
Quantum gas microscopy of interacting quantum matter with artificial gauge fields

Alexander Fabian Impetro



München 2025

Quantum gas microscopy of interacting quantum matter with artificial gauge fields

Dissertation an der Fakultät für Physik
Ludwig-Maximilians-Universität München

vorgelegt von

Alexander Fabian Impertro

aus Gerolstein

München, den 24. April 2025

Tag der mündlichen Prüfung: 05. Juni 2025

Erstgutachter: Prof. Immanuel Bloch

Zweitgutachter: Prof. Christof Weitenberg

Weitere Prüfungskommissionsmitglieder: Prof. Ulrich Schollwöck,
Prof. Alexander Högele

Zusammenfassung

Die analoge Quantensimulation mit ultrakalten Atomen in optischen Gittern bietet einen leistungsfähigen Ansatz zur Untersuchung stark korrelierter quantenmechanischer Vielteilchensysteme, deren Komplexität klassische Simulationen unmöglich macht. Diese Dissertation berichtet über wesentliche Fortschritte, die mithilfe eines neu entwickelten Quantengasmikroskops auf Basis ultrakalter Cäsiumatome erzielt wurden. Die Plattform wurde speziell entworfen, um exotische Quantenmaterie in großskaligen Systemen mit bislang unerreichter Abstimmbarkeit, Präzision und Kontrolle zu erforschen.

Zunächst stellen wir eine Methode basierend auf unüberwacht trainierten neuronalen Netzen vor, um die atomare Besetzung in optischen Gittern in einem Regime zu rekonstruieren, in dem der Gitterabstand deutlich unterhalb der optischen Auflösungsgrenze liegt. Das unüberwachte Training direkt auf experimentelle Fluoreszenzbilder erfasst den gesamten Abbildungsprozess akkurat, einschließlich komplexer korrelierter Emissionsprozesse, die insbesondere im Bereich kurzer atomarer Abstände auftreten. Wir zeigen, dass diese Methode eine hochpräzise Detektion ermöglicht und so den Weg für Quantensimulationen in bisher unerreichbaren Regimen ebnet.

Des Weiteren ergänzen wir die Detektions- und Kontrollfähigkeiten der Quantengasmikroskopie über reine Dichtemessungen hinaus, indem wir eine robuste und skalierbare Detektionstechnik für nichtdiagonale Observablen wie lokale Ströme und kinetische Energien implementieren. Durch die Unterteilung eines zweidimensionalen Gittersystems in eine Anordnung isolierter Doppelpotentialtöpfe mittels eines optischen Supergitters demonstrieren wir die direkte Messung und Manipulation von Kohärenzen sowohl mit globaler als auch lokaler Kontrolle in großen Systemen. Diese Erweiterung verbessert die experimentellen Fähigkeiten von Quantengasmikroskopen erheblich und ermöglicht neuartige Untersuchungen von Vielteilchensystemen mit wichtigen Anwendungen in topologischer Materie, Nichtgleichgewichtsdynamik und Supraleitung.

Abschließend untersuchen wir wechselwirkende Vielteilchenphasen in Gegenwart eines künstlichen Magnetfeldes auf optischen Flussleitern. Insbesondere realisieren wir die Mott-Meissner-Phase – einen Zustand, der wechselwirkungsinduzierte Lokalisierung mit chiralen Strömen kombiniert, welche durch das Magnetfeld erzeugt werden – in großskaligen bosonischen Flussleitern mit 48 Gitterplätzen bei halber Füllung. Durch die Kombination der lokalen Messung von Teilchenströmen mit einer flexiblen Kontrolle der Wechselwirkungsstärke enthüllen wir die entstehenden Gleichgewichtsströme mit vollständiger räumlicher Auflösung und untersuchen deren Abhängigkeit von der Wechselwirkungsstärke. Diese Ergebnisse liefern einen direkten experimentellen Nachweis der vorhergesagten Mott-Meissner-Phase und stellen so einen bedeutenden Fortschritt in Richtung der großskaligen Realisierung wechselwirkender topologischer Materie dar. Zusammen mit den vorher genannten Ergebnissen eröffnen wir damit eine Vielzahl neuer Wege für die analoge Quantensimulation mit ultrakalten Atomen.

Abstract

Analog quantum simulation with ultracold atoms in optical lattices represents a powerful approach to investigate strongly correlated quantum many-body systems, whose complexity renders classical simulation infeasible. This thesis reports key advances achieved with a newly developed quantum gas microscope based on ultracold cesium atoms, specifically tailored to explore exotic quantum phenomena in large-scale systems with unprecedented tunability, precision, and control.

First, we introduce a method based on unsupervised neural networks to reconstruct lattice occupations at atomic separations significantly smaller than the optical resolution limit. The unsupervised approach, trained directly on experimental fluorescence images, accurately captures the full imaging process including complex correlated emission processes arising in the short-spacing regime. We demonstrate that this method provides a high-fidelity detection, enabling quantum simulations in previously inaccessible optical regimes.

Secondly, we expand quantum gas microscopy beyond simple density measurements by implementing a robust and scalable detection technique for off-diagonal observables, such as local currents and kinetic energies. By partitioning a two-dimensional lattice system into an array of isolated double wells using an optical superlattice, we demonstrate direct measurement and manipulation of coherences with both large-scale global and local control. This enhancement broadens the experimental capabilities of quantum gas microscopes, allowing novel exploration of many-body systems with important applications in topological matter, non-equilibrium dynamics and superconductivity.

Lastly, we experimentally study interacting many-body states in the presence of an artificial gauge field on optical flux ladders. In particular, we realize the strongly interacting Mott-Meissner phase – a state combining interaction-induced localization with chiral currents induced by the gauge field – in large-scale bosonic flux ladders with 48 sites at half filling. By combining the local readout of currents with a flexible interaction control, we reveal the emerging equilibrium particle currents with full spatial resolution and study their interaction dependence, providing direct experimental evidence of the predicted Mott-Meissner phase. These results constitute a significant experimental advance towards the large-scale realization of interacting topological matter, and more broadly, they provide new avenues for the analog quantum simulation with ultracold atoms.

Contents

Introduction	1
1 Quantum simulation of topological many-body phases	6
1.1 Quantum simulation using ultracold atoms in optical lattices	6
1.2 Topological many-body phases	10
1.3 Floquet engineering of artificial gauge fields	13
2 Experimental setup	19
2.1 A quantum simulation platform based on ultracold cesium atoms	19
2.2 Overview of the apparatus and the cooling stages	22
2.3 Optical lattice setups	25
2.3.1 Vertical lattices	25
2.3.2 Horizontal superlattices	30
2.4 High-resolution imaging and potential shaping	32
2.5 Laser-assisted tunneling setup	34
3 Calibration and preparation techniques	37
3.1 Imaging and molasses	37
3.2 Initial state preparation	40
3.3 Vertical lattice parameters and loading	43
3.3.1 Shallow-angle vertical lattice	43
3.3.2 Steep-angle vertical lattice	43
3.4 Calibration of the horizontal lattice parameters	45
3.4.1 Tunnel coupling in a monochromatic lattice via CDW decay	46
3.4.2 Superlattice phase	47
3.4.3 Double well tunnel coupling via double well oscillations	48
3.5 Calibration of the interaction energy	49
3.6 Calibration of the laser-assisted tunneling	51
3.7 On-site potential disorder	54
4 Single-site reconstruction using unsupervised neural networks	56
4.1 Imaging and reconstruction in optical lattices	56
4.1.1 The optical process: Fluorescence imaging	57
4.1.2 The deconvolution problem	60

4.1.3	Experimental parameter regimes	61
4.1.4	Occupation reconstruction and common approaches	62
4.1.5	Reconstruction of experimental data with conventional approaches	66
4.2	Neural networks for image reconstruction	70
4.2.1	Brief introduction to machine learning and neural networks	70
4.2.2	Supervised vs. unsupervised learning	73
4.2.3	A simple network for single-site reconstruction	75
4.3	Autoencoder-based approach for single-site reconstruction	80
4.3.1	Idea and network architecture	80
4.3.2	Training procedure and hyperparameters	83
4.3.3	Reconstruction process	86
4.3.4	Performance evaluation using experimental data	88
4.3.5	Performance evaluation using simulated data	94
4.3.6	Interpretation and visualization of the learned transformation	96
4.3.7	Optimizations and further development	99
4.4	Conclusions	103
5	Local readout and control of orbital operators	105
5.1	Orbital operators for quantum simulation of many-body physics	105
5.2	Local readout and control scheme using optical superlattices	108
5.3	Experimental implementation of global operations	110
5.4	Extension to spatially-selective operations	115
5.5	Readout of correlation functions	120
5.6	Conclusion	123
6	Strongly-interacting Meissner phases in bosonic flux ladders	125
6.1	Flux ladders as a platform to study orbital physics	125
6.1.1	Effective models	129
6.1.2	Experimental implementation	131
6.2	Measurement of currents on real and complex bonds	131
6.3	Currents in isolated plaquettes with two interacting particles	133
6.4	Identification of parameter regimes with low heating rates	137
6.5	Strongly-interacting Meissner states	139
6.5.1	Adiabatic preparation of the interacting Meissner regime	139
6.5.2	Chiral currents in the Meissner phase	141
6.6	Exploring the phase diagram beyond the Meissner regime	143
6.6.1	Adiabatic preparation by plaquette coupling	143
6.6.2	Currents and density correlations beyond the Meissner regime	145
6.7	Future avenues	147
6.8	Conclusions	149
	Conclusions and outlook	151

Appendices	153
A Supporting material for the study of interacting flux ladders	153
References	157

Introduction

This year, 2025, marks the 100th anniversary of the discovery of quantum mechanics and is celebrated internationally as the *year of quantum*. Quantum mechanics has revolutionized our understanding of the physical world, and led to the development of numerous technologies that have transformed daily life, ranging from lasers [1] and microelectronics [2] to medical diagnostics [3–5]. The theory introduced fundamentally new laws governing the behavior of matter at the microscopic scale. Among the most striking quantum phenomena are those arising from the collective behavior of many interacting particles [6], spanning topics at the forefront of modern research such as superconductivity [7, 8], topological matter [9–11], and other exotic quantum phases [12, 13]. In the strongly interacting regime, mean-field and effective descriptions break down, and the system must be described by a full many-body wavefunction at the fundamental level of individual constituents. Simulating such systems using classical computers is a hopeless endeavor, as the required computational resources scale exponentially with the number of particles, bringing even few-particle systems beyond the reach of the most powerful supercomputers.

In line with Feynman’s original vision of emulating quantum systems with other quantum systems, the field of *quantum simulation* has emerged as a powerful approach to address this challenge and enable the study of many-body physics [14–18]. Facilitated by groundbreaking experimental advances such as laser cooling [19, 20] and Bose-Einstein condensation [21–23], alongside a seminal theoretical proposal by Dieter Jaksch and colleagues [24], ultracold atoms in optical lattices established the first experimental platform for quantum simulation [25, 26]. Such a system of quantum degenerate atoms trapped in a periodic optical potential naturally realizes the *Hubbard model* [27], which is a paradigmatic model for electrons in solids, thereby establishing a direct connection to condensed matter physics [28]. Cold atom simulators provide an unparalleled level of control over system parameters such as dimensionality, lattice geometry, tunneling rate and interaction strength [29]. At the same time, the system is highly isolated from the environment, can be realized with few defects, and allows high-precision probing using optical imaging techniques. A key development in this context was the invention of *quantum gas microscopy*, which enables the detection of individual atoms with single-lattice-site resolution [30, 31]. It pushes detection and control to the ultimate limit of individual constituents, granting access to novel microscopic observables such as correlation functions, full counting statistics and entanglement entropy [32–38]. At the same time, these microscopes allow the shaping of the potential landscape and the manipulation of the atomic states with microscopic precision [38, 39]. Leveraging these capabilities has led to a broad range of experiments, for example on quantum phase transitions [25, 40], transport [41–43], quantum

magnetism [44–46], low-temperature fermionic phases [47–50], and topological matter [51–55]. In fact, the field has matured to a point where it is now possible to envision the realization of exotic quantum phases without known analogs in nature [56].

Beyond optical lattice-based experiments, the past decade has seen the emergence of a wide range of quantum simulation platforms [14, 17, 57], each with unique advantages and limitations. Closely related are neutral atoms in tweezer arrays, which offer enhanced control over individual atoms and their arrangement [58, 59]. Other platforms are based on molecules [60], ions [61, 62], photons [63] or superconducting circuits [64]. Recently, the field has gained further momentum, fueled by the race for practical quantum advantage [65] and ultimately a universal quantum computer [66]. Toward this goal, the complementary approach of digital quantum simulation was developed, where dynamics are decomposed into a sequence of elementary quantum gate operations [67, 68]. While conceptually more flexible, analog quantum simulators such as neutral atoms in optical lattices hold the reign in coherence times and accessible system sizes [69].

The boundaries of analog quantum simulation are constantly being pushed by new experimental platforms. In this spirit, a new quantum gas microscope based on ultracold cesium atoms was initiated in our group in 2017 and completed during the course of this thesis. It is tailored to the study of strongly-correlated quantum phases with exotic properties, featuring large system sizes of several thousand atoms and an exquisite interaction tunability, directly addressing a few of the central challenges, and opportunities, in analog quantum simulation. Throughout this thesis, we will present several of the key advances that we made so far.

The first aspect addresses the very essence of a quantum gas microscope – the ability to detect the individual atoms in the optical lattice. How challenging this is depends crucially on the lattice spacing, i.e., the separation between individual atoms. With novel lattice geometries [70–72], state-dependent lattices [73] or superlattices [74, 75], as well as elements with a large mass or long-range interactions [76, 77], smaller spacings become increasingly relevant. We follow this approach, motivated by the large mass of cesium and the wish for fast tunneling rates. To this end, we developed an entirely new approach to reconstruct the lattice occupation using neural networks that are trained in an unsupervised fashion. Neural networks are powerful tools to detect patterns and extract salient features from complex data, and are now also increasingly being used in quantum physics [78–87]. We leverage this power to build a model of the imaging process directly from experimental data, including complex correlated emission processes that emerge in the short-spacing regime [88, 89]. With this, we are able to push the high-fidelity detection to a regime where the atomic separation is more than a factor of two smaller than the optical resolution, enabling quantum simulation experiments in previously inaccessible settings.

Secondly, we resolve a long-standing limitation of quantum gas microscopes: the detection of only local occupations. Beyond mere density measurements, a wealth of additional information is hidden in off-diagonal elements of the density matrix, i.e., coherences [33, 38]. Examples of observables encoded in these coherences are local current and kinetic energies (so-called *orbital operators*), which are of great interest in the study of topological phases or non-equilibrium dynamics [90, 91]. Their importance has already been highlighted in several proof-of-concept experiments [34, 35, 92, 93], and furthermore, recent theoretical

work has shown that coherences provide access to superconducting pairing correlations [94]. Here, we show that such off-diagonal observables can be measured with local resolution in a robust and scalable way by partitioning a two-dimensional lattice into an array of isolated double wells using an optical superlattice. Programmable operations applied in these double wells enable a rotation of the measurement basis, or alternatively, to precisely manipulate quantum states, which can be used for coherent state engineering or as a building block for fermionic quantum computing and quantum chemistry [95, 96]. With this, we significantly expand the toolbox offered by quantum gas microscopes, enabling novel studies even beyond optical-lattice-based quantum simulators [97, 98].

Lastly, we combine the full capabilities of our setup for the realization of interacting topological matter on analog quantum simulators. The study of topological matter such as fractional quantum Hall states [99, 100] is at the forefront of condensed matter physics due to their exotic properties [11, 101–103] and potential applications in fault-tolerant quantum computing [104, 105]. Their controlled realization on quantum simulators is a long-standing goal, as local detection and control are ideally suited to probe features such as edge currents, fractional excitations (*anyons*) or topological entanglement entropy [106–109]. However, their inherent fragility and experimental overhead due to the need to engineer topological properties, e.g. through an artificial gauge field [106, 110], have so far limited experimental implementations to non-interacting systems [51, 52] or interacting systems with just two particles [111, 112]. Here, we experimentally realize quasi-one-dimensional analogs of interacting topological matter – so-called *flux ladders* – with many interacting particles using a Floquet-engineered gauge field [113–115]. Exploiting the flexible interaction tunability of cesium, we engineer preparation paths to realize large ladder systems with unprecedented sizes of up to 48 sites at half filling. Combined with local current detection, we reveal equilibrium chiral currents and their characteristic interaction scaling, identifying hallmark features of the predicted Mott-Meissner phase. These results demonstrate a pathway towards the large-scale realization of interacting topological matter, and more generally, advance the frontiers of analog quantum simulation with ultracold atoms in optical lattices.

Outline

This thesis is divided into six chapters, which are structured as follows:

Chapter 1 starts with a review of quantum simulation based on ultracold atoms in optical lattices, followed by a survey of the interesting physics of topologically ordered quantum phases, and the opportunities of their implementation on quantum simulators. Furthermore, we detail how Floquet engineering can be used to engineer an artificial magnetic field in cold atom systems, which poses one route to realize topologically ordered matter.

Chapter 2 provides an overview of our experimental platform, detailing the different optical setups to control and detect two-dimensional atomic samples in optical lattices.

Chapter 3 gives a detailed account on calibration measurements and techniques that are important for the subsequent chapters, such as the Hubbard parameters.

Chapter 4 discusses the neural network-based reconstruction algorithm that we use to analyze our experimental data. We start with a review of the basic principles behind high-resolution fluorescence imaging and survey common approaches to reconstruct the lattice occupation, motivating why novel approaches are required in our resolution regime. After giving a brief introduction to neural networks, we present our approach in detail, illustrating the training and evaluation procedure, and benchmark its performance.

Chapter 5 presents a method to extend the measurement capabilities of quantum gas microscopes to off-diagonal elements of the density matrix. After a description of the measurement scheme based on optical superlattices, we benchmark the different operations, corresponding to local rotations of the measurement basis, both globally across large systems in parallel and with an additional local degree of control. Furthermore, we show how these operations can be used to manipulate quantum states, and how correlation functions for orbital operators can be measured.

Chapter 6 demonstrates the realization of strongly correlated many-body phases on bosonic ladder systems in the presence of a Floquet-engineered artificial magnetic field. We start with a theoretical introduction to the physics of the flux ladder system, and present their hallmark features that can be used to probe the many-body phases experimentally. A central quantity are local particle currents, whose measurement we first benchmark on plaquette systems with two interacting particles. We then present the experimental realization of Meissner states, revealing the emergence of chiral currents on top of a homogeneous density plateau. Furthermore, we probe the behavior of the chiral currents as a function of the interaction strength, where we find the predicted behavior of the strongly-correlated Mott-Meissner phase. Lastly, we explore the phase diagram beyond the Meissner regime, and conclude the chapter with a discussion of the current state-of-the-art in terms of accessible temperatures, and the prospects for future experiments.

Publications

The central results presented in this thesis have been published or appeared as preprints in the following references:

- **A. Impertro**, J. F. Wienand, S. Häfele, H. von Raven, S. Hubele, T. Klostermann, C. R. Cabrera, I. Bloch, and M. Aidelsburger, *An Unsupervised Deep Learning Algorithm for Single-Site Reconstruction in Quantum Gas Microscopes*, *Commun. Phys.* **6**, 1 (2023)
- **A. Impertro**, S. Karch, J. F. Wienand, S. Huh, C. Schweizer, I. Bloch, and M. Aidelsburger, *Local Readout and Control of Current and Kinetic Energy Operators in Optical Lattices*, *Phys. Rev. Lett.* **133**, 063401 (2024)
- **A. Impertro**, S. Huh, S. Karch, J. F. Wienand, I. Bloch, and M. Aidelsburger, *Strongly interacting Meissner phases in large bosonic flux ladders*, *Nat. Phys.* (2025)

During the course of this thesis, the following additional research articles have been published or appeared as preprints:

- J. F. Wienand, S. Karch, **A. Impertro**, C. Schweizer, E. McCulloch, R. Vasseur, S. Gopalakrishnan, M. Aidelsburger, and I. Bloch, *Emergence of Fluctuating Hydrodynamics in Chaotic Quantum Systems*, *Nat. Phys.* **20**, 1732–1737 (2024)
- S. Karch, S. Bandyopadhyay, Z.-H. Sun, **A. Impertro**, S. Huh, I. P. Rodríguez, J. F. Wienand, W. Ketterle, M. Heyl, A. Polkovnikov, I. Bloch, and M. Aidelsburger, *Probing Quantum Many-Body Dynamics Using Subsystem Loschmidt Echos*, [arXiv:2501.16995](https://arxiv.org/abs/2501.16995)

CHAPTER 1

Quantum simulation of topological many-body phases

In this chapter, we begin by providing a short introduction to the field of quantum simulation using ultracold atoms in optical lattices. Subsequently, we focus on topological many-body phases of matter, and give a short overview of the interesting physics that can be studied in these systems. Lastly, we show how topological band structures can be engineered for neutral atom systems.

1.1 Quantum simulation using ultracold atoms in optical lattices

The field of ultracold atoms has opened the door to studying a broad array of quantum phenomena under highly controlled laboratory conditions [15]. Two key advantages characterize these experiments: the ability to achieve extremely low temperatures where quantum effects dominate, and the precision with which atomic degrees of freedom can be manipulated and detected. Central to both aspects is the interaction between atoms and light. This interaction can either be resonant, driving specific atomic transitions, or off-resonant, generating optical potentials to trap and manipulate atoms [116]. Focusing on the off-resonant case, the interaction between an optical field and the atomic dipole moment induces an energy shift known as the AC-Stark shift, leading to a conservative dipole potential. This interaction is governed by the Hamiltonian $\hat{H} = -\mathbf{d} \cdot \mathbf{E}$, where \mathbf{d} denotes the atomic dipole moment and \mathbf{E} is the electric field of the incident light. For a single dominant atomic transition (linewidth Γ), the resulting potential $V(\mathbf{r})$ is proportional to the local intensity $I(\mathbf{r})$ of the optical field [116]:

$$V(\mathbf{r}) = -\frac{\text{Re}(\alpha)}{2\epsilon_0 c} I(\mathbf{r}) = \frac{3\pi c^2 \Gamma}{2\omega^3 \Delta} I(\mathbf{r}), \quad (1.1)$$

where α represents the atomic polarizability, ϵ_0 is the vacuum permittivity, c is the speed of light, ω is the frequency of the optical field, and Δ denotes the detuning from resonance. The sign of the detuning determines whether atoms experience an attractive force toward regions of high intensity (red detuning, $\Delta < 0$) or low intensity (blue detuning, $\Delta > 0$). By engineering the intensity profile of the optical field, diverse trapping geometries can be realized. For instance, a focused, red-detuned laser beam generates an attractive harmonic potential

known as an optical dipole trap. More complex geometries can be created using multiple laser beams, leading to a central application in quantum simulation: The interference of two lasers at the same frequency generates a standing-wave pattern of alternating high- and low-intensity regions, realizing a periodic potential known as an *optical lattice* [117]. In one dimension, the potential can be described by

$$V(x) = V_0 \cos^2(kx), \quad (1.2)$$

where V_0 is the lattice depth and k is the wave vector of the lattice. Let us now consider an atomic sample evolving in such a periodic potential, which can be described by the following second-quantized Hamiltonian [118]:

$$\hat{H} = \int d^3x \hat{\Psi}^\dagger(\mathbf{x}) \left(\frac{\hat{\mathbf{p}}^2}{2m} + V(\mathbf{x}) \right) \hat{\Psi}(\mathbf{x}) + \frac{g}{2} \int d^3x \hat{\Psi}^\dagger(\mathbf{x}) \hat{\Psi}^\dagger(\mathbf{x}) \hat{\Psi}(\mathbf{x}) \hat{\Psi}(\mathbf{x}). \quad (1.3)$$

Here, $\hat{\Psi}(\mathbf{x})$ are the bosonic field operators, $\hat{\mathbf{p}}$ is the momentum operator and m is the mass of the atoms. In addition to the kinetic and potential energy terms, we include a contact interaction with a strength g . In the low-energy regime, this interaction can be effectively described by s-wave scattering, characterized by the scattering length a , and is quantified by the coupling constant $g = (4\pi\hbar^2 a)/m$. Due to the space-periodic potential, we can apply the *Bloch theorem*, which tells us that the energy spectrum of this Hamiltonian splits up into discrete Bloch bands separated by band gaps [119]. The eigenstates are so-called Bloch states, characterized by a quasi-momentum \mathbf{q} and delocalized over the entire system. To bring the Hamiltonian into a more convenient form, we work in the basis of *Wannier functions* $w(\mathbf{x})$, which are exponentially localized around the individual lattice sites. Since they form an orthonormal basis, we can expand the field operators as $\hat{\Psi}(\mathbf{x}) = \sum_j w(\mathbf{x} - \mathbf{x}_j) \hat{a}_j$, where \hat{a}_j is the annihilation operator for an atom at site j . In combination with the assumption that the energy of the sample is sufficiently low that only the ground band is occupied, we can simplify the Hamiltonian to obtain

$$\hat{H} = - \sum_{i,j} J_{i,j} \hat{a}_i^\dagger \hat{a}_j + \frac{1}{2} \sum_{i,j,k,l} U_{i,j,k,l} \hat{a}_i^\dagger \hat{a}_j^\dagger \hat{a}_k \hat{a}_l + \sum_i \epsilon_i \hat{a}_i^\dagger \hat{a}_i, \quad (1.4)$$

with the tunnel matrix elements $J_{i,j} = \int d^3x w^*(\mathbf{x} - \mathbf{x}_i) \left(-\frac{\mathbf{p}^2}{2m} + V(\mathbf{x}) \right) w(\mathbf{x} - \mathbf{x}_j)$, the interaction matrix elements $U_{i,j,k,l} = \frac{g}{2} \int d^3x w^*(\mathbf{x} - \mathbf{x}_i) w^*(\mathbf{x} - \mathbf{x}_j) w(\mathbf{x} - \mathbf{x}_k) w(\mathbf{x} - \mathbf{x}_l)$, and on-site energy offsets ϵ_i . In a last step, we can make use of the fact that one typically works with sufficiently deep lattices, where we only have tunneling between nearest-neighbor sites $J_{i,j} \approx J \delta_{\langle i,j \rangle}$ with a uniform tunneling strength J , and the interaction is dominated by on-site terms $U_{i,j,k,l} \approx U \delta_{i,j} \delta_{i,k} \delta_{i,l}$ with an interaction strength U (*tight-binding limit*). With this, we arrive at the Hamiltonian

$$\hat{H} = -J \sum_{\langle i,j \rangle} \hat{a}_i^\dagger \hat{a}_j + \frac{U}{2} \sum_i \hat{n}_i (\hat{n}_i - 1) + \sum_i \epsilon_i \hat{n}_i, \quad (1.5)$$

which is known as the *Bose-Hubbard model*. Here, $\hat{n}_i = \hat{a}_i^\dagger \hat{a}_i$ is the number operator at site i . This model provides a simple, intuitive description of the essential physics of ultracold atoms in optical lattices, reduced to the interplay between two processes, nearest-neighbor tunneling and on-site interactions (see illustration for a two-dimensional lattice in Fig. 1.1a). Historically, the Hubbard model was however not introduced in the context of ultracold atoms, but rather as a model to describe the behavior of strongly correlated electrons moving in the crystal lattice of a solid. It is this correspondence that was one of the original motivations for the study of ultracold atoms in optical lattices as a quantum simulator for condensed matter systems [18, 24, 25]. Despite the conceptual simplicity, the ground-state phase diagram of the Bose-Hubbard model captures the paradigmatic quantum phase transition between a superfluid and a Mott insulator [120]. In particular, the phase of the system is determined by the ratio J/U between tunneling and interactions. For vanishing interactions ($U \ll J$), the tunneling term dominates the physics, and the system is in a *superfluid phase*, characterized by a delocalized wave function and long-range phase coherence. The ground state is a Bose-Einstein condensate (BEC) with all atoms in the $\mathbf{q} = 0$ state, which, in the case of a large system, can be expressed as a product state where each lattice site i is occupied by a coherent state $|\alpha\rangle_i$ with average particle number $\bar{n} = |\alpha|^2$ [121]:

$$|\Psi\rangle_{U/J=0} = \prod_i |\alpha\rangle_i = \prod_i \sum_n \frac{\alpha^n}{\sqrt{n!}} e^{-\bar{n}/2} |n\rangle_i. \quad (1.6)$$

In contrast to that, the ground state in the limit of strong, repulsive interactions ($U \gg J$) is a *Mott insulator*, where the atoms are localized at individual lattice sites. For an integer filling factor \bar{n} , the ground state is a tensor product of Fock states $|n\rangle_i$ with exactly $\bar{n} = n$ atoms on each site [121]:

$$|\Psi\rangle_{U/J \rightarrow \infty} = \prod_i \left(\frac{1}{\sqrt{n!}} \hat{a}_i^\dagger \right)^n |0\rangle. \quad (1.7)$$

This state exhibits no phase coherence between different lattice sites and is incompressible. The incompressibility originates from the existence of a finite energy gap (the *Mott gap*) in the excitation spectrum. For example, in the case of unity filling ($\bar{n} = 1$), where each site is occupied by exactly one atom, the smallest excitation is given by the addition of a second atom to one of the sites, which costs an energy U . The two phases are separated by a quantum phase transition at a critical ratio $(U/J)_c$ (in 2D, $(U/J)_c \approx 16.74$ [122]), at which the Mott gap suddenly opens (coming from the superfluid phase). In contrast to classical phase transitions, quantum phase transitions are driven by quantum fluctuations and occur also at zero temperature [123]. The direct observation of this transition in an ultracold atom experiment has been an important milestone [25], establishing the field of quantum simulation beyond discipline boundaries. Central advantages of such experiments are the high degree of control over the system parameters, the high cleanliness and homogeneity, and the possibility to directly observe the atomic distribution in the lattice using high-resolution imaging techniques.

A key development for neutral-atom based quantum simulators was the invention of *quantum gas microscopy*, which allows for the microscopic observation of individual atoms

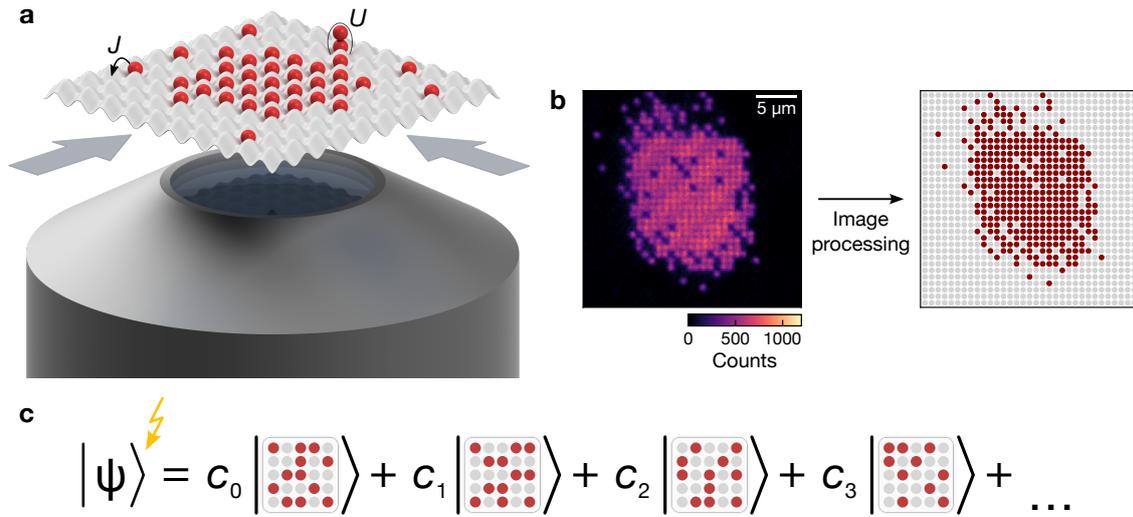


Figure 1.1 | Quantum gas microscopy. **a**, Neutral atoms in a two-dimensional optical lattice potential realize the Hubbard model, characterized by the interplay between nearest-neighbor tunneling (strength J) and on-site interactions (strength U). Using a high resolution objective, fluorescence images of the atomic distribution can be recorded. **b**, Fluorescence image, showing individual atoms as bright spots, arranged in a well-resolved periodic lattice structure (spacing $a = 767 \text{ nm}$). Using an image processing algorithm, the binary occupation in the lattice can be reconstructed. **c**, In the imaging process, the wave function is collapsed onto the Fock basis of the lattice sites. By taking many snapshots, information about the wave function can be extracted.

in the respective sites of the optical lattice [31, 38, 40]. As illustrated in Fig. 1.1a, the idea is to combine a high-resolution microscope objective with fluorescence imaging to probe the atomic distribution. To this end, the many-body wave function is first collapsed onto the local Fock basis of the lattice sites by abruptly ramping the lattices to a large depth, pinning the atoms in place. Next, the atoms are illuminated with near-resonant light, which causes them to scatter fluorescence photons that are collected by the objective and imaged onto a camera. The resulting image contains a snapshot of the density distribution, where each atom appears as a bright spot (see Fig. 1.1b). For further analysis, the integer occupation of each lattice site is typically extracted using a reconstruction method, yielding a sample of the density operator. As the many-body wave function is collapsed during the pinning step, the imaging process is destructive. However, by recording many such snapshots, each of which contains a unique sample of the density distribution (Fig. 1.1c), one can extract a wealth of information about the system. In particular, the sampling allows in principle to reconstruct the full joint probability distribution to observe a specific lattice occupation (e.g., the probabilities $|c_i|^2$ of the density patterns in Fig. 1.1c), and the evaluation of density-density correlations at arbitrary order and over arbitrary distances [32, 33, 38]. Furthermore, it is now also possible to measure off-diagonal elements (coherences) of the density matrix. This gives access to the complex amplitudes c_i , and enables probing particle currents, kinetic energies, and entanglement in many-body systems with local resolution (see Chapter 5 and Refs. [34–36, 92, 124]). Lastly, the high-resolution imaging system also allows the site-resolved manipulation using programmable light patterns to apply operations or shape the potential landscape with local control [38, 39].

1.2 Topological many-body phases

Symmetries and order parameters are cornerstone concepts in condensed matter physics and central to the classification of different phases of matter. Through symmetries, we can understand how classical and quantum many-body systems organize themselves into different phases, and how transitions between them occur. According to the Landau paradigm, the classification of phases is based on the presence or absence of symmetries in the system, which can be quantified by local order parameters [125, 126]. The behavior of such order parameters enables a detection of phase transitions, which often go in hand with the breaking of a symmetry, and to find universal features among seemingly different systems. The Landau theory proved to be extraordinarily successful, and seemed to capture all phases of matter and the transitions between them. However, attempts to explain the fractional quantum Hall effect [99, 100] and to devise models for high-temperature superconductivity [127] led to the realization that there are phases of matter that cannot be described by local order parameters or symmetry breaking. Instead, they exhibit an entirely new kind of order that is not associated with any local order parameter, but rather manifests in the distribution of global quantum entanglement. This new kind of order is known as *topological order* [128–130].

Topologically ordered phases of matter exhibit a variety of fascinating properties: a non-trivial ground state degeneracy [128, 131], fractional and non-abelian statistics of excitations (*anyons*) [101, 102, 132–135], topologically protected edge modes [136, 137], and long-range entanglement with a non-vanishing topological entanglement entropy [138–141]. Understanding these effects and their origin is a central goal of modern condensed matter physics, but has proven to be challenging both theoretically and in experimental implementations. An exciting avenue to study these phases experimentally is on quantum simulation platforms, motivated by the superior control, cleanliness and the access to novel microscopic observables (for example entanglement entropy) [106, 109].

The quantum Hall effect and lattice systems

Prototypical examples of topologically ordered phases can be found in the fractional quantum Hall effect (FQHE), which has been first discovered in two-dimensional electron systems subjected to a strong magnetic field [99, 142]. There, it was found that the Hall conductivity is quantized to rational fractions of e^2/h , where e is the electron charge and h is Planck's constant. The microscopic origin of this quantization is the interplay of the strong interactions between electrons and the topology of the Landau levels that are flattened due to the strong magnetic field. The FQHE is thus strictly a many-body phenomenon, in contrast to the integer quantum Hall effect, which can be understood in terms of the single-particle band structure [143].

The controlled preparation and study of strongly correlated many-body phases is a prime application for ultracold atoms, which would thus also motivate the study of quantum Hall physics. However, due to the lack of an electrical charge, neutral atoms do not couple to magnetic fields via the Lorentz force. The charge neutrality hence necessitates the engineering of an artificial magnetic field. To see ways how this could be achieved, it is instructive to look at the Harper-Hofstadter model [144, 145]. It was introduced to provide a microscopic

description of an electron moving in the two-dimensional ionic crystal lattice of a solid, subject to a uniform perpendicular magnetic field:

$$\hat{H} = -t \sum_{\langle i,j \rangle} \left[e^{i\phi_{i,j}} \hat{a}_i^\dagger \hat{a}_j + h.c. \right]. \quad (1.8)$$

The model describes nearest-neighbor hopping with a hopping amplitude t (the annihilation operator \hat{a}_i at site i can be bosonic or fermionic), and the magnetic field is introduced as a complex phase into the hopping amplitudes via the Peierls substitution [146]. The phases correspond to the Aharonov-Bohm phase acquired upon hopping to a nearest-neighboring site, and are given by

$$\phi_{i,j} = \frac{2\pi}{\Phi_0} \int_{\mathbf{x}_i}^{\mathbf{x}_j} \mathbf{A}(\mathbf{x}) \cdot d\mathbf{x}, \quad (1.9)$$

where $\Phi_0 = h/e$ is the magnetic flux quantum and \mathbf{A} is the vector potential. The wavefunction of a particle hopping around a plaquette in a closed loop hence acquires a non-trivial geometric phase as a result of the magnetic flux that pierces the plaquette. It is this phase, together with the periodicity of the Bloch bands in momentum space, that gives rise to a non-trivial topology of this system. As a result, the Harper-Hofstadter model (Eq. 1.8) can already capture the integer quantum Hall effect, with a quantized Hall conductance that is proportional to a topological invariant of the filled band, the Chern number [147, 148]. Moreover, the Harper-Hofstadter model can be extended to include interactions, such as the on-site Hubbard interaction (Eq. 1.5, i.e., one could call this the Harper-Hofstadter-Hubbard model). This composite model describes what are called *fractional Chern insulators* (FCI) – a lattice analog and generalization of fractional quantum Hall states [149, 150]. On the one hand, the similarity between the flat Bloch bands of the Harper-Hofstadter model and the continuum Landau levels gives rise to equivalent phases as in the continuum systems, novel bosonic variants and explicitly also the sought-after non-abelian phases such as for example the Moore-Read and Read-Razayi states [151–156]. On the other hand, the presence of the lattice structure and possibilities to tailor the Hamiltonian allows to enter the high-field regime and create entirely new FQH states with no counterpart in the low-field, continuum limit [156–160]. Lastly, the phase diagram of the above model also features interesting, strongly correlated phases beyond FQH such as chiral Mott insulators [161–165], chiral spin liquids [166, 167] and topological superfluids [168].

Realizing quantum Hall systems with ultracold atoms

In the past decade, several experimental techniques have been developed to simulate (fractional) quantum Hall physics with ultracold atoms, each with their respective advantages and limitations. By comparison between the Hubbard model (Eq. 1.5) and the Harper-Hofstadter model (Eq. 1.8), it is evident that one route involves ‘simply’ exchanging the real-valued tunnel couplings in an optical lattice with complex-valued ones. While it is not a trivial problem, after an initial theoretical proposal by Jaksch and Zoller [169], this has been achieved using *Floquet*

engineering [51, 52, 170], which is the approach we will follow in this thesis (see Section 1.3). It enables a direct realization of the Hofstadter model with tunable on-site interactions in real space, and is hence straightforwardly compatible with single-site resolved detection. However, a central limitation of Floquet engineering is the inherent susceptibility to heating, in particular, when strong interactions are present (see also discussion in Chapter 6).

A different approach, which has recently gained traction, is based on the use of synthetic dimensions [171, 172]. Here, one of the spatial dimensions of the lattice is replaced by momentum states or an internal degree of freedom [173–176]. The effective magnetic field is then introduced by coupling the states along the synthetic dimension for example via a Raman transition. In addition to being less sensitive to heating compared to Floquet engineering, this approach is also advantageous for studying edge states due to the inherently finite size of the synthetic dimension. Drawbacks are the experimental overhead due to the spin-resolved detection, which also leads to a more challenging microscopic detection, particularly for the readout of correlation functions and currents. For interacting systems, the use of multiple spin states typically results in many different interaction strengths depending on the involved states. Additionally, since atoms in different internal states are physically located on the same site of the real-space lattice, the interactions are infinitely ranged in the synthetic dimension, which naturally implements a different class of Hamiltonians.

A third noteworthy approach, albeit not directly related to lattice systems, is the use of rapidly rotating systems [177]. Here, the similarity between the Lorentz force and the Coriolis force is exploited to create an effective magnetic field, which has so far however mostly been used without microscopic detection [178, 179].

Opportunities in the quantum simulation of topologically ordered matter

While the implementation of fractional quantum Hall states in quantum simulation platforms is still in its infancy, with the state-of-the-art being two-particle Laughlin states [111, 112, 180, 181], the superior control and detection capabilities in these platforms open up a wealth of exciting opportunities. Firstly, the single-site and single-atom-resolved detection allows for a direct probing of the microscopic density correlations, and thus to reveal for example the correlated vortex motion in the Laughlin state [111, 112, 180, 181]. Through locally resolved density and current measurements (see Chapter 5), one can directly reveal edge states and their spatial distributions [55].

Furthermore, as motivated above, one of the central and most intriguing features of topologically ordered phases is the presence of anyonic excitations with fractional statistics. These could be probed in a controlled fashion for example on a fractional quantum Hall state by adding a small repulsive potential tweezer that causes an elementary quasi-hole excitation [91, 182, 183]. With the ability to manipulate the position of the potential via spatial light modulators or acousto-optic deflectors, one could then in principle move two such excitations around each other to implement a braiding operation. Achieving a manipulation of excitations is however technically rather challenging, as it has to be done without the creation of additional excitations. On the one hand, this would enable to study the exchange statistics through an interferometric measurement of the statistical phase [184]. On the other hand, when done

on a state with a non-Abelian braiding statistics, this would allow the implementation of a topological quantum gate, which constitutes the basic building block of a fault-tolerant, topological quantum computer [104, 105].

Lastly, a central advantage of quantum simulation platforms is the direct access to the distribution of entanglement and the entanglement entropy, which has particular importance for topologically ordered phases [138]. To this extent, one could prepare two identical copies of a 2D topologically ordered state, and then implement a beamsplitter operation between the copies via a vertical double well potential (see also Chapter 5). As demonstrated for 1D systems in Ref. [34], the resulting density distribution after the beamsplitter operation gives access to the entanglement entropy, which can be used to probe the entanglement structure over different length scales. An extension to large 2D systems will likely be challenging, but the increasing programmability of analog quantum simulators holds the promise to implement also further, more practical protocols to extract entanglement and proxies thereof.

In summary, this provides ways to study some of the most intriguing properties of topologically ordered phases at the ultimate limit of the individual constituents, opening up promising avenues for research that could lead to a deeper understanding of these exotic phases of matter.

1.3 Floquet engineering of artificial gauge fields

As discussed in the previous section, topological band structures can be emulated by engineering a magnetic field. In our experiment, we create an artificial magnetic field using Floquet engineering, where a system is periodically driven by an external field. As a result of the periodic drive, the long-term dynamics of the system can be described by an effective time-averaged Hamiltonian with new emergent properties. In the following, we will start by introducing the concept of Floquet engineering and then discuss how it can be used to generate artificial gauge fields in ultracold atom systems.

Floquet engineering

Floquet engineering via periodic modulation has emerged as a powerful tool for Hamiltonian engineering [114, 115]. It is based on the idea of applying a rapid, time-periodic modulation to a static system. By choosing suitable modulation parameters, the resulting dynamics of the system can be described by an effective time-independent Hamiltonian, which can exhibit novel, emergent properties that are not present in the original system. This is the essence of Floquet engineering, which has enabled the exploration of novel quantum matter across various platforms including photonics [185], neutral atoms in optical lattices [113, 114, 186], superconducting qubits [187, 188] as well as Rydberg atom arrays [189, 190], and increasingly, also solid-state systems [191, 192]. In the following, we will give a very brief introduction to the theoretical foundation of Floquet engineering, focusing on the aspects that are necessary to understand the techniques used in this thesis. For more comprehensive derivations and discussions, we refer the reader to specialized literature such as Refs. [113, 114, 186, 193].

The theoretical foundation underlying this technique is formulated in the Floquet theorem,

which applies to a time-periodic Hamiltonian satisfying $\hat{H}(t + T) = \hat{H}(t)$. The aim is to find a solution to the time-dependent Schrödinger equation

$$i\hbar \frac{\partial}{\partial t} |\psi(t)\rangle = \hat{H}(t) |\psi(t)\rangle \quad (1.10)$$

given the time-periodicity. In general, an abstract solution to this is given by the time-evolution operator $\hat{U}(t, t_0)$, which evolves a state from an initial time t_0 to a final time t via $|\psi(t)\rangle = \hat{U}(t, t_0) |\psi(t_0)\rangle$. It can be written as a time-ordered exponential

$$\hat{U}(t, t_0) = \mathcal{T} \exp\left(-\frac{i}{\hbar} \int_{t_0}^t \hat{H}(t') dt'\right), \quad (1.11)$$

where \mathcal{T} denotes the time-ordering operator. Now, by making use of the time-periodicity of the Hamiltonian, we can obtain further insights on the properties of $\hat{U}(t, t_0)$. In particular, one can show that [193–195]

$$\hat{U}(t + T, 0) = \hat{U}(t, 0) \hat{U}(T, 0), \quad \hat{U}(t + T, T) = \hat{U}(t, 0). \quad (1.12)$$

From this follows that

$$\hat{U}(nT, 0) = [\hat{U}(T, 0)]^n = [\hat{U}(T)]^n, \quad \text{with } \hat{U}(T, 0) \equiv \hat{U}(T) \text{ and } n \in \mathbb{N}. \quad (1.13)$$

Here, $\hat{U}(T)$ is the time-evolution operator over one period T . However, this must mean that the evolution of the system after multiples of the driving period T can be described by an effective time-independent Hamiltonian \hat{H}_{eff} :

$$\hat{U}(nT, 0) = \exp\left(-\frac{i}{\hbar} nT \hat{H}_{\text{eff}}\right). \quad (1.14)$$

The full time evolution is then given by

$$\hat{U}(t_f, t_i) = e^{-i\hat{K}(t_f)} e^{-\frac{i}{\hbar} \hat{H}_F(t_f - t_i)} e^{i\hat{K}(t_i)}, \quad (1.15)$$

where $\hat{K}(t)$ is a so-called *kick operator* that governs the evolution within a period of the drive. This is the general form of the Floquet theorem for a time-periodic Hamiltonian. It gives us a mathematical framework to separate the slow dynamics of the system, described by the effective Hamiltonian \hat{H}_{eff} , from the fast dynamics during one period of the drive (also called the *micromotion*). Starting from the full, time-dependent Hamiltonian $\hat{H}(t)$, it is in general not possible to obtain a closed form for the effective Hamiltonian. However, in the high-frequency limit, where the driving frequency $\omega = 2\pi/T$ is much larger than the energy scales of the system, one can obtain an approximate form for \hat{H}_{eff} perturbatively. One possibility is to use the *Floquet-Magnus expansion*, which allows expressing the effective Hamiltonian as a series in powers of the driving frequency [196, 197]. In particular, starting from a Hamiltonian of the form $\hat{H}(t) = \hat{H}_0 + \hat{V}(t)$, we first perform a unitary transformation to the rotating frame of the drive via

$$\hat{H}'(t) = \hat{R}(t) \hat{H}(t) \hat{R}^\dagger(t) - i\hbar \hat{R}(t) \partial_t \hat{R}^\dagger(t), \quad (1.16)$$

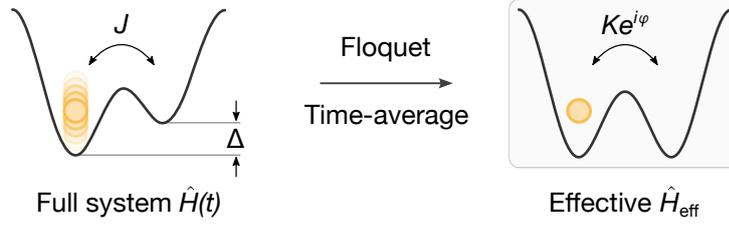


Figure 1.2 | Periodically driven double well. In a strongly tilted double well ($\Delta \gg J$), direct tunneling between the two sites is suppressed. By applying a periodic on-site modulation with phase φ , on resonance with the energy gap $\sqrt{\Delta^2 + 4J^2}$, an effective tunnel coupling is induced in the time-averaged system. The effective tunnel coupling K carries a complex phase factor $e^{i\varphi}$, which can be used to engineer an artificial magnetic field.

using the operator

$$\hat{R}(t) = \exp \left[i \sum_{\alpha} \alpha \omega t |\alpha\rangle \langle \alpha| + \frac{i}{\hbar} \int^t \hat{V}(t') dt' \right], \quad (1.17)$$

which is constructed specifically for a modulation that is resonant with an intrinsic term of the non-driven Hamiltonian [193]. The lowest order in the Floquet-Magnus expansion is simply the time-average of the transformed Hamiltonian over one period T of the drive, giving the effective Hamiltonian

$$\hat{H}_{\text{eff}} = \frac{1}{T} \int_0^T \hat{H}'(t) dt. \quad (1.18)$$

Using this relation, one can derive an approximate form for the effective Hamiltonian in the high-frequency limit. Note that if the driving frequency is not much larger than the energy scales of the system, higher-order terms in the Floquet-Magnus expansion must be considered to obtain an accurate description of the effective Hamiltonian.

Complex tunneling phase in a periodically driven double well

The fundamental building block that we will use to engineer the complex tunneling phases is a periodically driven double well (Fig. 1.2). The Hamiltonian governing the single-particle-dynamics of this system is given by

$$\hat{H}(t) = -J (|L\rangle \langle R| + |R\rangle \langle L|) + \Delta |R\rangle \langle R| + V_{\text{mod}} \cos(\omega t + \varphi) |L\rangle \langle L|. \quad (1.19)$$

Here, J is the tunnel coupling between the two wells, Δ is the energy difference, and V_{mod} , ω and φ are the amplitude, frequency and phase of the periodic modulation, respectively. In absence of the drive ($V_{\text{mod}} = 0$), the Hamiltonian has two eigenstates separated by an energy gap $E_{\text{gap}} = \sqrt{\Delta^2 + 4J^2}$. We now consider the case of a large tilt ($\Delta \gg J$), such that $E_{\text{gap}} \simeq \Delta$ and direct tunneling between the two sites is strongly suppressed. In this setting, we apply the periodic modulation on resonance with the energy gap, i.e., $\omega = E_{\text{gap}}/\hbar \simeq \Delta/\hbar$.

Evidently, this system is time-periodic with period $T = 2\pi/\omega$, and can be treated with the Floquet formalism introduced above to derive an approximate, effective Hamiltonian. To this end, we first transform the Hamiltonian into the rotating frame of the drive by applying the unitary transformation as specified in Eq. 1.17, which reads in this case

$$\hat{R}(t) = \exp \left[i\omega t |R\rangle\langle R| + i \frac{V_{\text{mod}}}{\hbar\omega} \sin(\omega t + \varphi) |L\rangle\langle L| \right]. \quad (1.20)$$

Applying this transformation, we obtain the Hamiltonian in the rotating frame as

$$\begin{aligned} \hat{H}'(t) &= -J (|L\rangle\langle R| e^{i\eta(t)} + |R\rangle\langle L| e^{-i\eta(t)}), \text{ with} \\ \eta(t) &= - \left[\omega t - \frac{V_{\text{mod}}}{\hbar\omega} \sin(\omega t + \varphi) \right]. \end{aligned} \quad (1.21)$$

Lastly, the effective Hamiltonian \hat{H}_{eff} is obtained to lowest order by averaging over one period T :

$$\begin{aligned} \hat{H}_{\text{eff}} &= \frac{1}{T} \int_0^T \hat{H}'(t) dt \\ &= -\frac{1}{2\pi} \int_0^{2\pi} \left[J |L\rangle\langle R| e^{-i\left[\tau - \frac{V_{\text{mod}}}{\hbar\omega} \sin(\tau + \varphi)\right]} + \text{h.c.} \right] d\tau \\ &=: -K e^{i\varphi} |L\rangle\langle R| + \text{h.c.}, \end{aligned} \quad (1.22)$$

where we have substituted $\tau = \omega t$ and used the definition of the first-order Bessel function of the first kind $\mathcal{J}_1(x) = \frac{1}{2\pi} \int_0^{2\pi} e^{-i(\tau - x \sin \tau)} d\tau$ to introduce the effective coupling $K = J \mathcal{J}_1\left(\frac{V_{\text{mod}}}{\hbar\omega}\right)$. Comparing this result to a bare double well Hamiltonian (Eq. 1.19), we find that the effect of the resonant, periodic modulation is to induce an effective tunnel coupling with strength K between the two sites. Importantly, the engineered coupling element now carries a complex phase factor $e^{i\varphi}$, which originates from the phase of the drive. This is the crucial ingredient that allows the engineering of an artificial magnetic field. Additionally, as shown in Fig. 1.2, the double well in the effective Hamiltonian is symmetric as a result of the resonant modulation (i.e., $\Delta_{\text{eff}} = 0$), and an effective tilt can be added by driving off-resonantly ($\hbar\omega \neq E_{\text{gap}}$).

Experimental implementation of an artificial gauge field using a running-wave lattice

We will now show how, based on periodically modulated double wells as elementary building blocks, an artificial magnetic field for ultracold atoms in an optical lattice can be engineered. In particular, we will be working with a ladder geometry, where each rung of the ladder is formed by a double well. The lattice configuration is shown in Fig. 1.3a, comprised of a tilted superlattice in the x -direction (short spacing $a_s = 383.5$ nm, long spacing $a_l = 2a_s = 767$ nm) and a monochromatic lattice in the y -direction (short spacing a_s). This realizes many independent copies of two-leg ladders. Additionally, we apply a strong tilt $\Delta \gg J_x$, such that direct tunneling along the rungs of each ladder is inhibited. As shown above, by applying a periodic potential modulation to each double well, resonant with the energy difference between the wells, we can induce an effective tunnel coupling with a complex phase factor. However, this modulation cannot be done globally in-phase (which would be the case for example when

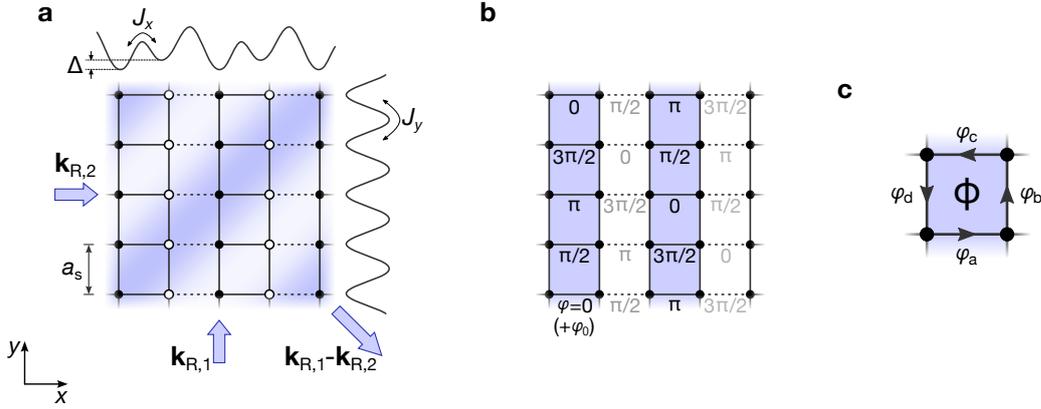


Figure 1.3 | Running-wave modulation scheme. **a**, Schematic of the two-dimensional lattice geometry, consisting of a tilted superlattice in the x -direction, and a monochromatic lattice in the y -direction. Superimposed is a diagonal running-wave lattice, created through the cross-interference of two laser beams (purple arrows) aligned with the axes of the physics lattice. **b**, The periodic modulation of each rung-double-well induces complex hopping phases along the rungs, which vary spatially as a result of the spatially varying phase of the running-wave lattice. By choice of the geometry, the phases change by $\pi/2$ per bond in both directions. Additionally, a global phase offset ϕ_0 is added to each bond phase, which randomly varies in each experimental realization. **c**, Zoom-in to one plaquette. The flux per plaquette Φ is determined by the sum of the bond phases φ_i , $i = \{a, b, c, d\}$ around the plaquette.

modulating the superlattice depths), as this would lead to a vanishing flux. To apply a periodic modulation with a spatially varying phase, we employ a running-wave lattice that diagonally traverses the lattice system (see Fig. 1.3a), created by interfering two laser beams aligned with the axes of the physics lattice (wave vectors $\mathbf{k}_{R,1} \simeq \mathbf{k}_{R,2} \equiv \mathbf{k}_R$, frequencies ω_i and phases ϕ_i , $i = \{1, 2\}$). This creates an attractive, time-dependent potential of the form

$$V_{\text{mod}}(\mathbf{r}, t) \propto V_0 \cos(\delta\mathbf{k} \cdot \mathbf{r} + \omega t + \phi_0), \quad (1.23)$$

where $\delta\mathbf{k} = \mathbf{k}_{R,1} - \mathbf{k}_{R,2}$ is the wave vector difference, $\omega = \omega_1 - \omega_2$ the frequency difference and $\phi_0 = \phi_1 - \phi_2$ a global phase offset. The time dependence ωt results in a time-periodic modulation of each rung double well, hence, by setting $\omega = \sqrt{\Delta^2 + 4J_x^2}/\hbar$, we can induce an effective tunnel coupling in the rung direction. In contrast, the scalar product $\delta\mathbf{k} \cdot \mathbf{r}$ leads to a spatially varying phase of the modulation throughout the lattice at each point in time, which crucially determines the resulting flux. The position \mathbf{r} of each lattice site can be expressed as $\mathbf{r} = ma_s\mathbf{e}_x + na_s\mathbf{e}_y$, where $\mathbf{e}_{x(y)}$ are the unit vectors in the $x(y)$ -direction, and we obtain

$$\delta\mathbf{k} \cdot \mathbf{r} = k_R a_s (m - n), \quad (1.24)$$

where $k_R = |\mathbf{k}_R|$. This means that the phase changes by $k_R a_s$ per bond in both directions, hence, with appropriate choice of the running-wave geometry with respect to the underlying physics lattice, an arbitrary flux can be engineered. In our case, we choose the wavelength of the running-wave beams to be $\lambda_R = 1534$ nm, which is equivalent to $k_R = \pi/(2a_s)$ and thus the tunneling phase changes by $|\delta\varphi_x| = |\delta\varphi_y| = \pi/2$ per bond (Fig. 1.3b). The resulting flux

per plaquette is then found by summing up the respective tunneling phases on the four bonds comprising the plaquette (Fig. 1.3c). Since there is no induced tunneling in the y -direction, we have $\varphi_b = \varphi_d = 0$, and hence, the magnitude of the flux per plaquette is

$$|\Phi| = |\varphi_c - \varphi_a| = \pi/2. \quad (1.25)$$

The sign of the flux (i.e., if the artificial magnetic field points in $+z$ - or $-z$ -direction) can be changed by inverting either the tilt or the frequency ($\Delta \rightarrow -\Delta$ or $\omega \rightarrow -\omega$, respectively).

Compared to the individual double well example, there are two minor differences: Firstly, in this scheme, both sites of a double well are modulated. Hence, the effective tunnel coupling K depends on the differential modulation depth, which introduces another factor of $\sin \delta\varphi_x$ and for our case of $\delta\varphi_x = \pi/2$ leads to

$$K = J_x \mathcal{J}_1 \left(\frac{\sqrt{2}V_0}{\hbar\omega} \right). \quad (1.26)$$

Secondly, while there is no driven tunneling in the y -direction, a treatment of the full 2D system with the Floquet-Magnus expansion shows that the time-periodic modulation leads to a small renormalization of the tunneling in the y -direction to [193]

$$J = J_y \mathcal{J}_0 \left(\frac{\sqrt{2}V_0}{\hbar\omega} \right). \quad (1.27)$$

However, in most experiments we work in the weak-driving limit, where $J \approx J_y$.

Floquet heating

A periodic modulation of the system necessarily implies that there is a constant energy input into the system, which in general leads to unwanted heating. This is one of the major challenges in Floquet engineering, as it is inevitable that the system will eventually heat up to a featureless, infinite temperature state [198, 199]. To observe the physics of interest in the engineered system, it is necessary that the heating rate is sufficiently small to enable the existence of a pre-thermal regime, where the system can be described by the effective Hamiltonian for a finite time [200–203]. The heating rate is crucially determined by the driving frequency ω in comparison to the energy scales of the system. In particular, the frequency has to be chosen low enough to avoid resonant excitations to higher orbital bands, but high enough with respect to the processes in the ground band such as tunneling and on-site interaction [204].

The time-periodicity of the system relaxes energy conservation, as the eigenstates of the Hamiltonian are associated with a quasi-energy that is defined only modulo $\hbar\omega$ (analogous to the quasi-momentum of the Bloch states), allowing the absorption of energy quanta $\hbar\omega$ from the drive. This can cause, particularly for interacting systems, photon-assisted inelastic scattering processes [205], motional excitations, particularly in a transverse direction [206], as well as collective and many-body excitations [207–210]. To manage these effects experimentally, a flexible control over the parameters of the drive and the confinement is paramount.

CHAPTER 2

Experimental setup

In this chapter, we will give an overview of the experimental apparatus used for the studies that are presented throughout this thesis. The earlier parts of the sequence, including pre-cooling, optical transport and evaporation, have already been described in detail by earlier theses, so they will only be summarized briefly here [211–213]. We will mostly focus on the optical lattice setups, the high-resolution imaging system, and the laser-assisted tunneling setup, which have been changed or newly implemented for the experiments presented in this thesis.

2.1 A quantum simulation platform based on ultracold cesium atoms

The first elements that were successfully laser-cooled and brought to quantum degeneracy were sodium and rubidium, both alkali metals. The recipe for their success was the simple level structure due to the single valence electron, featuring several closed transitions that facilitate laser cooling and trapping, as well as favorable and known scattering properties. In the following decades, the toolbox of ultracold atomic physics has been expanded to include ever more exotic species on the hunt for specific properties such as ultra-narrow transitions, strong dipolar interactions, or a large magnetic moment. Nevertheless, alkali metals are subject to continued interest, with new apparatuses based on them being built until today.

Among the alkali metals, cesium, with its single stable, bosonic isotope ^{133}Cs , holds a unique position: While it has long played a crucial role in physics – as the foundation of our primary frequency standard [214] and a key element in fundamental metrology [215, 216] – it was the last alkali metal to reach quantum degeneracy [217]. This delay was caused by its unusually complex low-energy scattering properties [216, 218, 219], which initially made it challenging to reach the required phase-space densities using standard magnetic trapping techniques. However, since the routine realization of Bose-Einstein condensation in cesium, a few distinct advantages over other alkali metals have become increasingly evident, ultimately motivating the construction of this new apparatus in 2017.

One of the key advantages that make cesium particularly appealing for our purposes is the very same complex scattering behavior that initially made it difficult to cool. Specifically, the

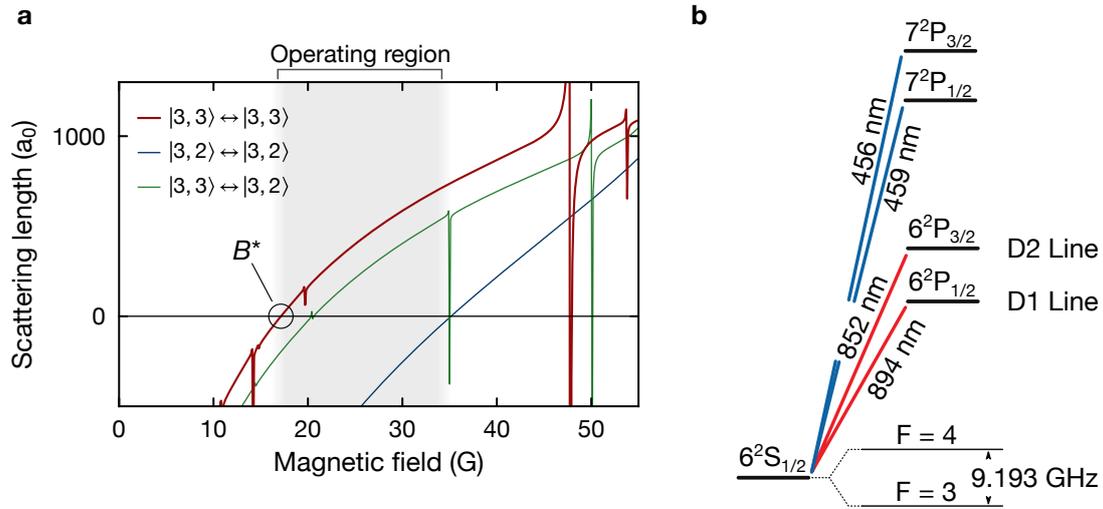


Figure 2.1 | Properties of cesium. **a**, Scattering length of selected spin states in the ground state manifold as a function of the magnetic field strength. Most experiments using a single spin component are performed with the absolute ground state $|F = 3, m_F = 3\rangle$, which has a zero-crossing of the scattering length at $B^* = 17.12$ G. We typically work at a field strength between 20 – 35 G, where the interactions are repulsive, and we can smoothly tune the scattering length between $100 a_0$ and $600 a_0$ (shaded region). **b**, Energy levels of the cesium atom. The ground state manifold ($6S_{1/2}$) consists of two hyperfine states, $|F = 3, m_F = 3\rangle$ and $|F = 4, m_F = 4\rangle$, which can be connected using a microwave transition. The D1 and D2 lines are the transitions used for laser cooling and trapping. There is an additional blue line to the $7P_{3/2}$ state, which is a future candidate for high-resolution imaging and single-site addressing.

presence of multiple narrow and broad Feshbach resonances¹ at low magnetic fields allows for exceptional tunability of the interatomic interaction strength. Fig. 2.1a shows the scattering length of a few selected spin states of cesium as a function of the magnetic field strength. We focus first on the absolute ground state $|F = 3, m_F = 3\rangle$, which we use to create a Bose-Einstein condensate and to perform most of our experiments involving a single spin component. The scattering length shows a pronounced dependence on the magnetic field strength, which is mainly determined by a broad Feshbach resonance centered at -11.7 G [221]. This causes the scattering length to smoothly vary from large negative values at zero field to around $1000 a_0$ at 45 G ($a_0 = 0.0529$ nm is the Bohr radius), with the zero-crossing occurring at $B^* = 17.12$ G [222]. Such a large tunability at low field strengths is unique among the alkali metals and a strong asset for the quantum simulation of interacting matter. It allows us to smoothly vary the interaction strengths from strongly attractive, over non-interacting, to strongly repulsive. At the same time, the low field strengths simplify the technical requirements on the coils and the field stabilization, and further enable fast changes of the interaction strength to study dynamical processes. Additional narrow Feshbach resonances, most notably at 19.84 G (5 mG width) and 47.97 G (120 mG width) [220], provide a second tuning knob to rapidly adjust the scattering length to extreme values.

¹A Feshbach resonance is a magnetically tunable scattering resonance between two particles, which occurs whenever the energy of the two free particles is tuned into resonance with a bound molecular state, and results in a strong variation of the s -wave scattering length [220].

We perform most of our experiments at a field strength between 20 – 35 G, where the interactions are repulsive, and the scattering length can be smoothly tuned between around $100 a_0$ and $600 a_0$. This six-fold variation in the interaction energy independent of the optical lattice parameters is a significant advantage for the study of strongly correlated matter, and is used in this thesis for example to investigate the behavior of local currents in the strongly-correlated Mott-Meissner phase over a wide range of interaction strengths (see Chapter 6). Additionally, the zero-crossing point is frequently used to turn off interactions, which is for example crucial to perform the local double-well manipulations presented in Chapter 5. Lastly, the existence of further spin states in the ground state manifold with equally favorable tunability, most notably the $|F = 3, m_F = 2\rangle$ state shown in Fig. 2.1a, opens up a promising path to realize spin models such as the Heisenberg XXZ model with tunable anisotropy [42].

One drawback following from the specific scattering properties of cesium is its relatively high loss rates. While two-body loss processes are prohibited in the absolute ground state ($|3, 3\rangle$), three-body loss rates are at least an order of magnitude higher than in other alkali metals. Due to Efimov physics, the three-body loss coefficient assumes a minimum at an offset field around 21 G [218], making this a favorable operating point. However, we observed that the loss rate can increase significantly at higher magnetic fields, particularly when the density is enhanced by a strong vertical confinement. This effect could impose practical limitations on the range of accessible interaction strengths.

A second interesting feature of cesium can be found in the internal level structure, which is shown in Fig. 2.1b. The main optical transitions used for laser cooling and trapping are the D1 and D2 lines, which connect the ground state manifold $6S_{1/2}$ to the $6P_{1/2}$ and $6P_{3/2}$ states. With their wavelengths being around 894 nm for the D1 line and 852 nm for the D2 line, cesium has by far the largest fine-structure splitting of 42 nm among all stable alkali metals. This opens up unique possibilities for the realization of state-dependent lattices with low heating rates. Particularly interesting for our experiment is a novel scheme to realize artificial gauge fields that is based on an anti-magic wavelength lattice². The anti-magic wavelength depends on the states involved and the polarization of the light, but generally lies between the D1 and D2 lines. Using the $|F = 3, m_F = 3\rangle$ and $|F = 3, m_F = 2\rangle$ states, as well as circularly polarized light, the anti-magic condition is fulfilled for example at 871.1 nm, which is detuned by around 20 nm from both lines, hence resulting in much lower off-resonant scattering rates than in other alkali metals. Further details on this scheme can be found in Section 6.7.

The last property of cesium that is noteworthy is the large mass of 133 atomic units, which has implications particularly for high-resolution imaging. Due to the large mass, the tunneling rate between adjacent lattice sites is significantly reduced compared to lighter elements, which requires small lattice spacings to keep the timescale for the dynamics fast enough compared to typical decoherence times. In combination with the long wavelength of the D2 line that is used for imaging, this results in a situation where the separation between lattice sites is significantly smaller than the imaging resolution. This makes the reconstruction of the occupation in the individual lattice sites particularly challenging, and is the main reason why we had to develop an entirely new reconstruction method (see Chapter 4). In the future, this could be partly

²At the anti-magic wavelength, two specific states experience a polarizability equal in magnitude but opposite in sign.

mitigated by performing fluorescence imaging using the blue transition to the $7P_{3/2}$ state, whose smaller wavelength would bring the resolution closer to the lattice spacing.

2.2 Overview of the apparatus and the cooling stages

An overview of the experimental apparatus is shown in Fig. 2.2. The main component is a large vacuum chamber, which is subdivided into the cesium oven section ①, a Zeeman slower ②, a magneto-optical trap (MOT) chamber ③, as well as the science chamber that consists of a glass cell ④ and a surrounding magnetic field coil and objective assembly. The glass cell has a dodecagonal shape, featuring 12 side viewports with 12 mm diameter as well as large, 30 mm diameter bottom and top viewports for high optical access. Additionally, a separate rubidium source and 2D MOT are connected to the MOT chamber, but have not been used so far. In the following, we will briefly walk through the different steps involved in the preparation of a cesium Bose-Einstein condensate. For detailed information on the pre-cooling, optical transport and evaporation stages, the reader is referred to Refs. [211–213].

Pre-cooling

An experimental realization starts in the oven section ①, where a small sample of cesium metal is continuously heated to around 70°C . This increases the vapor pressure sufficiently to create a beam of cesium atoms, which is collimated by two circular apertures and guided in the direction of the Zeeman slower ②. The Zeeman slower provides a deceleration along the propagation direction using a combination of magnetic fields and a counter-propagating laser beam (traveling in y -direction) which is red-detuned from the D2 line. The atoms are then captured in a 3D magneto-optical trap (MOT) in the center of the MOT chamber ③. The MOT is formed by six counter-propagating laser beams, red-detuned from the $F = 4 \rightarrow F' = 5$ transition, and a pair of anti-Helmholtz coils that provide a magnetic field gradient. Additionally, an optical repumper beam, resonant with the $F = 3 \rightarrow F' = 4$ transition, is used to maintain all atoms in the MOT cycling transition. After loading the MOT for five seconds, the cloud is first compressed (cMOT stage) and then undergoes an optical molasses cooling stage to lower the temperature to around $2\ \mu\text{K}$.

As a last pre-cooling step, we apply degenerate Raman sideband cooling (dRSC). Using an optical lattice with resolved vibrational levels, the sample is cooled by driving cycles consisting of a Raman transition with following spontaneous decay, which lowers the vibrational energy of the atoms. This works particularly well for cesium and results in around 2×10^7 atoms at a temperature of around $1\ \mu\text{K}$. Additionally, after dRSC we obtain a spin-polarized sample with almost all atoms in the $|F = 3, m_F = 3\rangle$ state. Since there is no trapping force after switching off the MOT, we further use a large volume optical dipole trap (called the reservoir trap) in conjunction with a magnetic levitation gradient to hold the atoms in the center of the MOT chamber. The reservoir trap is formed from a single beam (wavelength 1064 nm, 1 mm diameter, around 15 W power), which goes through the MOT chamber and is then recycled for a second pass in a bow-tie configuration. This concludes the pre-cooling stage.

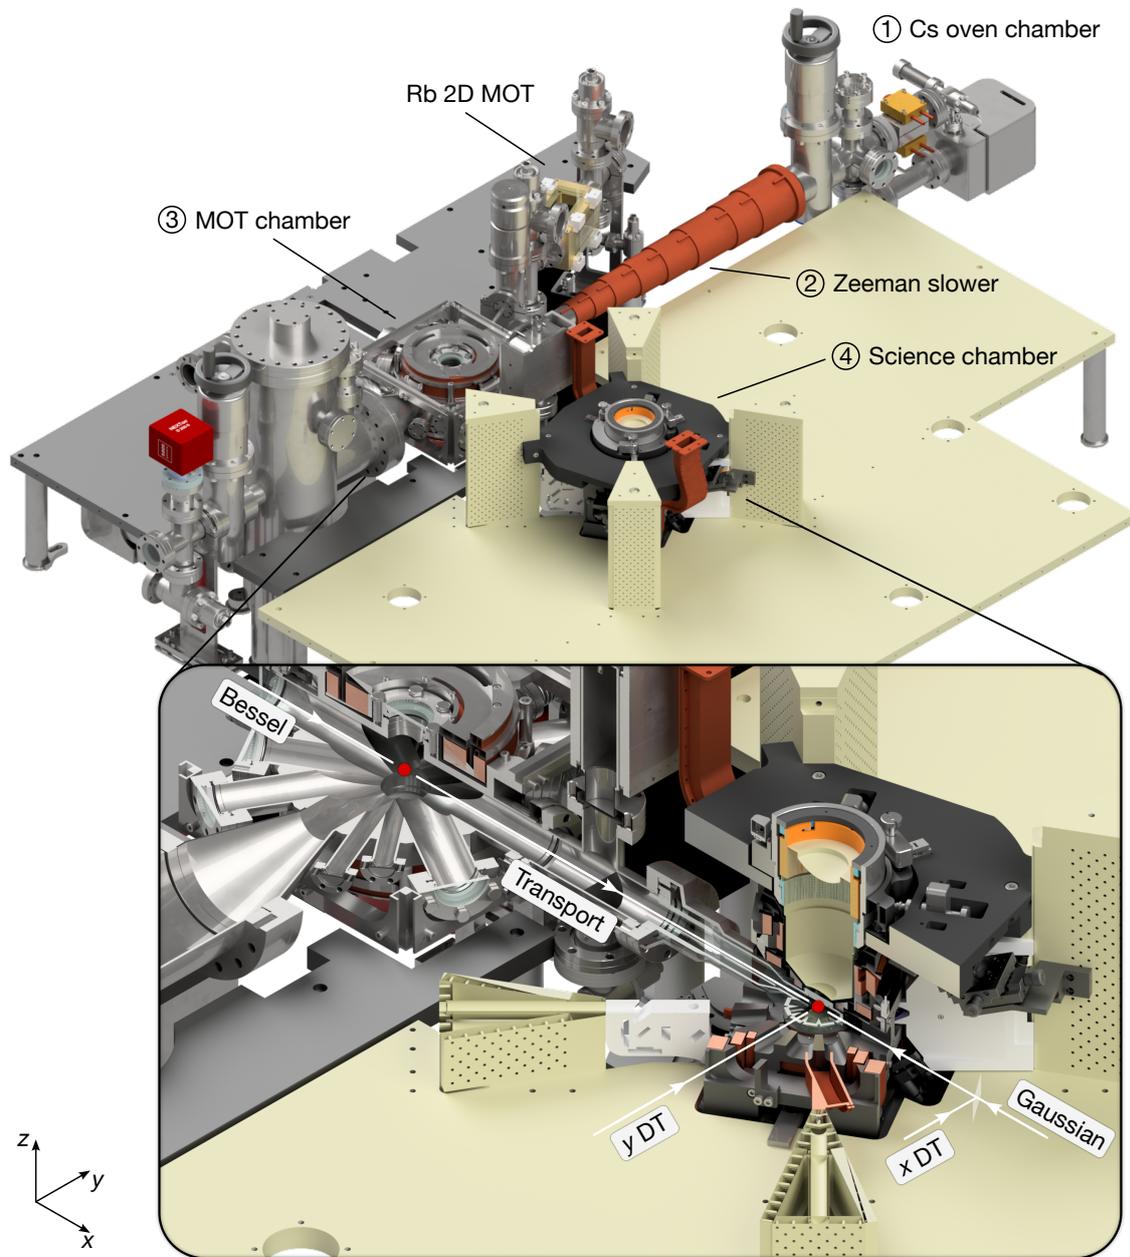


Figure 2.2 | The experimental apparatus. Overview of the experimental apparatus including the breadboards on the main level. The vacuum system consists of an oven section, which is connected to the MOT chamber via a Zeeman slower, and a science chamber, to which the atoms are optically transported from the MOT chamber. Around the science chamber is an integrated assembly that houses a set of magnetic field coils as well as two high-resolution objectives from below and above. Note that the top-level breadboard is not shown in this schematic. The zoom-in shows a cutout focusing on the transport path and the science chamber. The arrows denote the Gaussian and Bessel beams that form the transport lattice, as well as the x and y dipole traps (DT) that are used for the evaporation.

