

WEAKLY CORRELATED STUDIES OF STRONGLY CORRELATED SYSTEMS

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

Jordan Matthew Venderley

December 2019

© 2019 Jordan Matthew Venderley

ALL RIGHTS RESERVED

WEAKLY CORRELATED STUDIES OF STRONGLY CORRELATED SYSTEMS

Jordan Matthew Venderley, Ph.D.

Cornell University 2019

In this dissertation we employ and develop several different methods for studying strongly correlated electronic physics with the goal of realizing exotic quantum phenomena. In the first half, we use the density matrix renormalization group (DMRG) to probe the superconducting tendencies of a variety of strongly correlated systems. We start by investigating the triangular lattice Hubbard model as motivated by the organic salts, cobaltates, and recent Moiré superlattice materials. Here we find a clear transition from p -wave superconductivity at moderate on-site repulsion to d -wave superconductivity at strong on-site repulsion. Given the unusual tunability that Moiré superlattices offer in controlling the relative interaction strength, U/t , we thus provide a potential route for realizing a transition between d -wave and p -wave superconductivity via interlayer twist angles. We subsequently modify this model to mimic the band structure of the hole-doped transition metal dichalcogenides and establish evidence for spatially modulated superconductivity i.e. pair density wave ordering. Together, these works suggest that the interplay of frustration and moderately repulsive electronic interactions can be used to drive unconventional superconductivity.

The latter half of this work is largely a product of the growing synergy between the physics and machine learning communities. We first use supervised machine learning to identify quantum phase transitions in a disordered, transverse field Ising model with Ising-duality preserving

interactions. Specifically, we use a neural network trained on entanglement spectra to identify non-equilibrium, thermal-MBL phase transitions in this model and show that our method outperforms traditional regression schemes. Our approach has several additional advantages in that it offers a single framework for identifying multiple types of order, can enable a speedy exploration of large phase spaces by providing meaningful information from single disorder configurations, and has the potential to identify previously unknown phases. We go on to use unsupervised machine learning in order to identify quantum phase transitions in large volume, experimental single crystal x-ray diffraction data. In this setting, data analysis is becoming an increasingly prominent bottleneck with advancements in detector capabilities for X-ray and neutron scattering enabling researchers to collect hundreds of GB to several TB of data in the span of a few hours. To address this, we present a novel label-diffused Gaussian mixture model for clustering over temperature dependences of scattering intensities that allows us to readily identify phase transitions. Our algorithm is capable of analyzing hundreds of GBs of data in the span of minutes, offering the tantalizing possibility of real time analysis. Applications to several materials are discussed.

BIOGRAPHICAL SKETCH

Jordan Venderley was born and raised in the midwestern United States, growing up in the suburbs of Fort Wayne, Indiana. He attended Indiana University in Bloomington from 2010-2014 where he obtained a Bachelor of Science in math and a Bachelor of Science in physics. During this time he worked with Professor Srinivasan Iyengar in the theoretical chemistry department doing ab initio calculations on enzyme dynamics. In 2014, he moved to Ithaca, NY in pursuit of a Ph.D. in physics at Cornell University. Here, he joined the group of Professor Eun-Ah Kim to work in theoretical condensed matter physics. When not doing physics, he enjoys frequenting the gym, playing guitar, and wrestling with Cornell's wrestling club.

Dedicated to my parents and to my sister

ACKNOWLEDGEMENTS

First and foremost, I would like to thank my advisor Eun-Ah Kim for her guidance throughout my Ph.D. Her encouragement and support were vital to the successful completion of my degree. She played a large role in shaping me as a researcher, and I'm indebted to her for always suggesting interesting problems to work on throughout my time at Cornell.

I'd also like to thank my former undergraduate advisors and professors at Indiana University: Srinivasan Iyengar, Kevin Pilgrim, and Babak Seradjeh who were fundamental in laying the intellectual foundation that enabled me to get my Ph.D. They were also largely responsible for inspiring me to pursue it in the first place.

I'm grateful to my family, specifically my parents and sister, for their unwavering support during graduate school. I always look forward to spending time together during the holidays and Fort Wayne will always be my home.

I'd like to thank my current and former groupmates: Mike Matty, Aaron Hui, Eli Gerber, Peter Cha, Yi-Ting Hsu, Sam Lederer, and Yi Zhang for discussions of physics as well as the entire Cornell physics department for providing an environment conducive to doing research in condensed matter theory. I'd particularly like to thank Paul Ginsparg, Kilian Weinberger, and Andrew Wilson for our discussions on machine learning and Katja Nowack and Brad Ramshaw for keeping me up to date on progress in the experimental world.

Finally, I'd like to thank all of my friends. Those in Ithaca, specifically Josh Squires, Eddie Lee, Sal Lombardo, Mike Matty, and Eli Gerber, played an integral role in making my time in graduate school an enjoyable experience. I'd also like to thank my friends back home in Indiana: Tyler Weilbaker, Joey

Schenkel, Phil Tarnowski, Jon Schmidt, Marcus Wadell, and Daniel Carmody for keeping in touch and making my trips home more enjoyable.

All of these people have helped shape me into who I am today and for that I am grateful. I wouldn't be here without you.

TABLE OF CONTENTS

Biographical Sketch	iii
Dedication	iv
Acknowledgements	v
Table of Contents	vii
List of Tables	ix
List of Figures	x
1 Introduction	1
2 DMRG Studies of Superconductivity	5
2.1 The Density Matrix Renormalization Group	5
2.1.1 The Algorithm	6
2.1.2 Matrix Product State Formulation	11
2.1.3 Limitations of DMRG	14
2.2 Introduction to Superconductivity	16
2.2.1 Conventional (BCS) Superconductivity	16
2.2.2 Unconventional Superconductivity	19
2.3 Superconductivity on the Triangular Lattice Hubbard Model	21
2.3.1 Introduction	21
2.3.2 Model and Method	24
2.3.3 Results	27
2.3.4 Conclusion	33
2.3.5 Appendix: Phase Structure for Full Lattice	35
2.3.6 Appendix: Truncation Error	36
2.3.7 Appendix: Log-Linear Plots of Pair Amplitude	37
2.3.8 Appendix: Charge Density	37
2.4 Observation of Pair Density Wave in Spin-Valley Locked Systems	38
2.4.1 Introduction	39
2.4.2 Results	42
2.4.3 Discussion	47
2.4.4 Materials and Methods	49
2.4.5 Acknowledgements	52
2.4.6 Supplementary Materials I: Phase Structure Near the Pair-Field	52
2.4.7 Supplementary Materials II: DMRG convergence	54
2.4.8 Supplementary Materials III: Phase Structure for Attractive Interactions	55
2.4.9 Supplementary Materials IV: Triplet Phase Structure for Repulsive Interactions	57
2.4.10 Supplementary Materials V: Effects of Chemical Potential on the PDW	58

3	Machine Learning Many Body Localization	60
3.1	Foray into the World of Machine Learning	60
3.2	Introduction to Neural Networks	62
3.3	Introduction to Many-body Localization	70
3.4	Machine learning out-of-equilibrium phases of matter	72
3.4.1	Introduction	72
3.4.2	Model	75
3.4.3	Neural Network based Approach	77
3.4.4	Summary and Outlook	83
3.4.5	Acknowledgements:	85
3.4.6	Appendix: Neural Network Sensitivity to Unknown Phases	85
3.4.7	Appendix: Finite Size Scaling	87
4	Machine Learning for Single Crystal X-ray Diffraction Data	89
4.1	Bayesian Inference and Generative Models	90
4.2	Gaussian Mixture Model	92
4.3	Single Crystal X-ray Diffraction	94
4.4	ML for Accelerating Analysis of X-ray Data	96
4.4.1	Introduction	96
4.4.2	Preprocessing	98
4.4.3	Gaussian Mixture Model Approach	103
4.4.4	Charge Density wave detection in TiSe_2	105
4.4.5	Label Smoothing	107
4.4.6	Application to $(\text{Ca}_n\text{Sr}_{1-n})_3\text{Rh}_4\text{Sn}_{13}$	110
4.4.7	Conclusions	116
A	Appendices	118
A.1	The Singular Value Decomposition (SVD) and Properties	118
A.2	Derivation of EM algorithm for GMM and general proof of convergence.	119
	Bibliography	122

LIST OF TABLES

LIST OF FIGURES

2.1	DMRG scheme	7
2.2	Left sweep of finite DMRG scheme	11
2.3	MPS state as designated by the solid black line defined on a 2D lattice. The intrinsic 1D nature of the MPS causes nearest neighbor interactions as shown by the arrow to act as effective long-range couplings in the context of the MPS state.	15
2.4	Schematic for BCS pairing around a Fermi pocket between electrons with spin up at momentum \vec{k} and spin down at $-\vec{k}$. . .	17
2.5	Contour plot of the band structure. The Brillouin zone is marked by the dashed line. A typical Fermi surface for our hole-doped calculations is marked by the thick, black line, here $n = 0.65$. The color legend indicates the energy of the band at a given k-point (for $t=1$).	25
2.6	A portion of our lattice with the location of the pair edge-fields marked in red.	26
2.7	(a) SC phase plots for random edge-field, triplet channel, $U/t = 2$ and $n \sim 0.65$. The phase for each $\Delta_{ij}^{triplet}$ is represented by the color of the bond ij . (b) Amplitude of the SC response with p-wave edge-field (black) and d-wave edge-field (red) for $U/t = 2$. Only the parity channel corresponding to the parity of the edge field is shown. Reported amplitudes are normalized by the applied field ($V = 0.1$).	29
2.8	(a) SC phase plots for random edge-field, singlet channel, $U/t = 10$ and $n \sim 0.4$. The phase for each $\Delta_{ij}^{singlet}$ is represented by the color of the bond connecting site i and j . (b) Amplitude of the SC response with p-wave edge-field (black) and d-wave edge-field (red) for $U/t = 10$. Only the parity channel corresponding to the parity of the edge field is shown. Reported amplitudes are normalized by the applied field ($V = 0.1$).	32
2.9	(a) SC phase plots for random edge-field, singlet channel, $U/t = 2$ and $n \sim 0.65$ (b) SC phase plots for random edge-field, triplet channel, $U/t = 2$ and $n \sim 0.65$ (c) SC phase plots for random edge-field, singlet channel, $U/t = 10$ and $n \sim 0.4$ (d) SC phase plots for random edge-field, triplet channel, $U/t = 10$ and $n \sim 0.4$.	35
2.10	Truncation error for $L=18$ simulations with $U/t = 2$ and $U/t = 10$.	36
2.11	Log-lin plot of superconducting amplitudes presented in main text with $L=24$ for (a) $U/t = 2$ and (b) $U/t = 10$	37
2.12	(a) Charge density along sites of central rung for $U=10$ (with d-wave pair field on left edge) and $U=2$ simulations (with p-wave pair field on left edge). (b) Charge density for $U=10$ simulation (same data as in (a) but shown here as points) and the fitted effective model shown as a continuous plot.	38

2.13	Model and Fermi surface. (a) The spin-dependent staggered flux pattern for one spin component with $\pm\Phi$ flux per plaquette. An opposite flux pattern for the other spin component guarantees time-reversal symmetry. The arrows indicate the direction of positive phase hopping. (b) Our Fermi surface with $t_{i+\hat{x},i;\uparrow} = \frac{2}{\sqrt{3}}e^{i0.3\pi}$ and $\mu = 4.6$ in the tight-binding model in Eqs: (2.14) (2.15). Here the spin-valley locked, circular Fermi pockets are evident.	41
2.14	Evidence of PDW oscillations. (a) $\text{Arg}(\Delta_{\langle i,j \rangle}^{\text{singlet}})$ for all nearest-neighbors with $U = +2$ for our 3×36 lattice with periodic boundary conditions along the short direction and open boundary conditions along the long direction. For visibility, we truncate the plot so that only the third farthest from the edge field is shown. The line thickness is proportional to the pairing amplitude. (b) We plot the real and imaginary components of $\Delta_{ij}^{\text{singlet}}$ and $\Delta_{ij}^{\text{triplet}}$ for i,j along the middle rung of our lattice in order to present the phase oscillations.	44
2.15	Fourier decomposition of PDW and bond charge order. (a) Fourier transforms of the PDW and charge bond order. Zero momentum i.e. constant contributions and decay effects have been removed. (b) Depiction of pairing in momentum space. The regions demarcated by dashed lines are the approximate pairing regions.	48
2.16	Lattice and edge field. A depiction of our lattice. It is periodic in the short direction with 3 unit cells and has open boundaries in the long direction. The ellipses on the right signify that multiple lengths are studied: $L = 12, 18, 24, 36$. The edge field, shown as red lines, is a pair-field of the form given in equation 2.17. The nearest-neighbor hopping structure for spin up is also shown with the spin down hopping structure being the complex conjugate of that shown above.	50
2.17	Singlet phase structure near the edge-field. The singlet phase structure, $\text{Arg}(\Delta_{\langle i,j \rangle}^{\text{singlet}})$, for (a) attractive interactions, $U = -2$, and (b) repulsive interactions, $U = +2$, with random pair-edge-field and (c) repulsive interactions with uniform edge field. Here the $L = 36$ lattice is shown where now we plot only the third closest to the pair-field. As before, the lattice has periodic boundary conditions along the short direction and open boundary conditions along the long direction. Unlike previous plots of this kind, the line thickness here is constant since we only want to emphasize the phase change and the decay from the edge makes this more difficult to see.	53

2.18	DMRG convergence. The decrease in energy between sweeps for the DMRG calculation in the main text ($L = 36$, $U = +2$) demonstrating convergence.	54
2.19	Phase structure for attractive interactions. (a) $\Delta_{\langle ij \rangle}^{singlet}$, (b) $\text{Arg}(\Delta_{\langle ij \rangle}^{singlet})$ and (c) $\text{Arg}(\Delta_{\langle ij \rangle}^{triplet})$ for $U = -2$ along all directed nearest-neighbor bonds. Note that (a) and (b) presents the same phase data and (a) is given only to help the understanding of the new style of plotting. Here we provide the $L=36$ results, but only show the third of the lattice away from the edge probe. Recall our lattice has periodic boundary conditions along the short direction and open boundary conditions along the long direction. The arrows point in the direction of the phase and are also colored according to the argument. Although redundant, this is done to aid visibility. Note that this method of plotting possesses an additional redundancy in that $\Delta_{\langle ij \rangle}^{singlet} = \Delta_{\langle ji \rangle}^{singlet}$ and $\Delta_{ij}^{triplet} = -\Delta_{ji}^{triplet}$	56
2.20	Amplitude for attractive interactions. The real and imaginary components of $\Delta_{ij}^{singlet}$ and $\Delta_{ij}^{triplet}$ for the attractive Hubbard regime, $U = -2$, of the $L=36$ lattice with uniform edge field. Here i, j lie along the middle rung of our lattice.	57
2.21	Triplet phase structure for repulsive interactions. The triplet superconducting phase structure, $\text{Arg}(\langle c_{i\uparrow}^\dagger c_{j\downarrow}^\dagger + c_{i\downarrow}^\dagger c_{j\uparrow}^\dagger \rangle)$, for all nearest-neighbors with $U = +2$ for our 3×36 lattice with a uniform edge-field, periodic boundary conditions along the short direction, and open boundary conditions along the long direction. For visibility, we truncate the plot so that only the third farthest from the edge field is shown. The arrow points in the direction of the phase and is also colored according to the argument for visibility. Note that this method of plotting possesses a redundancy in that $\Delta_{ij}^{triplet} = -\Delta_{ji}^{triplet}$	58
2.22	Effect of chemical potential on PDW phase structure. The phase of the bond-centered singlet superconducting order parameter, $\Delta_{\langle i, j \rangle}^{singlet}$ for nearest-neighbors on the 3×36 lattice with random pair-edge-field where now a larger chemical potential of $\mu = 6.0$ is used rather than that in the main text, $\mu = 4.6$. Here, the line thickness is proportional to the amplitude. The lattice has periodic boundary conditions along the short direction and open boundary conditions along the long direction.	59
2.23	Effect of chemical potential on PDW dominant Fourier mode. The Fourier transforms of the $\mu = 4.6$ and $\mu = 6.0$ PDW. Zero momentum i.e. constant contributions and decay effects have been removed.	59

3.1	Cartoon of single layer feed forward neural network.	63
3.2	Activation functions for neural networks.	65
3.3	(a) A depiction of our neural network. (b) The 2-simplex codomain of our neural network outputs. Colored circles represent regions with maximal confidence; the white circles represent regions of maximal confusion. The green point represents an example output with its associated d_{\min} marked with a green line.	73
3.4	(a) The phase diagram with average neural network output plotted as an RGB parameter. Cuts marked by white lines in (a) are shown in (b) $\lambda = 1.0$ (c) $\overline{\log J} - \overline{\log h} = 0.8$, and (d) $\overline{\log J} - \overline{\log h} = -0.8$ The sampling width is 0.1 for each parameter.	80
3.5	(a) Our C -measure for extracting phase boundaries (defined in the main text) and (b) the average standard deviation of the entanglement entropy. The data in each has been normalized by the largest value in the parameter space for meaningful comparison. (c-e) The measures plotted in (a-b) along the cuts marked in white lines: (c) $\lambda = 1.0$, (d) $\overline{\log J} - \overline{\log h} = +0.8$, and (e) $\overline{\log J} - \overline{\log h} = -0.8$	82
3.6	The 2D phase diagram where the fully trained network has been tested on a single disorder realization.	83
3.7	(a) The phase diagram generated by a partially blind neural network trained only on the two MBL phases. The average neural network output plotted as an RGB parameter. (b) The intensity of the normed gradient of the SG-MBL confidence is plotted.	87
3.8	Finite size scaling of the phase diagram plots for (a) $L=8$, (b) $L=10$, and (c) $L=12$. All plots use 100 disorder configurations at all points in phase space.	88
4.1	Schematic for x-ray scattering of plane wave with wavevector k_{in} scattering off point r' with density $\rho(r')$ and hitting a detector at location r with wavevector k_{out}	95
4.2	Histogram (blue) of $\log \overline{I_q(T)}$ for a single unit-cell of TiSe_2 with background fit (orange) and truncation point described in the main text.	101
4.3	(a) All trajectories for a single Brillouin zone (b) post preprocessing trajectories (c) momenta above the threshold for single unit-cell of TiSe_2 have been colored light blue (d) momenta of high-intensity subset not connected to the Bragg locations have been colored orange.	103

4.4	(a) Logarithm of intensity at 193 K, $Q_l = -3.5$, $Q_h, Q_k \in [-2.5, -1.5]$ (b) Logarithm of intensity at 200 K, $Q_l = -3.5$, $Q_h, Q_k \in [-2.5, -1.5]$ where the CDW peak has disappeared (c) corresponding color coded classification of momenta with $Q_l = -3.5$, $Q_h, Q_k \in [-2.5, -1.5]$ by our clustering algorithm where Bragg-related momenta have been colored blue and CDW peaks have been colored red (d) logarithm of intensity plotted as opacity for a single 3D Brillouin zone centered at $(Q_h, Q_k, Q_l) = (-2, -2, -3.5)$ at 150K (e) corresponding color coded classification of momenta for a single 3D BZ centered at $(Q_h, Q_k, Q_l) = (-2, -2, -3.5)$ by our clustering algorithm where Bragg-related momenta have been colored blue and CDW peaks have been colored red	106
4.5	Clustering results for a single unit cell of TiSe_2 . (a) Color coded trajectories with red and black corresponding to different clustering labels. (b) Average trajectories associated with the clustering.	107
4.6	Kernel, $K(k, 0)$, showing the similarity between the origin and momenta in a 2D.	109
4.7	Experimental single crystal diffraction data for $\text{Sr}_3\text{Rh}_4\text{Sn}_{13}$ with $Q_h \in [-15, 15]$, $Q_k \in [-15, 15]$, $Q_l = 0$ at (a) $T = 300$ K and (b) $T = 30$ K. (c) Mean cluster trajectories outputted by our clustering algorithm. (d) Color-coded classification of the same slice of k-space shown in (a) and (b) showing the low-variance subset (light blue), Bragg peaks (blue), and CDW peaks (red).	111
4.8	(a) Experimental single crystal diffraction data trajectories for $\text{Sr}_3\text{Rh}_4\text{Sn}_{13}$ with $Q_h \in [-15, 15]$, $Q_k \in [-15, 15]$, $Q_l = 0$. (b) Rescaled trajectories after our preprocessing scheme.	112
4.9	(a) Experimental single crystal diffraction data for $\text{Sr}_3\text{Rh}_4\text{Sn}_{13}$ with $Q_h \in [-5, 0]$, $Q_k \in [-5, 0]$, $Q_l = 0$ at $T = 30$ K. (b) Momenta in this slice with average intensities large enough to pass our thresholding. (c) Momenta with trajectories having a high variance. (d) Color-coded, label smoothed clustering results showing the low-variance subset (light blue), Bragg peaks (blue), and CDW peaks (red).	113
4.10	(a) Experimental single crystal diffraction data for $\text{Sr}_3\text{Rh}_4\text{Sn}_{13}$ with $Q_h, Q_k, Q_l \in [-5, 0]$ at $T = 30$ K. The logarithm of the intensities is plotted as the opacity (b) Corresponding color-coded classification of momenta showing Bragg peaks (blue) and CDW peaks (red).	114

- 4.11 Clustering results of $\text{Sr}_3\text{Rh}_4\text{Sn}_{13}$ (a) without and (b) with label smoothing for the k-space slice with $Q_h \in [-5, 0]$, $Q_k \in [-5, 0]$, $Q_l = 0$. Average cluster trajectories for each case are shown in the upper left. The momenta has been color coded according to our classification results. White space corresponds to missing data or points below the thresholding cutoff. Light blue points correspond to the low-variance subset not clustered over. Red and dark blue colors are used for the two clusters of our GMM. . 115
- 4.12 (a) Extracted CDW order parameters for $(\text{Ca}_n\text{Sr}_{1-n})_3\text{Rh}_4\text{Sn}_{13}$ with $Q_h \in [-15, 15]$, $Q_k \in [-15, 15]$, $Q_l = 0$ for several dopings. (b) Phase diagram for $(\text{Ca}_n\text{Sr}_{1-n})_3\text{Rh}_4\text{Sn}_{13}$ as a function of doping and critical CDW temperatures. 116

CHAPTER 1

INTRODUCTION

Dirac famously said, “The underlying physical laws necessary for the mathematical theory of a large part of physics and the whole of chemistry are thus completely known, and the difficulty is only that the exact application of these laws leads to equations much too complicated to be soluble. It therefore becomes desirable that approximate practical methods of applying quantum mechanics should be developed, which can lead to an explanation of the main features of complex atomic systems without too much computation.”[31] This sentiment first expressed nearly a century ago still carries a lot of weight today even in spite of the massive increase in computational power since Dirac’s time. Given the recent death of Moore’s law, it is arguably even more relevant now than when it was first stated in 1929 since physicists may no longer make progress by exploiting the exponential growth in available compute.

In the context of condensed matter theory, this difficulty is quickly encountered when studying strongly correlated electronic systems, most notably when searching for the ground state of an interacting many-body system where eigenstates cannot properly be described by an effective single-particle picture. Given an arbitrary Hamiltonian, the determination of its ground state amounts to solving the time-independent Schrodinger equation $H\Psi = E_g\Psi$ where E_g is the lowest energy eigenvalue i.e. the ground state energy. While analytic forms for most physically reasonable Hamiltonians may be readily written down, the exponential growth of their Hilbert space with the number of physical sites precludes a general solution. For example, naively solving for even a simple spin- $\frac{1}{2}$ system with 50-sites requires diagonalizing

a matrix with more than 10^{15} basis vectors (2^{50} assuming no symmetries are exploited). Such a problem when taken at face value is computationally intractable. Although some headway may be made by working in different limits e.g. strong or weak coupling limits and applying perturbation theory, some of the greatest treasures of condensed matter theory, e.g. high- T_C superconductivity, strongly correlated quantum spin-liquids, and others, lie in this middle ground. Consequently, a great deal of effort has been expended by the community in order to access this parameter regime. This has resulted in the development of several different numerical methods such as quantum monte carlo, exact diagonalization, density matrix renormalization group (DMRG), dynamical mean field theory (DMFT), density matrix embedding theory (DMET), density functional theory (DFT), and others. Of these, DMRG has emerged as one of the most powerful methods for studying 1D and quasi-2D strongly correlated systems.

More recently, some distinct yet in many ways spiritually similar problems have emerged in the field of computer science: the problems of big data and model-selection. For example, consider the task of learning the “simple” function $f : \{0, 1\}^n \rightarrow \{0, 1\}$. It is clear that in order to truly “learn” this function, one must know the respective outputs for each of the exponentially many (2^n) inputs. This begs one to ask the question if any reasonable notion of “learning” is even possible when n becomes large. So it’s important to ask “how good” a learned rule performs after seeing a non-exhaustive subset of inputs and how well a given hypothesis generalizes to unseen examples. Machine learning in conjunction with statistical learning theory is precisely the field devoted to addressing this problem i.e. learning models in the context of a bias-variance tradeoff.

In this dissertation we employ a variety of different approaches to tackle the exponential scaling and big data problems encountered in physics. This work is largely a product of the growing synergy between the physics and computer science communities. It is my hope that the research contained in this dissertation has meaningfully contributed to the cross-fertilization of these fields.

In chapter 2 we use the density matrix renormalization group to probe superconductivity in strongly correlated systems. We first introduce the DMRG algorithm in section 2.1 to address the problem of exponential scaling problem of quantum Hilbert spaces and briefly introduce superconductivity in section 2.2. In section 2.3 we use DMRG to probe the dominant superconducting instabilities for the triangular lattice Hubbard model as motivated by the organic salts, cobaltates, and recent moire superlattice materials. Then in section 2.4 we further modify this model to mimic the spin-valley locking of the hole-doped transition metal dichalcogenides and examine its superconducting behavior.

In chapter 3 we use supervised machine learning to identify quantum phase transitions. We start by briefly introducing machine learning in section 3.1 and providing an elementary introduction to neural networks in section 3.2. After introducing the physics of many-body localization (MBL) in section 3.3, we use a neural network to identify non-equilibrium, thermal-MBL phase transitions in section 3.4 and show that our method outperforms traditional schemes.

In chapter 4 we use unsupervised machine learning to identify phase transitions in large volume experimental data. We start by discussing machine learning through the lens of Bayesian inference in section 4.1 and introducing

Gaussian mixture models in section 4.2. After providing a crash course on single crystal x-ray diffraction in section 4.3, we employ a modified Gaussian mixture model for accelerating the analysis of experimental single crystal x-ray diffraction in section 4.4.

CHAPTER 2
DMRG STUDIES OF SUPERCONDUCTIVITY

2.1 The Density Matrix Renormalization Group

The Density Matrix Renormalization Group (DMRG) is a powerful, non-perturbative method for studying strongly interacting systems [103, 143, 135, 133, 10, 63]. It has been used with great success to explore a diverse selection of strongly correlated phenomena highlighted by stripes, spin-liquids, and superconductivity [143, 135, 133, 10, 63, 158, 51, 59, 61, 62]. It was initially developed by Steve White in 1992 [131] as a generalization of Ken Wilson's numerical renormalization group and is one means of alleviating the numerical bottleneck accompanying the exponential scaling of the Hilbert space with the linear increase in physical sites. It has emerged as the gold-standard for numerical studies of 1D and quasi-2D strongly correlated systems at zero temperature especially in the presence of the fermion-sign problem. Although some people question whether it is "truly" a renormalization group approach since there is no notion of a UV or IR cutoff, it may be properly conceptualized as an RG procedure in the sense that it is a systematic process for removing "unimportant" degrees of freedom. Here there is some flow to a fixed point in the space of density matrices. Modern day DMRG is most properly and concisely described as a variational optimization scheme over the class of matrix product states with finite dimension. Alternatively, it is an algorithm for finding the approximate ground state that is "closest" to the true eigenstate of a Hamiltonian where "close" is taken with respect to the Frobenius norm so that if you construct a density matrix from the true ground state and retain only the

largest D singular values, you are left with the DMRG approximation.

However, while DMRG has seen much success over the past decades, it is not without its shortcomings. In particular, it struggles to extend to 2D and calculations in higher dimensions are infeasible. Although higher dimensional, tensor-based generalizations of DMRG have been developed such as the multi-scale entanglement renormalization ansatz (MERA) and projected entangled pair states (PEPS), their large computational cost has prevented their widespread adoption. Despite these limitations, DMRG has become an invaluable tool in the condensed matter theorist's repertoire, especially for studying strongly correlated systems.

We will first discuss the DMRG algorithm as it was initially conceived a.k.a. traditional DMRG and motivate its efficacy. We will subsequently put it in the more modern language of matrix product states and in the next chapter discuss its implementation in two projects studying the superconducting tendencies of the Hubbard model.

2.1.1 The Algorithm

We first describe the traditional infinite DMRG algorithm since it is conceptually easiest to state and its finite counterpart can be readily established afterwards. For the most part we follow the notation used in [103] which is nearly universal in the DMRG literature.

Assume that we have a 1D model Hamiltonian, H , for which we'd like to solve for the ground state with $\lim_{L \rightarrow \infty}$ at $T = 0$. Consider two system "blocks", call

them A and B (WOLOG the left and right blocks respectively), each containing a single physical site with local Hilbert space dimension of size d . These sites could be spins such as in the quantum Ising model so that $d = 2$ or particles with spin such as in the Hubbard model so that $d = 4$. Define the Hamiltonian with terms only contained within each block as H_A and H_B and call the set of basis vectors for these blocks $\{|\phi\rangle_A\}$ and $\{|\phi\rangle_B\}$. At this point, A and B are small enough that the Hamiltonians may be written down exactly but we will assume that there is an upper limit to the number of basis vectors that we are able to keep, call this limit D . Now add one new site each to A and B to build new Hamiltonians $H_{A\bullet}$ and $H_{\bullet B}$. Let the single site basis vectors be written as $\{|\sigma\rangle_A\}$ ($\{|\sigma\rangle_B\}$) so that the basis of $H_{A\bullet}$ ($H_{\bullet B}$) is $\{|\phi\rangle_A|\sigma\rangle_A\}$ ($\{|\sigma\rangle_B|\phi\rangle_B\}$). Now combine these into a “superblock” Hamiltonian $H_{A\bullet\bullet B}$ as in Fig. 2.1.

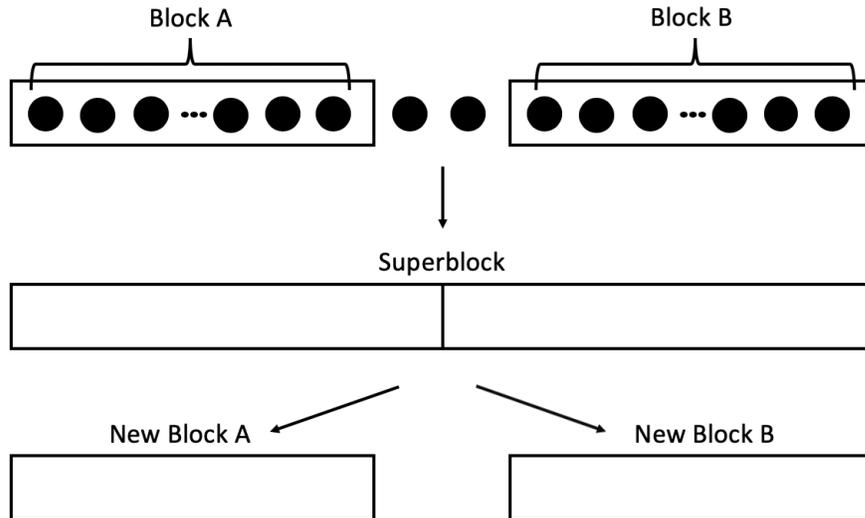


Figure 2.1: DMRG scheme

The idea of DMRG is to continue adding sites but to employ a decimation procedure to forestall the exponential growth of the Hilbert space. This allows us to only keep the “important” degrees of freedom for our system and find an

approximate ground state of our Hamiltonian in the large system size limit. The DMRG algorithm is simply a means of finding the “optimal” effective basis for approximating the ground state under the assumption that we have a limited number of degrees of freedom. Once we have constructed our superblock, we find its ground state, $|\psi\rangle$, with a sparse diagonalization method such as the Lanczos algorithm, Eqn. 2.1. By construction, this state is written in terms of the basis states we have specified for our blocks, see Eqn. 2.2.

$$E_0 = \frac{\langle \psi | H_{A \bullet \bullet B} | \psi \rangle}{\langle \psi | \psi \rangle} \quad (2.1)$$

$$|\psi\rangle = \sum_{\phi_A, \sigma_A, \sigma_B, \phi_B} \psi_{\phi_A \sigma_A \sigma_B \phi_B} |\phi\rangle_A |\sigma\rangle_A |\sigma\rangle_B |\phi\rangle_B = \sum_{i,j} \psi_{ij} |i\rangle_A |j\rangle_B \quad (2.2)$$

Now, we must apply the decimation procedure. This is readily done by using SVD to perform a Schmidt decomposition of the matrix coefficient ψ_{ij} in Eqn. 2.2. Those unfamiliar with SVD may see Appendix A.1 for details. The key to DMRG is to only retain the largest D singular values in this decomposition. The size of D here is set by the max size of the superblock for which the ground state can be efficiently found via sparse diagonalization, typically $D \sim 1000$.

$$\begin{aligned} |\psi\rangle &= \sum_{i,j} \psi_{ij} |i\rangle_A |j\rangle_B = \sum_{i,j,m} U_{im} \Sigma_{mm} (V^\dagger)_{mj} |i\rangle_A |j\rangle_B \\ &= \sum_m \sigma_m \left(\sum_i U_{im} |i\rangle_A \right) \left(\sum_j (V^\dagger)_{mj} |j\rangle_B \right) \\ &= \sum_m \sigma_m |m\rangle_A |m\rangle_B \end{aligned} \quad (2.3)$$

Physically, the largest singular values of our Schmidt decomposition

correspond to those states with the largest contribution to the bipartite von Neumann entanglement entropy. This is clear if we construct the reduced density matrix ρ_A by tracing out the B sub-block degrees of freedom as in Eqn. 2.4.

$$\rho_A = Tr_B |\psi\rangle\langle\psi| = \sum_m \sigma_m^2 |m\rangle\langle m|_A \quad (2.4)$$

We now project each block into its new respective effective basis and return to the first step of the algorithm, i.e. add a new site to each block. This procedure is continued until the DMRG energy converges.

The two block embedding approach of DMRG is successful because one block acts as a “bath” for the other so when we keep those states that have the largest contribution to the entanglement entropy we are keeping those states that are most relevant in the context of a larger system. Intuitively, it makes sense that these are the states one would want to keep when trying to access large system size properties with a limited number of degrees of freedom. Importantly, this truncation scheme is viable (as long as the system is gapped) because the entanglement spectra decays exponentially. Thus the contribution of the states being “thrown out” is relatively small and the DMRG ground state is a very good approximation to the true ground state. An unfortunate side effect of this is that DMRG struggles with gapless systems such as chiral spin liquids since the entanglement spectra decay rate is no longer exponential and the truncation weight associated with the discarded states becomes non-negligible. Of course, this fact hasn’t prevented people from trying to apply DMRG to these types of systems.

