Higgs Mode and Critical Temperature in the BCS-BEC Crossover

Dissertation
zur
Erlangung des Doktorgrades (Dr. rer. nat.)
der
Mathematisch-Naturwissenschaftlichen Fakultät
der
Rheinischen Friedrich-Wilhelms-Universität Bonn

vorgelegt von
Moritz Breyer
aus
Bonn

Bonn, 2023

Angefertigt mit Genehmigung der Mathematisch-Naturwissenschaftlichen Fakultät der Rheinischen Friedrich-Wilhelms-Universität Bonn

1. Gutachter:	Prof. Dr. Michael Köhl
2. Gutachter:	Prof. Dr. Stefan Linden
Tag der Promotion:	20.12.2023
Erscheinungsjahr:	2024

Abstract

For this thesis, analogue quantum simulations of the BCS-BEC crossover have been performed. To this end, an ultracold gas of fermionic ^{6}Li atoms was prepared in an optical dipole trap. The gas is in a balanced mixture of two hyperfine states which constitute the two spin states necessary for the implementation of the crossover. The strength of the attractive interaction between these states is controlled by an external magnetic field with the help of a Feshbach resonance.

New techniques to detect superfluidity, based on machine learning, were developed. These give a clear indication of the phase transition for arbitrary interaction strengths. By reconstructing the density distribution of the inhomogeneous gas, a phase diagram of the homogeneous BCS-BEC crossover could be established. It confirmed, for the first time, an increase of critical temperature from the unitary point towards the BEC limit.

With a newly developed coil, the interaction strength of the gas could be quenched to excite the Higgs mode of the superfluid. It is observed as a fast, interaction strength-dependent oscillation in the condensed fraction of the gas and constitutes the first observation of the Higgs mode in time domain in the whole BCS-BEC crossover. The oscillations exhibit a strong damping, which a model based on the local-density approximation attributes to dephasing caused by the inhomogeneity of the gas. Towards the BEC limit of the crossover, a stronger damping than expected from dephasing is observed indicating an additional instability of the Higgs mode. This has been predicted theoretically in this regime, where the system cedes to be particle-hole symmetric.

As a complementary method to the interaction quench, some first results of experiments in which the interaction strength is periodically modulated are presented. Here, a resonance in the response of the superfluid, matching some predictions for the Higgs mode, is observed.

Contents

CHAPTER 1

Introduction

Superconductivity is a phenomenon in which the resistance of a conductor suddenly vanishes if it is cooled below a critical temperature and the superconductor constitutes a distinct phase of matter. It was first observed over a hundred years ago in mercury which has a critical temperature of 4.2 K [1]. Superconductors have a wide range of applications. They are crucial to the construction of the strongest electromagnets employed in medical applications like nuclear magnetic resonance imaging and in large particle colliders. Superconductors are also used as quantum sensors for magnetic fields [2]. Superconducting qubits have already been constructed and are promising candidates for the realisation of quantum computers [3]. A big obstacle for applications of superconductivity are the prohibitively low temperatures needed to reach this phase. Even the superconductors with the highest known critical temperature at ambient pressure, the cuprates, do not surpass critical temperatures of 135 K [4] and therefore require cooling with liquid nitrogen. The reason behind the low critical temperatures lies in the quantum-mechanical nature of the superconducting phase which makes it very sensitive to perturbations like thermal excitations. Currently, our understanding of superconductivity is very limited [5] and advances could facilitate the development of superconductors with higher critical temperatures, ideally the highly sought after room-temperature superconductor. Immediately prior to submission of this thesis, a preprint reporting on the discovery of room-temperature superconductivity at ambient pressure in modified lead-apatite was published [6, 7]. What follows from this claim remains to be seen but, if verified, it could potentially be one of the most important technological breakthroughs in decades. Nevertheless, a better understanding of superconductivity will certainly be valuable in any case. Of especial interest is the poorly understood relation between the macroscopic phenomenon of superconductivity and the microscopic quantum mechanics governing the behaviour of electrons in a material.

Superfluidity of uncharged particles, an analogue to the superconductivity of electrons, was first observed clearly in ^{4}He [8]. In this case the origin of the macroscopic phenomenon in quantum mechanics is relatively obvious: ^{4}He atoms are bosons and superfluidity can be connected to the macroscopic wave function which forms upon Bose-Einstein condensation. As a rough analogy, this holds even though ^{4}He is significantly more complicated than a simple Bose-Einstein condensate (BEC) [9]. The electrons forming a superconductor are, however, fermions and cannot Bose-Einstein condense. Another, more complex mechanism must therefore cause superconductivity. It took roughly half a century until Bardeen, Cooper and Schrieffer (BCS) were able to solve this problem by a theory that includes electrons with attractive interactions [10]. The usual Coulomb repulsion between

electrons is modified in solid state materials due to interactions with phonons which may lead to attractively interacting electrons. At low enough temperatures, the attraction leads to the formation of a macroscopic wave function enabled by the pairing of electrons. The formation of the macroscopic wave function is the key factor for superconductivity and, in analogy to the BEC, is also called condensation.

The BCS theory describes what is known today as conventional superconductivity. The superconductors with the highest critical temperatures, called high T_c-superconductors, are however unconventional superconductors, which means that they have properties diverging from the predictions of BCS theory [11]. In fact, the origin of their superconductivity from a microscopic theory is still not understood [5]. The most important contributions likely come from special geometries, e.g. 2-D electron layers in the cuprates, pairing symmetries, d-wave pairing in the cuprates instead of the BCS s-wave pairing [12], and strong correlations of the electrons due to their interactions. Recently, evidence has been found that the reduced dimensionality is of lesser importance compared to interaction effects [13] making the latter an interesting field of study.

Conventional superconductors are characterised by a very weak attractive interaction. In this limit, the electron pairs are very fragile which results in very low critical temperatures. This regime can also be characterised by pair sizes much larger than the interparticle spacing. If the attraction strength is increased, the result is a stronger binding which leads to smaller pair sizes and higher critical temperatures. At some point the pairs will be much smaller than the interparticle spacing and form, for all practical purposes, composite bosons. Consequently, an increasing attraction strength can change the condensation mechanism for superconductivity or superfluidity from a BCS type to a BEC type. This process is not accompanied by a (quantum) phase transition but rather happens as a smooth transition known as the *BCS-BEC crossover* [14, 15]. Such a system is interesting in its own regard since it unifies the microscopically very different phenomena of BEC and BCS theory to a generalised pairing. As an application, the crossover's relevance is not limited to unconventional superconductivity [16]. The fairly general nature of the underlying model, fermions with attraction, makes it relevant for other fields like nuclear systems [17], neutron stars [18] and the hydrodynamics of quark-gluon plasmas [19] as well.

Despite continuous theoretical effort [17, 20], the crossover is not fully understood today, especially in the strongly interacting regime where pair size and interparticle distance are comparable. The main reasons for a limited understanding are the origin in many-body quantum physics whose large Hilbert spaces prohibit direct solutions on classical computers, strong interactions between the particles which cause perturbative approaches to be inaccurate, and the fermion sign problem in numerical approaches [21]. In condensed matter experiments, a study of the crossover is also difficult because the interaction strength is usually strongly constrained by the material and variations will be correlated with other effects [22]. Despite of this, the first condensed matter realisation of BCS-BEC crossovers have recently been reported [23, 24].

To gain a better understanding of the BCS-BEC crossover, other approaches are therefore of high interest. The limitations encountered with classical computers could be overcome with quantum computers, which tackle the complications of quantum mechanics by directly incorporating quantum mechanics into their principles of operation [25]. However, only small, noisy quantum computers with limited capabilities exist today and universal fault-tolerant versions, necessary to tackle many-body problems like the crossover, might not be available in the near future [26]. An alternative are quantum simulators, machines which employ quantum mechanics and are purpose-built to solve one (or maybe

several similar) quantum-mechanical problems [27]. Usually, the quantum simulator is a physical system with well controllable and widely tunable parameters which realises a specific theoretical model, in this thesis it will the BCS-BEC crossover. Measurements on the simulator will then provide results of the realised model. A platform that has proven to be well suited for the simulation of models from condensed matter physics are ultracold atoms [28, 29]. With several cooling techniques developed in the last decades [30, 31] large, gaseous samples of neutral atoms, usually on the order of 10^4 to 10^6, are brought down to temperatures on the order of 100 nK where the quantum nature of the atoms dominates their behaviour. Highly controllable electromagnetic fields can be employed to tune various properties of these systems, for example the atoms' interactions and internal states as well as the geometries in which these atoms move [32]. Some successes of this approach have been the realisations of the weakly repulsive Bose gas, where Bose-Einstein condensation [33, 34], the superfluidity of the condensate [35, 36], and a reduction of the critical temperature due to an external potential [37] have been observed, the Bose Hubbard model whose superfluid to Mott insulator transition was observed [38], and the BCS-BEC crossover [39, 40] and demonstration of its superfluidity [41].

In this thesis, an ultracold gas of ^{6}Li atoms is used as a quantum simulator for the BCS-BEC crossover. ^{6}Li atoms are extremely suitable for this purpose because their attraction strength is easily controlled with high accuracy and because they realise the crossover model almost perfectly. Since the first realisation almost two decades ago [39, 40], many aspects of the crossover have been studied [28, 42–44]. Nevertheless, open questions are still left.

Remarkably, the critical temperature for superfluidity throughout the crossover is one of these. This is surprising because it can be considered as the most important property by being the highest of all fermionic superfluids found so far. On the intrinsic temperature scale set by the Fermi temperature T_F, the critical temperature of the unitary Fermi gas, a special point in the middle of the crossover where the scattering length of the interaction diverges, is $T_\mathrm{c} = 0.167\,T_\mathrm{F}$ [42]. This is about three times higher than the critical temperature of the high T_c superconductors [23]. In the unitary limit, the precise determination of the critical temperature was possible due to the strongly simplified thermodynamics [45]. Away from the unitary case however, the critical temperature is less well known. On general grounds, a local maximum of the critical temperature is expected in the crossover [31]. Neither the precise attraction strength at which this maximum occurs nor its prominence are known. A big obstacle is the lack of a clear signal of superfluidity in the common observables of Fermi gases, usually density or momentum distributions.

To determine the critical temperature throughout the crossover, we therefore developed and applied novel detection techniques for condensation based on machine learning. Machine learning is a currently quickly developing programming paradigm with a vast range of applications including chatbots [46] and image and speech recognition [47]. Some examples for applications in physics are the optimisation of experiments [48], the detection of phase transitions [49] and calculations of the BCS-BEC crossover [50, 51]. Many more exist and new ones are published frequently at the moment. The key advantage of machine learning techniques is the possibility to analyse data without a pre-assumed model. In this sense, it is well suited to detect the faint signals of superfluidity in an unbiased way. The new detection techniques therefore allowed us to measure the critical temperature for a wide range of interaction strengths and establish the phase diagram with unprecedented accuracy.

The superfluid phase of the BCS-BEC crossover is an example of a system with spontaneously broken symmetry. Spontaneous symmetry breaking is a process encountered frequently in physics

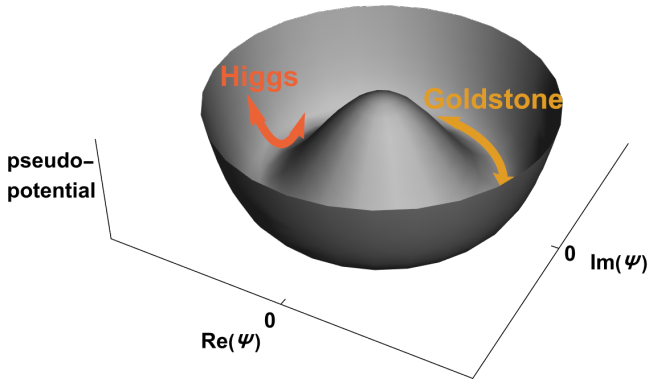

Figure 1.1: Pseudopotential for the superfluid below the critical temperature. Goldstone and Higgs modes, the two excitations arising due to the broken U(1) symmetry of the phase of the macroscopic wave function, are shown.

[52] with examples ranging from the profane, a pencil balanced on its tip falling over, to central to our understanding of the universe: the discovery of the Higgs boson confirmed that the universe is currently an electroweak condensate and therefore underwent spontaneous symmetry breaking at some point [53].

In the superfluid case, spontaneous symmetry breaking is related to the macroscopic wave function. Its amplitude in equilibrium is constrained to a specific value by the system's properties – density, temperature, interactions etc. – but this is not the case for the phase of the wave function which can take on an arbitrary value and is therefore called a continuous symmetry. Nevertheless, every superfluid takes on a fixed phase, which also influences observations [54, 55], upon formation and in this way breaks the symmetry spontaneously.

Spontaneously broken continuous symmetries give rise to two general classes of excitations. These are most intuitively identified from a graphical representation of the U(1) phase symmetry broken at the superfluid phase transition. The macroscopic wave function Ψ can be represented as a point in the pseudopotential shown in figure 1.1. Upon formation, Ψ takes on a random value in the circular minimum. From this equilibrium state, two general types of excitations are now possible: Goldstone modes along the minimum and Higgs modes perpendicular to it.

The Goldstone excitations of superfluids are observable as phonons and therefore well understood. Higgs modes on the other hand are much less studied in both condensed matter and ultracold atomic gases. In general, experimental studies of the Higgs mode are challenging due to difficulties in exciting the mode. As a consequence almost all observations are very recent, see [56] and references therein. In the BCS-BEC crossover only a single study, performed previously in our group, is available [57]. More studies are therefore necessary, especially since the Higgs mode can reveal important properties of the crossover. In this regard, two aspects are sticking out. Firstly, the frequency of the Higgs mode is a way to measure the amplitude of the macroscopic wave function $|\Psi|$ [58]. The evolution of this value through the crossover is not precisely known yet. Secondly, the Higgs mode can reveal the extend to which the system is particle-hole symmetric. Particle-hole symmetry is necessary for the stability of the Higgs mode in condensed matter systems [59] and a decay of this symmetry is expected to happen in the transition from the BCS to the BEC limit. Consequently, the Higgs mode is expected

to gradually become less stable throughout the crossover [60]. This decay has so far not been studied in detail.

For this thesis, the Higgs mode of a trapped gas of ultracold ^{6}Li was excited by a new method, an abrupt magnetic field change that leads to a quench of the attraction strength. The ensuing dynamics were observed in time domain revealing internal effects in the gas as well as a coupling to the modes of the external trapping potential. The new excitation method made it possible to study the evolution of these dynamics for a wide range of interaction strengths in the BCS-BEC crossover.

Outline of the thesis

In chapter 2, the general theoretical foundations of this thesis are briefly reviewed. Special attention is given to the BCS-BEC crossover, the Higgs mode and some smaller topics necessary to understand the investigations of this thesis. In chapter 3, the experimental platform used to produce the results of this thesis is described and the most relevant aspects of its working principles are introduced. Several methods to detect superfluidity are discussed in chapter 4. While some of them are established techniques, newly developed ones based on machine learning are also introduced. In chapter 5, the new detection schemes for superfluidity are used to measure the phase diagram of the BCS-BEC crossover with unprecedented precision. Chapter 6 contains the main results of this thesis. It presents and analyses the dynamics of a superfluid in the BCS-BEC crossover whose interaction strength has been quenched to excite the Higgs mode. In chapter 7, some very recent measurements in which the interaction strength in the crossover was modulated instead of quenched are presented. Finally, chapter 8 summarises the most important results and gives some perspectives on possible further studies.

CHAPTER 2

Theory of Fermions with Attractive Interactions

This chapter gives a brief review of the theoretical background for the topics investigated in this thesis. A key point is the condensation of fermions with attractive interactions into a superfluid, which is a generalisation of superconductivity. Two main properties of the condensate are discussed: the critical temperature for condensation and the Higgs mode, a unique excitation of the superfluid. The strength of the attractive interactions has a strong influence on both of them and this is the main focus of the chapter. Finally, the consequences of external potentials are briefly discussed.

2.1 Non-interacting fermions

Fermions are particles with half-integer spin S, i.e. $S = \frac{2n+1}{2}\hbar$ with $n \in \mathbb{N}$ and $\hbar$ the reduced Planck constant. Since the research in this thesis is inspired by physical systems comprised of electrons or neutrons, which both have $S = \frac{\hbar}{2}$, only spin-$\frac{1}{2}$ fermions are considered in this thesis. These fermions have two internal states – the two possible projections of the Spin onto a quantisation axis – which will be denoted as $|\uparrow\rangle$ and $|\downarrow\rangle$. From the spin-statistics theorem [61] follows the Pauli-principle: two fermions cannot occupy the same state, characterised by both internal and external degrees of freedom at any given time. This leads to the Fermi-Dirac distribution function $f(E,T)$ which gives the probability that a state with energy E, which will be related to the state's quantum numbers, is occupied at temperature T

$$f(E,T) = \frac{1}{\exp\left(\frac{E-\mu(T)}{k_\mathrm{B}T}\right)+1}. \tag{2.1}$$

Here, k_B denotes the Boltzmann constant and $\mu(T)$ is the chemical potential of the system, which also depends on temperature. At $T = 0$, all states up to the energy $\mu(T{=}0)$, which is also called Fermi energy E_F, are always occupied while all states above E_F are never occupied. Collectively, the states below E_F are called the Fermi-sea and the step in the distribution function is called the Fermi-edge or the Fermi-surface.

At non-zero temperatures, additional thermal energy is distributed among the particles such that some states above E_F have a probability of being occupied while some below E_F are no longer always occupied leading to a softening of the Fermi-edge. Since new particles can now be added to the system in states below the Fermi energy, μ the necessary energy to do so decreases with T. The exact behaviour of this decrease depends on the specific system [62]. The decrease of μ and the

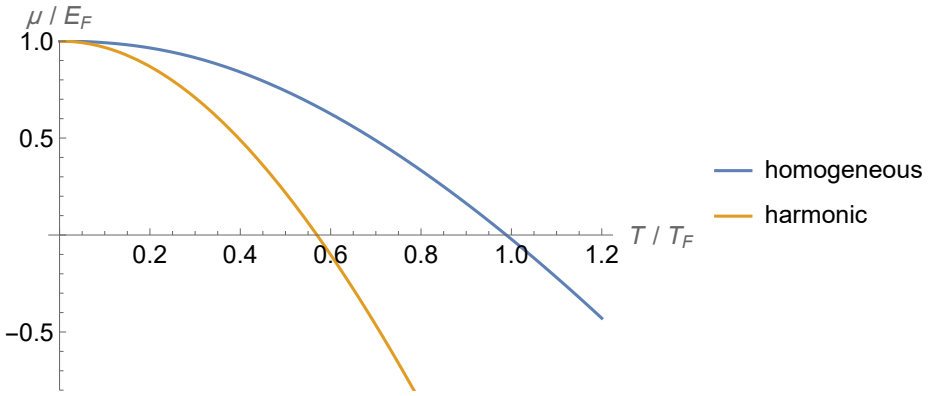

Figure 2.1: The temperature dependence of the chemical potential of a non-interacting three-dimensional Fermi gas for a system with homogeneous density and one with an external harmonic potential.

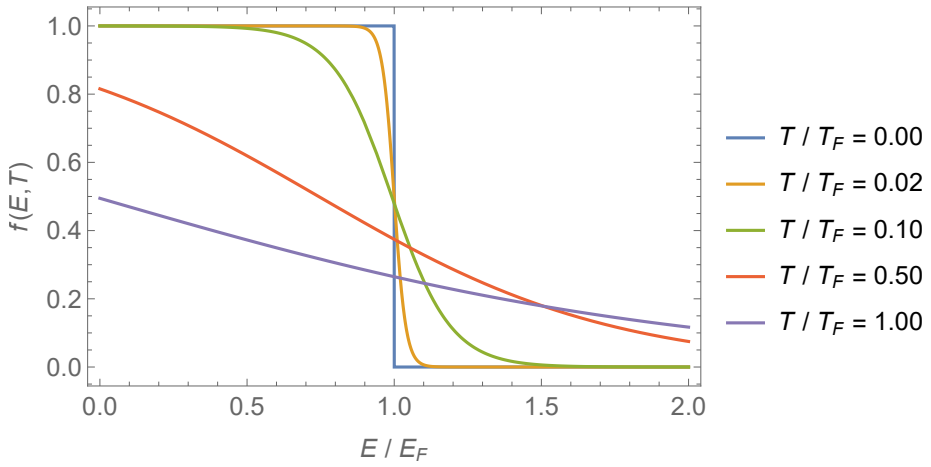

Figure 2.2: The Fermi-Dirac distribution function for a homogeneous Fermi gas at various temperatures. The increasingly smoother Fermi-edge with rising temperature is visible.

smoothening of the Fermi-edge are depicted in figures 2.1 and 2.2 with temperatures given on the natural scale $T_\mathrm{F} = E_\mathrm{F}/k_\mathrm{B}$. In the limit of high temperatures $T \gg T_\mathrm{F}$, the Fermi-Dirac distribution approaches the Boltzmann distribution of classical thermodynamics.

A special case is the Fermi gas with homogeneous density no interactions. In this situation the energy of a particle only depends on the absolute value of its momentum and consequently, the Fermi-edge at low temperatures does not only exist at E_F in energy space but also in momentum space at

$$k_\mathrm{F} = \sqrt{2mE_\mathrm{F}}/\hbar \tag{2.2}$$

where m denotes the mass of the fermions. In this situation, k_F is related to the density of the gas n by

$$k_\mathrm{F} = \left(3\pi^2 n\right)^{1/3}. \tag{2.3}$$

2.2 Fermions with weak attraction – BCS-theory

The non-interacting Fermi gas has, apart from the formation of the Fermi-sea, no particularly remarkable properties that arise when it is cooled towards zero temperature. This behaviour changes quite drastically once a small attractive interaction is introduced to the fermions.

A contact interaction between two fermions with a relative potential $V(\boldsymbol{r}) = v_0\delta(\boldsymbol{r})$, where $v_0 < 0$ and $\delta(\boldsymbol{r})$ is the Dirac delta function, leads, in momentum space, to the Hamiltonian

$$\hat{H} - \mu\hat{N} = \sum_{\boldsymbol{k},\sigma} \xi_{\boldsymbol{k}} \hat{c}^\dagger_{\boldsymbol{k},\sigma}\hat{c}_{\boldsymbol{k},\sigma} + v_0 \sum_{\boldsymbol{k},\boldsymbol{k}',\boldsymbol{q},\sigma,\sigma'} \hat{c}^\dagger_{\boldsymbol{k}+\boldsymbol{q},\sigma}\hat{c}^\dagger_{\boldsymbol{k}'-\boldsymbol{q},\sigma'}\hat{c}_{\boldsymbol{k}',\sigma'}\hat{c}_{\boldsymbol{k},\sigma} \tag{2.4}$$

where $\hat{c}_{\boldsymbol{k},\sigma}$ ($\hat{c}^\dagger_{\boldsymbol{k},\sigma}$) annihilates (creates) a fermion with momentum $\boldsymbol{k}$ and spin $\sigma = \{\uparrow,\downarrow\}$ and $\xi_{\boldsymbol{k}} = \frac{\hbar^2\boldsymbol{k}^2}{2m} - \mu$ is the dispersion relation of free fermions with mass m measured from the chemical potential while $\boldsymbol{q}$ is the momentum transferred during the scattering process.

For the experiments performed for this thesis, the only relevant case is that of *ultracold* scattering: a situation in which the energy of the scattering particles is so low that the lowest-energy partial wave, the s-wave, of the full scattering amplitude's partial wave expansion is the only relevant one [31]. For fermions, this situation has an important implication due to an alternative formulation of the spin-statistics theorem. It states that the total wave function of indistinguishable fermions has to be antisymmetric under particle exchange [61]. Since the s-wave scattering amplitude is symmetric, the spin wave function of two scattering fermions must be antisymmetric which leaves the spin singlet as the only option for spin-$\frac{1}{2}$ fermions. Scattering is therefore only possible between particles with opposite spin, so $\sigma = -\sigma'$.

If the momentum dependence of the scattering process is considered, the largest scattering probability is found for fermions with opposite momentum $\boldsymbol{k} = -\boldsymbol{k}'$. This is due to the fact that they have the largest number of final states that conserve momentum and energy. A closer examination reveals that it is reasonable to only consider these scattering processes [31]. Together, these approximations lead to the so-called "BCS-Hamiltonian"

$$\hat{H} - \mu\hat{N} = \sum_{\boldsymbol{k},\sigma} \xi_{\boldsymbol{k}} \hat{c}^\dagger_{\boldsymbol{k},\sigma}\hat{c}_{\boldsymbol{k},\sigma} + v_0 \sum_{\boldsymbol{k},\boldsymbol{k}'} \hat{c}^\dagger_{\boldsymbol{k}\uparrow}\hat{c}^\dagger_{-\boldsymbol{k}\downarrow}\hat{c}_{-\boldsymbol{k}'\downarrow}\hat{c}_{\boldsymbol{k}'\uparrow}. \tag{2.5}$$

An arbitrarily weak attractive potential does not support a bound state between the scatterers in three dimensions, but the presence of the Fermi-sea has an important implication if the system is cold enough: the effects of a weak interaction are confined to a narrow region around the Fermi-surface because any excitations from the inner Fermi sea would require more energy than available from the interaction. The density of states relevant for this problem $g(E)$ is therefore effectively just its value at the Fermi-surface $g(E_\mathrm{F})$ and thus a constant. A constant density of states is also encountered in two-dimensional systems which do support a bound state for arbitrarily weak attractive interactions. For these reasons, the many-body Hamiltonian 2.5 supports the pairing of fermions with opposite spin and momentum for all values $v_0 < 0$ and these are called Cooper pairs [31]. The total momentum and spin of these pairs are then both equal to zero.

For weak interactions, the Hamiltonian can be solved by a mean-field approach that looks at fluctuations around the non-zero expectation value of the pair annihilation operator $\langle\hat{c}_{-\boldsymbol{k}\downarrow}\hat{c}_{\boldsymbol{k}\uparrow}\rangle$. A Bogoliubov-transformation reveals that quasiparticles are the fundamental excitations of the system.

Their dispersion relation is given by $E_{\boldsymbol{k}} = \sqrt{\xi_{\boldsymbol{k}}^2 + |\Delta|^2}$ with

$$\Delta = v_0 \sum_{\boldsymbol{k}} \langle \hat{c}_{-\boldsymbol{k}\downarrow} \hat{c}_{\boldsymbol{k}\uparrow} \rangle \,. \tag{2.6}$$

The quasiparticles are the mixture of particles and holes characteristic for a Bogoliubov-transformation. The corresponding operators $\hat{\gamma}_{\boldsymbol{k},\uparrow}$ and $\hat{\gamma}_{\boldsymbol{k},\downarrow}$ are given by

$$\begin{aligned} \hat{\gamma}_{\boldsymbol{k},\uparrow} &= u_{\boldsymbol{k}}^* \hat{c}_{\boldsymbol{k},\uparrow} - v_{\boldsymbol{k}}^* \hat{c}_{-\boldsymbol{k},\downarrow}^\dagger \\ \hat{\gamma}_{-\boldsymbol{k},\downarrow}^\dagger &= u_{\boldsymbol{k}} \hat{c}_{-\boldsymbol{k},\downarrow}^\dagger + v_{\boldsymbol{k}} \hat{c}_{\boldsymbol{k},\uparrow}. \end{aligned} \tag{2.7}$$

The momentum-dependent coefficients $u_{\boldsymbol{k}}$ and $v_{\boldsymbol{k}}$ in the quasiparticle operators operators are given by

$$\begin{aligned} \left|u_{\boldsymbol{k}}\right|^2 &= \frac{1}{2}\left(1 + \frac{\xi_{\boldsymbol{k}}}{E_{\boldsymbol{k}}}\right) \\ \left|v_{\boldsymbol{k}}\right|^2 &= \frac{1}{2}\left(1 - \frac{\xi_{\boldsymbol{k}}}{E_{\boldsymbol{k}}}\right). \end{aligned} \tag{2.8}$$

The ground state wave function $|\Psi_{\mathrm{BCS}}\rangle$ can be found by requiring the absence of quasiparticles, i.e. $\hat{\gamma}_{\boldsymbol{k},\sigma} |\Psi_{\mathrm{BCS}}\rangle$=0. It has the form [9]

$$|\Psi_{\mathrm{BCS}}\rangle = \prod_{\boldsymbol{k}} \left(u_{\boldsymbol{k}} + v_{\boldsymbol{k}} \hat{c}_{\boldsymbol{k},\uparrow}^\dagger \hat{c}_{-\boldsymbol{k},\downarrow}^\dagger\right) |0\rangle \,. \tag{2.9}$$

It can be seen that $|v_{\boldsymbol{k}}|^2$ is the momentum distribution of the fermions while $|u_{\boldsymbol{k}}|^2$ is that of the holes and thus $|v_{\boldsymbol{k}}|^2 + |u_{\boldsymbol{k}}|^2 = 1$. In the limit of weak interactions considered for now, the distribution functions are assumed to be largely unchanged by the interaction and therefore $\mu = E_{\mathrm{F}}$ as long as finite temperatures are not considered. This also implies that $\Delta \ll E_{\mathrm{F}}$. Then, the quasiparticle dispersion relation has a non-zero minimum of $|\Delta|$ at k_{F} and for this reason Δ is called the gap of the system. This case is shown in figure 2.3.

An important consequence of the pairing is a modification of the fermions' distribution in momentum space: If higher momentum states are occupied, the particles can reduce their relative distance without violation of the Pauli-principle and therefore reduce the total energy by the gain in interaction energy. This leads to a softening of the Fermi edge as depicted in figure 2.4. The gap can also be related to the sum $\sum_{\boldsymbol{k}} u_{\boldsymbol{k}} v_{\boldsymbol{k}}$ [17] which is dominated by the softened region where neither $u_{\boldsymbol{k}}$ nor $v_{\boldsymbol{k}}$ are close to zero. From this it follows that the softened Fermi surface has a width in energy space that roughly corresponds to $|\Delta|$. Qualitatively, the softening due to pairing is very similar to the influence of temperature as shown above in figure 2.2. This will become relevant later in the thesis.

The ground state 2.9 of the system is a state in which all fermions are paired into Cooper pairs. It has a property called off diagonal long range order [9], defined as

$$\lim_{|\boldsymbol{r}-\boldsymbol{r}'|\to\infty} \langle \hat{c}_\uparrow^\dagger(\boldsymbol{r}) \hat{c}_\downarrow^\dagger(\boldsymbol{r}) \hat{c}_\downarrow(\boldsymbol{r}') \hat{c}_\uparrow(\boldsymbol{r}') \rangle \neq 0 \tag{2.10}$$

where $\hat{c}_\sigma(\boldsymbol{r})$ is the annihilation operator for a particle with spin σ at position $\boldsymbol{r}$. This means that the removal of a pair at position $\boldsymbol{r}'$ with the simultaneous addition of another one far away at $\boldsymbol{r}$ leaves the

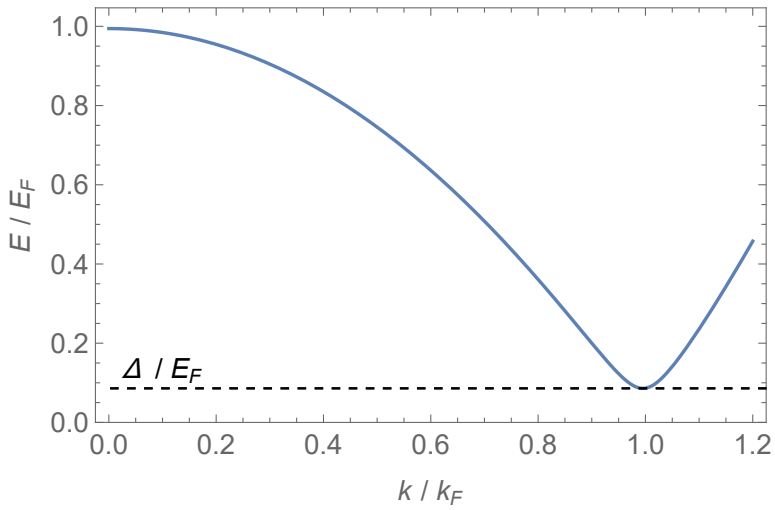

Figure 2.3: Quasiparticle dispersion relation for a weak attraction strength. The attraction leads to a gapped spectrum with minimal excitation energy $\Delta \ll E_{\mathrm{F}}$ at k_{F}.

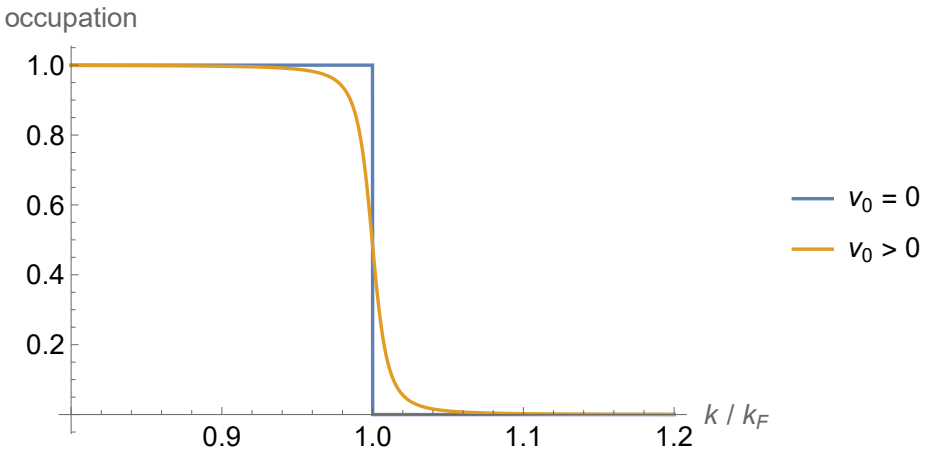

Figure 2.4: Softening of the Fermi edge at zero temperature due to pairing. Shown is the momentum distribution of fermions $|v_{\boldsymbol{k}}|^2$.

system in a state that still has a finite overlap with the initial state. Interpreted differently, this is a flow without an impact on the system. This is the frictionless flow which is the hallmark of superfluidity, or superconductivity in the case of charged fermions [63].

2.3 The superfluid phase transition in the BCS limit

In this section, the effects of temperature, which were not considered in the previous one, are considered. With increasing temperature, thermal quasiparticle excitations destroy the quantum coherence necessary to support a non-zero $|\Delta|$. The gap therefore decreases monotonically until Δ=0 is reached at the critical temperature T_c. Due to the similarity between equations 2.6 and 2.10, this coincides with the end of off diagonal long range order and therefore the system cedes to be superfluid and enters the normal phase at T_c.

A closer examination would reveal that this constitutes a second order phase transition as described by Ginzburg-Landau theory which approximates the free energy in proximity to the phase transition as a power series of a complex order parameter ψ. In this case, it is the macroscopic wave function of the superfluid. It has been shown that ψ can be equated to Δ [9]. It is indeed a complex quantity by being the expectation value of a non-hermitian operator. In a homogeneous system without external fields, the power series for the free energy takes the form

$$F = F_0 + \alpha|\Delta|^2 + \frac{\beta}{2}|\Delta|^4. \tag{2.11}$$

Above the critical temperature, the real parameters α and β obey $\alpha > 0$ and $\beta > 0$ so that the free energy is minimised for $\Delta = 0$ and becomes F_0, the free energy of the normal phase. Below T_c however, we have $\alpha < 0$ and $\beta > 0$ so that the free energy is minimised for a finite order parameter

$$|\Delta|^2 = -\frac{\alpha}{\beta}. \tag{2.12}$$

The absolute value of the order parameter is therefore constrained, but nothing fixes the phase of Δ. This situation below T_c is depicted graphically in figure 2.5. Going back to the definition of Δ in equation 2.6, it can be seen that the phase of Δ is related to the phase of the system's wave function. It is therefore possible to gauge the phase of Δ to an arbitrary value. Accordingly, the phase transition to the superfluid phase is an example of spontaneous symmetry breaking [56]. For the rest of this thesis, Δ is chosen to be real and $\Delta \geq 0$.

Two distinct excitations of the system become possible in the symmetry-broken phase. Intuitively, this can be seen by interpreting the free energy landscape in figure 2.5 as a pseudopotential in which the system "moves". One option is the movement along the circular valley corresponding to fluctuations in the order parameter's phase. The related excitations of the system in the superfluid case are phonons and their existence is a manifestation of Goldstone's theorem [52][1]. Since the pseudopotential is flat in this direction, no minimal energy is required to create a phonon, just as one would expect from their usual dispersion relation. This mode is called the phase mode or the (Nambu-)Goldstone mode. If the fermions in the system were electrically charged, the Goldstone mode would couple to electromagnetic

[1] In the previous section 2.2, phononic excitations with their characteristic density modulations were not found because a homogeneous density was assumed throughout.

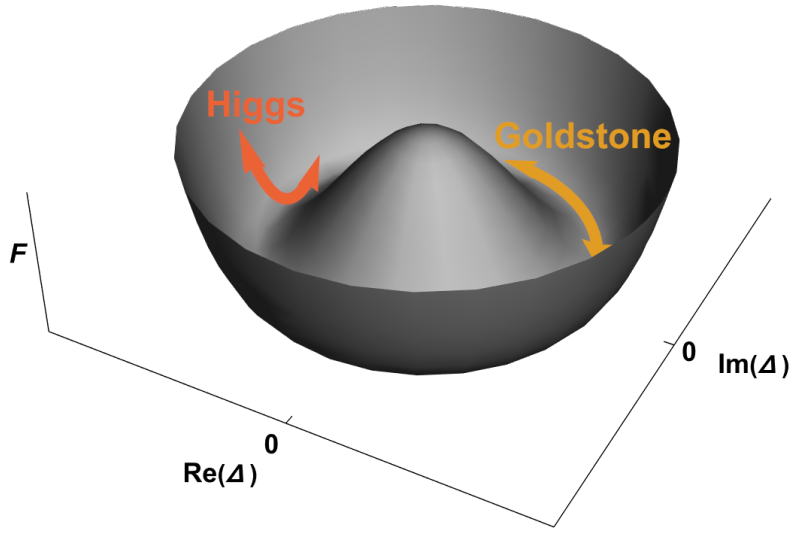

Figure 2.5: The free energy landscape of a superfluid with order parameter Δ. The two basic excitations connected to the spontaneously broken symmetry, Goldstone- and Higgs-mode, are illustrated.

excitations which would shift its minimal energy to the plasma frequency [59]. This has remarkable consequences like the Meißner effect for superconductors [56] but is not relevant for the studies with neutral fermions presented in this thesis.

The other possible excitation is perpendicular to the trench of minimal free energy and therefore a modification of the order parameter's amplitude. This resembles the excitation of a harmonic oscillator and consequently a finite amount of energy is necessary to excite this mode. It is called the amplitude or (Anderson-)Higgs mode and its minimum energy is given by 2Δ [53]. If this mode is excited, it also manifests itself as an oscillation of the order parameter's value with frequency $2\Delta/h$ and a characteristic damping of $t^{-1/2}$ [64]. A measurement of this oscillation is presented later in this thesis in chapter 6.

A closer examination of the Higgs mode's origin would reveal that the heuristic derivation via the pseudopotential misses a few important points: mainly the Higgs mode being only stable in Lorentz-invariant systems, a condition that is not fulfilled for any condensed matter system and therefore also not for the studies in this thesis. For weak interactions, the BCS Hamiltonian can however be rewritten in a form that reveals its equivalence to the relativistic Dirac-Hamiltonian. It therefore is effectively Lorentz-invariant and supports a Higgs mode [59, 65].

Excitations of the Higgs mode are challenging because it lacks the typical properties often used for excitations, for example charge, electric or magnetic moments or quantum numbers [56]. Instead, it is necessary to "shake" the superfluid out of its equilibrium [53]. For the superfluid, this involves a manipulation of Δ_{eq}, the equilibrium value of the order parameter. This is most easily done by changing the strength of the interaction. One option is a sudden quench of the system to a different interaction strength with a new value of Δ_{eq}. Another option is the time-dependent modulation of

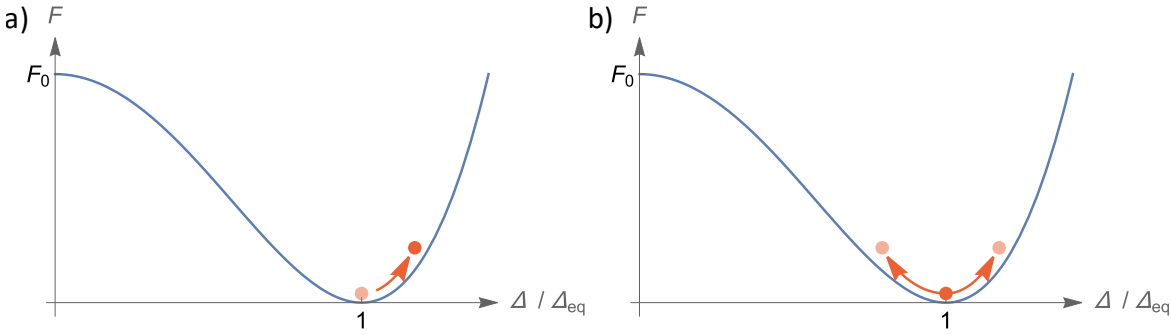

Figure 2.6: Schematic representation of the excitation mechanisms for the Higgs mode. a) A quench changes Δ_{eq} faster than Δ can follow. This excites the system to a higher point in the pseudopotential and therefore initiates the Higgs mode. b) A time-dependent modulation of Δ_{eq} drives the Higgs mode.

the interaction strength, and correspondingly Δ_{eq}, around a central value. A graphical representation of these excitation mechanisms is shown in figure 2.6. For a proper excitation of the Higgs mode, it is important that the manipulations happen on a faster time scale than the equilibration time of the system. Both of these excitation mechanism will be studied later in this thesis.

2.4 Fermions with arbitrary attractive interactions – the BCS-BEC crossover

This section extends the previous discussions, limited to weak attractive interactions, by the inclusion of arbitrarily strong, but still attractive, interactions. This makes it necessary to quantify the interaction strength more rigorously.

One could expect that the potential depth v_0 of the Hamiltonian 2.5 is the relevant quantity, but this turns out to be wrong. The reason lies in divergences arising in the calculation of the scattering amplitude if more than the first order Born-approximation is taken into account [31]. A regularisation depending on the exact shape of the potential would therefore in principle be necessary.

In the case of ultracold scattering, the range of the potential is however usually much shorter than the interparticle distance, so that the physical effect of s-wave scattering can be completely described by the scattering length a. A proper parameter for the interaction strength is therefore the dimensionless quantity $1/k_{\mathrm{F}}a$. The previously discussed BCS limit of arbitrarily weak attraction is reached for $1/k_{\mathrm{F}}a \to -\infty$ and in this limit the gap is related to the interaction parameter via

$$\Delta = \frac{8}{e^2} E_{\mathrm{F}} \exp\left(\frac{\pi}{2k_{\mathrm{F}}a}\right). \tag{2.13}$$

Likewise, the critical temperature is given by

$$T_{\mathrm{c,BCS}} = \frac{e^{\gamma}}{\pi}\frac{8}{e^2} T_{\mathrm{F}} \exp\left(\frac{\pi}{2k_{\mathrm{F}}a}\right) \tag{2.14}$$

with the Euler-Mascheroni constant γ. The first order correction to this result is a reduction by a factor $(4e)^{1/3} \approx 2.2$ known as the Gorkov-Melik-Barkhudarov (GMB) correction [66].

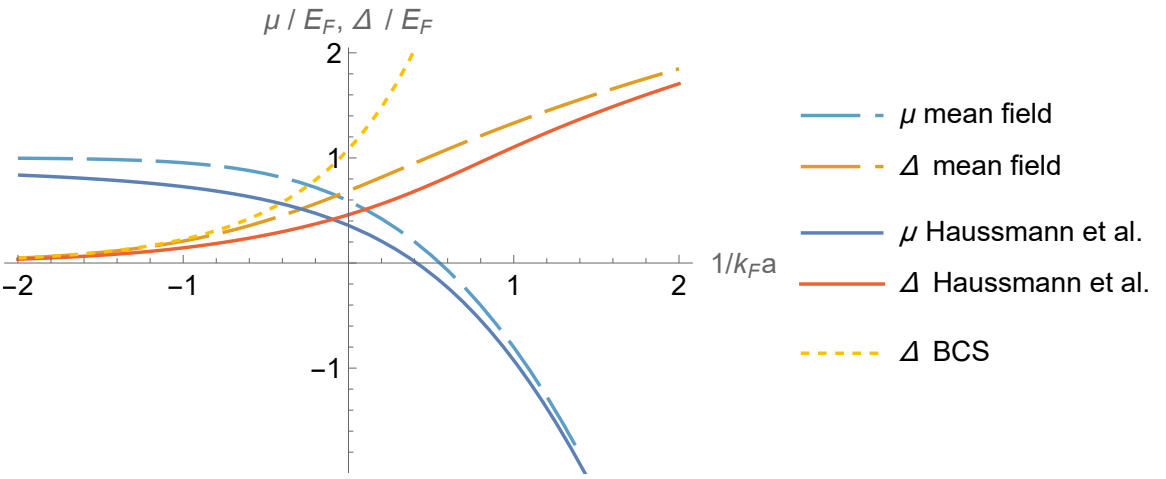

Figure 2.7: Different solutions for Δ and μ in the BCS-BEC crossover. For details on the mean field or Haussmann et al. solutions refer to [31] and [67], respectively. The gap calculated with equation 2.13 is denoted as Δ BCS.

When extending the mean-field BCS solution to stronger interactions, the chemical potential can no longer be considered equal to E_F because of the softening of the Fermi edge as shown in figure 2.4. It is then necessary to find a self-consistent solution for both Δ and μ. The result cannot be given by a simple relation like 2.13 any more although an analytic form can be obtained [31]. By simple means, only a numerical solution for Δ and μ in mean field approximation is easily calculated and this is presented in figure 2.7. The mean field solution is however only expected to be reliable in the limit of weak attractions or $1/k_\mathrm{F}a \ll -1$ and a calculation based on a variational many-body approach, performed in [67], indeed shows substantial deviations, as can be seen in the same figure.

The critical temperature in the BCS limit is equivalent to the onset of fermion pairing. This correspondence no longer holds for stronger attractions because a second pairing mechanism arises: A two-body bound state is supported by the interaction potential if the interaction gets strong enough. This bound state is present for $1/k_\mathrm{F}a > 0$ and has a binding energy that increases with a as $E_\mathrm{B} = -\hbar^2/(2ma^2)$ in first order approximation.

In the limit $1/k_\mathrm{F}a \to \infty$ of extremely strong attraction, the fermions will be tightly bound into dimers which can be regarded as composite bosons with mass $2m$. In an atomic context the dimers are also called molecules[2]. The scattering length between dimers has been shown to be $a_\mathrm{DD} = 0.6a$ [68]. In the limit $1/k_\mathrm{F}a \to \infty$, a_DD is mall and positive and therefore describes a weakly repulsive interaction between dimers. The critical temperature of superfluidity will no longer be connected to pairing at all in this case, but is instead equal to the critical temperature of Bose-Einstein condensation (BEC) of bosons with mass $2m$ and scattering length a_DD. It is given by [69, 70]

$$T_{\mathrm{c,BEC}} = \left(1 + \tilde{c}\, k_\mathrm{F}a\right) \frac{2\pi}{\left(6\pi^2 \zeta\left(\frac{3}{2}\right)\right)} T_\mathrm{F} \approx 0.218 \cdot (1 + 0.2\, k_\mathrm{F}a)\, T_\mathrm{F} \tag{2.15}$$

[2] These molecules are however in a highly excited electronic state.

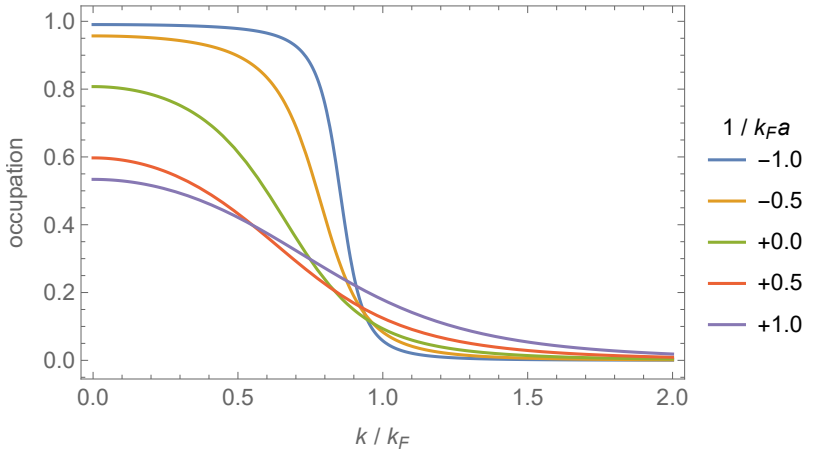

Figure 2.8: The smooth evolution of the fermionic momentum distribution in the crossover at $T = 0$.

with ζ denoting the Riemann zeta function and the constant $\tilde{c} \approx 1.3$. A connection from Fermi temperature to the seemingly unrelated phenomenon of Bose-Einstein condensation was introduced in the derivation of equation 2.15 by replacing the dimer density n_D, the relevant quantity for BEC formation, with the density of the constituent fermions $n = 2n_\mathrm{D}$, which is related to the Fermi momentum via equation 2.3.

Between the two limits of fermionic Cooper pairing and dimerised Bose-Einstein condensation no phase transition occurs [14]. Instead, there is a generalised form of pair condensation smoothly connecting them and this behaviour is called the BCS-BEC crossover. The generalised pairing is predominantly fermionic in the BCS regime $1/k_\mathrm{F}a < -1$ and shifts the pairing weight to a molecular type through the *strongly interacting regime* $-1 < 1/k_\mathrm{F}a < 1$ until it is predominantly molecular in the BEC regime $1/k_\mathrm{F}a > 1$. The smoothness of the crossover is most evident in the size of the pairs: it is significantly larger than the interparticle spacing in the BCS limit, gradually shrinks to become comparable to the spacing around $1/k_\mathrm{F}a = 0$ and continues to shrink to much smaller than the spacing in the BEC limit where the pairs are tightly bound.

The smoothness of the crossover is also reflected in the momentum distribution of the fermions in the ground state which is given by the same $|v_k|^2$ as in the BCS limit, see equation 2.8 and [31]. Figure 2.8 shows the evolution of this distribution through the crossover calculated with the values for μ and Δ from [67]. It becomes apparent that the increasing attraction strength leads to a continuously softening Fermi edge.

The best possible distinction between fermionic and bosonic behaviour is provided by the chemical potential. The quasiparticle dispersion relation $E_{\boldsymbol{k}}$ has its minimum at $|\boldsymbol{k}| = k_\mu$ which is defined as $\hbar^2 k_\mu^2/2m = \mu$. Around $1/k_\mathrm{F}a \approx 0.6$, the zero-temperature value of μ becomes negative and k_μ is no longer well-defined. The minimum in $E_{\boldsymbol{k}}$, the gap of the superfluid, is then located at $k = 0$ and given by $\sqrt{\Delta^2 + \mu^2}$. Such a dispersion relation is closer to a bosonic than a fermionic superfluid. Alternatively, it can be argued that the Fermi edge, located at μ, is no longer present if $\mu < 0$ and the system cedes to be fermionic and becomes bosonic at this point.

A special situation is encountered in the middle of the crossover at $1/k_\mathrm{F}a = 0$, called the unitary limit. Here, the scattering length diverges to $\pm\infty$ and can therefore no longer be a relevant length scale

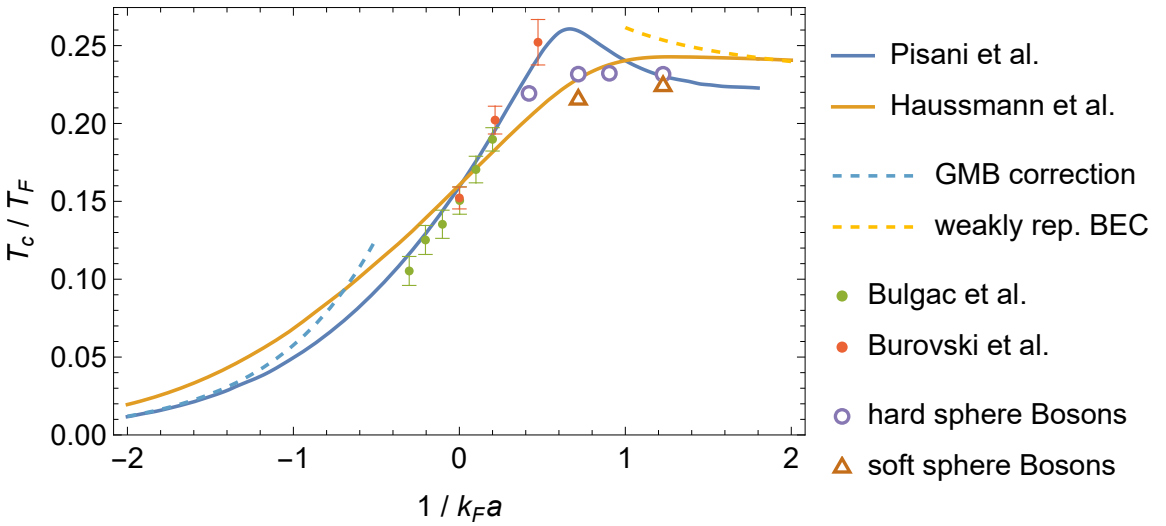

Figure 2.9: Several theoretical predictions of the critical temperature in the BCS-BEC crossover: An extended GMB correction by Pisani et al. [71], a variational many body approach by Haussmann et al. [67], fermionic quantum Monte Carlo calculations by Bulgac et al. [72] and Burovski et al. [73], equation 2.14 with GMB correction, equation 2.15 for the weakly repulsive Bose gas, and a quantum Monte Carlo calculation for hard and soft sphere bosons by Pilati et al. [74].

for the description of the many-body system. This is a big simplification [42, 45] and consequently, the many-body properties in the unitary case are relatively well known [42].

The generalised pairing encountered the crossover has the consequence that there is no longer a one-to-one correspondence of superfluidity and pairing as in the BCS regime. The critical temperature for superfluidity T_c will instead be lower than the temperature necessary to break the pairs. A calculation of the critical temperature through the crossover is no simple task because the strong interactions do not allow for a perturbative treatment. One interesting property is however certain: there must be a maximum in T_c/T_F in the crossover because this ratio increases when approaching $1/k_F a = 0$ from both limits, cf. equations 2.14 and 2.15. The exact location and pronouncement of the maximum are however not known. Several theoretical predictions are presented in figure 2.9. Measurements of the critical temperature throughout the crossover are presented in chapter 5.

2.4.1 Higgs mode with strong attractions

Above, it was stated that the BCS-Hamiltonian in equation 2.5 is effectively Lorentz-invariant and therefore supports a well-defined Higgs mode. The origin of this property lies in an effective particle hole symmetry of the BCS-Hamiltonian around the Fermi surface at energy μ [75]. Expressed as a formula, the fermion dispersion relation measured from the chemical potential $\xi_{\boldsymbol{k}}$ fulfils

$$-\xi_{\boldsymbol{k}_\mu - \delta\boldsymbol{k}} = \xi_{\boldsymbol{k}_\mu + \delta\boldsymbol{k}} \tag{2.16}$$

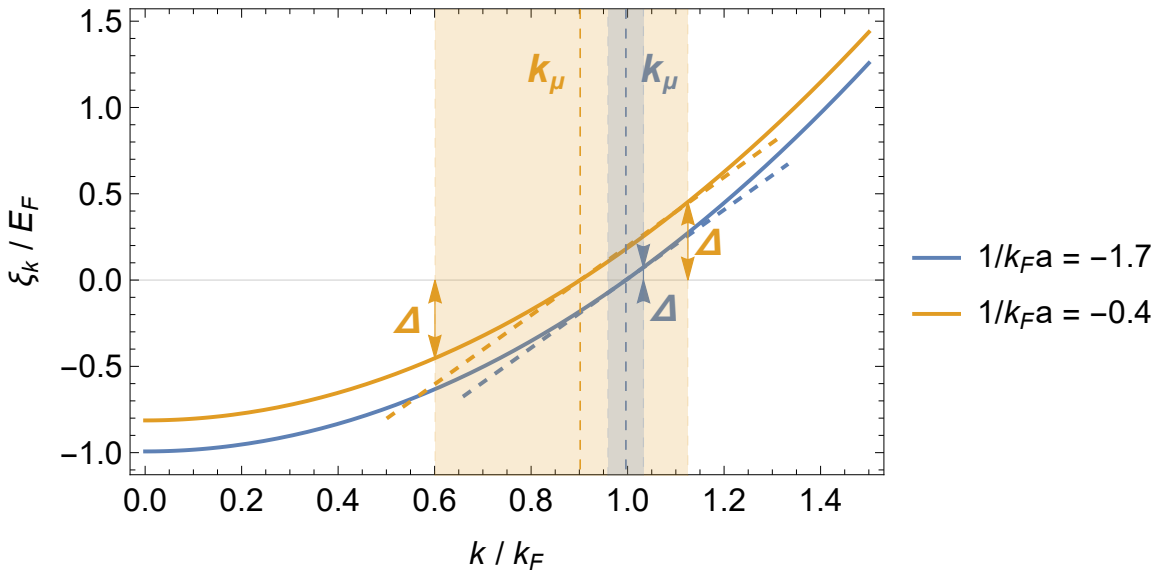

Figure 2.10: Violation of effective particle-hole symmetry in the free-fermion dispersion relation for strong attraction. The shaded areas mark the extend of softened Fermi edge for the respective interaction strengths. If the free fermion dispersion relation (full line) is well described by the linear approximation (dashed line) over the full shaded area of the corresponding interaction strength, the system is effectively particle-hole symmetric. This condition is fulfilled for weak (blue) but violated for strong attractions (orange). Figure inspired by [75].

with $\xi_{\boldsymbol{k}_\mu} = 0$ [3]. For free fermions, $\xi_{\boldsymbol{k}} = \hbar^2\boldsymbol{k}^2/2m - \mu$ does not fulfil this condition, so the BCS-Hamiltonian is in principle not particle hole symmetric. However, the system's low-temperature properties are dominated by the softened region around the Fermi surface with width $\sim \Delta$. If this region is small enough, a linear approximation $\xi_{\boldsymbol{k}} \approx v_\mathrm{F}(k - k_\mu)$ suffices for a proper description and the system is effectively particle-hole symmetric because it now fulfils condition 2.16 [75]. Figure 2.10 shows, at mean-field level and $T = 0$, that the linear approximation is valid in the BCS limit $1/k_\mathrm{F}a < -1$ but not in the regimes with strong attractions $1/k_\mathrm{F}a > -1$, where ΔE_F is not a small quantity.
In the particle-hole symmetric case, additional symmetries are broken for the ground state on top of the $U(1)$ symmetry breaking of Δ and they are the reason for a stable Higgs mode to occur [75].

A violation of particle-hole symmetry does not directly imply that the Higgs mode becomes unobservable. It rather causes a coupling of the Higgs- to the Goldstone mode which leads to a damped Higgs mode. The further the relevant dispersion relation deviates from a particle-hole symmetric one, the stronger the coupling and damping are expected to become. Through the BCS-BEC crossover, the Higgs mode is therefore expected to become less and less stable until it cedes to be observable. For zero temperature, it is predicted to still be observable at unitarity, but not at $1/k_\mathrm{F}a = 1$ [60]. Finite temperatures decrease the Higgs mode's instability as well, but [60] predicts that this effect is not particularly pronounced for $T = 0.9T_\mathrm{c}$ at $1/k_\mathrm{F}a = -1$. The finite temperature theory of the reference is however not applicable to the more interesting case around unitarity.

The frequency of the Higgs mode in the BCS limit was given by twice the gap $\omega_\mathrm{H} = 2\Delta/\hbar$. In

[3] This means that the dispersion relation does not change under a combined reflection along the axes $\xi_{\boldsymbol{k}} = 0$ and $\boldsymbol{k} = \boldsymbol{k}_\mu$, c.f. figure 2.10

the crossover, this equality only holds as long as $\mu \geq 0$ because the energy gap of the quasiparticles becomes $\sqrt{\Delta^2 + \mu^2}$ once $\mu < 0$. Consequently, the Higgs frequency is given by $\omega_{\mathrm{H}} = 2\sqrt{\Delta^2 + \mu^2}/\hbar$ in this case. However, the relation between gap and Higgs frequency is derived in the BCS limit and deviations in the case of strong interactions are possible but have not been much investigated [53].

2.5 External potentials

In the previous sections, most discussions were limited to homogeneous situations where the Hamiltonian contains only kinetic and, potentially, interaction energy. In the experiments presented in the later parts of this thesis, the fermions are placed in an external potential adding another term to the Hamiltonian. The potentials are three-dimensional and (approximately) harmonic so that the potential energy term takes the shape

$$\hat{V}(\boldsymbol{r}) = \frac{m}{2}\left(\omega_x^2 x^2 + \omega_y^2 y^2 + \omega_z^2 z^2\right), \tag{2.17}$$

with trapping frequencies ω_i in the three spatial directions. The Fermi energy for a large number N of fermions with two internal states is in this case calculated to

$$E_{\mathrm{F,harm}} = \hbar\bar{\omega}\,(3N)^{1/3} \tag{2.18}$$

with the geometric mean trapping frequency $\bar{\omega} = \sqrt[3]{\omega_x\omega_y\omega_z}$. Casting aside interactions for the moment, the existence of an external potential causes the fermions' energy to be split into kinetic and potential energy. The result is an inhomogeneous density distribution. This makes a theoretical description much more involved and the trapped gas is therefore usually discussed under certain approximations.

2.5.1 Thomas-Fermi approximation

If the level spacing $\hbar\omega_i$ is much smaller than the other energy scales $k_{\mathrm{B}}T$ and μ, it is possible to use the semi-classical Thomas-Fermi-approximation, which approximates the discrete states as a continuum [31], and write the Fermi-Dirac distribution 2.1 as

$$f\,(\boldsymbol{r}, \boldsymbol{p}) = \frac{1}{\exp\left(\beta\left(\frac{\boldsymbol{p}^2}{2m} + V(\boldsymbol{r}) - \mu\right)\right) + 1} \tag{2.19}$$

using the abbreviation $\beta = 1/k_{\mathrm{B}}T$. *in-situ* density and momentum distribution of a single spin state are then obtained by integrating $f\,(\boldsymbol{r}, \boldsymbol{p})$ over $\boldsymbol{p}$ or $\boldsymbol{r}$ respectively, resulting in

$$n(\boldsymbol{r}) = \int \frac{\mathrm{d}^3\boldsymbol{p}}{(2\pi\hbar)^3} f(\boldsymbol{r}, \boldsymbol{p}) = -\frac{1}{\lambda_{\mathrm{dB}}^3}\mathrm{Li}_{3/2}\left(-e^{\beta(\mu - V(\boldsymbol{r}))}\right) \tag{2.20}$$

and

$$n(\boldsymbol{p}) = \int \frac{\mathrm{d}^3\boldsymbol{r}}{(2\pi\hbar)^3} f(\boldsymbol{r}, \boldsymbol{p}) = -\frac{1}{\left(m\lambda_{\mathrm{dB}}\bar{\omega}\right)^3}\mathrm{Li}_{3/2}\left(-e^{\beta(\mu - \boldsymbol{p}^2/2m)}\right) \tag{2.21}$$

with the thermal de Broglie length $\lambda_{\mathrm{dB}} = h/\sqrt{2\pi m k_{\mathrm{B}} T}$ and the polylogarithm function $\mathrm{Li}_n(z)$. The former equation is valid for arbitrary external potentials while the latter only is for harmonic traps. Neither of the two observable distributions reveals the Fermi edge present in the energy distribution 2.19, because a part of the information is erased in the process of observation (integration). The resulting distributions are in shape often similar to Gaussian distributions due to $\lim_{z\to 0}\left[-\mathrm{Li}_{3/2}(-z)\right] = z$.

