


Simulating and machine learning quantum criticality in a nearly antiferromagnetic metal



Carsten Bauer
Dissertation

Simulating and machine learning quantum criticality in a nearly antiferromagnetic metal

Inaugural-Dissertation

zur

Erlangung des Doktorgrades

der Mathematisch-Naturwissenschaftlichen Fakultät

der Universität zu Köln

vorgelegt von

Carsten Bauer

aus Frankfurt am Main



Köln 2021

Berichtersteller:

Prof. Dr. Simon Trebst

Prof. Dr. Achim Rosch

Tag der mündlichen Prüfung:

27.11.2020

Abstract

In the last few decades, quantum criticality in itinerant electron systems has become a central focus of condensed matter physics. On the one hand, it represents a candidate mechanism for high-temperature superconductivity, an exciting state of matter with relevance to technical applications. On the other hand, it can lead to a breakdown of Fermi liquid theory, which is inherently fascinating from a theory perspective. Experimentally, there is ample evidence that quantum criticality is responsible for essential features in the phase diagrams of many materials, most notably the unconventional superconductors.

The formation of a comprehensive understanding of these quantum critical phenomena has, however, been significantly hampered by the fact that in many-fermion systems, fluctuations of a critical order parameter can couple to extensive gapless modes on a finite Fermi surface. This interplay, while giving rise to intriguing physical phenomena, also leads to strong electronic correlations, which are notoriously difficult to handle by analytic methods. For this reason, despite numerous attempts, only few definite statements about the properties of metals at the verge of quantum phase transitions could be made.

In this thesis, we investigate metallic quantum criticality by means of large-scale quantum Monte Carlo simulations and contribute unbiased, rigorous results to the discussion. Focusing on the case of antiferromagnetic, collinear, spin-density wave ordering, we present the numerically exact phase diagram of a spin-fermion model whose solution has so far been out of reach. In particular, we highlight the emergence of high-temperature d-wave superconductivity and rigorously establish the quantum critical properties of the model. Specifically, we determine the unusual low-energy scaling of order parameter correlations and study the decoherence of fermionic quasiparticles. In addition to generic energy dispersions, we also consider the case of (almost) local Fermi surface nesting near the hot spots.

Combining the Monte Carlo method with a quantum loop topography approach, we demonstrate that important features of quantum critical metals can be autonomously identified by machine learning current-current correlations. In comparison to conventional techniques, the introduced approach is characterized by a significantly improved numerical efficiency. This allows us to analyze the electronic transport characteristics of two quantum critical metals, including the spin-fermion model, over a large parameter range and leads to the identification of extended non-Fermi liquid regimes in their respective phase diagrams.

Kurzzusammenfassung

Die Erforschung von Quantenkritikalität in elektronischen Vielteilchensystemen ist in den letzten Jahren zu einem zentralen Fokus der Physik der kondensierten Materie geworden. Zum einen können kritische Fluktuationen zu Hochtemperatur-Supraleitung führen, einem faszinierenden Materie-Zustand mit großer Relevanz für technische Anwendungen. Zum anderen können sie zur Dekohärenz von Quasiteilchen und damit einem Zusammenbruch der Fermi-Flüssigkeits-Theorie führen. Zahlreiche experimentelle Hinweise deuten ferner darauf hin, dass Quantenkritikalität für wesentliche Merkmale der Phasendiagramme von Festkörpersystemen, wie insbesondere den unkonventionellen Supraleitern, verantwortlich ist.

Der Aufbau eines umfassenden, theoretischen Verständnisses dieser quantenkritischen Phänomene wird jedoch erheblich dadurch erschwert, dass in fermionischen Vielteilchensystemen quantenkritische Fluktuationen eines Ordnungsparameters an eine Fülle von niederenergetischen Moden auf einer Fermi-Fläche koppeln können. Dieses Wechselwirkung führt typischerweise zu starken elektronischen Korrelationen, die sich in theoretischen Untersuchungen einer Handhabung durch analytische Techniken entziehen. Trotz zahlreicher Forschungsstudien konnten aus diesem Grund bisher nur wenige definitive Aussagen über die Eigenschaften von Metallen am Rande von Quantenphasenübergängen gemacht werden.

In der vorliegenden Dissertation untersuchen wir quantenkritische Metalle mit Hilfe von großskaligen Quanten-Monte-Carlo-Simulationen und bringen wertvolle, rigorose Ergebnisse in die thematische Diskussion ein. Mit speziellem Fokus auf antiferromagnetische Spin-Dichte-Wellen-Ordnung präsentieren wir das numerisch exakte Phasendiagramm eines Spin-Fermionen-Modells, dessen Lösung bisher unzugänglich war. Darauf aufbauend diskutieren wir die Entstehung von Hochtemperatur-Supraleitung, untersuchen die Dekohärenz von fermionischen Quasiteilchen und extrahieren die quantenkritischen Eigenschaften des Modells, wie insbesondere die spezifische, funktionale Form der bosonischen Korrelationen. Hierbei betrachten wir neben generischen Energiedispersionen auch den Grenzfall einer lokalen Fermi-Flächen-Verschachtelung in der Nähe der sogenannten “hot spots”.

Im Anschluss kombinieren wir die Monte-Carlo-Methode mit einer neuen “Quanten-Topographie”-Technik und demonstrieren, dass essentielle quantenkritische Phänomene durch maschinelles Lernen von Strom-Strom-Korrelationen autonom identifiziert werden können. Dabei zeichnet sich der entwickelte Ansatz im Vergleich zu konventionellen Methoden durch eine deutlich verbesserte numerische Effizienz aus. Dies ermöglicht die Untersuchung der elektronischen Transporteigenschaften zweier quantenkritischer Metalle über große Parameterbereiche hinweg und führt zur robusten Identifikation von ausgedehnten “Nicht-Fermi-Flüssigkeits”-Bereichen in deren Phasendiagrammen.

Acknowledgements

First and foremost, I would like to express my sincere gratitude to Simon Trebst, my supervisor, for giving me the opportunity to work on exciting research projects and mentoring me as a PhD student. I am highly grateful for his support and, in particular, the great freedom with which I could pursue my physics, programming, and teaching interests. Second, I wish to thank Achim Rosch, not only for acting as a referee of this thesis, but also for his valuable comments and lively discussions in the condensed matter theory group seminar. I further thank Markus Braden for agreeing to chair the committee of my thesis defense.

Fortunately, in the course of my research, I had the pleasure of collaborating with many inspiring scientists around the world. Above all, I owe a great debt of gratitude to Erez Berg and Yoni Schattner. Our regular meetings over the last four years have been a great pleasure and have significantly impacted both this thesis as well as my general perception of physics. In particular, I am highly grateful for the enlightening discussions and the warmhearted hospitality during my extended research stays at the Weizmann Institute of Science, the University of Chicago, and Stanford University. In this context, I would like to also thank Merav Laniado, Beth Nakatsuka, and Noelle Rudolph, for coordinating my visits, Snir Gazit, Xiaoyu Wang, and Steven Kivelson, for interesting discussions, as well as the German Research Foundation for their financial support (CRC1238 and CRC183). Furthermore, I wish to thank Eun-Ah Kim and her group at Cornell University, especially Yi Zhang, Samuel Lederer, and George Trey Driskell for fruitful collaborations and stimulating exchanges about the interface of physics and machine learning.

My time at the University of Cologne has been a great pleasure, which is why I wish to kindly thank the people at the Institute of Theoretical Physics and especially all members of our research group (forgive me that I won't list all of you explicitly!). Many thanks go to Ciarán and Christoph for proofreading parts of this thesis and to my *excellent* office mates Vaishnavi, Christoph, and Kai for their pleasant company. I am particularly grateful to Ciarán for many interesting exchanges on the the sign problem and Zhong-Chao and Frederic for the in-house collaboration. I further wish to thank Mariela, Benita, and Clara for there cordial assistance in organizational matters and Andreas for his IT support. Of course, I can not go without also mentioning “die festen Körper”, Florian, Tim, Christoph, and Kai, who have always tried to defend the honors of the institute at the yearly University run.

Abschließend möchte ich mich bei meiner Familie, Oriella, Petra, Maris, und Rollin bedanken: Ohne euch würde ich diese Zeilen heute sicher nicht schreiben. Vielen Dank für eure Unterstützung und euer Verständnis!

This thesis is dedicated to Rollin and Petra

Contents

Outline	1
1 Quantum criticality in metals	3
1.1 Quantum phase transitions	3
1.1.1 Metallic quantum critical points	7
1.2 Unconventional superconductors	10
1.3 The spin-fermion model	14
1.4 Analytical insights	19
1.4.1 Fermi surface reconstruction	19
1.4.2 Hertz-Millis theory	21
1.4.3 Hierarchy of energy scales	22
1.4.4 Emergence of a control parameter	25
2 Determinant quantum Monte Carlo	27
2.1 Solving classical thermodynamics at random	27
2.2 Markov chain Monte Carlo	29
2.3 The Metropolis algorithm	32
2.3.1 Proposing moves	33
2.4 Determinant quantum Monte Carlo	34
2.5 Measuring physical observables	37
2.5.1 Equal-time Green's function	37
2.5.2 Time-displaced Green's function	38
2.5.3 Wick's theorem	39
2.6 Error estimation	40
2.6.1 Statistical error of direct sampling	40
2.6.2 Statistical error of Markov chain Monte Carlo	42
2.6.3 Functions of observables	44
2.6.4 Equilibration	45
3 Avoiding the fermion-sign problem	47
3.1 Enforced positive semidefiniteness	47
3.2 Circumventing the sign-problem of the spin-fermion model	48
3.3 Determinant factorization: Hubbard model	52
3.4 Systematic model classification	54
3.4.1 Theorem of T -invariant decomposition	55
3.4.2 Majorana time reversal symmetries	57
3.4.3 Majorana reflection positivity	60
3.4.4 A road to new sign-problem free models	61

4 Aspects of DQMC implementation	69
4.1 Stable linear algebra	70
4.1.1 Time slice matrix multiplications	70
4.1.2 Equal-time Green's function	75
4.1.3 Time-displaced Green's function	77
4.2 Efficient DQMC	81
4.2.1 Efficient adaptive local updates	82
4.2.2 Adaptive global updates	92
4.2.3 Checkerboard decomposition	92
4.2.4 Effective imaginary time slice propagators	96
4.3 Reducing finite-size effects	97
4.3.1 Artificial magnetic flux	97
4.3.2 Twisted boundary conditions	99
4.4 General remarks	100
4.4.1 Benchmark and open source	100
4.4.2 Supercomputing infrastructure	102
5 Numerically exact results for the AFM QCP	103
5.1 Fermi surface and model parameters	104
5.2 Phase diagram ($\lambda = 2$)	105
5.2.1 Antiferromagnetic crossover	105
5.2.2 Unconventional superconductivity	107
5.2.3 Charge-density wave correlations	112
5.3 Quantum critical correlations ($\lambda = 1$)	116
5.3.1 Antiferromagnetic correlations	118
5.3.2 Single-fermion correlations	122
5.4 The case of local nesting	124
5.4.1 Tuning to local nesting	125
5.4.2 Shift of the AFM QCP	126
5.4.3 Quantum criticality	127
5.5 First-order quantum phase transition	130
5.6 Discussion	132
6 Machine learning phases of matter	135
6.1 Artificial neural networks	136
6.2 Supervised learning	138
6.2.1 Backpropagation	139
6.2.2 Regularization and dropout layers	141
6.3 Convolutional neural networks	142
6.4 Discriminating phases of matter	143
6.4.1 Spin-density wave transition in the spin-fermion model	145
7 Quantum loop topography of metallic quantum criticality	147
7.1 Quantum loop topography	147
7.1.1 Current-current correlations	148
7.2 Probing superconducting transport	150
7.2.1 Attractive Hubbard model	150

7.2.2	Spin-fermion model	153
7.3	Mapping out non-Fermi liquid regimes	154
7.3.1	$O(2)$ symmetric AFM QCP	155
7.3.2	Ising-nematic QCP	160
7.4	Discussion	163
Concluding remarks		165
A Hubbard-Stratonovich transformation of the spin-fermion model		169
B Evaluation of the fermion trace		171
C Locating the isotropic AFM QCP		173
D Supplement to QLT classification		175
Bibliography		179
Erklärung		
Publikationen		

Outline

In this thesis, we report on large-scale determinant quantum Monte Carlo and machine learning studies of quantum critical phenomena in metals at the verge of antiferromagnetic spin-density wave transitions. Specifically, the thesis is structured as follows.

In **Chapter 1**, we provide a succinct conceptual introduction into the field of metallic quantum criticality. We discuss the crucial role of the Fermi surface in quantum phase transitions and highlight experimental evidence for quantum critical phenomena in the phase diagrams of many unconventional superconductors. We then motivate a spin-fermion model for the onset of antiferromagnetism in metals, the major target of the numerical efforts in this thesis, and summarize various analytical attempts to determine its low-energy properties.

Following this, we turn to a discussion of simulation techniques for numerically studying quantum critical metals. To that end, **Chapter 2** introduces the determinant quantum Monte Carlo (DQMC) method in a pedagogical, self-contained manner. Starting from classical statistical physics, we describe how the concept of Markov sampling can be utilized to efficiently evaluate expectation values of physical observables and outline how the famous Metropolis algorithm can be applied to quantum partition functions of metals. Furthermore, we discuss the numerical error associated with statistical sampling and means to estimate the same within DQMC. The focus of **Chapter 3** is the biggest limitation of quantum Monte Carlo simulations: the infamous fermion sign-problem. In particular, we address strategies to circumvent the occurrence of negative “probabilities” and systematically describe frameworks for categorizing sign-problem free Hamiltonians. This includes a novel semigroup approach, which we apply to identify a many-fermion model of a qualitatively new sign-problem free class. In **Chapter 4**, we discuss a selection of implementation aspects of DQMC. Specifically, we highlight how numerical instabilities arise in Green’s function computations and present a systematic assessment of stabilization schemes to keep the numerical error under control. We further describe various means to speed up Monte Carlo simulations and to reduce the impact of finite-size effects.

In **Chapter 5**, we provide a rigorous, unbiased DQMC study of the spin-fermion model. We present the numerically exact phase diagram of the effective theory and investigate the intertwined magnetic and electronic ordering tendencies in detail. A special focus of our large-scale simulations are the quantum critical properties of the model in the vicinity of the onset of antiferromagnetic order. Specifically, we establish the decoherence of fermion quasiparticles and provide the precise scaling-form of low-energy order parameter correlations. We consider regular and (almost) locally nested Fermi surfaces.

Turning to machine learning in **Chapter 6**, we introduce artificial and convo-

lutional neural networks and showcase that by means of a supervised training scheme, these statistical models can be utilized to discriminate phases of matter.

Finally, in **Chapter 7**, we introduce a novel quantum loop topography scheme for machine learning longitudinal transport and demonstrate that this approach can be utilized to probe quantum critical phenomena based on DQMC data. Specifically, we present studies of quantum many-fermion systems, including the spin-fermion model, in which we locate superconductivity transitions and map out extended non-Fermi liquid regimes.

As indicated by the corresponding publications, some of the presented results have been obtained in international collaborations. In particular, Erez Berg (Weizmann Institute of Science) and Yoni Schattner (Stanford University) were close partners in the DQMC study of Ch. 5. Furthermore, Eun-Ah Kim, Samuel Lederer, and George Trey Driskell (Cornell University) as well as Yi Zhang (Peking University) were collaborators in the machine learning studies of Ch. 7. For clarity, we have indicated the contributions by the author of this thesis in the respective introductions.

Note that we have published the scientific codes developed as part of this dissertation under open source licenses in Refs. [S1–S4]. This includes a state-of-the-art DQMC implementation and various single-purpose software libraries.

1 | Quantum criticality in metals

In this chapter, we provide a concise introduction into quantum criticality in metals with a special focus on antiferromagnetic (AFM) spin-density wave (SDW) ordering. To that end, we start with a review of the phenomenology of quantum phase transitions and, in particular, discuss how the presence of a finite Fermi surface, which fundamentally defines a system of itinerant fermions, qualitatively impacts the anticipated physical properties and aggravates analytical studies. Afterwards, we review experimental evidence for metallic quantum criticality in various classes of unconventional superconductors, such as copper- and iron-based compounds as well as “heavy-fermion” materials. We then turn to the theoretical description of the onset of antiferromagnetic order in itinerant electron systems and introduce the spin-fermion model [1–5], which is the major target of numerical simulations in this thesis. After highlighting its conceptual origin, we review the large body of analytical work on the model, including mean-field, perturbation, and Eliashberg theories. Finally, we discuss the particularly intriguing studies by Lee *et al.* [6, 7], in which the authors have identified an emergent control parameter and a novel strongly interacting fixed-point with unique characteristics. For more comprehensive introductions into quantum phase transitions and metallic quantum criticality we refer the interested reader to Refs. [1, 2, 8–16], which have influenced the compilation below. Note that the presentation in Sec. 1.4.3 is based on the review section of Ref. [P1], written by the author of this thesis.

1.1 Quantum phase transitions

Phase transitions mark a qualitative change of the physical properties of a system upon variation of an external control parameter. Classically, phase transitions are driven by thermal fluctuations. Any form of structural, magnetic, or electronic order, present at low temperatures is destroyed at the transition temperature T_c by these critical fluctuations. Examples include the melting of a (ice-like) solid, a permanent magnet losing its macroscopic ferromagnetism, and the superconductor-metal transition. Quantum phase transitions on the other hand, while similar in spirit, occur strictly at the absolute zero of temperature and are governed by quantum fluctuations, whose origin lies in Heisenberg’s uncertainty principle. Upon tuning a non-thermal parameter r , such as pressure, doping, or a microscopic coupling, the system transitions between two quantum ground states, typically of different symmetry, that meet at the quantum critical point (QCP) r_c [14]. The quantum paramagnet transition of the transverse field quantum Ising model [1], the quantum Hall transition [17], and the interaction-driven

Mott-insulator transition of metals [18] represent principle examples of quantum phase transitions.

Traditionally, phase transitions are classified into first-order and continuous (second-order) transitions [1, 8, 9, 11, 14]. A distinguishing feature of the former class is the coexistence of phases near the critical point, such as ice and water at $T_c = 0^\circ$ C. In this case, a finite amount of latent heat is absorbed from or released to the thermodynamic environment to change the state of the system. In contrast, when driven through a continuous phase transition the system, typically, breaks a microscopic symmetry, such as translation, spin-rotation, or a $U(1)$ symmetry associated with particle conservation, spontaneously and alters its character by self-organization of the relevant degrees of freedom¹. In the following, we concentrate on this second class of phase transitions which can be described by an order parameter: a thermodynamic quantity whose average is finite in the ordered phase and vanishes continuously at the phase boundary. In terms of fluctuations of the latter, the quantum phase transition is invariably signified by a diverging length scale of spatial correlations, [1, 8, 11, 19]

$$\xi \sim |r - r_c|^{-\nu}. \quad (1.1)$$

Here, $r - r_c$ is the tuning parameter deviation from criticality and ν is the critical correlation length exponent. In addition to this spatial scaling, the quantum critical point is characterized by a vanishing energy scale Δ , which may be identified with the typical energy of long-distance order parameter fluctuations [8] such as, if nonzero, an energy gap of low-lying excitations in the ordered state [1]. Consequently, there is an associated timescale $\tau_c \sim \Delta^{-1}$ which diverges upon approaching the critical point and, similar to Eq. (1.1), indicates the presence of long-range correlations in time, [1, 8, 19]

$$\tau_c \sim \xi^z \sim |r - r_c|^{-\nu z}. \quad (1.2)$$

Here, z is the dynamical critical exponent which relates the scaling in space and time.

The divergences in Eqs. (1.1) and (1.2) constitute the origin of quantum criticality [8, 19]. Near the critical point, order parameter fluctuations occur at all length and timescales such that the associated correlation functions extend over ever larger regions and their regular exponential decay turns into a slower power law scaling [11]. Concomitantly, the system becomes scale-invariant, $f(\lambda x) = (\lambda x)^c = \lambda^c x^c \sim f(x)$, and its critical properties are entirely characterized by the set of critical power law exponents – we have already touched upon ν and z as prominent examples. Figuratively speaking, the system statistically looks the same at all length scales beyond a microscopic lattice scale². A fascinating consequence of this emergent³ self-similarity is universality: The microscopic structure, that is, for instance, the nature of the elementary degrees of freedom

¹Here, the term self-organization should not be confused with its meaning in a different context: the tendency of a dynamical system to, effectively, tune itself to criticality.

²A beautiful visualization of the RG-flow of the two-dimensional Ising model and its scale invariance at criticality can be found in Ref. [20].

³The microscopic model generally isn't scale-invariant.

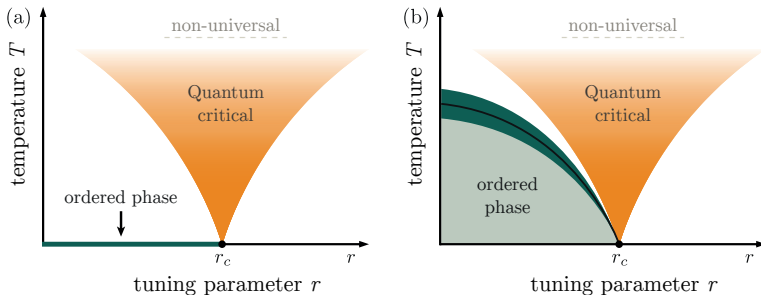


Figure 1.1: Quantum criticality in insulators. At the QCP (r_c) the system transitions from an ordered phase (green) to a quantum disordered ground state. This generally gives rise to an extended quantum critical region (orange) in which the scaling is governed by the QCP. The boundaries of this “fan” indicate crossovers and are determined by the condition $k_B T \sim \Delta \sim |r - r_c|^{\nu z}$. Energy scales associated with microscopic couplings act as universality cut offs at high temperatures. The ordered phase may be restricted to $T = 0$ (a) or extend to finite temperatures (b) in which case it is delimited by a classical phase transition (black) terminating in the QCP. In the envelope region around this transition line a classical description is sufficient. Reproduced from Refs. [1] and [8].

and their coupling strengths, is rendered unimportant by the critical divergences, Eqs. (1.1) and (1.2), such that vastly different physical systems share the same macroscopic critical behavior. Upon increasing the length and timescales of consideration, i.e. “zooming out”, the couplings in two microscopically distinct models get effectively renormalized and, if in the same basin of attraction, flow to the same universal fixed point - a formal implementation of this procedure is given by the renormalization group (RG) [17, 19]. Generally, only a few essential qualitative characteristics, such as the dimensionality of space and the symmetry of the order parameter, determine the long-wavelength critical scaling and allow for a categorization into universality classes [1]. For instance, in two-dimensions, the liquid-gas transition of a Lennard-Jones fluid and the classical Curie transition of an uniaxial magnet are characterized by precisely the same set of critical exponents⁴ [19, 21].

In view of the fact that, strictly speaking, quantum phase transitions only occur in the zero temperature limit, one may ask about their relevance for real experiments, which are inevitably conducted at finite, though possibly very low, temperatures. As it turns out, in spite of the fact that one generally needs a fine tuning to reach the vicinity of a quantum critical point at $T = 0$, the latter can govern the universal scaling in an extended region of the phase diagram [1, 8, 15], that is for a broad range of values for $|r - r_c|$ and T , Fig. 1.1. By comparing the energy scale Δ and the scale set by temperature, $k_B T$, [1, 8]

$$k_B T \sim \Delta \sim |r - r_c|^{\nu z}, \quad (1.3)$$

⁴See Ref. [21] for a visualization of this universality.




Figure 1.2: Visualization of the $d + 1$ -dimensional spacetime grid of the quantum-classical mapping. For illustration purposes, the d spatial dimensions have been combined into a single “space” dimension. The finite extent of the imaginary time axis is $\beta = \hbar/k_B T$. Shown are configurations of an $O(2)$ -symmetric order parameter in the ordered and quantum disordered phase left and right of a QCP at $T = 0$. Inspired by Ref. [10].

one can fundamentally identify a quantum critical region, Fig. 1.1, in which quantum and thermal effects are equally important [1]. In the latter, the system features significant footprints of quantum criticality that are blurred by thermal fluctuations [8]. This interplay gives rise to a variety of unusual physical properties such as unconventional power-laws and, for metallic systems, non-Fermi liquid behavior.

Apart from the quantum critical fan, we can conceptually distinguish two finite-temperature scenarios, Fig. 1.1, instances of both of which will appear in this thesis. As shown in Fig. 1.1a, one possibility is that the phase on the ordered side of the QCP exclusively exists at $T = 0$. Low-dimensional systems, in which thermal fluctuations fundamentally preclude the formation of long-range order at $T > 0$ according to the Mermin-Wagner theorem [1, 8, 17, 22] are examples of this case. Alternatively, there might be an ordered phase present at finite temperatures. In this scenario, the corresponding phase boundary terminates in the QCP, as illustrated in the schematic phase diagram in Fig. 1.1b. This $T > 0$ phase transition, which in our simulations will turn out to be useful as a practical guide for locating the QCP, is not really a property of the QCP and, at criticality, can be described by a purely classical theory [1].

A powerful concept for understanding and simulating quantum criticality – we will make the latter precise in Ch. 2 – is the quantum-classical mapping, Fig. 1.2, underlying the functional integral formulation of quantum statistical physics [1, 17, 19, 23]. The key idea is to interpret the functional integral of a d -dimensional quantum field theory as a conventional statistical ensemble of a classical field theory in $D = d + 1$ dimensions⁵ [1]. Hereby, the additional

⁵As indicated by Eq. (1.2), correlation functions of the classical theory might scale differently in space and imaginary time direction. In this case, the effective dimension is $D = d + z$, where z is the dynamical critical exponent [8].

dimension is the imaginary time τ . Compared to the d space dimensions it is distinguished in the sense that it generally is of finite extent: Its size is given by the inverse temperature⁶ $\beta = 1/k_B T$ of the quantum system and, to account for the trace in the partition function $Z = \text{Tr}[e^{-H/k_B T}]$ of a Hamiltonian H , it comes with periodic boundary conditions. According to the quantum-classical correspondence, quantum fluctuations of the original system may now simply be viewed as order parameter fluctuations along the imaginary time dimension of the associated classical theory. Note that, as per custom, we use the designation “quantum-classical” despite the fact that the resulting classical theories may be rather artificial and, due to their quantum origin, often involve complex-valued Boltzmann weights [1], Ch. 3.

As phase transitions and universality depend sensitively on the dimensionality of a system, the finiteness of the temporal axis⁷ manifests novel quantum critical behavior beyond what is captured by canonical classical theories. At the same time, one can directly deduce important classical limits of quantum criticality from this point of view. We notice that in the zero temperature limit $T \rightarrow 0$ the extent of the imaginary time dimension diverges, $\beta \rightarrow \infty$. This implies that the criticality of a d -dimensional quantum system at $T = 0$ may be identified with the criticality of a truly $(d + 1)$ -dimensional classical system [1, 8, 10]. Similarly, for $T \rightarrow \infty$ (or $\hbar \rightarrow 0$) the temporal dimension vanishes, thus representing a static limit, such that an effective d -dimensional classical theory suffices to capture the physics of the quantum critical system.

It is important to highlight that beyond the general thermodynamic considerations above, there are crucial limitations to the quantum-classical mapping. In particular, care has to be taken when considering quantum critical dynamics. Since the mapping only yields correlation functions in imaginary time, the computation of real time correlations, which are of particular relevance for experiments like nuclear magnetic resonance (NMR), inelastic neutron scattering, and angle-resolved photoemission spectroscopy (ARPES) [1], generally necessitates a form of analytic continuation. Unfortunately, the latter represents an ill-posed problem when combined with any kind of analytical or numerical approximation scheme and, if circumvented naively, leads to arbitrarily inaccurate and unphysical predictions [1]. Hence, any attempt to transfer established dynamical properties of classical $(d + 1)$ -dimensional systems at criticality to corresponding d -dimensional quantum systems is inherently flawed. In the determinant quantum Monte Carlo simulations in this thesis we will account for this fact by analyzing relevant imaginary time proxies instead.

1.1.1 Metallic quantum critical points

The critical behavior of a quantum system crucially depends on whether order parameter fluctuations are the only relevant low-energy degrees of freedom near the QCP. This is often the case in insulators where the critical theory is of

⁶Unless stated otherwise, we set $\hbar = 1$.

⁷For convenience, we will generally treat imaginary time as a “temporal” dimension similar to real time and only indicate the distinction when necessary.

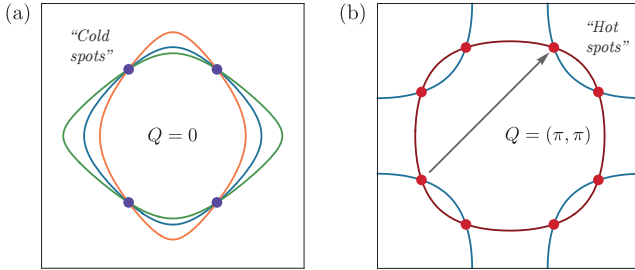


Figure 1.3: Fermi surfaces of metals at the verge of an Ising-nematic (a) and an antiferromagnetic spin-density wave (b) transition with ordering wave vectors $\mathbf{Q} = 0$ and $\mathbf{Q} = (\pi, \pi)$, respectively. The surrounding square indicates the momentum space Brillouin zone. In (a) the original Fermi surface (blue) is elongated in k_x or k_y direction as a consequence of the breaking of the rotational symmetry. Only a few cold spots (blue points) remain unaltered. For the AFM transition (b), the finite wave vector connects different hot spots (red points) at which low-energy fermions can scatter off order parameter fluctuations while remaining on the Fermi surface (blue). The red line is the Fermi surface shifted by \mathbf{Q} so that the hot spots correspond to intersection points. Inspired by Ref. [24].

Ginzburg-Landau-Wilson type (see Eq. (1.4)) and can be expressed in terms of these soft modes and based on symmetry considerations alone [1, 19, 25].

The physics is far richer in a metal where order parameter fluctuations can couple to gapless excitations on a finite Fermi surface, an inherently quantum object⁸ rooted in the Pauli principle. In general, we can distinguish between two classes of metallic QCPs based on whether the critical fluctuations carry a finite ordering wave vector \mathbf{Q} , Fig. 1.3. If $\mathbf{Q} = 0$, the collective modes induce scattering on the entire Fermi surface which, pictorially speaking, becomes “hot”. Only a number of “cold spots”, whose precise location is determined by the symmetry of the order parameter, are invariant under the continuous symmetry breaking at the QCP. Examples for quantum phase transitions that fall into this category are the Ising-nematic transition (shown in the Fig. 1.3), a Fermi liquid (FL) at the verge of a Pomeranchuk instability to broken rotational symmetry [26], and a spinon Fermi surface coupled to a gapless gauge boson as emergent in the low-energy sector of a $U(1)$ spin-liquid [15, 26]. A non-vanishing wave vector $\mathbf{Q} \neq 0$ on the other hand indicates the onset of density-wave order of e.g. spin, charge, or orbital degrees of freedom. In this virtually inverted scenario the order parameter fluctuations couple resonantly only to a discrete set of points on the Fermi surface: the “hot spots”. As shown in Fig. 1.3, these points are linked by the wave vector \mathbf{Q} and, in the low-energy limit, mark the most relevant patches of the Fermi surface. Among others, magnetic and charge-density wave QCPs are described by hot spot theories and have been studied in Refs. [27–32]. With the exception of Ch. 7, we will exclusively consider this second class of

⁸In light of the quantum-classical mapping, we note that there is no classical limit to a Fermi surface.

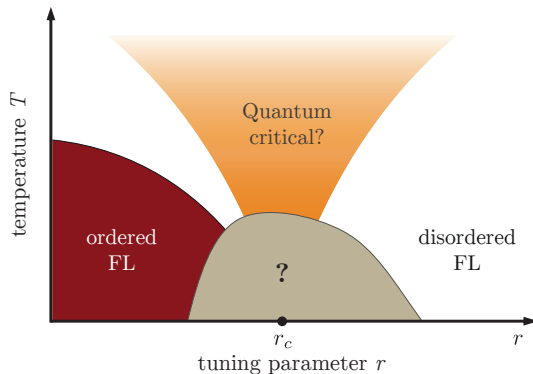


Figure 1.4: Schematic phase diagram of a metal near quantum criticality. Upon decreasing the value of the tuning parameter r , the system transitions from a disordered Fermi-liquid state (white) into an ordered Fermi-liquid state (red) such as a spin density-wave phase. Near the QCP, $r \approx r_c$, (black dot) strong quantum fluctuations can give rise to a quantum critical non-FL regime (orange) and induce the formation of a different kind of electronic order (beige), for example, superconductivity. Inspired by Refs. [1, 5, 25].

metallic quantum phase transitions in this thesis; specifically, variants of an antiferromagnetic spin-density wave transition with $Q = (\pi, \pi)$, Fig. 1.3. We note in passing that the ordering wave vector will always be commensurate with the lattice, a fact which, in principle, can impact the critical behavior and the universality classification [29, 33, 34].

In Fig. 1.4, we show a schematic phase diagram of a metal at the verge of a quantum phase transition into a metallic ordered phase. While conventional Fermi-liquids are remarkably stable against small perturbations, the critical quantum fluctuations in the vicinity of a QCP can cause a breakdown of Landau’s theory and lead to a state with unusual power-law scaling [1, 3–7, 15, 26, 35, 36] and a destabilized Fermi surface characterized by incoherent quasi-particle excitations [1, 24, 26, 35, 37]. As discussed in the previous section, this anomalous non-FL state is expected to govern the physical features of an emergent quantum critical regime at finite temperatures [15]. However, yet another intriguing possibility is that the critical order parameter fluctuations mediate singular inter-fermion interactions and, in this way, promote a secondary⁹ electronic order. Most notably, the collective modes may represent a pairing “glue” for high-temperature superconductivity [1, 4, 5, 31, 32, 38–42] – the focus of the next section – or induce charge-, spin-, and pair-density wave phases [5, 34, 43–48]. In this case, the QCP, and potentially even the entire incipient quantum critical fan, will be overshadowed by this emergent electronic order, Fig. 1.4.

Whether the formation of a novel non-Fermi liquid state is pre-empted by

⁹This emergent electronic order is generally distinct from the primary order that onsets at the QCP.

such a gapping out of the Fermi surface [41] and, if so, which of these various competing or cooperating orders dominates the physical behavior in the vicinity of a particular QCP is one of the central questions of any theoretical analysis. Similarly, if the secondary order is suppressed or entirely absent, obtaining a systematic understanding of the nature of the non-FL state and its universal properties is of central importance. However, as the physical richness of a quantum critical metal is rooted in the presence of strong inherent correlations, a theoretical analysis of the same poses a substantial challenge. Particularly interesting two dimensional systems are typically below the upper critical dimension and are characterized by pronounced infrared quantum fluctuations [15]. Despite numerous attempts over the past decades [1, 3, 4, 15, 34–36, 49, 50] few definite analytical statements about the properties of metallic QCPs could be made – we will discuss recent exceptional progress by Lee et. al [6, 7, 15] in Sec. 1.4.4. As a consequence, even basic aspects of the critical behavior are still under debate and ask for controlled results by complementary techniques [2, 25].

It is a central purpose of this thesis to provide the latter for the AFM QCP and to significantly improve the understanding of metallic quantum criticality in this case by means of numerically exact quantum Monte Carlo simulations, Ch. 5, and novel machine learning techniques, Ch. 7. With respect to the former, our work expands on a recent line of numerical efforts [2] that have targeted the low-energy properties of metals at the verge of quantum phase transitions involving the breaking of various magnetic [27–29, 31–33, 51], charge density wave [30], and nematic [30, 39, 40, 49, 52] orders. Complementary studies have focused on deconfined QCPs, where fermionic matter fields are coupled to Z_2 [45, 53–55] or $U(1)$ [56] lattice gauge theories.

1.2 Unconventional superconductors

Going beyond the theoretical appeal of metallic quantum criticality, the onset of order in a metal is of vital importance to the understanding of an intriguing class of correlated electron materials: the unconventional high-temperature superconductors [1, 25, 34, 57]. Since its discovery in 1911, superconductivity is arguably one of the most fascinating states of condensed matter. It is a striking example of a macroscopic collective quantum phenomenon in itinerant many-electron systems [17, 26]. The realization of a state of matter in which the charge carriers conduct electricity with zero resistance has quickly spawned dreams of bringing this alluring quantum feature to room temperature. However, as Bardeen, Cooper, and Schrieffer (BCS) derived their “Microscopic Theory of Superconductivity” [58] in 1957 – an achievement worthy of the physics nobel prize in 1972 – it seemed that superconductors couldn’t exceed a transition temperature of about $T_c \sim 30K$ [59], except when put under high pressure [60]. With the discovery of the copper oxide superconductors in 1986 [59, 61], which ranks among the major scientific highlights of the twentieth century [59], this picture changed entirely. The predicted upper limit for the critical temperature was spectacularly surpassed multiple times before the turn of the millennium with T_c values of up to $138K$ [16, 59], which is still the record at atmospheric




Figure 1.5: Phase diagram of $\text{BaFe}_2(\text{As}_{1-x}\text{P}_x)_2$. The exponent α of the T -dependence of the resistivity is indicated by the color shading. The temperatures T_s (yellow triangles), T_N (gray circles) and T_c (white squares) correspond to structural (nematic), spin density wave (SDW), and superconductivity (SC) transitions. The Weiss temperature θ (pink triangles) is determined by the nuclear magnetic resonance (NMR) relaxation rate. A steep increase of the effective mass m^*/m_b (right axis) is observed near the QCP in de Haas-van Alphen (light green squares) and specific-heat measurements (dark green diamonds). Reprinted with permission from Ref. [57].

pressure today [59]. Crucially, superconductivity in these systems as well as the later discovered iron-pnictide superconductors [57, 62] is not captured by the conventional BCS theory [16, 38, 59, 63] and must be emanating from a different microscopic mechanism. Despite an enormous research effort, manifested in over 100000 publications [16], this origin still remains elusive and represents one of the major challenges of condensed-matter physics [15, 16, 64].

In this thesis we approach unconventional superconductivity from the perspective of a metallic QCP scenario. To motivate the latter, we discuss experimental phase diagrams of an iron-pnictide superconductor and doped copper oxides below. We do not attempt to provide a review of the state of experimental research but rather try to highlight aspects of quantum criticality in these materials. More comprehensive reviews in this spirit can be found in Refs. [38, 57] for the Fe-based and in Ref. [16, 38] for the copper-based materials. We note in passing that although we focus on only two material classes, evidence for a metallic QCP also exists for heavy-fermion [38, 65, 66] and organic superconductors [67, 68].

Fig. 1.5 shows the phase diagram of the prototypical iron-pnictide superconductor $\text{BaFe}_2(\text{As}_{1-x}\text{P}_x)_2$ which is perhaps the most prominent case for quantum criticality. While the parent compound BaFe_2As_2 exhibits antiferromagnetic SDW order, the isovalent¹⁰ partner BaFe_2P_2 is a conventional Fermi liquid. Hence, there inevitably is a (quantum) phase transition present in between those

¹⁰In contrast to increasing or decreasing the carrier concentration, isovalent doping leads to a form of microscopic pressure due to different ion sizes [57].




Figure 1.6: Phase diagrams of the copper oxide compounds $\text{RE}_{2-x}\text{Ce}_x\text{CuO}_4$, where $\text{RE} = \text{La}, \text{Pr}, \text{Nd}$, and $\text{La}_{2-x}\text{Sr}_x\text{CuO}_4$. Upon electron- or hole-doping the antiferromagnetic Mott-insulating systems (AF) turn superconducting (SC) below a critical temperature T_c . In both scenarios there is a pseudogap regime (only shown for hole-doping because of copyright restrictions) entering the SC dome close to optimal doping. Reprinted with permission from Ref. [63], Copyright (2020) by the American Physical Society.

limits. Specifically, upon increasing the phosphorus content at $T = 0$, the metal enters a nodal superconducting state which, initially, coexists with the SDW order. Right above the¹¹ putative QCP, $x_c \approx 0.3$, at which the Curie-Weiss temperature associated with the magnetic susceptibility vanishes, the critical superconducting T_c finds its maximum value. Quantum oscillation and specific-heat measurements in the vicinity of x_c reveal a sharp increase of the effective electron mass which suggests the presence of strong fermionic correlations due to the interplay of quasi-particles on the Fermi surface and critical SDW modes. At finite-temperatures, a striking indication of metallic quantum criticality is the observed linear T -dependence of the resistivity, $\alpha \approx 1$, in an extended region centered around $x \approx x_c$. This distinct non-Fermi liquid behavior near the QCP is in stark contrast to the regular $\alpha = 2$ scaling observed in the $x \rightarrow 1$ doping limit and reminiscent of the quantum critical region in Fig. 1.4.

Turning to the cuprate superconductors $\text{La}_{2-x}\text{Sr}_x\text{CuO}_4$ and $\text{R}_{2-x}\text{Ce}_x\text{CuO}_4$, Fig. 1.6, the situation becomes much more complex. We start by noting that the antiferromagnetic order of the parent compounds is non-metallic and of a Mott-insulating kind. While in $\text{R}_{2-x}\text{Ce}_x\text{CuO}_4$ superconductivity occurs in direct proximity to this AFM state there is a notable separation between magnetic long range order and a SC phase in the hole-doped partner at lower band fillings. Interestingly, in both doping scenarios a pseudogap, indicating a loss of spectral weight [16, 25], is found to onset below temperatures T^* and seems to originate from near optimal doping x_{T_c} . Similar to the iron-pnictide, Fig. 1.5, linear- T scaling of the resistivity is observed in a marked “strange metal” regime above this point [16, 38, 70]. Taken together with the detection of hot spots in photoemission

¹¹We note that it has been questioned whether there is only a single QCP because of an additional structural (nematic) transition the sets in right above T_N [69].

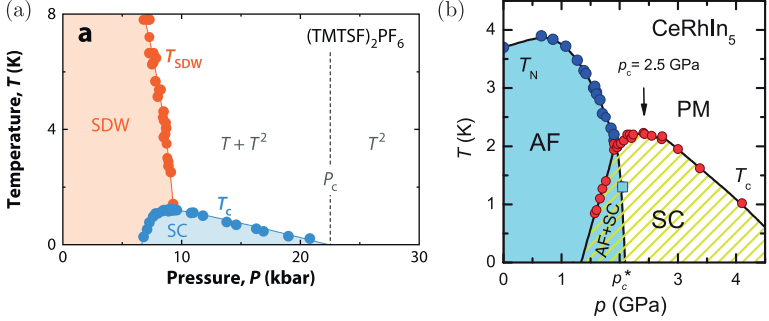


Figure 1.7: Phase diagrams of the organic superconductor $\text{TMTSF}_2\text{PF}_6$ (a) and the heavy-fermion material CeRhIn_5 (b). In both cases, unconventional superconductivity (SC) emerges below T_c in the vicinity of a phase of antiferromagnetic order (SDW/AF). The critical transition temperature associated with the latter is indicated by T_{SDW} or T_N . Subfigures (a) and (b) have been reprinted with permission from Ref. [16] and [76], respectively.

spectra [16, 71] this is indicative of a QCP within the superconducting dome [16, 38, 72, 73] which is likely to be linked directly to the AFM order (electron-doping) or to the pseudogap state (hole-doping). However, there is a (somewhat mysterious) uncertainty about the nature of the latter. Numerous experiments have identified a plethora of enhanced correlations, perhaps signaling the presence of intertwined orders [74], and provided evidence for various forms of spontaneous symmetry breaking at T^* , including the breaking of time reversal symmetry [16, 25]. Quantum oscillation experiments in strong magnetic fields – to suppress superconductivity – have revealed a diverging effective mass at two distinct doping values suggesting that there might be more than one QCP at play [75]. In any case, antiferromagnetic fluctuations, either short or long-ranged, seem to be omnipresent in the vicinity of superconducting phases across almost all classes of unconventional superconductors [16, 34, 76–79]. We show additional phase diagrams of organic and heavy fermion superconductors in Fig. 1.7, rendering the AFM QCP one of the most relevant candidates. Augmenting the latter by an order parameter coupling to an emergent \mathbb{Z}_2 gauge field has most recently been proposed as a theory for the pseudogap phase of the hole-doped cuprates [45, 80].

In this work, we will not try to resolve the nature of the QCP in the unconventional superconductors. Instead, we focus on the AFM QCP and perform extensive computer simulations to establish its pristine phenomenology in a rigorous (numerically exact) manner, Ch. 5, and by means of a modern machine learning approach, Ch. 7. This will serve as an important guide for theory and, hopefully, contribute a valuable benchmark to the experimental discussion. Let us highlight that previous quantum Monte Carlo studies have investigated scenarios in which the onset of metallic AFM order is of easy-axis [33] or easy-plane character [31, 32]. While we will consider the latter in the machine learning

study presented in Ch. 7, our main emphasis lies on the isotropic case, where the onset of antiferromagnetism breaks an $O(3)$ symmetry [81]. This focus aligns with the vast majority of analytic theories¹² [3–5, 34, 35, 50, 69, 82] and should be particularly relevant for systems where spin-orbit coupling is subordinate. From a technical point of view, our simulations of the isotropic AFM QCP come at an (up to) eight-fold increased computational cost – see Fig. 4.12 in Ch. 4 for a more thorough analysis.

1.3 The spin-fermion model

As we have argued above, there are two fundamental low-energy degrees of freedom near a continuous quantum phase transition in a metal: fluctuations of the order parameter and fermionic excitations on a Fermi surface. In the following, we formulate a quantum field theory for the AFM QCP: the spin-fermion model [1–5]. This minimal model describes the interplay of the gapless modes near the QCP and captures the universal aspects of the onset of antiferromagnetism in a two-dimensional system of itinerant electrons.

Our starting point is a quantum-classical version, Fig. 1.2, of the paradigmatic Ginzburg-Wilson-Landau (GWL) action [1, 8, 19, 83],

$$S_\phi = \int_0^\beta d\tau \int d^2x \left[\frac{1}{2c^2} (\partial_\tau \phi)^2 + \frac{1}{2} (\nabla \phi)^2 + \frac{r}{2} \phi^2 + \frac{u}{4} (\phi^2)^2 \right], \quad (1.4)$$

for a N_ϕ -component bosonic order parameter field ϕ . Here, c is a velocity, r and u are mass- and interaction-like coupling constants and the upper limit of the imaginary time integral is generally finite. The associated partition function is given by the functional integral

$$\mathcal{Z}_\phi = \int D\phi e^{-S_\phi} \quad (1.5)$$

over ϕ -field configurations periodic in τ [1]. For $N_\phi = 1, 2, 3$ respectively, the “ ϕ^4 -theory” above describes the universal physics of quantum Ising, XY, and Heisenberg models at long-wavelengths and low-energies [1, 84]. As discussed around Eqs. (1.1) and (1.2), it is mainly characterized by the symmetry of the order parameter and the number of space-time dimensions. Physically, the bosonic field $\phi(\tau, x)$ may be viewed as a coarse-grained average of local magnetic moments. As the microscopic degrees of freedom are subject to a hard length constraint, the magnitude of ϕ is “softly”¹³ restricted in Eq. (1.4) by the effect of terms involving higher order powers of ϕ^2 [1]. An explicit derivation of a 2-dimensional, rather than (2+1)-dimensional, version of S_ϕ from a microscopic classical Ising Hamiltonian is presented in Refs. [1, 19].

The relevance of the GWL action for understanding critical phenomena was already emphasized by Wilson and Fisher in the 1970s [83, 86]. Famously, the

¹²Note, however, that the number of order parameter components often only enters as a numerical factor.

¹³For $N_\phi \geq 2$ and in (2+1) dimensions one may alternatively consider a quantum nonlinear sigma model with a hard constraint which has the same universal properties [1, 84, 85].

theory displays spontaneous symmetry breaking: the $O(N)$ symmetry corresponding to rotations in order parameter space is spontaneously broken upon decreasing the tuning parameter r . A Landau-theory of this circumstance is obtained by applying the saddle-point approximation to the integral in Eq. (1.5) in which the integrand is estimated at the single field value ϕ_0 that minimizes S_ϕ . Clearly, for $r > 0$ (and assuming $u > 0$) the optimum is given by the constant solution $\phi_0 = 0$. In contrast, for $r < 0$ a finite ϕ -value will be preferential. Hence, we expect a quantum critical point at $r = r_c = 0$ that marks a continuous phase transition between a paramagnetic ($\langle\phi\rangle = 0$) and a magnetically ordered ($\langle\phi\rangle \neq 0$) phase [1]. Of course, fluctuation corrections associated with the quartic coupling u , or caused by the interplay with gapless fermions in the extended theory below, will alter this simple picture. In particular, coupling constants will generally get renormalized and the critical value r_c will deviate from zero. Note that time and space scale in the same way in S_ϕ and the dynamical critical exponent at the Wilson-Fisher fixed point¹⁴ is $z = 1$.

Turning to the metallic structure, we describe the kinetics of non-interacting itinerant spin-1/2 fermions by the following action,

$$S_\psi = \int_0^\beta d\tau \int d^2k \sum_s \psi_{\mathbf{s}\mathbf{k}}^\dagger (\partial_\tau + \epsilon_{\mathbf{k}} - \mu) \psi_{\mathbf{s}\mathbf{k}}. \quad (1.6)$$

Here, $s = \uparrow, \downarrow$ denotes spin and the ψ, ψ^\dagger are conjugated Grassmann valued fields representing fermions with an energy dispersion ϵ_k and chemical potential μ ¹⁵. Inspired by the cuprates, in which superconductivity originates from two-dimensional CuO_2 layers, we consider the dispersion

$$\epsilon_{\mathbf{k}} = -2t [\cos k_x + \cos k_y] - 2t' [\cos(k_x + k_y) + \cos(k_x - k_y)], \quad (1.7)$$

corresponding to nearest (t) and next-nearest (t') neighbor hopping on a square lattice. The resulting Fermi surface for $t = 1, t' = -0.5$, and $\mu = -0.5$ is shown in Fig. 1.3b. In terms of universality, we generally do not expect quantitative variations of the latter to affect the critical properties. We will vary the precise parameters throughout this thesis and will consider a qualitatively distinct local nesting scenario in Sec. 5.4.

The final step towards a quantum field theory for the AFM QCP is to couple the (Heisenberg) order parameter ϕ and the fermions ψ, ψ^\dagger together. In particular, this interaction must be designed with the targeted symmetry in mind. In order for the quantum phase transition of S_ϕ to induce antiferromagnetic spin ordering, that is break the $SU(2)$ spin symmetry of S_ψ , we must couple ϕ to the fermionic spin-density with an appropriate form factor. To lowest order in the fields, we consider the following contribution to the action

$$S_\lambda = \lambda \int_0^\beta d\tau \int d^2x \sum_{s,s'} e^{i\mathbf{Q}\cdot\mathbf{r}} \phi_{\mathbf{r}} \cdot \psi_{s\mathbf{r}}^\dagger \sigma_{ss'} \psi_{s'\mathbf{r}} + \text{h.c.} \quad (1.8)$$

¹⁴We consider $d = 2$ space dimensions which is below the upper critical dimension $d_c = 4$ [19].

¹⁵Throughout this thesis, we will label Pauli matrices by σ , use ψ, ψ^\dagger to indicate Grassmann fields, and denote fermionic Fock space operators as c, c^\dagger .

Here, \mathbf{r} is a position corresponding to the space integral, σ is a vector of the Pauli matrices, and λ is a Yukawa coupling parameter. The ordering wave vector is chosen as $\mathbf{Q} = (\pi, \pi)$. Because of the C_4 symmetry of $\epsilon_{\mathbf{k}}$ associated with spatial rotations there are eight hot spots at which the fermion-boson scattering is resonant, Fig. 1.3b.

In its entirety, the action of the spin-fermion model reads [2, 4, 5, 32, 34, 81]

$$\begin{aligned}
S_{\text{SF}} &= S_\phi + S_\psi + S_\lambda \\
&= \int_0^\beta d\tau \int d^2k \sum_s \psi_{s\mathbf{k}}^\dagger (\partial_\tau + \epsilon_{\mathbf{k}} - \mu) \psi_{s\mathbf{k}} \\
&+ \lambda \int_0^\beta d\tau \int d^2x \sum_{s,s'} e^{i\mathbf{Q}\cdot\mathbf{r}} \phi_{\mathbf{r}} \cdot \psi_{s\mathbf{r}}^\dagger \sigma_{ss'} \psi_{s'\mathbf{r}} + \text{h.c.} \\
&+ \int_0^\beta d\tau \int d^2x \left[\frac{1}{2c^2} (\partial_\tau \phi)^2 + \frac{1}{2} (\nabla \phi)^2 + \frac{r}{2} \phi^2 + \frac{u}{4} (\phi^2)^2 \right],
\end{aligned} \tag{1.9}$$

and the associated quantum partition function is given by

$$Z_{\text{SF}} = \int D(\psi, \psi^\dagger, \phi) e^{-S_{\text{SF}}}. \tag{1.10}$$

The spin-fermion model, Eq. (1.9), should be seen as an effective low-energy theory for the onset of antiferromagnetic spin-density wave order in a metal. It describes the universal long-wavelength physics of itinerant fermions coupled to antiferromagnetic order parameter fluctuations independent of their precise origin.

Physically, the effective action S_{SF} may arise from a microscopic many-fermion Hamiltonian of the generalized Hubbard-form

$$H = \sum_{\mathbf{k},s} \epsilon_{\mathbf{k}} c_{\mathbf{k}s}^\dagger c_{\mathbf{k}s} + \sum_{\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3, \mathbf{k}_4} U_{\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3, \mathbf{k}_4}^{s_1, s_2, s_3, s_4} c_{\mathbf{k}_1 s_1}^\dagger c_{\mathbf{k}_2 s_2}^\dagger c_{\mathbf{k}_3 s_3} c_{\mathbf{k}_4 s_4}, \tag{1.11}$$

in which $U_{\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3, \mathbf{k}_4}^{s_1, s_2, s_3, s_4}$ is a generic four-fermion interaction, upon integrating out high-energy modes at short length-scales, Fig. 1.8. Formally, one introduces a characteristic energy cutoff Λ , eliminates all degrees of freedom with energies $\Lambda < E < W$, where W is the fermionic bandwidth, and expects that the resulting low-energy theory captures the essential (universal) aspects of the system [4]. Of course, such a scheme, although conceptually appealing, is generally impossible to implement in a rigorous manner [5]. In fact, it is precisely this procedure of identifying the relevant effective degrees of freedom that represents one of the central tasks and difficulties of condensed matter theory [5, 19]. To obtain the spin-fermion model, Eq. (1.9), one assumes that the integration over high energies does not produce any singularities but instead yields conventional antiferromagnetic fluctuations. Hence, in this picture, it is the magnetic structure (spin) of the high-energy fermions themselves that gives rise to the order parameter field ϕ in Eq. (1.9). Given the regularity of the integral, the fermionic and bosonic propagators¹⁶ are taken to be of conventional Fermi liquid, [5, 17, 19]

$$G_0(\mathbf{k}, i\omega_n) = \frac{1}{i\omega_n - \epsilon_{\mathbf{k}} - \mu}, \tag{1.12}$$

¹⁶The Gaussian two-point functions of the spin-fermion model, Eq. (1.9).




Figure 1.8: Sketch of the formal procedure that may generate the effective spin-fermion model, Eq. (1.9), from a microscopic many-fermion Hamiltonian, Eq. (1.11). Fermionic degrees of freedom with high energies $\Lambda < E < W$ are integrated out and, effectively, give rise to antiferromagnetic fluctuations. The resulting spin-fermion model describes the physics below an artificial energy cutoff Λ . The fermionic bandwidth W represents a natural limiting scale at high energies. Reproduced from Ref. [4].

and Ornstein-Zernike type¹⁷, [5, 87, 88]

$$\chi_0(\mathbf{q}, i\omega_n) = \frac{\alpha}{r + (\mathbf{q} - \mathbf{Q})^2 + \omega_n^2/c^2}. \quad (1.13)$$

In these expressions, ω_n denotes fermionic and bosonic Matsubara frequencies, and α is a numerical constant.

It is worth highlighting that, in a slightly different manner, one can exactly derive an action of (almost) the form of S_{SF} from a purely fermionic Hamiltonian of the kind of Eq. (1.11). Specifically, we consider

$$H_{\text{tJ}} = - \sum_{ijs} \left(t_{ij} c_{is}^\dagger c_{js} + \text{h.c.} \right) + \tilde{J} \sum_i e^{i\mathbf{Q} \cdot \mathbf{r}_i} \mathbf{S}_i^2 - \mu \sum_i n_i \quad (1.14)$$

in which the

$$\mathbf{S}_i = \frac{1}{2} \sum_{ss'} c_{is}^\dagger \boldsymbol{\sigma}_{ss'} c_{is'} \quad (1.15)$$

are the conventional spin operators and \tilde{J} is a coupling constant. The key idea is to perform a Hubbard-Stratonovich (HS) transformation in the spin channel [17, 89] by applying the exact Gaussian integral identity


$$e^{\frac{A^2}{2\alpha}} = \sqrt{\frac{\alpha}{2\pi}} \int dx e^{-\frac{1}{2}\alpha x^2 - xA} \quad (1.16)$$

with the identification $A \sim c_{is}^\dagger [\sigma]_{s,s'} c_{is'}$ to the partition function of the system (see App. A). This way, at the cost of introducing an auxiliary bosonic field ϕ , one can decouple the spin-spin interaction term and replace it by a fermion bilinear that is only quadratic in the fermionic operators, compare Eq. (1.8). Hence, one is left with non-interacting fermions moving in and coupling to a field of ϕ fluctuations. This change in perspective is analogous to viewing the direct electromagnetic electron-electron interaction as being mediated by a bosonic

¹⁷In real space this corresponds to the standard form $\chi(\mathbf{r}, 0) \sim \frac{e^{-|\mathbf{r}|/\xi}}{|\mathbf{r}|}$ in which ξ is the magnetic correlation length.

Figure 1.9:

Quasi-classical illustration of the processes included in the Hubbard model, Eq. (1.17), on a two-dimensional square lattice. Itinerant electrons (dark circles) can move between neighboring lattice sites with the hopping amplitude t . If two electrons of opposite spin (arrows) are situated on the same site they interact with coupling strength U .



photon field. Setting $\tilde{J} = -2\lambda^2/r$ this HS transformation yields precisely the spin-fermion action, Eq. (1.9), up to gradient and quartic terms in S_ϕ . Apart from the philosophical aspects of the connection between the spin-fermion model and the Hamiltonian H_{tJ} , we note that the latter can readily be diagonalized exactly for small system sizes. This will serve as an accuracy benchmark for our quantum Monte Carlo simulations.

We want to point out that for $\mathbf{Q} = 0$ there is an exact identity between the Hamiltonian H_{tJ} and the Hubbard model (Fig. 1.9) [17, 90, 91]

$$H = -t \sum_{\langle i,j \rangle, s} (c_{is}^\dagger c_{js} + \text{h.c.}) + U \sum_i n_{i\uparrow} n_{i\downarrow} - \mu \sum_i n_i \quad (1.17)$$

if one only considers nearest neighbor hopping processes, replaces the spin-operators in Eq. (1.14) according to Eq. (1.15), and absorbs contributions proportional to the total electron density operator $n = \sum_i n_i = \sum_{i,s} c_{is}^\dagger c_{is}$ into the chemical potential [91]. One finds $U = -3\tilde{J}/2 = 4\lambda^2/3r$ and $\mu = \tilde{\mu} + U/2$ where $\tilde{\mu}$ is the chemical potential of H_{tJ} . Arguably, the Hubbard model ranks among the most famous and intensively studied models of solid-state physics [17, 90–96]. Despite its conceptual simplicity, illustrated in Fig. 1.9, it describes the effect of strong correlations in itinerant many-fermion systems. Upon increasing U/t at half-filling, $\langle n \rangle/N = 1$, it displays a quantum phase transition from a (paramagnetic) metallic to an antiferromagnetic Mott-insulating ground state [93, 95, 96]. In the infinite coupling limit, $U \rightarrow \infty$, and to second order in t the fermions are effectively localized and the Hubbard model reduces to the antiferromagnetic spin-1/2 quantum Heisenberg model, [1, 17]

$$H = J \sum_{\langle i,j \rangle} \mathbf{S}_i \cdot \mathbf{S}_j. \quad (1.18)$$

Here $J \sim t^2/U$ is a positive exchange constant [17].

The AFM Mott-insulating state serves as the point of departure in many theoretical efforts [85, 97–102] to understand the rich physics of the hole-doped cuprate superconductors [4]. These studies approach the emergence of superconductivity and “strange metal” behavior by decreasing the charge carrier density of Hubbard-like models away from half-filling, that is by “doping a Mott insulator” [98]. Figuratively speaking, this corresponds to starting at $x = 0$ in Fig. 1.6 and working one’s way rightwards to optimal doping [4]. In contrast, the metallic QCP perspective taken in this thesis departs from a Fermi liquid state and hence to approaching x_{T_c} from the right side of the phase diagram, $x \gtrsim 0.3$.

1.4 Analytical insights

As mentioned in Sec. 1.2, the spin-fermion model, Eq. (1.9), has been the central focus of numerous prominent analytical studies in the last few decades [1, 3, 4, 6, 7, 15, 34–36, 47, 49, 50, 103–109]. In spite of these tremendous efforts progress has been significantly hampered by the observation that conventional Feynman graph expansions are uncontrolled and flow to strong coupling at low energies [81]. This includes various $1/N$ expansions [3, 4, 34] and extensions of the problem to fractional dimensions [6, 110]. Only very recently Lee *et al.* [6, 7, 15] have made impressive progress in handling strong correlations and gaining control over calculations. In the following we highlight some of the essential aspects of this theoretical background.

1.4.1 Fermi surface reconstruction

Based on the expectation value of the order parameter field one can conceptually distinguish two phases of the spin-fermion model: The disordered Fermi liquid state, $\langle\phi\rangle = 0$, and the antiferromagnetic spin-density wave state with $\langle\phi\rangle \neq 0$. Within mean-field theory [1, 9], we can monitor the evolution of the Fermi surface across the quantum phase transition separating these two phases, Fig. 1.10. To that end, we replace ϕ by its average value $\phi = (0, 0, \bar{\phi})$, representing Néel ordering in z-direction¹⁸ (Fig. 1.10), in which case the fermionic part of the spin-fermion model, $S_\psi + S_\lambda$, is described by the mean-field Hamiltonian

$$H_{\text{MF}} = \sum_{\mathbf{k}, s} \epsilon_{\mathbf{k}} c_{\mathbf{k}s}^\dagger c_{\mathbf{k}s} + \lambda \sum_{\mathbf{k}, s, s'} c_{\mathbf{k}+\mathbf{Q}, s}^\dagger \bar{\phi} \sigma_{ss'}^z c_{\mathbf{k}, s'} - \mu \sum_i n_i. \quad (1.19)$$

Here momenta \mathbf{k} and $\mathbf{k} + \mathbf{Q}$ are coupled together. Diagonalizing the effective 2×2 Hamiltonian of an individual momentum and spin sector one finds the energy eigenvalues [1, 9]

$$E_{\mathbf{k}}^\pm = \frac{\epsilon_{\mathbf{k}} + \epsilon_{\mathbf{k}+\mathbf{Q}}}{2} \pm \left(\left(\frac{\epsilon_{\mathbf{k}} - \epsilon_{\mathbf{k}+\mathbf{Q}}}{2} \right)^2 + (\lambda \bar{\phi})^2 \right)^{1/2}. \quad (1.20)$$

In Fig. 1.10 we illustrate this band structure and the resulting transformation of the Fermi surface, obtained by filling up the energy states for different order parameter magnitudes $\bar{\phi}$. Upon developing antiferromagnetism, $\bar{\phi} \neq 0$, energy gaps appear and open up the Fermi surface at the hot spots. Concomitantly, particle and hole pockets associated with the two energy bands $E_{\mathbf{k}}^\pm$ emerge. Depending on the value of the chemical potential only one sort of pockets survives for $\bar{\phi} = 0.9$ before eventually being gapped out as well, $\bar{\phi} \gg 1$. In Fig. 1.10 we show the electron-doped case $\mu > -1$ (half-filling corresponds to $\mu = -1$). Experimentally, the reconstruction of the Fermi surface has been observed in, among others, ARPES [71] and quantum oscillation measurements [111] of the electron-doped cuprate $\text{Nd}_{2-x}\text{Ce}_x\text{CuO}_4$. It has further been demonstrated numerically that the phenomenology of Fig. 1.10 is correct beyond a mean-field treatment [32, 81].

¹⁸Because of the $O(3)$ symmetry, all order parameter directions are equivalent and we are free to choose ϕ to point in z-direction.

¹⁹For illustration purposes, we discretize ϕ and take it to be two-dimensional.