Note that the process of efficiently representing our blocks in terms of the newly found effective basis has not been discussed in the above exposition. Ultimately, this is just an exercise in storing enough basis transformations (depending on how long range the hoppings/interactions are) and chaining them together to write down the terms in $H_{A\bullet}$ and $H_{\bullet B}$ that must be introduced upon adding an additional site to the blocks. There is a large swath of literature on performing all of these operations in the most efficient way possible (both with respect to time and space complexity) but as this is not the focus of this thesis, we refer the interested reader to [103, 102] for extensive reviews.

A simple, bare-bones python implementation is provided on my github at https://github.com/jvend/traditional_DMRG/blob/master/DMRG.py where I use infinite DMRG to solve the 1D Ising model with transverse field. This is a nice proof of principle since exact results are known and the ground state energy may be expressed analytically as an elliptic integral of the second kind [90]. As expected, the DMRG energy converges to this exact result.

For a finite variant of traditional DMRG, one must first proceed identically to the infinite algorithm in order to grow the Hamiltonian to the desired size. Then a “sweeping” procedure is employed in which the system is divided into two sub-blocks A and B and the same local DMRG decimation procedure is applied as above but now sub-block A is grown and sub-block B is shrunken so that the total size of the system doesn’t change, see Fig. 2.2. In actuality, the B sub-block isn’t explicitly shrunken, but rather the block of that size from the previous sweep is reused. By iteratively performing these left and right DMRG sweeps through the system, the effective basis for representing our Hamiltonian is improved until convergence is reached.

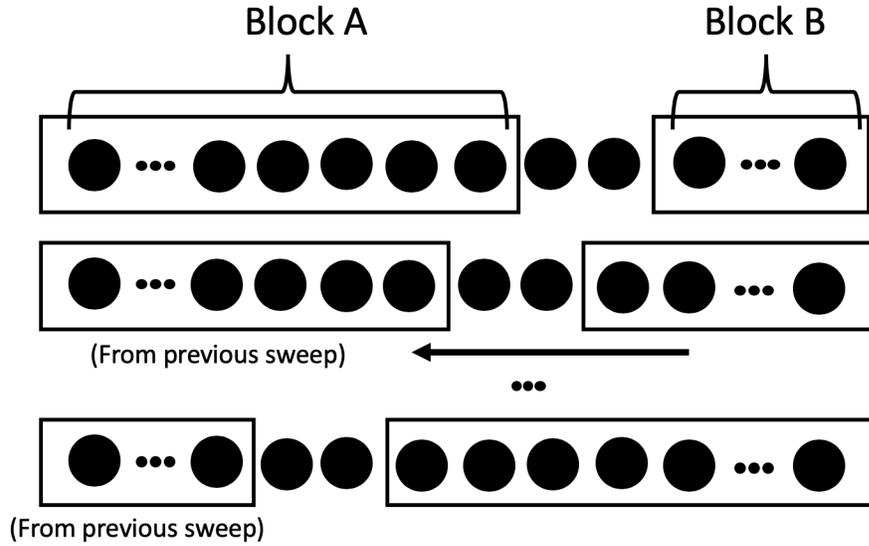


Figure 2.2: Left sweep of finite DMRG scheme

2.1.2 Matrix Product State Formulation

Modern DMRG implementations employ a matrix product state (MPS) scheme that offers several advantages over the traditional approach described above (even though the two schemes are actually equivalent). Probably, the most important consequence of the MPS approach towards DMRG is that it put DMRG on a proper theoretical foundation as a variational optimization scheme. The ansatz for this variational optimization is the the set of finite matrix product states, and the so-called bond-dimension, χ , of these matrix product states controls the representational power of this ansatz. However, there is another practical advantage in the case of finite DMRG. The MPS language allows one to compactly represent initial states and Hamiltonian operators so that it is no longer necessary to initialize the finite DMRG algorithm by first growing from a small system size.

In principle any wavefunction can be written as a MPS. Consider some wavefunction $|\psi\rangle$, defined on a n physical sites. Let the local Hilbert space on the i th site be described by the set of basis vectors $\{|\sigma_i\rangle\}$. Then $|\psi\rangle$ may be written in this basis as in Eqn. 2.5.

$$|\psi\rangle = \psi_{\sigma_1 \dots \sigma_n} |\sigma_1, \dots, \sigma_n\rangle \quad (2.5)$$

The MPS representation of of this state may be built by iteratively performing singular value decompositions on the coefficient $\psi_{\sigma_1 \dots \sigma_n}$. That is to say consider this coefficient as a 2D matrix with first index corresponding to the first local physical site and the second corresponding to all others i.e. $\psi_{\sigma_1, (\sigma_2 \dots \sigma_n)}$. Now perform an SVD on this matrix as in Eqn. 2.6.

$$\begin{aligned} \psi_{\sigma_1, (\sigma_2 \dots \sigma_n)} &= \sum_{a_1} U_{\sigma_1 a_1} \Sigma_{a_1} (V^\dagger)_{a_1, (\sigma_2 \dots \sigma_n)} = \sum_{a_1} A_{a_1}^{\sigma_1} \psi_{a_1 \sigma_2, (\sigma_2, \dots, \sigma_n)} \\ &= \sum_{a_1, a_2} A_{a_1}^{\sigma_1} U_{a_1 \sigma_2, a_2} \Sigma_{a_2} (V^\dagger)_{a_2, (\sigma_3 \dots \sigma_n)} = \sum_{a_1, a_2} A_{a_1}^{\sigma_1} A_{a_1 a_2}^{\sigma_2} \psi_{a_2 \sigma_3, (\sigma_2, \dots, \sigma_n)} \quad (2.6) \\ &= \dots = \sum_{a_1, \dots, a_n} A_{a_1}^{\sigma_1} A_{a_1 a_2}^{\sigma_2} \dots A_{a_{n-1} a_n}^{\sigma_{n-1}} A_{a_n}^{\sigma_n} \end{aligned}$$

After performing the SVD, reshape the matrix U consisting of the left singular vectors into the matrix A^{σ_1} where the elements of A^{σ_1} are $A_{a_1}^{\sigma_1}$. We have thus promoted the left index of U to an upper index corresponding to physical sites (it may now be thought of as an index over matrices) and kept the right index as an ‘‘auxiliary’’ index. The other two parts of the factorization ($\Sigma_{a_1} (V^\dagger)_{a_1, (\sigma_2 \dots \sigma_n)}$) are combined into $\psi_{a_1 \sigma_2, (\sigma_2, \dots, \sigma_n)}$ where we have again shifted the left most physical site into a left index with the new auxiliary index a_1 . Again perform an SVD on $\psi_{a_1 \sigma_2, (\sigma_2, \dots, \sigma_n)}$, and this time only promote the σ_2 part of the

left index of $U_{a_1\sigma_2,a_2}$ to an upper index to build matrices A^{σ_2} . We can continue this procedure to completely rewrite $\psi_{\sigma_1\dots\sigma_n}$ as a product of matrices (hence the name matrix product state): $A^{\sigma_1} \cdots A^{\sigma_n} = \sum_{a_1,\dots,a_n} A_{a_1}^{\sigma_1} A_{a_1 a_2}^{\sigma_2} \cdots A_{a_{n-1} a_n}^{\sigma_{n-1}} A_{a_n}^{\sigma_n}$.

At this point, one should question the utility of the MPS representation because all we've done at this point is rewrite a scalar in terms of a more complicated product of matrices. The advantage of this representation lies in assuming a maximum dimension to the matrices A^{σ_i} . This is the so-called bond dimension, χ , of the matrix product state. By keeping only the χ largest singular values in each SVD step of our MPS decomposition, we are left with a very controlled way of preventing the exponential growth of our Hilbert space. Specifically the number of degrees of freedom is now at most $n\chi^2$ as opposed to d^n where d is the local Hilbert space dimension, n is the number of physical sites, and χ is the MPS bond dimension. Notably, the MPS representation of a wavefunction is linear in the number of sites!

Now that we have explicated the MPS scheme, one can rewrite DMRG in this language. It should be clear given the sweeping nature of DMRG through physical sites that rewriting the wavefunction as a product over local matrices offers a very natural framework for this algorithm. We will not do this here, because although the MPS representation offers the proper formal language for thinking about these states, the conceptual premise of DMRG should already be evident by now. We refer the interested reader to [103] for a fairly exhaustive exposition.

2.1.3 Limitations of DMRG

We discuss in a little more detail the two primary limitations of DMRG: (1) its quasi-1D nature and (2) its applicability to gapped systems. As evidenced by Eqn. 2.6, matrix product states are intrinsically 1D. While one can write down higher dimensional tensor networks such as projected entangled pair states (PEPS) or tensor product states (TPS), the index contractions required by DMRG for sweeping through these higher dimensional objects are too computationally expensive to make them viable. In spite of their 1D nature, MPSs are still used to study two-dimensional systems by “snaking” a 1D MPS through the system as shown in Fig. 2.3. However, this snaking prevents true 2D calculations since it results in development of effective long-range couplings as shown by the arrow in the figure. These long-range couplings translate into requiring larger bond-dimensions for a good approximation and quickly become intractable. Consequently, only quasi-1D simulations are viable with DMRG where one lattice dimension is short with periodic boundary conditions and the other is long with open boundary conditions.

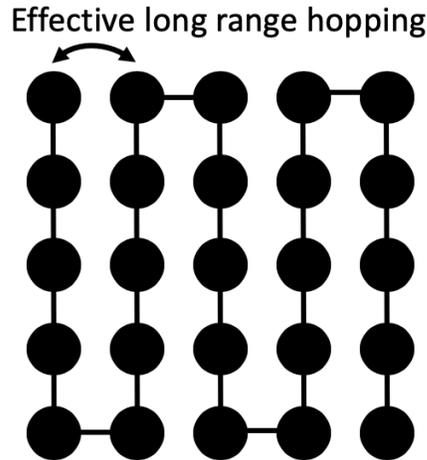


Figure 2.3: MPS state as designated by the solid black line defined on a 2D lattice. The intrinsic 1D nature of the MPS causes nearest neighbor interactions as shown by the arrow to act as effective long-range couplings in the context of the MPS state.

The second main limitation of DMRG is that it is only really suitable for studying gapped systems. Hastings showed using Lieb-Robinson bounds that gapped 1D systems have an entanglement entropy that scales with an area law and thus exhibit exponential decay in their entanglement spectra.[50] Within the MPS description, this exponential decay directly translates into an exponentially small truncation weight and consequently an efficient matrix product state approximation. However, we know from conformal field theory that as criticality is approached and the system becomes gapless, power-law correlations develop so that entanglement entropy diverges logarithmically with system size. Consequently, the entanglement spectra cannot decay exponentially and the truncation weight can become non-negligible. This directly implies difficulty in simulating gapless systems e.g. chiral spin-liquids using DMRG because the effectiveness of the underlying MPS ansatz becomes tenuous. Any DMRG studies at criticality must proceed with caution keeping extremely large bond dimensions and carefully performing finite-size scaling.

2.2 Introduction to Superconductivity

Superconductivity is a macroscopic quantum phenomena in which electron pairs condense into coherent bosonic quasiparticles known as Cooper pairs. These Cooper pairs are the source of several exotic superconducting properties notably dissipationless currents and bulk magnetic field expulsion i.e. Meissner-Ochsenfeld screening. Ultimately, these unusual properties arise because the Cooper pairs obey bosonic quantum statistics as opposed to the fermionic statistics of their constituent electrons. Although electrons naturally repel each other through Coulombic interactions, this electron-electron repulsion can be overcome if a sufficiently strong effective attractive interaction develops between electrons e.g. by electron-phonon interactions in the case of conventional superconductivity. For conventional superconductors, this attractive electron-electron interaction induces a Fermi surface instability and in turn produces a new quantum state of matter, the superconducting state.

2.2.1 Conventional (BCS) Superconductivity

The correct microscopic theory of conventional superconductivity was first proposed by John Bardeen, Leon Cooper, and John Schrieffer in 1957 and is known as BCS theory. It is a weak-coupling mean field theory in which the underlying mean field consists of pairs of annihilation operators with opposite spin and momenta, i.e. $\langle c_{-k\downarrow} c_{k\uparrow} \rangle$. This is shown schematically in Fig. 2.4.

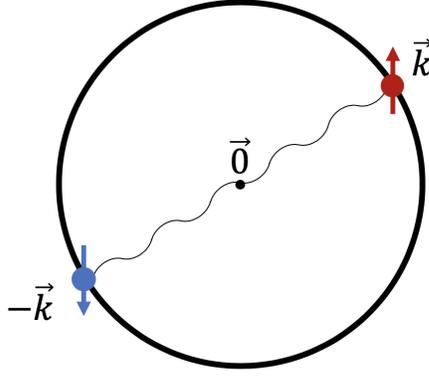


Figure 2.4: Schematic for BCS pairing around a Fermi pocket between electrons with spin up at momentum \vec{k} and spin down at $-\vec{k}$

If we ignore finite momentum scattering processes so that the Cooper pairs have zero net momentum and assume a constant interaction in k-space we may write down the following Hamiltonian:

$$H = \sum_{k,\sigma} \xi_k c_{k\sigma}^\dagger c_{k\sigma} + g \sum_{k,k'} c_{k\uparrow}^\dagger c_{-k\downarrow}^\dagger c_{-k'\downarrow} c_{k'\uparrow} \quad (2.7)$$

A mean field decomposition may be performed by rewriting the above Hamiltonian in terms of $b_k \equiv c_{-k\downarrow} c_{k\uparrow}$ and neglecting terms of $O(\delta b_k^2)$ where $\delta b_k \equiv b_k - \langle b_k \rangle$ to obtain the new Hamiltonian below:

$$H_{MF} = \sum_{k,\sigma} \xi_k c_{k\sigma}^\dagger c_{k\sigma} + g \sum_{k,k'} (\langle b_{k'} \rangle b_k^\dagger + \langle b_k^\dagger \rangle b_{k'} - \langle b_k^\dagger \rangle \langle b_{k'} \rangle) \quad (2.8)$$

This hamiltonian may be readily diagonalized by a Bogoliubov transformation: $c_{k\uparrow} = u_k^* \gamma_{k+} + v_k \gamma_{k-}^\dagger$ and $c_{-k\downarrow} = -v_k^* \gamma_{k+} + u_k \gamma_{k-}^\dagger$ where $|u_k|^2 + |v_k|^2 = 1$ and γ_α are fermions (called bogoliubons) obeying the standard anticommutation relations $\{\gamma_{k\alpha}^\dagger, \gamma_{k'\beta}\} = \delta_{kk'} \delta_{\alpha\beta}$. This results in

$$H_{MF} = \sum_{k,\sigma} E_k \gamma_{k\alpha}^\dagger \gamma_{k\alpha} + \left[\sum_k \xi_k - E_k + \Delta \langle b_k \rangle^* \right] \quad (2.9)$$

where $E_k \equiv \sqrt{\xi_k^2 + \Delta^2}$ and $\Delta \equiv -g \sum_k \langle b_k \rangle$ is the superconducting gap which must be determined self-consistently. Note that our assumption of a structureless interaction translates into a structureless i.e. spherically symmetric or s-wave superconducting gap. The associated ground state, the BCS state, is given in Eqn. 2.10. A complementary approach to that given above is to proceed variationally by treating the state in Eqn. 2.10 as an ansatz and showing that it yields a lower ground state energy than the Fermi gas.

$$|\Phi_{BCS}\rangle = \prod_k (u_k + v_k c_{k\uparrow}^\dagger c_{-k\downarrow}^\dagger) |0\rangle \quad (2.10)$$

The above formulation for s-wave superconductivity may be generalized to describe anisotropic, non-singlet Cooper pairs by replacing the constant interaction term g above with a momentum and spin dependent coefficient $V_{\sigma_3\sigma_4}^{\sigma_1\sigma_2}(k, k')$ where the σ are the spins for the operators in the interaction term. This so-called generalized BCS theory is typically agnostic about the underlying mechanism and proceeds by assuming some phenomenological form for the interaction.

Key to this generalized formulation is the idea of classification by symmetry and irreducible representations (irreps). When the electrons live on a lattice, the superconducting gap must transform according to the irreducible representations of the symmetry group of the underlying lattice. Importantly this allows one to make statements about what types of superconductivity are in principle allowed to occur for different lattices. This is formalized by

Schur's lemma of representation theory. Since certain representations such as $p_x + ip_y$, chiral p-wave superconductivity can host Majorana quasiparticles that can be used as a medium for topological quantum computations, understanding what irreps can occur in what contexts is highly important. Even though we must ultimately abandon this weak-coupling generalized BCS picture, this notion of classification by symmetry will be indispensable in our simulations of strongly correlated systems with our primary goal being to determine the symmetry of the dominant superconducting instability.

2.2.2 Unconventional Superconductivity

Even with the generalizations discussed above, BCS theory is insufficient for a proper description of the new generation of strongly correlated superconductors e.g. the cuprates and more recently Moire superlattices. Since BCS is intrinsically a weak-coupling mean-field theory, it is not applicable for describing the behavior of strongly correlated SCs. Specifically, while BCS superconductivity is an instability of the Fermi surface, the parent state for superconductivity in the cuprates and Moire superlattices is a Mott insulator (as opposed to a Fermi liquid). This Mott insulating state differs from a Fermi liquid in that it possesses strong correlations and lacks a Fermi surface (it's an insulator). Hole doping this Mott insulator induces the superconducting state. The non-BCS character of the cuprates is reinforced by the short correlation length of their Cooper pairs as opposed to the large coherence lengths characteristic of the BCS superconductors. This short coherence length translates into large superconducting fluctuations so that the $O(\delta b_k^2)$ terms of our mean-field theory cannot be neglected and the mean-field picture breaks

down. The importance of these fluctuations is further enhanced by the quasi-two-dimensionality of the band-structure. One of the best ways to gain insight into these strongly correlated physics is to proceed numerically via DMRG. We do this in the next two sections.

2.3 Superconductivity on the Triangular Lattice Hubbard Model

This section is adapted from a PRB Rapid Communications paper [122] with Eun-Ah Kim.

With the discovery of strong coupling physics and superconductivity in Moiré superlattices, it's essential to have an understanding of strong coupling driven superconductivity in systems with trigonal symmetry. The simplest lattice model with trigonal symmetry is the triangular lattice Hubbard model. Although the triangular lattice spin model is a heavily studied model in the context of frustration, studies of the hole-doped triangular lattice Hubbard model are rare. Here we use density matrix renormalization group (DMRG) to investigate the dominant superconducting channels in the hole-doped triangular lattice Hubbard model over a range of repulsive interaction strengths. We find a clear transition from p -wave superconductivity at moderate on-site repulsion strength ($U/t = 2$) at filling above $1/4$ ($n \sim 0.65$) to d -wave superconductivity at strong on-site repulsion strength ($U/t = 10$) at filling below $1/4$ ($n \sim 0.4$). The unusual tunability that Moiré superlattices offer in controlling U/t would open up the opportunity to realize this transition between d -wave and p -wave superconductivity.

2.3.1 Introduction

The discovery of superconductivity in van der Waals materials [15] with an (effective) triangular lattice structure has expanded the scope of correlation

physics in triangular lattice models. Nevertheless, recent literature motivated by the experiments[15, 14, 145] have for the most part been limited to mean-field theories[140, 32, 38] and perturbative renormalization group studies[57, 98]. Such approximations are unavoidable when one tries to capture the full band structure of Moiré superlattices such as magic angle twisted bilayer graphene, which require several bands and large unit-cells to accommodate the characteristic Moiré patterns that ultimately drive their correlated behavior.

However, there is increasing awareness that the effective band structure of a wider class of Moiré superlattice systems is much simpler[136, 137]. In particular, Ref. [136] showed a simple triangular lattice Hubbard model adequately captures hetero-transition metal dichalcogenide (TMD) bilayers. The observation of superconductivity and other strong coupling phenomena in these systems[70] motivates studying simple effective models with strong coupling approaches. In fact, Guo *et al.* [48] have already studied a three-band triangular lattice Hubbard model with modulated hopping using determinant quantum Monte Carlo in the context of Moiré superlattices. Similarly, Zheng *et al.* [157], studied a 2-orbital model on the honeycomb lattice with density matrix renormalization group (DMRG) in the insulating state. In this work, we study strong coupling driven superconductivity and investigate how the strength of the repulsive interaction affects the pairing symmetry in the triangular lattice Hubbard model.

Even before the discovery of superconductivity in twisted bilayer graphene, superconductivity in the organic salts [141, 142] and cobaltates[116] had generated interest in strong interaction driven superconductivity in triangular lattice systems. Here, enhancement of superconducting pair correlations by

Hubbard repulsion, has been found in triangular lattices at quarter-filling, thought to be appropriate for organics by some authors[45]. Gutzwiller projected t-J models with smaller Hilbert space predicted $d + id$ pairing for light to moderate hole-dopings away from half-filling[129, 127, 146]. For the Hubbard model, Chen *et al.* [21] used dynamical cluster quantum Monte Carlo to study strong repulsion with moderate hole-doping to find the dominant pairing susceptibility in the $d + id$ channel. These 2D irreps are of theoretical interest because of the non-zero Chern number associated with their quasi-particle spectra and the corresponding ability to support Majorana bound states. In particular, the Majorana's are expected to be better protected in the $(p \pm ip)$ paired state.

Here we use DMRG to study superconducting tendencies driven by moderate and strong repulsive interactions. DMRG has been used with great success to explore a diverse selection of strongly correlated phenomena highlighted by stripes, spin-liquids, and superconductivity [103, 143, 135, 133, 10, 63, 158, 51, 59, 61, 62, 60]. However, since DMRG is quasi-1D in nature, no true long-range order can be seen in the correlations. Thus in order to access our system's superconducting tendencies we implement a pair-edge-field motivated by the field-pinning approach first introduced in [134] and employed similarly in [121]. By studying how the system responds to the pair-edge field comparing responses in different superconducting channels, one can assess the model's propensity for superconducting instabilities in different channels.

2.3.2 Model and Method

The full model hamiltonian is $H_{\text{tot}} = H + H_{\text{edge}}$, where H is the standard Hubbard model on a triangular lattice:

$$H = -t \sum_{\langle i,j \rangle \sigma} c_{i\sigma}^\dagger c_{j\sigma} - \mu \sum_{i\sigma} n_{i\sigma} + U \sum_i n_{i\uparrow} n_{i\downarrow}, \quad (2.11)$$

and H_{edge} is the pair-edge field. The non-interacting component of our base Hamiltonian H has the spin-degenerate band structure shown in Fig. 2.5. As evidenced by the figure, this model has a D_6 lattice symmetry and T -reversal symmetry. At half-filling, earlier studies found a metal - spin-ordered Mott insulator transition with increasing U/t where the intermediate regime, $U/t \approx 8.3 - 10.6$ is a gapped spin-liquid [108, 115]. Here we set the chemical potential μ such that the system is hole-doped away from half-filling.

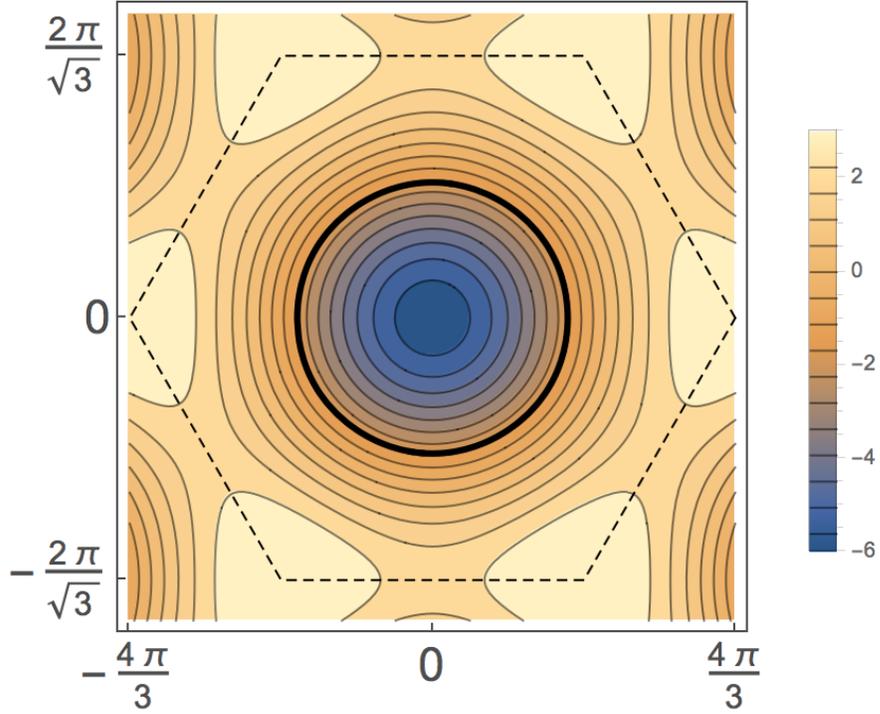


Figure 2.5: Contour plot of the band structure. The Brillouin zone is marked by the dashed line. A typical Fermi surface for our hole-doped calculations is marked by the thick, black line, here $n = 0.65$. The color legend indicates the energy of the band at a given k-point (for $t=1$).

The edge field we apply, H_{edge} , takes the form of

$$H_{edge} = \sum_{\langle i,j \rangle \in \text{Edge}, \sigma \neq \sigma'} V_{ij}^{\sigma\sigma'} c_{i\sigma} c_{j\sigma'} + h.c. \quad (2.12)$$

This edge field is depicted in Fig. 2.16 as red lines connecting sites adjacent to the left edge of our lattice. As is standard for DMRG, our calculation is performed on a cylinder where we use $L=18,24$ along the open direction and $W=3$ for the periodic direction. The introduction of a pair-field breaks the $U(1)$ gauge symmetry and the associated particle number conservation as well as the D_6 lattice symmetry of the Hubbard model in Eq. (2.14). Nevertheless, we can determine the susceptibility of the system to different pairing symmetries by observing the pairing response of the system under changes of the angular and

spin dependence of $V_{ij}^{\sigma\sigma'}$. The notable irreducible representations of D_6 include two 2-dimensional irreps E_1 and E_2 corresponding to linear combinations of triplet p-wave ($p \pm ip$) and singlet d-wave ($d \pm id$) basis functions respectively. We will look into the amplitude and phase of singlet and triplet pair fields.

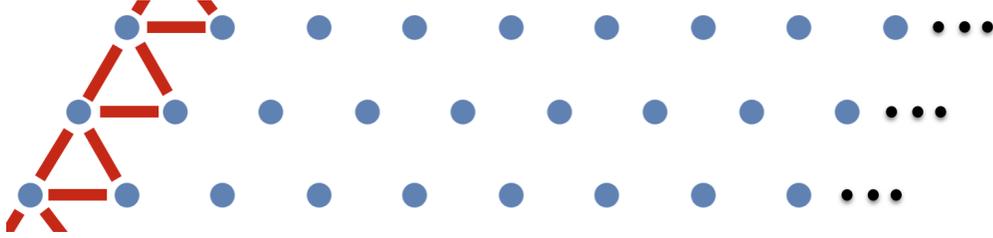


Figure 2.6: A portion of our lattice with the location of the pair edge-fields marked in red.

For our DMRG simulations we utilize the ITensor library developed by Miles Stoudenmire and Steve White [114]. We perform 14 sweeps with the final sweep containing up to 2500 states. The system is initialized by randomly sampling the even particle states, picking 10 states from each even particle sector and constructing a superposition of these states with coefficients randomly sampled from the interval $[-1,1]$. Proper implementation is ensured by double checking with exact diagonalization on small systems.

Note that since our calculation does not conserve particle number, the filling in the converged system has a non-trivial dependence on the interaction strength. Consequently, working at a specified filling is difficult as it is not directly set by the chemical potential. Thus for simplicity we set $\mu = 0$ for all simulations which corresponds to fillings $n \sim 0.65$ for $U/t = 2$ and $n \sim 0.4$ for $U/t = 10$ where n here is defined so that $n \in [0, 2]$ and half-filling corresponds to $n = 1$.

2.3.3 Results

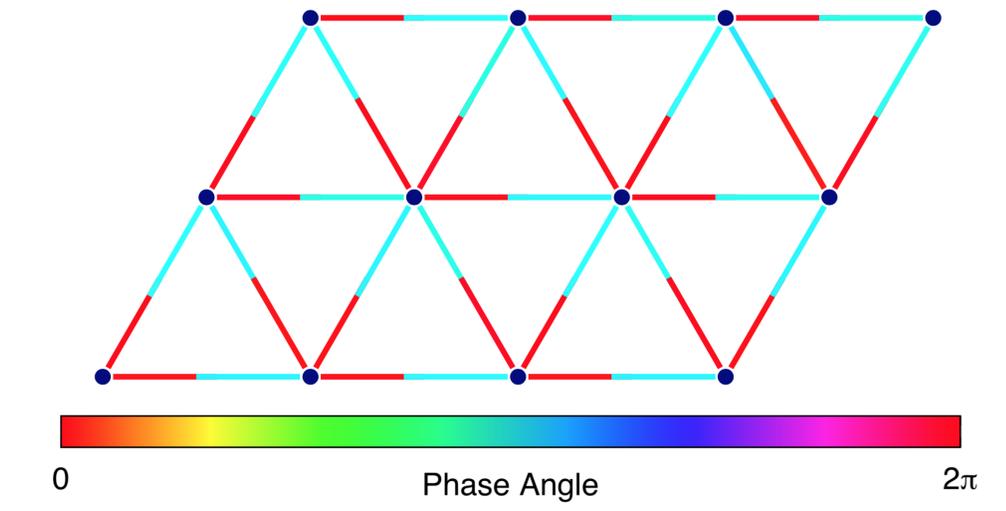
We assess the superconducting tendency in a particular channel by measuring the bond pair order parameter $\Delta_{ij} = \langle c_{i\uparrow}^\dagger c_{j\downarrow}^\dagger \rangle$, defined below for the singlet and triplet channels:

$$\begin{aligned}\Delta_{ij}^{singlet} &= \langle c_{i\uparrow}^\dagger c_{j\downarrow}^\dagger - c_{i\downarrow}^\dagger c_{j\uparrow}^\dagger \rangle \\ \Delta_{ij}^{triplet} &= \langle c_{i\uparrow}^\dagger c_{j\downarrow}^\dagger + c_{i\downarrow}^\dagger c_{j\uparrow}^\dagger \rangle.\end{aligned}\tag{2.13}$$

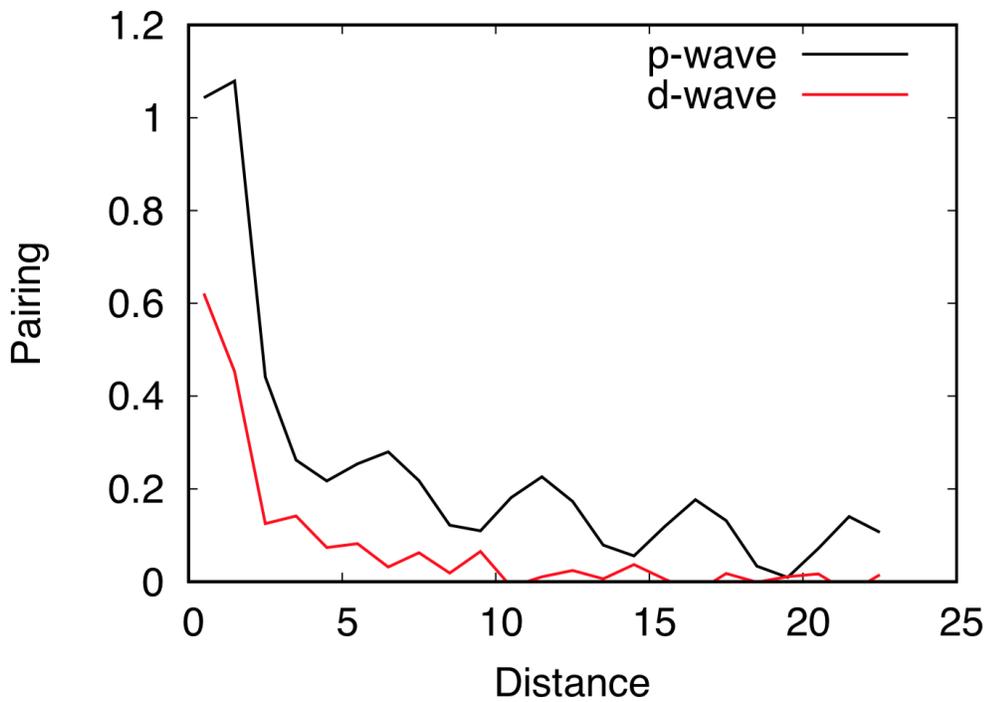
In order to let inherent pairing symmetry emerge, we study the phase structure of the bond pair order parameters in response to a *random phase* singlet and triplet edge-field with constant amplitudes. This allows the dominant pairing symmetry (assuming one exists) to develop naturally from the RG flow of DMRG. We look for phase coherence over the extent of the system and discern the symmetry of the phase coherent channel. In this context, phase coherence refers to the existence of extended superconducting domains with the same pairing symmetry i.e. an absence of domain walls or random phase perturbations. We then assess the strength of pairing tendency in each channel through the amplitude of the bond pair order parameter in response to the pair fields of *p*-wave and *d*-wave symmetries. We compare these responses in a system with moderately repulsive interaction of $U/t = 2$ and $n \sim 0.65$ to those in a system with strong repulsive interaction of $U/t = 10$ and $n \sim 0.4$.

For moderate repulsion of $U/t = 2$ and $n \sim 0.65$, we find a strong superconducting response with *p*-wave symmetry in the triplet channel and much weaker and phase incoherent superconducting response in the singlet channel. The dominance of *p*-wave pairing response for this moderate repulsion is evident from the phase structure established for the much of the system in response to the random edge field. We show the phase of the triplet channel

pair order parameter for $U/t = 2$ from the segment of the system in the bulk in Fig. 2.7(a).



(a)



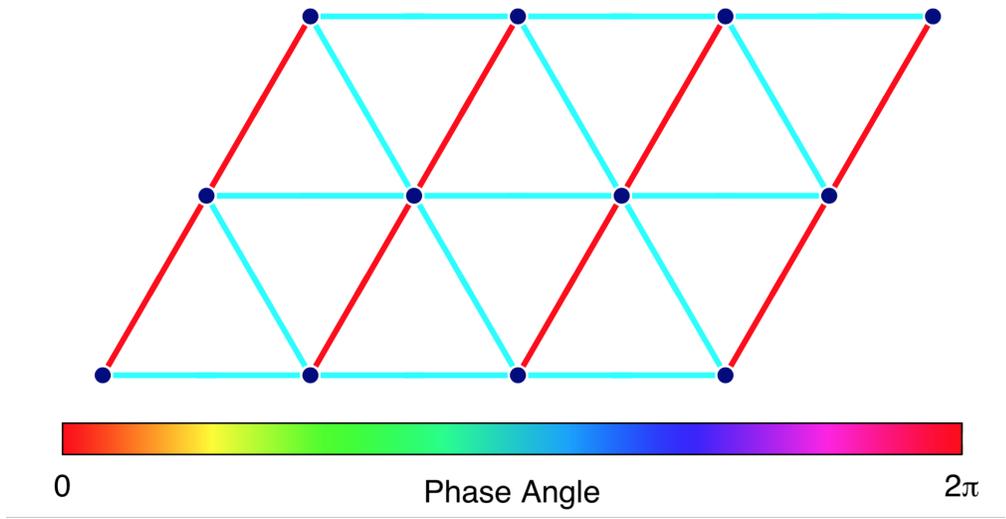
(b)

Figure 2.7: (a) SC phase plots for random edge-field, triplet channel, $U/t = 2$ and $n \sim 0.65$. The phase for each $\Delta_{ij}^{triplet}$ is represented by the color of the bond ij . (b) Amplitude of the SC response with p-wave edge-field (black) and d-wave edge-field (red) for $U/t = 2$. Only the parity channel corresponding to the parity of the edge field is shown. Reported amplitudes are normalized by the applied field ($V = 0.1$).

