Production and Characterisation of Dipolar Bose–Einstein Condensates

Péter Juhász

Pembroke College
University of Oxford

This thesis is submitted for the degree of

Doctor of Philosophy

Trinity 2023

Abstract

Production and Characterisation of Dipolar Bose–Einstein Condensates

Péter Juhász

Remarkable progress in the field of ultracold atoms has enabled the study of a great variety of topics in many-body quantum mechanics. The precise control of key parameters, such as interactions, temperature, density, internal and external degrees of freedom, dimensionality and the trapping geometry makes them a powerful and flexible experimental platform.

The ability to create degenerate samples of atoms which feature long-range and anisotropic dipole–dipole interactions besides the more conventional short-range and isotropic contact interactions drew considerable attention, enabling the creation of quantum droplets and a supersolid phase. This thesis reports on experimental and theoretical progress in investigating dipolar many-body quantum systems. We present an overview of our experimental apparatus and the techniques used for obtaining a Bose–Einstein condensate (BEC) of erbium. We then discuss our experimental sequence for producing a quantum degenerate gas, creating a quasi-pure BEC with 2.2×10^5 atoms. To optimise the production of erbium BECs, we explore density- and temperature-dependent losses in ^{166}Er and identify six previously unreported resonant loss features. Finally, to enable studies of density-dependent phenomena, we present a theoretical investigation of dipolar condensates in box-like traps, where we explore stability and how one can use it to replicate properties of an infinite, homogeneous system to study dipolar physics. We found that traps with hard walls trigger roton-like density oscillations even if the bulk of the system is far from the roton regime, so smoother potentials are better suited to recreate homogeneous conditions. This sets the ground for future experiments, where the atoms will be loaded into a box trap to enable studies of systems which are tightly trapped in one direction but homogeneous in the other two.

Acknowledgements

The University. *Azure, upon a book open proper leathered Gules garnished Or having on the dexter side seven seals of the last, the words* DOMINVS ILLVMINATIO MEA*, all between three open crowns two and one Or.*

Writing acknowledgements is perhaps one of the most important and most difficult parts of writing a thesis. There is a long road and a series of events, influenced by a large number of people, which leads to eventually producing a DPhil thesis.

It is then only right to begin acknowledgements with my family, who helped me in a variety of ways. In a large part it is thanks to them that I got this far and it is difficult to overestimate their impact. I am grateful for all their help.

Acknowledgements should then go to my research supervisor and other members of my research group I had the pleasure to work with. Prof. Robert Smith, who supervised my DPhil, is an exceptional experimentalist with a wide-ranging skill set that one can only envy. His experience shined through our discussions in group meetings and I was always impressed by the number of ways he could look at scientific problems, often offering insights and viewpoints that we (the students) did not consider. I should then thank Milan Krstajić next, who started off as a PhD student a year ahead of me and then stayed on as a postdoc, and whom I spent most time in the lab with. His mischievous sense of humour brought great joy and a good atmosphere to the lab and made everyone at ease immediately. This was compounded by his impeccable work ethics, often staying past the late hours of the night to complete a measurement, and by his generous willingness and aptitude to teach me and other students in the lab the experimental aspects of our work. It is fair to say that it was from him that I learnt the most about experimental methods and best practices, as well as experimental design (whether optics, electronics

or 3D modelling). Besides all that, as we came from a similar background, we had many interesting conversations. I should then turn to Lucas Hofer, whose wicked sense of humour, lacking any sense of shame or taboos, made any conversation memorable. As a fellow computing enthusiast, it was often through him that I learnt about the latest computing tools and machine-learning developments, some of which I then used myself, including for work presented in this thesis. Turning then to Anna Marchant, a postdoc in the first half of my DPhil, I should thank her for her *very English* sense of humour, as well as more generally for her always lending an insight into the British psyche. This was a most interesting experience and I learnt about the country I lived in through her a lot. Finally, I should thank Jiří Kučera and Gavin Lamb, two other DPhil students, for the time and kindness they offered during my time. Although we worked less together, it was always a pleasure to be around them.

I would also like to acknowledge all those who funded me during my DPhil, principally the Royal Society, but also the University of Oxford, Pembroke College, St John's College, President László Sólyom and the (Hungarian) National Talents Programme. I would like to particularly thank St John's and its fellows (most notably Dr Georg Viehhauser) for providing me with ample space and a very good company over the last year of my studies. I thoroughly enjoyed the privileges extended there which made a real difference.

Besides these principal contributors, there are a number of other people and entities who had an indirect impact, more than could ever be listed.

I shall thank the physics teachers who helped me along the way—it was from them that I learnt a large part of physics which I still rely on. In this regard, I would particularly like to thank Tibor Szokolai and Gábor Horváth of *Fazekas*, who had an essential contribution, and Máté Vigh and Péter Vankó, who generously organised extracurricular classes at *BME*. They all had a decisive impact on me and many other students who were lucky enough to be taught by them.

I would also like to thank all the teachers of *Piarista Gimnázium* and the fellows of Trinity College, Cambridge, who taught me either physics or otherwise. The knowledge acquired at these institutions is a strong, ongoing basis that I am able to build on. Both of them had an strong impact on me and I take great pride in being affiliated with both. Particular thanks goes to Trinity for not only providing first-class education but also for being a home to me during my undergraduate studies.

Besides physics, it is the people that made these years special. I would therefore like to thank all those who made the past years a time to remember fondly. Particularly notable is the contribution of Norbert Fehér and the other Trinity Hungarians (Eszter, Gábor, Roland), who proved to be good friends throughout the years. Besides them, an important help was Attila Szabó, who, being only one year ahead of me in his studies,

was always happy to help with some advice. A particular contribution of his to this thesis was providing me with his own thesis, which influenced many of the stylistic choices presented here. I am also grateful to my friends from Cambridge more generally, for making that time truly memorable. Equally important were the friends from school, who always provided me with good company when I was back home.

The contribution of Sára Mándoki also needs particular acknowledgement. Many DPhil students are prone to treating their thesis as their *magnum opus*,[1] and she was a key company in my drive for typographic perfection, and for making sure that I do not refer to lines in figures as green when in fact they are red.[2] I thank her for this and for her company throughout the years.

I would also like to thank the Oxford University Catholic Chaplaincy not only for providing me with accommodation during a large part of my studies but also with good company, which was especially important during the years of COVID. Equally formative were the contributions of Mgr Mark Langham and Fisher House during my undergraduate studies, and that of the Order of Malta later on, for being a spiritual home. Finally, I would like to acknowledge my colleagues at the Permanent Observer Mission of the Holy See to the United Nations, where I spent almost half a year as a break from my studies.[3] I enjoyed my time there very much.

There are many other names that come to mind who had a contribution towards this work and it would be impossible to list them all. As I always found that mottos were an extravagant way of self-expression, I decided to use that space as a brief armorial of those who supported me during my time.[4] Special thanks goes to Dr Malcolm Golin in this regard, for his help with blazoning the arms presented there. I hope the reader will enjoy these as a form of light entertainment while reading what is otherwise a heavy text.

[1] Me being no exception.

[2] Let me introduce you to the life of someone with colour blindness.

[3] Although it was perhaps one of the most exciting experiences to work on a figure for this thesis in the UN's General Assembly Hall.

[4] The drawings were adapted from various sources [1].

Contents

1
Introduction

The Royal Arms. *Quarterly, first and fourth Gules three lions passant gardant in pale Or armed and langued Azure (England), second quarter Or a lion rampant within a double tressure flory-counter-flory Gules (Scotland), third quarter Azure a harp Or stringed Argent (Ireland).*

Quantum theory, formulated in the early 20$^{\text{th}}$ century, set the foundations of our current understanding of the microscopic world. As an early theoretical result, Satyendra Nath Bose [2] and Albert Einstein [3] predicted almost 100 years ago that if a gas made of identical bosons is cooled down sufficiently, a large number of atoms will occupy the ground state of the quantum system, yielding a novel state of matter, a *Bose–Einstein condensate (BEC)*. Interestingly, even though the theory was formulated for a non-interacting system, Bose–Einstein condensation is an inherently many-body effect [4]. Following breakthrough research in laser cooling [5] and the development of the magneto–optical trap (MOT) [6, 7], the first BEC was realised some 70 years later, in 1995 in quick succession at JILA using ^{87}Rb [8] and at MIT using ^{23}Na [9] (an early claim at Rice University using ^{7}Li was ultimately confirmed two years later [10]). Reaching quantum degeneracy in a fermionic system followed in 1999 [11], soon after realising a BEC with spin-polarised H in 1998 [12], the original candidate system for creating a

BEC since the 1970s. These results were significant: the development of laser cooling and trapping methods, and the realisation of a Bose–Einstein condensate were rewarded by Nobel Prizes in 1997 [7, 13] and 2001 [14]. The creation of macroscopic quantum-correlated systems opened the possibility of studying a large variety of phenomena in many-body quantum physics, with applications ranging from atom interferometry [15] and precision measurements [16] to atomtronics [17].

Remarkably, quantum theory succeeded in explaining experimental observations already in its early days, such as black-body radiation [18] and the photoelectric effect [19]. However, to this day, the description of strongly-correlated quantum systems remains difficult, even if they consist of only a relatively small number of particles, and while various phenomena (e.g. superconductivity) have been studied extensively within condensed matter systems, a comprehensive description is still lacking. To tackle this problem, Richard Feynman proposed using a model system where particles can be added one-by-one, allowing the study of the properties of a complex quantum system via the well-understood model system, called the *quantum simulator* [20]. Essentially, this requires tuning the Hamiltonian of the quantum simulator to match a desired model, bringing it into a specific initial state and then looking at the subsequent evolution. This requires the precise experimental control of quantum states, and various platforms were identified for this purpose, including photons [21], trapped ions [22] and cold atoms in dipole traps [23].

Ultracold atoms provide a particularly attractive option for realising a quantum simulator and for studying many-body quantum physics more generally, given their high degree of controllability, the various detection possibilities and the extreme physical parameter regimes that can be explored [23]. The research that ensued after creating the first BECs yielded many prominent achievements [24]. The early days of the field saw the demonstration of the (matter-wave) interference of two BECs [25], confirming the spatial coherence of the quantum state [26], the realisation of the crossover from Bardeen–Cooper–Schrieffer-type (BCS) superfluidity to Bose–Einstein condensation [27–29], the observation of quantised vortices [30–32], the lattice-ordering of vortices [33] and other phenomena related to superfluidity [34, 35], and the observation of the Mott insulator transition in optical lattices [36]. Later on, building on these results, attention turned towards systems out of equilibrium [37], simulating condensed matter lattice models [38], exploring the dynamics of phase transitions [39], precision metrology [40, 41] and quantum information [42].

This wealth of research directions is supported by a wide range of tools and techniques that keep the field of ultracold atoms at the forefront of quantum simulation. Spatial confinement of atoms can be achieved with either magnetic fields [6] or non-resonant light

fields via the AC Stark effect [43]. In both cases, the imposed field induces a potential well (or, in the case of optical lattices, a periodic array of wells), which traps the atoms. Methods most commonly used in experiments, e.g. magnetic time-orbiting-potential (TOP) traps [44] and optical dipole traps using focused Gaussian beams [43], result in approximately harmonic trapping potentials. While this is sufficient for many applications, offering a large degree of freedom in terms of the trap aspect ratio, anharmonicity and tightness of the trap, it falls short of simulating models requiring a homogeneous density. Therefore, optical box traps have been developed [45–47], enabled by manipulating light intensity patterns. These open the door for numerous experiments that require uniform systems, such as exploring the dynamics of turbulent flow [48] and the phase transition to BEC [49]. Besides such bulk traps, optical lattices have also been exploited to realise solid-state crystal-like structures [50, 51], offering full control over lattice parameters such as the tunnelling rate. Combined with single-site-resolution imaging (a *quantum gas microscope* [52]), this is a particularly powerful configuration for quantum simulation.

In addition to the ability of tuning the trap geometry, the creation of systems with reduced dimensionality is also possible. By imposing sufficiently tight trapping along one or two directions, forcing the system into the ground state along these axes, dimensions can effectively be frozen out. This led to the observation of inherently one- [53, 54] and two-dimensional effects [55].

The strength and nature of interparticle interactions can also be tuned. Due to the diluteness of ultracold gases, interactions can be described by two-body collisions. Furthermore, due to the low energies involved, interactions of most species can be described by short-range contact interactions, whose strength and (attractive or repulsive) nature is determined by a single parameter, the s-wave scattering length (a_s). Feshbach resonances can be used to tune this parameter [56–58], with possibilities ranging from changing the interactions from repulsive to attractive, to switching them off completely or making them as strong as quantum mechanics allows (the so-called unitarity limit).

Recent progress expanded the spectrum of elements that can be cooled to quantum degeneracy (see Fig. 1.1). The condensation of more exotic elements with high magnetic moments (chromium (Cr) [59], dysprosium (Dy) [60], erbium (Er) [61], thulium (Tm) [62] and europium (Eu) [63]) added a new aspect to experiments [64, 65] by introducing long-range and anisotropic dipole–dipole interactions (DDI) on top of the simpler contact interactions present in all atomic species, which are inherently short-range and isotropic.

First results obtained with chromium demonstrated how the anisotropic nature of the dipolar interaction affects the stability [66] and the aspect ratio [67] of a trapped BEC, and the density patterns emerging after the collapse of a cloud brought to instability [68]. Further developments came with the condensation of dysprosium and erbium, as their

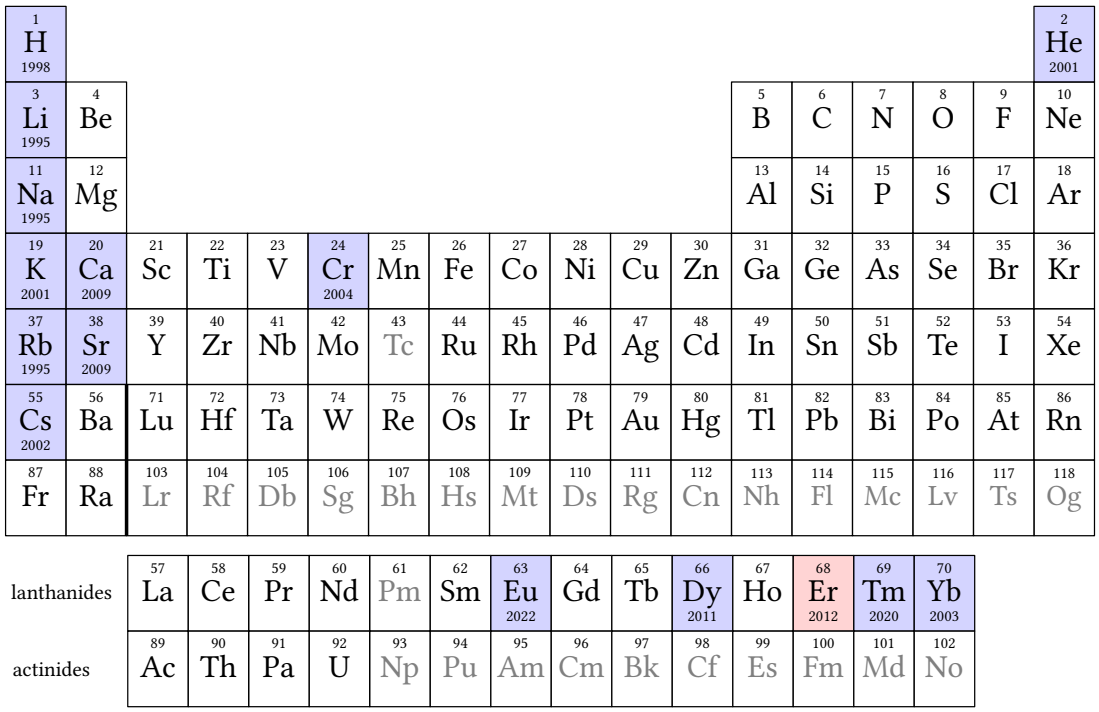

Figure 1.1. Bose–Einstein condensation of different atomic species. The periodic table of elements with those successfully condensed into a BEC highlighted. The fields also state the year when condensation was first reported. In our experiment we use erbium, highlighted with a different colour.

higher magnetic moments and larger atomic masses make dipolar interactions even more pronounced relative to the contact interactions. Furthermore, contact interactions can be tuned more readily in these elements due to the large number of accessible Feshbach resonances [61, 69, 70]. This eventually allowed the detection of the so-called roton minimum in the excitation spectrum of a strongly dipolar gas [71, 72], showing similarities with the strongly-correlated system of liquid helium [73–75], with the additional ability to tune the roton gap via tuning the contact interactions. The observation of quantum droplets, emerging from a dipolar cloud collapsing due to strong anisotropic interactions [76, 77], was also linked to the roton spectrum, eventually leading to the recent creation of the first *supersolids*,[1] an exotic state of matter combining the spatial density modulation of solids and the phase coherence of superfluids [80–82]. Besides these, the first observation of vortices was reported recently [83], while the dynamics of phase transitions to a dipolar BEC or a supersolid, as an extension of Kibble–Zurek theory [48, 49] to dipolar systems,

[1] Note that supersolids have been created earlier by coupling a BEC to optical cavities [78] and by inducing a 'stripe phase' in a BEC using lasers [79]. Compared to these systems, dipolar supersolids are special as the density modulation is induced by the internal interactions themselves, rather than an external factor.

remains to be explored.

Inspired by this success, there are now a number of experiments with ultracold dipolar atoms (mostly Er and Dy) around the world, including the ones at Innsbruck [61], Stuttgart [77], Florence [84], Stanford [60], Harvard [85], MIT [86], Bonn [87], Paris [88] and Hong Kong [89]. Furthermore, some groups are preparing experiments with atomic mixtures, adding another layer of possibilities, using dysprosium and potassium at Paris [90] and Innsbruck [91], and erbium and dysprosium also at Innsbruck [92]. Besides the work in laboratories, dipolar physics also inspired art: a stained glass church window depicting the dipolar BEC collapse has been created in Lohmar, Germany [64].

In our group, we are condensing erbium. Furthermore, as both the BEC phase transition and the formation of a supersolid are sensitive to the atomic density, we are implementing a trap providing a homogeneous atomic density to study these processes in depth. Combining the uniform trapping of atoms with long-range and anisotropic dipole–dipole interactions will yield a unique apparatus for exploring many-body physics, and to achieve this goal, over the course of my DPhil I studied the stability of dipolar BECs in box-like traps and their atom loss properties, to extend the lifetime of samples and to achieve a large atom number in the trap. Furthermore, another atomic species, potassium, is currently being added to the system, as the introduction of such impurities will enable the investigation of Bose polaron physics [93] and qubit decoherence in non-Markovian reservoirs [94, 95].

Finally, it is worth noting that other fields aimed at exploring quantum simulation with dipolar systems, including ultracold heteronuclear molecules possessing a large electric dipole moment [96–98], Rydberg atoms [99, 100] and trapped ions [22], are developing at an equally fast pace, sharing many of the experimental techniques with ultracold atoms. All of these systems, as well as ultracold atoms, can also be used to probe fundamental physics [101]. The advantage of these alternative systems for researching many-body dipolar physics is that they offer dipole moments up to two orders of magnitude larger than that of ultracold atoms, but they are limited by the short lifetime of samples and typically require more complex apparatus.

1.1 Thesis outline

In this thesis, I offer an overview of the work undertaken over the course of my DPhil, in collaboration with my research supervisor and colleagues. This included designing and building an experimental platform for investigating many-body quantum physics

in dipolar systems, as well as the theoretical investigation and measurement of such systems.

The thesis is structured as follows:

- **Chapter 1** is an introduction to the field this thesis is set in and provides a brief overview of the scientific results presented in the chapters that follow.
- **Chapter 2** is a relatively thorough, but necessarily incomplete review of the key properties of magnetic lanthanides in general and erbium in particular, the element we use in our experiment for exploring many-body quantum physics. It also presents an introduction into the main theoretical concepts one encounters while running ultracold atom experiments, including atomic interactions and the arising scattering properties, and a mean-field description of Bose–Einstein condensates via the Gross–Pitaevskii equation. This chapter is based on Refs. 64, 65, 102 and 103.
- **Chapter 3** describes the experimental apparatus used for producing an erbium Bose–Einstein condensate, built together with Dr Anna Marchant, Lucas Hofer, Jiří Kučera, Gavin Lamb and Dr Milan Krstajić, whose thesis [104] contains a thorough description of the apparatus, still in the building phase at that time. This was used as the basis for parts of the description presented in this thesis. Besides the daily running of the experiment, I was particularly responsible for designing and implementing the optical dipole trap, used for evaporation and transport, and for a variety of smaller pieces of instrumentation.
- **Chapter 4** discusses the attainment of an erbium BEC, including the experimental sequence and the details of the optical dipole trap to produce it. Results presented in this chapter rely on data taken together with Dr Milan Krstajić and Jiří Kučera.
- **Chapter 5** describes the measurement of the three-body loss coefficient of erbium and the investigation of its magnetic field and temperature dependence. Data were taken together with Dr Milan Krstajić and Jiří Kučera, who also derived additional contributions to the atom loss. Besides taking data and preparing a manuscript for publication, I was also responsible for data analysis.
- **Chapter 6** describes a theoretical study of the stability of dipolar BECs in box-like traps, with the aim of creating a system resembling an infinite, homogeneous one. The project was started by two undergraduate students, David Strachan and Edward Gandar, and was developed more fully by me and Dr Milan Krstajić. A description of some early results and the numerical techniques behind the simulation, still in the data-taking phase at that time, form a small part of his thesis [104]. I was

responsible for developing an independent numerical technique so that results can be checked against each other, data analysis and the preparation of a manuscript for publication.

- **Chapter 7** presents the conclusions of this thesis and a brief outlook towards future directions of research.

On a stylistic note, inline fractions of the form a/bc will be presented in this thesis, which should be read as $a/(bc)$, as the brackets have been omitted for readability. Furthermore, a^* marks the complex conjugate of a, while $a^\dagger$ is the Hermitian adjoint of operator a. The $\equiv$ symbol is used to note a definition, i.e. $a \equiv b$ means that b is defined by a. Finally, the figures aim to be colourblind-friendly by using appropriate colour palettes (see e.g. Ref. 105 about possible methods and the importance of this).

1.2 List of publications

Results presented in this thesis were published in the following articles:

- L. R. Hofer, M. Krstajić, P. Juhász, A. L. Marchant and R. P. Smith, 'Atom cloud detection and segmentation using a deep neural network', Mach. Learn.: Sci. Technol. 2, 045008 (2021), presented in Chapter 3.

- M. Krstajić, P. Juhász, J. Kučera, L. R. Hofer, G. Lamb, A. L. Marchant and R. P. Smith, 'Characterisation of three-body loss in ^{166}Er and optimised production of large Bose–Einstein condensates', arXiv, 2307.01245 (2023), presented in Chapter 5.

- P. Juhász, M. Krstajić, D. Strachan, E. Gandar and R. P. Smith, 'How to realize a homogeneous dipolar Bose gas in the roton regime', Phys. Rev. A **105**, L061301 (2022), presented in Chapter 6.

2
Ultracold erbium atoms

Hungary. *Per pale, the first barry of eight Gules and Argent, the second Gules, on a triple mount Vert a crown Or, issuant therefrom a double cross Argent.*

Our apparatus employs a highly magnetic rare-earth element, erbium, to explore many-body quantum physics with long-range interactions. We review the key properties of erbium in §2.1, with a particular emphasis on its electron structure which is the source of its high magnetic moment and enables the laser cooling of this element. In §2.2 we explore the interactions between erbium atoms, inducing exotic scattering properties. Finally, in §2.3 we revisit the theory of Bose–Einstein condensation and how dipole–dipole interactions affect it. This review of the theory behind dipolar ultracold-atom experiments is carried out with the aid of Refs. 64, 65 and 102.

2.1 Why erbium?

In our apparatus, we use a highly magnetic element called erbium, which possesses a considerable, permanent magnetic dipole moment. To study many-body quantum physics involving long-range and anisotropic interactions, we produce a Bose–Einstein condensate using this element. Besides erbium, chromium (Cr) and dysprosium (Dy) offer the possibility of similar studies [64, 65]. Furthermore, electric dipoles can be induced

in heteronuclear molecules [96–98] or Rydberg atoms [99, 100] in an electric field, and long-range interactions beyond the $1/r^3$ scaling can be achieved using optical cavity- or waveguide-mediated interactions [108–110], or phonon-mediated interactions in trapped-ion systems [22]. These systems often exhibit orders of magnitude larger dipole strengths than what is achievable with magnetic dipoles, but other limitations arise in these systems, e.g. short lifetimes, density limitations and rapid dissipation.

Erbium (chemical symbol Er) is a rare-earth metallic element in the lanthanide series of elements in the periodic table, with atomic number $Z = 68$ and average atomic mass $m = 167.26$ Da.[1] It is found principally in the minerals monazite and bastnäsite, and can be isolated by ion exchange and solvent extraction. It was discovered in a sample of gadolinite ore by Carl Gustaf Mosander in 1843 [111] and has a fascinating history [112]. Its name, along with the elements terbium (Tb), ytterbium (Yb) and yttrium (Y), is derived from the mine from where the ore originated: Ytterby, near Stockholm, Sweden. Interestingly, the originally separated parts of gadolinite were in fact a mixture of several oxides and pure erbium could only be isolated by Wilhelm Klemm and Heinrich Bommer as late as 1934. Ytterby proved quite fertile for the discovery of new elements. Other elements which were discovered using the ore mined there were named gadolinium (Gd, after gadolinite), scandium (Sc, after *Scandia*, the Latin name for Scandinavia), holmium (Ho, after *Holmia*, the Latin name for Stockholm), thulium (Tm, after *Thule*, an Ancient Greek word for Scandinavia), dysprosium (Dy, after *dysprositos*, the Greek word meaning 'hard to find') and lutetium (Lu, after *Lutetia*, the Latin name for Paris, where the ore was analysed). Altogether, ca. half of the lanthanide elements were discovered in Ytterby.

Despite being a rare-earth element, erbium has a variety of uses. Er^{3+} has a prominent transition at 2940 nm, which is highly absorbed by water and is thus well-suited for laser surgery and other medical applications. It is also used as a dopant in optical fibres, and erbium-doped fibre amplifiers are widely used in optical communications. Er_2O_3 has a pink colour and is sometimes used as a colourant for glass and porcelain. When alloyed with metals (e.g. vanadium), erbium lowers their hardness and improves their workability. Furthermore, the erbium–nickel alloy Er_3Ni has an unusually high specific heat capacity at liquid-helium temperatures and is therefore used in cryocoolers.

Within ultracold atomic physics, the interest in erbium is due to its large magnetic dipole moment, energy level structure favourable for laser cooling and the accessibility of both bosonic and fermionic isotopes. These make it suitable for implementing standard cooling and trapping techniques to reach low temperatures and quantum degeneracy, and hence for exploring many-body quantum phenomena with dipole–dipole interactions.

[1] The Da (dalton) is the unified atomic mass unit, $1\,\text{Da} \approx 1.66 \times 10^{-27}$ kg.

Isotope	^{162}Er	^{164}Er	^{166}Er	^{167}Er	^{168}Er	^{170}Er
Abundance	0.14%	1.60%	33.5%	22.9%	27.0%	14.9%
Statistics	boson	boson	boson	fermion	boson	boson

Table 2.1. Naturally occurring isotopes of erbium. Abundances and spin statistics of stable isotopes of erbium. Data taken from Ref. 114.

2.1.1 Basic properties

By appearance, erbium is a soft, pale, silvery metallic substance which easily oxidises and hence covers itself with a layer of oxide when exposed to air. Like many other lanthanides, it has a high melting point at 1529 °C and a boiling point at 2868 °C [113]. It is slightly toxic if ingested, and in dust form presents a fire and explosion hazard. It naturally occurs in six stable isotopes, five of them bosonic, as shown in Table 2.1. It has a large magnetic dipole moment of $7\mu_B$, where μ_B is the Bohr magneton.[2] Although a large magnetic moment is not common to all lanthanides, some of them have the highest magnetic moment of all elements. Dysprosium and terbium have the largest magnetic moments, $10\mu_B$, holmium has $9\mu_B$ and europium has $7\mu_B$, the same as erbium. However, we will see that it is not only the magnetic moment but also the atomic mass which influences the strength of dipole–dipole interactions. Therefore, even though chromium (Cr) also has a high magnetic moment of ca. $6\mu_B$, its significantly smaller atomic weight ($m = 52.0\,\mathrm{Da}$) yields a less pronounced dipolar interaction.

2.1.2 Electron configuration

The electron configuration of erbium can be written as

$$[\mathrm{Xe}]\, 4f^{12}\, 6s^2\,, \tag{2.1}$$

which is an example of a so-called submerged-shell configuration, where the valence electrons fill the $6s$ subshell while the $4f$ subshell remains partially filled. This configuration is responsible for the high magnetic moment of erbium, as the two holes in the $4f$ subshell greatly contribute to the total angular momentum of $J = 6$ in the ground state, yielding a high magnetic moment.

This ground state electron configuration can be understood via the Madelung rule and the aufbau principle, stating that a subshell with a lower $n + \ell$ will have a lower

[2] The Bohr magneton is defined as $\mu_B = e\hbar/2m_e \approx 9.274 \times 10^{-24}\,\mathrm{J\,T^{-1}}$, where e is the elementary charge, $\hbar$ is the reduced Planck constant and m_e is the electron mass.

energy and so will fill first, where n is the principal quantum number and ℓ is the orbital quantum number. Given the $4f$ subshell has $n + \ell = 7$ and the $6s$ subshell has $n + \ell = 6$, the latter will fill first, leaving two gaps in the former.

We can proceed with finding the ground state using LS coupling of the $4f$ electrons. According to Hund's first rule, the ground state will have maximum multiplicity.[3] Therefore, the two gaps will result in two unpaired electrons in the ground state with a spin quantum number $m_s = {}^1/_2$, yielding a total spin quantum number $S = 1$ and a multiplicity $2S + 1 = 3$. According to Hund's second rule, the ground state will also have the highest total orbital quantum number, so the unpaired electrons will have magnetic quantum numbers $m_\ell = +2$ and $+3$, yielding a total orbital quantum number $L = 5$. Finally, as per Hund's third rule, the ground state will maximise the total angular momentum quantum number J as the subshell is more than half-filled, so S and L will couple to give $J = S + L = 6$. Therefore, the term symbol of the ground state in standard spectroscopic notation can be written as

$$^3\mathrm{H}_6\,. \tag{2.2}$$

In general, the coupling of electrons in heavy elements is more complicated, as the spin–orbit interaction becomes comparable to the Coulomb interaction for electrons in higher orbitals. Therefore, the LS coupling scheme is no longer applicable and the jj coupling scheme needs to be used. A special case of this is called J_1J_2 coupling, relevant for most of the excited states of lanthanides [115]. In this scheme, LS coupling is assumed to work separately for the $4f$ subshell and the outer valence electrons, yielding J_1 and J_2, respectively. These then sum up to give the total angular momentum $\mathbf{J} = \mathbf{J}_1 + \mathbf{J}_2$, with corresponding quantum number J. The whole state is then denoted as $(J_1, J_2)_J$. In the ground state, the two $6s$ electrons are in the $L = 0$, $S = 0$ state, yielding $J_2 = 0$, and the $4f$ electrons give an overall $^3\mathrm{H}_6$ state with $L = 5$, $S = 1$ and $J_1 = 6$ as stated previously. Therefore, the final coupled state is $(6, 0)_6$.

Let us use the prominent 401 nm transition of erbium to demonstrate the electron coupling of an excited state. This state results from the excitation of an electron from the $6s$ orbital into the $6p$ orbital. Therefore, within this scheme, the $6s$ and $6p$ electrons couple together to a $^1\mathrm{P}_1$ state by LS coupling, and the rest of the electrons couple to $^3\mathrm{H}_6$ as before. These two states then couple to give the $(6, 1)_7$ state, and so this excited state can be written in full as

$$[\mathrm{Xe}]\, 4f^{12}\, (^3\mathrm{H}_6)\, 6s\, 6p\, (^1\mathrm{P}_1)\, (6, 1)_7\,. \tag{2.3}$$

[3] The multiplicity is defined as $2S + 1$, where S is the total spin angular momentum.

It is interesting to note that nuclei with an even number of protons and neutrons have no nuclear spin [115], so the bosonic isotopes of erbium likewise have no nuclear spin and hence no hyperfine structure. Nuclei with odd–even or even–odd number of protons–neutrons have a half-integer, and odd–odd nuclei an integer nuclear spin, and in accordance with this the fermionic isotope of erbium has a nuclear spin quantum number $I = {}^{7}/_{2}$ and a total angular momentum quantum number $F = {}^{19}/_{2}$.

2.1.3 Energy spectrum

The submerged-shell electron structure also gives rise to a rich spectrum of excited states. Figure 2.1 shows the energy spectrum—a total of 312 states with odd parity and 358 states with even parity have been discovered so far.[4] The ground state has even parity.

Understanding the excitation spectrum is important as laser cooling is most easily set up on a two-level system with a prominent, electric-dipole-allowed transition. However, due to the large number of energy levels in erbium, it is possible that an excited state decays with a large probability into an intermediate, lower-lying state before eventually decaying back to the ground state. Such intermediate states with long lifetimes compared to the initially excited state are called metastable states, and can be problematic as they deplete the electrons from the ground state and the target excited state, hence decreasing the efficiency of laser cooling. While metastable states can also be actively depleted using repump lasers (which pump the electrons from the metastable states to the target excited state), they are ideally avoided for technological simplicity.

In the case of erbium, the selection rules of electric dipole transitions restrict possible transitions from the ground state to an excited state with $J = 5$ or 7. To reduce the number of possible decay channels into metastable states, the excited state should not have too high an energy as the number of decay channels increases with energy, and so the probability of decaying into an intermediate state (the so-called branching ratio) can also potentially increase.

Various suitable transitions for laser cooling have been identified with wide (30 MHz) to ultra-narrow (2 Hz) linewidths [118]. While a broad transition allows the efficient operation of the Zeeman slower (ZS) due to the high scattering rate of photons, a narrow transition is better suited for operating a magneto–optical trap (MOT) due to the lower achievable temperature. In most erbium experiments, the transitions used for cooling are the ones at 401 nm [119] and 583 nm [120], with linewidths of 29.7 MHz and 190 kHz,

[4] The parity relates to the symmetry of the wave function: the parity operator $\hat{P}$ inverts the wave function through the origin. If $\hat{P}\Psi(\mathbf{r}) = \Psi(-\mathbf{r})$, the wave function is even, if $\hat{P}\Psi(\mathbf{r}) = -\Psi(-\mathbf{r})$, it is odd.

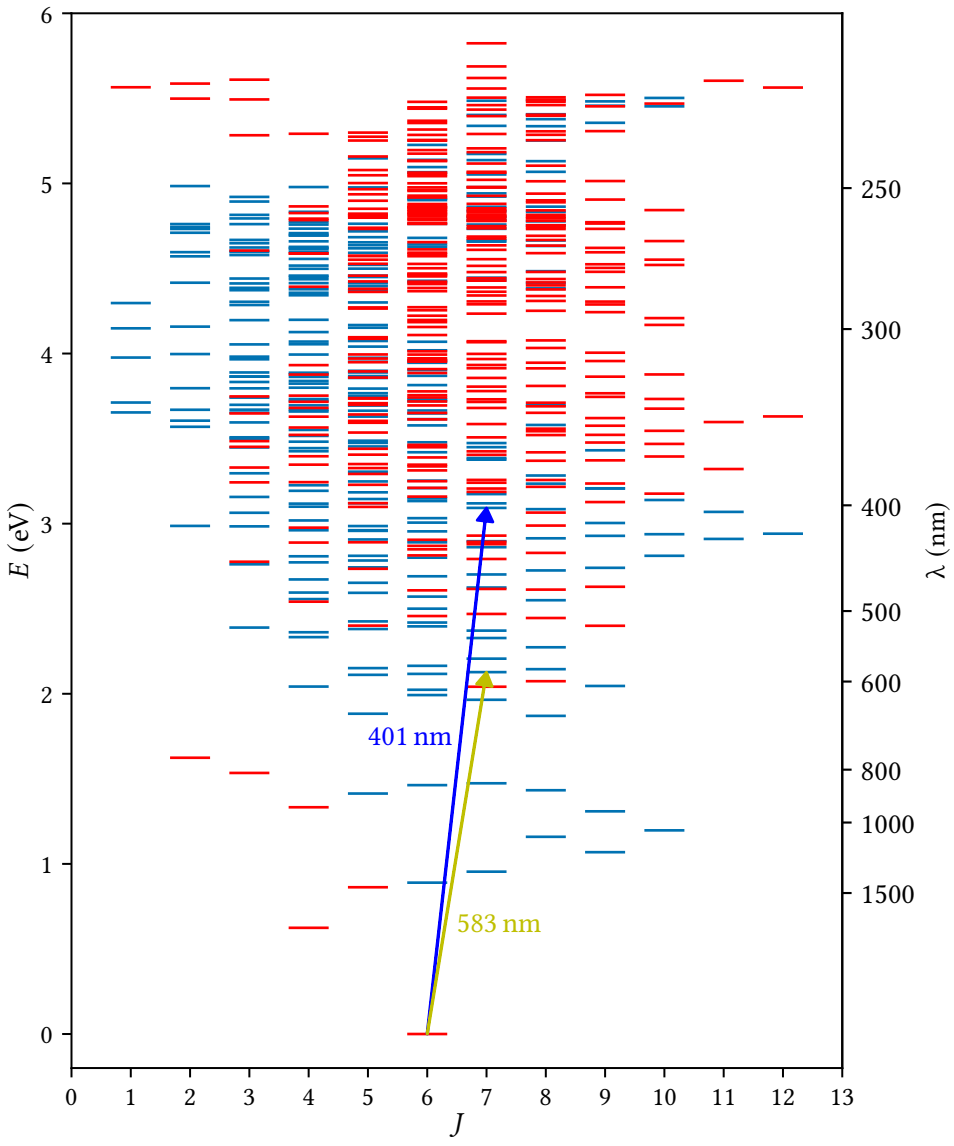

Figure 2.1. Energy levels of erbium. Energy levels of erbium arranged by their total angular momentum quantum number J. Even and odd parity levels are indicated in red and blue, respectively. The transitions used for laser cooling, at 401 nm and 583 nm, are highlighted. Spectroscopic data obtained from Refs. 116 and 117.

Transition	Blue	Yellow
Wavelength (λ)	400.91 nm	582.84 nm
Lifetime (τ)	5.4 ns	857 ns
Transition rate ($\Gamma = 1/\tau$)	1.85×10^8 s^{-1}	1.17×10^6 s^{-1}
Natural linewidth ($\Delta\nu = \Gamma/2\pi$)	29.5 MHz	186 kHz
Doppler temperature ($T_D = \hbar\Gamma/2k_B$)	707 µK	4.46 µK
Recoil temperature ($T_r = \hbar^2k^2/2mk_B$)	0.36 µK	0.17 µK
Saturation intensity ($I_{\text{sat}} = \pi hc\Gamma/3\lambda^3$)	59.8 mW cm^{-2}	0.12 mW cm^{-2}

Table 2.2. Laser cooling atomic transitions. Various properties of the transitions used for laser cooling. Here, k_B is the Boltzmann constant, ($\hbar$) h is the (reduced) Planck constant, c is the speed of light and $k = 2\pi/\lambda$ is the wavenumber. Data taken from Ref. 123.

respectively. These transitions are highlighted in Fig. 2.1 and their properties are tabulated in Table 2.2. The 401 nm (blue) transition[5] is the most prominent transition in erbium, where a 6*s* electron is excited into a 6*p* state and couples with the second 6*s* electron to the singlet state $^1P_1$. As the transition is quite wide, it has a large photon scattering rate, which makes it convenient for the initial stages of laser cooling (the Zeeman slower and transversal cooling) and imaging. The 583 nm (yellow) transition occurs when the 6*s* and the excited 6*p* electron couple to the triplet state $^3P_1$. As the transition is reasonably narrow, it is well-suited for operating the magneto–optical trap, giving a low Doppler temperature of 4.5 µK.

Due to a change in the charge distribution within the nuclei of different isotopes, the transition energies change linearly with the atomic mass. Therefore, the frequency of the laser addressing the transition needs to be changed accordingly. This is known as the isotope shift, and between neighbouring bosonic isotopes it is ca. 820 MHz for the 401 nm transition [124] and 980 MHz for the 583 nm transition [125], corresponding to an isotope shift of 410 MHz and 490 MHz per atomic mass unit, respectively.

In this thesis we concentrate on the bosonic isotopes, whose energy levels have no hyperfine structure due to the lack of nuclear spin. However, a discussion on the hyperfine structure of the energy levels of the fermionic isotope can be found in Ref. 126.

[5] We will refer to the colour associated with the wavelength of this transition as 'blue', as is the case in literature relating to erbium experiments, but we note that this wavelength is more often characterised as violet [121, 122].

2.1.4 Magnetic properties

Lanthanides are ferromagnetic, but given they have a Curie temperature below room temperature, this cannot be seen under ordinary circumstances [127]. However, they can form compounds with transition metals such as iron, nickel or cobalt, and create some of the strongest permanent magnets at room temperature. These so-called rare-earth magnets (e.g. neodymium magnets) can create permanent magnetic fields of up to 1.4 T [128] due to the high magnetic anisotropy of their crystalline structures and the high magnetic moment of the lanthanides within [129].

As mentioned in §2.1.1, some of the lanthanide elements have the highest magnetic moments in the periodic table as their electron structure results in exceptionally large quantum numbers. While most other elements have a much smaller magnetic moment, chromium (Cr) has a magnetic moment of ca. $6\mu_B$ and was in fact the first strongly dipolar element to be cooled down to quantum degeneracy.

The component of the magnetic moment along the direction of the quantisation axis is given by

$$\mu = m_J g_J \mu_B \,, \tag{2.4}$$

where m_J is the total angular momentum projection quantum number, i.e. the projection of the total angular momentum $\mathbf{J}$ along the quantisation axis (set by the direction of the external magnetic field $\mathbf{B}$), and g_J is called the Landé g-factor. Within LS coupling, the Landé g-factor can be calculated as

$$g_J = 1 + (g_s - 1)\frac{J(J+1) - L(L+1) + S(S+1)}{2J(J+1)} \,, \tag{2.5}$$

where $g_s = 2(1 + \alpha/2\pi + \cdots) \approx 2.002\,319\,3$ [130] is the electron spin g-factor and α is the fine-structure constant.[6] This formula does not include some known corrections [131], such as deviations from perfect LS coupling, a relativistic correction to the kinetic energy of high-orbit electrons and the so-called diamagnetic correction, relating to the changed interaction between electrons in a magnetic field. For the ground state of erbium ($S = 1$, $L = 5$, $J = 6$), Eq. (2.5) gives $g_J = 1.167\,053\,2$. The LS coupling correction to this is $\Delta g_{LS} = -0.001\,334\,7$, whereas the relativistic and diamagnetic corrections combined yield $\Delta g_{\text{rel.+diamag.}} = -0.001\,92$, resulting in the corrected value $g_{J,\,\text{calc.}} = 1.163\,798\,5$ [132]. This is in close agreement with the experimentally measured value of $g_{J,\,\text{exp.}} = 1.163\,801(1)$ [132]. Therefore, the magnetic moment of bosonic erbium in the energetically lowest state

[6] The fine-structure constant is defined as $\alpha = e^2 c\mu_0/2h \approx 1/137$, where e is the elementary charge and μ_0 is the permeability of free space.

with $m_J = -6$ is

$$\mu = m_J g_J \mu_B = -6.982\,806(6)\mu_B = -6.475\,861(6) \times 10^{-23}\,\mathrm{J\,T^{-1}}\,. \tag{2.6}$$

The case of excited states need to be approached slightly differently, using the J_1J_2 coupling scheme. The g-factor in this scheme can be calculated as

$$g_J = g_{J_1}\frac{J(J+1)+J_1(J_1+1)-J_2(J_2+1)}{2J(J+1)} + g_{J_2}\frac{J(J+1)+J_2(J_2+1)-J_1(J_1+1)}{2J(J+1)}\,, \tag{2.7}$$

where, g_{J_1} and g_{J_2} are the Landé g-factors of the respective parts according to Eq. (2.5). For both the 401 nm and the 583 nm transitions we have $J_1 = 6$, $J_2 = 1$ and $J = 7$. However, although the state of the electrons in the $4f$ subshell (corresponding to J_1) is $^3\mathrm{H}_6$ in both cases, the electrons in the $n = 6$ shell are in the singlet $^1\mathrm{P}_1$ state for 401 nm, yielding $g_{J_2} = 1$ for the blue, and in the triplet $^3\mathrm{P}_1$ state for 583 nm, yielding $g_{J_2} = 1.501$ for the yellow transition. Using Eq. (2.7), we get $g_{J,401} = 1.140$ and $g_{J,583} = 1.212$ theoretically, which are close to the experimentally measured values of $g_{J,401} = 1.160$ and $g_{J,583} = 1.195$ [117].

It is worth noting that given the fermionic isotope has a nonzero nuclear spin $\mathbf{I}$, the total electron angular momentum $\mathbf{J}$ couples with the nuclear spin to form the total angular momentum $\mathbf{F} = \mathbf{J} + \mathbf{I}$. This introduces a hyperfine structure into the energy levels and the Landé g-factor has to be calculated differently [126].

Besides calculating the magnetic moment, the g-factor also comes into play when the atom is subjected to a constant magnetic field $\mathbf{B}$. In this case, the energy levels with different m_J split according to the Zeeman effect. The energy shift of the $|J, m_J\rangle$ state in low magnetic fields[7] is given by

$$\Delta E_Z = m_J g_J \mu_B B\,. \tag{2.8}$$

Knowing the precise value of g_J is important, as the Zeeman shift is different for different energy levels, and the energy gap of a transition $|J_g, m_g\rangle \longrightarrow |J_e, m_e\rangle$ changes in a magnetic field according to

$$\Delta E_Z = \left(m_e g_e - m_g g_g\right)\mu_B B\,. \tag{2.9}$$

Furthermore, using radio frequency (RF) fields, one can deliberately drive the $m_J \longrightarrow m_J + 1$ transitions. Therefore, in an experiment, measuring the transition frequency can be used to calibrate magnetic fields once the relevant g-factor is known.

[7] For the bosonic isotopes, the range of 'low magnetic fields' extends to B fields higher than typically accessible in cold atom experiments, due to the unusually large spin–orbit coupling constant of erbium [131].

2.2 Atom–atom interactions

As it will be shown in §2.2.3, the length scales associated with scattering processes are on the order of $10^2 a_0 \approx 5\,\mathrm{nm}$, where a_0 is the Bohr radius.[8] As typical densities of ultracold atomic clouds are on the order of $n \approx 10^{20}\,\mathrm{m}^{-3}$, corresponding to an average interparticle distance $\langle d \rangle \approx n^{-1/3} \approx 200\,\mathrm{nm} \gg 10^2 a_0$ [24, 133], ultracold systems are usually comparatively dilute and are in the regime where scattering events can be described accurately by considering only two-body interactions.

The interaction potential U of two dipolar atoms, separated by a distance $\mathbf{r}$, has two main contributions: the van der Waals interaction U_{vdW} present in all species and the dipole–dipole interaction U_{dd}. Neglecting higher-order interactions (e.g. dipole–quadrupole or quadrupole–quadrupole), the interaction potential can be written as

$$U(\mathbf{r}) = U_{\mathrm{vdW}}(\mathbf{r}) + U_{\mathrm{dd}}(\mathbf{r})\,. \tag{2.10}$$

The dipole–dipole interaction (DDI) is intrinsically anisotropic as the sign and strength of the interaction depend on the orientation of the dipoles, and it is this anisotropy which gives rise to the complex physics of dipolar quantum systems. The van der Waals interaction is isotropic for simple atoms like alkalis, but it is anisotropic for atoms with complex electronic structures, including for lanthanides.

2.2.1 Van der Waals interaction

The attractive interaction of two atoms, known as the van der Waals force, is caused by mutually-induced electric multipoles and can be described by the potential

$$U(r) = -\frac{C_6}{r^6} - \frac{C_8}{r^8} - \frac{C_{10}}{r^{10}} - \cdots\,, \tag{2.11}$$

where r is the interatomic distance and the C_i are called dispersion coefficients. The first term is due to the induced dipole–dipole interaction. As other terms represent higher-order contributions, usually only the first term is retained to approximate the interaction, which is therefore described by the potential

$$U_{\mathrm{vdW}} = -\frac{C_6}{r^6}\,. \tag{2.12}$$

[8] The Bohr radius is defined as $a_0 = \hbar/m_e c\alpha \approx 5.29 \times 10^{-11}\,\mathrm{m}$.

This is the attractive part of the Lennard-Jones potential and is also called London dispersion.[9] The potential is attractive and is usually spherically symmetric. It is short-ranged as it drops off as r^{-6}, and is due to the interaction of the electrons of the two interacting atoms. We will show in §2.2.4 that the strength of the van der Waals interaction (the 'contact interaction') can be effectively controlled by changing the strength of an external magnetic field, via utilising Feshbach resonances.

Due to the complex and non-spherical electron structure of the lanthanides, their electrons interact in an anisotropic way. Therefore, in their case, the constant van der Waals coefficient C_6 needs to be replaced by a matrix of coefficients, C_6, representing the coupling between different pairs of total angular momentum eigenstates. Calculating these coefficients is difficult, although it has been done for some elements, including erbium. For this element, the isotropic part[10] of C_6 is $C_6 = 1723$ a.u. and the mean anisotropic component is $\Delta C_6 = 350$ a.u. [134]. While this anisotropy has profound consequences on the spectrum of Feshbach resonances, it can be shown that for ultracold collisions it can still be described by a single number, the unique and isotropic s-wave scattering length [135]. This is due to the short-range character of the potential and will be discussed in §2.2.3 in more detail.

2.2.2 Dipole–dipole interaction

The dipole–dipole interaction (DDI) is the interaction between two magnetic (or electric) dipoles, and is intrinsically long-range and anisotropic. As the contact interaction, usually encountered in ultracold systems, is isotropic and short-range, the introduction of the DDI greatly enriches the variety of physical phenomena in ultracold systems and hence gained significant interest in the last few years.

The interaction potential U_{dd} between two dipoles, pointing in directions $\mathbf{e}_1$ and $\mathbf{e}_2$ and separated by a distance $\mathbf{r}$, is given by

$$U_{\mathrm{dd}}(\mathbf{r}) = \frac{C_{\mathrm{dd}}}{4\pi} \frac{(\mathbf{e}_1 \cdot \mathbf{e}_2) r^2 - 3(\mathbf{e}_1 \cdot \mathbf{r})(\mathbf{e}_2 \cdot \mathbf{r})}{r^5}, \tag{2.13}$$

where C_{dd} is the dipolar coupling constant and $r = |\mathbf{r}|$. The coupling constant is $C_{\mathrm{dd}} = \mu_0 \mu_1 \mu_2$ for particles with a magnetic dipole moment μ (e.g. lanthanides) and $C_{\mathrm{dd}} = d_1 d_2 / \varepsilon_0$ for particles with an electric dipole moment d (e.g. heteronuclear molecules), where ε_0 is the permittivity of free space.

[9] The repulsive part becomes significant only at very small distances and is not considered here.

[10] Here a.u. is the 'atomic unit' within the Hartree atomic units, where lengths are expressed in terms of a_0 and energies are expressed in terms of hartrees (Ha), 1 Ha $= m_e c^2 \alpha^2$. Therefore, 1 a.u. $= a_0^6 m_e c^2 \alpha^2$.

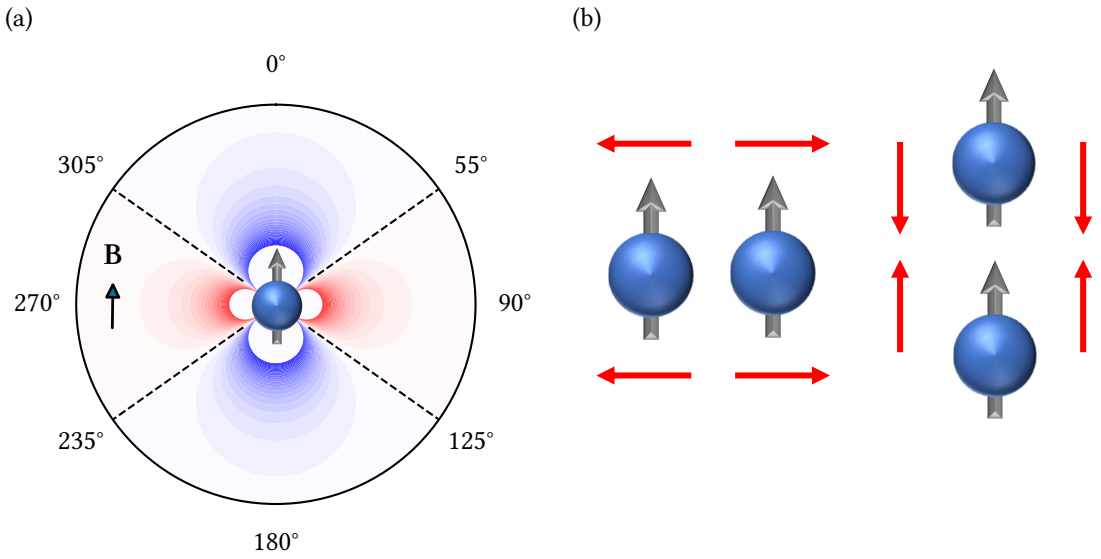

Figure 2.2. Dipole–dipole interaction. (a) The anisotropic profile of the interaction potential for a pair of parallel dipoles. The potential changes sign at the 'magic angle', $\theta_m \approx 54.7°$ (dashed lines). (b) Side-by-side dipoles ($\theta = \pi/2$) repel each other, head-to-tail dipoles ($\theta = 0$) attract each other.

In the presence of an external magnetic field, (magnetic) dipoles align themselves along the direction of the field (which we call z) and the atomic cloud becomes polarised. For two identical ($\mu_1 = \mu_2 = \mu$) and aligned dipoles, the interaction simplifies to

$$U_{\mathrm{dd}}(\mathbf{r}) = \frac{\mu_0}{4\pi}\frac{\mu^2 - 3(\hat{\mathbf{r}}\cdot\boldsymbol{\mu})^2}{r^3} = \frac{\mu_0\mu^2}{4\pi}\frac{1-3\cos^2\theta}{r^3}, \tag{2.14}$$

where $\boldsymbol{\mu}$ is the magnetic dipole moment (vector) of the (magnetically polarised) atoms, $\hat{\mathbf{r}} = \mathbf{r}/r$ and θ is the angle between the orientation of the dipoles and the interatomic axis (i.e. z and $\mathbf{r}$).

As θ changes from 0 to $\pi/2$, the interaction changes its sign, and at the so-called 'magic angle' $\theta_m = \cos^{-1}(1/\sqrt{3}) \approx 54.7°$, the DDI vanishes. For angles smaller than θ_m, the potential is negative and the dipoles attract each other, whereas above θ_m the potential is positive, and the dipoles repel each other. As the term $1-3\cos^2\theta$ ranges from -2 to 1, the maximum attraction is twice as strong as the maximum repulsion. The $\theta = 0$ configuration is called the head-to-tail configuration and the $\theta = \pi/2$ one is called the side-by-side configuration. It is interesting to note that the angular dependency of the DDI is the second-order Legendre polynomial $P_2(\cos\theta)$ and the interaction is therefore d-wave symmetric (in the nomenclature of partial wave decomposition). The anisotropy of the DDI is illustrated in Fig. 2.2.

In our system of magnetic dipoles, the intrinsic strength of the DDI cannot be tuned,

as it depends on the (fixed) magnetic dipole moment of the atoms.[11] However, one can reduce the effective strength of the DDI by rotating the dipoles rapidly [136]. Using a fast rotating external magnetic field which the dipoles are still able to follow adiabatically, from the DDI's perspective, the dipoles will be effectively aligned along the rotation axis but will only contribute with the relevant projection. This regime applies if the frequency of the rotating field Ω is much smaller than the Larmor frequency $\omega_L = \mu B/\hbar$ but is much higher than the typical frequencies of atomic motion, determined by the frequencies of the trapping potential. In this limit, the atoms experience a time-average of the DDI (cf. Eq. (2.14)), given by [65]

$$U_{\mathrm{dd}}(\mathbf{r}) = \frac{\mu_0 \mu^2}{4\pi} \frac{1 - 3\cos^2\theta'}{r^3} \frac{3\cos^2\beta - 1}{2}, \tag{2.15}$$

where θ' is the angle between the rotation and the interatomic axes, and β is the angle between the rotation axis and the magnetic field. The averaged interaction potential differs from the potential of aligned dipoles by a factor of $(3\cos^2\beta - 1)/2$. By varying β, this factor can be changed continuously from 1 to $-1/2$, thus allowing to invert the sign of the DDI or to cancel it completely when $\beta = \theta_m$, the magic angle.

The DDI is often presented in contrast to contact interactions, adding greatly to the wealth of phenomena that can be explored, as it is not only anisotropic but also long-range. However, it is worth noting that there is no unique definition for classifying interactions as short- or long-range. One way to classify the range of interactions is to check whether the chemical potential μ_c depends only on the density n, or also on the total number of particles N [137]. If $\mu_c = \mu_c(n)$, the interaction is short-range, if $\mu_c = \mu_c(N, n)$, the interaction is long-range. Following this definition, it can be shown that an interaction is short-ranged if it decays faster than r^{-D} at large distances in D dimensions. Therefore, the r^{-3} potential of the DDI is long-ranged in three dimensions (3D), but it is short-ranged in 1D and 2D. The van der Waals interaction (r^{-6}) is short-ranged in both 1D, 2D and 3D. We will see in §2.2.3 that the r^{-3} nature of the DDI leads to scattering properties significantly different from other, faster-decaying interactions.

2.2.3 Scattering

The variety of ultracold dipolar physics has its roots in the anisotropic, long-range interaction between atoms. This manifests itself in the relatively complex scattering behaviour of dipolar atoms, and it is this scattering behaviour which is directly responsible

[11] The situation is different e.g. for heteronuclear molecules, where the induced electric dipole moment depends on the strength of the external electric field.

for the diversity and complexity of physical phenomena. We will therefore now turn to reviewing the scattering properties of dipolar atoms. First, we present an overview of scattering theory and then examine how it applies to dipolar atoms. However, this is a complex topic and a much more in-depth discussion can be found in Refs. 102, 103, 133, 138 and 139.

We will only consider elastic collisions, in which the total kinetic energy is conserved, and we shall therefore assume that after collision the internal states of the particles are left unchanged. However, it is worth noting that the dipolar interaction does allow non-elastic, spin-changing collisions [140] and specialised cooling schemes have been realised based on this effect [141–143]. However, if the Zeeman splitting is large compared to the thermal energy and the atoms are in the lowest-energy spin state, such spin-flips are suppressed as the other spin states correspond to higher, kinetically inaccessible energies.

Let us consider the scattering of two particles with mass m_1 and m_2, interacting via a potential $U(\mathbf{r}_1, \mathbf{r}_2)$. For simplicity, we investigate collisions in the centre-of-mass frame, in which the problem of elastic collision amounts to the scattering of a single particle, with reduced mass $m_r = m_1 m_2/(m_1 + m_2)$ and relative momentum $\mathbf{p} = \hbar\mathbf{k} = \hbar(\mathbf{k}_1 - \mathbf{k}_2)$, off a scattering centre at the origin with an interaction potential $U(\mathbf{r}_1 - \mathbf{r}_2) = U(\mathbf{r})$ [103]. In this system, the wave function for the centre-of-mass motion is a plane wave, and that of the relative motion needs to satisfy the Schrödinger equation with the Hamiltonian $H = -\hbar^2\nabla^2/2m_r + U(\mathbf{r})$.

First of all, let us suppose that $U(\mathbf{r})$ either vanishes or is otherwise negligible outside some finite region (the 'interaction region'). Interaction with such a localised potential brings the simplification that outside the range of the potential, the energy eigenstates of the system take the form of plane waves, so a free particle can be represented as such there. Therefore, one can write the wave function for the relative motion as the sum of an incoming plane wave and a scattered wave,

$$\psi(\mathbf{r}) = e^{i\mathbf{k}\cdot\mathbf{r}} + \psi_{\mathrm{sc}}(\mathbf{r})\,. \tag{2.16}$$

The form of $\psi_{\mathrm{sc}}(\mathbf{r})$ is fixed by the idea that outside the interaction region an eigenstate should coincide with a solution of the free-particle Schrödinger equation (i.e. a plane wave), only deviating from such a solution inside. Therefore, one can write the scattered wave as an outgoing spherical wave at large distances, and the wave function for large r becomes

$$\psi(r) = e^{ikz} + f(\theta, \varphi, k)\frac{e^{ikr}}{r}\,, \tag{2.17}$$

where (r, θ, φ) are spherical coordinates and we chose z to point in the direction of $\mathbf{k}$. The function $f(\theta, \varphi, k)$ is called the scattering amplitude and it generally depends on

direction and energy. It is worth noting that although this form is actually not an exact solution to the free-particle Schrödinger equation (as it is not a plane wave), its use is justified as it is an asymptotic one.

The goal of scattering theory is to determine how particles are scattered (e.g. how often and what their angular distribution is). The physical significance of the scattering amplitude can be understood by considering the probability current of the scattered particle.[12] One can define the differential cross section $\mathrm{d}\sigma$, the ratio of the probability per unit time that the scattered particle will pass through a surface element $r^2\,\mathrm{d}\Omega$ to the probability current per unit area of the incoming wave. Here, $\mathrm{d}\Omega$ is a small element of solid angle. The probability current density of the incoming wave for the wave function in Eq. (2.17) is $\hbar k/m_r$, while the probability per unit time that the scattered particle will pass through the surface element $r^2\,\mathrm{d}\Omega$ is $(\hbar k/m_r)|f(\theta,\varphi,k)|^2\,\mathrm{d}\Omega$ for large distances r. Therefore, the differential cross section is

$$\mathrm{d}\sigma = |f(\theta,\varphi,k)|^2\,\mathrm{d}\Omega\,. \tag{2.18}$$

It is worth noting that the differential cross section is a function of θ and φ, although those arguments are normally not explicitly written. Furthermore, some define the differential cross section as $\mathrm{d}\sigma/\mathrm{d}\Omega$, but this notation is misleading: the differential cross section is not the derivative of anything; $\mathrm{d}\sigma$ is simply the probability per unit time for a scattered particle to pass through a solid angle element $\mathrm{d}\Omega$ at coordinates (θ,φ), normalised to the incoming probability flux.

Equation (2.18) is only valid for the scattering of non-identical particles, e.g. atoms occupying different states. For indistinguishable particles, the spatial part of the wave function must be symmetrised for a boson and antisymmetrised for a fermion (see Fig. 2.3 for an intuitive plot). Therefore, for identical particles, the differential cross section takes the form

$$\mathrm{d}\sigma = \left|\frac{f(\mathbf{k},\mathbf{k}') \pm f(\mathbf{k},-\mathbf{k}')}{\sqrt{2}}\right|^2 \mathrm{d}\Omega\,, \tag{2.19}$$

where the positive sign applies to bosons and the negative sign applies to fermions, and $\mathbf{k}'$ is the relative wave vector after collision ($\mathbf{k}$ and $\mathbf{k}'$ define θ and φ).

If the scattering field can be treated as a mere perturbation to the free particle, the scattering amplitude can be calculated as

$$f(\theta,\varphi,k) = -\frac{m_r}{2\pi\hbar^2}\int U(\mathbf{r}')e^{-i\mathbf{q}\cdot\mathbf{r}'}\,\mathrm{d}^3\mathbf{r}', \tag{2.20}$$

[12] The probability current associated with the wave function $\psi(\mathbf{r})$ is defined as $\mathbf{j} = (\psi^*\hat{\mathbf{p}}\psi - \psi\hat{\mathbf{p}}\psi^*)/2m_r$, where $\hat{\mathbf{p}} = -i\hbar\nabla$ is the momentum operator.

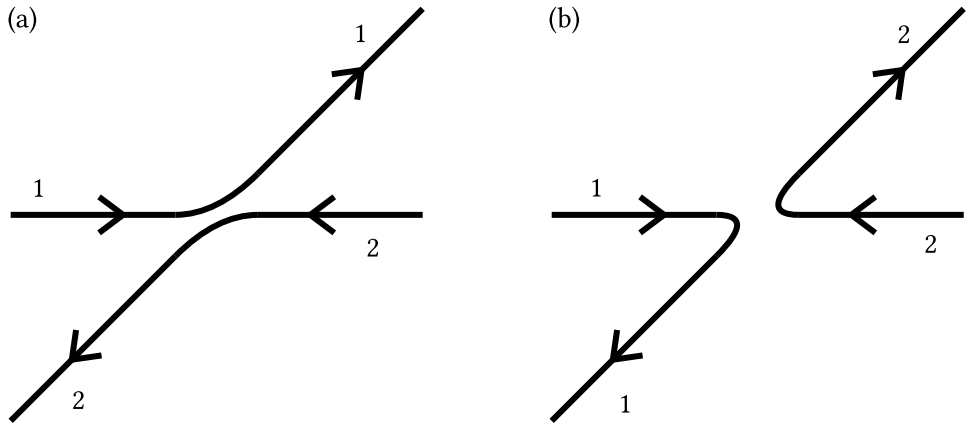

Figure 2.3. Indistinguishable scattering processes. The two processes involving two identical particles yield the same final state (up to parity), so their differential cross section has to be symmetrised accordingly.

where $\mathbf{q} = \mathbf{k}' - \mathbf{k}$ is the momentum transfer in the collision. This is called the (first) Born approximation. However, in some special cases, more accurate results can be obtained.

Let us therefore assume that the interaction between the atoms is spherically symmetric with range r_0. In this case, the solution of the Schrödinger equation has axial symmetry with respect to the direction of the incident particle, and the scattering amplitude depends on direction only through the scattering angle θ. To take advantage of the axial symmetry (and so that the Hamiltonian conserves L^2 and L_z), let us expand the wave function for the relative motion using the spherical harmonics $Y_\ell^m(\theta, \varphi)$. As here $L_z = 0$ and is conserved, only the $Y_\ell^0(\theta, \varphi) \propto P_\ell(\cos\theta)$ terms contribute and we can write ψ in terms of the Legendre polynomials $P_\ell(\cos\theta)$:

$$\psi = \sum_{\ell=0}^{\infty} P_\ell(\cos\theta) \frac{R_{k\ell}(r)}{kr}\,, \tag{2.21}$$

where the radial functions $R_{k\ell}$ satisfy the 1D Schrödinger equation

$$-\frac{\hbar^2}{2m_r}\frac{\mathrm{d}^2 R_{k\ell}}{\mathrm{d}r^2} + \left(U(r) + \frac{\hbar^2}{2m_r}\frac{\ell(\ell+1)}{r^2}\right) R_{k\ell} = E R_{k\ell}\,. \tag{2.22}$$

This is called partial wave analysis, where the $\ell = 0, 1, 2, \ldots$ terms are called $s, p, d, \ldots$ waves, i.e. the same nomenclature is used as in spectroscopy. The benefit of this expansion is that the partial waves, each carrying a definite angular momentum, scatter independently of each other (they are independent 'scattering channels').

For large distances outside the range of the potential, $r \gg r_0$, one can neglect the interaction $U(r)$ and the second (the so-called centrifugal) term in Eq. (2.22). Then, given

we effectively have a free particle, the solution coincides with a free wave up to a phase shift. The solution takes the general form

$$R_{k\ell}(r) = A_\ell \sin\left(kr - \ell\frac{\pi}{2} + \delta_\ell\right), \tag{2.23}$$

where δ_ℓ is the phase shift (see Ref. 144 for a good illustration). To find the phase shifts δ_ℓ, one must enforce that Eqs. (2.17) and (2.21) give the same result, and so in general δ_ℓ depends on the form of $U(r)$. By expanding Eq. (2.17) and ensuring consistency, one finds $A_\ell = (2\ell+1)i^\ell e^{i\delta_\ell}$ and

$$f(\theta,k) = \frac{1}{2ik}\sum_{\ell=0}^{\infty}(2\ell+1)\left(e^{2i\delta_\ell}-1\right)P_\ell(\cos\theta)\,. \tag{2.24}$$

The total scattering cross section is obtained by integrating the differential cross section over all solid angles, and it is given by

$$\sigma(k) = \int \mathrm{d}\sigma = 2\pi\int_0^\pi |f(\theta,k)|^2 \sin\theta\,\mathrm{d}\theta\,. \tag{2.25}$$

Substituting Eq. (2.24) into this, one finds

$$\sigma = \frac{4\pi}{k^2}\sum_{\ell=0}^{\infty}(2\ell+1)\sin^2(\delta_\ell)\,. \tag{2.26}$$

It is worth noting at this point the centrality of the phase shifts for deducing the physical behaviour. Once phase shifts are determined, all scattering properties can be inferred.

One should not forget that this result only holds for non-identical particles. For identical particles, due to the symmetry of the states, we need to use the symmetrised $\mathrm{d}\sigma$ (cf. Eq. (2.19)). Using that, one can observe that only even (odd) values of ℓ contribute to the differential cross section for bosons (fermions) due to the symmetry requirements. Furthermore, the factor $1/\sqrt{2}$ avoids double counting by effectively integrating over only half the total solid angle 4π, and so the cross section can be written

$$\sigma_{\text{ident.}} = \frac{8\pi}{k^2}\sum_{\substack{\ell\geq 0 \text{ even}\\ (\ell\geq 1 \text{ odd})}}^{\infty}(2\ell+1)\sin^2(\delta_\ell)\,. \tag{2.27}$$

We will now concentrate on the low-energy limit,

$$kr_0 \ll 1 \quad \text{and} \quad r \ll \frac{1}{k}, \tag{2.28}$$

relevant for ultracold-atom experiments.[13] In this limit, the energy of the relative particle is negligible and one can set $E = 0$. In the case of $\ell = 0$, Eq. (2.22) takes the form

$$\frac{\mathrm{d}^2 R_{k0}}{\mathrm{d}r^2} - \frac{2m_r}{\hbar^2} U(r) R_{k0} = 0\,. \tag{2.29}$$

The solution of Eq. (2.29) needs to match the asymptotic form Eq. (2.23) in the region

$$r_0 \ll r \ll \frac{1}{k}\,, \tag{2.30}$$

where both equations are applicable. In this region, one can also neglect the interaction term in Eq. (2.29) and the solution takes the linear form

$$R_{k0} = c_0(1 - \kappa r)\,, \tag{2.31}$$

where the value of κ depends on the choice of the potential $U(r)$. On the other hand, for $kr \ll 1$, the asymptotic solution can also be expanded, yielding the same form as Eq. (2.31) if we set $c_0 = e^{i\delta_0} \sin\delta_0$ and

$$\kappa = -k \cot\delta_0\,. \tag{2.32}$$

Therefore, for small k,

$$\delta_0 = -\frac{k}{\kappa}\,. \tag{2.33}$$

A similar investigation shows that the phase shifts at higher values of ℓ behave like $\delta_\ell \propto k^{2\ell+1}$ [139], and so at low temperatures the $\ell = 0$ term (the s-wave) dominates the scattering amplitude which can be written as

$$f(\theta, k) = \frac{\delta_0}{k} = -\frac{1}{\kappa} \equiv -a_s\,. \tag{2.34}$$

Therefore, at low temperatures the scattering amplitude is a constant, independent of angles or energies: $\lim_{k\to 0} f(\theta, k) = -1/\kappa$. In fact, $a_s = 1/\kappa$ is called the s-wave scattering length and is a very important parameter for ultracold-atom experiments, defining the (s-wave) interaction strength. Per Eq. (2.26), the scattering cross section for non-identical

[13] This is the limit where the energy E of the incoming particle is much smaller than the interaction $U(r)$ and the centrifugal term. The characteristic energy of the interaction can be estimated using Heisenberg's uncertainty principle. If we say that within the potential the uncertainty of the atom's position is $\Delta x \approx r_0$, then the uncertainty of its momentum is $\Delta p \approx \hbar/\Delta x = \hbar/r_0$. Therefore, the characteristic energy associated with the potential is $\Delta p^2/2m_r = \hbar^2/2r_0^2 m_r$. This energy needs to be much higher than the energy of the incoming particle, which yields the limit $kr_0 \ll 1$. On the other hand, comparing the centrifugal factor $\ell(\ell+1)/r^2$ with $2m_r E/\hbar^2$, one can deduce the other criterion, $r \ll 1/k$.

particles can therefore be written as

$$\sigma = 4\pi a_s^2 \,, \tag{2.35}$$

whereas per Eq. (2.27) for identical bosons it is

$$\sigma_B = 8\pi a_s^2 \tag{2.36}$$

and for identical fermions it is

$$\sigma_F = 0 \,. \tag{2.37}$$

The latter is a striking result—the scattering cross section for identical fermions is zero. This means that if identical fermions only interact via contact interactions, they cannot be cooled via evaporation which intrinsically relies on scattering for the thermalisation of the gas in the process.

Classically, the dominance of the s-wave term can be understood by examining the potential term in Eq. (2.22). Specifically, the second term in the effective potential presents a so-called 'centrifugal barrier', and a relative particle with an energy much smaller than the barrier is unable to reach small r as it is reflected on the barrier. Therefore, for $\ell > 0$, the short-range potential $U(r)$ inside the barrier has no effect. The case of $\ell = 0$ is special as there is no barrier then. One thus expects that the partial scattering amplitudes (or equivalently the phase shifts $\delta_\ell(k)$) will vanish as $k \to 0$ for all partial waves with $\ell > 0$. However, for certain potentials, this barrier can also lead to quasi-bound states in the potential well. If the incident relative particle has an energy close to the energy of such a quasi-bound state, a strong scattering resonance can be observed.[14] These so-called 'shape resonances' (depending on the shape of the scattering potential) may strongly enhance the contribution of $\ell > 0$ partial waves in an energy domain where one would naïvely expect pure $\ell = 0$ (s-wave) scattering [103, 139].

While our discussions have been relatively abstract so far, not fixing the form of the interaction potential, we reached the important result that the exact form of the potential does not actually matter—only the associated scattering length a_s. In other words, although the details of the interaction potential determine the scattering length, it is only a_s that matters in the end. This also simplifies the experimental effort: instead of having to probe the interaction potential at different inter-particle distances, it is sufficient to measure only the scattering length. Calculating a_s for an arbitrary potential is difficult in general, but can be done e.g. for a square well [103, 139]. It is worth noting

[14] This can be phenomenologically explained by the scattering particle spending an extended amount of time in the quasi-bound state, leading to an enhanced interaction.

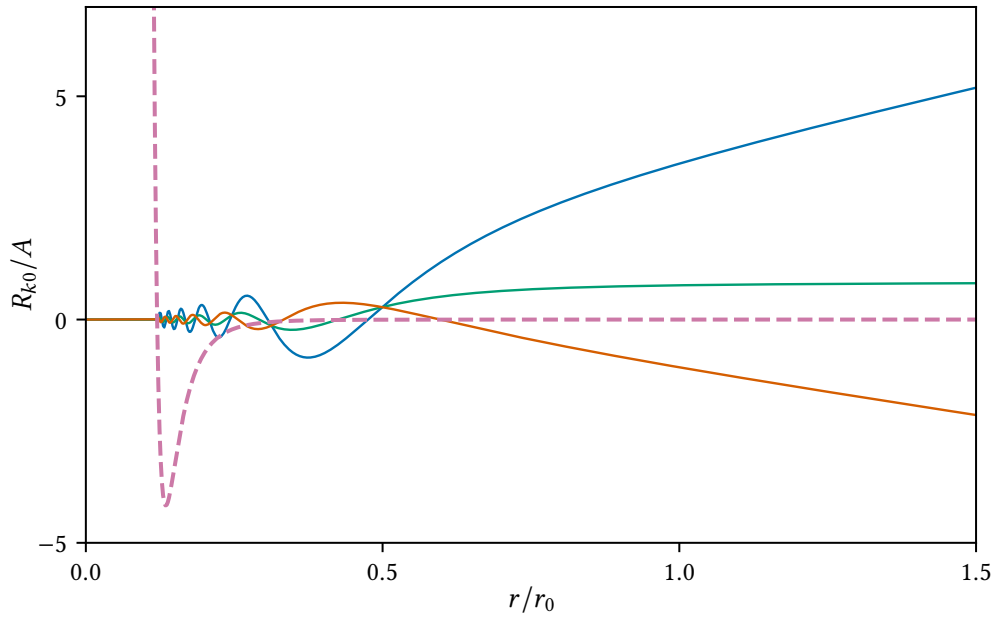

Figure 2.4. Radial part of the wave function. Solutions of Eq. (2.29) for a truncated van der Waals potential $-C_6/r^6$ (solid lines; up to a scaling factor A) [133]. The repulsive part of the Lennard-Jones potential is approximated with a hard core of radius r_c (the full potential is shown with a dashed line). The solutions shown are for $r_c = 0.118r_0$ (orange), $r_c = 0.121r_0$ (green) and $r_c = 0.122r_0$ (blue), where $r_0 = \sqrt[4]{2m_rC_6/\hbar^2}$ is the characteristic range of the potential. We see the solutions tend to a line above r_0 and are very sensitive to the size of the core; the corresponding scattering lengths (proportional to the negative inverse slope of the solutions for $r \gg r_0$) are $a_s = 0.4r_0$, $-33r_0$ and $-0.3r_0$, respectively.

that the scattering length can be positive and negative, and can be much larger than the interaction range, tending to infinity at the development of a new bound state of the scattering potential. A schematic behaviour of the solutions of Eq. (2.29) is shown in Fig. 2.4 for positive and negative values of the scattering length, respectively.

Remember that according to Eq. (2.21), the wave function at low energies can be written as

$$\psi = \frac{c_0}{k}\left(\frac{1}{r} - \frac{1}{a_s}\right). \tag{2.38}$$

In fact, this is an exact solution of the Schrödinger equation

$$\left(-\frac{\hbar^2}{2m_r}\nabla^2 + U(\mathbf{r})\right)\psi = 0 \tag{2.39}$$

if we assume the potential is

$$U(\mathbf{r}) = U_{\text{ps}}(\mathbf{r}) = \frac{2\pi\hbar^2 a_s}{m_r}\delta(\mathbf{r})\frac{\partial}{\partial r}r\,, \tag{2.40}$$

where $\delta(\mathbf{r})$ is the Dirac delta function.[15] U_{ps} is called a pseudo-potential—a stand-in potential which is much easier to work with, but which still gives exactly the same results as a more complicated $U(\mathbf{r})$ with the same scattering length. In fact, given only a_s is relevant for practical purposes, we can simplify this further by choosing a potential which gives a_s as its scattering length in the Born approximation at the zero energy limit. Such a potential is $g_s\delta(\mathbf{r}) = 2\pi\hbar^2 a_s\delta(\mathbf{r})/m_r$ which is called the zero-range contact interaction potential. It is worth noting that while the scattering cross section depends on a_s^2, the potential depends directly on a_s. This will be important later on when we consider Bose–Einstein condensates.

Scattering of dipolar atoms

Analysing the scattering of real atoms is much more complicated, as their interactions are neither necessarily finite-range nor isotropic. However, some of the analysis that we presented so far still remains valid [145]. Specifically, the expansion in Eq. (2.21) and the associated partial wave analysis is still possible.

We have previously discussed that for a finite-range potential, the phase shifts are $\delta_\ell \propto k^{2\ell+1}$ for small k. In our case, both the van der Waals interaction and the dipole–dipole interaction scale like a power law r^{-n} at large distances, so in the strict sense, these are not finite-range.[16] For such potentials, $\delta_\ell \propto k^{2\ell+1}$ only holds for $\ell < (n-3)/2$, and for higher-order partial waves $\delta_\ell \propto k^{n-2}$ [138]. This means that in the case of the van der Waals potential ($n \geq 6$), all phase shifts become negligible compared to δ_0 for ultracold temperatures ($k \to 0$). Therefore, the cross section is still dominated by the $\ell = 0$ term (s-wave scattering) and we can still describe scattering with a unique scattering length a_s, even if the van der Waals interaction is not isotropic.

On the other hand, for the DDI we have $n = 3$, meaning that all phase shifts scale as $\delta_\ell \propto k$, and so all partial waves need to be taken into account. Furthermore, due to its anisotropic nature, this interaction does not conserve L_z and so it mixes the scattering channels (the partial waves), and a simple description is no longer possible. Importantly, while for isotropic interactions the total scattering cross section does not depend on the direction of the incident relative particle, for the dipolar interaction it does. However, in the ultracold regime, the average scattering cross section (averaged across all incident directions) still takes a constant value, i.e. it does not depend on the collision energy [146].

[15] One can easily check this using the fact $\nabla^2(1/r) = -4\pi\delta(\mathbf{r})$.