2.5.2 Local-density approximation

Most of the physical properties of the Fermi gas are connected to the density of states and this is different for the trapped Fermi gas compared to the homogeneous gas. The trapped gas therefore has quantitatively different properties than the homogeneous case. This can already be seen in the temperature dependence of the chemical potential for the non-interacting case in figure 2.1. Qualitatively, μ decreases with increasing T in either case, but the actual values are quite different.

Analogously, the effects of interactions will be qualitatively very similar to the homogeneous case but quantitatively different. A striking example is the critical temperature in the BEC limit. It is given by [76]

$$T_{\mathrm{c}} = \hbar\bar{\omega} k_{\mathrm{B}}^{-1} \left(\frac{N}{2\,\zeta(3)}\right)^{1/3} = \frac{T_{\mathrm{F}}}{(6\,\zeta(3))^{1/3}} \approx 0.518\; T_{\mathrm{F}} \tag{2.22}$$

which is more than twice the value of the homogeneous case. Weak repulsive interactions between the dimers will reduce this temperature [77], in contrast to the homogeneous case.

In general, a combined theoretical treatment of the external potential and the interactions is extremely challenging and no reliable, rigorous calculations are available for the BCS-BEC crossover. A common way to disentangle the interactions from the external potential is the *local density approximation* (LDA). Like the Thomas-Fermi approximation, it makes use of the fact that the level spacing in the harmonic trap $\hbar\omega_i$ is much smaller than the relevant energies, which implies that the density and related quantities like the superfluid density change on length-scales much larger than the interparticle spacing [78]. The gas can then considered to be also *locally* in thermal equilibrium with density $n(\boldsymbol{r})$ and behaves locally like a homogeneous gas. The varying density in the trap is described by modifying the chemical potential with the external potential

$$\mu \;\rightarrow\; \mu - V(\boldsymbol{r}), \tag{2.23}$$

which can be considered as a formal definition of the local density approximation. The local density approximation can provide a prediction of the behaviour, for example condensation, of confined gases based on the (much simpler) homogeneous theory. Similarly, the comparison of observations in confined gases to homogeneous theories is possible.

An open question is the applicability of the local density approximation to non-equilibrium phenomena like the Higgs-mode: is there one common Higgs mode of the inhomogeneous gas or are there many (largely) independent local ones? This question will be investigated in chapter 6.

CHAPTER 3

A Quantum Simulator of the BCS-BEC Crossover

This chapter introduces the apparatus used to conduct the experiments presented in this thesis: an analogue quantum simulator of the BCS-BEC crossover using ultracold ^{6}Li atoms. The setup was conceived and constructed by Alexandra Behrle, Timothy Harrison, and Kuiyi Gao. It is presented in detail in the first two constructors' PhD theses [79, 80]. Here, the relevant aspects of the experiment in connection to the work presented in this thesis are discussed.

First, the mechanism that enables the simulation of arbitrary interaction strengths is discussed. Then follows an introduction of ^{6}Li, the atom used for the implementation of the BCS-BEC crossover. Afterwards, the process of cooling and preparing the atoms for the experiments is briefly discussed. For the studies of the Higgs mode, a new excitation capability was introduced into the setup and this is discussed in more detail. An important aspect is the extraction of results from the experiment. This is done by absorption imaging of the final sample and the intricacies of this process deserve some attention.

3.1 Feshbach resonances

Arguably the most important ingredient for a simulation of the BCS-BEC crossover is a well-controlled attraction between fermions of opposite spin. Dilute gases of ultracold atoms provide an extremely useful feature in this regard: their scattering properties can be easily manipulated by *Feshbach resonances*. The interatomic potential governing scattering is dominated by the van der Waals interaction. This is due to the large particle separations in dilute gases. The exact value of the potential then depends on the relative orientation of the spins of the scatterers' valence electrons[1]. At low temperatures, only the lower-energy configuration with the spins of the valence electrons of the scatterers in singlet configuration is accessible, this is called the "open channel" for scattering. The "closed channel" with electronic triplet configuration is not accessible due to its high energy at large distances. It may however support a bound state with an energy that is similar to that of the scattering

[1] Here, the discussion is limited to the most simple case of ground state Alkali atoms which is the relevant one for this thesis. The valence electrons therefore have $S = 1/2$ and $L = 0$ quantum numbers. A more general discussion can be found in [81].

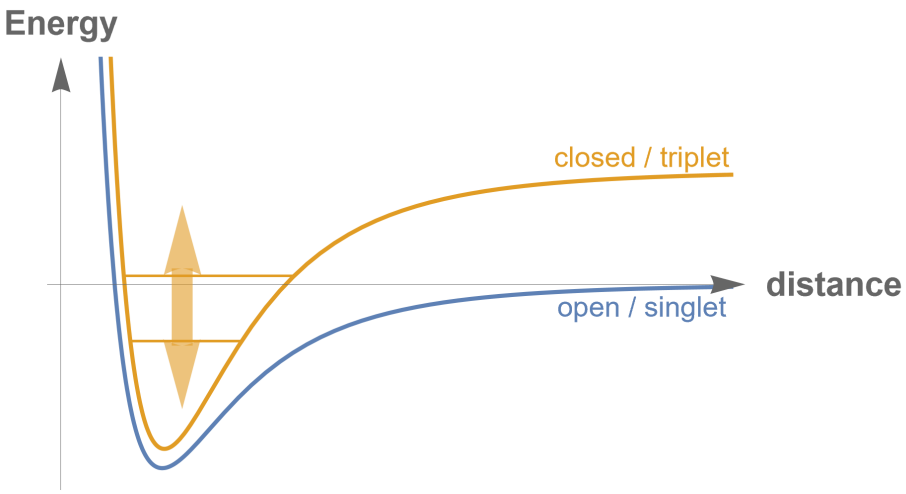

Figure 3.1: Principle of a Feshbach resonance. An external magnetic field can shift the bound states of the closed channel into resonance with the large distance energy of the open channel. This greatly enhances the scattering length.

atoms. This energy is chosen to be 0, the large distance limit of the open channel, see figure 3.1. The hyperfine interaction within the individual atoms mixes the electron's spin projections and therefore leads to a coupling between the two channels which can be strong even if the hyperfine interaction is weak [81]. This results in a strongly enhanced scattering between the atoms. Additionally, a two-body bound state becomes accessible for the scatterers if the bound state energy is slightly negative. Changing the energy difference between the open channel and the closed channel's bound state in a controlled manner is relatively simple because singlet and triplet configuration have different magnetic moments: an external magnetic field is therefore all that is needed to move the channels relatively to each other and shift the bound state in and out of resonance with the large distance energy limit of the open channel. As a consequence, the scattering length a changes with magnetic field strength B. Around the resonance, the relation can be given as the simple expression

$$a(B) = a_{\mathrm{bg}} \left(1 - \frac{\Delta}{B - B_0}\right) \tag{3.1}$$

where a_{bg} is a background scattering length, B_0 is the field at which the resonance occurs and Δ is the width of the resonance. The values of these parameters depend on atomic species, isotopes and even the specific spin configuration of the scatterers. From equation 3.1, it becomes evident that it is possible to tune the scattering length to arbitrary values. This is exactly what is needed for a simulation of the BCS-BEC crossover. It is important that equation 3.1 is only a local approximation of the scattering length around the resonance B_0. The full dependence is significantly harder to evaluate but is necessary to identify, for example, the zero-crossing of the scattering length.

The Feshbach resonance between ^{6}Li atoms in the ground state is characterised by an extremely strong coupling between open and closed channel. Strong in this context means that the energy associated with the coupling greatly exceeds the Fermi energy. This in turn leads to large values for a_{bg} and Δ. It can be shown that in such a case the Feshbach resonance is an almost perfect realisation of

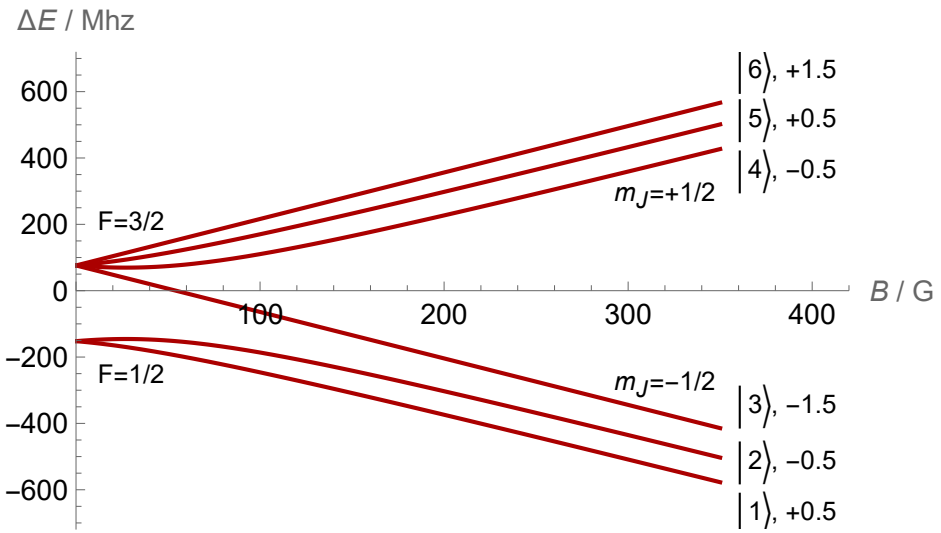

Figure 3.2: Hyperfine structure of the ^{6}Li ground state. The states are labelled with increasing energy and their respective quantum number m_F. In the low and high field regime, other good quantum numbers are shown. Figure adapted from [84]

the BCS-BEC crossover and the details of open and closed channel become irrelevant [31]. Therefore, ^{6}Li is the isotope of choice for our and many other quantum simulators of the BCS-BEC crossover [40, 42, 44, 82, 83].

3.2 Properties of ^{6}Li

Level structure

^{6}Li is an isotope of Lithium with nuclear spin $I = 1$. With a single valence electron, the ground state splits into hyperfine manifolds with total angular momentum $F = 1/2$ and $F = 3/2$. This makes ^{6}Li atoms fermionic particles, just as needed for a simulation of the BCS-BEC crossover. The hyperfine states are typically not labelled by their quantum numbers but are numbered with increasing energy from $|1\rangle$ to $|6\rangle$. This is due to the fact that the Feshbach resonance lies in the strong-field regime of the hyperfine interaction, see figure 3.2. A detailed derivation of the energy levels, together with the relevant isotope-specific constants can be found in [84].

Feshbach parameters

The scattering length is given by equation 3.1 with parameters depending on the hyperfine states chosen for the two fermionic states $|\uparrow\rangle$ and $|\downarrow\rangle$. In this thesis, a combination of $|1\rangle$ and $|2\rangle$ will be denoted as $|12\rangle$ and analogously for the other states. The parameters of the Feshbach resonance between several combinations of hyperfine states are listed in table 3.1.

The parametrisation of the Feshbach resonance from equation 3.1 is only a local approximation of the scattering length around the resonance B_0. The full dependence, shown in figure 3.3, is significantly harder to evaluate to evaluate but is necessary to identify, for example, the zero crossing

Table 3.1: Feshbach resonance parameters for the three lowest hyperfine states of ^{6}Li. Values as reported in [85][2].

states	B_0/G	Δ/G	a_{bg}/a_0
$\|12\rangle$	832.18 ± 0.08	−262.3 ± 0.3	−1582 ± 1
$\|13\rangle$	689.68 ± 0.08	−116.6 ± 0.3	−1770 ± 5
$\|23\rangle$	809.76 ± 0.05	−200.2 ± 0.5	−1642 ± 5

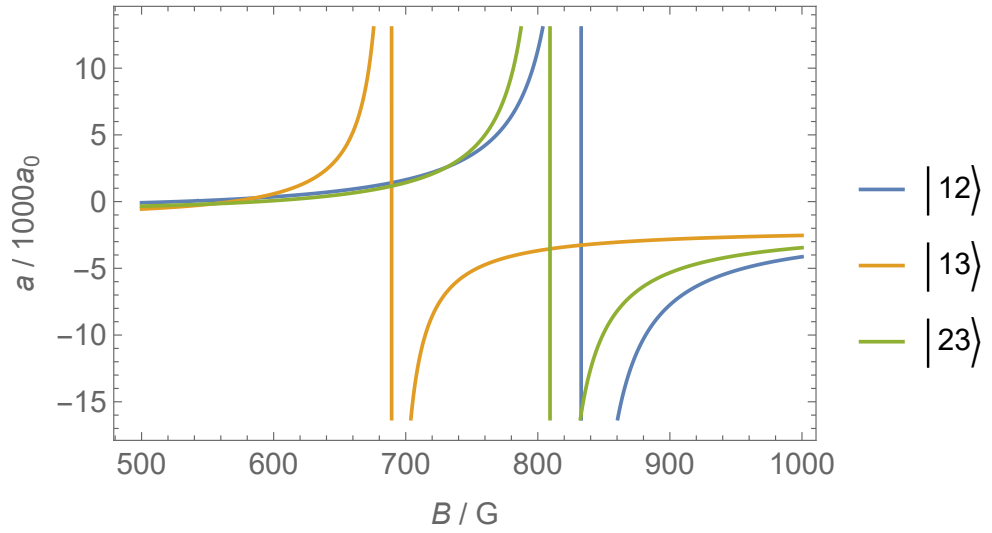

Figure 3.3: Full dependence of the scattering length between the lowest hyperfine states of ^{6}Li on the magnetic field, data from [85].

of the scattering length.

The precise form of the interaction potential as sketched in figure 3.1 is not relevant for the scattering of ^{6}Li because the associated range of the van der Waals potential, $R_{vdW} \approx 50\,a_0$ [86], is much smaller than the thermal deBroglie wavelength or the Fermi wavelength of the ultracold sample, which are both around $\sim 2000\,a_0$, with the Bohr radius a_0. Consequently, the interaction can be fully described by the scattering length alone implying that the interactions are universal, i.e. do not depend on the specifics of ^{6}Li.

All these properties combined make ^{6}Li an ideal atom to simulate the BCS-BEC crossover.

3.3 Preparation of quantum-degenerate ^{6}Li

In order to simulate the BCS-BEC crossover, an ultracold gas of ^{6}Li has to be prepared at temperatures below $T/T_F = 0.1$. This is only possible under ultra-high vacuum conditions and a 3D CAD model of the vacuum chamber is shown in figure 3.4. In our setup, we also create a cold gas of ^{23}Na. This is not strictly necessary to prepare an ultracold Li sample but greatly helps in one of the cooling steps, as

[2] The main body of [85] has a typo in the value for $\Delta_{|13\rangle}$. Here, the correct value, as deduced from the same reference's supplementary material, is listed.

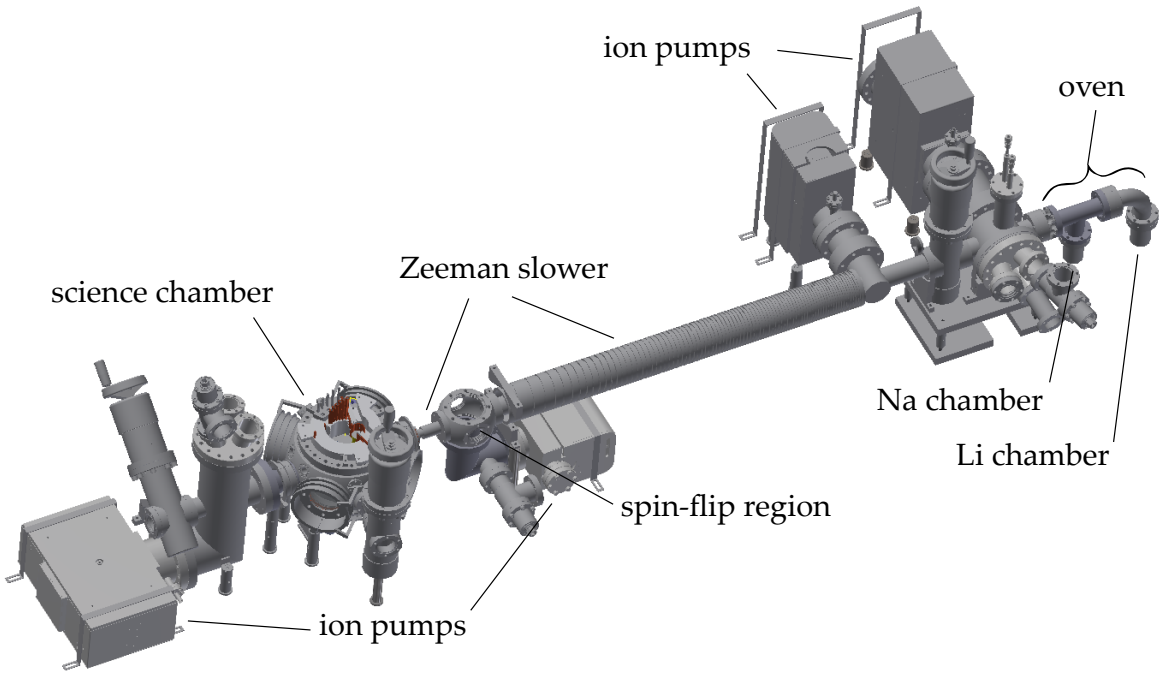

Figure 3.4: Sketch of the vacuum system. The atomic gas is created in the oven chamber at the right end. It travels through the Zeeman slower to the science chamber where it is trapped and cooled to quantum degeneracy, where the experiments are performed. Ion pumps maintain the vacuum in the various parts of the chamber with the pressure decreasing from oven to science chamber by three to four orders of magnitude.

explained below.

A full cycle of preparation, performing of the experiment, data read out, and cooldown of the magnetic field coils last around 25 s. In every cycle, the read out by imaging destroys the superfluid sample and the full process has to be started anew.

Almost all steps in the preparation require various forms of laser light, often with very specific characteristics. This light is generated in several homebuilt laser setups. An overview of these is provided in appendix A.

Source of atomic gases

The oven is used to form atomic gases of ^{6}Li and ^{23}Na by heating pure isotope samples in two separate chambers with individually controllable temperatures. These are usually 660 K for Lithium and 600 K for Sodium. A small hole allows Li to enter the Na chamber so that a mixture is created. A series of apertures is then used to form a collimated beam directed towards the science chamber. This beam contains both Lithium and Sodium atoms whose ratio can be controlled by varying the temperature of the respective chambers. More details on the oven can be found in [87].

Zeeman slower

Both elements in the beam are slowed down in a spin-flip Zeeman slower which cools the hot beams down to the mK regime, below the capture velocity of a magneto-optical trap (MOT). The transition

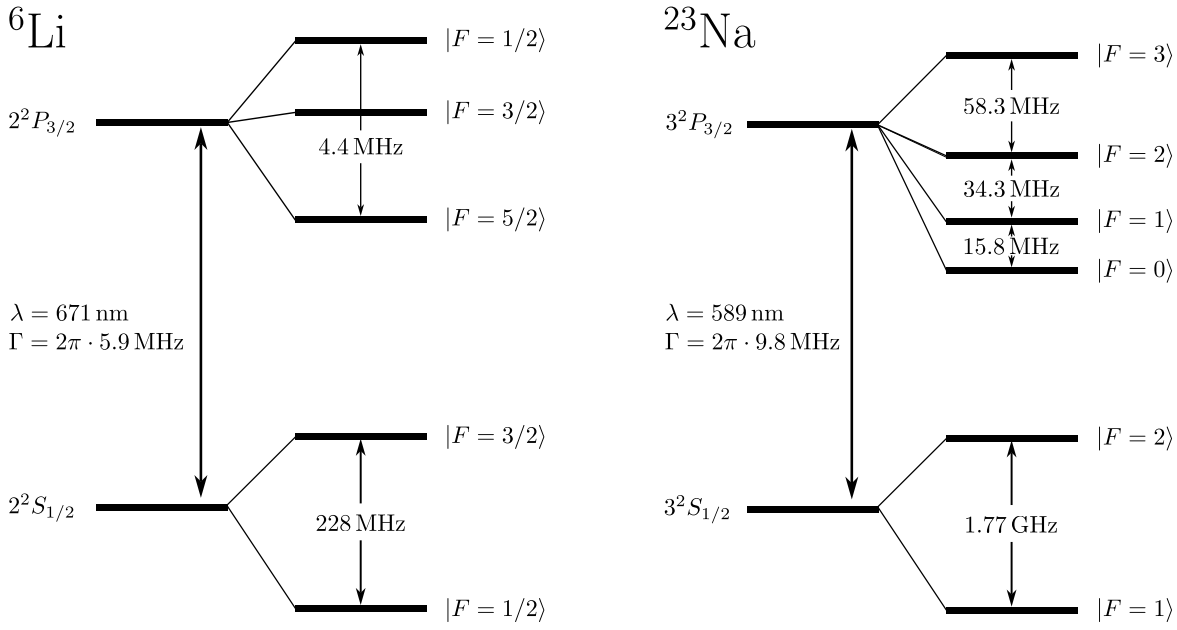

Figure 3.5: Hyperfine structure of the D_2-line for ^{6}Li and ^{23}Na without external fields. Figure not to scale and adapted from [84] for Li and [88] for Na.

used for cooling is the D_2-σ^+-transition into the maximally stretched spin state for both elements and repumping light is used to keep the atoms in the cycling transition. The level-scheme for these transitions is shown in figure 3.5. The two parts of the Zeeman slower also serve as differential pumping stages bridging the pressure difference from the oven chamber at 10^{-7} mbar to the science chamber below 10^{-10} mbar.

Magneto-optical trap

Both elements are caught in magneto-optical traps (MOTs) which are superimposed onto each other, but not of identical structure. While the Li-MOT corresponds more or less to the textbook example of a MOT with circularly polarised beams containing a small admixture of repumping light, the Na-MOT uses a slightly modified scheme: the dark-SPOT (dark spontaneous-force optical trap) [89]. In this scheme, the trapping beams are unaltered but the repumping light has a distinct spatial structure; it consists of a ring-shaped beam with a dark centre created by imaging a black dot onto the atoms. The repumping light therefore only affects atoms at the edge of the trapping region, i.e. those who are about to escape the MOT, and returns these back to cycling, trapping transition.

Inside the cloud, Na atoms falling into a dark state will stay in it for a longer time. because they are not affected by the rempumping light. This is beneficial for two reasons: Firstly, photons can be scattered by multiple ground-state atoms creating an effective repulsion between these atoms which limits the atom number in the MOT; secondly, collisions between ground and excited state atoms can release enough energy to expel atoms from the MOT which limits the density [89]. In the dark-SPOT, both processes are strongly suppressed, because dark state atoms do not scatter light and the mean excited state population is strongly reduced. The gain compared to a standard, "bright" MOT is a density increase by a factor ~ 100 [79].

The atomic flux rates from the oven and the two MOTs' parameters are chosen such that the ratio of Na to Li atoms is roughly 4:1 – around $2 \cdot 10^{10}$ Na atoms and $5 \cdot 10^9$ Li atoms.

Further laser cooling

After the MOTs have been completely loaded, the Li MOT is compressed by reducing detuning and power of the MOT beams. The compression reduces the phase space density of the Li atoms and therefore cools the gas. The compressed MOT in its final configuration is not used for loading because its capture velocity is much lower.

In the MOT, cooling is achieved by photon scattering. In this process, the photon is emitted into a random spatial direction and this randomness fundamentally limits the temperature in the MOT from below. The limit is given by the Doppler-temperature [90]

$$T_\mathrm{D} = \frac{1}{2}\frac{\hbar\Gamma}{k_\mathrm{B}} \tag{3.2}$$

with the natural line-width Γ of the cooling transition. This limit can be overcome by polarisation gradient cooling as happening in an optical molasses [91]. Such a molasses is implemented for Na by switching off the magnetic quadrupole field of the MOT while keeping the MOT-beams on, now with repumping light in all beams. For Li, an optical molasses cannot be performed because the excited hyperfine splitting is smaller than the natural line-width and therefore not resolvable. Instead, Li expands freely during the molasses phase and its expansion limits the molasses time to around 2 ms.

Loading into a magnetic trap

For further cooling, the atoms are loaded into a magnetic quadrupole trap created by magnetic field coils in anti-Helmholtz configuration. If the atoms are in a low-field seeking state, they are confined to the centre of the coils because the in all spatial directions linearly increasing field acts as a trapping potential. Optical pumping is performed after the molasses to transfer the atoms into the appropriate states; for Na this is the state $|F = 2, m_F = 2\rangle$, see figure 3.6, and for Li the state $|6\rangle$. These are particularly well suited because they both have a maximally stretched angular momentum which can therefore not change during collisions.

To load into the magnetic trap, the MOT coils are used to create the necessary magnetic field. This serves however only as a preliminary magnetic trap. By an elaborate procedure, described in detail in [80], the minimum of the magnetic trap potential is shifted from the MOT coils to the centre of the Feshbach coils, c.f. figure 3.7a). Several advantages are connected to the new position: Firstly, it is not located in the atomic beam emitted from the oven, which runs continuously[3] and would bombard the atoms in the magnetic trap. Secondly, the new position is significantly closer to the viewport of the vacuum chamber, so that it becomes much easier to interface the atoms with external apparatus. This feature is exploited to perform interaction quenches as described in section 3.4. Finally, the Feshbach

[3] A mechanical shutter to block the atomic beam is included in the design of the oven chamber but a mechanical defect inhibited its use [79]. Recently, a new, improved oven chamber, described in [87] and with a new shutter design, has been installed. The shutter has however not been tested so far.

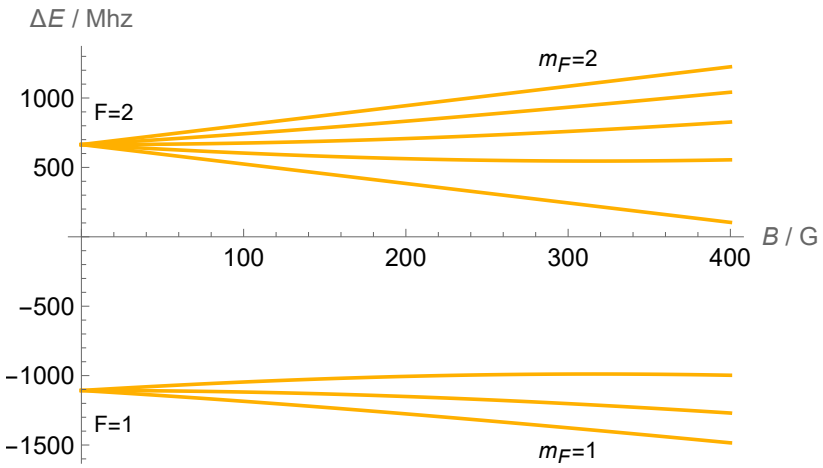

Figure 3.6: Hyperfine structure of the ^{23}Na ground state. Only the low field regime with the good quantum numbers $|F, m_F\rangle$ is relevant for the experiment. Figure adapted from [88].

coils are optimised to create strong magnetic fields or gradients over an extended amount of time, mainly by having a significantly better cooling system than the MOT coils.

Evaporative cooling in the magnetic trap

In the magnetic trap, the temperature of the cloud can be lowered by selectively removing the atoms carrying the highest amount of energy, a process called evaporative cooling. If the cloud is given enough time to re-thermalise after the removal, the average energy per atom, and therefore also the cloud's temperature, is decreased. The removal happens by driving the RF-transition from the Na atoms' $|F = 2, m_F = 2\rangle$ state to the $|F = 1, m_F = 1\rangle$ state – a high-field seeking state that is quickly ejected from the trap. The frequency of this transition increases with magnetic field strength as evident from figure 3.6. By starting the driving with a relatively high frequency, only the atoms at very high magnetic field strengths are transferred to the high-field seeking state. At high magnetic fields, the atoms in the $|F = 2, m_F = 2\rangle$ state carry a lot of potential energy and therefore only atoms carrying a lot of energy are emitted from the trap. The cooling cloud will quickly have very few atoms left resonant to the drive frequency and hence it becomes necessary to slowly lower the frequency to keep the cooling process running. In practice, the RF-frequency is lowered from 1 890 MHz to 1 772.5 MHz within 5 s. These settings have been found empirically by optimising for the coldest cloud of Li atoms at the end of an experimental cycle. The initial (final) evaporation frequency is resonant to the $|F = 2, m_F = 2\rangle \rightarrow |F = 1, m_F = 1\rangle$ transition at a magnetic field strength of 55.2 G (0.4 G). The low final magnetic field strength is reached very closely to the centre of the magnetic trap quadrupole field. This fits to the observation that all Na atoms are evaporated from the trap in this process.

The Li atoms are not affected by the RF-field but are still cooled sympathetically with the Na atoms. This is one advantage of preparing a mixture of Na and Li atoms: no trade-off between a high atom number and low temperature arises in this cooling step because only the (expandable) Na atoms are evaporated. An additional obstacle that is overcome by the addition of Na is the suppressed scattering

of identical fermions as explained in section 2.2. With all Li atoms in the $|6\rangle$ state, thermalisation of a pure Li cloud would be extremely inefficient because only collision with finite angular momentum (which are suppressed at low temperatures) are possible. The addition of bosonic Na, which interacts with Li, solves this problem and ensures efficient cooling.

The distinction between high- and low-field seeking states is crucial for the operation of the magnetic trap. Unwanted transitions between these states within the $F = 2$ manifold can occur in the centre of the trap where the magnetic field vanishes and thus no quantisation axis exists. These losses, called Majorana-losses, are suppressed by shining a repulsive laser beam ("plug laser") into the centre of the trap to expel all atoms from the loss region. The laser beam has a wavelength of 532 nm, blue-detuned to both the Na and Li D-line transitions, and a power of several watts.

Further evaporation in an optical dipole trap

For an effective simulation of the BCS-BEC crossover, the atoms need to be trapped independent of their internal state. To this end, they are transferred from the magnetic trap into an optical dipole trap where they are also cooled down to quantum degeneracy. The optical dipole trap is formed by two intersecting laser beams with a wavelength of 1 070 nm; no interference pattern is created by the two beams due to the laser's extremely short coherence length.

In order to overcome the Pauli-suppression of scattering between Li atoms, which limits thermalisation during the cooling procedure, a spin-mixture has to be created. For the experiments presented in this thesis, this is either a $|12\rangle$ or a $|13\rangle$ mixture. A first Landau-Zener sweep with a small homogeneous background magnetic field transfers all Li atoms from the $|6\rangle$ state to the $|1\rangle$ state. A second Landau-Zener sweep is then performed to bring all atoms into the coherent superposition $\frac{1}{\sqrt{2}}(|1\rangle + |2\rangle)$. If the desired mixture is $|13\rangle$, a third Landau-Zener sweep transforms the $|2\rangle$ part of the superposition into the $|3\rangle$ state. The Pauli-suppression of scattering still applies to the coherent superpositions so that the evaporation cannot start until the superposition has decohered. Complete decoherence is achieved by ramping the magnetic field, created by the Feshbach coils in Helmholtz-configuration, over $100 - 130$ ms close to the value B_0 of the Feshbach resonance of the respective mixture. The residual spatial inhomogeneity of the Feshbach field introduces random phases into the superposition state and therefore leads to full decoherence.

The decohered mixture close to the Feshbach resonance has two very strongly interacting components making evaporative cooling very efficient. In the dipole trap, evaporative cooling is performed by reducing the power of the laser beams over roughly 1 s. This reduces the trap depth continuously and allows the most energetic atoms to escape the trap thus cooling the gas. In the end, the cloud will be at a temperature around a few percent of the Fermi temperature and trapped in a potential $V(\boldsymbol{r})$ given by the laser beams' intensity distribution $I(\boldsymbol{r})$ [90]

$$V(\boldsymbol{r}) = \frac{\hbar}{8}\frac{\Gamma^2}{\Delta}\frac{I(\boldsymbol{r})}{I_{\mathrm{sat}}}. \tag{3.3}$$

Here, $\Delta = \omega - \omega_0$ denotes the detuning of the trapping laser's frequency ω from the atomic resonance ω_0 and I_{sat} the atoms' saturation intensity, for Li this is $7.59\,\mathrm{mW/cm^2}$ for the D_2- and $2.54\,\mathrm{mW/cm^2}$

for the D_1-line [84]. $I(\boldsymbol{r})$ is the sum of the two laser beams' intensity. Both have a Gaussian profile given by

$$I(x, y, z) = I_0 \exp\left(-\frac{2x^2}{w_x^2} - \frac{2y^2}{w_y^2}\right) \tag{3.4}$$

with z the spatial coordinate in the beam's propagation direction, x, y the spatial coordinates perpendicular to the propagation direction, I_0 the maximum intensity of the beam, and $w_{x,y}$ the two waists of the beam at position z. The divergence of the beam along z is too small to be relevant for the trapping potential experienced by the atoms. The full intensity distribution is then given by the sum of the Gaussian profiles with perpendicular propagation direction.

Often, the full shape of the potential as given by equations 3.3 and 3.4 is not necessary for an appropriate description, because the atoms accumulate in the potential minimum at the centre of the intersecting beams. The potential can then be approximated as harmonic:

$$V(x, y, z) = \frac{m}{2}\left(\omega_x^2 x^2 + \omega_y^2 y^2 + \omega_z^2 z^2\right) \tag{3.5}$$

with trap frequencies ω_i obtained from a Taylor series of equation 3.3. The frequencies therefore have a square-root dependence on the intensity or power of the beams. In the experiment, the three trap frequencies usually are in the range of $2\pi \cdot 80\,\mathrm{Hz}$ to $2\pi \cdot 280\,\mathrm{Hz}$.

Setting the interaction strength

After the evaporation is finished, the desired interaction strength of the two-component Fermi gas is set by adjusting the scattering length a via its magnetic field dependence 3.1. To disturb the gas as little as possible, the field is swept adiabatically from the evaporation setting to the final setting. With this step, the preparation is completed and the actual experiments as presented in the later sections can performed.

3.4 Fast manipulations of the interaction strength

In order to study the Higgs mode, the interaction strength has to be modified. This is most conveniently done by manipulating the magnetic field strength B around the Feshbach resonance in order to change a as given in 3.1. For the two excitation schemes presented in section 2.3, either a substantial change in B on a time scale shorter than the Fermi time, which is usually slightly below $10\,\mu\mathrm{s}$, or a modulation of B with a frequency around the Fermi frequency, usually between $10\,\mathrm{kHz}$ and $20\,\mathrm{kHz}$, have to be realised. The former realises a quench of the interaction strength while the latter corresponds to a modulation. Both scenarios are technically challenging, because the magnetic field has to be changed on very short timescales. The large Feshbach coils which set the interaction strength have a large self-inductance and mutual inductance with the rest of the setup thus inhibiting any rapid changes in magnetic field strength.

In order to overcome this limitation, a new, compact coil design, the quench coil, has been developed in the group shortly before the work on this thesis began. It is presented in [92]. The key idea is to insert an additional coil into the setup as shown in figure 3.7 a) to create a magnetic field that is either added to or subtracted from the Feshbach field. A single coil located outside the vacuum chamber

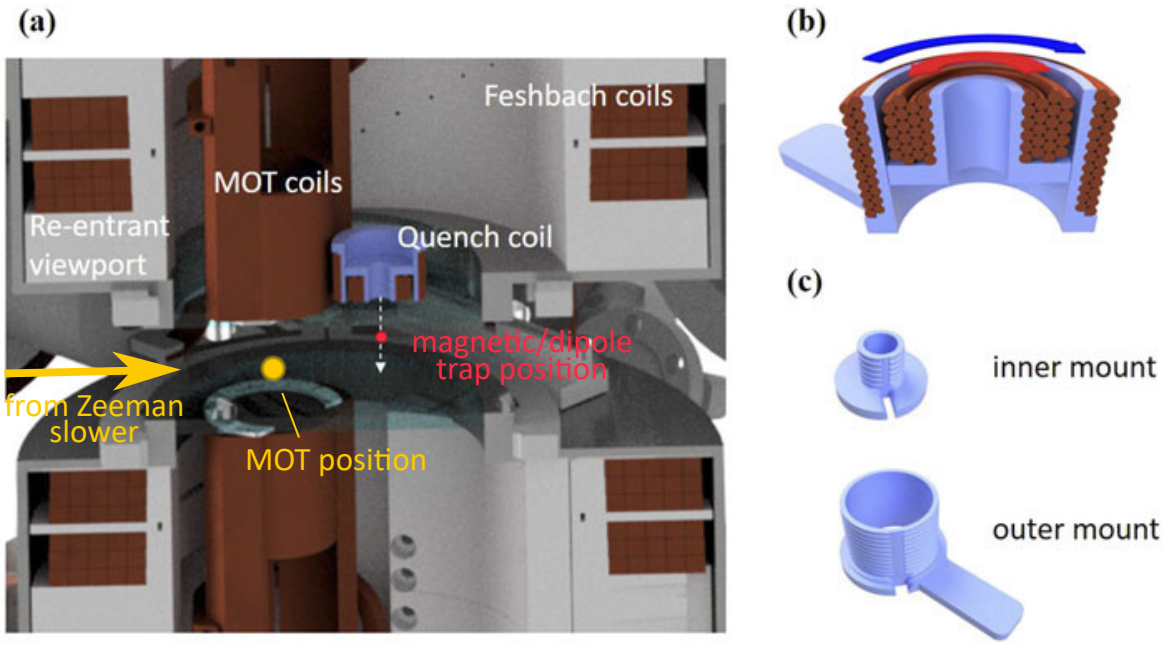

Figure 3.7: a) View into the science chamber. The MOT is located in the centre of the MOT coils which is overlapped with the atomic beam from the Zeeman slower. The magnetic and dipole traps, located in the centre of the Feshbach coils, are closer to the upper viewport's window such that the atoms are closer to external apparatus like the quench coil or imaging optics (not shown). b) Internal structure of the quench coil, it comprises of two coils with opposite winding whose gradients cancel at the atom position. c) The coils are wound onto 3D-printed plastic mounts. Image adapted from [92]

would always create a gradient at the atoms' position; a scenario that should be avoided because it leads to a force that could pull the atoms out of the trap. Usually, such a gradient is avoided by the placement of two coils in Helmholtz configuration but due to the large distance from the atoms to the lower viewport, relatively large coils with higher inductance would be needed making this approach impractical. Instead, the quench coil is made up out of two slightly different coils connected in series and inserted into each other through which the current runs in opposite directions, see figure 3.7 b). In this way, there is a point along the coil axis where the gradients cancel and this is the intended position for the atoms. Because the coils are of different size, there is a finite field strength at this point.

3.4.1 Quench coil characteristics

The quench coil's field was characterised by moving it around[4] and measuring the magnetic field strength at the atoms' position. They can be used to probe the field at their location by the magnetic field dependence of the $|1\rangle \rightarrow |2\rangle$ RF-transition frequency. This measurement is presented in detail in [93] and reveals a linear magnetic field strength dependence on the current with a gradient of (0.874 ± 0.008) G/A. The maximal current of 40 A then leads to a highest achievable field strength of (35.0 ± 0.3) G. Since the field is added to the much stronger homogeneous background field, the quench coil does effectively not change the direction of the total $\boldsymbol{B}$ field but rather just its magnitude.

[4] The quench coil is mounted on a stack of three translation stages.

In practice, it was found that the quench coil could not be moved close enough to the atoms to overlap the zero-gradient point with the atoms. Instead a relatively small gradient remains, at a field strength of 21.5 G it was measured to be (1.6 ± 0.5) G/cm. In principle, a slightly modified coil that would trade field strength for a further displaced field maximum could be built, but because it is simple to cancel the remaining gradient with spare layers in the Feshbach coils, the present coil was considered to be sufficient.

3.4.2 Alignment

The most thorough alignment process is the calibration procedure described above, which, however, takes a considerable amount of time. In practice, different methods are used and the most effective has proven to be the following: The atoms are released from the dipole trap with the quench coil field still on. If the atoms are located on the quench coil's symmetry axis, their movement perpendicular to this axis should be identical to the case without quench coil field. The biggest advantage of this method is that it immediately indicates the direction of misalignment making the process very efficient. With increasing current in the quench coil, the precision of the alignment can be increased.

3.4.3 Current control

Depending on the intended action, the current through the quench coil is controlled differently. For quenches, a control circuit designed by Dr. Akos Hoffmann and discussed in [92] is used. Its main feature is to enable the fast dissipation of the magnetic field energy in an RC snubber circuit. A high-speed power MOSFET driver and a power MOSFET allow shutting the current through the coil off with a switch time down to 30 ns. In the end, the switch time of the magnetic field is the relevant quantity. This has been measured with a pick-up coil to be 2.6 µs after insertion of the quench coil into the experimental setup. It is therefore, as desired, faster than the Fermi time $\tau_F = \hbar/E_F$ which typically has a value around 7 µs.

If a modulation of the magnetic field strength is desired, the control circuit is replaced by an amplifier circuit also designed by Dr. Akos Hoffmann. It was built to drive sinusoidal currents with an amplitude of up to 10 A (20 A peak-to-peak) and 40 kHz through the quench coil with the signal sourced from a programmable arbitrary function generator[5]. A characterisation of the circuit can be found in [93].

3.4.4 Temperature control

Because the quench coil is not actively cooled, it has to be protect from overheating by other means. An additional importance is given to this aspect by the fact that the quench coil mount consists of 3D-printed plastic and epoxy resin which could both melt and contaminate the vacuum chamber's viewport. The temperature of the quench coil is therefore measured with two Pt100 sensors and monitored with an interlock circuit. On top of that, both the quench control circuit and the modulation amplifier circuit include protective elements that limit the on-time of the quench coil to around 1 s. During quench measurements, the quench coil is typically switched on for 0.5 s per 25 s experiment

[5] AFG-2225 by GW Instek

cycle time. The temperature increases to up to 25 °C above room temperature under these conditions and then oscillates by around 1 °C during the experimental cycle. It has been observed that the initial warm-up slightly displaces the magnetic field generated by the quench coil. Alignment of and experiments with the quench coil are therefore always performed after the quench coils has fully thermalised.

3.5 Imaging of the atoms

3.5.1 General procedure

After the desired experiment has been performed, the sample has to be probed and this is done by absorption imaging, often with high intensities. Resonant light is shone onto the atoms and the shadow cast by photon scattering, usually called absorption in this context, is imaged onto a camera. To get a good absorption signal, a cycling transition of the D_2-line for one of the final hyperfine states is chosen. The other state is not resonant to this light and therefore not imaged but that state would create an identical absorption image since all experiments of this thesis treat both spin states identically. The recorded absorption pattern can be converted into column densities ρ, the cloud's density n integrated along the camera's line of sight, as detailed below. To get a more complete picture of the cloud's properties, imaging optics and beam paths exist in several perpendicular directions. The preferred imaging direction is along the direction of gravity, which coincides with the quench coil axis and is defined as the z-direction, because the imaging optics can be brought closest to the atoms along this axis, see figure 3.7.

The atoms can either be imaged *in-situ* to record the density distribution of the cloud or after release from the dipole trap and with a varying expansion time, a technique called time-of-flight (TOF) imaging. The release is implemented by rapidly extinguishing the trapping beams. The sudden removal of the external potential leads to an (almost) free movement of the atoms which causes their distribution to be more and more dominated by the momentum distribution, similar to equation 2.21. During the expansion, the atoms move in a potential given by the residual curvature of the external magnetic field. Perpendicular to the z-direction, this potential is harmonic with a frequency of around $\omega_e = 2\pi \cdot 16\,\mathrm{Hz}$ while it is anti-confining in the z-direction. After a quarter period expansion time ($T/4$), the atoms' momentum is converted to potential energy and the new density profile reflects the momentum distribution in the original potential [94]. However, this statement has to be taken with a grain of salt, because the atoms interact strongly in the beginning of the expansion and this influences the final distribution in a way that is hard to predict. Additionally, the very light Lithium atoms move relatively far in $T/4 \approx 15\,\mathrm{ms}$ which causes the final density to be relatively low and hard to detect.

If the atoms are imaged along the z-direction, they fall under the influence of gravity if an expansion time is applied. Two imaging paths therefore exist along this direction, one focused for *in-situ* imaging and one focused for an expansion time of 15 ms. The cameras used for imaging are EMCCD cameras by Andor, model iXon Ultra DV-897E-CSO-#BV for *in-situ* imaging and iXon Ultra DU-888U3-CS0-#BV for expanded clouds.

3.5.2 Prerequisites for modelling absorption imaging

Modelling the shadows observed in absorption imaging is easiest if the atoms are assumed to be stationary during imaging but this is not the case because radiation pressure accelerates the atoms

during illumination. However, the imaging pulse can be chosen short enough for the final image to be unaffected by radiation pressure. Following the discussion in [95], the moving atoms are subject to a Doppler shift causing the imaging light to become off-resonant. The acceleration a experienced by the atoms is given by the photon momentum $\hbar k$ times the scattering rate which, on resonance, leads to the relation

$$a = \frac{\hbar k}{m}\frac{\Gamma}{2}\frac{I}{I + I_{\mathrm{sat}}} \tag{3.6}$$

with the saturation intensity I_{sat}. Requiring the Doppler shift to be smaller than a fraction $f_{\mathrm{Doppler}} < 1$ of the intensity broadened line width gives an upper limit to the imaging pulse length [79]:

$$\tau < f_{\mathrm{Doppler}}\frac{2m}{\hbar k^2}\sqrt{s\left(\frac{1}{s}+1\right)^3} \tag{3.7}$$

with the saturation parameter $s = I/I_{\mathrm{sat}}$.

In principle, it is also possible that the atoms move out of the imaging system's depth of focus. However, since we do not use high resolution optics, the depth of focus is relatively large and the Doppler limit 3.7 dominates in all relevant imaging scenarios.

3.5.3 High-intensity absorption imaging

If a cloud of atoms is irradiated with light of considerable intensity, Beer's law for absorption has to be modified to include saturation effects [96]. The absorption of the cloud is then described by the relation

$$\sigma_0\rho(x,y) \equiv \sigma_0\int n(\boldsymbol{r})\mathrm{d}z = -\alpha\log\left(\frac{I_{\mathrm{f}}(x,y)}{I_{\mathrm{i}}(x,y)}\right) + \frac{I_{\mathrm{i}}(x,y) - I_{\mathrm{f}}(x,y)}{I_{0,\mathrm{sat}}}. \tag{3.8}$$

Here, $\sigma_0 = 3\lambda^2/2\pi$ is the resonant cross section at wavelength λ, ρ the clouds column density, obtained by integrating the full density $n(\boldsymbol{r})$ along the camera's line of sight z, I_{i} (I_{f}) the intensity before (after) the atoms, and $I_{0,\mathrm{sat}}$ the ideal saturation intensity. In reality, the observed saturation intensity is higher than $I_{0,\mathrm{sat}}$ by a factor α; the reasons for $\alpha > 1$ can be varied, from experimental imperfections like incorrectly polarised or off-resonant imaging light to atoms influencing each other. For the latter reason, α usually depends on the column density. Since α is an effective parameter describing imperfections, it cannot be predicted but has to be determined from measurements.

In principle, only I_{f} can be measured and I_{i} would have to be calibrated before the measurement. However, power drifts in the imaging beam make this method rather unreliable and there is a simpler way to determine I_{i}: after the atoms are imaged – the measurement of I_{f} – a time of ca. 17 ms is given for the atoms to disperse and then a second image, now without any absorbing atoms, is recorded to measure I_{i}. This second image is also called the light image and the cameras used for imaging are designed to record such a series of images efficiently.

The cameras do not measure the intensities directly but rather the charges accumulated on each pixel of the camera chip, the counts C. These counts are proportional to the intensity at the pixel's

position after subtraction of the camera's dark counts C_{dark}. Equation 3.8 can then be recast into

$$\rho\sigma_0 = -\alpha \log\left(\frac{C_\mathrm{f}}{C_\mathrm{i}}\right) + \frac{C_\mathrm{i} - C_\mathrm{f}}{C_{0,\mathrm{sat}}\tau/\tau_0} \tag{3.9}$$

for every pixel, with $C_{\mathrm{i/f}} = C_{\mathrm{light/atom}} - C_{\mathrm{dark}}$, $C_{0,\mathrm{sat}}$ the counts accumulated on the pixel for an intensity $I_{0,\mathrm{sat}}$ over a time τ_0, and τ the imaging pulse duration for atom and bright image.

3.5.4 α-calibration

The two terms that determine the column density, found on the right hand side of equation 3.8, scale differently with I_i. This can be exploited to determine α. In the limit $I_\mathrm{i} \ll I_{0,\mathrm{sat}}$ ($I_\mathrm{i} \gg I_{0,\mathrm{sat}}$), the column density is dominated by the first (second) term. The interpolation between these two regimes varies depending on α and can be probed by imaging identically prepared clouds with varying intensities. With the abbreviations

$$\begin{aligned} \lambda &= -\log\left(\frac{C_\mathrm{f}}{C_\mathrm{i}}\right) \\ \delta &= \frac{C_\mathrm{i} - C_\mathrm{f}}{C_{0,\mathrm{sat}}\tau/\tau_0} \end{aligned} \tag{3.10}$$

for the two terms, 3.9 can be rewritten as a function $\delta(\lambda)$

$$\delta = \rho\sigma_0 - \alpha\lambda \tag{3.11}$$

and α is obtainable by a linear fit. In principle, α can be obtained by this method for every pixel of the camera but in practice, pixels with similar column density are binned to determine a common α. Considering all bins together, a relation $\alpha(\rho\sigma_0)$ can be found. It is important to keep in mind that this relation is a property of the cloud under observation and not of the camera. An α-calibration therefore has to be performed for every measurement requiring a certain degree of precision in ρ. An example of this calibration is presented below in section 3.5.7.

3.5.5 Detection limit

Another limitation for imaging, originating from the camera, has been considered in [97]: the image with the atoms needs a sufficient signal to noise ratio. This becomes an issue if the density $\rho\sigma_0$ is large and the intensity $I_\mathrm{i}/(\alpha I_{0,\mathrm{sat}}$ = small. In this case C_f is very small and dominated by noise on the camera image like photon shot noise. Circumventing this issue sets a *lower* bound on the imaging pulse length [97]. The exact value is however difficult to predict a priori because a good knowledge of the cloud's and camera's properties is required. Nevertheless, there are scenarios in which the required pulse length for a good signal to noise ratio could exceed the Doppler limit 3.7. On our cameras, we observe that λ saturates at around 2.3 for regions with high $\rho\sigma_0$ when lowering I_i. We attribute this to the detection limit and discard the saturated images from α-calibrations.

3.5.6 Magnification

To calculate physical quantities like atom number N or column density ρ from the absorption images, the magnification of the different cameras has to be known. Equation 3.9 might give the impression that the magnification is irrelevant because it only consists of ratios of camera counts. However, $C_{0,\mathrm{sat}}$ is calculated from the counts accumulated on the camera from a beam with known power and the conversion to intensity is only possible if the (physical) pixel size and magnification are taken into account. In a low intensity setting, where only the first term of equation 3.9 is relevant, the column density ρ of atoms can indeed be calculated without knowledge of the magnification. Calculation of the atom number N however again requires knowledge of (physical) pixel size and magnification because ρ has to be integrated over the area imaged onto one pixel.

The magnification of the imaging systems is determined with the help of the multiple cameras installed in the experiment. Cameras whose imaging plane contains the direction of gravity can be calibrated by observing the free fall of a cloud of atoms in near-zero magnetic field. The more important cameras for imaging after time-of-flight and *in-situ* have their line of sight in the direction of gravity and therefore cannot be calibrated in this way. Instead, a cross-calibration with the other cameras, already calibrated by the action of gravity, is performed: clouds prepared in identical ways are imaged with different cameras. The real size of the cloud can be determined with an already calibrated camera and then used to calibrate a camera with previously unknown magnification. The resulting magnifications are $M = 4.32$ for the camera recording images after time-of-flight and $M = 7.26$ for *in-situ* images.

3.5.7 Example α-calibration

For the analysis of the dynamics following a quench of the interaction strength (see chapter 6) it is necessary to know the density distribution of the atoms in the dipole trap. Since the column densities can be relatively high in this case, high imaging intensities are necessary and an alpha-calibration is critical for an understanding of the absorption images.

The first quantity necessary for the calibration is the value of $C_{0,\mathrm{sat}}$, the counts recorded on a pixel if it is exposed to a beam with intensity $I_{0,\mathrm{sat}}$ for a time τ_0. To this end, a beam is sent onto the camera without any absorbing atoms present. The beam's power is measured with a power-meter before and after the vacuum chamber to get a good estimate for the power P at the atoms' position. From the total counts on the camera after different illumination times – with the dark counts subtracted –, the counts accumulated per unit time at the power P is found. This allows the identification of a fixed amount of energy with every count. Note that this energy is not the real energy per count because it does not consider the quantum efficiency of the camera or losses in the imaging path, instead, it is the energy at the atoms' position for every count. Using the information on magnification, physical pixel size on the camera chip and the theory value of $I_{0,\mathrm{sat}}$ [84], $C_{0,\mathrm{sat}}$ can be calculated for arbitrary values of τ_0. For $\tau_0 = 1\ \mu\mathrm{s}$, a value of $C_{0,\mathrm{sat}} = 33 \pm 4$ was found in this way for the camera used for *in-situ* imaging.

Some care has to be taken when calculating the pulse length τ used to record the absorption images: due to experimental imperfections it is not necessarily equal to the pulse length set in the experimental sequence. Comparing the counts on the camera for different pulse lengths as done in figure 3.8 indeed reveals that the real pulse length is (410 ± 10) ns shorter than its set value. As a check, this comparison is done for several imaging intensities and the discrepancy is always the same.

With the quantities $C_{0,\mathrm{sat}}$, τ_0, and τ understood, the terms λ and δ from equations 3.10 can be

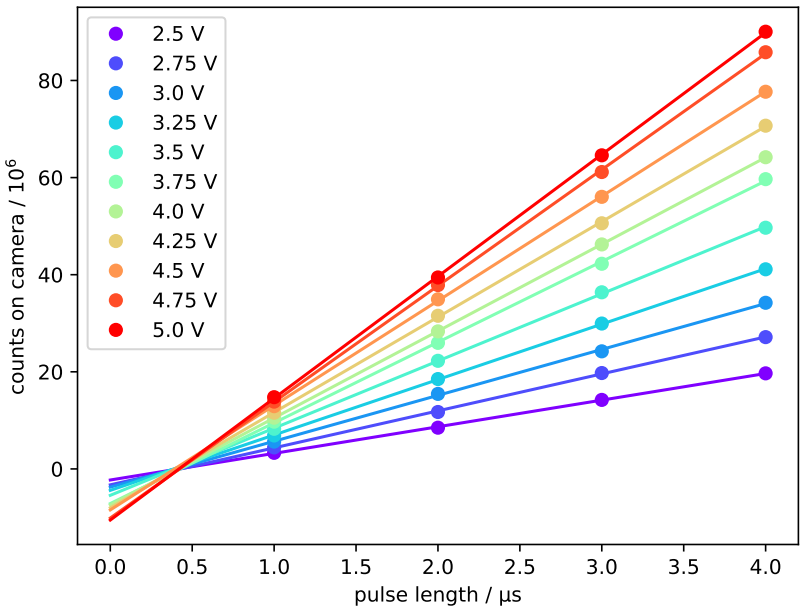

Figure 3.8: Recorded counts on the camera for different pulse lengths and imaging intensities. The intensity is set by a control voltage with a non-linear behaviour. The points are the measured counts on the camera while the lines are the corresponding linear fits. The lines meet at roughly zero counts but at a finite time indicating that the real pulse length is shorter than the set value given on the x-axis.

calculated directly from the atom and light images. If identically prepared clouds are imaged with varying intensity, different combinations of the terms λ and δ are recorded and the scaling of these terms reveals α via equation 3.11. Figure 3.9 shows an example for this procedure where pixels with similar column density have been binned together. This example has been recorded at a Feshbach field of 690 G corresponding to unitarity of the $|13\rangle$-resonance. Because the attractive interactions modify the density distribution of the cloud, identical measurements have been performed further on the BEC and BCS side of the crossover to check for an interaction dependence of α.

When α is plotted against the column density normalised by the cross section $\rho\sigma_0$, the other free parameter of the fits in figure 3.9, a linear relation between the quantities is found, see figure 3.10. This is consistent with the observations in [79, 95] where a linear dependence was found as well. However, both report a region of constant α for low values of $\rho\sigma$ and smaller values of α in general. These discrepancies are most likely connected to the relatively high densities encountered in this measurement.

The behaviour for the different magnetic field or interaction strengths is further analysed in figure 3.11 where gradient and offset of the function $\alpha(\rho\sigma)$ are plotted. The gradient $\frac{\mathrm{d}\alpha}{\mathrm{d}(\rho\sigma)}(B)$ is well described as a linear function of the magnetic field. The offset in contrast seems to be lower on the BEC side compared to unitarity and the BCS side. The difference is however relatively small and choosing a constant value $\alpha_0 \equiv \alpha(\rho\sigma = 0) = 3.0$ will not alter the results by a lot. α can therefore be

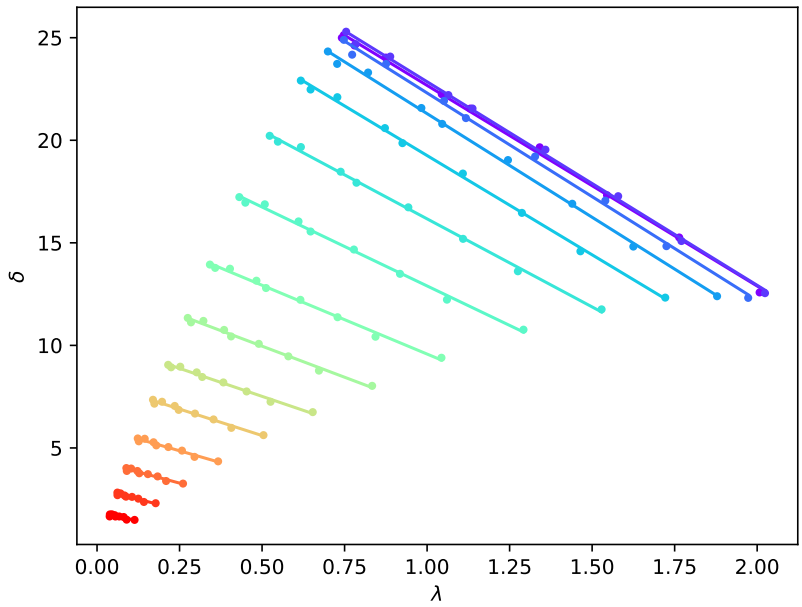

Figure 3.9: Determination of the factor α by a linear fit according to equation 3.11. The different colours indicate bins with similar column density. Bins around $\lambda = 0$, shown in red, are located far away from the centre of the cloud where the column densities are low. The bins with the highest values of λ or δ, shown in blue, are around the centre of the cloud where the column density is highest.

written as a function of $\rho\sigma$ and B as

$$\alpha(\rho\sigma, B) = \frac{\mathrm{d}\alpha}{\mathrm{d}(\rho\sigma)}(B)\ \rho\sigma + \alpha_0. \tag{3.12}$$

Combining this with equation 3.11 enables the elimination of α to calculate the column density from the easily available quantities δ and λ and the magnetic field strength alone:

$$\rho\sigma = \frac{\delta - \alpha_0\lambda}{1 + \frac{\mathrm{d}\alpha}{\mathrm{d}(\rho\sigma)}(B)\lambda} \tag{3.13}$$

It was also found that the calculated column density decreases with the imaging pulse length, most likely due to the Doppler effect from atoms accelerated by light pressure. It is found that the column density in the limit of instantaneous pulses is ca. 16 % higher than observed with the imaging pulse lengths of 2 µs (before pulse length correction) used later in chapter 6. A correction factor to remedy this behaviour is therefore included in the analysis.

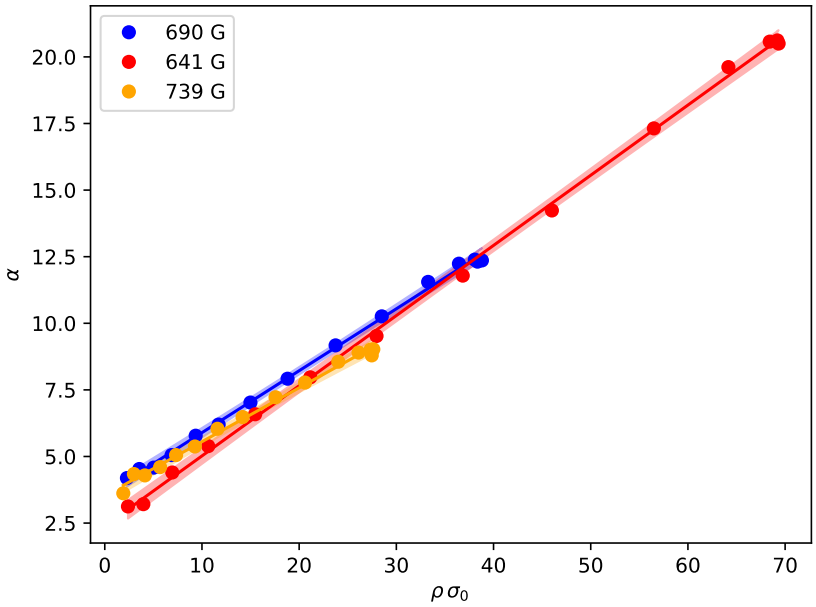

Figure 3.10: Empirically, it is found that α depends linearly on the column density $\rho\sigma_0$ with a small dependence on the interaction strength set by the magnetic field.