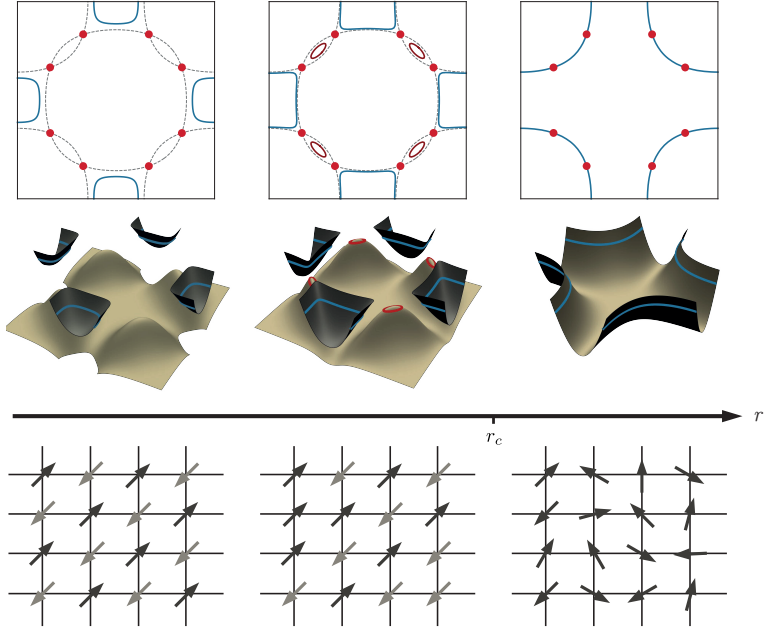


Figure 1.10: Reconstruction of the Fermi surface (solid lines) of a metal undergoing a SDW quantum phase transition. As the bosonic order parameter field ϕ (arrows¹⁹ in the bottom row) develops a finite expectation value, FS gaps open up at the hot spots (red points) and electron and hole pockets form. The three-dimensional plots in the second row show the energy dispersions $\epsilon_{\mathbf{k}}$, Eq. (1.7), (one band) and $E_{\mathbf{k}}^{\pm}$, Eq. (1.20), (two bands) of the disordered ($r > r_c$) and ordered ($r < r_c$) phase respectively. We have set $\lambda = 1$ and $\bar{\phi} = 0.9, 0.4, 0$ from left to right.

1.4.2 Hertz-Millis theory

The traditional approach to AFM metallic quantum criticality is due to Hertz and Millis [35, 50, 82]. Departing from the effective low-energy description, the strategy is to integrate out the fermions and keep only the leading order contribution to the bosonic self-energy. Evaluating the corresponding Feynman diagram, Fig. 1.11a, one finds that the propagator of the bosonic field ϕ takes the renormalized form [1, 34]

$$\chi^{-1}(\mathbf{q}, i\omega_n) = \chi^{-1}(\mathbf{q}, i\omega_n)_0 - \Pi(\mathbf{q}, i\omega_n) \quad (1.21)$$

$$= \tilde{r} + \mathbf{q}^2 + \frac{\omega_n^2}{c^2} + \gamma|\omega_n| \quad (1.22)$$

indicating Landau damping of the order parameter modes. Specifically, \tilde{r} is the renormalized tuning parameter and

$$\gamma = \frac{N\lambda^2}{\pi v_F^2 \sin \theta} \quad (1.23)$$

is a damping constant in which $N = 4$ is the number of hot spot pairs and $\theta \in [0, \pi]$ is the angle between the Fermi velocities (of magnitude v_F) at opposite hot spots related by \mathbf{Q} . The Hertz-Millis action at the metallic QCP then reads [1, 34, 50]

$$S_H = \int \frac{d^2q}{(2\pi)^3} T \sum_{\omega_n} \frac{1}{2} (\tilde{r} + \mathbf{q}^2 + \gamma|\omega_n|) |\phi|^2 + \frac{u}{4} \int d\tau d^2x (\phi^2)^2 + \dots \quad (1.24)$$

One proceeds by conducting a naive RG analysis. Because of the emergence of the Landau-damping term $\gamma|\omega_n|$ the quadratic part of S_H shows a scaling $\omega \sim q^2$. The dynamical critical exponent z has therefore increased from $z = 1$ (at the Wilson-Fisher fixed point) to $z = 2$ [1, 3, 4, 34, 112]. This implies that the effective dimension of the theory is $d + z = 4$ and the quartic coupling u is marginal. By power counting, all couplings u_n associated with higher-order terms ϕ^{2n} , indicated by the ellipses in Eq. (1.24), are irrelevant and may be neglected [34]. Hence, up to logarithmic corrections, the Hertz-Millis action seems to be governed by the Gaussian fixed point.

As has been pointed out by Abanov *et al.* in Ref. [36] this conclusion is incorrect [1, 4, 34]. The key shortcoming of the Hertz-Millis analysis is that it treats the couplings u_n as mere factors. Instead they turn out to be singular, non-local functions $u_n(\mathbf{q}, i\omega_n)$. Upon closer inspection, the dangerous procedure of integrating out the gapless fermionic degrees of freedom generates an infinite number of terms which all turn out to be marginal [36]. Consequently, the Gaussian fixed point is unstable.

A further cause for concern in the Hertz-Millis approach is that the feedback of the damped magnetic fluctuations on the fermions is entirely neglected. As we will see below, one generally finds that Fermi liquid theory breaks down because of this interplay. This raises questions about the self-consistency of the program above which started by integrating out an assumed Fermi liquid [39].

Despite the fact that Hertz-Millis theory is formally flawed some of its characteristic features, such as the Landau damping, will, remarkably, be confirmed by




Figure 1.11: Leading order contributions to the boson (a) and fermion (b) self-energies Π and Σ as well as the first order correction to the boson-fermion vertex (c). Solid and wavy lines indicate fermion and boson propagators G_0 and χ_0 , respectively. Black dots represent the bare Yukawa vertex λ .

our simulations of the AFM QCP presented in Ch. 5. Similar observations have been made in numerical studies of easy-axis and easy-plane antiferromagnetism [31–33]. Clearly, a controlled theory to unravel this mystery would be highly desirable.

1.4.3 Hierarchy of energy scales

To identify a number of crossover regimes and to understand the basic impact of the spin-fermion interactions, it is instructive to perform an analysis for small Yukawa coupling λ – despite the fact that, strictly speaking, the Yukawa coupling is a relevant parameter in the sense that conventional perturbation theory becomes unreliable at small energy scales [7]. Our discussion expands on the short review by the author of this thesis published in Ref. [P1] (Section III) and is largely inspired by Ref. [2].

Focusing on the antiferromagnetic order parameter first, we can extract an energy scale Ω_b , at which the nature of the dynamics of collective boson excitations changes due to interactions with surrounding fermions. As in Hertz-Millis theory, magnetic fluctuations can decay into particle-hole pairs, Fig. 1.11a, and, to leading order in λ , the propagator of the bosonic field ϕ takes the renormalized form of Eq. (1.22) [8]. By estimating when the dynamically generated contributions to the action (ω_n) dominate over bare ones (ω_n^2/c^2) one can identify the energy scale [2–4, 34, 112]

$$\Omega_b = \frac{N\lambda^2 c^2}{\pi v_F^2 \sin \theta}. \quad (1.25)$$

At frequencies small compared to Ω_b the dynamics of the order parameter modes is overdamped with a dynamical critical exponent $z = 2$ [1, 3, 4, 34, 112].

Next, we consider the lowest-order effect of the renormalized bosons on the fermion dynamics at the hot spots. To leading order in λ , Fig. 1.11b, one finds

[3, 34] that, at zero temperature, the fermions acquire a self-energy

$$\Sigma(i\omega_n) \sim \frac{in_b\lambda\sqrt{\sin\theta}}{\sqrt{N}}\sqrt{|\omega_n|}, \quad (1.26)$$

where $n_b = 3$ is the number of order parameter components. This implies that the fermions become strongly damped by the coupling to SDW fluctuations, with a damping rate that scales as $\sqrt{\omega}$, indicating a distinct deviation from ordinary Fermi liquid character ($\omega^2 \log(1/\omega)$ in two spatial dimensions). Taking the same approach as above, we estimate the energy scale Ω_f at which this breakdown occurs as

$$\Omega_f \sim \frac{n_b^2\lambda^2 \sin\theta}{N}. \quad (1.27)$$

For frequencies $\omega \ll \Omega_f$ the feedback from “dressed” bosons on fermion excitations at the hot spots is strong and leads to quasi-particle decoherence. Away from the hot spots FL theory survives and fermions remain well-defined.

Apart from entering a non-Fermi liquid regime, the system may eventually become unstable against pairing at sufficiently low temperatures [34, 113–115]. Close to the critical point collective spin fluctuations can generate attractive interactions, facilitate the formation of Cooper pairs, and enhance d-wave superconductivity [7, 15, 30, 32, 42, 52, 81, 116, 117]. At the hot spots, and for small λ , one can estimate the energy scale T_c of the onset of superconductivity within Eliashberg theory [4, 26, 112]. Working in the Landau-damped regime, i.e. using the boson propagator of Eq. (1.22), one solves self-consistently for the fermion self-energy at one-loop level: the augmented Green’s function

$$G^{-1}(\mathbf{k}, i\omega_n) = G_0^{-1}(\mathbf{k}, i\omega_n) - \Sigma^1(\mathbf{k}, i\omega_n) \quad (1.28)$$


is used in the computation of Σ^1 , the first order diagram shown in Fig. 1.11b. Vertex corrections, such as the diagram in Fig. 1.11c, are neglected. By parameterizing the self-energy in the Nambu spinor basis $\Psi_{\mathbf{k}}^\dagger = (\psi_{\mathbf{k}\uparrow}^\dagger, \psi_{-\mathbf{k}\downarrow}^\dagger)$ [17] as

$$\Sigma(\mathbf{k}, i\omega_n) = (1 - Z_{\mathbf{k}}^{-1})i\omega_n + \Delta_{\mathbf{k}}\sigma_1, \quad (1.29)$$

one obtains Eliashberg equations for the quasiparticle residue $Z_{\mathbf{k}}$ and the superconducting gap function $\Delta_{\mathbf{k}}$. Solving the latter – in the hot spot approximation [4, 112] – for the spin-fermion model predicts the superconducting susceptibility to diverge at a scale [112]

$$T_c \sim \frac{n_b^2\lambda^2 \sin\theta}{N}. \quad (1.30)$$

We summarize the hierarchy of energy scales obtained from a small λ analysis in Fig. 1.12. Note that the energy scales T_c and Ω_f are both of the order $O(\lambda^2/N)$ and there is, a priori, no parametrically large separation between a non-Fermi liquid and a superconducting regime. Similarly, note that the scales Ω_b and Ω_f are both of order λ^2 in the Yukawa coupling. This is in stark contrast to the case of an Ising-nematic QCP (where $\Omega_b \sim \lambda$ and $\Omega_f \sim \lambda^4$ [2]) which shows a

**Figure 1.12:**

Scaling regimes suggested by perturbation theory. The energy scales Ω_b , Ω_f , and T_c correspond to Eqs. (1.25), (1.27), and (1.30), respectively. There is a parametrically large (in the number of hot spot pairs N) Landau damped regime. For the spin-fermion model (1.9) with a Fermi surface as in Fig. 1.10, $N = 4$. Resized from Ref. [P1] under CC-BY 4.0.

parametrically large window of Landau-damped physics (with $z = 3$) [39, 40]. However, the energy scales are separated by their dependence on the number of hot spot pairs, $N = 4$, which may serve as a control parameter for an extended Fermi liquid regime with relaxational boson dynamics.

Finally, we inspect first-order corrections $\delta\Gamma$ to the spin-fermion vertex, Fig. 1.11c. Employing the same procedure as around Eqs. (1.25) and (1.27), one can extract an energy scale Ω_λ at which the dynamically generated higher-order interactions become comparable to the bare vertex, $\delta\Gamma \sim \lambda$. One finds [34]

$$\Omega_\lambda = \Omega_{\text{uv}} \exp\left(-\frac{\pi N}{\pi - \theta}\right), \quad (1.31)$$

where Ω_{uv} is an ultra-violet cutoff and $0 < \theta < \pi$. We note that the scale Ω_λ does not depend on the coupling strength λ – the naive λ^2 is compensated by the Landau-damping constant γ , Eq. (1.23). Instead, for a fixed number of hot spot pairs, the scale Ω_λ is only sensitive to the relative angle θ between the Fermi surfaces at the hot spots. We illustrate the dependence of Ω_λ on θ for a set of two-band Fermi surfaces, which we will describe and study in Ch. 5, in Fig. 1.13. In particular, we see a moderate increase of Ω_λ in the limit of (local) nesting, $\theta \rightarrow 0$.

Let us point out that Eq. (1.31) is special to the case of an $O(3)$ symmetric SDW order parameter [7]. Generically, the leading order correction to the boson-fermion vertex scales as $\delta\Gamma \sim 2 - n_b$ and vanishes precisely for easy-plane antiferromagnetism with $n_b = 2$. This can be directly deduced from Fig. 1.11c by noting that the Pauli matrices, sitting at the three bare vertices in the diagram, appear in the combination $\sigma_a \sigma_b \sigma_a$ – implicitly assuming Einstein’s sum convention. Using the basic properties $\{\sigma_a, \sigma_b\} = 2\delta_{ab}$ and $\sigma_a^2 = \mathbb{1}$ one immediately finds

$$\delta\Gamma \sim (2\delta_{ab} - \sigma_b \sigma_a) \sigma_a = 2\sigma_b - n_b \sigma_b \sim 2 - n_b. \quad (1.32)$$




Figure 1.13: Effect of nesting on the vertex energy scale Ω_λ given in Eq. (1.31). Here, θ is the relative angle between the Fermi velocities (as indicated in the insets). Perfect local nesting corresponds to $\theta = 0$ and anti-nesting to $\theta = \pi$. Resized from Ref. [P1] under CC-BY 4.0.

In fact, understanding that the argument above is more general, it is actually an infinite series of “ladder-like” vertex corrections that vanishes for $n_b = 2$. This indicates a qualitative difference in the perturbative structure of Heisenberg and XY antiferromagnetism in itinerant electron systems [7].

1.4.4 Emergence of a control parameter

Despite the general guidance provided by the perturbative treatment of the Yukawa coupling, this approach is formally uncontrolled [2, 4, 34, 50] and must be complemented by a more sophisticated analysis. The main route that has been pursued in the last decades [4, 34, 118] is to self-consistently solve for the Landau damping and fermion self-energy, similar to the Eliashberg approach above [26, 112]. In the generalized case of $N_f = \infty$ fermion flavors, this confirms the essential features of the perturbative analysis such as the hierarchy of energy scales in Fig. 1.12. However, departing from those extended theories, all attempts to taking the physical limit $N_f \rightarrow 1$ turned out to be logarithmically unstable [118]. In particular, it was observed [3, 4, 34, 103, 110] that, upon approaching low energies, the Fermi velocity \mathbf{v}_F gets renormalized in such a way that the Fermi surface near the hot spots flows towards local nesting, $\theta \rightarrow 0$ (see Fig. 1.12). This dynamically induced change of the FS shape is accompanied by the appearance of anomalous dimensions [34]. Furthermore, related to the renormalization of the damping factor γ in a $1/N_f$ expansion [34, 118], the dynamical critical exponent was found to shift away from the naive prediction $z = 2$.

Given the unreliability of the extension to $N_f = 1$, a novel theory that bridges this gap in a controlled manner is in great demand. A huge leap forward in this

direction has recently been made in form of the groundbreaking work by Lee *et al.* in Refs. [6, 7, 15]. Rather than treating the fermion-boson interaction as a perturbation, they employ an interaction-driven scaling scheme [119] in which the Yukawa coupling and the fermion kinetics are kept marginal – the irrelevance of the boson kinetics is justified by assuming that dynamical damping dominates in the low energy limit. Remarkably, by conducting a self-consistent analysis [7], Lee *et al.* managed to identify an emergent control parameter: the ratio v/c where v is the component of \mathbf{v}_F perpendicular to the ordering wave vector $\mathbf{Q} = (\pi, \pi)$ and c is the boson velocity. They could show that, to leading order in v/c , a (self-consistent) diagrammatic expansion is dominated by an infinite fractal series of Feynman diagrams, based on the two-loop boson and fermion self-energies [7]. Exploiting this structure, Lee *et al.* were able to derive a non-perturbative solution for the spin-fermion model which reveals a novel “strong coupling” fixed point at low energies.

Notably, their solution is characterized by heavily “dressed” bosonic modes [7]. At the “strong coupling” fixed point, the propagator of the order parameter fluctuations takes the following form²⁰ (up to logarithmic corrections), [6, 7]

$$\chi^{-1}(\mathbf{q}, i\omega_n) \sim \gamma' |\omega_n| + |q_x + q_y| + |q_x - q_y|, \quad (1.33)$$

in which γ' is a parameter. Note that the correlations of the bosons are highly anisotropic. Furthermore, time and space scale identically and the dynamical critical exponent is $z = 1$. Interestingly, because of this decreased value of z [118], the fermions remain coherent, even at the hot spots [7]. At low temperatures, Lee *et al.* find that the AFM fluctuations, Eq. (1.33), promote pairing and d-wave superconductivity. However, similar to BCS theory, the critical temperature T_c is exponentially small in $\sqrt{c/v}$ and, thus, only mildly enhanced. Compared to Fig. 1.12, these findings naively suggest the modified hierarchy of energy scales (we use tilde signs to indicate altered scales),

$$\tilde{\Omega}_f \ll \tilde{T}_c \ll \tilde{\Omega}_b. \quad (1.34)$$

Here, above the onset of superconductivity, there is a parameterically large $z = 1$ regime in which the SDW susceptibility is described by Eq. (1.33).

Arguably, a confirmation of those predictions with an unbiased method would be highly desirable and is one of the motivations for the numerical study reported in this thesis. Practically, a central question is whether the “strong coupling” fixed point is generic for an antiferromagnetic SDW transition in a metal, in the sense of having an extensive basin of attraction, or rather nonuniversal and limited to a narrow parameter range. In particular, the energy window for the $z = 1$ scaling might strongly depend on the bare value of v and require careful fine tuning [7, 118]. Previous numerical studies of metallic easy-plane antiferromagnetism have not found any indication for a $z = 1$ regime [31, 32]. To the end of answering questions like the above, we will focus on both “generic” as well as designed (almost) locally nested Fermi surfaces in our quantum Monte Carlo simulations in Ch. 5.

²⁰We note that in comparison to Eq. (16) of Ref. [7] we have $|q_x + q_y| + |q_x - q_y|$ in Eq. (1.33) instead of $|q_x| - |q_y|$. This is due to the fact that Ref. [7] is working with a coordinate system that is rotated (clockwise) by $\pi/4$, i.e. $\mathbf{k}_x \sim \mathbf{Q}$.

2 | Determinant quantum Monte Carlo

In this chapter, we describe the finite-temperature determinant quantum Monte Carlo (DQMC) method for numerically simulating itinerant many-fermion systems. We put our focus on quantum field theories that are of the form of the spin-fermion model, Eq. (1.9), where interactions are of Yukawa type and mediated by a bosonic field. Starting with a general presentation of the Monte Carlo technique for classical degrees of freedom we review the concept of a Markov process and discuss the famous Metropolis algorithm. Applying the quantum-classical mapping, Sec. 1.1, we transfer these techniques to quantum systems of itinerant fermions and derive the simulation technique underlying the results presented in Chs. 5 and 7. The key feature of the resulting DQMC method is that it is numerically exact – given sufficient computation time the systematic errors can be made arbitrarily small – and does not rely on the smallness of any expansion parameter. Furthermore, provided the absence of the famous sign problem, Ch. 3, it allows one to explore the relevant regions of exponentially large configuration spaces in polynomial time. Postponing a discussion of more technical implementation details such as numerical stabilization methods and efficient linear algebra computations to Ch. 4, statistical data analysis considerations will conclude this chapter. For more comprehensive introductions into quantum Monte Carlo (QMC) methods we refer the interested reader to Refs. [90, 120–123] which have served as valuable resources for this thesis.

2.1 Solving classical thermodynamics at random

Suppose we have a physical system with a finite number of classical, microscopic degrees of freedom. Arguably, the most prominent example would be the Ising model, named after the Cologne-born physicist Ernst Ising, in which classical binary spins $s \in \{\uparrow, \downarrow\}$ are situated on a two-dimensional lattice. Generally, we are interested in calculating the thermodynamic expectation value of an observable,

$$\langle X \rangle = \int_{\Omega} d\mathcal{C} X(\mathcal{C}) p_{\mathcal{C}}^B, \quad (2.1)$$

in which Ω denotes the d_Ω -dimensional space²¹ of all possible configurations of the microscopic variables²² and $X(\mathcal{C})$ is the observable evaluated for an element $\mathcal{C} \in \Omega$. For the Ising model, $\mathcal{C} = \{s_i\}$ could be a particular arrangement of the spins s_i and $X(\mathcal{C}) = M(\{s_i\})$ the corresponding magnetization. Given a fixed temperature, each configuration \mathcal{C} has a certain probability of occurrence in the canonical statistical ensemble²³. Therefore, the observable values $X(\mathcal{C})$ are weighted by Boltzmann probabilities

$$p_{\mathcal{C}}^B = \frac{\exp(-\beta E(\mathcal{C}))}{\mathcal{Z}}. \quad (2.2)$$

Here, $\beta = 1/T$ is inverse temperature, $E(\mathcal{C})$ is energy, and

$$\mathcal{Z} = \int_{\Omega} d\mathcal{C} \exp(-\beta E(\mathcal{C})) \quad (2.3)$$

is the partition function of the system, which serves as a normalization factor such that $\int_{\Omega} d\mathcal{C} p_{\mathcal{C}}^B = 1$.

Unfortunately, only very rarely can the integral in Eq. (2.1) be solved exactly. In fact, even computing the partition function, and therefore the Boltzmann probabilities, is typically intractable and one must rely on approximations and numerical methods to make progress. One idea that may come to mind is to use standard integration techniques such as Simpson's quadrature procedure. However, in this case the numerical error Δ_S scales with the number of points N , at which the integrand is explicitly evaluated, as $O(N^{-4/d_\Omega})$ [124]. Combined with the fact that the dimension d_Ω of the configuration space of a (quantum) statistical system typically grows exponentially with system size – a direct consequence of the tensor product property $\dim(H_1 \otimes H_2) = \dim H_1 \dim H_2$ – the numerical error quickly approaches unity for fixed N .

In this thesis, we will therefore take a different route for computing the expectation value $\langle X \rangle$ known as statistical sampling. The strategy is to choose the integrand evaluation points $\{\mathcal{C}_i \in \Omega\}$ randomly, for now according to a uniform distribution, such that the simple estimator

$$\bar{X} = \frac{1}{N} \sum_{i=1}^N X(\mathcal{C}_i) \quad (2.4)$$

converges to the desired integral, Eq. (2.1), by the law of large numbers. While we will postpone a more detailed discussion of the numerical error Δ_{MC} associated with the approximation $\langle X \rangle \approx \bar{X}$ to Sec. 2.6, it is understood that it is only of statistical nature and therefore famously goes as $\Delta_{\text{MC}} \sim O(N^{-1/2})$.

²¹Note that Ω , being the *sampling space* of configurations, is generally distinct from the quantum mechanical Hilbert space. For the spin-fermion model, Eq. (1.9), it is the space of the bosonic order parameter ϕ .

²²We implicitly assume that Ω is countable, although possibly infinite. This is not really a restriction since one can always choose an arbitrarily fine grid in the computation of Eq. (2.1).

²³Although we focus on the canonical ensemble in this general discussion we will later use the grand canonical ensemble with $p_{\mathcal{C}}^B = \exp(-\beta[E(\mathcal{C}) - \mu N(\mathcal{C})])/\mathcal{Z}$, where μ is the chemical potential and N is particle number.

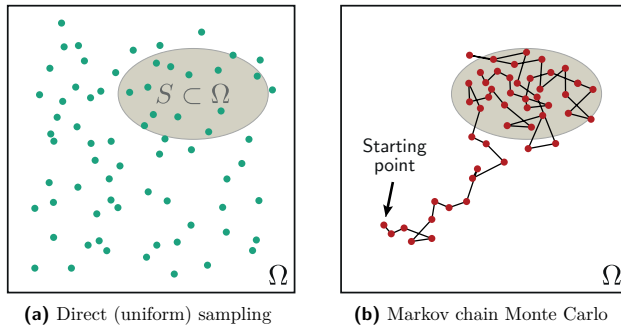


Figure 2.1: Schematic²⁴ visualization of direct statistical sampling (a) and Markov chain Monte Carlo (b). The surrounding box represents the, typically high-dimensional, configuration space Ω . The statistically relevant subset $S \subset \Omega$ with significant Boltzmann weight is illustrated by the gray ellipse. Configurations are independently drawn from a uniform distribution (a) and by performing a stochastic walk through Ω generating a Markov chain (b).

Importantly, it is independent of d_Ω and hence, from a scaling point of view, is expected to outperform Simpson’s quadrature for physical systems with $d_\Omega > 8$.

Eq. (2.4) represents the fundamental principle underlying all Monte Carlo methods which, due to their random nature, are named after the *Casino de Monte-Carlo* in Monaco. The goal of calculating the expectation value $\langle X \rangle$, Eq. (2.1), is reduced to two conceptually simple steps:

- Randomly choosing configurations $\{\mathcal{C}_i \in \Omega\}$.
- *Measuring* the observable X by computing and averaging $\{X_i = X(\mathcal{C}_i)\}$.

Even though we will be drastically adjusting the details of the sampling procedure in the following sections, the statistical nature will remain the key property allowing us to efficiently compute expectation values of observables.

2.2 Markov chain Monte Carlo

In spite of the superior scaling compared to standard numerical integration techniques, statistical sampling is not necessarily preferable to Simpson’s method. The most important reason for this apparent contradiction is that the scaling of the numerical error doesn’t make any statements about its magnitude. Given that valuable information about the specific integrand enters the regular Simpson’s construction, whereas the same is completely disregarded in the naive Monte Carlo approach, one should expect that the numerical prefactor in the error function is generally smaller in the former case. To obtain a result of similar

²⁴Note that, for simplicity, we assume that there is a rather sharply distinguished statistically relevant subset $S \subset \Omega$. In general the distribution may be much more complex.

accuracy, this quantitative difference needs to be compensated by a potentially much larger number of intermediate function evaluations overshadowing the *a priori* scaling advantage in a practical computation.

Fortunately, we can lift this limitation without sacrificing the scaling advantage. The key idea is to account for the non-uniform structure of the integrand I in an importance sampling process. Physically, the statistical ensemble describing the thermodynamics of the physical system often only covers a fraction of the exponentially large configuration space Ω because of the Boltzmann suppression present in Eq. (2.1). In an efficient computation of $\langle X \rangle$ one should therefore focus the sampling on those regions $S(p_{\mathcal{C}}^B) \subset \Omega$ that contribute with a finite thermodynamic weight. As indicated in Fig. 2.1a, naively sampling the integrand from a uniform distribution leads to many unnecessary function evaluations in sub spaces with $I \approx 0$. Unfortunately, the statistical importance of a configuration $\mathcal{C} \in \Omega$ is given by the generally inaccessible Boltzmann probability $p_{\mathcal{C}}^B$. The question therefore is: How can one sample configurations from an unknown distribution?

The answer to this question is Markov chain Monte Carlo (MCMC), which comes with a qualitative change of the sampling strategy. So far, we have considered sampling arrangements of the microscopic degrees of freedom independently from a uniform distribution. Figuratively speaking, this amounts to selecting configurations $\{\mathcal{C}_i\}$ by randomly throwing darts [125], Fig. 2.1a. In MCMC this direct sampling is replaced by a stochastic Markov walk through the space of configurations, as illustrated in Fig. 2.1b. Clearly, to realize a sampling by importance the navigation through Ω must, in some way, follow a set of rules reflecting the non-uniform distribution of the statistical weight. Leaving the formulation of an explicit algorithm to the next section, let us briefly discuss the concept of MCMC sampling in more abstract terms.

A Markov walk is a stochastic process that generates a sequence of configurations $\{\mathcal{C}_n \in \Omega\}$ associated with random variables²⁵ $\{C_n\}$ under the condition [120, 125]

$$P(C_n = \mathcal{C}_n | C_{n-1} = \mathcal{C}_{n-1}, \dots, C_1 = \mathcal{C}_1) = P(C_n = \mathcal{C}_n | C_{n-1} = \mathcal{C}_{n-1}), \quad (2.5)$$

that is the probability $P(C_n = \mathcal{C})$ in step n of the walk is only effected by C_{n-1} immediately preceding it. Consequently, we can describe the process by a stochastic matrix of transition probabilities,

$$W_{CC'} = P(C_n = \mathcal{C} | C_{n-1} = \mathcal{C}'). \quad (2.6)$$

Being a stochastic matrix, all entries are non-negative and furthermore

$$\sum_{\mathcal{C}} W_{CC'} = 1, \quad \forall \mathcal{C}' \in \Omega, \quad (2.7)$$

indicating column-wise normalization. Given a probability distribution $p_{\mathcal{C}'}^{(n-1)} = P(C_{n-1} = \mathcal{C}')$, a single step of the Markov walk can be expressed as a simple

²⁵For each value n , the random variable C_n models the statistical distribution of the possible values for \mathcal{C}_n .

matrix-vector product,

$$p_C^{(n)} = \sum_{C'} W_{CC'} p_{C'}^{(n-1)}. \quad (2.8)$$

The usefulness of Markov processes now stems from the powerful insight that, under quite general conditions, repeated application of Eq. (2.8) leads to convergence to a stationary distribution \tilde{p}_C [120],

$$\lim_{n \rightarrow \infty} p_C^{(n)} = \tilde{p}_C. \quad (2.9)$$

Remarkably, this convergence is independent of the starting distribution $p_C^{(1)}$ – we are hence free to start our Markov walk at any point $C_1 \in \Omega$. Combining Eqs. (2.8) and (2.9), one finds the stationary balance

$$\tilde{p}_C = \sum_{C'} W_{CC'} \tilde{p}_{C'}, \quad (2.10)$$

rendering \tilde{p}_C an eigenvector of the transition matrix $W_{CC'}$ with eigenvalue 1.

With regard to our initial goal of using the Monte Carlo estimator of Eq. (2.4) to compute the thermodynamic expectation value of an observable X the remaining question is clear: Can we make the Markov process converge specifically to the Boltzmann distribution of Eq. (2.2), i.e. $\tilde{p}_C = p_C^B$? Fortunately, the answer is positive. Using the column-wise normalization, Eq. (2.7), it is straightforward to check that the celebrated detailed-balance condition [120, 125, 126],

$$\frac{W_{CC'}}{W_{C'C}} = \frac{\tilde{p}_C}{\tilde{p}_{C'}}, \quad (2.11)$$

is sufficient for the stationary balance, Eq. (2.10). By simply setting \tilde{p}_C to the desired Boltzmann probability p_C^B on the right hand side of the equation we obtain a criterion for how to choose the transition probabilities $W_{CC'}$ such that the stationary limit of the associated Markov process will be guaranteed to converge to p_C^B . It is important to note that although the probabilities $\tilde{p}_C, \tilde{p}_{C'}$ themselves are only formally known, Eq. (2.2), the ratio $\tilde{p}_C/\tilde{p}_{C'}$ is readily computable since the normalization, the unknown partition function \mathcal{Z} , cancels out. Similarly, although the sampling space Ω may be (exponentially) large fulfilling detailed-balance is simple as it only relates two elements $C, C' \in \Omega$.

In summary, we have established that, under the detailed-balance condition, a long enough Markov walk through the space of configurations Ω will sample elements according to the unknown Boltzmann distribution. The corresponding Markov chain²⁶ of passed configurations $\{C_t \in \Omega\}$ therefore becomes a discrete approximation of the statistical ensemble. Combined with the measurement step in Eq. (2.4), this allows us to compute thermodynamic expectation values efficiently.

²⁶In physics literature, the term Markov chain is commonly used to refer to the sequence of configurations rather than the Markov process itself. We will adopt this convention in this thesis.

2.3 The Metropolis algorithm

The detailed-balance condition captures the essential idea of the MCMC scheme. However, given its abstract form, Eq. (2.11), it is still unclear how to practically apply it in a Monte Carlo simulation. In a seminal paper [127], Metropolis *et al.* at the Los Alamos scientific laboratory famously closed this gap in 1953 by putting forward what is now known as the *Metropolis algorithm*. The presented “general method, suitable for fast electronic computing machines, of calculating the properties of any substance” [127] has by today become a general purpose technique across many disciplines, such as chemistry, economics, and social sciences. According to the first millennium issue of *Computing in Science & Engineering*, it is the most influential algorithm of the 20th century [128, 129]. Over time, the algorithm has been refined and extended by many scientists, most importantly by Hastings in 1970 [130].

Metropolis *et al.* proposed to split up the transition probability matrix into two contributions [120, 127],

$$W_{C'C} = T_{C'C} A_{C'C}. \quad (2.12)$$

Here, $T_{C'C}$ is a symmetric stochastic matrix representing transition proposal probabilities and $A_{C'C}$ is an transition acceptance matrix with elements

$$A_{C'C} = \min \left\{ 1, \frac{\tilde{p}_{C'}}{\tilde{p}_C} \right\}. \quad (2.13)$$

It is straightforward to verify that this choice of $W_{C'C}$ respects the detailed-balance condition which, using the symmetry of T , simplifies to

$$\frac{A_{CC'}}{A_{C'C}} = \frac{\tilde{p}_C}{\tilde{p}_{C'}}. \quad (2.14)$$

Remarkably, this indicates that the stationary distribution \tilde{p}_C is (almost) completely independent of how we choose the transition probabilities $T_{C'C}$ – they only need to be probabilities in the sense that they are non-negative and (column-wise) sum up to unity.

Translating this rather formal analysis into an explicit procedure, the Metropolis algorithm for a Markov walk through the space of configurations is as follows:

1. Given the current configuration $C \in \Omega$, propose a move (a transition) to a different configuration $C' \in \Omega$.
2. Compute the acceptance probability $A_{C'C} = \min \left\{ 1, \frac{\tilde{p}_{C'}}{\tilde{p}_C} \right\}$.
3. Draw a random number $\alpha \in [0, 1]$ from a uniform distribution.
4. Accept the move if $\alpha \leq A_{C'C}$, otherwise reject it.
5. Start over.




Figure 2.2: Illustration of a local move for a two-dimensional field $\mathcal{C} \equiv \phi$ of $O(2)$ degrees of freedom. For more details, see Fig. 4.9.

For the thermodynamic Boltzmann distribution, Eq. (2.2), the probability ratio in $A_{\mathcal{C}'\mathcal{C}}$ is given by the exponential energy difference $\exp(-\beta(E(\mathcal{C}') - E(\mathcal{C})))$ of the current and proposed configuration. This implies that our Markov walker generally tends to move downhill in the energy landscape. However, at finite temperature, $\beta < \infty$, there is a non-vanishing probability for transitioning into a state of higher energy as a consequence of thermal fluctuations.

2.3.1 Proposing moves

As seen in Eq. (2.14), the basic Metropolis algorithm leaves a lot of freedom in proposing moves $\mathcal{C} \rightarrow \mathcal{C}'$, which asks for a bit of structuring. First, we have so far neglected the (few) formal preconditions [120] of the Markov process convergence, Eq. (2.9), the most important of which is *irreducibility*. From Fig. 2.1b, it is intuitively clear that only proposing moves for, say, going back and forth between two configurations $\mathcal{C}_A, \mathcal{C}_B$ will be pathological in that the Markov walker could never reach the statistically relevant region $S \subset \Omega$. Irreducibility, which in physics is often also referred to as ergodicity²⁷, expresses the idea that every configuration \mathcal{C}' must be reachable from every other configurations, by one or multiple transitions. Typically, this is ensured by consecutively proposing a sweep of *local moves* in which every degree of freedom is modified individually and only slightly²⁸, Fig. 2.2. Being, in a sense, the “quantum of moves” it is plausible that this strategy will, in principle, make every configuration $\mathcal{C} \in \Omega$ accessible. The same picture, however, also reveals a major disadvantage of local updates in cases where the equilibrium distribution is spread out in configuration space. To travel a larger distance in the latter, that is to go from a configuration $\mathcal{C} \in \Omega$ to an entirely different element $\mathcal{C}' \in \Omega$ one generally needs to perform multiple successive sweeps. While Eq. (2.9) still guarantees the convergence of the Markov walk in such a case, it doesn’t make any statement about how long – we interpret

²⁷Technically, ergodicity is a statement about the equilibrium distribution while irreducibility addresses the Markov transitions itself [120].

²⁸If the degree of freedom is continuous, the proposed value is drawn from a uniform distribution.

a sweep as the unit of *Monte Carlo time* – it takes to reach the stationary limit. A physical scenario where this slowdown is especially pronounced is given by a first order phase transition, for instance the transition between the two fully polarized Ising states - where all spins point upwards or downwards - as a function of an external magnetic field. In this case, transitioning between the separate regions of Ω associated with coexisting qualitatively distinct phases A and B by a sequence of only minor local modifications is inefficient and, according to the Metropolis acceptance, unlikely, since intermediate configurations lie higher in energy. Similarly, in the vicinity of a (second-order) quantum critical point the divergence of correlation length and correlation time, Eqs. (1.1) and (1.2), implies a critical slowing down of local Monte Carlo algorithms [1, 120, 122, 126, 131]. To counter both of these effects, we supplement sweeps of local moves by periodic *global updates* in which all degrees of freedom (or a significant subset) are modified collectively. For the Ising model in particular, various forms of cluster updates have been constructed and proven to improve the Monte Carlo convergence near the ferromagnet-paramagnet phase transition dramatically [132, 133]. However, for more sophisticated systems no such designed updates are available and we implement the general shifting scheme described in Sec. 4.2.2.

2.4 Determinant quantum Monte Carlo

With respect to the main focus of this thesis, the numerical simulation of the spin-fermion model Eq. (1.9), our previous discussion raises a central question: How to apply the stochastic Markov sampling idea to a quantum system of itinerant fermions coupled to an order parameter? More specifically asked, is it possible to identify a *classical*²⁹ sampling space Ω for the quantum system such that the relevant quantum partition function takes the simple form of Eq. (2.3)? Fortunately, the answer is positive. In fact, there are multiple different ways to rewrite the partition function of a quantum system such that it becomes amenable to classical Monte Carlo. Famous examples include World line [122] and continuous time quantum Monte Carlo [134]. In this thesis we focus on the determinant quantum Monte Carlo flavor [90, 122, 135–137] which is based on the quantum-classical mapping of Ch. 1. First introduced by Blankenbecler, Scalapino, and Sugar in Ref. [135], this method is also known as the BSS algorithm, and is the de facto standard for simulating itinerant electron systems.

Our starting point is an action of the form $S = S_B[\phi] + S_F[\psi, \psi^\dagger, \phi]$ with a free bosonic contribution S_B and a fermionic part S_F comprising hopping processes and a fermion-boson interaction. We note in passing that while the spin-fermion model, Eq. (1.9), naturally falls into this class, this is not the case for the Hubbard model in its standard formulation, Eq. (1.17), featuring direct fermion-fermion interactions. We will resolve this seeming limitation in Sec. 3.3 by artificially introducing an auxiliary field by means of a discrete Hubbard-Stratonovich transformation. Because of this necessity, the DQMC method is also known as “auxiliary field quantum Monte Carlo”.

²⁹We note that one dimension of \mathcal{C}_n will typically correspond to imaginary time and therefore be of quantum nature.

Structurally, the partition function associated with the action S is given by the function integral

$$\mathcal{Z} = \int D(\psi, \psi^\dagger, \phi) e^{-S_B[\phi] - S_F[\psi, \psi^\dagger, \phi]}, \quad (2.15)$$

where ψ and ψ^\dagger are conjugate fermion fields and ϕ is a N_ϕ -component bosonic field. Generally speaking, the task of a quantum Monte Carlo scheme is to identify the statistically relevant quantum field configurations and sample them efficiently. In DQMC, the strategy to accomplish this goal is to exactly integrate out the fermionic part of the theory, $S_F[\psi, \psi^\dagger, \phi]$. This is feasible since, for fixed ϕ , S_F is a free theory in the sense that it is quadratic in the fermionic fields and can thus be evaluated by standard Gaussian integration rules. However, the coupling to the bosonic field will lead to a parametric dependence on ϕ .

To derive the DQMC form of Eq. (2.15), we first split the imaginary time propagation $e^{-S_F[\psi, \psi^\dagger, \phi]}$ into a sequence of successive propagations. To that end, we discretize imaginary time as $\tau = l\Delta\tau$, where $\Delta\tau$ is a step size, l enumerates different time slices, and $\beta = M\Delta\tau$. An illustration of this discretization is given in Fig. 4.13 in Ch. 4. Rewriting S_F in terms of hopping and interaction coefficient matrices T and V_{ϕ_l} associated with fermion bilinear terms³⁰ the partition function then reads

$$\mathcal{Z} = \int D(\psi, \psi^\dagger, \phi) e^{-S_B} \left[\exp \left(-\Delta\tau \sum_{l=1}^M \psi^\dagger [T + V_{\phi_l}] \psi \right) \right]. \quad (2.16)$$

Next, we factorize kinetic and potential contributions by repeatedly applying a symmetric Trotter-Suzuki decomposition [138, 139],

$$e^{-\Delta\tau(T+V_{\phi_l})} \approx e^{-\frac{\Delta\tau}{2}T} e^{-\Delta\tau V_{\phi_l}} e^{-\frac{\Delta\tau}{2}T} + O(\Delta\tau^3), \quad (2.17)$$

where a systematic numerical error of the order $O(\Delta\tau^3)$ is due to the fact that T and V_{ϕ_l} generally do not commute. Writing the integration over the fermionic sector as a trace over many-fermion states in Fock space [90, 122], one finds

$$\mathcal{Z} = \int D\phi e^{-S_B} \text{Tr} \left[\prod_{l=M}^1 \hat{B}_l \right] + O(\Delta\tau^2). \quad (2.18)$$

Here, the operators

$$\hat{B}_l = e^{-\frac{\Delta\tau}{2}\psi^\dagger T \psi} e^{-\Delta\tau\psi^\dagger V_{\phi_l} \psi} e^{-\frac{\Delta\tau}{2}\psi^\dagger T \psi} \quad (2.19)$$

can be viewed as imaginary time slice propagators [122]. Note that the numerical error scales as $O(\beta\Delta\tau^3) = O(\Delta\tau^2)$ since each factor in the imaginary time product of extent β comes with an error of the order of $O(\Delta\tau^3)$.

Realizing that, for fixed ϕ , the trace in Eq. (2.18) corresponds to a non-interacting fermion problem, it can be explicitly evaluated to a matrix determinant (see App. B), reflecting the fermionic anticommutation relations, [32, 90, 120, 122, 135]

$$\mathcal{Z} = \int D\phi e^{-S_B} \det G_\phi^{-1} + O(\Delta\tau^2). \quad (2.20)$$

³⁰See Sec. 3.2 and Sec. 3.3 for the spin-fermion model and the Hubbard model, respectively.

As we will prove below, Sec. 2.5.1, here the fermion matrix G_ϕ is given by the equal-time Green's function [90, 120, 122]

$$G_\phi = \langle c_i c_j^\dagger \rangle_{\phi_l} = \left[\mathbb{1} + \prod_{l=M}^1 B_l \right]^{-1}, \quad (2.21)$$

with³¹

$$B_l = e^{-\frac{\Delta\tau}{2}T} e^{-\Delta\tau V_{\phi_l}} e^{-\frac{\Delta\tau}{2}T} \quad (2.22)$$

being the coefficient matrix equivalent of the operator \hat{B}_l .

Eq. (2.20) is the central result of the DQMC reformulation of the quantum partition function. By comparing it to Eq. (2.3), it becomes apparent that we may interpret the real-valued bosonic field ϕ as a Monte Carlo configuration, $\mathcal{C} \equiv \phi$, and identify the integral kernel $w_\phi = e^{-S_B} \det G_\phi^{-1}$ as the corresponding probability weight. Note that the Green's function G_ϕ is implicitly dependent on the bosonic field due to the presence of fermion-boson interactions – for simplicity we will drop the sublabeled ϕ from now on. With the above mapping in mind, the *quantum* version of the Metropolis algorithm amounts to a Markov chain sampling of field configurations $\{\phi_n\}$ with acceptance probabilities, Eq. (2.13),

$$A_{\phi'\phi} = \min \left\{ 1, e^{-\Delta S_B} \frac{\det G}{\det G'} \right\}. \quad (2.23)$$

As we will discuss in the next section, expectation values of observables, such as bosonic susceptibilities, may then be calculated directly from the Markov chain $\{\phi_n\}$ whereas fermionic imaginary time correlations are accessible via Wick's theorem when applied to the equal-time Green's function [17, 90, 122].

The most severe limitation of DQMC is the famous fermion sign-problem [90, 122, 140–147]. Clearly, the identification

$$w_\phi = e^{-S_B} \det G_\phi^{-1} \quad (2.24)$$

as a probability weight is only justified if the resulting weights are real and positive (sign-problem free). However, since the Green's function as well as its determinant are generally complex-valued it is far from obvious why and for which physical systems this should be the case. In Ch. 3, we will provide a comprehensive systematic discussion of this issue and will present sign-problem free variants of the models considered in this thesis: the spin-fermion model, Eq. (1.9), and the Hubbard model, Eq. (1.17).

It is important to note that the DQMC procedure is unbiased and numerically exact in that it does not rely on any approximation. Conceptually, the only systematic sources of error are the statistical nature of the sampling, Secs. 2.1 and 2.6.2, and the imaginary time discretization, Eq. (2.16). Hence, arbitrary precision can be reached by increasing the simulation time and the resolution $1/\Delta\tau$. However, as we will discuss in detail in Ch. 4, these desirable properties are associated with a high computational cost. An optimal implementation

³¹Similar to \hat{B}_l we will refer to the objects B_l as (imaginary) time slice propagators.

of local updates $\phi_{li} \rightarrow \phi'_{li}$, Fig. 2.2, scales as $O(\beta N^3)$ as a function of inverse temperature and system size, N , such that simulations of large systems close to the absolute zero of temperature are typically out of reach even on state of the art supercomputers.

2.5 Measuring physical observables

In the DQMC framework, the expectation value of an observable X , Eq. (2.1), reads

$$\langle X \rangle = \langle \langle X \rangle_\phi \rangle_{\text{MC}} = \frac{1}{Z} \int D\phi \langle X \rangle_\phi w_\phi, \quad (2.25)$$

where $\langle X \rangle_\phi$ represents the fermion average of X for a fixed field configuration ϕ . If the physical observable is bosonic, that is it is directly expressible in terms of the bosonic field ϕ , the computation of Eq. (2.25) is straightforward: one simply calculates $\langle X \rangle_\phi = X(\phi)$ along the Markov chain of field configurations $\{\phi_n\}$. How to evaluate $\langle X \rangle_\phi$ for a fermionic observable on the other hand is not as obvious. Below we will derive concrete expressions for equal-time and time-displaced Green's functions as well as for pairing and density-like imaginary time correlations.

2.5.1 Equal-time Green's function

As per custom, we define the time-dependent equal-time Green's function as³² [23]

$$G_l = \langle c_i c_j^\dagger \rangle_{\phi_l} = \frac{\text{Tr}[\hat{B}_{l-1} \dots \hat{B}_1 \hat{B}_M \dots \hat{B}_l c_{il} c_{lj}^\dagger]}{\text{Tr}[\hat{B}_{l-1} \dots \hat{B}_1 \hat{B}_M \dots \hat{B}_l]}, \quad (2.26)$$

where we have discretized imaginary time in the same way as in the derivation of Eq. (2.18) and have dropped spin indices for convenience. Focusing on the related expectation value $\langle c_i^\dagger c_j \rangle_{\phi_l}$ first, we introduce a source term $\eta c_{il}^\dagger c_{lj}$ and take a derivative with respect to the source η to write

$$\begin{aligned} \langle c_i^\dagger c_j \rangle_{\phi_l} &= \frac{d}{d\eta} \ln \text{Tr} [\hat{B}_{l-1} \dots \hat{B}_1 \hat{B}_M \dots \hat{B}_l e^{-\eta c_{il}^\dagger c_{lj}}] \Big|_{\eta=0} \\ &= \frac{d}{d\eta} \ln \det [\mathbb{1} + B_{l-1} \dots B_1 B_M \dots B_l e^{\eta A_{ij}}] \Big|_{\eta=0} \\ &= \frac{d}{d\eta} \text{Tr} [\ln (\mathbb{1} + B_{l-1} \dots B_1 B_M \dots B_l e^{\eta A_{ij}})] \Big|_{\eta=0}. \end{aligned} \quad (2.27)$$

Here, in the second equality, we have evaluated the Gaussian fermionic trace, similar to around Eq. (2.18), defined a matrix A_{ij} with elements $[A_{ij}]_{kl} = \delta_{ik} \delta_{jl}$,

³²Note that we define G_l as the fermionic expectation value for fixed bosonic field, i.e. $\langle c_i c_j^\dagger \rangle_{\phi_l}$, rather than the full thermal Green's function $\langle c_i c_j^\dagger \rangle = \langle \langle c_i c_j^\dagger \rangle_\phi \rangle_{\text{MC}}$.

and used the identity $\ln \det A = \text{Tr} \ln A$ in the last line. Executing the derivative and taking the $\eta \rightarrow 0$ limit, one finds

$$\begin{aligned} \langle c_i^\dagger c_j \rangle_{\phi_l} &= \lim_{\eta \rightarrow 0} \frac{1}{\eta} \{ \text{Tr} [\ln (\mathbb{1} + B_{l-1} \dots B_1 B_M \dots B_l (\mathbb{1} + \eta A_{ij}))] \\ &\quad - \text{Tr} [\ln (\mathbb{1} + B_{l-1} \dots B_1 B_M \dots B_l)] \}, \\ &= \text{Tr} [(\mathbb{1} + B_{l-1} \dots B_1 B_M \dots B_l)^{-1} B_{l-1} \dots B_1 B_M \dots B_l A_{ij}] \\ &= \left[(\mathbb{1} + B_{l-1} \dots B_1 B_M \dots B_l)^{-1} B_{l-1} \dots B_1 B_M \dots B_l \right]_{ij}. \end{aligned} \quad (2.28)$$

This immediately leads to

$$\begin{aligned} [G_l]_{ij} &= \delta_{ij} - \langle c_j^\dagger c_i \rangle_{\phi_l} \\ &= [\mathbb{1} + B_{l-1} \dots B_1 B_M \dots B_l]^{-1}, \end{aligned} \quad (2.29)$$

which for $l = 1$ or $l = M + 1$ coincides with Eq. (2.21) and retrospectively justifies the labeling of G_ϕ as fermionic Green's function. Note that as a direct consequence of translational invariance in imaginary time, the thermal Green's function $\langle G_l \rangle = \langle \langle c_i c_j^\dagger \rangle_{\phi_l} \rangle$ is independent of imaginary time.

For later convenience we introduce the following alternative continuum notation of Eq. (2.29),

$$G(\tau) \equiv \langle c_i(\tau) c_j(\tau)^\dagger \rangle_\phi = [\mathbb{1} + B(\tau, 0) B(\beta, \tau)]^{-1}, \quad (2.30)$$

where the time slice propagators $B(\tau_2, \tau_1)$ are defined in terms of their discrete partners as

$$B(\tau_2, \tau_1) \equiv \prod_{l=l_2}^{l_1+1} B_l \quad (2.31)$$

with $\tau_1 = l_1 \Delta\tau$ and $\tau_2 = l_2 \Delta\tau$. Note that in this notation $G(0) = G(\beta) = G_\phi = G_1 = G_{M+1}$.

2.5.2 Time-displaced Green's function

Utilizing the time ordering operator T , we define the time displaced Green's function as [23, 90, 122]

$$G(\tau_1, \tau_2) = \langle T c_i(\tau_1) c_j^\dagger(\tau_2) \rangle_\phi = \begin{cases} \langle c_i(\tau_1) c_j^\dagger(\tau_2) \rangle_\phi, & \tau_1 > \tau_2, \\ -\langle c_j^\dagger(\tau_2) c_i(\tau_1) \rangle_\phi, & \tau_2 > \tau_1. \end{cases} \quad (2.32)$$

Similar to the derivation presented in the previous section it is straightforward to prove [120, 122, 148] the following more practical expression of $G(\tau_1, \tau_2)$,

$$G(\tau_1, \tau_2) = \begin{cases} B(\tau_1, \tau_2) G(\tau_2), & \tau_1 > \tau_2, \\ -(1 - G(\tau_1)) B(\tau_2, \tau_1)^{-1}, & \tau_2 > \tau_1, \end{cases} \quad (2.33)$$

$$= \begin{cases} B_{l_1} \dots B_{l_2+1} G_{l_2+1}, & \tau_1 > \tau_2, \\ -(1 - G_{l_1+1}) (B_{l_2} \dots B_{l_1+1})^{-1}, & \tau_2 > \tau_1. \end{cases} \quad (2.34)$$

Given that both the equal-time Green's functions G_l as well as the propagators B_l are readily accessible in DQMC, the time displaced Green's function may, in principle, be calculated as per the relation above. However, as we will demonstrate in Sec. 4.1.3 in detail, a naive computation of $G(\tau_1, \tau_2)$ based on Eq. (2.34), that is via a one-sided propagation of the appropriate equal-time Green's function $G(\tau_2)$, generally leads to severe numerical instabilities and unreliable estimates. For this reason, the more sophisticated approach presented in Ch. 4 should be employed in finite machine-precision DQMC implementations.

2.5.3 Wick's theorem

To compute fermionic equal-time and time-displaced correlations, for example in the pairing, charge, or spin channel, we utilize Wick's theorem [17, 23], which holds for a fixed bosonic field configuration [90, 120, 122, 135, 149] and may be applied to fermionic expectation values $\langle \cdot \rangle_\phi$. In the following, we derive concrete expressions for general pairing- and density-like four-point correlations suited for computations within the DQMC scheme.

Pairing correlations:

We consider imaginary time correlation functions of the form $\langle c_\alpha(\tau) c_\beta(\tau) c_\gamma^\dagger c_\delta^\dagger \rangle_\phi$ in which $\alpha, \beta, \gamma, \delta$ are multi-indices consolidating space, spin, and (further) flavor indices. In this case, the relevant contractions are given by

$$\begin{aligned} \overbrace{\langle c_\alpha(\tau) c_\beta(\tau) c_\gamma^\dagger c_\delta^\dagger \rangle_\phi} &= \langle c_\alpha(\tau) c_\delta^\dagger \rangle_\phi \langle c_\beta(\tau) c_\gamma^\dagger \rangle_\phi, \\ \overbrace{\langle c_\alpha(\tau) c_\beta(\tau) c_\gamma^\dagger c_\delta^\dagger \rangle_\phi} &= -\langle c_\alpha(\tau) c_\gamma^\dagger \rangle_\phi \langle c_\beta(\tau) c_\delta^\dagger \rangle_\phi, \\ \overbrace{\langle c_\alpha(\tau) c_\beta(\tau) c_\gamma^\dagger c_\delta^\dagger \rangle_\phi} &= \langle c_\alpha(\tau) c_\beta(\tau) \rangle_\phi \langle c_\gamma^\dagger c_\delta^\dagger \rangle_\phi = 0, \end{aligned} \quad (2.35)$$

Here, the third contribution vanishes precisely because of the presence of an exact $U(1)$ -symmetry for a fixed field configuration ϕ : the DQMC method implicitly assumes fermion particle-number conservation³³ by asserting that the fermionic part of the action can be expressed in terms of kinetic and potential coefficient matrices T and V in Eq. (2.16). Importantly, this fact does not preclude the occurrence of superconductivity as finite anomalous expectation values may still emerge under the thermodynamic Monte Carlo averaging over different ϕ -field configurations. From the contractions in Eqs. (2.35) we conclude that Wick's theorem takes the form

$$\begin{aligned} \langle c_\alpha(\tau) c_\beta(\tau) c_\gamma^\dagger c_\delta^\dagger \rangle_\phi &= \langle c_\alpha(\tau) c_\delta^\dagger \rangle_\phi \langle c_\beta(\tau) c_\gamma^\dagger \rangle_\phi - \langle c_\alpha(\tau) c_\gamma^\dagger \rangle_\phi \langle c_\beta(\tau) c_\delta^\dagger \rangle_\phi, \\ &= G_{\alpha\delta}(\tau, 0) G_{\beta\gamma}(\tau, 0) - G_{\alpha\gamma}(\tau, 0) G_{\beta\delta}(\tau, 0), \end{aligned} \quad (2.36)$$

where in the second equality we have identified the time-displaced Green's functions, Eq. (2.32). As the latter are accessible in DQMC, this expression may be used to compute the imaginary time pairing correlation function.

³³We will shortly describe in Ch. 3 how to simulate models with explicit $U(1)$ -symmetry breaking pairing terms.

Density-like correlations:

We now turn to density-like correlation functions of the form $\langle c_\alpha^\dagger(\tau) c_\beta(\tau) c_\gamma^\dagger c_\delta \rangle_\phi$. Examples include charge, spin, and pair-density correlations as well as current-current correlations, if appropriate weighing factors are included. The relevant non-vanishing (see above) contractions are given by

$$\begin{aligned} \overbrace{\langle c_\alpha^\dagger(\tau) c_\beta(\tau) c_\gamma^\dagger c_\delta \rangle_\phi} &= \langle c_\alpha^\dagger(\tau) c_\beta(\tau) \rangle_\phi \langle c_\gamma^\dagger c_\delta \rangle_\phi, \\ \underbrace{\langle c_\alpha^\dagger(\tau) c_\beta(\tau) c_\gamma^\dagger c_\delta \rangle_\phi} &= \langle c_\alpha^\dagger(\tau) c_\delta \rangle_\phi \langle c_\beta(\tau) c_\gamma^\dagger \rangle_\phi. \end{aligned} \quad (2.37)$$

Except the last expectation value, these resulting two-point correlations aren't directly interpretable as Green's functions. However, using anticommutativity, $\{c_\alpha, c_\beta^\dagger\} = \delta_{\alpha\beta}$, and defining $\tilde{G}_{\alpha\beta}(\tau_1, \tau_2) = \delta_{\alpha\beta} - G_{\beta\alpha}(\tau_2, \tau_1)$, we find that Wick's theorem takes the form

$$\begin{aligned} \langle c_\alpha^\dagger(\tau) c_\beta(\tau) c_\gamma^\dagger c_\delta \rangle_\phi &= \langle c_\alpha^\dagger(\tau) c_\beta(\tau) \rangle_\phi \langle c_\gamma^\dagger c_\delta \rangle_\phi + \langle c_\alpha^\dagger(\tau) c_\beta \rangle_\phi \langle c_\gamma(\tau) c_\delta^\dagger \rangle_\phi \\ &= \tilde{G}_{\alpha\beta}(\tau, \tau) \tilde{G}_{\gamma\delta}(0, 0) + \tilde{G}_{\alpha\beta}(\tau, 0) G_{\gamma\delta}(\tau, 0). \end{aligned} \quad (2.38)$$

As in the case of pairing correlations, Eq. (2.36), the final expression only consists of time-displaced Green's functions and hence serves as a recipe for computing density-like correlations in DQMC.

2.6 Error estimation

Irrespective of the bosonic or fermionic nature of a physical observable X , estimating reliable error bounds for the corresponding Monte Carlo expectation value $\langle X \rangle$ takes some caution. Most importantly, statistical correlations between successive data points, which are inherent to the Markov process, Eq. (2.5), and may be enhanced by critical divergences, Eqs.(1.1) and (1.2), can easily lead to an overestimation of confidence intervals and, consequently, to a misinterpretation of the simulation results. In the following, we describe the methods used in this thesis to handle these issues as well as derive the famous statistical $O(M^{-1/2})$ scaling. Our presentation is largely influenced by Refs. [120], [126], and [131]. Concrete open-source implementations of all the discussed data analysis methods have been published in form of the Julia package `BinningAnalysis.jl` [S2] as a supplement to this thesis.

2.6.1 Statistical error of direct sampling

Formally, an observable X represents a mapping, $X : \Omega \rightarrow \mathbb{R}$, that assigns a value to each configuration of the sample space, i.e. $x = X(\mathcal{C} \in \Omega)$. Denoting the distribution of the values x as p_x^X – which should not be confused with the Boltzmann distribution of configurations, Eq. (2.2) – the expectation value of X is formally defined as

$$\langle X \rangle = \int dx x p_x^X. \quad (2.39)$$

A measure for the dispersion of p_x^X is given by the variance

$$\sigma_X^2 = \langle X^2 \rangle - \langle X \rangle^2. \quad (2.40)$$

These definitions, which assume knowledge of p_x^X , are of little practical use in a Markov chain Monte Carlo context where one has only access to a finite sample $\{x_i = X(\mathcal{C}_i) | i = 1, 2, \dots, N\}$. We therefore introduce the sample mean [120]

$$\bar{X} = \frac{1}{N} \sum_{i=1}^N x_i, \quad (2.41)$$

already seen in Eq. (2.4), and the sample variance [120]

$$\bar{\sigma}_X^2 = \frac{1}{N-1} \sum_{i=1}^N (x_i - \bar{X})^2. \quad (2.42)$$

These quantities are straightforward to calculate and are valid estimators for the expectation value and variance above in the sense that $\langle \bar{X} \rangle = \langle X \rangle$ and $\langle \bar{\sigma}_X^2 \rangle = \langle \sigma_X^2 \rangle = \sigma_X^2$. Consequently, we may approximate $\langle X \rangle \approx \bar{X}$ in our Monte Carlo simulations.

How can we estimate the error associated with this identification? Analogous to Eq. (2.40), we introduce the standard error of the sample mean as the root-mean-square deviation from the true expectation value,

$$\sigma_{\bar{X}} = \sqrt{\langle (\bar{X} - \langle X \rangle)^2 \rangle} = \sqrt{\langle \bar{X}^2 \rangle - \langle \bar{X} \rangle^2}. \quad (2.43)$$

In the case of a direct sampling scheme, Fig. 2.1a, data points are statistically independent, $\langle x_i x_j \rangle = \langle x_i \rangle \langle x_j \rangle = \langle X \rangle^2$, and $\sigma_{\bar{X}}$ takes the famous form [120, 126]

$$\begin{aligned} \sigma_{\bar{X}} &= \sqrt{\frac{1}{N^2} \left\langle \sum_i x_i^2 + \sum_{i \neq j} x_i x_j \right\rangle - \langle X \rangle^2} \\ &= \sqrt{\frac{1}{N} \left\langle \frac{1}{N} \sum_{i=1}^N x_i^2 \right\rangle - \frac{1}{N} \langle X \rangle^2} \\ &= \frac{\sigma_X}{\sqrt{N}}. \end{aligned} \quad (2.44)$$

It is important to note that while the standard deviation σ_X is an intrinsic property of the distribution function p_x^X the standard error $\sigma_{\bar{X}}$ is determined by the sampling process. Most importantly, only the latter characteristically vanishes in the limit $N \rightarrow \infty$ as $N^{-1/2}$.

Finally, given that the sample mean is essentially an accumulation, the distribution of \bar{X} is guaranteed to be Gaussian by the central limit theorem. We can therefore state the following “n-sigma” confidence intervals [120, 126],

$$\begin{aligned} P(\bar{X} - \sigma_{\bar{X}} < \langle X \rangle < \bar{X} + \sigma_{\bar{X}}) &= 65\%, \\ P(\bar{X} - 2\sigma_{\bar{X}} < \langle X \rangle < \bar{X} + 2\sigma_{\bar{X}}) &= 95\%, \\ P(\bar{X} - 3\sigma_{\bar{X}} < \langle X \rangle < \bar{X} + 3\sigma_{\bar{X}}) &= 99\%, \end{aligned} \quad (2.45)$$

indicating the probability that an interval of width $n\sigma_{\bar{X}}$ centered around \bar{X} contains the expectation value $\langle X \rangle$. As per custom, we will visualize one-sigma intervals as error bars in plots throughout this thesis.

2.6.2 Statistical error of Markov chain Monte Carlo

As we have seen in the derivation of Eq. (2.44), the characteristic $N^{-1/2}$ scaling of the Monte Carlo error assumes the absence of correlations between data points. This is the case in a direct sampling scheme, Fig. 2.1, where different draws are independent. The sequence of values generated by Markov chain Monte Carlo on the other hand is already correlated due to the defining property of the Markov process: step n is dependent on the previous step $n - 1$, Eq. (2.5). Additionally, pronounced thermal and quantum fluctuations in the vicinity of phase transitions can enhance correlations beyond this process-inherent dependence [1, 120, 122, 126, 131]. This naturally leads to the question of how the statistical analysis of the previous section is effected by the presence of statistical correlations.