[16] For the van der Waals interaction, a characteristic length scale can be derived by solving $\Delta p^2/2m_r = C_6/a_{\text{vdW}}^6$, where $\Delta p = \hbar/a_{\text{vdW}}$ as per the uncertainty principle, yielding $a_{\text{vdW}} = \sqrt[4]{2m_r C_6/\hbar^2}$. For ^{166}Er, $a_{\text{vdW}} \approx 150a_0$.

For dipolar scattering, this regime is reached when the collision energy is smaller than the natural energy scale of the dipolar interaction,

$$E_{\mathrm{dd}} = \frac{\mu_0 \mu^2}{4\pi a_{\mathrm{dd}}^3}, \tag{2.41}$$

where the natural length scale a_{dd} for two colliding (identical) dipoles is defined as[17]

$$a_{\mathrm{dd}} = \frac{m\mu_0 \mu^2}{12\pi\hbar^2}, \tag{2.42}$$

which is sometimes called the 'dipolar length'. The numerical prefactor in a_{dd} is chosen such that a three-dimensional homogeneous condensate becomes unstable to local density perturbations for $a_s \leq a_{\mathrm{dd}}$ [66]. For ^{166}Er, $a_{\mathrm{dd}} \approx 66 a_0$.

In the Born approximation (i.e. where the potential is a perturbation) and at small energies, the angle-dependence can be described exactly [147], and it can be shown that the average cross section arising from even partial waves is

$$\sigma_e = \frac{4\pi}{5} a_{\mathrm{dd}}^2, \tag{2.43}$$

and from odd partial waves it is

$$\sigma_o = \frac{12\pi}{5} a_{\mathrm{dd}}^2. \tag{2.44}$$

Therefore, the total average cross section for bosons is

$$\sigma_B = \frac{8\pi}{5} a_{\mathrm{dd}}^2 + 8\pi a_s^2, \tag{2.45}$$

and for fermions it is

$$\sigma_F = \frac{24\pi}{5} a_{\mathrm{dd}}^2. \tag{2.46}$$

This means that dipolar fermions have a finite cross section even at very low energies (in fact three times as large as purely dipolar bosons), and can therefore be cooled evaporatively [148]. It is also worth noting that dipolar scattering has also been studied beyond the Born approximation, indicating a small, temperature-dependent enhancement of the DDI [149].

Even with the dipolar interaction taken into account, the pseudo-potential formalism can also still be employed [145, 150–152]. Interestingly, besides the long-range contribution, the potential U_{dd} also generates a contribution to the (short-range) s-wave scattering

[17] This expression can be derived (up to numerical prefactors) by solving $\Delta p^2/2m_r = E_{\mathrm{dd}}$, where $\Delta p = \hbar/a_{\mathrm{dd}}$ as per the uncertainty principle.

length due to the coupling between different scattering channels [153–156]. As a result, changing the strength of the DDI also results in a change in the short-range part of interparticle interaction. Therefore, for a gas of identical, polarised dipolar atoms, the full pseudo-potential can be written as

$$U_{\mathrm{ps}}(\mathbf{r}) = \frac{4\pi\hbar^2}{m}\left(a_s^{\mathrm{eff}}\delta(\mathbf{r}) + \frac{3}{4\pi}a_{\mathrm{dd}}\frac{1-3\cos^2\theta}{r^3}\right), \tag{2.47}$$

where a_s^{eff} is the effective s-wave scattering length and we used the fact that for identical particles $m_r = m/2$, so $g_s = 4\pi\hbar^2 a_s^{\mathrm{eff}}/m$.

Given the choice of prefactors in a_{dd}, it can also be used to compare the relative strength of the contact and dipole–dipole interactions, by defining the relative dipole strength[18]

$$\varepsilon_{\mathrm{dd}} = \frac{a_{\mathrm{dd}}}{a_s}\,. \tag{2.48}$$

2.2.4 Feshbach resonances

We have already seen that when the energy of a bound state is resonant with the energy of the scattered particle (e.g. due to a shape resonance), the scattering cross section greatly increases. While shape resonances are induced by quasi-bound states in the collision channel due to the shape of the scattering potential itself, a resonance can also be induced by a bound state of another collision channel. This is called a Feshbach resonance [58]. This phenomenon has been exploited widely, due to the fact that the energy of the bound state and hence the scattering length can be tuned using an external magnetic field as we will shortly see. This means that the strength and sign of the contact interaction can be controlled relatively easily in experiments.

First theoretical investigations regarding resonances between scattering channels took place in the 1930s [157–159], with further work by Feshbach [160, 161] and Fano [162], and it is after them that the effect is named Fano–Feshbach resonance, or Feshbach resonance (FR) for short. Interestingly, Feshbach himself considered his name being attached to a well-known resonance phenomenon as mere atomic physics jargon [163, 164]. Fano's name is usually associated with the asymmetric line shape of such a resonance, called a 'Fano profile'.

Feshbach resonances are not exclusive to ultracold atoms—in fact, Feshbach himself thought about them in the context of nuclear physics, and they also appear in molecular

[18] From here on, we drop the 'eff' superscript from a_s^{eff} for readability and take a_s to include contributions from all kinds of interactions.

and chemical physics [58]. As we will see later, the resonance provides access to a weakly-bound (molecular) state of the atoms, leading to the formation of ultracold molecules, commonly referred to as Feshbach molecules [58]. Such molecules can then be transferred into a more deeply bound state via adiabatic transfer schemes [165]. FRs were first detected experimentally in the 1970s [166, 167], and gained traction after the demonstration of the tunability of interactions in a BEC in 1998 [56]. This was an important milestone for ultracold-atom research, and exploiting Feshbach resonances became a well-established and effective tool to tune the scattering length over a wide range, from strongly attractive via non-interacting to strongly repulsive regimes. While for single-valence-electron atoms (alkalis) methods to calculate a_s are well-established, predictive scattering models for multi-electron atoms are more difficult to develop as more approximations are needed. In the case of complex electronic structures, such as for lanthanides, the required treatment is highly non-trivial, involving many atomic potentials and complex interactions [65, 134]. While FRs offer great flexibility, ultracold-atom experiments usually try to avoid ramping the magnetic field very close to (or through) a resonance, as the large scattering length also greatly increases the rate of three-body collisions, which lead to atom loss from the trap.[19] This is discussed in detail for the case of erbium in Chapter 5.

We will now turn to the phenomenological explanation of Feshbach resonances, based on comprehensive reviews for dipolar [65] and non-dipolar atoms [58]. Let us assume that the atoms collide with energy E in a certain interaction channel. This channel is referred to as an open channel, as it is energetically possible. Let us also assume that a different interaction channel also exists (e.g. the atoms colliding in a different spin state), which is energetically not allowed for our collision energy E: this is referred to as a closed channel. If the interaction potential of the closed channel supports a bound state and a coupling exists between the open and closed channels, the atoms can transfer to this bound state and a resonance can occur—this is called a Feshbach resonance. The scenario is visualised in Fig. 2.5. The existence of a (quasi-)bound state close to the energy of the colliding atoms greatly affects the scattering cross section. One can think of this in the context of the atoms forming a virtual bound state temporarily, and the closer the energy of the atoms to this bound state is, the more this state affects the interaction properties. While in shape resonances the quasi-bound state is induced in the open channel by the shape of the interaction potential itself, Feshbach resonances occur due to the coupling of the colliding atoms to a bound state in a different interaction channel.

The coupling between the channels can be e.g. hyperfine in nature, and for dipolar atoms, the anisotropic interaction itself provides a coupling. (Remember that we observed

[19] It can be shown that the three-body loss coefficient L_3 scales as $L_3 \propto a_s^4$ close to a Feshbach resonance [168].

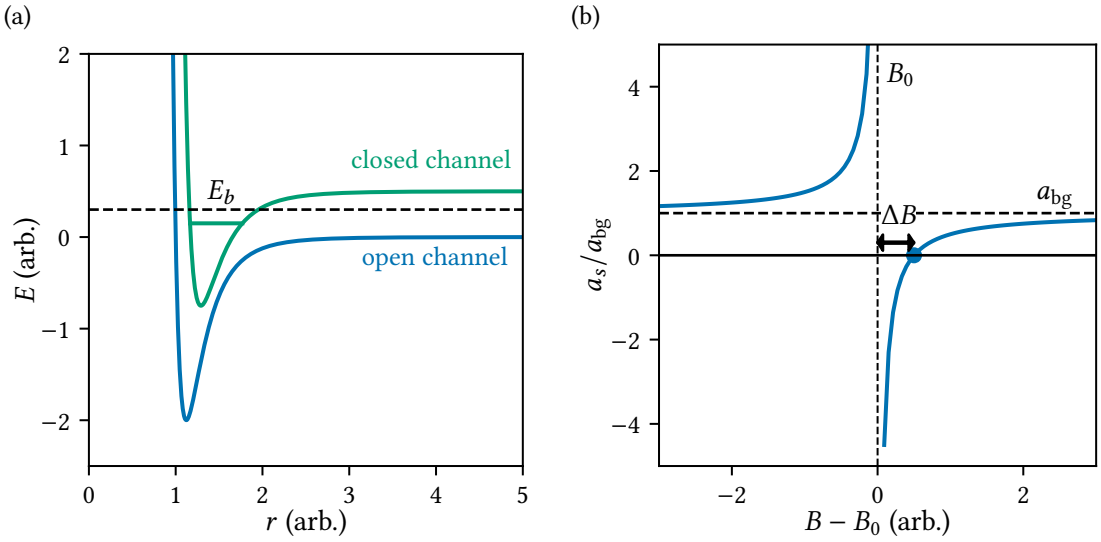

Figure 2.5. Mechanism of Feshbach resonances. (a) Interaction energy in the open (entrance) and the closed collision channels (blue and green lines, respectively), and the energy of the colliding particles (dashed line). In a magnetic Feshbach resonance, the relative energy $E_b(B)$ of a molecular bound state can be changed by tuning the magnetic field. When E_b is close to the energy of the colliding particles, the scattering length diverges. (b) Behaviour of the scattering length a_s near a Feshbach resonance (blue line). The resonance can be described by its position B_0 and width ΔB, defined as the separation between B_0 and the magnetic field B_{zero} for which $a_s = 0$ (also called the *zero crossing*, marked on the figure). Far from the resonance, a_s settles to the background value a_{bg}.

earlier that the anisotropy of the dipolar interaction mixes the collision channels.) It is worth noting that the bosonic isotopes of Er do not have a hyperfine structure, so FRs arise only due to the interaction anisotropy. However, given the large quantum numbers of the electronic configuration and the fact that $\ell > 0$ channels are involved even at ultracold temperatures, a large number of Feshbach resonances occur for magnetic lanthanides. Figure 2.6 shows the resonances for ^{166}Er and ^{168}Er, up to fields[20] of $B_0 = 8$ G. Unlike most other species used in ultracold-atom experiments, the Feshbach spectrum of Er is very rich even in this narrow field interval. Experimental measurements have been performed up to much higher fields, and a resonance number density of ca. 3 G^{-1} was observed for both isotopes [61, 126, 169].

For Cr, the ground state electron configuration is spherical, so assigning a particular FR to a particular channel configuration is easier (given the short-range van der Waals interaction is isotropic). On the other hand, lanthanides have an anisotropic electron

[20] In ultracold-atom experiments it is customary to express the magnetic field in gauss (G), 1 G = 10^{-4} T.

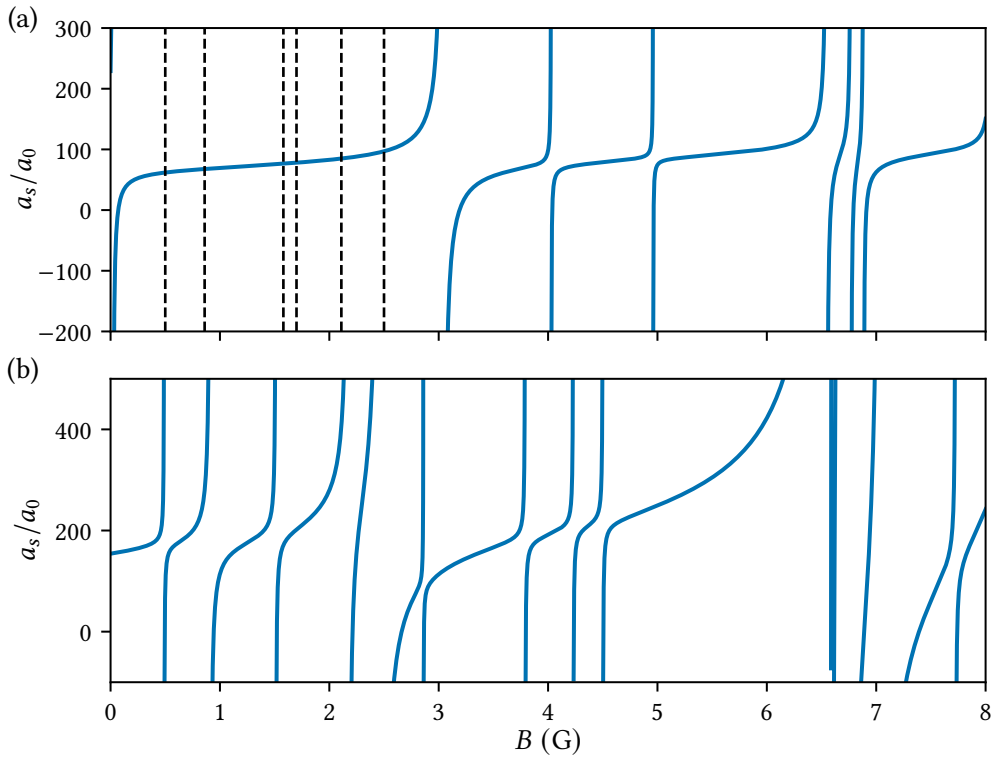

Figure 2.6. Feshbach resonances in ^{166}Er and ^{168}Er. Scattering length a_s as a function of the magnetic field B, for the two most abundant isotopes, (a) ^{166}Er and (b) ^{168}Er, for fields up to 8 G, using data from Refs. 126 and 170. The newly discovered loss features presented in Chapter 5 are shown with dashed lines (the scattering length has not been determined around these resonances yet).

configuration, which makes describing FRs more difficult. Moreover, it was shown that the arising anisotropy of the van der Waals interaction and the anisotropy of the dipolar interaction both significantly contribute to the number of FRs [171]. Furthermore, given the high number of possible channels, more than 40 channels need to be considered in simulations to reproduce the density of resonances (i.e. the number of FRs in a unit interval of magnetic field) reliably [169]. It is also worth noting that given the possible combinations of channels depend on the energy of the colliding particles, the number, exact shape and location of FRs are temperature-dependent. This is explored in Chapter 5 for Er.

If the states in the different channels have different magnetic moments (e.g. due to the differing spin configurations), applying an external magnetic field changes their relative energy (the energy gap) according to the Zeeman effect. If the open channel has a magnetic moment μ_o and the closed channel has μ_c, their relative energy changes according to $\Delta E = (\mu_o - \mu_c)B$ (cf. Eq. (2.9)). This means that by changing the magnetic field, the energy of a bound state in the closed channel can be made resonant with the energy of the atoms in the open channel, thereby affecting the scattering length. This has

been greatly exploited within ultracold atoms as it is a convenient way to control the magnitude and sign of interactions. It is also worth noting that the method of changing the relative energy need not be magnetic, it can also be e.g. optical, which opens the possibility of ultrafast and local control of interparticle interactions [172–174].

The scattering length in the vicinity of a resonance located at a magnetic field B_0 (where a_s diverges) with width ΔB can be approximated as [175]

$$a_s = a_{\mathrm{bg}}\left(1 - \frac{\Delta B}{B - B_0}\right), \tag{2.49}$$

where a_{bg} is the background value of a_s, away from the resonance (i.e. the scattering length associated with the open channel alone). The width of the resonance is defined as the change in the magnetic field needed from the resonance (at B_0) such that $a_s = 0$, and it relates to how strong the coupling is between the different channels. This form can be understood by expanding κ of Eq. (2.33) to first order in B around the resonance [102]. Given $\kappa = 0$ at the resonance (as $a_s = 1/\kappa$), we see $a_s \propto 1/(B - B_0)$ close to it. Away from the resonance, it takes the background value a_{bg}, so altogether it is usefully parametrised as in Eq. (2.49). The form of a_s near a resonance is shown in Fig. 2.5 and it is worth noting that both a_{bg} and Δ can be positive or negative. Furthermore, if there are more resonances present, the scattering length can be approximated as [176, 177]

$$a_s = a_{\mathrm{bg}} \prod_i \left(1 - \frac{\Delta B_i}{B - B_{0,i}}\right). \tag{2.50}$$

2.3 Dipolar Bose–Einstein condensates

Our experimental effort focuses on bosonic isotopes of erbium. A Bose gas, when cooled down sufficiently (typically to the order of 100 nK), undergoes a phase transition whereby the ground state of the system is occupied by a macroscopic number of atoms. This is called Bose–Einstein condensation (BEC). At very cold temperatures and weak interactions, the gas can be described by an order parameter, the so-called macroscopic wave function, with small excitations on top of the ground state. This approach provides the basis for exploring the influence of interactions and trapping geometry on the behaviour of the system theoretically, and in what follows we review some of the most important equations and results based on Refs. 64, 102 and 133.

2.3.1 Origins of Bose–Einstein condensation

The origins of Bose–Einstein condensation lie in the limitations of the available phase space at very low temperatures. For simplicity, we will first consider a free, non-interacting Bose gas (a gas made up of identical bosons) in three dimensions, with a fixed atom number N.

For bosons in thermodynamic equilibrium, the average occupation number of an energy level with energy ε_k is given by

$$\langle n_k \rangle = n(\varepsilon_k) = \frac{1}{e^{\beta(\varepsilon_k - \mu_c)} - 1}, \tag{2.51}$$

where μ_c is the chemical potential, $\beta = 1/k_B T$, k_B is the Boltzmann constant and T is the temperature. It is important to note that μ_c must be smaller than the lowest energy level ε_0, as otherwise for $\mu_c = \varepsilon_0$ the occupation number would be infinite, which would contradict our assumption that the atom number is finite. The density of states, i.e. the number of states within a unit interval of energy, is given by

$$g(\varepsilon) = \frac{V m^{3/2}}{\sqrt{2}\pi^2 \hbar^3} \sqrt{\varepsilon}, \tag{2.52}$$

where m is the atomic mass and V is the volume of the sample. The number of atoms that can be accommodated in these states can be expressed as

$$N_{\text{ex}} = \int_0^\infty n(\varepsilon) g(\varepsilon) \, \mathrm{d}\varepsilon \,. \tag{2.53}$$

To accommodate the largest possible number of atoms, we can take the limit $\mu_c \longrightarrow 0$ and obtain

$$N_{\text{ex}} = V \left(\frac{m k_B T}{2\pi\hbar^2} \right)^{3/2} \zeta\left(\frac{3}{2}\right), \tag{2.54}$$

where $\zeta(x)$ is the Riemann zeta function [178]. We see this is a finite number and $N_{\text{ex}} \propto T^{3/2}$.

This is an important result as if $N > N_{\text{ex}}$, some $(N - N_{\text{ex}})$ atoms cannot be accommodated in the excited states and have to occupy (condense into) the ground state. Combining the numerical prefactors to give

$$N_{\text{ex}} = N \left(\frac{T}{T_c} \right)^{\frac{3}{2}}, \tag{2.55}$$

where T_c is called the critical temperature, we find that the number of atoms in the ground

state is given by

$$N_0 = N\left(1 - \left(\frac{T}{T_c}\right)^{\frac{3}{2}}\right). \tag{2.56}$$

As the number of atoms in the ground state N_0 is proportional to the total number of atoms N, this is called macroscopic occupation. (Note that it is only the ground state that is macroscopically occupied.) It is important to note that Bose–Einstein condensation does not occur in all systems, as T_c crucially depends on the details of the system via $g(\varepsilon)$—here we have taken the simple example of non-interacting free particles in three dimensions. For example, it can be shown that non-interacting free atoms in two dimensions will never form a BEC. For neutral atoms in harmonic traps, T_c is on the order of 100 nK, i.e. atoms need to be cooled below this very cold temperature to obtain a BEC, which takes a significant experimental effort.

The need for this cold temperature can be understood in terms of the phase-space density $\rho = n\lambda_{\mathrm{dB}}^3$, where n is the number density of the atoms and $\lambda_{\mathrm{dB}} = h/\sqrt{2\pi m k_B T}$ is the thermal de Broglie wavelength. Equation (2.54) can be rewritten in terms of this quantity as

$$\rho_{\max} = \zeta\left(\frac{3}{2}\right) \approx 2.612\,, \tag{2.57}$$

where $\rho_{\max}$ is the maximum phase-space density before atoms start accumulating at the ground state for an unconfined, non-interacting Bose gas. If the phase-space density exceeds $\rho_{\max}$ (either via an increase in the number density or a reduction in temperature), Bose–Einstein condensation will occur.

We emphasise again that the value of $\rho_{\max}$ depends on the geometry and interactions in the system. For the important case of harmonic trapping, the density varies across the gas, so ρ does too. It can be shown that BEC will occur if the peak phase-space density (achieved at the centre of the trap) exceeds $\rho_{\max} = \zeta(3/2)$, which is the same result as for an unconfined gas (except that in that case the phase-space density is uniform across the gas) [133, 179].

We see $\rho \propto nT^{-3/2}$, so to achieve the critical ρ, we need to either increase the density or reduce the temperature. Given high densities lead to excessive three-body losses that destroy the condensate (see Chapter 5 for a thorough discussion), the density needs to be kept low. Therefore, to achieve the critical ρ, the temperature also needs to be extremely low. In fact, the required temperatures make ultracold-atom experiments the coldest places in the universe according to our current knowledge, as the temperature of outer space (2.7 K [180, 181]) and the coldest natural object ever observed (the Boomerang Nebula at 1 K [182], 5000 light years away from us in the constellation Centaurus) are

many orders of magnitude warmer than ultracold atoms.

2.3.2 Gross–Pitaevskii equation

We will now turn to the formal description of the condensate, using the formalism of second quantisation. Let us start by defining the so-called quantum field operator $\hat{\Psi}^\dagger$ ($\hat{\Psi}$), which creates (annihilates) a particle at position $\mathbf{r}$. It can be written via the single-particle states ψ_i as

$$\hat{\Psi}(\mathbf{r}, t) = \sum_i \psi_i \hat{a}_i \,, \tag{2.58}$$

where $\hat{a}_i^\dagger$ ($\hat{a}_i$) creates (annihilates) a particle in the single-particle state ψ_i. Given the ground state ψ_0 is macroscopically occupied (whereas the other states are not), it will be convenient to separate it out as

$$\hat{\Psi}(\mathbf{r}, t) = \psi_0 \hat{a}_0 + \sum_{i>0} \psi_i \hat{a}_i \,. \tag{2.59}$$

Given for the ground state $\left\langle \hat{a}_0^\dagger \hat{a}_0 \right\rangle = N_0 \gg 1$, the operators are of $O\left(\sqrt{N_0}\right)$. As their commutator $\left[\hat{a}_0, \hat{a}_0^\dagger\right] = 1 \ll N_0$, we can replace the operators with a complex number (c-number) $\sqrt{N_0}$, which is called the Bogoliubov approximation.[21] This is afforded by the fact that we have $N_0 \gg 1$ atoms in the BEC while all other states are occupied to a much smaller extent, and it amounts to treating the macroscopic component of the field operator as a classical field. We can then rewrite Eq. (2.59) as

$$\hat{\Psi}(\mathbf{r}, t) = \Psi_0(\mathbf{r}, t) + \delta\hat{\Psi}(\mathbf{r}, t) \,, \tag{2.60}$$

where $\Psi_0(\mathbf{r}, t) = \sqrt{N_0}\psi_0(\mathbf{r}, t)$. If one can neglect the non-condensed component $\delta\hat{\Psi}$ (which is the case for dilute Bose gases at very low temperatures), the system behaves like a classical object. This is analogous to the classical limit of quantum electrodynamics (QED), where the classical electromagnetic field replaces the microscopic description of photons.

Ψ_0 is called the macroscopic wave function (or condensate wave function), and plays the role of an order parameter for the phase transition to a BEC (the BEC exhibits long-range order). Expressing the macroscopic wave function as $\Psi_0 = \left\langle N-1 \middle| \hat{\Psi} \middle| N \right\rangle$ and recalling that the stationary states evolve in time according to $e^{-iEt/\hbar}$, we see that the time evolution of Ψ_0 is according to $e^{-i(E_N - E_{N-1})t/\hbar}$. Given $E_N - E_{N-1} \approx \partial E/\partial N = \mu_c$ for

[21] Recall that the commutation relations for bosons are $\left[\hat{a}_i, \hat{a}_j^\dagger\right] = \delta_{ij}$ and $\left[\hat{a}_i, \hat{a}_j\right] = \left[\hat{a}_i^\dagger, \hat{a}_j^\dagger\right] = 0$, where δ_{ij} is the Kronecker delta.

large N, we find

$$\Psi_0(\mathbf{r},t) = \Psi_0(\mathbf{r})e^{-i\mu_c t/\hbar}\,. \tag{2.61}$$

It is interesting to note that the time evolution of the order parameter is not governed by the energy of the condensate, as would happen for ordinary wave functions, but by the chemical potential which is a key parameter in the physics of BECs. It is important to remember at this point that the macroscopic wave function is not the many-body wave function of the system, hence we should not expect the same time evolution. We will expand on this important distinction later on.

Given $\left\langle \hat{\Psi}^\dagger \hat{\Psi} \right\rangle = N$, if all atoms are in the condensate and excitations can be neglected, we have

$$\int |\Psi_0(\mathbf{r},t)|^2 \,\mathrm{d}^3\mathbf{r} = N\,. \tag{2.62}$$

Therefore, $|\Psi_0(\mathbf{r},t)|^2$ equals the (number) density of the gas,

$$n(\mathbf{r},t) = |\Psi_0(\mathbf{r},t)|^2\,. \tag{2.63}$$

Let us now turn to the description of non-uniform condensates, which is the case for BECs formed in a trap. The Hamiltonian of the system in terms of the field operator is

$$\begin{aligned}\hat{H} = \int \frac{\hbar^2}{2m}\nabla\hat{\Psi}^\dagger(\mathbf{r})\nabla\hat{\Psi}(\mathbf{r})\,\mathrm{d}^3\mathbf{r} + \int V_{\mathrm{trap}}(\mathbf{r})\hat{\Psi}^\dagger(\mathbf{r})\hat{\Psi}(\mathbf{r})\,\mathrm{d}^3\mathbf{r} \\ + \frac{1}{2}\iint \hat{\Psi}^\dagger(\mathbf{r}')\hat{\Psi}^\dagger(\mathbf{r})V_{\mathrm{int}}(\mathbf{r}'-\mathbf{r})\hat{\Psi}(\mathbf{r}')\hat{\Psi}(\mathbf{r})\,\mathrm{d}^3\mathbf{r}'\,\mathrm{d}^3\mathbf{r}\,,\end{aligned} \tag{2.64}$$

where $V_{\mathrm{trap}}(\mathbf{r})$ is an external trapping potential and $V_{\mathrm{int}}(\mathbf{r})$ is the two-body interaction potential. In this Hamiltonian, the first term captures the kinetic energy, the second corresponds to the potential energy due to trapping and the third captures the contribution of interactions. Here we took advantage of the fact that we are considering dilute gases, where the range of interatomic forces is much less than the average interparticle distance, allowing one to consider only pairs of interacting particles, while configurations with three (or more) particles interacting simultaneously can be safely neglected.

To devise the equations governing the field $\Psi_0(\mathbf{r},t)$, we can use the Heisenberg equation to deduce

$$\begin{aligned}i\hbar\frac{\partial}{\partial t}\hat{\Psi}(\mathbf{r},t) &= \left[\hat{\Psi}(\mathbf{r},t),\hat{H}\right] \\ &= \left(-\frac{\hbar^2}{2m}\nabla^2 + V_{\mathrm{trap}}(\mathbf{r},t) + \int \hat{\Psi}^\dagger(\mathbf{r}',t)V_{\mathrm{int}}(\mathbf{r}'-\mathbf{r})\hat{\Psi}(\mathbf{r}',t)\,\mathrm{d}^3\mathbf{r}'\right)\hat{\Psi}(\mathbf{r},t)\,,\end{aligned} \tag{2.65}$$

where we used the commutation relations of the field operator $\hat{\Psi}$,

$$\left[\hat{\Psi}(\mathbf{r}), \hat{\Psi}^\dagger(\mathbf{r}')\right] = \delta(\mathbf{r}-\mathbf{r}'), \quad \left[\hat{\Psi}(\mathbf{r}), \hat{\Psi}(\mathbf{r}')\right] = \left[\hat{\Psi}^\dagger(\mathbf{r}), \hat{\Psi}^\dagger(\mathbf{r}')\right] = 0\,. \tag{2.66}$$

Proceeding similarly to Eq. (2.60), one can replace $\hat{\Psi}$ with Ψ_0 when the non-condensed component $\delta\hat{\Psi}$ can be neglected and obtain[22]

$$i\hbar\frac{\partial}{\partial t}\Psi(\mathbf{r},t) = \left(-\frac{\hbar^2}{2m}\nabla^2 + V_{\text{trap}}(\mathbf{r},t) + \int \Psi^*(\mathbf{r}',t)V_{\text{int}}(\mathbf{r}'-\mathbf{r})\Psi(\mathbf{r}',t)\,\mathrm{d}^3\mathbf{r}'\right)\Psi(\mathbf{r},t)\,. \tag{2.67}$$

Let us now consider the two-body interaction potential. Given the distance between the interacting particles is large, we can use the asymptotic expression for the wave function of their relative motion, whose form is fixed by the scattering amplitude $f(\theta,\varphi,k)$ as we have shown in §2.2.3 (cf. Eq. (2.17)). This implies that all macroscopic properties of the system can be expressed in terms of $f(\theta,\varphi,k)$, which for short-range interactions and low energies is determined by the s-wave scattering length a_s (cf. Eq. (2.34)). This underlines the importance of a_s when describing the properties of the condensate.

If the Born approximation is valid, we can reproduce the low-energy scattering properties of the full two-body potential $V_{\text{int}}(\mathbf{r})$ by using the (effective) pseudo-potential [150, 151] as presented in Eq. (2.47),

$$V_{\text{eff}}(\mathbf{r}'-\mathbf{r}) = \frac{4\pi\hbar^2 a_s(\mu)}{m}\delta(\mathbf{r}'-\mathbf{r})\frac{\partial}{\partial|\mathbf{r}'-\mathbf{r}|}|\mathbf{r}'-\mathbf{r}| + \frac{\mu_0}{4\pi}\frac{\mu^2-3(\mathbf{n}\cdot\boldsymbol{\mu})^2}{|\mathbf{r}'-\mathbf{r}|^3}\,, \tag{2.68}$$

where $\boldsymbol{\mu}$ is the magnetic dipole moment (vector) of the (magnetically aligned) atoms and $\mathbf{n} = (\mathbf{r}'-\mathbf{r})/|\mathbf{r}'-\mathbf{r}|$. Inserting this into Eq. (2.67) we finally obtain the (dipolar, time-dependent) Gross–Pitaevskii equation (GPE) [183–185],

$$i\hbar\frac{\partial}{\partial t}\Psi(\mathbf{r},t) = \left(-\frac{\hbar^2}{2m}\nabla^2 + V_{\text{trap}}(\mathbf{r},t) + g_s|\Psi(\mathbf{r},t)|^2 + \Phi_{\text{dd}}(\mathbf{r},t)\right)\Psi(\mathbf{r},t)\,, \tag{2.69}$$

where the short-range interaction parameter is $g_s = 4\pi\hbar^2 a_s/m$ and the dipolar contribution is

$$\Phi_{\text{dd}}(\mathbf{r},t) = \int |\Psi(\mathbf{r}',t)|^2 U_{\text{dd}}(\mathbf{r}'-\mathbf{r})\,\mathrm{d}^3\mathbf{r}', \tag{2.70}$$

where $U_{\text{dd}}(\mathbf{r}'-\mathbf{r})$ is according to Eq. (2.14). In the case of stationary states, where $\Psi(\mathbf{r},t)$ evolves in time according to Eq. (2.61), this leads to the time-independent GPE [186],

$$\mu_c\Psi(\mathbf{r}) = \left(-\frac{\hbar^2}{2m}\nabla^2 + V_{\text{trap}}(\mathbf{r}) + g_s|\Psi(\mathbf{r})|^2 + \Phi_{\text{dd}}(\mathbf{r})\right)\Psi(\mathbf{r})\,. \tag{2.71}$$

[22] From here on, we drop the index 0 for ease of notation.

The time-independent GPE can be simplified further when the Hamiltonian is dominated by the atomic interactions and the trapping potential, and the kinetic energy can be neglected [187]. This is called the Thomas–Fermi approximation, in which case the time-independent GPE becomes

$$\mu_c = V_{\text{trap}}(\mathbf{r}) + g_s|\Psi(\mathbf{r})|^2 + \Phi_{\text{dd}}(\mathbf{r})\,. \tag{2.72}$$

The benefit of this approximation is that it yields an equation which is much easier to solve than the full GPE, and in some cases, Eq. (2.72) is analytically solvable.

Equations (2.69) and (2.71) are the main theoretical tools for investigating dilute non-uniform Bose gases at low temperatures, as the time-independent GPE Eq. (2.71) can be used to describe stationary states (including the ground state) and the time-dependent GPE Eq. (2.69) can be used to explore the dynamics of the condensate. They are also the starting point for our considerations regarding dipolar gases in box-like potentials in Chapter 6. In most cases, the GPE has to be solved numerically for Ψ to give a self-consistent solution, which is not an easy task as it is not linear in $\Psi(\mathbf{r})$, Φ_{dd} is non-local and it involves an integral. Solutions of Ψ have to satisfy the normalisation condition Eq. (2.62), and for the time-independent case, the value of the chemical potential μ_c is fixed by this normalisation condition.

It is important to point out, that the GPE is not the Schrödinger equation, and $\Psi(\mathbf{r})$ is not the many-body wave function. $\Psi(\mathbf{r})$ is the order parameter of Bose–Einstein condensation, sometimes called the condensate (or macroscopic) wave function. This means that different solutions to the GPE, Ψ_a and Ψ_b, need not be orthogonal, but the corresponding many-body wave functions, Φ_a and Φ_b, do. Furthermore, we need to remember that $\mu_c = \partial E/\partial N \neq E/N$.

As a consequence of the diluteness condition, one can ignore correlations among the particles to first approximation, which is called the Hartree–Fock approximation. One can then write the many-body wave function of the system as

$$\Phi(\mathbf{r}_1, \mathbf{r}_2, \ldots, \mathbf{r}_N) = \prod_{i=1}^{N} \psi(\mathbf{r}_i) = \prod_{i=1}^{N} \frac{1}{\sqrt{N}} \Psi(\mathbf{r}_i)\,, \tag{2.73}$$

where $\psi(\mathbf{r})$ is the single-particle wave function. In this case, we see

$$\langle \Phi_a | \Phi_b \rangle = \left(\frac{1}{N} \int \Psi_a^* \Psi_b \, \mathrm{d}^3\mathbf{r} \right)^N , \tag{2.74}$$

which vanishes as $N \to \infty$, as $\langle \Psi_a | \Psi_b \rangle / N < 1$ except if $a = b$.

Equation (2.71) admits different solutions, and the solution with the lowest energy defines the macroscopic wave function of the ground state, which is a real function. As the energy of the system is given by

$$\begin{aligned} E &= \left\langle \hat{H} \right\rangle \\ &= \int \left(\frac{\hbar^2}{2m} |\nabla \Psi(\mathbf{r})|^2 + V_{\text{trap}}(\mathbf{r}) |\Psi(\mathbf{r})|^2 + \frac{g_s}{2} |\Psi(\mathbf{r})|^4 + \frac{1}{2} \Phi_{\text{dd}}(\mathbf{r}) |\Psi(\mathbf{r})|^2 \right) \mathrm{d}^3\mathbf{r}\,, \end{aligned} \tag{2.75}$$

we see that the phase of the generally complex Ψ does not affect the contributions of the trap potential or the interactions to the energy. Writing $\Psi(\mathbf{r}) = |\Psi(\mathbf{r})| e^{i\varphi(\mathbf{r})}$, such that $\varphi(\mathbf{r})$ is the phase of the macroscopic wave function, we find $|\nabla \Psi(\mathbf{r})|^2 = (\nabla |\Psi(\mathbf{r})|)^2 + |\Psi(\mathbf{r})|^2 (\nabla \varphi(\mathbf{r}))^2$. We see that any variation in φ would lead to an increase in energy, so for the ground state φ needs to be constant which we can choose to be 0, i.e. we can choose $\Psi(\mathbf{r})$ to be real. On the other hand, solutions of Eq. (2.71) corresponding to excited states are usually given by complex functions. It is worth noting that the ground state can be approximated by assuming it has a certain shape (e.g. a Gaussian) depending on some free parameters (e.g. its width), and minimising the energy with respect to these parameters. Furthermore, the GPE itself can also be derived by considering the variations of E subject to the constraint that the number of particles N (given by the normalisation condition Eq. (2.62)) is fixed. In that case, one can utilise the method of Lagrange multipliers to consider the variation of $E - \mu_c N$ with respect to Ψ and Ψ^*, where μ_c acts as the Lagrange multiplier.

We have referred to the lowest-energy solution of Eq. (2.71) as the ground state of the system, which is not, however, completely accurate. In fact, it is well known that the ground state of most physical systems interacting via interatomic potentials does not correspond to a gas but rather to a solid. Therefore, for such systems, the BEC gas phase is only a metastable configuration, where thermalisation is ensured by two-body collisions. However, the theory ignores three-body collisions which, in real systems, eventually drive the system into the solid phase. On the other hand, experiments carried out using various atomic species have proven that the BEC phase can be realised and that it survives for large enough times such that many relevant physical quantities can be measured. With this caveat in mind we will refer in the following to the lowest-energy BEC state as the ground state.

We should also reiterate that the GPE only applies if the system has a large number of atoms (such that one can talk about Bose–Einstein condensation), the gas is dilute (such that one can neglect interactions involving more than two atoms) and the temperature is essentially zero (so one can neglect excitations depleting the condensate). Current

experiments are able to create quasi-pure BECs with very small thermal fractions, so all these conditions can be satisfied. Furthermore, we can only use it to model phenomena taking place over distances much larger than the scattering length, as for microscopic distances the approximations needed to derive the GPE are no longer valid.

It is also worth pointing out that in Eq. (2.68) we separated the short- and long-range parts of the two-body potential into two different terms, which is not an obviously correct thing to do. Furthermore, we mixed together the pseudo-potential for the short-range interaction and the real potential for the dipole–dipole interaction (cf. Eq. (2.14)). Furthermore, a_s itself depends on μ, as explicitly noted in the equation. These points were a topic of much debate, and it was concluded that Eq. (2.68) is correct away from shape resonances [149–151, 153, 154, 188]. Further to this approximation, more rigorous effective potentials have been derived, which include a velocity dependence [189].

2.3.3 Excitations and stability

A large part of the attractiveness of using dipolar atoms in ultracold gas experiments is the effect of their anisotropic interactions on the condensate's excited states, which exhibit exotic properties. Let us now therefore consider excitations which are small compared to the stationary value of the order parameter, writing the order parameter in the form

$$\Psi(\mathbf{r}, t) = (\Psi_0(\mathbf{r}) + \delta\Psi(\mathbf{r}, t))\, e^{-i\mu_c t/\hbar}, \tag{2.76}$$

where $\Psi_0(\mathbf{r})$ is the order parameter of the ground state, μ_c is its chemical potential and $\delta\Psi \ll \Psi_0$ is a small perturbation. Ψ can solve the time-dependent GPE if we look for an oscillatory perturbation in the form

$$\delta\Psi(\mathbf{r}, t) = \sum_i \left(u_i(\mathbf{r})e^{-i\omega_i t} + v_i^*(\mathbf{r})e^{i\omega_i t}\right), \tag{2.77}$$

where ω_i are the (real) frequencies of the excitations described by u_i and v_i. This ansatz can be inserted into the time-dependent GPE, Eq. (2.69). Looking for perturbations on top of the ground state, such that Ψ_0 is real, when the GPE is expanded up to first order in $\delta\Psi$ and the terms are collected according to the separate frequency components $e^{\pm i\omega_i t}$, for terms with $e^{-i\omega_i t}$ we find[23]

$$\hbar\omega_i u_i = \left(\hat{H}_0 - \mu_c + \hat{C} + \hat{X}\right) u_i + \hat{X} v_i, \tag{2.78}$$

[23] From here on, we drop the explicit $\mathbf{r}$ dependence of the operators and functions for ease of notation.

whereas for terms with $e^{i\omega_i t}$ we find

$$-\hbar\omega_i v_i = \left(\hat{H}_0 - \mu_c + \hat{C} + \hat{X}\right) v_i + \hat{X} u_i \,, \tag{2.79}$$

where we assumed that ω_i are real and $\hat{H}_0 = -\hbar^2\nabla^2/2m + V_{\text{trap}}$. $\hat{C}$ is called the direct interaction term and $\hat{X}$ the exchange interaction term, whose actions on a function $f(\mathbf{r})$ are

$$\hat{C} f(\mathbf{r}) = f(\mathbf{r}) \int V_{\text{eff}}(\mathbf{r}' - \mathbf{r}) \Psi_0^2(\mathbf{r}') \, \mathrm{d}^3\mathbf{r}', \tag{2.80}$$

$$\hat{X} f(\mathbf{r}) = \Psi_0(\mathbf{r}) \int V_{\text{eff}}(\mathbf{r}' - \mathbf{r}) \Psi_0(\mathbf{r}') f(\mathbf{r}') \, \mathrm{d}^3\mathbf{r}'. \tag{2.81}$$

Equations (2.78) and (2.79) are called the Bogoliubov–de Gennes (BdG) equations. We can express this set of equations in matrix form as

$$\begin{pmatrix} \hat{H}_0 - \mu_c + \hat{C} + \hat{X} & \hat{X} \\ -\hat{X} & -\hat{H}_0 + \mu_c - \hat{C} - \hat{X} \end{pmatrix} \begin{pmatrix} u_i \\ v_i \end{pmatrix} = \hbar\omega_i \begin{pmatrix} u_i \\ v_i \end{pmatrix}, \tag{2.82}$$

which is equivalent to

$$\begin{pmatrix} 0 & \hat{H}_0 - \mu_c + \hat{C} \\ \hat{H}_0 - \mu_c + \hat{C} + 2\hat{X} & 0 \end{pmatrix} \begin{pmatrix} u_i + v_i \\ u_i - v_i \end{pmatrix} = \hbar\omega_i \begin{pmatrix} u_i + v_i \\ u_i - v_i \end{pmatrix}. \tag{2.83}$$

If we take the square of the matrix and define $u_i + v_i = f_i$ and $u_i - v_i = g_i$, we see

$$\begin{pmatrix} 0 & \hat{H}_0 - \mu_c + \hat{C} \\ \hat{H}_0 - \mu_c + \hat{C} + 2\hat{X} & 0 \end{pmatrix}^2 \begin{pmatrix} f_i \\ g_i \end{pmatrix} = (\hbar\omega_i)^2 \begin{pmatrix} f_i \\ g_i \end{pmatrix}. \tag{2.84}$$

Once we calculate the squared matrix, we find it defines a set of decoupled equations,

$$\left(\hat{H}_0 - \mu_c + \hat{C}\right)\left(\hat{H}_0 - \mu_c + \hat{C} + 2\hat{X}\right) f_i = (\hbar\omega_i)^2 f_i \,, \tag{2.85}$$

$$\left(\hat{H}_0 - \mu_c + \hat{C} + 2\hat{X}\right)\left(\hat{H}_0 - \mu_c + \hat{C}\right) g_i = (\hbar\omega_i)^2 g_i \,. \tag{2.86}$$

These equations in general need to be solved numerically. However, it is enough to solve only one of them, after which we can use Eq. (2.84) to obtain f_i from g_i and vice versa. Once that is done, we can obtain the original u_i and v_i as $u_i = (f_i + g_i)/2$ and $v_i = (f_i - g_i)/2$.

The solutions of the BdG equations Eq. (2.82) have some special properties which are worth pointing out. First of all, it can be shown that solutions must have a real frequency (as we originally assumed), and the occurrence of a complex frequency signals a dynamic instability of the system. Furthermore, looking at Eq. (2.77), it can be seen

that if $\{u_i, v_i\}$ is a solution with frequency ω_i, $\{v_i^*, u_i^*\}$ is an equivalent solution with frequency $-\omega_i$. Finally, note that $\omega = 0$ is always a solution with $u = \alpha\Psi_0$ and $v = -\alpha\Psi_0$, where α is a (complex) constant. In that case the order parameter becomes $\Psi(\mathbf{r}, t) = \Psi_0(\mathbf{r})(1 + \alpha - \alpha^*)e^{-i\mu_c t/\hbar} \approx \Psi_0(\mathbf{r})e^{-i\mu_c t/\hbar + \alpha - \alpha^*}$, which is really a gauge transformation in which the phase of the order parameter is changed by $(\alpha - \alpha^*)/i$. This transformation is not a physical excitation of the system. In fact, it can be shown that the energy increase associated with the presence of these excitations is

$$\Delta E = \sum_i \hbar\omega_i \int \left(|u_i|^2 - |v_i|^2\right) \mathrm{d}^3\mathbf{r}\,. \tag{2.87}$$

Therefore, the excitation energy spectrum of the condensate can be calculated by solving the BdG equations.

We can also formulate excitations in the scope of second quantisation [190, 191], where the Hamiltonian is given by Eq. (2.64), but we work to a higher-order expansion. Following the classical treatment, we write the field operator as

$$\hat{\Psi}(\mathbf{r}, t) = \left(\Psi_0(\mathbf{r}) + \delta\hat{\Psi}(\mathbf{r}, t)\right) e^{-i\mu_c t/\hbar} \tag{2.88}$$

with

$$\delta\Psi(\mathbf{r}, t) = \sum_i \left(u_i(\mathbf{r})\hat{b}_i e^{-i\omega_i t} + v_i^*(\mathbf{r})\hat{b}_i^\dagger e^{i\omega_i t}\right), \tag{2.89}$$

where $\hat{b}_i^\dagger$ ($\hat{b}_i$) are the creation (annihilation) operators associated with the excitations and u_i, v_i and ω_i are the solutions of the (classical) BdG equations. Similarly to the classical case, Eq. (2.87), it can be shown that the change in the Hamiltonian is given by

$$\Delta\hat{H} = \sum_i \hbar\omega_i \int \left(|u_i|^2 - |v_i|^2\right) \mathrm{d}^3\mathbf{r}\, \hat{b}_i^\dagger \hat{b}_i \tag{2.90}$$

plus a constant term arising from the commutation between $\hat{b}_i^\dagger$ and $\hat{b}_i$. As this term is the same for the ground state and the excited state, it can be neglected. Furthermore, it is conventional to normalise u_i and v_i according to

$$\int \left(|u_i|^2 - |v_i|^2\right) \mathrm{d}^3\mathbf{r} = 1\,, \tag{2.91}$$

such that Eq. (2.90) becomes

$$\Delta\hat{H} = \sum_i \hbar\omega_i \hat{b}_i^\dagger \hat{b}_i\,. \tag{2.92}$$

This implies that the system can be described by independent quasi-particles (elementary

excitations) with energy $\hbar\omega_i$. In that picture, the ground state of the system corresponds to the vacuum of quasi-particles. Furthermore, in this normalisation, $\hat{b}_i^\dagger$ and $\hat{b}_i$ satisfy the Bose commutation relations.

Generally, the BdG equations can only be solved numerically. However, in some simple cases, analytical solutions are possible, which provide valuable insights. Let us first consider the case of an unconfined, spatially homogeneous dipolar gas with (number) density n in equilibrium (i.e. $\Psi_0 = \sqrt{n}$), where the dipoles are polarised along a certain direction. In that case, the solutions are plane waves whose energy (normalised according to Eq. (2.91)) reads

$$\hbar\omega(\mathbf{k}) = \sqrt{\frac{\hbar^2k^2}{2m}\left(\frac{\hbar^2k^2}{2m} + 2n\tilde{V}_{\text{int}}(\mathbf{k})\right)}. \tag{2.93}$$

Here the dependence of ω on the wave vector associated with the plane wave, $\mathbf{k}$, is made explicit (the momentum of the excitation is $\hbar\mathbf{k}$ and $k = |\mathbf{k}|$). $\tilde{V}_{\text{int}}(\mathbf{k})$ is the Fourier transform[24] of the interactions which, for the case of contact and dipolar interactions, is given by

$$\tilde{V}_{\text{int}}(\mathbf{k}) = g_s + \frac{C_{\text{dd}}}{3}\left(3\cos^2\theta_{\mathbf{k}} - 1\right), \tag{2.94}$$

where $\theta_{\mathbf{k}}$ is the angle between $\mathbf{k}$ and the direction of the dipoles. It is worth noting that analogously to $g_s = 4\pi\hbar^2a_s/m$, it is also customary to define $g_{\text{dd}} = C_{\text{dd}}/3 = 4\pi\hbar^2a_{\text{dd}}/m$. Casting the solution in terms of the scattering length a_s and the dipolar length a_{dd} yields

$$\hbar\omega(\mathbf{k}) = \sqrt{\frac{\hbar^2k^2}{2m}\left(\frac{\hbar^2k^2}{2m} + \frac{8\pi\hbar^2}{m}n\big(a_s + a_{\text{dd}}\left(3\cos^2\theta_{\mathbf{k}} - 1\right)\big)\right)}. \tag{2.95}$$

There are some important features which are worth noting. First of all, it can be seen that for a purely contact-interacting gas ($a_{\text{dd}} = 0$), $\omega(\mathbf{k})$ does not depend on the direction of $\mathbf{k}$ and is only real for all k if $a_s \geq 0$. Specifically, if $a_s < 0$, small-k excitations render the gas unstable (as they yield an imaginary ω), which is the well-known result that a (non-dipolar) condensate is only stable for $a_s \geq 0$. Furthermore, for small k we find $\omega \approx k\sqrt{g_sn/m}$. This relation describes sound waves with phase velocity $\omega/k = \sqrt{g_sn/m}$, which are therefore called phonon modes. On the other hand, in the large-k limit one obtains $\hbar\omega = \hbar^2k^2/2m + g_sn$, which is the relation for a free particle. We see the transition between the regimes occurs at $\hbar^2k_{\text{tr}}^2/2m \approx g_sn$. This defines a length scale called the healing length,

$$\xi = \frac{1}{k_{\text{tr}}} = \sqrt{\frac{\hbar^2}{2mg_sn}}. \tag{2.96}$$

[24] The Fourier transform of a function $f(\mathbf{r})$ is defined here as $\tilde{f}(\mathbf{k}) = \int e^{-i\mathbf{k}\cdot\mathbf{r}}f(\mathbf{r})\,\mathrm{d}^3\mathbf{r}$.

This length scale emerges as an important one in a variety of phenomena, where it characterises the size of disturbances, e.g. the size of the core of vortices.

For a dipolar gas, we see that the energy of the excitation (and so the speed of sound) has a directional dependence (as it depends on the angle $\theta_{\mathbf{k}}$), and that the gas is unstable for $a_{\mathrm{dd}} > a_s$, or equivalently, $\varepsilon_{\mathrm{dd}} > 1$. In fact it is this relation that is used to define the prefactors of a_{dd} (in Eq. (2.42)), and hence $\varepsilon_{\mathrm{dd}}$. It is also worth noting that an unconfined purely dipolar gas ($a_s = 0$) is never stable, as certain angles $\theta_{\mathbf{k}}$ yield imaginary frequencies. It is interesting to note that the most unstable situation is the case of $\theta_{\mathbf{k}} = 90°$, when the direction of the wave vector is perpendicular to the orientation of the dipoles. At first sight, this might seem counterintuitive: as dipoles side-by-side repel each other, one could (wrongly) conclude that the most unstable phonons should correspond to $\mathbf{k}$ being parallel to the dipoles. However, for $\theta_{\mathbf{k}} = 90°$ it is the wavefront that is parallel to the dipoles—which causes a density increase along the direction of the dipoles, thereby enhancing the attractive part of their interaction. On the other hand, for $\theta_{\mathbf{k}} = 0°$ the wavefront (and the density increase) is perpendicular to the dipoles, which enhances the repulsive part of the interaction.

Another important case is that of a dipolar BEC tightly confined in a harmonic trap along the polarisation direction of the dipoles (let us call this $\hat{\mathbf{z}}$) and unconfined in the other two directions, such that the trapping potential is

$$V_{\mathrm{trap}}(\mathbf{r}) = \frac{1}{2} m\omega_z^2 z^2 \,. \tag{2.97}$$

If we approximate that the condensate has a Gaussian density profile along the trapping direction, it can be shown that the excitation spectrum for in-plane modes (i.e. no excitation along the tight direction) is given by [192]

$$\hbar\omega(k_\rho) = \sqrt{\frac{\hbar^2 k_\rho^2}{2m}\left(\frac{\hbar^2 k_\rho^2}{2m} + 2n_{2\mathrm{D}}\tilde{V}_{2\mathrm{D}}(k_\rho)\right)}, \tag{2.98}$$

where $n_{2\mathrm{D}}$ is the areal (number) density of the gas, k_ρ is the magnitude of the in-plane wave vector and $\tilde{V}_{2\mathrm{D}}$ is the Fourier transform of the interaction potential once the trapping direction has been integrated out, given by

$$\tilde{V}_{2\mathrm{D}}(k_\rho) = \frac{g_s}{\sqrt{2\pi}\ell_z} + \frac{g_{\mathrm{dd}}}{\sqrt{2\pi}\ell_z} F_\perp\left(\frac{k_\rho \ell_z}{\sqrt{2}}\right), \tag{2.99}$$

where $\ell_z = \sqrt{\hbar/m\omega_z}$ is the harmonic oscillator length, $F_\perp(x) = 2 - 3\sqrt{\pi}x e^{x^2}\mathrm{erfc}(x)$ and $\mathrm{erfc}(x) = 1 - \mathrm{erf}(x)$ is the complementary error function [178].

Equation (2.98) has some striking properties which sparked experimental and theoretical interest. Figure 2.7(a) shows an example spectrum, demonstrating the emergence of a plateau (a local maximum) followed by a dip (a local minimum) when $\varepsilon_{\mathrm{dd}}$ is high enough.[25] (This is due to the fact that $\tilde{V}_{\mathrm{2D}}$ becomes negative above $k_\rho \ell_z \approx \sqrt{2}$.) As liquid helium exhibits a similar excitation spectrum [73–75], following the nomenclature introduced by Landau, the finite-wavelength excitation minimum is called a 'roton' (note that nothing is rotating in our case) whereas the preceding local maximum is called a 'maxon', and the whole shape is referred to as the roton–maxon spectrum [71, 72, 193, 194]. It is worth noting that this spectrum only develops in the Thomas–Fermi regime, where the condensate (in the Thomas–Fermi approximation) assumes a parabolic profile [195]. (This profile can be estimated with a Gaussian, whose width is then varied to minimise the energy of the ground state [64].) However, in the strict quasi-2D case, where the condensate is assumed to be in the ground state of the harmonic oscillator along z, the roton spectrum actually never develops [196]—too tight a trapping prevents the dipoles from sampling the attractive part of the potential. As all these length scales are $O(\ell_z)$ and the shapes are not too dissimilar, this does not qualitatively modify the physics.

The softening of (i.e. a dip developing in) the excitation spectrum at intermediate wavelengths can be understood by looking at the form of Eq. (2.99), underlined by the geometry of perturbations shown in Fig. 2.7(b). For small k_ρ, the dipolar interactions are mostly repulsive and interactions lead to a phonon mode. On the other hand, for large k_ρ, the dipolar interaction is mostly attractive which decreases the energy of excitations, and if the s-wave part is weak enough compared to the dipolar part, it can create a local minimum in the excitation spectrum at intermediate wavelengths ($k_{\mathrm{rot}} \ell_z \approx 1$). However, as the momentum of excitations is increased, the kinetic part of the spectrum overpowers the dipolar contribution, leading to a free-particle-like solution as before.

As the particle density increases, the roton gap decreases. At a critical density the roton gap (the energy of the roton minimum) becomes zero, and using a different approximation to Eq. (2.98) it can be shown that $k_{\mathrm{rot}} = \sqrt{2}/\ell_z$ at this point [193]. As the density is increased further, the gas collapses as the excitation energy becomes imaginary. It is interesting to note that in this geometry even a purely dipolar gas can be stable, as the tight trapping stabilises the gas. Moreover, the dipolar interaction can stabilise a gas even with a slightly negative a_s, as the repulsive part of the potential compensates the attractive contact interactions. Finally, the roton minimum also has a profound effect on the Landau critical

[25] What 'high enough' means depends on the system parameters n_{2D}, m, ω_z and a_{dd}, but for realistic systems the roton minimum appears when $\varepsilon_{\mathrm{dd}}$ is $O(1)$.

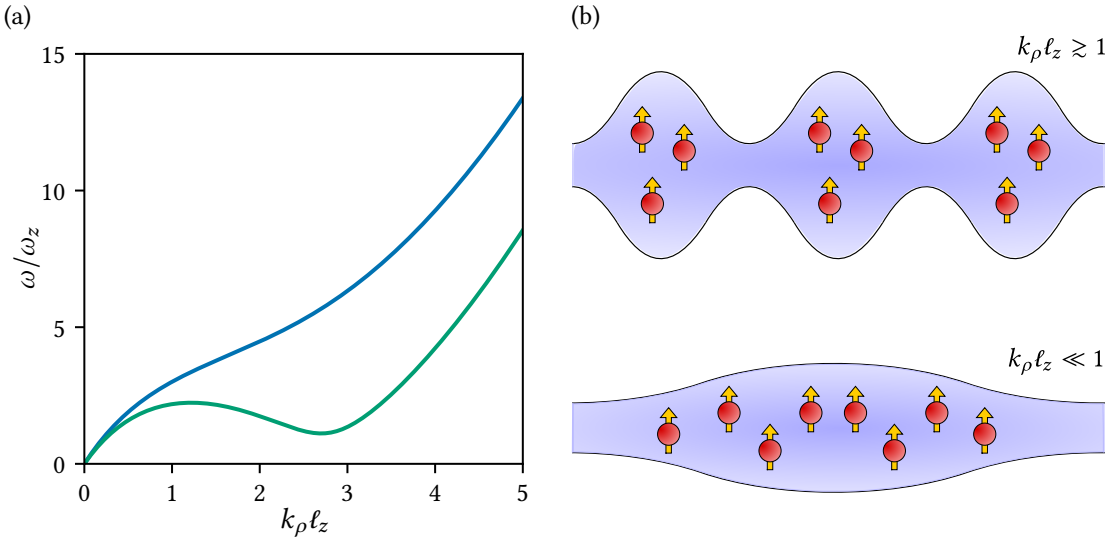

Figure 2.7. Roton minimum. (a) The excitation spectrum of a dipolar gas, harmonically trapped along the polarisation direction but free in the other two, for $n_{2D}\ell_z^2 = 300$, $a_{dd}/\ell_z = 0.01$, and $\varepsilon_{dd} = 1$ (blue) and $\varepsilon_{dd} = 2$ (green). For a less dipolar gas the spectrum features no local (roton) minimum, but for a more dipolar one it does. (b) A drawing showing the (exaggerated) effect of short- and long-wavelength excitations on the density distribution. While long-wavelength excitations mostly sample the repulsive ('side-by-side') part of dipolar interactions, short-wavelength ones sample the attractive ('head-to-tail') part more. This leads to a negative energy contribution, which at intermediate wavelengths can lead to a roton minimum. At very short wavelengths the kinetic energy overpowers the negative contribution of the dipolar interaction, leading to a free-particle form. Adapted from Ref. 104.

velocity, given by

$$v_c = \min_{\mathbf{k}} \left(\frac{\omega(\mathbf{k})}{\hat{\mathbf{v}} \cdot \mathbf{k}} \right) \tag{2.100}$$

for excitations propagating in an (infinite, uniform) superfluid in direction $\hat{\mathbf{v}}$ [73, 197]: as the roton minimum forms, the critical velocity is determined by the decreasing gap.

While the models above were developed in infinite (or semi-infinite) systems, some of their findings can be translated to fully trapped configurations, where the trapping changes the excitation spectrum. However, cylindrical traps with large aspect ratios ('pancake traps') have been formed in experiments and the rotonisation of the energy spectrum has been demonstrated. It is this energy spectrum that motivates our investigations in Chapter 6.

2.3.4 Beyond-mean-field corrections

The Gross–Pitaevskii equation presented in §2.3.2 is sometimes called a mean-field approach, as we effectively replaced the individually interacting bosons with non-interacting particles in the potential $V_{\text{trap}}+V_{\text{eff}}$, arising from the trap and the average interactions with all other bosons (the 'mean field'). However, it is possible to work to a higher accuracy, keeping up to quadratic terms in the Hamiltonian. This is in fact what we have already done when describing excitations: the Hamiltonian for an unconfined gas is essentially $\hat{H} = E_0 + \sum_i \hbar\omega_i \hat{b}_i^\dagger \hat{b}_i$.

The zero-point (quantum) fluctuations of these elementary excitations, described by $\hat{b}_i$, lead to a correction to the energy of the ground state. This has first been calculated by Lee, Huang and Yang for contact interactions [198, 199], and is therefore called the Lee–Huang–Yang (LHY) term. For an unconfined dipolar gas it is given by [200–202]

$$\Delta E_{\text{LHY}} = \frac{g_s}{2}\frac{N_0^2}{V}\frac{128}{15\sqrt{\pi}}\sqrt{n_0 a_s^3}\, Q_5(\varepsilon_{\text{dd}})\,, \tag{2.101}$$

where $Q_l(x) = \int_0^1 \left(1 - x + 3xu^2\right)^{l/2} \mathrm{d}u$, V is the volume the gas occupies (with cyclic boundary conditions), N_0 is the number of particles in the condensate as before and $n_0 = N_0/V$ is their density. It is worth noting that the energy change is set by the so-called gas parameter $n_0 a_s^3$. Furthermore, as an unconfined dipolar gas is unstable for $\varepsilon_{\text{dd}} > 1$, $Q_5(x)$ becomes complex for $x > 1$. However, the imaginary part is small for $\varepsilon_{\text{dd}} \lesssim 3$, and in current literature it is ignored in this regime without a complete breakdown of the theory. One can use Eq. (2.101) to calculate the change in the chemical potential,

$$\Delta\mu_{\text{LHY}} = \frac{\partial \Delta E_{\text{LHY}}}{\partial N_0} = \frac{32}{3\sqrt{\pi}} g_s n_0 \sqrt{n_0 a_s^3}\, Q_5(\varepsilon_{\text{dd}})\,, \tag{2.102}$$

which can be accounted for at the mean-field level. In the local density approximation (LDA), one assumes that the density of the gas changes slowly enough such that the energy shift ΔE_{LHY} can be calculated locally. Furthermore, if one assumes that the zero-point fluctuations of the excitations do not significantly change the number of atoms in the condensate, i.e. $N_0 = N$, the change in the chemical potential can be incorporated in the Gross–Pitaevskii equation. This extended time-dependent GPE is given by

$$i\hbar\frac{\partial}{\partial t}\Psi(\mathbf{r},t) = \Bigg(-\frac{\hbar^2}{2m}\nabla^2 + V_{\text{trap}}(\mathbf{r},t) + g_s|\Psi(\mathbf{r},t)|^2 + \Phi_{\text{dd}}(\mathbf{r},t) + \frac{32}{3\sqrt{\pi}} g_s a_s^{3/2} Q_5(\varepsilon_{\text{dd}})|\Psi(\mathbf{r},t)|^3\Bigg)\Psi(\mathbf{r},t)\,, \tag{2.103}$$

while the extended time-independent version is given by

$$\mu_c \Psi(\mathbf{r}) = \left(-\frac{\hbar^2}{2m}\nabla^2 + V_{\text{trap}}(\mathbf{r}) + g_s|\Psi(\mathbf{r})|^2 + \Phi_{\text{dd}}(\mathbf{r}) + \frac{32}{3\sqrt{\pi}} g_s a_s^{3/2} Q_5(\varepsilon_{\text{dd}})|\Psi(\mathbf{r})|^3 \right) \Psi(\mathbf{r}) \,. \quad (2.104)$$

Further extensions to the GPE are also possible, e.g. one can take into account three-body interactions leading to losses or non-zero temperatures [203].

While the effect of the LHY term is usually negligible for mean-field-stable condensates, it becomes significant when one approaches collapse. As the LHY term is repulsive and scales more strongly with density than the interaction terms, it can prevent collapse. In some cases, as the condensate collapses and its density increases, the LHY term prevents further collapse (and a further density increase), creating a *quantum droplet*. Such quantum droplets can exist even in the absence of trapping [204]. Furthermore, if these droplets are synchronised in phase, this leads to an exotic phenomenon called *supersolidity* [80, 81, 205–213]: a superfluid with spontaneous spatial density modulation. This is a counter-intuitive state of matter that combines the dissipationless flow of a superfluid with the crystal-like periodic density modulation of a solid [214], achieved via the self-organisation of the excited ultracold dipolar gas into phase-synchronised quantum droplets [76, 77, 81, 215]. However, at these large densities, three-body losses can become significant, limiting the lifetime of quantum droplets and supersolids.

Quantum fluctuations also give rise to a depletion of the number of atoms in the condensate, which is called quantum depletion. (This is what was neglected when deriving the extended GPE.) For a homogeneous dipolar BEC, this is given by

$$\frac{\delta n}{n_0} = \frac{8}{3\sqrt{\pi}} \sqrt{n_0 a_s^3} Q_3(\varepsilon_{\text{dd}}) \,, \quad (2.105)$$

where $\delta n = n - n_0$ is the depletion of the condensate. This is normally on the few percent level.

3
Experimental apparatus

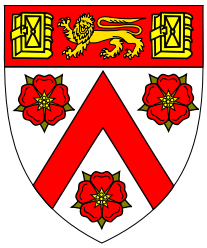

Trinity College, Cambridge. *Argent, a chevron between three roses Gules barbed and seeded proper and on a chief Gules a lion passant gardant between two closed books all Or.*

This Chapter describes the experimental apparatus, our platform for achieving Bose–Einstein condensation of erbium and other experiments. As several steps are required to produce a quantum-degenerate gas in a box trap, we first present a summary of the sequence of events that take place. After this overview, we turn to describing the individual components of our system. The principles behind their operation are presented with the help of Refs. 133 and 216 and key implementation details with the help of Ref. 104.

3.1 Overview of the experiment

To produce an ultracold erbium gas, we trap atoms in a vacuum chamber, shown in Fig. 3.1, in which they go through a sequence of cooling steps as shown in Fig. 3.2. Our vacuum system is designed with a 'three chamber' layout. Due to the high temperatures needed to achieve a significant vapour pressure of erbium, a solid erbium sample is first heated to 1150 °C in an effusion cell oven. This produces a collimated beam of hot atoms, directed along a tube that supports a large pressure differential and therefore protects the remaining sections from the potentially high outgassing rate of the heated parts of

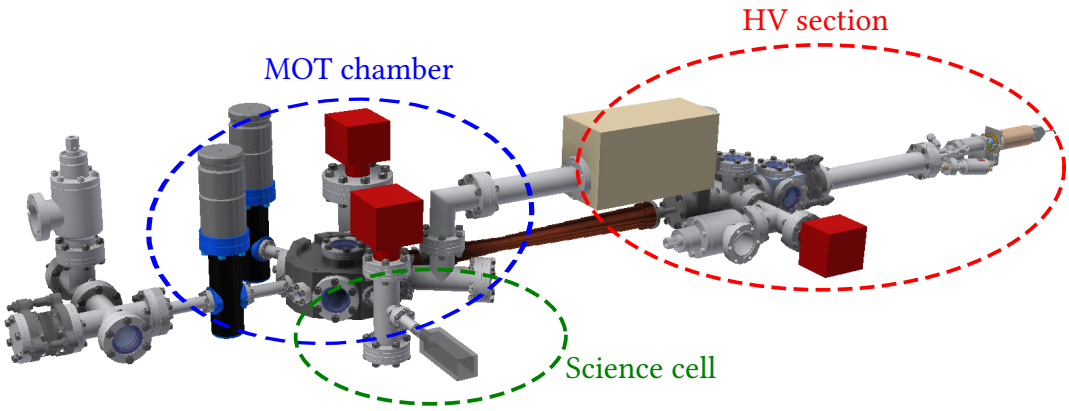

Figure 3.1. Vacuum system. The vacuum system consists of three separate parts. The high-vacuum (HV) section houses the erbium oven and the viewports for transversal cooling and probing of the atomic beam. It is separated from the remainder of the system with a low-conductance tube to prevent the high outgassing rate originating from the oven having an impact on the pressure in regions where the atom cloud is cooled. The MOT chamber is the central part of the experiment, where the main cooling steps take place (MOT and evaporative cooling). As it lacks good optical access for the box trap and measurement apparatus, we transport the atoms into an all-glass chamber, the 'science cell'. The box trap will be implemented there, along with high-resolution imaging and other beams to probe and manipulate the atoms. Adapted from Ref. 104.

the oven. As the atoms travel towards this tube, they first pass through a 2D molasses, operating in the transversal directions to the atom beam (this is called transversal cooling, TC). They are then slowed down along their propagating direction using a Zeeman slower (ZS), which has a capture velocity of ca. $350\,\mathrm{m\,s^{-1}}$ in our setup (i.e. it can slow atoms which enter the ZS below this speed), to less than $10\,\mathrm{m\,s^{-1}}$. These initial stages of cooling employ the broad 401 nm atomic transition.

The slowed atoms then arrive at the next chamber, where they are loaded into a narrow-line magneto–optical trap (MOT), operating on the 583 nm transition. As the atoms are simultaneously compressed and cooled (compressed-MOT configuration, cMOT) their temperature reduces to ca. 10 μK. They are then transferred into an optical dipole trap (ODT), formed using a high-power, 1030 nm laser beam. The main ODT beam (ODT1) is projected on the atoms using focus-tunable lenses, allowing the trap, formed in the focus of the beam, to be translated along the propagation direction. Prior to transport, an additional trapping beam (ODT2) is overlaid with the main beam, to enhance the axial confinement of the trap, in order to perform evaporative cooling and achieve Bose–Einstein condensation (detailed in Chapter 4).

In the next steps of the sequence, which have not been implemented yet, the atoms will be transported to a glass ('science') cell, providing better optical access. Here the atoms

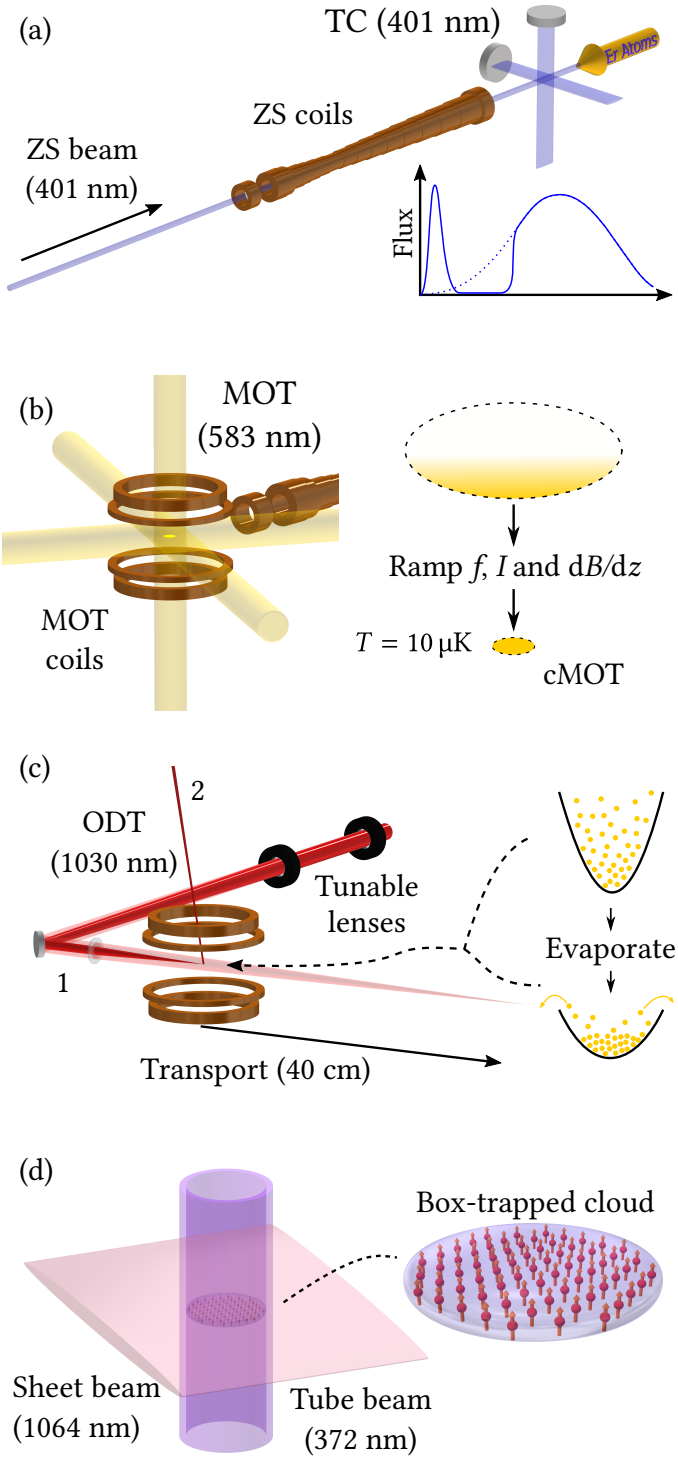

Figure 3.2. Experimental sequence. (a) After the atoms emerge from the oven, they are slowed in the transversal directions using an optical molasses (transversal cooling, TC). A Zeeman slower (ZS) then slows the atoms along their direction of propagation. (b) A narrow-line magneto–optical trap (MOT) is loaded with atoms slowed by the ZS. Low temperatures (ca. 10 μK) are obtained by ramping the trap into the compressed-MOT configuration (cMOT), which also creates better overlap with the optical dipole trap (ODT). (c) Atoms are transferred into the ODT, formed by a high-power 1030 nm laser projected onto the atoms by a pair of focus-tunable lenses. The trap can be moved between the two experimental chambers 40 cm apart by changing the focal length of the lens. A BEC is created after evaporative cooling in a crossed trap, formed by adding another ODT beam. (d) After transport, the atoms are transferred into the box trap made from a hollow, tubular, repulsive beam (372 nm) and an attractive, elliptical sheet beam (1064 nm). This creates a trap with a high aspect ratio and a homogeneous in-plane trapping potential. Adapted from Ref. 104.

will be transferred into a box trap, formed by the combination of a vertical, repulsive, hollow cylindrical beam, trapping the atoms in the horizontal direction, with a horizontal, attractive sheet beam for vertical confinement. The repulsive potential is provided by a 372 nm laser, which is blue-detuned compared to the dominant atomic transition, whereas the sheet beam will be provided by a (red-detuned) laser operating at 1064 nm.

3.2 Vacuum chamber

The vacuum chamber was designed to satisfy various requirements. First, we wanted to achieve ultra-high vacuum (UHV) conditions (pressures around 10^{-11} mbar) in sections of the apparatus where atoms are trapped, to ensure long atomic cloud lifetimes. Furthermore, we wanted to have a high atom flux for the quick and efficient loading of the MOT, trapping ca. 10^8 atoms within approximately 10 s. Finally, we wanted to have a high degree of optical access for the laser cooling and trapping beams, and for beams probing and manipulating the atoms in the experiments.

Our apparatus is a three-chamber system shown in Fig. 3.3. The first chamber is the high vacuum (HV) section, housing the erbium oven[1] (emitting a somewhat collimated beam of atoms) and transversal cooling. The atoms then move through a differential pumping tube while being slowed by a Zeeman slower. After this, they arrive at the MOT chamber,[2] where the MOT, the ODT and evaporative cooling are implemented. Finally, they are set to be transported to a rectangular glass cell,[3] providing a large degree of optical access for future experiments using this apparatus, which is where the optical box trap will be implemented. The MOT chamber and the glass 'science' cell are maintained at ultra-high vacuum (UHV) conditions in order to achieve long atomic cloud lifetimes (by minimising collisions with the background gas), and are therefore referred to as the UHV section. Besides these chambers, a fourth chamber is currently being constructed for adding potassium as a second atomic species to the experiment.

The effusion cell oven, shown in Fig. 3.4, consists of a tantalum crucible and two heated apertures which produce a collimated atomic beam. The crucible contains 8 g of solid erbium[4] in small pieces, and is heated to 1150 °C to increase the vapour pressure. The second aperture, which limits the flux of divergent atoms, has its temperature controlled independently by a separate heater filament and is heated to 1250 °C to prevent material build-up on the second aperture (this region is called the 'hot lip'). A third aperture

[1] Dual Filament Cell (DFC-40-10-WK-2B-SHE) by CreaTec.
[2] A custom 316LN stainless steel chamber by Scanwel.
[3] A 30 mm × 30 mm × 100 mm large cell with 2.5 mm thick walls, made from Borofloat glass by Precision Glassblowing.
[4] Distilled dendritic erbium from Alfa Aesar.

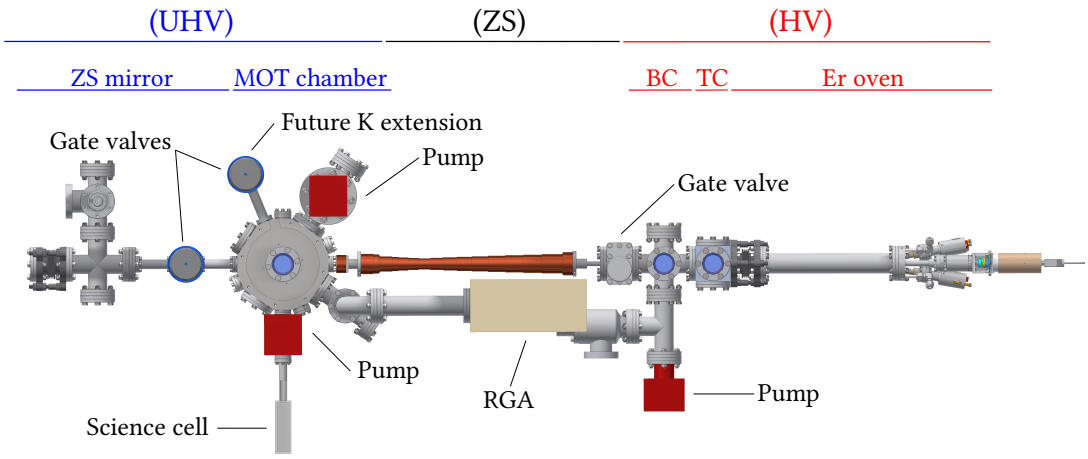

Figure 3.3. Vacuum chambers (top view). Atoms originate from the high vacuum section (HV) housing the erbium oven, a cube with viewports for transversal cooling (TC) and a cross with optical access for atomic beam characterisation (BC) and a vacuum pump. The Zeeman slower (ZS) tube maintains a pressure differential towards the ultra-high vacuum (UHV) section and carries the magnetic coils for operating the slower. The UHV section consists of the MOT chamber (for trapping and cooling the atoms), the science cell (for performing experiments) and connections for the ZS beam delivery system, the residual gas analyser (RGA) and the potassium extension (the potassium 2D MOT). Adapted from Ref. 104.

(a bored copper gasket) is installed after the oven to block those atoms which are too divergent to pass through the long and narrow differential pumping tube. The oven requires water cooling to prevent excessive heating of the surrounding chamber and to enhance the responsiveness of the temperature control loop. Cooling water is supplied via a closed-loop system built around a commercial water pump, and a secondary loop with an additional pump activated via an interlock circuit as a fail-safe. A shutter is also installed, which can block the atomic beam out of the oven. In our design, the oven is mounted on a port aligner, which allows any tilt of the oven to be corrected.

The oven is followed by a cubic chamber with viewports for transversal cooling (TC) and a 6-way cross with viewports for performing spectroscopy, atomic beam measurements and connecting the vacuum pump and the rough pumping port to the system. Finally, the atomic beam proceeds through the ZS differential pumping tube into the MOT chamber.

Our system uses three combined non-evaporable getter (NEG) and ion pump elements,[5] a choice motivated by their ability to pump all types of gases typically present

[5] Two NEXTorr D 100-5 pumps and a NEXTorr D 300-5 by SAES Getters.