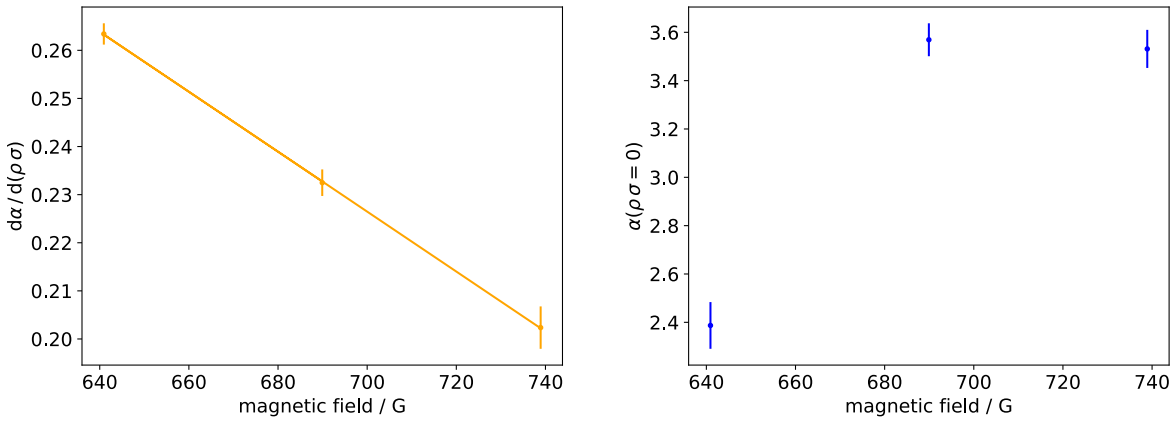

Figure 3.11: Magnetic field dependence of the linear relation $\alpha(\rho\sigma_0)$. The gradient (left) shows a linear dependence on the magnetic field/interaction strength while the situation is less clear for the offset (right).

CHAPTER 4

Detecting Superfluidity

In this chapter, different methods to detect superfluidity in ultracold Fermi gases are presented. The arguably most simple method to identify a superfluid, at least from a conceptual point of view, is based on the Landau criterion of superfluidity [98]: If an obstacle, e.g. a repulsive laser beam, is dragged through a superfluid, it does not create excitations below a critical velocity. This method has been successfully employed by other groups [35, 43] but gives little quantitative information on the system apart from the critical velocity itself. We choose a different approach and determine the fraction of atoms that are condensed.

In contrast to Bose-Einstein condensates with their clear signature of condensation in the momentum distribution, Fermi gases have no obvious signal revealing condensation, see chapter 2. Additional means are therefore necessary to determine the condensate fraction. A traditional method we use is the "rapid ramp" that has already been developed for the first experiments realising the BCS-BSC crossover [39, 40]. This method is not understood in detail yet. Therefore, we developed new methods based on convolutional neural networks and machine learning.

In this chapter, the rapid ramp technique is discussed first. Then, I give a short introduction to convolutional neural networks and machine learning. Finally, two applications of these concepts to detect condensation are presented.

4.1 The rapid ramp technique

If the momentum distribution is recorded with the standard method of free expansion as described in section 3.5, the condensed pairs break upon release from the trap and are therefore not visible as a zero-momentum peak on the recorded images. The instability of the pairs is caused by their origin in many-body effects which cease to play a role when the density diminishes during expansion. The two-body bound state that starts to play a role around unitarity is extremely weakly bound throughout the strongly interacting region of the crossover and therefore unstable as well. Since the pairs comprise of two fermions with opposite momentum and spin, it would be expected that such correlations between fermions can be detected. This is however challenging because the momentum distributions of both spin states have to be recorded in a single realisation of the experiment. This is not a straightforward task because the resonance frequencies for imaging differ. Additionally, only a projection of the full three-dimensional momentum distribution is recorded. Nevertheless, these

correlations have recently been detected in a two-dimensional system with 12 fermions [99].

An alternative approach that has been used already for the first observations of condensation in the BCS-BEC crossover [39, 40] prevents the pairs from breaking during the expansion. To this end, the attraction between fermions is rapidly increased at the moment they are released from the trap. If this is done fast enough, the fermion pairs are projected onto molecules which are strongly bound by two-body effects and thus stable even at low densities. The total momentum of the pairs is 0 and consequently they are observable as a zero-momentum peak on top of the thermal momentum distribution. The condensate fraction is then the defined as the fraction of atoms forming the zero-momentum peak, see figure 4.1. It can be calculated by fitting the bimodal function

$$\begin{aligned} f(x,y) =& A_{\mathrm{th}} \exp\left(-\frac{(x-x_0)^2}{2\sigma_x^2} - \frac{(y-y_0)^2}{2\sigma_y^2}\right) + \\ &+ A_{\mathrm{c}} \left(1 - \frac{(x-x_0)^2}{w_x^2} - \frac{(y-y_0)^2}{w_y^2}\right)^{3/2} + c \end{aligned} \tag{4.1}$$

to the optical density image. Here, the first term describes the thermal background and the second the condensate. x and y indicate the position on the image and all other quantities in equation 4.1 are fit parameters. The condensate fraction is then given by $N_{\mathrm{c}}/(N_{\mathrm{c}} + N_{\mathrm{th}})$ with

$$\begin{aligned} N_{\mathrm{c}} &= \int \mathrm{d}x\mathrm{d}y\, A_{\mathrm{c}} \left(1 - \frac{(x-x_0)^2}{w_x^2} - \frac{(y-y_0)^2}{w_y^2}\right)^{3/2} \\ N_{\mathrm{th}} &= \int \mathrm{d}x\mathrm{d}y\, A_{\mathrm{th}} \exp\left(-\frac{(x-x_0)^2}{2\sigma_x^2} - \frac{(y-y_0)^2}{2\sigma_y^2}\right). \end{aligned} \tag{4.2}$$

4.1.1 Experimental implementation

Experimentally, this method corresponds to a fast ramp of the magnetic field, hence the name rapid ramp. For the implementation it is necessary to quickly reduce the magnetic field strength from a value around the Feshbach resonance to a value far on the BEC side where the expansion happens. To get a faithful momentum distribution, we opt for a final field at the zero-crossing of the scattering length, see figure 3.3, therefore realising a non-interacting system during expansion. Together, this implies two crucial steps for the implementation. Firstly, the magnetic field has to be lowered quickly to a specific value. Secondly, the field has to be stabilised at that lower value. Both steps have to be performed in a reliable manner. Technically, this is a challenging task and the implementation involves several steps as described in [80]. Because a small modification was necessary for the investigations of this thesis, the steps are reviewed here as well. A simplified drawing of the circuit is depicted in figure 4.2.

1. Dipole trap is switched off.

2. After a 1 µs delay, the IGBTs 1 and 2 are turned off and the Feshbach coils discharge over the snubber circuit formed by diode and resistor. The discharge IGBT is turned on to short the

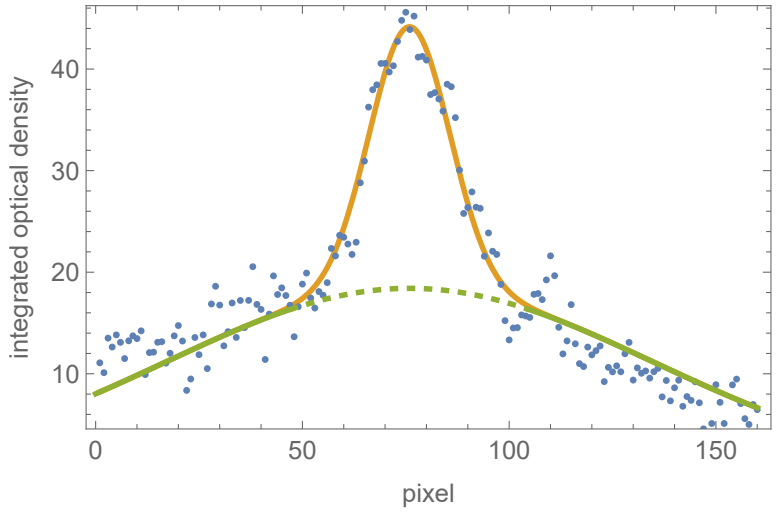

Figure 4.1: The bimodal density distribution after the rapid ramp with a clearly visible condensed peak (orange) on top of the thermal background (green). To improve the visibility of the background, the two dimensional absorption image was integrated along one axis.

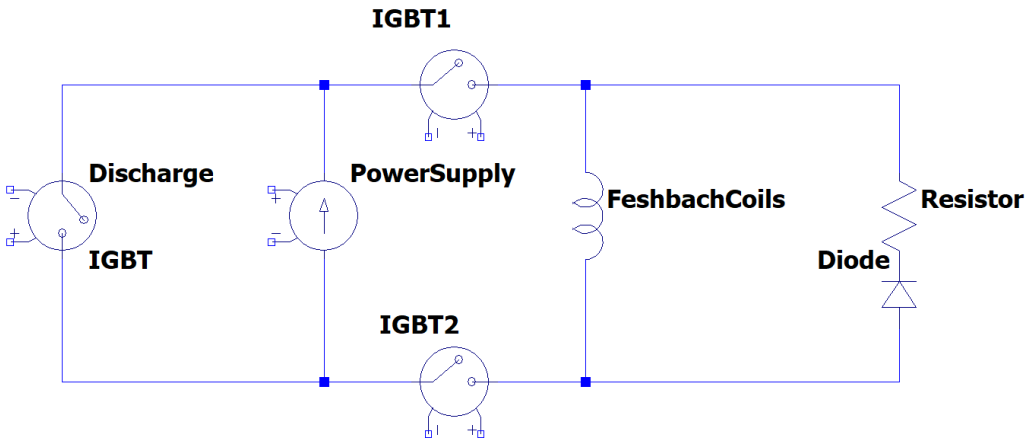

Figure 4.2: Simplified scheme of the circuit implementing the rapid ramp. The Feshbach coils create the magnetic field at the atoms' position. Resistor and diode form a snubber circuit used to dump the energy of the magnetic field. The necessary switching is done with the IGBTs 1,2 and the discharge IGBT. For full details refer to [80].

power supply for faster discharging.

3. A variable time $\tau_{\text{off}} \approx 5\,\mu\text{s}$ to $60\,\mu\text{s}$ later, the IGBTs 1 and 2 are turned back on. τ_{off} increases the more the magnetic field has to be ramped down.

4. At the same time, the control voltage of the power supply is set to a new value corresponding to the magnetic field value where the scattering length vanishes.

5. A variable time $\tau_{\text{discharge}} \approx 170\,\mu\text{s}$ to $260\,\mu\text{s}$ later, the discharge IGBT is turned off again.

6. The magnetic field is now at the zero-crossing of the Feshbach resonance and held there until around 5.2 ms after switch-off of the dipole trap.

7. Finally, the magnetic field is ramped back over approximately 10 ms to slightly below the Feshbach resonance[1]. Here, the atoms are no longer strongly bound and single atom imaging works reliably. The total expansion time is always 15 ms. This ensures that the kinetic energy is fully converted to potential energy by the residual potential created by the magnetic field curvature. The density distribution then reflects the initial momentum distribution of the pairs [94].

The current through the Feshbach coils during the rapid ramp is measured with a current transducer and the corresponding magnetic field created by the coils is shown in figure 4.3. An extremely fast drop, which suggests an equally fast drop of the magnetic field, is observed within the first $100\,\mu\text{s}$. However, eddy currents are created in other parts of the experimental setup when the coils are turned off and create a residual field at the atoms' position. This has been confirmed in [80] by measuring the resonance frequency of the imaging transition for various times during the rapid ramp. In reality, the magnetic field experienced by the atoms decays exponentially and reaches the molecular regime after around $25\,\mu\text{s}$ and the zero-crossing of the Feshbach resonance around $600\,\mu\text{s}$ after the start of the rapid ramp.

4.1.2 Rapid ramp with high temporal resolution

To observe the oscillations in the condensate fraction associated with the Higgs mode, the condensate fraction has to be measured with a time resolution in the µs regime. Consequently, the above steps have to be performed with µs precision. The switching of the IGBTs is unproblematic in this regard because the digital control channels have a time resolution of 100 ns. The analogue output card[2] used to control the power supply however only has a resolution of $100\,\mu\text{s}$. This is measured with regard to the time that has elapsed since the beginning of the experimental cycle. If step 4 of the above procedure, the switch of the scattering length from its value before the rapid ramp to zero, therefore shifts within the sequence by times which are not a multiple of $100\,\mu\text{s}$, the magnetic ramp will vary for different sequences. An example is shown in figure 4.4. Even though these changes are observed to be rather small, they do create a detectable feature in the condensate fraction with a $100\,\mu\text{s}$ period that could possibly obscure other dynamics.

The problem was solved by performing the switch associated with step 4 with a digital instead of an analogue channel. A TTL-controlled channel select circuit now switches between two analogue

[1] 745 G for a $|12\rangle$ and 666 G for a $|13\rangle$ mixture

[2] NI PCI-6733

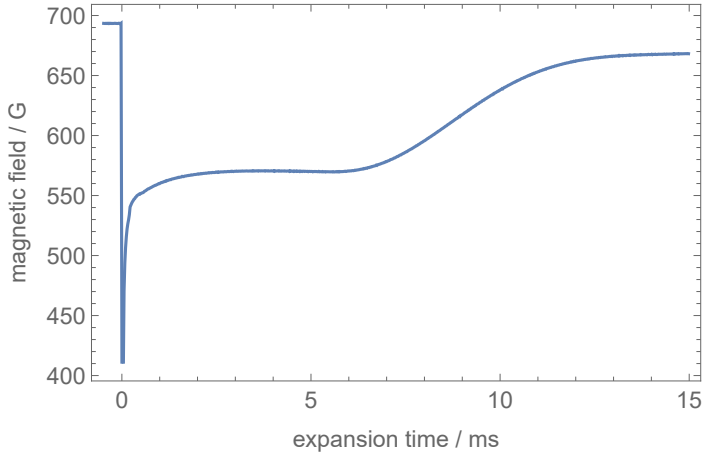

Figure 4.3: Magnetic field created by the Feshbach coils during the rapid ramp. The cloud expands for 15 ms before it is imaged. The field is calculated from the current running through the coils. This current has been measured with a current transducer.

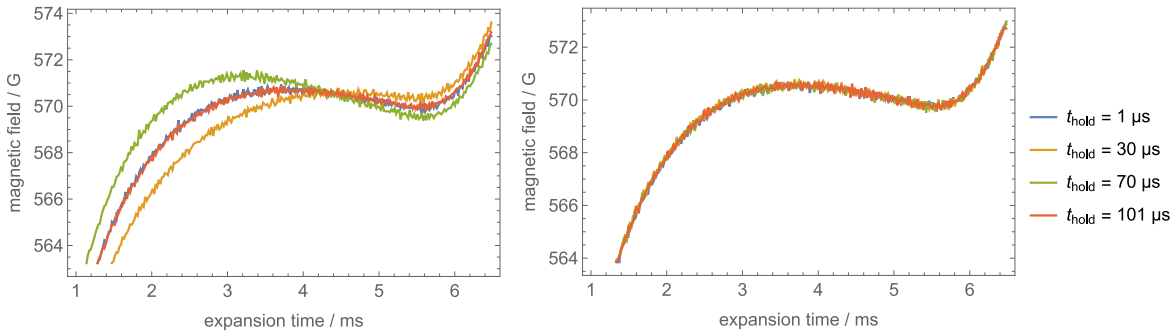

Figure 4.4: The magnetic field during the rapid ramp around the zero-crossing of the scattering length. Shown are experimental sequences with different lengths, adjusted by the variable t_{hold}. Left: without the channel select circuit, the rapid ramp is affected by the time resolution of 100 µs. Right: After inclusion of the channel select circuit, the necessary time resolution of 1 µs is achieved.

channels whose values correspond to the magnetic field strength before the rapid ramp and that of zero scattering length. The analogue channels can therefore remain at constant values up until step 7 and allow for a rapid ramp with the time-resolution of the digital channel. For a demonstration of the effect, see figure 4.4.

The circuit diagram of the channel select circuit can be found in appendix B. The circuit was designed to not affect the control voltage but a small but relevant influence was found in practice. This is most likely due to the fact that the range of 0 V - 5 V is mapped to 0 G - 1 000 G where a precision of a bit less than 1 G is needed. It could however be compensated for by a small adjustment to the formula calculating the voltage necessary for a specific magnetic field value. More details are given in appendix B.

The continuous final ramp to the imaging field (step 7) still suffers from the limited resolution of the analogue channels but the effect is much weaker for the slow ramp compared to an instantaneous switch. Additionally, the final ramp is expected to have little effect on the atoms because the gas has expanded far enough to be weakly interacting. No measures were therefore taken with regard to this minor imperfection.

4.1.3 Uncertainties of the rapid ramp

A big uncertainty is caused by the finite duration of the rapid ramp. Especially the exponential approach to the field corresponding to the zero-crossing of the scattering length is of concern. It implies that the ramp changes at some point from a projective to an adiabatic sweep. While the projective sweep is generally believed to preserve the condensate fraction, the adiabatic sweep should lead to a final condensate fraction that would depend much more on the entropy before the sweep. Since the real rapid ramp is not purely projective, the final condensate fraction cannot be assumed to be a perfect measure of the initial condensate fraction. Relative changes in the condensate fraction are however detectable with the rapid ramp if it is always performed in the same manner [100].

Despite of efforts in both experiment [31, 100] and theory [101, 102], a rigorous understanding of the rapid ramp including the finite ramp time is still not available. Of especial concern are non-condensed molecules in the "pseudogap" state close to, but above the phase transition, which could be wrongly interpreted as a condensate [101]. This problem is not only present for $a > 0$ but already as soon as the BCS limit is left. It has however been found experimentally that macroscopic structures like a vortex lattice “survive” the rapid ramp. It must therefore preserve a certain level of superfluid coherence [41]. The limited understanding of the rapid ramp makes the development of new detection techniques for superfluidity highly desirable.

4.2 Convolutional neural networks and machine learning

Inspired by the general recent success of machine learning with deep neural networks [47] and more specially the detection of a topological phase transition in a quantum gas with a convolutional neural network [49], the possibility to detect condensation in the momentum distribution of a Fermi gas with these techniques was explored. This could almost be a textbook example of a typical problem for machine learning because the condensation leaves a detectable mark in the momentum distribution for which however no model exists to disentangle it from the similar effects of temperature and external potential, see chapter 2. Neural networks work without underlying models and could therefore be able

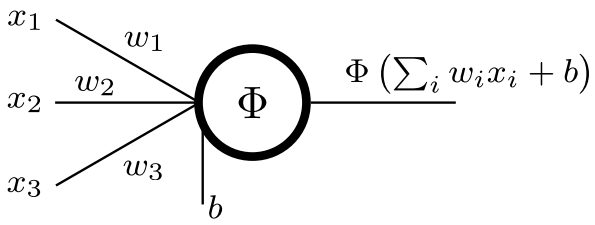

Figure 4.5: Graphical representation of an artificial neuron. The weighted sum of several inputs x_i, plus a possible bias b, is processed by the activation function Φ to generate an output.

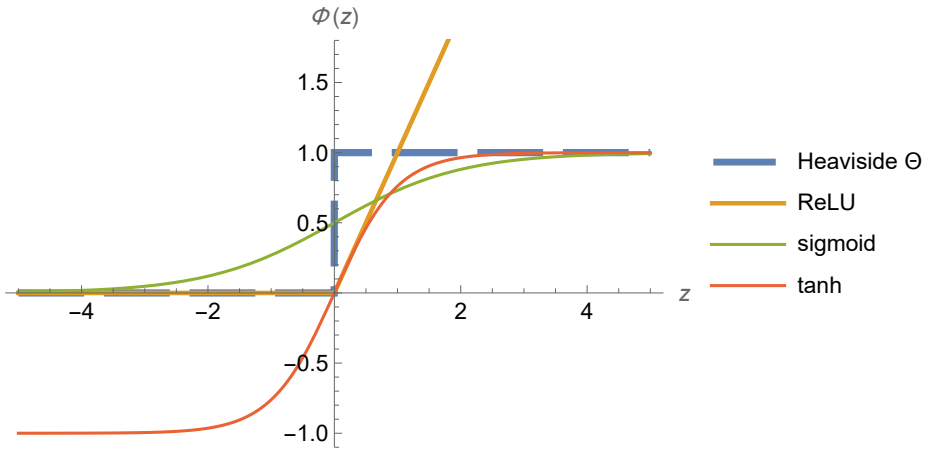

Figure 4.6: Some common options for an artificial neuron's activation function $\Phi(z)$.

to separate condensation from other effects.

This research was initiated by Martin Link in whose PhD thesis [103] a detailed discussion of the employed techniques can be found. An in-depth introduction to the field is provided in [104]. Here, only a rough overview of the basic principles necessary to understand our application is given.

4.2.1 Artificial neurons

The main building blocks of a deep neural network are artificial neurons. They mimic the behaviour of the brain and other nervous tissue in living beings. The working principle of such a neuron is depicted in figure 4.5. Several inputs x_i are added with individual weights w_i and possibly a constant bias b to form z, the input of the neuron,

$$z = \sum_i w_i x_i + b. \tag{4.3}$$

The neurons activation function $\Phi(z)$ then determines the output of the neuron. In the simplest case, this is just the Heaviside-Θ function. In general however, smooth step-like functions like $\tanh(z)$ or the sigmoid $\sigma(z) = (1 + e^{-z})^{-1}$ are preferred because of their non-vanishing gradient. This facilitates the training of neural networks, see below. Another alternative that has the advantage of allowing outputs exceeding 1 is the rectified linear unit $\mathrm{ReLU}(z) = \max(0, z)$. Different activation functions are plotted in figure 4.6 for comparison.

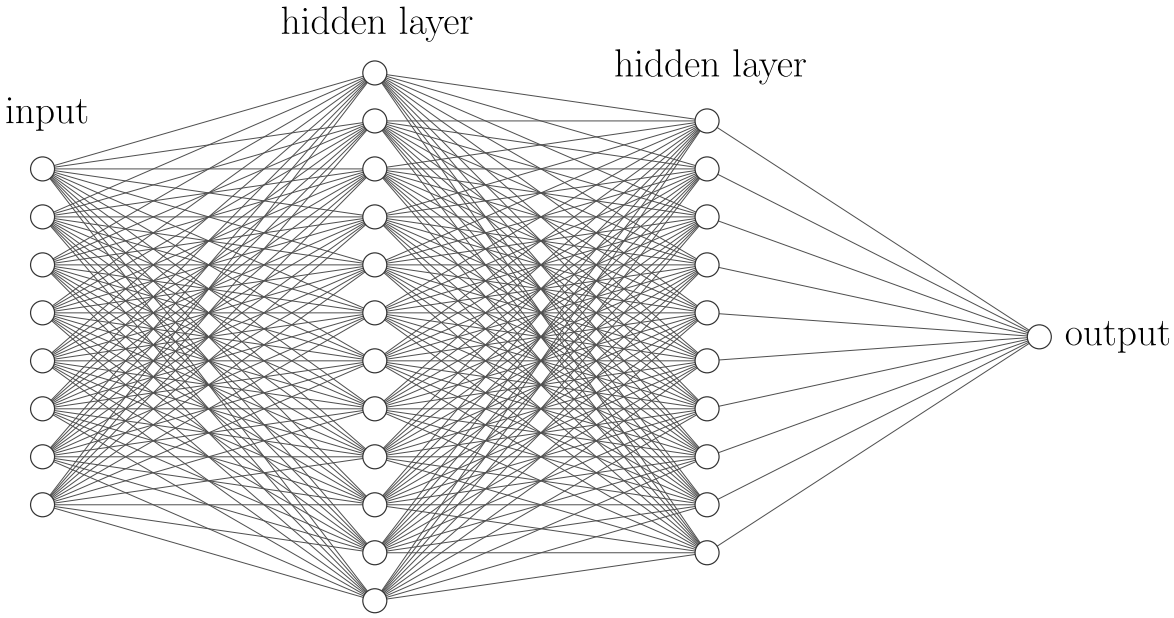

Figure 4.7: Scheme of a deep neural network calculating one number from an eight-dimensional input, for example an image with eight pixels. Every circle represents one neuron and every line has an individual weight. Figure created with the help of [105].

4.2.2 Deep neural networks

While a single artificial neuron seems to be a rather simple object, this significantly changes when many neurons are connected in a deep neural network. It consists of several layers of parallel neurons whose outputs are given as inputs to the neurons of the next layer. All neurons have an individual set of weights. If the output of any neuron is given as an input to every neuron of the subsequent layer, the network is called fully connected. A neural network is called deep if it has at least one layer between the input and output layers; an example is depicted in figure 4.7. If such a network is sufficiently large, almost any function can be approximated with suitable weights [106] making deep neural networks an extremely powerful tool for a vast range of applications.

4.2.3 Convolutional neural networks

For processing images, the architecture of convolutional neural networks has proven to be a successful approach [47]. Since we want to detect condensation from absorption images, this network type is the natural choice for us.

Convolutional neural networks are very similar to deep neural networks but instead of feeding the input directly to the fully connected layers as suggested by figure 4.7, the input image is first processed by (potentially several) pairs of convolutional and max pooling layers. Their working principle is depicted in figure 4.8. In a convolutional layer, a small filter is moved across the large image and combines the covered pixels into a single value according to the filter's weights. In this way, every position of the filter creates one output pixel. The image created by all output pixels is then the output of the layer. Different to the simplified scheme of figure 4.8, the output of a convolutional layer may have a higher dimension than its input [107]. The max pooling layer simply reduces the size of the

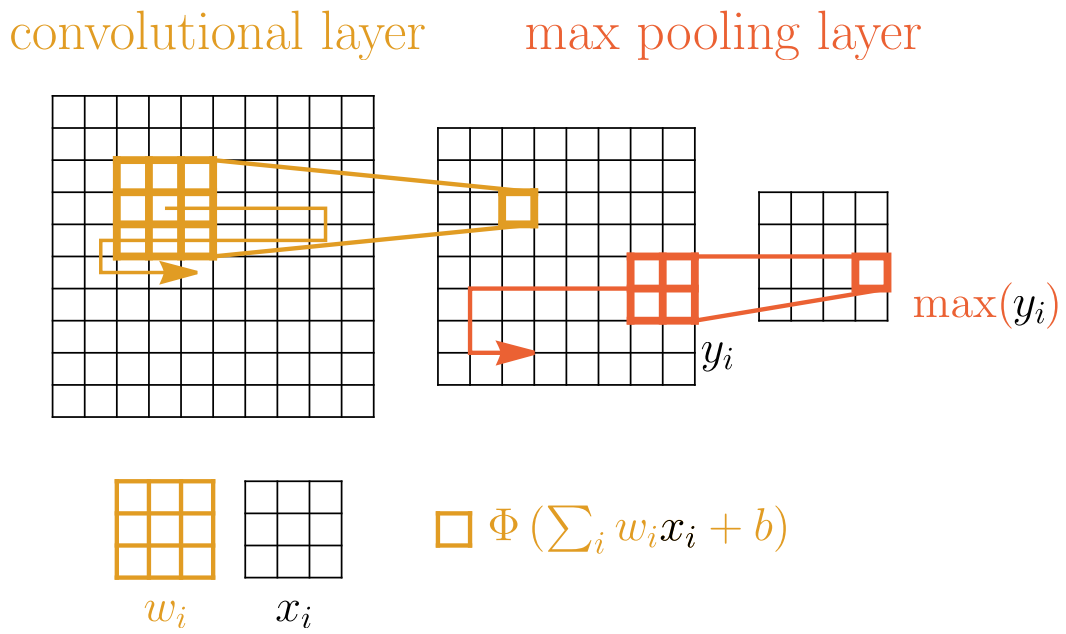

Figure 4.8: Working principle of a pair of convolutional and max pooling layer. Both work by moving a filter across the image to form a new, smaller image as the output. Every filter position yields one pixel in the output. Image adapted from [103]

image by just passing on the highest valued pixel of each filter position. In this way, the convolutional layers highlight features in the input image and reduce its size.

In a convolutional neural network, the output from the last layer is finally fed into a fully connected network similar to the one depicted in figure 4.7. Overall, the whole networks calculates a numerical quantity from the convoluted images.

4.2.4 Training neural networks

The big question is now how the correct weights for a specific task are chosen. Due to the complexity of large networks with millions or even more weights, no procedure to chose them a priori is known. Instead, neural networks are initialised with random weights and subsequently "learn" to perform the desired task by training. To this end, the network processes large samples of data for which the correct output is known – the "labelled" data. The output generated by the untrained network from the labelled data will differ almost certainly from the correct value. The error E, defined with a suitable metric as the distance between output and label, can however be reduced by iteratively tuning the weights $\boldsymbol{w}$ of the network. To this end, the gradient of the function $f : \boldsymbol{w} \to E$ is calculated numerically and adjusting the weights accordingly reduces E. The training of a neural network is therefore nothing but a gradient descent optimisation in a very high-dimensional space. The components of ∇f corresponding to weights in the output layer can be simply determined by the derivative of E with respect to those weights. For the weights in the hidden layers on the other hand, the chain rule has to be used, potentially multiple times. This means that the training procedure traces the network backwards and is therefore called backpropagation. Due to the extremely high dimension of the space on which the function f acts, backpropagation is only feasible with powerful modern computers.

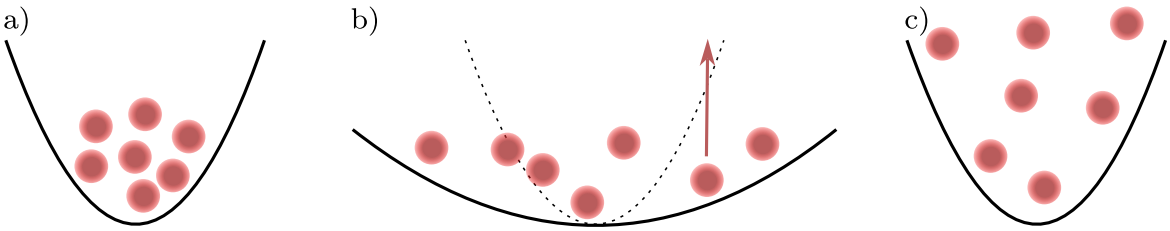

Figure 4.9: The heating scheme. a) The cold gas in the initial trap. b) In the shallow trap, the gas expands. Upon restoration of the original trap (dashed line), the atoms gain potential energy as indicated by the arrow. c) After thermalisation, the gas's temperature has increased.

A question that naturally arises is the effectiveness of training. More specifically, one can ask whether backpropagation finds the global minimum of f. In general, this question is difficult to answer but most likely only local minima in f are found. Experience however shows that this is usually sufficient for neural networks to perform well. In order to check that this also holds for any problem at hand, the labelled dataset can be split into a training and validation set, commonly in a ratio 90:10. The backpropagation optimisation is then performed with just the training set and only considered to be successful if the network can generate the correct output for the validation set equally well. In this case, it is assumed that the network has implemented a generalised model that correctly determines the relevant aspects from the data. As an additional upside, the validation set helps to prevent overfitting of the training set. If overfitting occurs, a decreasing error on the training set coincides with an increasing error on the validation set.

The neural networks whose results are described in this thesis were realised with the TensorFlow library [108] and trained with the Adam optimiser [109].

4.3 A neural network to determine the condensate fraction

As already mentioned in the introduction of this chapter, determining the condensate fraction from a Fermi gas's momentum distribution is similar to typical machine learning applications: An in principle detectable feature is present but no fittable model or other procedure to determine the condensate fraction exists. In order to train a convolutional neural network on this task, a labelled dataset has to be acquired first. Because the network is actively forced to reproduce a certain result, as given by the labelled dataset, this approach is called supervised training.

4.3.1 Generation of a labelled dataset

In order to have a well-working neural network, the training data should cover the full range of the problem at hand. For us, this first of all means that we need images of momentum distributions with widely varying condensate fractions including a substantial amount without any condensate at all. On top of that, the nature of condensation changes in the BCS-BEC crossover and this will be reflected in the resulting effect on the momentum distribution. The training set for a broadly applicable network should therefore include different interaction strengths covering an extensive region of the crossover.

To generate the dataset, clouds of $|12\rangle$-mixtures with six different interactions strengths, $1/k_{\mathrm{F}}a$ ranges from around -0.6 to roughly 1.6 in harmonic approximation, were prepared. While these clouds already have different condensate fractions, a much wider range can be covered if these cloud are heated to gradually deplete the condensate. The heating is performed similar to [110] by a sudden reduction of the trapping beams' power. The atomic cloud then expands in the new shallower trap for a variable time of up to 15 ms. A subsequent recompression of the trap by restoration of the original beam powers suddenly gives the atoms a higher potential energy, see figure 4.9. The amount of energy gained in this way increases with expansion time. After restoration of the original trap, the cloud is given a time of 50 ms to thermalise by converting the gained potential energy into temperature[3]. The result will be a cloud with lower condensate fraction or no condensate fraction at all if the critical temperature has been surpassed. In principle, the atom loss from the heating mechanism could cause the network to learn an unphysical relation between atom number and condensate fraction. The inclusion of different interaction strengths however helps in this regard because the atom number varies with the magnetic field. Additionally, the atom number also fluctuates from experimental shot to shot as well as on longer time scales due to slow drifts of the experimental performance.

The momentum distribution of the final clouds is recorded by imaging them after 5 ms of free expansion. This is a significantly shorter expansion than the $T/4 = 15$ ms needed for the residual harmonic potential to fully reveal the momentum distribution [94] but the signal to noise ratio of the images is greatly improved due to higher densities. The resulting distribution will still reflect the momentum distribution but is much more influenced by the initial in-trap density distribution. On top of this, even a $T/4$ expansion will not provide a faithful momentum distribution due to interactions in the initial phase of the expansion.

The labels necessary for training, in this case the condensate fractions of the imaged clouds, are generated by preparing clouds with identical interaction strengths and heating times as for the training images but then performing the rapid ramp technique from section 4.1 to obtain the clouds' condensate fraction. In this way, a dataset of 7895 labelled images is created [107].

4.3.2 Training of the network

With the labelled dataset, a convolutional neural network is trained to determine the condensate fraction from the fermionic momentum distribution detected after a time-of-flight of 5 ms, for a sketch of this procedure see figure 4.10. The network consists of three pairs of convolutional and max pooling layers followed by two internal layers, the first with 600 and the second with 300 neurons, and one output neuron whose output is the condensate fraction. In total, the network has 11601591 free parameters tuned during training. The exact composition was not found to be critical since small changes do not hamper the performance of the network. A slightly more detailed discussion is provided in the publication [107] while in-depth discussions of the training process are given in the PhD theses [93, 103]. An important point during training is the exclusion of the temperature range around the phase transition from the training data. In this way the neural network identifies the phase transition by itself and does not blindly follow the rapid ramp which might be inaccurate for very small condensate fractions. Far above the phase transition, the labels are simply set to zero instead of using the fitted condensate fraction. This is done because the fit always predicts a non-zero condensate

[3] This conversion will not be one-to-one because atoms are evaporated in the thermalisation process and carry some of the added energy away with them.

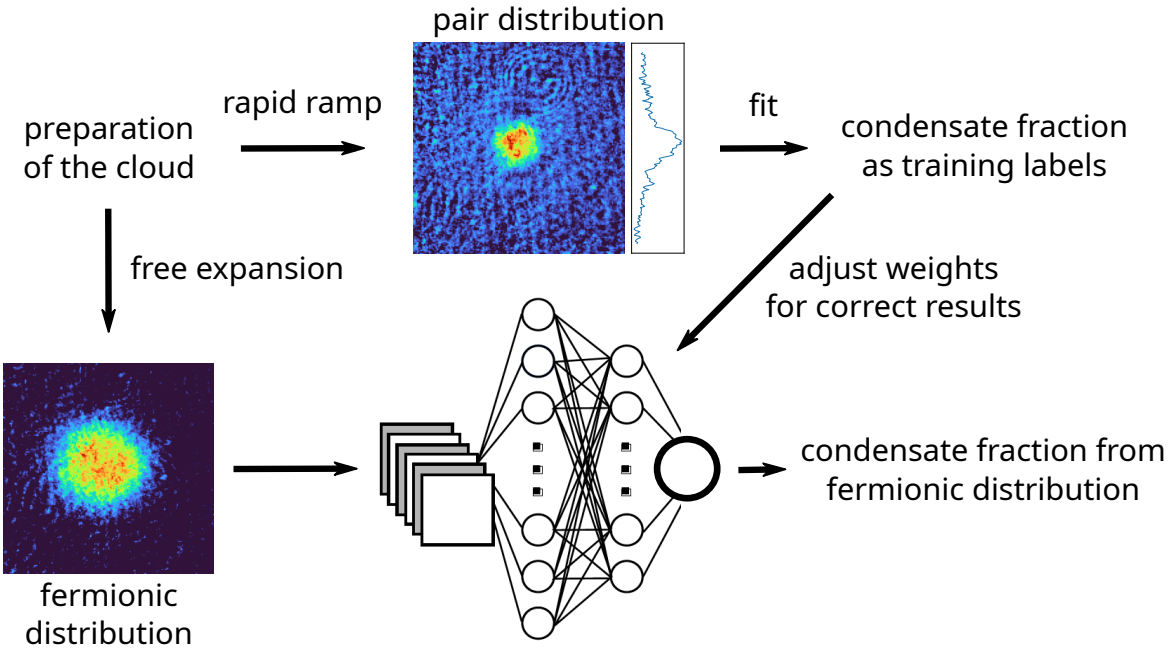

Figure 4.10: Scheme of the training procedure. The atomic clouds are either imaged after a free expansion or after a rapid ramp. With the fitted condensate fraction, the network can be adjusted to determine the condensate fraction from the fermionic momentum distribution.

fraction. The reason for this are experimental noise and imperfections which cause deviations from a pure thermal distribution.

After 15 iterations of gradient descent optimisation, the network can successfully predict the rapid ramp result for the condensate fraction from the fermionic momentum distribution. For a given hold time, the predicted condensate fractions scatter as much as those measured with the rapid ramp. The performance of the neural network is therefore equal to the rapid ramp and only limited by the training data.

In the normal phase, where the bimodal fit wrongly identifies experimental imperfections as a condensate, the neural network outperforms the rapid ramp because it unambiguously predicts a condensate fraction of zero.

4.3.3 Validation of the training

In order to ensure that the network has successfully learned to detect the features associated with condensation, we do not rely on the labelled validation data alone. As a first additional check, the network has to identify the condensate fraction of clouds with an interaction strength that was not included in the training set. Here, two cases can be distinguished: firstly, interaction strengths within the range spanned by the training data. This tests the interpolation capabilities of the network. Secondly, fermion attraction strengths that are either weaker or stronger than the weakest or strongest attraction in the labelled dataset are used to test the extrapolation capabilities. Extrapolation should be possible for the network, at least up to a certain degree, because the signal of condensation is expected to change little within the BCS and BEC limits. In both cases, we do not find significant deviations

from the results obtained with the rapid ramp technique.

A second check is to test the network with clouds of atoms in a $|13\rangle$-mixture where the atom number and background magnetic field are different. The latter will slightly affect the expansion of the cloud. Nevertheless, the network still predicts the condensate fraction correctly [103].

Together, we therefore conclude that the trained network has indeed found a way to extract the condensate fraction from the momentum distribution making it the first tool that can perform this task.

4.3.4 Advantages of the network

Apart from the novelty aspect of extracting the condensation signal from the momentum distribution, the trained network has the technical advantage of being able to replace the rapid ramp. This is experimentally beneficial because the correct values for the times τ_{off} and $\tau_{\text{discharge}}$ depend on the initial current through the coils and therefore have to be determined empirically for every magnetic field. If the condensate fraction for many different interactions strengths is supposed to be measured, usage of the network can reduce the necessary effort noticeably.

A second and much more substantial advantage of the network over the rapid ramp is the possibility to identify a clear phase transition. With the rapid ramp this is not possible because of noise in the bimodal momentum distributions. If the condensate fraction drops below a threshold, usually around 3%, the best fit for the bimodal function 4.1 no longer identifies the condensate peak but compensates deviations from a perfect thermal background instead. This effect usually results in condensate fractions of around 5% close to and above the phase transition. It is therefore not possible to precisely locate the phase transition with the rapid ramp. The neural network on the other hand is able to identify the absence of a condensate and does so consistently above a threshold heating time. This property will be exploited in the next chapter to construct a high-resolution phase diagram of the BCS-BEC crossover.

4.3.5 Limitations of the network

A general problem of supervised neural networks is that they learn all errors present in the training data. In our application, these are especially all the uncertainties associated with the rapid ramp, see section 4.1.3. The trained network therefore has to be considered as simply a variation of the rapid ramp technique even though it does not utilise the projection from many-body pairing to dimers directly. Another constraint comes from the reduced expansion time of the cloud: after 5 ms time-of-flight, the cloud's shape still depends on the trap geometry. The network is thus trained on a specific power and alignment of the trapping beams. After a re-alignment of the trapping beams, it was indeed necessary to train the network anew with a newly recorded labelled dataset.

Since the input image is pixel-based, a change in the magnification of the imaging system or the pixel size on the camera chip will cause the network to give wrong predictions of the condensate fraction. During the project, the camera had to be replaced by a new one with smaller pixel size. In this situation, it was possible to adapt the already trained network with a method called transfer learning [111]. The network optimised for the old camera was taken as the starting point for a new round of training; this time with labelled images taken by the new camera but rescaled to the old size. In this way, the network is already close to a good choice for the weights at the beginning of the training process and successful training becomes possible with a considerably smaller new labelled dataset, in our case, we used 2500 images.

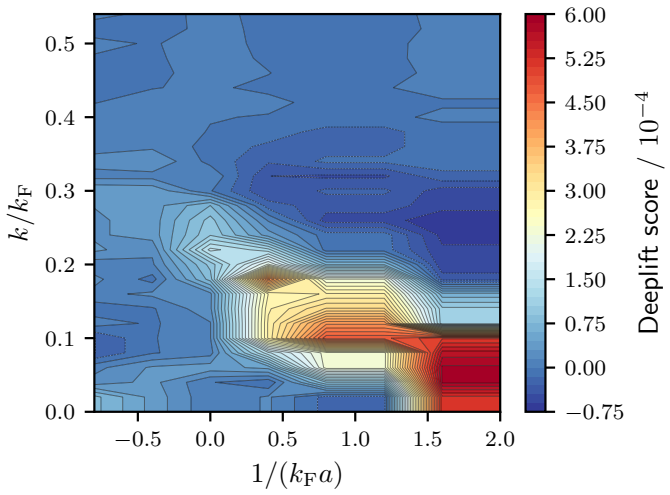

Figure 4.11: Analysis of the trained neural network. For the lowest achievable temperature and all interactions, the importance of different momenta to the determination of the condensate fraction is found with the help of [112]. Image from [107].

4.3.6 Insights provided by the network

The complexity of neural networks makes it extremely challenging to interpret their behaviour and decision making, hence they are often considered to be "black boxes". A possibility to gain more insight is finding the input neurons which have the largest influence on the result, as implemented in the DeepLIFT method [112]. In our case, this will tell us which regions of the cloud have the most relevance for determining the condensate fraction. Since we can assign a momentum to every pixel if we assume free expansion, it becomes possible to figure out which momenta carry, for the neural network, the most information with respect to the condensate fraction. The result is depicted in figure 4.11 with k_F determined from the Fermi energy in the harmonic approximation. It reveals that the neural network considers, for different interactions, roughly the regions where condensation is expected to have the strongest signature: on the BEC side only low momenta are important, as expected for a bosonic condensate, while higher momenta become relevant towards unitarity and towards the BCS limit where outer and central regions are equally important. This importance of high momenta for condensation is a distinctive feature of fermionic condensation. Higher momenta around $k = k_\mathrm{F}$ do not show up in figure 4.11 due to a small signal to noise ratio in the outermost regions of the cloud. Additionally, the momenta are calculated assuming, in contrast to the experiment, a free expansion of the cloud. The attractive interactions could lead to a slower expansion of the cloud and therefore an underestimation of k/k_F.

4.4 Unsupervised detection of superfluidity

4.4.1 Autoencoder networks

By taking full advantage of the versatility of neural networks, it could also be possible to detect superfluidity without enforcing a predetermined result as in the supervised approach. The big appeal of

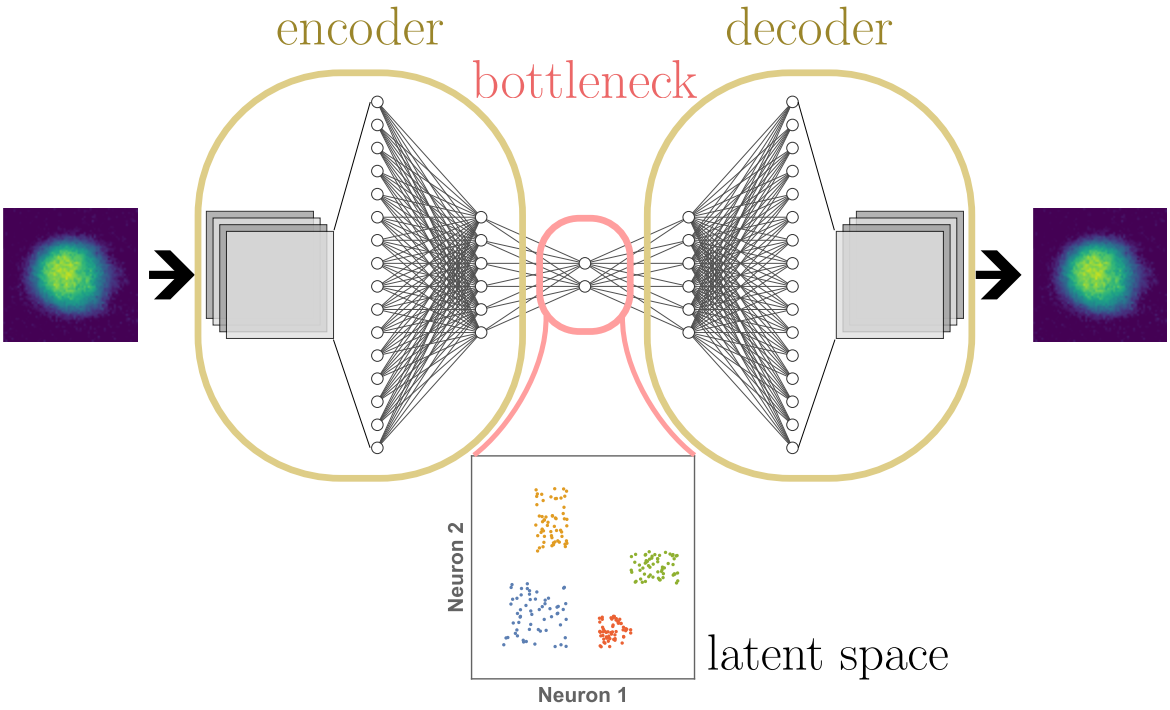

Figure 4.12: Working principle of the autoencoder. See main text for explanations.

such an *unsupervised* approach is the elimination of systematic errors that might have been introduced during the labelling process. Since some form of labelling is however still necessary for training, a different network architecture, relying on labels that are not equal to the result, has to be chosen.

The autoencoder architecture as sketched in figure 4.12 has been conceived for these purposes. The general idea is to propagate an input, in our case an image of a momentum distribution, through a neural network with an information bottleneck in such a way that the input can be reconstructed as the output. To this end, the first part of the network, the encoder, has to identify the defining features of every input and compress them in a form that they can be passed through the bottleneck, a layer with a very small number of neurons. The second half of the autoencoder, the decoder, is structurally the inverse of the encoder and receives the encoded information from the bottleneck to create an output that is supposed to be as close as possible to the input. Together, this makes it possible that the inputs themselves serve as their own training labels. Consequently, no additional information has to be provided to enable the usual backpropagation training and no systematic error can be introduced in the labelling process.

Once the autoencoder has been trained to faithfully recreate the inputs, the decoder is stripped off and the output of the bottleneck layer, called the latent space, is investigated further. It can now be considered as a low-dimensional representation of the input data, which highlights the differences between the individual inputs. All common features in the data can be discarded by the encoder because the decoder has learned to reconstruct these during training. Overall, similar inputs are then expected to lie close to each other in the latent space while very different ones are expected to be further apart. In this way, an obscured structure in the original data can reveal itself in latent space.

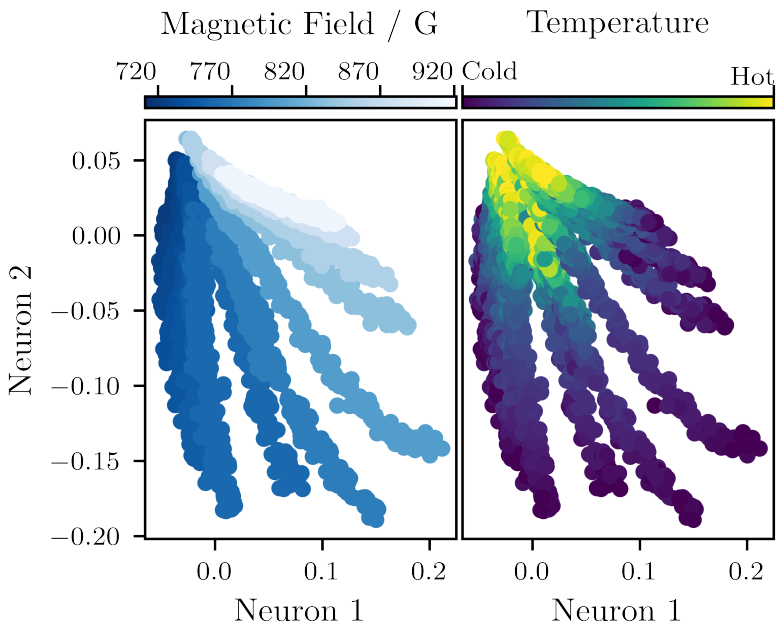

Figure 4.13: Latent space representation of momentum distributions covering a large range of temperatures and interaction strengths. Every image corresponds to one point. The left panel shows that the autoencoder sorts images with identical interaction strength in lines while the right panel reveals that the lines themselves are sorted by temperature. For clarity, temperature is only shown in a qualitative way because heating times vary for the different fields. Image from [113].

4.4.2 An autoencoder for momentum distributions

Applied to our system, we train an autoencoder network with 5031 images of momentum distributions recorded at various temperatures and interaction strengths, again with a ratio of 90:10 between training and validation data. For the autoencoder, we use the same data as for the network from section 4.3. The structure of the encoder is very similar to the network for determination of the condensate fraction from section 4.3. Two iterations of convolution and max pooling followed by two hidden layers with 512 neurons each and finally the bottleneck, which consists of two neurons. The decoder part of the network has the same structure as the encoder but reversed, so that it can reproduce the input from the output of the two bottleneck neurons. A more specific structure will be given in the future publication [113].

The latent space representation of a large number of momentum distributions is presented in figure 4.13. For every interaction strength, both the thermal and the superfluid phase are covered. It reveals that the data is sorted in the latent space by its two relevant physical quantities: different interaction strengths are separated into lines and along these lines, the clouds are ordered by their temperature. Even more so, the hottest clouds are clustered together and therefore considered to be fairly similar by the autoencoder. On the other hand, the coldest clouds lie far apart from each other which highlights their distinctness. This sorting mirrors qualitatively the expectations from simple physical arguments. If heated far enough, all momentum distributions should approach the Boltzmann distribution independent of interactions or particle statistics. In the zero-temperature limit on the other

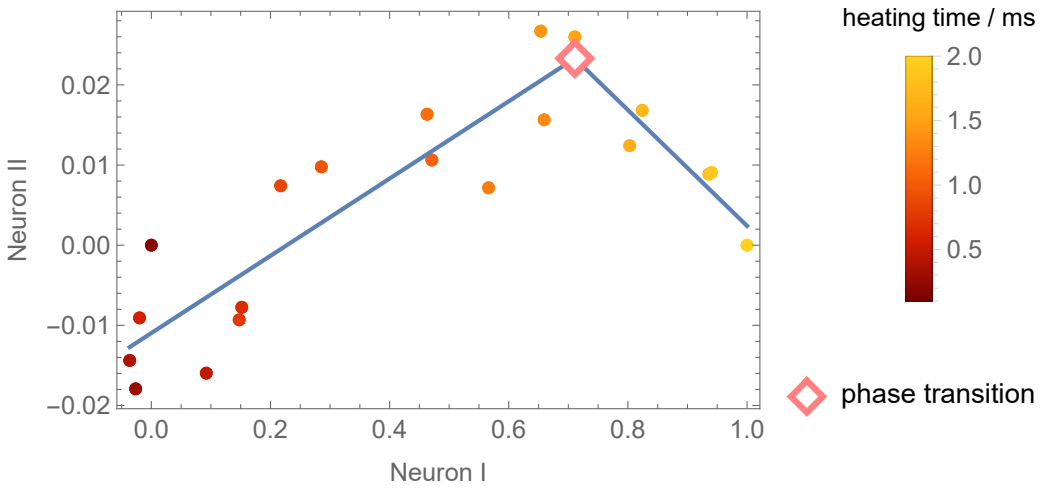

Figure 4.14: The latent space around the phase transition for $1/k_F a \sim -0.26$ after a linear transformation. The points are the average positions of around 13 images with identical preparation. The piecewise linear fit to identify the phase transition as well as the position of the phase transition are shown.

hand, the interactions have defining influence on the momentum distribution and a large continuous variation should be found through the crossover, as shown in section 2.4. This behaviour is indeed reflected in the latent space. It should be noted that these results are a consequence of the network and data structure alone and no information on magnetic fields, interaction strengths, temperature etc. is ever provided to the network.

The use of two neurons in the bottleneck is motivated by the presence of the two parameters varied in the data. Nevertheless, there is a possibility that the use of more neurons could reveal even more information. A network with a bottleneck consisting of three neurons was therefore trained as well but revealed no new information. Instead, all points lay on or very close to a common plane in the three-dimensional latent space indicating that no new information is encoded in the third neuron.

4.4.3 Signature of superfluidity

A straightforward question is whether the autoencoder is able to detect the superfluid phase transition. In figure 4.13, no related, obvious feature catches the eye. This however changes if only the region close to the phase transition, whose approximate location is known from the rapid ramp technique, is considered. There, the points can be separated into two regions; in both, the points are organised roughly along a line and sorted by temperature. The two regions are not distinguishable by a separation in the latent space but because the associated lines have different angles in the latent space, see figure 4.14. We interpret the two regions as different responses to the heating and the transition between the regions as a qualitative change of the system. Since the superfluid phase transition is expected to coincide with such a change, we identify the intersection between the regions as the phase transition. The linear nature of the regions enables the identification of the phase transition by fitting a piecewise linear function. To increase the robustness of the fit, the region around the phase transition is shifted and linearly transformed prior to fitting while all angles are preserved. An example of the fit around

the phase transition is shown in figure 4.14. The critical heating time is then determined from the average of the three points in latent space closest to the position of the phase transition. It is found that this heating time is equal or at least very close to the critical heating times found with the neural network trained with the rapid ramp data.

With this procedure, the phase transition can be determined over a wide range of interactions directly from the time-of-flight images of the cloud. Because no additional input is needed at any point in the whole process, this is a completely new way to determine the superfluid phase transition. Unlike previous methods like the rapid ramp, it does not rely on assumptions in the interpretation of the images that could lead to systematic errors. The biggest uncertainty comes from the not particularly pronounced signal of the phase transition in latent space. A closer look on this will be provided in the next chapter.

CHAPTER 5

Phase diagram of the BCS-BEC crossover

The neural network based techniques to identify superfluidity introduced in the previous chapter identify the phase transition clearly for a wide range of interaction strengths. They can therefore be employed to construct a phase diagram of the BCS-BEC crossover. So far only the crude phase diagrams of the first crossover realisations [39, 40] and one more refined measurement [114] are available. The latter however reports a fairly constant critical temperature in the strongly interacting regime, deviating significantly from theoretical expectations. A construction of the phase diagram with the new techniques could therefore provide important insights.

For the phase diagram, the phase transition has to be determined over a wide range of interaction strengths. The necessary experimental data are generated in the same way as the labelled dataset, see section 4.3.1. The trained networks are then used to identify the phase transition. For a meaningful phase diagram, it is necessary to convert the "lab units" magnetic field and heating time into the dimensionless quantities interaction strength $1/k_\mathrm{F}a$ and reduced temperature T/T_F. While a can be simply calculated from the magnetic field with the help of equation 3.1, a reliable determination of T and E_F (which in turn provides k_F and T_F) is significantly more involved and will be discussed in this chapter.

5.1 Complications of a trapped Fermi gas

The theoretical treatment of the harmonically trapped, strongly interacting Fermi gas is significantly more involved than the already challenging homogeneous gas. This is due to the presence of an extra length scale, the harmonic oscillator length $a_\mathrm{HO} = \sqrt{\hbar/m\bar{\omega}}$. The only theoretical predictions of the critical temperature [115, 116] originate from the time shortly after the first realisation of the BCS-BEC crossover in 2004 and, at unitarity, agreed relatively well with measurements of the time [110, 117]. Since then, the use of the local density approximation (LDA), see section 2.5.2, has become the standard treatment to connect the trapped case to the better understood theory of the homogeneous case. This has been especially fruitful at unitarity where the complete equation of state for the homogeneous case was measured in this way [42, 118, 119]. At unitarity, where $a \to \infty$, a drops out as a length scale leaving $1/k_\mathrm{F}$ as the only remaining scale. The system is therefore describable by a few universal functions that can be deduced from experimental data [45].

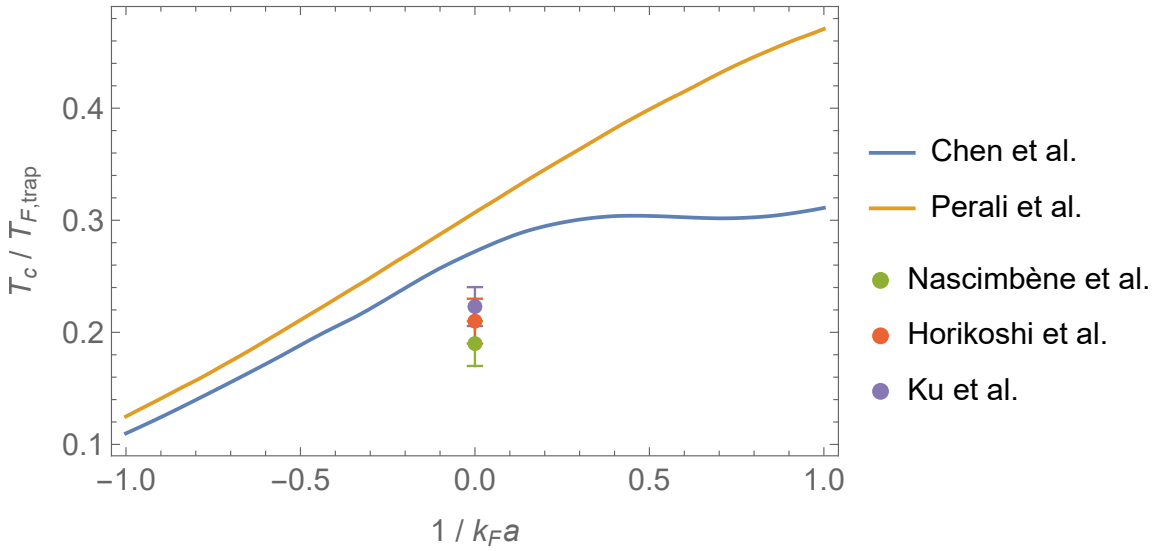

Figure 5.1: Comparison of theoretical predictions for the critical temperature in the BCS-BEC crossover – "Chen et al.' from [116] and "Perali et al." from [115] – with the values determined from the measured equations of state – "Nascimbène et al." from [118], "Horikoshi et al." from [119], and "Ku et al." from [42].

The shift from harmonic theory to LDA has an important implication, namely the definition of the Fermi energy changes. In the harmonic case, it is set by atom number and trap frequencies $E_{\mathrm{F,trap}} = \hbar\bar{\omega}(3N)^{1/3}$ while the density is decisive in the homogeneous case $k_{\mathrm{F}} = (3\pi^2 n)^{1/3}$. Since the measured homogeneous equations of state at unitarity allow for a complete description of the system, it is possible to calculate the critical temperature of the harmonically trapped unitary gas from them, see below. In figure 5.1, the resulting values are compared to the early theoretical predictions. It is apparent that the relatively old theories for a trapped gas overestimate the critical temperature. This is in contrast to the homogeneous theories from figure 2.9 which are in good agreement with the values determined from the equation of state measurements.

For a fruitful comparison with theories, the critical temperatures determined with the help of the neural networks should therefore be normalised with a T_{F} determined from the density of the cloud. Since the clouds are not homogeneous, a criterion to determine the correct T_{F} from the inhomogeneous density distribution has to be found. In the case of the phase transition, this is particularly easy because superfluidity will appear first where T/T_{F} is lowest and this is where the density is the highest, i.e. the centre of the cloud. k_{F} calculated from the central density is therefore a normalisation that allows the comparison of critical temperatures of trapped gases with homogeneous theories.

5.1.1 Critical temperature of the trapped unitary gas from the equation of state

While [118, 119] provide the critical temperature of the trapped unitary gas shown in figure 5.1, it had to be calculated for the more precise measurement of the equation of state [42]. Since this calculation is quite instructive and introduces some concepts needed later for the thermometry of the cloud, it is presented here.

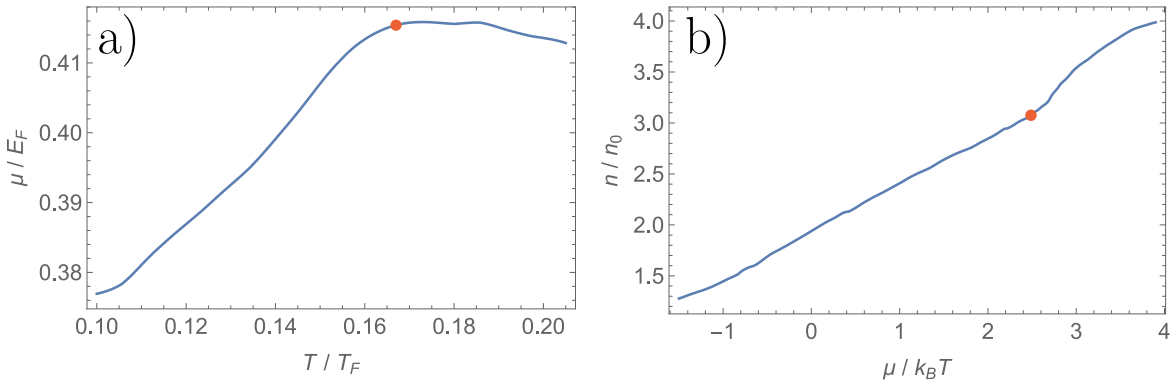

Figure 5.2: Measured equation of state of a homogeneous unitary Fermi gas from [42] with the red dot indicating the phase transition. a) The temperature dependence of the chemical potential. b) The density equation of state, normalised by the density of a non-interacting Fermi gas n_0 at the same value of $\mu/k_\mathrm{B}T$. Data extracted with [120].

The general idea of the calculation is to use the measured equation of state for the homogeneous gas from [42] to calculate the density distribution of a unitary Fermi gas at the critical temperature in a harmonic trap. Integration of the density profile will reveal the atom number and allow the calculation of $T_\mathrm{c}/T_{\mathrm{F,trap}}$.