To answer the same, we drop the assumption of statistical independence around Eq. (2.44). In this case, one finds that the Monte Carlo error of an observable X generally takes the form [120, 126]

$$\sigma_X^2 = \frac{\sigma_X^2}{N}(1 + 2\tau_X), \quad (2.46)$$

where τ_X is the *correlation time*, [120, 126]

$$\tau_X = \sum_{k=1}^{N-1} \left(1 - \frac{k}{N}\right) \frac{\chi_k}{\chi_0}. \quad (2.47)$$

As a function of the autocorrelation $\chi_{|i-j|} \equiv \chi_{ij} = \langle x_i x_j \rangle - \langle x_i \rangle \langle x_j \rangle$, τ_X serves as a measure for the presence of statistical correlations independent of their origin. Clearly, in the case $\tau_X = 0$ we recover the direct sampling result, Eq. (2.44). If, however, $\tau_X \geq 0$, the numerical error σ_X increases by a factor $\sqrt{1 + 2\tau_X}$. This implies that the correlated data set of length N effectively only contains information of $N/(1 + 2\tau_X) < N$ statistically independent samples. Consequently, ignoring the presence of correlations by using Eq. (2.44) can lead to a severe underestimation of the Monte Carlo error.

Binning analysis

To estimate the correlation time τ_X of an observable X from a finite Markov chain we employ a binning analysis. The approach is to divide the available data set into M bins, each of size m . Assuming that those batches are long enough to be statistically independent and to share a common variance, i.e. $m \gtrsim \tau_X$, we may apply Eq. (2.44) for uncorrelated samples on a block level,

$$\sigma_X^2 = \frac{\sigma_{\text{block}}^2}{M}. \quad (2.48)$$

Here the block variance σ_{block}^2 can readily be estimated from a sample block variance analogous to Eq. (2.42).

The difficulty of the error analysis now lies in the choice of the bin size m . On the one hand, one must assure $m \gtrsim \tau$ to justify Eq. (2.48). Without prior knowledge of τ_X , this suggests to choose m as large as possible. On the other

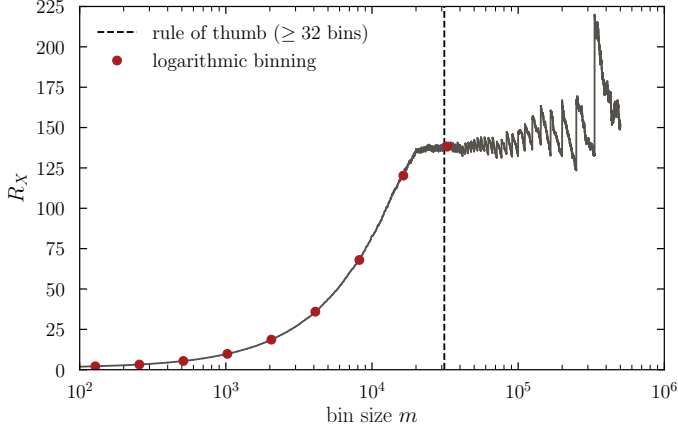


Figure 2.3: The dashed line indicates the maximal bin size leading to at least 32 bins, i.e. for which $M \geq 32$. The dots indicate bin sizes compatible with logarithmic binning, see Fig. 2.4. Inspired by Fig. 3.3 in Ref. [120].

hand, a firm computation of σ_{block}^2 requires a sufficiently large number of blocks $M = N/m$ as otherwise statistical fluctuations, due to a too small sample size, will spoil the estimate. A strategy for balancing this opposition is to compute the correlation error factor³⁴, [120]

$$R_X = 1 + 2\tau_X = \frac{m\sigma_{\text{block}}^2}{\sigma_X^2}, \quad (2.49)$$

where the second equality follows from a comparison of Eqs. (2.46) and (2.48), for all compatible³⁵ bin sizes, Fig. 2.3. As a function of the latter, R_X first increases, as a consequence of the presence of correlations, before converging to a stable plateau once different bins are statistically independent. At even larger bin sizes, visible fluctuations indicate strong statistical uncertainties due to a decreasing number of blocks. This last regime can be excluded by requiring a minimum number of bins M_{\min} . A common rule of thumb is to choose $M_{\min} \approx 32$ which is typically considered to be statistically sufficient [120, 126].

Estimates for τ_X and therefore the standard error $\sigma_{\bar{X}}$ are only reliable for bin sizes corresponding to the plateau of $R_X(m)$. For automation purposes, one may take $R_X(m_{M_{\min}})$ as an error estimate. However, the absence of a plateau must be excluded as it is typically indicative of an insufficiently long Markov chain and requires a continuation of the Monte Carlo simulation.


On a practical note, it is advantageous to bin individual observable values X_i logarithmically and calculate standard error estimates (and the sample mean)

³⁴Compared to monitoring $\sigma_{\bar{X}}$ directly, the quantity R_X has the advantage of being normalized to the variance. It may therefore be compared to unity.

³⁵We only consider bin sizes m for which $M = N/m$ is integer.

Figure 2.4:

Logarithmic binning of observables values during a Monte Carlo measurement process. The leaves of the binary tree (top row) represent the unmodified time series X_i . A node at a level n indicates an arithmetic mean of two values of the preceding layer, i.e. $X_i^{(n)} = (X_i^{(n-1)} + X_{i+1}^{(n-1)})/2$. Different levels correspond to logarithmic bin sizes $m_n = 2^{n-1}$ indicated in Fig. 2.3.



during the measurement process *on-the-fly*, Fig. 2.4. This way one can avoid storing the entire time series such that the binning analysis only has $O(\log_2 N)$ space complexity. The resulting subset of bin sizes, which corresponds to different levels of the binary tree in Fig. 2.4, is indicated in Fig. 2.3.

2.6.3 Functions of observables

We now consider the case in which an observable of interest G is a function of the expectation values of other observables,³⁶

$$G = g(\langle X \rangle, \langle Y \rangle) \quad (2.50)$$

In general, X and Y can be correlated such that error propagation is not straightforward. Prominent examples include the specific heat $C_V(T) = (\langle E^2 \rangle - \langle E \rangle^2)/T^2$, where $X = E$ and $Y = E^2$, and the Binder cumulant $B_O = \langle O^4 \rangle / \langle O^2 \rangle^2$ [150], with $X = O^2$ and $Y = O^4$. The questions we want to ask is: What is the statistical error associated with G and how to estimate it?

Formally, statistical correlations between X and Y are captured by the covariance matrix

$$\sigma_{X,Y} = \langle (X - \langle X \rangle)(Y - \langle Y \rangle) \rangle. \quad (2.51)$$

Taking the sample mean, Eq. (2.41), as an estimate for the expectation values and approximating $g(\langle X \rangle, \langle Y \rangle) \approx g(\bar{X}, \bar{Y})$ one finds the following squared error, [120]

$$(\Delta G)^2 \approx \left(\partial_X g|_{\bar{X}, \bar{Y}} \right)^2 \frac{\sigma_X^2}{N} + \left(\partial_Y g|_{\bar{X}, \bar{Y}} \right)^2 \frac{\sigma_Y^2}{N} + 2 \partial_X g|_{\bar{X}, \bar{Y}} \partial_Y g|_{\bar{X}, \bar{Y}} \frac{\sigma_{X,Y}}{N^2}. \quad (2.52)$$

Note that in the absence of cross-correlations, $\sigma_{X,Y} = 0$, and taking the identity for g , this expression reduces to two independent instances of Eq. (2.44).

Computing the full covariance matrix $\sigma_{X,Y}$ from finite data sets is generally a laborious task. To circumvent this cost, we rely on the Jackknife method [120,

³⁶We restrict the discussion to the bivariate case but note that it trivially extends to an arbitrary number of dependencies.

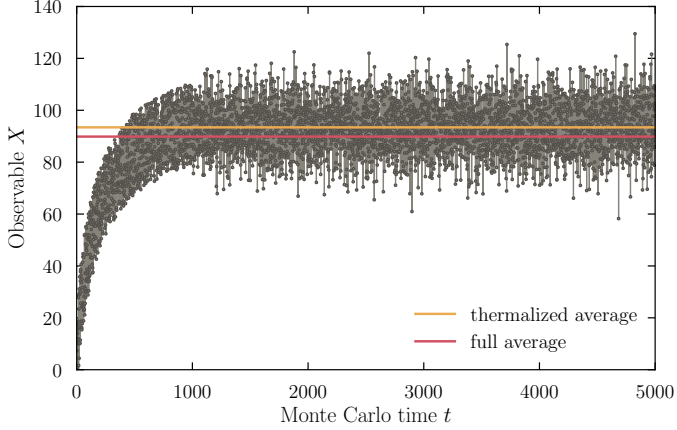


Figure 2.5: Thermalization of a Monte Carlo observable. Shown are the observable values, evaluated at the current position of the Markov walker in configuration space, as a function of Monte Carlo time. Horizontal lines indicate the observable average $\langle O \rangle$ computed using the entire time series (beige) and by disregarding the thermalization phase, $t \lesssim 1000$ (red).

123] to practically estimate the standard error ΔG in Monte Carlo simulations. The procedure is to resample the Markov chains of X and Y by successively evaluating g on all subsets with a single disregarded element i , i.e. computing

$$\bar{g} = \frac{1}{N} \sum_{i=1}^N g_{j \neq i} \quad (2.53)$$

with $g_{j \neq i} \equiv g(\bar{X}_{j \neq i}, \bar{Y}_{j \neq i})$. In this case, the Jackknife estimates for mean and standard error are given by [120]

$$\bar{G} = N g(\bar{X}, \bar{Y}) - (N-1) \bar{g}, \quad (2.54)$$

$$(\Delta G)^2 = \frac{N-1}{N} \sum_{i=1}^N (g_{j \neq i} - g(\bar{X}, \bar{Y}))^2. \quad (2.55)$$

Note that if the Markov chains associated with X and Y are by themselves correlated, the Jackknife method must be combined with the binning analysis of the previous section. In this case, a single block is disregarded in each Jackknife sample $g_{j \neq i}$.

2.6.4 Equilibration

An important property of convergence of the Markov process, Eq. (2.9), is its independence of the starting point $\mathcal{C}_1 \in \Omega$. However, given a randomly initialized Markov walker the region $S \subset \Omega$ of the configuration space associated with

the stationary distribution is generally only reached after a finite number of steps, Fig. 2.1. This fact, commonly known as *thermalization* or *equilibration*, is obvious in the time series of an observable. In the example shown in Fig. 2.5, it is only after about 1000 moves that the equilibrium plateau is reached and the observable starts to fluctuate around a well-defined mean. At shorter Monte Carlo times, artifacts of the initialization become apparent in form of a stark increase. For this reason, the thermalization phase should be omitted in all Monte Carlo measurements. For the case shown in Fig. 2.5, the relative deviation between a full average, based on the entire time series, and a thermalized average is approximately 4%.

3 | Avoiding the fermion-sign problem

Because of the combination of polynomial efficiency and a controlled systematic error, determinant quantum Monte Carlo is considered by many to be the gold standard of numerical methods for studying itinerant electron systems [151]. The major obstacle in applying it to arbitrary physical models is the fact that, in general, the DQMC determinant $\det G_\phi^{-1}$, Eq. (2.20), is complex valued and prohibits an interpretation of $p \equiv w_\phi \propto \det G_\phi^{-1}$, Eq. (2.24), as a probability weight. This is the infamous fermion-sign problem. Clearly, DQMC as described in Ch. 2 requires p to be real and positive semidefinite, i.e. $p \geq 0$. This leads to two natural questions: (1) Is it possible to circumvent the sign-problem for a specific model? (2) Which physical models are free of the sign-problem to begin with?

In the following, we touch upon both of these questions. As for (1), it turns out that the spin-fermion model introduced in Ch. 1 is sign-problematic and, in the form of Eq. (1.9), not amenable to DQMC. Following Ref. [81], we demonstrate below that the sign-problem can nonetheless be circumvented in this case by designing a two-band version of the model that shares the same low-energy physics³⁷. With regard to question (2), we provide a compact review of systematic classifications of sign-problem free models based on symmetries and beyond. In particular, we discuss an original semigroup approach, recently put forward in Ref. [144], that unifies all established categorization principles and, by extending them, describes a novel class of sign-problem free models. We report on an (unpublished) collaborative effort³⁸ with the author of Ref. [144], which has led to the first implementation of these novel insights in form of the identification of a new, sign-problem free, physical model and proof-of-principle DQMC simulations of the same.

3.1 Enforced positive semidefiniteness

The physical origin of the sign-problem in QMC simulations of itinerant many-electron systems is the exchange statistics of fermions. Since the latter anti-commute a sign is generated on every exchange of two fermions. If the number of these exchanges is odd a specific configuration acquires a total minus sign.

³⁷ Assuming that the vicinity of the hot spots governs the important low-energy properties.

³⁸ The author of Ref. [144] and the author of this thesis (CB) have created the pedagogical visualization in Fig. 3.3 and identified the specific models in Eqs. (3.59) and (3.52) in the novel $J_1^2 = J_2^2 = -1$ class. Frederic Freyer (FF) has performed DQMC simulations of model Eq. (3.59). CB and FF have written the DQMC code, Ref. [S4].

More generally speaking, the sign-problem can be traced back to the topological properties of worldlines, i.e. paths along imaginary time in the sense of the quantum-classical mapping, Fig. 1.2 [152]. It is hence not restricted to many-fermion systems but also occurs in worldline QMC³⁹ of spin-systems such as frustrated Heisenberg antiferromagnets [141, 153].

A formal solution to the sign-problem is to enforce the positivity of the statistical weights by ignoring the sign σ of the same in the Monte Carlo sampling and rewriting the expectation value, Eq. (2.1), as [120]

$$\langle X \rangle = \frac{\int d\mathcal{C} X(\mathcal{C}) w(\mathcal{C})}{\int d\mathcal{C} w(\mathcal{C})} = \frac{\int d\mathcal{C} X(\mathcal{C}) \sigma(\mathcal{C}) |w(\mathcal{C})|}{\int d\mathcal{C} \sigma(\mathcal{C}) |w(\mathcal{C})|} = \frac{\langle X \cdot \sigma \rangle_{\text{abs}}}{\langle \sigma \rangle_{\text{abs}}}. \quad (3.1)$$

Here $\langle \cdot \rangle_{\text{abs}}$ indicates a Monte Carlo sampling with the strictly positive semidefinite weights $|w_{\mathcal{C}}| = \sigma w_{\mathcal{C}}$ and $\langle \sigma \rangle_{\text{abs}}$ is the average of the sign. Although this reformulation is universally valid it is of little avail in practical simulations since $\langle \sigma \rangle_{\text{abs}}$ typically vanishes exponentially. As a function of system size N and inverse temperature β one finds [120, 140]

$$\langle \sigma \rangle_{\text{abs}} \propto \exp(-\beta N \Delta f), \quad (3.2)$$

where $\Delta f = f - f_{\text{abs}}$ is a free energy density difference due to sampling a modified statistical ensemble. While the average of σ approaches zero its variance is still set by the scale of the possible values $+1$ and -1 and the relative error of the sign scales as

$$\Delta_{\text{rel}} = \frac{\Delta \sigma}{\sigma} \propto \frac{\exp(\beta N \Delta f)}{\sqrt{M}}. \quad (3.3)$$

Hence, to keep Δ_{rel} reasonably small the exponential dependence on N and β must be compensated by an (exponential) increase of the number of Monte Carlo sweeps M . This effectively spoils the polynomial computational scaling and defeats the advantage of the Monte Carlo method.

An alternative approach to enforcing positive semidefinite statistical weights is a basis transformation [151, 154, 155]. For example, it is simple to see that the eigenbasis of an Hamiltonian H is always sign-problem free [151]: the determinant reduces to a product of eigenvalue exponentials that are real and non-negative due to the Hermiticity of H . Similar bases with $w_{\mathcal{C}} \geq 0, \forall \mathcal{C}$ have been found and utilized in studies of e.g. frustrated spin-chains [153]. However, in general, the search for a sign-problem free basis is NP-hard [141, 156] and therefore not efficiently feasible.

3.2 Circumventing the sign-problem of the spin-fermion model

Although there is no universal solution to the sign-problem a seminal work by Berg *et al.* [81] has demonstrated that a powerful strategy for solving the same

³⁹The sample space is taken to be all possible worldlines.





Figure 3.1: Fermi surfaces (solid lines) of the original (left) and the two-band spin-fermion model (right), Eqs. (1.9) and (3.4), respectively. In the vicinity of the hot spots both Fermi surfaces are locally identical. The surrounding box represents the first Brillouin zone and the dashed lines indicate the original FS. The parameters of the energy dispersions are as indicated below Eq. (1.7) (left) or given by $t_{xh} = t_{yv} = 1$, $t_{yh} = t_{xv} = 0.5$, $\mu_x = \mu_y = 0.5$ (right), where subscripts h, v indicate horizontal and vertical nearest-neighbor hopping, respectively.

for a specific set of sign-problematic models is to “deform” the original theory in such a way that the weights become non-negative while the low-energy physics is preserved. It is precisely this idea that allows us to circumvent the sign-problem of the spin-fermion model, which sits at the center of this thesis. The fact that the original low-energy theory, Eq. (1.9), is subject to the sign-problem is understood from the fact that the fermionic sector corresponds to the repulsive Hubbard model, Sec. 1.3, which except for special fillings [90], is famously known to be sign-problematic – we will consider the attractive and repulsive Hubbard models in detail below.

As discussed in Ch. 1, in the vicinity of a metallic SDW quantum critical point the collective order parameter modes with a commensurate antiferromagnetic ordering wave vector $\mathbf{Q} = (\pi, \pi)$ couple most resonantly to electrons near the hot spots, a distinct set of points on the Fermi surface, see Fig. 3.1. It is therefore expected that the behavior of the quantum critical metal at low energies is solely governed by Fermi surface patches around those hot spots. If this assumption is justified – it is precisely the working hypothesis of the vast majority of theoretical works on the spin-fermion model [3, 4, 34, 81] – one may arbitrarily deform the energy dispersion away from those distinguished points without affecting the quantum critical properties. As shown in Fig. 3.1, we utilize this freedom to introduce another artificial fermion flavor into the spin-fermion model while ensuring that the resulting two-band dispersion is tuned such that it precisely reproduces the original Fermi surface in the close vicinity of the hot spots. In contrast to many theoretical approaches in which the Fermi surface patches are linearized, for example Ref. [34], our scheme preserves the curvature of the Fermi surface in the sense of a Taylor series expansion around \mathbf{k}_{hs} up to an arbitrary desired order. The action of the obtained two-band spin-fermion model reads

Figure 3.2:

Illustration of the two-band spin-fermion model, (3.4), on a two-dimensional square lattice. Shown are the local Yukawa interactions (dashed lines) between the ψ_x and ψ_y fermion flavors as well as the magnetic order parameter ϕ (arrows). Modified from Ref. [P1] under CC-BY 4.0.



[81],

$$\begin{aligned}
 S_{\text{SFL}} &= S_\phi + S_\psi + S_\lambda \\
 &= \int_0^\beta d\tau \sum_{\mathbf{r}, \mathbf{r}'} \sum_{s, \alpha} \psi_{\alpha \mathbf{r} s}^\dagger [(\partial_\tau - \mu) \delta_{\mathbf{r} \mathbf{r}'} - t_{\alpha \mathbf{r} \mathbf{r}'}] \psi_{\alpha \mathbf{r}' s} \\
 &\quad + \lambda \int_0^\beta d\tau \sum_{\mathbf{r}} \sum_{s, s'} e^{i\mathbf{Q} \cdot \mathbf{r}} \phi_{\mathbf{r}} \cdot \psi_{x \mathbf{r} s}^\dagger \sigma_{ss'} \psi_{y \mathbf{r} s'} + \text{h.c.} \\
 &\quad + \int_0^\beta d\tau \sum_{\mathbf{r}} \left[\frac{1}{2c^2} (\partial_\tau \phi_{\mathbf{r}})^2 + \frac{1}{2} (\nabla \phi_{\mathbf{r}})^2 + \frac{r}{2} \phi_{\mathbf{r}}^2 + \frac{u}{4} (\phi_{\mathbf{r}}^2)^2 \right].
 \end{aligned} \tag{3.4}$$

Here $\alpha \in \{x, y\}$ is the additional fermion flavor index and we have discretized real space such that \mathbf{r}, \mathbf{r}' are discrete coordinates on a square lattice and ∇ is a lattice gradient⁴⁰. Note that the Yukawa term only couples fermions of different flavor, i.e. ψ_x and ψ_y , since the ordering wave vector \mathbf{Q} connects hot spots in different energy bands, Fig. 3.1. Interpreting α as a layer index, the interaction can be visualized as in Fig. 3.2. As we will see below this qualitative change, compared to Eq. (1.9), where S_λ is of intraband character, is responsible for the absence of the sign-problem in DQMC simulations of S_{SFL} .

Let us remark that the core ingredient of the “deformation scheme” outlined above is the existence of hot spots. It should therefore be generally applicable to models of metallic quantum phase transitions with a finite wave vector $\mathbf{Q} \neq 0$, i.e. describing the onset of density-wave order. Furthermore, important classes of unconventional superconductors, Sec. 1.2, such as the iron-pnictides and heavy fermion materials naturally feature multiple bands. In these cases, the onset of metallic order associated with interband wave vectors \mathbf{Q} can be efficiently studied without the introduction of an artificial fermion flavor.

Fermion matrix structure

To prove that the two-band spin-fermion model S_{SFL} is indeed free of the sign-problem we investigate the structure of the Green’s function G_ϕ , Eq. (2.21), and show that its determinant is strictly positive semidefinite for all configurations ϕ . The hopping and interaction matrices associated with the fermion bilinear

⁴⁰We approximate all lattice gradients by first order numerical derivatives.

terms in Eq. (3.4) are given by

$$[T]_{ij,\alpha\alpha',ss'} = \delta_{ss'}\delta_{\alpha\alpha'}(-t_{\alpha ij} - \mu\delta_{ij}), \quad (3.5)$$

$$[V]_{ij,\alpha\alpha',ss',l} = \lambda[\sigma_x]_{\alpha\alpha'}\delta_{ij}[\sigma \cdot \phi_{li}]_{ss'}. \quad (3.6)$$

Since T is diagonal in fermion flavor space, so is its matrix exponential,

$$e^{-\frac{\Delta\tau}{2}T} = \begin{pmatrix} x \uparrow & y \downarrow & x \downarrow & y \uparrow \\ e^{-\frac{\Delta\tau}{2}T_x^\uparrow} & & & \\ & e^{-\frac{\Delta\tau}{2}T_y^\downarrow} & & \\ & & e^{-\frac{\Delta\tau}{2}T_x^\downarrow} & \\ & & & e^{-\frac{\Delta\tau}{2}T_y^\uparrow} \end{pmatrix} \begin{pmatrix} x \uparrow \\ y \downarrow \\ x \downarrow \\ y \uparrow \end{pmatrix}. \quad (3.7)$$

To compute the matrix exponential of V one may apply Sylvester's formula [157],

$$f(a(\hat{n} \cdot \sigma)) = \mathbb{1} \frac{f(a) + f(-a)}{2} + \hat{n} \cdot \sigma \frac{f(a) - f(-a)}{2}, \quad (3.8)$$

with the identifications $f(x) = \exp(-x)$, $a = \Delta\tau|\phi|$, and $\hat{n} = \phi/|\phi|$. One finds that $\exp(-\Delta\tau V)$ takes the form

$$e^{-\Delta\tau V_i} = \begin{pmatrix} x \uparrow & y \downarrow & x \downarrow & y \uparrow \\ C & S & & R \\ \bar{S} & C & -R & \\ & -R & C & \bar{S} \\ R & & S & C \end{pmatrix} \begin{pmatrix} x \uparrow \\ y \downarrow \\ x \downarrow \\ y \uparrow \end{pmatrix}, \quad (3.9)$$

in which, with $\phi = (\phi^x, \phi^y, \phi^z)^T$, the block matrices C , S , and R are diagonal and given by [32]

$$\begin{aligned} C_{ij} &= \delta_{ij} \cosh(\lambda\Delta\tau|\phi_i|), \\ S_{ij} &= \delta_{ij} \frac{(i\phi_i^y - \phi_i^x)}{|\phi_i|} \sinh(\lambda\Delta\tau|\phi_i|), \\ R_{ij} &= -\delta_{ij} \frac{\phi_i^z}{|\phi_i|} \sinh(\lambda\Delta\tau|\phi_i|). \end{aligned} \quad (3.10)$$

Postponing a discussion of the absence of the sign-problem for the $O(3)$ symmetry case to Sec. 3.4.1 below, we note that for $N_\phi < 3$, i.e. $\phi^z = 0$, the block matrix R vanishes and the interaction matrix exponential simplifies to

$$e^{-\Delta\tau V_i} = \begin{pmatrix} X & \\ & \bar{X} \end{pmatrix}, \quad X = \begin{pmatrix} C & S \\ \bar{S} & C \end{pmatrix} \begin{pmatrix} x \uparrow \\ y \downarrow \end{pmatrix}. \quad (3.11)$$

Combined with Eq. (3.7) this implies that the Green's function is block diagonal as well. Under the condition [32]

$$T_\alpha^\uparrow = \bar{T}_\alpha^\downarrow, \quad \text{for } \alpha \in \{x, y\}, \quad (3.12)$$

which is trivially fulfilled if the hopping amplitudes $t_{\alpha ij}$ are real and spin-independent⁴¹, it decomposes into two complex conjugated blocks

$$G_\phi = \begin{pmatrix} \Gamma & \\ & \bar{\Gamma} \end{pmatrix}. \quad (3.13)$$

The DQMC determinant hence factorizes and the probability weight is strictly positive semidefinite,

$$w_\phi \propto \det G_\phi^{-1} = |\det \Gamma|^2 \geq 0. \quad (3.14)$$

Consequently, the two-band spin-fermion model S_{SFL} is sign-problem free.

Beyond indicating the absence of negative fermion-signs, Eq. (3.14) indicates that instead of having to operate with the full Green's function of size $fN \times fN$, where $f = 4$ indicates the total number of fermion flavors (2 band and 2 spin), it is sufficient to work with the four times smaller matrix Γ in DQMC. The full Green's function can always be obtained by Eq. (3.13), if necessary. As we will see around Eq. (4.33) in the next chapter this, theoretically, leads to a 8-fold speedup of local updates⁴². In the case of easy-axis ordering, where the order parameter $\phi = \phi_x$, one further has $S = \bar{S}$ such that the matrix Γ is real. This leads to an additional speedup of all finite precision linear algebra operations.

3.3 Determinant factorization: Hubbard model

The factorization of the Green's function determinant, as observed in Eq. (3.14), is a general scheme that appears in the sign-problem analysis of many physically relevant models. Arguably, the most prominent example is the Hubbard model, Eq. (1.17), introduced in Ch. 1. In the following we consider both the repulsive ($U > 0$) and attractive ($U < 0$) variants. DQMC simulations of the latter will be presented in Ch. 7 in the context of machine learning transport properties.

We start by repeating the definition of the Hubbard model, [158, 159]

$$H = -t \sum_{\langle i,j \rangle, s} (c_{is}^\dagger c_{js} + \text{h.c.}) + U \sum_i \left(n_{i\uparrow} - \frac{1}{2} \right) \left(n_{i\downarrow} - \frac{1}{2} \right) - \mu \sum_i n_i \quad (3.15)$$

where, compared to Eq. (1.17), we have shifted the chemical potential by $U/2$ such that half-filling corresponds to $\mu = 0$. To bring the Hamiltonian into DQMC form, that is to decouple the fermion-fermion interaction term, one could employ a regular Hubbard-Stratonovich transformation as discussed in Sec. 1.3 and App. A. In this case, the introduced auxiliary field ϕ would represent continuous degrees of freedom. However, as has been pointed out by Hirsch [160], it is more efficient to perform a modified HS transformation which renders ϕ to be of discrete Ising character, i.e. $\phi_{li} \in \{+1, -1\}$. Specifically, we apply the following

⁴¹We introduce an artificial magnetic flux in Sec. 4.3.1 where we will make sure that this condition is still satisfied.

⁴²A practical benchmark is provided in Fig. 4.12.

Hirsch decouplings of the onsite-interaction,⁴³ [90, 93, 160, 161]

$$\begin{aligned} e^{-U\Delta\tau(n_{i\uparrow}-\frac{1}{2})(n_{i\downarrow}-\frac{1}{2})} &= \gamma \sum_{\phi_{li}=\pm 1} e^{\phi_{li}\lambda(n_{i\uparrow}-1)}, \\ \gamma &= \frac{1}{2}e^{U\Delta\tau/4}, \quad \cosh(\lambda) = e^{-U\Delta\tau/2}, \end{aligned} \quad (3.16)$$

and

$$\begin{aligned} e^{-U\Delta\tau(n_{i\uparrow}-\frac{1}{2})(n_{i\downarrow}-\frac{1}{2})} &= \gamma \sum_{\phi_{li}=\pm 1} e^{\phi_{li}\lambda(n_{i\uparrow}-n_{i\downarrow})}, \\ \gamma &= \frac{1}{2}e^{-U\Delta\tau/4}, \quad \cosh(\lambda) = e^{U\Delta\tau/2}, \end{aligned} \quad (3.17)$$

corresponding to transformations in the density and spin channel, respectively. Note that to keep λ real, we must distinguish between the attractive and repulsive case and select the appropriate decoupling based on the sign of the interaction, i.e. Eq. (3.16) for $U < 0$ and Eq. (3.17) for $U > 0$. Choosing $\uparrow \equiv +1$, $\downarrow \equiv -1$ and factorizing the individual spin contributions for both cases we obtain

$$e^{|U|\Delta\tau(n_{i\uparrow}-\frac{1}{2})(n_{i\downarrow}-\frac{1}{2})} \propto \sum_{\phi_{li}=\pm 1} \prod_{s \in \{\uparrow, \downarrow\}} e^{\phi_{li}\lambda(n_{is}-\frac{1}{2})}, \quad U < 0, \quad (3.18)$$

$$e^{-U\Delta\tau(n_{i\uparrow}-\frac{1}{2})(n_{i\downarrow}-\frac{1}{2})} \propto \sum_{\phi_{li}=\pm 1} \prod_{s \in \{\uparrow, \downarrow\}} e^{s\phi_{li}\lambda n_{is}}, \quad U > 0. \quad (3.19)$$

Hence, the kinetic and interaction matrices after the Hirsch transformation are given by

$$\begin{aligned} [T]_{ij,ss'} &= \delta_{ss'} \begin{cases} -t & \text{if } i \text{ and } j \text{ are nearest neighbors,} \\ 0 & \text{otherwise,} \end{cases} \\ [V]_{ij,ss',l} &= -\frac{\delta_{ij}\delta_{ss'}}{\Delta\tau} \begin{cases} \lambda\phi_{li} - \Delta\tau\mu, & U < 0, \\ \lambda s\phi_{li} + \Delta\tau\mu, & U > 0, \end{cases} \end{aligned} \quad (3.20)$$

Let us now compare the attractive and repulsive scenario. Since different spin channels do not mix in either case, the Green's function is block diagonal, i.e. $[G]_{s,s'} = \delta_{ss'}\tilde{G}_s$, and, consequently, the DQMC determinant factorizes into separate spin contributions

$$\det G^{-1} = \det \tilde{G}_{\uparrow}^{-1} \tilde{G}_{\downarrow}^{-1}. \quad (3.21)$$

For negative U , one further observes that the interaction matrix V is independent of spin such that $\tilde{G}_{\uparrow} = \tilde{G}_{\downarrow}$. In this case, the statistical weight reads⁴⁴

$$w_{\phi} = e^{-\sum_{i,l} \lambda\phi_{li}} \left(\det \tilde{G}^{-1} \right)^2. \quad (3.22)$$

⁴³The identities are easily proven by comparing both sides for all possible eigenvalues of the density operators on both sides.

⁴⁴The bosonic factor in Eq. (3.22) is due to the constant $1/2$ in Eq. (3.18).

As per construction in Eq. (3.18), the Green's function is real and it trivially follows that the attractive Hubbard model is sign-problem free. For positive U on the other hand, there is an explicit spin dependence of V in Eq. (3.20). Hence, the generic repulsive Hubbard model is sign-problematic and not efficiently solvable by DQMC.

However, as is famously known, the situation is different at half-filling on a bipartite lattice, such as the square or honeycomb lattices [90, 92, 93, 95, 96, 122]. In this case, the Hubbard model, Eq. (3.15), is invariant under the particle-hole transformation⁴⁵

$$\text{PH} : \begin{cases} c_{i\uparrow} \rightarrow (-1)^i c_{i\uparrow}^\dagger, \\ c_{i\downarrow} \rightarrow (-1)^i c_{i\downarrow}^\dagger, \end{cases} \quad (3.23)$$

where $(-1)^i$ indicates alternating signs between the two different sublattices [93]. Note that the density operator transforms as $n_{is} \rightarrow 1 - n_{is}$ under PH and hence switches from particle to hole character. Focusing on the spin-up channel, one sees that the Hirsch transformed interaction term $H_{I\uparrow} = \phi_{li} \lambda n_{i\uparrow}$, Eq. (3.19), maps onto $-H_{I\uparrow} + \phi_{li} \lambda = H_{I\downarrow} + \phi_{li} \lambda$. This implies that $V_\uparrow \rightarrow V_\downarrow$ and, up to an unimportant bosonic factor, $\det G_\uparrow^{-1} \propto \det G_\downarrow^{-1}$ [93, 149]. Accordingly the model is sign-problem free.

Let us point out that one may employ another particle-hole transformation in only one of the spin channels, [145]

$$\text{PH}' : \begin{cases} c_{i\uparrow} \rightarrow c_{i\uparrow}, \\ c_{i\downarrow} \rightarrow (-1)^i c_{i\downarrow}^\dagger. \end{cases} \quad (3.24)$$

While the kinetic part of the repulsive Hubbard model is invariant under this transformation (for nearest neighbor hopping on a bipartite lattice) the chemical potential term transforms as $\mu n_i \rightarrow \mu(m_i + 1)$ where $m_i = n_{i\uparrow} - n_{i\downarrow}$ is the local magnetization. Hence, this term is only invariant for $\mu = 0$ and takes the role of a Zeeman field away from half-filling. The Hirsch transformed interaction term, Eq. (3.17), which is proportional to m_i , maps onto the density operator, $m_i \rightarrow n_i - 1$, in which both spin contributions enter with the same sign. The interaction matrix V , Eq. (3.20), hence becomes spin independent and the model is sign-problem free. In fact, performing the PH' transformation before the Hirsch decomposition one realizes that the interaction term in Eq. (3.15) changes sign, i.e. $U \rightarrow -U$. Hence, the repulsive Hubbard model at half-filling on a bipartite lattice is PH'-equivalent to the attractive Hubbard model and results obtained for one model can be transferred to the other.

3.4 Systematic model classification

To systematically identify and classify sign-problem free models we take the viewpoint of the Hamiltonian as an arbitrary Hermitian random matrix, i.e.

⁴⁵The term “particle-hole transformation” has different meanings in the literature [162]. Irrespective of details, we will use it for general transformations that involve a mapping of creation operators onto annihilation operators. Examples are given by Eqs. (3.23) and (3.24).

$H \in \mathbb{C}^{n \times n}$ and $H = H^\dagger$. The strategy then is to establish rigorous criteria for H and its coefficients that ensure the non-negativity of the statistical weights.

3.4.1 Theorem of T -invariant decomposition

The factorization of the determinant into equal or complex conjugated parts in the spin-fermion model⁴⁶ (for $N_\phi \leq 2$) and the attractive Hubbard model is no coincidence but a consequence of inherent antiunitary symmetries relating different fermion flavor channels, i.e. spin up and down for the latter. As stated in the following theorem, proven by Wu and Zhang in Ref. [142], the presence of an antiunitary symmetry that squares to the negative identity operator is sufficient to guarantee the absence of the sign-problem for an otherwise unrestricted⁴⁷ Hamiltonian.

Theorem 1 (Wu and Zhang’s theorem) *If there exists an antiunitary operator $T = UC$, where U is unitary and C is complex conjugation, such that*

$$THT^{-1} = H, \quad T^2 = U^2 = -\mathbb{1}, \quad (3.25)$$

then the eigenvalues λ_i of G^{-1} appear in complex conjugated pairs. Consequently, the fermion determinant is positive semidefinite,

$$\det G^{-1} = \prod_i |\lambda_i|^2 \geq 0. \quad (3.26)$$

Note that if G^{-1} is Hermitian, the statement above reduces to Kramers’ theorem of quantum mechanics [142, 163]. For this reason, the set of sign-problem free models covered by Theorem 1 is referred to as “Kramers class” [143].

Two important remarks are in order. Conventional DQMC, as described in Ch. 2, is formulated in terms of complex⁴⁸ fermions and requires the number of the latter to be conserved. In this context, Wu and Zhang’s theorem implicitly requires charge conservation [143, 147]. For a model to be in the Kramers class it must therefore not contain any pairing terms that explicitly break the $U(1)$ symmetry. Furthermore, it is important to realize that in the application of Theorem 1 the Hamiltonian H must be considered in – DQMC compatible – fermion bilinear form, that is, after any Hubbard-Stratonovich (or Hirsch) transformation. This is because a decoupling of an inter-fermion interaction term may reduce the number of symmetries. As a consequence, the sign-problem is dependent on the particular HS decomposition and, in this sense, not a unique property of a given physical model⁴⁹.

⁴⁶For convenience we will refer to the sign-problem free two-band variant, Eq. (3.4), as “spin-fermion model” in what follows.

⁴⁷See our remarks below the theorem.

⁴⁸Complex fermions in contrast to real Majorana fermions.

⁴⁹We have already touched upon the fact by noting that the sign-problem is basis dependent.

Application: Hubbard model

Let us revisit the sign-problem structure of the Hubbard model by solely considering symmetries. Starting with the attractive case, $U < 0$, we note that the Hamiltonian, Eq. (3.15), is time reversal symmetric. Representing complex conjugation by C , the corresponding antiunitary operator is $T_{\text{TR}} = -i\sigma_y C$ with $T^2 = -1$ [142, 143]. Here the Pauli matrix σ_y operates in spin space. It is trivial to see that $T_{\text{TR}} H T_{\text{TR}}^{-1} = H$ since all constants in H are real, kinetic and chemical potential terms are essentially density operators⁵⁰ that are clearly T_{TR} -invariant, and n_{\uparrow} and n_{\downarrow} appear symmetrically in the Hubbard interaction. If we decouple the latter in the density channel, Eq. (3.16), this symmetry is preserved as the auxiliary Hirsch field couples to the density operator. As a consequence, the conditions of Wu and Zhang's theorem are fulfilled implying the absence of the sign-problem.

For $U > 0$, the situation is different. Both Hirsch decouplings, Eqs. (3.16) and (3.17), break time reversal symmetry as they lead to terms that are odd under T_{TR} : either λ is imaginary, $T_{\text{TR}} i T_{\text{TR}}^{-1} = -i$, (density channel) or the Hirsch field couples to the local magnetization, $T_{\text{TR}} m_i T_{\text{TR}}^{-1} = -m_i$ (spin channel). Consequently, Theorem 1 does not apply, indicating⁵¹ the sign-problem of the repulsive Hubbard model.

Application: Spin-fermion model

Besides the attractive Hubbard model, the theorem by Wu and Zhang allows us to prove the absence of the sign-problem for the spin-fermion model, Eq. (3.4). This includes the case of isotropic antiferromagnetism, i.e. $N_{\phi} = 3$, left out in the discussion in Sec. 3.2. The key observation is that, in virtue of the two-band construction, the action S_{SFL} is invariant under the augmented time reversal symmetry $T = T_{\text{TR}} \kappa_z = -i\sigma_y \kappa_z C$, in which κ_z is a Pauli matrix in band space associated with the index α in Eq. (3.4). Assigning the following values to the spin and band indices,

$$\begin{aligned} s : \uparrow &\equiv +1, \downarrow \equiv -1 \\ \alpha : x &\equiv +1, y \equiv -1, \end{aligned} \tag{3.27}$$

all density-like terms transform as $\psi_{\alpha is}^{\dagger} \psi_{\alpha js} \rightarrow (\alpha s)^2 \psi_{\alpha is}^{\dagger} \psi_{\alpha js}$ and are trivially symmetric under T . For the fermion-boson interaction one may calculate the

⁵⁰The kinetic terms involve operators on different sites i and j but this is irrelevant since T does not effect space indices.

⁵¹This is not a proof since Theorem 1 doesn't state that the absence of an antiunitary symmetry implies a sign-problem.

invariance as follows,

$$\begin{aligned}
 T \sum_{\mathbf{r}ss'} e^{i\mathbf{Q}\cdot\mathbf{r}} \phi_{\mathbf{r}} \cdot \psi_{x\mathbf{r}s}^\dagger \sigma_{ss'} \psi_{y\mathbf{r}s'} T^{-1} &= - \sum_{\mathbf{r}ss'} e^{-i\mathbf{Q}\cdot\mathbf{r}} \phi_{\mathbf{r}} \cdot \psi_{x,\mathbf{r},-s}^\dagger [\phi_{\mathbf{r}} \cdot \sigma]_{ss'}^* \psi_{y,\mathbf{r},-s'} \\
 &= - \sum_{\mathbf{r}ss'} e^{i\mathbf{Q}\cdot\mathbf{r}} \phi_{\mathbf{r}} \cdot \psi_{x,\mathbf{r},-s}^\dagger [\phi_{\mathbf{r}} \cdot \sigma]_{ss'}^* \psi_{y,\mathbf{r},-s'} \\
 &= - \sum_{\mathbf{r}ss'} e^{i\mathbf{Q}\cdot\mathbf{r}} \phi_{\mathbf{r}} \cdot \psi_{x,\mathbf{r},-s}^\dagger [\phi_{\mathbf{r}} \cdot (-\sigma_x, \sigma_y, \sigma_z)^T]_{ss'} \psi_{y,\mathbf{r},-s'} \\
 &= - \sum_{\mathbf{r}ss'} e^{i\mathbf{Q}\cdot\mathbf{r}} \phi_{\mathbf{r}} \cdot \psi_{x\mathbf{r}s}^\dagger [\phi_{\mathbf{r}} \cdot (-\sigma)]_{ss'} \psi_{y\mathbf{r}s'} \\
 &= \sum_{\mathbf{r}ss'} e^{i\mathbf{Q}\cdot\mathbf{r}} \phi_{\mathbf{r}} \cdot \psi_{x\mathbf{r}s}^\dagger \sigma_{ss'} \psi_{y\mathbf{r}s'},
 \end{aligned}$$

where we have shifted the summation indices $s \rightarrow -s$ and $s' \rightarrow -s'$ in the fourth line. Combined with the fact that $T^2 = (-i)^2(\sigma_y)^2(\kappa_x)^2 = -1$ we may conclude from Theorem 1 that the spin-fermion model is an element of the Kramers class and free of the sign problem. This includes the case $N_\phi = 3$ which can not be explained by a simple factorization of the determinant.

The antiunitary symmetry $T = T_{\text{TR}}\kappa_x$ also sheds more light onto the two-band construction of Sec. 3.2. From a sign-problem point of view, the crucial aspect of the latter is the change from intraband to interband interactions⁵². This way, the odd transformation behavior of the spin density bilinear under regular time reversal, T_{TR} , is compensated by an additional minus sign associated with κ_x in T . This balance is explicit in the last step above in form of two canceling signs.

3.4.2 Majorana time reversal symmetries

The theorem by Wu and Zhang establishes the importance of antiunitary symmetries in guaranteeing non-negative statistical weights. As it turns out their approach can be generalized to Majorana “time reversal”⁵³ (MTR) symmetries which then serve as a fundamental classification principle for sign-problem free models [147].

First, we switch to a Majorana fermion basis by rewriting the regular complex fermion operators c_i in terms of real Majorana operators γ_{1i}, γ_{2i} , (we neglect spin for convenience)

$$c_i = \frac{1}{2}(\gamma_{1i} + i\gamma_{2i}). \quad (3.28)$$

Here the latter obey $\{\gamma_{\alpha i}, \gamma_{\beta j}\} = 2\delta_{\alpha\beta}\delta_{ij}$ and square to the identity, $\gamma_{\alpha i}^2 = 1$. Of course, the definition above is not unique and is understood to be modulo unitary transformations. An Hamiltonian consisting of arbitrary fermion bilinear terms, that is after a potential Hubbard-Stratonovich decomposition of any fermion-fermion interaction terms, may now be expressed in this new basis as

$$H = \frac{1}{4}\gamma^T A_i \gamma, \quad (3.29)$$

⁵²Despite our wording it is understood from the discussion in Sec. 1.3 that the fermion bilinear doesn’t represent a fermion-fermion but a fermion-boson interaction.

⁵³As is custom in the literature [143, 145, 147], we will use the terms “time reversal” and “antiunitary” synonymously. Note, however, that this includes symmetries T with $T^2 = +\mathbb{1}$.

Number of anticommuting MTR symmetries	Sign-problem free	Sign-problematic
0	none	$\{I\} = \text{no symmetry}$
1	none	$\{T_1^+\}$ $\{T_1^-\}$
2	$\{T_1^+, T_2^-\} = \text{Majorana class}$ $\{T_1^-, T_2^-\} = \text{Kramers class}$	$\{T_1^+, T_2^+\}$
3	$\{T_1^+, T_2^+, T_3^-\}$ $\{T_1^+, T_2^-, T_3^-\}$ $\{T_1^-, T_2^-, T_3^-\}$	$\{T_1^+, T_2^+, T_3^+\}$
≥ 4	all	none

Table 3.1: “Periodic table” of sign-problem free and sign-problematic symmetry classes. Based on the (number of) anticommuting MTR symmetries T_i^\pm the table indicates whether a given model is guaranteed to be free of the sign-problem. Individual models that fall into sign-problematic symmetry classes may or may not be sign-problem free. Reproduced from Ref. [147].

where $\gamma = (\gamma_1, \gamma_2)^T$ and A is a $2N \times 2N$ complex coefficient matrix⁵⁴ which due to the Hermitian property of H is skew-symmetric, i.e. $A = -A^T$.

As demonstrated in seminal works by Li *et al.* [143, 146, 147] categorizing Hamiltonians by their set of MTR symmetries allows one to identify all symmetry classes that *must* be sign-problem free, Table 3.1. Because Majorana operators are real, time reversal transformations can be represented as $T^\pm = J^\pm C$ with J^\pm being real orthogonal (instead of unitary) and either symmetric or skew-symmetric, i.e. $(T_i^\pm)^2 = (J_i^\pm)^2 = \pm 1$ ⁵⁵. By considering all combinations of such operators T_i^\pm with $\{T_i^\pm, T_j^\pm\} = \pm 2\delta_{ij}$ one can show [147] that all models with ≥ 2 MTR symmetries, at least one of which is skew-symmetric, are free of the sign problem. In particular, there are two (and only two⁵⁶) fundamental symmetry classes that are guaranteed to be sign-problem free [143, 147],

$$\begin{aligned} \{T_1^-, T_2^-\} & \quad (\text{Kramers class}) \\ \{T_1^+, T_2^-\} & \quad (\text{Majorana class}). \end{aligned}$$

The case of two negative⁵⁷ MTR symmetries corresponds to the Kramers class associated with Wu and Zhang’s theorem. By pairing up T_1^- and T_2^- one can define a charge operator that commutes with the Hamiltonian [147]. Hence, the conditions of Theorem 1 are fulfilled: particle number is conserved and there

⁵⁴If spin is incorporated we group the 4 Majorana fermion operators per site into two pairs $\gamma^{(1)} = (\gamma_1, \gamma_2)^T$, $\gamma^{(2)} = (\gamma_3, \gamma_4)^T$ and, without loss of generality, take the lattice size to be reduced by half, i.e. $N/2$, such that A is always of size $2N \times 2N$. See the supplement material of Ref. [145] for more details.

⁵⁵Since J is orthogonal, i.e. $JJ^T = 1$, the condition $(J^\pm)^2 = \pm 1$ implies $J^+ = (J^+)^T$ and $J^- = -(J^-)^T$, respectively.

⁵⁶Higher symmetry classes can be reduced to the two fundamental classes.

⁵⁷We denote a MTR symmetry represented by a skew-symmetric (symmetric) operator T^- (T^+ as negative (positive).)

exists an antiunitary symmetry that squares to minus one. Note that a single antiunitary symmetry T_1^- , i.e. being in the class $\{T_1^-\}$, doesn't generally suffice to ensure the absence of the sign-problem. This confirms the implicit requirement of particle number conservation in Theorem 1 [143, 147].

The second fundamental category is the Majorana class $\{T_1^+, T_2^-\}$ which comprises models that have both a positive and a negative MTR symmetry. In this case, the coefficient matrix A can be block-diagonalized into the form [147]

$$A = \begin{pmatrix} B & \\ & \bar{B} \end{pmatrix} \quad (3.30)$$

where B and \bar{B} correspond to eigensectors of the real symmetric operator $P = T_1^+ T_2^-$ with $P^2 = \mathbb{1}$. To utilize this structure one may implement a Majorana flavor of the DQMC scheme – following Ref. [146] we dub this variant “Majorana QMC” (MQMC). Specifically, instead of evaluating the fermion trace in the complex fermion basis, Eqs. (2.18) and (2.20), we compute it in the Majorana basis, Eq. (3.28), to find [144–146]

$$w = \text{Tr} [e^{H_1} e^{H_2} \dots e^{H_k}] = \det(\mathbb{1} + e^{A_1} e^{A_2} \dots e^{A_k})^{\frac{1}{2}}. \quad (3.31)$$

Here we use the index $i = 1, 2, \dots, k$ to enumerate all occurring exponential fermion bilinears which arise from a Trotter-decomposition of different contributions to H , Eq. (2.17), as well as the imaginary time product chain, Eq. (2.18). For clarity, we have further incorporated all factors of $\Delta\tau$ into the definition of the H_i and left the boson field dependence implicit. A proof of the identity above is given in App. A of Ref. [146]. Note the occurrence of a square root which, intuitively, originates from the fact that a Majorana fermion represents only half of a complex fermion.

For models in the Majorana class, the block diagonal structure of A implies that the determinant in Eq. (3.31) factorizes into complex conjugated parts and is strictly positive semidefinite. This feature may be viewed as the Majorana analogon of the factorizations in Sec. 3.3. Importantly, and in contrast to the latter, the block sectors associated with B and \bar{B} in Eq. (3.30) may not correspond to complex fermion flavors, such as spin as in Eq. (3.21). As first demonstrated in Ref. [146] this allows one to study spinless one-band fermion models – hosting only a single fermion flavor – that feature Dirac semimetal to CDW quantum phase transitions. In these cases, natural candidates for the antiunitary symmetry T_1^+ are particle-hole symmetries which, in contrast to conventional time reversal symmetry T_{TR} , square to the (positive) identity. An important example is the Majorana reflection symmetry (MRS) $\Theta^+ = \tau_x C$ where τ_x is a Pauli matrix acting in Majorana operator space, i.e. $\gamma_1 \leftrightarrow \gamma_2$ [143, 145–147]. In terms of complex fermions, Eq. (3.28), the effect of Θ^+ is

$$\Theta^+ : c_j \rightarrow i c_j^\dagger, \quad (3.32)$$

clearly indicating the particle-hole character. Combined with the additional MTR symmetry $T^- = i\tau_y C$ it is precisely this MRS that puts the spinless fermion model of Ref. [146] into the Majorana class and, consequently, guarantees the

absence of the sign-problem. Similar arguments can be made for $SU(N = \text{odd})$ negative- U Hubbard models⁵⁸ on bipartite lattices [146].

A key difference between the Kramers and Majorana class is that particle number conservation is only required in the former. Natural examples in $\{T_1^+, T_2^-\}$ that utilize this extra freedom are models with explicit pairing terms. This includes time reversal symmetric topological superconductors with triplet pairing of $(p + ip)$ and $(p - ip)$ kind for spin-up and spin-down electrons, respectively [147]. Here a pairing term $\Delta c_{is}^\dagger c_{js}^\dagger$ clearly breaks the $U(1)$ -symmetry.

3.4.3 Majorana reflection positivity

The MTR symmetry classification by Li *et al.*, Tbl. 3.1, is both elegant and instructive. However, it is not complete since individual models in the sign-problematic symmetry classes may still be sign-problem free for reasons other than symmetry. In Ref. [145] Wei *et al.* have shown that it is possible to derive a condition for the absence of the sign-problem for coefficient matrices A with only one MTR symmetry T_1^+ . The idea is to expand on the symmetry constraint associated with T_1^+ , which is taken to be the Majorana reflection symmetry Θ^+ above, by further requiring Majorana reflection positivity (MRP) [145]. Leaving a precise definition to Ref. [145], the major result of this approach is summarized in the following theorem.

Theorem 2 (MRP theorem) *The probability weight w , Eq. (3.31), is positive semidefinite if all the coefficient matrices A are Majorana reflection positive kernels, that is they can be brought into the form*

$$A = \begin{pmatrix} B & iC \\ -iC^T & \bar{B} \end{pmatrix} \quad (3.33)$$

where B is skew-symmetric, $B^T = -B$, and C is Hermitian and either positive or negative semidefinite, i.e. $C \geq 0$ or $C \leq 0$.

First, we note that for $C = 0$ the condition for the coefficient matrices A_i Eq. (3.33) coincides with Eq. (3.30) and we recover the Majorana symmetry class. Hence, the latter is entirely covered by Theorem 2. The important new degree of freedom lies in the off-diagonal part of A . Note that the positive or negative semidefiniteness of C can not be enforced by symmetry alone [143–145].

Examples of models that fall under Theorem 2 and are not part of the Majorana class include generalized Kane-Mele-Hubbard models with spin-orbit coupling and spin-flip terms [145]. The latter explicitly break the time reversal symmetry of the model indicating the absence of an additional T^- MTR symmetry besides Θ^+ . Furthermore, all sign-problem free models studied by the fermion bag [154] or split orthogonal group approach [164] can also be explained by the MRP theorem above [144].

⁵⁸The repulsive Hubbard model away from half-filling is part of the symmetry class $\{T_1^+, T_2^+\}$. A detailed proof can be found in the supplement materials of Ref. [147].

3.4.4 A road to new sign-problem free models

A novel approach to classifying sign-problem free models based on the identification of a Lie group structure [165] in the definition of the probability weight w , Eq. (3.31), has recently been put forward by Wei in Ref. [144]. As we will demonstrate below⁵⁹ this framework unifies the discussion above by covering both fundamental MTR symmetry classes, i.e. the Kramers and Majorana class, as well as the MRP extension of $\{T^+\}$. Even more importantly, it allows us to identify a new class of sign-problem free models within the (sign-problematic) $\{T^-\}$ symmetry class.

We start by considering all skew-symmetric coefficient matrices A_i , Eq. (3.31), that are invariant under the transformation $A \rightarrow J_1^T \bar{A} J_1$ where J_1 is real orthogonal and either symmetric or skew-symmetric. In other words, the A_i are elements of $\{T^+\}$ or $\{T^-\}$ with $T = J_1 C$ and $J_1^2 = 1$ or $J_1^2 = -1$, respectively. It is straightforward to see that this condition already implies the realness of the statistical weight,

$$w \xrightarrow{T} w' = \det(\mathbb{1} + e^{J_1^T \bar{A}_1 J_1} e^{J_1^T \bar{A}_2 J_1} \dots e^{J_1^T \bar{A}_k J_1})^{\frac{1}{2}} \quad (3.34)$$

$$= \det(\mathbb{1} + J_1^T e^{\bar{A}_1} J_1 J_1^T e^{\bar{A}_2} J_1 \dots J_1^T e^{\bar{A}_k} J_1)^{\frac{1}{2}} \quad (3.35)$$

$$= \det(\mathbb{1} + J_1^T e^{\bar{A}_1} e^{\bar{A}_2} \dots e^{\bar{A}_k} J_1)^{\frac{1}{2}} \quad (3.36)$$

$$= \det(\mathbb{1} + e^{\bar{A}_1} e^{\bar{A}_2} \dots e^{\bar{A}_k})^{\frac{1}{2}} \quad (3.37)$$

$$= \overline{\det(\mathbb{1} + e^{A_1} e^{A_2} \dots e^{A_k})^{\frac{1}{2}}} \quad (3.38)$$

$$= \bar{w}, \quad (3.39)$$

where we have used the orthogonality of J_1 and the cyclic property $\det(\mathbb{1} + ABC) = \det(\mathbb{1} + BCA)$.

Similar to the MRP scenario above, the strategy is now to establish non-symmetry conditions that ensure the non-negativity of w . To that end, we note that given a Hermitian matrix η with $\eta^2 = \mathbb{1}$, all the complex Majorana coefficient matrices A with the property

$$\eta A + A^\dagger \eta \leq 0 \quad (3.40)$$

constitute a Lie contraction semigroup (LCS) [144]: the elements are exponentials $g = e^A$ and the semigroup operation is the regular matrix multiplication. Note that this structure strongly resembles the second term in the MQMC determinant, Eq. (3.31). In terms of the group elements g the condition above reads $g^\dagger \eta g \leq \eta$. Interpreting the left hand side as a generalized “inner product” the elements g preserve (equality) or contract (inequality) the “length” of a vector associated with the metric η . Hence, it is clear that the η -unitary group, associated with pure “rotations”, is a subgroup of the contraction semigroup and corresponds to the case of strict equality in Eq. (3.40). It is generated by those matrices A that are invariant under the sign-preserving transformation $A \rightarrow \eta^\dagger \bar{A} \eta$, in which case

$$\eta A + A^\dagger \eta = \eta A - \eta A = 0. \quad (3.41)$$

⁵⁹Our presentation closely follows but extends the discussion in Ref. [144].

This suggests a polar decomposition of the semigroup elements, i.e. $g = g_U e^{A_0}$ where g_U is η -unitary and A_0 transforms as $A_0 \rightarrow -\eta^\dagger \bar{A}_0 \eta$ - i.e. it changes sign [144]. The latter implies

$$\eta A_0 + A_0^\dagger \eta = 2\eta A_0 \leq 0 \quad \Rightarrow \quad \eta A_0 = -\bar{A}_0 \eta \leq 0 \quad (3.42)$$

such that the A_0 form an invariant cone under action of the unitary group [144]. Importantly, for strict contraction elements one has $\eta A_0 < 0$ and all eigenvalues of g are, in magnitude, smaller than 1.

To utilize the above, we consider a real orthogonal matrix J_2 and specifically choose the metric η as $\eta = iJ_2$. Note that in contrast to J_1 above, J_2 must be skew-symmetric, i.e. $J_2^2 = -\mathbb{1}$, as the invariant cone would otherwise only contain the zero element (See App. A of Ref. [144] for details). Similarly, to avoid this trivial case we must require that J_1 and J_2 anticommute, i.e. $\{J_1, J_2\} = 0$ [144]. Postponing a proof of the non-negativity of w under these conditions we obtain the following theorem.

Theorem 3 (LCS theorem) *If there exist real orthogonal and anticommuting matrices J_1^\pm and J_2^- with $(J_1^\pm)^2 = \pm\mathbb{1}$ and $(J_2^-)^2 = -\mathbb{1}$ such that the Majorana coefficient matrices A fulfill the following conditions*

$$(J_1^\pm)^T A J_1^\pm = \bar{A}, \quad (3.43)$$

$$i(J_2^- A - \bar{A} J_2^-) \leq 0, \quad (3.44)$$

then the probability weight w , Eq. (3.31), is positive semidefinite.

The argument that Eq. (3.44) of the LCS theorem indeed implies non-negative statistical weights is based on the continuity of the Lie semigroup [144]. Note that Eq. (3.44) corresponds to the invariant cone condition in Eq. (3.42). As seen above, the eigenvalues of strict contraction elements $\exp(A)$ are smaller than 1. This also trivially holds for any matrix product chain of semigroup elements, i.e. $e^{A_1} e^{A_2} \dots e^{A_k}$, involving the same. As a consequence all eigenvalues of the matrix $M = \mathbb{1} + e^{A_1} e^{A_2} \dots e^{A_k}$ are non-zero which, when applied to Eq. (3.31), implies that the probability weight $w = \det(M)^{\frac{1}{2}}$ does not vanish. In this case, one is left with two possibilities: w is either positive or negative definite. Given that the conditions of Theorem 3 also include known sign-problem free symmetry classes – we will explicitly demonstrate this fact below – the Lie semigroup continuity rules out the second scenario⁶⁰. Therefore, w is generally positive semidefinite.

Let us highlight how the semigroup framework reduces to the classification schemes discussed in the previous sections. An illustration of these relations is provided in Fig. 3.3. In the case of strict equality Eqs. (3.43) and (3.44) represent two independent antiunitary symmetry conditions. We may therefore define the MTR operators $T_1^\pm = J_1^\pm C$ and $T_2^- = J_2^- C$ and conclude that, depending on the symmetry of J_1^\pm , the unitary subgroup of the contraction semigroup coincides with the Majorana or Kramers symmetry class, i.e. $\{T_1^\pm, T_2^-\}$ or

⁶⁰As a continuous function of the Majorana coefficients $[A]_{ij}$ the probability weight $w(A)$ never vanishes and is known to be non-negative for a particular choice of A .




Figure 3.3: Conceptual illustration of model space divided into sign-problem free classes. The inner circles (light and dark blue) correspond to the fundamental MTR symmetry classes of Table 3.1 and strict equality in Eq. (3.44) of Theorem 3. The outer circles (beige and red) represent the Lie contraction semigroups corresponding to $J_1^2 = +1$ and $J_1^2 = -1$ and to subsets of the symmetry classes $\{T_1^+\}$, $\{T_1^-\}$, respectively. The former case $J_1^2 = +1$ coincides with the MRP class covered by Theorem 2. The models in Eqs. (3.52) and (3.59) lie inside of the red circle but outside of the Kramers class.

$\{T_1^-, T_2^-\}$. As discussed in Sec. 3.4.2, the latter corresponds to the sign-problem free models covered by Wu and Zhang’s theorem, Th. 1. Note that for models in the intersection of $\{T_1^+, T_2^-\}$ and $\{T_1^-, T_2^-\}$, see Fig. 3.3, it is possible to choose either a symmetric J_1^+ or skew-symmetric J_1^- , if necessary in an appropriately chosen Majorana basis. Examples include the attractive Hubbard model and the repulsive Hubbard model on a bipartite lattice at half-filling. One can also recover the MRP form of the Majorana coefficient matrices, Eq. (3.33), by considering the case of a symmetric J_1^+ . Specifically choosing $J_1^+ = \tau_x$ and $J_2^- = i\tau_y$ [144] in the Majorana basis of Eq. (3.28) the first condition, Eq. (3.43), restricts models to the symmetry class $\{T_1^+\}$ and the second condition, Eq. (3.44) reduces to the semidefiniteness criterion of Theorem 2 for the off-diagonal matrix elements.

New sign-problem free models

Having established these connections, one realizes that the semigroup framework defines a novel class of sign-problem free models: the strict contraction elements in case of a skew-symmetric J_1^- , Fig. 3.3. These models are part of the symmetry class $\{T_1^-\}$ and, as a function of the coefficients, have the Kramers class as a limiting case (strict equality in Eq. (3.44)). To the end of deriving specific sign-problem free Hamiltonians it is natural to consider two complex fermion flavors and choose $J_1 = i\sigma_y$, i.e.

$$J_1 : \begin{cases} c_{i\uparrow} \rightarrow c_{i\downarrow} \\ c_{i\downarrow} \rightarrow -c_{i\uparrow} \end{cases}, \quad (3.45)$$

as in the conventional time reversal symmetry, T_{TR} in Sec. 3.4.1. In the adjusted Majorana basis

$$\begin{aligned} c_{i\uparrow} &= \frac{1}{2}(\gamma_{1i} + i\gamma_{2i}), \\ c_{i\downarrow} &= \frac{1}{2}(\gamma_{4i} + i\gamma_{3i}), \end{aligned} \quad (3.46)$$

the unitary matrix J_1 thus takes the form $J_1 = \tau_x \otimes i\tau_y$ and operates on the vector $\gamma = (\gamma_1, \gamma_2, \gamma_3, \gamma_4)^T$. The first condition of the LCS theorem, Eq. (3.43), then selects arbitrary bilinear Hamiltonians of the symmetry class $\{T_1^-\}$ which can generally be written as [144]

$$H = \frac{1}{4}\gamma^T A \gamma = H_h + H_p, \quad (3.47)$$

$$H_h = (c_{\uparrow}^{\dagger}, c_{\downarrow}^{\dagger}) M \begin{pmatrix} c_{\uparrow} \\ c_{\downarrow} \end{pmatrix} - (c_{\uparrow}, c_{\downarrow}) M^T \begin{pmatrix} c_{\uparrow}^{\dagger} \\ c_{\downarrow}^{\dagger} \end{pmatrix}, \quad (3.48)$$

$$H_p = (c_{\uparrow}, c_{\downarrow}) R J_1 \begin{pmatrix} c_{\uparrow} \\ c_{\downarrow} \end{pmatrix} - (c_{\uparrow}^{\dagger}, c_{\downarrow}^{\dagger}) S J_1 \begin{pmatrix} c_{\uparrow}^{\dagger} \\ c_{\downarrow}^{\dagger} \end{pmatrix}. \quad (3.49)$$

Here H_h and H_p are hopping-⁶¹ and pairing-like contributions, respectively. Furthermore, $J_1 = i\sigma_y$, as stated above, and M , R , and S are complex coefficient matrices such that RJ_1 and SJ_1 are skew-symmetric. Upon choosing $J_2 = i\sigma_x$, i.e.

$$J_2 : \begin{cases} c_{i\uparrow} \rightarrow ic_{i\downarrow} \\ c_{i\downarrow} \rightarrow ic_{i\uparrow} \end{cases}, \quad (3.50)$$

which in Majorana basis reads $J_2 = -i\tau_y \otimes \mathbb{1}_2$, the second condition of Theorem 3, Eq. (3.44), translates into the requirement that R and S are Hermitian and positive semidefinite⁶² [144], i.e.

$$R^{\dagger} = R \geq 0, \quad S^{\dagger} = S \geq 0. \quad (3.51)$$

Under these conditions, all complex fermion Hamiltonians of the form of Eq. (3.47) are guaranteed to be sign-problem free.

