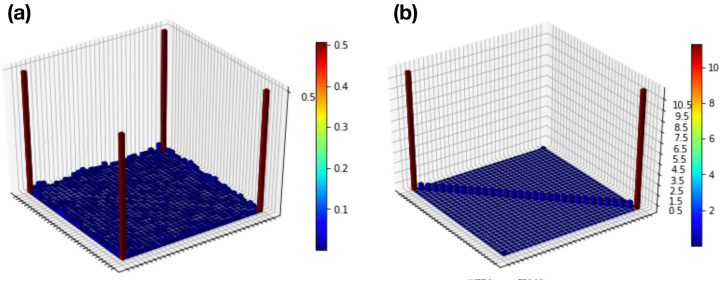

Figure A.2: Early time reconstruction of GHZ state. Measurements are taken at $T = 0.4$. (a) shows the unbiased reconstruction of GHZ density matrix with 100,000 classical snapshots, while (b) shows the reconstruction channel using unitary 2-design is highly biased for early time reconstruction.

In the main text, we derived the unbiased reconstruction channel $\mathcal{M}^{-1}(X)$ for whole range of dynamical time. We demonstrated this using 5 qubits GHZ states, which GUE random Hamiltonians. At early time $T = 0.4$, we collected 10^5 classical snapshots. In Fig. A.2(a), we showed our reconstruction channel give unbiased density matrix of 5 qubits GHZ states, while the reconstruction channel of unitary 2-design is high biased for early time reconstruction. Especially, the off-diagonal information in density matrix is missing. In practice, we need to measure the state and collect classical snapshots at short-time or intermediate-time scale. Firstly, we want to shrink the evolution time to reduce the total time of experiments. Secondly, long-time evolution posts difficulties in classical post-processing. Because simulate chaotic Hamiltonian dynamics is hard and inaccurate for long time. These reasons justify the need for an unbiased reconstruction channel for whole time range.

A.4.2 Numerical case studies of variance estimations

In the main text, Appendix A.2 and Appendix A.3, we derived and analyzed the upper bound of the variance of $\hat{o} = \mathrm{Tr}(O\hat{\rho})$. We numerically studied variance for several cases. We found our analytical calculation agrees well with numerical simulations, and the upper bound is tight for those cases.

Case 1: The observable only contains off-diagonal term $O = O_o$.

From the previous discussion, we know

$$\mathrm{Var}(\hat{o}) \leq \mathbb{E}[\hat{o}^2] = f_1(t)\,\mathrm{Tr}(O_o^2) + f_2(t)\,\mathrm{Tr}(O_o^2\rho) + f_3(t)\,\mathrm{Tr}(O_o^2\rho_d) \tag{A.30}$$

The $f_1(t)$ and $f_2(t)$ terms are regular and bounded, while $f_3(t)$ will give t^{-2} divergence as

$$\lim_{t\to 0} f_3(t) = \frac{D}{2t^2}. \tag{A.31}$$

So in early time, $\mathbb{E}[\hat{o}]$ will scale with D, and t^{-2}. To test these phenomenon, we prepare the system in the GHZ state, and $O_o = 0.5|\uparrow\cdots\uparrow\rangle\langle\downarrow\cdots\downarrow| + 0.5|\downarrow\cdots\downarrow\rangle\langle\uparrow\cdots\uparrow|$ to be the off-diagonal fidelity. We collected 10^4 classical snapshots, and numerically estimate $\mathrm{Var}(O_o)$ under various conditions. In Fig. A.3(a), we found the t^{-2} behavior agrees well with the numerical simulation. And in Fig. A.3(b), we also numerically confirmed the variance scales with Hilbert space dimension D.

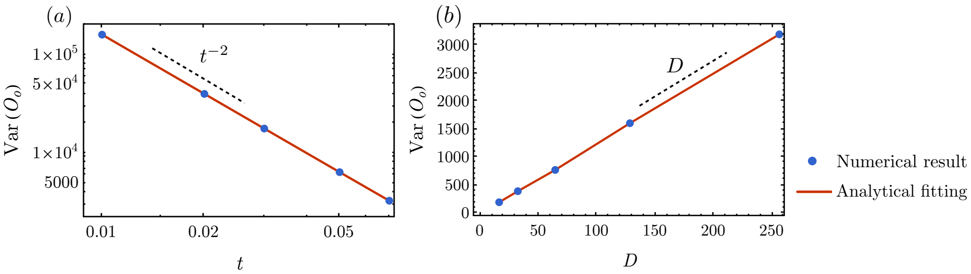

Figure A.3: Numerical tests on the scaling of variance for off-diagonal operators. (a) shows the t^{-2} scaling using 5 qubits GHZ states; (b) shows D scaling at fixed evolution time $T = 0.1$.

Case 2: The observable only contains diagonal term $O = O_d$, and $\mathrm{Tr}(O_d^2\rho) \simeq \mathrm{Tr}(O_d^2)/D$.

For diagonal operator O_d, the variance of estimation using shadow tomography is upper bounded by

$$\mathrm{Var}(\hat{o}_d) \leq \mathbb{E}[\hat{o}_d^2] = f_4(t)\,\mathrm{Tr}(O_d^2) + f_5(t)\,\mathrm{Tr}(O_d^2\rho) = (f_4(t) + f_5(t)/D)\,\mathrm{Tr}(O_d^2). \tag{A.32}$$

The behavior of the variance will dominated by $f_4(t)$. To test this, we prepared 9 qubits GHZ state as initial state, and we measure the diagonal operator $O_d = Z_1 I_2 \cdots I_9$. In Fig. A.4(a), the blue dots shows the numerical estimation of the variance of $\hat{o}_d$. Indeed, we see numerical results match $f_4(t)$ behavior. And to further quantify it, we use ansatz $c_1 f_4(t) + c_2 f_5(t) + c_3$ to fit the numerical results, where $c_{1,2,3}$ are fitting parameters. By minimizing the mean-square error, we got the best fitting results with $c_1 = 514.5$, $c_2 = 0.93$, and $c_3 = -1.742$. We see the value of c_1 matches its theoretical value $c_1 = \mathrm{Tr}(O_d^2) = 512$, and c_2 matches $c_2 = \mathrm{Tr}(O_d^2\rho) = 1$ really well. And the red curve in Fig. A.4(a) shows the fitting result.

Case 3: The observable only contains diagonal term $O = O_d$, and $\mathrm{Tr}(O_d^2\rho) \simeq \mathrm{Tr}(O_d^2)$.

If this is the case, then the variance behavior will be dominated by $f_5(t)$. To test this idea, we prepare the state in $|\psi\rangle = |\downarrow \cdots \downarrow\rangle$ with 9 qubits, and choose the diagonal operator as $O_d = |\downarrow \cdots \downarrow\rangle\langle\downarrow \cdots \downarrow|$. In Fig. A.4(b), we see the variance behavior is indeed resembles $f_5(t)$. We also use ansatz $c_1 f_4(t) + c_2 f_5(t) + c_3$ to fit the numerical results, where $c_{1,2,3}$ are fitting parameters. By minimizing the mean-square error, we got the best fitting results with $c_1 = 1.02$, $c_2 = 0.94$, and $c_3 = -1.02$. Again, the fitting result c_1 matches its theoretical result $c_1 = \mathrm{Tr}(O_d^2) = 1$, and c_2 matches $c_2 = \mathrm{Tr}(O_d^2\rho) = 1$ very well. From these tests, we can see our analytical upper bound estimation of variance matches the numerical results under various conditions, and the upper bound is tight in the numerical tests.

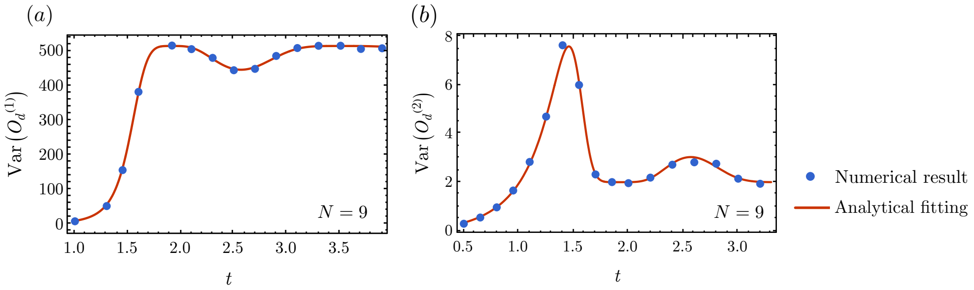

Figure A.4: Numerical tests of the variance of diagonal variables. We use 9 qubits GHZ states as our initial state. Subplot (a) and (b) shows different choices of diagonal observables corresponding to case 2, and case 3. Blue dots are numerical results, and the red curves are analytical fittings.

A.5 Variance of nonlinear functions

For nonlinear function involving k copies of ρ, such as $o_k = \mathrm{Tr}(O\rho^{\otimes k})$, we can estimate it using $\hat{o}_k = \mathrm{Tr}(O\hat{\rho}_1 \otimes \cdots \otimes \hat{\rho}_k)$, where $\hat{\rho}_i$ are independent samples of classical shadows. The efficiency is related to $\mathrm{Var}(\hat{o}_k) \leq \mathbb{E}[\hat{o}_k^2]$. Since $\hat{\rho}_i$ are independent random variables, we can have

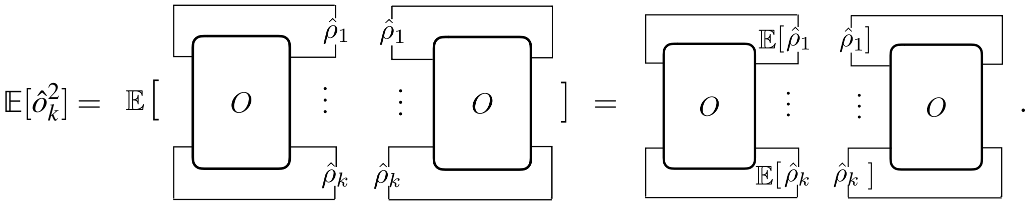

(A.33)

Using Eq.A.10, we can simplify the above equation as

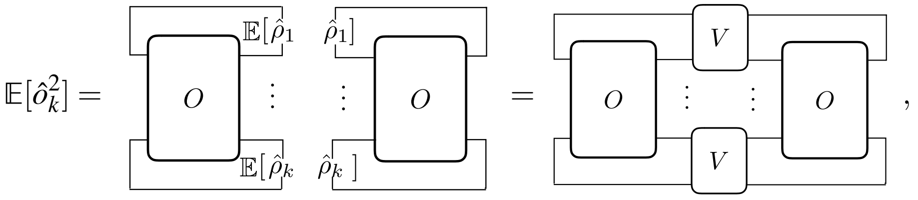

(A.34)

where each $V[\cdots]$ involves 7 P-diagrams as defined in Eq.A.12. From Eq.A.34, we can in principle calculate the upper bound for variance of nonlinear functions involving any order k of ρ:

$$\mathrm{Var}(\mathrm{Tr}(O\hat{\rho}^{\otimes k})) \leq \mathbb{E}[\mathrm{Tr}(O\hat{\rho}^{\otimes k})^2] = \sum_{\boldsymbol{\alpha}\in[1,8]^k} \prod_{\alpha_i} f_{\alpha_i}(t) P_{\alpha_i}, \tag{A.35}$$

where $\boldsymbol{\alpha} = (\alpha_1, \cdots, \alpha_k) \in [1,7]^k$ labels the P-diagram attached to each tensor leg. For convenience, we rename

$$\begin{aligned} f_1(t) + f_2(t) + f_3(t) &= F_{(o,o)}(t) \\ f_4(t) + f_5(t) &= F_{(d,d)}(t) \\ f_6(t) &= F_{(I,o)}(t) = F_{(o,I)}(t) \\ f_7(t) &= F_{(I,d)}(t) = F_{(d,I)}(t) \\ f_8(t) &= F_{(o,d)}(t) = F_{(d,o)}(t). \end{aligned} \tag{A.36}$$

With those names, we can formally rewrite Eq.A.35 as

$$\mathrm{Var}(\mathrm{Tr}(O\hat{\rho}^{\otimes k})) \leq \mathbb{E}[\mathrm{Tr}(O\hat{\rho}^{\otimes k})^2] \leq \sum_{\boldsymbol{\alpha}=(\vec{\alpha}_1,\cdots,\vec{\alpha}_k)} \mathrm{Tr}(O_{\alpha_1^L\alpha_2^L\cdots\alpha_k^L} O_{\alpha_1^R\alpha_2^R\cdots\alpha_k^R}) \prod_{\vec{\alpha}_i} F_{\vec{\alpha}_i}(t), \tag{A.37}$$

where $\boldsymbol{\alpha} = (\vec{\alpha}_1, \cdots, \vec{\alpha}_k)$ has k-component, and each $\vec{\alpha}_i = (\alpha_i^L, \alpha_i^R)$ is a two component vector where $\alpha_i^{L/R}$ labels the i-th tensor leg of left/right O operator. Each

$$\vec{\alpha}_i = (\alpha_i^L, \alpha_i^R) \in \{(o,o), (d,d), (I,o), (o,I), (I,d), (d,I), (o,d), (d,o)\}$$

, and the summation is over all combinations of $\boldsymbol{\alpha}$. In the last step, we use the fact, $\mathrm{Tr}(O\rho) \leq \mathrm{Tr}(O)$.

Appendix B

Classical shadow tomography with locally scrambled quantum dynamics

B.1 Entanglement feature and the reconstruction channel

In this appendix, we discuss the details about the derivation on measurement channel $\sigma = \mathcal{M}[\rho]$ and reconstruction channel $\rho = \mathcal{M}^{-1}[\sigma]$. With the notation developed in 6.2.1, the expected classical snapshot σ is expressed as

$$\sigma = \mathcal{M}[\rho] = \sum_{\hat{\sigma}\in\mathcal{E}_{\sigma}} \hat{\sigma}\,\mathrm{Tr}(\hat{\sigma}\rho)d^{N}. \tag{B.1}$$

By utilizing the assumption that prior ensemble $\mathcal{E}_{\sigma}$ is locally scrambled, i.e. $P(\hat{\sigma}) = P(V^{\dagger}\hat{\sigma}V), \forall V \in \mathrm{U}(d)^{N}$, we are free to insert local basis transformations V, and average it out. By doing so, we have

$$\sigma = \mathbb{E}_{V\in\mathrm{U}(d)^N}\mathbb{E}_{\hat{\sigma}\in\mathcal{E}_\sigma} V^\dagger\hat{\sigma}V\,\mathrm{Tr}(V^\dagger\hat{\sigma}V\rho)d^N$$

$$= \underset{V\in U(d)^N}{\mathbb{E}}\ \underset{\hat{\sigma}\in\mathcal{E}_\sigma}{\mathbb{E}}$$

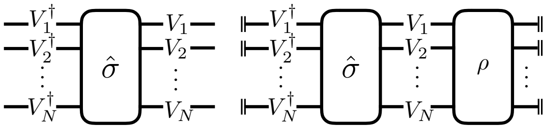

$$d^N \tag{B.2}$$

$$= \sum_{B,C\in 2^{\Omega_N}} d^N\,\mathrm{Wg}_{B,C}$$

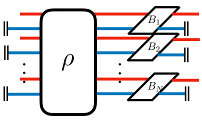

$$\Bigg(\underset{\hat{\sigma}\in\mathcal{E}_\sigma}{\mathbb{E}}$$

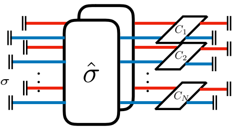

$$\Bigg),$$

where each B_i and C_i have two choices: swap operator (1) or identity operator (0), and $\mathrm{Wg}_{B,C} = (d^2-1)^{-N}(-1/d)^{|B\ominus C|}$ is the Weingarten function of regions B and C, where $B\ominus C = (B\setminus C)\cup(C\setminus B)$ denotes their symmetric difference. In the above tensor diagram, short parallel lines indicate the periodic boundary condition, and the summation of B and C is over all possible subregions of the N qudit system. As we can see, if we choose a subregion B to be the swap operators, then ρ will be traced out on the counter part $\bar{B}$. In addition, the identity operators (red lines) on $\bar{B}$ are inserted. So the first tensor diagram in Eq. (B.2) is the reduced density matrix embedded back into the total Hilbert space. We spoil the notation and use $\rho_B d^{\bar{B}} = (\mathrm{Tr}_{\bar{B}}\,\rho)\otimes\mathbb{1}_{\bar{B}}$ to denote the first tensor diagram, but one should remember the identity operators are supported in region $\bar{B}$. The tensor product $\otimes$ notation indicates that $(\mathrm{Tr}_{\bar{B}}\,\rho)$ and $\mathbb{1}_{\bar{B}}$ act separately in regions B and $\bar{B}$, which does *not* imply that B should be a consecutive region “in front of” $\bar{B}$ (as in the conventional notation). The second tensor diagram in Eq. (B.2) is the 2nd *entanglement feature* of the prior POVM $\mathcal{E}_\sigma$,

$$W^{(2)}_{\mathcal{E}_\sigma,C} \equiv \underset{\hat{\sigma}\in\mathcal{E}_\sigma}{\mathbb{E}} \mathrm{Tr}_C(\mathrm{Tr}_{\bar{C}}\,\hat{\sigma})^2 = \underset{\hat{\sigma}\in\mathcal{E}_\sigma}{\mathbb{E}} e^{-S_C^{(2)}(\hat{\sigma})}, \tag{B.3}$$

where $S_C^{(2)}(\hat{\sigma})$ denotes the 2nd Rényi entanglement entropy of the state $\hat{\sigma}$ in region C. The above tensor diagram representation is equivalent to Eq.6.7 in the main text.