Optical transport

After the pre-cooling stage, the atoms are transported over a distance of 43 cm to the science chamber ④. The transport is implemented using an optical running-wave scheme [223]. To this

end, we add a Bessel-shaped beam (generated using an axicon, 1064 nm wavelength, central spot size around $100\ \mu\text{m}$) from the MOT side (see inset in Fig. 2.2), causing atoms to accumulate in its central spot. Using this 'dimple trick', the density can be increased significantly, while retaining a low temperature as the atoms in the Bessel potential can thermalize with the surrounding reservoir cloud [217]. After a short thermalization time, the reservoir trap is adiabatically removed and the Bessel beam is ramped up to the full power of around 30 W. From the opposing side of the science chamber, we superimpose a Gaussian beam (wavelength 1064 nm, around 6.5 W power) which initially has zero detuning with the Bessel beam, forming a static, axial lattice. To initiate the transport, a relative detuning between the Bessel and Gaussian beam is introduced, causing the nodes of the transport lattice to move, carrying the atoms along. Due to the strong axial confinement along the transport direction, high accelerations and transport velocities can be applied, which allows for a fast transport of the atoms over the 43 cm distance in around 25 ms. At the end of the transport, we typically have around $1.0 - 1.5 \times 10^6$ atoms at a temperature of around $2\ \mu\text{K}$ in the science chamber.

Evaporation in the science chamber

In the science chamber, the atoms are loaded from the transport lattice into a crossed dipole trap, which is formed by two laser beams with a wavelength of 1064 nm (see inset of Fig. 2.2). The x -dipole trap has a waist of $100 \times 100\ \mu\text{m}$ and an initial power of 1.0 W, and the y -dipole trap has a waist of $370 \times 110\ \mu\text{m}$ (horizontal, vertical) and an initial power of 7.0 W. The waists are chosen to match the shape and size of the transport cloud. After loading the crossed dipole trap, we perform forced evaporation by ramping down the power of the dipole traps exponentially to $\sim 110\ \text{mW}$ in x -, and $2.7\ \text{W}$ in y -direction, respectively, over a duration of 6 s. At the end of the evaporation, we obtain a pure Bose-Einstein condensate with around 2×10^4 atoms.

Magnetic field coils and microwave

As shown in Fig. 2.2, the glass cell is enclosed by a cube structure that houses a set of magnetic field coils in direct proximity to the atoms. Specifically, there are two coil pairs aligned with the horizontal molasses beam (see Fig. 2.3), one coil pair perpendicular to it in the horizontal direction, and two further pairs in the vertical z -direction. The coils can be freely configured in Helmholtz or anti-Helmholtz configuration, to (simultaneously) generate offset and gradient fields in an arbitrary direction. Additionally, we employ an active magnetic field stabilization system, which uses three vectorial field sensors in proximity to the glass cell in conjunction with large compensation coils around the experiment table to stabilize the environmental field fluctuations to below $100\ \mu\text{G}$.

To drive microwave transitions, we employ a total of two microwave guides that point directly at the glass cell (see the copper waveguides in Fig. 2.2), one for cesium, and one for rubidium. The cesium waveguide is optimized to drive transitions between the $|F = 3, m_F = 3\rangle$ and $|F = 4, m_F = 4\rangle$ states of the ground state manifold at around 9.2 GHz. Further details on the coil design, the magnetic field stabilization, and the microwave setup can be found in Ref. [212].

2.3 Optical lattice setups

In this section, we will describe the vertical and horizontal optical lattice setups, which are added immediately after the preparation of the BEC. We currently have a total of two vertical lattices as well as two horizontal superlattices that form a square lattice geometry. The beam geometry of all important trap and lattice beams in the science chamber is shown in Fig. 2.3.

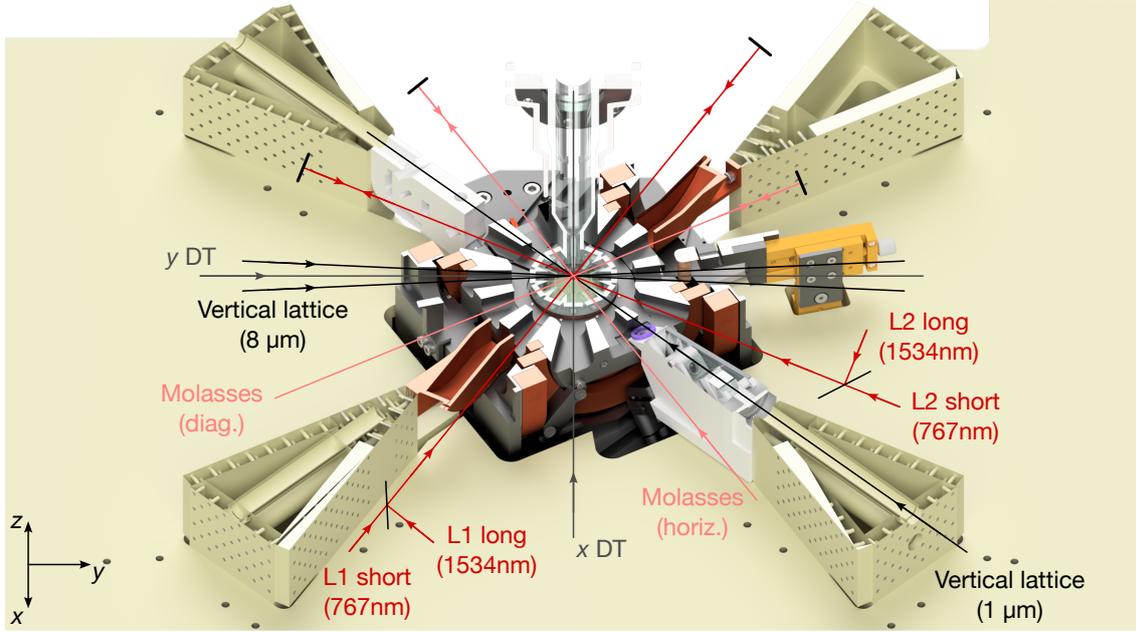


Figure 2.3 | Beam geometry in the science chamber. Overview of the beam geometry in the science chamber. There are two vertical lattices operating at 1064 nm. The long-spacing lattice (spacing $d = 8 \mu\text{m}$) is formed from two beams that co-propagate with the y -DT and interfere under a small half-angle of 3.6° . The short-spacing lattice (spacing $d = 1 \mu\text{m}$) enters through the bottom right pillar, and is formed through the interference of two beams that are out of the horizontal plane by 30° (indicated is the common-mode path). The two horizontal superlattice axes (L1 and L2) create a square geometry, and each consist of a short-spacing lattice (spacing $d = 383.5 \text{ nm}$) and a long-spacing lattice (spacing $d = 767 \text{ nm}$). Both colors of each axis are retro-reflected on a common mirror. For the fluorescence imaging, there are two retro-reflected molasses beams at 852 nm wavelength, one enters horizontally, the other diagonally at an angle of 30° out of the horizontal plane (shown is the projection into the horizontal plane).

2.3.1 Vertical lattices

The two vertical lattices in our experiment are formed in both cases by interfering two beams at an angle, which results in a one-dimensional optical lattice with a spacing of

$$d = \frac{\lambda}{2 \sin(\theta/2)}, \quad (2.1)$$

where λ is the wavelength of the light and θ is the angle enclosed by the two beams (i.e., $\theta/2$ is the half-angle which we often quote here). The central difference between the two lattices is the use of a different interference angle: The *shallow-angle vertical lattice* uses a small interference

angle to create a large lattice spacing, while the *steep-angle vertical lattice* uses a significantly larger angle to create a much smaller lattice spacing.

Shallow-angle vertical lattice

As can be seen in Fig. 2.3, the two beams forming the shallow-angle vertical lattice are co-propagating with the y -dipole trap beam. They initially propagate in parallel with a separation of 12.7 mm, and are then focused by an $f = 100$ mm focal length lens, resulting in an interference at the position of the atoms under a half-angle of 3.6° , and an elliptical waist of $(w_v, w_h) = (35, 170) \mu\text{m}$ ($1/e^2$ radius). This creates a one-dimensional, vertical lattice with a particularly large lattice spacing of $d = 8.4 \mu\text{m}$. The reason for this choice is that we load the lattice directly after the BEC preparation, and the large spacing allows us to load almost all atoms into a single plane of the lattice, resulting in a large 2D system that is required for the further steps of the sequence. We also use this lattice to provide the z -component of the deep pinning potential required during fluorescence imaging, yielding a maximal depth of around $500 \mu\text{K}$ at a power of 14 W. For a long time, this lattice was the only vertical lattice in the experiment, and has in particular been used for the experiments presented in Chapters 4 and 5. Further information on the beam forming and accompanying laser setup can be found in Refs. [211, 212].

Steep-angle vertical lattice

After the first experiments that we performed using exclusively the shallow-angle vertical lattice, we noticed a central shortcoming of the large spacing for the study of strongly-correlated systems: The large spacing results in a weak harmonic confinement along the vertical direction, where we obtain a vertical trap frequency of $\omega_z = 2\pi \times 1.2$ kHz for a typical power of 100 mW. Due to the existence of *confinement-induced resonances* (CIRs), the maximum on-site interaction energy is limited to the smallest trap frequency [224]. Hence, the maximum interaction energy is on the order of $U/h \sim 1$ kHz, which poses a significant limitation when working in the hard-core limit ($U \gg J$), as it restricts the use of large tunnel couplings. This also constrains experiments that require a wide tunability of the interaction strength. Experimentally, it is challenging to increase the trap frequency, and hence the upper limit on U , due to the unfavorable scaling of the harmonic oscillator length $a_{\text{HO}} = \sqrt{\hbar/(m\omega)} \sim P^{-1/4}$ with the lattice power P . Additionally, saturating the bound and working close to a CIR leads to increased atom loss and heating due to the high offset fields, as well as doubly-occupied sites becoming resonant with higher vibrational levels in the z -direction. In preparation for the study of strongly-correlated flux ladders (see Chapter 6), we hence integrated a second vertical lattice, which was luckily already planned for in the original design of the experiment.

The steep-angle vertical lattice is formed by two beams that are out of the horizontal plane by 30° , which results in a lattice spacing of $d = 1 \mu\text{m}$. The eight-fold smaller spacing yields vertical trap frequencies around $\omega_z = 2\pi \times 5$ kHz even for small powers ~ 20 mW, which is a large improvement over the shallow-angle lattice and allows us to reach the hard-core regime with significant tunnel couplings, while being sufficiently far away from any CIRs. Due to

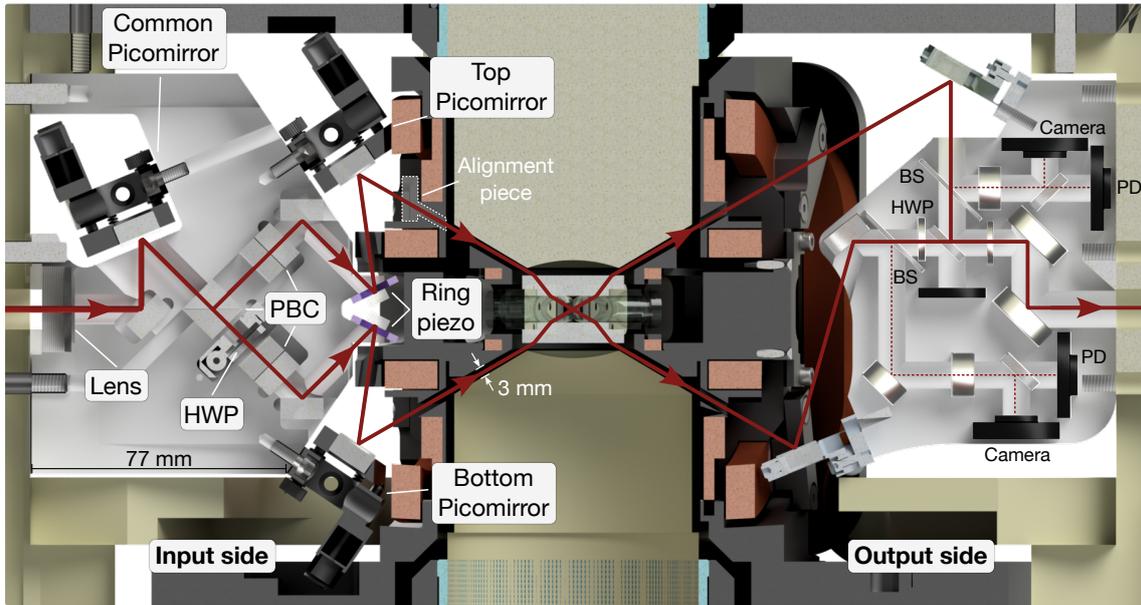


Figure 2.4 | Steep-angle vertical lattice (STVL) input and output assemblies. Cut through the science cell assembly along the propagation plane of the STVL. The input assembly is a monolithic Macor piece, in which a single input beam is focused using an $f = 250$ mm lens and split into two beams that are then guided to the atoms. Three picomotor mirrors (one common, two relative) allow fine alignment of the beams, and ring piezos are used to control the phase of the interference pattern. On the output side, a 3D printed mount captures the beams and guides them out of the experiment for safe dumping. Inside the mount, a small fraction of the beams is picked off for monitoring using cameras. Abbreviations: HWP - half-wave plate, PBC - polarizing beamsplitter cube, BS - beam sampler, PD - photodiode.

the large enclosed angle that is required to create such a small spacing, the beams have to be separated by a significant distance, which imposes stringent requirements on the beam splitting geometry to ensure a stable interference pattern. To this end, a single input beam is split into two beams and guided to the atoms using an integrated optics mount that is fabricated from a monolithic, shielded Macor³ piece (see Fig. 2.4). Tracing the beam path starting from the left, the incident collimated beam, which enters through a hole in one of the vertical pillars, first goes through an $f = 250$ mm lens that focuses it onto the atomic sample. Next follows a picomotor mirror that allows common-mode adjustments of the lattice, before the beam is split into the two components via a polarizing beamsplitter cube (PBC). The individual beams then each go through another PBC to clean up the polarization to be horizontal (i.e., into the plane in Fig. 2.4), and are finally guided to the atoms via a set of mirrors. One of the mirrors in each path is located in a picomotor mount, which allows fine adjustments of the individual beam pointing. Lastly, a set of ring piezo actuators is used to introduce a path length difference between the two arms, by which the phase of the interference pattern, i.e., the z -position of the potential nodes, can be adjusted. Further information on the design considerations for the mount can be found in Ref. [212].

On the output side after the glass cell, the beams have to be captured and guided out of

³Macor is a machinable glass-ceramic with excellent thermal expansion properties.

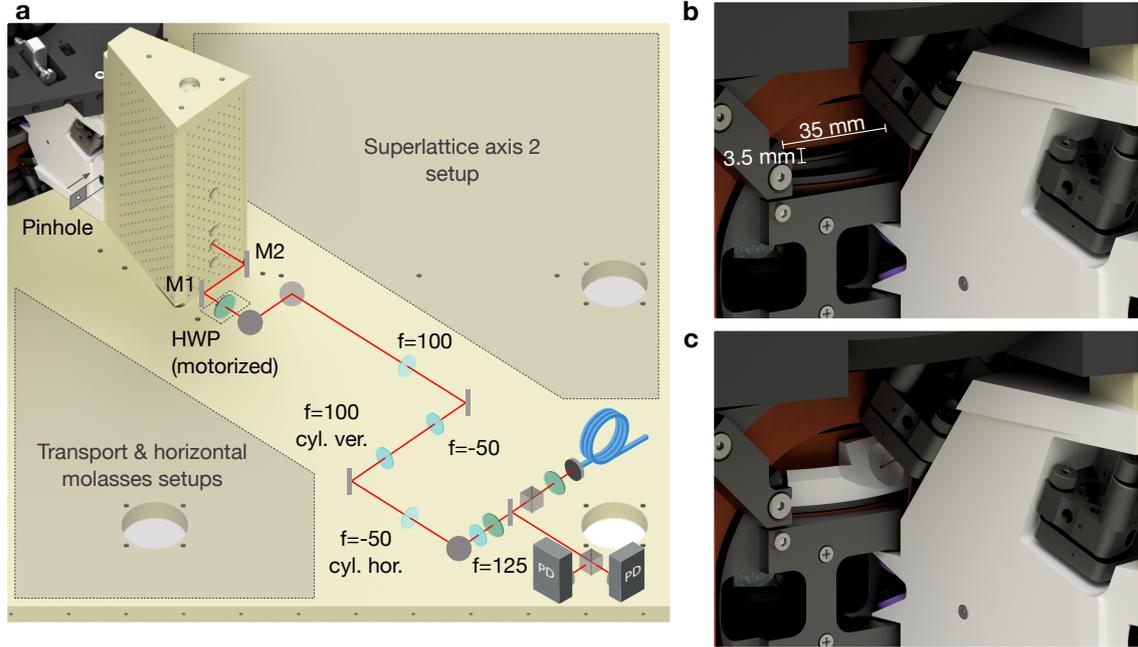


Figure 2.5 | Beam shaping and alignment of the STVL. **a**, Overview of the STVL beam forming setup. The 1064 nm light out of a photonic crystal fiber is polarization-cleaned, intensity-stabilized, and shaped elliptically using two telescopes. The motorized half-wave plate (HWP) enables the adjustment of the beam balancing. **b**, Geometry of the slits in the glass cell coil assembly through which the STVL beams have to be guided. **c**, 3D printed alignment piece for the STVL. The piece is inserted into the entrance and exit slits of the coil mount and has a small pinhole at the ideal beam position. A power meter or camera can additionally be mounted (via the recessed, SM05 compatible hole) for profiling the beam on the exit side.

the experiment for safe dumping. To this end, we manufacture a mount that is 3D printed out of the high-temperature resistant polymer Ultem. The beams are picked up and guided by a set of mirrors through the pillar opposite of the input side and then out of the experiment, where they are dumped. Additionally, the mount houses cameras using which we can monitor the beam profiles and position, allowing us to catch drifts or misalignments of the lattice beams during operation.

At the position of the atoms, the lattice beams have a waist of $(w_v, w_h) = (56, 107) \mu\text{m}$ ($1/e^2$ radius), which results in a radially isotropic in-plane confinement due to the projection under the 30° angle of incidence. To engineer the waists, we employ a beam forming setup in front of the pillar, which is shown in Fig. 2.5a. The light for the steep VL is generated from a Mephisto MOPA laser operating at a wavelength of 1064 nm and a maximum output power of 50 W. The light is delivered to the experiment breadboard via a photonic crystal fiber (NKT Photonics LMA-PM-15), and collimated using an $f = 18.4$ mm lens to a waist radius of 1.0 mm. Following polarization cleaning and a split-off of a small fraction of the light for intensity stabilization, the beam is shaped using a sequence of two telescopes. The first telescope consists of cylindrical lenses to set the correct ellipticity of 2 : 1, and the second telescope implements an additional $2\times$ magnification to reach the target waist size of $(w_v, w_h) = (1.60, 0.82)$ mm in front of the

pillar. The final $f = 250$ mm lens at the input of the STVL mount then focuses the beam to the position of the atoms. To set the beam balance between the two arms, a motorized half-wave plate (HWP) is added in front of the pillar, enabling a remote adjustment of the power ratio.

Coarse alignment of the steep-angle vertical lattice

Due to the tight geometric constraints as well as the lack of proper access to diagnose the beam pointing after it has entered the pillar, the alignment of the steep-angle vertical lattice has turned out to be particularly challenging. Here, we want to briefly outline the alignment procedure that helped us to bring the lattice into operation. A central aspect is to disentangle the different degrees of freedom (DOF), namely, the three picomotors inside the mount, as well as the mirrors in front of the pillar. If all mirrors are simultaneously touched on the single objective to have the beams emerge from the other side of the glass cell, one will never succeed due to the tight geometric constraints. It is thus essential to consequently start from the first element, align it on a reference point shortly after, and then work downstream through the optical path.

In particular, we use the mirror M1 (see Fig. 2.5a) to center the beam on the input hole of the pillar, and then adjust mirror M2 to center the beam on a 3D printed pinhole which we insert at the entrance of the STVL mount right before the $f = 250$ mm lens. After walking these DOF, the input beam should go centered through the tube in the pillar and enter the mount on its optical axis, and in particular, hit the common picomirror in the center, which forms the basis for the further alignment (cf. also Fig. 2.4). The remaining DOF afterward are the two individual picomotors, which are placed right in front of the entrance slits of the coil mount (see Fig. 2.4). Due to this proximity, their pointing minimally affects the beam position at the entrance slits, and hence we can use the relative picomotors to center the beams on the slits on the *exit* side. From this follows that we have to use the common picomirror to center the beams on the *entrance* slits. To facilitate this in light of the strongly elliptical shape of the slits (see slit geometry in Fig. 2.5b), we add 3D printed alignment pieces with a small pinhole at the ideal beam position, which we can slide into the slits. An example of the alignment piece is shown in Fig. 2.5c. Since the common picomotor affects both beams, we have to iterate with the last mirror in front of the pillar (M2) to have both the bottom and the top beam pass through the entrance alignment pieces. If the beams do not exit yet, a good observable for this is to look with an infrared-sensitive endoscope camera onto the surface of the alignment pieces and center the scattered light on the pinholes. Lastly, we use the individual picomotors to center the beams on the exit slits, which is again facilitated by the alignment pieces. Additionally, we found that it is best to watch the beams on the exit side right after the slits, which we do by mounting a power meter or a Raspberry Pi camera right on the alignment piece. The power meter is best for systematically maximizing the transmission, while the camera is useful to check the beam profile for artifacts that could hint at clipping, and to visually check the adjustment range that is free of clipping. To obtain an optimum alignment, we iterate between the individual picomotors and the common picomotor for several instances.

Using this strategy, we obtain beams that show no signs of clipping on a camera after the chamber and measure a total transmission of around 85% for both beams, as deduced by the fraction of power measured after the glass cell assembly compared to the power mea-

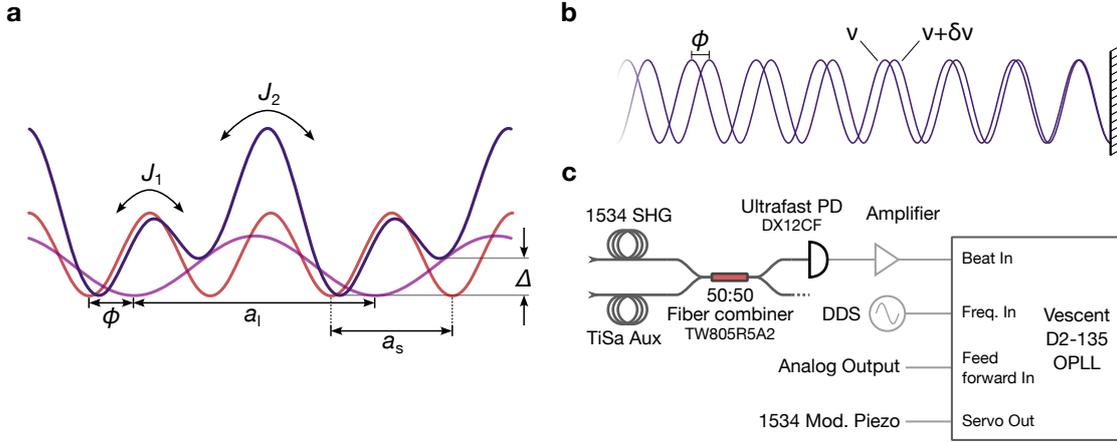


Figure 2.6 | Superlattice potential and frequency stabilization. **a**, Illustration of the superlattice potential created from a short-spacing lattice with spacing a_s , and a long-spacing lattice with spacing $a_l = 2a_s$, forming a double-well array. Via the depths and the phase offset Φ , the tunnel couplings J_1 and J_2 as well as the tilt Δ can be programmed. **b**, Principle of the phase tuning: By introducing a small frequency difference $\delta\nu$ between the two components, a small phase difference accumulates with every period between the retro-reflector and the atoms. Since there has to be an intensity node at the retro-reflector, the phase at its location is fixed, making the superlattice intrinsically phase-stable. **c**, Stabilization of the superlattice frequency difference: The short lattice light (TiSa) is interfered with frequency-doubled light from the long lattice laser, and the beat note is detected on a fast photodiode. The beat note frequency is stabilized to the target value via an optical phase-locked loop (OPLL), which controls the frequency modulation piezo of the long lattice laser.

sured right before the pillar. The remaining losses likely originate from the various optical surfaces in the STVL mount, as well as reflections on the glass cell surfaces due to the steep angle of incidence. Hence, this should not pose any problems to an operation even at high powers on the order of 10 W.

2.3.2 Horizontal superlattices

While the vertical lattices create a single two-dimensional system, the horizontal lattices play a key role in shaping the potential landscape that realizes the Hubbard-type Hamiltonians in our experiment. As shown in Fig. 2.3, we have two perpendicular superlattice axes (L1 and L2), each consisting of a retro-reflected optical lattice with wavelength $\lambda = 767$ nm (*short-spacing lattice*, spacing $a_s = 383.5$ nm) and a second retro-reflected lattice with $\lambda = 1534$ nm (*long-spacing lattice*, spacing $a_l = 767$ nm). The two axes are oriented at an angle of approximately 90° with respect to each other, resulting in a square lattice geometry. The potential of such a bi-chromatic superlattice in each direction can be written as

$$V(x) = V_s \cos^2(\pi x/a_s) + V_l \cos^2(\pi x/a_l + \phi/2), \quad (2.2)$$

with $V_{s(l)}$ being the potential depth of the short (long)-period lattice, and ϕ is a phase offset between the two components. Since the lattice spacings differ approximately by a factor of two, the resulting potential is a periodic array of double wells as shown in Fig. 2.6a. The shape

of the double wells can be programmed freely via the superlattice parameters. In particular, the intra-well tunnel coupling J_1 as well as the inter-well coupling J_2 can be adjusted by the depths of the short and long lattice $V_{s(l)}$, and the tilt Δ can be controlled by the phase offset ϕ and the potential depth. We distinguish between specific double-well configurations with respect to the phase: In the *symmetric* configuration ($\phi = 0$), the two wells have equal energy ($\Delta = 0$) and tunneling is resonant inside the well. In the *staggered* configuration, the two wells are energetically offset ($\Delta > 0$). The couplings are equal ($J_1 = J_2$), but direct tunneling is suppressed due to the tilt. For intermediate phases ($0 < \phi < \pi/2$), one obtains tilted double wells with unequal couplings ($J_1 \neq J_2$).

Experimentally, the short lattice light is generated from two titanium-sapphire (TiSa, Sirah Matisse, up to 8 W output power) lasers operating at 767 nm, where we use one laser for each axis to have sufficient power as the short lattices are also used to provide the pinning potential during fluorescence imaging. The long lattice light is provided by two infrared fiber amplifiers at 1534 nm (NKT Photonics Koheras Adjustik seed + Boostik HP amplifier for L1, around 6 W; Rio Lasers Orion seed + Nufern Nuamp amplifier for L2, around 3.5 W). For each lattice axis, the individual colors are delivered to the experiment table via optical fibers and, after the respective beam forming, combined right in front of the chamber via a dichroic mirror and focused onto the atoms using a common lens. In the atomic plane, the short lattice has a waist of $(w_v, w_h) = (45, 160) \mu\text{m}$, and the long lattice has a circular waist of $w = 140 \mu\text{m}$. Using these parameters, we can reach lattice depths up to around 220 μK during pinning for the short lattice (using around 2.5 W power after the fiber), and roughly $60 E_{r,l}$ in the long lattices for physics, where $E_{r,l} = \hbar^2/(8ma_l^2)$ is the recoil energy of the long lattice light. For further details on the optical setups, refer to Ref. [213].

Superlattice frequency stabilization

A precise control over the double well configuration requires an active stabilization of the superlattice phase. To this end, we make use of the fact that both colors are retro-reflected on the same mirror, which fixes the absolute phase of the interference pattern. As shown in Fig. 2.6b, the introduction of a small frequency difference between the two components causes a phase difference to accumulate with every period between the retro-reflector and the atoms. In this way, the superlattice phase at the position of the atoms can be controlled by adjusting the frequency difference between the two colors. The retro-reflection additionally makes the phase intrinsically robust, as it is only sensitive to the path between the retro-reflector and the atoms. In detail, the superlattice phase ϕ is given by

$$\phi = 2\pi \frac{L}{c} \delta\nu, \quad (2.3)$$

where L is the distance between the retro-reflector and the atoms, and $\delta\nu = \nu_{\text{long}} - \nu_{\text{short}}/2$ is the frequency difference between the two colors. In our case, we have $L_{L1} = 300 \text{ mm}$ and $L_{L2} = 200 \text{ mm}$, such that a frequency difference of $\delta\nu = 750 \text{ MHz}$ for L1 and $\delta\nu = 500 \text{ MHz}$ for L2, respectively, results in a phase shift of $\phi = \pi$ (i.e., two successive symmetric configurations).

A schematic of our setup to stabilize and control the frequency difference is shown in

Fig. 2.6c. In detail, we interfere around 1 mW of the light from the short lattice laser with an equal amount of frequency-doubled light from the long lattice laser using a fiber combiner. The resulting beat note, which oscillates at the difference frequency, is detected on a fast, fiber-coupled photodiode and fed into an optical phase-locked loop (OPLL, Vescent D2-135). The OPLL compares it to a reference frequency (which is internally multiplied by a factor 16 or 32) and controls the frequency modulation piezo of the long lattice laser to stabilize the beat note to this reference. The reference frequency is generated via a DDS module, which we can control using the timing system.

One limitation that we found with this scheme is that the bandwidth of the feedback loop is relatively small, taking around 500 ms to realize a full π phase shift. To circumvent this and allow fast jumps between different double well configurations (as used e.g. in Chapter 5.3), we make use of the feed-forward input of the OPLL controller (see Fig. 2.6c). Here, an analog signal is summed onto the feedback servo output, allowing changes that are limited only by the actuator bandwidth. To calibrate the necessary feed-forward voltage, we perform a slow ramp to the final frequency difference without any feed-forward and record the resulting change in the servo output voltage, which is then applied to the feed-forward input for the jump.

2.4 High-resolution imaging and potential shaping

We image the 2D atomic sample in the optical lattices with single-atom and single-site resolution using fluorescence imaging. To this end, our experiment is equipped with two identical high-resolution microscope objectives that image the atoms through the top and bottom viewports of the glass cell, respectively. The objectives are custom manufactured by Special Optics, have a numerical aperture of 0.8, and are diffraction limited at 852 nm (for imaging on the D2 line of cesium), at 780 nm (for imaging of rubidium) and at 456 nm (for blue-line imaging of cesium). Their working distance is around 10.6 mm, with an effective focal length of 25 mm.

As illustrated in Fig. 2.7, we built two similar imaging setups on the top and bottom breadboards, respectively. The first element is a mirror that brings the imaging information in the horizontal plane, followed by the tube lens, where we use an $f = 1000$ mm achromatic lens (Thorlabs ACT508-1000-B). The tube lens forms an image of the atomic plane after $1f$ distance, where we place an sCMOS camera with enhanced sensitivity at 852 nm (Teledyne Photometrics Kinetix, $6.5 \mu\text{m}$ pixel size). The choice of the tube lens leads to a magnification of around 40, which is the smallest value given typical focal lengths that still allows to faithfully sample the short lattice spacing without aliasing. In detail, one short lattice site (area 383.5×383.5 nm) is imaged onto an area of $15.5 \times 15.5 \mu\text{m}$ in the camera plane, corresponding to around 5.7 camera pixels spanning the area of one lattice site.

To perform the fluorescence imaging, we first freeze the atomic distribution by suddenly quenching the short horizontal lattices as well as the shallow-angle vertical lattice to full depth (around $220 \mu\text{K}$ in the horizontal directions, around $500 \mu\text{K}$ in the vertical direction). Then, we illuminate the atoms with light at 852 nm (red-detuned from the D2 line by about 70 MHz) to perform optical molasses cooling, which keeps each atom localized in its respective lattice site. As a by-product, the atoms scatter photons, which we collect as fluorescence light through the

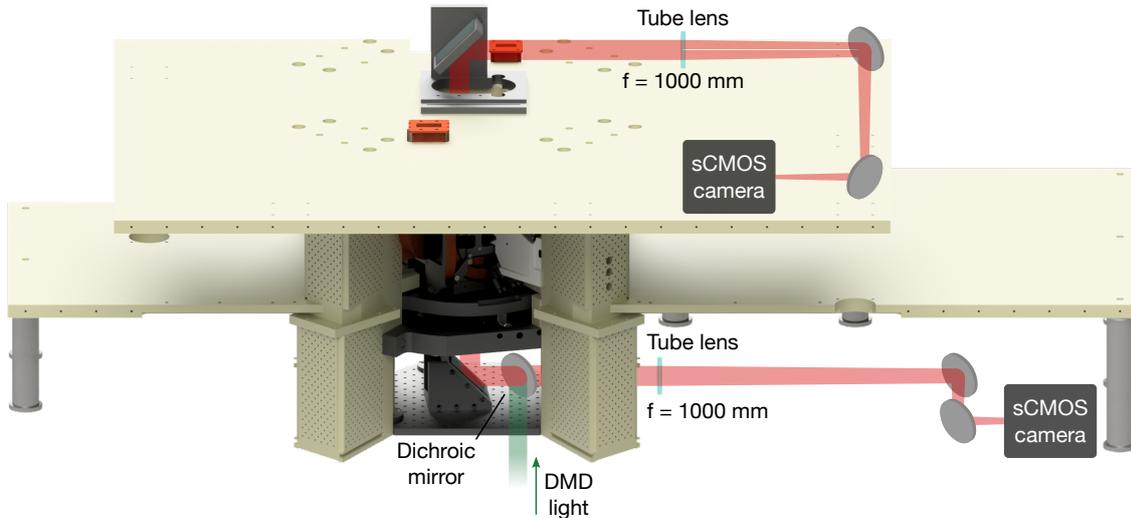


Figure 2.7 | High-resolution imaging setup. The setup features two high-resolution objectives from above and below the glass cell. The corresponding identical imaging setups, consisting of an $f = 1000$ mm tube lens and an sCMOS camera $1f$ after, are built on the bottom and the top breadboard, respectively. In the bottom setup, there is an additional dichroic mirror in the Fourier plane, through which we send the DMD light in to project programmable repulsive potentials into the atomic plane. Note that the front breadboard of the main level is not shown for clarity.

high-resolution objectives and image onto the cameras. The continuous cooling is necessary to allow for a sufficient number of scattered photons during the exposure time of 300 ms without heating the atoms out of their lattice sites. Experimentally, we have two retro-reflected, circularly polarized laser beams to drive the optical molasses cooling, one entering horizontally and the other at an angle of 30° out of the horizontal plane (see Fig. 2.3). Since the lattice structure of the molasses beams leads to spatial variations of the cooling and scattering rate that are not commensurate with the pinning lattice, this can cause large-spacing fringes in the fluorescence intensity which decreases the signal-to-noise ratio by broadening the intensity distribution. To minimize this effect, we modulate the interference pattern through piezos on the retro-reflectors and a modulation of the frequency detuning. For further information on the molasses setup and the optimization, refer to Ref. [213].

Initially, we had just the bottom objective and imaging setup installed, which was used for the experiments presented in Chapters 4 and 5. The imaging performance was however found to be far from diffraction limited with significant aberrations (see Section 3.1). Hence, we decided in early 2024 to additionally install the top objective using a newly developed alignment strategy, which was also shown to reach close to diffraction-limited performance in an optical test setup [225]. This enabled us to increase the resolution slightly and get rid of most aberrations. Since then, fluorescence imaging is performed exclusively using the top objective, in particularly also for the experiments presented in Chapter 6.

In addition to the fluorescence imaging, we also use the high-resolution objectives to project programmable optical potentials into the atomic plane using a digital micromirror device (DMD, Texas Instruments DLP6500, interface by bbs Bild- und Lichtsysteme GmbH,

resolution 1920×1080 pixels). The DMD is illuminated with incoherent light at $\lambda = 525$ nm from a multimode diode laser (Wavespectrum WSLX-525-005-400M-H, 5 W maximum output power, ~ 10 nm spectral width). Due to the wavelength of 525 nm, the resulting potentials are repulsive. The imaging optics for the DMD are built on the bottom breadboard separate from the fluorescence imaging optics, and implement a total demagnification of 160 including the objective. Such a large magnification is chosen to allow a gray-scaling, which enables continuous intensity variations in the projected patterns. As shown in Fig. 2.7, the DMD light is overlapped with the imaging light using a dichroic mirror in proximity to the periscope mirror that sends the imaging information upwards into the bottom objective.

The DMD has proven itself as a versatile tool and plays a central role in our experiment: We use it in virtually every sequence to provide a repulsive box confinement that enables the preparation of large, high quality Mott insulators. Furthermore, we have used the DMD to create local potentials that energetically detune certain lattice sites, which forms the basis for the local manipulations presented in Section 5.4. Lastly, the DMD can also be used to compensate the harmonic curvature coming from the horizontal and vertical lattice beams, which allowed to study the spread of density correlations across large distances [226]. Further information on the DMD setup, its calibration and the potential flattening algorithm can be found in Refs. [213, 227].

2.5 Laser-assisted tunneling setup

As presented in Section 1.3, we engineer artificial gauge fields using a laser-assisted tunneling scheme that is implemented via a running-wave lattice. Since the resulting magnetic flux value depends on the geometry of the running-wave lattice and the underlying physics lattice, the target flux of $\pi/2$ imposes restrictions on the lattice spacings and the angle between the beams. Based on the square physics lattice with a spacing of $767/2$ nm, we can realize a flux of $\pi/2$ using a 1D lattice with a spacing of $1534/\sqrt{2}$ nm = 1085 nm that is rotated by a 45° angle with respect to the physics lattice vectors. This spacing can be realized by interfering two 1534 nm beams under an angle of 90° . To orient the resulting 1D lattice diagonally with respect to the physics lattice, we have to send the beams in along the same axes as the physics lattice beams (see Fig. 2.3).

Experimentally, we realize this by roughly overlapping the beams for the running wave with the physics lattice beams under a small angle using D-shaped mirrors, and then focusing them through a common lens onto the atoms. Since the space around the glass cell is already very crowded, we assemble the beam forming optics for the running-wave lattice on small, separate breadboards that are mounted vertically on the pillars. This is shown as an example for lattice axis 1 in Fig. 2.8 (the breadboard for axis two has identical components, but a slightly adapted geometry). The breadboard features photodiodes for intensity stabilization, a lens to set a suitable beam waist, and a picomotor mirror for remote alignment. Having the beam forming optics on separate breadboards also has the advantage that the components can be placed and pre-aligned before the breadboard is put into the experiment, which greatly simplifies the installation process. Together with the $f = 150$ mm lens that is shared with the superlattice

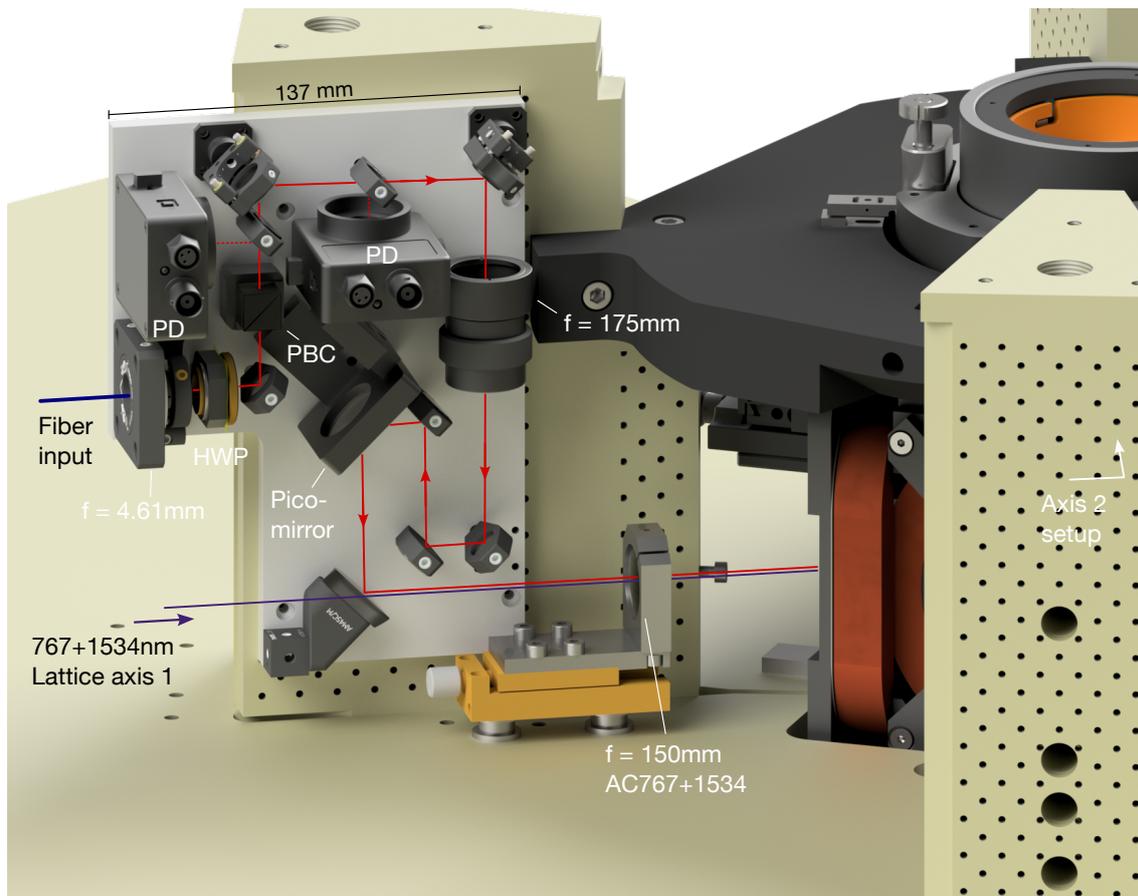


Figure 2.8 | Vertical breadboard for laser-assisted tunneling. The beams for the running-wave lattice used for laser-assisted tunneling are conditioned on small, separate breadboards that are mounted vertically on the pillars. Via D-shaped mirrors, the beams are overlapped under a small angle with the main superlattice beams, and focused through a common lens onto the atoms. Here shown is the breadboard for axis 1, featuring photodiodes for intensity stabilization, a lens and a picomotor mirror for remote alignment. The breadboard for the second axis is similar. Abbreviations: HWP - half-wave plate, PBC - polarizing beamsplitter cube, PD - photodiode.

beam, the running-wave beams are focused to a circular spot size of $w = 380 \mu\text{m}$, which is on purpose chosen to be relatively large to ensure a homogeneous modulation strength across the lattice as well as low sensitivity to drifts. The light for the running-wave lattice is generated from a single 1534 nm fiber amplifier (NKT Photonics Koheras Adjustik seed + Boostik HP amplifier), which is split into two paths on a separate breadboard, before going through individual AOMs and into the optical fibers that lead to the experiment table. The AOMs are driven from two channels of a DDS board, which allows us to program the kHz-level frequency difference between the two beams that is required for the laser-assisted tunneling scheme.

For initial alignment, we adjust each running-wave beam by eye to have the same distance of roughly 2 mm to the physics lattice beam in front of and after the chamber. Finer alignment is then done using an infinity-focused camera placed right after the superlattice retro mirror. Together with the lens on the retro side of the glass cell (see Fig. 2.3), the camera images the

atomic plane, which allows us to overlap the running-wave with the physics lattice beams at the location of the atoms. After this step, we adiabatically turned on the lattice with zero frequency detuning after preparing a cold superfluid sample and immediately observed how the atoms arrange in the attractive, diagonal lattice potential. We further attempted to optimize the lattice depth by briefly pulsing the lattice on and performing Bragg spectroscopy, but found that the lattice depth was already close to the optimum value, which is probably also a result of the large beam waist.

CHAPTER 3

Calibration and preparation techniques

In this chapter, we will describe the central calibration and preparation techniques used in the experiments throughout this thesis. We will begin with the calibration of the imaging parameters such as the lattice vectors, the size and shape of the point-spread function, as well as the hopping and loss rate. These parameters are of central importance for the reconstruction discussed in the following chapter. We will then describe the sequence for transferring the Bose-Einstein condensate into Mott-insulating states in various lattice geometries, which provide the initial states for most of our experiments. Afterward, calibration techniques for the microwave parameters, the vertical lattice alignment and the superlattice parameters are presented. Finally, we will discuss the calibration of the Hubbard parameters as well as the laser-assisted tunneling, and present methods to estimate the strength of on-site potential disorder.

3.1 Imaging and molasses

Individual atoms have a negligible extent compared to typical optical resolutions, allowing us to treat them as point sources. As such, they appear in a fluorescence image as the point-spread function (PSF) of the imaging system, rendering them ideal to characterize the imaging system and its parameters (for further details, see also Section 4.1). All important imaging properties can hence be extracted from experimental images of dilute samples, where the average atomic spacing is much larger than the width of the PSF. Experimentally, we create a dilute sample starting from a Bose-Einstein condensate in a single plane of the shallow-angle vertical lattice, which has been evaporated to very low atom numbers. We then add the horizontal short-spacing lattices and modulate their intensity for a short duration on resonance with the first band excitation, which heats the sample and causes the individual atoms to disperse throughout the field of view. When the sample is sufficiently dilute, we take a fluorescence image, exposing the camera for a duration of 300 ms. An example fluorescence image of a dilute sample is shown in Fig. 3.1a.

We begin with the bottom imaging system, and calibrate the lattice geometry as well as the magnification, which determines the pixel size in the atomic plane. To this end, we fit the positions of the individual atoms in many dilute images (typically at least 40), and compute the Fourier transform of the positions. A direct Fourier transform of the images is not possible

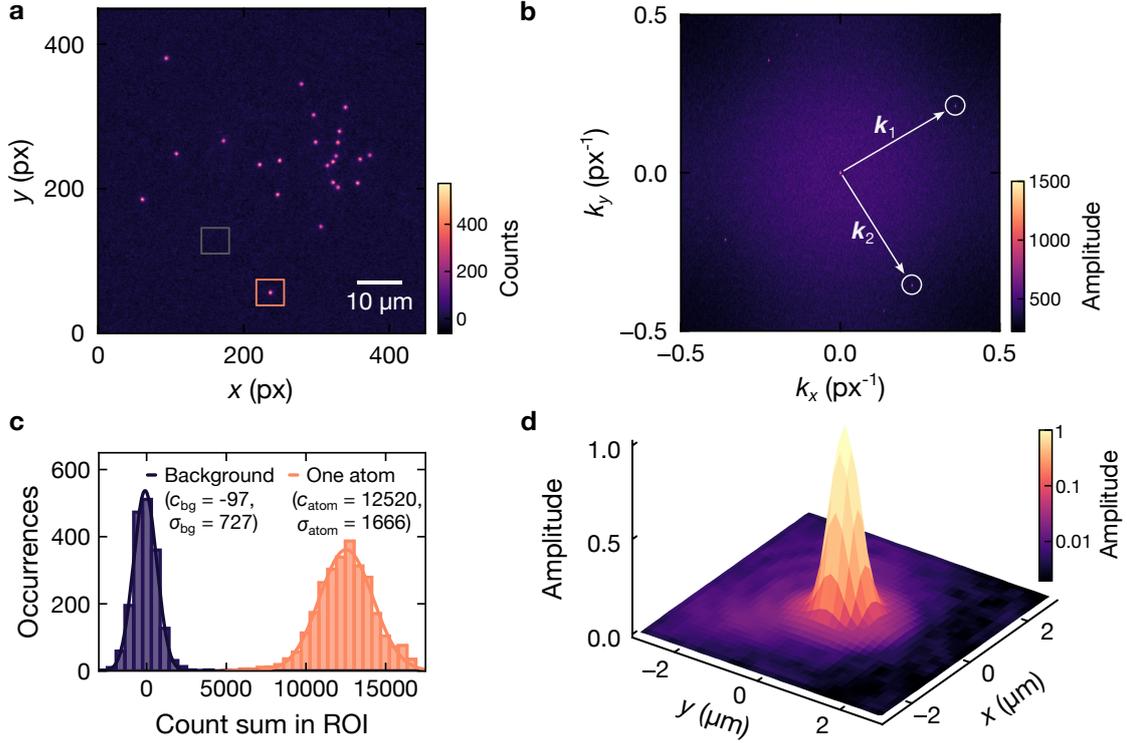


Figure 3.1 | Calibration of the imaging parameters. **a**, Fluorescence image of a dilute sample of atoms in the lattice, from which we extract further imaging parameters. The orange (gray) square indicates the small subsections used to evaluate the atom (background) signal distributions and the signal-to-noise ratio. **b**, Fourier transform of the fitted locations of many individual atoms, showing distinct peaks corresponding to the lattice vectors $\mathbf{k}_{1,2}$. From this, we calibrate the magnification and the lattice geometry. **c**, Signal and background count distributions, evaluated from crops ($1.6 \mu\text{m}$ width) containing exactly one (atom) or zero (background) atoms as indicated in (a). The solid lines denote Gaussian fits, and the fitted peak positions and widths are used to determine the signal-to-noise ratio. A background image without atoms was subtracted to shift the background peak center closer to zero. **d**, Averaged point-spread-function from many individual atoms. The logarithmic color scale reveals residual aberrations in form of a plateau around the central peak and long-range coma tails in the negative y -direction.

due to the large PSF width compared to the lattice spacing. In the resulting Fourier transform (see Fig. 3.1b), we find distinct peaks corresponding to the lattice vectors \mathbf{k}_1 and \mathbf{k}_2 . The lattice vectors determine the positions of the lattice sites (up to a phase), which is a key ingredient to later reconstruct the lattice occupation (see Chapter. 4). Furthermore, their length allows us to calibrate the magnification and the pixel size in the atomic plane. Using the camera pixel size of $6.5 \mu\text{m}$ and the known short lattice spacing of $a_s = 383.5 \text{ nm}$, we obtain a magnification of $M = 40.53$, corresponding to a pixel size of 160.4 nm in the atomic plane. The proximity of the lattice wave vectors to the edges of the Fourier transform (given by the Nyquist frequency [228]) also shows that a much smaller magnification would not be possible, as the lattice vectors would no longer be resolved.

The next parameter we can calibrate from such dilute images is the signal-to-noise ratio (SNR). To this end, we sample crops containing exactly one atom, and an equal number of

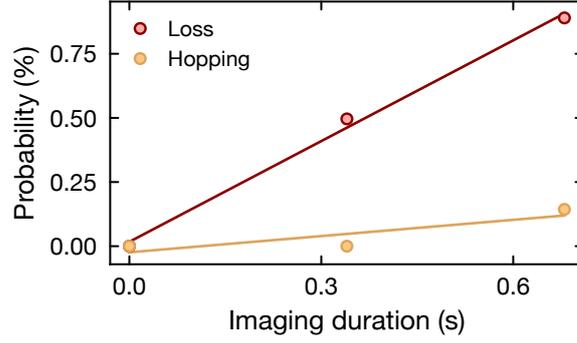


Figure 3.2 | Calibration of the loss and hopping rate. Probability of hopping and loss events as a function of the exposure duration, where each data point corresponds to one exposure. The probability is computed as the change fraction in the reconstructed occupation. The solid lines are linear fits to extract the hopping and loss rate.

crops from the background, and analyze the distribution of the count sums within the crops. As shown in Fig. 3.1c, this reveals a well-separated bimodal distribution, indicative of a high SNR. The width of the background peak is given by camera readout noise and scattered light. Furthermore, the position of the atom peak is determined by the number of scattered photons. Its width is fundamentally limited by shot noise, but is usually enhanced due to fluorescence inhomogeneities as well as hopping and loss processes during the imaging. To capture all of these contributions, we define the SNR as

$$\text{SNR} = \frac{c_{\text{atom}} - c_{\text{bg}}}{\sigma_{\text{atom}} - \sigma_{\text{bg}}}, \quad (3.1)$$

where c_{atom} and c_{bg} are the average counts of the atom signal and the background signal, respectively, and σ_{atom} and σ_{bg} are the respective standard deviations. We extract all four parameters through Gaussian fits of the background and signal peaks, and find an SNR of 5.2 for an exposure time of 300 ms. Based on this quantity, we can optimize the imaging parameters.

Lastly, since each atom appears as an individual point-spread function, we can immediately estimate the resolution and assess the presence of aberrations. To this end, we average the signals of 2000 isolated single atom signals, centering them with sub-pixel accuracy by ten-fold oversampling the image to obtain a high resolution reconstruction of the PSF. The result is shown in Fig. 3.1d, which exhibits a strong central peak that corresponds to a Rayleigh resolution of around 850 nm. The logarithmic color scale furthermore reveals a plateau around the central peak as well as long-range coma tails in the $-y$ -direction, which is indicative of residual aberrations. While their relative magnitude is rather small, they are still noticeable, especially for dense samples where the signals of many atoms overlap, which also has an impact on the reconstruction performance (see Section 4.3).

As discussed above, the occurrence of hopping and loss processes during the exposure reduces the SNR and complicates the reconstruction of the lattice occupation. Their influence is mostly determined by the cooling performance of the optical molasses as well as the depth of the pinning lattices. To quantify the hopping and loss rate, we again prepare dilute samples,

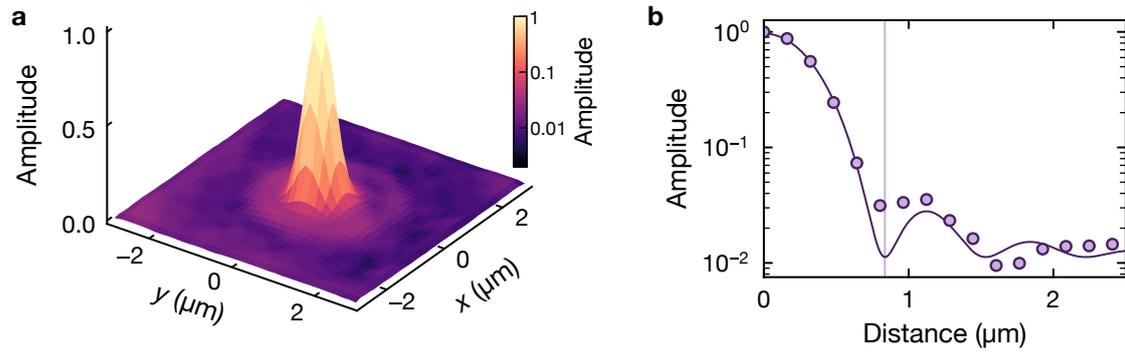


Figure 3.3 | PSF of the top imaging system. **a**, 3D representation of the averaged PSF of the top imaging system, with the key improvement being the absence of the long-range coma aberrations. The color scale is logarithmic. **b**, Central cut through the PSF along the y -direction. The solid line is a fit of an Airy function, which indicates a Rayleigh resolution of around 820 nm (vertical line).

but take several subsequent exposures of the same sample. By reconstructing and comparing the occupation between exposures, we can estimate the hopping and loss rate. A calibration of the loss and hopping rate is shown in Fig. 3.2. Here, we took three subsequent exposures with a duration of 300 ms, separated by a short delay of 40 ms. Using linear fits, we extract a hopping rate of 2.1(1.2) mHz and a loss rate of 13(1) mHz. This corresponds to a loss probability of around 0.44% and a hopping probability of 0.07% per atom and exposure. These quantities will limit the attainable reconstruction fidelity, and can be used to optimize the molasses cooling and the pinning lattice depth, as described in more detail in Ref. [213].

PSF of the top imaging system

As mentioned in Section 2.4, we additionally installed a second high-resolution imaging system from the top after we noticed that the remaining aberrations in the bottom imaging system were limiting the reconstruction performance. The averaged PSF of the top imaging system after alignment is shown in Fig. 3.3, which was determined in the same way as described above for the bottom imaging. We now find a PSF with a Rayleigh resolution of around 820 nm, which is roughly 50 nm smaller than the bottom imaging system (see cut in Fig. 3.3b). In addition, the remaining aberrations are significantly reduced, and in particular, we no longer find long-range coma aberrations.

3.2 Initial state preparation

In this section, we describe the sequence following the preparation of the Bose-Einstein condensate (BEC), which is used to create the Mott insulating states that are the starting point for most of our experiments. Schematically, this consists of loading the BEC into a single layer of the shallow-angle vertical lattice and an in-plane repulsive box potential, a second evaporation step to decrease the temperature, and finally the adiabatic ramp-up of the horizontal (super)lattices to load a Mott insulator.

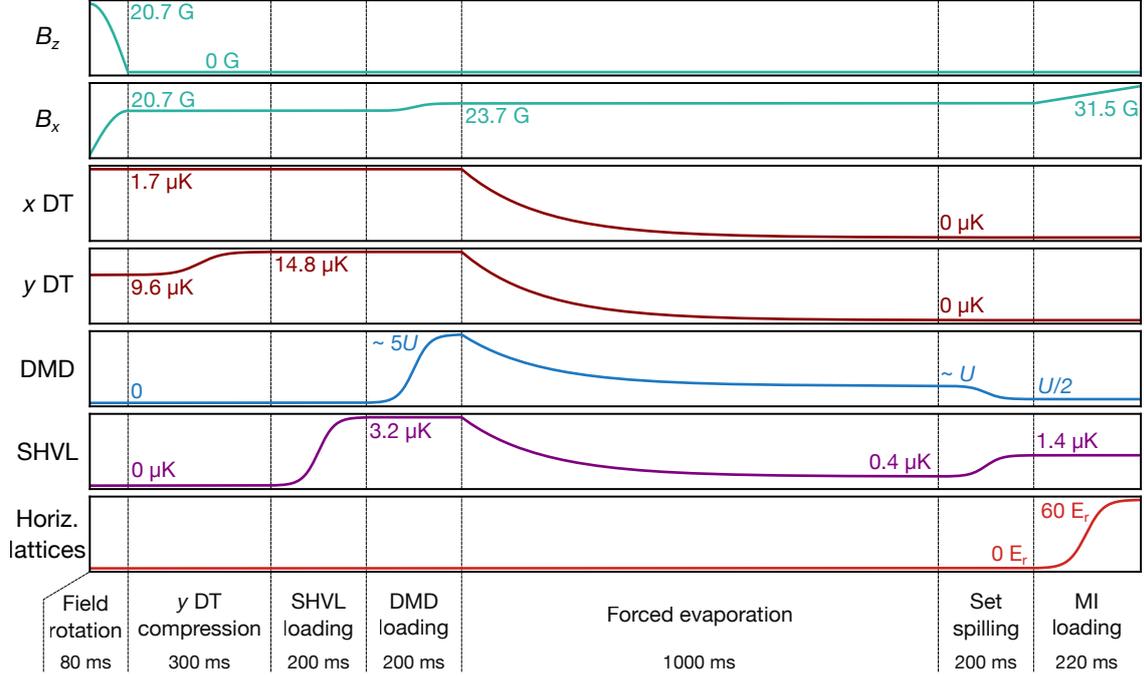


Figure 3.4 | Sequence diagram for the preparation of Mott-insulating initial states. Starting from a BEC in the crossed dipole trap, we first load the atoms into a single layer of the shallow-angle vertical lattice and the in-plane repulsive box potential. Then, a second evaporation in the SHVL is carried out to reduce the temperature. Finally, the horizontal lattices are ramped up to load a Mott insulator. The depth of the DMD box is given relative to a typical final on-site interaction energy in a short-short Mott insulator of around $U/h = 900$ Hz.

The full sequence diagram is shown in Fig. 3.4, starting from a BEC in the crossed ($x + y$) dipole traps. In a first step, we rotate the offset field from the z - into the x -direction, such that we can later use small horizontal field gradients to shift and correct the center of the harmonic in-plane confinement. Next, we load the atoms into a single layer of the shallow-angle vertical lattice ($8 \mu\text{m}$ spacing), which is done by first ramping up the depth of the y dipole trap and then adiabatically turning on the SHVL to a depth of around $3.2 \mu\text{K}$ (vertical trap frequency of around $2\pi \times 1.2$ kHz). Due to the vertically squeezed shape of the y dipole trap, the compression helps to reduce the size of the BEC in the vertical direction and hence to load most of the atoms ($\gtrsim 95\%$) into a single layer. After this, we also adiabatically turn on an in-plane repulsive box potential, which is projected using the DMD and will later provide the main horizontal confinement. Next follows a second evaporation step, which is necessary to compensate any heating that has occurred during the loading process and reduce the temperature of the sample sufficiently to create high-quality Mott-insulating states. To this end, we exponentially ramp down the depth of the vertical lattice, which causes the hottest atoms to leave the trap in the vertical direction and the remaining atoms to thermalize to a lower temperature. Simultaneously with the SHVL depth, we also remove the x and y dipole traps, such that the atoms are only confined in the box potential and the vertical lattice at the end of the evaporation. The final atom number in the box (i.e., the chemical potential), which has to be matched to the desired filling in the lattice, is set by the final depth of the vertical lattice. In a last preparation step, we increase the

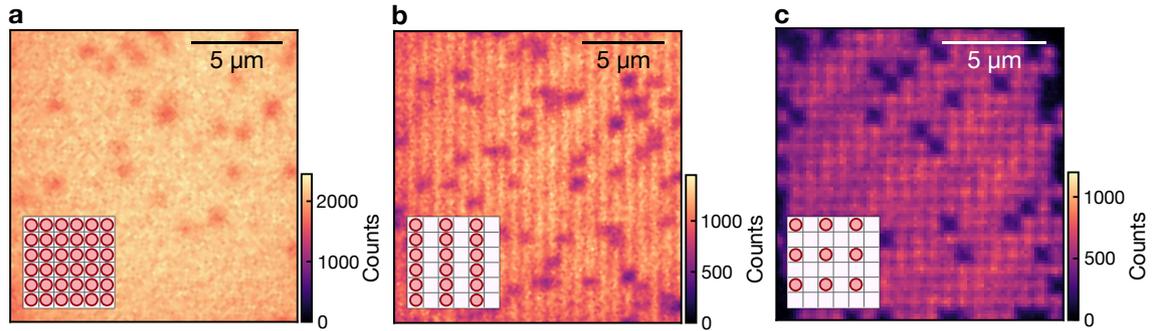


Figure 3.5 | Mott-insulating states in different geometries. **a**, Mott-insulator in a short-short lattice geometry, resulting in a unity-filling state with one atom per short lattice site. **b**, Mott-insulator in a rectangular lattice geometry, consisting of a short lattice in the vertical axis and a short+long lattice in the horizontal axis with a non-zero superlattice phase. The resulting tilt blocks every other site during loading, resulting in a charge-density wave state with a filling of $1/2$. **c**, Mott-insulator in a square lattice geometry with both superlattice axes in a tilted configuration, resulting in a filling of $1/4$.

SHVL depth back slightly corresponding to the final U that we want to obtain, and decrease the depth of the box potential such that the height of the walls correspond roughly to $U/2$ in the final Mott-insulating state. This is done to enable what we call the “spilling technique”, which makes the loading of the Mott insulator more robust and increases its quality substantially [43]. The idea behind this technique is to set the chemical potential after the second evaporation slightly higher than the number of atoms that would be required to fill the lattice in the box at the desired filling fraction. As a result, during the transition to the Mott-insulating state, surplus atoms can spill over the walls of the box, which is energetically favorable compared to creating doubly-occupied sites at an energy cost of U . The higher-than-necessary initial atom number reduces the number of holes and additionally makes the final filling less sensitive to fluctuations in the previous stages. The low-temperature superfluid is then finally transformed to a Mott insulator by adiabatically ramping up the horizontal lattices to a depth of around $60 E_r$, while we simultaneously increase the offset field to $B = 31.5$ G to increase the scattering length. At this point, the sample is in an approximate product state deep in the Mott-insulating regime with $U \gg J$, which is the starting point for most of our experiments.

Using our bi-chromatic superlattices, we can realize Mott insulators in a variety of lattice geometries, which results in different filling fractions with respect to the short-spacing lattices. The first possibility is to ramp up the two short-spacing lattices in the last step of Fig. 3.4, which results in a unity-filling state with one atom per short lattice site (see Fig. 3.5a). Here, we can reach fillings up to around 0.98 using the spilling technique. Secondly, we can replace one of the short-spacing lattices with a long-spacing lattice, resulting in a rectangular geometry where we can realize a filling of $1/2$ with respect to the short lattices. This configuration allows additionally to combine long and short lattices in one axis. For example, we can ramp up the short lattice to an intermediate depth on the order of $10 E_{r,s}$ in the symmetric configuration (superlattice phase $\phi = 0$), such that each double well is in the state $|+\rangle = \frac{1}{\sqrt{2}}(|L\rangle + |R\rangle)$, i.e., a product state of double wells with local superpositions between an atom in the left well ($|L\rangle$) and in the right well ($|R\rangle$). The best half-filling states are typically realized when we ramp both

short and long lattice in one axis to a large depth of around $60 E_{r,s}$ and $60 E_{r,l}$, respectively, in presence of a tilt. Here, the higher site of each double well is energetically blocked, resulting in a charge-density wave state as shown in Fig. 3.5b with typical fillings in the occupied sites up to around 0.95. Lastly, we can do the same with both superlattices, which results in a filling of $1/4$ with respect to the short lattices (see Fig. 3.5c, where both axes are in a tilted configuration). The superlattices hence give us access to specific products states beyond the standard unity filling with high quality and repeatability.

3.3 Vertical lattice parameters and loading

In this section, we discuss fine-alignment and calibration of the vertical lattice parameters. The latter consists most importantly of the trap frequencies as well as the correct loading into a single plane.

3.3.1 Shallow-angle vertical lattice

In the case of the shallow-angle vertical lattice (SHVL), we primarily have control over common degrees of freedom that affect both beams at the same time. This includes a mirror that moves the crossing point perpendicular to the axis of propagation, as well as the focusing lens, which shifts the crossing point along the propagation axis. We optimize the trap depth by loading the BEC into the lattice at a given depth and subsequently switching off all other traps, such that the remaining number of atoms is proportional to the trap depth of the vertical lattice. An optimal alignment is important for the second evaporation step, as this is controlled by the depth of the SHVL. After optimal alignment, we obtain trap frequencies of around $(\omega_x, \omega_y, \omega_z) = 2\pi \times (28, 8, 1250)$ Hz at a power of 100 mW, which is a typical power used to load the atoms into the lattice. Experimentally, the trap frequencies can be determined using a sloshing protocol, where we displace the atoms in the lattice via a magnetic field gradient and measure the resulting oscillation frequency in the harmonic potential.

Besides an optimal depth, it is crucial to load the atoms into a single plane of the lattice. To adjust this, we use a matter wave focusing protocol [229], where we suddenly turn off the vertical lattice and let the atoms expand in the harmonic confinement of the y dipole trap for a quarter oscillation period. After 15 ms time of flight, the population of the individual planes can be resolved in absorption imaging. We then optimize the vertical position of the y dipole trap to maximize the population in the central plane. Depending on the environmental conditions, this typically remains stable for at least a full day.

3.3.2 Steep-angle vertical lattice

For the steep-angle vertical lattice (STVL), we use the two individual picomotors (cf. Fig. 2.4) to optimize the alignment of the lattice. Here, we perform the alignment using fluorescence imaging, where we image a horizontal slice of the system. In detail, we load the atoms into the shallow-angle vertical lattice and evaporate as described in Fig. 3.4. After the evaporation, we adiabatically turn off the box potential while ramping up the power in one of the STVL

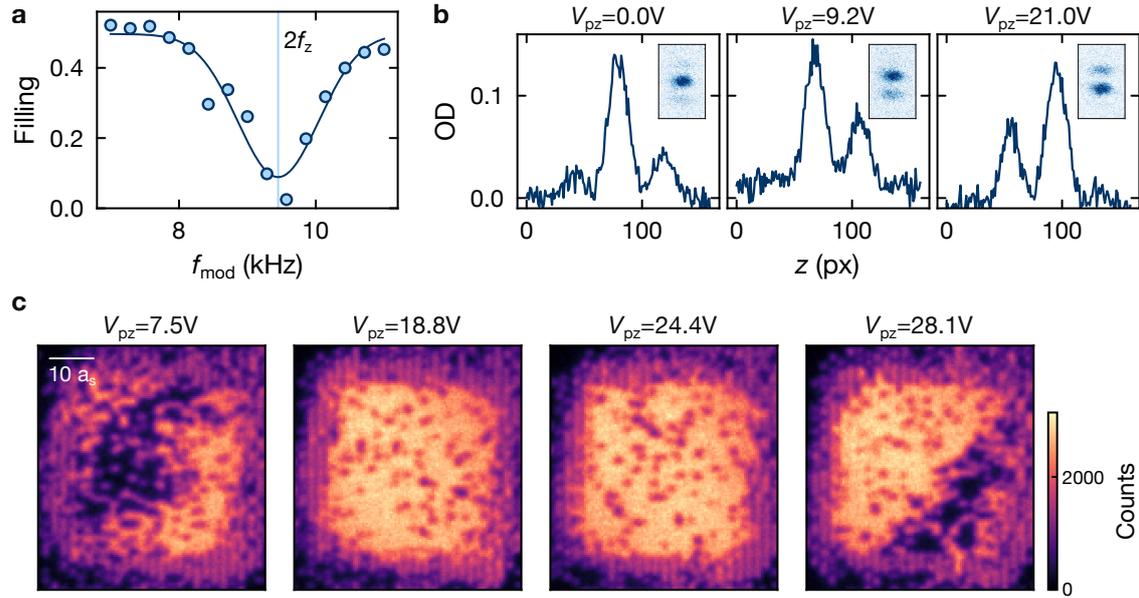


Figure 3.6 | Calibration of the steep-angle vertical lattice. **a**, Parametric heating in the steep-angle vertical lattice, showing a heating resonance at $2 \cdot (\omega_z/2\pi)$ (vertical line), corresponding to $\omega_z = 2\pi \times 4.7$ kHz. **b**, Matter wave focusing protocol to calibrate the loading for three different piezo voltages V_{pz} . A shift of the interference pattern by one plane corresponds to a piezo voltage of roughly 30 V. **c**, Doublon splitting protocol for different piezo voltages, starting from an $n = 2$ CDW state. When more than one plane is loaded, regions appear in the sample where doublons are not correctly split into two adjacent singlons. Note that the edge region shows an $n = 1$ CDW, originating from a reduced chemical potential during loading as a result of the finite wall sharpness of the box potential.

beams (via the motorized wave plate, all power is sent into one beam). The position of the atoms will then coincide with the location where the STVL beam cuts the horizontal plane defined by the SHVL. We then center the STVL beam onto the box potential via the individual picomotor. The same is repeated for the second beam, and afterward, we cross-check that the trap center is unchanged with both beams balanced. This concludes the fine-alignment, after which we obtain trap frequencies of $\omega_r = 2\pi \times 20$ Hz (radially isotropic) and $\omega_z = 2\pi \times 4.7$ kHz (vertically) at a power of roughly 20 mW. The radial trap frequency is measured as before using a sloshing protocol, while the vertical trap frequency is determined through parametric heating (see Fig. 3.6a for an example measurement).

Due to the significantly smaller lattice spacing of $1 \mu\text{m}$, we cannot load the BEC directly into a single plane of the STVL. Instead, we have to transfer the atoms from the SHVL, where the vertical extent of the cloud is reduced as a result of the stronger vertical confinement, and additionally load the STVL only after the second evaporation. To load a single plane correctly, a fringe of the STVL interference pattern has to be aligned with the central SHVL plane. We experimentally tune this by applying a voltage to one of the piezos in the STVL mount (see Fig. 2.4), which introduces a path length difference and shifts the interference pattern. The plane occupation can in principle be checked as before using the matter wave focusing protocol, when using the SHVL to provide the harmonic confinement for the focusing. This is shown as an example in Fig. 3.6b for a few different piezo voltages. The scan indicates that we can

smoothly adjust the relative phase of STVL and SHVL patterns, and that a shift by one STVL plane corresponds to a piezo voltage of roughly 30 V. The matter wave focusing protocol is not ideal however, as the limited SNR in absorption imaging allows us to probe the plane occupation only before the second evaporation, which is also the reason why not all atoms are loaded into the central plane in Fig. 3.6b (left panel). Since the sequence conditions at this point are not identical to the fluorescence sequence, we found that a calibration directly in fluorescence imaging is necessary. To this end, we use a *doublon splitting protocol*: We start from a CDW state prepared in the shallow vertical lattice (as described in Fig. 3.4 and Fig. 3.5b), but with two atoms in the lower site of each double well. After freezing the state, we transfer it to the steep vertical lattice by first ramping the STVL power up over 100 ms and then turning the SHVL off over 200 ms. Afterward, the doubly-occupied sites are split into two adjacent singly-occupied sites by first ramping down the short lattice of the superlattice axis (60 ms), which merges the double well into a long lattice site. Subsequently, the superlattice phase is ramped close to the symmetric point (100 ms, about $\pi/500$ phase offset), before we ramp the short lattice back up (60 ms to $40 E_{r,s}$), which splits the long lattice site into a double well. Due to the repulsive on-site interaction, also the doublon on the lower energy well is split into a double well with two singly-occupied sites. If however the planes are not aligned to one another, the repulsive interaction causes the doublons to split into two adjacent planes of the STVL during loading. In this case, the doublon splitting in the double wells of the horizontal superlattice does not work, as there is only a single atom in the double well of each STVL plane, which will end up in the same well due to the small applied tilt. During the fluorescence imaging, the depth of the STVL stays negligible, such that a doublon that was split into two adjacent STVL planes undergoes a light-assisted collision and is lost, which can be detected as an empty double well in the fluorescence image.

The result of such a doublon splitting protocol is shown in Fig. 3.6c for different piezo voltages. In the middle two images, the vertical lattice planes are correctly aligned, which causes all doublons of the $n = 2$ CDW state to split and the resulting cloud to show a homogeneous, unity filled plateau in the center. In contrast, in the left and right images, the planes are misaligned, which results in regions where doublons are not correctly split, and the lattice appears empty. This does not happen homogeneously across the sample as there is likely a small relative tilt between the planes of the steep and the shallow vertical lattices. The correct piezo voltage is hence in the middle of the range resulting in a homogeneous plateau, which is here around 6 V, or $\sim 0.2\pi$, wide. We typically re-check the loading once or twice per day, depending on the environmental conditions.