Note that besides the antiunitary symmetry associated with J_1 the contribution H_h conserves particle number. Hence we conclude, in accordance with Wu and Zhang's theorem, Th. 1, that Hamiltonians with $H_p = 0$ are part of the Kramers class. Rewriting the condition as $R = S = 0$, it is clear from our discussion above that these models correspond to the unitary subgroup of the LCS whereas strict contraction elements $H = H_p$ with $R > 0, S > 0$ form the invariant cone. In this light, a specific example of a sign-problem free model that lies in the symmetry

⁶¹A more precise description is that H_h is particle number conserving as, besides hopping terms, it may include spin-flip and chemical potential terms.

⁶²Note that via a Majorana basis transformation this condition may be turned into negative semidefiniteness, compare Theorem 2.

class $\{T_1^-\}$ and is part of the novel J_1^- semigroup is the repulsive Hubbard model at half-filling on a bipartite lattice with an explicit singlet-pairing term, [144]

$$H = -t \sum_{\langle i,j \rangle, s} c_{is}^\dagger c_{js} + t_\perp \sum_i \left(c_{i\uparrow}^\dagger c_{i\downarrow}^\dagger + \text{h.c.} \right) - U \sum_i \left(n_{i\uparrow} - \frac{1}{2} \right) \left(n_{i\downarrow} - \frac{1}{2} \right) \quad (3.52)$$

However, as elaborated around Eq. (3.31) conventional DQMC, Ch. 2, implicitly requires the absence of any explicit $U(1)$ symmetry breaking terms. For this reason the Hamiltonian above may only be treated by MQMC.

Particle-hole transformation

To derive a model of the J_1^- semigroup outside of the Kramers class which, at the same time, is amenable to DQMC simulation we must consider a different operator J'_1 . Taking $J_1 = \tau_x \otimes i\tau_y$ from above as a reference point, a useful strategy is to employ the particle-hole transformation PH', Eq. (3.24), where we drop the sublattice factor $(-1)^i$. For Majorana fermions this corresponds to the sign change $\gamma_3 \rightarrow -\gamma_3$ such that J_1 becomes $J'_1 = i\tau_y \otimes \tau_x$. Similarly, one finds that J_2 , Eq. (3.50), turns into $J_2 \rightarrow J'_2 = i\tau_y \otimes \tau_z$. In terms of complex fermions this reads

$$J'_1 : \begin{cases} c_{i\uparrow} \rightarrow c_{i\downarrow}^\dagger \\ c_{i\downarrow} \rightarrow -c_{i\uparrow}^\dagger \end{cases}, \quad J'_2 : \begin{cases} c_{i\uparrow} \rightarrow -ic_{i\downarrow}^\dagger \\ c_{i\downarrow} \rightarrow ic_{i\uparrow}^\dagger \end{cases}. \quad (3.53)$$

Under the transformation PH' the spin- \downarrow channel of the hopping term in Eq. (3.52) acquires a minus sign and, as seen in Sec. 3.3, the Hubbard interaction becomes attractive. More importantly, the contributions H_h and H_p in Eq. (3.47) mix and the singlet-pairing term in Eq. (3.52), associated with the invariant cone, turns into an onsite spin-flip term, i.e. $c_{i\uparrow}^\dagger c_{i\downarrow}^\dagger \rightarrow c_{i\uparrow}^\dagger c_{i\downarrow}$. Note that in the original case, Eq. (3.47), such a spin-flip term would have corresponded to the unitary subgroup. In summary, the PH'-transformed partner model of Eq. (3.52) reads

$$\begin{aligned} H' &= H_\uparrow^0 - H_\downarrow^0 + t_\perp \sum_i \left(c_{i\uparrow}^\dagger c_{i\downarrow} + \text{h.c.} \right) + U \sum_i \left(n_{i\uparrow} - \frac{1}{2} \right) \left(n_{i\downarrow} - \frac{1}{2} \right), \\ H_s^0 &= -t \sum_{\langle i,j \rangle} c_{is}^\dagger c_{js}. \end{aligned} \quad (3.54)$$

It is straightforward to verify that H' , when Hirsch decoupled in the spin channel⁶³, Eq. (3.17), fulfills the conditions of the LCS theorem with J'_1 and J'_2 as specified above. To demonstrate this, we note that the Hamiltonian has the general form

$$H = \begin{pmatrix} c_\uparrow^\dagger \\ c_\downarrow^\dagger \end{pmatrix} \begin{pmatrix} T & T_\perp \\ T_\perp & -T \end{pmatrix} \begin{pmatrix} c_\uparrow \\ c_\downarrow \end{pmatrix}, \quad (3.55)$$

⁶³As density terms in different spin channels must appear with opposite sign we need to decouple in the spin- rather than the density channel.

where $[T]_{ij}$ and $[T_\perp]_{ij}$ are coefficient matrices corresponding to conventional hopping and spin-flip terms, respectively. For generality, we allow T to be arbitrary and complex despite its simple structure in H_s^0 above. Choosing the Majorana fermion basis as in Eq. (3.46) and dropping irrelevant numerical prefactors, we find that the Hamiltonian may be written as $H \sim \gamma^T A \gamma$, where

$$A = i \begin{pmatrix} \text{Im}(T) & \text{Re}(T) & T_\perp & 0 \\ -\text{Re}(T) & \text{Im}(T) & 0 & -T_\perp \\ -T_\perp & 0 & -\text{Im}(T) & \text{Re}(T) \\ 0 & T_\perp & -\text{Re}(T) & -\text{Im}(T) \end{pmatrix}. \quad (3.56)$$

For $J_1 = J'_1 = i\tau_y \otimes \tau_x$, the first condition of the LCS theorem, Th. 3, is fulfilled as

$$J_1^T A J_1 = -A = \bar{A}. \quad (3.57)$$

Note that the equality only holds because T appears with opposite sign in the different spin channels of Eq. (3.55). Choosing $J_2 = J'_2 = i\tau_y \otimes \tau_z$, the second condition of Th. 3 becomes

$$i(J_2^- A - \bar{A} J_2^-) = T_\perp \delta_{\alpha\beta} \leq 0, \quad (3.58)$$

where $\alpha, \beta \in \{1, 2, 3, 4\}$ are Majorana fermion indices. Clearly, as per LCS theorem, this implies sign-problem free DQMC if T_\perp is negative semidefinite. Since a sign change of T_\perp in H may be absorbed into a redefinition of the Majorana basis, i.e. $\gamma_2 \rightarrow -\gamma_2$ and $\gamma_3 \rightarrow -\gamma_3$, without affecting the first condition of the theorem, the absence of the sign-problem is also guaranteed if T_\perp is positive semidefinite⁶⁴. This requirement is trivially fulfilled for the Hamiltonian H' , Eq. (3.54), since in this case $T_\perp = t_\perp \delta_{ij} \geq 0$. Note that as a function of the coefficient t_\perp , one has strict equality or inequality in Eq. (3.58), indicating the unitary subgroup of the LCS for $t_\perp = 0$. As a consequence DQMC simulations of H' are free of the sign-problem and for $t_\perp > 0$ in the novel J_1^- semigroup class⁶⁵, Fig. 3.3.

Interacting flat band model

It is clear from the demonstration above that beyond the requirement that hopping terms for spin- \uparrow and spin- \downarrow fermions are identical in magnitude and have opposite sign, the conditions of the LCS theorem do not impose any restrictions on the kinetic Hamiltonian. One may therefore derive further interesting sign-problem free models in the J_1^- semigroup by varying H_s^0 in Eq. (3.54). As an explicit example, one may consider the following triangular lattice bilayer system,

⁶⁴Compare this to the semidefiniteness condition in the MRP theorem, Th. 2. We have numerically verified this statement for the flat band in Eq. (3.59).

⁶⁵Note that at the time of this writing we can not rule out the possibility that H' lies in the intersection of the MRP and J_1^- classes in Fig. 3.3 since a Majorana basis transformation may turn J_1^- into J_1^+ .

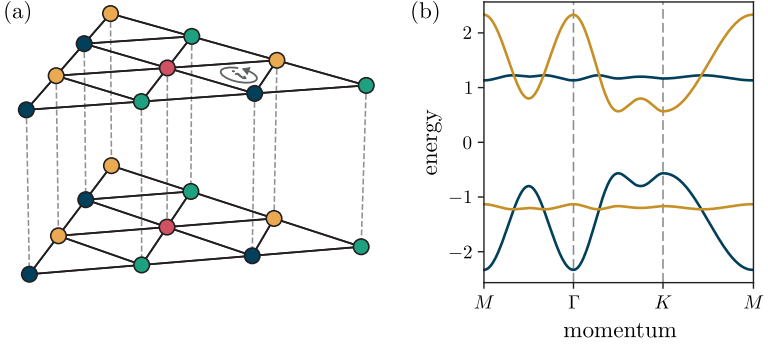


Figure 3.4: The triangular bilayer structure of model H_{FB} , Eq. (3.59), (a) and its non-interacting ($U = 0$) band structure along a high-symmetry path of the Brillouin zone (b). The parameters in (b) are $t = 1$, $t' = 0.2$, and $t_{\perp} = 0$ and different colors correspond to different spin-layers in (a). Each layer features a flat and a dispersive band according to the two-site unit-cell. Due to the occurrence of $-H_{\downarrow}^0$ the combined band structure is symmetric around $E = 0$.

Fig. 3.4a, based on the non-interacting Hamiltonian studied in Ref. [166],

$$\begin{aligned}
 H_{\text{FB}} &= H_{\uparrow}^0 - H_{\downarrow}^0 + t_{\perp} \sum_i \left(c_{i\uparrow}^{\dagger} c_{i\downarrow} + \text{h.c.} \right) + U \sum_i \left(n_{i\uparrow} - \frac{1}{2} \right) \left(n_{i\downarrow} - \frac{1}{2} \right), \\
 H_s^0 &= t \sum_{\langle i,j \rangle_{\gamma}} \left(\sigma_{\gamma} c_{js}^{\dagger} c_{is} + \text{h.c.} \right) + t'' \sum_{\langle\langle i,j \rangle\rangle} \left(c_{js}^{\dagger} c_{is} + \text{h.c.} \right) - \mu \sum_i n_i.
 \end{aligned} \tag{3.59}$$

Here the terms proportional to t and t'' represent nearest and third nearest neighbor hopping whereas the former is direction dependent and complex⁶⁶, $\gamma = x, y, z$. The 2×2 dimensional matrix structure of H_s^0 corresponds to the two-site unit cell of the triangular lattice [166]. As illustrated in Fig. 3.4b, each individual spin-layer features a flat and a dispersive band both of which are topological and come equipped with a non-trivial Chern number ± 1 [166]. For $U = 0$, Kourtis *et al.* [166] have observed the formation of a fractionalized Chern insulating state upon adding nearest-neighbor interactions V between the (spinless) fermions within a single sheet.

For specific parameter values, the sign-problem free model H_{FB} represents a platform for studying interactions between topological flat bands in a numerically exact manner. The physics associated with the latter has been in the spotlight of condensed matter research at least since the recent discovery of superconductivity in two-dimensional materials with moiré superlattices [167–171]. By tuning the chemical potential μ , which due to sign structure of $H_{\uparrow}^0 - H_{\downarrow}^0$ plays the role of a Zeeman field, one can readily shift the flat bands in Fig. 3.4b in opposite directions and make them overlap at the Fermi energy ($\mu \approx -2.4$). The interlayer spin-flip amplitude t_{\perp} can be seen as a magnetic field parallel to the layers and

⁶⁶For $\gamma = y$, σ_y is purely complex.

may be used to split and hybridize bands associated with different spin flavors. Although our initial DQMC simulations of H_{FM} do not show any indication of superconductivity – but instead a tendency towards ferromagnetic ordering – the same has been theoretically predicted and numerically observed [172–174] in similar flat band scenarios. The question of what sets the scale T_c in this case is an active matter of research [173, 175].

Beyond the specific Hamiltonians in Eqs.(3.52) and (3.54), the LCS scheme provides a general route towards identifying novel physical models amenable to sign-problem free DQMC simulations. While we have implemented the bilayer structure shown in Fig. 3.4a in terms of spin, one may just as well replace the spin indices in Eq. (3.55) by indices corresponding to, for instance, sublattice or pseudospin degrees of freedom. In this case, an onsite Hubbard term turns into a neighbor interaction between sites of different sublattices. Combined with the freedom of choosing the fermion kinetics, H_s^0 , this indicates the richness of the novel J_1^- semigroup class and the diversity of the physical models it contains.

4 | Aspects of DQMC implementation

Supplementing the general description of determinant quantum Monte Carlo in Ch. 2, it is the purpose of this chapter to present the most important numerical aspects of practical DQMC implementations. As we demonstrate below, numerical instabilities inherent to the DQMC method due to finite machine precision require careful stabilization of matrix products and inversions [P2, 90, 120, 122, 136, 176–178] in the computation of equal-time and time-displaced Green’s functions, Eqs. (2.29) and (2.34). We highlight the conceptual origin of these instabilities and present a systematic empirical assessment of various proposed stabilization schemes [122, 136, 176, 178, 179] with respect to accuracy and efficiency. Note that our presentation in Sec. 4.1 is a close adaption⁶⁷ of the corresponding publication by the author of this thesis, Ref. [P2]. Concrete implementations of the discussed techniques are provided in the open-source software library `StableDQMC.jl`, Ref. [S3].

The second major focus of this chapter is computational efficiency. For a naive DQMC implementation, the computational cost of a spatial lattice sweep of local updates increases with the fourth power of the number of lattice sites, $O(\beta N^4)$. This scaling severely limits the range of accessible system sizes, even when simulations are run on state of the art supercomputers, and therefore asks for systematic performance optimizations as well as techniques to reduce the impact of finite-size effects. In Sec. 4.2, we present an efficient algorithm for performing local and global DQMC updates based on intermediate Green’s function caching [90, 120, 122, 149, 161] and the checkerboard decomposition method [90, 122]. This way, the computational complexity of a DQMC sweep is reduced to the optimal $O(\beta N^3)$ scaling [90, 120, 122]. Furthermore, we discuss methods to improve the quality of simulations for a fixed system size, in particular the utilization of a magnetic flux-quantum [31, 32, 149, 161] and the imposition of twisted fermionic boundary conditions [149, 180, 181].

Finally, we will conclude this chapter with general notes on the extensive DQMC simulations underlying the physical results presented in Chs. 5 and 7 and the deployed DQMC implementation. Specifically, we will provide a performance benchmark of the open-source Julia code [S1] developed for this thesis as well as information about the utilized supercomputing facilities and scientific long-term data storage.

⁶⁷Apart from a general restructuring, we have extended the discussion of the time-displaced Green’s function in Sec. 4.1.3.

4.1 Stable linear algebra

Although conceptually straightforward, care has to be taken in the implementation of DQMC because of inherent numerical instabilities arising from ill-conditioned matrix exponentials, such as Eqs. (3.7) and (3.9). Over time, stabilization schemes [122, 136, 176, 178, 179] based on various matrix factorizations, such as the singular value decomposition (SVD) [124, 182], the modified Gram-Schmidt method [122, 124], and the QR decomposition [122, 124], have been proposed for lifting these numerical issues. In this section we review a subset of these techniques and systematically compare them with respect to accuracy and speed. Unless otherwise noted, our test platform is the one-dimensional Hubbard model, Eq. (3.15), at half-filling,

$$H = - \sum_{\langle i,j \rangle, s} \left(c_{is}^\dagger c_{js} + \text{h.c.} \right) + U \sum_i \left(n_{i\uparrow} - \frac{1}{2} \right) \left(n_{i\downarrow} - \frac{1}{2} \right), \quad (4.1)$$

where we set the hopping amplitude to unity, $t = 1$ ⁶⁸.

4.1.1 Time slice matrix multiplications

To showcase the first kind of numerical instability arising in the DQMC framework we focus on the non-interacting part of Eq. (4.1),

$$H = - \sum_{\langle i,j \rangle, s} \left(c_{is}^\dagger c_{js} + \text{h.c.} \right). \quad (4.2)$$

As seen in Eq. (2.21), the building block of the equal-time Green's function G_ϕ is a matrix chain multiplication of imaginary time slice matrices B_t . Because of the absence of interactions in Eq. (4.2), this chain reads⁶⁹

$$B(\beta, 0) \equiv B_M B_{M-1} \cdots B_1 = \underbrace{B B \cdots B}_{M \text{ factors}}. \quad (4.3)$$

As shown in Fig. 4.1, a naive computation of $B(\beta, 0)$ strikingly fails for $\beta \geq \beta_c \approx 10$. Leaving a discussion of the stabilization of the computation to the next section, let us highlight the origin of this instability. The eigenvalues of the non-interacting system are readily given by

$$\epsilon_k = -2t \cos(k), \quad (4.4)$$

such that energy values are bounded by $-2t \leq \epsilon_k \leq 2t$. A single positive definite time slice matrix $B = e^{-\Delta\tau T}$ therefore has a condition number [124] of the order of $\kappa \approx e^{4|t|\Delta\tau}$ and, consequently, $B(\tau, 0)$ has $\kappa \approx e^{4|t|M\Delta\tau} = e^{4|t|\beta}$. This implies that the numerical scales present in $B(\tau, 0)$ broaden exponentially at low temperatures $T = 1/\beta$ leading to inevitable roundoff errors due to finite machine precision.

⁶⁸We consider the canonical discrete Hirsch decoupling in the spin channel, Eq. (3.17).

⁶⁹For simplicity, we will take the B_t to be independent of imaginary time throughout our purely numerical analysis even in the presence of interactions. This amounts to working with a constant bosonic field $\phi = \text{const.}$




Figure 4.1: Numerical instabilities due to finite machine precision arising in the calculation of the time slice matrix chain product $B_M B_{M-1} \cdots B_1$ for free fermions in one dimension, Eq. (4.2). Different lines represent logarithmic singular values as observed in naive (green) and arbitrary precision computations (orange) for a $N = 4$ system. Due to (quasi) degeneracies only 5 out of 8 singular values are visually distinguishable. The dashed line (grey) indicates the expected floating point accuracy⁷⁰. Resized from Ref. [P2] under CC-BY 4.0.

We can estimate the expected inverse temperature of this breakdown for double floating-point precision according to the IEEE 754 standard [183], implemented in Julia in the data type `Float64`, by solving $\kappa(\beta) \sim 10^{-17}$ for β_c . One finds $\beta_c \approx 10$ in good agreement with what is observed in Fig. 4.1a. Switching to the data type `Float128`⁷¹ (quadruple precision) with $\beta_c \approx 20$ in Fig. 4.1b, the onset of roundoff errors is shifted to lower temperatures in accordance with expectations.

How can we get a handle on the numerical instabilities in a floating point precision computation? As has been realized [176] soon after the introduction of the DQMC method in 1981 [135], an efficient strategy is to keep the broadly different scales in the matrix exponentials separated throughout the computation (as much as possible) and only mix them in a final step, if necessary. To extract the scale information of a time slice matrix B_l we employ matrix decompositions of the form

$$B_l = UDX. \quad (4.5)$$

Here, U and X are matrices of the order of unity and D is a real diagonal matrix hosting the exponentially spread singular values⁷² of B_l . Various matrix decompositions can be used to obtain the factorization in Eq (4.5). In the following we will consider the two most popular choices deployed in DQMC simulations [90, 122, 176, 184–187].

⁷¹The datatype `Float128` is provided by the Julia package `Quadmath.jl`.

⁷²We will refer to the elements of D as singular values irrespective of the matrix decomposition.

Singular value decomposition

The singular value decomposition (SVD) is given by

$$B = UDV^\dagger, \quad (4.6)$$

where U and V^\dagger are unitary and D is real and diagonal. We will use and compare three heavily optimized routines [182] provided by LAPACK⁷³ [188] to compute the SVD of the time slice matrix B :

- **gesdd**: Divide-and-conquer (D&C)
- **gesvd**: Bidiagonal QR iteration (conventional)
- **gesvj**: Jacobi algorithm⁷⁴

QR decomposition

The QR decomposition reads

$$B = QR = UDT, \quad (4.7)$$

where Q is unitary and R is upper-triangular. In the second step, we have renamed $U = Q$ and split R into a real diagonal part $D = \text{diag}(R)$ and an upper-triangular part T with $\text{diag}(T) = 1$. We will consider the pivoted QR [124], which is deployed in the public DQMC implementations ALF [185] and QUEST [184], in form of LAPACK's **geqp3** in our analysis.

Using Eq. (4.5), we can stabilize the matrix multiplication of two time slice matrices B_i and B_j as follows [90, 122, 161, 176],

$$\begin{aligned} B_i B_j &= \underbrace{U_i D_i X_i}_{B_i} \underbrace{U_j D_j X_j}_{B_j} \\ &= U_i \underbrace{(D_i ((X_i U_j) D_j))}_{U' D' X'} X_j \\ &= U_r D_r X_r. \end{aligned} \quad (4.8)$$

Here, $U_r = U_i U'$, $D_r = D'$, $X_r = X' X_j$, and $U' D' X'$ indicates an intermediate matrix decomposition. If we follow this scheme, in which parentheses indicate the order of operations, largely different scales present in the diagonal matrices won't be additively mixed throughout the computation [90, 122, 161]. Specifically, note that the scale information is preserved in the matrix multiplication involving

⁷³We will report on results obtained with the LAPACK implementation OpenBLAS that ships with Julia. Qualitatively similar results have been found in an independent test based on Intel's Math Kernel Library (MKL).

⁷⁴To access **gesvj** in Julia we use the external package JacobiSVD.jl.

the well-conditioned unit-scale matrix $U = X_i U_j$: the diagonal matrices merely rescale the columns and rows of U ,

$$\begin{aligned}
 D_i U D_j &= \begin{bmatrix} \mathbf{S} & & \\ & \mathbf{S} & \\ & & \mathbf{s} \end{bmatrix} \underbrace{\begin{bmatrix} \mathbf{s} & \mathbf{s} & \mathbf{s} & \mathbf{s} \\ \mathbf{s} & \mathbf{s} & \mathbf{s} & \mathbf{s} \\ \mathbf{s} & \mathbf{s} & \mathbf{s} & \mathbf{s} \\ \mathbf{s} & \mathbf{s} & \mathbf{s} & \mathbf{s} \end{bmatrix}}_{\text{unit scale}} \begin{bmatrix} \mathbf{S} & & \\ & \mathbf{S} & \\ & & \mathbf{s} \end{bmatrix} \\
 &= \begin{bmatrix} \mathbf{S} & & \\ & \mathbf{S} & \\ & & \mathbf{s} \end{bmatrix} \begin{bmatrix} \mathbf{sS} & \mathbf{sS} & \mathbf{s}^2 & \mathbf{s_s} \\ \mathbf{sS} & \mathbf{sS} & \mathbf{s}^2 & \mathbf{s_s} \\ \mathbf{sS} & \mathbf{sS} & \mathbf{s}^2 & \mathbf{s_s} \\ \mathbf{sS} & \mathbf{sS} & \mathbf{s}^2 & \mathbf{s_s} \end{bmatrix} \\
 &= \begin{bmatrix} \mathbf{S}^2 \mathbf{s} & \mathbf{SSs} & \mathbf{Ss}^2 & \mathbf{Ss_s} \\ \mathbf{SSs} & \mathbf{S}^2 \mathbf{s} & \mathbf{Ss}^2 & \mathbf{Ss_s} \\ \mathbf{Ss}^2 & \mathbf{SSs} & \mathbf{s}^3 & \mathbf{s}^2 \mathbf{s_s} \\ \mathbf{Ss_s} & \mathbf{SSs} & \mathbf{s}^2 \mathbf{s_s} & \mathbf{s_s}^2 \end{bmatrix}.
 \end{aligned} \tag{4.9}$$

As we will see below, repeating the procedure, Eq. (4.8), allows us to obtain a numerically accurate UDX decomposition of the full time slice matrix chain $B(\tau, 0)$, Eq. (4.3). We note that, in practice, it is often unnecessary to stabilize every individual matrix-matrix product. Instead, one typically performs a mixture of naive and stabilized products for the sake of speed while still retaining numerical accuracy [122, 149, 187].

Benchmark

We test the accuracy of the matrix product stabilization procedure, Eq. (4.8), with respect to varying the SVD and QR factorization algorithms. Fig. 4.2 shows the logarithmic singular values of the time slice matrix chain $B(\beta, 0)$ as a function of inverse temperature $\beta = 1/T$. Clearly, the accuracy of the computed singular values shows a strong dependence on the chosen factorization algorithm. While the curves for the QR decomposition and Jacobi SVD seem to fall on top of the exact result⁷⁵, we observe large deviations for the conventional and D&C SVD algorithms. This effect is particularly pronounced at low temperatures, $\beta \gtrsim 25$. The fact that small scales are lost in these SVD variants, while large ones are still correct, can be understood from LAPACK's SVD error bounds [189]: The error is bounded relative to the largest singular value. Thus, large scales are computed to high relative accuracy and small ones may not be.

Turning to computational efficiency, we illustrate runtime cost measurements for the considered SVD variants relative to the QR decomposition in Fig. 4.3. We find that both the conventional SVD and Jacobi SVD are an order of magnitude slower than the QR decomposition while only the divide-and-conquer algorithm shows comparable speed. Among the SVD variants, the Jacobi SVD is the most costly by a large margin, having about twice the runtime of the conventional SVD for small matrix sizes.

⁷⁴We estimate the precision as $p = \log_{10}(2^{\text{fraction}})$, where fraction is the mantissa of a given binary floating point format. This gives $p \sim 16$ for `Float64` and $p \sim 34$ for `Float128`.

⁷⁵The exact $B(\tau, 0)$ can be determined in an arbitrary precision computation.

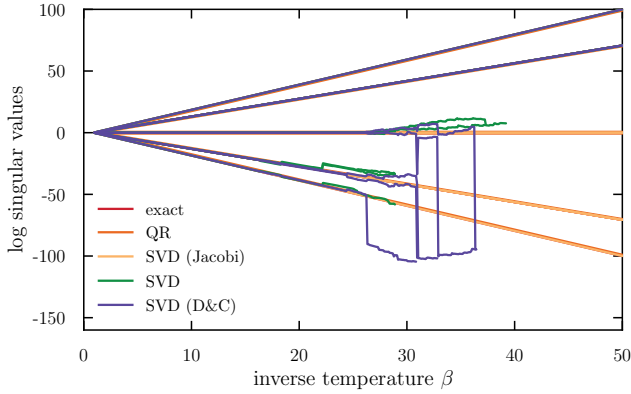


Figure 4.2: Comparison of matrix decompositions to heal the numerical instabilities in the calculation of the time slice matrix chain product $B_M B_{M-1} \cdots B_1$ for free fermions in one dimension, Eq. (4.2). Different lines represent logarithmic singular values as observed in stabilized computations. The QR (orange) and Jacobi SVD singular values (yellow) lie on top of the exact result (red) whereas both the regular SVD (green) and the divide-and-conquer SVD (purple) show large deviations at low temperatures $\beta \gtrsim 25$ ($\Delta\tau = 0.1$). Resized from Ref. [P2] under CC-BY 4.0.




Figure 4.3: Computational efficiency of the considered matrix decompositions. Shown is the runtime cost of the factorization of a complex matrix of size $N \times N$ by means of various SVD algorithms relative to the QR decomposition. Resized from Ref. [P2] under CC-BY 4.0.

4.1.2 Equal-time Green's function

Similar to the considerations above, a naive computation of the Green's function according to Eq. (2.21) is potentially unstable because of numerical roundoff errors due to finite machine precision. In particular, adding the identity to the ill-conditioned slice matrix chain $B(\tau, 0)$ will generally wash out small singular values and will lead to a non-invertible result such that the subsequent inversion in Eq. (2.21) is ill-defined. This clearly prohibits a safe calculation of the Green's function and asks for numerical stabilization techniques.

Inversion schemes

As for the time slice matrix products in Eq. (4.8), the strategy is to avoid an uncontrolled mixing of exponential scales during the computation. A straightforward scheme (`inv_one_plus`) to add the unit matrix and perform the inversion of $\mathbb{1} + B(\tau, 0)$ in a stabilized manner is given by [90, 122]

$$\begin{aligned}
 G &= [\mathbb{1} + UDX]^{-1} \\
 &= [U \underbrace{(U^\dagger X^{-1} + D)}_{udx} X]^{-1} \\
 &= [(Uu)d(xX)]^{-1} \\
 &= U_r D_r X_r,
 \end{aligned} \tag{4.10}$$

where $U_r = (xX)^{-1}$, $D_r = d^{-1}$, and $X_r = (Uu)^{-1}$. Here, the intermediate addition of unit scales and singular values (parentheses in the second line) is separated from the unitary rotations such that $U^\dagger X^{-1}$ only acts as a clean cutoff,

$$U^\dagger X^{-1} + D = \begin{bmatrix} s & s & s & s \\ s & s & s & s \\ s & s & s & s \\ s & s & s & s \end{bmatrix} + \begin{bmatrix} S & & & \\ & S & & \\ & & s & \\ & & & s \end{bmatrix} = \begin{bmatrix} S & s & s & s \\ s & S & s & s \\ s & s & S & s \\ s & s & s & S \end{bmatrix}. \tag{4.11}$$

As we will demonstrate for the time-displaced Green's function in Sec. 4.1.3, the procedure `inv_one_plus`, based on a single intermediate decomposition, will still fail to give accurate results for some of the matrix decompositions. For this reason, we consider another stabilization procedure (`inv_one_plus_loh`) put forward by Loh *et al.* [136, 176], in which one initially separates the scales of the diagonal matrix D into two factors $D_p = \max(D, 1)$ and $D_m = \min(D, 1)$,

$$D_p = \begin{bmatrix} S & & & \\ & S & & \\ & & s & \\ & & & s \end{bmatrix}, \quad D_m = \begin{bmatrix} s & & & \\ & s & & \\ & & s & \\ & & & s \end{bmatrix}, \tag{4.12}$$




Figure 4.4: Accuracy of the Green's function obtained from stabilized computations using the listed matrix decompositions and inversion schemes. Shown is $\Delta G = \log(\max(\text{abs}(G - G_{\text{exact}})))$ for $U = 0$ (solid) and $U = 1$ (alpha transparent) for the Hubbard model, Eq. (4.1). Resized from Ref. [P2] under CC-BY 4.0.

and performs two intermediate decompositions,

$$\begin{aligned}
 G &= [\mathbb{1} + UDX]^{-1} \\
 &= [\mathbb{1} + UD_m D_p X]^{-1} \\
 &= [(X^{-1} D_p^{-1} + UD_m) D_p X]^{-1} \\
 &= X^{-1} \underbrace{[D_p^{-1} (X^{-1} D_p^{-1} + UD_m)^{-1}]}_{\substack{udx \\ udx}} \\
 &= U_r D_r X_r,
 \end{aligned} \tag{4.13}$$

where $U_r = X^{-1}u$, $D_r = d$, and $X_r = x$.

Benchmark

Starting from a stabilized computation of $B(\beta, 0)$, Sec. 4.1.1, we calculate the equal-time Green's function by performing the inversion according to the schemes outlined above and by varying the applied matrix factorization. In Fig. 4.4a we show our findings for `inv_one_plus`, where we have taken the maximum absolute difference between the computed and the exact Green's function as an accuracy measure. At high temperatures and for $U = 0$, we observe that all decompositions lead to a good approximation of G_{exact} with an accuracy close to floating point precision. However, when turning to lower temperatures the situations changes dramatically. We find that only the QR decomposition and the Jacobi SVD produce the Green's function reliably. Compared to the other SVD variants, which fall behind by a large margin and fail to reproduce the exact result, they consistently show about optimal accuracy even in the presence of interactions. As displayed in Fig. 4.4b, switching to the inversion scheme




Figure 4.5: Efficiency of the stabilized Green’s function calculation using the listed matrix decompositions and inversion schemes. Shown are results for the non-interacting model, Eq. (4.2). Resized from Ref. [P2] under CC-BY 4.0.

`inv_one_plus_loh` generally improves the accuracy but deviations of the regular SVD and D&C SVD remain of the order of unity at the lowest temperatures.

Independent of the employed inversion scheme, matrix decompositions are expected to be the performance bottleneck in the Green’s function computation. We hence expect the speed differences apparent in Fig. 4.3 to dominate computation time benchmarks of the full Green’s function calculation as well. This anticipation is qualitatively confirmed in Fig. 4.5, which shows the runtime cost of the Green’s function computation for both inversion schemes and all matrix decompositions relative to the QR. While the divide-and-conquer SVD is in the same ballpark as the QR decomposition, the other SVD algorithms fall behind by a large margin (an order of magnitude) for both inversion procedures. Importantly, this apparent runtime difference is increasing with system size. The observation that the relative slowdown factor is larger for the inversion scheme `inv_one_plus_loh` can be understood from the fact that it requires one additional intermediate matrix decomposition.

In summary, these findings suggest that only the QR decomposition and the Jacobi SVD, irrespective of the inversion procedure, are suited for computing the equal-time Green’s function in DQMC reliably. Since the QR decomposition has a significantly lower runtime cost, it is to be preferred over the Jacobi SVD.

4.1.3 Time-displaced Green’s function

We now turn to the stabilization of the time-displaced Green’s function, which is required for computing time-displaced pairing and density-like correlations, as

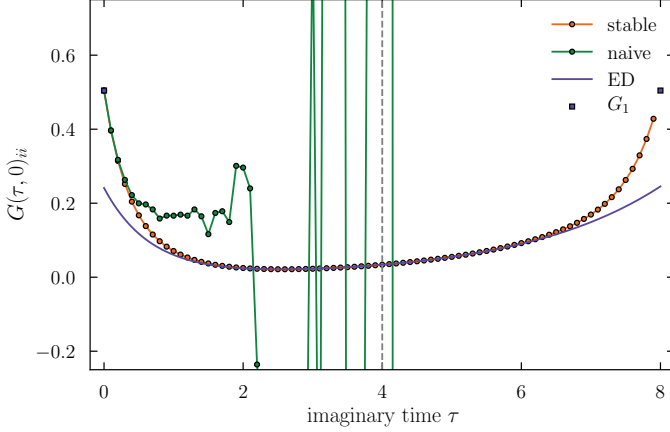


Figure 4.6: Onsite time-displaced Green's function $G(\tau, 0)_{ii}$ of the spin-fermion model in the form of Eq. (1.14), as obtained by a naive implementation of Eq. (4.14) (green), a stabilized computation (orange) using the QR decomposition and the inversion scheme `inv_sum_loh`, Eq. (4.17), as well as exact diagonalization (purple). The model parameters are $L = 2$, $\lambda = 1$, and $\beta = 8$. For convenience, we also show G_1 , Eq. (2.29), and highlight $\tau = \beta/2$ (dashed line). The ED curve deviates from the stabilized computation around $\tau \approx 0$, and $\tau \approx \beta$, because only the 20 smallest energy states have been considered in the diagonalization.

discussed in Ch. 2. In principle, Eq. (2.34), which we repeat here for convenience,

$$\begin{aligned}
 G(\tau_1, \tau_2) &= \begin{cases} B(\tau_1, \tau_2)G(\tau_2), & \tau_1 > \tau_2, \\ -(1 - G(\tau_1))B(\tau_2, \tau_1)^{-1}, & \tau_2 > \tau_1 \end{cases} \\
 &= \begin{cases} B_{l_1} \cdots B_{l_2+1}G_{l_2+1}, & \tau_1 > \tau_2, \\ -(1 - G_{l_1+1})(B_{l_2} \cdots B_{l_1+1})^{-1}, & \tau_2 > \tau_1, \end{cases} \quad (4.14)
 \end{aligned}$$

is a prescription for how to calculate $G(\tau_1, \tau_2)$ from the stabilized equal-time Green's function, Sec. 4.1.2. However, when $|\tau_1 - \tau_2|$ is large, a naive calculation of the matrix products in Eq. 4.14 will be numerically unstable. Furthermore, by first computing the equal-time Green's function, one inevitably mixes (and therefore loses) scale information in the last recombination step, $G = UDX$. We demonstrate the instability of a naive implementation of Eq. (4.14) for the spin-fermion model in Fig. 4.6.

For the reasons stated above, it is advantageous to rewrite Eq. (4.14) and com-

pute the time-displaced Green's function as (we focus on $\tau_1 > \tau_2$ for simplicity)

$$\begin{aligned}
 G(\tau_1, \tau_2) &= B_{l_1} \cdots B_{l_2+1} G_{l_2+1} \\
 &= B_{l_1} \cdots B_{l_2+1} [\mathbb{1} + B_{l_2} \cdots B_1 B_M \cdots B_{l_2+1}]^{-1} \\
 &= \left[\underbrace{B_{l_2+1}^{-1} \cdots B_{l_1}^{-1}}_{U_L D_L X_L} + \underbrace{B_{l_2} \cdots B_1 B_M \cdots B_{l_1+1}}_{U_R D_R X_R} \right]^{-1} \\
 &= [U_L D_L X_L + U_R D_R X_R]^{-1}.
 \end{aligned} \tag{4.15}$$

Inversion schemes

Similar to the equal-time Green's function in Sec. 4.1.2, one must be careful to keep the scales in D_L and D_R separated when performing the summation and inversion to avoid unnecessary floating point roundoff errors. As a first explicit procedure, we consider a simple generalization of Eq. 4.10 (`inv_sum`),

$$\begin{aligned}
 G(\tau_1, \tau_2) &= [U_L D_L X_L + U_R D_R X_R]^{-1} \\
 &= [U_L (\underbrace{D_L X_L X_R^{-1} + U_L^\dagger U_R D_R}_{udx}) X_R]^{-1} \\
 &= [(U_L u) d^{-1} (x X_R)]^{-1} \\
 &= U_r D_r X_r,
 \end{aligned} \tag{4.16}$$

where $U_r = (x X_R)^{-1}$, $D_r = d^{-1}$, and $X_r = (U_L u)^{-1}$.

Analogously, we can generalize the scheme by Loh *et al.* [136], Eq. 4.13, in which we split the scales into matrix factors $D_m = \min(D, 1)$, $D_p = \max(D, 1)$, (`inv_sum_loh`)

$$\begin{aligned}
 G(\tau_1, \tau_2) &= [U_L D_L X_L + U_R D_R X_R]^{-1} \\
 &= [U_L D_{Lm} D_{Lp} X_L + U_R D_{Rm} D_{Rp} X_R]^{-1} \\
 &= \left[U_L D_{Lp} \left(\underbrace{\frac{D_{Lm}}{D_{Rp}} X_L X_R^{-1} + U_L^\dagger U_R \frac{D_{Rm}}{D_{Lp}}}_{udx} \right) X_R D_{Rp} \right]^{-1} \\
 &= X_R^{-1} \underbrace{\frac{1}{D_{Rp}} [udx]^{-1} \frac{1}{D_{Lp}} U_L^\dagger}_{udx} \\
 &= U_r D_r X_r,
 \end{aligned} \tag{4.17}$$

where $U_r = X_R^{-1} u$, $D_r = d$, and $X_r = x U_L^\dagger$.

Benchmark

In Fig. 4.7, we show the logarithmic, maximal, absolute deviation of the time-displaced Green's function from the arbitrary precision result as a function of time-displacement τ at inverse temperature $\beta = 40$. Focusing on the inversion

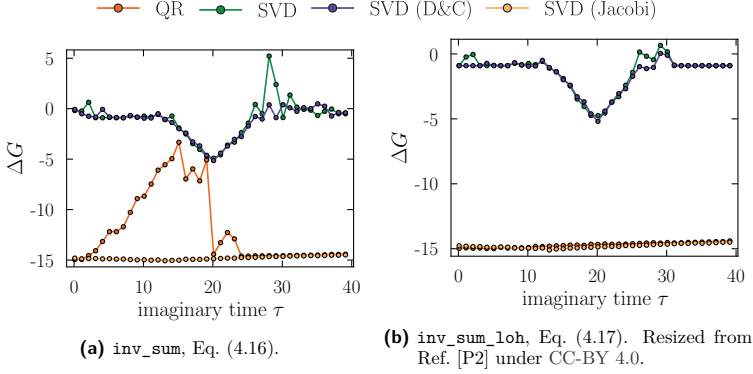


Figure 4.7: Accuracy of the time-displaced Green's function obtained from stabilized computations using the listed matrix decompositions and inversion schemes. Shown is $\Delta G = \log(\max(\text{abs}(G(\tau, 0) - G_{\text{exact}}(\tau, 0))))$ for $\beta = 40$.

scheme `inv_sum` first, Fig. 4.7a, both regular and D&C SVD clearly fail to capture the intrinsic scales sufficiently and errors much beyond floating point precision are visible. Although the QR decomposition systematically leads to equally or more accurate results for all considered imaginary times, it fails to be reliable at long times $\tau \sim \beta/2$ (the Green's function is anti-periodic in τ [17, 23]). Only the Jacobi-method based SVD leads to an accurate Green's function at all imaginary times. Switching to the inversion scheme `inv_sum_loh`, the situation changes, as illustrated in Fig. 4.7b. While the non-Jacobi SVDs still have insufficient accuracy, the result for the QR decomposition improves dramatically compared to `inv_sum` and leads to stable Green's function estimates up to floating point precision along the entire imaginary time axis.

Finally, we compare the computational runtime cost associated with both stable approaches: the Jacobi SVD combined with the regular inversion and the QR decomposition paired with `inv_sum_loh`. As shown in Fig. 4.8, we find that the latter is consistently faster for all considered system sizes. In relative terms, the SVD based approach falls behind by at least a factor of two and seems to display inferior scaling with system size.

Similar to the equal-time Green's function discussion (Sec. 4.1.2), our findings suggest that only the Jacobi SVD and the QR decomposition are reliable for computing time-displaced Green's functions in DQMC. In spite of the need to pair the QR decomposition with an inversion scheme of higher complexity, an observation that, to the best of our knowledge, has not been mentioned in the literature before, the QR-based approach appears to be preferable due to the superior computational efficiency.

We end this section by noting that Hirsch and Fye [90, 190, 191] – and similarly Assaad *et al.* [122] – have proposed an alternative method for stably computing the equal-time and time-displaced Green's functions based on a space-time matrix formulation of the problem. Although this technique has been

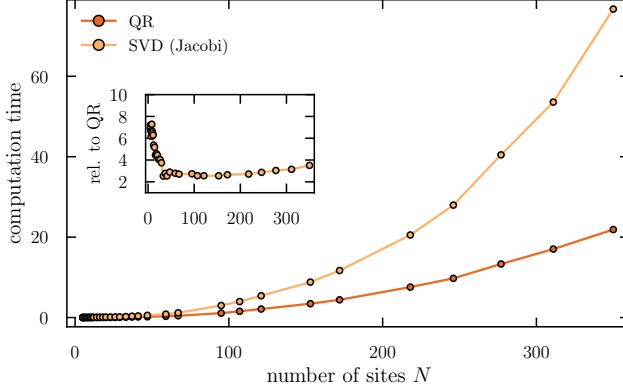


Figure 4.8: Efficiency of the time-displaced Green’s function obtained from stabilized computations using the QR decomposition in combination with the inversion scheme `inv_sum_loh`, Eq. (4.17) and the Jacobi SVD paired up with the regular inversion scheme `inv_sum`, Eq. (4.16). Measurements are taken over multiple runs at $\tau = \beta/2 = 20$. The inset shows the slowdown of the Jacobi SVD relative to the QR based approach. Resized from Ref. [P2] under CC-BY 4.0.

successfully deployed in many-fermion simulations of Hubbard models [94, 148, 159, 192], we exclude it from our discussion because of its subpar memory scaling: for a system composed of N lattice sites, f fermion flavors, and imaginary time extent M one needs to invert (naively a $O(n^3)$ operation [124]) a matrix which takes up $O((NMf)^2)$ memory. For the spin-fermion model with $f = 4$, Eq. (3.4), and system sizes of about $N \approx 100$ lattice sites, this approach turned out to be impractical for the large-scale simulations in this thesis.


4.2 Efficient DQMC

The determinant quantum Monte Carlo method, as presented in Ch. 2, requires the computation of the equal-time Green’s function and its determinant in every (proposed) step of the Markov walk, Eq. (2.23). All matrix operations involved in the process – multiplication, inversion, and determinant calculation – typically scale as $O(n^3)$ [124], where n is the linear matrix dimension, and a sweep of local updates, Fig. 2.2, across the space-time lattice of extent βN therefore has a computational complexity of $O(\beta N \cdot (fN)^3) = O(\beta f^3 N^4) = O(\beta N^4)$ ⁷⁶. In particular, the fourth-power scaling of the asymptotic runtime cost as a function of N severely limits the accessible system sizes - for the square lattice spin-fermion model, Eq. (3.4), $L = \sqrt{N} \approx 6$ would have been the practical upper bound for our simulations in this thesis on the available hardware, Sec. 4.4.

⁷⁶The Green’s function is generally a $Nf \times Nf$ matrix, where N is the number of spatial lattice sites and f is the number of fermion flavors (including spin).

Figure 4.9:

Illustration of a local update $\phi_{li} \rightarrow \phi'_{li}$ of a N_ϕ -component order parameter at lattice site i and imaginary time slice l . A move vector $\Delta\phi$, indicating a rotation and scaling of ϕ_{li} , is uniformly drawn at random from a N_ϕ -dimensional box of side length Δ_{box} . During the equilibration phase, the scale Δ_{box} is adapted to approximately realize a 50% acceptance rate. Inspired by Ref. [149].



In the following, we describe a series of strategies, applied in the simulations of Chs. 5 and 7, to improve the efficiency of Monte Carlo updates and linear algebra operations in DQMC. For simplification, we focus our discussion on the spin-fermion model but note that the presented concepts are largely model independent [120, 122, 148].

4.2.1 Efficient adaptive local updates

As shown in Fig. 4.9, we implement a local update at a space-time lattice site (l, i) by a small rotation and scaling of the bosonic field ϕ_{li} . Specifically, we draw a random vector $\Delta\phi$ from a box with side length Δ_{box} and propose the new field value

$$\phi'_{li} = \phi_{li} + \Delta\phi. \quad (4.18)$$

Generally, a small value of $\Delta\phi$ leads to a small change of the associated action and to Metropolis acceptance with high probability. However, in this case the autocorrelation time τ is large, since subsequent configurations in the Markov chain are highly correlated. In contrast, if $\Delta\phi$ is large, one expects a higher Monte Carlo rejection rate and a decrease of τ . To balance between these extremes and ensure ergodicity, Sec. 2.3.1, we dynamically adapt the scale Δ_{box} during the equilibration phase of our DQMC simulations, Sec. 2.6.4, such that a stable acceptance rate of about 50% is observed [122, 149].

To decide between accepting and rejecting individual proposals, Eq. (4.18), we must evaluate the DQMC Metropolis acceptance probability, Eq. (2.23). Focusing on the bosonic contribution first, we note that the change of the bosonic action ΔS_ϕ can be efficiently computed for the spin-fermion model as [149],

$$\Delta S_\phi = S'_\phi - S_\phi = \Delta S_{\phi 1} + \Delta S_{\phi 2} + \Delta S_{\phi 3}, \quad (4.19)$$

$$\Delta S_{\phi 1} = \frac{1}{c^2 \Delta \tau} (-(\phi_{l+1i} + \phi_{l-1i}) \cdot \Delta\phi_{li}), \quad (4.20)$$

$$\Delta S_{\phi 2} = \frac{\Delta \tau}{2} \left(4(|\phi'_{li}|^2 - |\phi_{li}|^2) - 2\Delta\phi_{li} \cdot \sum_{j \in N(i)} \phi_{lj} \right), \quad (4.21)$$

$$\Delta S_{\phi 3} = \Delta \tau \left(\frac{r}{2} \left(|\phi'_{li}|^2 - |\phi_{li}|^2 \right) + \frac{u}{4} \left(|\phi'_{li}|^4 - |\phi_{li}|^4 \right) \right), \quad (4.22)$$

in which we have replaced the spatial and imaginary time gradients in Eq. (3.4) by numerical lattice derivatives [32, 124].

Turning to the fermionic determinant ratio, it is advantageous [90, 120, 122, 161] to replace the Green's function G_ϕ in Eq. (2.23) by its time-dependent partner G_l , Eq. (2.29). This is justified since, due to the property $\det(1 + AB) = \det(1 + BA)$, the acceptance probability is invariant under this exchange,

$$\begin{aligned} A_{\phi' \phi} &= \min \left\{ 1, e^{-\Delta S_\phi} \frac{\det G_\phi}{\det G'_\phi} \right\} \\ &= \min \left\{ 1, e^{-\Delta S_\phi} \frac{\det(\mathbb{1} + B_M \dots B_{l+1} B'_l B_{l-1} \dots B_l)}{\det(\mathbb{1} + B_M \dots B_{l+1} B_l B_{l-1} \dots B_l)} \right\} \\ &= \min \left\{ 1, e^{-\Delta S_\phi} \frac{\det(\mathbb{1} + B_{l-1} \dots B_l B_M \dots B_{l+1} B'_l)}{\det(\mathbb{1} + B_{l-1} \dots B_l B_M \dots B_{l+1} B_l)} \right\} \\ &= \min \left\{ 1, e^{-\Delta S_\phi} \frac{\det G_l}{\det G'_l} \right\}. \end{aligned} \quad (4.23)$$

As we will see below, this form will allow us to exploit the space-time locality of the proposed update and reduce the cost associated with the determinant ratio computation, naively an expensive $O((fN)^3)$ operation, to $O(f \approx 1)$.

First, we define $R = \det G_l / \det G'_l$ and rewrite

$$\begin{aligned} R &= \det(G_l [\mathbb{1} + B_{l-2} \dots B_l B_M \dots B'_l]) \\ &= \det(G_l [\mathbb{1} + (G_l^{-1} - \mathbb{1}) B_l^{-1} B'_l]) \\ &= \det(G_l + (\mathbb{1} - G_l) B_l^{-1} B'_l) \\ &= \det \left(\underbrace{\mathbb{1} + (\mathbb{1} - G_l)(B_l^{-1} B'_l - \mathbb{1})}_{\equiv K_l} \right). \end{aligned} \quad (4.24)$$

The structure of the matrix K_l , representing the Green's function ratio for a local update, is visualized in Fig. 4.10. Introducing [32, 90, 122]

$$\Delta_l = B_l^{-1} B'_l - \mathbb{1} = e^{\Delta \tau V_l} e^{-\Delta \tau V'_l} + O(\Delta \tau^2) - \mathbb{1}, \quad (4.25)$$

in which the numerical error is of Trotter type, Eq. (2.17), we realize that under a local update $\phi_{li} \rightarrow \phi'_{li}$ only a single element (i, i) in every $N \times N$ flavor sector of the interaction matrix exponential $e^{\Delta \tau V_l}$ is being altered. As a consequence, only f^2 of the $(fN)^2$ entries of the matrix Δ_l are finite⁷⁷, Fig. 4.10. While preserving the order, we extract these non-vanishing elements into a reduced $f \times f$ matrix $\tilde{\Delta}_{li}$ which, using slice indexing notation⁷⁸, reads

$$\begin{aligned} \tilde{\Delta}_{li} &= \Delta_l[i:N:\text{end}] \\ &= e^{\Delta \tau V_l}[i:N:\text{end}] \cdot e^{-\Delta \tau V'_l}[i:N:\text{end}] - \mathbb{1}_{f \times f} + O(\Delta \tau^2). \end{aligned} \quad (4.26)$$

⁷⁷For the spin-fermion model with O(3) symmetry, 4 flavor sectors vanish exactly such that only $16 - 4 = 12$ elements are finite, Eq. (3.9).

⁷⁸We choose the Julia syntax $M[\text{start} : \text{step} : \text{stop}]$.




Figure 4.10: Illustration of the structure of the matrix K_l appearing in the determinant in Eq. (4.24). The local nature of an update at lattice site i and time slice l leads to a sparsity pattern (colored dots) which can be exploited by expanding the determinant of K_l by minors, $\det K_l = \det \tilde{K}_{li}$.

Crucially, as illustrated in Fig. 4.10, the matrix K_l , appearing in the determinant, inherits the sparsity of Δ_l . Analogous to $\tilde{\Delta}_{li}$ above, we define⁷⁹

$$\tilde{K}_{li} = \mathbb{1}_{f \times f} + \tilde{\Delta}_{li} (\mathbb{1}_{f \times f} - G_l[i:N:\text{end}]) . \quad (4.27)$$

The determinant in Eq. (4.24) may now be Laplace expanded by minors to give [32, 90, 120, 122, 161]

$$R = \det K_l = \det \tilde{K}_{li} . \quad (4.28)$$

Note that compared to the original K_l , the matrix \tilde{K}_{li} only is of size $f \times f$ and, importantly, independent of system size. The determinant can therefore be computed efficiently. If the number of flavors f is small enough, it may even be implemented explicitly, for example using Sarrus' scheme [193]. In this case, the cost associated with the Green's function determinant ratio computation according to Eq. (4.28) is strictly $O(1)$ – otherwise it scales as $O(f^3)$.

In-slice Green's function update

Although the determinant computation can be performed “on-site”, Eq. (4.28), the effect of a local update of the bosonic field according to Eq. (4.18) on the fermions is highly non-local and reflects the metallic nature of the system. Fortunately, the sparsity discussed above can nonetheless be utilized to derive an efficient update rule for the equal-time Green's function. We start by expressing G'_l through G_l ,

$$\begin{aligned} G'_l &= [\mathbb{1} + B_{l-1} \dots B_1 B_M \dots B_{l+1} B'_l]^{-1} \\ &= [\mathbb{1} + (G_l^{-1} - \mathbb{1}) B_l^{-1} B'_l]^{-1} \\ &= [\mathbb{1} + (\mathbb{1} - G_l) \underbrace{(B_l^{-1} B'_l - \mathbb{1})}_{\Delta_l}]^{-1} G_l . \end{aligned} \quad (4.29)$$

⁷⁹At this point, it might seem more natural to define $\tilde{K}_{li} = \mathbb{1}_{f \times f} + (\mathbb{1}_{f \times f} - G_l[i:N:\text{end}]) \tilde{\Delta}_{li}$. However, in Eq. (4.33) the chosen form will turn out to be more convenient. Note that both variants give the same determinant since $\det(\mathbb{1} + AB) = \det(\mathbb{1} + BA)$.

Since only f rows (and columns) of Δ_l are non-zero, Fig. 4.10, the second summand in this expression only has rank f (rather than fN) and allows for the application of the Sherman-Morrison-Woodbury formula [122, 124],

$$\left[\mathbb{1} + uv^\dagger\right]^{-1} = \mathbb{1} - u \left(1 + v^\dagger u\right)^{-1} v^\dagger \quad (4.30)$$

in which, in our context, u, v are matrices of size $N \times f$. Identifying

$$v^\dagger = \Delta_l[i:N:\text{end},:], \quad u = (\mathbb{1} - G_l)[:, i:N:\text{end}], \quad (4.31)$$

one obtains

$$\begin{aligned} G'_l &= \left[\mathbb{1} + uv^\dagger\right]^{-1} G_l \\ &= G_l - u \left[\mathbb{1} + v^\dagger u\right]^{-1} v^\dagger G_l \\ &= G_l - \left\{(\mathbb{1} - G_l)[:, i:N:\text{end}] \cdot \tilde{K}_{li}^{-1}\right\} \cdot \left\{\Delta_l[i:N:\text{end},:] \cdot G_l\right\}. \end{aligned} \quad (4.32)$$

Understanding that $\Delta_l[i:N:\text{end},:]G_l = \tilde{\Delta}_l G_l[i:N:\text{end},:]$, we can formulate the Green's function update rule [32, 90, 122]

$$G'_l = G_l + \left\{(G_l - \mathbb{1})[:, i:N:\text{end}] \cdot \tilde{K}_{li}^{-1}\right\} \cdot \left\{\tilde{\Delta}_l \cdot G_l[i:N:\text{end},:]\right\}, \quad (4.33)$$

where parentheses indicate the order of operations. Rather than having to invert a $fN \times fN$ matrix, Eq. (4.29), this update scheme only requires the inversion of the reduced $f \times f$ matrix \tilde{K}_{li} , which has already appeared in the determinant ratio calculation, Eq. (4.28). Furthermore, since the two factors (curly braces) are sparse, a computation of G'_l according Eq. (4.33) takes only $O(f^3 N^2)$ operations – ideally $O(1)$ for the inversion and $O(f^3 N^2)$ for each of the sparse matrix multiplications⁸⁰. The computation cost of a DQMC sweep of local updates across the entire space-time lattice of extent βN is therefore of the order of $O(\beta N \cdot f^3 N^2) \approx O(\beta N^3)$. Compared to a naive DQMC implementation, the optimizations therefore reduce the complexity by one power of system size. In Fig. 4.11, we demonstrate this $O(\beta N^3)$ scaling for the DQMC implementation used in this thesis [S1].

We note that in the presence of symmetries, the effective size of the Green's function matrix used in the DQMC formalism may be smaller than the physical $fN \times fN$. As mentioned around Eq. (3.13) in Ch. 3, for the O(2) and O(1) symmetric spin-fermion models, the number of fermion flavors, $f = 4$ is effectively reduced to $f^{\text{eff}} = 2$. From the scaling above, we hence theoretically expect an eightfold speedup compared to the O(3) symmetry case. Empirically, we observe a smaller but still significant performance difference between the O(2) and O(3) variants⁸¹, Fig. 4.12. We find that a DQMC sweep of local updates for a system of linear size $L = 16$ takes about six times longer when switching from two to three order parameter components.

⁸⁰We estimate the complexity of a sparse matrix multiplication of a $fN \times f$ matrix A with fN non-vanishing entries per column and a $f \times fN$ matrix B with f non-zero elements per column as $O(fN \cdot f \cdot fN) = O(f^3 N^2)$.

⁸¹We leave out the O(1) model in the comparison here for two reasons: First, the matrices are real in this case and therefore of a different data type. Second, it would be superior to implement a Z_2 order parameter.




Figure 4.11: Scaling of the computation time of a sweep of local updates with inverse temperature β and cubic system size N^3 . Markers indicate measurements of our DQMC implementation [S1] of the spin-fermion model, Eq. (3.4), for $L = 8, 10, 12$ and $\beta = 2, 3, 4, 5, 10, 20, 30, 40$. Simulations have been performed with a mean acceptance rate of 50%. The theoretical $O(\beta N^3)$ scaling (dashed line) is shown for comparison.




Figure 4.12: Comparison of the computation time of a sweep of local updates for the $O(3)$ and $O(2)$ symmetric spin-fermion models, Eq. (3.4). Shown is the relative factor $t_{O(3)}/t_{O(2)}$ such that unity corresponds to the case of equal sweep durations. An acceptance ratio of 50% has been strictly enforced to avoid statistical fluctuations. Time measurements have been obtained by averaging over 100 up-down sweeps at inverse temperature $\beta = 5$.




Figure 4.13: Discretization of imaginary time in our DQMC simulations. The interval $[0, \beta]$ is split into M blocks of size $\Delta\tau$. The index l enumerates all imaginary time slices (solid lines). We iterate downwards and upwards in imaginary time and perform u-turns at the boundaries. At every time slice, we compute the equal-time Green's function by propagation, Eqs. (4.34) and (4.35), or from scratch (red solid lines) and propose local updates along the spatial dimension (not shown).

Green's function propagation and caching

We implement the DQMC sweep of local update proposals on the space-time lattice by successively sweeping downwards and upwards in imaginary time, Fig. 4.13, and iterating over all spatial sites i within each time slice l . In the last section, we have shown that such a sweep can be efficiently performed in $O(\beta N^3)$. We thereby implicitly assumed the availability of the equal-time Green's function G_l in every imaginary time slice. Repeatedly computing the latter from scratch according to Eq. (2.29) is, however, a costly operation of the order of $O(\beta N^3)$ by itself. Below, we describe a systematic procedure which, by exploiting the locality of the update, improves the efficiency of DQMC and significantly reduces this cost in a practical simulation.

We start by highlighting an important relation between equal-time Green's functions at neighboring imaginary time slices. By symmetrically factoring out time slice propagations in the definition of G_{l-1} , Eq. (2.29), one sees that [90, 120, 122]

$$\begin{aligned}
 G_{l-1} &= [\mathbb{1} + B_{l-2} \dots B_1 B_M \dots B_{l-1}]^{-1} \\
 &= \left[B_{l-1}^{-1} \left(\underbrace{B_{l-1} B_{l-1}^{-1}}_{\mathbb{1}} B_{l-1} \dots B_1 B_M \dots B_l \right) B_{l-1} \right]^{-1} \\
 &= B_{l-1}^{-1} G_l B_{l-1},
 \end{aligned} \tag{4.34}$$




Figure 4.14: Visualization of the Green's function computation in a DQMC simulation. A stack of *UDT* decompositions, Eq. (4.7), of time slice matrix propagators $B(\tau, 0)$ (dark blue) and $B(\beta, \tau)^\dagger$ (light grey) is filled in down- and up-sweeps (vertical axis). Stack slots correspond to the stabilization time slices $\lambda \in \Lambda$, see Fig. 4.13, at which the Green's function is recomputed according to Eq. (4.40). If the entire stack is single colored, the matrix product chain in the latter is maximally unbalanced, Eq. (4.41).

and analogously,

$$G_{l+1} = B_l G_l B_l^{-1}. \quad (4.35)$$

Hence, the Green's functions can be propagated in imaginary time by multiplication with time slice matrices B_l .

In principle, this suggests the following sweeping strategy: compute the initial Green's function from the ground up, Eq. (2.29), update it, Eq. (4.33), when accepting local moves, and propagate it, Eqs. (4.34) and (4.35), when switching between time slices. However, as demonstrated in Sec. 4.1, successive multiplications of time slice matrices lead to severe numerical round off errors such that the propagation of the Green's function becomes unstable after a few steps [122]. Assuming that Δ_s factors of B_l can be multiplied without significant loss of accuracy, it is therefore necessary to recompute a fresh Green's function at every time slice $\lambda \in \Lambda = \{l \mid \text{mod}(l - 1, \Delta_s) = 0\}$, Fig. 4.13. Empirically, we find that values of the order of $\Delta_s \approx 10$ represent a reasonable compromise between floating-point precision and simulation speed for the spin-fermion model, Eq. (3.4), and the Hubbard model, Eq. (1.17), considered in this thesis.