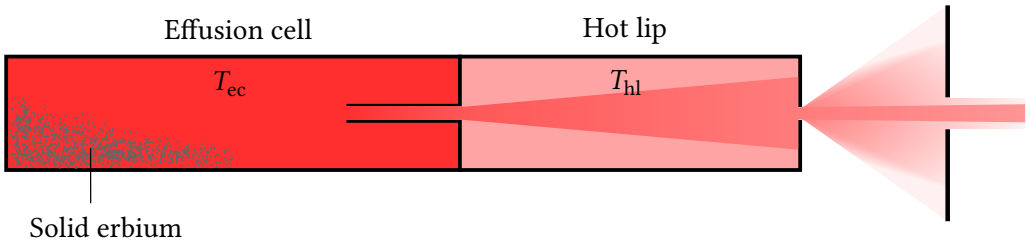

Figure 3.4. Effusion cell oven. A set of three apertures, two of them heated and integrated into the oven, are used to produce a collimated atomic beam. The first aperture is near the effusion cell where solid erbium is evaporated. The hot lip houses the second aperture, blocking the most divergent atoms, which is kept at a higher temperature to prevent material accumulation. A third aperture (a bored copper gasket) is mounted in front of the oven to block atoms which are too divergent to reach the MOT chamber (ca. 96% of the total flux). Adapted from Ref. 104.

in vacuum chambers. These pump molecular hydrogen especially efficiently, which is the dominant background component in stainless steel chambers. Pumps are placed next to the ZS tube, the MOT chamber and the glass cell, to provide the lowest possible pressure in the cell. Furthermore, a differential pumping tube, being able to support pressure differentials of at least two orders of magnitude, was installed between the oven and the other regions of the experiment where an ultra-high vacuum is needed. This was necessary to guard against the potentially high outgassing rates of the oven, which needs to be operated at a high temperature (above 1000 °C) to reach sufficient vapour pressure for producing an erbium atom beam with sufficient flux (cf. the melting point of erbium is 1529 °C).

The diameter of the tube presents a tradeoff between the possible pressure differential (the larger the diameter, the smaller the pressure differential) and the atom flux that can pass through the tube (the larger the diameter, the more atoms that can pass through). Furthermore, it affects the architecture of the ZS magnetic coils, as more windings or higher currents are required in order to produce the same field as the tube is made thicker. The placement and design of the oven apertures also needed to be optimised. These were achieved by numerical modelling of the pressures and the atomic beam [104], and by relying on the experience of groups working with similar elements [126, 135]. The atomic beam emerging from the oven was modelled using a Monte Carlo method to gain a better understanding of how the chamber geometry affects the atom flux reaching the MOT chamber and the efficiency of laser cooling. Pressures in the chamber were modelled assuming the molecular flow regime, estimating the equilibrium pressure in different locations in the system taking into account the pumping speeds, background gas sources and vacuum part conductances. As shown in Ref. 104, equilibrium pressures in the chambers

can be estimated by using the leak rates and pumping speeds to draw up an equivalent electric circuit where pressure corresponds to voltage, molecular conductivity and pumping speed to electric conductivity, molecular flow to electric current, volume to capacitance and the number of background molecules to charge. Besides determining the ZS tube diameter, this helped optimising the type, number and placement of pumps, and for our configuration the model gives 3×10^{-11} mbar in the MOT chamber and 1×10^{-11} mbar in the science cell. We are confident that we achieved similar conditions in our system as the pressure gauges of the ion pumps always maintained a reading below their detection level of 1×10^{-10} mbar. To minimise outgassing rates during usage, the apparatus was 'baked' (i.e. heated to a high temperature for a prolonged period of time) to increase the effusion rate of gases trapped in the steel before pumping it down completely.[6] For leak checking, the MOT chamber is connected to a residual gas analyser (RGA).[7]

The operating temperature of the oven also had to be decided. Using the atom beam model, it was chosen such that a sufficiently large atomic flux reaches the MOT chamber without draining the solid erbium sample too quickly. Too high a flux would also be problematic from the retro-reflected TC point of view. If the absorption rate of the photons is too high on their way to the reflecting mirror, the reflected beam would be significantly weaker, causing a transversal velocity bias. An aperture is therefore used to block those atoms which diverge too heavily to reach the MOT chamber but would increase the atom flux at the TC stage.

To implement the various laser-based methods for manipulating and probing atoms, a number of viewports[8] with appropriate anti-reflection coatings were installed on the chamber to provide optical access. The HV section has four ports for transversal cooling (for the two perpendicular retro-reflected beams) and three for probing the atomic beam. The MOT chamber has 14 ports, six used for the MOT beams, two for imaging, two for the Zeeman slower, two for the ODT and two spare (for observation). Furthermore, a separate rectangular glass cell was installed in which experiments will take place, as the large optical access leaves a large design flexibility for future experiments. (The drawback of this setup is that as there is a pump between the glass cell and the MOT chamber to reduce the pressure, atoms have to be transported over 40 cm between them.) Finally, the ZS laser beam is directed into the path of the atoms by a mirror[9] fixed to a port aligner within the vacuum chamber, as if it was directed directly through a viewport, the atoms

[6] Rough pumping was done with a turbo pump (Œrlikon Leybold TURBOVAC TW 70 H) backed with a scroll pump (Œrlikon Leybold SCROLLVAC SC 5 D).

[7] RGA100 by Stanford Research Systems.

[8] Manufactured by Torr Scientific.

[9] A custom, solid aluminium mirror with a UV-enhanced, polished and coated (F01 UV Enhanced Aluminium) surface by Thorlabs.

hitting the viewport would diminish its transparency. However, the deposition of erbium on the mirror does not affect its performance significantly [126, 135]. This section can be valved off in case mirror maintenance is required over the lifetime of the apparatus.

3.3 Magnetic field control

Magnetic fields are employed for a variety of purposes in our experiment. Laser cooling techniques (the MOT and the ZS) require particular field profiles to create a spatially variable detuning for the respective atomic transitions. Furthermore, homogeneous fields are required for maintaining the spin polarisation of atoms, tuning interactions, controlling the position of the MOT cloud and cancelling background fields. As erbium has a large number of Feshbach resonances (see Fig. 2.6), any spurious noise and offsets in the magnetic field need to be cancelled, with fields controlled ideally at the milligauss level.[10] As field stability is particularly important near the science cell and the MOT chamber, components in the immediate vicinity were manufactured using 316LN stainless steel, which has a particularly low magnetic permeability ($\mu_r - 1 \approx 10^{-3}$). Furthermore, the breadboard on which the experiment is assembled and the optical table on which it is mounted are made of non-magnetic materials (aluminium and 316L stainless steel, respectively).

We implemented a range of electromagnetic coils around the vacuum chamber. The coils installed so far are shown in Fig. 3.5 and their design details can be found in Refs. 104 and 217. The Zeeman slower system consists of six coils: three create the suitably-shaped magnetic field (the profile and bias coils), one adjusts the field at the end of the slowing trajectory and a pair of compensation coils cancel any residual field offsets and gradients in the MOT chamber originating from the rest of the ZS system.

The MOT chamber is equipped with four sets of coils. The gradient coils (in anti-Helmholtz configuration[11]) and the bias coils (in Helmholtz configuration) control the vertical field gradient and offset, respectively. The gradient is required for the operation of the MOT, whereas the offset is used to control the vertical position of the MOT and the scattering properties of the atoms (via tuning the field for exploiting Feshbach resonances). Furthermore, a pair of lower-current offset coils (for fine-tuning) and a single-loop RF antenna (coaxial with the other MOT coils, not shown in Fig. 3.5) are installed.

[10] This will be achieved through active feedback using the compensation coils, based on the signal from a high-resolution 3-axes magnetometer probe (Honeywell HMC2003).

[11] Two parallel coils are said to be in (anti-)Helmholtz configuration if the same current flows in the coils in the same (opposite) direction, producing a uniform field (gradient).

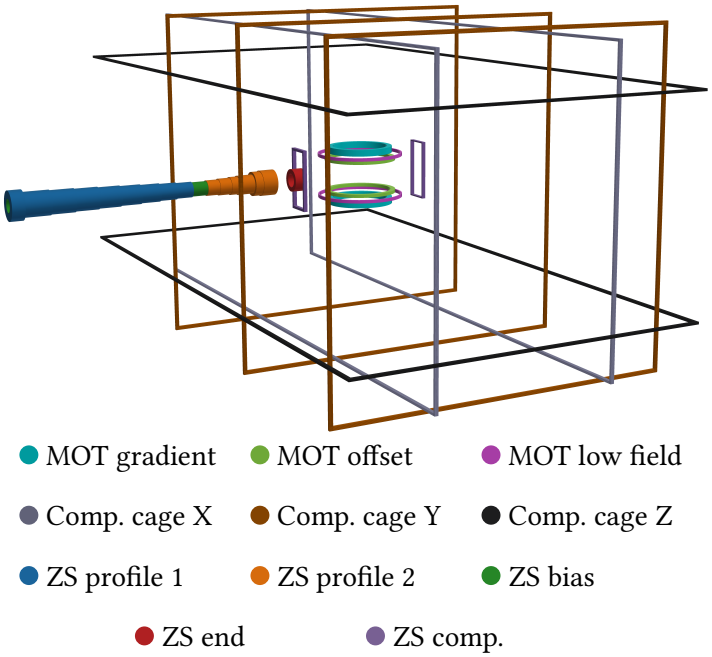

Figure 3.5. Electromagnetic coils. The Zeeman slower coil system includes three coils for producing the field profile ('profile 1', 'profile 2' and 'bias'), a coil for adjusting the field at the end of the slower and a pair of small compensation coils which mitigate any residual fields at the MOT position originating from the ZS. The MOT chamber is fitted with a pair of gradient coils for cooling and two offset coil pairs for fine and coarse field adjustments. Finally, a large compensation cage surrounds the chamber which can be used to apply uniform fields along all three directions. Adapted from Ref. 104.

The entire UHV section of the vacuum chamber is encapsulated in a compensation cage, consisting of seven large rectangular coils,[12] designed for the active compensation of magnetic field noise and for applying bias fields. As their geometry was optimised for field uniformity, yielding a large size, they have a slow frequency response, so a smaller (and faster) set of compensation coils for the science cell is being implemented. The coils around the science cell (currently being installed, not shown in Fig. 3.5) include two pairs of (large) coils to reach high-field Feshbach resonances of potassium and to rotate the dipoles to a given direction, and three pairs of (fast) coils to rapidly rotate the magnetic field, tuning the effective dipole strength [136].

[12] The direction along transport is addressed by three coils, as a pair of coils cannot provide a sufficiently uniform field along the full length between the MOT chamber and the science cell.

3.4 Laser cooling and trapping

Ultracold temperatures, required for reaching quantum degeneracy in atomic gases, are attained via standard laser cooling methods [5] to slow and capture hot atoms emerging from the oven maintained at 1150 °C, eventually producing a Bose–Einstein condensate below 1 μK. These systems and techniques are presented below.

3.4.1 Laser systems and locking

Blue laser

A blue (401 nm) laser is used for the initial laser cooling steps, transversal cooling (TC) and the Zeeman slower (ZS). This wavelength is used to address the corresponding atomic transition, whose large scattering rate makes it ideal for these initial stages (see Table 2.2). However, as the accompanying Doppler limit on the temperature is high (707 μK), further cooling is performed on a narrower transition.

Light is generated by a frequency-doubled titanium–sapphire laser[13] pumped by a diode-pumped solid-state laser[14] operating at 532 nm. The beams are delivered by a complex optical system, detailed in Ref. 104. It is worth mentioning that care needs to be taken to avoid high intensities on optical elements to avoid damage (which are particularly sensitive at near-UV wavelengths), and with setting up acousto-optic modulators (AOMs) as large powers can lead to heating-induced effects. (The AOMs are used for precise frequency and intensity control of the beams.) It is also worth mentioning that we encountered various issues with the 401 nm laser system, these are detailed in Ref. 104.

The frequency of the lasers employed for cooling need to be stabilised ('locked' to the atomic transition they address). This can be achieved by doing spectroscopy directly on the transition of interest using a sample of erbium (an *atomic reference*), or providing a stable external frequency reference (such as an optical cavity or a stabilised secondary laser source). In our case, there are two possibilities for an atomic reference: the beam emerging from the erbium oven and an erbium spectroscopic lamp.[15] While the former is readily accessible and has a relatively narrow distribution of transversal velocities, the ZS would interfere with any locking beams over the course of the experimental sequence. Furthermore, its relatively low atomic density makes locking to the narrower, 583 nm transition challenging. The other option consists of using a hollow cathode lamp (HCL)

[13] A SolsTiS system with an ECD-X module (the doubler) by M Squared Lasers.
[14] A 15 W Sprout-G by Lighthouse Photonics.
[15] Erbium vapour cells do not exist as impractically high temperatures are required for a significant vapour to be formed (the melting point of erbium is 1529 °C).

as a reference. This involves a pair of electrodes in an enclosure filled with a buffer gas, typically neon or argon. The cathode is shaped as a hollow cylinder, coated with erbium on the internal surface. When a high voltage (ca. 120 V DC for a typical current of 10 mA) is applied across the electrodes, ionised buffer gas atoms flow from the anode towards the cathode, sputtering erbium atoms off the coated surface into the region inside the cylinder. This provides an independent source of atoms, although with a larger velocity spread. Additionally, for this type of spectroscopy a see-through geometry is required, where a laser beam can be directed through the cylindrical cathode.

We currently use such an HCL[16] to perform modulation transfer spectroscopy [104, 218, 219] to lock our blue laser, which is relatively simple to implement and provides a Doppler-free, zero-offset signal which is also insensitive to the background magnetic fields. The locking signal is narrow enough compared to the width of the blue transition, but not compared to the yellow one. Over the course of building other techniques were explored (e.g. when the HCL malfunctioned), notably modulated fluorescence spectroscopy using the atomic beam in combination with a wavemeter, with a double feedback loop implemented on a RedPitaya data acquisition board[17] using the PyRPL software package [104, 220].

Yellow laser

A yellow (583 nm) laser system is used for the next stage of cooling, the magneto–optical trap (MOT). This narrower atomic transition is used for the MOT due to its lower Doppler temperature (4.5 μK).

Light is generated by a frequency-doubled diode laser with a tapered amplifier[18] and is delivered by an optical system described in Ref. 104. During the building of this system, special care needed to be taken to avoid reflections either from the MOT towards the laser locking system or vice versa, as these can be detrimental to the lock quality or the MOT lifetime (the lifetime of the atoms in the MOT).

Given the narrow linewidth of the transition (186 kHz) and the fact that the laser is principally used for cooling, achieving a robust and tight lock is especially important, as noise in the frequency leads to the heating of the atoms and a jitter of their position. Various techniques were tried to implement the yellow laser lock. We found that although modulation transfer spectroscopy (MTS) yields a reasonable lock using a good-quality hollow cathode lamp, it still leads to a considerable jitter in the MOT position. When we had to replace our HCL, we temporarily implemented shelving spectroscopy [104, 221], but the frequency stability was inferior even compared to the MTS locking setup.

[16] Manufactured by Photron.
[17] RedPitaya STEMlab 125-14.
[18] DL-TA-SHG pro by Toptica.

Therefore, an ultra-low expansion (ULE) optical cavity[19] is now used as a frequency reference, to lock the pre-doubled (1166 nm) light via the Pound–Drever–Hall (PDH) technique [222–224]. This relies on the fact that the sign of the *derivative* of the intensity of the light reflected from a cavity with respect to the light frequency changes as one crosses the resonance. Therefore, one can establish which side of the cavity resonance the light frequency is (and how far away it is), which provides a suitable locking signal that can be fed back to the laser. We vary the frequency of the laser (to establish the derivative) using an electro-optic modulator (EOM),[20] which is powered by an RF signal generator.[21] The reflected light from the cavity is captured using a photodiode[22] and it is fed into the laser locking module.[23] To minimise the dependence of the cavity length (and therefore its frequency) on temperature (which can vary slightly in the lab), the cavity spacer was manufactured out of a material whose coefficient of linear thermal expansion is zero slightly above room temperature (at 37 °C). Therefore, we stabilise the cavity at this temperature by using built-in electric heaters connected to an (external) PID controller to control the heating power, which monitors the cavity temperature using thermistors.

IR laser

The atoms are finally loaded into an optical dipole trap (ODT) for evaporative cooling. As the trap relies on the light being significantly red-detuned from the most prominent atomic transition (401 nm), the trap is formed using a 45 W, 1030 nm infrared (IR) fibre laser.[24] Besides trapping and cooling the atoms, it is also used for transporting them into the glass cell. Given this type of trap does not rely on light being resonant with an optical transition, this laser does not need to be locked.

3.4.2 Light scattering force

The operation of most laser cooling schemes can be understood via the light scattering force,[25] so we present a brief theoretical overview based on Ref. 216.

Let us consider a laser beam propagating in the opposite direction to that of a moving atom, which can absorb photons from this beam if it is resonant with an atomic transition. After the atom relaxes back to its ground state, the photon is re-emitted in a random

[19] Manufactured by Stable Laser Systems.
[20] NIR-MPX-LN-02-00-P-P-FA-FA by iXblue.
[21] DSG815 by Rigol.
[22] PDA05CF2 by Thorlabs.
[23] DigiLock 110 by Toptica.
[24] ALS-IR-1030-50-I-SF by Azurlight Systems.
[25] A notable exception is evaporative cooling, which will be discussed in Chapter 4.

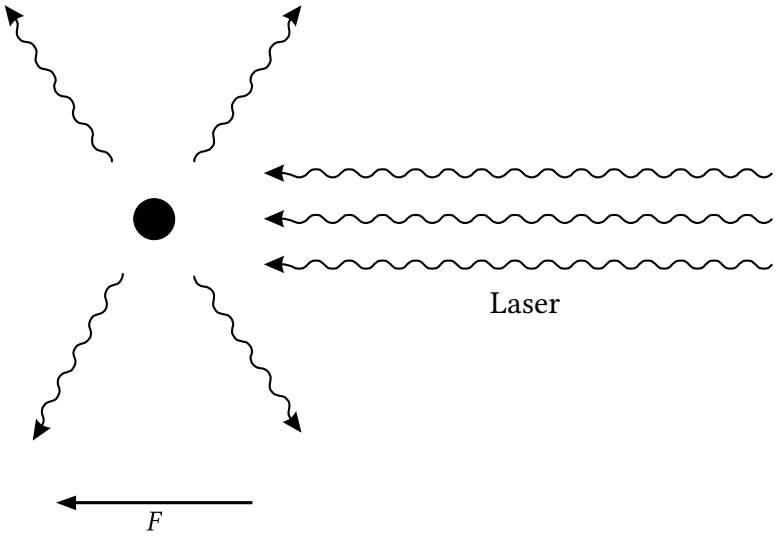

Figure 3.6. Scattering force. For an atom moving towards a laser beam, each absorbed photon gives the atom a kick in the direction opposite to its motion. As the photons are scattered in random directions, this results in a force on the atom. Figure adapted from Ref. 216.

direction. As the photons are emitted isotropically, their net effect on the momentum of the atom is zero. However, as the photons are absorbed from the direction opposite to the direction of propagation of the atom (the direction of propagation of the laser beam), the atom will slow down as shown in Fig. 3.6.

The scattering rate of the photons is given by

$$R = \frac{\Gamma}{2}\frac{s}{1+s+(2\delta/\Gamma)^2}\,, \tag{3.1}$$

where Γ is the transition rate of the respective atomic transition, $s = I/I_{\text{sat}}$ is the saturation parameter (I is the intensity of the laser and $I_{\text{sat}} = \pi hc\Gamma/3\lambda^3$ is the saturation intensity) and δ is the laser (angular) frequency detuning compared to the atomic transition. Given the atom absorbs a photon with momentum $\hbar k$ at each scattering event, where $k = 2\pi/\lambda$ is the wavenumber of the laser, the light scattering force is given by $F = \hbar kR$. It is worth noting that the closer the detuning is to zero, the larger the force is. The maximum force $F_{\text{max}} = \hbar k\Gamma/2$ occurs as $s \to \infty$, which corresponds to *half* the transition rate as half of the population is in the excited state at this point (this follows from Einstein's equations for radiation interacting with a two-level atom). Furthermore, note that as the atom slows down, the apparent frequency of the laser in the rest frame of the atom will change according to the Doppler effect. This effect can be compensated by changing either the detuning of the laser or the frequency of the atomic transition itself (e.g. via the Zeeman effect).

3.4.3 Transversal cooling

In order to slow the atoms in the transversal directions to their direction of propagation, to increase the atomic flux available for the ZS and passing through it, a 2D optical molasses [225] (transversal cooling, TC) is implemented after the oven using two orthogonal, elliptical, retro-reflected laser beams addressing the 401 nm transition.

While the light scattering force can be used to slow atoms with a counter-propagating laser beam, the difficulty in transversal cooling is that the atoms have a cylindrically symmetric velocity distribution around their average propagation direction (the ZS tube axis). However, if one takes two orthogonal pairs of counter-propagating laser beams, both red-detuned compared to the atomic transition (i.e. with a negative δ), one can create an 'optical molasses', slowing the atoms in all transverse directions (this is called transversal cooling). Detuning the beams is critical as this means that the atoms preferentially scatter photons from the counter-propagating laser source, as the magnitude of the apparent detuning reduces for the counter-propagating beam and increases for the co-propagating beam due to the Doppler effect. The resulting scattering rate is shown in Fig. 3.7. As the Doppler-shift of the angular frequency is given by kv_z for an atom propagating towards a laser beam with velocity v_z, the net force in the molasses is given by

$$F_{\mathrm{mol}} = F(\delta - kv_z) - F(\delta + kv_z) \approx 4\hbar k^2 s \frac{2\delta/\Gamma}{\left(1 + (2\delta/\Gamma)^2\right)^2} v_z \,, \tag{3.2}$$

where we assumed $s \ll 1$ (which is needed to be able to treat the beams acting independently on the atom) and $kv_z \ll \Gamma$ in the approximation. We see this is a viscous (linear damping) force, $F_{\mathrm{mol}} = -\alpha v_z$ if $\delta < 0$ (i.e. if the beam is red-detuned, matching our physical description), hence the term 'optical molasses'. This leads to an energy dissipation,

$$\left(\frac{\mathrm{d}E}{\mathrm{d}t}\right)_{\mathrm{mol}} = v_z F_{\mathrm{mol}} = -\alpha v_z^2 \,. \tag{3.3}$$

As we mentioned, the relaxing atom re-emits the photons in random directions. While these random kicks lead to $\Delta\langle v_z \rangle = 0$, they lead to heating over time t via $\Delta\langle v_z^2 \rangle \propto v_r^2 R t$, where $v_r = \hbar k/m$ is the recoil velocity. The randomness (Poissonian statistics) of photon absorption also leads to an increase in the velocity spread. These processes limit the temperatures that can be reached (the *Doppler limit*), and the apparent heating for two pairs of beams is given by [216]

$$\left(\frac{\mathrm{d}E}{\mathrm{d}t}\right)_{\mathrm{heat}} = \frac{10}{3} E_r R \,, \tag{3.4}$$

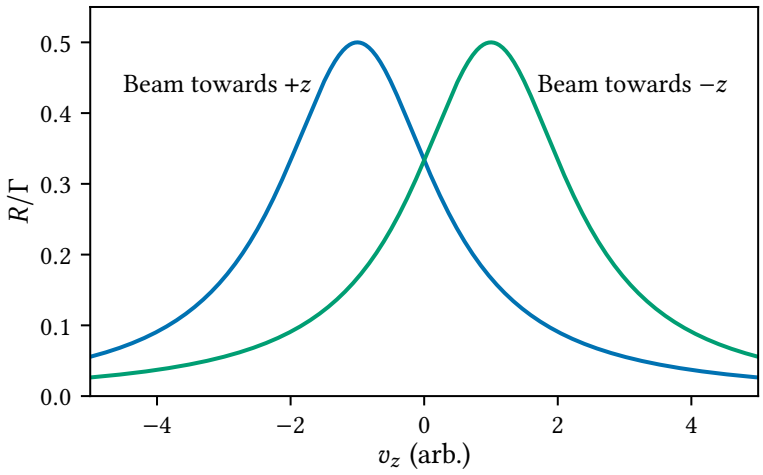

Figure 3.7. Operation of the optical molasses. Rate of absorption of photons from the two red-detuned laser beams as a function of the atomic velocity. The atoms always absorb more photons from the beam counter-propagating them.

where $E_r = mv_r^2/2$ is the recoil energy. Heating and cooling equilibrate when $(\mathrm{d}E/\mathrm{d}t)_{\mathrm{heat}} + (\mathrm{d}E/\mathrm{d}t)_{\mathrm{mol}} = 0$, which leads to

$$\left\langle v_z^2 \right\rangle_{\mathrm{min}} = \frac{5\hbar\Gamma}{24m} \frac{1 + (2\delta/\Gamma)^2}{-2\delta/\Gamma} \tag{3.5}$$

for $s \ll 1$. This function has a minimum at $\delta = -\Gamma/2$ with $\left\langle v_z^2 \right\rangle_{\mathrm{min}} = 5\hbar\Gamma/12m$. Using the equipartition theorem, $m\left\langle v_z^2 \right\rangle/2 = k_B T/2$, one can associate a temperature with this velocity, $T_{\mathrm{min}} = 5\hbar\Gamma/12k_B$. (For three pairs of beams $(\mathrm{d}E/\mathrm{d}t)_{\mathrm{heat}} = 4E_r R$, and the minimum temperature $T_D = \hbar\Gamma/2k_B$ is called the Doppler temperature.)

Given for our transition $\Gamma = 2\pi \times 29.5\,\mathrm{MHz}$, a frequency detuning of $-14.75\,\mathrm{MHz}$ should be optimal. Experimentally, we found optimal operation at $-12.5\,\mathrm{MHz}$, when the beam shape is close to the size of the atomic beam (the difference could be due to saturation effects), leading to a threefold increase in MOT loading numbers (compared to no TC). Optimisation of these parameters was aided by simulating their effect on the atomic beam, and the optical delivery system and key parameters can be found in Ref. 104. It is worth mentioning that a quarter-wave plate is inserted in the path of the beam to avoid the retro-reflected beam interfering with itself upon reflection. Furthermore, it is also interesting to note that as we have a problem with the anti-reflection coatings on the TC viewports, leading to a significant loss of light intensity upon passage through the chamber, we use slightly converging beams to re-establish the balance of forces in the TC setup, to avoid imprinting a transversal velocity bias. The faulty viewports are set to be replaced in the future.

3.4.4 Zeeman slower

As we have demonstrated earlier, photons scattered off atoms can slow them down. However, it takes time for enough photons to scatter to achieve the desired cooling. In many cold atom experiments, including ours, atoms are sourced from a high-temperature oven which emits a collimated beam. While the apertures (acting as collimators) ensure that the transversal velocity of the atoms is relatively small, due to the high temperature of the oven, the velocity of the atoms is ca. 420 m s^{-1} on average. Therefore, while sufficient transversal cooling can be achieved by just passing through the TC (the beam width along the propagation direction is 17 mm), it takes a considerable time (and due to their high velocity, a considerable path length) to slow the atoms sufficiently in the axial direction. This is especially important as the next cooling stage, the MOT, has a capture velocity of ca. 10 m s^{-1}.

Another important aspect is that the change in the Doppler shift of the slowing atoms (ca. 1 GHz at the initial average velocity, decreasing towards zero as the atoms stop) is much larger than the natural linewidth (29.5 MHz) of the 401 nm transition. Therefore, in order to ensure that the laser is resonant with the atomic transition during the whole duration of slowing, we continuously change the atomic transition frequency along the path to compensate the changing Doppler shift as the atoms slow down. This is achieved via the Zeeman effect, by using a suitably-shaped magnetic field $B(x)$, hence the term Zeeman slower [226].

To achieve a constant force on the atoms, the apparent detuning needs to be kept constant (assuming that other parameters, like the intensity of the beam are constant along the path, which is only an approximation due to the scattering of photons). Therefore, the change in the Zeeman shift needs to exactly cancel the change in the Doppler shift. As in our assumption s is constant (uniform profile, collimated beam, negligible absorption), the work–energy principle after the (constant) scattering force F acted over a path x on a particle with speed v_0 can be written as

$$Fx = -\hbar k R x = \frac{1}{2} m v^2(x) - \frac{1}{2} m v_0^2 , \tag{3.6}$$

resulting in

$$v(x) = v_0 \sqrt{1 - \frac{x}{x_0}} , \tag{3.7}$$

where

$$x_0 = \frac{m v_0^2}{2\hbar k R} = \frac{m v_0^2}{\hbar k \Gamma s} \left(1 + s + \left(\frac{2\delta}{\Gamma} \right)^2 \right) . \tag{3.8}$$

We see that x_0 is the length at which the atoms stop; if they are subjected to a force beyond this point, they will turn back. Furthermore, given a fixed x_0, the capture velocity above which atoms will not be stopped completely by the slower (up to the Doppler limit) is v_0.

The laser frequency in the rest frame of the atoms can be written as $\omega + kv$, where the laser frequency in the lab frame is ω and the Doppler shift is kv. Furthermore, with the Zeeman shift taken into account, the frequency of the atomic transition can be written as $\omega_0 + \mu_Z B/\hbar$, where ω_0 is the frequency of the unperturbed atomic transition and $\mu_Z = \left(m_e g_e - m_g g_g\right) \mu_B$ is the Zeeman factor given by the properties of the two states involved in the atomic transition (see Eq. (2.9)). Therefore, the total apparent detuning, which we aim to keep constant along the length of the slower, is given by $\delta = (\omega + kv) - (\omega_0 + \mu_Z B/\hbar) = \delta_0 + kv - \mu_Z B/\hbar$, where $\delta_0 = \omega - \omega_0$ is the laser detuning without any Zeeman shifts for an atom at rest in the lab frame. Therefore, the magnetic field profile of an ideal Zeeman slower is given by

$$B(x) = \frac{\hbar}{\mu_Z}(\delta_0 - \delta + kv(x)) = \frac{\hbar}{\mu_Z}\left(\delta_0 - \delta + kv_0\sqrt{1 - \frac{x}{x_0}}\right). \tag{3.9}$$

This magnetic field can be created with a suitable set of coils. In our system, we use one coil to generate the constant offset (the 'bias') and two to generate the spatially varying part (the 'profile').

When designing a realistic Zeeman slower, a number of other considerations need to be taken into account [104]. First, the total detuning δ should be negative, for the faster atoms to encounter a larger force and the slower atoms to feel a smaller force, leading to a robust operation. However, the larger the detuning, the larger the slowing path x_0 is, so a sensible compromise is needed. Furthermore, as the end of the Zeeman slower is not at the centre of the MOT and the light can turn atoms back after they leave the ZS, it is designed such that the atoms are not stalled completely. (The field at the end of the ZS can be adjusted with a separate 'end' coil, providing a kick at the end, and a pair of coils are used to cancel the field of the Zeeman slower at the MOT.) The exit velocity needs to be chosen such that the MOT can capture the atoms, but it is not so small that the atoms eventually turn back or that the atom beam expands so significantly due to the remaining transversal velocity component that many atoms hit the wall of the ZS tube or evade the MOT upon exit. Another consideration that is worth taking into account is that the magnitude of the magnetic field should be minimised in the ZS (to avoid large coils or high currents), so the bias needs to be chosen suitably. When the bias is negative and a zero-crossing of $B(x)$ occurs during the slowing path (as the profile is always positive), the ZS is said to be in the 'spin-flip' configuration (which is the case

in our system). While this is advantageous from the coils' point of view, care needs to be taken with the field profile as the atoms can become magnetically depolarised in the vicinity of the zero-crossing. The existence of this zero-crossing is also the reason behind using two coils for creating the field profile, one on each side of the crossing, to create fields with opposite directions.

In our case, the Zeeman slower has a length of 42 cm and its capture velocity is ca. 350 m s^{-1}. It slows atoms down to ca. 10 m s^{-1} with a slow-atom flux of 5×10^7 s^{-1}.

3.4.5 Magneto–optical trap

After the atoms exit the Zeeman slower, their velocity is on the order of 10 m s^{-1}. Further cooling of the atoms takes place using the 583 nm transition due to its narrower linewidth leading to a lower Doppler temperature ($T_D = 4.5\,\mu$K). Furthermore, as the forces acting in the TC and the ZS depend only on the velocity of the atoms but not on their position, they provide slowing but not trapping. Therefore, a magneto–optical trap (MOT) [227] is implemented on this transition, which also provides trapping besides cooling. This is important as the stable spatial manipulation of the cloud is necessary to optimise loading into the optical dipole trap, where evaporative cooling (the final cooling stage) takes place. Due to the high scattering rates involved, the MOT cloud is visible to the naked eye.

The MOT can be thought of as an enhancement of the combination of the optical molasses and the Zeeman slower. Compared to the ZS, the magnetic field here is used not to provide a constant detuning of the atomic transition but a position-dependent one, leading to a position-dependent light scattering force.

The operation of the MOT is shown in Fig. 3.8. Its basis is a 6-beam optical molasses (three counter-propagating pairs of beams with all pairs orthogonal to each other), and the operation principle is similar: the atoms preferentially absorb photons from their direction of propagation, i.e. from the beam propagating towards them. However, to achieve spatial trapping, we also need this to happen not only based on their velocity but also on their position, i.e. we need the atoms to always be pushed towards the intersection of the beams, the centre of the trap. This can be achieved by using a magnetic field, which acts on the different transitions differently. In our case, the 583 nm light addresses transitions from the ground state $\left|J, m_J\right\rangle = |6, -6\rangle$ to the excited states $|7, -7\rangle$, $|7, -6\rangle$ and $|7, -5\rangle$, which require a photon to be absorbed with polarisation σ^-, π and σ^+, respectively.[26] The Zeeman shift of the transitions is given by $\left(m_e g_e - m_g g_g\right) \mu_B B/\hbar$ as earlier, where the

[26] This can be seen by recalling that the angular momentum projection is $J_z = -\hbar$ for a σ^--polarised photon, $J_z = 0$ for a π-polarised photon and $J_z = +\hbar$ for a σ^+-polarised one.

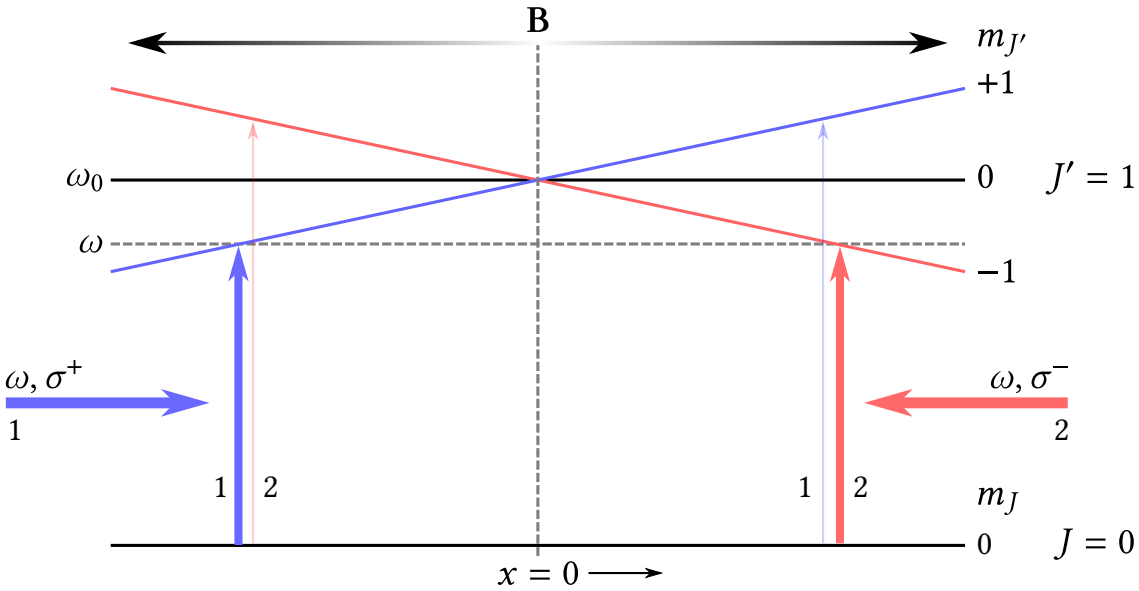

Figure 3.8. Operation of the MOT. Atoms are excited from the ground state $|J, m_J\rangle = |0, 0\rangle$ to the excited states $|J', m_{J'}\rangle = |1, -1\rangle$ (red), $|1, 0\rangle$ (black) and $|1, 1\rangle$ (blue). Due to the magnetic field gradient and the red-detuned laser beams, atoms at $x < 0$ are more likely to absorb photons from the σ^+-polarised beam (1), propagating in the $+x$ direction, while atoms at $x > 0$ are more likely to scatter from the σ^--polarised beam (2). This leads to a force with spatial dependence, confining the atoms near the origin. The same argument can be adopted for the ground state $|6, -6\rangle$ and the excited states $|7, -7\rangle$, $|7, -6\rangle$ and $|7, -5\rangle$. Figure adapted from Ref. 217.

Landé g-factor is given by Eq. (2.5) for the ground state and by Eq. (2.7) for the excited states. Given the g-factors of the excited states are different, the respective transitions will be detuned differently at the same magnetic field. Furthermore, for our case the Zeeman factors of the σ^- and σ^+ transitions have an opposite sign, so the detunings of the different transitions will change in the opposite sense.

Let us now focus on the dynamics in one spatial direction. A linearly changing magnetic field (a constant magnetic field gradient) with $B = 0$ at the trap centre will mean that a given transition will match a given detuning only on one side of the trap (the other transition will be resonant on the opposite side of the trap), given the B field is the opposite on the two sides of the trap and the transitions change with the field in the opposite sense. If we choose one of the laser beams to be σ^--polarised and the other to be σ^+ (both red-detuned compared to $B = 0$), this means that the atoms will absorb photons preferentially from opposing beams on each side of the trap (from the one corresponding to the transition which is resonant, as the transition corresponding to the opposite beam will be far-detuned and hence will scatter fewer photons), leading to spatial trapping.

While we only considered one direction for simplicity so far, the argument can be

extended to all three dimensions. In the most common configuration (including in our implementation), one pair of beams is vertical (the z-axis) and the others are horizontal, and the magnetic field gradient (generated using a pair of coils in the anti-Helmholtz configuration) is that of a quadrupole field,

$$\mathbf{B}(x, y, z) = G\left(-\frac{1}{2}\mathbf{x} - \frac{1}{2}\mathbf{y} + \mathbf{z}\right), \tag{3.10}$$

where $G = |\partial B/\partial z|$ is the magnitude of the field gradient. Without going into details, one can qualitatively argue some important effects. The size and shape of the MOT cloud can be estimated based on the properties of the surface on which the laser is resonant with a transition. In our case this is an ellipsoid, and as it can be seen from Fig. 3.8, if the magnitude of the detuning is decreased, the surface shrinks, so the cloud is compressed. Furthermore, it can be seen that if a constant offset field (a 'bias') is applied, the position of the MOT cloud will be shifted (as the centre of the trap is defined by the position where the field is zero).

The description of a narrow-line MOT (where the frequency corresponding to the single-photon recoil energy is on the order of the linewidth) is somewhat more involved than the general argument above [228, 229]. As in our case the MOT is operated in the regime where $\delta \gg \Gamma$, besides the light scattering force, the magnetic force and gravity are also significant. The force on a magnetic dipole in a non-uniform field is given by $\mathbf{F} = \nabla(\boldsymbol{\mu} \cdot \mathbf{B})$ [230]. Given the magnetic dipole moment of our atoms is $\boldsymbol{\mu} = m_g g_g \mu_B \hat{\mathbf{B}}$ (cf. Eq. (2.4)), the force in the ground state ($m_g = -6$) is given by

$$\mathbf{F} = -\mu \nabla |\mathbf{B}|\,, \tag{3.11}$$

where $\mu = \left|m_g g_g \mu_B\right|$. Since the atoms are in a high-field-seeking state, this force is anti-trapping. Along the vertical direction, in the region below the centre of the MOT, this leads to an effective gravitational force of

$$mg' = mg + G\mu\,, \tag{3.12}$$

where g (g') is the (effective) gravitational acceleration and m is the mass of the atoms as before. This leads to a sagging of the MOT under gravity, as the light scattering force needs to balance the magnetic and gravitational forces (i.e. the vertical position of the MOT can be controlled via the light intensity and the field gradient besides the field offset). Using a simple model [231], one can take the cloud to only interact with the beam pointing upwards, which has the benefit of pumping (spin-polarising) the atoms to the

ground state ($m_J = -6$) if the beam is σ^--polarised (i.e. the vertical magnetic field gradient is negative). Furthermore, it can be shown that the equilibrium temperature of the MOT is independent of the detuning to first approximation and that it increases with s [228].

With all these in mind, MOT loading takes place in three stages [104]. First, the ZS and the MOT beams are switched on at high intensity (for a suitably high MOT capture velocity), with the MOT light strongly red-detuned (such that the atoms sag in the MOT for spin-polarisation and to avoid interaction with the ZS beam). We typically wait for 15 s for the MOT to load. Once loading is completed, the ZS light and coils are turned off, and the MOT is held in place for a hold time of 1 s (in which the cloud equilibrates from the perturbation caused by the varying magnetic fields upon the turn-off of the ZS coils). Finally, the magnitude of the MOT light detuning and intensity are decreased in parallel to cool and compress the MOT, and a magnetic field bias is introduced vertically to raise the cloud to the height of the optical dipole trap beam. Details of this procedure and the optical system can be found in Ref. 104, but it is worth emphasising that the stability of MOT loading and cMOT (compressed MOT) position depends critically on the stability of the frequency and the polarisation of the laser beams. We typically achieve 10^8 atoms in the cMOT at a temperature of 10 μK (measured with the time-of-flight technique presented in §3.6.2). The lifetime of the MOT was measured to be 40 s, which is sufficiently longer than the period of the experimental cycle (ca. 20 s) and confirms we have a suitably high vacuum.

3.4.6 Optical dipole trap

The final stage of cooling makes use of the optical dipole trap (ODT) [43, 232], which relies on another aspect of the atom–light interaction, the AC Stark effect. As a simple but intuitive model, let us consider a neutral atom placed in an oscillating electric field $\mathbf{E}(\mathbf{r}, t)$. In this field, the centres of mass of the positive and negative charges are periodically separated by the field (they are dragged in opposite directions), which means that an oscillating electric dipole $\mathbf{p}(\mathbf{r}, t)$ is induced,

$$\mathbf{p}(\mathbf{r}, t) = \alpha \mathbf{E}(\mathbf{r}, t)\,, \tag{3.13}$$

where α is called the (complex) polarisability of the atom. Writing the dipole moment this way is justified by the fact that the wavelength of the laser is usually much larger than the atomic size. The atom can then be treated as a dipole which interacts with the driving field, with the interaction energy given by

$$U_{\text{dip}}(\mathbf{r}) = -\frac{1}{2}\langle \mathbf{p}(\mathbf{r}, t) \cdot \mathbf{E}(\mathbf{r}, t) \rangle = -\frac{1}{2\varepsilon_0 c} \operatorname{Re}(\alpha) I(\mathbf{r})\,, \tag{3.14}$$

where c is the speed of light and we used the fact that the field intensity can be written as $I(\mathbf{r}) = \varepsilon_0 c|\mathbf{E}(\mathbf{r})|^2/2$ where $\mathbf{E}(\mathbf{r})$ is the field amplitude. The averaging is justified by the fact that the field oscillates much faster than the frequencies of atomic motion. Furthermore, it is worth noting that the factor of $1/2$ is necessary as the dipole is an induced one, not a stationary one. Given the force on (any) electric dipole is $\mathbf{F} = (\mathbf{p} \cdot \nabla)\mathbf{E}$ [230], it can be shown that the interaction energy $U_{\text{dip}} = -\int \mathbf{F} \cdot \mathrm{d}\mathbf{r} = -\mathbf{p} \cdot \mathbf{E}/2$ for $\mathbf{p} \propto \mathbf{E}$ (the constant would be different for a different relationship between the field and the induced dipole).

When a dipole is driven below resonance (the field is red-detuned), it oscillates in phase with the field, but when it is driven above resonance (the field is blue-detuned), it oscillates out of phase. Below (above) resonance the interaction energy is negative (positive), so the atom is attracted to (repelled from) the region with the highest electric field intensity. A trap will be created for negative interaction energies, where the atoms are drawn towards the highest intensity (the focus of a Gaussian laser beam).

This can also be understood in a slightly more rigorous treatment. In the 'dressed-state' picture, the eigenstates of the atoms in the electric field are the unperturbed eigenstates mixed with small proportions of other states. All the states corresponding to transitions for which the laser is blue-detuned (red-detuned) increase (decrease) the energy of the atom. For a two-level system, if the light is far-detuned compared to a transition, the energy shift can be written as [216]

$$\Delta E = U_{\text{dip}} \approx \frac{\hbar\Gamma^2}{8\delta} s = -\frac{1}{2\varepsilon_0 c} \operatorname{Re}(\alpha) I(\mathbf{r}) \,. \tag{3.15}$$

The atoms also scatter the photons, which leads to a momentum transfer resulting in heating and atom loss. The scattering rate (Eq. (3.1)) for a far-detuned field is given by

$$R \approx \frac{\Gamma^3}{8\delta^2} s \,. \tag{3.16}$$

We see that the interaction energy is $U_{\text{dip}} \propto I/\delta$, while the scattering rate is $R \propto I/\delta^2$, so scattering (heating) is suppressed for large detunings given the same trap depth.

In general, the polarisability α has a tensorial character and depends on the atomic species, the light frequency and the angles of polarisation θ_p and the propagation direction θ_k of the electric field with respect to the quantisation axis (the direction of the magnetic field B). This is due to the fact that it accounts for the effects from all transitions, some of which are polarisation-dependent and are associated with anisotropic orbitals. The polarisability can be separated into three components (called scalar, vector and tensor),

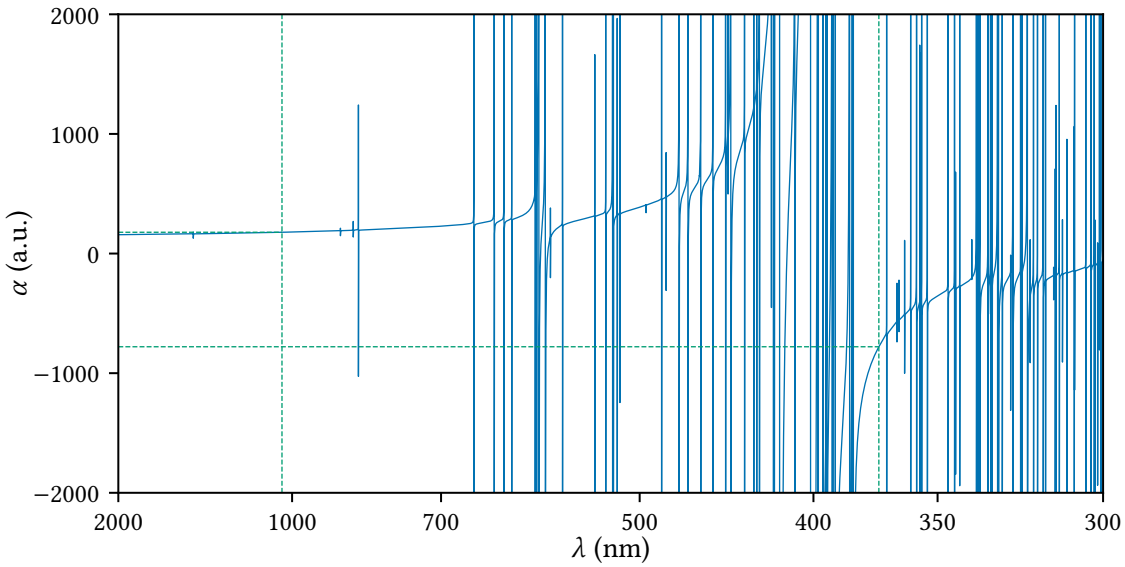

Figure 3.9. Polarisability of erbium. For the horizontal ODT beam, $\theta_k = \theta_p = 90°$ and this configuration is shown in the figure, using data from Ref. 234. We see that the many available transitions induce resonances in the polarisability; the finite peak heights associated with some transitions are due to the finite sampling rate of λ in the figure. The polarisability is highlighted for 1030 nm ($\alpha = 178$ a.u., i.e. a trap is formed at this wavelength, used for evaporation and transport), and for 372 nm ($\alpha = -779$ a.u., i.e. the potential is repulsive at this wavelength, used for the radial 'wall' of the optical box trap). Note that the horizontal axis scaling is reciprocal (goes as $1/\lambda$). Adapted from Ref. 104.

and can be written as [65]

$$\alpha(\omega) = \alpha_{\text{scal}}(\omega) + \mathcal{E}\cos\theta_k \frac{m_J}{2J}\alpha_{\text{vect}}(\omega) + \frac{3m_J^2 - J(J+1)}{J(2J-1)}\frac{3\cos^2\theta_p - 1}{2}\alpha_{\text{tens}}(\omega)\,, \tag{3.17}$$

where $\left|J, m_J\right\rangle$ is the (electronic) quantum state of the atom ($|6, -6\rangle$ for the ground state of erbium) and $\mathcal{E}$ is the ellipticity of the light polarisation. The polarisability has been calculated and measured[27] at different wavelengths for erbium [233, 234] and is shown in Fig. 3.9 for $\theta_p = \theta_k = 90°$. From the figure it is clear that an optical dipole trap is created for $\lambda = 1030$ nm, whereas a repulsive optical potential can be set up using light close to 372 nm.

To proceed with characterising the trap, we need to consider the intensity distribution of the light. In most cases a Gaussian beam is used for trapping, whose intensity distribution is given by

$$I(x, y, z) = \frac{2P}{\pi w_x(z) w_y(z)} e^{-2\left(\frac{x^2}{w_x^2(z)} + \frac{y^2}{w_y^2(z)}\right)}\,, \tag{3.18}$$

[27] The atomic unit (a.u.) of polarisability is defined as 1 a.u. $= 4\pi\varepsilon_0 a_0^3 = 1.65 \times 10^{-41}\,\text{C}^2\,\text{m}^2\,\text{J}^{-1}$.

where P is the total power in the beam and $w_x(z)$ and $w_y(z)$ are the beam radii (the beam propagates in the z-direction). For atomic clouds which are small compared to the size of the laser beam, we can use a harmonic approximation of the dipole potential (Eq. (3.14)) to extract the trap frequencies, which leads to

$$U_{\text{dip}} \approx U_{\text{dip}}(0) + \frac{1}{2}\frac{\partial^2 U_{\text{dip}}}{\partial x^2}\bigg|_0 + \frac{1}{2}\frac{\partial^2 U_{\text{dip}}}{\partial y^2}\bigg|_0 + \frac{1}{2}\frac{\partial^2 U_{\text{dip}}}{\partial z^2}\bigg|_0 \tag{3.19}$$

$$= -\frac{1}{2\varepsilon_0 c}\,\text{Re}(\alpha)\frac{2P}{\pi w_x w_y}\left(1 - \frac{2x^2}{w_x^2} - \frac{2y^2}{w_y^2} - \frac{1}{2}z^2\left(\frac{1}{z_{Rx}^2} + \frac{1}{z_{Ry}^2}\right)\right), \tag{3.20}$$

where w_x and w_y are the beam waists in the x and y direction, respectively, and z_{Rx} and z_{Ry} are the respective Rayleigh ranges.[28] We see the trap depth is given by

$$U_0 = U_{\text{dip}}(0) = -\frac{1}{2\varepsilon_0 c}\,\text{Re}(\alpha)\frac{2P}{\pi w_x w_y}, \tag{3.21}$$

which is sometimes expressed as a temperature, $T_0 = U_0/k_B$. By comparing the trapping potential (Eq. (3.20)) to the potential of a simple harmonic oscillator, $m(\omega_x^2 x^2 + \omega_y^2 y^2 + \omega_z^2 z^2)/2$, the trapping frequencies can be extracted as

$$\omega_x = \sqrt{\frac{4U_0}{m w_x^2}} = \sqrt{\frac{4\alpha P}{m\varepsilon_0 c\pi w_x^3 w_y}}, \tag{3.22}$$

$$\omega_y = \sqrt{\frac{4U_0}{m w_y^2}} = \sqrt{\frac{4\alpha P}{m\varepsilon_0 c\pi w_x w_y^3}}, \tag{3.23}$$

$$\omega_z = \sqrt{\frac{2U_0}{m z_{R_{\text{eff}}}^2}} = \sqrt{\frac{\alpha P\lambda^2\left(w_x^4 + w_y^4\right)}{m\varepsilon_0 c\pi^3 w_x^5 w_y^5}}, \tag{3.24}$$

where

$$z_{R_{\text{eff}}} = \frac{z_{Rx} z_{Ry}}{\sqrt{\frac{1}{2}\left(z_{Rx}^2 + z_{Ry}^2\right)}} = \sqrt{\frac{2}{w_x^4 + w_y^4}}\frac{w_x^2 w_y^2 \pi}{\lambda} \tag{3.25}$$

is the effective Rayleigh range.

Besides changing the size and the power of the ODT beam directly, the trap size and depth can also be changed by periodically moving (dithering) the trapping beam. If the beam is dithered at a frequency much higher than the trapping frequency, the atoms are not able to follow the motion of the light and instead experience an effective, time-averaged potential. Creating such an elongated trap is useful as a larger overlap

[28] The Rayleigh range is defined as $z_R = \pi w_0^2/\lambda$ for a beam with waist w_0.

with the MOT cloud can be achieved, enabling a higher ODT loading efficiency.

If the position of the beam is given by the periodic functions $x_D(t), y_D(t)$ with period T, the time-averaged potential is given by

$$U(x,y) = \frac{U_0}{T}\int_0^T e^{-\left(\frac{2(x_D(t)-x)^2}{w_x^2}+\frac{2(y_D(t)-y)^2}{w_y^2}\right)}\,\mathrm{d}t\,, \tag{3.26}$$

where U_0 is the undithered trap depth and $w_{x,y}$ are the undithered beam waists. To match the previous configuration, we would like the effective potential to be a Gaussian widened in one direction (let this be x, without loss of generality)—this means setting $U(x,y)$ to be a Gaussian with $w_x' = cw_x$. Given the total power is the same in the dithered beam, from Eq. (3.21) we see $U_0' = U_0/c$. Equation (3.26) can be solved in this case to yield[29]

$$x_D(t) = \begin{cases} \sqrt{\frac{c^2-1}{1}}w_x \operatorname{erf}^{-1}\left(\frac{4t}{T}-1\right) & \text{if } 0 \le t \le \frac{T}{2}\,, \\ \sqrt{\frac{c^2-1}{1}}w_x \operatorname{erf}^{-1}\left(3-\frac{4t}{T}\right) & \text{if } \frac{T}{2} \le t \le T\,, \end{cases} \tag{3.27}$$

$$y_D(t) = 0. \tag{3.28}$$

However, this dithering function diverges at $t = 0$ and $T/2$. Ref. 235 shows that it can be approximated as

$$x_D(t) = \begin{cases} \frac{2cw_x}{\pi}\arcsin\left(\frac{4t}{T}-1\right) & \text{if } 0 \le t \le \frac{T}{2}\,, \\ -\frac{2cw_x}{\pi}\arcsin\left(\frac{4t}{T}-3\right) & \text{if } \frac{T}{2} \le t \le T\,, \end{cases} \tag{3.29}$$

$$y_D(t) = 0. \tag{3.30}$$

A plot of the dithering functions and the resulting potentials can be seen in Fig. 3.10. While Eq. (3.26) cannot be solved analytically for arbitrary $U(x,y)$, it can be solved numerically to yield an arbitrary trap, e.g. one which is exactly harmonic towards its centre portion.

In our system, an acousto-optic modulator[30] (AOM) is used to dither the beam, which can also be used to control the laser power. An AOM works by inducing a sound wave travelling through a crystal, which periodically alters the refractive index of the medium. This acts as a layered structure from which light can be reflected into Bragg peaks. The first-order deflection angle of the AOM is governed by the Bragg condition

$$2\frac{v_s}{f_{\mathrm{AOM}}}\sin\theta_d = \lambda\,, \tag{3.31}$$

[29] I thank Tevž Lotrič for pointing out this solution.
[30] AOMO 3080-199 by Gooch & Housego.

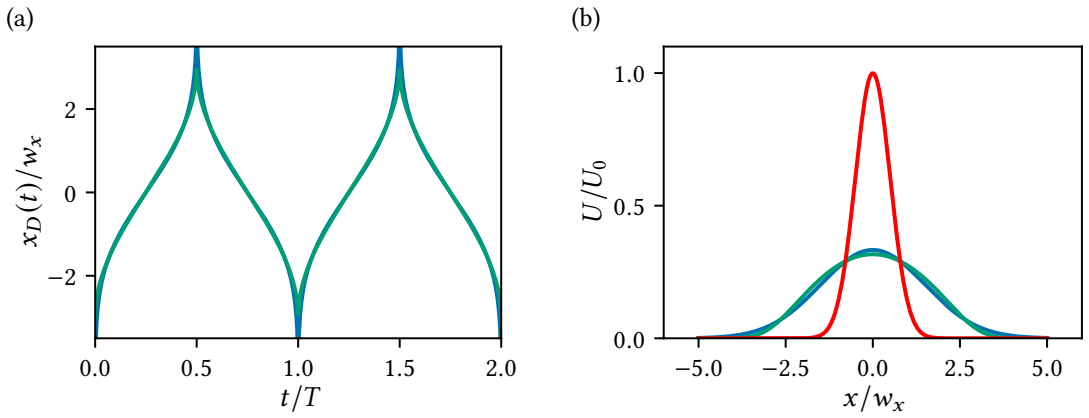

Figure 3.10. Beam dithering. (a) The dithering functions described in Eqs. (3.27) and (3.29) (blue and green, respectively) with $c = 3$. We see the functions are nearly identical and an arcsin function can be used to generate an approximately Gaussian potential with an increased waist. (b) The undithered potential (red) compared to the potential dithered via Eq. (3.29) (green) and via Eq. (3.27) (blue), the latter yielding the target Gaussian potential with waist $3w_x$. The arcsin-dithered potential is also close to Gaussian.

where θ_d is the deflection angle, λ is the wavelength of the light, f_{AOM} is the driving frequency and v_s is the speed of sound. Therefore, changing the driving frequency of the AOM affects the deflection angle and moves the beam slightly, which means that modulating the AOM frequency is an effective way of dithering the beam. It can also be shown that the diffraction efficiency depends on the intensity of the sound, which means that changing the sound intensity changes the diffracted laser intensity. This is useful in our case as changing the laser intensity using the laser itself changes other beam properties (e.g. its waist size) as well. However, the AOM has a finite bandwidth and the power in the diffracted beam depends on the angle of diffraction, which can distort the function we use to dither the beam.

Details of the optical system can be found in Ref. 104, and we are able to load 25% of our atoms from the cMOT into the ODT, where we use the horizontal dithering of the beam to achieve a better spatial overlap. It is worth mentioning that due to the low trapping frequency in the propagation direction of the horizontal ODT beam, we superimpose another (roughly orthogonal) ODT beam for evaporation, creating a 'crossed trap'.

Optical trapping will also be used to realise the box trap in the science cell. For that purpose, the trapping characteristics of a red-detuned, horizontal sheet beam (confining the atoms in the vertical direction) will be combined with a blue-detuned, vertical tube beam (repelling the atoms from the wall, thereby confining them in the intersection of the sheet beam and the tube). The sheet beam will be realised using a 1064 nm fibre

laser,[31] whereas the walls will be realised using a frequency-doubled titanium–sapphire laser[32] with an output of 372 nm, pumped by a diode-pumped solid-state laser[33] operating at 532 nm. As with the ODT, neither of these lasers need locking as we only need their detunings to be large.

3.5 Transport

In our system, a glass cell will be used to conduct experiments due to the better optical access it provides. Therefore, it is necessary to transport the atomic cloud over 40 cm. Sometimes this is done with overlapping magnetic coils [236] or by physically moving a coil pair [237, 238], but the state we use is not magnetically trappable. A different method is all-optical transport which can be performed by displacing the focus of a beam which the atoms are trapped in (an optical dipole trap). This can be achieved with the focusing lens mounted on an air-bearing translation stage [239], but it comes with the drawback of placing a cumbersome system close to the vacuum chamber, which bears the risk of transferring vibrations to the dipole trap or the optical table.

An alternative all-optical approach was demonstrated in Ref. 240, which involves using focus-tunable lenses: lenses that change their focal lengths in response to an applied current. As atoms are trapped in the focal point of the beam (in the ODT), they can be moved from the MOT chamber to the science cell by tuning the focal length of the lenses (an optical tweezer). As this eliminates the risk of transferring vibrations, this solution is being implemented in our system.

To transport the atoms, we need to translate the focal point of the trapping beam along the optical axis. In principle, a focus displacement can be achieved by using a single tunable lens focusing a collimated beam. However, increasing the focal length increases the waist size as well, thereby changing the trap frequencies and depth during transport. The system presented in Fig. 3.11 provides uniform trapping conditions over the full transport distance, which is preferable as only strong confinement and a large trap depth allow for fast transport. Figure 3.11(a) shows two beams which are focused behind the static lens L with focal length f at distances f (position A) and $2f$ (position B). As the waist size of the beams are equal if their divergences are, this requires beam diameters d and $2d$ at L, respectively (resulting in the same divergence $\theta = d/f$ for both). The beam focused at point A must be collimated before the lens, and it can be shown that the beam focused at point B must have the same divergence θ before passing L. Therefore, the two

[31] ALS-IR-1064-10-I-CP-SF by Azurlight Systems.
[32] A SolsTiS system with an ECD-X module (the doubler) by M Squared Lasers.
[33] An 8 W Sprout-G by Lighthouse Photonics.

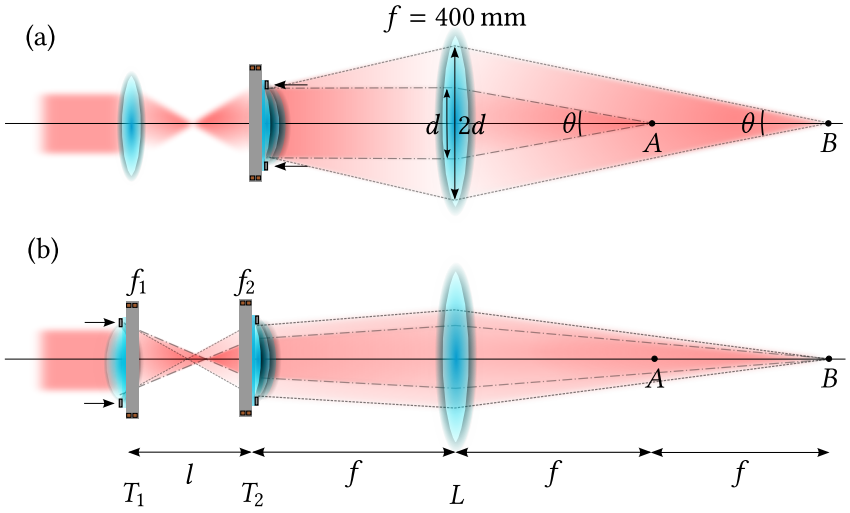

Figure 3.11. All-optical atom transport. (a) Transport at constant beam waist over a distance f. If the separation between the tunable lens T_2 and the static lens L equals the focal length f of the latter, the trap can be translated by tuning f_2, while maintaining the same divergence $\theta = d/f$ and thus the same waist size between points A and B. (b) Independent control over the waist size and position of the focus. Replacing the first lens with the tunable lens T_1 allows changing the beam size at T_2, resulting in a different divergence behind L. Illustration adapted from Ref. 240.

beams have the same size at a distance f before L, and placing a lens (let us call it T_2) with a tunable focus f_2 at this position allows the continuous transformation of one beam into the other, resulting in a moving focus with a constant waist.

This can also be shown using geometrical optics. By using the thin lens equation, $1/f_2 = 1/t_2 + 1/k_2$, and the fact that the incoming beam is collimated (so the object distance is $t_2 = \infty$), we see the distance of the image created by T_2 is $k_2 = f_2$. Therefore, if the incoming beam diameter is d, the beam divergence after T_2 is $\theta_2 = d/k_2 = d/f_2$. For the other lens, the object distance will be $f - f_2$, so using the thin lens equation again, $1/f = 1/(f - f_2) + 1/x_0$, we find the image distance is $x_0 = f(f_2 - f)/f_2$. The diameter at f is $d_f = \theta_2(f - f_2)$, so the divergence after f is $\theta = d_f/x_0 = -d/f$, which is independent of f_2 as expected.

The system as originally proposed [240] used a focus-tunable lens (FTL) which limited $f_2 > 0$. Therefore, the divergence could only be reduced and the beam had to be focused before crossing T_2, as can be seen in Fig. 3.11(a), achieved by a static lens placed in front of T_2, which in turn defines the waist size behind L. In our case, the tunable lenses[34] have a larger range and could also be set to produce $f_2 < 0$. This means that the beam does

[34] EL-16-40-TC-VIS-5D from Optotune.

not have to be focused before T_2 in our case, which is useful to avoid power fluctuations caused by dust particles traversing the focal point. However, the same principle applies as in the original system, and in order to gain independent control over both the position and waist size of the focus, another tunable lens (T_1) is placed in front of T_2, as shown in the extended system in Fig. 3.11(b).

The system was modelled in Mathematica using ray transfer matrix analysis [241] to examine how the laser beam propagates through the system. This enables us to compute the waist size and the focus position for any focal length combination (f_1, f_2). Ray transfer matrix analysis (also known as ABCD matrix analysis) is a type of ray tracing technique used in the design of some optical systems. It involves the construction of a ray transfer matrix which describes the optical system; tracing the light path through the system can then be performed by multiplying this matrix with a vector representing the light ray.

In the thin lens approximation, ray propagation over free space of length l and refraction at a lens with focal length f can be described by the ray transfer matrices

$$\mathsf{P}_l = \begin{pmatrix} 1 & l \\ 0 & 1 \end{pmatrix}, \quad \mathsf{L}_f = \begin{pmatrix} 1 & 0 \\ -1/f & 1 \end{pmatrix}. \tag{3.32}$$

The full system transfer matrix can then be written as

$$\mathsf{S}_x = \mathsf{P}_x \mathsf{L}_f \mathsf{P}_f \mathsf{L}_{f_2} \mathsf{P}_l \mathsf{L}_{f_1}, \tag{3.33}$$

where x is the distance after L and l is the distance between T_1 and T_2. Let us now assume a Gaussian beam with waist size w, waist position z_0 and wavelength λ propagating through the system, along the optical axis (the z-direction). The complex beam parameter $q(z)$ is defined as

$$\frac{1}{q(z)} = \frac{1}{R(z)} - i\frac{\lambda}{\pi w^2(z)}, \tag{3.34}$$

where $R(z)$ is the radius of curvature of the beam and $w(z)$ is the beam radius. From this, we see that the beam radius is given by

$$w(z) = \sqrt{\frac{\lambda}{\pi \operatorname{Im}(-1/q(z))}}, \tag{3.35}$$

and the focus position is given by the condition

$$\operatorname{Re}(q(z)) = 0. \tag{3.36}$$

Now, assuming the incoming beam is collimated at T_1 with waist w, which we choose to

be the origin of our coordinate-system (i.e. $z_0 = 0$ and so $R(0) = \infty$), the incoming beam parameter is

$$q_0 = i\frac{\pi w^2}{\lambda}\,. \tag{3.37}$$

We then use the ray transfer equation

$$q_t = \frac{Aq_0 + B}{Cq_0 + D}\,, \tag{3.38}$$

where the transfer matrix of the system is written as $\mathsf{S}_x = \left(\begin{smallmatrix} A & B \\ C & D \end{smallmatrix}\right)$, which acts on the light vector proportional to $\left(\begin{smallmatrix} q_0 \\ 1 \end{smallmatrix}\right)$. In our case we find

$$q_t(x) = x - f + f^2\left(\frac{1}{f_2} + \frac{\pi w^2 + i f_1 \lambda}{(f_1 - l)\pi w^2 - i f_1 l \lambda}\right). \tag{3.39}$$

We can use this to find the waist size at the trap, w_0, and the focus position x_0, using Eqs. (3.35) and (3.36), respectively, given $w_0 = w(x_0)$. This yields

$$w_0 = \frac{f|f_1|w\lambda}{\sqrt{(f_1 - l)^2(\pi w^2)^2 + (f_1 \lambda l)^2}}\,, \tag{3.40}$$

$$x_0 = f - f^2\left(\frac{1}{f_2} + \frac{(f_1 - l)(\pi w^2)^2 - f_1^2\lambda^2 l}{(f_1 - l)^2(\pi w^2)^2 + (f_1\lambda l)^2}\right). \tag{3.41}$$

We see w_0 is independent of f_2, so it stays constant throughout the transport as expected. Also note that apart from the last term in x_0, which is a small Gaussian-ray correction, it is the same as we found from geometrical optics. While this system has not been implemented yet, some design considerations can be found in Ref. 104.

An interesting question that needs to be investigated is the dynamics of transport, i.e. $x_0(t)$. We want the atoms transported as fast as possible (to avoid atom loss and heating due to background processes), but too fast a transport can also lead to atom loss (atoms not being able to follow the trap) and heating (atoms sloshing in the trap). This question has been addressed in Refs. 242 and 243, where it was found via utilising dynamical invariants (the so-called Lewis–Riesenfeld invariants) that $x_0(t)$ needs to be chosen both classically and quantum mechanically such that

$$x_0(t) = q_c(t) + \frac{1}{\omega_z^2}\ddot{q}_c(t)\,, \tag{3.42}$$

where $q_c(t)$ is an arbitrary function that satisfies $q_c(0) = \dot{q}_c(0) = \ddot{q}_c(0) = \dot{q}_c(T) = \ddot{q}_c(T) = 0$ and $q_c(T) = d$, with T the total time and d the total distance of transport. Given the six

constraints, an obvious candidate for q_c is a 5$^{\text{th}}$-order polynomial,

$$\frac{q_c(t)}{d} = 10\left(\frac{t}{T}\right)^3 - 15\left(\frac{t}{T}\right)^4 + 6\left(\frac{t}{T}\right)^5. \tag{3.43}$$

Another option is using a Fourier series with an added linear term,

$$\frac{q_c(t)}{d} = \frac{t}{T} - \frac{2}{3\pi}\sin\left(2\pi\frac{t}{T}\right) - \frac{1}{12\pi}\sin\left(4\pi\frac{t}{T}\right). \tag{3.44}$$

It is worth noting that while the functional forms of these functions are different, they both define S-shaped $x_0(t)$ curves. To test their behaviour, a simulation is being set up using the AtomECS package [244]; early results indicate that $x_0(t)$ defined via these functions indeed outperform simpler (e.g. 'bang-bang'—a kick at the beginning and the end of transport—or constant acceleration) ramps. Furthermore, it seems that $x_0(t)$ defined via the Fourier series is more robust against uncertainties in ω_z (i.e. ultimately the beam waists) than that defined by the 5$^{\text{th}}$-order polynomial.

3.6 Data acquisition

3.6.1 Control system

As we have seen before, operating the experiment requires precise control of various laser frequencies, magnetic fields, intensities etc. Our magnetic coils are attached to variable power supplies, whereas the lasers are routed through shutters and AOMs to control their frequencies and intensities. The control and time-synchronisation of the experiment are realised by using a controller system with 40 analogue and 32 digital output channels, all independent from each other.[35] These can be controlled from a computer using the Cicero Word Generator [245], a software which generates the experimental sequence and sends it to the hardware. The system also relies on an FPGA chip[36] which provides a variable-frequency clock signal for the output cards to precisely time the sequence.

3.6.2 Imaging

Data in ultracold atom experiments is taken via taking images of the atoms, as the (evolution of the) cloud shape contains important information. Imaging is done via two

[35] Three controller cards are mounted in a PXIe-1082 (8-slot PXI Express) chassis, the PXI-6733 (8-channel, 16-bit analogue output), the PXIe-6738 (32-channel, 16-bit analogue output) and the PXIe-6536 (32-channel digital output), all by National Instruments.

[36] XEM6001 from Opal Kelly. As the Cicero Word Generator was originally designed to work with the XEM3001, a minor code modification was necessary.

techniques, absorption imaging and fluorescence imaging, which look at the light absorbed and scattered by the atoms, respectively. To maximise these effects, the 401 nm (blue) transition is used, due to its large transition rate. To take the images, a CMOS camera[37] is used, triggered by a TTL signal from the digital control card. The images are then saved into a database with the experimental parameters used in the sequence.

Image processing can be automated with the assistance of a neural network which creates region-of-interest boxes (ROIs) centred on our clouds [106]. (One needs to select bounding (ROI) boxes around the cloud to be fitted to speed up the fitting process and to fit only the cloud of interest.) Images are then fitted with a Gaussian function using the ROI as seed via the JAX library [246]. Furthermore, the two steps can also be unified, training a neural network to directly extract cloud properties [247].

To achieve automatic ROI placement, we trained a neural network using data collected in the lab (including by me, among others) to produce the ROI boxes. Besides automating the fitting process, we found that this method performed better than conventional image fitting methods and yielded significant speed gains. Details of the training of the neural network is presented in §3.7.

Besides the atoms, sometimes the beam itself needs to be imaged for characterisation purposes. This was done by directing beams on the camera directly, or via neutral-density (ND) filters if the intensity was too high. However, imaging beams with very high intensities (which not even ND filters can withstand) is not trivial, and a technique to avoid aberrations is described in Appendix A.

Absorption imaging

Absorption imaging measures the optical density of the atomic cloud along the imaging direction, from which the areal density of the cloud can be extracted. Each measurement requires three images to be taken, one with the atoms in place casting a shadow (resulting in $I_a(x, y)$), one of the probe beam without any atoms ($I_p(x, y)$) and one without any light to take the background ($I_b(x, y)$). The optical density (OD) of the atomic cloud is given by

$$\mathrm{OD}(x, y) = \ln\left(\frac{I_p - I_b}{I_a - I_b}\right), \tag{3.45}$$

[37] DCC3260M by Thorlabs, to be changed to ORCA-Fusion by Hamamatsu for better performance. It is worth noting that the ORCA-Fusion is a CMOS camera, not an EMCCD which are traditionally used for low-light applications. We chose this camera after careful (numerical and experimental) evaluation against an Andor EMCCD camera, and we found the EMCCD only performed better in the single-photon (per pixel) regime. This can be explained by the extremely low read noise of the ORCA-Fusion; while the EMCCD has essentially no read noise, it does have a higher shot noise which cancels its advantage at higher photon numbers.

which is related to the areal (number) density of the atoms through the optical cross section σ via

$$n_{2D}(x, y) = \frac{1}{\sigma} \mathrm{OD}(x, y) \,. \tag{3.46}$$

The optical scattering cross section of a two-level system can be written as

$$\sigma = \frac{3\lambda^2}{2\pi} \frac{1}{1 + s + (2\delta/\Gamma)^2} \,. \tag{3.47}$$

As we use resonant light with $s \approx 0.1$, we can use $\sigma \approx 3\lambda^2/2\pi$. Furthermore, to ensure we indeed have a two-level system, we excite the closed transition $|6, -6\rangle \rightarrow |7, -7\rangle$ with a σ^--polarised beam. A magnetic field is therefore kept on during imaging to polarise the atoms.

Fluorescence imaging

Fluorescence imaging is performed by collecting the photons scattered by the atomic cloud and directing them on the camera, which can be used to determine the number of atoms present. This can be done the following way. Given the parameters of the imaging laser, one can calculate the photon scattering rate R using Eq. (3.1). N atoms scatter at a rate of NR, and given a camera exposure time t the number of scattered photons is given by NRt. The number of photons reaching the camera is proportional to the solid angle under which the camera is seen from the atomic cloud. With the knowledge of camera parameters (e.g. its quantum efficiency) the camera sensor counts can be converted to a photon number count, which can then be used to determine the number of atoms responsible for scattering that many photons in the given solid angle.

3.7 Publication: Atom Cloud Detection and Segmentation Using a Deep Neural Network

Machine Learning: Science and Technology **2**, 045008 (2021)
Submitted: 25 November 2020; Published: 15 July 2021
DOI: 10.1088/2632-2153/abf5ee

Lucas R. Hofer,[1] Milan Krstajić,[1,2] Péter Juhász,[1] Anna L. Marchant[1,3] and Robert P. Smith[1]

[1] *Clarendon Laboratory, University of Oxford, Parks Road, Oxford OX1 3PU, United Kingdom*
[2] *Cavendish Laboratory, University of Cambridge, J. J. Thomson Avenue, Cambridge CB3 0HE, United Kingdom*
[3] *Current address: RAL Space, Science and Technology Facilities Council, Rutherford Appleton Laboratory, Harwell, Didcot OX11 0QX, United Kingdom*

IOP Publishing Mach. Learn.: Sci. Technol. 2 (2021) 045008 https://doi.org/10.1088/2632-2153/abf5ee

MACHINE LEARNING
Science and Technology

OPEN ACCESS

RECEIVED
25 November 2020

REVISED
25 March 2021

ACCEPTED FOR PUBLICATION
8 April 2021

PUBLISHED
15 July 2021

PAPER

Atom cloud detection and segmentation using a deep neural network

Lucas R Hofer[1], Milan Krstajić[1,2], Péter Juhász[1], Anna L Marchant[1,3] and Robert P Smith[1,*]

[1] Clarendon Laboratory, University of Oxford, Parks Road, Oxford OX1 3PU, United Kingdom
[2] Cavendish Laboratory, University of Cambridge, J. J. Thomson Avenue, Cambridge CB3 0HE, United Kingdom
[3] Current address: RAL Space, Science and Technology Facilities Council, Rutherford Appleton Laboratory, Harwell, Didcot, Oxfordshire OX11 0QX, United Kingdom
[*] Author to whom any correspondence should be addressed.

E-mail: robert.smith@physics.ox.ac.uk

Keywords: ultracold quantum matter, machine learning, deep neural networks, Bayesian optimization, object detection, instance segmentation, image processing

Abstract

We use a deep neural network (NN) to detect and place region-of-interest (ROI) boxes around ultracold atom clouds in absorption and fluorescence images—with the ability to identify and bound multiple clouds within a single image. The NN also outputs segmentation masks that identify the size, shape and orientation of each cloud from which we extract the clouds' Gaussian parameters. This allows 2D Gaussian fits to be reliably seeded thereby enabling fully automatic image processing. The method developed performs significantly better than a more conventional method based on a standardized image analysis library (Scikit-image) both for identifying ROI and extracting Gaussian parameters.

1. Introduction

Deep neural networks (NNs) have revolutionized data analysis and led to automation of tasks that previously required human supervision. Image analysis has particularly benefited through the use of convolutional neural networks (CNNs) [1] and their derivatives which have allowed for image classification [2, 3], object detection [4, 5] and instance segmentation [6]. Although many of these NNs were developed for tasks such as facial recognition by social media networks [7, 8], they have also been used to identify laser modes [9], classify phases in condensed-matter systems [10, 11], reduce measurement errors for trapped-ion qubits [12] and process images from cold atom experiments [13–15]. In this work, we use an instance segmentation NN (see figure 1) to analyze experimental images containing atom clouds in magneto-optical traps (MOTs) and optical dipole traps (ODTs).

Neural networks consist of an input layer and an output layer with a number of intermediate hidden layers which are connected to one another via 'weights'. Rather than employing hard-coded algorithms, NNs learn to emulate data they encounter through training cycles, in which data is iteratively passed through the NN and the output compared to the 'ground truth'. The difference between the two is then used to update the weights between the NNs layers, thereby improving its accuracy. When employing supervised training, this requires a dataset which includes both input data and their associated ground truth values. For object detection NNs, these ground truth values are rectangular bounding boxes for each object, as well as labels classifying the object types in the bounding boxes. Instance segmentation NNs build upon object detection NNs by also requiring pixel-to-pixel segmentation masks in which image pixels comprising the object have mask values of one, whereas all other pixels have mask values of zero.

Our dataset consists of images of cold atom clouds in a MOT [16] and an ODT [17] (see figures 2(a)–(c)). Atom clouds in these traps form approximately Gaussian density distributions [18]. Fitting a cloud allows the parameters describing the distribution (Gaussian parameters) to be extracted and used to ascertain information such as the cloud's size and density. Furthermore, by using time-of-flight measurements [19] the temperature of the atoms can be determined.

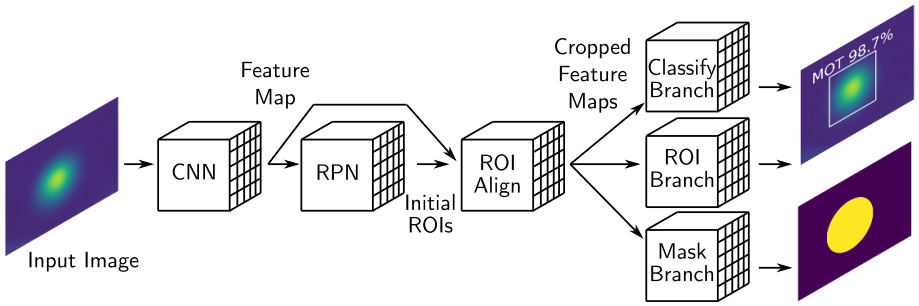

Figure 1. Overview of the neural network (NN). An experimental image is initially fed into a convolutional neural network (CNN) which produces a 'feature map' of the image. The region-proposal network (RPN) uses this feature map to generate regions-of-interest (ROI) where atom clouds are likely contained. In the ROI alignment stage, the CNNs feature map is cropped and resized for each cloud's ROI before being fed into three parallel branches. The first branch (ROI) generates refined ROIs, whereas the second branch (classify) classifies the type of atom cloud in the ROI. Finally, the third branch (mask) generates a segmentation mask corresponding to the $1/e^2$ contour of the atom cloud.