As a first step, the chemical potential of the homogeneous unitary gas at T_c has to be found. Its temperature dependence is given in [42] and shown in figure 5.2 a). At the phase transition, the ratio $\mu/k_\mathrm{B}T = 2.49$ is found. This ratio determines the density of the cloud via the density equation of state from figure 5.2 b). This density is the starting point for the calculation of the density profile becuase it is the density at the centre of the trap, where $V = 0$. For $V > 0$, the density is calculated according to the LDA by the replacement $\mu \to \mu - V$, which can be plugged into the density equation of state.

A problem arising is the density equation of state not being provided for $(\mu - V)/k_\mathrm{B}T \lesssim -1.3$ in [42]. In this region, the gas is however very hot, $T \approx 1.8T_\mathrm{F}$, and the equation of state can be approximated by a virial expansion. It consists of writing the grand potential Ω as the power series [121]

$$\Omega = -2\, k_\mathrm{B}T\tilde{V}\lambda_\mathrm{dB}^{-3}\left(z + b_2 z^2 + b_3 z^3 + \cdots\right) \tag{5.1}$$

with the volume of the homogeneous system $\tilde{V}$, the thermal deBroglie length λ_dB, the i-th order virial coefficient b_i and the factor 2 accounting for the spin degeneracy. z denotes the fugacity which is, in LDA, $z = \exp((\mu - V)/k_\mathrm{B}T)$. At unitarity, we have $b_2 = 3 \times 2^{-5/2} \approx 0.530$ [122] and $b_3 = -0.290953...$; this value has first been reported in [121] and is by now well established [123]. At unitarity, all b_i are independent of temperature. Since $n = -\tilde{V}^{-1}\partial\Omega/\partial\mu$, the virial expansion for the density equation of state n is given by

$$n = 2\,\lambda_\mathrm{dB}^{-3}\left(z + 2b_2 z^2 + 3b_3 z^3 + \cdots\right). \tag{5.2}$$

In figure 5.3, it becomes apparent that a 3rd order virial expansion is sufficient to extend the measured equation of state to the regime of very high T/T_F where $\mu/k_\mathrm{B}T \to -\infty$. The figure normalises n from equation 5.2 by the non-interacting equation of state $n_0 = -2\,\mathrm{Li}_{3/2}(-z)$.

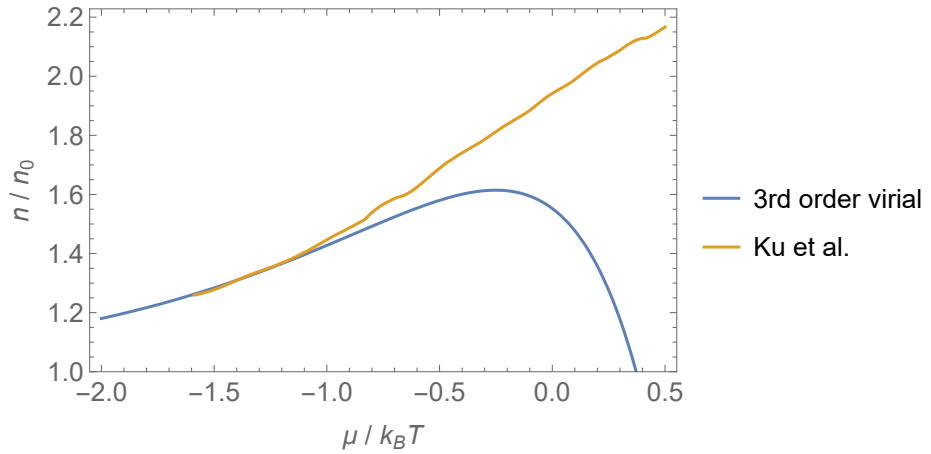

Figure 5.3: The 3rd order virial expansion covers the high T end of the measured density equation of state from figure 5.2 [42] and extends it to infinite T ($\mu/k_\mathrm{B}T \to -\infty$).

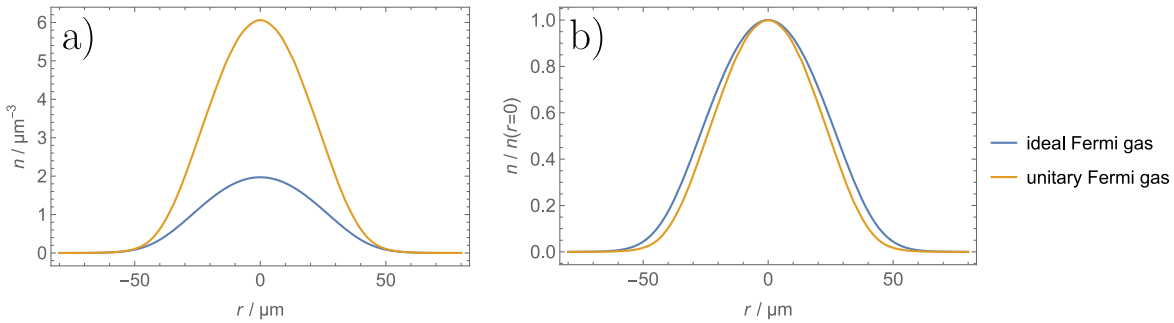

Figure 5.4: a) Cut along a weakly confining axis through the density distribution of a unitary Fermi gas at the critical temperature. For comparison, the density distribution of an ideal Fermi gas with identical temperature and chemical potential in the same trap is shown as well. b) Same as a) but normalised to the central density.

Now everything is together to calculate the density distribution in a harmonic trap. Figure 5.4a) shows a cut through the density distribution for a unitary Fermi gas consisting of ^{6}Li atoms in a potential with trapping frequencies $(\omega_x, \omega_y, \omega_z) = 2\pi \times (160\,\mathrm{Hz}, 160\,\mathrm{Hz}, 260\,\mathrm{Hz})$ and a temperature $T = T_\mathrm{c} = 214\,\mathrm{nK}$, parameters that are comparable to those in the experiment. For comparison, the density distribution of an ideal Fermi gas with identical T and μ is shown as well and highlights the strong influence of the interactions. In the wings of the cloud, where T/T_F calculated from the local density becomes very large, the two density distributions both approach the value given by the Boltzmann distribution. The ideal Fermi gas is a special case because $E_\mathrm{F,trap}$ and E_F determined from the central density are, by definition, identical. For the unitary Fermi gas, $E_\mathrm{F,trap}$ and E_F differ due to the density profile having a different shape compared to the ideal case, see figure 5.4b).

The atom number in the trap can be obtained by numerical integration of the full unitary density profile giving $N = 400\,000$. The Fermi energy of the trapped gas is then given by $E_\mathrm{F,trap} = \hbar\bar{\omega}(3N)^{1/3}$ which leads to a critical temperature of the harmonically trapped unitary Fermi gas of

$$T_\mathrm{c}/T_\mathrm{F,trap} = 0.223 \pm 0.017. \tag{5.3}$$

This is significantly lower than predicted by the theories in figure 5.1 but matches the estimation of [31], which was made prior to the measurement of the unitary equation of state, and is in agreement with the other measurements [118, 119].

5.2 The inverse Abel transform

Since we have to determine the Fermi energy for the phase diagram from the central density of the clouds, we need a way to measure it. While it is straightforward to image the cloud *in-situ*, such an image does not reveal the density $n(x, y, z)$ but only the column density $\rho(x, y)$, the density integrated along the camera's line of sight z, see also section 3.5. If the cloud's density distribution possesses spherical symmetry in the y, z-plane, the integration can be inverted with the inverse Abel transform [124]

$$n(x, r) = -\frac{1}{\pi} \int_r^\infty \mathrm{d}y' \frac{1}{\sqrt{y'^2 - r^2}} \frac{\partial \rho}{\partial y'}(x, y') \tag{5.4}$$

where the coordinate r denotes $\sqrt{y^2 + z^2}$. If the equidensity/equipotential lines in the y, z-plane form ellipses of the form $y^2/w_y^2 + z^2/w_z^2 = \text{const}$ rather than circles, the inverse Abel transform only has to be modified slightly [42] to

$$n(x, r) = -\frac{w_y}{w_z} \frac{1}{\pi} \int_r^\infty \mathrm{d}y' \frac{1}{\sqrt{y'^2 - r^2}} \frac{\partial \rho}{\partial y'}(x, y') \tag{5.5}$$

where the radial coordinate is now given by $r = \sqrt{y^2 + z^2 \cdot w_z^2/w_y^2}$. Such a situation occurs in a harmonic trap if the frequencies in y- and z-direction differ. Since our trap is mostly harmonic, we use equation 5.5 to convert the recorded column densities to real densities. This introduces a small systematic error which will be investigated later.

Performing the inverse Abel transform with experimental data comes with some subtleties because the integral diverges for $y' \to r$ and the derivative of ρ has to be determined numerically from discrete quantities. Both of these effects can enhance experimental noise present on ρ. We therefore usually take advantage of the elliptic symmetry of $\rho(x, y)$ and average it radially while respecting the ellipticity to reduce the noise. The inverse Abel transform is then performed with an algorithm developed by Andreas Kell which deals with some of the noise sources [93]. In this way, the central density, needed for normalisation in the phase diagram, is determined.

5.3 Thermometry

Performing a thermometry of an ultracold Fermi gas is a difficult problem because the small variations of the Fermi edge with temperature are difficult to observe. On top of this, interactions add even more complications. The cleanest way to measure the temperature would be the admixture of a small quantity of bosonic atoms which can serve as a "thermometer". They can be imaged independently and, if in thermal contact with the fermions, reveal the common temperature. This method is, for example, used in [118] where a small amount of bosonic ^{7}Li is added to a unitary Fermi gas of ^{6}Li atoms.

Our setup is in principle well equipped for this kind of thermometry because a small quantity of bosonic ^{23}Na can be loaded into the dipole trap with ease, only requiring a slightly shorter evaporation in the magnetic trap. However, the large mass difference between ^{23}Na and ^{6}Li causes a very different gravitational sag which physically separates the clouds in the final trap thus prohibiting any thermal contact.

We therefore opt for a different method based on the fact that the equation of state will approach the Boltzmann regime for very low densities as shown by the virial expansion 5.2. In this case, the density dependence on the external potential reads

$$n(V) = 2\,\lambda_{\mathrm{dB}}^{-3} \exp\left(\frac{\mu - V}{k_{\mathrm{B}}T}\right) \tag{5.6}$$

and the logarithm of the density depends linearly on the potential

$$\log(n)(V) = \log\left(2\,\lambda_{\mathrm{dB}}^{-3}\right) + \frac{\mu}{k_{\mathrm{B}}T} - \frac{V}{k_{\mathrm{B}}T} \equiv c - \frac{V}{k_{\mathrm{B}}T}. \tag{5.7}$$

This relation can be fitted to reveal the temperature of the cloud if the potential is known.

5.3.1 Calibration of the external potential

The chosen thermometry requires a very good knowledge of the trapping potential, especially in the wings of the cloud. Often, it is sufficient to characterise the trap by its harmonic frequencies because most of the atoms are located in the centre of the trap where the harmonic approximation is valid. The Boltzmann regime, however, is reached far away from the centre where the harmonic approximation is no longer valid. It is therefore necessary to consider the full trapping potential given by the two intersecting Gaussian beams. Their complete properties cannot be deduced from the trap frequencies because the frequencies are given by (in the directions where only one beam is confining)

$$\omega = \frac{2}{w}\sqrt{\frac{V_0}{m}} \tag{5.8}$$

where w is the waist of the Gaussian beam, V_0 the trap depth, and m the mass of a ^{6}Li atom. The frequency thus only provides the ratio $w/\sqrt{V_0}$ while a full characterisation of the beam would require a separate knowledge of w and V_0.

In principle, the beam properties are roughly known and can also be measured outside the vacuum chamber. However, the remaining uncertainties, exacerbated by the vacuum chamber's windows, are relatively large. It is therefore necessary to measure the potential *in-situ*. No straightforward procedure exists for this so several approaches had to be tested.

Anharmonicities

One approach is to measure the anharmonicity of the trapping potential. This can be done by exciting a sloshing mode of the cloud with varying amplitudes. The larger this amplitude gets the more this mode is affected by the regions of the trap not well described by the harmonic approximation. The

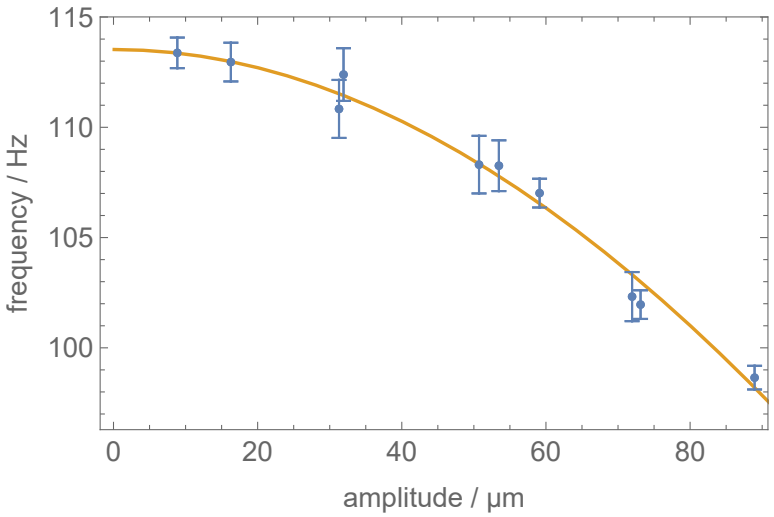

Figure 5.5: Amplitude dependence of the trap frequency. The points are measured with a cloud sloshing in the trap. The line is the simulated motion of a classical particle moving in a Gaussian beam potential whose parameters w and V_0 are fitted to the measured data.

result is an amplitude dependent frequency of the sloshing mode[1]. Since the functional form of the trapping potential is known, the amplitude dependence of the frequency can be measured and compared to numerical simulations of a point or extended mass in a Gaussian potential to characterise the Gaussian beam profiles. While a measurable decrease was observed, see figure 5.5, the achievable amplitudes did not suffice to make w and V_0 independent quantities in the shown fit, see figure 5.6. Larger amplitudes can be excited but lead to a breakup of the cloud and therefore do not provide a meaningful frequency.

Ideal Fermi gas

An alternative is the preparation of an ideal Fermi gas in the trap. It does not interact during expansion and therefore reveals, after an appropriate time-of-flight, a faithful momentum distribution which can be fitted with the momentum distribution in an external potential, equation 2.21, to determine the temperature. The *in-situ* density distribution of the ideal Fermi gas with known temperature is then sufficient to calibrate the potential by the equation of state 2.20. The challenging part of this method is the preparation of an ideal gas at temperatures low enough to achieve detectable densities in the final trap because the non-interacting gas does not thermalise during evaporation. An adiabatic sweep of the Feshbach magnetic field from close to the resonance, where cooling is possible, to $a = 0$ is not possible because it lies on the far end of the BEC side. Ramping the magnetic field there is found to cause strong atom loss, most likely due to collisional decays into low-lying molecular states.

A different path to an ideal gas is a spin-polarised gas where the minority component allows for thermalisation during evaporation. The final gas will then, especially if the Feshbach resonance is tuned to weak attraction, have an outer shell that is almost completely polarised and forms an ideal

[1] Strictly speaking, the movement becomes anharmonic and no longer has a single, well defined frequency. In the experiment, the deviation from a harmonic movement is however too small to be discernible from experimental noise.

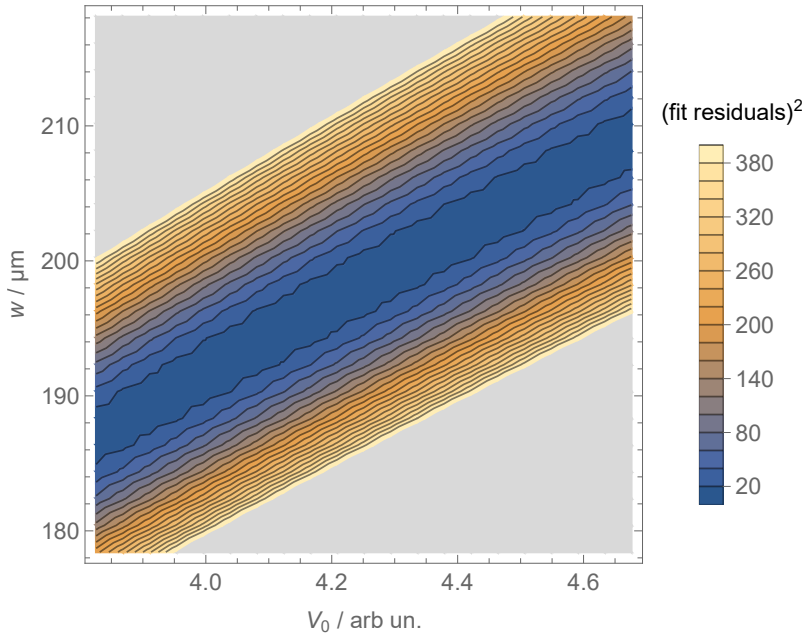

Figure 5.6: The squared residuals of the measured frequencies from figure 5.5 to the calculated ones for a range of combinations of V_0 and w. The extensive valley of a near-constant residual shows that V_0 and w cannot be independently determined from the data. The grey area is clipped away to highlight the extent of the valley.

Fermi gas. In reality this approach did not work well because the fraction of minority atoms necessary for a good evaporation, around 25 % to 30 %, was too large to allow the formation of a substantial fully polarised region.

A alternative approach to achieve full polarisation is the removal of the minority spin component with a pulse of resonant light. The majority component is unaffected by the light directly because the hyperfine splitting is significantly larger than the laser line width. An indirect effect can however arise from the interactions between the spin components and bring the remaining atoms into a non-thermal state. The final momentum distribution of the leftover atoms indeed turned out to be slightly anisotropic, in contrast to the prediction of equation 2.21. The calibration with an ideal Fermi gas was therefore considered to be unreliable.

Unitary equation of state

In the end, a method that brings together the strengths of the experiment with a good theoretical understanding proved to be successful: the creation of unitary Fermi gases and knowledge of their equation of state. If a cloud with a known temperature is prepared, the density equation of state can be inverted to determine the potential from the density distribution in LDA. In general, the temperature is not known but there is one exception. At the phase transition, which is identifiable with the neural networks, the centre of the cloud will be at the critical temperature. $T = T_c = 0.167\,T_F$, with T_F determined from the local density at the centre, is therefore the temperature of the whole cloud. The potential can then be reconstructed by inversion of the equation of state and the result is shown in figure 5.7. It is given as a one-dimensional potential of the radial coordinate r because radial averaging

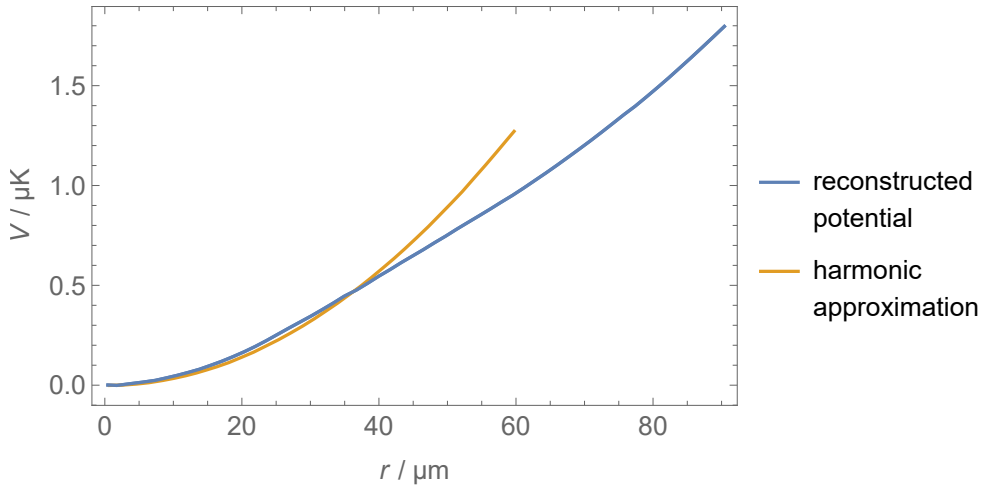

Figure 5.7: The external trapping potential reconstructed from the density distribution at unitarity and T_c. The harmonic approximation of this potential is also shown. Figure adapted from [93].

reduces the influence of experimental noise, see section 5.2.

This method has the conceptual drawback that the critical temperature at unitarity in the phase diagram will not be a measured quantity any more. Instead, it is assigned a fixed value a priori. The critical temperature at all other interaction strengths also relies on this value and is, taken strictly, only a relative measurement. Since both theoretical, see the references in figure 2.9, and experimental, see references in figure 5.1, values for T_c at unitarity agree well with each other, we consider this drawback to be acceptable.

5.3.2 Thermometry of strongly interacting gases

With the reconstructed potential, the temperature of arbitrary clouds can be determined from the Boltzmann wings. Unfortunately, higher order deviations from the linear relation 5.7 are very small just outside the Boltzmann regime as shown in figure 5.8. Data with noise covering only a limited potential range can therefore be wrongly assumed to be in the Boltzmann limit leading to wrong thermometry results.

The solution to this problem is the perturbative inclusion of interactions in the equation of state. This is provided exactly by the virial expansion 5.2 [125]. If the addition of another order in the expansion leaves the fitted temperature unchanged, the chosen equation of state can be considered appropriate. For the thermometry of the phase diagram measurements, it was necessary to include the 3rd order of the virial expansion and even the 4th order in extreme cases furthest on the BCS side.

The necessary virial coefficients are by now well known. For b_2, an analytic formula has been

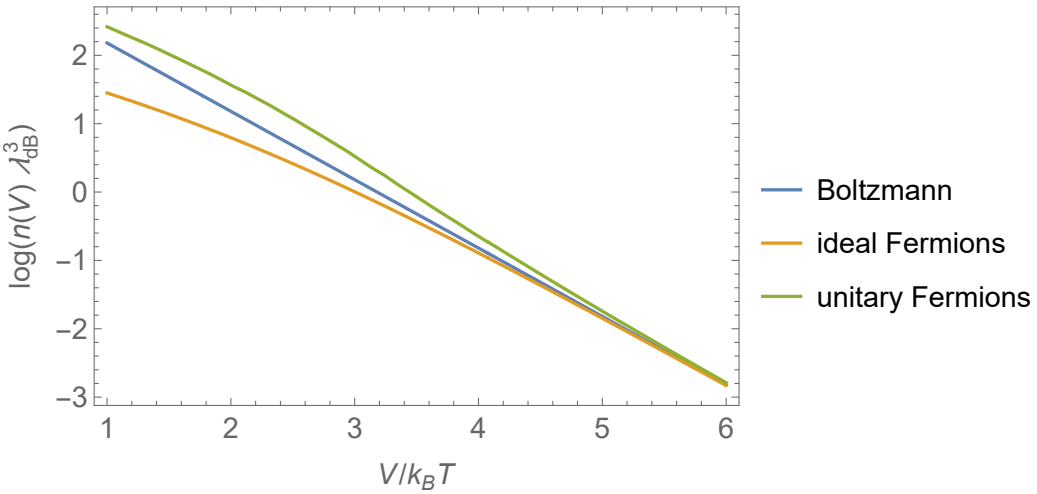

Figure 5.8: The convergence of the equation of state to the Boltzmann limit for ideal fermions and unitary fermions. The Boltzmann limit is approached very gently. It is therefore challenging to identify it unambiguously in real, noisy data.

found long ago[2] [122]

$$b_2(T) = \begin{cases} \frac{e^{\lambda^2}}{\sqrt{2}} \left[1 - \mathrm{erf}\left(|\lambda|\right)\right] - \frac{1}{4\sqrt{2}} & \text{for } \lambda \leq 0 \\ \sqrt{2} e^{\lambda^2} - \frac{e^{\lambda^2}}{\sqrt{2}} \left[1 - \mathrm{erf}\left(\lambda\right)\right] - \frac{1}{4\sqrt{2}} & \text{for } \lambda > 0 \end{cases} \tag{5.9}$$

with a thermal deBroglie length rescaled by the scattering length

$$\lambda = \frac{\lambda_{\mathrm{dB}}}{\sqrt{2\pi} a}. \tag{5.10}$$

For the third order coefficient, the expressions are significantly more involved and different approaches have been used to calculate them [121, 123, 126]. By now, the numerical results for these coefficients agree with each other. [123] also calculated the 4th order virial coefficient needed for the clouds furthest on the BCS-side of the crossover.

5.4 Phase diagram

Now, all the techniques necessary to construct the phase diagram of the homogeneous Fermi gas are together. A detailed discussion of the implementation for most steps is already presented in the PhD thesis of Andreas Kell [93], therefore only an outline of the procedure is given here. A new addition is a closer look at a systematic error, presented at the end of this section.

Experimentally, data are taken identically to the generation of the labelled training set, see section 4.3.1, but covering more interaction strengths at the cost of less repetitions. The number of repetitions can be reduced because the networks are already trained. The range of $1/k_{\mathrm{F}}a$ is limited from below

[2] Valid for zero-range interactions as realised to very high accuracy in ultracold ^{6}Li.

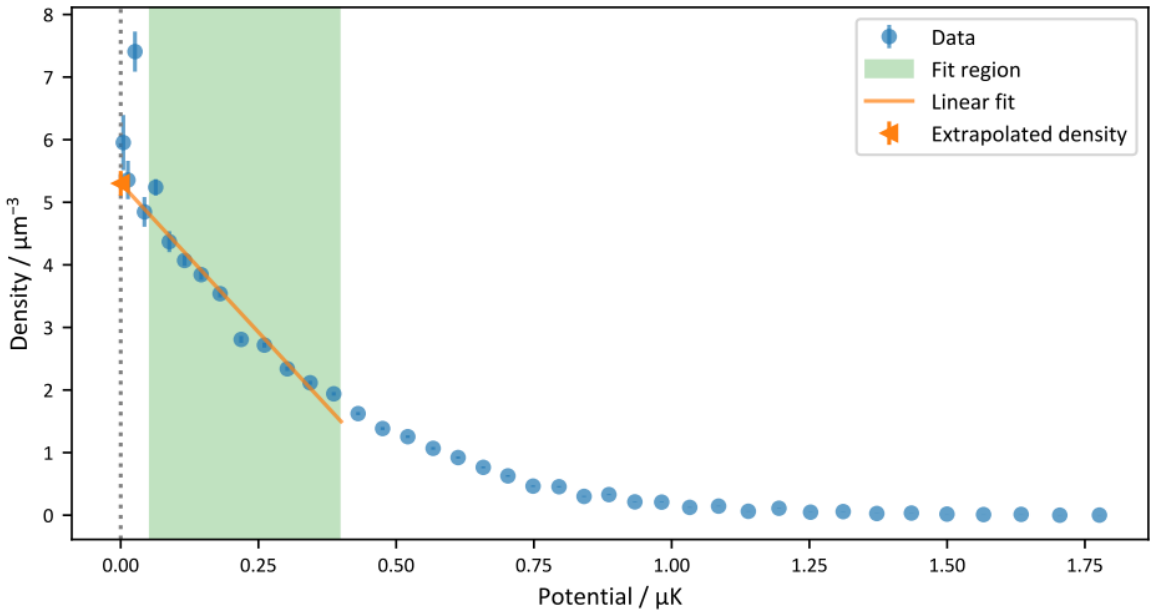

Figure 5.9: Determination of the central density. In the centre of the trap, the reconstructed density scatters too much to give a reliable value. A linear extrapolation is therefore used to determine the central density defining T_F and k_F. This is the density distribution of a gas at unitarity at the lowest achieved temperature. In contrast to the rest of this thesis, this figure shows the density of a single spin component which is half the total density. Figure from [93].

by the condition of being in the superfluid phase after initial preparation of the cloud and T_c quickly decreases in this limit. The highest $1/k_F a$ achievable is limited by atom loss from the trap observed in the BEC limit of the crossover.

For the supervised neural network, which predicts the condensate fraction of a cloud, the critical heating time is found by a piecewise linear fit to the condensate fraction as a function of the heating time in the vicinity of the phase transition. In this way, the heating time at which the condensate vanishes, the critical heating time, is found.

For the autoencoder, the phase boundary is also determined with a piecewise linear fit but this time in the latent space, see section 4.4.3. The phase boundary is converted to a critical heating time by averaging the heating time of the three closest points in latent space. The result of the autoencoder is found to be subject to statistical variations, caused by the randomly sampled training data and the random initialisation of the network together with stochastic gradient descent. 172 equivalent autoencoder networks were therefore trained [113]. Finally, the critical heating time was determined by the weighted average of the different autoencoders' results.

With the critical heating time established, the corresponding T/T_F and $1/k_F a$ have to be found. To this end, the ratio T/T_F is determined for a range of heating times t_{heat} around the phase transition such that T_c/T_F can be calculated from a fitted linear relation $T/T_F(t_{heat})$.

T_F and k_F are calculated from the central density of the cloud which is reconstructed with the inverse Abel transform. Because radial averaging of the recorded absorption images has little effect in the centre of the cloud, the central density is dominated by noise which has an experimental contribution

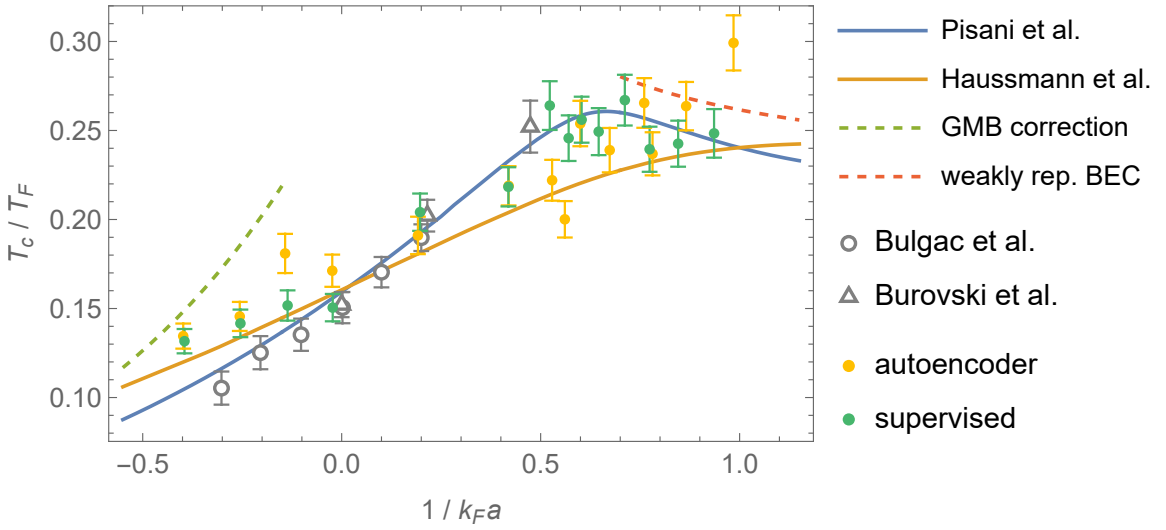

Figure 5.10: Critical temperature of the homogeneous Fermi gas in the BCS-BEC crossover. The results from the supervised analysis as well as from the autoencoder are shown. They are compared with the theoretical predictions from [71] (Pisani et al.), [67] (Haussmann et al.), [72] (Bulgac et al.), [73] (Burovski et al.) and the limiting results of the GMB correction and the weakly repulsive BEC.

from the image and a numeric one from the reconstruction algorithm. It can therefore not be read off but is instead found from a linear fit to the relation $n(V)$, which is then slightly extrapolated to the centre, see figure 5.9. The equation of state at unitarity, see figure 5.2, shows that, at least at this field, the critical density is indeed approached linearly for decreasing V (which corresponds to increasing μ in figure 5.2 due to the local density approximation).

The temperature is determined with the virial expansion of the equation of state as described in section 5.3.2 and a is determined from the magnetic field with equation 3.1.

The resulting critical temperatures are presented in figure 5.10. Both approaches reveal a critical temperature continuously rising from the lowest recorded attraction strength $1/k_\mathrm{F}a \approx -0.4$ up to at least $1/k_\mathrm{F}a = 0.5$. Above $1/k_\mathrm{F}a = 0.5$, T_c obtained from the supervised method levels off and is consistent with a local maximum predicted in this region by [71]. For a confirmation of this maximum, the uncertainty of our measurement is slightly too high. The main source of this uncertainty is discussed below in section 5.4.1.

For the autoencoder the situation above $1/k_\mathrm{F}a = 0.5$ is less clear because the results scatter much stronger. The behaviour is consistent with both a continued increase of T_c as well as a more constant value. This uncertainty reflects that the autoencoder approach is not as specialised for the detection of a condensate as the supervised one.

In general, however, the results from both methods agree well with each other and several theories, especially [71] and the quantum-Monte Carlo results [72, 73] for $1/k_\mathrm{F}a \geq 0$. For $1/k_\mathrm{F}a < 0$, [71, 72] expect a faster drop of T_c than observed but the measured results agree reasonably well with the theory from [67]. It should be noted that previous experiments have also observed relatively high critical temperatures towards the BCS limit of the crossover [39, 40, 114]. Possible reasons could be the trapping potential [115] and the formation of dimers already below $1/k_\mathrm{F}a = 0$ [40]. The rapid

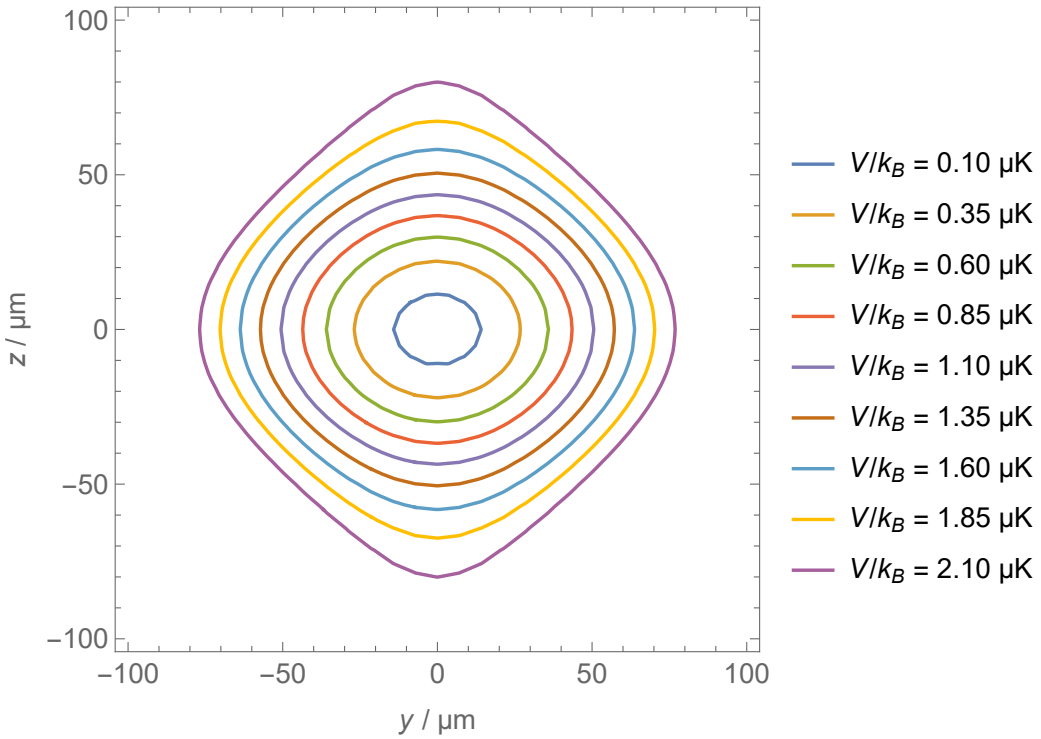

Figure 5.11: Equipotential lines in a cut through a crossed-beam dipole trap with parameters similar to the one in the experiment. Near the centre, the trap has an elliptic symmetry but this is increasingly violated when moving outwards. The two beams are propagating in x- and z-direction.

ramp could in principle also be a reason but if this was the case, the issue should not appear with the autoencoder method, which gives almost identical values for the smallest $1/k_{\mathrm{F}}a$.

5.4.1 Systematic error from the trap geometry

The analysis leading to the phase diagram of figure 5.10 makes extensive use of the inverse Abel transform under the assumption of elliptic symmetry. Figure 5.11 however shows that this assumption is not rigorously true for traps formed by two Gaussian beams. It can be seen that the equipotential and therefore also the equidensity lines deviate increasingly from ellipses the further away they are from the centre. Since the density in these regions is rather low, it is not immediately clear how strong the influence on the trap calibration and thermometry is. To shed some light on this, the full trap reconstruction and thermometry procedure was repeated with simulated density profiles in the potential of figure 5.11.

First, the reliability of the calibration of the external potential is investigated. To this end, the density distribution of a unitary Fermi gas in the trap at the critical temperature, which is chosen arbitrarily to be 160 nK, is calculated. The density is discretised on a grid that mimics the camera's resolution and then summed in one direction to obtain the column density. This procedure imitates the imaging procedure of the real experiment. Just like in the real analysis, radial averaging and the inverse Abel transform are then performed with the column density to reconstruct the full density profile in the trap.

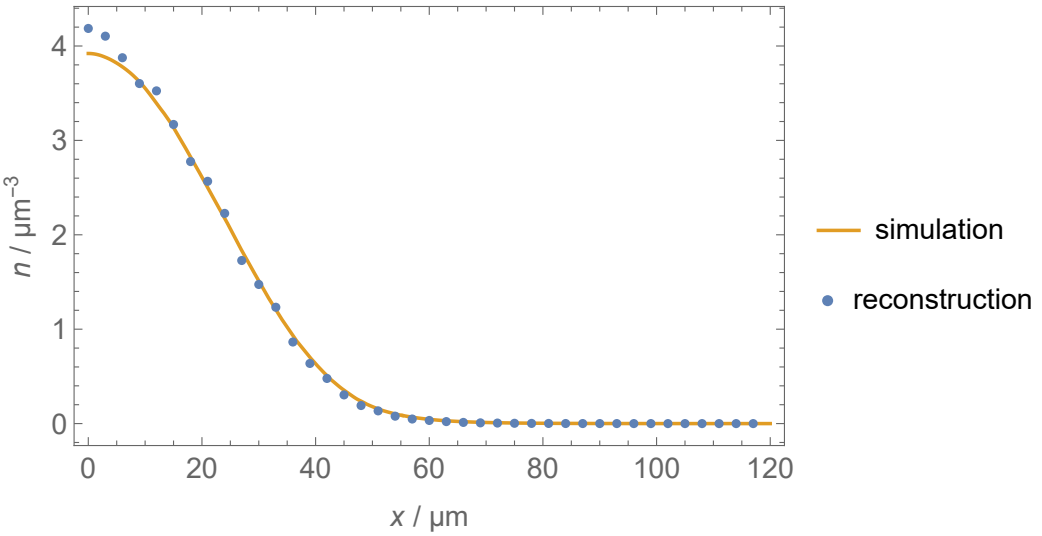

Figure 5.12: Reconstructed density of a unitary Fermi gas at T_c in the trap of figure 5.11. The simulated density distribution is shown for comparison. Note that the linear extrapolation of figure 5.9 would correspond to a parabolic extrapolation in this case where the density is plotted as a function of the position.

Figure 5.12 compares the reconstructed discretised density to the underlying simulated density and shows that the deviations are small except for the centre. The discrepancy in the centre should not have strong consequences because the real experimental results are too noisy to be reliable in this region. The linear extrapolation of the reconstructed density towards the centre in potential space (figure 5.9) becomes a parabolic extrapolation in real space due to the almost harmonic shape of the Gaussian potential in this region. It is found that the extrapolation overestimates the simulated density by around 5 % and therefore k_F and T_F will be overestimated as well. In this way, the temperature of the cloud is overestimated by around 3 % ($T \propto n^{2/3}$). The potential, reconstructed with the inverted equation of state, is finally also overestimated, especially in the wings of the cloud, as shown in figure 5.13.

While figure 5.13 suggests a substantial systematic error for our thermometry, the actual consequences are not so obvious because the data used for thermometry are processed with the same algorithms and therefore subject to the same systematic error. The deviations from the correct potential could very well be compensated by incorrectly reconstructed density distributions. To test this, clouds of the other case with a fully understood equation of state – the ideal Fermi gas – are simulated in the same potential. Their reconstructed density is then calculated in the same manner. The reconstructed density as a function of the reconstructed potential is then used for thermometry with different orders of the virial expansion. For the non-interacting case, the virial coefficients are given by $b_n^{(0)} = (-1)^{n+1} n^{-5/2}$ [123]. An example of this thermometry is shown in figure 5.14. The resulting temperatures deviate by less than 5 % from the "real" temperature of the cloud of 180 nK. The same procedure for a larger, hotter cloud with 500 nK gives deviations of no more to 5 % as well.

From these observations, it is concluded that the systematic error introduced by the non-ellipticity of the trapping potential results in a 5 % uncertainty of the critical temperature. This error turns out to dominate the uncertainty of the critical temperature as shown in figure 5.10.

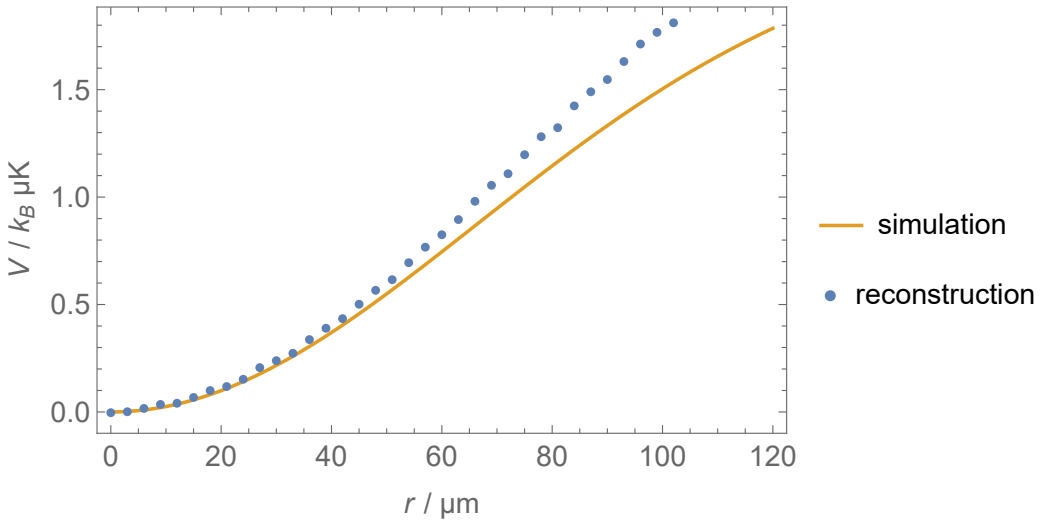

Figure 5.13: The external potential reconstructed from the density of 5.12 compared to the actual potential used for the simulation. Significant deviations are apparent in the wings of the cloud.

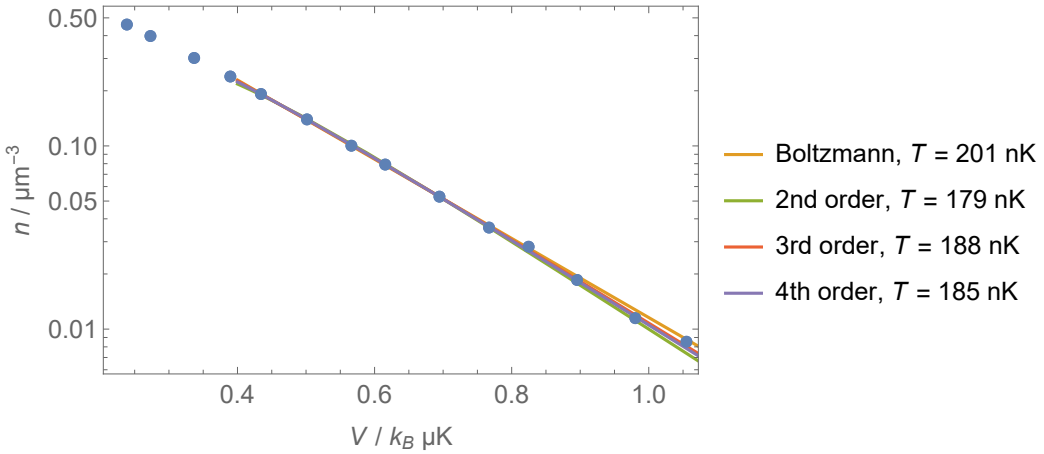

Figure 5.14: Thermometry with the reconstructed density of an ideal Fermi gas in the reconstructed potential. The fitted equations of state for several orders of the virial expansion are shown. The actual temperature used in the simulation to generate the data points was 180 nK.

CHAPTER 6

Interaction Quenches

This chapter presents the experimental results of interaction quenches in superfluids in the BCS-BEC crossover and these are the main results of this thesis. The overall goal of these experiments was to excite the Higgs mode of the superfluid and observe it in time domain. The oscillation of the order parameter Δ [64] can be detected as an oscillation of the condensate fraction and this has not been observed before. Quantitatively, the oscillation frequency is related to the order parameter and therefore an option to measure Δ [58]. It is also interesting to investigate the stability of the mode which is expected to deteriorate through the crossover towards the BEC limit [60], a consequence of the increasing violation of particle hole symmetry [75]. In this sense, the properties of the Higgs mode can help to classify the underlying system, as has been done recently for the unconventional charge density wave in RTe_3 [127].

Below, it will be shown that dynamics of the condensate fraction resembling the Higgs mode are indeed observed after an interaction quench. They are however strongly damped. The behaviour shows a strong dependence on the interaction parameter $1/k_{\mathrm{F}}a$ as expected for phenomena connected to Δ. A decreasing lifetime of the dynamics when moving from the BCS to the BEC regime is observed as well.

Previously, the Higgs mode in the BCS-BEC crossover has only been observed spectroscopically and with limited resolution in earlier studies in our group [57]. In a related system, the precursor of the many-body Higgs mode in a mesoscopic, two-dimensional fermionic superfluid has also been observed [128]. In superconductors, observations have long been elusive due to difficulties in exciting the mode. In $NbSe_2$, superconductivity coexists with a charge density wave and this allowed for the Higgs mode to be observed in Raman scattering [129]. With new developments in THz technology, the Higgs mode was detected in time domain in a BCS- [130] and a cuprate superconductor [131]. Examples of other observed Higgs modes are in magnetic systems [132], cold atom implementations of the Bose-Hubbard model [133, 134], and, most famously, the Higgs Boson [135, 136].

6.1 Experimental implementation of quenches

To study interaction quenches, a cloud of superfluid Lithium in a balanced $|13\rangle$-mixture is prepared in the crossed-beam dipole trap with trapping frequencies $(\omega_x, \omega_y, \omega_z) = 2\pi \cdot (137\,\text{Hz}, 214\,\text{Hz}, 85\,\text{Hz})$. The Feshbach magnetic field is then ramped adiabatically to the value corresponding to the scattering length a_f which will be the scattering length after the quench. The preparation of the quench concludes by slowly increasing the current through the quench coil from 0 to $I_\text{quench} \leq 40\,\text{A}$ over a time of 240 ms. The magnetic field of up to 35 G created in this way is either added or subtracted from the Feshbach magnetic field, depending on the current direction in the quench coil. The combined magnetic field of Feshbach and quench coils then sets the scattering length prior to the quench a_i.

Since the magnetic field of the quench coil has a small gradient at the atoms' position, the ramping of this field adds a tilt to the trapping potential leading to atom loss. We limit this loss by adding an additional magnetic field gradient, which compensates the quench coil gradient. This is done with a free layer of the Feshbach coils, connected in anti-Helmholtz configuration [80]. In the case of the quench coil field being added to the Feshbach field, the compensation gradient can be tuned to optimise atom number and condensate fraction, a measure for temperature, simultaneously. In the opposite case of a quench coil field subtracted from the Feshbach field, a trade-off between high atom number and low temperature is found by tuning the compensation gradient. This indicates that the tilted trap implements an efficient evaporative cooling scheme, an effect that has already been observed in [137]. In this case, we use the highest possible condensate fraction that still gives a good signal to noise ratio in absorption imaging. In most cases, the atom number is around $N = 500\,000$.

At the final interaction strength, the cloud is allowed to equilibrate for 60 ms before the current through the quench coil is suddenly switched off, see section 3.4 for details. Because the switching time of less than 3 µs [92] is shorter than the Fermi time t_F, usually around 7 µs, the quench is considered to be instantaneous.

The interaction parameter $1/k_\text{F}a$ decreases monotonically with the magnetic field strength around the Feshbach resonance, see equation 3.1. This means that the configuration in which the quench coil field is added to the Feshbach field realises a quench from weaker to stronger attraction between the fermions. On the other hand, a subtracted quench coil field realises the opposite quench. The $|13\rangle$-mixture of states was chosen because the corresponding Feshbach resonance has the narrowest width of the three possible hyperfine combinations. This choice therefore maximises the achievable quench strengths, which are limited by the current through the quench coil.

The dynamics following a quench are observed by letting the cloud evolve for a variable time t_hold in the dipole trap before imaging. Most of the time we are interested in the condensate fraction of the cloud which is determined with the rapid ramp method.

6.2 Short time dynamics after a quench

Since the expected frequency of the Higgs mode $\omega_\text{H} = 2\Delta/\hbar$ is on the order of the Fermi frequency ω_F in the strongly interacting regime and $\omega_\text{F} \sim 2\pi \cdot 15\,\text{kHz}$, we study the dynamics after a quench for hold times up to a few 100 µs. The observable signal of the mode is potentially rather small so roughly 50 repetitions for every hold time are averaged. Measuring 50 different hold times, the the 50 repetitions take around 17 hours making the stability of the experiment on short and long

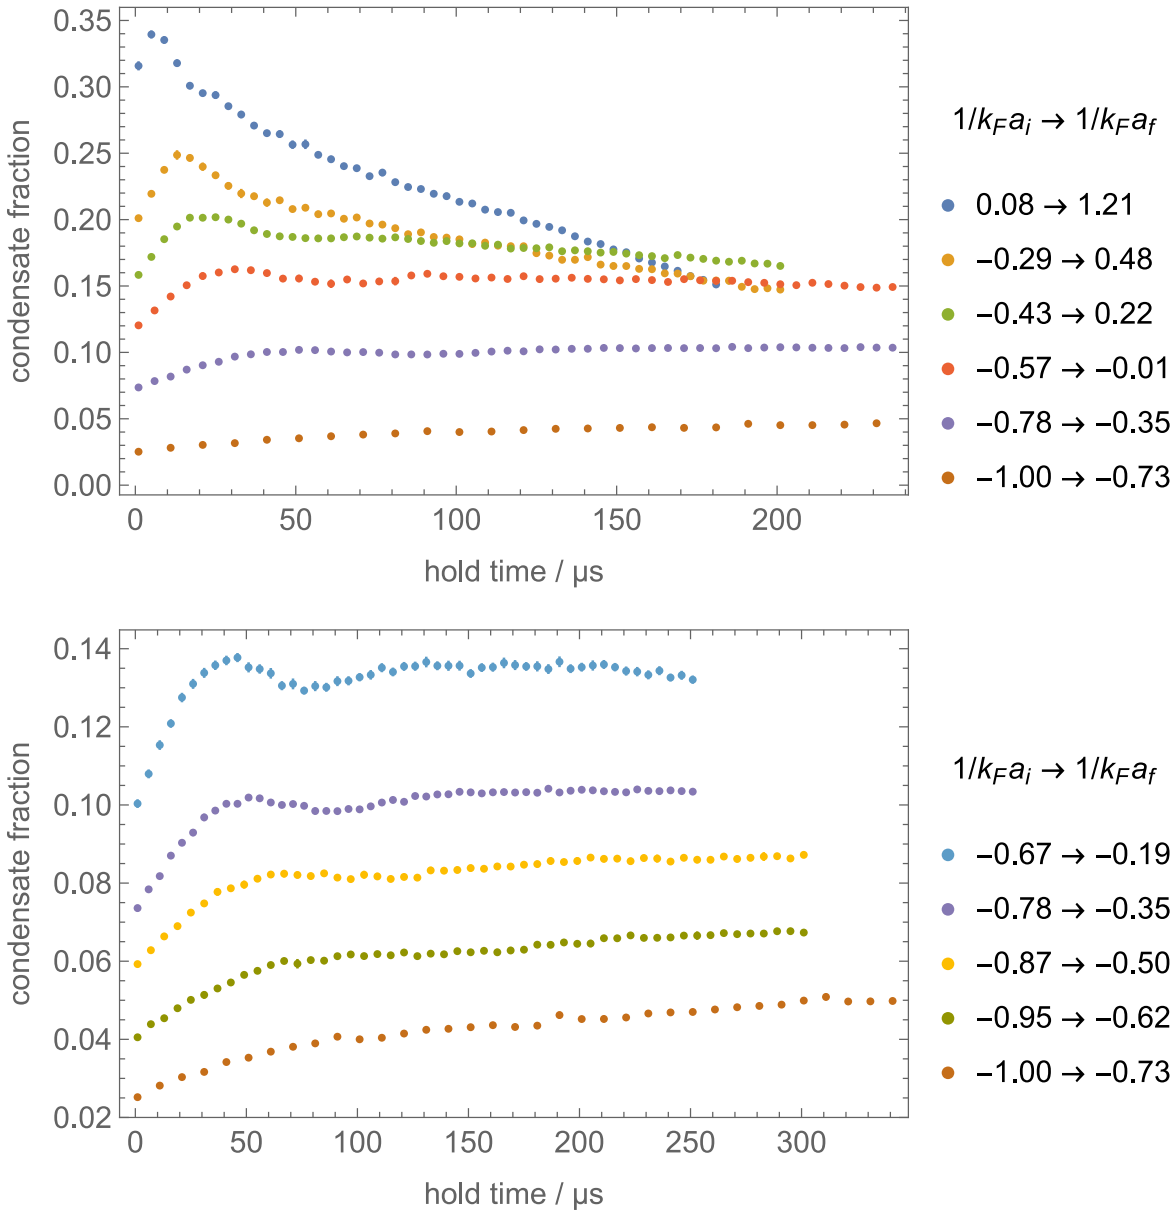

Figure 6.1: Examples of the observed dynamics in the condensate fraction after a quench of 35 G for a vast range of interaction strengths. k_F is determined from the Fermi energy in harmonic approximation. The error bars roughly equal the size of the points and indicate the standard error of the mean of around 50 repetitions. The lower panel contains only measurements with low condensate fraction to improve the visibility of the dynamics.

timescales a concern. To counter the influence of theses effects, the different hold times are recorded in a random order, which is reshuffled after every hold time has been recorded once. Additionally, it was ensured that atom number and condensate fraction show no strong drifts during the measurement.

Examples of the observed time evolution for quenches from weak to strong attractions with the maximal achievable quench strength of 35 G are shown in figure 6.1. For almost all interaction strengths, the condensate fraction rises immediately after the quench until it reaches a maximum; it is reached faster for stronger attractions and usually within the first 10s of µs. The subsequent evolution of the condensate fraction depends on the attraction strength and two qualitatively different regimes are observed. For final interaction parameters $1/k_\mathrm{F}a_\mathrm{f} \lesssim 0.25$ (calculated from the Fermi energy in harmonic approximation), the condensate fraction drops to a local minimum and then equilibrates to a final value larger than at the beginning of the quench, potentially with a faint second maximum. For final interaction parameters $1/k_\mathrm{F}a_\mathrm{f} \gtrsim 0.25$, the condensate fraction only decreases from the first maximum onwards. This firstly happens at a fast rate, comparable to the initial rise, and then slows down significantly with the slow decay lasting over the whole measurement range with maximal hold times between 150 µs and 200 µs. In general, the two regimes are not separated by a clear transition but rather connected by a smooth evolution of the qualitative behaviour. Overall, it is observed that the fast dynamics of initial rise, decay and potentially minimum have a clear dependence on the attraction strength and speed up considerably with increased attraction.

A unique behaviour is found for the quench closest to the BCS limit with $1/k_\mathrm{F}a_\mathrm{f} = -0.73$. In this case, no distinct dynamics are observed on a 10 µs timescale. Instead, the condensate fraction grows in a rather conventional manner for the whole observation range of 500 µs (not shown in its entirety in figure 6.1). Since the condensate fraction in this case is very small for all hold times, it is possible that the faster dynamics are not pronounced enough to be detectable.

As a working hypothesis for the following analysis, it is assumed that the initial fast dynamics with maximum and (potentially) minimum are connected to the Higgs mode of the superfluid, which is strongly damped, while the slower decay or equilibration for long hold times has a different cause, for example the external potential.

6.2.1 Quantifying the observations

For a quantitative analysis of the potential Higgs mode, numerical values have to be extracted from the observed data and this is most conveniently done with a fit. Since outside the regime of very weak interactions and homogeneous density [64] no simple functional prediction for the dynamics exists apart from the general expectation of potentially damped oscillations [138], the fit function has to be of empirical nature.

The Higgs mode is incorporated into the fit function as a simple harmonic oscillation with exponentially damped amplitude

$$f_\mathrm{H}(t) = A\, e^{-t/\tau_\mathrm{H}} \cos\left(\omega_\mathrm{H} t + \phi\right) \tag{6.1}$$

where A, τ_H, ω_H, are respectively the amplitude, lifetime, frequency of the Higgs mode and ϕ is the phase of the oscillation. In general, the mode is not necessarily harmonic [139] but it seems

unlikely that a deviation could be identified with the data at hand. Similarly, the damping is predicted to follow a power law $t^{-\kappa}$ with interaction-dependent κ instead of being exponential [140, 141] but this difference is in general difficult to observe. Additionally, the power law decay is a prediction for the homogeneous gas and not necessarily correct for our case. On top of this, the power law decay introduces a singularity for $t \to 0$ that would have to be dealt with.

The optimal choice for the slower background dynamics of the condensate fraction is less obvious because it is most likely determined by a combination of several different effects. Namely, the increase of the background condensate fraction from its value before the quench to a new equilibrium value, a possible "heating" effect when the cloud evolves towards a thermal state from the non-equilibrated state just after the quench, and the onset of collective modes in the trapping potential. These are caused by the stronger attraction which leads to a shrinking of the cloud and also a change in the external potential that happens when the quench coil field is switched off.

Inspired by the theoretical results of [138, 142], an initial rise is incorporated into the fit function. It is roughly supposed to model the equilibration towards the higher condensate fraction associated with stronger attraction. After a time $t_{\max}$ has elapsed, the background condensate fraction has reached its maximum and the rise is succeeded by a decrease which could be due to heating. Both rise and decay are modelled with an exponential behaviour resulting in the background function $b(t)$ given by

$$b(t) = \begin{cases} \eta_+ \left(1 - e^{-t/\tau_+}\right) + \eta_0 & \text{for } t \leq t_{\max} \\ e^{-(t-t_{\max})/\tau_-} \left(\eta_+ \left(1 - e^{-t_{\max}/\tau_+}\right) + \eta_0\right) & \text{for } t > t_{\max} \end{cases} \tag{6.2}$$

with η_0 the initial condensate fraction, η_+ the increase in condensate fraction approached after the quench with a rise time of τ_+, and τ_- the decay rate for hold times exceeding $t_{\max}$. $b(t)$ is set up to be continuous at $t_{\max}$. Its derivative on the other hand is not continuous at $t_{\max}$. We consider this to be acceptable because a more complicated background function would have to be chosen if we would require $b(t)$ to be continuously differentiable at $t_{\max}$. As a consequence, this would increase the number of free parameters of the fit function which we want to keep low.

The function fitted to the observed dynamics is then given by $f_{\mathrm{H}}(t) + b(t)$ and has 9 free parameters. If either no initial increase or no decay for long hold times are observed, the fit function can still be used because $t_{\max}$ is not restricted to lie within the observed hold times. Figure 6.2 shows that the function can be nicely fit to the observations for all interaction strengths regardless of the qualitative behaviour.

In figure 6.3, the two parts $f_{\mathrm{H}}(t_{\mathrm{hold}})$ and $b(t_{\mathrm{hold}})$ of the fitted function are shown independently for a few selected interaction parameters. It becomes obvious that the chosen $b(t_{\mathrm{hold}})$ is not an optimal choice for the background since the resulting functions are not necessarily smooth. Nevertheless, the present functions allows robust fits for all covered interaction strengths and the extraction of the most interesting parameters ω_{H} and τ_{H}.

Variations on the fit function were explored as well, examples are a linear instead of exponential evolution in $b(t_{\mathrm{hold}})$ or an exclusion of the initial rise from the fit. These gave very similar results but the stability of the fit was often found to be worse and usually included a few interaction strengths where the fit failed. Overall, the chosen function gave the most reliable results for the most interesting parameters ω_{H} and τ_{H}. The choice of the background function will also be revisited below in section 6.3.3.

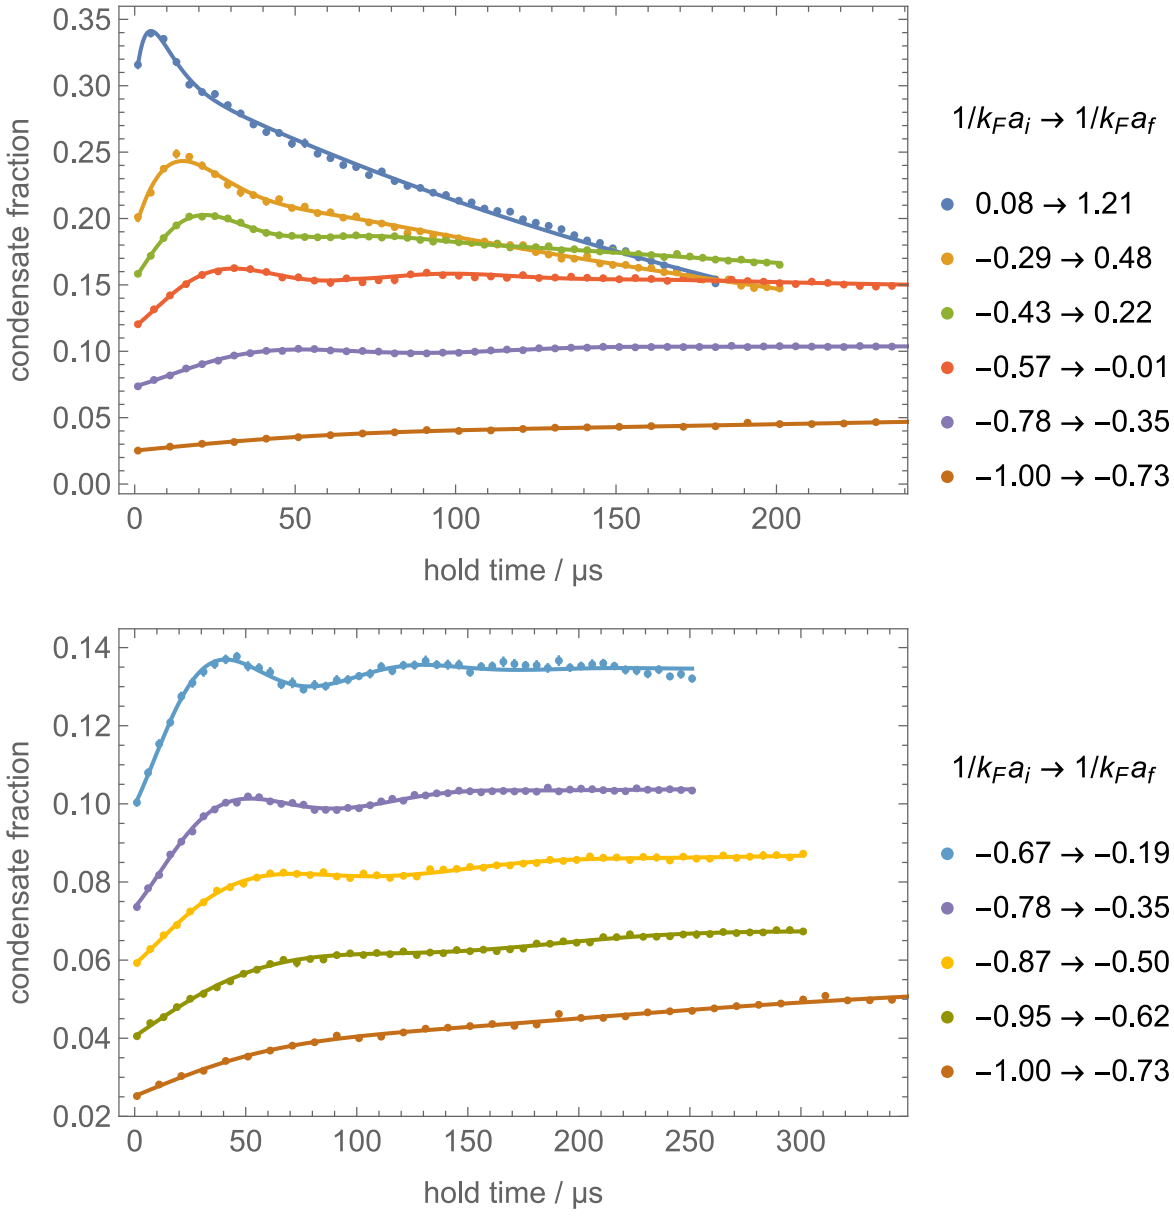

Figure 6.2: The observations from figure 6.1 together with the best fit of the function given by equations 6.1 and 6.2.