B.2 Variance estimation and sample complexity

In the main text, we relate the sample complexity M with the ρ-dependent shadow norm $\|O\|^2_{\mathcal{E}_{\sigma|\rho}}$, by

$$M \geq \|O\|^2_{\mathcal{E}_{\sigma|\rho}}/\varepsilon^2\delta. \tag{B.4}$$

However, the ρ-dependent shadow norm $\|O\|^2_{\mathcal{E}_{\sigma|\rho}}$ is generally complicated to evaluate. If we are not interested in the shadow norm for a specific state ρ, but rather the expectation of the shadow norm over an ensemble of states $\{V\rho V^\dagger\}$ that are similar to ρ by local basis transformations $V \in \mathrm{U}(d)^N$, we can actually define a ρ-independent shadow norm by averaging over V. The result is similar to Eq. (6.15)

$$\begin{aligned}\|O\|^2_{\mathcal{E}_\sigma} &\equiv \mathop{\mathbb{E}}_{V\in\mathrm{U}(d)^N} \|O\|^2_{\mathcal{E}_{\sigma|V\rho V^\dagger}} \\ &= \sum_{g,h\in S_3^N} \|O\|^2_g \mathsf{Wg}_{g,h} W^{(3)}_{\mathcal{E}_\sigma,h},\end{aligned} \tag{B.5}$$

where $\|O\|^2_g$ is inherited from Eq. (6.16)

$$\begin{aligned}\|O\|^2_g &\equiv \mathop{\mathbb{E}}_{V\in\mathrm{U}(d)^N} \|O\|^2_{V\rho V^\dagger,g} \\ &= \mathrm{Tr}\left((\mathcal{M}^{-1}[O]^{\otimes 2}\otimes \mathbb{1})\chi_g\right).\end{aligned} \tag{B.6}$$

Compared with Eq. (6.16), we can see that the ensemble average $\mathbb{E}_{V\in\mathrm{U}(d)^N}$ in Eq. (B.6) removes the ρ dependence by effectively replacing ρ with $d^{-N}\mathbb{1}$ (the prior density matrix that defines the prior POVM $\mathcal{E}_\sigma$). This explains the consistency in our notation that $\|O\|^2_{\mathcal{E}_\sigma} = \mathbb{E}_{\hat{\sigma}\in\mathcal{E}_\sigma}\hat{o}(\hat{\sigma})^2$ follows from essentially the same definition as in Eq. (6.14).

Note that the reconstruction map $\mathcal{M}^{-1}$ always commutes with the local basis transformation $V = \prod_i V_i$, i.e. $\mathcal{M}^{-1}[V^\dagger O V] = V^\dagger \mathcal{M}^{-1}[O]V$, because V_i acts on each qudit separately and hence does not interfere with the partial trace operation. This indicates that the norm

$\|O\|_g^2 = \|V^\dagger O V\|_g^2$ is invariant under the transformation V. This suggests us to define a locally scrambled ensemble $\mathcal{E}_O$ (or known as $U(d)^N$-twirling) associated with any given observable O

$$\mathcal{E}_O \equiv \{V^\dagger O V | V \in \mathrm{U}(d)^N\}, \tag{B.7}$$

such that $\|O\|_g^2$ in Eq. (B.6) can be redefined as its ensemble average

$$\begin{aligned}
\|O\|_g^2 &= \mathop{\mathbb{E}}_{O\in\mathcal{E}_O} \|O\|_g^2 \\
&= \mathop{\mathbb{E}}_{V\in\mathrm{U}(d)^N} \mathrm{Tr}\left((\mathcal{M}^{-1}[V^\dagger O V]^{\otimes 2} \otimes \mathbb{1})\chi_g\right) \\
&= \sum_{A,B,C,D\in 2^{\Omega_N}} d^{2N} r_A r_B \mathsf{Wg}_{C,D} W^{(2)}_{\mathcal{E}_O,D} \mathrm{Tr}\left(((\chi_C)_{A,B} \otimes \mathbb{1})\chi_g\right).
\end{aligned} \tag{B.8}$$

Here χ_C denotes the swap operator supported in region C that acts between the first two copies of the Hilbert space, and $(\chi_C)_{A,B}$ denotes the reduction of χ_C in region A and B respectively in the first and the second copies of the Hilbert space, which results in $(\chi_C)_{A,B} = \chi_{A\cap B\cap C} d^{|A\cap B\cap C|-|C|}$. The operator entanglement feature $W^{(2)}_{\mathcal{E}_O,D} = \mathbb{E}_{O\in\mathcal{E}_O} \mathrm{Tr}_D(\mathrm{Tr}_{\bar{D}}\, O)^2$ follows from the same definition given in Eq. (6.8). r_A, r_B are the reconstruction coefficients given by the solution of Eq. (6.12). Substitute Eq. (B.8) to Eq. (B.5), we can evaluate the summation of g,h given that $\sum_{g,h\in S_3^N} \mathrm{Tr}(\chi_{A\cap B\cap C}\chi_g) \mathsf{Wg}_{g,h} W^{(3)}_{\mathcal{E}_\sigma,h} = W^{(2)}_{\mathcal{E}_\sigma,A\cap B\cap C}$. The reduction of the 3rd entanglement feature to the 2nd entanglement feature is a consequence of the fact that ρ drops out from the tensor product in Eq. (B.6), such that only 2-fold Hilbert space is required to define $\|O\|^2_{\mathcal{E}_\sigma}$.

Thus we finally arrive at the expression for the operator shadow norm purely in terms of the entanglement features of $\mathcal{E}_\sigma$ and $\mathcal{E}_O$,

$$\|O\|^2_{\mathcal{E}_\sigma} = \sum_{A,B,C,D\in 2^{\Omega_N}} v_{A,B,C,D} W^{(2)}_{\mathcal{E}_\sigma,A\cap B\cap C} W^{(2)}_{\mathcal{E}_O,D}, \tag{B.9}$$

where the coefficient $v_{A,B,C,D}$ is given by

$$v_{A,B,C,D} = r_A r_B \left(\frac{d^2}{d^2-1}\right)^N d^{|A\cap B\cap C|-|C|}\left(-\frac{1}{d}\right)^{|C\ominus D|}. \tag{B.10}$$

B.3 Fidelity estimation for mixed state

Our method is not restricted to pure state. In variational quantum state preparation, even the target state is some pure state, noise in the preparation circuit could make the final state in experiments a mixed state. We can use the shallow circuit classical shadow tomography to efficiently estimate the quantum fidelity between final prepared quantum state ρ_P and the target quantum state ρ_T. Fast access to this quantity is crucial for variational quantum state preparation, error mitigation and etc. As an example, we consider the noisy preparation of a perfect GHZ state with Z errors occurs at probability p. The prepared state can be expressed as

$$\rho_P = (1-p)|\psi^+_{\mathrm{GHZ}}\rangle\langle\psi^+_{\mathrm{GHZ}}| + p|\psi^-_{\mathrm{GHZ}}\rangle\langle\psi^-_{\mathrm{GHZ}}|, \tag{B.11}$$

where $|\psi^\pm_{\mathrm{GHZ}}\rangle = \frac{1}{\sqrt{2}}(|0^{\otimes N}\rangle \pm |1^{\otimes N}\rangle)$. We compare the performance between random Pauli measurement and shallow circuit shadow tomography with 3 layers of local random unitaries. Experiments are performed on a 9-qubit system and 5000 classical snapshots are collected for both random Pauli shadow tomography and shallow circuits shadow tomography. The result is shown in Fig. B.1. As we can see, for 5000 experiments, the quantum fidelity estimated using random Pauli measurement has huge error bar, indicated by the blue shaded region. However, same amount of data collected after shallow circuit evolution can give accurate estimation of quantum fidelity, and error bar is almost four times smaller. Practically, this makes the usage of shallow circuits more appealing.

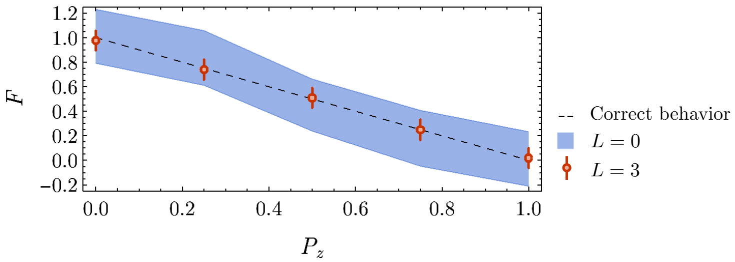

Figure B.1: Fidelity estimation between mix state and target state. 5000 experimental classical snapshots are prepared for both random Pauli measurement ($L = 0$) and shallow random unitary circuit ($L = 3$). Error bar indicates 3 standard deviation.

B.4 Approximated unitary ensemble and purification

In the main text, we have seen when the measurement channel $\mathcal{M}$ in data acquisition and the reconstruction channel $\mathcal{M}^{-1}$ in classical post-processing mismatch, the reconstructed density matrix $\frac{1}{M}\sum_{\hat{\sigma}\in\mathcal{E}_{\sigma|\rho}}\mathcal{M}^{-1}[\hat{\sigma}]$ may not be positive-definite. And it results in biased prediction of physical quantities. In Fig. 6.5 (a) and Fig. 6.9, we have seen the biased prediction of fidelity that is larger than one. In Fig. B.2, we plot the eigenvalues of reconstructed density matrix of 7-qubit GHZ state using DQIM ensemble with $T/T_{\text{Th}} = 1.38$. In the main text, we have seen the DQIM ensemble with one period of evolutional time or $T/T_{\text{Th}} = 1.38$ is not sufficient to achieve the local scrambling assumption, such that there is mismatch between data acquisition channel $\mathcal{M}$ and reconstruction channel $\mathcal{M}^{-1}$. We see the spectrum of density matrix contains some negative eigenvalues.

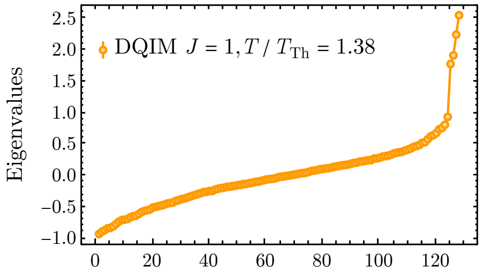

Figure B.2: Eigenvalues of reconstructed density matrix $\tilde{\rho}$ of 7-qubit GHZ state using mismatched channels. The unitary ensemble is single instance of DQIM ensemble with $J = 1$, and $T/T_{\text{Th}} = 1.38$. Under this condition, the unitary ensemble is not locally scrambled.

In addition, the approximate shadow tomography based on locally scrambling Hamiltonian evolution, such as DQIM ensemble or GUE2 ensemble, is approximately unbiased when local scrambling is approximately satisfied or frame potential gap is vanishingly small. In Fig. 6.9, we have seen they all can give unbiased prediction of quantum fidelity when $T \geq 10T_{\mathrm{Th}}$. We directly visualize the reconstructed density matrix using approximated DQIM ensemble in Fig. B.4 and Fig. B.5. As we see in Fig. B.5, at $T/T_{\mathrm{Th}} = 1.95$, the locally scrambling assumption is not satisfied, and reconstructed density matrix is biased. In contrast, at $T/T_{\mathrm{Th}} = 25.3$ (Fig. B.4), the reconstructed density matrix using a single instance of DQIM Hamiltonian is perfect, justifying the validity of our approach when the locally scrambling assumption is approximated satisfied.

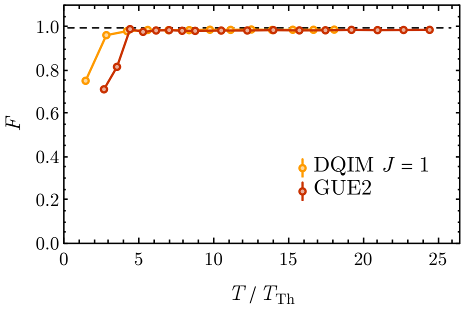

Figure B.3: Fidelity estimation of approximated unitary ensemble after purification. After around $T \sim 10T_{\mathrm{Th}}$, the fidelity is around 0.99. Same data are used as Fig. 6.9

Further more, for biased reconstruction, in order to make it positive definite, we can nonlinear project the reconstructed ρ to the convex set of physical states $C = \{\rho|\rho \succeq 1, \mathrm{Tr}(\rho) = 1\}$ by minimizing

$$\Pi_C(\sigma) = \arg\min_{\rho \in C} \mathrm{Tr}((\rho - \sigma)^2), \tag{B.12}$$

which is the method mentioned in Ref. [ASS21]. If we have more prior knowledge about the quantum state, such as it is a pure state, then we can further impose those assumptions into the projection. Here, as an illustration, we utilize the knowledge that the target quantum state is pure, and we project the reconstructed ρ to a pure state $\tilde{\rho}$ in C by choosing the eigenstate of ρ with the largest eigenvalue. As shown in Fig. B.3, for the approximated ensembles, the GUE2 and DQIM are biased in the short time region, and projected state $\tilde{\rho}$ has a fidelity less than one. And when locally scrambling assumption is approximately satisfied, the projected $\tilde{\rho}$ will have fidelity that is approximately 0.99. With these checks:

- unbiased prediction of physical quantities, see Fig. 6.9
- high fidelity of reconstructed density matrix projected back to physical space, see Fig. B.3

we confirm the approximated shadow tomography can perform unbiased reconstruction.

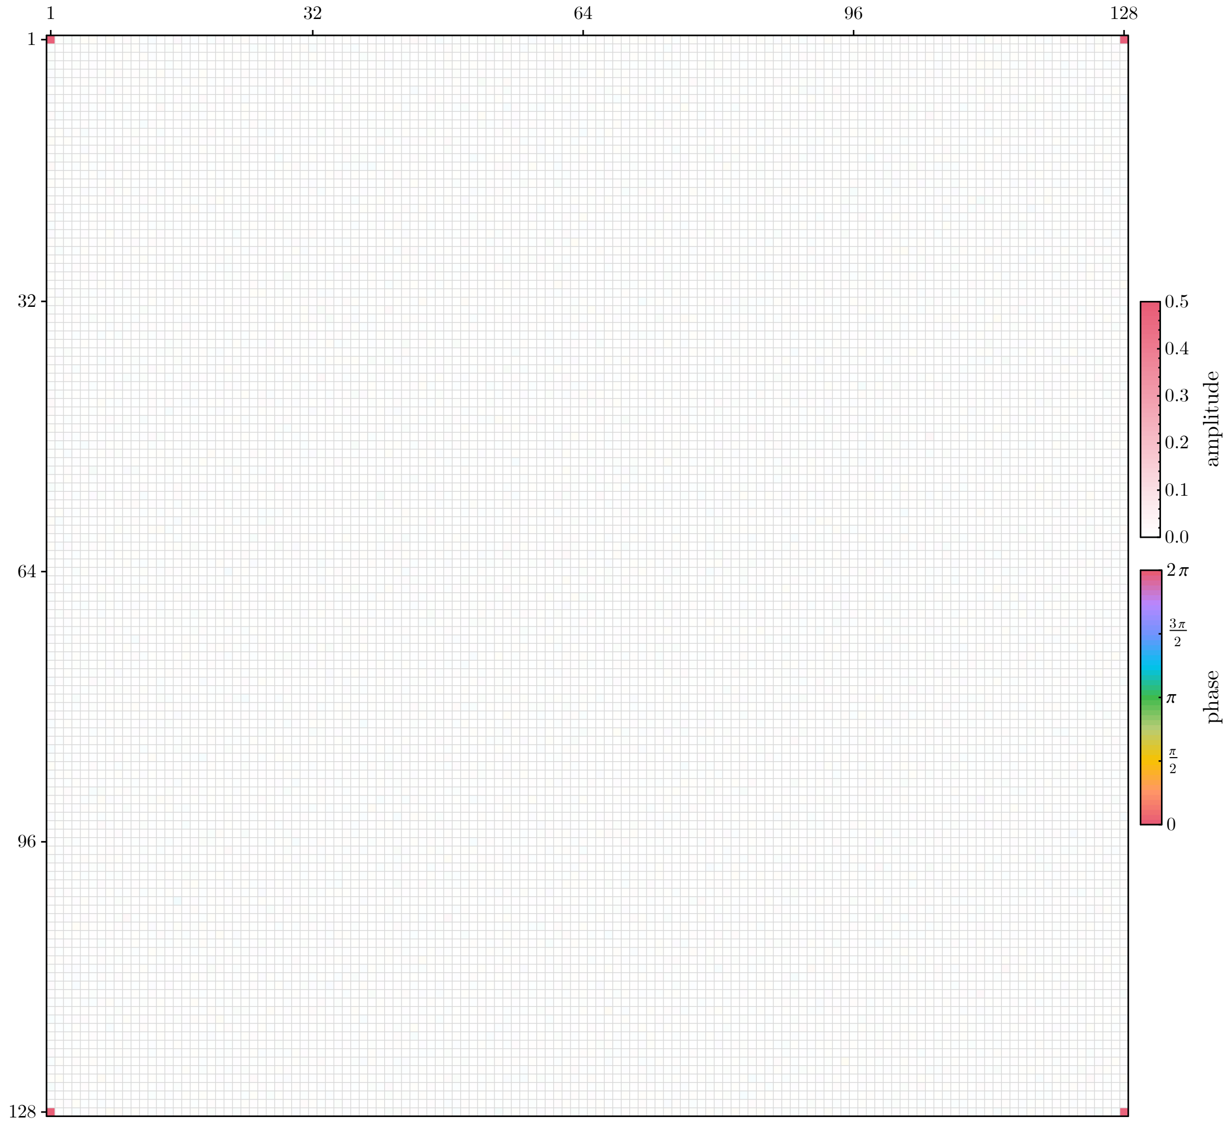

Figure B.4: Unbiased reconstruction of a 7-qubit GHZ density matrix, using a single instance of Hamiltonian in the DQIM ensemble at $T/T_{\mathrm{Th}} = 25.3$ (after the local scrambling condition is achieved).