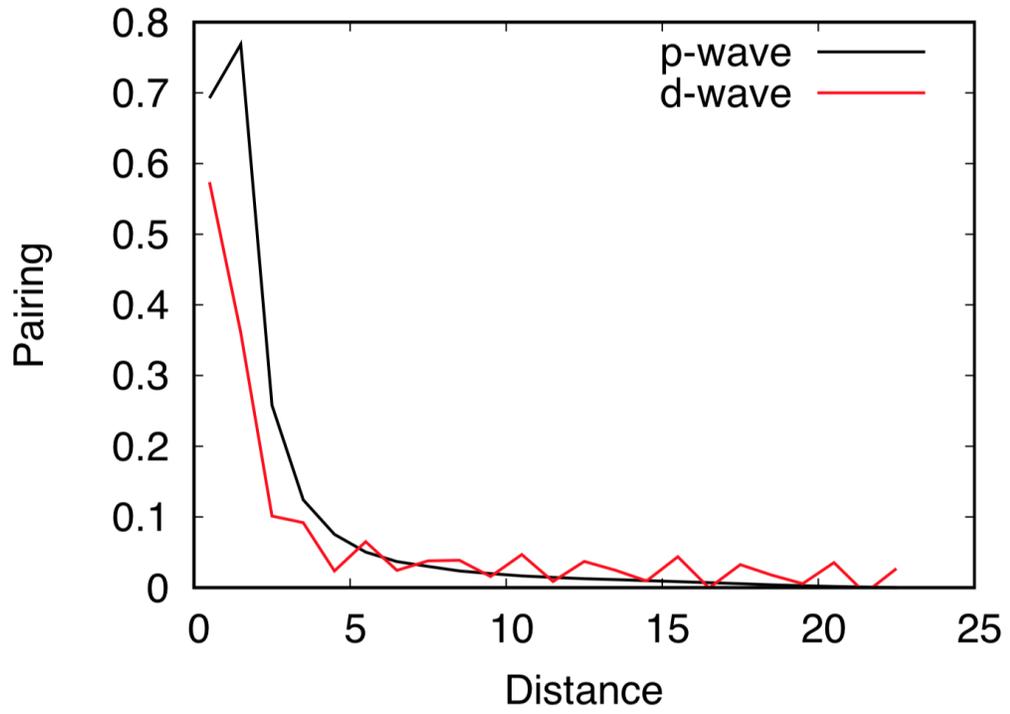
Note that in this phase plot, the phase undergoes a π -phase shift in between the drawn bonds since $\Delta_{ij}^{triplet}$ is odd parity i.e. $\Delta_{ij}^{triplet} = -\Delta_{ji}^{triplet}$. Here, the induced pairing symmetry about any given lattice point is clearly p-wave. We suspect that nodal SC occurs as a result of finite size effects and that its chiral counterpart $p \pm ip$ should emerge in the two-dimensional limit on the basis of Ginzburg-Landau theory in which nodeless pairing symmetries are energetically favorable since they fully gap the excitation spectrum and consequently maximize the condensation energy. Note that since a priori, the basis functions within a given irreducible representation are degenerate, the linear combination that maximizes the condensation energy should be preferred. However, this degeneracy does not strictly hold in the presence of finite size effects and the the edge field. Consequently, these small effects can bias the phase towards a different linear combination of eigenstates within this irreducible representation other than that which is expected for the 2D limit. However, in this context, the important result is the specific irreducible representation chosen by the ground state, not the specific linear combination of basis functions within that irrep. For the phase map of entire system, see Fig. 2.9 in the appendix. In the appendix we also show the phase disordered response in the singlet channel for $U/t = 2$ and discuss the DMRG convergence and truncation error.

Superconducting responses change qualitatively for strong repulsive interaction of $U/t = 10$ and $n \sim 0.4$. Here we see a response that is symmetry distinct from the above moderate U/t regime in response to the random phase edge fields. Namely, the induced superconductivity in the triplet channel is weak and phase incoherent while that in the singlet channel is dominant with d -wave symmetry. The truncated SC phase plot is given in Fig. 2.8(a). Again, we

expect $d \pm id$ in the 2D limit. The full phase map for both the singlet and triplet channel under random phase edge fields are given in the appendix 2.9 as is a discussion of the DMRG convergence and truncation error.



(a)



(b)

Figure 2.8: (a) SC phase plots for random edge-field, singlet channel, $U/t = 10$ and $n \sim 0.4$. The phase for each $\Delta_{ij}^{singlet}$ is represented by the color of the bond connecting site i and j . (b) Amplitude of the SC response with p-wave edge-field (black) and d-wave edge-field (red) for $U/t = 10$. Only the parity channel corresponding to the parity of the edge field is shown. Reported amplitudes are normalized by the applied field ($V = 0.1$).

In order to determine the strength of the response, we now turn away from the random edge-fields and instead employ pair-fields with p -wave and d -wave symmetries. Consistent with previous results, we find that for $U/t = 2$ and $n \sim 0.65$ the response in the p -wave channel is much stronger than the d -wave channel, see Fig. 2.7(b) while the opposite is true for $U/t = 10$ and $n \sim 0.4$ as seen in Fig. 2.8(b). This is reinforced by log-linear plots provided in the appendix. Note that the slightly stronger response of p -wave near the pinning field is curious, but we are ultimately interested in the induced bulk response where the approximate local crystal symmetry holds. Since the crystal symmetry is broken by the edge of our cylinder, a quantitative comparison of amplitudes along the edge is not meaningful. Thus the fact that p -wave is stronger near the edge in this case (and only slightly so) is likely due to edge effects.

2.3.4 Conclusion

Several recent works on this model at half-filling have focused on the realization of a chiral spin-liquid state at between the low U/t metal and the high U/t 120° Néel state [80, 108, 115]. In light of these spin-liquid works, investigating superconductivity in this model under hole-doping is especially interesting since the hole-doping of a spin-liquid, particularly a chiral one, is thought to yield exotic superconductivity [5, 75, 130]. Our results are consistent with this picture, though the study of the doping dependence is warranted.^{1 2}

The results discussed in this paper are particularly interesting in the context

¹Note that a previous study on this model [21] explored the regime $n \in [0.66, 1.0]$ where they found $d \pm id$ for $U/t = 8.5$ and $n \in [0.7, 1.0]$.

²An asymptotically-exact weak coupling solution also shows [93] p , d , and f -wave to be competitive but this is not our parameter regime.

of on-going experimental efforts in twisted TMD bilayers. Not only the hetero-TMD bilayers are well-described by the triangular lattice Hubbard model[136] the key parameter U/t is highly tunable in these systems through the twist angle. Our results suggest that the strength of the repulsion and filling in these systems would play an integral role in determining the pairing symmetry. Specifically, weak to moderate repulsion will likely facilitate $p \pm ip$ SC while stronger repulsion will tend towards $d \pm id$.

Note, that previous work by the authors [121] also explored superconductivity on the triangular lattice Hubbard model where they considered an effective model for hole-doped MoS₂. This material notably has spin-valley locking that was implemented through the incorporation of imaginary hoppings and some triangular warping. Here, we considered a model with conventional, real hoppings and find a transition between p-wave and d-wave order as opposed to the PDW order found previously. Together these works reinforce the picture that the underlying frustration of the triangular lattice offers a complex landscape for hosting superconductivity through strongly correlated physics.

ACKNOWLEDGEMENTS: We thank Allan MacDonald, Leon Balents, Liang Fu, Ashvin Vishwanath, Mike Zaletel, Donna Sheng, and Garnet Chen for helpful discussions. We acknowledge support from the National Science Foundation through the Platform for the Accelerated Realization, Analysis, and Discovery of Interface Materials (PARADIM) under Cooperative Agreement No. DMR-1539918.

2.3.5 Appendix: Phase Structure for Full Lattice

We provide the simulation results depicting the phase structure induced by the random phase edge field for the entire lattice as opposed to the truncated versions provided in the main text. We also provide the singlet and triplet components for each interaction regime studied.

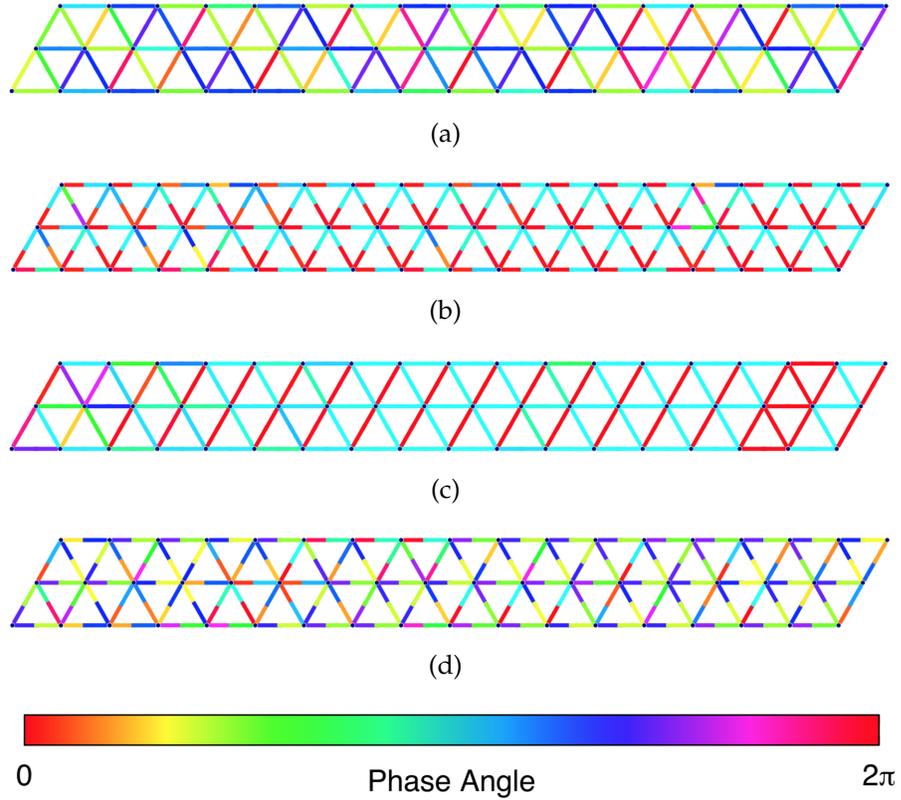


Figure 2.9: (a) SC phase plots for random edge-field, singlet channel, $U/t = 2$ and $n \sim 0.65$ (b) SC phase plots for random edge-field, triplet channel, $U/t = 2$ and $n \sim 0.65$ (c) SC phase plots for random edge-field, singlet channel, $U/t = 10$ and $n \sim 0.4$ (d) SC phase plots for random edge-field, triplet channel, $U/t = 10$ and $n \sim 0.4$.

2.3.6 Appendix: Truncation Error

We provide representative plots for truncation error of our $U/t = 2$ and $U/t = 10$ simulations in Fig. 2.10 below. It is well known from variational arguments that the error in the pairing strength is on the order of the square root of the truncation error. Recall that for the $U/t = 10$ calculations, the reported d-wave pairing strength is approximately 0.005. For these simulations, the truncation error is slightly above 1×10^{-8} so that the magnitude of the pairing strength is about 50 times the estimated error. For the $U/t = 2$ simulations, the reported pairing strength is around 0.02 for p-wave pairing while the truncation error is approximately 1×10^{-4} . Although, the signal to noise ratio is not as pronounced as the in the larger U case, the pairing strength is nevertheless still above the conventional error estimate. Furthermore, previous experience using DMRG to investigate superconductivity shows that in the case when superconductivity is not a feature of the system, the pairing strength decays much more rapidly and has no phase coherence i.e. the bulk superconducting phase looks random. Finally, note that this improvement in truncation error with increasing U/t is to be expected since the associated energetic gap increases with U/t .

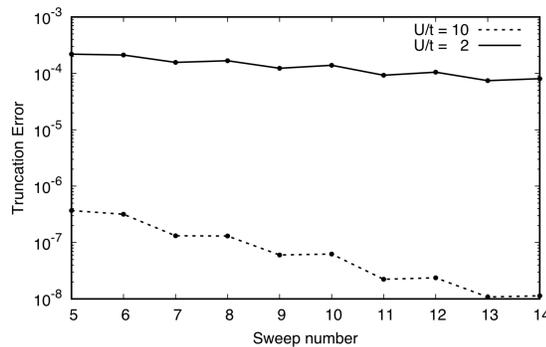


Figure 2.10: Truncation error for $L=18$ simulations with $U/t = 2$ and $U/t = 10$.

2.3.7 Appendix: Log-Linear Plots of Pair Amplitude

We present log-linear plots for the absolute value of the pairings below. The fact that d-wave pairing is dominant in the $U/t = 10$ regime can clearly be seen in 2.11(b).

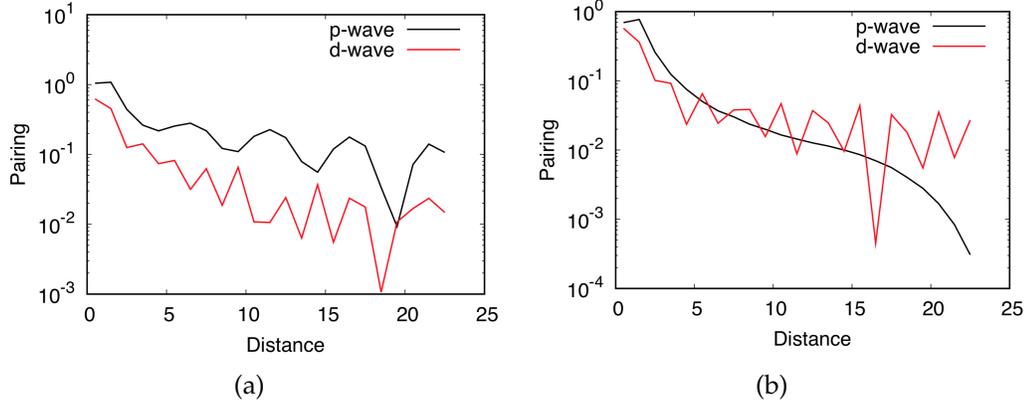


Figure 2.11: Log-lin plot of superconducting amplitudes presented in main text with $L=24$ for (a) $U/t = 2$ and (b) $U/t = 10$

2.3.8 Appendix: Charge Density

The charge density for the $U=10$ (with d-wave pair field on left edge) and $U=2$ simulations (with p-wave pair field on left edge) are provided below in Fig. 2.12. Modulation of the type seen for the $U=10$ case can directly be seen in non-interacting limit of a 1D tight-binding model where the charge modulations have the exact form $\frac{2}{L+1} \sum_{i=1}^N \sin\left(\frac{\pi i x}{L+1}\right)$. Of course our simulation is at strong coupling and is on a triangular lattice so we cannot expect to fully adhere to this, but if we consider an effective charge density of the form $A \frac{2}{L+1} \sum_{i=1}^N \sin\left(\frac{\pi i x}{L+1}\right)$ with free parameters A and N we can achieve a remarkably good fit with $A = 0.5, N = 14$ as evidenced below in Fig. 2.12(b).

However, the charge density of the $U=2$ case is complicated by the presence of appreciable pairing amplitude and the breaking of $U(1)$ gauge symmetry which together facilitate a direct coupling between the density and pairing. Within the context of a Ginzburg-Landau theory this allows for terms of the form $\rho_{+q}(\Delta_{-q} + \Delta_{+q}^*)$ through which charge and pairing interplay can occur. Nevertheless, the object of primary interest, the pair amplitude has structure consistent with previous systems in the literature. (DOI: 10.1016/j.physc.2012.04.004).

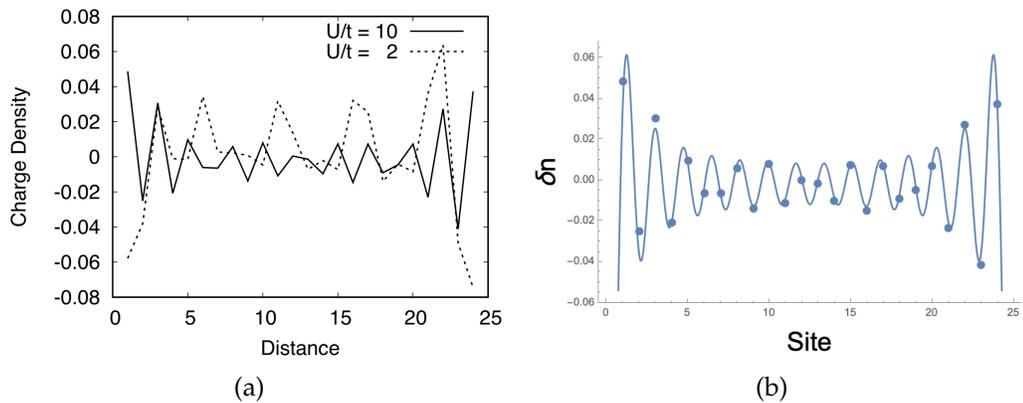


Figure 2.12: (a) Charge density along sites of central rung for $U=10$ (with d-wave pair field on left edge) and $U=2$ simulations (with p-wave pair field on left edge). (b) Charge density for $U=10$ simulation (same data as in (a) but shown here as points) and the fitted effective model shown as a continuous plot.

2.4 Observation of Pair Density Wave in Spin-Valley Locked Systems

This section is adapted from a Science Advances paper [123] with Eun-Ah Kim.

Cooper-pairs with a finite center-of-mass momentum form a remarkable state in which the superconducting order parameter is modulated periodically

in space. Although intense interest in such a ‘pair-density-wave’(PDW) state has emerged due to recent discoveries in high T_c superconductors, there is little theoretical understanding of the mechanism driving this exotic state. The challenge is that many competing states lie close in energy in seemingly simple models such as Hubbard model, in the strongly correlated regime[155]. Here we show that inversion symmetry breaking and the resulting spin-valley locking can promote PDW over more commonly found spin-stripe through frustration against magnetic order. Specifically, we find the first robust evidence for a PDW within density matrix renormalization group simulation of a simple fermionic model. Our results point to a tantalizing possibility in hole-doped group VI transition metal dichalcogenides, with spin-valley locked band structure and moderate correlations.

2.4.1 Introduction

Recent experimental and theoretical developments have brought a renaissance to the idea of a modulated superconducting state that spontaneously breaks translational symmetry (see Ref. [34] and references therein and Refs. [9, 11, 49, 112, 76, 33, 26, 128]). Earlier efforts towards realization of modulated superconductors[53, 156] or towards an interpretation of associated experiments[2] have relied on generating finite-momentum pairing using spin-imbalance under an (effective) magnetic field, in close keeping with the original proposals[41, 74] (FFLO). Alternatively, momentum space split, spinless fermions in the context of doped Weyl semi-metals[24] and hole-doped transition metal dichalcogenides[52] have been proposed as a platform for modulated superconductors due to pairing within fermi pockets centered

at finite crystal momentum. On the other hand, a modulated paired state proposed for cuprates requires a strong coupling mechanism.[9] Such a strong coupling driven state has been dubbed a PDW as a state distinct from FFLO-type superconductors.

The need for a strong coupling mechanism led to a search for the PDW state in numerical simulations. Numerous variational and mean-field studies have shown that pair-density wave type states are energetically competitive with uniform d-wave superconducting states in generalized t-J models, and it is thought that PDWs may become favorable in the presence of anisotropy.[40, 117, 92, 144, 58] Nevertheless, numerical evidence from the controlled approach of DMRG is lacking within simple fermionic models as the only evidence of PDW within DMRG was established in the 1D Kondo-Heisenberg model[10]. One signature difficulty in such a realization is that DMRG calculations on a Hubbard or t-J model on a square lattice with spin-rotation symmetry often find spin and charge stripe ground states instead of the PDW state. However, one could hope that frustrating spin order might nudge systems into a PDW state. Here we turn to a Hubbard model on the frustrated triangular lattice with broken inversion symmetry that captures the hole-doped monolayer group IV transition metal dichalcogenides (TMD's).

Rapidly growing interest in the monolayer group VI transition metal dichalcogenides (TMD's) has been fueled by the exotic possibilities driven by spin-orbit coupling and lack of centrosymmetry[100, 78, 138, 52, 36, 139, 106] as well as superconductivity in the n-doped TMD's[147, 64, 107, 100, 78, 138]. While the symmetry properties of the observed superconducting states remain unknown, the different translationally invariant superconducting channels for

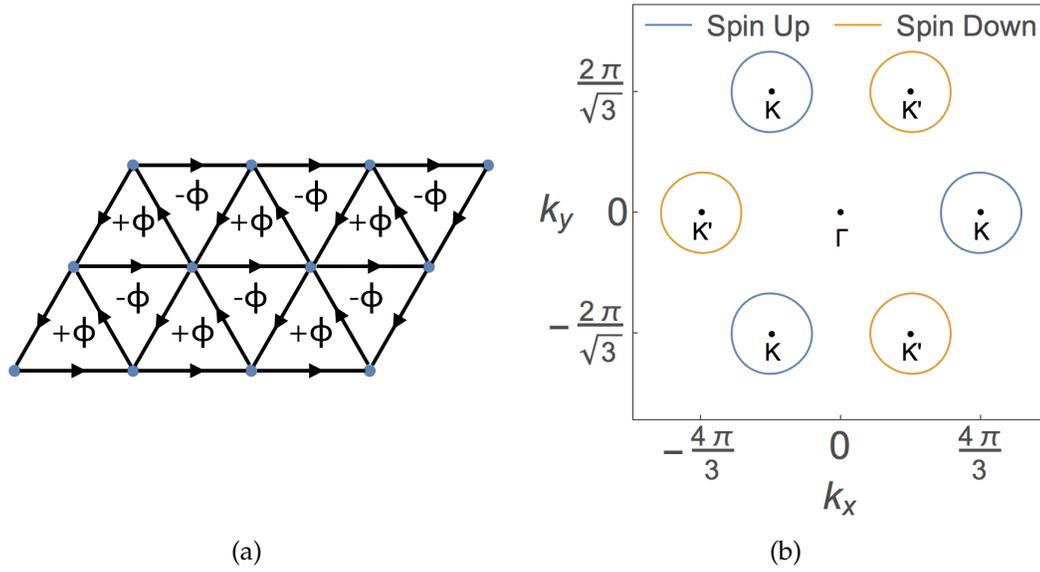


Figure 2.13: **Model and Fermi surface.** (a) The spin-dependent staggered flux pattern for one spin component with $\pm\Phi$ flux per plaquette. An opposite flux pattern for the other spin component guarantees time-reversal symmetry. The arrows indicate the direction of positive phase hopping. (b) Our Fermi surface with $t_{i+\hat{x},i;\uparrow} = \frac{2}{\sqrt{3}}e^{i0.3\pi}$ and $\mu = 4.6$ in the tight-binding model in Eqs: (2.14) (2.15). Here the spin-valley locked, circular Fermi pockets are evident.

the TMD's have been previously classified in mean-field studies [150, 111, 95]. Recently Hsu *et al.* [52] employed a weak-coupling RG approach to investigate a repulsive interaction driven pairing mechanism, predicting two topological superconducting instabilities with one of them being a spatially modulated intra-pocket state. However potentially strong correlation effects have largely been neglected despite the fact that the conduction electrons have substantial d -character. In this letter we use density matrix renormalization group (DMRG) calculations to study the effects of spin-orbit coupling on superconducting tendencies driven by repulsive interactions.

2.4.2 Results

In order to capture the spin-valley locked Fermi surfaces that occur in the valence band of group VI TMD's [100, 37, 97] in a one-band model we consider a nearest-neighbor tight-binding model on a triangular lattice with a staggered, spin-dependent magnetic flux of $\pm 0.9\pi$ per plaquette (see Figure 2.13(a)). This spin-dependent flux breaks the C_{6v} symmetry down to C_{3v} while preserving the time-reversal symmetry, mimicking the S_z preserving spin-orbit coupling present in the Kane-Mele model and generating two distinct spin-polarized pockets for our Fermi surface. Furthermore, the flux introduces a small amount of anisotropy in the pockets (see Fig 2.13(b)) analogous to that present in real materials such as MoS₂. [73] Finally, we include an on-site interaction U . Hence our model Hamiltonian is

$$H = - \sum_{\langle i,j \rangle} t_{ij,\sigma} c_{i\sigma}^\dagger c_{j\sigma} - \mu \sum_{i\sigma} c_{i\sigma}^\dagger c_{i\sigma} + U \sum_i n_{i\uparrow} n_{i\downarrow}, \quad (2.14)$$

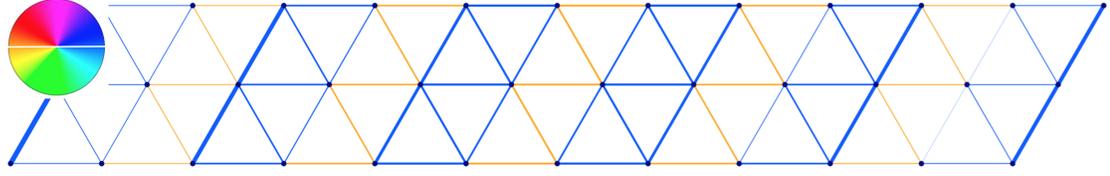
where $t_{ij,\sigma}$ is the spin-dependent complex nearest-neighbor hopping, μ is the chemical potential, and U is an on-site Hubbard interaction. The dispersion takes the form of

$$\epsilon_\sigma(\mathbf{k}) = -2 \sum_i \left[\text{Re}(t) \cos(\delta_i \cdot \mathbf{k}) + \text{Im}(t) \sin(\delta_i \cdot \mathbf{k}) \sigma_z \right] \quad (2.15)$$

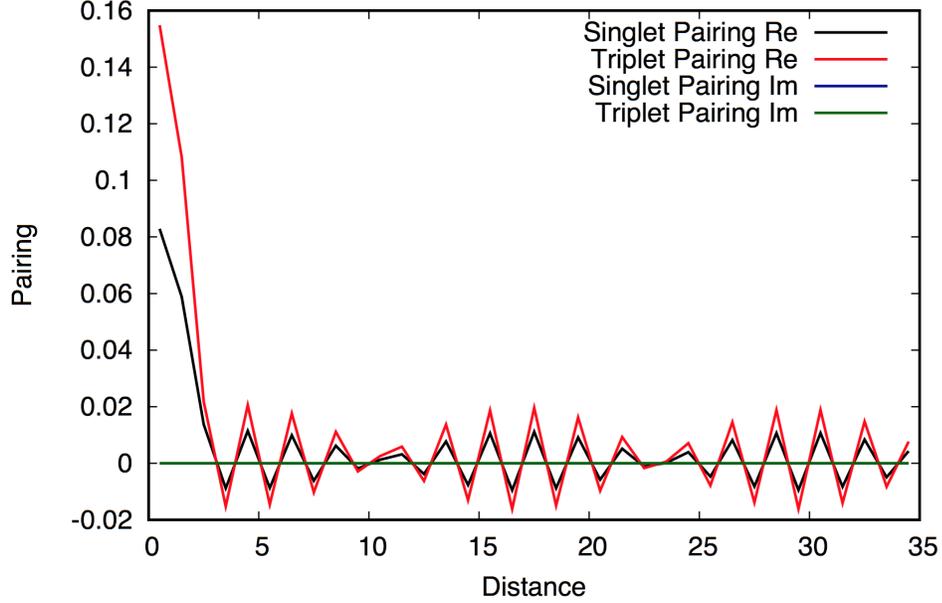
where $\delta_i \in \{\hat{\mathbf{x}}, -\frac{1}{2}\hat{\mathbf{x}} + \frac{\sqrt{3}}{2}\hat{\mathbf{y}}, -\frac{1}{2}\hat{\mathbf{x}} - \frac{\sqrt{3}}{2}\hat{\mathbf{y}}\}$, $\sigma_z = \pm$ for spin up and down respectively, we define t such that $t = t_{i+\hat{\mathbf{x}},i;\uparrow}$, and the lattice spacing has been set to 1.

Within this model, we probe the superconducting response to a pair-edge field. We initially consider a moderately attractive interaction with $U=-2$, where uniform pairing in the A1-irrep is expected. Data for the following statements may be found in Supplementary Materials III. First, we note that the phase disorder due to edge field effects quickly disappears upon moving

into the bulk [see Fig. A2.17(a) in Supplementary Materials I]. Inspecting the bond-singlet component of the superconducting order parameter $\Delta_{\langle ij \rangle}^{singlet}$ on the directed nearest-neighbor bonds away from the probe, we find a well-ordered uniform phase structure, Fig. A2.19(b). Moreover, zooming into the phase structure within a single unit-cell, we find the pair-field expectation value to be isotropic and definitively s-wave. The uniform and isotropic nature of the order parameter phase distribution is a robust property of our results in the negative U regime, that is insensitive to the profile of the edge fields and occurs for all system sizes studied. The triplet channel behaves analogously, displaying homogeneous f-wave pairing [see Fig. A2.19(c) in Supplementary Materials III].



(a)



(b)

Figure 2.14: **Evidence of PDW oscillations.** (a) $\text{Arg}(\Delta_{\langle i,j \rangle}^{\text{singlet}})$ for all nearest-neighbors with $U = +2$ for our 3×36 lattice with periodic boundary conditions along the short direction and open boundary conditions along the long direction. For visibility, we truncate the plot so that only the third farthest from the edge field is shown. The line thickness is proportional to the pairing amplitude. (b) We plot the real and imaginary components of $\Delta_{ij}^{\text{singlet}}$ and $\Delta_{ij}^{\text{triplet}}$ for i,j along the middle rung of our lattice in order to present the phase oscillations.

Armed with the attractive, $U < 0$, result that can serve as a reference, we now study the moderately repulsive Hubbard regime with $U = +2$. Note that estimates for the band-structure parameters of MoS_2 have the SOC to be 0.08 eV, hopping strength to be 0.5-1 eV, and on-site repulsion of the 4d Mo orbitals to be 2-10 eV [96, 97]. Thus with $U/t \approx 2$ for our calculations, we lie at the lower

end of these estimates for the effective interaction strength of these materials. Given the difficulty in estimating U/t for actual materials, the TMDs represent a realistic material setting for our simulation, at least to within the errors of these estimates. Earlier work using two-stage perturbative RG on a similar spin-valley locked model with repulsive U predicted superconductivity in the two dimensional E representation where some linear combination of p and d -wave symmetries occurs due to the lack of inversion symmetry[52]. Unexpectedly, our DMRG simulation in this repulsive interaction regime reveals a tendency to break translational symmetry along the length of the cylinder. Specifically the system forms a modulated paired state where both the singlet and triplet bond pair order parameters are everywhere real with modulation in their sign (see Fig. 2.14). From the symmetry perspective the observed state is analogous to the state proposed by [74]. Such modulation in the pair amplitude is evident in the plot of $\Delta_{\langle ij \rangle}^{singlet}$ for $U = +2$ in Fig. 2.14(a) where an anisotropic phase structure within the unit-cell is repeated with period 2. A similar unit-cell doubling of the phase is seen in the triplet channel [see Supplementary Materials IV]. We find this tendency to form a PDW is robust against changes in chemical potential although the periodicity changes appropriately based on the filling [see Supplementary Materials V]. For instance, increasing the chemical potential, μ , from $\mu = 4.6$ to $\mu = 6.0$ enlarges the unit-cell by an additional lattice site [see Fig. 2.22]. Although there has been much interest in modulated superconducting states, this is the first report of a strong coupling driven PDW within DMRG simulations of a simple fermionic model to the best of our knowledge.

In order to further understand this translational symmetry breaking, we examine the oscillations of the superconducting order parameter. Since $\Delta_{ij}^{singlet}$

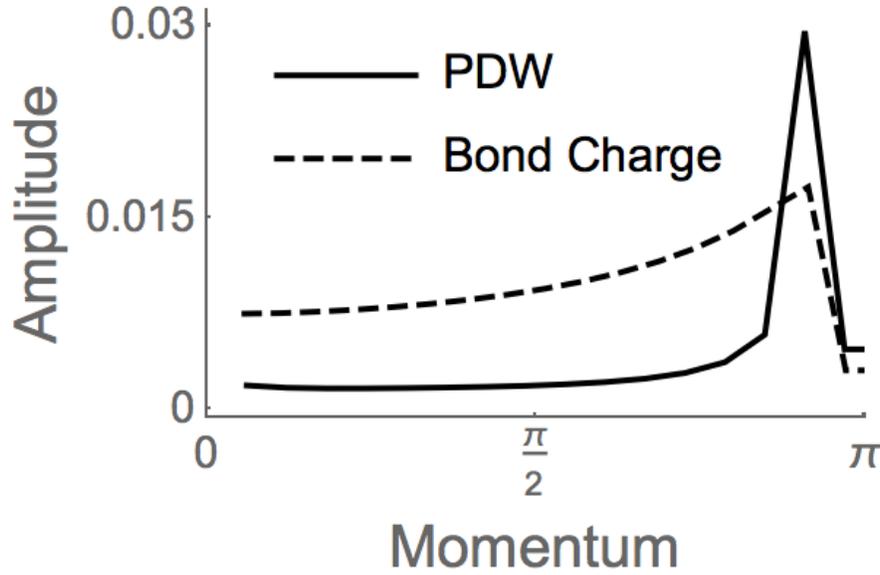
is characterized exclusively by π -phase shifts in the bulk, we may project it to real space after a gauge transformation and look at the decay properties of the pair field by plotting it for bonds directed along the middle rung, see Fig. 2.14(b). For attractive interactions, the singlet pairing strength falls off gradually as expected from the quasi-1D geometry of the system and exhibits slowly varying oscillations [see Supplementary Materials III Fig. 2.20] due to finite size effects induced by the open boundary conditions[110]. On the other hand, the pairing profile for the $U = +2$ simulation shows additional rapid oscillation about zero, see Fig. 2.14(b). These oscillations about zero occur for all repulsive Hubbard simulations near $U = +2$. Thus although the exact strength of the pairing response and the penetration depth of the edge field appear to have some dependence on the edge field profile and the length of the lattice, the PDW-type behavior reported has been observed for all system sizes and all edge-field types. Intriguingly, this plot strongly resembles the plot of the same quantity in the Kondo-Heisenberg model with PDW[10]

Fourier transforming these oscillations, we find a single dominant wavevector that is approximately twice the Fermi radius of a single pocket, $Q = 2k_F$, suggesting that the finite momentum of our Cooper pairs originates from interplay between the pockets. This picture is reinforced by probing the effect of increasing the chemical potential (decreasing the pocket radius). Here a PDW also develops and we see a shift in wavevector consistent with our change in pocket size see Fig. 2.23. A schematic of the finite momentum pairing is presented in Fig. 2.15(b). Since our PDW modulations are unidirectional and lie orthogonal to the applied edge field, our results are especially relevant to proximity induced superconductivity in the TMDs.