A region-of-interest (ROI) [20] centered on the atom cloud is used during fitting as objects in the image other than the atom cloud can cause an inaccurate fit (e.g. an atom cloud in another trap or extraneous noise). Additionally, decreasing the fit area can significantly decrease the fit time when using two-dimensional fitting. Manually determining the ROI is time-consuming when analyzing a large number of images and an algorithmic procedure is often employed, such as taking the 'center of mass' and then iteratively expanding the ROI around this point until the fractional enclosed 'power' exceeds some threshold. Another method involves performing connected component analysis on a binarized version of the image and then measuring the resulting regions using common image processing libraries [21]. However, if the image is noisy (e.g. contains fringing), the proposed ROIs will be inaccurate for these types of methods (section 6).

Here we propose a deep NN based approach to ROI determination in which a NN finds the ROI for each atom cloud in an image (see figure 1). Furthermore, the NN differentiates between clouds (classification) in a MOT and those in an ODT and also outputs a segmentation mask for each cloud from which Gaussian parameters are directly extracted. The classification feature is particularly useful when the cloud type is not *a priori* known from the experimental sequence (e.g. images containing both MOT and ODT clouds during the ODT loading). Although features such as the position or aspect ratio can be used for classification with additional manual or experimental input, the NN needs the image alone (beyond training on an appropriate dataset) to determine cloud types and so is easily adaptable to other cold atom experiments with different numbers of clouds or cloud types.

The rest of the article is arranged as follows: section 2 describes the experimental dataset used for NN training and validation, section 3 describes the training process, section 4 discusses Bayesian optimization (BO) [22, 23] of the NNs hyperparameters [23] and section 5 examines how the Gaussian parameters are calculated from the segmentation mask. In section 6 we compare our proposed NN method to a more conventional method of determining both ROIs and Gaussian parameters.

2. Experiment and dataset

To produce the ultracold atom clouds, erbium atoms are initially trapped and cooled to $\sim$20 μK in a narrow-line MOT [24] before being loaded into an ODT formed from a 30 W, 1030 nm laser beam focused down to a $\sim$40 μm waist (see figure 2). Optimization of the trap loading involves maximizing the atom number while minimizing the cloud temperature. The atom number is found by fitting the atom cloud in either a fluorescence or absorption image[4] with a two-dimensional (2D) Gaussian (see equation (1)) and then integrating under the curve. The cloud temperature can be determined from how the cloud width evolves during time-of-flight expansion.

The experimental dataset consists of 260 fluorescence and absorption images (1936 $\times$ 1216 pixels) along with their ROIs, labels and segmentation masks [25]. Of these, 130 images contain clouds released from the MOT with no ODT present (see figure 2(b)) and 130 images contain either just atoms released from the ODT (see figure 2(c)) or images where atoms released from the ODT and the MOT are both present (see

[4] Experimental images shown in the paper are processed from two to three raw images. Fluorescence images require both an image with atoms and a background image without atoms. Absorption images additionally require a probe image in which the probe beam is turned on, but no atoms are present.

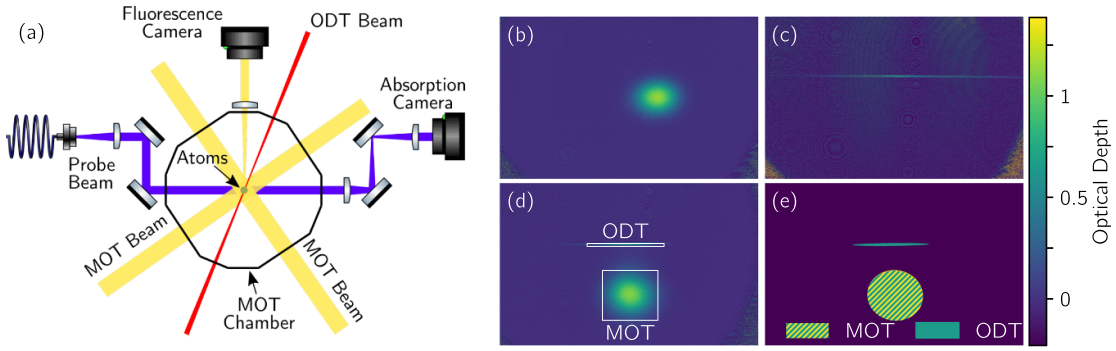

Figure 2. Construction of labelled experimental dataset. (a) Experimental setup. The MOT is formed from six 583 nm laser beams in conjunction with a quadrupole magnetic field. Atoms trapped in the MOT can be loaded into an ODT formed from a single 1030 nm laser beam. The fluorescence camera captures the light scattered by the atoms from the MOT beams. During absorption imaging, a 401 nm probe beam passes through the atom cloud which absorbs a portion of the light creating a shadow imaged by the camera. (b) Image of an atom cloud released from the MOT. (c) Image of an atom cloud released from the ODT. The optical depth is multiplied by six to increase ODT visibility. (d) Image with both MOT and ODT atom clouds. Each cloud is labelled and a region-of-interest (ROI) box is drawn around it (white lines). (e) Segmentation masks for the MOT (hatched yellow) and ODT (solid green) clouds in (d).

figure 2(d)). We manually label the atom clouds in the images as 'MOT' or 'ODT' and draw a ROI box at the clouds' edges which we define as the $1/e^2$ radii along the x and y axes (see figure 2(d)). This definition prevents the ROI boxes from overlapping when the MOT and ODT are both present; however, the ROI boxes are also easily expandable when analysis requires the wings of the distribution.

The manually drawn ROIs were expanded by a factor of two—excepting where the expanded ROIs would overlap—and the atom clouds inside fit with a 2D Gaussian:

$$I(x,y) = I_\mathrm{b} + I_0 e^{-2\left(\frac{[(x-x_0)\cos\theta+(y-y_0)\sin\theta]^2}{w_x^2}+\frac{[(y-y_0)\cos\theta-(x-x_0)\sin\theta]^2}{w_y^2}\right)}, \quad (1)$$

where $I(x,y)$ is the image intensity, I_b is the background intensity, I_0 is the peak intensity, x_0 and y_0 are the center coordinates, w_x and w_y are the $1/e^2$ radii along the major and minor axes and θ is the angular orientation of the distribution. To increase the accuracy of the ROIs used for training, the fit parameters were used to calculate the $1/e^2$ radii along the image axes [9] (previously estimated by eye) and the ROI boxes redrawn using these values. The process of fitting and redrawing the ROI boxes from the fit parameters was then completed a second time with subsequent iterations neglected due to an insignificant increase in accuracy.

A segmentation mask was generated for each atom cloud (see figure 2(e)) with the mask borders placed at the $1/e^2$ contour of the cloud—calculated from the final fit parameters; pixels within the $1/e^2$ contour were set to one, whereas pixels outside were set to zero. Finally, the dataset was randomly split into a training dataset with 200 images and a validation dataset with 60 images.

3. Neural network and training

We use the NN Mask R-CNN [26] to detect and bound the atom clouds, as well as to provide segmentation masks for each cloud (see figures 3(a, b)). The NN (see figure 1) begins with the experimental image being fed into a CNN base (ResNet-50 [27]). The CNNs outputted feature map [28] is then passed into a region-proposal network (RPN) which generates ROIs where objects are likely located. Next, these ROIs are cropped from the CNNs feature map in a ROI alignment stage and resized to uniform dimensions. The cropped feature maps are then fed into three parallel branches. The first applies a classifier to determine the object type and give the confidence of its prediction—which is helpful in determining whether to use the ROI in post-NN analysis. The second branch gives a more accurate ROI box prediction (see figure 3(a)) and finally the third outputs a segmentation mask for the object inside the ROI (see figure 3(b)). Since all three branches share the same base, computation speed is significantly increased [29] for both training and evaluation.

During training, the NN output is compared to the ground truth (i.e. the expected output from the training dataset) via a loss function; the loss is then back-propagated [30] through the NN to adjust the weights between layers and refine the NN model (see figure 4(a)). The loss function for the RPN stage is L1 loss [31], for the ROI branch it is Smooth L1 loss [32], the classifier uses categorical cross entropy loss [33] and lastly the mask branch utilizes binary cross entropy loss [34]. Although the loss can be separately

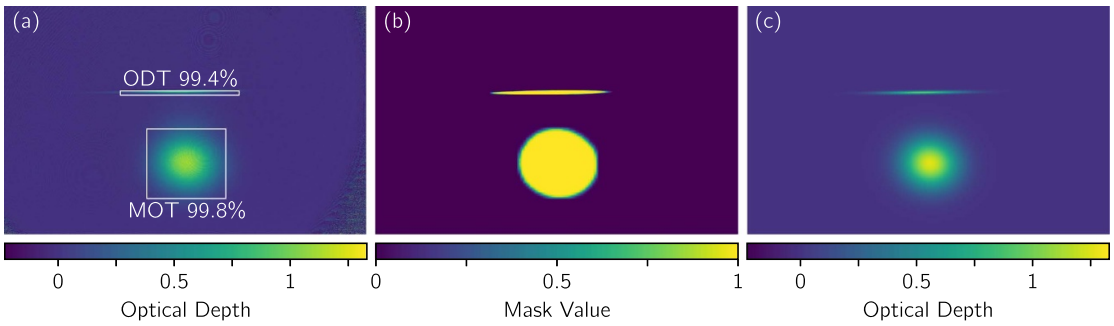

Figure 3. Neural network output. (a) Sample image from the validation dataset with the regions-of-interest (solid white lines), labels and confidence scores predicted by the NN. (b) The segmentation masks returned by the NN for (a) which, due to the internal structure of the NN, have values between zero and unity. (c) Reconstruction of (a) using equation (1) and the extracted 2D Gaussian parameters from (b).

back-propagated for each branch, a simpler approach is taken here in which the losses are summed together and then back-propagated to update the weights of the NN [26].

An epoch denotes a single cycle of training in which every image in the training dataset is passed through the NN, the loss calculated and the model weights updated. Increasing the number of training epochs can increase the NN's final accuracy so long as the NN does not overfit on the training data. However, due to finite computational resources, we restrict the training of an individual NN to 15 epochs—determined to be sufficient, see below—and instead use BO to increase the accuracy of the final trained model (see section 4).

Five hyperparameters, which tune the learning process, are set before the NN's training phase. The first is the *learning rate* which determines the size of the step the NN takes during stochastic gradient descent [35]. If the learning rate is too low, the NN will take too long to converge to the minimum of the loss function, whereas if the learning rate is too high, the NN will not be able to descend to the minimum, but will oscillate around it or diverge. Since larger learning rates are useful at the beginning of training and smaller learning rates are useful towards the end, a learning rate scheduler [36] is employed which decreases the learning rate by some scalar (*decay*, the second hyperparameter) after a fixed number of epochs (*step size*, the third hyperparameter). The fourth hyperparameter is the *momentum* which prevents the NN from getting stuck in a local minimum during training [37]. The last hyperparameter is the *batch size*—the number of images simultaneously passed through the NN—and, unlike the other hyperparameters which are tuned with BO, is fixed to four for all NNs.

The accuracy of the NN is evaluated (see figure 4(b)) using the Common Objects in Context (COCO) evaluation metric which entails calculating the intersection-over-union [38]:

$$\mathrm{IoU} = \frac{\mathrm{area}\left(\mathrm{ROI_p} \cap \mathrm{ROI_{gt}}\right)}{\mathrm{area}\left(\mathrm{ROI_p} \cup \mathrm{ROI_{gt}}\right)}, \tag{2}$$

where $\mathrm{ROI_p}$ is the ROI prediction from the NN and $\mathrm{ROI_{gt}}$ is the ground truth ROI. Values below a predetermined IoU threshold (e.g. $\mathrm{IoU} = 0.5$) are considered a false prediction—mislabelling the object class within the ROI is also a false prediction—whereas values above are considered a true prediction. A precision recall-curve [39] is then constructed and integrated to give the average precision (AP) for the given threshold (e.g. $\mathrm{AP_{50}}$ for the $\mathrm{IoU} = 0.5$ threshold). Finally, the AP is calculated for ten IoU thresholds (0.50–0.95 with a step size of 0.05) and averaged to give the mean average precision (mAP)—which is the metric generally reported for object detection. The mask AP values and mAP can be similarly calculated [40].

The NNs are trained in a Google Colab notebook [41] utilizing a GPU backend and implemented in PyTorch [42] using the pre-built Mask-RCNN model in the Torchvision package. Rather than training the NN from scratch, the model weights are loaded from a network pre-trained on the COCO train2017 dataset [43]—which significantly reduces the time required to train the network and increases the model's final accuracy [44]. Additionally, this transfer learning, complemented by data augmentation (in which images are randomly flipped horizontally during training), allows us to use a relatively small training dataset. After each training epoch, the updated NN is evaluated on the validation dataset which yields the mAP for both the object detection and the segmentation mask branches. These two values are averaged together to give the accuracy of the training epoch (average mAP) and the epoch with the highest average mAP determines the overall accuracy of the NN model. After using BO to determine the hyperparameters which give the highest accuracy, we retrain the NN with these parameters for 30 epochs (see figure 4(c)) and verify that the NNs accuracy becomes asymptotic at the 15 epoch mark—with no overfitting seen thereafter.

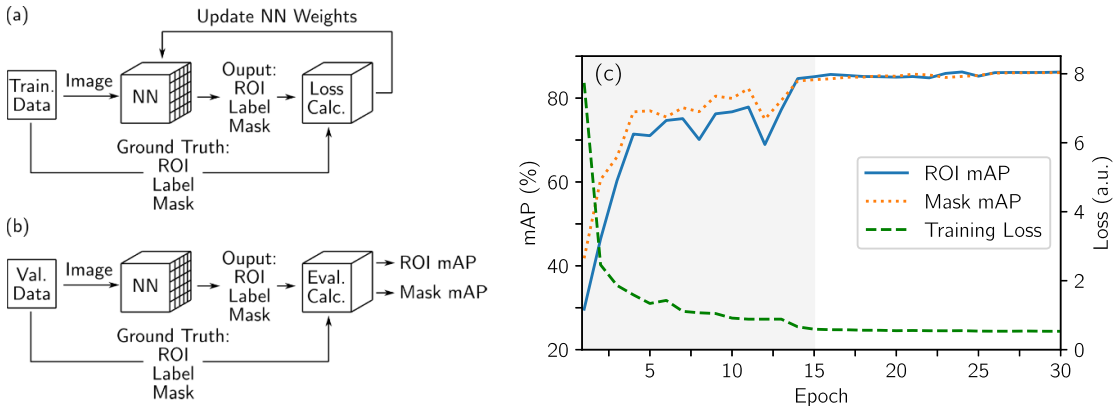

Figure 4. Training and validating the neural network. (a) During training of the NN, images from the training dataset are fed into the NN which outputs the regions-of-interest, labels and masks. These outputs are compared to the ground truth values and the loss calculated. The loss is then back-propagated through the NN and used to update the weights between the NNs layers. (b) When evaluating the NN, the validation dataset is used in a process similar to (a); however, the mean average precision (mAP) for both the object detection branch (ROI) and segmentation mask branch are calculated rather than the loss. (c) The mAP of the ROI branch (solid line), the mAP of the segmentation mask branch (dotted line) and the training loss (dashed line) as a function of the training epoch for the NN trained with the best set of hyperparameters (see table 1). The shaded region represents the number of epochs used during Bayesian optimization.

Table 1. The hyperparameter search space used during Bayesian optimization along with the hyperparameter values used to train the NN to the highest average mAP (see figures 4(a) and 5).

Parameters	Lower bound	Upper bound	Log scale	Best value
Learning rate	0.0001	0.009	Yes	0.0033
Momentum	0.7	0.92	No	0.86
LRS step size	3	15	No	13
LRS decay	0.001	1	Yes	0.049

4. Bayesian optimization of hyperparameters

The accuracy of the trained NN is sensitive to the hyperparameters used during training. In the past, the hyperparameters were tuned through grid search, random search [45] or by hand; however, in recent years, BO has been successfully employed to find the best set of hyperparameters [23, 46]. BO is particularly useful when trying to find the minimum (or maximum) of a function which is noisy and expensive to evaluate—such as NN training—thereby making a grid search of the parameter space impractical [22].

Bayesian optimization takes the function value (cost) at previously evaluated points and uses a Gaussian process to model the cost as a function of the parameter space [47]. The model also determines the confidence of its predictions in a given region of the parameter space: perfect certainty at evaluated points, low uncertainty near evaluated points and high uncertainty far from evaluated points. The BO loop then determines where to evaluate the function next by weighing the benefits of evaluating the function near the model's predicted minimum (or maximum) or evaluating the function in an unexplored region of the parameter space [48].

When optimizing our NN training with BO, the average mAP of the NN is evaluated as a function of the hyperparameter space which consists of the learning rate, momentum, and the learning rate scheduler (LRS) step size and decay (see table 1). As a warm start, the NN is initially trained and evaluated at five quasi-randomly (Sobol generated [49]) distributed points (see black squares in figures 5(a)–(f)). Further evaluation points (red circles in figures 5(a)–(f)) are iteratively determined by the BO loop. The Ax Python package is used for the BO loop as it provides a high level interface to the BoTorch [50] BO package.

With an increasing number of BO evaluation trials the best achieved average mAP rises and converges (see figure 5(g)). The best set of hyperparameters (see table 1) gives a mAP of 86.3% for the object detection branch—locating 88 of the 89 clouds—and a mAP of 85.8% for the mask branch. These mAP values are higher than those for similar NNs trained on the COCO validation dataset [26, 51] which is likely due to our NN only needing to classify two object types rather than the eighty in the COCO validation dataset [38], as well as the relatively simple features of the MOT and ODT clouds. The NN also performs well for clouds with a low signal-to-noise ratio (SNR) which we define as the cloud's peak amplitude divided by the standard

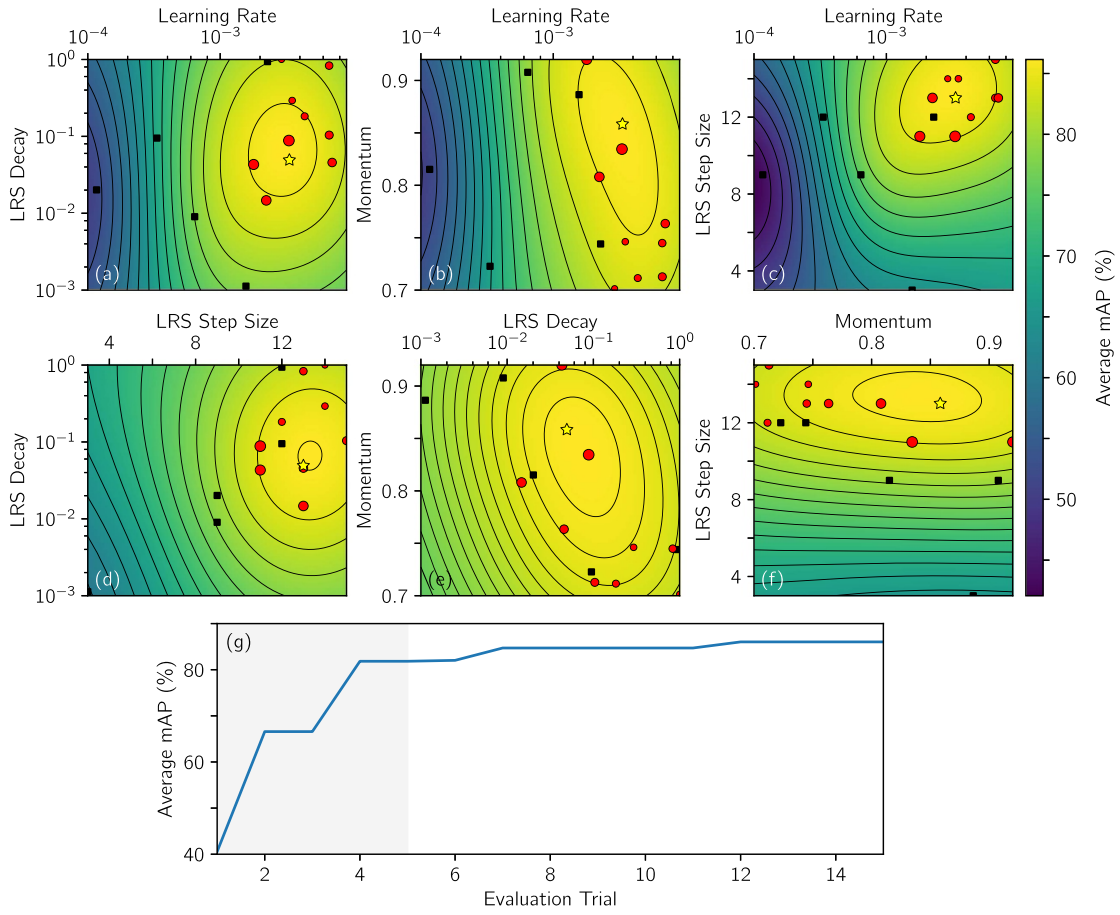

Figure 5. Bayesian optimization of hyperparameters. (a)–(f) The average mAP of the NN predicted by the Bayesian optimization (BO) model as a function of the hyperparameters. The plots are 2D slices through the 4D parameter space taken at the best position (yellow star) for those parameters not displayed. The NN is initially trained and evaluated at five quasi-random hyperparameter positions (black squares) after which the BO loop iteratively determines the evaluation points (red dots, size increases with trial). (g) The best achieved average mAP as a function of the evaluation trial; the quasi-random trial region is shaded in grey.

deviation of the image's background intensity; the validation dataset's nosiest cloud has an SNR of only 2.1 but returns an ROI IoU score of 0.86.

5. Gaussian parameter analysis

We can extract the parameters characterizing the cloud's 2D Gaussian distribution (see equation (1)) directly from the NNs segmentation mask output (see figure 6(a)). Taking the 1st moments of the mask yields the center coordinates $\{x_0, y_0\}$ of the atom cloud, whereas the $1/e^2$ radii $\{w_x, w_y\}$ and the angular orientation θ can be determined by calculating the 2nd moments [52] of the mask. The background intensity I_b is calculated by generating a histogram of the experimental image's pixel values (where there is no cloud) and then taking the peak position (see figure 6(c)). The amplitude of the 2D Gaussian I_0 is calculated by summing the image's intensity inside the $1/e^2$ contour to find the power (P) and then applying:

$$I_0 = \frac{2}{(1-e^{-2})}\left(\frac{P}{\pi w_x w_y} - I_b\right). \tag{3}$$

This method of extracting clouds' Gaussian parameters directly from the segmentation mask was applied to the validation dataset (see figure 3(c)) and the results compared to the Gaussian parameters extracted via a 2D fit of the experimental image (see figure 7, blue bars). Apart from θ, we normalized the differences to obtain a relative error. The cloud center $\{x_0, y_0\}$ and the cloud radii $\{w_x, w_y\}$ were divided by the fitted radii, whereas the peak amplitude I_0 and the background intensity I_b were both divided by the fitted peak amplitude.

The Gaussian parameters extracted from the NNs segmentation mask can either be used directly or as seed parameters for a conventional 2D fit—increasing the likelihood of fit convergence and reducing the

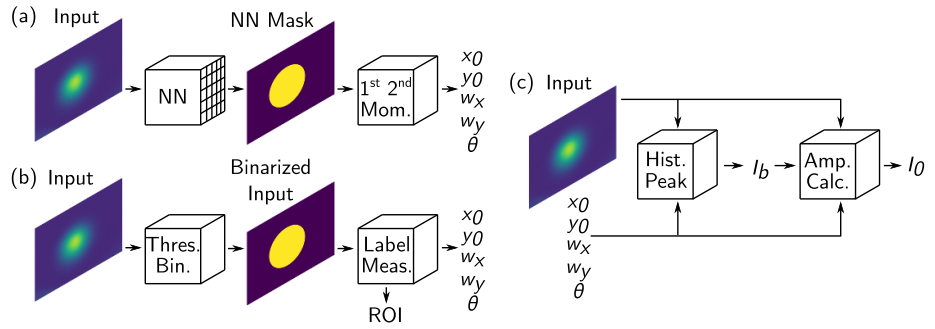

Figure 6. Extraction of Gaussian parameters. (a) NN extraction of the parameters is performed by first inputting the experimental image into the NN which returns a segmentation mask for each atom cloud. Applying the first and second moments (Mom.) directly to the segmentation mask yields the center coordinates of the atom cloud $\{x_0, y_0\}$, its $1/e^2$ radii $\{w_x, w_y\}$ and its angular orientation θ. (b) For conventional extraction the experimental image is binarized using a thresholding method (Thres. Bin.) and the binarized image regions labelled and measured (Label Meas.) using Scikit-image's regionprops function to yield the ROI as well as $\{x_0, y_0, w_x, w_y, \theta\}$. (c) Extracting the intensity parameters. The intensity offset I_b is determined by taking the intensity level corresponding to the peak of a histogram of the experimental image pixel values (hist. peak). The peak intensity I_0 is then calculated (see equation (3)) using the power inside the $1/e^2$ contour and the previously extracted parameters $\{x_0, y_0, w_x, w_y, \theta, I_b\}$.

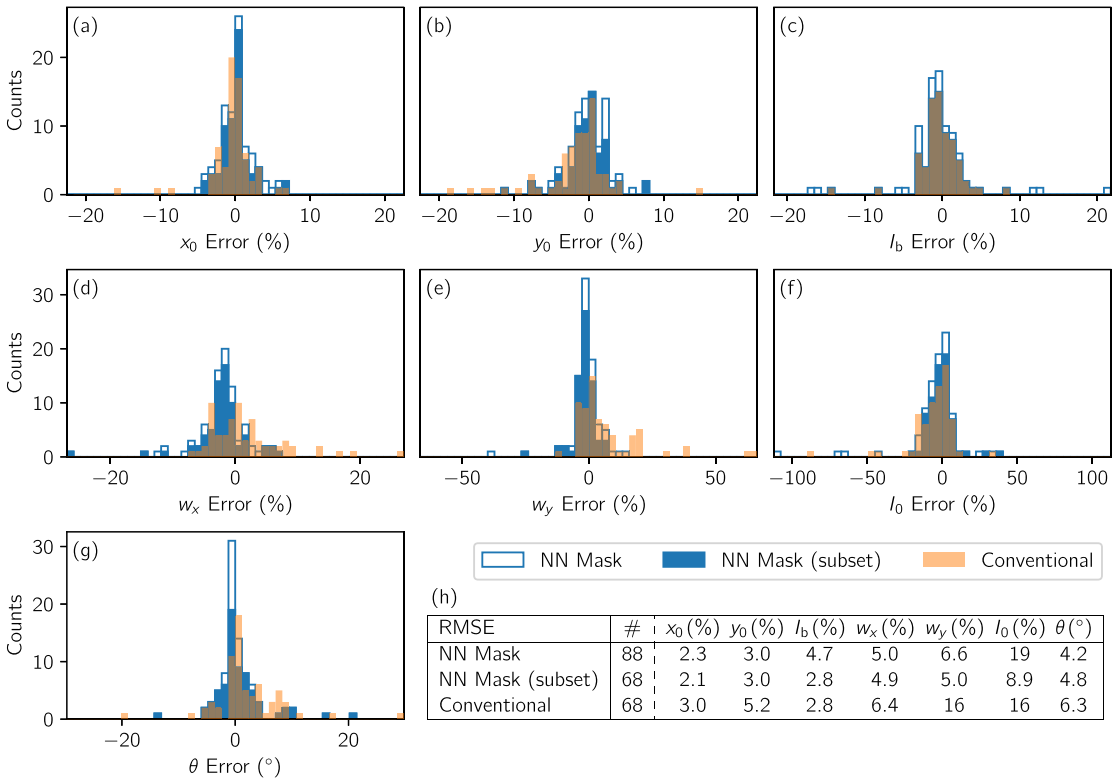

(h)

RMSE	#	x_0 (%)	y_0 (%)	I_b (%)	w_x (%)	w_y (%)	I_0 (%)	θ (°)
NN Mask	88	2.3	3.0	4.7	5.0	6.6	19	4.2
NN Mask (subset)	68	2.1	3.0	2.8	4.9	5.0	8.9	4.8
Conventional	68	3.0	5.2	2.8	6.4	16	16	6.3

Figure 7. Gaussian parameter errors. (a)–(g) Histograms of the extracted Gaussian parameter errors—calculated by comparison to the 2D fit values—for the atom clouds in the validation dataset. Both the relative errors for the neural network segmentation mask method (subset, filled blue) and the conventional method (orange) are shown for the 68 clouds which the conventional method locates. The relative errors for all 88 clouds the NN method locates are also shown (outlined blue). Cloud center coordinates $\{x_0, y_0\}$ and cloud radii $\{w_x, w_y\}$ are normalized by the fitted cloud radii. The peak amplitude and the background intensity $\{I_0, I_b\}$ are both normalized by the fitted peak amplitude. (h) The root-mean-square error (RMSE) of the Gaussian parameters and the number of clouds analyzed (#).

fitting time. Without a seed, the fit time for a full image (1936 × 1216 pixels) is approximately 15 s on an Intel Xeon 2.2 GHz processor core. For the validation dataset, this is reduced to an average of 3 s when using the ROIs proposed by the NN—expanded by a factor of two except where the expanded ROIs would overlap—which substantially increases the fit speed despite the NNs processing time of 0.17 s per image using a Nvidia Tesla V100 GPU. Extracting the Gaussian parameters takes an average of 0.18 s, but results in an average fit speed up of 2 s per region when used as a seed. Thus, our method of determining ROIs and seed parameters using a NN offers a significant speedup in conjunction with fitting.

6. Comparison with conventional analysis

For comparison with the NN, a conventional method is used to find the ROIs and Gaussian parameters of the clouds (see figure 6(b)). We utilize standard methods based on Python's Scikit-image (Skimage) package [21] to pre-process the images, determine ROIs for the clouds, and extract the Gaussian parameters.

During pre-processing, a histogram is taken of the image's pixel intensities. As the cloud areas are much smaller than that of the overall image an approximately Gaussian peak corresponding to the image's background level and noise is seen. The mean μ and standard deviation σ of a Gaussian fit to this peak are used to define a threshold value $I_{\text{thresh}} = \mu + 3\sigma$ [53] and the image is binarized by setting pixels with $I > I_{\text{thresh}}$ to one and the remainder to zero.

After pre-processing, connected component analysis [54] is applied to the binarized image which groups pixels of the same value together into 'regions' (e.g. an atomic cloud). These regions are fed into Skimage's regionprops method which returns both the ROI coordinates and the geometric parameters $\{x_0, y_0, w_x, w_y, \theta\}$ for each region. ROIs with an area $> \frac{1}{2}$ or $< \frac{1}{800}$ the image size (i.e. much larger or smaller than the size of the clouds under measurement) are removed and the two largest remaining ROIs—if there are more than one—are taken as the MOT and ODT ROIs since smaller regions generally coincide with remaining high noise regions.

Applying this method to the validation set locates 68 of the 89 clouds with an IoU $> 50\%$; this is significantly worse than the NN which locates 88 of the 89 clouds. The mAP is also calculated, but since the conventional method does not differentiate between MOT and ODT clouds, all the validation dataset regions are relabelled as 'cloud'. Even with this simplification, an mAP of only 11.1% is achieved—much lower than the NN's object detection mAP of 86.3%. On closer inspection, the conventional method fails on images with low SNR such as images with significant fringing.

To determine the Gaussian parameters of the 68 successfully bounded clouds we apply a numerical scaling factor (to account for our thresholding method [52]) to the $\{w_x, w_y\}$ returned from regionprops and then generate a binary mask from the resulting $\{x_0, y_0, w_x, w_y, \theta\}$ which is used to find the cloud's background intensity I_b and amplitude I_0 (see figure 6(c)) as in section 5. The parameters are normalized and compared to those from the 2D fit (see figure 7, orange bars). Since the conventional method only recognizes a subset of the clouds, we quantitatively compare against the NN segmentation method for the same cloud subset (shown as filled blue bars in figure 7); a direct comparison of the root mean squared errors is shown in figure 7(h).

When considering both object detection and Gaussian parameter extraction, the NN method significantly outperforms the conventional method. Additionally, the conventional method's efficacy depends on the SNR of the data, whereas the NN is more robust against high noise levels and fringing since it learns higher level features of the atom clouds. Furthermore, the conventional method requires manual determination of the best pre-processing procedures and therefore potentially needs further tuning for new data; however, when faced with new data the NN simply needs to be retrained with more labelled data, thus making it effective in a laboratory setting.

7. Conclusion

An instance segmentation NN (Mask R-CNN) was trained to identify ultracold atom clouds in MOTs and ODTs. The NN generates both a ROI and a segmentation mask for each cloud—corresponding to the cloud's $1/e^2$ radii—with a mAP of 86.3% and 85.8% on the ROI and the segmentation mask branches, respectively. We show that the Gaussian parameters describing the atom clouds' distributions can also be extracted directly from the segmentation masks. Both ROI determination and Gaussian parameter extraction via the NN are significantly more accurate than a conventional method based on Python's Scikit-image library.

With an appropriate training dataset these techniques could be applied to ultracold atom clouds in other traps such as optical lattices [55, 56] and box potentials [57]; they are also directly applicable to laser beam profiling and other machine vision applications which require analysis of one or more 2D Gaussian distributions.

In the future a custom NN could be created by adding a branch after the ROI alignment stage which would output the cloud parameters directly. This would enable the characterization of non-Gaussian density profiles, useful, for example, in the detection, identification and parameterization of the bimodal clouds seen when a Bose–Einstein condensate [58, 59] is present.

Data availability statement

The data that support the findings of this study are openly available at the following URL/DOI: https://doi.org/10.5287/bodleian:Y56kydgmj. Data will be available from 7 January 2021.

Acknowledgments

We thank Elliot Bentine, Shu Ishida and Jirka Kučera for helpful discussions and comments on the manuscript. This work was supported by EPSRC Grant No. EP/P009565/1, the John Fell Oxford University Press (OUP) Research Fund and the Royal Society. M K acknowledges funding from Trinity College, University of Cambridge.

Conflict of interest

The authors declare no conflicts of interest.

ORCID iDs

Lucas R Hofer https://orcid.org/0000-0002-5526-587X
Péter Juhász https://orcid.org/0000-0002-5187-730X
Anna L Marchant https://orcid.org/0000-0002-6350-4842

References

[1] Lecun Y, Bottou L, Bengio Y and Haffner P 1998 *Proc. IEEE* **86** 2278–324
[2] Krizhevsky A, Sutskever I and Hinton G E 2012 *Advances in Neural Information Processing Systems* pp 1097–105
[3] Simonyan K and Zisserman A 2015 (arXiv:1409.1556)
[4] Girshick R, Donahue J, Darrell T and Malik J 2014 *IEEE Conf. Computer Vision and Pattern Recognition* pp 580–7
[5] Redmon J, Divvala S, Girshick R and Farhadi A 2016 *IEEE Conf. Computer Vision and Pattern Recognition* pp 779–88
[6] Long J, Shelhamer E and Darrell T 2015 *IEEE Conf. Computer Vision and Pattern Recognition* pp 3431–40
[7] Taigman Y, Yang M, Ranzato M and Wolf L 2014 *IEEE Conf. Computer Vision and Pattern Recognition* pp 1701–8
[8] Schroff F, Kalenichenko D and Philbin J 2015 *IEEE Conf. Computer Vision and Pattern Recognition* pp 815–23
[9] Hofer L R, Jones L W, Goedert J L and Dragone R V 2019 *J. Opt. Soc. Am.* A **36** 936–43
[10] Ch'Ng K, Carrasquilla J, Melko R G and Khatami E 2017 *Phys. Rev.* X **7** 031038
[11] Carrasquilla J and Melko R G 2017 *Nat. Phys.* **13** 431–4
[12] Seif A, Landsman K A, Linke N M, Figgatt C, Monroe C and Hafezi M 2018 *J. Phys. B: At. Mol. Opt. Phys.* **51** 174006
[13] Picard L R, Mark M J, Ferlaino F and van Bijnen R 2019 *Meas. Sci. Technol.* **31** 025201
[14] Ness G, Vainbaum A, Shkedrov C, Florshaim Y and Sagi Y 2020 *Phys. Rev. Appl.* **14** 014011
[15] Miles C, Bohrdt A, Wu R, Chiu C, Xu M, Ji G, Greiner M, Weinberger K Q, Demler E and Kim E A 2020 (arXiv:2011.03474)
[16] Raab E, Prentiss M, Cable A, Chu S and Pritchard D E 1987 *Phys. Rev. Lett.* **59** 2631
[17] Chu S, Bjorkholm J, Ashkin A and Cable A 1986 *Phys. Rev. Lett.* **57** 314
[18] Pethick C J and Smith H 2008 *Bose–Einstein Condensation in Dilute Gases* (Cambridge: Cambridge University Press)
[19] Lett P D, Watts R N, Westbrook C I, Phillips W D, Gould P L and Metcalf H J 1988 *Phys. Rev. Lett.* **61** 169
[20] Muessel W, Strobel H, Joos M, Nicklas E, Stroescu I, Tomkovič J, Hume D B and Oberthaler M K 2013 *Appl. Phys.* B **113** 69–73
[21] van der Walt S, Schönberger J L, Nunez-Iglesias J, Boulogne F, Warner J D, Yager N, Gouillart E, Yu T and (The Scikit-image Contributors) 2014 *PeerJ* **2** e453
[22] Shahriari B, Swersky K, Wang Z, Adams R P and de Freitas N 2016 *Proc. IEEE* **104** 148–75
[23] Snoek J, Larochelle H and Adams R P 2012 *Advances in Neural Information Processing Systems* vol 25 pp 2951–9
[24] Frisch A, Aikawa K, Mark M, Rietzler A, Schindler J, Zupanič E, Grimm R and Ferlaino F 2012 *Phys. Rev.* A **85** 051401
[25] Hofer L R, Krstajić M, Juhász P, Marchant A L and Smith R P 2021 *Atom Cloud Detection and Segmentation Using a Deep Neural Network (Dataset) Version 1* (https://doi.org/10.5287/bodleian:Y56kydgmj)
[26] He K, Gkioxari G, Dollár P and Girshick R 2017 *IEEE Int. Conf. Computer Vision* pp 2980–8
[27] He K, Zhang X, Ren S and Sun J 2016 *IEEE Conf. Computer Vision and Pattern Recognition* pp 770–8
[28] Zeiler M D and Fergus R 2014 *Computer Vision—ECCV* pp 818–33
[29] Ren S, He K, Girshick R and Sun J 2015 *Advances in Neural Information Processing Systems* vol 28 pp 91–9
[30] Hecht-Nielsen R 1989 *Int. 1989 Joint Conf. Neural Networks* vol 1 pp 593–605
[31] Zhao H, Gallo O, Frosio I and Kautz J 2017 *IEEE Trans. Comput. Imaging* **3** 47–57
[32] Feng Z, Kittler J, Awais M, Huber P and Wu X 2018 *IEEE Conf. Computer Vision and Pattern Recognition* pp 2235–45
[33] Goodfellow I, Bengio Y and Courville A 2016 *Deep Learning* (Cambridge, MA: MIT Press)
[34] Keren G, Sabato S and Schuller B 2018 *IEEE Int. Conf. Data Mining* pp 227–36
[35] Bottou L 1999 *On-line Learning and Stochastic Approximations* (Cambridge: Cambridge University Press) pp 9–42
[36] Jastrzebski S, Kenton Z, Arpit D, Ballas N, Fischer A, Bengio Y and Storkey A 2018 (arXiv:1711.04623)
[37] Bengio Y 2012 *Practical Recommendations for Gradient-Based Training of Deep Architectures* (Berlin: Springer) pp 437–78
[38] Everingham M, Van Gool L, Williams C K, Winn J and Zisserman A 2010 *Int. J. Comput. Vis.* **88** 303–38
[39] Boyd K, Eng K H and Page C D 2013 *Joint European Conf. Machine Learning and Knowledge Discovery in Databases* pp 451–66
[40] Cordts M, Omran M, Ramos S, Rehfeld T, Enzweiler M, Benenson R, Franke U, Roth S and Schiele B 2016 *IEEE Conf. Computer Vision and Pattern Recognition* pp 3213–23

[41] Bisong E 2019 *Google Collaboratory* (Berlin: Springer) pp 59–64
[42] Paszke A *et al* 2017 *31st Conf. Neural Information Processing Systems*
[43] Lin T Y, Maire M, Belongie S, Hays J, Perona P, Ramanan D, Dollár P and Zitnick C L 2014 *Computer Vision—ECCV* pp 740–55
[44] Yosinski J, Clune J, Bengio Y and Lipson H 2014 *Advances in Neural Information Processing Systems* vol 27 pp 3320–8
[45] Bergstra J and Bengio Y 2012 *J. Mach. Learn. Res.* **13** 281–305
[46] Eggensperger K, Feurer M, Hutter F, Bergstra J, Snoek J, Hoos H and Leyton-Brown K 2013 *NIPS Workshop on Bayesian Optimization in Theory and Practice* vol 10 p 3
[47] Frazier P I 2018 (arXiv:1807.02811)
[48] Brochu E, Cora V M and de Freitas N 2010 (arXiv:1012.2599)
[49] Sobol' I M 1967 *Zh. Vychisl. Mat. Mat. Fiz.* **7** 784–802
[50] Balandat M, Karrer B, Jiang D R, Daulton S, Letham B, Wilson A G and Bakshy E 2020 (arXiv:1910.06403)
[51] Redmon J and Farhadi A 2018 (arXiv:1804.02767)
[52] Hofer L R, Dragone R V and MacGregor A D 2017 *Opt. Eng., Bellingham* **56** 043110
[53] ISO 11146-3 2004 *Lasers and Laser-Related Equipment—Test Methods for Laser Beam Widths, Divergence Angles and Beam Propagation Ratios* (International Organization for Standardization)
[54] Fiorio C and Gustedt J 1996 *Theor. Comput. Sci.* **154** 165–81
[55] Bloch I 2005 *Nat. Phys.* **1** 23–30
[56] Viebahn K, Sbroscia M, Carter E, Yu J C and Schneider U 2019 *Phys. Rev. Lett.* **122** 110404
[57] Gaunt A L, Schmidutz T F, Gotlibovych I, Smith R P and Hadzibabic Z 2013 *Phys. Rev. Lett.* **110** 200406
[58] Anderson M H, Ensher J R, Matthews M R, Wieman C E and Cornell E A 1995 *Science* **269** 198–201
[59] Davis K B, Mewes M O, Andrews M R, van Druten N J, Durfee D S, Kurn D and Ketterle W 1995 *Phys. Rev. Lett.* **75** 3969

4

Bose–Einstein condensation of erbium

University of Cambridge. *Gules, on a cross ermine between four lions passant gardant Or, a Bible fesswise Gules, clasped and garnished Or, the clasps in base.*

In this Chapter we present the attainment of a Bose–Einstein condensate of erbium in our apparatus. We first discuss evaporative cooling and methods for detecting and measuring BECs. With the knowledge of the loss landscape (presented in Chapter 5), we optimise the production of dipolar ^{166}Er BECs, achieving more than 2×10^5 atoms in the condensate, which we describe in detail.

4.1 Introduction

In order to produce a Bose–Einstein condensate of erbium, we employ standard laser cooling and trapping techniques as described in Chapter 3. Let us briefly revisit these, after which we can discuss our BEC production sequence in detail. In the initial steps, an atomic beam from a high-temperature effusion cell oven is slowed down using a Zeeman slower operating on the broad transition at 401 nm. The slow atoms are loaded into a narrow-line magneto–optical trap (MOT) operating on the atomic transition at 583 nm, and we typically capture 10^8 atoms after loading the MOT for 12 s. Afterwards, we ramp to a compressed MOT (cMOT) configuration in 600 ms, where reducing the light detuning and intensity causes simultaneous compression and cooling, resulting in a spin-polarised atomic sample at a temperature of 10 µK.

As the critical temperature for Bose–Einstein condensation is less than 1 µK, we need to cool the sample further. To achieve this, we transfer the atoms into an optical dipole trap (ODT) in which we perform evaporative cooling, presented in §4.2. In §4.3 we discuss techniques to detect and measure BECs and we present our experimental sequence in §4.4.

4.2 Evaporation

Recall that to achieve Bose–Einstein condensation, the peak phase-space density of the gas needs to reach $\rho = n\lambda_{\mathrm{dB}}^3 \approx 2.612$ (cf. §2.3). To calculate the peak ρ for our system, we make use of the density distribution of a thermal gas in a harmonic trap (well above the BEC transition temperature), given by

$$n(\mathbf{r}) = n(x, y, z) = n_0 \exp\left(-\frac{m\left(\omega_x^2 x^2 + \omega_y^2 y^2 + \omega_z^2 z^2\right)}{2k_B T}\right), \tag{4.1}$$

where ω_x, ω_y and ω_z are the trap frequencies and n_0 is the peak density [102, 133, 248]. If the number of the atoms in the trap is N, the peak density is fixed by the normalisation condition $N = \int n(\mathbf{r})\, \mathrm{d}^3\mathbf{r}$, which yields

$$n_0 = N\bar{\omega}^3 \left(\frac{m}{2\pi k_B T}\right)^{3/2}, \tag{4.2}$$

where $\bar{\omega} = \sqrt[3]{\omega_x \omega_y \omega_z}$ is the geometric mean of the trap frequencies. Therefore, the peak ρ can be calculated as

$$\rho_0 = n_0 \lambda_{\mathrm{dB}}^3 = N\left(\frac{\hbar\bar{\omega}}{k_B T}\right)^3. \tag{4.3}$$

After loading the ODT, our parameters yield $\rho_0 = 2.7 \times 10^{-4}$. To gain the required four orders of magnitude in phase-space density, we subject the atoms to the final stage of cooling, evaporation [133, 249], which takes place in the optical dipole trap (ODT). The surprisingly simple working principle of this technique is shown in Fig. 4.1. As the ODT trap depth is lowered (by decreasing the laser intensity), the most energetic atoms can escape the trap. As these atoms have a higher average energy than the average of all atoms, the average energy of the remaining atoms decreases. As the gas rethermalises via collisions, this leads to a lower temperature, i.e. cooling. Interestingly, this technique has no theoretical limit regarding the lowest achievable temperature, and in 2003, the group of Wolfgang Ketterle at MIT cooled a gas of ^{23}Na to a mere 0.5 nK above absolute

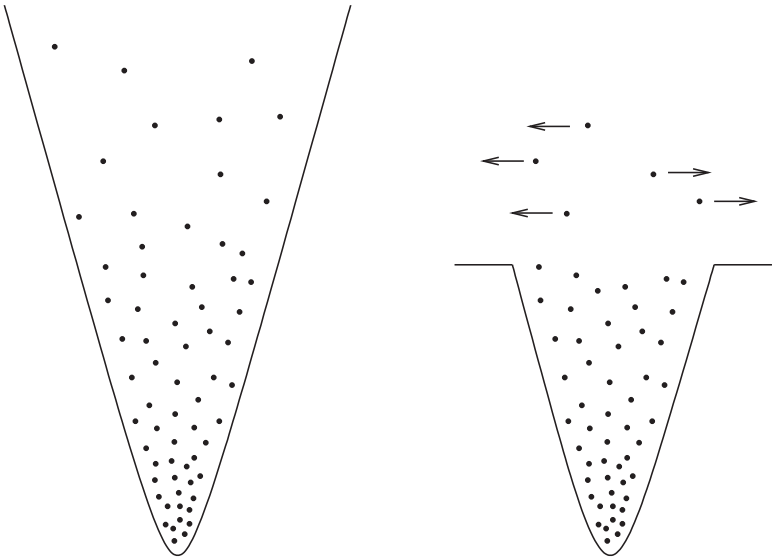

Figure 4.1. Evaporative cooling. As the trap is made shallower, the most energetic atoms escape leaving the remaining atoms colder on average, once they rethermalise via collisions. Figure adapted from Ref. 133.

zero [250]. However, we should recognise the fact that cooling comes at the expense of losing atoms from the system.

The method can be understood using a simple model [133]. Let us consider a small step in the evaporation process, during which the atom number changes by $\mathrm{d}N$ and the temperature by $\mathrm{d}T$. It can be shown that the total energy (kinetic and potential) of a thermal gas in a 3D harmonic trap is $E = 3Nk_BT$ [248]. (This can simply be understood via the equipartition theorem, as $k_BT/2$ energy resides in each of the three translational and the three vibrational degrees of freedom.) The energy of the atoms after an evaporation step can be written as the sum of the energy of the atoms before evaporation, $3Nk_BT$, and the energy taken away by the evaporated atoms, $\eta k_BT\,\mathrm{d}N$, where η is a free parameter relating to the average energy of the evaporated atoms, governed by how high in energy we 'cut' the trap. Therefore,

$$3k_B(N + \mathrm{d}N)(T + \mathrm{d}T) = 3Nk_BT + \eta k_BT\,\mathrm{d}N\,, \tag{4.4}$$

which leads to

$$\frac{\mathrm{d}T}{T} = \frac{\eta - 3}{3}\frac{\mathrm{d}N}{N}\,, \tag{4.5}$$

i.e. $T \propto N^{\beta}$ where $\beta = (\eta - 3)/3$. We see that for cooling to take place, we need $\beta > 0$, i.e. $\eta > 3$. Furthermore, for the largest fractional change in T, we see β should be as high as possible, but letting only the most energetic atoms escape the trap means that only a few atoms would leave each time and evaporation would take very long. This is problematic,

as there are other atom loss processes taking place at the same time which do not lead to cooling (e.g. collisions with the background gas, or three-body loss which leads to *heating*). Therefore, we need to evaporate quickly enough to not lose the atoms altogether.

To progress towards Bose–Einstein condensation, we also need to ensure that the phase-space density increases during evaporation. We can estimate the size of the cloud R by equating the potential energy with the kinetic energy (in a given direction), $m\omega^2 R^2/2 = k_B T/2$, which leads to $R \propto \sqrt{T}$. Therefore, $n \approx N/R^3 \propto N^{1-3\beta/2}$. As $\lambda_{\mathrm{dB}}^3 \propto T^{-3/2} \propto N^{-3\beta/2}$, we find $\rho = n\lambda_{\mathrm{dB}}^3 \propto N^{1-3\beta}$. Therefore, for the phase-space density to increase during evaporation, we need $\beta > 1/3$, i.e. $\eta > 4$.

Besides this, we also need to make sure that the collision rate is high enough to ensure that thermalisation takes place effectively. The collision rate is given by $\sqrt{2}n\sigma\bar{v}$, where σ is the collision cross section and $\bar{v}$ is the average velocity of the atoms. Given $\bar{v} \propto \sqrt{T} \propto N^{\beta/2}$, we find $n\sigma\bar{v} \propto N^{1-\beta}$. Therefore, to maintain efficient thermalisation, we need $\beta > 1$, i.e. $\eta > 6$. Such η can be achieved in experiments, and we can reach BEC with our apparatus after 7 s of evaporation.

4.3 Detection and measurement of Bose–Einstein condensates

To determine the phase-space density (and to confirm that we indeed achieved Bose–Einstein condensation) we measure T and N via absorption imaging after a time of flight (see §3.6.2). This yields us the areal density $n_{2\mathrm{D}}(x, y)$, from which the atom number can be deduced by integration.

The temperature of the gas can be determined from the time evolution of the shape of the spatial wings of the distribution ascribed to the thermal cloud, using the 'time of flight' (ToF) technique, where the atoms are released from the trap and imaged after a varying time t. (During this time the cloud expands and falls under gravity.) The density in the wings of the spatial distribution decays generally as $e^{-r^2/\sigma^2(t)}$, even if the gas is close to (or below) the critical temperature. (Note that from here on, we redefine σ to be a parameter describing the density distribution, and it is not the scattering cross section.) This evolves in time according to [251, 252]

$$\sigma^2(t) = \sigma^2(0) + \frac{k_B T}{m} t^2. \tag{4.6}$$

The parameter σ for a thermal cloud (far away from T_c) released from a harmonic trap can be determined by fitting a Gaussian function to the image (as the Gaussian density

distribution of a thermal cloud in a harmonic trap remains Gaussian after expansion). However, close to and below the critical temperature, the density distribution of the thermal component (the atoms not in the condensate) within the trap is given by [102]

$$n_{\mathrm{th}}(\mathbf{r}) = \frac{1}{\lambda_{\mathrm{dB}}^3} g_{3/2}\left(e^{-\beta U(\mathbf{r})}\right), \tag{4.7}$$

where $U(\mathbf{r})$ is the trap potential, $\beta = 1/k_B T$ as before and $g_{3/2}(x)$ is a so-called Bose function,[1] defined by $g_j(z) = \sum_{l=1}^{\infty} z^l/l^j$. When such an *in situ* density distribution is released from the trap, after free expansion the density distribution becomes [102]

$$n_{\mathrm{exp}}(\mathbf{r}, t) = \frac{1}{\lambda_{\mathrm{dB}}^3} g_{3/2}\left(e^{-\beta \tilde{U}(\mathbf{r},t)}\right) \prod_{i=x,y,z} \frac{1}{\sqrt{1+\omega_i^2 t^2}}, \tag{4.8}$$

where $\tilde{U}(\mathbf{r}, t) = \sum_{i=x,y,z} m\omega_i^2 r_i^2 / 2\left(1+\omega_i^2 t^2\right)$.

As in our measurements we can only measure the sum of the BEC and the thermal density distribution, we need to establish how the two components contribute to the sum. (It is important to note that the BEC in itself does not have a temperature, a temperature can only be ascribed to the whole system containing the BEC and the non-condensed thermal part.) In the trap, the shape of the BEC is given by a parabola within the Thomas–Fermi approximation (cf. Eq. (B.18)), and it can be shown that even after free expansion it retains a parabolic shape [64]. In our images, as we image along a certain direction, we measure the areal density along the direction of propagation of the imaging beam (let us call this z), so the 3D density distributions have to be integrated along this direction for fitting. These lead to a bimodal areal density distribution, which can be parametrised as [252]

$$n_{\mathrm{2D}}(x, y) = \frac{\tilde{n}_{\mathrm{th}}}{g_2(1)} g_2\left(e^{-x^2/2\sigma_x^2 - y^2/2\sigma_y^2}\right) + \tilde{n}_c \max\left(1 - \frac{x^2}{\tilde{R}_x^2} - \frac{y^2}{\tilde{R}_y^2}, 0\right)^{3/2}, \tag{4.9}$$

where the first term corresponds to the thermal component and the second one to the BEC. Here $\tilde{n}_{\mathrm{th}}$ and $\tilde{n}_c$ are the density amplitudes associated with the thermal component and the condensate, respectively, and $\sigma_{x,y}$ and $\tilde{R}_{x,y}$ characterise the size of the thermal component and the condensate after expansion, respectively ($\tilde{R}_{x,y}$ are called the scaled

[1] This function is more widely known as the polylogarithm function, with the more conventional notation $\mathrm{Li}_j(z) = g_j(z)$ outside quantum mechanics. It arises frequently within ultracold atoms as it is the intergral of the Bose–Einstein distribution,

$$\mathrm{Li}_j(z) = \frac{1}{\Gamma(j)} \int_0^{\infty} \frac{t^{j-1}}{e^t/z - 1}\, \mathrm{d}t\,.$$

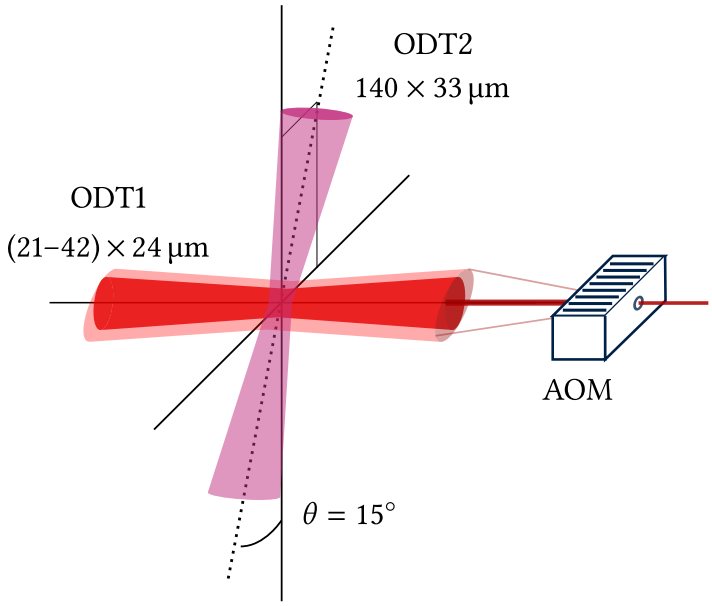

Figure 4.2. Crossed optical dipole trap. The ODT is implemented using two crossed, far-detuned beams at 1030 nm. The horizontal (ODT1) beam can be enlarged by dithering it with an AOM, e.g. to achieve a better overlap with the cMOT. The cross (ODT2) beam propagates 15° to the vertical.

Thomas–Fermi radii). It is this $\sigma_{x,y}$ that we need to use for determining the temperature via Eq. (4.6). For large x and y, i.e. $w = e^{-x^2/2\sigma_x^2 - y^2/2\sigma_y^2} \ll 1$, we find

$$g_2(w) = \int_0^\infty \frac{t}{e^t/w - 1}\,\mathrm{d}t \approx \int_0^\infty \frac{t}{e^t/w}\,\mathrm{d}t = w\,, \tag{4.10}$$

so the areal density of the thermal part still scales as $e^{-x^2/2\sigma_x^2 - y^2/2\sigma_y^2}$, the same as before the Bose enhancement.

4.4 Bose–Einstein condensation of ^{166}Er

Our optical dipole trap is implemented using two crossed (roughly orthogonal), far-detuned beams at 1030 nm, which we call ODT1 and ODT2, respectively (this is known as a 'crossed trap'). In our system, ODT1 provides the majority of trapping, and the effect of ODT2 becomes significant only towards the last stages of evaporation, when it ensures a sufficiently high trapping frequency in the propagation direction of the horizontal ODT beam to prevent the spread of atoms in that direction. The system is shown in Fig. 4.2.

Initially, the 21 µm × 24 µm waist ODT1 beam is superimposed onto the cMOT, with a total power of 21 W and with a 50 kHz spatial dithering applied in the horizontal direction using an acousto-optic modulator (AOM) [235], which broadens the horizontal (21 µm)

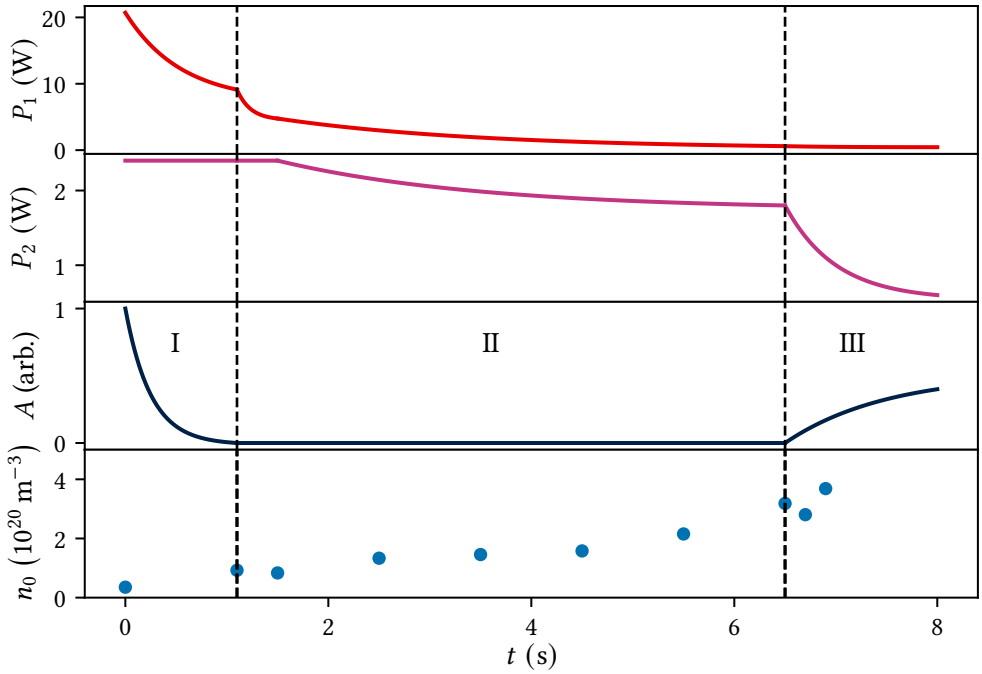

Figure 4.3. BEC production. The evaporation sequence consists of three parts. During phase I, the trap depth is lowered in parallel with ramping down the dithering. In phase II, the atoms are drawn into the crossing with the second ODT beam as the cooling continues. Finally, in phase III, the entire atomic density is in the crossed trap. P_1 and P_2 correspond to the power in ODT1 and ODT2, respectively, the dithering amplitude A is shown in arbitrary units (as it does not have a linear correspondence to the beam waist) and n_0 is the resulting peak atom number density.

waist by a factor of two. (This dithering is applied to increase the spatial overlap between the cMOT and the ODT.) We transfer 1.8×10^7 atoms into the ODT during the 40 ms overlap period. After the transfer, the temperature of the atoms is 37.5 µK, i.e. the sample heats up as it is moved from one trap to the other.

We then proceed with the evaporation sequence using a magnetic field of 1.4 G, as this corresponds to the region of lowest three-body loss across the temperature range encountered during evaporation (see Chapter 5). The BEC production sequence is shown in Fig. 4.3 and can be split into three stages. In stage I, we evaporate from the crossed trap while simultaneously reducing the horizontal trap depth and changing the aspect ratio by ramping down the dithering.[2] The ODT2 beam has a waist of 140 µm × 33 µm and an initial power of 2.4 W. In stage II, as we continue ramping down the ODT powers (both beams), the cooling continues and the remaining atoms go into the crossing of the ODT beams.

To characterise our cooling sequence towards quantum degeneracy, the peak phase-

[2] We change ODT powers and the dithering amplitude using exponential functions, to create a smooth ramp.

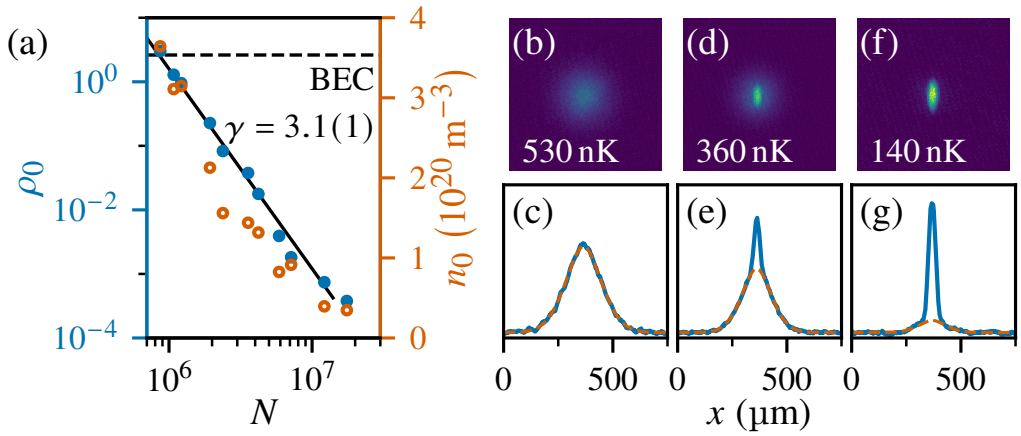

Figure 4.4. Evaporation to BEC. (a) Evolution of the peak phase-space density ρ_0 (blue circles) and atom number density n_0 (orange circles) with the total atom number N during evaporation. A linear fit on the logarithmic plot shows that $\gamma = 3.1(1)$ orders of magnitude in ρ_0 are gained at the expense of an order of magnitude reduction in the atom number. (b–f) Time-of-flight absorption images and summed density plots of atomic clouds: (b, c) just prior to condensation, corresponding to the highest-ρ_0 point in (a); (d, e) a partially condensed cloud; (f, g) a nearly-pure BEC with 2.2×10^5 atoms in the condensate at 140 nK. The time of flight was 18 ms for (b) and (d), and 24 ms for (f). The solid lines are the summed densities and the dashed lines are fits of the extended Bose distribution to the thermal component of the clouds, from which the temperature can be extracted. We see that a sharp peak appears on top of the thermal distribution, which confirms Bose–Einstein condensation in the cloud.

space density (given by Eq. (4.3)) is shown as a function of the atom number in the trap in Fig. 4.4(a) (note that the evaporation proceeds from right to left in this figure, as the atom number decreases during evaporation). After loading the ODT we have $N = 1.8 \times 10^7$ atoms at $T = 37.5\,\mu\text{K}$ with $\rho_0 = 3.8 \times 10^{-4}$ (the lowest-ρ point in Fig. 4.4(a)), whereas towards the end of the ramp, we have $N = 8.7 \times 10^5$ atoms at $T = 450\,\text{nK}$ with $\rho_0 = 3.02$ (the highest-ρ point in Fig. 4.4(a)). We see that during this process we lose ca. 95% of our atoms, but we gain four orders of magnitude in the peak phase-space density, and cross the boundary of Bose–Einstein condensation (at $\rho_0 \approx 2.612$).

The evaporation efficiency can be characterised by the order of magnitude increase in the phase-space density at the expense of an order of magnitude decrease in the atom number [249]. This is given by

$$\gamma = -\frac{\mathrm{d}(\ln(\rho_))}{\mathrm{d}(\ln(N))} = -\frac{N}{\rho_0}\frac{\mathrm{d}\rho_0}{\mathrm{d}N}\,. \tag{4.11}$$

Within the model presented in §3.4.6, for a harmonically trapped thermal gas we find the

surprisingly simple result $\gamma = \eta - 4$ via Eq. (4.5). In our case, by fitting Eq. (4.11) to the data we find $\gamma = 3.1(1)$, which leads to $\eta = 7.1(1)$. This confirms that we evaporate efficiently.

Once we evaporate further, we pass the BEC threshold and a condensate develops. From a detection point of view, the condensate manifests itself as a 'bump' on the Bose-enhanced Gaussian distribution of the thermal component, and its observation confirms Bose–Einstein condensation in the system. The development of the bimodal distribution (Eq. (4.9)) can be seen in Figs. 4.4(b)–4.4(g), which shows that an erbium BEC develops as the temperature is lowered (and the phase-space density is increased). We observe the onset of condensation at 500 nK towards the end of stage II, so in stage III of the evaporation, when the atoms fully reside in the crossed trap, we ramp the dithering up again to increase the trap size and thereby decrease the density of the gas. This is important as the large density of the condensate would otherwise lead to excessive three-body losses. At the end of stage III, at 140 nK we achieve a nearly pure condensate with 2.2×10^5 atoms. Therefore, we achieve higher atom numbers than what have been used so far (1.4×10^5 [207]) in experiments probing the roton regime of dipolar physics.

5

Characterisation of three-body loss processes

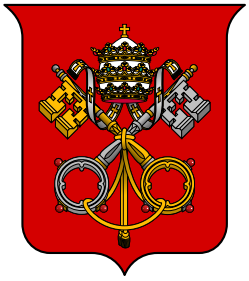

The Holy See. *Gules, two keys in saltire, that in bend Or and that in bend sinister Argent, wards pointing upwards and depicting the Cross of Our Lord therein, the handles interlaced of a cordon of the second, all surmounted by a tiara Argent triple crowned Or.*

The first experimental realisation of dipolar quantum droplets and supersolids in ultracold gases of highly magnetic atoms have been reported recently, and their exact structure and properties are the subject of intense research. However, studies have been limited by the achievable atom numbers and hindered by high three-body loss rates. In this Chapter we present a study of density-dependent atom loss in ultracold ^{166}Er, identifying several previously unknown features which display a strong temperature dependence, suggesting a higher partial-wave character. The detailed knowledge of the loss landscape enables the optimisation of the production of dipolar BECs, as presented in Chapter 4.

5.1 Introduction

The ability of precisely knowing and controlling the nature and strength of the interparticle interactions has been a key factor in the success of using degenerate ultracold-atom samples for studying many-body quantum phenomena. As shown in §2.2.4, the application of a magnetic field close to a Feshbach resonance is a highly versatile, convenient and well-established tool for tuning the sign and strength of *s*-wave contact interactions that typically dominate in ultracold gases [58]. Furthermore, Feshbach resonances can be utilised for the formation of diatomic molecules and the possibility of realising ultracold molecular ensembles holds great prospects for future research [96]. However, approaching a Feshbach resonance also leads to the enhancement of (detrimental) three-body loss processes, which result in heating and atom loss [168]. Therefore, knowing the location and width of Feshbach resonances and quantifying the associated losses are essential for designing and optimising ultracold-atom experiments.

The realisation of ultracold samples of highly magnetic erbium [61] and dysprosium atoms [60] has led to the discovery of dipolar quantum droplets [76, 77, 215] and a supersolid phase [81, 207, 253], which concurrently exhibits global phase order and spontaneous spatial density modulation. While first experiments were carried out in cigar-shaped traps leading to (relatively simple) one-dimensional (1D) spatial ordering, more recently droplet arrays and supersolids with two-dimensional (2D) ordering have also been observed [212, 254]. Theoretical works predict a plethora of possible novel patterns in 2D systems, including so-called honeycomb, labyrinthine and pumpkin phases [211, 213, 255–257]. However, reaching these exotic states requires degenerate samples with higher atom numbers than what have been used so far (1.4×10^5 [207]).[1]

The maximal achievable atom number in an experiment is often limited by three-body loss processes which limit the efficiency of evaporative cooling and result in losses at (or while approaching) the desired *s*-wave scattering length. In alkali atoms the (number) density of Feshbach resonances is typically between $0.1\,\mathrm{G}^{-1}$ and $0.01\,\mathrm{G}^{-1}$, and they are usually *s*-wave in character. However, in magnetic lanthanides (including erbium and dysprosium), the anisotropy of the van der Waals and the dipole–dipole interaction potentials leads to coupling between many scattering channels and the consequent abundance of Feshbach resonances [69, 169].

In this Chapter we present the theory behind three-body loss processes in §5.2, and in §5.3 we characterise losses for ^{166}Er for magnetic fields below 4 G, revealing the presence

[1] We note that Ref. 85 indicates that BECs with up to 10^6 atoms can be realised in that experiment, but that apparatus has not been used for exploring supersolid phases so far.

of six previously unreported resonant loss features which display a strong temperature dependence.

5.2 Loss processes

There are multiple processes which lead to atom loss from a trap. These can be characterised by their density dependence, which in turn depends on how many trapped atoms take part in the loss process. One-body loss arises e.g. due to a trapped atom colliding with an atom from the background gas (due to an imperfect vacuum), leading to a loss rate (per unit volume) proportional to $n(\mathbf{r})$, the number density of the trapped atoms. Two-body loss can arise due to non-elastic, spin-changing collisions between the trapped atoms [140], leading to loss proportional to $n^2(\mathbf{r})$. Finally, three-body loss arises when two particles form a molecule and a third takes away the excess energy and momentum,[2] leading to loss scaling with $n^3(\mathbf{r})$ [168]. The binding energy of the molecule formed during the three-body process is $E_{\text{bind}} = \hbar^2/ma_s^2 \approx k_B \times 100\,\mu\text{K}$ [138], which is much larger than the trap depth for a standard evaporation sequence (cf. the critical temperature is less than 1 μK). Therefore, when a molecule is formed and this energy is released as the kinetic energy of the particles participating in the collision, they leave the trap and are lost.

In a dilute gas, the leading loss term is one-body loss, which underlines the importance of creating ultra-high vacuum conditions, to minimise the number of background gas atoms interacting with the trapped gas. In our system, atoms are prepared in the lowest-energy spin-state at a low temperature, so two-body loss processes are energetically suppressed (see §2.2.3). While three-body loss can be negligible for a dilute gas, it can be significant for a denser sample, e.g. a condensate. Furthermore, three-body losses are also enhanced close to Feshbach resonances. Loss processes involving more atoms are also possible, but the likeliness of each process generally decreases with the number of atoms involved, so their contribution is normally small.[3]

The atom density loss rate for our system can therefore be written as

$$\dot{n}(\mathbf{r}) = -L_1 n(\mathbf{r}) - L_3 n^3(\mathbf{r})\,, \tag{5.1}$$

[2] It is important to note that it is this molecule formation which leads back to the true ground state of the ultracold erbium atom system, the solid phase.

[3] This is due to the fact that an increasing number of atoms need to be in the vicinity of each other, and for a dilute gas this is less and less likely. However, other phenomena can lead to the suppression of certain loss processes, and e.g. in our system even though two-body loss is suppressed, three-body loss still needs to be taken into account.

where $1/L_1 = \tau_1$ is the one-body lifetime and L_3 is the three-body loss coefficient.[4] The atom number loss can be found simply by

$$\dot{N} = \int \dot{n}(\mathbf{r})\, \mathrm{d}^3\mathbf{r}\,. \tag{5.2}$$

Interestingly, the three-body loss rate acquires a factor of $1/3!$ in a BEC, due to the quantum-correlated nature of the gas, accounting for the fact that density fluctuations are suppressed in a quantum gas [258] (the three-body process occurs between three identical bosons in the same single-particle state). The loss model for a BEC is presented in Appendix B for completeness.

5.3 Three-body loss

For our three-body loss measurements we prepare an ultracold, spin-polarised sample of ^{166}Er as described in §4.4; here we only note that the final stage of cooling is achieved by evaporation in the optical dipole trap (ODT), with the temperature T of the atom cloud set by the ODT depth. To measure clouds at different temperatures, we interrupt the normal evaporation sequence at different times and ramp up the depth of the ODT over 100 ms to achieve the desired temperature and to prevent any further evaporative cooling (and associated atom loss) during our measurements. We initiate the loss measurements by quenching the magnetic field B to the desired value.[5] To avoid ramping through wide resonances, for measurements above 3 G we evaporatively cool at 3.8 G, whereas for measurements below 3 G we cool at 1.4 G, the regions at which three-body loss is the smallest.

We measure the loss coefficient using thermal clouds at a range of temperatures by measuring the atom number N and temperature T as a function of the time t the atoms are held in the trap at a given magnetic field B using absorption imaging. The trap frequencies were (separately) measured by exciting dipole oscillations in the three perpendicular directions and τ_1 was independently determined to be $\tau_1 = 33.2(3)$ s from measurements of low-density clouds over much longer timescales (see Fig. 5.1).

5.3.1 Atom number decay

In Fig. 5.2 we show a typical $N(t)$ curve for a thermal gas. To understand the form of the curve and to extract the three-body loss coefficient, let us recall that for a thermal gas

[4] This equation is essentially a Taylor expansion of the loss rate in powers of n. The first-order term results in an exponential decay, so $1/L_1$ is the associated one-body lifetime.

[5] The magnetic field is calibrated using radio frequency spectroscopy within the ground state Zeeman manifold.

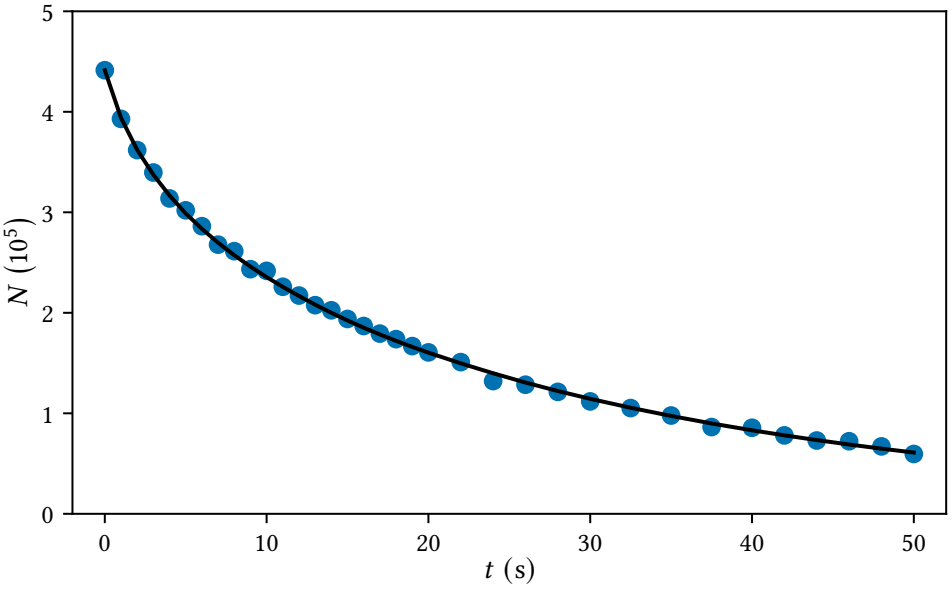

Figure 5.1. One-body lifetime. Measurement of the one-body lifetime τ_1 at a magnetic field of 1.4 G (to minimise three-body loss). Data points are indicated by blue circles, a fit of Eq. (5.3) with $\tau_1 = 33.2(3)$ s is shown with a black line.

in a harmonic trap (well above the BEC transition temperature), the density distribution is given by Eq. (4.1). Substituting this to Eq. (5.2), the evolution of the atom number is given by

$$\dot{N} = -\frac{1}{\tau_1} N - L_3 \left(\frac{m\bar{\omega}^2}{2\sqrt{3}\pi k_B} \right)^3 \frac{N^3}{T^3} \,. \tag{5.3}$$

The loss model can be extended by realising that due to the high kinetic energy of the lost particles, they can knock out further atoms on their way out of the trap. This contribution can be estimated with a simple approximation. If a collision product has a scattering cross section σ and velocity $v \gg \sqrt{k_B T/m}$ (i.e. the product is much faster than the gas particles), the average scattering rate with further gas particles is given by $\langle n \rangle \sigma v$. The average time the collision product spends in the cloud is on the order of $\langle r \rangle / v$, where $\langle r \rangle$ is the characteristic size of the cloud. Therefore, the collision product goes through $\langle n \rangle \langle r \rangle \sigma$ secondary collisions, and due to the high energies involved, the participating particles are also lost from the trap. This leads to

$$\dot{N} = -\frac{1}{\tau_1} N - L_3 \left(\frac{m\bar{\omega}^2}{2\sqrt{3}\pi k_B} \right)^3 \frac{N^3}{T^3} \left(1 + \frac{1}{3} \langle n \rangle \langle r \rangle (\sigma + \tilde{\sigma}) \right), \tag{5.4}$$

where the factor of $1/3$ accounts for the fact that the three-body collision event rate is one third of the original three-body atom loss rate (given each event leads to three atoms lost in our initial model) and we used the fact that the collision cross section for the

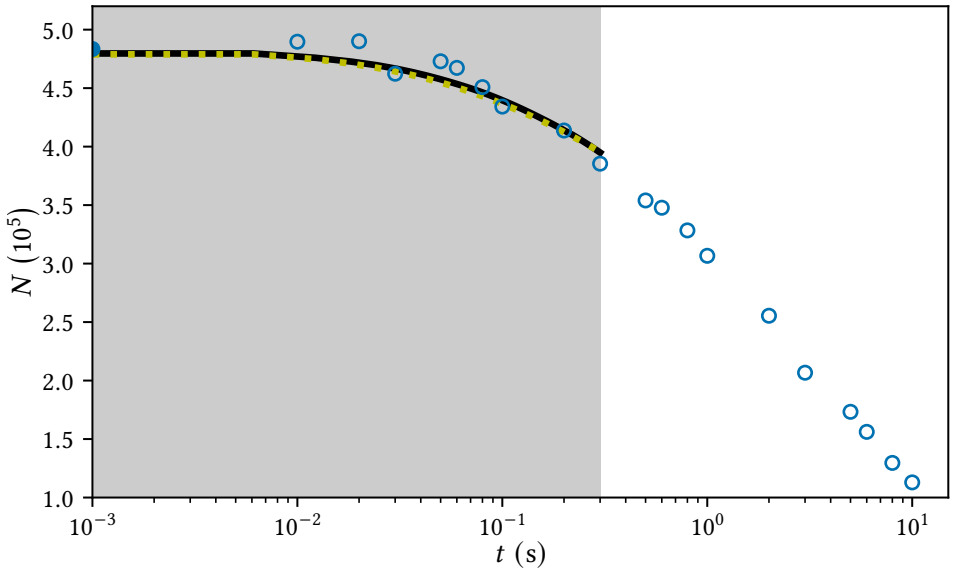

Figure 5.2. Time evolution of the number of atoms. Typical time evolution of the number of atoms in a thermal cloud at $B = 2.7$ G with an initial temperature of 0.5 µK. Empty markers represent data points and the filled marker is the initial number of atoms (its 0 time coordinate cannot be plotted on the logarithmic axis). The black line is a fit to the data according to Eq. (5.3), and the grey shading shows the points used for the fit, for which the temperature does not increase by more than 40% to avoid a fitting bias. The dashed line is a fit of Eq. (5.5).