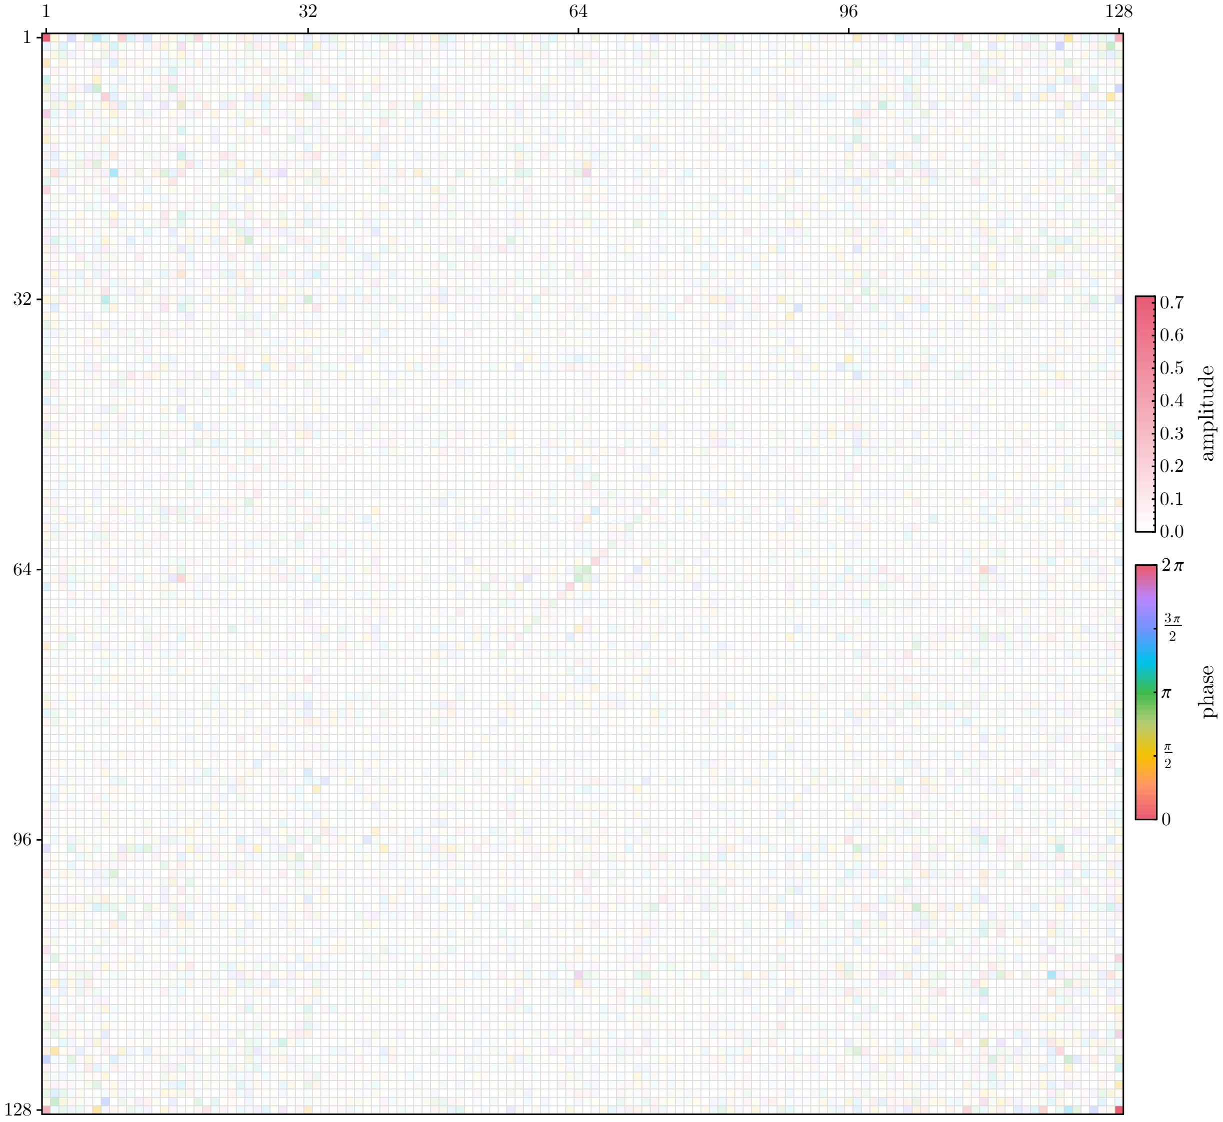

Figure B.5: Biased reconstruction of a 7-qubit GHZ density matrix, using a single instance of Hamiltonian in the DQIM ensemble at $T/T_{\mathrm{Th}} = 1.95$ (before the local scrambling condition is achieved).

Appendix C

Topological and symmetry-enriched random quantum critical points

C.1 Spectrum bifurcation renormalization group

In this appendix, we briefly summarize the spectrum bifurcation renormalization group (SBRG) method [YQX16, SYX16, SBYX17]. The basic idea of SBRG is to progressively identify conserved quantities by block diagonalization, and treat the off-diagonal term within second order perturbation theory. The general form of a qubit model Hamiltonian can be written as

$$H = \sum_{[\mu]} h_{[\mu]} \sigma^{[\mu]}, \tag{C.1}$$

where σ^{μ_i} ($\mu_i = 0, 1, 2, 3$) are Pauli matrices acting on the ith site, and $\sigma^{[\mu]} \equiv \sigma^{\mu_1} \otimes \sigma^{\mu_2} \otimes \sigma^{\mu_3} \cdots$ denotes a string of Pauli operators. The real coefficients $h_{[\mu]}$ are drawn from some random distribution. To progressively diagonalize the Hamiltonian in Eq.(C.1), we first find the leading energy scale term

$$H_0 = h_{[\mu]_{\max}} \sigma^{[\sigma]_{\max}}, \tag{C.2}$$

where $|h_{[\mu]_{\max}}|$ has the largest value among all the coefficients. We then block diagonalize term $\sigma^{[\mu]_{\max}}$ by a Clifford rotation R,

$$R^{\dagger}\sigma^{[\mu]_{\max}}R = \sigma^{3[0\cdots 0]}, \tag{C.3}$$

and all the other terms will be rotated by the same operator R, such that $H \to R^{\dagger}HR$. As we block diagonalize the leading energy term, the many-body spectrum bifurcates to high energy $E \sim |h_{[\mu]_{\max}}|$ sector, or to low energy $E \sim -|h_{[\mu]_{\max}}|$ sector. We now reduce the other leftover Hamiltonian to either sector. The Hamiltonian can be separated into the following three terms:

$$H = H_0 + \Delta + \Sigma, \tag{C.4}$$

where Δ are the terms that commute with the leading term H_0, and Σ are the terms that anti-commute with H_0. We will renormalize the off-diagonal term Σ using second order perturbation

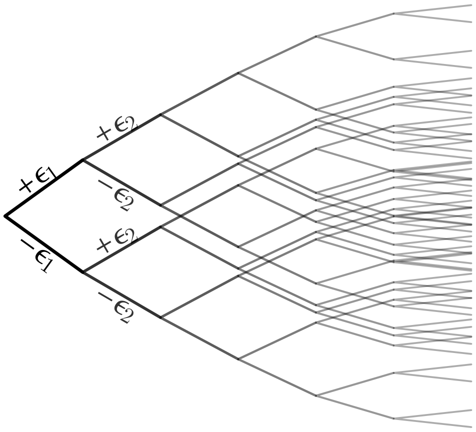

Figure C.1: **Spectrum bifurcation renormalization group.** Spectrum of energy generated by SBRG: at each RG steps, the RG bifurcates into a high- and low-energy branch, corresponding to $\tau_i = \pm 1$ respectively. Each many-body eigenstate corresponds to a leaf of the tree, uniquely labelled by the values $\{\tau_i = \pm 1\}$.

theory, $H \to S^\dagger H S$, with $S = \exp\left(-H_0\Sigma/2h_{[\mu]_\mathrm{max}}\right)$. The effective Hamiltonian reads

$$H = H_0 + \Delta + \frac{1}{2h^2_{[\mu]_\mathrm{max}}} H_0 \Sigma^2 + O\left(\frac{1}{2h^3_{[\mu]_\mathrm{max}}}\right). \tag{C.5}$$

And we will only keep terms generated up to 2nd order in perturbation theory. We then repeat this procedure at $O(N)$ times to fully diagonalize the original many-body Hamiltonian. At the end of RG flow, we will obtain the effective Hamiltonian

$$H_\mathrm{eff} = \sum_i \varepsilon_i \tau_i + \sum_{i<j} \varepsilon_{ij} \tau_i \tau_j + \cdots, \tag{C.6}$$

where $\tau_i = \pm 1$ are the emergent conserved quantities.

The whole RG transformation can be summarized as a unitary transformation,

$$U_\mathrm{RG} = R_1 S_1 R_2 S_2 \cdots = \prod_i R_i S_i, \tag{C.7}$$

and $H \to U^\dagger_\mathrm{RG} H U_\mathrm{RG} \approx H_\mathrm{eff}$. The eigenstates can be labeled by a vector of conserved quantities, i.e. $|\tau\rangle = |\tau_1 = 1, \tau_2 = -1, \cdots\rangle$. We can effectively calculate the expectation value of any physical observable $\hat{O}$ in the $\{\tau_i\}$ basis as

$$\langle\psi|\hat{O}|\psi\rangle = \langle\psi|U_\mathrm{RG}U^\dagger_\mathrm{RG}\hat{O}U_\mathrm{RG}U^\dagger_\mathrm{RG}|\psi\rangle = \langle\tau|\hat{O}_\mathrm{eff}|\tau\rangle. \tag{C.8}$$

The entanglement entropy can also be efficiently calculate using the stabilizer formalism[YQX16, FCY$^+$04].

We investigated the bulk property of Ising* transition as a sanity check. In Fig.C.2 (a) show the entanglement entropy $S_E = -\mathrm{tr}\rho_A \log \rho_A$ grows logarithmically with the size L_A of the entanglement region following $S_E(L_A) = (c_\mathrm{eff}/3)\log L_A$[RM07]. We observe the "effective central charge" $c_\mathrm{eff} = (0.51 \pm 0.02)\log 2$, which is consistent with the expectation $c_\mathrm{eff} = (1/2)\log 2$ for the infinite-randomness fixed point in the Ising universality class. Another signature of the infinite-randomness fixed point is the dynamical scaling $l \sim (\log t)^2 \sim (-\log \varepsilon)^2$ that relates the length scale l and the energy scale ε[Fis95]. We check this by examining the relation between the length l_a of the Pauli string τ_a and its corresponding energy coefficient ε_a in H_{MBL}. Fig. C.2 (b) confirms the scaling $\overline{\log \varepsilon_a} \sim -\sqrt{l_a}$. These evidences justify that the disordered Ising* does flow to an infinite-randomness fixed point.

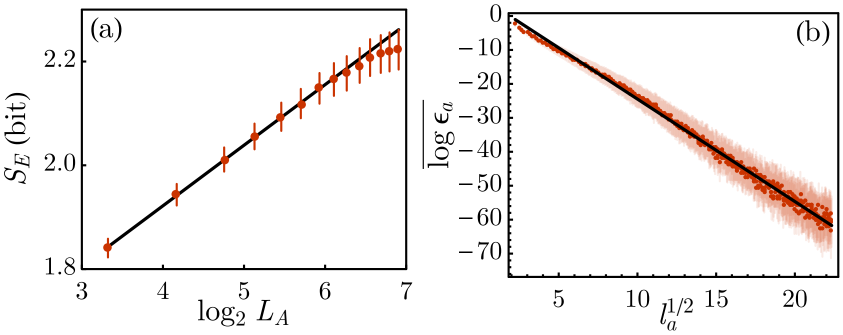

Figure C.2: **Bulk criticality of the Ising* fixed point** (a) Entanglement entropy scaling and (b) dynamical scaling (collected over 10^4 realizations) at the Ising* fixed point.

C.2 DMRG analysis of the symmetry-enriched random singlet phase

In order to study the stability of the symmetry-enriched random singlet phase discussed in the main text, we included perturbations to take the B spins away from integrability, to further couple the A and B spins, and we also added generic symmetry-preserving terms near the edges. In the following, we take J_i and Δ_i to be drawn from uniform random distributions, between $[0.1, 1.0]$ and $[0.3, 0.7]$, respectively; $g_B = 0.3$, and $V = -\sum_i 0.1 X_i^B X_{i+1}^B + 0.1 Z_i^A X_i^B Z_{i+1}^A$ where the first term breaks integrability of the B spins while the second term couples A and B spins before the unitary twist. We also add small fields $0.2(X_0^A + X_N^A) + (X_0^B + X_N^B)$ that preserve the

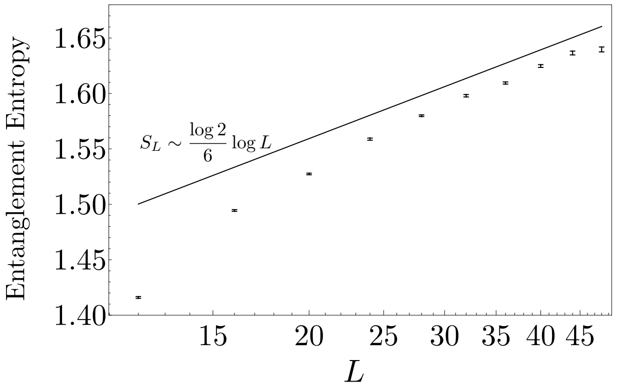

Figure C.3: **Entanglement scaling in the symmetry-enriched random singlet phase.** Bipartite entanglement entropy *vs* system size in the groundstate of the symmetry-enriched random singlet phase, computed using DMRG. The numerical results are consistent with the random singlet phase prediction $S \sim \log 2/6 \log L$ (solid lines).

$\mathbb{Z}_2^A \times \mathbb{Z}_2^B$ symmetry. (Note that those boundary terms break the $U(1)$ $\sum_i Z_i^A$ conservation at the boundary, but that symmetry is only required to tune the bulk to criticality, and does not play a role in protecting the edge modes.)

Because of the non-Abelian symmetry G_A of the model, the SBRG approach used to study eq. 1 in the main text does not apply here [1]. In order to establish the presence of topological edge modes and study their interplay with the coexisting quantum critical bulk fluctuations, we compute the low-energy spectrum using density-matrix renormalization group (DMRG) techniques [Whi93, ITe]. While DMRG cannot access system sizes as large as SBRG, it has the advantaged of being essentially numerically exact for reasonably small system sizes ($L \lesssim 50$ before numerical instabilities make the variational sweeps unpractical). We used the ITensor library [ITe] to perform DMRG simulations of the disordered symmetry-enriched random singlet Hamiltonian, using bond dimension $\chi \sim 500$. We considered system sizes $L = 2N = 12, 16, 20, \dots, 48$, where N is the number of spins per species A, B. The number of disordered realizations ranged from

---

[1] Contrary to the random Ising* universality class which extends to the whole spectrum (at least in the non-interacting case, where there is a "strong zero mode" [Fen12, Fen16]), here we expect the $\mathbb{Z}_2^A$ symmetry to be spontaneously broken in highly excited states, protected by MBL [VFPP16], resulting in very different physics than the quantum critical behavior at low energies.

~ 6000 (for the largest system size $L = 48$) to 3×10^4. We found that a large number of disorder realizations was necessary to observe the expected random singlet quantum critical behavior. However, the topological edge modes can be observed on individual disorder realizations.

We find an exponentially small in system size splitting between the two groundstates, $|\uparrow_L\downarrow_R\rangle \pm |\downarrow_L\uparrow_R\rangle$, where R and L denote the configurations of the right and left edge modes, as in the main text. The orientation of the edge modes can be detected from the spontaneous boundary magnetization in each eigenstate upon adding a small magnetic field. We observe numerically that those groundstates are separated from the first excited states $|\uparrow_L\uparrow_R\rangle \pm |\downarrow_L\downarrow_R\rangle$: critical bulk fluctuations separate ferro- and antiferromagnetic alignments of the edge modes, leading to a 2-fold groundstate degeneracy. Coexisting with these topological edge modes, we find bulk quantum critical properties characteristic of a random singlet phase. The typical bulk gap closes as $\Delta E_{\text{typical}} \equiv \mathrm{e}^{\overline{\log\Delta E}} \sim \mathrm{e}^{-\sqrt{L}}$, with as in the Ising case, a broad gap distribution leading to a different scaling for the average $\overline{\Delta E} \sim \mathrm{e}^{-L^{1/3}}$. This stretched-exponential gap closing is hallmark of an infinite randomness quantum critical point, and is apparent in our DMRG results as shown in the main text. We find other signatures of random-singlet criticality: in particular, the entanglement entropy (using open boundary conditions, with the entanglement cut in the middle of the system) grows logarithmically with system size L (Fig. C.3), with a prefactor compatible with the prediction $S \sim \frac{\log 2}{6}\log L$ [RM04].