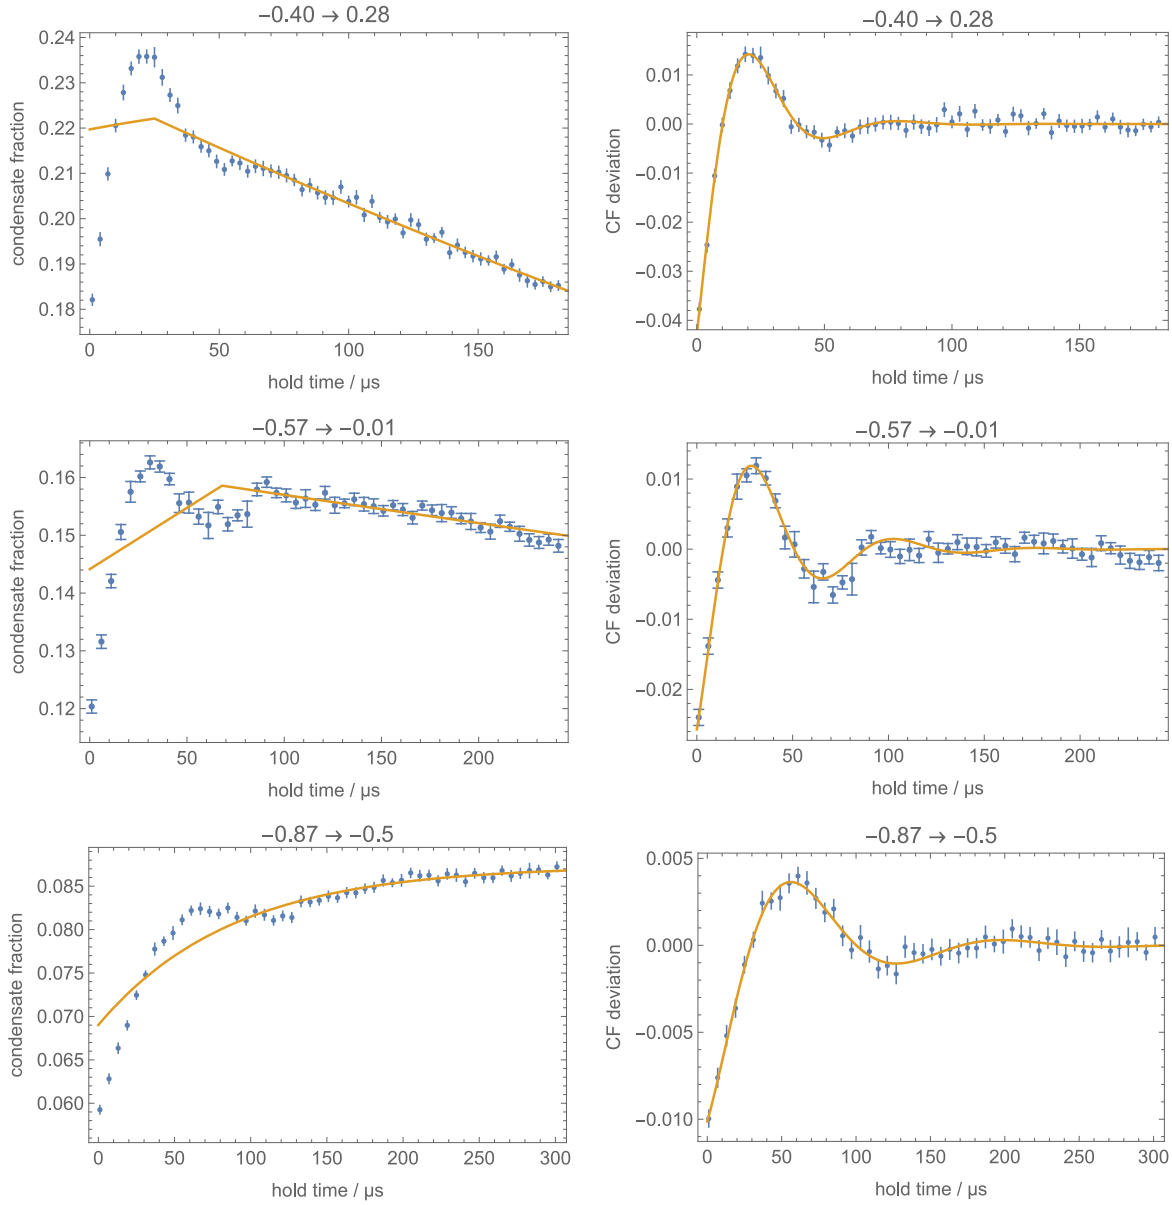

Figure 6.3: The fit function split into oscillation – equation 6.1, right panels – and background – equation 6.2, left panels – parts. The fitted background has been subtracted from the experimental data in the right panels. The numbers above each plot indicate $1/k_{\mathrm{F}}a_{\mathrm{i}} \rightarrow 1/k_{\mathrm{F}}a_{\mathrm{f}}$ in harmonic approximation. The quality of the background fits and their impact is discussed in the main text.

A closer look at the fitted oscillations in the right panels of figure 6.3 reveals that the oscillation's negative peak is not reached at $t_{\mathrm{hold}} = 0$ as naively expected but at small negative times of $\sim 10\,\mu\mathrm{s}$. This could be an indication for either non-harmonic behaviour, like the growth starting abruptly instead of smoothly after the quench, or an underestimated hold time. The latter is possible because the extinguishing of the dipole trap and initial expansion of the cloud take a finite amount of time that is not accounted for in t_{hold}. During these steps, the cloud however is probably still at a high enough density to keep the dynamics going. This leads to effectively longer hold times than indicated in the figures.

The fitted frequencies are in the range of 5 kHz to 18 kHz but become unreliable towards the BEC limit where no local minimum is observed. The fitted lifetimes are, where the oscillation can be fitted reliably, roughly between 10 µs and 50 µs.

6.2.2 Comparison with theoretical predictions

The comparison with theoretical results is difficult since no reliable predictions for our quench experiments in a trapped Fermi gas exist. Already the observability of the Higgs mode is disputed and ranges from observable [143] over an unclear case [142] to unobservable [144]. On top of this, the underlying theories are either of mean-field nature and therefore not accurate in the strongly interacting regime or use the local-density approximation whose validity in a non-equilibrium situation is unknown. In general, it is unclear how the Higgs mode is best described in an inhomogeneous system. Nevertheless, if the observed dynamics are connected to the Higgs mode, the fitted frequency should be on the order of twice the energy gap after the quench $\omega_{\mathrm{H}} \approx 2|\Delta|/\hbar$. This value can in general be smaller than the equilibrium value at the final interaction parameter $1/k_{\mathrm{F}}a_{\mathrm{f}}$ [58, 143]. In a trapped system, Δ can also be considered as a local inhomogeneous quantity making the relation to an observable frequency less obvious.

As a first test, we therefore compare the fitted frequencies ω_{H} to the predicted values for a homogeneous system. If the Higgs mode in the inhomogeneous case is qualitatively similar to the homogenous one, these frequencies should follow at least roughly similar trends. This comparison is therefore worthwhile even if it has no quantitative relevance. For the experimental frequencies, the Fermi frequency is chosen to be its value from the harmonic approximation $\hbar\omega_{\mathrm{F}} = \hbar\bar{\omega}(3N)^{1/3}$ while the Fermi frequency for the theoretical result is $\hbar\omega_{\mathrm{F}} = \hbar^2(3\pi^2 n)^{2/3}/2m$. Only measurements with $1/k_{\mathrm{F}}a_{\mathrm{f}} < 0.4$ are included in the comparison because the frequency could not be determined reliably for higher values of $1/k_{\mathrm{F}}a_{\mathrm{f}}$, see above. The comparison is presented in figure 6.4 and, surprisingly, the predictions match rather well with the observations; the variational many body approach [67] agrees for all observed interaction strengths and the renormalization group approach [145] for negative a_{f} and up to unitarity. The good agreement must however be considered to be accidental because several discrepancies are expected: Firstly, the theory does not account for the increased density in the trap centre that already led to a higher critical temperature in section 5.1.1. It therefore probably underestimates the equilibrium values of Δ and $\omega_{\mathrm{H}} = 2\Delta/\hbar$. Secondly, the harmonic approximation is an idealisation of the experiment that probably overestimates the Fermi energy because the Gaussian trap is shallower than its harmonic approximation. Lastly, there might be a reduction in the observed frequency due to a finite quench strength [143]. All of these effects modify the ratio Δ/E_{F}, might compensate each other to a certain degree and could together cause the good agreement with a not

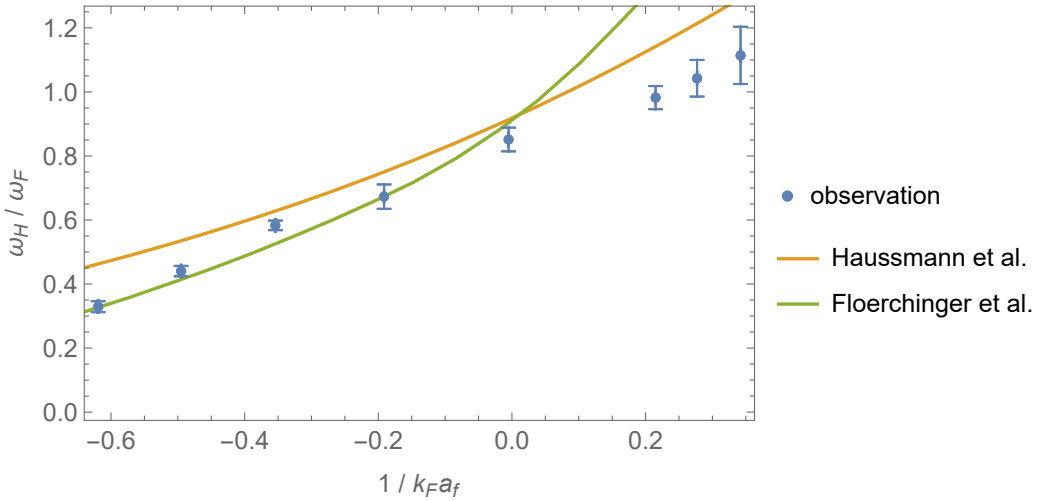

Figure 6.4: Comparison of the fitted Higgs frequencies with the predictions of homogeneous theories. Because the theories do not describe the inhomogeneous system realised in the experiment, at most qualitative agreement can be expected, see main text. The experimental data are normalised by the harmonic Fermi frequency. "Haussmann et al." from [67] and "Floerchinger et al." from [145].

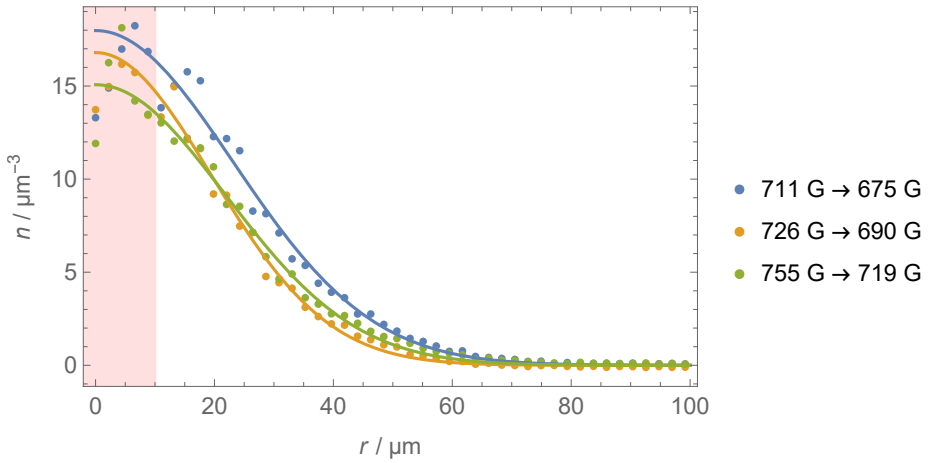

Figure 6.5: Reconstructed densities for selected quenches (points), 690 G is the setting for a unitary Fermi gas; 719 G (675 G) corresponds to a negative (positive) scattering length. The reconstructed densities are approximated by a Gaussian fit which does not take the noisy central region (shaded area) into account.

fully applicable theory.

Of highest concern out of these effects is the increased density due to the attraction. This effect is expected to intensify with $1/k_\mathrm{F}a$ and leads to higher *local* Fermi frequencies ω_F which would accelerate the general dynamics of the system and therefore also the fitted ω_H. In this sense, it cannot be fully excluded that the observed acceleration of the damped oscillation is not due to an increasing gap in the system Δ/E_F but instead caused by a higher ω_F only. To check this, the density distributions of the clouds are reconstructed with the inverse Abel transform, see section 5.2. Several examples of density distributions are shown in figure 6.5. To reduce experimental noise, the absorption images

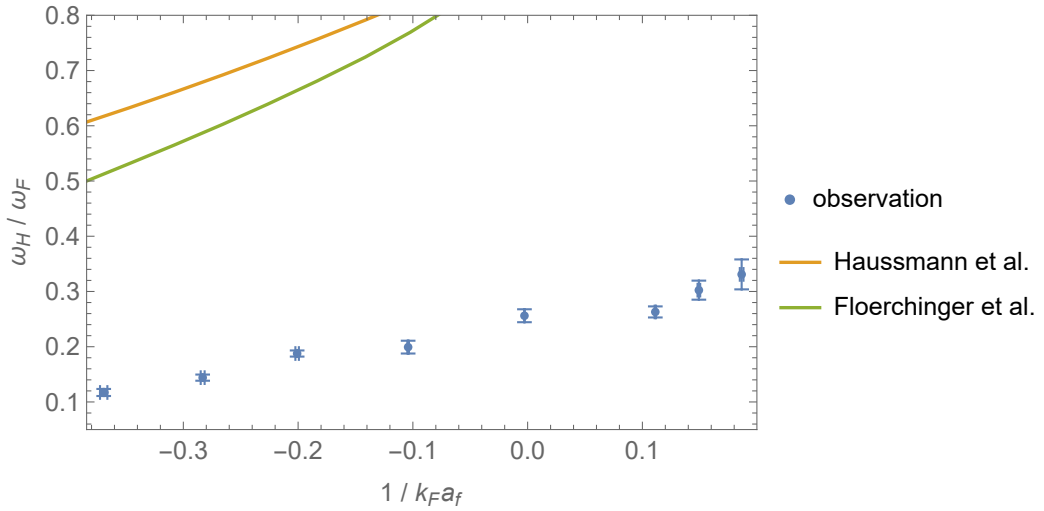

Figure 6.6: Comparison of the fitted Higgs frequency with the prediction of homogeneous theories. In contrast to figure 6.4, the fitted frequencies are normalised by the Fermi frequency obtained from the central density of the clouds. No quantitative agreement between observation and theory is expected here either, see main text.

have again been radially averaged leading to a one-dimensional density distribution $n(r)$. The full three-dimensional distribution $n(\boldsymbol{r})$ can be reconstructed by scaling $n(r)$ to reproduce the appropriate size of the elliptic cloud in every direction. For further calculations, $n(r)$ is approximated by a centred Gaussian fitted to the data. For the fit, the central region of the cloud was not taken into account because radial averaging has little effect causing experimental noise to still dominate there.

The frequencies fitted to the experimentally observed oscillation can then be normalised by the Fermi frequency obtained from the central density of the corresponding clouds. These values are shown in figure 6.6 and an increase of $\omega_\mathrm{H}/\omega_\mathrm{F}$ with $1/k_\mathrm{F}a$ (k_F also being calculated from the central density) is still obvious. The accelerating dynamics with attraction strength are therefore not only caused by an increased density but must have another reason as well. Quantitatively, there is a significant disagreement to the Higgs frequency predicted for a homogeneous gas with the central density. This is however expected since the quench affects the whole cloud and there is little reason to believe that the following dynamics are completely governed by a local quantity determined from one point in the cloud.

Altogether, the qualitative observation of accelerating dynamics with attraction strength supports the assumption that the observed phenomena are indeed connected to the Higgs mode but further investigation is necessary to shed light on the quantitative characteristics.

6.2.3 Pragmatic modelling of the inhomogeneous Higgs mode

The local-density model of [144] predicts, for a quench to unitarity with finite ramp speed, a result that is very similar to our observation: An initial rise followed by a possible faint oscillation. The (first) local minimum in the condensate fraction is reached after roughly $5t_\mathrm{F}$, the Fermi time in the harmonic oscillator, which is $t_\mathrm{F} = \hbar/E_\mathrm{F} \approx 10\,\mu\mathrm{s}$ in our case. Since we observe the minimum after a similar time, around 75 µs, this model is of high interest for our results. The general assumption of the model is that in the inhomogeneous system, many independent Higgs modes are excited and

evolve as their homogeneous analogues depending on the local density values of order parameter and Fermi energy. The resulting, observable dynamics are then characterised by the dephasing of the individual modes. Since we are able to measure the density in our clouds, we can perform this calculation, at least in a simplified form, with our parameters and compare it to the observations. The main motivation of this approach is to figure out whether the observed dynamics can be explained by a local density approximation or whether a more advanced calculation would be necessary. Especially the lower value of the observed ω_{H} compared to the expectation from the central density, see figure 6.6, is of interest in this regard because the model considers the whole density distribution of the cloud and can give an estimate for the strength of this reduction.

Several approximations will be necessary to apply the calculation to our experiment and it can therefore not be considered as a high precision calculation. While the exact implications of the approximations are difficult to judge, it is nevertheless expected that the results will have a quantitative and not only qualitative relevance.

First, the model will be discussed in a relatively abstract way to introduce the underlying reasoning. The precise implementation with its connection to measured quantities will follow afterwards.

For the original calculation in [144], the dynamics of the homogeneous Higgs modes are calculated with the time-dependent Bogoliubov-de Gennes equations, see also [58] and references therein. Since these equations have a certain degree of complexity, we make use of the already known result: a quench of the interaction strength induces oscillations of the local order parameter with frequency $2\Delta(\boldsymbol{r})/\hbar$ or $2\sqrt{\Delta^2(\boldsymbol{r}) + \mu^2(\boldsymbol{r})}/\hbar$ if $\mu < 0$ [141]. A significantly simplified description of the system is then achieved by modelling the Higgs mode as a large number of independent "oscillators" whose collective signal is measured. The oscillators may have different frequencies and damping but have the same phase due to the common excitation from the quench. In frequency space, the full spectrum of these oscillators $\mathscr{S}(\omega)$ is then simply given by the sum of all N individual spectra

$$\mathscr{S}(\omega) = \sum_{i}^{N} \gamma_i(\omega, \omega_{\mathrm{H}_i}) \tag{6.3}$$

where $\gamma_i(\omega, \omega_{\mathrm{H}_i})$ is the spectrum of the ith oscillator with the (local, homogeneous) Higgs frequency ω_{H_i}. From the density distribution, a relative number of oscillators for every Higgs frequency $N_{\omega_{\mathrm{H}}}$ can be deduced[1], making it possible to sum over frequencies instead of oscillators

$$\mathscr{S}(\omega) = \sum_{\omega_{\mathrm{H}}} N_{\omega_{\mathrm{H}}} \gamma(\omega, \omega_{\mathrm{H}}). \tag{6.4}$$

In this step, it was assumed that identical Higgs frequencies coincide with identical spectra, i.e. Higgs modes in identical configurations always look the same – something we take for granted. The sum is now over the continuous quantity ω_{H}, so it can be converted into the integral

$$\mathscr{S}(\omega) = \int \mathrm{d}\omega_{\mathrm{H}}\, g(\omega_{\mathrm{H}})\, \gamma(\omega, \omega_{\mathrm{H}}) \tag{6.5}$$

[1] Knowledge of the absolute number of oscillators is not necessary because it changes the full spectrum only by an irrelevant prefactor.

with the oscillator density of states

$$g(\omega_\mathrm{H}) = \frac{\partial \bar{N}(\omega_\mathrm{H})}{\partial \omega_\mathrm{H}} \tag{6.6}$$

and $\bar{N}(\omega_\mathrm{H}) = \int_0^{\omega_\mathrm{H}} \mathrm{d}\omega'_\mathrm{H} N_{\omega'_\mathrm{H}}$ the number of oscillators with Higgs frequency up to ω_H. In this way, $g(\omega_\mathrm{H})$ is the number of oscillators with a Higgs frequency in the range from ω_H to $\omega_\mathrm{H} + \mathrm{d}\omega_\mathrm{H}$. To make a connection to the reconstructed density $n(r)$ the chain rule is applied to $g(\omega_\mathrm{H})$

$$g(\omega_\mathrm{H}) = \frac{\partial \bar{N}(\omega_\mathrm{H})}{\partial \omega_\mathrm{H}} = \left.\frac{\partial \bar{N}(r)}{\partial r}\right|_{r=r(\omega_\mathrm{H})} \frac{\partial r(\omega_\mathrm{H})}{\partial \omega_\mathrm{H}}. \tag{6.7}$$

Here, the function $r(\omega_\mathrm{H})$ gives the radius at which the density $n(r)$, which defines the local values of k_F and Δ, leads to a Higgs frequency of ω_H and $\bar{N}(r)$ is the integrated number of oscillators up to the radius r. Because the reconstructed density is radially averaged, the ellipticity of the original cloud has to be considered when integrating the density of oscillators $n_\mathrm{o}(r)$ to obtain $\bar{N}(r) = \int_0^r \mathrm{d}r' n_\mathrm{o}(r') A(r')$ where $A(r)$ is the area of the equidensity surface with oscillator density $n_\mathrm{o}(r)$. The oscillator density $n_\mathrm{o}(r)$ itself has to be derived from the reconstructed atom density $n(r)$ by some model, see below.

Finally, the observable signal in time-domain can be calculated by a Fourier transform of the spectrum determined via equation 6.5. This model does not include an initial rise of the condensate fraction because the excitation of the Higgs mode is not taken into account, in contrast to the original calculation [144]. The slow decay of the condensate fraction observed for some interaction strengths is likewise not covered because concepts like thermalisation or in-trap dynamics are completely missing in this simple model. This is however also the case in [144]. Together, this implies that the dynamics resulting from the model have to be compared with the observations after subtraction of the background which were shown in figure 6.3.

Practical implementation of the model

After the relatively abstract introduction given in the previous section, the model is now presented more concretely and different options for the involved quantities are discussed. The basis for the calculation are always the density distributions $n(r)$ reconstructed with the inverse Abel transform. They define $k_\mathrm{F}(r) = (3\pi^2 n(r))^{1/3}$ and thus cause the interaction parameter $1/k_\mathrm{F}a$ to become a local quantity as well. Its absolute value $|1/k_\mathrm{F}a|$ is monotonically increasing towards the edge of the cloud, see figure 6.7. This shows that these regions are weakly interacting. Additionally, T_F becomes very small for low densities leading to a locally very hot ($T > T_\mathrm{F}$) gas, a behaviour that was used in the thermometry for the phase diagram in section 5.3. These regions are above the local critical temperature, therefore not superfluid and should not contribute to the Higgs mode.

To estimate the extend of the superfluid region, the absolute temperature of the cloud has to be compared to the local critical temperature. In contrast to the complex thermometry of chapter 5, a significantly simpler method, also used in [80], is chosen: By extrapolating the measured condensate fractions further to the BCS side, the interaction strength at which the condensate vanishes can be estimated. At the critical interaction strength, we have $T = T_\mathrm{c}(r = 0)$ and this value can be taken from a theoretical prediction; in this case we choose the extended Gorkov-Melik-Barkhudarov theory from [71] which agrees well with the measured phase diagram of section 5.4. It is then assumed that all clouds have this temperature. This is probably not fully correct because T can change during the adiabatic sweep to the various interaction strengths, but still an improvement compared to a

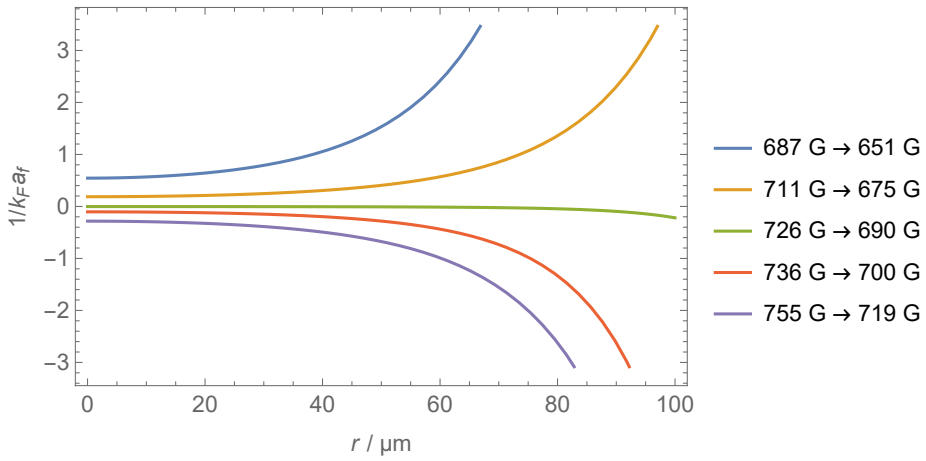

Figure 6.7: Radial dependence of the interaction parameter after the quench for various settings. It is clearly seen that the cloud is weakly interacting, $|1/k_\mathrm{F}a| > 1$, in the regions of low density at high values of r. A special case is the unitary gas at 690 G which has a constant interaction strength across the whole cloud. This is another manifestation of the conceptual simplicity of the unitary gas, see also [45].

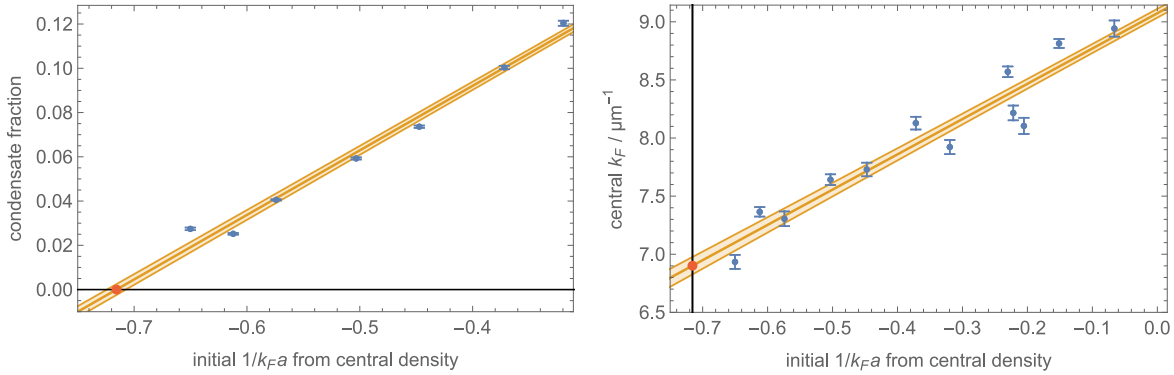

Figure 6.8: Estimation of the temperature. Left: by extrapolating the measured condensate fraction, the interaction strength at which the condensate disappears can be found. At this value of $1/k_\mathrm{F}a_\mathrm{i}$, indicated by the red point, the cloud is at its critical temperature. Right: Extrapolation of $k_\mathrm{F}(r = 0)$ to find its value at the critical interaction strength, again indicated by the red point.

$T = 0$ calculation as done in the original, purely theoretical approach [144]. The critical interaction strength on the BCS-side was determined to be $1/k_\mathrm{F}(r = 0)a_\mathrm{i} = -0.716 \pm 0.008$, see figure 6.8. The predicted critical temperature at this interaction is $T/T_\mathrm{F} = 0.0715 \pm 0.008$. Since the theory describes a homogeneous gas, T_F has to be calculated from the central density at the critical interaction strength because the superfluid appears in the centre first. As a consequence of the interaction-dependent chemical potential, the central density changes with interaction strength and the value at the critical interaction strength has to be found by extrapolation as well because no quench has been measured at the corresponding field. This extrapolation is also presented in figure 6.8 and, by putting everything together, leads to an absolute temperature of the cloud of $T = (137 \pm 4)$ nK.

To find the superfluid region, T has to be compared to the local value of $T_\mathrm{c}(r)$. It can be calculated from the local values $1/k_\mathrm{F}(r)a_\mathrm{f}$ and $T_\mathrm{F}(r)$, using the theory from [71] again. The superfluid region of

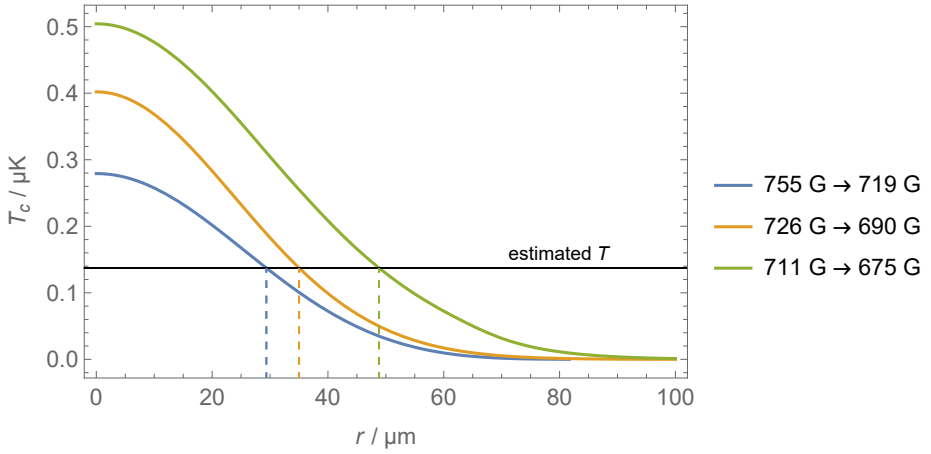

Figure 6.9: Determination of the superfluid region for different quench settings. The local critical temperature depends on $k_F(r)$ and up to a critical radius where $T = T_c$, indicated by the dashed lines, the cloud is in the superfluid phase.

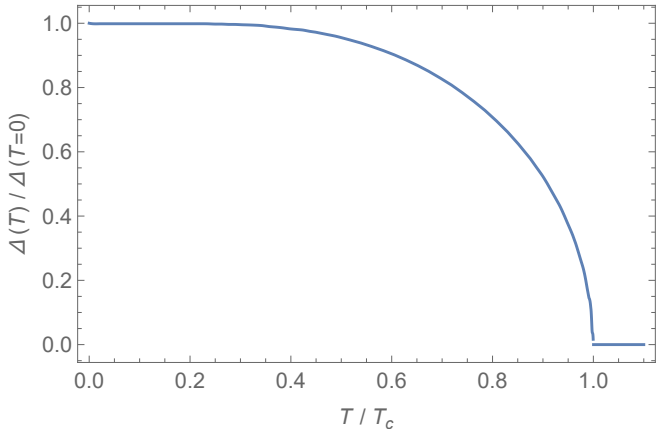

Figure 6.10: Temperature dependence of Δ in the BCS limit as given in [31].

the cloud is then given by $r < r_c$ with $T_c(r_c) = T$, as shown in figure 6.9.

With the region contributing to the Higgs mode settled, the radial dependence of the frequency $\omega_H(r)$ can be considered. Since essentially $\omega_H(r) \propto \Delta(r)$, it depends on $k_F(r)$ in two ways. Firstly, the local interaction strength $1/k_F a_f$ fixes Δ/E_F according to the theory from [67] and, secondly, E_F itself depends on k_F. The combined effect then gives the radial dependence $\omega_H(r)$. For $1/k_F a_f > 0.41$, the chemical potential becomes negative according to [67] and changes the Higgs frequency to $\omega_H(r) = 2\sqrt{\Delta^2(r) + \mu^2(r)}/\hbar$. The values for μ and Δ readily available from [67] are only those for $T = 0$. While these are most likely sufficient for a reasonable description, it would be interesting to also include temperature effects.

Unfortunately, the temperature dependence cannot be described by some simple function. For Δ however, the temperature dependence in the BCS limit is given in [31] and shown in figure 6.10. For the lack of other options, this relation is applied to the whole crossover. In the end, the whole

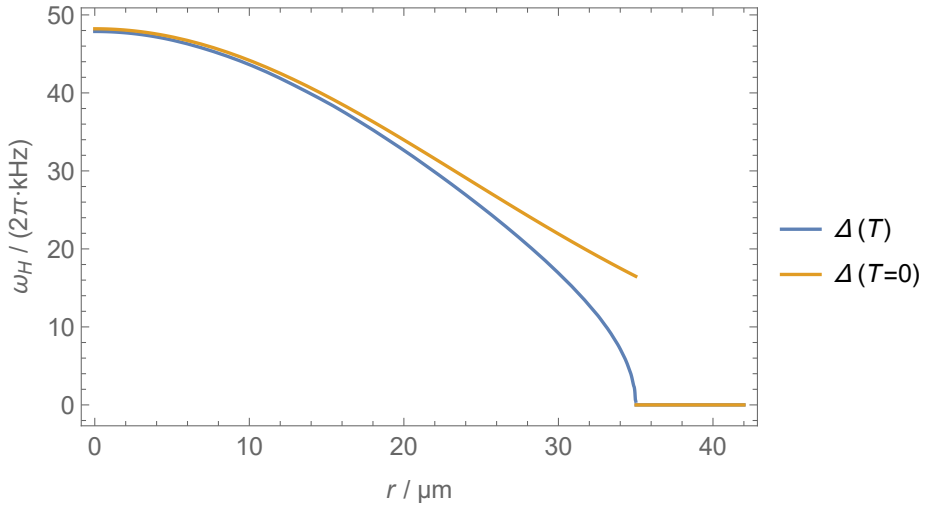

Figure 6.11: The radial dependence of the two options for ω_{H} in the simulation+. The value of Δ can either be taken to be its constant zero-temperature value $\Delta(T = 0)$ or with the BCS temperature dependence factored in $\Delta(T)$. Even in the constant case $\Delta(T = 0)$ the frequency changes because the absolute value of Δ depends on the local Fermi energy. Shown are.

calculation will be done for both the $T = 0$ and finite T value of Δ such that the respective results can be compared. For the temperature dependence of μ, no relation is easily available, except for the non-interacting case. Since $\mu < 0$ is only reached within the superfluid region for quenches relatively far to the BEC-side, where the experimental observations do not allow for a reliable frequency fit, the influence of μ on the Higgs frequency is of limited relevance for a comparison to the experimental results. No temperature dependence is therefore taken into account in the cases where it would in principle be needed. Further, at unitarity, the chemical potential has been measured to be relatively similar to its $T = 0$ value for $T \ll T_{\mathrm{c}}$ and only slightly increasing with $T \to T_{\mathrm{c}}$ [31, 119] – in agreement with the theoretical predictions of [67]. If this behaviour does not change drastically in the crossover region, the temperature-dependence of μ should have relatively little effect on the Higgs frequency.

The resulting spatial dependence of the Higgs frequency for a quench to unitarity is depicted in figure 6.11 including both the $T = 0$ and finite T values of Δ. It becomes apparent that the main effect of the temperature is a smooth interpolation of the Higgs frequency towards $\omega_{\mathrm{H}}(r) = 0$ at the critical radius. ω_{H} derived from $\Delta(T = 0)$ on the other hand has a discontinuity at r_c reflecting that the phase transition has been inserted rather artificially at this point. The derivative of the inverse function of $\omega_{\mathrm{H}}(r)$ is then $\partial r(\omega_{\mathrm{H}})/\partial\omega_{\mathrm{H}}$, the factor needed in equation 6.7 to calculate $g(\omega_{\mathrm{H}})$.

To calculate the full spectrum of the trapped Higgs mode, the relative strength of the varying frequencies has to be found; in this simulation, this is handled by the number of abstract “oscillators”. In the simplest case, their number is just equated to the number of atoms on the elliptic shell with constant $\omega_{\mathrm{H}}(r)$. Under the assumption that the Higgs mode signal is only created by the atoms forming the superfluid, this corresponds to a constant condensate fraction in the whole superfluid region. The precise value of the condensate fraction is irrelevant because it only adds a constant prefactor to the whole spectrum.

While conceptually simple, this approach ignores that the condensate fraction in reality depends on temperature. A possibility to model a varying condensate fraction is to consider Δ as the condensate

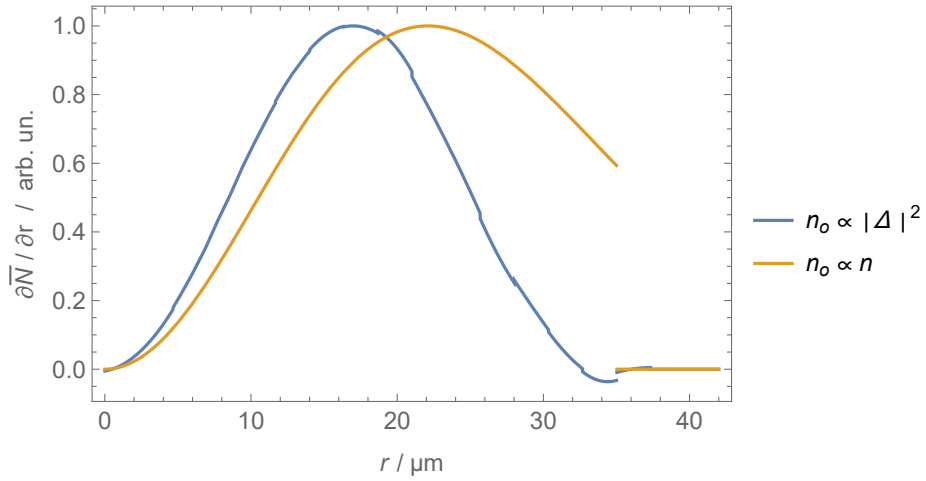

Figure 6.12: The shell density of oscillators $\partial\bar{N}(r)/\partial r$ for a constant ($n_\mathrm{o} \propto n$) or temperature dependent condensate fraction ($n_\mathrm{o} \propto |\Delta|^2$). The shell density is the density of oscillators integrated over an equidensity shell, i.e. the region with identical Higgs frequency. The functions in this plot are for a cloud quenched to unitarity. In the case of a temperature dependent condensate fraction, the shell density is not always continuous and even has some negative values. These artefacts are a consequence of the theoretical predictions for T_c and Δ only being known at discrete values and not as a smooth function.

"wave-function" leading to a condensate density $\propto |\Delta|^2$. The temperature dependent value $\Delta(T/T_\mathrm{c}(r))$ from figure 6.10 is chosen in this case because it smoothly interpolates to a oscillator density of 0 at the phase transition. This approach is in analogy to a Bose-Einstein condensate where the ground state wave function is identical to the order parameter. It can also be shown that $|\Delta|^2$ is indeed proportional to the superfluid density in the BEC limit of the BCS-BEC crossover [146]. In the BCS limit, the condensate density is instead proportional to $|\Delta|$ with a smooth interpolation between the two limits in the crossover. It was however found that for this calculation, both limits result in very similar spectra so only the BEC approximation is considered further. Overall, this leaves two options for the oscillator density $n_\mathrm{o}(r)$: either the atom density $n(r)$ or the "condensate density" $|\Delta|^2$. Of interest for the spectrum is now the quantity $\partial\bar{N}(r)/\partial r$, the oscillator density integrated over the elliptic shell with constant atom density $n(r)$. This quantity, calculated for both options of $n_\mathrm{o}(r)$, is depicted in figure 6.12. Again, the smooth interpolation towards the normal phase with the temperature-dependent condensate fraction is obvious. In this case, the Higgs mode is influenced stronger by inner shells even though they contain less atoms due to the smaller surface area.

The final bit missing to calculate the spectrum 6.5 is the individual oscillator spectrum $\gamma(\omega, \omega_\mathrm{H})$. The Higgs mode oscillates at a well-defined frequency $\omega_\mathrm{H} = 2\Delta/\hbar$ suggesting a Dirac delta line shape $\gamma(\omega, \omega_\mathrm{H}) = \delta(\omega - \omega_\mathrm{H})$ and that is indeed the simplest option. The line shape however also enables us to include the intrinsic decay of the Higgs mode into the simulation. Already in the first paper on the Higgs mode in superfluids [64], the decay was identified to follow a $t^{-1/2}$ power law. A Fourier transform reveals that the corresponding line shape is, in a slightly simplified form, given by

$$\gamma_\mathrm{H}(\omega, \omega_\mathrm{H}) = \sqrt{\frac{\pi}{8\omega_\mathrm{H}\,|\omega - \omega_\mathrm{H}|}}. \tag{6.8}$$

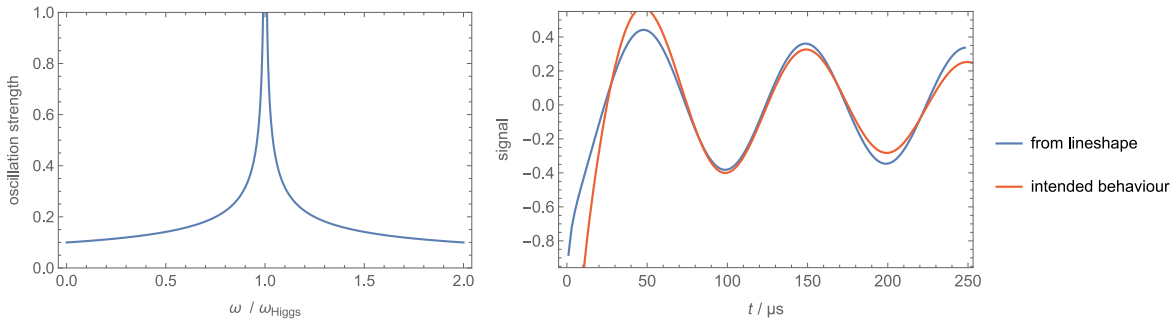

Figure 6.13: Spectrum γ_{H} of the individual oscillators for a power law decay (left) and the signal reconstructed from a sampled spectrum (right).

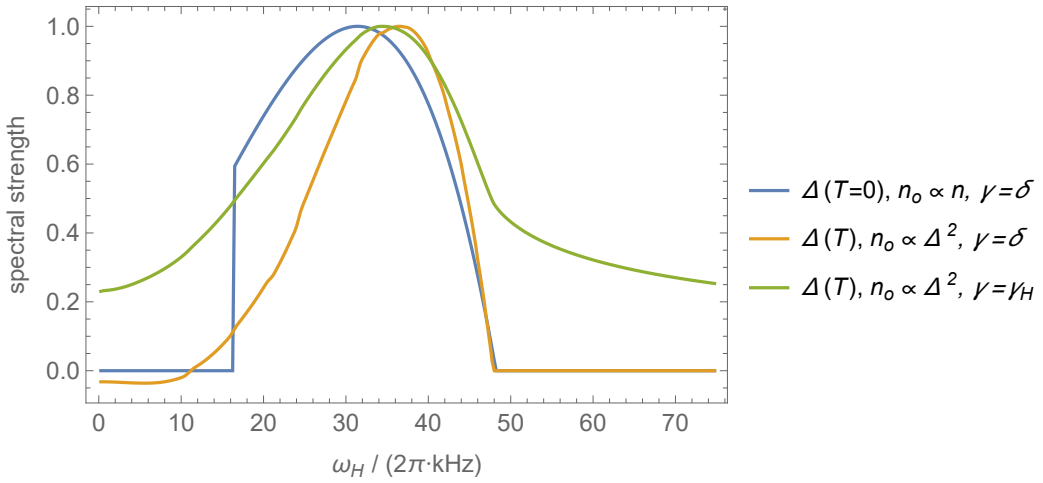

Figure 6.14: Comparison of different ways to calculate the spectrum of the Higgs mode in a cloud at unitarity. In the simplest case (blue), the condensate fraction is constant, the oscillator density is proportional to the atom density and the individual Higgs mode has a Dirac delta line shape leading to a truncated spectrum. If the temperature dependence of gap and condensate fraction is included (orange), the spectrum stays smooth at low frequencies; negative spectral strengths at low frequencies are an artefact of the limited theoretical knowledge, see also figure 6.12. With the decay of the Higgs mode included (green), the spectrum acquires long tails.

This decay is derived in the BCS limit and not expected to perfectly describe the strongly interacting region of the BCS-BEC crossover where the decay becomes faster and approaches the power law $t^{-3/2}$ in the BEC limit [140]. [141] however locates the transition to the faster decay in the region $0.5 \lesssim 1/k_{\mathrm{F}}a \lesssim 1.5$ implying that equation 6.8 describes most of our clouds quite well. Additionally, a $t^{-3/2}$ power law decay has a very broad spectrum which will be difficult to handle numerically. The spectrum of a single oscillator with $t^{-1/2}$ decay is depicted in figure 6.13 and reveals a sharp peak with extremely long tails. When sampling the spectrum for a discrete inverse Fourier transform to reconstruct the time evolution, it is difficult to cover both the peak with high resolution and the entire tails leading to small deviations of the reconstructed time evolution from the intended one, also shown in figure 6.13.

Finally, all components for the spectrum 6.5 are together and the spectra for various combinations of the different options can be calculated. A selection is presented in figure 6.14 and neatly summarises

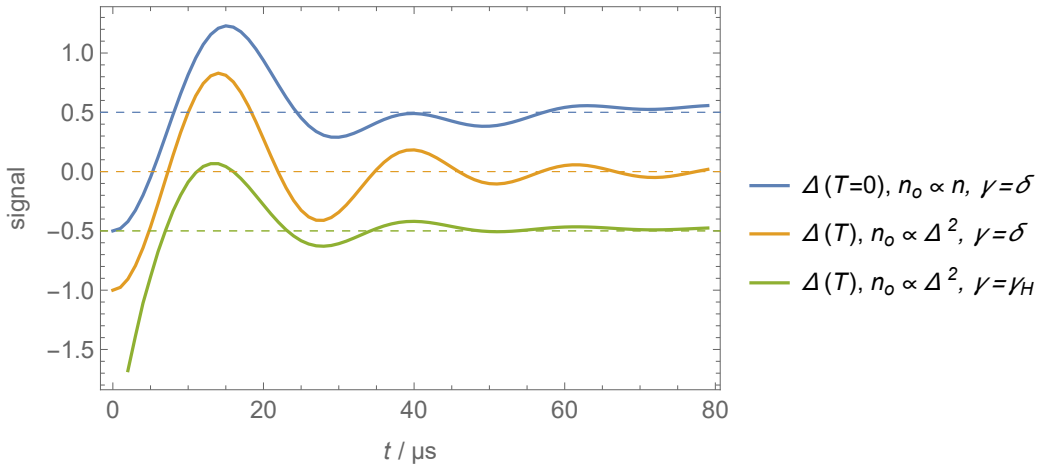

Figure 6.15: Higgs mode oscillations obtained from the spectra for quenches to unitarity shown in figure 6.14. For better visibility, the signals have been shifted by different offsets indicated by the dashed lines. In the experiment, the signal would be an oscillation of the condensate fraction. The background condensate fraction and its possible dynamics are not included in the model.

the previous observations. If condensate fraction and gap are constant and no Higgs decay is included, the spectrum is truncated at low frequencies where the phase transition is enforced. If condensate fraction and gap are chosen to be temperature-dependent, the spectrum becomes smooth on the low frequency side originating from the smooth approach to the phase transition. Also, the position of the maximum shifts to a higher frequency because the outer atoms, where a lower E_F leads to a lower ω_H, have a reduced impact due to a reduced condensate fraction. The inclusion of the Higgs mode's decay finally adds long tails to the spectrum, broadens it and slightly shifts the maximum to lower frequencies.

Results

With the spectra it is now possible to calculate the time domain signal of the Higgs mode in a (simplified) local density approximation. To this end, the spectrum is sampled at intervals of $\mathrm{d}\omega = 2\pi \cdot 250\,\mathrm{Hz}$ with a total of $N = 4000$ samples. The positive and negative frequency parts of the spectrum are chosen with a symmetric phase to enforce a $-\cos(\omega t)$ time evolution for every frequency. The collective time evolution of the Higgs mode is then obtained by a discrete inverse Fourier transform and has a resolution of $\mathrm{d}t = 2\pi/\mathrm{d}\omega\, N = 1\,\mu\mathrm{s}$. The resulting signals for the spectra from figure 6.14 are presented in figure 6.15. For all models, the signal shows a heavily damped oscillation whose second local maximum is already very weak. Especially the model including temperature-dependent gap and condensate fraction and decaying Higgs mode has a striking similarity to the observed signal. It has to be kept in mind that the model only predicts the relative strength of the signal and nothing can be inferred regarding the absolute amplitude of the oscillation in the experiment. It can however be stated that barely detectable first minima are predicted to most likely result in unobservable second maxima, in agreement with the observations. Quantitatively, an important difference is however found: in the calculation for a Higgs mode at unitarity, the first minimum is reached after around 25 µs to 30 µs with a weak dependence on the exact model. This is roughly twice as fast as the experimental

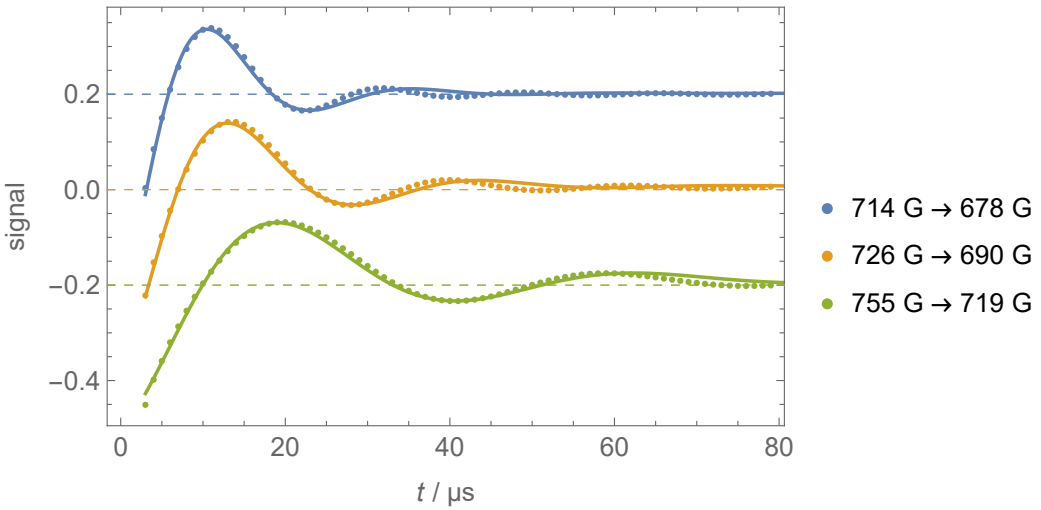

Figure 6.16: Calculated Higgs oscillations for several quench settings – the same which are shown in figure 6.3 – and the corresponding fits. The different settings are separated by an additional offset (dashed lines) for better visibility. All depicted time evolutions are calculated with temperature-dependent gap and condensate fraction and intrinsically decaying Higgs mode.

observation at unitarity, which has a local minimum after approximately 60 µs.

To investigate the quantitative discrepancy further, the calculated signals are fitted with the exponentially decaying oscillation $f_{\mathrm{H}}(t)$ from equation 6.1. It was found that the addition of a constant to $f_{\mathrm{H}}(t)$ leads to a better convergence of the fits in some cases. This is probably connected to the relatively strong low frequency component in the spectra calculated with $\gamma = \gamma_{\mathrm{H}}$. Some examples of these fits are shown in figure 6.16. All depicted time evolutions are calculated with temperature-dependent gap and condensate fraction and decaying Higgs mode; the quench settings are those whose experimental observations were shown in figure 6.3. The calculated dynamics are not perfectly described by the rather simple fit function f_{H} as visible in figure 6.16. Since the calculation is mainly performed to be compared to the experiment, which has a finite resolution, it is expected that these deviations would not be observable with real data. Overall, it can be seen that the Higgs dynamics accelerate from the BCS side of the crossover (high magnetic fields) to the BEC side (low magnetic fields), as expected. The discrepancy in the position of the first minimum between experiment and calculation on the other hand can be seen for all depicted interaction strengths.

The frequencies fitted to the observed and calculated dynamics are compared in figure 6.17. They are normalised to the Fermi frequency calculated from the central density, which has exactly the same value in both cases because the reconstructed experimental density is the basis of the calculation. The frequency obtained from the calculated time evolutions varies depending on the chosen model. Independent of this, it is however found that the calculated Higgs dynamics are about twice as fast as the observed ones independent of the interaction strength, confirming the observation made before for the quench to unitarity. This discrepancy in the frequencies is a first hint that a local density approximated Higgs mode cannot fully describe the observations from our trap. The Higgs frequency expected according to [67] for a homogeneous gas whose density equals the central density of our clouds is shown as a reference in the figure and even higher than the calculated ones but this comes at

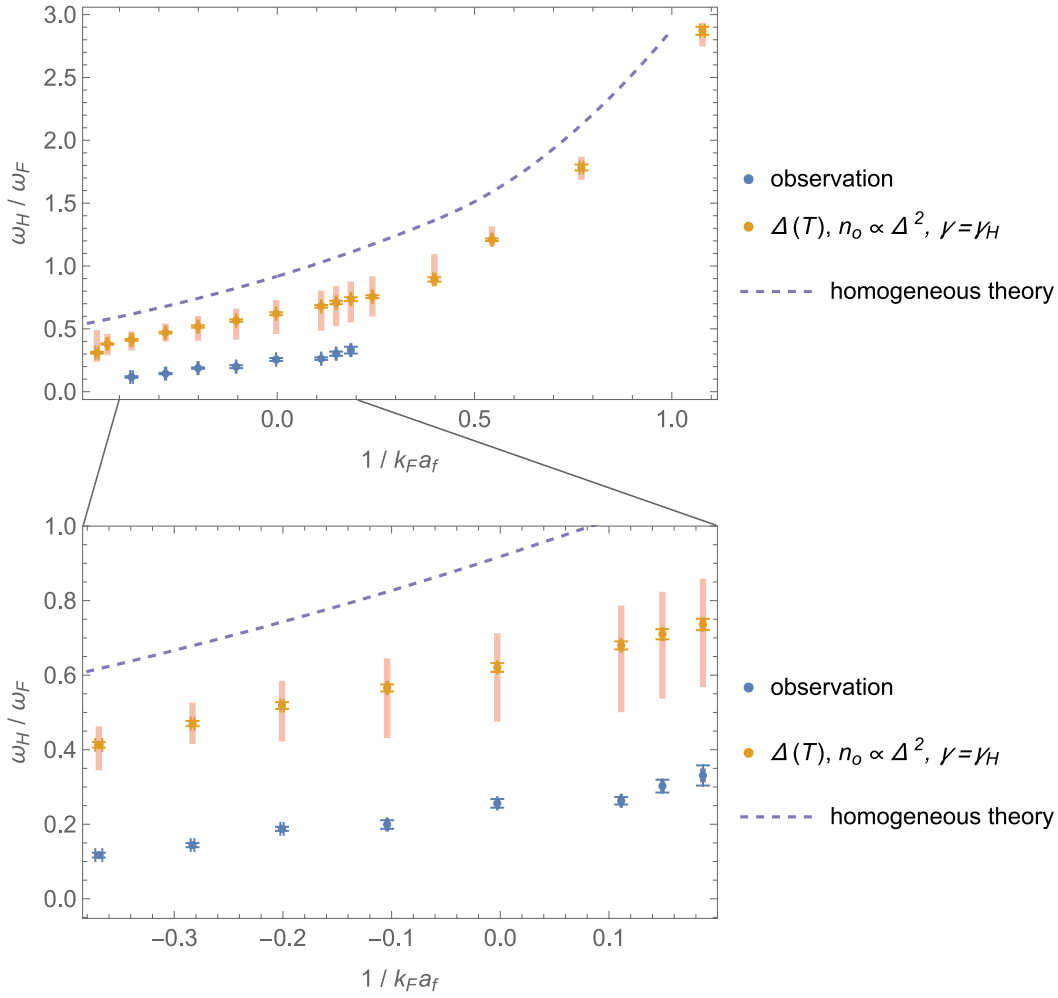

Figure 6.17: Comparison of the fitted Higgs oscillation frequencies for observation and local density calculation, normalised to the Fermi frequency calculated from the central density. Here, the model with temperature-dependent gap and condensate fraction and decaying Higgs mode is plotted. The bars around the calculated values denote the range of frequencies obtained from the various methods to calculate the spectrum. As a reference, the expected frequency for a homogeneous gas at $T = 0$ is given as well. The lower panel is a zoom into the region where it is possible to fit a frequency to the experimental observation.

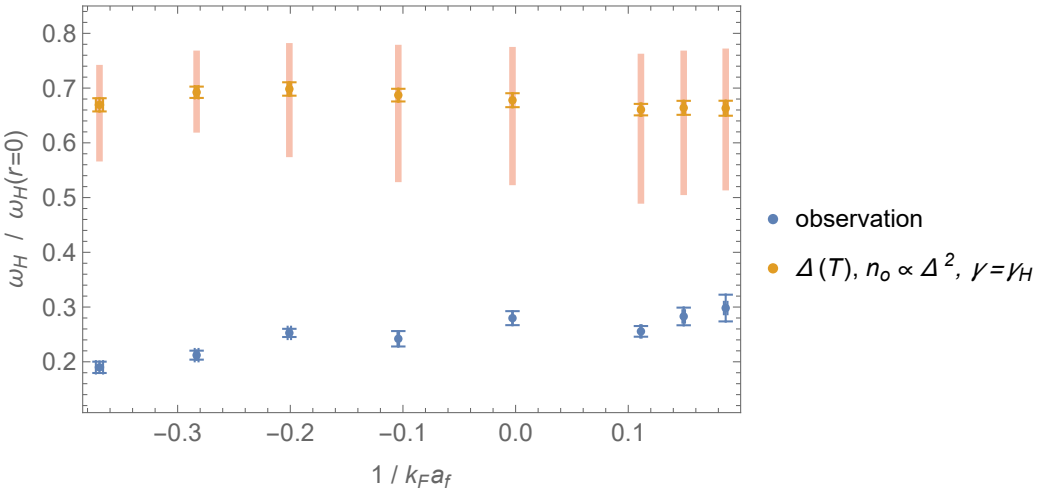

Figure 6.18: Fitted frequencies from experiment and calculation normalised by the Higgs frequency expected from a homogeneous gas with the central density at $T = 0$, according to [67]. The bars around the calculated values indicate the range of frequencies calculated with different models.

no surprise since only a tiny fraction of the clouds' volume can be considered to be roughly described by this case and most parts of the cloud oscillate at a lower frequency in the calculation.

Another normalisation for the fitted frequencies is the expected homogeneous Higgs frequency at the central density. This allows us to check how well the fitted frequencies follow the expected behaviour with varying interaction strength. The results of this normalisation are shown in figure 6.18 and reveal that the frequencies obtained from the calculation are related to the homogeneous Higgs frequency in the centre by a constant factor $\sim 2/3$ relatively independent of the interaction strength. For the observed frequencies the situation is less clear. On the one hand, an increase in the ratio $\omega_\mathrm{H}/\omega_\mathrm{H}(r = 0)$ is observed across the whole recorded interaction range. Since this is in contrast to the calculated behaviour, it would be an indication that the LDA approximation misses some aspects of the inhomogeneous Higgs mode. On the other hand, the observed ratio could also be interpreted as relatively constant for $1/k_\mathrm{F}a_\mathrm{f} > -0.2$ and dependent on the interaction strength only below this threshold. In this case, there would be no general discrepancy between calculation and observation in the dependence on $1/k_\mathrm{F}a_\mathrm{f}$ but only for the two weakest attractions. The clouds for these measurements are relatively close to the critical temperature, as evident from their low condensate fractions below 9 %, and high temperatures are predicted to influence the Higgs mode [60]. Overall, the general dependence of the observed frequencies on $1/k_\mathrm{F}a$ therefore matches the expectations for the Higgs mode. The absolute values of the observed frequencies are however significantly lower than expected.

The reliability of the frequency predicted by the calculation depends on an accurate reconstruction of the density in the trap. Due to the high optical depth and the only approximate radial symmetry of the trapping potential, this reconstruction is prone to systematic errors and it is indeed found that the reconstructed density profiles overestimate the atom number by around 50%[2]. While this appears to be a large error and overestimated densities do indeed lead to overestimated Fermi and Higgs frequencies in the calculation, the actual effect on the predicted frequency is not as drastic because the Fermi frequency

[2] The atom number considered to be the correct atom number is recorded after an expansion time of 5 ms. The reduced optical depth after expansion eliminates many sources of systematic errors.