Feshbach molecule ($\tilde{\sigma}$) and for the third particle (σ) is different. The characteristic size can be estimated via $3k_BT/2 \approx m\bar{\omega}^2\langle r\rangle^2/2$, whereas $\langle n\rangle$ can be estimated via Eq. (4.2). Putting all this together, the atom loss is given by

$$\dot{N} = -\frac{1}{\tau_1}N - L_3\left(\frac{m\bar{\omega}^2}{2\sqrt{3}\pi k_B}\right)^3 \frac{N^3}{T^3}\left(1 + G\frac{mN}{k_BT}\bar{\omega}^2(\sigma + \tilde{\sigma})\right), \tag{5.5}$$

where all factors of $O(1)$ (geometrical and otherwise) have been collected in the combined factor G. While G can be estimated (see Appendix B), $\sigma + \tilde{\sigma}$ is a new model parameter.

5.3.2 Heating

In Fig. 5.3 we show a typical $T(t)$ curve for a thermal gas. Besides an atom number decay, we see that the cloud also heats up. This heating is understood to be due to two processes. First, the loss rate is higher in the (higher density) central part of the trap, so lost atoms have a lower energy on average than the average of all the atoms in the trap, leading to 'anti-evaporation'. Second, secondary collisions also have an effect on the temperature evolution as well, via the lost energy of the secondary collision partners.

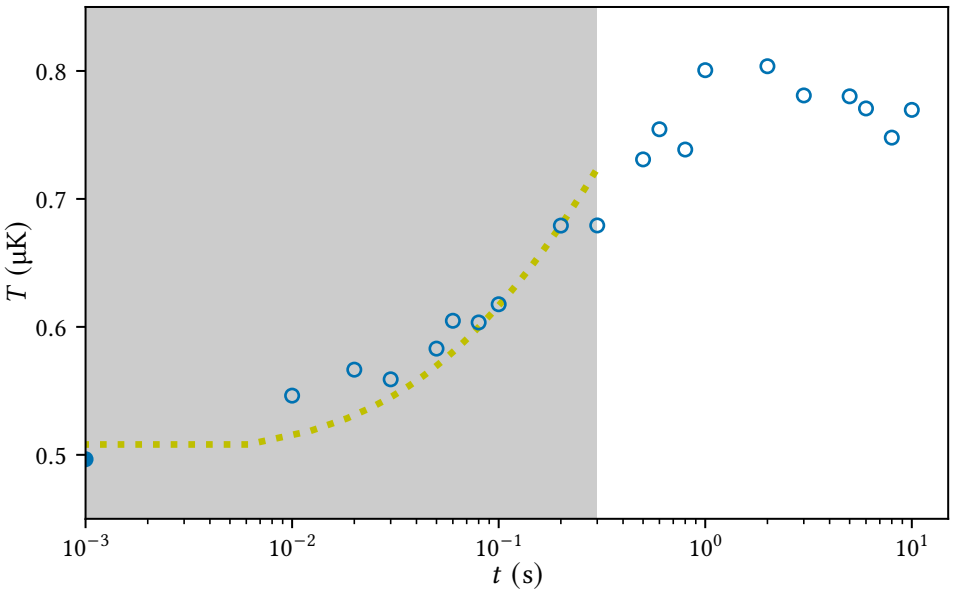

Figure 5.3. Time evolution of the temperature. Typical time evolution of the temperature of a thermal cloud at $B = 2.7\,\mathrm{G}$ with an initial temperature of 0.5 μK. Empty markers represent data points and the filled marker is the initial temperature of atoms (its 0 time coordinate cannot be plotted on the logarithmic axis). The dashed line is a fit to the data according to Eq. (5.9).

For a thermal cloud in a harmonic trap (with trapping potential $U(\mathbf{r}) = m\sum_i \omega_i^2 r_i^2/2$, where i labels the three spatial dimensions x, y and z), the temperature evolution can be obtained by considering the evolution of the energy of the cloud $E = 3Nk_BT$. On one hand, we see

$$\dot{E} = 3k_B(\dot{N}T + N\dot{T})\,. \tag{5.6}$$

However, using the loss rate and the energy lost at each event directly, we can also write

$$\dot{E} = \int \dot{n}(\mathbf{r})\left(U(\mathbf{r}) + \frac{3}{2}k_BT\right)\mathrm{d}^3\mathbf{r}\,. \tag{5.7}$$

Equating these two expressions for the energy loss and using Eq. (5.1) for $\dot{n}(\mathbf{r})$ leads to [168]

$$\dot{T} = \frac{L_3}{3}\left(\frac{m\bar{\omega}^2}{2\sqrt{3}\pi k_B}\right)^3\frac{N^2}{T^2}\,, \tag{5.8}$$

where we took the gas to be thermal with a density distribution given by Eq. (4.1). It is interesting to note that one-body loss does not lead to heating (as atoms at all energies are removed with an equal probability), but three-body loss does as expected.

Secondary collisions have an effect on the temperature evolution as well, via the lost energy of the secondary collision partners. (The average energy of the secondary

collision partners is not necessarily the same as the average energy in the cloud.) Using the extra number of lost particles in Eq. (5.5), we see

$$\dot{T} = \frac{L_3}{3}\left(\frac{m\bar{\omega}^2}{2\sqrt{3}\pi k_B}\right)^3 \frac{N^2}{T^2}\left(1 + G'\frac{mN}{k_B T}\bar{\omega}^2(\sigma + \tilde{\sigma})\right), \tag{5.9}$$

where we G' incorporates the (slightly different) numerical prefactors associated with the extra energy lost (see Appendix B for an estimation).

5.3.3 Fitting procedure and model comparison

To determine L_3, we measure N and T as a function of t. For a thermal cloud, the numerical solution of Eq. (5.3) is fit to the experimental data using the weighted least-squares method.[6] Fitting is done by varying the fit parameters to minimise

$$\chi^2 = \sum_i \frac{(N_{\text{data}}(t_i) - N_{\text{sol}}(t_i))^2}{\sigma_N^2(t_i)}, \tag{5.10}$$

where a 3% relative error was used[7] for estimating the standard deviation of the atom number measurements (i.e. $\sigma_N(t_i) = 0.03N_{\text{data}}(t_i)$). Furthermore, as our model assumes that L_3 is independent of temperature (which will be shown not to be the case), we only use those data points for which the temperature does not increase by more than 40% compared to the initial temperature, to avoid a large temperature change affecting our fit (this is shown in Figs. 5.2 and 5.3 as the shaded region).

To calculate the (numerical) solution of Eq. (5.3), $T(t)$ must be known. In these fits we use the measured T directly (by interpolating between the measured data points). However, when evaluating the extended loss model, we use its capability to predict both N and T (which both depend on each other), and so the fitting process is slightly more complicated as both of these are measured. In this case, fitting can be done by minimising the total $\chi^2 = \chi_N^2 + \chi_T^2$ by varying the fitting parameters, where χ_N^2 corresponds to the number data and χ_T^2 to the temperature data. (We measured $\sigma_T/T \approx \sigma_N/N \approx 3\%$ as well.) This approach is warranted by the fact that we really only have a single model which has an output of $N(t_i)$ and $T(t_i)$, and $\chi^2 = \sum_i(\text{data} - \text{prediction})^2/\sigma^2 = \chi_N^2 + \chi_T^2$.

[6] Fitting and numerical solution of the differential equations were done via the SCIPY package [259].

[7] This figure is based on the standard deviation of multiple measurements of the same data point—we found this was always ca. 3% of the data value itself. It is worth noting that as long as this is a small value (or constant across the data points), this does not have a considerable effect on the fitted L_3, it only affects the standard error of the fit.

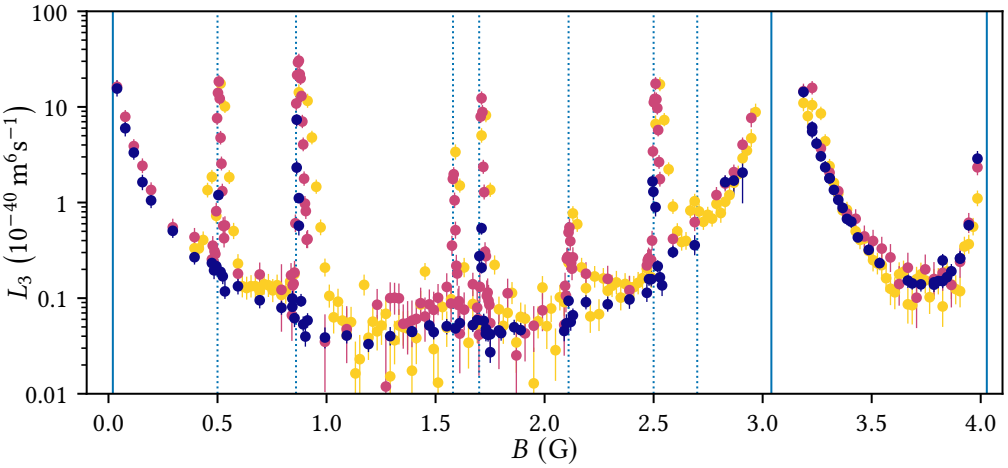

Figure 5.4. Three-body loss in ^{166}Er. Extracted three-body loss coefficients (L_3), and their standard errors. L_3 was measured at initial cloud temperatures of 0.5 μK, 1.5 μK and 4 μK. The colours ranging from darker to lighter represent temperature variation from colder to hotter samples. As the temperature increases, the peaks shift to higher B values and their widths increase, and indeed new peaks appear. Newly (previously) identified peaks are marked with dotted (solid) blue lines.

In both cases, the upper and lower standard error for the fit parameters are calculated by varying the given parameter up and down, respectively, to the point where χ^2 increases by 1, while optimising the other parameters [168, 260].

The models can be compared to each other by calculating the reduced χ^2, which we denote χ^2_r. (A better fit reduces χ^2_r; a good estimation of experimental errors should yield $\chi^2_r \approx 1$.) Without taking secondary collisions into account and using the measured temperatures directly, this is given by

$$\chi^2_r = \frac{\chi^2}{K-2}, \tag{5.11}$$

where K is the total number of measurements. This is reduced by 2 to account for the fact that we fix $N_{\mathrm{sol}}(0)$ to be the data point taken at $t = 0$ (i.e. this data point is not fitted) and that we have one fitting parameter (L_3, given τ_1 is fixed by earlier measurements). For the extended model, where we use its predictions for both N and T, the reduced χ^2 is given by

$$\chi^2_r = \frac{\chi^2_N + \chi^2_T}{2K-5}, \tag{5.12}$$

as here we fit $2 \times (K-1)$ data points (we fix $N_{\mathrm{sol}}(0)$ and $T_{\mathrm{sol}}(0)$ to be the data points taken at $t = 0$) and we have three fitting parameters (L_3, $G(\sigma + \tilde{\sigma})$ and $G'(\sigma + \tilde{\sigma})$).

The fits of the models are shown in Figs. 5.2 and 5.3. While they both fit the data well,

	1	2	3	4	5	6
B_0 (G)	0.50(3)	0.86(3)	1.58(3)	1.70(3)	2.11(3)	2.50(3)
Δ (mG)	6(12)	16(8)	12(7)	3(10)	18(12)	15(11)

Table 5.1. Newly detected Feshbach resonances. The locations B_0 and widths Δ of the resonances were determined by fitting Eq. (5.14) to the data taken at 1.5 µK. The width of the resonance is defined as the full width at half maximum (FWHM).

the non-extended model gives $\chi_r^2 = 0.68$ and the extended model gives $\chi_r^2 = 1.53$. This means that both models work well for predicting the atom number, but additionally we are able to reproduce the temperature evolution with the extended model.

Remember that the comparison of these models is not like-for-like, as we use the extended model to predict both the atom number and the temperature, while for the non-extended one we use the measured temperatures directly. This means that the extension predicts twice as many data points at the expense of two additional fit parameters. While secondary collisions do not yield a detectable difference in atom loss,[8] we found their contribution was essential to take into account for modelling the temperature evolution.[9]

5.3.4 Loss spectrum

Let us now discuss the spectrum of three-body coefficients. Figure 5.4 shows the measured three-body coefficients as a function of the magnetic field at various temperatures. Distinctive peaks in the coefficient represent a high atom loss, associated with Feshbach resonances. We observed that the position and width of the resonances change with temperature, and to explore this, L_3 was measured at a range of initial cloud temperatures. In Fig. 5.4, we show the data for initial temperatures of 0.5 µK, 1.5 µK and 4 µK, where colours ranging from dark to light denote a change in temperature from cold to hot.

For ^{166}Er, Feshbach resonances and associated loss features were reported[10] at 0.02(5) G, 3.04(5) G and 4.028 G, for magnetic fields below 4 G [170]. In our measurements, we identified six further features, whose position and width are listed in Table 5.1, for an initial cloud temperature of 1.5 µK.

[8] The two models give an L_3 within 2.4% of each other, while the error on the fitted value is 12%.

[9] It is worth noting that the extended model reaches $\chi_r^2 = 1.15$ on the $N(t)$ graph alone and $\chi_r^2 = 2.67$ on the $T(t)$ graph alone.

[10] The authors of Ref. 170 confirmed that the location of the highest resonance is 4.028 G and not 4.208 G, as can be seen from Fig. 3 of Ref. 170 (private communication). Furthermore, it is worth noting that Feshbach resonances were also reported at 0.121(1) G, 0.244(1) G, 0.366(1) G and 0.490(1) G [169], but these were later retracted [126].

5.3.5 Temperature dependence

To explore the temperature dependence further, we measured L_3 around the newly discovered resonance at 0.86(3) G for seven different initial temperatures, as shown in Fig. 5.5. As well as broadening with increasing temperature, the peak of L_3 shifts to higher B and decreases in amplitude.[11] To understand how temperature affects L_3, let us consider its dependence on the trimer collision energy E_3 and magnetic field B. This can be modelled by [69]

$$L_3(E_3, B) = \frac{48\pi^2\hbar^5}{\mu_3^3 E_3^2} \frac{(2\lambda+1)\Gamma(E_3)\Gamma_{\mathrm{br}}}{(E_3 - \mu(B-B_0))^2 + \Gamma_{\mathrm{tot}}^2(E_3)/4}, \tag{5.13}$$

where $\Gamma(E_3) = A_\lambda E_3^{\lambda+2}$, Γ_{br}, μ and B_0 are the entrance channel energy width, the trimer decay rate, the relative magnetic moment of the trimer and the entrance channel, and the position of the resonance, respectively. Furthermore, $\mu_3 = m/\sqrt{3}$ is the collisional reduced mass, $\Gamma_{\mathrm{tot}} = \Gamma(E_3) + \Gamma_{\mathrm{br}}$ and λ is the non-negative integer associated with the relative orbital angular momentum of the collision. We note that the long-range character of the dipole–dipole interaction allows non-s-wave entrance channels even at ultracold temperatures. At a particular temperature T, the collision energy E_3 is sampled from the Boltzmann distribution. Therefore, the measured L_3 is, in fact, an averaged value [69]:

$$L_3(T, B) = \frac{1}{2(k_B T)^3} \int_0^\infty E^2 L_3(E, B) e^{-E/k_B T} \, \mathrm{d}E \,. \tag{5.14}$$

In general, this integral cannot be solved analytically, so it was truncated at $E/k_BT = 10$ (an energy much higher than the average energy $3k_BT$) and evaluated numerically.

We found that our data were described well by a fit using $\lambda = 2$, denoting a d-wave entrance channel [261, 262], shown in Fig. 5.5(a).[12] We also show the dependence of the full width at half maximum of the resonances as a function of temperature in Fig. 5.5(b), and we found the width grew linearly with temperature. (This behaviour was confirmed for the other newly detected peaks as well.) Finally, we show the peak height as a function of temperature in Fig. 5.5(c), we observed these decayed with temperature.

[11] The peak position is also affected by the ODT light intensity, as described in Ref. 107, but this is not discussed in this thesis.

[12] Note, however, that other properties of this resonance point towards an s-wave character, as described in Ref. 107.

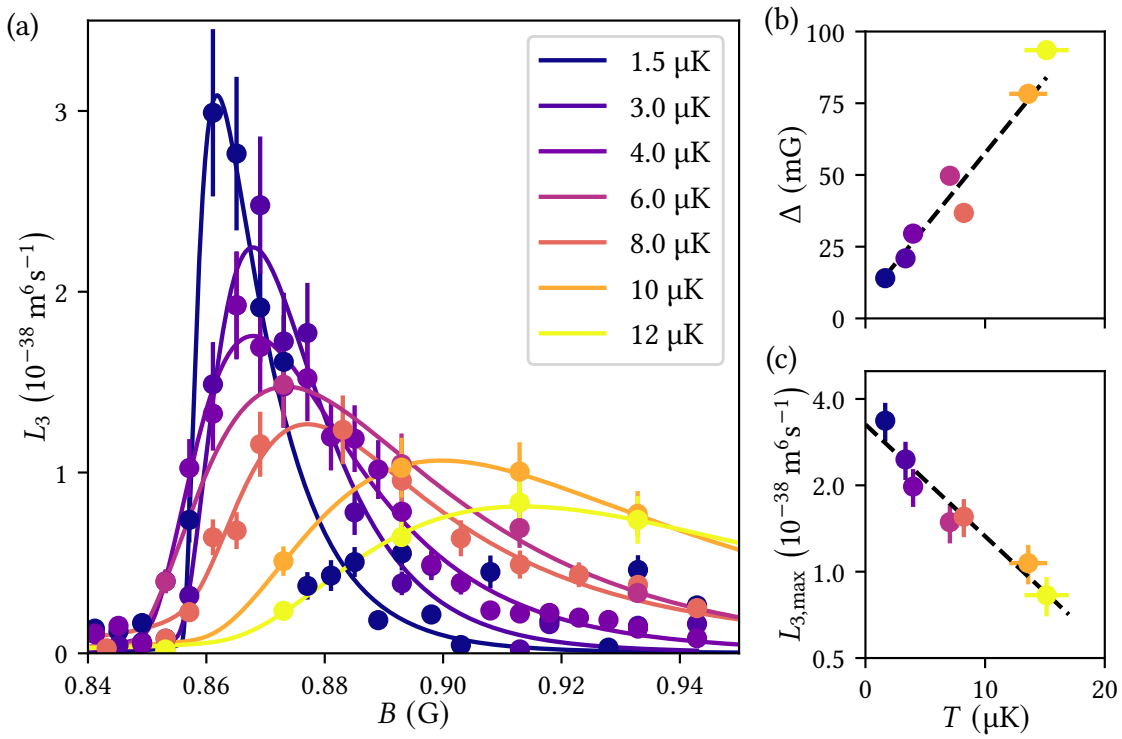

Figure 5.5. Temperature dependence of a Feshbach resonance. (a) The three-body loss coefficient L_3 is shown as a function of the magnetic field B around the newly discovered Feshbach resonance at 0.86(3) G for seven different initial temperatures (1.5 μK, 3 μK, 4 μK, 6 μK, 8 μK, 10 μK and 12 μK). The colours ranging from darker to lighter represent temperature variation from colder to hotter samples. Markers are the fitted loss coefficients L_3 with their errors, and lines of the same colour are the corresponding fits for a d-wave entrance channel. The peak of the resonance shifts to higher B-field values as the temperature increases. (b) The width (full width at half maximum) of the resonance against temperature. Marker colours correspond to the temperatures in (a). The width increases linearly with temperature (dashed line). (c) The peak value of L_3 against temperature. Marker colours correspond to the temperatures in (a). The peak decreases with increasing temperature, the dashed line is a guide to the eye.

5.4 Publication: Characterisation of three-body loss in ^{166}Er and optimised production of large Bose–Einstein condensates

Submitted for publication
DOI (arXiv): 10.48550/arXiv.2307.01245

Milan Krstajić,* Péter Juhász,* Jiří Kučera,* Lucas R. Hofer, Gavin Lamb, Anna L. Marchant and Robert P. Smith

Clarendon Laboratory, University of Oxford, Parks Road, Oxford OX1 3PU, United Kingdom

* The Authors contributed equally to this work.

Characterisation of three-body loss in ^{166}Er and optimised production of large Bose–Einstein condensates

Milan Krstajić,[*] Péter Juhász,[*] Jiří Kučera,[*] Lucas R. Hofer, Gavin Lamb, Anna L. Marchant,[†] and Robert P. Smith[‡]
Clarendon Laboratory, University of Oxford, Parks Road, Oxford, OX1 3PU, United Kingdom

Ultracold gases of highly magnetic lanthanide atoms have enabled the realisation of dipolar quantum droplets and supersolids. However, future studies could be limited by the achievable atom numbers and hindered by high three-body loss rates. Here we study density-dependent atom loss in an ultracold gas of ^{166}Er for magnetic fields below 4 G, identifying six previously unreported, strongly temperature-dependent features. We find that their positions and widths show a linear temperature dependence up to at least 15 μK. In addition, we observe a weak, polarisation-dependent shift of the loss features with the intensity of the light used to optically trap the atoms. This detailed knowledge of the loss landscape allows us to optimise the production of dipolar BECs with more than 2×10^5 atoms and points towards optimal strategies for the study of large-atom-number dipolar gases in the droplet and supersolid regimes.

I. INTRODUCTION

Precise knowledge and control of the nature and strength of interparticle interactions have been a key factor in the success of using degenerate ultracold-atom samples for studying many-body quantum phenomena. The application of a magnetic field close to a Feshbach resonance is a highly versatile and convenient tool for tuning the sign and strength of *s*-wave contact interactions that typically dominate in ultracold gases [1]. However, approaching a Feshbach resonance also leads to the enhancement of (detrimental) three-body processes, which result in atom loss and heating [2]. Knowing the location of Feshbach resonances and quantifying the associated loss features is thus essential for designing and optimising ultracold-atom experiments.

The realisation of ultracold samples of highly magnetic erbium [3] and dysprosium atoms [4], which interact via both long-range, anisotropic dipole–dipole interactions and tuneable contact interactions, has led to the discovery of dipolar quantum droplets [5–7] and a supersolid phase [8–10], which simultaneously exhibits a global phase order and a spontaneous spatial density modulation. While these first experiments were carried out in cigar-shaped traps leading to (relatively simple) one-dimensional (1D) spatial ordering, more recently droplet arrays and supersolids with two-dimensional (2D) ordering have also been observed [11, 12]. Theoretical works predict a plethora of novel patterns in 2D systems, including so-called honeycomb, labyrinthine and pumpkin phases [13–17]. However, reaching these exotic states requires degenerate samples with higher atom numbers than those used in these experiments so far (1.4×10^5 [10]).

The maximal achievable atom number in an experiment is often restricted by three-body loss processes, which limit the efficiency of evaporative cooling close to degeneracy and can greatly reduce the gas lifetime at (or while approaching) the desired *s*-wave scattering length. Moreover, in order to map out the parameter space of exotic dipolar phases, one needs to tune the relative strength of the contact and dipole–dipole interactions by controlling the strength of the magnetic field. The precise knowledge of the loss landscape as a function of the field strength is therefore paramount. Here we carefully characterise three-body loss in ^{166}Er for magnetic fields below 4 G, revealing the presence of six previously unreported resonant loss features which display a strong temperature dependence. In light of this, we describe our optimised procedure for the production of ^{166}Er Bose–Einstein condensates (BECs), containing more than 2×10^5 atoms.

II. THREE-BODY LOSS MEASUREMENTS

In alkali atoms the (number) density of Feshbach resonances is typically between $0.01\,\mathrm{G}^{-1}$ and $0.1\,\mathrm{G}^{-1}$ [1]. However, in magnetic lanthanides, including erbium and dysprosium, the anisotropy of the van der Waals and the dipole–dipole interaction potentials leads to coupling between many scattering channels and consequently to an abundance of Feshbach resonances [18–21], some of which show a strong temperature dependence [19, 22, 23]. Here we focus on ^{166}Er for magnetic fields below 4 G, where Feshbach resonances and associated loss features have been reported at 0.02(5) G, 3.04(5) G and 4.028 G [24].

For our measurements we prepare an ultracold, spin-polarised sample of ^{166}Er in an (approximately harmonic) optical dipole trap (ODT) formed from 1030 nm laser light. The experimental sequence is described in Sec. III, and the trap and gas parameters for all our loss measurements are given in the Supplemental Material [25]; here we only note that the final stage of cooling is achieved by evaporation in the ODT, with the temperature of the atom cloud controlled by the ODT depth. To produce clouds at different temperatures, we interrupt the normal evaporation sequence at different times and ramp up the depth of the ODT over 100 ms to prevent any further evaporative cooling (and associated atom loss) during our measurements. We initiate the loss measurements by quenching the magnetic field B [26] to the desired value in < 10 ms. To avoid ramping through wide resonances, for measurements above 3 G we evaporatively cool at 3.8 G, whereas for measurements below 3 G we cool at 1.4 G.

[*] M. K., P. J. and J. K. contributed equally to this work.
[†] Present address: STFC Rutherford Appleton Laboratory, Didcot, OX11 0QX, United Kingdom
[‡] robert.smith@physics.ox.ac.uk

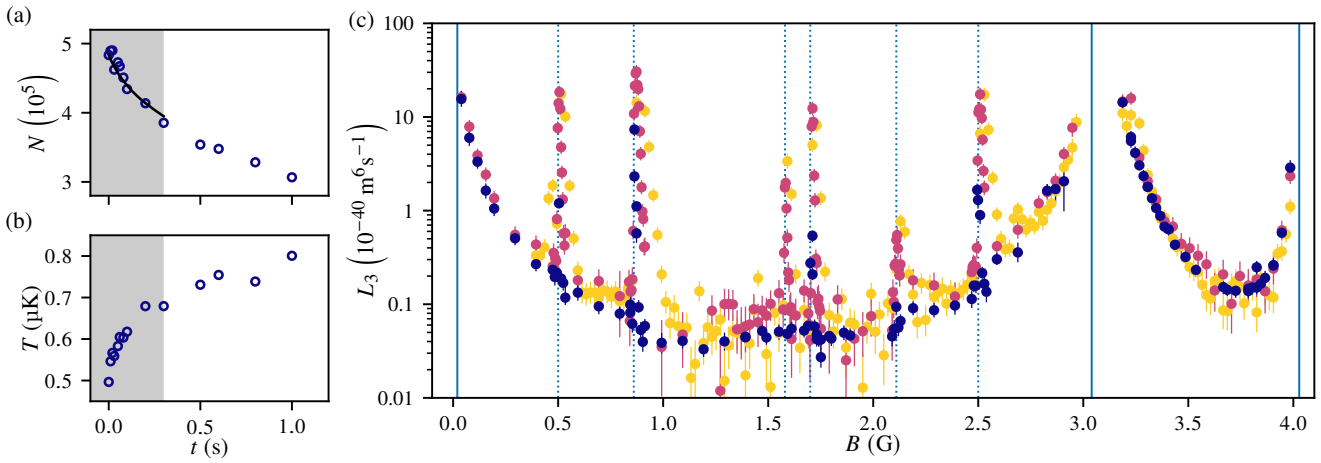

FIG. 1. Three-body loss in ^{166}Er. (a, b) Time evolution of the atom number N and temperature T, respectively, of a thermal cloud of ^{166}Er atoms at $B = 2.7$ G with an initial temperature of $T_i = 0.5$ μK. The shading denotes the region in which T is within 40% of T_i; in this region we fit $N(t)$ based on Eq. (2) to determine L_3 (solid line, see text). (c) Extracted three-body loss coefficients (L_3) for the different initial temperatures $T_i = \{0.5, 1.5, 4\}$ μK (blue, magenta and yellow points, respectively). In addition to the three previously identified Feshbach resonances (solid vertical lines), we identify six new loss features (dotted vertical lines). These new features show a noticeable temperature dependence; as T_i increases, the peaks shift to higher B and their widths also increase.

We use absorption imaging after a time-of-flight to measure the atom number N and temperature T as a function of the time t the atoms are held in the trap (at a given B). Examples of these $N(t)$ and $T(t)$ curves are shown in Figs. 1(a) and 1(b); here the initial temperature $T_i = 0.5$ μK and $B = 2.7$ G.

Let us first consider atom loss. As the atoms are prepared in the lowest Zeeman state at temperatures much lower than the sub-level splitting ($\approx 78\ \mu K\ G^{-1}$), two-body (spin relaxation) loss processes are energetically suppressed. The evolution of the atom number density in thermal samples can therefore be described by a combination of one- and three-body loss terms [2],

$$\dot{n}(\mathbf{r}) = -\frac{n(\mathbf{r})}{\tau_1} - L_3 n^3(\mathbf{r}), \qquad (1)$$

where τ_1 is the one-body lifetime (set by e.g. collisions with background gas atoms in an imperfect vacuum), L_3 is the three-body loss coefficient and $n(\mathbf{r})$ is the atom number density. For a thermal cloud (well above the BEC transition temperature), the atomic density distribution in a harmonic trap is Gaussian and Eq. (1) can be written as [2]

$$\frac{\dot{N}}{N} = -\frac{1}{\tau_1} - L_3 \left(\frac{m\bar{\omega}^2}{2\sqrt{3}\pi k_B T}\right)^3 N^2, \qquad (2)$$

where m is the atomic mass, $\bar{\omega}$ is the geometric mean of the trapping frequencies and k_B is the Boltzmann constant. The trapping frequencies were measured separately by exciting the cloud centre-of-mass oscillations in the three perpendicular directions and τ_1 was independently determined to be $\tau_1 = 33(1)$ s from measurements of low-density clouds for which three-body loss is negligible.

To determine $L_3(B)$ from our $N(t)$ measurements, we fit the numerical solution of Eq. (2) to our data [see solid line in Fig. 1(a)] using the corresponding measured $T(t)$ as an input. We only fit our data within the time interval over which the temperature stays within 40% of its initial value [gray shaded region in Figs. 1(a) and 1(b)] to limit any systematic errors arising from either (i) evaporative atom loss due to the finite trap depth or (ii) the fact that L_3 may depend on T [27].

Regarding the heating of the atom cloud [Fig. 1(b)], this can be understood to be due to two main processes. First, the loss rate is higher in the central (higher density) part of the trap, preferentially removing atoms with energy lower than the average energy in the cloud, leading to 'anti-evaporation' [2]. Second, the products of the three-body collision can have significant kinetic energy (acquired due to the released binding energy when two atoms form a molecule), which may be partially deposited in the cloud via secondary collisions.

Figure 1(c) shows the measured three-body coefficient as a function of the magnetic field for initial temperatures of 0.5 μK, 1.5 μK and 4 μK. In addition to the Feshbach resonances already reported [solid vertical lines in Fig. 1(c)], we observed six additional loss features (dotted vertical lines). These loss features both broaden and shift to higher B with increasing temperature. We note that at temperatures below 1 μK, where most previous measurements were performed [18, 24], these features become very narrow and can easily be missed.

To explore the temperature dependence further, we measured L_3 as a function of B around the newly discovered resonance at ≈ 0.86 G for several additional T_i values [see Figs. 2(a)–(c)]. Given the asymmetric shape of the loss features, for each T_i data series, $L_3(B)$ was fitted with a heuristic skewed Gaussian curve of the form

$$L_3(B) = A e^{-\frac{(B-B_c)^2}{2\sigma^2}} \left(1 + \mathrm{erf}\left(\frac{\alpha(B-B_c)}{\sqrt{2}\sigma}\right)\right) + C, \qquad (3)$$

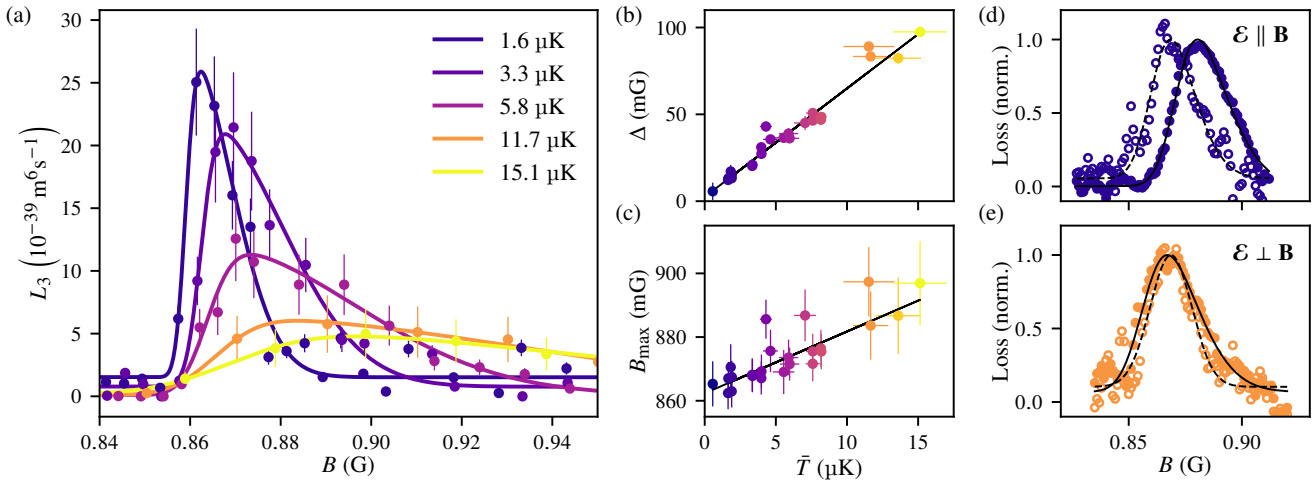

FIG. 2. Temperature dependence and light-shift of the loss feature at ≈ 0.86 G. (a) L_3 as a function of B for several decay series with different $\bar{T}$. Markers are experimental data points (with L_3 extracted in the same way as in Fig. 1) and lines are skewed Gaussian fits [see Eq. (3)]. (b) Peak width Δ as a function of $\bar{T}$. The line is a linear fit to the data and has a slope of 6.2(3) mG μK^{-1} and an intercept which is consistent with zero. (c) Peak position as a function of $\bar{T}$. The linear fit (line) has a slope of 2.0(3) mG μK^{-1} and an intercept of 861(1) mG. (d, e) Light-shift of the loss resonance for linearly polarised light with the polarisation vector $\mathcal{E}$ parallel and perpendicular to the dipole orientation, respectively. For both, we plot the normalised loss (see text) for clouds that are prepared at the same temperature (2 μK) in 1030 nm single-beam ODTs with a factor of four difference in laser intensity (filled points show the higher intensity); for $\mathcal{E} \parallel \mathbf{B}$ there is a noticeable light-shift in the resonance position, while the magnitude and width of the feature remains unchanged . Note that the data in (b) and (c) were corrected for this light-shift effect.

where B_c, σ, α, A and C are fitting parameters.

Figures 2(b) and 2(c) show, respectively, the peak width Δ (taken as twice the variance of the skewed Gaussian) and B_{max} [the location of the maximum of $L_3(B)$] as a function of the average temperature $\bar{T}$ of the decay series [28]. We observe that Δ grows linearly with temperature and so we parameterise the width of the resonance via a linear function, $\Delta = \Delta_0 + (d\Delta/dT)\,\bar{T}$, fitted to the data [solid line, Fig. 2(b)]. Note that all our extracted Δ_0 values are consistent with zero within our ±3 mG error bounds. Similarly, B_{max} also grows (approximately) linearly with temperature and so we fit the data using $B_{max} = B_0 + (dB_{max}/dT)\,\bar{T}$. The parameters of both these fits are tabulated in Table I for all the newly detected loss features. We also note that for the 0.86 G feature the maximum L_3 decreases with increasing temperature within our measured range [see Fig. 2(a)], however, for other peaks this trend is inconclusive.

The magnetic field dependence of the loss properties is due to the differential Zeeman shift arising from the difference in magnetic moments ($\delta\mu$) between the different scattering channels. However, it is also possible for light fields to exert similar differential shifts [29–32] due to the difference in polarisabilities ($\delta\alpha$) between scattering channels, which can, in some cases, also have vectorial and tensorial parts [33, 34]. To check if the optical field from our ODT causes such an effect, we measured the loss features for thermal clouds at the same temperature (2 μK) but for traps with two different light intensities (powers) and for the polarisation of the ODT light $\mathcal{E}$ either parallel ($\mathcal{E} \parallel \mathbf{B}$) or perpendicular ($\mathcal{E} \perp \mathbf{B}$) to the external magnetic field (and hence the spin-polarisation of the atoms) [35]. Data for the 0.86 G resonance is shown in Figs. 2(d) and 2(e); here, to identify the peak position we simply performed a two-point loss measurement [36]. For $\mathcal{E} \parallel \mathbf{B}$ we observe a significant (positive) shift of the loss feature with light intensity, whereas for $\mathcal{E} \perp \mathbf{B}$ the effect is much less noticeable and (if anything) has the opposite sign. Assuming that the resonance position shifts linearly with light intensity, one can extract a constant of proportionality between the light intensity I and the resonance peak shift which gives dB_0/dI for both orientations. These are tabulated in Table I for all the

B_0 (mG)	$\frac{dB_{max}}{dT}$ (mG μK^{-1})	$\frac{d\Delta}{dT}$ (mG μK^{-1})	$\left(\frac{dB_0}{dI}\right)_{\mathcal{E}\parallel\mathbf{B}}$ (G μm^2 W^{-1})	$\left(\frac{dB_0}{dI}\right)_{\mathcal{E}\perp\mathbf{B}}$ (G μm^2 W^{-1})
498(2)	3.0(4)	8.6(2.2)	1.5(3)	-0.5(2)
862(2)	1.9(3)	6.2(3)	1.7(5)	-0.2(2)
1571(4)	4.3(7)	4.6(1.5)	0.7(2)	0.1(3)
1705(3)	2.7(7)	5.3(1.0)	1.2(6)	-0.1(3)
2102(3)	5.6(7)	7.8(8)	2.6(9)	0.1(9)
2497(7)	3.8(1.6)	6.1(5)	1.9(3)	0.0(4)

Table I. Newly detected loss features and their properties. The position of the resonance B_0 for $T \to 0$ (and $I \to 0$) and the rate of change of the peaks' position (dB_{max}/dT) and width ($d\Delta/dT$) with temperature, obtained via a linear fit to the finite-temperature datasets. The shift in the resonance position with the intensity of linearly polarised 1030 nm laser light (dB_0/dI) is also tabulated, for dipoles aligned along ($\mathcal{E} \parallel \mathbf{B}$) and perpendicular to ($\mathcal{E} \perp \mathbf{B}$) the light polarisation.

newly detected loss features.

We note that, for all the loss features, the fact that the difference in the slopes $\mathrm{d}B_0/\mathrm{d}I$ between the two orthogonal light polarisations is of similar magnitude to either of the individual slopes suggests that the tensorial part of $\delta\alpha$ is of a similar magnitude to its scalar part. More quantitatively, using $\delta\alpha = 2\varepsilon_0 c\delta\mu\,\mathrm{d}B_0/\mathrm{d}I$, we find $\delta\alpha \sim 1$ a.u. using $\delta\mu \sim \mu_B$, which is the same order of magnitude as the tensorial part of the ground state polarisability at the wavelength of our ODT [33].

We now compare our findings to the predictions of a 'resonant trimer' model previously proposed in the context of temperature-dependent loss features in the lanthanides [19]. In this model the loss features are caused by resonances with 'closed channel' three-atom (trimer) bound states. We note that as the resonance is with a trimer (rather than a two-atom bound state as in more conventional Feshbach resonances), one would not expect these resonances to affect the (two-body) s-wave scattering length [37].

Within this model, some simple scalings emerge for $k_B T \gg \Gamma_{\mathrm{br}} \gg \Gamma(E)$, where Γ_{br} is the trimer decay rate (into an atom and dimer pair) and $\Gamma(E)$ is the collision energy dependent width of the trimer resonance. In this regime, one finds that $B_{\max} - B_0 = (\lambda+2)k_B T/\delta\mu$, $\Delta = 2\sqrt{3+\lambda}k_B T/\delta\mu$ and $L_3(B_{\max}) \propto T^{\lambda-1}$, where λ is related to the orbital angular momentum of the entrance channel ($\lambda = 0, 2$ for s-wave and d-wave respectively).

The first thing to note is the qualitative agreement between these predictions and our observed temperature dependencies of $B_{\max}$ and Δ. More quantitatively, the ratio of $\mathrm{d}\Delta/\mathrm{d}T$ and $\mathrm{d}B_{\max}/\mathrm{d}T$ is predicted to only depend on λ and be equal to 1.73 and 1.12 for $\lambda = 0$ and 2, respectively. For the 0.86 G resonance this ratio is 3.2(5), which is closer to the s-wave prediction [38]. An s-wave assignment would also be consistent with the fact that the maximum L_3 drops with temperature. For the other resonances, the larger error bars and the inconclusive trends of $L_3(B_{\max})$ against T make any λ assignment difficult.

Finally, we note that, despite the consistency of the 0.86 G feature with the resonant trimer model, we cannot rule out alternative models [22] which also predict regimes with linear T-dependence of both $B_{\max}$ and Δ, especially for the other less well-mapped loss features.

III. OPTIMISED BEC PRODUCTION

To produce erbium BECs, we employ standard laser cooling and trapping techniques and then use our knowledge of $L_3(B)$ to inform and optimise the evaporative cooling sequence.

In the initial steps, similarly to Ref. [3], an atomic beam emerging from a high-temperature effusion cell oven is slowed down using a Zeeman slower operating on the broad transition at 401 nm. The slow atoms are then loaded into a narrow-line magneto–optical trap (MOT) operating on the atomic transition at 583 nm. We typically capture 10^8 atoms after loading the MOT for 12 s. Afterwards, we ramp to a compressed MOT (cMOT) configuration in 600 ms, where reducing the light detuning and intensity causes simultaneous compression and cooling, resulting in a spin-polarised atomic sample at a temperature of 10 µK.

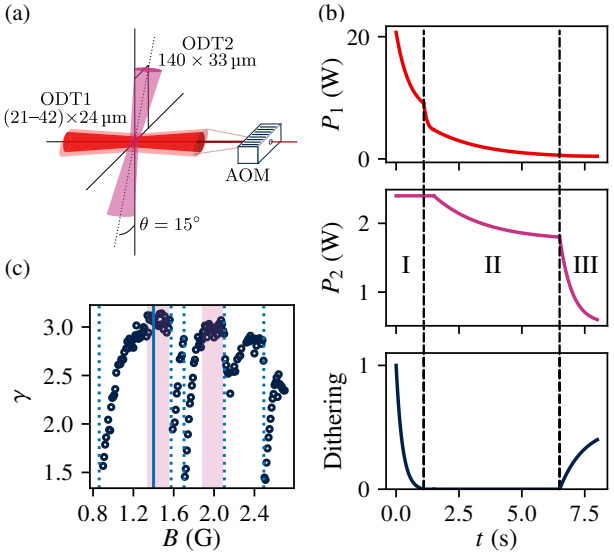

FIG. 3. Optimising evaporation. (a) Schematics of the crossed-beam optical dipole trap. The horizontal (ODT1) beam can be enlarged by dithering it with an AOM. The cross (ODT2) beam propagates at 15° to the vertical. (b) After loading the ODT from the cMOT, the evaporation sequence consists of three parts. During phase I, the power in ODT1 is lowered in parallel with ramping down the dithering. In phase II, the atoms are drawn into the crossing with the second ODT beam as the cooling continues. Finally, in phase III, with atoms residing entirely in the crossed region, we decompress the trap by ramping down ODT2 and dithering the ODT1 beam. P_1 and P_2 correspond to the power in ODT1 and ODT2, respectively, and the dithering amplitude is shown in arbitrary units. (c) Evaporation efficiency γ (see text) as a function of magnetic field. The dashed vertical lines show the $T \to 0$ positions of the loss features. The shaded regions denote optimal regions for evaporation and $B = 1.4$ G, which we use during evaporation, is indicated by a solid vertical line.

To cool the sample further, we transfer the atoms into an ODT, broadly following previous protocols, in which we perform evaporative cooling. As shown in Fig. 3(a), the ODT is implemented using two crossed, far-detuned beams at 1030 nm, which we call ODT1 and ODT2. Initially, the 21 µm × 24 µm waist ODT1 beam is superimposed onto the cMOT, with a total power of 21 W and with a 50 kHz spatial dithering applied using an acousto-optic modulator (AOM) [39], which broadens the horizontal (21 µm) waist by a factor of two. A 40 ms overlap of the dithered ODT1 beam with the cMOT results in 1.8×10^7 atoms being trapped in ODT1 at a temperature of ≈ 40 µK.

The next step is to evaporatively cool the atoms; it is here where a consideration of three-body losses becomes important. Maximising the evaporation efficiency requires minimising losses while still maintaining a sufficient rate of elastic two-body collisions which facilitate evaporation. To minimise three-body losses in the range of temperatures encountered during evaporation, one wants to be roughly in the middle of the range between the two (relatively broad) Feshbach resonances at 0 G and 3 G. Additionally, one wants to be as far as possible to the high-field side of any of the loss features, as they both move and

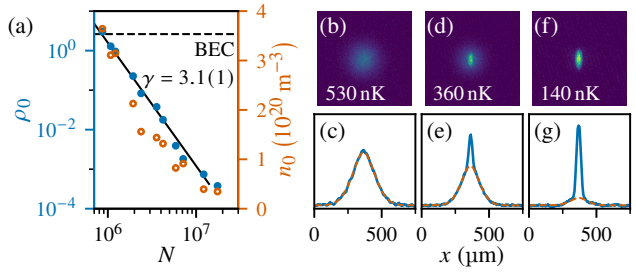

FIG. 4. Evaporation to BEC. (a) Evolution of the peak phase-space density (ρ_0, blue circles) and number density (n_0, orange circles) with the total atom number (N) during evaporation. A linear fit on the logarithmic plot shows that $\gamma = -\mathrm{d}(\ln\rho_0)/\mathrm{d}(\ln N) = 3.1(1)$. (b–f) Time-of-flight absorption images and summed density plots of atomic clouds: (b, c) just prior to condensation, corresponding to the highest ρ_0 point in (a), (d, e) a partially condensed cloud, (f, g) a nearly pure BEC with 2.2×10^5 atoms in the condensate. The dashed line in the summed density plots represents a fit of the extended Bose distribution to the thermal component of the cloud, from which the temperature can be extracted.

broaden towards higher B as T increases. This points towards choosing $B \approx 1.5$ G, on the right of the largest gap between loss features [see Fig. 1(c)]. On the other hand, the elastic collision rate, set by the s-wave scattering length a_s, increases as one approaches the 3 G resonance from below [24]. This favours higher B and points towards the regions around 2 G and 2.4 G.

To discern the optimal B for evaporation, in Fig. 3(c) we plot the efficiency (γ) of the evaporation ramp down to the point just above condensation as a function of B. Here $\gamma = -\mathrm{d}(\ln\rho_0)/\mathrm{d}(\ln N)$, where $\rho_0 = n_0\lambda_T^3$ is the peak phase-space density with $\lambda_T = \sqrt{2\pi\hbar^2/mk_BT}$ the thermal de Broglie wavelength. We see that there is indeed an optimal region around $B = 1.4$ G and so we perform our evaporation there, at which point $a_s = 73a_0$ [24]. We note that the region around 2 G is also suitable for evaporation, as although it suffers from greater three-body loss, it has a larger $a_s \approx 80a_0$. This observation is consistent with the 1.9 G field used previously for evaporation (see the supplemental material of Ref. [10]).

The evaporation sequence can be split into three stages [see Fig. 3(b)]. In stage I, in which the ODT2 contributes negligibly to the trapping, we simultaneously reduce the ODT1 power and ramp down its dithering. This leads to evaporation and a change in the trap aspect ratio, but avoids too much decompression. At the start of stage II, the ODT2 beam, with a waist of 140 μm × 33 μm and an initial power of 2.4 W, starts to have a noticeable effect and, as the cooling continues, the remaining atoms converge into the crossing of the ODT beams. In stage III, we employ the novel approach of broadening the ODT1 beam again by ramping up the dithering amplitude alongside significantly decreasing the power of ODT2. This lowers the trap depth and all trapping frequencies, and thereby reduces the atomic density and hence the rate of inelastic three-body collisions relative to the elastic two-body ones.

In Fig. 4(a) we show how the peak density n_0 and the peak phase-space density ρ_0 evolve with the falling N during the evaporation sequence. This highlights the growing density and justifies the need for our stage III decompression: at the end of stage II we reach $n_0 = 3 \times 10^{20}$ m^{-3} which gives a characteristic three-body lifetime at the centre of the cloud of only $1/L_3n_0^2 \approx 1$ s. We achieve efficient evaporation throughout the three stages, maintaining a steady increase of ρ_0 with efficiency $\gamma = 3.1(1)$; this results in the onset of condensation being reached with $N = 8 \times 10^5$ atoms and at a temperature of 500 nK. Finally, by evaporating further we achieve a nearly pure condensate with 2.2×10^5 atoms [see Figs. 4(b)–(g)].

IV. CONCLUSION

In conclusion, we have identified six new strongly temperature-dependent three-body loss features in ^{166}Er below 4 G. Both the position and width of these loss features increase linearly with temperature for $0.5 < T < 15$ μK; this is broadly consistent with a 'resonant trimer' model previously put forward to explain some loss features in lanthanide atoms [19].

Using our knowledge of the loss landscape to optimise the evaporation procedure enabled the production of large BECs of 2.2×10^5 atoms, providing a good starting point for the investigation of ultracold dipolar physics. Furthermore, these findings will enable the optimisation of atom numbers in existing and future experiments, and guide the way towards the experimental realisation of more exotic states, including honeycomb, labyrinthine and pumpkin phases. Moreover, precise knowledge of the three-body loss coefficient could enable the measurement of the atom number density, crucial for determining the structure of quantum droplets.

ACKNOWLEDGMENTS

We thank Nathaniel Vilas for contributions to the early stages of the experiment, and Raphael Lopes and Jean Dalibard for useful discussions. This work was supported by the UK EPSRC (grants no. EP/P009565/1 and EP/T019913/1). R. P. S. and P. J. acknowledge support from the Royal Society, P. J. acknowledges support from the Hungarian National Young Talents Scholarship, M. K. from Trinity College, Cambridge, J. K. from the Oxford Physics Endowment for Graduates (OXPEG) and G. L. from Wolfson College, Oxford.

[1] C. Chin, R. Grimm, P. Julienne, and E. Tiesinga, Feshbach resonances in ultracold gases, Rev. Mod. Phys. **82**, 1225 (2010).
[2] T. Weber, J. Herbig, M. Mark, H.-C. Nägerl, and R. Grimm, Three-Body Recombination at Large Scattering Lengths in an Ultracold Atomic Gas, Phys. Rev. Lett. **91**, 123201 (2003).
[3] K. Aikawa, A. Frisch, M. Mark, S. Baier, A. Rietzler, R. Grimm,

and F. Ferlaino, Bose-Einstein Condensation of Erbium, Phys. Rev. Lett. **108**, 210401 (2012).
[4] M. Lu, N. Q. Burdick, S. H. Youn, and B. L. Lev, Strongly Dipolar Bose-Einstein Condensate of Dysprosium, Phys. Rev. Lett. **107**, 190401 (2011).
[5] H. Kadau, M. Schmitt, M. Wenzel, C. Wink, T. Maier, I. Ferrier-Barbut, and T. Pfau, Observing the Rosensweig instability of a quantum ferrofluid, Nature **530**, 194 (2016).
[6] I. Ferrier-Barbut, H. Kadau, M. Schmitt, M. Wenzel, and T. Pfau, Observation of Quantum Droplets in a Strongly Dipolar Bose Gas, Phys. Rev. Lett. **116**, 215301 (2016).
[7] L. Chomaz, S. Baier, D. Petter, M. J. Mark, F. Wächtler, L. Santos, and F. Ferlaino, Quantum-Fluctuation-Driven Crossover from a Dilute Bose-Einstein Condensate to a Macrodroplet in a Dipolar Quantum Fluid, Phys. Rev. X **6**, 041039 (2016).
[8] F. Böttcher, J.-N. Schmidt, M. Wenzel, J. Hertkorn, M. Guo, T. Langen, and T. Pfau, Transient Supersolid Properties in an Array of Dipolar Quantum Droplets, Phys. Rev. X **9**, 011051 (2019).
[9] L. Tanzi, S. M. Roccuzzo, E. Lucioni, F. Famà, A. Fioretti, C. Gabbanini, G. Modugno, A. Recati, and S. Stringari, Supersolid symmetry breaking from compressional oscillations in a dipolar quantum gas, Nature **574**, 382 (2019).
[10] L. Chomaz, D. Petter, P. Ilzhöfer, G. Natale, A. Trautmann, C. Politi, G. Durastante, R. M. W. van Bijnen, A. Patscheider, M. Sohmen, M. J. Mark, and F. Ferlaino, Long-Lived and Transient Supersolid Behaviors in Dipolar Quantum Gases, Phys. Rev. X **9**, 021012 (2019).
[11] M. A. Norcia, C. Politi, L. Klaus, E. Poli, M. Sohmen, M. J. Mark, R. N. Bisset, L. Santos, and F. Ferlaino, Two-dimensional supersolidity in a dipolar quantum gas, Nature **596**, 357 (2021).
[12] T. Bland, E. Poli, C. Politi, L. Klaus, M. A. Norcia, F. Ferlaino, L. Santos, and R. N. Bisset, Two-Dimensional Supersolid Formation in Dipolar Condensates, Phys. Rev. Lett. **128**, 195302 (2022).
[13] D. Baillie and P. B. Blakie, Droplet Crystal Ground States of a Dipolar Bose Gas, Phys. Rev. Lett. **121**, 195301 (2018).
[14] Y.-C. Zhang, T. Pohl, and F. Maucher, Phases of supersolids in confined dipolar Bose-Einstein condensates, Phys. Rev. A **104**, 013310 (2021).
[15] J. Hertkorn, J.-N. Schmidt, M. Guo, F. Böttcher, K. S. H. Ng, S. D. Graham, P. Uerlings, T. Langen, M. Zwierlein, and T. Pfau, Pattern formation in quantum ferrofluids: From supersolids to superglasses, Phys. Rev. Res. **3**, 033125 (2021).
[16] J. Hertkorn, J.-N. Schmidt, M. Guo, F. Böttcher, K. S. H. Ng, S. D. Graham, P. Uerlings, H. P. Büchler, T. Langen, M. Zwierlein, and T. Pfau, Supersolidity in Two-Dimensional Trapped Dipolar Droplet Arrays, Phys. Rev. Lett. **127**, 155301 (2021).
[17] E. Poli, T. Bland, C. Politi, L. Klaus, M. A. Norcia, F. Ferlaino, R. N. Bisset, and L. Santos, Maintaining supersolidity in one and two dimensions, Phys. Rev. A **104**, 063307 (2021).
[18] A. Frisch, M. Mark, K. Aikawa, F. Ferlaino, J. L. Bohn, C. Makrides, A. Petrov, and S. Kotochigova, Quantum chaos in ultracold collisions of gas-phase erbium atoms, Nature **507**, 475 (2014).
[19] T. Maier, H. Kadau, M. Schmitt, M. Wenzel, I. Ferrier-Barbut, T. Pfau, A. Frisch, S. Baier, K. Aikawa, L. Chomaz, M. J. Mark, F. Ferlaino, C. Makrides, E. Tiesinga, A. Petrov, and S. Kotochigova, Emergence of Chaotic Scattering in Ultracold Er and Dy, Phys. Rev. X **5**, 041029 (2015).
[20] K. Baumann, N. Q. Burdick, M. Lu, and B. L. Lev, Observation of low-field Fano-Feshbach resonances in ultracold gases of dysprosium, Phys. Rev. A **89**, 020701(R) (2014).
[21] V. A. Khlebnikov, D. A. Pershin, V. V. Tsyganok, E. T. Davletov, I. S. Cojocaru, E. S. Fedorova, A. A. Buchachenko, and A. V. Akimov, Random to Chaotic Statistic Transformation in Low-Field Fano-Feshbach Resonances of Cold Thulium Atoms, Phys. Rev. Lett. **123**, 213402 (2019).
[22] Q. Beaufils, A. Crubellier, T. Zanon, B. Laburthe-Tolra, E. Maréchal, L. Vernac, and O. Gorceix, Feshbach resonance in d-wave collisions, Phys. Rev. A **79**, 032706 (2009).
[23] V. A. Khlebnikov, V. V. Tsyganok, D. A. Pershin, E. T. Davletov, E. Kuznetsova, and A. V. Akimov, Characterizing the temperature dependence of Fano-Feshbach resonances of ultracold polarized thulium, Phys. Rev. A **103**, 023306 (2021).
[24] A. Patscheider, L. Chomaz, G. Natale, D. Petter, M. J. Mark, S. Baier, B. Yang, R. R. W. Wang, J. L. Bohn, and F. Ferlaino, Determination of the scattering length of erbium atoms, Phys. Rev. A **105**, 063307 (2022), the authors of that work confirmed that the location of the highest resonance is 4.028 G and not 4.208 G, as can be seen from Fig. 3 of that work (private communication).
[25] See Supplemental Material for ODT powers, polarisations and associated trap frequencies, the average temperatures and initial atom numbers for all the data series used in Figs. 1 and 2.
[26] The magnetic field is calibrated to an accuracy of 1 mG using radio frequency (RF) spectroscopy within the ground state Zeeman manifold. However, the quoted magnetic fields are also affected by an error of ±5 mG between different datasets, due to long-term drifts between RF measurements.
[27] The 40% cutoff is chosen as a tradeoff between minimising systematic errors (with a lower cutoff) and random errors (by choosing a higher cutoff to include more data).
[28] Note that the average temperature of a decay series $\bar{T}$ is up to 20% higher than the T_i for the same series due to the heating associated with the three-body loss.
[29] D. M. Bauer, M. Lettner, C. Vo, G. Rempe, and S. Dürr, Combination of a magnetic Feshbach resonance and an optical bound-to-bound transition, Phys. Rev. A **79**, 062713 (2009).
[30] Z. Fu, P. Wang, L. Huang, Z. Meng, H. Hu, and J. Zhang, Optical control of a magnetic Feshbach resonance in an ultracold Fermi gas, Phys. Rev. A **88**, 041601(R) (2013).
[31] M. Cetina, M. Jag, R. S. Lous, J. T. M. Walraven, R. Grimm, R. S. Christensen, and G. M. Bruun, Decoherence of Impurities in a Fermi Sea of Ultracold Atoms, Phys. Rev. Lett. **115**, 135302 (2015).
[32] L. W. Clark, L.-C. Ha, C.-Y. Xu, and C. Chin, Quantum Dynamics with Spatiotemporal Control of Interactions in a Stable Bose-Einstein Condensate, Phys. Rev. Lett. **115**, 155301 (2015).
[33] J. H. Becher, S. Baier, K. Aikawa, M. Lepers, J.-F. Wyart, O. Dulieu, and F. Ferlaino, Anisotropic polarizability of erbium atoms, Phys. Rev. A **97**, 012509 (2018).
[34] T. Chalopin, V. Makhalov, C. Bouazza, A. Evrard, A. Barker, M. Lepers, J.-F. Wyart, O. Dulieu, J. Dalibard, R. Lopes, and S. Nascimbene, Anisotropic light shift and magic polarization of the intercombination line of dysprosium atoms in a far-detuned dipole trap, Phys. Rev. A **98**, 040502(R) (2018).
[35] Note that for linearly polarised light, there is no vector component of the polarisability [33] and so the angle between **B** and the direction of propagation of the light does not matter.
[36] In our two-point measurements the loss is $N(0)-N(t_{\mathrm{hold}})$, where t_{hold} is a fixed hold time, the peak loss is then normalised to 1.
[37] Consistent with this, we note that we found no significant change of the cross-thermalisation rate across the 0.86 G resonance.
[38] Note that as the assumption $\Gamma_{\mathrm{br}} \gg \Gamma(E)$ weakens, we numerically found that while the linear T-dependence of B_{max} and Δ still approximately hold, the ratio of the gradients becomes larger.
[39] C. Kohstall, *A New Toolbox for Experiments with Ultracold ^{6}Li*, Master's thesis, University of Innsbruck (2007).

6

Stability of a dipolar quantum gas in power-law potentials

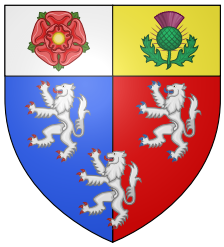

Pembroke College, Oxford. *Per pale Azure and Gules, three lions rampant two and one Argent, on a chief per pale Argent and Or, in the first a rose Gules seeded Or barbed Vert, in the second a thistle proper.*

Homogeneous quantum gases open up new possibilities for studying many-body phenomena and have now been realised for a variety of systems. For gases with short-range interactions the way to make the cloud homogeneous is, predictably, to trap it in an ideal (homogeneous) box potential. In this Chapter, we show that creating a close-to-homogeneous dipolar gas in the roton regime, when long-range interactions are important, actually requires trapping particles in soft-walled (inhomogeneous) box-like potentials, as hard walls trigger density oscillations near the wall even when the bulk of the system is not in the roton regime. We characterise how the optimum density distribution depends on the shape of the trapping potential and find it is controlled by the trap wall steepness.

6.1 Introduction

Traditionally, ultracold atomic gases are confined in harmonic trapping potentials, realised by Gaussian laser beams or magnetic traps. However, this does not reflect the continuous translational symmetry that many systems possess. This problem was overcome

by placing the atoms in a so-called optical box potential [45, 47], which had much success in systems with purely contact interactions [47]. As making a condensate homogeneous almost invariably makes the interpretation of experiments easier and the comparisons with theory more direct, this cleared the way for many interesting topics to be examined, including the dynamics of spontaneous symmetry breaking at a phase transition [49], the dynamics of strongly-interacting Bose–Einstein condensates (BECs) [263], turbulence in a quantum gas [48, 264], and first- and second-sound in a superfluid [265].

In dipolar experiments so far, the dipolar gases were confined in anisotropic, harmonic potentials; theoretically, most attention has focused on such fully harmonically trapped gases [80, 82, 209, 210, 212, 213, 253, 254, 256, 257, 266–272] and on homogeneous condensates [268, 273–278] which are harmonically confined along the polarisation direction but are unconfined in at least one of the other two (in-plane) directions. The natural way to create homogeneous conditions experimentally is to make the in-plane confinement box-like. Such traps are yet to be used for experimentally studying many-body phenomena in dipolar quantum gases, but theoretical studies involving ideal box traps have revealed non-trivial effects such as the accumulation of density near the box walls [279] and novel supersolid crystal structures [280].

An interesting consequence of the anisotropic nature of the dipolar force is the strong dependence of the stability of a trapped BEC on the trapping geometry and the interaction strength. As each atom in the cloud attracts other atoms within a conical domain around the dipole axis, dipolar interactions can render the cloud unstable. As we have seen in §2.3.3, this is the case for an unconfined, homogeneous dipolar gas, and also when the attractive component of the dipolar force overpowers the repulsive contact interactions, the confining potential and the quantum pressure. The case of flat (2D-like) BECs is particularly interesting, where strong dipolar interactions introduce a dip in the excitation spectrum at a finite momentum, the so-called ‘roton minimum’ [71, 72, 193, 194]. This leads to exotic phenomena, e.g. supersolidity [80, 205–213], a counter-intuitive state of matter that combines the dissipationless flow of a superfluid with the crystal-like periodic density modulation of a solid [214], achieved via the self-organisation of the excited dipolar quantum gas into quantum droplets [76, 77, 81, 215].

The point at which the roton minimum leads to instability has been predicted for infinite and homogeneous quasi-2D [273] and quasi-1D [268] systems. Dipolar clouds in cylindrical harmonic traps have also been extensively investigated: BECs in more oblate traps with respect to the axis of magnetic polarisation tend to be more stable [193, 281, 282], which can be understood by realising that the more oblate the trap is, the more the trapping potential squishes the atoms, i.e. the more they are forced to be in a side-by-side configuration where the DDI is repulsive (as opposed to a head-to-tail configuration

where the DDI would be attractive). This external force ensures that purely dipolar BECs can be made stable in oblate traps, where they can withstand even some attractive contact interactions [66]. The DDI also has notable consequences on the shape of the clouds. It was found that in a harmonic trap, clouds can assume a concave, 'red blood cell'-type density profile [266], and perturbations in the trapping potential can induce 'roton-like' density oscillations [267].

In this Chapter, I investigate the effect of the dipolar interaction on the stability and uniformity of a dipolar quantum gas confined in a box-like trap, and compare these results to the case of an infinite system. The aim is to explore the experimentally relevant parameter-space and find what parameters yield a system closest to the (solvable) infinite case close to collapse, in terms of uniformity of the density distribution and the strength of the largest possible dipolar interaction before collapse. While infinite systems cannot be produced in reality, creating a large trapped system with similar properties to the infinite system would enable us to check the theory against measurements.

Investigations are restricted to a dipolar gas cylindrically trapped along the radial (r) direction in a purely power-law (r^p) potential, polarised and harmonically trapped along the axis of the cylinder (z). This is an experimentally-motivated choice, as it smoothly interpolates between a harmonic potential and a perfect box trap, and is a good way to characterise a non-perfect box potential [45]. To be able to find the ground state, we numerically solve the dipolar Gross–Pitaevskii equation (Eq. (2.69)) and test the solutions for dynamical stability by solving the Bogoliubov–de Gennes (BdG) equations (Eqs. (2.78) and (2.79)). Given the cylindrical symmetry of our system, the 3D simulation (x, y, z) can be reduced to a 2D one (r, z). However, as the angular dependence is explicitly assumed to be non-existent, this prevents the detection of angular instability so the calculated ground states need to be tested for angular instabilities using the BdG approach. A similar method was used to explore stability in harmonic traps [266, 282], but also to investigate properties of supersolids [82, 254, 257, 280] and quantised vortices [270].

The simulation package consists of programs for finding the ground state by solving the GPE in a cylindrical geometry and for its subsequent testing for angular stability using the BdG equations. To solve the GPE, both the imaginary time propagation method [283–285] and the conjugate gradient method [286–289] were implemented. The program is implemented in Python [290], but takes advantage of several fast, low-level libraries. Data supporting the results presented in this chapter and the relevant publication, generated by this program, are openly available in Ref. 291.

The results of the simulation provide some interesting considerations that are relevant for the implementation of the optical box trap in our experiment. In particular, we showed that a hard wall induces large density oscillations in the cloud, which makes it both less

stable and less homogeneous. A less steep wall is better suited to achieve a homogeneous condensate, as compared to either a harmonic trap or a box trap with hard walls.

Section 6.2 explains the theoretical background of these calculations. However, to be able to actually solve the equations, various numerical techniques are required, which are presented in §6.3 (understanding these is not necessary for following the results, but are presented here for completeness). Results and an outlook to future directions are discussed in §6.4. Finally, as all continuous transformations need to be calculated on discrete grids, the discrete equivalents of various continuous transformations are given in Appendix C.

6.2 Theory

6.2.1 Dimensionless Gross–Pitaevskii equation

The starting point of our discussions is the Gross–Pitaevskii equation Eq. (2.69). A dipolar Bose–Einstein condensate with N atoms of mass m_a, magnetic dipole moment μ and s-wave scattering length a_s can be described by the macroscopic wave function $\Psi(\mathbf{r}, t)$. In the Hartree–Fock approximation Eq. (2.73), $\Psi(\mathbf{r}, t) = \sqrt{N}\psi(\mathbf{r}, t)$, where $\psi(\mathbf{r}, t)$ is the normalised single-particle wave function (i.e. $\int_V |\psi(\mathbf{r}, t)|^2 \, \mathrm{d}^3\mathbf{r} = 1$). Both the ground state and the dynamics of the condensate are described by the time-dependent Gross–Pitaevskii equation (GPE), which we can write in terms of $\psi(\mathbf{r}, t)$:

$$i\hbar \frac{\partial \psi(\mathbf{r}, t)}{\partial t} = \left(-\frac{\hbar^2}{2m_a}\nabla^2 + V_{\mathrm{trap}}(\mathbf{r}) + \frac{4\pi\hbar^2 a_s N}{m_a}|\psi(\mathbf{r}, t)|^2 + N\int_V U_{\mathrm{dd}}(\mathbf{r}' - \mathbf{r})|\psi(\mathbf{r}', t)|^2 \, \mathrm{d}^3\mathbf{r}' \right) \psi(\mathbf{r}, t)\,, \quad (6.1)$$

where the dipolar interaction potential $U_{\mathrm{dd}}(\mathbf{r}' - \mathbf{r})$ between two magnetic dipoles μ at $\mathbf{r}$ and $\mathbf{r}'$, both aligned along the z-axis, is given by Eq. (2.14). Note that here we neglect beyond-mean-field effects (quantum fluctuations), as their contribution is negligible up to the typical densities required for mean-field collapse; they only become significant at higher densities when they can arrest the collapse, leading to quantum droplets or supersolidity [77]. Also, while each atom only really interacts with $N - 1$ other atoms, given in the assumptions behind the GPE N is very large, it makes no practical difference to write N instead of $N - 1$ in the equation.

The ground state also satisfies the time-independent GPE, which in terms of ψ is

written as

$$\mu_c\psi(\mathbf{r}) = \left(-\frac{\hbar^2}{2m_a}\nabla^2 + V_{\text{trap}}(\mathbf{r}) + \frac{4\pi\hbar^2 a_s N}{m_a}|\psi(\mathbf{r})|^2 + N\int_V U_{\text{dd}}(\mathbf{r}'-\mathbf{r})|\psi(\mathbf{r}')|^2\,\mathrm{d}^3\mathbf{r}'\right)\psi(\mathbf{r})\,, \quad (6.2)$$

where μ_c is the chemical potential. As we have already noted, the ground state must have a constant (but otherwise arbitrary) phase [292]. One can choose this phase to be 0 and in the following with this choice we will choose the ground state to be real.

Equation (6.1) can be made dimensionless by introducing (arbitrary) units of measure for length and time, x_s and t_s, and can be further simplified by the appropriate choice of these units. The transformation into the dimensionless form can be done via a change of variables:

$$\mathbf{r} \to x_s\tilde{\mathbf{r}}\,, \quad t \to t_s\tilde{t}\,, \quad \psi \to x_s^{-\frac{3}{2}}\tilde{\psi}\,. \quad (6.3)$$

Applying these substitutions to Eq. (6.1) and choosing $t_s = m_a x_s^2/\hbar$ leads to

$$i\frac{\partial\tilde{\psi}(\tilde{\mathbf{r}},\tilde{t})}{\partial\tilde{t}} = \left(-\frac{1}{2}\tilde{\nabla}^2 + \frac{m_a x_s^2}{\hbar^2}V_{\text{trap}}(\mathbf{r}) + \frac{4\pi a_s N}{x_s}\left|\tilde{\psi}(\tilde{\mathbf{r}},\tilde{t})\right|^2 + N\frac{m_a\mu_0\mu^2}{4\pi\hbar^2 x_s}\int_{\tilde{V}}\frac{1-3\cos^2\theta}{|\tilde{\mathbf{r}}'-\tilde{\mathbf{r}}|^3}\left|\tilde{\psi}(\tilde{\mathbf{r}}',\tilde{t})\right|^2\mathrm{d}^3\tilde{\mathbf{r}}'\right)\tilde{\psi}(\tilde{\mathbf{r}},\tilde{t})\,. \quad (6.4)$$

We can now define the dimensionless contact and dipolar interaction strengths:

$$g = \frac{4\pi a_s N}{x_s}\,, \quad D = \frac{3a_{\text{dd}}N}{x_s}\,. \quad (6.5)$$

Finally, by absorbing the numerical coefficients into a dimensionless trapping potential $\tilde{V}_{\text{trap}}(\tilde{\mathbf{r}}) = m_a x_s^2 V_{\text{trap}}(\mathbf{r})/\hbar^2$, and by defining $\tilde{U}_{\text{dd}}(\tilde{\mathbf{r}}'-\tilde{\mathbf{r}}) = (1-3\cos^2\theta)/|\tilde{\mathbf{r}}'-\tilde{\mathbf{r}}|^3$, we reach the final, compact form of the time-dependent GPE:

$$i\frac{\partial\tilde{\psi}(\tilde{\mathbf{r}},\tilde{t})}{\partial\tilde{t}} = \left(-\frac{1}{2}\tilde{\nabla}^2 + \tilde{V}_{\text{trap}}(\tilde{\mathbf{r}}) + g\left|\tilde{\psi}(\tilde{\mathbf{r}},\tilde{t})\right|^2 + D\int_{\tilde{V}}\tilde{U}_{\text{dd}}(\tilde{\mathbf{r}}'-\tilde{\mathbf{r}})\left|\tilde{\psi}(\tilde{\mathbf{r}}',\tilde{t})\right|^2\mathrm{d}^3\tilde{\mathbf{r}}'\right)\tilde{\psi}(\tilde{\mathbf{r}},\tilde{t})\,, \quad (6.6)$$

and similarly, the time-independent GPE:

$$\tilde{\mu}_c\tilde{\psi}(\tilde{\mathbf{r}}) = \left(-\frac{1}{2}\tilde{\nabla}^2 + \tilde{V}_{\text{trap}}(\tilde{\mathbf{r}}) + g\left|\tilde{\psi}(\tilde{\mathbf{r}})\right|^2 + D\int_{\tilde{V}}\tilde{U}_{\text{dd}}(\tilde{\mathbf{r}}'-\tilde{\mathbf{r}})\left|\tilde{\psi}(\tilde{\mathbf{r}}')\right|^2\mathrm{d}^3\tilde{\mathbf{r}}'\right)\tilde{\psi}(\tilde{\mathbf{r}})\,. \quad (6.7)$$

For brevity later on, we can abbreviate the dipolar interaction potential as $\tilde{\Phi}_{\text{dd}}(\tilde{\mathbf{r}}) = \int_{\tilde{V}}\tilde{U}_{\text{dd}}(\tilde{\mathbf{r}}'-\tilde{\mathbf{r}})\left|\tilde{\psi}(\tilde{\mathbf{r}}')\right|^2\mathrm{d}^3\tilde{\mathbf{r}}'$.

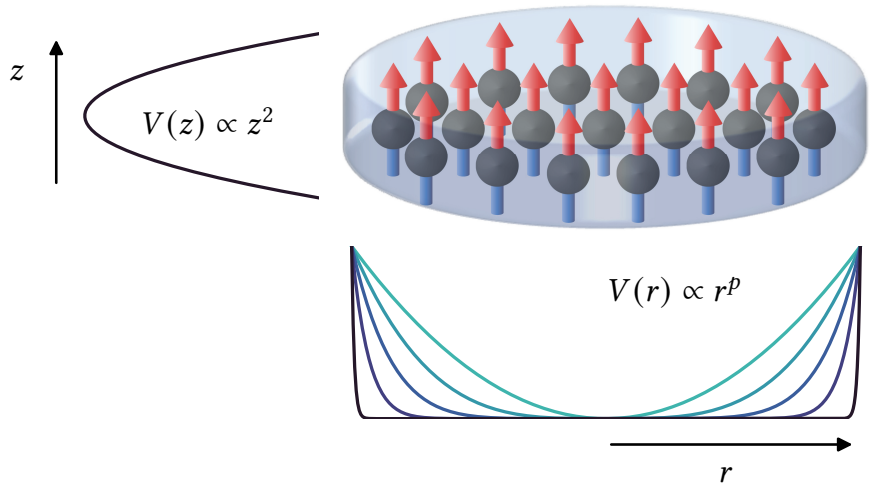

Figure 6.1. Trap geometry. We consider a gas of dipoles aligned along z which interact via dipole–dipole and contact interactions, and are confined in a 'pancake' trap by a harmonic potential along the axis (z^2) and a cylindrically symmetric power-law potential in the radial direction (r^p). Adapted from Ref. 46.

Given we are only interested in axially symmetric power-law traps along r which are harmonic along z, we can define

$$V_{\text{trap}}(\mathbf{r}) = V_{\text{trap}}(r, z) = \frac{1}{2}\kappa r^p + \frac{1}{2}m_a\omega_z^2 z^2 \,. \tag{6.8}$$

We can make the convenient choice of $x_s = \ell_z = \sqrt{\hbar/m_a\omega_z}$ (the harmonic oscillator length), which leads to $\tilde{V}_{\text{trap}}(\tilde{\mathbf{r}}) = \tilde{V}_{\text{trap}}(\tilde{r}, \tilde{z}) = \left((\tilde{r}/\gamma)^p + \tilde{z}^2\right)/2$ with $\gamma = \left(m_a\kappa/\hbar^2\right)^{-1/p} \times \ell_z^{-(1+2/p)}$. Choosing $x_s = \ell_z$ also means settling on $t_s = 1/\omega_z$, $g = 4\pi a_s N/\ell_z$ and $D = 3a_{\text{dd}}N/\ell_z$. We can then also express energies in units of $\hbar\omega_z$. Increasing γ means making the trap along r wider, i.e. making the condensate more oblate.

To find the ground state and excitations, we will also need to calculate the energy and the chemical potential. In dimensionless units these are:

$$\tilde{E}\left[\tilde{\psi}\right] = \int_{\tilde{V}} \left(-\tilde{\psi}^*(\tilde{\mathbf{r}})\frac{\tilde{\nabla}^2}{2}\tilde{\psi}(\tilde{\mathbf{r}}) + \tilde{V}_{\text{trap}}(\tilde{\mathbf{r}})\left|\tilde{\psi}(\tilde{\mathbf{r}})\right|^2 + \frac{g}{2}\left|\tilde{\psi}(\tilde{\mathbf{r}})\right|^4 + \frac{D}{2}\left|\tilde{\psi}(\tilde{\mathbf{r}})\right|^2 \int_{\tilde{V}} \tilde{U}_{\text{dd}}(\tilde{\mathbf{r}}' - \tilde{\mathbf{r}})\left|\tilde{\psi}(\tilde{\mathbf{r}}')\right|^2 \mathrm{d}^3\tilde{\mathbf{r}}'\right) \mathrm{d}^3\tilde{\mathbf{r}}\,, \tag{6.9}$$

$$\tilde{\mu}_c\left[\tilde{\psi}\right] = \int_{\tilde{V}} \left(-\tilde{\psi}^*(\tilde{\mathbf{r}})\frac{\tilde{\nabla}^2}{2}\tilde{\psi}(\tilde{\mathbf{r}}) + \tilde{V}_{\text{trap}}(\tilde{\mathbf{r}})\left|\tilde{\psi}(\tilde{\mathbf{r}})\right|^2 + g\left|\tilde{\psi}(\tilde{\mathbf{r}})\right|^4 + D\left|\tilde{\psi}(\tilde{\mathbf{r}})\right|^2 \int_{\tilde{V}} \tilde{U}_{\text{dd}}(\tilde{\mathbf{r}}' - \tilde{\mathbf{r}})\left|\tilde{\psi}(\tilde{\mathbf{r}}')\right|^2 \mathrm{d}^3\tilde{\mathbf{r}}'\right) \mathrm{d}^3\tilde{\mathbf{r}}\,. \quad (6.10)$$

The latter can be seen by writing Eq. (6.7) as $\tilde{\mu}_c\tilde{\psi}(\tilde{\mathbf{r}}) = \hat{G}\tilde{\psi}(\tilde{\mathbf{r}})$. This means $\tilde{\mu}_c\left|\tilde{\psi}(\tilde{\mathbf{r}})\right|^{+2} = \tilde{\psi}^*(\tilde{\mathbf{r}})\hat{G}\tilde{\psi}(\tilde{\mathbf{r}})$, which, due to the normalisation of ψ, leads to $\tilde{\mu}_c = \int_{\tilde{V}} \tilde{\psi}^*(\tilde{\mathbf{r}})\hat{G}\tilde{\psi}(\tilde{\mathbf{r}})\,\mathrm{d}^3\tilde{\mathbf{r}}$ as written above.

Theoretically, all we need to do to find the ground states is to solve Eq. (6.7). However, to be able to do that, various numerical techniques are needed, as described in §6.3.