C.3 Floquet Ising* criticality

In this appendix, we analyze the edge mode structure of the Floquet transverse Ising chain. We focus on the non-interacting limit for simplicity, and also consider uniform couplings as most features we want to illustrate are already apparent in this limit. Disorder is needed to many-body localize the phases in the presence of interactions, and to make the notion of Floquet phase structure meaningful away from the non-interacting limit. (Without disorder, Floquet systems typically heat up indefinitely and thermalize to infinite temperature.) Let us work in the Majorana representation with a pair of Majorana fermions $\gamma_j, \tilde{\gamma}_j$ per site j, and define

$$H_\alpha \equiv \frac{i}{2}\sum_n \tilde{\gamma}_n \gamma_{n+\alpha}. \tag{C.9}$$

Under time evolution, we have

$$\gamma_n(t) = e^{iH_\alpha t}\gamma_n e^{-iH_\alpha t} = \cos(t)\gamma_n - \sin(t)\tilde{\gamma}_{n-\alpha},$$

and

$$\tilde{\gamma}_n(t) \equiv e^{iH_\alpha t}\tilde{\gamma}_n e^{-iH_\alpha t} = \cos(t)\tilde{\gamma}_n + \sin(t)\gamma_{n+\alpha}.$$

If we introduce the unitaries

$$U_0 \equiv e^{-ihH_0} \qquad \text{and} \qquad U_1 \equiv e^{iJH_1}, \tag{C.10}$$

we have, after a Jordan-Wigner transformation

$$F = U_1 U_0 = \mathrm{e}^{-\frac{i}{2}\sum_i JZ_iZ_{i+1}}\mathrm{e}^{-\frac{i}{2}\sum_i hX_i}, \tag{C.11}$$

which is indeed the Floquet operator of the driven Ising chain.

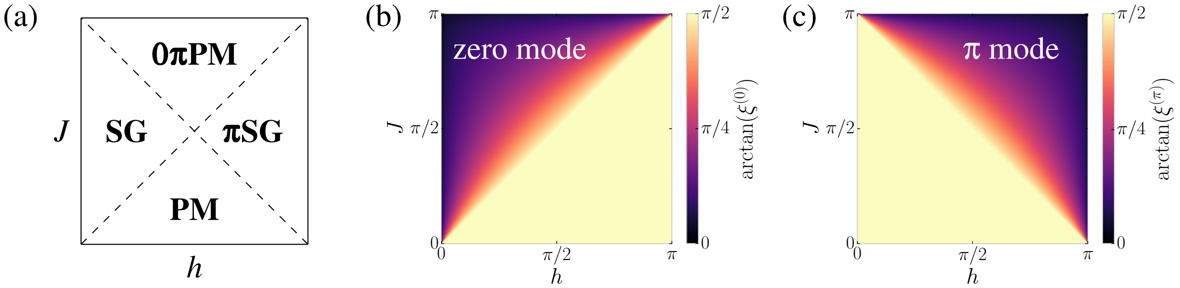

Figure C.4: **Topological edge modes at Floquet criticality.** (a) Schematic representation of the Floquet phase diagram of eq. (C.11). (b) The localization length of the zero mode; note that it remains *finite* (purple region) at the critical line between SG and 0πPM (c) Similarly, the localization length for the π mode remains finite at the transition between πSG and 0πPM.

Under Floquet stroboscopic evolution, we have

$$F^{\dagger}\gamma_n F = \begin{cases} \cos(J)\left(\cos(h)\gamma_n - \sin(h)\tilde{\gamma}_n\right) + \sin(J)\left(\cos(h)\tilde{\gamma}_{n-1} + \sin(h)\gamma_{n-1}\right) & \text{if } n > 1 \\ \cos(h)\gamma_1 - \sin(h)\tilde{\gamma}_1 & \text{if } n = 1 \end{cases} \tag{C.12}$$

$$F^{\dagger}\tilde{\gamma}_n F = \begin{cases} \cos(J)\left(\cos(h)\tilde{\gamma}_n + \sin(h)\gamma_n\right) - \sin(J)\left(\cos(h)\gamma_{n+1} - \sin(h)\tilde{\gamma}_{n+1}\right) & \text{if } n < N \\ \cos(h)\tilde{\gamma}_N + \sin(h)\gamma_N & \text{if } n = N. \end{cases} \tag{C.13}$$

From these expressions, we can construct the exact 0 and π modes for a half-infinite chain [YEM19]:

$$\gamma_L^{(0)} \equiv \sum_{n=1}^{\infty} \left(\frac{\tan(h/2)}{\tan(J/2)}\right)^{n-1} \left(\cos(h/2)\gamma_n - \sin(h/2)\tilde{\gamma}_n\right), \tag{C.14}$$

$$\gamma_L^{(\pi)} \equiv \sum_{n=1}^{\infty} \left(-\frac{1}{\tan(h/2)\tan(J/2)}\right)^{n-1} \left(\sin(h/2)\gamma_n + \cos(h/2)\tilde{\gamma}_n\right). \tag{C.15}$$

Here the labels 0 and π refer to the quasi-energies of those modes, that is, $F^{\dagger}\gamma_L^{(0)}F = \gamma_L^{(0)}$, and $F^{\dagger}\gamma_L^{(\pi)}F = \mathrm{e}^{i\pi}\gamma_L^{(\pi)}$.

Those edge modes are only meaningful when they are localized and normalizable. We plot the localization length of those modes as a function of J and h, and compare it to the known phase diagram of this model [KLMS16] (Fig. C.4). As expected, all four phases can be distinguished by their edge mode content: the paramagnet (PM) has no edge mode, the spin glass (SG) Ising symmetry-broken phase has a 0-Majorana mode, the π-spin glass (πSG) has a π-Majorana mode, and finally, the 0π-paramagnet (0πPM) has both a 0 and a π-Majorana mode. Remarkably, we find that the localization length of the zero mode remains *finite* at the critical line between SG and 0πPM. Similarly, the localization length for the π mode remains finite at the transition between πSG and 0πPM. In other words, the transitions out of the 0πPM phase are in the Ising* universality class. While we focused on the clean (uniform) case for simplicity, we have checked numerically that this carries over to the disordered case, and that both transition lines out of the 0πPM support exponentially localized edge modes. This Ising* critically is protected by a $\mathbb{Z}_2 \times \mathbb{Z}_2$ symmetry, where one of the $\mathbb{Z}_2$'s is emergent and inherited from time translation symmetry. Indeed, the SG-0πPM transition will have an Ising disorder operator μ (for the critical $\mathbb{Z}_2$ Ising symmetry) which is charged under the emergent $\mathbb{Z}_2$ time-translation symmetry (this protects the edge mode for the same reasons as the $\mathbb{Z}_2 \times \mathbb{Z}_2^T$-symmetric case studied in the main text); analogously, the Ising disorder operator μ for the πSG-0πPM transition corresponds to the time translation $\mathbb{Z}_2$ which is now charged under the explicit Ising symmetry $\prod_n X_n$. Disorder is essential to meaningfully discuss the phase structure of Floquet systems in the presence of interactions, so the relevant criticality in general is the random Ising* universality class discussed in the main text.

Appendix D

Machine learning holographic mapping by neural network renormalization group

D.1 Minimal Bulk Mutual Information Principle

The maximal real-space mutual information (maxRMI) principle proposed in Ref. [KR18, LRHK18] aims to maximize the mutual information between the coarse-grained field and the fine-grained field in the surrounding environment at a single RG step. In this section, we show that the maxRMI principle can be derived from our minimal bulk mutual information (minBMI) principle under certain assumptions.

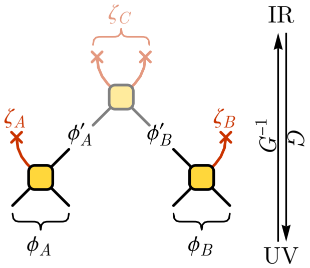

Figure D.1: Functional dependence of variables in the neural-RG network. Each block represents a bijective map.

Let us set up the problem based on Fig.D.1. Assuming ϕ_A and ϕ_B are field configurations in two neighboring regions A and B in the UV layer. Under one step of the RG transformation, ϕ_A gets mapped to the coarse grained variable ϕ'_A and the bulk variable ζ_A, and the mapping is bijective. Similarly, another bijection takes ϕ_B to ϕ'_B and ζ_B. Eventually, ϕ'_A and ϕ'_B will be mapped to the bulk field ζ_C in deeper IR layers. Therefore the random variables that appeared in Fig. D.1 are related by the following bijections f_A, f_B, f_C as

$$(\phi'_A,\zeta_A)=f_A(\phi_A),(\phi'_B,\zeta_B)=f_B(\phi_B),\zeta_C=f_C(\phi'_A,\phi'_B). \tag{D.1}$$

What are the information theoretical principles to guide the bijections f_A, f_B, f_C toward good RG transformations? We propose the minBMI principle that these bijections should minimize the mutual information among the bulk variables,

$$\min I(\zeta_A:\zeta_B)+I(\zeta_A:\zeta_C)+I(\zeta_B:\zeta_C). \tag{D.2}$$

Ref. [KR18, LRHK18] propose another principle, the maxRMI principle, that the RG transformation should maximize the mutual information between the coarse grained variable (such as ϕ'_A) and its environments (such as ϕ_B),

$$\max I(\phi'_A:\phi_B). \tag{D.3}$$

We can show that the objective of the maxRMI in Eq. (D.3) is consistent with the objective of the minBMI in Eq. (D.2) in the limit of UV-IR decoupling.

The minBMI principle aims to minimize mutual information among all bulk variables, both between different RG scales and within the same RG scale. Its objective has a broader scope than the maxRMI principle, because the later does not specify its objectives across the RG scales. So to make a connection between these two principles, one must first restrict the scope of the minBMI principle to a single layer. This can be achieved by assuming that there is no mutual information between bulk variables at different RG scales. In our setup, this corresponds to $I(\zeta_A,\zeta_B:\zeta_C)=0$, which factorizes the joint probability $p(\zeta_A,\zeta_B,\zeta_C)=p(\zeta_A,\zeta_B)p(\zeta_C)$ and decouples the bulk variables between UV and IR. As a result, the mutual information between any bulk variables across different RG scales vanishes $I(\zeta_A:\zeta_C)=I(\zeta_B:\zeta_C)=0$. This already minimizes the bulk mutual information across layers and reduces the minBMI objective in Eq. (D.2) to

$$\min I(\zeta_A:\zeta_B). \tag{D.4}$$

In this UV-IR decoupled limit, we can prove that $\max I(\phi'_A:\phi_B)$ and $\min I(\zeta_A:\zeta_B)$ are equivalent.

The proof starts by considering the mutual information between ϕ_A and ϕ_B. We can see that

$$\begin{aligned} I(\phi_A:\phi_B) &= I(\phi'_A,\zeta_A:\phi_B) \\ &= I(\phi'_A:\phi_B)+I(\zeta_A:\phi_B) \\ &= I(\phi'_A:\phi_B)+I(\zeta_A:\phi'_B,\zeta_B) \\ &= I(\phi'_A:\phi_B)+I(\zeta_A:\phi'_B)+I(\zeta_A:\zeta_B) \\ &= I(\phi'_A:\phi_B)+I(\zeta_A:\zeta_B). \end{aligned} \tag{D.5}$$

Here we have used the bijective property of f_A, f_B, f_C to obtain $I(\phi_A:\phi_B)=I(\phi'_A,\zeta_A:\phi_B)$, $I(\zeta_A:\phi_B)=I(\zeta_A:\phi'_B,\zeta_B)$ and $I(\zeta_A:\zeta_C)=I(\zeta_A:\phi'_A,\phi'_B)$. In the UV-IR decoupled limit, $I(\zeta_A:\zeta_C)=0$, so $I(\zeta_A:\phi'_A,\phi'_B)=0$, which further implies $I(\zeta_A:\phi'_A)=I(\zeta_A:\phi'_B)=0$. With these relations, all steps in Eq. (D.5) are justified. On the left hand side, $I(\phi_A:\phi_B)$ is determined by the field theory in the UV layer, which can be treated as a constant. For the given amount of information between regions A and B, Eq. (D.5) tells us that $I(\phi'_A:\phi_B)$ and $I(\zeta_A:\zeta_B)$ are competing for information resources. Therefore maximizing $I(\phi'_A:\phi_B)$ is equivalent to minimizing $I(\zeta_A:\zeta_B)$.

We can apply this argument layer by layer. Then to achieve the objective of the maxRMI principle, we need to minimize mutual information among bulk variables in the same RG scale, which is precisely the statement of the minBMI principle when restricted to each layer. In this sense, the maxRMI and minBMI principles are consistent. However, the minBMI principle actually relaxes the assumption that bulk variables at different RG scales are fully decoupled. Instead, we want to minimize mutual information among all bulk variables, including those across the scales. In this sense, the minBMI principle is more general than the maxRMI principle.

D.2 Design of Bijectors

We designed a set of symmetry-persevere bijectors to making sure that $U(1)$ symmetry of the boundary is preserved at each bijector. For the generative process, at each RG step, it takes the four complex degrees of freedom and they go through three layers of bijectors: $\mathcal{S}$, $\mathcal{O}$, and $\mathcal{A}$.

I. Scaling layer($\mathcal{S}$): At scaling layer, each complex variables ϕ_i is multiplied by a factor e^{λ_i}. The inverse and the Jacobian of this transformation can be obtained easily.

II. Orthogonal transformation layer($\mathcal{O}$): The orthogonal transformation in disentangler and decimator is in general an $O(4)$ transformation. In stead, we implemented it by stacking multiple $O(2)$ transformations. In Fig. D.2(a), each blue block represents the matrix:

$$M_{\text{blue}}(\theta_i) = \begin{pmatrix} \sin\theta_i & \cos\theta_i \\ \cos\theta_i & -\sin\theta_i \end{pmatrix}, \tag{D.6}$$

and the orange block in Fig. D.2(b) represents the matrix:

$$M_{\text{orange}}(\theta_i) = \begin{pmatrix} \cos\theta_i & -\sin\theta_i \\ \sin\theta_i & \cos\theta_i \end{pmatrix}. \tag{D.7}$$

θ_i in those blocks are training parameters. The arrangement of the type I and type II blocks are such designed that when $M_{\text{blue}}(\theta_i = \pi/4)$ and $M_{\text{orange}}(\theta_i = 0)$ the network reproduces the ideal EHM originally proposed in Ref. [Qi13]. We initialize the parameter to this ideal limit.

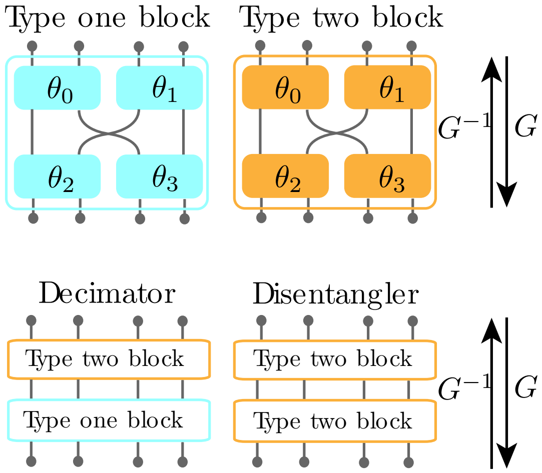

Figure D.2: Orthogonal transformation.

III. Non-linear layer($\mathscr{A}$) For non-linear part, we use the amplitude hyperbolic functions for complex field ϕ_i. In coarse-graining direction, it acts in the following,

$$\begin{aligned} \mathrm{Re}(\zeta) &= \sinh|\phi|\frac{\mathrm{Re}\,\phi}{|\phi|}; \\ \mathrm{Im}(\zeta) &= \sinh|\phi|\frac{\mathrm{Im}\,\phi}{|\phi|}. \end{aligned} \tag{D.8}$$

The corresponding inverse and Jacobian can be calculated easily.

D.3 Neural Network Training

All the training parameters of our neural RG network are contained in scaling bijectors and orthogonal transformation bijectors as illustrated in Appendix D.2. We imposed translation invariance of our network at each layer, due to translation invariance of the system at each energy scale. The total number of training parameters scale with $O(\log(N))$, where N is the size of boundary theory. The prefactor depends on the depth of bijector neural networks. In our case, the total number of training parameters are $24\log_2 N$. In order for faster convergence of the training, we first set learning rate for parameters contained in scaling bijectors as 10^{-2}, and gradually reduce it to 10^{-4}. And the learning rate for parameters contained in orthogonal transformation bijectors is always 10^{-4}.

D.4 Monte Carlo Sampling Efficiency

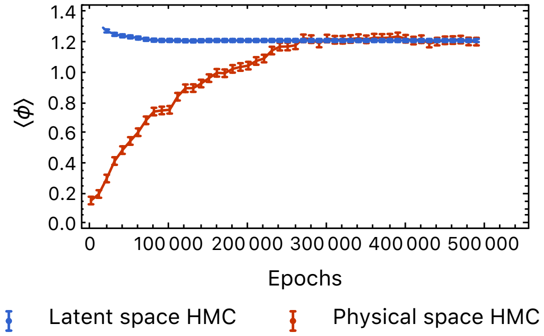

Figure D.3: MCMC result.

We tested numerical efficiency of our method by comparing convergence rate between Hamiltonian Monte Carlo(HMC) on the boundary system, and HMC in the bulk system. Both method are implemented using TensorFlow probability API with same parameters. The result is shown in Fig.D.3. As we can see, the HMC in the bulk system converges faster than the HMC on the boundary system.

D.5 Design of the Correlated Gaussian Prior

In finding the effective bulk field theory, we assume the bulk field is very massive. Under this assumption, higher-order interaction terms are irrelevant. Therefore, we use a correlated Gaussian distribution with positive definite kernel matrix K as our effective bulk field theory. We also assumed locality of our effective bulk field theory, which means K_{ij} is non-zero if and only if ζ_i and ζ_j are nearest neighbors in the bulk, including neighbors inter-scale and intra-scale. To further reduce the fitting parameters of matrix K, we also imposed translation invariance of the bulk field at each scale. This is reasonable, because our RG scheme also has the translation invariance at each scale.