3.4 Calibration of the horizontal lattice parameters

In this section, we discuss the calibration of the horizontal lattice parameters, which entails the tunnel couplings in monochromatic and superlattice configurations, as well as the superlattice phase and tilt.

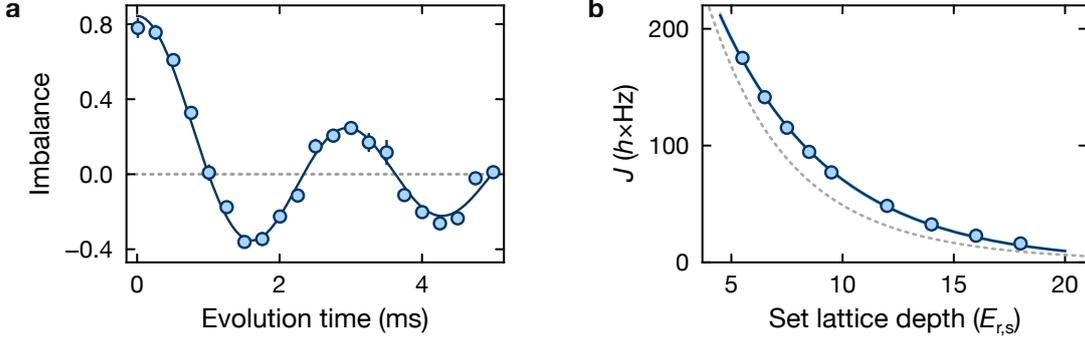


Figure 3.7 | Calibration of the tunnel coupling via CDW decay. **a**, Example CDW decay measurement for a set short lattice depth along axis 1 of $8.5 E_{r,s}$. The solid line is a fit of an exponentially damped Bessel function as defined in Eq. 3.3, yielding a tunnel coupling of $J/h = 94.6(8)$ Hz. The error bars denote the standard deviation across two repetitions and, if invisible, are smaller than the marker size. **b**, Calibrated tunnel coupling as a function of the programmed short lattice depth (using values from the parametric heating calibration). The markers denote the fit results from individual CDW decays as in (a), while the solid line is an exponential fit to the data. The gray dashed line shows the tunnel coupling as predicted from the band structure using the parametric heating depth calibration. The error bars denote the fit errors of the tunnel coupling, and are invisible due to the small size.

3.4.1 Tunnel coupling in a monochromatic lattice via CDW decay

To obtain a first rough calibration of the lattice depth in units of the recoil energy, which also determines the tunnel coupling via the band structure [121], we perform parametric heating as described in Ref. [226]. This is however not precise enough and underestimates the tunnel coupling by around 10 – 20%. To obtain a more accurate value, we measure the tunnel coupling directly using a charge-density wave decay protocol. Here, we start from an $n = 1$ CDW state prepared with the superlattice in the axis where we want to measure the tunnel coupling. With respect to the short lattice, this corresponds to a highly excited state, which will undergo characteristic relaxation dynamics. To probe these, we first remove the long lattice while both short lattices are kept deep to freeze the state. Then, we suddenly quench the short lattice depth in the axis of interest to a lower value (150 μ s ramp duration), which initiates the decay of the CDW state. The decay can be quantified by the density imbalance \mathcal{I} between even and odd sites,

$$\mathcal{I} = \frac{\langle \hat{n}_{\text{even}} \rangle - \langle \hat{n}_{\text{odd}} \rangle}{\langle \hat{n}_{\text{even}} \rangle + \langle \hat{n}_{\text{odd}} \rangle}, \quad (3.2)$$

where the averages run over all double wells in the sample, and it follows an exponentially damped Bessel function, described by the fit function [226]

$$\mathcal{I}(t) = A \mathcal{J}_0(4Jt/\hbar) \exp(-t/\tau), \quad (3.3)$$

where A is an amplitude factor, \mathcal{J}_0 is the zeroth-order Bessel function of the first kind, J is the tunnel coupling, and τ is the decay time. Fitting the experimentally measured imbalance decay hence enables a direct calibration of the tunnel coupling. An example of such a measurement is shown in Fig. 3.7a, where we find excellent agreement with the behavior predicted by Eq. 3.3.

We repeat this measurement for several different short lattice depths, which allows us

to calibrate the tunnel coupling as a function of the lattice depth. This calibration is shown in Fig. 3.7b, where we find an exponential dependence of the tunnel coupling on the programmed lattice depth (according to a previously performed parametric heating calibration) in the probed range. The exponential fit (solid line) allows to predict the tunnel coupling for arbitrary intermediate lattice depths. Additionally, we see that the tunnel couplings as predicted from the band structure using the parametric heating depth calibration are systematically lower (gray dashed line). By comparing the two exponential curves, we can finally extract a correction law that relates the programmed lattice depth to the actual lattice depth, where we obtain $V_s^{(\text{true})} = 0.82V_s^{(\text{set})} + 0.36E_{r,s}$ for lattice axis one. This fully calibrates the band structure of the short lattice.

3.4.2 Superlattice phase

A central quantity of the superlattice is the phase, which determines the configuration of the resulting double wells (see Section 2.3.2). We can experimentally calibrate the phase values corresponding to a symmetric configuration (i.e., vanishing tilt $\Delta = 0$) via a straightforward protocol: We prepare a Mott-insulator in a rectangular geometry consisting of a long lattice along the axis that is to be calibrated, and a short lattice in the perpendicular direction. Subsequently, we slowly ramp up the short lattice along the first axis over a duration of 30 ms, which adiabatically splits each long lattice site into a double well. Depending on the superlattice phase, the atom in each double well ends up either always in the left or in the right well, or – only in the case of a symmetric configuration – in a superposition of both. Due to the slow ramp, this is very sensitive to small tilts, which enables a precise calibration of the symmetric phase. Experimentally, we fit the resulting average imbalance as a function of the superlattice phase setpoint (which is programmed via an analog voltage that modulates the reference frequency for the frequency offset lock, see Section 2.3.2) with a sigmoid function. Such a fit allows us to determine the symmetric configuration with a phase uncertainty of around 1 mrad. We repeat this calibration for an adjacent symmetric point (i.e., a π phase offset), to additionally calibrate the slope of the phase dependence, which we assume to be linear. This allows us to predict the resulting double well tilt energy for arbitrary phase values between the two calibration points with sufficient precision for day-to-day operation. For a more accurate calibration, the tilt can also be measured spectroscopically using the running-wave lattice as described in Section 3.6. The superlattice phase is very sensitive to environmental conditions, in particular the ambient pressure that changes the refractive index of the air, which requires us to re-calibrate the phase every 20 – 30 min via an automatic calibration script.

For the experiments presented throughout this thesis, it is essential that the superlattice configuration is identical across the whole sample (i.e., a constant superlattice phase ϕ). However, even a small relative angle between the short and long lattice beams of a superlattice axis can result in a considerable phase gradient across the sample [230]. To zero this angle, we evaluate the imbalance in the phase calibration protocol as a function of the position for a large sample (up to 80×80 sites), with the phase set close to the symmetric point. If there is a phase gradient present, the resulting imbalance will vary across the cloud, which we can homogenize using a glass plate in the long lattice path to slightly change the beam angle. A

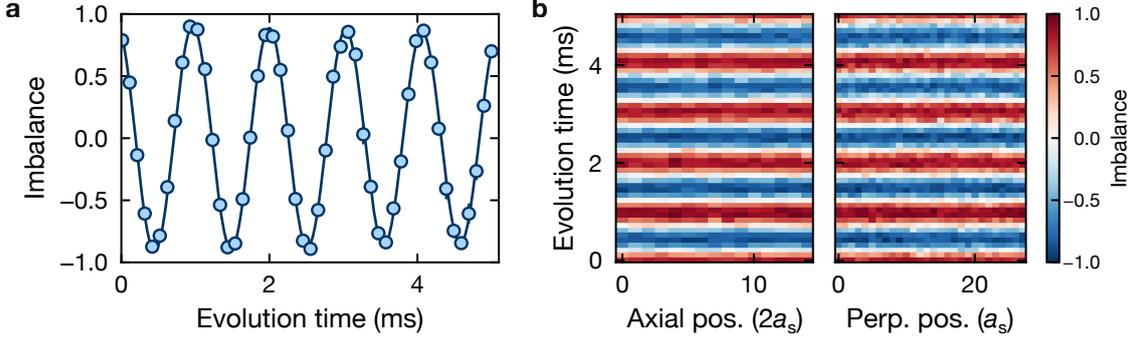


Figure 3.8 | Calibration of the double well tunnel coupling. **a**, Double well oscillation measurement for a long lattice depth of $45 E_{r,l}$ and a short lattice depth of $12 E_{r,s}$ along superlattice axis 2. The solid line is a fit of a sinusoidal function, which allows extracting the intra double well tunnel coupling via its oscillation frequency. **b**, Spatially resolved evaluation of the double well oscillation measurement shown in (a) as a function of the position along the superlattice axis (left panel) and perpendicular to it (right panel). We find homogeneous oscillations with a constant oscillation frequency across the sample.

correct zeroing can also be verified using (slow) double well oscillations, as described below, which should show a constant oscillation frequency across the sample. For further information on the superlattice phase calibration and stability estimates, see also Ref. [213].

3.4.3 Double well tunnel coupling via double well oscillations

We calibrate the tunnel coupling inside a symmetric double well using tunneling oscillations. To this end, we start from a CDW state prepared from a rectangular Mott insulator, where each double well is initialized with one particle in $|L\rangle$ (we arbitrarily set the lower energy well to be $|L\rangle$). After the density distribution is frozen out, we slowly ramp the superlattice phase to the symmetric configuration and set the long lattice depth to the target value. As soon as the phase has stabilized, we suddenly ramp down the short lattice of the superlattice axis to a lower value ($150 \mu\text{s}$ ramp duration), which initiates an oscillation of the particle between the two wells at a frequency of $2J_{\text{DW}}/h$ (corresponding to Rabi oscillations under the symmetric double well Hamiltonian $\hat{\mathcal{H}}_{\text{DW}} = -J_{\text{DW}}\hat{\sigma}_x$, where $\hat{\sigma}_x$ is the Pauli-X operator).

An example double well oscillation trace is shown in Fig. 3.8a, where we find clear sinusoidal oscillations, from which we can extract the tunnel coupling. In this case, we find a tunnel coupling of $J_{\text{DW}}/h = 484.3(5)$ Hz for a long lattice depth of $45 E_{r,l}$ and a short lattice depth of $12 E_{r,s}$ along superlattice axis 2. Double well oscillations are also useful to verify the homogeneity of the superlattice parameters across the sample through a spatial evaluation. As shown in Fig. 3.8b, we find homogeneous oscillations with a constant oscillation frequency across the sample along both spatial axes. If a significant phase gradient would be present, this would be visible as curved oscillation fronts as a function of the perpendicular position (right panel), since the locally varying tilt changes the oscillation frequency according to $f = \sqrt{4J_{\text{DW}}^2 + \Delta^2}/h$. This formula also shows that care must be taken to ensure that the superlattice phase is indeed set to the symmetric configuration, as otherwise the tunnel coupling will be overestimated. However, a detuned oscillation also leads to a decreased amplitude

and oscillations that are asymmetrically localized to the initial imbalance (in accordance with detuned Rabi oscillations), which can be used to cross-check the correct phase calibration.

We repeat this calibration as before for different (short) lattice depths, which allows a prediction of the tunnel coupling for arbitrary lattice depths. Furthermore, it can be used to calibrate the actual depth of the long lattice. Since the tunnel coupling is fully determined by the depths of the long and the short lattice, we can combine the double well oscillation measurement with the previously calibrated short lattice depths (through CDW decay) to fix the long lattice depth. As a result, the band structure of the superlattice is fully calibrated, enabling the use of a band structure calculation to predict Hubbard parameters and the energies of the Bloch bands for quasi-arbitrary lattice configurations.

3.5 Calibration of the interaction energy

After the tunnel coupling, the on-site interaction energy U is the second crucial parameter of the Hubbard model. It is centrally determined by the lattice depths which influence the size of the Wannier functions, as well as the scattering length. In the case of cesium, we can independently tune the scattering length with a magnetic field through a Feshbach resonance (see Fig. 2.1). As the magnetic field at the position of the atoms and the scattering length is not precisely known, we have to directly calibrate the interaction energy using the atoms. The protocol is based on the idea of adding a staggered superlattice potential to a unity-filled 1D chain, where we observe resonant formation of doubly-occupied sites as soon as the on-site interaction energy matches the staggered tilt added by the superlattice (see illustration in Fig. 3.9a). Using a known calibration of the superlattice tilt energy, this provides a direct measurement of the interaction energy. Compared to a direct spectroscopic calibration using parametric modulation via the lattice depth [231] or a superimposed running wave lattice, this method is more accurate as particles cannot delocalize over multiple sites, which can lower the resonance frequency in the latter methods by the order of the tunneling energy.

In the following, we will describe the calibration procedure as an example for the ladder geometry and the parameters used in Chapter 6, but the procedure is straightforwardly translated to other geometries. We start with a ladder geometry consisting of a long-short superlattice along axis 2 (*rung direction*) and a short lattice along axis 1 (*leg direction*), where the superlattice phase is set to realize a strong tilt, causing one leg of the ladder to be energetically higher by about $h \times 4.4$ kHz. The ladder is initialized at half filling with all atoms in the energetically lower leg, which is done by preparing a CDW state using the perpendicular superlattice. The rung lattice parameters are set to the calibrated values $(V_l, V_s, \phi_2) = (50.5 E_{r,l}, 9.4 E_{r,s}, 0.057\pi)$. Due to the strong tilt, the upper leg does not participate in the protocol. We then add the long lattice along the leg direction in a staggered configuration ($\phi_1 = \pi/2$) at a variable depth, which adds a staggered potential offset of the form $(-1)^i \Delta/2$, where i is the leg site index. As a last step, we quickly (within 200 μ s) reduce the short lattice depth along the leg direction to $7.7 E_{r,s}$, which initiates the dynamics. When the additional stagger potential Δ is close to the interaction energy U , we observe coherent production of doublons on adjacent sites $\langle i, j \rangle$ as $|1, 1\rangle_{\langle i, j \rangle} \leftrightarrow |2, 0\rangle_{\langle i, j \rangle} + |0, 2\rangle_{\langle i, j \rangle}$ [232], which we can detect as a drop in the atom number

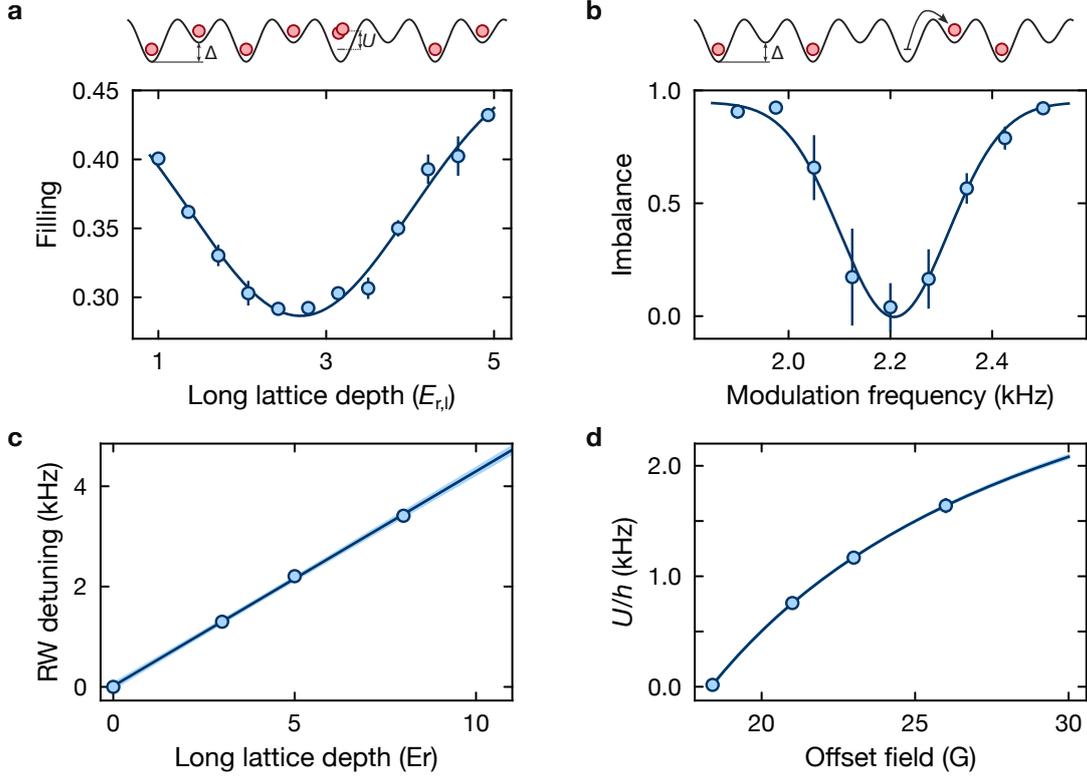


Figure 3.9 | Calibration of the interaction energy. **a**, Observation of the $\Delta = U$ resonance in the filling as a function of the staggered long lattice depth for an offset field of $B = 21.6$ G. The hold time after initiating the dynamics was 10 ms. The solid line is a fit of a Gaussian function, yielding a resonance position at $2.69(3) E_{r,l}$. **b**, Spectroscopic calibration of the stagger potential Δ using a running-wave modulation lattice for $V_l = 5 E_{r,l}$, where we find a resonance in the imbalance at $2.207(5)$ kHz through a Gaussian fit (solid line). In (a) and (b), the error bars denote the standard deviation across two repetitions. **c**, Calibration of the tilt into frequency units using the resonance position from tilt spectroscopy at several recoil depths as in (b). The solid line is a linear fit to the data, yielding the relation $\Delta/h = [0.428(7)V_l + 0.02(4)]$ kHz. The shaded area is the 1σ confidence interval of the fit. Note that the zero point is set manually. **d**, Calibration of the on-site interaction energy as a function of the offset field B in frequency units by combining several $\Delta = U$ resonance at different offset field as in (a) with the calibration of the tilt into frequency units from (c). The solid line is a fit of Eq. 3.4, the error bars denote the fit errors, and the shaded area is the 1σ confidence interval of the fit.

due to parity projection. An example scan of the staggered long lattice depth for an offset field of $B = 21.6$ G in the steep-angle vertical lattice (vertical confinement $2\pi \times 4.7$ kHz) is shown in Fig. 3.9a. As expected, we find a decrease in the observed atom number around a certain stagger potential depth, corresponding to the interaction resonance $U = \Delta$. The behavior is well-described by a Gaussian, which allows us to extract the resonance position at $2.69(3) E_{r,l}$. We repeat this calibration for different offset fields in the ranges of interest, which yields a calibration of the interaction energy as a function of the offset field in units of the long lattice recoil energy. In the next step, we need to calibrate the long lattice recoil energy into frequency units, which we do spectroscopically using the running wave lattice. Note that we cannot use the calibrated long lattice depth from double well oscillations, as the small recoil

depths used here are in a completely different parameter regime than the deep lattices used to isolate individual double wells. To this end, we prepare 1D chains in presence of a staggered potential offset Δ at half filling, with the lower energy sites initially populated (see illustration in Fig. 3.9b). The short lattice is set to the same depth of $7.7 E_{r,s}$ as before, and the long lattice to an adjustable value. Since for typical long lattice depths $\Delta \gg J$, the atoms initially remain localized on the lower-energy sites. We then turn on the running-wave lattice and modulate the system for a duration of 50 ms at a variable frequency. When the modulation frequency f matches the potential stagger Δ/h , the initially maximal imbalance between even and odd sites is equilibrated, which yields a calibration of the long lattice recoil energy in units of frequency.

We show such a tilt spectroscopy measurement for a long lattice depth of $V_l = 5 E_{r,l}$ in Fig. 3.9b, which reveals a clear dip in the imbalance that corresponds to the $\Delta = hf$ resonance. This measurement is repeated for different long lattice depths, using which we can establish a linear relationship between the tilt energy and the long lattice depth, as shown in Fig. 3.9c. Finally, we combine the calibration of the interaction energy as a function of the offset field with the calibration of the tilt energy into frequency units to obtain the interaction energy in frequency units. This is shown in Fig. 3.9d. We fit the data with the expected behavior of the s -wave scattering length under a magnetically tuned Feshbach resonance [220],

$$U(B) = U_0 \left(1 - \frac{\delta_B}{B - B_0} \right), \quad (3.4)$$

where U_0 is a prefactor, δ_B is the width, and B_0 the position, of the resonance. As can be seen, we find excellent agreement with the expected behavior, which allows us to predict the interaction energy for arbitrary offset fields. Note that the zero-interaction point was determined independently using a Hong-Ou-Mandel experiment [92, 233], where we interfere two particles in an isolated double well, for which a maximum contrast is observed at $U = 0$. Lastly, it is important to note that the interaction energy depends on the confinement parameters, and hence Fig. 3.9d is only valid for the specific lattice depths used during the calibration. However, we found that it works well to extrapolate the interaction energy to other lattice depths via the fully calibrated band structure, which was also independently verified using a few additional calibration points at different horizontal lattice depths (see Appendix A).

3.6 Calibration of the laser-assisted tunneling

As introduced in Section 1.3, we engineer artificial gauge fields by resonantly modulating strongly tilted double wells using a running-wave lattice, which results in an effective, complex-valued tunnel coupling. The central calibrations necessary for this are the resonance condition to drive the laser-assisted tunneling processes, the magnitude of the effective tunnel coupling, as well as the engineered flux, which we will describe in the following.

First, one has to select a suitable parameter regime. Considerations for this are in particular:

- The tilt Δ should be much larger than the bare tunnel coupling J_{DW} in the double well, which suppresses contributions to the micromotion on $\mathcal{O}(J_{DW}/\Delta)$.

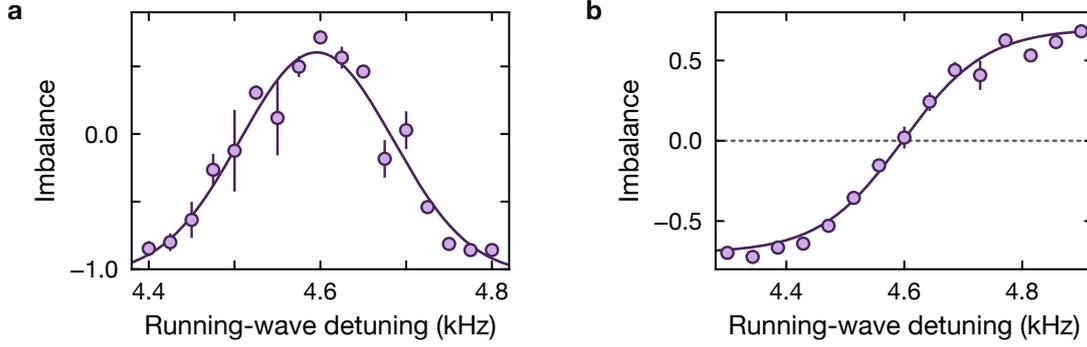


Figure 3.10 | Calibration of the running-wave detuning. **a**, Running-wave lattice detuning calibration using a modulation duration that corresponds roughly to a π pulse. At the resonance, the imbalance is inverted. The solid line is a fit of a Gaussian function, with a fitted center position of $f_{\text{res}} = 4.596(5)$ kHz. The error bars denote the standard deviation across two repetitions. **b**, Calibration including a tilt ramp to adiabatically prepare the single-particle ground state in the driven double wells. The solid line is a fit of a sigmoid function, whose center value corresponds to the resonance frequency, where we find $f_{\text{res}} = 4.597(6)$ kHz. The error bars denote the standard deviation across two repetitions.

- The bare tunnel coupling J_{DW} should be large enough to realize sufficiently high driven tunnel couplings $K \sim h \times 100$ Hz for low modulation strengths V_{mod} , as high drive strengths can lead to unwanted higher-order processes in the effective Hamiltonian and exacerbate heating.
- One has to avoid double well parameters where the modulation frequency is close to band structure resonances of the underlying tilted superlattice potential, which leads to single-particle excitations and heating.

According to these requirements, we typically choose superlattice depths of $V_l = 50.5 E_{r,l}$ and $V_s = 9.4 E_{r,s}$ along with a phase of $\phi = 0.057\pi$ in the axis where the driven tunneling should be realized (here: axis two). A band structure calculation predicts a bare tunnel coupling of $J_{\text{DW}}/h = 812$ Hz and a tilt of $\Delta/h = 4.4$ kHz for these parameters, which are in good agreement with the above conditions. In a next step, the resonance condition for the running-wave modulation has to be calibrated, which we expect at $f = \sqrt{\Delta^2 + 4J_{\text{DW}}^2}/h$. To this end, we prepare a CDW initial state using the same superlattice axis (see Fig. 3.4 and Fig. 3.5b), which prepares each double well with a single particle in $|L\rangle$. Afterward, we turn on the running-wave lattice for a duration that corresponds approximately to a π pulse, after which we expect a transfer to the initially unoccupied $|R\rangle$ sites, inverting the imbalance. When no prior knowledge of the expected π pulse duration is known, one can alternatively also modulate for a long duration (e.g. 50 ms), which works as well but with a smaller contrast. An example of such a measurement is shown in Fig. 3.10a, where we find a resonance at $f = 4.596(5)$ kHz, which is also in good agreement with the band structure prediction.

In most experiments, we want to prepare the ground state of the driven double wells from the $|L\rangle$ initial state. To do this adiabatically, we have to ensure that we are constantly in the instantaneous ground state throughout the ramp. This requirement is however broken when the modulation is turned on at resonance (even with initially zero drive strength), as the double

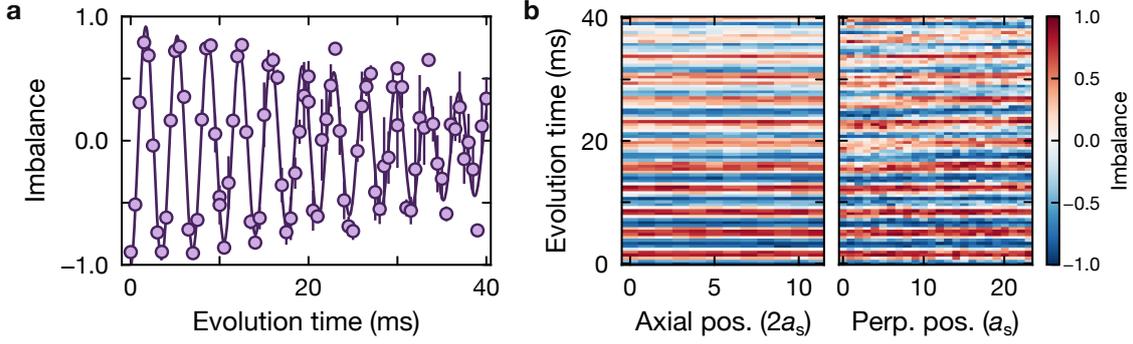


Figure 3.11 | Driven double well oscillations. **a**, Average imbalance as a function of the hold time, showing high-contrast driven tunneling oscillations. The solid line is a fit of an exponentially damped sine, yielding a tunnel coupling of $K/h = 141.9(3)$ Hz and a $1/e$ decay time of $43(7)$ ms. The error bars denote the standard deviation across three repetitions. **b**, Spatially-resolved evaluation of the driven oscillations from (a) as a function of the position along the superlattice axis (left panel) and perpendicular to it (right panel). We find homogeneous oscillations with a constant oscillation frequency across the sample.

well in the effective model is then symmetric, for which the initial state $|L\rangle$ is no longer the ground state. To alleviate this, we use a small additional tilt of 0.01π that is ramped to zero in conjunction with the ramp-up of the modulation strength over a duration of 30 ms. Since this additional tilt ramp could in principle also change the resonance condition slightly, we do a calibration measurement with the tilt ramp included whenever it is employed in an experiment. Fig. 3.10b shows a scan of the running-wave detuning including the additional tilt ramp. In contrast to the previous measurement, the imbalance on resonance is zero, which indicates a preparation of the balanced superposition state. For opposite final detuning, the initial state can then be transferred to the $|R\rangle$ state, similar to an adiabatic passage. The shape of the imbalance is well-described by a sigmoid function, which allows extracting the resonance frequency via its symmetry center.

After the running-wave detuning has been calibrated, we can determine the magnitude of the driven tunnel coupling through driven double-well oscillations. To this end, we start again from a CDW state where each double well is initialized with a single particle in $|L\rangle$, and the superlattice parameters are set to the values for which the resonance was calibrated. We then suddenly turn on the running-wave modulation on resonance over a duration of $150\ \mu\text{s}$, and measure the imbalance as a function of a subsequent hold time. The resulting driven oscillations are shown in Fig. 3.11a, where we find high-contrast oscillations with a frequency that corresponds to a tunnel coupling of $K = h \times 141.9(3)$ Hz, and a $1/e$ decay time of $43(7)$ ms. A spatially-resolved evaluation of the oscillations (Fig. 3.11b) confirms again that the superlattice and running-wave parameters are homogeneous across the sample.

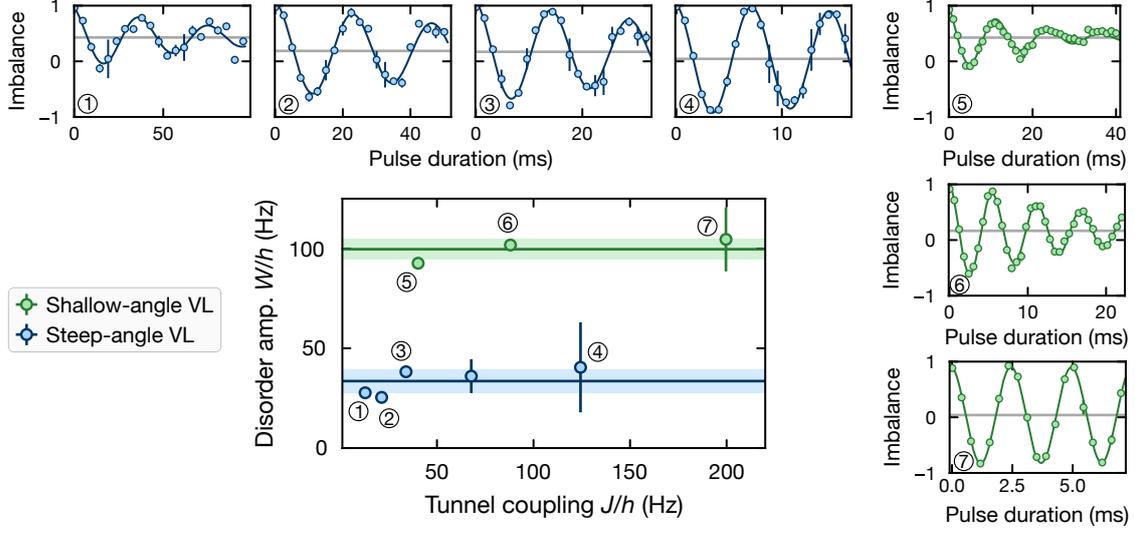


Figure 3.12 | On-site potential disorder estimation. Estimation of the on-site potential disorder via slow double-well oscillations in the shallow-angle (SHVL, green) and steep-angle (STVL, blue) vertical lattices. The detail plots show the raw imbalance oscillations for different tunnel couplings, where the solid lines are fits of exponentially damped sine functions to extract the tunnel coupling J and the steady-state imbalance (gray horizontal lines). The disorder strength is then computed according to Eq. 3.6 and shown in the central plot. The error bars are propagated from the standard errors of the fits, and the horizontal lines denote the average disorder strength, with the shaded area around it the 1σ confidence interval. Here, the STVL was at a power of 20 mW (corresponding to a vertical confinement of $2\pi \times 4.7$ kHz), while the SHVL was at 100 mW (corresponding to $2\pi \times 1.2$ kHz).

3.7 On-site potential disorder

Randomly distributed potential disorder is an important experimental limitation, caused by optical imperfections in the lattice beams or residual, unwanted light fields (e.g. from reflections). Depending on its strength, the random nature of the additional potential terms can lead to a broadening or washing out of the physics of interest, in particular, when the disorder strength is on the order of the relevant physical energy scales such as the tunnel coupling. The average disorder strength is hence an important energy scale that sets a boundary on the experimentally accessible regimes. To estimate the average strength of random on-site potential disorder across our system, we use double well oscillations at low tunnel couplings. The idea is that disorder induces local tilts, which in turn broaden the distribution of oscillation frequencies. Averaged over the sample, this leads to a dephasing of the oscillations as well as a localization toward the initial imbalance. Especially the latter can be detected experimentally with good accuracy when the tunnel coupling is comparable to the disorder strength. We model the imbalance during the double well oscillations by an exponentially damped sine function

$$\mathcal{I}(t) = \bar{\mathcal{I}} + \mathcal{I}_A e^{-t/\tau} \sin(\omega t + \phi), \quad (3.5)$$

where $\bar{\mathcal{I}}$ is the steady-state imbalance, \mathcal{I}_A is the oscillation amplitude, $\omega = \sqrt{4J_{\text{DW}}^2 + \Delta_w^2}/\hbar$ is the modified frequency due to the disorder, and ϕ is an initial phase. Assuming a white-noise

disorder distribution within $[-W, W]$, we can derive an analytical expression for the disorder strength W as a function of the steady-state imbalance $\bar{\mathcal{I}}$:

$$W/h = \sqrt{\frac{6\bar{\mathcal{I}}/\mathcal{I}(0)}{1 - \bar{\mathcal{I}}}}. \quad (3.6)$$

Based on this relation, we extract the disorder strength from double well oscillations at a few different tunnel couplings for both of our vertical lattices. The results are shown in Fig. 3.12, where the detail plots show the raw imbalance oscillations for the different tunnel couplings for the shallow-angle vertical lattice (green) and the steep-angle vertical lattice (blue). As predicted, we find in both cases a stronger localization towards the initial imbalance for lower tunnel couplings, which signifies in a non-zero steady-state imbalance (gray horizontal lines). However, this occurs at different tunnel couplings for the two lattices, which hints at a different disorder strength. Computing the disorder strength according to Eq. 3.6, we indeed find a significantly stronger disorder in the SHVL at an average strength of $W/h = 100(5)$ Hz compared to the STVL with $W/h = 34(6)$ Hz. This three-fold reduction is particularly advantageous in light of the fact that the STVL also has a much larger vertical confinement (here $2\pi \times 4.7$ kHz at 20 mW power compared to $2\pi \times 1.2$ kHz at 100 mW input power), which allows us to work at higher tunnel couplings, further decreasing the sensitivity to disorder.

CHAPTER 4

Single-site reconstruction using unsupervised neural networks

In this chapter, we present a novel method to reconstruct the local occupation in optical lattices with single-lattice-site resolution. The method is based on unsupervised neural networks, which are trained to learn the transformation from the intensity distribution of fluorescence images to local occupation. We devised this method to overcome the limited performance of current state-of-the-art techniques in a regime where the lattice spacing is significantly smaller than the resolution of the imaging system.

In the first section, we will give a brief overview of the image reconstruction problem, and survey current state-of-the-art techniques. Next, we will introduce neural networks and illustrate their application to image reconstruction. We will then present the details of our method, including the network architecture, the training procedure, and the evaluation of the reconstruction performance. Finally, we will discuss the results of the reconstruction and compare them to the performance of previous techniques, as well as elaborate on possible future improvements. Further information on initial work with supervised neural networks can be found in the Master's thesis of Sophie Häfele [234].

4.1 Imaging and reconstruction in optical lattices

Quantum gas microscopy has enabled the study of quantum many-body systems with unprecedented levels of control and detection [30, 31, 38]. Central to this technique is the ability to resolve the local occupation in an optical lattice with single-atom and single-site resolution. This is typically achieved by optically imaging the sample in the lattice and reconstructing the local occupation using a detection algorithm based on the acquired image. To extract more complex observables, such as full counting statistics and (multi-point) correlation functions [33], this reconstruction must be performed with high fidelity. However, as the lattice spacing decreases, the signals from neighboring atoms increasingly overlap, making the reconstruction more challenging. One approach to mitigate this is on the hardware side, for example by using an accordion lattice to dynamically increase the lattice spacing for imaging [235], or by employing quantum gas magnification [37]. Since these methods typically introduce

significant experimental overhead, an alternative is to address the problem on the software side by developing advanced reconstruction algorithms.

In the following, we will start by discussing the imaging process that underlies each experimental image, introducing the physical quantities that determine the image quality – and in turn, how challenging it will be to reconstruct the lattice occupation.

4.1.1 The optical process: Fluorescence imaging

Images in a quantum gas microscope are typically obtained using fluorescence imaging. Here, the sample is illuminated with a near-resonant laser beam, causing the atoms to scatter photons. The scattered photons are then collected by a high-resolution imaging system and imaged onto a camera for a finite exposure time. The chosen exposure times are usually on the order of tens to hundreds of milliseconds, comprising many spontaneous emission cycles for the atoms to emit a sufficiently high number of photons. To inhibit tunneling between lattice sites during the imaging process, the atoms are pinned in their respective lattice sites using a deep pinning lattice. Due to the near-resonant illumination, pairs of atoms on a single lattice site undergo inelastic light-assisted collisions [30, 31, 236], causing the immediate loss of both atoms at the beginning of the exposure (within 10 – 100 μs). Strictly speaking, one thus only measures the parity of the occupation in the lattice, with the exposed sample having either zero or one atom per lattice site (often called *parity projection*)¹.

To understand the quantities determining the properties of a recorded image, it is instructive to consider the different steps comprising the imaging process. As illustrated in Fig. 4.1, we begin with the original fluorescence intensity distribution in the atomic plane. Here, the individual atoms appear approximately as δ -like emitters, with a horizontal extent much smaller than a typical imaging resolution. The width can be estimated as the wave-packet size of an atom in the harmonic potential of each pinning lattice site, which is due to the continuous cooling during imaging of a thermal nature:

$$\sigma_{\text{atom}} = \sqrt{\frac{k_B T}{m \omega_{\text{lat}}^2}}. \quad (4.1)$$

Here, ω_{lat} is the harmonic oscillator frequency of a lattice site, m is the atomic mass of cesium, k_B is the Boltzmann constant, and T is the steady-state temperature during imaging. At a typical trap frequency of $\omega_{\text{lat}} = 2\pi \times 200$ kHz and a temperature on the order of $T \sim 10$ μK , we find $\sigma_{\text{atom}} \approx 20$ nm, which is significantly smaller than the lattice spacing of $a_{\text{lat}} = 383.5$ nm or typical imaging resolutions.

The original intensity distribution then passes through the high-resolution imaging system, consisting typically of a high numerical aperture objective lens in direct proximity to the atoms, and a second tube lens which together form a real image of the atomic plane at a chosen magnification. The action of the imaging system can be absorbed into a convolution with the

¹Recently, a few experiments have implemented a variation of this where the atoms are only subjected to a brief illumination pulse in free space without the use of a pinning lattice [235]. This results in very few captured photons per atom, and hence is only applicable for large lattice spacings, contrary to the short-spacing regime considered here. In this case, it is however possible to measure occupations larger than one.

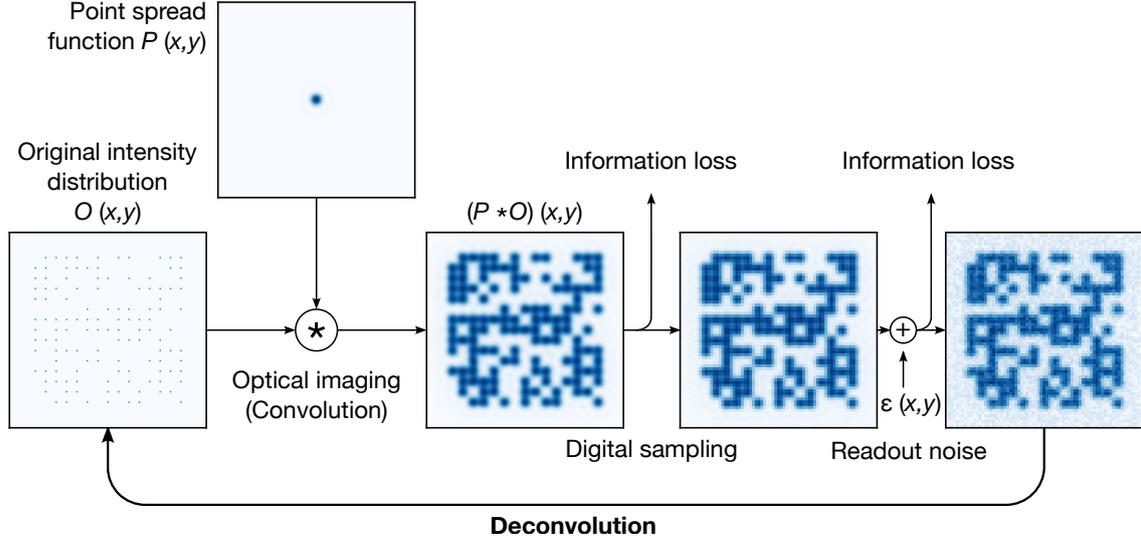


Figure 4.1 | Schematic of the imaging process. First, the original intensity distribution in the atomic plane passes through the optical imaging system, which can be expressed by a convolution with the point spread function of the imaging system. The resulting image is then digitized by the camera, which also adds electronic noise to the image. Both of these lead to a loss of information in the final image. Drawing adapted from [237].

point spread function (PSF) of the imaging system, which denotes the intensity distribution of the image of a point source [238]. It quantifies the spatial resolution of the imaging system and captures aberrations and imperfections. In absence of aberrations, the ideal, diffraction-limited PSF has the form of an Airy disk with a width (defined as the radial position of the first minimum of the disk) given by

$$r_{\text{PSF}} = \frac{0.61\lambda}{\text{NA}}, \quad (4.2)$$

where λ is the wavelength of the imaging light and NA is the numerical aperture of the objective lens. This formula is also known as the *Rayleigh criterion*, as it coincides with the minimal distance of two point sources that can still be resolved by the imaging system. In presence of aberrations or imperfections, the PSF can deviate from the ideal Airy disk, resulting in a broader PSF and a reduced spatial resolution. In addition, specific aberrations such as coma can cause additional challenges during image processing.

As the atoms in the pinning lattice can be treated as point emitters, each atom will appear as a PSF in the image plane, with a width given by Eq. 4.2. This broadening of the intensity distribution of each atom due to the finite spatial resolution of the imaging system is a central step in the imaging process. In particular, the ratio between the PSF width and the separation of neighboring atoms is one of the major factors determining the difficulty of reconstructing the local occupation in the lattice. We denote this as the β -parameter

$$\beta = \frac{r_{\text{PSF}}}{a_{\text{lat}}}. \quad (4.3)$$

For a β close to unity or below, reconstructing the occupation is typically a rather straightforward task, as the individual PSFs of neighboring atoms are well separated. As β increases, in particular close to $\beta = 2$ and higher, the PSFs of neighboring atoms overlap significantly, leading to a loss of information in the image. This is the regime where the reconstruction task becomes challenging, as the individual PSFs are no longer distinguishable in the image. To summarize, the image at this step consists of a convolution of the original intensity distribution in the atom plane with the PSF of the imaging system. Additionally, at this step also stray light can enter the imaging system and be superimposed on the image, which can lead to a significant background signal. In a next step, the image is digitized by the camera, which samples the intensity distribution at discrete pixel positions and converts the analog intensity values to digital values with a finite bit depth. Both of these inevitably lead to information loss, but in particular the finite pixel size causes an important consideration when designing the imaging system: According to the Nyquist-Shannon sampling theorem [228], the magnification of the imaging system must be chosen such that both one lattice period and the PSF width are sampled by at least two pixels. Otherwise, information is lost due to undersampling. At the same time, the magnification should not be chosen too large, as otherwise the signal from one atom is spread out over too many pixels, leading to a reduced signal-to-noise ratio (SNR). Lastly, the digitization process also introduces electronic noise to the image, which can be modeled as an additive noise term. See also Ref. [237] for a detailed discussion and modelling of different noise sources. In conclusion, the image recorded by the camera can be related to the original intensity distribution $O(x, y)$ by

$$I(x, y) = (P * O)(x, y) + \epsilon(x, y), \quad (4.4)$$

where we absorbed all noise sources into one common noise term $\epsilon(x, y)$, and $P(x, y)$ denotes the PSF of the imaging system. Besides these effects, there are also further elements of the fluorescence process itself that influence the SNR of an image. Centrally, the exposure duration determines the number of photons emitted by each atom. By increasing the exposure duration, the signal can be increased, but at the same time, the influence of hopping or loss processes during the imaging is enhanced. Here, an atom obtains enough thermal energy to overcome the potential barrier of the pinning lattice site, leading either to a complete loss of the atom or a re-capturing at a different lattice site. As this can happen stochastically during the imaging, it gives rise to atoms appearing at a fraction of their typical intensity, which increases the spread of the signal distribution and with this reduces the SNR. Furthermore, the width of the signal distribution can be increased by spatial inhomogeneities in the cooling light or the light that drives the fluorescence process, causing a locally varying scattering rate.

In total, we identify two key quantities that determine the quality of the recorded image and how challenging a reconstruction of the lattice occupation will be: (1) the β -parameter, which quantifies the overlap of atoms on neighboring lattice sites, and (2) the signal-to-noise ratio of the image.

4.1.2 The deconvolution problem

As discussed above, one of the central challenges affecting the reconstruction of the lattice occupation is the spread of the signal of individual atoms due to the finite spatial resolution of the imaging system. Most reconstruction methods thus aim to revert this effect, effectively reducing the β parameter to a value where the individual PSFs of neighboring atoms are well separated. This process is known as *deconvolution*, and aims to infer the original intensity distribution in the atomic plane from the recorded image (i.e., to invert the steps illustrated in Fig. 4.1). Mathematically, we can formulate this as follows: Given the recorded image $I(x, y)$, the goal is to infer the original intensity distribution $O(x, y)$ in the atomic plane. This corresponds to solving Eq. 4.4 for $O(x, y)$. In the absence of noise, this can be done for example in Fourier space as

$$O(x, y) = \mathcal{F}^{-1} \left[\frac{\mathcal{F}[I(x, y)]}{\mathcal{F}[P(x, y)]} \right], \quad (4.5)$$

using knowledge about the shape of the PSF. However, in presence of noise, the noise term has to be retained in the inverse problem

$$O(x, y) = \mathcal{F}^{-1} \left[\frac{\mathcal{F}[I(x, y)]}{\mathcal{F}[P(x, y)]} - \frac{\mathcal{F}[\epsilon(x, y)]}{\mathcal{F}[P(x, y)]} \right]. \quad (4.6)$$

Due to the existence of a cutoff frequency for the PSF, high-frequency noise components comparable or above the cutoff frequency can be amplified in the inversion. This makes the deconvolution an ill-posed problem, and has been a central challenge in image processing for many years [239]. Several approaches have been developed to address this problem, mainly in the context of astronomical imaging or microscopy. One immediate approach is to regularize the inversion by adding a penalty term that suppresses high-frequency noise components, such as the Tikhonov regularization [240]. Another method of stabilization is to use a Fourier filter that incorporates the SNR of the image, such as the Wiener filter [241]. Here, the idea is to modify the Fourier-space inversion and compute the deconvolved image via

$$O(x, y) = \mathcal{F}^{-1} \left[\frac{\mathcal{F}[I(x, y)]}{\mathcal{F}[P(x, y)]} \frac{|\mathcal{F}[P(x, y)]|^2}{|\mathcal{F}[P(x, y)]|^2 + K} \right], \quad (4.7)$$

where K is a parameter that depends on the SNR of the image and can be optimized numerically using e.g. a scheme as in Ref. [242]. The Wiener filter is still widely used today in different fields, and has even been applied in the context of quantum gas microscopy. As this approach however often leads to ringing artifacts [239], further methods have been developed, among which the Richardson-Lucy algorithm is one of the most prominent [243, 244]. This is an iterative method, based on maximum likelihood estimation with the goal to calculate the most likely original intensity distribution given the recorded image. Lastly, a more recent approach is the use of deep learning methods, which have been shown to be particularly powerful in tackling ill-posed inverse problems such as the deconvolution due to their ability to learn complex, non-linear transformations from data [245, 246]. This is also the approach that we will be following, since neural networks can be tailored well to make use of a-priori knowledge about our system, and we have access to high-quality training data directly from the experiment.

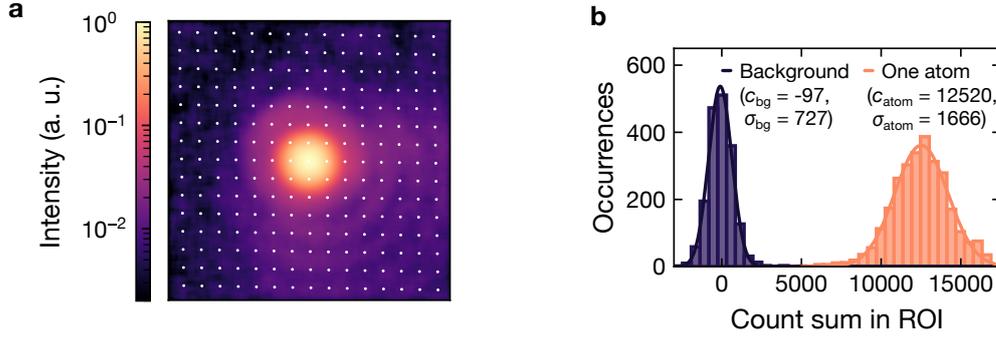


Figure 4.2 | Experimental imaging conditions. **a**, Experimentally measured point spread function, obtained by averaging around 2000 isolated single atom signals. The white dots denote the lattice site centers. The central peak spans several lattice sites, and a fit gives a Rayleigh width of around 850 nm. The logarithmic scale highlights a long-ranged, asymmetric tail of the PSF. **b**, Signal-to-noise ratio determined from the count sum distributions of crops containing exactly one atom (signal) or zero atoms (background). The solid lines denote Gaussian fits, from which we extract an SNR of 5.2 at an exposure time of 300 ms. The crops were obtained from sparse images with a crop width of 1.6 μm . A background image without atoms was subtracted to shift the background peak center closer to zero.

4.1.3 Experimental parameter regimes

To understand the need for novel reconstruction approaches, it is instructive to revisit the experimental imaging parameters in our apparatus and compare them to existing quantum gas microscopes. In Fig. 4.2, we show the experimentally measured point spread function (PSF) and the signal-to-noise ratio (SNR) of our imaging system, as obtained from fluorescence images of sparsely populated samples (see Section 3.1). The averaged PSF shows a round, central peak spanning several lattice sites. Through a 2D Gaussian fit, we extract a Rayleigh width of around 850 nm. This is significantly larger than the lattice spacing of $a_{\text{lat}} = 383.5$ nm, resulting in a β parameter of around 2.2. In addition, the logarithmic scale reveals a long-ranged, asymmetric tail of the PSF, protruding over more than 10 lattice sites into the bottom-right corner. This tail likely originates from residual coma aberrations in the imaging system, which we were not able to fully correct for during alignment. Even though the tail has only a fraction of the amplitude of the central peak, it still complicates the reconstruction significantly, as we will show later.

To determine the signal-to-noise ratio, we use the count sum distributions of crops containing exactly one atom (signal) or zero atoms (background). An experimentally determined count sum distribution for an exposure duration of 300 ms is shown in Fig. 4.2b. The achieved SNR is mainly determined by the average number of photons emitted by each atom during the exposure time, the spread of this number due to inhomogeneities or hopping and loss processes, as well as the width of the background signal. To capture these contributions, we define the SNR as

$$\text{SNR} = \frac{c_{\text{atom}} - c_{\text{bg}}}{\sigma_{\text{atom}} - \sigma_{\text{bg}}}, \quad (4.8)$$

where c_{atom} and c_{bg} are the average counts of the atom signal and the background signal, respectively, and σ_{atom} and σ_{bg} are the respective standard deviations. We extract all four

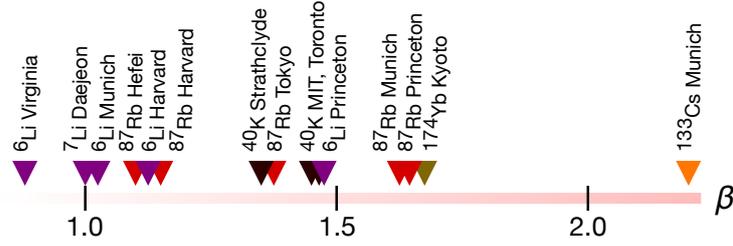


Figure 4.3 | Overview of β parameters among quantum gas microscopes. Most existing quantum gas microscopes work with a β parameter between 1.0 and 1.5, whereas our cesium experiment operates in a entirely different regime of $\beta > 2$. Non-exhaustive list, state as of mid 2022, data from Refs. [30, 31, 247–253, 255, 256].

parameters through Gaussian fits of the background and signal peaks, and find an SNR of 5.2 for an exposure time of 300 ms.

To put the measured parameters into context, we show an overview of the β parameter for a selection of quantum gas microscopes in Fig. 4.3 [30, 31, 247–256]. From the comparison, it is evident that most existing quantum gas microscopes work with a β parameter between 1.0 and 1.5. In contrast to that, the determined $\beta = 2.2$ puts our cesium experiment in an entirely different, previously unexplored regime. This makes the reconstruction of the lattice occupation a challenging task, and required us to investigate novel approaches for reconstruction and deconvolution methods. Of course, there is no such thing as a free lunch. Achieving a high-fidelity reconstruction at such a high β parameter requires us to work also at a particularly high signal-to-noise ratio. Comparing this between different experiments is not straightforward, as few experiments have carefully characterized the SNR. In general, one can however expect that experiments using optical molasses (often rubidium) during fluorescence images achieve an SNR comparable to ours due to a similar number of scattered photons, while Raman sideband cooling (often lithium) typically yields a far lower photon number, which is also reflected in the distribution of different elements in Fig. 4.3.

4.1.4 Occupation reconstruction and common approaches

It is important to distinguish the topics of deconvolution and reconstruction of the lattice occupation. For our application, these topics are closely related, but refer to different aspects of the problem. Reconstruction deals with the inference of the local occupation in the lattice from measured fluorescence images, i.e., to obtain a sample of the local density operator \hat{n}_i on each lattice site i . This is the ultimate goal of the reconstruction step and one of the main ingredients in quantum gas microscopy. Deconvolution, on the other hand, is a further step in the reconstruction process, which can be applied to enhance the reconstruction fidelity in challenging parameter regimes. To understand this distinction better, we can study the simplest method of reconstructing the lattice occupation without any deconvolution, which is the pixel count summation method.

Pixel count summation

In the pixel count summation method, a fluorescence image is first divided into regions of interest (ROIs) corresponding to the individual lattice sites. For this, a construction as for Wigner-Seitz cells is used, i.e., the area for each lattice site is delimited by the midpoint lines between neighboring lattice sites. Then, the pixel counts in each ROI are summed up and a histogram for the count sums is constructed. For a moderate β parameter and sufficient SNR, one obtains a bimodal distribution of the pixel counts, with one peak corresponding to the empty lattice sites and the other peak to the occupied lattice sites. The occupation on each site is then inferred via a threshold value separating the two peaks, which yields the reconstructed occupation with single-site resolution.

In Fig. 4.4, we investigate this based on simulated fluorescence images with different β parameters. For the simulation, we assume a similar SNR as in the experiment, and produce images at half filling, neglecting density-dependent effects (see Ref. [234] for details on the simulation). In the left-hand column of Fig. 4.4, we show one example image at half filling for each β . A larger β causes the signals of atoms on neighboring sites to overlap increasingly, making a distinction of empty and occupied sites by eye impossible for $\beta \gtrsim 1.5$. The increasing overlap of neighboring PSFs is also reflected in the pixel count sum distributions, shown in the right-hand column of Fig. 4.4. For small β , we observe two well-separated peaks, corresponding to empty and occupied sites, respectively. Based on this histogram, we can define a threshold value that separates the two peaks at their crossing point, and via this threshold infer the occupation in each lattice site. In such a regime of small β and sufficient SNR, this constitutes a straightforward method to reconstruct the lattice occupation with single-site resolution.

However, with increasing β , the two peaks move closer to each other and start to merge and gain in overlap. As a result, a growing number of errors is made when distinguishing empty and occupied sites via a threshold value, reducing the reconstruction fidelity. Specifically, empty sites with a count sum larger than the threshold value are misclassified as occupied sites (*false positive*), and vice versa occupied sites with a count sum below the threshold are misclassified as empty sites (*false negative*). Eventually, for our experimental condition with $\beta = 2.2$, the two peaks are no longer distinguishable in the combined pixel count sum distribution of all sites, and this method fails to provide a meaningful reconstruction. If one would still apply a threshold, around 20% of all sites would be incorrectly reconstructed. The merging of the two peaks occurs at a β parameter of around 2.0 for comparable SNRs, which marks a transition point where further image processing is unavoidable to enable single-site reconstruction. Note also that we have neglected the influence of density-dependent effects as well as other experimental imperfections in this simulation, which would further complicate the reconstruction task. Secondly, the count sum distribution is also sensitive to the SNR, hence, a lower SNR would lead to a merging of the peaks already for lower β values.

From this discussion, it is evident that we have to apply further techniques to enable reconstruction in the high- β regime and boost the reconstruction fidelity in the intermediate regimes. The solution to this is to apply a deconvolution method to the recorded images, which aims to revert the main driving force behind the merging of the peaks: the spillover of signal into neighboring lattice sites due to the finite spatial resolution of the imaging system. After a

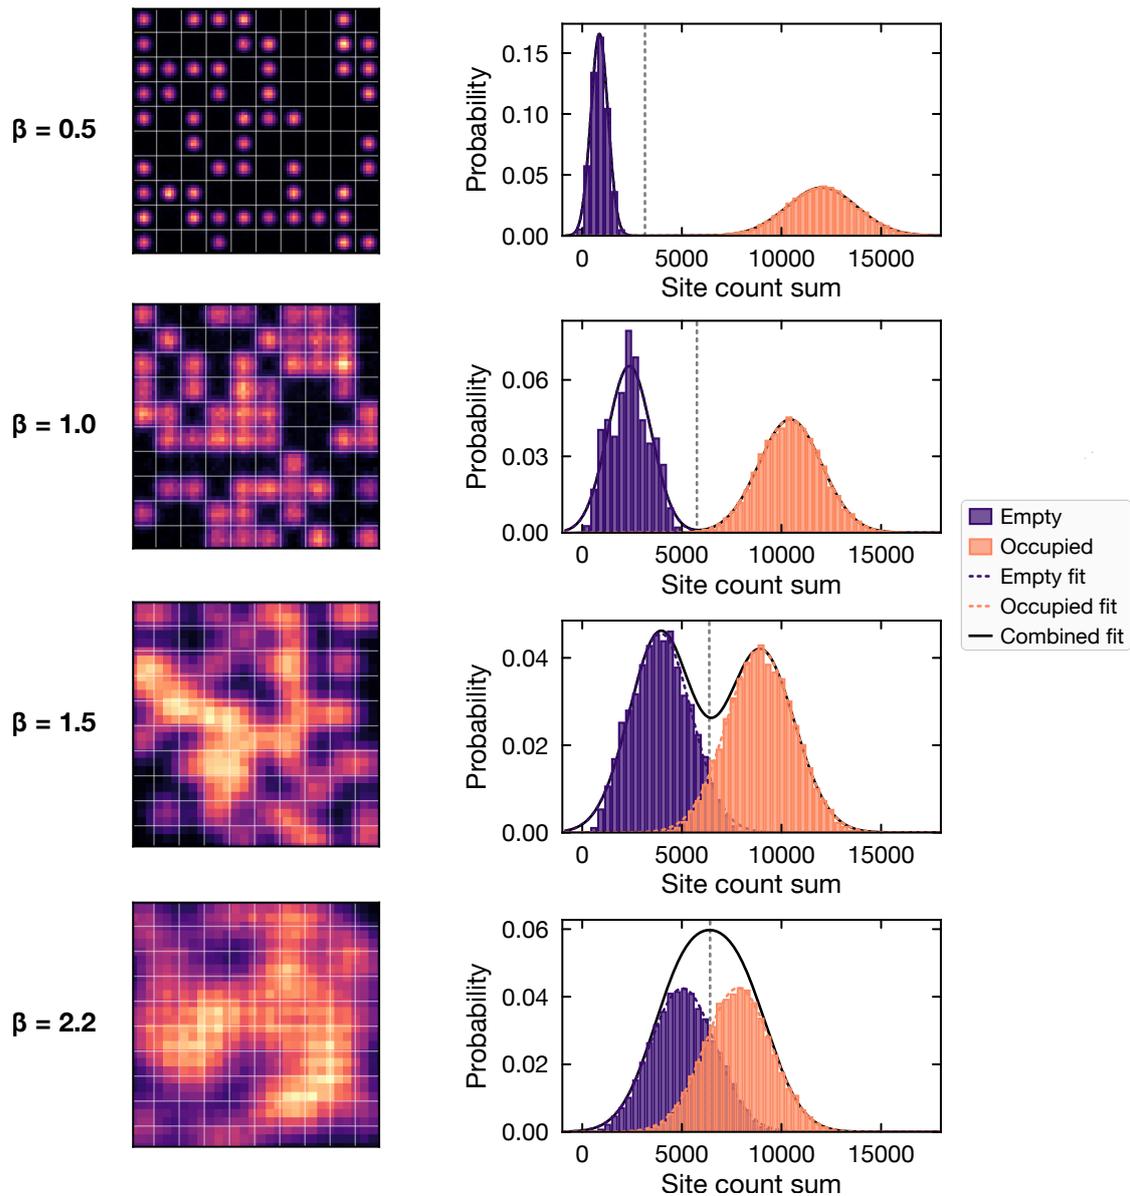


Figure 4.4 | Pixel count sum distributions for different β . Using a fluorescence simulation, we investigate how the distributions of the pixel count sums within each lattice site change as a function of the resolution-to-spacing ratio β . In the left column, one example image at half filling is shown for each β . Here, the light gray lines denote the lattice site borders. Note that the pixel size in the example images was decreased for the smaller resolutions for a better visualization. In the right column, the corresponding pixel count sum distributions are shown, each computed over around 25000 lattice sites at half filling. The simulation assumes a similar SNR as in the experiment. The solid purple (orange) lines denote Gaussian fits to the empty (occupied) site distributions, and the solid black line is the sum of the individual fits, showing the total distribution as it would be seen without prior information on the classes. The dashed line denotes the optimal reconstruction threshold.

successful deconvolution, we expect the image to be closer to the original intensity distribution corresponding to a significantly reduced β , effectively increasing the separation between the

peaks. This allows us to distinguish again via a threshold at a much higher fidelity. In the following, we will discuss a few common deconvolution methods that have been applied in the context of quantum gas microscopy.

Fitting of point spread functions

The first microscope experiments implemented a deconvolution based on fitting PSFs on each individual lattice site [30, 31]. This can be done either in a least-squares fashion [30], or by iteratively trying out different combinations of occupied and empty sites [31], both with the goal to minimize the difference between the measured image and a trial image constructed from PSFs placed with a variable amplitude on the lattice grid. After the deconvolution through the fit, each lattice site is assigned with a fitted PSF amplitude, and a histogram of these amplitudes should show a bi-modal distribution. As before, this allows to reconstruct the binary lattice occupation via a threshold. One central drawback of this method is that the fit can only be performed on small subsections of a fluorescence image at a time, as noise and spatial variations accumulating over the whole image would otherwise dominate the fit loss and prevent its convergence towards the correct result. This causes a significant computational overhead, with the method taking up to one minute for the reconstruction of a single image with several hundred to thousands of lattice sites. Additionally, this approach leads to poor fidelities for $\beta \gtrsim 1.5$, as the fit objective (average squared difference) becomes too insensitive when the PSFs of neighboring atoms overlap significantly, especially for a limited SNR [242].

Deconvolution via a weight kernel

A different approach that has been proposed first in Ref. [257] is based on deconvoluting the image with a linear kernel matrix (see also Section 4.2.3). This kernel is constructed using the PSF of the imaging system to have unity overlap with the signal on a site of interest, and minimum overlap with the signals on all neighboring sites. Denoting the kernel on site (i, j) as $k_{i,j}$ and a PSF on site (m, n) as $p_{m,n}$, the kernel matrix is constructed to satisfy [257]

$$\sum_{m,n} k_{i,j} p_{m,n} = \delta_{i,m} \delta_{j,n}, \quad (4.9)$$

where $\delta_{i,m}$ is the Kronecker delta. Due to the discrete translational symmetry, the kernel can be expressed as a linear combination of PSFs on all sites in the vicinity of the site of interest, with prefactors that are determined numerically [257] or analytically [258] according to Eq. 4.9. To implement the deconvolution, the image is then multiplied with kernel matrices centered on each lattice site, i.e., the image is convolved with the kernel at a step size equal to the lattice spacing. According to the above construction, the action of the kernel is to invert the spillover of signal into neighboring lattice sites, effectively reducing the β parameter. After the deconvolution, we are left with the signal amplitude on each site, which should – as before – exhibit a bi-modal distribution, enabling a reconstruction of the lattice occupation via a threshold.

The central advantage of this method is the computational simplicity, making it very fast (on the order of a few ms for the deconvolution itself) compared to iterative or fit-based methods.

However, we found that this method is also limited in its performance for the large- β regime that we are interested in. On the one hand, it is not able to capture density-dependent effects since it only assumes a single kernel that does not depend on the number of neighboring atoms. On the other hand, a large PSF size causes the kernel to have a large spatial extent, leading to an accumulation of noise from many neighboring sites, which in sum can quickly dominate the physically relevant signals. The method is hence mostly applicable to boost the fidelity in low- to intermediate- β regimes, but not to enable a reconstruction in the high- β regime.

Wiener and Richardson-Lucy deconvolution

A third set of methods is based on image restoration techniques such as the Wiener filter [241] or the Richardson-Lucy deconvolution algorithm [243, 244]. As discussed before, these are general deconvolution techniques, developed in the context of image deblurring and astronomical image analysis, but have been applied to quantum gas microscopy as well [242, 251]. These methods are well understood due to their widespread use, but a central disadvantage is that they are not tailored to the specific properties of the quantum gas microscopy images, and it is for example challenging to incorporate the discrete lattice structure or the binary nature of the lattice occupation. This severely limits the performance of these methods, and in fact, they have been found to rapidly decrease in fidelity in the high- β regime [242]. Building on these techniques, Ref. [237] proposed a method to extend Wiener filtering with a sophisticated noise model and a way to integrate the discrete lattice grid. This showed enhanced fidelities, but is complex to implement, requires detailed calibration and has so far only been demonstrated in 1D.

4.1.5 Reconstruction of experimental data with conventional approaches

To motivate the need for novel reconstruction approaches in our parameter regime, we apply a few of the previously introduced reconstruction approaches to our experimental data and benchmark their performance. To set the stage, let us first take a look at the site count sum distributions of the experimental data, according to which one could already reconstruct the lattice occupation in a low- β regime. In Fig. 4.5, we show the site count sum distributions for three different fillings. It is immediately evident that the distributions exhibit no separation

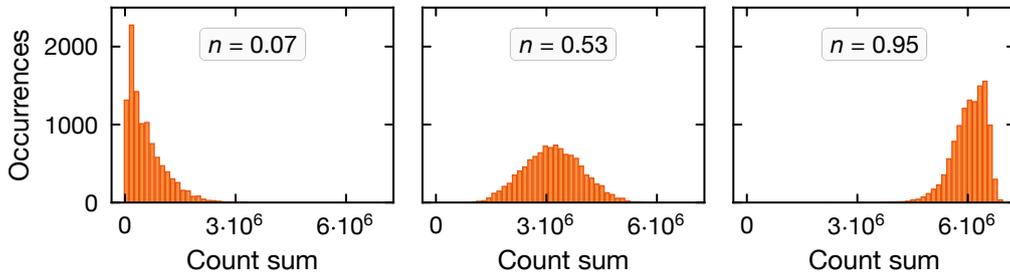


Figure 4.5 | Experimental site count sum distributions. Unprocessed distributions of the pixel count sums within a lattice site for three different average fillings from the validation set. In no case, a separation between empty and occupied sites is visible, which precludes a direct reconstruction of the lattice occupation without deconvolution.

between empty and occupied sites at any filling. As a result, a direct reconstruction of the lattice occupation via a thresholding of the site count sums is not possible, and an additional deconvolution is unavoidable.

Deconvolution with a linear kernel

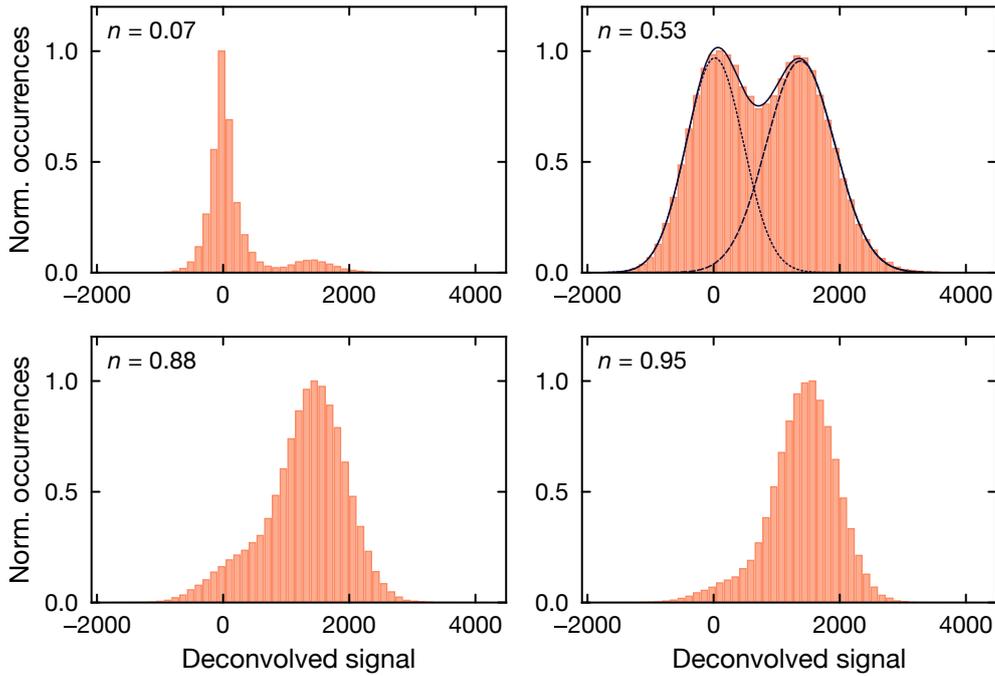


Figure 4.6 | Deconvolved signals for the linear kernel method. Distributions of the deconvolved signals for four different fillings from experimental data. A separation between empty and occupied sites is visible for low fillings up to around half filling, but vanishes for the highest fillings, where the deconvolution fails. In the panel for half filling, the solid line is a bi-modal Gaussian fit, capturing the individual distributions of empty (dotted line, left peak) and occupied (dashed line, right peak) lattice sites, respectively. From the overlap fraction, we estimate a fidelity of around 91% at half filling. The occurrences are normalized to the respective maximum in each plot, and the average filling of each bin was determined using the autoencoder reconstruction presented later.

We start by investigating the performance of a linear kernel deconvolution on our experimental data. Due to the conceptual simplicity and very fast computation, it is a popular approach and applicable across a wide range of experiments with moderate β -parameter. As introduced above, it is computed solely based on the experimentally measured PSF and the lattice grid. Attempting to construct the kernel for our parameter regime by inverting Eq. 4.9, we found that the large extent of the PSF causes the computational complexity of the inversion to explode, as PSFs on many neighboring sites have to be taken into account. A restriction to a smaller kernel size as proposed in Ref. [257] was not feasible either, as the resulting kernel would not be correct and have a significantly reduced performance. Instead, we resorted to determining the kernel from simulated fluorescence images via the linear kernel neural network as described in Section 4.2.3. The validity of this is based on the assumption that the simulation

uses the measured PSF and lattice geometry, which are the defining ingredients for the shape of the kernel. The kernel is then applied to the experimental data to perform the deconvolution.

Fig. 4.6 shows the deconvolved signal distributions for four different average fillings using experimental data. The kernel has a size of 9×9 lattice sites. As can be seen, this method is indeed able to separate the empty and occupied sites for low fillings up to around half filling, as is evident from the bimodal distribution of the amplitudes. However, the separation is only weakly pronounced, and we estimate a fidelity of around 91% at half filling through the overlap fraction (see also Section 4.2.3 for an evaluation on simulated data at half filling). Importantly, the separation vanishes for fillings above ~ 0.6 , and the deconvolution completely fails for the highest fillings near unity. In this case, no reconstruction is possible any longer. The reason for this is likely the onset of density-dependent effects in combination with the large extent of the PSF, which decrease the visibility of individual holes in high-density regions too strongly.

Deconvolution by fitting of point-spread functions

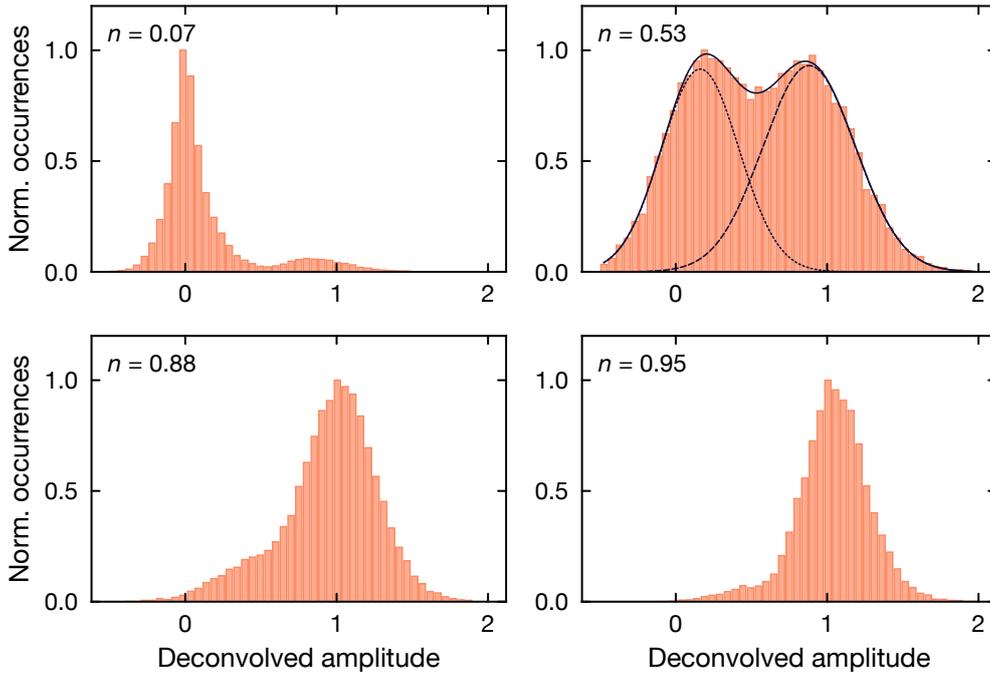


Figure 4.7 | Deconvolved amplitudes for the PSF fitting method. Distributions of the deconvolved PSF amplitudes α_n for four different fillings from experimental data. A separation between empty and occupied sites is visible for low fillings up to around half filling, but vanishes for the highest fillings, where the deconvolution fails. In the panel for half filling, the solid line is a bi-modal Gaussian fit, capturing the individual distributions of empty (dotted line, left peak) and occupied (dashed line, right peak) lattice sites, respectively. From the overlap fraction, we estimate a fidelity of around 90% at half filling. The occurrences are normalized to the respective maximum in each plot, and the average filling of each bin was determined using the autoencoder reconstruction presented later.

A second common approach is to perform the deconvolution by fitting of PSFs on a discrete lattice grid. Specifically, we take a subsection of a fluorescence image with a size of about

20×20 lattice sites, and then compose a trial image consisting of an array of PSFs located at pre-determined lattice sites \mathbf{x}_n with amplitudes α_n . The trial image is thus given by

$$I_{\text{trial}}(\mathbf{x}, \boldsymbol{\alpha}) = \sum_n \alpha_n \text{PSF}(\mathbf{x} - \mathbf{x}_n), \quad (4.10)$$

where \mathbf{x} is the location in the image. The amplitudes α_n are then determined by minimizing the squared difference with the input image. Such an approach is for example used in Refs. [30, 31, 252] in this or a similar form. For the creation of the trial image, we use the experimentally measured PSF. The optimization is performed through stochastic gradient descent (ADAM optimizer [259], learning rate 0.5, 200 iterations). The resulting distributions of the PSF amplitudes α_n (i.e., equivalent to the deconvolved site signals) are shown in Fig. 4.7 for four different fillings using experimental data. This method appears to perform similar to the linear kernel deconvolution above, successfully creating a separation between empty and occupied sites for low fillings up to around half filling. The fidelity at half filling is slightly lower, with an estimated value of around 90%, and the separation vanishes already for fillings above ~ 0.5 . The slightly worse performance is likely due to the fact that the combination of a broad PSF and a high filling fraction lead to small changes in the mean squared error of the fit during the optimization, preventing good convergence. For the highest fillings, the deconvolution completely fails, and no reconstruction is possible any longer.