6.2.2 Excitation spectrum

The excitation spectrum of the condensate can be calculated by the direct solution of the Bogoliubov–de Gennes (BdG) equations [282], Eqs. (2.78) and (2.79). In the following, we present the dimensionless version of the theory presented in §2.3.3. We can cast the macroscopic wave function of the excited state into a dimensionless form and write it in terms of the single-particle wave function,

$$\tilde{\psi}(\tilde{\mathbf{r}}, \tilde{t}) = \left(\tilde{\psi}_0(\tilde{\mathbf{r}}) + \delta\tilde{\psi}(\tilde{\mathbf{r}}, \tilde{t})\right) e^{-i\tilde{\mu}_c\tilde{t}}\,, \quad (6.11)$$

where $\tilde{\psi}_0(\tilde{\mathbf{r}})$ is the ground state (which is chosen to be a real function), $\tilde{\mu}_c$ is its chemical potential and $\delta\tilde{\psi}$ is a small perturbation as before. The perturbation is constrained to the form

$$\delta\tilde{\psi}(\tilde{\mathbf{r}}, \tilde{t}) = \sum_i \left(\tilde{u}_i(\tilde{\mathbf{r}})e^{-i\tilde{\omega}_i\tilde{t}} + \tilde{v}_i^*(\tilde{\mathbf{r}})e^{i\tilde{\omega}_i\tilde{t}}\right), \quad (6.12)$$

where $\omega_i = \tilde{\omega}_i/t_s$ are the frequencies of the excitations ($\hbar\omega_i$ is their energy), and $\tilde{u}$ and $\tilde{v}$ are normalised according to $N\int_{\tilde{V}}\left(\tilde{u}^2 - \tilde{v}^2\right)\mathrm{d}^3\tilde{\mathbf{r}} = 1$. This ansatz can be tried in the (dimensionless) time-dependent GPE Eq. (6.6), leading to the following set of equations when expanded up to first order in $\delta\tilde{\psi}$ and with terms collected according to frequency components ($e^{\pm i\omega_i t}$) (cf. Eqs. (2.85) and (2.86)):

$$\left(\hat{H}_0 - \tilde{\mu}_c + \hat{C}\right)\left(\hat{H}_0 - \tilde{\mu}_c + \hat{C} + 2\hat{X}\right)\tilde{f} = \tilde{\omega}^2\tilde{f}\,, \quad (6.13)$$

$$\left(\hat{H}_0 - \tilde{\mu}_c + \hat{C} + 2\hat{X}\right)\left(\hat{H}_0 - \tilde{\mu}_c + \hat{C}\right)\tilde{g} = \tilde{\omega}^2\tilde{g}\,, \quad (6.14)$$

where $\tilde{f} = \tilde{u} + \tilde{v}$, $\tilde{g} = \tilde{u} - \tilde{v}$ and the dimensionless operators are $\hat{H}_0 = -\tilde{\nabla}^2/2 + \tilde{V}_{\text{trap}}$ and

$$\hat{C}\tilde{f}(\tilde{\mathbf{r}}) = g\tilde{\psi}_0^2(\tilde{\mathbf{r}})\tilde{f}(\tilde{\mathbf{r}}) + D\tilde{f}(\tilde{\mathbf{r}}) \int_{\tilde{V}} \tilde{U}_{\text{dd}}(\tilde{\mathbf{r}}' - \tilde{\mathbf{r}})\tilde{\psi}_0^2(\tilde{\mathbf{r}}')\, \mathrm{d}^3\tilde{\mathbf{r}}', \tag{6.15}$$

$$\hat{X}\tilde{f}(\tilde{\mathbf{r}}) = g\tilde{\psi}_0^2(\tilde{\mathbf{r}})\tilde{f}(\tilde{\mathbf{r}}) + D\tilde{\psi}_0(\tilde{\mathbf{r}}) \int_{\tilde{V}} \tilde{U}_{\text{dd}}(\tilde{\mathbf{r}}' - \tilde{\mathbf{r}})\tilde{\psi}_0(\tilde{\mathbf{r}}')\tilde{f}(\tilde{\mathbf{r}}')\, \mathrm{d}^3\tilde{\mathbf{r}}'. \tag{6.16}$$

Note that the operators are redefined compared to §2.3.3 (they are cast in a dimensionless form) and that $\tilde{g}$ is not the dimensionless version of $g = 4\pi a_s N/x_s$ as defined in Eq. (6.5) (which is already dimensionless) but is the dimensionless version of $g = u - v$ as defined in §2.3.3.

Besides calculating excitations, this can also be used to test stability—if any excitation has an imaginary energy (i.e. $\omega^2 < 0$) associated with it, the condensate is unstable against that particular excitation mode. For stable condensates, all of ω, u and v are real [282]. Refs. 266 and 282 state that clouds which assume a concave, 'red blood cell'-shaped ground state are prone to angular instabilities. This can be understood by realising that in that case there exists a local maximum in the density of the atoms in the radial direction, which is extended along the angular direction, and therefore an angular excitation can render the gas unstable. A direct consequence of this is that energy-minimisation techniques that do not account for angular degrees of freedom (e.g. the ones we use) will overestimate the stability boundary in these cases. Therefore, all the calculated candidate ground states need to be checked using the BdG equations for angular stability. This can also be performed in a cylindrical coordinate system, given the angular (θ) dependence of any excitation can be written as $e^{-im\theta}$ ($m \in \mathbb{Z}$) in a cylindrically symmetric system. However, excitations corresponding to different phase windings (different orders m) must be checked separately.

6.2.3 Thomas–Fermi approximation

For the purposes of estimating various properties of the gas, we can make use of the Thomas–Fermi approximation, Eq. (2.72). In the limit where the trap aspect ratio (γ) is infinitely large, the gas is practically homogeneous in the radial direction, but still harmonically trapped along the axial one. In that limit, the findings of Ref. 273 apply, where they conclude that the dipolar interaction affects the shape of the condensate the same way as a contact interaction with a scattering length $2a_{\text{dd}}$ does. While the nature of these interactions are very different, due to the infinite nature of the system the direction-dependence of the dipolar interaction can be integrated out. Therefore, such a condensate has the same shape as one with an effective scattering length $a_{\text{eff}} = a_s + 2a_{\text{dd}}$ and in the

Thomas–Fermi approximation the time-independent GPE, Eq. (6.2), can be written as

$$\mu_c \psi(\mathbf{r}) = \left(V_{\text{trap}}(\mathbf{r}) + \frac{4\pi\hbar^2 a_{\text{eff}} N}{m_a} |\psi(\mathbf{r})|^2 \right) \psi(\mathbf{r}) \,. \tag{6.17}$$

We see this admits the solution

$$|\psi(\mathbf{r})|^2 = n(\mathbf{r}) = \begin{cases} \frac{\mu_c - V_{\text{trap}}(\mathbf{r})}{g_{\text{eff}}} & \text{if } \mu_c > V_{\text{trap}}(\mathbf{r}) \\ 0 & \text{otherwise} \end{cases} , \tag{6.18}$$

where we defined $g_{\text{eff}} = 4\pi\hbar^2 a_{\text{eff}} N/m_a$ and $n(\mathbf{r}) = n(r, z)$ is the 3D single-particle number density. This is an interesting result—the DDI can be treated effectively as a form of contact interaction in the infinite homogeneous system limit, where the density distribution does not change if we change a_s or a_{dd}, as long as g_{eff} is constant. To find n, we can take the ansatz

$$n(r, z) = n_0 \left(1 - \frac{r^p}{R_r^p} - \frac{z^2}{R_z^2} \right), \tag{6.19}$$

where R_r and R_z are called the Thomas–Fermi radii. Substituting this into Eq. (6.17), enforcing the condition that μ_c needs to be constant and that the density is normalised according to $\int_V n(\mathbf{r})\, \mathrm{d}^3\mathbf{r} = 1$, we find

$$R_r = \left(\frac{2g_{\text{eff}} n_0}{\kappa} \right)^{\frac{1}{p}}, \quad R_z = \sqrt{\frac{2g_{\text{eff}} n_0}{m_a \omega_z^2}}, \quad n_0 = \left(\alpha(p) \left(\frac{2g_{\text{eff}}}{\kappa} \right)^{\frac{2}{p}} \sqrt{\frac{2g_{\text{eff}}}{m_a \omega_z^2}} \right)^{\frac{-2p}{3p+4}}, \tag{6.20}$$

using the definitions of Eq. (6.8) and defining

$$\alpha(p) = \frac{\pi^{\frac{3}{2}} \Gamma\left(1 + \frac{2}{p}\right)}{\Gamma\left(\frac{5}{2} + \frac{2}{p}\right)}, \tag{6.21}$$

taking especial care with the integration limits given the piecewise definition of our function. Finally, we can now find the areal density distribution

$$n_{\text{2D,TF}}(r) = \int_{-\infty}^{\infty} N n(r, z)\, \mathrm{d}z = \frac{4}{3} N n_0 R_z \left(1 - \frac{r^p}{R_r^p} \right)^{\frac{3}{2}}. \tag{6.22}$$

In Ref. 273 the authors found that the areal density n_{2D} of a homogeneous gas can be described by $n_{\text{2D}} = \nu(\varepsilon_{\text{dd}}) \times \hbar\omega_z \ell_z / g_{\text{eff}}$, where the constant of proportionality $\nu(\varepsilon_{\text{dd}})$ is a (dimensionless) number that only depends on ε_{dd} ($\nu(\varepsilon_{\text{dd}})$ needs to be determined numerically). They found that a condensate remains stable as long as the areal density

does not exceed

$$n_{2D}^{\text{crit}} = \frac{\hbar\omega_z \ell_z}{g_{\text{eff}}} \nu_{\text{crit}}(\varepsilon_{\text{dd}}), \tag{6.23}$$

where the value of $\nu_{\text{crit}}(\varepsilon_{\text{dd}})$ is tabulated in Ref. 273. While this result is not directly applicable to our trapped system as $n_{2D}(r)$ changes in a confined system even in the Thomas–Fermi approximation, we can estimate that the stability boundary depends only on the maximum density (assuming a local collapse). If we therefore take $n_{2D} = n_{2D,TF}(0)$, then using the results of Eq. (6.20) and putting everything into dimensionless units (using the convention $x_s = \ell_z$), we can write ν in the Thomas–Fermi case as

$$\nu_{\text{TF}} = \frac{n_{2D,TF}(0) g_{\text{eff}}}{\hbar\omega_z \ell_z} = \frac{4}{3\alpha(p)} \left(\frac{\alpha(p)}{2\gamma^{\frac{6p}{4}} \left(g + \frac{8\pi D}{3}\right)} \right)^{\frac{4}{3p+4}} \left(g + \frac{8\pi D}{3}\right) \tag{6.24}$$

for a stable, confined dipolar gas (not necessarily close to collapse). This is helpful as Ref. 273 gives the stability boundary $\nu_{\text{crit}}(\varepsilon_{\text{dd}})$, so we can estimate where the stability boundary will lie in terms of D from the condition $\nu_{\text{TF}} = \nu_{\text{crit}}(\varepsilon_{\text{dd}})$, given p, γ and ε_{dd}.

We can also express the Thomas–Fermi radii in dimensionless units, which can be used to determine the size of the simulation grid (cf. §6.3.5). For a purely dipolar gas these are

$$R_r = \left(\frac{3\nu_{\text{TF}}}{2}\right)^{\frac{2}{3p}} \gamma \ell_z, \quad R_z = \left(\frac{3\nu_{\text{TF}}}{2}\right)^{\frac{1}{3}} \ell_z. \tag{6.25}$$

Finally, we will see that for high-p, high-γ boxes the density increases significantly near the trap wall. This is a consequence of the repulsive (and long-range) nature of the interaction between side-by-side dipoles and can be understood in the Thomas–Fermi approximation, in which $\tilde{V}_{\text{trap}}(\tilde{\mathbf{r}}) + g|\psi(\tilde{\mathbf{r}})|^2 + D\tilde{\Phi}_{\text{dd}}(\tilde{\mathbf{r}})$ must be constant and equal to the chemical potential within the cloud (cf. Eq. (6.7)). Note that here we do not assume the dipolar interaction can be substituted by an effective contact interaction, which is strictly only true for a condensate unconfined along r. In a sharp-walled trap, the contribution of the external potential is negligible. However, due to the long-range nature of the dipolar interactions, if the gas had a homogeneous density distribution, the dipolar term would be significantly reduced near the wall, so the density needs to increase to compensate. For less steep traps, the increasing V_{trap} compensates the decay of Φ_{dd} and so no density accumulation occurs near the edge of the trap.

This increase in density is prone to angular excitations. The wavelength of excitations is characterised by the roton wavelength $\lambda_{\text{rot}} = 2\pi/k_{\text{rot}}$, which is set by the condition $k_{\text{rot}}\ell_z \approx 1$ [193]. To estimate the order m of the lowest-energy angular excitations, one could expect that the roton wavelength will fit around the rim of the box m_{crit} times, i.e.

$2\pi R_r = m_{\text{crit}}\lambda_{\text{rot}}$, $m_{\text{crit}} \in \mathbb{Z}$. From here, we find

$$m_{\text{crit}} \approx \frac{R_r}{\ell_z} = \left(\frac{3\nu_{\text{TF}}}{2}\right)^{\frac{2}{3p}} \gamma\,. \tag{6.26}$$

This is useful as it gives us an idea of what order m to test for when calculating the angular BdG spectrum. In practice, we find testing above $1.25m_{\text{crit}}$ is not required. It is also worth noting that to be able to resolve radial and axial density oscillations, the wave function should be sampled fine enough to resolve the roton length scale.

6.2.4 Homogeneity

The purpose of creating a box trap is to obtain a condensate which behaves similarly to an infinite, perfectly homogeneous system. To quantify this, we need to establish a measure for that. Here we are aiming for a homogeneous areal density which is the same as the infinite system would have close to collapse, as we only implement box confinement in the radial direction and trapping along z is harmonic.

In a perfectly uniform cloud, each atom is surrounded by gas of the same local density. As the distribution becomes more and more inhomogeneous, the areal density takes an increasingly wider range of values across the cloud. A measure for homogeneity can be established via the probability distribution of finding an atom between a certain areal density $n_{2\text{D}}$ and $n_{2\text{D}} + \Delta n$, denoted by $Q(n_{2\text{D}}, \Delta n)$. We can find the corresponding (scaled) probability density function $P(n_{2\text{D}})$ of this distribution, where $P(n_{2\text{D}})\,\mathrm{d}n_{2\text{D}}/n_{2\text{D}}^{\text{crit}}$ gives the probability of finding a particle at a density between $n_{2\text{D}}$ and $n_{2\text{D}} + \mathrm{d}n_{2\text{D}}$:

$$\begin{aligned} P(n_{2\text{D}}) &= n_{2\text{D}}^{\text{crit}} \frac{\mathrm{d}Q(0, n_{2\text{D}})}{\mathrm{d}n_{2\text{D}}} \\ &= n_{2\text{D}}^{\text{crit}} \lim_{\Delta n_{2\text{D}} \to 0} \left(\frac{1}{\Delta n_{2\text{D}}} \frac{\int_{n_{2\text{D}} < n_{2\text{D}}(r) < n_{2\text{D}} + \Delta n_{2\text{D}}} n_{2\text{D}}(r) 2\pi r \,\mathrm{d}r}{\int_0^\infty n_{2\text{D}}(r) 2\pi r \,\mathrm{d}r} \right). \end{aligned} \tag{6.27}$$

The notation in this equation may look somewhat confusing—what we mean is that $n_{2\text{D}}$ is the value around which we evaluate the function $n_{2\text{D}}(r)$. The reason for introducing the scaling of $n_{2\text{D}}^{\text{crit}}$ is twofold: we are comparing $n_{2\text{D}}$ to $n_{2\text{D}}^{\text{crit}}$, so it is easier to look at the ratio, and to make $P(n_{2\text{D}})$ dimensionless.

A perfectly homogeneous cloud with areal density $n_{2\text{D}}^{\text{crit}}$ has a probability density distribution $P(n_{2\text{D}}) = \delta(n_{2\text{D}} - n_{2\text{D}}^{\text{crit}})$, where $\delta(x)$ is the Dirac delta function. However, a non-homogeneous-like gas has a larger spread. To quantify how similar the gas is to the infinite case, using the probability density function $P(n_{2\text{D}})$ we calculate the fraction of particles experiencing an areal density close to $n_{2\text{D}}^{\text{crit}}$, and we call this fraction H the

'homogeneity' parameter. Choosing the 'band-width' around n_{2D}^{crit} is somewhat arbitrary, and we found that calculating the fraction of atoms within 5% of n_{2D}^{crit} worked well. For a perfectly homogeneous gas $H = 1$, but for a significantly inhomogeneous gas $H < 1$, as not all the particles are at this density.

We should pay particular attention to the terminology here and emphasise that we are not looking for the most homogeneous system with any dipolar interaction strength. We are exploring how closely one can replicate an infinite homogeneous system at the critical boundary for the roton instability in a confined system. While one cannot realise infinitely large systems experimentally, recreating the same conditions in an extended but finite region should still reproduce infinite homogeneous system phenomena. Although a cloud with $P(n_{2D})$ being a delta function not centred around n_{2D}^{crit} would be perfectly homogeneous in the conventional sense of the word, it would not be like a homogeneous system at the roton instability boundary, and is therefore of no interest to us. Therefore, we only classify systems 'homogeneous' if a large proportion of their atoms are close to n_{2D}^{crit} (one could call these homogeneous-like systems).

6.3 Numerical methods

In our simulations, for each trap with a given $\{\gamma, p\}$ and for a given ε_{dd}, we solve the GPE and find the maximum value of D (g is fixed by $g = 4\pi D/3\varepsilon_{dd}$) for which a stable ground state can be found. Solving Eq. (6.7) can usually not be done analytically. Therefore, it is solved numerically, representing the wave function $\tilde{\psi}(\tilde{\mathbf{r}})$ on a discrete and finite grid (§6.3.5). Direct computation is still numerically intensive and results are sensitive to the strictness of the convergence criteria used. Therefore, we devise two strategies to calculate the ground state (to countercheck results, see §§6.3.1 and 6.3.2) and use a Fourier transform-based method to calculate the kinetic and dipolar terms (§6.3.3), taking advantage of the cylindrical and axial symmetries of the system (see §6.3.4). The computational implementation is described in §6.3.6.

6.3.1 Imaginary time propagation

There are many ways to solve the GPE, one of which is imaginary time propagation [283–285, 293, 294]. Consider a single-particle wave function $\psi(\mathbf{r}, t)$,[1] written as a

[1] Beware that $\psi(\mathbf{r}, t)$ is not the macroscopic wave function, but a general wave function which solves the Schrödinger equation.

superposition of energy eigenstates $\psi_n(\mathbf{r})$ with energy E_n,

$$\psi(\mathbf{r}, t) = \sum_n \psi_n(\mathbf{r}) e^{-iE_n t/\hbar} . \tag{6.28}$$

Performing a so-called Wick rotation, whereby we make the substitution $t \longrightarrow -i\tau$, we find

$$\psi(\mathbf{r}, -i\tau) = \sum_n \psi_n(\mathbf{r}) e^{-E_n \tau/\hbar} , \tag{6.29}$$

where τ is called 'imaginary time'. In this representation, as imaginary time progresses, eigenstates with higher energies decay faster than those with a lower energy, leaving behind the ground state which is the lowest energy state by definition. Therefore, if we make a guess of the ground state and propagate it in imaginary time, only the true ground state will eventually remain (which we will need to renormalise as it also decays).

Remember that the time-dependent GPE can be written in terms of the single-particle wave function in the Hartree–Fock approximation as

$$i\hbar \frac{\partial \psi(\mathbf{r}, t)}{\partial t} = \hat{G}\psi(\mathbf{r}, t) , \tag{6.30}$$

where $\hat{G}$ is the GPE operator (cf. Eq. (6.1)). The solution is readily given by the propagator

$$\hat{U}(t) = e^{-i\hat{G}t/\hbar} , \tag{6.31}$$

such that

$$\psi(\mathbf{r}, t) = \hat{U}(t)\psi(\mathbf{r}, 0) , \tag{6.32}$$

or, rewriting this using imaginary time,

$$\psi(\mathbf{r}, -i(\tau + \Delta\tau)) = \hat{U}(-i\Delta\tau)\psi(\mathbf{r}, -i\tau) . \tag{6.33}$$

Direct evaluation of the propagator is computationally costly due to the fact that $\hat{G}$ consists of components that do not commute (namely $\hat{T} = -\hbar^2\nabla^2/2m_a$ and $\hat{V} = V_{\text{trap}}(\mathbf{r}) + 4\pi\hbar^2 a_s N|\psi(\mathbf{r}, t)|^2/m_a + N\int_V U_{\text{dd}}(\mathbf{r}' - \mathbf{r})|\psi(\mathbf{r}', t)|^2 \, \mathrm{d}^3\mathbf{r}'$), so $e^{\hat{G}} = e^{\hat{T}+\hat{V}} \neq e^{\hat{T}}e^{\hat{V}}$. While $\hat{V}$ is diagonal in the position basis, $\hat{T}$ is diagonal in the momentum basis, so explicit evaluation of $\hat{T}$ in the position basis is difficult. Nevertheless, it can be shown that

$$\hat{U}_{\text{split}}(\Delta t) = e^{-\frac{i\hat{V}\Delta t}{2\hbar}} e^{-\frac{i\hat{T}\Delta t}{\hbar}} e^{-\frac{i\hat{V}\Delta t}{2\hbar}} \tag{6.34}$$

approximates the propagator $\hat{U}(\Delta t)$ with an error that is only $O(\Delta t^3)$. The other advant-

age of using this approximation is that the terms can be applied one by one and each can be calculated efficiently in their own basis. (This is the so-called split-step method).

Therefore, the method for finding the ground state is as follows: propagate an approximate ground state $\psi(\mathbf{r}, 0)$ in small $\Delta\tau$ imaginary time steps using $\hat{U}_{\text{split}}$ (so that the approximation holds with good precision), and renormalise ψ between each step. This needs to be continued until convergence on the energy $E[\psi]$ is reached.

Deciding the convergence criteria is important, as too loose criteria can lead to unstable states being flagged as stable ones, and too stringent criteria can lead to very long computation times. We have already shown in §2.3.3 that $\omega = 0$ is always a solution of the BdG equations if ψ is the ground state. Therefore, the lowest $m = 0$ eigenvalue of the BdG equations should be effectively 0 and the presence of such a neutral mode confirms the ground state has been reached [282]. This is a strict and precise convergence criterion to use, and this was our ultimate test to confirm the convergence of the ground state. However, calculating the BdG eigenvalues is computationally costly, so we implemented other, less stringent tests which need to be passed before the BdG test takes place. During imaginary time propagation, assuming an exponential convergence of the energy of the wave function, one can write the energy difference between the actual energy $E(\tau)$ and the converged energy E_0 as

$$\Delta E(\tau) = E(\tau) - E_0 = Ae^{-\lambda\tau}\,. \tag{6.35}$$

How quickly the energy converges indicates how far we are from the ground state. Between two imaginary time steps, the energy changes by $\delta E = (\mathrm{d}E/\mathrm{d}\tau)\,\delta\tau = -\lambda\Delta E\delta\tau$, so $\Delta E = -\delta E/\delta\tau\lambda$. We also have $\ln(-\mathrm{d}E/\mathrm{d}\tau) = \ln(\lambda A) - \lambda\tau$, and given $\mathrm{d}E/\mathrm{d}\tau \approx \delta E/\delta\tau$, we find in successive time steps $\ln(-\delta E_i/\delta t) - \ln(-\delta E_{i+1}/\delta t) = \ln(\delta E_i/\delta E_{i+1}) = \lambda\delta\tau$. Therefore, we find:

$$\Delta E = -\frac{\delta E_{i+1}}{\ln\left(\frac{\delta E_i}{\delta E_{i+1}}\right)}\,. \tag{6.36}$$

Before doing the BdG test, we require this ΔE to be less than some energy threshold E_{tol}. Furthermore, we check that $\Delta E\delta E < 0$, i.e. the algorithm is converging rather than diverging. Making larger time steps $\delta\tau$ has the benefit of converging faster, but given the split-step method yields an error in the energy of $O\left(\delta\tau^3\right)$, it comes at the expense of making larger errors. Therefore, after an iteration set with a certain $\delta\tau$ converged, we decrease our $\delta\tau$ by $\sqrt[3]{2}$ (such that the error is halved) and continue with this procedure until the energy change between successive $\delta\tau$ iteration sets is also smaller than E_{tol} (as the difference between the converged results gives an indication of how far we are from the result with $\delta\tau \to 0$). We then do the $m = 0$ BdG lowest eigenvalue test and lower E_{tol} until it is passed.

6.3.2 Preconditioned conjugate gradient method

As during code development finding stable convergence criteria for imaginary time propagation was difficult, another algorithm was developed to independently check results. Imaginary time propagation can be thought of as minimising the energy associated with ψ by changing it in small steps in a way that one goes down the gradient of the energy landscape $E[\psi]$ with each step. However, when one tries to find the deepest point of a valley, then instead of always just looking at their feet and taking a step in the steepest direction, one could hold their head high and look towards a particular direction, then go to the minimum along that direction, repeating this until one converges to the deepest point. In this case, the difficulty lies in choosing the directions wisely.

The conjugate gradient method [286] is similar to imaginary time propagation (i.e. steepest descent minimisation), only it does not necessarily follow the energy gradient at each step. Instead, additional information is drawn from the Hamiltonian $\hat{H}$ to construct a set of directions that efficiently leads towards the minimum when applied consecutively. The method was originally motivated by the problem of minimising a quadratic functional and is laid out in detail in Refs. 287–289. It can be faster than imaginary time propagation, and was used in Refs. 104, 266, 267 and 282.

Let ψ_n be the wave function before the n^{th} iteration. ψ_1 is again a guess, and the first step is taken along the direction of steepest descent. Afterwards, every iteration starts by calculating the residual of the Hamiltonian eigenvalue problem,

$$|r_n\rangle = \hat{H}|\psi_n\rangle - \left\langle \hat{H} \right\rangle \psi_n |\psi_n\rangle, \tag{6.37}$$

and the orthogonal projection ($|p_{n-1}\rangle$) of the previous direction of descent ($|d_{n-1}\rangle$) onto the space generated by $|\psi_{n-1}\rangle$,

$$|p_{n-1}\rangle = |d_{n-1}\rangle - \langle\psi_{n-1}||d_{n-1}\rangle|\psi_{n-1}\rangle. \tag{6.38}$$

After this, the new direction of descent is calculated according to

$$|d_n\rangle = -\hat{P}|r_n\rangle + \beta_n |p_{n-1}\rangle, \tag{6.39}$$

where $\hat{P}$ is called the preconditioner and β_n is the step size that mixes in $|p_{n-1}\rangle$, the 'conjugate direction'. There are various ways of choosing β_n, the most common of which

is the Polak–Ribière formula [295]:

$$\beta_n = \max\left(0, \beta_n^{\mathrm{PR}}\right), \tag{6.40}$$

$$\beta_n^{\mathrm{PR}} = \frac{(\langle r_n| - \langle r_{n-1}|)\hat{P}|r_n\rangle}{\langle \hat{P}\rangle r_{n-1}}, \tag{6.41}$$

which effectively restarts the conjugate gradient method whenever $\beta_n^{\mathrm{PR}} < 0$. Finally, the minimum of the energy along the new direction needs to be found. This can be done by writing the new wave function as

$$|\psi_{n+1}(\theta)\rangle = \cos(\theta)|\psi_n\rangle + \frac{\sin(\theta)}{\|p_n\|}|p_n\rangle, \tag{6.42}$$

which ensures the norm of the new wave function is equal to unity and that the direction of descent is orthogonal to the current state. The energy is then minimised with respect to θ, after which we can generate the wave function for the next step:

$$\theta_n = \arg\min_\theta(E[\psi_{n+1}(\theta)]), \tag{6.43}$$

$$|\psi_{n+1}\rangle = |\psi_{n+1}(\theta_n)\rangle. \tag{6.44}$$

These steps are repeated until the energy difference between the subsequent steps, $\delta E = E[\psi_{n+1}] - E[\psi_n]$, is small enough. We set this criterion to be $\delta E < 10^{-12}\hbar\omega_z$, which we found was enough to pass the $m = 0$ BdG test.

The algorithm works without using a preconditioner, i.e. choosing $\hat{P} = \hat{I}$ (the identity operator). However, using one speeds up the method considerably. One suggestion is based on the kinetic energy operator [287–289], which is diagonal in momentum space:

$$\hat{P}_T = \left(\alpha[\psi_n] - \frac{\nabla^2}{2}\right)^{-1}, \tag{6.45}$$

where α is defined as

$$\alpha[\psi_n] = \langle \hat{H}\rangle\psi_n. \tag{6.46}$$

This preconditioner becomes inefficient for large spatial grid sizes and strong interactions but was nevertheless used for finding the ground state of a dipolar gas in Refs. 266 and 282. Therefore, Refs. 287–289 define a preconditioner based on the potential and interaction terms, which are diagonal in position space:

$$\hat{P}_V = \left(\alpha[\psi_n] + V_{\mathrm{trap}} + \hat{H}_{\mathrm{int}}\right)^{-1}. \tag{6.47}$$

The problem with $\hat{P}_V$ is that its performance deteriorates for high-resolution grids. However, this can be overcome by using it in combination with $\hat{P}_T$, which leads to three possibilities,

$$\hat{P}_{VT} = \hat{P}_V \hat{P}_T \,, \tag{6.48}$$

$$\hat{P}_{TV} = \hat{P}_T \hat{P}_V \tag{6.49}$$

and the symmetric version,

$$\hat{P}_S = \hat{P}_V^{\frac{1}{2}} \hat{P}_T \hat{P}_V^{\frac{1}{2}} \,. \tag{6.50}$$

The last option should theoretically be a universally good choice, and it is the one we used in our simulations.

6.3.3 Calculation of the dipolar and the kinetic terms

Calculating the dipolar and the kinetic terms in Eq. (6.6) is computationally costly given the integration and differentiation involved. Using the convolution theorem, the dipolar term can be rewritten in a more convenient form:

$$D \int_{\tilde{V}} \tilde{U}_{\mathrm{dd}}(\tilde{\mathbf{r}}' - \tilde{\mathbf{r}}) \left|\tilde{\psi}(\tilde{\mathbf{r}}', \tilde{t})\right|^2 \mathrm{d}^3\tilde{\mathbf{r}}' = D\tilde{U}_{\mathrm{dd}} * \left|\tilde{\psi}\right|^2 = D\mathcal{F}^{-1}\left\{\mathcal{F}\left\{\tilde{U}_{\mathrm{dd}}\right\}\mathcal{F}\left\{\left|\tilde{\psi}\right|^2\right\}\right\}, \tag{6.51}$$

where $\mathcal{F}$ ($\mathcal{F}^{-1}$) denotes the (inverse) Fourier transform. Recall that the dipolar potential in its dimensionless form is

$$\tilde{U}_{\mathrm{dd}}(\tilde{\mathbf{r}}) = \frac{1 - 3\cos^2\theta}{|\tilde{\mathbf{r}}|^3} \,. \tag{6.52}$$

The Fourier transform of this can be written analytically and takes the form [281]

$$\mathcal{F}\left\{\tilde{U}_{\mathrm{dd}}\right\}(\tilde{\mathbf{k}}) = 4\pi\left(\cos^2\theta_k - \frac{1}{3}\right), \tag{6.53}$$

where $\cos\theta_k = \hat{\mathbf{z}} \cdot \tilde{\mathbf{k}}/\left|\tilde{\mathbf{k}}\right|$.

This result allows for a faster calculation of the dipolar term, given the Fourier transform of the dipolar potential only needs to be calculated once, and then calculating the dipolar interaction term is a matter of calculating a Fourier and an inverse Fourier transform, rather than evaluating an integral at every single point. However, it introduces a new issue—as our functions are represented on a finite and discrete grid, we need to use the discrete Fourier transform (DFT), which assumes the transformed function is periodic. This means there would be a contribution from alias images of the cloud outside the grid domain. To mitigate this, one can limit the range of the dipolar force by introducing a

cutoff, and make sure that a sufficiently large grid is used in the simulation [266, 279, 284]. There are multiple ways to introduce this cutoff, one of which is the spherical cutoff [282]

$$\tilde{U}_{\mathrm{dd,S}}(\tilde{\mathbf{r}}) = \begin{cases} \frac{1-3\cos^2\theta}{|\tilde{\mathbf{r}}|^3} & \text{if } |\tilde{\mathbf{r}}| < R_c \\ 0 & \text{otherwise} \end{cases}, \tag{6.54}$$

with its Fourier transform given by

$$\mathcal{F}\left\{\tilde{U}_{\mathrm{dd,S}}\right\}(\tilde{\mathbf{k}}) = 4\pi\left(\cos^2\theta_k - \frac{1}{3}\right)\left(1 + 3\frac{\cos\left(R_c\tilde{k}\right)}{R_c^2\tilde{k}^2} - 3\frac{\sin\left(R_c\tilde{k}\right)}{R_c^3\tilde{k}^3}\right), \tag{6.55}$$

where $\tilde{k} = \left|\tilde{\mathbf{k}}\right|$. This is a simple option, but is inefficient in cases when the cloud has a large aspect ratio, and so the cutoff sphere covers a lot of empty space—and consequentially the grid would be unnecessarily large along z. A possible solution for this is using a cylindrical cutoff instead,

$$\tilde{U}_{\mathrm{dd,C}}(\tilde{\mathbf{r}}) = \begin{cases} \frac{1-3\cos^2\theta}{|\tilde{\mathbf{r}}|^3} & \text{if } \tilde{r} < \tilde{r}_c \text{ and } \tilde{z} < Z_c \\ 0 & \text{otherwise} \end{cases}, \tag{6.56}$$

where $\tilde{\mathbf{r}} = (\tilde{r}, \tilde{z}, \varphi_r)$. The Fourier transform of this takes the form

$$\begin{aligned} \mathcal{F}\left\{\tilde{U}_{\mathrm{dd,C}}\right\}(\tilde{\mathbf{k}}) = {} & 4\pi\left(\cos^2\theta_k - \frac{1}{3}\right) \\ & + 4\pi e^{-Z_c\tilde{k}_r}\left(\sin^2(\theta_k)\cos\left(Z_c\tilde{k}_z\right) - \sin(\theta_k)\cos(\theta_k)\sin\left(Z_c\tilde{k}_z\right)\right) \\ & - 4\pi\int_{\tilde{r}_c}^{\infty}\int_0^{Z_c}\cos\left(\tilde{k}_z\tilde{z}\right)\frac{\tilde{r}^2 - 2\tilde{z}^2}{\left(\tilde{r}^2+\tilde{z}^2\right)^{\frac{5}{2}}}J_0\left(\tilde{k}_r\tilde{r}\right)\tilde{r}\,\mathrm{d}\tilde{z}\,\mathrm{d}\tilde{r}\,, \end{aligned} \tag{6.57}$$

where $\tilde{\mathbf{k}} = \left(\tilde{k}_r, \tilde{k}_z, \varphi_k\right)$. This is a semi-analytic expression and requires a costly numerical evaluation every time a grid is defined. Truncating the interaction in the radial direction can be avoided in our case, as we use the (discrete) Hankel transform to calculate the Fourier transform along r, which does not suffer from aliasing (see §6.3.4). Therefore, it is sufficient to implement the cutoff only along z [282], given by

$$\tilde{U}_{\mathrm{dd},z}(\tilde{\mathbf{r}}) = \begin{cases} \frac{1-3\cos^2\theta}{|\tilde{\mathbf{r}}|^3} & \text{if } \tilde{z} < Z_c \\ 0 & \text{otherwise} \end{cases}, \tag{6.58}$$

with its Fourier transform given by

$$\mathcal{F}\left\{\tilde{U}_{\mathrm{dd},z}\right\}(\tilde{\mathbf{k}}) = 4\pi\left(\cos^2\theta_k - \frac{1}{3}\right) + 4\pi e^{-Z_c\tilde{k}_r}\left(\sin^2(\theta_k)\cos\left(Z_c\tilde{k}_z\right) - \sin(\theta_k)\cos(\theta_k)\sin\left(Z_c\tilde{k}_z\right)\right). \quad (6.59)$$

The kinetic term, present in both the GPE and the BdG equations, can also be calculated via a Fourier transform. Due to the cylindrical symmetry of our system, the projection of the angular momentum along z is conserved and so all our excitations and wave functions can be written in the form $\tilde{\Psi}(\tilde{\mathbf{r}}) = \tilde{\psi}(\tilde{r},\tilde{z})e^{-im\varphi_r}$, where m is an integer.[2] The Fourier transform of such a function is $\mathcal{F}\left\{\tilde{\Psi}\right\} = \mathcal{F}\left\{\tilde{\psi}\right\}i^{-m}e^{-im\varphi_k}$. Expanding the definition of $\tilde{\nabla}^2$ and writing $\tilde{\Psi}$ as the inverse Fourier transform of $\mathcal{F}\left\{\tilde{\Psi}\right\}$, one can show

$$\tilde{\nabla}^2\tilde{\Psi} = (-1)^{m+1}\mathcal{F}^{-1}\left\{\left(\tilde{k}_r^2 + \tilde{k}_z^2\right)\mathcal{F}\left\{\tilde{\psi}\right\}\right\}. \quad (6.60)$$

This is faster than forming $\tilde{\nabla}^2$ in the position basis, as it only requires a multiplication and two Fourier transforms.

From here, the kinetic energy can also be easily calculated:

$$\int_{\tilde{V}}\tilde{\Psi}^*(\tilde{\mathbf{r}})\left(-\frac{\tilde{\nabla}^2}{2}\right)\tilde{\Psi}(\tilde{\mathbf{r}})\,\mathrm{d}^3\tilde{\mathbf{r}} = \frac{1}{(2\pi)^2}\int_{-\infty}^{\infty}\int_0^{\infty}\frac{\tilde{k}_r^2+\tilde{k}_z^2}{2}\left|\mathcal{F}\left\{\tilde{\psi}\right\}\right|^2\tilde{k}_r\,\mathrm{d}\tilde{k}_r\,\mathrm{d}\tilde{k}_z\,, \quad (6.61)$$

which only requires a single Fourier transform and an integration in k-space. Interestingly, this formula is independent of m. Furthermore, the kinetic part of our propagator (cf. §6.3.1) can also be expressed as

$$e^{-\frac{i\hat{T}\Delta t}{\hbar}}\psi = e^{\frac{\Delta\tilde{\tau}}{2}\tilde{\nabla}^2}\tilde{\psi}x_s^{-\frac{3}{2}} = \mathcal{F}^{-1}\left\{e^{-\frac{\Delta\tilde{\tau}}{2}\left(\tilde{k}_r^2+\tilde{k}_z^2\right)}\mathcal{F}\left\{\tilde{\psi}\right\}\right\}x_s^{-\frac{3}{2}}. \quad (6.62)$$

Calculating integrals and Fourier transforms on a grid is not trivial, and these discrete transforms are presented in Appendix C.

6.3.4 Cylindrical and axial symmetry

As pointed out before, due to the cylindrical symmetry of our system imposed by the shape of the trapping potential (Eq. (6.8)), the projection of the angular momentum along z is conserved. This means all our excitations and wave functions can be written in the

[2] Note the change of notation: $\tilde{\Psi}$ is now the wave function, whereas $\tilde{\psi}$ is its component after the (predetermined) angular dependence has been separated out. This is justified by the fact that it is $\tilde{\psi}(\tilde{r},\tilde{z})$ that we need to find.

form $\Psi(\mathbf{r}) = \psi(r,z)e^{-im\theta}$, $m \in \mathbb{Z}$. Ground states will have no angular momentum (as the phase needs to be constant), so $m = 0$ and $\psi(r,z,\theta) = \psi(r,z)$. In the case of excitations, they might assume any integer m.

For such functions, we can reduce a 3D simulation to a 2D one, given the dependence on θ is fixed. Furthermore, we expect ground states to be even around the $z = 0$ plane, so we only need to calculate $\psi(r,z)$ for $z \geq 0$. Excitations should also have a definite symmetry—they can be either even or odd along z, which means we can constrain our simulations to $z \geq 0$ in their case as well. These simplifications greatly speed up the simulation. It is interesting to note that in our simulations, the lowest-lying odd excitation always has an energy of exactly $\hbar\omega_z$, where the atoms slosh back and forth along z in unison, with frequency ω_z. This is called the Kohn mode [296], and can be understood by realising that if the interacting condensate is displaced without deformation, the interparticle interactions are unchanged. Therefore, the system experiences a restoring force proportional to the displacement (which is the same for all atoms), with a spring constant equal to that of the bare trap. Therefore, the condensate will oscillate with this frequency and without any distortion. This also means that it is always an even excitation that softens first.

For systems with such a symmetry, the Fourier transform of $\Psi(\mathbf{r})$,

$$\mathcal{F}\{\Psi\}(\mathbf{k}) = \int_V \Psi(\mathbf{r})e^{-i\mathbf{k}\cdot\mathbf{r}}\,\mathrm{d}^3\mathbf{r}\,, \tag{6.63}$$

and its inverse,

$$\Psi(\mathbf{r}) = \mathcal{F}^{-1}\{\mathcal{F}\{\Psi\}\}(\mathbf{r}) = \frac{1}{(2\pi)^3}\int_V \mathcal{F}\{\Psi\}(\mathbf{k})e^{i\mathbf{k}\cdot\mathbf{r}}\,\mathrm{d}^3\mathbf{k}\,, \tag{6.64}$$

can be expressed in terms of the so-called Hankel transform. The Hankel transform of order m of a function $f(r)$ is defined as

$$\mathcal{H}_m\{f\} = \int_0^\infty f(r)\,J_m(kr)\,r\,\mathrm{d}r\,, \tag{6.65}$$

with its inverse defined as

$$f(r) = \mathcal{H}_m^{-1}\{\mathcal{H}_m\{f\}\} = \int_0^\infty \mathcal{H}_m\{f\}\,J_m(kr)\,k\,\mathrm{d}k\,, \tag{6.66}$$

where J_m is the Bessel function of the first kind of order m. Using the 1st Bessel integral, $J_m(r) = \int_0^{2\pi} e^{i(r\cos\theta + m\theta)}\,\mathrm{d}\theta/2\pi i^m$, expanding the definition of the Fourier transform and doing the integral first with respect to θ, it can be shown that the Fourier transform of a

2D function of form $f(x,y) = f(r)e^{-im\theta}$ can be written as

$$\mathcal{F}\left\{f(r)e^{-im\theta}\right\} = 2\pi i^{-m} e^{-im\varphi}\, \mathcal{H}_m\{f(r)\}\,. \tag{6.67}$$

Note the special case of $m = 0$, where the Hankel transform of $f(r)$ is equivalent to the 2D Fourier transform up to a factor of 2π. This means that the full 3D Fourier transform of a function $f(x,y,z) = f(r,z)e^{-im\theta}$ can be written as

$$\mathcal{F}_{3D}\left\{f(r,z)e^{-im\theta}\right\} = 2\pi i^{-m} e^{-im\varphi}\, \mathcal{F}_z\{\mathcal{H}_m\{f(r,z)\}\}\,, \tag{6.68}$$

where $\mathcal{F}_z$ is the 1D Fourier transform along z and $\mathcal{H}_m$ is the Hankel transform along r as before.

Similarly, the inverse 2D Fourier transform of a function of the form $\tilde{f}(k_r)e^{-im\varphi}$ can be written as

$$\mathcal{F}^{-1}\left\{\tilde{f}(k_r)e^{-im\varphi}\right\} = \frac{i^m}{2\pi} e^{-im\theta}\, \mathcal{H}_m^{-1}\left\{\tilde{f}(k_r)\right\}. \tag{6.69}$$

Again, note the special case of $m = 0$, where the inverse Hankel transform of $\tilde{f}(k_r)$ is equivalent to the inverse 2D Fourier transform up to a factor of 2π. This means that the full inverse 3D Fourier transform of a function $\tilde{f}(k_x,k_y,k_z) = \tilde{f}(k_r,k_z)e^{-im\varphi}$ can be written as

$$\mathcal{F}_{3D}^{-1}\left\{\tilde{f}(k_r,k_z)e^{-im\varphi}\right\} = \frac{i^m}{2\pi} e^{-im\theta}\, \mathcal{F}_z^{-1}\left\{\mathcal{H}_m^{-1}\left\{\tilde{f}(k_r,k_z)\right\}\right\}, \tag{6.70}$$

where $\mathcal{F}_z^{-1}$ is the inverse 1D Fourier transform along z and $\mathcal{H}_m^{-1}$ is the inverse Hankel transform along r as before.

Rewriting the Fourier transform this way is useful, as when a function is represented on a discrete and finite grid, the discrete Fourier transform (DFT) assumes the function is periodic outside this domain (see Appendix C.1), but the discrete Hankel transform (DHT) does not (see Appendix C.4). As we saw before, a periodic extension is undesirable given that would yield an interaction between alias copies of our gas, so the usage of the Hankel transform is preferable.

Regarding parity along z, the Fourier series of an even function should only have cosine terms, while an odd function should only have sine terms. Therefore, to calculate the Fourier transform, we can use the discrete cosine transform (DCT) for even functions (see Appendix C.2) and the discrete sine transform (DST) for odd functions (see Appendix C.3), instead of the more computationally costly full discrete Fourier transform (DFT). More details about the discrete transforms are given in Appendix C.

6.3.5 Spatial discretisation

The wave function $\psi(r, z)$ is represented on a grid of size $N_r \times N_z$, spanning the region $[0, L_r] \times [-L_z/2, 0]$ (i.e. representing the wave function on $[0, L_r] \times [-L_z/2, L_z/2]$). Grid points are evenly distributed along z, but not along r. In the radial direction, grid points are placed at the zeros of the first order Bessel function of the first kind, to facilitate the efficient calculation of the Hankel transform (see Appendix C.4). The grid spacing along r is only significantly non-uniform for a few points closest to the origin, and the outer portion of the grid is approximately evenly distributed. Appendix C provides further details about the grid geometry and describes an accurate interpolation formula and methods for evaluating integrals with high precision. In our simulations, we used 256×65 points. The grid needs to be large enough to comfortably contain the gas, whose size can be estimated using the Thomas–Fermi approximation (Eq. (6.20)). Along r we used a grid size of $1.2R_r$ (R_r is the Thomas–Fermi radius), along z we used a grid size of $10\ell_z$, which is large enough to avoid interaction between phantom copies of the gas. We can use a constant grid size along z as R_z depends only weakly on ν (and hence $\varepsilon_{\mathrm{dd}}$). The chosen number of points ensured adequate sampling of the shortest relevant length scales (the oscillator length ℓ_z along z and the roton wavelength $\lambda_{\mathrm{rot}} \approx 2\pi\ell_z$ along r). We checked that our results were insensitive to the exact number of grid points.

6.3.6 Computing implementation

The program is implemented in the Python language [290]. As Python is known to be fairly slow, the program takes advantage of various efficient, parallel-computing methods written in C and Fortran under the hood. The Fourier transform is calculated by taking the discrete cosine/sine transform (DCT/DST) along z (based on whether the function is even/odd along z) and the discrete Hankel transform (DHT) along r. The DHT was implemented using the recipes from Refs. 297 and 298, using the NUMPY package [299] compiled against the MKL library to provide highly efficient linear algebra (BLAS) operations (e.g. matrix multiplication). The DCT and DST are implemented using the fast Fourier transform (FFT) via the MKL_FFT package, which provides a Python interface to MKL. We chose MKL_FFT as it seemed faster than both NUMPY and FFTW [300]. Some other parallel vector and matrix operations were implemented using the NUMBA package [301], and the whole execution is parallelised using the CONCURRENT.FUTURES Python module. The BdG eigenvalue problem is solved by using ARPACK [302] via the SCIPY package [259]. As the eigenvalues of the BdG equations are real for a stable condensate, to find the energy of the lowest-lying excitation it is sufficient to use the 'SR' (smallest real part) mode of

ARPACK, without a preconditioner, which avoids calculating the inverse of the BdG matrix. We also tried implementing various parts on a GPU using the PYTORCH [303] and JAX [304] packages, but they did not achieve a considerable speed gain, presumably because we use rather small arrays to represent our wave functions which all fit into the cache of our CPUs.

In terms of strategy, we search for the stability boundary by varying D for a given p, γ and $\varepsilon_{\mathrm{dd}}$. Using the Thomas–Fermi approximation, we can calculate ν_{TF}. We know the stability boundary should be reached when $\nu_{\mathrm{TF}} \approx \nu_{\mathrm{crit}}$, as given by Ref. 273, so we employ a binary search technique to find the critical D, the strongest dipolar interaction strength where a stable condensate can be found.

6.4 Results

The goal of the simulations is to design a trap in which the distribution of the atoms resembles that of an infinite homogeneous system (trapped along z but unconfined along r) as close as possible, at the roton instability boundary. More generally, we would like to get a better understanding of the effect of trap geometry on the homogeneity and stability of the ground state of our dipolar BEC. We look at the case of spin-polarised, strongly dipolar condensates (i.e. $\varepsilon_{\mathrm{dd}} \gg 1$), and we consider trapping in a power-law potential (cf. Eq. (6.8)). To illustrate general trends, three representative values of the box exponent, $p = 2$, 6 and 20 are used in the initial analysis, for a purely dipolar gas (i.e. $g = 0$) with $\gamma = 40$, representing harmonic trapping and box potentials with a moderate (soft-wall) and high steepness (hard-wall), respectively. The case of harmonic trapping is already well explored [266, 282]—indeed, we checked our results against these sources and found perfect agreement –, and our goal is to extend this analysis to higher values of p to explore a box potential-like configuration.

For each p, we investigate the stability boundary by increasing the strength of the interactions D and g (such that $\varepsilon_{\mathrm{dd}}$ is fixed) until the condensate collapses. The program uses the imaginary time propagation method to find the ground state, taking advantage of the cylindrical symmetry of the system. Results were counter-checked using the preconditioned conjugate gradient method [104]. Given the cylindrical setup, the ground states need to be checked for angular excitations using the BdG approach. To achieve numerical stability, a smaller initial timestep ($\Delta\tilde{\tau} \approx 10^{-5}$) was used for large exponents ($p \gtrsim 10$), but a larger initial timestep ($\Delta\tilde{\tau} \approx 10^{-3}$) was sufficient for small exponents ($p \lesssim 5$). Conversely, tighter convergence ($\tilde{E}_{\mathrm{tol}} \approx 10^{-6}$) was needed for small exponents, compared to $\tilde{E}_{\mathrm{tol}} \approx 10^{-3}$ for large exponents. It is worth mentioning that convergence for high aspect ratios ($\gamma \gtrsim 40$) could be achieved with larger $\tilde{E}_{\mathrm{tol}}$ as well.

The data obtained reveal interesting trends that influence our trap design. Very large aspect ratios and power laws converge to the infinite homogeneous case on the stability boundary, but in the experimentally accessible region a hard wall induces large density oscillations in the cloud which makes it both less stable and less homogeneous. A less steep wall seems to be better suited to achieve a homogeneous condensate, compared to either a harmonic trap or a box trap with hard walls.

In §6.4.1 we illustrate general trends in stability and homogeneity, and in §6.4.2 we present universal behaviour.

6.4.1 Homogeneity and stability against collapse

Given the anisotropic nature of the dipolar interaction, it is inevitably attractive for some of the atoms. Therefore, a cloud of dipolar atoms trapped in any trap will eventually become unstable as the interaction strength is increased. In each run of the simulation, we fix p, γ and $\varepsilon_{\mathrm{dd}}$, and find the ground state of a non-interacting gas ($D = g = 0$). We then continue by increasing the dipolar interaction strengths D and g (such that $\varepsilon_{\mathrm{dd}}$ is fixed) until instability is reached. Practically, as we know the stability boundary for an infinite system is given by ν_{crit} (cf. Eq. (6.23)), we do the search for the critical interaction strength D_{crit} in terms of ν_{TF}, as defined by Eq. (6.24). To set the grid size, we use Eq. (6.25).

In this subsection, we limit ourselves to exploring purely dipolar gases, but more general, universal behaviour will be presented in the next section. As we are interested in reproducing an infinite, fully homogeneous system, we present our analysis in terms of $n_{2\mathrm{D}}^{\mathrm{crit}}$, the critical density an infinite system would have (cf. Eq. (6.23)).

The critical density distributions for the different trap geometries, along with the corresponding probability density distributions (as defined in §6.2.4) are shown in Fig. 6.2. Results for $p = 2$ (Fig. 6.2(a)) match the data of existing publications on stability in harmonic traps [266, 282], but our simulations go up to larger aspect ratios.

Figures 6.2(a)–6.2(c) show examples of $n_{2\mathrm{D}}(r)/n_{2\mathrm{D}}^{\mathrm{crit}}$ for a purely dipolar gas ($g = 0$) with three different p's for $\gamma = 40$. In all cases, the gas becomes unstable when $n_{2\mathrm{D}}(r)$ reaches $n_{2\mathrm{D}}^{\mathrm{crit}}$ (or just above) somewhere in the trap, suggesting the local onset of the homogeneous roton instability (in a local density approximation picture). While for $p = 2$ and $p = 20$ the critical density is only reached at the trap centre and the trap edge respectively, for $p = 6$ it is reached across most of the gas simultaneously. To further highlight this, in Figs. 6.2(d)–6.2(f) we plot the corresponding probability density distributions $P(n_{2\mathrm{D}})$, where $P(n_{2\mathrm{D}})\,\mathrm{d}n_{2\mathrm{D}}/n_{2\mathrm{D}}^{\mathrm{crit}}$ gives the probability of finding a particle at a density between $n_{2\mathrm{D}}$ and $n_{2\mathrm{D}} + \mathrm{d}n_{2\mathrm{D}}$ (cf. §6.2.4). For a perfectly homogeneous system $P(n_{2\mathrm{D}})$ would be a delta function. For $p = 2$ we see that $P(n_{2\mathrm{D}})$ varies smoothly and only

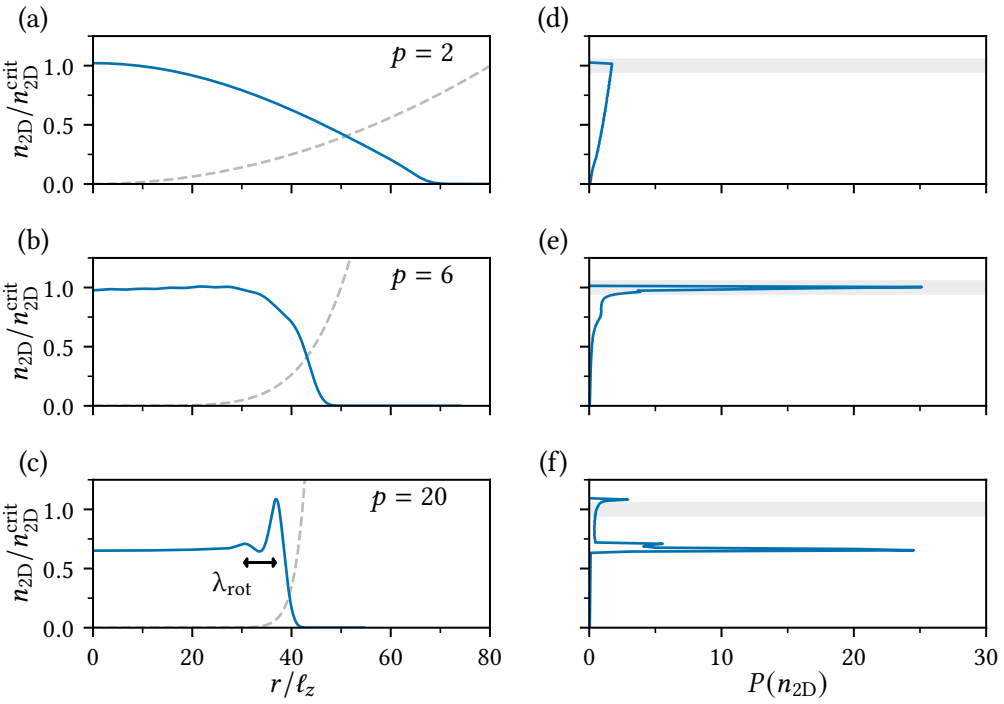

Figure 6.2. Critical density distributions of a purely dipolar gas. (a)–(c) Areal density distributions n_{2D} (solid blue lines) at the highest interaction strength at which a stable ground state can form in a power-law potential with $\gamma = 40$ and $p = 2, 6$ and 20, respectively (dashed grey lines). The n_{2D}'s are given relative to n_{2D}^{crit}, the critical density for the roton instability in an *infinite* flattened system; the trap potential is given relative to the chemical potential and shares the same axis. It can be seen that instability in the trapped system occurs when the maximum n_{2D} is close to n_{2D}^{crit}. For $p = 20$, this is due to a pronounced density oscillation near the trap wall, whose wavelength is close to λ_{rot} (see arrow). (d)–(f) Corresponding probability density distributions $P(n_{2D})$ of the areal density (see text). $P(n_{2D})$ is plotted on the horizontal axis such that the vertical axis is shared with plots (a)–(c), the grey shading denotes the region within 5% of n_{2D}^{crit}. Whereas for the $p = 6$ trap 63% of the atoms are within 5% of n_{2D}^{crit}, for both high and low p only a small fraction is. Adapted from Ref. 46.

a small fraction of the particles are near n_{2D}^{crit}. The distribution for $p = 20$ is very different, with a large peak corresponding to the bulk of the system at $n_{2D}/n_{2D}^{crit} \approx 0.6$, but still with only a small fraction near n_{2D}^{crit}. For $p = 6$, the peak corresponding to the bulk of the system sits at n_{2D}^{crit} and so the majority of the system approaches the roton instability simultaneously.

This is one of the main findings of our simulations, as this implies that to achieve the most homogeneous-like condensate (as measured by H) in an experimentally accessible region, it is better to implement a box with soft walls (moderate p) rather than one with perfectly rigid walls (large p). A gas with repulsive contact interactions, trapped in an optical box trap, reaches maximum uniformity for large values of the box exponent p, but for a dipolar gas the situation is different. The anisotropic nature of the DDI leads to the depletion of the atomic density in the middle and the accumulation of material towards the walls in the case of strong interactions at experimentally accessible aspect ratios. For larger aspect ratios, we recover the homogeneous case of Ref. 273, where the DDI has the same effect as contact interactions as far as stability is concerned, so a high p in that case is more favourable (see next section).

We can also examine the shape and spectrum of excitations. The effect of angular excitations is understandably more significant for biconcave clouds (when the maximum density is located away from the trap centre, along a rim), and less significant for convex clouds, so we will focus on the case of $p = 20$ to illustrate this effect.

The shape of the excitations can be visualised. Using Eq. (6.11), we can see that up to first order in the (small) excitation $\delta\psi$, the 3D density can be written as

$$n(\mathbf{r}, t) = |\psi(\mathbf{r}, t)|^2 = \psi_0^2(\mathbf{r}) + \psi_0(\mathbf{r})(\delta\psi(\mathbf{r}, t) + \delta\psi^*(\mathbf{r}, t)) , \tag{6.71}$$

given $\psi_0(\mathbf{r})$ can be chosen real. With the appropriate choice of $t = 0$, the components u and v of an excitation (cf. Eq. (6.12)) can be written as

$$u(\mathbf{r}) = u(r, z)e^{-im\varphi_r} , \quad v(\mathbf{r}) = v(r, z)e^{-im\varphi_r} , \tag{6.72}$$

where m is the angular order of the excitation. Therefore, we find

$$n(r, z, \varphi_r, t) = \psi_0^2(r, z) + 2\psi_0(r, z)(u(r, z) + v(r, z)) \cos(\omega t + m\varphi_r) . \tag{6.73}$$

To find the 2D-density, we can take the integral along z as usual.

The shape and energy spectrum of the angular excitations for $p = 20$, $\gamma = 40$ (i.e. the same condensate that was presented in Fig. 6.2(c)) are presented in Fig. 6.3. The areal density of the excited state (Fig. 6.3(a)) oscillates between the two extrema according to

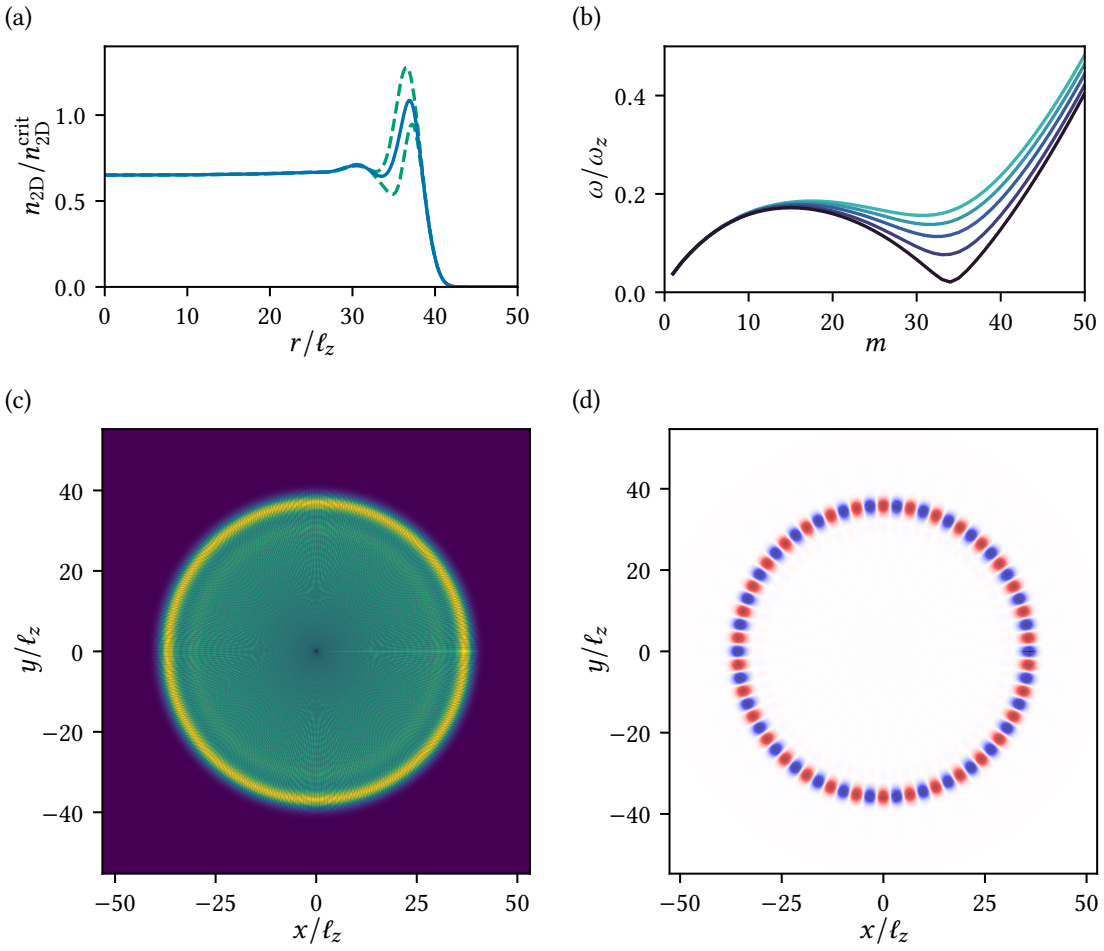

Figure 6.3. Shape and spectrum of excitations. (a) Areal density of the ground state (solid line) and the lowest energy excited state (dashed line) of the same condensate as presented in Fig. 6.2(c) ($p = 20$, $\gamma = 40$). The condensate oscillates between the two extrema according to $\cos(\omega t + m\varphi_r)$. The order of the lowest excitation is $m = 34$, so there are 68 (!) nodes around the rim. The excitation amplitude should be small compared to the ground state density, but here it is chosen relatively large for better visibility. (b) Excitation energy spectrum of this system as a function of the order of the excitation m, for a variety of interaction strengths (darker colour indicates stronger excitation). A roton–maxon-like spectrum can be seen with a pronounced dip at $m = 34$, where the excitation energy is almost 0 (i.e. the condensate is almost unstable). As the order m (and hence the momentum) of the excitation is increased further, the energy increases again. Furthermore, as the interactions get weaker, the dip in the excitation spectrum becomes smaller, signifying a more stable condensate. (c) 2D visualisation of the ground state areal density (the solid line in (a)). (d) 2D visualisation of the lowest-energy excitation (the dashed line in (a)), displaying the 68 nodes around the rim. We see the density of the condensate is only changed around the rim, and the excitation periodically increases (red) and decreases (blue) it.

$\cos(\omega t + m\varphi_r)$, where the order of the lowest energy mode is $m = 34$ (!). It can be seen that the excitation is concentrated on the density oscillation around the rim, i.e. it tries to make it so dense that the condensate collapses. To make this more easily visible, we show the distribution of the areal density of the ground state of the condensate and the excitation in 2D in Figs. 6.3(c) and 6.3(d).

The excitation energy spectrum (Fig. 6.3(b)) also reveals some interesting trends. When the excitation energy is plotted as a function of the order of the excitation m, a roton–maxon-like excitation spectrum can be seen with a pronounced dip at $m = 34$, where the excitation energy is almost 0 (i.e. the condensate is almost unstable). As the order m (and hence the momentum) of the excitation is increased further, the energy increases again. Furthermore, if ν_{TF} is decreased, the roton dip becomes less pronounced and so the condensate becomes more stable as expected.

6.4.2 Universal behaviour

To better quantify the power-law best suited for studying the physics of an infinite homogeneous system in the roton regime, we define a 'homogeneity' parameter H as the fraction of particles that experience an $n_{2\mathrm{D}}$ within 5% of $n_{2\mathrm{D}}^{\mathrm{crit}}$. We note this parameter quantifies how close the system is to a perfectly homogeneous system *at the roton instability*, and not (only) how uniform the density is across the sample. In Fig. 6.4(a), we plot H against the exponent p for aspect ratios $\gamma = 30, 50$ and 70 for a purely dipolar gas ($g = 0$, solid lines). For all three aspect ratios, H gradually increases with p up to some optimum p^* before dropping sharply with higher p as the peak in $P(n_{2\mathrm{D}})$ moves below $0.95\, n_{2\mathrm{D}}^{\mathrm{crit}}$. We see that p^* increases with γ; it is determined by the p at which significant density starts accumulating near the edge of the trap (cf. Fig. 6.2(c)). We have also checked that this behaviour is not specific to purely dipolar gases but also applies in the presence of (weak) contact interactions. We show the curves for $\varepsilon_{\mathrm{dd}} = 3$ (dashed lines), and see that repulsive contact interactions increase both p^* and the maximum H slightly (attractive contact interactions have the opposite effect).

One would expect density oscillations near the wall to somehow be controlled by the trap wall steepness, which not only depends on p but also on γ. We define the wall steepness as the gradient of the trap potential (relative to the chemical potential $\tilde{\mu}_c$) at half the chemical potential,

$$S = \left.\frac{\mathrm{d}\left(\tilde{V}_{\mathrm{trap}}/\tilde{\mu}_c\right)}{\mathrm{d}\tilde{r}}\right|_{\tilde{V}_{\mathrm{trap}}(\tilde{r},0)=\tilde{\mu}_c/2} = \frac{p}{2\gamma}\tilde{\mu}_c^{-1/p}\,. \tag{6.74}$$

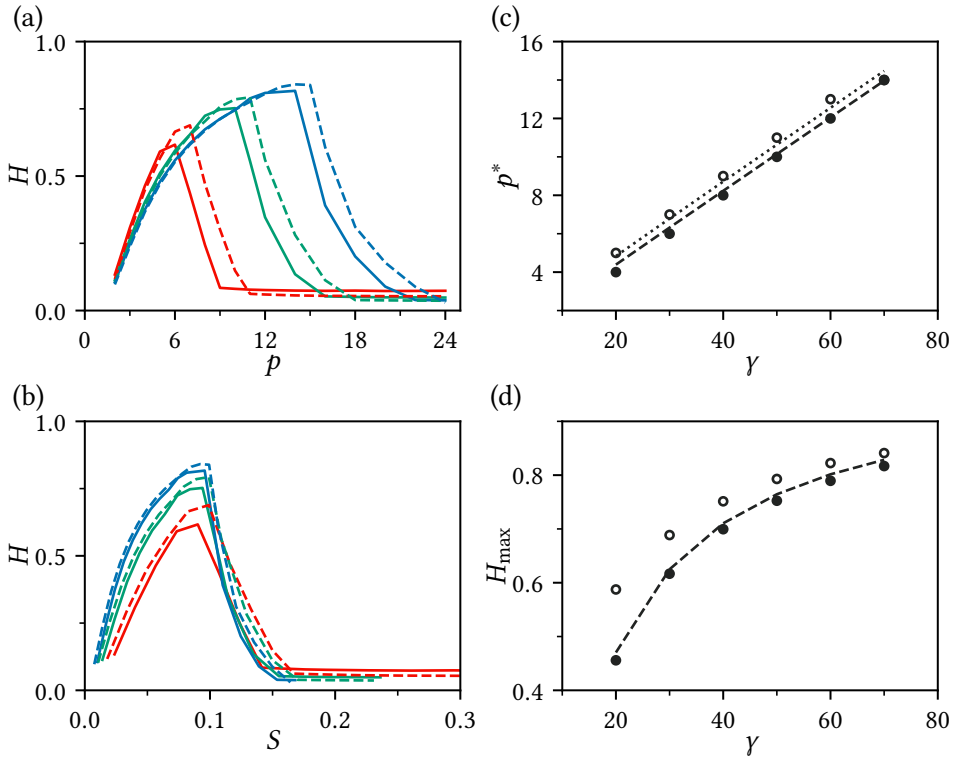

Figure 6.4. Optimum power-law. (a) The 'homogeneity' H (see text) as a function of the power-law exponent p for trapping potential aspect ratio $\gamma = 30$ (red, lower lines), 50 (green, middle lines) and 70 (blue, upper lines). The curves are plotted for a purely dipolar gas (solid lines) and for a gas with $\varepsilon_{dd} = 3$ (dashed lines). The optimum $p(= p^*)$ increases with γ. (b) The same data plotted against the trap wall steepness S (see text). The optimum H occurs at the same $S \approx 0.1(= S^*)$ for all γ. (c) The optimum power-law p^* (dots) and the prediction p^*_{pred} using $S^* = 0.1$ (lines) as a function of γ for a purely dipolar gas (filled dots, dashed line) and for one with $\varepsilon_{dd} = 3$ (empty dots, dotted line). (d) The maximum $H(= H_{\text{max}})$, achievable for a given γ (filled dots for a purely dipolar gas, empty dots for $\varepsilon_{dd} = 3$). The dashed line provides a simple interpretation of how H_{max} depends on γ (see text). Adapted from Ref. 46.

In Fig. 6.4(b), we plot H for the same aspect ratios as in Fig. 6.4(a) but now against S. Plotting this way reveals that the maximum H occurs at the same $S = S^* \approx 0.1$ for all three aspect ratios for a dipolar condensate with or without contact interactions. Given $\{p, \gamma, \tilde{\mu}_c\}$ uniquely defines S, we can invert Eq. (6.74) and use S^* to predict

$$p^*_{\text{pred}} = \frac{\ln \tilde{\mu}_c}{W_0\left(\frac{\ln \tilde{\mu}_c}{2S^*\gamma}\right)}, \tag{6.75}$$

where $W_0(x)$ is the Lambert W function. To check this, in Fig. 6.4(c) we plot p^* and p^*_{pred} for a range of γ and see there is very good agreement. This behaviour is in contrast to a gas with only contact interactions ($\varepsilon_{\text{dd}} = 0$), where homogeneity would monotonically increase with S, but would saturate when $\delta = \ell_z/S$ (the trap 'wall thickness') reaches the healing length $\xi = \ell_z/\sqrt{2\tilde{\mu}_c}$ [47] (for our parameters $\xi \approx 0.5\ell_z \ll \lambda_{\text{rot}}$). In our case, we reach the optimum H at $\delta \approx 10\ell_z \gg \xi$, which is close to the roton wavelength $\lambda_{\text{rot}} \approx 2\pi\ell_z$.

As shown in Fig. 6.4(d), as γ increases, the maximum H (H_{max}, achieved at the also growing p^*) increases towards 1, suggesting that the homogeneous limit can still in principle be approached if γ and p are increased together in a suitable way. The trend can be understood via a simple model (dashed line). If we assume the cloud consists of a homogeneous centre with radius $\gamma\ell_z - \lambda_{\text{rot}}$ and an inhomogeneous boundary with width λ_{rot}, we can estimate

$$H_{\text{max}} \approx \frac{(\gamma\ell_z - \lambda_{\text{rot}})^2\pi}{(\gamma\ell_z)^2\pi} = \left(1 - \frac{\lambda_{\text{rot}}}{\gamma\ell_z}\right)^2. \tag{6.76}$$

6.4.3 Experimental considerations

Finally, we consider the implications of our results for experimentally realising a close-to-homogeneous dipolar gas in the roton regime. Unlike for gases with solely repulsive contact interactions, the need for relatively soft walls means that the optics for creating an appropriate trap is unlikely to be a significant constraint. Instead, the limiting factor is likely to be the number of atoms required to fill a high-γ trap. For an approximately uniform gas in the roton regime $n_{\text{2D}} \approx 1/\pi\gamma^2\ell_z^2 \approx n^{\text{crit}}_{\text{2D}}$, which using Eq. (6.23) gives $\gamma^2 = 4Na_{\text{eff}}/v_{\text{crit}}(\varepsilon_{\text{dd}})\ell_z$. This shows that filling a large-γ trap requires ℓ_z to be small, but ℓ_z needs to be kept large enough to avoid high (3D) number densities which result in excessive three-body losses. The (dimensionful) peak density can be approximated via the Thomas–Fermi approximation. From Eq. (6.18) we see that the peak density will occur at $V_{\text{trap}} = 0$ at which point $n^{\text{max}}_{\text{3D}} \approx N\mu_c/g_{\text{eff}} = \tilde{\mu}_h(\varepsilon_{\text{dd}})\hbar\omega_z N/g_{\text{eff}} = \tilde{\mu}_h(\varepsilon_{\text{dd}})/4\pi a_{\text{eff}}\ell_z^2$, where $\tilde{\mu}_h(\varepsilon_{\text{dd}})$ is the (dimensionless) chemical potential tabulated in Ref. 273. Solving for ℓ_z and

inserting it into our expression for γ^2 gives

$$\gamma_{\max}^2 \approx \frac{8\sqrt{\pi}}{\nu_{\text{crit}}\sqrt{\tilde{\mu}_h}} N \sqrt{n_{3\text{D}}^{\max} a_{\text{eff}}^3}\,. \tag{6.77}$$

Therefore, with 10^5 erbium or dysprosium atoms (for which $a_{\text{dd}} \approx 100a_0$, and setting $a_s \approx 0$), if we limit $n_{3\text{D}} \lesssim 100\,\mu\text{m}^{-3}$, one could reach $\gamma_{\max} \approx 40$ with $\ell_z \approx 0.4\,\mu\text{m}$ (equivalent to a vertical trapping frequency of approximately 400 Hz), resulting in $H \approx 70\%$ for $p^* = 8$ (cf. $H \approx 10\%$ in a harmonic trap).

6.5 Publication: How to realize a homogeneous dipolar Bose gas in the roton regime

Physical Review A **105**, L061301 (2022)
Submitted: 5 January 2022; Published: 2 June 2022
DOI: 10.1103/PhysRevA.105.L061301

Péter Juhász,* Milan Krstajić,* David Strachan, Edward Gandar and Robert P. Smith

Clarendon Laboratory, University of Oxford, Parks Road, Oxford OX1 3PU, United Kingdom

* The Authors contributed equally to this work.

Letter

How to realize a homogeneous dipolar Bose gas in the roton regime

Péter Juhász,* Milan Krstajić,* David Strachan, Edward Gandar, and Robert P. Smith†
Clarendon Laboratory, University of Oxford, Parks Road, Oxford OX1 3PU, United Kingdom

(Received 5 January 2022; revised 15 March 2022; accepted 9 May 2022; published 2 June 2022)

Homogeneous quantum gases open up new possibilities for studying many-body phenomena and have now been realized for a variety of systems. For gases with short-range interactions the way to make the cloud homogeneous is, predictably, to trap it in an ideal (homogeneous) box potential. We show that creating a close to homogeneous dipolar gas in the roton regime, when long-range interactions are important, actually requires trapping particles in soft-walled (inhomogeneous) box-like potentials. In particular, we numerically explore a dipolar gas confined in a pancake trap which is harmonic along the polarization axis and a cylindrically symmetric power-law potential r^p radially. We find that intermediate p's maximize the proportion of the sample that can be brought close to the critical density required to reach the roton regime, whereas higher p's trigger density oscillations near the wall even when the bulk of the system is not in the roton regime. We characterize how the optimum density distribution depends on the shape of the trapping potential and find it is controlled by the trap wall steepness.

DOI: 10.1103/PhysRevA.105.L061301

The behavior of many-body quantum systems is governed by the interplay of the potential confining the particles and the interactions between them; ultracold gases allow for the fine control of both of these aspects. While in most ultracold-atom experiments interparticle interactions are short-ranged and isotropic, the realization of ultracold dipolar gases, using highly magnetic atoms [1–4], molecules [5], and Rydberg atoms [6], has introduced anisotropic, long-range dipole–dipole interactions, opening up many new avenues for research. In the case of degenerate Bose gases, the presence of dipole–dipole interactions has, for example, led to the study of roton physics [7,8] and the related discovery of a supersolid phase [9–11].

The term 'roton' was first coined in the context of liquid helium [12], where it describes excitations observed around a minimum in the excitation spectrum at nonzero momentum. Ultracold dipolar gases tightly confined along the polarization direction of the dipoles and held more loosely in (at least one of) the other two directions display a similar roton dispersion relation. In this case, the origin of the roton feature is the interplay of the anisotropic, long-range interactions and the tight confinement. As the strength of the interactions is increased, the roton minimum forms, deepens, and then reaches zero energy, causing the roton instability. In certain cases, this leads to the formation of quantum droplets [13–15] and, very close to the instability, a supersolid phase [9–11,16].

In the experiments so far, the dipolar gases were confined in anisotropic, harmonic potentials; theoretically, most attention has focused on such fully harmonically trapped gases [17–33] and on homogeneous condensates [34–40] which are harmonically confined along the polarization direction but are unconfined in at least one of the other two (in-plane) directions. The natural way to create homogeneous conditions experimentally is to make the in-plane confinement box-like. Box traps had much success in systems with purely contact interactions [41], as making a condensate homogeneous almost invariably makes the interpretation of experiments easier and the comparisons with theory more direct. Such traps are yet to be used for experimentally studying many-body phenomena in dipolar quantum gases, but theoretical studies involving ideal box traps have revealed nontrivial effects such as the accumulation of density near the box walls [42] and novel supersolid crystal structures [43].

In this Letter, we numerically explore the homogeneity of a dipolar gas, tuned close to the roton instability, in a flattened, cylindrically symmetric ('pancake') potential, with tight harmonic confinement along z (the direction of polarization of the dipoles) and a power-law potential r^p in the perpendicular plane (see Fig. 1). This choice is motivated by

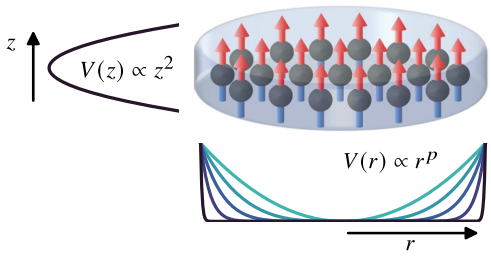

FIG. 1. Trap geometry. We consider a gas of dipoles aligned along z which interact via dipole–dipole and contact interactions, and are confined in a 'pancake' trap by a harmonic potential along the axis (z^2) and a cylindrically symmetric power-law potential in the radial direction (r^p).

*The authors contributed equally to this work.
†robert.smith@physics.ox.ac.uk

the fact that a general power-law potential smoothly interpolates between a harmonic potential ($p = 2$) and an ideal box potential ($p \to \infty$), and that experimentally relevant box-like traps are typically characterized as power-law potentials [44]. Additionally, Laguerre–Gaussian beams, often used to create optical box traps, can be used to controllably realize power-law potentials. We should emphasize that we are not looking for the most homogeneous system with any dipolar interaction strength, but we are exploring how closely one can replicate an infinite homogeneous system at the critical boundary for the roton instability. While one cannot realize infinitely large systems experimentally, recreating the same conditions in an extended but finite region should still reproduce infinite homogeneous system phenomena. We find that achieving the most homogeneous conditions within the roton regime requires an intermediate p which depends on the trap aspect ratio—this is in sharp contrast to systems with only contact interactions, where a higher p always leads to more homogeneous condensates [41]. We show that the optimum p for a given aspect ratio is determined by the box walls being soft enough so as not to trigger a roton-like instability at the edge significantly before it occurs in the bulk. We also examine how the optimum p depends on the aspect ratio and how homogeneous-system-like a sample could be produced within realistic experimental limitations.

We consider a bosonic gas of N atoms, each with mass m and magnetic dipole moment μ_m, confined in a pancake trap (see Fig. 1), and we work in dimensionless units where times are expressed in units of the inverse z-axis oscillator frequency $1/\omega_z$, energies in units of $\hbar\omega_z$, and lengths in units of the harmonic oscillator length $\ell_z = \sqrt{\hbar/(m\omega_z)}$. The trapping potential we are considering is given by

$$V_{\text{trap}}(r, z) = \frac{1}{2}\left(\left(\frac{r}{\gamma}\right)^p + z^2\right), \tag{1}$$

where γ characterizes the aspect ratio of the trap.