To improve the efficiency of stabilized Green's function calculations, it is useful to cache intermediate results of previous from scratch computations. Realizing that a local update in a time slice l only changes the corresponding factor B_l in Eq. (2.29), we may temporarily store partial time slice matrix product chains corresponding to imaginary time segments that have already been visited in

the DQMC sweep [122, 149, 194]. Concretely, we rewrite the Green's function definition in continuum notation, Eq. (2.30), as

$$G(\tau) = [\mathbb{1} + B(\tau, 0)(B(\beta, \tau)^\dagger)^\dagger]^{-1} \quad (4.36)$$

and cache the *UDT* decompositions, Eq. (4.7), of the contributions $B(\tau, 0)$ and $B(\beta, \tau)^\dagger$. Depending on whether we are sweeping downwards or upwards in imaginary time, Fig. 4.13, at every stabilization time $\tau_\lambda = \lambda \Delta_s \Delta \tau$ with $\lambda \in \Lambda$ we overwrite slots in a stack S of length $M_\lambda = M + 1$ with *UDT* representations of updated time slice propagators $B(\beta, \tau_\lambda)^\dagger$ or $B(\tau_\lambda, 0)$. A visualization of this procedure is provided in Fig. 4.14. In each step, we use partial results, that have been previously stored in the stack, by making use of the transitivity relations

$$\underbrace{B(\beta, \tau_{\lambda-1})^\dagger}_{S[\lambda]} = B(\tau_\lambda, \tau_{\lambda-1})^\dagger \underbrace{B(\beta, \tau_\lambda)^\dagger}_{S[\lambda+1]}, \quad (4.37)$$

$$\underbrace{B(\tau_\lambda, 0)}_{S[\lambda+1]} = B(\tau_\lambda, \tau_{\lambda-1}) \underbrace{B(\tau_{\lambda-1}, 0)}_{S[\lambda]}, \quad (4.38)$$

for downwards and upwards sweeps, respectively. Here, it is left implicit that multiplications are performed on *UDT* representations, for example

$$B(\tau_1, \tau_2) \cdot UDT = \underbrace{((B(\tau_1, \tau_2) \cdot U) \cdot D)}_{udt} T = U' D' T' \quad (4.39)$$

with $U' = u$, $D' = d$, and $T' = tT$. Given the *UDT* decompositions of $B(\tau_\lambda, 0)$ and $B(\beta, \tau_\lambda)^\dagger$ from the stack, we may compute the equal-time Green's function at a stabilization time slice $\lambda \in \Lambda$ according to

$$\begin{aligned} G(\tau_\lambda) &= [\mathbb{1} + B(\tau_\lambda, 0)(B(\beta, \tau_\lambda)^\dagger)^\dagger]^{-1} \\ &= [\mathbb{1} + U_L D_L T_L (U_R D_R T_R)^\dagger]^{-1} \\ &= \left[\mathbb{1} + U_L \underbrace{(D_L (T_L T_R^\dagger) D_R)}_{udt} U_R^\dagger \right]^{-1} \\ &= [\mathbb{1} + UDT]^{-1}, \end{aligned} \quad (4.40)$$

with $U = U_L u$, $D = d$, and $T = tU_R^\dagger$, which may then be substituted into Eq. (4.10).

In Algs. 1 and 2, we summarize the described Green's function propagation of downwards and upwards DQMC sweeps in schematic pseudocode. We note that our discretization of imaginary time leads to an asymmetric distribution of stabilization time slices over the interval $[0, \beta]$ as illustrated in Fig. 4.13. We formally include the time slices $l = 0$ and $l = M + 1$ in our DQMC procedure to account for this asymmetry and to obtain reliable Green's functions at the time boundaries of the simulation grid. There, the contributions to Eq. (4.36), and equivalently Eq. (4.40), are maximally unbalanced in the sense that

$$G = [\mathbb{1} + B(\beta, 0)]^{-1} = \left[\mathbb{1} + (B(\beta, 0)^\dagger)^\dagger \right]^{-1}, \quad (4.41)$$

and one of the propagators is the identity matrix. We hence special case these situations in the provided pseudo-implementation.

```
1 function greens_downwards()
2   l' = l - 1                                # time slice we're heading to
3   # stabilization time slice?
4   if l ∈ Λ
5     λ = (l' - 1) / Δs + 1                    # stack index
6     # upper boundary: l = M + 1, l' = M, λ = Mλ
7     if λ == Mλ
8       B_L = S[λ]                            # read B(β, 0)
9       B_R = 1                               # set BR = B(β, β)† = 1
10      S[Mλ] = B_L                           # overwrite last stack slot
11      compute_G(B_L, B_R)                    # Eq. (5.56)
12      l = l'                                 # update time slice index
13
14      # intermediate time slice
15      elseif Mλ > λ > 1
16        B_L = S[λ]                            # read B((λ - 1)ΔsΔτ, 0)
17        # S[λ] = B(β, (λ - 1)ΔsΔτ)† and S[λ + 1] = B(β, λΔsΔτ)†
18        S[λ] = compute_B_dagger(λΔsΔτ, (λ - 1)ΔsΔτ) * S[λ + 1]
19        B_R = S[λ]                            # BR = B(β, (λ - 1)ΔsΔτ)†
20        compute_G(B_L, B_R)                  # Eq. (5.56)
21        l = l'                                 # update time slice index
22
23        # lower boundary: l = 1, l' = 0, λ = 1
24      else
25        # S[1] = B(β, 0)† and S[2] = B(β, λΔsΔτ)†
26        S[1] = compute_B_dagger(ΔsΔτ, 0) * S[2]
27        l = l' = 0                             # update time slice index
28        greens_upwards()                      # switch direction (u-turn)
29      end
30
31    # regular time slice
32  else
33    propagate_G_downwards()                  # Eq. (5.53)
34    l = l'                                    # update time slice index
35  end
36 end
```

Algorithm 1: Green's function propagation and stack logic for a DQMC downwards sweep. The function `greens_downwards` safely computes G_l' at the next time slice $l' = l - 1$ by either propagating the current Green's function G_l , Eq. (4.34), or computing it from the stack, Eq. (4.40) and Fig. 4.14. When the lower imaginary time boundary is reached, $l' = 0$, the sweep direction is reversed (u-turn) and `greens_upwards` takes over. The indices are as visualized in Fig. 4.13 and as described in the main text.

```

1 function greens_upwards()
2    $l' = l + 1$  # time slice we're heading to
3   # stabilization time slice?
4   if  $l' \in \Lambda$ 
5      $\lambda' = (l' - 1) / \Delta_s + 1$  # stack index
6     # lower boundary:  $l = 0, l' = 1, \lambda' = 1$ 
7     if  $\lambda' == 1$ 
8        $B\_R = S[1]$  # read  $B(\beta, 0)^\dagger$ 
9        $B\_L = 1$  # set  $B_L = B(0, 0) = \mathbb{1}$ 
10       $S[1] = B\_L$  # overwrite first stack slot
11      compute_G( $B\_L, B\_R$ ) # Eq. (5.56)
12       $l = l'$  # update time slice index
13
14      # intermediate time slice
15      elseif  $1 < \lambda' < M_\lambda$ 
16         $B\_R = S[\lambda']$  # read  $B(\beta, (\lambda' - 1)\Delta_s\Delta\tau)^\dagger$ 
17        #  $S[\lambda'] = B((\lambda' - 1)\Delta_s\Delta\tau, 0)$  and  $S[\lambda' - 1] = B((\lambda' - 2)\Delta_s\Delta\tau, 0)$ 
18         $S[\lambda'] = \text{compute\_B}((\lambda' - 1)\Delta_s\Delta\tau, (\lambda' - 2)\Delta_s\Delta\tau) * S[\lambda' - 1]$ 
19         $B\_L = S[\lambda']$  #  $B_L = B((\lambda' - 1)\Delta_s\Delta\tau, 0)$ 
20        compute_G( $B\_L, B\_R$ ) # Eq. (5.56)
21         $l = l'$  # update time slice index
22
23        # upper boundary:  $l = M, l' = M + 1, \lambda' = M_\lambda$ 
24      else
25        #  $S[M_\lambda] = B(\beta, 0)$  and  $S[M_\lambda - 1] = B((M_\lambda - 2)\Delta_s\Delta\tau, 0)$ 
26         $S[M_\lambda] = \text{compute\_B}((M_\lambda - 1)\Delta_s\Delta\tau, (M_\lambda - 2)\Delta_s\Delta\tau) * S[M_\lambda - 1]$ 
27         $l = l' = M + 1$  # update time slice index
28        greens_downwards() # switch direction (u-turn)
29      end
30
31    # regular time slice
32  else
33    propagate_G_upwards() # Eq. (5.54)
34     $l = l'$  # update time slice index
35  end
36 end

```

Algorithm 2: Green's function propagation and stack logic for a DQMC upwards sweep. The function `greens_upwards` safely computes G'_l at the next time slice $l' = l + 1$ by either propagating the current Green's function G_l , Eq. (4.35), or computing it from the stack, Eq. (4.40) and Fig. 4.14. When the upper imaginary time boundary is reached, $l' = M + 1$, the sweep direction is reversed (u-turn) and `greens_downwards` takes over. The indices are as visualized in Fig. 4.13 and as described in the main text.

4.2.2 Adaptive global updates

We supplement the local updates discussed in Sec. 4.2.1 by global updates in which the bosonic field ϕ is altered at all space-time lattice points at once. This drastically improves the rate of convergence of the Markov walk and reduces the autocorrelation time τ , Eq. (2.47) [120, 122, 131, 149]. Specifically, we apply a uniform constant shift to all elements ϕ_{li} . As for the local updates, Fig. 4.9, we draw the modification vector from a box of width $\hat{\Delta}_{\text{box}}$ and, during Monte Carlo equilibration, tune the latter to achieve a Metropolis acceptance rate of approximately 50%.

Since all slices and lattice sites are involved in the update, we need to rebuild ($O(\beta N^3)$ operations) the entire time slice matrix stack, Fig. 4.14, to compute the equal-times Green's function G . Afterwards, given the UDT decomposition of G , we can write each determinant factor in the Metropolis acceptance rate as [149]

$$\det G \stackrel{\text{sign-free}}{=} |\det G| = |\det (UDT)^{-1}| = |\det D^{-1}| = \prod_{\nu} \sigma_{\nu}^{-1}, \quad (4.42)$$

where the σ_{ν} are the singular values of G^{-1} . The determinant ratio in Eq. (2.23) hence takes the form

$$R = \frac{\prod_{\nu} \sigma_{\nu}^{-1}}{\prod_{\mu} \sigma'_{\mu}{}^{-1}}. \quad (4.43)$$

Care has to be taken in the evaluation of this expression to avoid a numerical mixing of singular values of vastly different magnitude. Practically, we order the singular values by their absolute value, pair up σ_{ν} and σ'_{ν} of similar scale, and perform the division logarithmically,

$$\begin{aligned} R &= \frac{\prod_{\nu} \sigma_{\nu}^{-1}}{\prod_{\mu} \sigma'_{\mu}{}^{-1}} = \prod_{\nu} \frac{\sigma'_{\nu}}{\sigma_{\nu}} \\ &= \exp \left(\ln \left(\prod_{\nu} \frac{\sigma'_{\nu}}{\sigma_{\nu}} \right) \right) \\ &= \exp \left(\sum_{\nu} [\ln \sigma'_{\nu} - \ln \sigma_{\nu}] \right). \end{aligned} \quad (4.44)$$

4.2.3 Checkerboard decomposition

A central numerical operation in DQMC is the computation of matrix multiplications of the form $B_l M$ involving an arbitrary matrix M and a time slice matrix B_l . Examples are given in Eqs. (4.39) and (4.35). Generally, the B_l factor is dense. To see the reason, we consider its building blocks $e^{-\Delta\tau V_l}$ and $e^{-\frac{\Delta\tau}{2}T}$, Eq. (2.22). We first note that inter-particle interactions in condensed matter systems like the Hubbard model, Eq. (1.17), and the spin-fermion model, Eq. (1.9), are of on-site nature. In matrix representation, this implies that both V_l as well as its matrix exponential are diagonal and hence sparse. Hopping processes on the other hand, while also local in the sense of typically being restricted to neighboring lattice sites, generally lead to dense matrix exponentials

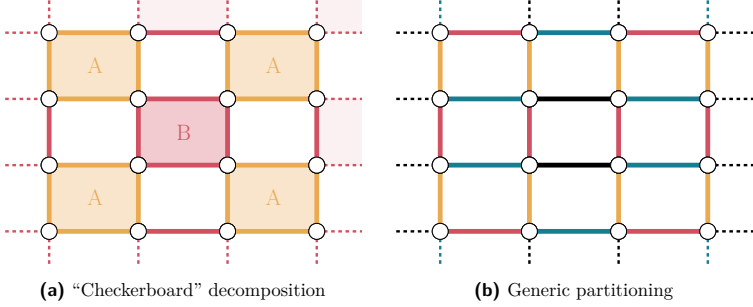


Figure 4.15: Decomposition of nearest neighbor hopping processes (bonds) on a two-dimensional square lattice into groups G (different colors) of commuting elements $T_g^{(i)}$. Shown are the square lattice specific “checkerboard” decomposition [122] with $G \in \{A, B\}$ (a) and a generic partitioning with $G \in \{1, 2, 3, 4\}$ (b). The latter has been produced using a graph-agnostic *edge coloring* algorithm, Alg. 3. Subfigure (a) is inspired by Ref. [122].

because of their purely off-diagonal character. Note that this is true in spite of the fact that the hopping matrix T itself is sparse and contains only $O(N)$ finite elements. With respect to the product above, the time slice matrix B_l inherits the denseness of $e^{-\frac{\Delta\tau}{2}T}$ and the number of numerical operations associated with a straightforward evaluation of the matrix-matrix multiplication scales as $O(N^3)$ [122, 124].

As pointed out in Refs. [90, 122, 149, 177], it is possible to effectively reduce this computational cost at the expense of introducing a systematic numerical error. The strategy is to divide the kinetic matrix T into smaller sub-blocks and replace the dense multiplications in

$$B_l M = e^{-\frac{\Delta\tau}{2}T} e^{-\Delta\tau V_{\phi(l)}} e^{-\frac{\Delta\tau}{2}T} M \quad (4.45)$$

by a series of successive sparse multiplications. This way, the dense matrix exponential $e^{-\frac{\Delta\tau}{2}T}$ is never computed explicitly and the sparsity of T is propagated to each of the individual sub-multiplications. Various implementations of this idea, differing in how the decomposition is designed, have been proposed in the literature [90, 122, 177]. Focusing on a square lattice with exclusive nearest-neighbor hoppings first, we discuss the most efficient variant in this case, the so-called *checkerboard decomposition* [122].

In the checkerboard decomposition [122] hopping processes are divided into two groups $G \in \{A, B\}$ as visualized in Fig. 4.15a. Each colored square represents a four-site hopping plaquette capturing the fermion movements along the bordering bonds. Explicitly, within a fixed fermion flavor sector, the associated $N \times N$

hopping matrices $T_g^{(i)}$ are of the form

$$T_g^{(i)} = \begin{pmatrix} & i_1 & i_2 & & i_3 & i_4 \\ & & & & & \\ & & -t_v & & -t_h & \\ -t_v & & & & & -t_h \\ & & & & & \\ -t_h & & & & -t_v & \\ & & -t_h & & -t_v & \\ & & & & & \end{pmatrix} \begin{pmatrix} i_1 \\ i_2 \\ i_3 \\ i_4 \end{pmatrix}, \quad (4.46)$$

where i_1, i_2, i_3, i_4 are linear indices specifying the lattice sites of the i -th plaquette in group g and t_h, t_v are (positive) hopping amplitudes in horizontal and vertical direction. Given the spatial separation between different plaquettes within a group, two elements $T_g^{(i)}$ and $T_g^{(j)}$ commute. Analytically, one finds that the corresponding matrix exponentials read [149]

$$e^{-\frac{\Delta\tau}{2}T_g^{(i)}} = \mathbb{1} + \begin{pmatrix} & i_1 & i_2 & & i_3 & i_4 \\ & & & & & \\ c_h c_v - 1 & c_h s_v & & c_v s_h & s_h s_v \\ c_h s_v & c_h c_v - 1 & & s_h s_v & c_v s_h \\ & & & & \\ c_v s_h & s_h s_v & & c_h c_v - 1 & c_h s_v \\ s_h s_v & c_v s_h & & c_h s_v & c_h c_v - 1 \\ & & & & \end{pmatrix} \begin{pmatrix} i_1 \\ i_2 \\ i_3 \\ i_4 \end{pmatrix}, \quad (4.47)$$

in which $c_\alpha = \cosh(\Delta\tau t_\alpha/2)$, $s_\alpha = \sinh(\Delta\tau t_\alpha/2)$, and $\alpha \in h, v$. Importantly, since the $T_g^{(i)}$ only represent hoppings between four sites, the individual matrix exponentials $e^{-\frac{\Delta\tau}{2}T_g^{(i)}}$ have only eight non-vanishing entries irrespective of the size of system and are therefore sparse in the $N \rightarrow \infty$ limit. Defining

$$T_A = \sum_{i=1}^{N/4} T_A^{(i)}, \quad T_B = \sum_{i=1}^{N/4} T_B^{(i)}, \quad (4.48)$$

we can approximate the dense sub-products of Eq. (4.45) as [122]

$$\begin{aligned} e^{-\frac{\Delta\tau}{2}T} M &= e^{-\frac{\Delta\tau}{2}T_A} e^{-\frac{\Delta\tau}{2}T_B} M + O(\Delta\tau^2) \\ &= \prod_{i=1}^{N/4} e^{-\frac{\Delta\tau}{2}T_A^{(i)}} \prod_{j=1}^{N/4} e^{-\frac{\Delta\tau}{2}T_B^{(j)}} M + O(\Delta\tau^2). \end{aligned} \quad (4.49)$$

```

1  function generic_partitioning(bonds::Set)
2      uncolored_bonds = copy(bonds)
3      coloring = Dict{<int>,<int>}()
4      color = 1
5
6      while !isempty(uncolored_bonds)
7          visited_sites = Set{<int>}()
8          for bond in uncolored_bonds
9              if !(bond.site1 in visited_sites ||
10                 bond.site2 in visited_sites)
11                  coloring[bond] = color
12                  delete!(uncolored_bonds, bond)
13                  push!(visited_sites, bond.site1)
14                  push!(visited_sites, bond.site2)
15              end
16          end
17          color += 1
18      end
19
20      return coloring
21  end

```

Algorithm 3: Schematic description of an *edge coloring* algorithm that colors the bonds (edges) of a lattice (graph) such that no two bonds sharing a common site (vertex) end up with the same color. For a square lattice with nearest neighbor connections, the algorithm leads to the coloring shown in Fig. 4.15b.

Here, each sparse factor $e^{-\frac{\Delta\tau}{2}T_g^{(i)}}$ may be multiplied to M individually, which is feasible in $O(N)$ operations [122, 124]. This implies that the overall computational complexity of the right hand side of Eq. (4.49) is $O(N \cdot N) = O(N^2)$ which, compared to the original $O(N^3)$ scaling, is significantly improved by one power of system size. The introduced numerical error is of Trotter type and due to the fact that T_A and T_B do not commute – the splitting between different group elements $T_g^{(i)}$ is exact since $[T_g^{(i)}, T_g^{(j)}] = 0$ as ensured by the decomposition scheme. It is of the same order as the already present systematic error associated with the discretization of imaginary time, Eq. (2.20), so that the checkerboard decomposition does not diminish the DQMC precision.

While the strategy underlying Eq. (4.49) is general, the specific two-group checkerboard decomposition is fine-tuned to the square lattice with nearest neighbor hopping. In particular, it is not compatible with the inclusion of higher order hopping processes, see Fig. 5.1, which we will utilize to construct locally nested Fermi surfaces in Ch. 5. In DQMC simulations of the latter, we implement the generic partitioning scheme shown in Alg. 3. Here, we make use of the observation that the task of identifying distinct groups of commuting hopping processes represents an *edge coloring* problem, in which the edges (bonds) of an arbitrary graph (lattice) shall be colored in such a way that no two edges sharing

a common vertex (lattice site) have the same color and the number of colors is minimal. Although finding an optimal edge coloring is known to be NP-hard [156], the simple Alg. 3 suffices to obtain significant speed ups in our DQMC use case. Applying it to a square lattice with nearest neighbor hopping, we obtain the four-group partitioning illustrated in Fig. 4.15b. In this case, Eq. (4.49) takes the general form

$$e^{-\frac{\Delta\tau}{2}T}M = \prod_{g \in G} \prod_{i \in g} e^{-\frac{\Delta\tau}{2}T_g^{(i)}} M + O(\Delta\tau^2). \quad (4.50)$$

Upon including further range second and third order hoppings (see Fig. 5.1 in Ch. 5), the number of groups increases to 14 (not shown). Although the scaling arguments discussed around Eq. (4.49) still hold in this case, having more distinct “checkerboard” groups implies that the numerical – system size independent – prefactor in the computational cost function of a practical DQMC simulation is larger.

4.2.4 Effective imaginary time slice propagators

To decrease the necessary total number of sparse matrix multiplications in our DQMC simulations, in particular in products described by Eqs. (4.49) and (4.50), we replace the time slice matrices B_l , Eq. (2.22), by modified effective time slice propagators \tilde{B}_l [149]. Focusing on the square lattice checkerboard decomposition, Eq. (4.50), for simplicity, we define

$$\tilde{B}_l = e^{-\frac{\Delta\tau}{2}T_B} e^{-\Delta\tau T_A} e^{-\frac{\Delta\tau}{2}T_B} e^{-\Delta\tau V_{\phi(l)}}, \quad (4.51)$$

which relates to the original time slice matrix as

$$\begin{aligned} B_l &= e^{-\frac{\Delta\tau}{2}T} e^{-\Delta\tau V_{\phi(l)}} e^{-\frac{\Delta\tau}{2}T} \\ &= e^{-\frac{\Delta\tau}{2}T_A} e^{-\frac{\Delta\tau}{2}T_B} e^{-\Delta\tau V_{\phi(l)}} e^{-\frac{\Delta\tau}{2}T_B} e^{-\frac{\Delta\tau}{2}T_A} + O(\Delta\tau^2) \\ &= e^{\frac{\Delta\tau}{2}T_A} e^{\frac{\Delta\tau}{2}T_B} \underbrace{\left(e^{-\frac{\Delta\tau}{2}T_B} e^{-\Delta\tau T_A} e^{-\frac{\Delta\tau}{2}T_B} e^{-\Delta\tau V_{\phi(l)}} \right)}_{\equiv \tilde{B}_l} \\ &\quad \times e^{-\frac{\Delta\tau}{2}T_B} e^{-\frac{\Delta\tau}{2}T_A} + O(\Delta\tau^2). \end{aligned} \quad (4.52)$$

Here, we have performed the checkerboard splitting of T into parts T_A and T_B as described in the previous section, Eq. (4.48), and the numerical Trotter error, stemming from $[T_A, T_B] \neq 0$, is hence of the same kind and order as in Eq. (4.49). Since an effective time slice propagator \tilde{B}_l includes one less kinetic matrix exponential than the corresponding B_l , we save one sparse matrix multiplication per application.

The crucial point is that the switch to effective time slice propagators does not spoil the numerical exactness of the DQMC procedure. To highlight this fact, we consider the effective equal-time Green’s function

$$\tilde{G}_\phi = \left[1 + \prod_{l=M}^1 \tilde{B}_l \right]^{-1}, \quad (4.53)$$

and realize that it is related to the original Green's function, Eq. (2.21), by

$$\begin{aligned} G_\phi &= [1 + B_M \cdots B_1]^{-1} \\ &= e^{\frac{\Delta\tau}{2}T_A} e^{\frac{\Delta\tau}{2}T_B} \left[1 + \tilde{B}_M \cdots \tilde{B}_1\right]^{-1} e^{-\frac{\Delta\tau}{2}T_B} e^{-\frac{\Delta\tau}{2}T_A} \\ &= e^{\frac{\Delta\tau}{2}T_A} e^{\frac{\Delta\tau}{2}T_B} \tilde{G}_\phi e^{-\frac{\Delta\tau}{2}T_B} e^{-\frac{\Delta\tau}{2}T_A}. \end{aligned} \quad (4.54)$$

In combination with the matrix identity

$$\det(1 + X^{-1}\tilde{B}_l X) = \det(1 + \tilde{B}_l X X^{-1}) = \det(1 + \tilde{B}_l), \quad (4.55)$$

this proves that the Metropolis acceptance, Eq. (2.23), is invariant under the operational substitution $B_l \rightarrow \tilde{B}_l$. However, the hopping matrix exponentials must be incorporated according to the last line of Eq. (4.54) when computing fermionic observables directly from the Green's function, Sec. 2.5.3.

4.3 Reducing finite-size effects

As discussed around Eq. (4.33), the scaling of the computational cost of local updates in DQMC is cubic in system size, $O(\beta N^3)$. As a consequence, depending on the number of fermion flavors and the performance of the available hardware, only small to moderately large physical systems are accessible in practical simulations. To reliably extrapolate results for physical observables to the thermodynamic limit, potential finite-size effects must be under control. In the following we describe two methods employed in this thesis to artificially improve the quality of DQMC simulations in this regard.

4.3.1 Artificial magnetic flux

As shown in Refs. [32, 122, 149, 161] threading a single magnetic flux quantum ϕ_0 through the simulation cell can drastically improve the accuracy of physical observables. For this reason, in most of our simulations, we introduce Peierls phases into the kinetic part of a given model [31, 32, 161],

$$\sum_{\alpha,s} \sum_{\langle i,j \rangle} t_{ij}^{\alpha s} e^{iA_{ij}^{\alpha s}} c_{i\alpha s}^\dagger c_{j\alpha s} + \text{h.c.} \quad (4.56)$$

Thereby, we choose the vector potential $\mathbf{A}^{\alpha s}(\mathbf{r}) = -B_{\alpha s} y \hat{x}$ in Landau gauge such that the magnetic field⁸² $\mathbf{B}^{\alpha s} = \nabla \times \mathbf{A}^{\alpha s} = B_{\alpha s} \hat{z}$ is perpendicular to the two-dimensional square lattice and the phases are given by $A_{ij}^{\alpha s} = \frac{2\pi}{\phi_0} \int_i^j \mathbf{A}^{\alpha s} \cdot d\mathbf{l}$. To retain the antiunitary symmetry $T = -i\sigma_y \kappa_z C$ that renders our model signproblem free, Sec. 3.4.1, we choose the following flavor and spin dependent amplitudes,

$$B_{x\uparrow} = B_{y\downarrow} = \frac{\phi_0}{L^2} = -B_{x\downarrow} = -B_{y\uparrow}. \quad (4.57)$$

Clearly, in the thermodynamic limit $L \rightarrow \infty$ the magnetic field vanishes and we recover the original system.

⁸²For simplicity, we will refer to \mathbf{B} as a magnetic field despite the fact that the (signs of the) amplitudes are unphysical in the sense that they are fermion flavor dependent, Eq. (4.57).




Figure 4.16: Inverse AFM susceptibilities $\chi_\gamma^{-1} = \int_\tau \sum_{\mathbf{r}} \langle \phi_{\mathbf{r},\gamma}(\tau) \phi_{0,\gamma}(0) \rangle$, with $\gamma \in \{x, y, z\}$ labeling the order parameter components, across the SDW transition. Simulations have been performed with a small perpendicular magnetic field, Eq. (4.57), a linear system size of $L = 12$ and $\lambda = 1$. The approximate location of the quantum critical point is $r_c \approx -1.89$, compare Fig. 5.14.

Note that our choice for B_{as} , Eq. (4.57), explicitly breaks the $SU(2)$ spin rotation symmetry of the model. However, the symmetry breaking field scales as L^{-2} , rendering it unimportant as long as the critical correlation length is smaller than the system size. Numerically, we find that for the spin-fermion model, Eq. (3.4), the parallel and perpendicular components of the antiferromagnetic susceptibility agree (up to error bars) for $r > r_c$ and only start to diverge at low temperatures on the ordered side of the QCP, Fig. 4.16.

In fixing the explicit phases A_{ij}^{as} we need to take into account that a square lattice with periodic boundary conditions is topologically equivalent to a torus. We therefore require (dropping band α and spin indices s for clarity) [122, 161]

$$e^{iA_{i+L\hat{x},j+L\hat{x}}c_{i+L\hat{x}}^\dagger c_{j+L\hat{x}}} \stackrel{!}{=} e^{iA_{ij}c_i^\dagger c_j}, \quad (4.58)$$

$$e^{iA_{i+L\hat{y},j+L\hat{y}}c_{i+L\hat{y}}^\dagger c_{j+L\hat{y}}} \stackrel{!}{=} e^{iA_{ij}c_i^\dagger c_j}, \quad (4.59)$$

for all lattice sites $i = (i_x, i_y)$ and $j = (j_x, j_y)$. This condition can be realized by implementing the following boundary conditions, [122, 161]

$$c_{i+L\hat{x}} = c_i, \quad (4.60)$$

$$c_{i+L\hat{y}} = c_i e^{i\frac{2\pi}{\phi_0} B L i_x}, \quad (4.61)$$

in which the extra phase factor in y-direction explicitly depends on i_x . Considering nearest, next-nearest, and next-next nearest neighbor hopping processes (see Fig. 5.1 in Ch. 5) and restoring fermion flavor indices one finds the following

Peierls phases.

Nearest neighbor:

$$A_{ij}^{\alpha s} = \begin{cases} -\frac{2\pi}{\phi_0} B_{\alpha s} i_y, & \leftarrow \text{hopping,} \\ +\frac{2\pi}{\phi_0} B_{\alpha s} i_y, & \rightarrow \text{hopping,} \\ 0, & \uparrow, \downarrow \text{hopping,} \\ -\frac{2\pi}{\phi_0} B_{\alpha s} L i_x, & \uparrow \text{hopping (boundary crossing),} \\ +\frac{2\pi}{\phi_0} B_{\alpha s} L i_x, & \downarrow \text{hopping (boundary crossing).} \end{cases} \quad (4.62)$$

Next-nearest neighbor:

$$A_{ij}^{\alpha s} = \begin{cases} -\frac{2\pi}{\phi_0} B_{\alpha s} \left(i_y + \frac{1}{2}\right), & \swarrow \text{hopping,} \\ +\frac{2\pi}{\phi_0} B_{\alpha s} \left(i_y + \frac{1}{2}\right), & \nearrow \text{hopping,} \\ -\frac{2\pi}{\phi_0} B_{\alpha s} \left(i_y + \frac{1}{2}\right), & \nwarrow \text{hopping,} \\ +\frac{2\pi}{\phi_0} B_{\alpha s} \left(i_y + \frac{1}{2}\right), & \searrow \text{hopping,} \\ +\frac{2\pi}{\phi_0} B_{\alpha s} \left(L i_x + \frac{1}{2}\right), & \swarrow \text{hopping (boundary crossing),} \\ -\frac{2\pi}{\phi_0} B_{\alpha s} \left(L i_x + \frac{1}{2}\right), & \nearrow \text{hopping (boundary crossing),} \\ +\frac{2\pi}{\phi_0} B_{\alpha s} \left(L i_x - \frac{1}{2}\right), & \nwarrow \text{hopping (boundary crossing),} \\ -\frac{2\pi}{\phi_0} B_{\alpha s} \left(L i_x - \frac{1}{2}\right), & \searrow \text{hopping (boundary crossing).} \end{cases} \quad (4.63)$$

Next-next-nearest neighbor:

$$A_{ij}^{\alpha s} = \begin{cases} -\frac{4\pi}{\phi_0} B_{\alpha s} i_y, & \leftarrow \text{hopping,} \\ +\frac{4\pi}{\phi_0} B_{\alpha s} i_y, & \rightarrow \text{hopping,} \\ 0, & \uparrow, \downarrow \text{hopping,} \\ -\frac{2\pi}{\phi_0} B_{\alpha s} L i_x, & \uparrow \text{hopping (boundary crossing),} \\ +\frac{2\pi}{\phi_0} B_{\alpha s} L i_x, & \downarrow \text{hopping (boundary crossing).} \end{cases} \quad (4.64)$$

4.3.2 Twisted boundary conditions

The artificial magnetic flux, introduced in the previous section, clearly breaks lattice translation symmetry. It must therefore be generally avoided in studies of momentum resolved fermionic observables such as the single-particle self-energy, Eq. (1.29). As an alternative measure to reduce finite-size effects in these cases, we follow Refs. [31, 32, 149, 180, 181] and run separate batches of simulations in which we apply twisted boundary conditions. Dropping fermion flavor indices for simplicity, these read

$$c_{i+L\hat{x}} = e^{i\varphi_{nx}} c_i, \quad (4.65)$$

$$c_{i+L\hat{y}} = e^{i\varphi_{ny}} c_i. \quad (4.66)$$

The φ_{n_γ} , with $\gamma \in x, y$, are twist angles [180] and may be interpreted to be due to an in-plane vector potential $\mathbf{A} = (n_x, n_y, 0)^T \phi_0 / 4L$ [149]. Note that $\mathbf{B} = \nabla \times \mathbf{A} = 0$ in contrast to the discussion of Sec. 4.3.1.

Specifically, we choose the angles

$$\varphi = \begin{pmatrix} \varphi_{n_x} \\ \varphi_{n_y} \end{pmatrix} = \frac{\pi}{2} \begin{pmatrix} n_x \\ n_y \end{pmatrix}, \quad (4.67)$$

in which we vary $n_\gamma = 1 \dots 4$ in each individual simulation. In this case, the discrete momenta associated with the twisted boundary conditions, Eq. (4.66), are given by [180]

$$\mathbf{k}_\mathbf{m} = (2\pi\mathbf{m} + \varphi)/L. \quad (4.68)$$

Here, the two-component index $\mathbf{m} \in \mathbb{Z}^2$ enumerates the L^2 distinct momenta in the Brillouin zone. Combining the results of simulations for all possible combinations of n_x and n_y , we obtain a 16 fold enhancement of the momentum resolution and, in this sense, probe an effectively larger system.

When fixing the twist angles φ , one must be careful to not reintroduce a sign problem, similar to the considerations taken around Eq. (4.57). For the spin-fermion model, we therefore supplement our choice in Eq. (4.67) and introduce orbital and spin dependent signs,

$$\varphi_{x\uparrow} = \varphi_{y\downarrow} = -\varphi_{x\downarrow} = -\varphi_{y\uparrow}. \quad (4.69)$$

This way, the antiunitary symmetry $T = -i\sigma_y\kappa_zC$, Sec. 3.4.1, is respected and a DQMC study can be conducted efficiently. However, it is important to note that this flavor dependency propagates to the momentum grid, Eq. (4.68), such that individual simulations deliver results at different $\mathbf{k}_\mathbf{m}$ -points for each fermion flavor. Only when combined with other DQMC runs does one obtain data for all flavors at the same momenta.

4.4 General remarks

Let us conclude this chapter with a few notes on the DQMC implementation of this thesis, Ref. [S1], and the supercomputing infrastructure utilized for the simulations of Chs. 5 and 7.

4.4.1 Benchmark and open source

As mentioned in Sec. 4.2.1, there is a high computational cost associated with the DQMC method. In practice, apart from the asymptotic $O(\beta N^3)$ scaling the total runtime of simulations is largely influenced by raw CPU power and the efficiency of the programming language – more precisely the produced binary code. For the DQMC implementation of this thesis, we have chosen the open-source, high-level programming language Julia [195, 196] which has proven to be capable of reaching state of the art performance in the field of numerical computing [197, 198]. For linear algebra operations we utilize the highly optimized Intel Math Kernel Library (MKL) [199].

In Fig. 4.17 we show a benchmark of our DQMC code against the implementations of Ref. [186] and Ref. [187] which were used in multiple studies of




Figure 4.17: Performance comparison of our Julia DQMC code [S1] and independent implementations in C++ [186] and Fortran [187]. Time measurements have been obtained by averaging over 100 up-down sweeps for the $O(3)$ symmetric spin-fermion model, Eq. (3.4), at $\beta = 5$ and are indicated by markers. In the simulations, an acceptance ratio of 50% has been enforced to avoid statistical fluctuations. Solid lines are guides to the eye, and dashed lines are linear least square fits. First appeared in Ref. [S1], used under the MIT license.

many-fermion systems, Refs. [2, 27, 31, 32, 39, 46, 81, 112, 200], and are written in the established “low-level” languages C++ and Fortran, respectively⁸³. We observe that, on average, our highly optimized Julia code is approximately 40% faster than the Fortran and about 80% faster than the C++ pendant. Note that upon exchanging Intel MKL for OpenBLAS [201] (Julia default) we observe a slowdown by $\sim 30\%$.

To the end of maximizing transparency and reproducibility, we have open sourced the DQMC implementation of this thesis in Ref. [S1]. Submodules handling the stabilization of Green’s function computations, Sec. 4.1, and statistical error estimation, Sec. 2.6, have been separately released as standalone Julia packages in Refs. [S2, S3] and may thus be readily integrated into foreign DQMC codes. Furthermore, as an attempt to improve the accessibility of DQMC for the wider physics community, we have published a more general DQMC framework in Ref. [S4], which has been derived from the specific implementation in Ref. [S1]. Among others, this framework has enabled the DQMC study of the sign-problem free flat band model of Sec. 3.4.4 in Ref. [202].

⁸³The benchmark has been conducted on the JUWELS supercomputer. The sweep duration is measured in single-core mode. The Fortran and C++ codes have been compiled with Intel compilers *icc* and *ifort* (version 19.0.3.199), respectively.

4.4.2 Supercomputing infrastructure

Being supported as part of the excellence project “Simulation of quantum-mechanical many-fermion systems” of the John von Neumann Institute for Computing, the numerical simulations of this thesis have been performed on the JUWELS cluster at the Jülich Supercomputing Centre of the Forschungszentrum Jülich and, partially, on the CHEOPS high-performance computing cluster of the University of Cologne. In terms of performance, the former ranks among the top 3 (top 40) supercomputing facilities in Germany (world-wide) [203]. The total computational effort of the presented work surpasses 10 million CPU hours and has involved the processing of more than 80 Terabytes of raw data⁸⁴, obtained in large scale DQMC simulations. To ensure the long-term availability of this raw data for future research projects, we have utilized the Tivoli Storage Manager (TSM) services of both supercomputing facilities and have backed up the scientific data on magnetic band drives.

⁸⁴This mostly comprises Markov chains of boson configurations as well as equal-time and time-displaced Green’s functions.

5 | Numerically exact results for the AFM QCP

Utilizing the techniques of the previous chapters, we present numerically exact results for the $O(3)$ symmetric metallic AFM QCP as obtained in extensive DQMC studies of the sign-problem free spin-fermion model, Eq. (3.4). The underlying large-scale simulations have been performed on the JUWELS supercomputer at the Forschungszentrum Jülich – ranking 40 on the TOP500 list [203] – and correspond to a total computing time of more than 10 million CPU hours.

In the following, we showcase the full phase diagram of the spin-fermion model, and discuss in detail the appearing intertwined magnetic and electronic ordering tendencies. In particular, we investigate the AFM crossover, as indicated by a growing magnetic correlation length, and the emergence of charge-density fluctuations as well as high-temperature superconductivity by calculating the superfluid density and the relevant susceptibilities. We then zoom into the close vicinity of the AFM QCP and study the quantum critical region by dissecting the inherent bosonic and fermionic correlations. Specifically, we determine the precise form of the low-energy propagator of the collective AFM modes and compare it to the analytical theories discussed in Ch. 1. Moreover, we investigate the fate of fermionic quasiparticles on the Fermi surface, e.g. in the vicinity of the hot spots, by considering (a proxy for) the spectral function and computing the Matsubara frequency dependence of the fermionic self-energy. As we will see below, our numerical results are in stark contrast to the theoretical $z = 1$ fixed-point discovered by Lee *et al.*, Sec. 1.4.4. As an attempt to verify physical traces of the latter, we report on an ultra-low temperature study of a series of Fermi surfaces, which approaches the limit of (almost) perfect local nesting near the hot spots. To the best of our knowledge, the considered temperature $T = 1/100$ is among the lowest ever accessed in finite-temperature DQMC simulations. Finally, we demonstrate that for certain parameters the onset of metallic antiferromagnetism in the spin-fermion model is marked by a first-order quantum phase transition – rather than a continuous one. Note that this chapter is largely based on the corresponding publication by the author of this thesis, Ref. [P1]⁸⁵.

⁸⁵Except for the twisted boundary condition analysis in Sec. 5.3.2, the author of this thesis has generated all presented results.

5.1 Fermi surface and model parameters

We consider the sign-problem free spin-fermion model, Eq. 3.4, on a two-dimensional square lattice. For convenience, we repeat the corresponding action here in a slightly different form⁸⁶,

$$\begin{aligned}
 S_{\text{SFL}} = & \int_0^\beta d\tau \sum_{\mathbf{k}} \sum_{s,\alpha} \psi_{\alpha\mathbf{k}s}^\dagger (\partial_\tau + \epsilon_{\mathbf{k},\alpha} - \mu) \psi_{\alpha\mathbf{k}s} \\
 & + \lambda \int_0^\beta d\tau \sum_{\mathbf{r}} \sum_{s,s'} e^{i\mathbf{Q}\cdot\mathbf{r}} \phi_{\mathbf{r}} \cdot \psi_{\mathbf{r}s}^\dagger \sigma_{ss'} \psi_{\mathbf{r}s'} + \text{h.c.} \\
 & + \int_0^\beta d\tau \sum_{\mathbf{r}} \left[\frac{1}{2c^2} (\partial_\tau \phi_{\mathbf{r}})^2 + \frac{1}{2} (\nabla \phi_{\mathbf{r}})^2 + \frac{r}{2} \phi_{\mathbf{r}}^2 + \frac{u}{4} (\phi_{\mathbf{r}}^2)^2 \right].
 \end{aligned} \tag{5.1}$$

With the exception of Secs. 5.4 and 5.5, we choose the following model parameters: a boson velocity of $c = 3$ and a quartic boson coupling of $u = 1$. As for the fermions, we consider the canonical energy dispersion

$$\epsilon_{\mathbf{k},\alpha} = -2[t_{\alpha h} \cos k_x + t_{\alpha v} \cos k_y], \tag{5.2}$$

associated with nearest neighbor hopping on a square lattice. In particular, we choose the chemical potential such that it has opposite signs for the two fermion flavors, i.e. $\mu_x = -\mu_y = -0.5$, and set the hopping amplitudes to $t_{xh} = t_{yv} = 1$ and $t_{yh} = t_{xv} = 0.5$. Here, as illustrated in Fig. 5.1, subscripts h, v indicate horizontal and vertical lattice directions, respectively.

Note that for this choice of parameters, the energy dispersion changes sign upon swapping the space dimensions and fermion flavors, i.e. $\epsilon_x(k_x, k_y) = -\epsilon_y(k_y, k_x)$. As a consequence, the spin-fermion model has a two-fold rotational symmetry, consisting of a combination of a $\pi/2$ real-space rotation, shifting momenta by \mathbf{Q} , a particle-hole transformation, and interchanging fermion flavors. Specifically, the action S_{SFL} , Eq. (5.1), is invariant under the transformation


$$\psi_{\alpha\mathbf{r}s} \rightarrow s e^{i\mathbf{Q}\cdot\mathbf{r}} \psi_{-\alpha R_{\pi/2}(\mathbf{r})-s}^\dagger, \tag{5.3}$$

where $R_{\pi/2}$ indicates a $\pi/2$ -rotation and we assume values for the spin and band indices s and α as in Eq. (3.27). This symmetry implies that the hot spots,

⁸⁶We have expressed the free fermion part in terms of an energy dispersion $\epsilon_{\mathbf{k},\alpha}$.

Figure 5.1:

Illustration of nearest (green), next-nearest (blue), and next-next-nearest neighbor (red) fermion hopping processes on the considered two-dimensional square lattice. Modified from Ref. [P1] under CC-BY 4.0.






Figure 5.2: Fermi surface with hot spots across the Brillouin zone. The black lines correspond to the Fermi surfaces of the two fermion flavors ψ_x (solid), ψ_y (dashed). One band has been shifted by $Q = (\pi, \pi)$ such that hot spot pairs (red) occur at crossing points. The energy of the ψ_x band is indicated by color shading in the background. Resized from Ref. [P1] under CC-BY 4.0.

where the AFM scattering is resonant, occur on the momentum space diagonals with $k_x = k_y$. This is apparent in Fig. 5.2, where we provide a visualization of the considered Fermi surface. As we will see explicitly below, a further consequence of the invariance above is that most fermionic correlation functions are C_4 symmetric.

5.2 Phase diagram ($\lambda = 2$)

In Fig. 5.3, we present the numerically exact phase diagram of the spin-fermion model obtained in large-scale DQMC simulations. We show results for a system of linear size $L = 12$, i.e. of 144 sites, and inverse temperatures in the range $3 \leq \beta \leq 40$. Note that the Yukawa coupling is set to $\lambda = 2$. In the following, we will systematically discuss the magnetic and electronic features of the phase diagram based on the corresponding susceptibilities.

5.2.1 Antiferromagnetic crossover

In two spatial dimensions, the $O(3)$ symmetry of the spin-fermion model associated with rotations of the AFM order parameter can not be spontaneously broken at finite temperatures according to Mermin and Wagner's theorem [22]. Magnetic correlations are therefore bound to be finite-ranged at any nonzero T and, instead of a classical phase transition, a crossover originates from the QCP at $T = 0$, Ch. 1. To highlight this aspect, we consider the AFM SDW




Figure 5.3: Phase diagram of the spin-fermion model, Eq. (5.1), for $\lambda = 2$ and parameters chosen as specified in Sec. 5.1. The correlation length ξ_{AFM} , Eq. (5.5), is shown for a $L = 12$ system (color coding, parameter grid indicated by grey points). In accordance with the Mermin-Wagner theorem there is no magnetic phase transition at finite temperatures but only a crossover originating from the AFM QCP at $r = r_c \approx 3.8$. For large tuning parameter values $r \gg r_c$ the system is in an ordinary Fermi-liquid phase. In the opposite limit, $r \ll r_c$ the SDW correlation length diverges as $T \rightarrow 0$. An extended dome-shaped d -wave superconducting phase is masking the QCP (blue). The maximal T_c is of the order of $E_F/20$. Resized from Ref. [P1] under CC-BY 4.0.

susceptibility

$$\chi(\mathbf{q}) = \int_{\tau} \sum_{\mathbf{r}} e^{-i\mathbf{q} \cdot \mathbf{r}} \langle \phi_{\mathbf{r}}(\tau) \cdot \phi_0(0) \rangle, \quad (5.4)$$

in which we measure momenta \mathbf{q} relative to the antiferromagnetic ordering wave vector $\mathbf{Q} = (\pi, \pi)$. Note that χ_0 can be efficiently computed from the (fast) Fourier transform of the order parameter field by means of the convolution theorem [124]. Anticipating that, at large length scales, the susceptibility is of conventional Ornstein-Zernike form [87], Eq. (1.13), we can define a correlation length of the antiferromagnetic fluctuations as [87, 88]

$$\xi_{\text{AFM}} = \frac{L}{2\pi} \sqrt{\frac{\chi(q=0)}{\chi(q=2\pi/L)} - 1}. \quad (5.5)$$

Here $q = 2\pi/L$ is the smallest non-vanishing momentum of a finite system of linear size L . In Fig. 5.3, we display ξ_{AFM} across the (r, T) -plane using a linear interpolation between a discrete set of r values and temperatures (indicated by dots in the figure). As apparent from the visualization, the data manifestly reveals a finite-temperature magnetic crossover that culminates into a quantum




Figure 5.4: Tuning parameter dependence of the inverse magnetic susceptibility, Eq. (5.4), for the spin-fermion model with $\lambda = 2$ close to the AFM QCP at inverse temperature $\beta = 40$. The black solid line indicates a linear fit of the DQMC data (markers) for the largest system size, i.e. $L = 12$. Extrapolating this fit to $T = 0$ we find the root $r_c \approx 3.8$.

critical point marking the onset of SDW order at $T = 0$. To determine the location of this AFM QCP, we perform a heuristic finite-size analysis (see App. C) and extrapolate the tuning-parameter dependence of the inverse susceptibility χ^{-1} to $T = 0$. As shown in Fig. 5.4, this procedure leads to the estimate $r_c \approx 3.8$.

5.2.2 Unconventional superconductivity

While the Mermin-Wagner theorem also holds for the continuous $U(1)$ symmetry breaking associated with the onset of superconductivity (SC), in this case there is the possibility of a Berezinskii-Kosterlitz-Thouless (BKT) transition [17, 204] for the corresponding two-component order parameter. To determine the critical temperature T_c across the phase diagram we follow Refs. [32, 137] and characterize the superconducting state by its diamagnetic linear response to a static magnetic field $B(\mathbf{q})$. In terms of the current density operator⁸⁷,

$$j_x(\mathbf{r}_i, \tau) = i \sum_{\alpha s} \underbrace{t_{ij}^{\alpha s} \psi_{\alpha i s}^\dagger \psi_{\alpha j s} - t_{ji}^{\alpha s} \psi_{\alpha j s}^\dagger \psi_{\alpha i s}}_{j_x^{\alpha i s}(\tau)}, \quad (5.6)$$

in which $\mathbf{r}_j = \mathbf{r}_i + \hat{\mathbf{x}}$, the latter reads [32, 137]

$$j_x(\mathbf{q}) = -\frac{4iF_{xx}(\mathbf{q})}{q_y}B(\mathbf{q}). \quad (5.7)$$

⁸⁷Without loss of generality we focus on the x -direction.

Here, the response function

$$F_{xx}(\mathbf{q}) = \frac{1}{4} [\Lambda_{xx}(q_x \rightarrow 0, q_y = 0) - \Lambda_{xx}(\mathbf{q})] \quad (5.8)$$

is given by the imaginary time current-current correlator

$$\begin{aligned} \Lambda_{xx} &= \sum_i \int d\tau e^{-i\mathbf{q}\cdot\mathbf{r}_i} \langle j_x(\mathbf{r}_i, \tau) j_x(0, 0) \rangle \\ &= - \sum_i \sum_{\alpha s, \alpha' s'} \int d\tau e^{-i\mathbf{q}\cdot\mathbf{r}_i} \langle j_x^{\alpha is}(\tau) j_x^{\alpha' 0 s'} \rangle. \end{aligned} \quad (5.9)$$

To evaluate Λ_{xx} in DQMC, we note that each expectation value in Eq. (5.9) unfolds to

$$\begin{aligned} \langle j_x^{\alpha is}(\tau) j_x^{\alpha' 0 s'} \rangle &= \langle t_{ij}^{\alpha s} t_{00'}^{\alpha' s'} \langle \psi_{\alpha is}^\dagger(\tau) \psi_{\alpha js}(\tau) \psi_{\alpha' 0 s'}^\dagger \psi_{\alpha' 0 s'} \rangle_\phi \\ &\quad - t_{ij}^{\alpha s} t_{0'0}^{\alpha' s'} \langle \psi_{\alpha is}^\dagger(\tau) \psi_{\alpha js}(\tau) \psi_{\alpha' 0 s'}^\dagger \psi_{\alpha' 0 s'} \rangle_\phi \\ &\quad - t_{ji}^{\alpha s} t_{00'}^{\alpha' s'} \langle \psi_{\alpha js}^\dagger(\tau) \psi_{\alpha is}(\tau) \psi_{\alpha' 0 s'}^\dagger \psi_{\alpha' 0 s'} \rangle_\phi \\ &\quad + t_{ji}^{\alpha s} t_{0'0}^{\alpha' s'} \langle \psi_{\alpha js}^\dagger(\tau) \psi_{\alpha is}(\tau) \psi_{\alpha' 0 s'}^\dagger \psi_{\alpha' 0 s'} \rangle_\phi \rangle_{\text{MC}}, \end{aligned} \quad (5.10)$$

where 0 and 0' indicate reference points \mathbf{r}_0 and $\mathbf{r}_{0'} = \mathbf{r}_0 + \hat{\mathbf{x}}$, respectively⁸⁸. Successively employing Wick's theorem in the form of Eq. (2.38) to each of the fermionic expectation values $\langle \cdot \rangle_\phi$ one finds the following expression involving only equal-time and time-displaced Green's functions,

$$\begin{aligned} \langle j_x^{\alpha is}(\tau) j_x^{\alpha' 0 s'} \rangle &= \left\langle \left[t_{ij}^{\alpha s} G_{\alpha js; \alpha is}(\tau) - t_{ji}^{\alpha s} G_{\alpha is; \alpha js}(\tau) \right] \right. \\ &\quad \times \left[t_{00'}^{\alpha' s'} G_{\alpha' 0 s'; \alpha' 0 s'}(0) - t_{0'0}^{\alpha' s'} G_{\alpha' 0 s'; \alpha' 0 s'}(0) \right] \\ &\quad - t_{ij}^{\alpha s} t_{00'}^{\alpha' s'} G_{\alpha' 0 s'; \alpha is}(0, \tau) G_{\alpha js; \alpha' 0 s'}(\tau, 0) \\ &\quad + t_{ij}^{\alpha s} t_{0'0}^{\alpha' s'} G_{\alpha' 0 s'; \alpha is}(0, \tau) G_{\alpha js; \alpha' 0 s'}(\tau, 0) \\ &\quad + t_{ji}^{\alpha s} t_{00'}^{\alpha' s'} G_{\alpha' 0 s'; \alpha js}(0, \tau) G_{\alpha is; \alpha' 0 s'}(\tau, 0) \\ &\quad \left. - t_{ji}^{\alpha s} t_{0'0}^{\alpha' s'} G_{\alpha' 0 s'; \alpha js}(0, \tau) G_{\alpha is; \alpha' 0 s'}(\tau, 0) \right\rangle_{\text{MC}}. \end{aligned} \quad (5.11)$$

As established by Scalapino *et al.* in Ref. [137], a reliable criterion for detecting superconductivity is whether the superfluid density of the system, [137]

$$\rho_s = \lim_{q_y \rightarrow 0} \lim_{L \rightarrow \infty} F_{xx}(q_x = 0, q_y), \quad (5.12)$$

surpasses the universal BKT value $\Delta\rho_s = 2T/\pi$ [32, 137]. Given the finite-size nature of our DQMC simulations, we can not compute ρ_s in the thermodynamic limit. Instead, we consider the system size dependent proxy [32]

$$\rho_s(L) = F_{xx}(q_x = 0, q_y = 2\pi/L), \quad (5.13)$$

which we denote by the same label for convenience. As in the definition of the magnetic correlation length, Eq. (5.5), the momentum $q = 2\pi/L$ is the

⁸⁸We average over these reference points in our computation of Λ_{xx} .




Figure 5.5: Temperature dependence of the superfluid density of the spin-fermion model, Eq. (5.1), for $\lambda = 2$. The colors indicate three different tuning parameter values $r = 4.2 \approx r_{T_c^{\max}}$ (blue), $r = 5.1 > r_{T_c^{\max}}$ (orange), and $r = 6.0 \gg r_{T_c^{\max}}$ (green). The dashed line indicates the critical BKT value $\Delta\rho_s = 2T/\pi$. Resized from Ref. [P1] under CC-BY 4.0.

smallest non-zero momentum of the finite system such that $\rho_s(L) \rightarrow \rho_s$ in the limit $L \rightarrow \infty$. Applying Scalapino’s criterion to the spin-fermion model, we identify the critical temperature T_c for each value of the tuning parameter r as the temperature for which $\rho_s(L_{\max}; r) = \Delta\rho_s$. Here, L_{\max} indicates the largest available system size, i.e. $L = 12$. In Figs. 5.5 and 5.6, we show the temperature and tuning parameter dependence of the superfluid density for three values of r and nine values of T , respectively. The biggest source of error in resolving $T_c(r)$ is the limitation to finite system sizes, see Fig. 5.5. However, based on the available data we anticipate that deviations originating from finite-size effects are typically $\lesssim 10\%$ of the determined transition temperature.

Employing the strategy above, we find that the spin-fermion model hosts an extended dome-shaped superconducting phase, as visualized in the phase diagram in Fig. 5.3. Notably, the maximal critical temperature $T_c^{\max} = \max(T_c(r)) \approx 0.09$ occurs close to the AFM QCP, which is indicative of the intertwined character of magnetic and electronic ordering: the critical antiferromagnetic fluctuations both facilitate and compete with superconductivity. Remarkably, compared to the Fermi energy E_F , the value of T_c^{\max} is of the order of $E_F/20$. This marks that the onset of superconductivity occurs at *high* temperatures. Compared to theory, we note that T_c^{\max} is in good agreement with an independent self-consistent Eliashberg calculation⁸⁹, Sec. 1.4.3, which gives $T_c^{\text{El}} \approx 0.08$. Anticipating that fermions at the hot spots dominate in the formation of superconductivity [112] and, furthermore, that the associated energy scale T_c is independent of the

⁸⁹Up to the different model parameters, Sec. 5.1, the calculation corresponds to the one presented in Ref. [112].




Figure 5.6: Superfluid density of the spin-fermion model, Eq. (5.1), as a function of the tuning parameter r for various temperatures and $\lambda = 2$. The dashed line indicates the critical BKT value $\Delta\rho_s = 2T/\pi$.

precise band structure, our results may be compared to those of Refs. [31, 32], which consider an $O(2)$ symmetric spin-fermion model featuring easy-plane antiferromagnetism. For the maximal critical temperatures, we find the ratio $T_c^{O(3)}/T_c^{O(2)} \approx 2.2$, which is consistent with a quadratic dependence of T_c on the number of order parameter components as suggested by weak-coupling theory, Eq. (1.30).

Pairing correlations

It is important to note that the superfluid density criterion is unbiased towards particular pairing symmetries. To investigate the nature of the superconducting state we define order parameters

$$\Delta_\eta(\mathbf{r}_i, \tau) = \psi_{xi\uparrow}^\dagger \psi_{xi\downarrow}^\dagger - \eta \psi_{yi\uparrow}^\dagger \psi_{yi\downarrow}^\dagger, \quad (5.14)$$

with $\eta = \pm 1$, and consider the zero-frequency pairing susceptibilities

$$P_\eta(\mathbf{r}_i) = \int d\tau \langle \Delta_\eta^\dagger(\mathbf{r}_i, \tau) \Delta_\eta(0, 0) \rangle. \quad (5.15)$$

Note that under a $\pi/2$ rotation $\Delta_\eta \rightarrow \eta \Delta_\eta^\dagger$ such that Δ_- and Δ_+ have d -wave and s -wave character, respectively. Similar to the current-current correlation function, the expectation value in Eq. (5.15) may be computed in DQMC via Wick's




Figure 5.7: Nature of the superconducting state. Momentum resolved equal-time pairing correlations across the first Brillouin zone with d -wave (left) and s -wave symmetry (right) close to the quantum critical point, $r \approx 4.2$, for $\lambda = 2$, $\beta = 20$, and $L = 12$. The pairing correlations have been normalized to system size. While the s -wave correlations are featureless there is a distinct peak in the d -wave channel. Taken from Ref. [P1] under CC-BY 4.0.

theorem in the form of Eq. (2.36). Specifically, we find the expression

$$\begin{aligned}
 \langle \Delta_{\eta}^{\dagger}(\mathbf{r}_i, \tau) \Delta_{\eta} \rangle = & \langle G_{xi\downarrow; x0\downarrow}(\tau, 0) G_{xi\uparrow; x0\uparrow}(\tau, 0) - G_{xi\downarrow; x0\uparrow}(\tau, 0) G_{xi\uparrow; x0\downarrow}(\tau, 0) \\
 & + \eta [G_{xi\downarrow; y0\downarrow}(\tau, 0) G_{xi\uparrow; y0\uparrow}(\tau, 0) - G_{xi\downarrow; y0\uparrow}(\tau, 0) G_{xi\uparrow; y0\downarrow}(\tau, 0)] \\
 & + \eta [G_{yi\downarrow; x0\downarrow}(\tau, 0) G_{yi\uparrow; x0\uparrow}(\tau, 0) - G_{yi\downarrow; x0\uparrow}(\tau, 0) G_{yi\uparrow; x0\downarrow}(\tau, 0)] \\
 & + G_{yi\downarrow; y0\downarrow}(\tau, 0) G_{yi\uparrow; y0\uparrow}(\tau, 0) - G_{yi\downarrow; y0\uparrow}(\tau, 0) G_{yi\uparrow; y0\downarrow}(\tau, 0) \rangle_{\text{MC}}.
 \end{aligned} \tag{5.16}$$

In Fig. 5.7, we show the Fourier-transformed correlation functions $P_{\eta}(\mathbf{q})$ across the first Brillouin zone. While the s -wave susceptibility is featureless, there is a pronounced peak in the sign-changing symmetry channel. This is strong evidence for a d -wave superconducting state. Note that the pairing correlations are C_4 -symmetric because of the corresponding rotational symmetry of the spin-fermion model, Eq. (5.3).

Since the SC transition is of BKT character, we anticipate that the correlation length ξ_{SC} is infinite in the entire superconducting phase. Accordingly, spatial correlation functions are expected to decay exponentially above T_c and with a power law below T_c [32], i.e.

$$\langle \Delta^{\dagger}(\mathbf{r}) \Delta(0) \rangle \sim \begin{cases} e^{-r/\xi_{\text{SC}}}, & T \geq T_c, \\ r^{-\eta(T)}, & T \leq T_c. \end{cases} \tag{5.17}$$

Here, $\eta(T)$ is a temperature dependent exponent – not to be confused with $\eta = \pm 1$ above – which, at the critical temperature, takes the universal BKT value $\eta(T_c) = 1/4$ [32]. Note that Eq. (5.17) implies a finite-size scaling of the pairing susceptibility $P_- = \sum_i P_-(\mathbf{r}_i)$ of the form $P_- \sim L^{2-\eta(T)}$ for $T \lesssim T_c$ ⁹⁰. In Fig. 5.8, we show the tuning parameter dependence of the pairing correlations for three different system sizes and temperatures: $T > T_c^{\max}$, $T \lesssim T_c^{\max}$, and $T < T_c^{\max}$. Upon cooling, we observe that P_- develops a pronounced, strongly system-size dependent peak in the vicinity of the AFM QCP, $r \gtrsim r_c$. In contrast, the s-wave susceptibility P_+ , shown in the insets in Fig. 5.8, is small in magnitude, featureless, and essentially system-size independent across the entire phase diagram. This confirms the d-wave symmetry of the superconducting state. Determining the transition line $T_c(r)$ based on the criterion $\eta(T_c(r)) \leq 1/4$ and a finite-size scaling analysis of the pairing correlation P_- at different temperatures, we find a critical temperature curve that is in reasonable agreement with the superfluid density study above.

5.2.3 Charge-density wave correlations

As discussed in Ch. 1, besides superconductivity, quantum critical fluctuations can, in principle, also promote other electronic orders, such as charge-density wave (CDW) states [43, 44, 47]. For example, for particle-hole symmetric electron dispersions, *d*-wave superconductivity and *d*-wave CDW are found to be degenerate due to an emergent $SU(2)$ symmetry [46]. In order to investigate the possibility of the presence of charge-density fluctuations we inspect the susceptibilities [32]

$$C_\eta(\mathbf{r}) = \int d\tau \langle \tilde{\Delta}_\eta^\dagger(\mathbf{r}_i, \tau) \tilde{\Delta}_\eta(0, 0) \rangle, \quad (5.18)$$

$$\tilde{\Delta}_\eta(\mathbf{r}_i, \tau) = \sum_{s=\uparrow, \downarrow} \psi_{xis}^\dagger \psi_{xis} + \eta \psi_{yis}^\dagger \psi_{yis}, \quad (5.19)$$

with $\eta = \pm 1$. As for the current-current correlation function Λ_{xx} , Eq. (5.9), we apply Wick's theorem in the form of Eq. (2.38) to obtain the following expression in terms of Green's functions accessible in DQMC,

$$\begin{aligned} \langle \tilde{\Delta}_\eta^\dagger(\mathbf{r}_i, \tau) \tilde{\Delta}_\eta(0, 0) \rangle &= \sum_{ss'} \langle [G_{xis;xis}(\tau) + \eta G_{yis;yis}(\tau)] \\ &\quad \times [G_{x0s';x0s'}(0) + \eta G_{y0s';y0s'}(0)] \\ &\quad - G_{x0s';xis}(0, \tau) G_{xis;x0s'}(\tau, 0) \\ &\quad - \eta G_{y0s';xis}(0, \tau) G_{xis;y0s'}(\tau, 0) \\ &\quad - \eta G_{x0s';yis}(0, \tau) G_{yis;x0s'}(\tau, 0) \\ &\quad - G_{y0s';yis}(0, \tau) G_{yis;y0s'}(\tau, 0) \rangle_{\text{MC}}. \end{aligned} \quad (5.20)$$

⁹⁰For temperatures slightly above T_c the correlation length ξ_{SC} may still be effectively infinite in the sense of being larger than the considered system size.

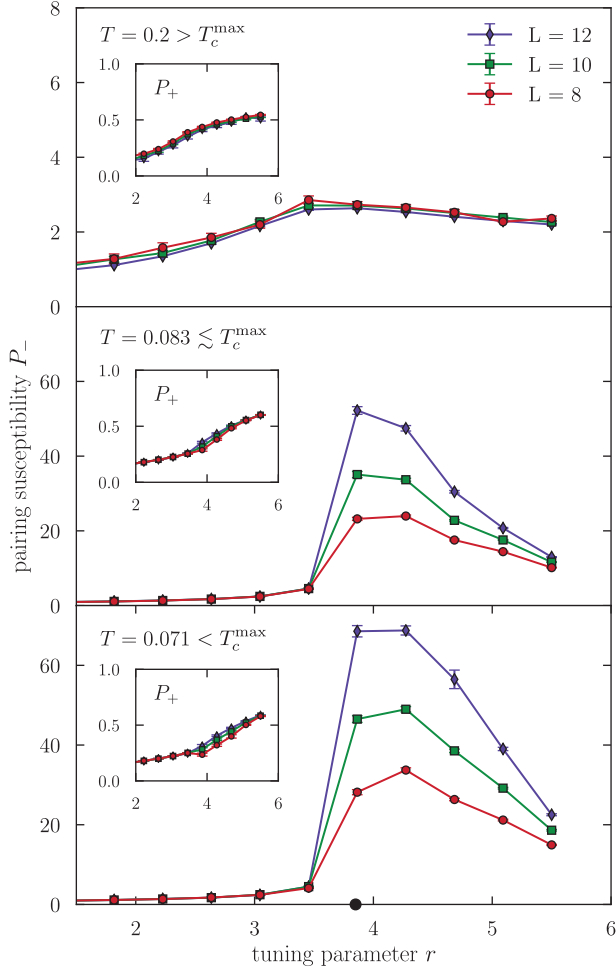


Figure 5.8: Zero-frequency pairing susceptibility $P_- = \sum_i P_-(\mathbf{r}_i)$ as a function of the tuning parameter r for three system sizes and temperatures, $T = 0.2 > T_c^{\max}$ (top), $T = 0.083 \lesssim T_c^{\max}$ (middle), and $T = 0.071 < T_c^{\max}$. The insets show the featureless susceptibility $P_+ = \sum_i P_+(\mathbf{r}_i)$ at the same temperatures. The black circle indicates the location of the putative QCP.