To ensure matrix K is positive definite, we decomposed matrix K into a set of positive semi-definite matrix and a mass term. Particularly,

$$K = \sum_{\langle ij \rangle} \lambda_{ij} (|i\rangle\langle i| + |j\rangle\langle j| - |i\rangle\langle j| - |j\rangle\langle i|) + m\mathbb{I}, \tag{D.9}$$

where $\mathbb{I}$ is the identity matrix, and λ_{ij} and m are positive numbers. This ensures matrix K we constructed is positive definite.

Bibliography

[AAB^{+}15a] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dan Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.

[AAB^{+}15b] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.

[Aar07] Scott Aaronson. The learnability of quantum states. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 463(2088):3089–3114, Sep 2007.

[Aar17] Scott Aaronson. Shadow Tomography of Quantum States. *arXiv e-prints*, page arXiv:1711.01053, November 2017.

[ACQ21] Dorit Aharonov, Jordan Cotler, and Xiao-Liang Qi. Quantum Algorithmic Measurement. *arXiv e-prints*, page arXiv:2101.04634, January 2021.

[ADH15] A. Almheiri, X. Dong, and D. Harlow. Bulk locality and quantum error correction in AdS/CFT. *Journal of High Energy Physics*, 4:163, April 2015.

[ADR82] Alain Aspect, Jean Dalibard, and Gérard Roger. Experimental test of bell's inequalities using time-varying analyzers. *Phys. Rev. Lett.*, 49:1804–1807, Dec 1982.

[AF18] Yoshihiko Abe and Masafumi Fukuma. Gradient flow and the renormalization group. *Progress of Theoretical and Experimental Physics*, 2018:083B02, Aug 2018.

[AG04] Scott Aaronson and Daniel Gottesman. Improved simulation of stabilizer circuits. *Phys. Rev. A*, 70:052328, Nov 2004.

[AK16] K.-I. Aoki and T. Kobayashi. Restricted Boltzmann machines for the long range Ising models. *Modern Physics Letters B*, 30:1650401, December 2016.

[AKLT88] Ian Affleck, Tom Kennedy, Elliott H. Lieb, and Hal Tasaki. Valence bond ground states in isotropic quantum antiferromagnets. *Comm. Math. Phys.*, 115(3):477–528, 1988.

[AKR$^+$21] Michael S. Albergo, Gurtej Kanwar, Sébastien Racanière, Danilo J. Rezende, Julian M. Urban, Denis Boyda, Kyle Cranmer, Daniel C. Hackett, and Phiala E. Shanahan. Flow-based sampling for fermionic lattice field theories. *Physical Review D*, 104(11), Dec 2021.

[AKS19] M. S. Albergo, G. Kanwar, and P. E. Shanahan. Flow-based generative models for markov chain monte carlo in lattice field theory. *Phys. Rev. D*, 100:034515, Aug 2019.

[AN17] Moritz August and Xiaotong Ni. Using recurrent neural networks to optimize dynamical decoupling for quantum memory. *Physical Review A*, 95(1), Jan 2017.

[And72] P. W. Anderson. More is different. *Science*, 177(4047):393–396, 1972.

[AR19] Scott Aaronson and Guy N. Rothblum. Gentle Measurement of Quantum States and Differential Privacy. *arXiv e-prints*, page arXiv:1904.08747, April 2019.

[ASS21] Atithi Acharya, Siddhartha Saha, and Anirvan M. Sengupta. Informationally complete POVM-based shadow tomography. *arXiv e-prints*, page arXiv:2105.05992, May 2021.

[AY20] A. A. Akhtar and Yi-Zhuang You. Multiregion entanglement in locally scrambled quantum dynamics. , 102(13):134203, October 2020.

[B13] C. Bény. Deep learning and the renormalization group. *ArXiv e-prints*, January 2013.

[BB17] Jacob Biamonte and Ville Bergholm. Tensor networks in a nutshell, 2017.

[BCA19] Yimu Bao, Soonwon Choi, and Ehud Altman. Theory of the Phase Transition in Random Unitary Circuits with Measurements. *arXiv e-prints*, page arXiv:1908.04305, Aug 2019.

[BEJ$^+$19] Tiff Brydges, Andreas Elben, Petar Jurcevic, Benoît Vermersch, Christine Maier, Ben P. Lanyon, Peter Zoller, Rainer Blatt, and Christian F. Roos. Probing Rényi entanglement entropy via randomized measurements. *Science*, 364(6437):260–263, April 2019.

[Bel64a] J. S. Bell. On the einstein podolsky rosen paradox. *Physics Physique Fizika*, 1:195–200, Nov 1964.

[Bel64b] J. S. Bell. On the einstein podolsky rosen paradox. *Physics Physique Fizika*, 1:195–200, Nov 1964.

[BFH$^+$18] James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs, 2018.

[BGL13] V. Balasubramanian, M. Guica, and A. Lawrence. Holographic interpretations of the renormalization group. *Journal of High Energy Physics*, 1:115, January 2013.

[BGM18] M. J. S. Beach, A. Golubeva, and R. G. Melko. Machine learning vortices at the Kosterlitz-Thouless transition. , 97(4):045207, January 2018.

[BH86] J. D. Brown and Marc Henneaux. Central charges in the canonical realization of asymptotic symmetries: an example from three-dimensional gravity. *Comm. Math. Phys.*, 104(2):207–226, 1986.

[BKL$^+$17] Fernando G. S. L. Brandão, Amir Kalev, Tongyang Li, Cedric Yen-Yu Lin, Krysta M. Svore, and Xiaodi Wu. Quantum SDP Solvers: Large Speed-ups, Optimality, and Applications to Quantum Learning. *arXiv e-prints*, page arXiv:1710.02581, October 2017.

[BKPV18] William Berdanier, Michael Kolodrubetz, S. A. Parameswaran, and Romain Vasseur. Floquet quantum criticality. *Proceedings of the National Academy of Sciences*, 115(38):9491–9496, 2018.

[BO15] Cédric Bény and Tobias J Osborne. The renormalization group via statistical inference. *New Journal of Physics*, 17(8):083005, aug 2015.

[BV14] Yasaman Bahri and Ashvin Vishwanath. Detecting majorana fermions in quasi-one-dimensional topological phases using nonlocal order parameters. *Phys. Rev. B*, 89:155135, Apr 2014.

[CBQA19] Soonwon Choi, Yimu Bao, Xiao-Liang Qi, and Ehud Altman. Quantum error correction and entanglement phase transition in random unitary circuits with projective measurements. *arXiv e-prints*, page arXiv:1903.05124, Mar 2019.

[CCMK17] Kelvin Ch'ng, Juan Carrasquilla, Roger G. Melko, and Ehsan Khatami. Machine learning phases of strongly correlated fermions. *Phys. Rev. X*, 7:031038, Aug 2017.

[CFS02] Carlton M. Caves, Christopher A. Fuchs, and Rüdiger Schack. Unknown quantum states: The quantum de Finetti representation. *Journal of Mathematical Physics*, 43(9):4537–4559, September 2002.

[CGLW12] Xie Chen, Zheng-Cheng Gu, Zheng-Xin Liu, and Xiao-Gang Wen. Symmetry-protected topological orders in interacting bosonic systems. *Science*, 338(6114):1604–1606, 2012.

[CGLW13] Xie Chen, Zheng-Cheng Gu, Zheng-Xin Liu, and Xiao-Gang Wen. Symmetry protected topological orders and the group cohomology of their symmetry group. *Phys. Rev. B*, 87:155114, Apr 2013.

[CGW11] Xie Chen, Zheng-Cheng Gu, and Xiao-Gang Wen. Complete classification of one-dimensional gapped quantum phases in interacting spin systems. *Phys. Rev. B*, 84:235128, Dec 2011.

[CGW$^+$20] Peter Cha, Paul Ginsparg, Felix Wu, Juan Carrasquilla, Peter L. McMahon, and Eun-Ah Kim. Attention-based Quantum Tomography. *arXiv e-prints*, page arXiv:2006.12469, June 2020.

[CHhH02] Murray Campbell, A.Joseph Hoane, and Feng hsiung Hsu. Deep blue. *Artificial Intelligence*, 134(1):57–83, 2002.

[CHLY17] Jordan Cotler, Nicholas Hunter-Jones, Junyu Liu, and Beni Yoshida. Chaos, complexity, and random matrices. *Journal of High Energy Physics*, 2017(11):48, November 2017.

[CHN18] Andrea Carosso, Anna Hasenfratz, and Ethan T. Neil. Nonperturbative Renormalization of Operators in Near-Conformal Systems Using Gradient Flows. , 121:201601, Nov 2018.

[Cho75] Man-Duen Choi. Completely positive linear maps on complex matrices. *Linear Algebra and its Applications*, 10(3):285–290, 1975.

[CJR$^+$12] Jianxin Chen, Zhengfeng Ji, Mary Beth Ruskai, Bei Zeng, and Duan-Lu Zhou. Comment on some results of Erdahl and the convex structure of reduced density matrices. *Journal of Mathematical Physics*, 53(7):072203–072203, July 2012.

[CJZZ12] Jianxin Chen, Zhengfeng Ji, Bei Zeng, and D. L. Zhou. From ground states to local Hamiltonians. , 86(2):022339, August 2012.

[CLP$^+$19] Juan Carrasquilla, Di Luo, Felipe Pérez, Ashley Milsted, Bryan K. Clark, Maksims Volkovs, and Leandro Aolita. Probabilistic Simulation of Quantum Circuits with the Transformer. *arXiv e-prints*, page arXiv:1912.11052, December 2019.

[CLV14] Xie Chen, Yuan-Ming Lu, and Ashvin Vishwanath. Symmetry-protected topological phases from decorated domain walls. *Nature Communications*, 5(1):3507, 2014.

[CLW11] Xie Chen, Zheng-Xin Liu, and Xiao-Gang Wen. Two-dimensional symmetry-protected topological orders and their protected gapless edge excitations. *Phys. Rev. B*, 84:235141, Dec 2011.

[CLZ$^+$18] Yuan Cao, Yu-Huai Li, Wen-Jie Zou, Zheng-Ping Li, Qi Shen, Sheng-Kai Liao, Ji-Gang Ren, Juan Yin, Yu-Ao Chen, Cheng-Zhi Peng, and Jian-Wei Pan. Bell test over extremely high-loss channels: Towards distributing entangled photon pairs between earth and the moon. *Phys. Rev. Lett.*, 120:140405, Apr 2018.

[CM17a] J. Carrasquilla and R. G. Melko. Machine learning phases of matter. *Nature Physics*, 13:431–434, February 2017.

[CM17b] Juan Carrasquilla and Roger G. Melko. Machine learning phases of matter. *Nature Physics*, 13(5):431–434, 2017.

[CMH$^+$21] Jordan S. Cotler, Daniel K. Mark, Hsin-Yuan Huang, Felipe Hernandez, Joonhee Choi, Adam L. Shaw, Manuel Endres, and Soonwon Choi. Emergent quantum state designs from individual many-body wavefunctions. *arXiv e-prints*, page arXiv:2103.03536, March 2021.

[CNPS19] Amos Chan, Rahul M. Nandkishore, Michael Pretko, and Graeme Smith. Unitary-projective entanglement dynamics. , 99(22):224307, Jun 2019.

[CPF$^+$10a] Marcus Cramer, Martin B. Plenio, Steven T. Flammia, Rolando Somma, David Gross, Stephen D. Bartlett, Olivier Land on-Cardinal, David Poulin, and Yi-Kai Liu. Efficient quantum state tomography. *Nature Communications*, 1(9):149, December 2010.

[CPF$^+$10b] Marcus Cramer, Martin B. Plenio, Steven T. Flammia, Rolando Somma, David Gross, Stephen D. Bartlett, Olivier Landon-Cardinal, David Poulin, and Yi-Kai Liu. Efficient quantum state tomography. *Nature Communications*, 1(1):149, 2010.

[CRBD18] Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David Duvenaud. Neural Ordinary Differential Equations. *arXiv e-prints*, page arXiv:1806.07366, Jun 2018.

[CŚ06] Benoît Collins and Piotr Śniady. Integration with Respect to the Haar Measure on Unitary, Orthogonal and Symplectic Group. *Communications in Mathematical Physics*, 264(3):773–795, Jun 2006.

[CSM$^+$21] Joonhee Choi, Adam L. Shaw, Ivaylo S. Madjarov, Xin Xie, Jacob P. Covey, Jordan S. Cotler, Daniel K. Mark, Hsin-Yuan Huang, Anant Kale, Hannes Pichler, Fernando G. S. L. Brandão, Soonwon Choi, and Manuel Endres. Emergent Randomness and Benchmarking from Many-Body Quantum Chaos. *arXiv e-prints*, page arXiv:2103.03535, March 2021.

[CT11] Meng Cheng and Hong-Hao Tu. Majorana edge states in interacting two-chain ladders of fermions. *Phys. Rev. B*, 84:094503, Sep 2011.

[CT17] G. Carleo and M. Troyer. Solving the quantum many-body problem with artificial neural networks. *Science*, 355:602–606, February 2017.

[CTMA18] Juan Carrasquilla, Giacomo Torlai, Roger G. Melko, and Leandro Aolita. Reconstructing quantum states with generative models. *arXiv e-prints*, page arXiv:1810.10584, October 2018.

[CTMA19] Juan Carrasquilla, Giacomo Torlai, Roger G. Melko, and Leandro Aolita. Reconstructing quantum states with generative models. *Nature Machine Intelligence*, 1(3):155–161, 2019.

[CW20] Jordan Cotler and Frank Wilczek. Quantum overlapping tomography. *Phys. Rev. Lett.*, 124:100401, Mar 2020.

[CYT$^+$18] Chun Chen, Wei Yan, C. S. Ting, Yan Chen, and F. J. Burnell. Flux-stabilized majorana zero modes in coupled one-dimensional fermi wires. *Phys. Rev. B*, 98:161106, Oct 2018.

[CYZF20a] Senrui Chen, Wenjun Yu, Pei Zeng, and Steven T. Flammia. Robust shadow estimation. *arXiv e-prints*, page arXiv:2011.09636, November 2020.

[CYZF20b] Senrui Chen, Wenjun Yu, Pei Zeng, and Steven T. Flammia. Robust shadow estimation. *arXiv e-prints*, page arXiv:2011.09636, November 2020.

[CZ19] Xiao Chen and Tianci Zhou. Quantum chaos dynamics in long-range power law interaction systems. , 100(6):064305, Aug 2019.

[DCLT19] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding, 2019.

[DHY$^+$21] Carlos M. Duque, Hong-Ye Hu, Yi-Zhuang You, Vedika Khemani, Ruben Verresen, and Romain Vasseur. Topological and symmetry-enriched random quantum critical points. *Phys. Rev. B*, 103:L100207, Mar 2021.

[Di4] Lajos Diósi. Three-party pure quantum states are determined by two two-party reduced states. *Phys. Rev. A*, 70:010302, Jul 2004.

[DM80] Chandan Dasgupta and Shang-keng Ma. Low-temperature properties of the random heisenberg antiferromagnetic chain. *Phys. Rev. B*, 22:1305–1319, Aug 1980.

[DR93] J. Darriet and L.P. Regnault. The compound y2banio5: A new example of a haldane gap in a s = 1 magnetic chain. *Solid State Communications*, 86(7):409 – 412, 1993.

[DRLB20] A. Dikme, N. Reichel, A. Laghaout, and G. Björk. Measuring the mermin-peres magic square using an online quantum computer, 2020.

[DSB16] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using Real NVP. *arXiv e-prints*, page arXiv:1605.08803, May 2016.

[DSDD19] Sankar Das Sarma, Dong-Ling Deng, and Lu-Ming Duan. Machine learning meets quantum physics. *Physics Today*, 72(3):48–54, Mar 2019.

[dVV00] J. de Boer, E. Verlinde, and H. Verlinde. On the holographic renormalization group. *Journal of High Energy Physics*, 8:3, August 2000.

[DZ18] Xiao Dong and Ling Zhou. Spacetime as the optimal generative network of quantum states: a roadmap to QM=GR? *arXiv e-prints*, page arXiv:1804.07908, Apr 2018.

[EBM18] Stavros Efthymiou, Matthew J. S. Beach, and Roger G. Melko. Super-resolving the Ising model with convolutional neural networks. *arXiv e-prints*, page arXiv:1810.02372, Oct 2018.

[EBN16] Dominic V. Else, Bela Bauer, and Chetan Nayak. Floquet time crystals. *Phys. Rev. Lett.*, 117:090402, Aug 2016.

[EG20a] Dax Enshan Koh and Sabee Grewal. Classical Shadows with Noise. *arXiv e-prints*, page arXiv:2011.11580, November 2020.

[EG20b] Dax Enshan Koh and Sabee Grewal. Classical Shadows with Noise. *arXiv e-prints*, page arXiv:2011.11580, November 2020.

[EKH$^+$20] Andreas Elben, Richard Kueng, Hsin-Yuan Robert Huang, Rick van Bijnen, Christian Kokail, Marcello Dalmonte, Pasquale Calabrese, Barbara Kraus, John Preskill, Peter Zoller, and Benoît Vermersch. Mixed-State Entanglement from Local Randomized Measurements. , 125(20):200501, November 2020.

[EN16] Dominic V. Else and Chetan Nayak. Classification of topological phases in periodically driven interacting systems. *Phys. Rev. B*, 93(20):201103, May 2016.

[EPR35] A. Einstein, B. Podolsky, and N. Rosen. Can quantum-mechanical description of physical reality be considered complete? *Phys. Rev.*, 47:777–780, May 1935.