At zero temperature, our dipolar gas is expected to form a Bose–Einstein condensate (BEC) described by the macroscopic wave function $\Psi(\mathbf{r}, t) = \sqrt{N}\psi(\mathbf{r}, t)$, where $\psi(\mathbf{r}, t)$ is the normalized single-particle wave function that obeys the (dimensionless) Gross–Pitaevskii equation (GPE):

$$i\frac{\partial\psi(\mathbf{r}, t)}{\partial t} = \left(-\frac{1}{2}\nabla^2 + V_{\text{trap}}(\mathbf{r}) + g_s|\psi(\mathbf{r}, t)|^2 + D\Phi_{\text{dd}}(\mathbf{r}, t)\right)\psi(\mathbf{r}, t). \tag{2}$$

Here, the strength of contact interactions is characterized by the parameter $g_s = 4\pi a_s N/\ell_z$, where a_s is the s-wave scattering length and the strength of dipolar interactions is characterized by the parameter $D = 3a_{\text{dd}}N/\ell_z$, where $a_{\text{dd}} = m\mu_0\mu_m^2/(12\pi\hbar^2)$ is the dipolar length and μ_0 is the permeability of free space. The form of the mean-field dipolar interaction potential is given by

$$\Phi_{\text{dd}}(\mathbf{r}, t) = \int \frac{1 - 3\cos^2\theta}{|\mathbf{r} - \mathbf{r}'|^3}|\psi(\mathbf{r}', t)|^2\, d^3\mathbf{r}', \tag{3}$$

where θ is the angle between $\mathbf{z}$ and $\mathbf{r} - \mathbf{r}'$. The relative strength of these interactions (compared to the contact interactions) is given by the ratio $\varepsilon_{\text{dd}} = a_{\text{dd}}/a_s$. Note that here we neglect quantum fluctuations, as their contribution is negligible up to the typical densities required for mean-field collapse; they only become significant at higher densities when they can arrest the collapse, leading to quantum droplets or supersolidity [14].

In the absence of an in-plane potential (i.e. $p \to \infty$ and $\gamma \to \infty$), a dipolar gas is predicted to develop a roton-like excitation spectrum, with a roton minimum for excitations of wavelength $\lambda_{\text{rot}} \approx 2\pi$ [7], which deepens with increasing dipolar interaction strength and reaches zero energy at the roton instability. The instability occurs when the single-particle areal density $n_{2\text{D}}(r) = \int_{-\infty}^{\infty}|\psi(\mathbf{r})|^2\, dz$ reaches a critical value [7,34] given by

$$n_{2\text{D}}^{\text{crit}}(\varepsilon_{\text{dd}}) = \frac{3\nu_{\text{crit}}(\varepsilon_{\text{dd}})}{4\pi D\left(\varepsilon_{\text{dd}}^{-1} + 2\right)}, \tag{4}$$

where the value of the dimensionless prefactor $\nu_{\text{crit}}(\varepsilon_{\text{dd}})$ is tabulated in Ref. [34]. Interestingly, up until the gas becomes unstable, the density distribution of the BEC has the same form as a gas with only contact interactions with an effective scattering length $a_{\text{eff}} = a_s + 2a_{\text{dd}} = a_{\text{dd}}(\varepsilon_{\text{dd}}^{-1} + 2)$ [34], or equivalently an effective interaction parameter $g_{\text{eff}} = 4\pi a_{\text{eff}}N/\ell_z$.

In our simulations, for each trap with given $\{\gamma, p\}$ and for a given ε_{dd}, we solve the GPE and find the maximum value of D (g_s is fixed by $g_s = 4\pi D/(3\varepsilon_{\text{dd}})$) for which a stable ground state can be found (see the Appendix for further information about our algorithm). As we aim to compare the resulting critical density distributions to the infinite (perfectly homogeneous) flattened system, we evaluate the r-dependent areal density $n_{2\text{D}}(r)$ and compare it to $n_{2\text{D}}^{\text{crit}}$, the density a perfectly homogeneous system would have at the roton instability.

Figures 2(a)–2(c) show examples of $n_{2\text{D}}(r)/n_{2\text{D}}^{\text{crit}}$ for a purely dipolar gas ($\varepsilon_{\text{dd}} \to \infty$) with three different p's for $\gamma = 40$. In all cases, the gas becomes unstable when $n_{2\text{D}}(r)$ reaches $n_{2\text{D}}^{\text{crit}}$ (or just above) somewhere in the trap, suggesting the local onset of the homogeneous roton instability (in a local density approximation picture). While for $p = 2$ and $p = 20$ the critical density is only reached at the trap center and the trap edge respectively, for $p = 6$ it is reached across most of the gas simultaneously. To further highlight this, in Figs. 2(d)–2(f) we plot the corresponding probability density distributions $P(n_{2\text{D}})$, where $P(n_{2\text{D}})\, dn_{2\text{D}}/n_{2\text{D}}^{\text{crit}}$ gives the probability of finding a particle at a density between $n_{2\text{D}}$ and $n_{2\text{D}} + dn_{2\text{D}}$. For a perfectly homogeneous system $P(n_{2\text{D}})$ would be a delta function. For $p = 2$ we see that $P(n_{2\text{D}})$ varies smoothly and only a small fraction of the particles are near $n_{2\text{D}}^{\text{crit}}$. The distribution for $p = 20$ is very different, with a large peak corresponding to the bulk of the system at $n_{2\text{D}}/n_{2\text{D}}^{\text{crit}} \approx 0.6$, but still with only a small fraction near $n_{2\text{D}}^{\text{crit}}$. For $p = 6$, the peak corresponding to the bulk of the system sits at $n_{2\text{D}}^{\text{crit}}$ and so the majority of the system approaches the roton instability simultaneously.

We note that the increase of density seen near the trap walls in high-p traps (cf. Fig. 2(c)) is a consequence of the repulsive (and long-range) nature of the interaction between side-by-side dipoles and can be understood in the Thomas–Fermi approximation, in which $V_{\text{ext}}(\mathbf{r}) + g_s|\psi(\mathbf{r})|^2 + D\Phi_{\text{dd}}(\mathbf{r})$ must be constant and equal to the chemical potential within the

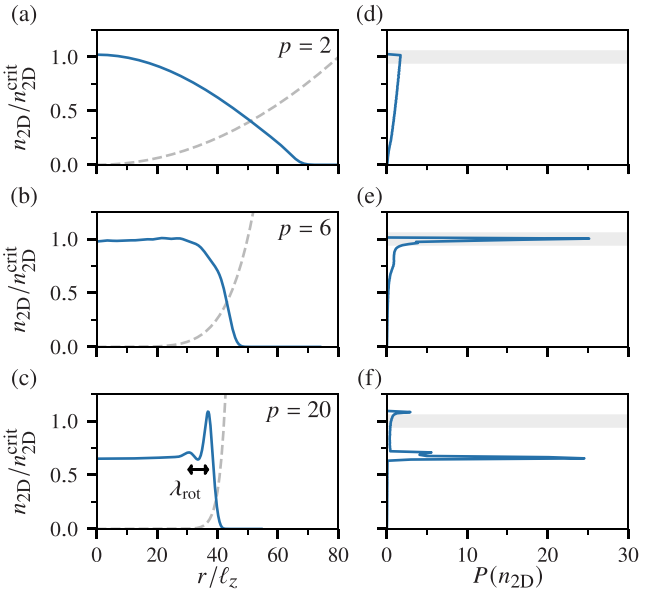

FIG. 2. Critical density distributions of a purely dipolar gas. (a)–(c) Areal density distributions n_{2D} (solid blue lines) at the highest interaction strength at which a stable ground state can form in a power-law potential with $\gamma = 40$ and $p = 2$, 6, and 20, respectively (dashed gray lines). The n_{2D}'s are given relative to n_{2D}^{crit}, the critical density for the roton instability in an *infinite* flattened system; the trap potential is given relative to the chemical potential and shares the same axis. It can be seen that instability in the trapped system occurs when the maximum n_{2D} is close to n_{2D}^{crit}. For $p = 20$, this is due to a pronounced density oscillation near the trap wall, whose wavelength is close to λ_{rot} (see arrow). (d)–(f) Corresponding probability density distributions $P(n_{2D})$ of the areal density (see text). $P(n_{2D})$ is plotted on the horizontal axis such that the vertical axis is shared with plots (a)–(c), the gray shading denotes the region within 5 % of n_{2D}^{crit}. Whereas for the $p = 6$ trap 63 % of the atoms are within 5 % of n_{2D}^{crit}, for both high and low p only a small fraction is.

cloud. In a sharp-walled trap, the contribution of the external potential is negligible. However, due to the long-range nature of the dipolar interactions, if the gas had a homogeneous density distribution, the dipolar term would be significantly reduced near the wall, so the density needs to increase to compensate. For less steep traps, the increasing V_{ext} compensates the decay of Φ_{dd} and so no density accumulation occurs near the edge of the trap.

To better quantify the power-law best suited for studying the physics of a homogeneous system in the roton regime, we define a 'homogeneity' parameter H as the fraction of particles that experience an n_{2D} within 5 % of n_{2D}^{crit}. We note this parameter quantifies how close the system is to a perfectly homogeneous system *at the roton instability*, and not (only) how uniform the density is across the sample. In Fig. 3(a), we plot H against the exponent p for aspect ratios $\gamma = 30$, 50, and 70 for a purely dipolar gas ($\varepsilon_{dd} \to \infty$, solid lines). For all three aspect ratios, H gradually increases with p up to some optimum p^* before dropping sharply with higher p as the peak in $P(n_{2D})$ moves below $0.95\, n_{2D}^{crit}$. We see that p^* increases with γ; it is determined by the p at which significant density starts accumulating near the edge of the trap (cf. Fig. 2(c)). We

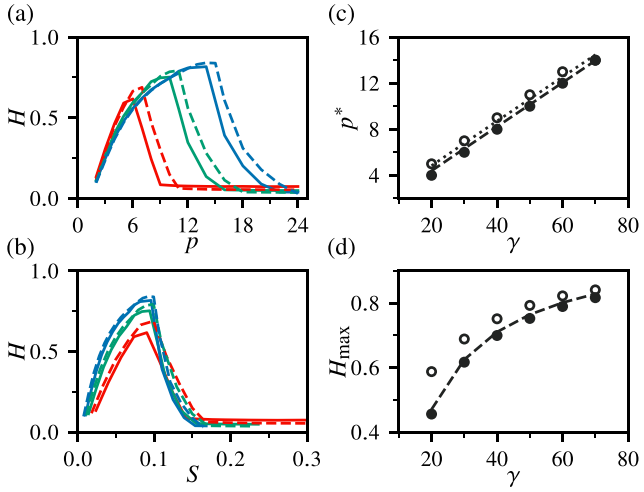

FIG. 3. Optimum power-law. (a) The 'homogeneity' H (see text) as a function of the power-law exponent p for trapping potential aspect ratio $\gamma = 30$ (red, lower lines), 50 (green, middle lines), and 70 (blue, upper lines). The curves are plotted for a purely dipolar gas (solid lines) and for a gas with $\varepsilon_{dd} = 3$ (dashed lines). The optimum $p(= p^*)$ increases with γ. (b) The same data plotted against the trap wall steepness S (see text). The optimum H occurs at the same $S \approx 0.1 (= S^*)$ for all γ. (c) The optimum power-law p^* (dots) and the prediction p^*_{pred} using $S^* = 0.1$ (lines) as a function of γ for a purely dipolar gas (filled dots, dashed line) and for one with $\varepsilon_{dd} = 3$ (empty dots, dotted line). (d) The maximum $H(= H_{max})$, achievable for a given γ (filled dots for a purely dipolar gas, empty dots for $\varepsilon_{dd} = 3$). The dashed line provides a simple interpretation of how H_{max} depends on γ (see text).

have also checked that this behavior is not specific to purely dipolar gases but also applies in the presence of (weak) contact interactions. We show the curves for $\varepsilon_{dd} = 3$ (dashed lines), and see that repulsive contact interactions increase both p^* and the maximum H slightly (attractive contact interactions have the opposite effect).

One would expect density oscillations near the wall to somehow be controlled by the trap 'wall' steepness, which not only depends on p but also on γ. We define the steepness as the gradient of the trap potential (relative to the chemical potential μ) at half the chemical potential:

$$S = \frac{d(V_{trap}/\mu)}{dr}\Big|_{V_{trap}(r,0)=\mu/2} = \frac{p}{2\gamma}\mu^{-1/p}. \tag{5}$$

In Fig. 3(b), we plot H for the same aspect ratios as in Fig. 3(a) but now against S. Plotting this way reveals that the maximum H occurs at the same $S = S^* \approx 0.1$ for all three aspect ratios for a dipolar condensate with or without contact interactions. Given $\{p, \gamma, \mu\}$ uniquely defines S, we can invert Eq. (5) and use S^* to predict

$$p^*_{pred} = \frac{\ln \mu}{W_0\left(\frac{\ln \mu}{2S^*\gamma}\right)}, \tag{6}$$

where $W_0(x)$ is the Lambert W function. To check this, in Fig. 3(c) we plot p^* and p^*_{pred} for a range of γ and see there is very good agreement. This behavior is in contrast to a gas with only contact interactions ($\varepsilon_{dd} = 0$), where

homogeneity would monotonically increase with S, but would saturate when $\delta = 1/S$ (the trap 'wall thickness') reaches the healing length $\xi = 1/\sqrt{2\mu}$ (for our parameters $\xi \approx 0.5 \ll \lambda_{\rm rot}$). In our case, we reach the optimum H at $\delta \approx 10 \gg \xi$, which is close to the roton wavelength $\lambda_{\rm rot} \approx 2\pi$.

As shown in Fig. 3(d), as γ increases, the maximum H ($H_{\rm max}$, achieved at the also growing p^*) increases toward 1, suggesting that the homogeneous limit can still in principle be approached if γ and p are increased together in a suitable way. The trend can be understood via a simple model (dashed line). If we assume the cloud consists of a homogeneous center with radius $\gamma - \lambda_{\rm rot}$ and an inhomogeneous boundary with a width $\lambda_{\rm rot}$, we can estimate $H_{\rm max} \approx (1 - \lambda_{\rm rot}/\gamma)^2$.

Finally, we consider the implications for experimentally realizing a close-to-homogeneous dipolar gas in the roton regime. Unlike for gases with solely repulsive contact interactions, the need for relatively soft walls means that the optics for creating an appropriate trap is unlikely to be a significant constraint. Instead, the limiting factor is likely to be the number of atoms required to fill a high-γ trap. For an approximately uniform gas in the roton regime $n_{\rm 2D} \approx 1/(\pi\gamma^2) \approx n_{\rm 2D}^{\rm crit}$, which using Eq. (4) gives $\gamma^2 = 4Na_{\rm eff}/(\nu_{\rm crit}(\varepsilon_{\rm dd})\ell_z)$. This shows that filling a large-γ trap requires ℓ_z to be small; however, ℓ_z needs to be kept large enough to avoid high (3D) number densities which result in excessive three-body losses. The (dimensionful) peak density can be obtained via the chemical potential and is given by $n_{\rm 3D}^{\rm max} \approx \mu_h(\varepsilon_{\rm dd})/g_{\rm eff} \times N/\ell_z^3 = \mu_h(\varepsilon_{\rm dd})/(4\pi a_{\rm eff}\ell_z^2)$, where $\mu_h(\varepsilon_{\rm dd})$ is the (dimensionless) chemical potential tabulated in Ref. [34]. Solving for ℓ_z and inserting into our expression for γ^2 gives

$$\gamma_{\rm max}^2 \approx \frac{8\pi^{1/2}}{\nu_{\rm crit}\mu_h^{1/2}} N\left(n_{\rm 3D}^{\rm max} a_{\rm eff}^3\right)^{1/2}. \tag{7}$$

Therefore, with 10^5 erbium or dysprosium atoms (for which $a_{\rm dd} \approx 100a_0$, and setting $a_s \approx 0$), if we limit $n_{\rm 3D} \lesssim 100\,\mu{\rm m}^3$, one could reach $\gamma_{\rm max} \approx 40$ with $\ell_z \approx 0.4\,\mu$m (equivalent to a vertical trapping frequency of $\approx 400\,$Hz), resulting in $H \approx 70\,\%$ for $p^* = 8$ (cf. $H \approx 10\,\%$ in a harmonic trap).

In conclusion, we have explored the homogeneity of a dipolar gas, tuned close to its stability boundary, in a flattened, cylindrically symmetric power-law potential. We found that a large exponent in the power-law triggers density oscillations near the trap wall which prevent the bulk of the trap achieving the density a perfectly homogeneous flattened system would have. An intermediate exponent is therefore more suitable, and we found its optimal value is determined by the trap wall steepness, which depends on both the aspect ratio and power-law exponent. These findings guide the way toward the experimental realization of such a homogeneous dipolar gas for the study of, for example, droplet arrays, novel supersolid phases, and critical phenomena.

Data supporting this publication are openly available in Ref. [45].

ACKNOWLEDGMENTS

We thank Tevž Lotrič for valuable discussions and Zoran Hadzibabic and Anna Marchant for comments on the manuscript. This work was supported by the UK EPSRC (Grants No. EP/P009565/1 and No. EP/T019913/1). R.P.S. and P.J. acknowledge support from the Royal Society, P.J. acknowledges support from the Hungarian National Young Talents Scholarship, and M. K. from Trinity College, Cambridge.

APPENDIX

We numerically solve the GPE using both the preconditioned conjugate gradient method [46] and imaginary time propagation with the split-operator technique [47] to cross-check our results. Our algorithm largely follows Ref. [48] with some differences laid out below, and is implemented in PYTHON using several highly efficient and parallelized lower-level libraries for the most computationally expensive parts [49]. Given the trap has axial symmetry, the 3D problem can be reduced to a 2D one computationally. We sample the wave function on a grid and calculate the kinetic and dipolar interaction terms in the GPE using a Hankel transform along r and a cosine transform along z, given the ground state is symmetric with respect to $z = 0$. To avoid interaction between phantom copies of the cloud along z due to the cosine transform, we employ a cutoff of the dipolar interaction in this direction (this is not a problem along r) [48]. The drawback of using a 2D grid is that it does not allow for instability due to angular excitations. To take these into account, we ensure that all angular excitations have a real positive energy using the Bogoliubov–de Gennes (BdG) formalism [48]. To find the largest interaction strength at which a stable condensate can be produced, we employ a binary search technique.

1. Grid

The grid needs to be large enough to comfortably contain the gas, whose size can be estimated using the Thomas–Fermi approximation. We calculate the Thomas–Fermi radii of a gas in our trap with an effective scattering length $a_{\rm eff} = a_s + 2a_{\rm dd}$ [34], and find (in our dimensionless units)

$$R_r = \left(\frac{3\nu_{\rm crit}(\varepsilon_{\rm dd})}{2}\right)^{\frac{2}{3p}}\gamma, \quad R_z = \left(\frac{3\nu_{\rm crit}(\varepsilon_{\rm dd})}{2}\right)^{\frac{1}{3}}, \tag{A1}$$

where $\nu_{\rm crit}$ is tabulated in Ref. [34].

Along z we use a uniform grid with a grid size of 10, which is large enough to avoid interaction between phantom copies of the gas. Note that we use a constant grid size as R_z depends only weakly on $\varepsilon_{\rm dd}$.

The grid along r is (slightly) nonuniform and is defined by $r_j = \alpha_j/\beta$, $j = 0, \ldots, N$ where α_j are the zeros of the first-order Bessel function $J_1(r)$ [50] and β is chosen to give an overall grid size of $1.2R_r$. The Hankel transform can be calculated on this grid with the same computational complexity as in Ref. [48], but additionally it samples the center of the trap.

This grid allows exact integration (for normalization and to calculate the energy and the chemical potential) and interpolation (for expressing the wave function on different grids during the calculation of excitations). Similarly to Ref. [48], for a function $f(r)$ sampled on this grid, it can be shown using

a Dini series expansion that

$$\int_0^\infty f(r) r\, dr = \frac{2}{\beta^2} \sum_{j=0}^{N} f(r_j) J_0^{-2}(\alpha_j). \tag{A2}$$

Exact integrals in k-space can be similarly calculated. Furthermore, like in Ref. [48], using a Dini series expansion again an exact interpolation formula can be derived [51]:

$$f(r) = 2r\beta J_1(r\beta) \sum_{j=0}^{N} \frac{1}{r^2\beta^2 - \alpha_j^2} J_0^{-1}(\alpha_j) f(r_j). \tag{A3}$$

The number of grid points are chosen to be 256×65 ($r \times z$) which ensures adequate sampling of the shortest relevant length scales (the oscillator length $\ell_z = 1$ along z and the roton wavelength $\lambda_{\rm rot} \approx 2\pi$ along r). We checked that our results were insensitive to the exact number of grid points.

2. Excitations

Excitations can in general be written in the form $f(r, \theta, z) = f(r, z)e^{-im\theta}$ [48], where $f(r, z)$ has a definitive symmetry (even or odd) with respect to $z = 0$ and m is the phase winding number of the excitation. For a stable wave function, excitations with any m must have a (real) positive energy and so excitations with a range of m need to be calculated. The highest m excitations that can lead to instability occur for high-p traps, when the peak density is along a ring near the wall. In this case, the lowest-lying excitation can be thought of as a buckling along this ring (an angular roton [17]), such that $m_{\rm crit} \approx 2\pi R_r/\lambda_{\rm rot}$. In practice, we find testing above $1.25 m_{\rm crit}$ is not required and that always even excitations soften first, as the lowest-lying odd excitation is the Kohn mode, with exactly $\hbar\omega_z$ energy [52].

We note that as the eigenvalues of the BdG equations are real for a stable condensate, to find the energy of the lowest-lying excitation it is sufficient to use the 'SR' (smallest real part) mode of ARPACK [48], without a preconditioner, which avoids calculating the inverse of the BdG matrix.

3. Convergence

To ensure our solution to the GPE has adequately converged, we:

(a) required the smallest $m = 0$ eigenvalue of the BdG equations to be (effectively) 0, as the presence of such a neutral mode confirms the ground state has been reached [48];

(b) independently applied both the preconditioned conjugate gradient method [46] and imaginary time propagation with the split-operator technique [47] and checked for consistency.

However, as calculating BdG eigenstates is numerically expensive, we implemented less stringent but numerically much less expensive tests before the $m = 0$ test takes place.

In the case of imaginary time propagation, convergence depends on both the time step size δt and the criteria for halting the imaginary time propagation. For a given δt we assume an exponential convergence (in imaginary time) of the wave function's energy, and require the energy difference from its infinite-time value to be below a certain threshold $E_{\rm tol}$. By considering the change in $\ln(-dE/dt)$ in successive time steps, this energy difference can be calculated using the energy difference δE_i between successive time steps, and the criteria amounts to requiring

$$\frac{\delta E_i}{\ln\left(\frac{\delta E_i}{\delta E_{i-1}}\right)} < E_{\rm tol}. \tag{A4}$$

Making larger time steps δt has the benefit of converging faster, but given the split-step method yields an error in the energy of $\mathcal{O}(\delta t^3)$, it comes at the expense of making larger errors. Therefore, after an iteration set with a certain δt converged, we decrease our δt by $\sqrt[3]{2}$ and continue with this procedure until the energy change between successive δt iteration sets is also smaller than $E_{\rm tol}$. We then do the $m = 0$ BdG lowest eigenvalue test and lower $E_{\rm tol}$ until it is passed.

For the preconditioned conjugate gradient method [46], we follow the approach in Ref. [53], with the choice of the combined (symmetric) preconditioner and the Polak–Ribière formula [54] to enforce the conjugacy criterion. The convergence of this method is determined by a single threshold, by ensuring the energy change between subsequent ground state candidates is not more than $\delta E_{\rm PCG} = 10^{-12}$. We found that this was sufficient to pass the $m = 0$ test.

[1] A. Griesmaier, J. Werner, S. Hensler, J. Stuhler, and T. Pfau, Bose-Einstein Condensation of Chromium, Phys. Rev. Lett. **94**, 160401 (2005).

[2] M. Lu, N. Q. Burdick, S. H. Youn, and B. L. Lev, Strongly Dipolar Bose-Einstein Condensate of Dysprosium, Phys. Rev. Lett. **107**, 190401 (2011).

[3] K. Aikawa, A. Frisch, M. Mark, S. Baier, A. Rietzler, R. Grimm, and F. Ferlaino, Bose-Einstein Condensation of Erbium, Phys. Rev. Lett. **108**, 210401 (2012).

[4] L. Chomaz, I. Ferrier-Barbut, F. Ferlaino, B. Laburthe-Tolra, B. L. Lev, and T. Pfau, Dipolar physics: A review of experiments with magnetic quantum gases, arXiv:2201.02672.

[5] J. L. Bohn, A. M. Rey, and J. Ye, Cold molecules: Progress in quantum engineering of chemistry and quantum matter, Science **357**, 1002 (2017).

[6] M. Saffman, T. G. Walker, and K. Mølmer, Quantum information with Rydberg atoms, Rev. Mod. Phys. **82**, 2313 (2010).

[7] L. Santos, G. V. Shlyapnikov, and M. Lewenstein, Roton-Maxon Spectrum and Stability of Trapped Dipolar Bose-Einstein Condensates, Phys. Rev. Lett. **90**, 250403 (2003).

[8] L. Chomaz, R. M. W. van Bijnen, D. Petter, G. Faraoni, S. Baier, J. H. Becher, M. J. Mark, F. Wächtler, L. Santos, and F. Ferlaino, Observation of roton mode population in a dipolar quantum gas, Nat. Phys. **14**, 442 (2018).

[9] F. Böttcher, J.-N. Schmidt, M. Wenzel, J. Hertkorn, M. Guo, T. Langen, and T. Pfau, Transient Supersolid Properties in an Array of Dipolar Quantum Droplets, Phys. Rev. X **9**, 011051 (2019).

[10] L. Tanzi, E. Lucioni, F. Famà, J. Catani, A. Fioretti, C. Gabbanini, R. N. Bisset, L. Santos, and G. Modugno, Observation of a Dipolar Quantum Gas with Metastable Supersolid Properties, Phys. Rev. Lett. **122**, 130405 (2019).

[11] L. Chomaz, D. Petter, P. Ilzhöfer, G. Natale, A. Trautmann, C. Politi, G. Durastante, R. M. W. van Bijnen, A. Patscheider, M. Sohmen, M. J. Mark, and F. Ferlaino, Long-Lived and Transient Supersolid Behaviors in Dipolar Quantum Gases, Phys. Rev. X **9**, 021012 (2019).

[12] L. D. Landau, J. Phys. USSR **11**, 91 (1947).

[13] H. Kadau, M. Schmitt, M. Wenzel, C. Wink, T. Maier, I. Ferrier-Barbut, and T. Pfau, Observing the Rosensweig instability of a quantum ferrofluid, Nature **530**, 194 (2016).

[14] I. Ferrier-Barbut, H. Kadau, M. Schmitt, M. Wenzel, and T. Pfau, Observation of Quantum Droplets in a Strongly Dipolar Bose Gas, Phys. Rev. Lett. **116**, 215301 (2016).

[15] L. Chomaz, S. Baier, D. Petter, M. J. Mark, F. Wächtler, L. Santos, and F. Ferlaino, Quantum-Fluctuation-Driven Crossover from a Dilute Bose-Einstein Condensate to a Macrodroplet in a Dipolar Quantum Fluid, Phys. Rev. X **6**, 041039 (2016).

[16] F. Böttcher, J.-N. Schmidt, J. Hertkorn, K. S. H. Ng, S. D. Graham, M. Guo, T. Langen, and T. Pfau, New states of matter with fine-tuned interactions: Quantum droplets and dipolar supersolids, Rep. Prog. Phys. **84**, 012403 (2021).

[17] S. Ronen, D. C. E. Bortolotti, and J. L. Bohn, Radial and Angular Rotons in Trapped Dipolar Gases, Phys. Rev. Lett. **98**, 030406 (2007).

[18] R. M. Wilson, S. Ronen, J. L. Bohn, and H. Pu, Manifestations of the Roton Mode in Dipolar Bose-Einstein Condensates, Phys. Rev. Lett. **100**, 245302 (2008).

[19] P. B. Blakie, D. Baillie, and S. Pal, Variational theory for the ground state and collective excitations of an elongated dipolar condensate, Commun. Theor. Phys. **72**, 085501 (2020).

[20] S. M. Roccuzzo, A. Gallemí, A. Recati, and S. Stringari, Rotating a Supersolid Dipolar Gas, Phys. Rev. Lett. **124**, 045702 (2020).

[21] Y.-C. Zhang, T. Pohl, and F. Maucher, Phases of supersolids in confined dipolar Bose-Einstein condensates, Phys. Rev. A **104**, 013310 (2021).

[22] E. Poli, T. Bland, C. Politi, L. Klaus, M. A. Norcia, F. Ferlaino, R. N. Bisset, and L. Santos, Maintaining supersolidity in one and two dimensions, Phys. Rev. A **104**, 063307 (2021).

[23] A. Gallemí, S. M. Roccuzzo, S. Stringari, and A. Recati, Quantized vortices in dipolar supersolid Bose-Einstein-condensed gases, Phys. Rev. A **102**, 023322 (2020).

[24] G. Natale, R. M. W. van Bijnen, A. Patscheider, D. Petter, M. J. Mark, L. Chomaz, and F. Ferlaino, Excitation Spectrum of a Trapped Dipolar Supersolid and its Experimental Evidence, Phys. Rev. Lett. **123**, 050402 (2019).

[25] J. Hertkorn, F. Böttcher, M. Guo, J. N. Schmidt, T. Langen, H. P. Büchler, and T. Pfau, Fate of the Amplitude Mode in a Trapped Dipolar Supersolid, Phys. Rev. Lett. **123**, 193002 (2019).

[26] L. Tanzi, S. M. Roccuzzo, E. Lucioni, F. Famà, A. Fioretti, C. Gabbanini, G. Modugno, A. Recati, and S. Stringari, Supersolid symmetry breaking from compressional oscillations in a dipolar quantum gas, Nature **574**, 382 (2019).

[27] M. Guo, F. Böttcher, J. Hertkorn, J.-N. Schmidt, M. Wenzel, H. P. Büchler, T. Langen, and T. Pfau, The low-energy Goldstone mode in a trapped dipolar supersolid, Nature **574**, 386 (2019).

[28] J. Hertkorn, J.-N. Schmidt, M. Guo, F. Böttcher, K. S. H. Ng, S. D. Graham, P. Uerlings, T. Langen, M. Zwierlein, and T. Pfau, Pattern formation in quantum ferrofluids: From supersolids to superglasses, Phys. Rev. Res. **3**, 033125 (2021).

[29] J. Hertkorn, J.-N. Schmidt, M. Guo, F. Böttcher, K. S. H. Ng, S. D. Graham, P. Uerlings, H. P. Büchler, T. Langen, M. Zwierlein, and T. Pfau, Supersolidity in Two-Dimensional Trapped Dipolar Droplet Arrays, Phys. Rev. Lett. **127**, 155301 (2021).

[30] P. Ilzhöfer, M. Sohmen, G. Durastante, C. Politi, A. Trautmann, G. Natale, G. Morpurgo, T. Giamarchi, L. Chomaz, M. J. Mark, and F. Ferlaino, Phase coherence in out-of-equilibrium supersolid states of ultracold dipolar atoms, Nat. Phys. **17**, 356 (2021).

[31] M. A. Norcia, C. Politi, L. Klaus, E. Poli, M. Sohmen, M. J. Mark, R. N. Bisset, L. Santos, and F. Ferlaino, Two-dimensional supersolidity in a dipolar quantum gas, Nature **596**, 357 (2021).

[32] M. N. Tengstrand, D. Boholm, R. Sachdeva, J. Bengtsson, and S. M. Reimann, Persistent currents in toroidal dipolar supersolids, Phys. Rev. A **103**, 013313 (2021).

[33] T. Bland, E. Poli, C. Politi, L. Klaus, M. A. Norcia, F. Ferlaino, L. Santos, and R. N. Bisset, Two-Dimensional Supersolid Formation in Dipolar Condensates, Phys. Rev. Lett. **128**, 195302 (2022).

[34] D. Baillie and P. B. Blakie, A general theory of flattened dipolar condensates, New J. Phys. **17**, 033028 (2015).

[35] S. M. Roccuzzo and F. Ancilotto, Supersolid behavior of a dipolar Bose-Einstein condensate confined in a tube, Phys. Rev. A **99**, 041601(R) (2019).

[36] Y.-C. Zhang, F. Maucher, and T. Pohl, Supersolidity Around a Critical Point in Dipolar Bose-Einstein Condensates, Phys. Rev. Lett. **123**, 015301 (2019).

[37] P. B. Blakie, D. Baillie, L. Chomaz, and F. Ferlaino, Supersolidity in an elongated dipolar condensate, Phys. Rev. Res. **2**, 043318 (2020).

[38] F. Ancilotto, M. Barranco, M. Pi, and L. Reatto, Vortex properties in the extended supersolid phase of dipolar Bose-Einstein condensates, Phys. Rev. A **103**, 033314 (2021).

[39] B. K. Turmanov, B. B. Baizakov, F. K. Abdullaev, and M. Salerno, Oscillations of a quasi-one-dimensional dipolar supersolid, J. Phys. B: At., Mol. Opt. Phys. **54**, 145302 (2021).

[40] S. Pal, D. Baillie, and P. B. Blakie, Infinite dipolar droplet: A simple theory for the macrodroplet regime, Phys. Rev. A **105**, 023308 (2022).

[41] N. Navon, R. P. Smith, and Z. Hadzibabic, Quantum gases in optical boxes, Nat. Phys. **17**, 1334 (2021).

[42] H.-Y. Lu, H. Lu, J.-N. Zhang, R.-Z. Qiu, H. Pu, and S. Yi, Spatial density oscillations in trapped dipolar condensates, Phys. Rev. A **82**, 023622 (2010).

[43] S. M. Roccuzzo, S. Stringari, and A. Recati, Supersolid edge and bulk phases of a dipolar quantum gas in a box, Phys. Rev. Res. **4**, 013086 (2022).

[44] 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).

[45] P. Juhász, M. Krstajić, D. Strachan, E. Gandar, and R. P. Smith, How to realise a homogeneous dipolar Bose gas in the roton regime (data), University of Oxford, doi:10.5287/bodleian:9RavjqbVB (2022).

[46] X. Antoine, A. Levitt, and Q. Tang, Efficient spectral computation of the stationary states of rotating Bose-Einstein condensates by preconditioned nonlinear conjugate gradient methods, J. Comput. Phys. **343**, 92 (2017).

[47] M. Feit, J. Fleck, and A. Steiger, Solution of the Schrödinger equation by a spectral method, J. Comput. Phys. **47**, 412 (1982).

[48] S. Ronen, D. C. E. Bortolotti, and J. L. Bohn, Bogoliubov modes of a dipolar condensate in a cylindrical trap, Phys. Rev. A **74**, 013623 (2006).

[49] Matrix multiplications are performed using MKL via NUMPY, Fourier transforms are calculated via MKL_FFT. Piecewise array operations are executed using NUMBA, a package which turns PYTHON code into parallelized machine code.

[50] K.-M. You, S.-C. Wen, L.-Z. Chen, Y.-W. Wang, and Y.-H. Hu, A quasi-discrete Hankel transform for nonlinear beam propagation, Chin. Phys. B **18**, 3893 (2009).

[51] This formula is not well defined for $r = \alpha_j/\beta$, but in that case $r = r_j$ and the known $f(r_j)$ can be directly used.

[52] A. L. Fetter and D. Rokhsar, Excited states of a dilute Bose-Einstein condensate in a harmonic trap, Phys. Rev. A **57**, 1191 (1998).

[53] X. Antoine, Q. Tang, and Y. Zhang, A preconditioned conjugated gradient method for computing ground states of rotating dipolar Bose-Einstein condensates via kernel truncation method for dipole-dipole interaction evaluation, Commun. Comput. Phys. **24**, 966 (2018).

[54] E. Polak and G. Ribière, Note sur la convergence de méthodes de directions conjuguées, Rev. Fr. Inform. Rech. Opér. **3**, 35 (1969).

7
Conclusion

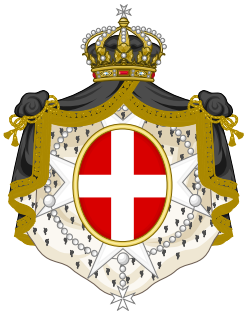

Sovereign Order of Malta. *Gules, a cross Argent.*

This thesis reported on three important steps towards realising homogeneous dipolar Bose–Einstein condensates.

First, after presenting an overview of our experimental apparatus and the techniques used therein, we reported our optimised experimental sequence to produce a quantum degenerate gas of erbium and demonstrated our ability to create a quasi-pure BEC with 2.2×10^5 atoms. This is a larger sample than what have been used so far in experiments probing the roton regime of dipolar physics and presents experimental progress towards probing exotic quantum phases.

Second, we explored density- and temperature-dependent losses in ^{166}Er and identified six previously unreported resonant loss features. It is with this knowledge that we were able to optimise our BEC production sequence. We found that losses in ^{166}Er have a strong temperature dependence, suggesting a higher partial-wave character. We also presented an extension of the loss model and found that it describes the temperature evolution of the atom cloud well.

Finally, we presented a theoretical study of strongly dipolar condensates in box-like traps. We explored how one can replicate a fully homogeneous system within a trap, and explored the stability of condensates in box-like traps. We found that traps with hard walls trigger roton-like density oscillations even if the bulk of the system is far from the roton

regime, so smoother potentials are better suited to recreate homogeneous conditions. This gives important implications for designing an optical box trap for dipolar atoms.

7.1 Outlook

The experimental and theoretical progress presented in this thesis set the ground for future experiments. To enable studies of systems tightly trapped in one direction but homogeneous in the other two, from an experimental point of view, the atoms have to be transported to the science cell where a box trap will be implemented. To this end, the optical setup to project the walls of the box trap (using a spatial light modulator) has been tested and the laser producing the attractive sheet beam has been sourced. The coils around the cell have also been designed and manufactured, and are currently being installed on the system. Finally, a system to project arbitrary optical potentials (e.g. for Bragg spectroscopy, using a digital micromirror device) has also been tested (see Ref. 104 and references therein). An extension to add potassium as a second species has also been designed and manufactured. The laser locking system has been implemented and the relevant vacuum parts have been attached to the chamber, and will be pumped down to sufficient vacuum conditions in the months to come.

The apparatus is set to be used for a variety of physical experiments. The long-range interactions, through their effect on the excitation spectrum, can affect the critical temperature of the gas. The roton-like excitation spectrum of a dipolar BEC, as explained in §2.3.3, also introduces a variety of phenomena which can be explored, one of which is its effect on the superfluid critical velocity. This can be deduced by measuring the excitation spectrum directly using Bragg spectroscopy [72, 305, 306] or by stirring the BEC with a blue-detuned laser [307]. This technique can also be used to investigate the role of vortices in this context. Furthermore, by tuning the roton gap via Feshbach resonances, the system can be driven to form quantum droplets, exhibiting supersolidity in some cases. To this end, the first quantum droplets have been recently realised with our system (see Fig. 7.1).

Another phenomenon which can be explored is the effect of long-range interactions on the dynamics of the BEC phase transition. The transition from a thermal gas to a BEC is a second-order phase transition, i.e. the order parameter (the macroscopic wave function) changes continuously, but the phase symmetry is broken as the system crosses the critical temperature. If the transition is crossed infinitely slowly (the so-called adiabatic limit), the entire condensate will assume a state with a unique phase. However, if the transition is crossed at a finite rate, domains with different order-parameter values appear due to the limited range over which correlations can spread before being frozen out by the transition.

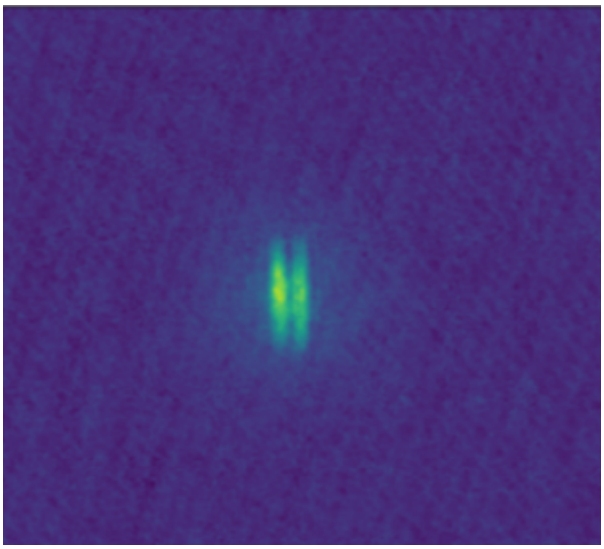

Figure 7.1. Quantum droplets. Absorption image of quantum droplets taken after a time of flight (false colour). Quantum droplets develop if the system is driven close to instability by decreasing the relative strength of contact interactions. Furthermore, if the droplets establish a common phase, the system can exhibit supersolid properties.

This mechanism is described by Kibble–Zurek (KZ) theory [39, 308, 309], which relies on the universal scaling behaviour of system parameters in the vicinity of the critical point. It predicts that the correlation length ξ of the system obeys a power-law scaling with distance from the critical point, $\xi \propto |\varepsilon|^{-\nu}$, where $\varepsilon = T/T_c - 1$. Furthermore, the freeze-out correlation length depends on the time scale of crossing the transition (see Fig. 7.2). There are similar scaling relations for the heat capacity, order parameter magnitude etc., each with its own scaling exponent, and it has been shown that the exponents are insensitive to microscopic properties of the system and are rather determined by the dimensionality and symmetries of the system and the order parameter, yielding certain universality classes [310, 311]. As the long-range nature of the dipole–dipole interaction could alter the mechanism of correlation propagation through the system, it could result in a different set of scaling exponents (and hence a different universality class).

Ultracold-atom platforms proved to be well-equipped for exploring this experimentally, due to the tunability of system parameters and the readily accessible critical phenomena [39]. Given that the BEC transition is sensitive to density, systems with a homogeneous density have a crucial advantage in these efforts, as all parts of the system cross the phase transition at the same time [48, 49, 312, 313]. The transition dynamics can be explored via temperature quenches, but also by changing the strength of contact interactions via Feshbach resonances (due to the dependence of T_c on interaction strengths). This process is, unlike evaporative cooling, reversible, so the reversibility of

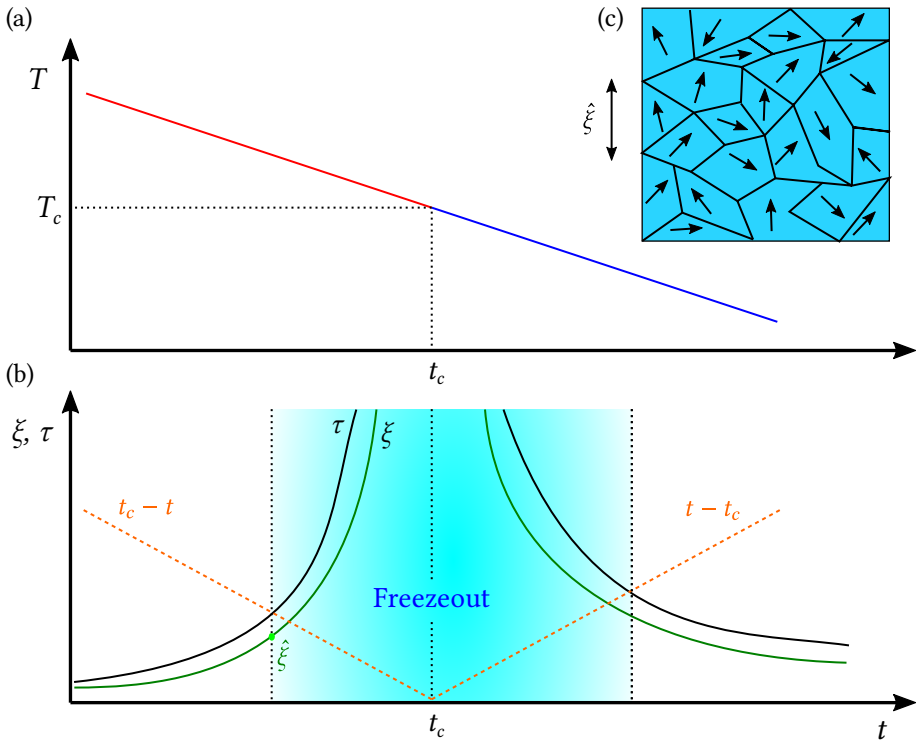

Figure 7.2. Kibble–Zurek theory. (a) A linear temperature (T) ramp in time, inducing a second-order phase transition when the temperature reaches T_c, at time t_c. (b) The correlation length ξ and the associated relaxation time τ (the time scale over which the system adjusts to a new equilibrium value of ξ following a change in temperature) as a function of time for a linear temperature ramp. As both diverge near the critical point, correlations freeze out when the relaxation time becomes longer than the time remaining to the crossing of the transition. (c) Freeze-out results in domains developing with a characteristic size defined by the finite freeze-out correlation length. Figure adapted from Ref. 39.

the transition and the presence of any hysteretic artefacts (arising from crossing the transition multiple times) could be examined. It is worth mentioning that the supersolid transition [80–82] could also be explored, which, being a quantum phase transition, could probe the quantum limit of KZ theory.

Physical phenomena in ultracold dipolar systems can be further enriched by adding another species. By adding potassium atoms to the platform, making it a dual-species experiment, the system will be capable of exploring impurity physics in dipolar bosonic systems. There are two directions of research which could be explored. First, Bose polarons could be investigated (a potassium impurity moving through an erbium BEC background), which can be made anisotropic via the dipole force [93]. As the complex and anisotropic electron structure of erbium yields a rich spectrum of Feshbach resonances in Er–Er collisions, we expect a similar abundance of Feshbach resonances for Er–K scattering. This, combined with the long-range and anisotropic dipolar force, should enable us to

flexibly tune the dispersion relation of the polarons.

Furthermore, the tunability of the excitation spectrum of the erbium BEC makes it a promising candidate for creating a controllable quantum reservoir for quantum computing. As the spectrum softens, the appearance of the roton minimum and the accompanying maxon leads to a high density of states at the roton and maxon frequencies. As shown in Refs. 94 and 95, the information flow from a single-atom qubit immersed in a dipolar condensate can be controlled by tuning the parameters of the reservoir, as the evolution of decoherence is dominated by the frequencies corresponding to the roton and the maxon features, leading to non-Markovian behaviour and the possibility of information flow back from the reservoir into the qubit system.

Appendices

A
High-power beams

Most of our optics and devices were not prepared to bear the high powers and intensities of the IR lasers that we use. In fact, another group have already abandoned using the same type of lenses and switched to fixed focal length lenses made of a different material (UV fused silica instead of N-BK7, as the former has a lower thermal expansion coefficient, causing a smaller thermal lensing effect) and also decided to use a transport system with lenses on a translation stage instead of the focus-tunable lenses (FTLs) precisely because they do not operate well under high power [314].

This was not easy to foresee as the FTLs have a damage threshold that is much higher than the laser intensities we use. In line with this, the lenses are not permanently damaged in our system but show severe thermal lensing, meaning they can only be used during the first few seconds after the laser reaches them and then they need to be left cooling for around 30 s. On top of this, these lenses have to be used with their optical axes vertical. This is not ideal, but in our experimental sequence we will only have high power on the lenses for a short time, so we are going ahead with this solution. However, alternatives are being considered, such as a low-vibration translation stage [239] and Moiré lenses,[1] a different type of tunable lens which is reportedly significantly more immune to aberrations and thermal lensing effects [315]. Regarding the fixed focal length lenses, using N-BK7 lenses seemed suitable for our case, and no difference could be seen when lenses made of UV fused silica were substituted for (the cheaper) lenses made of N-BK7.[2]

Another problem with characterising our setup was imaging high intensities (e.g. the ODT at full power) as most intensity filters break at such high powers and cubes do not work perfectly either.

These problems were not straightforward to circumvent, and the high-power imaging techniques are presented in Appendix A.1 along with a quantisation of beam quality presented in Appendix A.2. A few images of the various problems are presented in Fig. A.1.

[1] Manufactured by Diffratec.
[2] Both manufactured by Thorlabs.

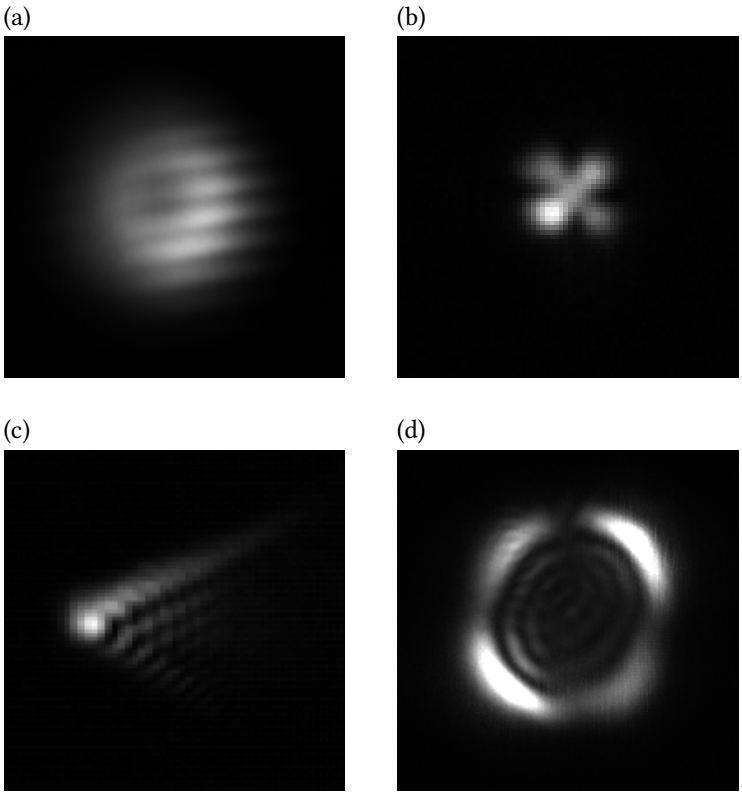

Figure A.1. Issues with the laser beam. (a) The beam coming out of the laser near its waist. The beam is doughnut-shaped, due to the cladding modes present. A Gaussian shape can be restored by filtering these modes out, putting an iris at the focus. (b) This beam shape is a high-power imaging artefact, caused by internal reflections and the finite extinction ratio when a cube and wave plate are used to reduce the power to image the beam. (c) Comatic aberration when the focus-tunable lens is used with its optical axis horizontal. This can be solved by turning the lens to have its optical axis vertical. (d) Thermal lensing, caused by the heating up of the focus-tunable lens. A severely distorted beam shape develops over 5 s, so the lens can only be used for a few seconds after which it needs to cool down, taking 30 s.

A.1 High-power beam imaging

It was not straightforward to image a 45 W beam, given no ordinary camera or filter can take such high power. As the filters can take up to 10 mW at the beam waists we use, we need to filter out as much as 99.98% of the light. Traditionally, to achieve this, a combination of a half-wave plate placed in front of a polarising beamsplitter cube is used. The wave plate rotates the polarisation of the light which is then split by the cube, so by rotating the wave plate, the transmitted light intensity can be controlled. In our case, the cube transmits horizontally polarised light, and since the laser polarisation is

vertical it has to be rotated by only a very small amount to get a suitably low power in the transmitted wave.

The problem with this is that there is a horizontally polarised part of the laser light that has a strange pattern, which then also gets transmitted through the cube. While the power in this originally horizontal part (power $H = 40\,\mathrm{mW}$) is much smaller than that of the vertical part (power $V = 45\,\mathrm{W}$), it is comparable to the transmitted portion of the vertical part we really want to look at, and so we observe a pattern that is really not Gaussian (cf. Fig. A.1(b)). We also tried to image the beam by trying to look at the transmitted part through a back-polished mirror, but this was not successful as the transmission ratio was too low.

The solution was introducing a beam sampler between the wave plate and the cube. To get the smallest amount of light through the system, we need to look at the p-polarised (i.e. horizontal) light as the reflectance of the beam sampler is much smaller for p-polarisation than it is for s-polarisation (i.e. vertical). The idea then is that we turn the wave plate so it rotates the polarisation of the laser by 90°, so a small part of V and a larger part of H is reflected. This is not a problem, since $V \gg H$, so the cube will now be able to filter out the H part which is now vertical, enabling us to look at the transmitted V part.

Mathematically, if the wave plate rotates the polarisation by θ, the power in the horizontal part of the beam after the wave plate (p-polarisation) will be $P_p = V \sin^2\theta + H\cos^2\theta$, while in the new vertical part (s-polarisation) it will be $P_s = V\cos^2\theta + H\sin^2\theta$. If the reflection coefficients for the two polarisations are R_p and R_s ($R_s \gg R_p$), the total reflected power is $P = R_p P_p + R_s P_s$. Due to the large power imbalance between the p and the s part, we have to take into account the finite extinction ratio of the cube, e. Assuming a horizontal transmittance close to 1, the total power transmitted through the cube is $P_t = R_p P_p + R_s P_s/e$. The coefficients $R_{p,s}$ are a function of the angle of incidence θ_i of the laser on the beam sampler and the refractive index of the beam sampler n, and can be written as

$$R_p = \left| \frac{\sqrt{1 - \left(\frac{\sin\theta_i}{n}\right)^2} - n\cos\theta_i}{\sqrt{1 - \left(\frac{\sin\theta_i}{n}\right)^2} + n\cos\theta_i} \right|^2, \tag{A.1}$$

$$R_s = \left| \frac{n\sqrt{1 - \left(\frac{\sin\theta_i}{n}\right)^2} - \cos\theta_i}{n\sqrt{1 - \left(\frac{\sin\theta_i}{n}\right)^2} + \cos\theta_i} \right|^2. \tag{A.2}$$

In theory, we can then make P_s arbitrarily small by adjusting θ_i, but for our beam

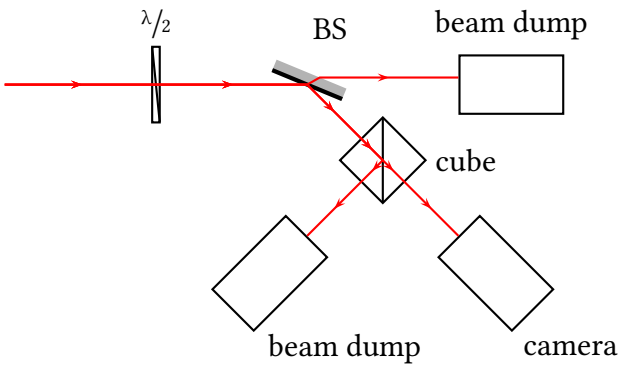

Figure A.2. System to image a high-power beam. The beam goes through a wave plate rotating its polarisation by $\theta = 90°$ and then reflects from a beam sampler (BS) at ca. $\theta_i = 55°$ angle of incidence. It then goes through a polarising beam splitter cube so the horizontal polarisation is kept and then it is imaged with a camera.

sampler[3] the smallest we can make R_p is 0.0022% with $R_s \approx 1000R_p$ at the same angle. We also see that the extinction ratio for the transmitted beam of our cube[4] is $e \approx 1000$.

If $\theta = 0°$, the transmitted power is $P_0 = R_pH + R_sV/e$, whereas if $\theta = 90°$, the transmitted power is $P_{90} = R_pV + R_sH/e$. These are roughly equal, as by substituting the known ratios we get $P \approx R_p(V + H)$ in both cases, which was confirmed by measurement. Since $V/H \approx 1000$, we get $P \approx R_pV$, so we are able to look at only the V part in both cases. However, we prefer to look at the $\theta = 90°$ case as then we can look at a directly transmitted V part rather than relying on it being transmitted due to the non-perfect finite extinction ratio of the cube.

With this technique we can get rid of all the high-power imaging artefacts. The system can be seen in Fig. A.2.

A.2 Beam quality

Due to observing various different non-Gaussian patterns, we had to devise a metric describing how close our beam is to a perfect Gaussian. While there are many different ways to do this, what we are really interested in is how well trapping would work, i.e. how high the peak intensity of the beam is compared to a Gaussian beam with the same total power.

[3] BSF10-B by Thorlabs.
[4] CCM1-PBS253/M by Thorlabs.

To measure the peak intensity, we look at our beams with a camera with square pixels of size a. The intensity distribution of a Gaussian beam is given by

$$I(x,y) = I_0 e^{-\frac{2x^2}{w_x^2}-\frac{2y^2}{w_y^2}}, \tag{A.3}$$

where I_0 is the peak intensity and $w_{x,y}$ are its waists. The brightness of each pixel depends on the power incident on the pixel, and at the central pixel we will record a power of

$$P_c = \int_{-\frac{a}{2}}^{\frac{a}{2}} \int_{-\frac{a}{2}}^{\frac{a}{2}} I(x,y)\,\mathrm{d}x\,\mathrm{d}y = \frac{1}{2} I_0 \pi w_x w_y \operatorname{erf}\left(\frac{a}{\sqrt{2}w_x}\right) \operatorname{erf}\left(\frac{a}{\sqrt{2}w_y}\right). \tag{A.4}$$

The total power in the beam is

$$P_0 = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} I(x,y)\,\mathrm{d}x\,\mathrm{d}y = \frac{1}{2} I_0 \pi w_x w_y, \tag{A.5}$$

so the relative power on the central pixel compared to the total power for a perfectly Gaussian beam is

$$P_{\mathrm{rel}} = \frac{P_c}{P_0} = \operatorname{erf}\left(\frac{a}{\sqrt{2}w_x}\right) \operatorname{erf}\left(\frac{a}{\sqrt{2}w_y}\right). \tag{A.6}$$

For the imaged beam, we can record the total power by summing all pixel values. To get the central power, we can select the brightest pixel and claim it is the centre of the beam, where atoms will be trapped, or if the beam is reasonably Gaussian, we can also fit a Gaussian to it and extract the peak power from that. The ratio of these two will be the experimentally measured $P_{\mathrm{rel,\,exp}}$ which we can compare to the theoretical P_{rel} we expect, and create the metric

$$I_{\mathrm{rel}} = \frac{P_{\mathrm{rel,\,exp}}}{P_{\mathrm{rel}}}. \tag{A.7}$$

If $I_{\mathrm{rel}} \approx 1$, it means the peak intensity is close to what we expect of a Gaussian beam, but if $I_{\mathrm{rel}} < 1$, it means the peak intensity is smaller than what we would like to see.

B

Extended three-body loss model

Here we estimate the geometrical prefactors G and G' for the extended loss model presented in Chapter 5 and also present the loss model for a BEC.

B.1 Three-body loss in a thermal gas

B.1.1 Event rate

Let us start by calculating the probability that a particle suffers a three-body collision and the associated rate of collisions. The event rate (per unit volume) of three particles colliding is proportional to $n^3(\mathbf{r})$. Writing the constant of proportionality as $L_3/3$ (which is equivalent to saying that the probability of a three-body event per particle in unit time is $P_{3B}(\mathbf{r}) = L_3 n^2(\mathbf{r})/3$) and recalling that three particles are lost from the trap at each event, one recovers the three-body loss rate (per unit volume)

$$\dot{n}_3(\mathbf{r}) = -L_3 n^3(\mathbf{r})\,. \tag{B.1}$$

In a thermal gas, this leads to the following equations for the evolution of atom number and temperature (cf. Eqs. (5.3) and (5.8)):

$$\dot{N}_0 = -L_3 \left(\frac{m\bar{\omega}^2}{2\sqrt{3}\pi k_B}\right)^3 \frac{N^3}{T^3}\,, \tag{B.2}$$

$$\dot{T}_0 = \frac{L_3}{3} \left(\frac{m\bar{\omega}^2}{2\sqrt{3}\pi k_B}\right)^3 \frac{N^2}{T^2}\,, \tag{B.3}$$

where the subscript denotes the fact that these do not take into account secondary collisions. While this simple model is valid in low-density regimes, for higher-density clouds, secondary collisions of the collision products with the rest of the atomic cloud need to be accounted for.

B.1.2 Secondary collisions

Let us consider the probability of a secondary collision given a 'projectile' particle of velocity $\mathbf{v}$, starting at position $\mathbf{r}_0$ and moving through a cloud of static 'targets' of (number) density n. In time Δt, the probability that the projectile will have hit a target on its way is given by $P_{\text{coll}} = \sigma n v \Delta t$, where σ is the scattering cross section. In time $t = \sum_i \Delta t_i$, the probability of avoiding a collision is

$$P_{\text{surv}} = \lim_{\Delta t \to 0} \prod_i (1 - \sigma n v \Delta t_i) = \lim_{\Delta t \to 0} (1 - \sigma n v \Delta t)^{t/\Delta t} = e^{-n\sigma v t}. \tag{B.4}$$

As harmonically trapped gasses are not homogeneous, one has to account for time-varying target density as the projectile travels through the cloud, leading to the modified expression

$$P_{\text{surv}} = \lim_{\Delta t \to 0} \prod_i e^{-\sigma v n(t_i)\Delta t_i} = \lim_{\Delta t \to 0} e^{-\sigma v \sum_i n(t_i)\Delta t_i} = e^{-\sigma v \int_0^t n(t')\,\mathrm{d}t'}. \tag{B.5}$$

If we assume that the projectile particle is moving along a straight line, i.e. it is minimally perturbed by the trapping potential, we can write $\mathbf{r}(t) = \mathbf{r}_0 + \mathbf{v}t$. Substituting $n(t) = n(\mathbf{r}(t))$ into Eq. (B.5) and averaging over initial velocities[1] gives

$$P_{\text{surv}}(\mathbf{r}_0, \bar{\omega}, T) = \exp\left(-\sigma n(\mathbf{r}_0) \sqrt{\frac{3\pi}{2m\beta \sum_i \omega_i^2}} e^{\beta \tilde{U}(\mathbf{r}_0)}\right), \tag{B.6}$$

where we defined $\tilde{U}(\mathbf{r}) = m \sum_i \omega_i^4 r_i^2 / 2 \sum_i \omega_i^2$ and $\beta = 1/k_B T$. We note that the survival probability is independent of the projectile velocity as expected.

We can extend the dilute-cloud loss model by considering three-body collision events in which one or both of the products encounter a cloud (target) particle on their way out of the trap. In all of the subsequent calculations we assume $P_{\text{surv}} \approx 1$, i.e. $P_{\text{coll}} = 1 - P_{\text{surv}} \ll 1$ so we can neglect terms of $O\left(P_{\text{coll}}^2\right)$.

B.1.3 Ideal three-body loss

The dominant contribution to atom loss comes from events when both the formed molecule and the remaining (singular) particle escape the cloud (a 'double escape'). The

[1] It is important to note that in deriving Eq. (B.6) we do not average velocities in a mathematically rigorous way. We simply set $v_x^2 = v_y^2 = v_z^2 = v^2/3$ in an equipartition-like assumption.

probability of such an event per particle in unit time is

$$P_{\mathrm{DE}}(\mathbf{r}_0) = P_{\mathrm{3B}} P_{\mathrm{esc}} \tilde{P}_{\mathrm{esc}} = P_{\mathrm{3B}}(1 - P_{\mathrm{coll}})(1 - \tilde{P}_{\mathrm{coll}}) \approx P_{\mathrm{3B}}(1 - P_{\mathrm{coll}} - \tilde{P}_{\mathrm{coll}}) \,, \tag{B.7}$$

where the tilde is used to distinguish between quantities referring to the singular particle (plain) and the molecule (tilde) and $\mathbf{r}_0$ is the position where the collision takes place. (Note that both P_{3B} and P_{coll} depend on position, but we did not mark their explicit position dependence for ease of notation.) This process leads to the particle and energy loss rates

$$\dot{N}_{\mathrm{DE}} = -3 \int n(\mathbf{r}_0) P_{\mathrm{3B}}(1 - P_{\mathrm{coll}} - \tilde{P}_{\mathrm{coll}}) \, \mathrm{d}^3\mathbf{r}_0 \,, \tag{B.8}$$

$$\dot{E}_{\mathrm{DE}} = -3 \int n(\mathbf{r}_0) \bar{E}(\mathbf{r}_0) P_{\mathrm{3B}}(1 - P_{\mathrm{coll}} - \tilde{P}_{\mathrm{coll}}) \, \mathrm{d}^3\mathbf{r}_0 \,, \tag{B.9}$$

where $\bar{E}(\mathbf{r}_0)$ is the mean energy of a particle at $\mathbf{r}_0$ and the factor of three accounts for the fact that all three particles participating in the collision are lost (i.e. three particles colliding at $\mathbf{r}_0$ have energy $3\bar{E}(\mathbf{r}_0)$ on average).

B.1.4 Higher-order corrections

The remaining two cases (where either the molecule or the singular particle undergoes a collision on its way out while the other leaves without colliding) mirror each other, so we will only describe one of the processes. The probability that both collide is $O\left(P_{\mathrm{coll}}^2\right)$ and can therefore be neglected.

The probability of singular particle collision and molecular escape is $P_{\mathrm{ME}} = P_{\mathrm{3B}} P_{\mathrm{coll}}(\mathbf{r}_0)$ $\times \tilde{P}_{\mathrm{esc}}(\mathbf{r}_0) \approx P_{\mathrm{3B}} P_{\mathrm{coll}}(\mathbf{r}_0)$. (The subscript 'ME' stands for 'molecular escape' and 'PE' will be used for 'particle escape'.) The effect of the collision on the particle number and energy of the cloud needs some careful thought. If we assume that the particles undergo elastic collisions, there are three possible outcomes of a collision of the singular 'bullet' particle with a slow particle of the cloud:

1. *High energy (HE).* The bullet hits the other particle 'head on', transferring most of its energy such that the bullet becomes slow and stays in the cloud, and the target particle becomes a bullet and leaves the cloud. In this case, the original bullet stays behind and heats up the cloud with the excess energy it has left.

2. *Knockout (KO).* The bullet transfers some of its energy to the target particle, turning it effectively into a bullet while itself remaining ballistic in nature. In this case, both the bullet and the target leave the cloud.

3. *Low energy (LE).* The bullet grazes the target particle transferring some energy, but not enough to knock it out. In this case, the target stays in the cloud and heats it up.

For a very fast bullet and an essentially stationary target of equal mass, $P_{\mathrm{LE}} \ll 1$, $P_{\mathrm{HE}} \ll 1$ and $P_{\mathrm{KO}} \approx 1$. This allows us to neglect terms of $O(P_{\mathrm{coll}}P_{\mathrm{HE}})$ and obtain the particle and energy loss rates

$$\dot{N}_{\mathrm{ME}} \approx -4 \int n(\mathbf{r}_0) P_{\mathrm{3B}} P_{\mathrm{coll}} \, \mathrm{d}^3\mathbf{r}_0 \,, \tag{B.10}$$

$$\dot{E}_{\mathrm{ME}} \approx -\int n(\mathbf{r}_0) P_{\mathrm{3B}} P_{\mathrm{coll}} \left(3\bar{E}(\mathbf{r}_0) + \bar{E}_{\mathrm{coll}}(\mathbf{r}_0)\right) \mathrm{d}^3\mathbf{r}_0 \,, \tag{B.11}$$

where $\bar{E}_{\mathrm{coll}}(\mathbf{r}_0)$ is the average cloud particle energy at the point of the secondary collision, subject to the three-body event taking place at $\mathbf{r}_0$. Note that the factor of four in Eq. (B.10) accounts for the fact that four particles leave the trap altogether in this case (the original three plus the one involved in the secondary collision), whereas the factors in Eq. (B.11) reflect the fact that the energy $3\bar{E}(\mathbf{r}_0)$ is lost via the three-body collision and $\bar{E}_{\mathrm{coll}}(\mathbf{r}_0)$ via the secondary collision. $\bar{E}_{\mathrm{coll}}(\mathbf{r}_0)$ can be obtained by integrating the probability of colliding in a small time interval $[t, t + \mathrm{d}t]$ and using the average particle energy at position $\mathbf{r}(t)$, leading to

$$\begin{aligned} \bar{E}_{\mathrm{coll}}(\mathbf{r}_0) &= \int_0^\infty n(t)\sigma v e^{-\sigma v \int_0^t n(t')\,\mathrm{d}t'} \left(U(\mathbf{r}(t)) + \frac{3}{2}k_B T\right) \mathrm{d}t \\ &= \underbrace{U(\mathbf{r}_0) + \frac{3}{2}k_B T}_{E(\mathbf{r}_0)} + \frac{1}{2}k_B T - \tilde{U}(\mathbf{r}_0) \,. \end{aligned} \tag{B.12}$$

Similar expressions hold for the case of particle escape and molecular collision (PE), with P_{coll} replaced by $\tilde{P}_{\mathrm{coll}}$. Combining the three contributions to particle and energy loss, i.e. $\dot{N} = \dot{N}_{\mathrm{DE}} + \dot{N}_{\mathrm{ME}} + \dot{N}_{\mathrm{PE}}$ and $\dot{E} = \dot{E}_{\mathrm{DE}} + \dot{E}_{\mathrm{ME}} + \dot{E}_{\mathrm{PE}}$, and evaluating the integrals finally give

$$\dot{N} = \dot{N}_0 \left(1 + \frac{3}{4\pi}\frac{mN}{k_B T}\frac{\bar{\omega}^6}{\bar{\Omega}^3}\frac{1}{\sqrt{\sum_i \omega_i^2}}(\sigma + \tilde{\sigma})\right) \equiv \dot{N}_0 \left(1 + \Theta\frac{N}{T}\right), \tag{B.13}$$

$$\dot{T} = \dot{T}_0 \left(1 + \Theta\frac{N}{T}\left(3\sum_i \frac{\omega_i^2}{\Omega_i^2} - 2\right)\right), \tag{B.14}$$

where $\bar{\Omega}$ is the geometric mean of $\Omega_i = \sqrt{4\omega_i^2 - \omega_i^4/\sum_j \omega_j^2}$. Therefore, the geometrical

factors can be expressed as

$$G = \frac{3}{4\pi}\frac{\bar{\omega}^4}{\bar{\Omega}^3}\frac{1}{\sqrt{\sum_i \omega_i^2}}, \tag{B.15}$$

$$G' = \frac{3}{4\pi}\frac{\bar{\omega}^4}{\bar{\Omega}^3}\frac{1}{\sqrt{\sum_i \omega_i^2}}\left(3\sum_i \frac{\omega_i^2}{\Omega_i^2} - 2\right). \tag{B.16}$$

B.2 Three-body loss in a Bose–Einstein condensate

In a BEC, due to the quantum-correlated nature of the gas, the three-body loss rate acquires a factor of $1/3!$ accounting for the fact that density fluctuations are suppressed in a quantum gas [258] (the three-body process occurs between three identical bosons in the same single-particle state), leading to

$$\dot{n}(\mathbf{r}) = -\frac{1}{\tau_1}n(\mathbf{r}) - \frac{1}{3!}L_3 n^3(\mathbf{r}). \tag{B.17}$$

In the Thomas–Fermi (TF) approximation, the density distribution in a harmonic trap is given by

$$n(\mathbf{r}) = \frac{15N}{8\pi R_x R_y R_z}\left(1 - \frac{x^2}{R_x^2} - \frac{y^2}{R_y^2} - \frac{z^2}{R_z^2}\right), \tag{B.18}$$

where R_x, R_y and R_z are the respective Thomas–Fermi radii [64, 187, 316, 317]. Substituting this to Eq. (B.17) leads to

$$\dot{N} = -\frac{1}{\tau_1}N - \frac{1}{3!}L_3\frac{75}{56\pi^2 R_x^2 R_y^2 R_z^2}N^3. \tag{B.19}$$

It can be shown that $R_x R_y R_z = \alpha N^{3/5}$, where α is a scaling parameter which can be determined by obtaining a self-consistent solution to the TF equations. Therefore, Eq. (B.19) can be rewritten more compactly as

$$\dot{N} = -\frac{1}{\tau_1}N - L_3\frac{25}{112\pi^2\alpha^2}N^{9/5}. \tag{B.20}$$

Furthermore, if one neglects one-body loss altogether, which is often valid as the high density of BECs leads to a more significant three-body loss, Eq. (B.20) can be solved analytically, leading to

$$N^{-\frac{4}{5}}(t) = N^{-\frac{4}{5}}(0) + \frac{5}{28\pi^2\alpha^2}L_3 t. \tag{B.21}$$

C

Discrete transforms and their efficient calculation

Given we run all our simulations on a computer with finite memory and computational power, we need to represent our wave functions on a discrete and finite grid. This means we can not use continuous Fourier (FT) and Hankel transforms (HT) or continuous integrals but we need to establish their discrete approximations.

C.1 Discrete Fourier transform

Let a function $f(x)$ be sampled on $[-L/2, L/2]$ at N evenly distributed points x_n, with

$$x_n = -\frac{L}{2} + \frac{nL}{N} \tag{C.1}$$

for $n \in \{0, \dots, N-1\}$, and let us define $f_n = f(x_n)$. Let us assume $f(x) = 0$ if $|x| > L/2$. In this case, the Fourier transform can be written as

$$\mathcal{F}\{f\}(k) = \int_{-\frac{L}{2}}^{\frac{L}{2}} f(x) e^{-ikx} \, \mathrm{d}x \approx \sum_{n=0}^{N-1} f_n e^{-ikx_n} \frac{L}{N} \, . \tag{C.2}$$

The discrete Fourier transform (DFT) of a sequence of N (complex) numbers $\{f_n\}$, mapping it onto another sequence of N (complex) numbers $\{F_l\}$, is defined as

$$F_l = \mathrm{DFT}\{f\}_l = \sum_{n=0}^{N-1} f_n e^{-2\pi i \frac{nl}{N}} \, , \tag{C.3}$$

By comparing Eqs. (C.2) and (C.3), one can show

$$\mathcal{F}\{f\}\left(\frac{2\pi l}{L}\right) = \frac{L}{N}(-1)^l \, \mathrm{DFT}\{f\}_l \, . \tag{C.4}$$

One could then conclude the DFT samples the Fourier transform at $k_l = 2\pi l/L$. However, due to the arbitrariness of the phase in the exponent in Eq. (C.3) (i.e. one could add arbitrary amounts of 2π to it), we can also write

$$k_l = \begin{cases} \frac{2\pi}{L} l & \text{if } 0 \leq l < \frac{N}{2} \\ \frac{2\pi}{L}(l-N) & \text{if } \frac{N}{2} \leq l \leq N-1 \end{cases}. \tag{C.5}$$

Similarly, if the Fourier transform of the function is band-limited to $|k| < \pi N/L$ (i.e. it is zero outside this band) and is sampled on $[-\pi N/L, \pi N/L]$ at

$$k_l = -\frac{\pi N}{L} + \frac{2\pi l}{L} \tag{C.6}$$

for $l \in \{0, \ldots, N-1\}$, the inverse Fourier transform can be written as

$$\frac{1}{2\pi} \int_{-\frac{\pi N}{L}}^{\frac{\pi N}{L}} \mathcal{F}\{f\}(k) e^{ikx} \, \mathrm{d}k \approx \frac{1}{2\pi} \sum_{l=0}^{N-1} F_l e^{ik_l x} \frac{2\pi}{L}, \tag{C.7}$$

where $F_l = \mathcal{F}\{f\}(k_l)$. The inverse discrete Fourier transform (IDFT) is defined as

$$f_n = \mathrm{DFT}^{-1}\{F\}_n = \frac{1}{N} \sum_{l=0}^{N-1} F_l e^{2\pi i \frac{nl}{N}} = \frac{1}{N} \mathrm{DFT}\{F^*\}_n^*, \tag{C.8}$$

where * represents complex conjugation. By comparing Eqs. (C.7) and (C.8) one can show

$$f\left(\frac{nL}{N}\right) = \frac{N}{L}(-1)^n \, \mathrm{DFT}^{-1}\{F\}_n \, . \tag{C.9}$$

Again, we could conclude the IDFT samples the function at $x_n = nL/N$, but by a similar argument as above, it can be shown this can also be written as

$$x_n = \begin{cases} \frac{L}{N} n & \text{if } 0 \leq n < \frac{N}{2} \\ \frac{L}{N}(n-N) & \text{if } \frac{N}{2} \leq n \leq N-1 \end{cases}, \tag{C.10}$$

which is what we have written originally in Eq. (C.1).

It is worth noting that a function cannot be (band-)limited in both r and k space at the same time. If we assume the Fourier transform is zero outside the domain that is represented by the N k_l points, we assume the function we are sampling is periodic outside the (real-space) sampling domain. This aliasing causes a problem with the dipolar interaction as in this case one gets a contribution from the interaction between the alias clouds. To

avoid this along r, the Hankel transform is used in that direction—see Appendix C.4. Along z, a cutoff is introduced to the dipole interaction and a relatively large grid is used.

It is also not trivial to calculate derivatives for a similar reason. A recipe for developing an interpolation function for calculating derivatives can be found in Ref. 318.

C.2 Discrete cosine transform

The discrete cosine transform (DCT) is a faster way of calculating the DFT when the function sampled is of even parity. Let us assume that our data points $\{f_n\}$ represent a function sampled at $\{x_n\} \in [-L/2, 0]$, which is even around $x_0 = -L/2$ and $x_{N-1} = 0$. The DCT (DCT-I) is defined as

$$F_l = \mathrm{DCT}\{f\}_l = f_0 + (-1)^l f_{N-1} + 2\sum_{n=1}^{N-2} f_n \cos\left(\frac{\pi n l}{N-1}\right). \tag{C.11}$$

It can be shown that the DCT is exactly equivalent to the DFT of the even extension of the series $\{f_n\}$, $\{g_m\}$, defined as

$$g_m = \begin{cases} f_m & \text{if } m \in \{0, \ldots, N-1\} \\ g_{2N-2-m} & \text{if } m \in \{N, \ldots, 2N-3\} \end{cases}. \tag{C.12}$$

It can be seen that g_m represents the even extension of our function f into the interval $[-L/2, L/2]$, sampling it at $2N-2$ points. Therefore,

$$\mathcal{F}\{f\}\left(\frac{2\pi l}{L}\right) = \frac{L}{2N-2}(-1)^l \,\mathrm{DFT}\{g\}_l = \frac{L}{2N-2}(-1)^l \,\mathrm{DCT}\{f\}_l, \tag{C.13}$$

i.e. the Fourier transform of an even function can be calculated using only half of the sampling interval.

The inverse discrete cosine transform (IDCT) is defined as

$$f_m = \mathrm{DCT}^{-1}\{F\}_m = \frac{1}{2N-2}\,\mathrm{DCT}\{F\}_m. \tag{C.14}$$

Using the relationship between the IDFT and the DFT (Eq. (C.8)), and that our functions are real and even over z (and so the DFT is also real), we find

$$f_l = \mathrm{DFT}^{-1}\{G\}_l = \frac{1}{2N-2}\,\mathrm{DFT}\{G\}_l = \frac{1}{2N-2}\,\mathrm{DCT}\{F\}_l = \mathrm{DCT}^{-1}\{F\}_l, \tag{C.15}$$

where G_l is sampling the Fourier transform on $[-\pi(2N-2)/L, \pi(2N-2)/L]$ and F_l is sampling the Fourier transform on $[-\pi(2N-2)/L, 0]$. Therefore,

$$f\left(\frac{nL}{2N-2}\right) = \frac{2N-2}{L}(-1)^n \,\mathrm{DFT}^{-1}\{G\}_n = \frac{2N-2}{L}(-1)^n \,\mathrm{DCT}^{-1}\{F\}_n \,. \tag{C.16}$$

We also see

$$f\left(\frac{nL}{2N-2}\right) = \frac{1}{L}(-1)^n \,\mathrm{DCT}\{F\}_n \tag{C.17}$$

by definition.

C.3 Discrete sine transform

The discrete sine transform (DST), similarly to the DCT, is a faster way of calculating the DFT when the function sampled is of odd parity. Let us again assume that our data points $\{f_n\}$ represent a function sampled at $\{x_n\} \in [-L/2, 0]$, which is odd around $x_0 = -L/2$ and $x_{N-1} = 0$ (i.e. $f_0 = f_{N-1} = 0$). The DST (DST-I) is defined as

$$F_l = \mathrm{DST}\{f\}_l = 2\sum_{n=1}^{N-2} f_n \sin\left(\frac{\pi n(l+1)}{N-1}\right). \tag{C.18}$$

It can be shown that $-i$ times the DST is exactly equivalent to the DFT of the odd extension of the series $\{f_n\}$, $\{g_m\}$, defined as

$$g_m = \begin{cases} f_m & \text{if } m \in \{0, \ldots, N-1\} \\ -g_{2N-2-m} & \text{if } m \in \{N, \ldots, 2N-3\} \end{cases}. \tag{C.19}$$

It can be seen that g_m represents the odd extension of our function f into the interval $[-L/2, L/2]$, sampling it at $2N-2$ points. Therefore,

$$\mathcal{F}\{f\}\left(\frac{2\pi l}{L}\right) = \frac{L}{2N-2}(-1)^l \,\mathrm{DFT}\{g\}_l = -\frac{iL}{2N-2}(-1)^l \,\mathrm{DST}\{f\}_l \,, \tag{C.20}$$

i.e. the Fourier transform of an odd function can be calculated using only half of the sampling interval.

The inverse discrete sine transform (IDST) is defined as

$$f_m = \mathrm{DST}^{-1}\{F\}_m = \frac{1}{2N-2}\,\mathrm{DST}\{F\}_m \,. \tag{C.21}$$

Using the relationship between the IDFT and the DFT (Eq. (C.8)) again and that our functions are real and odd over z (and so the DFT is purely imaginary), we find

$$f_l = \mathrm{DFT}^{-1}\{G\}_l = -\frac{1}{2N-2}\,\mathrm{DFT}\{G\}_l = \frac{i}{2N-2}\,\mathrm{DST}\{F\}_l = i\,\mathrm{DST}^{-1}\{F\}_l\,. \tag{C.22}$$

where G_l is sampling the Fourier transform on $[-\pi(2N-2)/L, \pi(2N-2)/L]$ and F_l is sampling the Fourier transform on $[-\pi(2N-2)/L, 0]$. Therefore,

$$f\left(\frac{nL}{2N-2}\right) = \frac{2N-2}{L}(-1)^n\,\mathrm{DFT}^{-1}\{G\}_n = i\frac{2N-2}{L}(-1)^n\,\mathrm{DST}^{-1}\{F\}_n\,. \tag{C.23}$$

We also see

$$f\left(\frac{nL}{2N-2}\right) = \frac{i}{L}(-1)^n\,\mathrm{DST}\{F\}_n \tag{C.24}$$

by definition.