Figure 5.9: Momentum dependence of charge-density correlations, Eq. (5.18), across the first Brillouin zone in the vicinity of the AFM QCP, $r \approx 4.2$, for $\lambda = 2$, $\beta = 20$, and $L = 12$. The peak at $\mathbf{q} = 0$ is excluded.




Figure 5.10: Momentum cut of the charge-density susceptibility C_- , Eq. (5.18), in the vicinity of the AFM QCP, $r \approx 4.2$, for $\lambda = 2$ and $\beta = 20$. The maxima positions are $q_y \approx 0.83\pi$ and $q_y = 1.17\pi$. The solid line is a guide to the eye based on the $L = 12$ data points.

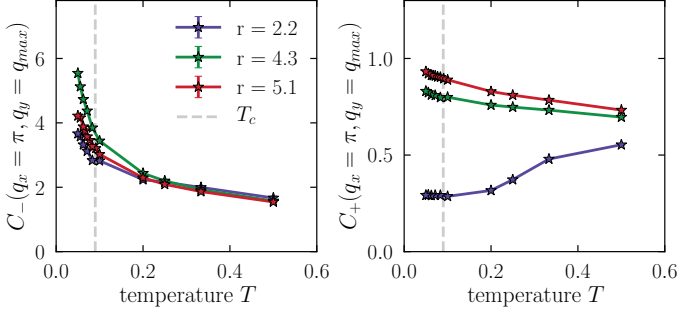


Figure 5.11: Charge-density correlations, Eq. (5.18), across the phase diagram for $\lambda = 2$ and $L = 10$. Shown is the temperature dependence of the maximum (excluding $\mathbf{q} = 0$) of the susceptibilities C_{\pm} , Eq. (5.18), for different tuning parameter values $r \ll r_c$ (purple), $r \gtrsim r_c$ (green), and $r \gg r_c$ (pink). The dashed, grey line indicates the temperature at which superconductivity sets in (for $r = 4.3 \gtrsim r_c$). Modified from Ref. [P1] under CC-BY 4.0.

In Fig. 5.9, we show the momentum dependence of C_+ and C_- in the vicinity of the QCP at temperature $T = 0.05$. We find that both C_+ and C_- are peaked close to (but not quite at) the corners of the Brillouin zone at wave vectors $\mathbf{q}_{-}^{\max} = (\pi, q_{\max} \approx 0.83\pi)$ and $\mathbf{q}_{+}^{\max} = (\pi, q_{\max} \approx 0.67\pi)$, respectively. A momentum space cut of C_- for fixed $q_x = \pi$ is provided in Fig. 5.10. As for the pairing susceptibilities, the apparent C_4 symmetry of both C_+ and C_- is due to the symmetry in Eq. (5.3). Note that the maximum of the susceptibility C_+ is much smaller than the maximum of C_-^{\max} indicating that CDW correlations are predominantly of d-wave character.

In Fig. 5.11, we show the temperature dependence of the maxima of the susceptibilities $C_{\pm}(\mathbf{q})$ for three different tuning parameter cuts through the phase diagram (Fig. 5.3). First, we notice that $C_+ < C_-$ consistently for all chosen parameters. Second, we observe that while C_+ is suppressed with decreasing temperature for $r \ll r_c$, indicating a competition with SDW fluctuations, C_- is enhanced for all tuning parameter values. Focusing on $C_-(T)$, we observe that the amplification is most pronounced close to the QCP, $r \gtrsim r_c$, and seems insensitive to the onset of superconductivity at $T = T_c$. Note that for the $O(2)$ symmetric spin-fermion model, i.e. antiferromagnetism of easy-plane character, C_- was found to drop sharply below T_c [32].

Investigating the tuning parameter dependence of CDW correlations in more detail, Fig. 5.12, we notice that C_- is moderately peaked right above the QCP, $r = r_c \approx 3.8$, and decreases upon going to smaller or larger values of r . In agreement with our assessment in Fig. 5.11, the magnitude of the maximum increases when cooling the system to lower temperatures. On the other hand, s-wave correlations indicated by C_+ do not show a distinct peak as a function of r but instead appear to be suppressed by the onset of SDW fluctuations for $r \lesssim r_c$. Noting that the system size dependence of both C_- and C_+ is very weak

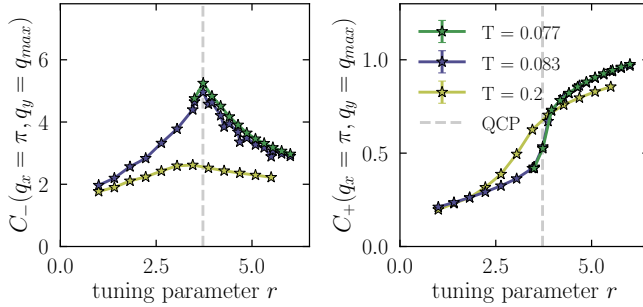


Figure 5.12: Charge-density correlations, Eq. (5.18), across the phase diagram for $\lambda = 2$ and $L = 10$. Shown is the tuning parameter dependence of the maximum (excluding $\mathbf{q} = 0$) of the susceptibilities C_{\pm} , Eq. (5.18), for different temperatures $T < T_c$ (green), $T \lesssim T_c$ (purple), and $T \gg T_c$ (yellow). The dashed, grey line indicates the position of the AFM QCP.


across the phase diagram (not shown but qualitatively similar to P_+ in Fig. 5.8), we conclude that despite a mild enhancement of d-wave CDW correlations at low temperatures there is no evidence for CDW quasi long-range order.

5.3 Quantum critical correlations ($\lambda = 1$)

We now turn to the low-energy quantum critical properties of the spin-fermion model. An obstacle in extracting the low energy scaling is the dome of high-temperature superconductivity masking the QCP, see the phase diagram presented in Fig. 5.3 for $\lambda = 2$. In order to uncover and probe the “bare” AFM QCP we must, in some way, lower the critical temperature T_c of the SC transition. In experiments, this suppression is typically implemented by applying a strong magnetic field [16, 205–207]. However, since the latter spoils time reversal invariance, it also inevitably breaks the $T = -i\sigma_y\kappa_zC$ symmetry discussed in Sec. 3.4.1 and would, thus, reintroduce a sign-problem into the DQMC simula-

Figure 5.13:

Superfluid density of the spin-fermion model, Eq. (5.1), as a function of the tuning parameter r for $T = 1/40 = 0.05$ and $\lambda = 1$. The dashed line indicates the critical BKT value $\Delta\rho_s = 2T/\pi$.






Figure 5.14: Phase diagram of the spin-fermion model, Eq. (5.1), for $\lambda = 1$. The correlation length ξ_{AFM} , Eq. (5.5), is shown for a $L = 12$ system (color coding, parameter grid indicated by grey points). In accordance with the Mermin-Wagner theorem there is no magnetic phase transition at finite temperatures but only a crossover originating from the QCP at $r = r_c$. For large tuning parameter values $r \gg r_c$ the system is in an ordinary Fermi-liquid phase. In the opposite limit, $r \ll r_c$ the SDW correlation length diverges as $T \rightarrow 0$. The quantum critical point, marking the onset of SDW order at $T = 0$, lies at $r_c \approx -1.89$, see App. C. Resized from Ref. [P1] under CC-BY 4.0.

tions. Instead, we take an alternative approach and reduce the superconducting T_c by considering the spin-fermion model at smaller⁹¹ Yukawa coupling. Clearly, one expects $T_c(\lambda)$ to be monotonously decreasing for $\lambda \rightarrow 0$ given that the non-interacting system is a conventional metal. According to Eliashberg theory, T_c is quadratic in the Yukawa coupling, Eq. (1.30), a scaling which has been numerically confirmed for the O(2) symmetric spin-fermion model in Ref. [31]. As indicated by the superfluid density in Fig. 5.13, choosing $\lambda = 1$, we find that the onset of superconductivity is suppressed below the lowest temperature considered in our DQMC simulations, i.e. $T_c < 1/40$. In this case, we obtain the phase diagram presented in Fig. 5.14. In the following, we will extract quantum critical correlations slightly above the (presumably) remaining SC phase, i.e. at temperature $T = 0.05$ and in close vicinity to the AFM QCP, $r \approx -1.74$ (see App. C).

It is important to note that while decreasing the Yukawa coupling has the desired effect of lowering T_c , it also, eventually, leads to a collapse of the hierarchy of energy scales discussed in Sec. 1.4.3. In particular, the putative quantum critical fan, Sec. 1.1.1, inevitably vanishes in the $\lambda \rightarrow 0$ limit. Its survival in form of non-Fermi liquid behavior at finite temperatures above the QCP

⁹¹Relative to $\lambda = 2$. Note that the chosen $\lambda = 1$ is still of the order of unity.




Figure 5.15: Charge-density correlations, Eq. (5.18), across the phase diagram for $\lambda = 1$ and $L = 10$. Shown is the temperature dependence of the maximum (excluding $\mathbf{q} = 0$) of the susceptibilities C_η defined in Eq. (5.18) for different tuning parameter values $r \ll r_c$ (purple), $r \gtrsim r_c$ (green), and $r \gg r_c$ (pink). Modified from Ref. [P1] under CC-BY 4.0.

must therefore be verified *a posteriori*. Moreover, anticipating that orders are intertwined [74], the suppression of superconductivity can, in principle, favor the formation of other electronic states. To rule out such an indirect promotion, we analyze the CDW susceptibilities, Eq. (5.18), across the $\lambda = 1$ phase diagram. As indicated in Fig. 5.15, we find that both C_- and C_+ have the same qualitative temperature and tuning parameter dependence as for $\lambda = 2$, c.f. Fig. 5.11. Compared to the latter, CDW correlations in the d-wave channel are reduced whereas the susceptibility C_+ appears to be slightly enhanced.

5.3.1 Antiferromagnetic correlations

We investigate the low temperature antiferromagnetic correlations by studying the bosonic SDW susceptibility

$$\chi(\mathbf{q}, i\omega_n; r, T) = \int_\tau \sum_{\mathbf{r}} e^{i\omega_n \tau - i\mathbf{q} \cdot \mathbf{r}} \langle \phi_{\mathbf{r}}(\tau) \cdot \phi_0(0) \rangle \quad (5.21)$$

for momenta \mathbf{q} , taken relative to the ordering wave vector \mathbf{Q} , and Matsubara frequencies $\omega_n = 2\pi nT$, cf. Eq. (5.4). Specifically, in our DQMC simulations, we consider linear system sizes in the range of $8 \leq L \leq 14$ and, as described above, probe the vicinity of the AFM QCP at temperature $T = 0.025$.

First, we illustrate the dependence of the inverse susceptibility $\chi^{-1}(\mathbf{q}, i\omega_n)$ on Matsubara frequency close to the QCP in Fig. 5.16. The data very visibly has linear character for small Matsubara frequencies ω_n , both for $\mathbf{q} = 0$ and small finite momenta $\mathbf{q} > 0$, with an apparent cusp at $\omega_n = 0$. To establish the presence of a $|\omega_n|$ -term in χ^{-1} we perform a linear regression for the small frequency data. The resulting fits to the form $a_1|\omega_n| + a_0$, displayed in Fig. 5.16 as black lines, are in good agreement with the data and confirm the linear Matsubara frequency dependence. This finding suggests overdamped dynamics




Figure 5.16: Frequency dependence of the magnetic susceptibility χ close to the quantum critical point ($r = -1.74$) at inverse temperature $\beta = 40$ and $\lambda = 1.0$. The different colors indicate different momenta. The solid lines are linear low-frequency fits. Resized from Ref. [P1] under CC-BY 4.0.




Figure 5.17: Momentum dependence of the magnetic susceptibility χ close to the quantum critical point ($r = -1.74$) at inverse temperature $\beta = 40$ and $\lambda = 1.0$. The different colors indicate the lowest few Matsubara frequencies. The solid lines are linear fits. Resized from Ref. [P1] under CC-BY 4.0.




Figure 5.18: Tuning parameter dependence of the inverse magnetic susceptibility, Eq. (5.21), close to the quantum critical point at inverse temperature $\beta = 40$ and $\lambda = 1.0$. We show a linear fit of the data (light green line) with root $r_c \approx -1.89$, an estimate for the location of the QCP. Resized from Ref. [P1] under CC-BY 4.0.

of the boson field ϕ due to interactions with the fermions, see Sec. 1.4.3. At vanishing momentum, finite-size effects are negligibly small (within error bars), as evident in the data collapse of $\chi^{-1}(\mathbf{q} = 0, i\omega_n)$ for different system sizes.

Turning to the momentum dependence of $\chi^{-1}(\mathbf{q}, i\omega_n)$ next, we find that the momentum dependence shown in Fig. 5.17 is consistent with a quadratic form \mathbf{q}^2 for small momenta \mathbf{q} . This holds both for $\omega_n = 0$ and small non-vanishing Matsubara frequencies ω_n . Similar to above we establish the presence of a q^2 term in $\chi^{-1}(\mathbf{q}, i\omega_n)$ by fitting the data to the form $a_1 \mathbf{q}^2 + a_0$. The results are indicated in Fig. 5.17 as black lines, showing good agreement with the DQMC data. In combination with the observed linear Matsubara frequency dependence, this provides strong evidence for a dynamical critical exponent $z = 2$.

Next, we illustrate the dependence of the inverse susceptibility $\chi^{-1}(\mathbf{q} = 0, i\omega_n = 0)$ on the tuning parameter r in Fig. 5.18. For tuning parameter values $r \geq r_c \approx -1.89$ we find that the data for different system sizes follows a linear dependence. Due to finite-size and finite-temperature effects the onset of a finite value of χ^{-1} does not appear to be as abrupt but instead is moderately continuous [31].

Finally, we inspect the temperature dependence $\chi^{-1}(T)$ close to the critical point in Fig. 5.19. Within the limits of our numerical accuracy, we find that the DQMC results are consistent with a T^2 term at low temperatures. We highlight this point by fitting the data to a second degree polynomial (black line) which is able to adequately capture the temperature trend over a broad range $0.025 < T \lesssim 0.5$. At higher temperatures, the situation changes and we find a linear T -dependence, as shown in the inset of Fig. 5.19.




Figure 5.19: Temperature dependence of the magnetic susceptibility χ at $r \approx -1.86$, close to the quantum critical point, for $\lambda = 1.0$. The black line is a quadratic fit, $f(x) = a_2x^2 + a_1x + a_0$ with $a_2 \approx 0.366$, $a_1 \approx -0.005$, and $a_0 \approx 0.009$, for system sizes $L > 8$. Resized from Ref. [P1] under CC-BY 4.0.

Taken together, these findings indicate that the inverse SDW susceptibility near the metallic AFM QCP (at low temperatures) has the form,

$$\chi^{-1}(\mathbf{q}, i\omega_n; r, T \rightarrow 0) = c_\omega |\omega_n| + c_q \mathbf{q}^2 + c_r(r - r_0). \quad (5.22)$$

Remarkably, this is precisely the functional dependence predicted by Hertz-Millis theory [50, 82], Sec. 1.4.2, in spite of the fact that the latter is understood to be formally uncontrolled [36]. Also, note that this agreement is in strong analogy to what has been observed in Refs. [31, 32] for the O(2) symmetric AFM QCP, describing the onset of SDW order of XY type.

However, as a function of temperature, our DQMC result for χ^{-1} deviates from Hertz-Millis theory. Most importantly, the SDW susceptibility seems to scale quadratically with temperature instead of showing a linear T -dependence. Furthermore, note that in Hertz-Millis theory the correlation length exponent ν , Eq. (1.1), takes the mean-field value $\nu = 1/2$ such that the combined critical exponent νz appearing in the scaling of temporal correlations, Eq. (1.2), is unity, i.e. $\nu z = 1$. On the other hand, the fixed- ξ contours of the magnetic crossover in the DQMC phase diagram, Fig. 5.14, seem to be non-linear and bended, which hints towards $\nu \neq 1/2$ ⁹².

Finally, in comparison to the studies by Lee *et al.* [6, 7, 15], Sec. 1.4.4, we note that there is no evidence for a $z = 1$ dynamical scaling. The extracted critical boson propagator, Eq. (5.22), is furthermore manifestly isotropic in \mathbf{q} and therefore inconsistent with the highly anisotropic form in Eq. (1.33). Given

⁹²Here we identify $1/\tau = \Delta \sim T_{\text{crossover}}$ and conclude from $\nu z = 1$ that the crossover from the AFM SDW region to the quantum critical Landau damped regime is linear in Hertz-Millis theory.

the generic nature of our DQMC study – the considered Fermi surface, Fig. 5.2, does not appear to be special in any relevant manner – this strongly suggests that the strong-coupling fixed point by Lee *et al.* does not, in general, govern the physics of the AFM QCP at temperatures $T \geq T_{\text{DQMC}}^{\text{min}} = 1/40 = 0.025$. Of course, based on our DQMC study, we can not rule out the possibility that traces of $z = 1$ scaling appear at even lower temperatures.

5.3.2 Single-fermion correlations

Having discussed the scaling of antiferromagnetic correlations we now turn to the fermionic spectral properties of the spin-fermion model. Given that the extraction of real-time dynamics from DQMC simulations involves a form of numerical analytic continuation, which is generally an ill-defined procedure, we can not access the desired spectral function $A_{\mathbf{k}}(\omega)$ directly. Instead we utilize the following relation, [31, 208]

$$G_{\mathbf{k}}(\tau, 0) = \int_{-\infty}^{\infty} d\omega \frac{e^{-\omega(\tau-\beta/2)}}{2 \cosh \beta\omega/2} A_{\mathbf{k}}(\omega). \quad (5.23)$$

Here $G_{\mathbf{k}}(\tau, 0)$ is the momentum space Fourier transform of the time-displaced Green's function, which is accessible in DQMC as discussed in Secs. 2.5.2 and 4.1.3, and we focus our discussion on a single flavor of ψ_x fermions. Recognizing that for $\tau = \beta/2$ the fraction in Eq. (5.23) has a Gaussian-like structure as a function of frequency, i.e. it is symmetric and smoothly peaked at $\omega = 0$ with a width of the order of $1/\beta$, one realizes that $G_{\mathbf{k}}(\beta/2, 0)$ corresponds to the spectral function integrated over a frequency window of size T [31].

In the right panel of Fig. 5.20, we show the time-displaced Green's function $G_{\mathbf{k}}(\beta/2, 0)$ over the first Brillouin zone and across the antiferromagnetic transition. We have increased the momentum space resolution by conducting simulations at different twisted boundary conditions as explained in Sec. 4.3.2. In the disordered Fermi liquid phase, $r > r_c$, Fig. 5.20a, the Fermi surface is distinctly visible as a continuous line of the approximately constant maxima of $G_{\mathbf{k}}(\beta/2, 0)$. However, upon moving to the vicinity of the AFM QCP we notice a significant loss of spectral weight at the hot spots on the momentum space diagonal. Away from the latter the Fermi surface peak of $G_{\mathbf{k}}(\beta/2, 0)$ remains comparably sharp. Going further to the ordered side of the QCP, $r < r_c$, we observe a reconstruction of the Fermi surface in form of the emergence of pronounced gaps at the hot spots. This is in qualitative agreement with the mean-field analysis of Sec. 1.4.1.

Next, we study the fate of the fermions at the hot spots and near the QCP by investigating the imaginary part of the self-energy

$$\Sigma(\mathbf{k}, i\omega_n) = G_0^{-1}(\mathbf{k}, i\omega_n) - G^{-1}(\mathbf{k}, i\omega_n) \quad (5.24)$$

$$= i\omega_n - \epsilon_{\mathbf{k}} - \mu - G^{-1}(\mathbf{k}, i\omega_n). \quad (5.25)$$

Here G_0 and G are the non-interacting, Eq. (1.12), and full Green's functions, respectively. To compute $\Sigma(\mathbf{k}, i\omega_n)$ on the Fermi surface, we combine the data from simulations with different twisted boundary conditions and identify the

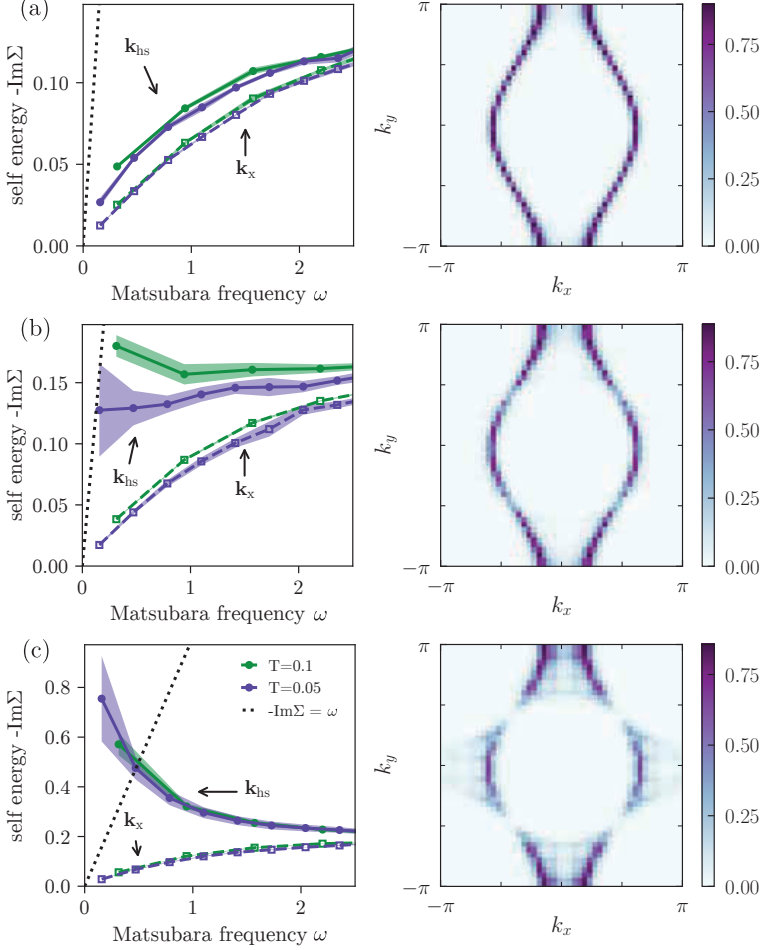


Figure 5.20: Imaginary part of the fermionic self-energy as a function of Matsubara frequency (left column) and time-displaced Green's function $G_{\mathbf{k}}(\beta/2, 0)$ (right column) for $r = -1.55 > r_c$ (a), $r = -1.95 \approx r_c$ (b), and $r = -1.95 < r_c$ (c). Shown are results for $\lambda = 1$, $L = 12$, and $\beta = 20$ (right panel). The dotted line indicates the non-Fermi-liquid crossover scale $-\text{Im}\Sigma(\omega_n) = \omega_n$. The solid lines represent the self-energy at the hotspots, $\mathbf{k} = \mathbf{k}_{\text{hs}}$, and dashed lines represent the self energy at the point of the Fermi surface with $k_y = 0$, namely $\Sigma(\mathbf{k} = \mathbf{k}_x = (k_F, 0))$. The data points are averaged over different twisted boundary conditions and the width of the line is the standard deviation of $\text{Im}\Sigma$ between the different boundary conditions, thus representing the uncertainty due to finite-size effects. Left column modified from Ref. [P1] under CC-BY 4.0.

local maxima of the $G_{\mathbf{k}}(\tau = \beta/2)$ at higher temperature $T = 0.2$ as the Fermi surface. The hot spots are then determined as the closest point to the intersection of the Fermi surfaces of the two bands ψ_x and ψ_y . Since finite-size effects are the dominating error source, we take the self-energy to be the average of the self-energies at the hotspot and the 4 adjacent momenta, corresponding to different boundary conditions, and take the standard deviation across the aforementioned momenta as the numerical error. The statistical error due to Monte Carlo sampling is entirely overshadowed by the latter.

As shown in the left panel of Fig. 5.20, we find that upon approaching the QCP the fermions at the hot spots lose their coherence as manifested in a substantial growth of the imaginary part of the self-energy. Remarkably, for $r \approx r_c$, Fig. 5.20b, the self-energy at the hotspots is very weakly dependent on Matsubara frequency and approximately constant. This implies a stark decrease of the quasiparticle residue⁹³

$$Z_{\mathbf{k}_{\text{hs}}}^{-1} = 1 - \left. \frac{\partial \Sigma(\mathbf{k}_{\text{hs}}, i\omega)}{\partial i\omega} \right|_{i\omega=0}, \quad (5.26)$$

Eq. (1.29), and hence marks a distinct deviation from ordinary Fermi liquid behavior in the vicinity of the AFM QCP. Note that the observed $\text{Im}\Sigma(\mathbf{k}_{\text{hs}}) \approx \text{const}$ stands in contrast to the $\sqrt{\omega}$ dependence suggested by one-loop perturbation theory, Sec. 1.4.3, but agrees with observations [31–33, 51] for the $O(N < 3)$ symmetric variants of the spin-fermion model, despite their distinct perturbative structure. We identify the lowest available temperature, $T = 0.05$, as the non-Fermi liquid crossover temperature, at which the dynamical contribution to the self-energy begins to dominate the bare frequency dependence of the Green's function, namely $\text{Im}\Sigma(\mathbf{k}, \omega_n = \pi T) = \pi T$. Considering fermions away from the hot spots, specifically at the point $\mathbf{k}_x = (k_F, 0)$ on the Fermi surface, we find that $-\text{Im}\Sigma(\mathbf{k}_x, i\omega_n)$ is steeply decreasing as a function of Matsubara frequency with an almost vanishing intercept. At small frequencies this decrease is approximately linear and indicates conventional Fermi liquid behavior as observed for large tuning parameter values on the entire Fermi surface, Fig. 5.20a. On the AFM SDW side of the QCP, $r < r_c$, Fig. 5.20c, the imaginary part of the self-energy tends to diverge, indicating the gapping out of the hot spots across the transition. Away from these gaps the self energy remains small, and tends to vanish at low frequencies.

5.4 The case of local nesting

Our findings above establish the phenomenology and the quantum critical properties of the AFM QCP in a metal with a generic Fermi surface, Fig. 5.2. At the considered temperatures, most of the observed characteristics are inconsistent with the recently proposed $z = 1$ fixed point by Lee *et al.* discussed in Sec. 1.4.4. Most notably, close to the QCP the fermions become incoherent at the hot spots and the AFM collective modes are characterized by an unambiguous dynamical

⁹³Formally, the quasiparticle residue is only well defined at $T = 0$. At finite temperatures the proxy $Z(T) = \omega_1 / \text{Im}G^{-1}(\omega_1)$ with $Z(T \rightarrow 0) \rightarrow Z$ may be considered instead [32].

scaling with $z = 2$. In the following, we address the question of whether the distinguished physics associated with the novel strong-coupling fixed point can be observed in a more fine tuned scenario.

5.4.1 Tuning to local nesting

One of the central features of the self-consistent theory by Lee *et al.* is the occurrence of dynamical nesting: In the low energy limit, the emergent control parameter v/c flows to zero such that patches of the Fermi surface associated with opposite hot spots become anti-parallel, Sec. 1.4.4. As an attempt to promote a flow to the strong coupling fixed point, we modify the band structure of the fermions in the spin-fermion model, Eq. (5.2), and consider a series of Fermi surfaces, which approaches the case of perfect local nesting. To design the latter, we include further range hopping terms into the energy dispersion,

$$\begin{aligned} \epsilon_{\mathbf{k},\alpha} = & -2 [t_{\alpha h} \cos k_x + t_{\alpha v} \cos k_y \\ & + t'_{\alpha d} \cos k_x \cos k_y \\ & + t''_{\alpha h} \cos^2 k_x + t''_{\alpha v} \cos^2 k_y] . \end{aligned} \quad (5.27)$$

Here, t , t' , and t'' are first, second, and third order nearest neighbor hopping amplitudes as illustrated in Fig. 5.1. Next, we evaluate the Fermi velocity $\mathbf{v}_\alpha = -\nabla_{\mathbf{k}} \epsilon_{\mathbf{k},\alpha}$ along the momentum diagonals $k_x = k_y = k_{\text{hs}}$. Retaining the relation $\epsilon_x(k_x, k_y) = -\epsilon_y(k_y, k_x)$ and the symmetry discussed in Sec. 5.1, we obtain

$$\begin{pmatrix} v_{x1} \\ v_{x2} \end{pmatrix} = 2 \sin(k_{\text{hs}}) \begin{pmatrix} t_{xh} + 2t''_{xh} \cos(k_{\text{hs}}) + t_{xd} \cos(k_{\text{hs}}) \\ t_{xv} + 2t''_{xv} \cos(k_{\text{hs}}) + t_{xd} \cos(k_{\text{hs}}) \end{pmatrix} = \begin{pmatrix} -v_{y2} \\ -v_{y1} \end{pmatrix} \quad (5.28)$$

where the integer index in $v_{\alpha i}$ enumerates the components of \mathbf{v}_α . The relative angle between the two Fermi velocities, Fig. 5.21, is given by

$$\sin(\theta) = -\frac{v_{x1}^2 - v_{x2}^2}{v_{x1}^2 + v_{x2}^2}, \quad (5.29)$$

which entails that nesting, $\theta = 0$, occurs for $v_{x1} = v_{x2}$ or, equivalently,

$$\cos(k_{\text{hs}}) = -\frac{1}{2} \frac{t_x - t_y}{t_{xx} - t_{yy}}. \quad (5.30)$$

Note that the next-nearest neighbor hopping amplitude t_{xd} drops out and, for simplicity, may be set to zero. Finally, solving for the chemical potential one finds the perfect nesting condition

$$\mu_x = -2 [(t_{xh} + t_{xv}) \cos(k_{\text{hs}}) + (t''_{xh} + t''_{xv}) \cos(k_{\text{hs}})^2] = -\mu_y. \quad (5.31)$$

For the particular parameter choice $t_{xh} = 1 = -t_{yv}$, $t''_{xv} = 0.45 = -t''_{yh}$ this implies that local nesting in the vicinity of the hot spots occurs precisely for a chemical potential of $\mu_0 = -0.77$. As illustrated in Fig. 5.21, we may thus generate a series of Fermi surfaces characterized by a decreasing nesting angle $\theta \in [\pi/2, 0]$ by defining the deviation $\tilde{\mu} = \mu - \mu_0$ and considering $\mu \in \{1.5, 1.2, 0.9, 0.6, 0.3, 0.0\}$. Note that for these values, $\tilde{\mu}$ is approximately proportional to the sine of the nesting angle, i.e. $\tilde{\mu} \sim \sin \theta$, and the magnitude of the Fermi velocity does not change significantly.

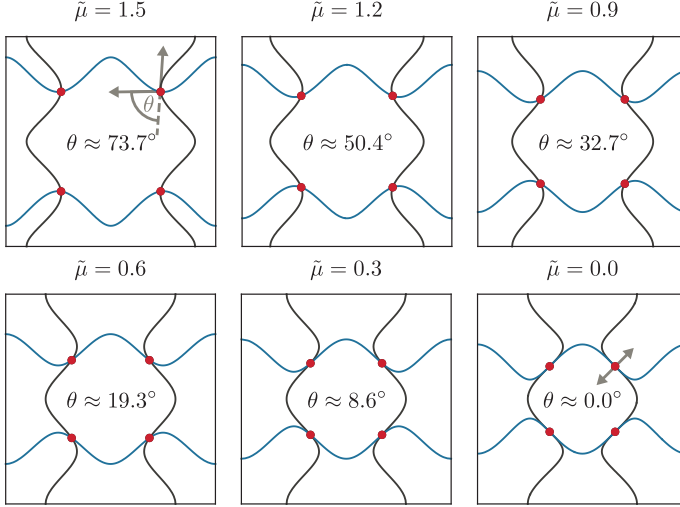


Figure 5.21: Series of Fermi surfaces obtained by tuning the chemical potential $\tilde{\mu} = \mu - \mu_0$ with $\mu_0 = -0.77$. For $\tilde{\mu} \rightarrow 0$ the ψ_x (dark grey) and ψ_y (blue) bands become locally nested in the vicinity of the hot spots (red points), i.e. $\theta \approx 0$. The underlying energy dispersion is given in Eq. (5.27) and the hopping amplitudes are as described below Eq. (5.31).

5.4.2 Shift of the AFM QCP

Performing DQMC simulations of the spin-fermion model for each of the Fermi surfaces in Fig. 5.21 and a Yukawa coupling of $\lambda = 1.0$, we investigate the dependence of magnetic correlations on the chemical potential $\tilde{\mu}$ and, thus, the nesting angle θ . In particular, we compute contour cuts of the AFM susceptibility, Eq. (5.4), as defined by the condition $\chi^{-1} = 0.01^{94}$. As shown in Fig. 5.22, these magnetic contours clearly bend to larger tuning parameter values upon decreasing the chemical potential $\tilde{\mu}$. This trend is especially pronounced at lower temperatures: Comparing the extremal cases $\tilde{\mu} = 1.5$ and $\tilde{\mu} = 0.0$, we observe a displacement $\Delta \approx 0.15$ for $T = 0.2$ and $\Delta \approx 0.38$ for $T = 0.05$. These findings indicate that upon approaching local nesting at the hot spots, $\theta \rightarrow 0$, the AFM QCP shifts to larger tuning parameter values. Note that this movement is consistent with the expectation of an effectively enhanced Yukawa coupling because of the increased low-energy scattering phase space.

⁹⁴In practice, we compute $\chi^{-1}(r)$ by linearly interpolating between the DQMC data points $\chi^{-1}(r_i)$ for a discrete set of values r_i and, if possible, determine the root of $f(r) = \chi^{-1}(r) - 0.01$.




Figure 5.22: Fixed contour cuts of the AFM susceptibility ($\chi^{-1} = 0.01$). Shown is the shift to larger tuning parameter values upon approaching perfect local nesting, $\tilde{\mu} \rightarrow 0$. Although in good approximation $\sin \theta$ (right axis) is proportional to $\tilde{\mu}$, the right axis should be seen as a specification of the nesting angle for the shown data points only. The considered system size is $L = 10$ and the Yukawa coupling is $\lambda = 1.0$.

5.4.3 Quantum criticality

We now turn to the low-energy properties of the spin-fermion model in the case of almost local nesting near the hot spots. Working close to but slightly away from the perfect nesting limit, $\tilde{\mu} = 0$, we focus on the chemical potential $\tilde{\mu} = 0.3$, i.e. we set $\mu_x = -0.47$ and $\mu_y = -\mu_x = 0.47$ for the ψ_x and ψ_y fermions, respectively. Note that compared to the generic Fermi surface considered in Sec. 5.2, which has $\theta \approx 36.9^\circ$, the relative angle between the Fermi velocities for these parameters is $\theta \approx 8.6^\circ$, Fig. 5.21, and is decreased by about a factor of four. To enhance the Matsubara frequency resolution, we perform simulations at the ultra-low temperature $T = 1/100$. Recalling the optimal $O(\beta N^3)$ scaling of the DQMC method, Sec. 4.2.1, this implies that in comparison to the studies above, where we have considered temperatures as low as $\beta = 40$, the computational cost is increased by a factor of 2.5. As for the remaining model parameters, we set the quartic boson coupling to zero, i.e. $u = 0$, and the Yukawa coupling to $\lambda = 1.0$.

We start off by discussing a potential superconducting transition. Numerically, by investigating the superfluid density, we find no indication of superconductivity down to the lowest temperature $T \geq 0.01$. This is qualitatively consistent with the prediction that in an almost locally nested system antiferromagnetic fluctuations will only slightly enhance T_c [4, 7, 112]. From a weak-coupling viewpoint, Eq. (1.30), the energy scale associated with the onset of superconductivity is expected to decrease as a function of the angle, $T_c \sim \sin \theta$. This relation has

been numerically confirmed for a generic Fermi surface coupled to an $O(2)$ order parameter in Refs. [31, 112]. Anticipating that superconductivity, arising from fermions at the hotspots, is independent of the precise band structure [112], one expects about a fourfold reduction of T_c for the almost locally nested Fermi surface compared to the one considered in the studies of the previous sections, Fig. 5.2.

Next, we turn to the critical AFM fluctuations in the vicinity of the QCP. Note that we estimate the location of the latter by extrapolating the tuning parameter dependence of the SDW susceptibility down to $T = 0$, see App. C. In Fig. 5.23, we show the Matsubara frequency dependence of the inverse SDW susceptibility $\chi^{-1}(\mathbf{q} = 0, i\omega_n)$. In contrast to the case of a generic Fermi surface, we notice a distinct curvature over the interval $0 < \omega_n < 1$, which is more pronounced at smaller Matsubara frequencies – the focus on ultra-low temperatures improves the resolution of this nonlinearity. For $\omega_n > 0$, data points obtained from DQMC simulations of differently sized systems fall on top of each other, indicating the absence of significant finite size effects.

Although this curvature might, in principle, stem from enhanced low-energy scattering in the vicinity of the hot spots, and, figuratively speaking, might be caused by relevant higher-order Feynman diagrams, it is instructive to compare our DQMC result to the non-interacting fermionic susceptibility Π_0 , Fig. 1.11a. Explicitly, the latter evaluates to the Lindhard form [17, 23]

$$\Pi_0(\mathbf{q}, i\omega_n) = \frac{2}{\beta N_\phi N} \sum_{\mathbf{k}\alpha} G_\alpha(\mathbf{k}) G_{-\alpha}(\mathbf{k} + \mathbf{q}) \quad (5.32)$$

$$= \frac{2}{N_\phi N} \sum_{\mathbf{k}\alpha} \frac{n_\alpha(\mathbf{k}) - n_{-\alpha}(\mathbf{k} + \mathbf{q})}{i\omega_n - \epsilon_{\alpha,\mathbf{k}} + \epsilon_{-\alpha,\mathbf{k}+\mathbf{q}}}, \quad (5.33)$$

where $n_\alpha(\mathbf{k}) = [\exp(\beta(\epsilon_\alpha(\mathbf{k}) - \mu_\alpha)) + 1]^{-1}$ is the Fermi distribution and the factor of 2 is due to spin. As illustrated in Fig. 5.23, the fermionic susceptibility $\Pi_0(\mathbf{q} = 0, i\omega_n)$ qualitatively shows the same trend as the SDW correlations $\chi^{-1}(\omega_n)$ for small Matsubara frequencies. In particular, noting that the pronounced curvature of the non-interacting fermionic susceptibility develops upon tuning to local nesting, Fig. 5.24, this strongly suggests that the observed non-linearity in the AFM correlations is primarily due to low-energy features of the band structure.

Finally, we show the dependence of the inverse SDW susceptibility on squared momentum in Fig. 5.25. For finite $\mathbf{q} > 0$ the DQMC results for $\chi^{-1}(\mathbf{q}, i\omega_n = 0)$ are consistent with a q^2 term, while a noticeable drop is visible at $\mathbf{q} = 0$. We perform a linear regression to establish the quadratic momentum dependence, illustrated in Fig. 5.25, and find good agreement with the numerical data over the range $0 < q^2 \leq 2$ (excluding $\mathbf{q} = 0$). Importantly, all finite- q data points collapse onto a single line and do not branch out as a function of squared momentum. We note that the momentum dependence is similar at higher temperature ($T = 0.025$, see inset to Fig. 5.25). The isotropic behavior of the boson propagator lies in stark contrast to an anisotropic $|q_x + q_y| + |q_x - q_y|$ term predicted at the novel $z = 1$ fixed point [7], Eq. (1.33). However, the apparent discontinuity at $q = 0$ hints towards features at low momenta that can not be resolved with the





Figure 5.23: Frequency dependence of the inverse AFM susceptibility close to the quantum critical point ($r = 1.625$) at ultra low temperature $T = 1/100$ for $\lambda = 1.0$ and the almost locally nested Fermi surface with $\tilde{\mu} = 0.3$, see Fig. 5.21. The solid line is the non-interacting fermionic susceptibility Π_0 . Resized from Ref. [P1] under CC-BY 4.0.

Figure 5.24:

Evolution of the Matsubara frequency dependence of the non-interacting fermionic susceptibility $\Pi_0(\mathbf{q} = 0, i\omega_n)$, Eq. (5.33), upon approaching local nesting near the hot spots, $\tilde{\mu} \rightarrow 0$. The values of the chemical potential $\tilde{\mu}$ correspond to the series of Fermi surfaces shown in Fig. 5.21 (lines in different grey scales).






Figure 5.25: Momentum dependence of the inverse AFM susceptibility close to the quantum critical point ($r = 1.625$) at ultra low temperature $T = 1/100$ ($T = 1/40$ in the inset) for $\lambda = 1.0$ and the almost locally nested Fermi surface with $\bar{\mu} = 0.3$, see Fig.5.21. The solid lines are linear fits. Resized from Ref. [P1] under CC-BY 4.0.

available system sizes. Hence, an anisotropy, as appearing in the critical theory by Lee *et al.*, may emerge at smaller momenta, i.e. $q < 2\pi/14$.

5.5 First-order quantum phase transition

In the DQMC studies presented above, we have considered the spin-fermion model, Eq. (5.1), to investigate the quantum critical properties associated with the onset of antiferromagnetism in metals. In this context, it is important to remark that it is always a possibility that the SDW quantum phase transition is of first order rather than continuous character. Specifically, considering the generic Fermi surface of Fig. 5.2, we find that setting the bare quartic boson coupling to zero, i.e. $u = 0$, leads to a region of coexistence of magnetic and non-magnetic states near the phase transition point. As shown in Fig. 5.26, this is indicated by a jump-like tuning parameter dependence of the magnetic susceptibility $\chi^{-1}(\mathbf{q} = 0, i\omega_n = 0; r)$. In particular, $\chi^{-1}(r)$ is continuous at temperatures $T \geq 0.1$ but very visibly becomes discontinuous upon cooling the system down to lower temperatures. To verify that this feature is indeed due to the coexistence of “phases” of similar energy, we inspect the (unnormalized) distribution function of the AFM susceptibility, which may be extracted from DQMC by binning the corresponding Markov chain values $\{\chi_j\}$. In Fig. 5.27, we show the obtained histogram, which clearly displays a double-peak structure. Here, the first peak at $\chi = 0$ corresponds to the disordered Fermi liquid phase whereas the one centered around $\chi \approx 2600$ indicates a state with finite AFM SDW correlations.




Figure 5.26: First order transition in the spin-fermion model occurring for the Fermi surface of Fig. 5.21 at low temperatures and for $u = 0$. Shown is the tuning parameter dependence of the inverse magnetic susceptibility χ^{-1} close to the quantum critical point for various temperatures. Resized from Ref. [P1] under CC-BY 4.0.




Figure 5.27: Distribution of the values of the antiferromagnetic susceptibility χ obtained along the DQMC Markov chain in the vicinity of the AFM QCP. The Fermi surface is as shown in Fig. 5.2 and the parameters are $r \approx 1.42$, $L = 10$, and $\beta = 40$.

While it is understood from Landau theory, that a vanishing quartic term favors the occurrence of a first order phase transition, we find that this is not the case for the almost locally nested system of Sec. 5.4. Here, despite the fact that we set $u = 0$ and that one might argue on general ground that nesting favors a first-order transition [209], our numerics shows no indication of a jump of the antiferromagnetic susceptibility (see Fig. C.2 in App. C). Instead the onset of SDW correlations is continuous down to temperatures as low as $T = 0.025$.

5.6 Discussion

In this chapter, we have deployed large-scale DQMC simulations to explore the phase diagram of itinerant fermions coupled to an isotropic AFM order parameter in an unbiased and rigorous manner. By inspecting fermionic and bosonic correlations, we were able to unveil the critical low-energy behavior associated with the $O(3)$ symmetric metallic AFM QCP up to numerical accuracy. Our main finding is that, over a broad range of temperatures, the critical SDW fluctuations are remarkably well described by Hertz-Millis theory. The dynamical exponent in this regime, within our accuracy, is $z = 2$. This is surprising in view of the fact that Hertz-Millis theory neglects infinitely many marginal couplings that are generated when integrating over the gapless fermion degrees of freedom at the Fermi surface, Sec. 1.4.2. Qualitatively similar observations have been made for metals with easy-plane XY [31, 32] and Z_2 [33] antiferromagnetic order, despite noticeable differences in the perturbative structure, which hints towards a generic property of the AFM QCP. A further important aspect of our

DQMC results is that, at finite-temperatures above the QCP, the fermions lose their coherence as a result of strong scattering. In particular, we find that the imaginary part of the self-energy at the hot spots is only weakly dependent on Matsubara frequency, in stark contrast to the behavior expected within Fermi liquid theory. On the ordered side of the QCP the self-energy seems to be diverging, which is consistent with a Fermi surface reconstruction and a gapping out of the hot spots.

Motivated by the recent prediction of a novel $z = 1$ fixed point [6, 7, 15] by Lee and collaborators, we considered a variant of the spin-fermion model where the Fermi surface in the vicinity of the hot spots is (almost) locally nested. In this case, the QCP is shifted to higher tuning parameter values and the superconducting T_c is strongly suppressed below the lowest temperature available in our study, $T \approx E_F/200$. Above this temperature, we find substantial deviations from the $z = 2$ Hertz-Millis behavior, but no evidence for the predicted $z = 1$ criticality. Notably, the momentum dependence of the SDW order parameter is isotropic up to numerical accuracy.

Considering the general phase diagram of the spin-fermion model, we have shown that the onset of collinear SDW order at $T = 0$ gives rise to a magnetic crossover at finite-temperatures and, most remarkably, is associated with the emergence of an extended phase of high-temperature superconductivity. The corresponding superconducting order parameter is unambiguously of d-wave character, i.e. it changes sign under $\pi/2$ -rotations. Our DQMC results reveal that the maximal critical temperature $T_c^{\max} \approx E_F/20$ (for $\lambda = 2$) occurs close to the QCP, i.e. for $r \approx r_c$, which is indicative of the intertwined character of magnetic and electronic ordering. Note that while we have detected finite CDW fluctuations, they are (partially) competing with superconductivity and seem to play a minor role. Overall, our findings bear close resemblance to the essential features of the phase diagrams of many unconventional superconductors, such as the electron-doped cuprates and the iron-pnictides. This suggests that metallic quantum criticality plays an important role in these materials and, in particular, provides numerical evidence for the hypothesis that critical AFM fluctuations are responsible for the unusually high T_c .

The presented DQMC work offers many starting points for follow-up studies. In particular, further investigations of potential manifestations of the $z = 1$ fixed point, for example by considering even lower temperatures or different band structures, are highly desirable to connect numerical and analytical results. Moreover, directly extracting the renormalized three-point vertex function from DQMC simulations and analyzing its Matsubara frequency dependence seems to be a viable approach to shed important light onto why Hertz-Millis theory and Eliashberg theory are capable of capturing large parts of the critical physics correctly, despite being formally uncontrolled. In view of the fact that besides AFM correlations and superconductivity a plethora of subtle phenomena has been observed in unconventional superconductors, it would also be interesting to consider variants of the spin-fermion model with multiple competing order parameters. In particular, a finite-temperature DQMC study of the interplay of AFM and nematic fluctuations could significantly improve our understanding of the iron-pnictide superconductors – an extension of the preliminary work in

Ref. [30], which has considered the strict $T = 0$ limit. Furthermore, it would be worthwhile to examine the non-Fermi liquid state above the QCP in more detail. Especially, verifying a linear-in- T resistivity over an extended region of the phase diagram, as observed in experiments, would be highly desirable. Although analytic continuation generally prevents the analysis of real-time dynamics, a promising memory matrix technique to partially circumvent this issue has recently been put forward in Ref. [210]. In the following chapters, we take an orthogonal approach and demonstrate that machine learning can be utilized to, among other aspects, reveal extended non-Fermi liquid regimes in the vicinity of metallic quantum critical points.

6 | Machine learning phases of matter

The past decade has seen a phenomenal triumph of machine learning (ML) techniques in almost all areas of industry. Considered by many to be as disruptive as the computer revolution of the last century [211], the rise of self-learning algorithms has, among others, revolutionized manufacturing, stirred up the service sector, and severely impacted financial trading. The strength of machine learning methods lies in recognizing patterns of interest in large amounts of complex data which may then be utilized in classification or decision making tasks. Famous applications include image recognition, email spam filtering, as well as autonomous driving.

In a seminal interdisciplinary work, Carrasquilla *et al.* [212] have, for the first time, transferred those novel ML ideas to condensed matter physics. In particular, they have demonstrated that the paramagnet-ferromagnet phase transition in the two-dimensional Ising model can be detected by combining classical Monte Carlo simulations with a purely statistical machine learning approach that, compared to a traditional finite-size scaling analysis, has no prior knowledge of the order parameter. Triggered by this finding, three conceptual lines of research have developed within condensed matter physics: First, the utilization of ML architectures, such as restricted Boltzmann machines (RBMs), as wave function parametrizations in variational algorithms to obtain ground states of quantum many-body systems [213–218]. Second, the embedding of ML models into quantum Monte Carlo simulations to reduce autocorrelation times [27, 219–223]. Third, the discrimination of phases of matter by employing artificial neural network (ANN) structures [P3–P5, 212, 224–228].

In this thesis, we focus on the latter and consider a novel ML approach to probe aspects of quantum criticality in itinerant electron systems. To prepare the corresponding discussion in Ch. 7, it is the purpose of the following sections to provide a concise introduction into the essentials of machine learning. More comprehensive treatises on the topic can be found in Refs. [211, 229–231]. Starting with a general review of neural networks and their building blocks, we discuss the concept of supervised training and, especially, the efficient computation of gradients. Afterwards, we introduce convolutional neural networks (CNNs), which have proven to be particularly powerful in classification tasks. Finally, we apply the ML concepts and showcase that they can be utilized to detect the antiferromagnetic finite-temperature BKT transition in the $O(2)$ symmetric spin-fermion model.




Figure 6.1: Visualization of a generalized perceptron that takes an input vector $\mathbf{x} = (x_1, x_2, \dots, x_n)^T$ and implements the function $f: \mathbb{R}^n \rightarrow \mathbb{R}$, $f(\mathbf{x}) = \sigma(\boldsymbol{\omega} \cdot \mathbf{x} + b)$, where $\boldsymbol{\omega}$ is a vector of weights, b is a bias, and $\sigma: \mathbb{R} \rightarrow \mathbb{R}$ is a non-linear activation function. For a regular perceptron, $\sigma = \Theta$ is the Heaviside step function.

6.1 Artificial neural networks

Historically, artificial neural networks, or simply “neural networks” (NNs), have been introduced as models for information processing in biological systems [232–234]. In the last century, their development was largely based on the sentiment that, essentially, the human brain is a highly complex, parallel computer that is capable of organizing its structural constituents, the neurons, to perform sophisticated tasks, such as perceptual recognition [231]. The modern field of machine learning, however, has detached from this biological realism [229] and, while the terminology is still significantly influenced by the original context, merely identifies neural networks as particularly efficient statistical models.

The fundamental building block of most artificial neural networks are perceptron neurons as illustrated in Fig. 6.1. Based on a given vector input $\mathbf{x} \in \mathbb{R}^n$, a perceptron produces a binary output $f(\mathbf{x}) \in [0, 1]$ by computing a weighted sum, adding a constant, and applying a Heaviside step function, i.e.

$$f(\mathbf{x}) = \Theta(\boldsymbol{\omega} \cdot \mathbf{x} + b). \quad (6.1)$$

Here, the vector of weights $\boldsymbol{\omega}$ and the “bias” b parameterize the perceptron and may be tuned such that the latter implements a meaningful operation – we will discuss this training process in more detail in the next section. Given the binary structure of the step function, one realizes that a perceptron can be utilized to implement a classification algorithm that divides a set of inputs into two categories. In particular, by noting that the argument of the Heaviside function is a regular linear model, as used in conventional linear regression, it can be proven that a perceptron is, in general, capable of correctly classifying data that is *linearly separable*⁹⁵ [231].

In practical applications, it is useful to substitute the Heaviside function as it is discontinuous, not derivable at $x = 0$, and has $\Theta'(x) = 0$ everywhere else – features that turn out to be suboptimal for training [235]. Concretely, the

⁹⁵Figuratively speaking, linear separability means that there exists a hyperplane – a decision surface – in the potentially high-dimensional data space such that data points corresponding to the two different classes lie on opposite sides of the plane.




Figure 6.2: Illustration of a maximally connected, feed-forward neural network. An input vector \mathbf{x} (green circles) is propagated through two layers of “hidden” neurons (red circles) before converging into one, or multiple, output neurons (blue circle). Effectively, the ANN implements a highly non-linear function $F : \mathbb{R}^n \rightarrow \mathbb{R}$, which is parameterized by biases, sitting at the nodes, and weights, attached to inter-neuron connections (black arrows).

functional form of a neuron is generalized to

$$f(\mathbf{x}) = \sigma(\boldsymbol{\omega} \cdot \mathbf{x} + b), \quad (6.2)$$

where σ is a continuous activation function, such as a sigmoid,

$$\sigma(x) = \frac{1}{1 + e^{-x}}, \quad (6.3)$$

i.e. a smooth step function, or a rectified linear unit (ReLU), $\sigma(x) = \max(0, x)$.


In an artificial neural network, many neurons are arranged in a layered structure as illustrated in Fig. 6.2. An input vector $\mathbf{x} \in \mathbb{R}^n$ is then processed by propagating information through multiple layers of hidden⁹⁶ neurons before converging to a neural output. Note that while different layers are connected in series, there are no links between neurons within a layer such that information can be processed in parallel. On a mathematical level, an ANN represents a function $F(x) : \mathbb{R}^n \rightarrow \mathbb{R}$, which takes the form of a systematic composition of multiple perceptron operations $f(x)$, Eq. (6.2). Furthermore, because involving nested nonlinear activation functions, $F(x)$ is highly nonlinear. It is precisely this property that enables complex pattern recognition and classification beyond linear separability.

We remark that the network in Fig. 6.2 is a maximally connected feed-forward network, in which the neurons of two layers are all linked to each other and the flow of information is uni-directional (from left to right in the figure). In contrast, recurrent neural networks (RNNs), for example, contain backward connections, i.e. loops, and, in this way, implement memory features [229, 234]. In condensed matter physics, RNNs have been utilized to learn phase transitions from dynamics [236].

⁹⁶One categorizes neurons into “visible” input and output neurons as well as “hidden” neurons sitting in intermediate layers.

Figure 6.3:

Schematic effective loss landscape as defined by the cost function of a neural network. Starting from a given point, corresponding to the initial values of the weights and biases, the gradient descent algorithm amounts to moving “downhill”, i.e. in the direction opposite to the local gradient. Taken from Ref. [237] with permission.



6.2 Supervised learning

Artificial neural networks represent nonlinear statistical models that may be utilized in regression and classification tasks. Fundamentally, they are universal function approximators [229] and, if designed appropriately, can be shown to be Turing complete [238]. Depending on the size of a network, the associated function $F(x)$ is characterized by a large number of free parameters: the weights and biases associated with individual neurons. In order to have a NN perform a particular task, these parameters must be optimized with respect to a given objective.

In the following, we focus on supervised learning, where the network is trained on input data $X = \{x_1, x_2, \dots, x_n\}$ that comes associated with desired outputs $Y = \{y_1, y_2, \dots, y_n\}$. In this case, the objective is to learn the relation $Y(X)$ and correctly reproduce the y_i for the training data. Quantitatively, we introduce a cost function C as a measure for the performance of the neural network. Popular choices include the quadratic form,

$$C = \frac{1}{2n} \sum_i [y_i - F(x_i)]^2, \quad (6.4)$$

and the cross-entropy [230]

$$C = -\frac{1}{n} \sum_i [y_i \ln F(x_i) + (1 - y_i) \ln (1 - F(x_i))]. \quad (6.5)$$

Note that the latter is related to the Kullback-Leibler divergence [229], and, in this sense, quantifies how much the distributions $Y(X)$ and $F(X)$ differ from each other. Implicitly depending on the network parameters, the cost function effectively defines a loss landscape [239], Fig. 6.3, whose global minimum corresponds to the optimal weight and bias configuration. The goal of the training process is therefore to find this minimum as efficiently as possible. Given the nested structure of the neural function $F(x)$ and the extensiveness of the number of free parameters, an analytic optimization of the weights and biases is rarely

feasible. Instead, one applies established numerical gradient descent methods [124, 229]. As indicated in Fig. 6.3, the fundamental strategy is to start from a random initial configuration and to iteratively approach the minimum by moving in the opposite direction of local gradients. Hence, in each step, the weights and biases of a neuron j in a layer l are updated according to

$$\omega_{jk}^l = \omega_{jk}^l - \eta \frac{\partial C}{\partial \omega_{jk}^l}, \quad (6.6)$$

$$b_j^l = b_j^l - \eta \frac{\partial C}{\partial b_j^l}, \quad (6.7)$$

in which η is a learning rate (step size). Although the iteration itself is straightforward to implement, the demanding part of the training is the computation of the – large number of – partial derivatives in Eqs. (6.6) and (6.7). Below, we will discuss the instructive backpropagation algorithm [229, 230], which addresses this tasks efficiently.

Before moving on, let us note in passing that numerical optimization of, potentially non-convex, objective functions is an entire separate research field by itself. In this thesis, we only focus on the basic concepts and refer the interested reader to the vast body of literature [124, 240].

6.2.1 Backpropagation

The partial derivatives in Eqs. (6.6) and (6.7) indicate how the value of the cost function is affected by a change of a weight ω_{jk}^l , connecting neurons j and k in layers l and $l - 1$, and a bias b_j^l . The idea of backpropagation is to evaluate these derivatives for all weights and biases in a layer by layer manner starting with the output layer of the neural network first. The reason for this reverse traversal is simple: The output neurons in the last layer, $l = L$, appear directly in the cost function C , see $F(x_i)$ in Eqs. (6.4) and (6.5) for a single output neuron, whereas hidden neurons only implicitly influence the neural output and the gradient. By repeated application of the chain-rule, one may nonetheless derive explicit relations between derivatives of subsequent layers and compute the network gradient iteratively.

To see this, we start by introducing the compact notation $a_j^l = \sigma(z_j^l)$, in which a_j^l is the output and z_j^l the weighted input of a neuron j in layer l [230]. Focusing on the bias derivative first, we evaluate

$$\begin{aligned} \frac{\partial C}{\partial b_j^l} &= \frac{\partial C}{\partial a_j^l} \frac{\partial a_j^l}{\partial b_j^l} \\ &= \frac{\partial C}{\partial a_j^l} \sigma'(z_j^l) \frac{\partial z_j^l}{\partial b_j^l} \\ &= \delta_j^l, \end{aligned} \quad (6.8)$$

where we have used that $\frac{\partial z_j^l}{\partial b_j^l} = 1$ and have introduced

$$\begin{aligned}\delta_j^l &= \frac{\partial C}{\partial z_j^l} \\ &= \frac{\partial C}{\partial a_j^l} \frac{\partial a_j^l}{\partial z_j^l} \\ &= \frac{\partial C}{\partial a_j^l} \sigma'(z_j^l).\end{aligned}\tag{6.9}$$

Note that δ_j^l may be interpreted as the “error” associated with the j -th neuron in layer l in the sense that a value close to zero indicates optimality [230]. Performing a similar evaluation of the derivative in Eq. (6.6), we find

$$\begin{aligned}\frac{\partial C}{\partial \omega_{jk}^l} &= \frac{\partial C}{\partial a_j^l} \frac{\partial a_j^l}{\partial \omega_{jk}^l}, \\ &= \frac{\partial C}{\partial a_j^l} \sigma'(z_j^l) \frac{\partial z_j^l}{\partial \omega_{jk}^l} \\ &= \delta_j^l a_k^{l-1},\end{aligned}\tag{6.10}$$

which, in contrast to Eq. (6.8), indicates a dependence on the neural output of the preceding layer.

Having expressed both derivatives in terms of the neural error δ_j^l , we now derive a relation between the errors δ_j^l and δ_j^{l+1} of subsequent layers. Recalling that $z_k^{l+1} = \sum_j \omega_{kj}^{l+1} \sigma(z_j^l) + b_k^{l+1}$, we obtain [230]

$$\begin{aligned}\delta_j^l &= \frac{\partial C}{\partial z_j^l} \\ &= \sum_k \frac{\partial C}{\partial z_k^{l+1}} \frac{\partial z_k^{l+1}}{\partial z_j^l} \\ &= \sum_k \frac{\partial z_k^{l+1}}{\partial z_j^l} \delta_k^{l+1} \\ &= \sum_k \omega_{kj}^{l+1} \sigma'(z_j^l) \delta_k^{l+1}.\end{aligned}\tag{6.11}$$

This is the central equation of the backpropagation algorithm. It represents an explicit prescription for how to propagate the neural errors δ_j^{l+1} of layer $l+1$ backwards to layer l . Combined with the fact that δ_j^L in the final layer L is straightforward to compute – the cost function directly depends on neural outputs $a_j^L = F_j(x)$ – one may repeatedly apply Eq. (6.11) as well as Eqs. (6.8) and (6.10) to compute all partial derivatives and, thus, the entire network gradient. In summary, backpropagation thus consists of two fundamental steps:

- Forward pass: Propagate an input through the network to obtain the neural outputs and evaluate the cost function.

- Backwards pass: Compute the errors δ_j^L and iteratively backpropagate them via Eq. (6.11) from the output to the input layer. Use Eqs. (6.8) and (6.10) to obtain the network gradient (and update the weights and biases as per Eqs. (6.6) and (6.7)).

In principle, to obtain a good numerical estimate, the procedure above should be applied to the entire training set X : The gradient then corresponds to the average over the individual samples. However, since this calculation is the computational bottleneck of the training, one aims to reduce the number of backpropagation invocations as much as possible. For this reason, in practice, one implements a stochastic version of gradient descent in which the network gradient is approximated based on small subsets of the training data, so-called mini batches. By choosing the size of the latter, one can trade numerical accuracy for computational performance.

Let us note that while the described backpropagation algorithm is characterized by conceptual simplicity and good efficiency, it requires the storing of the outputs of all neurons for the backward pass and only produces the first-order gradient. It is therefore not compatible with more sophisticated optimizers that utilize higher-order gradients in the updating of weights and biases, Eqs. (6.6) and (6.7). For these reasons, modern machine learning frameworks are built around more general and efficient automatic differentiation (AD) methods [241, 242]. In the latter, one typically constructs an adjoint computational graph for the backward pass to speed up gradient calculations. Based on these techniques, *differential programming* has recently emerged as a powerful generalization of machine learning that considers arbitrary code, beyond a network representation, as a derivable statistical model [242].

6.2.2 Regularization and dropout layers

In comparison to conventional statistical models, e.g. linear regression models, artificial neural networks host a much larger number of free parameters – quite often $O(10000)$. While this intuitively explains the superior representation capabilities of ANNs, it may also lead to the practical issue of overfitting. In this case, the network over specializes on the presented training data, i.e. it reproduces the desired output values y_i almost perfectly, but fails entirely when applied to an unknown test input after training.

As a counter-measure to overfitting, a dropout unit [243] can be integrated into the network architecture between two neural layers, Fig. 6.4. During the optimization, this component randomly disables a certain percentage – typically about 50% – of the inter-layer connections for every training input or mini batch⁹⁷. This randomization prevents that individual neurons learn and overly represent the same feature in the data. Effectively, one is averaging over different thinned out networks that may overfit in different ways [230].

A second, more direct method to prevent over specialization is L2 regularization. Here, instead of modifying the network structure one adjusts the cost function.

⁹⁷Alternatively, one may disable entire neurons [243].




Figure 6.4: Illustration of the effect of a dropout unit (grey rectangle) in a neural network. During the training, a certain percentage of inter-neuron connections (black arrows) is randomly selected and disregarded (dashed grey arrows). Once the network parameters are fixed, the dropout unit acts as the identity.

Specifically, one introduces a term proportional to the squared norm of the weights. For the cross-entropy cost function, Eq. (6.5), this reads [230]

$$C = -\frac{1}{n} \sum_i [y_i \ln F(x_i) + (1 - y_i) \ln (1 - F(x_i))] + \frac{\lambda}{2n} \sum_{\omega} \omega^2, \quad (6.12)$$

where λ is a regularization factor. Note that the new term comes with a positive sign and thus has the effect of penalizing large weights. This way, one keeps the magnitudes of the latter on a similar level and prevents the formation of singular network paths that dominate the neural output.