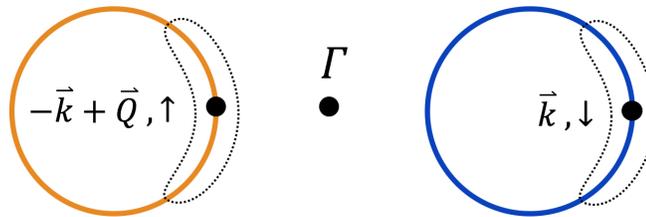
Finally, we compare oscillations in the singlet pairing strength and in the bond charge density. As the pairing amplitude profile of our $U = +2$ simulation has net pair amplitude on the whole system due to the edge field and our edge field explicitly breaks gauge symmetry, charge modulation of the same period ($Q = 2k_F$) is anticipated to be driven by the net component and the modulated pairing components[34]. From a Ginzburg-Landau perspective these add terms of the form $\rho_{-Q}(\Delta_0^*\Delta_{+Q} + \Delta_0\Delta_{-Q}^*)$ and $\rho_{-Q}(\Delta_{+Q} + \Delta_{-Q}^*)$ respectively to the free energy that account for the development of $2k_F$ CDW oscillations as opposed to $4k_F$ oscillations from $\rho_{-2Q}\Delta_{-Q}^*\Delta_{+Q}$ where ρ_Q and Δ_Q here correspond to the density and PDW order parameters with momentum Q . This is clearly seen in the Fourier decomposition of the PDW and bond charge order presented in Fig. 2.15(a) where both orders are dominated by the same Fourier mode. We remark that while both the attractive and repulsive charge bond densities have oscillations that dip below their mean values only the superconducting response of the repulsive case has oscillations about zero suggesting that the PDW oscillations are not finite size effects.

2.4.3 Discussion

In summary, we used DMRG to study the superconducting tendencies of a repulsive- U Hubbard model on a triangular lattice with spin-valley locking. These tendencies were probed by studying the pairing response profile in response to uniform and random pair fields along one edge. Our calculations indicate that the superconducting phase diagram of the model may be more complex than what was revealed by the previous perturbative RG study [52], with translational symmetry breaking superconducting states possibly in



(a)



(b)

Figure 2.15: **Fourier decomposition of PDW and bond charge order.** (a) Fourier transforms of the PDW and charge bond order. Zero momentum i.e. constant contributions and decay effects have been removed. (b) Depiction of pairing in momentum space. The regions demarcated by dashed lines are the approximate pairing regions.

competition with a uniform state. The PDW observed breaks translational symmetry with the superconducting order parameter alternating sign. This atypical pairing response may be related to the fact that Ising spin-orbit coupling and the triangular lattice conspire to frustrate any spin order including spin stripe. Indeed we fail to find any appreciable spin response to be induced by an S_z -edge field coexistent with our pair-field. Here the moment dies off rapidly away from the edge leaving no discernible moment and reaching

practically zero ($5e-5$ in units of the applied field) by site 4. It will be interesting to study whether the observed PDW state can be found in a truly two-dimensional setting using a different method such as density matrix embedding theory (DMET) where long-range order can be observed.[154]

2.4.4 Materials and Methods

DMRG is a powerful, non-perturbative method for studying strongly interacting systems [103, 143, 135, 133, 10, 63]. It has been used with great success to explore a diverse selection of strongly correlated phenomena highlighted by stripes, spin-liquids, and superconductivity [143, 135, 133, 10, 63, 158, 51, 59, 61, 62]. However, since DMRG is quasi-1D in nature no true long-range order can be seen in the correlations. Thus in order to access our system's superconducting tendencies we implement a pair-edge-field motivated by the field-pinning approach underlying several earlier studies [134, 113, 143, 10]. By biasing the system towards a particular superconducting state and studying the emergent symmetry of the appropriate order parameter in the bulk, one can gauge the model's propensity for various instabilities.

Even in the absence of the spin-valley locking special to our model, the geometric frustration of the triangular lattice is known to foster exotic phases. While a consensus has emerged that the ground state of the Heisenberg model is a 120° Néel antiferromagnet[132], the Hubbard model has been shown to have tendencies toward spin liquid and chiral d+id superconducting states.[99, 22]. It is notable that within the context of the Heisenberg model frustration can inhibit spin-stripes.

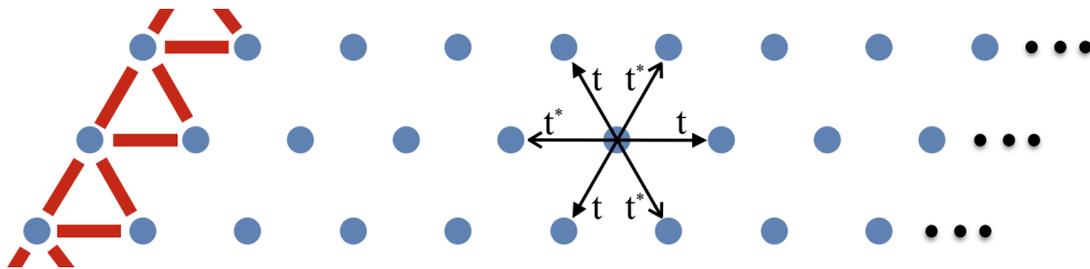


Figure 2.16: **Lattice and edge field.** A depiction of our lattice. It is periodic in the short direction with 3 unit cells and has open boundaries in the long direction. The ellipses on the right signify that multiple lengths are studied: $L = 12, 18, 24, 36$. The edge field, shown as red lines, is a pair-field of the form given in equation 2.17. The nearest-neighbor hopping structure for spin up is also shown with the spin down hopping structure being the complex conjugate of that shown above.

We emphasize that due to the lack of inversion symmetry, even and odd pairing components can coexist[109]. Thus the S_z preserving spin-orbit coupling allows for the mixing of $S_z = 0$ singlet and triplet states, i.e., the bond pair order parameter $\Delta_{ij} = \langle c_{i\uparrow}^\dagger c_{j\downarrow}^\dagger \rangle$ should be an admixture of singlet and triplet components:

$$\begin{aligned}\Delta_{ij}^{singlet} &= \langle c_{i\uparrow}^\dagger c_{j\downarrow}^\dagger - c_{i\downarrow}^\dagger c_{j\uparrow}^\dagger \rangle \\ \Delta_{ij}^{triplet} &= \langle c_{i\uparrow}^\dagger c_{j\downarrow}^\dagger + c_{i\downarrow}^\dagger c_{j\uparrow}^\dagger \rangle\end{aligned}\tag{2.16}$$

Note that since our system lacks translational symmetry due to open boundary conditions, the pairing symmetry is not constrained to transform under a single irreducible representation. Nevertheless, the real-space structure of these bond-centered order parameters provides insight into the nature of the dominant pairing state.

We carry out our DMRG simulations on a cylinder with 3 unit cells in the periodic direction and 12, 18, 24, and 36 unit cells in the non-periodic direction. The width is sufficiently large to sample both types of pockets in the Fermi surface but not so large as to make DMRG prohibitively expensive

for our available computational resources. We keep the band structure fixed and explore the effects of varying U . We investigate the superconducting susceptibility by applying a pair-field along one edge as illustrated above in Fig. 2.16. In order to reveal any inherent preferences for a particular superconducting channel, we consider two different phase structures for the edge field: a uniform field described by the A1 irrep and a random field i.e.

$$\Delta_{ij}^{edge} = Vc_{i\uparrow}^\dagger c_{j\downarrow}^\dagger e^{i\phi_{ij}} + \text{h.c.} \quad (2.17)$$

with the phase ϕ_{ij} chosen to transform under the A1 irrep or be random for $i \neq j$ and 0 for $i = j$. We remark that all results presented in this paper have been shown to be independent of the phase structure of the edge field applied [see Supplementary Materials I]. The strength of the pair-field was fixed to be $V=0.1$, about an order of magnitude less than the hopping amplitude which is consistent with that used in previous studies.[134, 113, 10] While the cylindricity of our geometry and the addition of a pair field break the C_{3v} symmetry of the lattice and translational invariance, if there is a well-defined structure to order parameter in the bulk, we expect to gain information about the inherent superconducting tendencies through the real space structure of the order parameters in Eq. (2.16).

Our DMRG calculation utilizes the iTensor library developed by Miles Stoudenmire and Steve White.[114] We perform up to 14 sweeps with a final bond dimension of $M=2500$. This is sufficient to obtain energy convergence to $O(10^{-7})$ for our repulsive calculations and $O(10^{-10})$ for attractive calculations [see Supplementary Materials II]. We focus exclusively on the inter-pocket instabilities of our model where we may exploit the conservation of the $S_z=0$ quantum number. S_z and fermion parity, N_2 , are conserved quantum numbers, but the U(1) particle number symmetry is broken by the pair field. As our

starting point, we construct a MPS that randomly samples the $S_z=N_2=0$ sector of our Hilbert space picking 10-30 states from each even particle number sector with the number of states depending on the system size. In addition, we use exact diagonalization (ED) to ensure correctness in the DMRG simulations with the energetics from the two methods agreeing to within machine precision for small 3×3 systems where ED is computationally tractable.

2.4.5 Acknowledgements

We thank Leon Balents, Erez Berg, Garnet Chan, Yi-Ting Hsu, Steve Kivelson, Kam Tuan Law, Patrick Lee, Steve White, and Hong Yao for useful discussions. E-AK acknowledges Simons Fellow in Theoretical Physics Award #392182 and DOE support under Award DE-SC0010313. E-AK is grateful to the hospitality of Kavli Institute of Theoretical Physics supported by NSF under Grant No. NSF PHY-1125915, where this work was completed. JV acknowledges support by the National Science Foundation (Platform for the Accelerated Realization, Analysis, and Discovery of Interface Materials (PARADIM)) under Cooperative Agreement No. DMR-1539918 and in part by NSF DMR-1308089.

2.4.6 Supplementary Materials I: Phase Structure Near the Pair-Field

Since in the main text we only provide the phase structure of the superconducting order parameter farthest from the pair-field, we now show the edge closest to the probe. We do this for both the random edge field and the

uniform edge field. Only the singlet plots are given as evidence but the triplet channel behaves analogously. These plots illustrate the effect of the pair-field in inducing whatever phase structure it has into the wavefunction and highlight the tendency for the pairing to settle into a dominant phase far from the edge field.

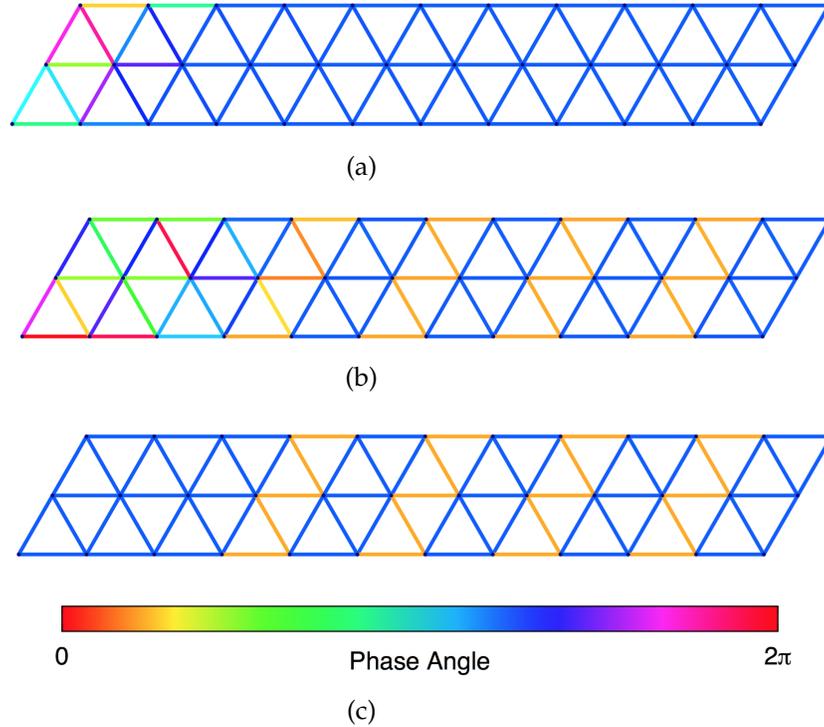


Figure 2.17: **Singlet phase structure near the edge-field.** The singlet phase structure, $\text{Arg}(\Delta_{\langle i,j \rangle}^{\text{singlet}})$, for (a) attractive interactions, $U = -2$, and (b) repulsive interactions, $U = +2$, with random pair-edge-field and (c) repulsive interactions with uniform edge field. Here the $L = 36$ lattice is shown where now we plot only the third closest to the pair-field. As before, the lattice has periodic boundary conditions along the short direction and open boundary conditions along the long direction. Unlike previous plots of this kind, the line thickness here is constant since we only want to emphasize the phase change and the decay from the edge makes this more difficult to see.

2.4.7 Supplementary Materials II: DMRG convergence

We check the convergence of our DMRG simulations by looking at the change in energy between sweeps. For all simulations we start our first sweep with a maximum bond dimension of 500 states and by the 14th sweep keep up to 2500 states, the maximum allowed by our current RAM limitations. We specify an SVD truncation error of 10^{-12} but in practice find that convergence within iTensor is limited to $\mathcal{O}(10^{-10})$. Typical convergence is achieved to $\mathcal{O}(10^{-7})$ for repulsive interactions and $\mathcal{O}(10^{-10})$ for attractive interactions. Below in Fig. 2.18, we plot the decrease in energy between sweeps for the $L = 36, U = +2$ calculation presented in the main text.

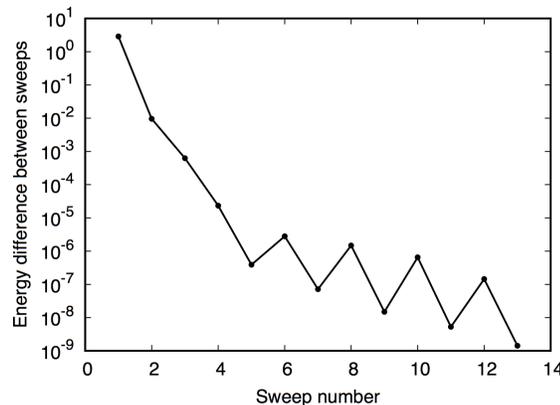


Figure 2.18: **DMRG convergence.** The decrease in energy between sweeps for the DMRG calculation in the main text ($L = 36, U = +2$) demonstrating convergence.

2.4.8 Supplementary Materials III: Phase Structure for Attractive Interactions

In order to establish a reference for our pair-edge field approach, we explore the effect of an attractive Hubbard interaction, $U = -2$, on the superconducting preferences. More specifically, we look at the bond-singlet and bond-triplet components of the superconducting order parameter along directed nearest-neighbor bonds. Here we find that regardless of the phase structure of the edge field, uniform or random, homogenous pairing is established in the bulk and any phase disorder due to the edge field quickly disappears upon moving away from the edge field [see Fig. A2.17(a) in Supplementary Materials I]. This induced, translationally invariant phase structure is found to be an admixture of s -wave and f -wave pairing, as expected for a conventional superconductor breaking parity symmetry. The uniform-A1 behavior of the bulk order parameter is robust, being insensitive to the profile of the edge fields and system size. Note that we have opted to use a different scheme for presenting the phase here as compared to the main text since the triplet component cannot be presented using the previous approach. We use this alternative style whenever plotting triplet phases but include it also for the singlet channel here along with the conventional style in order to help understand what's being presented.

We also provide the pairing strength along the middle rung analogous to the plot in Fig. 2.14(b). Here we plot the real and imaginary components of $\Delta_{ij}^{singlet}$ and $\Delta_{ij}^{triplet}$ for attractive Hubbard interactions with ij along the middle rung of our lattice. Comparison to the analogous plot in the main text for the repulsive case, Fig. 2.14(b), highlights the presence of oscillations about 0 in the repulsive case that are absent in the attractive case.

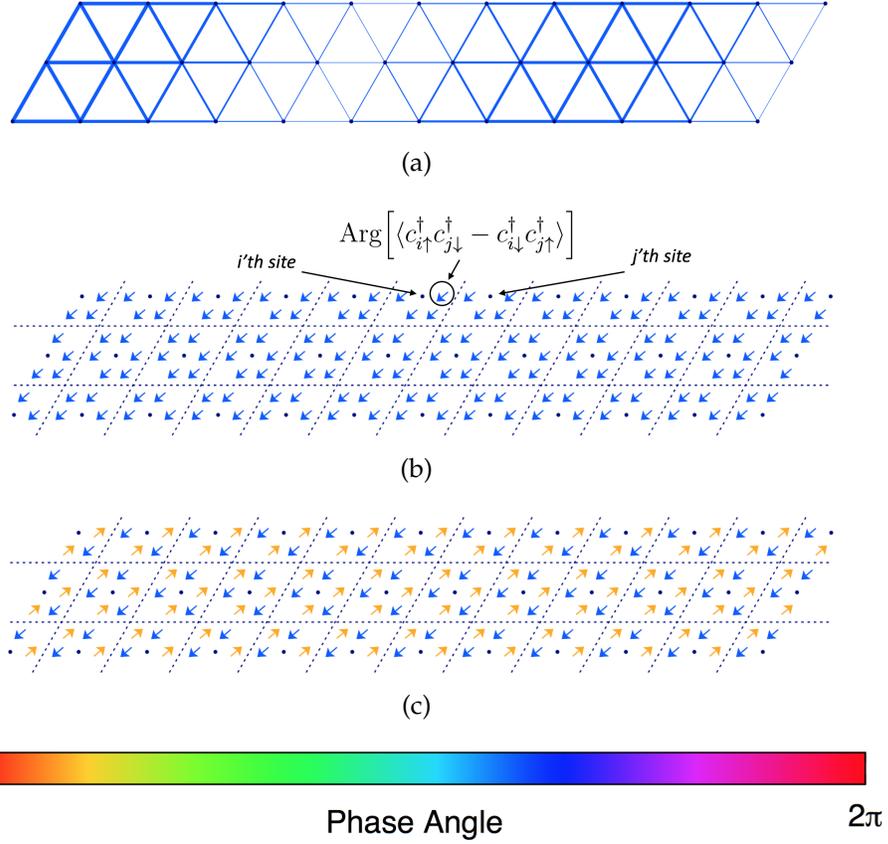


Figure 2.19: **Phase structure for attractive interactions.** (a) $\Delta_{\langle ij \rangle}^{singlet}$, (b) $\text{Arg}(\Delta_{\langle ij \rangle}^{singlet})$ and (c) $\text{Arg}(\Delta_{\langle ij \rangle}^{triplet})$ for $U = -2$ along all directed nearest-neighbor bonds. Note that (a) and (b) presents the same phase data and (a) is given only to help the understanding of the new style of plotting. Here we provide the $L=36$ results, but only show the third of the lattice away from the edge probe. Recall our lattice has periodic boundary conditions along the short direction and open boundary conditions along the long direction. The arrows point in the direction of the phase and are also colored according to the argument. Although redundant, this is done to aid visibility. Note that this method of plotting possesses an additional redundancy in that $\Delta_{\langle ij \rangle}^{singlet} = \Delta_{\langle ji \rangle}^{singlet}$ and $\Delta_{ij}^{triplet} = -\Delta_{ji}^{triplet}$.

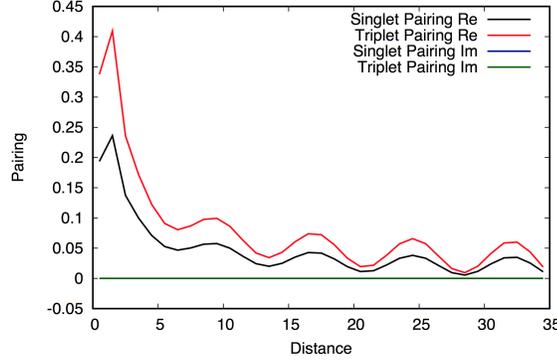


Figure 2.20: **Amplitude for attractive interactions.** The real and imaginary components of $\Delta_{ij}^{singlet}$ and $\Delta_{ij}^{triplet}$ for the attractive Hubbard regime, $U = -2$, of the $L=36$ lattice with uniform edge field. Here i,j lie along the middle rung of our lattice.

2.4.9 Supplementary Materials IV: Triplet Phase Structure for Repulsive Interactions

In Fig. 2.21 we provide the triplet phase plot of the $L=36$ lattice with uniform edge field for $U = +2$ case shown in the main text. Again, for the sake of visibility, only the third of the lattice farthest from the edge field is displayed. This result is qualitatively similar to the corresponding singlet case shown in the main text, Fig. 2.14(a), in that they both break translational symmetry along the length of the cylinder with a doubling of the unit-cell. Due to the breaking of translational symmetry, this system is not amenable to the kind of point-group symmetry analysis performed in Supplementary Materials II.

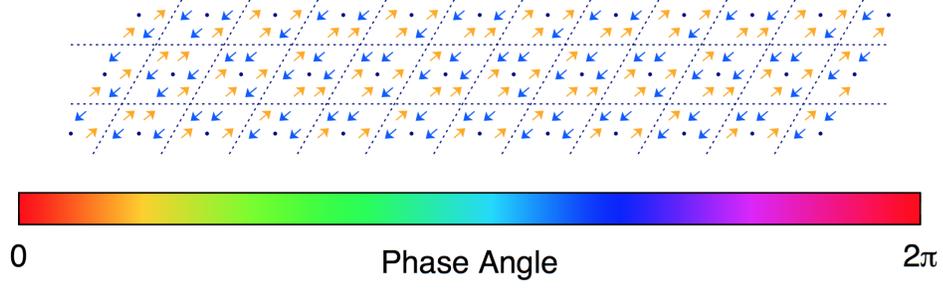


Figure 2.21: **Triplet phase structure for repulsive interactions.** The triplet superconducting phase structure, $\text{Arg}(\langle c_{i\uparrow}^\dagger c_{j\downarrow}^\dagger + c_{i\downarrow}^\dagger c_{j\uparrow}^\dagger \rangle)$, for all nearest-neighbors with $U = +2$ for our 3×36 lattice with a uniform edge-field, periodic boundary conditions along the short direction, and open boundary conditions along the long direction. For visibility, we truncate the plot so that only the third farthest from the edge field is shown. The arrow points in the direction of the phase and is also colored according to the argument for visibility. Note that this method of plotting possesses a redundancy in that $\Delta_{ij}^{\text{triplet}} = -\Delta_{ji}^{\text{triplet}}$.

2.4.10 Supplementary Materials V: Effects of Chemical Potential on the PDW

We explore the role of the chemical potential in the PDW-like structure seen in the repulsive Hubbard regime. To this end we shift the chemical potential from $\mu = 4.6$ to $\mu = 6.0$ for a fixed Hubbard interaction strength of $U = +2$. The resultant bond-centered superconducting pairing for the singlet channel is provided in Fig. 2.22. Herein we see a similar behavior to the $\mu = 4.6$ case in that there is translational symmetry breaking along the length of the cylinder. However, rather than a doubling of the unit-cell, we see a tripling of the unit-cell. This tripling also occurs in the triplet channel which is not shown. We note that the apparent randomness on the left side of the plot isn't due to edge field effects since this is the third of the lattice away from the probe and instead is an artefact of the small amplitude there relative to numerical convergence. Fourier

transforming the PDW order shows that this PDW is characterized by a single wavevector that is shifted from the wavevector of the PDW at $\mu = 4.6$. This is consistent with the suggestion that the wavevector is given by the Fermi pocket diameter as discussed in the main text.

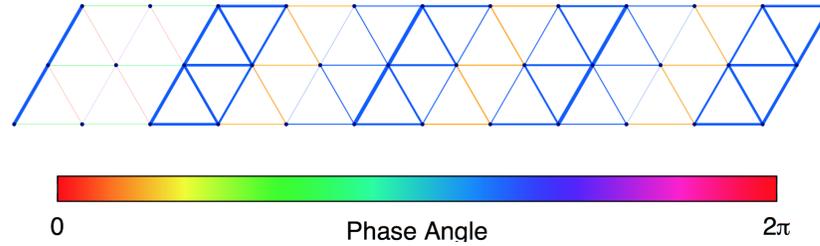


Figure 2.22: **Effect of chemical potential on PDW phase structure.** The phase of the bond-centered singlet superconducting order parameter, $\Delta_{\langle i,j \rangle}^{singlet}$ for nearest-neighbors on the 3×36 lattice with random pair-edge-field where now a larger chemical potential of $\mu = 6.0$ is used rather than that in the main text, $\mu = 4.6$. Here, the line thickness is proportional to the amplitude. The lattice has periodic boundary conditions along the short direction and open boundary conditions along the long direction.

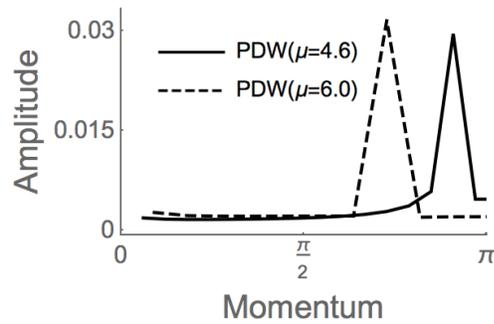


Figure 2.23: **Effect of chemical potential on PDW dominant Fourier mode.** The Fourier transforms of the $\mu = 4.6$ and $\mu = 6.0$ PDW. Zero momentum i.e. constant contributions and decay effects have been removed.

CHAPTER 3

MACHINE LEARNING MANY BODY LOCALIZATION

Neural-network-based machine learning is emerging as a powerful tool for obtaining phase diagrams when traditional regression schemes using local equilibrium order parameters are not available, as in many-body localized (MBL) or topological phases. Nevertheless, instances of machine learning offering new insights have been rare up to now. In this chapter, we show that a single feed-forward neural network can decode the defining structures of two distinct MBL phases and a thermalizing phase, using entanglement spectra obtained from individual eigenstates. For this, we introduce a simplicial geometry-based method for extracting multipartite phase boundaries. We find that this method outperforms conventional metrics for identifying MBL phase transitions, revealing a sharper phase boundary and shedding new insight on the topology of the phase diagram. Furthermore, the phase diagram we acquire from a single disorder configuration confirms that the machine-learning-based approach we establish here can enable speedy exploration of large phase spaces that can assist with the discovery of new MBL phases. We will first give an overview of supervised machine learning specifically neural networks and provide a short introduction to many-body localization.

3.1 Foray into the World of Machine Learning

Machine learning (ML) is a rapidly growing branch of computer science dedicated to building algorithms capable of learning from data. It has enabled us to tackle problems that are oftentimes easy for humans but difficult for computers. The prototypical example of this is the problem of image recognition

where identifying e.g. whether a dog is in a picture is trivial for a human being, but highly nontrivial for a machine. A priori, it's not obvious what a "dog" is or how to properly translate that information into an algorithm. The solution offered by machine learning is to leverage the massive amount of data being produced in the modern information age in order to essentially teach by example. The advent of ML has brought with it far-reaching societal consequences including but certainly not limited to self-driving cars, facial recognition, better drugs and alloys, and more accurate medical diagnoses. As it grows in sophistication, ML is continually changing the way we interact with machines and the world around us. Given the power of machine learning and its colossal impact, it's unsurprising that its influence and ideas have started to permeate the physics community.

There are two primary flavors of machine learning: supervised ML and unsupervised ML. Supervised machine learning is characterized by the act of "training" an algorithm to recognize and distinguish labeled data in such a way that it may generalize to properly identify unseen examples. It may ultimately be conceptualized as the act of function learning i.e. learning a mapping between inputs and their labels. Examples of supervised algorithms include neural networks, support vector machines (SVMs), decision trees, random forests, and conditional random fields (CRFs). Unsupervised machine learning on the other hand, seeks to learn without labeled data. Typically, unsupervised approaches fall into the class of clustering, manifold learning, or dimensional reduction where instead of learning some input-output relation, the algorithms seek to learn based on the structural properties of the data set. Some examples of unsupervised algorithms include k-means, latent Dirichlet allocation (LDA), Uniform Manifold Approximation and Projection (UMAP),

and Gaussian mixture models (GMM).

In the next section we will focus on particularly powerful type of supervised machine learning: neural networks. Later in dissertation we will discuss and use some unsupervised methods such as Gaussian mixture models (GMM) and Gaussian processes.

3.2 Introduction to Neural Networks

Artificial neural networks are one of the signature algorithms of machine learning and have single-handedly revolutionized the field of artificial intelligence, especially the subfields of speech and image recognition. The operational principles of these computational networks have strong biological ties and were designed to mimic neuron activity in brains. Roughly speaking, given some input, that input propagates through the network similar to signal propagation in systems of neurons via action potentials. That is to say, just as in the case of synapses whereby an incoming signal is passed onto connected neurons if the signal is “strong” enough, a neural network consists of connected layers of nodes that “fire” when the incoming signal exceeds some threshold. A schematic for a single layer feed forward neural network is shown in Fig. 3.1 where it can be visualized as a directed, weighted graph. The circles represent the nodes of the network while the lines represent the weighted connections between the nodes. Algorithmically, signal propagation through this network is implemented as interlaced matrix-vector multiplication and non-linear thresholding operations. The training of a neural network consists of learning the strengths of the connections between nodes that reproduce the

desired mapping between inputs and their labels.

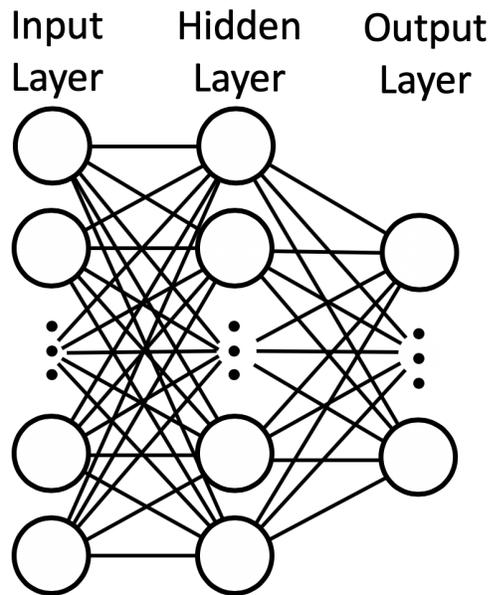


Figure 3.1: Cartoon of single layer feed forward neural network.

Fundamentally, NN's are general function approximators. While deep, convolutional neural networks are often deployed by practitioners, it's important to note that even a simple, single hidden layer feed-forward neural network is sufficient for representing any function. This is the so-called universal approximation theorem for neural networks. Notably, the nonlinearity of the activation function is essential for the representational power of neural networks since otherwise the network is simply a linear mapping between inputs and outputs. However, merely having a framework for representing any function is not useful, what makes neural networks so powerful is that there is also an algorithm, backpropagation, for efficiently learning these functions.

There are two main algorithms for neural networks: (1) the forward

pass algorithm through which an input is mapped to its label and (2) the backward pass algorithm which when combined with a forward pass constitute the backpropagation algorithm through which predicted labels are pushed backwards through the network in order to reinforce the network connections that correctly reproduce the label. Forward passes are used for predicting output labels while backpropagation is the algorithm that “trains” the network. Below, we describe each algorithm for a single hidden layer, fully connected, feed forward neural network. This is in some sense the “simplest” neural network, but as discussed above, even this basic network (with a large enough hidden layer) is capable of representing any function. Neural networks can be used for both classification and regression problems, but below, we will discuss them in the context of classification because we will subsequently go on use them for classifying quantum many-body phases of matter.

Consider the following problem: you are given a set of vectors, $\{\mathbf{x}_i\}_{i=1}^N$ where $\mathbf{x}_i \in \mathbb{R}^d \forall i$, and labels, $\{f(\mathbf{x}_i) = y_i\}_{i=1}^N$, $y_i \in \{1, \dots, K\}$ where K is the number of classes. The goal is learn a representation of the function $f : \mathbb{R}^d \rightarrow \text{Cat}(K)$, where $\text{Cat}(K)$ is the categorical distribution over K classes. Thus we want to learn a mapping over the space of inputs to a probability distribution over the space of classes. Neural networks implement this mapping, via a “forward pass”, as follows: given a single input vector \mathbf{x} with elements x_j , the strength of this signal propagated to the i th node of the 1st hidden layer is given by $z_i = \sum_j W_{ij}x_j + b_i$ where W_{ij} are the network weights of connections between the j th node of the previous layer to the i th node of the next. This incoming signal is then passed through an activation function that either allows the signal to continue to propagate or suppresses the signal. Examples of viable activation functions include the sigmoid function and rectified linear unit (ReLU) as depicted in Fig.

3.2. We proceed with a sigmoid activation function, σ , explicitly defined in Eqn. 3.1. After applying the activation function, the output $a_j = \sigma(z_j)$, is considered to be the new input for the subsequent layer. Eqn. 3.3 contains all the relevant equations for the forward pass where an upper, “layer”, index has been added to the weights, biases, and outputs. The final output is written as \hat{y}_j^L (assuming $L - 1$ hidden layers). Note, the final layer employs a softmax function, defined in Eqn. 3.2, as opposed to a sigmoid in order to ensure that the final outputs sum to unity and may be thought of as some probability vector over classes.

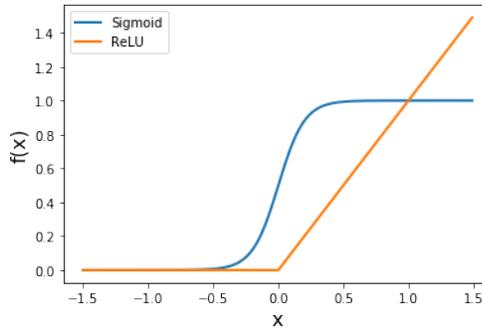


Figure 3.2: Activation functions for neural networks.

Sigmoid function:

$$\sigma(z) = \frac{1}{1 + e^{-z}} \quad (3.1)$$

Softmax function:

$$\sigma(\mathbf{z})_i = \frac{e^{z_i}}{e^{z_1} + \dots + e^{z_K}} \quad (3.2)$$

Forward Pass Equations:

$$\begin{aligned}
a_j^0 &= x_j \\
z_i^l &= \sum_j W_{ij}^l a_j^{l-1} + b_i^l \\
a_j^l &= \sigma(z_j^l) \\
\hat{y}_j^L &= \sigma(\mathbf{z}^L)_j
\end{aligned} \tag{3.3}$$

We now introduce the backward pass algorithm. Its inputs are the output i.e. predicted label of the forward pass and the true label. Central to its operation is a cost function for the neural network which is the measure used to gauge the neural network's performance. Example cost functions include the mean squared error, see Eqn. 3.4, and the cross-entropy a.k.a. log-loss, see Eqn. 3.5. Incorrect predictions contribute to a higher cost so by altering the network parameters to obtain a lower cost, we obtain a "better" neural network. Note that one must be careful with what "better" ultimately means and we will discuss this later in the context of generalization and overfitting. The appropriate network corrections are made by analytically calculating the partial derivatives of the cost function with respect to the network weights and biases and using these to do gradient descent. This is ultimately just an exercise in the application of the chain rule. The relevant equations have been derived and are given in Eqn. 3.6.

Mean squared error cost function:

$$C_{MSE} = -\frac{1}{n} \sum_i |y_i - \hat{y}_i|^2 \tag{3.4}$$

Cross-entropy cost function

$$C_{CE} = -\frac{1}{n} \sum_i [y_i \ln \hat{y}_i + (1 - y_i) \ln(1 - \hat{y}_i)] \tag{3.5}$$

where y_i and \hat{y}_i are respectively the ground truth label and the neural network prediction for the i th training example and n is the number of training examples.