C.4 Discrete Hankel transform

C.4.1 0^{th}-order Hankel transform

An efficient numerical approach to calculate the 0^{th} order Hankel transform using discrete data, which relies on the Dini series expansion, is laid out in Ref. 297. We will summarise its basic principles and give a direct recipe below.

Let us consider a function $f(r)$ where $f(r) = 0$ if $r > b$ for some b and $f'(b) = 0$. Let us also assume that the 0^{th} order Hankel transform of this function is band-limited, i.e. $\mathcal{H}_0\{f\}(k) = 0$ if $|k| > \beta$ for some β. The Dini series expansion of this function and its Hankel transform can be given as

$$f(r) = \frac{2}{b^2}\sum_{n=0}^{\infty}\mathcal{H}_0\{f\}\left(\frac{\alpha_n}{b}\right)J_0^{-2}(\alpha_n)\,J_0\left(\frac{\alpha_n r}{b}\right), \tag{C.25}$$

$$\mathcal{H}_0\{f\}(k) = \frac{2}{\beta^2}\sum_{n=0}^{\infty}f\left(\frac{\alpha_n}{\beta}\right)J_0^{-2}(\alpha_n)\,J_0\left(\frac{\alpha_n k}{\beta}\right), \tag{C.26}$$

where J_0 is the 0^{th} order Bessel function of the first kind and α_n are the real non-negative roots of the first derivate of J_0 (i.e. the non-negative roots of the 1^{st} order Bessel function

of the first kind, J_1) with $\alpha_0 = 0$. The quantities

$$G_m = \mathcal{H}_0\{f\}\left(\frac{\alpha_m}{b}\right)\left|J_0^{-1}(\alpha_m)\right|\beta\,, \tag{C.27}$$

$$F_n = f\left(\frac{\alpha_n}{\beta}\right)\left|J_0^{-1}(\alpha_n)\right|b \tag{C.28}$$

can be expressed as

$$G_m = \sum_{n=0}^{N} C_{mn}F_n\,, \tag{C.29}$$

$$F_n = \sum_{m=0}^{N} C_{nm}G_m\,, \tag{C.30}$$

where C_{nm} are elements of a transformation matrix C given by

$$C_{nm} = \frac{2}{S}\left|J_0^{-1}(\alpha_n)\right|\left|J_0^{-1}(\alpha_m)\right| J_0\left(\frac{\alpha_n\alpha_m}{S}\right), \tag{C.31}$$

where $S = \beta b$ and N is such that $\alpha_N = S$. To make the method self-consistent, one needs to ensure the transformation matrix C is orthogonal. This is well-achieved by choosing

$$S = 2\left|J_0^{-1}(\alpha_k)\right|\sqrt{1 + \sum_{n=1}^{N} J_0^{-2}(\alpha_n)\, J_0^2\left(\frac{\alpha_\kappa\alpha_n}{J_{N+1}}\right)}, \tag{C.32}$$

where J_{N+1} is the $(N+1)^{\text{th}}$ zero of J_0 and $\kappa = \lfloor N/4 \rfloor$, where $\lfloor x \rfloor$ is the floor function (returning the largest integer smaller or equal to x).

Practically, this means that given the values of $f(\alpha_n/\beta)$, we can calculate the Hankel transform $\mathcal{H}_0\{f\}(\alpha_m/b)$ and vice versa. Therefore, to be able to calculate the Hankel transform this way, we need to sample our function on a large enough grid covering $[0, b]$ such that we can assume our function is 0 outside the grid. We then need to choose the number of grid points, $N + 1$, from which we can calculate S. We need to make sure that N is large enough so that the Hankel transform is indeed band-limited by $\beta = S/b$. Once this is done, these directly define the location of our grid points, which in real space will be at $r_n = \alpha_n/\beta = \alpha_n b/S$ and in frequency space will be at $k_n = \alpha_n/b$ with $n = 0, 1, \ldots, N$. It is worth noting that $r_0 = k_0 = 0$, i.e. our function will be sampled at the origin this way.

With these, the procedure for calculating the Hankel transform and its inverse can be

given as

$$\mathcal{H}_0\{f\}(k_m) = \frac{|J_0(\alpha_m)|}{\beta} \sum_{n=0}^{N} C_{mn} f(r_n) \left|J_0^{-1}(\alpha_n)\right| b\,, \tag{C.33}$$

$$f(r_n) = \frac{|J_0(\alpha_n)|}{b} \sum_{m=0}^{N} C_{nm} \mathcal{H}_0\{f\}(k_m) \left|J_0^{-1}(\alpha_m)\right| \beta\,. \tag{C.34}$$

Therefore, calculating the Hankel transform is a matter of two multiplications and a matrix multiplication.

C.4.2 Hankel transform of order $m > 0$

The above is all very useful, but $\mathcal{H}_0$ is only relevant for calculating the Fourier transform of a function which has no angular (θ) dependence. For our general case, where the functions take the form $f(r)e^{-im\theta}$, a slightly different approach is needed [282, 298].

Let us assume again the function is limited to $r < b$ and its Hankel transform of order m is band-limited to $|k| < \beta$. Furthermore, let us sample the function at N points $r_n = \alpha_{mn}/\beta$ and its Hankel transform at $k_n = \alpha_{mn}/b$ with $n = 1, \ldots, N$. Here α_{mn} is the n^{th} root of J_m, the Bessel function of order m of the first kind. The function and its Hankel transform can be expanded as a Fourier–Bessel series,

$$f(r) = \frac{2}{b^2} \sum_{n=1}^{\infty} \mathcal{H}_m\{f\}\left(\frac{\alpha_{mn}}{b}\right) J_{m+1}^{-2}(\alpha_{mn})\, J_m\left(\frac{\alpha_{mn} r}{b}\right), \tag{C.35}$$

$$\mathcal{H}_m\{f\}(k) = \frac{2}{\beta^2} \sum_{n=1}^{\infty} f\left(\frac{\alpha_{mn}}{\beta}\right) J_{m+1}^{-2}(\alpha_{mn})\, J_m\left(\frac{\alpha_{mn} k}{\beta}\right). \tag{C.36}$$

Again, the quantities

$$G_l = \mathcal{H}_m\{f\}\left(\frac{\alpha_{ml}}{b}\right) \left|J_{m+1}^{-1}(\alpha_{ml})\right| \beta\,, \tag{C.37}$$

$$F_j = f\left(\frac{\alpha_{mj}}{\beta}\right) \left|J_{m+1}^{-1}(\alpha_{mj})\right| b\,, \tag{C.38}$$

are defined, which can be expressed as

$$G_l = \sum_{j=1}^{N} C_{lj} F_j\,, \tag{C.39}$$

$$F_j = \sum_{l=1}^{N} C_{jl} G_l\,, \tag{C.40}$$

where the transformation matrix C is now given by

$$C_{lj} = \frac{2}{S}\left|J_{m+1}^{-1}(\alpha_{ml})\right|\left|J_{m+1}^{-1}(\alpha_{mj})\right| J_m\left(\frac{\alpha_{ml}\alpha_{mj}}{S}\right). \tag{C.41}$$

Here, $S = b\beta$ again and N is such that $S = \alpha_{m,N+1}$. To make the method self-consistent, we again need to make sure C is orthogonal. This is best achieved by choosing $S = \alpha_{m,N+1}$. Therefore, the method is similar as before: decide on b and N, then calculate S and $\beta = S/b$, which defines the grid on which we can calculate the (inverse) Hankel transform. Note that in this case $r_1 > 0$, so the origin is not sampled. Therefore, there is an advantage in using a different method for $m = 0$, which samples $r = 0$, as we expect some of our ground states to have their peak density at the origin, and so to be most sensitive to excitations there.

With these, the procedure for calculating the Hankel transform and its inverse can be given as

$$\mathcal{H}_m\{f\}(k_l) = \frac{|J_{m+1}(\alpha_{ml})|}{\beta}\sum_{j=1}^{N} C_{lj} f(r_j)\left|J_{m+1}^{-1}(\alpha_{mj})\right| b\,, \tag{C.42}$$

$$f(r_j) = \frac{\left|J_{m+1}(\alpha_{mj})\right|}{b}\sum_{l=1}^{N} C_{jl}\mathcal{H}_m\{f\}(k_l)\left|J_{m+1}^{-1}(\alpha_{ml})\right|\beta\,. \tag{C.43}$$

Therefore, calculating the Hankel transform is again a matter of two multiplications and a matrix multiplication.

C.5 Integration

Sometimes we need to calculate integrals of our functions to calculate e.g. the energy. For the case of no dependence on θ, one can observe

$$I_r = \int_0^\infty f(r) 2\pi r\,\mathrm{d}r = 2\pi\,\mathcal{H}_0\{f\}(0)\,. \tag{C.44}$$

Using a Dini series expansion one can write

$$\mathcal{H}_0\{f\}(0) = \frac{2}{\beta^2}\sum_{n=0}^{\infty} f\left(\frac{\alpha_n}{\beta}\right) J_0^{-2}(\alpha_n) \tag{C.45}$$

and so deduce

$$I_r = \frac{4\pi}{\beta^2}\sum_{n=0}^{N} f\left(\frac{\alpha_n}{\beta}\right) J_0^{-2}(\alpha_n)\,, \tag{C.46}$$

given we can truncate the sum as $f(r) = 0$ for $r > \alpha_N/\beta$.

In some cases, e.g. when calculating dot products, we also need to calculate the integral of a product of two functions:

$$I_{r,2} = \int_0^\infty f(r)g(r)2\pi r \, \mathrm{d}r = 2\pi \, \mathcal{H}_0\{fg\}(0) \, . \tag{C.47}$$

Forming the Dini series for the product fg, one can immediately conclude

$$I_{r,2} = \frac{4\pi}{\beta^2} \sum_{n=0}^{N} f\left(\frac{\alpha_n}{\beta}\right) g\left(\frac{\alpha_n}{\beta}\right) J_0^{-2}(\alpha_n) \, , \tag{C.48}$$

as the functions f and g are represented on the same grid and therefore come with the same assumptions.

We can also calculate integrals in k-space. Similarly as before, one can observe

$$I_k = \int_0^\infty \mathcal{H}_0\{f\}(k_r) \, 2\pi k_r \, \frac{\mathrm{d}k_r}{(2\pi)^2} = \frac{f(0)}{2\pi} \, . \tag{C.49}$$

Using a Dini series expansion one can write

$$f(0) = \frac{2}{b^2} \sum_{n=0}^{\infty} \mathcal{H}_0\{f\}\left(\frac{\alpha_n}{b}\right) J_0^{-2}(\alpha_n) \tag{C.50}$$

and so deduce

$$I_k = \frac{1}{\pi b^2} \sum_{n=0}^{N} \mathcal{H}_0\{f\}\left(\frac{\alpha_n}{b}\right) J_0^{-2}(\alpha_n) \, , \tag{C.51}$$

given we can truncate the sum as $\mathcal{H}_0\{f\}(k) = 0$ for $k > \alpha_N/b$.

Integration along z can be done more simply (and less accurately) using a Riemann sum, as the grid is linear in that direction. As the function is only represented at $z \leq 0$, we need to take into account the symmetry of the function and take especial care not to add the $z = 0$ element twice to the integral.

C.6 Interpolation

We can also define an analytic interpolation function for a function represented on this grid. Writing the band-limited $f(r)$ as an inverse Hankel transform,

$$f(r) = \int_0^\beta \mathcal{H}_0\{f\}(k) \, J_0(k_r r) k_r \, \mathrm{d}k_r \, , \tag{C.52}$$

and substituting the Dini series expansion of $\mathcal{H}_0\{f\}(k)$ (Eq. (C.26)) into the integral, one can show

$$f(r) = 2r\beta J_1(r\beta) \sum_{n=0}^{N} \frac{1}{r^2\beta^2 - \alpha_n^2} J_0^{-1}(\alpha_n) f\left(\frac{\alpha_n}{\beta}\right). \tag{C.53}$$

This means we can calculate $f(r)$ at any r if the function is defined at points $r_n = \alpha_n/\beta$. The interpolation of $\mathcal{H}_0\{f\}(k)$ could also be done using a similar approach, and one can show

$$\mathcal{H}_0\{f\}(k) = 2bkJ_1(bk) \sum_{n=0}^{N} \frac{1}{b^2k^2 - \alpha_n^2} J_0^{-1}(\alpha_n) \mathcal{H}_0\{f\}\left(\frac{\alpha_n}{b}\right). \tag{C.54}$$

Special care needs to be taken for $r = \alpha_n/\beta$ ($k = \alpha_n/b$), when these formulae are not well-defined, but in that case $r = r_n$ ($k = k_n$) and the known $f(r_n)$ ($\mathcal{H}_0\{f\}(k_n)$) can be directly used.

This interpolation technique is used for interpolating a function represented on one grid to be represented on another. For example, calculating $m > 0$ excitations using the BdG equation requires the calculation of higher-order Hankel transforms, but the grid used along r changes with m so interpolation is needed.

Interpolation along z can be done more simply (and less accurately) using quadratic splines, as the grid is linear in that direction. This was done using the SCIPY package [259].

References

[1] The drawing of the arms of the University is used with permission from Oxford University Press, the arms of the University of Cambridge belongs to that university. All other drawings have been adapted from Wikipedia, some in the public domain and some under a Creative Commons licence. The original creators of those drawings are 'Thommy' (Hungary), 'Flanker' (Holy See), Mathieu Chaine (Order of Malta), 'ChevronTango' (Pembroke College), 'Sodacan' (Royal Arms) and 'Richie' (Trinity College).

[2] S. N. Bose, 'Plancks Gesetz und Lichtquantenhypothese', Z. Phys. **26**, 178 (1924).

[3] A. Einstein, 'Quantentheorie des einatomigen idealen Gases', Sitzungsber. Preuss. Akad. Wiss., 3 (1925); reprinted in D. Simon, ed., *Albert Einstein: Akademie-Vorträge* (Wiley, Weinheim, 2005), p. 237.

[4] A. J. Leggett, *Quantum Liquids: Bose condensation and Cooper pairing in condensed-matter systems* (Oxford University Press, Oxford, 2006).

[5] F. Schreck and K. van Druten, 'Laser cooling for quantum gases', Nat. Phys. **17**, 1296 (2021).

[6] S. Chu, L. Hollberg, J. E. Bjorkholm, A. Cable and A. Ashkin, 'Three-dimensional viscous confinement and cooling of atoms by resonance radiation pressure', Phys. Rev. Lett. **55**, 48 (1985).

[7] W. D. Phillips, 'Nobel Lecture: Laser cooling and trapping of neutral atoms', Rev. Mod. Phys. **70**, 721 (1998).

[8] 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).

[9] 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).

[10] C. C. Bradley, C. A. Sackett, J. J. Tollett and R. G. Hulet, 'Evidence of Bose-Einstein Condensation in an Atomic Gas with Attractive Interactions', Phys. Rev. Lett. **75**, 1687 (1995); 'Erratum: Evidence of Bose-Einstein Condensation in an Atomic Gas with Attractive Interactions [Phys. Rev. Lett. 75, 1687 (1995)]', Phys. Rev. Lett. **79**, 1170 (1997).

[11] B. DeMarco and D. S. Jin, 'Onset of Fermi Degeneracy in a Trapped Atomic Gas', Science **285**, 1703 (1999).

[12] D. G. Fried, T. C. Killian, L. Willmann, D. Landhuis, S. C. Moss, D. Kleppner and T. J. Greytak, 'Bose-Einstein Condensation of Atomic Hydrogen', Phys. Rev. Lett. **81**, 3811 (1998).

[13] 'The Nobel Prize in Physics 1997', NobelPrize.org, Press Release (1997).

[14] 'The Nobel Prize in Physics 2001', NobelPrize.org, Press Release (2001).

[15] A. D. Cronin, J. Schmiedmayer and D. E. Pritchard, 'Optics and interferometry with atoms and molecules', Rev. Mod. Phys. **81**, 1051 (2009).

[16] C. Gross, T. Zibold, E. Nicklas, J. Estève and M. K. Oberthaler, 'Nonlinear atom interferometer surpasses classical precision limit', Nature **464**, 1165 (2010).

[17] L. Amico, D. Anderson, M. Boshier, J.-P. Brantut, L.-C. Kwek, A. Minguzzi and W. von Klitzing, '*Colloquium*: Atomtronic circuits: From many-body physics to quantum technologies', Rev. Mod. Phys. **94**, 041001 (2022).

[18] M. Planck, *Vorlesung über die Theorie der Wärmestrahlung* (Verlag von J. A. Barth, Leipzig, 1906).

[19] A. Einstein, 'Über einen die Erzeugung und Verwandlung des Lichtes betreffenden heuristischen Gesichtspunkt', Ann. Phys. **322**, 132 (1905).

[20] R. P. Feynman, 'Simulating physics with computers', Int. J. Theor. Phys. **21**, 467 (1982).

[21] A. Aspuru-Guzik and P. Walther, 'Photonic quantum simulators', Nat. Phys **8**, 285 (2012).

[22] R. Blatt and C. F. Roos, 'Quantum simulations with trapped ions', Nat. Phys. **8**, 277 (2012).

[23] I. Bloch, J. Dalibard and S. Nascimbène, 'Quantum simulations with ultracold quantum gases', Nat. Phys. **8**, 267 (2012).

[24] I. Bloch, J. Dalibard and W. Zwerger, 'Many-body physics with ultracold gases', Rev. Mod. Phys. **80**, 885 (2008).

[25] M. R. Andrews, C. G. Townsend, H.-J. Miesner, D. S. Durfee, D. M. Kurn and W. Ketterle, 'Observation of Interference Between Two Bose Condensates', Science **275**, 637 (1997).

[26] I. Bloch, T. W. Hänsch and T. Esslinger, 'Measurement of the spatial coherence of a trapped Bose gas at the phase transition', Nature **403**, 166 (2000).

[27] M. Greiner, C. A. Regal and D. S. Jin, 'Emergence of a molecular Bose–Einstein condensate from a Fermi gas', Nature **426**, 537 (2003).

[28] M. W. Zwierlein, C. A. Stan, C. H. Schunck, S. M. F. Raupach, S. Gupta, Z. Hadzibabic and W. Ketterle, 'Observation of Bose-Einstein Condensation of Molecules', Phys. Rev. Lett. **91**, 250401 (2003).

[29] S. Jochim, M. Bartenstein, A. Altmeyer, G. Hendl, S. Riedl, C. Chin, J. Hecker Denschlag and R. Grimm, 'Bose-Einstein Condensation of Molecules', Science **302**, 2101 (2003).

[30] M. R. Matthews, B. P. Anderson, P. C. Haljan, D. S. Hall, C. E. Wieman and E. A. Cornell, 'Vortices in a Bose-Einstein Condensate', Phys. Rev. Lett. **83**, 2498 (1999).

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

[32] E. Hodby, G. Hechenblaikner, S. A. Hopkins, O. M. Maragò and C. J. Foot, 'Vortex Nucleation in Bose-Einstein Condensates in an Oblate, Purely Magnetic Potential', Phys. Rev. Lett. **88**, 010405 (2001).

[33] J. R. Abo-Shaeer, C. Raman, J. M. Vogels and W. Ketterle, 'Observation of Vortex Lattices in Bose-Einstein Condensates', Science **292**, 476 (2001).

[34] O. M. Maragò, S. A. Hopkins, J. Arlt, E. Hodby, G. Hechenblaikner and C. J. Foot, 'Observation of the Scissors Mode and Evidence for Superfluidity of a Trapped Bose-Einstein Condensed Gas', Phys. Rev. Lett. **84**, 2056 (2000).

[35] G. Hechenblaikner, E. Hodby, S. A. Hopkins, O. M. Maragò and C. J. Foot, 'Direct Observation of Irrotational Flow and Evidence of Superfluidity in a Rotating Bose-Einstein Condensate', Phys. Rev. Lett. **88**, 070406 (2002).

[36] 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).

[37] T. Langen, R. Geiger and J. Schmiedmayer, 'Ultracold Atoms Out of Equilibrium', Annu. Rev. Condens. Matter Phys. **6**, 201 (2015).

[38] M. Lewenstein, A. Sanpera, V. Ahufinger, B. Damski, A. Sen(De) and U. Sen, 'Ultracold atomic gases in optical lattices: mimicking condensed matter physics and beyond', Adv. Phys. **56**, 243 (2007).

[39] J. Beugnon and N. Navon, 'Exploring the Kibble–Zurek mechanism with homogeneous Bose gases', J. Phys. B: At. Mol. Opt. Phys. **50**, 022002 (2017).

[40] J. Ye, S. Blatt, M. M. Boyd, S. M. Foreman, E. R. Hudson, T. Ido, B. Lev, A. D. Ludlow, B. C. Sawyer, B. Stuhl and T. Zelevinsky, 'Precision Measurement Based on Ultracold Atoms and Cold Molecules', AIP Conf. Proc. **869**, 80 (2006).

[41] X. Zhang and J. Ye, 'Precision measurement and frequency metrology with ultracold atoms', Natl. Sci. Rev. **3**, 189 (2016).

[42] I. Bloch, 'Quantum coherence and entanglement with ultracold atoms in optical lattices', Nature **453**, 1016 (2008).

[43] R. Grimm, M. Weidemüller and Y. B. Ovchinnikov, 'Optical Dipole Traps for Neutral Atoms', in *Adv. At., Mol., Opt. Phys.* Vol. 42 (Academic Press, San Diego, 2000), p. 95.

[44] W. Petrich, M. H. Anderson, J. R. Ensher and E. A. Cornell, 'Stable, Tightly Confining Magnetic Trap for Evaporative Cooling of Neutral Atoms', Phys. Rev. Lett. **74**, 3352 (1995).

[45] 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).

[46] P. Juhász, M. Krstajić, D. Strachan, E. Gandar and R. P. Smith, 'How to realize a homogeneous dipolar Bose gas in the roton regime', Phys. Rev. A **105**, L061301 (2022).

[47] N. Navon, R. P. Smith and Z. Hadzibabic, 'Quantum gases in optical boxes', Nat. Phys. **17**, 1334 (2021).

[48] N. Navon, A. L. Gaunt, R. P. Smith and Z. Hadzibabic, 'Emergence of a turbulent cascade in a quantum gas', Nature **539**, 72 (2016).

[49] N. Navon, A. L. Gaunt, R. P. Smith and Z. Hadzibabic, 'Critical dynamics of spontaneous symmetry breaking in a homogeneous Bose gas', Science **347**, 167 (2015).

[50] D. Jaksch, C. Bruder, J. I. Cirac, C. W. Gardiner and P. Zoller, 'Cold Bosonic Atoms in Optical Lattices', Phys. Rev. Lett. **81**, 3108 (1998).

[51] A. Mazurenko, C. S. Chiu, G. Ji, M. F. Parsons, M. Kanász-Nagy, R. Schmidt, F. Grusdt, E. Demler, D. Greif and M. Greiner, 'A cold-atom Fermi–Hubbard antiferromagnet', Nature **545**, 462 (2017).

[52] C. Gross and W. S. Bakr, 'Quantum gas microscopy for single atom and spin detection', Nat. Phys. **17**, 1316 (2021).

[53] B. Paredes, A. Widera, V. Murg, O. Mandel, S. Fölling, I. Cirac, G. V. Shlyapnikov, T. W. Hänsch and I. Bloch, 'Tonks–Girardeau gas of ultracold atoms in an optical lattice', Nature **429**, 277 (2004).

[54] T. Kinoshita, T. Wenger and D. S. Weiss, 'Observation of a One-Dimensional Tonks-Girardeau Gas', Science **305**, 1125 (2004).

[55] Z. Hadzibabic, P. Krüger, M. Cheneau, B. Battelier and J. Dalibard, 'Berezinskii–Kosterlitz–Thouless crossover in a trapped atomic gas', Nature **441**, 1118 (2006).

[56] S. Inouye, M. R. Andrews, J. Stenger, H.-J. Miesner, D. M. Stamper-Kurn and W. Ketterle, 'Observation of Feshbach resonances in a Bose–Einstein condensate', Nature **392**, 151 (1998).

[57] Ph. Courteille, R. S. Freeland, D. J. Heinzen, F. A. van Abeelen and B. J. Verhaar, 'Observation of a Feshbach Resonance in Cold Atom Scattering', Phys. Rev. Lett. **81**, 69 (1998).

[58] C. Chin, R. Grimm, P. Julienne and E. Tiesinga, 'Feshbach resonances in ultracold gases', Rev. Mod. Phys. **82**, 1225 (2010).

[59] A. Griesmaier, J. Werner, S. Hensler, J. Stuhler and T. Pfau, 'Bose-Einstein Condensation of Chromium', Phys. Rev. Lett. **94**, 160401 (2005).

[60] M. Lu, N. Q. Burdick, S. H. Youn and B. L. Lev, 'Strongly Dipolar Bose-Einstein Condensate of Dysprosium', Phys. Rev. Lett. **107**, 190401 (2011).

[61] K. Aikawa, A. Frisch, M. Mark, S. Baier, A. Rietzler, R. Grimm and F. Ferlaino, 'Bose-Einstein Condensation of Erbium', Phys. Rev. Lett. **108**, 210401 (2012).

[62] E. T. Davletov, V. V. Tsyganok, V. A. Khlebnikov, D. A. Pershin, D. V. Shaykin and A. V. Akimov, 'Machine learning for achieving Bose-Einstein condensation of thulium atoms', Phys. Rev. A **102**, 011302(R) (2020).

[63] Y. Miyazawa, R. Inoue, H. Matsui, G. Nomura and M. Kozuma, 'Bose-Einstein Condensation of Europium', Phys. Rev. Lett. **129**, 223401 (2022).

[64] T. Lahaye, C. Menotti, L. Santos, M. Lewenstein and T. Pfau, 'The physics of dipolar bosonic quantum gases', Rep. Prog. Phys. **72**, 126401 (2009).

[65] L. Chomaz, I. Ferrier-Barbut, F. Ferlaino, B. Laburthe-Tolra, B. L. Lev and T. Pfau, 'Dipolar physics: a review of experiments with magnetic quantum gases', Rep. Prog. Phys. **86**, 026401 (2023).

[66] T. Koch, T. Lahaye, J. Metz, B. Fröhlich, A. Griesmaier and T. Pfau, 'Stabilization of a purely dipolar quantum gas against collapse', Nat. Phys. **4**, 218 (2008).

[67] T. Lahaye, T. Koch, B. Fröhlich, M. Fattori, J. Metz, A. Griesmaier, S. Giovanazzi and T. Pfau, 'Strong dipolar effects in a quantum ferrofluid', Nature **448**, 672 (2007).

[68] T. Lahaye, J. Metz, B. Fröhlich, T. Koch, M. Meister, A. Griesmaier, T. Pfau, H. Saito, Y. Kawaguchi and M. Ueda, 'd-Wave Collapse and Explosion of a Dipolar Bose-Einstein Condensate', Phys. Rev. Lett. **101**, 080401 (2008).

[69] T. Maier, H. Kadau, M. Schmitt, M. Wenzel, I. Ferrier-Barbut, T. Pfau, A. Frisch, S. Baier, K. Aikawa, L. Chomaz, M. J. Mark, F. Ferlaino, C. Makrides, E. Tiesinga, A. Petrov et al., 'Emergence of Chaotic Scattering in Ultracold Er and Dy', Phys. Rev. X **5**, 041029 (2015).

[70] T. Maier, I. Ferrier-Barbut, H. Kadau, M. Schmitt, M. Wenzel, C. Wink, T. Pfau, K. Jachymski and P. S. Julienne, 'Broad universal Feshbach resonances in the chaotic spectrum of dysprosium atoms', Phys. Rev. A **92**, 060702(R) (2015).

[71] L. Chomaz, R. M. W. van Bijnen, D. Petter, G. Faraoni, S. Baier, J. H. Becher, M. J. Mark, F. Wächtler, L. Santos and F. Ferlaino, 'Observation of roton mode population in a dipolar quantum gas', Nat. Phys. **14**, 442 (2018).

[72] D. Petter, G. Natale, R. M. W. van Bijnen, A. Patscheider, M. J. Mark, L. Chomaz and F. Ferlaino, 'Probing the Roton Excitation Spectrum of a Stable Dipolar Bose Gas', Phys. Rev. Lett. **122**, 183401 (2019).

[73] L. D. Landau, 'The Theory of Superfluidity of Helium II', J. Phys. (USSR) **5**, 71 (1941); reprinted in D. ter Haar, ed., *Collected Papers of L.D. Landau* (Pergamon Press, Oxford, 1965), p. 301.

[74] L. D. Landau, 'On the Theory of Superfluidity of Helium II', J. Phys. (USSR) **11**, 91 (1947); reprinted in D. ter Haar, ed., *Collected Papers of L.D. Landau* (Pergamon Press, Oxford, 1965), p. 466.

[75] D. G. Henshaw and A. D. B. Woods, 'Modes of Atomic Motions in Liquid Helium by Inelastic Scattering of Neutrons', Phys. Rev. **121**, 1266 (1961).

[76] H. Kadau, M. Schmitt, M. Wenzel, C. Wink, T. Maier, I. Ferrier-Barbut and T. Pfau, 'Observing the Rosensweig instability of a quantum ferrofluid', Nature **530**, 194 (2016).

[77] I. Ferrier-Barbut, H. Kadau, M. Schmitt, M. Wenzel and T. Pfau, 'Observation of Quantum Droplets in a Strongly Dipolar Bose Gas', Phys. Rev. Lett. **116**, 215301 (2016).

[78] J. Léonard, A. Morales, P. Zupancic, T. Esslinger and T. Donner, 'Supersolid formation in a quantum gas breaking a continuous translational symmetry', Nature **543**, 87 (2017).

[79] J.-R. Li, J. Lee, W. Huang, S. Burchesky, B. Shteynas, F. Ç. Top, A. O. Jamison and W. Ketterle, 'A stripe phase with supersolid properties in spin–orbit-coupled Bose–Einstein condensates', Nature **543**, 91 (2017).

[80] G. Natale, R. M. W. van Bijnen, A. Patscheider, D. Petter, M. J. Mark, L. Chomaz and F. Ferlaino, 'Excitation Spectrum of a Trapped Dipolar Supersolid and Its Experimental Evidence', Phys. Rev. Lett. **123**, 050402 (2019).

[81] F. Böttcher, J.-N. Schmidt, M. Wenzel, J. Hertkorn, M. Guo, T. Langen and T. Pfau, 'Transient Supersolid Properties in an Array of Dipolar Quantum Droplets', Phys. Rev. X **9**, 011051 (2019).

[82] J. Hertkorn, F. Böttcher, M. Guo, J. N. Schmidt, T. Langen, H. P. Büchler and T. Pfau, 'Fate of the Amplitude Mode in a Trapped Dipolar Supersolid', Phys. Rev. Lett. **123**, 193002 (2019).

[83] L. Klaus, T. Bland, E. Poli, C. Politi, G. Lamporesi, E. Casotti, R. N. Bisset, M. J. Mark and F. Ferlaino, 'Observation of vortices and vortex stripes in a dipolar condensate', Nat. Phys. **18**, 1453 (2022).

[84] G. Biagioni, N. Antolini, A. Alaña, M. Modugno, A. Fioretti, C. Gabbanini, L. Tanzi and G. Modugno, 'Dimensional Crossover in the Superfluid-Supersolid Quantum Phase Transition', Phys. Rev. X **12**, 021019 (2022).

[85] G. A. Phelps, A. Hébert, A. Krahn, S. Dickerson, F. Öztürk, S. Ebadi, L. Su and M. Greiner, 'Sub-second production of a quantum degenerate gas', arXiv, 2007.10807 (2020).

[86] W. Lunden, L. Du, M. Cantara, P. Barral, A. O. Jamison and W. Ketterle, 'Enhancing the capture velocity of a Dy magneto-optical trap with two-stage slowing', Phys. Rev. A **101**, 063403 (2020).

[87] J. Ulitzsch, D. Babik, R. Roell and M. Weitz, 'Bose-Einstein condensation of erbium atoms in a quasielectrostatic optical dipole trap', Phys. Rev. A **95**, 043614 (2017).

[88] T. Chalopin, T. Satoor, A. Evrard, V. Makhalov, J. Dalibard, R. Lopes and S. Nascimbene, 'Probing chiral edge dynamics and bulk topology of a synthetic Hall system', Nat. Phys. **16**, 1017 (2020).

[89] B. Seo, P. Chen, Z. Chen, W. Yuan, M. Huang, S. Du and G.-B. Jo, 'Efficient production of a narrow-line erbium magneto-optical trap with two-stage slowing', Phys. Rev. A **102**, 013319 (2020).

[90] R. Lopes, 'Radio-frequency evaporation in an optical dipole trap', Phys. Rev. A **104**, 033313 (2021).

[91] C. Ravensbergen, E. Soave, V. Corre, M. Kreyer, B. Huang, E. Kirilov and R. Grimm, 'Resonantly Interacting Fermi-Fermi Mixture of ^{161}Dy and ^{40}K', Phys. Rev. Lett. **124**, 203402 (2020).

[92] C. Politi, A. Trautmann, P. Ilzhöfer, G. Durastante, M. J. Mark, M. Modugno and F. Ferlaino, 'Interspecies interactions in an ultracold dipolar mixture', Phys. Rev. A **105**, 023304 (2022).

[93] B. Kain and H. Y. Ling, 'Polarons in a dipolar condensate', Phys. Rev. A **89**, 023612 (2014).

[94] B. Bylicka, D. Chruściński and S. Maniscalco, 'Non-Markovianity and reservoir memory of quantum channels: a quantum information theory perspective', Sci. Rep. **4**, 5720 (2014).

[95] J.-B. Yuan, H.-J. Xing, L.-M. Kuang and S. Yi, 'Quantum non-Markovian reservoirs of atomic condensates engineered via dipolar interactions', Phys. Rev. A **95**, 033610 (2017).

[96] L. D. Carr, D. DeMille, R. V. Krems and J. Ye, 'Cold and ultracold molecules: science, technology and applications', New J. Phys. **11**, 055049 (2009).

[97] J. L. Bohn, A. M. Rey and J. Ye, 'Cold molecules: Progress in quantum engineering of chemistry and quantum matter', Science **357**, 1002 (2017).

[98] S. A. Moses, J. P. Covey, M. T. Miecnikowski, D. S. Jin and J. Ye, 'New frontiers for quantum gases of polar molecules', Nat. Phys. **13**, 13 (2017).

[99] M. Saffman, T. G. Walker and K. Mølmer, 'Quantum information with Rydberg atoms', Rev. Mod. Phys. **82**, 2313 (2010).

[100] R. Löw, H. Weimer, J. Nipper, J. B. Balewski, B. Butscher, H. P. Büchler and T. Pfau, 'An experimental and theoretical guide to strongly interacting Rydberg gases', J. Phys. B: At. Mol. Opt. Phys. **45**, 113001 (2012).

[101] M. S. Safronova, D. Budker, D. DeMille, D. F. J. Kimball, A. Derevianko and C. W. Clark, 'Search for new physics with atoms and molecules', Rev. Mod. Phys. **90**, 025008 (2018).

[102] L. Pitaevskii and S. Stringari, *Bose-Einstein Condensation and Superfluidity*, International Series of Monographs on Physics (Oxford University Press, Oxford, 2016).

[103] J. Dalibard, 'Collisional dynamics of ultra-cold atomic gases', in *Bose-Einstein Condensation in Atomic Gases*, Proceedings of the International School of Physics "Enrico Fermi" 140 (IOS Press, Amsterdam, 1999), p. 321.

[104] M. Krstajić, 'Experimental Platform for a Box-Trapped Dipolar Quantum Gas', PhD thesis (University of Cambridge, 2021).

[105] J. R. Nuñez, C. R. Anderton and R. S. Renslow, 'Optimizing colormaps with consideration for color vision deficiency to enable accurate interpretation of scientific data', PLOS One **13**, e0199239 (2018).

[106] L. R. Hofer, M. Krstajić, P. Juhász, A. L. Marchant and R. P. Smith, 'Atom cloud detection and segmentation using a deep neural network', Mach. Learn.: Sci. Technol. **2**, 045008 (2021).

[107] M. Krstajić, P. Juhász, J. Kučera, L. R. Hofer, G. Lamb, A. L. Marchant and R. P. Smith, 'Characterisation of three-body loss in ^{166}Er and optimised production of large Bose–Einstein condensates', arXiv, 2307.01245 (2023).

[108] H. Ritsch, P. Domokos, F. Brennecke and T. Esslinger, 'Cold atoms in cavity-generated dynamical optical potentials', Rev. Mod. Phys. **85**, 553 (2013).

[109] F. Mivehvar, F. Piazza, T. Donner and H. Ritsch, 'Cavity QED with quantum gases: new paradigms in many-body physics', Adv. Phys. **70**, 1 (2021).

[110] V. D. Vaidya, Y. Guo, R. M. Kroeze, K. E. Ballantine, A. J. Kollár, J. Keeling and B. L. Lev, 'Tunable-Range, Photon-Mediated Atomic Interactions in Multimode Cavity QED', Phys. Rev. X **8**, 011002 (2018).

[111] C. G. Mosander, 'XXX. On the new metals, lanthanium and didymium, which are associated with cerium; and on erbium and terbium, new metals associated with yttria', Lond. Edinb. Dublin Philos. Mag. J. Sci. **23**, 241 (1843).

[112] C. Piguet, 'Extricating erbium', Nat. Chem. **6**, 370 (2014).

[113] W. M. Haynes, ed., *CRC Handbook of Chemistry and Physics*, 95th ed. (CRC Press, Boca Raton, 2014).

[114] J. K. Böhlke, J. R. de Laeter, P. De Bièvre, H. Hidaka, H. S. Peiser, K. J. R. Rosman and P. D. P. Taylor, 'Isotopic Compositions of the Elements, 2001', J. Phys. Chem. Ref. Data **34**, 57 (2005).

[115] B. G. Wybourne and L. Smentek, *Optical Spectroscopy of Lanthanides: Magnetic and Hyperfine Interactions* (CRC Press, Boca Raton, 2007).

[116] A. Kramida, Y. Ralchenko, J. Reader and the NIST ASD Team, *NIST Atomic Spectra Database*, version 5.10, National Institute of Standards and Technology (Gaithersburg, MD, 2022).

[117] W. C. Martin, R. Zalubas and L. Hagan, *Atomic Energy Levels—The Rare-Earth Elements*, National Bureau of Standards (United States, 1978).

[118] H. Y. Ban, M. Jacka, J. L. Hanssen, J. Reader and J. J. McClelland, 'Laser cooling transitions in atomic erbium', Opt. Express **13**, 3185 (2005).

[119] J. J. McClelland and J. L. Hanssen, 'Laser Cooling without Repumping: A Magneto-Optical Trap for Erbium Atoms', Phys. Rev. Lett. **96**, 143005 (2006).

[120] A. Frisch, K. Aikawa, M. Mark, A. Rietzler, J. Schindler, E. Zupanič, R. Grimm and F. Ferlaino, 'Narrow-line magneto-optical trap for erbium', Phys. Rev. A **85**, 051401(R) (2012).

[121] T. J. Bruno and P. D. N. Svoronos, *CRC Handbook of Fundamental Spectroscopic Correlation Charts* (CRC Press, Boca Raton, 2005).

[122] K. L. Kelly, 'Color Designations for Lights', J. Opt. Soc. Am. **33**, 627 (1943).

[123] E. A. Den Hartog, J. P. Chisholm and J. E. Lawler, 'Radiative lifetimes of neutral erbium', J. Phys. B: At. Mol. Opt. Phys. **43**, 155004 (2010).

[124] C. B. Connolly, Y. S. Au, S. C. Doret, W. Ketterle and J. M. Doyle, 'Large spin relaxation rates in trapped submerged-shell atoms', Phys. Rev. A **81**, 010702(R) (2010).

[125] W. Guo Jin, T. Horiguchi, M. Wakasugi and Y. Yoshizawa, 'Hyperfine Structure and Isotope Shift in Er I by the Atomic-Beam Laser Spectroscopy', J. Phys. Soc. Jpn. **59**, 3148 (1990).

[126] A. Frisch, 'Dipolar Quantum Gases of Erbium', PhD thesis (University of Innsbruck, 2014).

[127] S. Methfessel, 'Potential applications of magnetic rare earth compunds', IEEE Trans. Magn. **1**, 144 (1965).

[128] *Standard Specifications for Permanent Magnetic Materials*, Magnetic Materials Producers Association (Chicago, IL).

[129] J. Jensen and A. R. Mackintosh, *Rare Earth Magnetism: Structures and Excitations*, International Series of Monographs on Physics (Clarendon Press, Oxford, 1991).

[130] E. Tiesinga, P. J. Mohr, D. B. Newell and B. N. Taylor, 'CODATA recommended values of the fundamental physical constants: 2018', Rev. Mod. Phys. **93**, 025010 (2021).

[131] B. R. Judd and I. Lindgren, 'Theory of Zeeman Effect in the Ground Multiplets of Rare-Earth Atoms', Phys. Rev. **122**, 1802 (1961).

[132] J. G. Conway and B. G. Wybourne, 'Low-Lying Energy Levels of Lanthanide Atoms and Intermediate Coupling', Phys. Rev. **130**, 2325 (1963).

[133] C. J. Pethick and H. Smith, *Bose–Einstein Condensation in Dilute Gases*, 2nd ed. (Cambridge University Press, Cambridge, 2008).

[134] S. Kotochigova, 'Controlling interactions between highly magnetic atoms with Feshbach resonances', Rep. Prog. Phys. **77**, 093901 (2014).

[135] T. Maier, 'Interactions in a Quantum Gas of Dysprosium Atoms', PhD thesis (University of Stuttgart, 2015).

[136] S. Giovanazzi, A. Görlitz and T. Pfau, 'Tuning the Dipolar Interaction in Quantum Gases', Phys. Rev. Lett. **89**, 130401 (2002).

[137] G. E. Astrakharchik and Yu. E. Lozovik, 'Super-Tonks-Girardeau regime in trapped one-dimensional dipolar gases', Phys. Rev. A **77**, 013404 (2008).

[138] L. D. Landau and E. M. Lifshitz, *Quantum Mechanics: Non-Relativistic Theory*, 3rd ed. (Pergamon Press, Oxford, 1977).

[139] J. J. Sakurai and J. Napolitano, *Modern Quantum Mechanics*, 3rd ed. (Cambridge University Press, Cambridge, 2020).

[140] S. Hensler, J. Werner, A. Griesmaier, P. O. Schmidt, A. Görlitz, T. Pfau, S. Giovanazzi and K. Rzążewski, 'Dipolar relaxation in an ultra-cold gas of magnetically trapped chromium atoms', Appl. Phys. B **77**, 765 (2003).

[141] M. Fattori, T. Koch, S. Goetz, A. Griesmaier, S. Hensler, J. Stuhler and T. Pfau, 'Demagnetization cooling of a gas', Nat. Phys. **2**, 765 (2006).

[142] V. V. Volchkov, J. Rührig, T. Pfau and A. Griesmaier, 'Efficient demagnetization cooling of atoms and its limits', Phys. Rev. A **89**, 043417 (2014).

[143] B. Naylor, E. Maréchal, J. Huckans, O. Gorceix, P. Pedri, L. Vernac and B. Laburthe-Tolra, 'Cooling of a Bose-Einstein Condensate by Spin Distillation', Phys. Rev. Lett. **115**, 243002 (2015).

[144] L. S. Butcher, D. N. Stacey, C. J. Foot and K. Burnett, 'Ultracold collisions for Bose—Einstein condensation', Philos. Trans. R. Soc. A **357**, 1421 (1999).

[145] M. A. Baranov, 'Theoretical progress in many-body physics with ultracold dipolar gases', Phys. Rep. **464**, 71 (2008).

[146] J. L. Bohn, M. Cavagnero and C. Ticknor, 'Quasi-universal dipolar scattering in cold and ultracold gases', New J. Phys. **11**, 055039 (2009).

[147] J. L. Bohn and D. S. Jin, 'Differential scattering and rethermalization in ultracold dipolar gases', Phys. Rev. A **89**, 022702 (2014).

[148] K. Aikawa, A. Frisch, M. Mark, S. Baier, R. Grimm and F. Ferlaino, 'Reaching Fermi Degeneracy via Universal Dipolar Scattering', Phys. Rev. Lett. **112**, 010404 (2014).

[149] R. Ołdziejewski and K. Jachymski, 'Properties of strongly dipolar Bose gases beyond the Born approximation', Phys. Rev. A **94**, 063638 (2016); 'Erratum: Properties of strongly dipolar Bose gases beyond the Born approximation [Phys. Rev. A 94, 063638 (2016)]', Phys. Rev. A **95**, 049901 (2017).

[150] S. Yi and L. You, 'Trapped atomic condensates with anisotropic interactions', Phys. Rev. A **61**, 041604(R) (2000).

[151] S. Yi and L. You, 'Trapped condensates of atoms with dipole interactions', Phys. Rev. A **63**, 053607 (2001).

[152] J. Metz, 'Collapse of dipolar Bose-Einstein condensates for different trap geometries', PhD thesis (University of Stuttgart, 2010).

[153] S. Ronen, D. C. E. Bortolotti, D. Blume and J. L. Bohn, 'Dipolar Bose-Einstein condensates with dipole-dependent scattering length', Phys. Rev. A **74**, 033611 (2006).

[154] D. C. E. Bortolotti, S. Ronen, J. L. Bohn and D. Blume, 'Scattering Length Instability in Dipolar Bose-Einstein Condensates', Phys. Rev. Lett. **97**, 160402 (2006).

[155] I. E. Mazets and G. Kurizki, 'Modification of Scattering Lengths via Magnetic Dipole-Dipole Interactions', Phys. Rev. Lett. **98**, 140401 (2007).

[156] S. Müller, 'Stability and Collapse Dynamics of Dipolar Bose-Einstein Condensates in One-Dimensional Optical Lattices', PhD thesis (University of Stuttgart, 2013).

[157] O. K. Rice, 'Predissociation and the Crossing of Molecular Potential Energy Curves', J. Chem. Phys. **1**, 375 (1933).

[158] H. Beutler, 'Über Absorptionsserien von Argon, Krypton und Xenon zu Termen zwischen den beiden Ionisierungsgrenzen $^2P^0_{3/2}$ und $^2P^0_{1/2}$', Z. Phys. **93**, 177 (1935).

[159] U. Fano, 'Sullo spettro di assorbimento dei gas nobili presso il limite dello spettro d'arco', Nuovo Cimento **12**, 154 (1935).

[160] H. Feshbach, 'Unified theory of nuclear reactions', Ann. Phys. **5**, 357 (1958).

[161] H. Feshbach, 'A unified theory of nuclear reactions. II', Ann. Phys. **19**, 287 (1962).

[162] U. Fano, 'Effects of Configuration Interaction on Intensities and Phase Shifts', Phys. Rev. **124**, 1866 (1961).

[163] D. Kleppner, 'Professor Feshbach and His Resonance', Phys. Today **57**, 12 (2004).

[164] A. R. P. Rau, 'Historical Notes on Feshbach and Shape Resonances', Phys. Today **58**, 13 (2005).

[165] K. Bergmann, H. Theuer and B. W. Shore, 'Coherent population transfer among quantum states of atoms and molecules', Rev. Mod. Phys. **70**, 1003 (1998).

[166] D. Spence and T. Noguchi, 'Feshbach resonances associated with Rydberg states of the hydrogen halides', J. Chem. Phys. **63**, 505 (1975).

[167] H. C. Bryant, B. D. Dieterle, J. Donahue, H. Sharifian, H. Tootoonchi, D. M. Wolfe, P. A. M. Gram and M. A. Yates-Williams, 'Observation of Resonances near 11 eV in the Photodetachment Cross Section of the H^- Ion', Phys. Rev. Lett. **38**, 228 (1977); 'Erratum: Observation of Resonances near 11 eV in the Photodetachment Cross Section of the H^- Ion', Phys. Rev. Lett. **39**, 1236 (1977).

[168] T. Weber, J. Herbig, M. Mark, H.-C. Nägerl and R. Grimm, 'Three-Body Recombination at Large Scattering Lengths in an Ultracold Atomic Gas', Phys. Rev. Lett. **91**, 123201 (2003).

[169] A. Frisch, M. Mark, K. Aikawa, F. Ferlaino, J. L. Bohn, C. Makrides, A. Petrov and S. Kotochigova, 'Quantum chaos in ultracold collisions of gas-phase erbium atoms', Nature **507**, 475 (2014).

[170] A. Patscheider, L. Chomaz, G. Natale, D. Petter, M. J. Mark, S. Baier, B. Yang, R. R. W. Wang, J. L. Bohn and F. Ferlaino, 'Determination of the scattering length of erbium atoms', Phys. Rev. A **105**, 063307 (2022).

[171] A. Petrov, E. Tiesinga and S. Kotochigova, 'Anisotropy-Induced Feshbach Resonances in a Quantum Dipolar Gas of Highly Magnetic Atoms', Phys. Rev. Lett. **109**, 103002 (2012).

[172] P. O. Fedichev, Yu. Kagan, G. V. Shlyapnikov and J. T. M. Walraven, 'Influence of Nearly Resonant Light on the Scattering Length in Low-Temperature Atomic Gases', Phys. Rev. Lett. **77**, 2913 (1996).

[173] J. L. Bohn and P. S. Julienne, 'Prospects for influencing scattering lengths with far-off-resonant light', Phys. Rev. A **56**, 1486 (1997).

[174] R. Ciuryło, E. Tiesinga and P. Julienne, 'Optical tuning of the scattering length of cold alkaline-earth-metal atoms', Phys. Rev. A **71**, 030701(R) (2005).

[175] A. J. Moerdijk, B. J. Verhaar and A. Axelsson, 'Resonances in ultracold collisions of ^{6}Li, ^{7}Li, and ^{23}Na', Phys. Rev. A **51**, 4852 (1995).

[176] A. D. Lange, K. Pilch, A. Prantner, F. Ferlaino, B. Engeser, H.-C. Nägerl, R. Grimm and C. Chin, 'Determination of atomic scattering lengths from measurements of molecular binding energies near Feshbach resonances', Phys. Rev. A **79**, 013622 (2009).

[177] K. Jachymski and P. S. Julienne, 'Analytical model of overlapping Feshbach resonances', Phys. Rev. A **88**, 052701 (2013).

[178] I. N. Bronshtein, K. A. Semendyayev, G. Musiol and H. Muehlig, *Handbook of Mathematics* (Springer, Berlin, Heidelberg, 2007).

[179] Y. Castin, 'Bose-Einstein Condensates in Atomic Gases: Simple Theoretical Results', in *Coherent atomic matter waves – Ondes de matière cohérente*, edited by R. Kaiser, C. Westbrook and F. David, Les Houches – École d'Été de Physique Théorique 72 (Springer, Berlin, Heidelberg, 2001), p. 1.

[180] A. A. Penzias and R. W. Wilson, 'A Measurement of Excess Antenna Temperature at 4080 Mc/s', Astrophys. J. **142**, 419 (1965).

[181] 'The Nobel Prize in Physics 1978', NobelPrize.org, Press Release (1978).

[182] R. Sahai and L.-Å. Nyman, 'The Boomerang Nebula: The Coldest Region of the Universe?', Astrophys. J. **487**, L155 (1997).

[183] E. P. Gross, 'Structure of a Quantized Vortex in Boson Systems', Nuovo Cimento **20**, 454 (1961).

[184] L. P. Pitaevskii, 'Vortex Lines in an Imperfect Bose Gas', Sov. Phys. JETP **13**, 451 (1961).

[185] L. Erdős, B. Schlein and H.-T. Yau, 'Derivation of the Gross-Pitaevskii equation for the dynamics of Bose-Einstein condensate', Ann. Math. **172**, 291 (2010).

[186] E. H. Lieb, R. Seiringer and J. Yngvason, 'Bosons in a trap: A rigorous derivation of the Gross-Pitaevskii energy functional', Phys. Rev. A **61**, 043602 (2000).

[187] N. G. Parker and D. H. J. O'Dell, 'Thomas-Fermi versus one- and two-dimensional regimes of a trapped dipolar Bose-Einstein condensate', Phys. Rev. A **78**, 041601(R) (2008).

[188] G. E. Astrakharchik, J. Boronat, J. Casulleras, I. L. Kurbakov and Yu. E. Lozovik, 'Weakly interacting two-dimensional system of dipoles: Limitations of the mean-field theory', Phys. Rev. A **75**, 063630 (2007).

[189] A. Derevianko, 'Anisotropic pseudopotential for polarized dilute quantum gases', Phys. Rev. A **67**, 033607 (2003); 'Erratum: Anisotropic pseudopotential for polarized dilute quantum gases [Phys. Rev. A **67** , 033607 (2003)]', Phys. Rev. A **72**, 039901 (2005).

[190] N. Bogoliubov, 'On the Theory of Superfluidity', J. Phys. (USSR) **11**, 23 (1947); reprinted in D. Pines, *The Many-Body Problem* (W. A. Benjamin, New York, 1961), p. 292.

[191] A. L. Fetter, 'Nonuniform States of an Imperfect Bose Gas', Ann. Phys. **70**, 67 (1972).

[192] P. B. Blakie, D. Baillie and R. N. Bisset, 'Roton spectroscopy in a harmonically trapped dipolar Bose-Einstein condensate', Phys. Rev. A **86**, 021604(R) (2012).

[193] L. Santos, G. V. Shlyapnikov and M. Lewenstein, 'Roton-Maxon Spectrum and Stability of Trapped Dipolar Bose-Einstein Condensates', Phys. Rev. Lett. **90**, 250403 (2003).

[194] J.-N. Schmidt, J. Hertkorn, M. Guo, F. Böttcher, M. Schmidt, K. S. H. Ng, S. D. Graham, T. Langen, M. Zwierlein and T. Pfau, 'Roton Excitations in an Oblate Dipolar Quantum Gas', Phys. Rev. Lett. **126**, 193002 (2021).

[195] O. Dutta and M. Lewenstein, 'Unconventional superfluidity of fermions in Bose-Fermi mixtures', Phys. Rev. A **81**, 063608 (2010).

[196] U. R. Fischer, 'Stability of quasi-two-dimensional Bose-Einstein condensates with dominant dipole-dipole interactions', Phys. Rev. A **73**, 031602(R) (2006).

[197] Z.-Q. Yu, 'Landau criterion for an anisotropic Bose-Einstein condensate', Phys. Rev. A **95**, 033618 (2017).

[198] T. D. Lee, K. Huang and C. N. Yang, 'Eigenvalues and Eigenfunctions of a Bose System of Hard Spheres and Its Low-Temperature Properties', Phys. Rev. **106**, 1135 (1957).

[199] T. D. Lee and C. N. Yang, 'Many-Body Problem in Quantum Mechanics and Quantum Statistical Mechanics', Phys. Rev. **105**, 1119 (1957).

[200] R. Schützhold, M. Uhlmann, Y. Xu and U. R. Fischer, 'Mean-Field Expansion in Bose–Einstein Condensates with Finite-Range Interactions', Int. J. Mod. Phys. B **20**, 3555 (2006).

[201] A. R. P. Lima and A. Pelster, 'Quantum fluctuations in dipolar Bose gases', Phys. Rev. A **84**, 041604(R) (2011).

[202] A. R. P. Lima and A. Pelster, 'Beyond mean-field low-lying excitations of dipolar Bose gases', Phys. Rev. A **86**, 063609 (2012).

[203] N. P. Proukakis and B. Jackson, 'Finite-temperature models of Bose–Einstein condensation', J. Phys. B: At. Mol. Opt. Phys. **41**, 203002 (2008).

[204] M. Schmitt, M. Wenzel, F. Böttcher, I. Ferrier-Barbut and T. Pfau, 'Self-bound droplets of a dilute magnetic quantum liquid', Nature **539**, 259 (2016).

[205] A. Bühler and H. P. Büchler, 'Supersolid phase in atomic gases with magnetic dipole interaction', Phys. Rev. A **84**, 023607 (2011).

[206] F. Böttcher, J.-N. Schmidt, J. Hertkorn, K. S. H. Ng, S. D. Graham, M. Guo, T. Langen and T. Pfau, 'New states of matter with fine-tuned interactions: quantum droplets and dipolar supersolids', Rep. Prog. Phys. **84**, 012403 (2021).

[207] L. Chomaz, D. Petter, P. Ilzhöfer, G. Natale, A. Trautmann, C. Politi, G. Durastante, R. M. W. van Bijnen, A. Patscheider, M. Sohmen, M. J. Mark and F. Ferlaino, 'Long-Lived and Transient Supersolid Behaviors in Dipolar Quantum Gases', Phys. Rev. X **9**, 021012 (2019).

[208] M. Sohmen, C. Politi, L. Klaus, L. Chomaz, M. J. Mark, M. A. Norcia and F. Ferlaino, 'Birth, Life, and Death of a Dipolar Supersolid', Phys. Rev. Lett. **126**, 233401 (2021).

[209] P. Ilzhöfer, M. Sohmen, G. Durastante, C. Politi, A. Trautmann, G. Natale, G. Morpurgo, T. Giamarchi, L. Chomaz, M. J. Mark and F. Ferlaino, 'Phase coherence in out-of-equilibrium supersolid states of ultracold dipolar atoms', Nat. Phys. **17**, 356 (2021).

[210] J. Hertkorn, J.-N. Schmidt, F. Böttcher, M. Guo, M. Schmidt, K. S. H. Ng, S. D. Graham, H. P. Büchler, T. Langen, M. Zwierlein and T. Pfau, 'Density Fluctuations across the Superfluid-Supersolid Phase Transition in a Dipolar Quantum Gas', Phys. Rev. X **11**, 011037 (2021).

[211] J. Hertkorn, J.-N. Schmidt, M. Guo, F. Böttcher, K. S. H. Ng, S. D. Graham, P. Uerlings, H. P. Büchler, T. Langen, M. Zwierlein and T. Pfau, 'Supersolidity in Two-Dimensional Trapped Dipolar Droplet Arrays', Phys. Rev. Lett. **127**, 155301 (2021).

[212] M. A. Norcia, C. Politi, L. Klaus, E. Poli, M. Sohmen, M. J. Mark, R. N. Bisset, L. Santos and F. Ferlaino, 'Two-dimensional supersolidity in a dipolar quantum gas', Nature **596**, 357 (2021).

[213] J. Hertkorn, J.-N. Schmidt, M. Guo, F. Böttcher, K. S. H. Ng, S. D. Graham, P. Uerlings, T. Langen, M. Zwierlein and T. Pfau, 'Pattern formation in quantum ferrofluids: From supersolids to superglasses', Phys. Rev. Res. **3**, 033125 (2021).

[214] M. Boninsegni and N. V. Prokof'ev, '*Colloquium*: Supersolids: What and where are they?', Rev. Mod. Phys. **84**, 759 (2012).

[215] L. Chomaz, S. Baier, D. Petter, M. J. Mark, F. Wächtler, L. Santos and F. Ferlaino, 'Quantum-Fluctuation-Driven Crossover from a Dilute Bose-Einstein Condensate to a Macrodroplet in a Dipolar Quantum Fluid', Phys. Rev. X **6**, 041039 (2016).

[216] C. J. Foot, *Atomic Physics*, Oxford Master Series in Physics (Oxford University Press, Oxford, 2005).

[217] N. B. Vilas, 'Laser Cooling Towards a Dipolar Quantum Gas of Erbium Atoms', Master's thesis (University of Cambridge, 2018).

[218] G. Camy, Ch. J. Bordé and M. Ducloy, 'Heterodyne saturation spectroscopy through frequency modulation of the saturating beam', Opt. Commun. **41**, 325 (1982).

[219] J. H. Shirley, 'Modulation transfer processes in optical heterodyne saturation spectroscopy', Opt. Lett. **7**, 537 (1982).

[220] L. Neuhaus, R. Metzdorff, S. Chua, T. Jacqmin, T. Briant, A. Heidmann, P.-F. Cohadon and S. Deleglise, 'PyRPL (Python Red Pitaya Lockbox) — An open-source software package for FPGA-controlled quantum optics experiments', in 2017 Conference on Lasers and Electro-Optics Europe & European Quantum Electronics Conference (CLEO/Europe-EQEC) (2017).

[221] I. Manai, A. Molineri, C. Fréjaville, C. Duval, P. Bataille, R. Journet, F. Wiotte, B. Laburthe-Tolra, E. Maréchal, M. Cheneau and M. Robert-de-Saint-Vincent, 'Shelving spectroscopy of the strontium intercombination line', J. Phys. B: At. Mol. Opt. Phys. **53**, 085005 (2020).

[222] R. V. Pound, 'Electronic Frequency Stabilization of Microwave Oscillators', Rev. Sci. Instrum. **17**, 490 (1946).

[223] R. W. P. Drever, J. L. Hall, F. V. Kowalski, J. Hough, G. M. Ford, A. J. Munley and H. Ward, 'Laser phase and frequency stabilization using an optical resonator', Appl. Phys. B **31**, 97 (1983).

[224] E. D. Black, 'An introduction to Pound–Drever–Hall laser frequency stabilization', Am. J. Phys. **69**, 79 (2001).

[225] P. D. Lett, W. D. Phillips, S. L. Rolston, C. E. Tanner, R. N. Watts and C. I. Westbrook, 'Optical molasses', J. Opt. Soc. Am. B **6**, 2084 (1989).

[226] W. D. Phillips and H. Metcalf, 'Laser Deceleration of an Atomic Beam', Phys. Rev. Lett. **48**, 596 (1982).

[227] E. L. Raab, M. Prentiss, A. Cable, S. Chu and D. E. Pritchard, 'Trapping of Neutral Sodium Atoms with Radiation Pressure', Phys. Rev. Lett. **59**, 2631 (1987).

[228] T. H. Loftus, T. Ido, M. M. Boyd, A. D. Ludlow and J. Ye, 'Narrow line cooling and momentum-space crystals', Phys. Rev. A **70**, 063413 (2004).

[229] T. Chanelière, L. He, R. Kaiser and D. Wilkowski, 'Three dimensional cooling and trapping with a narrow line', Eur. Phys. J. D **46**, 507 (2008).

[230] T. H. Boyer, 'The force on a magnetic dipole', Am. J. Phys. **56**, 688 (1988).

[231] D. Dreon, L. A. Sidorenkov, C. Bouazza, W. Maineult, J. Dalibard and S. Nascimbene, 'Optical cooling and trapping of highly magnetic atoms: the benefits of a spontaneous spin polarization', J. Phys. B: At. Mol. Opt. Phys. **50**, 065005 (2017).

[232] A. Ashkin, 'Optical trapping and manipulation of neutral particles using lasers', Proc. Natl. Acad. Sci. U.S.A. **94**, 4853 (1997).

[233] M. Lepers, J.-F. Wyart and O. Dulieu, 'Anisotropic optical trapping of ultracold erbium atoms', Phys. Rev. A **89**, 022505 (2014).

[234] J. H. Becher, S. Baier, K. Aikawa, M. Lepers, J.-F. Wyart, O. Dulieu and F. Ferlaino, 'Anisotropic polarizability of erbium atoms', Phys. Rev. A **97**, 012509 (2018).

[235] C. Kohstall, 'A New Toolbox for Experiments with Ultracold ^{6}Li', Master's thesis (University of Innsbruck, 2007).

[236] M. Greiner, I. Bloch, T. W. Hänsch and T. Esslinger, 'Magnetic transport of trapped cold atoms over a large distance', Phys. Rev. A **63**, 031401(R) (2001).

[237] J. Goldwin, S. Inouye, M. L. Olsen, B. Newman, B. D. DePaola and D. S. Jin, 'Measurement of the interaction strength in a Bose-Fermi mixture with ^{87}Rb and ^{40}K', Phys. Rev. A **70**, 021601(R) (2004).

[238] H. J. Lewandowski, D. M. Harber, D. L. Whitaker and E. A. Cornell, 'Simplified System for Creating a Bose–Einstein Condensate', J. Low Temp. Phys. **132**, 309 (2003).

[239] T. L. Gustavson, A. P. Chikkatur, A. E. Leanhardt, A. Görlitz, S. Gupta, D. E. Pritchard and W. Ketterle, 'Transport of Bose-Einstein Condensates with Optical Tweezers', Phys. Rev. Lett. **88**, 020401 (2001).

[240] J. Léonard, M. Lee, A. Morales, T. M. Karg, T. Esslinger and T. Donner, 'Optical transport and manipulation of an ultracold atomic cloud using focus-tunable lenses', New J. Phys. **16**, 093028 (2014).

[241] A. E. Siegman, *Lasers* (Oxford University Press, Oxford, 1986).

[242] E. Torrontegui, S. Ibáñez, X. Chen, A. Ruschhaupt, D. Guéry-Odelin and J. G. Muga, 'Fast atomic transport without vibrational heating', Phys. Rev. A **83**, 013415 (2011).

[243] M. Murphy, L. Jiang, N. Khaneja and T. Calarco, 'High-fidelity fast quantum transport with imperfect controls', Phys. Rev. A **79**, 020301(R) (2009).

[244] X. Chen, M. Zeuner, U. Schneider, C. J. Foot, T. L. Harte and E. Bentine, 'AtomECS: Simulate laser cooling and magneto-optical traps', arXiv, 2105.06447 (2021).

[245] A. Keshet and W. Ketterle, 'A distributed, graphical user interface based, computer control system for atomic physics experiments', Rev. Sci. Instrum. **84**, 015105 (2013).

[246] L. R. Hofer, M. Krstajić and R. P. Smith, 'JAXFit: Trust Region Method for Nonlinear Least-Squares Curve Fitting on the GPU', arXiv, 2208.12187 (2022).

[247] L. R. Hofer, M. Krstajić and R. P. Smith, 'Measuring laser beams with a neural network', Appl. Opt. **61**, 1924 (2022).

[248] M. Ligare, 'Classical thermodynamics of particles in harmonic traps', Am. J. Phys. **78**, 815 (2010).

[249] W. Ketterle and N. J. van Druten, 'Evaporative Cooling of Trapped Atoms', in *Adv. At., Mol., Opt. Phys.* Vol. 37 (Academic Press, San Diego, 1996), p. 181.

[250] A. E. Leanhardt, T. A. Pasquini, M. Saba, A. Schirotzek, Y. Shin, D. Kielpinski, D. E. Pritchard and W. Ketterle, 'Cooling Bose-Einstein Condensates Below 500 Picokelvin', Science **301**, 1513 (2003).

[251] S. K. Ruddell, D. H. White, A. Ullah, D. Baillie and M. D. Hoogerland, 'Calorimetry of a harmonically trapped Bose gas', Phys. Rev. A **92**, 063622 (2015).

[252] W. Ketterle, D. S. Durfee and D. M. Stamper-Kurn, 'Making, probing and understanding Bose-Einstein condensates', in *Bose-Einstein Condensation in Atomic Gases*, Proceedings of the International School of Physics "Enrico Fermi" 140 (IOS Press, Amsterdam, 1999), p. 67.

[253] L. Tanzi, S. M. Roccuzzo, E. Lucioni, F. Famà, A. Fioretti, C. Gabbanini, G. Modugno, A. Recati and S. Stringari, 'Supersolid symmetry breaking from compressional oscillations in a dipolar quantum gas', Nature **574**, 382 (2019).

[254] T. Bland, E. Poli, C. Politi, L. Klaus, M. A. Norcia, F. Ferlaino, L. Santos and R. N. Bisset, 'Two-Dimensional Supersolid Formation in Dipolar Condensates', Phys. Rev. Lett. **128**, 195302 (2022).

[255] D. Baillie and P. B. Blakie, 'Droplet Crystal Ground States of a Dipolar Bose Gas', Phys. Rev. Lett. **121**, 195301 (2018).

[256] Y.-C. Zhang, T. Pohl and F. Maucher, 'Phases of supersolids in confined dipolar Bose-Einstein condensates', Phys. Rev. A **104**, 013310 (2021).

[257] E. Poli, T. Bland, C. Politi, L. Klaus, M. A. Norcia, F. Ferlaino, R. N. Bisset and L. Santos, 'Maintaining supersolidity in one and two dimensions', Phys. Rev. A **104**, 063307 (2021).

[258] E. A. Burt, R. W. Ghrist, C. J. Myatt, M. J. Holland, E. A. Cornell and C. E. Wieman, 'Coherence, Correlations, and Collisions: What One Learns about Bose-Einstein Condensates from Their Decay', Phys. Rev. Lett. **79**, 337 (1997).

[259] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov et al., 'SciPy 1.0: fundamental algorithms for scientific computing in Python', Nat. Methods **17**, 261 (2020).

[260] P. R. Bevington and D. K. Robinson, *Data Reduction and Error Analysis for the Physical Sciences*, 3rd ed. (McGraw-Hill, Boston, 2003), 320 pp.

[261] Q. Beaufils, A. Crubellier, T. Zanon, B. Laburthe-Tolra, E. Maréchal, L. Vernac and O. Gorceix, 'Feshbach resonance in d-wave collisions', Phys. Rev. A **79**, 032706 (2009).

[262] L. Fouché, A. Boissé, G. Berthet, S. Lepoutre, A. Simoni and T. Bourdel, 'Quantitative analysis of losses close to a d-wave open-channel Feshbach resonance in ^{39}K', Phys. Rev. A **99**, 022701 (2019).

[263] C. Eigen, J. A. P. Glidden, R. Lopes, E. A. Cornell, R. P. Smith and Z. Hadzibabic, 'Universal prethermal dynamics of Bose gases quenched to unitarity', Nature **563**, 221 (2018).

[264] N. Navon, C. Eigen, J. Zhang, R. Lopes, A. L. Gaunt, K. Fujimoto, M. Tsubota, R. P. Smith and Z. Hadzibabic, 'Synthetic dissipation and cascade fluxes in a turbulent quantum gas', Science **366**, 382 (2019).

[265] P. Christodoulou, M. Gałka, N. Dogra, R. Lopes, J. Schmitt and Z. Hadzibabic, 'Observation of first and second sound in a BKT superfluid', Nature **594**, 191 (2021).

[266] S. Ronen, D. C. E. Bortolotti and J. L. Bohn, 'Radial and Angular Rotons in Trapped Dipolar Gases', Phys. Rev. Lett. **98**, 030406 (2007).

[267] R. M. Wilson, S. Ronen, J. L. Bohn and H. Pu, 'Manifestations of the Roton Mode in Dipolar Bose-Einstein Condensates', Phys. Rev. Lett. **100**, 245302 (2008).

[268] P. B. Blakie, D. Baillie, L. Chomaz and F. Ferlaino, 'Supersolidity in an elongated dipolar condensate', Phys. Rev. Res. **2**, 043318 (2020).

[269] S. M. Roccuzzo, A. Gallemí, A. Recati and S. Stringari, 'Rotating a Supersolid Dipolar Gas', Phys. Rev. Lett. **124**, 045702 (2020).

[270] A. Gallemí, S. M. Roccuzzo, S. Stringari and A. Recati, 'Quantized vortices in dipolar supersolid Bose-Einstein-condensed gases', Phys. Rev. A **102**, 023322 (2020).

[271] M. Guo, F. Böttcher, J. Hertkorn, J.-N. Schmidt, M. Wenzel, H. P. Büchler, T. Langen and T. Pfau, 'The low-energy Goldstone mode in a trapped dipolar supersolid', Nature **574**, 386 (2019).

[272] M. N. Tengstrand, D. Boholm, R. Sachdeva, J. Bengtsson and S. M. Reimann, 'Persistent currents in toroidal dipolar supersolids', Phys. Rev. A **103**, 013313 (2021).

[273] D. Baillie and P. B. Blakie, 'A general theory of flattened dipolar condensates', New J. Phys. **17**, 033028 (2015).

[274] S. M. Roccuzzo and F. Ancilotto, 'Supersolid behavior of a dipolar Bose-Einstein condensate confined in a tube', Phys. Rev. A **99**, 041601(R) (2019).

[275] Y.-C. Zhang, F. Maucher and T. Pohl, 'Supersolidity around a Critical Point in Dipolar Bose-Einstein Condensates', Phys. Rev. Lett. **123**, 015301 (2019).

[276] F. Ancilotto, M. Barranco, M. Pi and L. Reatto, 'Vortex properties in the extended supersolid phase of dipolar Bose-Einstein condensates', Phys. Rev. A **103**, 033314 (2021).

[277] B. Kh. Turmanov, B. B. Baizakov, F. Kh. Abdullaev and M. Salerno, 'Oscillations of a quasi-one-dimensional dipolar supersolid', J. Phys. B: At. Mol. Opt. Phys. **54**, 145302 (2021).

[278] S. Pal, D. Baillie and P. B. Blakie, 'Infinite dipolar droplet: A simple theory for the macro-droplet regime', Phys. Rev. A **105**, 023308 (2022).

[279] H.-Y. Lu, H. Lu, J.-N. Zhang, R.-Z. Qiu, H. Pu and S. Yi, 'Spatial density oscillations in trapped dipolar condensates', Phys. Rev. A **82**, 023622 (2010).

[280] S. M. Roccuzzo, S. Stringari and A. Recati, 'Supersolid edge and bulk phases of a dipolar quantum gas in a box', Phys. Rev. Res. **4**, 013086 (2022).

[281] K. Góral and L. Santos, 'Ground state and elementary excitations of single and binary Bose-Einstein condensates of trapped dipolar gases', Phys. Rev. A **66**, 023613 (2002).

[282] S. Ronen, D. C. E. Bortolotti and J. L. Bohn, 'Bogoliubov modes of a dipolar condensate in a cylindrical trap', Phys. Rev. A **74**, 013623 (2006).

[283] M. Feit, J. Fleck and A. Steiger, 'Solution of the Schrödinger equation by a spectral method', J. Comput. Phys. **47**, 412 (1982).

[284] T. Bland, 'Elementary and topological excitations in ultracold dipolar Bose gases', PhD thesis (Newcastle University, 2018).

[285] A. L. Gaunt, 'Degenerate Bose Gases: Tuning Interactions & Geometry', PhD thesis (University of Cambridge, 2014).

[286] W. H. Press, S. A. Teukolsky, W. T. Vetterling and B. P. Flannery, *Numerical Recipes: The Art of Scientific Computing*, 3rd ed. (Cambridge University Press, Cambridge, 2007).

[287] J. Gaidamour, Q. Tang and X. Antoine, 'BEC2HPC: A HPC spectral solver for nonlinear Schrödinger and rotating Gross-Pitaevskii equations. Stationary states computation', Comput. Phys. Commun. **265**, 108007 (2021).

[288] X. Antoine, A. Levitt and Q. Tang, 'Efficient spectral computation of the stationary states of rotating Bose–Einstein condensates by preconditioned nonlinear conjugate gradient methods', J. Comput. Phys. **343**, 92 (2017).

[289] X. Antoine, Q. Tang and Y. Zhang, 'A Preconditioned Conjugated Gradient Method for Computing Ground States of Rotating Dipolar Bose-Einstein Condensates via Kernel Truncation Method for Dipole-Dipole Interaction Evaluation', Commun. Comput. Phys. **24**, 966 (2018).

[290] G. van Rossum and F. L. Drake, *Python 3 Reference Manual* (CreateSpace, Scotts Valley, 2009).

[291] P. Juhász, M. Krstajić, D. Strachan, E. Gandar and R. P. Smith, *How to realise a homogeneous dipolar Bose gas in the roton regime (data)* (University of Oxford, 2022).

[292] J. Mur-Petit, A. Polls and F. Mazzanti, 'The variational principle and simple properties of the ground-state wave function', Am. J. Phys. **70**, 808 (2002).

[293] M. L. Chiofalo, S. Succi and M. P. Tosi, 'Ground state of trapped interacting Bose-Einstein condensates by an explicit imaginary-time algorithm', Phys. Rev. E **62**, 7438 (2000).

[294] R. K. Kumar, L. E. Young-S., D. Vudragović, A. Balaž, P. Muruganandam and S. Adhikari, 'Fortran and C programs for the time-dependent dipolar Gross–Pitaevskii equation in an anisotropic trap', Comput. Phys. Commun. **195**, 117 (2015).

[295] E. Polak and G. Ribière, 'Note sur la convergence de méthodes de directions conjuguées', Rev. Fr. Inform. Rech. Opér. **3**, 35 (1969).

[296] A. L. Fetter and D. Rokhsar, 'Excited states of a dilute Bose-Einstein condensate in a harmonic trap', Phys. Rev. A **57**, 1191 (1998).

[297] K.-M. You, S.-C. Wen, L.-Z. Chen, Y.-W. Wang and Y.-H. Hu, 'A quasi-discrete Hankel transform for nonlinear beam propagation', Chin. Phys. B **18**, 3893 (2009).

[298] M. Guizar-Sicairos and J. C. Gutiérrez-Vega, 'Computation of quasi-discrete Hankel transforms of integer order for propagating optical wave fields', J. Opt. Soc. Am. A **21**, 53 (2004).

[299] C. R. Harris, K. J. Millman, S. J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N. J. Smith, R. Kern, M. Picus, S. Hoyer, M. H. van Kerkwijk, M. Brett et al., 'Array programming with NumPy', Nature **585**, 357 (2020).

[300] M. Frigo and S. Johnson, 'The Design and Implementation of FFTW3', Proc. IEEE **93**, 216 (2005).

[301] S. K. Lam, A. Pitrou and S. Seibert, 'Numba: a LLVM-based Python JIT compiler', in Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC (2015), pp. 1–6.

[302] R. Lehoucq and J. Scott, 'An evaluation of software for computing eigenvalues of sparse nonsymmetric matrices', 1996.

[303] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison et al., 'PyTorch: An Imperative Style, High-Performance Deep Learning Library', in Advances in Neural Information Processing Systems, Vol. 32, edited by H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox and R. Garnett (2019), pp. 8024–8035.

[304] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne and Q. Zhang, *JAX: Composable transformations of Python+NumPy programs*, 2018.

[305] J. Stenger, S. Inouye, A. P. Chikkatur, D. M. Stamper-Kurn, D. E. Pritchard and W. Ketterle, 'Bragg Spectroscopy of a Bose-Einstein Condensate', Phys. Rev. Lett. **82**, 4569 (1999).

[306] D. M. Stamper-Kurn, A. P. Chikkatur, A. Görlitz, S. Inouye, S. Gupta, D. E. Pritchard and W. Ketterle, 'Excitation of Phonons in a Bose-Einstein Condensate by Light Scattering', Phys. Rev. Lett. **83**, 2876 (1999).

[307] 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).

[308] T. W. B. Kibble, 'Topology of cosmic domains and strings', J. Phys. A: Math. Gen. **9**, 1387 (1976).

[309] W. H. Zurek, 'Cosmological experiments in superfluid helium?', Nature **317**, 505 (1985).

[310] P. M. Chaikin and T. C. Lubensky, *Principles of Condensed Matter Physics* (Cambridge University Press, Cambridge, 1995).

[311] P. C. Hohenberg and B. I. Halperin, 'Theory of dynamic critical phenomena', Rev. Mod. Phys. **49**, 435 (1977).

[312] L. Corman, L. Chomaz, T. Bienaimé, R. Desbuquois, C. Weitenberg, S. Nascimbène, J. Dalibard and J. Beugnon, 'Quench-Induced Supercurrents in an Annular Bose Gas', Phys. Rev. Lett. **113**, 135302 (2014).

[313] L. Chomaz, L. Corman, T. Bienaimé, R. Desbuquois, C. Weitenberg, S. Nascimbène, J. Beugnon and J. Dalibard, 'Emergence of coherence via transverse condensation in a uniform quasi-two-dimensional Bose gas', Nat. Commun. **6**, 6162 (2015).

[314] C. Politi, 'Optical dipole trap for an erbium and dysprosium mixture', Master's thesis (University of Pisa, 2017).

[315] G. Unnikrishnan, C. Beulenkamp, D. Zhang, K. P. Zamarski, M. Landini and H.-C. Nägerl, 'Long distance optical transport of ultracold atoms: A compact setup using a Moiré lens', Rev. Sci. Instrum. **92**, 063205 (2021).

[316] D. H. J. O'Dell, S. Giovanazzi and C. Eberlein, 'Exact Hydrodynamics of a Trapped Dipolar Bose-Einstein Condensate', Phys. Rev. Lett. **92**, 250401 (2004).

[317] C. Eberlein, S. Giovanazzi and D. H. J. O'Dell, 'Exact solution of the Thomas-Fermi equation for a trapped Bose-Einstein condensate with dipole-dipole interactions', Phys. Rev. A **71**, 033618 (2005).

[318] S. G. Johnson, 'Notes on FFT-based differentiation', 2011.