6.3 Convolutional neural networks

Although the basic ANNs discussed above, and illustrated in Fig. 6.2, represent powerful models that can learn a wide variety of patterns in data, it has proven to be useful to design more specialized machine learning architectures for particular tasks. For example, correlations in the input data are often known to be primarily short-ranged, in which case it is meaningful to equip the neural network with a notion of locality by embedding components beyond the maximal-connection principle. To that end, we introduce convolutional layers, which, as the name suggests, implement convolution operations. Specifically, the output of a neuron in such a layer is given by [230]

$$a_{j,k}^l = \sigma \left(\sum_{l=-K}^K \sum_{m=-K}^K \omega_{l,m}^l a_{j+l,k+m}^{l-1} + b^l \right). \quad (6.13)$$

Here, as illustrated in Fig. 6.5, a subset of the input vector a^{l-1} , which is interpreted as a two-dimensional matrix $a_{j,k}^{l-1}$, is multiplied by a kernel of size $(2K + 1) \times (2K + 1)$, corresponding to the shared weights $\omega_{l,m}^l$ of the layer. Effectively, the kernel acts as a two-dimensional filter, which is moved across the




Figure 6.5: Illustration of a convolution operation. A subregion of an input image (green) is multiplied element by element with a kernel (red) to produce an output value (blue), corresponding to the sum of the intermediate results.

entire input to produce an output, commonly referred to as *feature map*. The power of convolutions is well established in image processing, where by choosing Sobel, Laplace, or Gaussian convolution matrices, one can readily implement edge detection, sharpening, and blurring operations, respectively [244]. In the present context, the neural network learns and optimizes the kernel autonomously during the training process. We hence anticipate that convolutional layers allow for efficient feature extraction, independent of the precise type and origin of the input data.

In practice, one typically considers a stack of convolutional layers whose output is then fed into a maximally-connected ANN. Such a *deep*⁹⁸ convolutional neural network (CNN) architecture is visualized in Fig. 6.6. Here, max pooling layers have been inserted between the convolutional units. The former divide an input “image” into quadratic pools, i.e. subregions, and implement a dimensional reduction by replacing every individual pool by its maximum value. This down-sampling has at least two effects: On the one hand, it reduces the convolutional feature maps to their most prominent characteristics and, in this way, prevents overfitting. On the other hand, it represents an (irreversible) compression which lowers the computational cost associated with further data processing.

Convolutional neural networks have applications in many disciplines, including natural language and image processing [244]. In particular, they rank among the best machine learning classifiers in the MNIST benchmark [245] for hand-written digit recognition [246]. In the next section, we will demonstrate the capabilities of CNNs in the context of condensed matter physics.

6.4 Discriminating phases of matter

Before we discuss *how* to employ supervised machine learning for phase classification, let us address *why* this might be useful in the first place. After all, there already exist many established methods to characterize physical properties and

⁹⁸As per custom, a machine learning architecture is termed deep, or shallow, based on how many neural layers it contains.




Figure 6.6: Illustration of a convolutional neural network. An input “image” (colored matrix) is first processed by a series of convolutional filters and pooling layers, which, respectively, extract feature maps (red, green, and blue rectangles) and perform a dimensional reduction. The result is then fed into a maximally-connected network component with a hidden layer (pink circles) and an output layer (blue circles). Modified from Ref. [P3] under APS copyright.

detect phase transitions, such as, perhaps most importantly, finite-size scaling analyses of correlation functions. However, while conceptually valuable, these approaches have important disadvantages. Typically, they make assumptions about the nature of the phases of a physical system, for example, in requiring the knowledge of explicit order parameters. In this sense, they are biased and tailored to specific scenarios. In addition, their implementation often comes at a high computational cost and, in some cases, isn’t feasible at all – we will see an example in the next chapter. In contrast, machine learning methods represent general purpose tools that make no (or few⁹⁹) prior assumptions about the type of correlations in input data. Moreover, an evaluation of the neural function $F(x)$ essentially amounts to computing simple matrix-vector products, which can be efficiently parallelized on modern multi-core systems and graphical processing units (GPUs). We hence expect machine learning to be a valuable supplement to the existing body of numerical techniques.

To discriminate the phases of a physical system, one combines the idea of machine learning with quantum Monte Carlo. Formally, QMC may be seen as a data generation process that, for every point in a, potentially unknown, phase diagram, produces Markov chains of configurations, which, statistically, represent the thermodynamic system at hand, Ch. 2. Based on these samples, one may implement machine learning that tries to identify patterns and recognize changes of correlations in the data, for example, when crossing a phase boundary. Focusing on supervised learning in particular, the procedure is as follows: First, one trains a ML model at certain points $p^{(1)}, \dots, p^{(m)}$ in the phase diagram, which are known or, at least, expected¹⁰⁰ to mark qualitatively different physical states $s^{(1)}, \dots, s^{(m)}$. Specifically, the training data set is composed of the Markov chain elements $x_j^{(i)}$, generated at these points, and a discrete set of corresponding labels that indicates to which state $s^{(i)}$ an element $x_j^{(i)}$ belongs, c.f. X and Y in

⁹⁹For example, convolutional layers implement the assumption that input data might have translation invariance.

¹⁰⁰Consider the spin-fermion model as an example: By construction, it is clear that AFM SDW correlations are present for tuning parameter values $r \ll r_c$ but are absent for $r \gg r_c$.

Sec. 6.2. During the optimization, the ML architecture then learns to classify QMC samples based on distinct data characteristics associated with the different states $s^{(1)}, \dots, s^{(m)}$. After the training, one freezes the free parameters, e.g. the weights and biases of a neural network, and utilizes the ML model to classify QMC data generated at non-training points $p \notin \{p^{(1)}, \dots, p^{(m)}\}$. In this way, one can map out the phase diagram based on the learned characteristics of the states $s^{(1)}, \dots, s^{(m)}$. In particular, by tracking how the neural output changes along a straight path connecting, say, $p^{(1)}$ and $p^{(2)}$, one may identify a phase transition as the unique point for which the ML model indicates with equal confidence that the found correlations match those of the states $s^{(1)}$ and $s^{(2)}$ (see Fig. 6.7 below).

The power of supervised learning in the context of condensed matter physics has been demonstrated in numerous recent studies [212, 227, 247–250]. In particular, CNNs have proven to be capable of reliably discriminating phases in itinerant electron systems [227, 228] – we explicitly demonstrate this point below. Finally, let us note that the described supervised learning scheme can be turned into a semi-supervised approach by repeatedly applying it on a small scale [228]. Truly unsupervised machine learning studies of condensed matter systems based on, for instance, principle component analysis (PCA) have been conducted in Refs. [224, 251, 252].

6.4.1 Spin-density wave transition in the spin-fermion model

Let us demonstrate the supervised learning approach outlined above in the context of the onset of antiferromagnetism in the spin-fermion model, Eq. (3.4). To lower the computational cost of DQMC simulations and to facilitate the ML-based phase discrimination, we consider a two-component order parameter. In this case, the model has an $O(2)$ symmetry and features an AFM SDW phase transition of BKT type at finite temperatures [31, 32]. Specifically, we choose the model parameters as in Ref. [32], i.e. $u = 1$, $c = 2$, $\lambda = 3$, $t_{xh} = t_{yv} = 1$, $t_{xv} = t_{yh} = 0.5$, and $\mu = 0.5$.

In principle, there are two choices of what to use as data input for the machine learning: the bosonic configurations ϕ or snapshots of the equal-time Green's function G_ϕ , Eq. (2.21). Clearly, since for the spin-fermion model the former represent the order parameter, which is understood to develop a finite expectation value at the AFM transition, they appear as the superior alternative. However, for other systems, such as the Hubbard model, the ϕ -field is merely a Hubbard-Stratonovich (or Hirsch) degree of freedom, which may be of either charge or spin character, and hence not directly related to the transition at all. For this reason, we follow the study in Ref. [227], and generally prefer the the Green's function data over the bosonic configurations.

As for the ML architecture, we use the TensorFlow framework [253] to implement the CNN illustrated in Fig. 6.6, where in the convolutional layers we use kernels of size 5×5 and ReLU activation functions. Specifically, the first layer consists of 64 independent kernels and the second of 32. The pooling window size is chosen as 2×2 . In the subsequent fully connected neural network, we use 256 hidden and two final neurons. The output of the latter, which is normalized by a softmax activation to always sum up to unity, is interpreted as

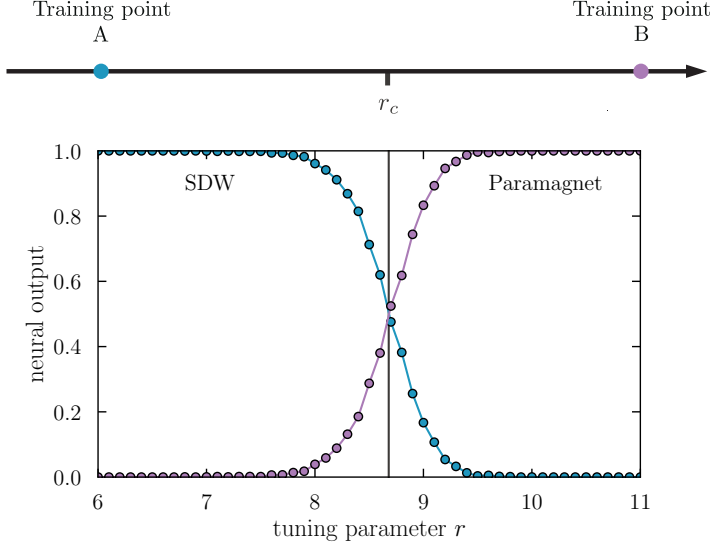


Figure 6.7: AFM SDW transition in the O(2) symmetric spin-fermion model at inverse temperature $\beta = 6$ as detected by the CNN, illustrated in Fig. 6.6. Shown is the normalized neural output of the network (blue and pink), which has been trained at points $r_A = 6.0$ and $r_B = 11.0$. The vertical solid line indicates the point of “maximal confusion” $r_{mc} \approx r_c \approx 8.67$, where both neurons indicate 50% confidence.

the networks confidence for the presence of a paramagnetic or antiferromagnetic state, respectively.

As illustrated in Fig. 6.7, we train the CNN based on 4000 Green’s function samples deep inside of each of the adjacent phases, that is at tuning parameter values $r_A = 6.0$ and $r_B = 11.0$ for the SDW phase and paramagnetic phase, respectively. In the stochastic optimization, we consider a cross-entropy cost function and a mini batch size of 128. Afterwards, we lock the weights and biases of the network and present it with Green’s function data at tuning parameter values between r_A and r_B . The neural output $F(r)$ of the CNN, obtained by this approach, is shown in Fig. 6.7. First, we note that the network successfully discriminates the SDW and paramagnetic phases as indicated by the near zero and unity outputs in the vicinity of the extremal training points. Furthermore, the neural curves clearly show that upon varying the tuning parameter r one transitions from one phase to the other, indicating the existence of a phase transition. Determining the position of the latter by the “maximal confusion” criterion $F(r) = 1/2$, we find the estimate $r_{mc} \approx 8.67$, which approximately agrees with the conventional finite-size scaling analysis conducted in Ref. [32]. This demonstrates that the machine learning approach is capable of discriminating phases of matter by extracting physical correlations from QMC data.

7 | Quantum loop topography of metallic quantum criticality

In this chapter, we demonstrate that machine learning, when combined with DQMC, can be utilized to study two essential aspects of metallic quantum criticality: superconductivity, the secondary, emergent electronic order, and the finite temperature quantum critical fan, typically associated with a breakdown of Fermi-liquid theory, Fig. 1.4. Specifically, we consider a quantum loop topography (QLT) [225] approach for probing transport in itinerant many-fermion systems, which is based on the autonomous learning of imaginary time current-current correlations, Eq. (5.9). We showcase that this physics-inspired ML technique is capable of recognizing superconductivity in the $O(2)$ symmetric spin-fermion model and, for comparison, the attractive Hubbard model in good agreement with established knowledge. As we will discuss, compared to traditional methods, the QLT approach only requires a fraction of the raw DQMC data and has a significantly lower computational cost. Because of this numerical efficiency, it is perfectly suited for identifying changes in electronic transport properties across entire phase diagrams. Focusing on Yukawa coupling strengths for which superconductivity is suppressed, we present such large scale QLT studies for the AFM QCP of the spin-fermion model and the Ising-nematic QCP considered in Ref. [39]. As we demonstrate, this analysis unambiguously reveals robust and distinguished fan-shaped regions in the phase diagrams associated with the metallic QCPs and strongly indicates the existence of extended non-Fermi liquid regimes. Note that the presented collaborative¹⁰¹ studies have been published in Refs. [P3–P5], from which we have compiled this chapter.

7.1 Quantum loop topography

Machine learning based on artificial neural networks has recently emerged as a powerful method for quantum state recognition in condensed matter physics [212, 224, 226–228, 247–252, 254–256], see Ch. 6. In particular, it has been demonstrated that convolutional neural networks can reliably discriminate phases of matter in itinerant electron systems [227, 228], Sec. 6.4. However, most of the employed ML architectures stem from applications outside of physics, such as image recognition (CNNs), email spam filtering (SVMs¹⁰²), and autonomous

¹⁰¹The author's contributions to these studies include generating and providing DQMC data for the AFM QCP, preprocessing the same for QLT, training and applying the CNNs in Ref. [P3], producing all figures in Ref. [P3] and Fig. 2 in Ref. [P4], as well as participating in the writing of the manuscripts, Refs. [P3, P4].

¹⁰²Support-vector machines, which are related to the perceptron discussed in Ch. 6.




Figure 7.1: Architecture of the quantum loop topography approach. A dimensional reduction of the full Green’s function data is obtained by only considering correlations along (short) spatial loops. For illustration purposes only four exemplary loops (yellow, red, green, purple) are shown. The resulting quantum loop vector field (colored lattices) are fed into a maximally connected feed-forward neural network. First appeared in Ref. [P4].

driving (RNNs), and are not tailor-made for a physical analysis. As a consequence, from a traditional scientist’s point of view, they often seem rather artificial and notoriously opaque. Even more importantly, one generally expects that these problem-agnostic setups need longer training – they need to learn a reliable discriminator from scratch – and, presumably, do not exhaust the full potential of ML for condensed matter physics.

Quantum loop topography [225], Fig. 7.1, is a numerical method which aims to improve upon these limitations by integrating information about the targeted physical response into the ML scheme. It stands out from other quantum state recognition approaches as a preprocessing step that selects and organizes DQMC data by using a physically motivated “loop topography” filter [P3]. The extracted loop vector field, which contains the dissected and preprocessed information, is then fed into a regular, shallow neural network, Fig. 7.1, to implement a quantum state discrimination. So far, variants of QLT have been successfully utilized to distinguish conventional and fractional Chern insulators from trivial insulators based on the topological Hall response of a system [225] and to detect \mathbb{Z}_2 quantum spin liquid phases by targeting Wilson loops [226]. In contrast, our studies below [P3, P4] represent the first implementation of QLT in the context of interacting itinerant electron systems and, in particular, metals close to quantum criticality.

7.1.1 Current-current correlations

To the end of being sensitive to superconductivity and a potential “strange metal” state, which is generally characterized by a linear temperature dependence of the electrical resistivity, Sec. 1.2, we target longitudinal transport and consider the zero-frequency current-current correlation function, cf. Eq. (5.9),

$$\Lambda_{xx}(\mathbf{r}_1, \mathbf{r}_2) = \int d\tau \langle j_x(\mathbf{r}_1, \tau) j_x(\mathbf{r}_2, 0) \rangle. \quad (7.1)$$




Figure 7.2: Illustration of the (i) triangular and (ii) quadrilateral loop operators employed in the QLT of the longitudinal transport. Taken from Ref. [P3] under APS copyright.

As discussed in Sec. 5.2.2, in the appropriate limit, the Fourier transform of Λ_{xx} is related to the superfluid density ρ_s , which indicates the onset of superconductivity by surpassing the critical BKT value [137].

To see that the current-current correlator can have an underlying loop structure, we consider a gapped Hamiltonian with a single flat band, i.e. $H' = -\Pi$, where $\Pi \equiv |g\rangle\langle g|$ projects onto the ground state $|g\rangle$ [P3]. In this case, one can compute the current-current correlation function Λ_{xx} at zero temperature¹⁰³, [P3]

$$\begin{aligned} \Lambda_{xx}(\mathbf{r}_1, \mathbf{r}_2) &= \langle g | j_x(\mathbf{r}_1) (1 - \Pi) j_x(\mathbf{r}_2) | g \rangle \\ &= \text{Tr} [\Pi j_x(\mathbf{r}_1) (1 - \Pi) j_x(\mathbf{r}_2)], \\ &= \sum_{\mathbf{r}_3 \mathbf{r}_4} P_{\mathbf{r}_2 \mathbf{r}_4} P_{\mathbf{r}_4 \mathbf{r}_1} P_{\mathbf{r}_1 \mathbf{r}_3} P_{\mathbf{r}_3 \mathbf{r}_2} (x_1 - x_4) (x_2 - x_3) \\ &\quad - \sum_{\mathbf{r}_4} P_{\mathbf{r}_2 \mathbf{r}_4} P_{\mathbf{r}_4 \mathbf{r}_1} P_{\mathbf{r}_1 \mathbf{r}_2} (x_1 - x_4) (x_2 - x_1). \end{aligned} \quad (7.2)$$

Here, $P_{\mathbf{r}'\mathbf{r}} \equiv \langle g | c_{\mathbf{r}'}^\dagger c_{\mathbf{r}} | g \rangle = \delta_{\mathbf{r}'\mathbf{r}} - \langle g | c_{\mathbf{r}} c_{\mathbf{r}'}^\dagger | g \rangle$ is essentially the Green's function, the position vector is $\mathbf{r}_i = (x_i, y_i)^T$, and we have used the definition of the current density operator, $j_x(\mathbf{r}) = -i[H(\mathbf{r}), \hat{x}]$, for the third equality. Note that the current-current correlation function Λ_{xx} decomposes into weighted sums of quadrilateral and triangular loops of two-point functions.

Based on the availability of these Green's functions in DQMC simulations, our QLT strategy is to approximate Λ_{xx} by short loop contributions. At each lattice site j , we build a loop vector by computing all chained products of two-point functions corresponding to loops with three, L_{jkl}^Δ , and four, L_{jklm}^\square , vertices that include the site j [P3]. Specifically, we consider the loop products

$$L_{jkl}^\Delta \equiv \tilde{P}_{jk} |_\alpha \tilde{P}_{kl} |_\beta \tilde{P}_{lj} |_\gamma, \quad (7.3)$$

$$L_{jklm}^\square \equiv \tilde{P}_{jk} |_{\alpha'} \tilde{P}_{kl} |_{\beta'} \tilde{P}_{lm} |_{\gamma'} \tilde{P}_{mj} |_{\delta'}, \quad (7.4)$$

where $\tilde{P}_{jk} |_\alpha$ indicates that we evaluate the two-point functions P_{jk} for individual Monte Carlo samples α . In particular, we impose a short-distance cutoff d_c and only take short loops into account. An illustration of the lowest order loop operators L_{jkl}^Δ and L_{jklm}^\square is provided in Fig. 7.2.

¹⁰³See the supplementary of Ref. [P3] for more details.

Clearly, the QLT approximation of the current-current correlation function based on only the contributions L_{jklm}^{\square} and L_{jkl}^{\triangle} built from single Monte Carlo samples, rather than full thermodynamic expectation values, is no substitute for a rigorous computation of Λ_{xx} , especially for a gapless system [P3]. Nonetheless, as we demonstrate below, QLT may serve as a valuable proxy in probing longitudinal transport in itinerant electron systems. In particular, while calculating the current-current correlator in DQMC requires the computation of the time-displaced Green's function, which is numerically and computationally demanding¹⁰⁴, Sec. 4.1.3, QLT is straightforward to implement based on the readily available equal-time Green's functions.

7.2 Probing superconducting transport

As demonstrated in Ch. 5, one of the prominent features of the onset of antiferromagnetism in a metal is the concomitant emergence of superconductivity from critical order parameter fluctuations. In the following, we will showcase that the QLT approach is capable of detecting a transition from regular metallic to superconducting transport. Specifically, to allow a benchmark of the method, we apply it to well established superconducting transitions of different symmetries in two metallic systems.

7.2.1 Attractive Hubbard model

First, we consider the prototypical negative- U Hubbard model [137] on a two-dimensional square lattice, which we already discussed in the context of the sign-problem in Ch. 3. For convenience, we repeat its Hamiltonian here,

$$H = -t \sum_{\langle i,j \rangle, s} (c_{is}^{\dagger} c_{js} + \text{h.c.}) - |U| \sum_i \left(n_{i\uparrow} - \frac{1}{2} \right) \left(n_{i\downarrow} - \frac{1}{2} \right) - \mu \sum_i n_i. \quad (7.5)$$

Being one of the conceptually simplest condensed matter models which, at the same time, is complex enough to host interesting physics, the attractive Hubbard model has been the target of various analytical and numerical studies over the last decades [137, 158, 159, 192, 257, 258]. Given the presence of an explicit attractive interaction between the itinerant fermions, it is understood from BCS theory that the model undergoes a superconductivity transition of Berezinskii–Kosterlitz–Thouless type at sufficiently low temperatures [158]. However, at half-filling, $\mu = 0$, there is a degeneracy between s-wave singlet pairing and the formation of charge-density wave order due to particle-hole symmetry [258], Sec. 3.3. This leads to an effective three-component order parameter and a suppression of T_c in two dimensions [258]. Only away from half-filling is this degeneracy broken in favor of SC order [158, 258]. For this reason, we choose the parameters as $t = 1$, $|U| = 8$, and tune the chemical

¹⁰⁴Empirically, we find that measuring the time-displaced Green's function has approximately the same computational complexity as performing the DQMC simulation in the first place. This is plausible, as one needs to construct stacks of time slice propagators similar to the one discussed in Sec. 4.2.1.




Figure 7.3: The neural outputs of the QLT and CNN architectures for superconducting transport as a function of inverse temperature β for the negative- U Hubbard model, Eq. (7.5). While the inputs for the deep CNN are the unprocessed Green’s functions, produced by DQMC, the QLT architecture only regards quantum loops as in Eqs. (7.3) and (7.4), Fig. 7.2. Both architectures are trained with samples at low temperature $\beta = 20$ and high temperature $\beta = 2$. The parameters are $|U| = 8$, $\langle n \rangle \approx 0.9$, and $L = 8$. Taken from Ref. [P3] under APS copyright.

potential such that the density $\langle n \rangle = \langle n_{\uparrow} \rangle + \langle n_{\downarrow} \rangle \approx 0.9$ is slightly below unity. In this case, the critical temperature takes the value $T_c \approx 0.1$ [137].

Focusing on a system of linear size $L = 8$, we conduct DQMC simulations of the attractive Hubbard model and use the obtained Green’s function data as input for the QLT. The latter is then fed into a shallow neural network with a single hidden layer composed of 40 neurons, Fig. 7.1. To gauge the internal parameters of the neural network, we train the latter deep in the normal phase, $\beta = 2$, and deep in the superconducting phase, $\beta = 20$, based on ~ 20000 DQMC samples for each temperature. Hereby, we use a cross-entropy cost function with L2 regularization to avoid over-training, Eq. (6.12), and a mini-batch size of 10 in the stochastic gradient descent optimization. The value of the single output neuron, which is normalized by a softmax activation function and, thus, lies in the interval $[0, 1]$, is interpreted as the networks confidence for recognizing superconductivity. The optimized network is then used to classify the electric transport of the system at temperatures between the extremal training points, i.e. over the range $0.05 \leq T \leq 0.5$.

Clearly, as shown in Fig. 7.3, the neural network, fed with compressed QLT data, is capable of reliably distinguishing normal and superconducting transport at high and low temperatures as indicated by neural outputs close to zero and one, respectively. At the critical temperature $T_c \approx 0.1$, the network classifies the Hubbard model as superconducting with a certainty as high as $\sim 90\%$. However, one observes that the increase of the neural output with decreasing temperature is smooth and steady, rather than abrupt. As a consequence, the point of




Figure 7.4: Mean-field superconducting transition of the attractive Hubbard model for $U = -8$ and $\mu = -0.5$. Comparison of the mean-field s -wave pairing gap Δ (green diamonds) and the neural output from the QLT approach (red dots). The vertical dashed line indicates the mean-field transition temperature $\beta_{\text{mf}} \sim 0.545$ [P3]. Taken from Ref. [P3] under APS copyright.

“maximal confusion”, Sec. 6.4, occurs at approximately $T_{\text{mc}} = \beta_{\text{mc}}^{-1} \sim 0.28$, which lies slightly above T_c . We attribute this finding to the presence of diamagnetic superconducting fluctuations, which are expected to onset before entering the superconducting phase.

To justify this interpretation, we consider a mean-field analysis of the attractive Hubbard model, in which fluctuations are neglected entirely. In this case, we expect that the QLT method is able to sharply detect the onset of superconductivity at the mean-field transition temperature $\beta_{\text{mf}} \sim 0.545$. Specifically, our approach is as follows: We solve the self-consistent gap equations at each inverse temperature β to obtain the s -wave pairing gap $\Delta(\beta)$ [P3]. Then, we perform Monte Carlo simulations of the corresponding Bogoliubov-de-Gennes mean-field Hamiltonian with the respective $\Delta(\beta)$ to generate the necessary QLT input [P3]. The phase diagram resulting from this procedure is presented in Fig. 7.4. Note that, here, the high and low temperature training points are chosen as $\beta = 0.3$ and $\beta = 0.8$, respectively. As apparent in the figure, the output of the neural network has a jump-like character near β_{mf} such that the location of the phase transition is accurately predicted, i.e. $\beta_{\text{mc}} \approx \beta_{\text{mf}}$. This suggests that QLT does indeed capture the onset of superconductivity at the correct critical temperature if superconducting fluctuations are absent.

Finally, we benchmark our QLT approach against a regular CNN-based phase discrimination scheme, Sec. 6.4. A direct comparison of the two methods for the attractive Hubbard model is shown in Fig. 7.3. Within the expected accuracy, we observe that both techniques produce essentially the same neural output and detect the onset of superconducting fluctuations based on the DQMC data. However, note that the CNN structure contains significantly more free parameters: About six times the number of hidden neurons and, even more importantly,




Figure 7.5: Output of the QLT (red) and CNN (blue) architectures indicating the detection of superconducting transport in the vicinity of the AFM QCP of the O(2) symmetric spin-fermion model, Eq. (7.6). The vertical lines indicate the critical transition temperature $\beta_c \sim 12.5$ (dashed), derived from the superfluid density measurements, and the onset of diamagnetic fluctuations $\beta_{dia} \sim 6.9$ (dotted) where the orbital magnetic susceptibility changes sign [32]. The inset (modified from [32]) illustrates the chosen finite-temperature scan cutting into the superconducting dome. Both ML architectures were trained at $\beta = 30$ for superfluid transport and $\beta = 5$ for metallic transport. Taken from Ref. [P3] under APS copyright.

the weights and biases associated with every convolutional layer. During the optimization step, the latter are dynamically adjusted based on the training data whereas the QLT filtering is entirely static. Furthermore, while the entire equal-time Green's functions (of size $2N \times 2N$) are available to the CNN, the QLT setup only regards quantum loops of short lengths, i.e. matrix inputs of size $\sim 2N \times d_c$. It hence achieves similar performance based on a significantly smaller data set, indicating the potential of the QLT approach.

7.2.2 Spin-fermion model

We now turn to superconductivity arising from quantum critical fluctuations in the vicinity of an O(2) symmetric AFM QCP [2, 31, 32, 81]. We consider the sign-problem free spin-fermion model of Eq. (3.4) with an easy-plane order parameter, i.e $N_\phi = 2$, repeated here for convenience,

$$\begin{aligned}
 S_{\text{SFL}} = & \int_0^\beta d\tau \sum_{\mathbf{r}, \mathbf{r}'} \sum_{s, \alpha} \psi_{\alpha \mathbf{r} s}^\dagger [(\partial_\tau - \mu) \delta_{\mathbf{r} \mathbf{r}'} - t_{\alpha \mathbf{r} \mathbf{r}'}] \psi_{\alpha \mathbf{r}' s} \\
 & + \lambda \int_0^\beta d\tau \sum_{\mathbf{r}} \sum_{s, s'} e^{i \mathbf{Q} \cdot \mathbf{r}} \phi_{\mathbf{r}} \cdot \psi_{x \mathbf{r} s}^\dagger \sigma_{ss'} \psi_{y \mathbf{r} s'} + \text{h.c.} \\
 & + \int_0^\beta d\tau \sum_{\mathbf{r}} \left[\frac{1}{2c^2} (\partial_\tau \phi_{\mathbf{r}})^2 + \frac{1}{2} (\nabla \phi_{\mathbf{r}})^2 + \frac{r}{2} \phi_{\mathbf{r}}^2 + \frac{u}{4} (\phi_{\mathbf{r}}^2)^2 \right].
 \end{aligned} \tag{7.6}$$

Specifically, we choose the nearest-neighbor fermion hopping amplitudes as $t_{xh} = t_{yv} = 1$ and $t_{xv} = t_{yh} = 0.5$. We further set the Yukawa coupling to $\lambda = 3$, the chemical potential to $\mu = 0.5$, $u = 1$, and the bare bosonic velocity to $c = 2$. The phase diagram of the spin-fermion model for these parameters is shown in the inset of Fig. 7.5 [32]. To apply the QLT analysis, we tune the system to the vicinity of the AFM QCP, $r = 10.35$, and consider a finite-temperature scan, $T \in [0.2, 0.03]$, cutting into the dome of d-wave superconductivity. The training points for the shallow neural network are taken as the boundaries of this temperature interval: $T = 1/30$ for superconducting and $T = 1/5$ for metallic transport.

The results of our QLT study are presented in Fig. 7.5. We find that the normal state and the d-wave superconducting state are reliably distinguished by the QLT approach as evident from extremal neural outputs at high and low temperatures, respectively. Right at the critical temperature $T_c \approx 0.08$, which has been determined in Ref. [32] by means of a superfluid density computation as discussed in Sec. 5.2.2, the shallow neural network identifies superconducting transport with about 90% confidence. As for the Hubbard model, the neural output is found to steadily increase over an extended temperature range. Notably, the latter coincides with a region above the BKT transition, $T_c \approx 0.08 < T < T_{\text{dia}} \approx 0.15$, which is characterized by pronounced diamagnetic orbital correlations [32]. This supports our assessment above, that QLT is even sensitive to the onset of superconducting fluctuations.

Comparing the QLT architecture to a conventional CNN, Fig. 7.5, we find good agreement between both ML approaches. As mentioned above, this is in spite of the fact that QLT has a lower complexity and only uses a fraction ($\sim d_c/L^2$) of the full Green's function input processed by the CNN. Moreover, we note that the detection of superconducting transport displayed in Fig. 7.5 has been obtained by using only $O(10)$ uncorrelated DQMC samples. This constitutes a huge reduction over the typically required samples in order to obtain statistically reliable DQMC measurements of the superfluid density [32, 259].

7.3 Mapping out non-Fermi liquid regimes

The findings above establish that our QLT+NN scheme is capable of efficiently detecting changes in transport properties of strongly correlated electron systems. In the following, we investigate [P4] whether the same technique is capable of identifying more subtle anomalous transport, such as the linear- T resistivity observed in unconventional superconductors, Sec. 1.2. To that end, we numerically focus on two types of metallic quantum critical points, the $O(2)$ symmetric AFM QCP of the spin-fermion model [31, 32] and the Ising-nematic QCP studied in Refs. [39, 40], and attempt to map out non-Fermi liquid regimes in the associated phase diagrams. Note that these representative QCPs cover both conceptual classes of metallic ordering discussed in Sec. 1.1.1. While the onset of antiferromagnetic order is characterized by a *finite* commensurate wave vector, $\mathbf{Q} = (\pi, \pi) \neq 0$, and, consequently, the opening of gaps at specific hot spots on the Fermi surface, the Ising-nematic order is of uniform $\mathbf{Q} = 0$ type, Fig. 1.3.




Figure 7.6: Phase diagram of the $O(2)$ symmetric spin fermion model overlaid with machine-learned Fermi liquid to non-Fermi liquid crossover. The color map shows the output of a shallow neural network with a single hidden layer of 10 neurons trained on QLT-filtered equal-time Green’s function data extracted from DQMC simulations – see App. D for the underlying data grid. A value of 1 (dark red) corresponds to the non-Fermi liquid, a value of 0 (dark blue) corresponds to the Fermi liquid, and intermediate values represent the crossover region. White boxes indicate the Fermi liquid training points, $(r = 0.3, T = 0.05)$ and $(r = 1.4, T = 0.05)$, and the black circle the same for the non-Fermi liquid, $(r = 0.7, T = 0.2)$. The red star indicates the location of the quantum critical point, i.e. $r_c = 0.62$ [31]. The solid black line indicates the finite-temperature SDW phase transition. The dashed black line shows the – largely suppressed – superconducting $T_c(r)$ [31]. First appeared in Ref. [P4].

7.3.1 $O(2)$ symmetric AFM QCP

To suppress superconductivity and uncover the bare AFM QCP, we consider the spin-fermion model, Eq. (7.6), at Yukawa coupling $\lambda = 1.5$, which, compared to Sec. 7.2.2, is reduced by a factor of two. We further set the boson velocity to $c = 3$ and the quartic boson coupling to $u = 1$. In this case, the QCP is found at the critical tuning parameter value $r_c = 0.62$ and is barely masked by a superconducting dome, see Fig. 7.6 or Fig. 2b in Ref. [31]. Given the easy-plane character of the order AFM order parameter, the system features a classical SDW phase transition of BKT character at finite-temperatures in agreement with the Mermin-Wagner theorem.

Three-point classification

As one anticipates Fermi liquid behavior away from the AFM QCP, i.e. for $r \ll r_c$ and $r \gg r_c$, and a putative non-Fermi liquid regime at finite temperatures

above $r_c = 0.62$, we employ a three-point supervised learning scheme for the QLT-fed shallow neural network, Fig. 7.6. Specifically, we train the latter based on 3200 input vectors at the representative points ($r = 0.3, T = 0.05$), ($r = 1.4, T = 0.05$), and ($r = 0.7, T = 0.2$). Here, we are guided by the fact that for $r = 0.7$ and $T = 0.05$ the non-Fermi liquid character has been explicitly demonstrated in Ref. [31] by extracting the Matsubara self-energy, cf. Sec. 5.3.2. While this method for detecting a novel non-Fermi liquid state is theoretically appealing and numerically exact, it requires the calculation of the time-displaced Green's function and is hence associated with a considerable computational cost. Importantly, this cost is incurred for every parameter point of the phase diagram. For this reason, in spite of the DQMC simulations of Ref. [31] having a scope of $O(10)$ million CPU hours on modern supercomputers, i.e. comparable to our computational efforts in this thesis, mapping out an extended non-Fermi liquid region has so far been out of reach.

As illustrated in Fig 7.6, the QLT-fed neural network manages to learn the qualitative features of the phase diagram of the spin-fermion model and strikingly reveals an extended fan-shaped non-Fermi liquid regime. Being anchored at only three points far away from the metallic QCP, it is particularly remarkable that the identified non-Fermi liquid state, as supported by the high temperature training point, extends to the lowest temperatures. Moreover, the recognized NFL region narrows upon decreasing T and clearly culminates into the AFM QCP. Note that the location of the latter was not provided to the ML model at any point and, thus, represents a highly non-trivial feature autonomously learned by the neural network. Focusing on the boundaries of the identified NFL regime, the QLT architecture reveals the expected steep rise of the NFL-FL crossover temperature away from the QCP without explicit prior knowledge. For small tuning parameter values, we find qualitative agreement with the classical SDW phase transition line at finite temperatures.

To probe the numerical stability of the obtained QLT phase diagram, Fig. 7.6, we investigate the robustness of the neural output against the choice of training points. As shown in Fig. 7.7, we find that the identification of the NFL regime is remarkably stable and shows only modest quantitative changes when the NFL training point is moved closer to or further from the ordered phase (panels a and b), and similarly when the disordered FL training point is shifted (panels c and d). This observation stands in stark contrast to neural networks trained on snapshots of classical order parameters, which require training points in the immediate vicinity of the phase boundary [212], or otherwise fail¹⁰⁵.

Binary classification

Given the remarkable results of the QLT-based three-point scheme, we further investigate the neural network's learning by individually considering the SDW-NFL phase transition, for $r \leq r_c$, and the NFL-FL crossover, at larger tuning parameter values. Specifically, we explore the phase diagram by applying a binary classification in which we take only two of the three training points considered above and, besides the NFL regime, train the shallow neural network

¹⁰⁵We have reproduced the numerical study in Ref. [212] and verified this statement explicitly.




Figure 7.7: Robustness of the QLT-based three-point non-Fermi liquid classification for the spin-fermion model, Eq. (7.6). The different panels indicate that the identified NFL regime varies only modestly when: (a) the non-Fermi liquid training point is moved closer to the ordered phase ($r = 0.5$); (b) the non-Fermi liquid training point is moved further from the ordered phase ($r = 0.9$); (c) the disordered Fermi liquid training point is moved closer to the QCP ($r = 1.1$); (d) the disordered Fermi liquid training point is moved further from the QCP ($r = 1.8$). Modified from the supplement of Ref. [P4].




Figure 7.8: Binary classification for the spin-fermion model. The neural network is trained on QLT-filtered DQMC data to distinguish (a) the AFM SDW phase from magnetic disorder and (b) the non-Fermi liquid regime (NFL) from the disordered Fermi liquid regime (FL). The training points are indicated by white boxes, $(r = 0.3, T = 0.05)$ in (a) and $(r = 1.1, T = 0.05)$ in (b), and black circles, $(r = 0.7, T = 0.2)$ in (a) and (b). First appeared in Ref. [P4].

only in the antiferromagnetically ordered phase or the disordered FL region. The results of this analysis are presented in Fig. 7.8.

As apparent from Fig. 7.8a, The neural network clearly identifies the entire magnetic phase transition at finite temperatures in almost quantitative agreement with the systematic finite-size analysis conducted in Refs. [31, 32]. Notably, this performance is achieved without any explicit reference to the AFM order parameter. Instead, it is solely based on short quantum loops of equal-time Green’s functions representing proxies for longitudinal transport, Sec. 7.1. Although in this respect the observed quality of detection is somewhat surprising, it seems that signatures of the opening of gaps at the hot spots on the Fermi surface are still included in the local QLT input [P3]. Remarkably, as shown in Fig 7.8b, a binary classification targeting the disordered Fermi liquid state reveals a marked non-Fermi liquid crossover in qualitative agreement with the three-point learning results in Fig. 7.6. This demonstrates that the feature is stable and independent of a referencing of the ordered phase. Note that compared to the well-defined SDW phase transition the neural output for the NFL-FL crossover appears more spread in the transient area – see the pronounced color bleeding in Fig. 7.8b.

Similar to the three-point discrimination above, we further study the robustness of the binary classification by systematically altering the location of the respective SDW, NFL, and FL training points. As illustrated in Fig. 7.9, we notice a remarkable stability for both transitions as indicated by only modest changes of the neural output. In particular, while a shift of the NFL reference point p_{NFL} to the close vicinity of the SDW phase boundary leads to the largest observed deviations, the magnetic phase transition is reliably recognized even when moving p_{NFL} far away from the QCP and to the lowest considered temperature $T = 0.05$, see bottom panel in Fig. 7.9a.

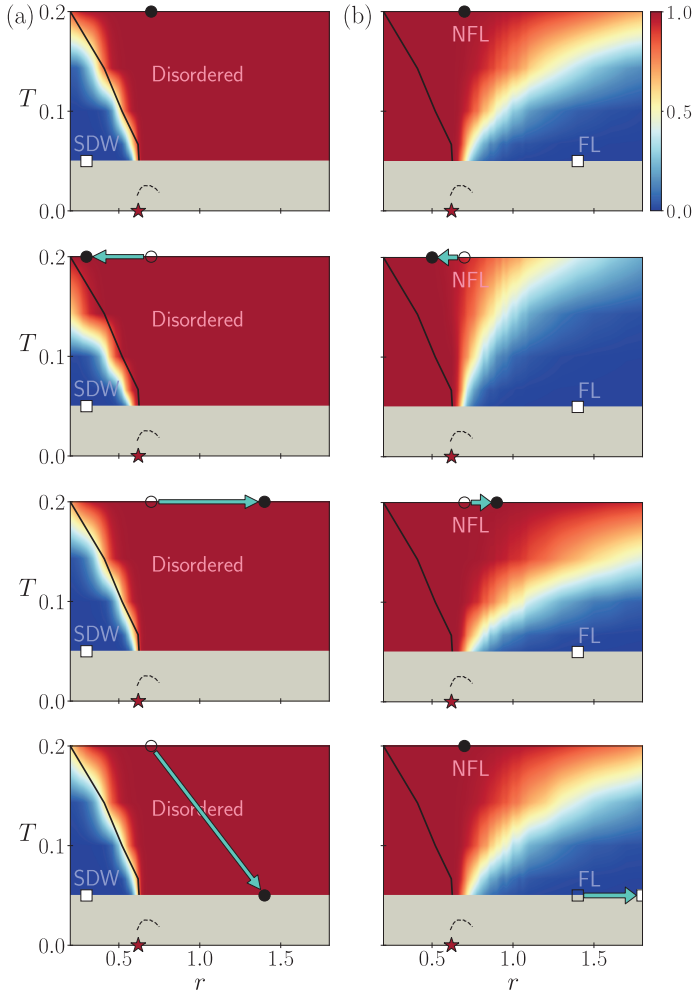


Figure 7.9: Robustness of the binary classification for the spin-fermion model. The panels show that the QLT-based discrimination of SDW order and disorder (a) as well non-Fermi liquid and Fermi liquid transport (b) varies only modestly when shifting the respective training points (cyan arrows). Modified from the supplement of Ref. [P4].

7.3.2 Ising-nematic QCP

As a prototype of a metallic quantum critical point associated with uniform $\mathbf{Q} = 0$ order, we consider the sign-problem free lattice model for Ising-nematic quantum criticality studied in Refs. [39, 40]. The Hamiltonian is $H = H_f + H_b + H_\alpha$, where


$$\begin{aligned} H_f &= -t \sum_{\langle i,j \rangle, s} c_{is}^\dagger c_{js} - \mu \sum_{is} c_{is}^\dagger c_{is}, \\ H_b &= V \sum_{\langle \langle i,j \rangle \rangle, \langle \langle k,l \rangle \rangle} \tau_{ij}^z \tau_{kl}^z - h \sum_{\langle i,j \rangle} \tau_{ij}^x, \\ H_\alpha &= \alpha t \sum_{\langle i,j \rangle, s} \tau_{ij}^z c_{is}^\dagger c_{js}. \end{aligned} \quad (7.7)$$

As illustrated in Fig. 7.10, the model's degrees of freedom are spin-1/2 fermions c_{is} that live on the sites i of a square lattice, and pseudospins τ_{ij}^a , with $a = x, y, z$, that are situated on the bonds connecting two neighboring sites i and j . While the fermions do not interact with each other, H_f , the pseudospins are governed by a $d = 2$ transverse field quantum Ising model, H_b . Here, V is a nearest-neighbor Ising interaction – $\langle \langle i, j \rangle \rangle, \langle \langle k, l \rangle \rangle$ denotes a pair of nearest-neighbor bonds – and h is a transverse field. Upon tuning h , the pseudospins undergo a quantum phase transition from a paramagnetic to an “antiferromagnetic” state that breaks $\pi/2$ rotational symmetry [1, 39]. The latter is indicated in Fig. 7.10: the z components of the pseudospins on horizontal bonds (white squares) differ from those on their neighboring vertical bonds (red squares). When coupled to the fermion bond density in H_α , this ordering implies a nematic transition at $h = h_c$ as the effective fermion-hopping becomes anisotropic.

Following Ref. [40], we set the dimensionless pseudospin-fermion coupling to $\alpha = 1.5$, the “antiferromagnetic” interaction to $V = 0.5$, and both the chemical potential as well as the hopping amplitude to $\mu = t = 1$. In this case, it has been established by conventional finite-size scaling and computations of the relevant susceptibility that there is a nematic finite-temperature phase transition leading to a metallic QCP at $h_c \approx 2.6$ [40], Fig. 7.11. As in the AFM spin-fermion model, the latter is covered by a broad superconducting dome with a high critical temperature T_c [40], indicated by the dashed line in Fig. 7.11. At temperatures $T > T_c$, there exists evidence for a non-Fermi liquid normal state in form of a linear-in- T resistivity $\rho(T)$ [2, 40]. However, the extraction of real-time dynamics

Figure 7.10:

Illustration of the two-dimensional square lattice model Eq. (7.7) hosting an Ising-nematic QCP. Itinerant fermions (dark circles) can move between neighboring lattice sites with the hopping amplitude t and interact with antiferromagnetically coupled Ising pseudospins (red and white squares). The latter are situated on the lattice bonds and subject to a transverse field (not shown). Based on Fig. 3 in Ref. [39].






Figure 7.11: Phase diagram of the Ising-nematic model, Eq. (7.7), overlaid with machine-learned Fermi liquid to non-Fermi liquid crossover. The color map shows the output of a shallow neural network with a single hidden layer of 10 neurons trained on QLT-filtered and nearest-neighbor equal-time Green’s function data extracted from DQMC simulations – see App. D for the underlying data grid. A value of 1 (dark red) corresponds to the non-Fermi liquid, a value of 0 (dark blue) corresponds to the Fermi liquid, and intermediate values represent the crossover region. White boxes indicate the Fermi liquid training points, $(h = 1.9, T = 0.17)$ and $(h = 4.1, T = 0.17)$, and the black circle the same for the non-Fermi liquid, $(h = 2.7, T = 0.5)$. The red star indicates the location of the quantum critical point, i.e. $h_c = 2.6$ [40]. The solid black line indicates the finite-temperature Ising-nematic phase transition. The dashed black line shows the superconducting $T_c(h)$ [40]. First appeared in Ref. [P4].

from imaginary time DQMC data necessarily involves a variant of analytic continuation – or the utilization of proxy quantities – and is therefore subject to considerable ambiguities.

Three-point classification

To implement the QLT-based three-point learning scheme, we choose a non-Fermi liquid anchor point at high temperatures above the Ising-nematic QCP, i.e. $(h = 2.7 \approx h_c, T = 0.5)$, and further train the neural network deep inside of the nematic and Fermi liquid states at small and large tuning parameter values, i.e. $h \ll h_c$ and $h \gg h_c$, respectively. In contrast to the study of the AFM QCP above, we supplement the QLT feature vectors with nearest-neighbor Green’s function data as we find *a posteriori* that this slightly improves the detection (see Ref. [P4] for a detailed comparison). In the generation of the input data, we consider a “cold start” of DQMC simulations for $h < h_c$ in which the pseudospin degrees of freedom are initialized in one of the classical ground states of H_b to




Figure 7.12: Binary classification for the Ising-nematic model. The neural network is trained on QLT-filtered DQMC data to distinguish (a) Ising-nematic order from disorder and (b) the non-Fermi liquid regime from the disordered Fermi liquid regime. The training points are indicated by white boxes, $(h = 1.9, T = 0.17)$ in (a) and $(h = 4.1, T = 0.17)$ in (b), and black circles, $(h = 2.7, T = 0.5)$ in (a) and (b). Modified from the supplement of Ref. [P4].

facilitate the convergence to a local minimum [P3].

As shown in Fig. 7.11, the neural network learns the nematic phase boundary in striking agreement with the conventional analysis [40], indicated by the solid line, down to the lowest temperatures. Similar to our findings for the AFM QCP, it identifies a NFL-FL crossover bending away from the QCP upon increasing T . In between those transitions, a marked NFL regime is recognized. Note that the latter extends down to the lowest considered temperatures, despite the NFL training point being located at $T = 0.5$, and leads to the vicinity of the metallic QCP at $T = 0$. We note that, in comparison to the AFM QCP, the efficient learning of the phase diagram using only three training points is particularly remarkable given that all three states – nematic, NFL, and FL – are gapless. As discussed in Ch. 1 and illustrated in Fig. 1.3, the Ising nematic order only elongates the Fermi surface and leaves the same gapless with well-defined quasi-particles. We hence conclude that the QLT-fed neural network has learned subtle changes in the fermionic correlations of the three states. Let us note in passing that, similar to our discussion around Fig. 7.7, we find that the phase diagram is robust against the choice of the NFL training point (see App. D).

Binary classification

In Fig. 7.12, we illustrate the neural output obtained by a binary classification in which only two of the training points in Fig. 7.11 are considered. We find that both the nematic ordering as well as the NFL-FL crossover are reproduced in almost quantitative agreement with the three-point learning scheme. This strongly indicates the robustness of our results. The latter is further demonstrated in App. D, where we observe only modest changes of the neural output under shifts of the two training points for both transitions.

7.4 Discussion

We have introduced a quantum loop topography scheme for machine learning longitudinal transport and have demonstrated that this approach, when applied to data obtained in DQMC simulations, can be utilized to efficiently probe essential aspects of metals near quantum criticality.

In particular, by simply processing readily available equal-time Green's functions, our QLT-based classification scheme could successfully identify the onset of superconductivity emerging in the vicinity of an AFM QCP (and in the attractive Hubbard model). Hereby, our results are in good agreement with rigorous computations of T_c based on the superfluid density [32]. However, while extracting the latter involves expensive and non-trivial computations of time-displaced Green's functions, see Ch. 4, our ML approach has a small numerical footprint and is straightforward to implement. In comparison to convolutional neural networks, which have previously been used to study phase transitions in itinerant electron systems [227, 228], the QLT architecture has the advantage that it has considerably fewer free parameters, is motivated by physics considerations, and requires only semi-local information as input.

Applying the novel ML scheme to two models of metallic QCPs, the O(2) symmetric spin-fermion model and an Ising-nematic theory, we could obtain detailed features of the associated phase diagrams such as the boundaries of the ordered phases. In both cases, the QLT approach clearly reveals the location of the QCP at $T = 0$, a non-trivial aspect autonomously learned from the raw DQMC data, and visually highlights the emergence of a striking, broad non-Fermi liquid regime. This is particularly remarkable in view of the fact that our analysis relies on only three, rather extremal training points. Even more notably is the observed robustness of the identified features under shifts of those training points, which, to the best of our knowledge, is unprecedented among ML classification studies of condensed matter systems – in many cases, e.g. Refs. [212, 227], the anchor points are fine-tuned.

Our results indicate that combining DQMC with physics-inspired ML techniques can be a powerful way to extract relevant features of phase diagrams of quantum many-body systems. The simplicity and stability of the QLT approach combined with its effectiveness in detecting subtle NFL behavior make it an promising tool for numerical condensed matter research. Looking ahead, we anticipate that it will be valuable in detecting other subtle, elusive states that are characterized by non-trivial transport properties.

Concluding remarks

In this thesis, we have investigated the onset of AFM SDW order in a metal by means of large-scale quantum Monte Carlo simulations and physics-inspired machine learning. Our presented numerical results establish the general phenomenology of the metallic AFM QCP, a long-standing problem that has been tackled by many physicists over decades. Most notably, we could rigorously demonstrate that critical SDW fluctuations can give rise to high-temperature d-wave superconductivity as well as an extended non-Fermi liquid regime. As discussed in Ch. 1, these features are characteristic for many unconventional superconductors and provide strong numerical evidence that antiferromagnetic quantum criticality plays a central role in those systems.

Despite the remarkable similarities between the obtained numerically exact phase diagram, Ch. 5, and those of high- T_c materials, Ch. 1, the considered spin-fermion model fails to capture all aspects of the rich physics of unconventional superconductors, in particular the plethora of subtle correlations in the hole-doped cuprates. Starting from the results of this thesis, it would be highly interesting to see whether the integration of additional terms into the effective theory could give rise to further complex behavior. In particular, it has been pointed out [45, 80, 260] that coupling the AFM order parameter to a \mathbb{Z}_2 gauge field could lead to a topological phase transition and a Fermi surface reconstruction prior to the onset of AFM order, as observed in ARPES measurements. Alternatively, in light of the presence of both magnetic and nematic correlations in some of the iron-pnictides, it would be worthwhile to investigate the interplay of multiple competing order parameters in a numerically exact manner [30]. Moreover, studying the spin-fermion model on frustrated lattices, considering more than two electron bands, and investigating the case of incommensurate AFM order are natural and highly interesting ways to extend our efforts in this thesis. First attempts to explore these directions have been published in Refs. [33, 51, 52, 259].

Ideally, our numerical DQMC results represent valuable feedback for theory. In particular, it is intriguing that the low-energy AFM modes are well-described by the uncontrolled Hertz-Millis theory. Obtaining a deeper analytical understanding of this phenomenon is highly desirable. In comparison to similar works on easy-axis and easy-plane antiferromagnetism [32, 33, 39], our findings demonstrate that many essential aspects of the AFM QCP are independent of the dimensionality of the order parameter – the maximal critical temperature is a notable exception. Given that the latter qualitatively changes the structure of a perturbative Feynman graph expansion, this asks for theoretical clarification. Furthermore, our DQMC study has unambiguously revealed that, within the accessed temperature regimes, the $z = 1$ fixed point discovered by Lee *et al.* [6, 7, 15] does not generally govern the physics of the metallic AFM QCP. Even for a fine-tuned

Fermi surface that is almost locally nested at the hot spots, we have not found clear indications of a strong-coupling fixed point. It would be highly interesting to see what causes this discrepancy between theory and numerical experiment. A follow-up DQMC study that, if feasible, considers even lower temperatures than we have done in this thesis could be a first step towards a solution.

In principle, it would also be desirable to study the properties of the AFM QCP in a finite magnetic field. As superconductivity is suppressed in this case, one could extract the bare critical properties for larger fermion-boson interactions, i.e. $\lambda \gg 1$, which might potentially cause further deviations from weak-coupling theory. Unfortunately, DQMC simulations of the spin-fermion model in a magnetic field are sign-problematic and, thus, exponentially inefficient. Finding ways to circumvent this limitation seems highly desirable. As we have demonstrated in Ch. 3, it is generally possible to write down interacting many-fermion Hamiltonians with complex phase factors that have strictly positive definite DQMC weights – the identified flat-band model represents an example. Based on our spadework, it would be worth exploring, if the novel semigroup approach to the sign-problem can be utilized to construct alternative models for the AFM QCP that allow for a magnetic field and yet are amenable to efficient DQMC. If successful, this could represent an important milestone for numerical simulations of metallic quantum criticality.

Apart from the sign-problem, a limiting aspect of DQMC is its high computational cost. As shown in Ch. 4, the generation of Markov chains scales as $O(\beta N^3)$, preventing studies of large systems at low temperatures. Similarly, measuring fermionic correlations based on time-displaced Green’s functions is almost equally expensive – a fact that often goes unmentioned – and significantly hampers desirable large-scale parameter studies of, for instance, the fermion self-energy. In this thesis, we have showcased that combining DQMC with modern machine learning represents a powerful way to partially circumvent this issue and probe important aspects of metallic quantum criticality. In particular, we have demonstrated that a novel quantum loop topography approach based on the learning of longitudinal transport correlations can reliably detect the onset of superconductivity and SDW order in the $O(2)$ symmetric spin-fermion model. Moreover, the relatively small computational footprint of this technique has enabled us to study changes in transport across wide parameter regions and, in this way, to map out an extended non-Fermi liquid regime. The latter has so-far been elusive, even in large-scale DQMC studies of the order of 10 million CPU hours. Remarkably, despite training the ML architecture at only two or three points, we could reproduce many essential aspects of the established phase diagram of the spin-fermion model, including the location of the QCP [31, 32].

Looking forward, a practical extension of work is to apply QLT to other effective models of metallic quantum criticality, that is beyond the Ising-nematic theory considered in Ch. 7, and systematically compare the features of different metallic QCPs. Furthermore, it should be viable to use QLT to identify other states without traditional representation that are defined by non-trivial transport properties. In view of the fact that ML methods, despite attempts to physically anchor them, are always accompanied by a certain opacity, it would be worthwhile to rigorously assess our results by means of orthogonal numerical techniques.

An interesting candidate could be the memory matrix approach introduced in Ref. [210].

Parallel to our work, various machine learning approaches for internally speeding up DQMC simulations have been put forward [219, 220, 222, 223, 261]. Essentially, the idea is to train a ML architecture based on small Markov chains of configurations such that it approximates the desired distribution function and may be used to propose efficient global update. While we have found that for the spin-fermion model, direct applications of these strategies do not significantly improve statistical convergence – continuous $O(3)$ order parameters seem to be more difficult to handle than \mathbb{Z}_2 counterparts – exploring this path further seems promising to reach larger system sizes in DQMC.

As part of this thesis, we have developed state-of-the-art DQMC codes and various software libraries focused on specific functionalities, such as the intricate stabilization of Green’s function computations [S3]. To improve the transparency of the presented research as well as to facilitate further numerical studies on quantum critical metals, we have published these implementations under open source licenses, see e.g. Refs. [S1–S4]. In contrast to existing frameworks [184, 185], our codes are written in a “high-level” programming language and, in this way, will hopefully make DQMC accessible to a wider audience.

A | Hubbard-Stratonovich transformation of the spin-fermion model

If we drop the gradient and quartic terms in S_ϕ , Eq. (1.4), the parts of the spin-fermion model, Eq. (1.9), involving the bosonic field ϕ read

$$S_\phi + S_\lambda = \int_0^\beta d\tau \int d^2x \left[\frac{r}{2} \phi^2 + \lambda \sum_{s,s'} e^{i\mathbf{Q}\cdot\mathbf{r}} \phi_{\mathbf{r}} \cdot \psi_{s\mathbf{r}}^\dagger \sigma_{ss'} \psi_{s'\mathbf{r}} \right]. \quad (\text{A.1})$$

For convenience we temporarily set $\mathbf{Q} = 0^{106}$. For every imaginary time τ and every position \mathbf{r} we may now consider the inverse Hubbard-Stratonovich transformation

$$\begin{aligned} \int_{-\infty}^{\infty} dx \exp \left(-\frac{1}{2} \alpha x^2 - Ax \right) &= \exp \left(\frac{A^2}{2\alpha} \right) \int_{-\infty}^{\infty} dx \exp \left[-\frac{1}{2} \alpha \left(x + \frac{A}{\alpha} \right)^2 \right] \\ &= \exp \left(\frac{A^2}{2\alpha} \right) \int_{-\infty}^{\infty} dy \exp \left(-\frac{1}{2} \alpha y^2 \right) \\ &= \sqrt{\frac{2\pi}{\alpha}} \exp \left(\frac{A^2}{2\alpha} \right) \end{aligned} \quad (\text{A.2})$$

where we have substituted $y = x + A/\alpha$ in the second line. Identifying $\alpha = r$ and $A = -\lambda \sum_{s,s'} \phi \cdot \psi_{s\mathbf{r}}^\dagger \sigma_{ss'} \psi_{s'\mathbf{r}}$ as well as absorbing the irrelevant numerical prefactor into the integration measure we find that the ϕ integration in the partition function \mathcal{Z}_{SF} , Eq. (1.10), gives

$$\begin{aligned} \int D(\phi, \psi, \psi^\dagger) e^{-S_\phi - S_\lambda - S_\psi} &= \int D(\psi, \psi^\dagger) e^{-S_\psi} \times \\ &\times \exp \left(- \underbrace{\left(-\frac{\lambda^2}{2r} \right)}_{\tilde{J} \text{ in Eq. (1.14)}} \left(\sum_{s,s'} \phi \cdot \psi_{s\mathbf{r}}^\dagger \sigma_{ss'} \psi_{s'\mathbf{r}} \right)^2 \right). \end{aligned} \quad (\text{A.3})$$

In Hamiltonian language this generated four-fermion interaction reads

$$H_J = \tilde{J} \sum_{n=1}^3 \sum_{iss'} \left(c_{is}^\dagger [\sigma_n]_{s,s'} c_{is'} \right)^2 = \tilde{J} \sum_i \mathbf{S}_i^2, \quad (\text{A.4})$$

¹⁰⁶In the two band version of the model no explicit term $e^{i\mathbf{Q}\cdot\mathbf{r}}$ is necessary. Instead one may simply shift the dispersion of one of the bands by $\mathbf{Q} = (\pi, \pi)$.

where we have used the definition of the spin operators, Eq. (1.15). In combination with the non-interacting fermion contribution S_ψ the total Hamiltonian is therefore given by

$$H_{tJ} = - \sum_{ijs} \left(t_{ij} c_{is}^\dagger c_{js} + \text{h.c.} \right) + \tilde{J} \sum_i \mathbf{S}_i^2 - \mu \sum_i n_i, \quad (\text{A.5})$$

which, after reintroduction of the ordering wave-vector term corresponds to Eq. (1.14).

B Evaluation of the fermion trace

In the following, we prove the identity

$$\text{Tr} \left[\prod_l e^{-\Delta\tau H(l)} \right] = \det \left(1 + \prod_l e^{-\Delta\tau h(l)} \right), \quad (\text{B.1})$$

for a non-interacting Hamiltonian $H(l)$ to fill the gap between Eqs. (2.18) and (2.20) of the main text. Closely following Ref. [90], we split the proof into two parts: First, we show that a time product of exponentials of $H(l)$ may be grouped into a single exponential of a bilinear form \tilde{H} . Then, we prove the identity

$$\text{Tr} [e^{-\Delta\tau \tilde{H}}] = \det (1 + e^{-\Delta\tau \tilde{H}}). \quad (\text{B.2})$$

Grouping products of exponentials

We start with a generic bilinear Hamiltonian

$$H = \sum_{ij} h_{ij} c_i^\dagger c_j. \quad (\text{B.3})$$

In the Heisenberg picture, the time dependence of a fermion creation operator is given by

$$c_i^\dagger(\tau) = \sum_j [e^{-\tau h}]_{ji} c_j^\dagger. \quad (\text{B.4})$$

Performing a basis transformation to operators

$$c_\alpha^\dagger = \sum_i \langle i|\alpha \rangle c_i^\dagger, \quad (\text{B.5})$$

one finds,

$$c_\alpha^\dagger(\Delta\tau) = \sum_{ij} [e^{-\Delta\tau h}]_{ji} \langle i|\alpha \rangle c_j^\dagger \quad (\text{B.6})$$

$$\equiv \sum_j U_{j\alpha} c_j^\dagger, \quad (\text{B.7})$$

with $U_{j\alpha}$ defined by the previous line. Now, considering a discretized imaginary time and $H(l)$ with time-dependent coefficients $h_{ij}(l)$, one has

$$c_\alpha^\dagger = \sum_{ij} \left[\prod_l e^{-\Delta\tau h(l)} \right]_{ji} \langle i|\alpha \rangle c_j^\dagger. \quad (\text{B.8})$$

Defining a coefficient matrix \tilde{h} element-wise through the relation

$$\left[e^{-\Delta\tau\tilde{h}} \right]_{ji} = \left[\prod_{\ell} e^{-\Delta\tau h(\ell)} \right]_{ji}, \quad (\text{B.9})$$

and considering the product of exponentials

$$B = \prod_l e^{-\Delta\tau H(l)}, \quad (\text{B.10})$$

one obtains the identity [90]

$$B c_{\alpha}^{\dagger} B^{-1} = \sum_{ij} \left[e^{-\Delta\tau\tilde{h}} \right]_{ji} \langle i|\alpha\rangle c_j^{\dagger}. \quad (\text{B.11})$$

This indicates that B effectively acts as a single exponential of the effective bilinear Hamiltonian \tilde{H} with coefficients \tilde{h}_{ij} .