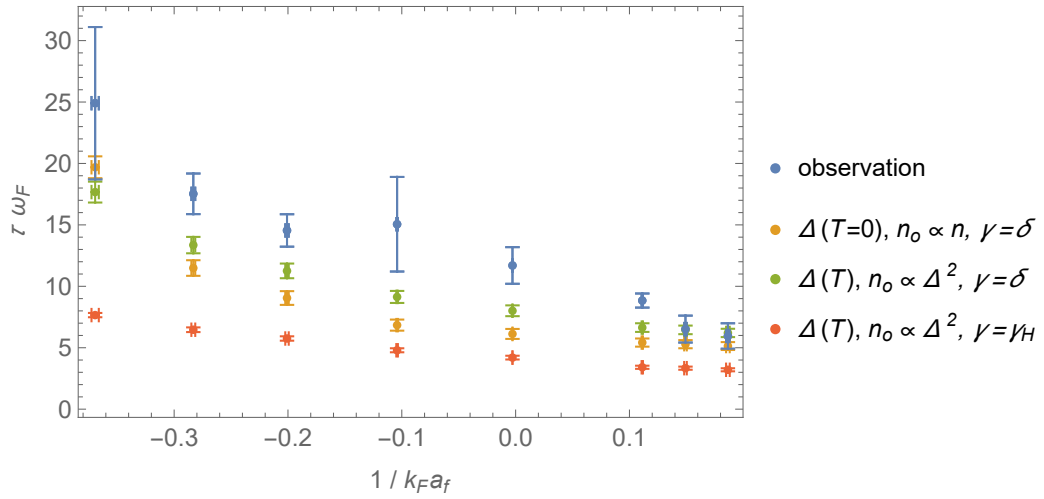

Figure 6.19: Fitted Higgs mode lifetimes from experiment and calculation. The normalisation factors k_F and ω_F are obtained from the central density of the clouds. The calculated lifetimes are obtained from different models for the spectrum, as explained in figure 6.14.

ω_F depends on the density n only as $\omega_\mathrm{F} \propto n^{2/3}$. To check the influence of the overestimated density, the most simple model ($T = 0$ value of Δ, oscillator density proportional to atom density and no intrinsic Higgs mode decay) is also applied to density distributions rescaled by a constant factor. This factor is chosen such that the rescaled density reproduces the correct atom number. It is found that the unaltered density distributions result in observable frequencies that are around 20% higher than those calculated from the rescaled density distributions. The systematic error from the density reconstruction could therefore explain calculated frequencies that would be 1.2 times larger than observed. In reality however, the calculated frequencies are 2 – 3 times larger than observed. Therefore, a systematic error in the density reconstruction is unlikely to explain the mismatch between calculated and observed frequencies.

Another quantity to be compared between experiment and calculation is the decay of the Higgs mode, quantified by its lifetime τ_H. From the observations (figure 6.1), it is immediately obvious that the very fast decay of the inhomogeneous oscillation does not follow the $t^{-1/2}$ prediction for a homogeneous gas in the BCS limit [64]. This is at least qualitatively consistent with the LDA calculations where the collective oscillation decays quickly due to dephasing between the independent Higgs modes. The values for τ_H obtained by fitting the experimental and calculated time evolutions are compared in figure 6.19. For almost all interaction strengths, the calculation predicts a faster decay than observed in the experiment, especially if an intrinsic decay of the Higgs mode is included ($\gamma = \gamma_\mathrm{H}$). This discrepancy can however be explained by the generally faster dynamics predicted by the LDA models: if the local oscillations are faster, a decay due to dephasing will also happen sooner. To eliminate this dependence on the absolute timescale of the dynamics, the lifetime is also normalised by the fitted oscillation frequency. In this way, the quality factor $\tau\,\omega_\mathrm{H}$ of the oscillation is calculated. Figure 6.20 shows the quality factor for calculations with several models and the experimental realisation. It is found that the most complex model – with temperature dependent gap and condensate fraction and an intrinsic decay of the Higgs mode – results in exactly the same quality factor as observed in the experiment for all interaction strengths with a fittable oscillation in

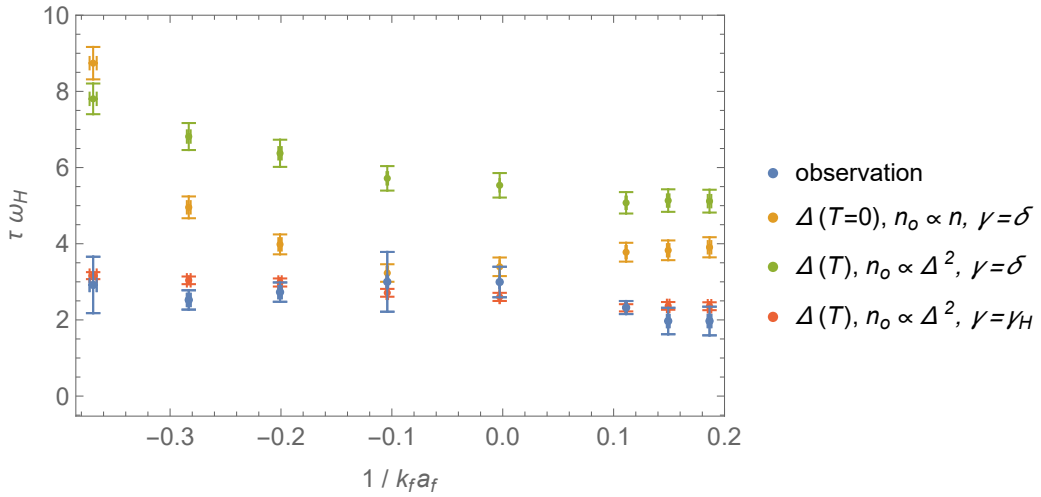

Figure 6.20: Quality factor of observed and calculated Higgs modes. There is a striking agreement between experimental observation and the model with temperature dependent gap and condensate fraction and an intrinsic decay of the Higgs mode ($\Delta(T), n_\mathrm{o} \propto \Delta^2, \gamma = \gamma_\mathrm{H}$)

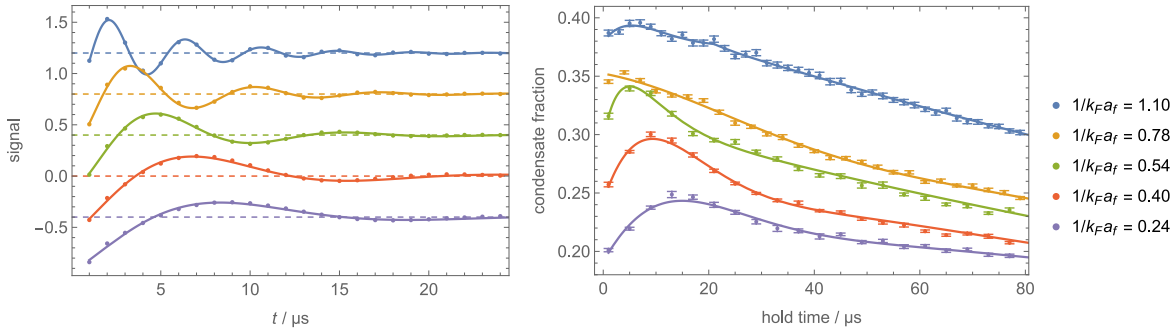

Figure 6.21: Calculated (left) and observed (right) dynamics after a quench to the BEC side with the corresponding fit. The final interaction strength is calculated from the central density of the clouds. The calculated dynamics are shown with an arbitrary offset (dashed lines) and do not include the background dynamics. The calculations predict improved oscillations towards the BEC limit of the crossover, contrary to the observation. The model for the calculation is the one with temperature dependent gap and condensate fraction and an intrinsic decay of the Higgs mode.

the experiment. It is therefore concluded that the decay of the oscillation observed in the experiment can be explained by the dephasing of local Higgs modes with different frequencies, at least in the range of interactions where an oscillation is still identifiable in the experimental data.

The good agreement in this range around unitarity raises the question if it continues further to the BEC side. Fits to the experimental data result in unreliable parameters in this regime with the previously used function because the fast dynamics decay too quickly. The comparison therefore has to be done more qualitatively. Figure 6.21 compares the calculated and observed dynamics. Close to unitarity, i.e. $1/k_\mathrm{F}a_\mathrm{f} < 0.5$, there is still qualitative agreement between calculation and observation, similar to the weaker attraction strengths considered above. For $1/k_\mathrm{F}a_\mathrm{f} > 0.5$ however, an increasing discrepancy is found: the fast dynamics become increasingly weaker until they are hardly visible

at all in the experimental data. The calculation on the other hand predicts an improvement of the oscillation's visibility for higher values of $1/k_{\mathrm{F}}a_{\mathrm{f}}$. This is an indication that the model is missing something in this interaction regime. A very plausible possibility is that this is the instability of the Higgs mode towards the BEC limit as discussed in [59, 60]. Even more so since a rule-of-thumb threshold for the beginning of bosonic behaviour is the zero-crossing of the single-particle chemical potential which lies around $1/k_{\mathrm{F}}a \approx 0.41$, coinciding with the onset of qualitative disagreement between observation and calculation. The necessity of $\mu > 0$ for an observable Higgs mode has been predicted by [147] and [60] also locates the critical interaction parameter for the observability of the Higgs mode in the region $0 < 1/k_{\mathrm{F}}a < 1$.

Summary

To quickly sum up the results from the calculation using the local density approximation, very good agreement to the observation was found regarding the damping of the Higgs mode if compared to its own frequency for interaction strengths not too far on the BEC side of the crossover. This suggests that the observed damping is due to the dephasing of local Higgs oscillations with differing frequencies. The dynamics predicted by the models are however faster than the observed ones by a factor of $\sim 2-3$; this includes both oscillation frequency and decay. For attraction strengths above $1/k_{\mathrm{F}}a \approx 0.5$, the local density models predict oscillations in contrast to the observation. A possible reason is the non-observability of the Higgs mode in the BEC limit which is not included in the model and expected to show up at vaguely this interaction strength.

6.2.4 Quench strength

The biggest question left unanswered by the local density model is the observed low frequency. In [148], it was found that the mean value of Δ after a quench is lower than in its equilibrium value with the difference increasing with quench strength. [58] showed that this phenomenon leads to a reduction of the Higgs frequency which is set by the time averaged value of Δ after the quench. This effect was not taken into account in the model. In [138], it was shown that even qualitatively different behaviours are possible for certain quench configurations giving rise to a full quench phase diagram. It is therefore necessary to take a closer look at the quench strength realised in the measurements.

Figure 6.22 places the experimental quenches in the phase diagram of [138] (which was however calculated for the homogeneous case) and reveals that all quenches lie within the region "II". This is the region in which the Higgs mode manifests itself as the usual damped oscillation. Our previous analysis was therefore justified. Region "I" is characterised by fast decay of Δ to 0 and region "III" by long-lived, non-harmonic oscillations with lower frequency than in region "II", see also [139]. In general, it is experimentally very challenging to reach regions "I" or "III" because the cloud either has to be initially in the far BEC limit (region "I") where particle loss from the trap is a limiting factor or in the BCS limit (region "III") where the exponentially suppressed critical temperature hampers the formation of a superfluid.

In [143], a similar phase diagram is calculated for an elongated harmonic trap and only covering interaction strengths $1/k_{\mathrm{F}}a \leq -0.3$ at zero temperature. There, a smoothly decreasing Higgs frequency is predicted for increasing quench strength. A comparison with our quench parameters is shown in figure 6.23. For the experimentally realised quenches covered by the data provided in [143], the observable Higgs frequency is expected to be above 83 % of the frequency calculated from the

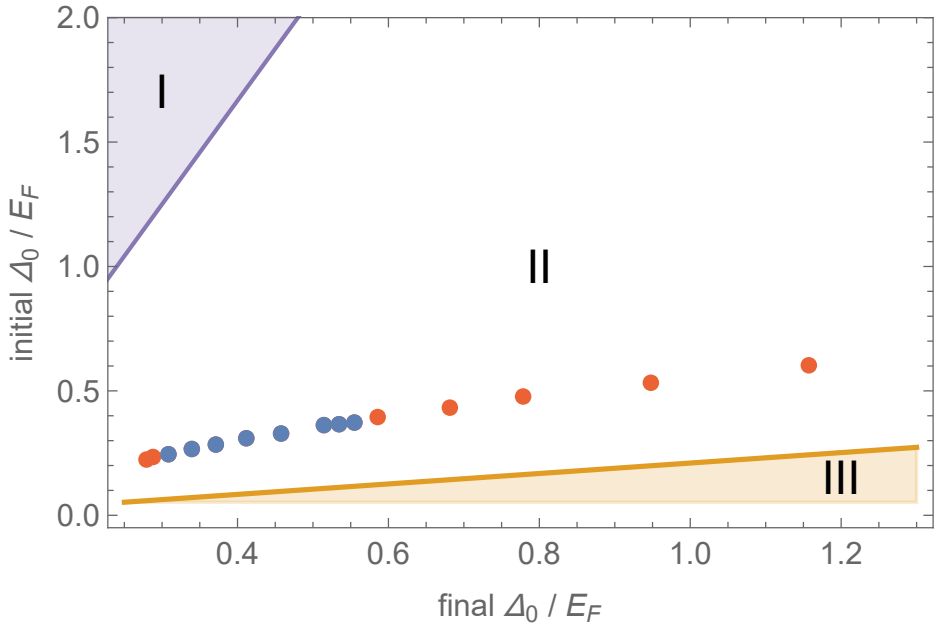

Figure 6.22: The quenches from section 6.2.1 placed in the quench phase diagram of [138]. The Fermi energy is calculated from the central density of the clouds and the equilibrium value Δ_0 from [67]. Quenches indicated by red dots did not allow for a fit of ω_{H}.

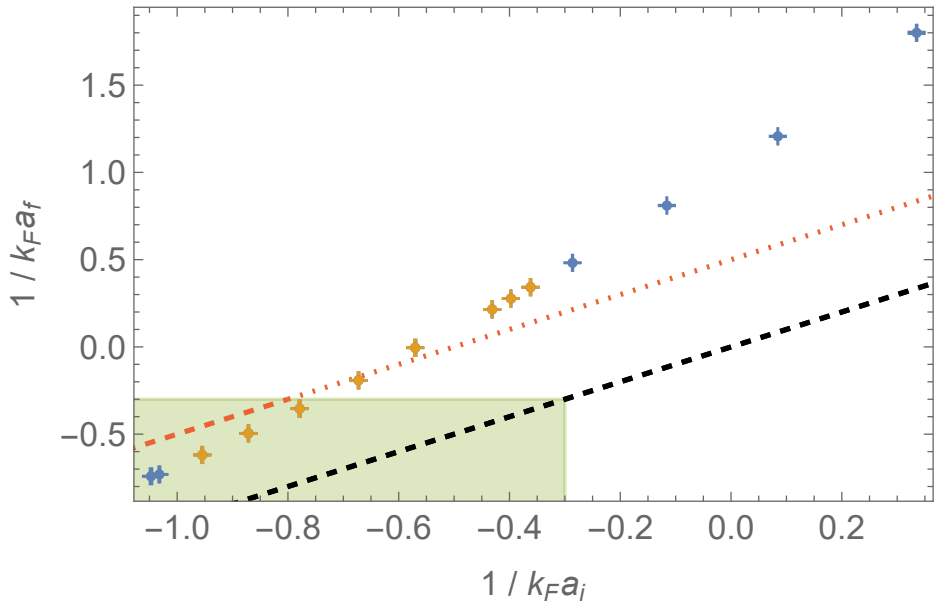

Figure 6.23: Comparison of the quench strength in harmonic normalisation with the phase diagram from [143] which only covers the greenly shaded area. For infinitesimally weak quenches (black dashed line), the Higgs frequency is predicted to be twice the equilibrium value Δ_0. Above this line, the observable frequency decreases continuously, reaching ca. 83 % at the red dashed line. The points are the experimental quenches, with the yellow points indicating the measurements which allowed for a reliable fit of the frequency.

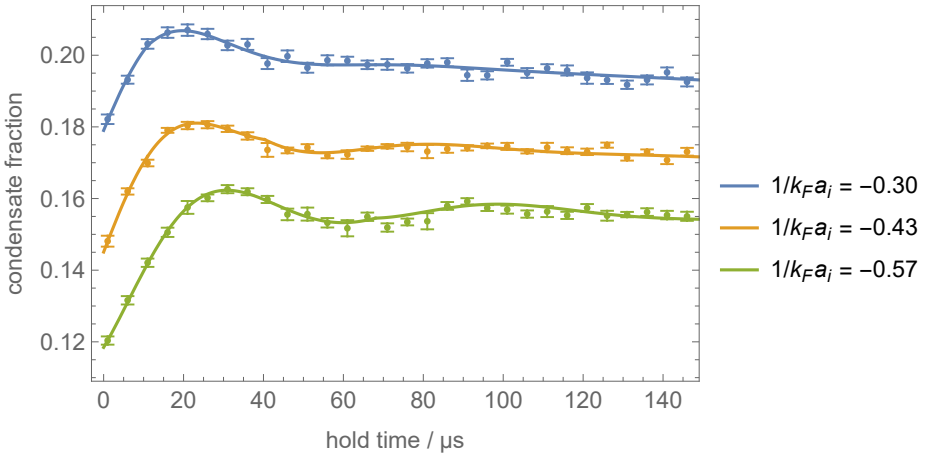

Figure 6.24: Observed dynamics after quenches to unitarity with varying quench strength, normalised by the harmonic oscillator E_{F}.

equilibrium value of the gap, Δ_0. This is in disagreement with the experiment where the observed frequency is below $\approx 50\,\%$ of the frequency calculated with Δ_0 (the result of the local density model), see figure 6.18. If the trends from [143] can be extrapolated through the whole crossover, a stronger reduction in the observable Higgs frequency is expected for our quenches with $1/k_{\mathrm{F}}a_{\mathrm{f}} > -0.3$. Nevertheless, the predicted decrease is still much smaller than the one found in the experiment. It is also noteworthy that the observable Higgs frequency in [143] seems to level off at around 66 % of $2\Delta_0/\hbar$ with increasing quench strength which is still above our observed ratios.

Experimentally, the dependence on quench strength can be tested by simply varying the current in the quench coil. The quenches discussed up to now were performed by sending the highest safe current of 40 A through the coil. In figure 6.24, the observations for quenches to unitarity with 40 A/30 A/20 A corresponding to 35 G/26 G/18 G are presented. The corresponding values for $1/k_{\mathrm{F}}a_{\mathrm{i}}$ are -0.57/-0.43/-0.30 while $1/k_{\mathrm{F}}a_{\mathrm{f}} = 0$ in all cases. Faster dynamics are indeed visible the weaker the quench gets. However, weaker quench strengths are performed with higher atom numbers because the quench coil field gradient affects the trap depth. If the observed frequency is properly normalised with either the harmonic or homogeneous Fermi frequency (calculated from the central density), the apparent acceleration for smaller quenches no longer holds. Instead, the observed frequencies are very similar, see figure 6.25. This strongly suggests that the quench strength is not the main cause for the frequency discrepancy between experiment and model. However, if we would place the 26 G quench in the phase diagram 6.23, it would only be slightly below the 83 % line while the 35 G quench lies slightly above it. It is therefore possible that the quench strength cannot be varied far enough to unambiguously observe the strength dependence of the frequency. The 18 G quench only excites a barely visible Higgs mode and therefore gives limited insights. This is reflected in the large uncertainty of the fitted frequency for the weakest quench in figure 6.25.

Weaker quench strengths were also tested for the smallest values of a_{f} at which no fittable oscillation was observed. The weaker quenches did not improve the visibility of the oscillation and it can therefore be ruled out that the increasing quench strength towards the BEC limit causes the observed decay of the Higgs mode.

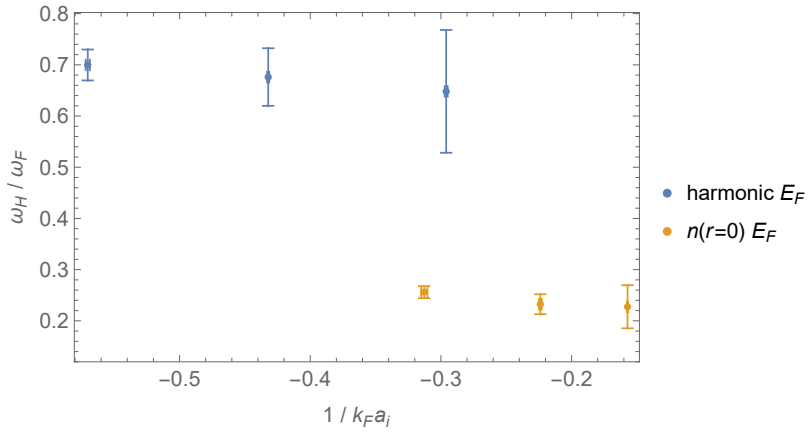

Figure 6.25: The fitted frequency to the observations from figure 6.24 for different initial interaction strengths/quench strengths to unitarity. Both harmonic and homogeneous normalisation from the central density are shown.

6.2.5 Quench direction

In the previous sections, only quenches from weak to strong attraction were investigated. This implies that Δ had a smaller value before the quench than its equilibrium value at the final interaction strength. Since the observed small frequencies are difficult to explain by the final interaction or quench strength, it seems natural to assume that this fact might be related in some other way to $1/k_\mathrm{F}a_\mathrm{i}$, the initial interaction strength. This is not directly suggested by any theories, which only predict a small decrease of the Higgs frequency due to a finite quench strength as discussed in the previous section. Additionally, the theoretical predictions for the equilibrium value of Δ at $1/k_\mathrm{F}a_\mathrm{i}$ are found to already predict a Higgs frequency higher than the observed ones. Nevertheless, there is experimentally a very simple way to investigate dependencies on the initial interaction strength in a different way than just the quench strength: reversing the quench direction. If the current direction in the quench coil is reversed, its magnetic field is subtracted from the Feshbach field instead of added. The resulting quench will then be from strong to weak attraction making $1/k_\mathrm{F}a_\mathrm{i}$ extremely different compared to the previous measurements.

The observations for several quenches from strong to weak attraction are shown in figure 6.26. For all quenches, the magnetic field strength was decreased by 35 G and therefore is the highest achievable quench strength. Again, heavily damped oscillations are visible, which slow down strongly towards the BCS side. This however has to be taken with care since not all measurements were taken with comparable atom numbers. The quenches furthest to the BCS side, $1/k_\mathrm{F}a_\mathrm{f} < -0.5$ for those depicted in figure 6.26, were performed with stronger evaporative cooling because colder samples were necessary to remain in the superfluid phase after the quench. The beam powers for the dipole trap were not changed but the compensation of the residual gradient of the quench coil. This procedure leads to a tilt of the trap during the ramp-up of the quench coil causing further cooling but also atom loss, all while preserving the harmonic parameters of the trap. The resulting dynamics slow down because the absolute values of all relevant quantities like E_F, Δ etc. decrease with the atom number.

It is observed that additionally to the oscillating behaviour, the condensate fraction drops quickly immediately after the quench, followed by a much slower decay for longer hold times. These

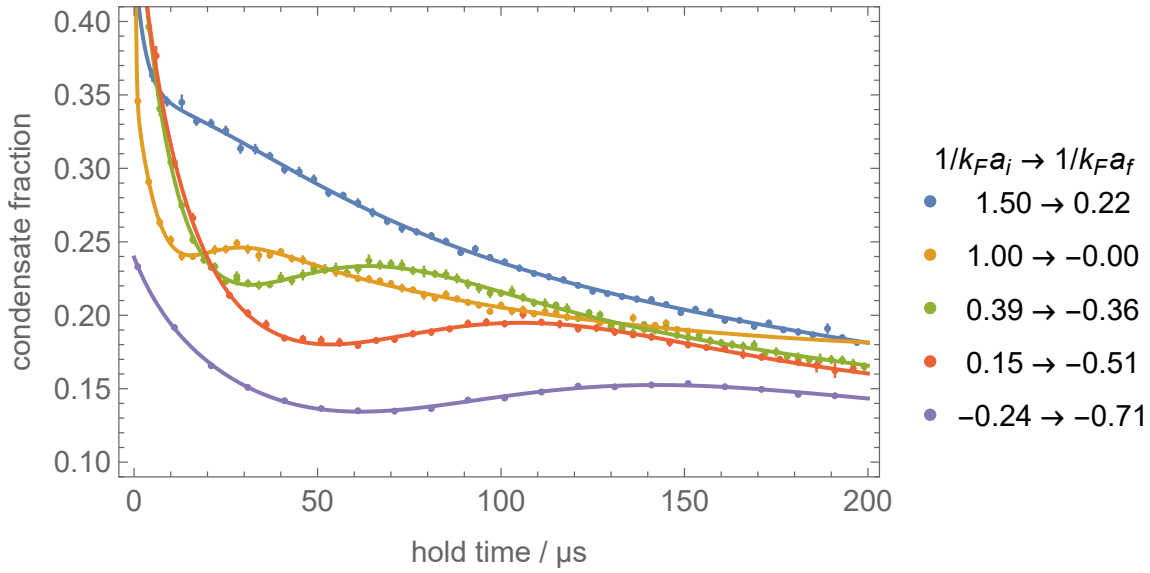

Figure 6.26: Observed dynamics after quenches from strong to weak attraction. The interaction parameters are calculated from the Fermi energy under harmonic approximation. The lines are the best fits of the function $f'(t)$, equation 6.9. The two quenches furthest to the BCS side are performed with reduced atom number resulting in slower dynamics, see main text.

observation motivate a fit with the function

$$f'(t) = A\, e^{-t/\tau_{\mathrm{H}}} \cos\left(\omega_{\mathrm{H}} t + \phi\right) + \Delta\eta_1\, e^{-t/\tau_1} + \Delta\eta_2\, e^{-t/\tau_2} + \eta_0 \tag{6.9}$$

which is very similar to the function used for quenches in the opposite direction, equations 6.1 and 6.2, with the difference that the background now consists of two exponential decays which reduce the condensate fraction by $\Delta\eta_{1,2}$ with decay times $\tau_{1,2}$. Also, the background is no longer split into two qualitatively different regions.

As before, the fitted frequency can be normalised in different ways. Figure 6.27 shows the fitted frequencies normalised by the Higgs frequency predicted for the central density at $T = 0$ and compares both quench directions. It is apparent that the quench from strong to weak attraction shows an even stronger reduction of the oscillation frequency than the previously considered quenches in the opposite direction. This result however has to be taken with care because the initial interaction strength sets the density distribution in the trap for the observed short hold times. The higher $1/k_{\mathrm{F}}a_{\mathrm{i}}$, the higher the density increase towards the centre becomes, leading in turn to a potentially stronger increase of $\omega_{\mathrm{H}}(r = 0)$ relative to other regions of the cloud. Overall, such an effect would reduce the ratio $\omega_{\mathrm{H}}/\omega_{\mathrm{H}}(r = 0)$ shown in figure 6.27.

To investigate this further, the expected observable frequencies under local density approximation are calculated from the reconstructed densities for these quenches in analogy to section 6.2.3. A temperature estimation of the clouds as before did not work in this case because the initial interaction strengths do not extend far enough towards the BCS limit. Also, the residual magnetic field gradient of the quench coil was compensated to different degrees for different quenches making the temperatures less comparable. The calculation can therefore only be done for a zero-temperature Higgs mode with

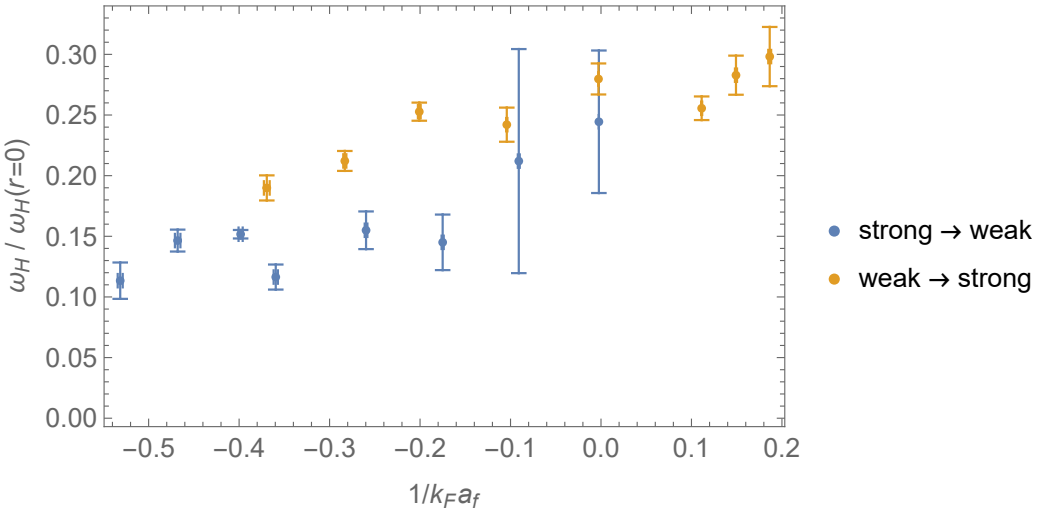

Figure 6.27: Fitted frequencies for both quench directions normalised by the predicted Higgs frequency for a homogeneous gas with the central density at $T = 0$. For quenches from strong to weak interaction, the reduction in frequency is even stronger than in the opposite direction which was discussed before. All quenches in this figure have a quench strength of 35 G.

the given density profiles. Similar to the quench towards stronger attractions, it is found that the observable frequency predicted by the calculation in local density approximation is larger than the actually observed one. The ratio of these frequencies is also very similar for both quench directions, around 2 to 3.

Around unitarity, the fitted frequency for strong → weak quenches has a high uncertainty, as seen in figure 6.27. This is most likely due to an insufficient measurement of the long-time decay of the condensate fraction which makes it difficult to disentangle the heavily damped oscillation and the two decays in the fit. In figure 6.26, it can be seen that the measurement of the quench to unitarity only extends to 150 µs and the fit extrapolates to a different long time behaviour compared to other measurements with similar interaction strength.

It is also possible to normalise the fitted frequency by the Fermi frequency of the harmonic trap. This is shown in figure 6.28 and in this normalisation both quench directions give very similar results. Compared to the homogeneous prediction, which, again, can only serve as a rough estimate of the expected frequency because it describes a different system, the agreement is in principle remarkably good, especially around unitarity. Further towards the BCS side, it appears that the harmonic observation decreases slightly faster than expected in the homogeneous case; possibly due to the finite temperature in the experiment, which is only slightly below the critical temperature in this limit.

A noteworthy feature is the shift of the region with observable oscillations depending on the quench direction. For quenches towards a specific limit of the crossover (BEC or BCS), it extends further towards that same limit than for quenches in the respective opposite direction. This is a strong indication that the initial state has a non-negligible influence on the dynamics. If the initial state is far on the BCS side, the relation $\Delta/k_B T$ before the quench is very small and, together with the disturbance caused by the quench, could lead to an unstable Higgs mode. It is therefore likely that the increased lower limit of $1/k_F a_f$ for observable oscillations when quenching from the BCS side is due to the

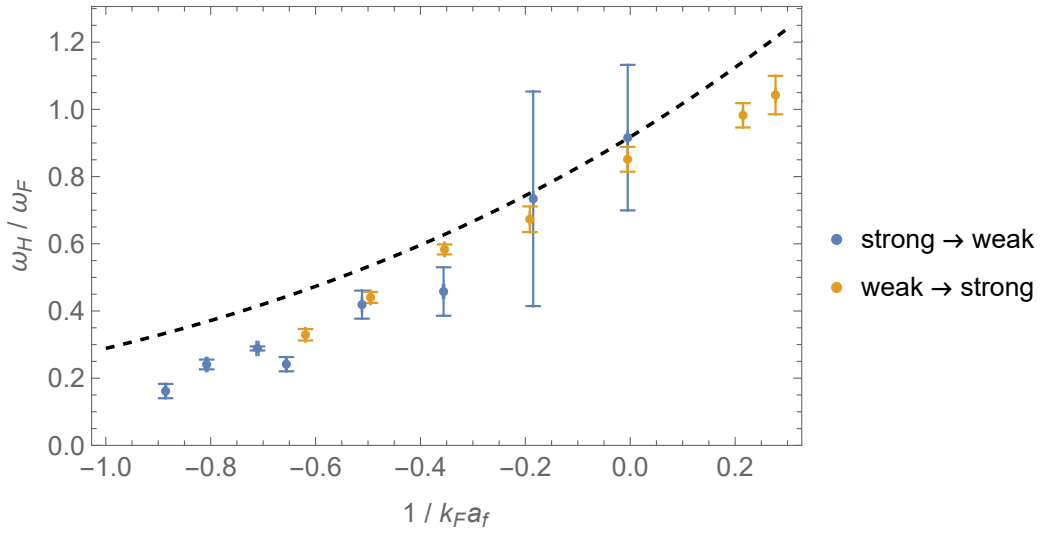

Figure 6.28: Fitted frequencies for both quench directions normalised by the Fermi frequency of the trap in harmonic approximation. The dashed line is the prediction for a homogeneous gas at $T = 0$ from [67].

lowest achievable temperature in the experiment and not an intrinsic feature of the Higgs mode. When quenching from the BEC side on the other hand, temperature should not be an issue. Instead, it is reasonable to assume that features of the initial state cause the Higgs mode to be unobservable for small positive values of $1/k_\mathrm{F}a_\mathrm{f}$. A possibility is the softening of the Fermi edge due to pairing for increasing attraction strength which leads to an instability of the Higgs mode, see section 2.4.1 and [75]. This softening should still be present immediately after the quench and could therefore obscure the Higgs mode at final interaction strengths where it should in principle be observable. For quenches from strong to weak attraction, $1/k_\mathrm{F}a_\mathrm{i} = 1.0$ is the largest value with a clearly observable oscillation. This is significantly larger than $1/k_\mathrm{F}a_\mathrm{f} = 0.34$, the largest value at which a clear oscillation was found for quenches from weak to strong attraction (both in harmonic approximation). The criterion for observability in a general quench is therefore seems not to be related to a single interaction parameter, for example the larger value of $1/k_\mathrm{F}a_\mathrm{i}$ and $1/k_\mathrm{F}a_\mathrm{f}$, but must consider both of these values.

6.2.6 Trap influence

A consistent feature so far has been the good agreement between the observed frequency and the prediction for a homogeneous gas at the same interaction strength and zero temperature. As mentioned before, this is most likely a coincidence caused by several systematic errors. Nevertheless, it is interesting to check how robust this agreement is for varying trap frequencies. To this end, quenches to unitarity with a quench strength of 35 G are performed for different trapping frequencies. The frequencies are set by the power of the dipole trap beams which are measured and regulated with photodiodes. The beam power is proportional to the photodiode voltage causing the trap frequencies to depend on the voltage according to $\omega_\mathrm{trap} \propto \sqrt{V_i}$ where $i \in \{H, V\}$ denotes the horizontal or vertical dipole trap beam respectively. From the measured trap frequencies of the initial setting $V_\mathrm{H} = 0.16\,\mathrm{V}$, $V_\mathrm{V} = 0.1\,\mathrm{V}$ used in the previous sections and the known beam shapes, the trap frequencies for other settings can be calculated.

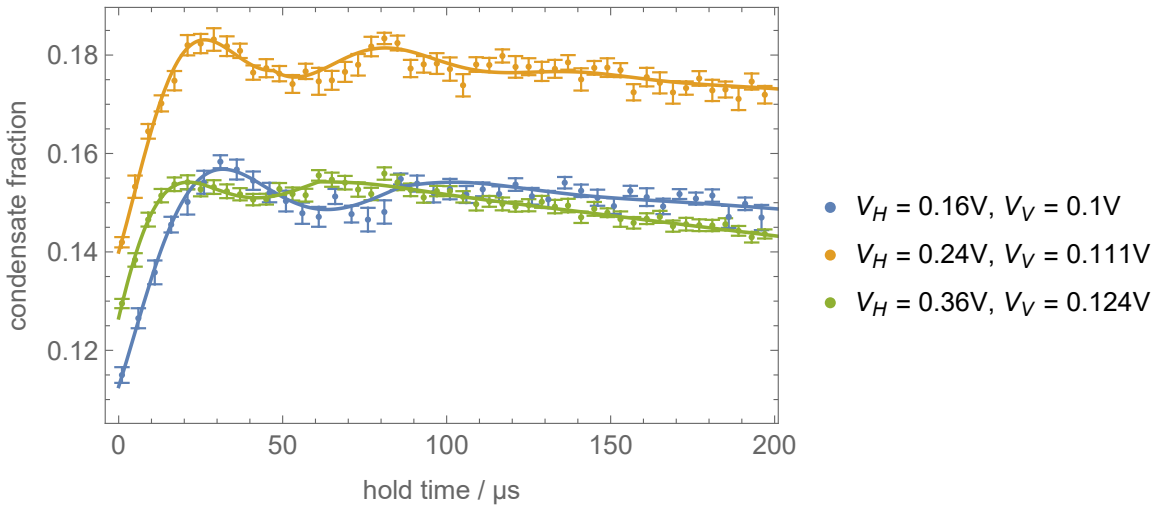

Figure 6.29: Dynamics following a quench to unitarity for different trap settings. With increasing trapping beam powers (photodiode voltages), an acceleration of the dynamics is visible. The beam powers are chosen to keep the atom number constant.

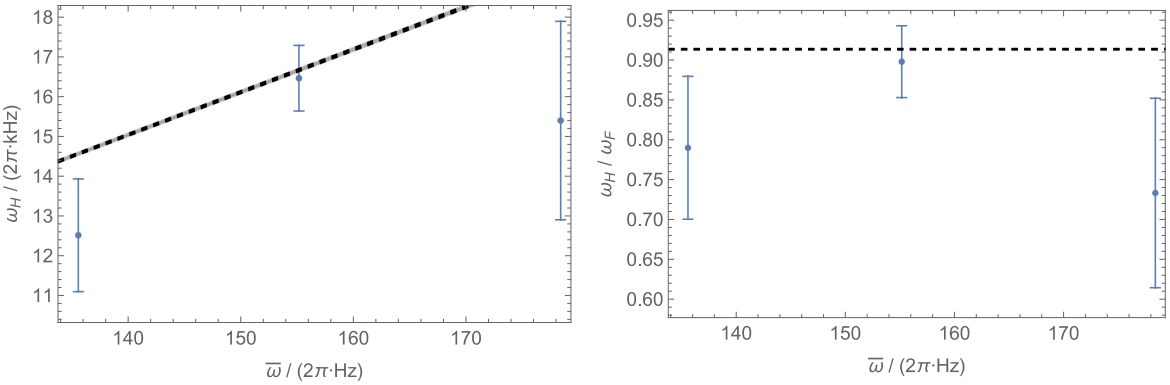

Figure 6.30: Fitted frequencies from figure 6.29, absolute (left) and normalised by the Fermi frequency in harmonic normalisation (right) as a function of the geometric mean of the trapping frequencies. In both cases, the dashed line is the prediction for a homogeneous gas with identical Fermi frequency.

The observations for different trap settings are shown in figure 6.29. Since the beam powers have been chosen to keep the atom number roughly constant around $N = 2.7 \cdot 10^5$, they can be compared without further normalisation. Higher beam powers/photodiode voltages increase the trapping frequencies and therefore E_F, which is proportional to the geometric mean of the trapping frequencies $\bar{\omega}$. The ratio $\omega_\mathrm{H}/E_\mathrm{F}$ is expected to be constant and should therefore lead to accelerated dynamics with increasing beam powers. This matches the observations. The frequencies are determined by fitting the same function as previously in section 6.2.1. The results are presented in figure 6.30 both in absolute and normalised values. They roughly follow the expected scaling with $\bar{\omega}$ but a definite conclusion is hard to draw because the values of ω_H cannot be determined precise enough. A stronger variation of $\bar{\omega}$ would be desirable but it is challenging to create a stable condensate for more extreme variations of the trap parameters. Further studies would be necessary to elucidate this point.

6.2.7 Quench of a warm superfluid

A possibility to reduce the influence of the inhomogeneities caused by the external potential is the confinement of the superfluid to the central region of the trap. This can be achieved by preparing a superfluid only slightly below T_c which, according to the local density approximation, will be concentrated around the centre of the trap where the variation in the density is small. The Higgs mode of this superfluid is therefore expected to be less affected by dephasing and potentially better observable.

Quenching a superfluid with a condensate fraction of a just few percent by 35 G to unitarity however did not result in any observable oscillation. Very recently, it has been shown theoretically that density fluctuations, which are created by finite temperatures, are detrimental to the oscillations associated with the Higgs mode [139]. This could explain the unobservability of the Higgs mode in this almost homogeneous case close to T_c.

6.2.8 Fourier spectrum

So far the observed frequencies have been determined by a fit to the time evolution of the condensate fraction. This fit has considerable uncertainties due to the strong damping of the oscillation and the slow background dynamics of the condensate fraction. An alternative way to find the frequency is the calculation of the spectra of the observed dynamics by a discrete Fourier transform. In these spectra, the Higgs frequency is expected to show up as a peak. Figure 6.31 shows several spectra and peaks are indeed visible, but only for a few quenches – all spectra not shown in figure 6.31 do not exhibit a peak. The frequencies obtained from fits to the time evolution are also shown in the figure and a rough agreement to the peak positions is found. Upon a closer look, it appears that the peaks in the spectra often tend to be at higher frequencies than the fitted ones. It is however difficult to actually confirm this trend because the sampling of the spectrum is extremely sparse with most peaks containing less than 5 data points.

This reflects the original intention of the measurement which was to observe the Higgs mode in time domain. Since the spacing of the points in the spectra is given by $\mathrm{d}\nu = \frac{1}{\mathrm{d}t\,N}$, where $\mathrm{d}t$ is the spacing of the equally spaced data points in time domain and N the total number of points, the spectral resolution can be improved by either increasing $\mathrm{d}t$ or by increasing N. The former is not ideal because it can obscure dynamics that are faster than expected while the latter increases the number of necessary experimental cycles. This would extend the already long measurement time with its ~ 50 repetitions making the measurement more susceptible to slow drifts in experimental performance.

An interesting observation is that some spectra show faint signs of higher harmonics of the main frequency. This could be a faint signal of multiple Higgs modes with frequencies $2\Delta/\hbar$, $4\Delta/\hbar$, $6\Delta/\hbar$, ... whose existence has indeed been predicted in [149].

6.2.9 Initial growth of the condensate

With the strong damping of the observed oscillatory behaviour, alternative quantities describing the dynamics following a quench become more interesting. One such quantity is the initial growth of the condensate immediately after the quench, i.e. before the first local maximum is reached. The growth rate can be found by fitting a linear function to the first 4 to 5 measured hold times as illustrated in figure 6.32.

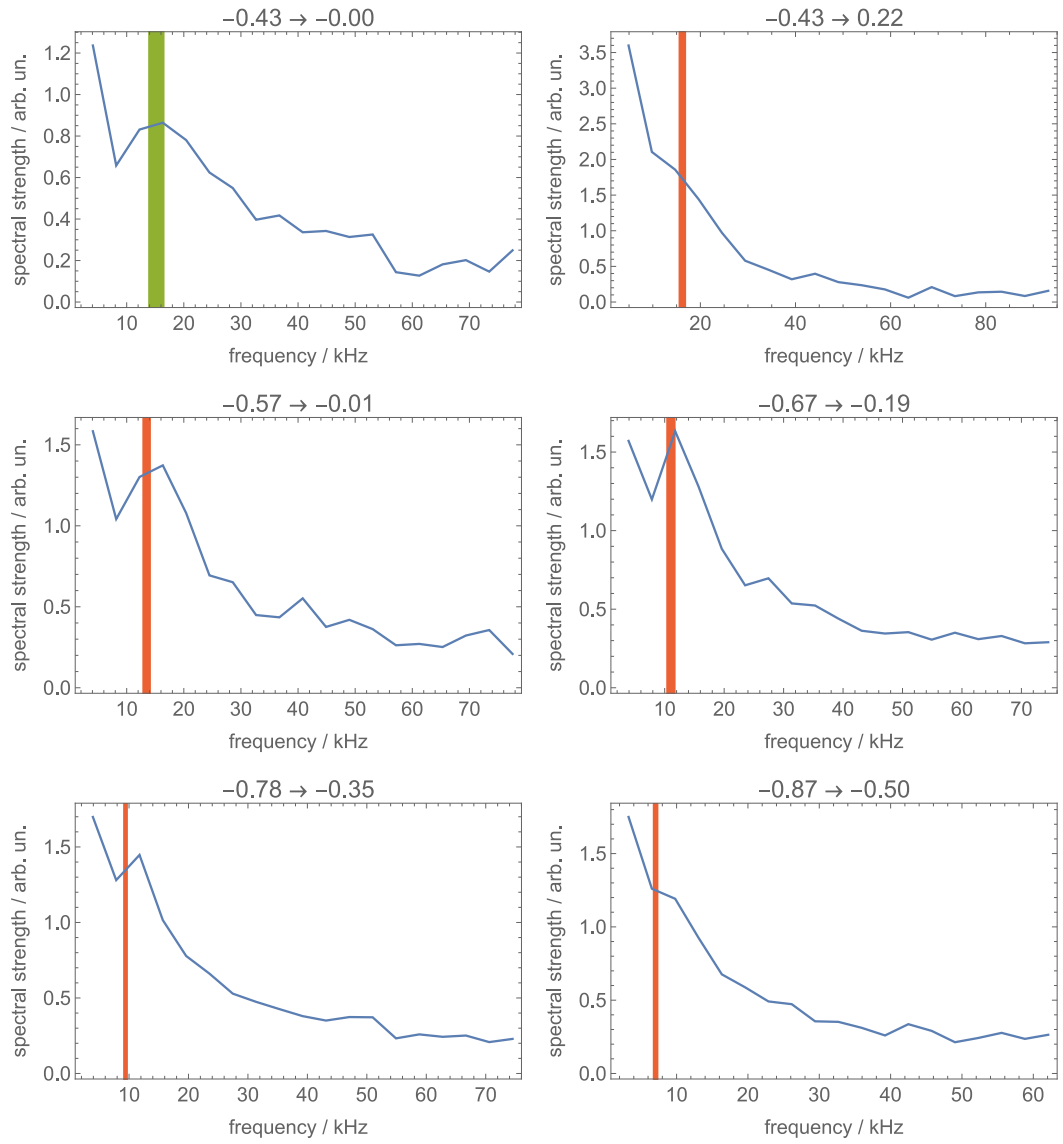

Figure 6.31: Fourier spectra calculated from the observed time evolutions. Every plot is labelled by the interaction quench parameters in harmonic approximation. The vertical line indicates the frequency obtained from the fits with the width indicating the uncertainty. A green (red) line denotes a quench strength of 26 G (35 G). The large zero-frequency component of the spectrum is not shown in these plots.

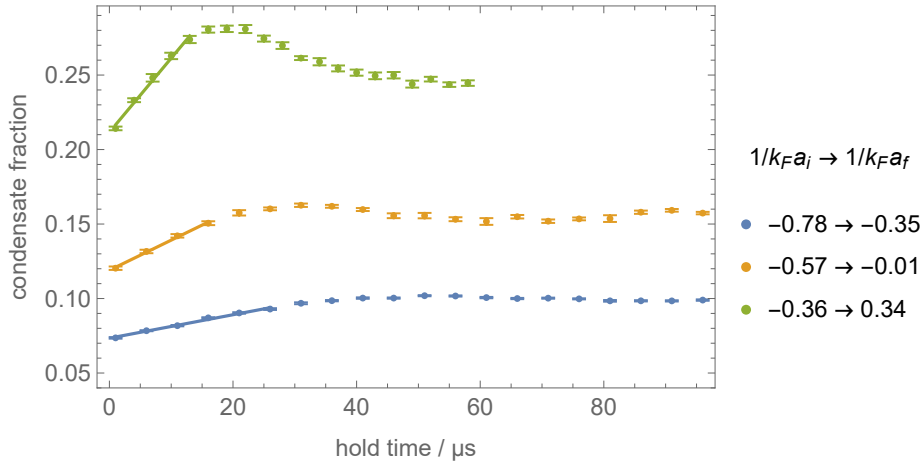

Figure 6.32: Linear fits to the initial rise of the condensate fraction. The interaction parameters of the quenches are in harmonic approximation.

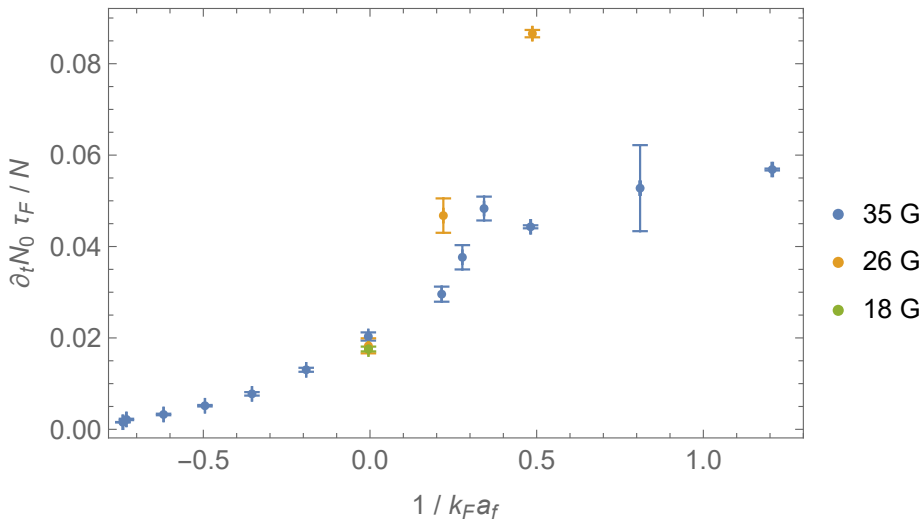

Figure 6.33: Interaction dependence of the growth of the condensed atom number normalised by Fermi time and total atom number; the Fermi energy is calculated in harmonic approximation.

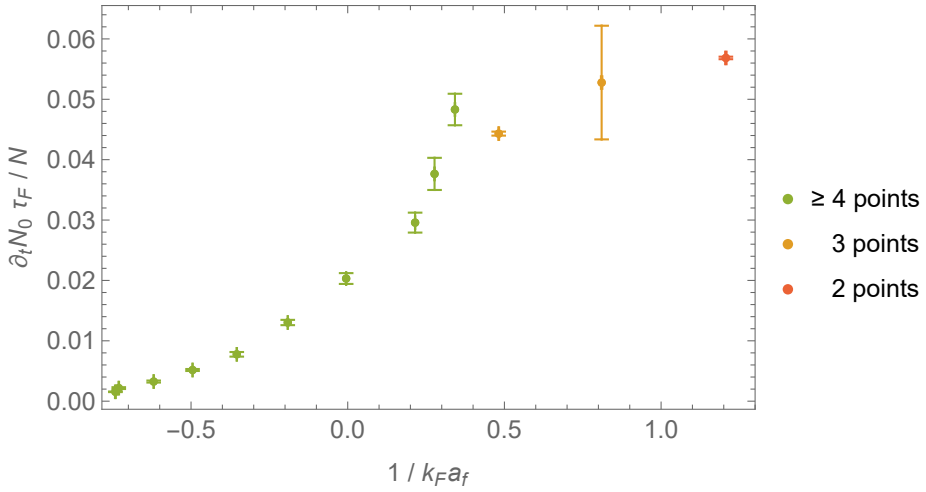

Figure 6.34: Number of data points included in the fit to determine the initial growth of the condensate for quenches with a strength of 35 G. The levelling off of the growth coincides with a reduction of the included data points. It is possible that this reduction leads to an underestimation of the initial growth, see main text.

To analyse the results, it is helpful to remember that the condensate fraction is given by N_0/N where N_0 is the number of atoms found in the condensed part after a rapid ramp and N the total number of atoms. For the short hold times considered here, N is a constant while N_0 changes with time. The growth rate $\partial_t N_0/N$ is therefore plotted against the final interaction parameter in figure 6.33. Since the dynamics of the system are expected to accelerate with higher atom number/Fermi energy, the growth rate is normalised by the Fermi time $\tau_\mathrm{F} = \hbar/E_\mathrm{F}$ to make the various measurements comparable. It is apparent that the condensate growth accelerates with increasing attraction strength from the BCS side up to $1/k_\mathrm{F}a = 0.4$ where the observed growth starts to level off.

In the regime $1/k_\mathrm{F}a \geq 0.5$, the result is however less reliable because the duration of the initial rise decreases, as can be seen in figure 6.2. Less data points therefore cover the initial rise and the hold time at which the growth rate slows down is difficult to identify. It is therefore possible that regions with slower growth are included in the fit as well and this would lead to an underestimation of the initial growth from the fit. In figure 6.34, it can be seen that the levelling off of the initial growth coincides with a reduction of the data points included in the fit to below 4 for quenches with 35 G. It can therefore not be excluded that the levelling off is an artefact caused by the finite sampling of the initial growth instead of an inherent feature of the dynamics.

Interpretation

In general, it is expected that the dynamics of the condensate accelerate with increasing attraction, as observed, because the time scale of the condensate is set by the value of the order parameter Δ which increases with attraction strength. It is however difficult to estimate this more quantitatively due to the system being in a strongly non-equilibrium state. For different quench strengths with identical $1/k_\mathrm{F}a_\mathrm{f}$, weaker quenches are expected to lead to faster dynamics because the initial attraction is stronger leading to a higher initial value of Δ. This behaviour is actually observed if the different quench strengths are compared in figure 6.33. For $1/k_\mathrm{F}a_\mathrm{f} > 0$, the condensate grows much faster after a weaker quench of 26 G compared to a quench of 35 G. However, almost no difference is

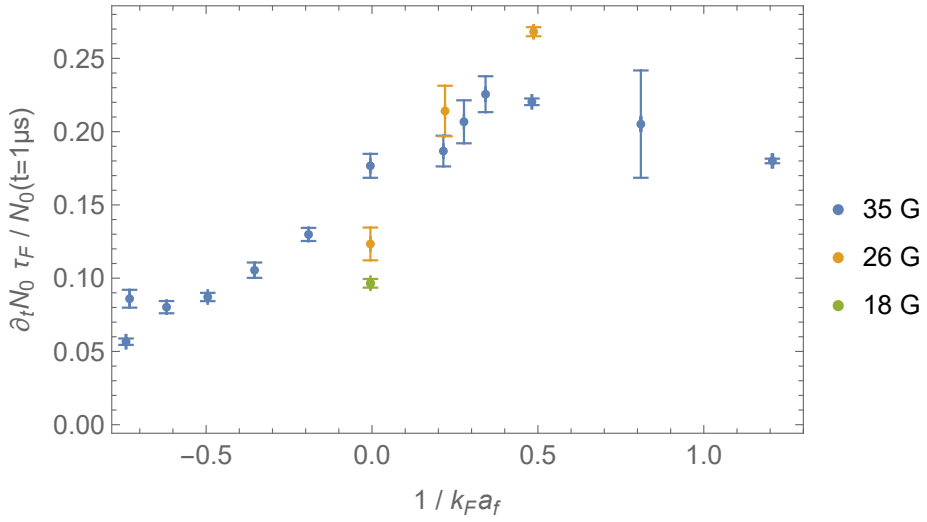

Figure 6.35: Initial growth of the condensed atom number normalised by Fermi time and the number of initially condensed atoms for different interaction and quench strengths.

observed for quenches to unitarity. Here, the condensate fractions are fairly similar indicating that Δ is similar which in turn should lead to similar timescales. Around $1/k_\mathrm{F}a_\mathrm{f} = 0.5$, the difference in the initial condensate fraction is much more pronounced with the 35 G quench having $N_0/N = 20\,\%$ immediately after the quench compared to 32 % for the 26 G quench.

To further investigate the connection to the initial condensate fraction it is also possible to normalise $\partial_t N_0$ by the number of initially condensed atoms $N_0(t_\mathrm{hold} = 1\ \mu\mathrm{s})$ and the Fermi time. The resulting interaction dependence is shown in figure 6.35 and is qualitatively similar to the normalisation with N. The evolution however appears to be much smoother over the full crossover. Additionally, the quench strength dependence becomes weaker.

Another aspect is introduced by the "shape" of the Feshbach resonance: quenches of identical B-field strength become stronger with increasing attraction, see figure 6.23, thus driving the system further away from equilibrium. The resulting state immediately after the quench should therefore have more quasiparticle excitations which in turn reduces Δ and slows down the resulting dynamics. There could therefore be a competition between attraction strength accelerating the dynamics and quench strength slowing them down. The observed levelling off of the growth towards the BEC limit could be a consequence of the quench strength gaining the upper hand. It is also an interesting fact that the levelling off happens roughly at the transition to negative chemical potentials, according to [67].

Overall, it is observed that the condensate grows faster with increasing attraction strength as expected. The details however appear to be relatively complicated and depending on a number of additional factors like quench strength and initial condensate fraction. Additionally, growth rates relatively far on the BEC side are difficult to observe because the initial linear growth slows down extremely quickly.

6.2.10 Comparison with other measurements

During the work on this thesis, the quench experiment was also performed using the alternative condensate detection method provided by the neural network from section 4.3. A detailed analysis of

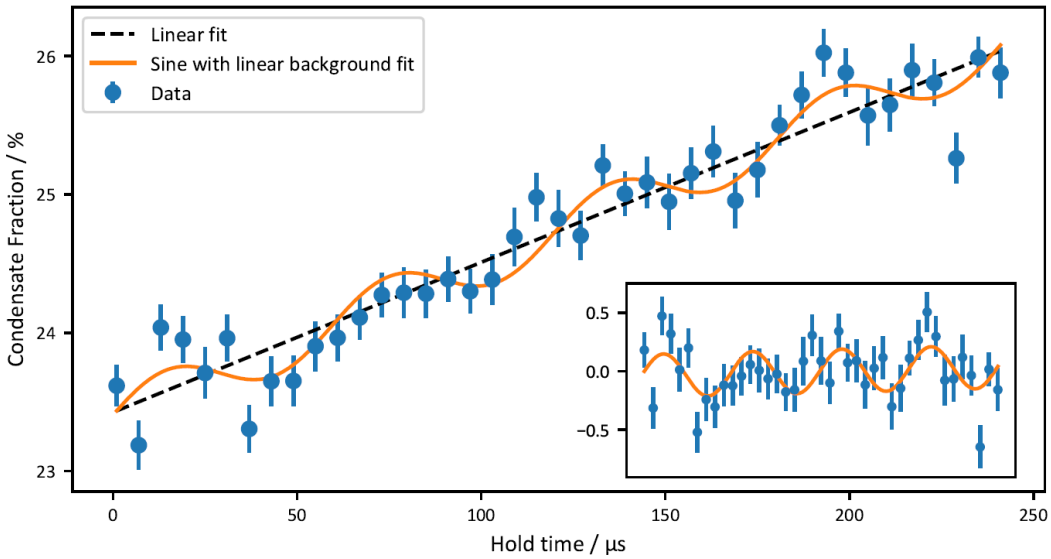

Figure 6.36: Dynamics of the condensate fraction after a quench from $1/k_\mathrm{F}a_\mathrm{i} = -0.44$ to $1/k_\mathrm{F}a_\mathrm{f} = 0$ as observed with the help of a neural network. The inset shows data and sine fit after subtraction of the linear fit. Figure from [93].

the results can already be found in the PhD thesis of Andreas Kell [93], therefore only a brief summary is given here.

Figure 6.36 shows the dynamics of the condensate fraction after an interaction quench to unitarity as observed by the neural network. An oscillation in a similar frequency range to those detected with the rapid ramp, first minimum after around 50 µs, is visible, see also figure 6.2. The signal is however significantly weaker compared to the rapid ramp detection and only just surpasses the noise level. For quenches to other interaction strengths, the signal to noise ratio tends to be even worse making consistently reliable determinations of the frequency difficult. Nevertheless, there is usually a frequency within 20 % of the theoretical expectation for a homogeneous gas for which an oscillation improves the simple linear background fit substantially. To check whether the fitted oscillations are just an artefact of the noisy data, a null-hypothesis test was employed. It revealed that there is only a 1.5 % probability of the observations resulting from random noise. Consequently, there is a 98.5 % probability that the Higgs mode was observed with this method. A meaningful life time of the Higgs mode could not be determined from these measurements due to the low signal to noise ratio.

The observations with rapid ramp and neural network exhibit considerable qualitative differences in the background evolution of the condensate fraction despite of the identical experimental protocol used in both cases. The neural network observes a steady but relatively slow grow of the condensate over the whole measurement range while the rapid ramp observes a fast intial increase followed by a saturation or decline, depending on the quench strength. Since the neural network has been trained and manages to reproduce the results of the rapid ramp in equilibrium, the discrepancy must be caused by the non-equilibrium state of the gas. It is not known how the neural network determines the condensate fraction or, phrased differently, which “assumptions” it employs. Depending on how many properties of the equilibrium state are taken for granted by the network’s algorithm, the estimation of the condensate fraction could produce wrong results in a non-equilibrium situation.

The rapid ramp on the other hand is not well enough understood to rule out unexpected behaviour in a non-equilibrium situation. The condensate fraction determined in this way could therefore differ from its analogue equilibrium value as well. The imperfections of either detection method would not necessarily impede the detection of the Higgs mode because only relative fluctuations of the condensate fraction have to be found. The observed oscillations however greatly differ in amplitude and damping (but are similar in frequency) indicating again that at least one of the detection methods struggles with the non-equilibrium situation. On a speculative level, it seems likely that the neural network has more troubles with the non-equilibrium situation because it was trained with equilibrium data only. Further investigations would however be necessary to clarify this point. One could for example start with very small deviations from equilibrium where neural network and rapid ramp might predict identical condensate fractions and then gradually increase the non-equilibrium.

A very similar quench experiment was performed at Swinburne University [150] in parallel to the experiments presented here. There, a Fermi gas was quenched from a normal state in the far BCS limit ($1/k_\mathrm{F}a_\mathrm{i} = -1.9$ in harmonic approximation) to unitarity, where a superfluid formed. The quench duration was varied and for the fastest quenches of 50 µs to 1 ms, a heavily damped oscillation of the condensate fraction similar to our observation is visible in the presented data. This is however not commented upon in the publication which has a slightly different focus.

Shortly before submission of this thesis, new results from the Swinburne experiment were published in [151]. For these, a superfluid gas was quenched from $1/k_\mathrm{F}a_\mathrm{i} = -0.18 \pm 0.02$ to unitarity within 50 µs. The condensate fraction was determined by Bragg spectroscopy with two tightly-focussed laser beams. Since the beam waists are much smaller than the cloud, the Bragg spectroscopy measures the local condensate fraction. This is a big difference to the rapid ramp method which can only determine the global condensate fraction. After the quench, they observed clear oscillations of the local condensate fraction with frequency and damping matching the expectations for the Higgs mode at unitarity quite well. Together, the results [150, 151] are another strong indication that the heavy damping of our oscillations is indeed caused by the inhomogeneous density, at least at unitarity, and that we have observed the Higgs mode.

6.3 Long time dynamics

So far, only the time evolution on "short" timescales has been discussed. "Short" in this context means everything that evolves much faster than the trap frequencies ω_i which lie in the range of $2\pi \cdot 80\,\mathrm{Hz}$ to $2\pi \cdot 220\,\mathrm{Hz}$. In this section, the time evolution on "long" timescales – those set by the trap frequencies, which are the lowest characteristic frequencies of the idealised system[3] – is analysed. This timescale is on the order of a few ms.