In summary, it is clear that both methods are not applicable in our resolution regime, as they fail to enable a reconstruction for fillings above ~ 0.5 , and only provide an insufficient fidelity for lower fillings. This highlights the need for novel reconstruction approaches that are tailored to the high- β regime and enable a high-fidelity reconstruction at all fillings.

4.2 Neural networks for image reconstruction

In this section, we will introduce the basic concepts of neural networks. As there is a lot of literature available on this topic, we will keep the discussion brief and focus on the aspects that are relevant for our application. The interested reader is referred to specialized literature such as Refs. [84, 260]. While machine learning is a vast field, we only need a basic understanding of the central principles here and additionally, we will see that many aspects of neural networks can be related to concepts we know from physics. After introducing the basic ingredients and describing the differences between supervised and unsupervised learning, we will construct a toy example to illustrate naturally how neural networks can be used for reconstructing the lattice occupation from fluorescence images.

4.2.1 Brief introduction to machine learning and neural networks

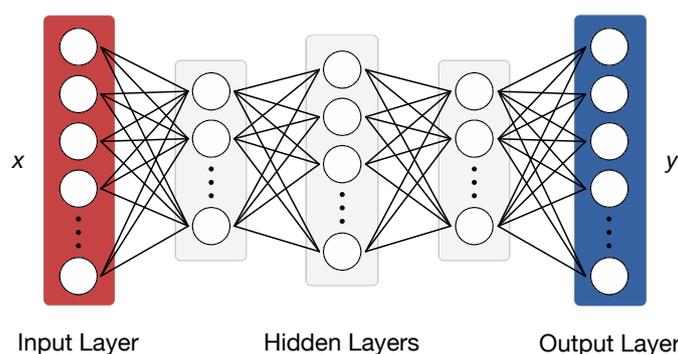


Figure 4.8 | Schematic of a feedforward neural network. A classic feedforward neural network consists of an input layer, one or more hidden layers, and an output layer. Each layer consists of interconnected nodes, called neurons. The connections between neurons in adjacent layers are characterized by weights, which determine the strength of the connection. The output of a neuron is given by the weighted sum of the inputs, passed through an activation function.

The field of machine learning deals with the process of building models from data using optimization and regression techniques. As a major subfield, neural networks are a class of machine learning algorithms that are inspired by the structure and function of the human brain [261]. Shown in Fig. 4.8 is the architecture of one of the classic neural network models that gave rise to the field of deep learning – the *feedforward neural network*. A feedforward neural network consists of interconnected nodes, called neurons, which are organized in layers. There is always an input layer that receives the input data, and an output layer that produces the output of the network. In between, there can be one or more hidden layers. Every neuron receives input signals, processes them, and passes the output to the next layer of neurons. The strength of the individual connections between neurons is characterized by weights. In this

architecture, the output of a neuron is determined by the weighted sum of the inputs, passed through an activation function. This means that the output y of a neuron is given by

$$y = f\left(\sum_i w_i x_i + b\right), \quad (4.11)$$

where x_i are all the inputs to this neuron, w_i are the connection weights, b is a bias term, and f is the activation function. Common choices for the activation function are the sigmoid function, with

$$f(x) = \sigma(x) = \frac{1}{1 + e^{-x}}, \quad (4.12)$$

the hyperbolic tangent function,

$$f(x) = \tanh(x) = \frac{e^x - e^{-x}}{e^x + e^{-x}}, \quad (4.13)$$

or the rectified linear unit (ReLU) function,

$$f(x) = \max(0, x) = \begin{cases} x & \text{if } x > 0, \\ 0 & \text{otherwise} \end{cases}. \quad (4.14)$$

The purpose of the activation function is to introduce non-linearity into the network, which allows it to approximate complex, non-linear relationships.

To have such a neural network learn a specific model or relationship, it needs to be trained. This learning or training process is typically done using a large dataset of input-output pairs. Through an optimization algorithm, the parameters of the network are adjusted with the goal to find an optimal set of parameters that minimize the difference between the predicted output of the network and the true output. This difference is quantified by a *loss function*, which measures the error of the network's prediction and is tailored to the task at hand. In the case of the feedforward network, the parameters that are optimized are the connection weights and bias values of all the individual neurons. The optimization is done iteratively using some (stochastic) gradient descent algorithm [262]. In each iteration, also called an *epoch*, the network is fed with batches of input data, which is passed through the network layer by layer to obtain a prediction (forward pass). The output is then compared to the true output via the loss function. In the subsequent backward pass, the gradient of the loss function with respect to the weights is calculated using the backpropagation algorithm [263]. Finally, this gradient is used to update the weights in a way that minimizes the loss function, where the change is scaled by a factor called the *learning rate*, which poses an important tuning knob.

After the training has converged, the network will have learned a function that maps the input data to the output data. This function can then be used to make predictions on new, unseen data, which is the goal of the machine learning process. Creating a model that generalizes well to unseen data is a non-trivial task and requires a careful optimization of the network architecture, the training data, and also a careful evaluation using validation data that

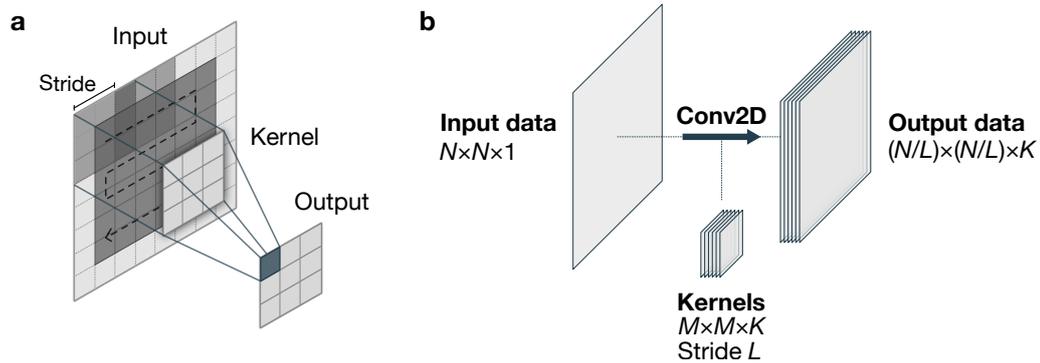


Figure 4.9 | Idea of a convolutional layer. **a**, In a convolutional layer, a discrete convolution operation with a set of filters is applied to the input data. The filters are swept across the input data, locally computing the inner product between the filter and the input data at each position. The result of each inner product then gives one pixel of the output image. **b**, The shape of the output is modified according to the stride and the number of filters. Note that edge effects and incommensurate dimensions can lead to deviations in the actual shape, which can be influenced by padding the data.

was not used during training. To optimize the training and generalization performance, one typically also has to tweak multiple *hyperparameters*, i.e., all the parameters that are not actively changed by the optimization algorithm during a training run. These can be, for example, the number of hidden layers, the number of neurons in each layer, the choice of activation function, or optimization parameters such as the learning rate and the batch size.

In other words: We can regard a neural network as a flexible, highly-parametrized fit function that is fitted to data using an optimization algorithm, with similarities to the least-squares fit of analytical or numerical models to data in physics. The main differences to traditional fitting methods are that the neural network can approximate arbitrary complex, non-linear functions, and that the optimization is done using a large dataset of input-output pairs. In fact, through the universal approximation theorem, it is known that a feedforward neural network with a single hidden layer and non-linear activation functions can approximate any continuous function on a compact domain to arbitrary accuracy, given a sufficiently large number of neurons in the hidden layer [264]. This highlights once more the power of neural networks in building models from (observational) data, in particular, when the governing relations underlying the data are not fully known or are very complex.

Since the introduction of the feedforward neural network, many further architectures have been developed, each with individual operating principles, and hence, applicable to different problems. Most important for us is the *convolutional neural network* (CNN), which is particularly well suited for image data [265, 266]. The basis of this architecture is the convolutional layer. In contrast to the fully connected layer, where weighted sums of the inputs to each neuron are computed, the convolutional layer applies a discrete convolution operation to the input data. This convolution is performed with a set of filters (also called *kernels*), whose entries are optimized during the training process. As illustrated in Fig. 4.9, each filter is swept across the input data, locally computing the inner product between the filter and the input data at each position. The result of each inner product then gives one pixel of the output image.

The sweep across the input data can also be done at a step size (*stride*) larger than one, which reduces the spatial resolution of the output image. After repeating this process for all filters, the individual outputs (also called *feature maps*) are stacked and finally passed through an activation function, forming the overall output of the convolutional layer.

A central advantage of the convolutional layer is that it allows weight sharing, i.e., the same filter is applied to all positions of the input data. This dramatically reduces the number of parameters in the network compared to e.g. a feedforward architecture, and is particularly useful for image data. Consider for example the detection of a specific shape in an image. In a feedforward network, each neuron would have to learn the shape independently for all locations in the image, while in a convolutional network, the same filter is applied to all locations. This is also one first example of how we can incorporate prior knowledge about the problem into the network architecture. In particular, the convolutional operation can be used to encode a (discrete) translational symmetry, which is often present in image data. This applies especially also to our application of fluorescence images from a quantum gas microscope, where the lattice sites are arranged in a regular grid and the PSF and imaging conditions are expected to remain similar across the lattice.

Another noteworthy class of neural networks are *recurrent neural networks* (RNNs), where information can also flow in loops from later to earlier layers, with applications in sequential data such as time series or text data [267]. This has no use in our concrete problem, but can be advantageous for example when analyzing the time evolution of a quantum many-body system. These different architectures, or specific components of them, are often also combined to form more complex, hybrid networks, which are tailored to specific problems. Recently, the increasing usage of machine learning techniques in the physical and other natural sciences has further led to the development of methods that no longer follow a mostly data-driven approach, but also incorporate detailed knowledge about the underlying system. In so-called physics-informed machine learning, and in particular, physics-informed neural networks (PINNs), one tries to incorporate knowledge about the underlying physical laws into various steps, for example, the architecture of the neural network, the loss function or the training data [268–272]. One particularly successful example is to constrain the outputs of the neural network to satisfy a set of differential equations that are known to describe the system, as it was applied for example in the context of fluid dynamics [273, 274]. As a result of this, the networks can often be trained using less data, generalize better to unseen data, are more robust to noise, and most importantly, have an increased interpretability as the solutions are constrained to be physically meaningful. We will also make use of this paradigm in our reconstruction method and try to incorporate prior knowledge about the system.

4.2.2 Supervised vs. unsupervised learning

Traditionally, the input-output pairs used for training a neural network are labeled, i.e., the training dataset consists of input examples with the corresponding true output. This approach is called *supervised learning*, and is the most common form of machine learning. It is simple, typically converges well and is also easy to validate as the true output is known, enabling the direct computation of accuracy metrics. However, there are instances where labeled data is

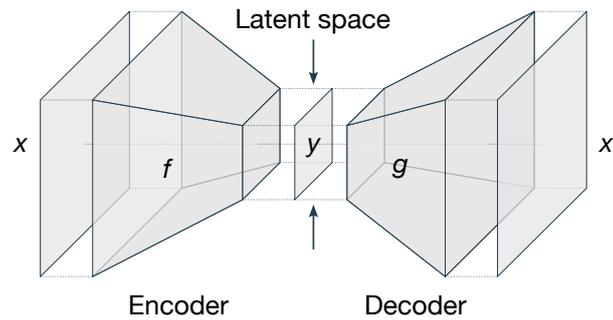


Figure 4.10 | Schematic of an autoencoder. An autoencoder consists of an encoder and a decoder. The encoder maps the input data to a lower-dimensional space, while the decoder maps the lower-dimensional representation back to the original input space. The network is trained to minimize the difference between the input data and the reconstructed data. The dimensionality is here illustrated as the size of 2D data.

not available, or the labeling process is too costly or time-consuming. In such cases, one can resort to *unsupervised learning*, where the network is trained on a dataset without labeled example data. This however changes the process of designing the network architecture and the training process significantly, as we are no longer dealing with a simple 'fit' to example observations, but rather the task of the network is more abstract and implicit. The goal of an architecture trained by unsupervised learning is then typically to find patterns or structure in the data itself. A classical example for this, developed and used way before the advent of neural networks, are clustering algorithms. They constitute a cornerstone of data analysis and are used to group data points into clusters based on some similarity measure, allowing to gain information about and find structures in completely unlabeled data [275]. A second important application of unsupervised learning is in the context of *dimensionality reduction*. Here, the goal is to find a compression of the input data to a lower-dimensional space, learning the salient features that make up the data while discarding irrelevant information [276]. This can be used for example to visualize high-dimensional data, to remove noise or irrelevant features, or as a pre-processing step to reduce the computational complexity of a subsequent algorithm or neural network. As we will see later, also our reconstruction problem can be formulated as a dimensionality reduction problem, making it a natural candidate for unsupervised learning.

To understand why dimensionality reduction can be performed in an unsupervised manner, let us look at the most common neural network architecture for this task: the *autoencoder*. As illustrated in Fig. 4.10, an autoencoder is a symmetrically structured neural network, consisting of an encoder and a decoder subnetwork [277, 278]. The encoder transforms the high-dimensional input data to a lower-dimensional space, while the decoder maps the lower-dimensional representation back to the original input space. The network is trained to minimize the difference between the input data and the reconstructed data, which allows a training entirely without labeled data. If the dimensionality of the intermediary space between encoder and decoder (also called the *latent space*) is chosen to be significantly smaller than the input space, the network is forced to learn to extract the most important features of the input data to allow information flow through this *bottleneck*. After a successful training, the

encoder and decoder parts are typically used separately. The output of the encoder generates a compressed representation of the input data, i.e., it implements the dimensionality reduction. On the other hand, the decoder can be used as a generative model, which can produce new data points by sampling from the latent space (e.g. in the context of variational autoencoders [279]). Note that in the case of a linear encoder and decoder, the autoencoder is equivalent to a principal component analysis (PCA), a standard algorithm for dimensionality reduction [280].

4.2.3 A simple network for single-site reconstruction

We will now present a simple example to illustrate how a neural network can be used to implement a deconvolution for a reconstruction of the lattice occupation from fluorescence images. To this end, we will employ perhaps the simplest neural network architecture possible: a network consisting only of a single input layer and a single output layer. As shown in Fig. 4.11, the input data is a crop of a fluorescence image around a lattice site whose occupation we want to determine. The reasoning behind this is that a reconstruction in the large- β regime is primarily complicated by the spillover of signal from adjacent lattice sites due to the large PSF. Hence, we need to include the information about the signal on neighboring sites to implement the deconvolution. This 2D input crop is flattened to a 1D vector and fed into the input layer, which has as many nodes as there are pixels in the crop. The output layer consists of a single neuron, whose output is a binary value, encoding the reconstructed occupation of the central lattice site in the input crop. In the language of machine learning, this corresponds to a binary classification problem.

Due to the structural simplicity of this network, we can write down the operation that the network implements in a closed form. The single output neuron computes the weighted sum of the input data and applies an activation function to produce the output. Choosing a sigmoid activation function, the output y of the network is hence given by

$$y = \sigma \left(\sum_i w_i x_i + b \right), \quad (4.15)$$

where x_i are the input pixel values, w_i are the connection weights, and b is a bias term. With this relation, we can intuitively understand the action of the network: The sigmoid function, similar to a step function, acts as a binarizer which maps the weighted sum of the input data to a binary occupation value. To perform this binarization for the correct signal values, the bias parameter b sets the reconstruction threshold (also called the decision boundary). Hence, the weights w_i determine the importance of each pixel value for the reconstruction. If the weights are fixed to

$$w_i = \begin{cases} 1 & \text{if pixel } i \text{ is part of the central lattice site,} \\ 0 & \text{otherwise} \end{cases}, \quad (4.16)$$

the network will effectively perform a reconstruction via the pixel count summation method as shown in Fig. 4.4 in the previous section. During the training, these weights will be adjusted to maximize the separation between the signal distributions of empty and occupied

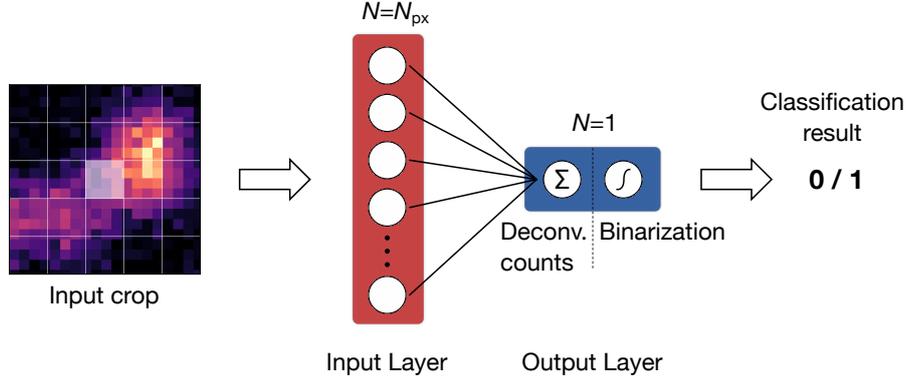


Figure 4.11 | Toy neural network architecture for reconstruction. The most simple neural network for reconstruction consists only of an input layer and an output layer with a single neuron. The input data is a crop of a fluorescence image around a lattice site whose occupation we want to reconstruct (shaded in white). The output neuron is figuratively split into two parts, corresponding to the computation of the weighted sum, and the application of the activation function. The output is a binary value, corresponding to the reconstructed occupation, and the network is trained in a supervised manner.

lattice sites. Hence, the weighted sum of the input data implements the deconvolution operation in this approach.

We train the network in a supervised manner using simulated fluorescence images where the true occupation is known. For the simulation, we use a diffraction limited PSF with a similar β -parameter of $\beta = 2.2$ as in the experiment, as well as a similar SNR (as calibrated in Fig. 4.2), and generate images at all fillings. Note that we neglect density-dependent effects for now. The large images are then cropped symmetrically around one central lattice site with a width of 9×9 lattice sites, and stored together with the true occupation of the central site, forming the training dataset. The network is trained using the Adam optimizer [259] with a binary cross-entropy loss function, which is a common choice for binary classification problems. The binary cross-entropy loss for a set of true labels y and predicted labels \hat{y} is given by

$$\mathcal{L}_{\text{BCE}}(y, \hat{y}) = -\frac{1}{N} \sum_{i=1}^N [y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i)], \quad (4.17)$$

where the sum runs over all N entries. For the training, we use a dataset containing around 4×10^5 crops at all fillings, an initial learning rate of 0.01 (using a learning rate scheduler that reduces the learning rate by a factor of 0.5 when reaching a plateau), and train for 30 epochs. The network converges quickly and reaches an overall accuracy of around 98.5% both on the training data and a separate validation set with about 5×10^4 crops. To better understand what the network does, it is instructive to visualize the weights after training. As the weights are applied to 2D input data, we reshape the flattened weight array to a matrix of the same size as the input crop. The resulting weight matrix is shown in Fig. 4.12a. We observe a very specific pattern that is reminiscent of a radially isotropic $\text{sinc}(x)$ function. Since the weight matrix has to implement the deconvolution, we can immediately understand why this pattern emerges: The central positive peak strongly weights the signal of the site that is

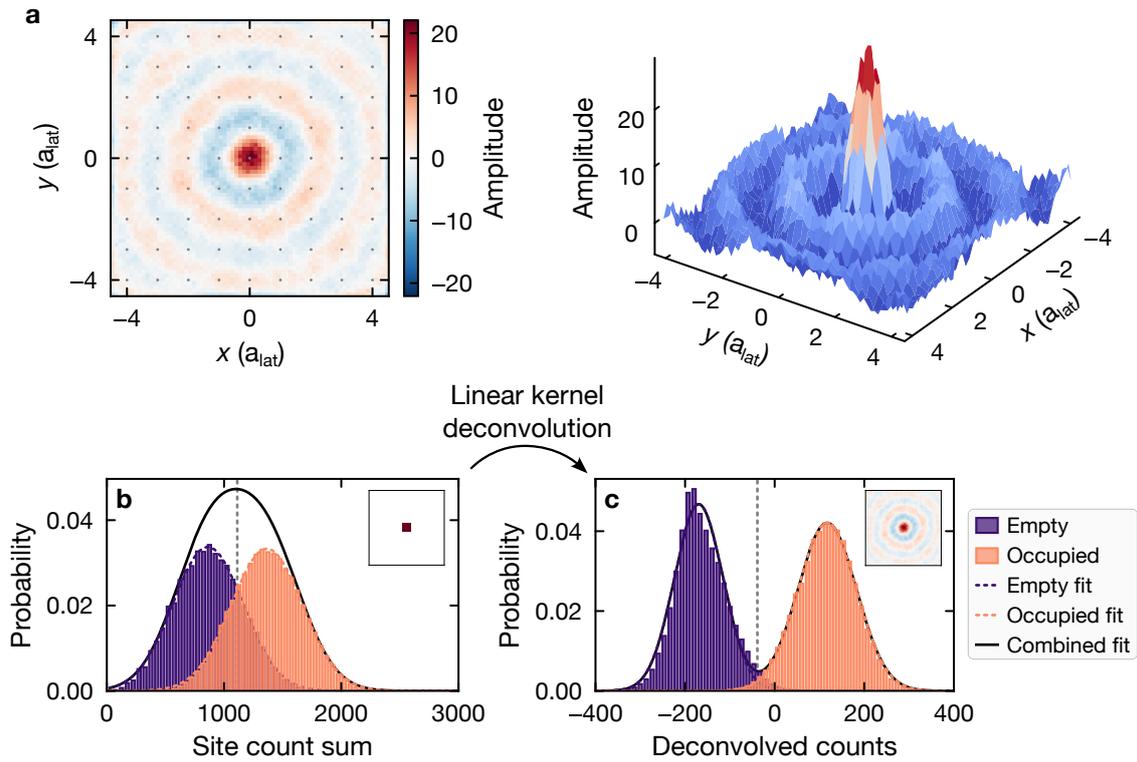


Figure 4.12 | Learned weight matrix and count distributions. **a**, Learned weights of the neural network after training, reshaped to the size of the input crop, illustrated as a 2D image (left) and a 3D projection (right). The weights exhibit a radially isotropic sinc pattern that cancels the signal spillover from neighboring sites. **b**, Site count sum distribution for the validation dataset, corresponding to a weighted sum with the weight matrix from Eq. 4.16, i.e., no applied deconvolution. Note that the histogram is restricted to images with an average filling fraction around 0.5 to obtain normal distributions. **c**, Distribution of the deconvolved counts, i.e., the weighted sum of the input crops with the learned weight matrix shown in (a). The separation between the distributions of empty and occupied sites is significantly enhanced compared to the simple sum weight matrix. For both distributions, the assignment to the empty and occupied classes was done according to the true occupation.

to be classified, while the negative lobes around it subtract the signal from the neighboring sites. Via this negative weighting, the network effectively cancels the spillover of signal from the neighboring sites, and hereby implements the deconvolution. Since the PSF in this regime spans multiple lattice sites, also the deconvolution kernel has a large spatial extent, spreading radially outward in a damped negative-positive alternation.

To further illustrate the applied deconvolution, we can study the distribution of the deconvolved counts, i.e., the weighted sum of the input data $\sum_i w_i x_i$, and compare between the learned weight matrix and the simple sum weight matrix (Eq. 4.16). In the latter case, this is equivalent to the pixel count sum method discussed before, and we obtain a strongly overlapping distribution of the signals from empty and occupied sites (Fig. 4.12b). In contrast, applying the learned weight matrix results in a significantly enhanced separation between the two distributions (Fig. 4.12c). It is particularly noteworthy that we have transformed a distribution which has no discernible separation into one with a clear separation between empty and

occupied sites, which allows for a reliable reconstruction. The following binary classification into the site occupation is performed by the network via the sigmoid activation function, which maps the weighted sum to a binary value. The correct threshold is also learned during the training and encoded in the bias term b . Due to the absence of a separation in the overall signal distribution for the site count sum case, a simple thresholding is not really meaningful. As we have prior information about the true occupation here, we see that the correct threshold would be in the center between the two distributions. In this case, reconstructing via a threshold results in a fidelity of around 80 %, both by comparing true and reconstructed occupation, and via the overlap fraction of the empty and occupied site distributions. In contrast to that, the deconvolution via the learned linear kernel matrix results in a significantly enhanced reconstruction fidelity of around 99 %, which is remarkable in light of the simplicity of this method. Interestingly, the deconvolution kernel learned by the network is equivalent to the one that would be obtained by constructing it from the PSF and the lattice geometry according to Eq. 4.9. This also shows that the network has learned to implement the deconvolution operation in a way that is consistent with the physical process that we want to reverse. Due to the simple architecture, we can fully interpret the network's operation and verify that it is physically reasonable. This will be an important aspect that we need to try to retain when we move to more complex neural network architectures.

One immediate drawback of this feedforward-type architecture is that we can only classify one site at a time, which adds a large computational overhead. As discussed in the previous sections, we can make use of the discrete translational symmetry of the problem and parallelize the deconvolution of multiple lattice sites without increasing the number of parameters by replacing the weighted sum of the feedforward neuron with a convolutional operation. A modified architecture that implements this is shown in Fig. 4.13. As an input, we use a larger subsection of the fluorescence image, here 32×32 lattice sites, which is convoluted with a 9×9 site kernel matrix. The kernel matrix is learned during the training process, and the convolution is performed with a stride equal to the number of pixels per lattice site, i.e., advancing by one lattice site. For each lattice site, the inner product between the kernel and the underlying input image area is computed, which is exactly equal to the weighted sum of the feedforward neuron in the previous architecture. The output of the convolutional layer is then a 24×24 pixel image, where each pixel corresponds to the deconvolved signal of one lattice site². The following steps of the network are the same as before: The output is offset with a global bias, setting the classification threshold, and each entry is passed through a sigmoid activation function to obtain the binary occupation. Training this network on the same dataset as before in a supervised manner, we find that the convolutional architecture learns an approximately identical kernel matrix as the feedforward network, and hence also achieves the same reconstruction fidelity, but with a significantly reduced computational overhead.

While such a linear kernel approach is a great example to illustrate the application of neural networks for deconvolution and reconstruction, its simplicity also limits its performance in an experimental scenario. In particular, the single linear transformation that is applied is not able to account for density-dependent effects, which are significant in the high- β regime. In the

²The output size is smaller since we exclude the outermost 4 sites to avoid edge effects.

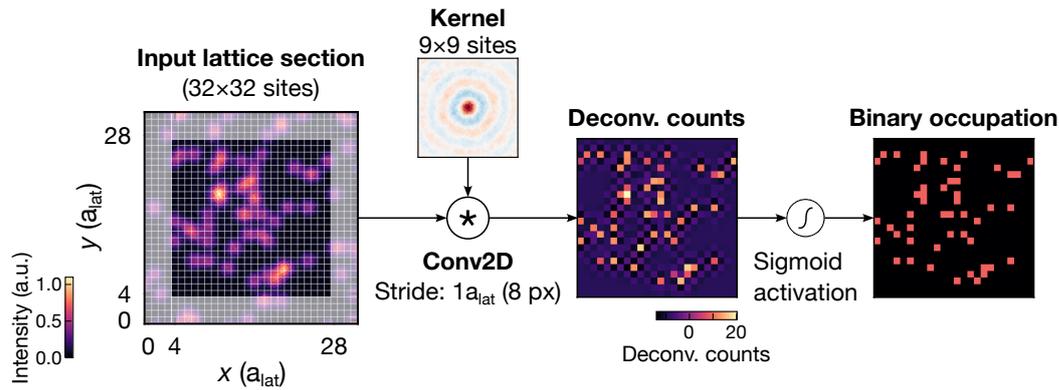


Figure 4.13 | Toy convolutional neural network architecture. To parallelize the deconvolution via a linear kernel matrix to multiple lattice sites, we replace the feedforward neuron with a convolutional layer. An input image (here: 32×32 sites) is convoluted with a 9×9 site kernel matrix, whose entries are learned. The gray shaded area is excluded to avoid edge effects. The stride is chosen to be equal to the number of pixels per lattice site (here: 8), i.e., advancing by one lattice site. As a result, we obtain a 24×24 pixel output image, where each pixel corresponds to the deconvolved signal of one lattice site. The output is then offset with a global bias, setting the classification threshold, and each entry is passed through a sigmoid activation function to obtain the binary occupation.

simulated data used to train this network, we have fully neglected these effects. Furthermore, this architecture can only be trained in a supervised manner, which is in our regime of interest only possible using simulated data. As the simulation captures the experimental conditions only approximately, the performance of the network on experimental data is additionally limited. In the following section, we will present a powerful neural network architecture that is able to overcome these limitations and can be trained in an unsupervised manner directly on experimental data. While the network is significantly more complex, many of the key concepts that we have developed in this toy example will carry over to the new architecture.

4.3 Autoencoder-based approach for single-site reconstruction

In this section, we will present our reconstruction method based on unsupervised neural networks in detail. We will start by introducing the general working principle and the network architecture, followed by a description of the training procedure and the hyperparameter optimization. We will then evaluate the performance of the reconstruction method and compare it to the performance of selected previous techniques. These parts will portray the technique as it was presented for the first time in our publication [281]. In the following part, we will discuss further optimizations that we have developed since then, and give an outlook on possible future improvements.

4.3.1 Idea and network architecture

Let us start by briefly reiterating the central considerations that underlie the choice of a neural network architecture for the reconstruction problem:

1. We want to make use of the ability of neural networks to efficiently approximate complex, non-linear transformations, as is the case for an inverse problem such as deconvolution [245, 246, 282].
2. The network should be trainable in an unsupervised manner, directly on experimental data. This avoids the reliance on a fluorescence simulation to provide pre-labeled training data, which would limit the performance of the network to the accuracy of the simulation (as was done in [83]). Additionally, in our resolution regime, it is not feasible to label the data by hand, nor do any other established reconstruction methods with high fidelity exist.
3. The architecture should allow the incorporation of as much prior knowledge about the system as possible. In particular, it should account for the discrete translational symmetry of the lattice, and that the lattice sites can be constrained to a fixed location in the processed images.
4. We want to be able to interpret the learned transformation and debug the correct operation of the network. One possibility is to provide access to the deconvolved signal distributions, which contain much more information about the reconstruction than a simple array of binary occupation values.

To introduce a suitable architecture for this problem, we start by reformulating the reconstruction task as a problem of dimensionality reduction, which is a classic application of unsupervised learning. In our setting, this corresponds to transforming a high-dimensional input fluorescence image (around 6 pixels per lattice site, each pixel contains an n -bit integer intensity value) to the low-dimensional lattice occupation (1 binary value per lattice site). As previously discussed, a suitable type of architecture for this is the autoencoder, which is shown schematically and adapted for this problem in Fig. 4.14. The task of the encoder subnetwork is to learn the deconvolution and reconstruction from an input fluorescence image to the

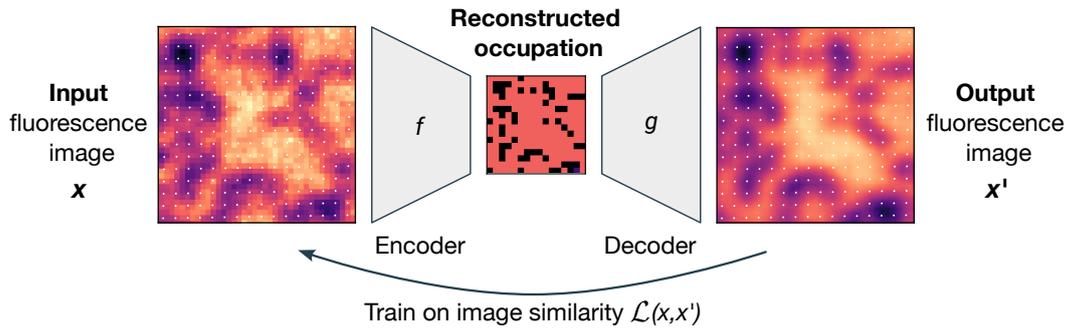


Figure 4.14 | Reconstruction as a dimensionality reduction problem. Cast as a dimensionality reduction problem, we can employ an autoencoder architecture to implement deconvolution and reconstruction with an architecture that can be trained in an unsupervised fashion. The training objective is to make the replicated fluorescence images as similar to the input images as possible, which enables training without labels.

corresponding 2D occupation matrix, which is situated in the latent space. The decoder sub-network then learns to simulate a fluorescence image corresponding to this occupation. The overall network is trained to minimize the difference between the input fluorescence image and the replicated image, which is possible without additional labels and hence enables unsupervised training. Via the linear structure, all information to replicate the input image must flow through the latent space, which forces the network to learn the most important features of the input data. Since the PSF and the lattice geometry as well as the imaging conditions are constant, what is left as the minimal information is indeed the lattice occupation. Our goal is now to shape the network and the training process in a way that this transformation is learned in a performant and robust fashion.

The detailed network topology of our autoencoder is shown in Fig. 4.15. As can be seen, we rely exclusively on convolutional operations to implement the encoder and decoder networks. On the one hand, this makes use of the discrete translational symmetry of the lattice, which reduces the required number of parameters as the optimal transformation should be identical throughout the lattice. On the other, the convolutional operations promote locality and ensure that the spatial order in the data is retained. If one were to use for example feedforward layers, it is likely that the sites in the bottleneck layer follow a different, random ordering than in the original image, which would render the reconstruction result useless.

As an input to the network, we use a subsection of a fluorescence image containing exactly 16×16 lattice sites at a shape of 256×256 pixels. Compared to the recorded images, this corresponds to an oversampling by a factor of ≈ 6.7 . We apply this oversampling such that we are able to cut the image subsections precisely and ensure that the lattice sites are always located at the same positions within the crop to an accuracy better than $\approx 0.1a_{\text{lat}}$. This constitutes important additional information that we implicitly feed in to enhance the performance in such an extreme resolution regime. Additionally, the images are rotated to align the lattice axes approximately with the horizontal and vertical axes of the image, which is important for the correct operation of the convolutional layers that traverse the image in a horizontal and vertical fashion.

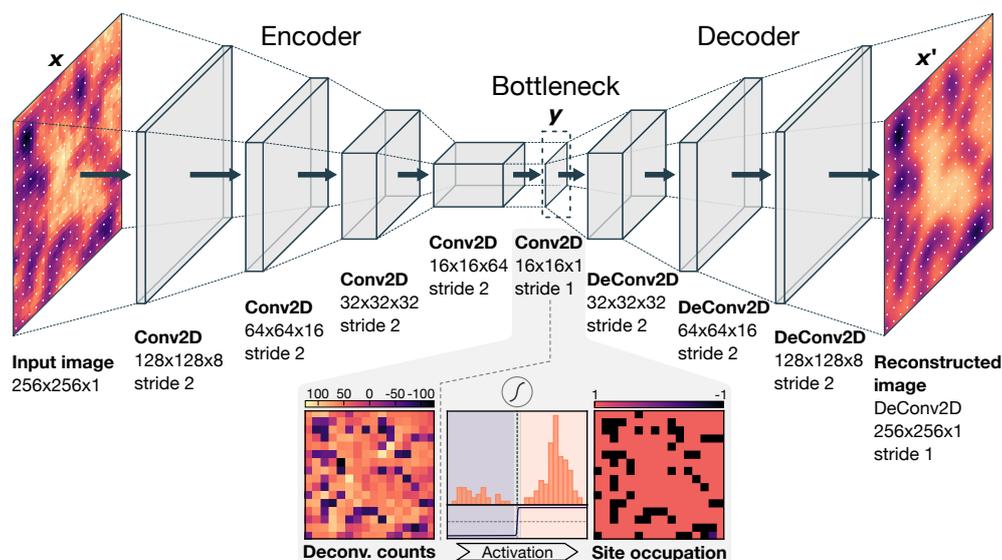


Figure 4.15 | Schematic of the convolutional autoencoder architecture for deconvolution and reconstruction. The autoencoder network consists of a stacked encoder and decoder, each constructed from several convolutional layers. The input image with 256×256 pixels size, cut to contain exactly 16×16 lattices sites, is first processed by the encoder. There, stride-2 convolutions with $\{8, 16, 32, 64\}$ filters subsequently downsample the input image. Each layer is followed by a ReLU activation function. At the end of the encoder, the values of the bottleneck layer are binarized through a tanh activation function to represent the lattice occupation. The bottleneck values before the binarization contain the deconvolved counts of each lattice site, which exhibit a bimodal distribution owing to the saturating effect of the tanh function. The bottleneck is followed by the decoder, which upsamples the occupation matrix in four transposed convolutional layers (stride 2, $\{32, 16, 8, 1\}$ filters) to yield a replicated fluorescence image corresponding to the occupation matrix in the bottleneck.

The input image is first processed by the encoder subnetwork, where it is downsampled by a set of four convolution layers, each with a stride of two pixels, multiple kernels as well as non-linear ReLU activation functions. The stride-2 convolutions reduce the size of the data in each step by a factor of two in both spatial directions. Hence, after the four convolutional layers, the input data has been reduced to a size exactly equal to the number of lattice sites in the input image, i.e., 16×16 pixels, which is an important design choice to promote the correct mapping. Compared to the simple toy network from the previous section, we split the downsampling into multiple layers, which was shown to lead to an improved performance in related applications [283–286]. The notion behind this is that each subsequent layer is able to learn more complex, high-level features by acting on the output of the previous layer. Additionally, the stride of two continuously changes the receptive field of kernels such that each layer is sensitive to correlations on different length scales [287]. By using ReLU activations after each layer, we enable the network to learn non-linear transformations.

The feature maps of the different filters in the last downsampling layer are then combined in a final stride-1 convolutional layer, whose output is a single 16×16 pixel array, forming the bottleneck. The last action in the bottleneck layer is the application of a tanh activation function, whose task is to binarize the output of the decoder to the values $-1, +1$, which we after the training want to interpret with the occupation of the lattice sites. This follows the

same approach as in the toy network: By doing the binarization as the very last step, where the data is already in the correct shape, we obtain access to the deconvolved signal distributions by looking at the values right before applying the activation function (cf. zoom-in below architecture in Fig. 4.15). After the bottleneck layer, the reconstructed occupation matrix is then processed by the decoder subnetwork. It essentially mirrors the encoder, consisting of four convolutional layers with stride-2 transposed convolutions, which subsequently upsample the data back to the original size of 256×256 pixels.

4.3.2 Training procedure and hyperparameters

The symmetric architecture of the autoencoder, producing a replication of an input fluorescence image, enables an unsupervised training directly on experimental data. The training objective is to minimize the difference between the input fluorescence image and the replicated image. To quantify this difference, we use the mean absolute error (MAE) as the loss function, which is defined as

$$\mathcal{L}_{L1}(x, x') = \frac{1}{N} \sum_{i=1}^N |x_i - x'_i|, \quad (4.18)$$

where x is the input image, x' is the replicated image, and the sum runs over all N pixels. If we were to train the network directly on this loss function, there would be no constraint on the values of the bottleneck layer, which in particular could take on arbitrary, continuous values between -1 and 1 . While the network would still learn to replicate the input image, the bottleneck would contain a reduced representation of the input data with properties that are difficult to interpret and not directly related to the lattice occupation. To enforce binary bottleneck values, we additionally augment the loss function by a regularization term that penalizes non-binary bottleneck values. The total composite loss function then reads

$$\mathcal{L}_{\text{total}}(x, x') = \mathcal{L}_{L1}(x, x') + \lambda \frac{1}{N_{\text{sites}}} \sum_{i=1}^{N_{\text{sites}}} (1 - |y_i|), \quad (4.19)$$

where y_i are the values in the bottleneck layer, and the sum runs over all N_{sites} lattice sites. The regularization strength λ determines the relative weight between the two terms and constitutes an important hyperparameter that needs to be optimized. This specific regularization term is just one possible choice to promote binary node values, other possibilities include for example a quadratic penalty term or a cross-entropy loss [288].

To train the network, a dataset containing a sufficient number of sample fluorescence images is required. Besides a minimum size, the dataset has to fulfill a few crucial criteria:

1. The data needs to be representative of all possible experimental conditions that it will later be applied to. In particular, the dataset should contain images at all fillings between zero (dilute clouds) and one (fully occupied lattice, i.e. a unity-filling Mott insulator).
2. The states underlying the images have to be uncorrelated and in absence of any spatial features that could be integrated into the learned transformation during the training.

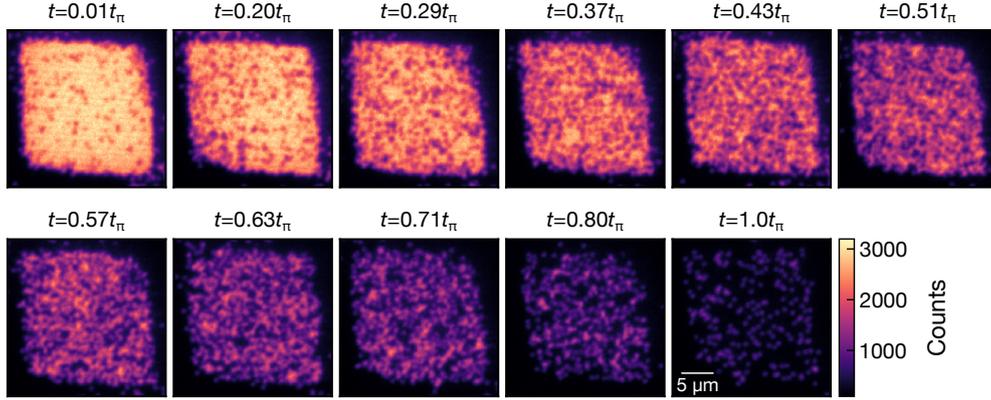


Figure 4.16 | Example fluorescence images from the training dataset. Starting from unity-filling Mott insulators in a 50×50 site box potential, we create samples with varying fillings between 0 and 1 using a combination of microwave and optical blowout pulses. The filling is tuned via the duration t of the microwave pulse, where t_π corresponds to the maximum state transfer (i.e., almost all atoms are removed). Since the atoms are removed randomly from the lattice, the occupation is homogeneous and uncorrelated. Note also that the color scale is the same in all images, which reveals the significantly enhanced brightness for the highest fillings due to density-dependent effects.

To prepare such a dataset, we start from unity-filling Mott insulator states, and then subsequently remove a varying fraction of atoms randomly from the system using a microwave pulse followed by an optical blowout pulse. In detail, the unity-filling Mott insulator (MI) is prepared in a single 2D layer of the shallow-angle vertical lattice, with horizontal confinement being provided by the repulsive DMD box potential (see Section 3.2). This results in a homogeneous occupation of the lattice within the 50×50 site large box at an average filling of up to $n \approx 0.98$. After loading the MI, all atoms are in the absolute hyperfine ground state $|F, m_F\rangle = |3, 3\rangle$. We then apply a microwave pulse resonant with the $|3, 3\rangle \rightarrow |4, 4\rangle$ transition, which induces microwave Rabi oscillations between the two hyperfine states. By varying the pulse duration between 0 and t_π , where t_π is the time corresponding to full state transfer, we can randomly transfer a variable fraction of atoms in the MI to the $|4, 4\rangle$ state. In a last step, the atoms in the $|4, 4\rangle$ state are removed by an optical blowout pulse that is resonant with the $|F = 4\rangle \rightarrow |F' = 5\rangle$ transition on the cesium-D2 line. Via this sequence, we can prepare a dataset with homogeneous samples at various fillings between zero and $n \approx 0.98$ that is free from any spatial correlations. A few example images for various pulse durations are shown in Fig. 4.16, demonstrating the homogeneous variation of the filling via the microwave pulse duration. We sample over 11 different pulse durations and record a total of around 300 images, corresponding to approximately 1.5 hours of measurement time.

The training data is then processed as follows: In a first step, the location of the lattice sites in each image is determined. To this extent, we use the previously determined lattice vectors and extract the phase with respect to the image origin by fitting the positions of a few isolated atoms (see Section 3.1 for details). The images are then rotated such that the lattice vectors are approximately aligned with the horizontal and vertical axes of the image (this is not perfectly possible as the lattices have a relative angle of around 88° , see e.g. lattice

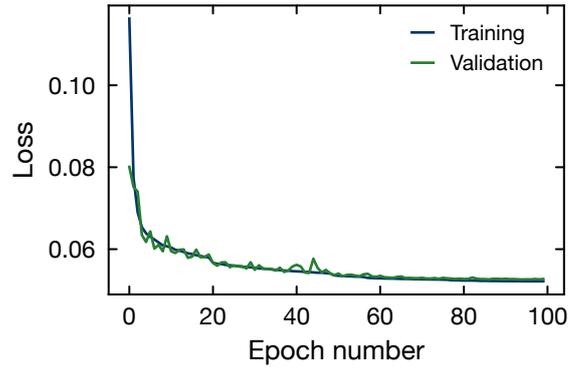


Figure 4.17 | Example training loss curve. Example evolution of the training and validation loss during network training. The final loss is around 0.053 at a regularization strength of $\lambda = 0.4$.

site positions in Fig. 4.14). Finally, the images are cut into subsections containing exactly 16×16 lattice sites, located at the same positions throughout all images. Constraining the atoms to a constant, discrete set of locations is important additional information that helps to overcome the resolution limit. When cutting the images, we allow a maximum overlap between adjacent crops, advancing only by one lattice site in each step. On the one hand, this results in a significantly larger training dataset for the same number of images. On the other, this implicitly encodes the discrete translational symmetry and the notion of fixed, individual lattice sites into the training data, as the data will contain multiple copies of the same lattice configuration that are shifted by an integer number of lattice sites. Via this method, around 300 recorded fluorescence images result in a comprehensive dataset of around 3×10^5 crops.

To train the network, we first randomly shuffle the dataset and then renormalize it using a simple min-max scaling to the range $[-1, 1]$, which is important to ensure convergence during training. The minimum and maximum values for scaling are taken across the whole dataset, and saved for later use to scale new data for the reconstruction process. We then split the dataset into a training (2×10^5 crops) and a validation (1×10^5 crops) set. The training is done using the ADAM optimizer [259]. The initial learning rate is set to 4×10^{-4} , and we additionally use a learning rate scheduler that reduces the learning rate by a factor of 0.5 when reaching a plateau in the validation loss (patience of 3 epochs, minimum learning rate 1×10^{-5}). The network is trained for a maximum of 100 epochs, but we allow early stopping if the validation loss does not improve for 10 epochs. A typical evolution of the training and validation loss is shown in Fig. 4.17. As can be seen, the network converges quickly, but a learning rate scheduling is necessary to obtain a minimum loss. Furthermore, the validation loss is consistently comparable with the training loss, which indicates that the network generalizes well to unseen data. A full training run takes around 1–2 hours on a single NVIDIA A40 GPU (training on a GPU is strictly necessary due to the heavy use of convolutional operations). Note that the network should be trained several times with different random weight initializations to select the best-performing model, which is generally advisable due to the stochastic nature of the training process.

Hyperparameter optimization

The convolutional autoencoder architecture as presented in Fig. 4.15 contains a few hyperparameters that need to be optimized to achieve the best possible reconstruction performance. As hyperparameters are, per definition, not varied during the training process, they need to be optimized manually by re-training the network for different hyperparameter choices and comparing the achieved accuracy. For our architecture, this includes in particular: the learning rate, the training dataset and batch sizes, the number of filters in the convolutional layers, the kernel sizes, and the regularization strength λ .

The learning rate is a crucial hyperparameter that determines whether a good optimum can be found during training. It is fortunately rather easy to fix. If the learning rate is too small, the optimizer will quickly get caught in a local optimum during the early stages of training, likely missing lower-loss optima. Conversely, a too high learning rate causes oscillations around the loss landscape, which can prevent the optimizer from converging. We hence start with a relatively large learning rate on the order of 10^{-2} and then decrease it, picking the largest value where the loss curve is smooth to ensure a fast convergence. In our case, we found that an initial learning rate of $4 - 8 \times 10^{-4}$ is a good choice. It is however also important to use the learning rate scheduler to allow smaller learning rates in the later stages of training, as otherwise a 'climbing down' into the bottom of a minimum is prevented by the large step size.

For the number of kernels and their size, we performed a random search of the hyperparameter space instead of varying all parameters due to the large amount of combinations. The number of filters showed no clear trend, due to which we settled on the initial choice shown in Fig. 4.15. The kernel sizes had a strong effect below some minimal size of around 8×8 , where the final loss was increased. Above that, changes were not significant, and we settled on the absolute optimal choice of 10×10 for the encoder and 22×22 for the decoder.

The regularization strength λ is probably the most critical hyperparameter, as it determines the relative weight between the reconstruction loss and the regularization term. A too small value will not enforce binary bottleneck values, while a too large value will prevent the network from learning the correct transformation. Because it modifies the loss function, it also trivially affects the scale of the loss values, which makes it difficult to optimize based on the final training loss. Instead, we found that the regularization strength should be optimized directly on the deconvolved count distributions. For $\lambda = 0$ or too small values, the deconvolved counts will not exhibit any separation. For increasing λ , a separation will appear with at first heavily skewed distributions, which become more normal distributed for larger values. With even larger λ , the training process will become ineffective again as the replication loss will be neglected. Empirically, we found that the smallest λ with well-separated, normally distributed deconvolved counts is the best choice, which in our case was around $\lambda = 0.4$ using the loss function in Eq. 4.19.

4.3.3 Reconstruction process

After a successful training of the network, only the encoder part is used to reconstruct new fluorescence images. The process for the reconstruction of a full recorded fluorescence image

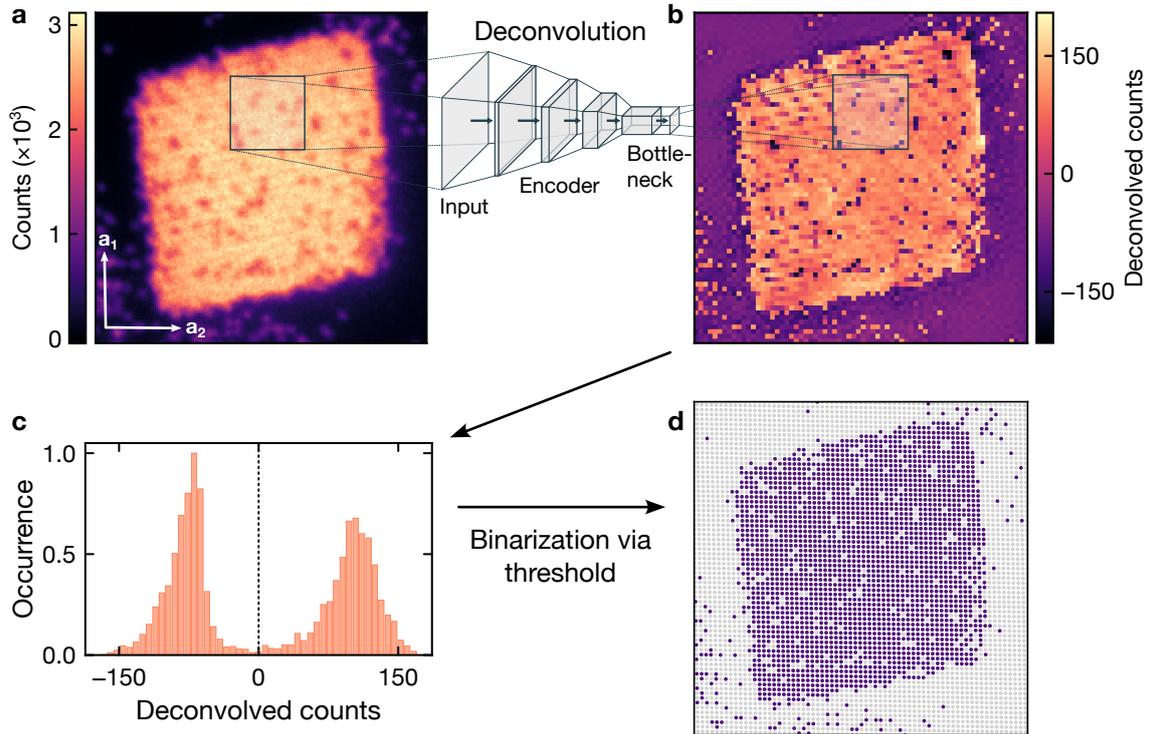


Figure 4.18 | Reconstruction of an experimental image. **a**, Raw fluorescence image of a unity-filling Mott insulator with around 2500 atoms (total image size 70×70 lattice sites, white arrows mark the lattice vectors). Crops of 16×16 lattice sites are cut from the image and fed into the trained encoder for deconvolution. **b**, The deconvolved signals for each lattice site are extracted from the bottleneck layer before binarization and reassembled according to their position in the full image. **c**, A histogram of the deconvolved signals reveals a bimodal distribution with a clear separation between empty and occupied lattices sites, allowing to set a threshold for binarizing the occupation. **d**, Applying the chosen threshold gives the reconstructed lattice occupation corresponding to the input image. Light gray circles denote empty and dark purple occupied sites.

is illustrated in Fig. 4.18, taking a unity-filling Mott insulator with around 2500 atoms as an example. In a first step, the raw fluorescence image is analyzed to determine the phase of the lattice and fix the locations of the lattice sites in the image. The image, here of size 70×70 lattice sites, is then cut into 16×16 lattice site crops, which are fed into the trained encoder (Fig. 4.18a). As for the training data, we advance only by one lattice site between each crop region, such that every site occurs in several reconstructed crops at different locations. The output of the encoder is intercepted before the binarization step, which yields a 16×16 matrix containing the deconvolved counts for each site in the crop. These matrices are then reassembled according to their position in the original image. During the reassembly, we average the deconvolved counts across overlapping crops, which enhances the signal-to-noise ratio after the deconvolution. Additionally, we only take the central 12×12 sites of each crop into account, as the border region suffers from edge effects due to the missing neighboring sites, which are important for the deconvolution. The reassembled array of the deconvolved site counts is shown in Fig. 4.18b, exhibiting a strongly enhanced contrast of individual holes in the high-density plateau of the

Mott insulator and vice versa of isolated atoms in the outside regions, which is a clear signature of the deconvolution. The effect of the deconvolution can also be seen in the histogram of the deconvolved counts, which shows a bimodal distribution with two well-separated peaks (Fig. 4.18c). These two peaks are correspondingly identified with empty and occupied lattice sites, respectively. In the full encoder network, the binarization is done by applying a tanh activation function to the deconvolved counts, whose decision boundary is at zero. A threshold of zero is also consistent with the histogram, and after applying this threshold to the matrix of deconvolved counts, we finally obtain the reconstructed lattice occupation corresponding to our input image (Fig. 4.18d). Note also that it is in principle possible to obtain a confidence metric for the reconstruction result directly from the output of the encoder. As the tanh activation function has a continuous output range between -1 and 1 , the absolute value of the output can be interpreted as a measure of the confidence. Deconvolved counts that fall on the slope of the tanh function have a lower confidence than those in the saturated regions (i.e., those that are farther away from zero), which can be used to exclude reconstruction results with low confidence from further analysis.

4.3.4 Performance evaluation using experimental data

A cross-check by eye of the reconstructed occupation on top of the original image is a useful first indicator if the training was successful and the deconvolution is working correctly. For example, during the first stages of development, we occasionally found that the occupation is globally shifted by one row/column, in particular when the hyperparameters were chosen suboptimally. However, since it is usually not possible to determine the true occupation by eye for all but the lowest fillings in our resolution regime, this does not allow for a quantitative evaluation of the reconstruction fidelity.

As a matter of fact, estimating the reconstruction fidelity directly from experimental data without knowledge of the true occupation is a highly-challenging task. In our resolution regime, there is currently no established method that would provide a reliable fidelity estimate. A benchmark against simulated data can also not alleviate this, as this would require a simulation that describes the experimental data to an accuracy much better than the fidelity that is to be benchmarked, which is in turn highly challenging on its own. This situation is additionally complicated by the fact that there are different categories of errors that can occur during reconstruction. On the one hand, there are statistical errors that arise from the finite SNR of the fluorescence images, which pose a fundamental upper bound on the fidelity that even an ideal reconstruction method could achieve. On the other hand, there are systematic errors that arise from the deconvolution process itself, which can result for example in a bias of the reconstructed occupation or artificial correlations. In particular the latter error category is challenging to quantify without a ground truth. To address this issue, we have developed a set of metrics that allow us to estimate the reconstruction fidelity directly using experimental data. While they do not provide a precise, absolute measure of the fidelity, they can be used to establish a rough estimate of the performance, and in particular, they are well-suited to compare different parameters and different reconstruction methods.

Fidelity estimation using the deconvolved signal distributions

A strong asset of our network architecture is the access to the deconvolved signals on each lattice site, whose distribution contains a large amount of information on the performance of the deconvolution. A high-fidelity reconstruction requires a bimodal distribution of the deconvolved signals, with two well-separated peaks corresponding to empty and occupied lattice sites. The degree of separation between the peaks can be quantified by the relative overlap area, which is defined as the fraction of the area under the curve of the empty sites that overlaps with the area under the curve of the occupied sites. When determining the occupation via a threshold value, the overlap fraction directly corresponds to the fraction of misclassified sites and is hence useful as a measure of the reconstruction fidelity. However, in practice this is based on a few assumptions: An accurate determination of the overlap fraction requires knowledge of the underlying functional form of the distributions. This, in turn, is strongly dependent on the distributions of noise and fluorescence inhomogeneities in the data, as well as the influence of effects such as density-dependent superradiance. Additionally, the non-linear transformation implemented by the network does not necessarily retain the shape of the distributions. We hence can only apply this metric when the observed distributions appear sufficiently well-described by normal distributions.

Since the reconstruction fidelity is expected to be dependent on the filling of the lattice, we evaluate the deconvolved signal distributions for different average fillings. This is shown for four different filling bins between 0 and 1 in Fig. 4.19 using data from the validation set. Across all fillings, we find well-separated bimodal distributions which enable us to perform a high-fidelity reconstruction via a binarization threshold. Comparing this to the completely merged distributions in the pixel count sum distributions of the experimental data (see Section 4.1.5), it is immediately clear that the network was able to learn a powerful deconvolution that enables a reconstruction even in such an extreme resolution regime. Starting from the lowest filling (upper left panel in Fig. 4.19), the two peaks exhibit a vanishing overlap, which is not surprising as the deconvolution is the least challenging for low fillings. As the filling increases, the deconvolution is expected to be more challenging, which is reflected in a slight increase of the overlap fraction. Indeed, one would naively expect the most challenging conditions around half filling due to a maximum possible number of nearest and next-nearest neighboring configurations that have to be distinguished. In the experiment however, the highest fillings are additionally complicated due to the emergence of density-dependent effects such as superradiance, which decrease the visibility of individual holes in high-filling regions. This effect can also be seen in the two lower panels of Fig. 4.19, where the distributions of empty sites are broadened and deviate more from a normal distribution.

Even though we do not have precise knowledge of the functional form underlying these distributions, especially the case of half filling appears to be well-described by a bimodal normal distribution. As shown in the upper right panel of Fig. 4.19, we can therefore fit two Gaussian functions to the distributions of empty and occupied sites, respectively, and use it to determine the overlap fraction. This overlap fraction can then be used as a rough estimate of the reconstruction fidelity by dividing it by the total area under the curves. From the fit, this yields an experimental reconstruction fidelity for half filling of around 99%.

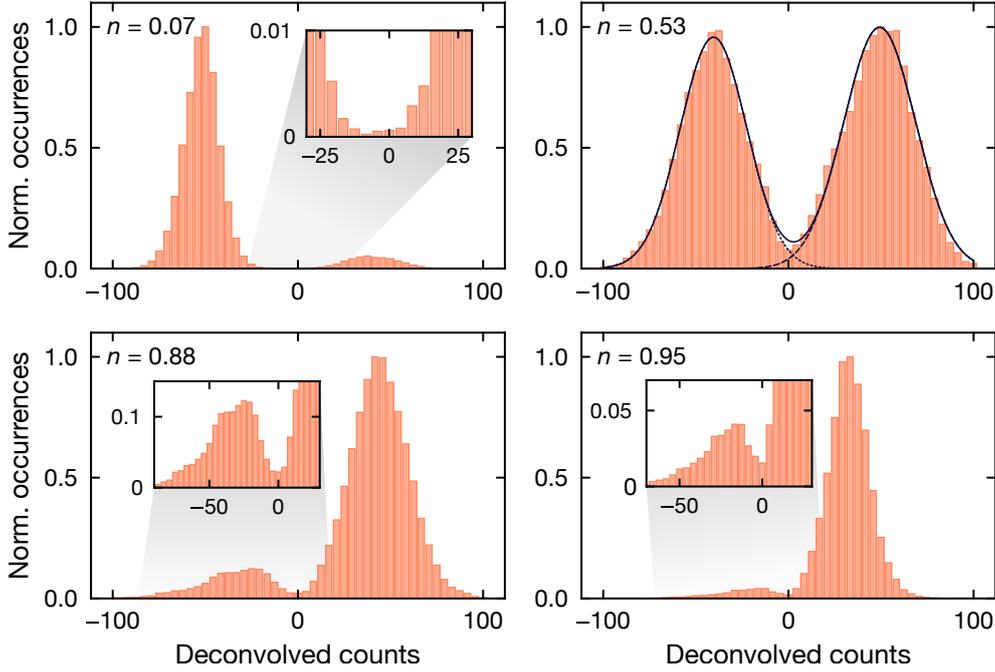


Figure 4.19 | Deconvolved count distributions for different fillings. The deconvolved signals for each lattice site are extracted from the bottleneck layer before binarization. Across all fillings we find a bimodal distribution with two well-separated peaks, corresponding to empty and occupied sites. Their overlap vanishes for small fillings and increases slightly towards higher fillings due to the more challenging deconvolution. The case of half filling is well-described by normal distributions, allowing to quantify the fidelity from the overlap through a fit of two Gaussians (solid line), capturing the distributions of empty (dotted line, left peak) and occupied (dashed line, right peak) sites, respectively. The insets show a zoom-in of the overlap region. The filling values were obtained directly from the reconstruction using a threshold of zero, and the occurrences are normalized to the respective maximum in each plot. Each histogram is computed from ~ 1500 crops with a size of 12×12 sites.

As mentioned before, the absence of the ground truth for the experimental data as well as the lack of knowledge of the precise functional form of the distributions precludes a precise, absolute measure of the fidelity. However, an evaluation of spatial correlations in the reconstructed occupation (see below) as well as a comparison with simulated data (see Section 4.3.5) suggests that we do not expect a large amount of systematic errors, which would lead to a skewing of the distributions in the overlap area and hence to an overestimate of the fidelity when assuming normal distributions. Furthermore, in light of the strong separation between the peaks, the true reconstruction fidelity cannot deviate significantly from this estimate, which especially for the challenging case of half filling is a strong indicator for a good performance of the network.

Fidelity estimation using multiple exposures

To obtain a second, independent estimate of the reconstruction fidelity, we devise a method based on a repeated exposure of the same atomic sample within a single experimental realization. Such a multiple exposure protocol has previously been used with dilute samples or when assuming a perfect reconstruction to quantify the probability of losses and hopping

events during the imaging process [31, 247–249]. However, here we show that this can also be used to estimate the reconstruction fidelity in the presence of hopping and loss processes during imaging. The method is based on the following idea: We prepare an atomic sample at a certain filling, and then subsequently take two exposures at the same imaging parameters. Both exposures are then independently reconstructed using the trained network, and we compare the reconstructed occupations. Whenever there is a discrepancy between the two reconstruction results, this can be either (i) due to a loss or hopping event that happened between the two images, or (ii) due to an incorrect reconstruction result that changes based on statistical fluctuations between the two exposures. By modeling this three-step process of reconstruction of the first image, loss or hopping event between images, and reconstruction of the second image, as a Bernoulli trial with the corresponding probabilities for each case, we can derive an estimate of the reconstruction fidelity as

$$\mathcal{F} = \frac{1}{2} \left(1 + \sqrt{\frac{1 - 2\delta}{1 - 2p_\delta(n)}} \right). \quad (4.20)$$

Here, δ is the difference in the reconstructed occupations between the two images, and $p_\delta(n)$ is the filling-dependent probability of a loss or hopping event between the two images. In absence of hopping and loss processes, the relation simplifies to

$$\mathcal{F} = \frac{1}{2} \left(1 + \sqrt{1 - 2\delta} \right). \quad (4.21)$$

If no difference between the first and the second image is detected ($\delta = 0$), the formula correctly predicts a perfect reconstruction fidelity of $\mathcal{F} = 1$. In contrast, when measuring the maximum possible difference $\delta = 0.5$ (equivalent to comparing two random images), Eq. 4.21 yields $\mathcal{F} = 0.5$, the lowest possible value for the reconstruction fidelity. A detailed derivation of this formula is presented in Refs. [213, 281]. The rate of hopping and loss processes can be independently calibrated using repeated exposures of very dilute samples with isolated atoms, where a reconstruction with a fidelity close to one is possible. As discussed in Section 3.1, we calibrate this rate to be $p_\delta(n) = n \cdot 5 \times 10^{-3}$ per exposure with a duration of 300 ms. Eq. 4.20 then enables an estimation of the reconstruction fidelity directly from the average difference δ in the reconstruction results of the two subsequent exposures. In Fig. 4.20, we show this estimation as a function of the average filling, both with and without the hopping and loss correction. The reconstruction fidelity shows a characteristic dependency on the filling that is in line with our expectations: For the lowest fillings ($n \lesssim 0.2$), we find a reconstruction fidelity of nearly 100%, which is expected as the deconvolution is least challenging in this regime. The fidelity then decreases towards higher fillings, with a minimum occurring around $n = 0.7$ at an uncorrected fidelity of 96.3(3)%. The experimentally observed minimum is shifted towards higher fillings than the naive expectation of 0.5 possibly due to the onset of density-dependent effects. While for even higher fillings, the fidelity increases again, density-dependent effects limit the highest achievable fidelity to slightly above 99%. The values that are corrected for hopping and loss processes are consistently higher than the uncorrected values, even though the difference is below 0.5% for all fillings. Nevertheless, these processes do occur in the images,

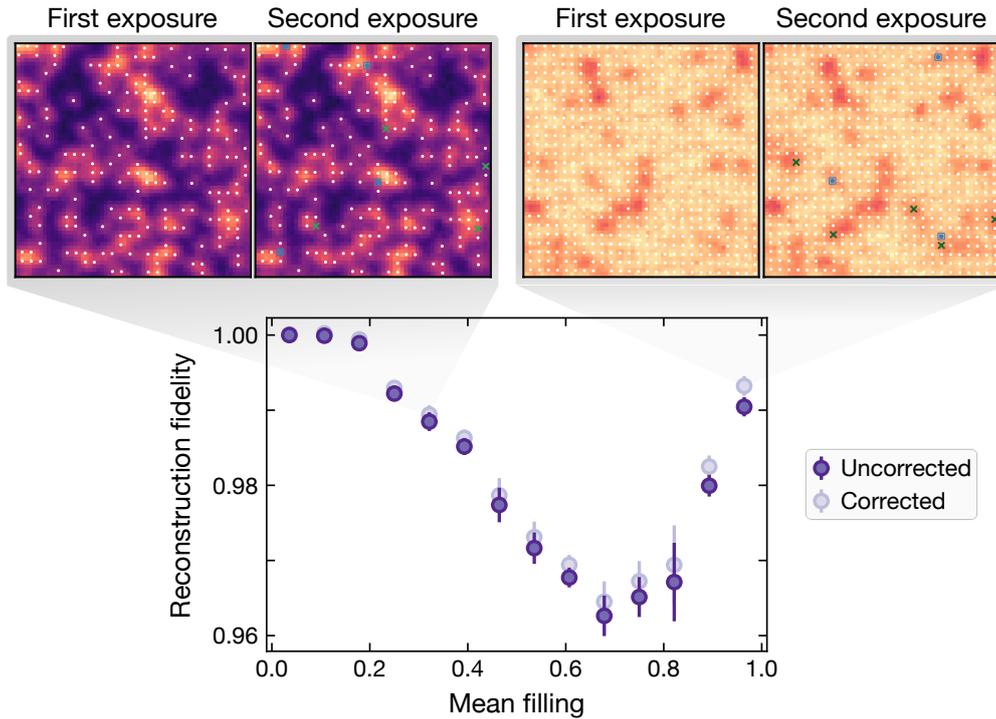


Figure 4.20 | Estimated fidelity using multiple exposures. We estimate the reconstruction fidelity using a multiple exposure protocol, where we take two subsequent images of the same atomic sample. The data points show the estimated fidelity via Eq. 4.20 as a function of the average filling. The uncorrected data points are obtained by neglecting hopping and loss, i.e., $p_\delta = 0$, while for the corrected points (shaded) we take the independently calibrated value $p_\delta(n) = n \cdot 5 \times 10^{-3}$ into account. The zoom-in panels show two typical images at the indicated mean filling with the respective first and second exposure. The blue squares in the second image indicate sites that changed from unoccupied to occupied, and the green crosses vice versa. The error bars denote the standard error of the mean.

as can be seen for example in the two exposures for the highest filling bin in Fig. 4.20. Here, there are a few instances where atoms are lost or hop between the two exposures, which appear as a hole in the high-filling plateau. This is correctly captured by our reconstruction, and if not corrected, would lead to an underestimation of the reconstruction fidelity.

While this analysis provides a quantitative estimate of the reconstruction fidelity for all possible fillings, it is again mostly sensitive to reconstruction errors stemming from statistical fluctuations due to the finite SNR. Any systematic errors arising in the deconvolution process itself would likely result in the same result in both exposures and hence would not be captured by this method. Assuming that the reconstruction fidelity is indeed limited by statistical errors, the two methods of fidelity estimation consistently estimate the reconstruction fidelity to be well above 95% for all fillings, which is comparable to fidelities of other state-of-the-art experiments with significantly better resolution [31, 242].

Analysis of density correlations

Since the previous methods are mostly sensitive to statistical errors, we additionally develop a method to estimate the presence of systematic errors in the reconstruction. Systematic errors are expected to show up as pattern-dependent biases towards certain reconstruction results. An easy example for this is an uncorrected signal spillover due to a large PSF, which would lead to a systematic overestimation of the occupation in the vicinity of occupied sites. As we lack an independent ground truth, it is challenging to investigate this directly by looking at the reconstruction of specific patterns. However, systematic errors are also expected to show up as artificial correlations in the reconstructed occupation. To investigate this systematically, we analyze data where the occupation is expected to be uncorrelated, as is the case for the training and validation data. If we find correlations in the reconstructed occupation in these datasets, this can be seen as an indicator for the presence of systematic reconstruction errors. To quantify correlations in the occupation, we compute the 1D connected density-density correlation function, defined as

$$C_d^{(c)} = \langle \hat{n}_i \hat{n}_{i+d} \rangle - \langle \hat{n}_i \rangle \langle \hat{n}_{i+d} \rangle, \quad (4.22)$$

where \hat{n}_i is the occupation of site i , d is the distance along one of the axes, and the brackets denote an average over all sites i and all crops in the dataset. The connected correlator $C_d^{(c)}$ is a measure of the correlation between the occupation of two sites separated by a distance d . If the occupation is uncorrelated, the connected correlator should vanish for all distances d . In contrast, if there are systematic errors in the reconstruction, we expect to find non-zero values of the connected correlator, which can be used to estimate the strength of the systematic errors.

Artificial correlations can arise due to the deconvolution, and as such, are expected to depend on the PSF and the local filling of the images, which influence the fidelity of the deconvolution. Since the PSF is not isotropic, we evaluate the correlations as a function of distance along the two lattice axes, and furthermore also as a function of the average filling. The result is shown in Fig. 4.21, evaluated on reconstructed snapshots of the validation dataset. Since the data is prepared by randomly removing atoms from unity-filling Mott insulators, we expect there to be no density correlations. The autocorrelator ($d = 0$) is simply given by the variance of the occupation, and hence trivially non-zero. All other distances should be uncorrelated, however, there appears to be a small positive correlation for $d = 1$ and a small negative correlation for $d = 2$, on the level of 2 – 3% in both axes. This is consistent across all fillings, but the influence seems the strongest around half filling. Larger distances show no significant correlations, besides some scatter due to the finite statistics on the level of 1%. The presence of artificial correlations at short distances is an indicator for systematic errors in the deconvolution process. While their strength is relatively small, this is an important limitation that needs to be taken into account when one wants to measure very small correlation signals. In this case, the artificial correlations need to be calibrated and corrected for, as was done in Ref. [226].

The precise origin of artificial density correlations after deconvolution and reconstruction is not well-understood. Their shape, alternating between positive and negative correlations, and the decay over a distance of two sites (i.e., the same extent as the PSF), however suggest

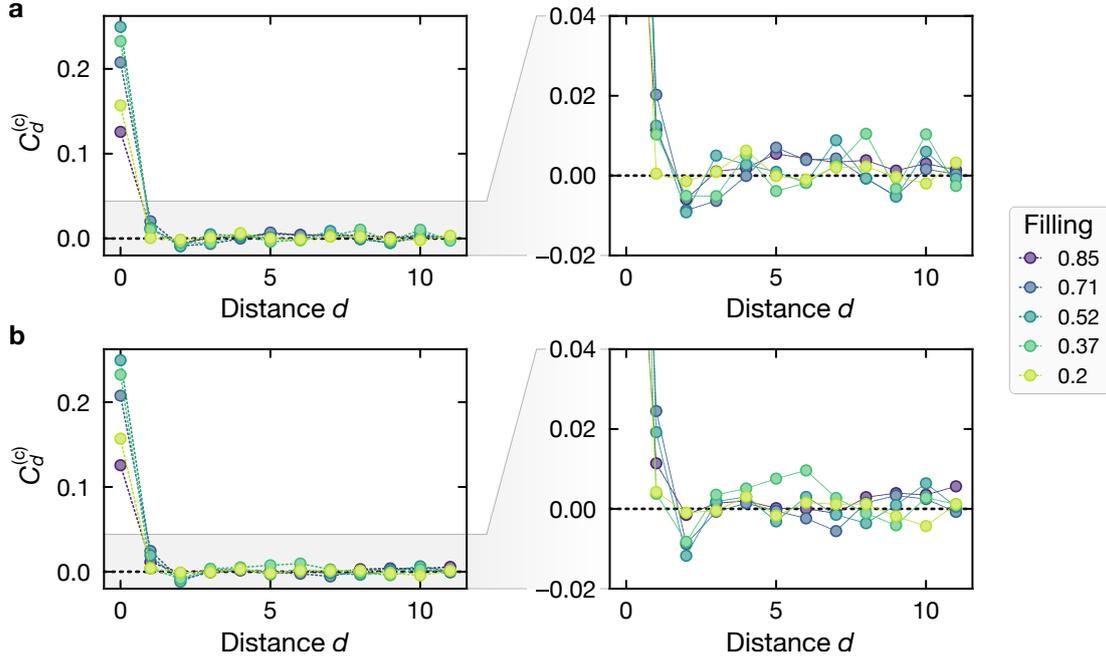


Figure 4.21 | Density-density correlations in the reconstructed occupation. For the validation dataset, we compute the connected density-density correlation function $C_d^{(c)}$ as a function of distance d along lattice axis 1 (a) and lattice axis 2 (b) for different fillings. The right panels show zoom-in plots for the correlator range close to 0. The data is assumed to be uncorrelated, but we find artificial correlations on the 2%-level for $d = 1$ and $d = 2$, which could arise from systematic errors in the reconstruction. The dashed lines are guides to the eye.

that they are related to an imperfect deconvolution itself. When the spillover is not exactly canceled, which is in general expected due to the finite SNR, the remaining dependence on the occupation of neighboring sites is likely to lead to such artificial correlations. This argument is further supported by the fact that other deconvolution approaches show similar correlation artifacts (see Section 4.1.5), where the strength of the artificial correlations seems to decrease with a higher fidelity. Lastly, as we will show in Section 4.3.7, a further optimization of the PSF, where we in particular removed the long-range tail, allowed us to completely eliminate artificial correlations of this kind.