[EVD$^+$18] A. Elben, B. Vermersch, M. Dalmonte, J. I. Cirac, and P. Zoller. Rényi entropies from random quenches in atomic hubbard and spin models. *Phys. Rev. Lett.*, 120:050406, Feb 2018.

[EVRZ19] A. Elben, B. Vermersch, C. F. Roos, and P. Zoller. Statistical correlations between locally randomized measurements: A toolbox for probing entanglement in many-body quantum states. , 99(5):052323, May 2019.

[EWL$^+$20] Sepehr Ebadi, Tout T. Wang, Harry Levine, Alexander Keesling, Giulia Semeghini, Ahmed Omran, Dolev Bluvstein, Rhine Samajdar, Hannes Pichler, Wen Wei Ho, Soonwon Choi, Subir Sachdev, Markus Greiner, Vladan Vuletic, and Mikhail D. Lukin. Quantum Phases of Matter on a 256-Atom Programmable Quantum Simulator. *arXiv e-prints*, page arXiv:2012.12281, December 2020.

[FCY$^+$04] David Fattal, Toby S. Cubitt, Yoshihisa Yamamoto, Sergey Bravyi, and Isaac L. Chuang. Entanglement in the stabilizer formalism. *arXiv e-prints*, pages quant–ph/0406168, June 2004.

[Fen12] Paul Fendley. Parafermionic edge zero modes inZn-invariant spin chains. *Journal of Statistical Mechanics: Theory and Experiment*, 2012(11):P11020, nov 2012.

[Fen16] Paul Fendley. Strong zero modes and eigenstate phase transitions in the XYZ/interacting majorana chain. *Journal of Physics A: Mathematical and Theoretical*, 49(30):30LT01, jun 2016.

[Fis92] Daniel S. Fisher. Random transverse field ising spin chains. *Phys. Rev. Lett.*, 69:534–537, Jul 1992.

[Fis94] Daniel S. Fisher. Random antiferromagnetic quantum spin chains. *Phys. Rev. B*, 50:3799–3821, Aug 1994.

[Fis95] Daniel S. Fisher. Critical behavior of random transverse-field ising spin chains. *Phys. Rev. B*, 51:6411–6461, Mar 1995.

[FK07] Liang Fu and C. L. Kane. Topological insulators with inversion symmetry. *Phys. Rev. B*, 76:045302, Jul 2007.

[FK11] Lukasz Fidkowski and Alexei Kitaev. Topological phases of fermions in one dimension. *Phys. Rev. B*, 83:075103, Feb 2011.

[FLNF11] Lukasz Fidkowski, Roman M. Lutchyn, Chetan Nayak, and Matthew P. A. Fisher. Majorana zero modes in one-dimensional quantum wires without long-ranged superconducting order. *Phys. Rev. B*, 84:195436, Nov 2011.

[Fuj16] Kazuo Fujikawa. The gradient flow in $\lambda\phi^4$ theory. *Journal of High Energy Physics*, 2016:21, Mar 2016.

[FVVY20] Ruihua Fan, Sagar Vijay, Ashvin Vishwanath, and Yi-Zhuang You. Self-organized error correction in random unitary circuits with measurement. 2020.

[GAW$^+$21] Xun Gao, Eric R. Anschuetz, Sheng-Tao Wang, J. Ignacio Cirac, and Mikhail D. Lukin. Enhancing generative models via quantum correlations, 2021.

[GCB$^+$18] Will Grathwohl, Ricky T. Q. Chen, Jesse Bettencourt, Ilya Sutskever, and David Duvenaud. FFJORD: Free-form Continuous Dynamics for Scalable Reversible Generative Models. *arXiv e-prints*, page arXiv:1810.01367, Oct 2018.

[GH19] Michael J. Gullans and David A. Huse. Dynamical purification phase transition induced by quantum measurements. *arXiv e-prints*, page arXiv:1905.05195, May 2019.

[GHST18] Hrant Gharibyan, Masanori Hanada, Stephen H. Shenker, and Masaki Tezuka. Onset of random matrix behavior in scrambling systems. *Journal of High Energy Physics*, 2018(7):124, Jul 2018.

[GKH$^+$20] Michael J. Gullans, Stefan Krastanov, David A. Huse, Liang Jiang, and Steven T. Flammia. Quantum coding with low-depth random circuits. *arXiv e-prints*, page arXiv:2010.09775, October 2020.

[GKKT18] Madalin Guta, Jonas Kahn, Richard Kueng, and Joel A. Tropp. Fast state tomography with optimal error bounds. *arXiv e-prints*, page arXiv:1809.11162, September 2018.

[GKP98] S. S. Gubser, I. R. Klebanov, and A. M. Polyakov. Gauge theory correlators from non-critical string theory. *Physics Letters B*, 428:105–114, May 1998.

[GLB17] K. Guther, N. Lang, and H. P. Büchler. Ising anyonic topological phase of interacting fermions in one dimension. *Phys. Rev. B*, 96:121109, Sep 2017.

[GLW$^+$16] Y. Gu, C. H. Lee, X. Wen, G. Y. Cho, S. Ryu, and X.-L. Qi. Holographic duality between $(2+1)$-d quantum anomalous Hall state and $(3+1)$-d topological insulators. *ArXiv e-prints*, May 2016.

[GMS18] Diego García-Martín and Germán Sierra. Five experimental tests on the 5-qubit ibm quantum computer. *Journal of Applied Mathematics and Physics*, 06(07):1460–1475, 2018.

[Got98a] Daniel Gottesman. The heisenberg representation of quantum computers, 1998.

[Got98b] Daniel Gottesman. The Heisenberg Representation of Quantum Computers. *arXiv e-prints*, pages quant–ph/9807006, July 1998.

[GS17] W.-C. Gan and F.-W. Shu. Holography as deep learning. *International Journal of Modern Physics D*, 26:1743020, 2017.

[GV12] Tarun Grover and Ashvin Vishwanath. Quantum Criticality in Topological Insulators and Superconductors: Emergence of Strongly Coupled Majoranas and Supersymmetry. *arXiv e-prints*, page arXiv:1206.1332, June 2012.

[GW09] Zheng-Cheng Gu and Xiao-Gang Wen. Tensor-entanglement-filtering renormalization approach and symmetry-protected topological order. *Phys. Rev. B*, 80:155131, Oct 2009.

[GZJ21] Roy J. Garcia, You Zhou, and Arthur Jaffe. Quantum scrambling with classical shadows. *arXiv e-prints*, page arXiv:2102.01008, February 2021.

[Hal83] F. D. M. Haldane. Nonlinear field theory of large-spin heisenberg antiferromagnets: Semiclassically quantized solitons of the one-dimensional easy-axis néel state. *Phys. Rev. Lett.*, 50:1153–1156, Apr 1983.

[Has05] Martin Hasenbusch. The two-dimensional XY model at the transition temperature: a high-precision Monte Carlo study. *Journal of Physics A Mathematical General*, 38:5869–5883, Jul 2005.

[Has19] Koji Hashimoto. AdS/CFT as a deep Boltzmann machine. *arXiv e-prints*, page arXiv:1903.04951, Mar 2019.

[HCY21] Hong-Ye Hu, Soonwon Choi, and Yi-Zhuang You. Classical Shadow Tomography with Locally Scrambled Quantum Dynamics. *arXiv e-prints*, page arXiv:2107.04817, July 2021.

[HG20] N. L. Holanda and M. A. R. Griffith. Machine learning topological phases in real space. *Phys. Rev. B*, 102:054107, Aug 2020.

[HHA$^+$21] Daniel C. Hackett, Chung-Chun Hsieh, Michael S. Albergo, Denis Boyda, Jiunn-Wei Chen, Kai-Feng Chen, Kyle Cranmer, Gurtej Kanwar, and Phiala E. Shanahan. Flow-based sampling for multimodal distributions in lattice field theory, 2021.

[HHJ$^+$17a] J. Haah, A. W. Harrow, Z. Ji, X. Wu, and N. Yu. Sample-optimal tomography of quantum states. *IEEE Transactions on Information Theory*, 63(9):5628–5641, 2017.

[HHJ$^+$17b] Jeongwan Haah, Aram W. Harrow, Zhengfeng Ji, Xiaodi Wu, and Nengkun Yu. Sample-optimal tomography of quantum states. *IEEE Transactions on Information Theory*, 63(9):5628–5641, 2017.

[HK10] M. Z. Hasan and C. L. Kane. Colloquium: Topological insulators. *Rev. Mod. Phys.*, 82:3045–3067, Nov 2010.

[HKP20] Hsin-Yuan Huang, Richard Kueng, and John Preskill. Predicting many properties of a quantum system from very few measurements. *Nature Physics*, 16(10):1050–1057, June 2020.

[HKT$^+$21] Hsin-Yuan Huang, Richard Kueng, Giacomo Torlai, Victor V. Albert, and John Preskill. Provably efficient machine learning for quantum many-body problems. *arXiv e-prints*, page arXiv:2106.12627, June 2021.

[HLWY20] Hong-Ye Hu, Shuo-Hui Li, Lei Wang, and Yi-Zhuang You. Machine learning holographic mapping by neural network renormalization group. *Phys. Rev. Research*, 2:023369, Jun 2020.

[HNO14] D. A. Huse, R. Nandkishore, and V. Oganesyan. Phenomenology of fully many-body-localized systems. , 90(17):174202, November 2014.

[HP97] M. Hasenbusch and K. Pinn. Computing the roughening transition of Ising and solid-on-solid models by BCSOS model matching. *Journal of Physics A Mathematical General*, 30:63–80, Jan 1997.

[HP11] Idse Heemskerk and Joseph Polchinski. Holographic and Wilsonian renormalization groups. *Journal of High Energy Physics*, 2011:31, Jun 2011.

[HR17] Fenner Harper and Rahul Roy. Floquet Topological Order in Interacting Systems of Bosons and Fermions. *Phys. Rev. Lett.*, 118(11):115301, Mar 2017.

[HSTT18a] Koji Hashimoto, Sotaro Sugishita, Akinori Tanaka, and Akio Tomiya. Deep learning and holographic QCD. , 98:106014, Nov 2018.

[HSTT18b] Koji Hashimoto, Sotaro Sugishita, Akinori Tanaka, and Akio Tomiya. Deep learning and the ads/cft correspondence. *Phys. Rev. D*, 98:046019, Aug 2018.

[HW17] L. Huang and L. Wang. Accelerated Monte Carlo simulations with restricted Boltzmann machines. , 95(3):035105, January 2017.

[HY21] Hong-Ye Hu and Yi-Zhuang You. Hamiltonian-Driven Shadow Tomography of Quantum States. *arXiv e-prints*, page arXiv:2102.10132, February 2021.

[HY22] Hong-Ye Hu and Yi-Zhuang You. Hamiltonian-driven shadow tomography of quantum states. *Phys. Rev. Research*, 4:013054, Jan 2022.

[IMR$^+$15] Fernando Iemini, Leonardo Mazza, Davide Rossini, Rosario Fazio, and Sebastian Diehl. Localized majorana-like modes in a number-conserving setting: An exactly solvable model. *Phys. Rev. Lett.*, 115:156402, Oct 2015.

[IR98] Ferenc Iglói and Heiko Rieger. Random transverse Ising spin chain and random walks. , 57(18):11404–11420, May 1998.

[ITe] *ITensor Library (version 2.1) http://itensor.org*.

[Jam72] A. Jamiołkowski. Linear transformations which preserve trace and positive semidefiniteness of operators. *Reports on Mathematical Physics*, 3(4):275–278, 1972.

[JGP17] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax, 2017.

[JKMW01] Daniel F. V. James, Paul G. Kwiat, William J. Munro, and Andrew G. White. Measurement of qubits. , 64(5):052312, November 2001.

[JLSL18] Hong-Chen Jiang, Zi-Xiang Li, Alexander Seidel, and Dung-Hai Lee. Symmetry protected topological luttinger liquids and the phase transition between them. *Science Bulletin*, 63(12):753 – 758, 2018.

[JSW19] Wenjie Ji, Shu-Heng Shao, and Xiao-Gang Wen. Topological Transition on the Conformal Manifold. *arXiv e-prints*, page arXiv:1909.01425, September 2019.

[JV19] N. G. Jones and R. Verresen. Asymptotic correlations in gapped and critical topological phases of 1d quantum systems. *Journal of Statistical Physics*, 175(6):1164–1213, Jun 2019.

[JYVL19] Chao-Ming Jian, Yi-Zhuang You, Romain Vasseur, and Andreas W. W. Ludwig. Measurement-induced criticality in random quantum circuits. *arXiv e-prints*, page arXiv:1908.08051, Aug 2019.

[KAAY20] Wei-Ting Kuo, A. A. Akhtar, Daniel P. Arovas, and Yi-Zhuang You. Markovian entanglement dynamics under locally scrambled quantum evolution. , 101(22):224202, June 2020.

[KAB$^+$20] Gurtej Kanwar, Michael S. Albergo, Denis Boyda, Kyle Cranmer, Daniel C. Hackett, Sébastien Racanière, Danilo Jimenez Rezende, and Phiala E. Shanahan. Equivariant flow-based sampling for lattice gauge theory. *Phys. Rev. Lett.*, 125:121601, Sep 2020.

[Kad66] Leo P. Kadanoff. Scaling laws for ising models near T_c. *Physics Physique Fizika*, 2:263–272, Jun 1966.

[KB15] Anna Keselman and Erez Berg. Gapless symmetry-protected topological phase of fermions in one dimension. *Phys. Rev. B*, 91:235309, Jun 2015.

[KBA18] Anna Keselman, Erez Berg, and Patrick Azaria. From one-dimensional charge conserving superconductors to the gapless haldane phase. *Phys. Rev. B*, 98:214501, Dec 2018.

[KC15] Nikolaos Kainaris and Sam T. Carr. Emergent topological properties in interacting one-dimensional systems with spin-orbit coupling. *Phys. Rev. B*, 92:035139, Jul 2015.

[KD18] Diederik P. Kingma and Prafulla Dhariwal. Glow: Generative Flow with Invertible 1x1 Convolutions. *arXiv e-prints*, page arXiv:1807.03039, Jul 2018.

[KDB$^+$13] Christina V. Kraus, Marcello Dalmonte, Mikhail A. Baranov, Andreas M. Läuchli, and P. Zoller. Majorana edge states in atomic wires coupled by pair hopping. *Phys. Rev. Lett.*, 111:173004, Oct 2013.

[KDK$^+$21] Niklas Käming, Anna Dawid, Korbinian Kottmann, Maciej Lewenstein, Klaus Sengstock, Alexandre Dauphin, and Christof Weitenberg. Unsupervised machine learning of topological phase transitions from experimental data. *Machine Learning: Science and Technology*, 2(3):035037, jul 2021.

[KHMZ$^+$20] Stefan Krastanov, Kade Head-Marsden, Sisi Zhou, Steven T. Flammia, Liang Jiang, and Prineha Narang. Unboxing quantum black box models: Learning non-markovian dynamics, 2020.

[Kit09] Alexei Kitaev. Periodic table for topological insulators and superconductors. *AIP Conference Proceedings*, 1134(1):22–30, 2009.

[KLMS16] Vedika Khemani, Achilleas Lazarides, Roderich Moessner, and S. L. Sondhi. Phase structure of driven quantum systems. *Phys. Rev. Lett.*, 116:250401, Jun 2016.

[KM04] J.P. Keating and F. Mezzadri. Random matrix theory and entanglement in quantum spin chains. *Communications in Mathematical Physics*, 252(1):543–579, 2004.

[KM05] C. L. Kane and E. J. Mele. Z_2 topological order and the quantum spin hall effect. *Phys. Rev. Lett.*, 95:146802, Sep 2005.

[KMS19] Vedika Khemani, Roderich Moessner, and S. L. Sondhi. A Brief History of Time Crystals. *arXiv e-prints*, page arXiv:1910.10745, October 2019.

[KPP$^+$] Byungmin Kang, S. A. Parameswaran, Andrew C. Potter, Romain Vasseur, and Snir Gazit. Infinite-randomness superuniversality at deconfinement transitions in two dimensions. *In preparation.*

[KPR20] Changjun Kim, Kyungdeock Daniel Park, and June-Koo Rhee. Quantum error mitigation with artificial neural network. *IEEE Access*, 8:188853–188860, 2020.

[KR18] M. Koch-Janusz and Z. Ringel. Mutual information, neural networks and the renormalization group. *Nature Physics*, 14:578–582, June 2018.

[Kra71] K. Kraus. General state changes in quantum theory. *Annals of Physics*, 64(2):311–335, June 1971.

[KSGC17] Nikolaos Kainaris, Raul A. Santos, D. B. Gutman, and Sam T. Carr. Interaction induced topological protection in one-dimensional conductors. *Fortschritte der Physik*, 65(6-8):1600054, 2017.

[KSH17a] Charles L. Kane, Ady Stern, and Bertrand I. Halperin. Pairing in luttinger liquids and quantum hall states. *Phys. Rev. X*, 7:031009, Jul 2017.

[KSH17b] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. *Commun. ACM*, 60(6):84–90, may 2017.

[KW13] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.

[KW14] Diederik P Kingma and Max Welling. Auto-encoding variational bayes, 2014.

[KWSDS11] J. P. Kestner, Bin Wang, Jay D. Sau, and S. Das Sarma. Prediction of a gapless topological haldane liquid phase in a one-dimensional cold polar molecular lattice. *Phys. Rev. B*, 83:174409, May 2011.

[KYL19] Jack Kemp, Norman Y. Yao, and Chris R. Laumann. Symmetry enhanced boundary qubits at infinite temperature. *arXiv e-prints*, page arXiv:1912.05546, December 2019.