6.3.1 Condensate fraction and density

The corresponding time evolution of the condensate fraction for different values of $1/k_\mathrm{F}a_\mathrm{f}$ is shown in figure 6.37. Very clear oscillations of the condensate fraction, albeit not of harmonic nature, are observed with a "period" of around 3 ms. After roughly two periods, the oscillations stop and the condensate fraction decays to a new equilibrium value which can also be zero. Since dynamics on

[3] Only the experimental imperfection of particle loss from the trap happens on even longer scales of a few seconds.

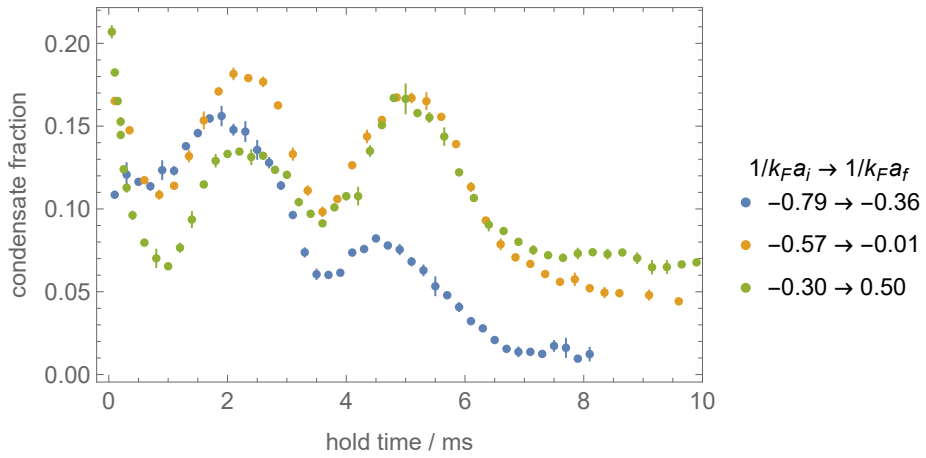

Figure 6.37: Time evolution of the condensate fraction after a quench on the timescale set by the trap frequency for different final interaction strengths. All quenches changed the magnetic field by 35 G and k_F is calculated in harmonic approximation.

the observed timescale are expected to be governed by the trapping potential, the time evolution of the density is recorded as well. In contrast to the previous density measurement, the inverse Abel transform is not employed in this case because a lot of time-consuming averaging would be necessary for every hold time. Therefore, an alternative quantity, similar to the density, is calculated from the *in-situ* size of the cloud. This alternative will be denoted here as the simplified density n' and is given by the inverse product of the cloud's widths σ_i along the three axes of the trapping potential $i = x, y, z$

$$n' = \left(\sigma_x \sigma_y \sigma_z\right)^{-1} . \tag{6.10}$$

The individual sizes are determined by a Gaussian fit to the recorded *in-situ* images. This approach relies on the atom number being constant during the observed time such that there is a one to one correspondence between cloud size and density. A constant atom number for all hold times has been verified experimentally.

The time evolution of the different σ_i after a quench to unitarity is shown in figure 6.38. It is observed that the cloud contracts immediately after the quench along all directions. This matches the expectations because an increased attraction strength should lead to a shrinking of the cloud. However, the external potential also plays a role because the inhomogeneous magnetic field of the quench coil reduces the trap frequencies. In turn, this implies that the trap frequencies increase with the quench and this is a second mechanism leading to a compression of the cloud. After a little less than 2 ms, the cloud starts to expand again in y- and z-direction, a behaviour reminiscent of the "breathing"-mode, a collective mode with zero angular momentum in the trapping potential that has been studied before in the BCS-BEC crossover both experimentally [152, 153] and theoretically [154, 155]. In our system, where all trapping frequencies are distinct[4], neither the angular momentum of the collective mode nor any of its components are conserved and consequently, the different collective modes like breathing, quadrupole, etc. are coupled [78, 155]. The frequencies of these various modes differ and the coupling

[4] $\omega_x = 2\pi \cdot 85\,\mathrm{Hz}$, $\omega_y = 2\pi \cdot 137\,\mathrm{Hz}$, and $\omega_z = 2\pi \cdot 214\,\mathrm{Hz}$

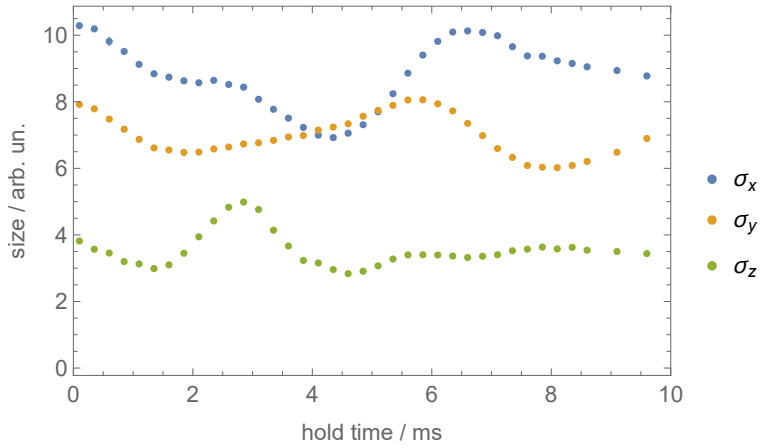

Figure 6.38: Time evolution of the cloud size along the three axes of the trapping potential after a quench from $1/k_\mathrm{F}a_\mathrm{i} = -0.57$ to $1/k_\mathrm{F}a_\mathrm{f} = -0.01$.

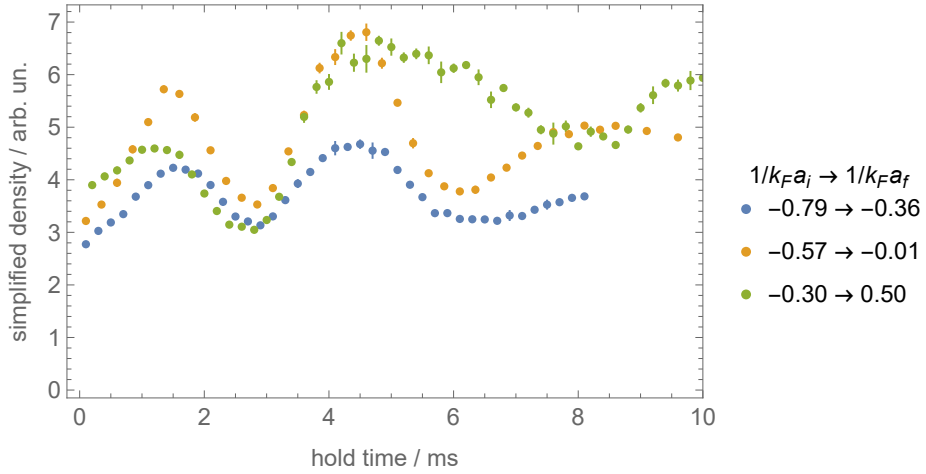

Figure 6.39: Time evolution of the simplified density n' for the quenches shown in figure 6.37.

should therefore cause an overall behaviour characterised by various frequencies and this indeed seems to match the observation in figure 6.38.

The next step in linking the condensate fraction to the collective density modes is the calculation of the simplified density n' from the cloud sizes with equation 6.10. The resulting time evolution of n' is shown in figure 6.39 and reveals oscillations very similar to those observed in the condensate fraction. The observed dynamics in the condensate fraction can therefore be explained as a consequence of the excitation of collective modes: The increased attraction between the atoms after the quench and the compression of the trap lead to a contraction of the cloud that initiates the breathing mode of the trapped gas which couples to other collective modes due to the anisotropic geometry. The result is an oscillation of the density not describable by a single frequency. The compression of the cloud changes the local density and therefore the local values of k_F and T_F. Consequently, the equilibrium value of the condensate fraction changes as well and the system will evolve towards the new equilibrium value. The variation in T_F seems to be decisive in this case because the behaviour at unitarity, where a variation in k_F does not change the interaction parameter $1/k_\mathrm{F}a$, is identical to that at finite scattering lengths.

6.3.2 Cross-correlation

From a comparison of the time evolution of condensate fraction in figure 6.37 and the density in figure 6.39, it can be seen that the evolution of the density precedes that of the condensate fraction, in agreement with the previous interpretation. A straightforward question is now the dependence of the condensate fraction's delay on the interaction strength. To determine the delay, the cross-correlation between density and condensate fraction is calculated. The cross-correlation $CC(\tau)$ between two quantities $A(t_i)$ and $B(t_i)$ sampled at the times t_i is defined as [156]

$$CC(\tau) = \frac{\sum_i A(t_i - \tau)B(t_i)}{\sqrt{\sum_i A(t_i)^2 \sum_i B(t_i)^2}} \tag{6.11}$$

where τ is an optional time-shift between A and B. If $A(t_i - \tau)$ is proportional to $\pm B(t_i)$, $CC(\tau)$ will take on the value ± 1. $CC(\tau) = 0$ on the other hand indicates that no linear trend relates the two quantities. Consequently, if the time evolution of B qualitatively follows that of A with some delay τ', the cross-correlation exhibits a local maximum at τ'. Since we observe a delay of the condensate fraction with respect to the density, a local maximum should be found in the cross-correlation for small positive delays if we equate B with the condensate fraction and A with the simplified density. To enhance the interesting oscillatory signal, the time-average of the two measured quantities is subtracted before calculation of the cross-correlation. The expected maxima are indeed found for all measured quenches and an example is shown in figure 6.40. The value of τ' is extracted by fitting the empirical function of a Gaussian peak on top of an exponentially decaying background to the cross-correlation data. The asymmetric background is related to the fact that data points have to be dropped to calculate CC for non-zero values of τ. In the case of $\tau > 0$, the longest hold times of the simplified density and the shortest of the condensate fraction are dropped. The qualitatively different behaviour of the quantities at long and short hold times then causes the asymmetric background. For delays of around -2.5 ms, a second peak appears in the cross-correlation due to the period of the observed oscillations. It is larger than the relevant peak at small positive values of τ due to the asymmetric background.

The dependence of the condensate fraction's delay τ', normalised by the Fermi time, on the

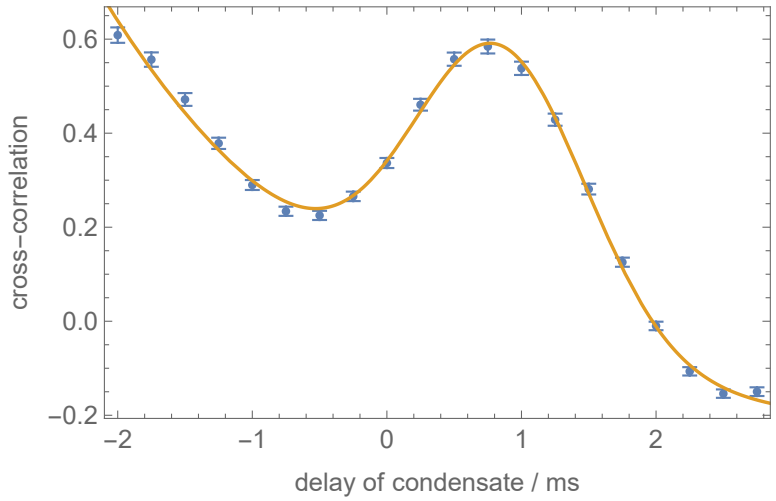

Figure 6.40: Cross-correlation between the long time dynamics of simplified density and condensate fraction for the quench to unitarity. The line is an empirical fit to extract the local maximum around 1 ms which is the delay of the condensate fraction to the density. The increase of CC towards negative delays is caused by the near-periodicity of the dynamics.

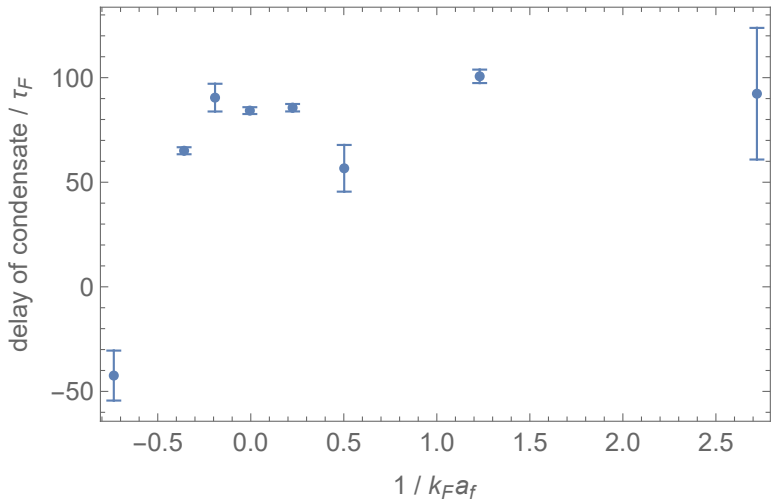

Figure 6.41: Interaction strength dependence of the condensate fraction's delay determined from the cross-correlation.

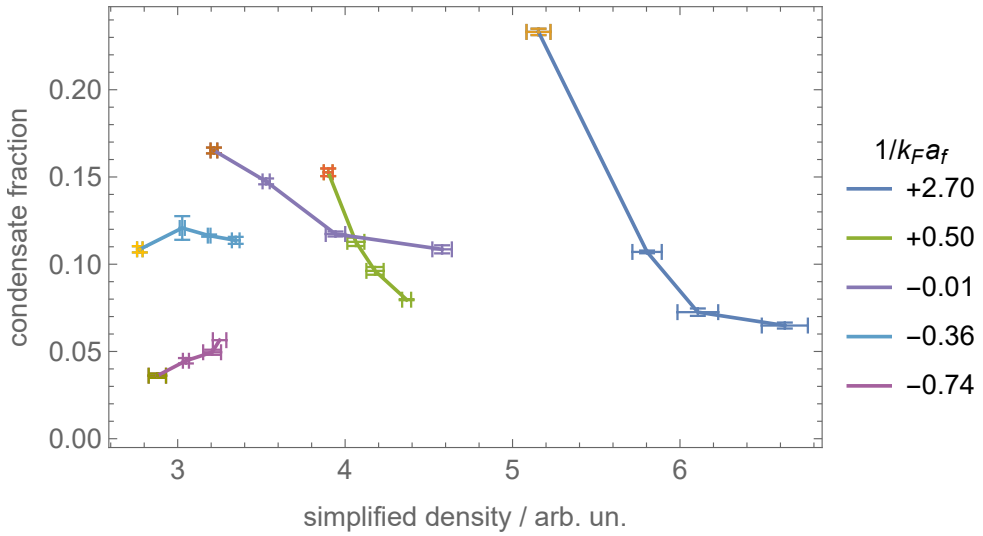

Figure 6.42: Evolution of the condensate fraction as a function of the density after the quench. Points with identical scattering lengths are connected in the order of their respective hold times. The shortest hold time, usually 100 µs, is highlighted in a different colour. The data shown cover the first 600 µs to 800 µs of the respective measurements.

interaction strength after the quench is shown in figure 6.41. It reveals an increase of the delay with $1/k_\mathrm{F}a_\mathrm{f}$ from the BCS limit up to unitarity from where it stays roughly constant. For $1/k_\mathrm{F}a_\mathrm{f} > 0.4$, the delay is however less reliable because the peak broadens into a plateau. Consequently, the local maximum is harder to determine as also reflected in the error-bars of figure 6.41. The observation of an increasing delay with attraction strength is counter-intuitive because the dynamics of the condensate should, at least in a very simple understanding, be related to the timescale set by the order parameter Δ and therefore be able to follow the density faster the higher $1/k_\mathrm{F}a_\mathrm{f}$. Another plausible expectation is a dependence on the effective interaction strength of the gas[5], peaked around unitarity, because it could enable a faster thermalisation. A possibility for the observed behaviour is a dependence of the delay on the quench strength $1/k_\mathrm{F}a_\mathrm{f} - 1/k_\mathrm{F}a_\mathrm{i}$ which increases with increasing $1/k_\mathrm{F}a_\mathrm{f}$ for constant magnetic field quench strengths, see section 6.2.4. A stronger quench should bring the system further away from equilibrium and therefore create more quasiparticles which could slow down the growth of the superfluid and lead to larger delays.

6.3.3 Initial compression

A remarkable feature from the interaction strength-dependent delay not mentioned so far is the negative delay observed for $1/k_\mathrm{F}a_\mathrm{f} = -0.74$, indicating that the evolution of the condensate fraction precedes that of the density. This can be understood by examining the effect of the initial contraction of the cloud on the condensate fraction. Figure 6.42 visualises this by showing the condensate fraction as a function of the simplified density n'. As already known from the previous considerations, the density always increases immediately after the quench. The condensate fraction however behaves differently depending on the interaction strength. For the final interaction furthest towards the BCS

[5] This decreases towards the BEC limit because the dimers only interact weakly with each other.

regime, $1/k_\mathrm{F}a_\mathrm{f} = -0.74$, the condensate fraction increases with the density while it decreases around unitarity and on the BEC side of the crossover.

The increasing condensate fraction with compression for $1/k_\mathrm{F}a_\mathrm{f} = -0.74$ is the origin of the negative delay of the condensate fraction. The response to the initial compression seems to vary smoothly with attraction strength: around $1/k_\mathrm{F}a_\mathrm{f} = -0.36$, the condensate fraction is almost unaffected by the compression and the decrease of the condensate fraction with compression seems to accelerate continuously towards the BEC limit.

The observed effect of the initial compression is a further indication that the local variation of $1/k_\mathrm{F}a$ due to the decreasing density is of lesser importance for the long time dynamics. $1/k_\mathrm{F}a$ always tends towards unitarity with increasing k_F and a qualitative change of behaviour in response to a compression would therefore be expected at unitarity and not at negative scattering lengths as observed. Instead, the temperature probably plays a more important role if it is even possible to speak of temperature in this non-equilibrium state. Similar to the compression of a non-interacting gas, the Fermi gas should heat up if compressed. This is even more reasonable if one considers that it is not coupled to a thermal bath but reaches its equilibrium temperature, set by the trap depth, through the very slow process of evaporation – happening on a timescale of seconds. Consequently, the compression is more or less adiabatic and heats the cloud thereby reducing the condensate fraction. At the same time however, the absolute value of the critical temperature changes due to the change in $1/k_\mathrm{F}a$ and T_F. Towards the BCS limit, where the critical temperature increases exponentially with $1/k_\mathrm{F}a$, the increase in T_c can be expected to surpass the increase in T caused by the compression. This results in a condensate growth with compression. Towards unitarity, the increase in T_c with $1/k_\mathrm{F}a$ slows down compared to the BCS limit while the heating caused by an adiabatic change in $1/k_\mathrm{F}a$ increases [67]. Together these effects could explain the observed decrease in condensate fraction with compression. In the BEC limit, $1/k_\mathrm{F}a$ has basically no effect on the critical temperature and the behaviour only depends on the compression-induced heating and the change of T_F, which affects the absolute value of T_c. Since T_F depends relatively weakly on the density, it is not surprising that the heating appears to be the dominant effect and reduces the condensate fraction.

Background of the fast dynamics

The evolution of the condensate fraction during the initial compression, happening during $t_\mathrm{hold} < 800\,\mu\mathrm{s}$, is also relevant in another context: it governs the background for the short time dynamics, which last for 150 µs at most. This opens the question whether the various observed behaviours of the background can be explained by the compression effect alone. Figure 6.43 shows that this is not the case. For a quench of 35 G to $1/k_\mathrm{F}a_\mathrm{f} = -0.36$, the condensate fraction stays roughly constant during the compression. The observed oscillation however is not centred around the value to which the condensate fraction tends after several hundred µs. Therefore the background is more involved than just the compression effects and the piecewise function from equation 6.2 is used. This however is not surprising because the quench increases the equilibrium condensate fraction and relaxation to this new value is expected to be fast and independent of any compression effects

6.3.4 Signs of non-equilibrium

In the previous considerations of the time-evolution of condensate fraction and density, the internal state of the gas was not analysed very systematically. A question arising here is whether the system

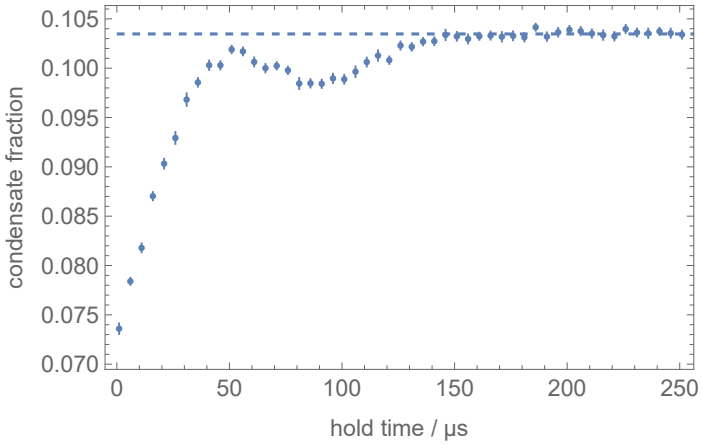

Figure 6.43: The observed short time dynamics with the constant background inferred from the initial compression of the cloud for a quench of 35 G to $1/k_{\mathrm{F}}a_{\mathrm{f}} = -0.36$.

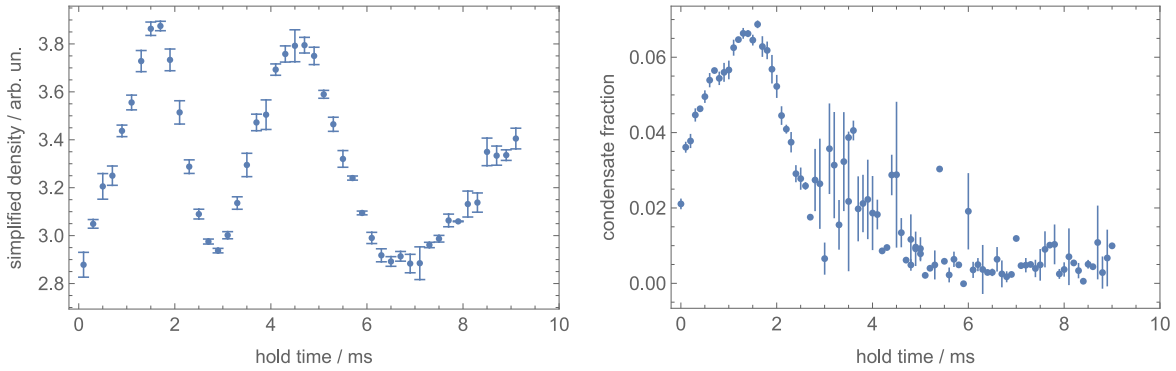

Figure 6.44: Time evolution of density and condensate fraction after a quench from $1/k_{\mathrm{F}}a_{\mathrm{i}} = -1.00$ to $1/k_{\mathrm{F}}a_{\mathrm{f}} = -0.74$. After a full oscillation, the condensate is depleted and does not re-form when the density reaches values that supported a condensate previously.

equilibrates quickly after the initial quench or stays in a non-equilibrium configuration for an extended amount of time. If the former case holds, the system should follow the induced collective mode as a thermal state. In the latter case however, the system would be in a non-thermal state during the collective mode and should exhibit deviations from thermal state characteristics. This would be a proof that the immediate quench effects influence the system over extended amounts of time.

A first evidence for the non-equilibrium situation is found in the behaviour after a quench relatively close to the BCS limit where the condensate fraction is very low. The induced collective mode destroys the initially present condensate around the point of maximal expansion after one period as shown in figure 6.44. When the cloud contracts again, the cloud reaches a configuration which supported a condensate previously but the condensate does not re-form. This implies that the system either heated up above its critical temperature or is in a non-thermal state inhibiting condensation. In either case, it can be concluded that the system had not fully equilibrated before the depletion of the condensate indicating that the quench affects the gas for relatively long times of several ms.

Another indication for a non-thermal state is found in the bimodal momentum distribution recorded

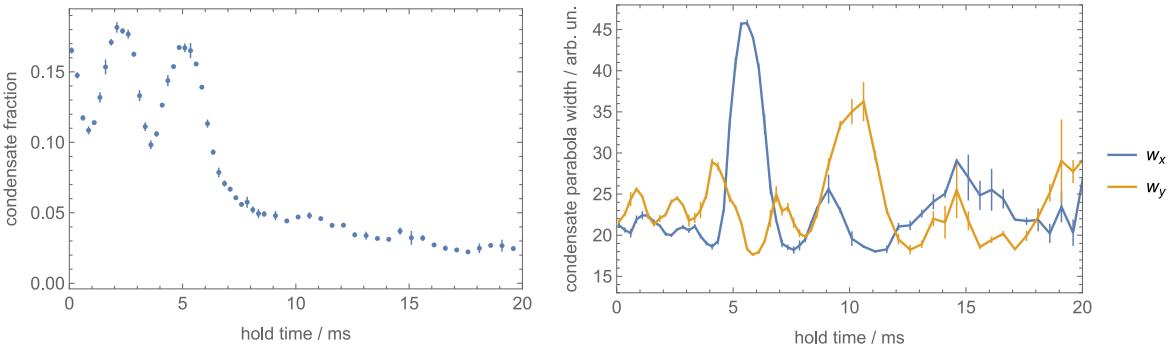

Figure 6.45: Time evolution of the condensed peak width in the momentum distribution after a rapid-ramp. The opposite phase of the oscillation of the two widths and the large amplitudes at longer hold times indicate a highly non-thermal state. For comparison, the time evolution of the condensate fraction is shown as well. The data are for a quench from $1/k_F a_i = -0.57$ to $1/k_F a_f = -0.01$

after the rapid ramp. The width of the condensed peak, the parameters w_x and w_y from equation 4.1, has a rather peculiar time evolution shown in figure 6.45. Initially, w_x and w_y oscillate opposed to each other reminiscent of a quadrupole oscillation. This oscillation is however twice as fast as the corresponding one in the condensate fraction which in turn is already faster than the oscillation of the cloud's size in x- and y-directions. The changing shape of the condensed peak can therefore not be a momentum signature of the collective mode but must be caused by internal dynamics of the cloud originating in its non-equilibrium configuration. Another feature contradicting the quadrupole origin is the initial equilibrium size of the condensed peak serving as an (almost) upper or lower bound for the size in x- and y-direction respectively during the first few oscillations. In a quadrupole oscillation, the symmetric configuration should equal the time-averaged one and this is not the case here.

After a few oscillations, an enormous increase in w_x is observed which coincides with a minimum in w_y resulting in an extremely asymmetric condensed peak. A situation that does not happen in equilibrium for fermionic s-wave superfluidity which should always have a symmetric momentum distribution[6]. Around the hold time at which the strong asymmetry is observed, the condensate fraction cedes its strong oscillation and starts to converge to a new, much lower equilibrium value with a much weaker residual oscillation. It is also found that the thermal background of the bimodal momentum distribution increases in width rather suddenly around the time the maximum in w_x is observed. This indicates that heating takes place at this time. Afterwards the thermal component's shape stays constant while another asymmetry of the condensed peak is observed, this time stretched in y-direction, before it returns to a more symmetric time evolution.

The clear deviation of the condensed peak behaviour from the equilibrium expectations demonstrates that the system is in a highly non-equilibrium state and exhibits rich internal dynamics which exceed those caused by the induced collective mode. The initial interaction quench therefore influences the system for prolonged times. This observation agrees with the predictions of [157] which state that a Fermi gas in the BCS-BEC crossover stays in a non-thermal state for prolonged times after a fast interaction quench. A relatively complicated collective mode caused by an interaction quench in a system trapped in only one of the three spatial dimensions has been predicted in [158]. The predicted

[6] It should however be noted that at least in the non-interacting BEC limit, the condensed peak can be asymmetric in equilibrium because it is the momentum distribution of the single particle ground state.

mode decays with time causing a heating of the system which matches our observation at hold times of around 8 ms at unitarity. The same reference also identifies a non-spherical Fermi surface during the dynamics, which could be the reason for the observed anisotropy of the condensed peak after a rapid ramp.

6.4 Conclusion

This very long chapter contained the most important results of this thesis and to wrap up, these are briefly summarized. The stand-out result is the observation of heavily damped oscillations in the condensate fraction after interaction quenches in the BCS-BEC crossover. The high frequency of these results, on the order of the Fermi frequency, shows that they originate from internal dynamics in the gas and not the external trapping potential. It is therefore concluded that they are the manifestation of the Higgs mode which has been excited by the interaction quench. These are the firs clear observations of the Higgs mode in time domain throughout the crossover. Compared to the first measurement utilising the neural network, the signal-to-noise ratio of the fast dynamics was greatly improved by using the rapid ramp technique for detection of the condensate. In [150], published during the work on this thesis, a similar behaviour is visible in the presented data which only cover the unitary Fermi gas. A clear signal of a local Higgs mode at unitarity was published shortly before submission of this thesis in [151].

The frequency of the oscillations, normalised by the Fermi frequency in harmonic approximation, matches the expectations from a homogeneous theory. While the good agreement is certainly coincidental due to the very different systems and uncertainties in the Fermi frequency arising from anharmonicities in the trapping potential, it signals that the observed frequencies are in the order of magnitude expected from the Higgs mode. The inhomogeneity of the gas can be theoretically accounted for in a local-density approximation but it fails to predict the observed frequency by a factor of roughly 3. One of the LDA models however correctly predicts the observed damping of the oscillation as a result of dephasing. It should however be noted that the applicability of the local-density approximation in non-equilibrium is uncertain. Towards the BEC limit, for interaction parameters $1/k_{\mathrm{F}}a_{\mathrm{f}} > 0.5$ (determined from the central density), the observed oscillations decay increasingly faster than predicted by the LDA calculation. This could be the instability of the Higgs mode due to the loss of particle-hole symmetry which has been predicted multiple time [59, 60, 75]. The instability has been seen in the earlier spectroscopic measurement [57] but the excitation method made an interpretation of the result difficult. In the measurements presented here, the disappearance of fast dynamics after the quench can be traced through the BEC side of the crossover.

On longer timescales, governed by the external trapping potential, it is observed that an interaction quench to stronger attraction causes the cloud to contract thereby initiating a collective mode. The ensuing density oscillation in turn affects the condensate fraction leading to fairly complicated dynamics as predicted in [158]. Additionally, a clear signature of a non-thermal internal state of the gas is found in the momentum distribution recorded after a rapid ramp. Therefore, the observations are likely to arise from a complex interplay of local, internal and global, trap related dynamics.

CHAPTER 7

Interaction Modulation

A modulation of a superfluid's interaction strength can be used to probe its excitations. With the help of the Feshbach resonance, such a modulation can be realised by a sinusoidal variation of the magnetic field. This incorporates a drive at a fixed frequency and forces the system to follow the external drive possibly creating excitations in this way. The modulation is therefore a complimentary approach to the quench presented in the previous chapter, which caused the system to evolve according to its own dynamics.

In a modulation scheme, excitations reveal themselves as features, often dips or peaks, in frequency space. Since these are easy to analyse, it is often the preferred way to determine the frequencies of excitations such as the Higgs mode. In a superfluid however, complications arise from the quasiparticle excitations if the Higgs mode is the object of interest. Atom number conservation, which is realised in our and most other cold-atom experiments, forces quasiparticles to be excited in pairs only – breaking a condensed pair creates two quasiparticles. Consequently, quasiparticle excitations have a threshold frequency of $2\Delta/\hbar$ which is identical to the Higgs mode's frequency. A clear distinction between these very different excitations is therefore not trivial. At the experimental platform of this thesis, a spectroscopic measurement of the Higgs mode was previously performed by modulating the population of one spin state [57]. The resulting signal could be identified as the Higgs mode after thorough theoretical modelling.

An interaction modulation with a time-dependent magnetic field has already been conducted in 2005, shortly after the first experimental realisations of the BCS-BEC crossover [159]. The frequency resolution was however not sufficient to investigate the Higgs mode, which was not of primary concern back then.

The project presented in this chapter is still ongoing and I was mainly involved in the discussion and interpretation of the already available results. It therefore includes only an overview of the experimental implementation and some first results.

7.1 Experimental implementation

To study the Higgs mode, the interaction strength has to be manipulated at the Higgs frequency which is expected to be in the regime of a few 10s of kHz. Similar to the quench experiments, a

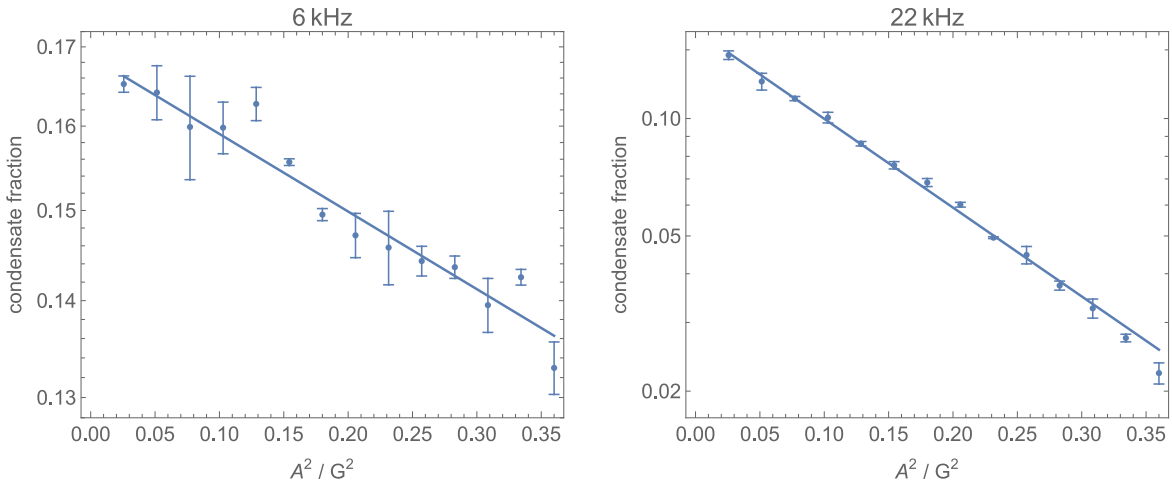

Figure 7.1: Amplitude dependence of the remaining condensate fraction after a modulation of 300 ms. The data were recorded at a magnetic field of 716 G ($1/k_\mathrm{F}a = -0.37$ in harmonic approximation). The line is an exponential decay fitted to the data, appearing as a straight line in the logarithmic plots. The modulation frequency of 6 kHz (22 kHz) in the left (right) panel is expected to be far below (above) the gap of the system.

$|13\rangle$-mixture of ^{6}Li is prepared in the crossed beam dipole trap. The quench coil can also be used for this modulation if a different driving circuit, purpose built for this application, is chosen. Details on the driving circuit can be found in [93]. A modulation of the magnetic field leads to a modulation of the attraction strength and therefore of the pair size. If the frequency of an excitation is hit in this way, the system heats up because quasiparticles are created; either by directly breaking pairs or indirectly by the excitation of the Higgs mode which subsequently decays into quasiparticles. The increase in temperature is in turn easily observable as a reduction in the condensate fraction, which is measured with the rapid ramp.

The important parameters of the sinusoidal modulation of the interaction strength are its frequency ω_mod, its amplitude A, and the modulation time t_mod. If the modulation time is kept constant, it is found that the condensate fraction decreases exponentially as

$$\frac{N_0}{N} \propto \exp\left(-\Gamma A^2\right) \tag{7.1}$$

in response to the drive, with the frequency-dependent condensate depletion rate Γ. This is shown in figure 7.1. For short t_mod, the decay is also proportional to $\exp(-t_\mathrm{mod})$. Both of these observations agree with the predictions from [160]. For longer t_mod of ~ 100 ms and more, deviations from the exponential dependence on t_mod are however observed. They are caused by evaporative cooling of the gas which increases the condensate fraction and acts against the reduction caused by the modulation. In absence of the drive, the growth due to evaporation is measured to be exponential with a time constant of ca. 180 ms. Additional deviations are found if the condensate is almost entirely depleted. Potential reasons are a modified response close to T_c and difficulties in detecting small condensate fractions reliably.

The quantity of interest in these measurements is the dependence of Γ on ω_mod and $1/k_\mathrm{F}a$, the mean interaction strength during the modulation. The best results are achieved with a modulation time

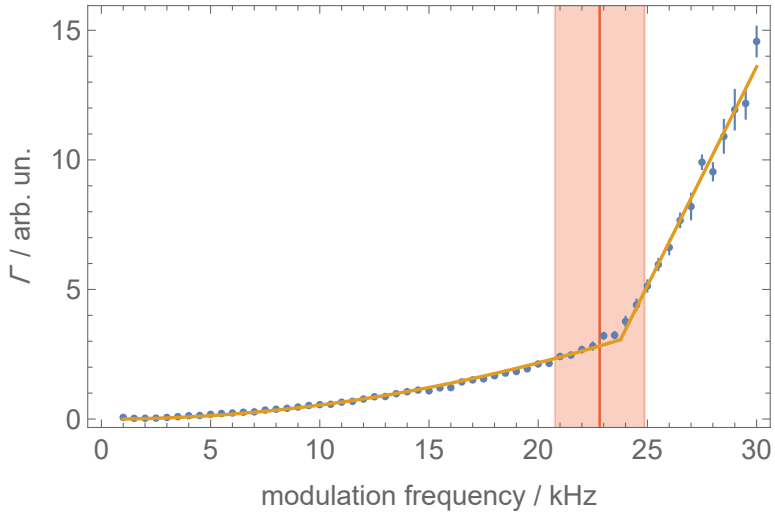

Figure 7.2: Frequency dependence of Γ at 660 G, corresponding to an interaction parameter of $1/k_F a = 0.82$ in harmonic approximation. The orange line is a piecewise fit, quadratic for low and linear for high frequencies. The red line is the binding energy of the dimers and the shaded region its uncertainty under the assumption of a total magnetic field uncertainty of 1 G.

of 300 ms. Keeping t_{mod} constant for all modulation frequencies allows for the use of equation 7.1 to determine Γ. A is not kept constant for the measurements because Γ can vary by some orders of magnitude depending on ω_{mod}. An adjustment of A is therefore necessary to always have a measurable drop in the condensate fraction without ever coming close to depleting it. In general, the amplitude is kept relatively low on the order of 0.1 G to 0.3 G.

7.2 Results

7.2.1 Excitation onset

In a homogeneous system, the gapped excitation spectrum of superfluids leads to an onset frequency corresponding to the minimum energy necessary for the excitation of a pair of quasiparticles. Above this frequency, there is a continuum of excitations with arbitrary energy because the momentum of the individual quasiparticles is not restricted as long as the total momentum of the pair is zero. In our inhomogeneous gas, the superfluid gap however varies spatially, according to the local density approximation, and an onset is therefore not necessarily expected. Nevertheless, a gap is still expected to appear at positive scattering lengths because even the thermal atoms are bound into dimers at these fields.

Figure 7.2 shows the frequency dependence at a magnetic field close to the BEC limit and an onset frequency for a strong depletion of the condensate is indeed found at a frequency that matches the dimer binding energy. Below this threshold frequency, the condensate fraction is still reduced by the modulation, an observation that has not been made in the old measurement [159]. The only excitations of the system in this frequency range should be phonons (density modulations). In general, it seems reasonable that sound modes can be excited by the interaction modulation because the induced oscillation of the pairs' sizes could lead to density fluctuations. It is not clear yet, why these excitations

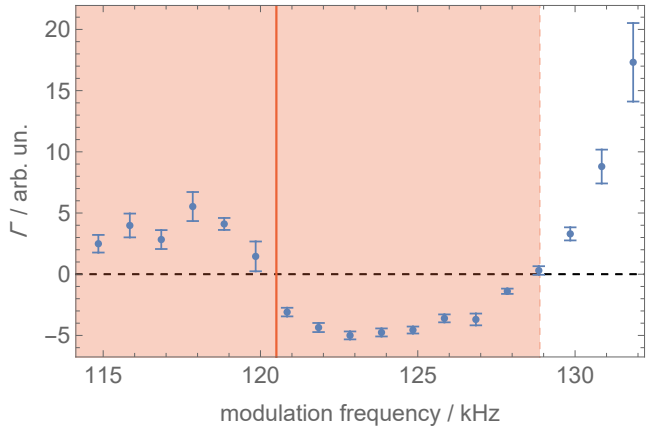

Figure 7.3: Frequency dependence of Γ around the dimer binding energy, indicated by the red vertical line, in the BEC limit. The shaded area is the uncertainty of the dimer binding energy under the assumption of a magnetic field uncertainty of 1 G. The interaction parameter has not been determined precisely but is $1/k_{\mathrm{F}}a \sim 1.5$ in harmonic approximation. Around the dimer binding energy, a dip in the depletion rate that even extends to negative values is visible.

were not observed in [159]. A possible explanation could be their different trap geometry with one weakly and two tightly confined axes compared to our three relatively similar axes. A lower signal to noise ratio in the old measurement could however also have simply obscured these excitations.

If the system is brought closer to unitarity, the onset frequency stays at the diminishing binding energy. It is however found that the frequency dependence of Γ at the binding energy becomes increasingly smoother. Already at positive scattering lengths, the threshold can no longer be reliably identified. It is therefore, at least at the current state, not possible to investigate the threshold's behaviour throughout the crossover.

In the far BEC limit, where the dimer binding energy is larger than the expected value of Δ, an unexpected feature has been found: a local minimum in Γ around the dimer dissociation frequency, see figure 7.3. This means that the heating effect by the modulation is diminished in a relatively narrow frequency range. Above the dissociation frequency, Γ still increases strongly, in a similar fashion as discussed above. The minimum deepens with increasing $1/k_{\mathrm{F}}a$ (decreasing magnetic field) and the minimal value of Γ can even be negative. This which corresponds to a growth of the condensate or a cooling of the cloud due to the modulation. The condensate fraction after modulation (CF) normalised by its value without modulation (CF_0) for various fields is shown in figure 7.4. The interaction parameters at these fields have not been measured yet. However, the lowest value of $1/k_{\mathrm{F}}a$ in the figure, reached at the highest field, is roughly $1/k_{\mathrm{F}}a = 1.4$ in harmonic approximation. It is clearly visible that the condensate growth caused by the modulation increases in strength the further one goes into the BEC limit.

So far, this feature is not understood. Since it is around the dimer binding energy, it could be caused by a process which selectively brakes thermal dimers into unpaired fermions which no longer take part in the Bose statistics. This constitutes an effective cooling of the molecular Bose gas since it removes the more energetic bosons. Factors that could potentially play a role in this process include

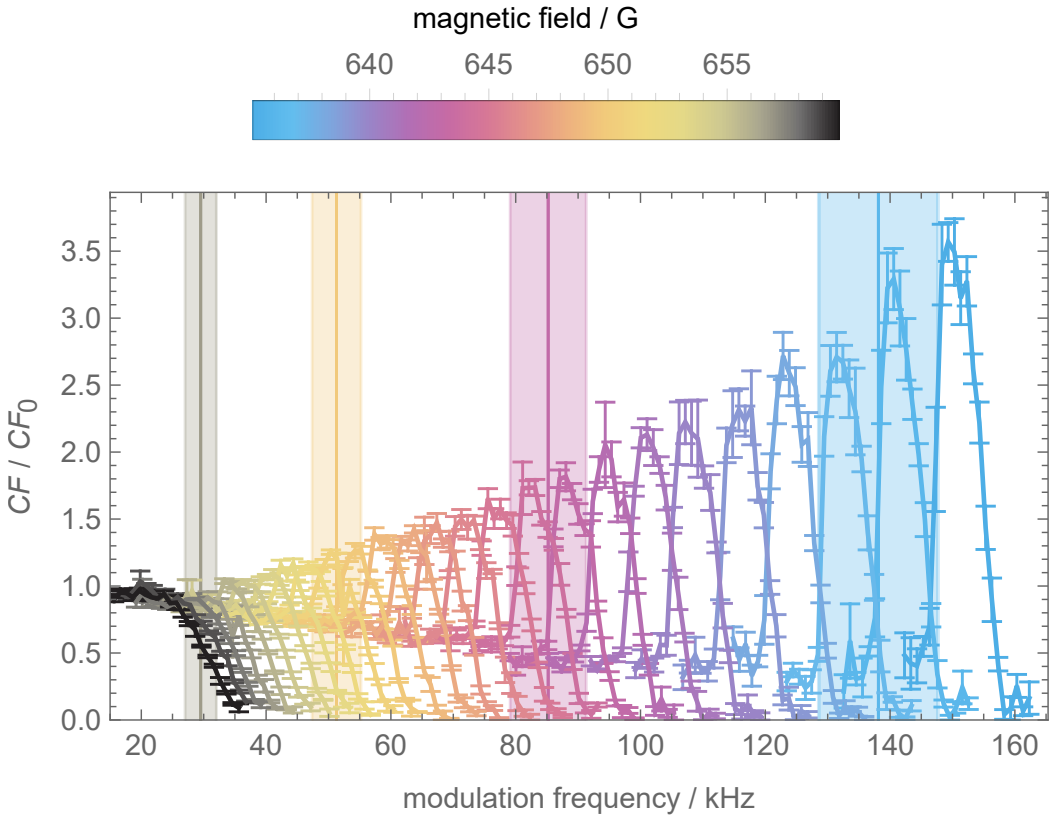

Figure 7.4: Increase of the condensate fraction due to modulation around the dimer binding frequency in the BEC limit. The vertical lines are the dimer binding energies for selected fields and at slightly lower frequencies than the maximal values of CF/CF_0 at the same field. The shaded region around the vertical lines is the uncertainty in the dimer binding energy caused by an uncertainty in the magnetic field value of 1 G.

interactions of the unpaired fermions with the surrounding molecules[1] and the spatial separation of the condensate in the centre of the trap from the thermal dimers around it. More research is however needed to elucidate these effects.

In figure 7.4, the dimer binding energy has been calculated by the simple relation $E_B = -\hbar^2/ma^2$ which underestimates the binding energy in the BEC limit [162]. The uncertainty caused by the magnetic field is however larger than this systematic error even for the largest binding energies considered here.

7.2.2 Resonance

The Higgs mode has a more or less well-defined frequency, depending on its stability, and should therefore only be excitable in a rather narrow frequency range. To investigate whether it can be observed by the interaction modulation, the modulation frequency is varied in smaller steps around the expected, $1/k_F a$-dependent frequency of the Higgs mode. A peaked feature is indeed found on

[1] The corresponding scattering length a_{am} is $a_{am} \approx 1.2a$ [161]

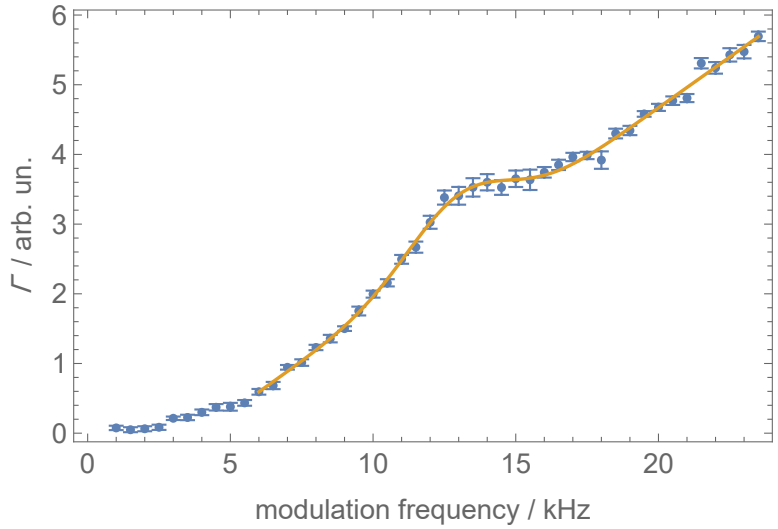

Figure 7.5: Observed resonance of Γ at 716 G, $1/k_{\mathrm{F}}a = 0.37$ in harmonic approximation. The orange line is a Gaussian peak on top of a linear background fitted to the data, it is only plotted in the frequency range that was included in the fit.

top of the background discussed in the previous section as shown in figure 7.5. The feature's central frequency ω_0, normalised by ω_{F} in harmonic approximation, is found to depend on the interaction strength. The frequencies can be determined by fitting a Gaussian peak on top of a linear background to the observed data. These frequencies are compared to theoretically expected Higgs frequencies of a homogeneous gas with identical Fermi energy in figure 7.6 and found to be consistently higher than the theoretical value. This could be an indication of the expected increase of the gap in a trap: the attraction between fermions causes a contraction of the cloud, an effect absent in the homogeneous case, leading to an increased density and thus a larger gap. The contraction and therefore also the increase in the gap are expected to increase towards the BEC limit and this increase is also observed for the fitted frequencies ω_0.

It is possible to compare the measured peak to a simulated spectrum of the Higgs mode, calculated in local-density approximation from the reconstructed density. The procedure is identical to the calculated dynamics from section 6.2.3 of the quench chapter. The calculation results are compared with the fitted peak centres in figure 7.7 and very good agreement is found at negative scattering lengths and unitarity. Towards the BEC limit, there is however an increasing disagreement between the location of the calculated and observed peak in the spectrum. Since the calculation considers the increased density due to attractions, another effect must lead to higher resonance frequencies towards the BEC limit.

Compared to the frequencies observed after an interaction quench, the central frequency of the peaked feature is consistently located at higher frequencies. The peak is also present well into the BEC-limit while a meaningful frequency could not be determined in the quench experiments for $1/k_{\mathrm{F}}a_f > 0.4$ in harmonic approximation.

From the frequency alone, the feature cannot be identified as the Higgs mode because the

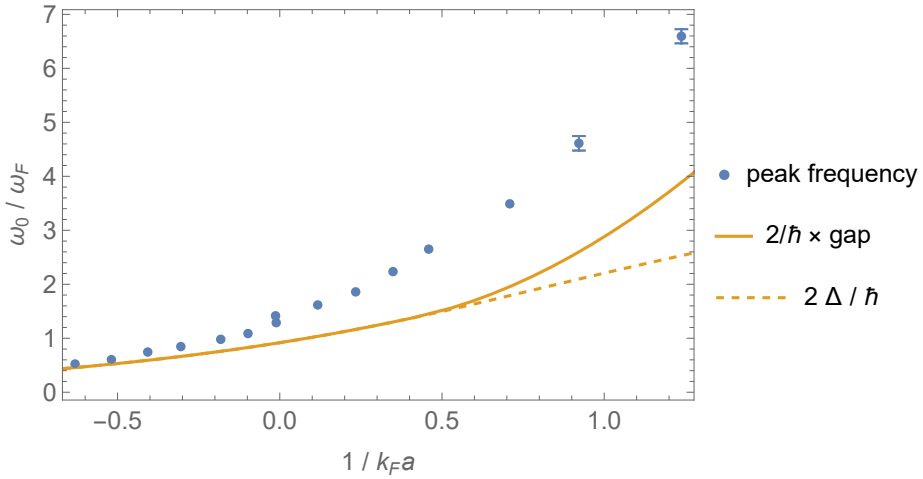

Figure 7.6: Central frequencies of the peaks ω_0 in the whole crossover normalised by the Fermi energy in harmonic approximation. They are compared to the Higgs frequency/quasiparticle breaking threshold expected from the theory for a homogeneous gas [67].

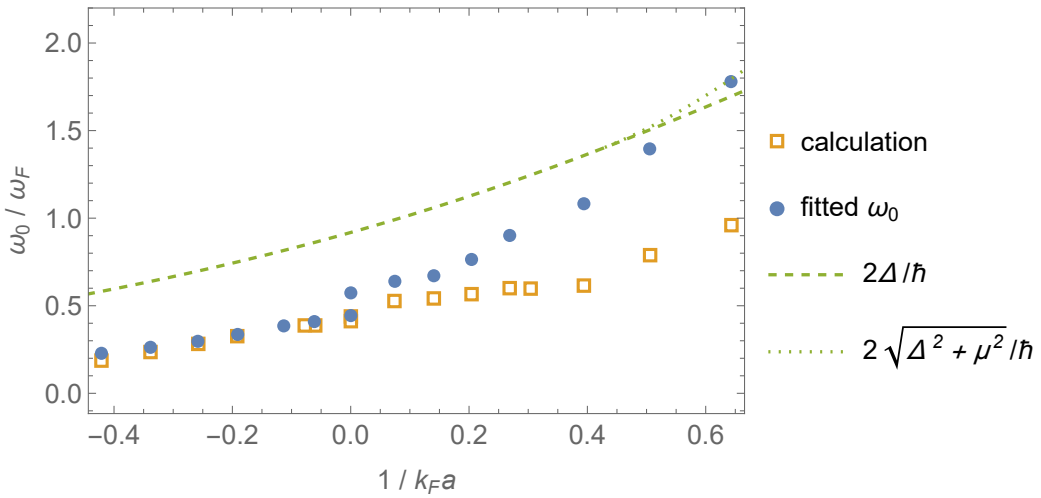

Figure 7.7: Comparison of the central frequency of the observed peaks compared to the main frequency of Higgs spectra calculated in local density approximation from the reconstructed density distributions. The theoretical predictions for a homogeneous gas [67] are provided for comparison.

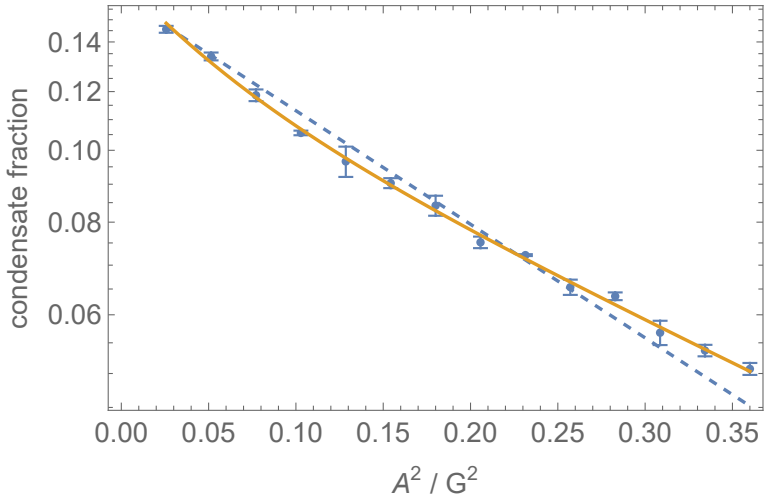

Figure 7.8: Amplitude dependence of the condensate fraction after a modulation at 13.5 kHz, close to the resonance frequency from figure 7.5. The orange line is the sum of two exponential decays fitted to the data while the dashed blue line is a single decay as in equation 7.1. The dashed line corresponds to the Γ shown in figure 7.5.

quasiparticle continuum should have the same threshold frequency. Since the quasiparticle dispersion relation has a minimum at the gap energy, there should be a peak-like feature in their density of states. This could in turn cause the excitation and condensate depletion rates to peak at this frequency, similar to the observations.

A peculiarity is found around the resonance in the amplitude dependence of the condensate fraction after modulations with constant duration. In contrast to the previously discussed exponential decay from equation 7.1 for frequencies above and below the resonance, the amplitude dependence around the resonance is better described as the sum of two exponential decays, see figure 7.8. This means that small modulation amplitudes disturb the superfluid much stronger than expected around the resonance frequency which manifests a qualitative change in the response to the modulation. This could be an indication that the observed resonance is indeed the Higgs mode and not caused by a maximum in the quasiparticle density of states. It should be noted that in the derivation of the resonance shown in figure 7.5, the double exponential behaviour around the resonance frequency was not included. Instead, Γ is calculated from a simple exponential fit, the dashed line in figure 7.8.

Another indicator of the Higgs mode should be its increasing instability towards the BEC limit which would manifest itself as a continuously broader peak in the depletion rate. Such a broadening is indeed observed for the resonances. However, the ratio of peak width to central frequency is constant for all observed interactions. More specifically, it is found that the full width at half-maximum (FWHM) and the central frequency ω_0 are related by roughly $\mathrm{FWHM}/\omega_0 = 0.39$. This ratio being fixed is consistent with a decay dominated by dephasing similar to the observations from the interaction quenches.

The width of the peak can be compared to the fitted lifetime τ of the oscillation following the quench. For a Lorentzian peak, the associated lifetime is given by $\mathrm{FWHM} = 2/\tau$. Under the assumption that this relation roughly holds for the observed peaks, the ratio FWHM/ω_0 from the modulation spectra can be compared to the combination $2/\tau\omega_\mathrm{H}$ of fit parameters from the measurement in time domain.

Figure 6.20 shows that $\tau\omega_{\mathrm{H}}$ usually lies in the range of 2 to 3, which implies that $2/\tau\omega_{\mathrm{H}}$ lies between 0.66 and 1. If the peak in the modulation spectrum corresponds to an oscillating mode, its property $\mathrm{FWHM}/\omega_0 = 0.39$ indicates that is slightly more long-lived than the oscillation following a quench. It is however also noted that the peak in the modulation spectrum has the property $2\omega_0/\mathrm{FWHM} = 5$. This is similar to the prediction of some models for local Higgs modes under local density approximation, see figure 6.20.

7.3 Conclusion

In this experiment, the interaction strength of a Fermi gas in the BCS-BEC crossover was modulated. Even though a similar experiment has been performed before [159], several new features have been found. The most important of these is probably a peak in the frequency dependence of the condensate depletion rate. The central frequency of the peak increases with attraction strength and roughly correspond to expectations of the Higgs mode. This approximate agreement is also found in the width of the peak but the expected increase of the width towards the BEC limit, where the Higgs mode becomes increasingly unstable, is not observed. It is also not possible to fully rule out a connection of the peak to quasiparticle excitations whose density of states is peaked at the frequency of the Higgs mode. A better theoretical understanding is probably necessary to resolve this uncertainty. An anomaly in the condensate fraction's dependence on the modulation amplitude could be a promising starting point in this regard.

Similar to the older experiment [159], we observe a strong depletion of the condensate fraction above a threshold frequency for positive scattering lengths sufficiently far away from the unitary limit. Additionally, we also detect excitations which lead to a much weaker heating of the condensate below the threshold frequency; these could be connected to the excitation of sound modes. [159] reports no such excitations but they could have been obscured by a low signal-to-noise ratio. It is however also possible that the different trap geometries have an effect on the possibility to excite sound modes.

The most unexpected feature found in with the modulation scheme so far is a very ineffective heating or even cooling just below the threshold frequency for fast depletion in the BEC limit where $E_{\mathrm{B}} \gg \Delta$. Possible explanations are still on a very vague and speculative level and more investigations are certainly necessary.

CHAPTER 8

Summary and Outlook

In this thesis, the BCS-BEC crossover was studied with a focus on the critical temperature and excitations in the superfluid phase. To this end, quantum simulations were performed with an ultracold gas of ^{6}Li atoms prepared in a balanced mixture of two spin states. It is confined in an optical dipole trap which leads to an inhomogeneous density distribution. We developed new detection techniques for superfluidity based on machine learning that are able to identify the superfluid phase directly from the atoms' momentum distribution. In a supervised machine learning approach, it was even possible to train a convolutional neural network to quantify the condensed fraction of the gas. The labelled dataset necessary for training in this case was generated with the traditional, not fully understood method of projecting the momentum distribution to tightly bound dimers to measure the condensed fraction ("rapid ramp" method). Even though the trained network "inherits" most conceptual difficulties from the rapid ramp, it still constitutes an improvement by being able to unambiguously identify the phase transition. A conceptual advancement away from the rapid ramp was achieved with an unsupervised machine learning approach based on the autoencoder network structure. This made it possible to detect the phase transition for a wide range of interaction strengths in the crossover for the first time without relying on a projection of the momentum distribution in any way. The critical temperatures detected with both methods agreed with each other. Unfortunately, the results of the unsupervised approach have a higher uncertainty than those from the rapid ramp so that it is not possible to judge the accuracy of the rapid ramp technique in general terms. It could be interesting to investigate whether other unsupervised machine learning approaches like the intrinsic dimension [163] can surpass the accuracy achieved by the autoencoder.

By reconstructing the density distribution of the gas and employing the local-density approximation, we determined the critical temperature of a homogeneous Fermi gas for various interaction strengths in the BCS-BEC crossover and observed for the first time an increase of the critical temperature from unitarity towards the BEC limit. The measured critical temperatures agreed with several theories at different interaction strengths but no theory matches in the whole crossover. A local maximum at small positive interaction parameters is in agreement with the results of the supervised approach but cannot be confirmed due to limitations from the thermometry. An improvement could therefore be achieved by better thermometry techniques. Especially suitable for our experiment should be the admixture of bosonic Na atoms whose temperature is much easier to determine. Right now, the mass difference between Na and Li with its different gravitational sag leads to a separation of the elements in the dipole trap. More involved trapping potentials should however be able to overcome this limitation,

for example by also including repulsive beams in the trapping scheme.

While most of the machine learning results were already presented in the theses of my collaborators [93, 103], the most important new result of this thesis is the unambiguous detection of fast dynamics in the condensed fraction of the gas after an interaction quench. This was possible due to an increased time resolution of our implementation of the rapid ramp technique. The most important feature of the fast dynamics is a quickly decaying oscillation of the condensed number of atoms. This is the expected signal of the Higgs mode of the superfluid phase and would be its first observation in time domain throughout the BCS-BEC crossover[1]. The expected frequency of the oscillation in an inhomogeneous situation is not clear but the observed frequencies are in the right frequency range and correctly scale with the interaction strength. Another observation supporting the claim of an observation of the Higgs mode is the increasing instability of the oscillations towards the BEC limit, where the Higgs mode is expected to be unobservable. In contrast to the earlier, spectroscopic measurement [57], the increasing instability is observed directly and does not have to be deduced by interpretation of the result.

With pragmatic modelling of the Higgs mode in local-density approximation, it was possible to explain the observed damping of the oscillation as the dephasing of independent Higgs modes around unitarity and towards the BCS limit. Towards the BEC limit the observed oscillations decayed faster than the predictions of the model, which did not include the increasing instability of the Higgs mode in this interaction range. This is another indication that the Higgs mode in the crossover has been observed. The observable frequency predicted by the model differs from the actually observed ones by a factor of 2 to 3 and this discrepancy is unlikely to be fully explained by the finite interaction quench strength which was not considered in the model. Nevertheless, the model does predict the observed scaling of the frequency with interaction strength relatively well.

The main conceptual complications in the interpretation of the interaction quench results arise from the inhomogeneity of the gas caused by the external trapping potential. A big improvement would therefore be the realisation of a homogeneous Fermi gas in our experiment. This is possible by employing spatially structured, repulsive trapping beams, usually a hollow beam and flat light "sheets", and has been demonstrated for bosons [164] and fermions [165].

Another interesting study would be the Higgs mode of the two-dimensional version of the BCS-BEC crossover. In reduced dimensions, spontaneous symmetry breaking is not possible at finite temperatures [166, 167] and the consequences for the Higgs mode are not well understood. Very recently, it has been predicted that, at least at $T = 0$, the Higgs mode in the two-dimensional crossover behaves very similar to its three-dimensional analogue [168]. Experimentally, a two-dimensional system can be created by tightly confining the atoms in one spatial direction. If thermal energy and chemical potential of the gas are much lower than the first excited state of the confining potential, no movement is possible in this direction and the system becomes effectively two-dimensional [13]. Several aspects of the two-dimensional crossover have already been studied, e.g. in [82, 83, 169], but the Higgs mode has not been one of them. Some preliminary studies on the two-dimensional gas have already been performed at our experimental platform [170].