4.3.5 Performance evaluation using simulated data

While it is not straightforward to compare between experimental and simulated data due to the limited precision of the simulation, an analysis of our architecture on simulated data can nevertheless provide valuable insights. Most importantly, the ground truth for the occupation is known, which allows us to quantify the reconstruction fidelity accurately and to investigate the presence of systematic errors. To implement a benchmark with simulated data, we try to mimic the experimental conditions using a basic fluorescence simulation. Specifically, we simulate fluorescence images at variable filling fractions with the experimentally calibrated lattice vectors, the PSF, and a similar SNR as calibrated in Sections 3.1 and 4.1.3. The simulated

images are then processed in the same way as the experimental data, including a rotation to align the lattice axes with the image axes, cutting into 16×16 lattice site crops, and rescaling to the range $[-1, 1]$. Using a training dataset consisting of around 7×10^4 crops, we train the network in an unsupervised manner as described in Section 4.3.2 directly on the simulated fluorescence images, not making use of the occupation labels yet.

To evaluate the deconvolution and reconstruction performance, we apply the encoder part of the trained network to a previously unseen validation set. As we have knowledge of the true underlying occupation, we can immediately evaluate the reconstruction fidelity by comparing the predicted to the true occupation. Specifically, we define the fidelity as the fraction of correctly reconstructed sites, i.e., $\mathcal{F} = N_{\text{correct}}/N_{\text{total}}$. This can be done for all sites in the image, as well as separately for the detection of atoms and holes, where the atom (hole) fidelity is accordingly defined as the fraction of correctly reconstructed occupied (empty) sites to the total number of occupied (empty) sites. These separate metrics are particularly useful in the regimes of minimum or maximum filling. For example close to unity filling, the challenging aspect is the correct detection of individual holes among a filled background. The overall fidelity would not account for this, as it is mostly determined by the contribution of the many filled sites, which are much easier to detect. The results are shown in Fig. 4.22a as a function of the average filling, split into the overall fidelity, the atom fidelity, and the hole fidelity. Remarkably, we find reconstruction fidelities well above 99% for all fillings, even though the simulation uses the same high β -parameter of 2.2 as in the experiment, which is a further proof of the high performance of the unsupervised autoencoder approach. The overall reconstruction fidelity exhibits the expected behavior, with a minimum around half filling. Interestingly, the hole fidelity monotonically decreases towards higher fillings. This indicates that it is indeed most difficult, even for simulated data, to identify individual holes in high-density regions. Compared to the experimentally estimated fidelities, the simulated data shows a significantly higher fidelity. The origin for this is most likely the absence of density-dependent effects as well as loss processes in the simulation, which further complicate the reconstruction in the experimental setting.

A second interesting avenue for benchmarking with simulated data is the analysis of the deconvolved count distributions. As for the experimental data, we extract the bottleneck values before binarization and analyze their distribution for different fillings. The result is shown in Fig. 4.22b, where we – in agreement with the fidelity evaluation – find well-separated bimodal distributions with a vanishing overlap for all fillings. A crucial advantage compared to the experimental evaluation is that the true underlying occupation allows us to sort the deconvolved counts into empty and occupied sites, which reveals the individual distributions. In Fig. 4.22b, this is indicated by the two different colors, corresponding to empty (purple) and occupied (orange) sites, respectively. As shown in the inset plots, the tails of the individual distributions indeed decay quickly, which confirms that the overlap fraction can be used as an estimate of the fidelity. However, the tails are not perfectly Gaussian, as is particularly visible in the case of $n = 0.9$, where the distribution of empty sites is slightly skewed. This leads to an overestimation of the fidelity when assuming normal distributions, which is a limitation and renders this method only usable as a rough estimator for the reconstruction fidelity.

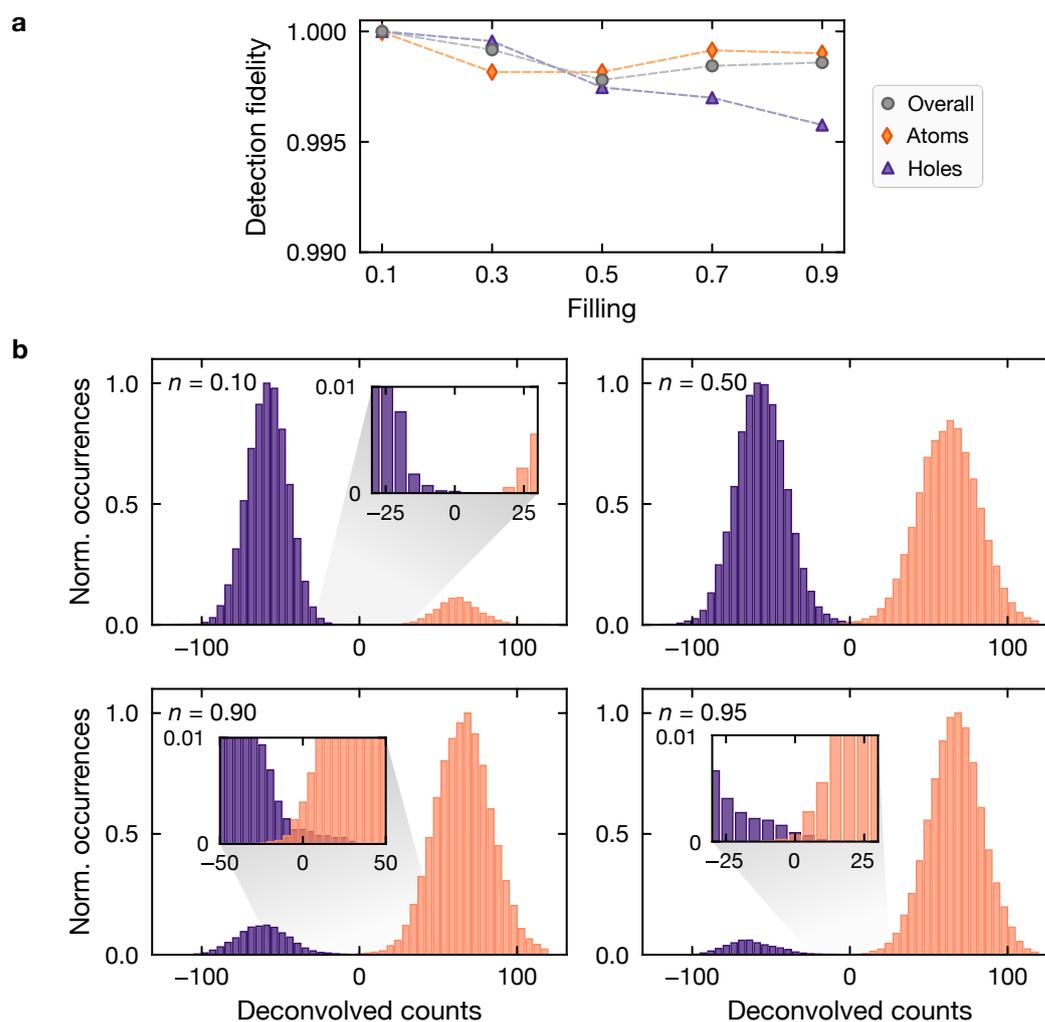


Figure 4.22 | Evaluation of the autoencoder on simulated data. **a**, Reconstruction fidelity as a function of the average filling, separate for all sites (overall fidelity), as well as the fidelity for the correct detection of atoms and holes. The fidelity is evaluated by comparing the predicted with the true occupation. **b**, Deconvolved count distributions for four different fillings. The deconvolved counts are extracted from the bottleneck values before binarization. Based on the true occupation labels, the deconvolved counts are sorted into their individual distributions for empty (purple) and occupied (orange) sites. The inset plots show a zoom-in into the overlap regions to highlight the tails of the distributions. The occurrences are normalized to the respective maximum in each plot.

4.3.6 Interpretation and visualization of the learned transformation

A crucial challenge when employing artificial neural networks for data analysis is to verify that the learned transformation is physically reasonable and captures the relevant features of the data and the problem. This is particularly difficult when dealing with highly-parametrized, deep architectures as is the case for our convolutional autoencoder, where it is not straightforward to write down the applied transformation in a closed form. To address this issue, we develop a set of methods that allow us to visualize and analyze the learned transformation, and to interpret the network in terms of the physical processes that are captured by the deconvolution.

Analysis of the decoder

We start by analyzing the transformation learned in the decoder part of the network. The decoder is typically not needed for the reconstruction process, as its task is to transform the occupation matrix in the bottleneck layer to a corresponding fluorescence image. It hence effectively performs a fluorescence simulation that closely replicates the imaging process underlying the experimental data. This poses the advantage that we can feed in tailored occupation matrices and analyze the output of the decoder systematically. While the decoder is not explicitly constrained to perform the inverse action of the encoder, the symmetric layout of the architecture implicitly results in this, which allows us to transfer our findings also to the encoder.

In Fig. 4.23, we study the transformation learned by the decoder by feeding in occupation matrices with specific patterns and analyzing the output. One interesting test is to set a single site to be occupied, which should result in an image of a single PSF at the corresponding lattice site. As shown in Fig. 4.23a, this is indeed the result, and the learned PSF qualitatively resembles the experimentally measured PSF. We can also repeat this test for any other lattice site, which consistently yields a PSF at the correct location. This is a strong indicator that one of the fundamental aspects of the problem, the assignment of an occupied site to a PSF at the respective location, is correctly encoded in the network after the unsupervised training. To investigate the learned PSF quantitatively, we show a cut through the center of the learned and measured PSF in Fig. 4.23b. The shape, width, and amplitude of the learned PSF are in excellent agreement with the independently measured PSF, which is remarkable in light of the fact that no a-priori information on the PSF was provided to the network.

Furthermore, we can analyze the decoder to verify whether the network was able to approximate density-dependent effects in the experimental data, which lead to an increased fluorescence signal for higher local filling fractions. To study this, we generate an occupation matrix with a block of 9×9 occupied sites and process it with the decoder. As shown in Fig. 4.23c, we obtain an output image that is significantly brighter than a corresponding linear superposition of measured PSFs on the occupied sites (i.e., a simple convolution of the occupation matrix with the measured PSF). This is a first qualitative indication that density-dependent effects are captured in the experimental data. Also, it is noteworthy that the decoder output features a pronounced tail towards the bottom right corner, which is a direct consequence of the interplay between long-ranged aberrations in the PSF and the density-dependent effects. To quantify this more systematically, we generate randomly occupied 16×16 site occupation matrices at different mean filling fractions and process them using the decoder as well as using a convolution with the measured PSF. As a metric for the brightness, we evaluate the average signal in the central 12×12 sites. This is then compared to the average fluorescence signal in experimental images with homogeneous fillings from the validation set. As shown in Fig. 4.23d, the decoder output shows a significantly higher brightness than the linear superposition of measured PSFs. The predicted brightness is in excellent agreement with the experimental data, which confirms that the network was indeed able to approximate the density-dependent effects in the experimental data. Given our situation where the interparticle spacing (383.5 nm) is significantly smaller than the emission wavelength (852 nm), cooperative effects such as superradiance are expected to play a significant role [88, 89], which could explain the about 22%

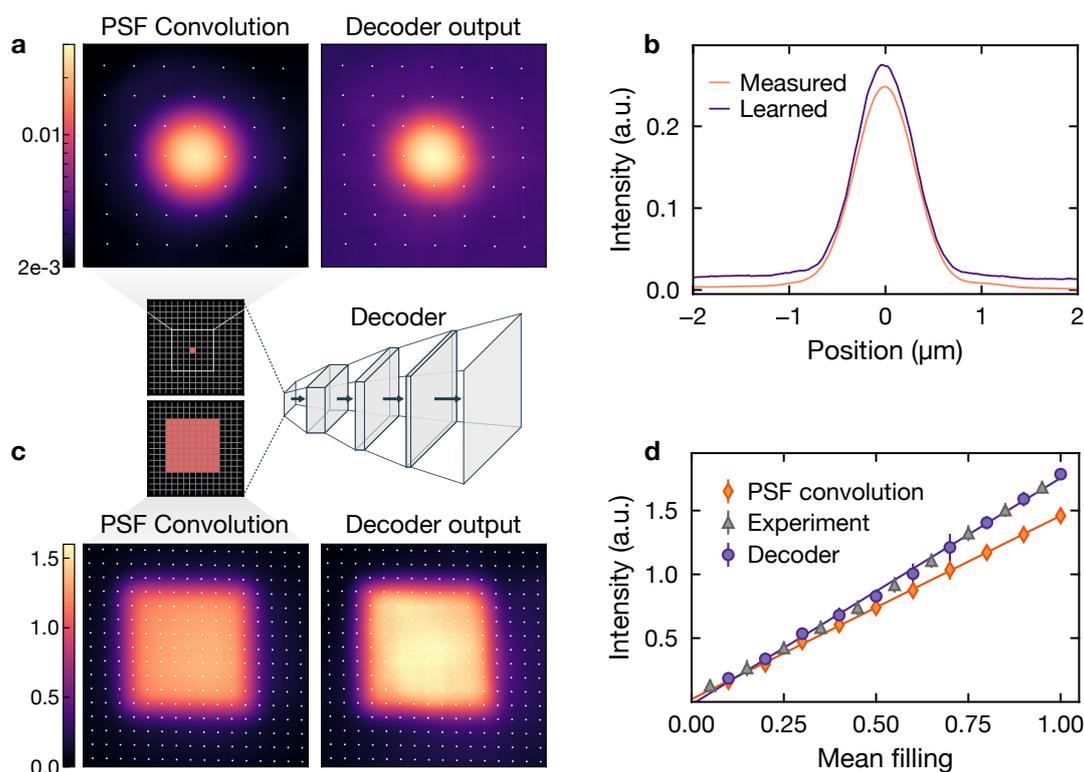


Figure 4.23 | Analysis of the decoder. By isolating the decoder part and feeding in specific occupation matrices, we can visualize the transformation learned by the network. **a**, If we set a single site to be occupied, we obtain an image of a single PSF at the corresponding lattice site (right). The learned PSF looks similar to the experimentally measured PSF (left). **b**, Horizontal cut through the center of the learned and measured PSF, showing a good agreement in shape, width and amplitude. **c**, Analysis of a group of 9×9 occupied sites. The prediction from a simple linear superposition of measured PSFs (left) is significantly less bright than the decoder output (right). **d**, Systematic comparison of the average brightness as a function of the mean filling between linear superposition of measured PSFs (orange diamonds), experiment (gray triangles) and the decoder output (purple circles). The experimental data is taken from the validation set with homogeneous fillings (~ 1500 crops per filling, 12×12 sites). For the decoder and the linear superposition, we generate randomly filled occupation matrices at a given mean filling and evaluate the brightness in the central 12×12 sites. The lines denote linear fits through the respective data points, and the error bars are the standard deviation. At unity filling, the decoder correctly predicts a 22% higher brightness than the linear superposition, capturing the density-dependent effects in the experimental data.

higher brightness of the experimental data as compared to a linear superposition of individual atom signals. The fact that the network is able to capture these effects is crucial to enable a high-fidelity reconstruction at high filling fractions. Linear approaches, such as the linear kernel deconvolution, are not able to account for this effect, which limits the fidelity in the high-filling regime (see Section 4.1.5) and highlights the importance of non-linear transformations.

Visualization of the encoder transformation

In a next step, we investigate the transformation learned by the encoder. As we do not have access to experimental fluorescence images with a known occupation, a systematic analysis as we did for the decoder is not straightforward. In fact, understanding and interpreting the transformations learned by neural networks is an active field of research [289]. For convolutional networks, a common approach is the visualization of the learned kernels as well as the intermediate output data of the different convolutional layers in form of 2D images.

We study this by tracking the evolution of a section of an experimental fluorescence image through the encoder. As shown in Fig. 4.24a, the overall output of the encoder, extracted from the bottleneck layer before binarization, corresponds to the deconvolved counts on each lattice site in the input image. For the relatively dilute input image, the strongly enhanced contrast is particularly visible, which is a direct signature of the implemented deconvolution. Fig. 4.24b shows a visualization of a selection of the different kernels and feature maps throughout the network. The kernels of the first layer (1) operate directly on the input image and show smooth structures that are reminiscent of the PSF. The output of this layer does not show distinct features yet, and could serve for example as a background subtraction or a segmentation step that marks the occupied regions. As we progress through the network, the outputs of the subsequent structures show the extraction of a combination of both large-scale and also increasingly small-scale features. The kernels of the last layer (5) are particularly interesting, as they exhibit structures where one central entry has a strong weight while the neighbors are oppositely weighted. This is very similar to the linear kernel deconvolution, where the negative weights are used to cancel the spillover of signal from neighboring sites (see Section 4.2.3 and Fig. 4.12). On the one hand, this serves as a further confirmation for a physically reasonable deconvolution. On the other hand, an interpretation of the working principle behind the network could be that the earlier layers are responsible for pre-processing and possibly the counting of the neighboring signal strength to capture density-dependent effects, while the later layers implement the actual deconvolution. Lastly, such a visualization can also be useful for optimizing the network architecture. For example, one can see that a significant number of feature maps and kernels are suppressed, which could be an indicator that the number of filters can be reduced, in turn decreasing the computational complexity of the network.

4.3.7 Optimizations and further development

At the time of writing this thesis, the autoencoder-based reconstruction has been in use in our experiment for more than two years and enabled a wide range of quantum simulation experiments [124, 226, 290, 291]. During this time, we have further optimized the network architecture and also the fluorescence imaging itself, which has allowed us to improve the reconstruction fidelity further and to eliminate systematic errors. In the following, we will discuss the changes and present their effect on the reconstruction performance, which reflects the current state-of-the-art.

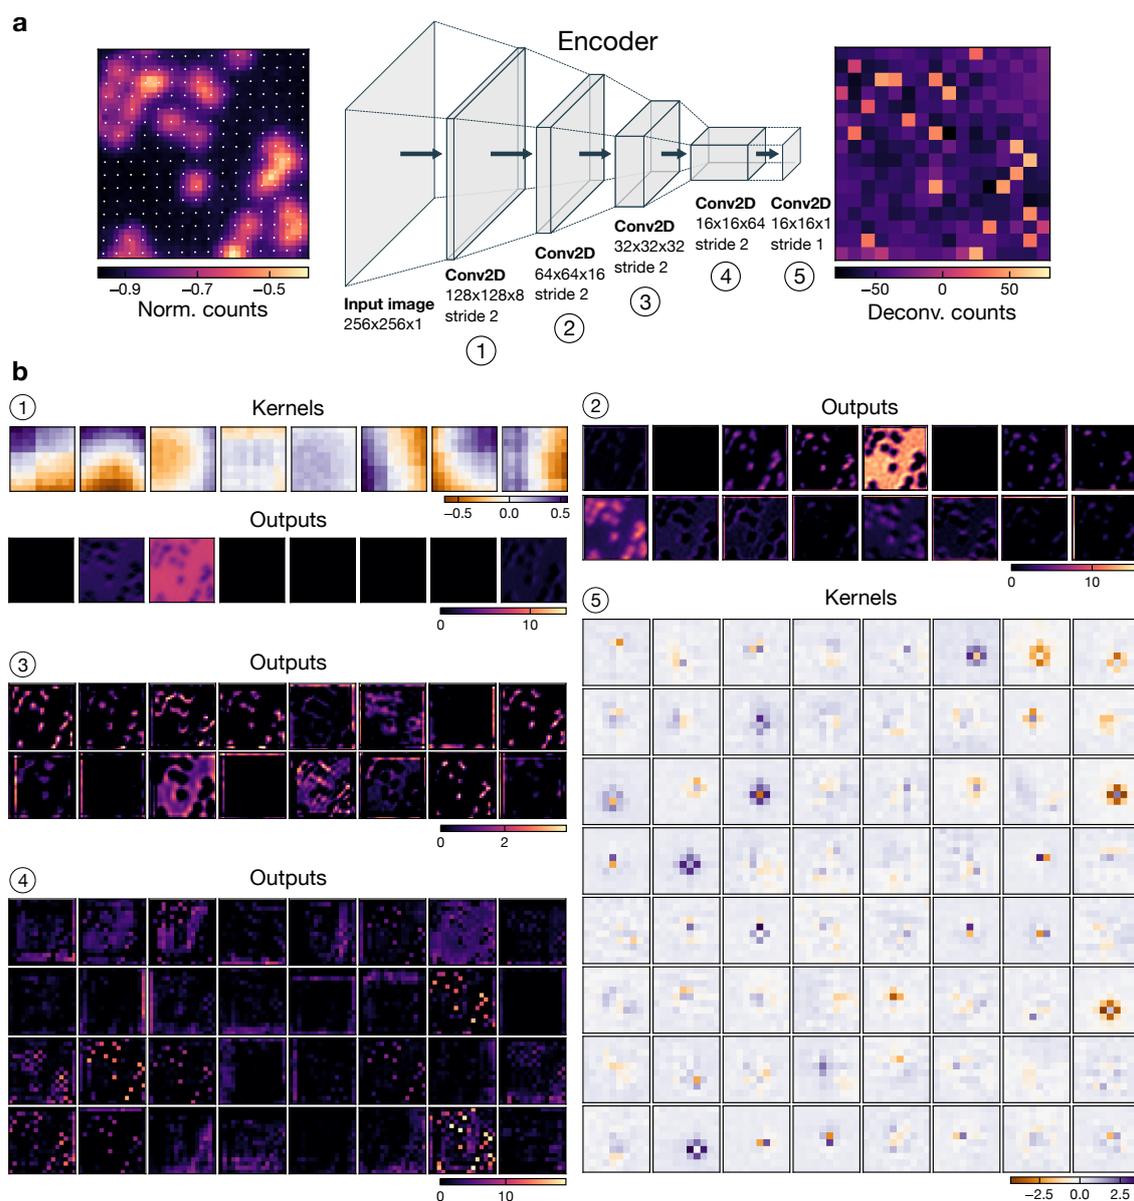


Figure 4.24 | Analysis of the learned kernels and feature maps of the encoder. **a**, Visualization of the learned transformation applied by the encoder to an experimental input subsection. The output values, extracted from the bottleneck layer before binarization, show a strongly enhanced contrast of the individual occupied sites in the input image, which is a clear signature of the deconvolution. **b**, Visualization of a selected number of kernels and feature maps (direct output of individual convolution layers) for the different steps of the encoder, which are marked by the encircled number. Especially noteworthy are the kernels of the last layer (5), whose alternating negative–positive weights are reminiscent of the linear kernel deconvolution, where the negative weights are used to cancel the spillover of signal from neighboring sites.

Changes to the network architecture and data processing

Due to the deviation of our lattice geometry from an ideal square lattice by about 2° , we cannot align the lattice vectors perfectly with the image axes by a simple rotation (see lattice sites in

Fig. 4.14). On the one hand, this could impact the performance of the network, as the centers of the sites are shifted throughout the image with respect to the kernels, which traverse the input horizontally/vertically. On the other hand, the angular deviation limits the maximum size of an input crop, as we have to prevent one row or column merging into the adjacent one at the end of the image. Even though the convolutional operations are in principle agnostic w.r.t. the input size, we had to always crop fluorescence images into 16×16 lattice site subsections, which was tedious and took a significant fraction of the duration for reconstructing a single image.

To address this, we instead resample the fluorescence images to a regular grid, where the width of one lattice site corresponds to exactly 16 pixels, and the lattice axes coincide with the image axes. We found that this resampling and interpolation operation (using a linear nearest-neighbor interpolation) does not influence the reconstruction performance. As a result, the architecture can now process images of arbitrary size within a single operation. To do this, the input layer of the network can be adjusted to a different shape even after training, with the only constraint that the fixed ratio of 16 between the width/height of the input layer and the bottleneck is maintained. This further also allows us to use larger subsections of 32×32 sites during training, which appeared to increase the reconstruction performance slightly. Lastly, the training crops can also be produced significantly faster, as they can be created simply using array slicing operations.

Optimization of the PSF

In the beginning of 2024, we decided to install a second microscope objective from the top of the vacuum chamber, together with the use of an improved alignment strategy. The goal was to decrease the size of the PSF and to reduce the long-range tail, which we believed to be a major source for the residual systematic errors in the reconstruction. As described in Section 3.1, we were indeed able to reduce the size of the PSF by a few percent from $d = 850$ nm to $d = 820$ nm, and, more importantly, to almost entirely eliminate the long-range tail. The further imaging parameters are roughly unchanged compared to the original data.

Current reconstruction performance

With these optimizations in place, we re-evaluate the reconstruction performance of the autoencoder on the experimental data. Fig. 4.25a shows the deconvolved count distributions for four different fillings from the validation set. Compared to the original data (Fig. 4.19), the distributions exhibit a similar separation between empty and occupied sites, hinting to a similar reconstruction fidelity. This is also confirmed by a fit of a bimodal normal distribution to extract the overlap fraction, which yields a fidelity above 99% at half filling. The separation at the highest filling bin is possibly even slightly increased, besides a small bridge between the distributions, which could originate from an increased hopping and loss rate compared to the original data.

A striking difference can be found when analyzing the density-density correlations in the reconstructed occupation. As shown in Fig. 4.25b, the artificial correlations at short distances are completely eliminated. This is a significant improvement compared to the original data

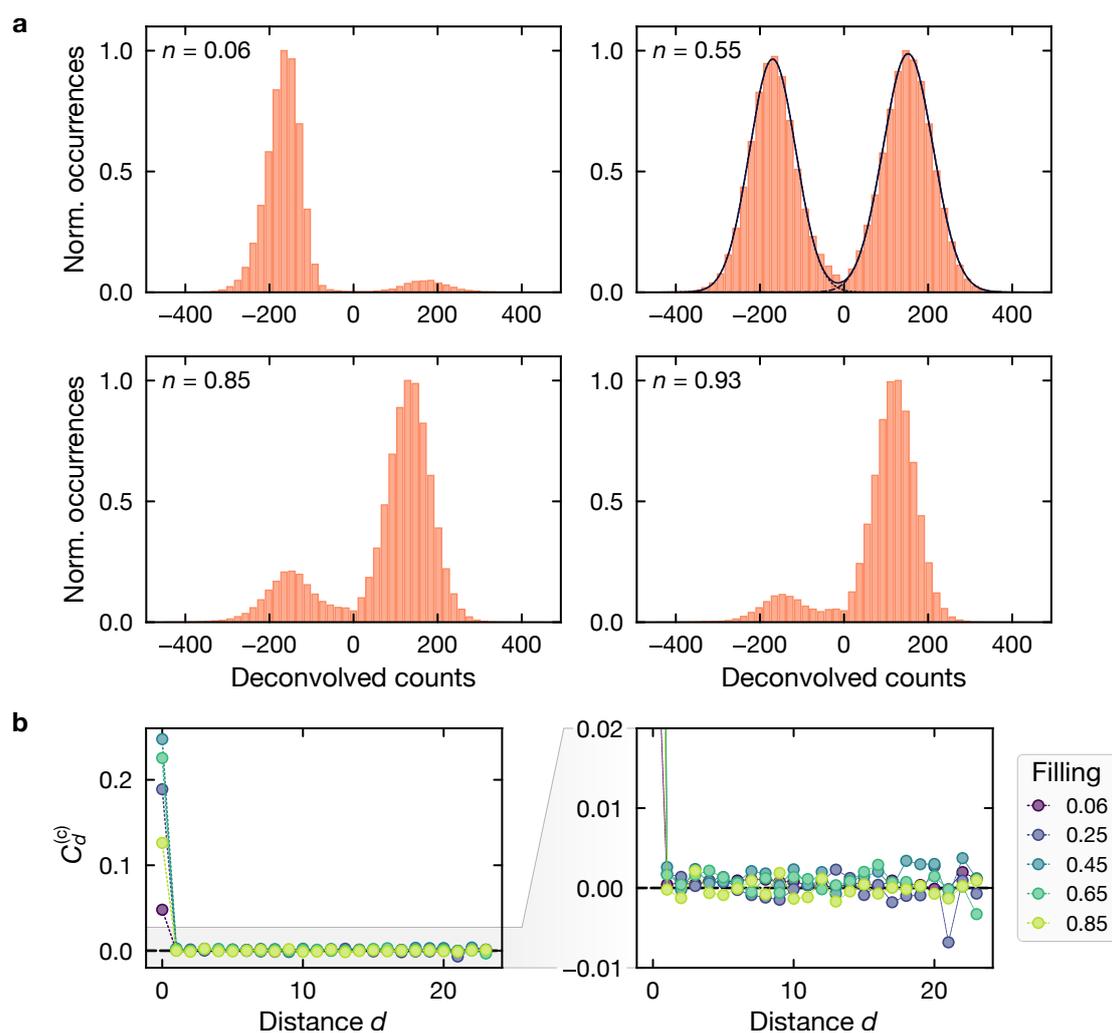


Figure 4.25 | Optimized autoencoder performance. **a**, Distributions of the deconvolved counts for four different fillings. The bi-modality of the distributions remains roughly unchanged with a clear separation across all fillings. A fit of a bimodal normal distribution (black lines, individual components shaded), confirm a fidelity above 99% at half filling. The occurrences are normalized to the respective maximum in each plot. **b**, Density-density correlations as a function of distance for different fillings, evaluated along lattice axis 1. The artificial correlations at short distances are fully eliminated, indicating the absence of systematic errors in the reconstruction. The dashed lines are guides to the eye.

(Fig. 4.21), where we found artificial correlations on the 2%-level for $d = 1$ and $d = 2$. Now, all correlators for $d > 0$ are below 0.5%, which is on the same level as the statistical fluctuations. This indicates that the residual systematic errors in the reconstruction have been eliminated, where the crucial ingredient was likely the removal of the long-range tail in the PSF and the slight decrease of the PSF size. The absence of artificial correlations is a crucial improvement, as it allows for a more precise measurement of small correlation signals. In particular, it recently allowed us to measure up to four-point density correlations and find excellent agreement with numerical simulations [291], which is a strong indicator of a high-fidelity reconstruction without systematic errors.

4.4 Conclusions

In this chapter, we have presented a powerful approach for the reconstruction of the atomic occupation in quantum gas microscopy images based on unsupervised neural networks. The deep, convolutional autoencoder architecture implements a flexible, non-linear transformation that has an inherent advantage in solving the inverse problem of deconvolution [246], while also capturing density-dependent effects [31, 89]. At the same time, the unsupervised nature enables a training directly with experimental data, eliminating the need for a detailed modeling of the imaging process or a fluorescence simulation whose accuracy ultimately limits the achievable performance. Based on the experimental data, we have developed a set of methods to evaluate the reconstruction performance despite the missing knowledge about the true underlying occupation, factoring in both statistical and systematic errors. This indicated high reconstruction fidelities above 95% for all fillings, previously seen only for significantly lower resolution-to-spacing ratios, and reconstructed occupations that are free from unphysical spatial correlations. For more than two years now, the autoencoder-based reconstruction has been in use in our experiment and enabled a wide range of quantum simulation experiments in an extreme resolution-to-spacing regime [124, 226, 290, 291]. Noteworthy, we have recently been able to detect up to four-point density correlations on the 10^{-3} level and find excellent agreement with numerical simulations [291], which requires a highly sensitive detection without systematic errors. Beyond our experiment, we expect our approach to enable new avenues for quantum simulation in exotic lattice configurations and extreme parameter regimes, as it becomes increasingly relevant e.g. for sub-wavelength lattices [292, 293] and heavy, dipolar quantum gases [77]. The applicability naturally extends also to related platforms featuring single-particle detection such as Rydberg-atom arrays [294, 295] and ion trap experiments [296].

Besides the high fidelity, a further strength is the single-pass, non-iterative nature of the reconstruction, which allows the reconstruction of a large fluorescence image with 10 000 lattice sites in around 50 ms on a standard consumer CPU, and even faster using a GPU. This is a large advantage in particular compared to iterative methods that can take up to a minute for an image of this size. A fast reconstruction time is crucial particularly in neutral atom quantum computing platforms for real-time feedback and control, e.g. for re-arrangement and mid-circuit operations [297, 298].

Future avenues for the development of the presented architecture include a systematic reduction of the network complexity, with the goal to find a minimal architecture that still captures the relevant features and allows a high reconstruction fidelity. This would shorten training and evaluation times, and also enhance the interpretability of the network. Besides a reduction of the number of kernels and their sizes, possible strategies include also a sharing of parameters between encoder and decoder [299]. Furthermore, it would be interesting to investigate a reconstruction of multiply-occupied sites with this architecture, as it can arise for example when splitting doubly-occupied sites into two adjacent layers of our short-spacing vertical lattice that are both within the objective's depth of focus [300]. Lastly, a relevant question is whether such an approach can also have advantages in the opposite regime of a high resolution but low SNR, as it is becoming increasingly important with tweezer-based platforms. Here, it could also be advantageous to employ a supervised approach (e.g. using only

the encoder part of the network) that is trained against another high-fidelity, but possibly slow, deconvolution method. A long term goal in the community should be a concise review of the performance of common approaches across different parameter regimes, in particular, as novel reconstruction methods are developed, such as the recently proposed Bayesian approach [301].

CHAPTER 5

Local readout and control of orbital operators

In this chapter, we present a method to read out and control orbital operators such as current and kinetic energy with local resolution in optical lattices. First, we start by illustrating the relevance of measuring observables beyond the local occupation in the context of simulating quantum many-body physics. We then explain our scheme to read out and control orbital operators using optical superlattices and demonstrate it experimentally, including both global and spatially-selective manipulations. Finally, we will show that the achieved single-shot readout allows for the measurement of correlation functions of these orbital observables.

5.1 Orbital operators for quantum simulation of many-body physics

With the advent of quantum gas microscopy, the study of many-body quantum systems with microscopic resolution has been possible through single-site resolved measurements. To probe the sample, fluorescence images are recorded, which yield a snapshot of the atomic density distribution. This corresponds to a readout of the local density operator \hat{n} , based on which information about the many-body quantum system is extracted [38]. In combination with the ability to measure non-local correlation functions as well as density fluctuations and full counting statistics on the single-atom and single-site level, this single observable already offers a powerful microscopic probe [33]. For example, in some of the first quantum gas microscopy experiments, local density distributions and density fluctuations were used to probe the quantum phase transition between a superfluid and a Mott insulator [31, 40]. Correlation functions of the local density are an important tool to characterize quantum phases, and have been used to detect (non-) local and string order [32, 33, 302], as well as to probe the build-up of correlations through quench experiments [226, 303]. By adding a spin-resolved density readout, the combined information about spin and density further opens a window into the physics of the Fermi Hubbard model [304–306] and the study of unconventional metals and high-temperature superconductivity [307].

However, a long-standing, central limitation of this detection technique is that it cannot detect off-diagonal elements (i.e., coherences) of the density matrix, in which a wealth of additional information about the quantum state remains hidden [33, 38]. Key examples of such off-diagonal observables include the current operator and the kinetic energy operator. A

measurement of the current operator would, for example, aid in the study of non-equilibrium dynamics by allowing one to probe information scrambling through bond-resolved local currents and off-diagonal correlations [308, 309], as well as the measurement of transport coefficients [310]. Spatially-resolved currents are also central fingerprints of interacting topological phases, where persistent currents can arise as a result of synthetic gauge fields [91, 165]. Lastly, a single-site-resolved detection of the full single-particle density matrix $\langle \hat{a}_i^\dagger \hat{a}_j \rangle$ is an important ingredient for Hamiltonian learning, which is a promising approach to benchmark analog quantum simulators [311–315].

A measurement of currents without spatial resolution has been achieved before both in a single-component system [93] and with a spin-degree of freedom [316]. With local resolution, experiments have attempted to implement atomic beam splitter operations either in two individual tweezers or using potentials engineered with a spatial light modulator, giving access to off-diagonal observables and information about particle-number statistics, indistinguishability, or the second-order Renyi entropy [34, 35, 92, 317, 318]. A notable further example includes a recent experiment where phase information was extracted using matter-wave imaging [319], although not yet demonstrated on the single-atom level. Beyond optical lattices, Ref. [54] measured coherences and off-diagonal string operators in a Rydberg array. However, the previous examples in the scope of itinerant lattice systems were merely proof-of-principle demonstrations, where the step allowing the readout of off-diagonal observables was either not scalable or not robust enough to serve as a standard readout technique in the neutral atom quantum simulation toolbox. In the following, we will showcase how optical superlattices can be used to do precisely that, i.e., to read out and control off-diagonal observables in a robust and scalable fashion. This is made possible by the ability of optical superlattices to create a high-quality, uniform array of double wells across many lattice sites, enabling parallel, high-fidelity manipulations between nearest neighbors. In the past, this has been used for example to generate many entangled atom pairs based on controlled collisions and exchange interactions [44, 320–323].

We will now begin by formally introducing the central observables of interest, based on the example of the Bose-Hubbard model, and then proceed to presenting the readout and control scheme using optical superlattices. As previously introduced in Section 1.1, the Bose-Hubbard model, which describes the interplay between kinetic energy and interaction energy in a tight-binding lattice, reads

$$\hat{H} = -J \sum_{\langle i,j \rangle} (\hat{a}_i^\dagger \hat{a}_j + h.c.) + \frac{U}{2} \sum_i \hat{n}_i(\hat{n}_i - 1), \quad (5.1)$$

where \hat{a}_i^\dagger and \hat{a}_i are the creation and annihilation operators at site i , $\hat{n}_i = \hat{a}_i^\dagger \hat{a}_i$ is the number operator, J is the nearest-neighbor tunnel coupling, and U is the on-site interaction energy.

The two characteristic terms that comprise the model contain only on-site and nearest-neighbor terms, respectively. While a measurement of local densities already captures the on-site contributions (the *diagonal* entries of the density matrix), and hence the interaction term, natural further observables that are characteristic for the system are those that live on bonds connecting two nearest-neighbor lattice sites (the *off-diagonal* entries of the density matrix).

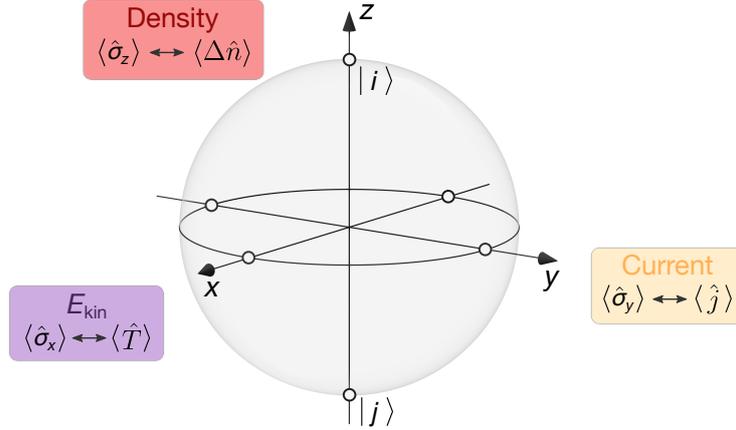


Figure 5.1 | The orbital Bloch sphere. The state space of a single particle on two nearest-neighbor lattice sites with labels i, j can be visualized as a Bloch sphere, where a measurement of the local density, current, and kinetic energy corresponds to a measurement along the $z, y,$ and x axis, respectively.

The first of this kind is the 'local kinetic energy', i.e., the individual spatial contributions of the kinetic energy term, defined as

$$\hat{T}_{i,j} = -J \left(\hat{a}_i^\dagger \hat{a}_j + \hat{a}_j^\dagger \hat{a}_i \right), \quad (5.2)$$

where i, j are two nearest-neighbor sites ($|i - j| = 1$). The second important observable of this kind can be found by looking at the local continuity equation, which relates the time derivative of the local density to the divergence of the local current, and for a lattice system reads:

$$\partial_t \hat{n}_i = - \sum_{j \in \text{NN}(i)} \hat{j}_{i \rightarrow j}. \quad (5.3)$$

Here, $\hat{j}_{i \rightarrow j}$ denotes the current between sites i and j , and $\text{NN}(i)$ is the set of nearest-neighbor lattice sites of site i . Using the Heisenberg equation of motion, the current operator for the Hubbard model (Eq. 5.1) can be written as [90]

$$\hat{j}_{i \rightarrow j} = -iJ \left(\hat{a}_i^\dagger \hat{a}_j - \hat{a}_j^\dagger \hat{a}_i \right). \quad (5.4)$$

As for the kinetic energy, the current operator is defined on bonds connecting nearest-neighbor lattice sites. In fact, within the restricted Hilbert space of the local two-site system, the density, current and kinetic energy operators form a complete basis. By combining measurements of these observables, one hence obtains access to the full single-particle density matrix $\langle \hat{a}_i^\dagger \hat{a}_j \rangle$, and can measure any local operator that can be expressed as a linear combination of these three.

In the following, we will use a simplified single-particle picture to illustrate the measurement of these observables. To this end, a single particle on the two sites i and j can be described by a two-level system, illustrated in the form of a Bloch sphere in Fig. 5.1. A measurement of the density, current, and kinetic energy then corresponds to a measurement of the z, y and x components of the Bloch vector, respectively. We refer to this picture as the *orbital Bloch sphere*, and the corresponding operators as *orbital operators*.

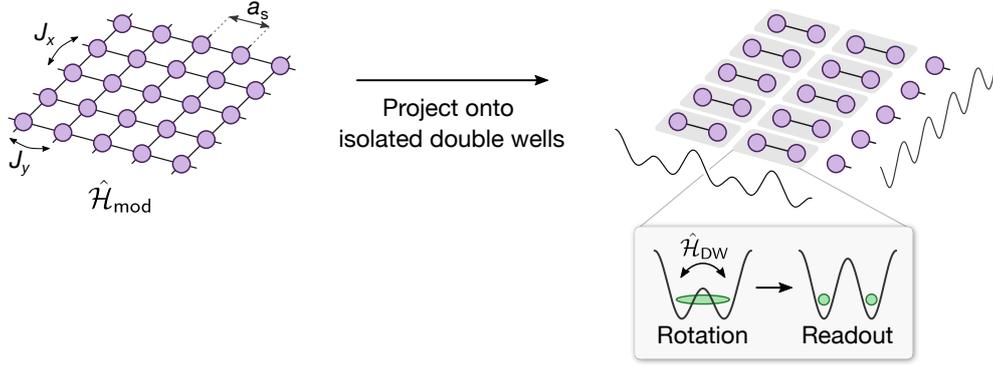


Figure 5.2 | Applying bond manipulations through double-well projection. To access off-diagonal observables for a quantum state in a 2D lattice system evolving under a model Hamiltonian $\hat{\mathcal{H}}_{\text{mod}}$, a deep period-two lattice is suddenly turned on, isolating individual bonds in the lattice along one axis. This projects the system onto an array of isolated double wells, allowing the application of programmable operations on each isolated bond.

5.2 Local readout and control scheme using optical superlattices

We now discuss how the previously introduced orbital operators can be measured with local resolution. The basis of our protocol is formed by the ability to isolate individual bonds in a 2D lattice system, which is experimentally implemented using an optical superlattice [74, 75]. As illustrated in Fig. 5.2, we start from a generic 2D optical lattice, where an experiment is carried out under some Hubbard-type model Hamiltonian $\hat{\mathcal{H}}_{\text{mod}}$. At the point where the quantum state should be probed, the first step of the measurement protocol is to suddenly turn on a deep lattice with a periodicity of two lattice sites along one horizontal axis. This partitions the lattice into a 2D array of decoupled double wells, isolating every other bond along the superlattice axis. We now require that we have the possibility to switch off on-site interactions ($U \rightarrow 0$) upon projection, which we can do using a magnetic Feshbach resonance. Note however that this is only required for the measurement, and the protocol applies for a generic many-body state irrespective of the interaction strength. In the absence of interactions, we can apply the single-particle Bloch sphere picture from before (see Fig. 5.1), where the two sites of each double well form a two-level system. If we now immediately after projection read out the local occupation from a fluorescence image, this would correspond to a measurement in the z -basis. A measurement of the other two orthogonal components can hence be achieved by rotating the measurement basis out of the z -direction and onto either x or y for kinetic energy and current, respectively. In other words, the goal is to find a protocol that allows us to map arbitrary orbital operators onto the local density, which we can subsequently read out using optical imaging. The isolated double wells provide the necessary operations to achieve this and allow us to rotate the state vector around the fundamental axes of the orbital Bloch sphere.

Specifically, after the projection, the state in each double well in the absence of on-site interactions ($U = 0$) evolves according to the double well Hamiltonian

$$\hat{\mathcal{H}}_{\text{DW}} = -J\hat{\sigma}_x - \frac{\Delta}{2}\hat{\sigma}_z, \quad (5.5)$$

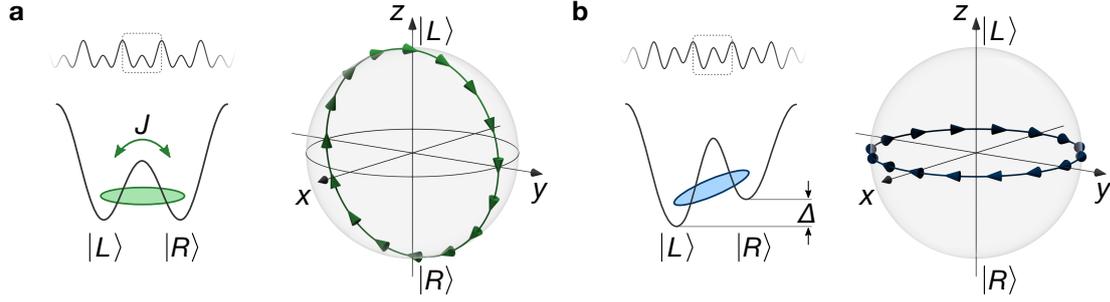


Figure 5.3 | Atomic operations in a double well. **a**, The evolution in a symmetric, coupled double well with coupling strength J drives a rotation around the x -axis. **b**, A tilted, decoupled double well with energy offset Δ implements a rotation around the z -axis of the Bloch sphere.

where $\hat{\sigma}_i$ ($i \in \{x, y, z\}$) are the Pauli operators acting on the two-level system, J is the tunnel coupling, and Δ is the energy offset between the two wells. The individual Pauli operators generate rotations around their respective axes. Hence, in a symmetric double well ($\Delta = 0$), a rotation around the x -axis of the Bloch sphere is implemented (Fig. 5.3a), while a strongly tilted, decoupled double well ($\Delta \gg J$) drives a rotation around the z -axis (Fig. 5.3b). In the following, we will denote a rotation by an angle of ϕ around the x -axis as X_ϕ , and a rotation around the z -axis as Z_ϕ .

We first discuss how the current operator can be measured using this protocol. Looking at the orbital Bloch sphere (Fig. 5.1), we see that the current operator is mapped onto the density operator by a 90° rotation around the x -axis. In more detail, after projecting the system into isolated, symmetric double wells at a tunnel coupling J_{DW} , each projected bond evolves according to $\hat{\mathcal{H}}_{\text{DW}} = -J_{\text{DW}}\hat{\sigma}_x$. Under this Hamiltonian, the left-right density difference evolves as a function of time according to [90, 93]

$$\begin{aligned} \hat{n}_R(t) - \hat{n}_L(t) = & [\hat{n}_R(0) - \hat{n}_L(0)] \cos(2J_{\text{DW}}t/\hbar) \\ & + i(\hat{a}_R^\dagger \hat{a}_L - \hat{a}_L^\dagger \hat{a}_R) \sin(2J_{\text{DW}}t/\hbar). \end{aligned} \quad (5.6)$$

We find two out-of-phase oscillating terms, where the first term is proportional to the initial density difference, while the second term encodes the bond current according to Eq. 5.4. Choosing an evolution time of $\tilde{t} = \hbar/(8J_{\text{DW}})$, equivalent to a $\pi/2$ or 90° rotation around the x -axis of the Bloch sphere, the first term drops out and we obtain

$$\hat{n}_R(\tilde{t}) - \hat{n}_L(\tilde{t}) = i(\hat{a}_R^\dagger \hat{a}_L - \hat{a}_L^\dagger \hat{a}_R) = \hat{j}_{L \rightarrow R}/J. \quad (5.7)$$

This means that the bond current can be directly read out from the density difference of the two sites after applying an $X_{\pi/2}$ pulse, enabling a single-shot sampling of the current operator.

A similar treatment can be applied for the measurement of the kinetic energy operator. A glance on the Bloch sphere (Fig. 5.1) reveals that the kinetic energy operator is mapped onto the density operator by a $\pi/2$ rotation around the y -axis. Such a rotation is not immediately available in our double well system, but can be emulated by concatenating a $\pi/2$ rotation around the z -axis with a $\pi/2$ rotation around the x -axis. The rotation around the z -axis is

applied by evolving in a strongly tilted double well, where the Hamiltonian reads approximately $\hat{\mathcal{H}}_{\text{DW}} = -\frac{\Delta}{2}\hat{\sigma}_z$. With this, the ladder operators evolve as

$$\begin{aligned}\hat{a}_R^\dagger(t) &= \hat{a}_R^\dagger(0)e^{i\Delta t/\hbar} \\ \hat{a}_L^\dagger(t) &= \hat{a}_L^\dagger(0).\end{aligned}\tag{5.8}$$

In particular, for a $Z_{\pi/2}$ rotation [$\tilde{t} = \hbar/(4\Delta)$], we obtain

$$\hat{a}_R^\dagger(\tilde{t}) = i\hat{a}_R^\dagger(0), \quad \hat{a}_L^\dagger(\tilde{t}) = \hat{a}_L^\dagger(0).\tag{5.9}$$

Inserting this into Eq. 5.7, which corresponds to a $(Z_{\pi/2}, X_{\pi/2})$ pulse sequence, we find

$$\hat{n}_R(\tilde{t}) - \hat{n}_L(\tilde{t}) = -\left(\hat{a}_R^\dagger\hat{a}_L + \hat{a}_L^\dagger\hat{a}_R\right) = \hat{T}_{L,R}/J.\tag{5.10}$$

In agreement with the Bloch sphere picture, this shows that the kinetic energy operator can be directly read out from the density difference of the two sites after a $(Z_{\pi/2}, X_{\pi/2})$ pulse sequence.

To summarize, we have shown how – in addition to the density – the kinetic energy and the current operator, and any linear combination of them, can be measured with local resolution using optical superlattices. A sampling of these observables is possible within a single experimental cycle, and the method is scalable to large system sizes due to the parallel realization of many high-quality double wells in a superlattice. We would further like to emphasize that the presented protocol applies to a generic many-body state, given that interactions can be switched off *during* the basis rotation. Additionally, as we will detail below, this protocol can also be used more generally to apply tailored, local manipulations to a quantum state, e.g. to engineer initial states with a complex spatial structure or containing spatial superposition states.

5.3 Experimental implementation of global operations

As a proof of principle, we demonstrate and benchmark the individual superlattice rotations that form the basis of the protocol. To allow benchmarking of the operations in a controlled manner, we use a simple model system, consisting of a 2D array of isolated double wells. Each double well is initialized with exactly one particle in the left well $|L\rangle$. The advantage of this initial state is that we can prepare it with high fidelity, and we are furthermore not sensitive to many-body effects. We obtain a maximum imbalance of typically 0.93(4), and during analysis, we post-select on double wells containing exactly one particle. In the Bloch sphere picture, the initial state corresponds approximately to the north pole with $\langle\hat{\sigma}_z\rangle = 1$ and $\langle\hat{\sigma}_x\rangle = \langle\hat{\sigma}_y\rangle = 0$.

We start by benchmarking rotations around the x -axis. With the given initial state, these are initiated by suddenly lowering the barrier between the two sites via the short lattice depth at a vanishing superlattice phase offset (see Fig. 5.4a). As shown in Fig. 5.4b, we observe high-contrast oscillations of the left-right density imbalance as a function of the pulse duration. For the chosen lattice parameters and as extracted from a sinusoidal fit, we find a tunnel coupling $J = \hbar \times 484.3(5)$ Hz and an experimental X_π pulse duration of $t_\pi^{(\text{exp})} = 449(3)$ μs . Within the first few oscillation cycles, there is a negligible decay of the oscillation contrast visible. To

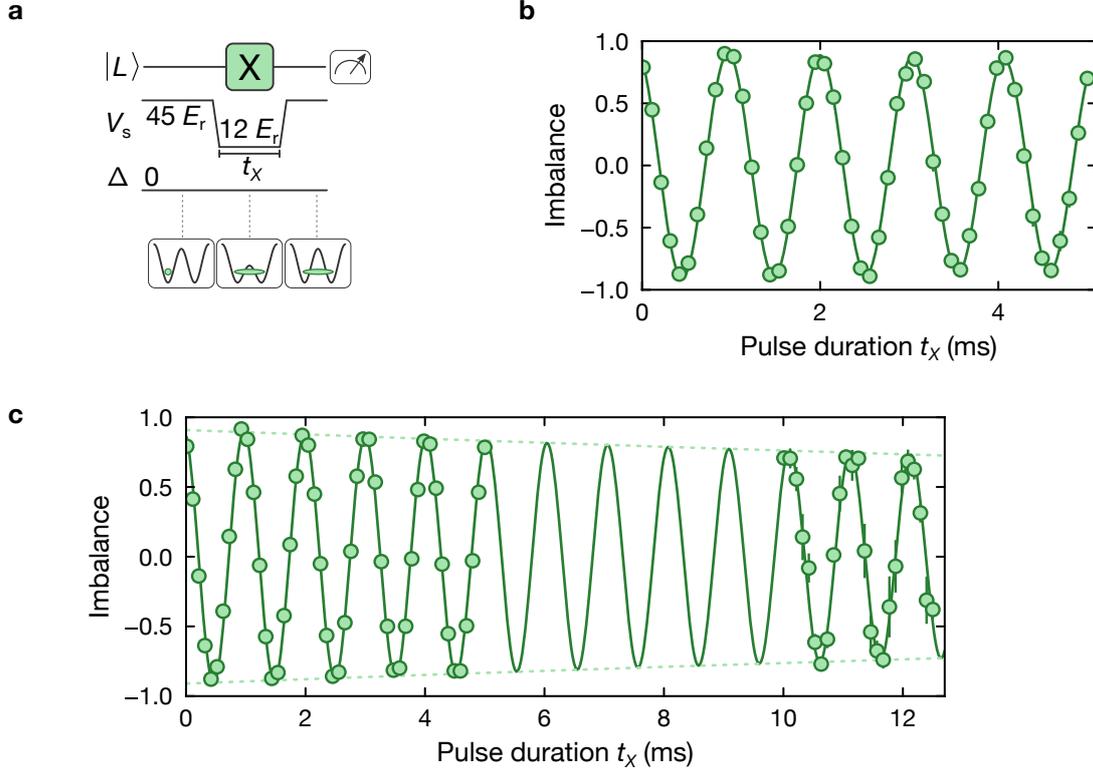


Figure 5.4 | Global X rotations. **a**, Schematic of the pulse sequence. Global X rotations are applied by evolving in symmetric double wells for a variable duration t_x . V_s denotes the depth of the short lattice, which sets the tunnel coupling J inside the well. **b**, Measured imbalance as a function of the pulse duration t_x . The solid line is a fit of a sine function, and the error bars denote the standard error of the mean across 6 repetitions (partly smaller than the marker size), evaluated in a ROI of 28×32 sites. **c**, Measured imbalance as a function of the pulse duration t_x for many cycles. The solid line is a fit of an exponentially decaying sine function, and the dashed lines show the decay envelope. The error bars denote the standard error of the mean across 20 repetitions for the shorter durations and 5 repetitions for the longer durations, respectively (partly smaller than the marker size), evaluated in a ROI of 28×32 sites.

further quantify the fidelity, we record the imbalance oscillation for longer durations (Fig. 5.4c), where we observe a slow decay of the envelope at a $1/e$ decay constant of $\tau = 57(13)$ ms. This decay constant corresponds to a fidelity of a single X_{π} pulse of $\mathcal{F} = 99.2(2)\%$, which is mostly limited by spatially-inhomogeneous potential energy variations which detune the double wells locally and modify the oscillation frequency according to $f = \sqrt{4J^2 + \Delta^2}/h$, as well as a residual coupling between adjacent double wells.

In the following, we demonstrate rotations around the z -axis. Since a z -rotation leaves the z -component of the Bloch vector unchanged, the most straightforward protocol to benchmark this operation is a *Ramsey sequence*. As illustrated in Fig. 5.5a, we implement this by embedding a z -pulse of variable duration t_z between two $X_{\pi/2}$ pulses. The first $X_{\pi/2}$ pulse moves the Bloch vector onto the equatorial plane, after which we quickly jump the superlattice phase away from the symmetric configuration while simultaneously decoupling the sites. The energy offset Δ between the sites causes the Bloch vector to rotate on the equator as the superposition

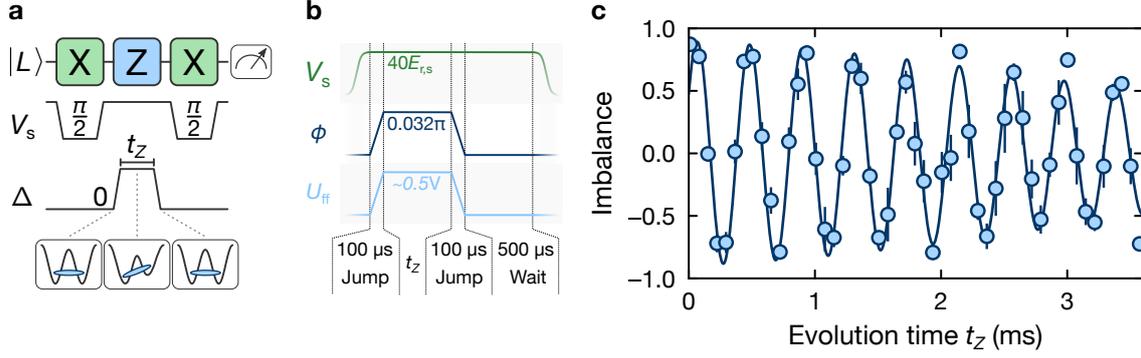


Figure 5.5 | Global Z rotations. **a**, Schematic of the Ramsey sequence to visualize global Z rotations that are applied by evolving in strongly tilted double wells for a variable duration t_z . **b**, Detailed sequence diagram, illustrating the timings for the feedforward technique, which allows us to jump the superlattice phase much faster than typical dephasing times. U_{ff} denotes the voltage applied to the feedforward input of the OPLL. **c**, Measured imbalance after the Ramsey sequence as a function of the pulse duration t_z . The solid line is a fit of a sine function with an envelope that describes the dephasing due to spatially varying energy offsets, and the error bars denote the standard error of the mean across 4 repetitions (partly smaller than the marker size), evaluated in a ROI of 16×30 sites.

of $|L\rangle$ and $|R\rangle$ evolves in the tilted double well. This evolution is finally probed using the second $X_{\pi/2}$ pulse, exhibiting oscillations that reveal the rotation of the state vector along the equator.

Experimentally, a challenging step is the sudden jump of the superlattice phase when changing between X and Z pulses. As described in Section 2.3.2, the intrinsic bandwidth of the frequency offset lock is limited to a change rate in the applied tilt of about $\delta\Delta/\delta T = (h \times 150 \text{ Hz}/1 \text{ ms})$ for the lattice depths used here, which is too slow compared to typical dephasing times. To overcome this limitation, we use a feedforward technique where a voltage is summed onto the output of the loop controller, allowing fast changes of the superlattice phase that are limited only by the actuator bandwidth. The full sequence to apply a z -rotation using the feedforward technique is shown in Fig. 5.5b. Immediately after applying the $X_{\pi/2}$ pulse, the superlattice phase is changed to the desired value in 100 μs , while simultaneously applying the previously calibrated analog voltage to the output of the loop controller. After evolving for a certain duration t_z , the superlattice phase is changed back to the symmetric configuration using the inverse of the previously applied ramp. While the feedforward technique allows for significantly faster phase changes, we found that we have to wait for around 0.5 ms for the phase to fully settle before applying the second $X_{\pi/2}$ pulse (note that this can be omitted after the first ramp, since it only leads to a z -rotation with a slightly modified initial speed that can be calibrated).

The resulting imbalance signal as a function of the Z pulse duration is shown in Fig. 5.5c. For the chosen superlattice phase offset of around 0.032π , we find sinusoidal oscillations corresponding to a tilt of $\Delta = h \times 2.406(5) \text{ kHz}$, exhibiting a weak damping over a few periods. The observed damping is primarily given by spatially varying energy offsets, which are most likely originating from random on-site potential disorder and locally modify the tilt Δ . This modification is linear in the energy offset as compared to a quadratically suppressed contribution for X rotations above, rendering Z rotations significantly more sensitive to spatial

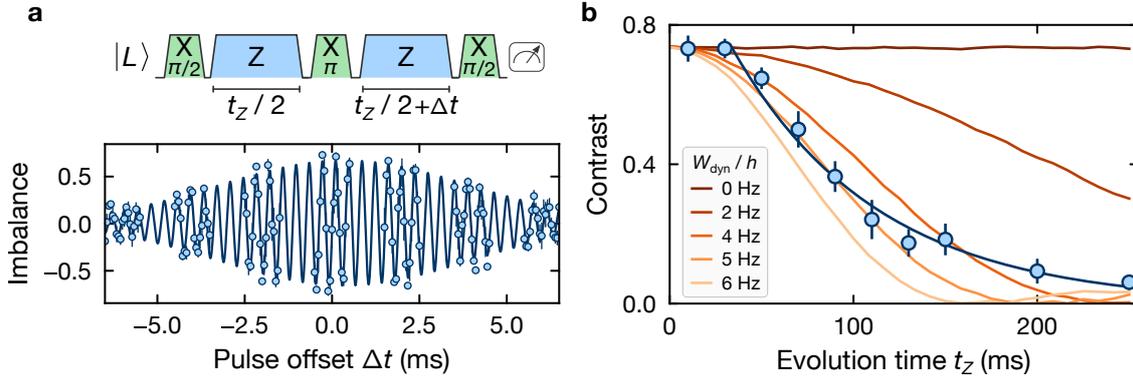


Figure 5.6 | Double well spin echo. **a**, Spin echo sequence used to cancel static potential inhomogeneities. An offset Δt in the duration of the second Z pulse serves to observe the revival of the imbalance contrast. Lower part: Example spin echo signal in the imbalance as a function of the pulse offset Δt for $t_z = 30$ ms. The solid line is a fit of a sinusoidal function with Gaussian envelope. **b**, Measured contrast of the spin echo signal as a function of the total Z pulse duration t_z . For each data point, the imbalance contrast was evaluated by varying the pulse offset and fitting the resulting imbalance oscillation. The solid blue line denotes an exponential fit omitting the initial data point, from which we derive T_2 as the duration at which the contrast decreases to $1/e$ of its initial value. The orange traces show a simulation of the spin echo contrast in presence of a dynamically varying, white-noise disorder of amplitude W_{dyn} between the two Z pulses, as indicated by the legend. The error bars denote the standard error of the fit. The data was evaluated in a ROI of 16×30 sites, averaging over 5 repetitions in each measurement.

inhomogeneities. The decay envelope in Fig. 5.5c is well-described by a numerical model assuming on-site white-noise disorder uniformly distributed in $[-W, W]$, yielding a disorder amplitude of $W = h \times 49(2)$ Hz. Alternatively, the decay can be approximated by a single exponential which allows us to extract a coherence time of $T_2^* = 6(1)$ ms ($1/e$ decay time). This coherence time is about one order of magnitude shorter than the dephasing time of X rotations, which highlights the increased sensitivity to spatial potential inhomogeneities. Additionally, it gives us an important timescale for developing experiment protocols. Whenever we are working with a quantum state where local phases or coherences between neighboring sites are important, and we have to freeze the system for a certain time to perform some change (e.g. a change in the offset field or in the superlattice phase), we have to do this much faster than this T_2^* time to preserve the coherences.

Despite the significant sensitivity of z -rotations to spatial inhomogeneities, the developed protocol allows for the quick application of large double well tilts, resulting in a fidelity well above 90% for a single Z_{π} pulse, which is sufficient for most measurement or manipulation applications. To enable high-fidelity manipulations beyond the timescale of single pulses, we can use composite pulse sequences such as a spin-echo sequence to cancel static spatial inhomogeneities. As illustrated in Fig. 5.6a, the idea behind a spin-echo is to extend the Ramsey sequence by an additional X_{π} pulse in the middle of the Z pulse. This pulse inverts the phase of the Bloch vector on the equator, effectively refocusing the evolution and canceling out the effect of spatial potential inhomogeneities. Varying the duration of the second half of the Z pulse, we observe a revival of the imbalance oscillations, directly visualizing the refocusing of

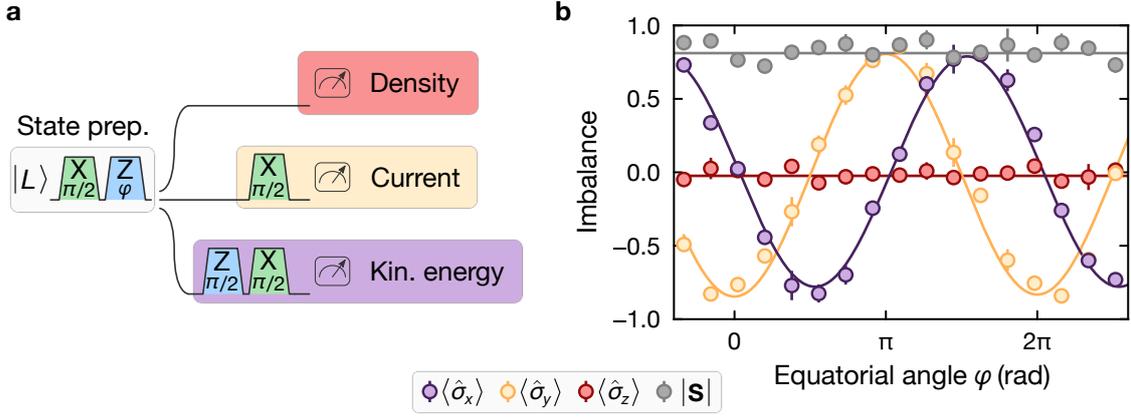


Figure 5.7 | Tomography of equatorial states. **a**, Sequence used to prepare single-particle states on the equator of the orbital Bloch sphere with a variable equatorial angle φ . All three components of the Bloch vector are measured by mapping them onto the density difference. **b**, Measured Bloch vector components as a function of the equatorial angle φ . The gray data points denote the calculated length of the Bloch vector $|\mathbf{S}|$. The solid lines are fits to the data (sine function for x - and y -components, constant function for z -component and $|\mathbf{S}|$), and the error bars denote the standard error of the mean across 3 repetitions (partly smaller than the marker size), evaluated in a ROI of 18×36 sites. The zero of the horizontal axis has been calibrated on the first minimum of the $\hat{\sigma}_y$ trace, and the 2π angle on the period of the corresponding oscillation.

the equatorial evolution (Fig. 5.6a). We repeat this sequence for different Z pulse durations t_Z and track the fitted maximum imbalance contrast of each spin-echo. As shown in Fig. 5.6b, we observe a smooth decay of the contrast with increasing t_Z on the order of 100 ms. By approximating the decaying part with a single exponential (solid blue line), we determine a T_2 time of 113(10) ms, which corresponds to around 270 cycles on the equator at the previously measured Δ . The T_2 time is more than an order of magnitude longer than the T_2^* time observed in the Ramsey sequence, demonstrating the effectiveness of the spin echo sequence in canceling out static spatial potential inhomogeneities. Remaining contributions to the decay are likely given by dynamic inhomogeneities, e.g. due to time-dependent fluctuations of the energy offset Δ , or by residual coupling between adjacent double wells. To describe the former, we perform a numerical simulation in which the white-noise disorder amplitude during the second Z pulse is changed by an additional value sampled from the interval $[-W_{\text{dyn}}, W_{\text{dyn}}]$. As can be seen from the orange traces in Fig. 5.6b, the observed decay is compatible with a dynamical disorder amplitude of $W_{\text{dyn}} = h \times 4 - 6$ Hz. This amplitude is rather small, even though the probed pulse durations on the order of 100 ms are expected to be sensitive to acoustic noise with frequencies of 10 – 1000 Hz that can modify the disorder through the pointing of laser beams or the phase of residual interference patterns. Since this is challenging to improve on the hardware side below this level, the most promising path to cancel also dynamical inhomogeneities is to use advanced pulse sequences such as multiple spin echoes [324, 325] or dynamical decoupling sequences [326].

As a simple example demonstrating our capability to measure in different bases, we perform a quantum state tomography of single-particle states on the equator of the orbital Bloch sphere. As shown in Fig. 5.7a, equatorial states of the form $|\psi\rangle = (|L\rangle + e^{i(\varphi+\pi/2)} |R\rangle)/\sqrt{2}$ are prepared

by applying an $X_{\pi/2}$ pulse followed by a Z_φ pulse to the initial state $|L\rangle$, with a variable Z pulse duration to tune the equatorial angle φ . Immediately after state preparation, we measure all three components of the Bloch vector $\mathbf{S} = (\langle\hat{\sigma}_x\rangle, \langle\hat{\sigma}_y\rangle, \langle\hat{\sigma}_z\rangle)$, i.e., the kinetic energy, the current, and the density, using the double well basis rotation protocol. As expected, we find that the density difference $\langle\hat{\sigma}_z\rangle$ remains zero independent of the angle (Fig. 5.7b). In contrast to that, the expectation values of the kinetic energy $\langle\hat{\sigma}_x\rangle$ and the current $\langle\hat{\sigma}_y\rangle$ show high-contrast oscillations with a relative phase shift of $\pi/2$. The measured Bloch vector length $|\mathbf{S}| = \sqrt{\langle\hat{\sigma}_x\rangle^2 + \langle\hat{\sigma}_y\rangle^2 + \langle\hat{\sigma}_z\rangle^2}$ is approximately constant at an average magnitude of 0.81(6). This length is limited mostly by the imperfect initial state preparation as well as the Z rotation used for state preparation, and the high oscillation contrast in the x and y components demonstrates the high fidelity of the basis rotations enabled by our protocol.

In this section, we have experimentally demonstrated the high-fidelity application of global bond operations. Both X and Z rotations are programmed using superlattice parameters, enabling an efficient, parallel application on all bonds along one direction of a 2D system. Moreover, the high-quality creation of double well structures using optical superlattices straightforwardly scales the protocol to large system sizes comprising thousands of lattice sites, rendering it a powerful new tool to extend the measurement capabilities of analog quantum simulators.

5.4 Extension to spatially-selective operations

In the previous section, we have shown how global operations can be applied to measure and control orbital operators with local resolution in a parallel fashion everywhere in a system. We now extend this protocol to allow for spatially-selective operations, enabling the application of truly local manipulations to a quantum state. At the heart of this lies the ability to selectively tilt certain double wells. To do this in a programmable fashion, we employ a digital micro-mirror device (DMD) to project arbitrary, repulsive potentials onto the lattice system. As before, the superlattice is used to partition the lattice into isolated double wells in a symmetric configuration, and the DMD is then patterned to project small repulsive potentials onto specific double wells to apply a tilt (see Fig. 5.8a). The advantage of this hybrid global-local approach is that it combines the high-quality bond partitioning of the superlattice with the local programmability offered by the DMD.

The repulsive potential features projected with the DMD effectively implement local Z rotations. These can be employed in two different ways, which we will detail in the following. First, we can apply local tilts during a global X rotation, which realizes locally detuned X rotations. As illustrated in Fig. 5.8a and identical to the benchmarking of global operations above, we start with single particles in isolated double wells initialized in $|L\rangle$. We then suddenly lower the barrier inside the double well via the short lattice depth to initiate global X rotations over $150\ \mu\text{s}$. At the same time, we ramp up the repulsive DMD potential with a mask consisting of stripes perpendicular to the superlattice direction. The stripe pattern has a periodicity of $4 a_s$, and each stripe has a programmed width of $0.6 a_s$, centered on one site of a double well (the width is chosen to give a sufficient tilt while minimizing leakage to adjacent sites). This pattern

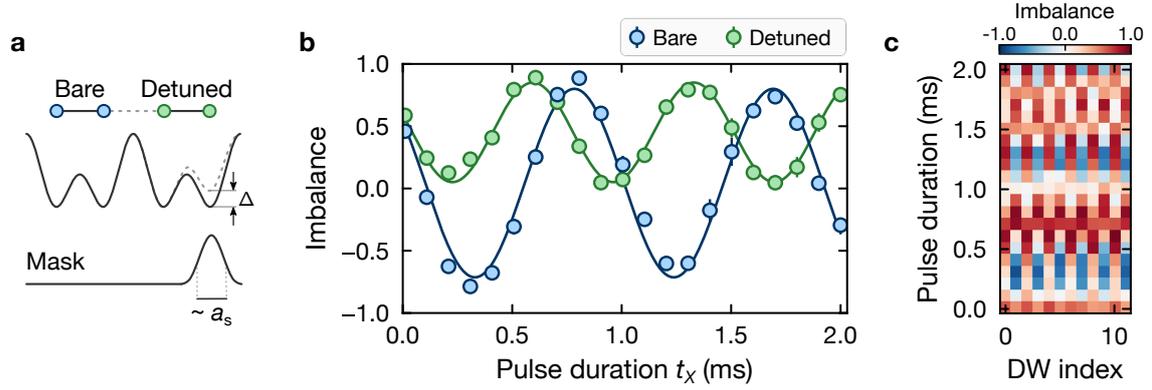


Figure 5.8 | Locally detuned X rotations. **a**, Idea of the local detuning protocol: Using a DMD, localized repulsive potentials are projected onto the system, tilting every other double well along a one-dimensional superlattice potential. Simultaneous to the local tilts, a global X rotation is applied. **b**, Measured imbalance in the unaffected (bare) and the tilted (detuned) double wells as a function of the X pulse duration. The solid lines are fits of a sine function, yielding a differential tilt of $\Delta = h \times 797(13)$ Hz. The error bars denote the standard error of the mean across 3 repetitions (partly smaller than the marker size), evaluated in a ROI of 24×24 sites. **c**, Spatially resolved measurement of the imbalance in the double wells along the superlattice axis using the same parameters as in **b**. The data is averaged over the perpendicular direction.

creates a local tilt in every other double well along the superlattice axis, resulting in subsequent locally detuned X rotations (Fig. 5.8a). We reveal these oscillations by evaluating the imbalance in even and odd double wells separately, each class averaged over the 2D system, as a function of the evolution time. As shown in Fig. 5.8b, we observe that one class oscillates with a high contrast and low frequency, while the other class shows a significantly lower contrast as well as a higher oscillation frequency. The former class corresponds to the un-tilted (or: *bare*) double wells, while the latter are the double wells subject to the DMD-imprinted tilt. The oscillation frequency of the tilted double wells is increased according to $f_{\text{osc}} = \sqrt{4J^2 + \Delta^2}/h$. From the fitted oscillation frequency, we extract a tunnel coupling of $J = h \times 550(5)$ Hz and a differential tilt of $\Delta = h \times 797(13)$ Hz. The maximum achievable tilt is experimentally limited by the available DMD power as well as the finite resolution of the DMD imaging system, which we estimate to broaden the tilt potentials to a width of around $\sim 1 - 2 a_s$. Nevertheless, it is noteworthy that we can achieve significant differential tilts up to around $h \times 1$ kHz despite our particularly small lattice spacing of 383.5 nm, which is sufficiently high to implement single rotations up to 2π that are fast compared to typical dephasing timescales. Lastly, we can additionally analyze this experiment with spatial resolution along the double well axis, averaging only over the perpendicular direction. As shown in Fig. 5.8c, this demonstrates that such local manipulations are possible with high spatial homogeneity across large systems. In particular, we find no discernible inhomogeneity across more than 20 lattice sites.