[L10] Martin Lüscher. Trivializing Maps, the Wilson Flow and the HMC Algorithm. *Communications in Mathematical Physics*, 293:899, Feb 2010.

[LB15] Nicolai Lang and Hans Peter Büchler. Topological states in a microscopic model of interacting fermions. *Phys. Rev. B*, 92:041118, Jul 2015.

[LBBH98] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.

[LBD$^+$89] Y. LeCun, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel. Backpropagation applied to handwritten zip code recognition. *Neural Computation*, 1(4):541–551, 1989.

[LBD$^+$90] Yann LeCun, Bernhard Boser, John Denker, Donnie Henderson, R. Howard, Wayne Hubbard, and Lawrence Jackel. Handwritten digit recognition with a back-propagation network. In D. Touretzky, editor, *Advances in Neural Information Processing Systems*, volume 2. Morgan-Kaufmann, 1990.

[LBH15] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015.

[LCF18] Yaodong Li, Xiao Chen, and Matthew P. A. Fisher. Quantum Zeno effect and the many-body entanglement transition. , 98(20):205136, Nov 2018.

[LCF19] Yaodong Li, Xiao Chen, and Matthew P. A. Fisher. Measurement-driven entanglement transition in hybrid quantum circuits. *arXiv e-prints*, page arXiv:1901.08092, Jan 2019.

[LG12] Michael Levin and Zheng-Cheng Gu. Braiding statistics approach to symmetry-protected topological phases. *Phys. Rev. B*, 86:115109, Sep 2012.

[LMH$^+$17a] B. P. Lanyon, C. Maier, M. Holzäpfel, T. Baumgratz, C. Hempel, P. Jurcevic, I. Dhand, A. S. Buyskikh, A. J. Daley, M. Cramer, M. B. Plenio, R. Blatt, and C. F. Roos. Efficient tomography of a quantum many-body system. *Nature Physics*, 13(12):1158–1162, 2017.

[LMH$^+$17b] B. P. Lanyon, C. Maier, M. Holzäpfel, T. Baumgratz, C. Hempel, P. Jurcevic, I. Dhand, A. S. Buyskikh, A. J. Daley, M. Cramer, M. B. Plenio, R. Blatt, and C. F. Roos. Efficient tomography of a quantum many-body system. *Nature Physics*, 13(12):1158–1162, December 2017.

[LPW02] N. Linden, S. Popescu, and W. K. Wootters. Almost every pure state of three qubits is completely determined by its two-particle reduced density matrices. *Phys. Rev. Lett.*, 89:207901, Oct 2002.

[LQ15] C. H. Lee and X.-L. Qi. Exact holographic mapping in free fermion systems. *ArXiv e-prints*, March 2015.

[LQMF17] J. Liu, Y. Qi, Z. Y. Meng, and L. Fu. Self-learning Monte Carlo method. , 95(4):041101, January 2017.

[LRHK18] Patrick M. Lenggenhager, Zohar Ringel, Sebastian D. Huber, and Maciej Koch-Janusz. Optimal Renormalization Group Transformation from Information Theory. *arXiv e-prints*, page arXiv:1809.09632, Sep 2018.

[LSH$^+$13] Nima Lashkari, Douglas Stanford, Matthew Hastings, Tobias Osborne, and Patrick Hayden. Towards the fast scrambling conjecture. *Journal of High Energy Physics*, 2013:22, Apr 2013.

[LSQ$^+$16] J. Liu, H. Shen, Y. Qi, Z. Y. Meng, and L. Fu. Self-Learning Monte Carlo Method in Fermion Systems. *ArXiv e-prints*, November 2016.

[LTR17] H. W. Lin, M. Tegmark, and D. Rolnick. Why Does Deep and Cheap Learning Work So Well? *Journal of Statistical Physics*, 168:1223–1247, September 2017.

[LV12] Yuan-Ming Lu and Ashvin Vishwanath. Theory and classification of interacting integer topological phases in two dimensions: A chern-simons approach. *Phys. Rev. B*, 86:125119, Sep 2012.

[LW02] N. Linden and W. K. Wootters. The parts determine the whole in a generic pure quantum state. *Phys. Rev. Lett.*, 89:277906, Dec 2002.

[LW18] Shuo-Hui Li and Lei Wang. Neural network renormalization group. *Phys. Rev. Lett.*, 121:260601, Dec 2018.

[LYL$^+$19] Zhihuang Luo, Yi-Zhuang You, Jun Li, Chao-Ming Jian, Dawei Lu, Cenke Xu, Bei Zeng, and Raymond Laflamme. Quantum simulation of the non-fermi-liquid state of sachdev-ye-kitaev model. *npj Quantum Information*, 5(1):53, 2019.

[Mal99] J. Maldacena. The Large-N Limit of Superconformal Field Theories and Supergravity. *International Journal of Theoretical Physics*, 38:1113–1133, 1999.

[McC69] Barry M. McCoy. Theory of a two-dimensional ising model with random impurities. iii. boundary effects. *Phys. Rev.*, 188:1014–1031, Dec 1969.

[MCD+21] C. Monroe, W. C. Campbell, L.-M. Duan, Z.-X. Gong, A. V. Gorshkov, P. W. Hess, R. Islam, K. Kim, N. M. Linke, G. Pagano, P. Richerme, C. Senko, and N. Y. Yao. Programmable quantum simulations of spin systems with trapped ions. *Rev. Mod. Phys.*, 93:025001, Apr 2021.

[MDH79] Shang-keng Ma, Chandan Dasgupta, and Chin-kun Hu. Random antiferromagnetic chain. *Phys. Rev. Lett.*, 43:1434–1437, Nov 1979.

[MDH01] Olexei Motrunich, Kedar Damle, and David A. Huse. Griffiths effects and quantum critical points in dirty superconductors without spin-rotation invariance: One-dimensional examples. *Phys. Rev. B*, 63:224204, May 2001.

[MDIR17] Arianna Montorsi, Fabrizio Dolcini, Rita C. Iotti, and Fausto Rossi. Symmetry-protected topological phases of one-dimensional interacting fermions with spin-charge separation. *Phys. Rev. B*, 95:245108, Jun 2017.

[Mer90] N. David Mermin. Quantum mysteries revisited. *American Journal of Physics*, 58(8):731–734, 1990.

[Mer94] N. David Mermin. Quantum mysteries refined. *American Journal of Physics*, 62(10):880–887, 1994.

[MHK20] Sanjay Moudgalya, David A. Huse, and Vedika Khemani. Perturbative instability towards delocalization at phase transitions between MBL phases. *arXiv e-prints*, page arXiv:2008.09113, August 2020.

[MMHF00] Olexei Motrunich, Siun-Chuon Mau, David A. Huse, and Daniel S. Fisher. Infinite-randomness quantum Ising critical fixed points. , 61(2):1160–1172, January 2000.

[MNZ04] Shuichi Murakami, Naoto Nagaosa, and Shou-Cheng Zhang. Spin-hall insulator. *Phys. Rev. Lett.*, 93:156804, Oct 2004.

[MPM17] Simon Milz, Felix A. Pollock, and Kavan Modi. An introduction to operational quantum dynamics. *Open Systems Information Dynamics*, 24(04):1740016, Dec 2017.

[MS14] P. Mehta and D. J. Schwab. An exact mapping between the Variational Renormalization Group and Deep Learning. *ArXiv e-prints*, October 2014.

[MSE+21] Cole Miles, Rhine Samajdar, Sepehr Ebadi, Tout T. Wang, Hannes Pichler, Subir Sachdev, Mikhail D. Lukin, Markus Greiner, Kilian Q. Weinberger, and Eun-Ah Kim. Machine learning discovery of new phases in programmable quantum simulator snapshots, 2021.

[NC10] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, 2010.

[NFJ$^+$20] Marcel Neugebauer, Laurin Fischer, Alexander Jäger, Stefanie Czischek, Selim Jochim, Matthias Weidemüller, and Martin Gärttner. Neural-network quantum state tomography in a two-qubit experiment. , 102(4):042604, October 2020.

[NOT18] Y. Nagai, M. Okumura, and A. Tanaka. Self-learning Monte Carlo method with Behler-Parrinello neural networks. *ArXiv e-prints*, July 2018.

[NRT12] M. Nozaki, S. Ryu, and T. Takayanagi. Holographic geometry of entanglement renormalization in quantum field theories. *Journal of High Energy Physics*, 10:193, October 2012.

[NRVH17] Adam Nahum, Jonathan Ruhman, Sagar Vijay, and Jeongwan Haah. Quantum entanglement growth under random unitary dynamics. *Phys. Rev. X*, 7:031016, Jul 2017.

[NSQ$^+$17] Y. Nagai, H. Shen, Y. Qi, J. Liu, and L. Fu. Self-learning Monte Carlo method: Continuous-time algorithm. , 96(16):161102, October 2017.

[NVH18] Adam Nahum, Sagar Vijay, and Jeongwan Haah. Operator spreading in random unitary circuits. *Phys. Rev. X*, 8:021014, Apr 2018.

[OC16] Gerardo Ortiz and Emilio Cobanera. What is a particle-conserving topological superfluid? the fate of majorana modes beyond mean-field theory. *Annals of Physics*, 372:357 – 374, 2016.

[ODC$^+$14] Gerardo Ortiz, Jorge Dukelsky, Emilio Cobanera, Carlos Esebbag, and Carlo Beenakker. Many-body characterization of particle-conserving topological superfluids. *Phys. Rev. Lett.*, 113:267002, Dec 2014.

[Ols95] Peter Olsson. Monte carlo analysis of the two-dimensional xy model. ii. comparison with the kosterlitz renormalization-group equations. *Phys. Rev. B*, 52:4526–4535, Aug 1995.

[ONE13] M. Ohliger, V. Nesme, and J. Eisert. Efficient and feasible state tomography of quantum many-body systems. *New Journal of Physics*, 15(1):015024, January 2013.

[Oru14] Roman Orus. A practical introduction to tensor networks: Matrix product states and projected entangled pair states. *Annals of Physics*, 349:117–158, Oct 2014.

[OT17] D. Oprisa and P. Toth. Criticality & Deep Learning II: Momentum Renormalisation Group. *ArXiv e-prints*, May 2017.

[OW16] Ryan O'Donnell and John Wright. Efficient quantum tomography. In *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '16, pages 899–912, New York, NY, USA, 2016. Association for Computing Machinery.

[PBTO12] Frank Pollmann, Erez Berg, Ari M. Turner, and Masaki Oshikawa. Symmetry protection of topological phases in one-dimensional quantum spin systems. *Phys. Rev. B*, 85:075125, Feb 2012.

[PGM$^+$19] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.

[PMV16] Andrew C. Potter, Takahiro Morimoto, and Ashvin Vishwanath. Classification of Interacting Topological Floquet Phases in One Dimension. *Physical Review X*, 6(4):041001, Oct 2016.

[PNGY21] Taylor L. Patti, Khadijeh Najafi, Xun Gao, and Susanne F. Yelin. Entanglement devised barren plateau mitigation. *Physical Review Research*, 3(3), Jul 2021.

[PR04] Matteo Paris and Jaroslav Rehacek. *Quantum state estimation*, volume 649. Springer Science & Business Media, 2004.

[PRA$^+$14] David Pekker, Gil Refael, Ehud Altman, Eugene Demler, and Vadim Oganesyan. Hilbert-Glass Transition: New Universality of Temperature-Tuned Many-Body Dynamical Quantum Criticality. *Phys. Rev. X*, 4(1):11052, 2014.

[Pre18a] John Preskill. Lecture notes for ph219/cs219, 2018.

[Pre18b] John Preskill. Quantum Computing in the NISQ era and beyond. *arXiv e-prints*, page arXiv:1801.00862, January 2018.

[PSV18] Daniel E. Parker, Thomas Scaffidi, and Romain Vasseur. Topological Luttinger liquids from decorated domain walls. , 97(16):165114, April 2018.

[PTBO10] Frank Pollmann, Ari M. Turner, Erez Berg, and Masaki Oshikawa. Entanglement spectrum of a topological phase in one dimension. *Phys. Rev. B*, 81:064439, Feb 2010.

[PVS19] Daniel E. Parker, Romain Vasseur, and Thomas Scaffidi. Topologically protected long edge coherence times in symmetry-broken phases. *Phys. Rev. Lett.*, 122:240605, Jun 2019.

[PYHP15a] Fernando Pastawski, Beni Yoshida, Daniel Harlow, and John Preskill. Holographic quantum error-correcting codes: toy models for the bulk/boundary correspondence. *Journal of High Energy Physics*, 2015:149, Jun 2015.

[PYHP15b] Fernando Pastawski, Beni Yoshida, Daniel Harlow, and John Preskill. Holographic quantum error-correcting codes: toy models for the bulk/boundary correspondence. *Journal of High Energy Physics*, 2015:149, Jun 2015.

[QFK18] Yihui Quek, Stanislav Fort, and Hui Khoon Ng. Adaptive Quantum State Tomography with Neural Networks. *arXiv e-prints*, page arXiv:1812.06693, December 2018.

[Qi13] X.-L. Qi. Exact holographic mapping and emergent space-time geometry. *ArXiv e-prints*, September 2013.

[QYY17] Xiao-Liang Qi, Zhao Yang, and Yi-Zhuang You. Holographic coherent states from random tensor networks. *Journal of High Energy Physics*, 2017(8):60, Aug 2017.

[RA17] Jonathan Ruhman and Ehud Altman. Topological degeneracy and pairing in a one-dimensional gas of spinless fermions. *Phys. Rev. B*, 96:085133, Aug 2017.

[RAS$^+$19] Andrea Rocchetto, Scott Aaronson, Simone Severini, Gonzalo Carvacho, Davide Poderini, Iris Agresti, Marco Bentivegna, and Fabio Sciarrino. Experimental learning of quantum states. *Science Advances*, 5(3):eaau1946, 2019.

[RB01] Robert Raussendorf and Hans J. Briegel. A one-way quantum computer. *Phys. Rev. Lett.*, 86:5188–5191, May 2001.

[RBA15] Jonathan Ruhman, Erez Berg, and Ehud Altman. Topological states in a one-dimensional fermi gas with attractive interaction. *Phys. Rev. Lett.*, 114:100401, Mar 2015.

[RDS$^+$15] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. Imagenet large scale visual recognition challenge, 2015.

[RDTHA12] J. Ruhman, E. G. Dalla Torre, S. D. Huber, and E. Altman. Nonlocal order in elongated dipolar gases. *Phys. Rev. B*, 85:125121, Mar 2012.

[RLR$^+$04] C. F. Roos, G. P. T. Lancaster, M. Riebe, H. Häffner, W. Hänsel, S. Gulde, C. Becher, J. Eschner, F. Schmidt-Kaler, and R. Blatt. Bell states of atoms with ultralong lifetimes and their tomographic state analysis. *Phys. Rev. Lett.*, 92:220402, Jun 2004.

[RM04] G. Refael and J. E. Moore. Entanglement entropy of random quantum critical points in one dimension. *Phys. Rev. Lett.*, 93:260602, Dec 2004.

[RM07] G. Refael and J. E. Moore. Entanglement entropy of the random $s = 1$ heisenberg chain. *Phys. Rev. B*, 76:024419, Jul 2007.

[RM09] G. Refael and J. E. Moore. Criticality and entanglement in random quantum systems. *Journal of Physics A Mathematical General*, 42(50):504010, December 2009.

[Roc18] Andrea Rocchetto. Stabiliser states are efficiently pac-learnable, 2018.

[RS18] J. F. Rodriguez-Nieva and M. S. Scheurer. Identifying topological order via unsupervised machine learning. *ArXiv e-prints*, May 2018.

[RY17] Daniel A. Roberts and Beni Yoshida. Chaos and complexity by design. *Journal of High Energy Physics*, 2017(4):121, Apr 2017.

[SAF$^+$11a] Pietro Smacchia, Luigi Amico, Paolo Facchi, Rosario Fazio, Giuseppe Florio, Saverio Pascazio, and Vlatko Vedral. Statistical mechanics of the cluster ising model. *Phys. Rev. A*, 84:022304, Aug 2011.

[SAF$^+$11b] W. Son, L. Amico, R. Fazio, A. Hamma, S. Pascazio, and V. Vedral. Quantum phase transition between cluster and antiferromagnetic states. *EPL (Europhysics Letters)*, 95(5):50001, 2011.

[Saf16a] M Saffman. Quantum computing with atomic qubits and rydberg interactions: progress and challenges. 49(20):202001, oct 2016.

[Saf16b] M Saffman. Quantum computing with atomic qubits and rydberg interactions: progress and challenges. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 49(20):202001, oct 2016.

[SBYX17] Kevin Slagle, Zhen Bi, Yi-Zhuang You, and Cenke Xu. Out-of-time-order correlation in marginal many-body localized systems. , 95(16):165136, April 2017.

[Sch35] E. Schrödinger. Discussion of probability relations between separated systems. *Mathematical Proceedings of the Cambridge Philosophical Society*, 31(4):555–563, 1935.

[Sch36] E. Schrödinger. Probability relations between separated systems. *Mathematical Proceedings of the Cambridge Philosophical Society*, 32(3):446–452, 1936.

[Sco08] A J Scott. Optimizing quantum process tomography with unitary2-designs. *Journal of Physics A: Mathematical and Theoretical*, 41(5):055308, jan 2008.

[Sen15] T. Senthil. Symmetry-protected topological phases of quantum matter. *Annual Review of Condensed Matter Physics*, 6(1):299–324, 2015.