Backward Pass Equations (via stochastic gradient descent):

$$\begin{aligned}
 \delta_j^l &\equiv \frac{\partial C}{\partial z_j^l} \\
 \delta^L &= \hat{y}^L - y \\
 \delta^l &= W^{l+1 T} \delta^{l+1} \odot \sigma'(z^l) \\
 \frac{\partial C}{\partial W_{jk}^l} &= \frac{1}{n} \sum_x \delta_j^l a_k^{l-1} \\
 \frac{\partial C}{\partial b_j^l} &= \frac{1}{n} \sum_x \delta_j^l \\
 W^l &\rightarrow W^l \left(1 - \frac{\eta \lambda}{n}\right) - \frac{\eta}{m} \sum_m \delta^l (a^{l-1})^T \\
 b^l &\rightarrow b^l - \frac{\eta}{m} \sum_m \delta^l
 \end{aligned} \tag{3.6}$$

where n is the training size number, m is the batch size number, η is the gradient descent rate, and λ is the regularization parameter. \odot denotes the Hadamard product i.e. the element-wise matrix product.

Together, a single application of the forward pass followed by a backward pass constitute a single backpropagation iteration. Training a neural network is the act of running backpropagation until convergence. However, as alluded to above, there are several subtleties in training neural networks. For example, we can imagine that a network only “memorizes” the examples it has previously seen and might not “learn” anything. Part of the “magic” of neural networks is that the learned functions actually generalize quite well. The network’s ability to generalize is gauged by dividing the training examples into a training set for use in backpropagation and a test set for checking how well the network performs on unseen examples. Typically a 4:1 ratio is used for splitting the

training and testing sets. One may also use a cross-validation set for optimizing neural network hyperparameters in order to get the “best” network possible.

While the training/test paradigm above allows us to obtain some estimate of the generalization error, it is not a priori obvious how to improve performance on the test set. Training neural networks in such a way as to properly mitigate overfitting to training data is essential for the viability of neural networks in real applications. Some progress may be made by appealing to Occam’s razor and assuming that the “true” function is the simplest one that can properly account for the data, modulo some noise. This is a standard regression practice and a manifestation of the bias-variance tradeoff. The model complexity may be constrained by using various regularization methods. One of the most common ways to implement regularization is to penalize large network weights by adding e.g. the Frobenius norm of the network weight matrices to the cost function, see Eqn. 3.7. This is known as L_2 regularization.

$$L_2 \text{ regularization} \\ C = C_0 + \frac{\lambda}{2n} \sum_{i,j,l} (W_{ij}^l)^2 \quad (3.7)$$

Another regularization technique is dropout in which random neurons are neglected during the training process in order to prevent individual neurons from “co-adapting” and to consequently train a more robust neural network. Conceptually, this may be thought of as inducing an effective model averaging. However, one of the most effective ways to prevent overfitting is to employ stochastic gradient descent during the optimization procedure. Instead of using the exact gradient to train the network, one may instead estimate the gradient by using only a subset of the training data i.e. a batch. This is not only is cheaper

than the exact gradient, but it surprisingly also yields improved performance. The stochastic nature of this process encourages it to find large basins in the cost landscape of the NN parameters as opposed to sharp minima. Intuitively, this makes sense since the optimization no longer caters to the details of a specific training set. To my knowledge, a proper theoretical understanding of neural network generalization capabilities remains an outstanding problem and is an active area of research.

I have written my own neural network implementation in Eigen (a C++ matrix library) for the MBL work below that can be found on my GitHub at https://github.com/jvend/Eigen_neural_network, but the serious practitioner should use either Google's Tensorflow or Facebook's PyTorch. These employ gradient descent optimization algorithms exist such as Adagrad and Adam that modify the learning rates while training for better convergence and use automatic differentiation for fast, automated, and elegant computations of derivatives. They are also powerful frameworks from which complicated, deep networks may be readily built and trained. Note that deep vanilla neural networks suffer from the so-called vanishing/exploding gradient issue where signals being propagated through the network either decay to zero or blow up, thwarting training attempts since information cannot readily be transferred from one side of the network to the other. This phenomenon occurs due to the $\sigma'(z)$ terms in Eqn. 3.6. When there are many layers, the network corrections contain many products of derivatives of the activation function so that the signal vanishes for sigmoid activation function and explodes for ReLU activation functions. However, this has not killed deep learning, rather this issue is circumvented by using convolutional neural networks and elaborate weight initialization schemes.

The network I've written for the MBL work utilizes the cross-entropy cost function, L^2 regularization, and a Gaussian weight/bias initialization as per Michael Nielson's suggestion in his book *Neural Networks and Deep Learning*. Weights were initialized with Gaussian random variables with mean = 0 and std. dev. = $\frac{1}{\sqrt{n_i}}$ where n_i is the number of inputs for the layer. Biases were initialized with Gaussian random variables with mean = 0 and std. dev. = 1. In my experience, a good starting point for the type of phase studies below is to use a single hidden layer with 30 neurons, $\eta = 0.5$, and $\lambda = 0.1$.

3.3 Introduction to Many-body Localization

Many-body localization (MBL) is a non-equilibrium quantum phenomena in which a system fails to "equilibrate under its own dynamics" i.e. it fails to quantum mechanically self-thermalize. More specifically, given some closed system, consider a small subset of the system, call it A , as being embedded so that the larger part serves as a bath or environment, call it B . Then a system is said to self-thermalize if $\rho_A(t) = \rho_A^{(eq)}(T) \forall A$ where ρ_A is the reduced density matrix associated with subset A obtained from tracing out the environmental degrees of freedom and $\rho^{(eq)}(T)$ is the thermal equilibrium according to the canonical ensemble, see Eqn. 3.8.

$$\begin{aligned} \rho_A(t) &= Tr_B \rho(t) \\ \rho^{(eq)}(T) &= \frac{\exp(-H/k_B T)}{Tr[\exp(-H/k_B T)]} \end{aligned} \quad (3.8)$$

That is to say that a system quantum mechanically self-thermalizes if the long-

time behavior of all its subsets, is the same as each subset's respective thermal equilibrium behavior in the context of the full system. Many-body localized systems are characterized by a violation of this behavior. Specifically, they violate the *Eigenstate Thermalization Hypothesis* (ETH) which is the hypothesis that every many-body eigenstate is quantum thermal in sense defined above. Information theoretically, this may be thought of as a failure to diffuse information as the system evolves in time. From a physics perspective this manifests as the slow (logarithmic as opposed to power-law) spreading of quantum entanglement and zero DC conductivity. For this reason, MBLs have piqued interest for applications as robust quantum memories.

Alternatively, MBLs may be thought of as a many-body generalization of Anderson localization. Anderson localization is a single-particle phenomena in which wavefunctions localize upon the introduction of quenched disorder. Heuristically, the disorder causes the wavefunction to interfere with itself and induces an exponential decay away from some focal point determined by the specific details of the disorder configuration. MBLs are similar in the sense that the MBL Hamiltonian defined in terms of e.g. interacting local Pauli spins $\{\sigma_i^{x,y,z}\}$ can be diagonalized via quasi-local unitary operators to be written in terms of "1-bits": $\tau_i^z = \sigma_i^z + \sum_{j,k,\alpha,\beta} f_{ijk}^{ab} \sigma_j^\alpha \sigma_k^\beta + \dots$ with $[\tau_i^z, \tau_j^z] = 0 \forall i, j$, see Eqn. 3.9.

$$H_{l\text{-bit}} = \sum_i h_i \tau_i^z + \sum_{i,j} J_{ij} \tau_i^z \tau_j^z + \sum_{i,j,k} K_{ijk} \tau_i^z \tau_j^z \tau_k^z + \dots \quad (3.9)$$

Since, $[\tau_i^z, H] = 0 \forall i$ and $[\tau_i^z, \tau_j^z] = 0 \forall i, j$, the many-body eigenstates are just the simultaneous eigenstates of these 1-bits. This differs from the Anderson localized case in that there still exists interactions between l - bits, although the interaction strength decays exponentially with separation. Importantly, the

conserved quantities and local integrals of motion corresponding to these l-bits imply an emergent integrability that manifests as area law entanglement for MBLs as opposed to the volume law expected for conventional thermal states. Entanglement entropy and adherence to volume vs area scaling laws is thus one conventional diagnostic for distinguishing MBL and thermal states. More refined statements about the entanglement and energy spectra may be made by appealing to random matrix theory and spectral statistics. Specifically, one expects the level statistics of the MBL and thermal energy spectra to follow a Poisson distribution and Gaussian orthogonal ensemble (GOE) respectively while the entanglement spectra follows a semi-Poisson distribution or GOE.

3.4 Machine learning out-of-equilibrium phases of matter

This section is adapted from a PRL paper [120] with Vedika Khemani and Eun-Ah Kim.

3.4.1 Introduction

The application of machine learning (ML)[65] to central questions in the theory of quantum matter is a rapidly developing research frontier. So far, efforts have been two-fold, focusing on: (1) representing states compactly[20, 29, 28, 77] and (2) identifying and classifying different phases of matter[12, 13, 152, 126, 16, 17, 118, 153, 101, 86]. The driving insight here is that the problems of theoretical interest are essentially those of regression in which an exponentially large amount of data must be condensed into a more accessible or meaningful

form, e.g. the labeling of wavefunctions with phases. As neural networks are universal function approximators and facilitate nonlinear regression, neural network based ML can effectively distill relevant information from complex data while taking it at face value. This is particularly appealing for phases outside the traditional regression scheme where a local order parameter may not be readily available. Such phases include topological phases and out-of-equilibrium *eigenstate phases* [55, 89] in the context of many-body localization (MBL) [4, 7, 88, 159, 85, 83, 3]. Although there has been recent progress in using ML for topological phases [152, 153, 86] and MBL phases [101, 118], extracting phase boundaries in these settings has been a challenging frontier. [101, 118] Moreover, the question of whether the same data and architecture can be used to discern multiple phases, especially multiple MBL phases has been unclear.

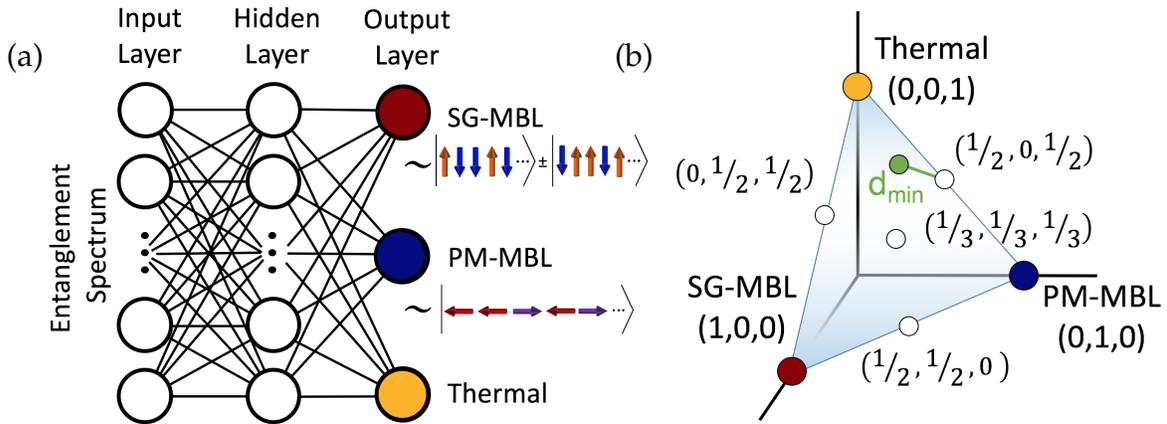


Figure 3.3: (a) A depiction of our neural network. (b) The 2-simplex codomain of our neural network outputs. Colored circles represent regions with maximal confidence; the white circles represent regions of maximal confusion. The green point represents an example output with its associated d_{\min} marked with a green line.

MBL generalizes the phenomenon of Anderson localization to the interacting setting, bringing out the interplay of disorder and interactions. Since MBL systems stay out of thermal equilibrium, they can display a host of rich

dynamical phenomena[54, 104, 56, 119, 67, 46]. Furthermore it is now known that different varieties of MBL phases (e.g. MBL paramagnets, symmetry breaking MBL phases, topologically ordered MBL phases etc.) — each showing different patterns of order in individual highly-excited many-body eigenstates — can be realized in a given system[55, 89, 8, 19, 6]. With experimental realizations of novel out-of-equilibrium states in MBL settings such as time crystals[66, 35, 124, 25, 151], it is all the more important to understand these out-of-equilibrium phases and the associated transitions. Moreover, we need efficient ways to study and discover new MBL phases *without an a priori* knowledge of the defining order-parameters.

Despite extensive research [88, 79, 125, 91, 47, 105, 149, 68, 69], a complete theoretical understanding of the MBL transitions is lacking, partially due to the absence of a comprehensive scheme for regression. Although entanglement entropy serves as a useful diagnostic of thermalization (excited eigenstates in the thermal phase are “volume-law” entangled, while they are only area-law entangled in the MBL phase[88, 8]), it appears to be too aggressive a regression since it traces out important entanglement correlations. The structure of these entanglement correlations is expected to be relevant for understanding the nature of the many-body “resonances” that drive the transition out of the MBL phase[125, 91, 69]. Moreover the entanglement spectra of individual many-body eigenstates must encode the structure of different MBL phases, even when the defining correlation functions are not known *a priori*. While there have been efforts to utilize the full entanglement spectra (ES)[43, 42], a complete understanding for how to interpret the ES has not yet been established. Alternatively there have been efforts to employ neural networks to extract relevant information from entanglement spectra [101, 118], but it has been

unclear whether ML has been able to offer any new insights thus far.

In this work, we take a first step towards a ML assisted study of MBL phase transitions. Our model is a disordered and interacting transverse-field Ising model (TFIM) which has two distinct many-body localized phases: (1) many-body localized spin-glass (MBL-SG) and (2) many-body localized paramagnetic (MBL-PM), in addition to a thermal phase. Using the entanglement spectra of individual eigenstates as our only input to a standard neural network, we are able to locate these phase boundaries with greater precision than standard methods for studying MBL transitions. To do this, we introduce a new geometric approach for interpreting neural network outputs for multipartite classification.

3.4.2 Model

The TFIM in the presence of disorder and interactions is a “canonical model” for studying novel eigenstate phases[55, 89, 71]. It has a well-studied non-interacting limit[39], and well-understood descriptions for the paramagnetic and spin-glass phases in the different limits. An Ising self-dual variant for an L site chain is [81]:

$$H = - \sum_{i=1}^L \left[J_i \sigma_i^z \sigma_{i+1}^z + h_i \sigma_i^x + \lambda (\bar{h} \sigma_i^x \sigma_{i+1}^x + \bar{J} \sigma_i^z \sigma_{i+2}^z) \right], \quad (3.10)$$

where σ_i^z are Pauli spin 1/2 matrices on site i . The couplings, $\{J_i\}$, and onsite fields, $\{h_i\}$, are drawn from log-normal distributions such that the standard deviation of their logarithms is $\delta(\log J) = \delta(\log h) = 1$. Note that this model is equivalent to a disordered interacting fermion model upon a Jordan Wigner transformation, where the interaction strengths are proportional to λ . Finally, \bar{h}

and \bar{J} denote the means of $\{J_i\}$ and $\{h_i\}$ ¹. The means $\overline{\log J}$, $\overline{\log h}$, and λ serve as tuning parameters that can be used to drive different phase transitions.

Let us first consider the axis $\lambda = 0$, which is equivalent to a disordered free-fermion model subject to Anderson localization. From the eigenstate order perspective, there are two distinct phases with respect to the global \mathbb{Z}_2 Ising symmetry of the model under spin flips $P = \prod_i \sigma_i^x$: the symmetry-broken spin-glass (SG) phase for $\bar{J} > \bar{h}$ and the paramagnetic (PM) phase for $\bar{J} < \bar{h}$. Deep in the SG phase, individual many-body eigenstates are macroscopic superpositions (i.e. Schrodinger “cat” states) in the σ^z basis with localized domain walls: $|\alpha\rangle \sim |\uparrow\downarrow\downarrow\uparrow\downarrow \dots\rangle \pm |\downarrow\uparrow\uparrow\downarrow\uparrow \dots\rangle$, and the connected correlation function of σ^z shows long-range order with

$$\langle \alpha | \sigma_i^z \sigma_j^z | \alpha \rangle_c \equiv \langle \alpha | \sigma_i^z \sigma_j^z | \alpha \rangle - \langle \alpha | \sigma_i^z | \alpha \rangle \langle \alpha | \sigma_j^z | \alpha \rangle = \pm |c_\alpha|, \quad (3.11)$$

$|c_\alpha| > 0$ even as $|i - j| \rightarrow \infty$. By contrast, for the equilibrium problem in the absence of disorder, a finite density of delocalized domain walls destroys long-range order at any finite temperature in 1D in accordance with Peierls-Mermin-Wagner theorems. Thus, the SG phase in 1D furnishes an example of a model where localization enables a new form of dynamical order that is disallowed in equilibrium[55, 89]. On the other hand, the eigenstates deep in the PM phase resemble product states in the σ^x basis, $|\alpha\rangle \sim |\rightarrow\leftarrow\leftarrow\rightarrow\leftarrow \dots\rangle$ without LRO, i.e., $\langle \alpha | \sigma_i^z \sigma_j^z | \alpha \rangle_c = 0$. The critical point between these two phases is at the Ising self-dual point, $\overline{\log J} = \overline{\log h}$, and the critical properties for $\lambda = 0$ are described by an infinite randomness fixed point[39].

Once $\lambda \neq 0$, a numerical study over a large number of disorder realizations looking at all the eigenstates is necessary to obtain the phase diagram that now

¹Scaling the interaction terms with \bar{J}, \bar{h} ensures that the interactions are not overwhelmed by the non-interacting terms which can be large on account of the log-normal distributions.

includes the thermal phase. With finite λ , the nature and mechanism of various phase transitions largely remain open questions since the existing theoretical understanding is limited to three extreme regimes in the phase space: (1) $J \gg h, \lambda$, (2) $h \gg J, \lambda$, (3) $\lambda \gg J = h$. In the limits (1) and (2), the Anderson localized SG and PM phases of the non-interacting system generalize to MBL versions of themselves[55, 89, 71]. On the other hand, in the strongly interacting limit, the system will be in a thermal phase with its excited states exhibiting volume law entanglement [87]. Finally, since our interactions were chosen to respect the Ising duality, we expect the phase diagram with non-zero λ to still be symmetric about $\overline{\log J} = \overline{\log h}$ (with small corrections for open boundary conditions). Nevertheless, the precise topology of the tri-partite phase boundary and the existence or absence of a direct MBL-MBL phase transition[27, 81] are hotly debated questions. On the other hand, most existing approaches for detecting phase boundaries rely on the standard deviation of the entanglement entropy (see Fig 3.5) with low resolution leaving the physics of the critical regime largely inaccessible.

3.4.3 Neural Network based Approach

In order to access the information in the entanglement spectra in a holistic manner, we build and employ a feed-forward neural network with a single hidden layer. Our hidden layer contains 200 neurons with sigmoid activation functions. We utilize a cross-entropy cost function with L2 regularization and use a softmax output layer with three neurons, each of which corresponds to one of three possible phases, namely the SG-MBL, the PM-MBL, and the thermal phase (see Fig. 3.3(a)).

We then generate the training and testing data for different disorder configurations of the model (3.10) on an open chain with 12-sites. Specifically, we use exact diagonalization to obtain all the eigenstates and take the middle-quarter of the eigenstates in each Ising symmetry sector to calculate the bipartite entanglement spectra for each eigenstate. The training set consists of three points in the phase space where the phase is known: $\overline{\log J} - \overline{\log h} = \pm 0.8$ with $\lambda = 0.2$ and $\overline{\log J} - \overline{\log h} = 0.0$ with $\lambda = 1.0$. We use 1000 disorder configurations labeled with each of the three points to train our network using a standard error function to an accuracy of over 90%. The fact that successful training could be reached already points to the fact that our network could extract qualitatively distinct information in the entanglement spectra of eigenstates in the three phases of interest. Additionally, in Appendix A, we look at the ability of a neural network to “discover” new phases by only training on the two MBL phases. In this case, although it knows nothing about the thermal phase, it is nevertheless able to recognize that it doesn’t fit within its learned phase paradigm.

Once the training is complete, we feed the entanglement spectra from each point in the phase space of $(\overline{\log J} - \overline{\log h}, \lambda)$ to the network. The neural network output is a triplet that may be thought of as the network’s confidence that the given input is in each phase. Note that all conventional measures require sampling thousands of disorder configurations. On the other hand, we find that averaging the neuron output over just 100 disorder configurations yields a satisfying phase diagram, paving the way for fast scans of large areas of phase space. The purpose of the averaging is to both look into the statistics, as well as to compare with the conventional measure on equal footing. In Fig. 3.4 (a) we plot the average neural network confidence output in the range of $\overline{\log J} - \overline{\log h} \in [-3.0, +3.0]$ and $\lambda \in [0.1, 2.0]$ by representing each component

of the triplet with three colors.

The phase diagram Fig. 3.4 (a) obtained by the neural network displays several satisfying features that are consistent with theoretical insights. First of all, the phase diagram is roughly symmetric about the line $\overline{\log J} - \overline{\log h} = 0.0$ and consistent with the Ising duality of the Hamiltonian Eq.(3.10). Furthermore, the upward curvature of the phase boundary between the MBL phases and the thermal phase is consistent with the fact that the non-interacting model is most delocalized near the SG-PM transition[39] and hence the transition is most susceptible to thermalization upon adding interactions near the $(\overline{\log J} - \overline{\log h} = 0.0, \lambda = 0)$ point. However, it is evident from the representative line cuts in Figs. 3.4 (b-d) that the variation of the confidence outputs is gradual and broad, masking the precise topology of the phase boundaries. Finite-size scaling is presented in Appendix B for $L=8, 10, 12$ showing the sharpening of the transition with increasing system size.

In order to more precisely study the topology of the phase diagram, we developed a protocol for extracting phase boundaries from multi-neuron outputs. This approach will extract the phase boundaries in an unbiased way that allows us to quantify the smoothness and width of the transition, and compare our results to conventional methods. Our approach is geometrically motivated and uses the fact that the neuron outputs sum to unity in a soft-max layer. Specifically, with a soft-max N -neuron output, the codomain of the neural network confidence output is a $(N - 1)$ -simplex embedded in the N -dimensional space of outputs. The points of maximal confusion constitute geometrically notable points on the $(N - 1)$ -simplex; for $N = 3$ these are the mid-points of the edges and the barycenter. Explicitly, in our present

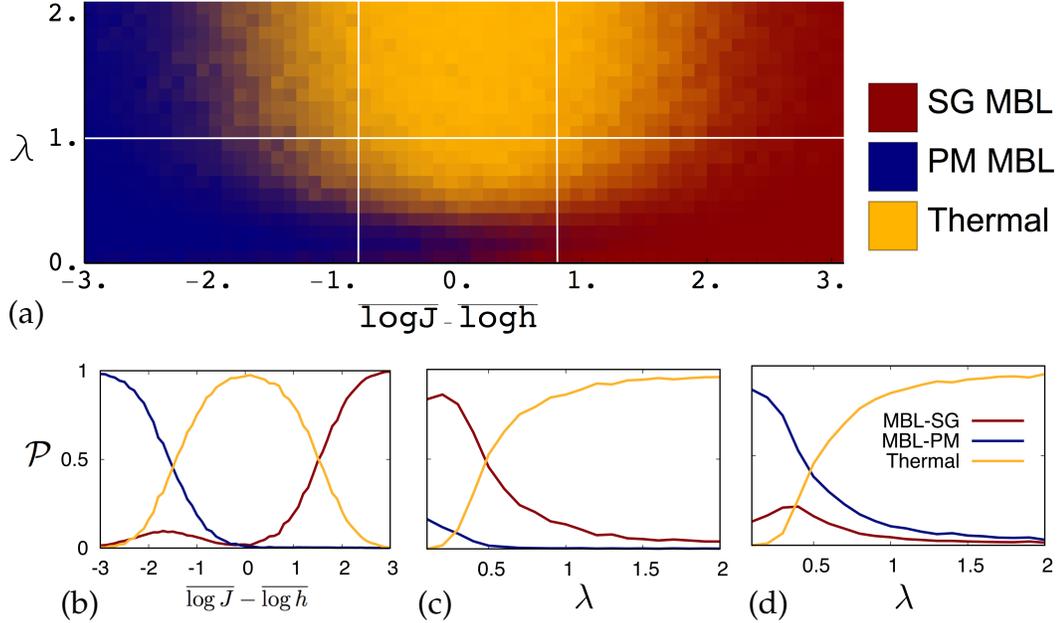


Figure 3.4: (a) The phase diagram with average neural network output plotted as an RGB parameter. Cuts marked by white lines in (a) are shown in (b) $\lambda = 1.0$ (c) $\overline{\log J} - \overline{\log h} = 0.8$, and (d) $\overline{\log J} - \overline{\log h} = -0.8$. The sampling width is 0.1 for each parameter.

case the codomain of our neural network is a 2-simplex and the points of maximum confusion that should naturally belong to the phase boundary[17] are $(1/2, 1/2, 0)$, $(1/2, 0, 1/2)$, $(0, 1/2, 1/2)$, and $(1/3, 1/3, 1/3)$ (see Fig. 3.3(b)). Now for any confidence triplet, one can measure the minimal distance d_{\min} to the set of maximal confusion points. Once we normalize this distance by the maximal possible distance of any point on the simplex to a point of maximal confusion, we obtain a continuous measure of confusion capable of extracting boundaries: $C \equiv 1 - \bar{d}_{\min}$, where \bar{d}_{\min} denotes the normalized distance. This measure of confusion ranges between $C = 1$ when the confidence corresponds to one of the maximal confusion points, and $C = 0$ when the network outputs a particular phase with 100% confidence.

Now at each point in the phase space, we take the average confidence triplet

to evaluate the confusion measure C as shown in Fig. 3.5(a). It is notable that our confusion measure allows us to establish phase boundaries in a manner that is native to the neural network approach. Surprisingly, the phase boundary detected by neural network has the topology of a “wishbone” with a visible phase boundary between two MBL phases (see Fig. 3.5(a)) at small λ . This warrants a more exhaustive study of this transition, including finite-size effects in order to probe the existence of a direct SG-MBL to PM-MBL transition.

The C -measure based extraction of the phase boundary can be contrasted with a more conventional entanglement entropy based approach[71]. Since the EE changes from area law to volume law upon transition from a MBL phase to a thermal phase, it is expected that the standard deviation of the EE in eigenstates peaks at the phase boundary[71]. Fig. 3.5(b) shows the standard deviation taken over all disorder samples and the middle quarter of the eigenstates from each sample. As expected, the standard deviation of the EE is peaked at the MBL-thermal boundaries and is aligned with our machine learning derived measure. However, two advantages of the neural-network C measure easily stand out. First, the EE-based approach cannot distinguish the boundary between the two area-law MBL phases (see the U -shaped phase boundary in Fig. 3.5(b)) whereas the neural network is successfully differentiating these (see the “wishbone” shaped phase boundary in Fig. 3.5(a)). For the MBL-SG problem, one can additionally construct an Edwards-Anderson spin-glass order parameter to single out the MBL-SG phase[71]. However, the ability of the neural network to distinguish between different MBL phases using just the ES and no other “prior knowledge” about order-parameters can prove useful for future studies of new MBL phases where order-parameters might be unknown. Second, the C measure reveals a markedly sharper phase boundary that enables a better study

of its topology (see the line cut comparisons in Fig. 3.5(c,d,e)).

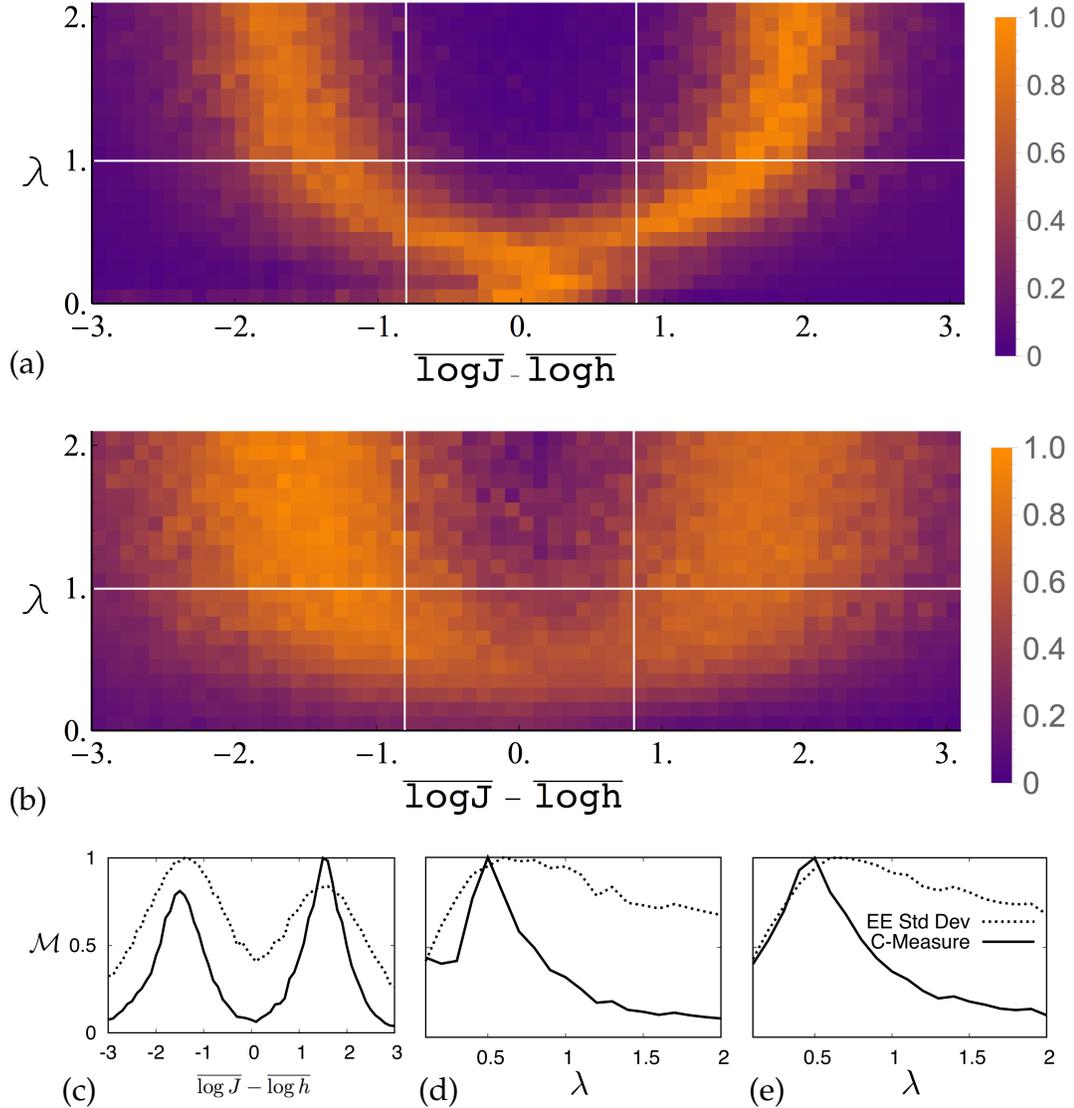


Figure 3.5: (a) Our C-measure for extracting phase boundaries (defined in the main text) and (b) the average standard deviation of the entanglement entropy. The data in each has been normalized by the largest value in the parameter space for meaningful comparison. (c-e) The measures plotted in (a-b) along the cuts marked in white lines: (c) $\lambda = 1.0$, (d) $\overline{\log J} - \overline{\log h} = +0.8$, and (e) $\overline{\log J} - \overline{\log h} = -0.8$.

Finally, we should remark on the neural networks' ability to see through the noise that is inevitable in studies of disorder effects. Although we have averaged over 100 different disorder configurations to gain statistics in the bulk

of this letter, Fig. 3.6 shows that the neural network can capture the coarse features of the phase diagram even for a single disorder realization. The fact that the neural network has gained a regression scheme alternate to the manual modelling of statistical distributions over disorder realizations implies that one can use it as a tool to quickly explore large areas of phase space to map out new non-equilibrium phase diagrams.

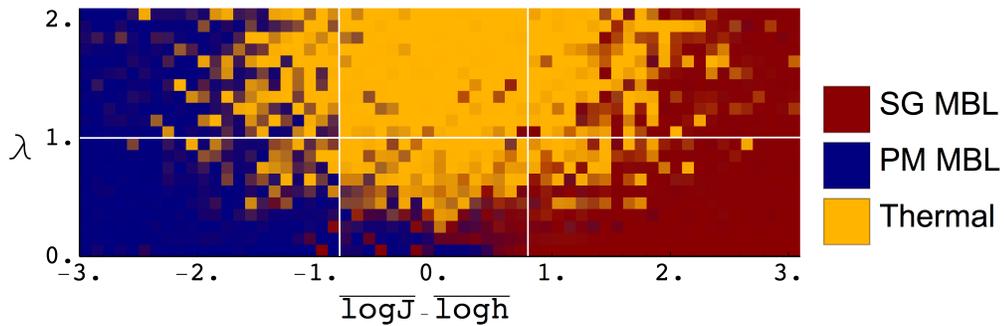


Figure 3.6: The 2D phase diagram where the fully trained network has been tested on a single disorder realization.

3.4.4 Summary and Outlook

Here we exploit the ability of neural networks to distill characteristic features from noisy data in order to extract information from the entanglement spectra associated with out-of-equilibrium phases. To this end, we built a neural network and employed it to process entanglement spectra from a transverse field Ising model with disorder, a poster-child model system that can be in one of three distinct out-of-equilibrium phases. Our neural network, being trained with typical data associated with three limiting points in the phase space, was able to output a phase diagram that is consistent with theoretical expectations. Moreover, using a simplicial geometry construction to quantify

network's degree of confusion, we were able to extract the phase boundary with significantly sharper resolution compared to entanglement entropy-based approaches. Any effort to better understand this transition and/or the possibility of an intervening sliver of thermal phase between the two MBL phases will benefit from a method for obtaining a sharper determination of phase boundaries, which our work provides.

The significance of what we have achieved is multi-faceted. First, we have demonstrated that a standard neural network based approach can give us a sharper look at the multi-partite phase boundary by using the geometric measure of confusion C that we introduced. This is the first example, to the best of our knowledge, that a neural network based approach in a standard setup outperformed conventional approach in terms of sharper phase boundaries². Our work paves the way for future studies on the nature of MBL phase transitions. Second, by having multiple neuron outputs, we were able to obtain tripartite phase diagram involving two distinct MBL phases with a *single* measurement. This is valuable even for MBL phases where there are known order-parameters [71] as in the model we considered. However, this multi-neuron output approach will be even more valuable when dealing with new out-of-equilibrium phases without *a priori* knowledge of suitable order parameters.

²We note that Ref. [101] found that a neural network can outperform conventional methods, after adding a non-trivial input that penalizes lack of confidence to the standard setup in the training process.

3.4.5 Acknowledgements:

E-AK and JV thank Yi Zhang for discussions. E-AK acknowledges the Simons Fellow in Theoretical Physics Award #392182 and DOE support under Award DE-SC0010313. E-AK is grateful to the hospitality of the Kavli Institute of Theoretical Physics supported by NSF under Grant No. NSF PHY-1125915, where this work was initiated. JV acknowledges NSF support under Award NSF DMR-1308089. VK thanks S. Moudgalya and D. Huse for an ongoing collaboration on the model studied in this paper. VK is supported by the Harvard Society of Fellows and the William F. Milton Fund.