To optimize the spatial homogeneity and the attainable differential tilts, it is crucial to align the projected DMD mask correctly to the optical lattice structure. The overall size calibration from DMD pixels to lattice sites, including the alignment of the base vectors, is done as described in Section 3.1 and Ref. [213]. These calibrations were found to be rather stable over the course

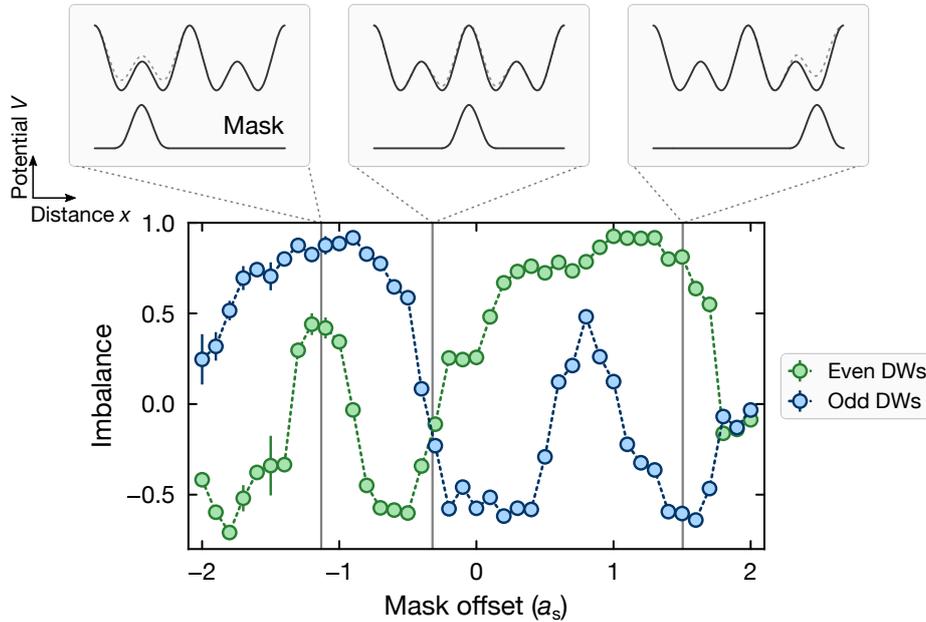


Figure 5.9 | Mask alignment. The correct alignment of the projected DMD mask with respect to the optical lattice is calibrated using locally detuned X rotations at a fixed pulse duration. The data points show the average imbalance in even and odd double wells as a function of the mask offset. The dashed lines are guides to the eye. The vertical lines together with the respective illustrations indicate a few example alignment situations of the DMD mask relative to the lattice structure. An optimal alignment is achieved when the difference in the imbalance of the two classes is maximized (here in four locations).

of several weeks. However, what has to be frequently recalibrated is a translational offset of the mask, which determines the centering of the projected potentials on the lattice sites. The offset is very sensitive to environmental parameters such as air humidity and pressure, which can drift on the order of one to several hours, depending on the specific conditions (rapid outside temperature changes, arrival of storm fronts). We calibrate this using locally detuned X rotations as done above in Fig. 5.8 for a fixed pulse duration. We use an identical mask, tilting every other double well along a 1D superlattice chain, and vary an offset that shifts the mask along one of the lattice axes. The pulse duration is chosen slightly shorter than the time for an X_π pulse in the absence of a tilt, i.e., around 0.25 ms for the parameters in Fig. 5.8b. In this case, no tilt results in maximum imbalance opposite to the initial state, while any applied tilt brings the imbalance closer to the initial imbalance, resulting in a maximum dynamic range and calibration sensitivity. The result of such a scan over the full periodicity of the pattern of $4 a_s$ is shown in Fig. 5.9. We observe a distinct behavior, which is symmetric roughly around the center of the scan range. At the symmetry point, the two double well classes are interchanged, corresponding to a situation where the tilt mask is aligned right in between two double wells (middle illustration). The optimal alignment is achieved when the difference in the imbalance of the two classes is maximized, which occurs at four locations in the scan (one location is shown in the rightmost illustration). In principle, either location can be chosen. However, we find that the pattern is not completely symmetric, with some locations showing a slightly broader plateau than others, which likely originates from non-isotropic aberrations in the DMD

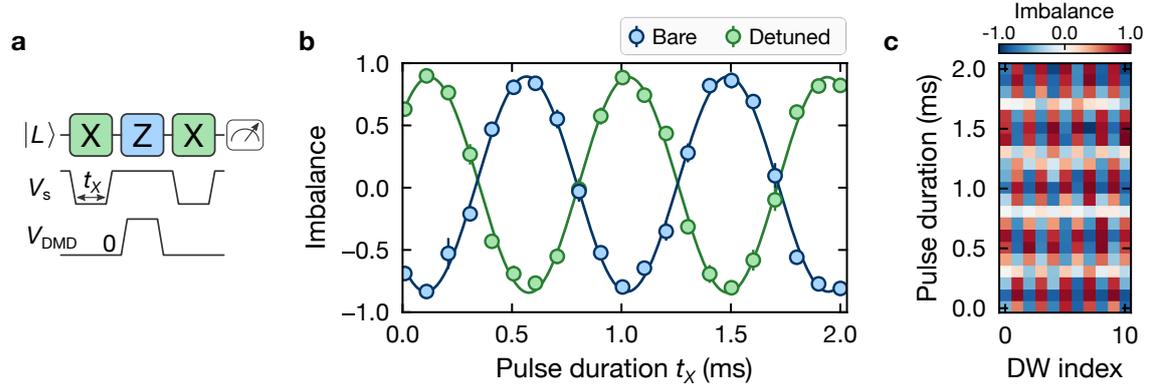


Figure 5.10 | Local Z rotations. **a**, Ramsey-type sequence used to calibrate and demonstrate local Z rotations imprinted with DMD potentials. The DMD is programmed with a mask consisting of stripes perpendicular to the superlattice axis, locally adding an energy offset to every fourth site along the superlattice axis. **b**, Measured imbalance in the unaffected (bare) and the tilted (detuned) double wells as a function of the first X pulse duration. The solid lines are fits of a sine function, yielding an imprinted relative phase shift of $1.04(2)\pi$. The error bars denote the standard error of the mean across 3 repetitions (partly smaller than the marker size), evaluated in a ROI of 24×24 sites. **c**, Spatially resolved measurement of the imbalance in the double wells along the superlattice axis using the same parameters as in **b**. The data is averaged over the perpendicular direction.

imaging system. We therefore choose the alignment with the broadest plateau and re-check correct alignment by taking a few data points around this location every ~ 30 min. Note that if a mask with a spatial structure in two dimensions is used, the offset along the second lattice axis needs to be calibrated as well (as is the case for the 2D state engineering below).

A second application of the local tilts is the realization of local Z rotations. To demonstrate this, we employ a Ramsey-type sequence, where we embed a local Z pulse between two global X pulses (Fig. 5.10a). Starting again with single particles in isolated double wells, initialized in $|L\rangle$, a first X pulse of variable duration creates a superposition of $|L\rangle$ and $|R\rangle$. We then pulse on the illumination intensity of the DMD for a variable duration (150 μs ramp duration) using the same mask as before, which applies a local Z_ϕ rotation to every other double well along the superlattice axis. The imprinted relative phase shift between the two double well classes is revealed by a second $X_{\pi/2}$ pulse. As shown in Fig. 5.10b, varying the pulse duration of the first X pulse results in high-contrast imbalance oscillations. The difference between the oscillation phases of the two double well classes directly corresponds to the imprinted relative phase, and can be extracted from the fits with high accuracy. For this example, we applied a Z_π rotation, which is confirmed by the fits as a phase difference of $\phi = 1.04(2)\pi$. In correspondence with the previous results, a spatially resolved analysis of this experiment proves a high spatial homogeneity also for pure Z rotations imprinted locally via the DMD (Fig. 5.10c).

The hereby demonstrated local control of the orbital operations constitutes a powerful tool to probe quantum states. By enabling the read-out of multiple orbital observables on different bonds in the system within a single experimental realization, it opens up the possibility to measure more complex observables, such as for example correlators between current and

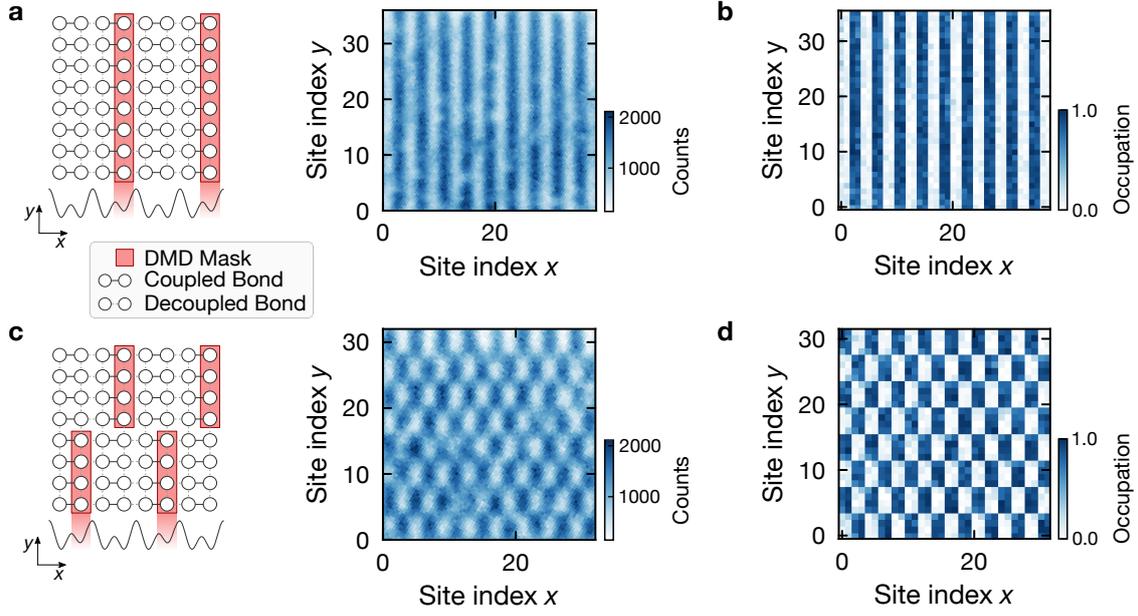


Figure 5.11 | State engineering. **a**, State engineering using locally detuned double well oscillations. The programmed mask is shown on the left, with the red shading marking the sites that are lifted in energy. On the right, a single fluorescence image of the resulting state is shown, consisting of alternating pairs of empty and occupied rows. **b**, Average reconstructed occupation for the state in (a) over 30 repetitions. **c**, State engineering using a mask (shown on the left) with a spatial variation in the direction perpendicular to the superlattice. This results in a checkerboard-like state of 4×2 site blocks (single fluorescence image shown on the right). **d**, Average reconstructed occupation for the state in (c) over 30 repetitions.

kinetic energy. Furthermore, it can be used to enhance measurement sensitivity through multi-quadrature measurements [327].

Besides the measurement capabilities, the local control can also be used for the precise, coherent engineering of spatially structured initial states. To demonstrate this, we employ locally detuned X rotations as presented in Fig. 5.8. We start from single particles in isolated double wells initialized in $|L\rangle$, which can be expressed along the superlattice axis by the occupation string

$$|\psi_i\rangle = |\dots 10101010 \dots\rangle. \quad (5.11)$$

Here, each digit denotes the occupation of a short lattice site, i.e., considering a single double well, $|L\rangle = |10\rangle$ and $|R\rangle = |01\rangle$. As a first example, we use the same mask as before (see Fig. 5.11a). We then set the DMD intensity such that a maximum in the detuned imbalance corresponds to a minimum in the bare imbalance during the X rotation, and fix the pulse duration to this point (e.g., around $t_X = 1.3$ ms in Fig. 5.8b). At this point, we have coherently transferred the initial state

$$|\psi_i\rangle = |\dots 10101010 \dots\rangle \rightarrow |\psi_f\rangle = |\dots 01100110 \dots\rangle, \quad (5.12)$$

which corresponds to a period-two charge density wave. Since the employed mask is trans-

lationally invariant in the direction perpendicular to the superlattice, we obtain a 2D state consisting of alternating pairs of empty and occupied columns (see Fig. 5.11a). Analyzing the reconstructed occupation shown in Fig. 5.11b, we find an average filling of 86(4) % in the occupied, and 7(2) % in the empty stripes. Within the error bar, this is consistent with the initial state quality, suggesting a high fidelity of the combined global and local operations.

To highlight the flexibility of this method, we add further spatial structure to the programmed DMD mask in the direction perpendicular to the superlattice axis. As shown in Fig. 5.11c, keeping the pulse sequence unchanged, this results in a checkerboard-like state of (4×2) site blocks. The reconstructed occupation in Fig. 5.11d shows an average filling of 84(5) % in the occupied, and 8(3) % in the empty blocks, which is consistent with the previous result and the initial state quality. The achievable preparation fidelity and robustness is mostly limited by the resolution and alignment of the DMD imaging system, which is particularly challenging for the small lattice spacing of 383.5 nm used in this experiment, as well as the available DMD power. For example, we found that the preparation fidelity is reduced when the mask is programmed with a higher spatial frequency in the perpendicular direction, e.g. for preparing a state with a (2×2) or (1×2) structure. In experiments that employ a larger lattice spacing or achieve a higher imaging resolution, high preparation fidelities can be expected even for small structures. Also, it is noteworthy that extended structures such as those we used here are advantageous for the achieved tilts, as interference effects between neighboring pixels of the DMD chip increase the intensity. When isolated tilts should be applied using this scheme, one likely needs significantly higher illumination intensities.

5.5 Readout of correlation functions

In this section, we will demonstrate that the single-shot sampling capability of our protocol can be used to measure correlation functions of orbital observables, and in particular current-current correlations. As a direct application, we use this technique to characterize the complex phase pattern imprinted by the running-wave lattice used to engineer artificial gauge fields in our experiment. As described in Section 1.3, the running-wave modulation scheme gives rise to a complex tunnel coupling with a spatially varying coupling phase. To probe this phase pattern, we prepare the ground state of a 2D lattice of isolated double wells, each filled with one particle, in presence of the running-wave modulation. The sequence to prepare this state starts from an initial state where each double well is initialized with one particle in $|L\rangle$ in presence of a strong tilt $\Delta \approx h \times 4.2$ kHz. Next, we slowly decrease the short lattice in x -direction to $9 E_{r,s}$ over 100 ms, which results an intra double well tunnel coupling of $J \approx h \times 650$ Hz, but the particle remains mostly localized in $|L\rangle$ since $\Delta \gg J$. In a final step, we adiabatically ramp up the depth of the running-wave lattice to induce a coupling magnitude of $K = h \times 150$ Hz, where the frequency difference of the individual laser beams is set to be in resonance with the double well

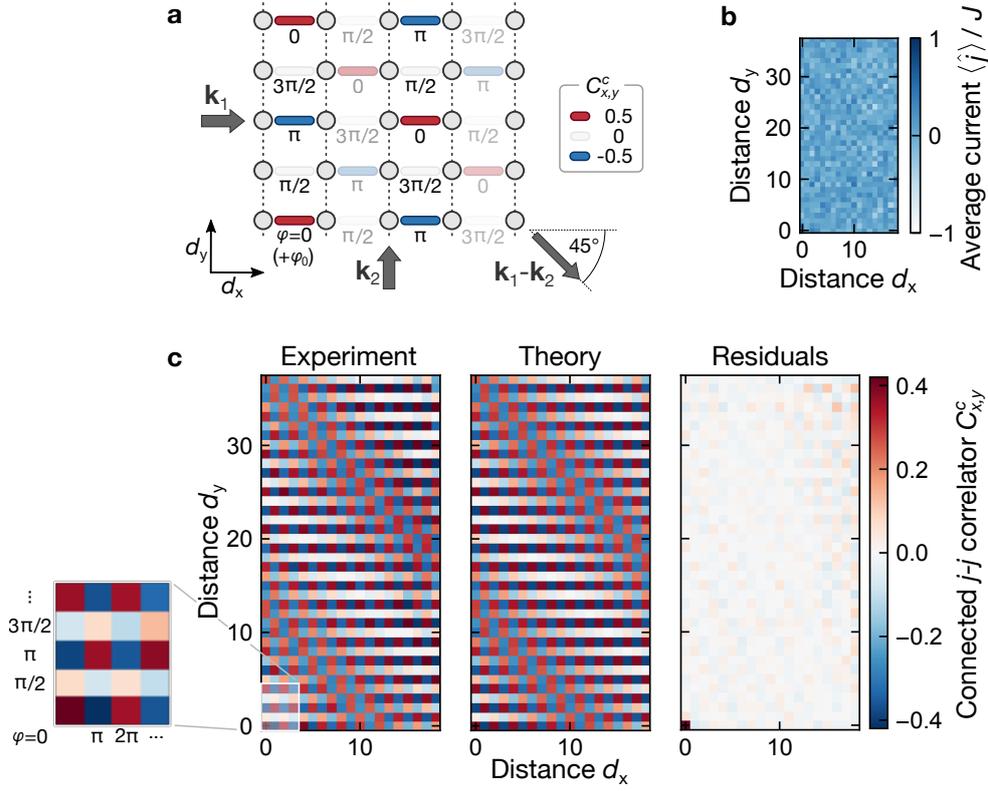


Figure 5.12 | Readout of current-current correlations. **a**, Using a running-wave lattice, we imprint a spatially varying phase pattern onto a 2D array of isolated, single particle double wells. The phases below each bond indicate the local bond phase, and the colored bonds indicate the expected current-current correlations relative to the bond at the origin. **b**, Averaged, spatially resolved current signal computed from 50 repetitions, showing a vanishing average current. **c**, Measured 2D connected current-current correlator evaluated from 50 averages.

energy gap $\sqrt{\Delta^2 + 4J^2}$. At this point, the system is in the ground state of each Floquet-dressed double well. The state of a double well at location (i, j) in the lattice is described by

$$|\psi_{i,j}\rangle = \frac{1}{\sqrt{2}} (|L\rangle + e^{i\varphi_{i,j}} |R\rangle), \quad (5.13)$$

where $\varphi_{i,j}$ is the phase of the complex tunnel coupling. The expected phase pattern is shown in Fig. 5.12a, where the phase increases by $\pi/2$ per bond in both spatial directions. To measure this phase pattern, a measurement in the density basis is not sufficient, as the density difference $\langle \hat{\sigma}_z \rangle$ vanishes identically for all bonds in the system. Instead, we can reveal the phase pattern by measuring the current operator, where the expectation value is directly related to the local bond phase via

$$\langle \hat{j}_{i,j} \rangle / J = \sin(\varphi_{i,j}). \quad (5.14)$$

To implement the current measurement experimentally, we rotate the measurement basis by applying a global $X_{\pi/2}$ pulse as described above. In preparation for the basis rotation, we

suddenly switch off the running-wave lattice and isolate the wells by increasing the short lattice depth to $30 E_{r,s}$ over $500 \mu\text{s}$. At the same time, we move the superlattice phase to the symmetric point using the feed-forward technique. Immediately after, the $X_{\pi/2}$ pulse is applied and the density is read out. The resulting spatially resolved current signal is shown in Fig. 5.12b, which, contrary to the naive expectation, shows a vanishing average current. Interestingly, a measurement of average currents in the system will not work due to the running-wave modulation. The reason for this lies in a global phase offset, which is added to all bond phases and fluctuates randomly between different experimental cycles, since the absolute phase of the running-wave lattice is not stabilized. Hence, the current signal is washed out when averaging over experimental repetitions. This is also the reason why a direct experimental measurement of such a phase pattern was not possible in previous experiments.

To overcome this limitation and reveal the phase pattern, we evaluate the connected current-current correlator, which is defined as

$$C_{x,y}^c = \langle \hat{j}_{i,j} \hat{j}_{i+d_x, j+d_y} \rangle - \langle \hat{j}_{i,j} \rangle \langle \hat{j}_{i+d_x, j+d_y} \rangle. \quad (5.15)$$

The connected correlator is insensitive to the global phase offset, as it directly probes the phase differences in the system. For the given phase pattern, the expected current correlations are indicated by the colored bonds in Fig. 5.12a. The experimentally measured 2D correlation function is shown in the leftmost panel of Fig. 5.12c. Focusing on small distances (see inset), we find a clear oscillatory pattern that is in excellent agreement with the expected phase pattern. For larger distances, we additionally find the appearance of Moiré patterns, which originate from slight angular misalignments between the lattice base vectors and the running-wave lattice. As a result of these deviations, the experimental phase pattern does not evolve exactly by $\pi/2$ between neighboring bonds. To model the observed current correlations, we compute the expected phase pattern including a set of angles allowing for deviations from the ideal beam geometry. To this extent, we express the location of a lattice site with 2D index (i, j) as

$$\mathbf{R} = ia_s \mathbf{e}_x + ja_s \mathbf{e}_y, \quad (5.16)$$

where a_s is the lattice constant of the short-period lattice and $\mathbf{e}_{x,y}$ are the unit vectors along the two axes. Additionally, we express the wave vectors of the running-wave modulation beams as

$$\mathbf{k}_1 = k_R [\cos(\theta_c - \theta_r) \mathbf{e}_x + \sin(\theta_c - \theta_r) \mathbf{e}_y] \quad (5.17)$$

$$\mathbf{k}_2 = k_R [-\sin(\theta_c + \theta_r) \mathbf{e}_x + \cos(\theta_c + \theta_r) \mathbf{e}_y], \quad (5.18)$$

where $k_R = \pi/(2a_s)$ is the magnitude of the wave vector of the running-wave, θ_c is a common-mode angle and θ_r a relative angle between the running-wave beams. This allows us to capture angular deviations from optimal alignment. The phase distribution $\varphi_{(i,j)}$ is then given by

$$\varphi_{(i,j)} = (\mathbf{k}_1 - \mathbf{k}_2) \cdot \mathbf{R} + \varphi_0, \quad (5.19)$$

where φ_0 is a global phase offset that is random within $[0, 2\pi)$ in every shot. To compare this model with the experimental observations, we fit the experimentally measured 2D correlator

to the 2D connected correlator of the sine of Eq. (5.19) multiplied with an amplitude factor. Additionally, we exclude the auto-correlator for the fit, as it is significantly stronger in the experiment, owing to the discretely distributed imbalance samples.

As shown in the middle and rightmost panels of Fig. 5.12c, we find excellent agreement between our theoretical model and the experimentally resolved phase distribution. The fit yields an amplitude of $A = 0.777(3)$, a common-mode angle of $\theta_c = 1.802(3)^\circ$ as well as a relative angle of $\theta_r = -0.558(8)^\circ$, where the uncertainties are the standard errors of the fit. We believe that the amplitude is mostly limited by initial state infidelities, an imperfect adiabatic ground state preparation, as well as detection infidelities. Note that an angular misalignment can in principle also change the engineered magnetic flux. However, the flux per plaquette is not significantly affected by the Moiré pattern, as it depends only on the phase difference between adjacent bonds in the y direction. In our case, the flux is changed from $\pi/2$ to around 0.51π .

The capability to measure correlation functions enabled us hence to, for the first time, take a glance at the spatially varying phase pattern that gives rise to artificial gauge fields engineered via laser-assisted tunneling. Besides this, such a measurement enables a very precise determination of the relative angles θ_r and θ_c and, in turn, to optimize the beam pointing and homogenize the phase distribution.

5.6 Conclusion

In this chapter, we have presented a technique to read out and control orbital observables such as the current and the kinetic energy with local resolution and within a single experimental realization. To do this, we extended previous ideas on atomic beam splitter operations [34, 35, 92, 93, 316, 317], combining basis rotations in superlattice-created double wells with the single-site resolution of our quantum gas microscope. The optical superlattices play a crucial role, as they allow us to partition large 2D systems into high-quality double well arrays and to apply global operations in a parallel fashion. This renders the presented protocol robust, high-fidelity and scalable to large system sizes, paving the way to establish it as a standard technique in the toolbox of quantum simulation. In doing so, we bridge a long-standing gap in quantum gas microscopy, where measurements were limited to local occupations in the lattice, now enabling also a measurement of coherences.

A locally-resolved measurement of the current operator will enable novel studies of strongly interacting topological phases with equilibrium currents [90, 91, 165], as we do in the following Chapter 6. Based on currents and current correlation functions, non-equilibrium dynamics and transport can be studied in unprecedented ways. For example, transport coefficients can be extracted from current correlation functions in the linear response regime [310, 328, 329], and the decay of current correlations can be used to identify universal behavior in quantum critical systems [330]. Note also that it is in principle possible to access expectation values of the form $\langle a_i^\dagger a_j \rangle$ for two non-nearest-neighbor sites i, j by concatenating multiple rotations with different double-well partitions (i.e., a π phase shift), or by interfering two identical copies of a 1D system [309].

Combining density, current and kinetic energy readout, a measurement of the full single

particle density matrix is realized. This has important applications in the context of quantum state tomography [331–336], measurements of the (many-body) energy spectrum [337, 338] and Hamiltonian learning [311, 313–315], which is becoming increasingly important with the advent of neutral-atom based analog and digital quantum computing platforms. In this context, the presented superlattice rotations can also be used as building blocks in lattice-based quantum computing protocols [323], for quantum chemistry [96], and for hybrid quantum computing approaches such as variational algorithms [95, 339–342]. Lastly, by combining the presented technique with an additional spin degree of freedom [316, 321], one can gain access to novel observables useful for studying low-temperature fermionic phases and superconductivity, such as the singlet fraction [343] and d -wave pairing correlations [94].

Besides global measurements of orbital observables, we also demonstrated the application of spatially-varying operations based on the local application of Z -rotations using digital micromirror devices. A local rotation of the measurement basis can be used for example to enhance the measurement sensitivity for metrology applications [327], as well as to access further nontrivial observables [54, 323] and correlators, e.g., between current and kinetic energy. Moreover, local manipulations can be used to engineer tailored initial states to study initial-state dependent thermalization in the presence of constraints, e.g., in the context of Hilbert-space fragmentation [344–346] or lattice gauge theories [347–351]. An important feature distinguishing our protocol from previous state engineering techniques, for example by energetically blocking or blowing out specific sites, is the use of coherent orbital operations. This allows us to engineer states with local coherent superpositions, i.e., states with a local current or a local kinetic energy, which was so far out of reach.

Future work concerning the method itself will involve the extension by composite pulse sequences [352–354] with the goal to optimize the fidelity, in particular of the Z -rotations. Additionally, we so far only considered the situation where the on-site interaction energy is switched off during the basis rotation, which can be achieved for example through the use of a Feshbach resonance. To probe many-body systems, it would be advantageous to alleviate this requirement through further studies on the influence of finite on-site interactions and development of sequences to compensate for them. First theoretical steps in this direction were already taken in Refs. [94, 355]. Lastly, the presented protocol is broadly applicable also to other platforms beyond cold atoms in optical lattices, with related work recently demonstrated on superconducting circuits [98] and Rydberg arrays [97].

CHAPTER 6

Strongly-interacting Meissner phases in bosonic flux ladders

In this chapter, we present the experimental realization of strongly-interacting bosonic phases in presence of a Floquet-engineered artificial gauge field on optical ladder systems. We will start by motivating why such flux ladders are an ideal platform to study orbital physics in the presence of strong interactions, and present our experimental implementation based on optical superlattices. Subsequently, we will discuss the extension of the previously presented local measurement of currents to bonds with a complex coupling element. This is benchmarked using isolated plaquettes with two particles, which also poses the first experimental measurement of particle currents with full spatial resolution in a cold atom system. We will then discuss the preparation of many-body ground states on extended ladder systems and present our results on the realization of strongly-interacting Meissner phases. Finally, we will show how the effective temperature can be benchmarked and discuss the prospects of reaching further phases in future experiments.

6.1 Flux ladders as a platform to study orbital physics

The microscopic study of interacting topological matter in the controlled environment of a quantum simulator has been a long-standing goal in the field [109, 110, 356]. Despite remarkable progress, a large-scale simulation of such systems remains highly challenging across various platforms, primarily arising from small energy scales as well as experimental challenges, e.g. resulting from the need to engineer a synthetic gauge field. It is hence crucial to identify minimal platforms that can capture the essential physics while being experimentally accessible. In this context, confining neutral atoms in a two-leg ladder geometry with a synthetic magnetic field (so-called *flux ladders*) realizes a minimal system to study the interplay of an orbital magnetic field and strong interactions. Two-leg ladder potentials can be readily implemented using optical superlattices, which simultaneously forms an ideal basis for the Floquet-engineering of an artificial gauge field as introduced in Section 1.3. During the past two decades, flux ladders have received considerable theoretical attention, with studies revealing an extraordinarily rich phase diagram that warrants exploration in its own right [165, 357–

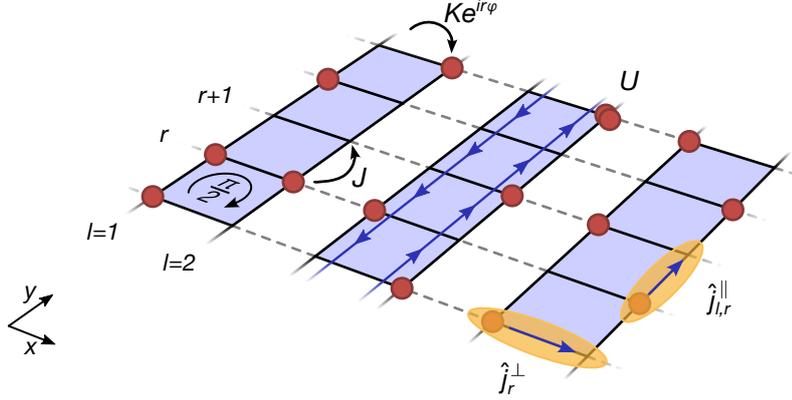


Figure 6.1 | Experimental implementation of the Hofstadter-Bose-Hubbard model. Using a superlattice in x -direction and a monochromatic lattice in y -direction combined with a synthetic gauge field, we realize the Hofstadter-Bose-Hubbard model on two-leg ladders. It is characterized by a real-valued tunnel coupling J along the leg direction, a complex-valued tunnel coupling $Ke^{ir\varphi}$ along the rung direction, and an on-site interaction energy U . Here, r denotes the rung index, and $l = 1, 2$ indexes the legs. Due to the synthetic magnetic field, persistent particle currents emerge in the system (middle ladder, blue arrows), which we can measure with local resolution.

366]. In the following, we will give an overview of the exotic physics that can be studied in optical flux ladders.

A system of bosonic particles on a two-leg ladder with a magnetic flux is described by the Hofstadter-Bose-Hubbard (HBH) model [144, 145]

$$\hat{\mathcal{H}} = \sum_{l,r} \left[-J \left(\hat{a}_{l,r+1}^\dagger \hat{a}_{l,r} + \text{h.c.} \right) + \frac{1}{2} U \hat{n}_{l,r} (\hat{n}_{l,r} - 1) \right] - K \sum_r \left(e^{ir\varphi} \hat{a}_{2,r}^\dagger \hat{a}_{1,r} + \text{h.c.} \right), \quad (6.1)$$

where $\hat{a}_{l,r}^\dagger$ and $\hat{n}_{l,r}$ are the bosonic creation and particle number operators for site $l = 1, 2$ of the r -th rung, J and K are the tunnel couplings in the leg and rung direction, respectively, U is the on-site interaction strength and φ denotes the Peierls phase (see Fig. 6.1). In the chosen gauge, the Peierls phase distribution leads to a flux ϕ per plaquette of $\phi = \varphi$, which in our experiment is fixed to $\phi = \pi/2$ (i.e., a quarter flux quantum). This Hamiltonian hosts a variety of different ground-state phases, depending on the coupling ratio K/J , the interaction strength U , the magnetic flux ϕ as well as the filling fraction $n = \langle \hat{n}_{l,r} \rangle$ [165, 357–366]. A central feature uniting all these phases for $\phi \neq 0$ are persistent particle currents that emerge from the presence of the synthetic gauge field. The operators describing such currents along the leg ($\hat{j}_{l,r}^{\parallel}$) and rung (\hat{j}_r^{\perp}) direction (see Fig. 6.1) are given by [90, 165]

$$\hat{j}_{l,r}^{\parallel} = iJ \left(\hat{a}_{l,r+1}^\dagger \hat{a}_{l,r} - \hat{a}_{l,r}^\dagger \hat{a}_{l,r+1} \right) \text{ and} \quad (6.2)$$

$$\hat{j}_r^{\perp} = iK \left(e^{-ir\varphi} \hat{a}_{1,r}^\dagger \hat{a}_{2,r} - e^{ir\varphi} \hat{a}_{2,r}^\dagger \hat{a}_{1,r} \right). \quad (6.3)$$

As motivated in Chapter 5, the orthogonality between current and density bases prevents

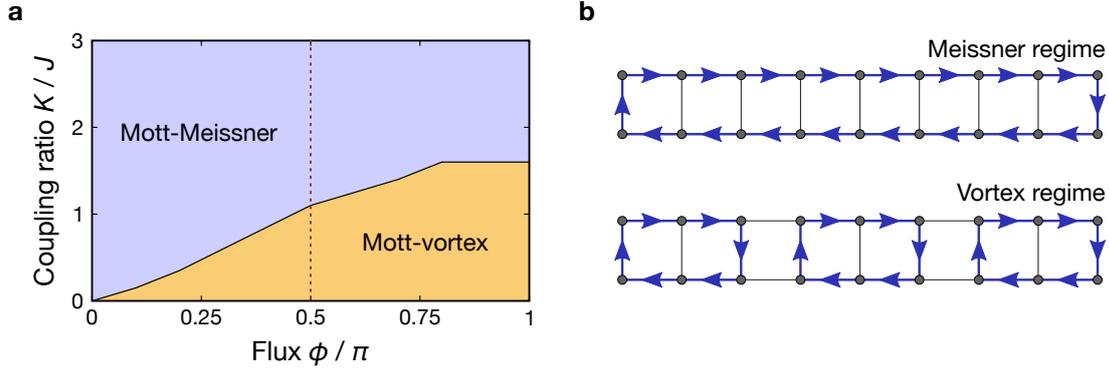


Figure 6.2 | Phase diagram and current profiles. **a**, Illustration of the phase diagram of the ladder-HBH model as a function of the coupling rate K/J and the flux per plaquette φ/π , as an example for half filling ($n = 0.5$) and hard-core bosons ($U \rightarrow \infty$). For soft-core or non-interacting bosons, the phase boundary can shift, but the phase diagram remains qualitatively the same. The dashed vertical line denotes the experimentally realized flux of $\pi/2$. **b**, Schematic current profiles in the Meissner and vortex regimes. The Meissner regime (upper drawing) is characterized by homogeneous chiral currents along the legs, and vanishing rung currents. In the vortex regime (lower drawing), the currents split up into several loops, separated by current vortices. As a result, the currents along the legs are modulated, and the rung currents are non-zero.

a detection of particle currents using standard fluorescence imaging. However, we can use our previously demonstrated double-well rotations to reveal these currents with single-bond resolution, which will be our main observable to probe the ground-state phases of the ladder-HBH model. Note that the leg current operator $\hat{j}_{l,r}^{\parallel}$ depends on the real-valued tunnel coupling J , and can be directly measured using the protocol presented in Chapter 5. In contrast, the rung current operator \hat{j}_r^{\perp} contains the complex-valued rung coupling $Ke^{ir\varphi}$, which requires an adaptation of the measurement protocol to account for the gauge field. We will discuss this in detail in Section 6.2.

The different ground-state phases of the HBH model can be classified according to characteristic spatial configurations of the particle currents. In particular, there are two main regimes across all fillings and interaction strengths: the *Meissner regime* and the *vortex regime*. As shown in Fig. 6.2a, the Meissner regime is reached below a critical flux $\phi < \phi_{\text{cr}}$, or conversely, for a fixed flux above a critical coupling ratio $(K/J) > (K/J)_{\text{cr}}$. It is characterized by homogeneous chiral currents along the legs, together with vanishing currents along the rung bonds (see Fig. 6.2b, upper drawing). In contrast, the vortex regime is present above a critical flux or below a critical coupling ratio. In the vortex regime, the macroscopic Meissner current loop breaks up into several smaller loops, separated by current vortices. As a result, the currents along the legs are modulated, and the rung currents are non-zero (see Fig. 6.2b, lower drawing). The names of these phases are inspired by an analogy to the circular surface currents in type-II superconductors, which fully screen an external magnetic field up to a critical field strength (Meissner effect), above which the currents break down into an Abrikosov vortex lattice [93, 367]. Qualitatively, the distinction into Meissner and vortex regimes exists for all interaction strengths, including the non-interacting limit. However, the precise location of the phase

boundary is modified by the presence of interactions, moving from $(K/J)_{\text{cr}} = \sqrt{2}$ [93, 367] for $U = 0$ to $(K/J)_{\text{cr}} \approx 1$ [165] for strong interactions.

In addition to the current configuration, the phase diagram further splits into superfluid and Mott-insulating regions depending on the filling fraction and the interaction strength. For low fillings $0 \leq n < 0.5$, the system is in a superfluid state regardless of the interaction strength, resulting in superfluid-vortex and superfluid-Meissner states. At half filling $n = 0.5$, i.e., one boson per rung, the system undergoes a superfluid-to-Mott-insulator transition at a critical interaction strength that depends on the specific coupling ratio K/J [165]. This results in Mott-vortex and Mott-Meissner states, which are characterized by a homogeneous density plateau, with superimposed current patterns as shown in Fig. 6.2. These Mott-insulating states at half filling are special in that they differ from the integer-filling Mott insulators that we know from the standard Bose-Hubbard model. Here, the repulsive interaction in combination with the rung coupling stabilizes a state where we have exactly one delocalized particle per rung, with the addition of a second particle to a rung penalized by a Mott gap. Due to the fractional filling, such states are also called *fractional Mott insulators* [368, 369]. Compared to weakly-interacting superfluid phases, strong on-site interactions suppress the currents with a characteristic scaling [165], which can be experimentally probed to study the phase diagram. Lastly, higher fillings $0.5 < n \leq 1$ are trivially related to the low-filling regime by particle-hole symmetry [165].

For completeness, we mention a few other interesting phases that emerge for specific parameter regimes. These are beyond the scope of the current work but constitute interesting prospects for future experiments. In particular, for quarter filling $n = 0.25$ and large fluxes $\phi \gtrsim 0.8\pi$, a fully gapped charge-density wave (CDW) phase emerges for strong interactions ($U/J \gtrsim 30$), where the homogeneous density plateau of the above Mott phases is broken up to form a period-two density modulation [165, 359]. Furthermore, there is a biased-ladder phase (BLP) predicted for intermediate interaction strengths $U/J \sim 2$ and fillings $n \lesssim 1$ [359, 370]. In the BLP, the \mathbb{Z}_2 inversion symmetry of the legs is spontaneously broken, resulting in a density imbalance between the two legs.

Connection to topologically ordered matter and quantum Hall states

While the ladder models studied here pose minimal quasi-1D implementations of the Hofstadter model, many of the signatures are closely related to their counterparts found in extended, two-dimensional quantum Hall systems. For the non-interacting case, Ref. [371] established an analytic correspondence between a ladder and a 2D Hofstadter model. In particular, by partitioning the 2D Hofstadter lattice into a sequence of two-leg ladders, they showed that the eigenstates of the ladder are exactly equal to the topological chiral edge states of the parent 2D model. This is a remarkable result, indicating that the chiral currents, which we can probe experimentally in a ladder system, are a direct consequence of the physics at the chiral edges of 2D Chern insulators.

For interacting systems, the precise connection between states on a two-leg ladder and the full-2D HBH model, in particular fractional quantum Hall (FQH) states, is less straightforward and depends on the specific parameters within the much richer phase diagram. Noteworthy,

several theoretical studies established a connection between the bosonic $\nu = 1/2$ Laughlin state and a similar, Laughlin-like state on a two-leg ladder [163, 372–375]. This correspondence is made via a coupled-wire picture, which already provided an intuitive understanding for the emergence of FQH states in 2D solid-state systems [376, 377]. Earlier studies in this regard found that such a Laughlin state is most likely stabilized under an additional rung-wise repulsive interaction [163, 372, 373], as it can be for example realized in a flux ladder implementation with synthetic dimensions. A later study found that a Laughlin-like state is also stable with on-site interactions only (i.e., exactly the setup in Eq. 6.1), in a small parameter regime between vortex and Meissner phases [374]. It can be distinguished from the surrounding phases by a characteristic chiral current profile, making use of the same techniques as we develop here. Lastly, Ref. [375] found an interesting experimental signature of the Laughlin state when the rung direction is subject to a periodic boundary condition. Here, the ground state exhibits CDW order (as in the Tao-Thouless state [378]), with two degenerate phases of the density modulation. Strikingly, the fractionally charged quasiparticle excitations of the 2D Laughlin state map onto domain walls between the two CDW phases (see also Ref. [379]). This could allow for a direct experimental observation of the fractional excitations despite the simple ladder geometry¹.

State-of-the-art and experimental feasibility

So far, the ground state physics of optical flux ladders have only been studied in the non-, or weakly-interacting limit [93, 174]. For realizations based on Floquet engineering, this is primarily due to the susceptibility of interacting Floquet systems to heating [198–200, 202, 203, 205, 206], which has so far limited experimental implementations to few-particle systems [111, 180, 181] or the study of quench dynamics [380]. Besides Floquet systems, flux ladders have also been realized using synthetic dimensions, where studies have however also been limited to dynamics such as measurements of the Hall response [175, 176, 381–384].

In the following experiments, we will demonstrate that optical flux ladders are an ideal platform to study the complex interplay between interactions and gauge fields, and that the ladder-HBH model can be realized using Floquet engineering with sufficiently low heating to explore the ground state physics. To this end, we will focus primarily on the interacting Meissner regime at half filling, whose finite excitation and mass gap enables us to engineer experimentally feasible adiabatic preparation sequences.

6.1.1 Effective models

To gain an intuition on the expected ground-state behavior across the phase diagram, it is instructive to derive approximate descriptions for the ladder-HBH model in the limiting cases deep in the Meissner and vortex regimes, respectively. We first discuss an approximate model valid in the strong rung coupling limit $K \gg J$, i.e., deep in the Meissner regime. As shown

¹Note however that these excitations are not topologically protected as in a full 2D system. The protection appears only when the number of coupled wires is increased significantly [375].

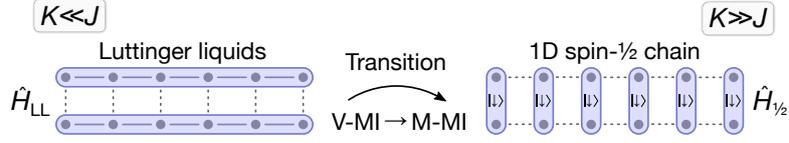


Figure 6.3 | Effective models of the flux ladder at half filling and strong interactions. Effective description of the flux ladder system in terms of two weakly coupled Luttinger liquids ($K \ll J$) and a 1D spin chain ($K \gg J, U \gg J$). The two regimes are separated by the transition between Mott-vortex (V-MI) and Mott-Meissner (M-MI) phases. Note that the spin chain description is only applicable in the strongly interacting limit, while the Luttinger liquid picture applies for any interaction strength.

first in Ref. [165], in the limit of strong interactions $U \gg J$ and $n \leq 0.5$, the system can be described by an effective spin-1/2 model consisting of the two pseudo-spin states

$$|\uparrow\rangle_r = |0, 0\rangle_r \quad (6.4)$$

$$|\downarrow\rangle_r = (|1, 0\rangle_r + e^{i\pi/2} |0, 1\rangle_r) / \sqrt{2}. \quad (6.5)$$

At half filling, the ground state is a product state of rung triplets, i.e., each rung r is in the state $|\downarrow\rangle_r$ (see Fig. 6.3, right drawing). The first excited states have a finite gap, which allows for a perturbative treatment to obtain an estimation of the chiral current, defined as the average difference between the two leg currents, $j_c = \frac{1}{2L} \left| \sum_r \langle \hat{j}_{1,r}^{\parallel} \rangle - \langle \hat{j}_{2,r}^{\parallel} \rangle \right|$. There are three fundamental excitations:

- The addition of one boson to a rung $|1, 1\rangle_r$, or the removal of all bosons from a rung $|0, 0\rangle_r$, costs an energy of K .
- The excitation of one rung r of the ladder to a singlet $(|1, 0\rangle_r - e^{i\pi/2} |0, 1\rangle_r) / \sqrt{2}$ costs an energy of $2K$.
- The addition of a second particle to an already occupied site costs an energy of $\sim U$.

Retaining these three fundamental excitations, one can obtain an approximate relation for the chiral current to second order in J as [165]

$$j_c = \frac{J^2(4K + U)^2}{2KU(2K + U)}. \quad (6.6)$$

This expression shows that the chiral current is suppressed both with increasing interactions U and increasing rung coupling K . In the opposite limit of weak rung coupling $J \ll K$, i.e., deep in the vortex regime, the system can be described in terms of two coupled Luttinger liquids (see Fig. 6.3, left drawing). We do not discuss this limit in detail here, as it is not the focus of our experiments. However, we note that the Luttinger liquids have a vanishing excitation gap, due to which it is expected that current modulations along the leg direction are rather fragile against finite temperature effects.

6.1.2 Experimental implementation

To experimentally realize the HBH model on a two-leg ladder geometry, we combine our bi-chromatic horizontal superlattices with a laser-assisted tunneling scheme based on an optical running-wave to engineer a synthetic magnetic field (see Section 1.3). Specifically, we use a bi-chromatic superlattice in x -direction (short spacing $a_s = 383.5$ nm, long spacing $a_l = 767$ nm) and a monochromatic short spacing lattice in y -direction (a_s), which realizes a chain of double wells along x that are coupled along the perpendicular direction. Using a deep long-period lattice along x , there is negligible tunneling between the double wells, realizing multiple independent two-leg ladders as shown in Fig. 6.1. In plane, the system is confined in a box potential projected via a DMD with a size of 40×40 lattice sites, constraining the dynamics to 20 ladder copies with a chain length of up to 40 sites each. To engineer a synthetic magnetic field, the double wells are initially strongly tilted, realizing a large energy offset between the two legs and suppressing direct tunneling along x . We then superimpose a 1D running-wave lattice (oriented diagonally to the base vectors, spacing 1534 nm/ $\sqrt{2}$), which resonantly restores tunneling between the legs as described in Section 1.3. In an effective Floquet picture, this cancels the energy offset between the legs (i.e., effective symmetric double wells), and realizes a complex-valued tunnel coupling $Ke^{i\varphi(l,r)}$ along the rung direction. Due to the choice of the running-wave geometry, the coupling phase is fixed to increase by $\pi/2$ per bond in the rung direction, resulting in a synthetic magnetic flux of $\pi/2$ per plaquette. In combination with the real-valued tunnel coupling J along the leg direction and a tunable on-site interaction energy U , this setup realizes the HBH model in Eq. (6.1) at a fixed flux of $\pi/2$.

6.2 Measurement of currents on real and complex bonds

As discussed above, a central signature of the HBH model are persistent particle currents in the system, which we can use to characterize the different ground-state phases. To measure these currents with local resolution, we employ our previously introduced measurement scheme based on local rotations of the measurement basis (see Section 5). A measurement of the currents along the leg direction (i.e., of the operator in Eq. 6.2) can be straightforwardly carried out using an $X_{\pi/2}$ rotation as benchmarked in Section 5.3. As shown in detail in Fig. 6.4a, we realize this by suddenly turning on an additional long lattice in the y -direction in 150 μ s to project the leg bonds into symmetrically coupled, isolated double wells. To disable the rung coupling, we simultaneously switch off the running-wave lattice and ramp up the short x -lattice to $35 E_{r,s}$. Additionally, we switch off on-site interactions by ramping the offset field to the non-interacting point and handing over from the steep-angle to the shallow-angle vertical lattice, reducing the vertical confinement from $\omega_z/2\pi = 4.7$ kHz to 800 Hz. The following time evolution under the DW Hamiltonian rotates the measurement basis, enabling a read-out of the current operator after a quarter rotation period. Before reading out the local occupation, the rotation is stopped by ramping up the short y -lattice in 150 μ s to $35 E_{r,s}$. Including the initial and freeze ramps, the measurement sequence has a total duration of 435 μ s at a DW coupling of $J_{DW}/h = 593(5)$ Hz. Note that due to the DW array structure, we straightforwardly measure only every second bond along the leg direction. To obtain full spatial resolution,

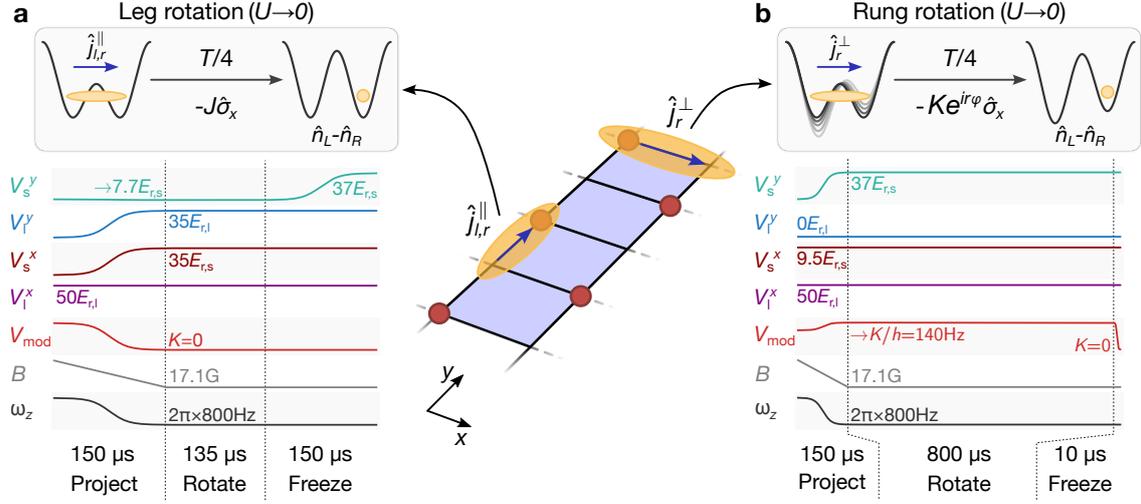


Figure 6.4 | Current measurement scheme and sequence. **a**, To measure the leg currents, bonds along the leg direction are isolated by ramping up an additional long lattice in the y -direction, switching off the rung coupling, and applying an $X_{\pi/2}$ rotation in the double wells to map the current onto the density difference, which can be read out after freezing. **b**, To measure the rung currents, individual rung bonds are isolated by ramping up the short lattice in the y -direction, and subsequently applying an $X_{\pi/2}$ rotation in the driven double wells to map the current onto the density difference, factoring in the complex coupling element. In both cases, interactions are switched off by ramping the offset field to the non-interacting point and handing over from the steep-angle to the shallow-angle vertical lattice.

one can repeat the measurement with a superlattice phase shift of π , which doubles the required number of snapshots.

In contrast to that, a measurement of the currents along the rung direction (i.e., of the operator in Eq. 6.3) requires an adaptation of the protocol to account for the gauge field. If we performed the same rotation in symmetric, real-coupled as for the leg bonds, we would measure an incorrect current operator (Eq. 6.2 and Eq. 5.7) that does not account for the complex coupling element. As a result, we would only measure the trivial, laser-induced phase $\varphi(l, r)$ instead of the emerging ground-state currents. To measure the correct rung current operator as defined in Eq. 6.3, we have to perform the basis rotation in the driven DWs, i.e., in DWs that have the same spatially-varying complex coupling phase as the rung bonds. In detail, the left-right density difference after projection into isolated rung bonds evolves under the Hamiltonian $\hat{\mathcal{H}}_r = -K_{\text{DW}} (e^{ir\varphi} \hat{a}_{2,r}^\dagger \hat{a}_{1,r} + \text{h.c.})$ as

$$\begin{aligned} \hat{n}_{1,r}(t) - \hat{n}_{2,r}(t) = & [\hat{n}_{1,r}(0) - \hat{n}_{2,r}(0)] \cos(2K_{\text{DW}}t/\hbar) \\ & + i (e^{-ir\varphi} \hat{a}_{1,r}^\dagger \hat{a}_{2,r} - e^{ir\varphi} \hat{a}_{2,r}^\dagger \hat{a}_{1,r}) \sin(2K_{\text{DW}}t/\hbar), \end{aligned} \quad (6.7)$$

where K_{DW} is the driven coupling strength during the basis rotation. In analogy to the basis rotation with a real coupling, a rotation for a quarter period [$T/4 = \hbar/(8K_{\text{DW}})$] maps the rung current operator onto the density difference. The experimental sequence to implement this rotation is shown in Fig. 6.4b. Since the periodic modulation is already on during the

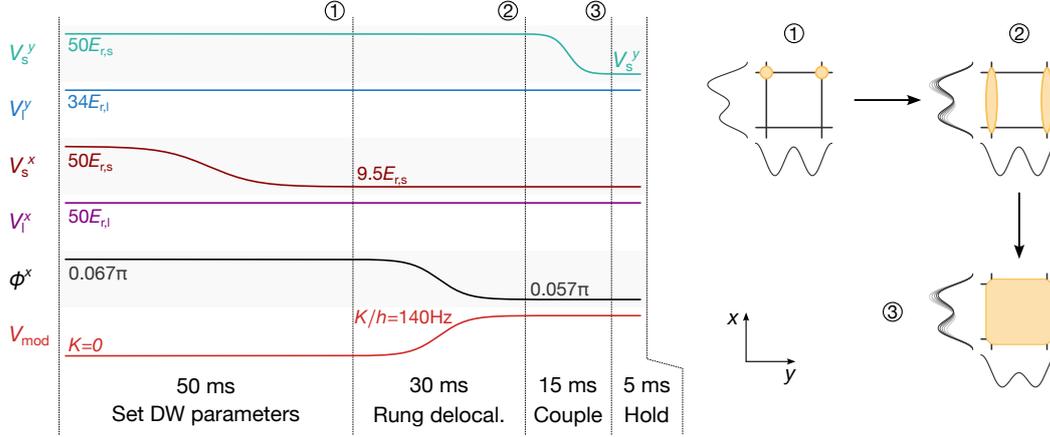


Figure 6.5 | Plaquette ground state preparation sequence. Starting from a CDW initial state in SL axis 2 (x -direction), we adiabatically prepare the ground state in isolated plaquettes by first delocalizing each particle along a rung bond, and subsequently coupling the rung bonds together to form the plaquette.

experiment, the projection step consists mainly of ramping up the short y -lattice to $37 E_{r,s}$ in $150 \mu\text{s}$ to isolate individual rung bonds. Additionally, we ramp the modulation depth to realize a constant coupling strength of around $K_{\text{DW}}/h = 140 \text{ Hz}$ during the rotation irrespective of the coupling set during the experiment, as well as ramp the offset field to the non-interacting point and hand over to the shallow-angle vertical lattice to switch off interactions. An experimental challenge in the Floquet basis rotation is the significantly smaller coupling element, which leads to a longer rotation duration close to 1 ms. Due to this, the rung current measurement is also more sensitive to residual tilts and interactions.

In summary, by extending our local current measurement scheme to basis rotations in Floquet-engineered double wells, these sequences provide access to snapshots of local particle currents on both leg and rung bonds with microscopic resolution, which we will benchmark in the following.

6.3 Currents in isolated plaquettes with two interacting particles

We begin by preparing the ground states of isolated plaquettes with two strongly interacting particles each (i.e., filling $n = 0.5$). Plaquettes pose the absolute minimal system that host persistent currents in presence of a synthetic gauge field. Furthermore, they provide an ideal system to benchmark the current measurement in a many-body phase, as the ground state has stable currents that circulate around all four bonds of the plaquette. At the same time, we can very accurately realize many isolated plaquettes by partitioning our 2D lattice using both perpendicular long lattices into a 2D plaquette array. The preparation sequence is shown in Fig. 6.5. We start from a product state with one particle per rung bond (here, rung and leg refer to the underlying ladder geometry in absence of the long lattice along y), prepared in a single layer of the steep-angle vertical lattice (cf. CDW preparation in Fig. 3.4). The vertical trap frequency is $\omega_z/2\pi = 4.7 \text{ kHz}$, and the offset field is kept constant at 22.3 G to maintain a

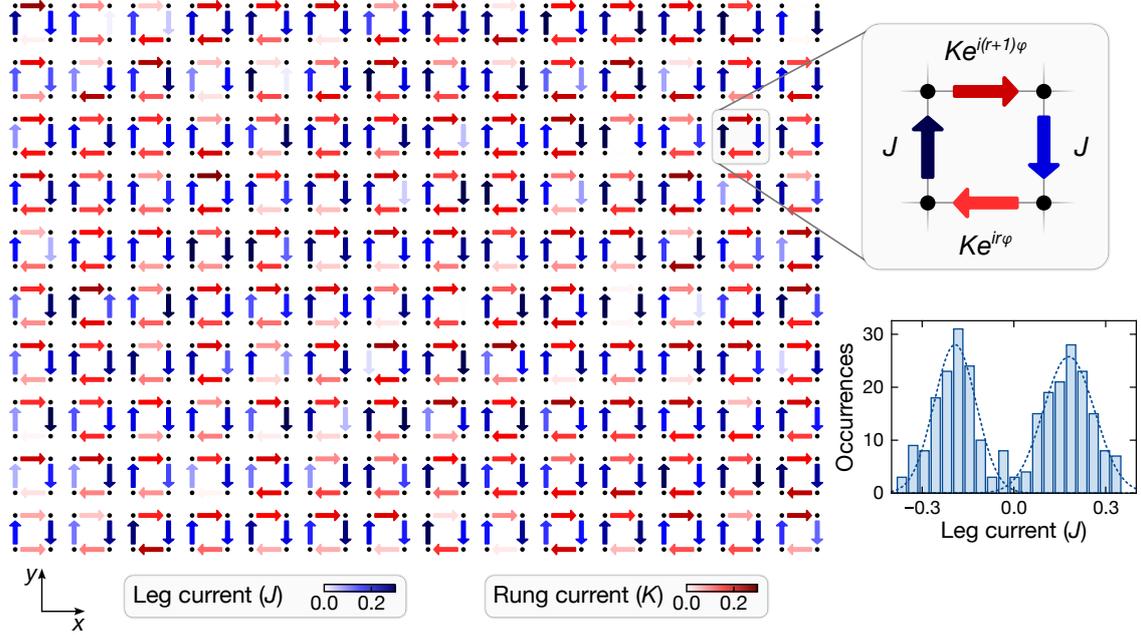


Figure 6.6 | Spatially-resolved map of currents in isolated plaquettes. Map of the ground state currents across a large array of 140 isolated plaquettes for $K/J \simeq 1.4$ and $U/K \simeq 10$. The direction of the current is indicated by the arrow, and the current magnitude is encoded in the color, where the leg currents are shaded in blue and the rung currents in red. The zoom-in shows an exemplary plaquette, indicating the orientation of the real (complex) tunnel couplings on the leg (rung) bonds as defined. The lower right-hand plot shows the distribution of the leg currents across the entire system. The left bonds have a mean current (1σ -deviation) of $0.18(8)J$ and the right bonds $-0.19(8)J$, respectively, as illustrated by the normal distributions (dashed line).

strong, repulsive on-site interaction. Both long lattices in the x - and y -direction are kept deep ($50 E_{r,l}$ and $34 E_{r,l}$, respectively) throughout the sequence to define the plaquette geometry and suppress tunneling between the plaquettes (residual coupling in the leg direction $J'/h < 5$ Hz, in the rung direction $K'/h \approx 1.5$ Hz). In a first step, we adiabatically decrease the depth of the short x -lattice in 50 ms to $9.5 E_{r,s}$, which sets a bare tunnel coupling of $J/h = 810$ Hz along the rung direction, but each particle remains localized in the energetically lower well due to the large tilt of $\Delta/h = 5.1$ kHz. Next, we adiabatically turn on the running-wave lattice in 30 ms at a modulation frequency of $\omega_{\text{mod}}/2\pi = 4.6$ kHz, while simultaneously ramping the superlattice phase to reduce the tilt to $\Delta/h = 4.3$ kHz. The combined modulation-tilt ramp brings the laser-assisted tunneling into resonance, while ensuring that the system remains in the instantaneous ground state. After the ramp, each particle is symmetrically delocalized across its rung bond at a driven coupling strength of $K/h = 140(1)$ Hz. In a final step, the short lattice in the leg direction is lowered over 15 ms to a variable depth, coupling the rung bonds together and transferring the system to the plaquette ground state at a final K/J . We optimized the sequence parameters by inverting the state preparation and maximizing the overlap with the initial state (quantified through atom number and imbalance) similar to Ref. [111]. After a hold time of 5 ms, we measure the ground state currents using the basis rotation sequences described above.

Equipped with the ability to reveal currents with local resolution, we study the spatial

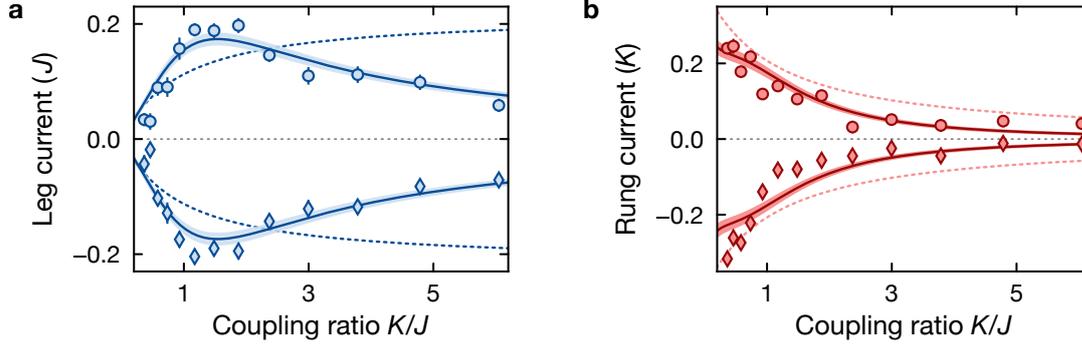


Figure 6.7 | Ground state phase diagram of the plaquette. Dependency of the leg (a) and rung (b) currents on the coupling ratio K/J for isolated plaquettes with $U/K \simeq 10$. The solid lines are fits of an ED simulation of the ideal currents, with the amplitude as a single free parameter, yielding $0.78(4)$ for the leg and $0.71(4)$ for the rung currents; the shaded area denotes the 1σ -confidence interval of the fit. The dashed lines indicate the currents in a non-interacting plaquette with the same fit amplitude. The data points are spatially averaged across 140 plaquettes, and the error bars denote the SEM over 200 snapshots, and if not visible, are smaller than the marker size. The numerical simulations take into account the reduced flux in isolated plaquettes of $0.71(2) \times \pi/2$, which we calibrated using cyclotron orbits as first demonstrated in Ref. [170].

distribution of the ground state currents in a large 2D plaquette array. Fig. 6.6 shows a current map across 140 isolated plaquettes, evaluated in a central sub-region of the full sample, for a coupling ratio of $K/J \simeq 1.4$ and an interaction strength of $U/K \simeq 10$. We observe a homogeneous distribution of currents across the entire system, circulating around all four bonds of each plaquette in the same direction as a result of the homogeneous magnetic flux threading each plaquette. By nature of a measurement with full spatial resolution, which requires a large amount of snapshots, there is some amount of statistical noise in the current distribution. However, a detailed analysis of the current distribution, shown as an example for the leg currents in Fig. 6.6, reveals that the variance is consistent with the expected projection noise at the experimental sampling of 200 snapshots per bond. In detail, the left bonds have a mean current (1σ -deviation) of $0.18(8)J$ and the right bonds $-0.19(8)J$, respectively. The current measurement itself is hence not expected to introduce a substantial amount of excess noise, and a smaller variance can simply be achieved by increasing the number of samples.

While the plaquettes are few-body systems with only two particles, one can already find a strong influence of the inter-particle interactions. To this end, we study the ground-state phase diagram by tracking the behavior of the bond currents as a function of the coupling ratio K/J . The coupling ratio is varied by changing the final depth of the short lattice in the leg direction. As a result, also the interaction strength changes slightly in the range of $U/K = 8.5(3) \dots 11.03(3)$, at an average of $\bar{U}/K = 9.8$. The measured dependency of the leg currents on the coupling ratio is shown in Fig. 6.7a. After an initial rise, we observe a maximum around $K/J \approx 1.5$, followed by a suppression of the currents towards higher K/J . This behavior is in excellent agreement with a numerical simulation based on exact diagonalization of the two-particle plaquette ground state, and we measure around $78(4)\%$ of the ideal current magnitude. The suppression of the currents at high K/J is a direct consequence of the strong on-site interactions, which

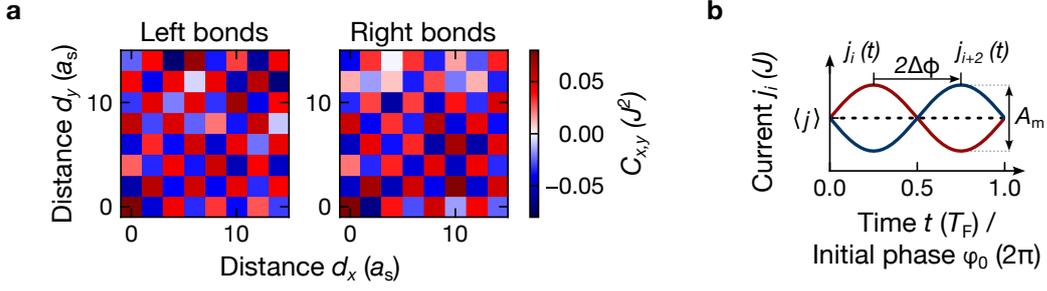


Figure 6.8 | Current-current correlations across different plaquettes. **a**, Measured current-current correlation function $C_{x,y}$ for the leg currents, separated into correlations between all left bonds and all right bonds in the 2D plaquette array. Even though different plaquette copies should be uncorrelated ($C = 0$), the Floquet micromotion causes strong, long-ranged checkerboard correlations with an average magnitude of $0.04(2) J^2$. **b**, Schematic illustration of the evolution of the current operator expectation value within one Floquet period for two bonds of the same type in adjacent plaquettes. Since the initial phase φ_0 randomly fluctuates between experimental realizations, we average over all time points within one Floquet period.

lead to a localization of the particles on the rung bonds and a corresponding suppression of the leg currents. Comparing this to the non-interacting case (dashed lines in Fig. 6.7a), we observe markedly different behavior, where the currents increase monotonically with the coupling ratio. We repeat this for the currents on the rung bonds (Fig. 6.7b), finding similarly good agreement with the numerical prediction at a slightly smaller current magnitude of around 71(4) % of the ideal value. The measured current amplitudes in both cases are likely limited by the finite switching speed of the offset field, causing a residual non-zero U during the basis rotation², as well as non-adiabaticities and Floquet heating during the state preparation. In summary, the large magnitude of the measured currents together with the excellent reproduction of the interacting behavior proves the capability of our local current measurement scheme to probe both types of bond currents in a many-body system with full spatial resolution.

The single-shot sampling of the current operators furthermore allows us to measure current-current correlation functions, as previously demonstrated outside a physical system in Section 5.5. Applying this for the first time to a physical system using the example of the plaquette states, this reveals an important implication arising from the underlying Floquet scheme. In detail, we use the measured plaquette ground-state currents (Fig. 6.6), but instead of evaluating the average current, we compute the connected current-current correlation function $C_{x,y}^c = \langle \hat{j}_{i,j} \hat{j}_{i+d_x, j+d_y} \rangle - \langle \hat{j}_{i,j} \rangle \langle \hat{j}_{i+d_x, j+d_y} \rangle$, where $d_{x(y)}$ is the distance along the rung (leg) direction. In principle, individual plaquette copies should not exhibit any mutual correlations. However, as is shown in Fig. 6.8a for the leg bonds, we experimentally find strong, long-ranged correlations across the system. When looking at the left and right bonds individually, we observe a checkerboard pattern, meaning that a specific bond is always anti-correlated with the corresponding bond in the nearest neighboring plaquettes. The origin of these unexpected correlations can be found in the micromotion – the difference between the effective model and

²In Ref. [212], we determined that a jump to the non-interacting point from an up to 10 G higher offset field results in residual scattering lengths below $50 a_0$ after $150 \mu\text{s}$.

the full Hamiltonian during one Floquet cycle (cf. Section 1.3). As illustrated in Fig. 6.8b, the micromotion causes observables such as the current to oscillate around the value given by the effective model. While this does not yet lead to spatial correlations by itself, the phase of the micromotion is determined by the phase of the periodic drive, which varies spatially. Since the drive phase advances by $\pi/2$ per bond, the micromotion of a bond in two adjacent plaquettes (corresponding to $2a_s$ distance) is π out-of-phase, which causes the observed checkerboard pattern in the connected correlator. Note that the micromotion is not visible in averaged observables, as the random initial phase φ_0 in every experimental realization corresponds to probing the micromotion at random times during a Floquet cycle, averaging out the effect. The observed correlations have an average magnitude of $0.04(2)J^2$, which is comparable to correlations expected from a physical signal with a strength of $0.2J$. This has important implications for experiments where correlations are used to probe physical signatures. However, it should be possible to calibrate out these micromotion-induced correlations by averaging over π -out-of-phase bonds before evaluating the correlation function. Furthermore, the strength of the micromotion can be tuned by the parameters of the Floquet drive, as well as the specific choice of observable (e.g., densities are expected to be less sensitive than currents).

6.4 Identification of parameter regimes with low heating rates

To study the low-energy properties of a model of interest, one has to experimentally prepare the many-body ground state. A typical strategy for this is to start from a (ground) state of an initial Hamiltonian that is readily prepared with low temperature in the experimental platform, and then adiabatically transfer the system to the target Hamiltonian and its ground state. Such an *adiabatic preparation* is a highly challenging task by itself, requiring the development of preparation sequences with large spectral gaps, allowing a fast state transfer that is robust against technical heating and noise. In the context of Floquet-engineered systems, this is further complicated by drive-induced heating due to the resonant modulation. Throughout the entire preparation path, one has to avoid heating resonances between different Bloch bands or losses mediated by interactions, which strongly limits the accessible parameter space [204]. In the following, we will present a strategy to experimentally characterize suitable parameter regimes with low heating rates. Based on this information, we can then develop a preparation sequence that allows us to adiabatically prepare many-body ground states of the ladder-HBH model.

In presence of the periodic modulation, there are two predominant heating channels that can limit state preparation and lifetimes: Firstly, atoms can be excited to higher Bloch bands in one of the horizontal lattices (parametric heating), where they experience enhanced tunneling and the transferred kinetic energy causes losses and collisional heating. This is a single-particle effect, which exists also for vanishing interactions or dilute systems. Secondly, in a many-body system with interactions, the modulation can induce many-body excitations and a resonant formation of multiply-occupied sites. These excitations can be lost via two- or three-body loss processes, depositing large amounts of energy into the system.

Due to the separability of the underlying horizontal lattice geometry, band resonances can be investigated independently in the two axes. The axis where the laser-assisted tunneling is

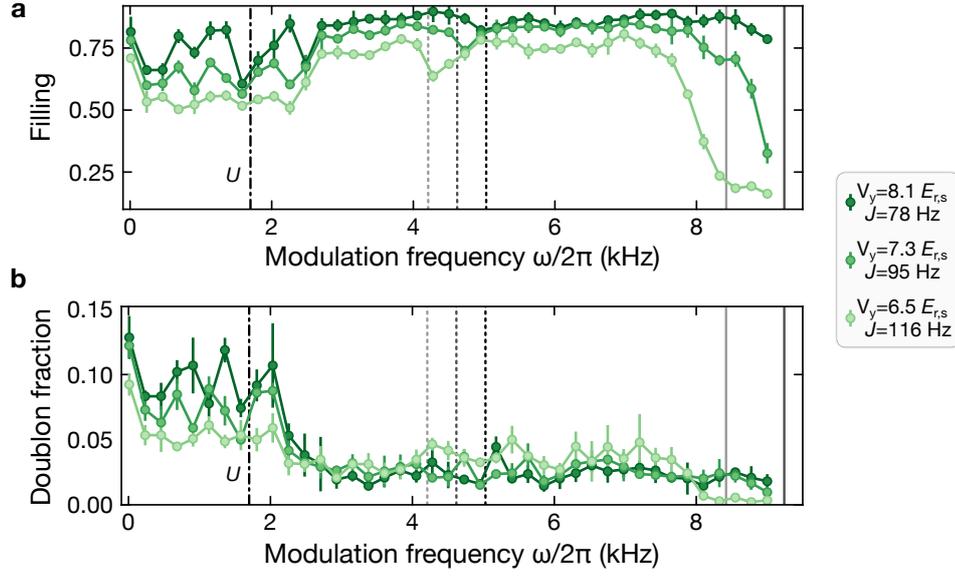


Figure 6.9 | Modulation loss spectroscopy. We probe the location of band resonances in the leg direction as well as interaction effects using modulation spectroscopy in isolated 1D chains for different leg lattice depths V_y , which tunes J along the chain. The offset field is 23 G, corresponding to $U/h \sim 1.7$ kHz. **a**, Average filling after modulating for a total duration of 90 ms. **b**, Average fraction of doubly-occupied sites after the modulation. In both plots, each data point has been averaged across a region of 30×32 sites and 3 repetitions, with the error bars denoting the standard deviation (note that every second chain was empty during the modulation and only used for resolving doublons). The solid traces are guides to the eye. The vertical solid (dashed) lines denote (half of) the transition frequency from the first to the third Bloch band. The vertical dot-dashed line marks the interaction energy U .

implemented (here, SL2 or x) is typically rather easy to fix, as it is operated at a constant depth. To this end, we perform driven double-well oscillations with single particles (see Section 3.6), and monitor the atom number remaining in the sample as a function of the hold time. If unsuitable parameters have been chosen, the atom number decreases with time (on the order of 50 ms), and the lattice depths or the tilt have to be adapted slightly. What we found to be much more restrictive for the experiments is the band structure of the monochromatic lattice along the leg (y) direction. We probe this experimentally in a controlled setting by performing modulation spectroscopy in isolated 1D chains. The chains are at unity filling in presence of strong repulsive on-site interactions, which allows us to simultaneously investigate interaction-mediated losses. Specifically, we start from a CDW prepared using SL2, such that every other chain is empty and can be used later to resolve doublons. Next, we adiabatically lower the short lattice depth along the chain axis to a final J while maintaining a large repulsive on-site interaction $U \gg J$. The running-wave modulation is then turned on for a total duration of 90 ms at a variable modulation frequency and depth. Afterward, we freeze the system again by ramping up the short lattice along the chain direction. Before reading out the density, we apply a doublon splitting sequence using the empty chains to resolve occupations up to $n = 2$ (see Section 3.3.2). Fig. 6.9 shows the results for three different short lattice depths, corresponding to three different tunnel couplings J . Focusing first on the higher frequency range above ~ 4 kHz,

we observe two loss resonances in the average filling (Fig. 6.9a) for each leg lattice depth, one at around 8 – 10 kHz, and one between 4 – 5 kHz. The higher frequency resonances correspond to excitations from the ground to the third Bloch band, which has a large matrix element and is the dominant loss channel. The resonances at half the frequency correspond to two-photon excitations of the transition from the first to the third Bloch band. They have a decreased loss rate due to the smaller coupling strength, but can still limit the lifetime and preparation fidelity at lower leg lattice depths. The vertical lines in Fig. 6.9 indicate the transition frequencies estimated from the calibrated band structure (see Section 3.4), which are in good agreement with the observed resonances. As can be seen, the resonances shift to lower frequencies with decreasing lattice depth, and also become broader due to the increased bandwidth.

In addition to the band resonances, we also observe interaction-mediated loss processes at frequencies in the vicinity of the on-site interaction energy U (here $U/h \sim 1.7$ kHz). While there appear several distinct resonances for the deepest leg lattice, a shallower lattice depth leads to their merging into one broad loss plateau. The interaction-mediated origin of these loss processes also manifests itself in an increased production rate of doubly-occupied sites (Fig. 6.9b). In addition to losses, driving these processes also leads to further unwanted terms in the effective Hamiltonian.

In summary, this modulation spectroscopy provides us with a direct experimental verification of parameter regimes with low heating rates. The strongly-interacting limit, which is of interest for the studies throughout this chapter, render the entire low-frequency range up to around 3 kHz inaccessible. From above, the usable frequency range is bounded by band resonances. Experimentally, we found that frequencies above ~ 5 kHz require tilts where the laser-assisted tunneling does not work well, and hence the remaining usable frequency range is between ~ 4 kHz and ~ 5 kHz. Unfortunately, as we saw in Fig. 6.9a, this also imposes an upper bound on the tunnel coupling J that can be realized in the system. Overall, we found that leg lattice depths between $7.5 - 8 E_{r,s}$ (corresponding to $J/h \approx 80 - 90$ Hz) and a modulation frequency right below the corresponding band resonance ($f_{\text{mod}} \approx 4.6$ kHz) pose an optimal compromise between a high tunnel coupling and sufficiently low heating rates over experimental timescales.

6.5 Strongly-interacting Meissner states

In this section, we will present an adiabatic preparation sequence that allows us to prepare the strongly interacting Meissner regime at half filling in extended flux ladders. Using this sequence, we will characterize the ground state properties of the Meissner phase through local current measurements.