Finally, this thesis presented some first results from experiments in which the interaction strength between fermions in the BCS-BEC crossover was modulated. The spectrum of the condensate

[1] Similar observations at unitarity only were reported during the work on this thesis in [150] and immediately before submission in [151].

depletion rate showed a previously unobserved peak in the frequency range where the Higgs mode should be located. For scattering lengths $a \leq 0$, the observed resonance matches the expectations from a local-density model of the Higgs mode. Additionally, an altered dependence of the condensate depletion rate on the modulation amplitude around the resonance indicates that indeed some special response of the superfluid happens at these frequencies. Nevertheless, a full understanding of the observed spectra is not achieved yet and limited by lacking theoretical predictions. Again, the conceptually simpler homogeneous Fermi gas could help to clarify things in this case.

A surprising result of the modulation experiments has been a growth of the condensate fraction if the modulation frequency is close to the dimer binding energy in the BEC limit. Here, the dimer binding energy is the highest characteristic energy scale of the system. This feature has not been understood yet and more studies are certainly necessary for an explanation. It has, for example, not been checked yet whether the amplitude dependence of the depletion rate has some anomalies at these frequencies and interaction strengths. It is also not clear whether the growth by modulation is a consequence of the gases inhomogeneity or an intrinsic feature of the superfluid, a question that could also be answered by performing the same experiment with a homogeneous Fermi gas.

Overall, a key result is the observation of the Higgs mode in the BCS-BEC crossover in time domain and its increasing instability towards the BEC limit. Especially the observation of the instability could improve the theoretical understanding of the crossover because it can provide information on the degree of particle-hole symmetry present at various interaction strengths. Alternatively, one could also ask about the amount of symmetry necessary for an observable Higgs mode. The measurements can also serve as a benchmark for possible theoretical descriptions of the inhomogeneous Higgs mode. These are still lacking and therefore a quantitative explanation of the experimental results is not possible at the moment.

Another key result of this thesis is the measurement of the critical temperature for a wide range of interaction strengths in the BCS-BEC crossover. The measured temperatures agree well with theoretical predictions and are compatible with the maximum proposed in [71]. The measurement uncertainties however do not allow for a confirmation of this maximum whose precise characteristics therefore remain an open question.

A detection of the phase transition in a novel approach based on unsupervised machine learning was able to confirm the results obtained with a technique based on the rapid ramp over a large range of interaction strengths. Considering the applicability of the rapid ramp, the accuracy of the new approach is not high enough to draw general conclusions and more experiments are necessary to clarify this issue further.

APPENDIX A

Laser Systems

This appendix gives a brief summary on the setups used for generation of the laser light necessary for trapping, cooling and imaging the atoms as described in chapter 3. Detailed descriptions of the different setups can be found in earlier theses and will be referenced.

Laser system for Lithium atoms

The laser source for all 671 nm light are two homebuilt external cavity diode lasers with commercially available laser diodes[1]. Light from one of the diode lasers is amplified through a series of two tapered amplifiers[2] and used for the six beams of the Lithium MOT with around 16 mW per beam. An admixture of repumping light is added with the help of an electro-optic modulator (EOM).

Some of the light is used to lock the laser's frequency to the atomic reference of a Lithium-spectroscopy cell. The design of the cell is described in detail in [171]. A recent upgrade has been the replacement of the cell's viewports, which have to be heated to around 180 °C to prevent Lithium condensation on the windows. This temperature is very close to the old fused silica viewports' maximum rating of 200 °C and the vacuum broke when this temperature was greatly surpassed – most likely by a faulty power supply regulation. Now, viewports with sapphire windows that can withstand up to 400 °C have been installed. The new windows are birefringent and this can in principle affect the generation of a lock signal by mixing polarisations. After the installation of the new viewports however, the lock signal was immediately recovered without any noticeable effect of the birefringence.

The light from the second diode laser is amplified by a single tapered amplifier (model as above) and used to generate the light for the Zeeman slower, optical pumping and imaging. The various slightly differing frequencies for the beams and their respective repumping admixtures are adjusted with the help of acousto-optic modulators. The second diode laser is locked to the first diode laser via an offset lock. To image the atoms at the high magnetic fields of the Feshbach resonance close to 1 000 G, the laser has to be detuned substantially from the low-field resonance; for the ground state of the imaging transition this is evident from figure 3.2. The offset between the lasers in the lock can therefore be tuned over around 2.2 GHz with the help of an EOM. This sweep is performed slow enough to keep the second laser in lock with the first one.

[1] Eagleyard EYP-RWE-670-00702-100-SOT02-0000

[2] Eagleyard EYP-TPA-670-00500-2003-CMT02-0000

The details of the technical implementation and sketches of the laser system can be found in [79] and [80].

Laser system for Sodium atoms

The 589 nm light for trapping and cooling Sodium atoms is generated by a single external cavity diode laser[3] running at 1 178 nm. It is amplified with the help of a Raman fibre amplifier[4] and frequency doubled to 589 nm with the help of a lithium triborate crystal placed inside a bow-tie cavity. The cavity's output is split into several paths which provide the light for Zeeman slower, MOT, and optical pumping beams and imaging light. The correct frequencies are again set with the help of AOMs and EOMs. A small portion of the 589 nm light is used to lock the diode laser to the Na D_2 line with the help of a spectroscopy cell.

The laser system is described in detail in [79] and the spectroscopy cell in [172].

Plug laser

The repulsive light beam for the magnetic trap is generated with a commercial laser system[5], coupled through a photonic crystal fibre[6] and focussed onto the trap centre. The power of the beam can be stabilised with an AOM but this feature has not been used recently because the output power of the laser has degraded from 15 W to 10.4 W over time. Neither the reduced maximum power nor the unregulated final power after the fibre have had an observable impact on the experiment's performance. Details about this can be found in [79, 80].

Dipole trap laser

The laser source for the dipole trap is a commercially available fibre laser at 1 070 nm with an output power of up to 50 W[7]. The light is split into two paths for the two dipole trap beams and coupled into two fibres; one a photonic crystal fibre optimised for high powers[8] and the other a standard single mode optical fibre[9] which has been working well for many years now. The power in both paths is regulated with an AOM and a PID controller. A good power stabilisation is crucial for these beams because the evaporative cooling in the dipole trap relies on the controlled reduction of laser power. Additionally, fluctuating beam powers after evaporation would heat the cloud again. The construction of this laser system is documented in [173].

[3] Toptica ECDL DL pro
[4] 8 W 1 178 nm Raman Fiber Amplifier by MPB Communications Inc.
[5] Sprout-G Laser by Lighthouse Photonics
[6] NKT Photonics, Photonic Crystal Fiber Patch Cable LMA-PM-15-S/A, 2m, SMA905/FC APC
[7] IPG Photonics, model YLR-50-LP
[8] NKT LMA-PM-15
[9] Thorlabs P3-1064PM-FC-10

APPENDIX B

Channel Select Circuit

Circuit diagram

On the following two pages the circuit diagram of the channel select circuit is shown for reference.

Power supply calibration

The voltage from an analogue output of the control computer sets the current output of the power supplies that creates the Feshbach field at the atoms. The current needed for a specific magnetic field strength has been found by calibrating the Feshbach coils during the construction of the experiment [80]. Without the channel select circuit, the correct control voltage V for a given current I is given by the function $V(I) = 0.011729258 \frac{\mathrm{V}}{\mathrm{A}} I$.

After installation of the channel select circuit, it was immediately found that this formula is no longer valid to the necessary accuracy because the imaging light was no longer resonant to the atoms. This is a clear sign that a wrong magnetic field had been set. By measuring the current through the coils with a current transducer it was found that the current was still proportional to the control voltage but the proportionality factor had to be adjusted leading to $V(I) = 0.0117807021 \frac{\mathrm{V}}{\mathrm{A}} I$. With this small adjustment, the atom number was brought back to its old value and the channel select circuit could be used without any further modifications.

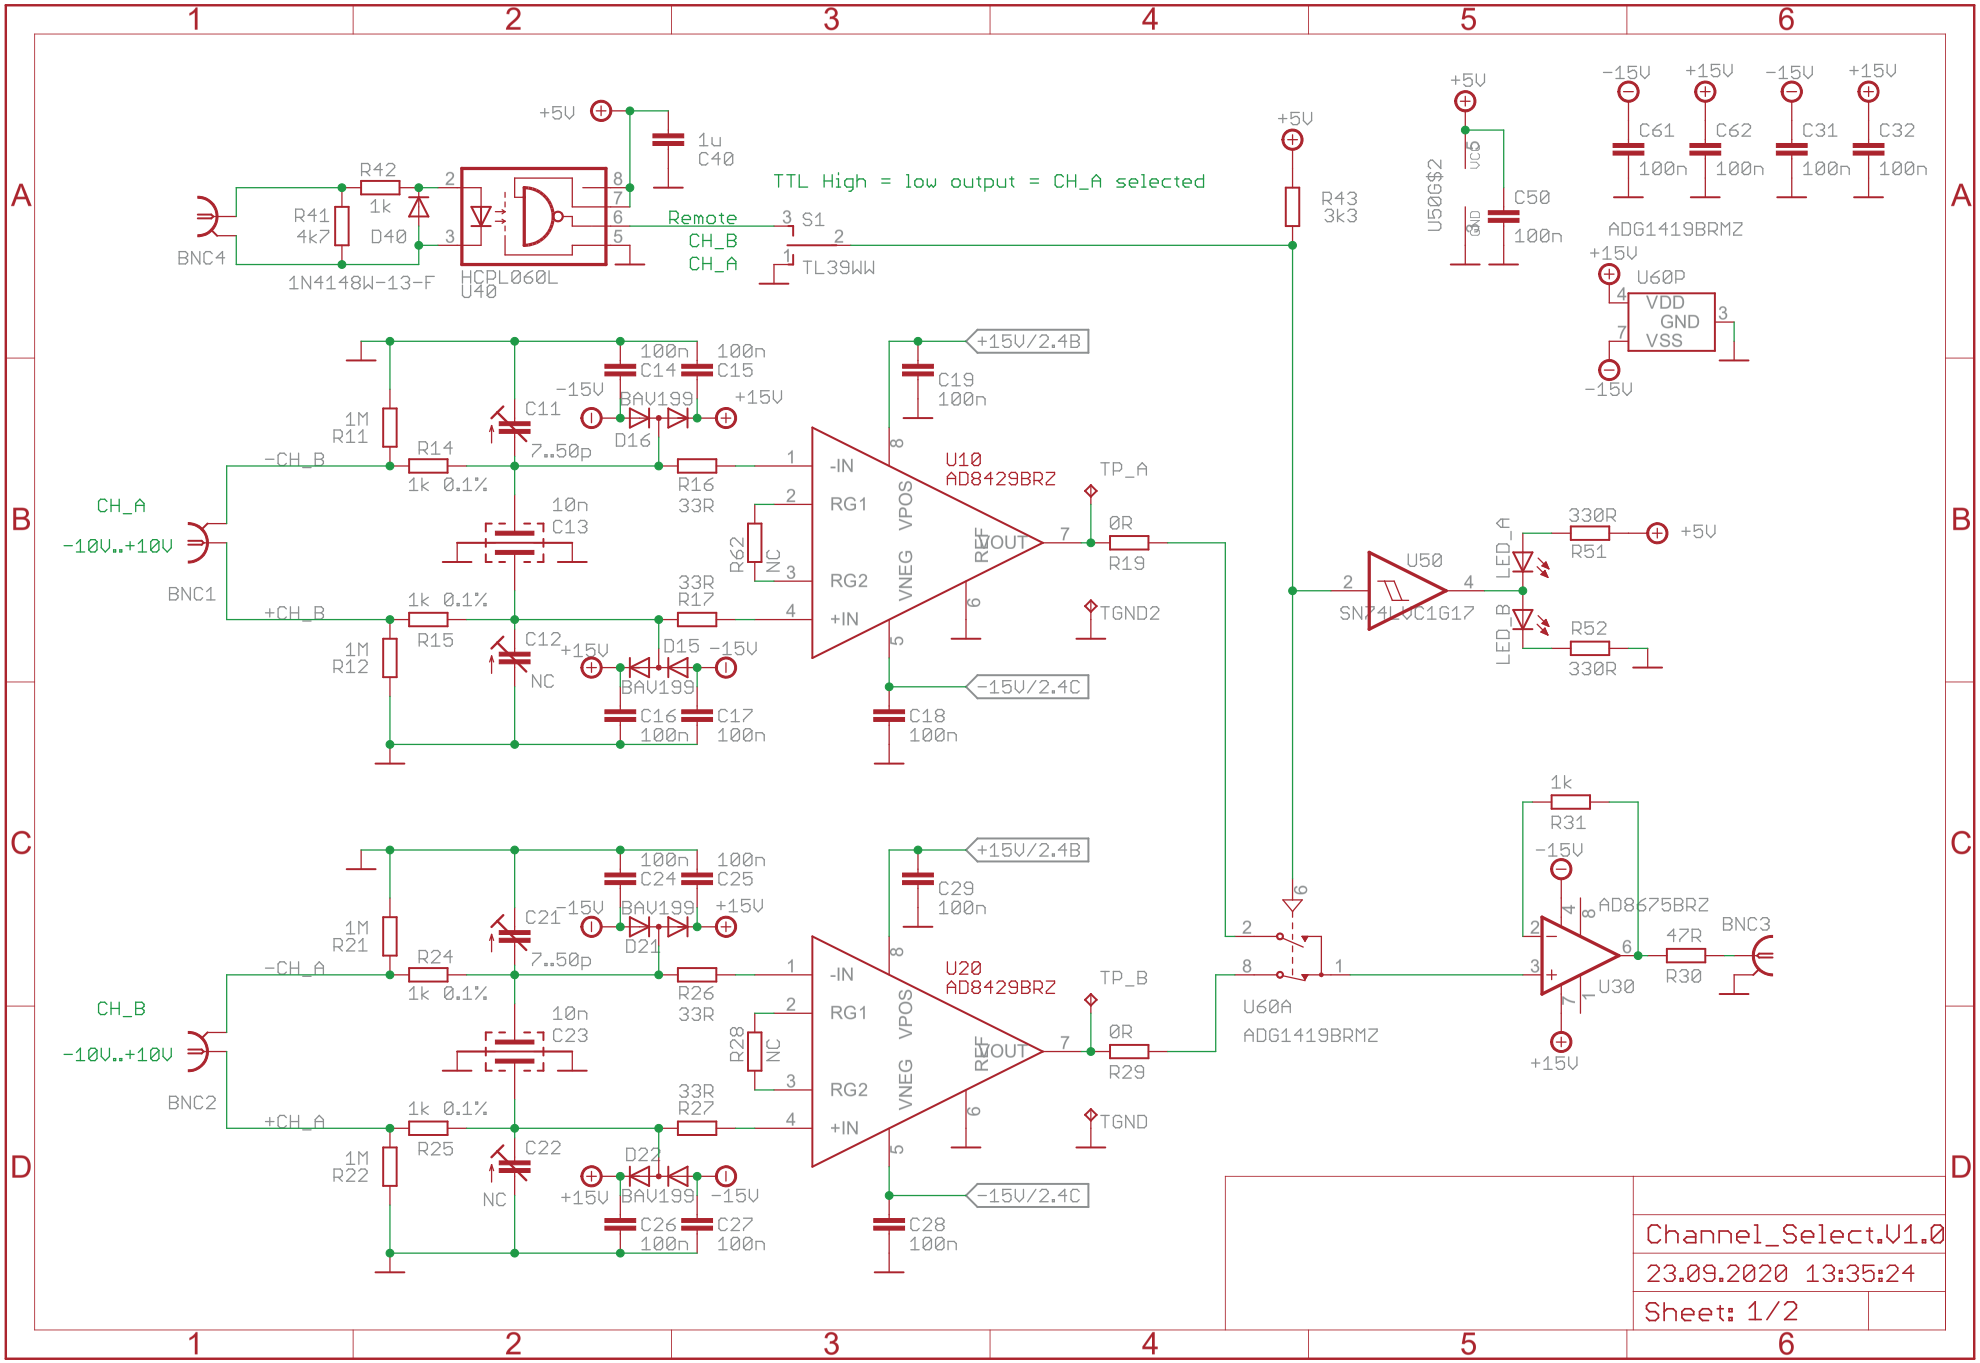
TTL High = low output = CH_A selected
Remote
CH_B
CH_A
S1
TL39WW
HCPL060L
U40
1N4148W-13-F
D40
R41 4k7
R42 1k
BNC4
+5V
1u C40
R43 3k3
U50G$2
C50 100n
ADG1419BRMZ
U60P
VDD GND VSS
C61 100n
C62 100n
C31 100n
C32 100n
-15V
+15V
CH_A
-10V..+10V
BNC1
-CH_B
+CH_B
1M R11
1M R12
R14 1k 0.1%
R15 1k 0.1%
C11 7..50p
C12 NC
10n C13
100n C14
100n C15
C16 100n
C17 100n
BAV199
D16
D15
R16 33R
R17 33R
R62 NC
U10
AD8429BRZ
-IN
RG1
RG2
+IN
VPOS
VNEG
REF
VOUT
C19 100n
C18 100n
+15V/2.4B
-15V/2.4C
TP_A
0R R19
TGND2
U50
SN74LVC1G17
LED_A
LED_B
330R R51
330R R52
CH_B
-10V..+10V
BNC2
-CH_A
+CH_A
1M R21
1M R22
R24 1k 0.1%
R25 1k 0.1%
C21 7..50p
C22 NC
10n C23
100n C24
100n C25
C26 100n
C27 100n
D21
D22
R26 33R
R27 33R
R28 NC
U20
AD8429BRZ
C29 100n
C28 100n
TP_B
0R R29
TGND
U60A
ADG1419BRMZ
1k R31
AD8675BRZ
47R R30
BNC3
U30
Channel_Select.V1.0
23.09.2020 13:35:24
Sheet: 1/2

Power:

AD8429 max 10 mA (High imp. load) -> 10mA x2
AD8675: 35 mA short circuit +3 mA -> 40 mA
LED: 10 mA

Needed:
+15V 60mA
-15V 60mA
+5V 20mA

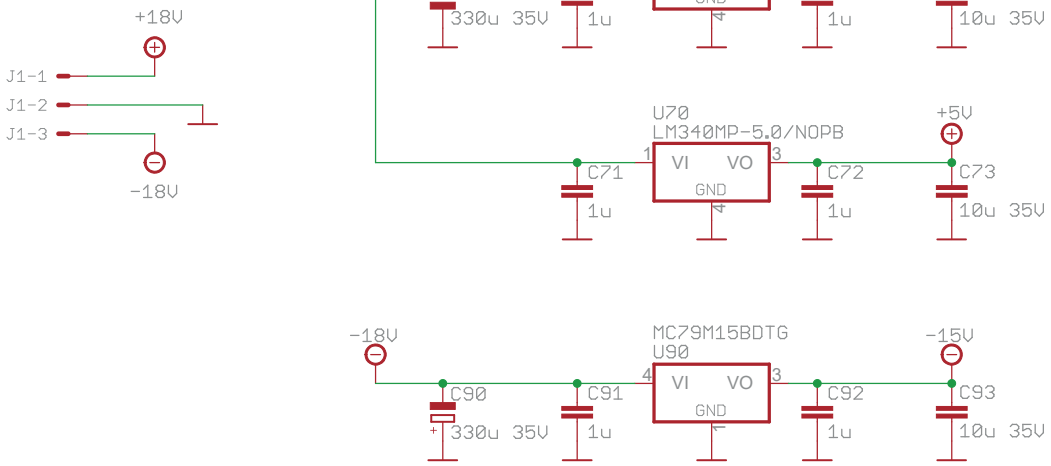

Channel_Select.V1.0	
23.09.2020 13:35:24	
Sheet: 2/2	

Bibliography

[1] D. van Delft and P. Kes, *The discovery of superconductivity*, Physics Today **63**, 38 (2010) (cit. on p. 1).

[2] R. L. Fagaly, *Superconducting quantum interference device instruments and applications*, Review of Scientific Instruments **77** (2006), 101101 (cit. on p. 1).

[3] M. Kjaergaard, M. E. Schwartz, J. Braumüller, P. Krantz, J. I.-J. Wang, S. Gustavsson and W. D. Oliver, *Superconducting Qubits: Current State of Play*, Annual Review of Condensed Matter Physics **11**, 369 (2020) (cit. on p. 1).

[4] A. Bussmann-Holder and H. Keller, *High-temperature superconductors: underlying physics and applications*, Zeitschrift für Naturforschung B **75**, 3 (2020) (cit. on p. 1).

[5] X. Zhou, W.-S. Lee, M. Imada, N. Trivedi, P. Phillips, H.-Y. Kee, P. Törmä and M. Eremets, *High-temperature superconductivity*, Nature Reviews Physics **3**, 462 (2021) (cit. on pp. 1, 2).

[6] S. Lee, J. Kim, H.-T. Kim, S. Im, S. An and K. H. Auh, *Superconductor $Pb_{10-x}Cu_x(PO_4)_6O$ showing levitation at room temperature and atmospheric pressure and mechanism*, 2023, arXiv: 2307.12037 [cond-mat.supr-con] (cit. on p. 1).

[7] S. Lee, J.-H. Kim and Y.-W. Kwon, *The First Room-Temperature Ambient-Pressure Superconductor*, 2023, arXiv: 2307.12008 [cond-mat.supr-con] (cit. on p. 1).

[8] P. Kapitza, *Viscosity of Liquid Helium below the λ-Point*, Nature **141**, 74 (1938) (cit. on p. 1).

[9] J. F. Annett, *Superconductivity, superfluids and condensates*, Oxford: Oxford University Press, 2004, ISBN: 0198507550 (cit. on pp. 1, 10, 12).

[10] J. Bardeen, L. N. Cooper and J. R. Schrieffer, *Theory of Superconductivity*, Phys. Rev. **108**, 1175 (1957) (cit. on p. 1).

[11] Y. J. Uemura et al., *Basic similarities among cuprate, bismuthate, organic, Chevrel-phase, and heavy-fermion superconductors shown by penetration-depth measurements*, Phys. Rev. Lett. **66**, 2665 (1991) (cit. on p. 2).

[12] C. Proust and L. Taillefer, *The Remarkable Underlying Ground States of Cuprate Superconductors*, Annual Review of Condensed Matter Physics **10**, 409 (2019) (cit. on p. 2).

[13] L. Sobirey, H. Biss, N. Luick, M. Bohlen, H. Moritz and T. Lompe, *Observing the Influence of Reduced Dimensionality on Fermionic Superfluids*, Phys. Rev. Lett. **129**, 083601 (2022) (cit. on pp. 2, 134).

[14] A. Legett, *Modern Trends in the Theory of Condensed Matter*, (Springer-Verlag, Berlin, 1980), pp. 13–27 (cit. on pp. 2, 16).

[15] P. Nozières and S. Schmitt-Rink, *Bose condensation in an attractive fermion gas: From weak to strong coupling superconductivity*, Journal of Low Temperature Physics **59**, 195 (1985) (cit. on p. 2).

[16] Q. Chen, K. Levin and J. Stajic, *Applying BCS–BEC crossover theory to high-temperature superconductors and ultracold atomic Fermi gases (Review Article)*, Low Temperature Physics **32**, 406 (2006) (cit. on p. 2).

[17] G. C. Strinati, P. Pieri, G. Röpke, P. Schuck and M. Urban, *The BCS–BEC crossover: From ultra-cold Fermi gases to nuclear systems*, Physics Reports **738**, 1 (2018) (cit. on pp. 2, 10).

[18] N. Chamel and P. Haensel, *Physics of Neutron Star Crusts*, Living Reviews in Relativity **11**, 10 (2008) (cit. on p. 2).

[19] J. E. Thomas, *Superfluidity in the picture*, Nature **442**, 32 (2006) (cit. on p. 2).

[20] W. Zwerger, ed., *The BCS-BEC Crossover and the Unitary Fermi Gas*, Springer Berlin, Heidelberg, 2011, ISBN: 978-3-642-21977-1 (cit. on p. 2).

[21] E. Y. Loh, J. E. Gubernatis, R. T. Scalettar, S. R. White, D. J. Scalapino and R. L. Sugar, *Sign problem in the numerical simulation of many-electron systems*, Phys. Rev. B **41**, 9301 (1990) (cit. on p. 2).

[22] B. Keimer, S. A. Kivelson, M. R. Norman, S. Uchida and J. Zaanen, *From quantum matter to high-temperature superconductivity in copper oxides*, Nature **518**, 179 (2015) (cit. on p. 2).

[23] Y. Nakagawa, Y. Kasahara, T. Nomoto, R. Arita, T. Nojima and Y. Iwasa, *Gate-controlled BCS-BEC crossover in a two-dimensional superconductor*, Science **372**, 190 (2021) (cit. on pp. 2, 3).

[24] Y. Suzuki, K. Wakamatsu, J. Ibuka, H. Oike, T. Fujii, K. Miyagawa, H. Taniguchi and K. Kanoda, *Mott-Driven BEC-BCS Crossover in a Doped Spin Liquid Candidate* $\kappa-(\mathrm{BEDT-TTF})_4\mathrm{Hg}_{2.89}\mathrm{Br}_8$, Phys. Rev. X **12**, 011016 (2022) (cit. on p. 2).

[25] R. P. Feynman, *Simulating physics with computers*, International Journal of Theoretical Physics **21**, 467 (1982) (cit. on p. 2).

[26] K. Bharti et al., *Noisy intermediate-scale quantum algorithms*, Rev. Mod. Phys. **94**, 015004 (2022) (cit. on p. 2).

[27] T. H. Johnson, S. R. Clark and D. Jaksch, *What is a quantum simulator?*, EPJ Quantum Technology **1**, 10 (2014) (cit. on p. 3).

[28] I. Bloch, J. Dalibard and W. Zwerger, *Many-body physics with ultracold gases*, Rev. Mod. Phys. **80**, 885 (2008) (cit. on p. 3).

[29] I. Bloch, J. Dalibard and S. Nascimbène, *Quantum simulations with ultracold quantum gases*, Nature Physics **8**, 267 (2012) (cit. on p. 3).

[30] W. D. Phillips, *Nobel Lecture: Laser cooling and trapping of neutral atoms*, Rev. Mod. Phys. **70**, 721 (1998) (cit. on p. 3).

[31] W. Ketterle and M. W. Zwierlein, *Making, probing and understanding ultracold Fermi gases*, La Rivista del Nuovo Cimento **31**, 247 (2008) (cit. on pp. 3, 9, 14–16, 19, 23, 46, 63, 88, 89).

[32] F. Schäfer, T. Fukuhara, S. Sugawa, Y. Takasu and Y. Takahashi, *Tools for quantum simulation with ultracold atoms in optical lattices*, Nature Reviews Physics **2**, 411 (2020) (cit. on p. 3).

[33] 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) (cit. on p. 3).

[34] 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) (cit. on p. 3).

[35] C. Raman, M. Köhl, R. Onofrio, D. S. Durfee, C. E. Kuklewicz, Z. Hadzibabic and W. Ketterle, *Evidence for a Critical Velocity in a Bose-Einstein Condensed Gas*, Phys. Rev. Lett. **83**, 2502 (1999) (cit. on pp. 3, 41).

[36] K. W. Madison, F. Chevy, W. Wohlleben and J. Dalibard, *Vortex Formation in a Stirred Bose-Einstein Condensate*, Phys. Rev. Lett. **84**, 806 (2000) (cit. on p. 3).

[37] F. Gerbier, J. H. Thywissen, S. Richard, M. Hugbart, P. Bouyer and A. Aspect, *Critical Temperature of a Trapped, Weakly Interacting Bose Gas*, Phys. Rev. Lett. **92**, 030405 (2004) (cit. on p. 3).

[38] 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) (cit. on p. 3).

[39] C. A. Regal, M. Greiner and D. S. Jin, *Observation of Resonance Condensation of Fermionic Atom Pairs*, Phys. Rev. Lett. **92**, 040403 (2004) (cit. on pp. 3, 41, 42, 59, 70).

[40] M. W. Zwierlein, C. A. Stan, C. H. Schunck, S. M. F. Raupach, A. J. Kerman and W. Ketterle, *Condensation of Pairs of Fermionic Atoms near a Feshbach Resonance*, Phys. Rev. Lett. **92**, 120403 (2004) (cit. on pp. 3, 23, 41, 42, 59, 70).

[41] M. W. Zwierlein, J. R. Abo-Shaeer, A. Schirotzek, C. H. Schunck and W. Ketterle, *Vortices and superfluidity in a strongly interacting Fermi gas*, Nature **435**, 1047 (2005) (cit. on pp. 3, 46).

[42] M. J. H. Ku, A. T. Sommer, L. W. Cheuk and M. W. Zwierlein, *Revealing the Superfluid Lambda Transition in the Universal Thermodynamics of a Unitary Fermi Gas*, Science **335**, 563 (2012) (cit. on pp. 3, 17, 23, 59–63).

[43] W. Weimer, K. Morgener, V. P. Singh, J. Siegl, K. Hueck, N. Luick, L. Mathey and H. Moritz, *Critical Velocity in the BEC-BCS Crossover*, Phys. Rev. Lett. **114**, 095301 (2015) (cit. on pp. 3, 41).

[44] S. Hoinka, P. Dyke, M. G. Lingham, J. J. Kinnunen, G. M. Bruun and C. J. Vale, *Goldstone mode and pair-breaking excitations in atomic Fermi superfluids*, Nature Physics **13**, 943 (2017) (cit. on pp. 3, 23).

[45] T.-L. Ho, *Universal Thermodynamics of Degenerate Quantum Gases in the Unitarity Limit*, Phys. Rev. Lett. **92**, 090402 (2004) (cit. on pp. 3, 17, 59, 87).

[46] J. Wang, *ChatGPT: A test drive*, American Journal of Physics **91**, 255 (2023) (cit. on p. 3).

[47] Y. LeCun, Y. Bengio and G. Hinton, *Deep learning*, Nature **521**, 436 (2015) (cit. on pp. 3, 46, 48).

[48] P. B. Wigley et al., *Fast machine-learning online optimization of ultra-cold-atom experiments*, Scientific Reports **6**, 25890 (2016) (cit. on p. 3).

[49] 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*, Nature Physics **15**, 917 (2019) (cit. on pp. 3, 46).

[50] W. T. Lou, H. Sutterud, G. Cassella, W. M. C. Foulkes, J. Knolle, D. Pfau and J. S. Spencer, *Neural Wave Functions for Superfluids*, 2023, arXiv: 2305.06989 [cond-mat.quant-gas] (cit. on p. 3).

[51] J. Kim, G. Pescia, B. Fore, J. Nys, G. Carleo, S. Gandolfi, M. Hjorth-Jensen and A. Lovato, *Neural-network quantum states for ultra-cold Fermi gases*, 2023, arXiv: 2305.08831 [cond-mat.quant-gas] (cit. on p. 3).

[52] A. J. Beekman, L. Rademaker and J. van Wezel, *An introduction to spontaneous symmetry breaking*, SciPost Phys. Lect. Notes, 11 (2019) (cit. on pp. 4, 12).

[53] D. Pekker and C. Varma, *Amplitude/Higgs Modes in Condensed Matter Physics*, Annual Review of Condensed Matter Physics **6**, 269 (2015) (cit. on pp. 4, 13, 19).

[54] W. J. Kwon, G. D. Pace, R. Panza, M. Inguscio, W. Zwerger, M. Zaccanti, F. Scazza and G. Roati, *Strongly correlated superfluid order parameters from dc Josephson supercurrents*, Science **369**, 84 (2020) (cit. on p. 4).

[55] N. Luick, L. Sobirey, M. Bohlen, V. P. Singh, L. Mathey, T. Lompe and H. Moritz, *An ideal Josephson junction in an ultracold two-dimensional Fermi gas*, Science **369**, 89 (2020) (cit. on p. 4).

[56] R. Shimano and N. Tsuji, *Higgs Mode in Superconductors*, Annual Review of Condensed Matter Physics **11**, 103 (2020) (cit. on pp. 4, 12, 13).

[57] A. Behrle, T. Harrison, J. Kombe, K. Gao, M. Link, J.-S. Bernier, C. Kollath and M. Köhl, *Higgs mode in a strongly interacting fermionic superfluid*, Nature Physics **14**, 781 (2018) (cit. on pp. 4, 75, 121, 123, 134).

[58] R. G. Scott, F. Dalfovo, L. P. Pitaevskii and S. Stringari,
Rapid ramps across the BEC-BCS crossover: A route to measuring the superfluid gap,
Phys. Rev. A **86**, 053604 (2012) (cit. on pp. 4, 75, 82, 85, 98).

[59] C. M. Varma, *Higgs Boson in Superconductors*,
Journal of Low Temperature Physics **126**, 901 (2002) (cit. on pp. 4, 13, 98, 121).

[60] X. Han, B. Liu and J. Hu,
Observability of Higgs mode in a system without Lorentz invariance,
Phys. Rev. A **94**, 033608 (2016) (cit. on pp. 5, 18, 75, 95, 98, 121).

[61] A. Jabs, *Connecting Spin and Statistics in Quantum Mechanics*,
Foundations of Physics **40**, 776 (2010) (cit. on pp. 7, 9).

[62] B. Cowan, *On the Chemical Potential of Ideal Fermi and Bose Gases*,
Journal of Low Temperature Physics **197**, 412 (2019) (cit. on p. 7).

[63] O. Penrose and L. Onsager, *Bose-Einstein Condensation and Liquid Helium*,
Phys. Rev. **104**, 576 (1956) (cit. on p. 12).

[64] A. F. Volkov and S. M. Kogan, *Collisionless relaxation of the energy gap in superconductors*,
Zh. Eksp. Teor. Fiz. **65**, 2038 (1973) (cit. on pp. 13, 75, 78, 90, 96).

[65] Y. Nambu, *Quasi-Particles and Gauge Invariance in the Theory of Superconductivity*,
Phys. Rev. **117**, 648 (1960) (cit. on p. 13).

[66] L. P. Gorkov and T. M. Melik-Barkhudarov,
Contribution to the theory of superfluidity in an imperfect Fermi gas,
Sov. Phys. JETP **13**, 1018 (1961) (cit. on p. 14).

[67] R. Haussmann, W. Rantner, S. Cerrito and W. Zwerger,
Thermodynamics of the BCS-BEC crossover, Phys. Rev. A **75**, 023610 (2007)
(cit. on pp. 15–17, 70, 82, 83, 88, 89, 93, 95, 99, 104, 110, 118, 129).

[68] D. S. Petrov, C. Salomon and G. V. Shlyapnikov, *Weakly Bound Dimers of Fermionic Atoms*,
Phys. Rev. Lett. **93**, 090404 (2004) (cit. on p. 15).

[69] P. Arnold and G. Moore,
BEC Transition Temperature of a Dilute Homogeneous Imperfect Bose Gas,
Phys. Rev. Lett. **87**, 120401 (2001) (cit. on p. 15).

[70] V. A. Kashurnikov, N. V. Prokof'ev and B. V. Svistunov,
Critical Temperature Shift in Weakly Interacting Bose Gas, Phys. Rev. Lett. **87**, 120402 (2001)
(cit. on p. 15).

[71] L. Pisani, A. Perali, P. Pieri and G. C. Strinati,
Entanglement between pairing and screening in the Gorkov-Melik-Barkhudarov correction to the critical temperature throughout the BCS-BEC crossover, Phys. Rev. B **97**, 014528 (2018)
(cit. on pp. 17, 70, 86, 87, 135).

[72] A. Bulgac, J. E. Drut and P. Magierski,
Quantum Monte Carlo simulations of the BCS-BEC crossover at finite temperature,
Phys. Rev. A **78**, 023625 (2008) (cit. on pp. 17, 70).

[73] E. Burovski, E. Kozik, N. Prokof'ev, B. Svistunov and M. Troyer, *Critical Temperature Curve in BEC-BCS Crossover*, Phys. Rev. Lett. **101**, 090402 (2008) (cit. on pp. 17, 70).

[74] S. Pilati, S. Giorgini and N. Prokof'ev, *Critical Temperature of Interacting Bose Gases in Two and Three Dimensions*, Phys. Rev. Lett. **100**, 140405 (2008) (cit. on p. 17).

[75] S. Tsuchiya, D. Yamamoto, R. Yoshii and M. Nitta, *Hidden charge-conjugation, parity, and time-reversal symmetries and massive Goldstone (Higgs) modes in superconductors*, Phys. Rev. B **98**, 094503 (2018) (cit. on pp. 17, 18, 75, 104, 121).

[76] L. Pitaevskii and S. Stringari, "The Ideal Bose Gas in the Harmonic Trap", *Bose-Einstein Condensation and Superfluidity*, Oxford University Press, 2016, ISBN: 9780198758884 (cit. on p. 20).

[77] S. Giorgini, L. P. Pitaevskii and S. Stringari, *Condensate fraction and critical temperature of a trapped interacting Bose gas*, Phys. Rev. A **54**, R4633 (1996) (cit. on p. 20).

[78] S. Giorgini, L. P. Pitaevskii and S. Stringari, *Theory of ultracold atomic Fermi gases*, Rev. Mod. Phys. **80**, 1215 (2008) (cit. on pp. 20, 113).

[79] Alexandra Bianca Behrle, *Driving a Strongly Interacting Superfluid out of Equilibrium*, PhD thesis: Rheinische Friedrich-Wilhelms-Universität Bonn, 2018 (cit. on pp. 21, 26, 27, 34, 37, 138).

[80] Timothy Joseph Harrison, *Measuring the Gap and Investigating Non-equilibrium in the BEC-BCS Crossover*, PhD thesis: Rheinische Friedrich-Wilhelms-Universität Bonn, 2017 (cit. on pp. 21, 27, 42–44, 76, 86, 138, 139).

[81] C. Chin, R. Grimm, P. Julienne and E. Tiesinga, *Feshbach resonances in ultracold gases*, Rev. Mod. Phys. **82**, 1225 (2010) (cit. on pp. 21, 22).

[82] M. G. Ries, A. N. Wenz, G. Zürn, L. Bayha, I. Boettcher, D. Kedar, P. A. Murthy, M. Neidig, T. Lompe and S. Jochim, *Observation of Pair Condensation in the Quasi-2D BEC-BCS Crossover*, Phys. Rev. Lett. **114**, 230401 (2015) (cit. on pp. 23, 134).

[83] K. Hueck, N. Luick, L. Sobirey, J. Siegl, T. Lompe and H. Moritz, *Two-Dimensional Homogeneous Fermi Gases*, Phys. Rev. Lett. **120**, 060402 (2018) (cit. on pp. 23, 134).

[84] M. E. Gehm, *Preparation of an optically-trapped degenerate Fermi gas of ^{6}Li: Finding the route to degeneracy*, PhD thesis: Duke University, North Carolina, 2003 (cit. on pp. 23, 26, 30, 36).

[85] G. Zürn, T. Lompe, A. N. Wenz, S. Jochim, P. S. Julienne and J. M. Hutson, *Precise Characterization of ^{6}Li Feshbach Resonances Using Trap-Sideband-Resolved RF Spectroscopy of Weakly Bound Molecules*, Phys. Rev. Lett. **110**, 135301 (2013) (cit. on p. 24).

[86] J. Mitroy and M. W. J. Bromley, *Semiempirical calculation of van der Waals coefficients for alkali-metal and alkaline-earth-metal atoms*, Phys. Rev. A **68**, 052714 (2003) (cit. on p. 24).

[87] D. Eberz,
Design and Construction of a Multi-Species Effusive Atomic Oven for Lithium and Sodium,
Master thesis: Rheinische Friedrich-Wilhelms-Universität Bonn, 2020 (cit. on pp. 25, 27).

[88] D. A. Steck, *Sodium D Line Data*,
available online at `http://steck.us/alkalidata` (revision 2.2.1, 21 November 2019)
(cit. on pp. 26, 28).

[89] W. Ketterle, K. B. Davis, M. A. Joffe, A. Martin and D. E. Pritchard,
High densities of cold atoms in a dark spontaneous-force optical trap,
Phys. Rev. Lett. **70**, 2253 (1993) (cit. on p. 26).

[90] C. J. Foot, *Atomic Physics*, Oxford: Oxford University Press, 2005, ISBN: 978-0-19-850696-6
(cit. on pp. 27, 29).

[91] J. Dalibard and C. Cohen-Tannoudji,
Laser cooling below the Doppler limit by polarization gradients: simple theoretical models,
J. Opt. Soc. Am. B **6**, 2023 (1989) (cit. on p. 27).

[92] A. Kell, M. Link, M. Breyer, A. Hoffmann, M. Köhl and K. Gao, *A compact and fast magnetic coil for the manipulation of quantum gases with Feshbach resonances*,
Review of Scientific Instruments **92**, 093202 (2021) (cit. on pp. 30–32, 76).

[93] Andreas Kell, *Measuring Excitations of a Strongly Interacting Superfluid Fermi Gas*,
PhD thesis: Rheinische Friedrich-Wilhelms-Universität Bonn, 2022
(cit. on pp. 31, 32, 51, 63, 67–69, 111, 124, 134).

[94] P. A. Murthy and S. Jochim,
Direct imaging of the order parameter of an atomic superfluid using matterwave optics, 2019,
arXiv: `1911.10824 [cond-mat.quant-gas]` (cit. on pp. 33, 44, 51).

[95] L. Chomaz, *Coherence and superfluidity of Bose gases in reduced dimensions: from harmonic traps to uniform fluids*, Theses: Ecole Normale Supérieure, 2014 (cit. on pp. 34, 37).

[96] G. Reinaudi, T. Lahaye, Z. Wang and D. Guéry-Odelin,
Strong saturation absorption imaging of dense clouds of ultracold atoms,
Opt. Lett. **32**, 3143 (2007) (cit. on p. 34).

[97] M. Horikoshi, A. Ito, T. Ikemachi, Y. Aratake, M. Kuwata-Gonokami and M. Koashi,
Appropriate Probe Condition for Absorption Imaging of Ultracold ^{6}Li Atoms, Journal of the Physical Society of Japan **86**, 104301 (2017), arXiv:1608.07152 [cond-mat.quant-gas]
(cit. on p. 35).

[98] L. Landau, *Theory of the Superfluidity of Helium II*, Phys. Rev. **60**, 356 (1941) (cit. on p. 41).

[99] M. Holten, L. Bayha, K. Subramanian, S. Brandstetter, C. Heintze, P. Lunt, P. M. Preiss and S. Jochim, *Observation of Cooper pairs in a mesoscopic two-dimensional Fermi gas*,
Nature **606**, 287 (2022) (cit. on p. 42).

[100] M. W. Zwierlein, C. H. Schunck, C. A. Stan, S. M. F. Raupach and W. Ketterle,
Formation Dynamics of a Fermion Pair Condensate, Phys. Rev. Lett. **94**, 180401 (2005)
(cit. on p. 46).

[101] E. Altman and A. Vishwanath,
Dynamic Projection on Feshbach Molecules: A Probe of Pairing and Phase Fluctuations,
Phys. Rev. Lett. **95**, 110404 (2005) (cit. on p. 46).

[102] S. Matyjaśkiewicz, M. H. Szymańska and K. Góral,
Probing Fermionic Condensates by Fast-Sweep Projection onto Feshbach Molecules,
Phys. Rev. Lett. **101**, 150410 (2008) (cit. on p. 46).

[103] Martin Link,
Exploring Non-Equilibrium in Ultracold Fermi Gases and Machine Learning in Physics,
PhD thesis: Rheinische Friedrich-Wilhelms-Universität Bonn, 2021
(cit. on pp. 47, 49, 51, 53, 134).

[104] M. A. Nielsen, *Neural Networks and Deep Learning*, Determination Press, 2015
(cit. on p. 47).

[105] A. Lenail, *NN-SVG*, `alexlenail.me/NN-SVG/index.html` (cit. on p. 48).

[106] K. Hornik, M. Stinchcombe and H. White,
Multilayer feedforward networks are universal approximators, Neural Networks **2**, 359 (1989)
(cit. on p. 48).

[107] M. Link, K. Gao, A. Kell, M. Breyer, D. Eberz, B. Rauf and M. Köhl,
Machine Learning the Phase Diagram of a Strongly Interacting Fermi Gas,
Phys. Rev. Lett. **130**, 203401 (2023) (cit. on pp. 48, 51, 54).

[108] Martín Abadi et al., *TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems*,
Software available from tensorflow.org, 2015 (cit. on p. 50).

[109] D. P. Kingma and J. Ba, *Adam: A Method for Stochastic Optimization*, 2017,
arXiv: `1412.6980 [cs.LG]` (cit. on p. 50).

[110] J. Kinast, A. Turlapov, J. E. Thomas, Q. Chen, J. Stajic and K. Levin,
Heat Capacity of a Strongly Interacting Fermi Gas, Science **307**, 1296 (2005)
(cit. on pp. 51, 59).

[111] L. Pratt and B. Jennings, *A Survey of Transfer Between Connectionist Networks*,
Connection Science **8**, 163 (1996) (cit. on p. 53).

[112] A. Shrikumar, P. Greenside and A. Kundaje,
Learning Important Features Through Propagating Activation Differences, Proceedings of the
34th International Conference on Machine Learning, PMLR 70, 3145-3153, 2017
(cit. on p. 54).

[113] D. Eberz, M. Link, A. Kell, M. Breyer, K. Gao and M. Köhl, *Detecting the phase transition in a strongly-interacting Fermi gas by unsupervised machine learning*,
submitted to Physical Review A, 2023 (cit. on pp. 56, 69).

[114] Y. Inada, M. Horikoshi, S. Nakajima, M. Kuwata-Gonokami, M. Ueda and T. Mukaiyama,
Critical Temperature and Condensate Fraction of a Fermion Pair Condensate,
Phys. Rev. Lett. **101**, 180406 (2008) (cit. on pp. 59, 70).

[115] A. Perali, P. Pieri, L. Pisani and G. C. Strinati,
BCS-BEC Crossover at Finite Temperature for Superfluid Trapped Fermi Atoms,
Phys. Rev. Lett. **92**, 220404 (2004) (cit. on pp. 59, 60, 70).

[116] Q. Chen, C. A. Regal, M. Greiner, D. S. Jin and K. Levin,
Understanding the superfluid phase diagram in trapped Fermi gases,
Phys. Rev. A **73**, 041601 (2006) (cit. on pp. 59, 60).

[117] L. Luo, B. Clancy, J. Joseph, J. Kinast and J. E. Thomas,
Measurement of the Entropy and Critical Temperature of a Strongly Interacting Fermi Gas,
Phys. Rev. Lett. **98**, 080402 (2007) (cit. on p. 59).

[118] S. Nascimbène, N. Navon, K. J. Jiang, F. Chevy and C. Salomon,
Exploring the thermodynamics of a universal Fermi gas, Nature **463**, 1057 (2010)
(cit. on pp. 59, 60, 63).

[119] M. Horikoshi, S. Nakajima, M. Ueda and T. Mukaiyama,
Measurement of Universal Thermodynamic Functions for a Unitary Fermi Gas,
Science **327**, 442 (2010) (cit. on pp. 59, 60, 63, 89).

[120] A. Rohatgi, *Webplotdigitizer: Version 4.6*, https://automeris.io/WebPlotDigitizer, 2022
(cit. on p. 61).

[121] X.-J. Liu, H. Hu and P. D. Drummond, *Virial Expansion for a Strongly Correlated Fermi Gas*,
Phys. Rev. Lett. **102**, 160401 (2009) (cit. on pp. 61, 68).

[122] D. Lee and T. Schäfer, *Cold dilute neutron matter on the lattice. I. Lattice virial coefficients and large scattering lengths*, Phys. Rev. C **73**, 015201 (2006) (cit. on pp. 61, 68).

[123] Y. Hou and J. E. Drut,
Fourth- and Fifth-Order Virial Coefficients from Weak Coupling to Unitarity,
Phys. Rev. Lett. **125**, 050403 (2020) (cit. on pp. 61, 68, 72).

[124] L. Montgomery Smith, D. R. Keefer and S. Sudharsanan,
Abel inversion using transform techniques,
Journal of Quantitative Spectroscopy and Radiative Transfer **39**, 367 (1988) (cit. on p. 63).

[125] X.-J. Liu, *Virial expansion for a strongly correlated Fermi system and its application to ultracold atomic Fermi gases*, Physics Reports **524**, 37 (2013) (cit. on p. 67).

[126] X. Leyronas, *Virial expansion with Feynman diagrams*, Phys. Rev. A **84**, 053633 (2011)
(cit. on p. 68).

[127] Y. Wang et al., *Axial Higgs mode detected by quantum pathway interference in RTe3*,
Nature **606**, 896 (2022) (cit. on p. 75).

[128] L. Bayha, M. Holten, R. Klemt, K. Subramanian, J. Bjerlin, S. M. Reimann, G. M. Bruun,
P. M. Preiss and S. Jochim,
Observing the emergence of a quantum phase transition shell by shell, Nature **587**, 583 (2020)
(cit. on p. 75).

[129] M.-A. Méasson, Y. Gallais, M. Cazayous, B. Clair, P. Rodière, L. Cario and A. Sacuto,
Amplitude Higgs mode in the $2H - NbSe_2$ *superconductor*, Phys. Rev. B **89**, 060503 (2014)
(cit. on p. 75).

[130] R. Matsunaga, Y. I. Hamada, K. Makise, Y. Uzawa, H. Terai, Z. Wang and R. Shimano,
Higgs Amplitude Mode in the BCS Superconductors $Nb_{1-x}Ti_xN$ *Induced by Terahertz Pulse Excitation*, Phys. Rev. Lett. **111**, 057002 (2013) (cit. on p. 75).

[131] K. Katsumi, N. Tsuji, Y. I. Hamada, R. Matsunaga, J. Schneeloch, R. D. Zhong, G. D. Gu, H. Aoki, Y. Gallais and R. Shimano, *Higgs Mode in the d-Wave Superconductor* $Bi_2Sr_2CaCu_2O_{8+x}$ *Driven by an Intense Terahertz Pulse*, Phys. Rev. Lett. **120**, 117001 (2018) (cit. on p. 75).

[132] C. Rüegg, B. Normand, M. Matsumoto, A. Furrer, D. F. McMorrow, K. W. Krämer, H.-U. Güdel, S. N. Gvasaliya, H. Mutka and M. Boehm, *Quantum Magnets under Pressure: Controlling Elementary Excitations in* $TlCuCl_3$, Phys. Rev. Lett. **100**, 205701 (2008) (cit. on p. 75).

[133] U. Bissbort, S. Götze, Y. Li, J. Heinze, J. S. Krauser, M. Weinberg, C. Becker, K. Sengstock and W. Hofstetter, *Detecting the Amplitude Mode of Strongly Interacting Lattice Bosons by Bragg Scattering*, Phys. Rev. Lett. **106**, 205303 (2011) (cit. on p. 75).

[134] M. Endres, T. Fukuhara, D. Pekker, M. Cheneau, P. Schauβ, C. Gross, E. Demler, S. Kuhr and I. Bloch, *The 'Higgs' amplitude mode at the two-dimensional superfluid/Mott insulator transition*, Nature **487**, 454 (2012) (cit. on p. 75).

[135] CMS Collab., *Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC*, Physics Letters B **716**, 30 (2012) (cit. on p. 75).

[136] ATLAS Collab., *Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC*, Physics Letters B **716**, 1 (2012) (cit. on p. 75).

[137] C.-L. Hung, X. Zhang, N. Gemelke and C. Chin, *Accelerating evaporative cooling of atoms into Bose-Einstein condensation in optical traps*, Phys. Rev. A **78**, 011604 (2008) (cit. on p. 76).

[138] E. A. Yuzbashyan, M. Dzero, V. Gurarie and M. S. Foster, *Quantum quench phase diagrams of an s-wave BCS-BEC condensate*, Phys. Rev. A **91**, 033628 (2015) (cit. on pp. 78, 79, 98, 99).

[139] A. Barresi, A. Boulet, G. Wlazłowski and P. Magierski, *Generation and decay of Higgs mode in a strongly interacting Fermi gas*, 2023, arXiv: 2303.13394 [cond-mat.quant-gas] (cit. on pp. 78, 98, 106).

[140] V. Gurarie, *Nonequilibrium Dynamics of Weakly and Strongly Paired Superconductors*, Phys. Rev. Lett. **103**, 075301 (2009) (cit. on pp. 79, 91).

[141] J. Tokimoto, S. Tsuchiya and T. Nikuni, *Excitation of Higgs Mode in Superfluid Fermi Gas in BCS–BEC Crossover*, Journal of the Physical Society of Japan **88**, 023601 (2019), arXiv:1810.09110 [cond-mat.quant-gas] (cit. on pp. 79, 85, 91).

[142] J. Tokimoto, S. Tsuchiya and T. Nikuni, *Higgs Mode in a Trapped Superfluid Fermi Gas*, Journal of Low Temperature Physics **187**, 765 (2017) (cit. on pp. 79, 82).

[143] S. Hannibal, P. Kettmann, M. D. Croitoru, V. M. Axt and T. Kuhn,
Persistent oscillations of the order parameter and interaction quench phase diagram for a confined Bardeen-Cooper-Schrieffer Fermi gas, Phys. Rev. A **98**, 053605 (2018)
(cit. on pp. 82, 98–100).

[144] S. Musolino,
Nonequilibrium dynamics of strongly correlated quantum gases: From few to many,
PhD thesis: Eindhoven University of Technology, 2021, ISBN: 978-90-386-5336-5
(cit. on pp. 82, 84–87).

[145] S. Floerchinger, M. M. Scherer and C. Wetterich,
Modified Fermi sphere, pairing gap, and critical temperature for the BCS-BEC crossover,
Phys. Rev. A **81**, 063619 (2010) (cit. on pp. 82, 83).

[146] L. Salasnich, N. Manini and A. Parola,
Condensate fraction of a Fermi gas in the BCS-BEC crossover,
Phys. Rev. A **72**, 023621 (2005) (cit. on p. 90).

[147] H. Kurkjian, S. N. Klimin, J. Tempere and Y. Castin,
Pair-Breaking Collective Branch in BCS Superconductors and Superfluid Fermi Gases,
Phys. Rev. Lett. **122**, 093403 (2019) (cit. on p. 98).

[148] E. A. Yuzbashyan and M. Dzero,
Dynamical Vanishing of the Order Parameter in a Fermionic Condensate,
Phys. Rev. Lett. **96**, 230404 (2006) (cit. on p. 98).

[149] S. Tsuchiya, *Emergence of multiple Higgs modes due to spontaneous breakdown of a $\mathbb{Z}_2$ symmetry in a superconductor*, 2021, arXiv: 2104.09744 [cond-mat.supr-con]
(cit. on p. 106).

[150] P. Dyke, A. Hogan, I. Herrera, C. C. N. Kuhn, S. Hoinka and C. J. Vale,
Dynamics of a Fermi Gas Quenched to Unitarity, Phys. Rev. Lett. **127**, 100405 (2021)
(cit. on pp. 112, 121, 134).

[151] D. Ahmed-Braun, *From few- to many-body physics in strongly interacting quantum gases: A lot about the little and a little about the lot*,
PhD thesis: Eindhoven University of Technology, 2023, ISBN: 978-90-386-5740-0
(cit. on pp. 112, 121, 134).

[152] M. Bartenstein, A. Altmeyer, S. Riedl, S. Jochim, C. Chin, J. H. Denschlag and R. Grimm,
Collective Excitations of a Degenerate Gas at the BEC-BCS Crossover,
Phys. Rev. Lett. **92**, 203201 (2004) (cit. on p. 113).

[153] J. Kinast, A. Turlapov and J. E. Thomas,
Breakdown of hydrodynamics in the radial breathing mode of a strongly interacting Fermi gas,
Phys. Rev. A **70**, 051401 (2004) (cit. on p. 113).

[154] Y. Castin,
Exact scaling transform for a unitary quantum gas in a time dependent harmonic potential,
Comptes Rendus Physique **5**, 407 (2004) (cit. on p. 113).

[155] G. De Rosi and S. Stringari,
Collective oscillations of a trapped quantum gas in low dimensions,
Phys. Rev. A **92**, 053617 (2015) (cit. on p. 113).

[156] B. F. Chao and C. H. Chung, *On Estimating the Cross Correlation and Least Squares Fit of One Data Set to Another With Time Shift*, Earth and Space Science **6**, 1409 (2019)
(cit. on p. 115).

[157] Johannes Amani Kombe, *Nonequilibrium Dynamics of Correlated Fermi Gases*,
PhD thesis: Rheinische Friedrich-Wilhelms-Universität Bonn, 2020 (cit. on p. 120).

[158] A. Bulgac and S. Yoon,
Large Amplitude Dynamics of the Pairing Correlations in a Unitary Fermi Gas,
Phys. Rev. Lett. **102**, 085302 (2009) (cit. on pp. 120, 121).

[159] M. Greiner, C. A. Regal and D. S. Jin,
Probing the Excitation Spectrum of a Fermi Gas in the BCS-BEC Crossover Regime,
Phys. Rev. Lett. **94**, 070403 (2005) (cit. on pp. 123, 125, 126, 131).

[160] J. Plata, *Magnetic-modulation spectroscopy of an atomic Fermi gas in the BCS-BEC crossover: Dissociation spectra in the Bose-Einstein condensate regime*,
Phys. Rev. A **74**, 013603 (2006) (cit. on p. 124).

[161] D. S. Petrov, *Three-body problem in Fermi gases with short-range interparticle interaction*,
Phys. Rev. A **67**, 010703 (2003) (cit. on p. 127).

[162] B. Gao, *Binding energy and scattering length for diatomic systems*,
Journal of Physics B: Atomic, Molecular and Optical Physics **37**, 4273 (2004) (cit. on p. 127).

[163] T. Mendes-Santos, X. Turkeshi, M. Dalmonte and A. Rodriguez,
Unsupervised Learning Universal Critical Behavior via the Intrinsic Dimension,
Phys. Rev. X **11**, 011040 (2021) (cit. on p. 133).

[164] A. L. Gaunt, T. F. Schmidutz, I. Gotlibovych, R. P. Smith and Z. Hadzibabic,
Bose-Einstein Condensation of Atoms in a Uniform Potential,
Phys. Rev. Lett. **110**, 200406 (2013) (cit. on p. 134).

[165] B. Mukherjee, Z. Yan, P. B. Patel, Z. Hadzibabic, T. Yefsah, J. Struck and M. W. Zwierlein, *Homogeneous Atomic Fermi Gases*, Phys. Rev. Lett. **118**, 123401 (2017) (cit. on p. 134).

[166] N. D. Mermin and H. Wagner, *Absence of Ferromagnetism or Antiferromagnetism in One- or Two-Dimensional Isotropic Heisenberg Models*, Phys. Rev. Lett. **17**, 1133 (1966)
(cit. on p. 134).

[167] P. C. Hohenberg, *Existence of Long-Range Order in One and Two Dimensions*,
Phys. Rev. **158**, 383 (1967) (cit. on p. 134).

[168] D. Phan and A. V. Chubukov,
Following the Higgs mode across the BCS-BEC crossover in two dimensions,
Phys. Rev. B **107**, 134519 (2023) (cit. on p. 134).

[169] L. Sobirey, N. Luick, M. Bohlen, H. Biss, H. Moritz and T. Lompe,
Observation of superfluidity in a strongly correlated two-dimensional Fermi gas,
Science **372**, 844 (2021) (cit. on p. 134).

[170] A. Kell, *Optical Trap for Ultracold 2D Fermi Gas*,
Master thesis: Rheinische Friedrich-Wilhelms-Universität Bonn, 2017 (cit. on p. 134).

[171] C. Linse, *Construction of a ^{6}Li spectroscopy cell and Doppler-free spectroscopy*,
Bachelor thesis: Rheinische Friedrich-Wilhelms-Universität Bonn, 2014 (cit. on p. 137).

[172] M. Zawierucha, *Design and Construction of a Sodium Spectroscopy Cell and a Lock-in Amplifier for Frequency Locking a Laser*,
Bachelor thesis: Rheinische Friedrich-Wilhelms-Universität Bonn, 2015 (cit. on p. 138).

[173] J. Andrijauskas, *Optical dipole trap laser setup in a new quantum gas experiment*,
Master thesis: Rheinische Friedrich-Wilhelms-Universität Bonn, 2015 (cit. on p. 138).

Acknowledgements

I am indebted to Michael Köhl who allowed me to spend all these years in his research group. Here, I came in contact with a wide range of research topics and learned so much more than I could have imagined. Not only the diverse interesting physical systems but also all the technologies necessary to run experiments in quantum optics. His scientific input is always extremely helpful and on point. I am also grateful for his trust in us running the experiment properly and producing meaningful results.

A huge thank you to the other "Humphries", especially Andreas Kell and Daniel Eberz who made up the team in the decisive phase of my PhD. These experiments are best run in a team and with you this was always pleasant and went extremely smoothly. I am also very happy to have been introduced to the experiment by Kuiyi Gao, with his phenomenal knowledge of quantum gas experiments, and Martin Link, who opened completely new avenues for the experiment by his interest in machine learning. I would also like to thank Benjamin Rauf and Valentin Jonas for their contributions to the experiment during my time. Additionally, this work would not have been possible if the other constructors of the experiment, Alexandra Behrle, Timothy Harrison, Justas Andrijasukas, Max Zawierucha and Christoph Linse had not build such a well working machine.

Akos Hoffmann has been an invaluable help in explaining and repairing all the electronics of the experiment. Thank you also for quickly providing solutions to every technical problem we encounter, be it electronics, water cooling or chemistry.

I would like to thank Tina Naggert, who is still missed, Janek Fleper and Beate Sund for their help in various administrative tasks. Thank you to Dietmar Haubrich for ensuring that the institute's infrastructure always fulfills the demands of our experiment and handling all occuring problems promptly.

Of tremeondous importance are all the other members of the group, past and present, who are not only colleagues but also friends. It was a pleasure to meet you on a daily basis, even more so during the pandemic with very limited other contacts. It is always great fun to spend an odd evening together with you. Thank you to Thorsten Langerfeld, Jonas Schmitz, Santhosh Surendra, Valentin Jonas, Andrea Bergschneider, Nicola Wurz, Marcell Gall, Pascal Kobel, Felix Rönchen, Jeffrey Chan, Moritz Scharfstädt, Nick Klemmer, Tobias Kree, Janek Fleper, Steffen Gohlke, Ralf Berner, Andreas von Haaren, and Jens Samland. An extra thank you to everyone who was around helping to fix water leaks at night.

A special thank you goes to Andreas Kell, Daniel Eberz, Andrea Bergschneider, Thorsten Langerfeld and Santhosh Surendra for proofreading parts of the thesis.

Lastly, I would also like to thank my family for believing in me and supporting me financially even though I could never enthuse you fully for quantum mechanics.