3.4.6 Appendix: Neural Network Sensitivity to Unknown Phases

We address the ability of our neural network to recognize phases that it hasn't seen before. Although the network only has a knowledge of whatever we teach it, it should nevertheless be capable of recognizing if a new phase does not fit in its current phase paradigm. This should manifest in two ways: (1) a discontinuity in the NN confidence along the transition point into the unknown phase and (2) an extended region of confusion in the presence of the unknown phase. We test this claim by training our network exclusively on the different MBL phases so that the network has no knowledge of the thermal region. A plot analogous to that in the main text is presented below in Fig. 3.7(a) for this new, partially blind network. Note that we have used yellow instead of red to represent the SG-MBL state since it is easier to see the unknown phase with this color. Here we see that while the MBL parts of the diagram look identical to

those in the original plot, the thermal region is now characterized by a yellow that is darker than that characterizing the SG-MBL. As plotted, this darker color signifies a decrease in confidence relative to that in the MBL phase. In other words, the network is saying that while the thermal region might look more SG-MBL than PM-MBL, it is decidedly not a SG-MBL.

We may now extract the aforementioned discontinuity demarcating the thermal region by computing the gradient of the SG-MBL confidence output. This gradient may be normed by the max value and plotted as an intensity plot as seen in Fig. 3.7(b). Here, there is a clear black line right where we know that the boundary between the thermal and MBL phases is. Thus, although our network knows nothing about the thermal phase, it is nevertheless able to recognize that this phase is not one of the MBL phases it has not been trained on. To our knowledge the ability of partially blind networks to recognize unknown phases has not been previously explored in literature due to the lack of studies on multipartite phase transitions. The phase classification problem in the literature has mostly boiled down to the binary question of “is this phase the same as that phase?”. It is remarkable that the network trained with “this” and “that” possibility recognized a non-binary answer of “none-of-the-above” and hence discovered the thermal phase that it had not been trained on.

Finally, we note that we were unable to produce similar results by training on only one of the MBL phases and the thermal phase. This however is to be expected because the network simply recognizes the MBL vs thermal paradigm and the MBL phase not trained on nicely fits in this paradigm. Thus although the network was unable to “discover” that the phase was a symmetry distinct MBL, it did indeed recognize that it was MBL.

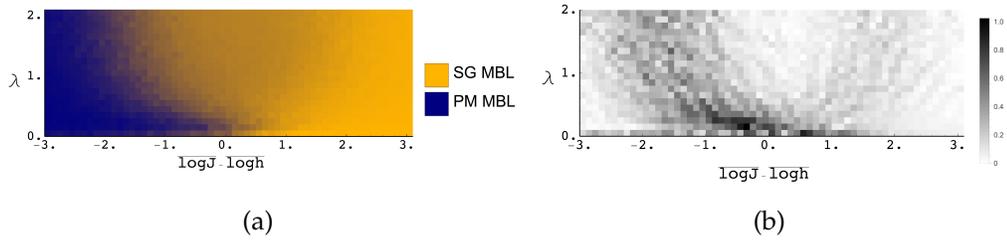


Figure 3.7: (a) The phase diagram generated by a partially blind neural network trained only on the two MBL phases. The average neural network output plotted as an RGB parameter. (b) The intensity of the normed gradient of the SG-MBL confidence is plotted.

3.4.7 Appendix: Finite Size Scaling

We perform finite size scaling and study bipartite entanglement energies for $L=8,10,12$. We consider a comparison for a fixed number of disorder configurations between each system size, specifically $N = 100$ disorder configurations as done in the main text, see Fig. 3.8 and we see a drastic smoothing of the phase diagram as one increases system size and a tightening of the transition region as one would expect. This ability of the network to diagnose MBL-to-thermal transitions with far greater precision than conventional measures is quite valuable given the large numbers of open questions about the MBL transition since it suggests our network is able to “glean” several features about the transition beyond the best accepted diagnostics of this transition. We plan to investigate in a future work what the network is learning with the hope that this will teach us more about the fundamental mechanisms driving the MBL transitions - a major open question. This should also help with understanding the nature of MBL-MBL transitions, another open question.

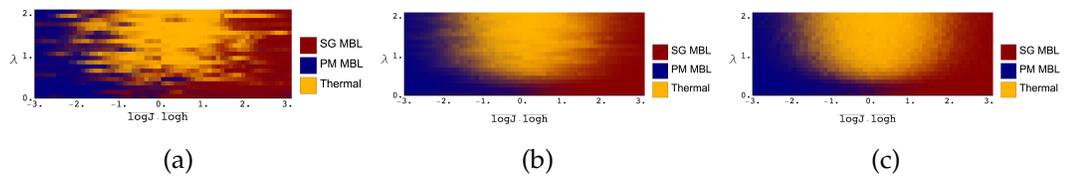


Figure 3.8: Finite size scaling of the phase diagram plots for (a) $L=8$, (b) $L=10$, and (c) $L=12$. All plots use 100 disorder configurations at all points in phase space.

CHAPTER 4
MACHINE LEARNING FOR SINGLE CRYSTAL X-RAY DIFFRACTION
DATA

Data analysis is becoming an increasingly prominent bottleneck for many experimental fronts of quantum matter research. In particular, advancements in detector capabilities for X-ray and neutron scattering have enabled researchers to collect hundreds of GB to several TB of data in the span of a few hours. In this chapter, we present a novel unsupervised machine learning approach for accelerating the analysis of temperature dependent single crystal X-ray diffraction data. Our method employs a Gaussian mixture model to cluster over the temperature dependence of scattering intensities and readily identify phase transitions. A sophisticated thresholding method is developed to leverage sparsity for speed and correlations can be included by implementing label diffusion between nearby momenta and between Brillouin zones. Our algorithm is capable of analyzing hundreds of GBs of data in the span of minutes, offering the tantalizing possibility of real time analysis. Applications to several materials are discussed. This work is part of a large collaboration with experimentalists at Argonne national lab: Ray Osborn and Matthew Krogstad, experimentalists at Cornell: Jacob Ruff, computer scientists at Cornell: Kilian Weinberger, Andrew Gordon Wilson, Geoff Pleiss, and Varsha Kishore, and theoretical physicists in our group at Cornell Eun-Ah Kim and Michael Matty. At the time of the writing of this thesis, this work is unpublished and all the

results below constitute preliminary results.

4.1 Bayesian Inference and Generative Models

Machine learning is ultimately a means for learning functions from data (supervised learning) or extracting structure from data (unsupervised learning). Since any model applied to real data is inherently some approximation of reality, probability theory and statistics are powerful tools that enable us to learn such models in the face of uncertainty. Because they provide us with the framework to quantify our beliefs, they are the natural language for machine learning.

In the case of supervised learning, we would like to model the conditional probability $p(y|x)$ or joint probability $p(y, x)$ between inputs x and outputs y . Similarly, in the case of unsupervised learning, we are interested in learning $p(x)$. Methods that directly model $p(y, x)$ and $p(x)$ are known as generative models while those that model $p(y|x)$ are referred to as discriminative models. The operational principles of these two different approaches are just as their respective names suggest. Generative approaches build a model capable of “generating” random new examples of data that are consistent with what it has previously seen. Discriminative models can predict or “discriminate” the appropriate label associated with a particular input. Generative methods can also predict labels given some instance of x , but this requires the application of Bayes theorem, Eqn. 4.1.

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)} \quad (4.1)$$

Once we have a model in mind, we may condition the above probabilities

on observed data and optimize the underlying hyperparameters of our model in order to learn from the data. Specifically consider the unsupervised case where we have seen some data $S = \{\mathbf{x}_i\}_{i=1}^N$ and we'd like to approximate $p(x) \forall x \in X$ where X is the set of all possible inputs. S may be thought of as our training set even though it has no labels, since it is unsupervised. The appropriate generalizations to the supervised case should be clear in everything below.

If we have some hypothesis or model for describing our data, Bayesian inference allows us to update our beliefs after seeing some data via Eqn. 4.1. Let θ refer to all model hyperparameters. Then, we can calculate the probability of these hyperparameters after having seen some data by calculating the posterior:

$$p(\theta|S) = \frac{p(S|\theta)p(\theta)}{\int d\theta' p(S|\theta')p(\theta')} \quad (4.2)$$

We can then use this posterior to predict the probability of new data, call it x^* , via the predictive posterior:

$$p(x^*|S) = \int d\theta p(x^*|\theta)p(\theta, S) \quad (4.3)$$

Unfortunately, using the full probability distributions for model hyperparameters in order to perform some kind of “model averaging” can be expensive and it is often convenient to use a single optimized value. Two common methods of optimization include Maximum Likelihood estimation (MLE) and Maximum A Posteriori (MAP) estimation. MLE picks out the theta that maximizes the likelihood. This is often used unwittingly when e.g. fitting gaussians to data since the gaussian mean/variance is explicitly chosen to match that of the sampled data. The log-likelihood, $\log P(S|\theta)$, is commonly

used in place of the likelihood for numerical stability and is theoretically equivalent since the logarithm is monotonic.

$$\theta_{MLE} = \arg \max_{\theta} P(S|\theta) \quad (4.4)$$

MAP differs from MLE in that it maximizes the posterior probability:

$$\begin{aligned} \theta_{MAP} &= \arg \max_{\theta} P(\theta|S) \\ &= \arg \max_{\theta} \frac{P(S|\theta)P(\theta)}{P(S)} = \arg \max_{\theta} P(S|\theta)P(\theta) \end{aligned} \quad (4.5)$$

Note that the only difference between *MAP* and *MLE* is the inclusion of the prior in the former optimization. Consequently, if using a uniform prior, these two estimates are equivalent. Thus *MLE* may be seen as a special instance of *MAP*.

4.2 Gaussian Mixture Model

Now that we have introduced the prerequisite language above, we introduce the Gaussian mixture model that will be the workhorse for our unsupervised learning of x-ray data. A Gaussian mixture model is a generative approach that models the density of training points as having been identically and independently generated by a collection of multivariate gaussians. Specifically, given a training set, $S = \{x_i\}_{i=1}^N$ with $x_i \in \mathbb{R}^m \forall i$, we assume that the likelihood, $p(S|\theta)$ with hyperparameters $\theta = \{\pi, \mu, \Sigma\}$ is given by:

$$\begin{aligned}
p(S|\theta) &= \prod_i p(x_i|\pi, \mu, \Sigma) = \prod_i \sum_{z_i} p(x_i|z_i, \mu, \Sigma)p(z_i|\pi) \\
&= \prod_i \sum_k \mathcal{N}(x_i|\mu_k, \Sigma_k)\pi_k
\end{aligned} \tag{4.6}$$

$$\log p(S|\theta) = \sum_i \log \left[\sum_k \mathcal{N}(x_i|\mu_k, \Sigma_k)\pi_k \right]$$

$$\begin{aligned}
z_i \sim \text{Cat}(\pi) \quad \text{so } p(z_i = k|\pi) = \pi_k \quad \text{with } \sum_k \pi_k = 1 \\
x_i|z_i = k \sim \mathcal{N}(x_i|\mu_k, \Sigma_k)
\end{aligned} \tag{4.7}$$

$$\mathcal{N}(x_i|\mu_k, \Sigma_k) = (2\pi)^{-m/2}(\text{Det}\Sigma)^{-1/2} \exp \left[\frac{1}{2}(x_i - \mu_k)^\dagger \Sigma_k^{-1}(x_i - \mu_k) \right]$$

above, $\mathcal{N}(x_i|\mu_k, \Sigma_k)$ is a multivariate normal distribution defined in Eqn. 4.7, with mean μ_k and covariance Σ_k . Note that the product over i , is a consequence of our i.i.d. assumption and conveniently transforms to a sum for the log-likelihood. The z_i 's are latent (a.k.a hidden) variables drawn from a categorical distribution introduced to facilitate a decomposition over clusters. Unfortunately, while these z_i 's allow for a nice decomposition, the fact that they are unobserved makes the associated optimization problem nontrivial.

In order to solve this optimization problem, we must proceed self-consistently. That is to say we make a good initial guess for the hyperparameters and then use that guess to infer the cluster probabilities of each training point. We can do this by applying Bayes' theorem:

$$\begin{aligned}
p(z_i = k|x_i, \mu, \Sigma, \pi) &= \frac{p(x_i|z_i = k, \mu, \Sigma)p(z = k|\pi)}{\sum_{k'} p(x_i|z_i = k', \mu, \Sigma)p(z = k'|\pi)} \\
&= \frac{\mathcal{N}(x_i|\mu_k, \Sigma_k)\pi_k}{\sum_{k'} \mathcal{N}(x_i|\mu_{k'}, \Sigma_{k'})\pi_{k'}}
\end{aligned} \tag{4.8}$$

These cluster probabilities may now be used to calculate new estimates of the hyperparameters via weighted averages where we have defined $w_i^k \equiv p(z_i = k|x_i, \mu, \Sigma, \pi)$:

$$\begin{aligned}
 \pi_k &= \frac{1}{N} \sum_i w_i^k \\
 \mu_k &= \frac{1}{\sum_i w_i^k} \sum_i w_i^k x_i \\
 \Sigma_k &= \frac{1}{\sum_i w_i^k} \sum_i w_i^k (x_i - \mu_k)(x_i - \mu_k)^\dagger
 \end{aligned} \tag{4.9}$$

This iterative procedure is continued until convergence is reached. Convergence to a local optimum is guaranteed. This is explicitly shown in Appendix A.2 along with a derivation of the results in Eqn 4.9.

4.3 Single Crystal X-ray Diffraction

Before applying our machine learning approach, we introduce some basic x-ray scattering theory. Specifically we will sketch how the diffraction pattern is really just a Fourier transform of the electron density and how this allows us to readily identify some quantum electronic phase transitions e.g. CDW melting.

Single crystal x-ray diffraction is an experimental technique in which a monochromatic beam of x-rays (typically produced at an accelerator) is shined on a material in order to probe its electronic structure. These x-rays weakly scatter off the electronic density of the material and interfere in order to produce a diffraction pattern. This diffraction pattern is really just a Fourier transform

of the electronic density of the material. By using a detector to measure the intensity of the outgoing photons at different locations, experimentalists can consequently deduce the internal structure of the compound being studied. More specifically, assume that we have an incident plane wave, $e^{i\vec{k}_{in}\cdot\vec{r}'}$ scatter off a point r' with electronic density $\rho(r')$ as in Fig. 4.1.

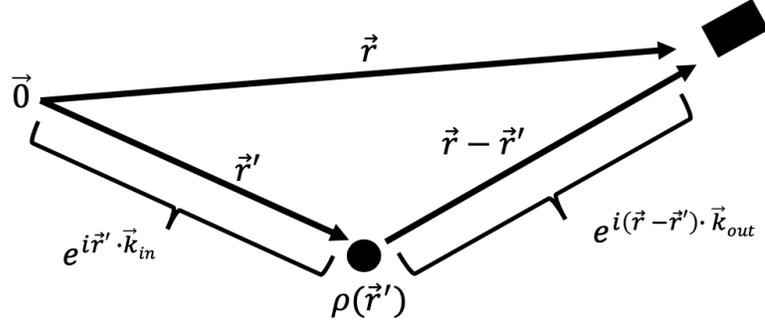


Figure 4.1: Schematic for x-ray scattering of plane wave with wavevector k_{in} scattering off point r' with density $\rho(r')$ and hitting a detector at location r with wavevector k_{out} .

We consider only elastic scattering and work in the weak scattering limit, ignoring multiple scatterings. In this first order approximation (really the Born approximation), we only need to track the phase change acquired upon scattering at each point in space. If our plane wave scatters off r' , with wavevector k' and is measured at detector location r , then the contribution to the outgoing wavefunction is proportional to $\rho(r')e^{ik_{in}\cdot r'}e^{ik_{out}\cdot(r-r')}$. Summing up all scattering contributions, we find:

$$\begin{aligned}\psi_{out}(r) &\propto \int d^3r' \rho(r') e^{ik_{in}\cdot r'} e^{ik_{out}\cdot(r-r')} \\ &\propto \rho(k) e^{ik_{out}\cdot r} \text{ with } k \equiv k_{out} - k_{in} \\ \implies I(k) &= |\psi_{out}(k)|^2 \propto |\rho(k)|^2\end{aligned}\tag{4.10}$$

Note that a more rigorous derivation can be obtained by solving the Schrodinger equation for quantum scattering off a potential and invoking asymptotic limits and the weak scattering/Born approximation.

Now assume that the material under study has some CDW phase transition with wavevector Q . Then as a function of temperature:

$$\begin{aligned}\rho(r, T) &= \rho_0 + \Delta_Q(T) \cos(Q \cdot r) \\ \rho(Q, T) &\propto \Delta_Q(T) \\ I(Q, T) &\propto |\Delta_Q(T)|^2\end{aligned}\tag{4.11}$$

Thus single crystal x-ray scattering provides a ready probe of electronic order parameters. Consequently, if we can extract meaningful, systematic changes in the intensity as a function of temperature we can identify quantum phase transitions. This the goal of our machine learning approach.

4.4 ML for Accelerating Analysis of X-ray Data

4.4.1 Introduction

Thanks to dramatic leaps in x-ray detector technology, it is now possible to collect highly comprehensive x-ray datasets in short amounts of time. As a result, data analysis has become a critical bottleneck in reciprocal space studies of single crystals due to the sheer volume of data available. This is because while improvements in detector technology have enabled the rapid collection of data, a parallel development in data analysis has been lacking. Current

approaches for investigating this data often require highly devoted researchers to manually comb through it looking for evidence of new physics. This is time-intensive and is becoming increasingly infeasible as datasets continue to grow in size. However, machine learning offers a computationally fast solution that is also robust against human error and biases. Such an approach is especially important for large datasets with incommensurability where peak locations are not known a priori or when searching for unknown orderings. By delegating this task to a machine, we may place the rate of data analysis on equal footing with that of data acquisition and ultimately accelerate the pace of x-ray diffraction research.

Here we introduce an unsupervised machine learning approach for identifying phase transitions in single crystal x-ray diffraction data. Our method leverages ideas from time-series clustering in order to classify different points in k-space based on the temperature evolution of their respective intensities, $I_q(T)$. As a first order approach, it is convenient to ignore correlations between momenta and assume these “trajectories” to be independently generated. Such trajectories are directly related to the order parameters of the charge sector so that the onset of order is reflected in the temperature dependence of the intensities. Thus by clustering over them we may identify the presence of different phase transitions, the momenta that drive those phase transitions a.k.a. the ordering vectors, and the characteristic temperature dependence of the orderings. We later extend our approach to account for the appropriate correlations by incorporating label diffusion between correlated momenta on-the-fly in our log-likelihood optimization.

We will benchmark our approach on bulk 1T-TiSe₂. This layered transition

metal dichalcogenide material serves as a good proof of principle since it possesses a well-established CDW transition around 200 K yet remains actively studied today. [30] Current work on the material focuses on its chiral ordering transition, superconductivity with Cu intercalation, and exciton condensation [1, 18, 72]. We then apply our method to the new material $(\text{Ca}_n\text{Sr}_{1-n})_3\text{Rh}_4\text{Sn}_{13}$ in order to discover the composition dependence of its associated structural phase transitions. The remainder of this section proceeds as follows: (1) We introduce the preprocessing required to make our problem tractable, (2) we introduce the clustering algorithm, (3) we benchmark our approach on TiSe_2 , and (4) we apply our approach to $(\text{Ca}_n\text{Sr}_{1-n})_3\text{Rh}_4\text{Sn}_{13}$.

4.4.2 Preprocessing

A signature difficulty in the analysis of X-ray diffraction data is the existence of physics at several different intensity scales. This is only further exacerbated when probing low-intensity features where the signal-to-noise ratio can be small. If one is to employ thresholding as part of some preprocessing, it is imperative to be careful in order to avoid thresholding-out any important physics. Nevertheless, thresholding is extremely useful for mitigating the influence of noise and for reducing dataset size since most single crystal x-ray diffraction patterns are sparse. Consequently, we propose a new thresholding methodology for isolating the physically relevant regions of k-space.

A naive way to cluster the type of datasets offered by single crystal x-ray diffraction is to apply an i.i.d. assumption and directly try to cluster the associated trajectories, $I_q(T)$, so that each q -point is classified according to its

functional temperature dependence. However, such an attempt is immediately thwarted by the existence of a continuum of trajectories spanning over a large temperature range as seen in Fig. 4.3(a) so that getting any meaningful clustering is difficult. The standard way of dealing with this is to use feature scaling a.k.a. standardization in which one removes the mean for each trajectory and then normalizes it by dividing by its standard deviation. However, the dominant features of x-ray diffraction data are usually relatively well-localized peaks and most trajectories may be attributed to background fluctuations and thermal diffuse scattering. These trajectories have small, finite means and variances so that conventional standardization amplifies the underlying experimental error and noise, thereby spoiling any immediate attempt at clustering. On the other hand, failing to standardize makes it difficult to cluster over different energy scales since low-intensity variations can be washed out by larger ones. Thus some cutoff is needed in order to avoid clustering over noise while maintaining the ability to cluster over different energy scales.

In order to properly threshold our data, we exploit the statistical properties of our trajectories' average intensities, $\log \overline{I_{\mathbf{q}}(T)}$. Here, the average is performed over temperature so that a single average intensity is obtained for each \mathbf{q} . Several properties of our data make it advantageous to examine the statistics of $\log \overline{I_{\mathbf{q}}(T)}$ rather than $\overline{I_{\mathbf{q}}(T)}$, most notably its positive semi-definiteness and large range. Since the dominant features our data are naturally sparse and the background trajectories are characterized by possessing small means and variances, we should expect the distribution of $\overline{I_{\mathbf{q}}(T)}$ to be sharply peaked near some relatively small background value. Looking at the logarithm, $\log \overline{I_{\mathbf{q}}(T)}$, broadens this peak allowing us to resolve the finer structural details of this low-intensity background. To first order, we find the distribution of $\log \overline{I_{\mathbf{q}}(T)}$ to

be well-characterized by a bulk background contribution that is approximately normally distributed at low intensities with sparsely distributed high intensity contributions. This can be seen in when looking at a the distribution of $\log \overline{I_q(T)}$ for a single unit-cell of TiSe_2 in Fig. 4.2. In order to separate these high intensity features from rest of the data, we take advantage of their sparsity relative to the background. Specifically, we minimize the Kullback-Leibler divergence D_{KL} :

$$D_{KL}(p(x)||q(x)) = \sum_{x \in X} p(x) \ln \frac{p(x)}{q(x)} \quad (4.12)$$

between the distribution of $\log(\overline{I_q(T)})$ with a high intensity cutoff and a gaussian. Information theoretically, the Kullback-Leibler divergence minimizes the information loss associated with approximating the distribution $p(x)$ by $q(x)$. In this context, the minimization optimally chooses a high-intensity cutoff so that the distribution of the remaining $\log \overline{I_q(T)}$ looks closest to a normal distribution. This is illustrated by applying our procedure to a single unit-cell of TiSe_2 in Fig. 4.2. Optimization is performed via gradient descent. Note that optimizing with this sliding cutoff is necessary and a gaussian cannot be directly fitted because the distribution $\log \overline{I_q(T)}$ is heavy tailed. Directly fitting with a gaussian yields a higher cutoff susceptible to missing important low-intensity features.

After thresholding, we sometimes find it convenient to rescale by dividing by the mean and subtracting one. This is advantageous over subtracting the mean and dividing by the standard deviation because it allows us to implement another thresholding step in which we only cluster over high variance trajectories. In particular, it bolsters the model's ability to cluster over distinct functional behavior because clusters can no longer be smoothly

connected to the origin. However, in the case of the TiSe_2 , we found it sufficient to simply subtract the mean.

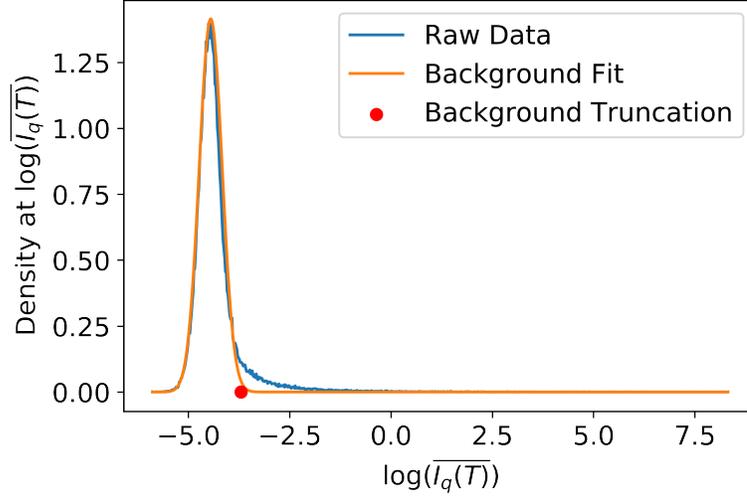


Figure 4.2: Histogram (blue) of $\log \overline{I_q(T)}$ for a single unit-cell of TiSe_2 with background fit (orange) and truncation point described in the main text.

For large volumes of data, the average local intensity scale can change substantially for Brillouin zones far from $\mathbf{q} = 0$ so that any global thresholding method will with over-threshold in some regions and under-threshold in others. This may be circumvented by independently thresholding local regions. However, local thresholding can result in instabilities if one tries to threshold regions lacking sufficient statistics such as zones lying on the boundary of where data was collected. It is also possible, although rare, for the KL-optimization to get stuck at a local maxima that doesn't correspond to the background peak. This is especially problematic when thresholding over tens of thousands of Brillouin zones where the thresholding can't be checked by hand. We are currently in the process of making this thresholding process more sophisticated by using machinery such as Gaussian processes and random forests to infer a continuous thresholding scheme based on the local thresholds. Such a scheme

would be capable of identifying potential mistakes in thresholding and provide appropriate thresholds for boundary regions. This will yield a smoother, more robust thresholding scheme.

Having separated the physically meaningful high-intensity features from the background, we may now again try to cluster. However, in the case of our TiSe_2 data, there is a significant amount of streaking, a detector artifact in which Bragg intensity can be smeared throughout k -space and it is convenient to separate out the Bragg contribution. Since we know precisely the locations of the Bragg peaks, (in our chosen basis they lie at all integer tuples), we may extract these features by examining the connected components of this subset and selecting those connected to the expected Bragg locations. If the Bragg peak locations were not known, we could still use a flood-fill or DBSCAN to find all connected components and the largest connected components would likely be the Bragg peaks. In practice we find this to effectively isolate the Bragg peaks from our high-intensity subset. The remaining, non-Bragg high intensity points can now be probed through clustering. Results for this procedure are shown in Fig. 4.3. Here, Fig. 4.3(a) highlights the primary difficulty preventing a direct clustering of our trajectories in that there is a rough “continuum” of trajectories with no clear clustering. Our preprocessing yields the subset of trajectories in Fig. 4.3(b) which is much more tractable. Fig. 4.3(c) shows the momenta (colored blue) associated with our high-intensity for a single unit-cell and Fig. 4.3(d) shows the subset of those high-intensity trajectories that are not associated with the Bragg peaks (colored orange). Note that the explicit isolation of Bragg peaks is typically not needed with cleaner experimental data where streaking is not as prominent of an issue. Apparently, this dataset was one the first one collected on a new beamline and experimental methods have been refined to significantly

reduce the severity of this streaking. In particular the TiSe_2 streaks had non-negligible experimental errors in the temperature dependence that interfered with clustering.

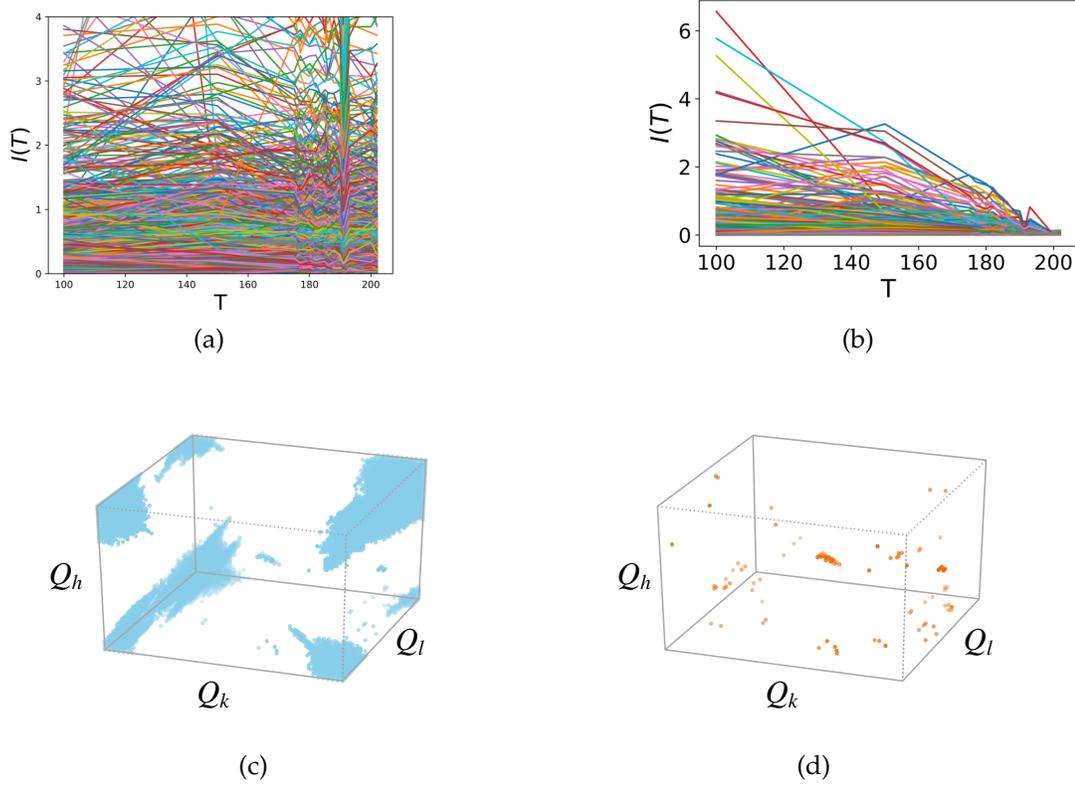


Figure 4.3: (a) All trajectories for a single Brillouin zone (b) post preprocessing trajectories (c) momenta above the threshold for single unit-cell of TiSe_2 have been colored light blue (d) momenta of high-intensity subset not connected to the Bragg locations have been colored orange.

4.4.3 Gaussian Mixture Model Approach

Now that we have properly preprocessed our data we may apply unsupervised machine learning to identify and cluster over the distinct temperature-dependent behaviors of our intensities. As a first order approach, we ignore

correlations between different reciprocal space points and treat the temperature evolution for each \mathbf{q} at the measured temperatures $(T_1, T_2, \dots, T_{d_T})$ as a single point in a d_T -dimensional space. Let us denote such points as $\vec{I}_{\mathbf{q}}$ where the functional temperature dependence has been written as an implicit vector over I . Then our momentum-labeled trajectories, $\vec{I}_{\mathbf{q}}$, are independently distributed in \mathbb{R}^{d_T} . This independence assumption combined with the approximate translational invariance of the lattice lends itself to powerful approaches where parallelization may readily be exploited.

We now adopt a Gaussian mixture model (GMM) approach [82], to cluster $\vec{I}_{\mathbf{q}}$. This model assumes that each point has been independently generated by one out of several high-dimensional multivariate normal distributions. Specifically, we will assume K -number of clusters that are specified by a set of hyperparameters $(\pi, \mu, \Sigma) \equiv \{(\pi_1, \mu_1, \Sigma_1), \dots, (\pi_K, \mu_K, \Sigma_K)\}$ so that $p(\vec{I}_{\mathbf{q}}|\pi, \mu, \Sigma)$ is a weighted sum over K multivariate normal distributions with means μ_k , covariances Σ_k , and mixing weights π_k . When we treat each $\vec{I}_{\mathbf{q}}$ to be independent, the log-likelihood for the entire set of trajectories, $\mathbf{I} \equiv \{\vec{I}_{\mathbf{q}} : \vec{I}_{\mathbf{q}} \text{ in preprocessed subset}\}$ is given by

$$\log p(\mathbf{I}|\mu, \Sigma, \pi) = \sum_{\mathbf{q}} \log \left[\sum_{k=1}^K \pi_k \mathcal{N}(\vec{I}_{\mathbf{q}}; \mu_k, \Sigma_k) \right], \quad (4.13)$$

where $\mathcal{N}(\vec{I}_{\mathbf{q}}; \mu_k, \Sigma_k)$ is a multivariate gaussian conditioned on d^T -dimensional mean μ_k and $d^T \times d^T$ -dimensional covariance Σ_k . Now the unsupervised ML task will be to find the hyperparameters (π, μ, Σ) that maximize the above log-likelihood. This objective can be readily optimized through the expectation maximization (EM) algorithm [94] where it is guaranteed to converge to a stationary point of the likelihood. In practice, we find a diagonal covariance matrix to yield the best performance.

4.4.4 Charge Density wave detection in TiSe_2

Applying our clustering approach to the diffraction data of bulk 1T- TiSe_2 we successfully detect the charge density order as evidenced by Fig. 4.4. Raw intensity data for a slice containing the unit cell in Fig. 4.4(a) at low temperature and Fig. 4.4(b) at high temperature where the CDW has clearly melted. The raw data for a full unit-cell at a single temperature is presented in Fig. 4.4(d) where the intensities have been shown as an opacity. Our algorithm's output is presented in Fig. 4.4(c) and Fig. 4.4(e) for the slice and full unit-cell respectively. In Fig. 4.4(c), we can see that our approach was successfully able to detect and distinguish CDW peaks (red), Bragg (blue), and background (white). Similar results are evident in Fig. 4.4(e) for the full 3D unit-cell where the same colors have been used.