Evaluating the trace

Using the result above, we are now in the position to prove Eq. (B.1) by evaluating the fermion trace in the eigenbasis of \tilde{H} , i.e.

$$\tilde{H} = \sum_{\alpha} \epsilon_{\alpha} c_{\alpha}^{\dagger} c_{\alpha}. \quad (\text{B.12})$$

This way, one obtains the following chain of equalities, [90]

$$\begin{aligned} \text{Tr} \left[\prod_l e^{-\Delta\tau \mathcal{H}(l)} \right] &= \text{Tr} \left[e^{-\Delta\tau \tilde{\mathcal{H}}} \right] \\ &= \text{Tr} \left[e^{-\Delta\tau \sum_{\alpha} \epsilon_{\alpha} c_{\alpha}^{\dagger} c_{\alpha}} \right] \\ &= \text{Tr} \left[\prod_{\alpha} e^{-\Delta\tau \epsilon_{\alpha} c_{\alpha}^{\dagger} c_{\alpha}} \right] \\ &= \prod_{\alpha} \left(1 + e^{-\Delta\tau \epsilon_{\alpha}} \right) \\ &= \det \left(1 + e^{-\Delta\tau \tilde{h}} \right) \\ &= \det \left(1 + \prod_l e^{-\Delta\tau h(l)} \right). \end{aligned} \quad (\text{B.13})$$

C Locating the isotropic AFM QCP

Given the $O(\beta N^3)$ asymptotic scaling of DQMC simulations, a brute force search for the QCP is associated with a high computational cost. Scanning a large range of tuning parameter values at low temperatures is therefore a suboptimal approach. Instead, a more efficient method for estimating the position of the quantum critical point, r_c , is to track a finite temperature magnetic transition down to $T = 0$ [31, 32]. Although the spin-fermion model with a three-component order parameter does not spontaneously break the continuous $O(3)$ symmetry at finite temperatures (Mermin-Wagner theorem) one can nonetheless use this strategy to obtain a rough estimate of the QCP location. Concretely, we extract the scaling exponent α from a finite size scaling analysis, $\chi \sim L^\alpha$, and heuristically define a crossover regime by the condition $0.5 < \alpha < 1.5$, see Fig. C.1 for the case $\lambda = 2$. Successively following the crossover down to the absolute zero of temperature and, for a given temperature, taking the mean of the tuning parameter boundaries (left and right) as an estimate, we obtain the approximate QCP location $3.55 < r_c \approx 3.85 < 4.15$.




Figure C.1: Heuristic crossover regime (grey) in the $\lambda = 2$ phase diagram, Fig. 5.3. The median of the crossover regime (dark grey line) extrapolated to $T = 0$ is an estimate for the QCP location. Within the limits of this procedure, one finds $3.55 < r_c \approx 3.85 < 4.15$.




Figure C.2: Tuning parameter dependence of the inverse magnetic susceptibility, Eq. (5.21), close to the quantum critical point in the case of almost local nesting near the hot spots. The inverse temperature is $\beta = 40$ and $\lambda = 1.0$. We show a linear fit of the data (black line) with root $r_c \approx 1.41$, an estimate for the location of the QCP.

Building upon this knowledge, we can improve this estimate by resolving the tuning parameter dependence of the inverse magnetic susceptibility χ^{-1} across the crossover regime. As shown in Fig. 5.4 of the main text, we observe a linear dependence which we can extrapolate to $\chi^{-1} = 0$, a diverging susceptibility. For our model parameters, we find $r_c \approx 3.8$. Applying the same technique to the spin-fermion model with an almost locally nested Fermi surface, Sec. 5.4, we find $r_c \approx 1.41$. The corresponding tuning parameter dependence is shown in Fig. C.2.

D | Supplement to QLT classification

Data grids

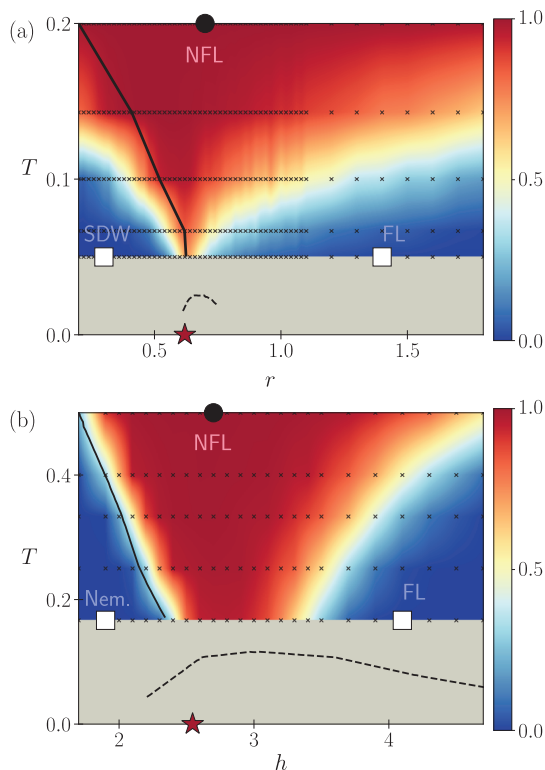


Figure D.1: Data interpolation for both SDW and nematic models. Both panels show the corresponding plots from the main text, Fig. 7.6 and Fig. 7.11 with the data points overlaid as black crosses. In order to make the color plots, a bilinear spline is used to interpolate the values in between the available data points. Modified from the supplement of Ref. [P4].

Robustness of classification for nematic model

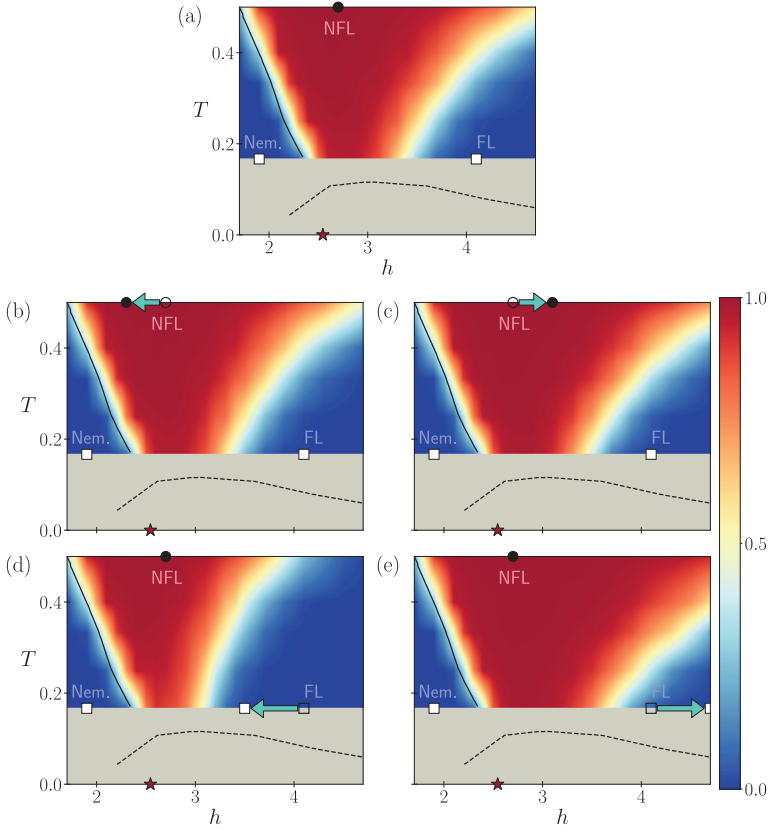


Figure D.2: Robustness of the QLT-based **three-point classification** for the nematic model, Eq. (7.7). Panel (a) is the color plot in the main text, Fig. 7.11: the nematic ordered training point is located at $h = 1.9, T = 0.167$, the non-Fermi liquid training point at $h = 2.7, T = 0.5$, and the disordered Fermi liquid training point at $h = 4.1, T = 0.167$. Panels (b) and (c) move the non-Fermi liquid training point to $h = 2.3$ and $h = 3.1$ respectively at the same temperature. Panels (d) and (e) move the disordered Fermi liquid training point to $h = 3.5$ and $h = 4.7$ respectively at the same temperature. Modified from the supplement of Ref. [P4].

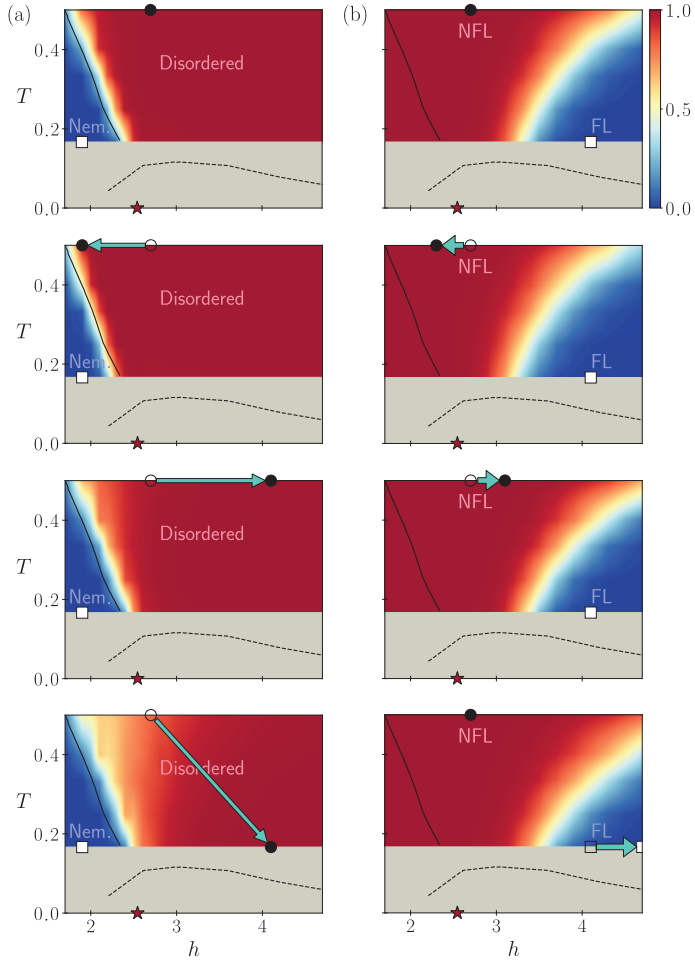


Figure D.3: Robustness of the **binary classification** for the nematic model, Eq. (7.7). The panels show that the discrimination of Nematic order and disorder (a) as well non-Fermi liquid and Fermi liquid transport (b) varies only modestly when shifting the respective training points (cyan arrows). Modified from the supplement of Ref. [P4].

Bibliography

- [P1] C. Bauer, Y. Schattner, S. Trebst, and E. Berg, Hierarchy of energy scales in an $O(3)$ symmetric antiferromagnetic quantum critical metal: A Monte Carlo study, *Physical Review Research* **2** (2), 023008, 2020.
- [P2] C. Bauer, Fast and stable determinant quantum Monte Carlo, *SciPost Physics Core* **2** (2), 11, 2020.
- [P3] Y. Zhang, C. Bauer, P. Broecker, S. Trebst, and E.-A. Kim, Probing transport in quantum many-fermion simulations via quantum loop topography, *Physical Review B* **99** (16), 161120, 2019.
- [P4] G. T. Driskell, S. Lederer, C. Bauer, S. Trebst, and E.-A. Kim, Observation of non-Fermi liquid physics in a quantum critical metal via quantum loop topography, 2020, (submitted to *Nature Communications*). arXiv: 2007.07898.
- [P5] C. Bauer and S. Trebst, Machine Learning Transport Properties in Quantum Many-Fermion Simulations, in *NIC Symposium 2020*, Vol. 50, Forschungszentrum Jülich GmbH Zentralbibliothek, Verlag, 2020.
- [S1] C. Bauer, *julia-sdw-dqmc: Determinant quantum Monte Carlo in Julia*, Version 1.0, Zenodo, 2020.
- [S2] C. Bauer and F. Freyer, *BinningAnalysis.jl: Standard error estimation tools*, Version 0.3.3, Zenodo, 2020.
- [S3] C. Bauer, *StableDQMC.jl: Numerical stabilization routines for determinant quantum Monte Carlo*, Version 0.1.3, Zenodo, 2020.
- [S4] C. Bauer, *MonteCarlo.jl: Classical and quantum Monte Carlo simulations in Julia*, Version 0.1, Zenodo, 2020.
- [1] S. Sachdev, *Quantum Phase Transitions*. Cambridge University Press, 2011.
- [2] E. Berg, S. Lederer, Y. Schattner, and S. Trebst, Monte Carlo Studies of Quantum Critical Metals, *Annual Review of Condensed Matter Physics* **10** (1), 63–84, 2019.
- [3] A. Abanov and A. V. Chubukov, Spin-Fermion Model near the Quantum Critical Point: One-Loop Renormalization Group Results, *Physical Review Letters* **84** (24), 5608–5611, 2000.
- [4] A. Abanov, A. V. Chubukov, and J. Schmalian, Quantum-critical theory of the spin-fermion model and its application to cuprates: Normal state analysis, *Advances in Physics* **52** (3), 119–218, 2003.
- [5] A. V. Chubukov, D. Pines, and J. Schmalian, A spin fluctuation model for d-wave superconductivity, 2002. arXiv: 1904.12872.

- [6] P. Lunts, A. Schlieﬀ, and S.-S. Lee, Emergence of a control parameter for the antiferromagnetic quantum critical metal, *Physical Review B* **95** (24), 245109, 2017.
- [7] A. Schlieﬀ, P. Lunts, and S.-S. Lee, Exact Critical Exponents for the Antiferromagnetic Quantum Critical Metal in Two Dimensions, *Physical Review X* **7** (2), 021010, 2017.
- [8] M. Vojta, Quantum phase transitions, *Reports on Progress in Physics* **66** (12), 2069–2110, 2003.
- [9] L. D. Carr, *Understanding Quantum Phase Transitions*. CRC Press, 2011.
- [10] S. L. Sondhi, S. M. Girvin, J. P. Carini, and D. Shahar, Continuous quantum phase transitions, *Review of Modern Physics* **69**, 315–333, 1997.
- [11] L. P. Kadanoff, Phases of Matter and Phase Transitions; From Mean Field Theory to Critical Phenomena, *Journal of Physics: Condensed Matter* **12**, 1496, 2009.
- [12] L. P. Kadanoff, More is the Same; Phase Transitions and Mean Field Theories, *Journal of Statistical Physics* **137** (5), 777–797, 2009.
- [13] L. P. Kadanoff, Theories of Matter: Infinities and Renormalization, 2010. arXiv: 1002.2985.
- [14] P. W. Anderson, In praise of unstable fixed points: the way things actually work, *Physica B: Condensed Matter* **318** (1), 28–32, 2002.
- [15] S.-S. Lee, Recent Developments in Non-Fermi Liquid Theory, *Annual Review of Condensed Matter Physics* **9** (1), 227–244, 2018.
- [16] L. Taillefer, Scattering and Pairing in Cuprate Superconductors, *Annual Review of Condensed Matter Physics* **1** (1), 51–70, 2010.
- [17] A. Altland and B. Simons, *Condensed Matter Field Theory*, 2nd Ed. Cambridge University Press, 2010.
- [18] N. Mott, *Metal-insulator transitions*. CRC Press, 2004.
- [19] P. Kopietz, L. Bartosch, and F. Schütz, *Introduction to the functional renormalization group*. Springer, 2010.
- [20] D. Ashton, The Renormalisation Group, <https://www.youtube.com/watch?v=MxRddFrEnPc>, [Online; accessed 24-June-2020], 2012.
- [21] D. Ashton, Universality at the critical point, <https://www.youtube.com/watch?v=Kd4UvhUsBAU>, [Online; accessed 24-June-2020], 2011.
- [22] N. D. Mermin and H. Wagner, Absence of ferromagnetism or antiferromagnetism in one- or two-dimensional isotropic Heisenberg models, *Physical Review Letters* **17** (22), 1133–1136, 1966.
- [23] J. W. Negele and H. Orland, *Quantum many-particle systems*. CRC Press, 2018.
- [24] E. Berg, Critical Metals: Lessons from quantum Monte Carlo studies, <https://physics.stanford.edu/events/video-erez-bergs-critical-metals-lessons-quantum-monte-carlo-studies>, [Online; Colloquium given at Stanford University], 2018.

-
- [25] Y. Schattner, *Quantum phase transitions in metals*, Ph.D. dissertation, Weizmann Institute of Science, 2017, (unpublished).
 - [26] D. Chowdhury and E. Berg, The unreasonable effectiveness of Eliashberg theory for pairing of non-Fermi liquids, *Annals of Physics* **417** (1990), 168125, 2020.
 - [27] X. Y. Xu, K. Sun, Y. Schattner, E. Berg, and Z. Y. Meng, Non-Fermi-liquid at (2+1)d ferromagnetic quantum critical point, *Physical Review X* **7** (3), 031058, 2017.
 - [28] A. V. Chubukov and D. L. Maslov, Spin conservation and fermi liquid near a ferromagnetic quantum critical point, *Physical Review Letters* **103** (21), 216401, 2009.
 - [29] A. V. Chubukov, C. Pépin, and J. Rech, Instability of the Quantum-Critical Point of Itinerant Ferromagnets, *Physical Review Letters* **92** (14), 147003, 2004.
 - [30] Z. X. Li, F. Wang, H. Yao, and D. H. Lee, Nature of the effective interaction in electron-doped cuprate superconductors: A sign-problem-free quantum Monte Carlo study, *Physical Review B* **95** (21), 214505, 2017.
 - [31] M. H. Gerlach, Y. Schattner, E. Berg, and S. Trebst, Quantum critical properties of a metallic spin-density-wave transition, *Physical Review B* **95** (3), 035124, 2017.
 - [32] Y. Schattner, M. H. Gerlach, S. Trebst, and E. Berg, Competing Orders in a Nearly Antiferromagnetic Metal, *Physical Review Letters* **117** (9), 097002, 2016.
 - [33] Z. H. Liu, G. Pan, X. Y. Xu, K. Sun, and Z. Y. Meng, Itinerant quantum critical point with fermion pockets and hotspots, *Proceedings of the National Academy of Sciences* **116** (34), 16760–16767, 2019.
 - [34] M. A. Metlitski and S. Sachdev, Quantum phase transitions of metals in two spatial dimensions. II. Spin density wave order, *Physical Review B* **82** (7), 075128, 2010.
 - [35] A. J. Millis, Nearly antiferromagnetic fermi liquids: An analytic eliashberg approach, *Physical Review B* **45** (22), 13047–13054, 1992.
 - [36] A. Abanov and A. Chubukov, Anomalous Scaling at the Quantum Critical Point in Itinerant Antiferromagnets, *Physical Review Letters* **93** (25), 255702, 2004.
 - [37] S. Sachdev, Antiferromagnetism and Superconductivity, <http://qpt.physics.harvard.edu/talks/jerusalem16a.pdf>, [Online; Talk given at the 34th Jerusalem Winter School in Theoretical Physics, New Horizons in Quantum Matter, Jerusalem], 2016.
 - [38] S. Sachdev and B. Keimer, Quantum criticality, *Physics Today* **64** (2), 29–35, 2011.
 - [39] Y. Schattner, S. Lederer, S. A. Kivelson, and E. Berg, Ising Nematic Quantum Critical Point in a Metal: A Monte Carlo Study, *Physical Review X* **6** (3), 031028, 2016.

- [40] S. Lederer, Y. Schattner, E. Berg, and S. A. Kivelson, Superconductivity and non-Fermi liquid behavior near a nematic quantum critical point, *Proceedings of the National Academy of Sciences* **114** (19), 4905–4910, 2017.
- [41] M. A. Metlitski, D. F. Mross, S. Sachdev, and T. Senthil, Cooper pairing in non-Fermi liquids, *Physical Review B* **91** (11), 2015.
- [42] T. Moriya, Y. Takahashi, and K. Ueda, Antiferromagnetic Spin Fluctuations and Superconductivity in Two-Dimensional Metals -A Possible Model for High T_c Oxides, *Journal of the Physical Society of Japan* **59** (8), 2905–2915, 1990.
- [43] M. A. Metlitski and S. Sachdev, Instabilities near the onset of spin density wave order in metals, *New Journal of Physics* **12** (10), 105007, 2010.
- [44] Y. Wang and A. Chubukov, Charge-density-wave order with momentum $(2Q,0)$ and $(0,2Q)$ within the spin-fermion model: Continuous and discrete symmetry breaking, preemptive composite order, and relation to pseudogap in hole-doped cuprates, *Physical Review B* **90** (3), 035149, 2014.
- [45] S. Sachdev, E. Berg, S. Chatterjee, and Y. Schattner, Spin density wave order, topological order, and Fermi surface reconstruction, *Physical Review B* **94** (11), 115147, 2016.
- [46] X. Wang, Y. Wang, Y. Schattner, E. Berg, and R. M. Fernandes, Fragility of Charge Order Near an Antiferromagnetic Quantum Critical Point, *Physical Review Letters* **120** (24), 247002, 2018.
- [47] K. B. Efetov, H. Meier, and C. Pepin, Pseudogap state near a quantum critical point, *Nature Physics* **9** (7), 442–446, 2013.
- [48] Y. Wang, D. F. Agterberg, and A. Chubukov, Interplay between pair- and charge-density-wave orders in underdoped cuprates, *Physical Review B* **91** (11), 115103, 2015.
- [49] M. A. Metlitski and S. Sachdev, Quantum phase transitions of metals in two spatial dimensions. I. Ising-nematic order, *Physical Review B* **82** (7), 075127, 2010.
- [50] J. A. Hertz, Quantum critical phenomena, *Physical Review B* **14** (3), 1165–1184, 1976.
- [51] Z. H. Liu, X. Y. Xu, Y. Qi, K. Sun, and Z. Y. Meng, Itinerant quantum critical point with frustration and a non-Fermi liquid, *Physical Review B* **98** (4), 045116, 2018.
- [52] P. T. Dumitrescu, M. Serbyn, R. T. Scalettar, and A. Vishwanath, Superconductivity and nematic fluctuations in a model of doped FeSe monolayers: Determinant quantum Monte Carlo study, *Physical Review B* **94** (15), 155127, 2016.
- [53] S. Gazit, M. Randeria, and A. Vishwanath, Emergent Dirac fermions and broken symmetries in confined and deconfined phases of Z_2 gauge theories, *Nature Physics* **13** (5), 484–490, 2017.
- [54] C. Chen, X. Y. Xu, Y. Qi, and Z. Y. Meng, Metal to orthogonal metal transition, *Chinese Physics Letters* **37** (4), 047103, 2020.

- [55] F. F. Assaad and T. Grover, Simple fermionic model of deconfined phases and phase transitions, *Physical Review X* **6** (4), 041049, 2016.
- [56] X. Y. Xu, Y. Qi, L. Zhang, F. F. Assaad, C. Xu, and Z. Y. Meng, Monte Carlo Study of Lattice Compact Quantum Electrodynamics with Fermionic Matter: The Parent State of Quantum Phases, *Physical Review X* **9** (2), 021022, 2019.
- [57] T. Shibauchi, A. Carrington, and Y. Matsuda, A quantum critical point lying beneath the superconducting dome in iron pnictides, *Annual Review of Condensed Matter Physics* **5** (1), 113–135, 2014.
- [58] J. Bardeen, L. N. Cooper, and J. R. Schrieffer, Microscopic Theory of Superconductivity, *Physical Review* **106** (1), 162–164, 1957.
- [59] B. Keimer, S. A. Kivelson, M. R. Norman, S. Uchida, and J. Zaanen, From quantum matter to high-temperature superconductivity in copper oxides. *Nature* **518** (7538), 179–86, 2015.
- [60] A. P. Drozdov, M. I. Erements, I. A. Troyan, V. Ksenofontov, and S. I. Shylin, Conventional superconductivity at 203 kelvin at high pressures in the sulfur hydride system, *Nature* **525** (7567), 73–76, 2015.
- [61] J. G. Bednorz and K. A. Müller, Possible high T_c superconductivity in the Ba-La-Cu-O system, *Zeitschrift für Physik B Condensed Matter* **64** (2), 189–193, 1986.
- [62] Y. Kamihara, H. Hiramatsu, M. Hirano, R. Kawamura, H. Yanagi, T. Kamiya, and H. Hosono, Iron-Based Layered Superconductor: LaOFeP, *Journal of the American Chemical Society* **128** (31), 10012–10013, 2006.
- [63] D. J. Scalapino, A common thread: The pairing interaction for unconventional superconductors, *Reviews of Modern Physics* **84** (4), 1383–1417, 2012.
- [64] N. P. Armitage, P. Fournier, and R. L. Greene, Progress and perspectives on electron-doped cuprates, *Reviews of Modern Physics* **82** (3), 2421–2487, 2010.
- [65] H. V. Löhneysen, A. Rosch, M. Vojta, and P. Wölfle, Fermi-liquid instabilities at magnetic quantum phase transitions, *Reviews of Modern Physics* **79** (3), 1015–1075, 2007.
- [66] P. Gegenwart, Q. Si, and F. Steglich, Quantum criticality in heavy-fermion metals, *Nature Physics* **4** (3), 186, 2008.
- [67] M. Dressel, Quantum criticality in organic conductors? Fermi liquid versus non-Fermi-liquid behaviour, *Journal of Physics: Condensed Matter* **23** (29), 293201, 2011.
- [68] A. G. Lebed, *The physics of organic superconductors and conductors*. Springer, 2008.
- [69] R. M. Fernandes, S. Maiti, P. Wölfle, and A. V. Chubukov, How Many Quantum Phase Transitions Exist Inside the Superconducting Dome of the Iron Pnictides? *Physical Review Letters* **111** (5), 057001, 2013.
- [70] J. A. N. Bruin, H. Sakai, R. S. Perry, and A. P. Mackenzie, Similarity of Scattering Rates in Metals Showing T-Linear Resistivity, *Science* **339** (6121), 804–807, 2013.

- [71] H. Matsui, T. Takahashi, T. Sato, K. Terashima, H. Ding, T. Uefuji, and K. Yamada, Evolution of the pseudogap across the magnet-superconductor phase boundary of $\text{Nd}_{2-x}\text{Ce}_x\text{CuO}_4$, *Physical Review B* **75** (22), 224514, 2007.
- [72] E. G. Moon and S. Sachdev, Quantum critical point shifts under superconductivity: Pnictides and cuprates, *Physical Review B* **82** (10), 104516, 2010.
- [73] S. Sachdev, Where is the quantum critical point in the cuprate superconductors? *Physica Status Solidi B* **247** (3), 537–543, 2010.
- [74] E. Fradkin, S. A. Kivelson, and J. M. Tranquada, Colloquium : Theory of intertwined orders in high temperature superconductors, *Reviews of Modern Physics* **87** (2), 457–482, 2015.
- [75] B. J. Ramshaw, S. E. Sebastian, R. D. McDonald, *et al.*, Quasiparticle mass enhancement approaching optimal doping in a high- T_c superconductor, *Science* **348** (6232), 317–320, 2015.
- [76] G. Knebel, D. Aoki, and J. Flouquet, Antiferromagnetism and superconductivity in cerium based heavy-fermion compounds, *Comptes Rendus Physique* **12** (5), 542–566, 2011.
- [77] M. V. Kartsovnik, T. Helm, C. Putzke, *et al.*, Fermi surface of the electron-doped cuprate superconductor $\text{Nd}_{2-x}\text{Ce}_x\text{CuO}_4$ probed by high-field magnetotransport, *New Journal of Physics* **13** (1), 015001, 2011.
- [78] P. Monthoux, D. Pines, and G. G. Lonzarich, Superconductivity without phonons, *Nature* **450** (7173), 1177–1183, 2007.
- [79] E. M. Motoyama, G. Yu, I. M. Vishik, O. P. Vajk, P. K. Mang, and M. Greven, Spin correlations in the electron-doped high-transition-temperature superconductor $\text{Nd}_{2-x}\text{Ce}_x\text{CuO}_{4\pm\delta}$, *Nature* **445** (7124), 186–189, 2007.
- [80] S. Sachdev, Emergent gauge fields and the high-temperature superconductors, *Philosophical Transactions of the Royal Society A* **374** (2075), 20150248, 2016.
- [81] E. Berg, M. A. Metlitski, and S. Sachdev, Sign-Problem-Free Quantum Monte Carlo of the Onset of Antiferromagnetism in Metals, *Science* **338** (6114), 1606–1609, 2012.
- [82] A. J. Millis, Effect of a nonzero temperature on quantum critical points in itinerant fermion systems, *Physical Review B* **48** (10), 7183–7196, 1993.
- [83] H. Kleinert and V. Schulte-Frohlinde, *Critical Properties of ϕ^4 - Theories*. World Scientific, 2001.
- [84] S. Chakravarty, B. I. Halperin, and D. R. Nelson, Two-dimensional quantum Heisenberg antiferromagnet at low temperatures, *Physical Review B* **39** (4), 2344–2371, 1989.
- [85] A. V. Chubukov, S. Sachdev, and J. Ye, Theory of two-dimensional quantum Heisenberg antiferromagnets with a nearly critical ground state, *Physical Review B* **49** (17), 11919–11961, 1994.
- [86] K. G. Wilson and M. E. Fisher, Critical Exponents in 3.99 Dimensions, *Physical Review Letters* **28** (4), 240–243, 1972.

-
- [87] L. Ornstein and F. Zernike, Die linearen Dimensionen der Dichteschwankungen, *Physikalische Zeitschrift* **19**, 134, 1918.
 - [88] P. M. Chaikin and T. C. Lubensky, *Principles of Condensed Matter Physics*. Cambridge University Press, 1995.
 - [89] P. Kopietz, *Bosonization of interacting fermions in arbitrary dimensions*. Springer, 2008.
 - [90] R. R. dos Santos, Introduction to quantum Monte Carlo simulations for fermionic systems, *Brazilian Journal of Physics* **33** (1), 36–54, 2003.
 - [91] E. Fradkin, *Field Theories of Condensed Matter Physics*. Cambridge University Press, 2013.
 - [92] H. Tasaki, The Hubbard model - an introduction and selected rigorous results, *Journal of Physics: Condensed Matter* **10** (20), 4353–4378, 1998.
 - [93] J. E. Hirsch, Two-dimensional Hubbard model: Numerical simulation study, *Physical Review B* **31** (7), 4403–4419, 1985.
 - [94] J. E. Hirsch and S. Tang, Antiferromagnetism in the two-dimensional Hubbard model, *Physical Review Letters* **62** (5), 591, 1989.
 - [95] C. N. Varney, C.-R. Lee, Z. J. Bai, S. Chiesa, M. Jarrell, and R. T. Scalettar, Quantum Monte Carlo study of the two-dimensional fermion Hubbard model, *Physical Review B* **80** (7), 075116, 2009.
 - [96] S. White, D. Scalapino, R. Sugar, E. Loh, J. Gubernatis, and R. Scalettar, Numerical study of the two-dimensional Hubbard model, *Physical Review B* **40** (1), 506–516, 1989.
 - [97] M. Vojta and S. Sachdev, Charge order, superconductivity, and a global phase diagram of doped antiferromagnets, *Physical Review Letters* **83** (19), 3916–3919, 1999.
 - [98] P. A. Lee, N. Nagaosa, and X.-G. Wen, Doping a Mott insulator: Physics of high-temperature superconductivity, *Reviews of Modern Physics* **78** (1), 17–85, 2006.
 - [99] E. Dagotto, A. Nazarenko, and A. Moreo, Antiferromagnetic and van Hove Scenarios for the Cuprates: Taking the Best of Both Worlds, *Physical Review Letters* **74** (2), 310–313, 1995.
 - [100] S.-C. Zhang, A Unified Theory Based on SO(5) Symmetry of Superconductivity and Antiferromagnetism, *Science* **275** (5303), 1089–1096, 1997.
 - [101] J. E. Hirsch, S. Tang, E. Loh, and D. J. Scalapino, Pairing interaction in two-dimensional CuO₂, *Physical Review Letters* **60** (16), 1668–1671, 1988.
 - [102] P. W. Anderson, The Resonating Valence Bond State in La₂CuO₄ and Superconductivity, *Science* **235** (1), 1196–1198, 1986.
 - [103] S. A. Maier and P. Strack, Universality in antiferromagnetic strange metals, *Physical Review B* **93** (16), 165114, 2016.
 - [104] C. M. Varma, L. Zhu, and A. Schröder, Quantum critical response function in quasi-two-dimensional itinerant antiferromagnets, *Physical Review B* **92** (15), 155150, 2015.

- [105] C. M. Varma, Quantum Criticality in Quasi-Two-Dimensional Itinerant Antiferromagnets, *Physical Review Letters* **115** (18), 186405, 2015.
- [106] A. L. Fitzpatrick, S. Kachru, J. Kaplan, and S. Raghu, Non-Fermi-liquid behavior of large- N_B quantum critical metals, *Physical Review B* **89** (16), 165114, 2014.
- [107] J. Lee, P. Strack, and S. Sachdev, Quantum criticality of reconstructing Fermi surfaces in antiferromagnetic metals, *Physical Review B* **87** (4), 045104, 2013.
- [108] D. F. Mross, J. McGreevy, H. Liu, and T. Senthil, Controlled expansion for certain non-Fermi-liquid metals, *Physical Review B* **82** (4), 045121, 2010.
- [109] D. Dalidovich and S.-S. Lee, Perturbative non-Fermi liquids from dimensional regularization, *Physical Review B* **88** (24), 245106, 2013.
- [110] S. Sur and S.-S. Lee, Quasilocal strange metal, *Physical Review B* **91** (12), 125136, 2015.
- [111] T. Helm, M. V. Kartsovnik, M. Bartkowiak, N. Bittner, M. Lambacher, A. Erb, J. Wosnitza, and R. Gross, Evolution of the Fermi Surface of the Electron-Doped High-Temperature Superconductor $\text{Nd}_{2-x}\text{Ce}_x\text{CuO}_4$ Revealed by Shubnikov-de Haas Oscillations, *Physical Review Letters* **103** (15), 157002, 2009.
- [112] X. Wang, Y. Schattner, E. Berg, and R. M. Fernandes, Superconductivity mediated by quantum critical antiferromagnetic fluctuations: The rise and fall of hot spots, *Physical Review B* **95** (17), 174520, 2017.
- [113] A. Abanov, A. V. Chubukov, and A. M. Finkel'stein, Coherent vs incoherent pairing in 2D systems near magnetic instability, *Europhysics Letters* **54** (4), 488–494, 2001.
- [114] Y. Wang and A. V. Chubukov, Superconductivity at the Onset of Spin-Density-Wave Order in a Metal, *Physical Review Letters* **110** (12), 127001, 2013.
- [115] A. Abanov, A. V. Chubukov, and M. R. Norman, Gap anisotropy and universal pairing scale in a spin-fluctuation model of cuprate superconductors, *Physical Review B* **78** (22), 220507, 2008.
- [116] K. Miyake, S. Schmitt-Rink, and C. M. Varma, Spin-fluctuation-mediated even-parity pairing in heavy-fermion superconductors, *Physical Review B* **34** (9), 6554–6556, 1986.
- [117] D. J. Scalapino, E. Loh, and J. E. Hirsch, D-wave pairing near a spin-density-wave instability, *Physical Review B* **34** (11), 8190–8192, 1986.
- [118] A. V. Chubukov, A jump in our understanding of quantum criticality in metals, https://doi.org/10.36471/JCCM_September_2017_03, [Online; accessed 20-June-2020], 2017.
- [119] S. Sur and S.-S. Lee, Chiral non-Fermi liquids, *Physical Review B* **90** (4), 045121, 2014.
- [120] J. Gubernatis, N. Kawashima, and P. Werner, *Quantum Monte Carlo Methods: Algorithms for Lattice Models*. Cambridge University Press, 2016.

- [121] F. Assaad and H. Evertz, World-line and Determinantal Quantum Monte Carlo Methods for Spins, Phonons and Electrons, in *Computational Many-Particle Physics*, H. Fehske, R. Schneider, and A. Weiße, Eds., Springer Berlin Heidelberg, 2008.
- [122] F. Assaad, Quantum Monte Carlo Methods on Lattices: The Determinantal Approach, in *Quantum Simulations of Complex Many-Body Systems: From Theory to Algorithms*, J. Grotendorst, D. Marx, and A. Muramatsu, Eds., Springer Berlin Heidelberg, 2002.
- [123] W. Krauth, *Statistical Mechanics: Algorithms and Computations*. Oxford University Press, 2006.
- [124] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in C: The Art of Scientific Computing*, 2nd Ed. Cambridge University Press, 1992.
- [125] H. G. Katzgraber, Introduction to Monte Carlo Methods, Springer Undergraduate Texts in Mathematics and Technology, 2009. arXiv: 0905.1629.
- [126] V. Ambegaokar and M. Troyer, Estimating errors reliably in Monte Carlo simulations of the Ehrenfest model, *American Journal of Physics* **78** (2), 150–157, 2010.
- [127] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, Equation of state calculations by fast computing machines, *The Journal of Chemical Physics* **21** (6), 1087–1092, 1953.
- [128] J. Dongarra and F. Sullivan, Guest editors introduction to the top 10 algorithms, *Computing in Science Engineering* **2** (1), 22–23, 2000.
- [129] B. Cipra, The best of the 20th century: Editors name top 10 algorithms, *SIAM News* **33** (4), 2000.
- [130] W. K. Hastings, Monte Carlo sampling methods using Markov chains and their applications, *Biometrika* **57** (1), 97–109, 1970.
- [131] M. Weigel and W. Janke, Error estimation and reduction with cross correlations, *Physical Review E* **81** (6), 066701, 2010.
- [132] U. Wolff, Collective monte carlo updating for spin systems, *Physical Review Letters* **62** (4), 361–364, 1989.
- [133] R. H. Swendsen and J.-S. Wang, Nonuniversal critical dynamics in monte carlo simulations, *Physical Review Letters* **58** (2), 86–88, 1987.
- [134] E. Gull, A. J. Millis, A. I. Lichtenstein, A. N. Rubtsov, M. Troyer, and P. Werner, Continuous-time monte carlo methods for quantum impurity models, *Reviews of Modern Physics* **83** (2), 349–404, 2011.
- [135] R. Blankenbecler, D. J. Scalapino, and R. L. Sugar, Monte Carlo calculations of coupled boson-fermion systems. I, *Physical Review D* **24** (8), 2278–2286, 1981.
- [136] E. Y. Loh, J. E. Gubernatis, R. T. Scalettar, S. R. White, D. J. Scalapino, and R. L. Sugar, Numerical Stability and the Sign Problem in the Determinant Quantum Monte Carlo Method, *International Journal of Modern Physics C* **16** (8), 1319–1327, 2005.

- [137] D. J. Scalapino, S. R. White, and S. Zhang, Insulator, metal, or superconductor: The criteria, *Physical Review B* **47** (13), 7995–8007, 1993.
- [138] H. F. Trotter, On the Product of Semi-Groups of Operators, *Proceedings of the American Mathematical Society* **10** (4), 545–551, 1959.
- [139] M. Suzuki, Quantum statistical monte carlo methods and applications to spin systems, *Journal of Statistical Physics* **43** (5-6), 883–909, 1986.
- [140] E. Y. Loh, J. E. Gubernatis, R. T. Scalettar, S. R. White, D. J. Scalapino, and R. L. Sugar, Sign problem in the numerical simulation of many-electron systems, *Physical Review B* **41** (13), 9301–9307, 1990.
- [141] M. Troyer and U.-J. Wiese, Computational complexity and fundamental limitations to fermionic quantum Monte Carlo simulations, *Physical Review Letters* **94** (17), 170201, 2005.
- [142] C. Wu and S. C. Zhang, Sufficient condition for absence of the sign problem in the fermionic quantum Monte Carlo algorithm, *Physical Review B* **71** (15), 155115, 2005.
- [143] Z.-X. Li and H. Yao, Sign-Problem-Free Fermionic Quantum Monte Carlo: Developments and Applications, *Annual Review of Condensed Matter Physics* **10** (1), 337–356, 2019.
- [144] Z.-C. Wei, Semigroup Approach to the Sign Problem in Quantum Monte Carlo Simulations, 2017. arXiv: 1712.09412.
- [145] Z. C. Wei, C. Wu, Y. Li, S. Zhang, and T. Xiang, Majorana Positivity and the Fermion Sign Problem of Quantum Monte Carlo Simulations, *Physical Review Letters* **116** (25), 250601, 2016.
- [146] Z.-X. Li, Y.-F. Jiang, and H. Yao, Solving the fermion sign problem in quantum Monte Carlo simulations by Majorana representation, *Physical Review B* **91** (24), 241117, 2015.
- [147] Z.-X. Li, Y.-F. Jiang, and H. Yao, Majorana-Time-Reversal Symmetries: A Fundamental Principle for Sign-Problem-Free Quantum Monte Carlo Simulations, *Physical Review Letters* **117** (26), 267002, 2016.
- [148] R. R. Dos Santos, Magnetism and pairing in Hubbard bilayers, *Physical Review B* **51** (21), 15540, 1995.
- [149] M. H. Gerlach, *Quantum monte carlo studies of a metallic spin-density wave transition*, Ph.D. dissertation, Universität zu Köln, 2016.
- [150] K. Binder and D. Heermann, *Monte Carlo Simulation in Statistical Physics : An Introduction*. Springer, 2010.
- [151] D. Hangleiter, I. Roth, D. Nagaj, and J. Eisert, Easing the Monte Carlo sign problem, 2019. arXiv: 1906.02309.
- [152] M. Iazzi, A. A. Soluyanov, and M. Troyer, Topological origin of the fermion sign problem, *Physical Review B* **93** (11), 115102, 2016.
- [153] T. Nakamura, Vanishing of the negative-sign problem of quantum Monte Carlo simulations in one-dimensional frustrated spin systems, *Physical Review B* **57** (6), R3197(R), 1998.

- [154] S. Chandrasekharan and A. Li, Fermion bag solutions to some sign problems in four-fermion field theories, *Physical Review D* **85** (9), 091502, 2012.
- [155] S. Chandrasekharan and U.-J. Wiese, Meron-Cluster Solution of Fermion Sign Problems, *Physical Review Letters* **83** (16), 3116–3119, 1999.
- [156] S. Mertens, Computational complexity for physicists, *Computing in Science & Engineering* **4** (3), 31–47, 2002.
- [157] M. A. Nielsen and I. L. Chuang, *Quantum computation and quantum information*. Cambridge University Press, 2000.
- [158] R. T. Scalettar, E. Y. Loh, J. E. Gubernatis, A. Moreo, S. R. White, D. J. Scalapino, R. L. Sugar, and E. Dagotto, Phase diagram of the two-dimensional negative-U Hubbard model, *Physical Review Letters* **62** (12), 1407–1410, 1989.
- [159] D. J. Scalapino, S. R. White, and S. C. Zhang, Superfluid density and the Drude weight of the Hubbard model, *Physical Review Letters* **68** (18), 2830–2833, 1992.
- [160] J. E. Hirsch, Discrete Hubbard-Stratonovich transformation for fermion lattice models, *Physical Review B* **28** (7), 4059–4061, 1983.
- [161] F. F. Assaad, Depleted Kondo lattices: Quantum Monte Carlo and mean-field calculations, *Physical Review B* **65** (11), 115104, 2002.
- [162] M. R. Zirnbauer, Particle-Hole Symmetries in Condensed Matter, 2020. arXiv: 2004.07107.
- [163] Théorie générale de la rotation paramagnétique dans les cristaux, *Proceedings Koninklijke Akademie van Wetenschappen* **33**, 959–972, 1930.
- [164] L. Wang, Y.-H. Liu, M. Iazzi, M. Troyer, and G. Harcos, Split Orthogonal Group: A Guiding Principle for Sign-Problem-Free Fermionic Simulations, *Physical Review Letters* **115** (25), 250601, 2015.
- [165] J. Hilgert and K.-H. Neeb, *Lie Semigroups and their Applications*. Springer, 1993.
- [166] S. Kourtis, J. W. F. Venderbos, and M. Daghofer, Fractional Chern insulator on a triangular lattice of strongly correlated t_{2g} electrons, *Physical Review B* **86** (23), 235118, 2012.
- [167] L. Balents, C. R. Dean, D. K. Efetov, and A. F. Young, Superconductivity and strong correlations in moiré flat bands, *Nature Physics* **16** (7), 725–733, 2020.
- [168] X. Lu, P. Stepanov, W. Yang, *et al.*, Superconductors, orbital magnets and correlated states in magic-angle bilayer graphene, *Nature* **574** (7780), 653–657, 2019.
- [169] G. Chen, A. L. Sharpe, P. Gallagher, *et al.*, Signatures of tunable superconductivity in a trilayer graphene moiré superlattice, *Nature* **572** (7768), 215–219, 2019.
- [170] Y. Cao, V. Fatemi, S. Fang, K. Watanabe, T. Taniguchi, E. Kaxiras, and P. Jarillo-Herrero, Unconventional superconductivity in magic-angle graphene superlattices, *Nature* **556** (7699), 43–50, 2018.

- [171] M. Yankowitz, S. Chen, H. Polshyn, Y. Zhang, K. Watanabe, T. Taniguchi, D. Graf, A. F. Young, and C. R. Dean, Tuning superconductivity in twisted bilayer graphene, *Science* **363** (6431), 1059–1064, 2019.
- [172] N. B. Kopnin, T. T. Heikkilä, and G. E. Volovik, High-temperature surface superconductivity in topological flat-band systems, *Physical Review B* **83** (22), 220503, 2011.
- [173] J. S. Hofmann, E. Berg, and D. Chowdhury, Strong-coupling superconductivity and pseudogap in topological flat bands: a quantum Monte Carlo study, 2019. arXiv: 1912.08848.
- [174] T. T. Heikkilä, N. B. Kopnin, and G. E. Volovik, Flat bands in topological media, *JETP Letters* **94** (3), 233–239, 2011.
- [175] V. J. Emery and S. A. Kivelson, Importance of phase fluctuations in superconductors with small superfluid density, *Nature* **374** (6521), 434–437, 1995.
- [176] E. Y. Loh, J. E. Gubernatis, R. L. Scalettar R. T. and Sugar, and S. R. White, Stable Matrix-Multiplication Algorithms for Low-Temperature Numerical Simulations of Fermions, in *Interacting Electrons in Reduced Dimensions*, D. K. Baeriswyl D. and Campbell, Ed., Springer US, 1989, pp. 55–60.
- [177] E. Y. Loh and J. Gubernatis, Stable Numerical Simulations of Models of Interacting Electrons in Condensed-Matter Physics, in *Electronic Phase Transitions*, W. Hanke and Y. V. Kopaev, Eds., Vol. 32, Elsevier, 1992, pp. 177–235.
- [178] Z. Bai, C. Lee, R. C. Li, and S. Xu, Stable solutions of linear systems involving long chain of matrix multiplications, *Linear Algebra and its Applications* **435** (3), 659–673, 2011.
- [179] S. Sorella, S. Baroni, R. Car, and M. Parrinello, A novel technique for the simulation of interacting fermion systems, *Europhysics Letters* **8** (7), 663–668, 1989.
- [180] C. Lin, F. H. Zong, and D. M. Ceperley, Twist-averaged boundary conditions in continuum quantum Monte Carlo algorithms, *Physical Review E* **64** (1), 016702, 2001.
- [181] D. Poilblanc, Twisted boundary conditions in cluster calculations of the optical conductivity in two-dimensional lattice models, *Physical Review B* **44** (17), 9562–9581, 1991.
- [182] J. Dongarra, M. Gates, A. Haidar, J. Kurzak, P. Luszczek, S. Tomov, and I. Yamazaki, The Singular Value Decomposition: Anatomy of Optimizing an Algorithm for Extreme Scale, *SIAM Review* **60** (4), 808–865, 2018.
- [183] D. Goldberg, What every computer scientist should know about floating-point arithmetic, *ACM Computing Surveys* **23** (1), 5–48, 1991.
- [184] R. C. Lee, S. Chiesa, C. N. Varney, *et al.*, *QUEST: Quantum electron simulation toolbox*, 2010.
- [185] M. Bercx, F. Goth, J. S. Hofmann, and F. F. Assaad, The ALF (Algorithms for Lattice Fermions) project release 1.0. Documentation for the auxiliary field quantum Monte Carlo code, *SciPost Physics* **3** (2), 2017.

- [186] M. H. Gerlach, *Mazhgerlach/detqmc: First release of detqmc*, Version v1.0, 2017.
- [187] Y. Schattner and E. Berg, Unpublished DQMC code, 2019.
- [188] E. Anderson, Z. Bai, C. Bischof, *et al.*, *LAPACK Users' Guide*, 3rd Ed. Society for Industrial and Applied Mathematics, 1999.
- [189] S. Blackford, Error Bounds for the Singular Value Decomposition, <http://www.netlib.org/lapack/lug/node96.html>, [Online; accessed 15-May-2020], 1999.
- [190] R. M. Fye and J. E. Hirsch, Monte carlo study of the symmetric anderson-impurity model, *Physical Review B* **38** (1), 433–441, 1988.
- [191] J. E. Hirsch, Stable monte carlo algorithm for fermion lattice systems at low temperatures, *Physical Review B* **38** (16), 12023–12026, 1988.
- [192] T. Paiva and R. R. Dos Santos, Charge-density waves in the Hubbard chain: Evidence for $4k_F$ instability, *Physical Review B* **61** (20), 13480, 2000.
- [193] S. Großmann, *Mathematischer Einführungskurs für die Physik*. Springer Vieweg, 2000.
- [194] P. Broecker and S. Trebst, Numerical stabilization of entanglement computation in auxiliary-field quantum Monte Carlo simulations of interacting many-fermion systems, *Physical Review E* **94** (6), 063306, 2016.
- [195] J. Bezanson, A. Edelman, S. Karpinski, and V. B. Shah, Julia: A Fresh Approach to Numerical Computing, *SIAM Review* **59** (1), 65–98, 2017.
- [196] J. M. Perkel, Julia: come for the syntax, stay for the speed, *Nature* **572**, 141–142, 2019, [Online; accessed 15-May-2020].
- [197] C. Rackauckas, ODE Solver Multi-Language Wrapper Package Work-Precision Benchmarks (MATLAB, SciPy, Julia, deSolve (R)), https://benchmarks.juliadiffeq.org/html/MultiLanguage/wrapper_packages.html, [Online; accessed 15-May-2020].
- [198] X.-Z. Luo, J.-G. Liu, P. Zhang, and L. Wang, Yao.jl: Extensible, Efficient Framework for Quantum Algorithm Design, 2019. arXiv: 1912.10877.
- [199] Intel, Intel Math Kernel Library, <https://software.intel.com/content/www/us/en/develop/tools/math-kernel-library.html>, [Online; accessed 08-August-2020], 2020.
- [200] A. Klein, Y. Schattner, E. Berg, and A. V. Chubukov, Normal state properties of quantum critical metals at finite temperature, 2020. arXiv: 2003.09431.
- [201] M. K. Zhang Xianyi, OpenBLAS, <https://www.openblas.net>, [Online; accessed 08-August-2020], 2020.
- [202] F. Freyer, *A Determinant Quantum Monte Carlo Study of a Sign Problem Free Flat Band Model*, M.S. thesis, University of Cologne, 2020, (unpublished).
- [203] E. Strohmaier, J. Dongarra, H. Simon, and M. Meuer, TOP500 LIST, <https://www.top500.org/lists/top500/list/2020/06/>, [Online; accessed 5-June-2020], 2020.

- [204] J. M. Kosterlitz and D. J. Thouless, Ordering, metastability and phase transitions in two-dimensional systems, *Journal of Physics C: Solid State Physics* **6** (7), 1181–1203, 1973.
- [205] R. Daou, N. Doiron-Leyraud, D. LeBoeuf, *et al.*, Linear temperature dependence of resistivity and change in the Fermi surface at the pseudogap critical point of a high-T_c superconductor, *Nature Physics* **5** (1), 31–34, 2009.
- [206] R. A. Cooper, Y. Wang, B. Vignolle, *et al.*, Anomalous Criticality in the Electrical Resistivity of $\text{La}_{2-x}\text{Sr}_x\text{CuO}_4$, *Science* **323** (5914), 603–607, 2009.
- [207] N. Doiron-Leyraud, C. Proust, D. LeBoeuf, J. Levallois, J. B. Bonnemaïson, R. Liang, D. A. Bonn, W. N. Hardy, and L. Taillefer, Quantum oscillations and the Fermi surface in an underdoped high-T_c superconductor, *Nature* **447** (7144), 565–568, 2007.
- [208] N. Trivedi and M. Randeria, Deviations from Fermi-Liquid Behavior above T_c in 2D Short Coherence Length Superconductors, *Physical Review Letters* **75** (2), 312–315, 1995.
- [209] B. L. Altshuler, L. B. Ioffe, and A. J. Millis, Critical behavior of the $T=0$ $2k_F$ density-wave phase transition in a two-dimensional Fermi liquid, *Physical Review B* **52**, 5563–5572, 8 1995.
- [210] X. Wang and E. Berg, Scattering mechanisms and electrical transport near an ising nematic quantum critical point, *Physical Review B* **99** (23), 235136, 2019.
- [211] G. Carleo, I. Cirac, K. Cranmer, L. Daudet, M. Schuld, N. Tishby, L. Vogt-Maranto, and L. Zdeborová, Machine learning and the physical sciences, *Reviews of Modern Physics* **91** (4), 45002, 2019.
- [212] J. Carrasquilla and R. G. Melko, Machine learning phases of matter, *Nature Physics* **13** (5), 431–434, 2017.
- [213] G. Carleo and M. Troyer, Solving the quantum many-body problem with artificial neural networks, *Science* **355** (6325), 602–606, 2017.
- [214] D.-L. Deng, X. Li, and S. Das Sarma, Machine learning topological states, *Physical Review B* **96** (19), 195145, 2017.
- [215] D.-L. Deng, X. Li, and S. Das Sarma, Quantum entanglement in neural network states, *Physical Review X* **7** (2), 021021, 2017.
- [216] I. Glasser, N. Pancotti, M. August, I. D. Rodriguez, and J. I. Cirac, Neural-network quantum states, string-bond states, and chiral topological states, *Physical Review X* **8** (1), 011006, 2018.
- [217] G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. Melko, and G. Carleo, Neural-network quantum state tomography, *Nature Physics* **14** (5), 447–450, 2018.
- [218] X. Gao and L.-M. Duan, Efficient representation of quantum many-body states with deep neural networks, *Nature Communications* **8** (1), 662, 2017.

- [219] S. Li, P. M. Dee, E. Khatami, and S. Johnston, Accelerating lattice quantum monte carlo simulations using artificial neural networks: Application to the holstein model, *Physical Review B* **100** (2), 020302, 2019.
- [220] J. Liu, H. Shen, Y. Qi, Z. Y. Meng, and L. Fu, Self-learning Monte Carlo method and cumulative update in fermion systems, *Physical Review B* **95** (24), 241104, 2017.
- [221] H. Shen, J. Liu, and L. Fu, Self-learning monte carlo with deep neural networks, *Physical Review B* **97** (20), 205140, 2018.
- [222] J. Liu, Y. Qi, Z. Y. Meng, and L. Fu, Self-learning monte carlo method, *Physical Review B* **95** (4), 041101, 2017.
- [223] X. Y. Xu, Y. Qi, J. Liu, L. Fu, and Z. Y. Meng, Self-learning quantum Monte Carlo method in interacting fermion systems, *Physical Review B* **96** (4), 2017.
- [224] E. P. L. van Nieuwenburg, Y.-h. Liu, and S. D. Huber, Learning phase transitions by confusion, *Nature Physics* **13** (5), 435–439, 2017.
- [225] Y. Zhang and E.-A. Kim, Quantum Loop Topography for Machine Learning, *Physical Review Letters* **118** (21), 216401, 2017.
- [226] Y. Zhang, R. G. Melko, and E.-A. Kim, Machine learning Z_2 quantum spin liquids with quasiparticle statistics, *Physical Review B* **96** (24), 245119, 2017.
- [227] P. Broecker, J. Carrasquilla, R. G. Melko, and S. Trebst, Machine learning quantum phases of matter beyond the fermion sign problem, *Scientific Reports* **7** (1), 8823, 2017.
- [228] P. Broecker, F. F. Assaad, and S. Trebst, Quantum phase recognition via unsupervised machine learning, 2017. arXiv: 1707.00663.
- [229] C. Bishop, *Pattern recognition and machine learning*. Springer, 2006.
- [230] M. A. Nielsen, *Neural Networks and Deep Learning*. Determination Press, 2015.
- [231] S. O. Haykin, *Neural Networks: A Comprehensive Foundation*, 2nd Ed. Pearson, 1999.
- [232] F. Rosenblatt, *Principles of Neurodynamics: Perceptrons and the Theory of Brain Mechanisms*. Spartan Books, 1962.
- [233] W. S. McCulloch and W. Pitts, A logical calculus of the ideas immanent in nervous activity, *The bulletin of mathematical biophysics* **5** (4), 115–133, 1943.
- [234] P. Dayan and L. Abbott, *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*. MIT Press, 2001.
- [235] P. Bröcker, *Disentangling and machine learning the many-fermion problem*, Ph.D. dissertation, Universität zu Köln, 2018.
- [236] E. van Nieuwenburg, E. Bairey, and G. Refael, Learning phase transitions from dynamics, *Physical Review B* **98** (6), 060301, 2018.
- [237] A. Amini, A. Soleimany, S. Karaman, and D. Rus, Spatial uncertainty sampling for end-to-end control, 2018. arXiv: 1805.04829.

- [238] H. T. Siegelmann and E. D. Sontag, Turing computability with neural nets, *Applied Mathematics Letters* **4** (6), 77–80, 1991.
- [239] H. Li, Z. Xu, G. Taylor, C. Studer, and T. Goldstein, Visualizing the Loss Landscape of Neural Nets, 2017. arXiv: 1712.09913.
- [240] M. J. Kochenderfer and T. A. Wheeler, *Algorithms for Optimization*. MIT Press, 2008.
- [241] A. G. Baydin, B. A. Pearlmutter, A. A. Radul, and J. M. Siskind, Automatic differentiation in machine learning: A survey, 2015. arXiv: 1502.05767.
- [242] M. Innes, Don't unroll adjoint: Differentiating ssa-form programs, 2018. arXiv: 1810.07951.
- [243] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, Dropout: A simple way to prevent neural networks from overfitting, *The Journal of Machine Learning Research* **15** (1), 1929–1958, 2014.
- [244] R. E. W. Rafael C. Gonzalez, *Digital image processing*, 4th Ed. Pearson, 2018.
- [245] C. J. B. Yann LeCun Corinna Cortes, The MNIST database of handwritten digits, <http://yann.lecun.com/exdb/mnist/>, [Online; accessed 07-September-2020], 2020.
- [246] D. Cireřan, U. Meier, and J. Schmidhuber, Multi-column deep neural networks for image classification, 2012. arXiv: 1202.2745.
- [247] K. Ch'ng, J. Carrasquilla, R. G. Melko, and E. Khatami, Machine Learning Phases of Strongly Correlated Fermions, *Physical Review X* **7** (3), 031038, 2017.
- [248] F. Schindler, N. Regnault, and T. Neupert, Probing many-body localization with neural networks, *Physical Review B* **95** (24), 245134, 2017.
- [249] I. A. Iakovlev, O. M. Sotnikov, and V. V. Mazurenko, Supervised learning approach for recognizing magnetic skyrmion phases, *Physical Review B* **98** (17), 174411, 2018.
- [250] X.-Y. Dong, F. Pollmann, and X.-F. Zhang, Machine learning of quantum phase transitions, *Physical Review B* **99** (12), 121104, 2019.
- [251] L. Wang, Discovering phase transitions with unsupervised learning, *Physical Review B* **94** (19), 195105, 2016.
- [252] C. Wang and H. Zhai, Machine learning of frustrated classical spin models. i. principal component analysis, *Physical Review B* **96** (14), 144432, 2017.
- [253] M. Abadi, A. Agarwal, P. Barham, *et al.*, TensorFlow: Large-scale machine learning on heterogeneous systems, Software available from tensorflow.org, 2015.
- [254] T. Ohtsuki and T. Ohtsuki, Deep Learning the Quantum Phase Transitions in Random Two-Dimensional Electron Systems, *Journal of the Physical Society of Japan* **85** (12), 123706, 2016.
- [255] T. Ohtsuki and T. Ohtsuki, Deep learning the quantum phase transitions in random electron systems: Applications to three dimensions, *Journal of the Physical Society of Japan* **86** (4), 044708, 2017.

-
- [256] M. J. S. Beach, A. Golubeva, and R. G. Melko, Machine learning vortices at the kosterlitz-thouless transition, *Physical Review B* **97** (4), 045207, 2018.
 - [257] A. Moreo and D. J. Scalapino, Two-dimensional negative-U Hubbard model, *Physical Review Letters* **66** (7), 946–948, 1991.
 - [258] T. Paiva, R. R. dos Santos, R. T. Scalettar, and P. J. H. Denteneer, Critical temperature for the two-dimensional attractive Hubbard model, *Physical Review B* **69** (18), 184501, 2004.
 - [259] Z.-X. Li, F. Wang, H. Yao, and D.-H. Lee, What makes the T_c of monolayer FeSe on SrTiO₃ so high: a sign-problem-free quantum Monte Carlo study, *Science Bulletin* **61** (12), 925, 2016.
 - [260] S. Sachdev and D. Chowdhury, The novel metallic states of the cuprates: Topological Fermi liquids and strange metals, *Progress of Theoretical and Experimental Physics* **2016** (12), 2016.
 - [261] L. Huang and L. Wang, Accelerated Monte Carlo simulations with restricted Boltzmann machines, *Physical Review B* **95** (3), 035105, 2017.

Erklärung

Ich versichere, dass ich die von mir vorgelegte Dissertation selbständig angefertigt, die benutzten Quellen und Hilfsmittel vollständig angegeben und die Stellen der Arbeit – einschließlich Tabellen, Karten und Abbildungen –, die anderen Werken im Wortlaut oder dem Sinn nach entnommen sind, in jedem Einzelfall als Entlehnung kenntlich gemacht habe; dass diese Dissertation noch keiner anderen Fakultät oder Universität zur Prüfung vorgelegen hat; dass sie – abgesehen von unten angegebenen Teilpublikationen – noch nicht veröffentlicht worden ist, sowie, dass ich eine solche Veröffentlichung vor Abschluss des Promotionsverfahrens nicht vornehmen werde. Die Bestimmungen der Promotionsordnung sind mir bekannt. Die von mir vorgelegte Dissertation ist von Prof. Dr. Simon Trebst betreut worden.

Köln, der 7. Januar 2021

(Carsten Bauer)

(Eine Liste der Teilpublikationen findet sich auf der nächsten Seite.)

Publikationen

(Peer-reviewed)

- [P1] **C. Bauer**, Y. Schattner, S. Trebst, and E. Berg, Hierarchy of energy scales in an $O(3)$ symmetric antiferromagnetic quantum critical metal: A Monte Carlo study, *Physical Review Research* **2** (2), 023008, 2020.
- [P2] **C. Bauer**, Fast and stable determinant quantum Monte Carlo, *SciPost Physics Core* **2** (2), 11, 2020.
- [P3] Y. Zhang, **C. Bauer**, P. Broecker, S. Trebst, and E.-A. Kim, Probing transport in quantum many-fermion simulations via quantum loop topography, *Physical Review B* **99** (16), 161120, 2019.
- [P4] G. T. Driskell, S. Lederer, **C. Bauer**, S. Trebst, and E.-A. Kim, Observation of non-Fermi liquid physics in a quantum critical metal via quantum loop topography, 2020, (submitted to *Nature Communications*). arXiv: 2007.07898.

(Book contributions / proceedings)

- [P5] **C. Bauer** and S. Trebst, Machine Learning Transport Properties in Quantum Many-Fermion Simulations, in *NIC Symposium 2020*, Vol. 50, Forschungszentrum Jülich GmbH Zentralbibliothek, Verlag, 2020.

(Selected open-source publications)

- [S1] **C. Bauer**, *julia-sdw-dqmc: Determinant quantum Monte Carlo in Julia*, Version 1.0, Zenodo, 2020.
- [S2] **C. Bauer** and F. Freyer, *BinningAnalysis.jl: Standard error estimation tools*, Version 0.3.3, Zenodo, 2020.
- [S3] **C. Bauer**, *StableDQMC.jl: Numerical stabilization routines for determinant quantum Monte Carlo*, Version 0.1.3, Zenodo, 2020.
- [S4] **C. Bauer**, *MonteCarlo.jl: Classical and quantum Monte Carlo simulations in Julia*, Version 0.1, Zenodo, 2020.

In the last few decades, quantum criticality in itinerant electron systems has become a central focus of condensed matter physics. On the one hand, it represents a candidate mechanism for high-temperature superconductivity. On the other hand, it can lead to a breakdown of Fermi liquid theory. The formation of a comprehensive understanding of metallic quantum criticality has, however, been significantly hampered by the fact that in many-fermion systems, fluctuations of a critical order parameter can couple to extensive gapless modes on a finite Fermi surface. This interplay, while giving rise to intriguing physical phenomena, leads to strong electronic correlations, which are notoriously difficult to handle by analytic methods.

In this thesis, we investigate metallic quantum criticality by means of large-scale quantum Monte Carlo simulations and contribute unbiased, rigorous results to the discussion. Focusing on antiferromagnetic spin-density wave ordering, we present the numerically exact phase diagram of a spin-fermion model whose solution has so far been out of reach. In particular, we highlight the emergence of high-temperature d-wave superconductivity and rigorously establish the quantum critical properties of the antiferromagnetic quantum critical point. Combining the Monte Carlo method with a quantum loop topography approach, we demonstrate that important features of quantum critical metals can be autonomously identified by machine learning of current-current correlations. This allows us to analyze the electronic transport characteristics of two quantum critical metals, including the spin-fermion model, over a large parameter range and leads to the identification of extended non-Fermi liquid regimes in their respective phase diagrams.