6.5.1 Adiabatic preparation of the interacting Meissner regime

As introduced in Section 6.1, the interacting Meissner regime is obtained when the rung coupling is larger than the leg coupling, $K/J \gtrsim 1$. Within the Meissner regime, no gap closings occur, and the many-body gap increases monotonically with K/J . Out of these considerations,

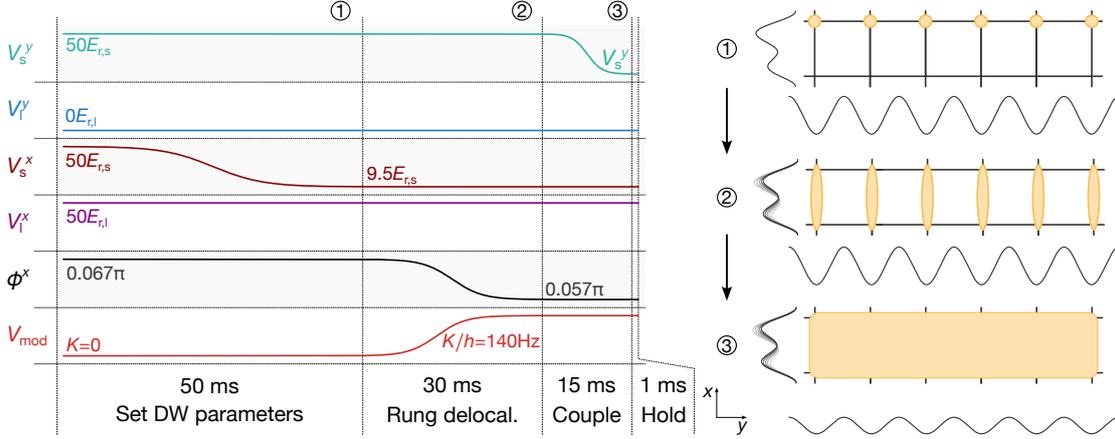


Figure 6.10 | Rung coupling sequence to access the Meissner regime. We start from a CDW in SL axis 2, providing a product state with one particle per rung. By ramping up the running-wave modulation, we delocalize each particle across a rung bond, which corresponds to a Meissner-like state in the $K/J \rightarrow \infty$ limit (for strong interactions). Finally, this state is connected to the Meissner regime at any K/J by ramping down the leg lattice and connecting the rung bonds together.

we devise what we call the *rung coupling sequence*, which provides a fast pathway to the entire Meissner regime starting from experimentally accessible initial states. Specifically, we start from a product state with one particle per rung, which is delocalized across both sites in the presence of the complex tunnel coupling. For the case of strong on-site interactions, this initial state corresponds already to a Meissner-like state in the $K/J \rightarrow \infty$ limit. Hence, it can readily be connected to the remainder of the Meissner phase using a single ramp of the leg coupling to any target $K/J \gtrsim 1$. This sequence has several advantages: Firstly, starting from a product state, we can use low-entropy initial states at a precise, programmable filling. Secondly, the density distribution is already identical to the target one, avoiding the need for a slow mass redistribution. Lastly, the single leg coupling ramp combined with a consistently maintained mass and excitation gap enables a fast preparation that is less sensitive to Floquet or technical heating.

The detailed preparation sequence is shown in Fig. 6.10. Similar to the plaquette initial state preparation, we start from a CDW state along x with one particle in the lower well of each rung, in a single layer of the steep-angle vertical lattice ($\omega_z = 2\pi \times 4.7$ kHz). The offset field is set already while the state is still frozen and kept constant throughout the sequence to realize a target interaction strength in the ladder. After setting the double well parameters for the laser assisted tunneling ($V_s^x = 9.5 E_{r,s}$, $V_l^x = 50 E_{r,l}$, $\Delta/h = 5.1$ kHz, residual coupling between ladders $\sim h \times 1.5$ Hz), the single particle in each rung is delocalized across the rung bond by ramping up the running-wave modulation in 30 ms to a driven coupling strength of $K/h = 140(1)$ Hz (including the extra tilt ramp, modulation frequency $\omega_{\text{mod}}/2\pi = 4.6$ kHz). In a final step, the rungs are coupled together by ramping down the short lattice in the leg direction in 15 ms to a variable depth, which transfers the system to a ladder in the Meissner regime at a final K/J .

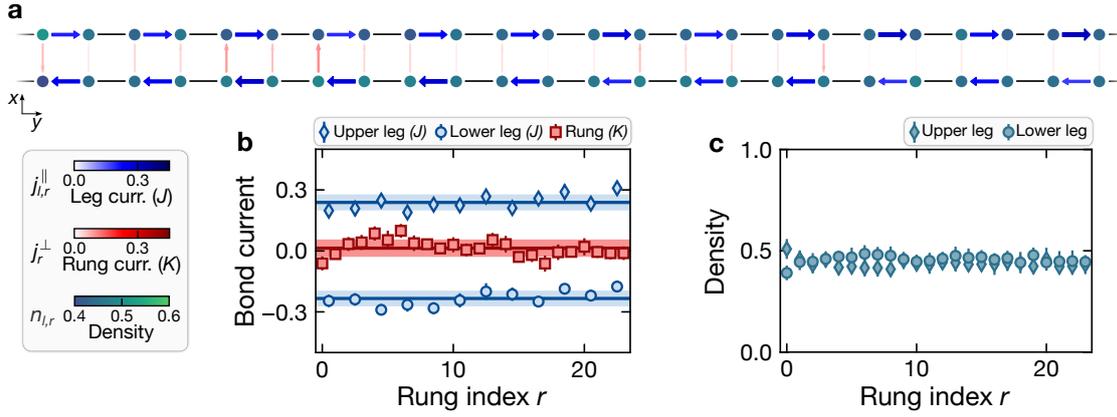


Figure 6.11 | Spatially resolved currents and densities in the Mott-Meissner phase. **a**, Spatially resolved distribution of density, leg current and rung current on ladders with 48 sites in the Meissner regime for $K/J = 1.98(5)$ and $U/J = 11.02(5)$. The width and the color of each arrow is given by the average magnitude of the respective bond current. **b**, Locally-resolved bond currents for the state in (a). The average currents are $0.24(4)J$ on the upper leg, $-0.23(3)J$ on the lower leg, and $0.01(4)K$ on the rungs, respectively, indicated also by the horizontal lines. **c**, On-site densities for the state in (a), yielding a homogeneous density profile with an average density of $0.45(2)$. In all subplots, each data point was averaged over 140 repetitions and 14 ladders, and the error bars denote the SEM over the 140 repetitions (if invisible, they are smaller than the marker size).

6.5.2 Chiral currents in the Meissner phase

Using the rung coupling sequence, we prepare a ladder in the Meissner regime and probe the current and density distribution with local resolution. In Fig. 6.11, this is shown for a Meissner state with 48 sites at half filling and $K/J = 1.98(5)$, $U/J = 11.02(5)$. To avoid edge effects as a result of the finite wall sharpness of the DMD potential, we restrict the evaluation to a central 24×28 region of the whole 40×40 site system³. As can be seen in the spatially-resolved map in Fig. 6.11a (averaged over 14 adjacent ladder copies), we find strong, chiral currents along the leg bonds, uniformly distributed across the ladder, accompanied by strongly suppressed currents on the rungs, as it is characteristic for the Meissner regime. A detailed evaluation reveals an average current of $0.24(4)J$ on the upper leg, $-0.23(3)J$ on the lower leg, and $0.01(4)K$ on the rungs, respectively (Fig. 6.11b). In addition, we find a homogeneous filling of on average $0.45(2)$ across the ladder without any imbalance between the legs, where the slight deviation from ideal half filling originates mostly from an imperfect initial state (Fig. 6.11c). This combination of strong chiral currents and a homogeneous density is a hallmark feature of the Mott-Meissner phase, arising from the interplay of the artificial magnetic field and the interaction-induced localization.

A second key feature arising from the strongly correlated nature of the state is a characteristic suppression of the chiral current with increasing interaction strength U as well as coupling ratio K/J [165]. This is distinctly different from the non-interacting case, where the chiral current in the Meissner regime is approximately constant. Thanks to the exquisite interaction

³In the months leading to these measurements, the DMD projection system developed aberrations that were not straightforward to remove and decreased the wall sharpness significantly. Prior to that, we were also able to find homogeneous currents in a larger ROI.

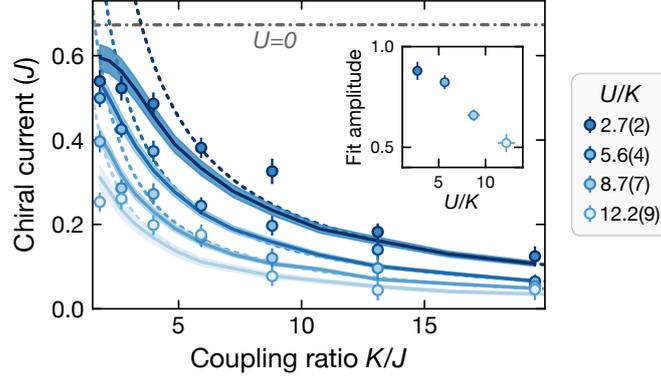


Figure 6.12 | Interaction scaling of the chiral current in the Meissner regime. Suppression of the chiral current with increasing interaction strength U and coupling ratio K/J . The solid lines are fits of a DMRG simulation of the ideal chiral current with the amplitude as a single free parameter, and the shaded areas denote the 1σ -confidence interval of the fits. The dashed lines are the perturbative approximations using the effective spin-1/2 model from Eq. 6.6, scaled to the same fit amplitude. For comparison, the gray dot-dashed trace indicates the non-interacting current from an ED simulation at the same amplitude as the lowest U/K measurement. The inset shows the fit amplitude as a function of U/K . The legend indicates the average U/K for each curve, with the uncertainty denoting the 1σ -variation throughout the K/J range. Each data point is averaged over 60 repetitions and 14 ladders with 48 sites. In all plots, the error bars denote the SEM across the 60 repetitions, and if invisible, are smaller than the marker size.

tunability via our magnetic Feshbach resonance, we can experimentally probe this behavior over a wide range of interaction strengths and compare the observation to the predictions of the effective spin-1/2 model from Eq. 6.6. As before, we prepare interacting Meissner states at half filling with the rung coupling sequence, where we set a fixed rung coupling $K/h = 155(1)$ Hz, and vary the final J to tune K/J for different interaction strengths. The resulting dependence of the chiral current j_c on the coupling ratio K/J for four different interaction strengths is shown in Fig. 6.12. Note that the tuning of K/J via the leg lattice depth causes U to vary slightly throughout each trace, which is labeled by the average U (see Appendix A for the calibration). In detail, we find that the chiral current is suppressed both for increasing interaction strength and for higher coupling ratios. This behavior is in stark contrast to the non-interacting case, where the chiral current remains approximately constant with K/J (cf. gray dot-dashed trace in Fig. 6.12). Additionally, we find good agreement of the scaling with a zero-temperature DMRG simulation (solid lines), as well as the perturbative approximation from the effective spin-1/2 model for $K/J \gtrsim 5$ (Eq. 6.6). As for the plaquette benchmark, the measured current amplitude is likely limited by a residual non-zero U during the basis rotation, which has a stronger influence for a larger initial U (see inset). The highest magnitude is observed for the lowest U/K value with about 88(5)% of the ideally predicted current strength.

This characteristic interaction scaling, together with the observed spatial current and density distributions, provide direct experimental evidence of the Mott-Meissner phase.

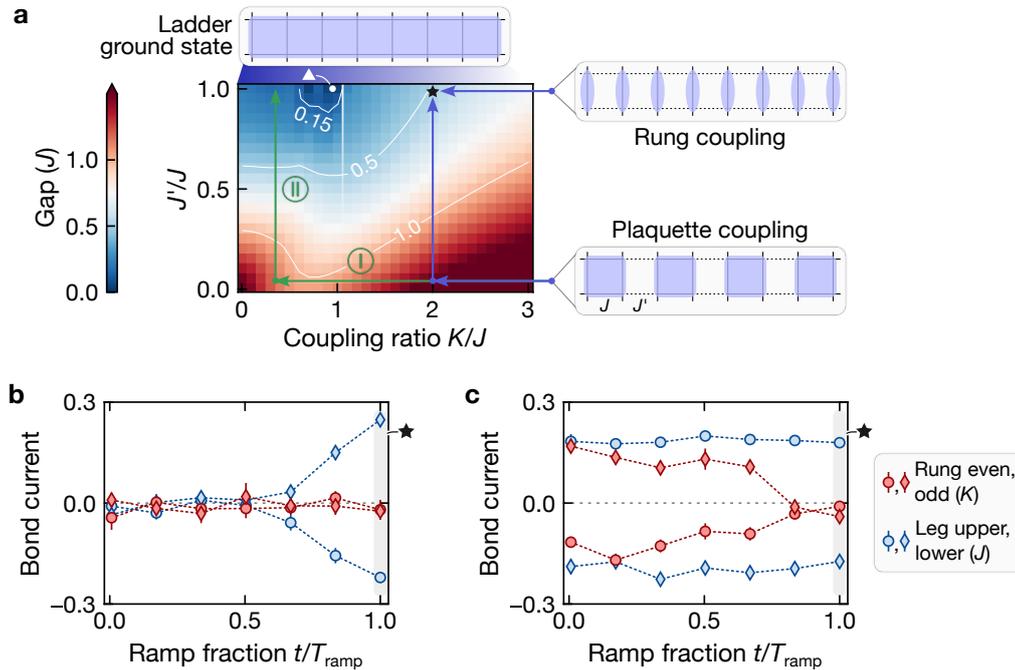


Figure 6.13 | Adiabatic pathways to explore the phase diagram at half filling. **a**, Simulated many-body gap diagram as a function of the coupling ratio K/J and the inter-plaquette coupling J'/J for $U/J = 10$. The solid purple arrows indicate the paths taken by the rung and plaquette coupling sequences to prepare the example Meissner state in (b) and (c), denoted by the black star. The green arrows show the path for the preparation of a vortex state with the plaquette sequence, circumventing the gap closing (white triangle). **b**, **c**, Evolution of the currents during the final ramp with duration T_{ramp} in the rung coupling (b) and in the plaquette coupling (c) sequences for an example Meissner state with $K/J = 1.98(5)$ and $U/J = 11.02(5)$ [marked by the black star in (a)]. The dashed lines are guides to the eye. Each data point is averaged over 30 repetitions and 14 ladders with 48 sites. The error bars denote the SEM across the 30 repetitions, and if not visible, are smaller than the marker size.

6.6 Exploring the phase diagram beyond the Meissner regime

As demonstrated above, the rung coupling sequence provides an ideal pathway for the efficient preparation of strongly-interaction Meissner states. However, it cannot be used to probe the entire phase diagram of the ladder-HBH model at half filling, since the Meissner and vortex regimes are separated by a phase transition and a corresponding closing of the many-body gap at $(K/J)_{\text{cr}} \approx 1$. To explore the entire phase diagram, we develop a second adiabatic preparation sequence that allows us to prepare both Meissner and vortex states at half filling, called the *plaquette coupling sequence*.

6.6.1 Adiabatic preparation by plaquette coupling

A simulated many-body gap diagram at half filling is shown in Fig. 6.13a, indicating the gap-closing point at $(K/J)_{\text{cr}} \approx 1$ (white triangle). Since the rung coupling sequence starts by construction from $K/J \rightarrow \infty$ and monotonically decreases K/J , a direct connection to the vortex regime is not possible. The idea behind the plaquette coupling sequence is to introduce a

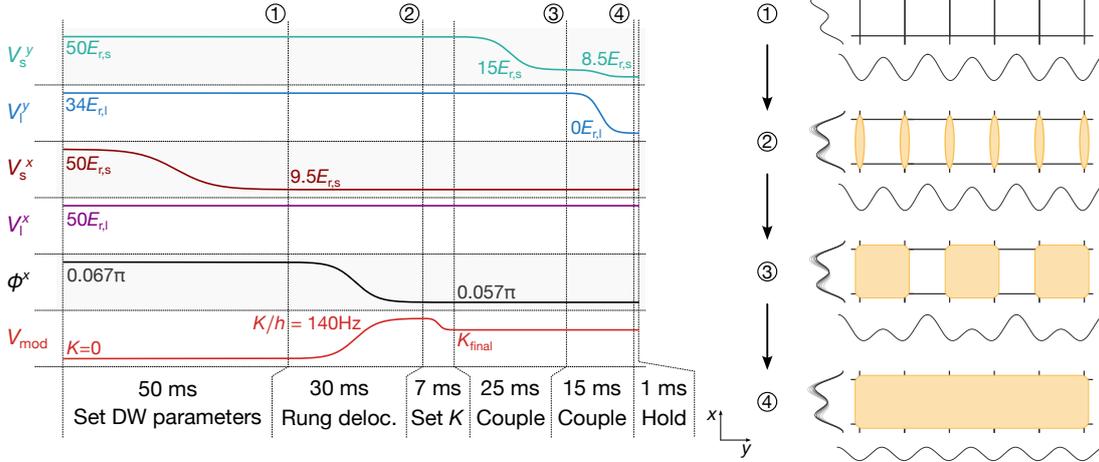


Figure 6.14 | Plaquette coupling sequence. In a first step, the ground states in an array of isolated plaquettes with two particles each at a target K/J is prepared. The plaquettes are then connected to a ladder by ramping down the long lattice along the leg direction, keeping the coupling ratio K/J approximately constant throughout the ramp.

second tuning parameter J' , which opens a finite gap for all K/J and allows us to adiabatically prepare the entire phase diagram. This tuning parameter J' corresponds to a staggered tunnel coupling along the leg direction, which we can introduce using an additional long lattice (vertical axis in Fig. 6.13a). The case of $J' \approx 0$ corresponds to isolated plaquettes, where we can prepare any K/J between 0 and ∞ adiabatically (horizontal path ①). In a second step, the plaquettes are coupled together by increasing $J'/J \rightarrow 1$ at a constant K/J (vertical path ②), which transforms the plaquette array into a fully coupled ladder system. One useful feature of this sequence is that it can be used to prepare both Meissner and vortex states. To compare the two sequences, we prepare an example Meissner state and track the evolution of the bond-resolved currents during the final ramp (coupling the rung bonds or the plaquettes together, respectively). In Fig. 6.13b+c, this is shown for a final state with $K/J = 1.98(5)$ and $U/J = 11.02(5)$. For the rung coupling sequence (Fig. 6.13b), we start without any currents on leg and rung bonds, and only the leg currents gradually build up with opposing sign. In contrast, the plaquette coupling sequence (Fig. 6.13c) starts with strong, opposing currents on both leg and rung bonds (due to the circular plaquette current). Upon connecting the plaquettes to a ladder, the rung currents gradually vanish, while the leg currents remain approximately constant. Due to the longer preparation path, the resulting chiral current is slightly smaller for the plaquette coupling sequence. Hence, for experiments where exclusively the Meissner regime is probed one should use the rung coupling sequence.

A detailed sequence diagram for the plaquette coupling protocol is shown in Fig. 6.14. In the first part, we prepare the ground states in an array of isolated plaquettes with two particles each, following an almost identical protocol as for the plaquette benchmark (Fig. 6.5). The only difference up to this point is a short ramp to a final K after the initial delocalization ramp to a fixed, initial $K/h = 140(1)$ Hz, ensuring that the particle is fully delocalized across the rung bond irrespective of the final rung coupling. By tuning the coupling ratio K/J in the

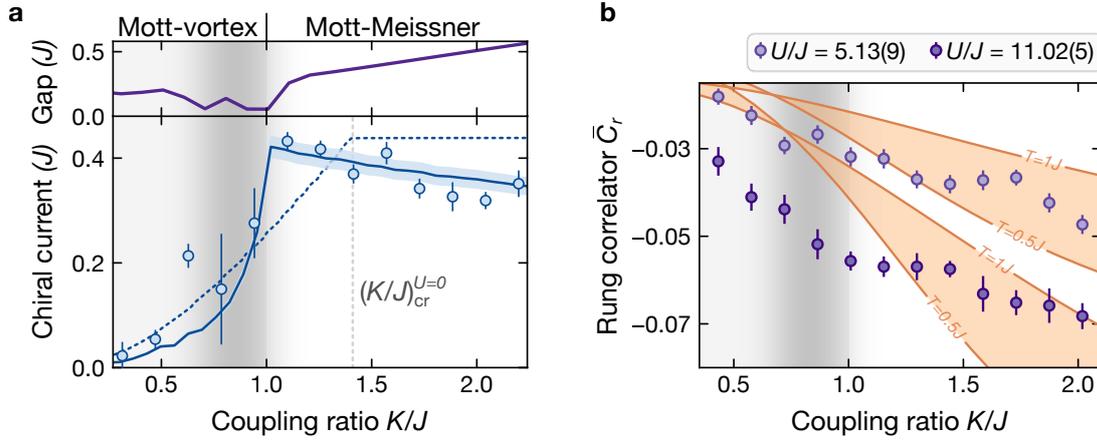


Figure 6.15 | Currents and density correlations between Meissner and vortex regimes. **a**, Average chiral current as a function of K/J for $U/J = 11.02(5)$ and $J/h = 71(1)$ Hz, prepared using the plaquette coupling sequence. The solid line is a fit of the expected chiral current from a DMRG simulation with a scaling factor A as a free parameter, yielding $A = 0.57(3)$. The blue shaded area denotes the 1σ -confidence interval of the fit. The dashed trace shows the non-interacting current, scaled down to the same amplitude, and the vertical line denotes the critical point without interactions. The upper panel indicates the many-body gap across the phase diagram. Each data point is averaged over 80 repetitions and 14 ladders with 48 sites. **b**, Enhancement of the average rung-wise density anti-correlations with increasing interaction energy U and coupling ratio K/J . The orange shaded areas indicate finite-temperature exact diagonalization simulations (2×6 sites) of the density correlations from $k_B T = 0.5 J$ (lower line) to $k_B T = 1 J$ (upper line) for both interaction energies. Each data point was averaged over 30 repetitions and 14 ladders with 48 sites. The error bars denote the SEM, and if not visible, are smaller than the marker size.

plaquette coupling sequence via K , we can make use of the largest possible energy scales in the regime where $K \lesssim J$, as the maximum J is limited by band resonances (see Section 6.4). The leg coupling in the plaquettes is set to roughly the same value as the final leg coupling in the ladders, such that a constant K/J can be maintained throughout the final coupling ramp. In a final step, the plaquettes are connected to a ladder by ramping down the long lattice along the leg direction in 15 ms, together with a slight ramp of the short lattice to maintain a constant leg coupling J . This concludes the plaquette coupling sequence, yielding a fully coupled ladder at half filling and a target coupling ratio K/J .

6.6.2 Currents and density correlations beyond the Meissner regime

The plaquette coupling sequence allows us to probe the phase boundary between the interacting Meissner and vortex regimes. In the presence of strong on-site interactions, the Meissner regime is predicted to extend down to a critical coupling ratio $(K/J)_{\text{cr}} \approx 1$, at which the system undergoes a phase transition to the vortex phase. For the non-interacting case on the other hand, it was previously shown that the transition occurs at a significantly higher value of $(K/J)_{\text{cr}} = \sqrt{2}$ [93, 367]. A striking signature of the phase transition to the vortex regime is a sudden drop of the chiral current upon crossing the critical point. To experimentally probe this behavior, we prepare states with varying K/J and constant, strong interactions of $U/J = 11.02(5)$ using the plaquette coupling sequence, and measure the chiral current

around the phase transition. The results are shown in Fig. 6.15a. Indeed, as we tune K/J below the critical point, we observe a sudden drop of the chiral current, signaling the transition into the vortex regime. The behavior of the current around the phase transition is in good agreement with a zero-temperature DMRG simulation (solid line), with the measured location of the critical point reproducing the predicted value for a strongly interacting system. In contrast to that, a non-interacting system would exhibit a distinctly different behavior (dashed trace), which is a good indication that we probe the ground state properties of the strongly-interacting system. The measured current amplitude is about 57(3)% of the ideal value, where the reduction compared to the rung-coupling sequence is likely due to the longer preparation path, as well as overall slightly smaller energy scales. Additionally, it is worth noting that we observe enhanced fluctuations in the measured current amplitude right below the phase transition, which is compatible with the small many-body gap (see upper panel in Fig. 6.15a) that renders the system highly susceptible to technical heating.

In addition to the chiral currents, we also investigate density-density correlations in the Meissner and vortex regimes. These density correlations can be used to reveal the fractional-Mott-insulating nature of the Meissner phase. Specifically, we expect the ground state for large interactions and in the strong rung coupling limit ($K \gg J$, i.e., deep in the Meissner regime) to be a product state of rung triplets (cf. effective model in Fig. 6.3). The presence of exactly one particle per rung, with the addition of a second particle penalized by an energy on the order of K , results in strong density anti-correlations across the rungs. To experimentally probe this behavior, we measure the average rung-wise density correlations, defined by the correlator $C_r = \langle \hat{n}_{1,r} \hat{n}_{2,r} \rangle - \langle \hat{n}_{1,r} \rangle \langle \hat{n}_{2,r} \rangle$, as a function of the coupling ratio K/J for two different interaction energies. For an ideal fractional Mott-insulator, i.e., exact half filling and infinitely strong interactions, the correlator would assume an extremal value of $C_r = -0.25$. As shown in Fig. 6.15b, we find that the density anti-correlations are indeed enhanced with increasing interaction strength U and coupling ratio K/J , which is an effect of the cross-over to the Mott regime. In comparison to the theoretical expectation, the transition is less sharp than expected, and the correlator strength is weakened. This is due to the slight deviation from half-filling (where one strictly does not expect a perfect Mott state), finite U and, most importantly, a finite temperature as a result of the imperfect state preparation and experimental heating.

As we can easily calibrate the exact filling and the interaction strength, we can make use of the sensitivity of the correlators to a finite temperature to provide a rough estimate of the effective temperature of our many-body state⁴. To this end, we conduct small-scale ED simulations on ladders with 2×6 sites, revealing that the rung-wise density correlations gradually decay as the temperature increases from zero to the order of the leg tunnel coupling. We note that the simulation also incorporates initial state imperfections and parity projection. As illustrated by the orange shaded areas in Fig. 6.15b, comparing the correlator strength with the simulation suggests a temperature on the order of $k_B T \sim J$ in the Meissner phase, which also aligns with the observed chiral current magnitudes (see Appendix A.4 for simulated currents at finite T). A temperature of this scale is also consistent with the predicted elementary

⁴In principle, the currents also show a pronounced dependency on the temperature. However, since the additional basis rotation is involved in the current measurement, it is less straightforward to deduce whether deviations from the ideal value are due to finite temperature or rather due to imperfections in the basis rotation.

excitations of the spin model with energy K , as at this level, the formation of doubly-occupied or empty rungs becomes possible, both of which drive the correlator closer to zero. In the vortex regime ($K < J$), the effective temperature is likely higher due to the smaller gap and lower overall energy scales. Furthermore, the gapless excitations along the leg direction rapidly suppress any current modulation or rung current patterns, making the direct observation of vortices challenging at the currently achievable experimental temperatures (cf. Appendix A.4).

6.7 Future avenues

The results presented in this chapter demonstrate the experimental feasibility of realizing a large-scale interacting many-body phase in the presence of a Floquet-engineered gauge field. A central remaining challenge is a further reduction of the temperature, which we currently estimate to be on the order of $0.5 - 1J$. To allow the exploration of additional phases, in particular gapless phases such as vortex states or precursors of Laughlin states, it is necessary to achieve temperatures on the order of $0.1J$ or below. In our case, the minimal temperature is most likely limited by Floquet heating (see Section 1.3), which fundamentally causes a finite lifetime of the system, and in turn also constrains ramp times for adiabatic state preparation. As motivated in Section 6.4, it is challenging to further reduce this significantly since the modulation frequency is – by construction of the Floquet scheme – always on the same order of magnitude as interband transitions and characteristic energy scales such as the interaction energy U . A study of the more fragile phases would hence benefit from a different scheme to engineer a gauge field. In the following, we briefly discuss strategies that should allow us in the future to reach significantly lower temperatures.

Specifically for the study of ladder systems, it could be worthwhile to explore the feasibility of a synthetic dimension approach with cesium atoms [173–175]. Here, we would use a single one-dimensional lattice along the leg direction, and replace the rung direction with two internal hyperfine states. The two spin states are coupled via two-photon Raman transitions, establishing a ‘tunneling’ term along the synthetic rung direction with a spatially periodic phase factor due to the momentum imparted in the transitions. Suitable spin states would be the states $|F = 3, m_F = 2\rangle =: |\uparrow\rangle$ and $|F = 3, m_F = 3\rangle =: |\downarrow\rangle$ from the ground manifold due to their small magnetic field sensitivity ($\Delta E \approx 350 \text{ kHz/G}$), which is important to maintain spin coherence and spatial homogeneity. Due to the typical optical couplings on the MHz-level, it is straightforward to reach large effective tunnel couplings on the order of $\hbar \times 500 \text{ Hz}$ with an excited state detuning of more than $\sim 500 \text{ GHz}$ [385]. Here, off-resonant scattering is strongly reduced, and the lifetime is expected to be above 1 s. As discussed before, disadvantages of synthetic dimensions are state-dependent interaction strengths (i.e., a different interaction strength on the two legs) and infinite-range interactions along the rung direction. However, the two chosen spin states exhibit several low-field Feshbach resonances that enable to address this challenge. In particular, for offset fields between 56 – 60 G, one can find conditions where the intra-species scattering lengths are approximately equal $a_{\uparrow\uparrow} \approx a_{\downarrow\downarrow} \approx 1000a_0$, and the inter-species scattering length vanishes ($a_{\uparrow\downarrow} \approx 0$) [386]. This would allow to realize the same Hamiltonian as in a real-space implementation, albeit with a reduced interaction tunability,

given that three-body losses at higher offset fields can be managed. Lastly, this implementation is straightforwardly compatible with the locally-resolved detection of densities and currents: When the 1D chains are created using our long-spacing lattice in the perpendicular axis, the two spin states can be separated into two different sites using a horizontal field gradient and the short-spacing lattice. However, this detection technique is limited to two-leg ladders, preventing the scalability of the synthetic dimension approach to larger systems.

In the long term, we envision to use a Raman-assisted tunneling scheme as proposed by Jaksch and Zoller [169], and later refined by Gerbier and Dalibard [292]. It works fully in real space, and is hence straightforwardly compatible with a locally-resolved detection even for large 2D systems. The idea is to use two spin states such as the states $|\uparrow\rangle$ and $|\downarrow\rangle$ from above, which are trapped in a 1D retroreflected antimagic lattice [387]. At the antimagic wavelength (around $\lambda = 871$ nm using σ^+ light), the two states experience an opposite light shift, resulting in an effective lattice with spacing $\lambda/4$, where the two spin states are trapped on adjacent sites. Tunneling between two adjacent antimagic lattice sites, which involves a spin flip, is induced by a two-photon Raman transition similar to the synthetic dimension approach. Due to the short spacing of $\lambda/4$, large effective tunnel couplings of several 100 Hz can be reached, which provides an intrinsic robustness to disorder and heating. By appropriate choice of the Raman wave-vectors in relation to the lattice axis, one can engineer a spatially-dependent complex tunnel coupling. In combination with a normal monochromatic lattice in the perpendicular direction, this allows realizing both staggered and homogeneous flux distributions.

Cesium is particularly well suited for this scheme due to the large fine-structure splitting, resulting in low off-resonant scattering rates for the antimagic lattice. Additionally, since the energy scale for the induced tunnel coupling is now given by a hyperfine transition (\sim MHz), substantially smaller heating rates than in the Floquet approach are expected. Nevertheless, in terms of the interparticle interactions, similar considerations as for the synthetic dimension approach apply, which requires further tests to ensure the experimental feasibility. Also note that the $\lambda/4$ separation between atoms complicates the imaging, which could – at least for the study of two-leg ladders – be relaxed by keeping columns between adjacent ladders empty.

6.8 Conclusions

In this chapter, we presented the preparation and study of large-scale interacting many-body phases on bosonic flux ladders. By leveraging the excellent interaction tunability of our platform, we devised adiabatic sequences to prepare the strongly interacting Mott-Meissner phase with up to 48 sites at half filling. In the context of many-body systems with a Floquet-engineered artificial gauge field, this constitutes a substantial advance beyond the state of the art of just two particles [111, 380], reaching a genuine many-body regime. Employing our readout technique for local currents, we revealed the key microscopic features of the Mott-Meissner phase: persistent chiral currents along the ladder legs, strongly suppressed rung currents, and an interaction-induced suppression of the chiral current. Additionally, we observed density anti-correlations across the rung bonds, which are a direct signature of fractional-Mott-insulating behavior at half filling. Our experiments also constitute the first measurement of local particle currents with microscopic resolution in a many-body system of neutral atoms. Lastly, we compared our measurements with numerical simulations of small systems, allowing us to estimate the effective temperature of the many-body state. With this, we set a new benchmark for interacting, periodically driven quantum systems and provide an important reference point for future theoretical and experimental efforts.

These experiments establish an entirely new platform for many-body interacting matter with artificial gauge fields, featuring high tunability and microscopic resolution. As a result, a wide range of exciting research directions becomes accessible:

Dynamics in time-reversal-symmetry-broken many-body phases

An immediate direction is to use this platform to study dynamics and transport in time-reversal-symmetry-broken many-body phases. To this end, we can employ the current measurement scheme as a novel probe of non-equilibrium dynamics, as it is directly applicable to time-resolved experiments. Potential applications include measurements of the Hall response [364–366], chiral Bloch oscillations [388, 389], optical conductivities, and quench dynamics [390–394]. Studying quenches across phase transitions may also provide insights into different phases when adiabatic ground-state preparation remains out of reach [363]. Furthermore, it would be interesting to realize spatially inhomogeneous flux distributions, for example by partitioning the system into two regions with different flux values. This can be implemented by projecting the running-wave modulation with a DMD, enabling the realization of magnetic barriers and the study of the resulting equilibrium and transport properties [395].

Exploring the phase diagram of interacting flux ladders

With the Mott-Meissner phase, we have realized a first phase from the rich phase diagram of interacting flux ladders. As introduced in Section 6.1, other theoretically predicted phases include vortex states, symmetry-broken states such as biased-leg and CDW phases, chiral Mott/superfluid phases, and Laughlin-like states, all of which remain unexplored in experiments [161, 165, 358, 359]. In the vortex regime, it would be particularly relevant to investigate vortex liquid and vortex lattice states and identify their distinct experimental signatures. A

vortex liquid is expected to exhibit current vortices at varying locations across experimental realizations, detectable via current-current correlation functions. In contrast, a vortex lattice state should exhibit pinned vortex patterns, akin to a Mott insulator of current vortices [358, 359].

Since many of these phases exhibit smaller many-body gaps than the Meissner phase, further efforts are required to reduce the effective temperature of the system. This could be achieved by employing techniques to suppress Floquet heating [396–398], and by exploring further preparation strategies such as shortcuts to adiabaticity and counterdiabatic driving protocols [399, 400]. For instance, vortex states might be prepared by initializing isolated 1D tubes that are later coupled, thereby avoiding gap closings during the preparation (similar to [160]). Alternatively, the low-frequency Floquet-engineering scheme could be replaced by a Raman-assisted tunneling protocol (see Section 6.7), which is expected to induce less heating. In the context of ladder geometries, a synthetic-dimension approach [174] may also be advantageous, with the two legs encoded in distinct spin states, such as $|F = 3, m_F = 2\rangle$ and $|F = 3, m_F = 3\rangle$. Single-site resolution of occupations and currents could still be maintained by splitting the spin states using a double well in the rung direction. Lastly, exploring the full phase diagram requires tunable control over the flux. Using the Floquet scheme, this can be achieved either by adjusting the angle of incidence of the running-wave beams for fixed flux values or by projecting the running-wave through one of the objectives, which would also allow dynamically varying the flux [111, 380].

Interacting topological matter in 2D

The ingredients demonstrated here – the combination of interactions and artificial gauge fields – form the foundation for realizing topologically ordered matter, such as fractional Chern insulators (FCIs) [150]. Beyond verifying Laughlin-like states in two-leg ladders [374], flux ladders serve as a promising starting point for constructing larger 2D systems. For example, one could merge the ground states of several adjacent ladders into a full 2D system [160, 401, 402]. Here, a key experimental challenge includes the engineering of preparation sequences: As system sizes increase, finite-size gaps that aid adiabatic preparation vanish. Preparation timescales also grow with the particle number, imposing stringent requirements on coherence times and heating rates. Moreover, since a transition from a topologically trivial to a non-trivial state necessarily involves a gap closing, adiabatic preparation may not be viable. As an alternative, strategies such as open-system [403] or dissipative [404–406] state engineering may therefore become increasingly relevant. Upon overcoming these challenges, theoretical studies suggest that most features of Laughlin physics already emerge in three-leg ladders [407], and adding more legs may quickly stabilize non-Abelian phases [408]. Given that we have already demonstrated control over many interacting particles, this could provide a realistic path toward large-scale quantum simulations of FCIs with hundreds of atoms.

Conclusions and outlook

In this thesis, we presented several experimental developments that culminated in a new analog quantum simulation platform based on ultracold cesium atoms, as well as studies of many-body phases in optical flux ladders. To conclude, we summarize the key results and provide a high-level outlook on the future of our experiment. For a detailed discussion of the individual projects, the reader is referred to the conclusions of the respective chapters.

The first two parts focused on the development of new detection capabilities for neutral-atom quantum simulators. First, we introduced a method to reconstruct site-resolved occupations in optical lattices using unsupervised neural networks. This method enables a fast, high-fidelity reconstruction in regimes where atomic separations are significantly smaller than the optical resolution limit. To date, our approach remains unsurpassed, achieving reconstruction fidelities above 95% despite an extreme resolution-to-spacing ratio exceeding two. For us, it was the key to obtaining single-site resolution and has proven robust across multiple quantum simulation experiments, including the detection of subtle higher-order correlation functions [124, 226, 290, 291]. The method is applicable to any platform with single-particle detection, with its high-speed operation being particularly relevant for neutral atom quantum computing platforms [297, 298].

Second, we implemented a robust and scalable detection and control technique for off-diagonal observables such as local currents and kinetic energies using optical superlattices. This approach offers both global and local control, bridging a long-standing gap in the capabilities of quantum gas microscopes. Promising applications include the study of topological matter [90, 91, 165], non-equilibrium dynamics [310], and superconductivity [94]. Additionally, the available local manipulations can be used for coherent state engineering and as building blocks for (hybrid) quantum computing approaches [95, 96].

In the third part, we experimentally realized large-scale bosonic flux ladders with up to 48 sites. By combining local readout of currents with flexible interaction control, we found experimental evidence for the strongly correlated Mott-Meissner phase, characterized by chiral particle currents atop a Mott-localized density plateau. A key component enabling this study was the implementation of a new short-spacing vertical lattice with a stronger confinement, which, in combination with the broad Feshbach resonance of cesium, allowed us to tune the interaction strength over a wide range. With this, we established a new benchmark for the size of many-body phases with artificial gauge fields in neutral atom quantum simulators, opening new directions for large-scale studies of topologically ordered matter.

All these projects directly address the strengths of analog quantum simulators based on

neutral atoms, either by enhancing or harnessing them. Featuring large system sizes, high tunability, long coherence times, and access to novel observables, these platforms hold promise for achieving practical quantum advantage in the study of exotic quantum matter [65, 69]. A central challenge in the field, and one we also aim to address in our platform, is the reduction of entropy and temperature to access low-energy phases. To this end, we are improving laser setups and superlattice frequency locks to reduce intensity and frequency noise, and we are implementing entropy engineering schemes such as the one presented in Ref. [322], based on a spin-dependent superlattice. This approach aims to realize low-entropy 1D systems and, if successful, could be extended by incorporating a second atomic plane acting as a reservoir, increasing the thermal contact for full 2D systems [409, 410].

We envision several further research directions that exploit the unique capabilities of our experimental platform. Specifically, our system is well-suited for quantum simulation of the sine-Gordon model, a relativistic field theory emerging in the relative degree of freedom in two-leg ladders [411]. Utilizing our current detection capabilities, we could directly probe its complex quasiparticle excitation spectrum, which features topological excitations such as solitons [412]. Furthermore, the broad interaction tunability of cesium across multiple spin states enables the realization of spin models with widely tunable anisotropy [42]. Lastly, our 2D superlattices enable the engineering of models featuring constrained dynamics, resulting in Hilbert space fragmentation [345, 413, 414], and allow the realization of lattice gauge theories [351, 415, 416]. These directions offer a wide range of opportunities for experiments at the frontier of quantum many-body physics, all within reach of our experimental platform.

Appendices

Appendix A Supporting material for the study of interacting flux ladders

In this appendix, we discuss additional calibration measurements and simulations that were carried out for the interacting flux ladders presented in Chapter 6.

A.1 Variation of the on-site interaction energy with the leg lattice depth

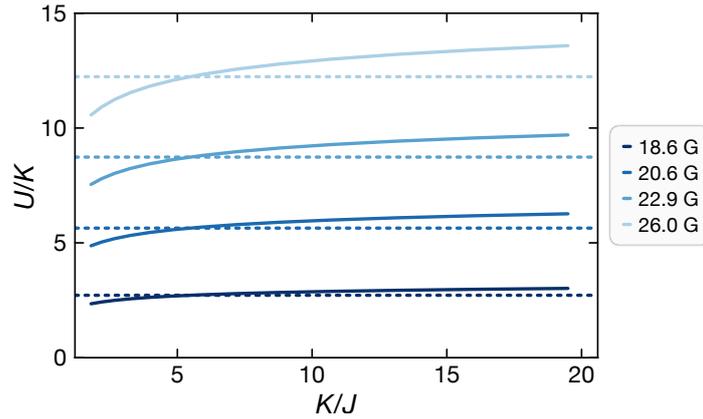


Figure A.1 | Variation of the on-site interaction energy with the leg lattice depth. The variation of the calibrated interaction energy as a function of the leg lattice depth is extrapolated using a band structure calculation (solid lines) for different offset fields (i.e., scattering lengths). The dashed horizontal lines indicate the average U for each offset field over the whole range of K/J .

We perform the spectroscopic calibration of the interaction energy U as described in Section 3.5 primarily for one specific lattice depth. As the interaction energy varies with the horizontal lattice parameters via the size of the Wannier functions, we need to extrapolate the interaction energy for different lattice depths. To do so, we make use of the fully calibrated band structure, which in particular requires knowledge of the z confinement, horizontal lattice depths and the superlattice phase. In detail, the spectroscopic calibration at a given lattice depth calibrates the scattering length a_s for the corresponding offset field. We then fix the scattering length and extrapolate the interaction energy U for different lattice depths using the band structure calculation. This is shown in Fig. A.1 as an example for the interaction scaling experiment in Fig. 6.12, where we tuned the coupling ratio K/J via the leg lattice depth (here $K/h = 155(1)$ Hz). To cross-check the extrapolation, we perform an independent spectroscopic U

calibration for different lattice parameters, which agrees with the extrapolated values within the measurement uncertainty.

A.2 Current lifetime and reference in real ladders

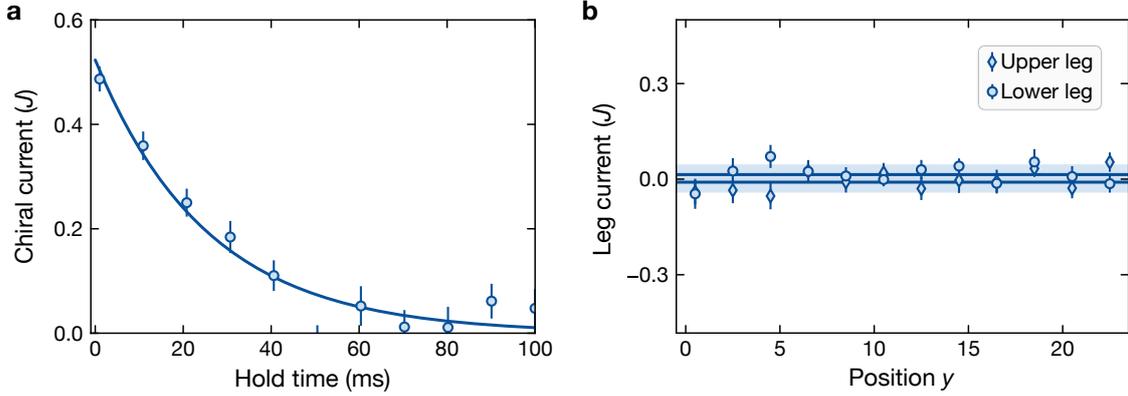


Figure A.2 | Current lifetime and reference measurement in real ladders. **a**, Lifetime of the chiral current in the Meissner phase for $K/J = 2.19(3)$ and $U/J = 11.02(5)$. The solid line is an exponential fit with a $1/e$ -decay time of $26(3)$ ms. The error bars denote the standard error of the mean over 30 repetitions (averaged over 14 ladders with 48 sites), and if invisible, are smaller than the marker size. **b**, Spatially resolved leg current measurement in a ladder without flux. The horizontal lines and the surrounding shading illustrate the spatial average current and its 1σ -uncertainty across the ladder. The average current is $-0.01(3)J$ for the upper leg and $0.01(3)J$ for the lower leg, both consistent with zero. The error bars denote the standard error of the mean over 80 repetitions (evaluation region 28×28 sites), and if invisible, are smaller than the marker size.

Due to Floquet and technical heating, the ground state phases exhibit a finite lifetime, which can be seen in a decay of the equilibrium currents over time. To quantify the lifetime, we prepare a Meissner state using the rung coupling sequence and vary the hold time between state preparation and current measurement. Fig. A.2a shows the measured chiral current as a function of the hold time for a Meissner state with $K/J = 2.19(3)$ and $U/J = 11.02(5)$. The decay of the current is well-described by a single exponential with a $1/e$ -decay time of $26(3)$ ms. As the lifetime is on a similar order with typical ramp durations, this highlights the importance of choosing suitable parameter regimes where heating is suppressed.

Additionally, we perform a reference current measurement in a ladder with real couplings along the leg and rungs. The corresponding ground state does not exhibit equilibrium currents, which should be reflected in a vanishing measured current. To this end, we prepare a ladder with only real-valued tunnel couplings in its ground state at half-filling. We start with the same initial state – one particle per rung bond, which we then delocalize in a symmetric, bare double well at a rung tunnel coupling of around $J_{\perp}/h = 165$ Hz. In the final step, we couple the rungs by increasing the leg tunnel coupling to $J/h = 71(1)$ Hz, realizing parameters similar to those used in the flux ladders, with an interaction energy of approximately $U/h = 890$ Hz. We then measure the leg currents with spatial resolution using the standard leg current measurement sequence. As shown in Fig. A.2b, we observe vanishing currents across all bonds of the ladder

with a spatial average of $-0.01(3)J$ for the upper leg, and $0.01(3)J$ for the lower leg, respectively. Both the individual bond currents and the spatial averages are consistent with zero, as expected.

A.3 Details on the numerical simulations

To benchmark our measurements and compare with theoretical predictions, we perform several different numerical simulations. In detail, small system numerics such as plaquettes, non-interacting ladders, as well as the finite temperature simulations are performed using exact diagonalization (ED). For this, we use the Python library QuSpin [417]. We account for parity projection in all simulations by evaluating observables in a parity-projected Fock basis with all on-site occupations taken mod 2.

The numerics for larger interacting systems are performed using tensor network algorithms implemented in the Python library TeNPy [418]. All ground-state wave functions are computed using the density matrix renormalization group (DMRG) algorithm [419]. We typically impose a local cutoff of $n_{\max} = 3$ bosons per site to reduce computational cost in the interacting regime and truncate at a bond dimension between $\chi = 600 \dots 1024$ depending on the specific parameters. For each case, we verify convergence by comparing results at different values of χ . To account for parity projection, we simulate the rotation to the current basis by evolving the ground-state wave function in isolated double wells using the time-dependent variational principle (TDVP) algorithm. Afterward, we sample typically around 10^4 density snapshots using the built-in sampling algorithm. The occupation samples are then taken mod 2, and via the left-right density difference, give access to the parity-projected current.

A.4 Finite-temperature simulations

To estimate the effective temperature of the system as described by the effective Hamiltonian – in our case, the ladder-HBH-model – we perform ED numerics on small systems and evaluate the finite temperature expectation value of the rung-wise density-density correlator [35, 420]. Assuming that the many-body statistics are governed by the canonical ensemble, the expectation value of an observable A at temperature T is calculated as

$$\langle A \rangle = \sum_{\Omega} e^{-E_{\Omega}/k_B T} \langle \psi_{\Omega} | A | \psi_{\Omega} \rangle \quad (\text{A.1})$$

where ψ_{Ω} denotes the eigenstate, which is labelled by Ω , E_{Ω} is the corresponding eigenenergy, and k_B denotes the Boltzmann constant. For each K/J and U/J , we obtain a full eigenspectrum in the Fock space basis considering up to three bosons per site and six particles in 2×6 ladder system. Then, the observable A is given by

$$\langle A \rangle = \sum_{\Omega} e^{-E_{\Omega}/k_B T} \sum_j |c_{\Omega,j}|^2 \langle \dots, n_k, \dots | A | \dots, n_k, \dots \rangle, \quad (\text{A.2})$$

with each eigenstate $|\psi_{\Omega}\rangle = \sum_j c_{\Omega,j} |\dots, n_k, \dots\rangle$ constrained to a fixed particle number $\sum_k n_k = 6$. We evaluate the connected rung density correlator including parity projection, $\langle \mathcal{P}_{1,r} \mathcal{P}_{2,r} \rangle -$

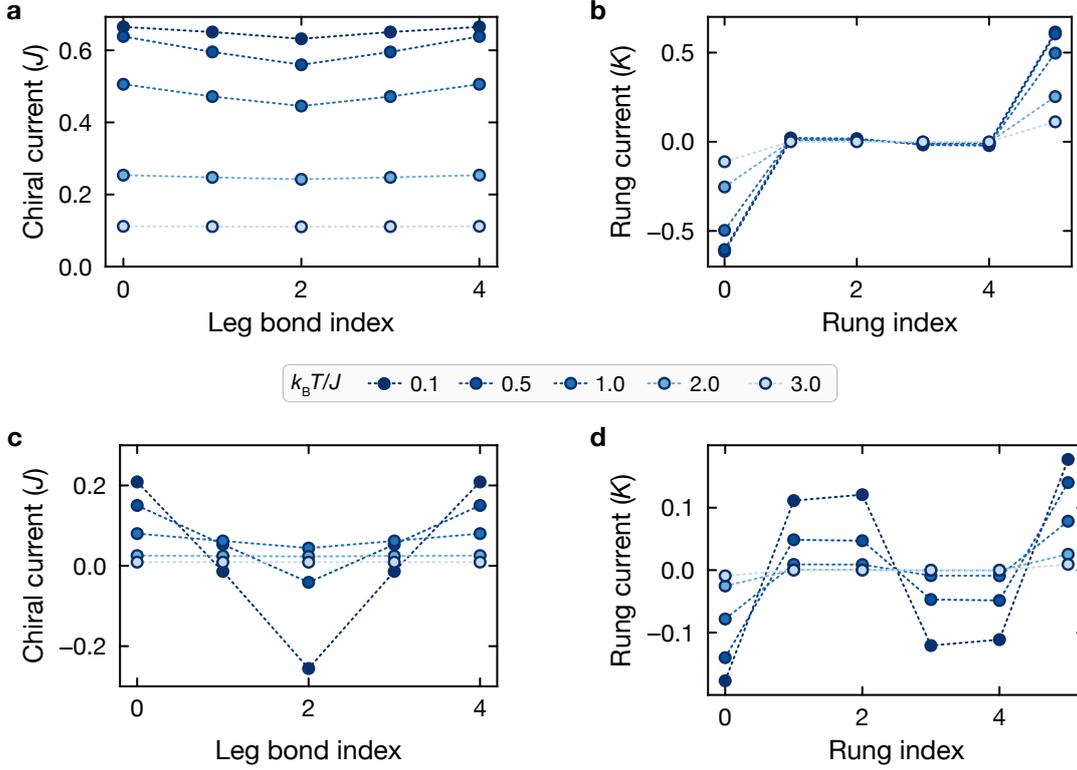


Figure A.3 | Simulated currents for finite temperature. Small-scale ED simulation on 2×6 sites at half filling, showing the chiral leg current and rung current distributions for $K/J = 2$ (Meissner phase, **a** and **b**) and $K/J = 0.5$ (vortex phase, **c** and **d**) for different temperatures.

$\langle \mathcal{P}_{1,r} \rangle \langle \mathcal{P}_{2,r} \rangle$, where $\mathcal{P}_{l,r}$ is the parity-projected particle number operator for the site $l = 1, 2$ of the r -th rung.

To compare with the experimentally measured density correlations, we simulate the system using the calibrated Hubbard parameters for five and six particles, including parity projection as described above, at several temperatures. We then average the results for five and six particles to obtain predictions for our experimental filling of approximately 0.45. As shown in Fig. 6.15, the measured correlation magnitude in the range $K/J = 1 \dots 2$ is consistent with a temperature on the order of $k_B T \sim J$. In the vortex regime, the temperature is likely higher due to the smaller many-body gap.

Additionally, we investigate the temperature dependence of the current distribution using the same finite-temperature numerical simulation described above. We set $U/J = 11$ and consider two coupling ratios in the Meissner ($K/J = 2$) and vortex ($K/J = 0.5$) regimes, respectively. In the Meissner regime, at $k_B T = J$, the current distribution retains its characteristic homogeneous profile, but the current is reduced to approximately 70%, consistent with the experimentally measured strength. In contrast, in the vortex regime, the modulation amplitude of the current distribution nearly vanishes at $k_B T = J$, making experimental detection significantly more difficult. This highlights that a central objective is to further reduce the effective temperature of the system.

References

- [1] T. H. Maiman, *Stimulated Optical Radiation in Ruby*, [Nature](#) **187**, 493 (1960).
- [2] J. E. Lilienfeld, *Device for Controlling Electric Current*, US1900018A (Mar. 1933).
- [3] I. I. Rabi, J. R. Zacharias, S. Millman, and P. Kusch, *A New Method of Measuring Nuclear Magnetic Moment*, [Phys. Rev.](#) **53**, 318 (1938).
- [4] E. M. Purcell, H. C. Torrey, and R. V. Pound, *Resonance Absorption by Nuclear Magnetic Moments in a Solid*, [Phys. Rev.](#) **69**, 37 (1946).
- [5] F. Bloch, *Nuclear Induction*, [Phys. Rev.](#) **70**, 460 (1946).
- [6] P. W. Anderson, *More Is Different*, [Science](#) **177**, 393 (1972).
- [7] B. Keimer, S. A. Kivelson, M. R. Norman, S. Uchida, and J. Zaanen, *From Quantum Matter to High-Temperature Superconductivity in Copper Oxides*, [Nat. Phys.](#) **518**, 179 (2015).
- [8] X. Zhou, W.-S. Lee, M. Imada, N. Trivedi, P. Phillips, H.-Y. Kee, P. Törmä, and M. Eremets, *High-Temperature Superconductivity*, [Nat. Rev. Phys.](#) **3**, 462 (2021).
- [9] M. Z. Hasan and C. L. Kane, *Colloquium: Topological Insulators*, [Rev. Mod. Phys.](#) **82**, 3045 (2010).
- [10] F. D. M. Haldane, *Nobel Lecture: Topological Quantum Matter*, [Rev. Mod. Phys.](#) **89**, 040502 (2017).
- [11] X.-G. Wen, *Colloquium: Zoo of Quantum-Topological Phases of Matter*, [Rev. Mod. Phys.](#) **89**, 041004 (2017).
- [12] L. Balents, *Spin Liquids in Frustrated Magnets*, [Nature](#) **464**, 199 (2010).
- [13] B. Keimer and J. E. Moore, *The Physics of Quantum Materials*, [Nat. Phys.](#) **13**, 1045 (2017).
- [14] I. Buluta and F. Nori, *Quantum Simulators*, [Science](#) **326**, 108 (2009).
- [15] I. Bloch, J. Dalibard, and S. Nascimbène, *Quantum Simulations with Ultracold Quantum Gases*, [Nat. Phys.](#) **8**, 267 (2012).
- [16] J. I. Cirac and P. Zoller, *Goals and Opportunities in Quantum Simulation*, [Nat. Phys.](#) **8**, 264 (2012).
- [17] I. M. Georgescu, S. Ashhab, and F. Nori, *Quantum Simulation*, [Rev. Mod. Phys.](#) **86**, 153 (2014).
- [18] C. Gross and I. Bloch, *Quantum Simulations with Ultracold Atoms in Optical Lattices*, [Science](#) **357**, 995 (2017).
- [19] S. Chu, L. Hollberg, J. E. Bjorkholm, A. Cable, and A. Ashkin, *Three-Dimensional Viscous Confinement and Cooling of Atoms by Resonance Radiation Pressure*, [Phys. Rev. Lett.](#) **55**, 48 (1985).
- [20] W. D. Phillips, *Nobel Lecture: Laser Cooling and Trapping of Neutral Atoms*, [Rev. Mod. Phys.](#) **70**, 721 (1998).

- [21] M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman, and E. A. Cornell, *Observation of Bose-Einstein Condensation in a Dilute Atomic Vapor*, *Science* **269**, 198 (1995).
- [22] C. C. Bradley, C. A. Sackett, J. J. Tollett, and R. G. Hulet, *Evidence of Bose-Einstein Condensation in an Atomic Gas with Attractive Interactions*, *Phys. Rev. Lett.* **75**, 1687 (1995).
- [23] K. B. Davis, M.-O. Mewes, M. R. Andrews, N. J. van Druten, D. S. Durfee, D. M. Kurn, and W. Ketterle, *Bose-Einstein Condensation in a Gas of Sodium Atoms*, *Phys. Rev. Lett.* **75**, 3969 (1995).
- [24] D. Jaksch, C. Bruder, J. I. Cirac, C. W. Gardiner, and P. Zoller, *Cold Bosonic Atoms in Optical Lattices*, *Phys. Rev. Lett.* **81**, 3108 (1998).
- [25] M. Greiner, O. Mandel, T. Esslinger, T. W. Hänsch, and I. Bloch, *Quantum Phase Transition from a Superfluid to a Mott Insulator in a Gas of Ultracold Atoms*, *Nature* **415**, 39 (2002).
- [26] I. Bloch, *Quantum Simulations Come of Age*, *Nat. Phys.* **14**, 1159 (2018).
- [27] J. Hubbard, *Electron Correlations in Narrow Energy Bands*, *Proc. R. Soc. Lond. Ser. A Math. Phys. Sci.* **276**, 238 (1963), JSTOR: 2414761.
- [28] D. Jaksch and P. Zoller, *The Cold Atom Hubbard Toolbox*, *Ann. Phys.* **315**, 52 (2005).
- [29] F. Schäfer, T. Fukuhara, S. Sugawa, Y. Takasu, and Y. Takahashi, *Tools for Quantum Simulation with Ultracold Atoms in Optical Lattices*, *Nat. Rev. Phys.* **2**, 411 (2020).
- [30] W. S. Bakr, J. I. Gillen, A. Peng, S. Fölling, and M. Greiner, *A Quantum Gas Microscope for Detecting Single Atoms in a Hubbard-regime Optical Lattice*, *Nature* **462**, 74 (2009).
- [31] J. F. Sherson, C. Weitenberg, M. Endres, M. Cheneau, I. Bloch, and S. Kuhr, *Single-Atom-Resolved Fluorescence Imaging of an Atomic Mott Insulator*, *Nature* **467**, 68 (2010).
- [32] M. Endres, M. Cheneau, T. Fukuhara, C. Weitenberg, P. Schauß, C. Gross, L. Mazza, M. C. Bañuls, L. Pollet, I. Bloch, and S. Kuhr, *Observation of Correlated Particle-Hole Pairs and String Order in Low-Dimensional Mott Insulators*, *Science* **334**, 200 (2011).
- [33] M. Endres, M. Cheneau, T. Fukuhara, C. Weitenberg, P. Schauß, C. Gross, L. Mazza, M. C. Bañuls, L. Pollet, I. Bloch, and S. Kuhr, *Single-Site- and Single-Atom-Resolved Measurement of Correlation Functions*, *Appl. Phys. B* **113**, 27 (2013).
- [34] R. Islam, R. Ma, P. M. Preiss, M. Eric Tai, A. Lukin, M. Rispoli, and M. Greiner, *Measuring Entanglement Entropy in a Quantum Many-Body System*, *Nature* **528**, 77 (2015).
- [35] A. M. Kaufman, M. E. Tai, A. Lukin, M. Rispoli, R. Schittko, P. M. Preiss, and M. Greiner, *Quantum Thermalization through Entanglement in an Isolated Many-Body System*, *Science* **353**, 794 (2016).
- [36] A. Lukin, M. Rispoli, R. Schittko, M. E. Tai, A. M. Kaufman, S. Choi, V. Khemani, J. Léonard, and M. Greiner, *Probing Entanglement in a Many-Body-Localized System*, *Science* **364**, 256 (2019).
- [37] L. Asteria, H. P. Zahn, M. N. Kosch, K. Sengstock, and C. Weitenberg, *Quantum Gas Magnifier for Sub-Lattice-Resolved Imaging of 3D Quantum Systems*, *Nature* **599**, 571 (2021).
- [38] C. Gross and W. S. Bakr, *Quantum Gas Microscopy for Single Atom and Spin Detection*, *Nat. Phys.* **17**, 1316 (2021).
- [39] C. Weitenberg, M. Endres, J. F. Sherson, M. Cheneau, P. Schauß, T. Fukuhara, I. Bloch, and S. Kuhr, *Single-Spin Addressing in an Atomic Mott Insulator*, *Nature* **471**, 319 (2011).
- [40] W. S. Bakr, A. Peng, M. E. Tai, R. Ma, J. Simon, J. I. Gillen, S. Fölling, L. Pollet, and M. Greiner, *Probing the Superfluid-to-Mott Insulator Transition at the Single-Atom Level*, *Science* **329**, 547 (2010).

-
- [41] A. Sommer, M. Ku, G. Roati, and M. W. Zwierlein, *Universal Spin Transport in a Strongly Interacting Fermi Gas*, *Nature* **472**, 201 (2011).
- [42] P. N. Jepsen, J. Amato-Grill, I. Dimitrova, W. W. Ho, E. Demler, and W. Ketterle, *Spin Transport in a Tunable Heisenberg Model Realized with Ultracold Atoms*, *Nature* **588**, 403 (2020).
- [43] D. Wei, A. Rubio-Abadal, B. Ye, F. Machado, J. Kemp, K. Srakaew, S. Hollerith, J. Rui, S. Gopalakrishnan, N. Y. Yao, I. Bloch, and J. Zeiher, *Quantum Gas Microscopy of Kardar-Parisi-Zhang Superdiffusion*, *Science* **376**, 716 (2022).
- [44] S. Trotzky, P. Cheinet, S. Fölling, M. Feld, U. Schnorrberger, A. M. Rey, A. Polkovnikov, E. A. Demler, M. D. Lukin, and I. Bloch, *Time-Resolved Observation and Control of Superexchange Interactions with Ultracold Atoms in Optical Lattices*, *Science* **319**, 295 (2008).
- [45] P. Scholl, M. Schuler, H. J. Williams, A. A. Eberharter, D. Barredo, K.-N. Schymik, V. Lienhard, L.-P. Henry, T. C. Lang, T. Lahaye, A. M. Läuchli, and A. Browaeys, *Quantum Simulation of 2D Antiferromagnets with Hundreds of Rydberg Atoms*, *Nature* **595**, 233 (2021).
- [46] M. Xu, L. H. Kendrick, A. Kale, Y. Gang, G. Ji, R. T. Scalettar, M. Lebrat, and M. Greiner, *Frustration- and Doping-Induced Magnetism in a Fermi-Hubbard Simulator*, *Nature* **620**, 971 (2023).
- [47] A. Mazurenko, C. S. Chiu, G. Ji, M. F. Parsons, M. Kanász-Nagy, R. Schmidt, F. Grusdt, E. Demler, D. Greif, and M. Greiner, *A Cold-Atom Fermi-Hubbard Antiferromagnet*, *Nature* **545**, 462 (2017).
- [48] X. Li, S. Wang, X. Luo, Y.-Y. Zhou, K. Xie, H.-C. Shen, Y.-Z. Nie, Q. Chen, H. Hu, Y.-A. Chen, X.-C. Yao, and J.-W. Pan, *Observation and Quantification of the Pseudogap in Unitary Fermi Gases*, *Nature* **626**, 288 (2024).
- [49] H.-J. Shao, Y.-X. Wang, D.-Z. Zhu, Y.-S. Zhu, H.-N. Sun, S.-Y. Chen, C. Zhang, Z.-J. Fan, Y. Deng, X.-C. Yao, Y.-A. Chen, and J.-W. Pan, *Antiferromagnetic Phase Transition in a 3D Fermionic Hubbard Model*, *Nature* **632**, 267 (2024).
- [50] T. Chalopin, P. Bojović, S. Wang, T. Franz, A. Sinha, Z. Wang, D. Bourgund, J. Obermeyer, F. Grusdt, A. Bohrdt, L. Pollet, A. Wietek, A. Georges, T. Hilker, and I. Bloch, *Probing the Magnetic Origin of the Pseudogap Using a Fermi-Hubbard Quantum Simulator*, [arXiv:2412.17801](https://arxiv.org/abs/2412.17801).
- [51] M. Aidelsburger, M. Atala, M. Lohse, J. T. Barreiro, B. Paredes, and I. Bloch, *Realization of the Hofstadter Hamiltonian with Ultracold Atoms in Optical Lattices*, *Phys. Rev. Lett.* **111**, 185301 (2013).
- [52] H. Miyake, G. A. Siviloglou, C. J. Kennedy, W. C. Burton, and W. Ketterle, *Realizing the Harper Hamiltonian with Laser-Assisted Tunneling in Optical Lattices*, *Phys. Rev. Lett.* **111**, 185302 (2013).
- [53] G. Jotzu, M. Messer, R. Desbuquois, M. Lebrat, T. Uehlinger, D. Greif, and T. Esslinger, *Experimental Realization of the Topological Haldane Model with Ultracold Fermions*, *Nature* **515**, 237 (2014).
- [54] G. Semeghini et al., *Probing Topological Spin Liquids on a Programmable Quantum Simulator*, *Science* **374**, 1242 (2021).
- [55] C. Braun, R. Saint-Jalm, A. Hesse, J. Arceri, I. Bloch, and M. Aidelsburger, *Real-Space Detection and Manipulation of Topological Edge Modes with Ultracold Atoms*, *Nat. Phys.* **20**, 1306 (2024).
- [56] B. Fauseweh, *Quantum Many-Body Simulations on Digital Quantum Computers: State-of-the-art and Future Challenges*, *Nat. Commun.* **15**, 2123 (2024).

-
- [57] E. Altman et al., *Quantum Simulators: Architectures and Opportunities*, *PRX Quantum* **2**, 017003 (2021).
- [58] H. Bernien, S. Schwartz, A. Keesling, H. Levine, A. Omran, H. Pichler, S. Choi, A. S. Zibrov, M. Endres, M. Greiner, V. Vuletić, and M. D. Lukin, *Probing Many-Body Dynamics on a 51-Atom Quantum Simulator*, *Nature* **551**, 579 (2017).
- [59] A. M. Kaufman and K.-K. Ni, *Quantum Science with Optical Tweezer Arrays of Ultracold Atoms and Molecules*, *Nat. Phys.* **17**, 1324 (2021).
- [60] S. L. Cornish, M. R. Tarbutt, and K. R. A. Hazzard, *Quantum Computation and Quantum Simulation with Ultracold Molecules*, *Nat. Phys.* **20**, 730 (2024).
- [61] R. Blatt and C. F. Roos, *Quantum Simulations with Trapped Ions*, *Nat. Phys.* **8**, 277 (2012).
- [62] S.-A. Guo et al., *A Site-Resolved Two-Dimensional Quantum Simulator with Hundreds of Trapped Ions*, *Nature* **630**, 613 (2024).
- [63] A. Aspuru-Guzik and P. Walther, *Photonic Quantum Simulators*, *Nat. Phys.* **8**, 285 (2012).
- [64] A. A. Houck, H. E. Türeci, and J. Koch, *On-Chip Quantum Simulation with Superconducting Circuits*, *Nat. Phys.* **8**, 292 (2012).
- [65] A. J. Daley, I. Bloch, C. Kokail, S. Flannigan, N. Pearson, M. Troyer, and P. Zoller, *Practical Quantum Advantage in Quantum Simulation*, *Nature* **607**, 667 (2022).
- [66] J. Preskill, *Quantum Computing in the NISQ Era and Beyond*, *Quantum* **2**, 79 (2018).
- [67] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition*, 1st ed. (June 2012).
- [68] L. Henriot, L. Beguin, A. Signoles, T. Lahaye, A. Browaeys, G.-O. Reymond, and C. Jurczak, *Quantum Computing with Neutral Atoms*, *Quantum* **4**, 327 (2020).
- [69] R. Trivedi, A. Franco Rubio, and J. I. Cirac, *Quantum Advantage and Stability to Errors in Analogue Quantum Simulators*, *Nat. Commun.* **15**, 6507 (2024).
- [70] R. Yamamoto, H. Ozawa, D. C. Nak, I. Nakamura, and T. Fukuhara, *Single-Site-Resolved Imaging of Ultracold Atoms in a Triangular Optical Lattice*, *New J. Phys.* **22**, 123028 (2020).
- [71] D. Garwood, J. Mongkolkiattichai, L. Liu, J. Yang, and P. Schauss, *Site-Resolved Observables in the Doped Spin-Imbalanced Triangular Hubbard Model*, *Phys. Rev. A* **106**, 013310 (2022).
- [72] J. Mongkolkiattichai, L. Liu, D. Garwood, J. Yang, and P. Schauss, *Quantum Gas Microscopy of Fermionic Triangular-Lattice Mott Insulators*, *Phys. Rev. A* **108**, L061301 (2023).
- [73] C. Robens, J. Zopes, W. Alt, S. Brakhane, D. Meschede, and A. Alberti, *Low-Entropy States of Neutral Atoms in Polarization-Synthesized Optical Lattices*, *Phys. Rev. Lett.* **118**, 065302 (2017).
- [74] J. Sebby-Strabley, M. Anderlini, P. S. Jessen, and J. V. Porto, *Lattice of Double Wells for Manipulating Pairs of Cold Atoms*, *Phys. Rev. A* **73**, 033605 (2006).
- [75] S. Fölling, S. Trotzky, P. Cheinet, M. Feld, R. Saers, A. Widera, T. Müller, and I. Bloch, *Direct Observation of Second-Order Atom Tunnelling*, *Nature* **448**, 1029 (2007).
- [76] S. Baier, D. Petter, J. H. Becher, A. Patscheider, G. Natale, L. Chomaz, M. J. Mark, and F. Ferlaino, *Realization of a Strongly Interacting Fermi Gas of Dipolar Atoms*, *Phys. Rev. Lett.* **121**, 093602 (2018).
- [77] L. Chomaz, I. Ferrier-Barbut, F. Ferlaino, B. Laburthe-Tolra, B. L. Lev, and T. Pfau, *Dipolar Physics: A Review of Experiments with Magnetic Quantum Gases*, *Rep. Prog. Phys.* **86**, 026401 (2022).

-
- [78] P. B. Wigley, P. J. Everitt, A. van den Hengel, J. W. Bastian, M. A. Sooriyabandara, G. D. McDonald, K. S. Hardman, C. D. Quinlivan, P. Manju, C. C. N. Kuhn, I. R. Petersen, A. N. Luiten, J. J. Hope, N. P. Robins, and M. R. Hush, *Fast Machine-Learning Online Optimization of Ultra-Cold-Atom Experiments*, *Sci. Rep.* **6**, 25890 (2016).
- [79] A. D. Tranter, H. J. Slatyer, M. R. Hush, A. C. Leung, J. L. Everett, K. V. Paul, P. Vernaz-Gris, P. K. Lam, B. C. Buchler, and G. T. Campbell, *Multiparameter Optimisation of a Magneto-Optical Trap Using Deep Learning*, *Nat. Commun.* **9**, 4360 (2018).
- [80] J. Carrasquilla, G. Torlai, R. G. Melko, and L. Aolita, *Reconstructing Quantum States with Generative Models*, *Nat. Mach. Intell.* **1**, 155 (2019).
- [81] B. S. Rem, N. Käming, M. Tarnowski, L. Asteria, N. Fläschner, C. Becker, K. Sengstock, and C. Weitenberg, *Identifying Quantum Phase Transitions Using Artificial Neural Networks on Experimental Data*, *Nat. Phys.* **15**, 917 (2019).
- [82] A. Bohrdt, C. S. Chiu, G. Ji, M. Xu, D. Greif, M. Greiner, E. Demler, F. Grusdt, and M. Knap, *Classifying Snapshots of the Doped Hubbard Model with Machine Learning*, *Nat. Phys.* **15**, 921 (2019).
- [83] L. R. B. Picard, M. J. Mark, F. Ferlaino, and R. van Bijnen, *Deep Learning-Assisted Classification of Site-Resolved Quantum Gas Microscope Images*, *Meas. Sci. Technol.* **31**, 025201 (2019).
- [84] G. Carleo, I. Cirac, K. Cranmer, L. Daudet, M. Schuld, N. Tishby, L. Vogt-Maranto, and L. Zdeborová, *Machine Learning and the Physical Sciences*, *Rev. Mod. Phys.* **91**, 045002 (2019).
- [85] J. Carrasquilla, *Machine Learning for Quantum Matter*, *Adv. Phys.* **X 5**, 1797528 (2020).
- [86] E. Khatami, E. Guardado-Sanchez, B. M. Spar, J. F. Carrasquilla, W. S. Bakr, and R. T. Scalettar, *Visualizing Strange Metallic Correlations in the Two-Dimensional Fermi-Hubbard Model with Artificial Intelligence*, *Phys. Rev. A* **102**, 033326 (2020).
- [87] N. Käming, A. Dawid, K. Kottmann, M. Lewenstein, K. Sengstock, A. Dauphin, and C. Weitenberg, *Unsupervised Machine Learning of Topological Phase Transitions from Experimental Data*, *Mach. Learn.: Sci. Technol.* **2**, 035037 (2021).
- [88] C. Weitenberg, P. Schauß, T. Fukuhara, M. Cheneau, M. Endres, I. Bloch, and S. Kuhr, *Coherent Light Scattering from a Two-Dimensional Mott Insulator*, *Phys. Rev. Lett.* **106**, 215301 (2011).
- [89] S. J. Masson and A. Asenjo-Garcia, *Universality of Dicke Superradiance in Arrays of Quantum Emitters*, *Nat. Commun.* **13**, 2285 (2022).
- [90] S. Keßler and F. Marquardt, *Single-Site-Resolved Measurement of the Current Statistics in Optical Lattices*, *Phys. Rev. A* **89**, 061601 (2014).
- [91] B. Wang, X. Dong, and A. Eckardt, *Measurable Signatures of Bosonic Fractional Chern Insulator States and Their Fractional Excitations in a Quantum-Gas Microscope*, *SciPost Phys.* **12**, 095 (2022).
- [92] A. M. Kaufman, B. J. Lester, C. M. Reynolds, M. L. Wall, M. Foss-Feig, K. R. A. Hazzard, A. M. Rey, and C. A. Regal, *Two-Particle Quantum Interference in Tunnel-Coupled Optical Tweezers*, *Science* **345**, 306 (2014).
- [93] M. Atala, M. Aidelsburger, M. Lohse, J. T. Barreiro, B. Paredes, and I. Bloch, *Observation of Chiral Currents with Ultracold Atoms in Bosonic Ladders*, *Nat. Phys.* **10**, 588 (2014).
- [94] D. K. Mark, H.-Y. Hu, J. Kwan, C. Kokail, S. Choi, and S. F. Yelin, *Efficiently Measuring d -Wave Pairing and beyond in Quantum Gas Microscopes*, [arXiv:2412.13186](https://arxiv.org/abs/2412.13186).