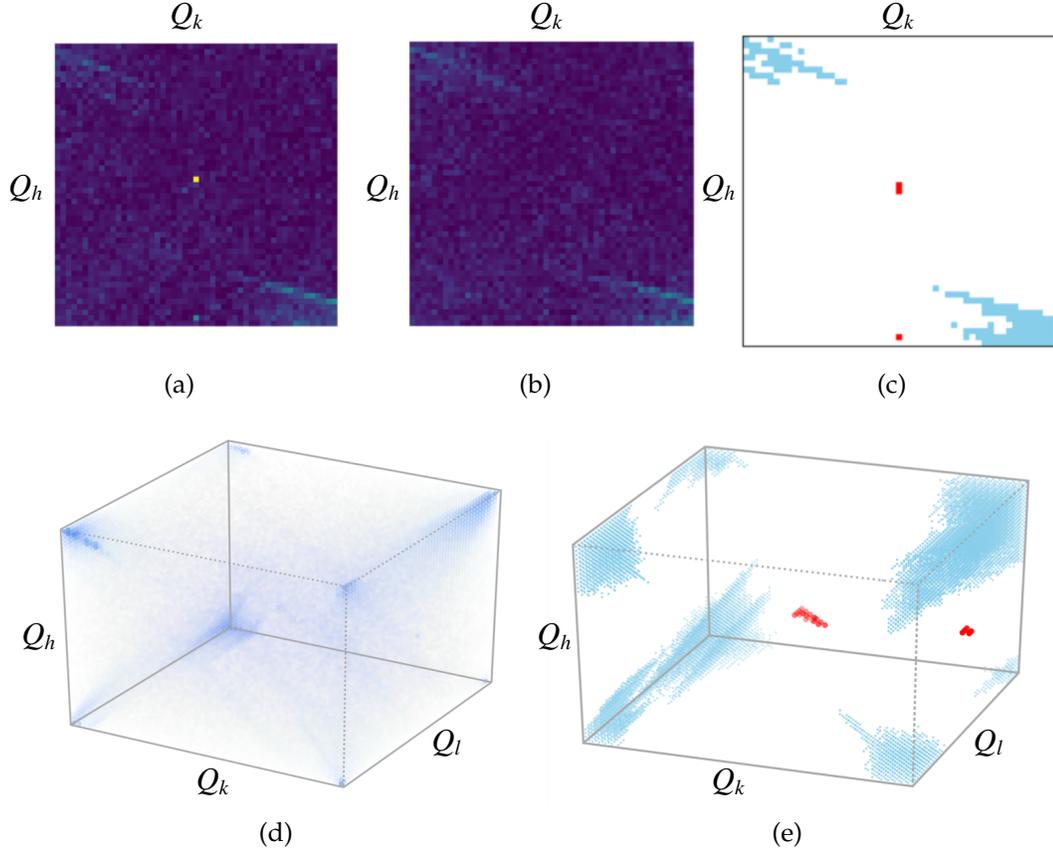


Figure 4.4: (a) Logarithm of intensity at 193 K, $Q_l = -3.5$, $Q_h, Q_k \in [-2.5, -1.5]$ (b) Logarithm of intensity at 200 K, $Q_l = -3.5$, $Q_h, Q_k \in [-2.5, -1.5]$ where the CDW peak has disappeared (c) corresponding color coded classification of momenta with $Q_l = -3.5$, $Q_h, Q_k \in [-2.5, -1.5]$ by our clustering algorithm where Bragg-related momenta have been colored blue and CDW peaks have been colored red (d) logarithm of intensity plotted as opacity for a single 3D Brillouin zone centered at $(Q_h, Q_k, Q_l) = (-2, -2, -3.5)$ at 150K (e) corresponding color coded classification of momenta for a single 3D BZ centered at $(Q_h, Q_k, Q_l) = (-2, -2, -3.5)$ by our clustering algorithm where Bragg-related momenta have been colored blue and CDW peaks have been colored red

The CDW peaks discovered by our approach have ordering vectors $\mathbf{Q}_{CDW} \in \{(\pi, 0, \pi), (\pi, \pi, \pi)\}$, equivalent momenta in our hexagonal basis. The characteristic behavior of their trajectories can be seen in 4.4.4 where we present the clustering results for the non-Bragg high-intensity trajectories. Fig. 4.5(a) shows the cluster assignments while Fig. 4.5(b) shows the means and variances for the

associated clusters. Prior to clustering we have removed (zeroed) the means of all trajectories so that our clustering is not simply thresholding the means. The melting of the CDW order is seen in Fig. 4.5(b) where the CDW cluster mean systematically decreases to the background value around 200 K, the known critical temperature for this order. Although here we only present results for a single unit cell, we can readily apply our method to the entire dataset which occupies around 200 GB. Even with this much data, our algorithm can successfully classify the full dataset in around five minutes.

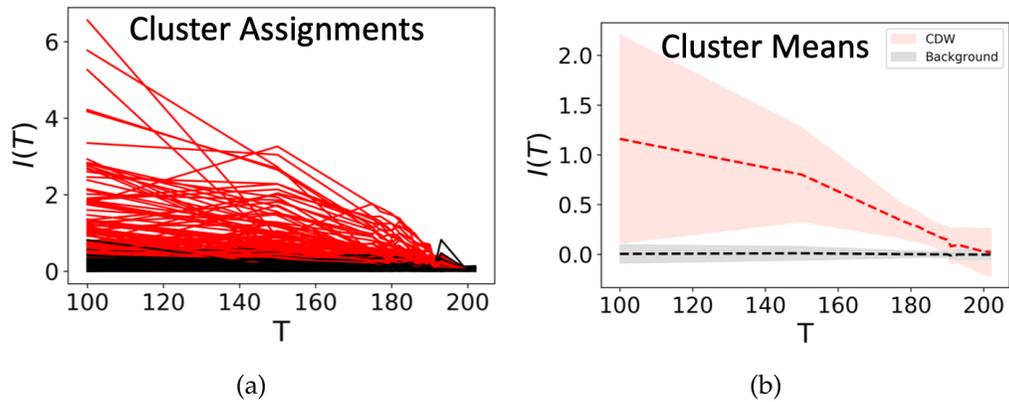


Figure 4.5: Clustering results for a single unit cell of TiSe_2 . (a) Color coded trajectories with red and black corresponding to different clustering labels. (b) Average trajectories associated with the clustering.

4.4.5 Label Smoothing

While the vanilla GMM algorithm used above is able to effectively identify the CDW phase transition, it nevertheless possesses some undesirable properties, most notably it employs an i.i.d. assumption between momenta. Clearly, this assumption is not true since the intensities of nearby momenta are correlated and so are the intensities associated with different unit cells. Exploiting these correlations should result in a more robust algorithm. Here we incorporate

label smoothing as a first order approach for incorporating these correlations by allowing labels to diffuse between neighboring points and between unit-cells. It ultimately results in cleaner, smoother classifications that better align with intuition.

Typical label smoothing is a semi-supervised method in which there exists a ground truth for certain points. These labelings are then “clamped” and diffused through the rest of the system. Here, we lack a bona fide ground truth and so instead incorporate label smoothing dynamically in between the E and M steps of our EM algorithm. Physically, this adds a diffusive “force” to our update scheme that encourages a similar labeling of nearby points and points differing by a reciprocal lattice vector. Convergence in this modified EM method occurs when an equilibrium is reached between this diffusion and the GMM clustering.

Our label smoothing requires us to construct a weighted graph connecting similar momenta in order for diffusion to occur. This may be done by computing the following kernel:

$$K(k, k') = \exp \left[- \sum_i \sin^2 \left(\frac{Q_i}{2} \cdot (k - k') \right) / \ell^2 \right] \quad (4.14)$$

where the Q_i are the reciprocal basis vectors and ℓ is the relevant length scale for the local correlations. The structure of this kernel is readily seen in Fig. 4.6 where $K(k, 0)$ is plotted as an intensity for a 2D slice.

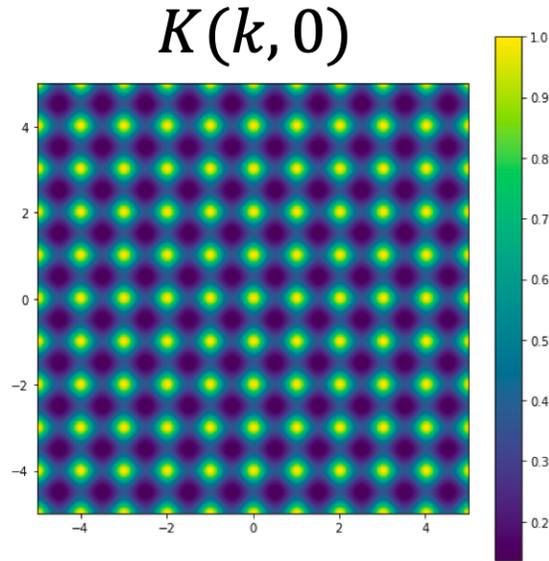


Figure 4.6: Kernel, $K(k, 0)$, showing the similarity between the origin and momenta in a 2D.

This kernel is really just a weighted adjacency matrix. By incorporating a cutoff in the weights, we may exploit the sparsity of our system for fast matrix-vector multiplication. When handling large datasets, this cutoff is essential since the full kernel is too large to be stored in any reasonable amount of RAM. Define A to be the matrix associated with this kernel after having normalized the rows i.e. it is row stochastic so that $\sum_j A_{ij} = 1$. Now define P to be the matrix consisting of cluster probabilities calculated by the E-step. Specifically, let the first index correspond to the different momenta and the second to the cluster probabilities so that P is also row stochastic. Then the product AP is also row stochastic since $\sum_{jk} A_{ij} P_{jk} = \sum_j A_{ij} (1) = 1$. So by multiplying P by A , we generate a new set of diffused cluster probabilities. The strength of this diffusion can be controlled by the number of matrix multiplications. However, note that we cannot simply apply A until $A^n P$ converges, because the largest eigenvector of A is just the constant vector. In practice, we find that even a single application of A between

E- and M-steps is sufficient for obtaining smooth labelings. Results are shown for $(\text{Ca}_n\text{Sr}_{1-n})_3\text{Rh}_4\text{Sn}_{13}$ below.

4.4.6 Application to $(\text{Ca}_n\text{Sr}_{1-n})_3\text{Rh}_4\text{Sn}_{13}$

$(\text{Ca}_n\text{Sr}_{1-n})_3\text{Rh}_4\text{Sn}_{13}$ is a member of the so-called quasi-skutterudites, a class of materials known to exhibit structural phase transitions as a function of temperature and pressure and to superconduct at low temperatures $T_C < 10\text{K}$. $(\text{Ca}_n\text{Sr}_{1-n})_3\text{Rh}_4\text{Sn}_{13}$ is thought to possess an ambient pressure structural quantum critical point at doping $n = 0.9$ as evidenced by electrical resistivity and heat capacity experiments.[44] Since such structural QCPs are rare, there is a large interest studying this material and the role of its quantum fluctuations in the context of superconductivity. [148, 23] Here we are interested in probing the phases of $(\text{Ca}_n\text{Sr}_{1-n})_3\text{Rh}_4\text{Sn}_{13}$ both as a function of temperature and doping.

The existence of a CDW phase transition can be seen from looking at a high and low temperature slices of the experimental single crystal diffraction data for $\text{Sr}_3\text{Rh}_4\text{Sn}_{13}$ with $Q_h \in [-15, 15]$, $Q_k \in [-15, 15]$, $Q_l = 0$ in Fig. 4.11. The associated clustering results are shown in the same figure. Here we can see that the algorithm readily detects the CDW order (colored red) and distinguishes it from the Bragg peaks (colored blue). The extracted CDW mean cluster trajectory clearly displays a nice order parameter in the temperature dependence with $T_{CDW} \sim 132\text{K}$. The results are consistent with what one should expect from transport experiments on $(\text{Sr,Ca})_3\text{Ir}_4\text{Sn}_{13}$. [44]

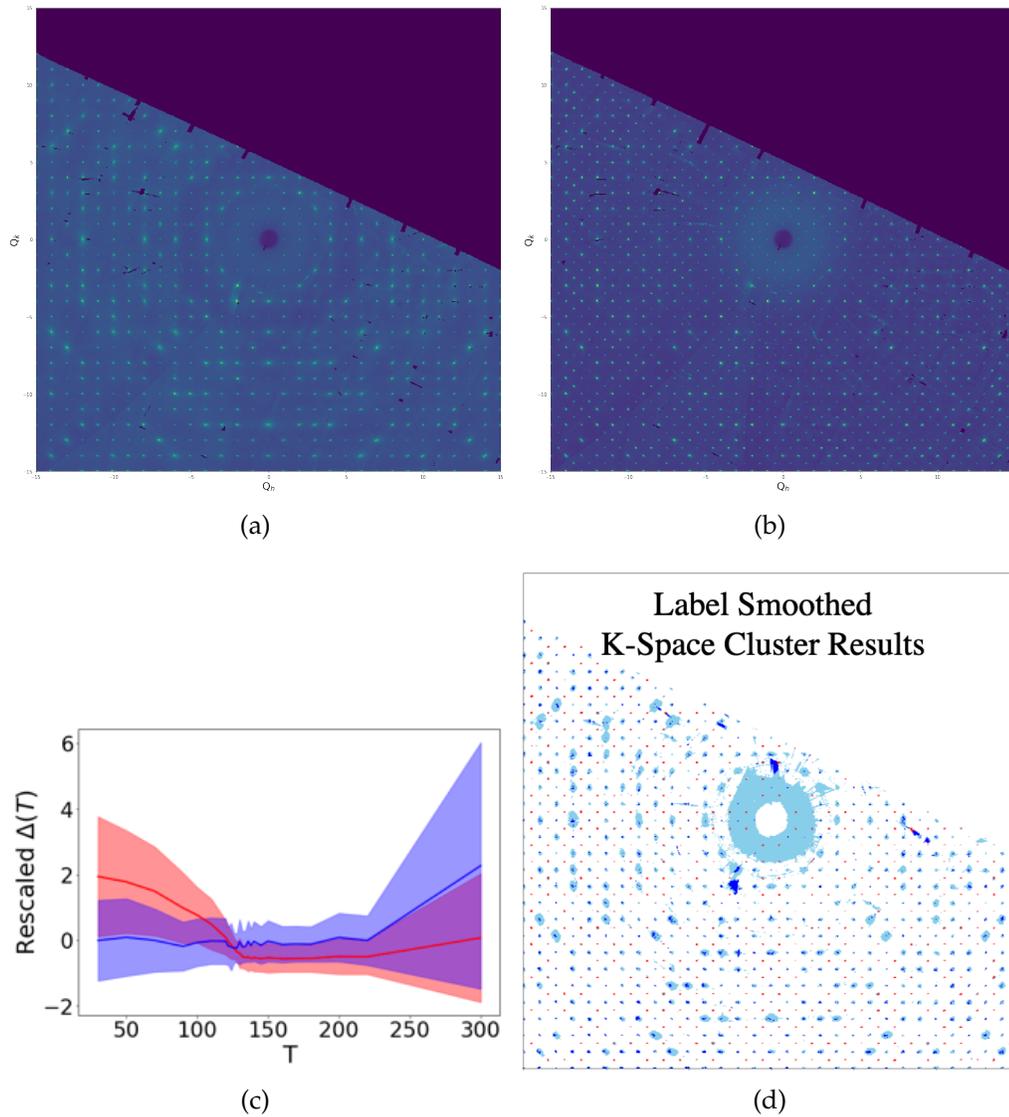


Figure 4.7: Experimental single crystal diffraction data for $\text{Sr}_3\text{Rh}_4\text{Sn}_{13}$ with $Q_h \in [-15, 15]$, $Q_k \in [-15, 15]$, $Q_l = 0$ at (a) $T = 300$ K and (b) $T = 30$ K. (c) Mean cluster trajectories outputted by our clustering algorithm. (d) Color-coded classification of the same slice of k-space shown in (a) and (b) showing the low-variance subset (light blue), Bragg peaks (blue), and CDW peaks (red).

The associated trajectories of this slice are plotted in Fig. 4.8(a). Similar to the case with TiSe_2 the raw trajectories span over a large range and some preprocessing must be performed before clustering. Here we threshold as before and rescale by dividing by the mean and subtracting one. The rescaled

trajectories after preprocessing are shown in Fig. 4.8(b). We further restrict ourselves to those trajectories that have a standard deviation greater than 0.5 after being rescaled. In practice, this helps isolate those momenta with large changes in temperature and helps facilitate the proper identification of phase transitions. Note that even after preprocessing, it's not obvious from just looking at Fig. 4.8(b) that there is any meaningful clustering. The fact that we can readily extract order parameters from these trajectories is a testament to the power of our approach. Our preprocessing scheme is shown in detail in Fig. 4.9 where we have zoomed in on a section of k-space and shown the output of each step of our algorithm.

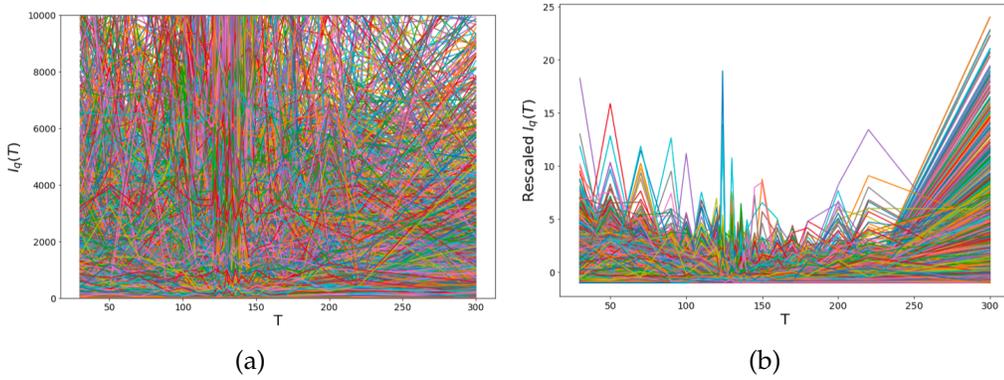


Figure 4.8: (a) Experimental single crystal diffraction data trajectories for $\text{Sr}_3\text{Rh}_4\text{Sn}_{13}$ with $Q_h \in [-15, 15]$, $Q_k \in [-15, 15]$, $Q_l = 0$. (b) Rescaled trajectories after our preprocessing scheme.

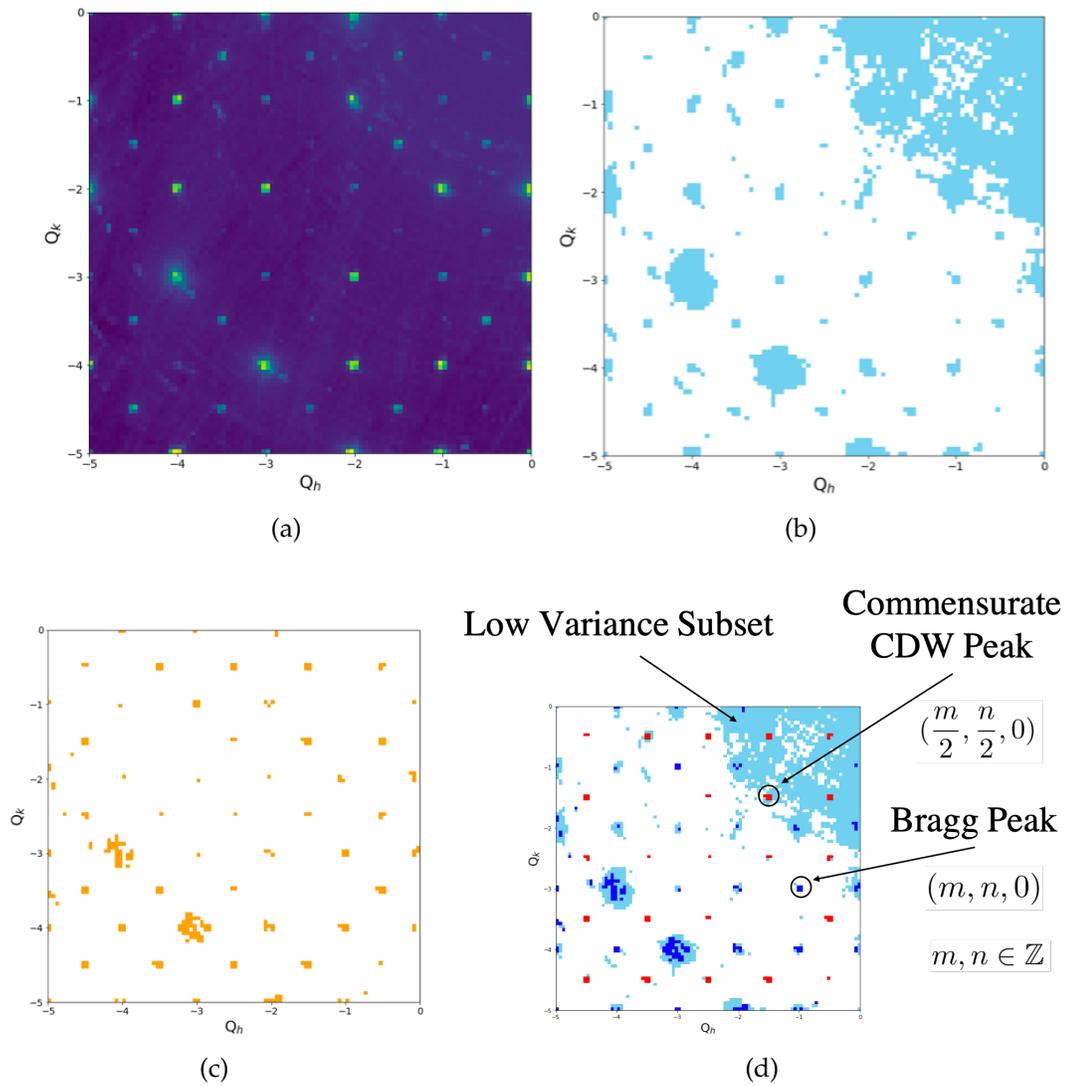


Figure 4.9: (a) Experimental single crystal diffraction data for $\text{Sr}_3\text{Rh}_4\text{Sn}_{13}$ with $Q_h \in [-5, 0]$, $Q_k \in [-5, 0]$, $Q_l = 0$ at $T = 30$ K. (b) Momenta in this slice with average intensities large enough to pass our thresholding. (c) Momenta with trajectories having a high variance. (d) Color-coded, label smoothed clustering results showing the low-variance subset (light blue), Bragg peaks (blue), and CDW peaks (red).

Of course we are not limited to studying single slices and are able to analyze large volumes of data. In Fig. 4.10 we show the classification for a large 3D portion of k -space. The results are consistent with one would expect from the 2D slice results.

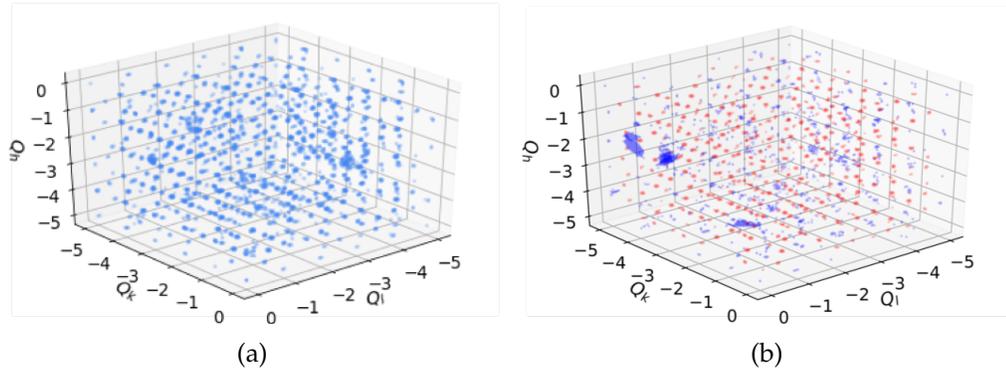


Figure 4.10: (a) Experimental single crystal diffraction data for $\text{Sr}_3\text{Rh}_4\text{Sn}_{13}$ with $Q_h, Q_k, Q_l \in [-5, 0]$ at $T = 30$ K. The logarithm of the intensities is plotted as the opacity (b) Corresponding color-coded classification of momenta showing Bragg peaks (blue) and CDW peaks (red).

We now compare the label unsmoothed and smoothed clustering results for $n = 0$ in Fig. 4.11. In the left, unsmoothed image you can see several sharp peaks with both red and blue labels. Such behavior is clearly a shortcoming of our i.i.d. assumption in this context and is undesirable. Specifically, since the CDW and Bragg peaks correspond to distinct physics we would like this distinction to be fully reflected in the classification, i.e. the momenta associated with a single peak should be similarly classified. For our vanilla GMM, these effects can arise if e.g. the outer parts of the peaks have a smaller variance than the center. Nevertheless, by allowing our labelings to diffuse as described above, these intra-peak misclassifications are nearly entirely eliminated as shown in Fig. 4.11(b). The mean cluster trajectories are also shown as subsets. We find that the inclusion of label smoothing unsurprisingly increases the variance of the mean trajectories but the important qualitative properties remain unchanged.

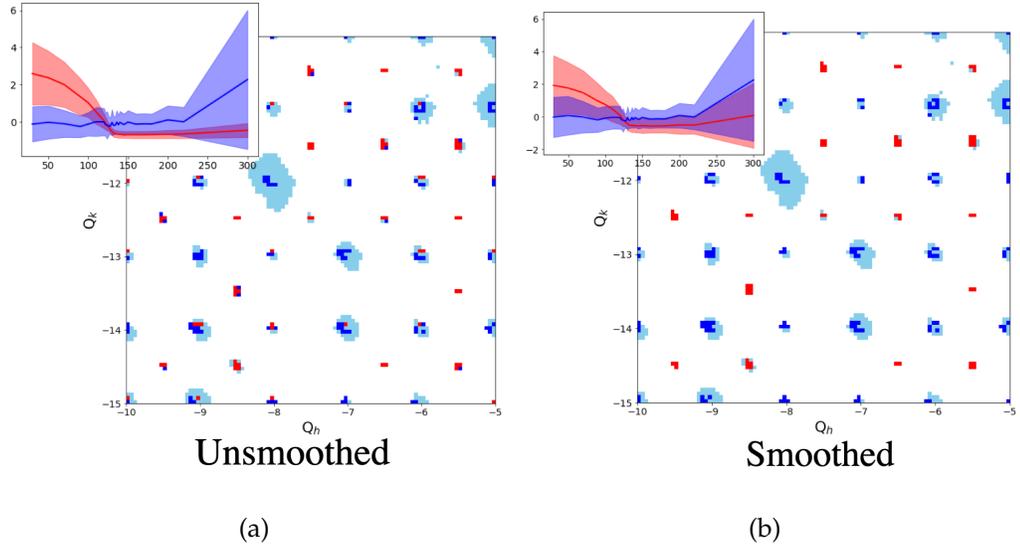


Figure 4.11: Clustering results of $\text{Sr}_3\text{Rh}_4\text{Sn}_{13}$ (a) without and (b) with label smoothing for the k-space slice with $Q_h \in [-5, 0]$, $Q_k \in [-5, 0]$, $Q_l = 0$. Average cluster trajectories for each case are shown in the upper left. The momenta has been color coded according to our classification results. White space corresponds to missing data or points below the thresholding cutoff. Light blue points correspond to the low-variance subset not clustered over. Red and dark blue colors are used for the two clusters of our GMM.

While the above results have all been for $n = 0$ where we found $T_{CDW} \sim 132K$, our algorithm is easily applied to other dopings and we find for $n = 0.1$ $T_{CDW} \sim 102K$, for $n = 0.6$ $T_{CDW} \sim 57K$ and for $n = 0.65$ $T_{CDW} \sim 49K$. This approach allows us to readily extract the CDW order parameters and construct a phase diagram as seen below in Fig. 4.12. This decrease in critical temperature upon increasing doping is consistent with transport measurements. At the time of writing, we don't have enough dopings to reliably extrapolate to zero temperature, but this can be done with more data.

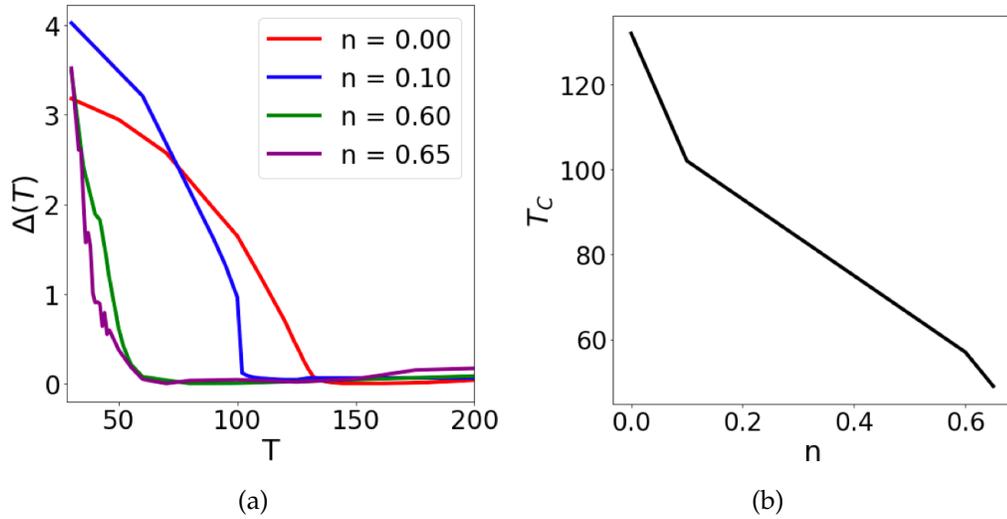


Figure 4.12: (a) Extracted CDW order parameters for $(\text{Ca}_n\text{Sr}_{1-n})_3\text{Rh}_4\text{Sn}_{13}$ with $Q_h \in [-15, 15]$, $Q_k \in [-15, 15]$, $Q_l = 0$ for several dopings. (b) Phase diagram for $(\text{Ca}_n\text{Sr}_{1-n})_3\text{Rh}_4\text{Sn}_{13}$ as a function of doping and critical CDW temperatures.

4.4.7 Conclusions

Although results have only presented for a single Brillouin zone for TiSe_2 and 125 Brillouin zones for $(\text{Ca}_n\text{Sr}_{1-n})_3\text{Rh}_4\text{Sn}_{13}$, our algorithm is readily capable of analyzing the entire datasets due to its speed and high parallelizability. We can apply our approach to study the entire 200 GB dataset of TiSe_2 in around 5 min. These preliminary results show that our GMM approach can efficiently discover quantum phase transitions with dramatic acceleration relative to laborious human searches that can require domain expertise. Importantly, our clustered results are easily interpretable because of the learned mean, μ_k , and variance, Σ_k . Furthermore, the speed of our approach offers the tantalizing potential for real-time applications in which data analysis is performed concurrently with acquisition. This offers the potential for feedback mechanisms in which “important” regions can be detected on-the-fly so that experimentalists can

further refine their search mid-experiment thereby saving valuable beam time.

APPENDIX A
APPENDICES

A.1 The Singular Value Decomposition (SVD) and Properties

The singular value decomposition is a matrix factorization method for matrices defined over the real or complex numbers, see Eqn. A.1.

$$A = U\Sigma V^\dagger \tag{A.1}$$

In Eqn. A.1, if A is an $m \times n$ matrix, U is the $m \times m$ unitary matrix whose columns are the left singular vectors, Σ is a $\min(m, n) \times \min(m, n)$ diagonal matrix whose elements are the singular values of the decomposition, and V is the $n \times n$ unitary matrix whose columns are the right singular vectors. It may be thought of as generalizing the eigendecomposition of square, diagonalizable matrices to non-square, complex matrices. The relation to the eigendecomposition may be derived by comparing the SVD of A to the eigendecomposition of AA^\dagger . Since AA^\dagger is hermitian, it is diagonalizable, and its eigenvectors are the left singular vectors of M . The hermiticity of AA^\dagger also implies that its eigenvalues are real and the singular values of A are precisely the square roots of these eigenvalues.

However, SVD is not just a more general factorization scheme. It is also an extremely powerful means of obtaining optimal low-rank matrix approximations. Specifically, the optimal approximation of a matrix A with rank r by a matrix A' with rank r' is obtained by applying the SVD to A and keeping only the r' largest singular values. Here optimality is taken with respect

to the Frobenius norm i.e. the sum of all matrix elements of squared. This is known as the Eckart–Young theorem.

A.2 Derivation of EM algorithm for GMM and general proof of convergence.

We follow derivations in [84, 82]. First recall Jensen’s inequality: for convex function f and random variable X , $\mathbb{E}[f(X)] \geq f(\mathbb{E}[X])$ where for strictly convex functions, equality holds iff $X = \mathbb{E}[X]$ almost surely. Let $\ell(\theta)$, denote the model log-likelihood and X be our dataset with $x_i \in X$. Then

$$\begin{aligned} \ell(\theta) &= \log p(X; \theta) = \sum_i \log p(x_i; \theta) = \sum_i \log \sum_{z_i} p(x_i, z_i; \theta) \\ &= \sum_i \log \sum_{z_i} q_i(z_i) \frac{p(x_i, z_i; \theta)}{q_i(z_i)} \geq \sum_{i, z_i} q_i(z_i) \log \frac{p(x_i, z_i; \theta)}{q_i(z_i)} \equiv \tilde{\ell}(q, \theta) \end{aligned} \quad (\text{A.2})$$

where $q_i(z_i)$ is some distribution over a random variable z_i (in our case this will be the cluster assignment) s.t. $\sum_{z_i} q_i(z_i) = 1$ and we have used Jensen’s inequality. In order for this bound to be tight, $X = \mathbb{E}[X] \implies q_i(z_i) = p(z_i|x_i; \theta)$. Tightness of this bound implies that improving $\tilde{\ell}(q, \theta)$ necessarily improves $\ell(\theta)$ but since theta is unknown, we will have to make a guess, θ_t , and improve it iteratively. This iterative prescription is known as expectation maximization (EM). It consists of an E-step, where $q_i^t \leftarrow p(z_i|x_i; \theta_t)$ and an M-step $\theta^{t+1} \leftarrow \operatorname{argmax}_{\theta} \tilde{\ell}(q^t, \theta)$.

We now derive the EM algorithm for the GMM. The E-step follows directly

from the model likelihood and Bayes' theorem:

$$w_i^k \equiv p(z_i = k | x_i; \pi_k, \mu_k, \Sigma_k) = \frac{\pi_k \mathcal{N}(x_i | \mu_k, \Sigma_k)}{\sum_k \pi_k \mathcal{N}(x_i | \mu_k, \Sigma_k)} \quad (\text{A.3})$$

$$\mathcal{N}(x_i | \mu_k, \Sigma_k) \equiv \frac{1}{(2\pi)^{n/2}} \frac{1}{\sqrt{\det \Sigma_k}} e^{-\frac{1}{2}(x_i - \mu_k)^\dagger \Sigma_k^{-1} (x_i - \mu_k)}$$

For the M-step, we must find $\{\pi, \mu, \Sigma\}$ that optimizes our lower log-likelihood bound:

$$\tilde{\ell}(\{w_i^k, \pi_k, \mu_k, \Sigma_k\}) = \sum_{i,k} w_i^k \log \left[\frac{\pi_k \mathcal{N}(x_i | \mu_k, \Sigma_k)}{w_i^k} \right] + \lambda (1 - \sum_k \pi_k) \quad (\text{A.4})$$

where λ is a Lagrange multiplier constraining the mixing weights to sum to unity.

Solving for the mixing weights:

$$0 = \partial_{\pi_j} \tilde{\ell} = \sum_{i,k} w_i^k \frac{1}{\pi_k} \delta_{jk} - \lambda \sum_k \delta_{jk} \implies \lambda = \frac{1}{\pi_j} \sum_i w_i^j$$

$$\lambda = \lambda \sum_k \pi_k = \sum_{i,k} w_i^k = \sum_i 1 \equiv m \quad (\text{A.5})$$

$$\implies \pi_j = \frac{1}{m} \sum_i w_i^j$$

Solving for the mean:

$$0 = \partial_{\mu_l} \tilde{\ell} = 2 \sum_i w_i^l \Sigma_l^{-1} (x_i - \mu_l)$$

$$\implies \mu_l = \frac{1}{\sum_i w_i^l} \sum_i w_i^l x_i \quad (\text{A.6})$$

Solving for the covariance is a little trickier. First note the following matrix identities for symmetric invertible matrix A :

$$\begin{aligned}\partial(\log(\det A)) &= \text{Tr}(A^{-1}\partial A) \\ \partial A^{-1} &= -A^{-1}(\partial A)A^{-1}\end{aligned}\tag{A.7}$$

Now, when solving for the covariance we promote the covariance cluster index to an upper index so that the lower indices refer to the matrix elements:

$$\begin{aligned}0 &= \partial_{\Sigma_{mn}^l} \tilde{\ell} = \sum_{i,k} w_i^l \partial_{\Sigma_{mn}^l} \left[\log \det \Sigma^k + (x_i - \mu_k)^\dagger (\Sigma^k)^{-1} (x_i - \mu_k) \right] \\ &= \sum_{i,k} w_i^l \left[\delta_{lk} \text{Tr} \{ (\Sigma^{k-1})_{rs} \delta_{sm} \delta_{in} \} - \delta_{lk} \sum_{ps} (x_i - \mu_k)_p^\dagger \left\{ \sum_{qr} (\Sigma^{k-1})_{pq} \delta_{mq} \delta_{nr} (\Sigma^{k-1})_{rs} \right\} (x_i - \mu_k)_s \right] \\ &= \sum_i w_i^l \left[\Sigma_{nm}^{l-1} - \sum_{p,s} (x - \mu)_p^\dagger \Sigma_{pm}^{l-1} \Sigma_{ns}^{l-1} (x - \mu)_s \right] \\ &= \sum_i w_i^l \left[\Sigma^{l-1} - \Sigma^{l-1} (x_i - \mu_i) (x_i - \mu_i)^\dagger \Sigma^{l-1} \right] \\ 0 &= \sum_i w_i^l \left[\Sigma^l - (x_i - \mu_i) (x_i - \mu_i)^\dagger \right] \\ \implies \Sigma_l &= \frac{1}{\sum_i w_i^l} \sum_i w_i^l (x_i - \mu_i) (x_i - \mu_i)^\dagger\end{aligned}\tag{A.8}$$

Note that all quantities derived about have the same form as one would expect from standard regression but with each data point x_i having a cluster weight w_i^k .