[SHFDS11] Jay D. Sau, B. I. Halperin, K. Flensberg, and S. Das Sarma. Number conserving theory for topologically protected degeneracy in one-dimensional fermions. *Phys. Rev. B*, 84:144509, Oct 2011.

[SHP$^+$21] Rhine Samajdar, Wen Wei Ho, Hannes Pichler, Mikhail D. Lukin, and Subir Sachdev. Quantum phases of rydberg atoms on a kagome lattice. *Proceedings of the National Academy of Sciences*, 118(4):e2015785118, Jan 2021.

[SK68] E. Specker Simon Kochen. The problem of hidden variables in quantum mechanics. *Indiana Univ. Math. J.*, 17:59–87, 1968.

[Ske02] Kostas Skenderis. Lecture notes on holographic renormalization. *Classical and Quantum Gravity*, 19(22):5849–5876, nov 2002.

[SLK$^+$21] G. Semeghini, H. Levine, A. Keesling, S. Ebadi, T. T. Wang, D. Bluvstein, R. Verresen, H. Pichler, M. Kalinowski, R. Samajdar, and et al. Probing topological spin liquids on a programmable quantum simulator. *Science*, 374(6572):1242–1247, Dec 2021.

[SMM$^+$22] Stefan H. Sack, Raimel A. Medina, Alexios A. Michailidis, Richard Kueng, and Maksym Serbyn. Avoiding barren plateaus using classical shadows, 2022.

[SMY$^+$20] Rahul Sahay, Francisco Machado, Bingtian Ye, Chris R. Laumann, and Norman Y. Yao. Emergent ergodicity at the transition between many-body localized phases. *arXiv e-prints*, page arXiv:2008.08585, August 2020.

[SPA13] M. Serbyn, Z. Papić, and D. A. Abanin. Local Conservation Laws and the Structure of the Many-Body Localized States. *Physical Review Letters*, 111(12):127201, September 2013.

[SPV17] Thomas Scaffidi, Daniel E. Parker, and Romain Vasseur. Gapless Symmetry-Protected Topological Order. *Physical Review X*, 7(4):041048, October 2017.

[SQC$^+$21] Armands Strikis, Dayue Qin, Yanzhu Chen, Simon C. Benjamin, and Ying Li. Learning-based quantum error mitigation. *PRX Quantum*, 2:040330, Nov 2021.

[SRFL09] Andreas P. Schnyder, Shinsei Ryu, Akira Furusaki, and Andreas W. W. Ludwig. Classification of topological insulators and superconductors. *AIP Conference Proceedings*, 1134(1):10–21, 2009.

[SRN19] Brian Skinner, Jonathan Ruhman, and Adam Nahum. Measurement-Induced Phase Transitions in the Dynamics of Entanglement. *Physical Review X*, 9(3):031009, Jul 2019.

[SRS19] Marcin Szyniszewski, Alessandro Romito, and Henning Schomerus. Entanglement transition from variable-strength weak measurements. *arXiv e-prints*, page arXiv:1903.05452, Mar 2019.

[SSW$^+$20] Pascal Scholl, Michael Schuler, Hannah J. Williams, Alexander A. Eberharter, Daniel Barredo, Kai-Niklas Schymik, Vincent Lienhard, Louis-Paul Henry, Thomas C. Lang, Thierry Lahaye, Andreas M. Läuchli, and Antoine Browaeys. Programmable quantum simulation of 2D antiferromagnets with hundreds of Rydberg atoms. *arXiv e-prints*, page arXiv:2012.12268, December 2020.

[sup] *See Supplemental Material for details about SBRG, Ising* Floquet criticality, and additional numerical results.*

[Suz71] Masuo Suzuki. Relationship among exactly soluble models of critical phenomena. i*)2d ising model, dimer problem and the generalized xy-model. *Progress of Theoretical Physics*, 46(5):1337, 1971.

[SV88] M. M. Salomaa and G. E. Volovik. Cosmiclike domain walls in superfluid ^3B: Instantons and diabolical points in (k,r) space. *Phys. Rev. B*, 37:9298–9311, Jun 1988.

[Swi12a] B. Swingle. Constructing holographic spacetimes using entanglement renormalization. *ArXiv e-prints*, September 2012.

[Swi12b] B. Swingle. Entanglement renormalization and holography. , 86(6):065007, September 2012.

[Swi13] B. Swingle. A simple model of many-body localization. *ArXiv e-prints*, July 2013.

[SYX16] Kevin Slagle, Yi-Zhuang You, and Cenke Xu. Disordered XYZ spin chain simulations using the spectrum bifurcation renormalization group. , 94(1):014205, July 2016.

[SZK$^+$21] G.I. Struchalin, Ya. A. Zagorovskii, E.V. Kovlakov, S.S. Straupe, and S.P. Kulik. Experimental estimation of quantum state properties from classical shadows. *PRX Quantum*, 2:010307, Jan 2021.

[SZYZ20] Huitao Shen, Pengfei Zhang, Yi-Zhuang You, and Hui Zhai. Information scrambling in quantum neural networks. *Physical Review Letters*, 124(20), May 2020.

[TM16] G. Torlai and R. G. Melko. Learning thermodynamics with Boltzmann machines. , 94(16):165134, October 2016.

[TM18] Giacomo Torlai and Roger G. Melko. Latent Space Purification via Neural Density Operators. , 120(24):240503, June 2018.

[TMC$^+$18a] G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. Melko, and G. Carleo. Neural-network quantum state tomography. *Nature Physics*, 14:447–450, May 2018.

[TMC$^+$18b] Giacomo Torlai, Guglielmo Mazzola, Juan Carrasquilla, Matthias Troyer, Roger Melko, and Giuseppe Carleo. Neural-network quantum state tomography. *Nature Physics*, 14(5):447–450, 2018.

[TPB11] Ari M. Turner, Frank Pollmann, and Erez Berg. Topological phases of one-dimensional fermions: An entanglement point of view. *Phys. Rev. B*, 83:075102, Feb 2011.

[Tsv11] A. M. Tsvelik. Zero energy Majorana modes in superconducting wires. *arXiv e-prints*, page arXiv:1106.2996, June 2011.

[TT17] A. Tanaka and A. Tomiya. Towards reduction of autocorrelation in HMC by machine learning. *ArXiv e-prints*, December 2017.

[TTv$^+$19] Giacomo Torlai, Brian Timar, Evert P. L. van Nieuwenburg, Harry Levine, Ahmed Omran, Alexander Keesling, Hannes Bernien, Markus Greiner, Vladan Vuletić, Mikhail D. Lukin, Roger G. Melko, and Manuel Endres. Integrating Neural Networks with a Quantum Simulator for State Reconstruction. , 123(23):230504, December 2019.

[VA13] Ronen Vosk and Ehud Altman. Many-body localization in one dimension as a dynamical renormalization group fixed point. *Phys. Rev. Lett.*, 110:067204, Feb 2013.

[VA14] Ronen Vosk and Ehud Altman. Dynamical quantum phase transitions in random spin chains. *Phys. Rev. Lett.*, 112:217204, May 2014.

[VBA85] Pankratov O. A. Volkov B. A. Two-dimensional massless electrons in an inverted contact. *JETP Letters*, 42:178, 1985.

[VED$^+$18] B. Vermersch, A. Elben, M. Dalmonte, J. I. Cirac, and P. Zoller. Unitary n-designs via random quenches in atomic hubbard and spin models: Application to the measurement of rényi entropies. *Phys. Rev. A*, 97:023604, Feb 2018.

[Ver20] Ruben Verresen. Topology and edge states survive quantum criticality between topological insulators. *arXiv e-prints*, page arXiv:2003.05453, March 2020.

[VFPP16] Romain Vasseur, Aaron J. Friedman, S. A. Parameswaran, and Andrew C. Potter. Particle-hole symmetry, many-body localization, and topological edge modes. *Phys. Rev. B*, 93:134207, Apr 2016.

[Vid07a] G. Vidal. Classical simulation of infinite-size quantum lattice systems in one spatial dimension. *Phys. Rev. Lett.*, 98:070201, Feb 2007.

[Vid07b] G. Vidal. Entanglement Renormalization. *Physical Review Letters*, 99(22):220405, November 2007.

[Vid07c] G. Vidal. Entanglement Renormalization. *Physical Review Letters*, 99(22):220405, November 2007.

[Vid08a] G. Vidal. Class of quantum many-body states that can be efficiently simulated. *Phys. Rev. Lett.*, 101:110501, Sep 2008.

[Vid08b] G. Vidal. Class of Quantum Many-Body States That Can Be Efficiently Simulated. *Physical Review Letters*, 101(11):110501, September 2008.

[VJP18] Ruben Verresen, Nick G. Jones, and Frank Pollmann. Topology and Edge Modes in Quantum Critical Chains. , 120(5):057001, February 2018.

[vKKS16] C. W. von Keyserlingk, Vedika Khemani, and S. L. Sondhi. Absolute stability and spatiotemporal long-range order in floquet systems. *Phys. Rev. B*, 94:085112, Aug 2016.

[vLH17] E. P. L. van Nieuwenburg, Y.-H. Liu, and S. D. Huber. Learning phase transitions by confusion. *Nature Physics*, 13:435–439, February 2017.

[VMP17] Ruben Verresen, Roderich Moessner, and Frank Pollmann. One-dimensional symmetry protected topological phases and their transitions. *Phys. Rev. B*, 96:165124, Oct 2017.

[VPP15] R Vasseur, A C Potter, and S A Parameswaran. Quantum Criticality of Hot Random Spin Chains. *Phys. Rev. Lett.*, 114(21):217201, 2015.

[VPYL18] Romain Vasseur, Andrew C. Potter, Yi-Zhuang You, and Andreas W. W. Ludwig. Entanglement Transitions from Holographic Random Tensor Networks. *arXiv e-prints*, page arXiv:1807.07082, Jul 2018.

[VR89] K. Vogel and H. Risken. Determination of quasiprobability distributions in terms of probability distributions for the rotated quadrature phase. *Phys. Rev. A*, 40:2847–2849, Sep 1989.

[vS16] C. W. von Keyserlingk and S. L. Sondhi. Phase structure of one-dimensional interacting Floquet systems. II. Symmetry-broken phases. *Phys. Rev. B*, 93(24):245146, Jun 2016.

[VTJP19] Ruben Verresen, Ryan Thorngren, Nick G. Jones, and Frank Pollmann. Gapless topological phases and symmetry-enriched quantum criticality. *arXiv e-prints*, page arXiv:1905.06969, May 2019.

[Wan16] Lei Wang. Discovering phase transitions with unsupervised learning. *Phys. Rev. B*, 94:195105, Nov 2016.

[WAV20] Brayden Ware, Dmitry Abanin, and Romain Vasseur. Perturbative instability of non-ergodic phases in non-Abelian quantum chains. *arXiv e-prints*, page arXiv:2010.10550, October 2020.

[WBC11] Christopher J. Wood, Jacob D. Biamonte, and David G. Cory. Tensor networks and graphical calculus for open quantum systems. *arXiv e-prints*, page arXiv:1111.6950, November 2011.

[WBC15] Christopher J. Wood, Jacob D. Biamonte, and David G. Cory. Tensor networks and graphical calculus for open quantum systems, 2015.

[Wei78] Don Weingarten. Asymptotic behavior of group integrals in the limit of infinite rank. *Journal of Mathematical Physics*, 19(5):999–1001, 1978.

[Whi93] Steven R. White. Density-matrix algorithms for quantum renormalization groups. *Phys. Rev. B*, 48:10345–10356, Oct 1993.

[WHW$^+$17] Jun Wang, Zhao-Yu Han, Song-Bo Wang, Zeyang Li, Liang-Zhu Mu, Heng Fan, and Lei Wang. Scalable Quantum Tomography with Fidelity Estimation. *arXiv e-prints*, page arXiv:1712.03213, December 2017.

[Wil83] Kenneth G. Wilson. The renormalization group and critical phenomena. *Rev. Mod. Phys.*, 55:583–600, Jul 1983.

[Wil92] Ronald J. Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8(3):229–256, 1992.

[Wit98a] E. Witten. Anti-de Sitter Space, Thermal Phase Transition, And Confinement In Gauge Theories. *Adv. Theor. Math. Phys.*, 2:505–532, March 1998.

[Wit98b] Edward Witten. Anti-de Sitter space and holography. *Advances in Theoretical and Mathematical Physics*, 2:253–291, Jan 1998.

[WMD21] David F. Wise, John J.L. Morton, and Siddharth Dhomkar. Using deep learning to understand and mitigate the qubit noise environment. *PRX Quantum*, 2(1), Jan 2021.

[WS21] Chenan Wei and Tigran A. Sedrakyan. Optical lattice platform for the sachdev-ye-kitaev model. *Phys. Rev. A*, 103:013323, Jan 2021.

[WXPH17] Zhiyuan Wang, Youjiang Xu, Han Pu, and Kaden R. A. Hazzard. Number-conserving interacting fermion models with exact topological superconducting ground states. *Phys. Rev. B*, 96:115110, Sep 2017.

[XLC$^+$19] Tao Xin, Sirui Lu, Ningping Cao, Galit Anikeeva, Dawei Lu, Jun Li, Guilu Long, and Bei Zeng. Local-measurement-based quantum state tomography via neural networks. *npj Quantum Information*, 5:109, November 2019.

[XS18] Shenglong Xu and Brian Swingle. Locality, Quantum Fluctuations, and Scrambling. *arXiv e-prints*, page arXiv:1805.05376, May 2018.

[XX18] Qian Xu and Shuqi Xu. Neural network state estimation for full quantum state tomography. *arXiv e-prints*, page arXiv:1811.06654, November 2018.

[YAM20] Daniel J. Yates, Alexander G. Abanov, and Aditi Mitra. Lifetime of almost strong edge-mode operators in one-dimensional, interacting, symmetry protected topological phases. *Phys. Rev. Lett.*, 124:206803, May 2020.

[YEM19] Daniel J. Yates, Fabian H. L. Essler, and Aditi Mitra. Almost strong $(0, \pi)$ edge modes in clean interacting one-dimensional floquet systems. *Phys. Rev. B*, 99:205419, May 2019.

[YG18] Yi-Zhuang You and Yingfei Gu. Entanglement features of random Hamiltonian dynamics. , 98(1):014309, Jul 2018.

[YHQ16] Z. Yang, P. Hayden, and X.-L. Qi. Bidirectional holographic codes and sub-AdS locality. *Journal of High Energy Physics*, 1:175, January 2016.

[YPPV17] N. Y. Yao, A. C. Potter, I.-D. Potirniche, and A. Vishwanath. Discrete time crystals: Rigidity, criticality, and realizations. *Phys. Rev. Lett.*, 118:030401, Jan 2017.

[YQX16] Yi-Zhuang You, Xiao-Liang Qi, and Cenke Xu. Entanglement holographic mapping of many-body localized system by spectrum bifurcation renormalization group. , 93(10):104205, March 2016.

[YYQ18a] Y.-Z. You, Z. Yang, and X.-L. Qi. Machine learning spatial geometry from entanglement features. , 97(4):045153, February 2018.

[YYQ18b] Yi-Zhuang You, Zhao Yang, and Xiao-Liang Qi. Machine learning spatial geometry from entanglement features. *Phys. Rev. B*, 97:045153, Jan 2018.

[ZC19] Tianci Zhou and Xiao Chen. Operator dynamics in a Brownian quantum circuit. , 99(5):052212, May 2019.

[ZEPS18] Kai Zhou, Gergely Endrődi, Long-Gang Pang, and Horst Stöcker. Regressive and generative neural networks for scalar field theory. *arXiv e-prints*, page arXiv:1810.12879, Oct 2018.

[ZEW18] Linfeng Zhang, Weinan E, and Lei Wang. Monge-Amp\‘ere Flow for Generative Modeling. *arXiv e-prints*, page arXiv:1809.10188, Sep 2018.

[ZHK$^+$17] J. Zhang, P. W. Hess, A. Kyprianidis, P. Becker, A. Lee, J. Smith, G. Pagano, I. D. Potirniche, A. C. Potter, A. Vishwanath, N. Y. Yao, and C. Monroe. Observation of a discrete time crystal. *Nature*, 543(7644):217–220, 2017.

[ZL18] Rui-Xing Zhang and Chao-Xing Liu. Crystalline symmetry-protected majorana mode in number-conserving dirac semimetal nanowires. *Phys. Rev. Lett.*, 120:156802, Apr 2018.

[ZLW18] W. Zhang, J. Liu, and T.-C. Wei. Machine learning of phase transitions in the percolation and XY models. *ArXiv e-prints*, April 2018.

[ZN18] Tianci Zhou and Adam Nahum. Emergent statistical mechanics of entanglement in random unitary circuits. *arXiv e-prints*, page arXiv:1804.09737, Apr 2018.

[ZPH$^+$17] J. Zhang, G. Pagano, P. W. Hess, A. Kyprianidis, P. Becker, H. Kaplan, A. V. Gorshkov, Z. X. Gong, and C. Monroe. Observation of a many-body dynamical phase transition with a 53-qubit quantum simulator. *Nature*, 551:601–604, Nov 2017.

[ZRM20] Andrew Zhao, Nicholas C. Rubin, and Akimasa Miyake. Fermionic partial tomography via classical shadows. *arXiv e-prints*, page arXiv:2010.16094, October 2020.

[ZW17] Long Zhang and Fa Wang. Unconventional Surface Critical Behavior Induced by a Quantum Phase Transition from the Two-Dimensional Affleck-Kennedy-Lieb-Tasaki Phase to a Néel-Ordered Phase. , 118(8):087201, February 2017.

[ZZL20] You Zhou, Pei Zeng, and Zhenhuan Liu. Single-Copies Estimation of Entanglement Negativity. , 125(20):200502, November 2020.