-
- [95] D. González-Cuadra, D. Bluvstein, M. Kalinowski, R. Kaubruegger, N. Maskara, P. Naldesi, T. V. Zache, A. M. Kaufman, M. D. Lukin, H. Pichler, B. Vermersch, J. Ye, and P. Zoller, *Fermionic Quantum Processing with Programmable Neutral Atom Arrays*, *PNAS* **120**, e2304294120 (2023).
- [96] F. Gkritis, D. Dux, J. Zhang, N. Jain, C. Gogolin, and P. M. Preiss, *Simulating Chemistry with Fermionic Optical Superlattices*, *PRX Quantum* **6**, 010318 (2025).
- [97] G. Bornet, G. Emperauger, C. Chen, F. Machado, S. Chern, L. Leclerc, B. Gély, Y. T. Chew, D. Barredo, T. Lahaye, N. Y. Yao, and A. Browaeys, *Enhancing a Many-Body Dipolar Rydberg Tweezer Array with Arbitrary Local Controls*, *Phys. Rev. Lett.* **132**, 263601 (2024).
- [98] B. Du, R. Suresh, S. López, J. Cadiente, and R. Ma, *Probing Site-Resolved Current in Strongly Interacting Superconducting Circuit Lattices*, *Phys. Rev. Lett.* **133**, 060601 (2024).
- [99] D. C. Tsui, H. L. Stormer, and A. C. Gossard, *Two-Dimensional Magnetotransport in the Extreme Quantum Limit*, *Phys. Rev. Lett.* **48**, 1559 (1982).
- [100] R. B. Laughlin, *Anomalous Quantum Hall Effect: An Incompressible Quantum Fluid with Fractionally Charged Excitations*, *Phys. Rev. Lett.* **50**, 1395 (1983).
- [101] A. Kitaev, *Anyons in an Exactly Solved Model and Beyond*, *Ann. Phys.* **321**, 2 (2006).
- [102] A. Stern, *Anyons and the Quantum Hall Effect—A Pedagogical Review*, *Ann. Phys.* **323**, 204 (2008).
- [103] J. Liang, Z. Liu, Z. Yang, Y. Huang, U. Wurstbauer, C. R. Dean, K. W. West, L. N. Pfeiffer, L. Du, and A. Pinczuk, *Evidence for Chiral Graviton Modes in Fractional Quantum Hall Liquids*, *Nature* **628**, 78 (2024).
- [104] A. Y. Kitaev, *Fault-Tolerant Quantum Computation by Anyons*, *Ann. Phys.* **303**, 2 (2003).
- [105] C. Nayak, S. H. Simon, A. Stern, M. Freedman, and S. Das Sarma, *Non-Abelian Anyons and Topological Quantum Computation*, *Rev. Mod. Phys.* **80**, 1083 (2008).
- [106] N. Goldman, J. C. Budich, and P. Zoller, *Topological Quantum Matter with Ultracold Gases in Optical Lattices*, *Nat. Phys.* **12**, 639 (2016).
- [107] M. Iqbal et al., *Non-Abelian Topological Order and Anyons on a Trapped-Ion Processor*, *Nature* **626**, 505 (2024).
- [108] M. Iqbal et al., *Topological Order from Measurements and Feed-Forward on a Trapped Ion Quantum Computer*, *Commun. Phys.* **7**, 1 (2024).
- [109] S. Nascimbene, *Simulating quantum Hall physics in ultracold atomic gases: prospects and challenges*, *C. R. Phys.* **26**, 317 (2025).
- [110] N. R. Cooper, J. Dalibard, and I. B. Spielman, *Topological Bands for Ultracold Atoms*, *Rev. Mod. Phys.* **91**, 015005 (2019).
- [111] J. Léonard, S. Kim, J. Kwan, P. Segura, F. Grusdt, C. Repellin, N. Goldman, and M. Greiner, *Realization of a Fractional Quantum Hall State with Ultracold Atoms*, *Nature* **619**, 495 (2023).
- [112] P. Lunt, P. Hill, J. Reiter, P. M. Preiss, M. Gałka, and S. Jochim, *Realization of a Laughlin State of Two Rapidly Rotating Fermions*, *Phys. Rev. Lett.* **133**, 253401 (2024).
- [113] N. Goldman and J. Dalibard, *Periodically Driven Quantum Systems: Effective Hamiltonians and Engineered Gauge Fields*, *Phys. Rev. X* **4**, 031027 (2014).
- [114] A. Eckardt, *Colloquium: Atomic Quantum Gases in Periodically Driven Optical Lattices*, *Rev. Mod. Phys.* **89**, 011004 (2017).

-
- [115] C. Weitenberg and J. Simonet, *Tailoring Quantum Gases by Floquet Engineering*, *Nat. Phys.* **17**, 1342 (2021).
- [116] R. Grimm, M. Weidemüller, and Y. B. Ovchinnikov, *Optical Dipole Traps for Neutral Atoms*, in *Advances In Atomic, Molecular, and Optical Physics*, Vol. 42, edited by B. Bederson and H. Walther (Jan. 2000), 95.
- [117] P. S. Jessen and I. H. Deutsch, *Optical Lattices*, in *Advances In Atomic, Molecular, and Optical Physics*, Vol. 37, edited by B. Bederson and H. Walther (Jan. 1996), 95.
- [118] M. Lewenstein, A. Sanpera, and V. Ahufinger, *Ultracold Atoms in Optical Lattices: Simulating Quantum Many-Body Systems* (Mar. 2012).
- [119] N. W. Ashcroft and N. D. Mermin, *Solid State Physics* (New York, 1976).
- [120] M. P. A. Fisher, P. B. Weichman, G. Grinstein, and D. S. Fisher, *Boson Localization and the Superfluid-Insulator Transition*, *Phys. Rev. B* **40**, 546 (1989).
- [121] I. Bloch, J. Dalibard, and W. Zwerger, *Many-Body Physics with Ultracold Gases*, *Rev. Mod. Phys.* **80**, 885 (2008).
- [122] B. Capogrosso-Sansone, Ş. G. Söyler, N. Prokof'ev, and B. Svistunov, *Monte Carlo Study of the Two-Dimensional Bose-Hubbard Model*, *Phys. Rev. A* **77**, 015602 (2008).
- [123] S. Sachdev, *Quantum Phase Transitions*, 2nd ed. (Cambridge, 2011).
- [124] A. Impertro, S. Karch, J. F. Wienand, S. Huh, C. Schweizer, I. Bloch, and M. Aidelsburger, *Local Readout and Control of Current and Kinetic Energy Operators in Optical Lattices*, *Phys. Rev. Lett.* **133**, 063401 (2024).
- [125] L. Landau, *On the Theory of Phase Transitions*, *Ukr. J. Phys.* **53** (2008).
- [126] P. C. Hohenberg and A. P. Krekhov, *An Introduction to the Ginzburg–Landau Theory of Phase Transitions and Nonequilibrium Patterns*, *Phys. Rep.* **572**, 1 (2015).
- [127] J. G. Bednorz and K. A. Müller, *Possible High-Tc Superconductivity in the Ba-La-Cu-O System*, *Z. Physik B - Cond. Mat.* **64**, 189 (1986).
- [128] X. G. Wen, *Vacuum Degeneracy of Chiral Spin States in Compactified Space*, *Phys. Rev. B* **40**, 7387 (1989).
- [129] X. G. Wen, *Topological Orders in Rigid States*, *Int. J. Mod. Phys. B* **04**, 239 (1990).
- [130] X.-G. Wen, *Quantum Field Theory of Many-Body Systems: From the Origin of Sound to an Origin of Light and Electrons* (Sept. 2007).
- [131] X. G. Wen and Q. Niu, *Ground-State Degeneracy of the Fractional Quantum Hall States in the Presence of a Random Potential and on High-Genus Riemann Surfaces*, *Phys. Rev. B* **41**, 9377 (1990).
- [132] F. Wilczek, *Magnetic Flux, Angular Momentum, and Statistics*, *Phys. Rev. Lett.* **48**, 1144 (1982).
- [133] D. Arovas, J. R. Schrieffer, and F. Wilczek, *Fractional Statistics and the Quantum Hall Effect*, *Phys. Rev. Lett.* **53**, 722 (1984).
- [134] D. P. Arovas, R. Schrieffer, F. Wilczek, and A. Zee, *Statistical Mechanics of Anyons*, *Nucl. Phys. B* **251**, 117 (1985).
- [135] X. G. Wen, *Non-Abelian Statistics in the Fractional Quantum Hall States*, *Phys. Rev. Lett.* **66**, 802 (1991).
- [136] X.-G. Wen, *Topological Orders and Edge Excitations in Fractional Quantum Hall States*, *Adv. Phys.*, 10.1080/00018739500101566 (1995).

-
- [137] X. Chen, Z.-X. Liu, and X.-G. Wen, *Two-Dimensional Symmetry-Protected Topological Orders and Their Protected Gapless Edge Excitations*, *Phys. Rev. B* **84**, 235141 (2011).
- [138] A. Kitaev and J. Preskill, *Topological Entanglement Entropy*, *Phys. Rev. Lett.* **96**, 110404 (2006).
- [139] M. Levin and X.-G. Wen, *Detecting Topological Order in a Ground State Wave Function*, *Phys. Rev. Lett.* **96**, 110405 (2006).
- [140] X. Chen, Z.-C. Gu, and X.-G. Wen, *Local Unitary Transformation, Long-Range Quantum Entanglement, Wave Function Renormalization, and Topological Order*, *Phys. Rev. B* **82**, 155138 (2010).
- [141] T. Grover, *Entanglement Entropy and Strongly Correlated Topological Matter*, *Mod. Phys. Lett. A* **28**, 10.1142/S0217732313300012 (2013).
- [142] H. L. Stormer, *Nobel Lecture: The Fractional Quantum Hall Effect*, *Rev. Mod. Phys.* **71**, 875 (1999).
- [143] K. V. Klitzing, G. Dorda, and M. Pepper, *New Method for High-Accuracy Determination of the Fine-Structure Constant Based on Quantized Hall Resistance*, *Phys. Rev. Lett.* **45**, 494 (1980).
- [144] P. G. Harper, *Single Band Motion of Conduction Electrons in a Uniform Magnetic Field*, *Proc. Phys. Soc. A* **68**, 874 (1955).
- [145] D. R. Hofstadter, *Energy Levels and Wave Functions of Bloch Electrons in Rational and Irrational Magnetic Fields*, *Phys. Rev. B* **14**, 2239 (1976).
- [146] R. Peierls, *Zur Theorie des Diamagnetismus von Leitungselektronen*, *Z. Physik* **80**, 763 (1933).
- [147] D. J. Thouless, M. Kohmoto, M. P. Nightingale, and M. den Nijs, *Quantized Hall Conductance in a Two-Dimensional Periodic Potential*, *Phys. Rev. Lett.* **49**, 405 (1982).
- [148] F. D. M. Haldane, *Model for a Quantum Hall Effect without Landau Levels: Condensed-Matter Realization of the "Parity Anomaly"*, *Phys. Rev. Lett.* **61**, 2015 (1988).
- [149] N. Regnault and B. A. Bernevig, *Fractional Chern Insulator*, *Phys. Rev. X* **1**, 021014 (2011).
- [150] Z. Liu and E. J. Bergholtz, *Recent Developments in Fractional Chern Insulators*, in *Encyclopedia of Condensed Matter Physics (Second Edition)*, edited by T. Chakraborty (Oxford, Jan. 2024), 515.
- [151] A. S. Sørensen, E. Demler, and M. D. Lukin, *Fractional Quantum Hall States of Atoms in Optical Lattices*, *Phys. Rev. Lett.* **94**, 086803 (2005).
- [152] R. N. Palmer and D. Jaksch, *High-Field Fractional Quantum Hall Effect in Optical Lattices*, *Phys. Rev. Lett.* **96**, 180407 (2006).
- [153] M. Hafezi, A. S. Sørensen, E. Demler, and M. D. Lukin, *Fractional Quantum Hall Effect in Optical Lattices*, *Phys. Rev. A* **76**, 023613 (2007).
- [154] R. N. Palmer, A. Klein, and D. Jaksch, *Optical Lattice Quantum Hall Effect*, *Phys. Rev. A* **78**, 013609 (2008).
- [155] G. Möller and N. R. Cooper, *Composite Fermion Theory for Bosonic Quantum Hall States on Lattices*, *Phys. Rev. Lett.* **103**, 105303 (2009).
- [156] S. A. Parameswaran, R. Roy, and S. L. Sondhi, *Fractional quantum Hall physics in topological flat bands*, *C. R. Phys.* **14**, 816 (2013).
- [157] E. Kapit and E. Mueller, *Exact Parent Hamiltonian for the Quantum Hall States in a Lattice*, *Phys. Rev. Lett.* **105**, 215303 (2010).
- [158] G. Möller and N. R. Cooper, *Fractional Chern Insulators in Harper-Hofstadter Bands with Higher Chern Number*, *Phys. Rev. Lett.* **115**, 126401 (2015).

-
- [159] J. Behrmann, Z. Liu, and E. J. Bergholtz, *Model Fractional Chern Insulators*, *Phys. Rev. Lett.* **116**, 216802 (2016).
- [160] Y.-C. He, F. Grusdt, A. Kaufman, M. Greiner, and A. Vishwanath, *Realizing and Adiabatically Preparing Bosonic Integer and Fractional Quantum Hall States in Optical Lattices*, *Phys. Rev. B* **96**, 201103 (2017).
- [161] A. Dhar, T. Mishra, M. Maji, R. V. Pai, S. Mukerjee, and A. Paramekanti, *Chiral Mott Insulator with Staggered Loop Currents in the Fully Frustrated Bose-Hubbard Model*, *Phys. Rev. B* **87**, 174501 (2013).
- [162] M. P. Zaletel, S. A. Parameswaran, A. Rüegg, and E. Altman, *Chiral Bosonic Mott Insulator on the Frustrated Triangular Lattice*, *Phys. Rev. B* **89**, 155142 (2014).
- [163] A. Petrescu and K. Le Hur, *Chiral Mott Insulators, Meissner Effect, and Laughlin States in Quantum Ladders*, *Phys. Rev. B* **91**, 054520 (2015).
- [164] I. Vasić, A. Petrescu, K. Le Hur, and W. Hofstetter, *Chiral Bosonic Phases on the Haldane Honeycomb Lattice*, *Phys. Rev. B* **91**, 094502 (2015).
- [165] M. Piraud, F. Heidrich-Meisner, I. P. McCulloch, S. Greschner, T. Vekua, and U. Schollwöck, *Vortex and Meissner Phases of Strongly Interacting Bosons on a Two-Leg Ladder*, *Phys. Rev. B* **91**, 140406 (2015).
- [166] V. Kalmeyer and R. B. Laughlin, *Equivalence of the Resonating-Valence-Bond and Fractional Quantum Hall States*, *Phys. Rev. Lett.* **59**, 2095 (1987).
- [167] G. Möller and N. R. Cooper, *Condensed Ground States of Frustrated Bose-Hubbard Models*, *Phys. Rev. A* **82**, 063625 (2010).
- [168] N. Lang and H. P. Büchler, *Topological States in a Microscopic Model of Interacting Fermions*, *Phys. Rev. B* **92**, 041118 (2015).
- [169] D. Jaksch and P. Zoller, *Creation of Effective Magnetic Fields in Optical Lattices: The Hofstadter Butterfly for Cold Neutral Atoms*, *New J. Phys.* **5**, 56 (2003).
- [170] M. Aidelsburger, M. Atala, S. Nascimbène, S. Trotzky, Y.-A. Chen, and I. Bloch, *Experimental Realization of Strong Effective Magnetic Fields in an Optical Lattice*, *Phys. Rev. Lett.* **107**, 255301 (2011).
- [171] O. Boada, A. Celi, J. I. Latorre, and M. Lewenstein, *Quantum Simulation of an Extra Dimension*, *Phys. Rev. Lett.* **108**, 133001 (2012).
- [172] J. Argüello-Luengo, U. Bhattacharya, A. Celi, R. W. Chhajlany, T. Grass, M. Płodzień, D. Rakshit, T. Salamon, P. Stornati, L. Tarruell, and M. Lewenstein, *Synthetic Dimensions for Topological and Quantum Phases*, *Commun. Phys.* **7**, 1 (2024).
- [173] A. Celi, P. Massignan, J. Ruseckas, N. Goldman, I. B. Spielman, G. Juzeliūnas, and M. Lewenstein, *Synthetic Gauge Fields in Synthetic Dimensions*, *Phys. Rev. Lett.* **112**, 043001 (2014).
- [174] M. Mancini, G. Pagano, G. Cappellini, L. Livi, M. Rider, J. Catani, C. Sias, P. Zoller, M. Inguscio, M. Dalmonte, and L. Fallani, *Observation of Chiral Edge States with Neutral Fermions in Synthetic Hall Ribbons*, *Science* **349**, 1510 (2015).
- [175] B. K. Stuhl, H.-I. Lu, L. M. Ayccock, D. Genkina, and I. B. Spielman, *Visualizing Edge States with an Atomic Bose Gas in the Quantum Hall Regime*, *Science* **349**, 1514 (2015).
- [176] T.-W. Zhou, G. Cappellini, D. Tusi, L. Franchi, J. Parravicini, C. Repellin, S. Greschner, M. Inguscio, T. Giamarchi, M. Filippone, J. Catani, and L. Fallani, *Observation of Universal Hall Response in Strongly Interacting Fermions*, *Science* **381**, 427 (2023).

- [177] K. W. Madison, F. Chevy, W. Wohlleben, and J. Dalibard, *Vortex Formation in a Stirred Bose-Einstein Condensate*, *Phys. Rev. Lett.* **84**, 806 (2000).
- [178] R. J. Fletcher, A. Shaffer, C. C. Wilson, P. B. Patel, Z. Yan, V. Crépel, B. Mukherjee, and M. W. Zwierlein, *Geometric Squeezing into the Lowest Landau Level*, *Science* **372**, 1318 (2021).
- [179] B. Mukherjee, A. Shaffer, P. B. Patel, Z. Yan, C. C. Wilson, V. Crépel, R. J. Fletcher, and M. Zwierlein, *Crystallization of Bosonic Quantum Hall States in a Rotating Quantum Gas*, *Nature* **601**, 58 (2022).
- [180] L. W. Clark, N. Schine, C. Baum, N. Jia, and J. Simon, *Observation of Laughlin States Made of Light*, *Nature* **582**, 41 (2020).
- [181] C. Wang, F.-M. Liu, M.-C. Chen, H. Chen, X.-H. Zhao, C. Ying, Z.-X. Shang, J.-W. Wang, Y.-H. Huo, C.-Z. Peng, X. Zhu, C.-Y. Lu, and J.-W. Pan, *Realization of Fractional Quantum Hall State with Interacting Photons*, *Science* **384**, 579 (2024).
- [182] M. Račiūnas, F. N. Ünal, E. Anisimovas, and A. Eckardt, *Creating, Probing, and Manipulating Fractionally Charged Excitations of Fractional Chern Insulators in Optical Lattices*, *Phys. Rev. A* **98**, 063621 (2018).
- [183] T. Comparin, A. Opler, E. Macaluso, A. Biella, A. P. Polychronakos, and L. Mazza, *Measurable Fractional Spin for Quantum Hall Quasiparticles on the Disk*, *Phys. Rev. B* **105**, 085125 (2022).
- [184] B. Paredes, P. Fedichev, J. I. Cirac, and P. Zoller, $\frac{1}{2}$ -Anyons in Small Atomic Bose-Einstein Condensates, *Phys. Rev. Lett.* **87**, 010402 (2001).
- [185] T. Ozawa, H. M. Price, A. Amo, N. Goldman, M. Hafezi, L. Lu, M. C. Rechtsman, D. Schuster, J. Simon, O. Zilberberg, and I. Carusotto, *Topological Photonics*, *Rev. Mod. Phys.* **91**, 015006 (2019).
- [186] M. Bukov, L. D'Alessio, and A. Polkovnikov, *Universal High-Frequency Behavior of Periodically Driven Systems: From Dynamical Stabilization to Floquet Engineering*, *Adv. Phys.* **64**, 139 (2015).
- [187] P. Roushan et al., *Chiral Ground-State Currents of Interacting Photons in a Synthetic Magnetic Field*, *Nat. Phys.* **13**, 146 (2017).
- [188] I. T. Rosen et al., *A Synthetic Magnetic Vector Potential in a 2D Superconducting Qubit Array*, *Nat. Phys.* **20**, 1881 (2024).
- [189] P. Scholl et al., *Microwave Engineering of Programmable XXZ Hamiltonians in Arrays of Rydberg Atoms*, *PRX Quantum* **3**, 020303 (2022).
- [190] L. Zhao, M. D. K. Lee, M. M. Aliyu, and H. Loh, *Floquet-Tailored Rydberg Interactions*, *Nat. Commun.* **14**, 7128 (2023).
- [191] J. W. McIver, B. Schulte, F.-U. Stein, T. Matsuyama, G. Jotzu, G. Meier, and A. Cavalleri, *Light-Induced Anomalous Hall Effect in Graphene*, *Nat. Phys.* **16**, 38 (2020).
- [192] T. Weitz, D. M. B. Lesko, S. Wittigshlager, W. Li, C. Heide, O. Neufeld, and P. Hommelhoff, *Lightwave-Driven Electrons in a Floquet Topological Insulator*, [arXiv:2407.17917](https://arxiv.org/abs/2407.17917).
- [193] M. Aidelsburger, *Artificial Gauge Fields with Ultracold Atoms in Optical Lattices*, Springer Theses (Cham, 2016).
- [194] W. R. Salzman, *Quantum Mechanics of Systems Periodic in Time*, *Phys. Rev. A* **10**, 461 (1974).
- [195] S. Rahav, I. Gilary, and S. Fishman, *Effective Hamiltonians for Periodically Driven Systems*, *Phys. Rev. A* **68**, 013820 (2003).

-
- [196] F. Casas, J. A. Oteo, and J. Ros, *Floquet Theory: Exponential Perturbative Treatment*, *J. Phys. A: Math. Gen.* **34**, 3379 (2001).
- [197] A. Eckardt and E. Anisimovas, *High-Frequency Approximation for Periodically Driven Quantum Systems from a Floquet-space Perspective*, *New J. Phys.* **17**, 093039 (2015).
- [198] A. Lazarides, A. Das, and R. Moessner, *Equilibrium States of Generic Quantum Systems Subject to Periodic Driving*, *Phys. Rev. E* **90**, 012110 (2014).
- [199] L. D'Alessio and M. Rigol, *Long-Time Behavior of Isolated Periodically Driven Interacting Lattice Systems*, *Phys. Rev. X* **4**, 041048 (2014).
- [200] P. Ponte, A. Chandran, Z. Papić, and D. A. Abanin, *Periodically Driven Ergodic and Many-Body Localized Quantum Systems*, *Ann. Phys.* **353**, 196 (2015).
- [201] M. Bukov, S. Gopalakrishnan, M. Knap, and E. Demler, *Prethermal Floquet Steady States and Instabilities in the Periodically Driven, Weakly Interacting Bose-Hubbard Model*, *Phys. Rev. Lett.* **115**, 205301 (2015).
- [202] S. Thanasilp, J. Tangpanitanon, M.-A. Lemonde, N. Dangniam, and D. G. Angelakis, *Quantum Supremacy and Quantum Phase Transitions*, *Phys. Rev. B* **103**, 165132 (2021).
- [203] Y.-G. Zheng et al., *Efficiently Extracting Multi-Point Correlations of a Floquet Thermalized System*, [arXiv:2210.08556](https://arxiv.org/abs/2210.08556).
- [204] G. Sun and A. Eckardt, *Optimal Frequency Window for Floquet Engineering in Optical Lattices*, *Phys. Rev. Res.* **2**, 013241 (2020).
- [205] T. Bilitewski and N. R. Cooper, *Scattering Theory for Floquet-Bloch States*, *Phys. Rev. A* **91**, 033601 (2015).
- [206] T. Bilitewski and N. R. Cooper, *Population Dynamics in a Floquet Realization of the Harper-Hofstadter Hamiltonian*, *Phys. Rev. A* **91**, 063611 (2015).
- [207] M. Bukov, M. Heyl, D. A. Huse, and A. Polkovnikov, *Heating and Many-Body Resonances in a Periodically Driven Two-Band System*, *Phys. Rev. B* **93**, 155132 (2016).
- [208] M. Reitter, J. Näger, K. Wintersperger, C. Sträter, I. Bloch, A. Eckardt, and U. Schneider, *Interaction Dependent Heating and Atom Loss in a Periodically Driven Optical Lattice*, *Phys. Rev. Lett.* **119**, 200402 (2017).
- [209] T. Boulier, J. Maslek, M. Bukov, C. Bracamontes, E. Magnan, S. Lellouch, E. Demler, N. Goldman, and J. V. Porto, *Parametric Heating in a 2D Periodically Driven Bosonic System: Beyond the Weakly Interacting Regime*, *Phys. Rev. X* **9**, 011047 (2019).
- [210] K. Wintersperger, M. Bukov, J. Näger, S. Lellouch, E. Demler, U. Schneider, I. Bloch, N. Goldman, and M. Aidelsburger, *Parametric Instabilities of Interacting Bosons in Periodically Driven 1D Optical Lattices*, *Phys. Rev. X* **10**, 011030 (2020).
- [211] T. M. Klostermann, *Construction of a caesium quantum gas microscope*, [PhD thesis](#), Ludwig-Maximilians-Universität München (2022).
- [212] H. von Raven, *A new Caesium quantum gas microscope with precise magnetic field control*, [PhD thesis](#), Ludwig-Maximilians-Universität München (2022).
- [213] J. F. Wienand, *Quantum gas microscopy of fluctuating hydrodynamics in optical ladders*, [PhD thesis](#), Ludwig-Maximilians-Universität München (2024).
- [214] G. Santarelli, P. Laurent, P. Lemonde, A. Clairon, A. G. Mann, S. Chang, A. N. Luiten, and C. Salomon, *Quantum Projection Noise in an Atomic Fountain: A High Stability Cesium Frequency Standard*, *Phys. Rev. Lett.* **82**, 4619 (1999).

- [215] M. Snadden, J. McGuirk, P. Bouyer, K. Haritos, and M. Kasevich, *Measurement of the Earth's Gravity Gradient with an Atom Interferometer-Based Gravity Gradiometer*, *Phys. Rev. Lett.* **81**, 971 (1998).
- [216] C. Chin, V. Leiber, V. Vuletić, A. J. Kerman, and S. Chu, *Measurement of an Electron's Electric Dipole Moment Using Cs Atoms Trapped in Optical Lattices*, *Phys. Rev. A* **63**, 033401 (2001).
- [217] T. Weber, J. Herbig, M. Mark, H.-C. Nägerl, and R. Grimm, *Bose-Einstein Condensation of Cesium*, *Science* **299**, 232 (2003).
- [218] P. J. Leo, C. J. Williams, and P. S. Julienne, *Collision Properties of Ultracold Cs-133 Atoms*, *Phys. Rev. Lett.* **85**, 2721 (2000).
- [219] A. J. Kerman, C. Chin, V. Vuletić, S. Chu, P. J. Leo, C. J. Williams, and P. S. Julienne, *Determination of Cs-Cs Interaction Parameters Using Feshbach Spectroscopy*, *C. R. Acad. Sci. - S. IV - Phys.* **2**, 633 (2001).
- [220] C. Chin, R. Grimm, P. Julienne, and E. Tiesinga, *Feshbach Resonances in Ultracold Gases*, *Rev. Mod. Phys.* **82**, 1225 (2010).
- [221] C. Chin, V. Vuletić, A. J. Kerman, and S. Chu, *High Resolution Feshbach Spectroscopy of Cesium*, *Phys. Rev. Lett.* **85**, 2717 (2000).
- [222] M. Gustavsson, E. Haller, M. J. Mark, J. G. Danzl, G. Rojas-Kopeinig, and H.-C. Nägerl, *Control of Interaction-Induced Dephasing of Bloch Oscillations*, *Phys. Rev. Lett.* **100**, 080404 (2008).
- [223] T. Klostermann, C. R. Cabrera, H. von Raven, J. F. Wienand, C. Schweizer, I. Bloch, and M. Aidelsburger, *Fast Long-Distance Transport of Cold Cesium Atoms*, *Phys. Rev. A* **105**, 043319 (2022).
- [224] E. Haller, M. J. Mark, R. Hart, J. G. Danzl, L. Reichsöllner, V. Melezhik, P. Schmelcher, and H.-C. Nägerl, *Confinement-Induced Resonances in Low-Dimensional Quantum Systems*, *Phys. Rev. Lett.* **104**, 153203 (2010).
- [225] I. Pérez Ramos, *Single-Site Addressing and Imaging of Cesium Atoms in a Quantum Gas Microscope*, MA thesis, Ludwig-Maximilians-Universität München (2023).
- [226] J. F. Wienand, S. Karch, A. Impertro, C. Schweizer, E. McCulloch, R. Vasseur, S. Gopalakrishnan, M. Aidelsburger, and I. Bloch, *Emergence of Fluctuating Hydrodynamics in Chaotic Quantum Systems*, *Nat. Phys.* **20**, 1732 (2024).
- [227] S. P. Hubele, *Potential Shaping Using a DMD and High-Resolution Imaging of Cesium Atoms in Optical Lattices*, MA thesis, Ludwig-Maximilians-Universität München (2022).
- [228] C. Shannon, *Communication in the Presence of Noise*, *Proc. IRE* **37**, 10 (1949).
- [229] P. A. Murthy, D. Kedar, T. Lompe, M. Neidig, M. G. Ries, A. N. Wenz, G. Zürn, and S. Jochim, *Matter-Wave Fourier Optics with a Strongly Interacting Two-Dimensional Fermi Gas*, *Phys. Rev. A* **90**, 043611 (2014).
- [230] M. Lohse, C. Schweizer, H. M. Price, O. Zilberberg, and I. Bloch, *Exploring 4D Quantum Hall Physics with a 2D Topological Charge Pump*, *Nature* **553**, 55 (2018).
- [231] M. N. Rispoli, *Microscopy of Correlations at a Non-Equilibrium Phase Transition*, PhD thesis, Harvard University (2019).
- [232] F. Meinert, M. J. Mark, E. Kirilov, K. Lauber, P. Weinmann, M. Gröbner, A. J. Daley, and H.-C. Nägerl, *Observation of Many-Body Dynamics in Long-Range Tunneling after a Quantum Quench*, *Science* **344**, 1259 (2014).

-
- [233] C. K. Hong, Z. Y. Ou, and L. Mandel, *Measurement of Subpicosecond Time Intervals between Two Photons by Interference*, *Phys. Rev. Lett.* **59**, 2044 (1987).
- [234] S. A. Häfele, *Deep Learning for Quantum Gas Microscopy of Cesium Atoms in Optical Lattices – Below the Resolution Limit*, MA thesis, Ludwig-Maximilians-Universität München (2022).
- [235] L. Su, A. Douglas, M. Szurek, A. H. Hebert, A. Krahn, R. Groth, G. A. Phelps, O. Markovic, and M. Greiner, *Fast Single Atom Imaging in Optical Lattice Arrays*, [arXiv:2404.09978](https://arxiv.org/abs/2404.09978).
- [236] N. Schlosser, G. Reymond, I. Protsenko, and P. Grangier, *Sub-Poissonian Loading of Single Atoms in a Microscopic Dipole Trap*, *Nature* **411**, 1024 (2001).
- [237] A. Alberti, C. Robens, W. Alt, S. Brakhane, M. Karski, R. Reimann, A. Widera, and D. Meschede, *Super-Resolution Microscopy of Single Atoms in Optical Lattices*, *New J. Phys.* **18**, 053010 (2016).
- [238] J. W. Goodman, *Introduction to Fourier Optics* (2005).
- [239] J. L. Starck, E. Pantin, and F. Murtagh, *Deconvolution in Astronomy: A Review*, *PASP* **114**, 1051 (2002).
- [240] D. L. Phillips, *A Technique for the Numerical Solution of Certain Integral Equations of the First Kind*, *J. ACM* **9**, 84 (1962).
- [241] N. Wiener, *Extrapolation, Interpolation, and Smoothing of Stationary Time Series* (1964).
- [242] A. La Rooij, C. Ulm, E. Haller, and S. Kuhr, *A Comparative Study of Deconvolution Techniques for Quantum-Gas Microscope Images*, [arXiv:2207.08663](https://arxiv.org/abs/2207.08663).
- [243] W. H. Richardson, *Bayesian-Based Iterative Method of Image Restoration*, *J. Opt. Soc. Am.* **62**, 55 (1972).
- [244] L. B. Lucy, *An Iterative Technique for the Rectification of Observed Distributions*, *Astron. J.* **79**, 745 (1974).
- [245] K. H. Jin, M. T. McCann, E. Froustey, and M. Unser, *Deep Convolutional Neural Network for Inverse Problems in Imaging*, *IEEE Trans. Image Process.* **26**, 4509 (2017).
- [246] S. Arridge, P. Maass, O. Öktem, and C. B. Schönlieb, *Solving Inverse Problems Using Data-Driven Models*, *Acta Numer.* **28**, 1 (2019).
- [247] L. W. Cheuk, M. A. Nichols, M. Okan, T. Gersdorf, V. V. Ramasesh, W. S. Bakr, T. Lompe, and M. W. Zwierlein, *Quantum-Gas Microscope for Fermionic Atoms*, *Phys. Rev. Lett.* **114**, 193001 (2015).
- [248] M. F. Parsons, F. Huber, A. Mazurenko, C. S. Chiu, W. Setiawan, K. Wooley-Brown, S. Blatt, and M. Greiner, *Site-Resolved Imaging of Fermionic ${}^6\text{Li}$ in an Optical Lattice*, *Phys. Rev. Lett.* **114**, 213002 (2015).
- [249] E. Haller, J. Hudson, A. Kelly, D. A. Cotta, B. Peaudecerf, G. D. Bruce, and S. Kuhr, *Single-Atom Imaging of Fermions in a Quantum-Gas Microscope*, *Nat. Phys.* **11**, 738 (2015).
- [250] G. J. A. Edge, R. Anderson, D. Jervis, D. C. McKay, R. Day, S. Trotzky, and J. H. Thywissen, *Imaging and Addressing of Individual Fermionic Atoms in an Optical Lattice*, *Phys. Rev. A* **92**, 063406 (2015).
- [251] A. Omran, M. Boll, T. A. Hilker, K. Kleinlein, G. Salomon, I. Bloch, and C. Gross, *Microscopic Observation of Pauli Blocking in Degenerate Fermionic Lattice Gases*, *Phys. Rev. Lett.* **115**, 263001 (2015).
- [252] R. Yamamoto, J. Kobayashi, T. Kuno, K. Kato, and Y. Takahashi, *An Ytterbium Quantum Gas Microscope with Narrow-Line Laser Cooling*, *New J. Phys.* **18**, 023016 (2016).

- [253] D. Mitra, P. T. Brown, E. Guardado-Sanchez, S. S. Kondov, T. Devakul, D. A. Huse, P. Schauf, and W. S. Bakr, *Quantum Gas Microscopy of an Attractive Fermi–Hubbard System*, *Nat. Phys.* **14**, 173 (2018).
- [254] M.-D. Li et al., *High-Powered Optical Superlattice with Robust Phase Stability for Quantum Gas Microscopy*, *Opt. Express* **29**, 13876 (2021).
- [255] J. Yang, L. Liu, J. Mongkolkittichai, and P. Schauss, *Site-Resolved Imaging of Ultracold Fermions in a Triangular-Lattice Quantum Gas Microscope*, *PRX Quantum* **2**, 020344 (2021).
- [256] K. Kwon, K. Kim, J. Hur, S. Huh, and J.-y. Choi, *Site-Resolved Imaging of a Bosonic Mott Insulator of Li-7 Atoms*, *Phys. Rev. A* **105**, 033323 (2022).
- [257] D. Greif, M. F. Parsons, A. Mazurenko, C. S. Chiu, S. Blatt, F. Huber, G. Ji, and M. Greiner, *Site-Resolved Imaging of a Fermionic Mott Insulator*, *Science* **351**, 953 (2016).
- [258] C. Roos, Unpublished.
- [259] D. P. Kingma and J. Ba, *Adam: A Method for Stochastic Optimization*, in *International Conference on Learning Representations (ICLR) 2015* (2015).
- [260] S. Raeisi and S. Raeisi, *Machine Learning for Physicists: A Hands-on Approach* (Nov. 2023).
- [261] Y. LeCun, Y. Bengio, and G. Hinton, *Deep Learning*, *Nature* **521**, 436 (2015).
- [262] L. Bottou and O. Bousquet, *The Tradeoffs of Large Scale Learning*, in *Adv. Neural Inf. Process. Syst.* Vol. 20 (2007).
- [263] T. P. Lillicrap, A. Santoro, L. Marris, C. J. Akerman, and G. Hinton, *Backpropagation and the Brain*, *Nat. Rev. Neurosci.* **21**, 335 (2020).
- [264] G. Cybenko, *Approximation by Superpositions of a Sigmoidal Function*, *Math. Contr. Sig. Sys.* **2**, 303 (1989).
- [265] Y. LeCun, B. Boser, J. Denker, D. Henderson, R. Howard, W. Hubbard, and L. Jackel, *Handwritten Digit Recognition with a Back-Propagation Network*, in *Adv. Neural Inf. Process. Syst.* Vol. 2 (1989).
- [266] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner, *Gradient-Based Learning Applied to Document Recognition*, *Proc. IEEE* **86**, 2278 (1998).
- [267] A. Tealab, *Time Series Forecasting Using Artificial Neural Networks Methodologies: A Systematic Review*, *FCIJ* **3**, 334 (2018).
- [268] L. Yang, X. Meng, and G. E. Karniadakis, *B-PINNs: Bayesian Physics-Informed Neural Networks for Forward and Inverse PDE Problems with Noisy Data*, *J. Comput. Phys.* **425**, 109913 (2021).
- [269] Z. Burns and Z. Liu, *Untrained, Physics-Informed Neural Networks for Structured Illumination Microscopy*, *Opt. Express* **31**, 8714 (2023).
- [270] N. Namaki, M. Eslahchi, and R. Salehi, *The Use of Physics-Informed Neural Network Approach to Image Restoration via Nonlinear PDE Tools*, *Comput. Math. Appl.* **152**, 355 (2023).
- [271] C. Banerjee, K. Nguyen, C. Fookes, and K. George, *Physics-Informed Computer Vision: A Review and Perspectives*, *ACM Comput. Surv.* **57**, 1 (2025).
- [272] C. Banerjee, K. Nguyen, O. Salvado, T. Tran, and C. Fookes, *PINNs for Medical Image Analysis: A Survey*, [arXiv:2408.01026](https://arxiv.org/abs/2408.01026).
- [273] M. Raissi, Z. Wang, M. S. Triantafyllou, and G. E. Karniadakis, *Deep Learning of Vortex-Induced Vibrations*, *J. Fluid Mech.* **861**, 119 (2019).

-
- [274] S. Cai, Z. Mao, Z. Wang, M. Yin, and G. E. Karniadakis, *Physics-Informed Neural Networks (PINNs) for Fluid Mechanics: A Review*, *Acta Mech. Sin.* **37**, 1727 (2021).
- [275] D. Xu and Y. Tian, *A Comprehensive Survey of Clustering Algorithms*, *Ann. Data. Sci.* **2**, 165 (2015).
- [276] L. van der Maaten, E. Postma, and J. van den Herik, *Dimensionality Reduction: A Comparative Review*, *J. Mach. Learn. Res.* **10**, 36 (2009).
- [277] J. Masci, U. Meier, D. Cireşan, and J. Schmidhuber, *Stacked Convolutional Auto-Encoders for Hierarchical Feature Extraction*, in *Artif. Neural Netw. Mach. Learn. – ICANN 2011*, edited by T. Honkela, W. Duch, M. Girolami, and S. Kaski (2011), 52.
- [278] Y. Bengio, A. Courville, and P. Vincent, *Representation Learning: A Review and New Perspectives*, *IEEE Trans. Pattern Anal. Mach. Intell.* **35**, 1798 (2013).
- [279] D. P. Kingma and M. Welling, *Auto-Encoding Variational Bayes*, in *International Conference on Learning Representations (ICLR) 2014* (2014).
- [280] I. T. Jolliffe and J. Cadima, *Principal Component Analysis: A Review and Recent Developments*, *Philos. Trans. A Math. Phys. Eng. Sci.* **374**, 20150202 (2016).
- [281] A. Impertro, J. F. Wienand, S. Häfele, H. von Raven, S. Hubele, T. Klostermann, C. R. Cabrera, I. Bloch, and M. Aidelsburger, *An Unsupervised Deep Learning Algorithm for Single-Site Reconstruction in Quantum Gas Microscopes*, *Commun. Phys.* **6**, 1 (2023).
- [282] M. Genzel, J. Macdonald, and M. März, *Solving Inverse Problems With Deep Neural Networks – Robustness Included?* *IEEE Trans. Pattern Anal. Mach. Intell.* **45**, 1119 (2023).
- [283] M. Bianchini and F. Scarselli, *On the Complexity of Neural Network Classifiers: A Comparison Between Shallow and Deep Architectures*, *IEEE Trans. Neural Netw. Learn. Syst.* **25**, 1553 (2014).
- [284] K. Simonyan and A. Zisserman, *Very Deep Convolutional Networks for Large-Scale Image Recognition*, in *3rd Int. Conf. Learn. Represent. ICLR 2015* (2015).
- [285] A. Krizhevsky, I. Sutskever, and G. E. Hinton, *ImageNet Classification with Deep Convolutional Neural Networks*, *Commun. ACM* **60**, 84 (2017).
- [286] Q. Nguyen and M. Hein, *Optimization Landscape and Expressivity of Deep CNNs*, in *Proc. 35th Int. Conf. Mach. Learn.* (July 2018), 3730.
- [287] J. T. Springenberg, A. Dosovitskiy, T. Brox, and M. Riedmiller, *Striving for Simplicity: The All Convolutional Net*, in *3rd Int. Conf. Learn. Represent. ICLR 2015* (2015).
- [288] R. Salakhutdinov and G. Hinton, *Semantic Hashing*, *IJAR* **50**, 969 (2009).
- [289] M. D. Zeiler and R. Fergus, *Visualizing and Understanding Convolutional Networks*, in *Comput. Vis. – ECCV 2014*, edited by D. Fleet, T. Pajdla, B. Schiele, and T. Tuytelaars (2014), 818.
- [290] A. Impertro, S. Huh, S. Karch, J. F. Wienand, I. Bloch, and M. Aidelsburger, *Strongly Interacting Meissner Phases in Large Bosonic Flux Ladders*, *Nat. Phys.*, **1** (2025).
- [291] S. Karch, S. Bandyopadhyay, Z.-H. Sun, A. Impertro, S. Huh, I. P. Rodríguez, J. F. Wienand, W. Ketterle, M. Heyl, A. Polkovnikov, I. Bloch, and M. Aidelsburger, *Probing Quantum Many-Body Dynamics Using Subsystem Loschmidt Echos*, *arXiv:2501.16995*.
- [292] F. Gerbier and J. Dalibard, *Gauge Fields for Ultracold Atoms in Optical Superlattices*, *New J. Phys.* **12**, 033007 (2010).
- [293] M. Łącki, M. A. Baranov, H. Pichler, and P. Zoller, *Nanoscale “Dark State” Optical Potentials for Cold Atoms*, *Phys. Rev. Lett.* **117**, 233001 (2016).

- [294] M. Saffman, T. G. Walker, and K. Mølmer, *Quantum Information with Rydberg Atoms*, *Rev. Mod. Phys.* **82**, 2313 (2010).
- [295] M. Morgado and S. Whitlock, *Quantum Simulation and Computing with Rydberg-interacting Qubits*, *AVS Quant. Sci.* **3**, 023501 (2021).
- [296] C. D. Bruzewicz, J. Chiaverini, R. McConnell, and J. M. Sage, *Trapped-Ion Quantum Computing: Progress and Challenges*, *Appl. Phys. Rev.* **6**, 021314 (2019).
- [297] K. Singh, S. Anand, A. Pocklington, J. T. Kemp, and H. Bernien, *Dual-Element, Two-Dimensional Atom Array with Continuous-Mode Operation*, *Phys. Rev. X* **12**, 011040 (2022).
- [298] F. Gyger, M. Ammenwerth, R. Tao, H. Timme, S. Snigirev, I. Bloch, and J. Zeiher, *Continuous Operation of Large-Scale Atom Arrays in Optical Lattices*, *Phys. Rev. Res.* **6**, 033104 (2024).
- [299] S. Kim, S.-H. Chu, Y.-J. Park, and C.-Y. Lee, *A Tied-Weight Autoencoder for the Linear Dimensionality Reduction of Sample Data*, *Sci. Rep.* **14**, 26801 (2024).
- [300] T. Hartke, B. Oreg, N. Jia, and M. Zwierlein, *Doublon-Hole Correlations and Fluctuation Thermometry in a Fermi-Hubbard Gas*, *Phys. Rev. Lett.* **125**, 113601 (2020).
- [301] M. Cheneau, R. Journet, M. Boffety, F. Goudail, C. Kulcsár, and P. Trouvé-Peloux, *Fast, Accurate, and Predictive Method for Atom Detection in Site-Resolved Images of Microtrap Arrays*, [arXiv:2502.08511](https://arxiv.org/abs/2502.08511).
- [302] M. Rispoli, A. Lukin, R. Schittko, S. Kim, M. E. Tai, J. Léonard, and M. Greiner, *Quantum Critical Behaviour at the Many-Body Localization Transition*, *Nature* **573**, 385 (2019).
- [303] M. Cheneau, P. Barmettler, D. Poletti, M. Endres, P. Schauß, T. Fukuhara, C. Gross, I. Bloch, C. Kollath, and S. Kuhr, *Light-Cone-like Spreading of Correlations in a Quantum Many-Body System*, *Nature* **481**, 484 (2012).
- [304] T. A. Hilker, G. Salomon, F. Grusdt, A. Omran, M. Boll, E. Demler, I. Bloch, and C. Gross, *Revealing Hidden Antiferromagnetic Correlations in Doped Hubbard Chains via String Correlators*, *Science* **357**, 484 (2017).
- [305] G. Salomon, J. Koepsell, J. Vijayan, T. A. Hilker, J. Nespolo, L. Pollet, I. Bloch, and C. Gross, *Direct Observation of Incommensurate Magnetism in Hubbard Chains*, *Nature* **565**, 56 (2019).
- [306] J. Koepsell, J. Vijayan, P. Sompet, F. Grusdt, T. A. Hilker, E. Demler, G. Salomon, I. Bloch, and C. Gross, *Imaging Magnetic Polarons in the Doped Fermi-Hubbard Model*, *Nature* **572**, 358 (2019).
- [307] P. A. Lee, N. Nagaosa, and X.-G. Wen, *Doping a Mott Insulator: Physics of High-Temperature Superconductivity*, *Rev. Mod. Phys.* **78**, 17 (2006).
- [308] B. Swingle, G. Bentsen, M. Schleier-Smith, and P. Hayden, *Measuring the Scrambling of Quantum Information*, *Phys. Rev. A* **94**, 040302 (2016).
- [309] A. Bohrdt, C. B. Mendl, M. Endres, and M. Knap, *Scrambling and Thermalization in a Diffusive Quantum Many-Body System*, *New J. Phys.* **19**, 063001 (2017).
- [310] J. De Nardis, B. Doyon, M. Medenjak, and M. Panfil, *Correlation Functions and Transport Coefficients in Generalised Hydrodynamics*, *J. Stat. Mech.* **2022**, 014002 (2022).
- [311] N. Wiebe, C. Granade, C. Ferrie, and D. G. Cory, *Hamiltonian Learning and Certification Using Quantum Resources*, *Phys. Rev. Lett.* **112**, 190501 (2014).
- [312] M. Holzäpfel, T. Baumgratz, M. Cramer, and M. B. Plenio, *Scalable Reconstruction of Unitary Processes and Hamiltonians*, *Phys. Rev. A* **91**, 042129 (2015).

-
- [313] J. Wang, S. Paesani, R. Santagati, S. Knauer, A. A. Gentile, N. Wiebe, M. Petruzzella, J. L. O'Brien, J. G. Rarity, A. Laing, and M. G. Thompson, *Experimental Quantum Hamiltonian Learning*, *Nat. Phys.* **13**, 551 (2017).
- [314] J. Carrasco, A. Elben, C. Kokail, B. Kraus, and P. Zoller, *Theoretical and Experimental Perspectives of Quantum Verification*, *PRX Quantum* **2**, 010102 (2021).
- [315] W. Yu, J. Sun, Z. Han, and X. Yuan, *Robust and Efficient Hamiltonian Learning*, *Quantum* **7**, 1045 (2023).
- [316] C. Schweizer, M. Lohse, R. Citro, and I. Bloch, *Spin Pumping and Measurement of Spin Currents in Optical Superlattices*, *Phys. Rev. Lett.* **117**, 170405 (2016).
- [317] J. Sebby-Strabley, B. L. Brown, M. Anderlini, P. J. Lee, W. D. Phillips, J. V. Porto, and P. R. Johnson, *Preparing and Probing Atomic Number States with an Atom Interferometer*, *Phys. Rev. Lett.* **98**, 200405 (2007).
- [318] G.-X. Su, H. Sun, A. Hudomal, J.-Y. Desaulles, Z.-Y. Zhou, B. Yang, J. C. Halimeh, Z.-S. Yuan, Z. Papić, and J.-W. Pan, *Observation of Many-Body Scarring in a Bose-Hubbard Quantum Simulator*, *Phys. Rev. Res.* **5**, 023010 (2023).
- [319] J. C. Brüggengjürgen, M. S. Fischer, and C. Weitenberg, *A Phase Microscope for Quantum Gases*, [arXiv:2410.10611](https://arxiv.org/abs/2410.10611).
- [320] M. Anderlini, P. J. Lee, B. L. Brown, J. Sebby-Strabley, W. D. Phillips, and J. V. Porto, *Controlled Exchange Interaction between Pairs of Neutral Atoms in an Optical Lattice*, *Nature* **448**, 452 (2007).
- [321] S. Trotzky, Y.-A. Chen, U. Schnorrberger, P. Cheinet, and I. Bloch, *Controlling and Detecting Spin Correlations of Ultracold Atoms in Optical Lattices*, *Phys. Rev. Lett.* **105**, 265303 (2010).
- [322] B. Yang, H. Sun, C.-J. Huang, H.-Y. Wang, Y. Deng, H.-N. Dai, Z.-S. Yuan, and J.-W. Pan, *Cooling and Entangling Ultracold Atoms in Optical Lattices*, *Science* **369**, 550 (2020).
- [323] W.-Y. Zhang et al., *Scalable Multipartite Entanglement Created by Spin Exchange in an Optical Lattice*, *Phys. Rev. Lett.* **131**, 073401 (2023).
- [324] H. Y. Carr and E. M. Purcell, *Effects of Diffusion on Free Precession in Nuclear Magnetic Resonance Experiments*, *Phys. Rev.* **94**, 630 (1954).
- [325] S. Meiboom and D. Gill, *Modified Spin-Echo Method for Measuring Nuclear Relaxation Times*, *Rev. Sci. Instrum.* **29**, 688 (1958).
- [326] L. Viola, E. Knill, and S. Lloyd, *Dynamical Decoupling of Open Quantum Systems*, *Phys. Rev. Lett.* **82**, 2417 (1999).
- [327] A. L. Shaw, R. Finkelstein, R. B.-S. Tsai, P. Scholl, T. H. Yoon, J. Choi, and M. Endres, *Multi-Ensemble Metrology by Programming Local Rotations with Atom Movements*, *Nat. Phys.* **20**, 195 (2024).
- [328] H.-M. Chun, Q. Gao, and J. M. Horowitz, *Nonequilibrium Green-Kubo Relations for Hydrodynamic Transport from an Equilibrium-like Fluctuation-Response Equality*, *Phys. Rev. Res.* **3**, 043172 (2021).
- [329] P. Pegolo, E. Drigo, F. Grasselli, and S. Baroni, *Transport Coefficients from Equilibrium Molecular Dynamics*, *J. Chem. Phys.* **162**, 064111 (2025).
- [330] S. Schweers, G. M. Schütz, and P. Maass, *Correlations of Density and Current Fluctuations in Single-File Motion of Hard Spheres and in Driven Lattice Gas with Nearest-Neighbor Interaction*, [arXiv:2502.15469](https://arxiv.org/abs/2502.15469).

- [331] K. Vogel and H. Risken, *Determination of Quasiprobability Distributions in Terms of Probability Distributions for the Rotated Quadrature Phase*, *Phys. Rev. A* **40**, 2847 (1989).
- [332] D. Gross, Y.-K. Liu, S. T. Flammia, S. Becker, and J. Eisert, *Quantum State Tomography via Compressed Sensing*, *Phys. Rev. Lett.* **105**, 150401 (2010).
- [333] M. Cramer, M. B. Plenio, S. T. Flammia, R. Somma, D. Gross, S. D. Bartlett, O. Landon-Cardinal, D. Poulin, and Y.-K. Liu, *Efficient Quantum State Tomography*, *Nat. Commun.* **1**, 149 (2010).
- [334] J.-H. Zheng, B. Irsigler, L. Jiang, C. Weitenberg, and W. Hofstetter, *Measuring an Interaction-Induced Topological Phase Transition via the Single-Particle Density Matrix*, *Phys. Rev. A* **101**, 013631 (2020).
- [335] S. N. Hearth, M. O. Flynn, A. Chandran, and C. R. Laumann, *Efficient Local Classical Shadow Tomography with Number Conservation*, *Phys. Rev. Lett.* **133**, 060802 (2024).
- [336] A. Prakash and B. H. Madhusudhana, *A Protocol to Characterize Errors in Quantum Simulation of Many-Body Physics*, *Phys. Rev. Res.* **6**, 043127 (2024).
- [337] P. Roushan et al., *Spectroscopic Signatures of Localization with Interacting Photons in Superconducting Qubits*, *Science* **358**, 1175 (2017).
- [338] Z.-C. Xiang et al., *Simulating Chern Insulators on a Superconducting Quantum Processor*, *Nat. Commun.* **14**, 5433 (2023).
- [339] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'Brien, *A Variational Eigenvalue Solver on a Photonic Quantum Processor*, *Nat. Commun.* **5**, 4213 (2014).
- [340] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, *The Theory of Variational Hybrid Quantum-Classical Algorithms*, *New J. Phys.* **18**, 023023 (2016).
- [341] R. de Keijzer, O. Tse, and S. Kokkelmans, *Pulse Based Variational Quantum Optimal Control for Hybrid Quantum Computing*, *Quantum* **7**, 908 (2023).
- [342] A. Michel, S. Grijalva, L. Henriët, C. Domain, and A. Browaeys, *Blueprint for a Digital-Analog Variational Quantum Eigensolver Using Rydberg Atom Arrays*, *Phys. Rev. A* **107**, 042602 (2023).
- [343] D. Greif, T. Uehlinger, G. Jotzu, L. Tarruell, and T. Esslinger, *Short-Range Quantum Magnetism of Ultracold Fermions in an Optical Lattice*, *Science* **340**, 1307 (2013).
- [344] P. Sala, T. Rakovszky, R. Verresen, M. Knap, and F. Pollmann, *Ergodicity Breaking Arising from Hilbert Space Fragmentation in Dipole-Conserving Hamiltonians*, *Phys. Rev. X* **10**, 011047 (2020).
- [345] V. Khemani, M. Hermele, and R. Nandkishore, *Localization from Hilbert Space Shattering: From Theory to Physical Realizations*, *Phys. Rev. B* **101**, 174204 (2020).
- [346] S. Moudgalya, A. Prem, R. Nandkishore, N. Regnault, and B. A. Bernevig, *Thermalization and Its Absence within Krylov Subspaces of a Constrained Hamiltonian*, in *Memorial Volume for Shoucheng Zhang* (Nov. 2020), 147.
- [347] E. Zohar, J. I. Cirac, and B. Reznik, *Quantum Simulations of Lattice Gauge Theories Using Ultracold Atoms in Optical Lattices*, *Rep. Prog. Phys.* **79**, 014401 (2015).
- [348] M. Dalmonte and S. Montangero, *Lattice Gauge Theory Simulations in the Quantum Information Era*, *Contemp. Phys.* **57**, 388 (2016).
- [349] M. C. Bañuls et al., *Simulating Lattice Gauge Theories within Quantum Technologies*, *Eur. Phys. J. D* **74**, 165 (2020).

-
- [350] M. Aidelsburger et al., *Cold Atoms Meet Lattice Gauge Theory*, *Philos. Trans. R. Soc. A Math. Phys. Eng. Sci.* **380**, 20210064 (2021).
- [351] J. C. Halimeh, M. Aidelsburger, F. Grusdt, P. Hauke, and B. Yang, *Cold-Atom Quantum Simulators of Gauge Theories*, *Nat. Phys.* **21**, 25 (2025).
- [352] S. Wimperis, *Broadband, Narrowband, and Passband Composite Pulses for Use in Advanced NMR Experiments*, *J. Magn. Reson.* **109**, 221 (1994).
- [353] H. K. Cummins, G. Llewellyn, and J. A. Jones, *Tackling Systematic Errors in Quantum Logic Gates with Composite Rotations*, *Phys. Rev. A* **67**, 042308 (2003).
- [354] H. L. Gevorgyan and N. V. Vitanov, *Ultra-high-Fidelity Composite Rotational Quantum Gates*, *Phys. Rev. A* **104**, 012609 (2021).
- [355] A. A. Stepanenko and M. D. Liberto, *Vortex Bound States in Dimerized π -Flux Optical Lattices: Characterization, State Preparation and Current Measurement*, [arXiv:2410.06184](https://arxiv.org/abs/2410.06184).
- [356] M. Aidelsburger, S. Nascimbene, and N. Goldman, *Artificial gauge fields in materials and engineered systems*, *C. R. Phys.* **19**, 394 (2018).
- [357] M. Di Dio, S. De Palo, E. Orignac, R. Citro, and M.-L. Chiofalo, *Persisting Meissner State and Incommensurate Phases of Hard-Core Boson Ladders in a Flux*, *Phys. Rev. B* **92**, 060506 (2015).
- [358] S. Greschner, M. Piraud, F. Heidrich-Meisner, I. P. McCulloch, U. Schollwöck, and T. Vekua, *Spontaneous Increase of Magnetic Flux and Chiral-Current Reversal in Bosonic Ladders: Swimming against the Tide*, *Phys. Rev. Lett.* **115**, 190402 (2015).
- [359] S. Greschner, M. Piraud, F. Heidrich-Meisner, I. P. McCulloch, U. Schollwöck, and T. Vekua, *Symmetry-Broken States in a System of Interacting Bosons on a Two-Leg Ladder with a Uniform Abelian Gauge Field*, *Phys. Rev. A* **94**, 063628 (2016).
- [360] S. Greschner and T. Vekua, *Vortex-Hole Duality: A Unified Picture of Weak- and Strong-Coupling Regimes of Bosonic Ladders with Flux*, *Phys. Rev. Lett.* **119**, 073401 (2017).
- [361] M. Buser, F. Heidrich-Meisner, and U. Schollwöck, *Finite-Temperature Properties of Interacting Bosons on a Two-Leg Flux Ladder*, *Phys. Rev. A* **99**, 053601 (2019).
- [362] R. Citro, S. De Palo, N. Victorin, A. Minguzzi, and E. Orignac, *Spectral Function of a Boson Ladder in an Artificial Gauge Field*, *Cond. Mat.* **5**, 15 (2020).
- [363] M. Buser, C. Hubig, U. Schollwöck, L. Tarruell, and F. Heidrich-Meisner, *Interacting Bosonic Flux Ladders with a Synthetic Dimension: Ground-state Phases and Quantum Quench Dynamics*, *Phys. Rev. A* **102**, 053314 (2020).
- [364] M. Buser, S. Greschner, U. Schollwöck, and T. Giamarchi, *Probing the Hall Voltage in Synthetic Quantum Systems*, *Phys. Rev. Lett.* **126**, 030501 (2021).
- [365] M. Buser, U. Schollwöck, and F. Grusdt, *Snapshot-Based Characterization of Particle Currents and the Hall Response in Synthetic Flux Lattices*, *Phys. Rev. A* **105**, 033303 (2022).
- [366] R. Citro, T. Giamarchi, and E. Orignac, *Hall Response in Interacting Bosonic and Fermionic Ladders*, [arXiv:2404.16973](https://arxiv.org/abs/2404.16973).
- [367] E. Orignac and T. Giamarchi, *Meissner Effect in a Bosonic Ladder*, *Phys. Rev. B* **64**, 144515 (2001).
- [368] F. Crépin, N. Laflorencie, G. Roux, and P. Simon, *Phase Diagram of Hard-Core Bosons on Clean and Disordered Two-Leg Ladders: Mott Insulator–Luttinger Liquid–Bose Glass*, *Phys. Rev. B* **84**, 054517 (2011).

- [369] S.-L. Zhu, Z.-D. Wang, Y.-H. Chan, and L.-M. Duan, *Topological Bose-Mott Insulators in a One-Dimensional Optical Superlattice*, *Phys. Rev. Lett.* **110**, 075303 (2013).
- [370] R. Wei and E. J. Mueller, *Theory of Bosons in Two-Leg Ladders with Large Magnetic Fields*, *Phys. Rev. A* **89**, 063617 (2014).
- [371] D. Hügél and B. Paredes, *Chiral Ladders and the Edges of Quantum Hall Insulators*, *Phys. Rev. A* **89**, 023619 (2014).
- [372] E. Cornfeld and E. Sela, *Chiral Currents in One-Dimensional Fractional Quantum Hall States*, *Phys. Rev. B* **92**, 115446 (2015).
- [373] M. Calvanese Strinati, E. Cornfeld, D. Rossini, S. Barbarino, M. Dalmonte, R. Fazio, E. Sela, and L. Mazza, *Laughlin-like States in Bosonic and Fermionic Atomic Synthetic Ladders*, *Phys. Rev. X* **7**, 021033 (2017).
- [374] A. Petrescu, M. Piraud, G. Roux, I. P. McCulloch, and K. Le Hur, *Precursor of the Laughlin State of Hard-Core Bosons on a Two-Leg Ladder*, *Phys. Rev. B* **96**, 014524 (2017).
- [375] M. Calvanese Strinati, S. Sahoo, K. Shtengel, and E. Sela, *Pretopological Fractional Excitations in the Two-Leg Flux Ladder*, *Phys. Rev. B* **99**, 245101 (2019).
- [376] C. L. Kane, R. Mukhopadhyay, and T. C. Lubensky, *Fractional Quantum Hall Effect in an Array of Quantum Wires*, *Phys. Rev. Lett.* **88**, 036401 (2002).
- [377] J. C. Y. Teo and C. L. Kane, *From Luttinger Liquid to Non-Abelian Quantum Hall States*, *Phys. Rev. B* **89**, 085101 (2014).
- [378] R. Tao and D. J. Thouless, *Fractional Quantization of Hall Conductance*, *Phys. Rev. B* **28**, 1142 (1983).
- [379] J. Flavin and A. Seidel, *Abelian and Non-Abelian Statistics in the Coherent State Representation*, *Phys. Rev. X* **1**, 021015 (2011).
- [380] M. E. Tai, A. Lukin, M. Rispoli, R. Schittko, T. Menke, Dan Borgnia, P. M. Preiss, F. Grusdt, A. M. Kaufman, and M. Greiner, *Microscopy of the Interacting Harper–Hofstadter Model in the Two-Body Limit*, *Nature* **546**, 519 (2017).
- [381] T. Chen, C. Huang, B. Gadway, and J. P. Covey, *Quantum Walks and Correlated Dynamics in an Interacting Synthetic Rydberg Lattice*, *Phys. Rev. Lett.* **133**, 120604 (2024).
- [382] Q. Liang, Z. Dong, J.-S. Pan, H. Wang, H. Li, Z. Yang, W. Yi, and B. Yan, *Chiral Dynamics of Ultracold Atoms under a Tunable $SU(2)$ Synthetic Gauge Field*, *Nat. Phys.* **20**, 1738 (2024).
- [383] T. Chen, C. Huang, I. Velkovsky, T. Ozawa, H. Price, J. P. Covey, and B. Gadway, *Interaction-Driven Breakdown of Aharonov–Bohm Caging in Flat-Band Rydberg Lattices*, *Nat. Phys.* **21**, 221 (2025).
- [384] T.-W. Zhou, T. Beller, G. Masini, J. Parravicini, G. Cappellini, C. Repellin, T. Giamarchi, J. Catani, M. Filippone, and L. Fallani, *Measuring Hall Voltage and Hall Resistance in an Atom-Based Quantum Simulator*, arXiv:2411.09744.
- [385] R. Wei and E. J. Mueller, *Magnetic-Field Dependence of Raman Coupling in Alkali-Metal Atoms*, *Phys. Rev. A* **87**, 042514 (2013).
- [386] M. D. Frye, B. C. Yang, and J. M. Hutson, *Ultracold Collisions of Cs Atoms in Excited Zeeman and Hyperfine States*, *Phys. Rev. A* **100**, 022702 (2019).
- [387] N. Baldelli, C. R. Cabrera, S. Julià-Farré, M. Aidelsburger, and L. Barbiero, *Frustrated Extended Bose-Hubbard Model and Deconfined Quantum Critical Points with Optical Lattices at the Antimagic Wavelength*, *Phys. Rev. Lett.* **132**, 153401 (2024).

-
- [388] Y. Zheng, S. Feng, and S.-J. Yang, *Chiral Bloch Oscillation and Nontrivial Topology in a Ladder Lattice with Magnetic Flux*, *Phys. Rev. A* **96**, 063613 (2017).
- [389] A.-X. Zhang, W. Zhang, Y.-H. Qin, X.-W. Hu, X. Qiao, and J.-K. Xue, *Nonlinear Bloch Dynamics and Spin-Wave Generation in a Bose-Hubbard Ladder Subject to Effective Magnetic Field*, *Phys. Rev. E* **106**, 044215 (2022).
- [390] X. Qiao, X.-B. Zhang, Y. Jian, A.-X. Zhang, and J.-K. Xue, *Quench Dynamics of Two-Leg Ladders with Magnetic Flux*, *Physica A: Stat. Mech. Appl.* **576**, 126062 (2021).
- [391] Y. Jian, A.-X. Zhang, X. Qiao, J.-C. Liang, Z.-F. Yu, and J.-K. Xue, *Defect Induced Nonequilibrium Quantum Dynamics in an Interacting Bose-Hubbard Flux Ladder*, *New J. Phys.* **25**, 043025 (2023).
- [392] M. K. Giri, B. Paul, and T. Mishra, *Flux-Induced Reentrant Dynamics in the Quantum Walk of Interacting Bosons*, *Phys. Rev. A* **108**, 063319 (2023).
- [393] W. Huang and Y. Yao, *Spatial Inversion Symmetry Breaking of Vortex Current in a Biased-Ladder Superfluid*, *Phys. Rev. Res.* **6**, 013037 (2024).
- [394] M. K. Giri, B. Paul, and T. Mishra, *Flux-Enhanced Localization and Reentrant Delocalization in the Quench Dynamics of Two Interacting Bosons on a Bose-Hubbard Ladder*, *Phys. Rev. A* **109**, 043308 (2024).
- [395] L.-L. Mi, J. Wang, J.-P. Yao, A.-X. Zhang, and J.-K. Xue, *Magnetic Induced Reflection, Resistance and Symmetry Breaking of Atomic Matter Wave in a Two-Leg Ladder*, *Phys. Lett. A* **541**, 130424 (2025).
- [396] K. Viebahn, J. Minguzzi, K. Sandholzer, A.-S. Walter, M. Sajnani, F. Görg, and T. Esslinger, *Suppressing Dissipation in a Floquet-Hubbard System*, *Phys. Rev. X* **11**, 011057 (2021).
- [397] Y. Wang, A.-S. Walter, G. Jotzu, and K. Viebahn, *Topological Floquet Engineering Using Two Frequencies in Two Dimensions*, *Phys. Rev. A* **107**, 043309 (2023).
- [398] Y. Chen, Z. Zhu, and K. Viebahn, *Mitigating Higher-Band Heating in Floquet-Hubbard Lattices via Two-Tone Driving*, [arXiv:2410.12308](https://arxiv.org/abs/2410.12308).
- [399] B. Wang, X.-Y. Dong, F. N. Ünal, and A. Eckardt, *Robust and Ultrafast State Preparation by Ramping Artificial Gauge Potentials*, *New J. Phys.* **23**, 063017 (2021).
- [400] P. M. Schindler and M. Bukov, *Counterdiabatic Driving for Periodically Driven Systems*, *Phys. Rev. Lett.* **133**, 123402 (2024).
- [401] B. Michen, C. Repellin, and J. C. Budich, *Adiabatic Preparation of Fractional Chern Insulators from an Effective Thin-Torus Limit*, *Phys. Rev. Res.* **5**, 023100 (2023).
- [402] F. A. Palm, J. Kwan, B. Bakkali-Hassani, M. Greiner, U. Schollwöck, N. Goldman, and F. Grusdt, *Growing Extended Laughlin States in a Quantum Gas Microscope: A Patchwork Construction*, *Phys. Rev. Res.* **6**, 013198 (2024).
- [403] B. Wang, M. Aidelsburger, J. Dalibard, A. Eckardt, and N. Goldman, *Cold-Atom Elevator: From Edge-State Injection to the Preparation of Fractional Chern Insulators*, *Phys. Rev. Lett.* **132**, 163402 (2024).
- [404] S. Diehl, E. Rico, M. A. Baranov, and P. Zoller, *Topology by Dissipation in Atomic Quantum Wires*, *Nat. Phys.* **7**, 971 (2011).
- [405] J. C. Budich, P. Zoller, and S. Diehl, *Dissipative Preparation of Chern Insulators*, *Phys. Rev. A* **91**, 042117 (2015).
- [406] Z. Liu, E. J. Bergholtz, and J. C. Budich, *Dissipative Preparation of Fractional Chern Insulators*, *Phys. Rev. Res.* **3**, 043119 (2021).

- [407] F. A. Palm, S. Mardazad, A. Bohrdt, U. Schollwöck, and F. Grusdt, *Snapshot-Based Detection of $\nu = \frac{1}{2}$ Laughlin States: Coupled Chains and Central Charge*, *Phys. Rev. B* **106**, L081108 (2022).
- [408] F. A. Palm, C. Repellin, N. Goldman, and F. Grusdt, *Absence of Gapless Majorana Edge Modes in Few-Leg Bosonic Flux Ladders*, *Phys. Rev. Res.* **7**, L012001 (2025).
- [409] A. Griessner, A. J. Daley, S. R. Clark, D. Jaksch, and P. Zoller, *Dark-State Cooling of Atoms by Superfluid Immersion*, *Phys. Rev. Lett.* **97**, 220403 (2006).
- [410] A. Kantian, S. Langer, and A. J. Daley, *Dynamical Disentangling and Cooling of Atoms in Bilayer Optical Lattices*, *Phys. Rev. Lett.* **120**, 060401 (2018).
- [411] E. Wybo, M. Knap, and A. Bastianello, *Quantum Sine-Gordon Dynamics in Coupled Spin Chains*, *Phys. Rev. B* **106**, 075102 (2022).
- [412] E. Wybo, A. Bastianello, M. Aidelsburger, I. Bloch, and M. Knap, *Preparing and Analyzing Solitons in the Sine-Gordon Model with Quantum Gas Microscopes*, *PRX Quantum* **4**, 030308 (2023).
- [413] L. D'Alessio, Yariv Kafri, Anatoli Polkovnikov, and M. and Rigol, *From Quantum Chaos and Eigenstate Thermalization to Statistical Mechanics and Thermodynamics*, *Adv. Phys.* **65**, 239 (2016).
- [414] S. Moudgalya, B. A. Bernevig, and N. Regnault, *Quantum Many-Body Scars and Hilbert Space Fragmentation: A Review of Exact Results*, *Rep. Prog. Phys.* **85**, 086501 (2022).
- [415] M. Aidelsburger et al., *Cold Atoms Meet Lattice Gauge Theory*, *Phil. Trans. R. Soc. A.* **380**, 20210064 (2022).
- [416] A. D. Meglio et al., *Quantum Computing for High-Energy Physics: State of the Art and Challenges. Summary of the QC4HEP Working Group*, *PRX Quantum* **5**, 037001 (2024).
- [417] P. Weinberg and M. Bukov, *QuSpin: A Python Package for Dynamics and Exact Diagonalisation of Quantum Many Body Systems. Part II: Bosons, Fermions and Higher Spins*, *SciPost Phys.* **7**, 020 (2019).
- [418] J. Hauschild and F. Pollmann, *Efficient Numerical Simulations with Tensor Networks: Tensor Network Python (TeNPy)*, *SciPost Phys. Lect. Notes*, **5** (2018).
- [419] U. Schollwöck, *The Density-Matrix Renormalization Group*, *Rev. Mod. Phys.* **77**, 259 (2005).
- [420] B. M. Spar, E. Guardado-Sanchez, S. Chi, Z. Z. Yan, and W. S. Bakr, *Realization of a Fermi-Hubbard Optical Tweezer Array*, *Phys. Rev. Lett.* **128**, 223202 (2022).

Acknowledgements

During the course of my PhD, I have been fortunate to work with many talented and inspiring people. Pursuing such an endeavor is always a team effort, and I would like to take this opportunity to express my gratitude to all those who have supported me along the way and allowed this thesis to come to fruition.

First and foremost, I would like to thank Monika Aidelsburger and Immanuel Bloch for giving me the opportunity to pursue a PhD in the cesium lab, and for their great supervision throughout the years. I am inspired by their fascination for physics, the depth of their knowledge and their sense for what are the important topics. To Immanuel, I am especially grateful for the amazing group that he has built, which continuously attracts the most talented and motivated people, and where the resources necessary to pursue ambitious projects are always available. To Monika, I am extremely grateful for the immeasurable amount of time she has dedicated to the lab, which oftentimes was the key to fast progress, and for her constant stream of great ideas and in-depth knowledge of the relevant aspects. It has been a true privilege to work with both of you, in an environment where we were always free to pursue our own ideas and actively shape the lab's projects.

I would also like to thank Christof Weitenberg, Ulrich Schollwöck and Alexander Högele for taking part in my thesis committee. Special thanks go to Christof for agreeing to be the second reviewer for this thesis.

Importantly, I would like to express my gratitude to all members of the cesium lab, past and present. Every single person has contributed their unique skills and ideas, which in sum have enabled many successful projects, and I have learned a great deal from each of you. Julian Wienand, with whom I have spent most of my time in the lab, and who has become a good friend. Thank you for being always so highly motivated (except when it comes to objectives), for an inspiring week in Leysin, and for reminding me of the importance to frame things in a 'sexy' way. Till Klostermann and Hendrik von Raven for setting up an ambitious yet robust apparatus, and for great company during the early stages. Cesar Cabrera for his dedication and motivation to advance the lab, and for being such a cool and social spirit. Christian Schweizer for his amazing coding and electronics skills, knowledge of all things 'current', and for creating many way too funny moments, thanks to which I now know where the rabbit is buried. Sophie Häfele, Scott Hubele and Ignacio Pérez for being incredibly motivated master students and for enriching the lab atmosphere with their funny and interesting personalities. Simon Karch for many interesting and funny discussions, an inspiring passion for fundamental physics, and his attention to detail. SeungJung Huh for his constant support and great ideas on all our projects, and his kind and patient spirit. And lastly, our newest member, Irene Prieto Rodríguez, for bringing in fresh energy and her insane motivation to learn everything about the lab and conduct cool experiments.

Beyond the cesium lab, I would like to thank all members of the Bloch and Aidelsburger groups at LMU and MPQ for creating an inspiring, friendly and collaborative atmosphere. I am grateful to Bodo Hecker and Reinhard Grottenthaler for their support and interesting discussions on electronics, to Thomas Großhauser, Jürgen Aust, Anton Mayer and their teams of the LMU and MPQ machine shops for their amazing work in manufacturing mechanical parts, and

to Ildiko Kecskesi, Kristina Schuldt and Annika Herlitzek for their administrative support. Special thanks go to Ildiko for her constant support, advice and interesting kitchen conversations.

I would also like to extend my thanks to various theory collaborators who brought in helpful advice, among them notably Botao Wang, Felix Palm, Paul Schindler and Yudong Sun.

Several of my colleagues and conference acquaintances have become good friends, and I hope to keep them in my life for the time to come. Besides the already mentioned people, I would particularly like to highlight: Daniel Adler, who has been a central part of my physics experience for many years, and provided an ear and helpful advice in various moments. I really hope we can keep up our traditions in the future! Nelson Darkwah Oppong for being an inspiring, incredibly knowledgeable mentor and friend, and for his useful advice and many funny discussions. Celine Castor, who managed to instantly fascinate me with her most amazing personality. Additionally, I would like to thank all my friends outside of physics for their support and unforgettable moments, both in the mountains and the city.

Lastly, but most importantly, I would like to thank my family, specifically my wonderful partner Nina, for their unconditional support and love, which was invaluable particularly during the difficult moments. I am so grateful to have you all in my life!

Thank you!