BIBLIOGRAPHY

- [1] Superconductivity in cuprates. *Nature Physics*, 2(8):544–550, 2006.
- [2] D. F. Agterberg and R. P. Kaur. Magnetic-field-induced helical and stripe phases in Rashba superconductors. *Phys. Rev. B*, 75:064511, Feb 2007.
- [3] Ehud Altman and Ronen Vosk. Universal dynamics and renormalization in many-body-localized systems. *Annual Review of Condensed Matter Physics*, 6(1):383–409, 2015.
- [4] P. W. Anderson. Absence of diffusion in certain random lattices. *Phys. Rev.*, 109:1492–1505, Mar 1958.
- [5] P. W. ANDERSON. The resonating valence bond state in La_2CuO_4 and superconductivity. *Science*, 235(4793):1196–1198, 1987.
- [6] Yasaman Bahri, Ronen Vosk, Ehud Altman, and Ashvin Vishwanath. Localization and topology protected quantum coherence at the edge of hot matter. *Nat Commun*, 6, July 2015.
- [7] D. M. Basko, I. L. Aleiner, and B. L. Altshuler. Metal-insulator transition in a weakly interacting many-electron system with localized single-particle states. *Annals of Physics*, 321(5):1126–1205, May 2006.
- [8] Bela Bauer and Chetan Nayak. Area laws in a many-body localized state and its implications for topological order. *Journal of Statistical Mechanics: Theory and Experiment*, 2013(09):P09005, 2013.
- [9] E. Berg, E. Fradkin, E.-A. Kim, S. A. Kivelson, V. Oganesyan, J. M. Tranquada, and S. C. Zhang. Dynamical layer decoupling in a stripe-ordered high- T_c superconductor. *Phys. Rev. Lett.*, 99:127003, Sep 2007.
- [10] Erez Berg, Eduardo Fradkin, and Steven A. Kivelson. Pair-density-wave correlations in the kondo-Heisenberg model. *Phys. Rev. Lett.*, 105:146403, Sep 2010.
- [11] Erez Berg, Eduardo Fradkin, Steven A Kivelson, and John M Tranquada. Striped superconductors: how spin, charge and superconducting orders intertwine in the cuprates. *New Journal of Physics*, 11(11):115004, 2009.

- [12] Peter Broecker, Fakher F. Assaad, and Simon Trebst. Quantum phase recognition via unsupervised machine learning. 07 2017.
- [13] Peter Broecker, Juan Carrasquilla, Roger G. Melko, and Simon Trebst. Machine learning quantum phases of matter beyond the fermion sign problem machine learning quantum phases of matter beyond the fermion sign problem. 08 2016.
- [14] Yuan Cao, Valla Fatemi, Ahmet Demir, Shiang Fang, Spencer L. Tomarken, Jason Y. Luo, Javier D. Sanchez-Yamagishi, Kenji Watanabe, Takashi Taniguchi, Efthimios Kaxiras, Ray C. Ashoori, and Pablo Jarillo-Herrero. Correlated insulator behaviour at half-filling in magic-angle graphene superlattices. *Nature*, 556:80 EP –, 03 2018.
- [15] Yuan Cao, Valla Fatemi, Shiang Fang, Kenji Watanabe, Takashi Taniguchi, Efthimios Kaxiras, and Pablo Jarillo-Herrero. Unconventional superconductivity in magic-angle graphene superlattices. *Nature*, 556:43 EP –, 03 2018.
- [16] Giuseppe Carleo and Matthias Troyer. Solving the quantum many-body problem with artificial neural networks. *Science*, 355(6325):602–606, 2017.
- [17] Juan Carrasquilla and Roger G. Melko. Machine learning phases of matter. *Nat Phys*, 13(5):431–434, 05 2017.
- [18] John-Paul Castellan, Stephan Rosenkranz, Ray Osborn, Qing’an Li, K. E. Gray, X. Luo, U. Welp, Goran Karapetrov, J. P. C. Ruff, and Jasper van Wezel. Chiral phase transition in charge ordered $1t$ - $1t_2$. *Phys. Rev. Lett.*, 110:196404, May 2013.
- [19] Anushya Chandran, Vedika Khemani, C. R. Laumann, and S. L. Sondhi. Many-body localization and symmetry-protected topological order. *Phys. Rev. B*, 89:144201, Apr 2014.
- [20] Jing Chen, Song Cheng, Haidong Xie, Lei Wang, and Tao Xiang. On the equivalence of restricted boltzmann machines and tensor network states. 01 2017.
- [21] Kuang Shing Chen, Zi Yang Meng, Unjong Yu, Shuxiang Yang, Mark Jarrell, and Juana Moreno. Unconventional superconductivity on the triangular lattice hubbard model. *Phys. Rev. B*, 88:041103, Jul 2013.

- [22] Kuang Shing Chen, Zi Yang Meng, Unjong Yu, Shuxiang Yang, Mark Jarrell, and Juana Moreno. Unconventional superconductivity on the triangular lattice hubbard model. *Physical Review B*, 88(4), 2013.
- [23] Y. W. Cheung, Y. J. Hu, M. Imai, Y. Tanioku, H. Kanagawa, J. Murakawa, K. Moriyama, W. Zhang, K. T. Lai, K. Yoshimura, F. M. Grosche, K. Kaneko, S. Tsutsui, and Swee K. Goh. Evidence of a structural quantum critical point in $(\text{ca}_x\text{sr}_{1-x})_3\text{rh}_4\text{sn}_{13}$ from a lattice dynamics study. *Phys. Rev. B*, 98:161103, Oct 2018.
- [24] Gil Young Cho, Jens H. Bardarson, Yuan-Ming Lu, and Joel E. Moore. Superconductivity of doped weyl semimetals: Finite-momentum pairing and electronic analog of the ^3he - a phase. *Phys. Rev. B*, 86:214514, Dec 2012.
- [25] S. Choi, J. Choi, R. Landig, G. Kucsko, H. Zhou, J. Isoya, F. Jelezko, S. Onoda, H. Sumiya, V. Khemani, C. von Keyserlingk, N. Y. Yao, E. Demler, and M. D. Lukin. Observation of discrete time-crystalline order in a disordered dipolar many-body system. *Nature*, 543:221–225, March 2017.
- [26] Zhehao Dai, Ya-Hui Zhang, T. Senthil, and Patrick Lee. Pair density wave, charge density wave and vortex in high T_c cuprates. *ArXiv e-prints*, 2018.
- [27] Wojciech De Roeck and Francois Huveneers. Stability and instability towards delocalization in many-body localization systems. *Phys. Rev. B*, 95:155129, Apr 2017.
- [28] Dong-Ling Deng, Xiaopeng Li, and S. Das Sarma. Exact machine learning topological states. 09 2016.
- [29] Dong-Ling Deng, Xiaopeng Li, and S. Das Sarma. Quantum entanglement in neural network states. *Physical Review X*, 7(2), 2017.
- [30] F. J. Di Salvo, D. E. Moncton, and J. V. Waszczak. Electronic properties and superlattice formation in the semimetal tise_2 . *Phys. Rev. B*, 14:4321–4328, Nov 1976.
- [31] Paul Adrien Maurice Dirac and Ralph Howard Fowler. Quantum mechanics of many-electron systems. *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*, 123(792):714–733, 1929.

- [32] J. F. Dodaro, S. A. Kivelson, Y. Schattner, X. Q. Sun, and C. Wang. Phases of a phenomenological model of twisted bilayer graphene. *Phys. Rev. B*, 98:075154, Aug 2018.
- [33] Stephen D. Edkins, Andrey Kostin, Kazuhiro Fujita, Andrew P. Mackenzie, Hiroshi Eisaki, Shin-Ichi Uchida, Subir Sachdev, Michael J. Lawler, Eun-Ah Kim, J. C. Seamus Davis, and Mohammad H. Hamidian. Magnetic-field Induced Pair Density Wave State in the Cuprate Vortex Halo. *ArXiv e-prints*, 2018.
- [34] Steven A. Kivelson Eduardo Fradkin and John M. Tranquada. Colloquium: Theory of intertwined orders in high temperature superconductors. *Reviews of Modern Physics*, 87(2):457–482, 2015.
- [35] Dominic V. Else, Bela Bauer, and Chetan Nayak. Floquet time crystals. *Phys. Rev. Lett.*, 117:090402, Aug 2016.
- [36] L. Bawden et al. Spin-valley locking in the normal state of a transition-metal dichalcogenide superconductor. *Nat Commun*, 7, 05 2016.
- [37] Shiang Fang, Rodrick Kuate Defo, Sharmila N. Shirodkar, Simon Lieu, Georgios A. Tritsaridis, and Efthimios Kaxiras. Ab initio tight-binding hamiltonian for transition metal dichalcogenides. *Physical Review B*, 92(20), 2015.
- [38] M. Fidrysiak, M. Zegrodnik, and J. Spałek. Unconventional topological superconductivity and phase diagram for an effective two-orbital model as applied to twisted bilayer graphene. *Phys. Rev. B*, 98:085436, Aug 2018.
- [39] Daniel S. Fisher. Critical behavior of random transverse-field ising spin chains. *Phys. Rev. B*, 51:6411–6461, Mar 1995.
- [40] Arno P. Kampf Florian Loder, Siegfried Graser and Thilo Kopp. *Phys. Rev. Lett.*, 107(187001), 2011.
- [41] Peter Fulde and Richard A. Ferrell. Superconductivity in a strong spin-exchange field. *Physical Review*, 135(3A):A550–A563, 1964.
- [42] S. D. Geraedts, N. Regnault, and R. M. Nandkishore. Characterizing the many-body localization transition through the entanglement spectrum. *ArXiv e-prints*, May 2017.

- [43] Scott D. Geraedts, Rahul Nandkishore, and Nicolas Regnault. Many-body localization and thermalization: Insights from the entanglement spectrum. *Phys. Rev. B*, 93:174202, May 2016.
- [44] S. K. Goh, D. A. Tompsett, P. J. Saines, H. C. Chang, T. Matsumoto, M. Imai, K. Yoshimura, and F. M. Grosche. Ambient pressure structural quantum critical point in the phase diagram of $(\text{Ca}_x\text{Sr}_{1-x})_3\text{Rh}_4\text{Sn}_{13}$. *Phys. Rev. Lett.*, 114:097002, Mar 2015.
- [45] Niladri Gomes, W. Wasanthi De Silva, Tirthankar Dutta, R. Torsten Clay, and S. Mazumdar. Coulomb-enhanced superconducting pair correlations and paired-electron liquid in the frustrated quarter-filled band. *Phys. Rev. B*, 93:165110, Apr 2016.
- [46] Sarang Gopalakrishnan, Markus Müller, Vedika Khemani, Michael Knap, Eugene Demler, and David A. Huse. Low-frequency conductivity in many-body localized systems. *Phys. Rev. B*, 92:104202, Sep 2015.
- [47] T. Grover. Certain General Constraints on the Many-Body Localization Transition. *ArXiv e-prints*, May 2014.
- [48] Huaiming Guo, Xingchuan Zhu, Shiping Feng, and Richard T. Scalettar. Pairing symmetry of interacting fermions on a twisted bilayer graphene superlattice. *Phys. Rev. B*, 97:235453, Jun 2018.
- [49] M. H. Hamidian, S. D. Edkins, Sang Hyun Joo, A. Kostin, H. Eisaki, S. Uchida, M. J. Lawler, E. A. Kim, A. P. Mackenzie, K. Fujita, Jinho Lee, and J. C. Séamus Davis. Detection of a cooper-pair density wave in $\text{Bi}_2\text{Sr}_2\text{CaCu}_2\text{O}_{8+x}$. *Nature*, 532(7599):343–347, 04 2016.
- [50] M B Hastings. An area law for one-dimensional quantum systems. *Journal of Statistical Mechanics: Theory and Experiment*, 2007(08):P08024–P08024, aug 2007.
- [51] Yin-Chen He, D. N. Sheng, and Yan Chen. Chiral spin liquid in a frustrated anisotropic kagome heisenberg model. *Physical Review Letters*, 112(13), 2014.
- [52] Yi-Ting Hsu, Abolhassan Vaezi, Mark H. Fischer, and Eun-Ah Kim. Topological superconductivity in monolayer transition metal dichalcogenides. *Nature Communications*, 8:14985 EP –, 04 2017.

- [53] Hui Hu and Xia-Ji Liu. Mean-field phase diagrams of imbalanced fermi gases near a feshbach resonance. *Phys. Rev. A*, 73:051603, May 2006.
- [54] David A. Huse, Rahul Nandkishore, and Vadim Oganesyan. Phenomenology of fully many-body-localized systems. *Phys. Rev. B*, 90:174202, Nov 2014.
- [55] David A. Huse, Rahul Nandkishore, Vadim Oganesyan, Arijeet Pal, and S. L. Sondhi. Localization-protected quantum order. *Phys. Rev. B*, 88:014206, Jul 2013.
- [56] John Z. Imbrie. On many-body localization for quantum spin chains. *Journal of Statistical Physics*, 163(5):998–1048, 2016.
- [57] H. Isobe, N. F. Q. Yuan, and L. Fu. Unconventional Superconductivity and Density Waves in Twisted Bilayer Graphene. *ArXiv e-prints*, May 2018.
- [58] Guillaume Roux Javier Almeida and Didier Poilblanc. Pair density waves in coupled doped two-leg ladders. *Physical Review B*, 82(4), 2010.
- [59] Hong-Chen Jiang, Matthew S. Block, Ryan V. Mishmash, James R. Garrison, D. N. Sheng, Olexei I. Motrunich, and Matthew P. A. Fisher. Non-fermi-liquid d-wave metal phase of strongly interacting electrons. *Nature*, 493(7430):39–44, 01 2013.
- [60] Hong-Chen Jiang and Thomas P. Devereaux. Superconductivity in the hubbard model and its interplay with charge stripes and next-nearest hopping t' . *ArXiv e-prints*.
- [61] Hong-Chen Jiang, Zhenghan Wang, and Leon Balents. Identifying topological order by entanglement entropy. *Nat Phys*, 8(12):902–905, 12 2012.
- [62] Hong-Chen Jiang, Hong Yao, and Leon Balents. Spin liquid ground state of the spin- $\frac{1}{2}$ square J_1 - J_2 heisenberg model. *Phys. Rev. B*, 86:024424, Jul 2012.
- [63] Shenghan Jiang, Andrej Mesaros, and Ying Ran. Chiral spin-density wave, spin-charge-chern liquid, and d+id superconductivity in 1/4-doped correlated electronic systems on the honeycomb lattice. *Physical Review X*, 4(3), 2014.

- [64] Sanghyun Jo, Davide Costanzo, Helmuth Berger, and Alberto F. Morpurgo. Electrostatically induced superconductivity at the surface of ws_2 . *Nano Letters*, 15(2):1197–1202, 2015.
- [65] M.—[nbsp]—I. Jordan. Machine learning: Trends, perspectives, and prospects. *Science*, 349:255–260, 2015.
- [66] V. Khemani, A. Lazarides, R. Moessner, and S. L. Sondhi. Phase Structure of Driven Quantum Systems. *Physical Review Letters*, 116(25):250401, June 2016.
- [67] V. Khemani, R. Nandkishore, and S. L. Sondhi. Nonlocal adiabatic response of a localized system to local manipulations. *Nature Physics*, 11:560–565, July 2015.
- [68] Vedika Khemani, S. P. Lim, D. N. Sheng, and David A. Huse. Critical properties of the many-body localization transition. *Phys. Rev. X*, 7:021013, Apr 2017.
- [69] Vedika Khemani, D. N. Sheng, and David A. Huse. Two universality classes for the many-body localization transition. *Phys. Rev. Lett.*, 119:075702, Aug 2017.
- [70] Philip Kim. Unpublished.
- [71] Jonas A. Kjäll, Jens H. Bardarson, and Frank Pollmann. Many-body localization in a disordered quantum ising chain. *Phys. Rev. Lett.*, 113:107204, Sep 2014.
- [72] Anshul Kogar, Melinda S. Rak, Sean Vig, Ali A. Husain, Felix Flicker, Young Il Joe, Luc Venema, Greg J. MacDougall, Tai C. Chiang, Eduardo Fradkin, Jasper van Wezel, and Peter Abbamonte. Signatures of exciton condensation in a transition metal dichalcogenide. *Science*, 358(6368):1314–1317, 2017.
- [73] Andor Kormányos, Guido Burkard, Martin Gmitra, Jaroslav Fabian, Viktor Zólyomi, Neil D. Drummond, and Vladimir Fal’ko. $k \cdot p$ theory for two-dimensional transition metal dichalcogenide semiconductors. *2D Materials*, 2(022001), 2015.
- [74] A. Larkin and Y. Ovchinnikov. Inhomogeneous state of superconductors. *Sov. Phys. JETP*, 20(762), 1965.

- [75] R. B. LAUGHLIN. The relationship between high-temperature superconductivity and the fractional quantum hall effect. *Science*, 242(4878):525–533, 1988.
- [76] Patrick A. Lee. Amperean pairing and the pseudogap phase of cuprate superconductors. *Phys. Rev. X*, 4:031017, Jul 2014.
- [77] Junwei Liu, Yang Qi, Zi Yang Meng, and Liang Fu. Self-learning monte carlo method. *Physical Review B*, 95(4), 2017.
- [78] J. M. Lu, O. Zheliuk, I. Leermakers, N. F. Q. Yuan, U. Zeitler, K. T. Law, and J. T. Ye. Evidence for two-dimensional ising superconductivity in gated mos2. *Science*, 350(6266):1353–1357, 2015.
- [79] David J. Luitz, Nicolas Laflorencie, and Fabien Alet. Many-body localization edge in the random-field heisenberg chain. *Phys. Rev. B*, 91:081103, Feb 2015.
- [80] Ryan V. Mishmash, Iván González, Roger G. Melko, Olexei I. Motrunich, and Matthew P. A. Fisher. Continuous mott transition between a metal and a quantum spin liquid. *Phys. Rev. B*, 91:235140, Jun 2015.
- [81] Sanjay Moudgalya, Vedika Khemani, and D. A. Huse. *in preparation*.
- [82] Kevin P. Murphy. *Machine learning : a probabilistic perspective*. MIT Press, Cambridge, Mass. [u.a.], 2013.
- [83] Rahul Nandkishore and David A. Huse. Many-body localization and thermalization in quantum statistical mechanics. *Annual Review of Condensed Matter Physics*, 6(1):15–38, 2015.
- [84] Andrew Ng. Cs229 lecture notes. *CS229 Lecture notes*, 2017.
- [85] Vadim Oganesyan and David A. Huse. Localization of interacting fermions at high temperature. *Phys. Rev. B*, 75:155111, Apr 2007.
- [86] Tomoki Ohtsuki and Tomi Ohtsuki. Deep learning the quantum phase transitions in random two-dimensional electron systems. *Journal of the Physical Society of Japan*, 85(12):123706, 2016.
- [87] Don N. Page. Average entropy of a subsystem. *Phys. Rev. Lett.*, 71:1291–1294, Aug 1993.

- [88] Arijeet Pal and David A. Huse. Many-body localization phase transition. *Phys. Rev. B*, 82:174411, Nov 2010.
- [89] 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*, 4:011052, Mar 2014.
- [90] Pierre Pfeuty. The one-dimensional ising model with a transverse field. *Annals of Physics*, 57(1):79 – 90, 1970.
- [91] Andrew C. Potter, Romain Vasseur, and S. A. Parameswaran. Universal properties of many-body delocalization transitions. *Phys. Rev. X*, 5:031033, Sep 2015.
- [92] Marcin Raczkowski, Manuela Capello, Didier Poilblanc, Raymond Frésard, and Andrzej M. Oleś. Unidirectional d -wave superconducting domains in the two-dimensional t - j model. *Phys. Rev. B*, 76:140505, Oct 2007.
- [93] S. Raghu, S. A. Kivelson, and D. J. Scalapino. Superconductivity in the repulsive hubbard model: An asymptotically exact weak-coupling solution. *Phys. Rev. B*, 81:224505, Jun 2010.
- [94] Richard Redner and Homer Walker. Mixture densities, maximum likelihood and the em algorithm. *SIAM Rev.*, 26:195–239, Apr. 1984.
- [95] R. Roldán, E. Cappelluti, and F. Guinea. Interactions and superconductivity in heavily doped mos_2 . *Phys. Rev. B*, 88:054515, Aug 2013.
- [96] R. Roldán, E. Cappelluti, and F. Guinea. Interactions and superconductivity in heavily doped mos_2 . *Phys. Rev. B*, 88:054515, Aug 2013.
- [97] Habib Rostami, Ali G. Moghaddam, and Reza Asgari. Effective lattice hamiltonian for monolayer mos_2 : Tailoring electronic structure with perpendicular electric and magnetic fields. *Physical Review B*, 88(8), 2013.
- [98] Bitan Roy and Vladimir Juricic. Unconventional superconductivity in nearly flat bands in twisted bilayer graphene. *ArXiv e-prints*, page arXiv:1803.11190, March 2018.

- [99] Peyman Sahebsara and David Sénéchal. Hubbard model on the triangular lattice: Spiral order and spin liquid. *Physical Review Letters*, 100(13), 2008.
- [100] Yu Saito, Yasuharu Nakamura, Mohammad Saeed Bahramy, Yoshimitsu Kohama, Jianting Ye, Yuichi Kasahara, Yuji Nakagawa, Masaru Onga, Masashi Tokunaga, Tsutomu Nojima, Youichi Yanase, and Yoshihiro Iwasa. Superconductivity protected by spin-valley locking in ion-gated mos2. *Nat Phys*, 12(2):144–149, February 2016.
- [101] Frank Schindler, Nicolas Regnault, and Titus Neupert. Probing many-body localization with neural networks. *Physical Review B*, 95(24), 2017.
- [102] U. Schollwöck. The density-matrix renormalization group. *Rev. Mod. Phys.*, 77:259–315, Apr 2005.
- [103] Ulrich Schollwöck. The density-matrix renormalization group in the age of matrix product states. *Annals of Physics*, 326(1):96–192, Jan 2011.
- [104] Maksym Serbyn, Z. Papic, and Dmitry A. Abanin. Local conservation laws and the structure of the many-body localized states. *Phys. Rev. Lett.*, 111:127201, Sep 2013.
- [105] Maksym Serbyn, Z. Papic, and Dmitry A. Abanin. Criterion for many-body localization-delocalization phase transition. *Phys. Rev. X*, 5:041047, Dec 2015.
- [106] Wen-Yu Shan, Hai-Zhou Lu, and Di Xiao. Spin hall effect in spin-valley coupled monolayers of transition metal dichalcogenides. *Phys. Rev. B*, 88:125301, Sep 2013.
- [107] Wu Shi, Jianting Ye, Yijin Zhang, Ryuji Suzuki, Masaro Yoshida, Jun Miyazaki, Naoko Inoue, Yu Saito, and Yoshihiro Iwasa. Superconductivity series in transition metal dichalcogenides by ionic gating. *Scientific Reports*, 5:12534 EP –, 08 2015.
- [108] Tomonori Shirakawa, Takami Tohyama, Jure Kokalj, Sigetoshi Sota, and Seiji Yunoki. Ground-state phase diagram of the triangular lattice hubbard model by the density-matrix renormalization group method. *Phys. Rev. B*, 96:205130, Nov 2017.
- [109] Manfred Sgrist and Kazuo Ueda. Phenomenological theory of unconventional superconductivity. *Rev. Mod. Phys.*, 63:239–311, Apr 1991.

- [110] Stefan A. Söffing, Michael Bortz, Imke Schneider, Alexander Struck, Michael Fleischhauer, and Sebastian Eggert. Wigner crystal versus Friedel oscillations in the one-dimensional Hubbard model. *Physical Review B*, 79(19), 2009.
- [111] Evan Sosenko, Junhua Zhang, and Vivek Aji. Unconventional superconductivity and anomalous response in hole-doped transition metal dichalcogenides. *Phys. Rev. B*, 95:144508, Apr 2017.
- [112] Rodrigo Soto-Garrido, Yuxuan Wang, Eduardo Fradkin, and S. Lance Cooper. Higgs modes in the pair density wave superconducting state. *Phys. Rev. B*, 95:214502, Jun 2017.
- [113] E.M. Stoudenmire and Steven R. White. Studying two dimensional systems with the density matrix renormalization group. *Annual Review of Condensed Matter Physics*, 3:111–128, March 2012.
- [114] Miles Studenmire and Steven R. White. <http://itensor.org>.
- [115] Aaron Szasz, Johannes Motruk, Michael P. Zaletel, and Joel E. Moore. Observation of a chiral spin liquid phase of the Hubbard model on the triangular lattice: a density matrix renormalization group study. *ArXiv e-prints*.
- [116] Kazunori Takada, Hiroya Sakurai, Eiji Takayama-Muromachi, Fujio Izumi, Ruben A. Dilanian, and Takayoshi Sasaki. Superconductivity in two-dimensional CoO₂ layers. *Nature*, 422:53 EP –, 03 2003.
- [117] H. Tsuchiura and M. Ogata. Anti-phase superconducting domain structures in the $t-t'-t''-j$ model. *Journal of Physics: Conference Series*, 150(5), 2009.
- [118] Evert P. L. van Nieuwenburg, Ye-Hua Liu, and Sebastian D. Huber. Learning phase transitions by confusion. *Nat Phys*, 13(5):435–439, 05 2017.
- [119] R. Vasseur, S. A. Parameswaran, and J. E. Moore. Quantum revivals and many-body localization. *Phys. Rev. B*, 91:140202, Apr 2015.
- [120] Jordan Venderley, Vedika Khemani, and Eun-Ah Kim. Machine learning out-of-equilibrium phases of matter. *Phys. Rev. Lett.*, 120:257204, Jun 2018.
- [121] Jordan Venderley and Eun-Ah Kim. *ArXiv e-prints*.

- [122] Jordan Venderley and Eun-Ah Kim. Density matrix renormalization group study of superconductivity in the triangular lattice hubbard model. *Phys. Rev. B*, 100:060506, Aug 2019.
- [123] Jordan Venderley and Eun-Ah Kim. Evidence of pair-density wave in spin-valley locked systems. *Science Advances*, 5(3), 2019.
- [124] 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.
- [125] Ronen Vosk, David A. Huse, and Ehud Altman. Theory of the many-body localization transition in one-dimensional systems. *Phys. Rev. X*, 5:031032, Sep 2015.
- [126] Lei Wang. Discovering phase transitions with unsupervised learning. *Physical Review B*, 94(19), 2016.
- [127] Qiang-Hua Wang, Dung-Hai Lee, and Patrick A. Lee. Doped $t - j$ model on a triangular lattice: Possible application to $\text{Na}_x\text{CoO}_2 \cdot y\text{H}_2\text{O}$ and $\text{Na}_{1-x}\text{TiO}_2$. *Phys. Rev. B*, 69:092504, Mar 2004.
- [128] Yuxuan Wang, Stephen D. Edkins, Mohammad H. Hamidian, J. C. Seamus Davis, Eduardo Fradkin, and Steven A. Kivelson. Pair Density Waves in Superconducting Vortex Halos. *ArXiv e-prints*, 2018.
- [129] Tsutomu Watanabe, Hisatoshi Yokoyama, Yukio Tanaka, Jun-ichiro Inoue, and Masao Ogata. Variational monte carlo studies of pairing symmetry for the t - j model on a triangular lattice. *Journal of the Physical Society of Japan*, 73(12):3404–3412, 2004.
- [130] X. G. Wen, Frank Wilczek, and A. Zee. Chiral spin states and superconductivity. *Phys. Rev. B*, 39:11413–11423, Jun 1989.
- [131] Steven R. White. Density matrix formulation for quantum renormalization groups. *Phys. Rev. Lett.*, 69:2863–2866, Nov 1992.
- [132] Steven R. White and A. L. Chernyshev. Néel order in square and triangular lattice heisenberg models. *Physical Review Letters*, 99(12), 2007.
- [133] Steven R. White and D. J. Scalapino. Competition between stripes and pairing in a t - t' - j model. *Physical Review B*, 60(2):R753–R756, 1999.

- [134] Steven R. White and D. J. Scalapino. Pairing on striped t-t'-j lattices. *Physical Review B*, 79(22), June 2009.
- [135] Steven R. White and Douglas J. Scalapino. Density matrix renormalization group study of the striped phase in the 2d t-j model. *Physical Review Letters*, 80(6):1272–1275, 1998.
- [136] Fengcheng Wu, Timothy Lovorn, Emanuel Tutuc, and A. H. MacDonald. Hubbard model physics in transition metal dichalcogenide moiré bands. *Phys. Rev. Lett.*, 121:026402, Jul 2018.
- [137] Fengcheng Wu, Timothy Lovorn, Emanuel Tutuc, Ivar Martin, and A. H. MacDonald. *ArXiv e-prints*.
- [138] Xiaoxiang Xi, Zefang Wang, Weiwei Zhao, Ju-Hyun Park, Kam Tuen Law, Helmuth Berger, Laszlo Forro, Jie Shan, and Kin Fai Mak. Ising pairing in superconducting nbse2 atomic layers. *Nat Phys*, 12(2):139–143, 02 2016.
- [139] Di Xiao, Gui-Bin Liu, Wanxiang Feng, Xiaodong Xu, and Wang Yao. Coupled spin and valley physics in monolayers of mos₂ and other group-vi dichalcogenides. *Phys. Rev. Lett.*, 108:196802, May 2012.
- [140] Cenke Xu and Leon Balents. Topological superconductivity in twisted multilayer graphene. *Phys. Rev. Lett.*, 121:087001, Aug 2018.
- [141] Minoru Yamashita, Norihito Nakata, Yuichi Kasahara, Takahiko Sasaki, Naoki Yoneyama, Norio Kobayashi, Satoshi Fujimoto, Takasada Shibauchi, and Yuji Matsuda. Thermal-transport measurements in a quantum spin-liquid state of the frustrated triangular magnet -(bedt-ttf)₂cu₂(cn)₃. *Nature Physics*, 5:44 EP –, 11 2008.
- [142] Satoshi Yamashita, Yasuhiro Nakazawa, Masaharu Oguni, Yugo Oshima, Hiroyuki Nojiri, Yasuhiro Shimizu, Kazuya Miyagawa, and Kazushi Kanoda. *Nature Physics*, 4:459 EP –, 04 2008.
- [143] +Simeng Yan, +David+A. Huse, and +Steven+R." White. Spin-liquid ground state of the s = 1/2 kagome heisenberg antiferromagnet. *Science*, 332(6034):1173–1176, 2011.
- [144] Kai-Yu Yang, Wei-Qiang Chen, T. M. Rice, M. Sigrist, and Fu-Chun Zhang. Nature of stripes in the generalized t-j model applied to the cuprate superconductors. *New Journal of Physics*, 11(055053), 07 2009.

- [145] M. Yankowitz, S. Chen, H. Polshyn, K. Watanabe, T. Taniguchi, D. Graf, A. F. Young, and C. R. Dean. Tuning superconductivity in twisted bilayer graphene. *ArXiv e-prints*, August 2018.
- [146] Bing Ye, Andrej Mesaros, and Ying Ran. Ferromagnetism and d+id superconductivity in 1/2 doped correlated systems on triangular lattice. *ArXiv e-prints*.
- [147] J. T. Ye, Y. J. Zhang, R. Akashi, M. S. Bahramy, R. Arita, and Y. Iwasa. Superconducting dome in a gate-tuned band insulator. *Science*, 338(6111):1193–1196, 2012.
- [148] Wing Chi Yu, Yiu Wing Cheung, Paul J. Saines, Masaki Imai, Takuya Matsumoto, Chishiro Michioka, Kazuyoshi Yoshimura, and Swee K. Goh. Strong coupling superconductivity in the vicinity of the structural quantum critical point in $(\text{Ca}_x\text{Sr}_{1-x})_3\text{Rh}_4\text{Sn}_{13}$. *Phys. Rev. Lett.*, 115:207003, Nov 2015.
- [149] Xiongjie Yu, David J. Luitz, and Bryan K. Clark. Bimodal entanglement entropy distribution in the many-body localization transition. *Phys. Rev. B*, 94:184202, Nov 2016.
- [150] Noah F. Q. Yuan, Kin Fai Mak, and K. T. Law. Possible topological superconducting phases of mos_2 . *Phys. Rev. Lett.*, 113:097001, Aug 2014.
- [151] 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*, 543:217–220, March 2017.
- [152] Yi Zhang and Eun-Ah Kim. Quantum loop topography for machine learning. *Physical Review Letters*, 118(21), 2017.
- [153] Yi Zhang, Roger G. Melko, and Eun-Ah Kim. Machine learning. 05 2017.
- [154] Bo-Xiao Zheng and Garnet Kin-Lic Chan. Ground-state phase diagram of the square lattice hubbard model from density matrix embedding theory. *Physical Review B*, 93(3), 2016.
- [155] Bo-Xiao Zheng, Chia-Min Chung, Philippe Corboz, Georg Ehlers, Ming-Pu Qin, Reinhard M. Noack, Hao Shi, Steven R. White, Shiwei Zhang, and Garnet Kin-Lic Chan. Stripe order in the underdoped region of the two-dimensional hubbard model. *Science*, 358(6367):1155–1160, 2017.

- [156] Zhen Zheng, Ming Gong, Xubo Zou, Chuanwei Zhang, and Guangcan Guo. Route to observable fulde-ferrell-larkin-ovchinnikov phases in three-dimensional spin-orbit-coupled degenerate fermi gases. *Phys. Rev. A*, 87:031602, Mar 2013.
- [157] Zhu Zheng, D. N. Sheng, and Liang Fu. *ArXiv e-prints*.
- [158] Zhenyue Zhu and Steven R. White. Spin liquid phase of the $s = \frac{1}{2} J_1 - J_2$ heisenberg model on the triangular lattice. *Phys. Rev. B*, 92:041105, Jul 2015.
- [159] M. Znidaric, T. Prosen, and P. Prelovsek. Many-body localization in the heisenberg xxz magnet in a random field. *Phys. Rev. B*, 77(6):064426, February 2008.